

Hybrid Classical-Quantum AI Approach for Detecting Cyberattacks in Vehicles

Final Report

By

Shaozhi Li
Clemson University

Sumanta Tewari, Clemson University
Mashrur Chowdhury, Clemson University
M Sabbir Salek, Clemson University
Vaneet Aggarwal, Purdue University
Satish Ukkusuri, Purdue University
Gurcan Comert, Benedict College

December 2025



**NATIONAL CENTER FOR TRANSPORTATION
CYBERSECURITY AND RESILIENCY (TraCR)**





DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by the National Center for Transportation Cybersecurity and Resiliency (TraCR) under Grant No. 69A3552344812 and 69A3552348317 which is headquartered at Clemson University, South Carolina, USA, from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

Non-exclusive rights are retained by the U.S. DOT.

CONTACTS

For more information:

M Sabbir Salek
National Center for Transportation
Cybersecurity and Resiliency (TraCR)
Clemson University
Phone: 865-207-0537
Email: msalek@clemson.edu

Sumanta Tewari
Department of Physics and Astronomy
Clemson University
Phone: 864-656-5321
Email: stewari@clemson.edu

TraCR
Clemson University
One Research Dr
Greenville, SC 29607
tracr@clemson.edu



ACKNOWLEDGMENT

This work is based upon the work supported by the National Center for Transportation Cybersecurity and Resiliency (TraCR), a U.S. Department of Transportation National University Transportation Center headquartered at Clemson University, Clemson, South Carolina, USA. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of TraCR. The U.S. Government assumes no liability for the contents or use thereof.



Technical Report Documentation Page

1. Report No. 14		2. Government Accession No. N/A		3. Recipient's Catalog No. N/A	
4. Title and Subtitle Hybrid Classical-Quantum AI Approach for Detecting Cyberattacks in Vehicles				5. Report Date: Dec 2025	
				6. Performing Organization Code: N/A	
7. Author(s) Shaozhi Li, Ph.D.; https://orcid.org/0000-0002-5432-6802 Sumanta Tewari, Ph.D.; https://orcid.org/0000-0002-9979-6306 Yao Wang, Ph.D.; https://orcid.org/0000-0003-1736-0187 Mashrur Chowdhury, Ph.D.; https://orcid.org/0000-0002-3275-6983 M Sabbir Salek, Ph.D.; https://orcid.org/0000-0001-7326-3694 Vaneet Aggarwal, Ph.D.; https://orcid.org/0000-0001-9131-4723 Satish Ukkusuri, Ph.D.; https://orcid.org/0000-0001-8754-9925 Gurcan Comert, Ph.D.; https://orcid.org/0000-0002-2373-5013				8. Performing Organization Report No. 14	
9. Performing Organization Name and Address National Center for Transportation Cybersecurity and Resiliency (TraCR), Clemson University, 414 A One Research Dr, Greenville, SC 29607				10. Work Unit No. N/A	
				11. Contract or Grant No. 69A3552344812 and 69A3552348317	
12. Sponsoring Agency Name and Address U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology, 1200 New Jersey Avenue, SE, Washington, DC 20590				13. Type of Report and Period Covered Final Report, 01/01/2024 - 05/31/2025	
				14. Sponsoring Agency Code OST-R	
15. Supplementary Notes Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.					
16. Abstract Machine learning has become central to self-driving vehicles, enabling tasks such as object detection and anomaly recognition. Yet classical methods often face challenges of high complexity, large parameter counts, and limited training efficiency. To address these, we propose quantum and quantum-inspired machine learning approaches that enhance both efficiency and resilience. We first develop a quantum-inspired weight-constrained neural network that requires substantially fewer parameters than classical counterparts, reducing energy costs for training and inference. Combined with a novel dropout-based defense, this model achieves robustness under strong adversarial attacks, with an average accuracy reduction of about 20%. We also explore hybrid quantum-classical convolutional neural networks using angle encoding and quantum activation functions, demonstrating faster training than classical convolutional neural networks. Motivated by these findings, we introduce quantum-inspired activation functions for classical models, which consistently lower training steps compared to standard Tanh. To capture long-range dependencies, we present the Quantum Non-Local Neural Network (QNL-Net), which integrates classical dimensionality reduction with quantum circuits. QNL-Net achieves state-of-the-art binary classification on MNIST and CIFAR-10 while requiring fewer qubits. Beyond supervised learning, we propose the Quantum Natural Policy Gradient (QNPG) algorithm, achieving improved sample complexity $O(\epsilon^{-1.5})$, surpassing the classical lower bound of $O(\epsilon^{-2})$. Finally, we formulate the Vehicle Routing Problem with Arc Interdiction as a QUBO and demonstrate solutions on D-Wave quantum annealers, showing near-optimal performance on small instances. These results highlight the potential of quantum methods to reduce complexity, accelerate training, and strengthen resilience, paving the way for safer and more efficient AI-driven vehicles.					
17. Keywords Quantum activation function, weight-constrained neural network, adversarial attack			18. Distribution Statement No restrictions.		
19. Security Classif. (of this report) Unclassified		20. Security Classif. (of this page) Unclassified		21. No. of Pages 32	22. Price N/A



TABLE OF CONTENTS

DISCLAIMER	2
CONTACTS	2
ACKNOWLEDGMENT.....	3
EXECUTIVE SUMMARY	7
CHAPTER 1	9
Introduction.....	9
CHAPTER 2	12
Literature Review.....	12
2.1 The development of quantum machine learning techniques.....	12
2.2 The development of quantum-inspired techniques.....	12
2.3 The resilience of quantum machine learning models.....	13
CHAPTER 3	14
Methods.....	14
3.1 Hybrid quantum-classical convolutional neural network.....	14
3.2 Hybrid quantum-classical neural network.....	14
3.3 Weight-constrained neural network.....	14
3.4 Quantum Non-local Neural Network for Image Classification	14
3.5 Quantum Reinforcement Learning	16
3.6 Arc Interdiction Vehicle Routing Problem using Quantum Annealing.....	17
CHAPTER 4	21
Results.....	21
4.1 Quantum-inspired activation functions.....	21
4.2 Quantum-inspired weight-constrained neural network.....	22
4.3 Quantum Non-local Neural Network for Image Classification	26
4.4 Quantum Reinforcement Learning	27
4.5 Arc Interdiction Vehicle Routing Problem using Quantum Annealing.....	27
CHAPTER 5	29
Conclusions.....	29
REFERENCES	31



List of Tables

Table 1. A summary of the training steps of CNNs using the Tanh activation function and quantum-inspired activation function. 21
Table 2. Summaries of the training results of CNN-QNL-Net model on MNIST and CIFAR-10 datasets. 26

List of Figures

Figure 1. Quantum neural network. (a) The architecture of a hybrid quantum-classical neural network. (b) The architecture of a hybrid quantum-classical neural network. (c) The architecture of a weight-constrained fully connected neural network. 15
Figure 2. The architecture of the quantum circuit. This quantum circuit includes two R_y rotation layers and one CNOT layer. 22
Figure 3. Quantum circuits. (a) Quantum circuit used in Figure 1(a). (b) The circuit for the convolutional layer. (c) The circuit for the pooling layer. 22
Figure 4. The accuracy (ACC) of the weight-constrained fully connected neural network (FNN) and the weight-constrained convolutional neural network (CNN) for different values of N and r . Here, N and r represent variables in the combination formula $C(N, r)$. Panel (a) and panel (b) plot the results of the MNIST and the FMNIST datasets. Panel (c) and panel (d) plot the results of the CIFAR and the traffic datasets. The gray dashed line denotes the ACC of the standard neural network. 24
Figure 5. The accuracy (ACC) under adversarial attack. The top panel shows the MNIST image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panels (c) and (d) plot the ACC of the hybrid quantum-classical neural network (HNN) and the weight-constrained fully connected neural network (FNN) with a dropout policy. Here, p denotes the dropout probability. The middle panel shows the Fashion MNIST (FMNIST) image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panels (g) and (h) plot the ACC of the HNN and the weight-constrained FNN with a dropout policy. The bottom left panel shows the CIFAR image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panel (k) plots the ACC of the weight-constrained convolutional neural network (CNN) with a dropout policy. The bottom right panel shows the traffic image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panel (n) plots the ACC of the weight-constrained CNN with a dropout policy. 25
Figure 6. A comparison between results from the QUBO approach and the upper bound approach. 28



EXECUTIVE SUMMARY

Machine learning techniques have been widely adopted in self-driving vehicles to recognize objects and detect anomaly signals. However, these machine learning techniques suffer from many issues, including a vast number of variables and low training efficiency. To address these issues, we develop quantum machine learning approaches to enhance the performance of AI systems in vehicles. For example, we developed a quantum-inspired weight-constrained neural network to reduce the complexity of AI models. Our weight-constrained neural network requires substantially fewer variables than standard classical neural networks, thereby significantly reducing the energy cost for both training and inference, which is an essential factor for image classification tasks in smart vehicles. Moreover, we develop a novel dropout-based approach to enhance the resilience of AI models in vehicles using the weight-constrained neural network. Combined with this novel dropout-based defense, the model achieves robustness against strong adversarial attacks, with an average accuracy reduction of about 20%. Furthermore, when the adversarial attack is removed, the accuracy remains largely unchanged. These results suggest that our weight-constrained neural network with dropout not only reduces complexity but also enhances resilience, both of which are essential for safe driving.

