

SymQNet: VQE Hamiltonian Estimation for Molecular Optimization with Quantum Computing

Yash V. Tomar
West Lafayette, Indiana

"Significantly Faster Drug Candidate Screening"

Drug Discovery is Broken

\$2.6B Average Drug Cost¹
12yr Time to Market²
90% Failure Rate³

Delayed treatment has been proven to **increase mortality rates**⁴.

Faster **molecular optimization (MO)** can significantly speed up drug candidate screening.

MO's Flawed Role

Molecular Optimization identifies most **stable configuration** to accurately evaluate a candidate^{5,6,7}.

- Stable structures are most commonly found configuration of candidate in nature⁵.
- Important because it **accurately** determines how a potential candidate **behaves** in biological systems^{5,6}

Born-Oppenheimer approximation implies most stable structure lies at lowest ground-state energy⁵.

Models such as DFT attempt to **minimize ground-state energy** to arrive at the most stable structure⁵.

MO methods have limits in accuracy^{5,6}

Struggle with weak *intermolecular forces* (e.g., London-Dispersion)⁷

Solution: Quantum MO (QMO)

Uses **Variational Quantum Eigensolver (VQE)** to approximate wavefunction better.^{8,9}

VQE is a Potential Solution

VQE reduces synthesis by 30-50%, saving **6-12 months**¹⁰

- Thousand-year** long problems *potentially* solved in **hours**¹⁰

Hamiltonian is operator that encodes molecule's energy¹¹

Calculating Hamiltonian Energy is **impossible**

Mathematically impossible to diagonalize & directly calculate the energy of a structure¹¹

Estimating Energy is "Next Best Bet"

VQE^{8,9} uses various methods to try estimate the energy on Hamiltonian

Choosing measurements becomes important

Optimally choosing what parameter to measure could speed up MO

Estimation Requires **Slow** Measurements

To be accurate, methods must continuously measure parameters (potentially **millions of times**) to reduce variation^{11,12}

Hypothesis based on RL



Choose Measurement
Usually chosen poorly, wasting shots later

Conduct Measurement
...on chosen qubit, expensive but essential.

Pain Point

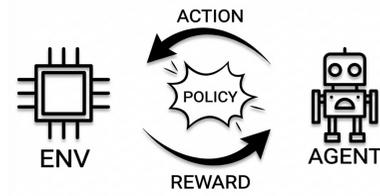
Hypothesis: If canonical **information** in each measurement shot can be **maximized** using reinforcement learning, then the number of estimation shots for a Hamiltonian's energy can be **significantly reduced** because maximizing information reduces uncertainty in the energy estimate more efficiently, allowing convergence with **fewer measurements**.

Key Challenge: Engineer a way around measurements' noisy and probabilistic nature. Such conditions are **unfavorable** for learning.

Methods: Problem Framed as POMDP and Architecture Iteratively Engineered

Reinforcement Learning Offers Reward Based Strategy

Reinforcement learning trains an "agent" to take **optimal actions** based on observations of its environment guided by a reward signal^{14,15}



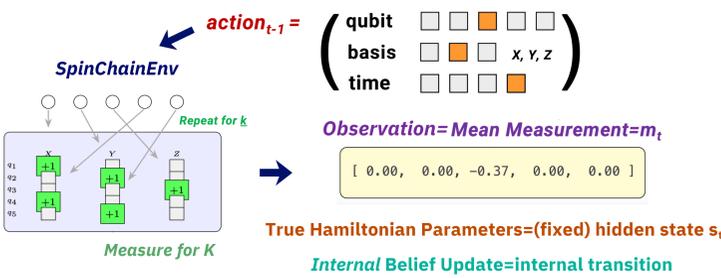
Infogain Reward

$$r_t = \max(0, H_{t-1} - H_t)$$

where H is *entropy* of belief and t represents the time index

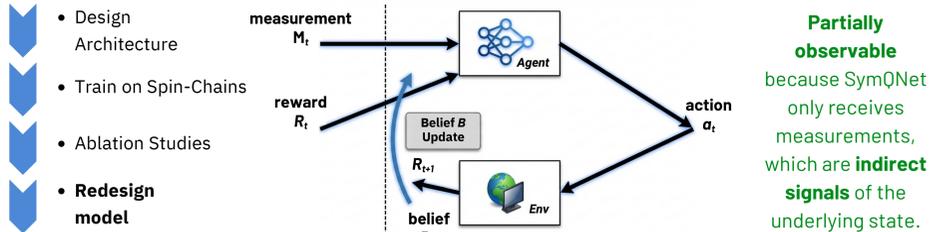
Core Idea: Learn a **strategy** to choose which measurements give the **most information** to estimate energy **faster**.

