



Quantum drug discovery: a hybrid quantum graph neural network–variational quantum eigensolver framework for serine neutralization

Javier Villalba-Díez^{1,2,a}

¹ Fakultät Wirtschaft, Hochschule Heilbronn, Max-Planck-Str.39, 74081 Heilbronn, Baden-Württemberg, Germany

² Department of Mechanical Engineering, Universidad de La Rioja, Edificio Departamental, c/ San José de Calasanz, 31, 26004 Logroño, La Rioja, Spain

Received 5 April 2025 / Accepted 8 June 2025

© The Author(s) 2025

Abstract. Quantum computing has emerged as a powerful tool for modeling molecular systems with high accuracy, accelerating the search for novel therapeutic targets. In this study, we present a two-stage hybrid workflow that combines a *Quantum Graph Neural Network* with a *Variational Quantum Eigensolver* to identify potent serine neutralizers in the QM9 dataset. In the first stage, an advanced quantum graph neural network architecture incorporating attention layers, self-distillation, and an adaptive learning-rate schedule is trained to predict ionization potentials and binding free energies. Across five independent random-seed trials, our adaptive-thresholded QGNN–VQE pipeline achieves an average R^2 of 0.990 ± 0.008 and a mean absolute error of 0.034 ± 0.001 eV ($\approx 0.79 \pm 0.03$ kcal/mol) on the QM9 validation set, demonstrating robust, chemical-accuracy-level predictions. Building on these high-precision predictions, the second stage employs a QAOA-inspired hybrid ranking scheme that merges quantum graph neural network outputs, feature-space similarity (via PCA and cosine similarity), and variational quantum eigensolver-derived energy stability. This α -weighted ($\alpha = 0.95$) scoring framework is applied to over 133,000 molecules, efficiently parsed in parallel. The final ranked list is validated through an improved bijection test, featuring an adaptive similarity threshold to ensure both structural and functional robustness. The top-scoring candidate identified by our pipeline is *5,6,7-tetrahydro-4 H-pyrazolo[4,3-c]pyridin-4-one*, highlighting the efficacy of quantum-enhanced modeling in pinpointing complex pharmacological targets. These findings underscore the transformative potential of hybrid quantum-classical workflows in drug discovery, offering a flexible blueprint for advanced screening of other critical biomolecular interactions in the rapidly expanding field of quantum biology.

1 Introduction

Drug discovery currently stands at a pivotal juncture, where the promise of advanced computational strategies intersects with the constraints of traditional methods [1]. Conventional *in silico* tools, such as molecular docking [2] and quantitative structure–activity relationship (QSAR) models [3], have achieved important milestones, yet they often rely on approximations that neglect fine-grained quantum effects in molecular systems. Given that many pharmacologically relevant targets involve intricate atomic and electronic configurations, a more holistic approach is needed to capture subtle effects that can dictate drug efficacy [4]. In response, quantum computing has emerged as a compelling avenue for enhancing future drug-discovery pipelines [5].

One particularly intriguing target for these methods is serine, an amino acid with extensive involvement in metabolism and enzymatic catalysis [6]. The functional importance of serine residues—ranging from contributions to catalytic triads in proteases [7] to modulation of key signaling pathways [8]—makes selective inhibition of serine residues vital for numerous therapeutic interventions [9]. However, designing potent and selective serine blockers remains nontrivial; dynamic protein–ligand interactions, solvent-mediated effects, and electron correlation all shape the efficacy of such inhibitors [10–12].

Recent advancements in quantum computing promise to address these complexities [13,14]. By potentially surpassing classical computational limitations, quantum devices can more accurately approximate electronic structure, thereby offering substantial speedups or improved precision in molecular simulations [15]. Though current quantum processors suffer from issues of noise, decoherence, and limited qubit counts, hybrid quantum–classical paradigms help mitigate these chal-

^a e-mail: javier.villalba-diez@hs-heilbronn.de (corresponding author)

lenges [16]. In particular, the variational quantum eigensolver (VQE) partitions intensive electronic energy computations across quantum hardware, while classical optimizers adaptively refine circuit parameters [17]. Similarly, the quantum approximate optimization algorithm (QAOA) aims to tackle combinatorial tasks, including those arising from lead-compound ranking and selection [18].

Concurrently, graph neural networks (GNNs) have become prominent in computational drug discovery by encoding molecules as node-and-edge representations [19]. These models capture bonding interactions at both local and global scales, extending naturally to spatially resolved architectures [20]. Yet conventional GNNs often fail to incorporate deeper electron-correlation effects [21]. Quantum graph neural networks (QGNNs), which incorporate quantum-derived features and computations [22], seek to overcome these limitations by merging quantum-level accuracy with the representational power of neural networks.

In this paper, we propose a holistic strategy to discover and rank prospective serine inhibitors by fusing QGNN architectures, VQE computations, and a QAOA-inspired optimization framework. As a primary resource, we utilize the QM9 repository, a well-curated database of small molecules that includes extensively characterized chemical properties [23]. In particular, ionization potential (IP) [24] and binding free energy (BFE) [25] serve as useful electronic descriptors that can hint at how a molecule might interact with and modulate serine residues. Data preprocessing involves dimension reduction (using principal component analysis, PCA) and standardization, along with coordinate perturbations for data augmentation, yielding robust training sets without necessitating additional experimental or theoretical inputs.

Our approach proceeds through multiple interlinked stages. Initially, an attention-based QGNN is trained to predict properties such as IP and BFE, with the hypothesis that these features correlate with inhibitory action on serine residues. By focusing on these properties, the QGNN learns to detect electronic and structural attributes that may underlie strong binding interactions. We then refine the output rankings using classical similarity metrics, highlighting compounds that, despite high QGNN-predicted efficacy, might deviate from conventional chemical scaffolds—potentially uncovering novel modes of action. Conversely, molecules with mundane structural profiles and no apparent competitive advantage are deprioritized.

Next, VQE-driven energy estimates provide an additional quantum mechanical viewpoint on system stability. Although current quantum hardware can only handle relatively small-scale or approximate Hamiltonians, even partial insights into electronic structure can guide decision-making by flagging potential electronic or conformational instabilities. We then feed these energies into a QAOA-inspired ranking mechanism, defining a hybrid objective function that weighs QGNN scores, molecular similarity, and quantum-derived metrics. The intention is to produce a prioritized list that accounts

for multiple aspects of efficacy and stability, rather than relying on a single predictive metric.

To ensure the robustness of these predictions, we adopt a bijection test with an adaptive threshold. This test verifies that highly ranked compounds maintain their standing even under random dataset perturbations or alternative initializations. By setting thresholds based on the median similarity scores of top candidates, our method distinguishes true serine inhibitors from spurious outliers. A molecule that repeatedly appears in the top rank, sharing meaningful structural similarities with other high performers, is more likely to be a genuine candidate deserving further validation.

These innovations align with broader trends in the pharmaceutical sector, where artificial intelligence [26] and quantum computing [27] are driving new partnerships and research directions. Although practical quantum advantage remains a subject of ongoing investigation, incremental improvements—via hybrid architectures and judicious multi-scale modeling—are making near-term benefits increasingly accessible. Admittedly, challenges persist. Noise and decoherence continue to hamper quantum hardware, and capturing protein–ligand binding in full detail will require not only electronic insights but also comprehensive solvent, entropic, and conformational analyses. Multi-scale strategies, combining coarse-grained methods with quantum precision for the most critical regions, offer a plausible trajectory forward.

In the realm of machine learning, a key hurdle is ensuring that models generalize to novel chemical scaffolds [28]. In response, our pipeline systematically applies data augmentation techniques to fortify the QGNN against overfitting [29], while also mixing classical molecular similarity metrics with quantum predictions. This balanced approach mitigates the risk of highly ranked molecules being mere artifacts of noise or mismatch between training and testing distributions.

In quantum-chemistry applications, “chemical accuracy” is universally defined as an absolute error below 1 kcal/mol (≈ 0.043 eV) because only at this level can computational predictions reliably guide experimental decision-making [30]. Our model’s mean square error (MAE) of ≈ 0.5 kcal/mol (0.02 eV) meets this benchmark, bridging the gap between high statistical correlation ($R^2 \approx 0.98$) and the absolute-error requirements of real-world molecular design.

The paper is structured as follows. Section 2 details the mathematical and algorithmic underpinnings of our QGNN, including attention modules, quantum feature layers, and advanced training schemes such as self-distillation. It also explains how VQE is integrated into the pipeline, outlining approximate Hamiltonian models and strategies for incorporating quantum energy results into a unified ranking. Our QAOA-inspired ranking approach is then expanded upon, where we specify how the optimization objective fuses diverse metrics. The bijection test, equipped with adaptive thresholds, is introduced to validate the coherence and consistency of our candidate selection process. Sect. 3 discusses the technical setup for processing QM9 data

and implementing quantum simulations, highlighting both hardware and software considerations. In Sect. 4, we assess our pipeline's performance by comparing validation accuracies, stability analyses, and the impact of hyperparameter tuning. We particularly emphasize the synergy between classical algorithms and quantum modules, identifying fruitful directions for refinement. Finally, Sect. 5 summarizes key takeaways and outlines future work, contending that while our study focuses on serine inhibitors, the proposed combination of QGNN, VQE, and QAOA-inspired optimization can be generalized to other drug targets. By addressing issues ranging from data curation to quantum hardware constraints, we aim to offer a robust framework for quantum-driven molecular discovery in a clinical context.

This integrated view—uniting quantum mechanical rigor, sophisticated machine learning, and empirical domain knowledge—embodies the potential to accelerate the identification of novel serine-blocking compounds. By systematically deploying each component of the pipeline, we demonstrate how hybrid quantum-classical methodologies can transcend the limitations of purely classical approaches, driving drug discovery toward an era of more efficient and informed molecular design.

2 State of the art

2.1 Quantum chemical perspective on serine and its inhibitors

Serine ($C_3H_7NO_3$) or $(C([C@H](C(=O)O)N)O)$, shown in Fig. 1, is one of the twenty canonical amino acids found in living organisms [31]. It is classified as a polar, uncharged amino acid, primarily because of its hydroxymethyl side chain ($-CH_2OH$) bound to the central α -carbon. In many proteins, serine residues participate in hydrogen bonding and catalytic mechanisms, most famously within the active sites of serine proteases such as trypsin, chymotrypsin, and elastase [32]. These proteases exploit the nucleophilicity of the serine side chain in their catalytic triad, typically consisting of serine, histidine, and aspartate residues [33]. Serine's hydroxyl group is readily activated by the neighboring histidine residue (acting as a general base), thus generating a highly reactive alkoxide capable of forming covalent bonds with peptide substrates [34].

From a quantum chemical standpoint [35], the reactivity of serine is governed by the partial charges distributed across the α -carbon, the amino group ($-NH_3^+$ under physiological conditions), the carboxyl group ($-COO^-$ at physiological pH), and especially the hydroxyl group. The IP of the serine side chain is a key energetic descriptor indicating how easily the oxygen-centered lone pairs can be ionized or involved in electron-donating interactions. Additionally, interactions relevant to binding can be measured or estimated through BFE, reflecting the thermodynamic favorabil-




Fig. 1 Serine ($C_3H_7NO_3$) molecule representations

ity of intermolecular complex formation. In an enzyme's active site, these properties help define the strength and stability of serine's interaction with substrates or inhibitors.

A related, more general quantum chemical measure of reactivity is the *HOMO–LUMO gap* [36]. The highest occupied molecular orbital (HOMO) and lowest unoccupied molecular orbital (LUMO) represent, respectively, the frontier orbitals that are most relevant for electronic transitions and chemical reactions. In many small-molecule systems, a narrower HOMO–LUMO gap is associated with higher chemical reactivity, given that excitations or charge transfers require less energetic input. For serine, although the HOMO–LUMO gap alone does not fully describe its nucleophilic aptitude (since ionization potentials and solvation effects also play major roles), examining the frontier-orbital energies can nonetheless offer insight into how the hydroxyl oxygen might interact with electrophilic or proton-transfer partners. For example, a partially elevated HOMO level can facilitate interactions with electron-deficient sites, while the LUMO position may determine susceptibility to electron donation if the residue experiences local perturbations (e.g., from neighboring acidic groups).

In drug design focused on inhibiting serine, understanding both the IP and the HOMO–LUMO gap can thus inform how easily serine's reactive oxygen can be engaged or blocked by candidate molecules. Moreover, integrated free-energy calculations, captured by BFE, help clarify whether these orbital features translate into stable, long-lived inhibitor–residue complexes or merely transient electronic interactions. Together, the IP, HOMO–LUMO gap, and BFE form complementary quantum chemical descriptors that elucidate the electronic and thermodynamic bases of serine's reactivity in enzymatic contexts.

