

UNIVERSITÀ DEGLI STUDI  
DEL SANNIO Benevento

Department of Engineering  
Piazza Roma, 21  
Benevento, Italy

PH.D. THESIS  
XXXVIII Cycle

FRANCESCO MAURO

# Quantum Machine Learning for Earth Observation: Methods, Applications, and Future Directions

**Advisor**

Prof. Silvia Liberata Ullo  
University of Sannio

**Co-Advisor**

Prof. Paolo Gamba  
University of Pavia

**Ph.D. Coordinator**

Prof. Massimiliano Di Penta  
University of Sannio

---

# Contents

|                                                                                |           |
|--------------------------------------------------------------------------------|-----------|
| <b>List of Figures</b>                                                         | <b>7</b>  |
| <b>List of Tables</b>                                                          | <b>11</b> |
| <b>1 Introduction</b>                                                          | <b>23</b> |
| 1.1 Motivation and Research Context . . . . .                                  | 23        |
| 1.1.1 The growing importance of EO in science and society . . . . .            | 23        |
| 1.1.2 The challenges of Big EO Data and traditional AI approaches . . . . .    | 24        |
| 1.1.3 QC as a disruptive technology . . . . .                                  | 24        |
| 1.2 Problem Statement . . . . .                                                | 25        |
| 1.2.1 Limitations of classical ML methods for EO . . . . .                     | 25        |
| 1.2.2 Quantum Computing: Paradigm, Limitations, and Potential for EO . . . . . | 26        |
| 1.3 Objectives of the Dissertation . . . . .                                   | 27        |
| 1.3.1 Scientific questions addressed . . . . .                                 | 27        |
| 1.3.2 Expected contributions . . . . .                                         | 27        |
| 1.4 Research Methodology and Approach . . . . .                                | 28        |
| 1.4.1 Integration of Classical and Quantum Methods . . . . .                   | 28        |
| 1.4.2 Experimental Validation on EO Datasets . . . . .                         | 28        |
| 1.4.3 Benchmarking and comparative analysis . . . . .                          | 28        |
| 1.5 Structure of the Dissertation . . . . .                                    | 28        |
| 1.5.1 Outline of Chapters and Logical Flow . . . . .                           | 28        |
| Bibliography . . . . .                                                         | 30        |
| <b>2 Quantum Machine Learning for Earth Observation</b>                        | <b>31</b> |
| 2.1 Introduction . . . . .                                                     | 31        |
| 2.2 QML: Foundations and Key Concepts . . . . .                                | 32        |
| 2.2.1 Quantum Information Principles . . . . .                                 | 34        |
| 2.2.2 From Classical to Quantum Machine Learning . . . . .                     | 36        |
| 2.2.3 Data Encoding . . . . .                                                  | 36        |
| 2.2.4 Quantum Machine Learning Paradigms . . . . .                             | 37        |
| 2.2.5 Training Quantum Models . . . . .                                        | 38        |
| 2.3 Literature review on QML4EO . . . . .                                      | 38        |
| 2.3.1 Review Methodology . . . . .                                             | 38        |
| 2.3.2 Collected Contributions . . . . .                                        | 40        |
| 2.4 Application Domains and Research Directions in QML4EO . . . . .            | 47        |
| 2.4.1 Emerging Patterns in EO Sensors and Tasks . . . . .                      | 47        |
| 2.4.2 Common Quantum Approaches and Technological Trends . . . . .             | 48        |
| 2.4.3 Emerging Use Cases and Research Directions . . . . .                     | 49        |

---

|          |                                                                                          |           |
|----------|------------------------------------------------------------------------------------------|-----------|
| 2.4.4    | Geographic and Institutional Ecosystem . . . . .                                         | 50        |
| 2.5      | Discussion . . . . .                                                                     | 50        |
| 2.5.1    | Methodological Weak Spots . . . . .                                                      | 51        |
| 2.5.2    | Hardware Reality Check . . . . .                                                         | 51        |
| 2.6      | Conclusions and Future Prospects . . . . .                                               | 51        |
|          | Bibliography . . . . .                                                                   | 53        |
| <b>3</b> | <b>Quanv4EO: Empowering Earth Observation by means of Quanvolutional Neural Networks</b> | <b>61</b> |
| 3.1      | Introduction . . . . .                                                                   | 61        |
| 3.2      | Related works on Quanvolution . . . . .                                                  | 63        |
| 3.3      | Datasets . . . . .                                                                       | 64        |
| 3.4      | Methodology . . . . .                                                                    | 68        |
| 3.4.1    | Encoding of Input Image Patches . . . . .                                                | 69        |
| 3.4.2    | Lazy Training Regime . . . . .                                                           | 70        |
| 3.4.3    | Parallelization and Optimization . . . . .                                               | 70        |
| 3.4.4    | Processing Time Analysis . . . . .                                                       | 72        |
| 3.4.5    | Feature Maps Consistency . . . . .                                                       | 73        |
| 3.5      | Experimental Results . . . . .                                                           | 75        |
| 3.5.1    | Framework and Model Configurations . . . . .                                             | 75        |
| 3.5.2    | Results on Data Science datasets . . . . .                                               | 76        |
| 3.5.3    | Results on EO datasets . . . . .                                                         | 77        |
| 3.6      | Discussion and Conclusions . . . . .                                                     | 84        |
|          | Bibliography . . . . .                                                                   | 86        |
| <b>4</b> | <b>Quantum Graph Neural Networks</b>                                                     | <b>91</b> |
| 4.1      | Introduction . . . . .                                                                   | 91        |
| 4.2      | Preliminaries on GNNs . . . . .                                                          | 93        |
| 4.2.1    | Definitions and Notations . . . . .                                                      | 93        |
| 4.2.2    | Architecture of GNNs . . . . .                                                           | 93        |
| 4.3      | QGNNs Models . . . . .                                                                   | 95        |
| 4.3.1    | Fully and Hybrid Quantum . . . . .                                                       | 95        |
| 4.3.2    | Quantum Assisted GNNs . . . . .                                                          | 100       |
| 4.4      | Applications and Potential for Quantum Advantage . . . . .                               | 101       |
| 4.4.1    | Earth Observation . . . . .                                                              | 101       |
| 4.4.2    | High-Energy Physics . . . . .                                                            | 103       |
| 4.4.3    | Molecular Chemistry and Biology . . . . .                                                | 105       |
| 4.4.4    | Complex Systems . . . . .                                                                | 107       |
| 4.4.5    | Quantum Mechanics . . . . .                                                              | 108       |
| 4.5      | Challenges and Limitations . . . . .                                                     | 109       |
| 4.5.1    | Noise and Decoherence . . . . .                                                          | 110       |
| 4.5.2    | Scalability and Adaptability . . . . .                                                   | 111       |
| 4.5.3    | Lack of Performance Guarantees . . . . .                                                 | 111       |
| 4.5.4    | Barren Plateaus . . . . .                                                                | 112       |
| 4.5.5    | Initialization Strategies . . . . .                                                      | 112       |
| 4.6      | Conclusion . . . . .                                                                     | 113       |
|          | Bibliography . . . . .                                                                   | 114       |

