



PhD Thesis Proposal

Research Title:

**Adaptive Training Strategies for Quantum Neural Networks in
Financial Fraud Data Classification**

Researcher's Name: Mohammed Omar Bin Makhashen

University ID: 242001540

Faculty: Information Technology

Year: 2025

1. Introduction

Placed within the current context of rapid digital transformation, Artificial Intelligence (AI) is increasingly considered as a primary source of innovation in all industries, among which the financial sector. But such futuristic technologies not only work as an opportunity creator but since now there is every possibility that these technologies pave the way for a new breed of threats as well. Artificial intelligence-based financial fraud has become increasingly sophisticated, evasive and challenging to traditional security defense, leading to huge economic losses every year. The reality provides the motivation to develop new generation computational models which are able to capture these complexities and open a window to explore unconventional paradigms like quantum computing.

In this respect, Quantum Machine Learning (QML) becomes particularly a new frontier to explore. Quantum Neural Networks (QNNs)³ are at the core of this field, and they are conjecturally infinitely powerful to learn very complex data structures that cannot be learned by their classical-cloud equivalents (Abbas et al., 2021). This work examines the study's independent variable: Adaptive Training Strategies to these networks. Instead of concentrating efforts mainly on building new architectures, we assume that the way the model is trained is the most important factor. Therefore, the main contribution is a novel training strategy, the "Imbalance-Aware Quantum Natural Gradient" (IA-QNG), that is proposed and its direct effectiveness is tested.

The dependent variable in this study is the Financial Fraud Classification Performance. This is the task from which we desire to enhance the results: the model's performance in correctly discriminating between legitimate and fraudulent transactions. Due to the pathology of the data, this performance will be evaluated with robust statistics adapted to imbalanced data, namely the F1-Score and the Area Under the Precision-Recall Curve (AUC-PR).

This situation presents a perplexing performance paradox in the recent literature. Comparison studies indicate the superior performance of simpler quantum models like Quantum Support Vector Machine (QSVM) against QNN in this particular setting (Innan et al., 2023). This paradox indicates that there is a deserve gap between the representational power of QNNs and the Trainability Crisis.

There are two related things exacerbating that crisis, and one is what I would call the healing burden. The first one is the so-called barren plateau effect (McClean et al., 2018) and the other one is the influence of intrinsic noise of NISQ devices, such which has been demonstrated to generate "noise-induced barren plateaus" (Wang et al., 2021). This fact emphasizes the importance of the mediating variable in our model: Trainability Robustness. We predict that the adaptive training strategy (independent variable) not only has a direct influence on the performance (dependent variable) but that this direct effect assumes adaptive training is better

at retaining informative gradients in the face of noise, i.e. the explanatory mechanism that connects cause and effect.

Moreover, such dependence cannot be explained independent of the statistical properties of the data. This is where the moderating factor, the Degree of Data Imbalance, takes it into play. It is this variable that says “in what conditions”. We expect that the superiority of our approach should not be uniform but more significant with the degree of data imbalance. In other words, as the problem gets “harder” this variable affects the relationship between the training strategy and the final performance more, so that it conditions and strengthens this dependency.

Adaptive training strategies have arisen as a promising field to combat these multivariate challenges (Liu et al., 2023) many still being heuristic. In contrast, the explicit use of QNG is a powerful and principled technique for utilizing the true information geometry of the problem to improve convergence (Stokes et al., 2020). However, this optimizer is not adapted to the statistical properties of the data in its raw form, a crucial limitation in case of a strong imbalance.

This is the precise junction that this study is on. Today, the hole in the literature is not deploying QNNs for the detection of fraud but providing a training framework to address the so-called dual trainability crisis due both to the nature of the hardware (noise) and the nature of the falsifications (imbalance). Thus, the main aim of this thesis is to create and test an adaptive retraining for quantum neural networks (QNNs) strategy, guided by the quantum natural gradient, to mitigate the performance decrease of QNNs on class-imbalanced financial fraud detection data. This thesis addresses a fundamental question: How far can we go in adapting the geometric power of QNG to be combined with a statistical knowledge of the data to relieve the trainability crisis and close the reported performance gap, and thereby enable us to proceed towards a practical quantum advantage in this critical financial use case?

References

- Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & Woerner, S. (2021). The power of quantum neural networks. *Nature Computational Science*, 1(6), 403-409. <https://doi.org/10.1038/s43588-021-00084-1>
- Innan, N., Khan, M. A.-Z., & Bennai, M. (2023). Financial fraud detection: A comparative study of quantum machine learning models. *International Journal of Quantum Information*, 22(02), 2350044. <https://doi.org/10.1142/S0219749923500442>

Liu, S., Zhang, S.-X., Jian, S.-K., & Yao, H. (2023). Training variational quantum algorithms with random gate activation. *Physical Review Research*, 5(3), L032040. <https://doi.org/10.1103/PhysRevResearch.5.L032040>

McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812. <https://doi.org/10.1038/s41598-018-35904-5>

Stokes, J., Izaac, J., Killoran, N., & Carleo, G. (2020). Quantum Natural Gradient. *Quantum*, 4, 269. <https://doi.org/10.22331/q-2020-05-25-269>

Wang, S., Fontana, E., Cerezo, M., Sone, A., Cincio, L., & Coles, P. J. (2021). Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1), 6906. <https://doi.org/10.1038/s41467-021-27045-6>

2. Statement of the Problem

The worldwide financial industry is increasingly threatened by financial fraud, the impact of which amounts to enormous monetary losses. Classical and quantum detection systems, including the well-known existing ones, have proven largely inadequate for handling the pathological properties of fraud data, a domain characterized by extreme class imbalance.

In this background, while Quantum Neural Network (QNN) is widely believed to be theoretically appealing because of the enhanced expressibility (Abbas et al., 2021), recent research results have reported a disturbing performance paradox. Despite the effectiveness of simpler quantum models such as Quantum Support Vector Machines (QSVM) in classification, (Zeguendry et al., 2023) the practical behavior of QNNs in severe imbalanced scenarios is largely unknown. It is alarmingly the case that even seminal works on the power of quantum classifiers tend to neglect this practical concern altogether (Du et al., 2023). This discrepancy indicates that the problem is not simply of representational power, but instead, that these networks suffer a "Trainability Crisis".

This crisis manifests as "barren plateaus" (McClean et al., 2018), a phenomenon heavily intensified by the intrinsic noise in NISQ machines, as seen to yield "noise-induced barren plateaus" (Wang et al., 2021). Adaptive training (Liu et al., 2023) has emerged as a promising way to address this problem, but most are heuristic. The Quantum Natural Gradient (QNG) optimizer, a more properly principled tool which otherwise also exploits the actual information-geometry of the problem (Stokes et al., 2020) seemed to be a better candidate. However its vanilla version is still agnostic to data statistics. These combined issues are the key bottleneck to reach a practical QA in finance (Mancilla and Pere, 2022).

This work hence aims to address a crucial issue:

to what extent a hybrid optimizer exploiting the geometric power of the quantum natural gradient (QNG) and adaptive training principles can potentially solve the joint trainability crisis induced by data imbalance and hardware noise and lead to a feasible quantum advantage in certain real-life problems such as fraud detection?

The paper aims to: To be able to answer the above-mentioned question, the research sets out to address the following specific sub-questions:

- How could the Quantum Natural Gradient (QNG) be rectified to be "imbalance-aware" applying extra focus to the geometry of the inaccessible class (fraud) during the learning phase?

- How successful is our adapted training strategy (Adapted QNG) in enhancing the effectiveness of a QNN, according to robust metrics such as F1-Score and AUC-PR, relative to a strong baseline based classical XGBoost models, QSVM, and QNNs trained using classical optimizers?

References

Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & Woerner, S. (2021). The power of quantum neural networks. *Nature Computational Science*, 1(6), 403-409.

<https://doi.org/10.1038/s43588-021-00084-1>

Du, Y., Yang, Y., Tao, D., & Hsieh, M. (2023). Problem-Dependent Power of Quantum Neural Networks on Multiclass Classification. *Physical Review Letters*, 131(14), 140601.

<https://doi.org/10.1103/PhysRevLett.131.140601>

Liu, S., Zhang, S.-X., Jian, S.-K., & Yao, H. (2023). Training variational quantum algorithms with random gate activation. *Physical Review Research*, 5(3), L032040.

<https://doi.org/10.1103/PhysRevResearch.5.L032040>

Mancilla, J., & Pere, C. (2022). A Preprocessing Perspective for Quantum Machine Learning Classification Advantage in Finance Using NISQ Algorithms. *Entropy*, 24(11), 1656.

<https://doi.org/10.3390/e24111656>

McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812.

<https://doi.org/10.1038/s41467-018-07090-4>

Stokes, J., Izaac, J., Killoran, N., & Carleo, G. (2020). Quantum Natural Gradient. *Quantum*, 4, 269.

<https://doi.org/10.22331/q-2020-05-25-269>

Wang, S., Fontana, E., Cerezo, M., Sone, A., Cincio, L., & Coles, P. J. (2021). Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1), 6906.

