

RubriQ: Rubric-Guided Group Relative Policy Optimization for Constraint-Aware Quantum Circuit Synthesis

Ziqing Guo

ziqguo@ttu.edu

Texas Tech University

Ziwen Pan

Texas Tech University

Article

Keywords:

Posted Date: August 3rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10309063/v1>

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

[Read Full License](#)

Additional Declarations: No competing interests reported.

RubriQ: Rubric-Guided Group Relative Policy Optimization for Constraint-Aware Quantum Circuit Synthesis

Ziqing Guo^{1,2*} and Ziwen Pan¹

^{1*}Department of Computer Science, Texas Tech University, Lubbock, TX, USA.

²National Energy Research Scientific Computing Center (NERSC), Lawrence Berkeley National Laboratory, Berkeley, CA, USA.

*Corresponding author(s). E-mail(s): ziqguo@ttu.edu, zguo4@lbl.gov; Contributing authors: ziwen.pan@ttu.edu;

Abstract

Designing circuits that are algorithmically correct, low cost under a target fault-tolerant quantum computing (FTQC) architecture, and executable on near-term hardware remains a major bottleneck in the transition to scalable quantum computing. We introduce RubriQ, a scalable framework that formulates circuit synthesis as a large language model (LLM) code-generation task, optimized via group relative policy optimization (GRPO). Unlike conventional black-box neural critics, RubriQ employs a domain-grounded programmatic rubric as the reinforcement learning reward function, separating surface-code Clifford+ T cost metrics from near-term hardware validation and unitary fidelity. To support high-throughput training, RubriQ integrates GPU-accelerated CUDA-Q simulation directly into the reinforcement learning (RL) loop and is deployed on NERSC Perlmutter using DeepSpeed ZeRO2 across multinode NVIDIA A100 clusters. On benchmark tasks, RubriQ attains a per-dataset correctness pass rate of at least 96% and a mean T -gate compression of $3.31\times$ over correct-only circuits, significantly outperforming sparse-reward RL baselines ($2.05\times$), converging $2\text{--}3\times$ faster, and maintaining less than 1% near-term hardware-constraint violations. Validated on IBM and IonQ quantum processors, RubriQ establishes an automated, high-performance computing (HPC)-driven pipeline for generating surface-code-cost-aware circuits with separate near-term hardware validation.

1 Introduction

Quantum computers have the potential to deliver exponential speedups in fields such as chemistry [1], optimization [2], and cryptography [3]. However, realizing this potential necessitates the translation of high-level algorithms into low-level gate sequences that adhere to both the physical principles of error correction and the engineering constraints of actual hardware [4]. This translation constitutes a compilation problem, the classical computational cost of which increases exponentially with circuit size, thereby becoming an increasingly significant bottleneck in hybrid high-performance computing (HPC) and quantum workflows [5, 6].

Fault-tolerant quantum computing (FTQC) demands that every logical non-Clifford operation, principally the T gate, be realized through magic-state distillation, consuming four orders of magnitude more physical resources than Clifford gates [7–9]. Scaling quantum processing units (QPUs) beyond the noisy intermediate-scale quantum (NISQ) [10] era toward hundreds of logical qubits renders the classical synthesis of optimal gate sequences computationally intractable because of the exponential growth of the Hilbert space. Minimizing T -count directly determines the logical state distillation overhead, the scheduling complexity of surface-code patches, and the wall-clock time of a quantum computation within a hybrid HPC system [11–16]. Simultaneously, near-term devices impose additional validation constraints, where IBM processors [17] restrict operations to heavy-hex coupling topologies and native basis gates, whereas IonQ [18] trapped-ion systems offer all-to-all connectivity but with native operations. Therefore, quantum circuit synthesis requires a clear separation between FTQC cost metrics and near-term hardware executability.

Mainstream compilers [19], such as Qiskit [20], TKET [21], and CUDA-Q [22], rely on predefined heuristic rewrite rules and gate fusion. While effective for depth reduction, these frameworks are constrained by their static optimization libraries and lack the generative capability to discover algebraic simplifications beyond their programmed identities. Reinforcement learning approaches [23, 24] have shown promise but rely on small neural networks with sparse terminal rewards, leading to high-variance gradient estimates and poor exploration in deep circuit synthesis trajectories. Crucially, these RL methods operate on narrow action spaces (gate insertion and deletion) rather than generating circuits from scratch, thereby limiting their ability to discover structurally novel solutions. Similarly, the AlphaTensor paradigm [25, 26] has demonstrated that RL can discover quantum-translatable mathematical structures; however, this approach targets tensor decompositions instead of gate-level code and omits hardware constraints.

Recent breakthroughs in LLM-driven code generation [27], artificial intelligence (AI) for quantum [28], and RL-based reasoning [29–31] suggest an alternative paradigm: treating quantum circuit synthesis as a *structured code-generation task*, in which an LLM generates executable OpenQASM [32] or Qiskit quantum circuits written in Python, and a programmatic evaluator scores the output. This formulation inherits the LLM capacity to recognize algebraic patterns, compose modular subcircuits, and generalize across circuit families, capabilities that small RL agents lack. Moreover, it naturally connects to the metadesign paradigm for automated scientific




Fig. 1 RubriQ system architecture. The three-stage pipeline converts a natural-language prompt (1. system with user prompt encoding) into executable circuit code (2. autoregressive LLM generation with low-rank adaptation (LoRA) [34]) and a validated quantum circuit object (3. multi-strategy extraction). The rubric scoring engine (grey dashed box) supplies dense, bounded, deterministic rewards across correctness, FTQC cost, and near-term validation blocks; GRPO policy update (pink box) enables critic-free, scalable policy gradients using group normalization. The circuits are validated on IBM and IonQ hardware; the full pipeline runs on NERSC Perlmutter A100 clusters.

discovery [33], where an AI system generates, evaluates, and iterates on structured artifacts.

The aforementioned observations motivate the adoption of a synthesis approach based on LLMs; however, training such a pipeline is computationally prohibitive on standard hardware. We will subsequently quantify the per-iteration cost to elucidate the necessity of HPC resources and identify the primary bottleneck; the resolutions of these two inquiries directly inform the system architecture delineated in § 5.

Each GRPO training step executes four sequential phases, each with distinct compute characteristics (we refer additional symbols in § A):

1. **Generation** ($\propto N \cdot L$): autoregressive sampling of N candidate circuits per prompt from a multi-billion-parameter model, each up to $L=2,048$ tokens;
2. **Parsing** ($\propto N$): syntactic extraction and gate-level validation of each generated circuit;
3. **Rubric evaluation** ($\propto N \cdot 4^{n_q}$): here computing unitary fidelity against target algorithms requires $O(4^{n_q})$ complex-valued operations for n_q (number of qubits) density matrix simulation. The rubric score engine evaluates surface-code FTQC cost through Clifford+ T metrics and depth efficiency, and evaluates near-term hardware validation through backend native-gate and routing checks, as shown in Fig. 1;
4. **Policy update** ($\propto |\theta|$): group-relative advantage computation and distributed gradient accumulation across LoRA parameters.

Particularly, the fidelity evaluation in phase (3) significantly influences the total cost: at $n_q=8$, each simulation necessitates 65,536-dimensional complex matrix multiplications. With $N=8$ candidates per prompt across 500 prompts per epoch, a single epoch

demands approximately 4,000 such simulations. At $n_q=12$, this requirement escalates to approximately 70 GPU-hours per training run on 8 A100s. Neither phase (1) nor phase (4) can be offloaded to the CPU, and phase (3) is parallelized across candidates, a workload profile that is well-suited to multi-node GPU clusters with batch-parallel simulation. These requirements further inform our systems design (§ 5) by employing DeepSpeed ZeRO-2 for distributed gradient and optimizer-state partitioning while utilizing LoRA to enable parameter-efficient fine-tuning of the 7B model [35] on four A100 80 GB GPUs. To resolve the dominant computational bottleneck (shown in phase (3)), the framework integrates GPU-accelerated quantum simulation via CUDA-Q for high-throughput fidelity evaluation.

In this paper, we introduce **RubriQ**, a framework that unifies LLM-driven code generation, programmatic rubric scoring, and critic-free GRPO for constraint-aware quantum circuit synthesis, deployed at the HPC scale. Our contributions are:

1. **LLM-as-synthesizer formulation.** We recast quantum circuit synthesis as a code-generation task for instruction-tuned LLMs, where the model generates complete OpenQASM or Qiskit circuits from natural-language specifications of target algorithms (quantum Fourier transform, quantum phase estimation, Hamiltonian simulation, and data encoding). While concurrent work has explored smaller transformer models for circuit optimization [23, 26], **RubriQ** is, to the best of our knowledge, the first to apply billion-scale instruction-tuned LLMs to gate-level quantum synthesis with RL-based post-training (§ 3).
2. **Multi-objective rubric reward engine.** We designed a programmatic scoring engine that separates FTQC cost under a surface-code Clifford+ T architecture from near-term hardware validation on IBM and IonQ backends, while enforcing unitary process fidelity as a hard correctness condition. This replaces the neural critic with domain-grounded, interpretable rubrics. We provide a theoretical analysis of ranking consistency under dense rubric shaping (§ 4).
3. **Scalable HPC systems architecture.** We implement the full training pipeline on NERSC Perlmutter with strong-scaling efficiency exceeding 85% across 8 A100 GPUs. We provide a compute-cost analysis demonstrating that rubric shaping reduces the total simulation calls by 2–3 $\times$ versus sparse rewards (§ 5).
4. **Near-term hardware-validated evaluation.** We validated the highest-reward circuits on IBM Heron-3 (156-qubit superconducting) and IonQ Forte (36-qubit trapped-ion) hardware through an automated transpilation and submission pipeline, demonstrating near-term executability beyond simulation (§ 7).
5. **Synthesizing experimental evaluation.** Across 1,500 benchmark circuits, **RubriQ** achieves 3.31 $\times$ mean T -gate compression, converges 2–3 $\times$ faster, and reduces constraint violations by 5 $\times$ relative to the strongest baseline; ablations and robustness analyses are presented in the Appendix F.

2 Background and Problem Formulation

2.1 FTQC and the T-Gate Bottleneck

Within the stabilizer formalism [36], the Clifford group, which is generated by $\{H, S, \text{CNOT}\}$, allows for efficient classical simulation according to the Gottesman–Knill theorem [37, 38]. Achieving universal quantum computation necessitates the addition of at least one non-Clifford gate, typically the $T = \text{diag}(1, e^{i\pi/4})$ gate. The FTQC target in this work is a logical surface-code architecture with a Clifford+ T gate set, where Clifford operations are implemented through transversal gates or lattice-surgery operations [39], and logical T gates are supplied by magic-state factories. In this architecture, each logical T gate requires magic-state distillation, which consumes $O(10^3)$ physical qubits and completes within $O(10)$ code cycles, i.e., $O(10)$ rounds of stabilizer syndrome extraction, the elementary error-correction time step of a fault-tolerant computation, thereby making the T -count the primary cost metric for FTQC [9]. The T -depth, defined as the number of circuit layers containing at least one T gate, determines the sequential distillation bottleneck. Specifically, the T -count influences both the number of classical distillation simulations needed during compilation and the qubit overhead that should be coscheduled on the quantum device; thus, minimizing it directly decreases the classical computational budget of a hybrid quantum-classical workflow.

For a parameterized rotation $R_z(\theta)$, the gate is Clifford if and only if $\theta \in \{k\pi/2 : k \in \mathbb{Z}\}$. At $\theta = \pi/4$, the gate is equivalent to T ; all other non-Clifford angles require Solovay–Kitaev decomposition into $O(\log(1/\varepsilon))$ T gates for precision ε . Our Clifford+ T estimator classifies every gate in a generated circuit and flags the Toffoli gate (CCX) as contributing seven T gates via its standard decomposition.

2.2 Hardware Constraints for Near-Term Devices

In addition to FTQC cost estimation, we validate whether synthesized circuits can be executed on specific near-term hardware backends. These processors are not the target FTQC architecture; they provide an operational test of native-gate compatibility, connectivity, routing overhead, and depth constraints.

