



axioms

IMPACT
FACTOR
1.6

Article

Variational Quantum Eigensolver for Clinical Biomarker Discovery: A Multi-Qubit Model

Juan Pablo Acuña González, Moisés Sánchez Adame and Oscar Montiel

Special Issue

A Century of Quantum Mechanics: Mathematical Foundations and Computational Applications

Edited by

Prof. Dr. Oscar Montiel Ross



<https://doi.org/10.3390/axioms15010023>

Article

Variational Quantum Eigensolver for Clinical Biomarker Discovery: A Multi-Qubit Model

Juan Pablo Acuña González , Moisés Sánchez Adame  and Oscar Montiel * 

Instituto Politécnico Nacional, CITEDI-IPN, Tijuana 22435, BC, Mexico; jagonzalez@citedi.mx (J.P.A.G); msanchez@citedi.mx (M.S.A.)

* Correspondence: oross@citedi.mx

Abstract

We formalize an inverse, data-conditioned variant of the Variational Quantum Eigensolver (VQE) for clinical biomarker discovery. Given patient-encoded quantum states, we construct a task-specific Hamiltonian whose coefficients are inferred from clinical associations and interpret its expectation value as a calibrated energy score for prognosis and treatment monitoring. The method integrates coefficient estimation, ansatz specification with basis rotations, commuting-group measurements, and a practical shot budget analysis. Evaluated on public infectious disease datasets under severe class imbalance, the approach yields consistent gains in balanced accuracy and precision–recall over strong classical baselines, with stability across random seeds and feature ablations. This variational energy scoring framework bridges Hamiltonian learning and clinical risk modeling, offering a compact, interpretable, and reproducible route to biomarker prioritization and decision support.

Keywords: variational quantum eigensolver; quantum machine learning; quantum biomarkers; clinical data; multi-qubit systems

MSC: 68T05; 81P68; 92C50

1. Introduction

Quantum machine learning (QML) has emerged as a promising interdisciplinary field at the confluence of quantum computing, machine learning, and quantum information [1,2]. Fueled by the potential for computational speedups and the ability to process data in quantum mechanical ways [3,4], QML explores how to take advantage of quantum computers to enhance classical data analysis. Early works in the field laid the foundation for a new theoretical framework for learning [5], while more recent contributions have focused on the practical implementation of algorithms and their potential to achieve a quantum advantage [6].

QML has been applied to various domains, from methods based on quantum kernels to improve predictive accuracy [7] to applications in data analysis such as breast cancer detection and image classification [8]. Although the field holds immense promise, it is not without challenges, most notably the issue of barren plateaus, which can hinder the training of deep quantum neural networks [9].

Recent reviews highlight the growing potential of QML in biological and biomedical domains [10]. Applications in the biological sciences span various areas, including bioinformatics, drug discovery, and medical diagnostics [11]. For example, QML has been



Academic Editor: Angel Ricardo Plastino

Received: 12 November 2025

Revised: 5 December 2025

Accepted: 24 December 2025

Published: 27 December 2025

Copyright: © 2025 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

explored for drug discovery through qualitative structure activity relationship (QSAR) predictions, demonstrating that quantum classifiers can outperform classical ones, particularly when working with limited data and a reduced number of features [12]. In bio-informatics, a quantum annealer has been trained on transcription data, showing a slight advantage in classification and nearly equal ranking performance compared to classical methods on simplified datasets [13]. Similarly, annealing algorithms of the Ising type, inspired by quantum processors, have shown superior classification performance in small training sets of omics cancer data, a finding particularly relevant for rare diseases or limited clinical data [14].

Further studies have showcased the utility of QML in medical diagnostics. A quantum machine learning framework, QProteoML, has been proposed to predict drug sensitivity in multiple myeloma using high-dimensional protein data, outperforming classical machine learning algorithms [15]. For breast cancer prediction, QML has been identified as a cost-effective and efficient strategy, with specific algorithms such as Pegasos-QSVC excelling in recall metrics and quantum neural networks (QNNs) achieving high precision for binary classification of genomic sequence data [16,17]. Research also explores the use of quantum-inspired machine learning for the prediction of metastasis in breast cancer, employing least-squares methods via quantum measurements [18]. The mapping of quantum algorithms to applications in electronic health records (EHRs), omics, and imaging has also been detailed, along with the challenges of discovery, validation, and adaptation of quantum computing in these fields [19]. This body of work underscores the interdisciplinary nature of QML and its potential to address complex challenges in molecular biology and medicine.

A key area of research within QML is the development of quantum algorithms for tasks such as classification and regression. The variational quantum eigensolver (VQE) [20] is a prime example of a hybrid quantum-classical algorithm that has been adapted from its traditional use in molecular ground state estimation to solve optimization problems [21]. As VQE gained popularity, a variety of research has focused on improving its performance and extending its capabilities, particularly for noisy intermediate-scale quantum (NISQ) devices. Recent advances include methods to handle constraints and ensure the algorithm finds a true eigen-state [22], as well as new approaches for exploring the Hilbert space to find non-orthogonal solutions [23]. To address limitations in circuit complexity, researchers have proposed techniques such as clustering qubits to create shallower circuits [24] and measurement-based schemes that use entangled resource states [25]. Other optimizations include qubit-efficient strategies that can represent complex states with fewer physical qubits by using sequential measurements [26] and accelerated VQE algorithms that interpolate between VQE and quantum phase estimation (QPE) [27]. Furthermore, VQE has been shown to be effective for complex problems, with successful implementations on the Hamiltonian with hundreds of Pauli terms [28]. These methodological advances, along with comprehensive reviews of VQE's components and extensions [29], have cemented its position as a central tool in near-term quantum computing.

Beyond its traditional application in finding the ground state of the molecular Hamiltonian, VQE and its counterpart, the variational quantum classifier (VQC), have found an expanding role in the biological and biomedical sciences, particularly in addressing complex optimization and classification problems. For example, VQE has been leveraged to improve molecular predictions in drug discovery, demonstrating that its application in molecular dynamics simulations can lead to improved accuracy compared to classical methods [30]. This is complemented by its use to identify candidate inhibitors through synergy with quantum graph neural networks, which can model molecular structures [31]. In a similar vein, researchers have applied VQE to find the ground state of molecular systems of genes, a critical step in understanding genetic interactions [32].

VQE and VQC have also been explored for a variety of clinical and medical informatics tasks. For example, they show promise in optimizing clinical trial site selection by finding the ground state of a Hamiltonian that represents the optimization problem, potentially improving high-dimensional classical parameter spaces [33]. In medical diagnostics, VQC has been successfully applied to problems such as classifying diabetes using an 8- and 4-qubit circuit with rotation and controlled-NOT (*CNOT*) entanglement layers [34] and predicting asthma with high accuracy using the quantum approximate optimization algorithm (QAOA) for feature selection [35]. Furthermore, a hybrid autoencoder and VQC framework has been proposed for the classification of anomalies in biological samples, demonstrating its potential to identify irregularities [36]. These applications highlight the ability of VQE and VQC to bridge quantum and molecular systems, a capability that makes them particularly promising for medical applications where complex, high-dimensional data are prevalent [37]. Such advances show the potential of these variational methods to transform biomedical data analysis by moving beyond traditional computational paradigms to harness quantum-mechanical principles for more efficient and accurate results.

A gap persists in the literature regarding the application of quantum computing to immunology and epidemiology, especially in the context of the human immunodeficiency virus (HIV) and co-infections. Immunology itself generates data of extraordinary complexity—ranging from protein interaction networks to single-cell sequencing—which often exceed the representational and computational capacity of classical approaches. To address this, Basu et al. [38] have proposed a quantum framework for cell-centric therapeutics, which employs quantum algorithms to infer missing links in protein interaction networks by mapping them into exponentially large Hilbert spaces. This approach offers a novel foundation for precision medicine and provides a computational pathway to model the intricate interplay of viral and host proteins, immune cells, and co-morbidity in HIV pathogenesis.

Complementing this progress, quantum-mechanical perspectives of immune signaling have begun to emerge. One such model proposes that T-cell activation operates through quantized energy transfers during protein phosphorylation, mediated by receptor phosphorylation cycles initiated by pulses in peptide complexes [39]. This framework suggests that fundamental immunological processes may unfold in discrete quantum units enabling more accurate modeling of disease progression in conditions such as HIV, where immune dysfunction is central. Similarly, the concept of “quantum micro-RNA (mi-RNA) immunity” has been advanced as a lens to interpret viral manipulation of host gene regulation [40]. By treating aberrant mi-RNA regulation as a process driven by entanglement, this model connects quantum principles to the nonlinear multivariate dynamics of viral immunodeficiency, opening avenues for QML applications to viral genomics and cancer biology.

On an epidemiological scale, the first empirical steps towards quantum applications in population health have been demonstrated by Dan Roosan and his research group in their study of QAOA for HIV cluster forecasting [41]. Using QAOA, the method achieved a clustering accuracy of 92% in 1.6 s, outperforming classical density-based approaches. In parallel, a hybrid neural network reached 94% prediction precision for HIV prevalence, surpassing a purely classical counterpart. In particular, quantum Bayesian networks revealed the causal roles of housing instability and stigma in driving cluster emergence and expansion, highlighting structural determinants of disease. This study demonstrates the potential of quantum computing to move beyond applications at a molecular scale towards modeling at the population level, thereby informing prevention strategies, resource allocation for pre-exposure prophylaxis (PrEP), and interventions addressing structural inequities.

Building on this emerging landscape, Choppara and Lokesh [42] worked with predictions of genetic mutations in viral proteins, with a specific focus on severe acute respiratory

syndrome coronavirus 2 (SARS-CoV-2) as a model system. Leveraging quantum features for dimensionality reduction and embedding them within a quantum long short-term memory (QLSTM) architecture, they demonstrate the model's ability to capture complex, nonlinear relationships in genomic data. QLSTM outperforms classical deep learning models in mutation prediction accuracy while also providing biologically meaningful insights into immune evasion and viral transmission.

Researchers at Merrimack College, led by Roosan, have made significant contributions to the field, with a particular focus on applying quantum computing and QML to complex problems in biology and clinical medicine. Their work spans diverse quantum algorithms and architectures at the molecular scale, demonstrating how quantum methods for complex bio-molecular interactions can predict protein–ligand binding affinity by mapping molecular interactions to a high-dimensional Hilbert space, capturing complex relationships with fewer parameters on par with or even improving classical models [43]. This is complemented by quantum simulation for protein structural analysis, which investigated the structural stability of proteins implicated in neuro-degenerative disorders and identified regions for protein–protein interactions [44]. In addition, a quantum regression model was developed to predict the toxicity of herbal compounds, integrating molecular data to achieve lower prediction errors than classical linear regression and random forest models [45]. Their research also highlights how a hybrid QNN can outperform a classical liquid neural network (LNN) in drug re-purposing by leveraging amplitude encoding and feature maps to analyze drug databases [46].

Expanding on these advances, Roosan's group further demonstrated the power of a quantum variational transformer model for improved cancer classification, using a hybrid approach to capture intricate genomic patterns more effectively than classical transformer models, resulting in reduced misclassifications for complex cancer types [47]. They also developed a prototype for drug re-purposing using omics data, integrating quantum principal component analysis (QPCA) and quantum kernel methods to achieve tighter clustering and better performance compared to classical approaches, particularly when combined with large language models (LLMs) [48]. Together, these studies establish a foundational context for the present work, showcasing the ability of QML methods to address complex challenges in biological and clinical research.

In this work, we extend a previously proposed approach using VQE for clinical biomarker discovery [49]. Our method inverts the traditional VQE paradigm: instead of minimizing the energy of a fixed Hamiltonian, we fix the quantum state (representing a patient's clinical profile) and search for a Hamiltonian whose ground state energy closely matches the state's expectation value. This "inverse VQE" offers a novel lens for identifying complex biomarkers with multiple relationships that are highly predictive of a patient's condition.

Our prior work demonstrated the feasibility of this approach using a single-qubit model, successfully identifying key biomarkers related to HIV treatment response and comorbidities. However, a single qubit's limited degrees of freedom restricts its ability to capture the multifaceted nature of clinical data. In this work, we generalize the model to a multi-qubit system, allowing for the simultaneous encoding of multivariate patient features and, critically, the inclusion of multiple qubit interaction terms in the Hamiltonian. The resulting increase in expressivity and complexity enables a deeper exploration of the relationships between biomarkers, including those that are non-classical in nature.

The specific clinical data used come from a cohort of patients with HIV and tuberculosis (TB) co-infection. We analyze key clinical markers such as viral load and CD4 count differences, along with quantum game theoretic features. We hypothesize that a multi-qubit model, by accommodating interaction terms like $I \otimes I$ and $Z \otimes Z$, will reveal

more intricate biomarker relationships and provide a more robust and clinically relevant set of Hamiltonian.

2. Materials and Methods

2.1. Data Acquisition and Feature Engineering

This study leverages the `qphen` dataset, a unique resource for applying quantum-inspired algorithms to clinical data related to HIV. The dataset is publicly available through the `qvirus` package [50] on the Comprehensive R Archive Network (CRAN) [51]. Data were collected from a longitudinal cohort of 213 individuals with HIV and TB co-infection at a clinic in Guerrero, Mexico, from 2018 to 2024, according to institutional and national research ethics guidelines. The protocol was approved by the Local Health Research Committee No. 1102 of the General Regional Hospital No. 1.

The dataset integrates classical clinical measurements with engineered features derived from a quantum game theoretic model of HIV phenotypes [52]. Key variables include differences in viral load (`v1_diff`) and CD4 T lymphocyte counts (`cd4_diff`), along with binary indicators of TB co-infection (`TB_1`) and genotypic drug resistance (`GR_1`).

A stringent data cleaning process was necessary to prepare the initial cohort consisting of 213 patient records for the quantum game theoretic model. This cleaning process excluded records with only a single longitudinal value, as these data points did not allow the derivation of change rates or the temporal dynamics that represent the competition between lymphocytes and viral load over time, features that proved critical for the model. Following this cleaning, the base utilized for all subsequent analyses was reduced to 176 records, which were used for feature engineering and modeling procedures.

A central component of the feature engineering process is the *Nearest Payoff Algorithm*, which bridges the observed viral dynamics in a study cohort with the vast state space of quantum game simulations. The quantum game is defined by a Hilbert space \mathcal{H} , an initial state Λ , and a set of strategies $\{s_j\}$ and corresponding payoffs $\{\pi_j\}$. Two infectious phenotypes are considered: a low-replication phenotype v and a high-replication phenotype V , encoded as qubits $|0\rangle$ and $|1\rangle$, respectively.

