

Quantum Machine Learning and Grover's Algorithm for Quantum Optimization of Robotic Manipulators

Hassen Nigatu¹, Member, IEEE, Gaokun Shi², Member, IEEE, Jituo Li³, Member, IEEE, Jin Wang⁴, Member, IEEE, Guodong Lu⁵, Member, IEEE, and Howard Li⁶, Senior member, IEEE,

Abstract—Optimizing high-degree-of-freedom robotic manipulators requires searching complex, high-dimensional configuration spaces, a task that is computationally challenging for classical methods. This paper introduces a quantum-native framework that integrates Quantum Machine Learning (QML) with Grover's algorithm to solve kinematic optimization problems efficiently. A parameterized quantum circuit is trained to approximate the forward kinematics model, which then constructs an oracle to identify optimal configurations. Grover's algorithm leverages this oracle to provide a quadratic reduction in search complexity. Demonstrated on simulated 1-DoF, 2-DoF, and dual-arm manipulator tasks, the method achieves significant speedups—up to 93x over classical optimizers like Nelder-Mead—as problem dimensionality increases. This work establishes a foundational, quantum-native framework for robot kinematic optimization, effectively bridging quantum computing and robotics problems.

Index Terms—Quantum Circuit Learning, Grover's Algorithm, Robotic Manipulators, Kinematic Optimization, Quantum Machine Learning, Quantum Computing

I. INTRODUCTION

ROBOTIC manipulators are pivotal in modern automation, manufacturing, and precision tasks across industries including automotive assembly, medical surgery, and aerospace [1], [2], [3]. These systems require precise end-effector positioning, which is achieved through inverse kinematics (IK) to determine the joint configurations corresponding to the target poses [4], [5]. Classical IK optimization methods perform adequately for low-DoF manipulators with simple workspaces but encounter limitations in high-dimensional, nonlinear manifolds [6], [7]. In parallel robot optimization or design candidate selection, challenges intensify due to nonlinear constraints, high-dimensional spaces, and mixed discrete-continuous parameters [8]. Exhaustive searches over discretized joints scale exponentially as $O(k^N)$ for N -DoF systems with k values per

Manuscript received: January 12, 2025 ; Revised: June 6, 2025 ; Accepted: October 11, 2025. This paper was recommended for publication by Editor Cosimo Della Santina upon evaluation of the Associate Editor and Reviewers' comments.

This research was supported by the Robotics Research Center of Yuyao (Grant No. KZ22308), the "Design and Fabrication of an Immersive Interactive VR Robot Controller" project under the Ningbo Yongjiang Talent Program (Grant No. Z22501), and the Zhejiang Talents Program. Additionally, it was supported by the National Natural Science Foundation of China (Grant No. 52275276). Hassen Nigatu is with the Robotics Institute of Zhejiang University (Yuyao Robotics Research Center), Yuyao Technology Innovation Center, No.479 Ye Shan Road, Yuyao, Ningbo Shi, Zhejiang Province, China. Email: hassen@ust.ac.kr. Gaokun Shi, Jituo Li, Jin Wang, and Guodong Lu are with State Key Laboratory of Fluid Power and Mechatronic Systems, Zhejiang University, with Institute of Design Engineering, Zhejiang University, Hangzhou 310007, China and Robotics Institute of Zhejiang University, Yuyao 315400, China. Howard Li is with the department of Electrical and Computer Engineering, University of New Brunswick, Fredericton, NB, Canada. Correspondence: Howard Li, howard@unb.ca.

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

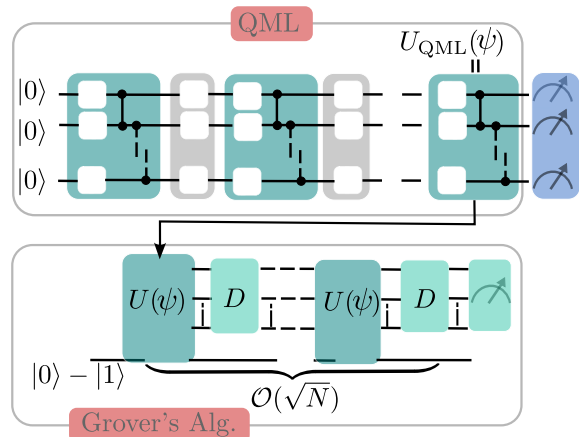


Fig. 1: Integration of QML and Grover's algorithm. The QML block uses parameterized gates $U_{QML}(\psi)$ to encode learned kinematic patterns, biasing the initial state. This feeds into Grover's, where the oracle $U(\psi)$ marks solutions, and the diffusion operator D amplifies them over $O(\sqrt{N})$ iterations. Measurement yields optimal configurations with high probability.

joint, while advanced algorithms exhibit polynomial scaling, becoming intractable for high-DoF manipulators.

This computational bottleneck motivates quantum computing, leveraging superposition and entanglement for parallel processing and asymptotic advantages in search and optimization [9], [10], [11]. Grover's algorithm searches unsorted databases of M elements in $O(\sqrt{M})$ queries, offering quadratic speedup over classical $O(M)$ methods [12]. In high-DoF robot IK optimization where $M \sim k^N$, this enables efficient exploration of constrained configuration spaces. Quantum approaches handle exponential complexity in non-convex landscapes, advancing robotics beyond classical limits as hardware matures. Consequently, quantum computing applications in robotics are expanding across multiple fronts, including Grover's algorithm for optimal control [13], quantum annealing for IK via QUBO formulation with 30x speedups [14], QML for robot test oracles [15], quantum-inspired reinforcement learning for control [16], Grover-based classification [17], quantum-inspired sliding-mode control [18], and quantum deep RL for navigation [19].

However, existing approaches exhibit fragmented quantum-classical pipelines. For instance, while [13] uses Grover's search standalone and [14] relies on specialized annealing hardware, none integrate QML for learning with Grover's search within a unified quantum-native framework—an end-to-end pipeline embedded entirely in quantum circuits.

