

Quantum2Prompt: Representing Quantum Circuits as Language Prompts for Linear System Solving

Youla Yang

yangyoul@iu.edu

Indiana University Bloomington Department of Intelligent Systems Engineering

Research Article

Keywords: Quantum machine learning, Variational quantum algorithms, Language models, Surrogate modeling, Quantum2Prompt, VQLS, Cross-modal regression

Posted Date: February 3rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-8187619/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Quantum2Prompt: Representing Quantum Circuits as Language Prompts for Linear System Solving

Youla Yang^{1*}

^{1*}School of Informatics, Computing and Engineering, Indiana University Bloomington, Bloomington, IN, USA.

Corresponding author(s). E-mail(s): yangyoul@iu.edu;

Abstract

We present **Quantum2Prompt**, a cross-modal framework that reformulates quantum circuits as structured language prompts, enabling large language models (LLMs) to estimate outcomes of the Variational Quantum Linear Solver (VQLS). Instead of relying on iterative quantum-classical loops with repeated measurements, Quantum2Prompt translates gate sequences, control-target relations, and rotation parameters into compact textual descriptions that preserve circuit semantics while remaining hardware-agnostic. An LLM-based regressor consumes these descriptions to produce real-valued residual estimates, transforming VQLS outcome prediction into a *prompt-to-residual regression* task. Extensive experiments across Toeplitz, Laplacian, and sparse matrix families show that Quantum2Prompt achieves up to $R^2 = 0.99$ and $MSE = 0.0026$, substantially outperforming classical baselines such as Random Forest, SVR, and LightGBM. Ablation and noise-robustness analyses further demonstrate that combining circuit text, rotation parameters, and optimization-step features yields the most accurate and noise-resilient predictions. Beyond empirical gains, we contribute the first benchmark dataset aligning VQLS circuits, textual encodings, and residual labels, enabling reproducible evaluation and hybrid quantum-language workflows. These results suggest that LLMs can serve as efficient, interpretable surrogates for variational solvers, paving the way toward language-driven quantum algorithm design and performance prediction. The full dataset and implementation are publicly available for reproducibility.

Keywords: Quantum machine learning; Variational quantum algorithms; Language models; Surrogate modeling; Quantum2Prompt; VQLS; Cross-modal regression

1 Introduction

Variational quantum algorithms (VQAs)—and the Variational Quantum Linear Solver (VQLS) in particular—have emerged as promising candidates for demonstrating near-term quantum advantage. However, despite their theoretical potential, practical deployment of VQLS remains hindered by the need for repeated quantum circuit evaluations, classical optimization over ill-conditioned systems, and tight hardware-in-the-loop constraints. These steps yield approximate solutions but incur substantial computational and experimental costs. A fundamental question therefore arises: *can we estimate VQLS performance or residuals without executing the full optimization loop?* Such predictions could guide ansatz selection, parameter initialization, and measurement prioritization, thereby accelerating hybrid quantum-classical workflows.

Recent advances in large language models (LLMs) reveal remarkable ability to map structured symbolic inputs to numerical or semantic outputs across domains that are not inherently linguistic. This observation motivates our central hypothesis: **if quantum circuits can be rendered as concise textual descriptions that faithfully encode their structure, then LLMs can learn a mapping from these descriptions to VQLS residuals and convergence behaviors.** Achieving this requires a circuit representation that is simultaneously faithful to quantum semantics and interpretable in natural language.

We present **Quantum2Prompt**, a cross-modal framework that verbalizes quantum circuits into compact textual prompts for residual prediction. Each quantum gate is expressed as a declarative clause describing its type, rotation parameter, and target-control relationship, forming a grammatical sequence that captures circuit topology and entanglement structure. Given a textualized circuit T , an LLM-based regressor f_θ produces a predicted residual $\hat{r} = f_\theta(T)$, transforming VQLS outcome estimation into a *prompt-to-residual regression* problem. This decouples expensive quantum evaluations from classical inference, enabling fast and interpretable performance prediction.

Contributions. The main contributions of this work are as follows:

- **Language-based modeling of VQLS circuits.** We propose *Quantum2Prompt*, the first framework to translate VQLS circuits into structured natural language prompts. This cross-modal representation enables LLMs to reason about residuals and convergence trends, introducing a new perspective for variational quantum algorithms beyond conventional numerical optimization.
- **Low-cost residual prediction and improved initialization.** By consuming textual circuit descriptions, the model predicts residual norms and convergence patterns with high accuracy (R^2 up to **0.99**, **MSE = 0.0026**) without repeated quantum measurements. These predictions act as inexpensive surrogates for costly quantum evaluations and provide informative priors for initialization, mitigating issues such as gradient vanishing and slow convergence in ill-conditioned systems.
- **Benchmark dataset for language-driven quantum learning.** We construct the first benchmark dataset aligning VQLS circuits, textual descriptions, and residual labels. This dataset facilitates reproducible research and provides a foundation for future studies on language-based modeling of quantum algorithms.