---



---

|          |                                                                  |            |
|----------|------------------------------------------------------------------|------------|
| <b>5</b> | <b>Quantum Diffusion Models</b>                                  | <b>125</b> |
| 5.1      | Introduction . . . . .                                           | 125        |
| 5.2      | Background . . . . .                                             | 127        |
| 5.2.1    | Fundamentals of Diffusion Models . . . . .                       | 127        |
| 5.3      | Related Works . . . . .                                          | 128        |
| 5.3.1    | Generative Adversarial Networks for EO Image Synthesis . . . . . | 128        |
| 5.3.2    | Diffusion Models in Remote Sensing . . . . .                     | 128        |
| 5.3.3    | Quantum Generative Models in Machine Learning . . . . .          | 129        |
| 5.4      | Methodology . . . . .                                            | 129        |
| 5.4.1    | Quantum Latent Diffusion Model . . . . .                         | 130        |
| 5.4.2    | QVCU-Net . . . . .                                               | 131        |
| 5.4.3    | QCU-Net . . . . .                                                | 132        |
| 5.4.4    | Variational Quantum Circuits and Ansatz Design . . . . .         | 133        |
| 5.4.5    | Conditioning Evaluation Protocol . . . . .                       | 134        |
| 5.5      | Experimental Setup . . . . .                                     | 134        |
| 5.5.1    | Software, libraries and training configurations . . . . .        | 135        |
| 5.5.2    | Hyperparameters tuning and models' size . . . . .                | 136        |
| 5.5.3    | Evaluation Metrics . . . . .                                     | 137        |
| 5.5.4    | EuroSAT dataset . . . . .                                        | 137        |
| 5.6      | Results . . . . .                                                | 138        |
| 5.6.1    | Main results . . . . .                                           | 138        |
| 5.6.2    | Ablation Study . . . . .                                         | 143        |
| 5.6.3    | Additional studies . . . . .                                     | 143        |
| 5.7      | Conclusion . . . . .                                             | 145        |
|          | Bibliography . . . . .                                           | 147        |
| <b>6</b> | <b>Future Prospects and Conclusions</b>                          | <b>153</b> |
| 6.1      | Mapping the findings to the Research Questions . . . . .         | 153        |
| 6.2      | Main Outcomes . . . . .                                          | 155        |
| 6.2.1    | From field overview to actionable gaps . . . . .                 | 155        |
| 6.2.2    | Hybrid architectures are the practical path . . . . .            | 155        |
| 6.2.3    | Where Does This Leave "Quantum Advantage"? . . . . .             | 155        |
| 6.3      | Current Limitations . . . . .                                    | 156        |
| 6.3.1    | Encoding and information bottlenecks . . . . .                   | 156        |
| 6.3.2    | Scalability and the simulation-to-hardware gap . . . . .         | 156        |
| 6.3.3    | Benchmarking and Reproducibility . . . . .                       | 157        |
| 6.4      | Future Prospects and Research Directions . . . . .               | 157        |
| 6.5      | Final Conclusions . . . . .                                      | 158        |

---

# List of Figures

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |    |
|-----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 1.1 | <b>Visual representation of the research context.</b> . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 26 |
| 2.1 | Number of papers on QML4EO over the last 8 years, distributed across different journals. Please note that data for 2025 represent an estimate and include publications publicly available up to June 15, 2025. Nevertheless, a steady positive growth of publications in the field is expected, e.g., following major events such as LPS 2025 and IGARSS 2025, and from papers submitted in 2025 and currently under review. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                    | 33 |
| 2.2 | High level scheme representing the process of training a quantum machine learning model using a classical optimizer. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 38 |
| 2.3 | Geographical distribution of the contributions identified in our study, based on the affiliation of the first author. The map highlights a strong European presence, with Germany, Italy, and Poland leading in the number of contributions. The right-hand box provides a detailed numerical breakdown by country, illustrating the global interest in this emerging research area. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                            | 39 |
| 2.4 | Classification of quantum computing libraries and platforms based on their quantum-classical nature, implementation status (simulated or deployed), and associated hardware (e.g., IBM, D-Wave, Rigetti). Each marker represents a specific software-hardware combination; the number displayed inside the marker indicates the number of reviewed works adopting that configuration, while the marker area is scaled proportionally to the same count to provide an immediate visual indication of concentration density. The figure also highlights a strong bias toward platforms with unknown or unspecified quantum hardware, suggesting a current emphasis on algorithm development rather than hardware-specific implementations. Jitter has been applied to the $(x, y)$ scatter coordinates to improve readability and reduce overlap. . . . . | 41 |
| 2.5 | Identified quantum machine learning methodologies in the field of earth observation. From left to right: 1) Hybrid quantum neural networks, divided into A) early hybrid scheme, where QC is added internally to an AI module; B) late hybrid scheme, where QC is applied to AI-reduce set of feature; C) fully quantum scheme, where no classical AI is involved, which represent the aimed solution in the future; 2) Quantum Kernels; 3) Quantum Annealing; 4) Quantum Encoding Methods. . . . .                                                                                                                                                                                                                                                                                                                                                     | 42 |
| 2.6 | Distribution of QML applications across EO sensor modalities (left) and task types (right). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 47 |
| 3.1 | Input study area depicting the entire urban area of Tunis, acquired through Sentinel-1 SAR data. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 65 |

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |     |
|------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 3.2  | Area of Interest (AOI) selection and turbidity distribution. The left panel shows the satellite image of the selected AOI, with the red bounding box marking the analysis area. The right panel displays the corresponding turbidity map from Copernicus HR-OC data, where colors indicate turbidity levels (low values in blue, high values in red). . . . .                                                                                                                                                                                                                  | 67  |
| 3.3  | Quanvolution Scheme, where a quantum kernel of $2 \times 2$ qubits is shown (top). Features extracted with the quantum kernel on Sentinel-1 and Sentinel-2 images (bottom). . . . .                                                                                                                                                                                                                                                                                                                                                                                            | 68  |
| 3.4  | Quanvolution Image Processing Time by varying size of the input image, stride, number of qubits, kernel size, and number of output channels. The top left plot shows how timing is influenced by the number of qubits, image size, and stride. The top right shows how timing is influenced by the number of qubits, image size, and kernel size. The bottom plot shows how timing is influenced by the number of qubits, image size, and number of output feature maps. . . . .                                                                                               | 72  |
| 3.5  | Quanvolution processing time with a configuration using 8/16 qubits, a $2 \times 2$ kernel size, and 4 output feature maps. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                            | 73  |
| 3.6  | Quanvolutional feature maps generated using a random quantum circuit with 16 qubits, for different kernel sizes ( $4 \times 4$ , $3 \times 3$ , and $2 \times 2$ ). . . . .                                                                                                                                                                                                                                                                                                                                                                                                    | 74  |
| 3.7  | Quanvolutional feature maps obtained from a random quantum circuit with 9 qubits, shown for kernel sizes $3 \times 3$ and $2 \times 2$ . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                               | 74  |
| 3.8  | Quanvolutional feature maps produced by a custom quantum circuit with 16 qubits, combining a highly entangling layer with a random layer, for kernel sizes $3 \times 3$ and $2 \times 2$ . . . . .                                                                                                                                                                                                                                                                                                                                                                             | 74  |
| 3.9  | Comparison between quanvolutional and classical convolutional feature maps (same spatial scale). Despite the absence of training, quanvolutional features exhibit richer and more structured representations than classical convolutional features. . . . .                                                                                                                                                                                                                                                                                                                    | 74  |
| 3.10 | Features clustering with Quanvolution (top) and with CNN (bottom) - Binary Case (considered as example Annual Crop and River classes of EuroSAT). . . . .                                                                                                                                                                                                                                                                                                                                                                                                                      | 79  |
| 3.11 | Features clustering with Quanvolution (top) and with CNN (bottom) - Multiclass Case . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | 79  |
| 3.12 | Comparison of quanvolved feature maps, prediction, and corresponding GT for a random sample from the test set. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                         | 83  |
| 4.1  | Number of publications on QGNNs over the years. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | 92  |
| 4.2  | Illustration of GNN and QGNN framework. In the classical path, graph data is encoded, optimized, and trained to produce the output. In the quantum path, quantum encoding is applied, followed by circuit optimization (purely quantum or hybrid), and measurements yield the output. . . . .                                                                                                                                                                                                                                                                                  | 94  |
| 4.3  | Proposed taxonomy for QGNNs. QGNN approaches are categorized into two primary groups: Fully/Hybrid Quantum and Quantum Assisted. The Fully/Hybrid Quantum group includes subcategories such as Quantum Graph Neural Networks, Recurrent, Convolutional, Time Series Convolutional, and Equivariant models, with references to key works in the field. The Quantum Assisted group highlights methods leveraging quantum-classical hybrid architectures for specific applications. Each category is annotated with relevant citations to provide context and background. . . . . | 95  |
| 4.4  | Distribution of QGNN categories across three main areas: (a) application domains of the analyzed works, including Quantum Hamiltonian Dynamics (QHD), finance, complex systems, Molecular Chemistry and Biology (MC&B), HEP, and other categories; (b) types of approaches used, distinguishing between hybrid, fully quantum, and quantum-assisted methods; (c) simulation platforms utilized, such as PennyLane, Qulacs, OriginQ, Cirq, Qiskit, TensorFlowQ, and unspecified cases. . . . .                                                                                  | 102 |

