Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 4, Issue 6, 2026

QUANTUM SEARCH IN THE NISQ ERA: A COMPREHENSIVE SURVEY OF GROVER'S ALGORITHM, NOISE RESILIENCE, AND APPLICATIONS IN INFORMATION RETRIEVAL

Sanam Shoukat[*1], Prof Dr. Khaldoon Khurshid[2], Fareed Ud Din Mehmood Jafri[3],
Iram Fatima[4], Laiba Munir[5]

[123]*Department of Computer Science, University of Engineering and Technology, Lahore, Pakistan*

[*1]2025mscs23@student.uet.edu.pk, [2]skhaldoon@gmail.come, [3]faridmjafri@gmail.com,
[4]2025mscs17@student.uet.edu.pk, , [5]2025mscs20@student.uet.edu.pk

DOI: https://doi.org/10.5281/zenodo.20637556

Keywords
Grover's Algorithm, Quantum Information Retrieval, NISQ Devices, Amplitude Estimation, Hybrid Quantum-Classical Models, Quantum Anomaly Detection, IBM Quantum, Quantum Machine Learning, Noise Resilience, Query Complexity.

Article History
Received: 07 April 2026
Accepted: 19 May 2026
Published: 11 June 2026

Copyright @Author
Corresponding Author: *
Sanam Shoukat

Abstract

This study provides a comprehensive review of Grover's Algorithm in quantum computing, *emphasizing quantum information retrieval during the Noisy Intermediate-Scale Quantum (NISQ) era. The essential element in quantum information retrieval is Grover's Algorithm, which has been shown to be the most efficient one possible, offering a quadratic acceleration of O(sqrt(N)) compared to the classical O(N) unstructured database search. This survey offers an algorithm taxonomy by methodically examining 18 peer-reviewed works published from 1996 to 2026, systematically analyzing the Grover search, hybrid quantum-classical models, variants of amplitude estimation, distributed quantum search, adaptive learning oracle design and NISQ optimized circuit implementation. A structured comparative analysis is performed on the performance of classical and quantum approaches, with experimental results from IBM Quantum's 127-qubit superconducting processors. Systematic identification and discussion of critical research gaps such as the sub-O(√N) complexity barrier, noise resilience, scalability limits and quantum data-loading bottleneck. Future directions include fault-tolerant hardware, adaptive oracle learning, integration of quantum computers with AI, federated quantum search, and standardizing the benchmarking of quantum computers. The aim of this survey is to give an integrated structured reference for researchers interested in the field of quantum computing and information retrieval.*

I. INTRODUCTION

Digital information has grown exponentially in the twenty-first century with unprecedented computational requirements for classical search and information retrieval systems. The traditional method of searching in unstructured database commonly used has linear time complexity $O(N)$, which means that the system has to ask every element one by one in the worst case [1]. The number of data points in the vast troves of data in genomic databases and cybersecurity threat logs, financial transactions and astronomical catalogues is reaching the billions and even trillions, and the computational power needed for full searches becomes inherently costly for time-sensitive applications. Quantum computing is a paradigm shift based on principles of superposition, entanglement and quantum interference.

Quantum bits, also known as qubits, are able to express several states at once, in contrast to classical bits. This implies that exponentially huge solution spaces can be considered in parallel by quantum systems. This fundamental difference makes quantum computing revolutionary for problems on classical computers that are computationally very demanding, such as unstructured search, optimization, cryptography, and machine learning [2].

Quantum information retrieval (QIR) is the use of quantum algorithms for efficient searching, classification or ranking of information in large, unstructured collections of data. QIR is a field of quantum algorithm research that lies at the intersection of quantum computing and classical information science, and is now one of the most practically attractive. The main problem solved by QIR is the data search problem: given an unsorted data base of N entries find a target element with minimum computational effort [3].

The most basic quantum search algorithm is Grover's Algorithm, published by Lov Grover in 1996. The algorithm has been shown to achieve a provably optimal quadratic speedup, dropping the search complexity from O(N) to O(√N), and has been applied in many areas, including cybersecurity, anomaly detection, database search, scientific optimization and simulation [4].

