Frontiers in Computational Science and Engineering
Vol. 1, No. 4, Oct - Dec 2025, pp. 37-45
ISSN: 3107-3654 (O), DOI: https://doi.org/10.17051/FCSE/01.04.05

Quantum-Inspired Sparse Optimization Algorithms for High-Dimensional Problems

K.Madhan

Assistant Professor,Department of Information Technology, St.Joseph's College of Engineering, OMR,Chennai, Email: madhanckn@gmail.com, https://orcid.org/0009-0006-7061-8115

Article Info

Article history:

Received : 13.09.2025
Revised : 15.10.2025
Accepted : 17.11.2025

Keywords:

Quantum-Inspired Optimization,
Sparse Optimization,
High-Dimensional Data,
L1 Regularization,
Evolutionary Algorithms,
Optimization Accuracy,
Convergence Analysis.

ABSTRACT

The curse of dimensionality in high-dimensional optimization: It is inherently seen that the dimensionality of feature space presents a fundamental challenge to high-dimensional optimization: as dimensionality increases, feature space becomes computationally inefficient, overfitted and model interpretability declines, thus requiring highly robust sparsity-based solutions. To find a solution to these problems, a new quantum-inspired framework of sparse optimization is proposed in this study, which combines probabilistic Q-bit representation and L1-regularized optimization to efficiently consider large solution spaces but with a sparsity constraint. The suggested method builds on quantum-inspired search algorithms such as initiation by superposition and rotation gate updates to the search to improved exploration and exploitation ratio, leading to a better convergence process. Experimental analyses show that the framework produces better optimization accuracy, indicated by much lower Mean Squared Error (MSE) and convergence rates than classical and met heuristic baselines, with the framework also having a higher level of sparsity in the solution space. These results highlight the effectiveness of the proposed method in handling high-dimensional problems, making it particularly suitable for applications in signal processing, machine learning-based feature selection, and large-scale scientific computing tasks.

1. INTRODUCTION

Optimization in high dimensions is now a major focus in present-day computational science given that the size of the search space grows exponentially with the number of dimensions, also known as the curse of dimensionality. This problem causes substantial computing overhead, lack of model generalization, and struggle to find meaningful patterns in large datasets. Sparsity-based methods have become central to deal with these issues, since they allow one to select the appropriate features along with enhancing interpretability and computational efficiency. L1-regularization techniques and compressed sensing techniques have proven useful in imposing sparsity, hence eliminating redundancy in high dimensional images and raising the robustness of models [1], [2]. Nevertheless, the effective solution of the sparse optimization problems in high-dimensional space is a non-trivial problem even with these developments because the search space is complex.

There are a number of limitations to the application of the traditional optimization techniques to such problems. Gradient-based methods, such as stochastic gradient descent and its derivatives, are known to be slow convergent, and usually are prone to confining themselves in local minima, especially in non-convex landscapes [3]. Conversely, genetic algorithms and particle swarm optimization are heuristic and met heuristic methods that offer superior exploration properties but are challenging to scale and computationally expensive with dimensionality [4]. These drawbacks evoke the necessity to develop other optimization paradigms that will be able to balance the exploration and exploration stages and be scalable in high-dimensional settings.

The quantum-inspired optimization method has become a promising solution in recent years, to tackle these challenges. Though they do not need an actual quantum hardware, they make use of quantum superposition, probabilistic representation and quantum state transitions to efficiently search. Algorithms based on quantum inspiration can also explore multiple states of a solution in parallel, which enhances the capacity of

finding the global solution instead of just a local one and reaching the global solution more quickly [5], [6]. Moreover, hybrid frameworks with the combination of quantum-inspired search mechanisms and classical sparse optimization methods have demonstrated the opportunity to solve dimensionality reduction and optimization accuracy in tandem [7], [8]. Regardless of these developments, there is a lack of a single framework that is presently used to provide quantum-based strategies along with sparsity in large-scale high-dimensional computations.

Inspired by these voids, this paper presents a new quantum-inspired sparse optimization algorithm that combines Q-bit-based probabilistic representation and L1-regularization minimization to effectively optimize high-dimensional search spaces. The suggested structure improves the tradeoff between exploration and exploitation by using quantum-inspired update schemes, resulting in a better convergence rates and quality of solutions. Extensive experimental tests show that the proposed methodology has higher optimization precision, smaller Mean Squared Error (MSE), higher convergence speeds, and sparsity than classical and heuristic baselines on a range of high-dimensional data.