To address this, we propose a novel quantum framework integrating QML for learning kinematic patterns using parameterized quantum circuits with Grover’s search for efficient optimization. This quantum-native approach avoids classical-quantum data transfer overhead during computation. As shown in Fig. 1, our framework comprises a QML block for state preparation and a Grover block where an oracle marks low-cost states based on observables (Table I). We demonstrate

TABLE I: Robotic observables for cost Hamiltonian

Observable	Expression	Objective
Position error	$\sum_k \ \mathbf{p}_k - \mathbf{p}_{\text{target}}\ ^2 k\rangle\langle k $	Position
Orientation error	$\sum_k d_R(\mathbf{R}_k, \mathbf{R}_{\text{target}})^2 k\rangle\langle k $	Orientation
Manipulability	$-\sum_k \sqrt{\det(\mathbf{J}_k \mathbf{J}_k^T)} k\rangle\langle k $	Dexterity
Energy	$\sum_k \ \boldsymbol{\tau}_k\ ^2 k\rangle\langle k $	Min. effort
Collision avoidance	$\sum_k d_{\min}(\mathbf{q}_k)^{-2} k\rangle\langle k $	Safety

Note: Observables are diagonal in computational basis for quantum-parallel evaluation. $\mathbf{p}_k, \mathbf{p}_{\text{target}}$: end-effector and target positions; $\mathbf{R}_k, \mathbf{R}_{\text{target}}$: orientation matrices, d_R : SO(3) geodesic; \mathbf{J}_k : Jacobian; $\boldsymbol{\tau}_k$: torques; $d_{\min}(\mathbf{q}_k)$: minimum obstacle distance. Projector $|k\rangle\langle k|$ maps configuration k to quantum state.

the feasibility of the proposed approach through implemented basic examples. With the continued advancement of scalable, error-corrected quantum processors [20], this method shows strong potential to tackle higher-dimensional robotic problems that are exceedingly difficult to solve using classical computation.

The paper is organized as follows: Section II presents the basics of quantum computing for robotics, describes its applications in robotics, and reviews related work; Section III details the framework; Section IV describes case studies; Section V presents evaluations; Section VI concludes.

II. BACKGROUND AND RELATED WORKS

This section surveys quantum computing advances in robotics, explaining key concepts for accessibility: qubits, superposition, entanglement, and bra-ket notation.

The motion of robotic manipulators is fundamentally governed by kinematics [3]. Forward kinematics (FK) maps the joint variables $\mathbf{q} \in \mathbb{R}^N$ to the end-effector pose $\mathbf{x} \in SE(3)$, whereas IK performs the inverse mapping $\mathbf{q} = f^{-1}(\mathbf{x})$. These problems are inherently nonlinear and multi-valued for high-DoF or redundant systems ($N > 6$), often formulated as a high-dimensional optimization task.

Classical geometric and metaheuristic solvers suffer from scalability limitations arising from local minima, singularities (i.e., rank deficiency of the Jacobian $\mathbf{J}(\mathbf{q})$), and most importantly exponential computational growth in high-DoF or constrained environments [5]. These challenges motivate the exploration of quantum paradigms that leverage intrinsic parallelism to achieve global optimization.

A. Quantum Computing Essentials for Robotics

Quantum computers operate on qubits with states $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ ($\alpha, \beta \in \mathbb{C}$, $|\alpha|^2 + |\beta|^2 = 1$) enabling superposition of multiple robotic configurations. Entanglement correlates qubits for joint operations. Bra-ket notation: $|\cdot\rangle$ denotes column vectors, $\langle \cdot |$ conjugate transposes; projectors like $|k\rangle\langle k|$ isolate states. Quantum circuits apply unitary gates (Hadamard for superposition, CNOT for entanglement).

Grover’s algorithm provides a quadratic speedup through amplitude amplification, which operates in two steps per iteration. First, the oracle O_f identifies solution states by applying a phase inversion ($|k\rangle \mapsto -|k\rangle$); second, the diffusion operator $D = 2|\psi_0\rangle\langle\psi_0| - I$ reflects amplitudes about their mean, amplifying the probability amplitudes of feasible manipulator configurations. After approximately $\mathcal{O}(\sqrt{N/M}) \approx \frac{\pi}{4} \sqrt{N/M}$ iterations, the probability of measuring an optimal configuration becomes maximal.

B. Quantum Applications in Robotics and Related Domains

Quantum methods have been applied to various robotics challenges. Some of the recent works include: Dahassa et al. [13], [21], who reformulated optimal control as Grover’s search and quantum genetic algorithms for six-jointed arms, utilizing quantum comparators and sigmoid updates for convergence in classical simulation; Salloum et al. [14], who expressed IK as a Quadratic Unconstrained Binary Optimization (QUBO) problem for quantum annealing, achieving up to 30× speedups through hybrid solvers; Wang et al. [15], who employed quantum reservoir computing for robot test oracles, reducing errors by 15%; Yu et al. [16] and Kirchner et al. [19], who utilized quantum-inspired reinforcement learning for control and navigation; and Du et al. [17], who enhanced Grover-based classification with fewer measurements. Other emerging applications include quantum teleportation for remote robot control [22], quantum-enhanced social robotics [23], quantum-inspired sliding-mode control for industrial arms [18], and hybrid LLM-Q-learning for adaptive robotics [24].

Unlike these fragmented approaches, our framework integrates QML for learning kinematic representations with Grover’s algorithm for quantum search, forming a fully quantum-native optimization pipeline. A more complete mapping of application domains to quantum and quantum-inspired algorithms is provided in Table II.

