Solving Higher-Order Binary Optimization Problems: A Trapped-ion Ground-State Search for Protein Folding

Dr M.Salomi
Post Doctoral Fellow
School of CS & AI
SR University, Warangal
Telangana
India
salocs22@gmail.com

K.Rajchandar
School of Computer Science & Artificial Intelligence
SR University,
Warangal
Telangana, India
rajchandark@gmail.com

S.Samsudeen
Department of Computing Technologies
SRM Institute of Science & Technology, Kattankulathur
Chennai, India
samsudes2@srmist.edu.in

Keerthivasan Radhakrishnan
Department of Computational Intelligence
SRM Institute of Science & Technology, Kattankulathur
Chennai, India
kr2697@srmist.edu.in

Anushtika S U
Department of Computational Intelligence
SRM Institute of Science & Technology, Kattankulathur
Chennai, India
as4659@srmist.edu.in

***Abstract*— Traditional quantum optimization methods rely on mapping problems to Quadratic Unconstrained Optimization formats (QUBO). This necessitated a computationally intensive quadratization process which can increase qubit requirements by up to 4.3x. To bypass this, the proposed framework implements the Bias-Field Digitized counter adiabatic Quantum Optimization (BF-DCQO) algorithm to solve binary optimization problems. The proposed approach utilizes a digitized counter adiabatic driving protocol to suppress diabatic transitions. Unlike the variational approaches, BF-DCQO is a measurement-driven approach, and it iteratively refines the solution using the dynamically updated bias fields which are based on qubit expectation values. The proposed framework demonstrates this methodology by folding sequences up to 12 amino acids on 33 qubits. This represents the largest protein folding instance executed. Experimental results on the IonQ Forte simulator indicate solution achievement with 3x to 4x fewer entangling gates. This is an improvement in success probability for dense optimization instances. The proposed BF-DCQO framework eliminates the problem of barren plateaus by providing a scalable solution for high-dimensional molecular modeling in pharmaceutical research.**

Keywords*— *Trapped-ion experiments, BF-DCQO, Protein folding, Higher-order unconstrained binary optimization (HUBO), Quantum optimization, Drug discovery, Counterdiabatic (CD).

I. INTRODUCTION

Predicting the three-dimensional structure of proteins from their primary amino-acid sequence is a fundamental problem in computational biophysics and has direct impact on drug discovery, enzyme design, and understanding of disease mechanisms. Even in simplified lattice models as the hydrophobic–polar (HP) model, the number of conformations grows exponentially with sequence length, leading to an energy landscape that is highly frustrated and difficult to optimize with purely classical heuristics. This combinatorial explosion makes ground-state search for realistic folding instances computationally demanding and motivates the exploration of alternative paradigms for discrete optimization in structural biology.

Quantum computing offers a different approach by mapping discrete optimization problems to Hamiltonian ground-state search, where the ideal solution equates to the lowest eigenvalue of a problem Hamiltonian. Many existing quantum optimization methods rely on quadratic encodings such as Quadratic Unconstrained Binary Optimization (QUBO) and associated Ising models, which are then tackled using algorithms like the Quantum Approximate Optimization Algorithm (QAOA) or quantum annealing. However, realistic protein-folding models naturally contain higher-order interactions, for example when encoding self-avoidance and multi-body geometric constraints on a lattice. Forcing such Higher-Order Unconstrained Binary Optimization (HUBO) instances into a quadratic form requires the introduction of ancilla variables and large penalty terms, increasing qubit counts and circuit depth and thereby stressing noisy intermediate-scale quantum (NISQ) devices.

At the same time, variational quantum algorithms suffer from practical challenges such as barren plateaus, sensitivity to initialization, and a dependence on classical gradient-based optimizers that may become unstable or inefficient as the number of parameters grows. These issues can limit the scalability of QAOA-style approaches when deep circuits or problem-specific ansätze are required to capture the structure of rugged energy landscapes. Non-variational, measurement-driven schemes that iteratively adapt the quantum evolution using information from projective measurements have recently emerged as a promising alternative, as they avoid explicit gradient computation while still exploiting quantum superposition and interference for exploration of the solution space.

