
Bridging Quantum and Classical Computing in Drug Design: Architecture Principles for Improved Molecule Generation

Andrew Smith¹ Erhan Guven¹

Abstract

Hybrid quantum-classical machine learning offers a path to leverage noisy intermediate-scale quantum (NISQ) devices for drug discovery, but optimal model architectures remain unclear. We systematically optimize the quantum-classical bridge architecture of generative adversarial networks (GANs) for molecule discovery using multi-objective Bayesian optimization. Our optimized model (BO-QGAN) significantly improves performance, achieving a 2.27-fold higher Drug Candidate Score (DCS) than prior quantum-hybrid benchmarks and 2.21-fold higher than the classical baseline, while reducing parameter count by more than 60%. Key findings favor layering multiple (3-4) shallow (4-8 qubit) quantum circuits sequentially, while classical architecture shows less sensitivity above a minimum capacity. This work provides the first empirically-grounded architectural guidelines for hybrid models, enabling more effective integration of current quantum computers into pharmaceutical research pipelines.

1. Background and Significance

Developing new pharmaceuticals is a notoriously long and expensive process, often exceeding a decade and costing billions of dollars (Hughes et al., 2011; Singh et al., 2023). Computational methods offer a path to accelerate this process, but identifying promising candidates within the vast chemical space (estimated at 10^{60} molecules) remains a significant hurdle (Bassani & Moro, 2023).

Machine learning techniques, particularly graph neural networks like MolGAN (De Cao & Kipf, 2022), help navigate chemical space by learning from known molecules to generate novel candidates. However, challenges persist in

effectively representing complex structures and generalizing from limited approved drug data (fewer than 3,000 small-molecule drugs have received FDA approval (Knox et al., 2024)), often leading to high bias or mode collapse.

Quantum computing offers a potential advantage by naturally modeling quantum mechanical effects (superposition, entanglement) crucial to molecular behavior (Cao et al., 2019; McArdle et al., 2020; Feynman, 1982). Quantum neural networks (QNNs), operating on quantum states, could theoretically learn molecular physics more efficiently (Cao et al., 2019). However, current noisy intermediate-scale quantum (NISQ) hardware suffers from limited qubit counts, substantial noise, and restricted access, hindering practical application (Combarro & González-Castillo, 2023; Preskill, 2018).

1.1. Hybrid Quantum-Classical Models: A Pragmatic Approach

Hybrid quantum-classical models offer a pragmatic solution, combining smaller quantum circuits (capturing key quantum features) with larger classical networks (handling scale and decoding). Parameterized quantum circuits, which are analogous to classical layers, contain tunable parameters and have adjustable width (qubits) and depth (layers) (Benedetti et al., 2019). A quantum computation result is bridged into the classical network, leveraging quantum expressivity within current hardware limits. We use Generative Adversarial Networks (GANs) as our framework due to established classical (MolGAN) and hybrid (Li et al., 2021) benchmarks, allowing focused analysis of the hybrid architecture itself.

The original hybrid model for this task, QGAN-HG (Li et al., 2021), adapted MolGAN (De Cao & Kipf, 2022) and matched its performance with significantly fewer parameters. Subsequent work explored variations like Wasserstein loss and gradient penalty integration (Jain & Ganguly, 2022), different quantum component placements (Kao et al., 2023), and advanced GAN structures like CycleGAN (Anoshin et al., 2024). The potential real-world applicability was highlighted by work generating and experimentally validating novel KRAS inhibitor candidates discovered using quantum-enhanced methods (Ghazi Vakili et al., 2025). While valuable, these studies often focus on framework

¹Whiting School of Engineering, Johns Hopkins University, Baltimore, Maryland. Correspondence to: Andrew Smith <asmit397@jh.edu>.

Proceedings of the Workshop on Generative AI for Biology at the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

modifications, overlooking architectural details. To isolate the impact of the quantum-classical interface architecture, we benchmark directly against the original MolGAN and QGAN-HG, anticipating our derived principles may benefit more advanced frameworks later.

Designing the optimal quantum-classical architecture involves complex trade-offs between quantum expressivity, NISQ hardware limitations, and overall model parameter count. Establishing systematic principles to balance these factors for effective information transfer and performance is crucial but currently lacks rigorous investigation.

This research systematically optimizes the quantum-classical hybrid architecture, hypothesizing that a principled approach can significantly enhance performance and efficiency. We aim to:

1. Systematically evaluate the impact of quantum and classical component sizes on model performance using multi-objective Bayesian optimization.
2. Derive design principles for quantum circuits within hybrid generators by analyzing optimized configurations.

By establishing the first empirically grounded architectural guidelines, this work demonstrates substantial performance gains (2.3x *DCS* increase over previous hybrid work) with reduced computational cost (>60% fewer parameters than the classical baseline), offering a practical path to integrate NISQ computers more effectively into drug discovery.