The game's dynamics are governed by unitary operators. For a two-phenotype system, the initial state is prepared as follows:

$$|U\rangle = \frac{1}{\sqrt{2}}(I^{\otimes 2} + i\sigma_x^{\otimes 2}), \quad (1)$$

where I is the Identity, and σ_x is the Pauli-X gate. Subsequent quantum moves, including the Pauli gates and the Hadamard gate H , are applied to generate superposition states that reflect different phenotype interactions. After applying the appropriate combination of quantum gates to evolve the initial state $|U\rangle$ to the final state $|\psi_f\rangle$, the expected payoffs π_v and π_V are computed as a weighted sum of the probability amplitudes of the final state (the quantum outcomes) and the classical payoff parameters:

$$\pi_v = \alpha|\langle\psi_f|00\rangle|^2 + \beta|\langle\psi_f|01\rangle|^2 + \gamma|\langle\psi_f|10\rangle|^2 + \theta|\langle\psi_f|11\rangle|^2, \quad (2)$$

and

$$\pi_V = \alpha|\langle\psi_f|00\rangle|^2 + \gamma|\langle\psi_f|01\rangle|^2 + \beta|\langle\psi_f|10\rangle|^2 + \theta|\langle\psi_f|11\rangle|^2, \quad (3)$$

Because $|\psi_f\rangle$ is a normalized quantum state, these measurement probabilities satisfy the following equation:

$$|\langle\psi_f|00\rangle|^2 + |\langle\psi_f|01\rangle|^2 + |\langle\psi_f|10\rangle|^2 + |\langle\psi_f|11\rangle|^2 = 1.$$

Consequently, the expected payoffs π_v, π_V are convex combinations of the classical parameters $\alpha, \beta, \gamma, \theta$, weighted by the corresponding quantum probabilities. Unitary evolution redistributes the amplitudes among the four outcomes without creating or removing the probability mass.

The classical payoff parameters $\alpha, \beta, \gamma, \theta$ correspond to the outcomes of the four possible interactions ($v \times v, v \times V, V \times v, V \times V$), capturing biologically relevant properties such as resource consumption, replication rate, and survival. These parameters are not normalized—they are scalar values that represent the payoff for each interaction. Only the quantum probabilities are normalized, ensuring the following:

$$\sqrt{|\langle \psi_f | 00 \rangle|^2 + |\langle \psi_f | 01 \rangle|^2 + |\langle \psi_f | 10 \rangle|^2 + |\langle \psi_f | 11 \rangle|^2} = 1$$

The nearest payoff algorithm maps an observed viral load difference from a patient, L_{obs} , to the closest quantum expected payoff previously simulated, p_j^* , from a set Π of all possible payoffs. This process yields both a numeric prediction p_j^* and a phenotypic characterization N_j^* for each patient, based on the combination of the quantum strategy that generated the closest payoff. By embedding these quantum game theoretic outcomes into clinically interpretable features, the approach enhances the dataset with descriptors that capture nontrivial interactions between viral phenotypes, providing a richer basis for downstream analyses such as multi-qubit Hamiltonian modeling.

The key engineered features of the qphen dataset include:

- **Quantum Rewards:** The variable `nearest_payoff` provides numerical prediction of the phenotype of a patient. It is the value of the closest simulated quantum payoff, or expected outcome, that best fits the patient's viral load difference.
- **Quantum Strategy Encoding:** This category is crucial for understanding the quantum model's interpretation of a patient's viral dynamics. The variables here encode the specific quantum gates used as strategies by two distinct HIV phenotypes, a low-replication phenotype (v) and a high-replication phenotype (V).
 - `phen_1` indicates whether the phenotype associated with the nearest payoff is v or V .
 - `str1_2`: represents the primary strategy of the first phenotype, with a binary encoding that corresponds to the Pauli-X gate (X), the T -gate (T), or the H gate.
 - `str1_3`: is the alternative strategy for the first phenotype, also using binary encoding.
 - `str2_2`: is the primary strategy of the second phenotype. Its binary encoding corresponds to one of several quantum gates: H , I , Phase gate (S), T , X , Pauli- Y gate (Y), or Pauli- Z gate (Z).
 - `str2_3` to `str2_7`: are alternative strategies for the second phenotype. Each variable is a binary indicator (0 or 1) for the presence or absence of a specific quantum gate of the set $\{H, I, S, T, X, Y, Z\}$.
- **Batch Indicator:** `batch_1` distinguishes predictions performed on the full dataset (0) versus subsets processed by batch (1).

Other features in the dataset are as follows:

- **Phenotypic Clustering:** Variables `classification_2` through `classification_4` are derived from a k-means clustering of standardized viral load and CD4 differences, which groups patients by their phenotypic behavior.
- **Clinical Differences:** Variables such as `v1_diff` and `cd_diff` are raw differences in viral load and CD4 T lymphocyte counts. Standardized versions (`vlogs_diff_mean` and `cds_diff_mean`) are also included.
- **Binary Clinical Indicators:** `TB_1` and `GR_1` are binary variables that indicate the presence of TB co-infection and genotypic drug resistance, respectively.

For example, the payoff label $V_X_H_b$ translates to $phen_1 = V$ (high-replication phenotype), $str1_2 = X$ (the first strategy is the Pauli- X gate), $str2_2 = H$ (the second strategy is the Hadamard gate), and $batch_1 = 1$ (indicating a batch-processed subset).

This synthesis of classical clinical data with features derived with quantum games of phenotypes provides a robust foundation for biomarker discovery and patient stratification using the multi-qubit variational approach.

2.2. Multi-Qubit VQE Algorithmic Methods

This section presents the theoretical and algorithmic principles underlying the extension of the VQE to multi-qubit systems. As a hybrid quantum-classical algorithm, VQE enables practical simulations on near-term quantum hardware by preparing shallow quantum circuits while delegating computationally intensive optimization tasks to a classical processor. At its core, VQE leverages quantum simulation and the Schrodinger equation to define a Hamiltonian operator $\hat{\mathcal{H}}$, whose eigenvalues correspond to the energy levels of the system. The variational principle ensures that the expectation value of any trial state provides an upper bound to the energy of the ground state.

For an n -qubit system, VQE uses a parameterized quantum trial state, or ansatz, $U(\vec{\theta})$:

$$|\psi(\vec{\theta})\rangle = U(\vec{\theta})|0\rangle^{\otimes n}. \tag{4}$$

The Hamiltonian is expressed as a linear combination of Pauli tensor products:

$$\hat{\mathcal{H}} = \sum_i c_i P_i, \tag{5}$$

where each P_i is a tensor product of Pauli matrices (I, X, Y, Z), and c_i are real coefficients. The expectation value of the Hamiltonian is

$$\langle \hat{\mathcal{H}} \rangle = \langle \psi(\vec{\theta}) | \hat{\mathcal{H}} | \psi(\vec{\theta}) \rangle = \sum_i c_i \langle \psi(\vec{\theta}) | P_i | \psi(\vec{\theta}) \rangle = E(\vec{\theta}), \tag{6}$$

with $\langle P_i \rangle$ obtained from projective measurements on the quantum device. By the variational principle, for any normalized trial state $\langle \psi | \psi \rangle = 1$,

$$\langle \hat{\mathcal{H}} \rangle \geq E_0, \tag{7}$$

where E_0 is the true energy of the ground state. Minimizing $E(\vec{\theta})$ over trial states provides the best approximation of E_0 .

In multi-qubit systems, Hamiltonians are decomposed into single-qubit operators (acting on one qubit) and interaction terms (acting on multiple qubits). For a two-qubit system, single-qubit terms include $Z \otimes I$ or $X \otimes I$, whereas the interaction terms include $Z \otimes Z$ or $X \otimes X$, which induce correlations between the qubits. The identity operator I acts as a placeholder for qubits that are not directly involved in a term.

Expectation values are measured using either analytical or sampling-based approaches. The analytical method involves simulation where the quantum state $|\psi\rangle$ can be represented as a density matrix $\rho = |\psi\rangle\langle\psi|$, and the exact expectation value of a Pauli operator P is as follows:

$$\langle P \rangle = \text{Tr}(\rho P). \tag{8}$$

The mathematical notation $\langle \hat{\mathcal{A}} \rangle = \langle \psi | \hat{\mathcal{A}} | \psi \rangle$ is the standard and simplest way to represent the expectation value of an observable $\hat{\mathcal{A}}$ with respect to a quantum state $|\psi\rangle$.

Alternatively, the sampling-based method is used on real quantum hardware, where expectation values are probabilistically estimated from repeated measurements (N_{shots}).

$$\hat{E}_i = \frac{1}{N_{shots}} \sum_{k=1}^{N_{shots}} e_k, \quad e_k \in \{\pm 1\}, \tag{9}$$

where e_k is the eigenvalue of P_i for the k -th shot after appropriate basis rotations. Specifically, to measure multi-qubit Pauli terms, the state must first be transformed with a basis change unitary operator U_{basis} to map the observable to the standard computational (Z) basis:

Step 4 in Algorithm 1 converts the raw results of the classical measurement into the eigenvalues of the Pauli operator being measured. A projective measurement in the Z-basis yields bit strings of 0s and 1s. For the Pauli operator Z, the state $|0\rangle$ corresponds to the eigenvalue 1, while $|1\rangle$ corresponds to the eigenvalue -1 . Thus, each measured bit is assigned to its associated eigenvalue. For Pauli multi-qubit terms, the eigenvalue of a full bit string $b_1 b_2 \dots b_N$ is obtained by multiplying the eigenvalues of the individual qubits. This mapping ensures that the statistical average of the eigenvalues across many shots yields the correct expectation value: $\langle P \rangle = (1) \times \text{Pr}(\text{eigval}_1) + (-1) \times \text{Pr}(\text{eigval}_{-1})$. This assignment to its associated eigenvalues is essential because the quantum computer outputs classical bits, while the VQE energy calculation requires Pauli eigenvalues. The mapping bridges this gap, enabling the statistical estimation of the expectation values from the measurement data.

$$\langle P_i \rangle = \langle \psi | U_{basis}^\dagger P_i U_{basis} | \psi \rangle_{Z\text{-measurement}}. \tag{10}$$

Algorithm 1 Sampling-Based Pauli Expectation Value Estimation

- 1: Requires: Quantum state $|\psi\rangle$, Pauli operator P , shots N
 - 2: Apply basis rotations for measurement in the Z-basis
 - 3: Compute probability distribution from state vector
 - 4: Draw N samples to simulate measurement outcomes
 - 5: Map samples to eigenvalues ± 1
 - 6: Average eigenvalues to estimate $\langle P \rangle$
 - 7: **return** Estimated expectation value
-

The analytical method provides exact benchmarks for algorithm validation, whereas the sampling-based method simulates realistic hardware execution, accounting for statistical fluctuations.

The core of our hybrid quantum-classical approach involves preparing a patient-specific quantum state that encodes the clinical status. This is achieved through two main components: amplitude encoding (the mapping of classical features to rotation angles) and the ansatz circuit (the parameterized quantum circuit structure).

2.2.1. Amplitude Encoding: Mapping Clinical Data to Angles

For our multi-qubit system, four key clinical features are assigned to the four rotation angles $(\theta_1, \phi_1, \theta_2, \phi_2)$ of the ansatz circuit. This technique, known as feature embedding, utilizes Min–Max Normalization to scale each clinical value x into the rotation range $[0, 2\pi]$, ensuring that the entire feature space is represented by the quantum Hilbert space.

We use Min–Max normalization rather than standardization because the rotation angles of the quantum circuit must lie in the finite interval $[0, 2\pi]$ to represent valid qubit rotations. Standardization (zero mean; unit variance) could produce values outside this range, which cannot be directly mapped to rotation angles without additional clipping or rescaling. Min–Max normalization ensures that all clinical features are proportionally scaled to the full rotation range, preserving relative differences between patients while fully using the Hilbert space of the qubits.

The four encoded features are as follows: viral load mean difference (v1_diff), CD4 mean difference (cd_diff), TB co-infection (TB_1), and genotypic drug resistance (GR_1). Each feature $x \in \{v1_diff, cd_diff, TB_1, GR_1\}$ is normalized using the full range of the dataset:

$$x^* = \frac{x - \min(x)}{\max(x) - \min(x)} \cdot 2\pi, \quad x \in \{v1_diff, cd_diff, TB_1, GR_1\}. \quad (11)$$

where x^* defines one of the four rotation angles. The resulting normalized features define the rotation parameters of the ansatz as

$$\theta_1 \leftarrow (v1_diff)^*, \quad \phi_1 \leftarrow (cd_diff)^*, \quad \theta_2 \leftarrow (TB_1)^*, \quad \phi_2 \leftarrow (GR_1)^*.$$


2.2.2. Ansatz Definition and Circuit Diagram

This work adopts a fully specified and reproducible two-qubit ansatz tailored to the clinical-data encoding scheme used in our inverse VQE method. The ansatz consists of three components:

1. State preparation layer: each qubit receives one rotation around the x -axis and one around the y -axis, parameterized by clinical features.
 - Qubit 1: $R_x(\theta_1)R_y(\phi_1)$
 - Qubit 2: $R_x(\theta_2)R_y(\phi_2)$
2. Entangling layer: a single CNOT gate couples the qubits.
3. No additional variational layers: unlike deep VQE ansatze, our model keeps the circuit intentionally shallow to preserve interpretability and reduce barren plateau risk when operating in a data-conditioned (inverse) regime.

Thus, the ansatz is fully determined and identical across all patients, except for the rotation parameters, which come from their normalized clinical features.

The circuit initializes both qubits in the ground state $|0\rangle|0\rangle$ and applies two sequential single-qubit rotations (R_x and R_y) on each qubit, followed by a CNOT gate for entanglement. The circuit diagram is as follows:



The circuit architecture is identical for all patients and only the rotation parameters change according to their clinical attributes.

2.2.3. Resulting Amplitudes and Entangled State

Each sequence of single-qubit rotations transforms the basis state $|0\rangle$ into a superposition state whose amplitudes are functions of the input clinical features. Applying feature-dependent rotations produces the amplitudes $\alpha_i = f(\theta_i, \phi_i) = \cos(\frac{\phi_i}{2}) \cos(\frac{\theta_i}{2}) - i \sin(\frac{\phi_i}{2}) \sin(\frac{\theta_i}{2}), \beta_i = g(\theta_i, \phi_i) = \sin(\frac{\phi_i}{2}) \cos(\frac{\theta_i}{2}) - i \cos(\frac{\phi_i}{2}) \sin(\frac{\theta_i}{2})$. Here, θ_i and ϕ_i are the real-valued angles encoding the clinical features of qubit i , while α_i, β_i are simply the resulting complex amplitudes produced by this rotation sequence, leading to the following:

$$R_x(\theta_i)R_y(\phi_i)|0\rangle = \alpha_i|0\rangle + \beta_i|1\rangle. \quad (12)$$

Before the entanglement step (CNOT), the state of the two-qubit system is a simple tensor product. The CNOT gate then acts to couple the features, mapping the state $|10\rangle$ to $|11\rangle$. Before the entangling operation, the system is in the tensor product:

$(\alpha_1|0\rangle + \beta_1|1\rangle) \otimes (\alpha_2|0\rangle + \beta_2|1\rangle)$. The CNOT couples the two embedded feature sets by mapping $|10\rangle \rightarrow |11\rangle$.

Generally, this describes the state of a multi-qubit register, explicitly written in terms of the applied rotation gates (R_x, R_y) and the clinical features they encode. The input clinical features (x_1, x_2) are mapped to the circuit parameters defining the two-qubit quantum register state $\psi_{patient}$

This results in the final, entangled patient state:

$$|\psi_{patient}(\theta, \phi)\rangle = CNOT_{1,2}R_x(\theta_i)R_y(\phi_i)|0\rangle^{\otimes N}, \tag{13}$$

The state in Equation (13) provides a structured quantum encoding of the patient’s viral, immunological, and co-infection profile. It captures both individual clinical contributions and their latent interactions, an essential property for the Hamiltonian-based search performed in the next stage of the model.

