Noisy Qubits, Hard Problems: A Systematic Review and Taxonomy of Quantum Optimization Beyond Toy Benchmarks

Abstract

Quantum optimization has become a leading application for near-term quantum computing, and yet many publications compare algorithms against idealized assumptions and small toy benchmarks. This limits the interpretability, reproducibility, and practical relevance of reported performance gains, particularly in the noisy intermediate-scale quantum (NISQ) era. In this work, we present a systematic literature review that investigate quantum optimization beyond toy benchmarks. Following established SLR protocols, we analyze the literature along multiple methodological dimensions, including algorithmic approach, benchmark realism, encoding strategies, hybrid quantum-classical workflows, hardware and noise modeling, evaluation metrics, and reporting practices. We introduce a unified taxonomy that captures the interaction between problem formulation, encoding overhead, noise-aware execution, and hybrid optimization loops. In addition, we propose a reproducibility checklist and scoring rubric to assess reporting completeness and experimental rigor across studies. Instead of developing new quantum optimization algorithms or making theoretical quantum advantage claims, the main contribution of this work is that it provides a methodological analysis on benchmarking realism and encoding or evaluation practices as well as reproducibility rigor in NISQ-era quantum optimization studies. Our review uncovers a continued chasm between algorithmic innovation and evaluation maturity, with the number of publications and diversity of methods growing without corresponding growth in standardized benchmarks, strong classical baselines, noise-consistent evaluation, or reproducibility.

Keywords: Quantum Optimization, Noisy Quantum Devices, NISQ, QAOA, Quantum Annealing, Toy Benchmark, Benchmarking

1 Introduction

Quantum computing has been developed as a promising paradigm to solve computational problems that are classically intractable, especially in the field of combinatorial and continuous optimization [1], [2] [3]. Problems which underlie well-known scientific, engineering, and industrial problems (like scheduling, routing, portfolio optimization, and resource allocation) are at the heart of science and could achieve significant practical impact even from marginal performance gains [4],[5],[6],[7]. As such, quantum optimization has emerged as one of the most actively investigated application areas in the NISQ era of quantum computing.

Modern quantum hardware, in fact, works in the so-called Noisy Intermediate-Scale Quantum (NISQ) devices with a limited qubit counts and short coherence times, gate errors, and few connectivity [8],[9],[10]. These limitations impose severe restrictions on the depth and reliability of quantum circuits, leading to fundamental questions about the practical benefits of quantum optimization algorithms on actual devices [11]. Despite rapid growth in the literature, there remains considerable uncertainty regarding what has genuinely been achieved beyond proof-of-concept demonstrations

This systematic literature review provides an in-depth look at the landscape of quantum optimization research beyond toy benchmarks paying particular attention to how noise, scalability, and benchmarking choices impact claims around performance.

1.1 Background and Motivation

Optimization has been considered an obvious target for quantum advantage since the early days, with early suggestions like quantum annealing and adiabatic quantum computation, through algorithms based on gates, like QAOA. Theoretically, these approaches propose alternative ways to navigate through energy landscapes complex enough to preclude navigation by classical heuristics [12].

In reality, however, the ability of a quantum optimization device is severely restricted by the limitations of NISQ-era hardware [13]. Noise builds up quickly against circuit depth, going through variational optimization landscapes becomes progressively hard to train, and algorithmic behavior is often dominated by hardware-specific effects. Therefore, most of the reported successes are for very small-scale, highly structured, or artificially simplified problems, commonly known as toy benchmarks [14],[15].

The growing gap between theoretical promise and experimental reality motivates a careful reassessment of the field. Rather than asking whether quantum optimization will eventually outperform classical methods, a more immediate and practical question must be addressed: What has actually been demonstrated under realistic noise and scaling conditions, and how meaningful are these results for real-world optimization problems?

1.2 Limitations of Existing Studies

Although the quantum optimization literature is extensive, there are several systematic limitations of previous studies. First, a heavy dependence on toy benchmarks, such as small size Max-Cut graphs, random instances of the QUBO, hand-crafted problems to advantage some quantum algorithms. Although helpful for early testing, such benchmarks are typically unable to reflect the complexity, constraints and heterogeneity involved in practical optimization problems. Second, the noise assumption in many studies is idealized or incomplete. The reported results are often coming from noiseless simulators or simplified noise models that do not truly capture device-specific error mechanisms such as correlated noise, crosstalk, and measurement bias.

This curtails the generalizability of performance improvements reported. Third, scalability analysis is often missing or superficial. Evaluation of performance often involves small fixed-sized problems, with little discussion on how algorithmic behavior evolves as a function of the number of variables, constraints, and circuit length. The lack of such an analysis makes it hard to examine whether observed advantages remain or disappear when increasing problem sizes. Finally, there are significant variations in the reporting of outcomes among studies, which makes comparisons difficult. Key details such as encoding strategies, hardware parameters, optimization budgets, and evaluation metrics are frequently under-specified. This paper does not propose new quantum optimization algorithms but rather fosters a systematic categorization of such techniques, reflects on encoding and benchmarking practices, and seeks out methodological gaps that hinder the interpretability and reproducibility of existing work.

1.3 Research Objectives and Contributions

Unlike prior surveys that emphasize idealized simulations or small toy instances, this PRISMA-guided systematic literature review focuses on quantum optimization studies evaluated under NISQ-relevant conditions, including explicit noise models and/or hardware-informed constraints. We introduce a multi-axis taxonomy covering:

- Problem realism and instance scale
- Encoding strategy (QUBO/Ising/CQM)
- Noise/hardware treatment
- Hybrid quantum–classical workflow design. Beyond descriptive synthesis

we extract standardized fields and conduct longitudinal trend analysis and a reproducibility audit (classical baselines, metrics definitions, hardware/noise details, repetitions, and artifact availability), identifying concrete methodological bottlenecks and evidence gaps that currently limit practical claims of advantage.

2 Fundamentals of Quantum Optimization

This section review is used to offer a brief concept introduction to optimization in the quantum world, which does not repeat standard textbook knowledge but is enough to put into necessary perspective literature being reviewed. The emphasis is on problem definition, algorithmic positions of strength, and the hardware limits that condition realizable performance in the NISQ era.

2.1 Optimization Problems in Classical vs. Quantum Computing

As in many real-world optimization problems, including scheduling, routing, graph partitioning, and resource allocation, belong to the class of combinatorial optimization problems where the goal is to find an optimal assignment of discrete variables subject to a set of constraints. Quantum optimization algorithms are not a direct solver for general optimization problems; rather, they solve only select classes of problems.

Rather, the instances of classical problems have to be expressed in mathematical formats that can run on quantum hardware, such as Ising or more generally Quadratic Unconstrained Binary Optimization (QUBO) [16] [17]. In these formulations, the problem variables are embedded as binary or spin variables, constraints are enforced by adding penalty terms and the objective function is associated with an energy landscape such that its ground state corresponds to the solution [18]. This translation allows quantum execution but imposes concrete limitations. The process of coding constraints can vastly enhance problem density and distance between energy scales, leading this way to a much stronger influence of noise and detuning on the system. As a result, the quality of the problem encoding is an important factor influencing the success of quantum optimization techniques, and this theme reappears in many of the papers we review.

A common approach to bridge classical problems with quantum hardware is to reformulate optimization tasks into standard mathematical forms such as *Quadratic Unconstrained Binary Optimization (QUBO)* and *Ising Hamiltonians*. In the QUBO formulation, the objective function is expressed as:

$$\min_{\mathbf{x}\in\{0,1\}^n} \mathbf{x}^T Q \mathbf{x},$$

where Q is a real-valued matrix encoding linear and quadratic interactions between binary variables. Similarly, the Ising formulation represents problems using spin variables $s_i \in \{-1, 1\}$:

$$H(s) = \sum_i h_i s_i + \sum_{i<j} J_{ij} s_i s_j,$$

where h_i and J_{ij} denote local fields and coupling strengths, respectively. These formulations enable direct implementation on quantum annealers and gate-based quantum devices.

2.2 Quantum Optimization Algorithms

There have been multiple algorithmic paradigms proposed for quantum optimization, reflecting different viewpoints in terms of assumptions about capabilities, noise tolerance and scalability.

2.2.1 Quantum Approximate Optimization Algorithm (QAOA)

The Quantum Approximate Optimization Algorithm (QAOA) is one of the most widely studied gate-based quantum optimization methods [19],[20],[21]. It is a variational method that iteratively applies problem and mixer Hamiltonians with adjustable parameters optimized in a classical outer loop [22]. Nevertheless, its implementable performance on NISQ machines is limited by the circuit depth, the optimization challenge on parameters and the accumulation of noise [23]. As problem sizes grow, useful levels of performance are often only achievable with even deeper circuits that soon run beyond the coherence window of current hardware[24]. As a result, many QAOA studies remain limited to shallow depths and small-scale benchmarks. It alternates between applying a *problem Hamiltonian* and a *mixing Hamiltonian*, parameterized by angles $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$:

$$|\psi(\boldsymbol{\gamma},\boldsymbol{\beta})\rangle = \prod_{k=1}^{p} e^{-i\beta_k H_M} e^{-i\gamma_k H_C} |+\rangle^{\otimes n}.$$

2.2.2 Quantum Annealing

In contrast, quantum annealing is an alternative approach based on a different computational model and usually executed on dedicated analog hardware [6],[25]. Optimization problems are translated into an Ising Hamiltonian, and the system evolves adiabatically from an initial Hamiltonian to a problem Hamiltonian with slow rates in ideal case up to its ground state. Some of the annealing-based techniques have the advantage of relatively low control resources, and they have been realized on a larger number of qubits than gate-based techniques [13]. However, they provide only a small degree of flexibility in terms of how intermediate states are affected and are more specifically focused than other reconfiguration schemes [1]. Furthermore, discerning true quantum effects from classical thermal behavior is an outstanding challenge in presence of realistic noise. Quantum annealing (QA) is an analog quantum optimization approach that solves Ising-formulated problems by adiabatically evolving a quantum system from an initial Hamiltonian to a problem Hamiltonian:

$$H(t) = A(t)H_0 + B(t)H_P,$$

where H_0 is a transverse-field Hamiltonian and H_P encodes the optimization objective.

2.2.3 VQE-based and Hybrid Algorithms

Beyond QAOA and annealing, a growing body of work explores hybrid quantum–classical optimization frameworks, often inspired by the Variational Quantum Eigensolver (VQE) [26] [27]. These approaches mix parametrized quantum circuits with classical optimizers and often rely on problem-specific heuristics, decomposition techniques or even classical preprocessing [28] [29]. Hybrid schemes are typically less sensitive to noise and hardware limitations, as they delegate much of the computational load to classical resources. Yet, this robustness is achieved at the expense of additional classical overhead, which makes it hard to determine if any observed improvements are due to quantum or classical contributions [30]. Such ambiguity provides an additional challenge to claims of quantum advantage and underscores the need for rigorous benchmarking and reporting.

2.3 NISQ Hardware Characteristics

The practical behavior of all quantum optimization algorithms is fundamentally shaped by the characteristics of NISQ hardware [31]. Current devices are limited by gate errors, decoherence, measurement noise, and device-specific effects such as crosstalk. These errors grow too quickly with the depth of circuits, which would cap the complexity of the algorithm that can be executed [32] [33]. Furthermore, nontrivial overhead due to qubit connectivity restrictions is introduced in the encoding of the problem and circuit compilation processes, recurrently involving extra swap operations that contribute more exposure to noise [34]. Thus, nominal problem size does not necessarily equate to actual computing capabilities.