---

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |     |
|-----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 5.1 | Architecture of the proposed quantum latent diffusion model. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 130 |
| 5.2 | Class-conditioned <i>Quantum Diffusion Model</i> . <b>Top:</b> forward chain $q(x_t x_{t-1})$ (green) corrupts $x_0$ into standard Gaussian noise $x_T$ ; the learned reverse $p_\theta(x_{t-1} x_t, y)$ (blue) denoises step-wise under label $y$ . <b>Bottom:</b> UNet denoiser with classical Conv/ResNet/Attention blocks plus two quantum residual variants: <i>QuanResNet</i> (red) replaces the <i>first</i> conv by a quanvolution, while <i>QResNet</i> (purple) at the bottleneck replaces <i>both</i> convs with a VQC acting on $2 \times 2 \times 3$ features, applied to a fraction $\rho$ of channels. . . . . | 132 |
| 5.3 | Hybrid quantum–classical ResNet. <i>Top:</i> comparison between the classical ResNet block and the QuanResNet block, in which one convolutional layer is replaced by a Quanvolutional (quantum) layer. <i>Bottom:</i> schematic of the HQConvAnsatz variational quantum circuit: angle encoding of feature maps, parameterized rotations, entangling gates, and measurements to generate quantum-enhanced features. . . . .                                                                                                                                                                                                   | 135 |
| 5.4 | FID scores over training epochs for different generative models. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 139 |
| 5.5 | Qualitative comparison between GT and generated images on the EuroSAT dataset. Columns correspond (left to right) to the ten land-cover classes: Annual Crop, Forest, Herbaceous Vegetation, Highway, Industrial, Pasture, Permanent Crop, Residential, River, and Sea Lake. Rows show GT reference images (first row), outputs of the classical model (second row), and outputs of the quantum models (QVCU-Net and QCU-Net) (third and fourth rows). . . . .                                                                                                                                                                | 140 |
| 5.6 | Convergence trends over training iterations for different models and evaluation metrics (FID, KID, IS, and Accuracy). Each plot shows the evolution of performance across iterations for the classical baseline, the QVCU-Net and the proposed QCU-Net. In addition to FID (shown in the main text), similar trends are evident for the other metrics. Notably, qCU-Net not only achieves superior scores but also demonstrates a significantly faster convergence compared to the other approaches. . . . .                                                                                                                  | 141 |
| 5.7 | Per-channel intensity distributions (densities; common bins in $[0, 1]$ ) for <i>AnnualCrop</i> (top) and <i>River</i> (bottom). Vertical dashed and dotted lines denote mean and median. We compare the distribution for the GT, Classic model and the QCU-Net model. . . . .                                                                                                                                                                                                                                                                                                                                                | 142 |

---

---

# List of Tables

|      |                                                                                                                                                                                                    |     |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 2.1  | Example of single-qubit gate operations: $X$ , $Y$ , $Z$ , $H$ , and rotations $R_x(\theta)$ , $R_y(\theta)$ , $R_z(\theta)$                                                                       | 35  |
| 2.2  | Representative QML applications in Earth Observation (EO). . . . .                                                                                                                                 | 48  |
| 3.1  | Processing time for Quanvolution . . . . .                                                                                                                                                         | 71  |
| 3.2  | Comparisons with SOTA on MNIST data classification. (Classical Model rank) . . . . .                                                                                                               | 77  |
| 3.3  | Comparisons with SOTA on Fashion MNIST data classification. (Classical Model rank) . . . . .                                                                                                       | 77  |
| 3.4  | Results of BIRCH Clustering for QuanvNN and CNN feature maps.<br>. . . . .                                                                                                                         | 78  |
| 3.5  | Results of AutoML applied to QCNN and CNN feature maps.<br>. . . . .                                                                                                                               | 80  |
| 3.6  | Results of AutoDL applied to QCNN and CNN feature maps.<br>. . . . .                                                                                                                               | 80  |
| 3.7  | Comparisons with SOTA on EuroSAT dataset classification . . . . .                                                                                                                                  | 81  |
| 3.8  | Simulation results for different quantum circuit configurations and classical baseline (AI4EDoET) . . . . .                                                                                        | 82  |
| 3.9  | Proposed model's average scores on the testing dataset: (a) GT, (b) Input with speckle, (c) Sebastianelli et Al. [22] and (d) Proposed <b>QSpeckleFilter</b> .<br>. . . . .                        | 83  |
| 3.10 | Performance comparison between the standard Attention U-Net and the Quantum-assisted Attention U-Net using different quantum configurations. . . . .                                               | 84  |
| 3.11 | Performance comparison of models (5 epochs). The kernel size value is the same for all the models and it is set to $3 \times 3$ . The number of qubits for the quantum models is set to 9. . . . . | 84  |
| 4.1  | Comparison of GNNs and QGNNs. . . . .                                                                                                                                                              | 96  |
| 4.2  | Summary of Applications and Potential Quantum Advantages in QGNNs. . . . .                                                                                                                         | 103 |
| 4.3  | Summary of Challenges, Descriptions, and Potential Solutions in QGNNs. . . . .                                                                                                                     | 110 |
| 5.1  | List of models's size in terms of trainable parameters and the delta in percentage with respect to the classical model (hybrid configurations). . . . .                                            | 136 |
| 5.2  | List of models's size in terms of trainable parameters and the delta in percentage with respect to the classical model (latent configurations). . . . .                                            | 136 |
| 5.3  | Comparison of generative models by parameters, ansatz, and metrics (FID, KID, IS). . . . .                                                                                                         | 138 |
| 5.4  | Comparisons of best classical and quantum models after 60,000 iterations on the EuroSAT dataset . . . . .                                                                                          | 140 |
| 5.5  | Comparison of performance metrics for each class across Classical Model, QVCU-Net, and QCU-Net. Best values per row are highlighted in bold. . . . .                                               | 141 |
| 5.6  | Ablation study on position analysis for QCU-Net after 60,000 iterations on the EuroSAT dataset . . . . .                                                                                           | 143 |

---

|     |                                                                                                                                                     |     |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------|-----|
| 5.7 | Ablation study on ansatz choice for QCU-Net after 60,000 iterations on the EuroSAT dataset                                                          | 143 |
| 5.8 | Per-class comparison of precision, recall, and F1-score between only rotation with no entanglement (N) and HQConv (E) circuits for QCU-Net. . . . . | 144 |
| 5.9 | Results with hyperparameters tuning for the classical baseline. . . . .                                                                             | 144 |