2.2.4. Inverse/Data-Conditioned VQE

Rather than the standard VQE, where ansatz parameters are optimized for a fixed Hamiltonian, here the paradigm is inverted: the patient-specific quantum state is fixed, and the algorithm searches for a Hamiltonian whose ground state best matches this state. The objective function is the delta metric:

$$\Delta(\hat{\mathcal{H}}_{trial}, |\psi_{patient}\rangle) = |\langle \psi_{patient} | \hat{\mathcal{H}}_{trial} | \psi_{patient} \rangle - E_{\min}(\hat{\mathcal{H}}_{trial})|, \tag{14}$$

with E_{\min} representing the minimal eigenvalue of the Hamiltonian trial. The optimization problem is

$$\hat{\mathcal{H}}_{best} = \arg \min_{\hat{\mathcal{H}}_{trial}} \Delta(\hat{\mathcal{H}}_{trial}, |\psi_{patient}\rangle), \tag{15}$$

which identifies the Hamiltonian that most accurately represents the patient’s encoded state and highlights the most informative features.

In standard VQE, the Hamiltonian ($\hat{\mathcal{H}}$) is fixed by the physical problem (e.g., a molecular structure), and the algorithm optimizes the parameters of a quantum trial state ($\psi(\vec{\theta})$) to find the state that minimizes the expectation value $\langle \psi(\vec{\theta}) | \hat{\mathcal{H}} | \psi(\vec{\theta}) \rangle$.

In this inverse VQE method, the patient-specific quantum state $|\psi_{patient}\rangle$, which encodes the patient’s clinical features, is fixed. The objective is to find $\hat{\mathcal{H}}$, where $\Delta(\hat{\mathcal{H}}_{trial} | \psi_{patient})$ quantifies the alignment. The task specific Hamiltonian is constructed as a linear combination of Pauli terms (P_j), which correspond to the specific single-qubit or multi-qubit interaction Equation (5). The core of the process, detailed in Algorithm 2, is how the coefficients (c_j) are generated. The c_j coefficients are explicitly inferred from clinical associations within the dataset by computing the statistical correlations between the features encoded in P_j . The algorithm then performs a classical search to find $\hat{\mathcal{H}}_{best}$ whose minimal eigenvalue (E_{\min}) is best aligned with the expectation value of the fixed patient state. The $\hat{\mathcal{H}}_{best}$ obtained by the search in this HIV application were often single-term operators. The Results section provides concrete examples linking the final coefficient to a clinical correlation.

This framework integrates the measurement of multi-qubit Pauli terms with a stochastic Hamiltonian optimization that highlights patient-specific feature correlations. The minimum delta is analogous to the ground state energy in conventional VQE, while the selected coefficients indicate the most informative clinical features.

Using the correlation between features as the coefficient of the corresponding Pauli term, Algorithm 2 ensures that the discovered $\hat{\mathcal{H}}_{best}$ is the most parsimonious and inter-

pretable biomarker combination that defines the patient's state. This is a crucial point that reflects the essence of novelty and the necessity of the proposed method. The Hamiltonian, though sometimes resulting in a single Pauli term (which is classically trivial to measure), is discovered through a non-trivial process that leverages the core strength of VQE for clinical biomarker discovery, which classical methods cannot replicate easily.

Algorithm 2 Optimal Clinical Hamiltonian Search for Multi-Qubit Systems

```

1: Requires: Patient record, clinical dataset, threshold  $\varepsilon$ , max iterations  $N_{\max}$ 
2: Initialize  $\Delta_{\min} \leftarrow \infty$ , best Hamiltonian  $\hat{\mathcal{H}}_{\text{best}}$ , feature set  $F_{\text{best}}$ ,  $i \leftarrow 0$ 
3: Encode patient features into  $(\theta_1, \phi_1, \theta_2, \phi_2)$  Equation (11)
4: Prepare ansatz  $|\psi(\theta_1, \phi_1, \theta_2, \phi_2)\rangle$ 
5: while  $\Delta_{\min} > \varepsilon$  and  $i < N_{\max}$  do
6:    $i \leftarrow i + 1$ 
7:   Generate trial Hamiltonian  $\hat{\mathcal{H}}_{\text{trial}} = \sum_j c_j P_j$  using feature correlations  $c_j = \text{cor}(f_j, f'_j)$ 
8:   Evaluate  $E_{\text{est}} = \sum_j c_j \langle P_j \rangle$  via measurement
9:   Compute  $E_{\min}(\hat{\mathcal{H}}_{\text{trial}})$  classically
10:   $\Delta = |E_{\text{est}} - E_{\min}|$ 
11:  if  $\Delta < \Delta_{\min}$  then
12:     $\Delta_{\min} \leftarrow \Delta$ 
13:     $\hat{\mathcal{H}}_{\text{best}} \leftarrow \hat{\mathcal{H}}_{\text{trial}}$ 
14:     $F_{\text{best}} \leftarrow \{f_j, f'_j\}$ 
15:  end if
16: end while
17: return  $\hat{\mathcal{H}}_{\text{best}}, F_{\text{best}}, \Delta_{\min}, i$ 

```

The proposed method is an inverse VQE where the goal is Hamiltonian learning, not just state preparation. The complete search space for a relative clinical Hamiltonian includes all possible k -local P_j on the N qubits used to encode patient features. Even for a small number of features ($N = 4$ or $N = 8$), this space is vast. Although the final $\hat{\mathcal{H}}_{\text{best}}$ may be simple (e.g., $Z_1 Z_2$), the search process that identifies this term as the most salient clinical biomarker is highly non-trivial.

The VQE framework provides a physically interpretable energy score for clinical risk. The final result of the method is the expectation value $\langle \psi_{\text{patient}} | \hat{\mathcal{H}}_{\text{best}} | \psi_{\text{patient}} \rangle$. In the context of the inverse VQE, this value is interpreted as a calibrated energy score for the prognosis risk of a patient. This is a compact, single-number metric directly derived from quantum mechanics.

When the optimal Hamiltonian simplifies to a single Pauli term, it means that the most relevant biomarker for the clinical task is the interaction (correlation) between the two features encoded by qubits 1 and 2, quantified by the coefficient c_j . VQE enables the selection of this most parsimonious and biologically meaningful combination from the complex feature space.

Although the inverse VQE may result in Hamiltonians that simplify to single terms, the VQE formulation is necessary for scalability and future generalization. The method is built to generalize to systems where the optimal clinical risk is represented by a complex entangling Hamiltonian that involves many-body interactions. Such complex Hamiltonians are not trivial to handle classically, and their expectation values are most naturally obtained using VQE commuting-group measurements.

The optimal Hamiltonian search described in Algorithm 2 is extended into a global binary classification routine to predict clinical status (e.g., TB co-infection or drug resistance) throughout a patient cohort Algorithm 3. In this classification context, the Hamiltonian expectation value for a patient x serves as the direct classification score $E_{\text{score}}(x)$ of Equation (16).

$$E_{score}(x) = \langle \psi(\phi(x)) | \hat{\mathcal{H}} | \psi(\phi(x)) \rangle = \sum_{k=1}^N C_k \langle \psi | \hat{P}_k | \psi \rangle. \tag{16}$$

The prediction \hat{y} is then a binary decision derived from this quantum score by applying a classically optimized threshold τ :

$$\hat{y} = \begin{cases} 1, & \text{if } E_{score}(x) \geq \tau, \\ 0, & \text{if } E_{score}(x) < \tau. \end{cases} \tag{17}$$

The goal shifts from finding a Hamiltonian that minimizes the Δ metric via Equation (14) for feature selection to finding the optimal pair $(\hat{\mathcal{H}}^*, \tau^*)$ that minimizes a selected classification loss function, $\mathcal{L}(\hat{y}, y)$, across the entire dataset.

The global optimization of $\hat{\mathcal{H}}$ and τ are driven by minimizing one of the three standard loss metrics in the iteration space.

1. Simple Error (\mathcal{L}_{Simple}): Minimizes the total number of misclassified instances, equivalent to $1 - Accuracy$, where the accuracy is the correct classification rate.

$$\mathcal{L}_{Simple} = \frac{1}{N_{total}} \sum_j \mathbb{I}(\hat{y}_j \neq y_j).$$

- \hat{y}_j : The predicted class (0 or 1) for the j -th instance;
- y_j : The true class (0 or 1) for the j -th instance;
- N_{Total} : The total number of instances (patients) in the dataset;
- \mathbb{I} : The indicator function, which equals 1 if the inside condition is true ($\hat{y}_j \neq y_j$, i.e., misclassifications) and 0 otherwise.

2. Balanced Error ($\mathcal{L}_{Balanced}$): Minimization of error based on the average of sensitivity (true positive, TP, rate) and specificity (true negative, TN, rate), making it robust against class imbalance.

$$\mathcal{L}_{Balanced} = 1 - \frac{1}{2}(\text{Sensitivity} + \text{Specificity}).$$

- **Sensitivity** : $\frac{TP}{TP+FN}$;
- **Specificity** : $\frac{TN}{TN+FP}$.

3. Penalized Error ($\mathcal{L}_{Penalized}$): Imposes a higher cost on false negatives (FNs) compared to false positives (FPs), often critical in medical diagnostics to avoid missing a condition.

$$\mathcal{L}_{Penalized} = \frac{2 \cdot FN + FP}{N_{total}}.$$

- **FN**: The total count of FNs (instances that are actually positive, $y = 1$, but were predicted to be negative, $\hat{y} = 0$);
- **FP**: The total count of FPs (instances that are actually negative, $y = 0$, but were predicted positive, $\hat{y} = 1$).

The core routine involves a stochastic global search where, at each iteration, a random Hamiltonian is generated, all patient scores are computed, and a classical grid search is performed over the expanded range of possible scores (from $-N$ to $+N$, where N is the number of Pauli terms) to find the threshold τ that minimizes the chosen loss function.

Algorithm 3 Global Multi-Qubit Classifier Search

```

1: Requires: Clinical Dataset  $D$ , Target Status  $y$ , Loss Function  $L$ , Max Iterations  $M$ ,
  Number of Pauli Terms  $N$ 
2: Initialize  $L_{\min}^* \leftarrow \infty, \hat{\mathcal{H}}^* \leftarrow \text{Null}, \tau^* \leftarrow \text{Null}$ 
3: Pre-calculate normalization bounds for feature encoding
4: Define threshold grid  $T \in \{-N, N\}$ 
5: for  $i=1$  to  $M$  do
6:   Generate Trial Hamiltonian  $\hat{\mathcal{H}}^{(i)}$ :
7:   i. Randomly select  $N$  Pauli operators  $\{P^k\}_{k=1}^N$ 
8:   ii. Assign coefficients  $C_k$  from random clinical feature correlations
9:   Compute Scores  $E^{(i)}$ :
10:  For all patients  $x_j \in D$ :
11:  Compute  $E_{\text{score}}(x_j)$  via the quantum ansatz and measurement routine (Equation (4))
12:  Optimize Threshold  $\tau$  (Grid Search):
13:  Find  $\tau^{(i)} = \operatorname{argmin}_{\tau \in T} L(\text{Classification}(E^{(i)}, \tau), y)$ 
14:  Compute  $L^{(i)} \leftarrow L(\text{Classification}(E^{(i)}, \tau^{(i)}), y)$ 
15:  if  $L^{(i)} < L_{\min}^*$  then
16:     $L_{\min}^* \leftarrow L^{(i)}$ 
17:     $\hat{\mathcal{H}}^* \leftarrow \hat{\mathcal{H}}^{(i)}$ 
18:     $\tau^* \leftarrow \tau^{(i)}$ 
19:  end if
20: end for
21: return Optimal Classifier  $(\hat{\mathcal{H}}^*, \tau^*)$  and  $L_{\min}^*$ 

```

The output of the routine is the optimal clinical Hamiltonian $\hat{\mathcal{H}}^*$ and the optimal threshold τ^* that together minimize the predefined classification loss function. This final pair constitutes the predictive quantum classifier for the target clinical status.

2.3. Classical Benchmarks

To rigorously assess the performance and stability of the multi-qubit VQE approach, its results were benchmarked against a suite of strong classical machine learning models. This comparative analysis was conducted under a unified and robust validation scheme designed to address class imbalance, ensure reproducibility under random conditions, and evaluate the robustness of the features.

The following models were selected as classical baselines for binary classification of TB_1 and GR_1:

1. **Regularized Logistic Regression (LogReg):** A strong linear baseline, implemented using the `brulee` package (equivalent to the `glmnet` engine), which incorporates L2 regularization (penalty) to prevent overfitting and optimize the loss function.
2. **Multi-Layer Perceptron (MLP):** A single-hidden-layer neural network, trained using the `brulee` package, which serves as a classical deep learning baseline that explicitly incorporates class weights to handle the data imbalance during optimization.

The entire modeling process, including feature processing, hyperparameter tuning, and final evaluation, was nested within a rigorous, multi-seed validation loop to ensure result stability. The analysis was replicated on $N = 10$ distinct random seeds. For each seed, a new 70/30 train/test split was performed, and the subsequent repeated cross-value (CV) and final evaluation were executed. The final metrics reported are the mean and standard deviation in these N replicates. Hyperparameter tuning for all classical baselines (LogReg, MLP) was performed using a 10-fold CV repeated 3 times on the training set. The selection of models during CV was based on maximization of the Area Under the Precision–Recall Curve (AUCPR), which is the most reliable metric for unbalanced data. The modeling pipeline utilized the `tidymodels` framework in R. All numeric predictors

were standardized ($\mu = 0, \sigma = 1$), and zero-variance features were removed using the `step_zv` function.

Crucially, to address the severe class imbalance in TB_1, the Synthetic Minority Over-sampling Technique (SMOTE) was applied within the CV folds of the training set to balance the training data for LogReg models. The MLP model was trained using explicit inverse class weights (inversely proportional to class frequency) instead of SMOTE, a common technique for gradient-based models, to directly influence the loss function toward the minority class.

To assess feature redundancy and the contribution of quantum-derived features, an ablation study was conducted. The complete validation pipeline (multi-seed CV and final test evaluation) was repeated for three distinct feature sets:

- A **Full Feature Set** (using all available clinical and quantum-derived features).
- A **No Quantum Features Set** (excluding features related to VQE payoffs, nearest payoff, and quantum structure).
- A **No Classical Features Set** (excluding raw clinical differences, ratios, and counts).

The comparison of performance metrics across these ablated sets provides direct evidence of the feature stability and incremental value of the proposed features.

3. Results

The search algorithm was applied to three distinct clinical scenarios to identify the optimal low dimensional clinical Hamiltonian $\hat{\mathcal{H}}_{best}$. This Hamiltonian serves as a personalized quantum biomarker, where its energy in the ground state E_{min} provides a measurable risk score, and its Pauli operator structure defines the most critical interaction (entanglement) between clinical features.

The objective function and the optimization process in this inverse VQE variant are closer to a feature selection/biomarker discovery problem than a standard quantum chemistry calculation.

The optimal Hamiltonian refers to the result of a data-conditioned classical search (Algorithm 2). In this variant, the goal is to find the $\hat{\mathcal{H}}_{best}$ that best satisfies the criterion for an ideal clinical biomarker: the patient-encoded quantum state $|\psi_{patient}\rangle$ must be the ground state of the discovered $\hat{\mathcal{H}}_{best}$. The algorithm achieves this by minimizing the Δ metric in Equation (15), with the delta distance given in Equation (14).