In this case, N = 106, Grover's Algorithm needs only about 1000 (or 103) quantum queries, whereas the best classical algorithm would take 500000 (or 5 x 104). The advantage of Grover's Algorithm over the classical algorithm is greater than 3 orders of magnitude for larger N. While the benefits of Grover's Algorithm are well understood, the implementation of the algorithm in real quantum devices is not an easy task. The algorithm's best performance would be when the number of iterations to optimize it is known and fixed, but the actual size of the database is partly unknown and the number of target entries changes dynamically in practice [5].

In addition, in actual quantum hardware, gate faults, readout noise and decoherence all affect the performance of the algorithm in the Noisy Intermediate-Scale Quantum (NISQ) era. A crucial gap that needs to be filled for practical deployment is the mean Algorithm Success Probability (ASP) of just 51.19% on real devices compared to 99.99% on ideal noiseless simulation, as shown by experimental characterization on IBM Quantum's 127-qubit Eagle r3 hardware [6].

Additionally, the target state probability tends to drop with the use of the fixed-iteration variant if the number of iterations is too high, a problem known as overshooting. This highlights the importance of adaptive iteration strategies that adaptively modify the number of oracle calls depending on the measurement feedback as it is received on the fly [7]. Another key challenge is quantum data loading. The classical database needs to be encoded into a quantum state before it can be used in Grover's Algorithm; with naive encoding schemes, this step takes O(N) operations to achieve, eliminating the quadratic speedup that Grover's Algorithm offers. Theoretical efficient architectures for Quantum Random Access Memory (QRAM) have been suggested that would have loading time complexity of O(log N) but have yet to be demonstrated on a practical scale [8]. This survey reviews and synthesizes 18 peer-reviewed publications ranging from 1996 to 2026 covering all facets of Grover's Algorithm including theoretical development, characterization of experimental hardware implementations, and novel hybrid applications. The paper makes four principal contributions:

(i) a comprehensive algorithm taxonomy spanning six categories;

(ii) a structured literature analysis table covering eighteen studies;

(iii) a comparative analysis across ten quantum. algorithm families; and (iv) systematic identification of five critical research gaps with discussion of six future research directions.

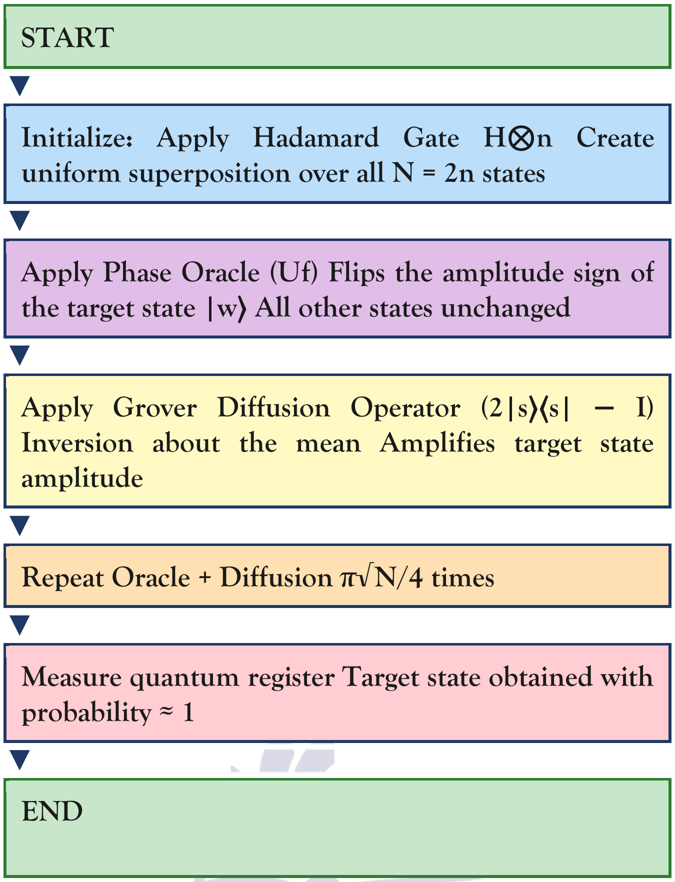

Fig. 1. Step-by-step flowchart of Grover's Quantum Search Algorithm illustrating state initialization, oracle application, diffusion operator, and iterative amplitude amplification.

