Stabilizing Grover Search via Density-Aware Phase Control

Simone Faro
Department of Mathematics and Computer Science
University of Catania
Catania, Italy
simone.faro@unict.it

Gianmarco La Rosa
Department of Mathematics and Computer Science
Università di Palermo
Palermo, Italy
gianmarco.larosa@unipa.it

Arianna Pavone
Department of Mathematics and Computer Science
Università di Palermo
Palermo, Italy
ariannamaria.pavone@unipa.it

Abstract

Grover's algorithm provides a quadratic speedup for unstructured search and represents one of the most fundamental primitives of quantum computation. However, its practical performance can be sensitive to the precise choice of the iteration count, which depends on the number of marked elements. This sensitivity becomes more pronounced when the density of solutions increases, where even small inaccuracies in estimating the number of marked states may lead to overshooting of the target subspace and a corresponding drop in success probability. In this work we introduce a density-aware parametrization of Grover search in which the phase rotations of the Grover operator are computed explicitly as a function of the solution density. The resulting construction can be interpreted as a density-dependent instantiation of Høyer's arbitrary-phase amplitude amplification framework. The resulting operator performs a controlled rotation in the two–dimensional search subspace and naturally interpolates between sparse and dense regimes of the search space. We provide a theoretical analysis of the resulting dynamics and show that the proposed parametrization stabilizes the amplification process across a wide range of solution densities. Numerical simulations demonstrate that the method maintains near–optimal success probability for densities where the standard Grover algorithm becomes unstable, while also exhibiting improved robustness with respect to errors in estimating the number of marked states. These results suggest that density-aware phase parametrizations offer a simple yet effective strategy for extending the practical applicability of Grover-style quantum search procedures.

CCS Concepts

• **Theory of computation → Quantum information theory**.

Keywords

Grover's algorithm, quantum search, quantum computing

ACM Reference Format:
Simone Faro, Gianmarco La Rosa, and Arianna Pavone. 2026. Stabilizing Grover Search via Density-Aware Phase Control. In *The 35th International Symposium on High-Performance Parallel and Distributed Computing (HPDC '26), July 13–16, 2026, Cleveland, OH, USA*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3806645.3816163

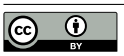

HPDC '26, Cleveland, OH, USA

ACM ISBN 979-8-4007-2640-8/26/07
https://doi.org/10.1145/3806645.3816163

1 Introduction

Grover's algorithm is one of the most celebrated results in quantum computing. Introduced in the mid-1990s, it provides a quadratic speedup for the problem of unstructured search, allowing a marked element to be found in a search space of size N using $O(\sqrt{N})$ oracle queries [2, 6, 7]. Beyond its original formulation, Grover's algorithm has become a fundamental building block in quantum algorithm design and underlies the more general framework of amplitude amplification [3]. In addition, it has also inspired a variety of quantum search techniques based on different computational paradigms, including quantum walk based search algorithms [1, 5].

At the heart of the algorithm lies a simple geometric principle. The Grover operator performs a sequence of rotations in a two–dimensional subspace spanned by the uniform superposition of marked and unmarked states. Each iteration increases the amplitude of the solution subspace through constructive interference, and after a suitable number of iterations the probability of measuring a marked element becomes close to one.

Despite its conceptual elegance and theoretical optimality [12], the practical behavior of Grover's algorithm is sensitive to the density of solutions in the search space. In particular, the algorithm assumes that the number of marked elements M is small compared with N. When the density of solutions increases, the amplification process becomes progressively less stable. The rotation angle grows with $\sqrt{M/N}$, and the success probability becomes highly sensitive to the precise number of iterations. As a consequence, small inaccuracies in the estimation of M may lead to significant performance degradation due to overshooting of the target subspace.

These limitations have motivated several extensions of Grover's algorithm based on generalized phase rotations and amplitude amplification operators. In particular, Høyer introduced a formulation of amplitude amplification with arbitrary phase shifts [9], while Long proposed phase-matching techniques that allow exact success probabilities under suitable parameter choices [10, 11]. These results highlight the crucial role played by phase parameters in controlling the dynamics of the search process.

In this work we revisit Grover search from a density-aware perspective. Our approach derives an explicit rule for selecting the phase parameters of the Grover operator as a function of the solution density. While previous frameworks establish general conditions under which arbitrary phase rotations implement amplitude amplification, the parametrization proposed here provides a simple closed-form prescription for choosing the phases directly from the problem parameters.

The resulting operator performs a controlled rotation in the two–dimensional search subspace and adapts smoothly to different

regimes of the search space, ranging from sparse to dense distributions of solutions. Our analysis suggests that this density-dependent phase selection mitigates the sensitivity of the amplification process to the choice of the iteration count, particularly when the density of marked states is not negligible.

In these regimes the standard Grover iteration becomes increasingly sensitive to the exact value of the optimal iteration number, and small inaccuracies in estimating the number of marked states may lead to overshooting of the target subspace and a corresponding drop in success probability. By explicitly adjusting the phase parameters as a function of the solution density, the proposed approach maintains near-optimal success probability across a wider range of densities and exhibits improved robustness with respect to errors in estimating the number of marked states.