<https://doi.org/10.1038/s41467-021-27045-6>

3. Significance of the Research

Importance of the work This work is important as it introduces a targeted computational framework for a key challenge for quantum neural networks to be practically useful. It shows up in the following ways:

3.1. Theoretical (Scientific) Significance:

- **Setting a New Paradigm in Quantum Optimization:** This article offers the new idea of imbalance-awareness in quantum geometric optimization protocol as its main result. This is a major empirical breakthrough that differs from previous heuristics and provides a new frontier for researchers to develop training algorithms that exploit both the geometry of the problem and statistical properties of the data at the same time.
- **Addressing The Trainability Crisis in a Crucial Context:** The research directly aims to solve the well-known "barren plateau" problem, by adapting a strong theoretical tool (QNG) to consider an under-studied context: class-imbalance impacted by stochasticity. This ensures a fairer comparison for the estimated role of quantum neural networks.

3.2. Practical (Applied) Significance:

- **Demonstrating a Practical Proof-of-Concept:** We prove, through empirical results, that meaningful computations of the training-hypothesis yield superior performance to other alternatives. This provides a clear roadmap towards constructing better quantum models for important applications, like financial fraud detection.
- **Creating a Better Prototype:** The developed methodology and its resulting insights can be transferred to the construction of prototype fraud detection systems that are significantly more accurate and robust to noise, presenting an important step towards a working practical quantum advantage.

4. Objectives of the Research

In order to achieve the aim of this study and to respond to its research questions, the researchers have: Main Objective: The overall aim of this study is to reach a final goal, supported by two specific objectives.

- **Objective:** We will propose and design an adaptive training procedure for quantum neural networks using the quantum natural gradient to handle the performance drop due to class imbalance in the financial fraud detection data.

To reach the overall aim, the following specific aims will be pursued:

1. To use the above intuition to derive and construct an optimizer that exploits the imbalance of the cost in the Fubini-Study metric, that is, and Imbalance-Aware Quantum Natural Gradient (IA-QNG) optimizer by defining and formulating the exact mathematical corrections to the Fubini-Study metric which will make the optimizer more sensitive to the minority class.
2. Experimental validation that the proposed model significantly out-performs possible alternatives, by comparing it quantitatively using F1-Score and AUC-PR metrics, and systematically to a strong baseline that includes classical (XGBoost) and standard quantum (QSVM, and QNN trained with standard optimizers) models.

5. Research Questions and Hypotheses

The following research springboards from the pivotal issue of the problem statement, which is also the focal research question:

- **Main Question:** How far can one go with the development of a hybrid optimizer, combining the geometrical power of the Quantum Natural Gradient (QNG) with the adaptive training, to address the dual trainability crisis (emanating from the data itself, in presence of noise), towards an effective practical quantum advantage in real-world financial applications, such as fraud detection?

To satisfy this main purpose, and consistent with the study's objectives, the following research hypotheses will be tested:

- **First Hypothesis (H1):**
 - *Related Objective:* We seek to validate the efficiency of the proposed training approach (IA-QNG) in making the performance of quantum neural network better than using classic methods.
 - *Hypothesis (H1):* There is a significant positive correlation between the use of the developed adaptive training strategy (IA QNG) as an independent variable and the improvement of the classification performance of financial fraud data (measured by F1-Score) as a dependent variable compared to the standard training strategies.
- **Second Hypothesis (H2):**
 - *Related Objective:* To quantify the amount of the documented performance gap between quantum neural networks and other quantum models the proposed training strategy can close.
 - *Hypothesis (H2):* The employment of the introduced IA-QNG leads to a significant improvement in the average classification capability* of QNN in comparison with the QSVM method; *measured by the F1-Score.

6. Proposed Study Model


The following model illustrates the comprehensive conceptual framework of this study, embodying the complex and hypothesized relationships between the research variables. This model is designed to visually reflect the hypotheses, clarifying not only *if* the proposed strategy works, but also *how*, *why*, and *under what conditions*.

Code snippet

```
None
graph TD
    subgraph "Conceptual Model of the Study"
        IV["Independent Variable: Training Strategy  
  
Levels:  
1. IA-QNG Strategy  
2. Standard QNN Strategy  
3. QSVM Strategy"]
        Mediator["Mediating Variable: Trainability  
Robustness  
  
Indicator:  
Gradient decay rate due to noise"]
        DV["Dependent Variable: Classification  
Performance  
  
Metrics:  
1. F1-Score  
2. AUC-PR"]
        Moderator["Moderating Variable: Degree of  
Imbalance  
  
Indicator:  
Minority class ratio in the data"]

        IV -- "Main Effect (H1 & H2)" --> DV
        IV -- "Improves" --> Mediator
        Mediator -- "Positively affects" --> DV
        Moderator -- "Moderates the strength of the relationship" --> IV
    end

    style IV fill:#f0f6ff,stroke:#005cc5,stroke-width:2px
    style DV fill:#f0fff4,stroke:#22863a,stroke-width:2px
    style Mediator fill:#fff7e6,stroke:#f6a821,stroke-width:2px
    style Moderator fill:#fafbfc,stroke:#6a737d,stroke-width:2px
```



Explanation of the Model's Components

- **Independent Variable: "Training Strategy"** This is the core of your innovation. It is the variable you control and whose effect you test by comparing your proposed strategy (IA-QNG) against the baseline strategies (standard QNN and QSVM).
- **Dependent Variable: "Classification Performance"** This is the ultimate result of your measurements that you compare to assess if your actions turned out differently. It will be quantitatively measured using two metrics F1-Score and AUC-PR to reach a fair evaluation when dealing with data imbalance.
- **Mediating Variable: "Trainability Robustness"**
 - **Methodological Role:** It is the explanatory mechanism in the model, the "how?" The way of thinking about this is that your method doesn't magically give better thresholds but introduces a method for preserving strong and meaningful gradients of a loss function in a quantum-mechanical setting (this is just a hypothesis not a conclusion) this in turn allows better performance. It is the connection between premise and consequence.
- **Moderating Variable: "Degree of Imbalance"**
 - **Methodological Role:** This variable state what to "do in what conditions". It's not that it creates the effect, but that it modulates the effect that exists. The model suggests that the advantage of your strategy is not so strong, but it gets bigger and more meaningful when the data is increasingly imbalanced. In other words, The more the problem is "harder" (i.e higher imbalance) your proposed solution, the more it turn out to be the best (in a broad sense I mean). This interplay modifies and strengthens the relationship between your strategy and performance.

7. Study Population and Sample

In line with the computational and experimental research type of this work, the term "study population" and "study sample" are defined methodologically in a way that best comply with the nature of the work. The population is not individuals; it is the complete universe of relevant information and the sample is the procedures employed looking for and handling that universal to prepare the data so that the results achieved by those looking are valid, replicable and reliable.

1. Study Population (Datasets)

In order to more completely accomplish the research goals, a dual-axis approach is being employed, with two sets of data (with different and complementary characteristics):

- **Primary Population (for Development and Initial Validation): PaySim Dataset** The model will be developed and tested within this domain. Since it is fake data with known fraud behavior, it allows a perfect (controlled) environment in order to assess the performance of the proposed model in learning the patterns of interest. Its main characteristics are:
 - **Size:** Over 6.3 million records.
 - **Inbalance:** Very exceptional (only ~0.13% of transactions are fraud), this presents a challenge for traditional model evaluation applications of different learning algorithms.
 - **Characteristics:** Low-dimensional and interpretable feature space (11 features).
- **Generalization Population (for Robustness and Realism Testing): IEEE-CIS Dataset** We are going to test at the end of the process with this dataset to ensure that the model isn't overfit and can generalize its performance on messier, more complex, real world type data. If the model could work here, it would demonstrate that the model is not just fitting the PaySim data. Its main characteristics are:
 - **Source:** Real-world e-commerce transaction data.
 - **Skew:** Vast (fraud is substantial (~3.5% of transactions)).
 - **Features:** The high-dimensional (431) feature space, which evaluate the model for dealing with complex representations.

This methodology guarantees that the model is evaluated under two different sources of challenge, a statistical challenge (weak signal in PaySim) and a representational challenge (complex features in IEEE-CIS).

Attribute	Primary Population (PaySim)	Generalization Population (IEEE-CIS)
Source	Synthetic (Simulator)	Real-world (E-commerce)

Total Transactions	~6.3 Million	~590k (Training)
Imbalance Ratio	~0.13% (Severe)	~3.5% (Significant)
Number of Features	11	431

2. Study Sample and Methods of Selection (The Method of the Experiment)

In this investigation, a static "sample" is not chosen. Instead, we follow a strict experimental 1 protocol to produce multiple training and testing sets in order to obtain statistically reliable results.

- **Data Sample Methodology: Stratified K-Fold Cross-Validation**

- **Procedure:** We will use a stratified k-fold cross-validation procedure (where k=5 or 10) over each dataset. The data is partitioned into k equal subsets in each round, the model is trained on k-1 subsets and tested on a remaining single subset.
- **Methodological Explanation:** Stratification is necessary for this experiment. This means that at each train and test fold the proportion of fraudulent transactions is preserved from the original dataset. Such a step is necessary to avoid the creation of a test set, which is quite important as it is the case here, which does not contain any fraud sample, so as to keep the performance estimate valid. This process is computationally equivalent to making several typical draws from the population, and consequently allows one to carry out rigorous statistical tests on the performance of the model.