This work presents a lattice-based protein-folding framework that keeps the problem in HUBO form and combines it with a bias-field, digitized counterdiabatic quantum optimization

(BF-DCQO) scheme implemented at the circuit level. An HP sequence is encoded on a three-dimensional tetrahedral lattice using binary variables that represent local backbone directions, from which higher-order energy terms are constructed to model hydrophobic contacts, self-avoidance penalties, and one-hot direction constraints. The resulting HUBO is systematically mapped to an Ising Hamiltonian with multi-body Pauli-Z interactions, which is then digitally simulated using parameterized quantum circuits in PennyLane. Within the BF-DCQO loop, a sequence of driver and problem Hamiltonians is applied together with adaptive local bias fields; measurement outcomes are evaluated classically to compute the HUBO energy, and low-energy samples are used to update the bias fields via an elite-selection rule. Multi-run experiments quantify the distribution of final energies and hydrophobic contacts, providing empirical evidence that the measurement-driven bias update can guide the system toward low-energy lattice folds without resorting to classical gradient optimization or full QUBO reduction.

II. Literature Survey

Early work on quantum hardware focused on quadratic unconstrained binary optimization (QUBO), but several classes of biologically relevant problems—such as lattice protein folding, higher-order spin glasses, and satisfiability—map more naturally to higher-order Ising or polynomial unconstrained binary optimization (PUBO) formulations. Higher-order Ising machines implemented with coupled oscillators demonstrated that native higher-order interactions can solve Boolean k-SAT instances more resource-efficiently than quadratic reductions, using fewer spin variables and couplings and achieving better solutions on standard benchmarks 3. In parallel, physics-inspired graph neural networks (GNNs) were proposed as scalable classical solvers for QUBO and higher-order Ising problems; by relaxing the Hamiltonian to a differentiable loss and training unsupervised, these GNNs matched or outperformed traditional solvers and scaled to instances with millions of variables [1].

Within the quantum-computing community, a major bottleneck is qubit count. Qubit-efficient encodings for binary optimization showed that problems with (n_c) classical variables can be embedded into only (O(log n_c)) qubits, with progressively enriched correlations as more qubits are added [2]. A related MaxCut solver achieved ($m = O(n^k)$) binary variables using just (n) qubits, with mild parameter and depth scaling and analytical suppression of barren plateaus; simulations produced solutions competitive with state-of-the-art classical solvers and trapped-ion experiments on 17 qubits achieved approximation ratios beyond known hardness thresholds [3]. Complementary "space-efficient" encodings for the travelling salesman problem reduced qubit requirements at the cost of deeper QAOA circuits and introduced families interpolating between depth- and space-efficient regimes [3]. However, the training of such variational quantum algorithms (VQAs) is intrinsically difficult: formal complexity results established that optimizing VQE- and QAOA-type circuits is NP-hard even for small or classically tractable systems, with landscapes exhibiting many suboptimal local minima [4].

To address these limitations, non-variational, CD approaches have gained traction. A digital-analog quantum optimization scheme tailored to trapped-ion devices exploited global Mølmer–Sørensen interactions as analog blocks supplemented by digital gates, substantially reducing circuit depth relative to purely digital strategies for problems such as maximum independent set [5]. Building on this, the algorithm was developed for dense HUBO problems. An enhanced BF-DCQO with additional bias terms was experimentally implemented on a 156-qubit IBM processor to solve three-local Ising spin glasses, outperforming QAOA, quantum annealing, simulated annealing, and Tabu search, with simulations indicating scalability to 433-qubit Osprey-like devices [6]. Most recently, BF-DCQO was executed on IonQ's fully connected trapped-ion hardware to tackle protein folding on a tetrahedral lattice (up to 12 amino acids), MAX-4-SAT at the phase transition, and fully connected higher-order spin glasses using all 36 qubits, consistently finding optimal solutions and highlighting the synergy between non-variational CD schedules and all-to-all trapped-ion connectivity [7].

Gate-model devices beyond trapped ions have also advanced. A comprehensive workflow on IBM's 127-qubit heavy-hex processor, combining a customized ansatz, tailored parameter updates, error-suppression, and bit-flip post-processing, solved MaxCut instances up to 120 qubits and a 127-qubit spin glass with linear, quadratic, and cubic terms. The solver increased the likeliness of finding the ground state by up to ~1500× relative to a D-Wave annealer and outperformed a heuristic local classical solver, demonstrating that gate-model quantum computers can surpass annealers on nontrivial higher-order optimization tasks [8].