In addition to the quantum-inspired weight-constrained neural network, we explore a hybrid quantum-classical convolutional neural network utilizing angle encoding. Our observations indicate that the hybrid model trains faster than its classical counterpart, which we attribute to the influence of the quantum activation function. Motivated by this insight, we develop quantum-inspired activation functions. Classical neural networks incorporating these functions consistently require fewer training steps compared to those using standard Tanh activation functions. This finding is particularly significant, as quantum-inspired activation functions can be directly implemented in classical neural networks to enhance efficiency. Consequently, image recognition systems in autonomous vehicles can be trained with substantially fewer computational resources. To further enhance pattern recognition in computer vision, we developed the Quantum Non-Local Neural Network (QNL-Net). Traditional non-local operations in computer vision, which compute pairwise relationships across all elements in a feature set, face computational challenges due to their quadratic complexity in both time and memory. QNL-Net overcomes these limitations by integrating classical dimensionality reduction techniques with quantum circuits, enabling more efficient computations in a quantum-enhanced feature space. Performance analysis on the MNIST and CIFAR-10 datasets demonstrates that QNL-Net achieves state-of-the-art accuracy in image classification among quantum classifiers while utilizing fewer qubits. This result confirms the effectiveness of integrating quantum circuits with classical neural networks, establishing a new benchmark in quantum machine learning and underscoring the potential for scalable, efficient quantum-enhanced models in real-world applications.

Additionally, we developed quantum reinforcement learning (QRL) to improve sample efficiency in sequential decision-making. This work introduces the Quantum Natural Policy Gradient (QNPG) algorithm, which replaces classical random sampling with deterministic quantum gradient estimation. By incorporating quantum oracles for Markov Decision Process (MDP) queries, QNPG achieves a sample complexity of $O(\epsilon^{-1.5})$, surpassing the classical lower bound of $O(\epsilon^{-2})$. The key to this improvement is the quantum evaluation oracle, which leverages quantum parallelism, amplitude amplification, and entanglement. Unlike classical reinforcement learning,



quantum methods introduce challenges due to the deterministic nature of operations, necessitating novel bias-variance trade-offs. This work presents the first quantum speedup for parameterized model-free infinite-horizon MDPs, providing theoretical guarantees on global convergence. The findings highlight the potential of quantum computing to enhance the efficiency of reinforcement learning for complex decision-making tasks, such as autonomous vehicle navigation.

NP-hard optimization problems are difficult for classical computers to solve efficiently, whereas quantum annealing offers a promising alternative. In this project, we used the Arc Interdiction Vehicle Routing Problem (VRP) as a case study of an NP-hard problem, focusing on minimizing traversal costs for a fleet of vehicles. The Arc Interdiction variant introduces an adversary who increases the cost of selected routes. We reformulate this computationally intractable problem as a quadratic unconstrained binary optimization (QUBO) problem to leverage the capabilities of quantum computing. This bilevel optimization consists of an upper-level adversarial interdiction and a lower-level vehicle routing problem. The QUBO formulation enables solutions via quantum annealing, as demonstrated on D-Wave's Quantum Processing Units. Numerical results indicate that while small instances yield near-optimal solutions, larger instances exhibit performance degradation, highlighting both the potential and the current limitations of quantum annealing. Hybrid solvers enhance efficiency, yet scalability remains a challenge, necessitating further advancements to improve the solution quality of the large-scale optimization problems.



CHAPTER 1

Introduction

Cybersecurity in vehicles is essential to safe driving. Our modern vehicles utilize various electronic control units (ECUs) to control different subsystems, including headlights, climate control, engine control, and the smart key that unlocks and starts the vehicle. ECUs are connected through in-vehicle networks, such as the Controller Area Network (CAN) buses. Secure communication between different ECUs is one of the guarantees of safe driving.

CAN in our current vehicles is vulnerable to high risks of cyberattacks due to a lack of efficient and built-in security features, such as authentication. Furthermore, it has been found that attacks on vehicles can be performed via the On-Board Diagnostics II (OBD-II) port, the infotainment system, Bluetooth, RFID car keys, and wireless communication channels. With numerous attack surfaces, modern vehicles are vulnerable to various dangers. Unlike malicious attacks on laptops, malicious attacks on vehicles could result in catastrophic consequences, including the loss of human life. One solution to enhance security is to adopt intrusion detection and mitigation techniques. Intrusion detection systems (IDSs) are used to identify intrusions in computer network systems. However, traditional network security techniques are often not applicable to vehicular networks due to the heterogeneous nature of nodes, varying speeds, and intermittent connections. Therefore, an effective and efficient IDS that can operate in in-vehicle networks is necessary. Currently, the deep learning method is one of the most powerful techniques to detect anomalous signals to ensure safe driving.

In addition to detecting anomalous signals, machine learning techniques have also been widely adopted in self-driving vehicles to identify objects, such as traffic signals, pedestrians, and other vehicles. These models enable real-time decision-making, allowing autonomous systems to navigate safely through complex environments. In our semi-unmanned driving system, a machine learning-based image identification system is employed to continuously monitor nearby vehicles and issue alerts. This system enhances driving safety by reducing human error and providing proactive collision avoidance mechanisms. Moreover, by integrating sensor fusion techniques that combine visual data with LiDAR and radar inputs, the system can achieve higher accuracy and robustness in various weather and lighting conditions. Given the critical role of machine learning in vehicular applications, ensuring both the efficiency and security of these models is paramount. Computational efficiency is essential to ensure real-time performance, thereby minimizing latency in decision-making processes. At the same time, security concerns, such as adversarial attacks and model vulnerabilities, must be addressed to prevent potential malfunctions or malicious exploits. Quantum computing presents new opportunities to enhance the efficiency of AI models and improve their resilience against adversarial attacks. Recent advancements in quantum machine learning (QML) have demonstrated advantages such as higher prediction accuracy and reduced computational costs compared to classical machine learning methods. These benefits make QML a promising candidate for applications in autonomous driving. Motivated by these advantages, it is essential to explore how quantum techniques can be leveraged to improve image identification systems in vehicles. However, due to the current limitations of integrating quantum computers into vehicular systems, we must develop quantum-inspired classical approaches. These approaches aim to capture key quantum advantages, such as enhanced feature selection and computational



efficiency, while remaining feasible for real-world implementation in vehicles. Additionally, quantum machine learning models select features differently than their classical counterparts, potentially providing greater robustness against adversarial attacks designed for classical systems. Investigating the resilience of QML models under adversarial conditions and developing new defense mechanisms based on quantum principles could open new pathways for ensuring the safety of autonomous driving systems. By integrating quantum-inspired methods into vehicular AI, we can enhance both performance and security, paving the way for safer and more efficient transportation.

With these motivations in mind, this report introduces two quantum-inspired techniques to enhance AI models. The first technique involves developing a quantum-inspired activation function designed to reduce the number of training steps required for any classical neural network (NN) models. In a hybrid quantum-classical NN, we observed that the hybrid model trains faster than a purely classical model. By analyzing the underlying mathematical structure of the hybrid model, we identified that this improved performance stems from the quantum-inspired activation function. Notably, this activation function can be expressed using classical mathematical formulations, making it easily implementable in conventional neural networks to enhance training efficiency. The second technique focuses on a quantum-inspired weight-constrained NN, which could contain substantially fewer variables than a standard NN. Our research initially explored a quantum NN with amplitude encoding, which inherently requires fewer parameters than classical NNs. By delving into the underlying mathematical principles, we identified the root of this advantage and leveraged it to develop a weight-constrained NN. Furthermore, we introduced a novel defense mechanism based on this weight-constrained structure to enhance the model's resilience against adversarial attacks. Our work not only reduces the computational cost of classical NN models in vehicular systems but also improves their robustness, paving the way for more efficient and secure AI-driven autonomous driving technologies.

In computer vision, non-local neural networks are powerful tools for capturing long-range dependencies, but their quadratic complexity in time and memory limits scalability. To address this, we present a hybrid quantum-classical framework, the Quantum Non-Local Neural Network (QNL-Net), which leverages quantum parallelism and entanglement to efficiently process high-dimensional feature representations. By embedding data into quantum-enhanced feature spaces, QNL-Net enables more expressive modeling of spatial dependencies while maintaining scalability for large-scale image classification tasks.

Reinforcement learning (RL) is another domain where quantum computing can provide transformative advantages, particularly in environments requiring fast adaptation and robust decision-making. Classical policy gradient methods often face high sample complexity and slow convergence, constraining their effectiveness in real-world applications such as traffic optimization and autonomous control. To mitigate these challenges, we introduce the Quantum Natural Policy Gradient (QNPG) algorithm, a hybrid quantum-classical RL framework that integrates quantum oracles for trajectory sampling and gradient estimation. This design reduces variance in policy updates, enhances convergence rates, and highlights the potential of quantum methods to improve the efficiency of adaptive learning systems.

Beyond learning and perception tasks, quantum computing also offers significant promise in



combinatorial optimization, especially in transportation and logistics. The Vehicle Routing Problem (VRP), when combined with adversarial interdiction, becomes a highly complex bilevel optimization challenge that classical solvers struggle to handle efficiently. We present a novel reformulation of the Arc Interdiction VRP as a Quadratic Unconstrained Binary Optimization (QUBO) problem, solvable on quantum annealers such as D-Wave's processors. This approach enables the exploration of large solution spaces in parallel, providing new insights into the feasibility of quantum annealing for resilient routing in adversarial settings and demonstrating its potential impact on critical infrastructure protection.