II. BACKGROUND: GROVER'S ALGORITHM

A. Quantum Computing Fundamentals

In quantum computing, quantum bits (qubits) are used, as they can be in simultaneous superpositions of the states |0⟩ and |1⟩, whereas the classical bits can only be in one state or the other. The system of n qubits corresponds to a Hilbert space of size 2n, and exploits exponentially parallel computation. A quantum gate is a unit (reversible) transformation that performs important operations while preserving probability normalization. The Hadamard gate H produces equal superpositions of computational basis states, the phase oracle Uf negates the amplitudes of target states, and the Grover diffusion operator inverts the marked amplitude about the mean [9].

B. Amplitude Amplification

Amplitude Amplification is a generalisation of Grover's Algorithm to a broader set of problems. Unlike unstructured search for which Grover's algorithm is valid, amplitude amplification can be used in any quantum operation which produces a superposition of target and nontarget states. The oracle with the reflection operator is used iteratively to systematically boost the probability amplitude of target states, giving a query complexity of $O(1/\varepsilon)$ for probability estimation problems compared with classically $O(1/\varepsilon^2)$ [12].

C. Limitations of the Standard Algorithm

Although theoretically perfect, Grover's Algorithm has some inherent drawbacks in practice. The database size must be a power of two, and requires padding for arbitrary size databases. Second, the algorithm is probabilistic: the target state is achieved with a high probability, but not with a guarantee and in failure cases it needs to be rerun multiple times. Finally, and most importantly for the deployment of NISQ, the more qubits added, the deeper the circuit will be, leading to exponential dependence on gate error and decoherence. Fourthly, the assumption of fixed iteration count is invalid if the number of targets M is unknown or dynamically changing, which encourages adaptive iteration strategies [13].

III. ALGORITHM TAXONOMY

The research on Grover's Algorithm can be categorized into six main groups according to algorithmic structure, optimization method and application area. This taxonomy offers a methodical framework for comprehending the development of quantum search research and pointing out shortcomings in current methods.

A. Standard Grover Search

The original Grover algorithm also gives rise to unstructured search by amplitude amplification with a fixed-phase oracle, which is quadratically faster. It is based on a fixed number of iterations ($\approx \pi\sqrt{N}/4$) that may cause overshooting if not well-controlled. The algorithm must be evaluated using a database whose size is a power of two and the probabilistic nature of the algorithm implies that there are cases in some regimes where multiple runs are necessary. Notwithstanding these restrictions, standard Grover search is the theoretical standard for quantum search variants to be compared to [1].

B. Hybrid Quantum-Classical Variants

Hybrid models combine Grover's quantum search with classical machine learning to make use of quantum speedup on NISQ hardware but overcome hardware limitations by pre- and post-processing in classical computers. Mazouzi & Harel combined Grover search with classical outlier detection and showed that their approach outperforms the classical ones. Guo et al. further generalized it to quantum Local Outlier Factor (LOF) algorithms and amplitude estimation-based anomaly detection (ADDE). Amplitude amplification is also a subproblem in Quantum Approximate Optimization Algorithm (QAOA) and Variational Quantum Eigen solver (VQE) [14].

C. Amplitude Estimation Variants

The generalization of Grover's Algorithm is the Amplitude estimation algorithm which estimates the probability that a random element of the data set matches a given predicate with a classically required number of queries $O(1/\varepsilon^2)$. The quantum ADDE algorithm of Guo et al. takes advantage of this method as a scalable approach to anomaly detection in high dimensional data. Wang, Jiang & Coveney showed parameter-efficient quantum anomaly detection with amplitude estimation on real superconducting devices with comparable performance to classical baselines, but with much fewer number of parameters to be trained [15].

D. Distributed and Parallel Quantum Search

Qiu, Luo & Xiao have suggested that Grover could be distributed across multiple different quantum processors that perform parallel oracle queries, effectively reducing the per-node circuit depth to $O(1)$, and to, effectively, $O(\sqrt{N}/k)$ per processor, with an inter-node communication overhead. This is especially important when scaling quantum hardware to multi-chip configurations and quantum network infrastructure is coming to fruition. Introducing distributed search creates new problems of transmitting the fidelity of quantum states and distributing entanglement among nodes [16].