Finally, the presence of challenging gate depth and execution time limitations hampers the practical feasibility of deep variational circuits and long annealing schedules. Together, these limitations account for why many studies of quantum optimization are restricted to small or highly structured problems and why exceeding toy benchmarks continues to be a fundamentally unsolved problem.

3 Methodology

This study adopts a Systematic Literature Review (SLR) methodology to ensure an objective, repeatable, and bias-free summary of the state of quantum optimization beyond toy problems. Due to varied methodologies and the fast pace of development in NISQ-era research, a review protocol is necessary to systematically find, evaluate, and classify the relevant work.

The overall process adheres to the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines, which have been considered a standard of high-impact review articles [35].

3.1 Review Design and PRISMA Framework

The review was performed by applying a four-step strategy according that followed the PRISMA flow diagram: identification, screening, eligibility, and inclusion. Every step was described explicitly such that it can be traceable and reproducible independently [36].

Identification: Academic databases were searched using predefined search questions to identify appropriate records for review.

Screening: Duplicates were removed, and title/abstract screening was applied to eliminate obviously irrelevant studies.

Eligibility: Languages Eligible for Full-Text Assessment (Explicit Inclusion and Exclusion Criteria)

Inclusive criteria: Studies that met all of the criteria were included for qualitative synthesis and synthesis.

A flowchart of the process, with the number of records retained and reasons for exclusion at each stage is shown in Figure 1(PRISMA flow diagram). This structured approach ensures that the final study set is both comprehensive and methodologically defensible.

3.2 Data Sources and Search Strategy

3.2.1 Databases and Publishers

To obtain a wide representative sample of the quantum optimization literature and minimize search bias, searches were performed in multiple data sources, including

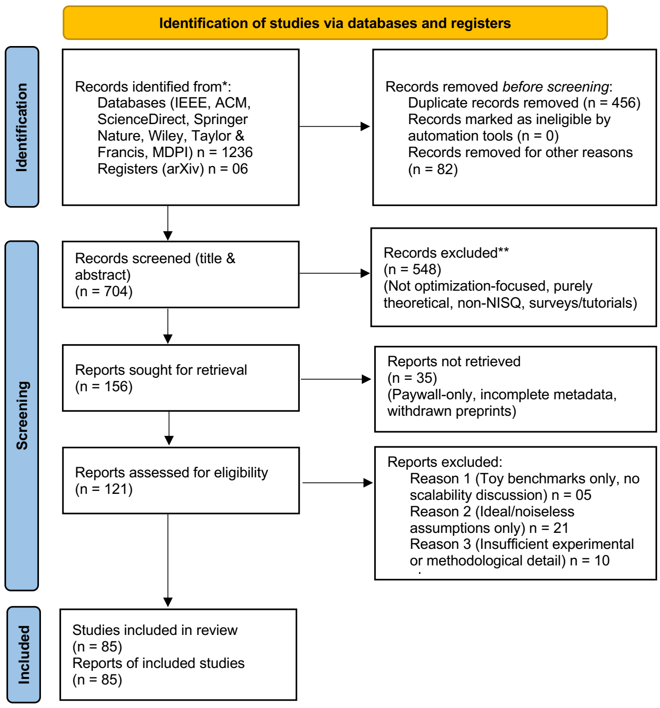

Fig. 1 PRISMA flow diagram

both peer-reviewed articles and high-quality preprints. Specifically, the following were searched: IEEE Xplore (journals and conferences in quantum computing, optimization and engineering), ACM Digital Library (algorithmic or computer science-focused literature), ScienceDirect/Elsevier (interdisciplinary journals covering quantum information, optimization and applied physics), SpringerLink (journals and edited volumes in quantum technologies and applied mathematics), Wiley Online Library (interdisciplinary work related to quantum science and optimization research), Taylor & Francis Online (quantum computing as well as optimization studies), MDPI (Journals of Quantum Physics are available at MDPI), and arXiv (to include emerging and rapidly evolving research not yet formally published).

3.2.2 Search Strings and Time Span

Search queries were composed by combining Boolean operators, specific keywords and domain terminology on the topics of quantum optimization, NISQ devices, benchmarking methodologies, and noise-aware evaluations. Keywords were mapped to titles, abstracts, and author-assigned keywords as appropriate depending on the functionality of the database. The search was restricted to papers in 2019 and onward, due to this being the time when NISQ machine became practical accessible, and experimental analysis paid more attention on hardware. This time frame accommodates modern capabilities of quantum computing, while minimizing over-representation of early entirely theoretical work. Table 1 shows the publisher names and their corresponding search strings for quantum optimization literature, formatted across multiple lines for readability.

Table 1 Search strings for Quantum Optimization Literature by publisher.

Publisher	Search String
IEEE	`"All Metadata":quantum optimization AND "All Metadata":NISQ OR "All Metadata":noisy intermediate-scale quantum AND "All Metadata":Noisy Qubits`
ACM	`[All: "quantum optimization"] AND [All: nisq] OR [All: "noisy intermediate-scale quantum"]`
ScienceDirect	`"quantum optimization" AND ("NISQ" OR "noisy intermediate-scale quantum") AND "Noisy Qubits"`
Springer	`"quantum optimization" AND ("NISQ" OR "noisy intermediate-scale quantum") AND "Noisy Qubits"`
Wiley	`"quantum optimization" AND (NISQ OR "noisy intermediate-scale quantum") AND ("real-world" OR "practical") AND "QAOA" AND ("toy problem" OR "toy benchmark") OR "hard problem" AND ("real-world" OR scalability)`
Taylor & Francis	`[All: quantum optimization] AND [All: noisy qubits] AND [All: nisq]`
MDPI	`"quantum optimization" AND ("NISQ" OR "noisy intermediate-scale quantum") AND "Noisy Qubits"`

3.3 Inclusion and Exclusion Criteria

Table 2 presents the inclusion and exclusion criteria for quantum optimization studies. It includes criteria such as the requirement for studies to focus on quantum optimization algorithms, incorporate assumptions about noisy intermediate-scale quantum (NISQ) devices, and include performance evaluation or benchmarking. Studies that are excluded are those that focus solely on pure quantum physics without an optimization focus, rely on ideal, noise-free hardware, lack experimental or simulation results, or deal only with toy problems (10 variables). Additionally, studies published before 2015 or in non-English languages are excluded from the selection.

3.4 Screening and Selection Process

The process of screening and selection of studies involved multiple stages to ensure methodological quality and reduce the bias. After the initial searching of databases and registers, 1242 records were sourced. After removing duplicate records and records being excluded for irrelevance or document type, 704 records were included in the screening process. Title and abstract screening then excluded 548 records that were evidently not the target type of study in quantum optimization in the NISQ era.

Table 2 Inclusion and Exclusion Criteria for Quantum Optimization Studies

Inclusion Criteria	Exclusion Criteria
Studies on quantum optimization algorithms	Pure quantum physics without optimization focus
NISQ or noisy quantum device assumptions	Ideal, noise-free hardware only
Benchmarking or performance evaluation	No experimental or simulation results
Real-world or scalable problem instances	Only toy problems (¡10 variables)
Published 2019–present	Non-English publications

Full text of the 156 articles was requested, and 121 were successfully retrieved and included in the assessment of eligibility resolved. Full texts were screened against predefined inclusion and exclusion criteria, focusing in particular on benchmark realism, noise and hardware constraints accounted for, and the amount of methodological details provided. At this level, 36 papers were rejected, as they relied only on toy benchmarks, idealized no-noise assumptions, or their experiments, and their writing was not clear enough.

Finally, a total of 85 studies met all the inclusion criteria and were included in the qualitative synthesis. The entire process of screening and study selection and the number of studies at each stage are presented in the PRISMA flow diagram (Figure 1).

3.5 Data Extraction and Analysis Procedure

A structured data extraction framework was utilized on each of the final set of studies in order to facilitate systematic analysis and minimize subjective interpretation. For each study, relevant information was consistently collected: optimization problem domain, algorithmic paradigm, problem encoding mechanism (QUBO or Ising formulation), benchmark properties, noise and hardware model assumptions, and the chosen performance metrics. Further experimental details on setup, choice of parameters and evaluation method were also taken from the source where provided.

Since we have a significant heterogeneity in studied problems, hardware platforms, benchmark designs and evaluation criteria among those studies, it was not possible to perform a meta-analysis using numerical summaries of performance estimates. Rather than being statistically pooled, the data were synthesized in a qualitative narrative way utilizing straightforward subject summation, comparative subgroups, and thematic analysis. The studies were classified according to the proposed taxonomy dimensions, allowing structured analysis and comparison of algorithms, noise models, and benchmark classes. The findings of this analysis are the foundation for the comparative tables, taxonomy figures and gap identification as presented below.

A structured data extraction form was designed to capture relevant information from each selected study. Extracted data were synthesized using qualitative and quantitative methods. Cross-study comparisons were conducted using the (Table 3). This systematic approach ensures that the review provides a comprehensive and unbiased assessment of quantum optimization beyond toy benchmarks.

The methodological quality of the selected publications was thoroughly evaluated based on a predefined scorecard assessing problem clearness, benchmark realism, noise

Table 3 Extracted attributes from selected quantum optimization studies

Attribute	Observed Values
Algorithm type	QAOA, VQE, QA, EA-QC, Bayesian QAOA
Hardware used	IBM, Rigetti, IonQ, D-Wave, simulators
Noise model	Depolarizing, readout, crosstalk
Error mitigation	DD, ZNE, pulse optimization
Problem class	Max-Cut, SVP, TSP, portfolio
Benchmark type	Toy / Hard / Real-world
Dataset	Synthetic, finance, energy
Evaluation mode	Hardware / simulator
Encoding method	QUBO, PUBO, CQM
Optimizer	COBYLA, SPSA, Bayesian optimization
Scalability test	Qubit / variable growth
Reporting quality	Partial / complete
Baseline used	Classical heuristics, tensor networks
Hardware control	Default / pulse-level
Integration level	HPC / standalone
Multi-objective	Yes / No
Constraint handling	Penalty / native
Reproducibility	Open / partial
Cost model	Reported / missing

model availability, scalability and reproducibility. For each of the criteria, a three-point scale (ranging from 0 to 2) was used to maintain consistency, objectivity and transparency when assessing the strength of the reviewed evidence shown in table 4.

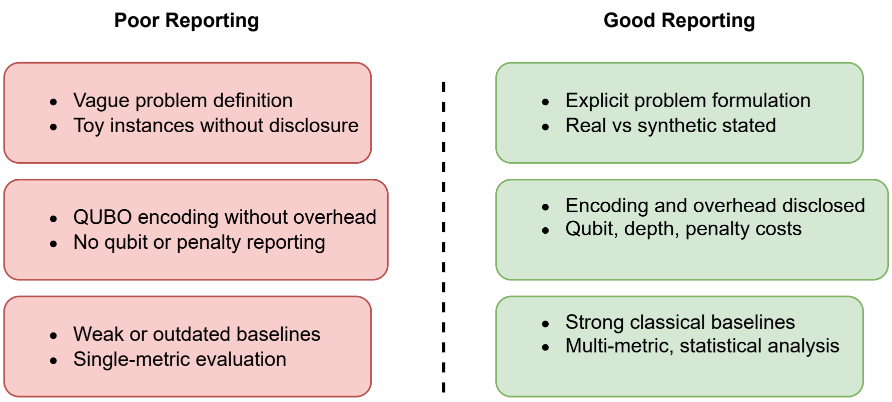

Fig. 2 Illustrative comparison of poor versus good reporting practices in quantum optimization studies

Figure 2 provides an illustrative comparison between poor and good reporting practices in quantum optimization studies. On the left, poor reporting is characterized by vague problem definitions, toy instances without proper disclosure, and incomplete

Table 4 Quality Assessment (QA) rubric used to evaluate the rigor and reliability of quantum optimization studies.

