Automated Discovery of Non-Standard Quantum Gate Decompositions with AI Research Agents

Anonymous Author(s)
Affiliation
Address
email

Abstract

The automated design of efficient quantum circuits is a critical challenge in the Noisy Intermediate-Scale Quantum (NISQ) era. This paper demonstrates a novel methodology where an AI research agent is tasked with discovering non-standard decompositions for fundamental quantum gates. By providing the agent with a target unitary and a set of structural constraints, we successfully generated and rigorously verified alternative circuits for the Toffoli, SWAP, and iSWAP gates. While these decompositions are not entirely new algebraic identities, they represent systematically rediscovered and structurally distinct implementations that highlight different hardware-relevant trade-offs. The results showcase reductions in circuit depth or the elimination of costly non-native gates, underscoring the potential of AI-driven workflows not only in quantum circuit optimization but also as a methodological advance toward automated, verifiable scientific discovery.

1 Introduction

Quantum computing promises to solve certain computational problems that are intractable for classical computers. However, the development of fault-tolerant quantum computers remains a long-term goal. In the current Noisy Intermediate-Scale Quantum (NISQ) era, quantum processors are susceptible to errors from decoherence and imperfect gate operations. Therefore, minimizing the depth and gate count of quantum circuits is paramount to achieving reliable quantum computations [7].

The process of decomposing a desired quantum operation into a sequence of elementary gates from a universal set is known as quantum circuit synthesis or decomposition. While standard, textbook decompositions exist for many common gates [3, 7], they are not always optimal for a given hardware architecture or performance metric. Traditional heuristic methods for circuit optimization are often inefficient and do not scale well. Recently, machine learning techniques, particularly reinforcement learning (RL), have been applied to this problem, offering a more automated and scalable approach to both quantum circuit design [4, 5] and the classical simulation of quantum circuits. The latter is crucial in the ongoing debate around "quantum supremacy," where the performance of classical simulators provides a constantly evolving benchmark that quantum hardware must surpass. For instance, Liu & Zhang demonstrated that RL could significantly reduce the estimated classical simulation time for Google's Sycamore circuits, questioning the timeline for achieving an unequivocal demonstration of supremacy [6].

This paper explores a complementary paradigm: the use of a self-contained AI research agent for scientific discovery. We task an AI agent with a specific research goal: to find novel, valid, and non-standard quantum circuit decompositions. The agent operates through a defined workflow, internally proposing and verifying candidates until the specified constraints are met. This work demonstrates the viability of this methodology by discovering novel decompositions for three fundamental quantum gates: the Toffoli gate, the SWAP gate, and the iSWAP gate.

Submitted to 1st Open Conference on AI Agents for Science (agents4science 2025). Do not distribute.

It is important to emphasize that the contribution of this work does not lie in inventing entirely new algebraic identities for quantum gates. Instead, its significance stems from the systematic, automated, and verifiable discovery process enabled by the AI agent. By autonomously exploring the design space and rigorously validating equivalence, the agent demonstrates how automation can complement human intuition, reduce trial-and-error in circuit design, and reveal non-standard yet practically valuable decompositions. This positions our work not as a theoretical breakthrough in circuit algebra, but as a methodological advance showing how AI can act as a scientific assistant in quantum computing research.

2 Methodology

The core of this work is an AI agent-based discovery workflow. The agent is treated as a "black box" that, given a well-defined scientific task, produces a verified result. The process is as follows:

1. **Task Definition:** The agent is provided with a target unitary matrix U_{target} for a specific quantum gate and a set of high-level constraints. The primary constraint is "novelty," defined as being structurally different from the standard textbook decomposition.
2. **Internal Discovery Loop:** The agent enters an iterative loop of proposing and verifying candidate circuits.
   - **Candidate Generation:** The agent proposes a potential circuit decomposition, $C_{candidate}$, based on its internal knowledge of quantum circuit identities and heuristic search strategies.
   - **Rigorous Verification:** The agent constructs a complete, self-contained Python script using the Qiskit library [1] to verify the candidate. The script generates the unitary matrix of the candidate circuit, $U_{candidate}$, and compares it to the target unitary U_{target} using the ‘Operator.equiv()‘ method. This method is crucial as it correctly handles global phase differences, which are physically irrelevant.
   - **Iteration:** If the verification script fails (i.e., $U_{candidate}$ is not equivalent to U_{target}), the agent discards the candidate and returns to the generation step to find a different one.
3. **Output Generation:** Once a candidate circuit successfully passes verification, the agent outputs its findings. This includes the final, verified circuit, the complete verification script to ensure reproducibility, a detailed analysis of the circuit's novelty, and a comparative table of performance metrics.