This report is organized into five chapters. Chapter 2 presents a comprehensive literature review on quantum machine learning, outlining key developments and existing research. Chapter 3 details the methods and experimental setups used in this study. Chapter 4 discusses the results of the proposed quantum-inspired techniques, analyzing their performance and implications. Finally, Chapter 5 concludes the report with key observations, implications for future research, and recommendations.



CHAPTER 2

Literature Review

This literature review contains the development of quantum machine learning and quantum-inspired techniques. The resilience of quantum AI models will also be presented.

2.1 The development of quantum machine learning techniques

Many supervised quantum machine learning methods have been developed. For example, Carsten Blank *et al.* developed a quantum circuit with tailored quantum kernel for classification (Blank *et al.*, 2020). Christa Zoufal *et al.* developed quantum generative adversarial networks for learning and loading random distributions (Zoufal, Lucchi and Woerner, 2019). Patrick Rebentrost *et al.* developed a quantum support vector machine for big data classification (Rebentrost, 2014). El Amine Cherrat *et al.* developed quantum vision transformer models (Cherrat *et al.*, 2024). All these models use quantum circuits to extract essential features. Unlike the kernel method used in classical machine learning, the quantum model encodes data in the high-dimensional Hilbert space and utilizes a series of quantum gates to construct nonlinear quantum kernel functions. Broadly, these quantum kernel methods fall into two categories: explicit and implicit quantum models. Explicit models utilize measurements of a quantum state controlled by encoded data and parameterized gates. The recently proposed data reuploading model also belongs to this category (Pérez-Salinas *et al.*, 2020; Moreira *et al.*, 2023). Implicit models weigh the inner products of quantum states, which are controlled by the encoded input data. A representative example of the implicit model is the quantum support vector machine (Rebentrost, 2014; Kusumoto *et al.*, 2021; Jäger and Krems, 2023). Compared to classical machine learning, the superior performance of quantum methods has been witnessed in some case studies. For example, Fan *et al.* implemented a hybrid quantum-classical convolutional neural network to classify earth observation data and found that quantum models can accelerate the convolutional operation in comparison with their classical counterparts (Fan *et al.*, 2024). Naim Ajlouni *et al.* applied the hybrid quantum CNN to diagnose medical images and demonstrated that the hybrid method has higher accuracy and faster speed compared to the classical method (Ajlouni *et al.*, 2023). Although quantum advantages have been observed in these applications, the underlying origins of these advantages remain unclear.

2.2 The development of quantum-inspired techniques

Understanding the origins of superior performance in quantum machine learning is crucial for optimizing quantum models and potentially enriching classical algorithms. Notably, quantum-inspired classical algorithms have been leveraged to accelerate linear algebra computations, demonstrating the advantages of translating quantum insights into classical contexts (Chia, Lin and Wang, 2018; Gilyén, Lloyd and Tang, 2018; Arrazola *et al.*, 2020; Tang, 2021). Beyond these advancements, quantum-inspired machine learning methods have also been proposed. For example, Wooseop Hwang introduced a quantum-inspired binary classifier that exhibits excellent performance (Hwang *et al.*, 2024). Chen Ding *et al.* developed a quantum-inspired support vector machine designed for processing large-scale data (Ding, Bao and Huang, 2022). Furthermore, inspired by Kerenidis and Prakash's quantum recommendation system, Ewin Tang devised a quantum-inspired recommendation algorithm that generates recommendations exponentially



faster than standard classical systems (Tang, 2019). Additionally, Trung Q. Duong et al. proposed quantum-inspired machine learning techniques for 6G networks, enhancing network security (Duong *et al.*, 2022).

These developments underscore the potential of quantum-inspired approaches in bridging the gap between quantum and classical machine learning, paving the way for more efficient computational techniques across various domains.

2.3 The resilience of quantum machine learning models

The advantage of quantum protocols lies in their high-dimensional Hilbert space, which enables efficient processing of large datasets for classification tasks. However, counterintuitive geometrical properties that emerge in high dimensions make classification problems susceptible to adversarial attacks. Nana Liu and Peter Wittek demonstrated that the perturbation required for an adversary to induce a misclassification scales inversely with dimensionality (Liu, 2020). This finding highlights a tradeoff between the security of classification algorithms against adversarial attacks and the computational advantages of quantum methods. Moreover, Weiyuan Gong and Dong-Ling Deng discovered the existence of universal adversarial examples capable of fooling multiple quantum classifiers. They proved that for a set of k classifiers, each receiving input data of n qubits, an $O(\ln[k]/2^n)$ increase in perturbation strength suffices to achieve a moderate universal adversarial risk (Gong and Deng, 2022). Given these insights, developing strategies to enhance the resilience of quantum machine learning methods is crucial. Recently, several approaches have been proposed. For instance, Yuxuan Du et al. suggested leveraging quantum noise to protect quantum classifiers from adversaries (Du, 2021). Additionally, Sirui Lu applied quantum adversarial learning to improve robustness (Lu, 2020). Similar to classical neural networks, the vulnerability of quantum machine learning to adversarial attacks remains an open challenge, necessitating further research and development.



CHAPTER 3

Methods

3.1 Hybrid quantum-classical convolutional neural network

Figure 1 (a) shows the architecture of a hybrid quantum-classical neural network. We take the handwritten image as an example to explain this hybrid neural network. This hybrid network consists of four layers, with 16 neurons in the first hidden layer and 64 neurons in the second hidden layer. The input data is encoded into a quantum circuit using amplitude encoding, and the output of the quantum circuit is stored in the first hidden layer. The detailed structure of the quantum circuit is presented in the results section.

3.2 Hybrid quantum-classical neural network

Figure 1 (b) shows the architecture of a hybrid quantum-classical convolutional neural network. This hybrid network consists of one convolutional layer, one flattened layer, and two fully connected layers. Unlike the classical convolutional neural network, the convolutional operation in the hybrid model is performed through a quantum circuit. For the convolutional operation, a small patch of an image is taken and encoded into the quantum circuit via the angle encoding. Therefore, the feature selection is made through quantum circuits rather than the standard classical activation function, such as ReLU and Tanh.

3.3 Weight-constrained neural network

Figure 1(c) shows the architecture of a weight-constrained neural network. This network consists of three fully connected layers, the same as a standard fully connected neural network. Unlike the standard fully connected neural network, the weights in the first fully connected layer are not independent but correlated. Here, we build the weight matrix using a combinational approach. For example, to generate K different weights, we first define N different variables, denoted as θ . Then, a combination process is performed where r variables are chosen from these N variables, resulting in $\mathcal{C}(N, r)$ possible combinations. Labeling the k -th combination as $\theta^{(k)}$, the k -th weight w_k is defined as:

$$w_k = \prod_{i \in \text{even}, i < r} \cos(\theta_i^{(k)}) \prod_{j \in \text{odd}, j < r} \sin(\theta_j^{(k)}) \quad (1)$$

This combination strategy enables constructing thousands of weights using only a few dozen variables. For example, if we set $N = 20$ and $r = 5$, we can build $\mathcal{C}(20, 5) = 15,504$ weights using only 20 variables.

3.4 Quantum Non-local Neural Network for Image Classification

Non-local operations play a crucial role in computer vision, enabling the capture of long-range dependencies through weighted sums of features across the input, surpassing the constraints of traditional convolution operations that focus solely on local neighborhoods. Non-local operations

typically require computing pairwise relationships between all elements in a feature set, leading to quadratic complexity in terms of time and memory. Due to the high computational and memory demands, scaling non-local neural networks to large-scale problems can be challenging. In this work, we introduce a hybrid quantum-classical scalable non-local neural network, the Quantum Non-Local Neural Network (QNL-Net), to enhance pattern recognition. The proposed QNL-Net relies on inherent quantum parallelism to allow the simultaneous processing of a large number of input features, enabling more efficient computations in quantum-enhanced feature space and involving pairwise relationships through quantum entanglement.