We support these findings with an extensive experimental evaluation based on quantum circuit simulations. The experiments analyze the behavior of the algorithm across different solution densities, compare the number of required iterations with the standard Grover search, and evaluate the robustness of the approach under imperfect knowledge of the number of solutions.

Taken together, these results suggest that density-aware phase parametrizations provide a simple and practical mechanism for selecting Grover phase parameters and for improving the robustness of Grover-style search procedures across different density regimes.

2 Background

This section reviews the main concepts underlying the analysis presented in this work. We first recall the standard formulation of Grover's search and discuss its behavior as a function of the solution density. We then briefly review the general framework of amplitude amplification and the extension to arbitrary phase rotations introduced by Høyer, which provides the theoretical foundation for the density-aware parametrization proposed in this paper.

2.1 Grover Search

Consider an unstructured search problem over a domain of size $N = 2^n$. Let $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function identifying the marked elements, and let M denote the number of inputs x such that $f(x) = 1$. The goal is to find one marked element using access to a quantum oracle for f.

Grover's algorithm starts from the uniform superposition of all basis states

$$|\psi_0\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle.$$

The oracle is implemented as a phase operator

$$O_f|x\rangle = (-1)^{f(x)}|x\rangle,$$

which flips the phase of the marked states while leaving the others unchanged.

The second ingredient of the algorithm is the diffusion operator

$$D = 2|\psi_0\rangle\langle\psi_0| - I,$$

which performs an inversion about the mean amplitude of the state vector. The Grover iteration is defined as the composition

$$G = DO_f.$$

A convenient way to analyze the algorithm is to restrict attention to the two-dimensional subspace spanned by the states

$$|g\rangle = \frac{1}{\sqrt{M}} \sum_{f(x)=1} |x\rangle, \qquad |b\rangle = \frac{1}{\sqrt{N-M}} \sum_{f(x)=0} |x\rangle,$$

representing the uniform superpositions of marked and unmarked states respectively. The initial state can be written in this basis as

$$|\psi_0\rangle = \sqrt{\frac{M}{N}}\,|g\rangle + \sqrt{\frac{N-M}{N}}\,|b\rangle.$$

Let

$$\sin^2\theta = \frac{M}{N}.$$

Under this representation the Grover operator acts as a rotation by angle 2θ in the plane spanned by $|g\rangle$ and $|b\rangle$. After k iterations the state becomes

$$|\psi_k\rangle = \sin((2k+1)\theta)\,|g\rangle + \cos((2k+1)\theta)\,|b\rangle.$$

Consequently, the probability of measuring a marked element after k iterations is

$$P_k = \sin^2((2k+1)\theta).$$

The success probability is maximized when the rotation approaches the marked subspace, which occurs for approximately

$$k \approx \frac{\pi}{4\theta} - \frac{1}{2} \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}.$$

Thus Grover's algorithm finds a marked element using $O(\sqrt{N/M})$ oracle queries, achieving a quadratic improvement over classical randomized search.

2.2 Practical Limitations of Grover Search

Although Grover's algorithm provides an optimal quadratic speedup for unstructured search problems, its standard formulation exhibits practical limitations when the number of marked elements becomes large relative to the size of the search space. In particular, the effectiveness of the amplitude amplification mechanism deteriorates as the ratio M/N approaches $1/2$, where $N = 2^n$ denotes the size of the search space and M the number of valid solutions.

The behavior of the algorithm can be understood through its geometric interpretation. Grover's iteration acts as a rotation in the two–dimensional subspace spanned by the uniform superpositions of solution and non–solution states. The rotation angle per iteration is given by 2θ, where $\sin^2\theta = M/N$.

As the ratio M/N increases, the angle θ becomes larger and the number of iterations required to reach the solution subspace decreases. When M/N approaches $1/2$, the distinction between marked and unmarked states becomes less pronounced, and the amplification process becomes increasingly sensitive to the precise number of iterations.

An even more problematic situation arises when the number of solutions exceeds half of the search space. In this regime the majority of basis states correspond to valid solutions, and the reflection about the mean amplitude performed by the diffusion operator no

longer amplifies the desired states in a meaningful way. Instead, the interference mechanism may lead to oscillatory or nearly uniform output distributions, thereby reducing the advantage provided by amplitude amplification.

For small values of M/N, the algorithm achieves success probabilities close to one when the number of iterations is appropriately chosen. However, as the ratio approaches $1/2$, the amplification effect becomes weaker and the success probability decreases.

For values of M/N beyond this threshold, the algorithm provides little advantage over classical random sampling. In fact, when a large fraction of the search space consists of valid solutions, a simple measurement of the uniform superposition already yields a high probability of success. This observation highlights that the standard Grover iteration is primarily designed for sparse search problems and may behave suboptimally in high-density regimes.