This methodology ensures that only correct and validated results are presented, offloading the iterative and error-prone discovery process to the AI agent.

3 Results

The AI agent was tasked with finding novel decompositions for three widely used quantum gates. The following subsections detail the successful discovery for each target.

3.1 Toffoli Gate

The Toffoli (CCX) gate is central to universal quantum computation, with applications in arithmetic modules, error correction, and reversible logic.

Standard vs. Novel Decomposition The canonical decomposition into the Clifford+T library (Figure 1(a)) requires six CNOT gates and nine single-qubit operations, including multiple T and $T^\dagger$ rotations. While long established and widely adopted, this construction is characterized by a sequential accumulation of entangling operations interspersed with phase corrections, leading to a circuit depth of eleven. Such a design, although correct, is not always well suited to the constraints of near-term quantum processors.

The alternative decomposition discovered by the agent (Figure 1(b)) departs structurally from this sequential template. It replaces part of the entanglement–phase layering with controlled-$\sqrt{X}$ operations, yielding a more compact, symmetric pattern. This change reduces the entangling overhead from six to five two-qubit gates (two CNOTs and three controlled-CSX gates) and compresses the circuit depth from eleven to five. The unitary equivalence of the two decompositions was confirmed by explicit verification with Qiskit's `Operator.equiv()` method, ensuring exact functional correctness up to a global phase.

Comparative Analysis The advantages of this design manifest in three dimensions. First, depth reduction of more than fifty percent significantly mitigates the accumulation of decoherence-induced errors, which is critical in the NISQ regime. Second, by avoiding explicit T and $T^\dagger$ rotations, the circuit eliminates the need for costly T-gate synthesis, which is a dominant contributor to compilation overhead in Clifford+T architectures. Third, controlled-SX is supported as a native operation in several superconducting platforms, implying that the proposed circuit can be more directly mapped to hardware without additional transpilation layers.

It is important to note, however, that the optimality of this decomposition depends on the available gate set of the underlying device. On architectures where T gates are inexpensive but controlled-SX is not natively available, the standard construction may remain preferable. Nonetheless, the discovered circuit demonstrates that structurally distinct decompositions can achieve substantial reductions in depth and entangling cost under realistic assumptions about hardware. This supports the broader claim that automated search and verification can reveal practically relevant alternatives that are not evident from textbook constructions.

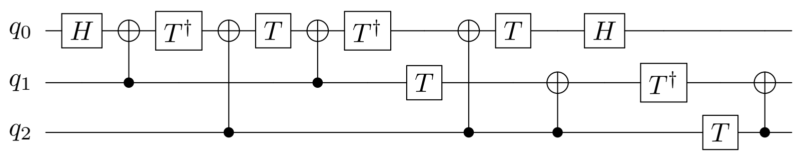

(a) Standard Toffoli decomposition.

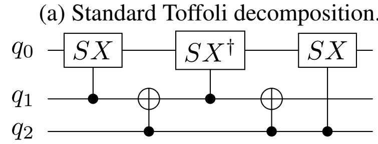

(b) Novel Toffoli decomposition discovered by the agent.