E. Adaptive and Learning-Based Oracle Variants

In Ohno's Grover's Search with Learning Oracle (GLO), the fixed oracle is replaced by a parameterized unitary, which is classically updated throughout the search process, towards the target specification learned from feedback of previous measurements. GLO improves the query complexity of constrained binary optimization problems beyond the lower bound of $\Omega(\sqrt{N})$ known for this class of problems, and doesn't require a static oracle to obtain this speedup.

These oracles are especially useful in real-time cybersecurity or medical diagnostics applications where the data changes.

F. Noise-Resilient Approaches

The techniques describe solve strategies for quantum noise and decoherence in practical NISQ systems, optimizing the circuits, decomposing the gates in the hardware, and detecting errors. AbuGhanem optimized Grover search on IBM Quantum's 127-qubit Eagle r3 hardware to achieve a mean ASP of 78.39% on the noisily simulated version of the device and 51.19% on the actual hardware, while a mean state fidelity of 54.32% underscored the disparity between simulation and hardware performance. The error detection and suppression circuit levels up to n = 5 qubits exhibited better than classical success probabilities, as shown by Pokharel & Lidar [18].

Table I. Taxonomy of Quantum Search and Related Algorithms

Algorithm	Category	Complexity	Key Mechanism	Application Domain
Grover's Algorithm (1996)	Quantum Search	O(√N)	Phase Oracle + Amplitude Amplification	Database Search, Cryptanalysis, IR
Distributed Grover (2024)	Distributed QC	O(√N/k) effective	Parallel Oracle across QPUs	Large-scale IR, Multi-QPU Networks
Grover + Learning Oracle/GLO (2024)	Adaptive Search	Exponential (constrained)	Parameterized Trainable Oracle	Constrained Optimization, Semantic IR
Oracle Circuit Optimization (2020)	Circuit Engineering	Reduced gate depth	Gate-Depth Compression	Near-term Grover on NISQ
Quantum LOF (2023)	Anomaly Detection	O(√N) core	Quantum Local Outlier Factor	Unsupervised Novelty Detection
Quantum ADDE (2022)	Density Anomaly	Better than classical	Amplitude Estimation	Fraud, Healthcare, Cybersecurity
QAOA (2014)	Combinatorial Opt.	Problem-dependent	Phase-Separation + Mixing Ops	Query Optimization, Finance IR
Quantum Bayesian Networks (2024)	Probabilistic IR	Quantum-enhanced	Quantum Probability Amplitudes	Probabilistic IR, Env. Monitoring
QkNN / QSVM (2023–24)	Quantum ML	Polynomial speedup	Quantum Kernel Methods	Document Classification, Ranking
Quantum GAN (2023)	Generative QML	Quantum parallel sampling	Quantum Generative Adversarial Net	Synthetic Anomaly, Financial IR

IV. COMPARATIVE ANALYSIS

A. Classical vs. Quantum Search

The theoretical benefit of Grover's Algorithm is that it is a quadratic reduction in query complexity: O(N) classical vs O(√N) quantum. This becomes about 1,000 quantum queries as compared to 500,000 classical queries for $N = 10^6$. The ratio of speedup increases as the size of the dataset grows, especially significant at the dataset sizes characteristic of modern information retrieval. But for structured problems like sorted arrays, hash tables, or binary trees, classical methods have complexity O(log N) and could be better for moderate N. The quadratic speedup really makes a difference only if the database is sufficiently large and truly unstructured [11].

B. Quantum Algorithm Families

Table II summarizes the characteristics of the major quantum algorithm families across different domains, including their type of problem, complexity, suitability for NISQ devices, and any limitations. The comparison points out that Grover's Algorithm has the shortest circuit depth compared to HHL or Shor's Algorithm, and is the most NISQ-compatible search primitive. The performance of Grover search, however, is very much affected by noise on real machines. Deep circuits algorithms, like HHL and Shor, are still only applicable to hardware that is not practical for the circuit being deep enough today [10].