Understanding and addressing this limitation has motivated several extensions of Grover's algorithm based on modified phase rotations and generalized amplitude amplification operators. Examples include phase-matching techniques, exact search algorithms, and fixed-point quantum search procedures designed to avoid overshooting phenomena [8, 10, 11]. These approaches aim at controlling the rotation dynamics more precisely, allowing the search procedure to remain effective across a wider range of solution densities.

2.3 Amplitude Amplification

Grover's search algorithm can be understood as a special case of the more general framework of *amplitude amplification*, introduced by Brassard, Høyer, Mosca, and Tapp [3]. This framework extends Grover's method by allowing a broader class of quantum procedures in which the probability amplitude of desirable states can be amplified through repeated unitary transformations.

Let A be a unitary operator preparing an initial state $|\psi\rangle = A|0\rangle$.

Assume that the Hilbert space can be partitioned into two orthogonal subspaces corresponding to *good* and *bad* states according to a Boolean predicate $f(x)$. The initial state can therefore be decomposed as $|\psi\rangle = \sqrt{a}\,|g\rangle + \sqrt{1-a}\,|b\rangle$, where $|g\rangle$ and $|b\rangle$ denote the normalized superpositions of good and bad states respectively, and a represents the initial success probability.

Amplitude amplification proceeds by repeatedly applying the operator $Q = -AS_0A^{-1}S_f$, where S_f flips the phase of the good states and S_0 flips the phase of the all-zero state. As in Grover's algorithm, the operator Q acts as a rotation in the two–dimensional subspace spanned by $|g\rangle$ and $|b\rangle$. Each iteration increases the amplitude associated with the good states, thereby increasing the probability of observing a valid solution upon measurement.

Grover's search algorithm corresponds to the special case in which A prepares the uniform superposition over all basis states. More generally, amplitude amplification provides a quadratic speedup for any procedure that produces a valid solution with non-zero probability.

This general framework reveals that Grover's algorithm is not an isolated construction but rather a specific instance of a broader class of amplitude amplification processes. In the following subsection we discuss an important extension of this framework that allows arbitrary phase rotations, leading to a more flexible control of the amplification dynamics.

2.4 Amplification with Arbitrary Phases

A particularly important generalization of Grover's algorithm was introduced by Høyer [9], who showed that the amplitude amplification process can be implemented using arbitrary phase rotations.

In this formulation the oracle and diffusion operators are replaced by phase operators

$$O(\phi), \qquad D(\varphi),$$

where ϕ and φ denote rotation angles applied to the marked states and to the initial state respectively. The generalized Grover operator becomes

$$Q(\phi, \varphi) = D(\varphi)O(\phi).$$

Høyer showed that the operator $Q(\phi, \varphi)$ continues to act as a rotation in the two–dimensional search subspace provided that the phases satisfy the relation

$$\tan(\varphi/2) = \tan(\phi/2)(1 - 2a), \tag{1}$$

where $a = M/N$ denotes the fraction of marked states.

This result reveals that Grover's algorithm corresponds to a specific point in a larger family of amplitude amplification procedures parameterized by the phases ϕ and φ. By selecting appropriate phase values it is possible to control the rotation performed in the search subspace and therefore influence the convergence properties of the algorithm.

The framework of arbitrary-phase amplitude amplification provides the theoretical foundation for several subsequent extensions of Grover search, including phase-matching techniques and exact search algorithms [10, 11]. In the next section we exploit this flexibility to derive a density-aware parametrization of the Grover operator in which the rotation phases are computed explicitly as a function of the ratio between marked and unmarked states.

2.5 Domain Expansion

A related line of work addresses the degradation of Grover's algorithm in high-density solution spaces. In particular, Faro and Marino [4] proposed a technique based on a virtual expansion of the input domain, in which auxiliary qubits are introduced in order to reduce the effective density of marked states in an extended search space. By selectively filtering the solution set in the expanded domain, the method restores the geometric conditions required for constructive amplitude amplification even when the number of solutions exceeds half of the original search space.

The approach proposed in this work is complementary in nature. Rather than modifying the structure of the search space, we adjust the phase parameters of the Grover operator itself, deriving a density-dependent rule that stabilizes the rotation dynamics of amplitude amplification across different solution densities.

To better position the proposed approach within the existing literature, Table 1 summarizes several representative extensions of Grover search and highlights their main characteristics in terms of phase selection, control of overshooting phenomena, and query complexity.

Table 1: Comparison of approaches addressing stability issues in Grover-style search.

Method	Phase choice	Overshoot control	Query complexity	Density-aware
Grover (1996)	fixed (π)	no	$O(\sqrt{N/M})$	no
Høyer (2000)	arbitrary phases	framework only	$O(\sqrt{N/M})$	implicit
Long (1999–2001)	phase matching	exact search	$O(\sqrt{N/M})$	no
Grover (2005)	phase sequence	fixed-point	$> O(\sqrt{N/M})$	no
Faro–Marino (2025)	fixed (π)	yes	$O(\sqrt{N/M})$	indirect
Faro-LaRosa-Pavone (this work)	closed-form $\phi(p)$	yes	$O(\sqrt{N/M})$	yes

3 Density-Robust Grover Search

In this section we introduce a density-aware parametrization of the Grover operator in which the phase rotations are computed explicitly as a function of the fraction of marked states in the search space. The goal is to obtain a search procedure whose behavior remains stable across a wide range of solution densities.