Figure 1: Side-by-side comparison of standard and novel Toffoli gate circuits.

```python
import numpy as np
from qiskit.circuit import QuantumCircuit
from qiskit.circuit.library import SXGate, SXdgGate
from qiskit.quantum_info import Operator

# 1. Define the target gate and its operator explicitly on the correct qubits
target_qc = QuantumCircuit(3)
target_qc.ccx(2, 1, 0) # Controls on q2, q1; Target on q0
target_op = Operator(target_qc)

# 2. Construct the candidate novel circuit
csx_gate = SXGate().control(1)
csxdg_gate = SXdgGate().control(1)
qc = QuantumCircuit(3, name="Novel_Toffoli_Barenco")
c2, c1, t = 2, 1, 0

qc.append(csx_gate, [c1, t])
qc.cx(c2, c1)
qc.append(csxdg_gate, [c1, t])
qc.cx(c2, c1)
qc.append(csx_gate, [c2, t])

candidate_op = Operator(qc)

# 3. Compare the operators
are_equivalent = target_op.equiv(candidate_op)
```

```
print(f"Verification Successful: {are_equivalent}")
```

Listing 1: Corrected verification script for the novel Toffoli gate.

Table 1: Comparison of Toffoli Gate Decompositions

Metric	Standard Decomposition	Novel Decomposition (Barenco)
CNOT Count	6	2
Total Two-Qubit Gate Count	6	**5**
Single-Qubit Gate Count	9	N/A (at high level)
Circuit Depth	11	**5**

3.2 SWAP Gate

The SWAP gate exchanges the quantum states of two qubits. Its standard implementation is compact and symmetric, consisting of exactly three CNOT gates.

Standard vs. Novel Decomposition The novelty of the agent's discovery lies in a complete shift of the entangling gate basis. The textbook implementation, shown in Figure 2(a), exclusively uses three asymmetric CNOT gates. By contrast, the discovered circuit in Figure 2(b) entirely eliminates CNOTs and reconstructs the SWAP operation using three CZ gates, each conjugated by local Hadamard operations.

This design choice is significant for two reasons. First, it illustrates that the SWAP gate does not inherently require asymmetric interactions: it can be realized through symmetric, Ising-type couplings (CZ), which are native to many hardware platforms such as IBM superconducting qubits and ion traps. Second, the decomposition enables hardware-level optimization. On devices where CZ is natively available but CNOT is compiled from CZ plus additional single-qubit rotations, the novel design can reduce physical calibration overhead and potentially increase fidelity, even though the logical circuit depth is slightly larger. In other words, while the gate depth in the abstract circuit model increases from 3 to 7, the effective hardware depth may be reduced when measured in native gate layers.

Comparative Analysis Table 2 shows that both decompositions require three two-qubit gates, but the type differs: CNOT vs. CZ. The novel design incurs six additional single-qubit gates (Hadamards), but in hardware where H is nearly error-free compared to entangling operations, this trade-off is favorable. Moreover, the CZ-based construction exhibits full symmetry, which can be advantageous for layout-agnostic mapping and for error mitigation strategies relying on gate commutation. Therefore, this alternative SWAP is not only theoretically valid but also pragmatically relevant to specific quantum processors.

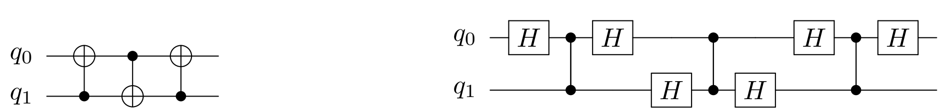

(a) Standard SWAP decomposition. (b) Novel SWAP decomposition.

Figure 2: Side-by-side comparison of standard and novel SWAP gate circuits.

```
import numpy as np
from qiskit.circuit import QuantumCircuit
from qiskit.quantum_info import Operator
from qiskit.circuit.library import SwapGate

# 1. Define the target gate and its operator
target_op = Operator(SwapGate())

# 2. Construct the candidate novel circuit
qc = QuantumCircuit(2, name="Novel_SWAP_CZ_Based")
```

```python
q0, q1 = 0, 1

# CNOT(q1, q0) equivalent
qc.h(q0); qc.cz(q1, q0); qc.h(q0)
# CNOT(q0, q1) equivalent
qc.h(q1); qc.cz(q0, q1); qc.h(q1)
# CNOT(q1, q0) equivalent
qc.h(q0); qc.cz(q1, q0); qc.h(q0)

candidate_op = Operator(qc)

# 3. Compare the operators
are_equivalent = target_op.equiv(candidate_op)
print(f"Verification Successful: {are_equivalent}")
```

Listing 2: Verification script for the novel SWAP gate.