QA ID	Assessment Criterion	Scoring
QA1	Problem clearly defined (objective, constraints, problem class explicitly stated)	0 = No, 1 = Partial, 2 = Clear
QA2	Problem scale reported (original size + post-encoding size)	0 = No, 1 = Partial, 2 = Full
QA3	Algorithm and encoding fully described (reproducible formulation)	0 = No, 1 = Partial, 2 = Full
QA4	Hardware / simulator details reported (backend, noise model, calibration)	0 = No, 1 = Partial, 2 = Full
QA5	Noise treatment addressed (noise model, mitigation, or resilience by design)	0 = No, 1 = Aware, 2 = Mitigated
QA6	Baseline comparison included and fairly configured	0 = Weak, 1 = Reasonable, 2 = Strong
QA7	Evaluation metrics comprehensive (quality + cost + variance)	0 = Single metric, 1 = Limited, 2 = Multi-metric
QA8	Scalability evidence provided (trend, asymptotic, or large instance)	0 = None, 1 = Limited, 2 = Clear
QA9	Reproducibility enabled (code, seeds, configs available)	0 = No, 1 = Partial, 2 = Full
QA Total	**Maximum Score**	**18**

details on QUBO encoding, qubits, and penalty reporting. Furthermore, weak or outdated baselines and the use of single-metric evaluation are common in studies with poor reporting. On the right, good reporting practices emphasize explicit problem formulation, the distinction between real and synthetic datasets, and the full disclosure of encoding and overhead costs, including qubit, depth, and penalty costs. Strong classical baselines and a multi-metric statistical analysis are also key components for high-quality reporting in quantum optimization research. This comparison provides an example of how transparency and rigor could be enhanced in quantum optimization research.

4 Results and Synthesis

4.1 Temporal and Topical Trends

In this subsection, we study the temporal trends of the number of studies in recent years and identify the key topics that have emerged within the field of study. From 2019 to 2025, the overall number of studies has also seen a dramatic increase, indicating that research on this area is attracting more interest. The trend does over time show a regular increase, with particularly large jumps in 2023 and 2025. This trend should be better over the temporal scaling, meaning that the field has been receiving more and more attention, possibly driven by changes in research needs, technology developments or societal priorities.

Most of the studies concerning subject trends, look to focus on a few topics that are popular for some years. The sudden increase in the number of studies post 2022 however may suggest that those issues became more prevalent. Therefore researchers have been more involved with those topics. By comparing the data on topical trends with the temporal pattern, it is possible to gain some valuable insights into how various concepts have evolved over time in this field. And reveal where the gaps are in current research and suggest areas for future investigation. The temporal trends observed in the number of studies from 2019 to 2025 are summarized in Table 3

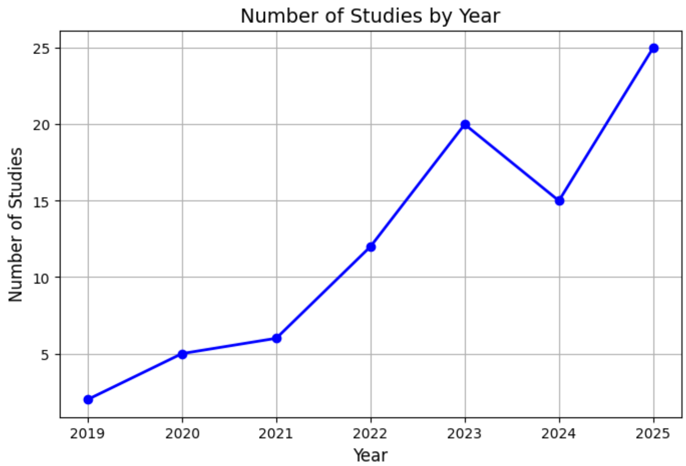

Fig. 3 Temporal Distribution of Studies from 2019 to 2025

4.2 Taxonomy of Quantum Optimization Beyond Toy Benchmarks

In this subsection, we explore the taxonomy of quantum optimization, categorizing it into three major dimensions: Noise, Algorithms, and Benchmarks. This classification goes beyond toy benchmarks and attempts to structure the various factors and components that influence the development and effectiveness of quantum optimization techniques.

- **Noise:** Quantum algorithms are often limited by noise, a feature that is present in quantum systems. Several kinds of noise such as readout errors, gate errors, crosstalk decoherence (T1/T2), deteriorate the precision and accuracy of quantum operations [37] [38] [39]. To counteract these challenges, noise mitigation and design techniques, including noise-aware design and mitigation strategies, are employed [40]. State Preparation and Measurement Errors (SPAM) occur during the preparation and measurement phases of quantum algorithms and can significantly impact the accuracy of results [41]. Effective noise management is critical for scaling quantum optimization algorithms beyond toy problems and for achieving practical real-world solutions.
- **Algorithms**: Quantum optimization is powered by several distinct algorithmic approaches. Notable among these are

- **Quantum Annealing**: A quantum computational method for optimization problems seeking the minimum energy configuration [25].
- **Variational Quantum Algorithms (VQAs)**: These types of algorithms merge classical optimizers and quantum resources for tackling optimization problems. Some of the most important methods include the Variational Quantum Eigensolver (VQE) and Quantum Approximate Optimization Algorithm (QAOA) [42].

- **Benchmarks**: Benchmarking quantum optimization is crucial for assessing the algorithm's performance across various scales:
  - **Toy Benchmarks**: These are small-scale models used to test basic optimization tasks, helping researchers understand algorithmic behavior and identify potential improvements.
  - **Real Benchmarks**: These focus on large-scale, real-world applications and application domains such as energy, finance, and telecommunications, which are illustrative rather than exhaustive examples of potential real-world optimization settings.
  - **Scaled Benchmarks**: This category involves problems scaled by factors such as the number of variables, qubits, circuit depth, runtime, or approximation ratios to better represent real-world complexity.

The taxonomy presented in this section helps in understanding the different factors that contribute to quantum optimization, highlighting areas where research is focused on overcoming noise challenges, refining algorithmic approaches, and testing them in increasingly complex benchmarks. The detailed taxonomy figure 4 provides a visual summary of these key elements. The diagram below offers a comprehensive look at the interrelations among Noise, Algorithms, and Benchmarks, allowing for a deeper understanding of the current landscape of quantum optimization beyond simple toy problems.

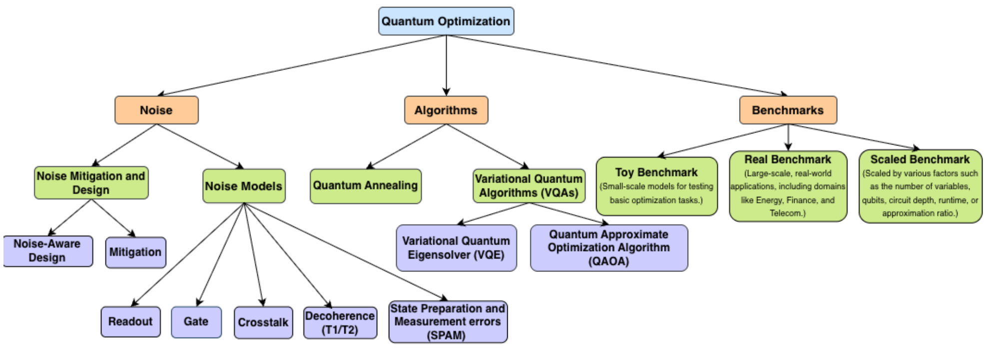

Fig. 4 Taxonomy of Quantum Optimization Beyond Toy Benchmarks

4.3 Classification by Algorithmic Approach

Quantum optimization is studied within the considered corpus of work using a wide variety of algorithmic paradigms, each corresponding to distinct assumptions regarding controllability, noise immunity and scaling in the NISQ regime [43] [44]. The following subsection seeks to categorize the reviewed studies according to their underlying algorithmic strategy and to compare and contrast their common usages, claimed benefits and potential methodological weaknesses.

The gate-based variational approaches, in particular the QAOA and related VQAs, represent the predominant class of methods in the reviewed works [45] [46]. These algorithms intersperse parameterized quantum circuits with classical optimizers to iteratively minimize a problem-dependent cost function [47] [48]. The majority of studies using gate-based techniques have considered low circuit depths due to constraints on NISQ hardware [49]. Shallow-depth realizations suppress decoherence and gate error accumulation but restrict expressivity, and performance plateaus with the increasing problem size or complexity. Gate-based methods are usually tested on small- and medium-sized combinatorial optimization problems in idealized or, at best, noise-adapted simulations [50] [51].

Quantum annealing stands as a unique optimization paradigm, which encodes problems into Ising Hamiltonians and applies analogue quantum dynamics to locate low-energy solutions [1] [52]. In the corpus, annealing methods continue to be primarily used for combinatorial optimization with QUBO or Ising problems. One of the strengths of such annealing approaches is the fact that they can tackle bigger nominal problems than gate-model devices [53].

Hybrid quantum-classical algorithms have become a practical solution to the challenges faced by existing quantum technologies [54]. They combine the QC subsystem, like variational circuits or annealing steps with more extensive classical optimization flows as shown in Fig. 5. Classical components solve problem decomposition, preprocessing, or constraint handling (post-optimization refinement) in the tested studies and quantum components solve a restricted subproblem (for instance, providing heuristic guidance) [55] [56]. This sharing of tasks helps in making the system more robust against noise and opens possibilities for running experiments on larger problem instances. But hybrid approaches make it difficult for you to assign blame. A number of the works demonstrate enhancement over classical baselines, but without a well-defined definition for the quantum component [57].

Table 5 provides a summary and a comparison among some quantum optimization algorithms (with also details of the main typical use cases and known limitations). It contains well-known algorithms such as the Quantum Approximate Optimization Algorithm (QAOA), Variational Quantum Eigensolver (VQE), and Quantum Annealing (QA) among others. We provide a brief description of the respective fields of applications: combinatorial optimization, molecular optimization, and large-scale QUBO optimization, underlining their potentiality to tackle real-world problems in different contexts such as financials, logistics, and scheduling. The table also gives an idea about the weaknesses of their algorithms in terms of scaled benchmark.

Figure 5 depicts the general procedure for runtime variational quantum optimization workflow, consisting of a fixed-problem encoding and iterative parameter optimization. The problem is then encoded to a quantum cost representation (such as QUBO or Ising), for which the parameterized quantum circuit ansatz is created. The circuit is compiled down and executed on a quantum device or simulator, potentially including error mitigation methods. Measurement results are fed to a classical optimizer that performs parameter updates. If this stopping criterion is not satisfied, then the optimized parameters are returned to the circuit, enabling a closed-loop hybrid quantum–classical optimization workflow. Hybrid quantum-classical methods are considered as a workflow pattern throughout this work and not as an isolated algorithmic class. More precisely, the name denotes optimization pipelines that include quantum subroutines as part of a classical control loop or pre/postprocessing heuristics.