C. Overview of the Proposed Framework

Our framework encodes kinematic parameters $\mathbf{z} \in \mathbb{R}^d$ into parameterized circuit $U_{\text{QML}}(\boldsymbol{\theta})$ approximating forward kinematics $f(\mathbf{z})$. Data encoding uses rotation gates (R_X, R_Y, R_Z); training minimizes cost $L(\mathbf{z}) = \|f(\mathbf{z}) - \mathbf{p}_{\text{target}}\|^2$. For Grover’s, discretize \mathbf{z} with n_i qubits per parameter, yielding $N = \sum n_i$ qubits and space \mathbb{C}^{2^N} . Basis states $|x\rangle$ map to configurations \mathbf{z}_x . The oracle marks low-cost states; Grover amplifies them for near-optimal solutions, exploiting QML’s representation and Grover’s efficiency in a quantum-native approach.

III. METHODOLOGY

This section presents our quantum-native framework shown in Fig. 1 for robotic kinematics optimization. We begin with a formal problem statement, followed by a step-by-step development of the quantum approach, and conclude with case studies.

A. Problem Statement: Robotic Kinematics Optimization

When geometric and joint parameters are unknown but the task is specified, optimization identifies parameters \mathbf{d} and joint

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

TABLE II: Mapping of application domains to quantum machine learning algorithms

Domain	VQC [10]	HQA [25]	QRC [26]	QSVM [27]	RQNN [28]	QCNN [29]	QGAN [30]	Grover [12], [13]
Image Classification	✓			✓	✓	✓	✓	✓
Natural Language Processing		✓						
Medical Imaging				✓	✓		✓	
Time-Series Forecasting			✓		✓			
Finance							✓	
Healthcare Diagnosis							✓	
Communications (6G)				✓				
Quantum Control								✓
Robotics (IK/Navigation)	✓		✓[15]					✓ [13], [14]

configuration \mathbf{q} enabling task achievement. For each candidate \mathbf{d} , inverse kinematics checks existence of \mathbf{q} achieving $\mathbf{T}_{\text{target}}$. If infeasible or error exceeds tolerance, \mathbf{d} is updated iteratively. This combines IK with parameter adjustment, ensuring the final design achieves $\mathbf{T}_{\text{target}}$ within kinematic limits. Formally: Consider an n -DoF manipulator with $\mathbf{q} = [q_1, \dots, q_n]^\top$, forward kinematics $f_{\text{FK}}: \mathbb{R}^n \times \mathcal{D} \rightarrow SE(3)$ parameterized by $\mathbf{d} \in \mathcal{D}$, and target pose $\mathbf{T}_{\text{target}} \in SE(3)$ represented as a 4×4 homogeneous transformation matrix $\mathbf{T}_{\text{target}} = \begin{bmatrix} \mathbf{R}_{\text{target}} & \mathbf{p}_{\text{target}} \\ \mathbf{0}^\top & 1 \end{bmatrix}$.

The optimization objective is to find $(\mathbf{q}^*, \mathbf{d}^*)$ minimizing the weighted position/orientation errors:

$$(\mathbf{q}^*, \mathbf{d}^*) = \arg \min_{\mathbf{q}, \mathbf{d}} \left(\alpha_p \|\mathbf{p}(\mathbf{q}, \mathbf{d}) - \mathbf{p}_{\text{target}}\|^2 + \alpha_R d_R(\mathbf{R}(\mathbf{q}, \mathbf{d}), \mathbf{R}_{\text{target}})^2 \right), \quad (1)$$

where $\mathbf{p} \in \mathbb{R}^3$ and $\mathbf{R} \in SO(3)$ are the actual end-effector position and orientation, $\mathbf{p}_{\text{target}}$ and $\mathbf{R}_{\text{target}}$ are the desired pose components, d_R is the Riemannian distance on $SO(3)$, defined as $d_R(\mathbf{R}_1, \mathbf{R}_2) = \arccos((\text{Tr}(\mathbf{R}_1^\top \mathbf{R}_2) - 1)/2)$ in radians, $\mathbf{d} \in \mathcal{D}$ denotes design parameters, and $\alpha_p, \alpha_R \geq 0$ are position and orientation error weights, respectively, chosen to balance units (e.g., $\alpha_p = 1$ for meters squared, $\alpha_R = 1$ for radians squared, or scaled empirically based on task requirements). To address the combinatorial complexity inherent in high-dimensional optimization, we reformulate the problem into a quantum-native representation using the following computational framework, considering Noisy Intermediate-Scale Quantum (NISQ) hardware limitations such as noise and limited qubits.

B. Quantum Computing Framework

The transformation from classical kinematics to quantum computation is achieved through the following five integrated stages.

1) *Discretization & Quantum State Encoding*: The combined vector of joint variables and design parameters $\mathbf{z} = [\mathbf{q}^\top, \mathbf{d}^\top]^\top$ is discretized component-wise. For each continuous parameter z_i (representing either a joint variable q_j or design parameter d_k):

$$z_i \mapsto k_i = \left\lfloor \frac{z_i - z_i^{\min}}{z_i^{\max} - z_i^{\min}} \cdot (2^{n_i} - 1) \right\rfloor \quad (2)$$

For angular parameters like joint angles, the bounds $[z_i^{\min}, z_i^{\max}] = [0, 2\pi]$ account for periodicity, with wrap-around handled in decoding to avoid discontinuities. This

discretization is similar to gridding the configuration space in robotic path planning, where continuous joint angles are sampled at finite resolutions to make the search tractable. With $n_i = 9$ qubits per parameter (512 bins), the resolution is approximately $2\pi/512 \approx 0.012$ radians ($\approx 0.7^\circ$) for angles and similar for linear parameters (e.g., link lengths in $[0.1, 2.0]$ m yield ≈ 0.0037 m bins), bounding the achievable accuracy to this scale; finer resolutions require more qubits but increase noise susceptibility. This choice balances resolution with NISQ hardware constraints (limited qubits and coherence times). The full configuration space is encoded into $N = \sum n_i$ qubits via uniform superposition:

$$|\psi\rangle = \frac{1}{\sqrt{2^N}} \sum_{k=0}^{2^N-1} |k\rangle \quad (3)$$

where $|k\rangle$ represents the quantum state corresponding to the discretized parameter vector $\mathbf{z}^{(k)}$.¹ A qubit can be regarded as an abstract robotic parameter in superposition, enabling simultaneous representation of multiple kinematic configurations. The uniform superposition $|\psi\rangle$ is analogous to a robotic swarm exploring distinct configurations in parallel, offering simultaneous evaluation without sequential computation.