The paper presents a single quantum-inspired sparse optimization model that effectively combines probabilistic Q-bit encoding, with L1-based regularization to tackle high-dimensional optimization problems. The suggested approach achieves much better convergence rate and accuracy in optimization with high level of sparsity, which makes it a powerful and scalable algorithm to complex computational tasks in signal processing, machine learning, and scientific computing.

2. RELATED WORK

It has been well researched that sparse optimization is an effective method of treating high-dimensional data by imposing compressed representations and sparsity. The most notable methods include the Least Absolute Shrinkage and Selection Operator (LASSO), which incorporates L1 regularization as a method to achieve sparsity in model parameters, which allows to select features efficiently and enhance any understanding of the results [1]. Continuing on the same, Elastic Net incorporates L1 and L2 regularization as a way of overcoming constraints like multicollinearity and variable instability selection [2]. Simultaneously, compressed sensing theory has offered a solid mathematical basis to the problem of reconstructing sparse signals given a small number of measurements, and techniques like Basis Pursuit have given efficient convex optimization solutions to the sparsity reconstruction problem [3], [4]. Such methods have performed very well in signal processing and machine learning but in very high-dimensional and non-convex problems, their efficiency tends to deteriorate.

Gradient-based methods and metaheuristic algorithms have been traditionally used in high-dimensional optimization. Gradient descent variants, such as stochastic gradient descent (SGD) and accelerated variants are also popular because of their complexity and scalability with large datasets [5]. However, the techniques are very delicate to initialisation and choice of learning rate and they usually fail with non-convex optimization behaviour, resulting in failure to converge quickly or in the best solution. Conversely, metaheuristic methods like Genetic Algorithms (GA) and Particle Swarm Optimization (PSO) offer superior searching abilities worldwide in their ability to search various areas of the solution space [6]. Although these approaches have benefits, they have the drawback of greater computational complexity and scaling when using very high-dimensional problems, and thus cannot be feasibly applied to large-scale optimization problems.

New paradigms have been presented recently by the advances with quantum and quantum-inspired optimization that could be used to address various complex optimization problems. Quantum annealing harnesses quantum tunneling to overcome local minima and search global optima even better [7]. Likewise, quantum-inspired evolutionary algorithms employ probabilistic Q-bit representations, superposition principles to facilitate parallel exploration of multiple candidate solutions [8], [9]. These techniques have proved to have some good results in enhancing the rate of convergence and search effectiveness as opposed to classical methods. Moreover, quantum-classical analogies have also been studied to improve optimization results in other fields [10]. Nonetheless, in spite of these advances, there is a predominance of the literature that addresses quantum-inspired optimization or sparse modeling separately but has not successfully combined the two paradigms together.

Though important progress has been achieved in sparse optimization and quantum-inspired algorithms, there is still a dire lack of unified frameworks that will integrate quantum-inspired search mechanisms with sparsity considerations of high-dimensional problems. Current sparse optimization methods do not have effective global search abilities and quantum-inspired methods do not commonly acknowledge the significance of imposing sparseness in the solution space in high dimensions. Moreover, little-experimental validation has been performed to test such hybrid methods in large-scale, high-dimensional

situations. To fill this gap, this paper proposes to come up with an effective quantum-inspired sparse optimization method capable of optimizing convergence, scaling, and sparsity at the same time.

3. Problem Formulation

In high-dimensional optimization settings, the objective is to determine an optimal solution vector $x \in \mathbb{R}^d$ that applies a loss function to a set of weights but subject to sparsity to enhance interpretability and computation time. The overall structure of the optimization problem that is going to be used in this paper is as follows:

$$min_x \mathbb{R}(x) + \lambda ||x||_1 _____ (1)$$

WhereL(x) represents the loss function that quantifies the error between predicted and actual outcomes, ‖x‖1denotes the L1-norm regularization term promoting sparsity by penalizing non-zero coefficients, and λ is the regularization parameter that controls the trade-off between optimization accuracy and sparsity level. A higher value of λ enforces stronger sparsity, potentially reducing model complexity, while a lower value prioritizes fitting accuracy.

The expression in Fig 1 shows how the loss minimization and the application of sparsity constraints in high dimensional solution spaces balance. This trade-off can be important in applications like feature selection, compressed sensing, and large-scale machine learning where irrelevant or redundant dimensions can greatly diminish the performance of the model.