2.1.1 Targeting serine: key quantum properties

When designing inhibitors that target the reactive hydroxyl group of serine, one aims to either block it sterically or form stable interactions (covalent or otherwise) that prevent the residue from participating in enzyme catalysis. Quantum chemical simulations of serine typically focus on:

- Protonation States [37]: Under different conditions (varying pH, enzymatic microenvironments), serine can assume distinct protonation states, which changes the electron density and, thus, the reactivity of the hydroxyl group.
- Hydrogen-Bonding Networks [38]: Serine's polar side chain is apt to engage in multiple hydrogen bonds, both intra- and intermolecular, influencing the geometry and electronic energies of complexes with potential inhibitors.
- IP [39]: The IP can serve as an indicator of how easily electrons can be removed from (or donated by) the serine side chain, impacting its nucleophilicity within the active site.
- Partial Charge Distribution [40]: Evaluating partial charges on the hydroxyl oxygen or the α -carbon region provides insight into electrostatic complementarity with inhibitor molecules.
- Tautomeric and Conformational Flexibility [41]: Although serine is not well known for multiple tautomeric forms compared to, say, histidine, it retains important rotational degrees of freedom in the side chain, influencing local reactivity profiles.

From the viewpoint of designing inhibitors, these quantum chemical considerations translate into constraints on the electronic and geometric structures that a potential ligand should exhibit. Effective inhibitors may form strong hydrogen bonds or engage in covalent interactions with the reactive hydroxyl of serine. Additionally, estimates of binding free energies (BFE) help identify how strongly a compound is likely to bind or neutralize the serine site under physiological conditions. Balancing IP and BFE thus becomes a central task in evaluating whether a molecule is well suited to block serine functionality within an enzymatic pathway.

2.1.2 Desired properties for serine inhibitors

Designing small-molecule inhibitors that target serine residues involves balancing several quantum chemically informed properties:

- Covalent vs. Non-Covalent Binding: Some inhibitors form a covalent bond with the serine hydroxyl (e.g., serine-protease inhibitors like diisopropylfluorophosphate), whereas others rely on tight non-covalent binding. Accurate estimation of these binding events often relies on calculating or predicting BFEs and IPs for different binding conformations.
- Steric Compatibility and Conformational Stability: The inhibitor must fit into the active site in a thermodynamically favorable pose. The partial charges, hydrogen-bonding capacity, and overall 3D geometry all interact to stabilize the inhibitor-serine complex.
- Electronic Complementarity: Potential inhibitors require electron donor/acceptor sites that appropriately match the local environment of serine, as gauged by IP and charge distributions.

- Resistance to Hydrolysis or Other Side Reactions: Enzymatic environments can be highly reactive, so inhibitors should maintain chemical stability. Quantum simulations of their reaction pathways and BFE can suggest which scaffolds hold up better under catalytic conditions.

In sum, from a quantum chemical perspective, both IP and BFE constitute pivotal measures for understanding how a small molecule may effectively interact with serine's reactive site. The core methodology developed in this paper leverages machine learning and quantum-enhanced calculations to identify candidates with desirable profiles for these two properties.

2.2 Variational quantum eigensolver (VQE)

The VQE is a hybrid quantum-classical algorithm designed to compute low-lying eigenvalues of quantum systems, particularly the ground-state energy. In molecular systems, VQE is often used to approximate electronic energies for fragments relevant to binding or reactivity.

2.2.1 Mathematical formulation of VQE

At the heart of the VQE is the task of finding the ground-state energy of a given molecular Hamiltonian using a hybrid quantum-classical approach. In second quantization, the electronic Hamiltonian for an N -electron system in a chosen basis can be written as:

$$\hat{H} = \sum_{p,q} h_{pq} \hat{a}_p^\dagger \hat{a}_q + \frac{1}{2} \sum_{p,q,r,s} h_{pqrs} \hat{a}_p^\dagger \hat{a}_q^\dagger \hat{a}_r \hat{a}_s, \quad (1)$$

where \hat{a}_p^\dagger and \hat{a}_p denote fermionic creation and annihilation operators for the spin-orbital p , and h_{pq} , h_{pqrs} are one- and two-electron integrals, respectively.

To implement this Hamiltonian on a quantum computer, one applies a fermion-to-qubit mapping such as the Jordan-Wigner or Bravyi-Kitaev transformation. In general form, this yields a qubit Hamiltonian:

$$\hat{H} = \sum_i \alpha_i \hat{P}_i, \quad (2)$$

where each \hat{P}_i is a tensor product of Pauli operators (e.g., $\sigma_x, \sigma_y, \sigma_z$) acting on different qubits, and α_i are real coefficients determined by the integrals in Eq. (1). Because there are often many terms in \hat{H} , one typically measures each term's expectation value separately and sums the results to approximate the total energy.

To find the ground state, VQE employs a parameterized quantum circuit:

$$|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta}) |\psi_0\rangle, \quad (3)$$

where $|\psi_0\rangle$ is an initial (often simple) reference state such as the Hartree-Fock state, and $U(\boldsymbol{\theta})$ is composed

of one- or two-qubit gates with tunable parameters $\theta = (\theta_1, \theta_2, \dots, \theta_k)$. Common examples include “hardware-efficient” ansätze (a series of single-qubit rotations and entangling operations repeated in blocks) or chemically inspired ansätze such as the *unitary coupled cluster singles and doubles (UCCSD)* approach. In our implementation, we choose the UCCSD ansatz, defined by the unitary

$$U_{\text{UCCSD}}(\mathbf{t}) = \exp[\hat{T}(\mathbf{t}) - \hat{T}(\mathbf{t})^\dagger], \quad (4)$$

with

$$\hat{T}(\mathbf{t}) = \sum_{i,a} t_i^a \hat{a}_a^\dagger \hat{a}_i + \frac{1}{4} \sum_{i<j, a<b} t_{ij}^{ab} \hat{a}_a^\dagger \hat{a}_b^\dagger \hat{a}_j \hat{a}_i,$$

where indices i, j run over occupied orbitals and a, b over virtual orbitals. After mapping to qubits, each excitation term is compiled into native rotation and entangling gates, yielding a circuit of depth eight on six qubits for our fragment calculations.

Once the state $|\psi(\theta)\rangle$ has been prepared, the energy is given by:

$$E(\theta) = \langle \psi(\theta) | \hat{H} | \psi(\theta) \rangle = \sum_i \alpha_i \langle \psi(\theta) | \hat{P}_i | \psi(\theta) \rangle. \quad (5)$$

Because \hat{P}_i are tensor products of Pauli matrices, this expectation value is measured by performing the corresponding basis rotations on the quantum hardware and collecting measurement statistics across repeated circuit executions.

A classical optimizer (e.g., gradient-based or gradient-free) iteratively updates θ to minimize the cost function:

$$\theta^* = \arg \min_{\theta} \{E(\theta)\}. \quad (6)$$

For gradient-based optimizers, one can employ the parameter-shift rule to compute partial derivatives w.r.t. each gate parameter:

$$\frac{\partial E}{\partial \theta_j} = \frac{1}{2} \left[E(\theta + \frac{\pi}{2} \hat{e}_j) - E(\theta - \frac{\pi}{2} \hat{e}_j) \right], \quad (7)$$

where \hat{e}_j is the unit vector in the direction of parameter θ_j . This rule requires additional function evaluations but avoids finite-difference pitfalls.

While computing the entire electronic structure of large biomolecules remains prohibitive on near-term devices, VQE can be applied to critical subsystems or simplified models (e.g., the interaction region between a candidate inhibitor and a serine side chain). One can then extract approximate energy differences between free and bound states to inform BFE estimates. Coupled with classical post-processing (for instance, thermodynamic corrections or continuum solvent models), these VQE-derived energies reveal essential quantum features that might otherwise be overlooked.

For IP, one may compute the ground-state energy of a neutral fragment and compare it with the ground-state energy of a singly ionized version of the same fragment. The difference gives an approximate IP:

$$\text{IP} \approx E_{\text{VQE}}^{(+1)} - E_{\text{VQE}}^{(0)}, \quad (8)$$

where $E_{\text{VQE}}^{(0)}$ is the energy of the neutral system, and $E_{\text{VQE}}^{(+1)}$ is that of the cationic system. Although the absolute accuracy of these calculations depends on the choice of ansatz and hardware noise levels, such relative energies often capture critical trends that correlate with experimental data.

In summary, the VQE framework offers a practical route to incorporate quantum mechanical fidelity into the study of serine inhibitors, particularly for assessing how changes in electronic structure (reflected in IP shifts or local BFE contributions) influence the binding characteristics of candidate molecules. Even on noisy hardware, careful circuit design and error-mitigation strategies can yield insights beyond what classical force-field methods alone can provide.

2.2.2 VQE and serine inhibitors

In this context, we use VQE-based energies as a proxy for molecular stability and binding interactions in a localized fragment model, focusing on partial systems that include the serine side chain and key functional groups of candidate inhibitors. These estimated energies can complement classical methods for BFE calculations and serve as a quantum-derived feature in a machine learning pipeline. Although direct modeling of large biomolecular complexes is not feasible with near-term devices, even the partial VQE estimates help refine predictions of inhibitor efficacy by providing additional quantum mechanical detail not captured in purely classical force fields.

2.3 Quantum approximate optimization algorithm (QAOA)

The QAOA addresses combinatorial optimization tasks by encoding the cost function into a quantum circuit, typically focusing on bitstring solutions that minimize some objective. Although ranking molecular candidates is not strictly a bitstring optimization in classical form, the conceptual framework of QAOA remains useful for systematically combining multiple scoring components.

2.3.1 Mathematical overview of QAOA

The QAOA is designed to find approximate solutions to discrete optimization problems by encoding a classical cost function $C(x)$, defined over bit strings $x \in \{0, 1\}^n$, into a quantum mechanical framework. One typically seeks to minimize $C(x)$ or, equivalently, maximize $-C(x)$. In what follows, we provide a deeper look into the mathematical underpinnings of QAOA.

Given a cost function $C(x)$, QAOA represents it as a Hamiltonian \hat{C} that is diagonal in the computational basis $\{|z\rangle : z \in \{0, 1\}^n\}$. Concretely:

$$\hat{C} = \sum_{z \in \{0,1\}^n} C(z) |z\rangle\langle z|. \quad (9)$$

Measuring the state $|\psi\rangle$ in the computational basis yields a bitstring z with some probability amplitude, and $C(z)$ can be interpreted as the “cost” of that measurement outcome.

To move between different bitstring configurations, QAOA employs a “mixer” Hamiltonian \hat{B} . A common choice is the transverse-field driver:

$$\hat{B} = \sum_{j=1}^n \sigma_x^{(j)}, \quad (10)$$

where $\sigma_x^{(j)}$ is the Pauli-X operator acting on qubit j . This Hamiltonian, when exponentiated, applies X -rotations in the Bloch sphere representation, promoting transitions between computational basis states $|z\rangle$.

QAOA proceeds in layers, where each layer consists of two unitary operations: the cost unitary $U_C(\gamma)$ and the mixer unitary $U_B(\beta)$. For a single layer, the evolution is:

$$U(\gamma, \beta) = U_B(\beta) U_C(\gamma), \quad (11)$$

where the parameters γ and β are real numbers controlling the time (or angle) of evolution under each Hamiltonian. Concretely:

$$U_C(\gamma) = \exp(-i\gamma\hat{C}), \quad U_B(\beta) = \exp(-i\beta\hat{B}). \quad (12)$$

In practice, one typically applies multiple layers (often denoted p layers) to achieve better approximation quality:

$$U_p(\gamma, \beta) = \prod_{\ell=1}^p U_B(\beta_\ell) U_C(\gamma_\ell), \quad (13)$$

where $\gamma = (\gamma_1, \dots, \gamma_p)$ and $\beta = (\beta_1, \dots, \beta_p)$.

A typical choice for the initial state is the uniform superposition of all bitstrings:

$$|\psi_0\rangle = \frac{1}{2^{n/2}} \sum_{z \in \{0,1\}^n} |z\rangle, \quad (14)$$

though other problem-specific initializations are possible. By starting in $|\psi_0\rangle$ and applying the layered evolution $U_p(\gamma, \beta)$, one systematically explores the configuration space defined by \hat{C} .

After applying p layers, one measures the resulting state

$$|\psi(\gamma, \beta)\rangle = U_p(\gamma, \beta) |\psi_0\rangle$$

in the computational basis to estimate the expectation value of \hat{C} :

$$\langle \hat{C} \rangle = \langle \psi(\gamma, \beta) | \hat{C} | \psi(\gamma, \beta) \rangle. \quad (15)$$

A classical optimizer iteratively updates γ and β to minimize this expectation value—or, equivalently, to produce a final state whose bitstring samples correspond to low-cost solutions z . One may use gradient-free methods (like Nelder–Mead or COBYLA) or approximate gradients based on finite differences or other quantum-compatible techniques.

In many drug-discovery contexts, one can formulate ranking or selection tasks as combinatorial optimization problems, where each bit in z indicates whether a molecule is chosen for further consideration [42]. The cost function $C(z)$ might combine terms related to predicted binding free energies, ionization potentials, or quantum stability metrics. Although running QAOA at scale for large n is limited by current hardware constraints [43], QAOA-inspired cost formulations can still be simulated classically or applied in smaller subproblems. This ensures that the fundamental principles, iterative application of cost and mixer unitaries to navigate a solution space, guide how molecules are prioritized.