---

---

# Acronyms

**AI** Artificial Intelligence.

**AIS** Automatic Identification System.

**AN** Autonomous Navigation.

**AOI** Area Of Interest.

**ATR** Automatic Target Recognition.

**AUC** Area Under the ROC Curv.

**AutoDL** Automated Deep Learning.

**AutoML** Automated Machine Learning.

**B2B** Band-to-Band.

**BIRCH** Balanced Iterative Reducing and Clustering using Hierarchies.

**BO** Bayesian Optimization.

**BOA** Bottom-Of-Atmosphere.

**BRACS** BReAst Cancer Subtyping.

**BRISK** Binary Robust Invariant Scalable Keypoints.

**CAM** Channel Attention Module.

**CCA** Cloud Cover Assessment.

**CCD** Charge-Coupled Device.

**CDA** Crater Detection Algorithm.

**CG** Coarse Georeferencing.

**CLAHE** Contrast Limited Adaptive Histogram Equalization.

**CMA** Crater Matching Algorithm.

**CMV** Curved-Moving Vessel.

**CNN** Convolutional Neural Network.

---

**COG** Speed Over Ground.

**CONIDA** Peruvian Space Agency.

**ConvGNN** Convolutional Graph Neural Network.

**COTS** Commercial Off-The-Shelf.

**CPU** Central Processing Unit.

**CSC** Coarse Spatial Coregistration.

**CV** Cross-Validation.

**DARTS** Differentiable Architecture Search.

**DDPM** Denoising Diffusion Probabilistic Model.

**DEM** Digital Elevation Map.

**DL** Deep Learning.

**DM** Diffusion Model.

**DN** Digital Number.

**DNN** Deep Neural Network.

**E2E** End-to-End.

**ED** Euclidean Distance.

**EDU-QGCs** Equivariantly Diagonalizable Unitary Quantum Graph Circuits.

**EGNN** Equivariant Graph Neural Network.

**EHQGCs** Equivariant Hamiltonian Quantum Graph Circuits.

**EKF** Extended Kalman Filter.

**EO** Earth Observation.

**EPnP** Efficient Perspective-n-Point.

**EQGCNN** Equivariant Quantum Graph Convolutional Neural Network.

**EQGNN** Equivariant Quantum Graph Neural Network.

**ESA** European Space Agency.

**FC** Fully-Connected.

**FFT** Fast Fourier Transform.

**FID** Fréchet Inception Distance.

**FLOAT16** 16-bit Floating Point.

---

**FOV** Field Of View.

**FPGA** Field Programmable Gate Array.

**FPGAs** Field Programmable Gate Arrays.

**FPN** Feature Pyramid Network.

**FQConv** Flat Quantum Convolutional.

**GAE** Graph Autoencoder.

**GAN** Generative Adversial Network.

**GCL** Graph Convolutional Layer.

**GCNN** Graph COnvolutional Neural Network.

**GEE** Google Earth Engine.

**GHZ** Greenberger–Horne–Zeilinger.

**GIS** Geographic Information System.

**GNN** Graph Neural Network.

**GPS** Global Positioning System.

**GPU** Graphics Processing Unit.

**GPUs** Graphics Processing Units.

**GRD** Ground Range Detected.

**GRSS** Geoscience and Remote Sensing Society.

**GT** Ground Truth.

**GTs** Graph Transformers.

**HEP** High-Energy Physics.

**HL-LHC** High-Luminosity Large Hadron Collider.

**HOG** Histogram of Oriented Gradients.

**HPC** High Performance Computing.

**HQConv** Hierarchical Quantum Convolutional.

**HQGCNN** Hybrid Quantum Graph Convolutional Neural Network.

**HQGNN** Hybrid Quantum-Graph Neural Network.

**HR-OC** High-Resolution Ocean Colour.

**HSI** Hyperspectral Image.

---

**ICA** Independent Component Analysis.

**IEEE** Institute of Electrical and Electronics Engineers.

**IMO** International Maritime Organization.

**IoT** Internet of Things.

**IoU** Intersection over Union.

**IS** Inception Score.

**ISA** Israeli Space Agency.

**IUU** Illegal, Unreported, and Unregulated.

**KID** Kernel Inception Distance.

**KNN** K-Nearest Neighbors.

**L0** Level-0.

**L1A** Level-1A.

**L1B** Level-1B.

**L1C** Level-1C.

**L2A** Level-2A.

**LCLF** Lunar Centered Lunar Fixed.

**LDM** Latent Diffusion Model.

**LEO** Low Earth Orbit.

**LHC** Large Hadron Collider.

**LOLA** Lunar Orbiter Laser Altimeter.

**LRO** Lunar Reconnaissance Orbiter.

**LROC** Lunar Reconnaissance Orbiter Camera.

**LULC** Land-Use-Land-Cover.

**LUT** Look-Up Table.

**MCC** Matthews Correlation Coefficient.

**MDA** Maritime Domain Awareness.

**MF** Multi-Frequency.

**ML** Machine Learning.

**MLP** Multilayer Perceptron.

---

**MM** Multi-Mission.

**MMD** Maximum Mean Discrepancy.

**MMSI** Maritime Mobile Service Identity.

**MPA** Marine Protected Area.

**MPGNN** Message Passing Graph Neural Network.

**MPS** Matrix Product State.

**MSA** Maritime Situational Awareness.

**MSI** MultiSpectral Instrument.

**MSSWD** Multi-Spectral Ship Wake Dataset.

**MTF** Modulation Transfer Function.

**NAS** Neural Architecture Search.

**NCS** Neural Compute Stick.

**NIR** Near-InfraRed.

**NISQ** Noisy Intermediate-Scale Quantum.

**NN** Neural Network.

**NTK** Neural Tangent Kernel.

**OA** Overall Accuracy.

**OD** Object Detection.

**ONI** Ocean Niño Index.

**OPNAV** Optical Navigation.

**ORB** Oriented FAST and Rotated BRIEF.

**PCA** Principal Component Analysis.

**PCC** Pearson Correlation Coefficient.

**PEA** Position Estimation Algorithm.

**PQC** Parameterized Quantum Circuit.

**PSF** Point Spread Function.

**PSNR** Peak Signal-to-Noise Ratio.

**PTQ** Post-Training Quantization.

**PyRawS** Python for RAW Sentinel-2 Data.

---

**QA** Quantum Annealing.

**QAGNN** Quantum-Assisted Graph Neural Network.

**QAOA** Quantum Approximate Optimization Algorithm.

**QAT** Quantization-Aware Training.

**QC** Quantum Computing.

**QCNN** Quantum Convolutional Neural Network.

**QConvGNNs** Quantum Convolutional Graph Neural Networks.

**QCU-Net** Quanvolutional Conditioned U-Net.

**QDM** Quantum Diffusion Model.

**QFT** Quantum Fourier Transform.

**QGAN** Quantum Generative Adversarial Network.

**QGDM** Quantum Generative Diffusion Model.

**QGNN** Quantum Graph Neural Network.

**QLDM** Quantum Latent Diffusion Model.

**QLSTM** Quantum Long Short-Term Memory Algorithm.

**QML** Quantum Machine Learning.

**QML4EO** Quantum Machine Learning for Earth Observation.

**QNN** Quantum Neural Network.

**QRAM** Quantum Random Access Memory.

**QRecGNNs** Quantum Recurrent Graph Neural Networks.

**QSGCN** Quantum Spatial Graph Convolutional Neural Network.

**QSVM** Quantum Support Vector Machine.

**QuanyNN** Quanvolutional Neural Network.

**QUBO** Quadratic Unconstrained Binary Optimization.

**QUEST** Quantum Earth Science and Technology.

**QVCU-Net** Quantum Vertex Conditioned U-Net.

**RADC** Residual-Attention-Dilated-Convolution.

**RANSAC** Random Sample Consensus.

**RCS** Radar Cross-Section.



---

**RecGNN** Recurrent Graph Neural Network.

**RMSE** Root Mean Squared Error.

**RNN** Recurrent Neural Network.

**ROI** Region of Interest.

**RS** Remote Sensing.

**S-2** Sentinel-2.

**SAR** Synthetic Aperture Radar.

**SDGs** Sustainable Development Goals.

**SIFT** Scale-Invariant Feature Transform.

**SNR** Signal-to-Noise Ratio.

**SOTA** State-of-the-Art.

**SQFL** slimmable quantum federated learning.

**SQNN** Slimmable quantum neural network.

**SSIM** Structural Similarity Index Measure.

**STGNN** Spatio Temporal Graph Neural Network.

**SVM** Support Vector Machine.

**TC** Technical Committee.

**TCAS** Traffic Collision Avoidance System.

**TRL** Technology Readiness Level.

**TS-QGCNN** Quantum Time Series Convolutional Graph Neural Network.

**TTN** Tree Tensor Network.

**ViT** Vision Transformers.

**VQA** Variational Quantum Algorithm.

**VQC** Variational Quantum Circuit.

**VQE** Variational Quantum Eigensolver.

---

# Chapter 1

## Introduction

This chapter outlines the motivation, research context, and primary objectives of this dissertation. It begins by discussing the pivotal role of Earth Observation (EO) in science and society, followed by an analysis of the scalability challenges posed by big EO data. Subsequently, Quantum Computing (QC) is introduced as a disruptive technology with the potential to overcome these computational barriers.