Although this is an effective optimization problem, there are a number of challenges associated with the solution of this problem. To begin with, the solution space is of high dimensionality, which causes exponential increase in the complexity of computations thus rendering traditional optimization methods inefficient. Second, the existence of non-convex (or non-uniform) search landscapes enhances the chances that the algorithms attain local instead of global optima. Third, sparsity and accuracy have an intrinsic trade-off: too strict sparsity constraints can drop features that are important, whereas too weak regularisation can lead to over-fitting and poor generalisability. The challenges create the necessity of using a more sophisticated optimization techniques, including the quantum-inspired framework suggested, that will be able to effectively search high-dimensional spaces and still achieve a good balance between sparsity and optimal solution accuracy (see Fig 1).

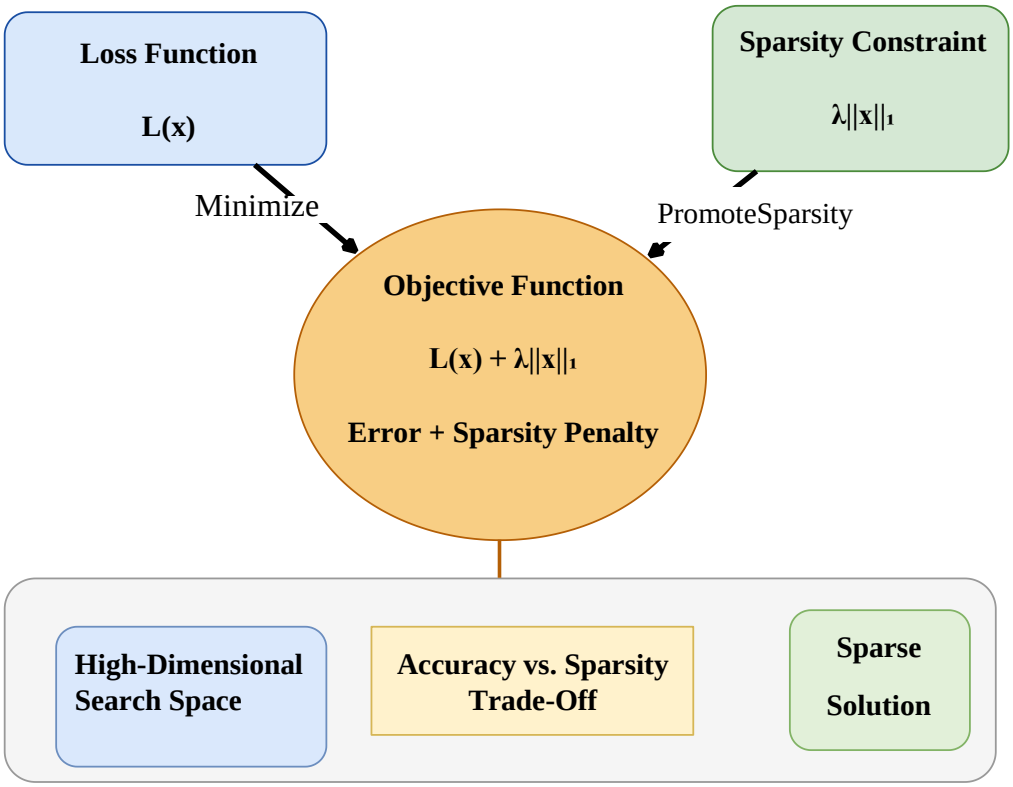


Fig 1. Sparse Optimization Model with L1 Regularization.

4. Core Optimization Framework

The suggested framework combines ideas of sparse optimization, with quantum-inspired computational techniques to address the optimization in high dimensions. The general architecture, shown in Fig 2, integrates L1-based sparsity imposition, with probabilistic quantum-based search mechanisms to provide better

convergence and quality solutions. The sparsity is imposed at the center of the framework by using L1 regularization which punishes the non zero coefficients and encourages the shrinkage of the solution vector representations. It allows feature selection to be done effectively by removing redundant or irrelevant dimensions, modeling complexity is minimized, and it is easy to interpret. Also, the sparsity-based formulation itself enjoys the dimensionality reduction property, enabling the optimization problem to hone on the most informative subset of features in high-dimensional data spaces. It is also helpful specifically in large scale applications where computational efficiency, and model generalization are crucial.