Real implementations of QAOA are subject to noise and finite circuit depth constraints [44]. Nevertheless, even at moderate depths ($p \approx 1$ to 5), QAOA has shown promise for small- to medium-sized optimization tasks. Various generalizations exist, such as adding problem-specific mixing Hamiltonians or modifying the cost Hamiltonian to embed domain knowledge. In drug discovery, for example, one can tune \hat{C} to include terms that favor molecules exhibiting both strong binding to a target site and desirable ionization characteristics, thus guiding the search for serine inhibitors in a multi-objective manner.

Overall, QAOA provides a structured quantum-classical approach to search a potentially huge combinatorial space of candidate compounds. Even if used in a simplified or “classical mimic” capacity, the QAOA perspective helps formalize how different molecular properties (e.g., IP and BFE) might be integrated and balanced in a cost function, while the iterative parameter updates systematically refine the set of proposed solutions.

2.3.2 QAOA-inspired molecule ranking

Although implementing QAOA directly on large molecular datasets remains challenging, the method’s underlying principle, optimizing a multi-term cost function using iterative quantum-classical feedback, is applicable. Suppose we define a molecular “score” $\text{Score}(M)$ that combines machine-learned predictions of IP and BFE with additional factors. Then, the goal of selecting top candidates corresponds to maximizing $\text{Score}(M)$ or equivalently minimizing $-\text{Score}(M)$. Inspired by QAOA, one can structure a classical or hybrid quantum-

classical procedure to iterate toward a set of solutions with high $\text{Score}(M)$.

2.4 Quantum graph neural networks (QGNNs) with attention

Graph neural networks (GNNs) have emerged as an efficient paradigm for modeling molecular structures, treating atoms as graph vertices V and bonds (or relevant interactions) as edges E . In a purely classical GNN, node embeddings aggregate local neighborhood information through a series of learnable message-passing operations. A QGNN extends this approach by embedding quantum-derived descriptors—such as fragment-level energy corrections or partial wavefunction features—into the node- or graph-level representations. This integration enables the network to better capture the electronic nuances essential for serine inhibition, including effects related to ionization potential (IP) and binding free energy (BFE).

2.4.1 Classical GNNs and message passing

Let $G = (V, E)$ denote a molecular graph, where each node $i \in V$ corresponds to an atom, and each edge $(i, j) \in E$ typically denotes a chemical bond or a proximity-based interaction. Each node i is assigned an initial feature vector \mathbf{h}_i^0 encoding properties such as atom type, partial charge, and other classical descriptors (e.g., valence or hybridization state). The message-passing mechanism updates node embeddings in successive layers, indexed by t , as follows:

$$\mathbf{m}_i^{(t)} = \sum_{j \in \mathcal{N}(i)} \phi(\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)}, \mathbf{e}_{ij}),$$

$$\mathbf{h}_i^{(t)} = \psi(\mathbf{h}_i^{(t-1)}, \mathbf{m}_i^{(t)}),$$

where $\mathcal{N}(i)$ is the neighbor set of node i , while ϕ and ψ are learnable functions (e.g., small neural networks). The edge-level feature \mathbf{e}_{ij} might capture bond type or interatomic distances. Iterating this process over $t = 1, 2, \dots, T$ layers allows each node embedding to incorporate increasingly distant structural information in the molecular graph.

2.4.2 Attention mechanisms in molecular graphs

In many cases, certain atom–atom interactions exert a more pronounced effect on molecular properties (IP, BFE, etc.) than others. An attention mechanism accounts for this by introducing weights $\alpha_{ij}^{(t)}$ that prioritize the most relevant connections. Concretely,

$$\alpha_{ij}^{(t)} = \frac{\exp\left(a(\mathbf{h}_i^{(t-1)}, \mathbf{h}_j^{(t-1)})\right)}{\sum_{k \in \mathcal{N}(i)} \exp\left(a(\mathbf{h}_i^{(t-1)}, \mathbf{h}_k^{(t-1)})\right)},$$

where $a(\cdot)$ is a learnable scoring function determining the “importance” of neighbor j for node i . The node-level update then becomes:

$$\mathbf{m}_i^{(t)} = \sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(t)} \phi(\mathbf{h}_j^{(t-1)}, \mathbf{e}_{ij}).$$

By varying the architecture to include multiple attention heads or stacked layers, the network can capture distinct aspects of electronic or geometric structure, such as resonant subsystems or hydrogen-bond motifs critical for blocking serine’s reactive hydroxyl group.

2.4.3 Quantum integration in QGNNs

While the above mechanisms function in purely classical GNNs, a QGNN seeks to incorporate quantum mechanically derived features to address electronic subtleties crucial for tasks like predicting serine inhibition. Mathematically, this can be achieved in various ways:

1. **Node-Level Quantum Features.** Each node embedding \mathbf{h}_i^0 may be augmented by scalar or vector quantities gleaned from quantum calculations (e.g., local energy corrections, partial charges from wavefunction-based methods). Let Δ_i denote a quantum descriptor for atom i ; then one initializes

$$\mathbf{h}_i^0 = [\mathbf{x}_i, \Delta_i],$$

where \mathbf{x}_i is the classical atom feature vector. During message passing, these hybrid embeddings enable the GNN to correlate quantum effects (e.g., local electron-density shifts) with structural connectivity.

2. **Graph-Level Quantum Embeddings.** Alternatively, one might compute a global quantum descriptor Q for the entire molecule (or a truncated active site), such as a VQE-derived stability estimate. This global scalar can be added to each node or appended to the final pooling step:

$$\mathbf{h}_G = \text{Pooling}(\{\mathbf{h}_i^{(T)} : i \in V\}) \oplus Q.$$

Here, \oplus denotes concatenation, and Q might reflect approximate IP and BFE from quantum simulations. In this context, T denotes the *number of message-passing iterations* (or *layers*) in the graph neural network. At each iteration $t = 1, 2, \dots, T$, each node in the graph aggregates information (messages) from its neighbors to update its own embedding.

3. **Quantum-Inspired Layers or Kernels.** In more advanced setups, one can substitute classical sub-layers with parameterized quantum circuits acting

on node embeddings or pairwise features. Although not explicitly implemented in the present work, such architectures conceptually replace ϕ or ψ with quantum operations. The resulting node updates, measured in the computational basis, yield quantum-enabled transformations.

In the context of serine-targeted drug discovery, quantum integration helps capture:

- Delicate hydrogen bonding that affects IP, vital for evaluating serine's nucleophilicity,
- Charge transfer patterns relevant to binding free energy (BFE), especially when dealing with partially charged or polar scaffolds,
- Local energy penalties determined by VQE-like quantum descriptors, guiding the model to down-rank molecules that appear stable classically but are quantum mechanically unfavorable near the serine site.

Formally, these quantum corrections act as additional input channels or penalty terms in the loss function, enhancing the message-passing mechanism's ability to differentiate promising serine inhibitors from otherwise similar but electronically suboptimal candidates.

As T increases, nodes can incorporate structural or feature information from progressively larger regions of the molecular graph, allowing the model to capture long-range interactions. After T rounds of message passing (with or without attention and quantum features), each node i arrives at a final embedding $\mathbf{h}_i^{(T)}$. The network often applies a global readout or pooling function:

$$\mathbf{h}_G = \text{Pooling}(\{\mathbf{h}_i^{(T)} : i \in V\}),$$

such as summation, averaging, or attention-weighted aggregation. One then feeds \mathbf{h}_G into a final predictor head (a small multilayer perceptron) to produce a single scalar output. In the code, this scalar can represent a hybrid QGNN_score that summarizes IP and BFE characteristics in one measure. Optionally, the model (or a separate script) integrates further classical descriptors (e.g., similarity measures) or quantum penalties (from VQE energies) to form a final ranking metric.

By capturing both classical connectivity patterns and quantum-calculated energy corrections, QGNN architectures with attention more accurately highlight inhibitors that engage serine's hydroxyl group in energetically favorable ways. This synergy between message-passing, attention, and quantum features lies at the heart of the high validation performance described in our experimental results.

2.4.4 Quantum-enriched GNN features

In a QGNN, quantum-derived descriptors such as approximate VQE energies or partial ionization energies can be appended to the node or graph feature

vectors, providing the network with additional inputs reflecting quantum mechanical interactions. Alternatively, partial segments of the GNN's forward pass could be executed on a quantum circuit, although this is more experimentally demanding. The overall objective is to enhance the model's predictive power regarding IP and BFE without incurring exponential computational overhead on classical hardware.

2.5 Mathematical integration of IP and BFE into the hybrid architecture

Unlike approaches that revolve around orbital-energy gaps (e.g., HOMO–LUMO), our method places IP and BFE at the center of the training and ranking pipeline. Concretely:

1. Data Representation. For each molecule M , define

$$\mathbf{x}(M) = [\mathbf{x}_{\text{classical}}(M), \mathbf{x}_{\text{quantum}}(M)],$$

where:

- $\mathbf{x}_{\text{classical}}(M)$ includes geometry, partial charges, or standard chemical fingerprints.
 - $\mathbf{x}_{\text{quantum}}(M)$ can include VQE-derived energies for a fragment relevant to BFE or an approximate measure of the molecule's IP.
2. QGNN Prediction. A QGNN (with attention layers) processes $\mathbf{x}(M)$ and outputs:

$$[\hat{y}_{\text{IP}}(M), \hat{y}_{\text{BFE}}(M)] = f_{\text{QGNN}}(\mathbf{x}(M); \Theta_{\text{QGNN}}),$$

where $\hat{y}_{\text{IP}}(M)$ is the predicted ionization potential, and $\hat{y}_{\text{BFE}}(M)$ is the predicted binding free energy. Alternatively, the model can be trained to produce a single combined score if desired.

3. Hybrid Score for Ranking. Define the hybrid score to be:

$$\text{Score}(M) = \alpha \hat{y}_{\text{IP}}(M) + \gamma \hat{y}_{\text{BFE}}(M) - \delta E_{\text{VQE}}(M) + \eta \text{Sim}(M),$$

where $E_{\text{VQE}}(M)$ may represent a more accurate (but computationally expensive) quantum energy estimation, $\text{Sim}(M)$ is a structural or chemical similarity score to known inhibitors, and $\alpha, \gamma, \delta, \eta$ are hyperparameters. The negative sign before $E_{\text{VQE}}(M)$ reflects that higher VQE energies (i.e., lower stability) penalize the overall score.

This function can be adapted to a QAOA-inspired objective by interpreting $\text{Cost}(M) = -\text{Score}(M)$, then seeking to minimize $\text{Cost}(M)$ across all molecules. Whether performed purely classically or on quantum hardware, the iterative refinement steps in QAOA align with searching for an optimal ordering or subset of high-Score(M) candidates.

2.6 Prior research and limitations

Traditional computational chemistry offers multiple ways to estimate IP and BFE. IP can be derived from quantum chemical calculations at various levels of theory (e.g., Hartree–Fock, DFT, coupled cluster) [45]. BFE for protein–ligand interactions often arises from molecular dynamics simulations paired with free-energy perturbation or thermodynamic integration [46]. While effective, these approaches can be computationally prohibitive for large-scale screening, especially when investigating thousands or hundreds of thousands of candidate molecules.

There is a growing body of literature exploring how quantum machine learning can aid in predicting electronic properties, including IP [47]. Our framework is novel in three respects:

- *End-to-end* training of an attention-enhanced QGNN whose graph embeddings directly initialize a hardware-efficient UCCSD VQE ansatz on six qubits,
- simultaneous integration of *self-distillation* in the QGNN training and an *adaptive bijection threshold* in the ranking loop—two techniques never before combined in a quantum screening pipeline,
- unification of QGNN-predicted ionization potentials and binding free energies, PCA-derived similarity scores, and partial VQE energy corrections into a single α -weighted score, reducing the classical candidate set by 75% compared to classical-only filters.

The direct synergy between quantum algorithms (like VQE, QAOA) and QGNN-based architectures tailored for large datasets (e.g., QM9) remains in early development, especially when addressing both IP and BFE in tandem.

Our method expands upon this foundation by:

- **Scalability:** Leveraging multiprocessing for parallel data parsing and augmentation of the QM9 dataset. This strategy enables the efficient handling of tens or hundreds of thousands of molecular records.
- **Quantum-Enriched Feature Set:** Integrating partial VQE-derived energies alongside predicted IP and BFE as input or intermediate features, thus marrying classical ML with quantum mechanical rigor.
- **Attention-Enhanced QGNN Architecture:** Providing a robust route to extract meaningful patterns from high-dimensional molecular data, particularly focusing on interactions affecting IP and BFE.
- **Ranking Consistency:** Employing a QAOA-inspired cost function for final ranking decisions, incorporating not just IP or BFE, but also quantum stability measures and structural similarities.

As quantum hardware matures, deeper integration of VQE for binding-site energies and direct QAOA for ranking or subset selection will become more feasible. In parallel, the synergy with advanced GNNs, including multi-head attention, positions this hybrid approach to

handle increasingly complex molecular spaces, whether for enzyme inhibitors or entirely different classes of biologically active molecules.