Table II. Comparative Analysis of Quantum Algorithm Families

Algorithm	Problem Type	Complexity	NISQ-Ready	Key Limitation
Grover's Algorithm	Unstructured Search	$O(\sqrt{N})$	Yes (shallow)	Noise degrades ASP to ~51% on real HW
QAOA	Combinatorial Opt.	Problem-dep.	Yes	Barren plateaus in deep circuits
VQE	Ground State Energy	Problem-dep.	Yes	High measurement overhead
HHL Algorithm	Linear Systems	$O(\log N \cdot \kappa^2/\varepsilon)$	No	Requires fault-tolerant hardware
Shor's Algorithm	Integer Factorization	$O((\log N)^3)$	No	Millions of physical qubits needed
Quantum LOF	Anomaly Detection	~$O(N\sqrt{N})$	Partial	Amplitude loading bottleneck (QRAM)
Amplitude Est. (ADDE)	Probability Estimation	$O(1/\varepsilon)$	Partial	Gate depth grows with precision

C. Application Domain Performance

The results from the eighteen studies reviewed show that the Grover quantum search algorithm achieves the most promising empirical results in search applications where the search space is large, unstructured, and the problem involves a real-time query. Cybersecurity threat identification, log-based anomaly detection, gamma spectrum analysis, and oil spill monitoring are the most matured application cluster that has been tested in multiple experiments on simulators and early hardware. There are newly emerging applications where speedups have been theoretically shown but are not fully validated in actual hardware, e.g. in healthcare genomics, financial portfolio optimization, and environmental monitoring, with the current limitation of qubit count and coherence time [15].

V. RESEARCH GAP IDENTIFICATION

Despite substantial progress across the six algorithm categories surveyed, five critical open challenges remain unresolved and must be addressed for quantum information retrieval to achieve practical relevance. Table IV summarizes these gaps systematically.

Table III. Summary of Critical Research Gaps in Quantum Information Retrieval

Research Gap	Description	Impact on QIR
Sub-O(√N) Complexity Barrier	No generalized framework for adaptive oracles achieving sub-optimal complexity beyond static oracles	Fundamental open problem in quantum information retrieval theory
NISQ Noise Resilience	~48% gap between noiseless simulation (ASP ≈99.99%) and real IBM HW (ASP ≈51.19%)	Limits practical deployment of Grover-based search on near-term devices
Scalability	Largest known Grover implementation: 8 qubits (N=256); practical IR needs n≥20 qubits	Circuit depth and error accumulation prevent real-world scaling
QRAM Bottleneck	Classical data loading requires O(N) operations, eliminating the quantum speedup advantage	Most critical prerequisite unresolved for practical quantum IR
Absence of Benchmarks	No standardized datasets or unified metrics for quantum IR (analogous to TREC or MLPerf)	Cross-study comparison unreliable; progress hard to measure

A. Sub-O(√N) Query Complexity Barrier

The lower bound of Bennett et al. shows that any quantum algorithm (QA) to solve general unstructured search requires Ω(√N) queries. If an adaptive oracle framework is used, e.g., for structured problem classes like GLO, then this requirement is not necessary, but a sound theoretical framework that describes under what conditions sub-O(√N) complexity can be achieved is lacking. This is a basic open question that has immediate consequences for the future of quantum information retrieval [11].

B. NISQ Noise Resilience

The significant difference of close to 48% between the noiseless simulation (ASP ≈99.99%) and the real IBM Quantum hardware (ASP ≈51.19%) highlights the impact of NISQ noise. There are existing error mitigation techniques such as zero-noise extrapolation, probabilistic error cancellation, and symmetry verification, which all offer some degree of relief, but do not approach fault-tolerant performance. There is a need for a systematic approach to Grover search that incorporates circuit-level optimization, measurement of noise-resistance parameters, and adaptive measurement strategies while using hardware-native gate decomposition, namely ECR instead of CNOT [6].