### 1.1 Motivation and Research Context

EO has become indispensable for climate monitoring, urban planning, and natural resource management. The increasing availability of data from satellite missions has created unprecedented opportunities for scientific advancement and policy support. However, the exponential growth of data poses significant challenges regarding scalability and efficient analysis.

#### 1.1.1 The growing importance of EO in science and society

EO plays an increasingly critical role in both scientific inquiry and societal development. Driven by the rapid expansion of satellite missions, ground-based monitoring networks, and the convergence with advanced data analytics and Artificial Intelligence (AI), EO has evolved from a specialized scientific tool into a fundamental asset for understanding and managing our planet. In parallel, the rise of dense satellite constellations and the advent of the "New Space" economy have significantly increased the spatial, spectral, and temporal granularity of observations. This evolution turns EO into a data-rich infrastructure for continuous global monitoring, capable of capturing rapid environmental changes at unprecedented scales.

From a research perspective, EO enables the study of major Earth system processes, such as land-use change, atmospheric dynamics, ocean circulation, biodiversity patterns, urban expansion and the monitoring of built environments [1]. By assimilating EO data into computational models and Digital Twins of the Earth, scientists can effectively track climate change, natural hazards and resource consumption [2]. For example, multi-decadal satellite time series enable consistent assessments of phenomena such as deforestation, glacier retreat, and surface water dynamics, providing evidence of anthropogenic impacts on the biosphere.

However, the impact of EO extends beyond academia; it is essential for policy-making, economic planning, and daily operations. Governments and international institutions increasingly rely on EO-derived information to design environmental policies, monitor progress toward Sustainable Development Goals (SDGs), and evaluate climate adaptation and mitigation strategies [3]. At the global level, EO products increasingly support evidence-based reporting frameworks and the verification of environmental

indicators used in international climate and sustainability agendas.

In practical terms, EO data are vital for disaster management, providing near real-time mapping of floods, wildfires, and earthquakes to guide emergency response efforts [4]. Moreover, commercial sectors are leveraging these data for precision agriculture to optimize food security, renewable energy forecasting, infrastructure monitoring and maritime surveillance, including ship detection and tracking for logistics, environmental monitoring, and maritime security. Furthermore, the field is undergoing a democratization process. Initiatives such as the European Copernicus programme and NASA's open-data policy have made vast amounts of EO data freely available [5]. This broader access, coupled with cloud-based computing platforms that negate the need for local data storage, has expanded the user base well beyond EO specialists. This has amplified the demand for automated, scalable, and trustworthy analytics that can translate raw observations into actionable information.

Consequently, EO is no longer merely a scientific instrument; it has become a critical component of the global knowledge infrastructure. It supports both fundamental discovery and tangible societal needs, from sustainability and resilience to informed decision-making in an era of climatic uncertainty.

### **1.1.2 The challenges of Big EO Data and traditional AI approaches**

With numerous Remote Sensing (RS) missions currently operational, carrying multispectral, hyperspectral, and radar sensors, the volume of data generated is immense. The volume of data transmitted and stored from this continuously growing stream constitutes over 150 terabytes per day [6]. This rapid data growth has pushed researchers to seek suitable AI-based techniques for data compression, retrieval, and large-scale interpretation.

Despite remarkable progress, traditional AI pipelines face persistent bottlenecks when deployed at planetary scale. Training and maintaining high-performing models often requires substantial computational resources and energy, while inference over global archives remains costly due to the size and heterogeneity of EO data. In addition, many operational applications demand near real-time processing, robust generalization across sensors and geographies, and reliable uncertainty estimates, requirements that are still challenging to meet with purely classical approaches.

These constraints motivate the exploration of alternative computing paradigms. As classical hardware scaling slows and the computational burden of modern models continues to grow, Quantum AI applied to EO is attracting significant interest for its potential to address these computational bottlenecks.

### **1.1.3 QC as a disruptive technology**

QC is widely recognized as a disruptive technological force capable of reshaping scientific inquiry, industrial innovation, and societal infrastructures. Unlike classical computing, which encodes information in binary states of 0 and 1, QC leverages the principles of quantum mechanics, most notably superposition, entanglement, and quantum interference, to represent and manipulate information in fundamentally new ways. This paradigm shift allows quantum processors to explore exponentially large state spaces and solve classes of problems that are intractable or inefficient for classical high-performance systems [7].

The disruptive nature of QC stems from its potential to surpass the limits of Moore's Law and traditional scaling. While classical architectures are approaching physical and economic barriers, quantum platforms promise unprecedented computational power for optimization, simulation, and Machine Learning (ML) tasks. For instance, a 300-qubit system can theoretically represent more states than the number of particles in the observable universe, opening opportunities for novel algorithmic approaches [7].

Space agencies and research organizations are actively exploring how QC can accelerate breakthroughs in strategic domains. The European Space Agency (ESA) emphasizes that quantum technologies, in-

cluding communication, sensing, and computing, will serve as the basis for the next generation of space infrastructure. In particular, the integration of QC with EO and Digital Twin Earth initiatives is viewed as a pathway to process vast satellite data streams, optimize climate models, and enable real-time decision support for sustainability challenges [8, 9]. By coupling QC with AI, ESA envisions hybrid pipelines that enhance prediction accuracy and resource efficiency in applications such as climate monitoring, disaster management, and planetary exploration.

Nevertheless, the trajectory of QC is not without obstacles. Current devices, often referred to as Noisy Intermediate-Scale Quantum (NISQ) systems, are limited by decoherence, high error rates, and restricted qubit counts. This technological immaturity implies that, in the medium term, QC will likely function as an accelerator within hybrid architectures where classical CPUs and GPUs collaborate with quantum processors [7]. Yet, even at this nascent stage, the prospect of rethinking algorithmic design, data encoding, and problem formulation around quantum principles represents a profound paradigm shift.