8. Proposed Statistical Methods

In order to perform a rigorous and dependable quantitative assessment of the effectiveness of the proposed training strategy, as well as to test the research hypotheses, a thorough statistical protocol - specifically tailored for comparing the performance of computational models on class-imbalanced sets - will be adopted.

1. Experimental Design and Statistical Reliability

In the current study, the independent variable is the Training Strategy, with the following levels:

- **Experimental Level:** The quantum neural network obtained with the IA-QNG strategy trained in the experiments (IA-QNG).
- **Control Groups:**
 1. The quantum neural network with standard training optimizer (QNN-Adam).
 2. The model is the Quantum Support Vector Machine (QSVM).
 3. The classical baseline model eXtreme Gradient Boosting model (XGBoost).

A Stratified 10-Fold Cross-Validation approach will be used in order to assure the consistency and stability of the performance values (dependent variable) of the models. The latter step is not just a validation procedure, but it is a condition of validity of the experiment because:

- **No bias from random split:** It makes certain that the model performance does NOT depend upon a single split of the data.
- **Dealing with Imbalance:** Stratification ensures that the class imbalance is preserved across the training and testing folds, which is necessary for a fair evaluation.
- **Creating a Distribution:** This results in a distribution with 10 values (F1 or AUC-PR) per model, which can be tested statistically, instead of relying on point-estimates comparison.

2. Statistical Protocol for Hypothesis Testing

Analysis of these performance distributions will proceed in a two-step fashion:

- **Step 1: Omnibus Test**
 - **Method: One-Way Repeated Measures ANOVA.**
 - **Objective:** This step is to solve the following question generally asked by developers: “Whether or not, the mean performance of the four training strategies are statistically different? We use a repeated measures design, because we test the same “subjects” (the data folds) on all conditions (the models).
- **Step 2: Post-Hoc Paired Comparisons for Direct Hypothesis Testing**

- **Method:** If the ANOVA test had clear difference, then a series of Paired Samples t-tests will be conducted to test each hypothesis individually, with a Bonferroni Correction to target the level of significance.
- **To Test Hypothesis 1 (H1):** Considering a direct comparison between the performance distribution of IA-QNG and the distribution of QNN-Adam. This is the so crucial test of whether your innovation is better than the oven.
- **To Test Hypothesis 2 (H2):** A direct comparison between the performance distribution of IA-QNG and of QSVM will be done. This comparison seeks to demonstrate that your strategy could close the performance gap with the best known quantum strategy.

3. Tools and Software Used

In view of the computational flavour of the data and analyses of interest, we will implement this complete statistical protocol with the Python programming language, especially leveraging its established scientific libraries (SciPy and Statsmodels), the academic and industry standard ones for data science having thorough implementations for all statistical tests considered above.

9. Proposed Study Structure (Research Outline)

Preliminary Chapter: General Framework of the Study

- **Section One: Introduction to the Study**
 - Subsection One: Introduction and Significance of the Study
 - Subsection Two: Research Problem and Questions
 - Subsection Three: Research Objectives and Hypotheses
 - **Section Two: Research Methodology and Plan**
 - Subsection One: Research Methodology and Conceptual Model
 - Subsection Two: Study Population and Sample (Datasets and Experimental Protocol)
 - Subsection Three: Proposed Statistical Methods
 - Subsection Four: Structure of the Dissertation
-

Chapter One: Theoretical Foundations of Quantum Machine Learning in Fraud Detection

- **Section One: Quantum Neural Networks and the Trainability Crisis**
 - Subsection One: The Mathematical Structure of Quantum Neural Networks
 - Subsection Two: Analysis of the Barren Plateau Phenomenon
 - Subsection Three: The Impact of Hardware Noise on the Training Landscape (Noise-Induced BPs)
 - **Section Two: Machine Learning Challenges in Financial Fraud Data**
 - Subsection One: The Problem of Class Imbalance
 - Subsection Two: Baseline Quantum Models (The Quantum Support Vector Machine - QSVM)
-

Chapter Two: Design and Development of the Adaptive Training Strategy (IA-QNG)

- **Section One: The Geometric Foundation: Quantum Natural Gradient (QNG)**
 - Subsection One: The Fubini-Study Metric and Information Geometry
 - Subsection Two: The Mechanism and Limitations of the Standard QNG Optimizer
- **Section Two: Constructing the Imbalance-Aware Optimizer (IA-QNG)**
 - Subsection One: Mathematical Derivation for Modifying the Geometric Metric

- Subsection Two: The Final Algorithm and Application Protocol for the Proposed Model
-

Chapter Three: Experimental Design and Analysis of Results

- **Section One: The Experimental Environment**

- Subsection One: Characterization of Datasets (PaySim and IEEE-CIS) and Preprocessing
- Subsection Two: Quantum Simulation Settings (Qiskit) and Application of Noise Models

- **Section Two: Presentation and Analysis of Statistical Results**

- Subsection One: Results of Testing the First Hypothesis (IA-QNG vs. QNN-Adam)
 - Subsection Two: Results of Testing the Second Hypothesis (IA-QNG vs. QSVM)
 - Subsection Three: Comprehensive Comparative Analysis of All Models' Performance
-

Chapter Four: Discussion of Results and General Conclusions

- **Section One: Interpretation and Discussion of Results**

- Subsection One: Discussing the Implications of the IA-QNG Strategy's Superiority
- Subsection Two: Analysis of the Model's Behavior Under Noise and Imbalance

- **Section Two: Conclusions and Recommendations**

- Subsection One: Theoretical and Applied Conclusions of the Study
- Subsection Two: Limitations of the Study and Future Research Directions
- Subsection Three: Final Conclusion of the Dissertation

10. Research Terms

To maintain the accuracy and consistency of the main concepts of the study, the following terms are clearly defined in concept and operation.

Quantum Neural Networks (QNNs)

- **Conceptual:** Hybrid quantum-classical models for machine learning that rely on a parameterized quantum circuit, $U(\theta)$, as a (trainable) layer of the model, with its parameters, θ , being determined by a classical computer through the minimization of some cost function (Cerezo et al., 2021).
- **Practical:** A Variational Quantum Classifier (VQC) for which the cost function is realized by measuring the expectation value of an observable, M , such that:

$$L(\theta) = \langle \psi_{in} | U^\dagger(\theta) M U(\theta) | \psi_{in} \rangle$$

Adaptive Training Strategies

- **Theoretical:** A class of optimization algorithms that adapt their algorithms automatically during training to mitigate issues related to complex optimization landscapes, e.g., the barren plateau problem (Kübler et al., 2020).
- **Operational (independent) Variable:** It is the one being studied, denoted in this work by the "ImbalanceAware Quantum Natural Gradient (IA-QNG)" strategy. This is an alteration to the usual QNG update rule:

$$\theta_{t+1} = \theta_t - \eta \cdot g^+ \cdot \nabla L(\theta)$$

where g^+ is the pseudo-inverse of the Quantum Fisher Information (QFI) metric, and η is the learning rate.

Classification Performance

- **Intuitive:** The degree of ability of a machine learning model to apply its predictions to new, previously unseen data points. For the case of imbalanced dataset, for which the dataset we used is one, the best metric to be used is F1-Score because it equally balances the Precision and Recall (Shalev-Shwartz & Ben-David, 2014).
- **Operational:** This represents the dependent variable of the research task, which is to be particularly calculated through F1-Score and provided in the equation as:

$$F_1 = 2 \cdot \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

along with AUC-PR (Area Under the Precision Recall curve) as the other evaluation metric.

Quantum Noise Levels

- **Introduction:** The sum of unwanted disturbances that results in the quantum decoherence, which degrades the computational accuracy of Noisy Intermediate-Scale Quantum (Preskill 2018) devices.
- **Mediating:** This is the variable that we propose to mediate, and is computed by simulating the built-in noise models in Qiskit Aer.

Dataset Imbalance Ratio

- **Terminological:** An attribute of a dataset in which the number of instances belonging to the minority class is much lower than that of the other classes in the dataset (the majority class) (Ali et al., 2015).
- **Operational:** These are the moderating terms in the analysis, which are quantified as the minority-class-to-total-sample ratio.

References

- Ali, A., Shamsuddin, S. M., & Ralescu, A. L. (2015). Classification with class imbalance problem: A review. *Journal of Big Data*, 2(1), 3. <https://doi.org/10.1186/s40537-015-0019-8>
https://www.researchgate.net/publication/288228469_Classification_with_class_imbalance_problem_A_review
- Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational Quantum Algorithms. *Nature Reviews Physics*, 3(9), 625-644. <https://doi.org/10.1038/s42254-021-00348-9>
- Kübler, J. M., Arrasmith, A., Cincio, L., & Coles, P. J. (2020). An Adaptive Optimizer for Measurement-Frugal Variational Algorithms. *Quantum*, 4, 263. <https://doi.org/10.22331/q-2020-05-11-263>
- Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>
- Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.

11. Literature Review

Introduction

This chapter presents the review of the literature-action that supports the theoretical and practical aspects of the study. The purpose of this review is not simply to document what has been done, but to construct an evolving logical argument in order to expose the void in understanding which this dissertation seeks to fill. We start by a review of works that characterize the dimension of the main problem at play the 'performance paradox' for QNNs and their training crisis beyond noise and class imbalance. Then, it would conduct an analysis of the current research attempts to deal with these very challenges -- resulting in the identification of an obvious "white space" in the extant literature, on which this paper hooks its originality and differentiation.