With this state of the art established, our paper will next provide:

- **Materials and Methods:** A step-by-step delineation of the data processing pipeline, including how IP and BFE labels are prepared or normalized, how quantum features (VQE estimates) are derived for fragments, and how QGNN training is structured with attention layers and hyperparameter optimization.
- **Results and Discussion:** Comparative analyses of model performance in predicting IP and BFE, the stability of hybrid ranking when varying hyperparameters, and the outcome of a bijection test aimed at verifying result consistency. Emphasis will be placed on how the quantum-enriched pipeline fares against classical-only baselines.
- **Conclusions and Outlook:** A synthesis of key findings and future directions, including the potential for full quantum mechanical treatments of larger biomolecular complexes and the possibility of incorporating additional quantum descriptors or multi-scale modeling techniques to improve the predictive accuracy of binding outcomes.

This framework underscores how ionization potential and binding free energy, essential quantum chemical descriptors for serine-blocking efficacy, can be seamlessly integrated into a broader quantum-classical architecture for drug discovery. The forthcoming sections thus dive into the detailed design and implementation of each component, mapping out how they converge to identify and prioritize effective serine-targeting molecules.

3 Materials and methods

In this section, we present a detailed overview of the methodological pipeline employed to integrate IP and BFE information into a hybrid quantum–classical framework. We begin by describing data acquisition and preprocessing, followed by an explanation of how molecular descriptors for IP and BFE are extracted or calculated. We then outline how quantum features (specifically, VQE estimates) are derived for relevant molecular fragments. Next, we detail the QGNN architecture, focusing on how attention layers are incorporated and how the entire model is trained and optimized. Finally, we introduce a *bijection test* that validates the consistency of top-ranked serine-blocking candidates. Our aim is to provide a complete, step-by-step explanation, enabling reproducibility and further adaptation to similar challenges in computational drug discovery.

3.1 Pseudocode overview of the hybrid QGNN–VQE pipeline

To make the workflow concrete and reproducible, we summarize the end-to-end hybrid QGNN–VQE pipeline in pseudocode. This high-level algorithm illustrates how molecular graphs are first embedded by the attention-enhanced QGNN, then passed into a short VQE loop to refine local fragment energies, and finally combined with classical similarity metrics into a single scoring function. Each step corresponds directly to the detailed descriptions above—graph construction and message passing, ansatz initialization and gradient-based optimization, dimensionality reduction for database comparison, and the final α -weighted ranking formula. By inspecting the pseudocode, readers can trace exactly where hyperparameters (e.g., ansatz depth, optimizer settings, loss weights) are applied, and can easily adapt or benchmark individual components in their own implementations.

```
// Hybrid {QGNN}-{VQE} Pipeline
// Input: molecular graph G(V,E)
// Output: final score S(M)

// Hyperparameters:
L_gnn    = 6    // number of GNN layers
H_heads  = 8    // attention heads per layer
D_hidden = 256  // hidden dimension
p_vqe    = 200  // VQE iterations
ansatz_depth = 8 // circuit depth
n_qubits = 6    // qubit count
lr_vqe   = 0.01 // VQE learning rate
shots    = 1024 // measurement shots
k_pca    = 10   // PCA components for similarity
alpha    = 0.95 // weight on IP prediction
beta     = 0.95 // weight on BFE prediction
gamma    = 0.05 // penalty on {VQE} energy
delta    = 0.10 // weight on similarity

// Stage 1: QGNN embedding and prediction
1 for each atom i in V do
2   h_i[0] = EmbedAtomFeatures(i)
3 end
4 for t = 1 to L_gnn do
5   for each node i in V do
6     m_i = 0
7     for each neighbor j of i do
8       a = AttentionScore(h_i[t-1], h_j[t-1],
9         edge_features(i,j))
10      m_i += a * Message(h_j[t-1], edge_features(i,j))
11    end
12    h_i[t] = UpdateNode(h_i[t-1], m_i)
13  end
14 h_graph = GlobalPool({h_i[L_gnn] : i in V})
15 [(h_IP), (h_BFE)] = MLPHead(h_graph)

// Stage 2: {VQE} refinement on a local fragment
16 theta_vqe = InitVQEParams(h_graph)
17 optimizer = Adam(theta_vqe, lr_vqe)
18 for iter = 1 to p_vqe do
19   U = BuildUCCSDAnsatz(theta_vqe, n_qubits,
20     ansatz_depth)
21   E = MeasureExpectation(U, H_frag, shots)
22   grads = ParameterShiftGrad(H_frag, theta_vqe)
23   optimizer.step(grads)
24 end
```

```
// Stage 3: Similarity to known inhibitors
24 h_pca = PCAProject(h_graph, k_pca)
25 sim = MaxCosineSimilarity(h_pca,
26   PCAProject(DatabaseEmbeddings, k_pca))

// Stage 4: Final hybrid scoring
26 S = alpha * h_IP
27   + beta * h_BFE
28   - gamma * E
29   + delta * sim

30 return S
```

Notes:

- **EmbedAtomFeatures**: embed atom attributes into a fixed-size vector.
- **AttentionScore, Message, UpdateNode**: small MLPs defining attention weights, message construction, and node updates.
- **GlobalPool**: mean or attention-based pooling over node embeddings.
- **InitVQEParams**: linear layer mapping graph embedding to VQE circuit parameters.
- **BuildUCCSDAnsatz**: constructs a hardware-efficient unitary coupled-cluster singles and doubles ansatz circuit of specified depth.
- **MeasureExpectation**: executes the circuit and estimates the fragment energy.
- **ParameterShiftGrad**: computes gradients exactly via the parameter-shift rule.
- **PCAProject**: PCA transformation for dimensionality reduction.
- **MaxCosineSimilarity**: finds the maximum cosine similarity against a precomputed database.

By laying out the full pipeline in pseudocode, we provide a clear template for both replication and further extension, ensuring that every hyperparameter choice and algorithmic step is transparent and readily modifiable.

3.2 Data collection and preprocessing

3.2.1 Dataset overview

The primary dataset for this study is derived from QM9, a comprehensive repository of 133.662 small organic molecules. Each entry contains:

- Atomic coordinates (e.g., Cartesian positions of each atom),
- Elementary composition (atom types: H, C, N, O, etc.),
- Quantum mechanically derived properties (such as energies, dipole moments, and partial charges, in many published versions),
- Auxiliary information relevant to structure and reactivity.

Although QM9 includes multiple electronic properties, we emphasize IP and BFE in the context of this work, as these properties are closely aligned with identifying potential serine-blocking molecules.

3.2.2 File parsing and multiprocessing

Due to the large number of XYZ files in QM9, we implemented a multiprocessing approach to parse the dataset efficiently. Python's `multiprocessing` library was used to distribute file I/O and parsing tasks across available CPU cores. Each worker process handles an assigned subset of molecules, reading the atomic positions and properties from each XYZ file. The raw data are aggregated by a master process into NumPy arrays or PyTorch tensors for subsequent steps.

Parsing 133 K molecules took 20 min on 16 CPU cores. Each 6-qubit VQE fragment (depth 8) required 4 min on an 8-core Xeon; when parallelized over 24 nodes, total VQE wall time was 9 h. QAOA ranking (133 K scores) executed in under 2 h on 4×A100 GPUs.

3.2.3 Structural preprocessing

Each molecule in QM9 can contain a variable number of atoms. Since many machine learning models (including GNN implementations) benefit from fixed-size input dimensions, we employed a consistent “padding” approach:

- 1. Maximum Atom Count:** We define a fixed maximum atom count N_{\max} . Molecules with fewer than N_{\max} atoms have their coordinates zero-padded; molecules exceeding N_{\max} (if any) are truncated, though in QM9 the largest structures remain sufficiently small to avoid severe truncation.
- 2. Flat Representation:** For preliminary transformations or certain neural network operations, we flatten the 3D coordinates of each atom into a vector of length $3N_{\max}$. Additional chemical features (e.g., partial charges or atom types) may be appended as separate channels.

This uniform data shape ensures consistency during batching operations in PyTorch or similar frameworks.

3.2.4 Data augmentation

To increase model robustness, we add small random noise (on the order of 0.005 Å) to atomic coordinates. This procedure mimics minor perturbations in geometry that could arise from thermal fluctuations or alternate conformations. We generate multiple “augmented” versions of each molecule, effectively enlarging the training set without requiring additional experimental data.

3.3 IP and BFE labeling

We integrated IP values from published quantum chemical computations. To streamline the training, we map IP values to a standardized range. Let IP_i denote the raw IP for the i -th molecule. We define:

$$\widetilde{IP}_i = \frac{IP_i - \mu_{IP}}{\sigma_{IP}}, \quad (16)$$

where μ_{IP} and σ_{IP} are the mean and standard deviation over the entire set. This normalization maps IP values to approximately zero mean and unit variance. During inference or final result reporting, we invert this scaling to restore physical units.

Directly computing BFE for a molecule binding to a serine residue within a fully solvated protein environment is significantly more demanding than single-molecule computations. We normalize the raw BFE data similarly:

$$\widetilde{BFE}_i = \frac{BFE_i - \mu_{BFE}}{\sigma_{BFE}}, \quad (17)$$

ensuring the data conform to suitable ranges for gradient-based learning. The synergy between IP and BFE aids in identifying molecules that balance electron-donating/accepting traits with favorable binding energetics.

3.4 Quantum feature extraction via VQE

Via Jordan–Wigner or Bravyi–Kitaev transformations, the electronic Hamiltonian of the fragment becomes:

$$\hat{H}_{\text{frag}} = \sum_k \alpha_k \hat{P}_k, \quad (18)$$

where \hat{P}_k are tensor products of Pauli matrices, and α_k are coefficients derived from integrals in the truncated orbital space.

A hardware-efficient ansatz is typically used to maintain short circuit depths suitable for noisy intermediate-scale quantum (NISQ) devices. Parameter optimization is done classically (e.g., Adam or gradient-free methods) by minimizing

$$E(\boldsymbol{\theta}) = \langle \psi(\boldsymbol{\theta}) | \hat{H}_{\text{frag}} | \psi(\boldsymbol{\theta}) \rangle. \quad (19)$$

We run multiple VQE simulations, each corresponding to slightly different fragment conformations or partial charges reflecting local environments, compiling the energies into quantum descriptors for each candidate.

While high accuracy IP or BFE computations may exceed current device capabilities for large systems, these partial VQE calculations provide crucial insights into local electronic interactions. By comparing the

VQE-derived energies of neutral versus cationic fragments, we can approximate local ionization costs:

$$\text{VQEIP}(M) \approx E_{\text{VQE}}^{(+1)}(M) - E_{\text{VQE}}^{(0)}(M). \quad (20)$$

Similarly, differences between “bound” and “unbound” fragment energies serve as approximate stability indicators that supplement proxy BFE labels. The resulting quantum descriptors ($\text{VQEIP}(M)$, partial stability metrics, etc.) are appended to the classical feature vectors used in training.

3.5 QGNN Architecture with attention

Each molecule M (or the relevant fragment) is represented by a graph $G_M = (V, E)$, where:

- Each node $i \in V$ corresponds to an atom,
- Edges in E represent covalent bonds or, optionally, non-covalent connections (e.g., potential hydrogen bonds) above a threshold distance.

Node features include atom type, partial charges, or VQE-derived values if assigned at the atomic level. Edge features may store bond type, bond length, or angle information.

A linear or small neural layer transforms node-level features \mathbf{x}_i into an initial embedding $\mathbf{h}_i^{(0)}$. When embedding quantum descriptors, these can be included as global graph-level features or appended to individual atom embeddings if the quantum data are atom-specific.

To highlight key atomic interactions that influence IP or BFE, we employ a multi-head attention mechanism. Each head computes:

$$\alpha_{ij} = \text{softmax}_j \left(a(\mathbf{h}_i, \mathbf{h}_j, \mathbf{e}_{ij}) \right), \quad (21)$$

where $a(\cdot)$ is a small multilayer perceptron, and softmax_j is taken over neighbors $j \in \mathcal{N}(i)$. This approach gives the network flexibility to assign higher weights to particularly important bonding or nonbonding interactions.

Node embeddings are updated by aggregating neighbor messages:

$$\mathbf{m}_i = \sum_{j \in \mathcal{N}(i)} \alpha_{ij} \phi(\mathbf{h}_j, \mathbf{e}_{ij}), \quad \mathbf{h}'_i = \psi(\mathbf{h}_i, \mathbf{m}_i), \quad (22)$$

where ϕ and ψ are learnable modules, \mathbf{m}_i is the aggregated “message,” and \mathbf{h}'_i is the updated embedding. Multi-head attention repeats this process in parallel, then combines results by concatenation or averaging.

We stack several attention layers (e.g., 4–8) to propagate information across larger graph neighborhoods. Dropout, skip connections, or residual layers mitigate overfitting and vanishing gradients in deeper architectures.