## 1.2 Problem Statement

### 1.2.1 Limitations of classical ML methods for EO

In recent years, State-of-the-Art (SOTA) ML techniques have significantly advanced the field of EO. Deep Learning (DL) models, particularly Transformer-based architectures and large pre-trained foundation models, have achieved remarkable results across key tasks such as land-cover mapping, scene classification, and change detection [10, 11]. These models, exemplified by Vision Transformers (ViT) and novel state-space networks like Mamba, demonstrate an unprecedented ability to capture complex spatial-temporal patterns and leverage massive datasets [12]. Such advancements have markedly improved the extraction of semantic information from the exponentially growing volume of satellite data.

However, despite these successes, the scalability of classical ML approaches faces fundamental scalability challenges when applied to planetary-scale EO data. A primary bottleneck lies in the computational complexity of modern architectures; for instance, the self-attention mechanism in Transformers scales quadratically ( $O(N^2)$ ) with sequence length. When applied to high-resolution or hyperspectral satellite imagery, which translates into massive token sequences, this complexity results in prohibitive memory and computational costs, effectively hindering the processing of large-scale scenes in reasonable timeframes [13].

Furthermore, the challenge is not merely computational but also mathematical. EO data is inherently high-dimensional and multi-modal, often characterized by complex, non-linear correlations that are difficult for classical algorithms to disentangle efficiently. Conventional optimization methods, such as stochastic gradient descent, frequently struggle with the non-convex loss landscapes of deep networks, risking convergence to suboptimal local minima. Additionally, training massive foundation models to capture these distributions demands immense energy resources, raising significant concerns regarding the sustainability and cost-effectiveness of purely classical deployment [14].

These limitations suggest that classical computing paradigms may be approaching practical efficiency limits regarding their efficiency in modeling complex Earth systems. To overcome the "curse of dimensionality" and the computational bottlenecks of current DL, it is necessary to explore fundamentally different processing frameworks. This context motivates the investigation of QC, which offers a novel mathematical basis for handling high-dimensional data and optimization problems.

### 1.2.2 Quantum Computing: Paradigm, Limitations, and Potential for EO

The computational and mathematical challenges encountered by classical ML have catalyzed interest in QC as a transformative paradigm. While classical computers process information in binary bits, strictly limited to two states (0 or 1), QC leverages the fundamental principles of quantum mechanics, specifically superposition and entanglement, to operate on qubits. This shift allows a quantum processor to represent and manipulate an exponential number of states ( $2^N$  for  $N$  qubits) simultaneously. Conceptually, this architecture addresses the "curse of dimensionality" inherent in EO data by naturally operating in high-dimensional Hilbert spaces, potentially offering exponential speedups for specific linear algebra operations and optimization tasks that are intractable for classical solvers. These ideas have naturally led to Quantum Machine Learning (QML), a research area that explores how QC can be integrated into ML pipelines to improve data representations and learning efficiency.

However, realizing this theoretical advantage in practice presents significant challenges. We are currently in the NISQ era, a term coined to describe the current generation of quantum hardware [15]. NISQ devices are characterized by a limited number of qubits and a lack of fault tolerance; they are highly susceptible to environmental noise and decoherence, which limits the depth of the circuits that can be executed reliably. Unlike ideal fault-tolerant quantum computers, which require massive error correction overheads not yet available, current processors introduce noise that can distort the results of computations, making the direct application of complex, deep algorithms difficult.




Figure 1.1: Visual representation of the research context. The dissertation is situated at the intersection of EO characterized by big data challenges and diverse societal applications like climate monitoring and urban planning and QC, leveraging principles such as superposition and entanglement. The overlapping area defines the specific focus on QML methods applied to EO tasks.

Despite these hardware constraints, the application of QC to EO remains a compelling frontier. The specific nature of EO problems, such as the need to classify multi-spectral pixels or optimize resource allocation in satellite constellations, often aligns with the capabilities of variational quantum algorithms, which are designed to be resilient to moderate levels of noise. By exploiting the vast expressivity of the quantum state space, QC offers a novel approach to feature extraction and kernel estimation, promising to reveal complex correlation patterns in satellite imagery that classical kernels might miss. Thus, while full quantum advantage remains a long-term goal, the exploration of hybrid quantum-classical protocols

offers a pragmatic pathway to enhance current EO capabilities.

## 1.3 Objectives of the Dissertation

The primary aim of this research is to bridge the gap between classical RS analytics and the emerging capabilities of QC. To visually outline the scope of this work, Figure 1.1 illustrates the convergence of these domains, defining the specific intersection of QML applied to EO tasks that forms the focus of this dissertation.

### 1.3.1 Scientific questions addressed

This dissertation aims to investigate the potential and limitations of QML methods when applied to EO data. Specifically, the study is guided by the following research questions:

- **RQ1:** In which contexts do QML models provide a tangible advantage over traditional DL approaches? This question seeks to identify problem domains, data characteristics, or computational settings where quantum-enhanced methods outperform or complement classical solutions.
- **RQ2:** How can hybrid quantum–classical architectures be effectively integrated into EO tasks? This addresses the design of model pipelines that leverage quantum circuits alongside classical components, ensuring scalability and robustness given current hardware limitations.
- **RQ3:** Can quantum-enhanced feature extraction improve the representational power for EO data? This investigates whether mapping data into high-dimensional quantum Hilbert spaces (e.g., via quanvolutional layers) enables the capture of spatial-spectral patterns that are challenging for classical kernels.
- **RQ4:** To what extent can quantum probabilistic models enhance the generation of synthetic EO data? This focuses on assessing whether the inherent probabilistic nature of quantum mechanics offers superior capabilities in modeling complex data distributions for tasks such as data augmentation.

### 1.3.2 Expected contributions

This work aims to provide the following contributions to the field of Quantum Machine Learning for Earth Observation (QML4EO):

- **Novel architectures:** Development and evaluation of new QML-based models tailored to EO tasks, including quanvolutional layers, hybrid quantum–classical networks, and Quantum Diffusion Model (QDM).
- **Benchmarking on EO datasets:** Comprehensive testing on existing real-world EO datasets and newly created benchmark datasets designed specifically for the proposed methods.
- **Critical analysis:** A rigorous assessment of the added value of quantum-enhanced approaches compared to SOTA classical ML/DL methods, focusing on efficiency, scalability, and generalization capabilities.
- **Future perspectives:** A forward-looking discussion on the role of QML4EO, highlighting opportunities, open challenges, and potential research directions.

## 1.4 Research Methodology and Approach

### 1.4.1 Integration of Classical and Quantum Methods

A central component of this dissertation is the exploration of hybrid approaches that integrate classical DL architectures with quantum layers. This design choice is motivated by the current technological landscape, characterized by NISQ devices. Since fully quantum models remain largely impractical due to constraints in qubit count and fidelity, hybrid methods provide a pragmatic pathway to leverage quantum advantages in realistic scenarios.

In this framework, quantum layers are embedded as modular components within classical pipelines. These layers typically consist of Parameterized Quantum Circuit (PQC), which encode classical data into quantum states to perform transformations that are challenging to replicate classically. The outputs, once measured, are processed by classical neural networks handling large-scale learning tasks.

The integration strategy is driven by two main factors:

1. **Architectural Variety:** Various hybrid architectures are investigated, including quantum convolutional layers for feature extraction and hybrid quantum-classical diffusion models for generative tasks.
2. **Strategic Enhancement:** The goal is not to replace classical DL entirely, but to strategically enhance it with quantum modules that introduce novel representational capabilities. This reflects a balanced perspective: acknowledging present-day hardware constraints while actively testing potential contributions in scientifically relevant applications.