1. Prevalence and Context Studies

This series of studies sets out the "problem" and the "research estrangement" this study intends to challenge on different levels.

1.1. Innan et al. (2023)

1.1.1. Previous Study Data

- **Researcher(s):** Innan, N., Khan, M. A.-Z., & Bennai, M.
- **Year of Publication:** 2023
- **Title of Study:** "Financial Fraud Detection: A Comparative Study of Quantum Machine Learning Models"
- **Type:** Peer-reviewed journal article (International Journal of Quantum Information)

1.1.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Task 1: Comparison of performance of four quantum machine learning models (such as QNN and QSVM) for financial fraud data classification.
- **Sample: Type of Samples:** A standard synthetic financial dataset (like PaySim) with heavy class imbalance. **Population: Data in.** Finance: Financial transactional records for detecting fraud.
- **Methodology: Method:** A comparative experimental study. **Major Tool:** the performance of each model was trained and test on the same dataset to ensure its performance validation.

1.1.3. Main Results and Recommendations

- **Main Results:** (1) The Quantum Support Vector Machine (QSVM) model exhibited unequivocal superiority and had the highest F1-Score of all the models. (2) Quantum Neural Networks (QNNs) proved challenging to train and performed poorly for imbalanced classes.
- **Main Recommendations:** (1) The paper called for further investigations to make the performance and trainability of QNNs better for tasks with imbalanced data.

1.1.4. Commentary, Critique, and Link to the Current Study

- **Strengths:** This paper is abundantly important as it contributes to validate and even better proves existence the “performance paradox” in a crucial and particular situation.
- **Weaknesses** The paper successfully described the problem, however, it stopped short at providing a solution.
- **Similarities:** Both works deal with the same applied problem (quantum fraud detection) and compare the performance of QNNs and QSVMs.
- **Differences and novelty:** The work of (Innan et al. were content with describing the problem, and we pick off from where they left, driven to solve the problem itself.
- **How This Study Will Fill the Gap:** Our study will directly address the gap identified in this study. It responds to their primary suggestion by presenting a novel training scheme (IA-QNG) orientated for specifically compensating experimentally observed limitations of the QNN, considering the performance of QNN is insufficient by the empirical studies.

1.2. Wang et al. (2021)

1.2.1. Previous Study Data

- **Researcher(s):** Wang, S., Fontana, E., Cerezo, M., et al.
- **Year of Publication:** 2021
- **Title of Study:** "Noise-induced barren plateaus in variational quantum algorithms"
- **Type:** Peer-reviewed journal article (Nature Communications)

1.2.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: How Noise-to-Hardware Communication as an Isolated Factor Affects the Training Landscape of VQAs. Objective 2: The theory and experimental observation of the "Noise-Induced Barren Plateaus" (NIBPs).
- **Sample:** Sample Type: computer simulations of variational quantum circuits under assorted noise models. Study Subjects: Variational quantum algorithms on NISQ processors.
- **Methodology: Method of research:** Mathematical analysis and numerical simulation. **Main Result:** The gradient variance goes to zero exponentially with the number of qubits if the noise strength of the circuit is beyond a critical point.

1.2.3. Main Results and Recommendations

- **Main Results:** (1) The indisputable evidence that noise alone can cause the barren plateau phenomenon, even for circuits that are potentially learnable in the absence of noise.(2) This effect is critically determined by the “depth” of the noise (with noise at deeper gates in the circuit worse).
- **Main Recommendations:** 1) There is a need to accomplish the task of training algorithms which are adaptive or are robust to noise.

1.2.4. Commentary, Critique, and Link to the Current Study

- **Strengths:** This work is the missing link that bridges the gap between abstract mathematical difficulties (barren plateaus) and the actual physical constraints of NISQ hardware.
- **Weaknesses:** The weakness in this work is that it only gives an algorithm to identify and prove the problem.

- **Commonalities:** Both papers investigate the influence of noise on trainability.
- **Differences and Uniqueness:** Although (Wang et al. argues that the noise kills the trainability, our work tries to address this problem.
- **How the Current Study Will Address the Gap:** Our study would address the gap by answering this study's call directly. Inspired by the information geometry of the parameter space the optimization process is designed to be inherently more robust and is thus expected to be well-suited to traverse distorted and flatten cost landscapes induced by noise.

1.3. Schumann, M., et al. (2023)

1.3.1. Previous Study Data

- **Researcher(s):** Schumann, M., Wilhelm, F. K., & Ciani, A.
- **Year of Publication:** 2023
- **Title of Study:** "Emergence of noise-induced barren plateaus in arbitrary layered noise models"
- **Type:** Peer-reviewed journal article (Physical Review A)

1.3.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Goal 1: Generalize recent work on "Noise-Induced Barren Plateaus" (NIBPs) introduced by (Wang et al., 2021). Objective 2: To show that this behaviour is not restricted to special noise models, but rather is a common property of noisy layered quantum circuits.
- **Sample:** Sample Type: Completely mathematical and theoretical investigation supported by numerical simulations of variational quantum circuits. Study Population: Variational quantum algorithms in the presence of general noise models.
- **Methodology:** Methodology: A thorough theoretical work grounded in quantum information theory techniques. Main Tool: It was shown that the vanishing gradient is an unavoidable consequence of the existence of a non-unitary noise layer in the circuit.

1.3.3. Main Results and Recommendations

- **Main Results:** (1) It was shown that noise-induced barren plateaus are a universal and generic effect that occurs in every layered noise model. (2) The paper found that the form and location of the noise in the circuit are important in determining the flatness of the training landscape.
- **Main Recommendations:** (1) The work brought out the need to develop training algorithms, which are noise adaptive or noise tolerant as a matter of design.

1.3.4. Commentary, Critique, and Link to the Current Study

- **Strengths:** This is a strong theoretical generalization, which removes any suspicion of the seriousness and commonness of the noise problem in VQAs.
- **Weakness:** It's all about characterizing and generalizing the problem theoretically; there is no algorithmic solution given.
- **Comparisons:** Both papers deal with the same physical-computational phenomenon – the impact of noise on trainability).
- **Differences and Uniqueness:** (Schumann et al.) aims at proving the inevitability of the problem we reason about how to solve it in practice.

- **How the Current Study Will Address the Gap:** Our study will address the gap by directly addressing the recommendation made by this study. Subsequently, we demonstrate here that noise is an unavoidable nuisance in such measurement problems: our study then presents a candidate solution (IA-QNG), better able to resist the destructive effects of noise by its recourse to information geometry.

1.4. Kashif, M., & Al-Kuwari, S. (2023)

1.4.1. Previous Study Data

- **Researcher(s):** Kashif, M., & Al-Kuwari, S.
- **Year of Publication:** 2023
- **Title of Study:** "The impact of cost function globality and locality in hybrid quantum neural networks on NISQ devices"
- **Type:** Peer-reviewed journal article (Scientific Reports)

1.4.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Task 1: Conduct an experimental study of the impact of using "global" and "local" cost functions on the performance and trainability of hybrid quantum neural networks.
- **Sample:** Type of samples: Simulations from a classification task with realistic noise models imposed. Study Population: NISQ- Devices: Hybrid quantum neural networks.
- **Methodology:** MATERIAL AND METHODS: Comparative experimental study. Main Tool: A QNN with the same architecture was trained twice, one with global costfunction and one with local cost function, their training curves and final accuracy were compared.

1.4.3. Main Results and Recommendations

- **Main Results:** (1) It was found that local cost functions are much robust compared to global one against noise and there are not many barren plateaus. (2) The performance of the global cost functions deteriorates at a fast rate in NISQ devices.
- **Main Recommendations:** (1) The work suggested to employ local cost functions as a realistic and trustable option for current quantum applications.

1.4.4. Commentary, Critique, and Link to the Current Study

- **Strengths:** The study gains a thorough understanding of one of the main mechanisms behind trainability crisis.
- **Weaknesses:** The solution proposed here (use of local cost functions) might be restricting the model from capturing more intricate, long-range correlations which could be present in the data.
- **Commonalities:** Both the papers considers the problem of how to handle the trainability on a noisy environment.
- **Differences and Distinctiveness:** (Kashif & Al-Kuwari) focuses on altering "what" cost function is optimized, our work focuses on changing "how" that optimum is achieved (the optimizer).
- **How the Current Study will Fill in the Gap:** Our study will fill in a significant gap by addressing the question: Can a powerful, geometrically-aware optimizer such as IA-QNG work to successfully train

a model which utilizes a challenging “global” cost function, and by doing so, overcome the constraints identified within this study to enable the leveraging of greater model capacity?

1.5. Mancilla, J., & Pere, C. (2022)

1.5.1. Previous Study Data

- **Researcher(s):** Mancilla, J., & Pere, C.
- **Year of Publication:** 2022
- **Title of Study:** “A preprocessing perspective for quantum machine learning classification advantage in finance using NISQ algorithms”
- **Type:** Peer-reviewed journal article (Entropy)

1.5.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: To investigate how pre-processing data can enhance the performance of quantum classifiers over financial tasks.
- **Sample:** Sample Type: A credit scoring data of real financial world. Study Population: Quantum machine learning in financial domain: The finance sector is one of the first adopters of quantum machine learning.
- **Methodology:** Research Method: An experimental approach. Primary Set-Up: LDA was used as a preprocessing for data prepared in a VQA model.