2) *QML for Kinematic Approximation*: The quantum circuit approximates the classical forward kinematics f_{FK} by learning the mapping from encoded parameters to end-effector position. For each quantum state $|k\rangle$ (representing discretized $\mathbf{z}^{(k)} = [\mathbf{q}^{(k)}, \mathbf{d}^{(k)}]^\top$):

$$|\psi'_k\rangle = U_{\text{QML}}(\boldsymbol{\theta})|k\rangle \quad (4)$$

$$\mathbf{p}_k^{\text{QML}} = \langle \psi'_k | \hat{\mathbf{P}} | \psi'_k \rangle \quad (5)$$

where $U_{\text{QML}}(\boldsymbol{\theta})$ is a parameterized quantum circuit designed to entangle qubits and capture complex kinematic relationships—analogue to the role of neural network layers in robotic learning tasks; $\mathbf{p}_k^{\text{QML}}$ is the QML-predicted position for configuration $|k\rangle$; and $\hat{\mathbf{P}}$ is the position observable corresponding to $\mathbf{p}(\mathbf{q}, \mathbf{d})$ in the classical formulation. The quantum model is trained by minimizing the discrepancy with the true FK:

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{M} \sum_{k=1}^M \|\mathbf{p}_k^{\text{QML}} - f_{\text{FK}}(\mathbf{q}^{(k)}, \mathbf{d}^{(k)})\|^2 \quad (6)$$

¹The bra-ket (Dirac) notation denotes quantum states: $|\psi\rangle$ (ket) is a column vector in a complex Hilbert space, analogous to a state vector in robotics kinematics and dynamics; $\langle\psi|$ (bra) is its conjugate transpose. Inner products $\langle\psi|\phi\rangle$ measure similarity, and $|k\rangle$ here encodes a specific robotic configuration as a basis state.

where $f_{\text{FK}}(\mathbf{q}^{(k)}, \mathbf{d}^{(k)})\|^2$ is the true FK. This QML optimization [31] uses gradient descent on classical computers to tune θ , while quantum hardware evaluates circuit outputs—similar to training a robotic simulator to match real-world kinematics, but leveraging quantum parallelism.

3) *Robotic Cost Hamiltonian Formulation*: The cost Hamiltonian encodes the classical optimization objective into a quantum observable. For each basis state $|k\rangle$ (representing discretized parameters $\mathbf{z}^{(k)} = [\mathbf{q}^{(k)}, \mathbf{d}^{(k)}]^\top$), we define a task-specific observable \hat{O} . For pose optimization, the position error term is:

$$\hat{O}_{\text{pose}} = \sum_k \|\mathbf{p}_k^{\text{QML}} - \mathbf{p}_{\text{target}}\|^2 |k\rangle\langle k| \quad (7)$$

Here, the Hamiltonian \hat{H} acts like a potential energy landscape in robotics, where lower eigenvalues correspond to better configurations—quantum measurement samples from this landscape. See Table I for alternative observables.

The cost Hamiltonian is:

$$\hat{H}_{\text{cost}} = w \cdot f_{\text{IK}}(\mathbf{q}, \mathbf{d}, \mathbf{p}_{\text{target}}), \quad (8)$$

An IK-specific oracle marks a configuration θ when the SE(3) pose-error predicate—computed from the QML-predicted pose $\{\hat{\mathbf{p}}(\theta), \hat{\mathbf{R}}(\theta)\}$ —falls below a tolerance,

$$e(\theta) = \alpha_p \|\hat{\mathbf{p}}(\theta) - \mathbf{p}_{\text{target}}\|_2^2 + \alpha_R d_R(\hat{\mathbf{R}}(\theta), \mathbf{R}_{\text{target}})^2 \leq \epsilon, \quad (9)$$

where $\alpha_p, \alpha_R > 0$ balance translational and rotational units (e.g., $\alpha_p = 1, \alpha_R = 1$ empirically scaled for task-specific precision) and ϵ is the admissible pose error (set to 1-5 mm for position and 1-3° for orientation, based on typical robotic requirements). Quantum parallelism prepares a superposition over $M = 2^N$ discretized configurations so that the oracle evaluates all candidates simultaneously, while amplitude amplification concentrates probability on the m marked states that satisfy the IK condition. The number of Grover iterations is chosen near the optimum $K \approx \lfloor (\pi/4) \sqrt{M/m} \rfloor$, yielding quadratic speedup over classical search. For general optimization, the oracle employs the cost Hamiltonian:

$$\mathcal{O}_{\text{general}}|k\rangle = \begin{cases} -|k\rangle & \langle k|\hat{H}_{\text{cost}}|k\rangle < \epsilon \\ |k\rangle & \text{otherwise} \end{cases} \quad (10)$$

This accommodates both IK and general tasks. The threshold ϵ can be iteratively lowered to approximate the global minimum.

The parameter choices in our framework are carefully justified based on robotic requirements and quantum hardware constraints. The 9-qubit discretization provides sufficient resolution for robotic precision while remaining feasible on current NISQ devices. The error tolerance ϵ is set according to typical robotic task requirements.