Table 5: Comparison of quantum optimization algorithms

Study	Algorithm	Typical Use Case	Key Limitations
[16, 24, 34, 58–60]	Quantum Approximate Optimization Algorithm (QAOA)	Combinatorial optimization (Max-Cut, MIS, TSP, scheduling)	Noise sensitivity; parameter scaling; QUBO overhead
[4, 9, 15, 22, 26, 61]	Variational Quantum Eigensolver (VQE)	Molecular optimization and variational formulations	Circuit depth growth; classical optimization cost
[20, 62]	Fixed-Angle / Shallow QAOA	Noise-aware optimization on NISQ devices	Reduced expressibility
[15, 60]	Noise-Mitigation-Enhanced VQAs	Accuracy improvement under hardware noise	High sampling overhead; limited scalability
[12, 63]	Constraint-Preserving VQAs	Optimization with strict feasibility constraints	Problem-specific design; limited portability
[13, 25]	Quantum Annealing (QA)	Large-scale QUBO optimization (finance, logistics, MIMO)	Restricted to Ising/QUBO; analog noise
[6, 64]	Hybrid Quantum Annealing Workflows	Industry-oriented optimization with classical preprocessing	Vendor dependence; tuning complexity
[5, 65]	Coherent / Analog Ising Machines (CIM)	Industrial Ising optimization (energy systems, unit commitment)	Non-universal; special-purpose hardware
[66, 67]	Hybrid Quantum–Classical Optimization	Real-world optimization via decomposition	Classical overhead dominates runtime
[68]	Distributed Quantum Optimization (Multi-QPU)	Scaling variational optimization beyond single QPU	Communication noise and latency

Continued on next page

Table 5 – continued from previous page

Study	Algorithm	Typical Use Case	Key Limitations
[69, 70]	Evolutionary Algorithms with Quantum Operators	NP-hard problems (TSP, scheduling)	Marginal quantum advantage at small scales
[71, 72]	Grover-Based Search	Exact or approximate solutions to NP-complete problems	Extremely fragile under noise
[73–76]	Quantum Neural Networks (QNNs)	Optimization-inspired learning and decision tasks	Barren plateaus; noise sensitivity
[77]	Quantum Support Vector Machine (QSVM)	Kernel-based classification in optimization pipelines	Quantum advantage degrades under noise
[78, 79]	Fault-Tolerant Optimization Algorithms	Logical-level optimization and cryptography	Require millions of physical qubits
[1, 24]	Higher-Order (PUBO / HOBO) QAOA	Reduced qubit count via expressive encodings	Higher circuit depth; measurement cost
[80]	Tensor-Network-Assisted Quantum Optimization	Structured polynomial optimization problems	High classical computational overhead

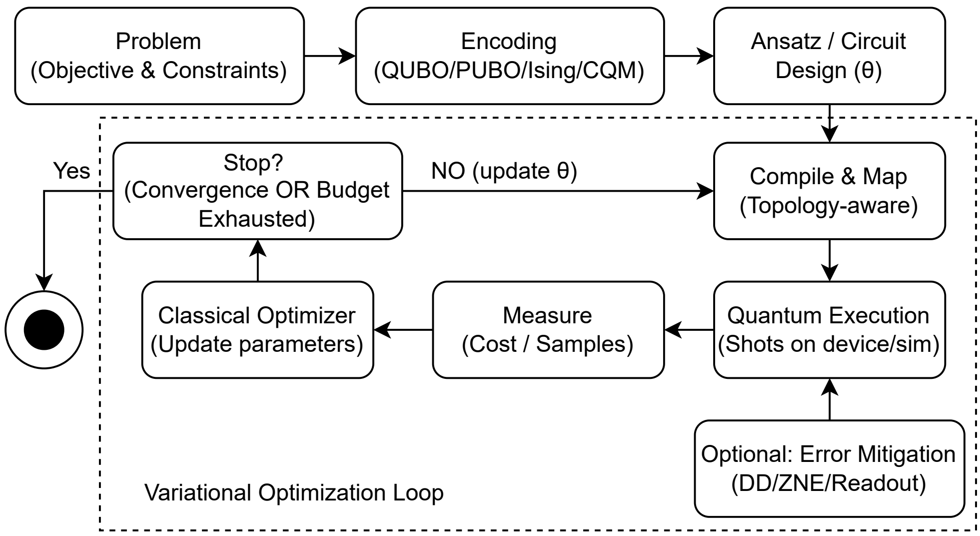

Fig. 5 Hybrid quantum–classical workflow for NISQ-era quantum optimization

4.4 Classification by Hardware and Noise Model

4.4.1 Simulator-based Studies

Many studies rely on noise-free or noisy simulators to evaluate algorithmic behavior. Simulations are, however, commonly known to better perform in comparison to the actual hardware. Noise models usually incorporated depolarizing noise, readout errors, and crosstalk estimates. However simulator has quite low fidelity compared to real device behavior [40] [26] [20] [22] [81].

4.4.2 Real NISQ Hardware Experiments

Experiments are carried out on IBM, Rigetti, IonQ, and D-Wave machines. These experiments offer realistic insight into the performance, subject to restrictions in qubit numbers, connectivity, and calibration variation. Hardware experiments typically exhibit worse solution qualities than simulations, due to the introduced noise and compilation overhead [31] [32] [61].

4.4.3 Noise-aware vs Noise-agnostic Approaches

Noise-independent works are based on the case of perfect implementation and neglect hardware noise. In contrast, noise-aware approaches incorporate:

- Error-mitigation procedures (ZNE, dynamical decoupling),
- Noise-adaptive parameter tuning,
- Compilation-aware circuit optimization.

The noise-aware designs consistently outperform the noise-agnostic counterparts on real hardware, validating the importance of hardware-adaptive optimization [82] [83] [84] [85].

4.5 Encoding Strategies and Problem Mapping

Encoding the problem is the most important interface between an abstract optimization model and an implementable quantum program. In the NISQ era, therefore, the choice of encoding strategy directly impacts scaled benchmark optimization beyond toy benchmarks [86] [87] [88]. This subsection consolidates encoding techniques found in the SLR corpus and discusses their consequences with respect to problem realism and experimental veracity.

4.5.1 Problem Reformulation and Discretization

Unfortunately, most of the optimization problems encountered in practice are not defined in quantum-compatible expressions. Therefore, the reviewed literature often starts with problem reformulation, which can include the discretization of continuous variables, the relaxation of hard constraints, or a reduction to canonical optimization representations [89]. Although these reformulations allow for a quantum implementation, they generally lead to either approximation errors or exponential blowup of the problem size. There are a few studies that recognize that the granularity of discretization is in fact a trade-off that exists between solution accuracy and hardware practicality. The downside of this trade-off is seldom studied in a systematic way, which makes it hard to interpret advances that are claimed.

4.5.2 QUBO and Ising Encodings

The encoding landscape is dominated by Quadratic Unconstrained Binary Optimization (QUBO) and Ising model formulations, as both are directly compatible to be solved on gate-based variational algorithms or quantum annealers [60]. These encodings offer a common mathematical language but often need to add auxiliary variables for higher-order terms or constraints. The literature surveyed indicates that the QUBO/Ising encoding may require significant qubit overhead, particularly for constrained or structured problems [90] [2]. Furthermore, tuneability of the successive penalty parameter is a common issue, and various studies have reported sensitivity towards penalty weights, which may severely influence solution quality as well as convergence characteristics [91].

4.5.3 Constraint Handling Strategies

Handling constraints is a major difference between toy benchmarks and real-life optimization problems. Over the SLR corpus, trade-offs are handled through penalization (also through slack terms or classical preprocessing hybrid). Penalty-based encoding is still the most popular because it is easy to implement, but the noise sensitivity and the ruggedness of the energy landscape are amplified [92]. Mixed approaches, where constraint satisfaction is shifted to classical solvers before a quantum optimization step, lead to slightly better stability at the expense of transparency between the two contributions. This tension emphasizes the necessity of constraint-aware quantum encodings of higher principles.

4.5.4 Encoding Pipeline and Execution Flow

The basic encoding process, as reported by the reviewed studies, was described in Figure 6. It is not just a linear development but the mapping between classical and quantum variables, and handling constraints can be fed directly to quantum encoding stage as well as refined by feedback from classical evaluation in multiple rounds of feedback influencing reformulation choice.

This pipeline serves to highlight that encoding decisions are not just technical details but key design choices with a formative impact on the results of one's experiments. Research that more directly traces this pipeline is also more easily reproduced and interpretable.

Figure 5 illustrates the iterative runtime interaction between classical optimization routines and quantum executions, whereas Figure 6 focuses on the offline problem formulation and encoding pipeline that precedes any algorithmic execution.

4.5.5 Comparative Analysis of Encoding Strategies

Table 6 presents a comparison of different types of encoding in the context of quantum optimization. Encoding plays a vital role in transcribing classical problems to

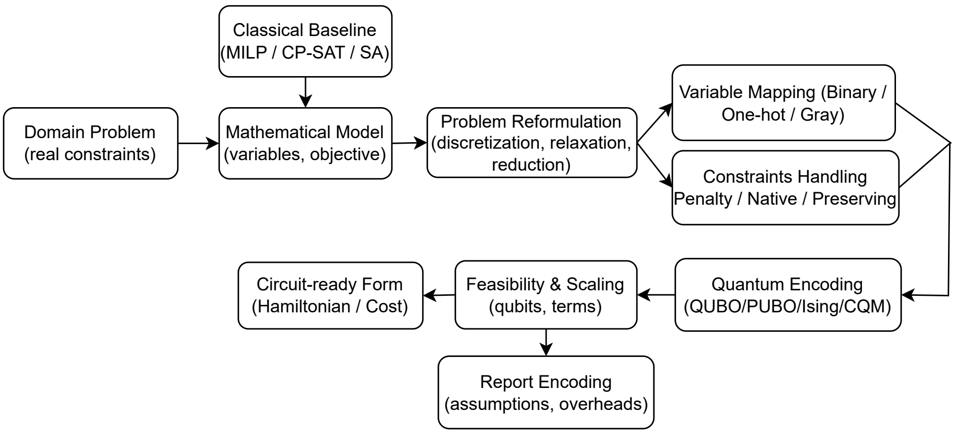

Fig. 6 Encoding pipeline for quantum optimization beyond toy benchmarks

quantum devices, impacting the performance and scalability of the quantum optimization algorithms. Various encoding methods and their characteristics, advantages and disadvantages are summarized in the following table. We hope to shed some light on these strategies and their associated advantages, trade-offs and how they respond to different quantum optimization problems.