3.1. The Core Result: Discovered Clinical Hamiltonians (Biomarker Prioritization)

This section focuses on the output of the inverse VQE search (Algorithm 2) and directly addresses the claim of biomarker discovery for three distinct patient scenarios. The Hamiltonian terms obtained for each scenario represent the most parsimonious interaction, i.e., the most predictive pair of features, identified by the quantum-driven search. The coefficient c_j is the clinical association (correlation) that best forces the patient's state ($|\psi_{patient}\rangle$) to act as the ground state of the resulting Hamiltonian.

For two encoded clinical qubits, the algorithm searches the restricted Pauli operator basis $\mathcal{P} = \{I, X, Y, Z\} \otimes \{I, X, Y, Z\}$, which contains 16 candidate Pauli strings.

A clinical Hamiltonian is defined as any linear combination:

$$\hat{\mathcal{H}}(\mathbf{c}) = \sum_{j=1}^{16} c_j P_j, \quad P_j \in \mathcal{P},$$

where c_j are training parameters and each P_j acts as a pairwise candidate interaction between clinical features. The search is not unconstrained: we restrict the algorithm to one-term Hamiltonians during the clinical search phase to enforce an interpretable

and sparse biomarker. Thus, the optimization domain of Equation (5) reduces to the following equation:

$$\hat{\mathcal{H}} = c\mathcal{P}_\gamma, \quad P_\gamma \in \mathcal{P}.$$

This guarantees that each discovered Hamiltonian identifies exactly one dominant interaction between two encoded clinical features to capture the main dynamics for a specific profile of patients. For all Pauli strings \mathcal{P}_γ and coefficients c , the Algorithm 2 solves a specific version of Equation (15):

$$(\gamma^*, c^*) = \operatorname{argmin}_{\gamma, c} \Delta c \mathcal{P}_\gamma.$$

The optimal Hamiltonian, denoted as

$$\hat{\mathcal{H}}_{best} = c^* \mathcal{P}_{\gamma^*},$$

is thus defined by its minimal delta distance, not by global physical optimality. In this context, $\hat{\mathcal{H}}_{best}$ is the best fit of the biomarker under search constraints; it is the Hamiltonian that statistically explains the patient-induced quantum state with the greatest fidelity.

Because the Hamiltonian is constrained to $c\mathcal{P}_\gamma$, the optimal coefficient is obtained analytically:

$$c^* = \frac{\langle \Psi_{patient} | \mathcal{P}_\gamma | \Psi_{patient} \rangle}{\lambda_{\min}(\mathcal{P}_\gamma)}$$

where $\lambda_{\min}(\mathcal{P}_\gamma)$ is the lowest eigenvalue of the Pauli operator. Thus, coefficients reflect the empirical correlation between the two clinical features encoded by the operator.

3.1.1. Stable Patient Scenario

This scenario establishes the clinical baseline, representing a patient with maximum resilience (no TB: $TB_1 = 0$; no drug resistance: $GR_1 = 0$). The patient’s quantum state $|\psi_{stable}\rangle$ exhibits an ideal outcome state ($\theta_2 = 0, \phi_2 = 0$) despite the active internal dynamics on qubit 1 ($\theta_1, \phi_1 = 0$). The search algorithm identified the simplest Hamiltonian, certifying the decoupling of the two-qubit system:

$$\hat{\mathcal{H}}_{best} = 0.2171959 \cdot I \otimes I.$$

This Identity operator serves as the quantum stability identifier. The analytical delta ($\Delta = 0$) confirms that $|\psi_{stable}\rangle$ is an exact eigen-state of $\hat{\mathcal{H}}_{best}$, resulting in a constant intrinsic energy $E_{\min} = 0.2171959$.

The result $I \otimes I$ implies that the internal dynamics (qubit 1) do not structurally affect the ideal clinical outcome (qubit 2). The constant coefficient C is the system’s baseline energy, derived from the correlation of the highest competing non-Identity term ($Z_n \otimes Z_{str2_7}$), which is successfully neutralized by the patient’s state.

The principal clinical advantage of the $I \otimes I$ Hamiltonian is its zero measurement cost. Following standard VQE shot budgeting, the expected number of measurements n_{expect} is proportional to the variance of the Pauli terms $Var(P_\gamma)$ [21]. Since $Var(I \otimes I) = 0$, the prognostic shot budget is $n_{expect} = 0$. This demonstrates the ultimate cost-efficiency achieved when the state of a patient certifies maximum clinical resilience Table 1.

Table 1. Data-driven clinical Hamiltonian for the stable patient scenario. The inverse VQE search identifies the Hamiltonian $\hat{\mathcal{H}}_{best}$ that minimizes the clinical delta metric Δ , ensuring that the patient-encoded quantum state $|\psi_{stable}\rangle$ acts as its ground state. In this scenario, the search converges to the Identity operator, indicating complete structural decoupling between features and establishing a constant intrinsic health energy E_{min} . The coefficient C is not analytically “optimal” but is the *data-derived* value that best neutralizes the competing $Z \otimes Z$ term under the search criterion. The table summarizes the discovered Hamiltonian, patient encodings, latent risk structure, and the clinical implications of the resulting zero-variance operator.

Metric	Value	Search Parameter	Clinical Implication
Optimal Hamiltonian	$0.2171959 \cdot I \otimes I$	1 term	Defines the Intrinsic Base Energy of Health (E_{min})
Analytic Delta (Δ)	0	N/A	Exact Eigen-space Alignment
Min Delta (Comp.)	2.78×10^{-17}	100,000 shots	Minimal computational search error
Qubit 1 Encoding	$\theta_1 \approx 3.04, \phi_1 \approx 3.64$	N/A	Active dynamics (non zero VL/CD4 differentials)
Qubit 2 Encoding	$\theta_2 = 0, \phi_2 = 0$	N/A	Ideal outcome state ($ 0\rangle$)
Latent Risk Coefficient C	+0.2171959	Cor(n, str2_7)	Energy derived from the rejected $Z \otimes Z$ risk term
Feature Z_1	n (Phenotypic Robustness)	n=119	System inertia embedded in the clinical cohort
Feature Z_2	str_2_7 (Z gate strategy)	N/A	Most aggressive (high-risk) viral strategy
Prognostic Shots n_{expect}	0	$Var(I \otimes I) = 0$	Zero Measurement Cost (Eliminates Finite Sampling Error)
Shots for Latent ZZ risk ($\epsilon = 10^{-3}$)	$\approx 47,710$	$n_{expect} = C^2/\epsilon^2$	Budget required to measure the neutralized competing term

3.1.2. TB Co-Infected Patient Scenario

This scenario examines a patient with TB co-infection (TB_1=1). The patient’s state $|\psi_{TB}\rangle$ is characterized by structural stress, evidenced by a high-risk outcome qubit encoding ($\theta_2 \approx 2\pi, \phi_2 = 0$) despite the current stable viral load dynamics (vlogs_diff_mean ≈ 0.010).

The optimal search algorithm identified a single entanglement operator that certifies this instability:

$$\hat{\mathcal{H}}_{best} = -0.1041764 \cdot Z_{str1_3} \otimes Z_{vlogs_diff_mean}.$$

This operator $Z \otimes Z$, with $E_{min} \approx -0.1041764$, is the quantum strategy interaction identifier Table 2. The low analytical delta ($\Delta \approx 2.78 \times 10^{-5}$) confirms a strong alignment between the patient state and this Hamiltonian’s eigen-state, meaning structural instability by establishing a measurable, inverse correlated entanglement between key features.

The Hamiltonian functions as a feature reduction filter, isolating the crucial interaction between Z_{str1_3} (the latent high-risk Hadamard strategy for the low replicating viral phenotype) and $Z_{vlogs_diff_mean}$ (the active biomarker of viral load dynamics).

Unlike the $I \otimes I$ scenario, the $Z \otimes Z$ operator requires a non zero shot budget to overcome the finite sampling error (FSE) inherent to quantum measurement.

The 10^7 shots used during the search phase represented an optimization cost to ensure that the Hamiltonian was reliably identified with minimal FSE. The prognostic shots are the operational cost for the final single measurement. The shots required for clinical grade precision ($\epsilon = 10^{-3}$) are calculated as follows:

$$n_{expect} = \frac{(-0.1041764)^2}{(10^{-3})^2} \approx 10853.$$

This cost quantifies the added complexity introduced by the co-infection. For resource constrained NISQ deployment, relaxing the tolerance to $\epsilon = 10^{-2}$ reduces the cost significantly:

$$n_{expect} = \frac{(-0.1041764)^2}{(10^{-2})^2} \approx 108$$

This demonstrates the critical trade-off between clinical precision and hardware budget [53].

Table 2. Data-driven clinical Hamiltonian for the TB co-infected scenario. The inverse VQE search identifies a negative $Z \otimes Z$ interaction as the dominant term, reflecting strong anti-correlated structure between the patient’s feature encoding and the underlying tuberculosis-associated instability. The coefficient C is the data-derived weight that best aligns the entangled patient state with this coupled Hamiltonian, revealing how strategic viral behavior and biomarker dynamics jointly shape the energetic landscape. The table summarizes the recovered Hamiltonian, feature encodings, latent interaction structure, and the hardware budget needed to resolve this clinically relevant ZZ coupling with high confidence.

Metric	Value	Search Parameter	Clinical Implication
Optimal Hamiltonian	$-0.1041764 \cdot Z \otimes Z$	1 term	Defines the Anti-Correlated Energy of Instability (E_{\min})
Analytic Delta (Δ)	2.78×10^{-5}	N/A	Good Approximate Eigenspace Alignment
Min Delta (Comp.)	3.47×10^{-5}	10^7 shots	High fidelity computational search convergence
Qubit 1 Encoding	$\theta_1 \approx 2.14, \phi_1 \approx 2.95$	N/A	Active dynamics (non zero VL/CD4 changes)
Qubit 2 Encoding	$\theta_2 \approx 2\pi, \phi_2 = 0$	N/A	High risk, non ideal outcome state.
Interaction Coefficient C	-0.1041764	$\text{Cor}(\text{str1_3}, \text{vlogs})$	Inverse correlation (Suppression of strategic risk by dynamics)
Feature Z_1	str1_3 (low replication’s H strategy)	$\text{str1_3}=0$ (low replication plays X)	Latent strategic risk factor
Feature Z_2	vlogs_diff_mean (viral load change)	≈ 0.010	Active biomarker coupled to the strategy
Prognostic Shots $n_{\text{expect}}(\epsilon = 10^{-3})$	$\approx 10,853$	ZZ variance ≈ 1	Budget for high assurance clinical prognosis
Prognostic Shots $n_{\text{expect}}(\epsilon = 10^{-4})$	$\approx 1,085,300$	$n_{\text{expect}} = C_2/\epsilon^2$	Budget for cost-efficient NISQ tracking

3.1.3. Drug Resistance (GR) Scenario

This scenario investigates a patient with resistance to treatment ($\text{GR}_1=1$). The patient’s state $|\psi_{GR}\rangle$ is highly unstable, characterized by a significant increase in viral load and a decrease in CD4. The encoding shows that the instability is mainly contained within the internal dynamics (qubit 1: $\theta_1 \approx 2.69, \phi_1 \approx 1.77$).

The optimal search identified an entanglement operator which defines the instability:

$$\hat{\mathcal{H}} = -0.102438 \cdot Z_{\text{str2_5}} \otimes Z_{\text{payoffs}}$$

This $Z \otimes Z$ operator is the quantum conditional payoff identifier. The highly accurate analytical delta ($\Delta \approx 1.11 \times 10^{-6}$) validates the Hamiltonian as the minimal energy operator, verifying a deep, inversely correlated structural instability linked to viral fitness (Table 3).

The Hamiltonian’s features are $Z_{\text{str2_5}}$ (the latent defensive X strategy for the high replicating phenotype) and Z_{payoffs} (the classical viral fitness score/reproductive capacity).

The negative coefficient ($C \approx -0.102$) shows that the patient’s drug resistance is defined by an inverse correlation in which the maximum classical survival capacity of the viral phenotype (Z_{payoffs}) suppresses its necessity to adopt a critical defensive strategy ($Z_{\text{str2_5}}$). This is an active strategic instability.

The prognosis requires measuring $\langle \hat{\mathcal{H}}_{\text{best}} \rangle$: a value near $E_{\min} \approx -0.102$ confirms the persistence of this conditional instability, while a shift towards zero (loss of inverse correlation) would signal a phase transition in the viral strategy, requiring intervention.

As in the TB scenario, the $Z \otimes Z$ operator requires a substantial budget for a high-assurance prognosis. The 10^7 search shots again serve as a high-fidelity optimization step. The operational budget is calculated for $\epsilon = 10^{-3}$:

$$n_{expect} = \frac{(-0.102438)^2}{(10^{-3})^2} \approx 10494.$$

This $\approx 10,500$ shot requirement quantifies the cost of measuring the instability introduced by drug resistance. For routine monitoring, relaxing the precision to $\epsilon = 10^{-2}$ drastically reduces the shot budget:

$$n_{expect} = \frac{(-0.102438)^2}{(10^{-2})^2} \approx 105.$$

This illustrates that the clinical Hamiltonian provides clinicians with a direct and quantifiable mechanism to choose the cost-effectiveness strategy based on the required risk tolerance ϵ .

Table 3. Data-driven clinical Hamiltonian for the drug-resistant (GR) scenario. The inverse VQE procedure converges to a negative $Z \otimes Z$ coupling, indicating a fitness-modulated antagonistic interaction between resistance-associated strategies and the patient’s encoded clinical state. The coefficient C represents the optimal data-driven weight that neutralizes competing fitness signals in the Hamiltonian search, capturing the energetic signature of drug resistance pressure. This table details the resulting Hamiltonian, patient-specific feature encodings, the latent strategy–fitness interaction, and the hardware requirements for reliably estimating the associated ZZ risk term in NISQ settings.

Metric	Value	Search Parameter	Clinical Implication
Optimal Hamiltonian	$-0.102438 \cdot Z \otimes Z$	1 term	Defines the Conditional Payoff Energy of Instability (E_{min})
Analytic Delta (Δ)	1.11×10^{-6}	N/A	High Approximate Eigenspace Alignment
Min Delta (Comp.)	1.73×10^{-4}	10^7 shots	Minimal computational search error
Qubit 1 Encoding	$\theta_1 \approx 2.69, \phi_1 \approx 1.77$	N/A	Unstable, active dynamics
Qubit 2 Encoding	$\theta_2 \approx 0, \phi_2 \approx 2\pi$	N/A	Minimal outcome state shift
Interaction Coefficient C	-0.102438	$Cor(str2_5, payoffs)$	Inverse Correlation (Strategy suppression vs. fitness)
Feature Z_1	$str2_5$ (high replication X strategy)	$str2_5=0$ (high replication plays I)	Latent defensive strategic threat
Feature Z_2	payoffs (Viral Fitness)	≈ -0.165	Active biomarker proxy for reproductive capacity
Prognostic Shots n_{expect} ($\epsilon = 10^{-3}$)	$\approx 10,494$	ZZ variance ≈ 1	Budget for high assurance clinical prognosis
Prognostic Shots n_{expect} ($\epsilon = 10^{-2}$)	≈ 105	$n_{expect} = C^2 / \epsilon^2$	Budget for cost-efficient NISQ tracking

3.2. The Clinical Outcome: VQE Energy Score for Prognosis

This section explains how the discovered Hamiltonian is used for prediction and introduces the core clinical output—the VQE energy score Algorithm 3. This score is a compact, single-value, interpretable metric derived from quantum simulation, which is the primary benefit over complex classical classifiers.