Let $N = 2^n$ denote the size of the search space and let M be the number of marked elements. We define the solution density as

$$p = \frac{M}{N}.$$

As discussed in the previous sections, the dynamics of Grover's algorithm can be analyzed within the two–dimensional subspace spanned by the normalized superpositions of marked and unmarked states. In this representation the algorithm performs a sequence of rotations whose angle depends on the parameter

$$\theta = \arcsin\left(\sqrt{p}\right).$$

In the standard Grover algorithm the phase inversions correspond to the choice $\phi = \varphi = \pi$. Within the more general amplitude amplification framework, however, the phases can be selected in order to control the rotation performed in the search subspace.

In our approach we explicitly compute the phase parameter as a function of the density p. First we select the number of Grover iterations as

$$k = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor. \tag{2}$$

This value corresponds to the largest integer number of rotations that does not overshoot the target solution subspace.

To determine the appropriate phase parameter, we recall that in the two–dimensional representation the state evolution can be interpreted geometrically as a sequence of rotations starting from the initial angle θ with respect to the marked subspace. After k Grover iterations the standard algorithm reaches the angle

$$(2k+1)\theta.$$

The optimal situation corresponds to reaching exactly the marked subspace, which occurs when the total rotation equals $\pi/2$. However, because k must be an integer, the standard Grover iteration generally cannot achieve this value exactly, leading to the well-known overshooting phenomenon.

To compensate for this discretization effect, we modify the phase of the Grover operator so that each iteration produces a slightly different effective rotation angle. Let α denote the effective rotation angle induced by a single iteration of the generalized Grover operator. After k iterations the total rotation becomes

$$\theta + 2k\alpha.$$

We choose α so that the final state lies exactly on the marked subspace, i.e.,

$$\theta + 2k\alpha = \frac{\pi}{2}.$$

Solving for α yields

$$\alpha = \frac{\pi/2 - \theta}{2k}.$$

Within the amplitude amplification framework with arbitrary phases, the effective rotation angle depends on the phase parameters applied by the oracle and diffusion operators. Following the standard analysis of generalized Grover operators, the relation between the phase ϕ and the induced rotation can be expressed in terms of the initial angle θ. Imposing that the rotation produced by a single iteration equals α leads to the phase choice

$$\phi = 2\arcsin\left(\frac{\sin(\pi/(4k+2))}{\sin\theta}\right), \tag{3}$$

In our construction we use the symmetric choice $\phi = \varphi$, resulting in the generalized Grover operator

$$Q(\phi,\phi) = D(\phi)O(\phi). \tag{4}$$

With this parametrization the operator performs a controlled rotation in the two–dimensional search subspace that approaches the marked component without overshooting the optimal point.

Intuitively, the phase parameter ϕ compensates for the variation of the rotation angle induced by the density of marked states. When the density is small, the value of ϕ approaches π and the algorithm reduces to the standard Grover search. As the density increases, the phase rotation adjusts accordingly, maintaining a stable amplification behavior even in regimes where the standard Grover operator becomes inefficient. This explicit density-aware parametrization provides a simple way to adapt the Grover iteration to the structure of the search space while preserving the geometric interpretation of amplitude amplification.

Finally, observe that, as in most amplitude amplification schemes, this parametrization assumes that an estimate of the number of marked states is available. The robustness experiments in Section 5 analyze the impact of estimation errors on the performance of the algorithm.

3.1 Relation with Høyer's Amplification

The density-aware parametrization introduced in the previous section can be naturally interpreted within the Høyer's general framework of amplitude amplification with arbitrary phases [9].

In Høyer's formulation the Grover operator is generalized by replacing the phase inversions of the oracle and the diffusion operator with arbitrary phase rotations. The resulting amplification operator can be written as in eq. (4), where ϕ and φ denote the phases applied to the marked states and to the initial state respectively.

Høyer showed that the operator $Q(\phi, \varphi)$ acts as a rotation in the two–dimensional search subspace spanned by the superpositions of marked and unmarked states provided that the two phases satisfy the relation given in eq. (1).

This result reveals that Grover's algorithm corresponds to a particular point in a larger family of amplitude amplification procedures parameterized by the phases ϕ and φ. In particular, the standard Grover operator is obtained for the choice $\phi = \varphi = \pi$.

The parametrization proposed in this work can be viewed as a constructive instantiation of this framework. While Høyer's result provides the general condition under which arbitrary phase rotations implement a valid amplitude amplification operator, it does not prescribe a specific rule for selecting the phases in terms of the problem parameters.

In contrast, our approach derives an explicit density-dependent rule for choosing the phase parameters. Given the solution density $p = M/N$ and the corresponding angle $\theta = \arcsin(\sqrt{p})$, we first determine the number of Grover iterations as in eq. (2), which corresponds to the largest integer number of rotations that does not overshoot the marked subspace. We then select the phase parameter given in eq. (3), and use the symmetric choice $\varphi = \phi$.