Table 6: Encoding strategies for quantum optimization

Study	**Encoding Strategy**	**Description**	**Advantages**	**Limitations / Trade-offs**	**Used In**
[6, 13, 16, 25, 58–60]	QUBO / Ising	Quadratic unconstrained binary optimization mapped to Ising Hamiltonians	Compatible with most NISQ devices and annealers; mature tooling	High qubit and constraint overhead; scalability limits	QAOA, QA, hybrid annealing
[1, 24, 34]	Higher-Order Binary Optimization (PUBO / HOBO)	Direct encoding of higher-order cost functions	Lower qubit count; preserves structure	Deeper circuits; complex measurements	HOBO-QAOA, polynomial optimization

Continued on next page

Table 6 – continued from previous page

Study	Encoding Strategy	Description	Advantages	Limitations / Trade-offs	Used In
[12, 63]	Constraint-Preserving Encoding	Feasibility encoded directly into ansatz or Hamiltonian	Eliminates penalties; guarantees feasibility	Problem-specific; limited generality	Routing, scheduling
[73–76]	Angle / Feature Map Encoding	Real-valued features mapped to rotation angles	Compact; efficient for ML-style optimization	Noise sensitivity; limited expressibility	QNNs, QSVM
[93]	Graph-Based Encoding	Graph embeddings tailored to hardware topology	Reduced SWAP overhead; hardware-aware	Topology-dependent; poor portability	Graph QAOA
[8, 65]	Physical / Analog Ising Encoding	Variables mapped to physical degrees of freedom	Noise-tolerant; naturally scalable	Non-universal; limited flexibility	CIMs, photonic solvers
[69, 70]	Hybrid Classical–Quantum Encoding	Problem decomposed into classical and quantum parts	Scales to real-world instances	Classical preprocessing dominates runtime	Hybrid TSP
[66, 68]	Distributed Encoding	Problem partitioned across multiple QPUs	Scales beyond single-device limits	Communication noise; orchestration cost	Distributed QAOA
[79, 94]	Logical / Fault-Tolerant Encoding	Encoded at logical qubit level	Asymptotic noise robustness	Infeasible on near-term hardware	FT optimization
[80]	Tensor-Network Encoding	Mapped to tensor-network structures	Efficient correlation capture	High classical computational cost	Tensor-QAOA

4.6 Comparative Analysis Across Studies

Table 7 8 9 provides a comparison of several quantum optimization works, discussing main performances characteristics, hardware configurations and problems-solving techniques. The table categorizes and arranges the studies according to problem, encoding approach, baseline methodology, and related hardware employed. In addition, we show the particular metrics tested for in each work, like cost function, accuracy, low power (LP), and speed computation, to give an overall picture of the variety of methods developed so far. This comparison involves studies using various quantum algorithms such as QAOA, QNN, and VQE and experiments with simulators or real quantum devices. In a nutshell, the table provides valuable information for both strengths and limitations of different classes of quantum optimization algorithms across different areas, from small-molecule-scale to network-based large-scale problems.

The overall synthesis across the three comparison dimensions shows that it is the methodological choices rather than the novelty of the algorithm that have a significant impact on reported performance. If studies involve realistic encoding-based noise-aware evaluation and strong baselines, one often observes a more conservative but solid result. On the other hand, optimistic statements often accompany unrealistic assumptions and little to no benchmarking stringency. This integrative scoping review serves as the empirical basis for identifying research needs and formulating pragmatic directions of future research, which are outlined in the next section.

Table 7: Main comparison matrix of quantum optimization studies beyond toy benchmarks

Study	Problem	Problem Size	Encoding	Solver / Algorithm	Hardware / Simulator	Baselines	Metrics	Key Results
[1]	Hadamard matrix construction	Medium–Large	High-order Hamiltonian	QAOA	Simulator	Annealing methods	Error rate	Better scaling than QA
[6]	Portfolio optimization	Medium–Large	CQM/ QUBO	Quantum annealing	D-Wave Advantage	Classical MO solvers	Pareto quality	Outperforms classical heuristics
[8]	Unit commitment (power grid)	Large	QUBO $\rightarrow$ Ising	Hybrid CIM optimization	CIM hardware	MILP solvers	Cost, feasibility	Demonstrated industrial feasibility
[9]	Differential equations	≤ 3 variables	Feature maps	Variational solver	IBM QPU	Classical DE solver	Approximation error	Works only for tiny systems
[13]	Supply-chain VRP	Large	QUBO/ PUBO	Hybrid QA workflow	D-Wave hybrid	Classical heuristics	Cost, feasibility	Real industrial deployment
[20]	Lattice SVP	Medium dimension	Angle encoding	Fixed-angle QAOA	Simulator	Classical sieving	Success rate	Asymptotic improvement predicted
[24]	Continuous optimization	Medium	PUBO	QAOA	Simulator	QUBO QAOA	Energy gap	PUBO uses fewer qubits
[25]	MIMO detection	Large (48×48)	QUBO	Quantum annealing	D-Wave 2000Q	Sphere decoding	BER, runtime	First large-scale QC success
[26]	Molecular vibrations	Small molecules	Second-quantized	VQE (CHC, UVCC)	Simulator	Classical quantum chemistry	Energy error	CHC reduces noise sensitivity

Table 8: Main comparison matrix of quantum optimization studies beyond toy benchmarks (Cont.)

Study	Problem	Problem Size	Encoding	Solver / Algorithm	Hardware / Simulator	Baselines	Metrics	Key Results
[28]	VLSI placement	Small circuits	Min-cut encoding	VQE/QML	Simulator	Kernighan–Lin	Cut quality	Comparable on small benchmarks
[34]	TSP (NP-hard)	Medium	HOBO	QAOA	Gate-based NISQ	QUBO-QAOA	Qubit count, depth	Drastic qubit reduction
[58]	Max-Cut, MIS	Small graphs	Ising/QUBO	BO-QAOA	Simulator	Standard QAOA	Convergence, energy	Faster convergence than vanilla QAOA
[59]	MIS, scheduling	Small graphs	Ising	QABOA (2 mixers)	Simulator	Classical BO	Convergence speed	Better exploration–exploitation
[60]	Max-Cut, MIS	Small graphs	QUBO	QAOA + QEM	NISQ model	No-mitigation QAOA	Energy error	QEM improves accuracy at cost
[65]	Ising optimization (analog)	Large	Physical Ising	CIM / photonic solver	Photonic hardware	Simulated annealing	Energy, speed	Orders-of-magnitude speedup
[68]	Distributed optimization	Medium–Large	Circuit partitioning	Distributed QAOA/VQE	Multi-QPU (conceptual)	Single-QPU VQE	Fidelity, latency	Scales theoretically; noise dominates
[69]	Traveling Salesman Problem	$\leq$ 15 cities	Binary / Gray code	EA + quantum operators	Hybrid QC simulator	Classical EA	Tour length	Outperforms classical EA at small size

Table 9: Main comparison matrix of quantum optimization studies beyond toy benchmarks (Cont.)

Study	Problem	Problem Size	Encoding	Solver / Algorithm	Hardware / Simulator	Baselines	Metrics	Key Results
[72]	k-Clique (NP-complete)	Small graphs	State-preparation	Grover-based search	IBM QPU	Classical exact solvers	Success probability	Rapid degradation under noise
[74]	Road-network optimization	Small graph	QUBO	QNN / VQA	Simulator	Classical heuristics	Cost, accuracy	Quantum comparable at small scale
[79]	Logical routing	Very large	Braiding paths	Surface-code scheduler	FT model	Naive routing	Latency	Orders-of-magnitude gains
[80]	Polynomial optimization	Medium	Non-orthogonal	VQE	Gate-based NISQ	QUBO-VQE	Energy, qubits	Multiple vars per qubit
[95]	Quantum networks	Medium	Network-state	VQO	IBM QPU / sim	Classical routing	Nonlocality	Noise-aware VQO superior
[96]	Max-Cut	Medium graphs	Graph embeddings	GNN-assisted QAOA	Simulator	Random init QAOA	Energy, convergence	Hybrid ML improves training

4.7 Classification by Benchmark Type

Benchmark selection plays a central role in shaping the interpretation and generalizability of results in quantum optimization research. In the studies reviewed, the literature surveys show that the choice of benchmarks shapes not only enacted performance results, but it also embodies assumptions about scalability, noise tolerance, and applicability. This subsection classifies the benchmarks employed in the SLR corpus and analyses how the realism of the benchmark influences methodological strength and comparative credibility.

4.7.1 Toy vs. Beyond Toy Benchmarks

Many papers still rely on toy benchmarks, like small Max-Cut instances, randomly generated QUBO problems, or low-dimensional synthetic data sets [24][97]. They are appealing because they are easy to understand, analytically tractable, and run on simulators or even limited hardware illusions. Accordingly, toy problems have served as a useful mechanism for algorithm development and methodology investigation. Nonetheless, statistical comparison shows that many overrely on benchmarks and hence lack external validity in their claimed performance. Toy domains typically don't have real constraint structures, diverse costs, and the scaling behavior that is found in actual optimization problems [53]. Further, as they are relatively small, classical solvers can easily find near-optimal solutions with a fully exhaustive search, making the reported quantum speedups less relevant.

In contrast, benchmarks that move beyond toy settings—including scaled synthetic instances and application-inspired problems—introduce additional complexity through larger problem sizes, structured constraints, and noise sensitivity [98]. Indeed, experiments with such benchmarks tend to result in more cautious performance reports but provide stronger evidence about practical feasibility. Of particular relevance to the work described here, beyond-toy simulations often require hybrid quantum–classical algorithms that typically combine advanced encoding techniques and measurement-level noise-aware simulations. The review also suggests that the move from toy to beyond-toy benchmarks is somewhat lopsided among algorithmic paradigms. However, current gate-based variational methods are mostly evaluated on toy or weakly scaled instances, while hybrid and annealing-based approaches often interact with structured benchmarks [99] [46]. This imbalance highlights the importance of greater usage of beyond-toy benchmarks across all algorithmic classes. In total, the dichotomy of toy versus beyond-toy benchmarks is a key dimension along which to evaluate the state of maturity and trustworthiness of quantum optimization research. Advancing towards realistic benchmarking is important for bridging the gap between theoretical promise and practical usability in the NISQ era.

Table 10 summarizes the "toy vs. beyond-toy" benchmarks that appear in quantum optimization research. It classifies various studies according to the size of instances, having realistic constraints (or not) and hardware/noise evaluation. The table differentiates small-scale problems, which are mostly based on synthetic data sets and idealized models, from large-scale ones with real-world challenges involving hard constraints and hardware considerations. Works with quantum simulators or real quantum processors are compared considering their evaluation of noise and use of error mitigation tools. The table evidences the evolution of quantum optimization computation from toy problems (small instances with few constraints) to large-scale, real-world applications, showing the growing complexity and inclusion of noise-aware protocols and error mitigation techniques in deployed quantum hardware.

Table 10 Toy vs. Beyond-Toy Benchmarks in Quantum Optimization

Study	**Instance Size Threshold**	**Presence of Realistic Constraints/-Datasets**	**Hardware/Noise Evaluation**
[11], [14], [43]	Small ($\leq$ 50 variables/qubits)	Synthetic datasets, idealized problems	Simulated, no hardware or noise evaluation
[46], [77], [86]	Small ($\leq$ 50 qubits)	No constraints or idealized constraints	Simulated models only, no noise mitigation
[40], [74]	Small to medium ($\leq$ 50 qubits)	No hard constraints	Simulated, no noise considerations
[61], [65], [66]	Medium to large ($\geq$ 100 qubits)	Real-world constraints (e.g., resource limits)	Evaluated on real quantum hardware (IBM QPU, D-Wave) with noise model
[68], [76]	Medium to large ($\geq$ 100 variables)	Hard constraints (budget, resource allocation)	Hardware-based testing, noise mitigation techniques included
[15], [20], [59], [100]	Large ($\geq$ 100 qubits)	Real datasets from industry or academia (e.g., finance, logistics)	Tested on quantum hardware, noise-aware evaluation
[9], [25], [28]	Large to very large ($\geq$ 100 qubits)	Real-world datasets, industry-standard constraints	Hardware evaluation with error mitigation and noise management
[6], [12], [13]	Very large ($\geq$ 100 qubits)	Real-world complex problems with hard constraints	Extensive hardware testing, noise-aware protocols, scalable to large instances

4.8 Evaluation Metrics and Benchmarking

Benchmarking quantum optimization algorithms in the NISQ era is heavily dependent on the choice of performance metrics and benchmark protocols. In contrast to classical optimization, where uniform metrics and benchmark suites are well accepted now, the literature covered varies significantly in terms of performance measurements and reporting and interpreting results.