### 1.4.2 Experimental Validation on EO Datasets

Experimental validation is carried out on a combination of existing benchmark datasets and newly developed resources tailored for EO. This dual strategy ensures comparability with established baselines while maintaining relevance to real-world scenarios.

All experiments are assessed using standard quantitative metrics. Classification tasks employ accuracy, precision, and recall. Generative and data augmentation performance is measured via Fréchet Inception Distance (FID), Kernel Inception Distance (KID), and Inception Score (IS). For image filtering tasks, pixel-wise quality is evaluated using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM).

### 1.4.3 Benchmarking and comparative analysis

QML models are systematically benchmarked against well-established classical baselines across a diverse set of EO tasks, including Land-Use-Land-Cover (LULC) classification, building segmentation, water quality estimation, speckle filtering, and El Niño forecasting. This benchmarking serves a dual purpose: it provides a rigorous evaluation of relative performance and highlights specific strengths and limitations regarding accuracy, interpretability, and efficiency.

## 1.5 Structure of the Dissertation

### 1.5.1 Outline of Chapters and Logical Flow

This dissertation is organized into six chapters, creating a narrative that progresses from context to specific contributions, and finally to broader reflections.



- **Chapter 1: Introduction.** This chapter outlines the research motivation, objectives, and scientific questions, positioning the study within the academic and applied contexts of QC for EO.
- **Chapter 2: Background.** A comprehensive review of the SOTA is provided, covering EO, QC, and QML4EO, establishing the necessary theoretical foundations.
- **Chapters 3, 4, 5: Methodological Contributions.** These chapters present the core methodological advances, including quanvolutional approaches, quantum graph methods, and QDMs. The rationale, design, and implementation are detailed, followed by a systematic comparison against classical baselines to highlight the added value for EO tasks.
- **Chapter 6: Future Directions and Conclusions.** This chapter discusses open challenges and future research avenues, exploring potential extensions and operational integrations. It concludes by summarizing findings and discussing the broader implications for the fields of QC and EO.

## Bibliography

- [1] G. B. Bonan and S. C. Doney, “Climate, ecosystems, and planetary futures: The challenge to predict life in Earth system models,” *Science*, vol. 359, no. 6375, p. eaam8328, 2018.
- [2] S. Liang and J. Wang, *Advanced remote sensing: terrestrial information extraction and applications*. Academic Press, 2019.
- [3] J. Holloway and K. Mengersen, “Statistical machine learning methods and remote sensing for sustainable development goals: A review,” *Remote Sensing*, vol. 10, no. 9, p. 1365, 2018.
- [4] G. Santilli, C. Vendittozzi, C. Cappelletti, S. Battistini, and P. Gessini, “CubeSat constellations for disaster management in remote areas,” *Acta Astronautica*, vol. 145, pp. 11–17, 2018.
- [5] R. Ramachandran, K. Bugbee, and K. Murphy, “From open data to open science,” *Earth and Space Science*, vol. 8, no. 5, p. e2020EA001562, 2021.
- [6] A. Sebastianelli, D. A. Zaidenberg, D. Spiller, B. L. Saux, and S. L. Ullo, “On Circuit-Based Hybrid Quantum Neural Networks for Remote Sensing Imagery Classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 565–580, 2022.
- [7] E. W. Bethel, M. G. Amankwah, J. Balewski, R. Van Beeumen, D. Camps, D. Huang, and T. Perciano, “Quantum computing and visualization: A disruptive technological change ahead,” *IEEE Computer Graphics and Applications*, vol. 43, no. 6, pp. 101–111, 2023.
- [8] European Space Agency, “Quantum technologies,” [https://www.esa.int/Enabling\\_Support/Space\\_Engineering\\_Technology/Space\\_Optoelectronics/Quantum\\_Technologies](https://www.esa.int/Enabling_Support/Space_Engineering_Technology/Space_Optoelectronics/Quantum_Technologies), 2023, accessed: 2025-09-26.
- [9] —, “Digital twin earth: quantum computing and ai take centre stage at esa’s -week,” [https://www.esa.int/Applications/Observing\\_the\\_Earth/Digital\\_Twin\\_Earth\\_quantum\\_computing\\_and\\_AI\\_take\\_centre\\_stage\\_at\\_ESA\\_s\\_Ph-week](https://www.esa.int/Applications/Observing_the_Earth/Digital_Twin_Earth_quantum_computing_and_AI_take_centre_stage_at_ESA_s_Ph-week), 2025, accessed: 2025-09-26.
- [10] A. Xiao, W. Xuan, J. Wang, J. Huang, D. Tao, S. Lu, and N. Yokoya, “Foundation models for remote sensing and earth observation: A survey,” *IEEE Geoscience and Remote Sensing Magazine*, 2025.
- [11] R. Wang, L. Ma, G. He, B. A. Johnson, Z. Yan, M. Chang, and Y. Liang, “Transformers for remote sensing: A systematic review and analysis,” *Sensors*, vol. 24, no. 11, p. 3495, 2024.
- [12] X. Ma, X. Zhang, and M.-O. Pun, “Rs 3 mamba: Visual state space model for remote sensing image semantic segmentation,” *IEEE Geoscience and Remote Sensing Letters*, vol. 21, pp. 1–5, 2024.
- [13] T. Shinohara, “ViT-Koop: Vision-Transformer-Koopman Operators for Efficient Time-Series Forecasting of Earth-Observation Data,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 2814–2823.
- [14] Z. Huang, H. Yan, Q. Zhan, S. Yang, M. Zhang, C. Zhang, Y. Lei, Z. Liu, Q. Liu, and Y. Wang, “A survey on remote sensing foundation models: From vision to multimodality,” *arXiv preprint arXiv:2503.22081*.
- [15] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.

## Chapter 2

# Quantum Machine Learning for Earth Observation

Traditional ML and DL models face major challenges in EO due to the increasing volume and complexity of data. These challenges include computational intensity, energy consumption, and data management constraints. QML emerges as a promising paradigm for overcoming these barriers by harnessing quantum phenomena such as superposition, entanglement, and interference. In fact, QML is an emerging interdisciplinary field that leverages QC principles to address complex computational challenges. This chapter aims at offering a detailed analysis of the state-of-the-art and new trends in QML application to EO. This chapter methodically examines important advancements in the field of QML4EO (including publicly available publications up to June 15, 2025), by exploring the fundamental ideas of QML. The methodology involved systematic searches across prominent scientific databases, using carefully formulated queries to ensure broad coverage and high relevance. In addition, the changing institutional and geographic environment of QML4EO research is addressed, highlighting major centers of contribution and the role of initiatives such as the Institute of Electrical and Electronics Engineers (IEEE) Geoscience and Remote Sensing Society (GRSS) Quantum Earth Science and Technology (QUEST) Technical Committee (TC). In order to transform QML from a specialized research area into a complementary and essential tool for advancing EO, this chapter critically analyzes current limitations and provides an outlook on future directions, highlighting the need for reliable hardware, improved algorithmic design, and standardized protocols [1].

## 2.1 Introduction

The development of AI has profoundly affected EO with its ability to extract meaningful information from a growing amount of data [2]. However, the integration of big data with AI-based EO is not without challenges. Despite advances brought forward by algorithms ML and especially DL in analyzing data EO, large-scale ML can be time-consuming and energy-intensive due to inherent limitations related to data management, processing, and effective resource utilization [3]. These limitations can be mitigated by improving computational capability and efficiency. QC has emerged as a pioneering solution that leverages principles of quantum physics such as superposition, entanglement and interference to perform calculations [4].