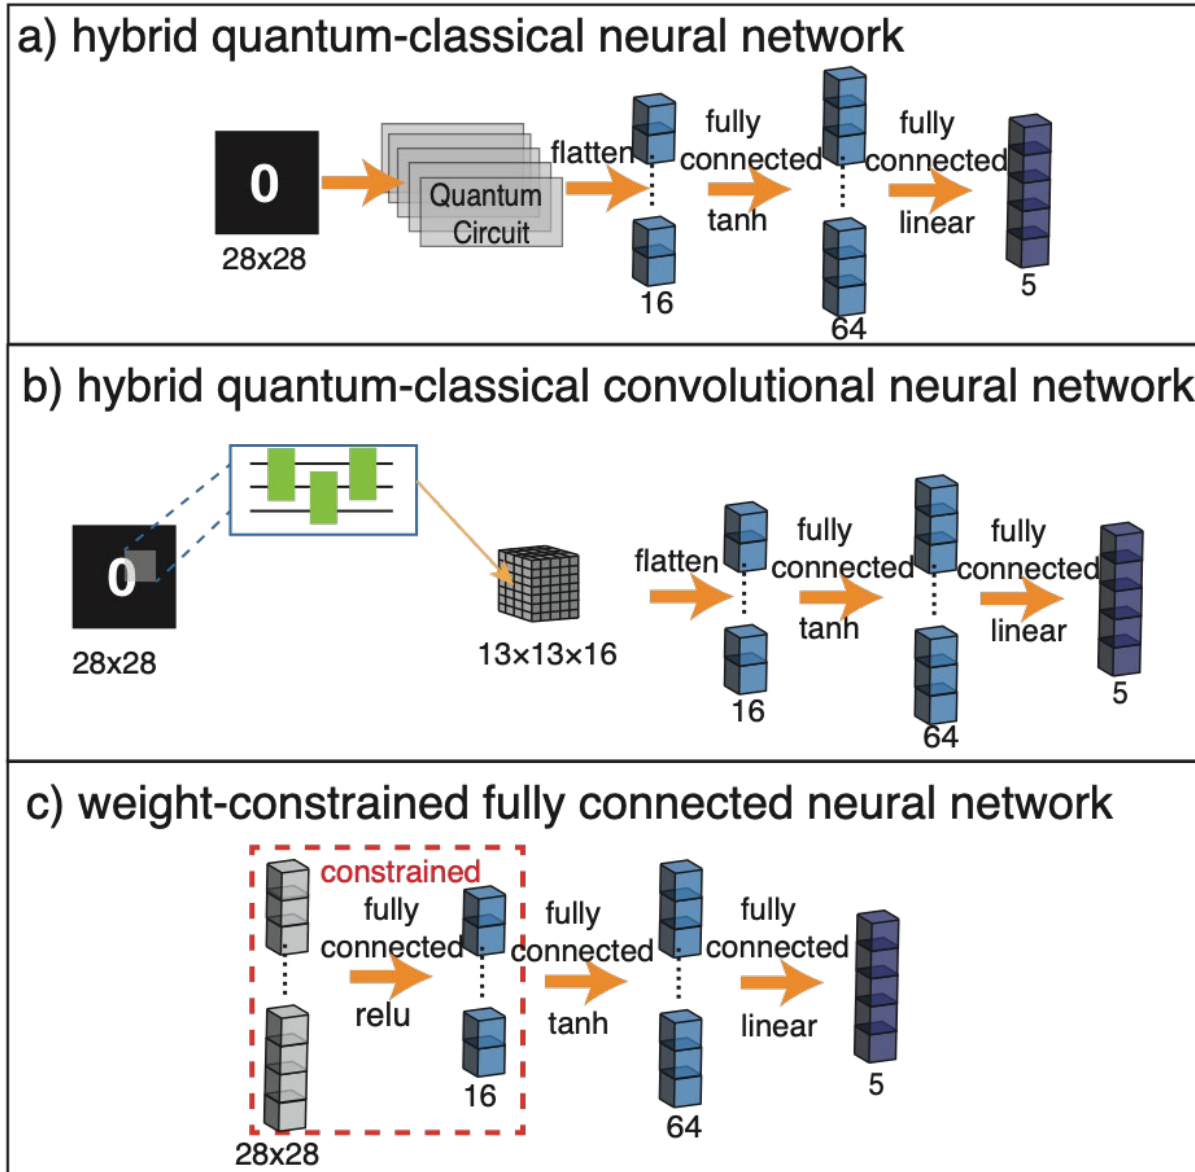


Figure 1. Quantum neural network. (a) The architecture of a hybrid quantum-classical neural network. (b) The architecture of a hybrid quantum-classical neural network. (c) The architecture of a weight-constrained fully connected neural network.



The Quantum Non-local Neural Network (QNL-Net) architecture is designed as a hybrid quantum-classical framework for image classification, integrating classical dimensionality reduction techniques with quantum circuits to efficiently capture long-range dependencies in data. The QNL-Net consists of three primary components: an encoder, a Variational Quantum Circuit (VQC), and a measurement process. Classical data is first encoded into quantum states through the encoder, which applies Hadamard gates and parameterized rotation gates to create superposition and entanglement among qubits. This encoding step transforms the input data into a high-dimensional quantum feature space, enabling more efficient representation and analysis of complex patterns.

Following the encoding phase, the VQC processes the quantum-enhanced features using parameterized quantum gates that can be optimized via classical methods. The VQC includes three distinct ansatzes—each employing five parameterized rotation gates (R_x , R_y , R_z) and three controlled-NOT (CX) gates arranged in unique configurations (cyclic, reverse linear chain, and mixed pattern). These configurations ensure comprehensive entanglement among all qubits, capturing non-local dependencies within the data. The rotation gates allow precise manipulation of quantum states by rotating them around various axes on the Bloch sphere, facilitating spatial transformations that are crucial for emulating classical non-local operations. Each layer of the VQC contains five parameters, and multiple repetitions of these layers enhance the model's expressiveness.

Finally, the measurement component evaluates the quantum state at qubit 0 in the Pauli-Z basis, while the remaining qubits are measured using the Identity operation. This measurement yields an expectation value that corresponds to the output of the quantum circuit. To further refine this output, a fully connected classical layer with a single learnable parameter is added post-measurement. This classical computation fine-tunes the quantum output, optimizing it for binary classification tasks.

3.5 Quantum Reinforcement Learning

Reinforcement learning (RL) is an emerging paradigm for autonomous decision-making, particularly in environments where real-time adaptation to dynamic and uncertain conditions is required. Classical RL methods, such as Natural Policy Gradient (NPG), have shown promise in applications like autonomous vehicle control and traffic routing optimization. However, these methods often suffer from high sample complexity and slow convergence due to reliance on stochastic gradient estimation.

Quantum computing offers new avenues for improving the efficiency of RL algorithms by leveraging quantum parallelism and amplitude amplification. We introduce a novel quantum-enhanced reinforcement learning framework, Quantum Natural Policy Gradient (QNPG), that significantly improves the sample complexity over classical NPG while maintaining global convergence guarantees.

The QNPG algorithm builds upon classical policy gradient methods by integrating quantum oracles for Markov Decision Process (MDP) queries. These oracles allow for deterministic gradient estimation through quantum superpositioned trajectories, reducing the variance inherent in classical sampling techniques.



Key components of the methodology include:

1. Quantum Transition Oracle (U_p): This oracle, at step t , returns the superposition over $s' \in S$ according to $P(s'|s_t, a_t)$, the probability distribution of the next state given the current state $|s_t\rangle$ and action $|a_t\rangle$ is defined as:

$$U_P : |s_t\rangle|a_t\rangle|0\rangle \rightarrow |s_t\rangle|a_t\rangle \sum_{s' \in S} \sqrt{P(s'|s_t, a_t)}|s'\rangle \quad (2)$$

2. Quantum Initial State Oracle (U_0): Provides coherent access to the initial state distribution.
3. We also assume the ability to construct a unitary Π that coherently implements a policy π_θ : Let $\pi_\theta : S \times A \rightarrow [0, 1]$ be a reinforcement learning policy acting in a state-action space $S \times A$ and parametrized by a vector $\theta \in \mathbb{R}^d$ (that can be encoded with finite precision as $|\theta\rangle$). We say that the policy is quantum-evaluable if we can construct a unitary satisfying:

$$\Pi : |\theta\rangle |s\rangle |0\rangle \mapsto |\theta\rangle |s\rangle \sum_{a \in A} \sqrt{\pi_\theta(a|s)} |a\rangle \quad (3)$$

With access to quantum oracles for both the environment and the policy, it becomes possible to design subroutines capable of generating superpositions of trajectories with fixed length within the environment and calculating the returns associated with these trajectories.

The key construct is the gradient estimator using the quantum oracles above. Our quantum NPG algorithm can be segregated into a classical outer loop and a novel quantum inner loop. In its outer loop, the policy parameters are updated K number of times following: $\theta_{k+1} = \theta_k + \eta \omega_k$, where $\{\omega_k\}$ denotes the natural policy gradient estimates. For the inner loop, the estimates $\{\omega_k\}$ are calculated by iterating the quantum stochastic gradient descent (SGD) algorithm H number of times. Each quantum SGD iteration consists of a quantum mini-batch approach, where the stochastic gradients are obtained via quantum mean estimation using quantum samples.

We utilize a Quantum Variance Reduction method, which leverages Quantum Mean Estimation as a subroutine. For each policy parameter θ_k and each inner-loop index h , the variance reduction is applied to reduce the variance of the gradient estimators while maintaining unbiasedness.

3.6 Arc Interdiction Vehicle Routing Problem using Quantum Annealing

The Vehicle Routing Problem (VRP) is a well-established NP-hard optimization problem that seeks to determine an efficient routing allocation for a fleet of vehicles over a graph, minimizing the total cost of traversal. Arc Interdiction adds another layer of complexity by introducing an adversary capable of disrupting some routes along the graph with the aim of increasing the total cost. Given the complexity of VRP, solving large instances efficiently remains computationally intractable for classical algorithms. Quantum computers, with their ability to explore vast solution spaces in parallel, offer a promising avenue for accelerating optimization for such problems. This



work formulates the bilevel optimization corresponding to Arc Interdiction VRP as a Quadratic Binary Unconstrained Optimization (QUBO), providing a novel approach to solving the problem using quantum hardware. Numerical experiments conducted using D-Wave's Quantum Processing Units demonstrate the feasibility of the proposed formulation and provide insights into the performance of quantum annealing in handling interdiction scenarios.