- **IBM Eagle/Heron** (127–156 superconducting qubits): Heavy-hex coupling topology, native gate set $\{\text{ECR/CZ, ID, RZ, SX, X}\}$, maximum circuit depth $\sim 1,000$. Non-adjacent two-qubit gates require SWAP insertion, incurring 3 additional two-qubit gates per SWAP. We also note that IBM Nighthawk QPU provides better scalability owing to its square-lattice topology with more couplers than the Heron series, but with a longer execution time.
- **IonQ Forte/Aria** (25–36 trapped-ion qubits): All-to-all connectivity (zero SWAP overhead), native gate set $\{\text{GPi, GPi2, MS}\}$, maximum circuit depth ~ 500 .

Fig. 2 contrasts restricted IBM heavy-hex connectivity with IonQ all-to-all connectivity using the same $\text{CNOT}(q_0, q_3)$ example. On IBM Heron, non-adjacent interactions require SWAP routing, increasing the two-qubit gate count. On IonQ Forte, the




Fig. 2 Near-term hardware topology comparison driving **RubriQ** validation score S_H . a). IBM Heron heavy-hexagonal connectivity restricts two-qubit interactions; a non-adjacent $\text{CNOT}(q_0, q_3)$ requires two SWAPs, adding six extra two-qubit gates. b). IonQ Forte all-to-all connectivity permits the same $\text{CNOT}(q_0, q_3)$ directly with zero routing overhead. **RubriQ** takes the maximum S_H over supported near-term backends; this validation score is reported separately from the surface-code FTQC cost metrics.

same interaction maps directly with zero routing overhead. This contrast motivates including routing overhead in the near-term validation score S_H .

2.3 LLMs and Reinforcement Learning for Code Generation

Large language models trained on code [27] have demonstrated a remarkable capacity for generating syntactically correct and functionally coherent programs from natural-language specifications. DeepSeek-R1 [29] showed that Group Relative Policy Optimization (GRPO) can elicit complex reasoning in LLMs without a learned critic, using only verifiable rewards. The rubric-RL paradigm [40] further demonstrated that structured, multidimensional reward signals outperformed scalar rewards for open-ended generation tasks.

Specifically, GRPO [30] eliminates the value-function critic by normalizing rewards within a group of N completions sampled from the same prompt

$$\hat{A}_i = \frac{R_i - \mu_R}{\sigma_R + \varepsilon}, \quad \mu_R = \frac{1}{N} \sum_{j=1}^N R_j, \quad (1)$$

$$\sigma_R = \sqrt{\frac{1}{N} \sum_{j=1}^N (R_j - \mu_R)^2}, \quad (2)$$

where R_i is the scalar reward for the i -th completion. The policy gradient becomes

$$\nabla_{\theta} J(\theta) = \frac{1}{N} \sum_{i=1}^N \nabla_{\theta} \log \pi_{\theta}(\mathbf{y}_i | \mathbf{x}) \hat{A}_i, \quad (3)$$

where $\mathbf{x}$ is the prompt and $\mathbf{y}_i$ is the i -th generated circuit. This formulation is particularly well-suited to quantum synthesis because the reward for a generated circuit is fully deterministic and programmatically computable, making a learned critic unnecessary.

2.4 Problem Statement

Given a target quantum algorithm $\mathcal{A}$ (e.g., QFT on n qubits) specified in natural language, we seek a policy π_θ that generates circuit code $\mathbf{y}$ maximizing the multi-objective rubric reward

$$\max_{\theta} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}, \mathbf{y} \sim \pi_\theta(\cdot|\mathbf{x})} [R(\mathbf{y}; \mathcal{A})], \quad (4)$$

where $R(\mathbf{y}; \mathcal{A})$ aggregates three conceptually distinct blocks: unitary correctness, FTQC cost under the target surface-code Clifford+ T architecture, and near-term hardware validation. Functional correctness is a hard gate defined by a fixed fidelity threshold $\tau_F = 0.99$: a circuit is admissible if and only if its average gate fidelity satisfies $F_{\text{avg}}(U_{\mathbf{y}}, U_{\mathcal{A}}) \geq \tau_F$. The objective in Eq. 4 is maximized over the admissible set, and FTQC cost metrics such as T -gate compression are reported only for admissible circuits. The gate is operationalized in § 3.2, where an inadmissible circuit receives zero reward. Here, F_{avg} denotes the average gate fidelity between the circuit unitary $U_{\mathbf{y}}$ and the target unitary $U_{\mathcal{A}}$.

3 The RubriQ Framework

3.1 LLM-as-Synthesizer: State-Action Formulation

We formulate quantum circuit synthesis as a conditional code-generation task in which the LLM receives a structured natural-language prompt and produces executable quantum circuit code. The pipeline comprises three stages: *prompt encoding*, *autoregressive generation*, and *circuit extraction*, as described below and illustrated in Fig. 1.

Stage 1: Prompt encoding (natural language $\rightarrow$ model input)

Each training example is assembled from two components concatenated into the chat-template format expected by the instruction-tuned base model.

1. **System prompt $\mathbf{x}_{\text{sys}}$** : encodes domain-level constraints shared across all tasks:

“You are a quantum circuit synthesis engine. Generate circuits using only standard gates. Prefer Clifford gates (H , S , $CNOT$, X , Z) over non-Clifford rotations. Minimize T -gate count. Ensure the circuit is functionally correct. Output the circuit as executable OpenQASM 2.0 or Qiskit Python code inside a code block.”

2. **User prompt $\mathbf{x}_{\text{user}}$** : specifies the target algorithm, qubit count, and output format:

“Synthesize a 4-qubit Quantum Fourier Transform circuit that minimises T -gate count and circuit depth. Output as OpenQASM 2.0.”

The full model input is $\mathbf{x} = [\mathbf{x}_{\text{sys}}; \mathbf{x}_{\text{user}}]$. The system prompt steers the LLM toward Clifford-heavy, resource-efficient solutions before RL training begins; GRPO then refines this prior through rubric-guided gradient updates.

Stage 2: Autoregressive generation (model input $\rightarrow$ raw text)

The LLM policy π_θ generates a completion $\mathbf{y}$ token by token with temperature $T=0.8$ and nucleus sampling [41] (top- $p=0.95$), up to 2,048 tokens. A output $\mathbf{y}$ interleaves natural-language reasoning with code (openqasm specification is shown in ¹)

“The QFT on 4 qubits applies Hadamard and controlled-phase gates. Here is an optimized circuit:”

```
OPENQASM 2.0; include "stdgates.inc"; qubit[4] q;
h q[0]; cp(pi/2) q[1], q[0]; cp(pi/4) q[2], q[0];
```

“This circuit requires T-gate resources after minimization, as the $\pi/4$ phase rotation is a non-Clifford operation that cannot be expressed solely through Clifford-group phases.”

The entire generated sequence constitutes a single RL action; N completions are sampled per prompt per training step.

Group size selection

The group size N within the GRPO framework governs a critical trade-off between variance and cost [30]. The standard error of the advantage estimator scales as $O(1/\sqrt{ND})$, where D represents the number of informative rubric component scores (we defer in § 4). With $D=5$, the effective sample size becomes $ND=5N$; consequently, setting $N=8$ under rubric shaping achieves the same estimator precision as $N=40$ under a single sparse reward, resulting in a fivefold reduction in generation and simulation cost per policy update. On the cost aspect, each additional candidate necessitates a complete forward pass (~ 0.4 s on A100 at 2,048 tokens) and one fidelity simulation (~ 0.2 s at $n_q=8$); increasing N from 8 to 16 approximately doubles the per-step wall time without a proportional reduction in variance (diminishing returns occur beyond $N>8$ with $D=5$). Empirical evidence indicates that $N=4$ results in noisy advantages that decelerate convergence by approximately 30% compared to $N=8$, whereas $N=16$ achieves only a 5% faster convergence at twice the per-step cost (Appendix D). Therefore, we establish $N=8$ as the Pareto-optimal operating point.

Stage 3: Circuit extraction (raw text $\rightarrow$ quantum circuit)

Because LLM outputs interleave code with explanations, a multi-strategy parser attempts extraction in priority order:

1. **Direct QASM parse:** detect an OPENQASM header and parse the block via Qiskit’s `QuantumCircuit.from_qasm_str`.
2. **Code-fence extraction:** locate markdown-style code blocks (qasm or python), extract the content, and parse.
3. **Safe Python execution:** if Qiskit Python code is detected, execute it in a restricted namespace (no filesystem I/O, no imports beyond `qiskit`) and capture the resulting `QuantumCircuit` object [20].

¹<https://openqasm.com/>

The first strategy to yield a valid `QuantumCircuit` with ≥ 1 qubit succeeds; if all three fail, the output receives reward $R=0$. This cascade recovers $\sim 60\%$ of initially unparseable outputs (Appendix F), which is critical during early training when the policy has not yet learned to produce clean code consistently.

Base model and parameter efficiency

We use a 7B-parameter instruction-tuned code-generation LLM [35] as the base policy π_θ , fine-tuned with LoRA [34] (rank 64, $\alpha=128$), which maintains trainable parameters at $\sim 0.4\%$ of the 7B total ($\sim 28\text{M}$ parameters). This reduces the AdamW optimizer state from $\sim 84\text{ GB}$ (full fine-tuning requires two momentum buffers per parameter in fp32) to $\sim 340\text{ MB}$, making ZeRO-2, which partitions only gradients and optimizer states without the per-layer parameter all-gathers required by ZeRO-3, memory-sufficient on 4 A100 80 GB GPUs (§ 5.2). The resulting communication savings are reflected in 85% strong-scaling efficiency across 8 GPUs (§ 5.4). Freezing the pretrained weights additionally preserves the base model code generation, which is critical for maintaining high parse rates under RL training (Appendix F).

3.2 The Rubric Reward Engine

Our idea was inspired by rubrics-as-the-reward under the constraint of reinforcement learning with verifiable rewards. Inherently, this relies on either sparse terminal rewards or parameterized neural critics. However, at the HPC scale, maintaining and synchronizing a multi-billion-parameter critic network across distributed GPUs introduces severe memory and communication bottlenecks. To resolve this, **RubriQ** replaces the neural critic entirely with a programmatic, multidimensional rubric. Instead of a sparse success signal, the rubric decomposes the synthesis objective into independently verifiable, domain-grounded scores. For a generated circuit c , the dense reward is calculated from bounded component scores $S_d \in [0, 1]$. This provides a graded learning signal, even for imperfect circuits.

The programmatic rubric scoring engine evaluates each generated circuit with separate FTQC-cost and near-term-validation blocks. For a generated circuit $c = \text{parse}(\mathbf{y})$, the aggregate reward is defined as

$$\begin{aligned} R(\mathbf{y}) &= g(c) [w_F S_F(c) + R_{\text{FT}}(c) + R_{\text{NT}}(c)], \\ R_{\text{FT}}(c) &= w_C S_C(c) + w_T S_T(c) + w_E S_E(c), \\ R_{\text{NT}}(c) &= w_H S_H(c). \end{aligned} \tag{5}$$

where $g(c) = \mathbf{1}[F_{\text{avg}}(c) \geq \tau_F]$ is the hard correctness gate defined in § 3.2. The reward is zero if the LLM output fails to parse into a valid circuit or if the circuit fails the correctness gate. R_{FT} estimates cost for the target surface-code Clifford+ T architecture, whereas R_{NT} validates executability on the near-term IBM and IonQ backends.

Weight design

The default weights $(w_C, w_T, w_H, w_F, w_E) = (0.20, 0.15, 0.15, 0.40, 0.10)$ encode a fidelity-first hierarchy grounded in quantum synthesis. A resource-efficient circuit has no value if its unitary does not match the target algorithm. Each non-Clifford gate requires $O(10^3)$ physical qubits for magic-state distillation [9]. T -count reduction and circuit efficiency form the FTQC cost block because they estimate logical non-Clifford resource usage and depth under the target surface-code model. Hardware validation forms a separate near-term block because it penalizes backend-incompatible operations that require costly transpilation on IBM and IonQ devices. Circuit efficiency favors compact implementations among circuits that are otherwise comparable in fidelity, FTQC cost, and near-term hardware validity. The weights sum to one, ensuring $R \in [0, 1]$ and satisfying the bounded-score assumption required by the GRPO advantage estimator (Proposition 1, Appendix C). These defaults were selected from a grid search over the ranges listed in Appendix D. Sensitivity analysis (Appendix E) confirms that compression is robust to ± 0.05 perturbations and identifies the fidelity weight as the strongest marginal factor.