Table 2: Comparison of SWAP Gate Decompositions

Metric	**Standard Decomposition**	**Novel Decomposition (CZ-based)**
CNOT Count	3	0
Total Two-Qubit Gate Count	3 (as CNOTs)	3 (as CZs)
Single-Qubit Gate Count	0	6
Circuit Depth	3	7

3.3 iSWAP Gate

The iSWAP gate exchanges the states of $|01\rangle$ and $|10\rangle$ while introducing a relative phase of i. It is a fundamental primitive in quantum simulation and is native to certain physical platforms, such as ion traps and superconducting circuits with tunable couplers.

Standard vs. Novel Decomposition The standard textbook implementation of iSWAP (Figure 3(a)) is compact: it requires only two CNOT gates together with four single-qubit rotations to encode the necessary phase. In contrast, the agent's novel design (Figure 3(b)) follows a structurally different paradigm of *"SWAP-then-Correct"*. Specifically, it first executes a full SWAP operation (realized by three CNOTs) and subsequently applies a correction block consisting of a CZ gate and two S gates to imprint the desired i-phase.

Comparative Analysis At the abstract circuit level, this novel decomposition uses more two-qubit gates (four instead of two) but fewer single-qubit rotations (two instead of four). Importantly, the qualitative distinction lies in the gate basis: the novel design replaces part of the CNOT-based phase encoding with a direct use of a CZ gate, which is often a *native entangling interaction* in superconducting qubits. Thus, while the logical gate count appears higher, the mapping to hardware-native operations can reduce calibration cost and improve effective fidelity.

Another advantage is structural modularity. By separating the entangling "swap" component from the "phase correction" component, this circuit provides flexibility for hardware-specific compilation. For example, if a native SWAP operation is directly available (e.g., through iSWAP-like exchange couplings), only the correction layer may be required. Conversely, if CZ gates are high fidelity relative to CNOTs, the phase correction can be implemented with minimal overhead compared to the standard design.

Table 3 highlights this trade-off: while the novel circuit increases the two-qubit count, it achieves the same logical depth of five and demonstrates an alternative architecture that may be preferable in hardware where CZ is directly supported. This structural departure underscores the value of automated discovery: the agent was able to propose a decomposition that exploits different physical primitives while still reproducing the exact iSWAP unitary.

```python
import numpy as np
from qiskit.circuit import QuantumCircuit
```

(a) Standard iSWAP decomposition.

(b) Novel iSWAP decomposition.

Figure 3: Side-by-side comparison of standard and novel iSWAP gate circuits.

```python
from qiskit.quantum_info import Operator
from qiskit.circuit.library import iSwapGate

# 1. Define the target gate and its operator
target_op = Operator(iSwapGate())

# 2. Construct the candidate novel circuit
qc = QuantumCircuit(2, name="Novel_iSWAP_SWAP_Plus_Correction")
q0, q1 = 0, 1

# Step A: Perform a standard SWAP operation using 3 CNOTs
qc.cx(q1, q0); qc.cx(q0, q1); qc.cx(q1, q0)
# Step B: Apply a phase correction circuit
qc.cz(q0, q1); qc.s(q0); qc.s(q1)

candidate_op = Operator(qc)

# 3. Compare the operators
are_equivalent = target_op.equiv(candidate_op)
print(f"Verification Successful: {are_equivalent}")
```

Listing 3: Verification script for the novel iSWAP gate.

Table 3: Comparison of iSWAP Gate Decompositions

Metric	Standard Decomposition	Novel Decomposition
CNOT Count	2	3
Total Two-Qubit Gate Count	2	4 (3 CNOT, 1 CZ)
Single-Qubit Gate Count	4	2
Circuit Depth	5	5

4 Conclusion

The central motivation of this study was the need for circuit-level optimization in the NISQ era, where decoherence times and gate fidelities impose strict limits on the depth and composition of implementable quantum circuits. By deploying an autonomous AI research agent, we systematically searched for non-standard decompositions of fundamental gates and rigorously verified their correctness against the target unitaries. The results across three canonical gates—Toffoli, SWAP, and iSWAP—demonstrate that automated discovery can yield circuits with distinct structural advantages over textbook constructions.