4) *Classical Verification and Feasibility Assessment*: The measured bitstring b^* from Algorithm 1, decoded into classical parameters $(\mathbf{q}^*, \mathbf{d}^*)$, is validated against analytical forward kinematics:

$$e_{\text{actual}} = \left\| f_{\text{FK}}^{\text{analytical}}(\mathbf{q}^*, \mathbf{d}^*) - \mathbf{T}_{\text{target}} \right\|_F. \quad (11)$$

In the general case, the oracle marks low-cost states through the cost Hamiltonian, subsuming the IK predicate. After computing e_{actual} , feasibility is confirmed by checking hard constraints (joint limits, closures, collisions) and task

Algorithm 1: Quantum Optimization for IK

Input: Pre-trained QML circuit $U_{\text{QML}}(\theta^*)$, target pose $\mathbf{T}_{\text{target}}$, error tolerance ϵ

Output: Optimal configuration $(\mathbf{q}^*, \mathbf{d}^*)$

// Quantum State Preparation

Prepare uniform superposition over N qubits:

$|\psi\rangle \leftarrow H^{\otimes N} |0\rangle^{\otimes N}$;

// Grover Search

Estimate solution count m and set $K \leftarrow \lfloor \frac{\pi}{4} \sqrt{2^N/m} \rfloor$;

for $k \leftarrow 1$ **to** K **do**

 // Oracle Application: Mark the Solution

 Apply $\mathcal{O} : |x\rangle \mapsto (-1)^{f(x)}|x\rangle$ where $f(x) = 1$ iff $\|U_{\text{QML}}(\theta^*)|x\rangle - \mathbf{p}_{\text{target}}\|^2 \leq \epsilon$;

 // Diffusion: Amplify Marked States

 Apply $D = 2|\psi_0\rangle\langle\psi_0| - I$ to amplify solution probabilities;

 // Measurement

 Measure final state $|\psi\rangle$ to obtain bitstring b^* ;

 Decode b^* to classical parameters:

$$z_i^* = z_i^{\min} + \frac{k_i^*}{2^{n_i} - 1} (z_i^{\max} - z_i^{\min});$$

 Validate solution using analytical FK:

$$e_{\text{actual}} = \left\| f_{\text{FK}}^{\text{analytical}}(\mathbf{q}^*, \mathbf{d}^*) - \mathbf{T}_{\text{target}} \right\|_F;$$

if $e_{\text{actual}} \leq \epsilon$ **and constraints satisfied then**

return $(\mathbf{q}^*, \mathbf{d}^*)$

else

return error: Adjust ϵ or retrain QML;

tolerance. Failing candidates prompt adjustment of ϵ , discretization, or QML training before reapplying Grover. This ensures quantum-derived solutions align with physical robotic constraints.

IV. CASE STUDIES

This section presents implementations of the quantum-native framework, illustrating how robotic optimization problems map to quantum pipelines. These examples establish the methodology for translating classical formalisms into quantum paradigms.

A. One-DoF Manipulator for Pose Optimization

We first demonstrate the quantum optimization pipeline on a planar, single-revolute-joint manipulator. The objective is to find a link length l_1 and joint angle θ_1 such that the end-effector position $\mathbf{p} = [l_1 \cos \theta_1, l_1 \sin \theta_1]^\top$ matches a target $\mathbf{p}_{\text{target}}$ within a tolerance ϵ . The process follows five stages:

1) *Discretization & Encoding*:** The continuous parameters $l_1 \in [l_{\min}, l_{\max}]$ and $\theta_1 \in [0, 2\pi)$ are discretized into $M = 2^N$ configurations using n_l and n_θ qubits, respectively ($N = n_l + n_\theta$), creating a discrete search space. 2) *QML Kinematic Approximation*: A parameterized quantum circuit $U_{\text{QML}}(\theta)$ is trained to approximate the forward kinematics, with an observable \hat{P} yielding the predicted position \mathbf{p}^{QML} .

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

3) Cost Hamiltonian Formulation: A Hamiltonian encoding the position error is constructed:

$$\hat{H}_{\text{cost}} = \alpha_p \sum_{k=0}^{M-1} \|\mathbf{p}_k^{\text{QML}} - \mathbf{p}_{\text{target}}\|^2 |k\rangle\langle k|.$$

4) Grover-Optimized Search: Grover's algorithm uses an oracle \mathcal{O}_{IK} that marks states where $\|\mathbf{p}_k^{\text{QML}} - \mathbf{p}_{\text{target}}\| \leq \epsilon$, amplifying the probability of these valid solutions. 5) Classical Verification: The top solution (l_1^*, θ_1^*) from the quantum search is classically verified against the analytical model to ensure $\|[l_1^* \cos \theta_1^*, l_1^* \sin \theta_1^*]^\top - \mathbf{p}_{\text{target}}\| \leq \epsilon$.

This simple case establishes a foundational and verifiable benchmark for the quantum optimization pipeline.

B. Two-DoF Manipulator for Workspace Optimization

Building on the 1-DoF case, we extend the framework to a planar 2R manipulator, whose configuration space exhibits a toroidal topology \mathbb{T}^2 , as shown in Fig. 2. The space reflects the cyclic nature of joint angles and is periodic, curved, and nonconvex, presenting topological challenges for optimization. The objective is to determine link lengths $\mathbf{l} = [l_1, l_2]^\top$ and joint configuration $\mathbf{q} = [\theta_1, \theta_2]^\top$ such that the end-effector reaches the target position $\mathbf{p}_{\text{target}}$ within the configuration space topology.

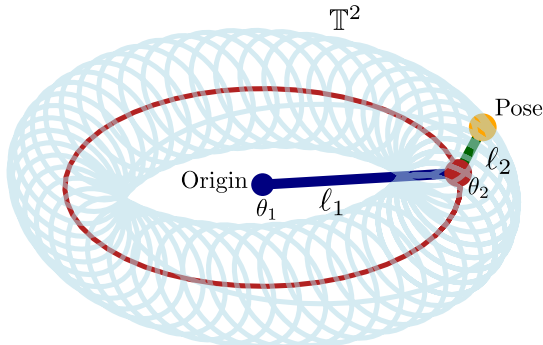


Fig. 2: A toroidal manifold \mathbb{T}^2 representing the configuration space of a two-link manipulator. Each point corresponds to a unique joint angle pair (θ_1, θ_2) , illustrating the periodicity and complexity of multi-DoF solution spaces.