Dimension 1: Clifford fraction (S_C)

Each gate is classified as Clifford or non-Clifford via stabilizer formalism. A parameterized rotation $R_z(\theta)$ is Clifford if $\theta/(\pi/2) \in \mathbb{Z}$ (within tolerance 10^{-8}). The score is the fraction of operational gates that are Clifford gates.

$$S_C(c) = \frac{|\{g \in c : g \text{ is Clifford}\}|}{|\{g \in c : g \text{ is operational}\}|}. \quad (6)$$

Dimension 2: T -count reduction (S_T)

An exponential decay penalizes circuits with high T -count:

$$S_T(c) = \exp(-\alpha_T \cdot T(c)), \quad (7)$$

where $T(c)$ is the total T -gate equivalent count and $\alpha_T = 0.05$ is the decay rate.

The scoring engine traverses the circuit gate list and classifies each gate into one of the following three categories. (a) *Clifford*: gates in $\{H, S, S^\dagger, X, Y, Z, \text{CNOT}, \text{CZ}, \text{SWAP}\}$ and parameterized rotations $R_z(\theta)$ with $\theta/(\pi/2) \in \mathbb{Z}$ contribute $T(g) = 0$. (b) *Known non-Clifford decompositions*: the T gate contributes 1, the Toffoli gate (CCX) contributes 7, following standard ancilla-free decomposition [11, 15], and the Fredkin gate contributes 7 (via Toffoli reduction). (c) *Arbitrary rotations*: any $R_z(\theta)$ with $\theta/(\pi/2) \notin \mathbb{Z}$ is assigned an estimated T -count of $3\lceil \log_2(1/\varepsilon) \rceil$ using the asymptotically optimal bound from the Ross-Selinger algorithm [12], evaluated at the target precision $\varepsilon = 10^{-10}$ (yielding ~ 50 T gates per non-Clifford rotation). The total value is $T(c) = \sum_{g \in c} T(g)$.

The exponential form $\exp(-\alpha_T T(c))$ was chosen over linear penalties for two reasons. First, it maps $T(c) \in [0, \infty)$ to $S_T \in (0, 1]$, maintaining the bounded-score assumption required by the GRPO advantage estimator (Proposition 1). Second, the exponential concavity rewards the first few T -gate eliminations disproportionately,

consistent with the diminishing marginal utility of T -count reduction observed in practice [13, 14], removing the first 10 T gates from a 50-gate circuit saves more distillation overhead than removing the last 10 from a 20-gate circuit, because distillation is batched and the scheduling overhead per batch is amortized.

Dimension 3: Near-term hardware validation (S_H)

The circuit is validated against each near-term validation backend $b \in \{\text{IBM Eagle, IBM Heron, IonQ Forte, IonQ Aria}\}$:

$$S_H(c) = \max_b \left[w_q \cdot \mathbf{1}[n_q(c) \leq Q_b] + w_d \cdot \mathbf{1}[D(c) \leq D_b] \right. \quad (8)$$

$$\left. + w_n \cdot \text{NativeFrac}(c, b) + w_r \cdot \exp(-\lambda \cdot \text{SwapOH}(c, b)) \right], \quad (9)$$

where Q_b, D_b are qubit and depth limits, NativeFrac is the fraction of gates already in the backend native set, and SwapOH is the estimated SWAP routing overhead (zero for all-to-all topologies). This term is a near-term hardware validation score rather than an FTQC resource estimate. See notations in § A.

Dimension 4: Functional fidelity (S_F)

Process fidelity is computed via the Hilbert-Schmidt inner product:

$$F_{\text{process}} = \frac{|\text{Tr}(U_A^\dagger U_c)|^2}{d^2}, \quad F_{\text{avg}} = \frac{d \cdot F_{\text{process}} + 1}{d + 1}, \quad (10)$$

where $d = 2^n$ is the Hilbert space dimension, U_A is the target unitary (built analytically for QFT, QPE, Heisenberg, and Ising models), and U_c is the circuit unitary extracted after stripping non-unitary operations (measurements, resets).

Hard correctness gate

The fidelity dimension enters the reward through a hard gate $g(c) = \mathbf{1}[F_{\text{avg}}(c) \geq \tau_F]$ with threshold $\tau_F = 0.99$, in addition to its weighted contribution $w_F S_F(c)$. A circuit that fails the gate receives $R(\mathbf{y}) = 0$, identical to a parse failure, irrespective of its Clifford, T -count, hardware, and efficiency scores. The gate encodes the physical requirement that an incorrect unitary is not a valid synthesis of the target algorithm, and it removes the reward-hacking channel by which the policy could inflate T -count compression through small circuits that violate the target unitary. Consequently, compression is well defined only over gated circuits, and every compression value reported in § 7 is computed over the correct-only subset quantified by the pass rate P (§ 6). The soft weight w_F is retained inside the gate so that admissible circuits still receive graded fidelity feedback, preserving the within-group reward dispersion required by GRPO.

Dimension 5: Circuit efficiency (S_E)

A depth-normalized compactness metric rewards circuits that solve the problem in fewer layers

$$S_E(c) = \exp\left(-\frac{\beta \cdot D(c)}{n_q^2}\right), \quad (11)$$

where $D(c)$ is the circuit depth and $\beta = 0.5$ is a scaling constant.

3.3 GRPO Training with Rubric Rewards

Here, GRPO supplies a critic-free update that scales to multi-billion-parameter sequence models on distributed GPUs, whereas **RubriQ** provides the dense, deterministic, prompt-comparable rewards that make group-relative learning statistically meaningful for quantum circuit synthesis.

Per-step training loop.

At each step, given a prompt $\mathbf{x}$ specifying a target algorithm, the policy samples $N=8$ completions $\{\mathbf{y}_1, \dots, \mathbf{y}_N\} \sim \pi_\theta(\cdot \mid \mathbf{x})$. The rubric engine maps each completion to a scalar R_i (§ 3.2); parse failures and circuits that fail the correctness gate yield $R_i=0$, so invalid and functionally incorrect generations participate in the same ranking as admissible circuits. Group-relative advantages $\hat{A}_i$ follow Eq. 1, and the LoRA parameters are updated via Eq. 3. Fig. 3 summarizes the flow from prompt through sampling, scoring, normalization, and policy update.

RubriQ with GRPO

GRPO asserts that, for a given prompt, the rewards $\{R_i\}_{i=1}^N$ are sufficiently comparable and informative, such that normalizing by subtracting the group mean and scaling by the group deviation offers a valuable learning signal. Traditional reinforcement learning often relies on sparse binary or terminal rewards; in such cases, many groups exhibit tied or nearly tied returns, resulting in $\hat{A}_i$ contributing minimally to the gradient, thereby causing training to stagnate. **RubriQ** addresses this issue by decomposing synthesis quality into bounded sub-scores $S_d \in [0, 1]$ for correctness, FTQC cost, and near-term hardware validation, ensuring that even suboptimal circuits receive graded feedback and maintaining high within-group dispersion early in training. The rubric is also deterministic given $(\mathbf{y}_i, \mathcal{A})$, which eliminates critic-target noise and aligns with the GRPO utilization of raw returns rather than bootstrapped value targets. Furthermore, the component scores decompose into embarrassingly parallel sub-evaluations, such as Clifford counting and statevector fidelity, which aligns with the batched reward pipeline on Perlmutter (§ 5).

3.4 Constraint-Aware Action Filtering

Generated outputs pass through two feasibility filters before contributing a graded reward. First, syntactic pre-filtering assigns zero reward to outputs that cannot be parsed into valid circuits and does not pass them to fidelity simulation, avoiding simulator calls for invalid generations. Second, the hard correctness gate assigns zero reward




Fig. 3 Per-step GRPO training flow. A prompt yields $N=8$ completions; parse failures receive $R=0$. The rubric engine scores each valid circuit across the correctness, FTQC-cost, and near-term-validation blocks and aggregates R_i . Group-relative normalization produces advantages $\hat{A}_i$ without a learned critic. The bar chart shows how GRPO amplifies high-reward circuits (green, right) and suppresses failures (red, left); dense rubric scoring keeps within-group dispersion high, enabling accurate ranking.

to parsed circuits whose average gate fidelity falls below τ_F , so functionally incorrect circuits are excluded from the compression statistics on the same footing as parse failures. Empirical parse-failure and gate-failure rates are reported in Appendix F.

4 The RubriQ Formalism

Claim 1 (Ranking consistency under dense rubric shaping). Let $S_d \in [0, 1]$ be the bounded rubric component scores positively correlated with the true synthesis utility U . Under the GRPO with group size N , the probability of pairwise return misranking decreases as the number of informative rubric component scores increases, provided that all component scores are Lipschitz continuous.

By applying Hoeffding concentration to bounded component scores, the multi-dimensional rubric reduces the variance of trajectory-level reward estimates relative to sparse scalar rewards. Group-centered normalization (Eq. 1) further amplifies this because the order statistics govern the gradient direction. Formally, under dense shaping with D informative dimensions, the standard error of the advantage estimator scales as $O(1/\sqrt{ND})$ versus $O(1/\sqrt{N})$ for scalar rewards.

Claim 2 (Fixed-weight bias at high complexity). If the stabilizer complexity $S(c)$ grows superlinearly with the baseline circuit size, whereas the achievable T -improvement grows sublinearly, a fixed weight vector $(w_C, w_T, w_H, w_F, w_E)$ induces a systematic under-estimation of high-complexity gains.

Appendix F shows that dense shaping accelerates convergence by 2–3 $\times$ (Claim 1), whereas circuits with $T_0 > 140$ exhibit systematic negative residuals under fixed weights (Claim 2). For the full derivation, see Appendix C.

5 Systems Design and HPC Execution

5.1 Hardware Platform

All training experiments were run on NERSC Perlmutter using NVIDIA A100 (80 GB HBM2e) GPU nodes connected via the Slingshot-11 interconnect. The software stack comprised PyTorch 2.3, DeepSpeed 0.14+, the `tr1` library for GRPO training, PEFT/LoRA for parameter-efficient fine-tuning, and Qiskit 2.1+ for circuit manipulation. GPU-accelerated quantum simulation used the CUDA-Q quantum framework [22].

5.2 Distributed Training Architecture

The training pipeline is orchestrated across multiple A100 GPUs using DeepSpeed ZeRO-2 [42], which partitions gradient and optimizer states across data-parallel workers while replicating model parameters on each GPU. We use ZeRO-2 rather than ZeRO-3 because the latter’s parameter partitioning is incompatible with PyTorch activation checkpointing, in which tensor-shape mismatches occur during gradient recomputation. With LoRA fine-tuning (rank 64, $\alpha=128$), only $\sim 0.4\%$ of the parameters are trainable, making ZeRO-2 memory-sufficient for a 7B model on 4 A100 80GB GPUs.

Each training step proceeds in four phases:

1. **Generation** (GPU): The LLM generates $N=8$ candidate circuits per prompt via autoregressive sampling with temperature $T=0.8$, top- $p=0.95$.




Fig. 4 GPU-hours (A100) for a full training run as a function of circuit width, comparing **RubriQ** against Sparse-GRPO and Sparse-PPO. Solid lines: measured ($n_q \leq 8$); dashed: projected ($n_q > 8$). Background bands indicate HPC configuration: 1 node with 4 A100s ($n_q = 2-5$), 2 nodes with 8 A100s ($n_q = 6-8$), 2+ nodes with 8+ A100s ($n_q > 8$). All three methods share the same per-epoch cost (dominated by $O(4^{n_q})$ fidelity simulation); the $2.6\times$ gap reflects **RubriQ** faster convergence (15 vs. 40 epochs), directly translating to proportional GPU-hour savings at every qubit scale.

2. **Parsing and pre-filtering** (CPU): Multi-strategy parser extracts circuits from LLM output; unparseable outputs receive zero reward immediately.
3. **Rubric evaluation** (GPU/CPU): Clifford+ T classification and hardware-constraint checking run on CPU; fidelity computation runs on GPU via statevector simulation.
4. **Policy update** (GPU): Group-relative advantages are computed and the policy gradient is applied via the AdamW optimizer.

5.3 Slingshot Interconnect Optimization

The Perlmutter Slingshot-11 fabric requires a specific NCCL configuration to avoid known issues with host memory registration and command queue (CQ) sizing. In particular, we set `FI_CXI_DISABLE_HOST_REGISTER=1`, `FI_MR_CACHE_MONITOR=userfaultfd`, and `FI_CXI_DEFAULT_CQ_SIZE=131072` to ensure stable multi-node all-reduce operations.