The multi-qubit VQE framework was deployed as a global binary classifier to predict two critical clinical statuses: genotypic drug resistance status ($GR_1=1$) and TB co-infection (TB_1). The algorithm utilizes a single-term Hamiltonian ($N = 1$) $\hat{H}^* = C_1 \cdot \hat{P}_1$, optimized over $M = 10,000$ iterations to minimize the loss function of the chosen classification. In both scenarios, the dataset is severely imbalanced, making balanced accuracy the key performance metric.

The target class, GR_1 , consists of 8 positive cases of 176 patients (prevalence $\approx 4.5\%$). The four features encoded in the two-qubit ansatz were vl_diff , cd_diff , TB_1 , and $nearest_payoff$.

The multi-qubit quantum classifier was optimized to minimize three distinct loss functions: simple error, balanced error, and penalized error. The performance of the optimal classifier ($\hat{\mathcal{H}}^*, \tau^*$) found for each loss function is presented in Table 4.

Table 4. Performance metrics for the optimal single-term VQE clinical Hamiltonian ($N = 1$) predicting genotypic drug resistance status (GR_1). The table reports accuracy, sensitivity, specificity, and balanced accuracy for classifiers optimized under three loss functions: simple, balanced, and penalized. Balanced accuracy highlights performance under severe class imbalance and identifies the best predictive model.

Loss Function	Accuracy	Sensitivity (TPR)	Specificity (TNR)	Balanced Accuracy
Simple Error	0.960	0.125	1.000	0.563
Balanced Error	0.909	0.875	0.923	0.899
Penalized Error	0.955	0.375	0.982	0.679

The classifier optimized with the balanced error loss function yielded the best performance, achieving a balanced accuracy of 89.9%. This model correctly identified seven out of the eight drug-resistant patients (Sensitivity = 0.875) while maintaining a high true negative rate (Specificity = 0.923). The poor performance of the simple error model (Sensitivity = 0.125) highlights the need to use loss functions that are robust to class imbalance in medical diagnostics.

The second task aimed to predict the co-infection status TB_1, which is even more imbalanced, with only 6 positive cases out of 176 patients (prevalence \approx 3.4%). For this prediction, the encoded features were modified to be v1_diff, cd_diff, GR_1 (using the final status as predictor) and nearest_payoff. The performance of the optimal classifier ($\hat{\mathcal{H}}^*, \tau^*$) for the prediction TB_1 is shown in Table 5.

Table 5. Performance metrics for the optimal single-term VQE clinical Hamiltonian ($N = 1$) predicting TB co-infection status (TB_1). Metrics include accuracy, sensitivity, specificity, and balanced accuracy for classifiers optimized using simple, balanced, and penalized loss functions. The balanced error loss function achieves the highest balanced accuracy, demonstrating robust prediction under extreme class imbalance.

Loss Function	Accuracy	Sensitivity (TPR)	Specificity (TNR)	Balanced Accuracy
Simple Error	0.972	0.167	1.000	0.583
Balanced Error	0.881	0.667	0.894	0.827
Penalized Error	0.977	0.333	1.000	0.667

Similarly to the scenario GR_1, the classifier optimized for simple error yielded a misleading accuracy > 0.97 , but had low sensitivity (0.167), correctly identifying only one of six TB-co-infected patients. The penalized error also failed to achieve robust sensitivity (Sensitivity = 0.333).

The best predictive performance for TB_1 was achieved using the balanced error loss function, with a balanced accuracy of 82.7%. This model successfully identified four out of six co-infected patients (Sensitivity = 0.667) while maintaining a high specificity of 0.894.

Although the balanced accuracy for TB_1 (82.7%) is lower than that for GR_1 (89.9%), this result is highly significant given the extreme imbalance and the difficulty of predicting TB co-infection using standard clinical markers. The single-term quantum classifier demonstrates an ability to leverage quantum encoded features and their correlations to distinguish between the minority and majority classes with substantial predictive power.

3.3. Comparative Analysis with Classical Baselines

To rigorously evaluate the strength of the VQE approach, its performance was benchmarked against standard classical models (LogReg and MLP) trained using inverse class weights to mimic the “Balanced Error” optimization and mitigate class imbalance. The results of these classical baselines are summarized in Table 6.

Table 6. Comparative performance of weighted classical baselines (LogReg and MLP) versus the optimal single-term VQE model (Balanced Error) for predicting GR_1 and TB_1 clinical statuses. Metrics include accuracy, sensitivity, specificity, and balanced accuracy. The table highlights the limitations of classical models under severe class imbalance and the ability of the VQE classifier to achieve high balanced accuracy by correctly identifying minority-class cases while preserving specificity.

Target	Model	Accuracy	Sensitivity (TPR)	Specificity (TNR)	Balanced Accuracy
GR_1	LogReg (Weighted)	0.906	0.000	0.960	0.480
	MLP (Weighted)	0.925	0.000	0.980	0.490
	VQE (Balanced Error)	0.950	0.840	0.960	0.900
TB_1	LogReg (Weighted)	0.774	0.333	0.800	0.567
	MLP (Weighted)	0.623	0.667	0.620	0.644
	VQE (Balanced Error)	0.881	0.667	0.894	0.827

The classical models demonstrated a high overall accuracy ($\geq 90\%$ for GR_1 and between 62.3% and 77.4% for TB_1) but showed a significant bias toward the majority class for GR_1. Specifically, for the variable GR_1, both LogReg and MLP failed to correctly identify any patient in the minority class (Sensitivity = 0.000), resulting in a balanced accuracy below the random chance level (50%).

For the TB_1 variable, the classical baselines showed better, but still limited, performance. The LogReg achieved a balanced accuracy of 56.7%, while the MLP achieved 64.4%. Although the MLP attained the same sensitivity (0.667) as the VQE, this was achieved by significantly sacrificing specificity (0.620). In contrast, the VQE model with balance error loss significantly outperformed all classical methods, achieving a balanced accuracy of 82.7% while maintaining both high sensitivity (0.667) and high specificity (0.894). These results suggest that the VQE approach effectively utilizes complex clinical and quantum-derived features to better resolve the minority class without compromising the correct classification of the majority class, a key requirement for reliable clinical risk scoring.

Although the balanced accuracy for TB_1 (82.7%) is lower than that for GR_1 (90.0%), the consistent and significant out-performance of the VQE model over all classical baselines, both unbalanced targets, validates the utility of quantum-derived features in this clinical domain. The VQE’s ability to maintain high sensitivity while preserving high specificity suggests that the feature encoding and Hamiltonian learning successfully capture latent, clinically relevant correlations that are inaccessible to conventional models, providing a superior basis for reliable clinical risk scoring.

3.4. Computational Complexity Analysis of Quantum Algorithms

The computational cost of our proposed Inverse VQE method is distributed between quantum operations (state preparation; measurement) and classical optimization (Hamiltonian search; classification). We provide an asymptotically rigorous breakdown below, defining the key variables:

- N : Number of qubits (encoding $2N$ patient features in this study).
- L : Maximum number of Pauli terms in a trial Hamiltonian (L_{best} for the final Hamiltonian).
- M : Total number of patients in the dataset.
- T : Total number of optimization iterations (for the classical search loop).
- S : Number of measurement shots used for expectation value estimation.

The Algorithm 1 defines the preparation of the patient-encoded quantum state $|\psi_{patient}\rangle$ and is executed once per patient for the calculation of the energy score. Its complexity is mainly driven by the number of gates in the fixed shallow data-encoding circuit (Table 7).

The Algorithm 2 is the central classical optimization routine that runs for T iterations on the dataset M to find the optimal Hamiltonian $\hat{\mathcal{H}}_{best}$ (Table 8).

The Algorithm 3 uses the final $\hat{\mathcal{H}}_{best}$ to classify all M samples by calculating their VQE energy score E_{score} . This is the inference phase (Table 9).

The overall complexity reveals critical trade-offs for future scalability. The quantum execution time (state preparation and measurement for one term) scales $\mathbf{O}(\text{poly}(N))$ on a quantum device and linearly with the measurement shots S . The classification inference phase (Algorithm 3) scales linearly with the size of the dataset M and the number of terms L_{best} , validating its practicality to score large-scale patients once the Hamiltonian is found. The Hamiltonian search phase (Algorithm 2) has a dependence $\mathbf{O}(8^N)$ due to the classical computation of the minimum eigenvalue $E_{min}(\hat{\mathcal{H}}_{trial})$. This is the primary classical bottleneck and limits the current method to low-qubit systems ($N \leq 4$).

For small clinically relevant datasets ($M \sim 100\text{--}200$) and a small number of qubits ($N = 2$), the inverse VQE is computationally practical on classical simulators, while maintaining strong predictive power and interpretable results.

Table 7. Computational complexity of Algorithm 1. The complexity is polynomial in the number of qubits (N) on a quantum device due to the shallow circuit depth, making state preparation scalable.

Component	Description	Complexity	Note
Classical Pre-processing	Feature normalization and angle calculation (θ_i).	$O(N)$	Linear in the number of features ($2N$) encoded.
Quantum State Preparation	Consists of $2N$ rotational gates and $N - 1$ CNOT gates (for entanglement).	$O(\text{poly}(N))$	Exponential in N ($O(2^N)$) using state-vector simulation, but polynomial in N on a real quantum device since the circuit depth is $O(N)$.
Total Complexity (Per Patient)	Time to prepare one $ \psi_{patient}\rangle$ state.	$O(\text{poly}(N))$	Dominated by the shallow encoding circuit depth.

Table 8. Computational complexity of Algorithm 2. The $O(8^N)$ term (minimum eigenvalue calculation) is the primary classical bottleneck, limiting scalability in the number of qubits (N).

Component	Description	Complexity	Note
Outer Loop	Total number of optimization iterations over the search space.	$O(T)$	T is typically small (e.g., $T \sim 1000$ to 2000).
Trial Hamiltonian generation	Construct $\hat{\mathcal{H}}_{trial} = \sum_j c_j P_j$ (including correlation computation).	$O(L \cdot M)$	Dominated by calculating L feature correlation coefficients over the M patients.
Expectation Value Calculation	Calculating $\langle \psi_{patient} \hat{\mathcal{H}}_{trial} \psi_{patient} \rangle$. $\hat{\mathcal{H}}_{trial}$ has L terms.	$O(L \cdot S)$	Uses sampling for L terms per patient run.
Minimum Eigenvalue Calculation	Finding $E_{min}(\hat{\mathcal{H}}_{trial})$ classically for the $2^N \times 2^N$ matrix.	$O(8^N)$	For a general matrix, this is $O((2^N)^3)$. For the 2-qubit case ($N = 2$), this is $O(4^3)$, which is constant and fast.
Total Complexity	Combining T iterations over M training patients.	$O(T \cdot M \cdot (L \cdot S + 8^N))$	The 8^N term is the ultimate bottleneck for large N . However, the current $N = 2$ implementation runs polynomially in the data size M and iterations T .

Table 9. Computational complexity of Algorithm 3. The total complexity $O(M \cdot L_{\text{best}} \cdot S)$ scales linearly with the dataset size (M), demonstrating efficient inference for large clinical cohorts.

Component	Description	Complexity	Note
Score Calculation (Per Patient)	Calculating $E_{\text{score}} = \langle \psi_{\text{patient}} \hat{\mathcal{H}}_{\text{best}} \psi_{\text{patient}} \rangle$. $\hat{\mathcal{H}}_{\text{best}}$ has L_{best} terms.	$O(L_{\text{best}} \cdot S)$	Dominated by L_{best} expectation value estimations, each requiring a quantum run.
Total Complexity (Inference)	Calculating the score for all M patients in the dataset.	$O(M \cdot L_{\text{best}} \cdot S)$	This complexity is linear in the size of the dataset (M) and the number of terms L_{best} , making it highly desirable for large clinical datasets.

4. Discussion

The variational quantum approach, extended here to a multi-qubit framework, has demonstrated significant potential as an interpretable feature discovery and classification tool for complex clinical data, specifically in the highly challenging domain of HIV/TB co-infection and drug resistance.

The present work successfully extends the methodology proposed in previous single-qubit analyses, which were limited to encoding only two clinical features into the parameters of a state $|\psi(\theta, \phi)\rangle$ [49]. Although the single-qubit model validated the core idea that the ground state of a clinically informed Hamiltonian can act as a classification boundary, its practical usefulness was constrained by the limited descriptive capacity of a single qubit. With only two amplitudes, the model cannot encode or disentangle the multivariate, nonlinear relationships that characterize real biological systems.

The transition to a multi-qubit ansatz with an entangling *CNOT* gate (as shown in Algorithm 2) addresses this limitation. By encoding four clinical features (*v1_diff*, *cd_diff*, *nearest_payoff*, and *TB_1* or *GR_1*) in state $|\psi(\theta_1, \phi_1, \theta_2, \phi_2)\rangle$, the model explores a much richer Hilbert space. The introduction of the entangling *CNOT* gate is crucial, as it allows the quantum state to encode non-linear correlations between features, which are believed to be essential for accurate prediction of multiple factor diseases and drug resistance. This architectural expansion is the critical step that allows the identification of a single-term Pauli operator, $\hat{\mathcal{H}}^* = C_1 \cdot \hat{P}_1$, that acts as a robust classifier.

The success of the multi-qubit model is deeply linked to the advanced feature engineering derived from the quantum game of infectious phenotypes [52]. By converting complex biological dynamics, such as the competition of HIV-1 phenotypes, into single clinically meaningful, one-hot encoded variables (e.g., *nearest_payoff*, *str1_3*, or *str2_5*), the model achieves two primary goals. The first is interpretable through dimensionality reduction. This abstraction limits the effective feature space, enabling the optimal Hamiltonian to be found as a minimal single-term operator. A Hamiltonian like $-0.1041764 \cdot Z \otimes Z$ (for the TB scenario) or $0.2171959 \cdot I \otimes I$ (for the stable patient) is highly interpretable, as its single term and coefficient immediately indicate the key clinical biomarker combination driving the patient’s state towards the minimal energy. The second goal relates to the identification of clinical biomarkers. The results of the multi-qubit model (Tables 1–3) demonstrate how the model pinpoints the specific, most relevant pair of clinical and quantum derived features that define the patient’s state (i.e., the variables coupled by the optimal Pauli operator \hat{P}_1).

A key aspect enabling the high performance of this model is the incorporation of a feature derived from quantum game theory: *nearest_payoff*. This is a powerful example of feature engineering where complex, multi-dimensional information (specifically, the competitive dynamics between HIV-1 phenotypes) is abstracted into a single, clinically relevant measure.

By including this one-hot encoded quantum strategy (as a quantum-derived feature), the total number of necessary clinical features for robust classification is kept low, which

in turn limits the complexity of the optimal Hamiltonian found by the optimal clinical Hamiltonian search algorithm. This reduction in the feature space for the Hamiltonian improves clinical interpretations, directly highlighting the most influential pairwise correlation for the target outcome, offering a clear interpretable biomarker. Fewer terms in the Hamiltonian lead to fewer measurements required to estimate the expectation value $\langle \psi_{patient} | \hat{\mathcal{H}}_{trial} | \psi_{patient} \rangle$, minimizing the effects of sampling error, a critical challenge in the NISQ era [53].

The classification results validated the efficacy of the multi-qubit model. Given the severe class imbalance in the dataset (positive cases were 4.5% for GR_1 and 3.4% for TB_1), the standard simple error loss function proved clinically misleading, favoring high accuracy and specificity at the expense of nearly zero sensitivity (e.g., Sensitivity = 0.125 for GR_1).