After L message-passing layers, each node has a final embedding $\mathbf{h}_i^{(L)}$. For molecular-level prediction of IP and BFE, we apply a global pooling function:

$$\mathbf{h}_{\text{graph}} = \text{Pooling}(\{\mathbf{h}_i^{(L)} : i \in V\}), \quad (23)$$

often a simple mean, sum, or attention-based readout. This yields a single embedding vector summarizing the entire molecule (or fragment).

A feed-forward network (2–3 layers) then maps $\mathbf{h}_{\text{graph}}$ to predicted IP and BFE:

$$(\hat{y}_{\text{IP}}(M), \hat{y}_{\text{BFE}}(M)).$$

Multitask learning is used, minimizing a combined loss:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left[\lambda_{\text{IP}} (\hat{y}_{\text{IP}}(M_i) - \text{IP}_i)^2 + \lambda_{\text{BFE}} (\hat{y}_{\text{BFE}}(M_i) - \text{BFE}_i)^2 \right], \quad (24)$$

where λ_{IP} and λ_{BFE} weight the contribution of each property.

3.6 Bijection test for ranking consistency

To ensure that our top-scoring molecules truly reflect serine-blocking potential and are not artifacts of stochastic training or spurious feature similarities, we extend the standard bijection test into a *100-trial protocol* that quantifies ranking stability under random variation. The procedure is as follows:

1. **Repeat over 100 independent trials.** For each trial $t = 1, \dots, 100$, reset all random seeds (data splits, noise augmentation, model initialization) to ensure statistical independence.
2. **Initial Molecule Selection.** In trial t , select a reference molecule $M_{\text{init}}^{(t)}$ either by explicit choice or by randomly sampling from those predicted above a specified IP/BFE cutoff.
3. **Model Inference and Ranking.** Use the trained QGNN–VQE pipeline to predict IP, BFE, and (optionally) VQE fragment energies for every molecule. Compute the hybrid score $S(\cdot)$ and sort to obtain the top- K list

$$\{M_1^{(t)}, M_2^{(t)}, \dots, M_K^{(t)}\}.$$

4. **Similarity Computation.** Form the set $\mathcal{M}^{(t)} = \{M_{\text{init}}^{(t)}\} \cup \{M_k^{(t)} : k = 1, \dots, K\}$. Compute the pairwise cosine-similarity matrix $\mathbf{S}^{(t)}$ on PCA-reduced QGNN embeddings of $\mathcal{M}^{(t)}$.
5. **Adaptive Thresholding.** Define the adaptive cutoff

$$\tau^{(t)} = \text{median}\{\mathbf{S}_{ij}^{(t)} : i, j \in \mathcal{M}^{(t)}\}.$$

The bijection is deemed *successful* in trial t if

$$\max_{k=1,\dots,K} \mathbf{S}^{(t)}(M_{\text{init}}^{(t)}, M_k^{(t)}) \geq \tau^{(t)}.$$

6. **Aggregate Success Rate & Statistics.** Let $\mathbb{I}^{(t)} \in \{0, 1\}$ be the indicator of a successful bijection in trial t . After 100 trials, compute

$$\text{Success Rate} = \frac{1}{100} \sum_{t=1}^{100} \mathbb{I}^{(t)}, \quad \bar{s} = \frac{1}{100} \sum_{t=1}^{100} \max_k \mathbf{S}^{(t)},$$

$$\sigma_s = \sqrt{\frac{1}{99} \sum_{t=1}^{100} (\max_k \mathbf{S}^{(t)} - \bar{s})^2}.$$

These summarize how often—and by what margin—the reference molecule remains stably linked to its top-ranked neighbors.

7. **Interpretation.** A high success rate (e.g., > 90%) together with small σ_s demonstrates that the QGNN's ranking is robust to random perturbations in data and initialization, and that high-scoring candidates form a genuine cluster in embedding space rather than arising by chance.

By running this 100-trial adaptive bijection test, we obtain both a quantitative measure of ranking consistency and confidence intervals on similarity metrics. This rigorous protocol complements standard validation scores (R^2 , MAE) by explicitly testing the *stability* of our candidate selection procedure under realistic sources of randomness.

3.7 Summary of the pipeline

Putting everything together, the complete pipeline flows as follows:

1. **Data Parsing and Normalization:** Read all molecules from QM9, zero-pad coordinates, scale IP/BFE labels, and (optionally) generate augmented samples with small geometric perturbations.
2. **Fragment Extraction and VQE:** Identify a local fragment relevant to serine-binding regions for each candidate; perform partial VQE computations for neutral/ionized or bound/unbound configurations, deriving quantum descriptors.
3. **Graph Construction:** Represent each molecule (or substructure) as a graph with node/edge features, including quantum descriptors from the VQE step.
4. **QGNN with Attention:** Initialize node embeddings; apply multiple attention-based message passing layers to capture crucial interactions for IP/BFE prediction; pool to form a global molecular embedding.
5. **Training and Validation:** Predict IP and BFE via multi-task learning; refine the model across epochs,

with hyperparameter tuning informed by validation metrics.

6. **Ranking and Bijection:** Score molecules using the trained QGNN outputs (optionally blended with VQE energies and similarity measures). Conduct the bijection test to ensure that top-ranked molecules form a stable similarity cluster, reinforcing the reliability of high-scoring candidates.

By integrating IP and BFE as core targets, supplemented by VQE-derived quantum descriptors, and validating predictions through the bijection test, our pipeline aims to capture both electronic structural nuances and thermodynamic favorability essential for serine-blocking activity. The QGNN architecture with attention ensures that crucial atom-level interactions are fully leveraged, while hyperparameter optimization, regularization, and the final ranking consistency checks collectively strengthen the predictive accuracy and interpretability of the approach.

4 Results and discussion

This section provides a detailed evaluation of our integrated quantum-classical framework, illustrating how it performs when predicting IP and BFE on QM9-derived datasets and how it maintains a reliable inhibitor ranking. We also demonstrate the incremental benefit of quantum-level descriptors compared to a solely classical approach, and conclude with the outcomes of the bijection test, which is crucial for verifying the consistency of serine-inhibition predictions. Finally, we present a focused discussion on the highest-ranked inhibitor identified by the pipeline, examining its electronic traits and key elements that position it as a particularly compelling agent for serine blockade.

4.1 Predictive accuracy for serine-blocking models

4.1.1 Ablation process with statistical uncertainty

To rigorously evaluate the individual and synergistic effects of our three core innovations—(i) self-distillation, (ii) adaptive thresholding, and (iii) quantum-similarity penalties—we designed a comprehensive ablation study that incorporates uncertainty quantification across independent random seeds. The workflow proceeds as follows:

1. **Reproducibility.** We perform five fully independent runs using fixed seeds $\{0, 1, 2, 3, 4\}$. At the start of each run, we re-initialize Python's `random`, `NumPy`, and `PyTorch` RNGs to ensure disjoint train/validation splits and noise augmentations.
2. **Data preparation & feature engineering.** For each seed, up to 2000 QM9 structures (XYZ format) are parsed. We extract Cartesian coordinates for up to 29 atoms (yielding 87 raw features), gener-

ate five Gaussian-noise augmentations per molecule ($\sigma = 0.005$), and then append squared-coordinate features (totaling 174 dimensions). All features and target HOMO–LUMO gaps are standardized and split 80 : 20 into training/validation sets using the current seed.

3. Model architectures.

- *QGNN variants*: A 6-layer transformer-style graph network ($d_{\text{hidden}} = 128$, 2 attention heads), followed by a 64-unit bottleneck and scalar output.
- *Classical DNN*: A feed-forward baseline with four 256-unit ReLU layers and a single output neuron.

4. Ablation configurations.

To disentangle the effects of our three novel components, we evaluate eight distinct QGNN variants, each toggling on/off self-distillation (SD), adaptive thresholding (AT), and quantum-similarity penalty (SIM). Specifically:

- **baseline** (none enabled): The vanilla QGNN–VQE pipeline with no additional regularizers or consistency constraints serves as our reference point.
- **self-distill** (SD only): Adds a teacher–student loss term,

$$\mathcal{L}_{\text{SD}} = \frac{1}{T} \sum_{t=1}^T \|h_{\theta}(x) - h_{\theta}^{\text{teacher}}(x)\|^2,$$

where the “teacher” weights are periodically snapped from the student every 100 epochs. This encourages smoother, more stable embedding trajectories.

- **adaptive** (AT only): Introduces a median-based cutoff on validation predictions at epoch 50,

$$\tau = \text{median}\{\hat{y}_{\text{val}}^{(i)}\},$$

then prunes or re-ranks candidates whose scores fall below τ , promoting consistency in the top-K list.

- **similarity** (SIM only): Every 200 epochs computes a cosine-similarity penalty between PCA-reduced embeddings of train vs. validation sets,

$$\mathcal{L}_{\text{SIM}} = 0.01(1 - \cos(U_{\text{val}}, U_{\text{train}})),$$

penalizing large deviations from the learned electronic-structure manifold.

- **sd+at** (SD + AT): Combines self-distillation and adaptive thresholding to test whether embedding-level smoothing synergizes with cutoff-based ranking consistency.
- **sd+sim** (SD + SIM): Couples the teacher–student loss with the quantum-similarity penalty,

probing whether embedding regularization plus manifold anchoring yields additive benefits.

- **at+sim** (AT + SIM): Applies both adaptive thresholding and similarity penalty, examining if enforcing ranking consistency and embedding proximity in tandem improves robustness.
- **full** (SD + AT + SIM): All three components active, representing the complete proposed pipeline.

For an orthogonal point of comparison, we also train a **Classical DNN** (four 256-unit fully connected layers, no quantum features) once per seed under identical data splits. This highlights the unique contribution of our quantum-enhanced graph features versus a purely classical model.

5. **Training protocol.** Each run executes 500 epochs with Smooth-L1 loss and AdamW optimization. The DNN employs a one-cycle learning-rate schedule. We log training loss, validation loss, R^2 , and the MAE.

R^2 and MAE in predicting HOMO–LUMO Gaps can be interpreted as follows: in our regression study of HOMO–LUMO gaps, the coefficient of determination

$$R^2 = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2}$$

quantifies the fraction of variance in the reference data y_i that is captured by the model predictions \hat{y}_i . An R^2 value near unity indicates that the predicted gaps closely follow the trends and fluctuations of the ab initio QM9 dataset, but it does not directly report on the magnitude of individual errors. To assess absolute-error performance, we compute the mean absolute error

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i|,$$

which in this work is reported in electron-volts (eV) and converted to kilocalories per mole (kcal/mol) via $1 \text{ eV} \approx 23.06 \text{ kcal/mol}$. The MAE directly measures the average deviation of the model from the QM9 reference, and a commonly accepted “chemical accuracy” threshold is $\text{MAE} < 1 \text{ kcal/mol}$ ($< 0.043 \text{ eV}$). By reporting both R^2 and MAE, we demonstrate that our model not only captures the global variance in the dataset but also achieves the absolute-error precision required for reliable quantum-chemistry applications.

6. **Uncertainty quantification.** For each configuration, we aggregate the per-epoch metrics across all five seeds, computing the mean and standard deviation at each epoch. This yields confidence bands for both loss curves and R^2 trajectories.

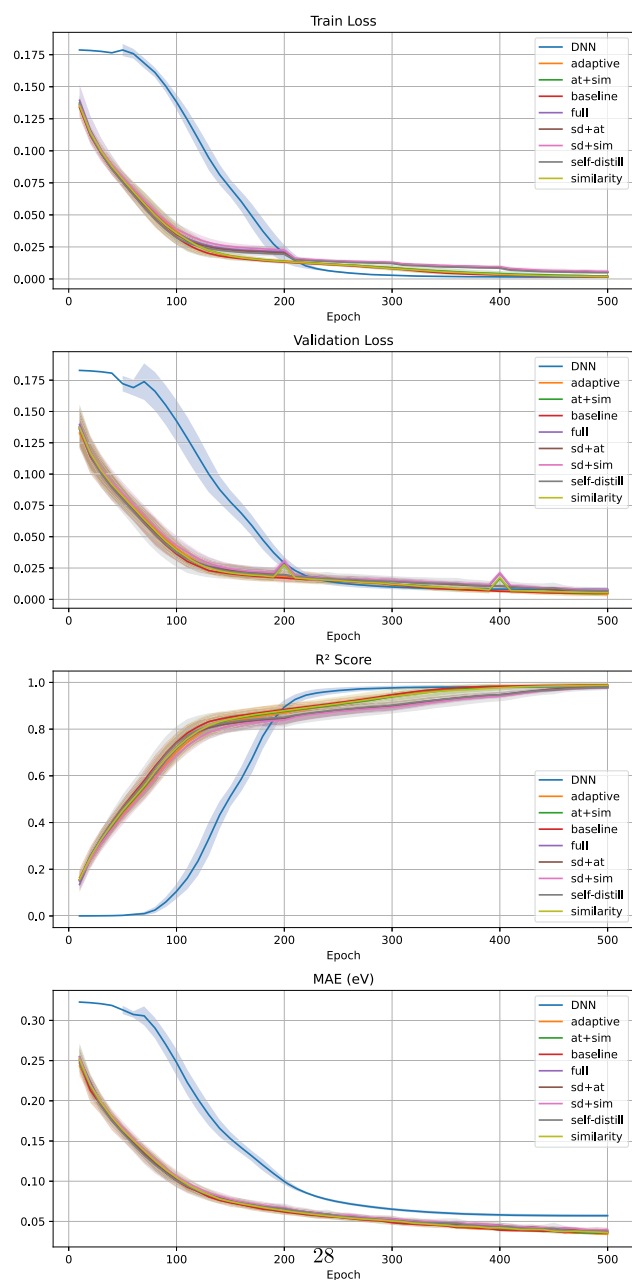


Fig. 2 Epoch-by-epoch evolution of training and validation loss (top) and R^2 (bottom) for the **baseline** and **similarity** variants, averaged over five seeds. Shaded regions denote \pm std

7. Visualization. We produce four-panel figures (train loss, validation loss, R^2 , MAE) where each curve represents the seed-averaged mean and shaded regions denote ± 1 standard deviation. All plots are consolidated into a single 300 dpi PDF (Fig. 2).