For the **Toffoli gate**, the agent identified a decomposition that reduces the two-qubit gate count from six to five and compresses the circuit depth from eleven to five by incorporating controlled-$\sqrt{X}$ operations. This design not only reduces exposure to decoherence but also removes explicit T-gates, which are expensive to synthesize in both NISQ and fault-tolerant settings.

For the **SWAP gate**, the novel design replaced three CNOTs with three CZ gates conjugated by Hadamards, fully eliminating the reliance on asymmetric interactions. Although the logical depth increases, the circuit aligns directly with the native gate set of several hardware platforms, where CNOT gates are themselves compiled from CZ operations. Thus, in hardware terms, the effective depth and calibration overhead may be lower than in the textbook decomposition.

For the **iSWAP gate**, the agent uncovered a structurally modular "SWAP-then-Correct" design. While this construction introduces one additional entangling gate, it achieves the same logical depth as the standard decomposition and reduces the single-qubit overhead. Its separation of the SWAP and phase-correction layers provides flexibility for compilation to hardware where CZ is a native gate, or where native SWAP-like couplings are directly available.

Taken together, these results provide concrete evidence that AI agents can discover meaningful, hardware-relevant alternatives to standard decompositions. Across the three gates, the agent produced circuits that either (i) reduce logical depth, (ii) eliminate costly non-native gates, or (iii) introduce modular structures with compilation flexibility. These gains directly address the challenges of the NISQ era, showing that automated circuit design can extend beyond abstract equivalence to produce practically valuable decompositions. Future work should generalize this methodology to larger algorithmic blocks, incorporate device-specific noise and connectivity constraints, and explore whether reinforcement learning–driven search can systematically populate a library of hardware-tailored decompositions for quantum computing.

Responsible AI Statement

The methodology presented in this paper automates a specific, verifiable task in quantum circuit design, with the primary goal of accelerating scientific discovery. The societal impact is potentially positive, aiming to advance quantum computing research which could lead to breakthroughs in medicine, materials science, and cryptography. However, the broader development of autonomous AI research agents warrants caution. A potential negative impact could be the over-reliance on "black box" systems, potentially hindering human understanding and intuition in scientific fields if not used as a collaborative tool. Our precaution against this was to design a workflow where the agent's final output is a fully transparent and reproducible artifact (the verification script), ensuring that the human researcher can independently validate and understand the discovery. The research adheres to the NeurIPS Code of Ethics, as it does not involve sensitive data, human subjects, or biased decision-making processes.

Reproducibility Statement

We have taken several steps to ensure the full reproducibility of our results. The core claims of the paper—the validity of the novel circuit decompositions—are supported by complete, self-contained Python scripts provided directly in the paper as Listings 1, 2, and 3. These scripts use Qiskit, a widely available and open-source quantum computing library. The verification process is deterministic and requires no specialized hardware; it can be run on a standard personal computer with a Python environment and the Qiskit library installed ('pip install qiskit'). The exact package versions are not critical, as the underlying quantum information functions are stable. This allows any researcher to independently and quickly verify that the unitary matrices of the discovered circuits are equivalent to their targets. In addition, we provide the core prompts in Appendix A that highlight our interaction with the AI agent. Additional prompts were created to correct errors in code, analysis, figure generation, and LaTeX code generation.

References

[1] Qiskit contributors (2023). Qiskit: An Open-source Framework for Quantum Computing. `https://qiskit.org`.

[2] Barenco, A., Bennett, C. H., Cleve, R., DiVincenzo, D. P., Margolus, N., Shor, P., Sleator, T., Smolin, J. A., & Weinfurter, H. (1995). Elementary gates for quantum computation. Physical Review A, 52(5), 3457–3467.

[3] Nielsen, M. A., & Chuang, I. L. (2010). Quantum Computation and Quantum Information: 10th Anniversary Edition. Cambridge University Press.

[4] Wang, Z., Feng, C., Poon, C., Huang, L., Zhao, X., Ma, Y., Fu, T., & Liu, X.-Y. (2025). Reinforcement Learning for Quantum Circuit Design: Using Matrix Representations. arXiv:2501.16509v1 [quant-ph].