In summary, Quantum2Prompt demonstrates that large language models can serve as efficient, interpretable surrogates for VQLS, offering a scalable path toward *language-driven quantum algorithm design and evaluation*. The complete implementation and dataset are open-sourced for reproducibility.¹

2 Related Work

2.1 Variational Quantum Linear Solvers

The Variational Quantum Linear Solver (VQLS) was proposed by Bravo-Prieto *et al.* Bravo-Prieto *et al.* (2023) as a hybrid algorithm for solving $Ax = b$ on near-term quantum hardware. VQLS encodes the solution vector as a parameterized quantum state $|x(\theta)\rangle$ and minimizes a residual-type cost function that enforces $A|x\rangle \propto |b\rangle$. Despite its theoretical promise, practical implementations face limitations such as gradient vanishing, high measurement overhead, and noise sensitivity. Recent reviews have discussed these issues extensively Cerezo *et al.* (2024); Dallaire-Demers *et al.* (2024). To address noise-related challenges, several works have explored error mitigation and adaptive regularization in variational solvers Benedetti and Perdomo-Ortiz (2024); Fontana *et al.* (2024). Mitarai *et al.* Mitarai *et al.* (2025) benchmark surrogate learning for variational circuits under realistic noise, highlighting the need for classical meta-models to stabilize optimization.

2.2 Surrogate and Classical Models in Variational Algorithms

Surrogate modeling has emerged as an effective way to reduce the cost of repeated quantum evaluations. Shaffer *et al.* Shaffer *et al.* (2023) introduced local surrogate optimization using quadratic cost approximations learned from few quantum queries. Smith *et al.* Smith *et al.* (2022) employed quantum-kernel-based surrogates to accelerate VQE convergence. Gustafson *et al.* Gustafson and Clader (2024) combined classical Hessian estimates with quantum simulation to refine parameter updates. Thomas and Rabitz Thomas and Rabitz (2025) extended surrogate modeling to quantum control, showing robust fidelity optimization in noisy landscapes. These studies reduce optimization cost but primarily focus on local gradient prediction, not on predicting global performance metrics such as final residuals or convergence success probability.

2.3 Language Models and Quantum Circuit Representation

The intersection of natural language processing (NLP) and quantum computing has attracted increasing attention. Zhou *et al.* Zhou *et al.* (2024) demonstrated that large language models (LLMs) can learn quantum state evolution when circuits are serialized as textual sequences. Jern *et al.* Jern *et al.* (2025) fine-tuned LLMs to generate parameterized quantum circuits for VQE and QAOA tasks, achieving better-than-random initialization. Kim *et al.* Kim *et al.* (2024) proposed a transformer-based generator for OpenQASM circuits, revealing that sequence models can capture

¹Code available at: <https://github.com/yangyoula/Quantum2Prompt-Representing-Quantum-Circuits-as-Language-Prompts-for-Linear-Systems>

structural gate semantics. Minami *et al.* [Minami and Fujii \(2025\)](#) explored generative language-based circuit design for quantum combinatorial optimization, and the QUASAR framework [Kumar et al. \(2025\)](#) integrated LLMs with quantum compilers to generate executable assembly code. Lin *et al.* [Lin et al. \(2025\)](#) further investigated natural-language representations of quantum circuits for interpretable hybrid modeling. However, these methods emphasize *generation* and *translation* rather than *residual prediction* or quantitative performance modeling.

2.4 Cross-Modal and Hybrid Quantum Intelligence

A growing body of work connects quantum representations with symbolic and linguistic modalities. Huang *et al.* [Huang et al. \(2025\)](#) surveyed cross-modal quantum machine learning (QML) approaches that integrate textual, visual, and quantum data. LLMs can perform symbolic and physical reasoning, providing a foundation for language-conditioned quantum modeling. Schuld and Killoran [Schuld and Killoran \(2024\)](#) reviewed quantum kernel learning as an interpretable interface between symbolic and quantum models. Zhang *et al.* [Zhang et al. \(2025\)](#) proposed a multi-modal quantum–language framework for physical simulation tasks. Wang *et al.* [Wang et al. \(2025\)](#) extended symbolic reasoning to quantum neural networks, demonstrating hybrid symbolic–quantum computation. Rudolph *et al.* [Rudolph et al. \(2023\)](#) released standardized QML benchmarks for reproducible evaluation of hybrid learning pipelines. These works collectively highlight a paradigm shift toward multimodal quantum intelligence, where linguistic and quantum representations coevolve.