Action Defined as Selecting Measurement (with Qubit and Basis)



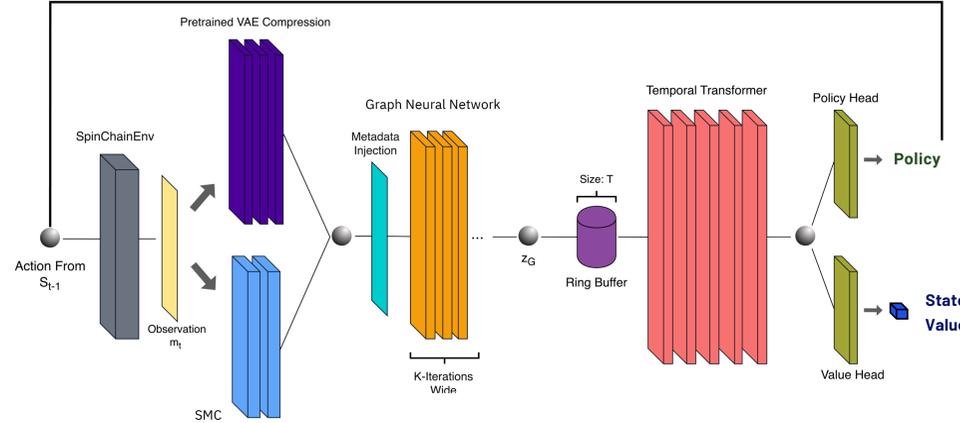
Bayesian Sequential Monte Carlo (SMC) is our **estimator**. It's what internally estimates energies and tells us how uncertain it is about parameters¹⁶. **SymQNet's job is to accurately guide SMC.**

Engineering Cycle & Partially Observable Markov-Decision Process (POMDP) Framing



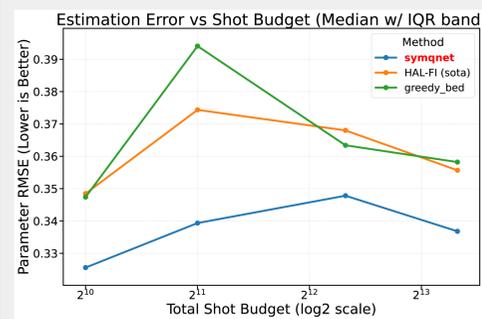
Partially observable because SymQNet only receives measurements, which are **indirect signals** of the underlying state.

Architecture Iteratively Designed to Account for Tough Learning Conditions



Results: Evaluation of Performance, Scaling Behavior, VQE Integration, and Reproducibility

SymQNet Offers Decisive Win over State-of-the-art Baselines



Decisive Win Over SOTA

Achieves **lower Parameter RMSE** than HAL-FI (SOTA) and GreedyBED across all shot budgets (2^{10} - 2^{13}).

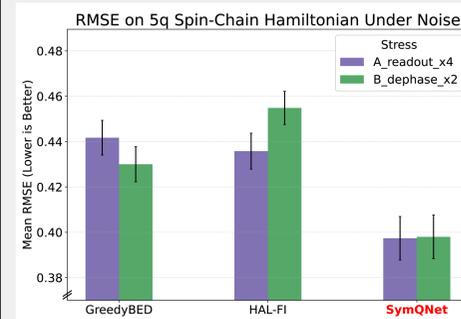
Statistically significant: **p < 0.05** via paired-Wilcoxon rank test (Holm-corrected)

- Softens exponential error growth in downstream Hamiltonian use.

Same budget. Better Results.

Shot Budget Benchmarking versus SOTA

Noise Stress-Test Shows Support for Hypothesis



Robust Under Realistic Noise

SymQNet outperforms both GreedyBED and HAL-FI under **readout noise (x4)** and **dephasing noise (x2)**.