This parametrization directly determines the phase values as a function of the solution density and the desired number of Grover iterations, yielding an operator that performs the appropriate rotation toward the marked subspace without overshooting.

From this perspective, the proposed algorithm can be interpreted as a density-aware realization of Høyer's arbitrary-phase amplitude amplification in which the phase parameters are obtained through a simple closed-form rule derived from the geometry of the process.

4 Theoretical Analysis

Let $N = 2^n$ denote the size of the search space and let M be the number of marked elements. Define the solution density $p = M/N$.

As discussed in the previous sections, Grover's algorithm operates within the two–dimensional subspace spanned by the normalized superpositions of marked and unmarked states. In this representation the state evolution can be interpreted as a sequence of rotations characterized by the angle $\sin^2\theta = M/N$.

After k Grover iterations the state can be written as

$$|\psi_k\rangle = \sin((2k+1)\theta)\,|g\rangle + \cos((2k+1)\theta)\,|b\rangle.$$

The standard Grover algorithm achieves optimal amplification when the number of iterations is chosen so that the rotation approaches the marked subspace without overshooting. Our parametrization explicitly adjusts the phase of the Grover operator to preserve this property across different solution densities.

We analyze the resulting behavior in three regimes.

4.1 Sparse Regime

Consider first the case in which the number of marked elements is small compared to the size of the search space ($M \ll N$). In this regime we have

$$\theta \approx \sqrt{\frac{M}{N}}.$$

Consequently, the optimal number of Grover iterations is

$$k \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}.$$

For small values of θ the phase parameter computed by our method satisfies $\phi \approx \pi$, and therefore the generalized Grover operator $Q(\phi, \phi)$ reduces to the standard Grover operator. This shows that in the sparse regime the proposed algorithm behaves identically to the original Grover search and preserves its optimal quadratic speedup.

4.2 Intermediate Regime

We now consider the regime in which the number of marked elements becomes comparable to the number of unmarked states ($M \approx N/2$). In this case the rotation angle approaches

$$\theta \approx \frac{\pi}{4}.$$

The number of required Grover iterations decreases rapidly, and the standard Grover algorithm becomes highly sensitive to the exact choice of the iteration count. A small deviation from the optimal value may cause the state to overshoot the solution subspace, leading to a significant decrease in success probability.

The phase parametrization introduced in this work mitigates this effect by modifying the rotation induced by the amplification operator. The resulting operator performs a controlled rotation that approaches the marked subspace without overshooting, thereby maintaining stable amplification behavior even when the density of solutions is relatively large.

4.3 Dense Regime

Finally, consider the case in which the majority of the search space consists of marked elements ($M \to N$). In this regime we have

$$\theta \to \frac{\pi}{2}.$$

The optimal number of iterations therefore approaches zero. Intuitively, this reflects the fact that the initial uniform superposition already contains a high probability of measuring a valid solution.

In this regime amplitude amplification becomes unnecessary. Our parametrization naturally captures this behavior: as the density p increases, the computed value of k decreases and the algorithm gradually reduces the number of amplification steps.

Thus the proposed parametrization smoothly interpolates between the three regimes of operation. It preserves the optimal behavior of Grover search when solutions are sparse, stabilizes the amplification process when the density of solutions increases, and naturally degenerates to direct sampling when the majority of states are marked.

As shown by the orange curve, the resulting algorithm exhibits a more regular behavior across different density regimes. While the algorithm naturally degenerates to direct sampling when the number of marked states becomes large, it avoids the abrupt oscillations that characterize the standard Grover operator.

4.4 Query Complexity

Before turning to the experimental evaluation, it is useful to clarify that the proposed density-aware parametrization does not alter the asymptotic query complexity of Grover search.

Theorem 4.1. *Let N be the size of the search space and M the number of marked elements. The density-aware Grover parametrization described above finds a marked element using $O(\sqrt{N/M})$ oracle queries.*

Proof. Let $p = M/N$ and $\theta = \arcsin(\sqrt{p})$. The algorithm selects

$$k = \left\lfloor \frac{\pi}{4\theta} - \frac{1}{2} \right\rfloor.$$

Since $\theta = \arcsin(\sqrt{p}) \approx \sqrt{p}$ for small p, we obtain

$$k \approx \frac{\pi}{4}\sqrt{\frac{N}{M}}.$$

Each iteration applies one oracle call and one diffusion operator. Therefore the total number of oracle queries is

$$O\left(\sqrt{\frac{N}{M}}\right).$$

The additional phase parameters introduced by the density-aware parametrization only modify the rotation angle of the Grover operator and do not change the asymptotic query complexity. □

Thus, the proposed parametrization does not alter the optimal query complexity of Grover search, but provides a simple mechanism for selecting phase parameters that improves robustness with respect to density variations and estimation errors.

5 Experimental Evaluation

In this section we evaluate the practical behavior of the proposed density-aware Grover parametrization through numerical simulations. The experiments aim at analyzing how the algorithm behaves across different regimes of solution density and how it compares with the standard Grover search.