1.5.3. Main Results and Recommendations

- **Main Results:** (1) Application of LDA resulted in remarkable enhancement of the quantum classifier performance. (2) The LDA-boosted quantum model achieved better performance than classical baselines in certain settings.
- **Main Recommendations:** (1) To investigate the communication of classical data, quantum algorithms to reach the ultimate performance.

1.5.4. Commentary, Critique, and Link to the Current Study

- **Strengths:** The study demonstrates that there is a real interest in quantum advantage in finance, and that it is feasible.
- **Weaknesses:** It only focuses on “inputs” of the model and supposes a standard training algorithm should work well.
- **Commonalities:** Both papers apply quantum machine learning for financial data.
- **Differences and uniqueness:** The work in (Mancilla & Pere) is centralized around the model “inputs” (data encoding), whereas our work concentrates on the model “engine” (the training algorithm).
- **How the Current Study Will Address this Gap:** Our study will address a complementary gap. We claim that, even if we have the most suitable encoding, the performance will be bounded if the engine itself (IA-QNG) is not able to deal with the data (imbalance) and hardware (noise) characteristics, such that our contribution would be an essential piece to extract the maximally potentiality of quantum machine learning.

2. Research on Solution Paths (Aspects Unique to This Research)

In this collection the proposed solution is compared with others reported by literature, denoting this research novelty.

2.1. 2.1. Data-Level Solutions (A Comparison with Our Algorithmic Approach)

2.1.1. Study of Li, W., et al. (2023)

2.1.1.1. Previous Study Data

- **Researcher(s):** Li, W., Chen, J., Cao, J., et al.
- **Year of Publication:** 2023
- **Title of Study:** "EID-GAN: Generative adversarial nets for extremely imbalanced data augmentation"
- **Type (Thesis, Published Paper, Book, etc.):** Peer-reviewed journal article (Neurocomputing)

2.1.1.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** 3.1 EID-GAN: The Enhanced Imbalanced Data Generative Adversarial Network Rank 1: Our first application is to propose EID-GAN, an original GAN-based framework designed to address the issue of extreme data imbalance. Goal 2: Obtain superior and diverse augmented synthetic samples for the minority class, such that classifiers with the augmented data outperform those with original samples.
- **Sample:** Samples type: standard imbalanced datasets. Sample Size: The method was tested on numerous datasets of different sizes and natures. Study Population: Multiple datasets from different fields (Non restricted to finance domain) were used to illustrate the model with the general performance.
- **Methodology:** Methodology: Experimental comparative design was followed. Data Collection Instrument: Not applicable (used existing datasets). EID-GAN, the proposed method was the main method that was compared against others including SMOTE.

2.1.1.3. Main Results and Recommendations

- **Main Results of the Study:** (1) The EID-GAN model was able to improve classifier performance in a greater extent than other data generation methods (SMOTE, other GAN variants). (2) The proposed model produced diverse and realistic synthetic samples, while maintaining the true distribution of the minority class and addressing typical problems such as mode collapse.
- **Main Recommendations:** 1) Extending this framework to real world problems with varying complexity such as financial fraud detection and medical diagnosis.

2.1.1.4. Commentary, Critique, and Link to the Current Study

- **Advantages of the Previous Study:** This study is state-of-the-art in classical data-level approaches, and a strong and novel classical algorithm for synthetic data generation that outperforms the competition.
- **Limitations or Unanswered Questions:** Although it has its strengths, the study is in the end completely classical. Moreover, no matter how sophisticated, data generation process

can also add artificial noise or unrealistic patterns and demands a time-consuming preprocessing stage.

- **Common of the Current Study:** Our current research and the previous study are both focused on the same issue i.e., how to enhance classification models with highly imbalanced databases.
- **Differences and Uniqueness of Our Study vs. Other Studies:**
 - The work of (Li et al.) that tackled the pace problem through changing and enriching the data, I address the pace problem at the algorithm level.
 - The study by (Li et al.) is based on a classical technique (GANs), while in my work I employ the quantum-native technique taking advantage of the special property of quantum state space.
- **How the Current Study will Fill the Gap:** My study will fill the gap by offering a more sophisticated and more basic solution; rather than transforming the data to make it fit, as in previous methods readily available with current quantum machine learning models, my study will modify the quantum model's learning algorithm (by IA-QNG) to fit the original sparse-fashioned data distribution. This method circumvents potential bias related to data acquisition and provides a more stable and scientific method.

2.1.2. Huang, Z., et al. (2022)

2.1.2.1. Previous Study Data

- **Researcher(s):** Huang, Z., Sang, Y., Sun, Y., & Lv, J.
- **Year of Publication:** 2022
- **Title of Study:** "A neural network learning algorithm for highly imbalanced data classification"
- **Type (Thesis, Published Paper, Book, etc.):** Peer-reviewed journal article (Information Sciences)

2.1.2.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: We seek to contribute a new training algorithm for neural networks that models data vagueness and is well-suited to severe class imbalanced data. Objective 2: Attempt to control the class imbalance problem at the algorithmic level, and without using resample technique, which could cause data distortion.
- **Sample:** Example: Standard datasets with different level of class imbalance. Size of the Sample: It was tested on numerous datasets from the popular repositories (UCI and Kaggle). Study Population: Binary classification tasks across different disciplines to convey the overall usefulness of our framework.
- **Methodology: Research Design:** An experimental, comparative method was employed.
Statistical Software: Python-Distributed-Version 0.0.1a for performance comparison of power iteration and Sinkhronized-GD All data reported in this paper can be found at <https://github>. The

main instrument is the developed algorithm which is based on the modified process of the gradient update. It takes into account the directions of the gradients coming from the two classes, by means of a controlled "rotation" of the gradient of the minority class samples; this rotation increases the impact of the minority class samples.

2.1.2.3. Main Results and Recommendations

- **Main Results of the Study:** (1) The proposed algorithm overall significantly outperformed most of the classical methods (e.g. SMOTE, and cost sensitive learning) on most of the datasets. (2) It was especially effective in enhancing the performance metrics that are important for imbalanced cases, such as G-mean and F1-Score.
- **Main Recommendations:** (1) This study suggested that this kind of gradient modification should be further extended to more complicated deep network structures.

2.1.2.4. Commentary, Critique, and Link to the Current Study

- **Novelty of the Previous Work:** This work constitutes the current state of the art of classical algorithmic solutions for the class imbalance problem, and introduces a new and effective approach that has shown to be better than others.
- **Weaknesses or Areas Not Covered:** The study is completely inside the classical model and does not investigate if there exist some limitations on the classical world that is only possible to go through them if we believe in new computational paradigms such as the quantum model.
- **Similarity to the Current Study:** The two studies aim to tackle the same fundamental challenge (class imbalance) by adapting the algorithm rather than its learning input.
- **Differences and Novelty of This Work:**
 - While (Huang et al.) considered modifying the Euclidean gradient in a classical parameter space, which is not geometrically endowed, the Quantum Natural Gradient is a modification of the gradient in the quantum Hilbert space that has a drastically different geometric structure.
 - The method of (Huang et al.) is indeed heuristic, and my work is principled based on quantum information physics.
- **How This Study Will Address The Gap:** This study will fill a major gap by shaping this pamphlet into one high classical baseline. My central research question is: Can a quantum-driven method using the quantum geometry of the state space, in other words, IA-QNG, reveal better training trajectories than even these sophisticated classical methods? Beating that threshold would then be compelling evidence for a real practical quantum advantage in this class of problems.

2.2. 2.2. Architecture-Level Solutions (Comparison with Our Training Approach)

These are attempts to design a new “car.” In this light, we take a complementary view to the previous one: instead of designing the car (i.e., a QNN architecture), we design a better “engine” (IA-QNG) that is plug-and-play into any car.

2.2.1. Ubale et al. (2025)

2.2.1.1. Previous Study Data

- **Researcher(s):** Ubale, R., K., S. K., Deshpande, S., & Byrd, G. T.
- **Year of Publication:** 2025 (Preprint)
- **Title of Study:** "Toward Practical Quantum Machine Learning: A Novel Hybrid Quantum LSTM for Fraud Detection"
- **Type (Thesis, Published Paper, Book, etc.):** Preprint paper published on arXiv.

2.2.1.2. 2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: To develop a novel hybrid architecture, which fuses classical Long Short-Term Memory (LSTM) networks with quantum features. Objective 2: Enhance the performance of the model in detecting fraud in financial data involving sequential (temporal) relation.
- **Sample:** Sample: Time-series financial transaction. Sample Size: Not explicitly stated in the abstract, but it was used on fraud data. Population of Study: A time series of digital payment transactions.
- **Methodology:** Method: The research method would be experimental that is building and testing performance of the proposed hybrid architecture (Quantum LSTM). Instruments For Data Collection: NA (pre-collected datasets were used). The key feature is the architecture that has been developed, in which the quantum part deals with specific information about the data, before passing it to the classical LSTM.