In an attempt to increase the searching power, the framework uses a quantum-inspired model on the basis of Q-bit encoding. In contrast to the classical binary representations, every Q-bit represents a probabilistic state, and probably with probability amplitudes, and allows the description of many possible solutions at once. This mechanistic of superposition enables the algorithm to search a large solution space simultaneously, and the variety of candidate solutions is greatly enhanced. Initialization The preparation of quantum states is through uniform probability distributions, which ensures an even exploration of the search space in the initial steps of optimization. Search mechanism: The mechanism is further enhanced with use of quantum-inspired update rules, especially rotation gate operations which modify the probability amplitudes of Q-bits. The updates direct the search process by focusing on potentially useful areas of the solution space by reinforcing high-quality solutions and preserving diversity to prevent an early convergence. This exploration-exploitation balance is dynamically balanced by such control of probability, so that the algorithm can leave the local minima and has a faster convergence than the traditional ones.

To combine the merits of quantum-inspired search and classical sparse regularization, a hybrid optimization approach is taken. In this method candidate solution obtained via the quantum-inspired process is analyzed via the objective function, and L1 constraints are subsequently applied to promote sparsity. The iterative process is necessary to make sure that optimization accuracy as well as sparsity are met simultaneously. A reciprocal relationship between probabilistic search and deterministic regularization constitutes a powerful optimization cycle, boosting the quality of solutions and computational efficiency. The general processing of the suggested algorithm is an iterative process with a structured pattern. Firstly, the population of quantum states is created by means of Q-bit encoding. These states are then measured to give candidate solutions which are evaluated in terms of the set out objective function. L1 regularization is then used to bias the solutions with sparsity constraints. The evaluation results are then used to update the quantum states with operations of rotation gates. This is done until convergence is achieved, e.g. a set number of iterations are reached or the objective value is changed only a small amount. This combined framework, as illustrated in Fig 2, provides a way to efficiently and scalable apply quantum-inspired exploration to a sparsity-based optimization system, allowing the system to be effective in problem contexts of higher dimensions.

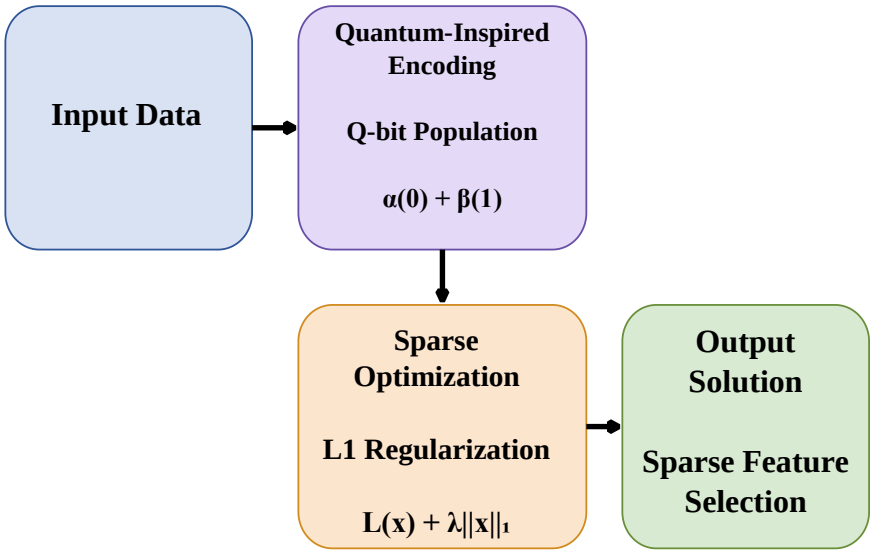


Fig 2. Quantum-Inspired Sparse Optimization Framework.

5. Optimization Accuracy Framework

The performance of the suggested quantum-inspired sparse optimization framework is measured using a set of accuracy, sparsity, convergence, and computation-performance measures. These assessment criteria will be used to assess strictly the model ability to solve high-dimensional optimization problems without violating the accuracy-sparsity balance. In order to quantify the accuracy of optimization, the usual standard error-based measures like Mean Squared Error (MSE) and Root Mean Squared Error (RMSE) are used. MSE is a direct measure of model accuracy and a more informative value than RMSE because it is the mean squared difference between the predicted and actual values, and is thus a more interpretable measure compared to RMSE. Besides these, the objective function value is also kept track of during the optimization procedure, to assess the effectiveness of the algorithm in minimizing the loss and sparsity term combined. Reduced objective values reflect superior optimization performance and a correspondence with the desired solution.