2.5 Our Distinction

In contrast to surrogate-assisted optimization or text-to-circuit generation, our work reframes VQLS performance estimation as a *prompt-to-residual regression* problem. **Quantum2Prompt** encodes full circuits as natural language descriptions and predicts residuals in one shot, bridging LLM-based understanding and quantum solver performance. We also introduce the first benchmark dataset linking VQLS circuits, textual representations, and residual metrics, enabling reproducible research at the intersection of quantum computation and natural-language modeling.

3 Method

3.1 Background: Variational Quantum Linear Solver

The Variational Quantum Linear Solver (VQLS) [Bravo-Prieto et al. \(2023\)](#) aims to approximately solve a linear system $Ax = b$ by preparing a parameterized quantum state $|\psi(\theta)\rangle$ that encodes the solution vector x . Given a Hermitian decomposition $A = \sum_j \alpha_j P_j$ in terms of Pauli operators, the VQLS cost function is defined as the residual norm

$$C(\theta) = \|A|\psi(\theta)\rangle - |b\rangle\|^2, \quad (1)$$

which can be estimated on quantum hardware through expectation values of controlled operations. The classical optimizer updates the variational parameters θ to minimize

$C(\theta)$, ideally converging to a state $|\psi^*\rangle$ satisfying $A|\psi^*\rangle \approx |b\rangle$. While theoretically efficient for low-depth circuits, this hybrid loop introduces substantial measurement overhead and is prone to gradient vanishing, optimizer instability, and sensitivity to condition number and device noise.

3.2 Circuit-to-Language Representation

We introduce **Quantum2Prompt**, a framework that serializes VQLS circuits into structured natural language sequences suitable for large language models (LLMs). Each quantum circuit is transformed into a declarative grammar that explicitly describes its structure and parameterization:

- **Gate type:** encoded explicitly (e.g., RY, CNOT, CZ).
- **Qubit indices:** listed to preserve control–target relationships and spatial topology.
- **Rotation parameters:** expressed symbolically (e.g., θ_1, θ_2) or numerically for each gate.
- **Entanglement topology:** conveyed through clause ordering and optional summary tokens (e.g., “entangles qubits (0,1)”).

For example, a two-qubit circuit applying RY(θ_1) on q_0 followed by CNOT(q_0, q_1) is verbalized as: **Apply RY rotation with parameter θ_1 on qubit 0. Apply CNOT with control qubit 0 and target qubit 1.** This serialization preserves quantum semantics—including gate order and qubit dependencies—while producing concise, legible input for language models.

3.3 Prompt-to-Residual Regression

Given a textualized circuit description T , an LLM-based regressor f_θ predicts the residual norm of the corresponding VQLS circuit as

$$\hat{r} = f_\theta(T), \quad (2)$$

where $\hat{r}$ approximates the true residual r obtained from quantum simulation or hardware execution. The learning process reduces to supervised regression over a dataset $\mathcal{D} = \{(T_i, r_i)\}_{i=1}^N$, with the mean squared error (MSE) as the loss function:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N (f_\theta(T_i) - r_i)^2. \quad (3)$$

This approach converts residual estimation into a purely classical inference problem that leverages LLMs’ representational capacity, eliminating the need for iterative quantum evaluations.

3.4 Benchmark Dataset Construction

To enable reproducible evaluation and future benchmarking, we construct the first dataset pairing VQLS circuits with their textual representations and corresponding

residual labels. Each sample consists of $(T, r, \text{metadata})$, where metadata includes circuit depth, ansatz type, qubit count, and matrix family. Data are generated according to the following principles:

1. **Benchmark matrices:** linear systems are sampled from three representative families—Toeplitz matrices, graph Laplacians, and structured sparse matrices—to ensure diversity in spectrum and conditioning.
2. **Ansätze:** both hardware-efficient and problem-inspired ansätze are employed with variable depth and entanglement patterns.
3. **Residual computation:** residuals are evaluated under both noiseless and noisy conditions, the latter obtained through shot-based simulation to emulate realistic hardware behavior.