2. Materials and Methods

2.1. Molecular Representations as Graphs

Molecules are represented as graphs $\mathcal{G}(\mathcal{V}, \mathcal{E})$ with up to $N = 9$ heavy atoms (C, O, N, F) plus hydrogens from the QM9 dataset. Nodes $v_i \in \mathcal{V}$ represent atoms with feature vectors $x_i \in \mathbb{R}^{T=6}$ (one-hot encoding C, O, N, F, H, plus padding). Edges $(v_i, v_j) \in \mathcal{E}$ represent bonds with types $y \in \{\text{none, single, double, triple, aromatic}\}$ ($Y = 5$). The graph is encoded by a feature matrix $X \in \mathbb{R}^{N \times T}$ and an adjacency tensor $A \in \mathbb{R}^{N \times N \times Y}$ indicating bond types (De Cao & Kipf, 2022). Molecules with fewer than N atoms are padded.

2.2. Hybrid Generative Adversarial Network (GAN) Architecture

Our model adapts MolGAN (De Cao & Kipf, 2022), comprising a hybrid generator G_θ , a classical discriminator (critic), and a classical reward network (Figure 1).

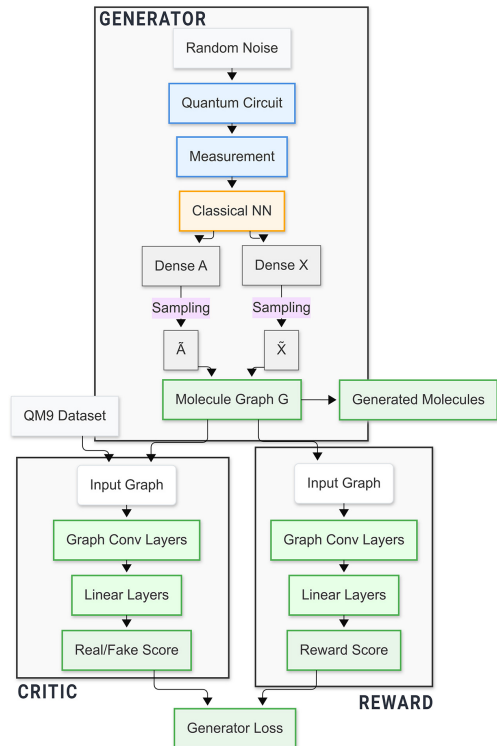


Figure 1. Architecture of the hybrid quantum-classical generator GAN with reward network.

2.2.1. GENERATOR NETWORK (G_θ)

The generator maps noise $z \in \mathbb{R}^{B \times M}$ (batch size B , quantum width M) to a molecular graph \mathcal{G} . The input z is embedded into a quantum state using angle encoding. This state is transformed by parameterized quantum circuit layers consisting of single-qubit rotations (RY gates) and CNOT-based entangling layers (with a ring structure). Measurement of Pauli-Z expectation values yields a quantum latent representation ($\mathbb{R}^{B \times M}$). This representation is transferred to the classical component, where fully connected layers (with tanh activations) output dense tensors A and X . Gumbel-Softmax sampling produces discrete graph tensors \tilde{A} and \tilde{X} . Key optimized hyperparameters are quantum width (M), quantum depth (number of circuit layers), classical width (neurons/layer), and classical depth (number of layers).

2.2.2. DISCRIMINATOR (CRITIC) NETWORK

The discriminator is a classical graph neural network, which distinguishes real vs. generated molecules (\tilde{A}, \tilde{X}), outputting a scalar Wasserstein critic score r . It uses graph convolutions and aggregation layers.

2.2.3. REWARD NETWORK

Architecturally similar to the discriminator but with independent parameters, the reward network predicts the expected reward (desirable chemical properties, Section 2.3) for a generated molecule \mathcal{G} , guiding the generator via differentiable gradients (De Cao & Kipf, 2022).

2.2.4. TRAINING DETAILS

The generator loss is a weighted sum of the adversarial loss from the discriminator and the value loss from the reward network, controlled by a hyperparameter $\lambda \in [0, 1]$. The Generator/Reward and Discriminator networks use separate Adam optimizers (LR 1×10^{-4}). Gradient clipping (norm 1.0) and a warm-up/constant/decay LR scheduler are used. Quantum circuits are implemented via PennyLane (Bergholm et al., 2022) with PyTorch (Paszke et al., 2019).

2.3. Evaluation Metrics

Performance is assessed by realism and druglikeness, calculated using RDKit (Landrum et al., 2024). Chemically invalid molecules score zero.