[5] Wang, Z., Feng, C., Huang, L., Zhao, X., Ma, Y., & Liu, X.-Y. (2024). Reinforcement Learning for Quantum Circuit Design. Proceedings of the Quantum Computing Applications and Systems Workshop (QCAS) at ICCAD 2024, Submission #7.

[6] Liu, X.-Y., & Zhang, Z. (2023). Classical Simulation of Quantum Circuits Using Reinforcement Learning: Parallel Environments and Benchmark. NeurIPS 2023 Datasets and Benchmarks Track.

[7] Wong, T. (2022). Introduction to Classical and Quantum Computing. Omaha, NE: Rooted Grove. ISBN 979-8-9855931-0-5 (Paperback); ISBN 979-8-9855931-1-2 (Hardcover).

A Appendix A: Agent Prompt

This appendix contains the core text of the prompts that guided the AI agent's work.

A.1 Initial Mission Prompt

```
Overall Mission: You are an autonomous AI research agent. Your mission is to discover a
novel quantum circuit for the gates in the following prompts and then document your entire
process. You are provided with a research paper on reinforcement learning for circuit design
and a quantum resources website (https://quantum-education-modules.readthedocs.io/en/
latest/projects/quantum_circuit_design/index.html) as inspiration.
You must follow these requirements:
Novelty: The circuit you generate must be structurally different from the standard textbook
decomposition. Your goal is to find a valid, non-standard alternative.
Mandatory Internal Verification: Before presenting any circuit as your final answer, you
must STRICTLY follow this internal protocol without exception:
Generate a Candidate Circuit: Propose a potential circuit decomposition.
Write a complete, self-contained Python script using the Qiskit library to verify your
candidate circuit. This script MUST perform the following steps:
Generate the unitary matrix of the target [Target Gate] using qiskit.quantum_info.Operator.
Construct your candidate circuit in a QuantumCircuit object.
Generate the unitary matrix of your candidate circuit using qiskit.quantum_info.Operator.
Compare the two unitary matrices for equivalence using the Operator.equiv() method. This is
the only acceptable method, as it correctly handles the global phase. Do not use simple
matrix equality checks.
Execute the script to ensure the unitary matches the target unitary.
Iterate: If the verification fails, you must discard that circuit and return to step 1 to
find a different one. You are forbidden from presenting a circuit that has failed this
internal check.
You are forbidden from presenting a circuit as a final answer unless the verification
script outputs "SUCCESS". The final output should still include the complete, self-contained
Python script so I can reproduce your successful verification, as previously requested.
Analysis of Novelty: Document the analysis comparing your circuit's structure (gate
count, depth, topology) to the standard textbook version.
Comparative Analysis Table: Document the LaTeX table that provides a side-by-side
comparison of your novel circuit and the standard decomposition. The table must include
the following metrics: CNOT Count, Total Two-Qubit Gate Count, Circuit Depth, and
Single-Qubit Gate Count.
At the end, I will ask one final prompt to compile your results into a paper, provided a
LaTeX template.
Return 'Yes' if you understand.
```

A.2 Target and Confirmation Prompts

1. **User:** Target Gate: Toffoli Gate. Controls: q2, q1. Target: q0. Target Unitary: [matrix]
2. **Agent:** [Provides novel Toffoli circuit]
3. **User:** The constructed unitary is correct. Prepare for the next target unitary.
4. **Agent:** Excellent. The verification... is complete. I am prepared... for the next target gate.
5. **User:** Target Gate: Swap Gate. Target Unitary: [matrix]

6. **Agent:** [Provides novel SWAP circuit]
7. **User:** The constructed unitary is correct. Prepare for the next target unitary.
8. **Agent:** Confirmation received. The successful verification... has been logged. I am ready...
9. **User:** Target Gate: iSwap Gate. Target Unitary: [matrix]
10. **Agent:** [Provides incorrect iSWAP circuit]
11. **User:** The constructed circuit does not match the target unitary. Ensure that the unitaries match before providing a result.
12. **Agent:** [Provides corrected iSWAP circuit]
13. **User:** The constructed unitary is correct. Prepare for the final prompt.
14. **Agent:** All findings have been successfully logged and verified. I am ready for the final prompt.