All experiments were conducted using quantum circuit simulations based on the Qiskit framework. To obtain exact probabilities and avoid sampling noise, we used statevector simulations of the quantum circuits. Given a search space of size $N = 2^n$ and a set of M marked elements, the oracle was implemented as a phase-flip operator acting on the marked states. For each experiment, both the standard Grover algorithm and the proposed density-aware variant were executed using the same oracle.

The evaluation focuses on three aspects of the search process.

First, we analyze the *success probability as a function of the solution density*. This experiment measures the probability of observing a marked element after running the algorithm for different values of M in a fixed search space. The results provide insight into how the amplification mechanism behaves as the ratio M/N varies from sparse to dense regimes.

Second, we examine the *number of iterations required by the algorithms*. Since the number of Grover iterations directly determines the circuit depth, comparing this quantity helps clarify how the proposed parametrization adapts the search procedure to different densities of solutions.

Finally, we investigate the *robustness of the algorithms with respect to errors in estimating the number of marked states*. In practical scenarios the exact value of M may not be known in advance. To evaluate this effect, we repeat the search procedure using an estimated value $\widehat{M}$ and measure the resulting success probability as a function of the estimation error. This experiment allows us to assess how sensitive the algorithms are to inaccuracies in the estimation of the solution density.

In all experiments we compare three baselines: the standard Grover algorithm using the optimal number of iterations, the proposed density-aware Grover parametrization, and the uniform random sampling, which provides a natural baseline with success probability M/N.

In the first experiment we analyze the success probability of the search algorithms as a function of the density of marked states. The search space size was fixed to $N = 2^8$, and for each value of $M \in \{1, \ldots, N-1\}$ an oracle marking M states was generated. Both the standard Grover algorithm and the proposed density-aware parametrization were executed using the same oracle.

The probability of measuring a marked state was computed using exact statevector simulation, which provides the theoretical success probability of the corresponding quantum circuit. For comparison, we also report the baseline success probability obtained by uniform random sampling, which corresponds to selecting an element uniformly at random from the search space and therefore equals M/N.

Figure 1 shows the resulting success probabilities as the ratio M/N varies from sparse to dense solution regimes.

As expected, the standard Grover algorithm performs extremely well when the number of marked states is small. However, as the density of solutions increases, the behavior of the algorithm becomes increasingly sensitive to the precise choice of the iteration count. In these regimes even small deviations from the optimal number of iterations may lead to overshooting of the target subspace, producing oscillations and occasional drops in success probability.

The density-aware Grover parametrization proposed in this work exhibits a significantly more stable behavior. For sparse regimes the algorithm achieves success probabilities close to one, similarly to the standard Grover search. However, the parametrized phase rotation maintains near-optimal amplification across a much wider range of solution densities.

When the number of marked states becomes very large, the advantage of amplitude amplification naturally disappears and the algorithm gradually approaches the behavior of uniform random sampling. This transition reflects the fact that when most states are solutions, performing additional amplitude amplification steps becomes unnecessary.

The second experiment analyzes how the number of Grover iterations varies as a function of the solution density. As in the

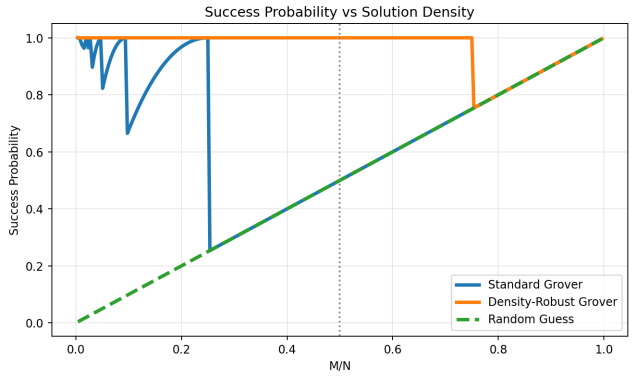

Figure 1: Success probability as a function of the solution density M/N for a search space of size $N = 2^8$. The blue curve corresponds to the standard Grover algorithm executed with the optimal number of iterations, while the orange curve represents the density-aware Grover parametrization proposed in this work. The dashed green line denotes the baseline success probability obtained by uniform random sampling. The vertical dashed line marks the threshold $M/N = 0.5$, which separates the sparse and dense solution regimes.

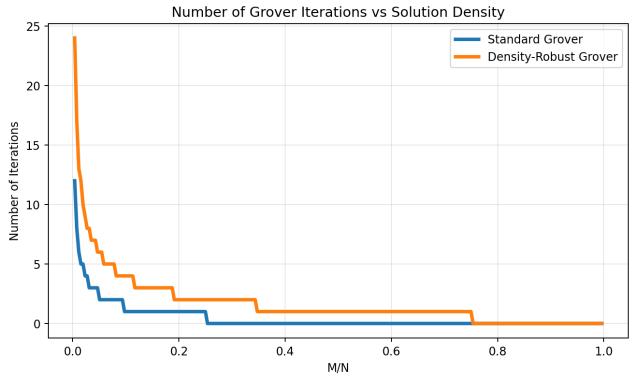

Figure 2: Number of Grover iterations as a function of the solution density M/N for a search space of size $N = 2^8$. The blue curve corresponds to the standard Grover algorithm, while the orange curve shows the number of iterations used by the proposed density-aware Grover parametrization.

previous experiment, the search space size was fixed to $N = 2^8$, and for each value of $M \in \{1, \ldots, N-1\}$ an oracle marking M states was generated.