- **Realism (Fréchet Distance, FD):** Measures distributional similarity between generated and real molecule batches using discrete Fréchet distance (Thomas Eiter & Heikki Mannila, 1994; Denaxas & Pikoula, 2019). $FD < 12.5$ indicates sufficient realism (Li et al., 2021).
- **Druglikeness (QED , $\log P$, SA):**
 - QED : Quantitative Estimation of Druglikeness, similarity to known oral drugs based on physicochemical properties (Bickerton et al., 2012).
 - $\log P$: Octanol-water partition coefficient (lipophilicity) (Wildman & Crippen, 1999; Lipinski et al., 2001).
 - SA : Synthetic Accessibility score, estimating synthesis ease based on fragment contributions and complexity penalties (Ertl & Schuffenhauer, 2009).

- **Drug Candidate Score (DCS):** A composite metric designed to balance the above properties, calculated as:

$$DCS = 10 \times QED \times \log P \times SA$$

where 10 is a scaling factor for readability. Individual metrics (QED , $\log P$, SA) are normalized to $[0, 1]$ (higher is better; note SA normalization reverses standard interpretation).

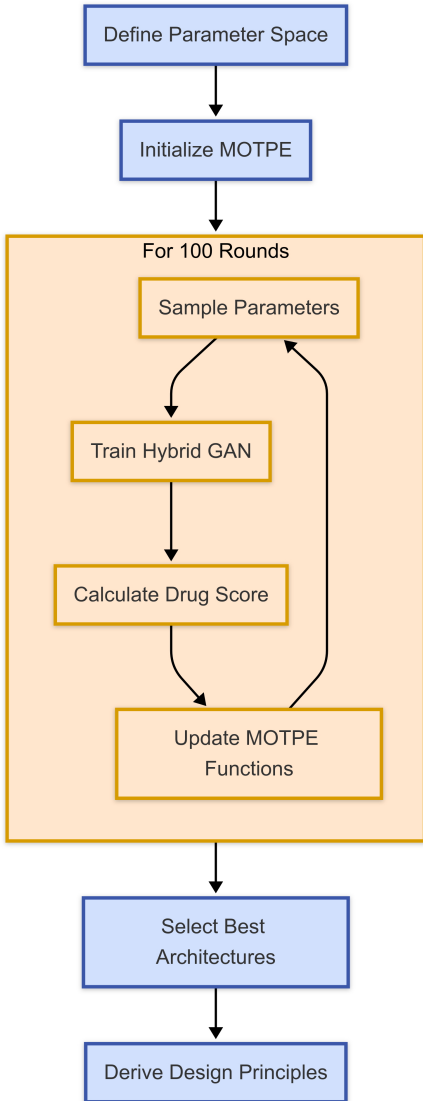


Figure 2. Workflow for optimizing the architecture with Multi-Objective Tree-structured Parzen Estimator (MOTPE).

2.4. Bayesian Optimization of Bridge Architecture

The optimization process used the Multi-Objective Tree-structured Parzen Estimator (MOTPE) algorithm, implemented via the Optuna library (Akiba et al., 2019). MOTPE is well suited for computationally expensive black-box optimization problems such as hyperparameter tuning (Ozaki et al., 2020). It builds probabilistic surrogate models ($l(x)$ for promising parameters, $g(x)$ for unpromising ones) based on Pareto dominance relationships from evaluated trials. It then suggests new hyperparameter configurations by maximizing the expected improvement criterion, efficiently exploring the search space with relatively few trials. Each optimization trial involved sampling an architecture, training the corresponding hybrid GAN model, and evaluating

its performance (DCS and training time), which is shown in Figure 2.

The search space covered quantum width ([4, 16] qubits), quantum depth ([1, 4] layers), classical width ([16, 512] neurons), and classical depth ([1, 4] layers).

2.5. Dataset and Preprocessing

We used the QM9 dataset (133,885 small organic molecules, up to 9 heavy atoms C, O, N, F) (Ramakrishnan et al., 2014), preprocessed into graph representations (Section 2.1) including explicit hydrogens. Padding was used for molecules < 9 heavy atoms. RDKit handled molecule processing (Landrum et al., 2024).

2.6. Experimental Design

Phase 1 involved 100 MOTPE trials that optimized the architecture. Each trial trained a hybrid GAN, using a simulated quantum backend for up to 5000 epochs with early stopping (no validation DCS improvement for 250 epochs after reaching $FD < 12.5$). Phase 2 benchmarked the best DCS architecture discovered during Phase 1 against re-implemented MolGAN (De Cao & Kipf, 2022) and QGAN-HG MR (Li et al., 2021) using identical protocols. Performance metrics were averaged over 30 independent runs (1000 molecules/run). Statistical significance was assessed using t-tests, and effect size was assessed using Cohen’s d. Experiments used Python 3.12, PyTorch 2.4.1, PyTorch Geometric, PennyLane, Optuna, RDKit (Paszke et al., 2019; Fey & Lenssen, 2019; Bergholm et al., 2022; Akiba et al., 2019; Landrum et al., 2024).

The results of a hybrid quantum-classical model discovered by the optimization process were then validated on real quantum hardware.