There is a large variety of metrics used for good performance in the reviewed studies. The common objectives are quality of the solution (such as approximation ratio or optimality gap), success probability, average energy value, and runtime or circuit depth in scaled benchmark [101]. In hardware experiments, other metrics such as the shot count, the execution time or the error rates are sometimes provided. Benchmarking methodologies are highly disparate in the literature. Toy benchmarks and instances generated at random are still highly prevalent, especially for simulation-based research. Although these benchmark problems allow for controlled experimentation, they are standardly structurally simple to be representative of real-world optimization. Classical baselines are included inconsistently. Some also compare quantum algorithms to classical state-of-the-art heuristics, while some use very basic baselines or do not even consider a classical comparison. Lack of strong classical baselines undermines the interpretability of observed enhancement and complicates interpretation of true quantum contribution [81]. There is an important difference between noise-aware and idealized evaluation methods. Noise-aware works explicitly include hardware noise models, error mitigation protocols, or runs on actual NISQ machines [102] [100]. Contrary, idealized evaluations consider noise-free or simplified error models.

The choice of classical baseline is crucial for shaping the achieved performance improvement. Comparing weak or non-competitive classical heuristics can significantly inflate a claimed quantum advantage, while the best-known classical solvers might be able to obtain comparable or better results. Therefore, any claims of advantage or efficiency should be taken in the context of how impressive the baseline being compared to is and a rigorous comparison with classical optimization techniques.

Table 11 summarizes the key performance metrics commonly used in quantum optimization studies, categorized into five main areas: solution quality, computational cost, hardware efficiency, Robustness, and scalability. The table tabulates other metrics like approximation ratio, energy error, number of shots, circuit depth fidelity, and problem size along with a brief description and the typical usage in areas like quantum optimization. Those metrics are essential to assess the performance of quantum optimization algorithms and hence their fitness for various application contexts, such as combinatorial optimization, scheduling, facility location, or quantum chemistry. Sample studies for each type of metric are also given to illustrate their utility in the quantum optimization landscape. This extensive comparison is intended as a road map for furthering quantum algorithms with the specific aspiration to optimize based on such parameters.

Table 11: Performance metrics used in quantum optimization studies.

Category	**Metric**	**Description**	**Used For**	**Representative Studies**
Solution Quality	Approximation ratio	Ratio between obtained and optimal solution	QAOA, QA	Max-Cut; portfolio optimization
	Energy error	Deviation from ground-state energy	VQE, QAOA	Quantum chemistry; Ising models
	Objective value	Final cost-function value	QUBO / PUBO	Scheduling; routing

Continued on next page

Table 11 continued from previous page

Category	Metric	Description	Used For	Representative Studies
	Constraint violation rate	Fraction of infeasible solutions	Constrained optimization	Facility location; finance
Computational Cost	Number of shots	Circuit execution count	Hardware experiments	NISQ studies
	Wall-clock runtime	End-to-end execution time	Hybrid workflows	HPC–QC integration
	Classical optimizer calls	Number of optimization iterations	Variational algorithms	QAOA; VQE
Hardware Efficiency	Circuit depth	Number of gate layers	Noise sensitivity analysis	Hardware benchmarks
	CNOT / SWAP count	Two-qubit gate overhead	Compilation analysis	Mapping optimization
	Qubit count	Number of active qubits	Scalability analysis	Clique; SVP
Robustness	Fidelity	Correctness of quantum state	Error mitigation	Dynamical decoupling
	Success probability	Probability of correct outcome	Search and annealing	Grover; QA
	Output variance	Stability across repeated runs	Noisy execution	NISQ hardware
Scalability	Problem size	Number of decision variables	Benchmarking	TSP; Max-Cut
	Qubit scaling	Growth of qubits with problem size	EFTQC analysis	Large circuits
	Compilation time	Preprocessing overhead	Compiler evaluation	QAOA compilers

5 Discussion and Future Research Directions

A dominant gap across the literature is the continued reliance on toy or weakly scaled synthetic benchmarks. Although they provide a suitable platform for extensive experimentation, these benchmarks do not capture the structural richness, diverse constraints, and sensitivity to noise of real-life optimization problems. Furthermore, due to the lack of a set of standard benchmark suites, it is not easy to compare different approaches across studies. Noise and the importance of its modelling Despite noise directly being part of the NISQ regime, it is ignored or oversimplified in many reviewed works. This discrepancy between assumptions made in evaluation and real hardware behavior introduces overly optimistic performance claims that do not hold for practical devices. The problem of encoding becomes a major barrier for the scalability of quantum optimization. To the contrary, many investigations underestimate qubit overhead, penalty parameter sensitivity and constraint-driven landscape distortion. Constraint handling is often treated as an implementation detail rather than a core research challenge. Hybrid quantum–classical approaches are more commonly deployed, but the performance gains reported on them often rely on classical preprocessing or postprocessing modifications. In a lot of cases, the marginal gain from the quantum element is unknown or unmeasured.

The community could benefit from the creation of open, standardized, and application-driven benchmark repositories with reference to classical baselines and well-established evaluation protocols. It would be desirable for future works to employ hardware-compatible noise models, provide calibration details of the used device, and test robustness under different levels of noise. Simulator-to-hardware benchmarking should become part of common practice. Constraint-aware and resource-effective encoding schemes are necessary, which directly trade off expressivity, qubit cost, and noise resilience. There is a dearth of comparative studies that examine encoding trade-offs among different classes of problems. Transparency in performance attribution whilst highlighting the contribution of quantum subroutines should be given priority in future work, including ablation studies. Accurate representation of classical and quantum computational difficulty is crucial to soundly assessing performance. The roadmap outlined in this section targets improvements in methodological rigor, benchmarking realism, and evaluation maturity, rather than near-term claims of quantum advantage.

Table 12 summarizes the map from identified research gaps in quantum optimization to associated future research directions and expected impacts. These gaps are classified into different research areas like standard benchmarks, multi-objective optimization and noise handling in the table. Each gap is coupled with a forward-looking proposal, including the creation of open benchmark suites, the development of Pareto-aware quantum solvers and the standardization of APIs for quantum computing accelerators. The anticipated effects of these directions are discussed as well, drawing again the importance not only towards solution quality but also robustness, scalability and fair comparisons. This mapping serves as a road map for addressing current challenges in quantum optimization, promoting large-scale adoption, and advancing the field toward more effective and reliable solutions.

Filling this gap is crucial in taking quantum optimization from adventurous demonstrations to work of practical relevance with a sound scientific basis. If algorithmic innovation, realistic benchmarking and transparent reporting are successfully aligned in future work, the extent to which quantum methods deliver advantages for hard optimization problems within NISQ constraints can be more credibly assessed.

6 Threats to Validity

This section discusses potential threats that may affect the validity of the findings reported in this systematic review.

Table 12 Mapping of research gaps to future research directions

Gap ID	Research Gap	Future Research Direction	Expected Impact
G1	Lack of standardized hard benchmarks	Create open benchmark suites with noise profiles	Comparable evaluation
G2	Missing full cost models	Define cost-per-solution metrics	Realistic performance claims
G3	Fragmented noise handling	Unified noise-aware optimization stacks	Robust performance
G4	No co-design of compilation and tuning	Joint compiler-optimizer frameworks	Higher solution quality
G5	Limited multi-objective optimization	Develop Pareto-aware quantum solvers	Industrial relevance
G6	Weak baseline comparisons	Mandatory quantum-inspired baselines	Fair advantage claims
G7	Poor HPC integration	Standard QC accelerator APIs	Large-scale adoption
G8	Underreported control effects	Open calibration datasets	Reproducibility
G9	Weak NISQ to FTQC transition path	EFTQC readiness levels	Roadmap to scaling
G10	Encoding sensitivity ignored	Encoding benchmarking protocols	Fair comparisons

6.1 Selection Bias

Bias of selection itself may be introduced from digital libraries used, search strings applied and inclusion criteria. Despite the use of several large databases (ACM, IEEE, SpringerLink, ScienceDirect, arXiv and Google Scholar), it is likely that relevant studies were not included due to variations in indexation or terminology as well as access issues. We thus used:

- Broad keyword combinations,
- Boolean search operators,
- Backward and forward snowballing.

However, the final set of selected papers may not represent the complete body of existing research.

6.2 Publication Bias

Publication bias is a recognized issue in systematic reviews whereby studies with positive outcomes are more likely to be published than those showing no or ambiguous results. The claimed speedup in quantum computing (quantum hypothesis) may therefore be exaggerated. A part of the conflict was mitigated by:

- Including arXiv preprints,
- Based on the results of studies with potential bias or false-negative findings,
- Avoiding venue-based filtering.

Nonetheless, unpublished negative findings and industry survey results are not well represented.

6.3 Reproducibility Concerns

Reproducibility is a major concern in quantum computing research due to:

- Rapidly evolving hardware,
- Calibration-dependent performance,
- Limited availability of source code,
- Incomplete reporting of experimental parameters.

Many studies do not provide sufficient details regarding:

- Qubit layout,
- Noise profiles,
- Shot counts,
- Compilation strategies.

This limits the ability to independently reproduce results and validate claims.

7 Conclusion

This systematic review examined the current status of quantum optimization research by this community, with particular emphasis on works that go beyond toy benchmarks and idealized assumptions. An analysis of a number of works allowed us to systematize the synthesis with algorithmic methodologies, benchmark types, encoding techniques, hybrid quantum-classical routines, evaluation choices and reproducibility patterns. Unlike previous surveys that predominantly focus on algorithmic advancement, the focus of this survey is the methodological settings under which performance gains are reported and how much such gains carry over to practical NISQ regimes.

Our findings suggest a discrepancy between practice and research: Hybrids and quantum optimization algorithms are well represented in our collected works, while simultaneously they appear unevenly considered by prospective evaluation. A large percentage of papers still depend on small-scale or artificial benchmarks, simplified noise models, non-robust classical baselines, and incomplete experimental details. The reproducibility estimation we propose in this work extends the observation that important information encoding parameters, optimizer configuration, random seeds, and code is often not reported all over again (preventing validation by a third party) and are not commonly available among studies.

The taxonomy and the comparative synthesis synthesized in this review clarify how benchmark realism, encoding overhead, noise sensitivity, and hybrid execution interact to influence reported results. Crucially, approaches that use more realistic baselines and account for the noise in evaluation results often produce conservative but believable estimates of performance, reiterating how methodology is above all essential when taking into account isolated performance claims. The identified research gaps indicate the usefulness of standardized benchmarks and resource-aware encoding frameworks, transparent performance attribution in hybrid approaches and community-driven reproducibility standards.