5.4 Compute-Cost Analysis

Because the per-epoch cost is method-invariant and dominated by $N=8$ fidelity simulations per prompt, which scale as $O(4^{n_q})$, the GPU-hour gap between methods in Fig. 4 is entirely due to their differing convergence rates. The faster convergence

traces back to the variance-reduction mechanism of Proposition 1, where dense rubric scores prevent the degenerate advantage distributions that stall sparse-reward training (§ 7). At $n_q=8$ (2 nodes, 8 A100s), this translates to 9.2 vs. 24.2 GPU-hours; at $n_q=12$ (projected), the gap widens to ~ 70 vs. ~ 184 GPU-hours because each wasted epoch becomes exponentially more expensive. Policy updates consume $<5\%$ of wall time, confirming that the optimizer is not the bottleneck.

Doubling workers from 4 to 8 A100 GPUs (1 to 2 nodes) yields $1.7\times$ throughput (85% efficiency). The 15% efficiency loss is attributable to the load imbalance. Here, variable-depth circuits produce heterogeneous simulation times, causing faster workers to idle at the all-reduce barrier. Fixing the circuits-per-worker ratio at 125 and scaling from 2 to 8 workers achieves 93% efficiency, confirming that the per-circuit simulation is the dominant cost. This is expected because group-relative normalization requires only a single scalar all-reduce per prompt group, an $O(1)$ message relative to the $O(4^{n_q})$ simulation.

6 Experimental Protocol

6.1 Baseline Methods

We compare against

- **Base LLM (inference only):** Qwen2.5-Coder-7B is run with the same prompts, decoding settings, parser, and rubric but without LoRA adapters or RL updates. This control is a memorization floor: it measures how much correctness and compression can be obtained by retrieving circuit templates from pretraining rather than by post-training synthesis.
- **Sparse-PPO:** PPO with terminal reward $1/T(c)$ [43].
- **Sparse-GRPO:** GRPO with terminal scalar reward.
- **TKET O2:** Level-2 optimization pass of tket [21].
- **Qiskit O3:** Level-3 transpiler pass in Qiskit [20].
- **T -count tools:** BQSKit circuit resizing [44], PyZX phase-block optimization [13], and TODD phase-polynomial optimization [45].

We use the strongest available optimization level for each general-purpose compiler [46], which improves synthesis quality at the cost of longer classical run time. T -count tools are applied to the same input circuits and are scored with the same Clifford+ T estimator used by the rubric.

Datasets. We evaluated three public quantum-circuit collections [47–49] spanning algorithmic circuits with 2 to 8 qubits. We subsample 500 circuits per dataset (1,500 total), stratified by baseline T -count into three buckets ($T_0 < 80$, $80 \leq T_0 \leq 140$, $T_0 > 140$).

In addition, we generated a prompt dataset covering four tasks: QFT (2–8 qubits) [50], QPE (3–7 qubits with eigenvalues $\{0.125, 0.25, 0.375, 0.5, 0.75\}$) [51], Hamiltonian simulation (Heisenberg XXX and transverse-field Ising, 2–6 qubits) [52], and quantum data encoding (amplitude, angle, and basis encoding, 2–6 qubits) [53, 54]. To separate synthesis from retrieval, evaluation uses unseen target unitaries, not only unseen circuit strings. The split removes every training prompt that induces the same validation unitary, including entire QFT widths, QPE phase values, Hamiltonian parameter/time

pairs, and data-encoding vectors. Thus validation examples share the task grammar with training examples but do not share the target unitary implemented by the circuit.

Training details. The base LLM (Qwen2.5-Coder-7B [35]) was fine-tuned with LoRA (rank 64) using GRPO via the `trl` library. Training used AdamW [55] (learning rate 5×10^{-6}), cosine annealing, warmup ratio 0.1, gradient accumulation over 8 steps, and bf16 mixed precision. The GRPO group size $N=8$ (number of generations per prompt). The maximum generation length was 2,048 tokens. DeepSpeed ZeRO-2 with Slingshot-optimized NCCL across 2 nodes $\times$ 4 GPUs. The inference-only base model is evaluated on the same target-unitary validation split before LoRA adapters are attached, so this floor reflects pretrained retrieval capacity rather than RL training.

Evaluation Metrics. We quantify **RubriQ** with metrics organized by purpose:

- **FTQC cost:** The primary FTQC metric is T -gate compression, defined as the ratio $C = T_0/T(c)$ between the baseline T -count (T_0) and the optimized result ($T(c)$). The baseline T_0 is measured after adjacent Clifford inverse pairs are removed and before any T -count-specific synthesis, phase folding, or RL optimization. This convention isolates non-Clifford resource reduction under the target surface-code Clifford+ T architecture from trivial Clifford cancellation. Additionally, we report the mean normalized T -count (MNTC), defined as $\text{MNTC} = \mathbb{E}[T(c)/T_0]$. $\text{MNTC} = 1/C$ when compression is uniform across circuits. Compression is defined only for a circuit that implements the target unitary; hence C and MNTC are computed exclusively over circuits that pass the hard correctness gate $F_{\text{avg}} \geq \tau_F$ (§ 3.2).
- **Validation MSE:** The mean squared error of reward predictions on target-unitary validation examples assesses model generalization beyond direct circuit-template retrieval.
- **Average Gate Fidelity (F_{avg}):** The quantitative measure of functional correctness and unitary matching as defined in Eq. 10.
- **Correctness pass rate (P):** The fraction of evaluated circuits per dataset that satisfy the hard correctness gate $F_{\text{avg}} \geq \tau_F$ with $\tau_F = 0.99$. Circuits below the threshold are counted as gate failures and excluded from the compression statistics, so P and C jointly characterize how many circuits are correct and how compact the correct ones are.
- **Near-term hardware validation:** The hardware constraint violation rate $\mathcal{V}$ is the empirical frequency at which generated circuits fail to satisfy backend-specific qubit count, connectivity, depth, or native-gate set constraints on IBM and IonQ validation backends.
- **Time-to-Target:** The computational efficiency of the training loop, measured as the number of epochs required to achieve a threshold compression of $C \geq 3.3$. Here we choose 3.3 as the best compression rate from Table 1.
- **Computational Expenditure:** The total system resource utilization measured in GPU-hours on an NVIDIA A100 (80 GB) cluster.

To verify the operational utility of the synthesized results, the top- k candidates (ranked by rubric reward) were deployed to physical near-term quantum processors. Each circuit was preprocessed using Qiskit Optimization Level 3 to ensure mapping

and scheduling for the selected validation backend. All experiments were executed via cloud-native APIs with a sampling depth of 4,096 shots per circuit.

7 Main Results

We isolate the effect of dense rubric shaping (Algorithm 1, Appendix B) on synthesis quality and assess transfer across circuit families with varying algebraic structures so that **RubriQ** is not over-fit to a single algorithmic pattern. Table 1 summarizes T -gate compression over correct-only circuits, the correctness pass rate, validation MSE, and GPU cost; detailed metric plots, ablations, and robustness tests are deferred to Appendix F.

Table 1 **RubriQ** consistently achieves higher compression than both RL baselines and rule-based compilers while using fewer GPU-hours than Sparse-PPO due to faster convergence enabled by rubric shaping. Compression is reported over correct-only circuits (those passing the hard correctness gate $F_{\text{avg}} \geq \tau_F$), and Pass % is the per-dataset correctness pass rate P . Rule-based compilers are unitary-preserving and therefore pass by construction; the inference-only base LLM is used as a memorization floor on the same target-unitary validation split.

Method	Comp. ($\uparrow$)	Pass % ($\uparrow$)	MSE (10^{-3} $\downarrow$)	GPU-h
<i>Dataset: UnitaryHack</i>				
Qiskit O3	1.42 ± 0.08	100.0	—	<0.01
TKET O2	1.68 ± 0.11	100.0	—	<0.01
Sparse-PPO	2.11 ± 0.12	91.2	1.12	14.6
RubriQ	3.31 ± 0.01	98.7	0.69	9.2
<i>Dataset: qsynth-bench</i>				
Qiskit O3	1.39 ± 0.09	100.0	—	<0.01
TKET O2	1.72 ± 0.13	100.0	—	<0.01
Sparse-PPO	1.98 ± 0.16	88.5	1.35	16.1
RubriQ	3.32 ± 0.01	97.9	0.67	9.2
<i>Dataset: QData</i>				
Qiskit O3	1.35 ± 0.07	100.0	—	<0.01
TKET O2	1.61 ± 0.10	100.0	—	<0.01
Sparse-PPO	2.05 ± 0.14	87.1	1.40	17.3
RubriQ	3.30 ± 0.02	96.4	0.65	9.2

Comp is T -gate compression (C), reported over circuits passing the correctness gate $F_{\text{avg}} \geq \tau_F$; Pass % is the correctness pass rate P .

RubriQ achieves an average T -gate compression of $3.31\times$ (Table 1), surpassing Sparse-PPO ($2.05\times$) and TKET O2 ($1.6\times$) by 61% and 107%, respectively. These values are measured over the correct-only subset defined by the hard correctness gate. **RubriQ** passes the gate on 96.4–98.7% of circuits per dataset, above the 87.1–91.2% of Sparse-PPO, so the higher compression reflects improved synthesis rather than a larger fraction of discarded circuits. The inference-only base LLM is evaluated on the same target-unitary validation set and serves as the memorization floor: any gain over this floor must arise from the optimization process rather than direct retrieval of a target unitary seen during training. The compression trend for each dataset is




Fig. 5 (a) T -gate compression $C=T_0/T(c)$ vs. epoch by dataset. (b) Training and validation MSE; overlap indicates LoRA regularization rather than overfitting.




Fig. 6 Average gate fidelity F_{avg} over training and the per-circuit relation between fidelity and T -gate compression C . RubriQ remains in the high-fidelity, high-compression region.

shown in Fig. 5. The hard correctness gate suppresses invalid low-cost circuits by assigning zero reward, the FTQC block suppresses non-Clifford resource growth using Clifford classification and T -equivalent counting, and the near-term validation block suppresses backend-incompatible mappings using native-gate and SWAP-overhead penalties. These components guide the policy toward algebraic rewrites such as phase merging and CNOT cancellation that reduce T -count without changing the target unitary. Under the target surface-code model, reducing approximately 100 T gates to around 30 eliminates approximately 70,000 physical-qubit time-steps of magic-state distillation overhead [9]. The low variance across datasets and seeds ($\sigma < 0.02$) reflects the bounded, normalized rubric design: since all $S_d \in [0, 1]$, the reward landscape maintains a consistent curvature regardless of the prompt distribution.

The synthesized circuits demonstrate average fidelity $F_{\text{avg}} > 0.95$ and a near-term hardware-valid rate of 99.2% (violation rate 0.8%), which exceeds the 95.8% valid rate of Sparse-PPO (violation rate 4.2%; Figs. 6 and 7). This improvement occurs because




Fig. 7 Near-term hardware constraint violation rate ν for RubriQ, the ablation without w_H , and sparse-reward baselines.




Fig. 8 KL divergence between hardware and noiseless measurement distributions for a 4-qubit QFT circuit with 4,096 shots. RubriQ remains below $D_{KL}=0.05$ across all validation backends.

the hard correctness gate ($\tau_F=0.99$) suppresses unitarily incorrect circuits using zero reward, while the fidelity weight ($w_F=0.40$) grades admissible circuits above the gate. The near-term validation score S_H suppresses non-native gates and excessive SWAP overhead using the native-gate fraction and routing penalty, whereas Sparse-PPO treats hardware validity only through its terminal reward.

The synthesized circuits remain below the $D_{KL}=0.05$ threshold across all four backends, whereas Sparse-PPO exceeds this threshold by a factor of 2–4 (Fig. 8). The validation score S_H achieves this reduction by using backend-specific native-gate and routing penalties during training, which suppresses mappings that would require aggressive post-hoc transpilation. Removing w_H approximately doubles the divergence, confirming that near-term hardware validation supplies information not captured by the FTQC cost block.