For the standard Grover algorithm, the number of iterations was computed using the well-known optimal formula $k = \lfloor \frac{\pi}{4\theta} - \frac{1}{2} \rfloor$, where $\sin^2 \theta = M/N$. For the density-aware parametrization proposed in this work, the number of iterations was determined according to the phase-matching strategy described in Section 3.

Figure 2 illustrates how the number of iterations varies as the density of marked states increases.

For both algorithms, the number of required iterations decreases as the density of solutions increases. This behavior reflects the geometric nature of amplitude amplification: when the number of marked states grows, the initial amplitude of the solution subspace becomes larger, and fewer rotations are required to reach the target.

However, the density-aware parametrization generally uses a larger number of iterations in the sparse regime compared with the standard Grover algorithm. This additional flexibility allows the algorithm to adjust the phase of the Grover operator in order to maintain stable amplification across different density regimes.

As the density of solutions increases further, both algorithms gradually reduce the number of iterations, eventually reaching zero when the probability of sampling a solution from the initial superposition becomes sufficiently large. In this regime, amplitude amplification becomes unnecessary and the search procedure effectively reduces to direct sampling.

In many practical scenarios the exact number of marked states is not known in advance and must be estimated. Quantum amplitude estimation techniques provide a natural framework for obtaining such estimates within quantum algorithms [3]. Since both the number of Grover iterations and the phase parameters depend on the value of M, errors in this estimation may significantly affect the performance of the algorithm.

To evaluate this aspect, we performed a set of experiments in which the oracle marks a fixed number of states M, while the algorithms are executed using an estimated value $\widehat{M}$. The relative estimation error is defined as

$$\epsilon = \frac{\widehat{M} - M}{M}.$$

For each value of ϵ we computed the success probability of both the standard Grover algorithm and the proposed density-aware parametrization. The search space size was fixed to $N = 2^8$, and the experiments were repeated for several representative values of M covering sparse, intermediate, and dense regimes.

The results are shown in Figure 3. The behavior of the standard Grover algorithm is highly sensitive to the accuracy of the estimate $\widehat{M}$. When the number of iterations is computed using an incorrect value of M, the amplification process may overshoot the target subspace, leading to abrupt drops in success probability. In several cases the performance rapidly collapses to the level of uniform random sampling.

In contrast, the density-aware Grover parametrization exhibits a significantly more stable behavior. Although the success probability decreases gradually as the estimation error grows, the algorithm maintains high success across a wide range of estimation errors.

This robustness can be observed across different density regimes. In the sparse regime ($M = 8$) the proposed algorithm behaves similarly to the standard Grover search near the optimal estimate but degrades more smoothly when the estimate becomes inaccurate. For intermediate and dense regimes ($M = 64$, $M = 128$, and $M = 200$), the difference becomes even more pronounced: while the standard Grover algorithm often collapses to random performance, the density-aware parametrization continues to provide high success probabilities over a broad range of estimation errors.

These results suggest that the proposed parametrization not only stabilizes the behavior of Grover search across different solution densities, but also improves its robustness in situations where the number of marked states is not known exactly.

