



Universiteit
Leiden
The Netherlands

Hybrid quantum-classical metaheuristics for automated machine learning applications

Von Dollen, D.J.

Citation

Von Dollen, D. J. (2025, November 18). *Hybrid quantum-classical metaheuristics for automated machine learning applications*. Retrieved from <https://hdl.handle.net/1887/4282905>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4282905>

Note: To cite this publication please use the final published version (if applicable).

Chapter 7

Conclusions

7.1 Summary

Chapter 1 In this chapter, we provide a high-level introduction to this work and introduce some of the challenges in machine learning. We discuss some of the theoretical speedups for quantum computing as motivation for our research work. We outline our research questions and give an overview of publications in contribution to this thesis.

Chapter 2 We cover the topics and preliminaries of our research work as part of this chapter for use in later chapters. These preliminaries are foundational to the subsequent work upon which this thesis is built.

Chapter 3 We introduce the feature selection problem in machine learning and investigate how near-term quantum computers may be applied in context to the feature selection problem in response to the first research question:

RQ1 Can we observe qualitative or quantitative performance gains in using quantum, quantum inspired, or hybrid quantum-classical algorithms to identify and search correlated data?

This topic also covered the second research question:

RQ2 Can access to quantum or quantum inspired computation generate qualitative or quantitative performance gains in preprocessing steps such

7.1. Summary

as feature selection and feature engineering in building machine learning models?

In response to these research questions, we develop a method for feature selection, framing the selection process as a QUBO, using a heuristic of maximum relevancy-minimum redundancy. We test classical correlation and distance measures, including the Pearson correlation coefficient, the generalized mean information coefficient, the maximal information coefficient, and mutual information.

We encode these distance measures into a QUBO and empirically test our method of feature selection against other well-known methods such as recursive feature elimination and greedy selection. We validate our selection methods for regression estimators in the Friedman 1 dataset [70] and apply our method to the problem of price prediction using the UCI auto dataset [183]. The estimators that we test include multiple linear regression and gradient boosted regression trees. We show that our method outperforms and achieves significant results in comparison to greedy selection, recursive feature elimination, and using all features in application towards these supervised learning tasks. For the Friedman 1 dataset, our method achieves the lowest scores (quantum-assisted Pearson's correlation coefficient-Multiple Linear Regression) with the lowest average $MAE = 2.27$ (p-value=0.00001). For the UCI Auto dataset, our method using gradient boosted trees (Quantum-Assisted Pearson correlation coefficient gradient boosted trees) achieves the lowest scores $MAE = 1471$ (p-value = 0.001).

We identify additional benefits of our method, from the explainability for machine learning models using this selection method to the automatic discovery of subset selection sizes for feature selection. We note that depending on the choice of estimator, or method, may help improve the runtime of the give estimator, by selecting $k < m$ features for which to train and test the estimator. Additionally, we identify this feature selection method as an application area for quantum computing towards an automated machine learning process.

Chapter 4 We develop a method for defining selection operators for evolutionary algorithms as QUBOs and study solving these with quantum devices and heuristics. This is in response to the third research question:

RQ3 Can access to fully quantum, quantum inspired, or hybrid quantum-classical computation generate qualitative or quantitative performance gains in optimization metaheuristics?

Building on our method of selection via QUBO, we show that a selection operator for evolutionary algorithms, such as a genetic algorithm, can be formulated. Starting with pseudo-Boolean optimization of perturbed and noisy objective functions, we study the efficacy of our approach and show significantly increased performance gains in quality diversity and fitness over ten out of fifteen test functions. We also test the classical, hybrid quantum-classical, and quantum samplers and characterized the effect of scaling linear and quadratic terms for the QUBO. We show that the hybrid quantum-classical sampler achieved the best results by examining the distributions of minimum energies of solutions returned by sampling. We identify how our method may be used for automated machine learning processes such as model selection and hyperparameter optimization.

Chapter 5 We extend our method from Chapter 4 to evolution strategies and examine a variety of learning tasks and the hyperparameter optimization problem. This was also in response to the third research question:

RQ3 Can access to fully quantum, quantum inspired, or hybrid quantum-classical computation generate qualitative or quantitative performance gains in optimization metaheuristics?

Building on work from the previous chapter, we extend our quantum-enhanced selection operator to continuous domains, from genetic algorithms to evolution strategies. We first test our method over a variety of multimodal black-box functions. Our method achieved the best average fitness in 94% of test cases, though the magnitude of improvement varied between problems. We then apply our method towards hyperparameter optimization of deep neural networks over 34 tabular data benchmarks, and show that our method achieves best performance in the form of validation accuracy and rank. We also test our method towards the application of neuroevolution of feedforward neural networks and show that our method outperforms baselines in reinforcement learning applications.