We define a directed graph $G = (V, A)$ where $V = \{0\} \cup C$ defines the set of nodes. Here, $\{0\}$ denotes the depot node, and C denotes the set of customer nodes. $A = \{(i, j) \mid i, j \in V\}$ defines the set of arcs between the nodes, with each arc having a non-negative cost of traversal $c_{ij} \geq 0$. We assume a set of K vehicles, each starting and ending at the depot node. The goal is to determine a set of feasible routes for vehicles such that each customer node in C is visited exactly once while minimizing the total cost.

Furthermore, the scenario of arc interdiction introduces an adversary that attempts to increase the cost of a subset of arcs. The interdicted arcs incur an additional cost of $d_{ij} \geq 0$. The adversary is constrained by a total interdiction budget B , which limits the number of arcs that can be interdicted. It is further assumed that the resources required to interdict any given arc $(i, j) \in A$ is 1 unit.

The problem can thus be viewed as a bi-level optimization, where the lower-level problem corresponds to finding the vehicle routes that minimize the total cost for a given interdiction configuration, and the upper-level problem involves identifying the set of interdicted arcs that maximizes the adversary's impact on the routing cost. The variables and problem parameters for the proposed formulation are defined as follows:

- x_{ij} : Binary variable indicating if arc (i, j) was interdicted.
- y_{ijk} : Binary variable indicating if arc (i, j) has been traversed by vehicle $k \in K$.
- c_{ij} : Cost of traversing arc (i, j) .
- d_{ij} : Penalty imposed on an interdicted arc (i, j) .
- B : Total Interdiction Budget.

The problem formulation for the Vehicle Routing Problem with Interdiction is as follows:

Upper-level formulation

$$\max_{x \in X} V(x)$$

subject to

$$\sum_{(i,j) \in A} x_{ij} \leq B$$

(4)



Lower-level formulation

$$V(x) = \min_{y \in Y} \sum_{i \in V} \sum_{j \in V} \sum_{k \in K} (c_{ij} + d_{ij}x_{ij})y_{ijk}$$

subject to

$$\begin{aligned} \sum_{k \in K} \sum_{i \in V} y_{ijk} &= 1 \quad \forall j \in V/\{0\} \\ \sum_{i \in V} y_{ijk} &= \sum_{i \in V} y_{jik} \quad \forall j \in V/\{0\}, k \in K \\ \sum_{j \in V} y_{0jk} &= 1 \quad \forall k \in K \\ \sum_{i \in V} y_{i0k} &= 1 \quad \forall k \in K \\ x_{ij} &\in \{0, 1\} \quad \forall i, j \in V \\ y_{ij} &\in \{0, 1\} \quad \forall i, j \in V \end{aligned} \tag{5}$$

The objective functions correspond to the upper-level and lower-level optimization problems, respectively. A penalty of d_{ij} is added to the cost of an interdicted arc (i, j) . The upper-level problem is subject to a constraint that limits the number of interdicted arcs based on the resource budget. The formulation ensures that each customer is only visited once through the first equation. The flow conservation constraint ensures that the number of vehicles entering a node is equal to the number of vehicles leaving the node. The last two equations enforce that each vehicle starts and finishes at the depot node.

To solve the problem using quantum annealers, we reformulate the bi-level optimization as a single-level QUBO formulation. We begin by transforming the bi-level problem into a single-level integer programming problem. We can fix the variables x of the upper-level problem, relax the integral constraints, take the dual of the lower-level problem, and then release the fixed variables x arriving at the following single-level maximization problem:

$$\max_{x, \pi, \lambda, \mu, \nu} \sum_{j \in V/\{0\}} \pi_j + \sum_{k \in K} (\mu_k + \nu_k)$$

subject to

$$\sum_{(i,j) \in A} x_{ij} \leq B$$

$$\pi_j + \lambda_{jk} - \lambda_{ik} + \delta_{i0}\mu_k + \delta_{j0}\nu_k \leq c_{ij} + d_{ij}x_{ij} \quad \forall i, j \in V; k \in K$$

$$\delta_{i0} = \begin{cases} 1, & \text{if } i = 0 \\ 0, & \text{otherwise} \end{cases}$$

$$y_{ij} \in \{0, 1\} \quad \forall i, j \in V \tag{6}$$

where π , λ , μ , and ν are the dual variables corresponding to the lower-level constraints.



The next step would be to transform this single-level integer programming problem into a QUBO formulation. To achieve this, we need to express the objective function and constraints in terms of binary variables, which is already the case for the interdiction decision variables x_{ij} . We derive the respective bounds for the dual variables, and they can be further decomposed into several binary variables through bit-wise mapping. Furthermore, we also convert the constraints to penalty terms in the objective functions, choosing sufficiently large penalty terms P1 and P2. The final QUBO formulation is given by:

$$\begin{aligned} \max_{x, \pi, \lambda, \mu, \nu, r, s} \quad & \sum_{j \in V \setminus \{0\}} \pi_j + \sum_{k \in K} (\mu_k + \nu_k) \\ & + P_1 \sum_{(i,j) \in A} \sum_{k \in K} (c_{ij} + d_{ij} x_{ij} \\ & - (r_{ijk} + \pi_j + \lambda_{jk} - \lambda_{ik} + \delta_{i0} \mu_k + \delta_{j0} \nu_k))^2 \\ & + P_2 \sum_{(i,j) \in A} (B - (s_{ij} + x_{ij}))^2 \end{aligned}$$

subject to

$$\begin{aligned} \pi_j, \lambda_{jk}, \mu_k, \nu_k, r_{ijk} \in [0, \max_{(i,j) \in A} \{c_{ij} + d_{ij}\}]; \\ s_{ij} \in [0, B] \end{aligned} \tag{7}$$

We introduce slack variables r and s to handle the inequality constraints. We use binary encoding to transform all integer variables z (where $z \in \{\pi_j, \lambda_{jk}, \mu_k, \nu_k, r_{ijk}, s_{ij} \} \forall i, j, k$) as follows:

$$z = \sum_{b=0}^{\lceil \log_2(U_z + 1) \rceil - 1} 2^b z_b \tag{8}$$

where, U_z is the upper bound on the integer variable z and $(z_b \in \{0, 1\})$ are binary variables corresponding to each bit b .

This methodology outlines the transformation process of the Arc Interdiction Vehicle Routing Problem into a Quadratic Unconstrained Binary Optimization format suitable for solving using quantum annealing techniques.



CHAPTER 4 Results

In this chapter, we first introduce two quantum-inspired techniques: the quantum-inspired activation function and quantum-inspired weight-constrained neural networks. A classical neural network enhanced with these techniques demonstrates superior performance compared to a standard neural network. Additionally, we develop a quantum nonlocal neural network to improve image identification in vehicles. We also explore quantum reinforcement learning to accelerate decision-making. Finally, we apply D-Wave quantum computers to solve arc interdiction vehicle routing problems.

4.1 Quantum-inspired activation functions

In this project, we develop a hybrid quantum-classical convolutional neural network for image classification. The architecture of this hybrid system is shown in Figure 1(b). Figure 2 plots the architecture of the quantum circuit that we adopted. This quantum circuit has two R_y rotation layers and one CNOT layer. The classical data is encoded into this quantum circuit by setting the rotation angle in the first R_y layer as $\theta_i = x_i\pi$, where x_i denotes the classical data. We applied this hybrid model to identify images, such as the MNIST hand-written images, Fashion images, and alphabet letters.

We observed that the hybrid model has a faster training speed compared to the classical model. To understand this better performance, we analyze the mathematical expression of this quantum circuit, which is given by,

$$o = \prod_{i \in \text{even}} \cos(\theta_i + \phi_i). \tag{9}$$

We attribute the better performance of the hybrid model to the quantum activation function. Table 1 summarizes the training steps of the hybrid model and a classical model using the Tanh activation function. Additionally, we present results of quantum-inspired activation functions. Compared to the model using the Tanh activation function, the model using the quantum-related activation functions needs a lower number of training steps.

Table 1. A summary of the training steps of CNNs using the Tanh activation function and quantum-inspired activation function.

Dataset	Function #1	Function #2	Function #3	Baseline function, Tanh
MNIST	1	7	6	18
FMNIST	1	14	3	35
Letter	2	5	3	20
Function #1: $\prod_{i \in \text{even}} \cos(\theta_i + \phi_i)$; Function #2: $\prod_{i \in \text{even}} [\cos\theta_i + \cos\phi_i] + \prod_{i \in \text{odd}} [\sin\theta_i + \sin\phi_i]$; and Function #3: $\prod_{i \in \text{even}} [\cos\theta_i + \cos\phi_i]$				




Figure 2. The architecture of the quantum circuit. This quantum circuit includes two R_y rotation layers and one CNOT layer.

This result demonstrates that the hybrid quantum-classical neural network can reduce the training cost for image identification, therefore reducing the cost in developing image identification system in vehicles.

4.2 Quantum-inspired weight-constrained neural network

4.2.1 Hybrid quantum-classical neural network

In this project, we investigate the performance of the hybrid quantum-classical neural network using amplitude encoding. The architecture of the hybrid quantum-classical neural network is shown in Figure 1(a). Figure 3 illustrates the quantum circuit adopted in the hybrid model. The input data is encoded through the input wave function.