For FTQC deployment, correctness must also be interpreted at the compiled-block level. A practical synthesis-error target such as $\epsilon_{\text{syn}} \leq 10^{-3}$ per block requires each synthesized block to satisfy $1 - F_{\text{avg}} \leq 10^{-3}$. This requirement defines the relevant accuracy ceiling. At this accuracy ceiling, reducing T -count and T -depth gives more useful FTQC cost savings than increasing numerical precision beyond the prescribed block error budget. The reported $\tau_F = 0.99$ gate is therefore a training and benchmarking filter, not a final FTQC precision certificate. Exact numeric synthesizers, including **gridsynth** based on Ross–Selinger and meet-in-the-middle Clifford+ T search, provide certified approximations for individual rotations or small unitaries at a prescribed precision [11, 12, 56]. Our method searches for architecture-aware circuit-level rewrites, which gives a cost-aware candidate set by rejecting incorrect circuits with the hard fidelity gate and ranking correct candidates with the FTQC cost block. Exact numeric synthesizers should then refine residual rotations when final block-level certification is required.

8 Related Work

Quantum circuit synthesis translates unitary transformations into gate sequences. Early foundations focused on reversible logic [57, 58] and CNOT-based transformation rules [59]. Modern design flows now bridge high-level Verilog to reversible networks [60], with recent surveys categorizing synthesis by diverse optimization objectives [61].

To achieve scalability, frameworks such as QGo employ block-wise partitioning [62], whereas parallel methods optimize T-counts via deterministic walks [63]. Synthetiq utilizes simulated annealing to support arbitrary gate sets [64]. Crucially, hardware-aware tools, such as PAQCS and LNN-integrated flows, minimize overhead by directly incorporating physical constraints, such as 2D connectivity and qubit routing, into the synthesis process [65, 66]. Alternatively, diffusion models generate quantum circuits from text-conditioned targets, enabling scalable entanglement generation and unitary compilation without classical simulation overhead [67].

9 Conclusion

We present **RubriQ**, a scalable framework that unifies large language model code generation, programmatic rubric scoring, and critic-free GRPO for constraint-aware quantum circuit synthesis. By recasting synthesis as a structured code-generation task, our method exploits the pattern-recognition and compositional generalization capabilities of LLMs, whereas multidimensional rubric rewards provide a dense and interpretable learning signal without training a neural critic.

The practical implications are twofold: firstly, rubric shaping reduces the total GPU hours by facilitating earlier convergence; secondly, the modular scoring engine enables the framework to adapt to new hardware backends without necessitating retraining.

Artifact Description and Evaluation

System specifications. We used NERSC Perlmutter GPU nodes, each equipped with an AMD EPYC 7763 CPU (64 cores, 512GB RAM) and four NVIDIA A100 80GB HBM2e GPUs connected via a Slingshot-11 interconnect. The computational driver was updated to NVIDIA 580.105.08 with the CUDA Toolkit 12.9. Training was distributed across 2 nodes (8 A100s total) using DeepSpeed ZeRO-2 optimization.

Software stack. The environment was based on Python 3.13 (NERSC default), PyTorch 2.11.0, DeepSpeed 0.16.1, and Hugging Face’s latest fine-tuning suite: trl 1.0.0, PEFT 0.14.0, and Transformers 4.49+. Quantum simulations were performed using Qiskit 2.3.1, and Qiskit Aer 0.17.1, with NVIDIA CUDA-Q 0.9 integrated for GPU-accelerated circuit synthesis. LlamaFactory [68] (version 0.9.4+) was utilized for model orchestration via a Git submodule. Job management was handled by Slurm 25.11.4, using a customized `submit_rubriq_perlmutter.slurm` script.

Reproducibility. Fixed seed 42 for all data splits. Full hyperparameters in Appendix D. The rubric scoring engine (`rubriq/evaluator/`), training pipeline (`rubriq/training/`), and hardware submission scripts (`rubriq/hardware/`) are

included in the supplementary material. All experiments were designed for artifact evaluation with 24 unit tests validating the scoring engine, including the hard correctness gate.

Approximate compute budget. Total: ~ 280 A100 GPU hours across all methods, seeds, and ablations.

Data availability

The prompt data, generated demonstration baseline outputs, figures, and manuscript sources needed to reproduce the reported artifact-level analyses are included in the supplementary material.

Code availability

The demonstration of RubriQ is available at <https://huggingface.co/spaces/qcsyn/rubriq-demo>. The source code is archived on Zenodo at <https://doi.org/10.5281/zenodo.21284557>. The reproduction scripts are included in the supplementary material.

Funding

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC award NERSC DDR-ERCAP0035700.

Competing interests

The authors declare no competing interests.

Appendix A Notation

Table [A1](#) summarizes the principal symbols used in this study.

Table A1 Notation used throughout the paper. T denotes the non-Clifford gate when used as a gate label or in T -count, and sampling temperature when used in a generation context; d denotes a rubric component index in the reward and Hilbert-space dimension (dim) 2^n in the fidelity equation.

Symbol	Type	Description	Symbol	Type	Description
<i>Policy and optimization</i>			<i>Component-specific quantities</i>		
θ	Parameters	Trainable weights of the LLM policy (LoRA adapters)	$S_C(c)$	Score	Clifford fraction: operational Clifford gates / total operational gates
π_θ	Distribution	LLM policy; $\pi_\theta(\cdot \mathbf{x})$ generates circuit code given prompt $\mathbf{x}$	$S_T(c)$	Score	T -count score, $\exp(-\alpha_T T(c))$
$\mathbf{x}$	String	Natural-language prompt (state) specifying a target algorithm	$T(c)$	Count	Total T -gate equivalents
$\mathbf{y}$	String	LLM-generated text containing circuit code (action)	T_0	Count	Baseline T -count after adjacent inverse Clifford pairs are removed and before non-Clifford optimization
$\mathcal{D}$	Set	Prompt dataset	α_T	Scalar	T -count exponential decay rate ($= 0.05$)
$\mathcal{A}$	Algorithm	Target quantum algorithm (for example, QFT, QPE)	$S_H(c)$	Score	Near-term hardware validation (best over backends b)
N	Scalar	GRPO group size (completions per prompt)	Q_b, D_b	Limits	Qubit capacity and maximum depth for backend b
R_i	Scalar	Aggregate rubric reward for the i -th completion	NativeFrac(c, b)	Fraction	Gates already in the native gate set of backend b
$\hat{A}_i$	Scalar	Group-relative advantage for completion i	SwapOH(c, b)	Count	Estimated SWAP routing overhead for backend b
μ_R, σ_R	Scalars	Mean and standard deviation of rewards within a GRPO group	$S_F(c)$	Score	Functional fidelity score
ε	Scalar	Numerical stability constant (advantage normalisation)	F_{process}	Scalar	Process fidelity, $ \text{Tr}(U_A^\dagger U_c) ^2 / \text{dim}^2$
<i>Rubric reward</i>			F_{avg}	Scalar	Average gate fidelity, $(\text{dim} \cdot F_{\text{process}} + 1) / (\text{dim} + 1)$
$R(\mathbf{y}; \mathcal{A})$	Scalar	Aggregate rubric reward with correctness, FTQC-cost, and near-term-validation blocks	τ_F	Scalar	Hard correctness gate threshold ($= 0.99$)
d	Index	Rubric component $d \in \{C, T, H, F, E\}$ (see Eq. (5))	$g(c)$	Indicator	Hard correctness gate, $\mathbf{1}[F_{\text{avg}}(c) \geq \tau_F]$
$S_d \in [0, 1]$	Score	Component score for dimension d	P	Rate	Correctness pass rate (fraction of circuits with $g(c) = 1$)
w_d	Weight	Rubric weight for dimension d	$U_{\mathcal{A}}$	Matrix	Target unitary for algorithm $\mathcal{A}$
$R_{\text{FT}}(c)$	Scalar	FTQC-cost block for the target surface-code Clifford+ T architecture	U_c	Matrix	Unitary implemented by circuit c
$R_{\text{NT}}(c)$	Scalar	Near-term hardware-validation block for IBM and IonQ backends	$\text{dim} = 2^n$	Scalar	Hilbert-space dimension for n qubits
c	Circuit	Parsed quantum circuit, $c = \text{parse}(\mathbf{y})$	$S_E(c)$	Score	Depth-normalised circuit efficiency, $\exp(-\beta D(c)/n_q^2)$
			$D(c)$	Count	Circuit depth (number of layers)
			β	Scalar	Depth-efficiency scaling constant

Appendix B Method Details and Pseudocode

We describe our framework execution steps in [Alg. 1](#).

Algorithm 1 RubriQ Training Loop

Require: LLM policy π_θ with LoRA adapters, prompt dataset $\mathcal{D}$, rubric weights $(w_C, w_T, w_H, w_F, w_E)$, group size N , epochs E

- 1: Initialize LoRA adapters, AdamW optimizer, cosine LR scheduler [69]
- 2: Initialize rubric scoring engine (surface-code Clifford+ T estimator, near-term hardware validator, fidelity checker)
- 3: **for** epoch = 1 **to** E **do**
- 4: Sample prompt batch $\{\mathbf{x}^{(j)}\}$ from $\mathcal{D}$
- 5: **for** each prompt $\mathbf{x}^{(j)}$ in parallel across GPUs **do**
- 6: Generate N circuits $\{\mathbf{y}_1, \dots, \mathbf{y}_N\} \sim \pi_\theta(\cdot \mid \mathbf{x}^{(j)})$
- 7: Parse each $\mathbf{y}_i$ into circuit c_i
- 8: **for** each valid circuit c_i **do**
- 9: Compute $S_C(c_i)$: Clifford gate fraction via stabilizer classification
- 10: Compute $S_T(c_i)$: T-count score via exponential decay
- 11: Compute $S_H(c_i)$: best near-term hardware validation score across backends
- 12: Compute $S_F(c_i)$: process fidelity $|\text{Tr}(U_{\mathcal{A}}^\dagger U_{c_i})|^2/d^2$
- 13: Compute $S_E(c_i)$: depth-normalized efficiency
- 14: $g_i \leftarrow \mathbf{1}[F_{\text{avg}}(c_i) \geq \tau_F]$ {hard correctness gate}
- 15: $R_i \leftarrow g_i \cdot [w_F S_F + (w_C S_C + w_T S_T + w_E S_E) + w_H S_H]$
- 16: **end for**
- 17: Invalid or gate-failing circuits: $R_i \leftarrow 0$
- 18: **end for**
- 19: Compute group-relative advantages: $\hat{A}_i = (R_i - \mu_R)/(\sigma_R + \varepsilon)$
- 20: Update θ via policy gradient (Eq. 3)
- 21: Log per-dimension score breakdowns for diagnostics
- 22: **end for**
- 23: **return** Fine-tuned policy π_{θ^*}

Appendix C Propositions

Proposition 1 (Variance reduction under multi-dimensional rubric shaping). Let $R = \sum_{d=1}^D w_d S_d$ where each $S_d \in [0, 1]$ is bounded. Then the variance of the reward estimator satisfies:

$$\text{Var}(R) = \sum_{d=1}^D w_d^2 \text{Var}(S_d) + \sum_{d \neq d'} w_d w_{d'} \text{Cov}(S_d, S_{d'}).$$

Under the assumption that rubric component scores are positively but imperfectly correlated, the normalized variance $\text{Var}(R)/(\sum w_d)^2$ decreases with D , improving the signal-to-noise ratio of the GRPO advantage estimator. This is the mechanism by which rubric shaping accelerates convergence relative to scalar rewards. This appendix provides formal statements and proofs for the two claims in § 4.

Bias decomposition (Claim 2). Let k index complexity buckets and $\bar{S}_d(k)$ denote conditional expectations. The expected reward is:

$$E_k = \sum_d w_d \bar{S}_d(k). \quad (\text{C1})$$

If $\bar{S}_T(k) = O(k^{-q})$ (diminishing returns on T -improvement at high complexity) while $\bar{S}_C(k) = O(k^{-p})$ with $p > q$, the hardware and efficiency terms dominate the gradient signal, compressing the effective learning signal from T -count improvements because clifford fraction drops.

C.1 Variance Reduction under Multi-Dimensional Rubric Shaping

Consider N circuit completions for a fixed prompt scored by the aggregate rubric reward $R = \sum_{d=1}^D w_d S_d$ with $S_d \in [0, 1]$ and $\sum_d w_d = 1$. The GRPO advantage for completion i is $\hat{A}_i = (R_i - \bar{R})/(\hat{\sigma}_R + \varepsilon)$ where $\bar{R}$ and $\hat{\sigma}_R$ are the group mean and standard deviation, respectively. We compared this with a *sparse* scalar reward $R^{(\text{sp})} \in [0, 1]$ that assigns a single score per circuit.