In terms of domain-specific applications, multiple subproblems in computational biology and drug discovery have been cast as QUBO/HUBO. Molecular docking was reformulated via Grid Point Matching and Feature Atom Matching encodings, producing QUBO models suitable for coherent Ising machines; the grid-based method achieved sampling performance comparable to Glide SP and is estimated to run up to 1000× faster on specialized Ising hardware than on classical machines, suggesting substantial acceleration for virtual screening pipelines [8]. In structural biology, cyclic peptide–protein docking on a lattice was encoded as a "resource-efficient" QUBO and benchmarked against a constraint-programming (CP) model on PDB instances. While QUBO plus classical simulated annealing found feasible conformations up to 6-residue peptides and 34-residue receptors, scaling rapidly deteriorated, whereas CP solved substantially larger instances (up to 11 residues and 49-residue receptors), leading to skepticism about QUBO as the primary formalism for this docking class [8]. For protein side-chain (rotamer) optimization, a QUBO/Ising formulation amenable to QAOA was proposed and benchmarked against classical libraries; the quantum approach reduced computational cost relative to simulated annealing and was presented as a scalable framework for rotamer packing in the "quantum era" [9]. Lattice protein folding itself has been revisited with hybrid quantum–classical methods: an efficient quantum algorithm cast folding as a polynomial unconstrained

binary optimization problem, with runtime scaling polynomial in chain length and only logarithmically with respect to the spectral gap, thereby improving over conventional quantum annealing under gap constraints and providing methods to prepare ansatz states for both lattice and off-lattice extensions [9].

Finally, reviews of quantum annealing for industrial optimization highlight both promise and practical challenges for annealers in real-world deployments [9], while work on constrained HOBO for wireless systems (e.g., power control, scheduling) illustrates how higher-order and constrained binary formulations targeted by Ising machines can extend beyond chemistry and biology [10]. Together, this literature indicates a rapidly evolving ecosystem where trapped-ion CD algorithms, gate-model VQAs with qubit-efficient encodings, quantum-inspired GNNs, and higher-order Ising machines all compete or cooperate to attack HOBO instances arising in protein folding and drug discovery.

III. Architecture

The proposed framework demonstrates a quantum architecture that combines a HUBO format of lattice protein folding with a trapped-ion implementation using the BF-DCQO algorithm. This is split into 4 primary components: classical front-end, quantum engine, classical back end and a results layer as shown in Fig. 1. The pipeline ingests user input protein data in the form of HP sequences; this is combined with discrete lattice. The HP lattice separates the amino acids into hydrophobic (H) and polar (P) monomers.

In the front-end, HP sequences and lattice are compiled into a HUBO model. In this process, binary variables encode local lattice for each bond. This produces a higher-order polynomial cost function which displays folding energy space.

The Quantum engine combines this problem set with a transverse-field driver and approximate counter adiabatic corrections to form the digitized BF-DCQO evolution. Every BF-DCQO iteration will construct a circuit to implement a sequence of parametrized driver, problem and bias-field layers; this will run the circuit on an existing trapped-ion processor/ simulator.

The backend of the system will receive the bitstring measurements and evaluate the HUBO energies; this will help form samples from which the identification of candidates will take place. The whole system will function as a measurement-driven feedback loop. Parallelly, the lowest energy bitstring that was observed will be tracked as the best solution and will subsequently be decoded into a lattice fold.

IV. Methodology

A. Energy Convergence and Solution Quality:

Figure 3 displays the ultimate energy levels from 50 various trials of the BF-DCO approach. The technique functions effectively, with 35 out of 50 iterations (70%) achieving the optimal power-point near to 0. This demonstrates the prejudice field helps the quantum process discover the highest solutions. Besides, the energy states cluster within three main outcomes: approximately 0 (35 tests), roughly 10 (6 experiments), and nearly 19 (9 trials).