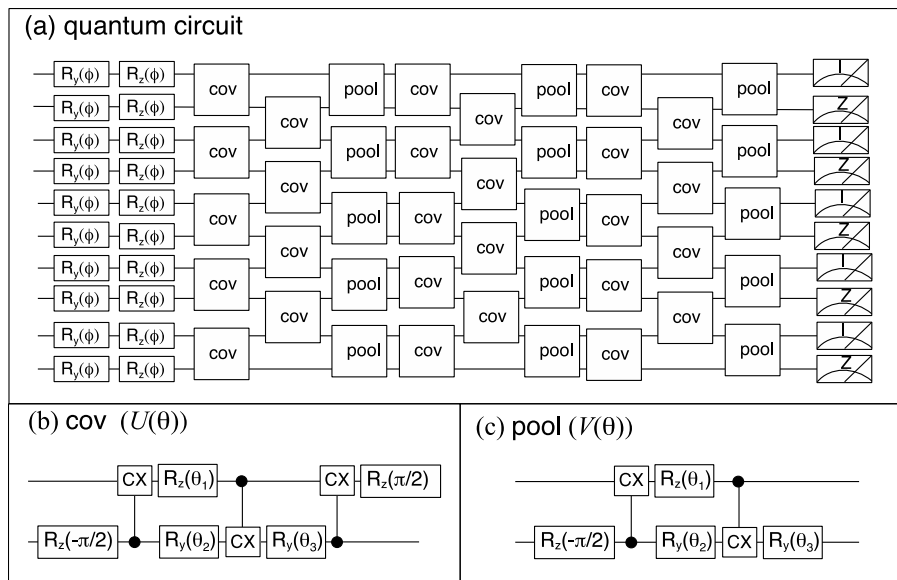


Figure 3. Quantum circuits. (a) Quantum circuit used in Figure 1(a). (b) The circuit for the convolutional layer. (c) The circuit for the pooling layer.



We observe that this hybrid model achieves a similar accuracy to the classical fully connected neural network. However, the hybrid model requires substantially fewer variables. We investigate the underlying mathematics of the hybrid model and uncover the reason for this reduction in variables. Building on this discovery, we adapt it to the classical neural network and develop a weight-constrained neural network, which requires fewer variables than the standard neural network.

4.2.2 Weigh-constrained neural network

In the weight-constrained neural network, we build the weight matrix using a combinational approach. For example, to generate K different weights, we first define N different variables, denoted as θ . Then, a combination process is performed where r variables are chosen from these N variables, resulting in $C(N, r)$ possible combinations. Labeling the k -th combination as $\theta^{(k)}$, the k -th weight w_k is defined as

$$w_k = \prod_{i \in \text{even}, i < r} \cos(\theta_i^{(k)}) \prod_{j \in \text{odd}, j < r} \sin(\theta_j^{(k)}) \quad (10)$$

This combination strategy enables constructing thousands of weights using only a few dozen variables. For example, if we set $N = 20$ and $r = 5$, we can build $C(20, 5) = 15504$ weights using only 20 variables.

To investigate the performance of the weight-constrained strategy, we perform simulations using different values of N and r . Figure 4 presents the accuracy (ACC) for the MNIST, FMNIST, CIFAR, and traffic sign datasets. In these simulations, the weight-constrained FNN is applied to the MNIST and FMNIST datasets, whereas the weight-constrained CNN is used for the CIFAR and traffic datasets. The gray dashed line represents the ACCs of their respective classical standard counterparts. For each pair of (N, r) , the displayed result reflects the optimal outcomes of 16 independent simulations. For weight-constrained FNN, the ACC generally decreases as r increases, which is attributed to the lower expressibility of the network at large r values. Furthermore, ACC does not show a significant dependence on the value of N .

The performance of the weight-constrained neural network, compared to the standard classical neural network, varies depending on the characteristics of the dataset. For example, the weight-constrained FNN achieves a larger ACC on the MNIST dataset compared to the standard FNN, but it exhibits 1% smaller ACC on the FMNIST dataset. The weight-constrained CNN achieves a comparable or larger ACC on the CIFAR and traffic datasets relative to the standard CNN, while its ACC for the FMNIST test dataset is 1% smaller. This minor variation in ACC is not critical as it can be mitigated by adjusting other hyperparameters, such as the activation function and the number of neurons in the hidden layers. The key takeaway from these simulations is that classical neural networks using the weight-constrained method can be effectively trained to produce results that are close to the results of standard neural networks.

4.2.3 Dropout enhanced resilience

In this study, we use the fast gradient sign method (FGSM) to generate adversarial samples. For

the FGSM method, each pixel is modified by $\epsilon \nabla_x L$, where L is the loss function and ϵ represents the strength of the attack. Therefore, a positive value of ϵ increases the loss function, reducing accuracy. To visualize the adversarial attack, we plot original images and their corresponding adversarially perturbed versions in Figure 5 (a), (b), (e), (f), (i), (j), (l), (m).

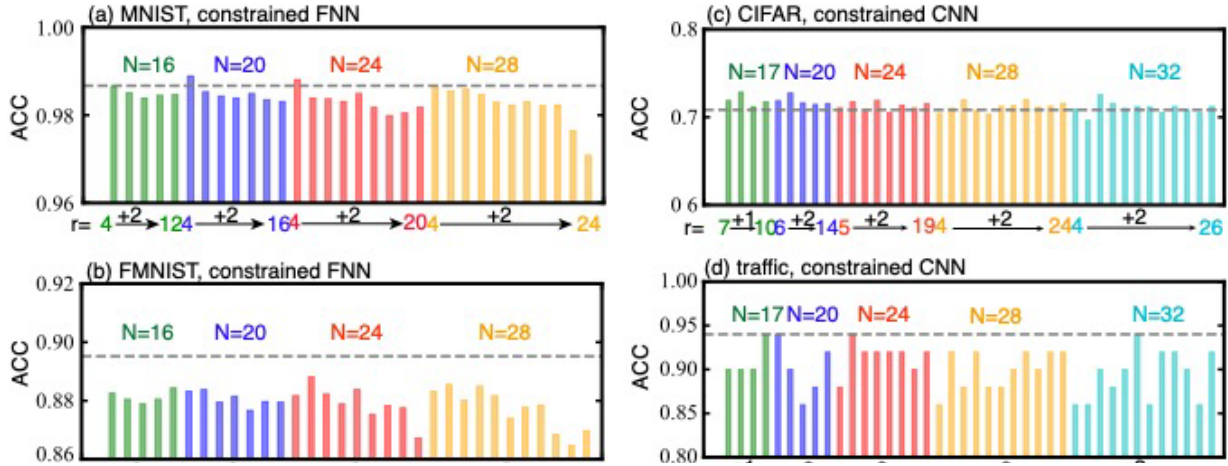


Figure 4. The accuracy (ACC) of the weight-constrained fully connected neural network (FNN) and the weight-constrained convolutional neural network (CNN) for different values of N and r . Here, N and r represent variables in the combination formula $\mathcal{C}(N, r)$. Panel (a) and panel (b) plot the results of the MNIST and the FMNIST datasets. Panel (c) and panel (d) plot the results of the CIFAR and the traffic datasets. The gray dashed line denotes the ACC of the standard neural network.

To enhance the resilience of the quantum machine learning model, we propose randomly dropping R_z gates within the quantum circuit of a trained model. In this approach, dropout is immediately added when the model is trained, denoting that attackers have full access only to the loss function of the model after the dropout has been applied. Consequently, adversarial samples are generated using the trained model with the dropout mechanism in place. Figures 5(c) and 5(g) plot the ACC of the HNN model under adversarial attack for the MNIST and FMNIST test datasets. Here, p denotes the probability of dropping out. Without the dropout policy ($p = 0$), the ACC for both datasets falls rapidly as the attack strength increases ϵ , eventually approaching zero at $\epsilon = 0.4$. This result demonstrates the weak resilience of quantum models under adversarial attacks. However, the application of the dropout strategy significantly improves resilience. For example, with $p = 0.08$, the ACC for the MNIST dataset drops from 0.968 to 0.567, and for the FMNIST dataset, it drops from 0.774 to 0.555 when ϵ increases from 0 to 0.1. It is also observed that the ACC for unattacked samples is linearly suppressed by the dropout probability. For example, increasing the probability of dropping out from 0 to 0.08 reduces the ACC by approximately 2% for the MNIST dataset and by 11.2% for the FMNIST dataset. These results highlight that while the dropout method effectively enhances the adversarial resilience of quantum neural networks, it is crucial to carefully select the dropout probability to strike a balance between maintaining ACC for unattacked samples and improving resilience under attack.



The dropout method can also be applied to enhance the adversarial resilience of classical weight-constrained neural networks. Dropping R_z gate within a quantum circuit corresponds to changing elements of the weight matrix. In a weight-constrained neural network, the weight is constructed from the continuous product of the trigonometric function. We can modify the weight by randomly dropping an angle and replacing the trigonometric function related to that angle with a value of one. Figure 5(d) and Figure 5(h) illustrate the ACC of the weight-constrained FNN for the MNIST and FMNIST test datasets, respectively. It is found that a small nonzero dropout probability ($p = 0.001$) can significantly improve adversarial resilience and make the change of the ACC under attacks tiny. For example, when $p = 0.001$ is adopted, the ACC for the FMNIST dataset is only decreased by 0.07 as the attack strength ϵ increases from 0 to 0.2. Compared to the hybrid quantum-classical model, the weight-constrained FNN requires a smaller value of p to achieve enhanced adversarial resilience. This ensures that the ACC for unattacked samples remains almost unaffected when such a small p is adopted. Furthermore, Figures 5(k) and 5(n) depict the ACC of the weight-constrained CNN, which is trained on the CIFAR and traffic datasets. Similarly, the dropout method with $p = 0.001$ substantially improves adversarial resilience. The oscillatory behavior observed in Figure 5(n) can be attributed to the small sample size of the traffic test dataset (50 samples), which lacks sufficient statistical robustness.