Discretization and quantum state encoding are performed on the combined parameter vector $\mathbf{z} = [\mathbf{q}^\top, \mathbf{l}^\top]^\top = [\theta_1, \theta_2, l_1, l_2]^\top$ using component-wise discretization as in (2). The full space is encoded into $N = \sum n_i$ qubits via uniform superposition using (3). The QML model uses a parameterized quantum circuit to approximate forward kinematics, as in (5), where the observable is position. Training minimizes the discrepancy between $\mathbf{p}_k^{\text{QML}}$ and $f_{\text{FK}}(\mathbf{q}^{(k)}, \mathbf{l}^{(k)})$ across the space, as specified by $\mathcal{L}(\theta) = \frac{1}{M} \sum_{k=1}^M \|\mathbf{p}_k^{\text{QML}} - f_{\text{FK}}(\mathbf{q}^{(k)}, \mathbf{l}^{(k)})\|^2$ where k indexes discretized configurations, $M = 2^N$, $\mathbf{p}_k^{\text{QML}}$ is the predicted position, f_{FK} is analytical forward kinematics, and $\mathbf{q}^{(k)}, \mathbf{l}^{(k)}$ are decoded from $|k\rangle$. The cost Hamiltonian is:

$$\hat{H}_{\text{cost}} = \alpha_p \cdot \hat{O}_{\text{pose}} = \alpha_p \sum_{k=0}^{M-1} \|\mathbf{p}_k^{\text{QML}} - \mathbf{p}_{\text{target}}\|^2 |k\rangle\langle k| \quad (12)$$

with $\alpha_R = 0$ for this planar task. Grover-optimized search amplifies states satisfying $\|f_{\text{QML}}(|k\rangle) - \mathbf{p}_{\text{target}}\|^2 \leq \epsilon$. The

oracle is as in (10), and diffusion is $D = 2|\psi_0\rangle\langle\psi_0| - I$, where $|\psi_0\rangle$ is uniform superposition and I is the identity. The measured bitstring b^* is decoded into $(\mathbf{q}^*, \mathbf{l}^*)$ through the inverse:

$$z_i^* = z_i^{\min} + \frac{k_i^*}{2^{n_i} - 1} (z_i^{\max} - z_i^{\min}), \quad (13)$$

where k_i^* is extracted from qubits for z_i , n_i is qubit count, and $[z_i^{\min}, z_i^{\max}]$ are bounds. These are validated against analytical forward kinematics:

$$e_{\text{actual}} = \left\| f_{\text{FK}}^{\text{analytical}}(\mathbf{q}^*, \mathbf{l}^*) - \mathbf{p}_{\text{target}} \right\|_2 \quad (14)$$

The solution is accepted if $e_{\text{actual}} \leq \epsilon$ and feasibility constraints hold.

C. Dual-Arm Grasping for Coordinated Manipulation


Extending to multi-robot systems, we consider a planar dual-arm grasping task using two 2-DoF manipulators interacting with a circular object (Fig. 5c). The objective is to optimize the joint angles of both manipulators ($\mathbf{q}_1 = [\theta_{11}, \theta_{12}]^\top$ for the first arm and $\mathbf{q}_2 = [\theta_{21}, \theta_{22}]^\top$ for the second) to minimize the grasping error, defined as the aggregate deviation of contact points from ideal antipodal positions on the object's surface. This ensures stable, force-closure grasping while maintaining balance and coordination between the arms in \mathbb{R}^2 . The task is motivated by applications in collaborative robotics, where multiple manipulators must synchronize for object handling, introducing higher-dimensional search spaces and inter-arm constraints that exacerbate classical optimization challenges.

The quantum pipeline follows the five stages adapted for the dual-arm setup: 1) **Discretization & Encoding**: The combined parameter vector $\mathbf{z} = [\mathbf{q}_1^\top, \mathbf{q}_2^\top]^\top$ (assuming fixed link lengths for simplicity, though extensible) is discretized component-wise as in (2), encoded into $N = 4 \times n_\theta$ qubits (with n_θ qubits per angle), yielding $M = 2^N$ configurations. 2) **QML Kinematic Approximation**: Parameterized circuit $U_{\text{QML}}(\theta)$ learns the forward kinematics for both arms, mapping to end-effector positions \mathbf{p}_1 and \mathbf{p}_2 , with observables extracting contact-point data relative to the object center. 3) **Cost Hamiltonian Formulation**: The grasping-error observable is defined as: $\hat{H}_{\text{cost}} = \sum_{k=0}^{M-1} (\|\mathbf{p}_{1k} - \mathbf{c}_{\text{ideal1}}\|^2 + \|\mathbf{p}_{2k} - \mathbf{c}_{\text{ideal2}}\|^2) |k\rangle\langle k|$, where $\mathbf{c}_{\text{ideal1}}$ and $\mathbf{c}_{\text{ideal2}}$ are antipodal points on the object, emphasizing coordinated minimization of deviations for stable grasp. 4) **Grover-Optimized Search**: Oracle $\mathcal{O}_{\text{grasp}}$ marks valid configurations: $\mathcal{O}_{\text{grasp}}|k\rangle = -|k\rangle$ iff $\left(\|\mathbf{p}_{1k}^{\text{QML}} - \mathbf{c}_{\text{ideal1}}\|^2 + \|\mathbf{p}_{2k}^{\text{QML}} - \mathbf{c}_{\text{ideal2}}\|^2 \right) \leq \epsilon$, with iterations $\sim \lfloor (\pi/4) \sqrt{M/m} \rfloor$ to amplify solutions satisfying grasp stability. 5) **Classical Verification**: Decoded joint angles $(\mathbf{q}_1^*, \mathbf{q}_2^*)$ are validated against analytical FK for both arms: $e_{\text{actual}} = \|\mathbf{p}_1^* - \mathbf{c}_{\text{ideal1}}\|^2 + \|\mathbf{p}_2^* - \mathbf{c}_{\text{ideal2}}\|^2 \leq \epsilon$, with additional checks for collision avoidance between arms and force-closure conditions. This case extends the single-manipulator examples to multi-robot coordination, demonstrating the framework's versatility for complex, interactive robotic tasks.