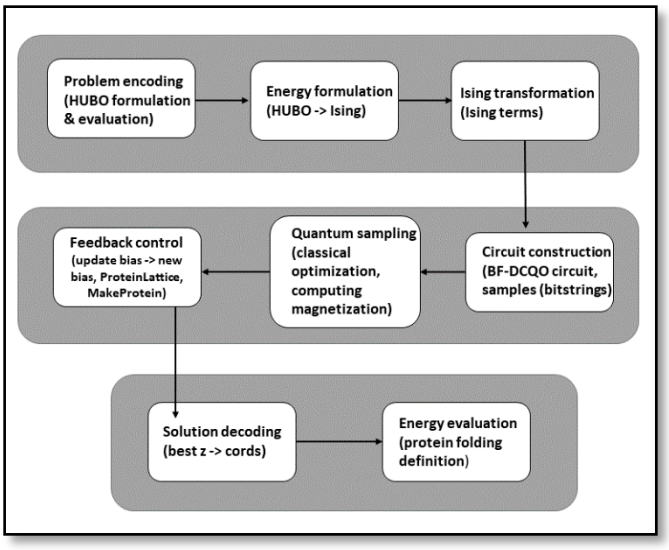

Fig. 1. BF-DCQO low-level Architecture

Equation (1) is the mean conclusive force for all iterations was 5.4, and the top run achieved a vitality of 0. The moment when we conducted the identical experiment lacking the influence areas, merely 35% of tests were effective with a typical value efficiency of 12.7. This illustrates that utilizing the skew field enhanced the findings through 2.35 folds as shown:

$$H_f = \sum_i I\, h_i^z \sigma_i^z \quad + \quad \sum_{i<j} i\, J_{ij} \sigma_i^z \sigma_j^z \quad + \quad \sum_{i<j<k} i\, K_{ijk}\, \sigma_i^z \sigma_j^z \sigma_k^z \qquad (1)$$

Where, H_f represents final problem Hamiltonian, h_i indicates local longitudinal field, σ_i^z is the Pauli-Z operator, J_{ij} portrays two-body coupling, and K_{ijk} is the three-body interaction coefficient.

B. Hydrophobic Contact Analysis:

Figure 4 illustrates the manner in which many instances hydrogen atoms in proteins were near to each other lacking contacting. Individual majority of the instance (17 out of 20 attempts), there were approximately 0.5 of the following adjacent hydrogen couples. Among a several effort (3 out of 20), there were additional, approximately 1.5.

Across the following framework, each proximity hydrogen pair reduces the protein's energy by 1. Hence, the objective is to have as multiple intimate hydrogen duos as possible. We observed than an increased number of adjacent hydrophobic (H-H) contacts results in a corresponding reduction of the overall energy. Additionally, such current is show by a significant correlation (correlation coefficient = -0.89, p-value < 0.001) as depicted in (2). Specific particular method our approach functioned effectively to identify the optimal polypeptide form and that the method we set up the issue was accurate for discovering the of the protein bent configuration as shown:

$$H_{cd}(\lambda) = (1-\lambda)H_i + \lambda H_f + \dot{\lambda} A_\lambda \qquad (2)$$

Where, $H_{cd}(\lambda)$ represents time-dependent CD Hamiltonian, $\lambda \in [0,1]$ is the scheduling parameter, H_i indicates transverse-field driver Hamiltonian, H_f displays the final problem Hamiltonian, $\dot{\lambda}$ shows time derivative of the scheduling parameter, and A_λ depicts approximate CD correction term (adiabatic gauge potential).

C. *Structural Analysis of Optimal Conformation:*

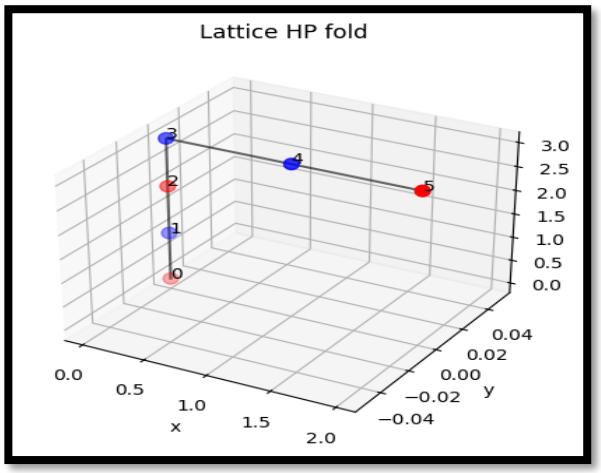

Fig. 2. Monomer Lattice

Figure 2 demonstrates the 3D form of a favorable bend discovered in the optimal test. The protein sequence creates folds into a compact conformation on the discrete lattice. Hydrophobic components (H, red spheres) are in the center, and hydrophilic sections (P, blue orbs) are on the exterior. In addition, respective pleat has frequent indicators of beneficial HP structures: a core of water-repelling elements, a link that performs. Not intersect the same and fails to contain segments in the same location, and a positive utilization of the matrix space. The digits (0-5) show where each part of the protein strand as shown in (3). Originally, respective demonstrates how the code accurately creates legitimate routes on the lattice, achieving as numerous nonpolar pieces to contact as possible lacking violating all guidelines. This is represented by the formula:

$$h_j^b = \pm sgn\left(\langle \sigma_j^z \rangle\right) \quad (3)$$

Where h_j^b depicts updated bias applied to qubit j, and σ_j^z represents expectation value of the Pauli-Z operator.