By consolidating evidence across algorithmic paradigms, encoding strategies, benchmark types, and evaluation practices, our review establishes a methodology toward the quantification of quantum optimization research beyond toy benchmarks. Emphasizing benchmarking realism, reproducibility, and evaluation rigor, the proposed taxonomy and synthesis aim to support more transparent, comparable, and practically grounded studies, thereby contributing to a more credible assessment of quantum optimization methods in the NISQ era.

References

[1] Suksmono, A.B.: A quantum approximate optimization method for finding hadamard matrices. Scientific Reports **15**(1), 33254 (2025)

[2] Saxena, P., Sabek, I., Spedalieri, F.: Constrained quadratic model for optimizing join orders. In: Proceedings of the 1st Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications, pp. 38–44 (2024)

[3] Serrano, M.A., Cruz-Lemus, J.A., Perez-Castillo, R., Piattini, M.: Quantum software components and platforms: Overview and quality assessment. ACM Computing Surveys **55**(8), 1–31 (2022)

[4] Smith, K.N., Ravi, G.S., Murali, P., Baker, J.M., Earnest, N., Javadi-Cabhari, A., Chong, F.T.: Timestitch: Exploiting slack to mitigate decoherence in quantum circuits. ACM Transactions on Quantum Computing **4**(1), 1–27 (2022)

[5] Fu, W., Xie, H., Chen, C., Bie, Z.: Quantum-embedded robust optimization for resilience-constrained unit commitment. IEEE Transactions on Power Systems (2025)

[6] Kuo, S.-Y., Lee, K.-L., Chou, Y.-H., Shen, J.-Y., Kuo, S.-Y.: Quantum annealing for bi-objective weighted portfolio optimization in real-world financial markets. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 2441–2448 (2025)

[7] Paler, A., Oumarou, O., Basmadjian, R.: On the realistic worst-case analysis of quantum arithmetic circuits. IEEE Transactions on Quantum Engineering **3**, 1–11 (2022)

[8] Ling, J., Zhang, Q., Geng, G., Jiang, Q.: Hybrid quantum annealing decomposition framework for unit commitment. Electric Power Systems Research **238**, 111121 (2025)

[9] Schillo, N., Sturm, A.: Variational quantum algorithms for differential equations on a noisy quantum computer. IEEE Transactions on Quantum Engineering

(2025)

[10] Islam, I., Jha, V., Thomas, S., Egan, K.F., Nobel, A., Kim, S., Chaudhary, M., Ogundele, S., Kneidel, D., Phillips, B., *et al.*: Quantum circuit synthesis using fuzzy-logic-assisted genetic algorithms. Algorithms **18**(4), 178 (2025)

[11] Hsu, N.-W., Wang, C.-C., Hsu, C.-H., Tu, C.-H., Hung, S.-H.: Toward cost-effective quantum circuit simulation with performance tuning techniques. Connection Science **36**(1), 2349541 (2024)

[12] He, Z., Shaydulin, R., Herman, D., Li, C., Raymond, R., Sureshbabu, S.H., Pistoia, M.: Parameter setting heuristics make the quantum approximate optimization algorithm suitable for the early fault-tolerant era. In: Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design, pp. 1–7 (2024)

[13] Weinberg, S.J., Sanches, F., Ide, T., Kamiya, K., Correll, R.: Supply chain logistics with quantum and classical annealing algorithms. Scientific Reports **13**(1), 4770 (2023)

[14] Ugya, A.Y., Meguellati, K.: Quantum technology a tool for sequencing of the ratio dss/dna modifications for the development of new dna-binding proteins. Egyptian Journal of Basic and Applied Sciences **9**(1), 308–323 (2022)

[15] Jose, S.T., Simeone, O.: Error-mitigation-aided optimization of parameterized quantum circuits: Convergence analysis. IEEE Transactions on Quantum Engineering **3**, 1–19 (2022)

[16] Lv, L., Yan, B., Wang, H., Ma, Z., Fei, Y., Meng, X., Duan, Q.: Using variational quantum algorithm to solve the lwe problem. Entropy **24**(10), 1428 (2022)

[17] Peduri, A., Bhat, S., Grosser, T.: Qssa: an ssa-based ir for quantum computing. In: Proceedings of the 31st ACM SIGPLAN International Conference on Compiler Construction, pp. 2–14 (2022)

[18] Ittah, D., Häner, T., Kliuchnikov, V., Hoefler, T.: Qiro: A static single assignment-based quantum program representation for optimization. ACM Transactions on Quantum Computing **3**(3), 1–32 (2022)

[19] Fakhimi, R., Validi, H.: Quantum approximate optimization algorithm (qaoa). In: Encyclopedia of Optimization, pp. 1–7. Springer, ??? (2023)

[20] Prokop, M., Wallden, P.: Heuristic time complexity of nisq shortest-vector-problem solvers. arXiv preprint arXiv:2502.05284 (2025)

[21] Eichhorn, D., Pett, T., Osborne, T., Schaefer, I.: Quantum computing for feature model analysis: Potentials and challenges. In: Proceedings of the 27th ACM

International Systems and Software Product Line Conference-Volume A, pp. 1–7 (2023)

[22] Meng, F.-X., Li, Z.-T., Yu, X.-T., Zhang, Z.-C.: Quantum circuit architecture optimization for variational quantum eigensolver via monto carlo tree search. IEEE Transactions on Quantum Engineering **2**, 1–10 (2021)

[23] Palsberg, J.: Toward a universal quantum programming language. XRDS: Crossroads, The ACM Magazine for Students **26**(1), 14–17 (2019)

[24] Stein, J., Chamanian, F., Zorn, M., Nüßlein, J., Zielinski, S., Kölle, M., Linnhoff-Popien, C.: Evidence that pubo outperforms qubo when solving continuous optimization problems with the qaoa. In: Proceedings of the Companion Conference on Genetic and Evolutionary Computation, pp. 2254–2262 (2023)

[25] Kim, M., Venturelli, D., Jamieson, K.: Leveraging quantum annealing for large mimo processing in centralized radio access networks. In: Proceedings of the ACM Special Interest Group on Data Communication, pp. 241–255. ACM, ??? (2019)

[26] Somasundaram, R., Jayaharish, R., *et al.*: Quantum computing for molecular vibrational energies: A comprehensive study. Materials Today Quantum **6**, 100031 (2025)

[27] Saini, R.K.: Quantum machine learning: Combining quantum algorithms with classical ai techniques for improved learning models. quantum **10**, 14

[28] Turtletaub, I., Li, G., Ibrahim, M., Franzon, P.: Application of quantum machine learning to vlsi placement. In: Proceedings of the 2020 ACM/IEEE Workshop on Machine Learning for CAD, pp. 61–66 (2020)

[29] Griend, A., Nurminen, J.K.: Quantmark: A benchmarking api for vqe algorithms. IEEE Transactions on Quantum Engineering **3**, 1–6 (2022)

[30] Marzoug, N., Halab, K., El Meslouhi, O., Abou Elassad, Z.E., Akhloufi, M.A.: Quantum-enhanced dual-backbone architecture for accurate gastrointestinal disease detection using endoscopic imaging. BioMedInformatics **5**(3), 51 (2025)

[31] Kusyk, J., Saeed, S.M., Uyar, M.U.: Survey on quantum circuit compilation for noisy intermediate-scale quantum computers: Artificial intelligence to heuristics. IEEE Transactions on Quantum Engineering **2**, 1–16 (2021)

[32] Cheng, B., Deng, X.-H., Gu, X., He, Y., Hu, G., Huang, P., Li, J., Lin, B.-C., Lu, D., Lu, Y., *et al.*: Noisy intermediate-scale quantum computers. Frontiers of Physics **18**(2), 21308 (2023)

[33] Ali, S., Yue, T.: Modeling quantum programs: Challenges, initial results, and

research directions. In: Proceedings of the 1st ACM SIGSOFT International Workshop on Architectures and Paradigms for Engineering Quantum Software, pp. 14–21 (2020)

[34] Glos, A., Krawiec, A., Zimborás, Z.: Space-efficient binary optimization for variational quantum computing. npj Quantum Information **8**(1), 39 (2022)

[35] Haddaway, N.R., Page, M.J., Pritchard, C.C., McGuinness, L.A.: Prisma2020: An r package and shiny app for producing prisma 2020-compliant flow diagrams, with interactivity for optimised digital transparency and open synthesis. Campbell systematic reviews **18**(2), 1230 (2022)

[36] Page, M.J., McKenzie, J.E., Bossuyt, P.M., Boutron, I., Hoffmann, T.C., Mulrow, C.D., Shamseer, L., Tetzlaff, J.M., Akl, E.A., Brennan, S.E., et al.: The prisma 2020 statement: an updated guideline for reporting systematic reviews. bmj **372** (2021)

[37] Khanal, B., Rivas, P.: A modified depolarization approach for efficient quantum machine learning. Mathematics **12**(9), 1385 (2024)

[38] Martinez, J.E.: Decoherence and quantum error correction for quantum computing and communications. arXiv preprint arXiv:2202.08600 (2022)

[39] Chai, X., Gu, Y., Zhuang, W., Qian, P., Xiao, X., Liu, D.E.: Comparisons among the performances of randomized-framed benchmarking protocols under t1, t2 and coherent error models. arXiv preprint arXiv:2309.15482 (2023)

[40] Zheng, H., Li, H., Zhang, J., Wang, G., Guo, J., Wang, J.: Quantum neural network-based approach for optimizing road network selection. Geocarto International **40**(1), 2471108 (2025)

[41] Yu, H., Wei, T.-C.: Efficient separate quantification of state preparation errors and measurement errors on quantum computers and their mitigation. Quantum **9**, 1724 (2025)

[42] Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S.C., Endo, S., Fujii, K., McClean, J.R., Mitarai, K., Yuan, X., Cincio, L., *et al.*: Variational quantum algorithms. Nature Reviews Physics **3**(9), 625–644 (2021)

[43] Melnikov, A., Kordzanganeh, M., Alodjants, A., Lee, R.-K.: Quantum machine learning: from physics to software engineering. Advances in Physics: X **8**(1), 2165452 (2023)

[44] Rehman, J.U., Ulum, M.S., Shaffar, A.W., Hakim, A.A., Abdullah, Z., Al-Hraishawi, H., Chatzinotas, S., Shin, H., *et al.*: Evolutionary algorithms and quantum computing: recent advances, opportunities, and challenges. IEEE Access **13**, 16649–16670 (2025)

[45] Xu, Y., Wang, X., Wang, Z.: Gate-based variational quantum algorithm for truss structure size optimization problem. AIAA Journal **62**(12), 4824–4833 (2024)

[46] Wang, Y., Ding, Y., Wang, J., Chen, X.: Digital quantum simulation of nonadiabatic geometric gates via shortcuts to adiabaticity. Entropy **22**(10), 1175 (2020)

[47] Gelman, M.: A survey of methods for mitigating barren plateaus for parameterized quantum circuits. arXiv preprint arXiv:2406.14285 (2024)

[48] Qiu, C., Zhu, P., Wei, L.: A beam search framework for quantum circuit mapping. Entropy **27**(3), 232 (2025)

[49] Niu, S., Suau, A., Staffelbach, G., Todri-Sanial, A.: A hardware-aware heuristic for the qubit mapping problem in the nisq era. IEEE Transactions on Quantum Engineering **1**, 1–14 (2020)

[50] Riaz, F., Abdulla, S., Suzuki, H., Ganguly, S., Deo, R.C., Hopkins, S.: Accurate image multi-class classification neural network model with quantum entanglement approach. Sensors **23**(5), 2753 (2023)