2.2.1.3. 3. Main Results and Recommendations

- **Main Results of the Study:** (1) Early investigation has proved that combining with quantum behavior of device components may benefit and to improve the capability of LSTM models for pattern recognition in longitudinal data. (2) The proposed model achieved comparable performance to its purely classical alternatives.
- **Main Recommendations:** (1) The research suggests the necessity of more extensive experiments with larger and more difficult examples. (2) It also suggested to enhance the performance of the quantum part of the model for saving the computational resources.

2.2.1.4. Comment and Critique and Link to the Study at Hand

- **Strengths of the Prior Study:** The study introduces a substantial architectural advance and investigates the application of quantum computing for time series data, which is an important and rapidly maturing area.
- **Weaknesses** or Points that are Not Tackled: In the paper, only the model structure is varied rather than the model's trainability. The paper does not touch upon how this new architecture would be impacted by barren plateaus, noise or how it would deal with class imbalance.

- **Similarities** to this work: Both are studies of the same applied problem (financial fraud detection) and seek to optimize quantum model performance.
- **Differences** and Novelty of the Current Work:
 - While the work of (Ubale et al.) studies a solution through Architecture Design, my work studies a solution based on Training Algorithm Engineering.
 - They think that the “car’s chassis” (itself) has to be changed, where I think designing a “better engine” (but still using that same chassis) is the more principled and wider-scoping approach.
- **How This Study Will Fill the Gap:** My study will fill the gap by showing that the training algorithm is the most important thing. Rather than creating a new “car” concept, I’m working to develop a better “engine” (IA-QNG) that you can drop into almost any “car” (i.e. QNN architecture including a Quantum LSTM) and make it perform well in unfavorable conditions. This is a more general and generally applicable contribution.

2.2.2. Das et al. (2025)

2.2.2.1. Previous Study Data

- **Researcher(s):** Das, S., Meghanath, A., Behera, B. K., et al.
- **Year of Publication:** 2025 (Preprint)
- **Title of Study:** “QFDNN: A Resource-Efficient Variational Quantum Feature Deep Neural Networks for Fraud Detection and Loan Prediction”
- **Type:** Preprint paper published on arXiv.

2.2.2.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** i) Objective 1 -Proposal of a new hybrid architecture (QFDNN) able to extract features of financial data by including a quantum layer. Objective 2: To further process quantum-extracted features by a classical deep neural network for the final classification, seeking for higher accuracy under the constraint on resources.
- **Sample: Domain:** FinanceRelated Datasets: Used datasets for fraud detection and loan prediction. Size of Dataset: Not provided explicitly, but implied in application on real world financial data. Study Population: Bank customers (study population) for credit and fraud related tasks.
- **Methodology: Methodology:** An experimental research approach consisting of the construction and performance assessment of the proposed hybrid architecture. Primary Instrument: Proposed QFDNN model.

2.2.2.3. Main Results and Recommendations

- **Main Results: Results** (1) Our architecture presented effective performance on finance classification tasks. (2) Feature wise quantum layer is potentially more resource friendly compared to fully quantum model.
- **Main Recommendations:** (1) The study proposed trying more complex quantum architectures and learning on larger datasets.

2.2.2.4. Commentary, critique, and link to current study

- **Strengths:** It introduces a modern and important application of hybrid models in finance and proposes a novel architecture.
- **Weaknesses:** The architecture design is the only innovation in the paper, given a vanilla training algorithm is enough. The authors did not study the trainability crisis of the quantum part of the model.
- **Commonalities:** The two works tackle the same application (fraud detection) and are both interested in using quantum computation.
- **Differences and Uniqueness:** Although in the work of (Das et al.) works on re-arranging the "structure" of the model, while we make quite a different, more root, assumption: The problem lies in the "training algorithm".
- **How the Proposal Will Fill the Gap:** Our proposal will fill an essential gap by trying to show that for any QNN architecture (including QFDNN), its performance can be substantially increased when it is trained with a smarter, data-aware strategy (IA-QNG). This renders our contribution more fundamental and more applicable.

2.3. Solutions Based on Problem Reframation

These solutions tell us "how" to wander across the loss landscape. Our work is more theoretical as it does so to the "geometry" of this landscape itself by IA-QNG, altering not just the "length of an optimal step," but its definition.

2.3.1. Park, G., & Huh, J. (2023)

2.3.1.1. Previous Study Data

- **Researcher(s):** Park, G., & Huh, J.
- **Year of Publication:** 2023
- **Title of Study:** "Variational Quantum One-Class Classifier"
- **Type:** Peer-reviewed journal article (Machine Learning: Science and Technology)

2.3.1.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Task 1: To implement a VQOC model for the anomaly detection scenario. Aims 2: Generalisation of the model to problems with training data for only one class (normal) available.
- **Sample:** Experimental Setting: You used the standard anomaly detection data sets (e.g, modified version of fraud data). Study Population: General anomaly detection problems.
- **Methodology:** Research Method: An experimental approach.
-
- **Main Tool:** The model was trained to find a close boundary of the normal data distribution in the quantum feature space. All new data points falling beyond this boundary are considered anomalous (fraudulent).

2.3.1.3. Main Results and Recommendations

- **Main Results:** (1) The model had high ability of anomaly detection for samples that were far from the training distribution. (2) For low dimensional tasks, it compared favorably with classical one-class classifiers.
- **Main Recommendations:** 1) To enhance the scalability of the model and to utilize more powerful quantum kernel methods.

2.3.1.4. Discussion, Critique, and Link to the Present Study

- **Strengths:** The article provides an elegant solution in the context of quantum mechanics to the asymmetry problem - by nearly framing it away.
- **Challenges:** To a wide extend, this method relies on the fact that all valid transactions fall under one learnable distribution and hence may miss on the less obvious patterns in the fraud class itself.
- **Similarities:** Both works are targetting the task of imbalanced data within a quantum setting.
- **Difference and Uniqueness:** The paper (Park & Huh) converts the problem into "anomaly detection" from "binary classification" but addresses the original problem directly.
- **How the Current Study Will Fill the Gap:** Our study will fill the gap by staying in the supervised binary classification framework, so that we can learn from the few fraud examples available. This makes our method more useful in a situation where we are trying to detect fraud of a particular kind and not just everything which is "abnormal".

3. Theoretical underpinning to the proposed solution

These analyses constitute a theoretical underpinning for our technical decisions and explain the significance of our solution.

3.1. Stokes et al. (2020)

3.1.1. Previous Study Data

- **Researcher(s):** Stokes, J., Izaac, J., Killoran, N., & Carleo, G.
- **Year of Publication:** 2020
- **Title of Study:** "Quantum Natural Gradient"
- **Type:** Peer-reviewed journal article (Quantum)

3.1.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Task 1: Definition of the QNG optimizer as a principled optimisation scheme We introduce a formal definition of the Quantum Natural Gradient (QNG) optimiser, including an algorithm for its implementation. Objective 2: Prove that training of VQAs can be expedited by orders of magnitude using this strategy rooted in the information geometry of quantum state space.
- **Sample:** Kind of Data: Benchmark problem in quantum chemistry and physics (e.g., for determining molecular ground states). Study Population: General variational quantum algorithm Optimisation Problem.

- **Methodology:** Method: A theoretical and computational study. Primary Tool: Derivation of the Fubini-Study metric mathematically, followed by preconditioning the gradient with the latter and simulation verification.

3.1.3. Main Results and Recommendations

- **Main Results:** (1) The QNG optimizer combines much faster convergence and a better stability than standard gradient-based optimizers, e.g. Adam. (2) QNG leads more directly to the optimal solution as it takes the actual geometry of the problem.
- **Main Recommendations:** (1) This work suggests the applicability of QNG to the variety of quantum problems and investigates an efficient approximation of it.

3.1.4. Commentary, Critique, and Relation to the Present Study

- **Pros:** Experimental demonstration of QNG being not just-another-optimizer but information physics based approach.
- **Weaknesses:** The presentation of QNG was somewhat of a letdown in that it was left in the abstract and theoretical world; no guidance whatsoever on how it would actually behave or perform "in the field" when faced with real world issues like extreme class imbalance.
- **Connections:** We are coming up based on a strong theoretical tool in this work.
- **Difference and Uniqueness:** Whereas (Stokes et al. presented a "data-agnostic" optimizer, its novelty lies in our "data-aware" optimizer.
- **How the Study Will Address the Gap:** The current study will fill an important gap between theory and practice. We do not throw away this strong theoretical underpinning, instead we solidify it with our novel adaptation and specialization of it to an open practical problem, so that the geometrically powerful optimizer can be also the statistically intelligent one.

3.2. Liu, S., et al. (2023)

3.2.1. 1. Previous Study Data

- **Researcher(s):** Liu, S., Zhang, S.-X., Jian, S.-K., & Yao, H.
- **Year of Publication:** 2023
- **Title of Study:** "Training variational quantum algorithms with random gate activation"
- **Type:** Peer-reviewed journal article (Physical Review Research)

3.2.2. 2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Goal 1: Develop a novel training algorithm using "random gate activation." Objective 2: Addressing the issues of barren plateaus and local minima together.
- **Sample: Type of Sample:** Variational quantum simulation problems (e.g., ground state of a Hamiltonian). Study Population: VQAs in all their glory.
- **Methodology:** Method of Research: Computational experimental research. Main Tool: An adaptive procedure is adopted where in each training step only a random subset of the parameters of the circuit are activated and trained; thus effectively enforcing a small number of active parameters at any time.