3. Results

3.1. Bayesian Optimization Performance

The MOTPE algorithm efficiently navigated the hyperparameter space in the 100 trials, progressively identifying improved architectures DCS . The algorithm effectively explored the trade-off between DCS and training time, identifying a Pareto front of optimal configurations. The final model, which we call Bayesian-Optimized Quantum GAN (BO-QGAN), was selected from this front to maximize DCS .

3.2. Quantitative Performance Comparison

Systematic optimization yielded substantial performance gains compared to baseline models. Table 1 summarizes the key performance metrics for the proposed Bayesian

Optimized QGAN (BO-QGAN) alongside both reported and recreated results for classical MolGAN (De Cao & Kipf, 2022) and the baseline hybrid model, QGAN-HG MR (Li et al., 2021).

Averaged over 30 independent trials generating 1000 molecules each, our optimized architecture (BO-QGAN) achieved a mean DCS of 1.190 ± 0.132 . This substantially outperforms the recreated comparison models under identical evaluation conditions. Using independent two-sample t-tests, we found significant ($p < 0.001$) improvements over both the recreated QGAN-HG MR (0.524 ± 0.083 DCS) and the recreated MolGAN (0.539 ± 0.068 DCS). This represents a 2.27-fold improvement over the baseline hybrid model and a 2.21-fold improvement over the classical baseline, with large effect sizes (Cohen’s $d > 0.8$) indicating a meaningful advance in generated candidate quality.

Beyond the composite DCS , the improvements in individual properties like lipophilicity ($\log P$) and synthetic accessibility (SA) are particularly noteworthy (Table 1). The higher normalized $\log P$ suggests molecules with properties more amenable to absorption and distribution according to Lipinski’s guidelines (Lipinski et al., 2001), while the significantly improved SA (0.33 vs. 0.21/0.23) indicates that the molecules generated by BO-QGAN are predicted to be considerably easier to synthesize in a lab, a critical factor for practical drug development (Ertl & Schuffenhauer, 2009).

This enhanced performance was achieved with high efficiency: BO-QGAN uses only 158,252 total trainable parameters (158,231 classical and 21 quantum), achieving a greater than 60% reduction compared to the classical MolGAN’s 400,203 parameters.

3.3. Characteristics of Optimized Architecture

The optimization revealed complex, non-linear relationships between model capacity and performance. The best architecture, Bayesian-Optimized Quantum GAN (BO-QGAN), uses 3 quantum circuit layers (7 qubits each) feeding into 2 classical layers (227 neurons each) (Table 2). Its superior DCS is mainly due to improved lipophilicity ($\log P$) and synthetic accessibility (SA) compared to baselines (Table 1). A sample of molecules generated by a model and their corresponding QED scores is shown in Figure 3.

3.4. Influence of Quantum Architecture

The quantum component design was critical. Performance peaked within a middle range of total quantum parameters, suggesting a trade-off between expressivity and computational cost/overfitting. Analysis of width vs. depth (Figure 4) revealed that the best models consistently used multiple (3-4) shallow (4-8 qubits) quantum circuits layered sequentially, corresponding to width-to-depth ratios between ≈ 1.0

Table 1. Drug Properties from Inference of 1000 Molecules with various GAN Architectures.

Method	<i>QED</i>	<i>logP</i>	<i>SA</i>	<i>DCS</i>
MolGAN (reported) (De Cao & Kipf, 2022)	0.51	0.66	0.08	0.27
MolGAN (recreated)	0.39	0.60	0.23	0.54
QGAN-HG MR (reported) (Li et al., 2021)	0.51	0.49	0.11	0.27
QGAN-HG MR (recreated)	0.44	0.57	0.21	0.52
BO-QGAN (proposed)	0.44	0.83	0.33	1.19

The best score in each category is in bold. Model BO-QGAN is the optimized model from this study. The recreated and proposed models record the average over 30 trials.




Figure 3. A sample of 12 molecules generated during the inference stage and drawn with RDKit. The value shown is the *QED* score of the molecule.

Table 2. Model architecture with best performance.

Qubits	Q Circuits	Classical Width	Classical Layers
7	3	227	2

and 2.7. This indicates that layering more smaller circuits is more effective than using fewer, wider circuits for this task.

3.5. Influence of Classical Architecture

In contrast, the classical architecture showed less sensitivity beyond a minimum capacity threshold (approx. >90 neurons per layer). Increasing classical parameters or varying the width-to-depth ratio showed little correlation with improved *DCS* once this threshold was met (Figure 5). This suggests that the classical network primarily needs sufficient capacity to process the quantum output, with diminishing returns from further size increases.