[51] Jun, Y.-M., Choi, I.-C.: Optimal multi-bit toffoli gate synthesis. IEEE Access **11**, 27342–27351 (2023)

[52] Volpe, D., Orlandi, G., Turvani, G.: Improving the solving of optimization problems: A comprehensive review of quantum approaches. Quantum Reports **7**(1), 3 (2025)

[53] Heng, S., Kim, D., Kim, T., Han, Y.: How to solve combinatorial optimization problems using real quantum machines: A recent survey. IEEE Access **10**, 120106–120121 (2022)

[54] Boyle, A.O., Nikandish, R.: A hybrid quantum-classical generative adversarial network for near-term quantum processors. IEEE Access (2024)

[55] Munawar, G., Maulidevi, N.U., Surendro, K.: Proposed feature dimensionality reduction method in a predictive model for carbon neutrality: A hybrid quantum-classical approach. In: Proceedings of the 2025 14th International Conference on Software and Computer Applications, pp. 154–160 (2025)

[56] AlSagri, H.S., Kumar, A., Jilani Saudagar, A.K., Kumar, A., Raja, L.: Optimizing resource allocation in precision farming using quantum enhanced algorithms and quantum sensor networks. Journal of Cloud Computing **14**(1), 47 (2025)

[57] Niu, S., Todri-Sanial, A.: Effects of dynamical decoupling and pulse-level optimizations on ibm quantum computers. IEEE Transactions on Quantum Engineering **3**, 1–10 (2022)

[58] Tibaldi, S., Vodola, D., Tignone, E., Ercolessi, E.: Bayesian optimization for qaoa. IEEE Transactions on Quantum Engineering **4**, 1–11 (2023)

[59] Kim, J.E., Wang, Y.: Quantum approximate bayesian optimization algorithms with two mixers and uncertainty quantification. IEEE Transactions on Quantum Engineering **4**, 1–17 (2023)

[60] CHIV, K., HOUR, L., LEE, S.: Strategies for noise-resilient quantum approximate optimization algorithms: A review and classification of error mitigation (2025)

[61] Baiardi, A., Christandl, M., Reiher, M.: Quantum computing for molecular biology. ChemBioChem **24**(13), 202300120 (2023)

[62] Zhu, Y., Zhou, Y., Cheng, J., Jin, Y., Li, B., Niu, S., Liang, Z.: Compiler optimizations for qaoa. In: Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design, pp. 1–7 (2024)

[63] Nakada, H., Tanahashi, K., Tanaka, S.: Inductive construction of variational quantum circuit for constrained combinatorial optimization. IEEE Access (2025)

[64] Gargiulo, R., Rizzi, M., Zeier, R.: Computing classical partition functions: from onsager and kaufman to quantum algorithms. In: Proceedings of Recent Advances in Quantum Computing and Technology, pp. 20–50 (2024)

[65] Stroev, N., Berloff, N.G.: Analog photonics computing for information processing, inference, and optimization. Advanced Quantum Technologies **6**(9), 2300055 (2023)

[66] Barral, D., Cardama, F.J., Diaz-Camacho, G., Faílde, D., Llovo, I.F., Mussa-Juane, M., Vázquez-Pérez, J., Villasuso, J., Piñeiro, C., Costas, N., *et al.*: Review of distributed quantum computing: from single qpu to high performance quantum computing. Computer Science Review **57**, 100747 (2025)

[67] Salman, U.T., Wang, Z., Hansen, T.M.: Quircy: Quantum integrated resiliency for power systems. IEEE Access (2025)

[68] Caleffi, M., Amoretti, M., Ferrari, D., Illiano, J., Manzalini, A., Cacciapuoti, A.S.: Distributed quantum computing: a survey. Computer Networks **254**, 110672 (2024)

[69] Saini, R., Mani, A., Prasad, M., Bhattacharyya, S., Platos, J.: Evolutionary algorithm for the traveling salesman problem with innovative encoding on hybrid quantum-classical machines. IEEE Access (2025)

[70] Morais, A., Osaba, E., Pastor, I., Oregui, I.: Comparative analysis of classical

and quantum-inspired solvers: A preliminary study on the weighted max-cut problem. arXiv preprint arXiv:2504.05989 (2025)

[71] Chen, T.-H., Hung, W.-H.: Feasibility and limitations of generalized grover search algorithm-based quantum asymmetric cryptography: An implementation study on quantum hardware. Electronics **14**(19), 3821 (2025)

[72] Metwalli, S.A., Le Gall, F., Van Meter, R.: Finding small and large k-clique instances on a quantum computer. IEEE Transactions on Quantum Engineering **1**, 1–11 (2020)

[73] Inglesant, P., Ten Holter, C., Jirotka, M., Williams, R.: Asleep at the wheel? responsible innovation in quantum computing. Technology Analysis & Strategic Management **33**(11), 1364–1376 (2021)

[74] Miranda, E.R., Siegelwax, B.N.: Advancing generative ai for music with photonics. International Journal of Parallel, Emergent and Distributed Systems, 1–20 (2025)

[75] Watabe, M., Shiba, K., Chen, C.-C., Sogabe, M., Sakamoto, K., Sogabe, T.: Quantum circuit learning with error backpropagation algorithm and experimental implementation. Quantum Reports **3**(2), 333–349 (2021)

[76] Sajadimanesh, S., Rad, H.A., Faye, J.P.L., Atoofian, E.: Nr-qnn: Noise-resilient quantum neural network. IEEE Access **13**, 40185–40197 (2025)

[77] Kashif, M., Al-Kuwari, S.: The unified effect of data encoding, ansatz expressibility and entanglement on the trainability of hqnns. International Journal of Parallel, Emergent and Distributed Systems **38**(5), 362–400 (2023)

[78] Oonishi, K., Tanaka, T., Uno, S., Satoh, T., Van Meter, R., Kunihiro, N.: Efficient construction of a control modular adder on a carry-lookahead adder using relative-phase toffoli gates. IEEE Transactions on Quantum Engineering **3**, 1–18 (2021)

[79] Hua, F., Chen, Y., Jin, Y., Zhang, C., Hayes, A., Zhang, Y., Zhang, E.Z.: Autobraid: A framework for enabling efficient surface code communication in quantum computing. In: MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 925–936 (2021)

[80] Bermejo, P., Orús, R.: Variational quantum non-orthogonal optimization. Scientific Reports **13**(1), 9840 (2023)

[81] Park, S., Baek, H., Kim, J.: Quantum split learning for privacy-preserving information management. In: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, pp. 4239–4243 (2023)

[82] Duque, A., Freire, P., Manuylovich, E., Stoliarov, D., Prilepsky, J., Turitsyn, S.: Improving analog neural network robustness: A noise-agnostic approach with explainable regularizations. arXiv preprint arXiv:2409.08633 (2024)

[83] Mourgias-Alexandris, G., Moralis-Pegios, M., Tsakyridis, A., Simos, S., Dabos, G., Totovic, A., Passalis, N., Kirtas, M., Rutirawut, T., Gardes, F., *et al.*: Noise-resilient and high-speed deep learning with coherent silicon photonics. Nature communications **13**(1), 5572 (2022)

[84] Khan, M.U., Kamran, M.A., Khan, W.R., Ibrahim, M.M., Ali, M.U., Lee, S.W.: Error mitigation in the nisq era: Applying measurement error mitigation techniques to enhance quantum circuit performance. Mathematics **12**(14), 2235 (2024)

[85] Attisara, L.A., Kumar, S.A.: Reinforcement learning for optimizing large qubit array based quantum sensor circuits. Discover Applied Sciences (2026)

[86] Chen, L., Li, T., Chen, Y., Chen, X., Wozniak, M., Xiong, N., Liang, W.: Design and analysis of quantum machine learning: a survey. Connection Science **36**(1), 2312121 (2024)

[87] Li, A., Stein, S., Krishnamoorthy, S., Ang, J.: Qasmbench: A low-level quantum benchmark suite for nisq evaluation and simulation. ACM Transactions on Quantum Computing **4**(2), 1–26 (2023)

[88] Liu, H., Spedalieri, F., Sabek, I.: A demonstration of q2o: Quantum-augmented query optimizer. Proceedings of the VLDB Endowment **18**(12), 5439–5443 (2025)

[89] Stein, O., Oldenburg, J., Marquardt, W.: Continuous reformulations of discrete–continuous optimization problems. Computers & chemical engineering **28**(10), 1951–1966 (2004)

[90] Mandal, A.K., Chakraborty, B.: Quantum computing and quantum-inspired techniques for feature subset selection: A review. Knowledge and Information Systems **67**(3), 2019–2061 (2025)

[91] Liu, J., Moraglio, A.: A framework for automatically setting multiple penalty weights in quadratic unconstrained binary optimization. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 575–578 (2025)

[92] Novák, V., Zelinka, I., Snášel, V.: Optimization strategies for variational quantum algorithms in noisy landscapes. arXiv preprint arXiv:2506.01715 (2025)

[93] Alvarez, G., Bennink, R., Irle, S., Jakowski, J.: Gene expression programming for quantum computing. ACM Transactions on Quantum Computing **4**(4), 1–14

(2023)

[94] Mohammadisiahroudi, M., Wu, Z., Augustino, B., Carr, A., Terlaky, T.: Improvements to quantum interior point method for linear optimization. ACM Transactions on Quantum Computing **6**(1), 1–24 (2025)

[95] Doolittle, B., Bromley, R.T., Killoran, N., Chitambar, E.: Variational quantum optimization of nonlocality in noisy quantum networks. IEEE Transactions on Quantum Engineering **4**, 1–27 (2023)

[96] Liang, Z., Liu, G., Liu, Z., Cheng, J., Hao, T., Liu, K., Ren, H., Song, Z., Liu, J., Ye, F., *et al.*: Graph learning for parameter prediction of quantum approximate optimization algorithm. In: Proceedings of the 61st ACM/IEEE Design Automation Conference, pp. 1–4 (2024)

[97] Lin, Y.-C., Wang, C.-C., Tu, C.-H., Hung, S.-H.: Towards optimizations of quantum circuit simulation for solving max-cut problems with qaoa. In: Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing, pp. 1487–1494 (2024)

[98] Amico, M., Zhang, H., Jurcevic, P., Bishop, L.S., Nation, P., Wack, A., McKay, D.C.: Defining best practices for quantum benchmarks. In: 2023 IEEE International Conference on Quantum Computing and Engineering (QCE), vol. 1, pp. 692–702 (2023). IEEE

[99] Calderer, S.R.: Quantum algorithms: Bridging the gap between theory and experiment. PhD thesis, Universitat de Barcelona (Spain) (2024)

[100] Pelofske, E., Bärtschi, A., Eidenbenz, S.: Quantum volume in practice: What users can expect from nisq devices. IEEE Transactions on Quantum Engineering **3**, 1–19 (2022)

[101] Bengtsson, A., Vikstål, P., Warren, C., Svensson, M., Gu, X., Kockum, A.F., Krantz, P., Križan, C., Shiri, D., Svensson, I.-M., *et al.*: Improved success probability with greater circuit depth for the quantum approximate optimization algorithm. Physical Review Applied **14**(3), 034010 (2020)

[102] Shafique, M.A., Munir, A., Latif, I.: Quantum computing: Circuits, algorithms, and applications. IEEE Access **12**, 22296–22314 (2024)