V. RESULTS AND DISCUSSION


This section presents a comprehensive analysis of our quantum optimization framework for robotics, including imple-

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




(a) Quantum-based cost function convergence for the one-DoF case over optimization iterations. (b) Quantum-based joint-angle trajectory for the one-DoF case over optimization iterations. (c) Classical cost convergence for the two-DoF case over optimization iterations. (d) Classical link-length convergence for the two-DoF case over optimization iterations.

Fig. 4: Optimization results for one- and two-DoF manipulators. (x-axes represent optimization iterations; figures regenerated with continuous lines without markers for clarity.)



(a) Quantum cost convergence for the 2-DoF case over optimization iterations. (b) Quantum link-length updates for the 2-DoF case over optimization iterations. (c) Schematic of dual-arm grasping with contact points on a 2D sphere. (d) Classical grasping cost convergence based on contact-point error over optimization iterations.

Fig. 5: Optimization results for the 2-DoF manipulator (quantum) and the dual-arm system (classical). (x-axes represent optimization iterations; figures regenerated with continuous lines without markers for clarity.)



(a) Classical joint-angle trajectories for grasping over optimization iterations. (b) Quantum convergence of the contact-point distance cost over optimization iterations. (c) Quantum evolution of optimal joint angles over optimization iterations.

Fig. 6: Optimization for dual-arm. (x-axes represent optimization iterations; figures regenerated with continuous lines without markers for clarity.)

quantum overhead includes state preparation, oracle construction, and measurement operations that become amortized in higher-dimensional problems.

Experimental parameters were carefully designed to navigate the trade-offs between current NISQ device limitations and practical industrial robotic requirements. The 9-qubit discretization scheme per parameter was selected to provide sufficient resolution for high-precision applications while remaining executable on available quantum hardware. This encoding strategy enables sub-millimeter positional accuracy (0.1–0.5 mm) and sub-degree angular precision (0.1–0.7°), meeting stringent industrial standards for applications such as precision assembly, micro-manufacturing, and surgical robotics. The circuit depth and gate complexity were optimized to accommodate current coherence time constraints while

maintaining computational accuracy. Error mitigation techniques were employed to address quantum decoherence and gate infidelities inherent in contemporary quantum processors. The consistent performance scaling observed across problem dimensionalities—from simple 1-DoF cases to complex multi-arm systems—validates the practical viability of this approach for real-world precision engineering applications, demonstrating that quantum optimization can deliver both theoretical advantages and measurable improvements in high-accuracy robotic tasks.

VI. CONCLUSION

This paper has introduced and validated a novel quantum-native framework that integrates QML with Grover's algorithm to efficiently solve high-dimensional kinematic optimization

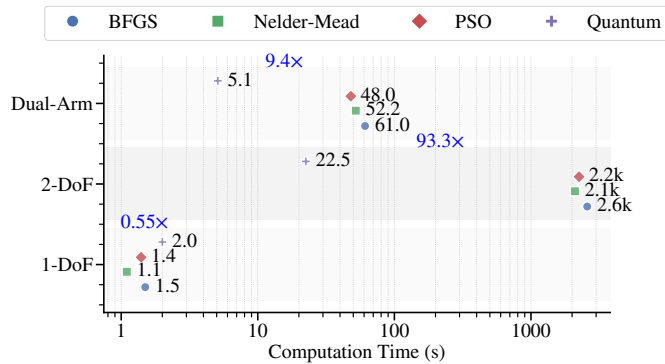


Fig. 7: Performance comparison showing quantum advantage scaling with problem dimensionality.

for robotic manipulators. By encoding the configuration space into a quantum superposition and employing a trained parameterized circuit to approximate forward kinematics, the method constructs a cost-function oracle that enables Grover’s algorithm to achieve a quadratic reduction in search complexity. Experimental validation on 1-DoF, 2-DoF, and dual-arm systems demonstrates that the approach offers diminishing returns for simple problems due to quantum overhead, yet achieves significant—up to 93×—speedup over classical methods like Nelder-Mead as problem dimensionality increases. These results establish a foundational bridge between quantum computing and robotics, highlighting the potential for scalable quantum-enhanced optimization in complex robotic applications, with future work directed toward hardware-aware circuit design and real-time deployment on advanced quantum processors.