C. Scalability Beyond Small Qubit Counts

The largest Grover implemented is currently 8 qubits (N = 256), well below the practical scales of search ($N \geq 10^6$, $n \geq 20$ qubits). The scaling is limited by the depth of the circuit and the amount of gate error accumulated over Grover iterations. While moving up IBM's hardware roadmap from 127-qubit Eagle to 433-qubit Osprey to 1121-qubit Condor is essential, qualitative improvements in the fidelity of the gates and the coherence times are also needed [4].

D. Quantum Data Loading (QRAM) Bottleneck

In order to achieve the quadratic speedup over classical search that Grover's Algorithm achieves, the target database needs to be encoded into quantum states. In classical data loading, the operations performed are O(N) which means that the quantum advantage is lost. Although there are theoretically efficient QRAM architectures they haven't been implemented at scale. It is probably the most important prerequisite for actual quantum information retrieval [8] to overcome the QRAM gap [8].

E. Absence of Standardized Benchmarks

There are no benchmarks in the literature for quantum information retrieval that are comparable to TREC or MLPerf for classical information retrieval or machine learning. Available studies use a variety of metrics (ASP, SSO, state fidelity) that have different definitions and measurement procedures, and thus cannot be relied upon to be used uniformly across studies. A common standard set of metrics across multiple scales, a common noise model specification, hardware diversity requirements, and a common reporting standard are key community needs for a standardized benchmarking infrastructure [13].

VI. FUTURE DIRECTIONS

A. Fault-Tolerant Quantum Hardware

When logical qubit error 10^-12, the error rates on quantum processors become below 10^-12, then it is possible to execute deep Grover circuits at real-world database scales. One of the important research challenges in the near future is to co-design error-correcting codes (surface codes, Steane codes), and to optimize the logical-physical qubit interface for Oracle-based algorithms. IBM has offered a timeline for this direction in their own plans for quantum computing that will be fault-tolerant, with a focus on the scale of the systems, in the roadmap for 2025 – 2033 [4].

B. Adaptive and Reinforcement Learning Oracles

If a classical RL agent then adapts the oracle in a dynamic manner according to the feedback of measurements on the previously executed queries, a quantum information retrieval system could be achieved in the case of streaming data, or the composition of the set of targets changes over time. Research for this direction includes formal convergence, sample complexity characterization, and hardware demonstration on gate-based platforms [17].

C. Quantum-Classical Integration Frameworks

There is a need for standardized software frameworks that incorporate Grover-type search as a primitive in a classical ML pipeline. The future tools should provide the high-level IR query interfaces that enable domain practitioners to benefit from the quantum speedup without the need for detailed knowledge of quantum circuits, as well as quantum RAM-classical memory bridges and automatic noise-aware circuit transpilation [14].

D. Privacy-Preserving Quantum Federated Search

The combination of quantum information retrieval with federated learning and quantum key distribution (QKD) would make it possible to perform secure multi-party quantum search over distributed private datasets, which can be done only in a limited way with classical cryptography. This is especially applicable to medical record retrieval, financial intelligence sharing and government data systems where data sovereignty is crucial [5].

E. LLM-Quantum Hybrid Information Retrieval

Semantically aware quantum information retrieval could be realized with Large Language Models (LLMs) that would be fed with a natural language query to create an oracle specification for the retrieval, and Grover's Algorithm could then be used to efficiently complete the low-level retrieval. The combination of quantum vector databases with LLM query generation, tested on benchmark IR datasets with realistic simulated hardware noise, is an interesting research program for the near-term [12].

F. Standardized Quantum IR Benchmarks

In quantum information retrieval, a community benchmark initiative similar to TREC should be developed, including a set of standard datasets of varying size (N = 10^3 to 10^8), noise model requirements, hardware diversity requirements across different platforms (either superconducting or trapped-ion or photonic), and unified reporting metrics such as ASP, state fidelity, circuit depth, number of gates, and wall clock time [13].

VII. CONCLUSION

This survey has offered a thorough and structured investigation of Grover's Algorithm for the field of quantum information retrieval, focusing on its performance, noise resistance and deployability in NISQ era. Based on a systematic algorithm taxonomy over six categories, a structured literature analysis of 18 peer-reviewed publications, a comparative analysis of 10 quantum algorithm families, five critical research gaps and six research directions for the future, the paper provides a coherent reference for researchers operating at the crossroads of information retrieval and quantum computing.