The optimization based on the balanced error loss function delivered highly relevant diagnostic results. For the status GR_1, a balanced accuracy of 89.9 successfully identified 7/8 drug-resistant patients (Sensitivity = 0.875). For the status TB_1, a balanced accuracy of 82.7, successfully identified 4/6 co-infected patients (Sensitivity = 0.667).

This performance confirms that the proposed quantum classifier can effectively distinguish between minority and majority classes by focusing on the underlying feature correlations, thereby providing a clinically valuable result that overcomes the limitations of traditional models in imbalanced settings.

Although performance metrics validate the hybrid quantum–classical methodology, it is essential to acknowledge specific limitations regarding the clinical cohort and data collection process. First, the dataset is derived exclusively from a single regional HIV clinic in Guerrero, Mexico. This homogeneity, while useful for controlling local treatment and epidemiological factors during training, inherently limits the generalization (external validity) of the final model. The precise numerical values of the optimal Hamiltonian coefficients may be specific to the regional strain variants or local clinical protocols and may not directly translate to cohorts from other geographical areas.

Secondly, the use of longitudinal clinical records (2018–2024) introduces potential selection biases. Patients who maintain consistent follow-up and remain in the cohort over extended periods tend to be more treatment-adherent or may represent a subpopulation with inherently better long term outcomes. Furthermore, the necessary data cleaning process, which excluded records lacking multiple longitudinal measurements, further restricted the final analytical base to those patients with documented temporal dynamics, potentially biasing the model toward compliant, multi-visit individuals.

The rigorous computational complexity analysis (Tables 7–9) highlights the crucial hybrid nature of the inverse computational cost of the VQE.

The inference phase Algorithm 3, which performs the classification in new patients, exhibits excellent scaling: $\mathcal{O}(M \cdot L_{best} \cdot S)$, making it linear in the size of the clinical cohort M . This confirms the method's efficiency for practical, large-scale deployment once the Hamiltonian is found.

In contrast, the Hamiltonian search phase Algorithm 2, which is performed only once during training, faces a significant classical bottleneck. This is due to the necessity of calculating the minimum eigenvalue, $E_{\min}(\hat{\mathcal{H}}_{trial})$, which requires classical diagonalization and scales exponentially as $\mathcal{O}(8^N)$. Although the current $N = 2$ model remains computationally practical, this exponential scaling is the primary factor limiting the current methodology to low-qubit systems. Future work must address this classical component to enable a fully scalable, high-qubit search.

To address these limitations, our immediate future work will focus on external validation of the trained model. The next step is to test the structural validity of the quan-

tum biomarker discovered, which means the specific pair of clinical features and the Pauli operator that produces the optimal classification (e.g., $Z \otimes Z$ coupling `v1_diff` and `nearest_payoff` for TB). The plan is to apply the trained classifier to distinct independent cohorts from other states in Mexico or from international public repositories, such as those made available through the NIH. This external validation will confirm whether the identified feature relationships (\hat{P}_i) are universal clinical biomarkers or are merely artifacts of the local training cohort, thus establishing the model's true external reliability for precision medicine applications.

The successful application of this multi-qubit entangling ansatz to a real world imbalanced biomedical classification problem is a powerful demonstration of the potential of variational quantum algorithms in the NISQ era [4]. The model leverages the strengths of quantum computation, the capacity for expressive state preparation, and the efficient estimation of Hamiltonian expectation values, while keeping the circuit depth and the number of required qubits minimal.

For addressing the computational bottleneck, future work must focus on mitigating the $\mathcal{O}(8^N)$ classical diagonalization step to enable a fully scalable, high-qubit search, like replacing the classical diagonalization with a QPE or a variant of VQE itself to find E_{\min} directly on a quantum device, or to develop novel classical heuristics that bypass the full diagonalization, such as using a low-rank approximation of $\hat{\mathcal{H}}$ or focusing the search to highly sparse, pre-selected Hamiltonians.

The R function `classify_global_GR_multi_qubit` of Appendix A.1 implements the core training loop for the inverse VQE. The current parallelization strategy is well chosen, but has clear limitations and opportunities for acceleration.

The only explicitly parallelized section of the routine is Step 2: Calculate Scores in Parallel, using the R function `mclapply`. Each patient's score calculation is independent, making it an ideal candidate for parallel execution across multiple CPU cores. For a single patient, the calculation was iterated sequentially through each Hamiltonian term (k in 1 to `num_hamiltonian_terms`). Each term requires independent quantum simulation and measurement.

To scale beyond the limitations of local multi-core CPUs, the focus should shift to optimizing this inner loop. The state-vector simulation in the helper function, which relies heavily on matrix multiplication, is highly amenable to Multi-GPU acceleration. This moves the main linear algebra burden from the CPU to the GPU, offering possible orders of magnitude acceleration in the $\mathcal{O}(2^N)$ component of the simulation time, allowing faster processing of the entire patient batch.

As quantum hardware continues to improve, the move from this two-qubit system to models with $n > 2$ qubits will naturally pave the way for further, more accurate models. Higher-qubit ansatzes will allow the encoding of even more clinical and omics features, while the fundamental clinical Hamiltonian principle will continue to serve as a powerful engine for identifying optimal, interpretable, and parsimonious biomarker combinations. This hybrid quantum-classical approach thus represents a promising direction for achieving the goal of precision medicine through quantum machine learning.

Author Contributions: Conceptualization; methodology; software; validation; formal analysis; investigation; resources; data curation; writing—original draft preparation; writing—review and editing; visualization; supervision; project administration; J.P.A.G., M.S.A. and O.M. All authors have read and agreed to the published version of the manuscript.

Funding: Instituto Politécnico Nacional SIP20253768, Secretaría de Ciencia, Humanidades, Tecnología e Innovación CF-2023-I-108.

Institutional Review Board Statement: This study was conducted in accordance with the Declaration of Helsinki and approved by the Institutional Review Board (Ethics Committee) of the Hospital General Regional No. 1 “Vicente Guerrero” del Instituto Mexicano del Seguro Social (IMSS). Approval was granted by the Comité Local de Investigación en Salud 1102, with COFEPRIS registration number 17 CI 12 001 066, CONBIOÉTICA registration CONBIOÉTICA 12 CEI 002 2018082, and institutional registration code R-2023-1102-029. Approved Date: 26 June 2023.

Informed Consent Statement: Patient consent was waived due to the retrospective nature of this study and the use of anonymized clinical records. All data were depersonalized prior to analysis, with removal of names and demographic identifiers to ensure confidentiality. No identifiable patient information is included in the dataset.

Data Availability Statement: The anonymized dataset used in this study is distributed through the CRAN package `qvirus` and its associated dataset `qphen`. Both are publicly available under the MIT License with redistribution permitted subject to the terms of the included LICENSE file. Access is provided via CRAN <https://cran.r-project.org/package=qvirus> (accessed on 5 December 2025) and the GitHub repository <https://github.com/juan-acu/qvirus> (accessed on 5 December 2025). Redistribution of raw clinical records is restricted due to privacy and ethical considerations; only anonymized and derived datasets are shared.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

QML	Quantum Machine Learning
QSAR	Qualitative Structure-Activity Relationship
QNNs	Quantum Neural Networks
EHR	Electronic Health Records
VQE	Variational Quantum Eigensolver
NISQ	Noisy Intermediate-Scale Quantum
QPE	Quantum Phase Estimation
VQC	Variational Quantum Classifier
CNOT	Controlled-Not
QAOA	Quantum Approximate Optimization Algorithm
HIV	Human Immunodeficiency Virus
miRNA	MicroRNA
PrEP	Pre-Exposure Prophylaxis
SARS-CoV-2	Severe Acute Respiratory Syndrome Coronavirus 2
QLSTM	Quantum Long Short-Term Memory
LQNN	Liquid Neural Network
QPCA	Quantum Principal Component Analysis
LLMs	Large Language Models
TB	Tuberculosis
CRAN	Comprehensive R Archive Network
TP	True Positive
TN	True Negative
FN	False Negative
FP	False Positive
FSE	Finite Sampling Error
LogReg	Regularized Logistic Regression
MLP	Multi-Layer Perceptron
CV	Cross-Validation
AUCPR	Area Under the Precision–Recall Curve

Appendix A

The complete implementation of our multi-qubit VQE energy-score routine (Appendix A.1), together with all helper functions required to run the simulations (Appendix A.2). These routines correspond exactly to the pipeline used in our experiments, including state preparation, amplitude encoding, basis rotations for Pauli-term measurement, and the parallelized score-search procedure. All functions are implemented in R, leveraging the parallel and qvirus packages for quantum-state manipulation and multi-core execution.

Appendix A.1. Main VQE Energy Score Routine

```

classify_global_GR_multi_qubit <- function(data_corr , data_enc ,
target , threshold_grid , max_iter = 2000 ,
loss_fn = c("simple" , "balanced" , "penalized") , ncores = 2 ,
num_hamiltonian_terms = 4 , shots = 1000 ,
v1m , v1M , cdm , cdM , tbm , tbM , c4m , c4M) {

  # Parameter validation and setup
  loss_fn <- match.arg(loss_fn)
  feature_names <- names(data_corr)

  # Assumes create_pauli_matrices() is defined and available
  pauli_terms_available <- names(create_pauli_matrices())[ -1]
  # Excludes "II"

  min_error <- Inf
  best_threshold <- NULL
  final_pred <- rep(0 , length(target))
  best_hamiltonian_terms <- NULL
  # Initialize to avoid errors if min_error isnt updated

  for (i in 1:max_iter) {

    # 1. Select Hamiltonian (Pauli Terms and Coefficients)
    selected_paulis <- sample(pauli_terms_available ,
num_hamiltonian_terms , replace = FALSE)

    P_terms <- list()
    C_coeffs <- numeric(num_hamiltonian_terms)

    for (k in 1:num_hamiltonian_terms) {
      term <- selected_paulis[k]
      feature_pair <- sample(feature_names , 2 , replace = FALSE)
      # Calculate correlation on the full feature set (data_corr)
      coeff <- cor(data_corr [[ feature_pair [1]]] ,
data_corr [[ feature_pair [2]]] ,
use = "complete.obs")

      P_terms [[k]] <- list(term = term ,
feature_pair = feature_pair ,
coefficient = coeff)
      C_coeffs [k] <- coeff
    }
  }
}

```

```

}

# 2. Calculate Scores in Parallel
scores <- unlist(mclapply(1:nrow(data_enc), function(j) {

  # Step 2a: Encode the record into angles
  angles_df <- encode_record_multi(data_enc[j,],
  vlm, vLM, cdm, cdM, tbm, tbM, c4m, c4M)

  # Step 2b: Generate the ansatz
  qc_ansatz <- two_qubit_ansatz(angles_df$theta1_new,
  angles_df$phi1_new, angles_df$theta2_new,
  angles_df$phi2_new)

  current_score <- 0

  # Step 2c: Sum the expectation values for all Hamiltonian
  for (k in 1:num_hamiltonian_terms) {
    term <- selected_paulis[k]
    coeff <- C_coeffs[k]

    basis_changed_qc <- qc_ansatz

    # --- Basis Change ---

    # Qubit 0 (First character of the Pauli string)
    if (substr(term, 1, 1) == "X") {
      basis_changed_qc <- h.qc(basis_changed_qc, 0)
      # Apply Hadamard for X-measurement
    } else if (substr(term, 1, 1) == "Y") {
      basis_changed_qc <- sdag.qc(basis_changed_qc, 0)
      # Apply S_dagger
      basis_changed_qc <- h.qc(basis_changed_qc, 0)
      # Apply Hadamard for Y-measurement
    }

    # Qubit 1 (Second character of the Pauli string)
    if (substr(term, 2, 2) == "X") {
      basis_changed_qc <-
      h.qc(basis_changed_qc, 1)
      # Apply Hadamard for X-measurement
    } else if (substr(term, 2, 2) == "Y") {
      basis_changed_qc <-
      sdag.qc(basis_changed_qc, 1) # Apply S_dagger
      basis_changed_qc <- h.qc(basis_changed_qc, 1)
      # Apply Hadamard for Y-measurement
    }

    # Calculate the expectation value
    term_expectation <- coeff *

```

```

        two_qubit_pauli_expectation(basis_changed_qc, term,
        shots)
        current_score <- current_score + term_expectation
    }

    return(current_score)
}, mc.cores = ncores))

# 3. Evaluate grid of thresholds
errors <- sapply(threshold_grid, function(th) {
  pred <- ifelse(scores >= th, 1, 0)
  TP <- sum(pred == 1 & target == 1)
  TN <- sum(pred == 0 & target == 0)
  FP <- sum(pred == 1 & target == 0)
  FN <- sum(pred == 0 & target == 1)

  if (loss_fn == "simple") {
    return(mean(pred != target))
  } else if (loss_fn == "balanced") {
    # Handles division by zero
    sens <- ifelse(TP + FN > 0, TP / (TP + FN), 0)
    spec <- ifelse(TN + FP > 0, TN / (TN + FP), 0)
    return(1 - 0.5 * (sens + spec))
  } else if (loss_fn == "penalized") {
    return((2 * FN + FP) / length(target))
  }
})

best_idx <- which.min(errors)
current_error <- errors[best_idx]

if (current_error < min_error) {
  min_error <- current_error
  best_threshold <- threshold_grid[best_idx]
  best_hamiltonian_terms <- P_terms
  final_pred <- ifelse(scores >= best_threshold, 1, 0)
}
}
return(list(
  min_error = min_error,
  best_hamiltonian_terms = best_hamiltonian_terms,
  best_threshold = best_threshold,
  predictions = final_pred,
  loss_fn = loss_fn
))

}
library(parallel)
library(qvirus)
qphen_data <- qphen

```

```

target_data <- qphen_data$TB_1
vlm <- min(qphen_data$vl_diff)
vM <- max(qphen_data$vl_diff)
cdm <- min(qphen_data$cd_diff)
cdM <- max(qphen_data$cd_diff)
tbm <- min(qphen_data$nearest_payoff)
tbM <- max(qphen_data$nearest_payoff)
c4m <- min(qphen_data$GR_1)
c4M <- max(qphen_data$GR_1)
#c4m <- min(qphen_data$TB_1)
#c4M <- max(qphen_data$TB_1)
data_corr <- qphen_data[, -c(1, 23)]
#data_corr <- qphen_data[, -c(1, 24)]
data_enc <- data_corr
threshold_grid <- seq(-1, 1, length.out = 1000)
result <- classify_global_GR_multi_qubit(
  data_corr = data_corr,
  data_enc = data_enc,
  target = target_data,
  threshold_grid = threshold_grid,
  max_iter = 10000,
  loss_fn = "balanced",
  ncores = 20,
  num_hamiltonian_terms = 1, shots = 10000,
  vlm = vlm, vM = vM, cdm = cdm, cdM = cdM,
  tbm = tbm, tbM = tbM, c4m = c4m, c4M = c4M
)