Figure 5. The accuracy (ACC) under adversarial attack. The top panel shows the MNIST image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panels (c) and (d) plot the ACC of the hybrid quantum-classical neural network (HNN) and the weight-constrained fully connected neural network (FNN) with a dropout policy. Here, p denotes the dropout probability. The middle panel shows the Fashion MNIST (FMNIST) image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panels (g) and (h) plot the ACC of the HNN and the weight-constrained FNN with a dropout policy. The bottom left panel shows the CIFAR image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panel (k) plots the ACC of the weight-constrained convolutional neural network (CNN) with a dropout policy. The bottom right panel shows the traffic image and the image under an adversarial attack with an attack intensity $\epsilon = 0.2$. Panel (n) plots the ACC of the weight-constrained CNN with a dropout policy.



We note that the success of our approach on the traffic dataset is critical to ensuring the safety of self-driving vehicles. In these vehicles, CNN-based regression models are widely used for object recognition. However, CNNs have been shown to be highly susceptible to adversarial samples. To address this issue, numerous studies have focused on improving the adversarial resilience of CNNs. Our work introduces a novel method to improve CNN resilience, thus improving the safety and reliability of self-driving vehicles.

4.3 Quantum Non-local Neural Network for Image Classification

We benchmark our proposed QNL-Net with other quantum counterparts for binary classification with the MNIST and CIFAR-10 data sets. The simulation findings showcase our QNL-Net, which achieves cutting-edge accuracy levels in binary image classification among quantum classifiers while utilizing fewer qubits.

Table 2 shows that the CNN-QNL-Net model achieved an average test accuracy of 99.96% on the MNIST dataset for binary classification tasks involving digits 0 and 1 and 93.98% on the CIFAR 10 dataset for binary classification tasks involving digits 2 and 8. This performance is consistent across different ansatz configurations, with Ansatz-0, Ansatz-1, and Ansatz-2 all achieving good test accuracies. The PCA-QNL-Net model also performed well but lagged behind the CNN-QNL-Net, where the difference is significant for CIFAR-10.

Table 2. Summaries of the training results of CNN-QNL-Net model on MNIST and CIFAR-10 datasets.

Dataset	Ansatz	Model	Learning Rate	Average Train Accuracy (%)	Average Test Accuracy (%)
MNIST (0, 1)	0	CNN-QNL-Net	1×10^{-4}	99.97 ± 0.02	99.96 ± 0.03
	1	CNN-QNL-Net	1×10^{-4}	99.96 ± 0.02	99.95 ± 0.02
	2	CNN-QNL-Net	1×10^{-4}	99.96 ± 0.03	99.95 ± 0.04
	0	PCA-QNL-Net	1.5×10^{-4}	99.65 ± 0.17	99.54 ± 0.16
	1	PCA-QNL-Net	1.5×10^{-4}	99.24 ± 0.19	99.18 ± 0.34
	2	PCA-QNL-Net	1.5×10^{-4}	99.67 ± 0.23	99.59 ± 0.21
CIFAR-10 (2, 8)	0	CNN-QNL-Net	3×10^{-4}	94.20 ± 0.77	93.54 ± 0.66
	1	CNN-QNL-Net	3×10^{-4}	94.13 ± 0.45	93.98 ± 0.37
	2	CNN-QNL-Net	3×10^{-4}	94.21 ± 0.32	93.76 ± 0.14
	0	PCA-QNL-Net	4×10^{-4}	81.94 ± 1.51	81.16 ± 1.09
	1	PCA-QNL-Net	4×10^{-4}	81.79 ± 0.34	80.95 ± 0.35
	2	PCA-QNL-Net	4×10^{-4}	81.67 ± 0.73	80.86 ± 0.74

In addition, the CNN-QNL-Net model not only outperformed models like QTN-VQC and Hybrid TTN-MERA but did so using significantly fewer qubits—only four qubits compared to 12 and 8 qubits, respectively. Overall, these empirical results validate the effectiveness of integrating quantum circuits with classical neural network architectures. The QNL-Net models, especially the CNN-QNL-Net, demonstrate competitive performance in binary classification tasks while utilizing fewer resources. These findings establish a new benchmark in quantum machine learning, highlighting the potential for scalable and efficient quantum-enhanced models in practical applications.



4.4 Quantum Reinforcement Learning

The Quantum Natural Policy Gradient (QNPG) algorithm replaces the random sampling used in classical Natural Policy Gradient (NPG) estimators with a deterministic gradient estimation approach, enabling seamless integration into quantum systems. While this modification introduces a bounded bias in the estimator, the bias decays exponentially with increasing truncation levels. This work demonstrates that the proposed QNPG algorithm achieves a sample complexity of $O(\epsilon^{-1.5})$ for queries to the quantum oracle, significantly improving the classical lower bound of $O(\epsilon^{-2})$ for queries to the MDP. More formally, we have that:

Let $\{\theta_k\}$ be the policy parameters generated by the proposed Algorithm, π^* be the optimal policy, and let J_{ρ}^* denote the optimal value of $J_{\rho}(\cdot)$ corresponding to an initial distribution ρ . The proposed algorithm achieves the following bound: $J_{\rho}^* - 1/K \sum_k E[J_{\rho}(\theta_k)] \leq \epsilon_{\text{bias}} + \epsilon$, where ϵ_{bias} is the gap due to neural network parametrization. This results in a sample complexity of $O(\epsilon^{-1.5})$ and an iteration complexity of $O(\epsilon^{-1})$.


The results of this work are threefold. First, we construct a quantum oracle for NPG gradient estimation by leveraging fundamental oracles from the Markov Decision Process (MDP), ensuring efficient computation in quantum environments. Second, we modify the classical Natural Policy Gradient (NPG) algorithm into a deterministic setting, enabling seamless integration into quantum systems while maintaining bounded gradient estimation bias. Third, the proposed approach achieves a sample complexity of $O(\epsilon^{-1.5})$, surpassing the classical lower bound of $O(\epsilon^{-2})$, demonstrating a significant improvement in efficiency.

To the best of our knowledge, this is the first work to demonstrate quantum speedups for parameterized model-free infinite-horizon Markov Decision Processes (MDPs). The results indicate that the QNPG algorithm offers a promising approach to harnessing quantum advantages in practical RL applications, particularly in scenarios with large state and action spaces where parameterized policies are necessary.

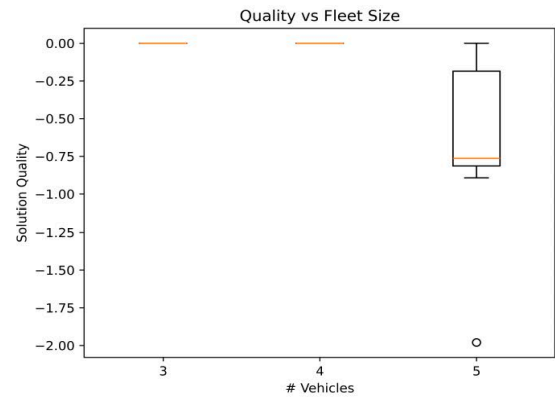
4.5 Arc Interdiction Vehicle Routing Problem using Quantum Annealing

In Figure 6, we compare the QUBO approach to the upper bound (UB), obtained by solving Integer optimization. The solution quality is represented as a percentage loss to the UB. The results indicate that the solution quality of QUBO is comparable to that of UB, while it will require less time on a quantum computer.

Our results highlight both the potential and current limitations of quantum annealing for Network Interdiction problems. For smaller problem sizes, the hybrid solver achieves near-optimal solutions. However, as problem complexity increases, whether due to a larger number of nodes or vehicles, the solution quality of the quantum approach declines. This indicates that while quantum annealing shows promise, further work is required to enhance its scalability and effectiveness for larger, more complex instances.



(a) Solution quality against number of nodes



(b) Solution quality against number of vehicles

Figure 6. A comparison between results from the QUBO approach and the upper bound approach.



CHAPTER 5

Conclusions

Quantum machine learning offers significant advantages in advancing machine learning technologies. In classical machine learning, a major bottleneck is the sheer scale of many AI models, which are often excessively large. For example, large language models incorporate billions or even trillions of variables to achieve remarkable capabilities in recognizing complex image patterns and interpreting audio. The immense scale of these AI models presents numerous challenges, including high memory requirements, overfitting, convergence difficulties, and complex hyperparameter tuning. These issues hinder the development of self-driving vehicles. To address this problem, we have developed a quantum-inspired weight-constrained neural network to reduce the complexity of AI models. Our weight-constrained neural network requires substantially fewer variables than standard classical neural networks, thereby significantly reducing the computational cost for both training and inference.