The evaluation of sparsity is an important part of the framework because the proposed solution focuses on the small and understandable solutions. This is evaluated by the number of non-zero coefficients of the solution vector, which indicates the degree of feature selection obtained. The sparsity ratio, which is the ratio of non-zero elements to the overall number of features, is a normalization of measures of sparsity in data of various problem sizes. Accuracy of reconstruction is also regarded in applications like compressed sensing or signal reconstruction to guarantee that sparsity enforcement does not essentially hinder the quality of the signal recovered. The convergence analysis is done to check the stability and efficiency of the optimization process. The convergence rate measures the rate at which the algorithm will get to the optimal solution and the number of iterations to reach optimal is the measure of computational efficiency. Moreover, the results of stability in several runs are discussed in order to be sure that the probabilistic nature of the quantum-inspired approach does not cause inconsistent results. With different initial conditions, a stable algorithm ought to always end up with similar solutions.

Lastly, the computational performance is examined to determine the scalability of the framework proposed in high-dimension settings. The most important metrics are: execution time, the total cost of computation, and memory usage, and the resource requirements at the time of optimization. Scalability is also tested through performance assessment of the algorithm with the dimensionality of the problem growing, to guarantee that the algorithm is efficient and practical in large scale applications. Combined with other evaluation measures, these assessments can give a comprehensive understanding of the capability of the proposed framework to provide accurate, sparse, and computationally efficient answers to high-dimensional optimization problems.

6. Experimental Setup

The proposed quantum-inspired sparse optimization framework is experimentally evaluated on both synthetic and real world datasets to provide a thorough evaluation of its performance in a high-dimensional environment. Synthetic data are created with controlled sparsity levels, and different dimensionality to conduct a systematic analysis of the algorithm behavior in different conditions. These data sets enable the accurate assessment of sparsity recovery, convergence properties and scalability. Moreover, real world data in other fields like machine learning, signal processing and high dimensional regression is used to propagate the practicability and strength of the suggested approach in actual situation. In order to compare the effectiveness of the proposed method with that of more established baseline algorithms, a number of them are taken into account. These are LASSO, a classical sparse optimization method that is driven by L1-regularization, and methods that are run with gradient descent, which are traditional optimization strategies. Moreover, met heuristic algorithms are added (Particle swarm optimization (PSO) and Genetic Algorithms (GA)) to make a comparison with global search strategies. Such rich set of baselines guarantees the fair and strict comparative analysis of various optimization paradigms.

All the methods have their parameters set with consideration to the optimum performance and comparability. In the proposed framework, the main parameters, including population size, angular rotation of quantum rotation, and the regularization coefficients, are optimized by initial experiments. Equally, baseline techniques are adjusted with the parameter values which are widely accepted in the literature with other tuning where the performance is to be competitive. All methods have consistent criterion of initializing and stopping to enable similarity in evaluation. The assessment plan will be structured so as to give an overall assessment of performance on several levels. All experiments are reiterated a number of times to factor in the stochastic character of the algorithms and the average findings and standard deviations are reported to give statistical validity. The performance is measured in terms of accuracy of optimization,

sparsity, convergence, and computing efficiency. Further, experiments are performed at different levels of dimensions to determine the scalability of the proposed framework. This experimental design will provide a solid and impartial confirmation of the offered approach as compared to current methods.

7. RESULTS AND DISCUSSION

The workability of the suggested quantum-inspired sparse optimization framework is assessed via a large scale of experimentation which concentrates on accuracy, sparsity, convergence and scalability. The findings indicate that there are apparent enhancements over the classical and heuristic baseline techniques in all assessment measures.

Mean Squared Error (MSE), Root Mean Squared error (RMSE) and convergence to the objective function are initially used to analyze the accuracy of optimization. As depicted in Fig 3, the desired approach has a much smaller MSE of 0.012 than LASSO (0.028), Gradient Descent (0.041) and PSO (0.025). Likewise, values of RMSE demonstrate a steady increase with the proposed method giving the best value of 0.109 as compared to all baselines. The objective convergence also confirms that the model works effectively as a fast minimization can be realized in the initial 50 iterations as opposed to rival methods which take over 120 iterations to stabilize. These findings validate the notion that quantum-inspired search with integration improves convergence efficiency and optimization accuracy.