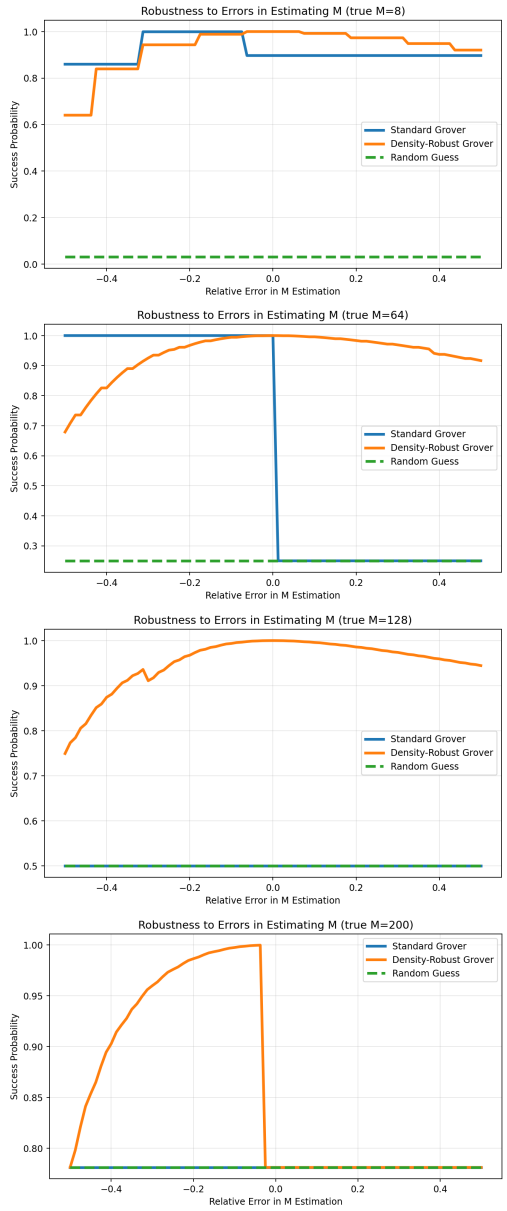

Figure 3: Robustness of the search algorithms with respect to errors in estimating the number of marked states. The plots show the success probability as a function of the relative estimation error $\epsilon = (\widehat{M} - M)/M$ for a search space of size $N = 2^8$. Each panel corresponds to a different number of marked states ($M = 8$, $M = 64$, $M = 128$, and $M = 200$). The blue curve represents the standard Grover algorithm, the orange curve corresponds to the density-aware Grover parametrization proposed in this work, and the dashed green line shows the baseline success probability obtained by uniform random sampling.

6 Discussion and Conclusions

We introduced a density-aware parametrization of the Grover operator in which the phase rotations are computed from the solution density. The resulting operator performs a controlled rotation in the two–dimensional search subspace and adapts naturally to different search regimes. The construction can be interpreted as a concrete instance of Høyer's arbitrary-phase amplitude amplification framework, providing a simple closed-form rule for selecting the phase parameters from the solution density and the desired number of iterations. Our analysis shows that the parametrization preserves the optimal quadratic query complexity of Grover search while improving the stability of the amplification process for higher solution densities. In sparse regimes the algorithm reduces to the standard Grover operator, whereas for denser search spaces the phase-adjusted rotations mitigate overshooting effects. Quantum circuit simulations support these findings, showing smoother amplification dynamics and improved robustness when the number of marked states is only approximately known.

Funding

This work was partially funded by GNCS, Istituto Nazionale di Alta Matematica "Francesco Severi", and the Q-Struct EUROSTART 2025 project, University of Palermo.

References

[1] Andris Ambainis. 2004. Quantum search algorithms. *SIGACT News* 35, 2 (2004), 22–35. doi:10.1145/992287.992296

[2] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. 1998. Tight Bounds on Quantum Searching. *Fortschritte der Physik* 46, 4-5 (1998), 493–505. doi:10.1002/(SICI)1521-3978(199806)46:4/5<493::AID-PROP493>3.0.CO;2-P

[3] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. 2002. Quantum amplitude amplification and estimation. 53–74 pages. doi:10.1090/conm/305/05215

[4] Simone Faro and Francesco Pio Marino. 2025. Scaling Grover's Search for Large Solution Spaces. In *Proceedings of the 34th International Symposium on High-Performance Parallel and Distributed Computing, HPDC 2025*, Robert W. Wisniewski and Ivona Brandic (Eds.). ACM, 43:1–43:8. doi:10.1145/3731545.3744149

[5] Simone Faro, Arianna Pavone, and Caterina Viola. 2024. Quantum Path Parallelism: A Circuit-Based Approach to Text Searching. In *Theory and Applications of Models of Computation - 18th Annual Conference, TAMC 2024, Hong Kong, China, May 13-15, 2024, Proceedings (Lecture Notes in Computer Science)*, Xujin Chen and Bo Li (Eds.). Springer, 247–259. doi:10.1007/978-981-97-2340-9_21

[6] Lov K. Grover. 1996. A Fast Quantum Mechanical Algorithm for Database Search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Philadelphia, Pennsylvania, USA) *(STOC '96)*. ACM, 212–219. doi:10.1145/237814.237866

[7] Lov K. Grover. 1997. Quantum Mechanics Helps in Searching for a Needle in a Haystack. *Physical Review Letters* 79, 2 (1997), 325–328. doi:10.1103/PhysRevLett.79.325

[8] Lov K. Grover. 2005. Fixed-Point Quantum Search. *Physical Review Letters* 95, 15 (2005), 150501. doi:10.1103/PhysRevLett.95.150501

[9] Peter Høyer. 2000. Arbitrary Phases in Quantum Amplitude Amplification. *Physical Review A* 62, 5 (2000), 052304. arXiv:quant-ph/0006031 doi:10.1103/PhysRevA.62.052304

[10] Gui Lu Long. 2001. Grover Algorithm with Zero Theoretical Failure Rate. *Physical Review A* 64, 2 (2001), 022307. arXiv:quant-ph/0106071 doi:10.1103/PhysRevA.64.022307

[11] Gui Lu Long, Yan Song Li, Wei Lin Zhang, and Li Niu. 1999. Phase matching in quantum searching. *Physics Letters A* 262, 1 (Oct. 1999), 27–34. doi:10.1016/s0375-9601(99)00631-3

[12] Christof Zalka. 1999. Grover's Quantum Searching Algorithm is Optimal. *Physical Review A* 60, 4 (1999), 2746–2751. doi:10.1103/PhysRevA.60.2746