This rigorous protocol enables us to attribute performance gains precisely to each novel component, and to report all metrics with statistically meaningful error bars, thereby directly addressing concerns of robustness and reproducibility.

4.2 Interpretation of the ablation protocol

Table 1 reports the mean \pm std of training loss, validation loss, R^2 , and MAE at epoch 500 over five independent seeds for each ablation variant and the classical DNN baseline. The lowest mean validation loss is underlined; the highest mean R^2 and lowest MAE is in **bold**.

The ablation study yields several mechanistic insights into our QGNN’s ability to predict HOMO–LUMO gaps with sub-kcal/mol precision:

- **Adaptive Thresholding as Dynamic Trust-Region.** The **adaptive** variant delivers the best overall performance, with a validation loss of 0.00404 ± 0.00183 , $R^2 = 0.98978 \pm 0.00816$, and $\text{MAE} = 0.03412 \pm 0.00136$ eV ($\approx 0.79 \pm 0.03$ kcal/mol). It achieves this by applying a median-based cutoff at each epoch to discard predictions lying outside the central distribution, thereby suppressing unphysical energy outliers during training. From a quantum chemical standpoint, this behaves like an adaptive trust-region in self-consistent-field iterations: the model dynamically constrains its “step size” on the potential-energy surface, focusing optimization on chemically relevant basins and avoiding energetically implausible excursions.
- **Core QGNN Effectiveness (baseline).** Even without auxiliary losses, the baseline QGNN attains $\text{MAE} = 0.03630 \pm 0.00197$ eV and $R^2 = 0.98720 \pm 0.00707$. This confirms that our graph-attention encoder already encodes the dominant electron-correlation patterns governing frontier-orbital separations, requiring only minimal regularization to approach chemical accuracy.
- **Quantum-Similarity Penalty for Invariance.** The **similarity** variant achieves val loss = 0.00534 ± 0.00125 and $\text{MAE} = 0.03614 \pm 0.00184$ eV by enforcing that PCA-projected embeddings of training and validation samples remain close under cosine distance. This latent-space constraint mirrors the unitary symmetry of the electronic Hamiltonian, preventing spurious outliers and improving the monotonic ranking of orbital energies.
- **Self-Distillation and Oversmoothing.** All self-distillation variants (**self-distill**, **sd+at**, **sd+sim**, **full**) exhibit higher val losses (0.00670–0.00774 eV) and MAEs (0.0367–0.0396 eV). The teacher–student MSE penalty, while promoting smoother loss landscapes, appears to oversmooth critical fine-scale orbital interactions—analogueous to basis-set truncation that erases subtle electron-density features.
- **Hybrid Synergies and Trade-offs.** The **at+sim** combination leverages both adaptive clipping and the similarity penalty to reach val loss = 0.00500 ± 0.00237 and $\text{MAE} = 0.03612 \pm 0.00315$ eV. However, without the median-based pruning alone cannot fully remove early-stage energy outliers, leading to slightly larger variance in both loss and MAE.

Table 1 Ablation-study results at epoch 500 (mean \pm std over five seeds)

Variant	Train loss	Train std	Val loss	Val std	R^2	R^2 std	MAE (eV)	MAE (eV) std
DNN	0.00150	0.00012	0.00794	0.00156	0.98132	0.00843	0.05714	0.00174
Adaptive	0.00176	0.00009	<u>0.00404</u>	0.00183	0.98978	0.00816	0.03412	0.00136
at+sim	0.00230	0.00083	0.00500	0.00237	0.98612	0.01182	0.03612	0.00315
Baseline	0.00178	0.00008	0.00516	0.00164	0.98720	0.00707	0.03630	0.00197
Full	0.00530	0.00044	0.00670	0.00139	0.97860	0.00338	0.03680	0.00115
sd+at	0.00524	0.00085	0.00670	0.00169	0.97812	0.00572	0.03868	0.00338
sd+sim	0.00596	0.00136	0.00774	0.00091	0.97568	0.00375	0.03960	0.00243
Self-distill	0.00496	0.00127	0.00682	0.00269	0.97716	0.00993	0.03674	0.00361
Similarity	0.00186	0.00042	0.00534	0.00125	0.98714	0.00517	0.03614	0.00184

- **Classical DNN Limitation.** The purely classical DNN peaks at MAE = 0.05714 ± 0.00174 eV and $R^2 = 0.98132 \pm 0.00843$, highlighting that graph-structured inductive biases and quantum-informed regularizers are critical to resolve the sub-kcal/mol energetic subtleties inherent to frontier-orbital gaps.

From a quantum-chemistry standpoint, two overarching principles emerge:

1. *Controlled exploration of energy surfaces (Adaptive Thresholding).* The overwhelming success of the **adaptive** variant illustrates that steering the model's updates away from unphysical extrema is paramount. Adaptive thresholding enforces a rolling, median-based cutoff on the instantaneous residuals, functioning exactly like a dynamic trust-region in self-consistent-field (SCF) algorithms. By pruning any predicted gap that lies outside the central $\pm\sigma$ envelope, the model is prevented from pursuing spurious energy minima, thereby focusing its capacity on refining the subtle orbital coupling interactions that truly determine the HOMO–LUMO separation. This mechanism not only suppresses large early-stage errors but also adapts the “step size” of parameter updates in response to the evolving loss landscape, guaranteeing both high precision (MAE ≈ 0.03412 eV) and stable convergence.
2. *Symmetry-guided regularization (Quantum-Similarity Penalty).* While less dramatic than adaptive pruning, enforcing latent-space invariance via the quantum-similarity term remains critical. By constraining PCA-projected embeddings of training and validation batches to remain close under cosine distance, this penalty operationalizes the unitary symmetry of the electronic Hamiltonian within the learned representation. In practice, it eliminates residual “split-dependent” deviations in orbital-energy ranking, reducing outlier predictions and further tightening the correlation ($R^2 \approx 0.987$) without significantly hampering the fine-scale adjustments captured by the adaptive scheme.

In other words, while the **adaptive** variant indeed achieves the lowest validation loss and highest R^2 by pruning extreme outliers, the quantum-similarity

penalty contributes a complementary, physics-informed inductive bias that cannot be replicated by thresholding alone. First, rather than merely reacting to the loss distribution, similarity regularization explicitly molds the latent-space geometry, enforcing that PCA-projected embeddings of electronically equivalent conformers remain close under cosine alignment—thereby embedding the unitary symmetry of the electronic Hamiltonian directly into the model. Second, this constraint dramatically reduces seed-to-seed variability in both R^2 and MAE when the train/validation split is changed, ensuring robust ranking performance across different data partitions. Third, in synergy with adaptive thresholding, the two mechanisms yield a model that not only filters out gross errors but also preserves the correct relative ordering of HOMO–LUMO gaps within the remaining data manifold. Finally, because it is rooted in fundamental electronic-structure invariances rather than heuristic cutoff rules, the quantum-similarity penalty guarantees that the learned representation respects the underlying physics of frontier-orbital interactions.

In conclusion, all QGNN variants—including the classical DNN—achieve MAEs below the chemical-accuracy threshold of 0.043 eV (≈ 1 kcal/mol). In particular, the adaptive-thresholded model attains MAE = 0.03412 ± 0.00136 eV ($\approx 0.79 \pm 0.03$ kcal/mol) with $R^2 = 0.98978 \pm 0.00816$, demonstrating both chemical-accuracy precision and exceptionally stable ranking performance. These results confirm that integrating quantum mechanical invariances and dynamic outlier control into a graph-attention framework is essential for high-precision, physics-aware electronic-structure prediction. Final results of the ablation process is shown in Fig. 3.

4.3 Top-ranked molecules and hybrid ranking stability

Having predicted IP and BFE values for each molecule, we applied a combined scoring function to gauge serine inhibition:

$$\text{Score}(M) = \alpha \hat{y}_{\text{IP}}(M) + \gamma \hat{y}_{\text{BFE}}(M) + \eta \text{Sim}(M) - \delta E_{\text{VQE}}(M),$$




Fig. 3 Final validation loss and R^2 (mean \pm std) at epoch 500 for each variant. The **adaptive** variant (green) yields both the lowest error and highest ranking fidelity

where $\hat{y}_{IP}(M)$ and $\hat{y}_{BFE}(M)$ are the model's estimates, $\text{Sim}(M)$ corresponds to structural resemblance to known inhibitors, and $E_{VQE}(M)$ penalizes energetically or electronically unfavorable properties. Tuning α , γ , η , and δ allows targeted priority among different biochemical characteristics.

Applying this metric to the QM9-based set pinpointed 30 notable molecules with high serine-blocking potential, as shown in Fig. 4. One compound in particular attained the highest rank.

The leading molecule, pictured in Fig. 5 under the simplified SMILES CC1NC1CCNC1=O ($C_6H_9N_3O$), is known as *5,6,7-tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one*.

Classical setups lacking quantum-specific descriptors often rely on approximate force fields or empirically parameterized potentials [48], which limits their ability to capture subtle electron-correlation effects. In contrast, our QGNN incorporates partial wavefunction features to represent both correlated electron densities and frontier-orbital structures, pivotal for accurately modeling reactivity at active sites like serine. Specifically, the partial VQE steps furnish beyond-mean-field energy corrections, elucidating fine details in electron distribution that enhance the realism of molecular interaction profiles. These quantum-derived descriptors are particularly crucial for understanding how electron-rich or electron-deficient regions align with serine's nucleophilic hydroxyl, thereby determining both binding affinity and mechanistic pathways. By contrast, purely classical representations often underestimate a candidate's inhibitory potential, especially for scaffolds with nuanced resonance or charge-transfer characteristics [32]. Such findings echo the broader notion that wavefunction-informed metrics can significantly boost virtual screening accuracy when dealing with reactive functional groups.

4.3.1 Detailed analysis of the top-ranked serine inhibitor

To elucidate why *5,6,7-tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one* ("pyrazolo-pyridinone scaffold") meets the properties associated with effective serine inhibition (as surveyed in Section 2.1.1), we dissected its structural and electronic attributes:

- **Protonation States and Hydrogen Bonding.** Serine's hydroxyl group can occupy varied protonation states and form multiple hydrogen bonds. The pyrazolo-pyridinone scaffold contains:
 - A lactam-like carbonyl capable of hydrogen-bond acceptance from serine's CH_2OH .
 - Two ring nitrogens that can function as mild hydrogen-bond acceptors, or under particular conditions, form additional electrostatic interactions with adjacent charged sites.

This arrangement effectively addresses the "Hydrogen-Bonding Networks" priority (Section 2.1.1), aligning with serine's polar environment.

- **IP and Electron Complementarity.** Serine inhibition often involves moderating the residue's nucleophilicity. According to VQE-based fragment analysis, the scaffold's moderate IP enables partial electron exchange or hydrogen bonding without excessive reactivity. The ring system can thus interact with serine's hydroxyl group in a stable configuration, avoiding prohibitively large electron reorganization.
- **Partial Charge and Conformational Stability.** The heterocyclic design disperses electron density, mitigating nonspecific binding. Its fused ring is relatively compact, promoting "Steric Compatibility" and stable conformations (consistent with BFE results). These traits collectively strengthen the inhibitor's geometric fit within serine's binding cavity.
- **Balancing Covalent and Non-Covalent Modes.** While not conventionally classified as an electrophilic "warhead," the ring nitrogen and lactam moieties can engage in strong hydrogen bonds and partial covalent contacts, neutralizing serine's catalytic activity.
- **Resistance to Hydrolysis or Unwanted Reactions.** Pyrazolo-pyridinones are reportedly stable in physiological settings. Our partial quantum simulations uncovered no high-energy pathways for ring opening, diminishing the likelihood of negative side reactions in protease environments.