```

Appendix A.2. Helper Functions

```

create_pauli_matrices <- function() {
PauliX <- matrix(c(0, 1, 1, 0), 2, 2)
PauliY <- matrix(c(0, -1i, 1i, 0), 2, 2)
PauliZ <- matrix(c(1, 0, 0, -1), 2, 2)
Identity <- diag(2)
pauli_matrices <- list(
  "II" = kronecker(Identity, Identity),
  "XI" = kronecker(PauliX, Identity),
  "YI" = kronecker(PauliY, Identity),
  "ZI" = kronecker(PauliZ, Identity),
  "IX" = kronecker(Identity, PauliX),
  "IY" = kronecker(Identity, PauliY),
  "IZ" = kronecker(Identity, PauliZ),
  "XX" = kronecker(PauliX, PauliX),
  "YY" = kronecker(PauliY, PauliY),
  "ZZ" = kronecker(PauliZ, PauliZ),
  "XY" = kronecker(PauliX, PauliY),
  "XZ" = kronecker(PauliX, PauliZ),
  "YX" = kronecker(PauliY, PauliX),
  "ZX" = kronecker(PauliZ, PauliX),
  "YZ" = kronecker(PauliY, PauliZ),

```

```

    "ZY" = kronecker(PauliZ, PauliY)
  )
  return(pauli_matrices)
}

encode_record_multi <- function(new_record, vl_diff_min,
vl_diff_max, cd_diff_min, cd_diff_max,
tb_1_min, tb_1_max, np_min, np_max) {
  theta1_new <- (new_record$vl_diff - vl_diff_min) /
(vl_diff_max - vl_diff_min) * 2 * pi
  phi1_new <- (new_record$cd_diff - cd_diff_min) /
(cd_diff_max - cd_diff_min) * 2 * pi
  theta2_new <- (new_record$GR_1 - tb_1_min) /
(tb_1_max - tb_1_min) * 2 * pi
  #theta2_new <- (new_record$TB_1 - tb_1_min) /
#(tb_1_max - tb_1_min) * 2 * pi
  phi2_new <- (new_record$nearest_payoff - np_min) /
(np_max - np_min) * 2 * pi
  return(data.frame(theta1_new, phi1_new, theta2_new, phi2_new))
}

two_qubit_ansatz <- function(theta1, phi1, theta2, phi2) {
  qc_obj <- qc('two-qubit ansatz')
  qc_obj <- reg.qc(qc_obj, 2)
  qc_obj <- rx.qc(qc_obj, 0, theta1)
  qc_obj <- ry.qc(qc_obj, 0, phi1)
  qc_obj <- rx.qc(qc_obj, 1, theta2)
  qc_obj <- ry.qc(qc_obj, 1, phi2)
  qc_obj <- cx.qc(qc_obj, 0, 1)
  return(qc_obj)
}

rx <- function(x, ...) UseMethod("rx") # Generic
rx.default <- function(x, ...) .NotYetImplemented()
rx.qc <- function(qc_obj, idx, theta) {
  qc_obj$psi <- oapply1.State(qc_obj$psi, RotationX(theta), idx)
  # Assuming RotationX(theta) operator
  invisible(qc_obj)
  #return(qc_obj)
}

RotationX <- function(theta, d = 1) {
  X <- Operator(matrix(c(0, 1, 1, 0), nrow = 2, byrow = TRUE))
  mat <- cos(theta / 2) * diag(2) + 1i *
sin(theta / 2) * X$data
  return(Operator(kpow(Operator(mat), d)$data))
}

ry <- function(x, ...) UseMethod("ry") # Generic
ry.default <- function(x, ...) .NotYetImplemented()

```

```

ry.qc <- function(qc_obj, idx, theta) {
  qc_obj$psi <- oapply1.State(qc_obj$psi, RotationY(theta), idx)
  # Assuming RotationY(theta) operator
  invisible(qc_obj)
  #return(qc_obj)
}

RotationY <- function(theta, d = 1) {
  Y <- Operator(matrix(c(0, -1i, 1i, 0), nrow = 2,
    byrow = TRUE))
  mat <- cos(theta / 2) * diag(2) +
    1i * sin(theta / 2) * Y$data
  return(Operator(kpow(Operator(mat), d)$data))
}

cx <- function(x, ...) UseMethod("cx") # Generic
cx.default <- function(x, ...) .NotYetImplemented()
cx.qc <- function(self, ctl_idx, tgt_idx) {
  # For a CNOT gate, the gate_obj is implicitly PauliX()
  pauli_x_gate <- PauliX()

  # Call the fast application function on the state object
  self$psi <- oapplyc.State(self$psi, pauli_x_gate,
    ctl_idx, tgt_idx)

  invisible(self)
}

PauliX <- function(d = 1) {
  mat <- Operator(matrix(c(0, 1, 1, 0), nrow = 2,
    byrow = TRUE))
  return(Operator(kpow(mat, d)$data))
}

oapplyc <-
function(self, gate_obj, ctl_idx, tgt_idx) UseMethod("oapplyc")
oapplyc.default <-
function(self, gate_obj, ctl_idx, tgt_idx) .NotYetImplemented()
oapplyc.State <- function(self, gate_obj, ctl_idx, tgt_idx) {
  # Map 0-based user indices
  classical_ctl_idx <- self$nbits - ctl_idx - 1
  classical_tgt_idx <- self$nbits - tgt_idx - 1

  # Calculate bitmasks for control and target qubits
  ctl_mask <- bitwShiftL(1, classical_ctl_idx)

  # Get gate matrix elements
  g00 <- gate_obj$data[1, 1]
  g01 <- gate_obj$data[1, 2]
  g10 <- gate_obj$data[2, 1]

```

```

g11 <- gate_obj$data[2, 2]

# Create a copy of the state data to avoid race conditions
temp_data <- self$data

# Iterate and apply the 2x2 matrix to relevant amplitude pairs
for (i in 0:(bitwShiftL(1, self$nbits) - 1)) {
  # Check if the control qubit is 1 AND the target qubit is 0
  if (bitwAnd(i, ctl_mask) != 0 &&
      bitwAnd(i, bitwShiftL(1, classical_tgt_idx)) == 0) {
    idx1 <- i + 1
    idx2 <- bitwXor(i, bitwShiftL(1, classical_tgt_idx)) +
      1

    # Perform the matrix multiplication on the amplitude pair
    t1 <- g00 * temp_data[idx1] + g01 * temp_data[idx2]
    t2 <- g10 * temp_data[idx1] + g11 * temp_data[idx2]

    self$data[idx1] <- t1
    self$data[idx2] <- t2
  }
}
invisible(self)
}

oapply1 <- function(self, gate_obj, index) UseMethod("oapply1")
oapply1.default <-
function(self, gate_obj, index) .NotYetImplemented()
oapply1.State <- function(self, gate_obj, index) {
  # To maintain qubit ordering, index needs to be reversed.
  # Qubit 0 is the topmost (MSB)
  # classical bit corresponds to the user's 'index'.
  classical_bit_idx <- self$nbits - index - 1

  pow_2_index <- bitwShiftL(1, classical_bit_idx)
  # Equivalent to 2^classical_bit_idx

  # Extract gate elements (R matrices are 1-indexed)
  g00 <- gate_obj$data[1, 1]
  g01 <- gate_obj$data[1, 2]
  g10 <- gate_obj$data[2, 1]
  g11 <- gate_obj$data[2, 2]

  total_elements <- bitwShiftL(1, self$nbits)
  # Equivalent to 2^self$nbits

  # Loop through the state vector, applying the gate efficiently
  # Blocks of elements that are not affected by this qubit
  for (g_val in seq(from = 0, to = total_elements - 1,
    by = bitwShiftL(1, classical_bit_idx + 1))) {

```

```

# The inner loop iterates through elements within each block
for (i_val in seq(from = g_val,
to = g_val + pow_2_index - 1)) {
  # R uses 1-based indexing for vectors
  idx_current <- i_val + 1
  idx_swap <- i_val + pow_2_index + 1

  # Perform the transformation
  t1 <- g00 * self$data[idx_current] +
g01 * self$data[idx_swap]
  t2 <- g10 * self$data[idx_current] +
g11 * self$data[idx_swap]

  self$data[idx_current] <- t1
  self$data[idx_swap] <- t2
}
}
invisible(self) # Return the modified object
}

sdag <- function(x, ...) UseMethod("sdag") # Generic
sdag.default <- function(x, ...) .NotYetImplemented()
sdag.qc <- function(qc_obj, idx) {
  # S_dagger matrix: [[1, 0], [0, -i]]
  mat <- matrix(c(1, 0, 0, -1i), nrow = 2, byrow = TRUE)
  op <- Operator(mat)
  qc_obj$psi <- oapply1.State(qc_obj$psi, op, idx)
  return(qc_obj)
}

h <- function(x, ...) UseMethod("h") # Generic
h.default <- function(x, ...) .NotYetImplemented()
h.qc <- function(qc_obj, idx) {
  # Check if idx is a single value or a vector
  if (length(idx) == 1) {
    # If it's a single value, apply the gate directly.
    qc_obj$psi <- oapply1(qc_obj$psi, Hadamard(), idx)
  } else {
    # Index and apply the gate.
    # Note: R's loops will handle the order correctly
    for (qubit_idx in idx) {
      qc_obj$psi <- oapply1(qc_obj$psi, Hadamard(), qubit_idx)
    }
  }
  invisible(qc_obj)
}

Hadamard <- function(d = 1) {
  mat <- Operator(1 / sqrt(2) * matrix(c(1, 1, 1, -1),
nrow = 2, byrow = TRUE))

```

```

    return(Operator(kpow(mat, d)$data))
  }

two_qubit_pauli_expectation <-
function(qc_obj, pauli_op_string, shots = 1000) {
  psi <- qc_obj$psi$data
  probabilities <- abs(psi)^2
  eigenvalues_list <- list(
    "II" = c(1, 1, 1, 1),
    "XI" = c(1, 1, -1, -1),
    "YI" = c(1, 1, -1, -1),
    "ZI" = c(1, 1, -1, -1),
    "IX" = c(1, -1, 1, -1),
    "IY" = c(1, -1, 1, -1),
    "IZ" = c(1, -1, 1, -1),
    "XX" = c(1, -1, -1, 1),
    "YY" = c(1, -1, -1, 1),
    "ZZ" = c(1, -1, -1, 1),
    "XY" = c(-1, 1, 1, -1),
    "XZ" = c(1, -1, -1, 1),
    "YX" = c(-1, 1, 1, -1),
    "ZX" = c(1, -1, -1, 1),
    "YZ" = c(-1, 1, 1, -1),
    "ZY" = c(-1, 1, 1, -1)
  )
  eigenvalues <- eigenvalues_list[[pauli_op_string]]
  if (is.null(eigenvalues)) {
    stop("Invalid Pauli operator string.")
  }
  outcomes <- sample(0:3, size = shots, replace = TRUE,
    prob = probabilities)
  counts <- table(factor(outcomes, levels = 0:3))
  expectation_value <- sum(eigenvalues * (counts / shots))
  return(expectation_value)
}

qc <- function(name = NULL) {
  circuit <- list(
    name = name,
    psi = NULL, # Will be set by arange
    nbits = 0,
    .global_reg = 0
  )
  class(circuit) <- "qc"
  return(circuit)
}

reg <-
function(qc_obj, size, it, name) UseMethod("reg") # Generic
reg.default <-

```

```

function(qc_obj, size, it, name) .NotYetImplemented()
reg.qc <- function(qc_obj, size, it = 0, name = NULL) {
  new_reg <- Reg(size = size, it = it, global_reg = qc_obj$nbits)
  if (is.null(qc_obj$psi)) {
    qc_obj$psi <- psi(new_reg)
  } else {
    qc_obj$psi <- State(qc_obj$psi * psi(new_reg))
    # Kronecker product to combine states
  }
  qc_obj$nbits <- qc_obj$nbits + nbits(new_reg)
  return(qc_obj)
}

int_to_bitstring <- function(num, size) {
  if (num < 0 || size < 0)
    stop("Los numeros y el tam. deben ser no negativos.")
  if (num == 0) return(paste(rep(0, size), collapse = ""))
  binary_chars <- character(size)
  current_num <- num
  for (i in seq_len(size)) {
    bit <- current_num %% 2
    binary_chars[size - i + 1] <- as.character(bit)
    current_num <- floor(current_num / 2)
    if (current_num == 0 && i < size) {
      for (j in (size - i):1) binary_chars[j] <- "0"
      break
    }
  }
  paste(binary_chars, collapse = "")
}

Reg <- function(size, it = 0, global_reg = NULL) {
  if (length(size) != 1 || !is.numeric(size) ||
      size <= 0)
    stop("size debe ser un entero positivo.")
  size <- as.integer(size)
  if (is.null(global_reg) || !is.numeric(global_reg) ||
      global_reg < 0) {
    stop("global_reg debe ser un entero no negativo.")
  }
  global_reg <- as.integer(global_reg)
  reg_obj <- list(
    size = size,
    global_idx = seq(global_reg, global_reg + size - 1),
    val = rep(0, size)
  )
  reg_obj$val <- bits

  class(reg_obj) <- "Reg"
  return(reg_obj)
}

```

```

}

nbits <- function(x) UseMethod("nbits")
nbits.default <- function(x) .NotYetImplemented()
nbits.Reg <- function(x) {
  return(x$size)
}
nbits.qc <- function(x) {
  return(x$nbits)
}

psi <- function(x) UseMethod("psi")
psi.default <- function(x)
psi.Reg <- function(x) {
  do.call(bitstring.numeric, as.list(x$val))
}

bitstring <- function(x, ...) UseMethod("bitstring")
bitstring.default <- function(x, ...) .NotYetImplemented()
bitstring.qc <- function(qc_obj, ...) {
  bits <- list(...)
  if (length(bits) == 0) {
    stop("bitstring_method.qc requires at least one bit.")
  }
  new_bitstring_state <- bitstring.numeric(unlist(bits))
  if (is.null(qc_obj$psi)) {
    qc_obj$psi <- new_bitstring_state
  } else {
    qc_obj$psi <- qc_obj$psi * new_bitstring_state
  }
  qc_obj$nbits <- qc_obj$nbits + length(bits)
  return(qc_obj)
}
bitstring.numeric <- function(..., verbose = FALSE) {
  args <- list(...)
  if (length(args) == 1 && is.numeric(args[[1]])) {
    bits <- args[[1]]
  } else {
    bits <- unlist(args)
  }
  if (length(bits) == 0) {
    stop("Need to specify at least 1 qubit.")
  }

  nbits <- length(bits)
  idx <- bits2val(bits) + 1
  if (verbose) {
    cat("Bits:", bits, "\n")
    cat("Decimal index:", idx - 1, "\n")
  }
}

```

```

    tensor_size <- 2 ^ nbits
    amplitudes <- complex(real = rep(0, tensor_size),
                          imaginary = rep(0, tensor_size))
    amplitudes[idx] <- complex(real = 1, imaginary = 0)
    amplitudes_array <- array(amplitudes, dim = c(tensor_size, 1))
    tensor_obj <- Tensor(amplitudes_array)
    state_obj <- State(tensor_obj)
    return(state_obj)
  }

bits2val <- function(bits) {

  n <- length(bits)
  sum(bits * 2 ^ ((n - 1):0))
}

Tensor <- function(input_array) {
  tensor <- list(
    data = input_array
  )
  class(tensor) <- "Tensor"
  return(tensor)
}

`*.Tensor` <- function(x, y) {
  new_data <- kronecker(x$data, y$data)
  new_tensor <- Tensor(new_data)
  new_state <- Operator(new_tensor$data)
  return(new_state)
}

`+.Tensor` <- function(op1, op2) {
  return(Operator(op1$data + op2$data))
}

`-.Tensor` <- function(op1, op2) {
  return(Operator(op1$data - op2$data))
}

kpow <- function(state, n) UseMethod("kpow")
kpow.default <- function(state, n) .NotYetImplemented()
kpow.Tensor <- function(x, n) {
  if (!inherits(x, "Tensor")) stop("Must be a Tensor object.")
  if (n < 0) stop("Exponent must be non-negative.")
  if (n == 0) {
    identity <- diag(1)
    return(Tensor(identity))
  }
  result <- x
  if (n >= 2) {