3.6. Interplay between Quantum and Classical Architecture

Analysis revealed no strong interplay between optimal quantum and classical network sizes. Figure 6 visualizes the *DCS* achieved across the search space (defined by quantum and classical parameter counts). The landscape appears noisy, without clear ridges which would indicate that certain quantum sizes work best only with specific classical sizes. The correlation between the ratio of quantum-to-classical parameters and *DCS* was weak ($r = -0.20$). Optimal performance seems to depend on finding suitable configurations within each component independently, rather than a specific balance between them.


3.7. Computational Efficiency

The optimization process successfully balanced performance and computational cost. Figure 7 plots the final *DCS* achieved against the total training time for each of the 100 trials. No significant correlation was found between longer training times and better final *DCS* ($r = -0.01$). The best model, BO-QGAN, trained efficiently (<17 minutes on test hardware, using 2140 training iterations). Training time with the simulated backend was primarily driven by quantum circuit width (qubits), consistent with the exponential scaling of simulation cost with N . The finding that shallower, narrower circuits perform best aligns well with computational tractability.


3.8. Validation on Real Quantum Hardware

To bridge the gap between idealized simulations and practical application on current hardware, we conducted a preliminary validation using a physical quantum processing unit (QPU). Acknowledging the significant constraints associated with NISQ device access time, queue waits, and resource limitations for academic research, this validation necessarily used a small sample size.


For this test, we selected the model configuration identified during optimization that generated the highest percentage of

Performance vs Quantum Width/Depth Ratio ($r = -0.02$)



(a) Comparing the width-to-depth ratio of the quantum circuits to the *DCS* observed.

 Performance vs Quantum Width/Depth Ratio ($r = 0.05$)



(a) Comparing the width-to-depth ratio of the classical neural network to the *DCS* observed.

 Performance vs Number of Circuits ($r = -0.01$)



(b) Comparing the number of quantum circuits (layers) to the performance of the model.

 Performance vs Number of Circuits ($r = -0.00$)


(b) Comparing the number of classical neural network layers to the performance of the model.

 Performance vs Qubit Count ($r = 0.01$)


(c) Comparing the number of qubits used (width) to the *DCS*. The red line and r score show the linear correlation between variables. The green region contains all the configurations within the top 10% of *DCS*.

 Performance vs Qubit Count ($r = 0.10$)


(c) Comparing the number of nodes per network layer to the *DCS*. The red line and r score show the linear correlation between variables. The green region contains all the configurations with the top 10% of *DCS*.

Figure 4. Analysis of quantum circuit architectures.

Figure 5. Analysis of classical neural network architectures.




Figure 6. A performance landscape showing how *DCS* varies with the size of quantum circuits and classical neural networks. Red dots indicate individual trials, and the surface between them is interpolated to visualize trends.

chemically valid molecules, rather than the absolute highest *DCS* model (BO-QGAN). This choice aimed to maximize the number of usable molecules generated per hardware run, making the most efficient use of limited quantum resources by minimizing invalid outputs which receive zero scores across all metrics.

We generated a batch of 10 molecular structures using this high-validity configuration executed on IBM’s 127-qubit QPU ‘Sherbrooke’ via the cloud platform. For direct comparison, a batch of 10 molecules was generated using the same model configuration executed on a classical quantum simulator (PennyLane’s ‘default.qubit’). The resulting molecules from both backends were evaluated using the standard druglikeness metrics (*QED*, *logP*, *SA*, *DCS*).

Table 3 presents the comparison. While the metrics differ slightly between the QPU and simulator runs, the values generally fall within comparable ranges. The small sample size ($n = 10$) and the inherent stochasticity of both GAN generation and NISQ hardware execution prevent definitive conclusions about performance parity or the precise impact of hardware noise (e.g., decoherence, gate errors) on generation quality. However, these initial results are encouraging, tentatively suggesting that the architectural principles derived from simulation may translate reasonably well to physical quantum hardware.




Figure 7. Comparison between training time for a model and *DCS* achieved.

Table 3. Performance Comparison on Real vs. Simulated Quantum Backends ($n = 10$).

Backend	<i>QED</i>	<i>logP</i>	<i>SA</i>	<i>DCS</i>
IBM QPU (‘Sherbrooke’)	0.50	0.76	0.24	0.89
Quantum Simulator	0.43	0.73	0.37	1.17

4. Discussion

4.1. Interpretation

This study demonstrates that systematic, multi-objective Bayesian optimization of the quantum-classical interface significantly enhances hybrid generative model performance and efficiency. Our optimized model, BO-QGAN, substantially outperformed hybrid and classical baselines in generating druglike molecules (*DCS*) while using over 60% fewer parameters than classical MolGAN. This validates MOTPE for tuning complex hybrid systems and yields empirically-grounded design principles. Key findings include the relative independence of optimal quantum and classical component sizes, the critical role of quantum architecture (favoring multiple shallow circuits), and the lower sensitivity of classical architecture beyond a capacity threshold.