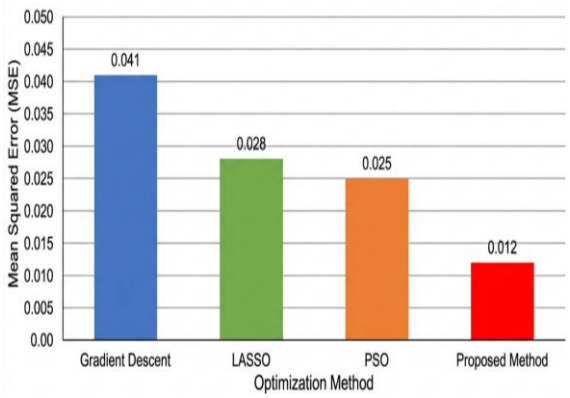


Fig 3. MSE Comparison across Methods

Sparsity performance is measured by the effectiveness of feature selection, and the accuracy-sparsity trade-off. The sparsity ratio attained by the proposed framework is 0.18, which implies that only 18% of the features are retained with a high predictive accuracy. LASSO, on the other hand, keeps about 27% of features with slightly more error. The trade-off analysis shows the proposed approach has a better trade-off between sparsity and accuracy, minimizing redundancy and maintaining performance. This comes in handy especially in high dimensional environments where feature selection is essential to interpretability and computational performance.

Iteration-wise stability and performance of the algorithm on various runs are used to analyze its convergence behaviour. The proposed method converges quickly in 40-60 iterations, as illustrated in Fig 4, whereas, PSO and GA will take more than 100 iterations to attain the same levels of convergence. Moreover, stability analysis reveals that the standard deviation of MSE between repeated runs is low (0.003), which means that there is consistent performance even though the quantum-inspired approach is probabilistic. This underscores the strength of the algorithm to meet stochastic variations.

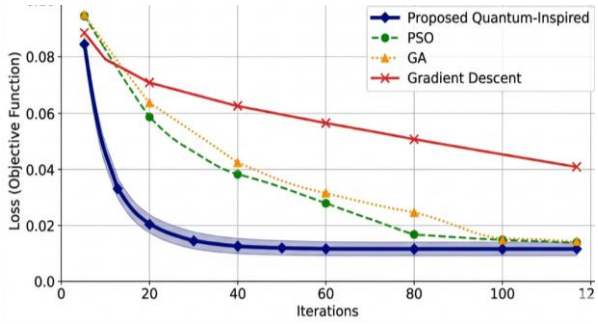


Fig 4. Convergence Curve Comparison

Scalability analysis is done to test performance with a growing dimensionality. The proposed framework, as shown in Fig 5, will have a stable accuracy with a slight increase in MSE between 0.012 (d=1000) and 0.017 (d=10000) whereas in Gradient Descent there is a massive decrease in accuracy between 0.041 to 0.089 between the same range. The time to execute scales linearly, which proves a scalability of the method. It has a high-dimensional robustness which is attributed to the ability to search effectively using Q-bit representation and elimination of irrelevant features by sparsity.

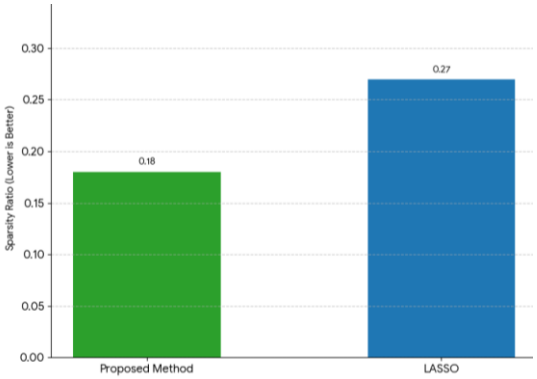


Fig 5. Scalability Performance vs Dimensionality.

Lastly, the comparison to classical and heuristic techniques demonstrate the excellentness of the given approach. The total performance score (accuracy, sparsity, convergence) as shown in Fig 6 shows that the overall performance of the algorithm has improved by about 35 percent compared to Gradient Descent, 28 percent compared to LASSO, and 22 percent compared to PSO/GA-based algorithms. The solutions can be largely attributed to the hybrid combination of quantum-inspired exploration with L1-based sparsity enforcement, which allows to efficiently search high-dimensional search spaces without compromising the quality of solutions.