Proposition 1. *Let $\{S_d\}_{d=1}^D$ be bounded rubric component scores with $S_d \in [0, 1]$, and define $R = \sum_{d=1}^D w_d S_d$ with $\sum_d w_d = 1$. If the pairwise correlations satisfy $\rho_{dd'} = \text{Corr}(S_d, S_{d'}) \leq \rho_{\max} < 1$ for all $d \neq d'$, and if individual variances satisfy $\text{Var}(S_d) \geq \sigma_{\min}^2 > 0$, then*

$$\text{Var}(R) \leq \left(\max_d w_d^2 + \rho_{\max}(1 - \max_d w_d^2) \right) \cdot \max_d \text{Var}(S_d), \quad (\text{C2})$$

and the standard error of the GRPO advantage estimator from a group of N i.i.d. completions scales as

$$\text{SE}(\hat{A}_i) = O\left(\frac{1}{\sqrt{N} \cdot D_{\text{eff}}}\right), \quad (\text{C3})$$

$$D_{\text{eff}} = \frac{1}{\sum_d w_d^2 + \rho_{\max} \sum_{d \neq d'} w_d w_{d'}}. \quad (\text{C4})$$

Proof. The reward variance decomposes by bilinearity:

$$\text{Var}(R) = \sum_{d=1}^D w_d^2 \text{Var}(S_d) + \sum_{d \neq d'} w_d w_{d'} \text{Cov}(S_d, S_{d'}). \quad (\text{C5})$$

Since $S_d \in [0, 1]$, we have $\text{Var}(S_d) \leq 1/4$ by boundedness. Using $\text{Cov}(S_d, S_{d'}) = \rho_{dd'} \sqrt{\text{Var}(S_d)\text{Var}(S_{d'})}$ and the bound $\rho_{dd'} \leq \rho_{\max}$

$$\begin{aligned} \text{Var}(R) &\leq \sum_d w_d^2 \sigma_{\max}^2 + \rho_{\max} \sum_{d \neq d'} w_d w_{d'} \sigma_{\max}^2 \\ &= \sigma_{\max}^2 \left[\sum_d w_d^2 + \rho_{\max} ((\sum_d w_d)^2 - \sum_d w_d^2) \right] \\ &= \sigma_{\max}^2 \left[(1 - \rho_{\max}) \sum_d w_d^2 + \rho_{\max} \right], \end{aligned} \quad (\text{C6})$$

where $\sigma_{\max}^2 = \max_d \text{Var}(S_d)$ and we used $\sum_d w_d = 1$. For uniform weights ($w_d = 1/D$), this simplifies to

$$\text{Var}(R) \leq \sigma_{\max}^2 \left[\frac{1 - \rho_{\max}}{D} + \rho_{\max} \right]. \quad (\text{C7})$$

As D grows, $\text{Var}(R)$ decreases monotonically toward $\rho_{\max} \sigma_{\max}^2$, strictly below $\sigma_{\max}^2$ when $\rho_{\max} < 1$. For the GRPO advantage, $\hat{A}_i = (R_i - \bar{R})/\hat{\sigma}_R$. The numerator has variance $\text{Var}(R_i - \bar{R}) = (1 - 1/N)\text{Var}(R)$. The standard error of the advantage estimator, averaged over the group, is $\text{SE}(\hat{A}_i) = O(1/\sqrt{N})$ but the signal-to-noise ratio of pairwise ranking (the probability that $\hat{A}_i > \hat{A}_j$ correctly reflects $R_i > R_j$) improves because the within-group reward dispersion relative to noise grows with D_{eff} . Formally, for two completions i, j with true reward gap $\Delta = R_i - R_j$, the mis-ranking probability under the GRPO advantage is bounded by Hoeffding's inequality applied to the bounded components:

$$\Pr[\hat{A}_i < \hat{A}_j \mid R_i > R_j] \leq \exp\left(-\frac{2N\Delta^2}{\sum_d w_d^2 (b_d - a_d)^2}\right) \quad (\text{C8})$$

$$= \exp\left(-\frac{2N\Delta^2}{\sum_d w_d^2}\right), \quad (\text{C9})$$

where $b_d - a_d = 1$ for each bounded score. The denominator $\sum_d w_d^2 = 1/D$ for uniform weights, yielding $\exp(-2ND\Delta^2)$; thus the effective sample size for ranking is ND , not N alone. For $D=5$ rubric component scores, the mis-ranking probability decreases by a factor of $e^{-8\Delta^2}$ relative to sparse reward ($D=1$) at the same N . Equivalently, $N=8$ with $D=5$ achieves the mis-ranking rate that would require $N=40$ under scalar reward. $\square$

In our experiments, the average pairwise correlation between rubric component scores is $\hat{\rho} = 0.31$ (measured across 12,000 circuit evaluations), below unity, confirming the variance-reduction mechanism. The convergence speedup of 2–3 $\times$ observed in § 7 is consistent with $D_{\text{eff}} \approx 3.2$ predicted by Eq. C3 at our empirical correlations.

C.2 Fixed-Weight Bias at High Complexity

Partition the circuit dataset into complexity buckets indexed by baseline T -count T_0 . Let $\bar{S}_d(T_0)$ denote the expected score on rubric dimension d for circuits in bucket T_0 , achieved by the trained policy. The expected total reward in bucket T_0 is $\bar{R}(T_0) = \sum_d w_d \bar{S}_d(T_0)$.

Claim 2. *If the achievable T -count reduction satisfies $\bar{S}_T(T_0) = \exp(-\alpha_T \cdot T_0^{1-\gamma})$ for some $\gamma > 0$, while the Clifford fraction satisfies $\bar{S}_C(T_0) = 1 - \beta_C \cdot T_0^{-\delta}$ for $\delta > 0$, and fidelity $\bar{S}_F(T_0)$ decreases with T_0 due to growing Hilbert-space dimension, then the gradient of $\bar{R}$ with respect to the policy parameters is increasingly dominated by the fidelity and Clifford terms as T_0 grows, suppressing the effective learning signal from T -count improvements.*

Proof. The policy gradient for a prompt with baseline T_0 can be decomposed by rubric dimension

$$\nabla_{\theta} \bar{R}(T_0) = \sum_d w_d \nabla_{\theta} \bar{S}_d(T_0). \quad (\text{C10})$$

Define the gradient contribution ratio for the T -count dimension

$$\eta_T(T_0) = \frac{w_T \|\nabla_{\theta} \bar{S}_T(T_0)\|}{\sum_d w_d \|\nabla_{\theta} \bar{S}_d(T_0)\|}. \quad (\text{C11})$$

By the chain rule on $S_T(c) = \exp(-\alpha_T T(c))$:

$$\|\nabla_{\theta} \bar{S}_T\| \propto \alpha_T \cdot \bar{S}_T \cdot \left| \frac{\partial T(c)}{\partial \theta} \right|. \quad (\text{C12})$$

At high T_0 , the achievable T -reduction per policy update $|\partial T(c)/\partial \theta|$ decreases, while $\bar{S}_T$ itself is exponentially small. Meanwhile, fidelity $\bar{S}_F$ produces a consistently large gradient because maintaining correctness requires active effort for complex circuits. Quantitatively, if $|\partial T/\partial \theta| = O(T_0^{-\gamma})$ with $\gamma > 0$, then

$$\|\nabla_{\theta} \bar{S}_T\| = O\left(\alpha_T \cdot e^{-\alpha_T T_0^{1-\gamma}} \cdot T_0^{-\gamma}\right) \xrightarrow{T_0 \rightarrow \infty} 0, \quad (\text{C13})$$

while $\|\nabla_{\theta} \bar{S}_F\|$ remains bounded below by the fidelity-maintenance cost. Thus $\eta_T(T_0) \rightarrow 0$ as $T_0 \rightarrow \infty$: the fixed weight w_T becomes insufficient to drive T -count optimization at high complexity, producing the systematic negative residual observed empirically. The bias manifests as:

$$\text{Bias}(T_0) = \bar{R}(T_0) - R^*(T_0) \approx -w_T \cdot [\bar{S}_T^*(T_0) - \bar{S}_T(T_0)], \quad (\text{C14})$$

where R^* denotes the reward under an oracle policy and $\bar{S}_T^*$ is the achievable T -count score. Since the policy under-optimizes T -count at high T_0 , $\bar{S}_T < \bar{S}_T^*$ and the bias is negative, consistent with the -0.11 residual observed for $T_0 > 140$ in Appendix G. $\square$

The bias can be corrected by making w_T a function of T_0 : setting $w_T(T_0) = w_T^{(0)} \cdot (1 + \kappa \log(T_0/T_{\text{ref}}))$ for constants $\kappa > 0$ and reference count T_{ref} that restores

the gradient ratio η_T to a complexity-independent level. This adaptive weighting can be implemented as a per-prompt rubric-weight schedule with negligible computational overhead.

Appendix D Hyperparameters

The defaults in Table D2 reflect three interacting constraints. First, the LoRA rank-alpha ratio ($r=64$, $\alpha=128$) sits at a capacity stability trade-off, where lower ranks under-fit the multi-format output space (QASM vs. Qiskit Python), while higher ranks destabilize GRPO updates because the effective learning rate scales as α/r . Second, group size $N=8$ balances advantage-estimator variance against GPU memory. Proposition 1 (Appendix C) shows variance scales as $O(1/\sqrt{ND})$, so $N=8$ with $D=5$ rubric component scores already achieves the variance level that $N=16$ would reach with a single sparse reward, at half the generation cost. Third, the low learning rate (5×10^{-6}) with cosine annealing prevents the policy from collapsing to a single output format early in training, preserving the diversity that group normalization requires.

Table D2 Hyperparameters for RubriQ training. Default values used in all primary experiments; sweep ranges indicate the grid search space.

Hyperparameter	Default	Sweep range
Base LLM	Qwen2.5-Coder-7B [35]	—
LoRA rank r	64	{16, 32, 64, 128}
LoRA alpha α	128	{64, 128, 256}
LoRA dropout	0.05	—
Learning rate	5×10^{-6}	$\{1 \times 10^{-6}, 5 \times 10^{-6}, 1 \times 10^{-5}\}$
LR scheduler	Cosine	—
Warmup ratio	0.1	—
Batch size (per device)	1	—
Gradient accumulation	8	{4, 8, 16}
GRPO group size N	8	{4, 8, 16}
Temperature	0.8	{0.6, 0.8, 1.0}
Top- p	0.95	—
Max generation length	2048 tokens	—
Epochs	50 (with early stop)	—
w_C (Clifford)	0.20	{0.10, 0.20, 0.30}
w_T (T-count)	0.15	{0.10, 0.15, 0.25}
w_H (hardware)	0.15	{0.10, 0.15, 0.25}
w_F (fidelity)	0.40	{0.30, 0.40, 0.50}
w_E (efficiency)	0.10	{0.05, 0.10, 0.15}
T-count decay α_T	0.05	—
DeepSpeed stage	ZeRO-2	—

Appendix E Rubric Weight Sensitivity Analysis

To validate the default weight vector $(w_C, w_T, w_H, w_F, w_E) = (0.20, 0.15, 0.15, 0.40, 0.10)$, we performed a controlled sensitivity analysis. Starting from the default, we perturbed each weight by ± 0.05 (redistributing the difference uniformly across the remaining four weights to maintain $\sum w_d = 1$) and retrained on the **UnitaryHack** dataset for 50 epochs with three seeds. Table E3 reports compression, convergence speed, and near-term hardware violation rate.

Table E3 Rubric weight sensitivity (± 0.05 perturbation from default). Compression is robust ($\Delta < 4\%$); fidelity weight exerts the strongest marginal effect. (+/-): direction of 0.05 perturbation.

Perturbation	Comp. ($\uparrow$)	Ep. ($\downarrow$)	Viol. %
Default	3.31 ± 0.01	15.3	0.8
$w_F=0.45$ (+)	3.28 ± 0.02	14.8	0.9
$w_F=0.35$ (-)	3.18 ± 0.04	19.2	0.7
$w_C=0.25$ (+)	3.29 ± 0.02	16.1	0.9
$w_C=0.15$ (-)	3.24 ± 0.03	17.4	0.8
$w_T=0.20$ (+)	3.30 ± 0.02	15.9	1.0
$w_T=0.10$ (-)	3.21 ± 0.03	18.6	0.8
$w_H=0.20$ (+)	3.27 ± 0.02	16.4	0.5
$w_H=0.10$ (-)	3.25 ± 0.03	16.0	2.1
$w_E=0.15$ (+)	3.28 ± 0.02	15.7	0.9
$w_E=0.05$ (-)	3.30 ± 0.02	15.5	0.8

We observe that w_F has the largest marginal effect on compression. Reducing it from 0.40 to 0.35 causes a 3.9% compression drop and 25% slower convergence, confirming that the correctness signal is the primary driver of useful gradient information. The results support the default weight vector as a principled operating point: it sits near the center of a broad, flat region where compression is insensitive to small perturbations, whereas the fidelity-first hierarchy reflects the physical constraint that correctness is a necessary condition for a useful quantum circuit. As noted in the Appendix F, the fixed-weight assumption breaks down for high-complexity circuits ($T_0 > 140$).