Statistically significant ($\pm 2SEM$).
p < 0.05 via paired-Wilcoxon test (Holm-corrected)

Gains maintained under realistic conditions: **ready for deployment.**

Deployment: `pip install symqnet-molopt`

Noise Stress-Test (Readout & Dephase) for Parameter Estimation

Molecular Application Shows Significant Speedup and Accuracy Gains

46.7%

Shot Reduction

Fewer shots to reach success threshold

21.2%

MAE Reduction

Better mean energy error reduction vs. baseline

4/4

Molecules Won

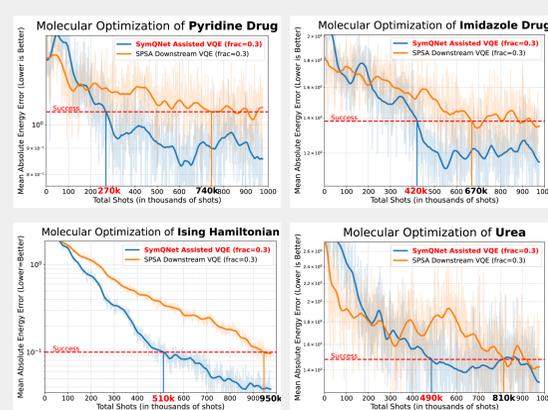
Better solution on Pyridine, Imidazole, Urea & Ising ($p < 0.05$)

1.88x

Drug Discovery Speedup

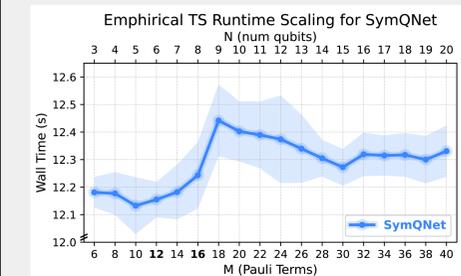
More molecules screened under the same budget

Energy Estimation with SymQNet on VQE Job versus Standard Pauli-Term Estimation



VQE MO Jobs with Energy Estimation on SymQNet vs Regular

Empirical Runtime Scaling Tests Show O(1) Scaling



Why Should I Care?

Cost doesn't explode as molecules get larger. **Significant applicability potential.**

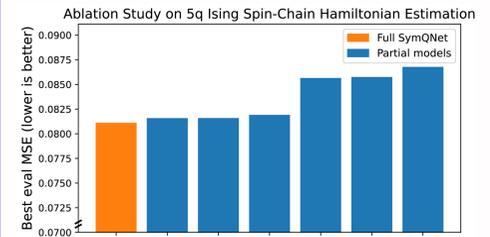
Reference: Baseline methods typically scale as $O(N^4)$ and $O(M)$.¹⁷

Empirically Flat Runtime Scaling (up to 20 qubits) and Reproducibility

Ablation Studies on Prototype

Ablation Studies show Maximal Component Utilization

Provides experimental **evidence SymQNet learned**.



Paired t-test reveals $p < 0.05$ (Holm-corrected).

Component	Impact
No Metadata	Removing metadata injection degrades performance (context matters!).
No Graph	Removing qubit graph structure hurts as spatial relationships are informative.
Reduced Particles	Fewer SMC particles (1.6x) reduces Bayesian estimation quality.

Ablation studies systematically remove or degrade each component of SymQNet to verify its contribution. Across 30 seeds, **SymQNet (Full)** achieves the **lowest MSE**, confirming that every design choice contributes meaningfully to performance.

Conclusions & Project Recap

From results, hypothesis appears to be supported!

Maximizing information from each shot with RL can lead to significant speedup of QMO

Drug Discovery Speedup

1.88x more molecules screened per budget. 46.7% shot reduction. 21.2% better energy error.

Beats state-of-the-art

Lower RMSE than HAL-FI & GreedyBED under same shot budget and under quantum noise ($p < 0.05$).

Empirical O(1) Scaling

Improves from $O(N^4)$, $O(M)$ to $O(1)$ (empirically).

Deployed: `pip install symqnet-molopt`

Next Steps and My Vision

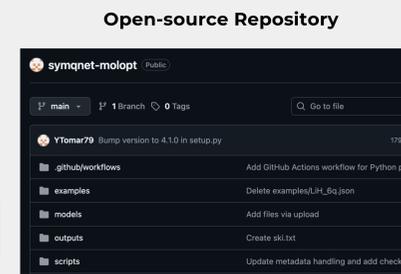
Next Steps for SymQNet

1 Higher-Dimensional Hamiltonians
Train on larger, more complex molecular systems

2 Real-Life Deployment
Trade off reward with device cost on actual quantum hardware

My Vision of Accelerating Medicine

- Faster drug development** rate with lower costs
- Better drugs** with lower clinical failure rates



SymQNet is currently **patent-pending**