3.2.3. 3. Main Results and Recommendations

- **Main Results:** (1) The method effectively ameliorated the problem of barren plateaus by restricting the effective parameter number in each training step. (2) The activations would be random during training, helping the optimizer to get out of local minima.
- **Main Recommendations:** (1) To examine dynamic and adaptive training strategies to circumvent the issue of trainability.

3.2.4. Commentary, Critique, and Link to This Study

- **Strengths:** This study represents invaluable evidence of a paradigm shift towards dynamic and adaptive rehabilitation tools in the field.
- **Weaknesses:** The proposed method is stochastic and heuristic, with a random aspect 'hope to' find a better path.
- **Commonalities:** • Both works propose using "adapting the training process" as a solution to the trainability crisis.
- **Differences and Uniqueness:** Our algorithm is similar to (Liu et al.) is stochastic, the IA-QNG method we present is actually non-stochastic but deterministic and geometrically informed.
- **How the Current Study Will Fill the Gap:** We intend to fill a large gap in the research by presenting a fundamentally more principled solution. Rather than randomly guessing, in our research we exploit the fundamental mathematical structure of the problem to find the ideal search direction. Thereby, our proposed solution cannot only be considered a "clever trick," but an approach which is challenging the geometric nature of the problem.

3.3. Sannia, A., et al. (2024)

3.3.1. Previous Study Data

- **Researcher(s):** Sannia, A., Tacchino, F., Tavernelli, I., & Barkoutsos, P.
- **Year of Publication:** 2024
- **Title of Study:** "Engineered dissipation to mitigate barren plateaus"
- **Type:** Peer-reviewed journal article (npj Quantum Information)

3.3.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: To investigate a novel ansatz to make the barren plateau effect disappear, not through modifications to the optimizer, but to the quantum system-environment (dissipative) interactions.
- **Sample:** Sample Type: Theoretical problems and simulation of variational quantum circuits. Population of Study: Variational quantum algorithms that are subject to barren plateaus.
- **Methodology:** Method: A theoretical as well as computational experimental method is developed which here differs from the usual treatment of reducing the dynamics of the quantum systems to either a small or finite Hilbert space.

3.3.3. Main Results and Recommendations

- **Main Results:** The results demonstrate that the engineered dissipation can be utilized to mitigate the vanishing of gradients. (2) This approach offers a different route for coping with trainability difficulties.

- **Main Recommendations:**(1) The study suggested to investigate combining the system-level modification coupled with improved optimizers (i.e., hybrid techniques).

3.3.4. Commentary, review, and link to current study

- **Strengths:** This work is at the cutting edge of how to approach the trainability crisis from the point of view of physics.
- **Weaknesses:** This is still a hardware level approach, not easy to scale across all platforms or problems.
- **Commonalities:** Both works consider the same issue (the barren plateau) and have the same goal (its resolution).
- **Differences and Uniqueness:** Whereas the work of (Sannia et al. concentrates on changing the physics of the quantum system, our work follows an obviously distinct and actually complementary direction, it considers the same problem but in the lens of information geometry and data science.
- **How to Fill This Gap in the Current Study:** our study will fill an important gap by offering a software-based solution. We do not apply changes to the quantum "device" but to the "intelligence" that trains it (IA-QNG); this is an original and more generalizable research vertex.

3.4. Cerezo, M., & Coles, P. J. (2023)

3.4.1. Previous Study Data

- **Researcher(s):** Cerezo, M., & Coles, P. J., et al.
- **Year of Publication:** 2023
- **Title of Study:** "The expressibility and trainability trade-off in quantum neural networks"
- **Type:** Peer-reviewed journal article (npj Quantum Information)

3.4.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Objective 1: Theoretical In-Depth Analysis for the Core Trade-off Relationship between Expressivity and Learnability in the Quantum Neural Network.
- **Sample:** Type of Sample: Non applicable (theoretical study). Population: Quantum neural networks as a whole.
- **Methodology:** Method of Approach: A mathematical and information-theoretic approach relating properties of a quantum circuit to the likelihood of finding a barren plateau.

3.4.3. Main Results and Recommendations

- **Main Results:** (1) The experiment showed that the more expressive the circuit, the higher is the probability of the loss landscape to be a barren plateau. (2) This suggests a fundamental tradeoff: strong circuits are difficult to train, and circuits that are easier to train may be too weak.
- **Main Recommendations:** (1) Strategies (in circuit design or optimization) that attempt to achieve a good compromise in this trade-off.

3.4.4. Discussion, Critique, and Link to the Current Study

- **Strengths:** This is a significant theoretical contribution that offers a greater context on why the lack of trainability crisis is foundational.

- **Drawbacks:** It is limited to the theoretical characterization of the problem and do not provide the complete practical solution.
- **Common Ground:** Both papers identify the trainability crisis as a critical issue.
- **Differences and Novelty:** In theory, the "trade-off" is proposed here, and our work provides a practical resolution.
- **How the Current Study Will Fill the Gap Our study** will fill the gap by offering pragmatic guidance (IA-QNG) for "better balance" between the two considerations. Rather than resorting to a wildly expressive circuit that may in fact be untrainable, we employ a standard circuit and aim to devise a smart optimizer that can successfully traverse its problematic landscape of loss, achieving a more flexible balance between trainability on one hand, and requisite expressibility on the other.

3.5. Garcia, A., & Gacon, J. (2024)

3.5.1. Previous Study Data

- **Researcher(s):** Garcia, A., & Gacon, J.
- **Year of Publication:** 2024
- **Title of Study:** "Mitigating barren plateaus in quantum neural networks"
- **Type:** Peer-reviewed journal article (Communications Physics)

3.5.2. Summary of the Previous Study

- **Main Objectives and Hypotheses:** Aim 1: Proposing and testing an alternative technique to address the barren plateau issue, by changing the initial setting of the parameters of the quantum circuit.
- **Sample:** Type of sample: Variational quantum circuits simulated by a classical computer, both for benchmark classification problems and theoretical problems. Population: Quantum-trained neural networks with trainability issues.
- **Methodology:** Research Design: A computer experiment design. The workhorse: A method was introduced in which each layer was initialized using initialized circuit parameters layers whose training is decoupled to prevent their gradient from vanishing, and then subsequently learning the entire circuit globally.

3.5.3. Main Results and Recommendations

- **Main Results:** The proposed layer-wise initialization method has a strong ability to escape from barren plateaus at the early training stage. (2) The approach caused faster convergence and better performance than the standard random initialization.
- **Main Recommendations:** (1) The authors suggested testing this approach on more compound quantum layout (not limited to nearest neighbor) and on a real quantum device.

3.5.4. Commentary, Critique, and Link toThe Present Study

- **Strengths:** This paper is a crucial step toward addressing the trainability crisis via understanding of the initialization.
- **Weakness** The solution as suggested merely promises a "good start" for optimization, but cannot guarantee the optimization path itself to be optimal after initialization.

- **Connections:** Both works are on the same topic (barren plateaus) and focused on the enhancement of trainability.
- **Differences and novelty:** (Garcia & Gacon) study is an analysis about the beginning of the optimization (Initial state), here we studied the optimization trajectory.
- **Contribution of the Present Study: Our study** is expected to fill in the gap by providing a more flexible and complete approach than previous work. Despite (Garcia & Gacon)'s study that guarantees that we start the training in a "good place" on the loss landscape, We believe that our IA-QNG optimizers represents an end-to-end solution that can impact all the training at any point, searching for the best road to the minima, not only from this starting condition.

General comment on previous studies and derivation of research gap

From the critical and methodological examination above, it can be seen that there is a distinct and precise research gap existing in studies. Although in recent literature the dimensions of the problem have been identified, different solutions have been proposed, thus casting a wide room for innovation. The research context can be characterised as follows:

First, we are in the fortunate position that studies have effectively meaningfully quantified the problematic nature of the situation: papers such as (Innan et al., 2023) have enshrined a performance-oriented 'hnce known as the 'trainability paradox'' in the realm of [financial] fraud whereas uniquely important studies like (Wang et al., 2021) and (Schumann et al., 2023) have pinned down just how this crisis of trainability conflicts with reality, highlighting how noise inherent to NISQ devices becomes, as costly as you wish it away to be, both extremely costly to minimize and to op-timize away to the levels required compatible with an effective theory of everything.

Second, belonging to this clearly defined optimization problem, solution paths were sprouted in their own directions with their own techniques and drawbacks:

- **Data-Level Solutions** (e.g., Li, W., et al., 2023): This family captures techniques that rectify the issue by altering the input data (such as using a GAN to produce data that obeys the model), rather than simply addressing the lack in the learning algorithm.
- **Architecture-Level Solutions** (e.g., Ubale et al., 2025; Das et al., 2025): Such methods lie on the design of new model structures, suggesting that the answer lies in modifying the "chassis" of a model rather than its "engine" during training.
- **Problem-Reframing Solutions** (e.g., Park & Huh, 2023): These avoid the main problem by encoding it in an alternative representation (e.g., a one-class classification).
- **Optimization-Level Solutions** (e.g., Liu, S., et al., 2023; Sannia et al., 2024): These provide novel solutions but they are either heuristic and stochastic and/or concentrate in modifying the physics of the system, rather than the algorithm information-based behaviour.