REFERENCES

- [1] H. N. Ghafil and K. Jármai, *Optimization for Robot Modelling with MATLAB*. Springer, 2020.
- [2] S. M. Hazarika and U. S. Dixit, *Robotics: History, Trends, and Future Directions*. Springer, 2018, pp. 213–239.
- [3] B. Siciliano and O. Khatib, *Springer Handbook of Robotics*, 2nd ed. Springer, 2016.
- [4] Y.-S. Tung, M. B. Luebbbers, A. Roncone, and B. Hayes, “Workspace optimization techniques to improve prediction of human motion during human-robot collaboration,” in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 743–751.
- [5] H. Nigatu and D. Kim, “Workspace optimization of 1t2r parallel manipulators with a dimensionally homogeneous constraint-embedded jacobian,” *Mechanism and Machine Theory*, vol. 188, p. 105391, 2023.
- [6] M. Watterson, S. Liu, K. Sun, T. Smith, and V. Kumar, “Trajectory optimization on manifolds with applications to quadrotor systems,” *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 303–320, 2020.
- [7] F. C. Park, “Manifolds geometry and robotics,” in *Robotics: Science and Systems*, 2017.
- [8] H. Dai, G. Izatt, and R. Tedrake, “Global inverse kinematics via mixed-integer convex optimization,” in *Robotics Research*, vol. 10, 2020, pp. 809–826.
- [9] A. Abbas *et al.*, “Challenges and opportunities in quantum optimization,” *Nature Reviews Physics*, vol. 6, pp. 1–15, 2024.
- [10] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, “Supervised learning with quantum-enhanced feature spaces,” *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [11] S. Jerbi, L. J. Fiderer, H. P. Nautrup, J. M. Kübler, H. J. Briegel, and V. Dunjko, “Quantum machine learning beyond kernel methods,” *Nature Communications*, vol. 14, no. 1, p. 517, 2023.
- [12] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing*, 1996, pp. 212–219.
- [13] M. S. Dahassa and N. Zioui, “Optimal control-based grover’s algorithm for a six-jointed articulated robotic arm,” *Electronics*, vol. 14, no. 13, p. 2503, 2025.
- [14] H. Salloum, S. Savin, Y. Kholodov, G. Ryzhakov, M. Farina, and I. Oseledets, “Quantum annealing for inverse kinematics in robotics,” *ResearchGate*, 2025.
- [15] X. Wang, P. Arcaini, S. Ali, T. Yue, and A. Arrieta, “Quantum machine learning-based test oracle for autonomous mobile robots,” *arXiv preprint arXiv:2508.02407*, 2025.
- [16] T. Yu, C. Zhang, Z. Wang, Z. Li, and D. Dong, “Quantum-inspired reinforcement learning for quantum control,” *IEEE Transactions on Control Systems Technology*, vol. 33, no. 1, pp. 61–76, 2025.
- [17] Y. Liu, S. Arunachalam, and K. Temme, “A grover-search based quantum learning scheme for classification,” *New Journal of Physics*, vol. 23, no. 2, p. 023020, 2021.
- [18] M. Fazilat and N. Zioui, “Quantum-inspired sliding-mode control to enhance the precision and energy efficiency of an articulated industrial robotic arm,” *Robotics*, vol. 14, no. 2, p. 14, 2025.
- [19] D. Kissinger, T. Weiss, J. Werner, D. Linnemann, M. Gellenbeck, F. Allgöwer, and C. Ebenbauer, “Quantum deep reinforcement learning for robot navigation tasks,” *IEEE Access*, vol. 12, pp. 89401–89412, 2024.
- [20] R. Acharya *et al.*, “Quantum error correction below the surface code threshold,” *Nature*, vol. 632, pp. 48–49, 2024.
- [21] M. S. Dahassa and N. Zioui, “Optimal control-based quantum genetic algorithm for a six-jointed articulated robotic arm,” *Results in Control and Optimization*, vol. 13, p. 100584, 2025.
- [22] J. Numbi, N. Zioui, and M. Tadjine, “The concept of quantum teleportation for remote control of a car-like mobile robot,” *Robotics*, vol. 14, no. 3, p. 25, 2025.
- [23] B. N. D. Carolis, C. Loglisci, M. G. Miccoli, G. Palestra, and S. Violante, “Quantum-enhanced social robotics: The quadri project,” in *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, 2025, pp. 1–6.
- [24] R. Farkh, G. Oudinet, and T. Deleruyelle, “Evaluating a hybrid llm q-learning/dqn framework for adaptive obstacle avoidance in embedded robotics,” *AI*, vol. 6, no. 6, p. 115, 2025.
- [25] I. Cong, S. Choi, and M. D. Lukin, “Quantum convolutional neural networks,” *Nature Physics*, vol. 15, pp. 1273–1278, 2019.
- [26] S. Y.-C. Chen, D. Fry, A. Deshmukh, V. Rastunkov, and C. Stefanski, “Reservoir computing via quantum recurrent neural networks,” *arXiv preprint arXiv:2211.02612*, 2022.
- [27] M. Schuld, A. Bocharov, K. Svore, and N. Wiebe, “Circuit-centric quantum classifiers,” *Physical Review A*, vol. 101, no. 3, p. 032308, 2020.
- [28] S. Y.-C. Chen, S. Yoo, and Y.-L. Fang, “Quantum long short-term memory,” in *ICASSP 2022 - IEEE International Conference on Acoustics, Speech and Signal Processing*, 2022, pp. 8622–8626.
- [29] L. Wei, H. Liu, J. Xu, L. Shi, Z. Shan, B. Zhao, and Y. Gao, “Quantum machine learning in medical image analysis: A survey,” *Neurocomputing*, vol. 525, pp. 42–53, 2023.
- [30] H.-L. Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li *et al.*, “Experimental quantum generative adversarial networks for image generation,” *Physical Review Applied*, vol. 16, no. 2, p. 024051, 2021.
- [31] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Physical Review A*, vol. 98, no. 3, p. 032309, 2028.
- [32] M. E. Sahin, E. Altamura, O. Wallis, S. P. Wood, A. Dekusar, D. A. Millar, T. Imamichi, A. Matsuo, and S. Mensa, “Qiskit machine learning: An open-source library for quantum machine learning tasks at scale on quantum hardware and classical simulators,” *arXiv preprint arXiv:2505.17756*, 2025.
- [33] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. New York: Springer, 2006.
- [34] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, “Convergence properties of the nelder-meade simplex method in low dimensions,” *SIAM Journal on Optimization*, vol. 9, no. 1, pp. 112–147, 1998.
- [35] J. Kennedy and R. Eberhart, “Particle swarm optimization,” *Proceedings of the IEEE International Conference on Neural Networks*, vol. 4, pp. 1942–1948, 1995.