```

```

    for (i in 2:n) {
      result <- result * x
    }
  return(result)
}

Operator <- function(input_array) {
  dims <- dim(input_array)
  nbits <- log2(dims[1])
  if (!isTRUE(all.equal(nbits, as.integer(nbits)))) {
    stop("Does not correspond to a valid qubit space.")
  }
  tensor_obj <- Tensor(input_array)
  operator <- list(
    data = tensor_obj$data ,
    nbits = as.integer(nbits)
  )
  class(operator) <- c("Operator", "Tensor")
  return(operator)
}

State <- function(tensor_obj) {
  if (!inherits(tensor_obj, "Tensor")) {
    stop("State must be created from a Tensor object.")
  }
  data_length <- length(tensor_obj$data)
  nbits <- log2(data_length)
  state <- tensor_obj
  state$nbits <- as.integer(nbits)
  class(state) <- c("State", class(tensor_obj))
  return(state)
}

```

References

- Schuld, M.; Petruccione, F. *Supervised Learning with Quantum Computers*; Springer: Berlin/Heidelberg, Germany, 2018.
- Wittek, P. *Quantum Machine Learning: What Quantum Computing Means to Data Mining*; Elsevier: Amsterdam, The Netherlands, 2014.
- Martín-Guerrero, J.D.; Lamata, L. Quantum Machine Learning: A tutorial. *Neurocomputing* **2022**, *470*, 457–461. [[CrossRef](#)]
- Zhang, Y.; Ni, Q. Recent Advances in Quantum Machine Learning. *Quantum Eng.* **2020**, *2*, e34. [[CrossRef](#)]
- Ciliberto, C.; Herbster, M.; Ialongo, A.D.; Pontil, M.; Rocchetto, A.; Severini, S.; Wossnig, L. Quantum machine learning: A classical perspective. *Proc. R. Soc.* **2018**, *474*, 20170551. [[CrossRef](#)] [[PubMed](#)]
- Cerezo, M.; Verdon, G.; Huang, H.Y.; Cincio, L.; Coles, P.J. Challenges and opportunities in quantum machine learning. *Nat. Comp. Sci.* **2022**, *2*, 567–576. [[CrossRef](#)]
- Huang, H.Y.; Broughton, M.; Mohseni, M.; Babbush, R.; Boixo, S.; Neven, H.; McClean, J.R. Power of data in quantum machine learning. *Nat. Commun.* **2021**, *12*, 2631. [[CrossRef](#)]
- Zeguendry, A.; Jarir, Z.; Quafafou, M. Quantum Machine Learning: A Review and Case Studies. *Entropy* **2023**, *25*, 287. [[CrossRef](#)]
- McClean, J.R.; Boixo, S.; Smelyanskiy, V.N.; Babbush, R.; Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **2018**, *9*, 4812. [[CrossRef](#)]
- Maheshwari, D.; Garcia-Zapirain, B.; Sierra-Sosa, D. Quantum Machine Learning Applications in the Biomedical Domain: A Systematic Review. *IEEE Access* **2022**, *10*, 80463–80484. [[CrossRef](#)]
- Baiardi, A.; Christandl, M.; Reiher, M. Quantum Computing for Molecular Biology. *ChemBioChem* **2024**, *24*, e202300120. [[CrossRef](#)]

12. Chiang, W.Y.; Kao, P.Y.; Yeh, T.L.; Yang, Y.C.; Lin, Y.C.; Zhavoronkov, A. Enhancing Drug Discovery: Quantum Machine Learning for QSAR Prediction with Incomplete Data. *arXiv* **2025**, arXiv:2501.13395. [[CrossRef](#)]
13. Li, R.Y.; Di Felice, R.; Rohs, R.; Lidar, D.A. Quantum annealing versus classical machine learning applied to a simplified computational biology problem. *npj Quantum Inf.* **2018**, *4*, 14. [[CrossRef](#)] [[PubMed](#)]
14. Li, R.Y.; Gujja, S.; Bajaj, S.R.; Gamel, O.E.; Cilfone, N.; Gulcher, J.R.; Lidar, D.A.; Chittenden, T.W. Quantum processor-inspired machine learning in the biomedical sciences. *Patterns* **2021**, *2*, 100246. [[CrossRef](#)] [[PubMed](#)]
15. Priyadharshini, M.; Raju, B.D.; Banu, A.F.; Kumar, P.J.; Muruges, V.; Rybin, O. A quantum machine learning framework for predicting drug sensitivity in multiple myeloma using proteomic data. *Sci. Rep.* **2025**, *15*, 26553. [[CrossRef](#)] [[PubMed](#)]
16. Prajapati, J.B.; Paliwal, H.; Prajapati, B.G.; Saikia, S.; Pandey, R. Quantum Machine Learning in Prediction of Breast Cancer. In *Quantum Computing: A Shift from Bits to Qubits. Studies in Computational Intelligence*; Pandey, R., Srivastava, N., Singh, N.K., Tyagi, K., Eds.; Springer: Singapore, 2023; pp. 351–382.
17. Singh, N.; Pokhrel, S.R. Modeling Quantum Machine Learning for Genomic Data Analysis. *arXiv* **2025**, arXiv:2501.08193. [[CrossRef](#)]
18. Pomarico, D.; Fanizzi, A.; Amoroso, N.; Bellotti, R.; Biafora, A.; Bove, S.; Didonna, V.; La Forgia, D.; Pastena, M.I.; Tamborra, P.; et al. A Proposal of Quantum-Inspired Machine Learning for Medical Purposes: An Application Case. *Mathematics* **2021**, *9*, 410. [[CrossRef](#)]
19. Flöther, F.F.; Blankenberg, D.; Demidik, M.; Jansen, K.; Krishnakumar, R.; Krishnakumar, R.; Laanait, N.; Parida, L.; Saab, C.; Utros, F. How quantum computing can enhance biomarker discovery for multi-factorial diseases. *Patterns* **2025**, *6*, 101236. [[CrossRef](#)]
20. Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.H.; Zhou, X.Q.; Love, P.J.; Aspuru-Guzik, A.; O'Brien, J.L. A variational eigenvalue solver on a photonic quantum processor. *Nat. Commun.* **2014**, *5*, 4213. [[CrossRef](#)]
21. McClean, J.R.; Romero, J.; Babbush, R.; Aspuru-Guzik, A. The theory of variational hybrid quantum-classical algorithms. *New J. Phys.* **2016**, *18*, 023023. [[CrossRef](#)]
22. Kuroiwa, K.; Nakagawa, Y.O. Penalty methods for a variational quantum eigensolver. *Phys. Rev. Res.* **2021**, *3*, 013197. [[CrossRef](#)]
23. Huggins, W.J.; Lee, J.; Baek, U.; O'Gorman, B.; Whaley, K.B. A non-orthogonal variational quantum eigensolver. *New J. Phys.* **2020**, *22*, 073009. [[CrossRef](#)]
24. Zhang, Y.; Cincio, L.; Negre, C.F.A.; Czarnek, P.; Coles, P.J.; Anisimov, P.M.; Mniszewski, S.M.; Tretiak, S.; Dub, P.A. Variational quantum eigensolver with reduced circuit complexity. *npj Quantum Inf.* **2022**, *8*, 96. [[CrossRef](#)]
25. Ferguson, R.R.; Dellantonio, L.; Al Balushi, A.; Jansen, K.; Dür, W.; Muschik, C.A. A measurement-based variational quantum eigensolver. *Phys. Rev. Lett.* **2021**, *126*, 220501. [[CrossRef](#)] [[PubMed](#)]
26. Liu, J.G.; Zhang, Y.H.; Wan, Y.; Wang, L. Variational quantum eigensolver with fewer qubits. *Phys. Rev. Res.* **2021**, *1*, 023025. [[CrossRef](#)]
27. Wang, D.; Higgott, O.; Brierley, S. Accelerated Variational Quantum Eigensolver. *Phys. Rev. Lett.* **2019**, *122*, 140504. [[CrossRef](#)]
28. Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J.M.; Gambetta, J.M. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **2017**, *549*, 242–246. [[CrossRef](#)]
29. Tilly, J.; Chen, H.; Cao, S.; Picozzi, D.; Setia, K.; Li, Y.; Grant, E.; Wossnig, L.; Rungger, I.; Booth, G.H.; et al. The Variational Quantum Eigensolver: A review of methods and best practices. *Phys. Rep.* **2022**, *986*, 1–128. [[CrossRef](#)]
30. Narnavaram, K.M.; Lo, D.C.T.; Zhang, X.; Deng, B. Quantum Synthetic Molecular Dynamics: Advanced Medical Innovations through Entangled Phenomena of Nucleic Acids. *EasyChair Preprint* **2025**, in press.
31. Villalba-Díez, J. Quantum drug discovery: A hybrid quantum graph neural network–variational quantum eigensolver framework for serine neutralization. *Eur. Phys. J. D* **2025**, *79*, 81. [[CrossRef](#)]
32. Roosan, D.; Khan, R.; Nirzhor, S.; Khou, T.; Hai, F. Classifying Hotspots Mutations for Biosimulation with Quantum Neural Networks and Variational Quantum Eigensolver. *arXiv* **2025**, arXiv:2507.00072. [[CrossRef](#)]
33. Doga, W.R.; Bose, A.; Sahin, M.E.; Bettencourt-Silva, J.; Pham, A.; Kim, E.; Andress, A.; Saxena, S.; Parida, L.; Robertus, J.L.; et al. How can quantum computing be applied in clinical trial design and optimization? *Trends Pharmacol. Sci.* **2024**, *45*, 880–891. [[CrossRef](#)]
34. Khan, W.R.; Kamran, M.A.; Khan, M.U.; Ibrahim, M.U.; Kim, K.S.; Ali, M.U. Diabetes Prediction Using an Optimized Variational Quantum Classifier. *Int. J. Intell. Syst.* **2025**, *2025*, 1351522. [[CrossRef](#)]
35. Noor, F.; Nabi, G. Quantum-Enhanced Feature Selection and Classification for Asthma Diagnosis Using a Variational Quantum Classifier. *Preprints* **2025**, 2025081219. [[CrossRef](#)]
36. Yadalam, P.K.; Natarajan, P.M.; Saeed, M.H.; Ardila, C.M. Variational Approaches for Drug-Disease-Gene Links in Periodontal Inflammation. *Int. Dent. J.* **2025**, *75*, 185–194. [[CrossRef](#)] [[PubMed](#)]
37. Rajput, N.K.; Bansal, R. Quantum State Preparation for Medical Data: Comprehensive Methods, Implementation Challenges, and Clinical Prospects. *arXiv* **2025**, arXiv:2508.05063. [[CrossRef](#)]
38. Basu, S.; Born, J.; Bose, A.; Capponi, S.; Chalkia, D.; Chan, T.A.; Doga, H.; Flöther, F.F.; Getz, G.; Goldsmith, M.; et al. Towards quantum-enabled cell-centric therapeutics. *arXiv* **2023**, arXiv:2307.05734.

39. Manjili, M.H.; Manjili, S.H. The quantum model of T-cell activation: Revisiting immune response theories. *Scand. J. Immunol.* **2024**, *100*, e13375 [[CrossRef](#)]
40. Fujii, Y.R. The Quantum microRNA Immunity in Human Virus-Associated Diseases: Virtual Reality of HBV, HCV and HIV-1 Infection, and Hepatocellular Carcinogenesis with AI Machine Learning. *Arch. Clin. Biomed. Res.* **2020**, *4*, 89–129.
41. Roosan, D.; Nirzhor, S.; Khan, R.; Hai, F.; Haidar, M.R. Quantum Approximate Optimization Algorithm for Spatiotemporal Forecasting of HIV Clusters. *arXiv* **2025**, arXiv:2507.00848. [[CrossRef](#)]
42. Choppara, P.; Lokesh, B. Leveraging Quantum LSTM for High-Accuracy Prediction of Viral Mutations. *IEEE Access* **2025**, *12*, 25282–25300. [[CrossRef](#)]
43. Roosan, D.; Khan, R.; Khou, T.; Nirzhor, S.; Hai, F.; Provencher, B. Bridging Classical Molecular Dynamics and Quantum Foundations for Comprehensive Protein Structural Analysis. *arXiv* **2025**, arXiv:2506.20830. [[CrossRef](#)]
44. Roosan, D.; Samrose, S.; Khan, R.; Nirzhor, S.; Provencher, B. Quantum analysis of protein-ligand binding by integrating structural resolution, sequence homology, and ligand properties. In Proceedings of the Obesity, Fitness and Wellness Week, Atlanta, GA, USA, 19 July 2025.
45. Roosan, D.; Nirzhor, S.; Khan, R. Integrating Pharmacokinetics and Pharmacodynamics Modeling with Quantum Regression for Predicting Herbal Compound Toxicity. In Proceedings of the 2025 World Congress in Computer Science, Computer Engineering, and Applied Computing, Las Vegas, NV, USA, 22 July 2025.
46. Roosan, D.; Khou, T.; Nirzhor, S.; Khan, R.; Hai, F.; Essien-Aleksi, I.; Baskys, A. Harnessing Quantum and Liquid Neural Networks for Drug Repurposing in Neurology. In *Management Science and Industrial Engineering*; Gaol, F.L., Xu, Y., Dessouky, Y., Eds.; Advances in Transdisciplinary Engineering; IOP Press: Bai Island, Indonesia, 2025; pp. 29–36.
47. Roosan, D.; Khan, R.; Ashakin, M.R.; Khou, T.; Nirzhor, S.; Haider, M.R. Quantum Variational Transformer Model for Enhanced Cancer Classification. In Proceedings of the 7th International conference on Industrial Engineering and Artificial Intelligence, Krabi, Thailand, 23–25 April 2026.
48. Roosan, D.; Nirzhor, S.; Khan, R.; Hai, F. Quantum Gradient Optimized Drug Repurposing Prototype for Omics Data. In Proceedings of the 14th International Conference on Data Science, Technology and Applications, Bilbao, Spain, 10–12 June 2025.
49. Acuña González, J.P.; Sánchez Adame, M.; Montiel Ross, O.H. Variational Quantum Eigensolver for Clinical Biomarker Classification: A Hybrid Quantum-Classical Approach to HIV/TB Co-Infection and Drug Resistance Prediction. In Proceedings of the Seminar of Artificial Intelligence and Quantum Computing, Tijuana, Mexico, 16 October 2025.
50. Acuña González, J.P. *qvirus: Quantum Computing for Analyzing CD4 Lymphocytes and Antiretroviral Therapy*, version 0.0.4; CRAN: Vienna, Austria, 2025.
51. R Core Team. *R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2025. Available online: <https://www.R-project.org/> (accessed on 21 September 2025).
52. Ozluer Baser, B. Analyzing the Competition of HIV-1 Phenotypes with Quantum Game Theory. *GU J. Sci.* **2022**, *35*, 1190–1198.
53. Crawford, O.; Straaten, B.V.; Wang, D.; Parks, T.; Campbell, E.; Brierley, S. Efficient quantum measurement of pauli operators in the presence of finite sampling error. *Quantum* **2021**, *5*, 385. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.