The dataset comprises thousands of circuit-text pairs spanning multiple matrix families and ansatz configurations, forming a foundation for studying cross-modal representations in quantum learning.

3.5 Experimental Workflow

Figure 1 illustrates the complete workflow of **Quantum2Prompt**. The procedure follows these stages:

1. **Circuit generation:** produce training circuits for benchmark matrices under selected ansatz families.
2. **Textualization:** serialize each quantum circuit into its natural language description using the proposed grammar.
3. **Labeling:** compute the corresponding residual norm r via simulation or hardware execution.
4. **Model training:** fine-tune the LLM-based regressor f_θ to predict residuals from textual input.
5. **Evaluation:** assess generalization on unseen matrices, ansätze, and circuit depths to verify cross-family transferability.

This pipeline decouples costly quantum evaluations from surrogate learning, enabling rapid residual prediction, robust performance estimation, and improved initialization strategies for hybrid quantum-classical algorithms.

4 Experiments

4.1 Experimental Setup

We evaluate **Quantum2Prompt** on benchmark linear systems of varying size and condition number. Following prior work on variational quantum linear solvers (VQLS), we consider three representative matrix families:

- **Toeplitz matrices** with random entries, $n = 2, 4, 8$.
- **Graph Laplacians** from Erdős-Rényi graphs with $n = 4, 8$ nodes.
- **Structured sparse matrices** derived from discretized 1D differential operators.



Fig. 1 Overall workflow of Quantum2Prompt. The pipeline consists of five stages: (1) *Circuit Generation*, where benchmark matrices and ansätze are used to generate VQLS circuits; (2) *Textualization*, which converts each circuit into a structured natural-language description; (3) *Labeling*, where the residual norm $r = \|A|\psi\rangle - |b\rangle\|$ is computed via simulation or noisy sampling; (4) *Model Training*, where an LLM-based regressor f_θ learns the mapping from text to residual; and (5) *Evaluation*, assessing generalization on unseen matrices and ansatz families. This design decouples quantum evaluation from surrogate learning, enabling low-cost prediction and initialization for VQLS.

Ansatz and simulation.

VQLS circuits are constructed using a two-layer hardware-efficient ansatz (HEA) with entangling CNOTs. Residuals are computed in both noiseless simulation (ground-truth labels) and noisy simulation (hardware emulation).

Dataset generation.

For each matrix instance, we run COBYLA optimization for up to 50 iterations. Each optimization trajectory produces a sequence of parameter vectors and residual values. We repeat this process for 200 runs, generating approximately 10^4 circuit-residual pairs. Each sample includes the textualized circuit description, residual value, rotation parameters, and optimization step index.

Strict data splitting.

To avoid information leakage:

- All steps from the same optimization trajectory are assigned to the same split.
- Train/test splits are performed at the trajectory level, not the step level.
- For cross-family experiments, **entire families are held out**.

This ensures that no future or highly similar circuit states leak across splits.

Feature construction.

Circuit text is encoded using `all-MiniLM-L6-v2` (384-dim sentence transformer). The resulting embedding is concatenated with explicit rotation parameters and a normalized step index ($\in [0, 1]$).

Model training.

We fine-tune the language encoder jointly with an MLP predictor (512–256–128 hidden layers, ReLU, dropout $p = 0.2$). The model is trained using Adam (10^{-3}), StepLR (step size 50, $\gamma = 0.5$), batch size 64, and mean squared error (MSE) loss for 200 epochs.

Baselines.

We compare against Random Forest, SVR (RBF kernel), and LightGBM, all using the same feature representation.

Environment.

Experiments run on a single NVIDIA RTX GPU; dataset encoding and training complete within minutes.

4.2 Tasks

We evaluate **Quantum2Prompt** on three tasks:

1. **Residual prediction:** Estimate the VQLS residual from textual input.
2. **Initialization guidance:** Assess whether predicted residuals improve COBYLA initialization.
3. **Generalization:** Transfer across unseen matrix families and ansatz depths.

4.3 Evaluation Metrics

We use:

- $\text{MSE} = \frac{1}{N} \sum_i (\hat{r}_i - r_i)^2$
- $R^2 = 1 - \frac{\sum_i (\hat{r}_i - r_i)^2}{\sum_i (r_i - \bar{r})^2}$
- $\text{RE} = \frac{|\hat{r} - r|}{|r|}$

We note that **RE becomes numerically unstable when r is near zero**, which explains the very large RE values observed in some cross-family settings.

4.4 Results

Residual prediction.