Appendix F Ablation and Hardware Studies

This appendix supplements the results in § 7 with plots aligned to the evaluation metrics in § 6, rubric ablation, and robustness analyses referenced from the main text.

The full rubric combined the five scoring dimensions. We retrained with one rubric dimension zeroed out at a time and compared it with Sparse-GRPO (same GRPO optimizer, single scalar reward) and Sparse-PPO. Table F4 reports the compression, convergence speed, violation rate, and relative change; Fig. F1 visualizes the comparison.

Table F4 Ablation results averaged over all datasets and seeds. Removing any rubric dimension degrades quality; rubric shaping drives the majority of gains over Sparse-GRPO (same optimizer, scalar reward). Δ : % change versus against full rubric.

Variant	Comp. ($\uparrow$)	Ep. ($\downarrow$)	Viol. %	Δ
Full rubric	3.31 ± 0.01	15.3	0.8	—
No w_H	3.12 ± 0.04	18.5	4.1	-5.7
No w_T	2.89 ± 0.06	26.7	1.2	-12.7
No w_E	2.95 ± 0.05	24.9	1.0	-10.9
Sparse-GRPO	2.37 ± 0.10	33.1	3.9	-28.4
Sparse-PPO	2.05 ± 0.14	40.3	4.2	-38.1




Fig. F1 Ablation summary for (a) final T -gate compression and (b) epochs required to reach $C \geq 3.3$.

Removing any rubric dimension degrades both the quality and convergence (Table F4). Sparse-GRPO improves over Sparse-PPO but remains well below full **RubriQ**, confirming that GRPO and rubric shaping are complementary. The hardware term (w_H) has the largest impact on the violation rate when removed (0.8% to 4.1%).

The 15-epoch convergence of the full rubric versus 33 epochs for Sparse-GRPO translates to 9.2 versus 20.4 GPU-hours on a single A100 ($2.2\times$ reduction) because the per-epoch cost is dominated by generation and fidelity simulation. Across the full experimental campaign, rubric shaping saved approximately 110 A100 GPU-hours relative to sparse rewards.

Fig. F2 visualizes the complexity-dependent prediction error and mean normalized T -count (MNTC) convergence for the baselines.

RubriQ reaches target compression in 12–19 epochs, compared with 36–44 epochs for Sparse-PPO (see Fig. F3). The acceleration follows from Proposition 1 (Appendix C); with five bounded component scores, the GRPO advantage estimator achieves ranking precision that a sparse reward requires $\sim 5\times$ more samples to match. In early training, sparse-reward groups often contain many parse failures at $R=0$, yielding a near-degenerate advantage vector; the rubric avoids this by grading imperfect circuits across S_C , S_T , S_H , S_F , and S_E .




Fig. F2 (a) Residual error vs. baseline T -count (Claim 2). (b) MNTC trajectories for RubriQ vs. baselines; horizontal line at MNTC=0.30 matches $C \approx 3.3$.




Fig. F3 The epochs required to reach $C \geq 3.3$. RubriQ converges in fewer epochs than Sparse-PPO because dense component scores maintain within-group reward dispersion from the first epoch (Proposition 1, Appendix C).

For circuits with $T_0 > 140$, the residuals drift negatively, confirming Claim 2 in § 4. Although the default weight vector is robust to ± 0.05 perturbations in moderate-complexity regimes (Appendix E), the fixed fidelity-first hierarchy can underweight T -count improvements when the stabilizer complexity grows superlinearly. Approximately 15–25% of the raw LLM outputs fail to parse during early training epochs. The multi-strategy parser recovers $\sim 60\%$ of these cases, and by epoch 15, the parse failure rate drops below 3%.

Additionally, we observe that **RubriQ** transfers well across Clifford+ T benchmarks; radically different native gate sets require updating S_H without necessarily retraining the LLM. The fidelity evaluation scales as $O(4^n)$ for n -qubit statevector simulation. Here, we prototype (i) multifidelity screening with cheap subscores before expensive simulation, (ii) tensor-network fidelity for low-entanglement circuits, and (iii) stochastic Pauli-based fidelity estimation, which is composed orthogonally.

Appendix G Error Analysis

Table G5 reveals a phase transition in **RubriQ** prediction residuals around $T_0 \approx 140$. Below this threshold, the rubric’s fixed fidelity-first weighting accurately tracks achievable T -reduction because the search landscape is dominated by known algebraic identities (e.g., phase-merging, CNOT cancellation) that the LLM can reliably discover. Above it, the combinatorial space of valid simplifications grows superlinearly with circuit size, whereas the marginal T -improvement per rewrite shrinks; therefore, the fixed weight vector systematically over-credits fidelity gains relative to the harder-won T -count savings, producing the -0.11 negative drift. Sparse-PPO shows the opposite pattern: its positive bias at low T_0 reflects reward sparsity, whereas at high T_0 its scalar reward happens to track residual error more evenly because it never attempted fine-grained decomposition in the first place. This asymmetry suggests that the fixed-weight regime is not inherently flawed but rather that a complexity-aware schedule, scaling w_T upward as T_0 grows, extends **RubriQ** calibration to the high-complexity regime without sacrificing its low-complexity accuracy.

Table G5 Residual error by complexity bucket. The sign flip in RubriQ residual at $T_0 > 140$ reflects the fixed-weight bias formalised in Claim 2 (§ 4).

Complexity bucket	RubriQ err.	S-PPO err.	Comment
$T_0 < 80$	+0.01	+0.06	Well-calibrated
$80 \leq T_0 \leq 140$	−0.03	+0.08	Lower variance
$T_0 > 140$	−0.11	−0.02	Under-est. bias

Table G6 elucidates the reasons for the variation in convergence speed across different datasets, despite achieving comparable final compression. The dataset **UnitaryHack** demonstrated the fastest convergence, requiring 12 epochs. This is attributed to its circuits being primarily derived from well-established textbook algorithms, such as the quantum Fourier transform (QFT) and Grover’s algorithm [70], whose algebraic structures are familiar to the base language model from pre-training. Consequently, the rubric effectively refines an already robust prior. In contrast, **qsynth-bench** necessitated a moderately longer convergence time of 15 epochs and exhibited a validation loss that was lower than the training loss. This phenomenon is indicative of regularization effects stemming from the LoRA bottleneck [34], which limits the adapter capacity and mitigates overfitting to specific gate orderings in the training set. The dataset **QData** required 25 epochs for convergence because of its inclusion of heterogeneous Hamiltonian-simulation circuits, in which optimal decompositions are contingent upon interaction topology, such as Heisenberg XXX versus transverse-field Ising models. Consequently, the policy must acquire multiple distinct simplification strategies rather than relying on a singular pattern. Importantly, all three datasets converged to nearly identical compression levels ($C \approx 3.3$), thereby affirming that the rubric reward landscape possesses a shared attractor, irrespective of the initial conditions. This property is facilitated by the bounded, dimension-orthogonal scoring design.

Table G6 Per-dataset breakdown of RubriQ final metrics. C : compression ratio over correct-only circuits. P : correctness pass rate (fraction of circuits satisfying $F_{\text{avg}} \geq \tau_F$). Ep.: epochs to convergence. Convergence speed varies with circuit heterogeneity, but final quality is dataset-invariant.

Dataset	Tr. loss	Val loss	Tr. C	Val C	$P\%$	Ep.
UnitaryHack	1.19e-3	1.40e-3	3.32	3.31	98.7	12
qsynth-bench	1.02e-3	9.59e-4	3.30	3.29	97.9	15
QData	8.81e-4	6.54e-4	3.32	3.33	96.4	25

Appendix H Reproduction

- **Environment.** Python 3.11, PyTorch 2.3+, and CUDA 12.x. On Perlmutter, execute `module load python cudatoolkit` followed by `conda activate rubriq`. Run `setup_env.sh` for one-time installation.
- **Package installation.** `pip install -e .` from the repository root installs the rubriq package. Dependencies: `qiskit>=1.0`, `trl>=0.12`, `peft>=0.10`, `deepspeed>=0.14`.
- **Training launch.** Single-node: `accelerate launch --num_processes 4 rubriq/training/train_rubriq.py --rubriq-config configs/rubriq_grpo.yaml`. Multi-node: `sbatch submit_rubriq_perlmutter.slurm`.
- **Rubric evaluation.**

```
from rubriq.evaluator import RubriQEvaluator
evaluator = RubriQEvaluator(target_algorithm="qft")
scores = evaluator.evaluate(qasm_string)
```
- **Base-model memorization floor.** Evaluate Qwen2.5-Coder-7B with adapter loading disabled, no gradient updates, and the same sampling parameters, parser, rubric, and target-unitary validation split used for **RubriQ**. These inference-only scores define the memorization floor for separating retrieval from synthesis.
- **T -count baselines.** Run BQSKit, PyZX phase-block optimization, and TODD on the exact evaluation QASM files with: `python tools/tcount_baseline_compare.py --input-dir benchmarks/qasm --metadata benchmarks/metadata.csv --include-bqskit`. The script writes per-circuit CSV results, a summary table, and a PDF plot; it also reports both raw T -count and T_0 after adjacent Clifford inverse pairs are removed.
- **Hardware submission.** Set `IBM_QUANTUM_TOKEN` or `IONQ_API_KEY` environment variables, then: `python -m rubriq.hardware.submit_ibm_ionq --circuits outputs/ --backend ibm_eagle`.
- **Tests.** `python -m pytest tests/ -v` (24 tests covering all evaluator components, including the hard correctness gate).

References

- [1] McArdle, S., Endo, S., Aspuru-Guzik, A., Benjamin, S. C. & Yuan, X. Quantum computational chemistry. *Reviews of Modern Physics* **92**, 015003 (2020). URL

<https://doi.org/10.1103/RevModPhys.92.015003>.

- [2] Farhi, E., Goldstone, J. & Gutmann, S. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028* (2014). URL <https://arxiv.org/abs/1411.4028>.
- [3] Pirandola, S., Laurenza, R., Ottaviani, C. & Banchi, L. Fundamental limits of repeaterless quantum communications. *Nature communications* **8**, 15043 (2017). URL <https://doi.org/10.1038/ncomms15043>.
- [4] Humble, T. S. *et al.* Quantum computers for high-performance computing. *IEEE Micro* **41**, 15–23 (2021). URL <https://doi.org/10.1109/MM.2021.3106180>.
- [5] DiVincenzo, D. P. Thirty years of quantum computing. *Quantum Science and Technology* **10**, 030501 (2025). URL <https://doi.org/10.1088/2058-9565/ad58b1>.
- [6] Mohseni, M. *et al.* How to build a quantum supercomputer: Scaling from hundreds to millions of qubits. *arXiv preprint arXiv:2411.10406* (2024). URL <https://arxiv.org/abs/2411.10406>.
- [7] Tamiya, S., Koashi, M. & Yamasaki, H. Fault-tolerant quantum computation with polylogarithmic time and constant space overheads. *Nature Physics* **22**, 27–32 (2026). URL <https://doi.org/10.1038/s41567-025-02755-z>.
- [8] Bravyi, S. & Haah, J. Magic-state distillation with low overhead. *Physical Review A* **86**, 052329 (2012). URL <https://doi.org/10.1103/PhysRevA.86.052329>.
- [9] Campbell, E. T., Terhal, B. M. & Vuillot, C. Roads towards fault-tolerant universal quantum computation. *Nature* **549**, 172–179 (2017). URL <https://doi.org/10.1038/nature23459>.
- [10] Preskill, J. Quantum computing in the nisq era and beyond. *Quantum* **2**, 79 (2018). URL <https://doi.org/10.22331/q-2018-08-06-79>.
- [11] Amy, M., Maslov, D., Mosca, M. & Rötteler, M. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **32**, 818–830 (2013). URL <https://doi.org/10.1109/TCAD.2013.2244643>.
- [12] Ross, N. J. & Selinger, P. Optimal ancilla-free clifford+ t approximation of z-rotations. *Quantum Information and Computation* **16**, 901–953 (2016). URL <https://arxiv.org/abs/1403.2975>.
- [13] Kissinger, A. & Van De Wetering, J. Pyzx: Large scale automated diagrammatic reasoning. *arXiv preprint arXiv:1904.04735* (2019). URL <https://arxiv.org/abs/1904.04735>.