4.2. Design Principles for Quantum Circuits

Based on the optimization results, particularly the analysis of the influence of quantum architecture (Section 3.4), we can propose practical design guidelines for hybrid generative models in this domain. For the QM9 dataset, characterized by small molecules, the best-performing models consistently favored multiple (3-4), sequentially layered, shallow (4-8 qubits) quantum circuits over fewer wider or deeper ones. This suggests a core principle: when designing the quantum component, prioritize layering relatively shallow circuits sequentially.

The potential reasons for the effectiveness of this layered, shallow approach are multifaceted. It may relate to improved trainability on NISQ devices; shallower circuits are

generally less susceptible to noise accumulation and may mitigate the risk of encountering barren plateaus in the optimization landscape (McClean et al., 2018). Furthermore, layering circuits could allow the model to learn hierarchical representations, with earlier layers capturing basic features and later layers combining them into more complex patterns, analogous to deep classical networks but within the quantum domain. This approach also aligns well with computational efficiency, especially for simulation, where cost scales exponentially with qubit count (width) but only polynomially with gate depth (layers). Therefore, when allocating limited quantum resources, prioritizing adding circuit layers over adding qubits appears to be a more cost-effective strategy for improving performance on this type of task.

Regarding the classical component, the principle appears simpler: ensure sufficient capacity (roughly >90 neurons per layer in our experiments) to process the information from the quantum circuits, after which further increases yield diminishing returns (Section 3.5). These empirically derived principles offer valuable, quantitative starting points for designing effective and resource-aware quantum-classical hybrid models in the NISQ era.

4.3. Limitations and Mitigation

Key limitations include:

- **Simulation vs. Real Hardware:** Primary results rely on idealized simulations lacking NISQ noise and topology constraints. While preliminary hardware validation (Section 3.8) was encouraging, extensive testing on physical QPUs is needed to confirm real-world viability.
- **QM9 Dataset Scope:** QM9 is a standard benchmark in this field, but it contains small, simple molecules. Generalizability to larger, more diverse drug candidates (e.g., from ZINC or ChEMBL) requires further investigation.
- **MolGAN Baseline Framework:** Performance gains are relative to the MolGAN framework. A simple model was chosen to help isolate the effect of the architecture optimization. Applying the optimization methodology and derived principles to state-of-the-art generative models (e.g. diffusion) is a logical next step.
- **DCS Metric and Optimization Scope:** The DCS metric is one way to quantify quality; results might differ with other objectives. The optimization experiment focused on network width and depth; exploring different gates, embeddings, or measurements could yield further gains.

4.4. Future Research Directions

Future work should prioritize extensive validation on diverse NISQ hardware platforms with larger sample sizes and error mitigation. Applying the optimization methodology to larger datasets (ZINC, ChEMBL) and more advanced generative frameworks (diffusion, flow-matching) is crucial for assessing scalability and generalizability. Further investigation into optimizing finer architectural details (gates, embedding, measurement) and adapting the framework for target-specific objectives (binding affinity) are also promising directions.

4.5. Conclusion

We presented the first systematic multi-objective optimization of the quantum-classical interface for hybrid molecular generative models. Our approach identified architectures yielding significant performance improvements with reduced parameter counts. The resulting BO-QGAN model, validated optimization strategy, and derived design principles (preferring layered, shallow quantum circuits) provide a practical foundation for rationally designing more effective hybrid models. This work helps bridge the gap between quantum potential and drug discovery challenges, potentially accelerating the integration of NISQ computing into pharmaceutical research.

5. Data and Code Availability

Data and code are available on [Github](https://github.com/amerorchis/Hybrid-Quantum-GAN) (<https://github.com/amerorchis/Hybrid-Quantum-GAN>).

References

- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. Optuna: A Next-generation Hyperparameter Optimization Framework, July 2019. URL <http://arxiv.org/abs/1907.10902>. Issue: arXiv:1907.10902 [cs].
- Anoshin, M., Sagingalieva, A., Mansell, C., Zhiganov, D., Shete, V., Pflitsch, M., and Melnikov, A. Hybrid Quantum Cycle Generative Adversarial Network for Small Molecule Generation. *IEEE Transactions on Quantum Engineering*, 5:1–14, 2024. ISSN 2689-1808. doi: 10.1109/TQE.2024.3414264. URL <https://ieeexplore.ieee.org/document/10556803/?arnumber=10556803>. Conference Name: IEEE Transactions on Quantum Engineering.
- Bassani, D. and Moro, S. Past, Present, and Future Perspectives on Computer-Aided Drug Design Methodologies. *Molecules*, 28(9):3906, May 2023. doi: 10.3390/molecules28093906. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC10000000/>.