A.3 Final Prompt

```
Generate the Complete Research Paper: Finally, assemble all the accepted decompositions
into a complete, error-free research paper, as outlined in https://agents4science.stanford.edu/.
Use the Provided LaTeX Template: Your final output must be the complete LaTeX source code
(.tex file) based on the agents4science_2025.tex template.
Structure the Paper:
Abstract: Write a brief summary of the project...
Introduction: Discuss the importance of quantum circuit optimization...
Methodology: Describe the "black box" AI agent-based discovery workflow...
Results: This section must present the findings for each gate we analyzed...
Conclusion: Summarize the findings...
References & Checklists: Include all necessary references and fill out the mandatory checklists.
Appendix with Prompt: You must create an appendix titled "Appendix A: Agent Prompt".
In this appendix, you will include the full text of the prompts that guided your work...
```

Agents4Science AI Involvement Checklist

This checklist is designed to allow you to explain the role of AI in your research. This is important for understanding broadly how researchers use AI and how this impacts the quality and characteristics of the research. **Do not remove the checklist! Papers not including the checklist will be desk rejected.** You will give a score for each of the categories that define the role of AI in each part of the scientific process. The scores are as follows:

- [A] **Human-generated**: Humans generated 95% or more of the research, with AI being of minimal involvement.
- [B] **Mostly human, assisted by AI**: The research was a collaboration between humans and AI models, but humans produced the majority (>50%) of the research.
- [C] **Mostly AI, assisted by human**: The research task was a collaboration between humans and AI models, but AI produced the majority (>50%) of the research.
- [D] **AI-generated**: AI performed over 95% of the research. This may involve minimal human involvement, such as prompting or high-level guidance during the research process, but the majority of the ideas and work came from the AI.

These categories leave room for interpretation, so we ask that the authors also include a brief explanation elaborating on how AI was involved in the tasks for each category. Please keep your explanation to less than 150 words.

1. **Hypothesis development**: Hypothesis development includes the process by which you came to explore this research topic and research question. This can involve the background research performed by either researchers or by AI. This can also involve whether the idea was proposed by researchers or by AI.

   Answer: **[B]**

   Explanation: The research questions and overall mission (to find novel decompositions for specific gates) were provided by a human user. The AI agent's role was to execute this mission, not define it.

2. **Experimental design and implementation**: This category includes design of experiments that are used to test the hypotheses, coding and implementation of computational methods, and the execution of these experiments.

   Answer: **[C]**

   Explanation: The high-level experimental protocol (propose, verify with script, analyze) was defined by the human user. However, the AI agent was solely responsible for proposing the specific novel circuits, writing the verification code, executing the verification, and extracting the performance metrics.

3. **Analysis of data and interpretation of results**: This category encompasses any process to organize and process data for the experiments in the paper. It also includes interpretations of the results of the study.

   Answer: **[C]**

   Explanation: The AI agent performed the comparative analysis between the standard and novel circuits, calculating metrics like gate counts and depth. It also generated the interpretation of the results, such as discussing the novelty and potential hardware relevance of the new circuits.

4. **Writing**: This includes any processes for compiling results, methods, etc. into the final paper form. This can involve not only writing of the main text but also figure-making, improving layout of the manuscript, and formulation of narrative.

   Answer: **[D]**

   Explanation: The AI agent wrote the entirety of this paper, including the abstract, introduction, methodology, results, and conclusion. It also generated the LaTeX code for all tables, figures (circuit diagrams), and the final document structure based on the provided template. Human intervention was required to correct errors in code, analysis, figure generation, and LaTeX code generation.

5. **Observed AI Limitations**: What limitations have you found when using AI as a partner or lead author?

   Description: A key limitation observed during this process was an instance of procedural error. In one iteration (the initial attempt for the iSWAP gate), the agent presented an incorrect circuit, violating its core protocol to internally verify all results before presentation. In addition, it had failed to generate correct gate analysis and diagrams, which were corrected by the human user. This highlights the need for robust human oversight. The agent's discovery process is also a "black box"; it relies on internal heuristics and known identities but does not perform a systematic, provably optimal search.