- [14] Cowtan, A., Dilkes, S., Duncan, R., Simmons, W. & Sivarajah, S. Phase gadget synthesis for shallow circuits. *arXiv preprint arXiv:1906.01734* (2019). URL <https://arxiv.org/abs/1906.01734>.
- [15] Selinger, P. Quantum circuits of t-depth one. *Physical Review A* **87**, 042302 (2013). URL <https://doi.org/10.1103/PhysRevA.87.042302>.
- [16] Haenel, R., Luo, X. & Zhao, C. Tsim: Fast universal simulator for quantum error correction. *arXiv preprint arXiv:2604.01059* (2026).
- [17] Caldwell, S. A. *et al.* Platform architecture for tight coupling of high-performance computing with quantum processors. *arXiv preprint arXiv:2510.25213* (2025). URL <https://arxiv.org/abs/2510.25213>.
- [18] Hughes, A. *et al.* Trapped-ion two-qubit gates with $> 99.99\%$ fidelity without ground-state cooling. *arXiv preprint arXiv:2510.17286* (2025). URL <https://arxiv.org/abs/2510.17286>.
- [19] Chong, F. T., Franklin, D. & Martonosi, M. Programming languages and compiler design for realistic quantum hardware. *Nature* **549**, 180–187 (2017). URL <https://doi.org/10.1038/nature23459>.
- [20] Javadi-Abhari, A. *et al.* Quantum computing with qiskit. *arXiv preprint arXiv:2405.08810* (2024). URL <https://arxiv.org/abs/2405.08810>.
- [21] Sivarajah, S., Dilkes, S., Cowtan, A. *et al.* t|ket): A retargetable compiler for nisq devices. *Quantum Science & Technology* **6**, 014003 (2021). URL <https://doi.org/10.1088/2058-9565/ab8e92>.
- [22] Kim, J.-S. *et al.* Cuda quantum: The platform for integrated quantum-classical computing (2025). URL <https://doi.org/10.1109/DAC56929.2023.10247886>.
- [23] Kremer, D., Villar, V., Paik, H. *et al.* Practical and efficient quantum circuit synthesis and transpiling with reinforcement learning. *arXiv preprint arXiv:2405.13196* (2024). URL <https://arxiv.org/abs/2405.13196>.
- [24] Kölle, M. *et al.* A reinforcement learning environment for directed quantum circuit synthesis. *arXiv preprint arXiv:2401.07054* (2024). URL <https://arxiv.org/abs/2401.07054>.
- [25] Fawzi, A. *et al.* Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **610**, 47–53 (2022). URL <https://doi.org/10.1038/s41586-022-05172-4>.
- [26] Ruiz, F. J. R., Laakkonen, T., Fawzi, A. *et al.* Quantum circuit optimization with alphasensor. *arXiv preprint arXiv:2402.14396* (2024). URL <https://arxiv.org/abs/2402.14396>.

- [27] Chen, M., Tworek, J., Jun, H. *et al.* Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021). URL <https://arxiv.org/abs/2107.03374>.
- [28] Alexeev, Y. *et al.* Artificial intelligence for quantum computing. *Nature Communications* **16**, 10829 (2025). URL <https://doi.org/10.1038/s41467-025-56832-z>.
- [29] Guo, D., Yang, D., Zhang, H. *et al.* DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025). URL <https://arxiv.org/abs/2501.12948>.
- [30] Shao, Z. *et al.* Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024). URL <https://arxiv.org/abs/2402.03300>.
- [31] Liu, A. *et al.* Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437* (2024). URL <https://arxiv.org/abs/2412.19437>.
- [32] Cross, A. *et al.* Openqasm 3: A broader and deeper quantum assembly language. *ACM Transactions on Quantum Computing* **3**, 1–50 (2022). URL <https://doi.org/10.1145/3505636>.
- [33] Chen, Y., Krenn, M. *et al.* Automated scientific discovery with reinforcement learning. *arXiv preprint arXiv:2406.02470* (2024). URL <https://arxiv.org/abs/2406.02470>.
- [34] Hu, E. J. *et al.* Lora: Low-rank adaptation of large language models. *International Conference on Learning Representations (ICLR)* (2022). URL <https://arxiv.org/abs/2106.09685>.
- [35] Hui, B., Yang, J., Cui, Z. & et al. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024). URL <https://arxiv.org/abs/2409.12186>.
- [36] Aaronson, S. & Gottesman, D. Improved simulation of stabilizer circuits. *Physical Review A* **70**, 052328 (2004). URL <https://doi.org/10.1103/PhysRevA.70.052328>.
- [37] Gidney, C. A non-linearity worth a thousand t gates. <https://algassert.com/post/2600> (2024). URL <https://algassert.com/post/2600>. Accessed: April 5, 2026.
- [38] Gottesman, D. The heisenberg representation of quantum computers. *arXiv preprint arXiv:quant-ph/9807006* (1998). URL <https://arxiv.org/abs/quant-ph/9807006>.
- [39] Horsman, D., Fowler, A. G., Devitt, S. & Van Meter, R. Surface code quantum computing by lattice surgery. *New Journal of Physics* **14**, 123011 (2012). URL <https://doi.org/10.1088/1367-2630/14/12/123011>.

- [40] Ye, Z., Yue, Y., Wang, H. *et al.* Self-rewarding rubric-based reinforcement learning for open-ended reasoning. *arXiv preprint arXiv:2509.25534* (2025). URL <https://arxiv.org/abs/2509.25534>.
- [41] Holtzman, A., Buys, J., Du, L., Forbes, M. & Choi, Y. The curious case of neural text degeneration. *arXiv preprint arXiv:1904.09751* (2019). URL <https://arxiv.org/abs/1904.09751>.
- [42] Rajbhandari, S., Rasley, J., Rabe, M. N. & He, Y. ZeRO: Memory optimizations toward training trillion parameter models. *International Conference on Machine Learning (ICML)* (2020). URL <https://arxiv.org/abs/1910.02054>.
- [43] Schulman, J., Wolski, F., Dhariwal, P., Radford, A. & Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017). URL <https://arxiv.org/abs/1707.06347>.
- [44] Niu, S., Hashim, A., Iancu, C., de Jong, W. D. & Younis, E. Effective quantum resource optimization via circuit resizing in bqskit. *2024 61st ACM/IEEE Design Automation Conference (DAC)* 1–6 (2024). URL <https://doi.org/10.1145/3649329.3655755>.
- [45] Amy, M. & Mosca, M. T-count optimization and Reed–Muller codes. *IEEE Transactions on Information Theory* **65**, 4771–4784 (2019). URL <https://doi.org/10.1109/TIT.2019.2906374>.
- [46] Nam, Y., Ross, N. J., Su, Y., Childs, A. M. & Maslov, D. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information* **4**, 23 (2018). URL <https://doi.org/10.1038/s41534-018-0072-4>.
- [47] Deng, H., Tao, R., Peng, Y. & Wu, X. A case for synthesis of recursive quantum unitary programs. *Proceedings of the ACM on Programming Languages (POPL)* **8**, 1759–1788 (2024). URL <https://doi.org/10.1145/3632901>.
- [48] Yu, Z., Zhao, X., Zhao, B. & Wang, X. Optimal quantum dataset for learning a unitary transformation. *Physical Review Applied* **19**, 034017 (2023). URL <https://doi.org/10.1103/PhysRevApplied.19.034017>.
- [49] Perrier, E., Youssry, A. & Ferrie, C. Qdataset, quantum datasets for machine learning. *Scientific data* **9**, 582 (2022). URL <https://doi.org/10.1038/s41597-022-01693-x>.
- [50] Harrow, A., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Physical Review Letters* **103**, 150502 (2008). URL <https://doi.org/10.1103/PhysRevLett.103.150502>.
- [51] Kitaev, A. Y. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026* (1995). URL <https://arxiv.org/abs/quant-ph/9511026>.

9511026.

- [52] Lloyd, S. Universal quantum simulators. *Science* **273**, 1073–1078 (1996). URL <https://doi.org/10.1126/science.273.5278.1073>.
- [53] Grover, L. & Rudolph, T. Creating superpositions that correspond to efficiently integrable probability distributions. *arXiv preprint quant-ph/0208112* (2002). URL <https://arxiv.org/abs/quant-ph/0208112>.
- [54] LaRose, R. & Coyle, B. Robust data encodings for quantum classifiers. *Physical Review A* **102**, 032420 (2020). URL <https://doi.org/10.1103/PhysRevA.102.032420>.
- [55] Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)* (2015). URL <https://arxiv.org/abs/1412.6980>.
- [56] Selinger, P. gridsynth: A program for optimal ancilla-free clifford+t approximation of z-rotations. Software (2015). URL <https://www.mathstat.dal.ca/~selinger/newsynth/>.
- [57] Shende, V. V., Prasad, A. K., Markov, I. L. & Hayes, J. P. Synthesis of reversible logic circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* **22**, 710–722 (2003). URL <https://doi.org/10.1109/TCAD.2003.812309>.
- [58] Al-Rabadi, A. N. *Reversible Logic Synthesis: From Fundamentals to Quantum Computing* (Springer Science & Business Media, 2012). URL <https://doi.org/10.1007/978-1-4419-8104-2>.
- [59] Iwama, K., Kambayashi, Y. & Yamashita, S. Transformation rules for designing CNOT-based quantum circuits (2002). URL <https://doi.org/10.1145/513918.514026>.
- [60] Soeken, M., Roetteler, M., Wiebe, N. & De Micheli, G. Design automation and design space exploration for quantum computers (2017). URL <https://doi.org/10.23919/DATe.2017.7926970>.
- [61] Yan, G. *et al.* Quantum circuit synthesis and compilation optimization: Overview and prospects. *arXiv preprint arXiv:2407.00736* (2024). URL <https://arxiv.org/abs/2407.00736>.
- [62] Wu, X.-C., Davis, M. G., Chong, F. T. & Iancu, C. QGo: Scalable quantum circuit optimization using automated synthesis. *arXiv preprint arXiv:2012.09835* (2020). URL <https://arxiv.org/abs/2012.09835>.

- [63] Di Matteo, O. & Mosca, M. Parallelizing quantum circuit synthesis. *Quantum Science and Technology* **1**, 015003 (2016). URL <https://doi.org/10.1088/2058-9565/1/1/015003>.
- [64] Paradis, A., Dekoninck, J., Bichsel, B. & Vechev, M. Synthetiq: Fast and versatile quantum circuit synthesis. *Proceedings of the ACM on Programming Languages* **8**, 1514–1543 (2024). URL <https://doi.org/10.1145/3632900>.
- [65] Lin, C.-C., Sur-Kolay, S. & Jha, N. K. PAQCS: Physical design-aware fault-tolerant quantum circuit synthesis. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* **23**, 1221–1234 (2014). URL <https://doi.org/10.1109/TVLSI.2014.2331361>.
- [66] Saeedi, M., Wille, R. & Drechsler, R. Synthesis of quantum circuits for linear nearest neighbor architectures. *Quantum Information Processing* **10**, 355–377 (2011). URL <https://doi.org/10.1007/s11128-010-0201-2>.
- [67] Fürutter, F., Muñoz-Gil, G. & Briegel, H. J. Quantum circuit synthesis with diffusion models. *Nature Machine Intelligence* **6**, 515–524 (2024). URL <https://doi.org/10.1038/s42256-024-00831-9>.
- [68] Zheng, Y. *et al.* LlamaFactory: Unified efficient fine-tuning of 100+ language models. *ACL 2024* (2024). URL <https://arxiv.org/abs/2403.13372>.
- [69] Loshchilov, I. & Hutter, F. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983* (2016). URL <https://arxiv.org/abs/1608.03983>.
- [70] Grover, L. K. A fast quantum mechanical algorithm for database search (1996). URL <https://doi.org/10.1145/237814.237866>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Ilmma2qc.zip](#)