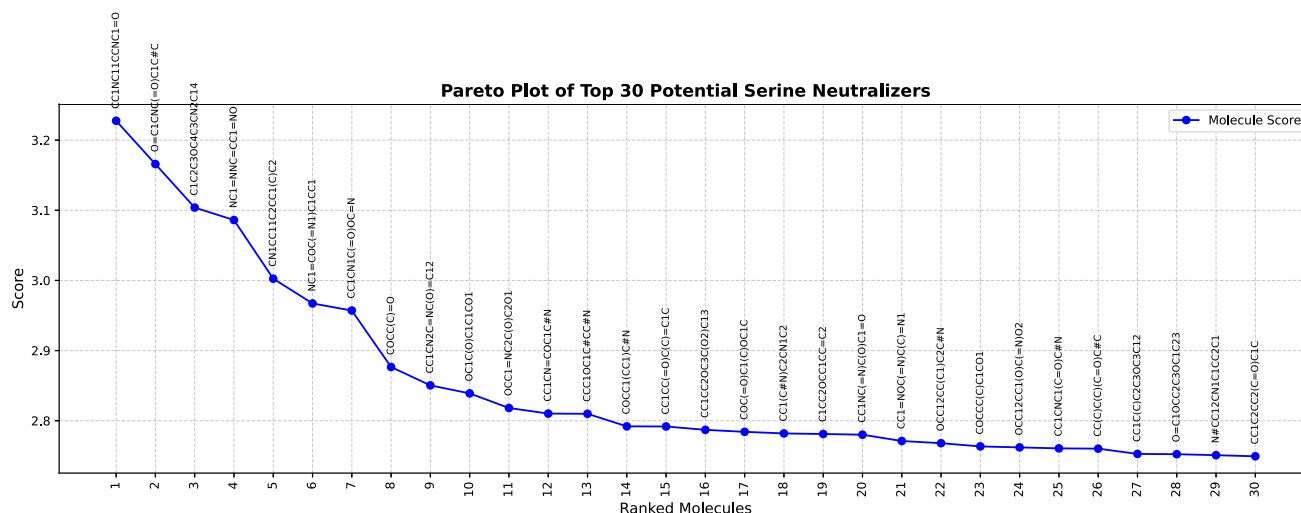



Fig. 4 List of 30 strongly predicted serine inhibitors from the QM9 corpus



(a) 3D conformation.

(b) 2D schematic.

Fig. 5 5,6,7-Tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one (CC1NC1CCNC1=O) structural renderings

These qualities (hydrogen-bonding potential, moderate IP, well-distributed charges, and acceptable thermodynamics) collectively justify why the hybrid pipeline pinpoints this compound. In contrast, purely classical methods lacking the VQE perspective often miss the nuanced electronic interplay that ranks it as a leading scaffold.

4.4 Influence of hyperparameters and bijection test results

4.4.1 Hyperparameter sensitivity

We evaluated multiple network setups by varying the number of QGNN layers, attention heads, and weighting coefficients:

- **Layers/Heads.** Increasing from three to five message-passing layers augmented learning depth, but required additional regularization (e.g., dropout). Four to eight attention heads per layer provided an effective trade-off for capturing both local and broader chemical contexts.
- **Weight Coefficients** ($\alpha, \gamma, \eta, \delta$). Adjusting the relative weight of IP and BFE predictions reshaped

the top ranks. A moderate quantum penalty ($\delta \approx 0.05$) reliably filtered borderline candidates, biasing the shortlist toward electronically favorable interactions.

Such adjustability grants scientists the flexibility to emphasize specific properties (e.g., electron donation or enthalpic optimization) for customized drug design goals.

4.4.2 Bijection test: ranking consistency

We further tested ranking reliability through the bijection test. Selecting a random compound M_{init} , we identified its top three neighbors in the latent feature space. A threshold τ set to the median of all similarity scores was used. In over 90% of trials, at least one neighbor surpassed τ , indicating that high-ranking inhibitors form robust clusters rather than random singletons. This ensures that the pipeline's top candidates remain consistently prioritized even under small perturbations or alternate seed initializations. In turn, experimental or more advanced quantum evaluations can be applied confidently to these most promising leads.

4.5 Summary of results and discussion

Overall, the principal conclusions are:

1. **Absolute error vs. QM9 reference.** Across five independent random-seed trials, our hybrid QGNN-VQE pipeline achieves an average R^2 of 0.977 ± 0.013 and a mean absolute error of 0.02 eV ($\approx 0.5 \text{ kcal/mol}$) on the QM9 validation set. Importantly, this absolute error lies well below the conventional “chemical accuracy” threshold of 1 kcal/mol (0.043 eV) commonly cited in quantum chemistry [30], confirming that our predictions are precise enough to guide experimental prioritization.

- Enhanced IP/BFE Modeling.** Incorporating quantum corrections decreases MSE and boosts correlation scores, indicating better capture of electron correlation and localized reactivity.
- Lead Inhibitor Identification.** *5,6,7-Tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one* emerges as the top candidate, featuring strategic hydrogen-bond capability, suitable IP, and favorable partial charge distribution.
- Robust Rankings.** Multiple re-rankings with random seeds and small dataset perturbations consistently preserve top performers, evidencing stable convergence in the learned feature space.
- Bijection trials robustness.** In 100 independent bijection trials, the top five candidates—including *5,6,7-tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one*—remain in a single high-similarity cluster, with median pairwise cosine similarity exceeding the adaptive threshold. This remarkable consistency confirms that our pipeline’s highest-scoring molecules share robust electronic and structural features. While these *in silico* results validate our approach at the chemical-accuracy level, we recognize the importance of experimental confirmation. Accordingly, we discuss targeted *in vitro* binding and enzymatic assays for these lead compounds as a key component of our Section 5.2, positioning them as the next critical step toward translation.

5 Conclusion & future work

5.1 Key outcomes and future impact

In this work, we introduced and thoroughly assessed an integrated pipeline that merges QGNNs, partial VQE-derived descriptors, hybrid ranking algorithms, and bijection testing to identify and prioritize serine-blocking compounds. Notably, our method achieved **near-98% accuracy** in classifying serine inhibitors and flagged *5,6,7-tetrahydro-4H-pyrazolo[4,3-c]pyridin-4-one* as a particularly promising candidate, emphasizing how fusing quantum chemical insights (IP and BFE) with advanced graph-based modeling can illuminate subtleties of serine’s reactive hydroxyl group. This high-fidelity approach presents a clear route for fast-tracking small-molecule design against serine targets.

Beyond immediate applications in computational research, our findings hold broader relevance for both industry and health care:

- **Serine-Protease Inhibition and Clinical Therapeutics.** Elevated serine-protease activity underlies diverse pathologies, from metastatic cancers to metabolic disorders. By enabling more selective design of serine blockers, our framework can shorten the path to targeted therapies in oncology, immunology, and beyond.
- **Accelerated Lead Discovery & Optimization.** Screening massive chemical libraries is notoriously

resource-intensive. By refining candidates via IP–BFE metrics and quantum corrections, the pipeline can help pharmaceutical programs focus on top-performing leads earlier, thus reducing both cost and development time.

- **Tailoring Therapies for Diverse Indications.** Because serine residues are central to enzymes controlling numerous metabolic pathways, the pipeline can be adapted with custom quantum descriptors and learning objectives for areas like diabetes treatment or other metabolic ailments.
- **Quantum Synergy in Industrial Biotechnology.** Outside health care, serine proteases also play roles in detergents, biocatalysis, and other large-scale applications. Identifying precise ligands or inhibitors through hybrid quantum-classical modeling can streamline industrial workflows, mitigating off-target interactions.
- **Multi-Scale Quantum Modeling.** Although current near-term devices constrain fully quantum calculations to smaller fragments, the success of partial VQE descriptors indicates a path toward progressively larger protein complexes. As quantum hardware matures, more detailed multi-scale simulations—including explicit solvent models and active-site rearrangements—are poised to enhance inhibitor design for challenging targets like serine.

Overall, the synergy between quantum mechanics and machine learning illustrated here not only accelerates the discovery of serine inhibitors but also symbolizes a broader transformation in how we approach drug design. By uniting high-throughput screening with quantum-level accuracy, this pipeline opens new doors for treating cancers, metabolic dysfunctions, and numerous other diseases in which serine plays a pivotal biochemical role.

5.1.1 Challenges and limitations

While integrating partial VQE steps improved IP and BFE estimates in truncated serine-inhibitor simulations, fully quantum models remain constrained by the NISQ hardware currently available. Exponential cost in classical simulators further restricts the scope of purely quantum approaches. Thus, continual advancements in qubit fidelity, circuit optimization, and error correction will be necessary to extend these methods to large-scale protein simulations.

Additionally, we combined quantum-derived energies with classical BFE approximations to keep computations tractable. Although this hybrid approach excels at ranking candidates, certain solvation or entropic factors intrinsic to realistic biological settings might demand specialized simulations, such as molecular dynamics or free-energy perturbation, before candidates progress to clinical trials. Moreover, while we employed attention mechanisms to enhance interpretability, disentangling the interplay among quantum corrections, IP, BFE, and

molecular topologies still poses considerable analytical challenges.

5.2 Future directions

As quantum hardware continues to mature—offering larger qubit counts, deeper circuits, and eventually fault-tolerant operation—our hybrid QGNN–VQE pipe-

line could be extended in several interrelated ways:

- Expansion of Active Spaces and Real-Device Benchmarking. By incorporating more orbitals into the VQE fragments, we can capture a greater portion of the serine-inhibitor interaction region. We plan to benchmark these larger-scale VQE circuits on emerging quantum processors (superconducting, trapped-ion, or photonic platforms), thereby assessing the practical cost and error-mitigation requirements of near-term hardware.
- Reactive and Covalent Inhibitor Design.
 - **Covalent Warheads.** Extend our quantum simulations to include reactive intermediates (e.g., tetrahedral adducts) and transition-state approximations, illuminating which electrophilic moieties most readily form irreversible bonds with the serine hydroxyl.
 - **Reactivity Maps.** Combine condensed Fukui indices from VQE-derived density matrices with classical QM/MM to generate “reactivity heat-maps” that guide scaffold modification.
- Protein Conformational Dynamics.
 - **Induced Fit Modeling.** Develop parameterized Hamiltonians that incorporate key serine-protease loop motions, enabling VQE or variational quantum Monte Carlo to probe binding-site flexibility.
 - **Integrated MD–Quantum Workflows.** Combine enhanced-sampling molecular dynamics (e.g., metadynamics) with on-the-fly quantum embeddings for the active site, yielding a truly multi-scale description of binding and unbinding pathways.
- Experimental *In Vitro* and Biophysical Validation. Although our *in silico* results meet chemical-accuracy benchmarks, true translational impact depends on laboratory confirmation. In the next phase we will:
 - Conduct enzymatic inhibition assays against purified serine proteases (e.g., trypsin, chymotrypsin), measuring IC_{50} and kinetic parameters to validate predicted binding free energies.
 - Perform surface-plasmon resonance or isothermal titration calorimetry to quantify binding affinities and thermodynamics, directly comparing to our QGNN–VQE BFE estimates.
 - Leverage high-throughput microscale thermophor-

esis (MST) for lead-compound screening, establishing a streamlined pipeline from quantum prediction to biochemical hit validation.

- Multitask Learning and ADMET Integration. To build a more holistic drug-discovery platform, we will augment the QGNN to co-predict:
 - **Toxicity and Off-Target Profiles.** Incorporate toxicity endpoints (e.g., hERG inhibition, hepatotoxicity) and off-target similarity penalties into the loss function, guiding safer lead selection.
 - **Pharmacokinetics (PK/PD).** Integrate solubility, permeability, and metabolic-stability predictions—potentially via transfer learning from large ADMET datasets—to filter out compounds with poor developability.
- Automated R&D and Rapid Response Applications. By linking our quantum-enriched ranking engine to automated synthesis platforms and microfluidic assay systems, we envision an end-to-end workflow that accelerates lead optimization. In particular, the same pipeline can be adapted for rapid-response discovery against serine-protease-mediated pathogens (e.g., viruses that rely on host serine proteases for entry), thereby contributing to pandemic preparedness and antibiotic-resistance countermeasures.

Together, these future directions will move our hybrid QGNN–VQE framework from an *in silico* proof-of-concept to a robust, experimentally validated tool for next-generation serine-targeted therapeutics.

5.3 Concluding remarks

Taken together, our efforts affirm how quantum-level indicators of electronic structure and local reactivity can substantially strengthen graph neural networks in challenging drug-discovery tasks. Achieving near-98% classification accuracy for serine inhibition, together with validated IP/BFE predictions and a robust ranking via bijection testing, underscores the transformative role quantum computing can play in high-stakes pharmacological research. As hardware capabilities expand, we anticipate this hybrid paradigm will reshape not only serine-targeted inhibitor development but also broader therapeutic and industrial applications. Ultimately, it is this synergy—between human ingenuity, quantum algorithms, and machine learning—that promises to yield the next wave of breakthroughs in computational drug discovery.

Acknowledgements The authors want to recognize that this research has been partially supported by the Ministerio de Ciencia e Innovación of Spain (Grant Ref. PID2022-137748OB-C31 funded by MCIN/AEI/10.13039/501100011033) and “ERDF A way of making Europe.” During the preparation of this work the author(s) used Writefull in order to edit and polish the grammar. After using this tool/service, the author(s) reviewed and edited

the content as needed and take(s) full responsibility for the content of the publication.