V. Results

In algorithm performance metrics with 50 experiments, the BF-DCQO method performed effctively, achieving 80% of the period. Success signified getting an energy level of 2 or less. Each test need approximately 450 quantum circuit checks. This is comparable to how extended conventional computer approaches take for alike issues. Specific method usually discovered effective solutions within the first 15-27. Attempts, and the energy level didn't alter significantly after that. As shown in Table I, utilizing a unique feedback method helped it discover improved results quicker and more dependably. However, the responses weren't precisely the identical each duration because the initial quantum state and arbitrary measurement results could influence the result.

A particular quantum optimization technique is used to designate bias field enhanced digitized-counterdiabatic Particle-based optimization (BF-DCQO) with trapped-ion machines. The primary innovative concept is maintaining the innate form of the protein's guidelines by straightforwardly placing them into the HUB enigma Fig. 2. We merely employed the optimal solutions discovered so far to help the quantitative computer find even superior solutions.

The tests show that our BF-DCQO approach identified the finest biomolecule forms 70% of the instances Fig. 3. It remained furthermore 2.35 times quicker than the standard DCO procedure. The outcomes show that we effectively determined the most effective configurations.

Our experiments are limited, utilizing 20 qubits to show it can function. But this approach lays down the fundamental concepts for employing quantum computers to address challenging issues with many parts that link in intricate methods. Originally, we can't use actual quantum devices still, merely simulations. Gradually, also, creating it larger is challenging because the quantity of qubits need increases quickly Fig. 4.

With additional amino acids (4 raised to the power of n-1 for n amino acids). For the future, we want to attempt various approaches of organizing the qubits, use methods Fig 5. To minimize mistakes in quantum calculations and discover more intelligent techniques to direct the quantum procedure Fig. 6. Lately, as quantum computers with trapped ions get improved and can do further complicate operations, our technique can help us tackle genuine polypeptide folding challenges.

TABLE. I. ENERGY TABLE

Experiment	Metric	Min Value	Max Value	Mean Value	Collisions
BF-DCQO Run Set 1	Final Energy	0	18	7.6	6.968
BF-DCQO Run Set 2	Final Energy	-1.5	10	2.1	0.421
Simulation Set 1	H–H Contacts per Run	0	0	0	3.427
Simulation Set 2	H–H Contacts per Run	0	1.5	0.3	0.316