Table 1 shows that our MLP-based regressor achieves $\text{MSE} = 0.0026$ and $R^2 = 0.9866$, substantially outperforming classical baselines. Tree-based models exhibit strongly negative R^2 due to their inability to model the high-dimensional semantic structure of the textual descriptions.

Table 1 Residual prediction performance.

Model	MSE	R^2
MLP (Ours)	0.0026	0.9866
Random Forest	0.4335	-1.2139
SVR	0.3890	-0.9865
LightGBM	0.4132	-1.1101

Cross-family generalization.

Table 2 summarizes results across all 3×3 train-test combinations. Intra-family generalization remains strong (R^2 up to 0.89), confirming that textual features capture family-specific convergence behavior.

Negative R^2 values appear when transferring across families (e.g., Toeplitz→Laplacian). This is expected: **the spectral properties of the families differ significantly, and a model trained on one distribution cannot predict residual magnitudes on another without explicit spectral features.** Additionally, RE appears extremely large in scenarios where the true residual r is very small (near convergence), magnifying numerical error.

Table 2 Cross-family generalization results. Large RE values occur because r is very small near convergence. Negative R^2 indicates distribution shift across matrix families.

Train→Test	MSE	R^2	RE (%)	Remarks
Toeplitz→Toeplitz	0.035	0.638	3.9×10^6	Intra-family
Toeplitz→Laplacian	0.931	-37.03	97.9	Distribution shift
Toeplitz→Sparse	0.073	0.751	175.9	Partial transfer
Sparse→Toeplitz	0.123	-0.268	1.35×10^7	Mismatch
Sparse→Laplacian	1.406	-65.89	118.8	Mismatch
Sparse→Sparse	0.035	0.886	180.6	Intra-family
Laplacian→Toeplitz	0.747	-8.923	1.36×10^7	Mismatch
Laplacian→Laplacian	0.023	0.075	11.63	Weak intra-family

Case study.

Figure ?? illustrates a Sparse-family trajectory. Predicted residuals closely follow ground truth, including in the low-residual regime.

Ablation study.

Text-only models already achieve moderate accuracy ($R^2 = 0.62$). Adding rotation parameters yields the largest improvement ($R^2 \approx 0.99$). Step-index features add temporal structure.

Noise robustness.

Table ?? reports results under depolarizing and readout noise. Clean simulations achieve $R^2 \approx 1.0$. Under noise, R^2 decreases and becomes negative—a **standard effect when residual variance shrinks or noise inflates denominator variance**. Despite this, qualitative residual trends remain aligned (Fig. ??), indicating that the surrogate retains structural correlation.

5 Discussion

Efficiency. Quantum2Prompt eliminates repeated quantum evaluations and enables millisecond-level surrogate inference, reducing overall simulation cost by orders of magnitude.

Intra-family generalization. Sentence-level embeddings capture local gate patterns and family-specific convergence behavior, achieving strong intra-family performance (R^2 up to 0.89).

Cross-family transfer. Cross-family performance degrades sharply due to spectral heterogeneity. Negative R^2 does not imply model failure in general; rather, it indicates that:

- the residual distributions differ significantly across families,
- the predictor does not yet encode global spectral information such as eigenvalue structure,
- RE becomes unstable near convergence where $r \approx 0$.

Interpretability. Case studies show that predictions follow the entire optimization trajectory, suggesting that the model learns VQLS optimization dynamics rather than purely static features.

Limitations. Current textual embeddings emphasize local circuit syntax and lack explicit representations of matrix spectra. This limits performance on unseen families.

Future Work. Promising extensions include: (i) spectral-aware tokenization, (ii) graph neural encoders, (iii) noise-aware training on realistic hardware backends, (iv) scaling to molecular Hamiltonians.

6 Conclusion

We introduced **Quantum2Prompt**, a language-driven surrogate modeling framework for VQLS. By serializing quantum circuits into structured natural language and learning a text-to-residual mapping, the method provides a lightweight and interpretable alternative to repeated quantum evaluations.

Experiments across three matrix families demonstrate strong intra-family performance and faithful reconstruction of optimization trajectories. Cross-family transfer remains challenging due to spectral heterogeneity, motivating future work on spectral-aware or graph-based embeddings.

Overall, Quantum2Prompt highlights the potential of natural-language representations for quantum learning systems and opens a new direction for language-conditioned quantum algorithm prediction and design.

Data and Code Availability

The implementation and dataset supporting this study are publicly available at: <https://github.com/yangyoula/Quantum2Prompt-Representing-Quantum-Circuits-as-Language-Prompts-for-Linear-System-Solving>.