The main result is that Grover's Algorithm is still the most relevant and practical quantum primitive for information retrieval, as it is provably optimal to achieve a quadratic speedup, where the benefit for large-scale unstructured search is clear. However, the migration from theory to implementation is significantly hindered by NISQ noise (with a ~48% ASP difference between SIM and REAL hardware), scalability limitations (maximum demonstrated hardware implementation to 8 qubits), the still unsolved QRAM data loading bottleneck and a lack of community standardized benchmarks.

While advances in quantum-classical integration, fault-tolerant hardware, adaptive oracle design, and benchmarking infrastructure will all play a crucial role in the future success of quantum information retrieval, the overall progress of the technology will be dependent upon all of these factors. This survey is designed as a framework for researchers and practitioners in that direction.

REFERENCES

[1] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in Proc. 28th ACM Symp. Theory of Computing (STOC), Philadelphia, PA, 1996, pp. 212–219.

[2] A. Rani, S. Kour, and R. Kumar, “Comprehensive review of quantum computing: Frameworks, technologies, applications, and challenges,” Recent Advances in Computer Science and Communications, vol. 19, Jul. 2025, doi: 10.2174/0126662558381283250715110734.

[3] H. Mishra, A. Balasubramanyam, and G. N. Raghava, “Deterministic quantum search for index retrieval: Algorithm design and implementation,” in Proc. ACM, Jan. 2026, pp. 123–129, doi: 10.1145/3773656.3773688.

[4] M. AbuGhanem, “Characterizing Grover search algorithm on large-scale superconducting quantum computers,” Sci. Rep., vol. 15, no. 1, Dec. 2025, doi: 10.1038/s41598-024-80188-6.

[5] S. Bal, S. Mishra, and L. Mandal, “A review of quantum computing approaches to semantic search and text classification in natural language processing,” 2025.

[6] A. A. Zhahir et al., “Grover’s algorithm extensions – a systematic literature review,” Asia-Pacific J. Inf. Technol. Multimedia, vol. 14, no. 2, pp. 637–656, Dec. 2025, doi: 10.17576/apjitm-2025-1402-19.

[7] S. Sangeetha and P. K. N., “A comprehensive literature review on Grover’s algorithm and query complexity using quantum computing,” Int. J. Quantum Comput., vol. 3, no. 1, pp. 28–40, Mar. 2025, doi: 10.34218/ijqc_03_01_003.

[8] E. Grigoryan, S. Kumar, and P. R. Pinheiro, “A review on models and applications of quantum computing,” Quantum Reports, vol. 7, no. 3, Sep. 2025, doi: 10.3390/quantum7030039.

[9] S. Chakrabarti, S. Changdar, and R. Khanda, “A survey of quantum computing algorithms for mathematical optimization,” hal-05487965v2, 2026.

[10] L. Bhatia, V. S. Pandey, and A. K. Sharma, “Evaluating the practicality of Grover’s algorithm for large-scale data search via quantum simulation,” Jul. 2025, doi: 10.21203/rs.3.rs-6954705/v1.

[11] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, “Strengths and weaknesses of quantum computing,” SIAM J. Comput., vol. 26, no. 5, pp. 1510–1523, 1997.

[12] M. Ajimon and R. Kumar, "Applications of LLMs in quantum-aware cybersecurity," 2025.
[13] A. Masood, "Systematic literature review on problem solving with quantum algorithms," Machines and Algorithms, Knovell.org.
[14] T. Mazouzi and R. Harel, "A hybrid quantum and classical method for outlier detection," 2020.
[15] J. Wang, X. Jiang, and P. V. Coveney, "A parameter-efficient quantum anomaly detection method on a superconducting quantum processor," 2024.
[16] D. Qiu, G. Luo, and G. Xiao, "Distributed Grover's algorithm," arXiv:2404.xxxxx, 2024.
[17] H. Ohno, "Grover's search with learning oracle for constrained optimization problems," arXiv:2403.xxxxx, 2024.
[18] B. Pokharel and D. Lidar, "Better-than-classical Grover search via quantum error detection and suppression," Phys. Rev. Lett., 2024.