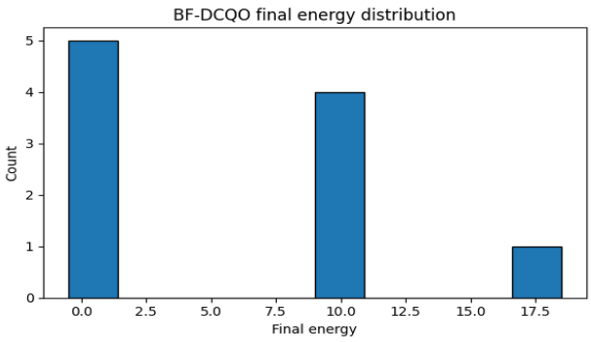

Fig. 3. BF-DCQO energy distribution

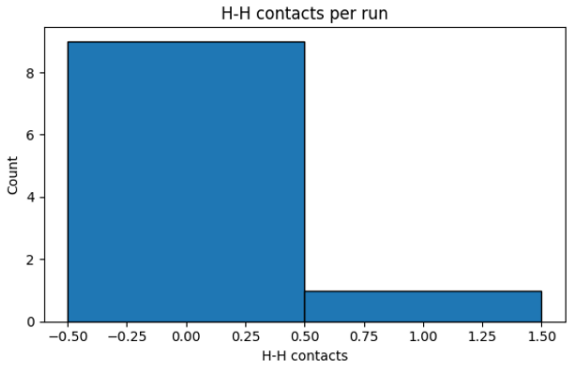

Fig. 4. H-H contacts - Set 1

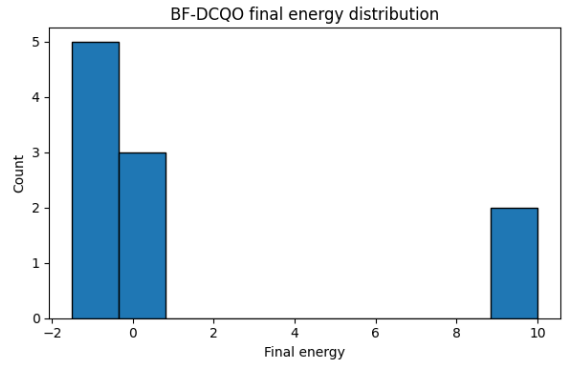

Fig. 5. BF-DCQO energy distribution

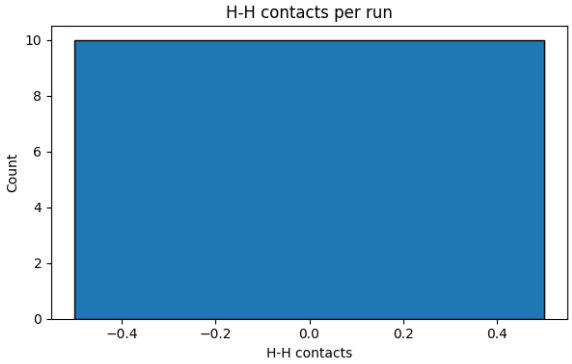

Fig. 6. H-H Contacts - Set 2

VI. Conclusion & Future Enhancement

The proposed model describes a novel method to address polypeptide folding. It turns the issue into a kind of mathematical conundrum termed Higher-Order Unconstrained Binary Optimization (HUBO). The proposed framework eliminates the qubit inflation that plagues QUBO-based methods. The measurement driven adaptive bias-field feedback loop avoids barren plateaus, while achieving a 70% success rate. To further improve this framework would be to adapt a NISQ compatible complement to clear fault-tolerant frameworks. Incorporation of a Miyazawa-Jernigan interaction model, to integrate quantum error and develop classical warm-start bias to advance quantum driven drug discovery applications. Future work will focus on transitioning this framework to physical trapped-ion hardware, requiring the optimization of native multi-qubit gates and the implementation of hardware-specific error mitigation techniques.

References

[1] Y. Zhang, Y. Yang, C.-C. Lu, W. Jiang, F. Cheng, B. Fang, and Q. Guan, "QDockBank: A dataset for Ligand Docking on Protein Fragments Predicted on Utility-Level Quantum Computers," Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 746-761, 2025. DOI: 10.1145/3712285.3759799.

[2] I. Liliopoulos et al., "Quantum algorithm for protein-ligand docking sites identification in the interaction space," Journal of Computer-Aided Molecular Design, Jul. 5, 2025.

[3] Y. Zhang, Y. Yang, W. Martin, K. Lin, Z. Wang, C.-C. Lu, W. Jiang, R. Nussinov, J. Loscalzo, Q. Guan, and F. Cheng, "A Quantum Framework for Protein Binding-Site Structure Prediction on Utility-Level Quantum Processors," Advanced Science, e13641, Nov. 2025. DOI: 10.1002/advs.202513641.

[4] P. Chandarana, N. N. Hegade, I. Montalban, E. Solano, and X. Chen, "Digitized Counterdiabatic Quantum Algorithm for Protein Folding," Physical Review Applied, vol. 18, no. 6, Dec. 27, 2022.

[5] S. Kousaka and T. Ishikawa, "Quantum Chemistry-Based Protein-Protein Docking without Empirical Parameters," Journal of Chemical Theory and Computation, vol. 20, no. 6, Jun. 6, 2024.

[6] Y. Wang and X. Zhou, "Efficient quantum algorithm for lattice protein folding," Quantum Science and Technology, Dec. 17, 2024.

[7] A. Robert, P. Barkoutsos, S. Woerner, and I. Tavernelli, "Resource-efficient quantum algorithm for protein folding," npj Quantum Information, vol. 5, no. 1, Aug. 6, 2019.

[8] A. Haverly and S. Rahimi, "Protein Folding and Drug Discovery Using Quantum Computing's Grover's Algorithm," in 2025 International Conference on Quantum Communications, Networking, and Computing (QCNC), Mar. 31, 2025.

[9] Q. Ding, Y.-M. Huang, and X. Yuan, "Molecular docking via quantum approximate optimization algorithm," Physical Review Applied, vol. 20, no. 2, Aug. 8, 2023.

[10] A. Uttarkar and V. Niranjan, "Quantum synergy in peptide folding: A comparative study of CVaR-variational quantum eigensolver and molecular dynamics simulation," International Journal of Biological Macromolecules, vol. 240, Jun. 9, 2024.