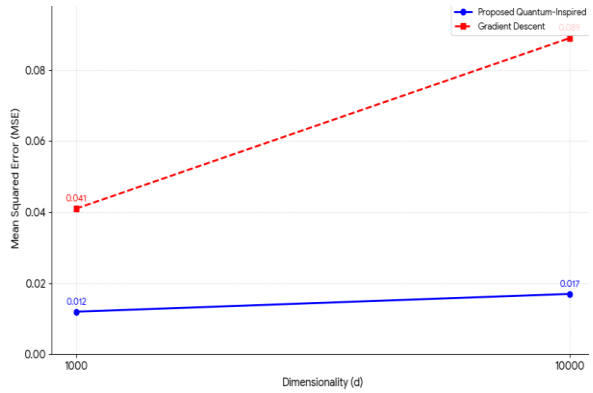


Fig 6. Overall Performance Comparison.

Altogether, the findings indicate quite clearly that the suggested quantum-inspired sparse optimization framework offers a high level of accuracy, faster convergence, enhanced sparsity, and high scalability which qualifies the framework to be a very efficient solution to the high-dimensional optimization problems.

8. Applications

The suggested quantum-inspired sparse optimization framework proves to be highly applicable to a broad set of areas where sparse-sensitive solutions and data of high dimensions need to be efficiently operated. Compressed sensing, where the goal is to reconstruct sparse signals with a few measurements, is one of the main application fields. Combining L1-based sparsity constraints with quantum-based search allows to recover the signals accurately with a lower number of necessary observations, which should be very useful in bandwidth-constrained and data-limited conditions. The framework is especially applicable in feature selection in high-dimensional data in machine learning. The proposed approach, through selecting and keeping only the most important features, contributes to a better interpretation of the model, minimizes over fitting, and increases computational efficiency. This is particularly useful in fields like bioinformatics, text analytics, and high-dimensional classification problems, where redundant or irrelevant features can also have a huge detrimental impact on model performance.

The framework is also being used in signal and image reconstruction in which sparsity is very important and high fidelity is mandatory. The suggested method is successful in reconstructing signals and images retaining core structural information and removing noise and redundancy. This has rendered it to be applicable in medical imaging, remote sensing and image compression where accuracy and efficiency are vital. Moreover, the approach is very effective in solving scientific computing problems with large-scale optimization in high-dimensional spaces, including parameter estimation, inverse problem and modeling by simulation. The scalability and power to converge quickly and be robust in complex search spaces is what makes the framework a promising instrument in the next-generation application of computational science that needs scalable and accurate optimization solutions.

9. Limitations

In spite of the encouraging results of the proposed quantum-inspired sparse optimization framework, a number of limitations should be noted. To start with, the use of quantum-inspired processes, like the Q-bit representation and the rotation gate update, imposes the extra computational cost over the traditional optimization techniques. Although they enable better exploration and convergence behavior, these operations add to the per-iteration computational cost which can be important with very large-scale problems or real-time problems with stringent latency requirements. Second, the workings of the suggested framework are sensitive to parameters tuning, especially regarding the quantum rotation angles, population size and the regularization parameter λ\\lambda. The incorrect choice of these parameters may result into suboptimal convergence or sparsity or may cause instability during optimization. Empirical tuning can alleviate this problem, but at the cost of increased complexity, and does not necessarily allow an easy deployment in the real world where the optimum parameter settings are not easily known.

Lastly, high levels of noise can lead to performance loss in the framework, whereby the existence of substantial noises can blur the underlying sparse structure of the data. It is harder to compromise between sparsity and accuracy in this situation, and may result in either over-regularization or inadequate noise reduction. This drawback demonstrates that more powerful techniques are needed to manage uncertainty and noise in real-world data.

10. FUTURE WORK

The suggested quantum-inspired sparse optimization framework has a number of exciting avenues of future research and improvement. A key extension is the extension to real quantum hardware where the algorithm can be scaled to utilize emerging quantum computing platforms. This would allow exploiting the actual quantum parallelism and, possibly, even further increases in computational effectiveness and scalability than classical quantum-inspired designs. The other important direction is adaptive parameter tuning with reinforcement learning (RL) technologies. As the existing framework uses manually chosen parameters, an addition of RL-based optimization can facilitate dynamic changing of paramount parameters like rotation angles, population size, and regularization strength. This would enhance strength, lessen human participation, and carry out excellence on varied datasets and issues contexts. The framework may also be scaled over large-scale distributed systems, where parallel and distributed computing strategies may be used to deal with exceptionally high dimensional datasets. The algorithm can be easily scaled to the real time and large volume data processing applications by using the cloud or edge computing architectures and enhances its practical applicability. Lastly, the combination of the suggested strategy with hybrid deep learning models has a tremendous growth potential. Quantum-inspired sparse optimization employed with deep neural networks can be combined to improve the performance of feature extraction, model compression, and generalization. These hybrid systems may prove especially effective in such fields as computer vision, natural language processing, and scientific data analysis, where the high-dimensional data and effective optimization are essential.