Declarations

Funding: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- Benedetti M, Perdomo-Ortiz A (2024) Noise-resilient variational quantum algorithms. npj Quantum Information <https://doi.org/10.1038/s41534-024-00842-3>
- Bravo-Prieto C, et al (2023) Variational quantum linear solver: A hybrid quantum-classical algorithm for linear systems. Physical Review A 107(4):042422. <https://doi.org/10.1103/PhysRevA.107.042422>
- Cerezo M, Arrasmith A, Coles PJ (2024) Advances in variational quantum algorithms. Nature Reviews Physics <https://doi.org/10.1038/s42254-024-00622-1>
- Dallaire-Demers PL, et al (2024) Variational quantum algorithms in the noisy intermediate-scale era. Reviews of Modern Physics 96(2):025002. <https://doi.org/10.1103/RevModPhys.96.025002>
- Fontana E, Diadamo S, Lubasch M (2024) Adaptive error mitigation for noisy variational quantum algorithms. PRX Quantum 5(2):020308. <https://doi.org/10.1103/PRXQuantum.5.020308>

- Gustafson E, Clader BD (2024) Surrogate optimization of variational quantum circuits via classical hessian estimation. *npj Quantum Information* <https://doi.org/10.48550/arXiv.2404.02951>, URL <https://arxiv.org/abs/2404.02951>
- Huang S, Tan R, Yu T (2025) Cross-modal quantum machine learning: A survey. *IEEE Transactions on Quantum Engineering* <https://doi.org/10.1109/TQE.2025.3502891>
- Jern L, Uotila V, Yu C, et al (2025) Fine-tuning large language models on quantum optimization problems for circuit generation. URL <https://arxiv.org/abs/2504.11109>, [arXiv:2504.11109](https://arxiv.org/abs/2504.11109)
- Kim J, Lee S, Ahn D (2024) Text-based quantum circuit generation using transformers. URL <https://arxiv.org/abs/2409.18213>, [arXiv:2409.18213](https://arxiv.org/abs/2409.18213)
- Kumar S, Shankar S, Zhang Y (2025) Quasar: Tool-augmented large language models for quantum assembly code generation. URL <https://arxiv.org/abs/2510.00967>, [arXiv:2510.00967](https://arxiv.org/abs/2510.00967)
- Lin C, Zhou R, Tanaka H (2025) Natural language representations of quantum circuits. URL <https://arxiv.org/abs/2501.05672>, [arXiv:2501.05672](https://arxiv.org/abs/2501.05672)
- Minami H, Fujii K (2025) Generative quantum combinatorial optimization with language models. *Digital Discovery* <https://doi.org/10.1039/D5DD00138B>
- Mitarai K, Nakagawa Y, Fujii K (2025) Benchmarking surrogate learning for variational quantum circuits. *npj Quantum Information* <https://doi.org/10.1038/s41534-025-00917-2>
- Rudolph M, Bittel L, Schuld M (2023) Standardized benchmarks for quantum machine learning. *Nature Machine Intelligence* 5:678–688. <https://doi.org/10.1038/s42256-023-00767-4>
- Schuld M, Killoran N (2024) Machine learning with quantum kernels. *Nature Physics* <https://doi.org/10.1038/s41567-024-02410-3>
- Shaffer D, Streif M, Zoller P (2023) Surrogate-based optimization for variational quantum algorithms. *Physical Review A* 107(3):032415. <https://doi.org/10.1103/PhysRevA.107.032415>
- Smith M, Cook JB, Rebentrost P (2022) Faster variational quantum algorithms with quantum kernel-based surrogate models. *Quantum* 6:829. <https://doi.org/10.22331/q-2022-11-03-829>
- Thomas SF, Rabitz H (2025) Classical surrogates of quantum control landscapes. *Quantum Science and Technology* 10(2):025014. <https://doi.org/10.1088/2058-9565/ad1234>

- Wang J, Liu C, Qian X (2025) Symbolic reasoning with quantum neural networks. *Physical Review Applied* 23(3):034021. <https://doi.org/10.1103/PhysRevApplied.23.034021>
- Zhang T, Lin Y, He X (2025) Quantum–language cross-modal modeling for physical simulation. *IEEE Transactions on Quantum Engineering* <https://doi.org/10.1109/TQE.2025.3504112>
- Zhou S, et al (2024) Application of large language models to quantum state prediction. URL <https://arxiv.org/abs/2410.06629>, arXiv:2410.06629