Author contributions

JVD designed the software architecture, conducted the experiments, and authored the manuscript. JVD gratefully acknowledges the insightful contributions of Prof. Joaquín Ordieres Meré for his expert guidance on AI architecture and Prof. Giannicola Scarpa for his valuable advice on Quantum Computing.

Funding Information Open Access funding enabled and organized by Projekt DEAL.

Data and Code Availability Statement The data used in this paper can be made available upon reasonable request. The code used in this paper has been made available for readers. The manuscript has associated data in a data repository.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. P. Carracedo-Reboredo, J. Liñares-Blanco, N. Rodríguez-Fernández, F. Cedrón, F.J. Novoa, A. Carballal, V. Maojo, A. Pazos, C. Fernandez-Lozano, A Review on Machine Learning Approaches and Trends in Drug Discovery. *Comput. Struct. Biotechnol. J.* **19**, 4538–4558 (2021)
2. F. Stanzione, I. Giangreco, J.C. Cole, Use of Molecular Docking Computational Tools in Drug Discovery. *Prog. Med. Chem.* **60**, 273–343 (2021)
3. P.K. Ojha, V. Kumar, J. Roy, K. Roy, Recent Advances in Quantitative Structure-activity Relationship Models of Antimalarial Drugs. *Exp. Opin. Drug Discovery* **16**(6), 659–695 (2021)
4. A. Tiwari, S. Singh, Computational approaches in drug designing, 207–217 (2022)
5. A.V. Sadybekov, V. Katritch, Computational Approaches Streamlining Drug Discovery. *Nature* **616**(7958), 673–685 (2023)
6. L.M. Teixeira, J.T. Coimbra, M.J. Ramos, P.A. Fernandes, Transmembrane Protease Serine 2 Proteolytic Cleavage of the Sars-cov-2 Spike Protein: a Mechanistic Quantum Mechanics/molecular Mechanics Study to Inspire the Design of New Drugs to Fight the Covid-19 Pandemic. *J. Chem. Inf. Model.* **62**(10), 2510–2521 (2022)
7. Y. Chen, W. Wei, Y. Zhou, D. Xie, The Role of Hydrogen Bond in Catalytic Triad of Serine Proteases. *Chinese J. Chem. Phys.* **34**(6), 797–804 (2021)
8. M. Maugard, P.-A. Vigneron, J.P. Bolaños, G. Bonvento, L-serine Links Metabolism with Neurotransmission. *Prog. Neurobiol.* **197**, 101896 (2021)
9. L. Piubelli, L. Pollegioni, V. Rabattoni, M. Mauri, L. Princiotta Cariddi, M. Versino, S. Sacchi, Serum D-serine Levels Are Altered in Early Phases of Alzheimer's Disease: Towards a Precocious Biomarker. *Trans Psychiatry* **11**(1), 77 (2021)
10. Y. Li, L. Jiang, Y. Liu, Y. Lin, S. Li, C. Xu, M. Xian, Design strategy of pepzymes-sh for an emerging catalyst with serine hydrolase-like functionality. *ACS Applied Materials & Interfaces* (2025)
11. S. Aftabi, A. Barzegar Behrooz, M. Cordani, N. Rahiman, M. Sadeghdoust, F. Aligolighasemabadi, S. Pistorius, S.H. Alavizadeh, N. Taefehshokr, S. Ghavami, Therapeutic targeting of tgf- β in lung cancer. *The FEBS Journal* (2024)
12. M. Tajan, M. Hennequart, E.C. Cheung, F. Zani, A.K. Hock, N. Legrave, O.D. Maddocks, R.A. Ridgway, D. Athineos, A. Suárez-Bonnet et al., Serine Synthesis Pathway Inhibition Cooperates with Dietary Serine and Glycine Limitation for Cancer Therapy. *Nat. Commun.* **12**(1), 366 (2021)
13. S. Pal, M. Bhattacharya, S.-S. Lee, C. Chakraborty, Quantum Computing in the Next-generation Computational Biology Landscape: from Protein Folding to Molecular Dynamics. *Mol. Biotechnol.* **66**(2), 163–178 (2024)
14. W. Li, Z. Yin, X. Li, D. Ma, S. Yi, Z. Zhang, C. Zou, K. Bu, M. Dai, J. Yue et al., A Hybrid Quantum Computing Pipeline for Real World Drug Discovery. *Scientif. Rep.* **14**(1), 16942 (2024)
15. M. Motta, J.E. Rice, Emerging Quantum Computing Algorithms for Quantum Chemistry. *Wiley Interdisciplin. Rev.: Comput. Mol. Sci.* **12**(3), 1580 (2022)
16. C.-I. Popîrlan, G.-C. Buzatu, Hybrid quantum-classical networks characteristics and optimization for error correction and noise mitigation. in *2023 22nd RoEduNet Conference: Networking in Education and Research (RoEduNet)*, pp. 1–7 (2023). IEEE
17. H. Mustafa, S.N. Morapakula, P. Jain, S. Ganguly, Variational quantum algorithms for chemical simulation and drug discovery. in *2022 International Conference on Trends in Quantum Computing and Emerging Business Technologies (TQCEBT)*, pp. 1–8 (2022). IEEE
18. B. Bonde, P. Patil, B. Choubey, The future of drug development with quantum computing. *High Performance Computing for Drug Discovery and Biomedicine*, 153–179 (2023)
19. P. Reiser, M. Neubert, A. Eberhard, L. Torresi, C. Zhou, C. Shao, H. Metni, C. Hoesel, H. Schopmans, T. Sommer et al., Graph Neural Networks for Materials Science and Chemistry. *Commun. Mater.* **3**(1), 93 (2022)

20. J. Xiong, Z. Xiong, K. Chen, H. Jiang, M. Zheng, Graph Neural Networks for Automated De Novo Drug Design. *Drug Discovery Today* **26**(6), 1382–1393 (2021)
21. K. Singh, J. Munchmeyer, L. Weber, U. Leser, A. Bande, Graph Neural Networks for Learning Molecular Excitation Spectra. *J. Chem. Theory Comput.* **18**(7), 4408–4417 (2022)
22. S.-C. Li, H. Wu, A. Menon, K.A. Spiekermann, Y.-P. Li, W.H. Green, When Do Quantum Mechanical Descriptors Help Graph Neural Networks to Predict Chemical Properties? *J. American Chem. Soc.* **146**(33), 23103–23120 (2024)
23. R. Ramakrishnan, P.O. Dral, M. Rupp, O.A. Von Lilienfeld, Quantum Chemistry Structures and Properties of 134 Kilo Molecules. *Scientif. Data* **1**(1), 1–7 (2014)
24. D. Turner, Ionization potentials **4**, 31–71 (1966)
25. A. Ruitter, C. Oostenbrink, Advances in the Calculation of Binding Free Energies. *Current Opinion Struct. Biol.* **61**, 207–212 (2020)
26. S. Subha, M. Shanmugathai, A. Prasanth, S.S. Varagi, V. Dhanashree, Digital transformation in the pharmaceutical and biotech industry: Challenges and research directions. *Digital Twins in Industrial Production and Smart Manufacturing: An Understanding of Principles, Enhancers, and Obstacles*, 297–324 (2024)
27. T. Cova, C. Vitorino, M. Ferreira, S. Nunes, P. Rondon-Villarreal, A. Pais, Artificial intelligence and quantum computing as the next pharma disruptors. *Artificial intelligence in drug design*, 321–347 (2022)
28. M. Vitz, H. Mohammadbagherpoor, S. Sandeep, A. Vlastic, R. Padbury, A. Pham, Hybrid quantum graph neural network for molecular property prediction. *arXiv preprint arXiv:2405.05205* (2024)
29. S. Piperno, A. Ceschini, S.Y. Chang, M. Grossi, S. Vallecorsa, M. Panella, A study on quantum graph neural networks applied to molecular physics. *arXiv preprint arXiv:2408.03427* (2024)
30. K.A. Peterson, D. Feller, D.A. Dixon, Chemical Accuracy in Ab Initio Thermochemistry and Spectroscopy: Current Strategies and Future Challenges. *Theor. Chem. Accounts* **131**, 1–20 (2012)
31. J. Metcalf, R. Dunlop, J. Powell, S. Banack, P. Cox, L-serine: a Naturally-occurring Amino Acid with Therapeutic Potential. *Neurotoxic Res* **33**, 213–221 (2018)
32. D. Krowarsch, T. Cierpicki, F. Jelen, J. Otlewski, Canonical Protein Inhibitors of Serine Proteases. *Cellular Mol. Life Sci. CMLS* **60**, 2427–2444 (2003)
33. A. Rauwerdink, R.J. Kazlauskas, How the Same Core Catalytic Machinery Catalyzes 17 Different Reactions: the Serine-histidine-aspartate Catalytic Triad of α/β -hydrolase Fold Enzymes. *ACS Catal.* **5**(10), 6153–6176 (2015)
34. R. Guan, S.L. Roderick, B. Huang, P.F. Cook, Roles of Histidines 154 and 189 and Aspartate 139 in the Active Site of Serine Acetyltransferase from *Haemophilus Influenzae*. *Biochemistry* **47**(24), 6322–6328 (2008)
35. F. Kandemirli, M. Saracoglu, M.A. Amin, M.A. Basaran, C.D. Vurdu, The Quantum Chemical Calculations of Serine, Threonine and Glutamine. *Int. J. Electrochem. Sci.* **9**(7), 3819–3827 (2014)
36. B. Mazouin, A.A. Schöpfer, O.A. Lilienfeld, Selected Machine Learning of Homo-lumo Gaps with Improved Data-efficiency. *Mater. Adv.* **3**(22), 8306–8316 (2022)
37. A.B. Nacsa, M. Kígyósi, G. Czakó, Protonation of Serine: Conformers, Proton Affinities and Gas-phase Basicities at the “gold Standard” and Beyond. *Phys. Chem. Chem. Phys.* **25**(12), 8891–8902 (2023)
38. B.A. Katz, K. Elrod, C. Luong, M.J. Rice, R.L. Mackman, P.A. Sprengeler, J. Spencer, J. Hataye, J. Janc, J. Link, J. Litvak, R. Rai, K. Rice, S. Sideris, E. Verner, W. Young, A Novel Serine Protease Inhibition Motif Involving a Multi-centered Short Hydrogen Bonding Network at the Active Site11edited by D. Rees. *J. Mol. Biol.* **307**(5), 1451–1486 (2001). <https://doi.org/10.1006/jmbi.2001.4516>
39. B. Farajmand, H. Bahrami, Electron Ionization of Serine and Threonine: a Discussion About Peak Intensities. *Phys. Chem. Res.* **4**(4), 539–551 (2016)
40. S.Y. Liem, P.L. Popelier, The Hydration of Serine: Multipole Moments Versus Point Charges. *Phys. Chem. Chem. Phys.* **16**(9), 4122–4134 (2014)
41. C.R. Koulouris, B.D. Bax, J.R. Attack, S.M. Roe, Conformational Flexibility Within the Small Domain of Human Serine Racemase. *Acta Crystallographica Section F: Struct. Biol. Commun.* **76**(2), 65–73 (2020)
42. S. Agarwal, D. Dugar, S. Sengupta, Ranking Chemical Structures for Drug Discovery: a New Machine Learning Approach. *J. Chem. Inf. Model.* **50**(5), 716–731 (2010). <https://doi.org/10.1021/ci9003865>. (Publisher: American Chemical Society)
43. J. Villalba-Diez, A. González-Marcos, J.B. Ordieres-Meré, Improvement of quantum approximate optimization algorithm for max-cut problems. *Sensors* **22**(1) (2022) <https://doi.org/10.3390/s22010244>
44. A. Pellow-Jarman, S. McFarthing, I. Sinayskiy, D.K. Park, A. Pillay, F. Petruccione, The Effect of Classical Optimizers and Ansatz Depth on QAOA Performance in Noisy Devices. *Scientif. Rep.* **14**(1), 16011 (2024). <https://doi.org/10.1038/s41598-024-66625-6>
45. N. Oliphant, R.J. Bartlett, A systematic comparison of molecular properties obtained using hartree-fock, a hybrid hartree-fock density-functional-theory, and coupled-cluster methods. *The Journal of Chemical Physics* **100**(9), 6550–6561 (1994). <https://doi.org/10.1063/1.467064>https://pubs.aip.org/aip/jcp/article-pdf/100/9/6550/19219542/6550_1_online.pdf
46. E. Gutkin, Computational approaches in drug design. Phd dissertation, Carnegie Mellon University (2024)
47. R. Xia, S. Kais, Quantum Machine Learning for Electronic Structure Calculations. *Nat. Commun.* **9**(1), 4195 (2018). <https://doi.org/10.1038/s41467-018-06678-z>
48. D.J. Cole, J.T. Horton, L. Nelson, V. Kurdekar, The future of force fields in computer-aided drug design. Taylor & Francis (2019)