Chapter 6 In this chapter, we examine the role of quantum data, by creating stilted data from a projected quantum kernel and studying quantum kernel methods in the context of sparse Gaussian process regression for Bayesian optimization. We apply these methods to the hyperparameter optimization problem. This is in response to the third research question:

7.1. Summary

RQ3 Can access to fully quantum, quantum inspired, or hybrid quantum-classical computation generate qualitative or quantitative performance gains in optimization metaheuristics?

and the fourth research question:

RQ4 Can we observe a performance difference between classical and quantum models for training and inference on classical and quantum data?

In response to these research questions, we identify that our method may be applied as a means to induce sparsity for Gaussian process regression as a surrogate model for Bayesian optimization. We show a reduction in the asymptotic complexity of fitting the Gaussian process from $O(n^3)$ to $O(m^3)$. We acknowledge the additional cost of constructing and solving the QUBO, and identify that these costs may be offset in settings where the evaluation of the underlying objective function may be prohibitively expensive. We develop a few new methods in addition to our approach; we use our QUBO-enhanced method to filter random Fourier features which we generate by sampling from the spectral density of the Gaussian kernel. We also use a quantum feature map as the kernel and apply our QUBO filtering method to reduce the size of the Gram matrix for the Gaussian process.

We tested these methods over classical data sets and show performance gains in the form of a lower validation mean squared error MSE for the sparse estimators for the surrogate model. We created stilted datasets using a projected quantum kernel to explore the best-case scenarios. On these quantum-engineered datasets, quantum feature map methods achieved lower validation MSE , suggesting domains where quantum approaches might prove beneficial if naturally quantum-generated data becomes available. This approach was designed to facilitate controlled experimentation and illustrate where quantum machine learning models might show applicability to quantum-generated data sets. Quantum feature maps also exhibited slower eigenvalue decay rates compared to classical kernels on varying length scales, which may indicate an enhanced ability to capture complex data relationships.

The search for features, models, and hyperparameter spaces can be complex and high-dimensional. This thesis shows that near-term quantum computers such as the D-Wave quantum annealer, quantum kernel methods, and quantum inspired methods incorporating QUBO may be applied towards subroutines for hybrid quantum-classical heuristics to improve this search as part of an AutoML process. For feature selection, we introduced a maximum relevancy-minimum redundancy formulation that accommodates multiple correlation measures, offering explainability benefits by directly

encoding relevance and redundancy in the optimization objective. The framework extends to evolutionary computation through QUBO-based selection operators applicable to both discrete genetic algorithms and continuous evolution strategies, which we show application areas for reinforcement learning and neuroevolution. We also developed several techniques for quantum-enhanced sparse Gaussian process regression, including methods that filter random Fourier features and apply quantum feature maps combined with QUBO-based dimensionality reduction to compress Gram matrices while attempting to preserve predictive performance. Through empirical studies employing QUBO-based techniques and quantum kernel feature maps across three different domains related to automated machine learning, we observed performance improvements in a majority of test cases. We acknowledge that some of the experiments in this thesis were conducted using classical simulation rather than quantum hardware, due to limited access to quantum processing units. These results suggest potential benefits, but could also benefit from future validation with error-corrected quantum hardware.

An important pattern emerged: hybrid quantum-classical methods often outperformed both purely classical heuristics and fully quantum approaches. This suggests that the value may lie in the formulation of the QUBO problem itself, as classical optimization of these quantum-inspired formulations proved to be effective. In evolutionary algorithm applications, hybrid samplers tended to achieve better distributions of minimum energies compared to classical simulated annealing or direct quantum annealing, suggesting that practical quantum advantage may lie in thoughtful integration rather than complete replacement of classical computation.

Integrating these components, we developed an AutoML pipeline that spanned quantum-enhanced feature selection, quantum-enhanced evolutionary optimization, QUBO-enhanced hyperparameter tuning, and quantum-inspired surrogate modeling, with each stage undergoing empirical validation on benchmark problems and real-world datasets. The resulting framework provides both theoretical motivation and implementation guidance for hybrid quantum-classical AutoML systems, illustrating how near-term quantum devices might augment rather than replace classical machine learning infrastructure and offering a template for quantum enhancement of automated machine learning workflows.

7.1. Summary