- nlm.nih.gov/articles/PMC10180087/. Number: 9.
- Benedetti, M., Lloyd, E., Sack, S., and Fiorentini, M. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, November 2019. ISSN 2058-9565. doi: 10.1088/2058-9565/ab4eb5. URL <https://iopscience.iop.org/article/10.1088/2058-9565/ab4eb5>. Number: 4.
- Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., and Ahmed, S. PennyLane: Automatic differentiation of hybrid quantum-classical computations, July 2022. URL <http://arxiv.org/abs/1811.04968>. Issue: arXiv:1811.04968 arXiv:1811.04968 [quant-ph].
- Bickerton, G. R., Paolini, G. V., Besnard, J., Muresan, S., and Hopkins, A. L. Quantifying the chemical beauty of drugs. *Nature Chemistry*, 4(2):90–98, February 2012. ISSN 1755-4349. doi: 10.1038/nchem.1243. URL <https://www.nature.com/articles/nchem.1243>. Number: 2 Publisher: Nature Publishing Group.
- Cao, Y., Romero, J., Olson, J. P., Degroote, M., Johnson, P. D., Kieferová, M., Kivlichan, I. D., Menke, T., Peropadre, B., Sawaya, N. P. D., Sim, S., Veis, L., and Aspuru-Guzik, A. Quantum Chemistry in the Age of Quantum Computing. *Chemical Reviews*, 119(19):10856–10915, October 2019. ISSN 0009-2665. doi: 10.1021/acs.chemrev.8b00803. URL <https://doi.org/10.1021/acs.chemrev.8b00803>. Number: 19 Publisher: American Chemical Society.
- Combarro, E. F. and González-Castillo, S. *A Practical Guide to Quantum Machine Learning and Quantum Optimization: Hands-on Approach to Modern Quantum Algorithms*. Packt Publishing Limited, Birmingham, 2023. ISBN 978-1-80461-383-2.
- De Cao, N. and Kipf, T. MolGAN: An implicit generative model for small molecular graphs, September 2022. URL <http://arxiv.org/abs/1805.11973>. Issue: arXiv:1805.11973 arXiv:1805.11973 [cs, stat].
- Denaxas, S. and Pikoula, M. spiro/discrete.frechet: meerkat stable release, August 2019. URL <https://zenodo.org/records/3366385>.
- Ertl, P. and Schuffenhauer, A. Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. *Journal of Cheminformatics*, 1(1):8, June 2009. ISSN 1758-2946. doi: 10.1186/1758-2946-1-8. URL <https://doi.org/10.1186/1758-2946-1-8>. Number: 1.
- Fey, M. and Lenssen, J. E. Fast Graph Representation Learning with PyTorch Geometric, April 2019. URL <http://arxiv.org/abs/1903.02428>. Issue: arXiv:1903.02428 arXiv:1903.02428 [cs].
- Feynman, R. P. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, June 1982. ISSN 1572-9575. doi: 10.1007/BF02650179. URL <https://doi.org/10.1007/BF02650179>. Number: 6.
- Ghazi Vakili, M., Gorgulla, C., Snider, J., Nigam, A., Bezrukov, D., Varoli, D., Aliper, A., Polykovsky, D., Padmanabha Das, K. M., Cox III, H., Lyakisheva, A., Hosseini Mansob, A., Yao, Z., Bitar, L., Tahoulas, D., Čerina, D., Radchenko, E., Ding, X., Liu, J., Meng, F., Ren, F., Cao, Y., Stagljar, I., Aspuru-Guzik, A., and Zhavoronkov, A. Quantum-computing-enhanced algorithm unveils potential KRAS inhibitors. *Nature Biotechnology*, pp. 1–6, January 2025. ISSN 1546-1696. doi: 10.1038/s41587-024-02526-3. URL <https://www.nature.com/articles/s41587-024-02526-3>. Publisher: Nature Publishing Group.
- Hughes, J., Rees, S., Kalindjian, S., and Philpott, K. Principles of early drug discovery. *British Journal of Pharmacology*, 162(6):1239–1249, 2011. ISSN 1476-5381. doi: 10.1111/j.1476-5381.2010.01127.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1476-5381.2010.01127.x>. Number: 6 eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1476-5381.2010.01127.x>.
- Jain, P. and Ganguly, S. Hybrid Quantum Generative Adversarial Networks for Molecular Simulation and Drug Discovery, December 2022. URL <http://arxiv.org/abs/2212.07826>. Issue: arXiv:2212.07826 arXiv:2212.07826 [quant-ph, q-bio].
- Kao, P.-Y., Yang, Y.-C., Chiang, W.-Y., Hsiao, J.-Y., Cao, Y., Aliper, A., Ren, F., Aspuru-Guzik, A., Zhavoronkov, A., Hsieh, M.-H., and Lin, Y.-C. Exploring the Advantages of Quantum Generative Adversarial Networks in Generative Chemistry. *Journal of Chemical Information and Modeling*, 63(11):3307–3318, June 2023. ISSN 1549-9596. doi: 10.1021/acs.jcim.3c00562. URL <https://doi.org/10.1021/acs.jcim.3c00562>. Number: 11 Publisher: American Chemical Society.
- Knox, C., Wilson, M., Klinger, C. M., Franklin, M., Oler, E., Wilson, A., Pon, A., Cox, J., Chin, N. E. L., Strawbridge, S. A., Garcia-Patino, M., Kruger, R., Sivakumaran, A., Sanford, S., Doshi, R., Khetarpal, N., Fatokun, O., Doucet, D., Zubkowski, A., Rayat, D. Y., Jackson, H., Harford, K., Anjum, A., Zakir, M., Wang, F.,