The void to be filled by this study is at the very convergence of those two trends. To the author's knowledge, there is no available such scheme which combines together all of the following qualities:

1. **An algorithm-level answer:** It is a correct to the deficiencies in the learning "engine" And not related to the data nor to the simple architecture.
2. **A Geometry-Informed Solution:** It uses the deep mathematical structure of the problem (by QNG), rather than heuristics.
3. **A Data-Dependent Solution:** By taking into account the statistic properties of the data (class imbalance), it does not operate "blindly" with respect to them.

To this end, in this dissertation, I propose to fill this apparent gap by developing and testing a new adaptive training protocol (IA-QNG) which addresses the dual trainability crisis from a joint geometric and statistical viewpoint. This is exactly the main novelty and original contribution of this work.

12. Study Implementation Plan

The action plan for implementation of the study is composed of five general sequential phases :

1. **Phases: 1. Theoretical and Methodological Premises (Preliminary and Initial Chapters)**
 - Finish the coverage of basic concepts with a thorough theoretical review.
 - Complete the full methodology of the investigation.
2. **Phase 2: Algorithm Development (Chapter Two)**
 - Conclude the proof and the final expression for the optimizer algorithm equation (IA-QNG).
 - Implement the programmatic prototype for the optimizer and the quantum circuits using Qiskit or PennyLane.
3. **Phase 3 : Experiment and data collection Issue 3 (Chapitre Trois)**
 - Experimental Environment Setup We setup the experimental environment to load and pre-process the datasets (PaySim and IEEE-CIS) and set the noise models on the simulator.
 - Run the experiments of all models (proposed and control) systematically, according to the stratified cross-validation procedure.
 - Save the raw performance scores for each of the cross-validation fold.
4. **Phase 4: Statistical Analysis of Results (Chapter Three)**
 - Apply the accepted statistical methods (ANOVA and t-tests) applied to the collected performance distributions.
 - Test H1 and H2 and identify the statistical significance of the results.
5. **Phase 5: Final write-up, discussion (Chapter Four and Entire Dissertation)**
 - Discuss the theoretical and practical interpretation of the statistical findings.
 - Compose the summary, the recommendations, and finally, the limitations and projections of the study.
 - Write and revise each chapter of your dissertation for submission.

13. References

1. Books

Nielsen, M. A., & Chuang, I. L. (2010). *Quantum computation and quantum information: 10th anniversary edition*. Cambridge University Press.

Schuld, M., & Petruccione, F. (2021). *Machine learning with quantum computers*. Springer.

2. Published Journal Articles

2.1. Foundational Papers

Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202. <https://doi.org/10.1038/nature23474>

Das, S., Meghanath, A., Behera, B. K., Mumtaz, S., Al-Kuwari, S., & Farouk, A. (2025). *QFDNN: A resource-efficient variational quantum feature deep neural networks for fraud detection and loan prediction*. arXiv preprint arXiv:2504.19632. <https://doi.org/10.48550/arXiv.2504.19632>

Liu, X., Liu, G., Huang, J., Zhang, H.-K., & Wang, X. (2022). *Mitigating barren plateaus of variational quantum eigensolvers*. arXiv. Retrieved from <https://arxiv.org/abs/2205.13539>

McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), 4812. <https://doi.org/10.1038/s41467-018-07090-4>

Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. *Quantum*, 2, 79. <https://doi.org/10.22331/q-2018-08-06-79>

Rebentrost, P., Mohseni, M., & Lloyd, S. (2014). Quantum support vector machine for big data classification. *Physical Review Letters*, 113(13), 130503. <https://doi.org/10.1103/PhysRevLett.113.130503>

Stokes, J., Izaac, J., Killoran, N., & Carleo, G. (2020). Quantum Natural Gradient. *Quantum*, 4, 269. <https://doi.org/10.22331/q-2020-05-25-269>

Ubale, R., K., S. K., Deshpande, S., & Byrd, G. T. (2025). *Toward practical quantum machine learning: A novel hybrid quantum LSTM for fraud detection*. arXiv preprint arXiv:2505.00137. <https://doi.org/10.48550/arXiv.2505.00137>

2.2. Recent Research Supporting the Proposal

Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & Woerner, S. (2021). The power of quantum neural networks. *Nature Computational Science*, 1(6), 403–409. <https://doi.org/10.1038/s43588-021-00084-1>

Ali, A., Shamsuddin, S. M., & Ralescu, A. L. (2015). Classification with class imbalance problem: A review. *Journal of Big Data*, 2(1), 3. <https://doi.org/10.1186/s40537-015-0019-8>

Cerezo, M., Coles, P. J., Bultrini, L., et al. (2023). The expressibility and trainability trade-off in quantum neural networks. *npj Quantum Information*, 9(1), 76.

<https://doi.org/10.1038/s41534-023-00744-2>

Cerezo, M., Arrasmith, A., Babbush, R., Benjamin, S. C., Endo, S., Fujii, K., McClean, J. R., Mitarai, K., Yuan, X., Cincio, L., & Coles, P. J. (2021). Variational Quantum Algorithms. *Nature Reviews Physics*, 3(9), 625–644. <https://doi.org/10.1038/s42254-021-00348-9>

Cybulski, J. L., & Nguyen, T. (2023). Impact of barren plateaus countermeasures on the quantum neural network capacity to learn. *Quantum Information Processing*, 22(9), 348.

<https://doi.org/10.1007/s11128-023-04187-8>

Du, Y., Yang, Y., Tao, D., & Hsieh, M. (2023). Problem-Dependent Power of Quantum Neural Networks on Multiclass Classification. *Physical Review Letters*, 131(14), 140601.

<https://doi.org/10.1103/PhysRevLett.131.140601>

Huang, Z., Sang, Y., Sun, Y., & Lv, J. (2022). A neural network learning algorithm for highly imbalanced data classification. *Information Sciences*, 610, 1111-1126.

<https://doi.org/10.1016/j.ins.2022.08.074>

Innan, N., Khan, M. A.-Z., & Bennai, M. (2023). Financial fraud detection: A comparative study of quantum machine learning models. *International Journal of Quantum Information*, 22(02), 2350044.

<https://doi.org/10.1142/S0219749923500442>

Kashif, M., & Al-Kuwari, S. (2023). The impact of cost function globality and locality in hybrid quantum neural networks on NISQ devices. *Scientific Reports*, 13(1), 16409.

<https://iopscience.iop.org/article/10.1088/2632-2153/acb12f>

Kübler, J. M., Arrasmith, A., Cincio, L., & Coles, P. J. (2020). An Adaptive Optimizer for Measurement-Frugal Variational Algorithms. *Quantum*, 4, 263.

<https://doi.org/10.22331/q-2020-05-11-263>

Liu, S., Zhang, S.-X., Jian, S.-K., & Yao, H. (2023). Training variational quantum algorithms with random gate activation. *Physical Review Research*, 5(3), L032040.

<https://doi.org/10.1103/PhysRevResearch.5.L032040>

Li, W., Chen, J., Cao, J., Ma, C., Wang, J., Cui, X., & Chen, P. (2023). EID-GAN: Generative adversarial nets for extremely imbalanced data augmentation. *Neurocomputing*, 527, 184-195.

<http://dx.doi.org/10.1109/TII.2022.3182781>

Mancilla, J., & Pere, C. (2022). A Preprocessing Perspective for Quantum Machine Learning Classification Advantage in Finance Using NISQ Algorithms. *Entropy*, 24(11), 1656.

<https://doi.org/10.3390/e24111656>

Park, G., & Huh, J. (2023). Variational quantum one-class classifier. *Machine Learning: Science and Technology*, 4(1), 015006. <https://doi.org/10.1088/2632-2153/acafd5>

Sannia, A., Tacchino, F., Tavernelli, I., & Barkoutsos, P. (2024). Engineered dissipation to mitigate barren plateaus. *npj Quantum Information*, 10(1), 81. <https://doi.org/10.1038/s41534-024-00875-0>

Schumann, M., Wilhelm, F. K., & Ciani, A. (2023). Emergence of noise-induced barren plateaus in arbitrary layered noise models. *Physical Review A*, 108(2), 022419. <https://doi.org/10.48550/arXiv.2310.08405>

Stewart, K., & Neftci, E. (2022). Meta-learning spiking neural networks with surrogate gradient descent. *Neuromorphic Computing and Engineering*, 2(4), 044004. <https://doi.org/10.48550/arXiv.2201.10777>

Wang, S., Fontana, E., Cerezo, M., Sone, A., Cincio, L., & Coles, P. J. (2021). Noise-induced barren plateaus in variational quantum algorithms. *Nature Communications*, 12(1), 6906. <https://doi.org/10.1038/s41467-021-27045-6>

Zeguendry, A., Jarir, Z., & Quafafou, M. (2023). Quantum Machine Learning: A Review and Case Studies. *Entropy*, 25(2), 287. <https://doi.org/10.3390/e25020287>

3. Theses and Dissertations

Beer, K. (2022). *Quantum neural networks* [Doctoral dissertation, Leibniz University Hannover]. <https://doi.org/10.15488/11896>

4. Institutional Reports & Online Sources

McKinsey & Company. (2023). *Quantum technology use cases as fuel for value in finance*. McKinsey Global Institute. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/tech-forward/quantum-technology-use-cases-as-fuel-for-value-in-finance>