Moreover, we develop a novel dropout-based approach to enhance the resilience of AI models in vehicles using the weight-constrained neural network. With this approach, the accuracy of the weight-constrained neural network decreases by only 20% under a strong adversarial attack. Furthermore, when the adversarial attack is removed, the accuracy remains largely unchanged. These results suggest that our weight-constrained neural network with dropout not only reduces complexity but also enhances resilience, both of which are essential for safe driving.

In addition to the quantum-inspired weight-constrained neural network, we investigate a hybrid quantum-classical convolutional neural network using angle encoding. We observe that the hybrid model achieves a faster training speed compared to the classical model, which we attribute to the quantum activation function. Inspired by this discovery, we develop quantum-inspired activation functions. Classical neural networks utilizing these quantum-inspired activation functions consistently require fewer training steps compared to those using standard Tanh activation functions. This finding is significant because quantum-inspired activation functions can be directly implemented in classical neural networks to enhance efficiency. As a result, image identification systems in vehicles can be trained using significantly fewer resources.

To enhance pattern recognition, we developed a quantum non-local neural network (QNL-Net) by leveraging quantum parallelism and entanglement to efficiently capture long-range dependencies in data. Traditional non-local operations in computer vision, which compute pairwise relationships across all elements in a feature set, face challenges due to their quadratic complexity in time and memory. QNL-Net addresses these limitations by integrating classical dimensionality reduction techniques with quantum circuits, enabling more efficient computations in a quantum-enhanced feature space. Performance analysis on the MNIST and CIFAR-10 datasets demonstrates that QNL-Net achieves state-of-the-art accuracy in binary image classification among quantum classifiers while using fewer qubits. Specifically, the CNN-QNL-Net model achieved an average test accuracy of 99.96% on MNIST and 93.98% on CIFAR-10, outperforming other quantum models like QTN-VQC and Hybrid TTN-MERA with significantly fewer qubits (4 compared to 12 and 8, respectively). These results validate the effectiveness of integrating quantum circuits



with classical neural networks, establishing a new benchmark in quantum machine learning and highlighting the potential for scalable, efficient quantum-enhanced models in practical applications.

We also developed quantum reinforcement learning (QRL) to improve sample complexity in sequential decision-making. This work introduces the Quantum Natural Policy Gradient (QNPG) algorithm, which replaces classical random sampling with deterministic quantum gradient estimation. By integrating quantum oracles for Markov Decision Process (MDP) queries, QNPG achieves a sample complexity of $O(\epsilon^{-1.5})$, surpassing the classical lower bound of $O(\epsilon^{-2})$. The key enabler of this speedup is the quantum evaluation oracle, utilizing quantum parallelism, amplitude amplification, and entanglement. Unlike classical RL, quantum methods introduce challenges due to the deterministic operations, requiring novel bias-variance trade-offs. This work demonstrates the first quantum speedup for parameterized model-free infinite-horizon MDPs, providing theoretical guarantees on global convergence. The results highlight the potential of quantum computing to enhance RL efficiency in complex decision-making tasks in self-driving vehicles.

NP-hard optimization problems are challenging to solve using classical computers because the solution space grows exponentially with the number of input parameters, whereas quantum annealing offers a feasible approach to addressing them. In this project, we considered the Arc Interdiction Vehicle Routing Problem (VRP) as a case study, which involves an NP-hard optimization challenge that seeks to minimize the traversal costs for a fleet of vehicles. The Arc Interdiction variant introduces an adversary who increases the cost of selected routes. We reformulated this problem, computationally intractable for classical methods, as a quadratic unconstrained binary optimization (QUBO) problem to leverage quantum computing. The bilevel optimization consists of an upper-level adversarial interdiction and a lower-level vehicle routing problem. The QUBO formulation enables solving it via quantum annealing, as demonstrated on D-Wave's Quantum Processing Units. Numerical results show that while small instances achieve near-optimal solutions, larger instances exhibit degraded performance, highlighting both the potential and current limitations of quantum annealing. Hybrid solvers improve efficiency, yet scalability remains a challenge, requiring further advancements to enhance solution quality for large-scale problems.



REFERENCES

- Ajlouni, N. *et al.* (2023) ‘Medical image diagnosis based on adaptive Hybrid Quantum CNN’, *BMC Medical Imaging*, 23(1), p. 126. Available at: <https://doi.org/10.1186/s12880-023-01084-5>.
- Arrazola, J.M. *et al.* (2020) ‘Quantum-inspired algorithms in practice’, *Quantum*, 4, p. 307. Available at: <https://doi.org/10.22331/q-2020-08-13-307>.
- Blank, C. *et al.* (2020) ‘Quantum classifier with tailored quantum kernel’, *npj Quantum Information*, 6(1), pp. 1–7. Available at: <https://doi.org/10.1038/s41534-020-0272-6>.
- Cherrat, E.A. *et al.* (2024) ‘Quantum Vision Transformers’, *Quantum*, 8, p. 1265. Available at: <https://doi.org/10.22331/q-2024-02-22-1265>.
- Chia, N.-H., Lin, H.-H. and Wang, C. (2018) ‘Quantum-inspired sublinear classical algorithms for solving low-rank linear systems’. arXiv. Available at: <https://doi.org/10.48550/arXiv.1811.04852>.
- Ding, C., Bao, T.-Y. and Huang, H.-L. (2022) ‘Quantum-Inspired Support Vector Machine’, *IEEE Transactions on Neural Networks and Learning Systems*, 33(12), pp. 7210–7222. Available at: <https://doi.org/10.1109/TNNLS.2021.3084467>.
- Du, Y. (2021) ‘Quantum noise protects quantum classifiers against adversaries’, *Physical Review Research*, 3(2). Available at: <https://doi.org/10.1103/PhysRevResearch.3.023153>.
- Duong, T.Q. *et al.* (2022) ‘Quantum-Inspired Machine Learning for 6G: Fundamentals, Security, Resource Allocations, Challenges, and Future Research Directions’, *IEEE Open Journal of Vehicular Technology*, 3, pp. 375–387. Available at: <https://doi.org/10.1109/OJVT.2022.3202876>.
- Fan, F. *et al.* (2024) ‘Hybrid Quantum-Classical Convolutional Neural Network Model for Image Classification’, *IEEE Transactions on Neural Networks and Learning Systems*, 35(12), pp. 18145–18159. Available at: <https://doi.org/10.1109/TNNLS.2023.3312170>.
- Gilyén, A., Lloyd, S. and Tang, E. (2018) ‘Quantum-inspired low-rank stochastic regression with logarithmic dependence on the dimension’. arXiv. Available at: <https://doi.org/10.48550/arXiv.1811.04909>.
- Gong, W. and Deng, D.-L. (2022) ‘Universal adversarial examples and perturbations for quantum classifiers’, *National Science Review*, 9(6), p. nwab130. Available at: <https://doi.org/10.1093/nsr/nwab130>.
- Hwang, W. *et al.* (2024) ‘Quantum-inspired classification via efficient simulation of Helstrom measurement’. arXiv. Available at: <https://doi.org/10.48550/arXiv.2403.15308>.
- Jäger, J. and Krems, R.V. (2023) ‘Universal expressiveness of variational quantum classifiers and quantum kernels for support vector machines’, *Nature Communications*, 14(1), p. 576. Available at: <https://doi.org/10.1038/s41467-023-36144-5>.



Kusumoto, T. *et al.* (2021) ‘Experimental quantum kernel trick with nuclear spins in a solid’, *npj Quantum Information*, 7(1), pp. 1–7. Available at: <https://doi.org/10.1038/s41534-021-00423-0>.

Liu, N. (2020) ‘Vulnerability of quantum classification to adversarial perturbations’, *Physical Review A*, 101(6). Available at: <https://doi.org/10.1103/PhysRevA.101.062331>.

Lu, S. (2020) ‘Quantum adversarial machine learning’, *Physical Review Research*, 2(3). Available at: <https://doi.org/10.1103/PhysRevResearch.2.033212>.

Moreira, M.S. *et al.* (2023) ‘Realization of a quantum neural network using repeat-until-success circuits in a superconducting quantum processor’, *npj Quantum Information*, 9(1), pp. 1–7. Available at: <https://doi.org/10.1038/s41534-023-00779-5>.

Pérez-Salinas, A. *et al.* (2020) ‘Data re-uploading for a universal quantum classifier’, *Quantum*, 4, p. 226. Available at: <https://doi.org/10.22331/q-2020-02-06-226>.

Rebentrost, P. (2014) ‘Quantum Support Vector Machine for Big Data Classification’, *Physical Review Letters*, 113(13). Available at: <https://doi.org/10.1103/PhysRevLett.113.130503>.

Tang, E. (2019) ‘A quantum-inspired classical algorithm for recommendation systems’, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. New York, NY, USA: Association for Computing Machinery (STOC 2019), pp. 217–228. Available at: <https://doi.org/10.1145/3313276.3316310>.

Tang, E. (2021) ‘Quantum Principal Component Analysis Only Achieves an Exponential Speedup Because of Its State Preparation Assumptions’, *Physical Review Letters*, 127(6). Available at: <https://doi.org/10.1103/PhysRevLett.127.060503>.

Zoufal, C., Lucchi, A. and Woerner, S. (2019) ‘Quantum Generative Adversarial Networks for learning and loading random distributions’, *npj Quantum Information*, 5(1), pp. 1–9. Available at: <https://doi.org/10.1038/s41534-019-0223-2>.