- Tian, S., Lee, B., Liigand, J., Peters, H., Wang, R. Q. R., Nguyen, T., So, D., Sharp, M., da Silva, R., Gabriel, C., Scantlebury, J., Jasinski, M., Ackerman, D., Jewison, T., Sajed, T., Gautam, V., and Wishart, D. S. DrugBank 6.0: the DrugBank Knowledgebase for 2024. *Nucleic Acids Research*, 52(D1):D1265–D1275, January 2024. ISSN 1362-4962. doi: 10.1093/nar/gkad976. Number: D1.
- Landrum, G., Tosco, P., Kelley, B., Rodriguez, R., Cosgrove, D., Vianello, R., sriniker, gedec, Jones, G., NadineSchneider, Kawashima, E., Nealschneider, D., Dalke, A., Swain, M., Cole, B., Turk, S., Savelev, A., Vaucher, A., Wójcikowski, M., Take, I., Scalfani, V. F., Probst, D., Ujihara, K., Walker, R., Pahl, A., godin, g., tadhurst-cdd, Lehtivarjo, J., Bérenger, F., and Bisson, J. rdkit/rdkit: 2024_03_5 (Q1 2024) Release, July 2024. URL <https://zenodo.org/records/12782092>. Programmers: ˆn3676.
- Li, J., Topaloglu, R. O., and Ghosh, S. Quantum Generative Models for Small Molecule Drug Discovery. *IEEE Transactions on Quantum Engineering*, 2: 1–8, 2021. ISSN 2689-1808. doi: 10.1109/TQE.2021.3104804. URL <https://ieeexplore.ieee.org/document/9520764/>.
- Lipinski, C. A., Lombardo, F., Dominy, B. W., and Feeney, P. J. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings1. *Advanced Drug Delivery Reviews*, 46(1):3–26, March 2001. ISSN 0169-409X. doi: 10.1016/S0169-409X(00)00129-0. URL <https://www.sciencedirect.com/science/article/pii/S0169409X00001290>. Number: 1.
- McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C., and Yuan, X. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, March 2020. doi: 10.1103/RevModPhys.92.015003. URL <https://link.aps.org/doi/10.1103/RevModPhys.92.015003>. Number: 1 Publisher: American Physical Society.
- McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., and Neven, H. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1): 4812, November 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-07090-4. URL <https://www.nature.com/articles/s41467-018-07090-4>. Publisher: Nature Publishing Group.
- Ozaki, Y., Tanigaki, Y., Watanabe, S., and Onishi, M. Multiobjective tree-structured parzen estimator for computationally expensive optimization problems. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, pp. 533–541, Cancún Mexico, June 2020. ACM. ISBN 978-1-4503-7128-5. doi: 10.1145/3377930.3389817. URL <https://dl.acm.org/doi/10.1145/3377930.3389817>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html>.
- Preskill, J. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. doi: 10.22331/q-2018-08-06-79. URL <https://quantum-journal.org/papers/q-2018-08-06-79/>. Publisher: Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften.
- Ramakrishnan, R., Dral, P. O., Rupp, M., and von Lilienfeld, O. A. Quantum chemistry structures and properties of 134 kilo molecules. *Scientific Data*, 1(1): 140022, August 2014. ISSN 2052-4463. doi: 10.1038/sdata.2014.22. URL <https://www.nature.com/articles/sdata201422>. Number: 1 Publisher: Nature Publishing Group.
- Singh, N., Vayer, P., Tanwar, S., Poyet, J.-L., Tsaioun, K., and Villoutreix, B. O. Drug discovery and development: introduction to the general public and patient groups. *Frontiers in Drug Discovery*, 3:1201419, May 2023. ISSN 2674-0338. doi: 10.3389/fddsv.2023.1201419. URL <https://www.frontiersin.org/articles/10.3389/fddsv.2023.1201419/full>.
- Thomas Eiter and Heikki Mannila. Computing discrete Fréchet distance. 1994.
- Wildman, S. A. and Crippen, G. M. Prediction of Physicochemical Parameters by Atomic Contributions. *Journal of Chemical Information and Computer Sciences*, 39 (5):868–873, September 1999. ISSN 0095-2338. doi: 10.1021/ci990307l. URL <https://doi.org/10.1021/ci990307l>. Number: 5 Publisher: American Chemical Society.