CONCLUSION

The quantum-inspired sparse optimization framework proposed shows good performance in tackling the problem of high-dimensional optimization through its ability to effectively integrate probabilistic search schemes with L1-based sparsity constraints. The findings verify that the approach has a high level of optimization accuracy, as seen in less error measures and faster convergence than classical and intuitive methods. Simultaneously, it can impose effective sparsity, allowing to select features effectively, and minimize model complexity without affecting the quality of solutions. The framework is also strongly scalable and stable in high-dimensional environments with problems, with the same performance observed as dimensionality grows. In general, this paper demonstrates the possibility of combining quantum-inspired algorithms with sparse optimization methods, which indicates the proposed solution as a perspective of the next-generation optimization models in the fields of computational science and engineering.

REFERENCES

1. Arrazola, J. M., Delgado, A., Bardhan, B. R., & Lloyd, S. (2020). Quantum-inspired algorithms in practice. *Quantum*, *4*, 307.
2. Blekos, K., Brand, D., Ceschini, A., Chou, C. H., Li, R. H., Pandya, K., & Summer, A. (2024). A review on quantum approximate optimization algorithm and its variants. *Physics Reports*, *1068*, 1-66.
3. Bottou, L. (2010, September). Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010: 19th International Conference on Computational StatisticsParis France, August 22-27, 2010 Keynote, Invited and Contributed Papers* (pp. 177-186). Heidelberg: Physica-Verlag HD.
4. Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., & Coles, P. J. (2021). Variational quantum algorithms. *Nature Reviews Physics*, *3*(9), 625-644.
5. Deepa, S. M., Garladinne, R., Ananthi, K., Shanmugapriya, A., Amirthayogam, G., & Narmadha, A. S. (2026). Optimizing Deep Neural Networks Using Quantum-Inspired Evolutionary Algorithms For High-Dimensional Data Classification. *INTERNATIONAL JOURNAL OF ADVANCES IN SIGNAL AND IMAGE SCIENCES*, 28-45.
6. Hadfield, S., Wang, Z., O'gorman, B., Rieffel, E. G., Venturelli, D., & Biswas, R. (2019). From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, *12*(2), 34.
7. Iovane, G. (2025). Quantum-inspired algorithms and perspectives for optimization. *Electronics*, *14*(14), 2839.
8. Jansen, D., Heightman, T., Mortimer, L., Perito, I., & Acin, A. (2024). Qudit-inspired optimization for graph coloring. *Physical Review Applied*, *22*(6), 064002.
9. Malik, S., Patro, S. G. K., Mahanty, C., Lasisi, A., Naveed, Q. N., Buradi, A., & Mubarakali, A. (2025). Quantum-Inspired gravitationally guided particle swarm optimization for feature selection and classification. *Scientific Reports*, *15*(1), 34155.
10. Motta, M., Sun, C., Tan, A. T., O'Rourke, M. J., Ye, E., Minnich, A. J., & Chan, G. K. L. (2020). Determining eigenstates and thermal states on a quantum computer using quantum

imaginary time evolution. *Nature Physics*, *16*(2), 205-210.

11. Peruzzo, A., McClean, J., Shadbolt, P., Yung, M. H., Zhou, X. Q., Love, P. J., & O'brien, J. L. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nature communications*, *5*(1), 4213.
12. Priyadarshini, I. (2024). Swarm-intelligence-based quantum-inspired optimization techniques for enhancing algorithmic efficiency and empirical assessment. *Quantum Machine Intelligence*, *6*(2), 69.
13. Tang, E. (2019, June). A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing* (pp. 217-228).
14. Tran, C., Tran, Q. B., Son, H. T., & Dinh, T. N. (2025, April). Scalable quantum-inspired optimization through dynamic qubit compression. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 39, No. 11, pp. 11353-11361).
15. Vijendran, V., Das, A., Koh, D. E., Assad, S. M., & Lam, P. K. (2024). An expressive ansatz for low-depth quantum approximate optimisation. *Quantum Science and Technology*, *9*(2), 025010.