Interest in implementing ML methodologies based on QC led to the creation of a new research field known as QML [5, 6]. Initial studies show that quantum-enhanced models can potentially explore larger solution spaces, leading to faster convergence, better generalization, and improved performance, especially for large and complex datasets [7].

The question of whether QC technologies and methods are compatible with EO tasks naturally arises. Interest in the field has been growing, with research primarily focusing on hybrid quantum-classical model architectures. This is motivated by the current capabilities of quantum hardware, which are limited in size and coherence time. Although significant, these advances only partially address the broader challenges and opportunities in deploying quantum-enhanced models in EO [8].

This research direction, summarized as QML4EO, has only emerged significantly in 2017 and is therefore relatively new. The total number of contributions is low, but has increased markedly in the last two years, as shown in Fig. 2.1.

This phenomenon is partially motivated by the fact that we are seeing a similar sharp increase in pure QC research, suggesting that significant breakthroughs may be close, and that we may be currently entering the «Quantum computers' "ChatGPT moment"»<sup>1</sup>. In contrast to more established research subjects that were yielding progressive results, there is a noticeable recent spike in research across various EO sub-fields, indicating a growing interest within the community towards this innovative and transformative computing paradigm. Notwithstanding the relatively short history, this field of research is motivated by the extraordinary potential impacts that this technology could have on EO.

This chapter aims at offering a less technical and more accessible perspective on this niche topic. Although existing reviews [9, 10] provide valuable information, they have limited scope and do not cover the latest advances in QML4EO in a comprehensive way. This chapter seeks to bridge this gap by presenting a broader, more inclusive, up-to-date overview, highlighting both opportunities and challenges. By doing so, we hope to foster interdisciplinary collaboration and drive further research in this emerging field.

The chapter is structured as follows:

- In Sect. 2.2, the fundamental aspects of quantum information and QML are introduced.
- In Sect. 2.3, an exhaustive summary of the QML4EO literature is included, highlighting the chosen review methodology and the different classes of methods.
- In Sect. 2.4, information on application domains, trends, research ecosystems is extrapolated from the QML4EO literature.
- In Sect. 2.5 and 2.6, general discussions and conclusions are drawn on the review, and future prospects are proposed.

## 2.2 QML: Foundations and Key Concepts

With the emergence of NISQ devices, which operate with limited quantum resources and without full error correction, QML has gained traction as a near-term strategy for solving problems in optimization, pattern recognition, and time-sensitive signal analysis. Preskill et al. [11] highlights how such devices may soon outperform classical computers in specific tasks, despite their inherent noise, and identifies hybrid quantum-classical algorithms (e.g, Quantum Approximate Optimization Algorithm (QAOA), Variational Quantum Eigensolver (VQE)) as particularly promising for early applications. In this context, QML becomes especially relevant for EO, where quantum-enhanced models could improve the classification of hyperspectral imagery, tracking of dynamic climate variables, and modeling of complex geophysical processes [5, 6].

While a comprehensive treatment of quantum mechanics is beyond the scope of this thesis (refer to standard texts such as Rieffel & Polak [12] and Nielsen & Chuang [4] for a complete introduction),

---

<sup>1</sup>«Quantum computers' "ChatGPT moment"» <https://www.bcg.com/capabilities/digital-technology-data/emerging-technologies/expert-insights/jean-francois-bobier>




Figure 2.1: Number of papers on QML4EO over the last 8 years, distributed across different journals. Please note that data for 2025 represent an estimate and include publications publicly available up to June 15, 2025. Nevertheless, a steady positive growth of publications in the field is expected, e.g., following major events such as LPS 2025 and IGARSS 2025, and from papers submitted in 2025 and currently under review.

here, a foundational overview of the concepts strictly necessary to understand how quantum algorithms manipulate data is provided.

### 2.2.1 Quantum Information Principles

Quantum states are represented as vectors in a complex vector space, more specifically a Hilbert space. A common notation is Dirac's bra-ket notation:  $|\psi\rangle$  (pronounced "ket psi") denotes a column vector representing a quantum state, while  $\langle\psi|$  (pronounced "bra psi") denotes its conjugate transpose (a row vector). The inner product between two states is written as  $\langle\phi|\psi\rangle$ , which yields a complex number. At the heart of quantum computing lies the qubit (quantum bit), a two-level quantum system that generalizes the classical bit. For a qubit, the Hilbert space is two-dimensional, spanned by the computational basis states  $|0\rangle$  and  $|1\rangle$ , analogous to the standard basis vectors in  $\mathbb{C}^2$ . Unlike a classical bit that exists in either 0 or 1, a qubit can exist in a superposition of states  $|0\rangle$  and  $|1\rangle$ :

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \quad \text{where } |\alpha|^2 + |\beta|^2 = 1. \quad (2.1)$$

This state resides on the surface of the Bloch sphere<sup>2</sup>, a geometrical representation where each point encodes a possible pure state. The manipulation of qubits is governed by unitary operations, which preserve the unit norm of the state vector. Geometrically, unitary operations correspond to rotations and reflections on the Bloch sphere, maintaining the constraint that the qubit state remains normalized. This normalization becomes essential once we consider measurement, where the amplitudes of a state acquire a direct probabilistic interpretation.

*Quantum measurement* is a process that extracts classical information from quantum states. When measuring a quantum state, the outcome is probabilistic, with probabilities determined by the Born rule:  $P(\text{outcome}) = |\langle\text{outcome}|\psi\rangle|^2$ . An *observable* is a measurable quantity (such as position or energy) represented by a Hermitian operator. Upon measurement, the quantum state is projected onto one of the observable's eigenstates, and the corresponding eigenvalue is observed. For example, the Pauli-Z operator  $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$  has eigenstates  $|0\rangle$  and  $|1\rangle$ . When measuring the  $Z$  observable on a qubit  $|\psi\rangle$ , the outcomes correspond to the eigenstates  $|0\rangle$  and  $|1\rangle$  occurring with probabilities  $|\alpha|^2$  and  $|\beta|^2$  respectively. QML algorithms estimate output probabilities and expectation values through repeated circuit executions and measurements. Each execution yields a single measurement outcome drawn from the quantum probability distribution, and statistical properties are estimated by aggregating results across many runs (shots).

Another key concept in quantum information is *entanglement*, which captures non-separable correlations between subsystems that cannot be reproduced by product states. In some settings, such correlations are incompatible with classical local models and can increase the expressive capacity of QML models when learning joint dependencies. For two qubits, there exist four maximally entangled Bell states. One example,  $|\Phi^+\rangle$ , is given by:

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \quad (2.2)$$

This specific state exhibits perfect correlation, measuring the first qubit always yields the same result as measuring the second qubit. This non-classical correlation cannot be explained by classical information. In EO applications, entanglement has the potential to enable quantum models to capture joint dependencies between multiple features (e.g., correlations between spectral bands, or correlation between different physical variables) more efficiently than classical representations.

*Quantum superposition* is the principle that a quantum system can exist in a linear combination of multiple basis states, each associated with a probability amplitude. *Interference* governs how these amplitudes combine, enabling algorithms to enhance selected outcomes and suppress others. Quantum amplitudes

<sup>2</sup><https://quantum.cloud.ibm.com/learning/it/courses/general-formulation-of-quantum-information/density-matrices/bloch-sphere>

can cancel (destructive interference) or reinforce (constructive interference) depending on their relative phases. In practice, this mechanism is exploited to concentrate measurement probability mass into desired results.

Quantum circuits consist of sequences of quantum gates, i.e., unitary transformations applied to qubits, to implement computations [13]. Unlike classical logic gates, quantum gates are reversible and operate on complex-valued amplitudes. A common distinction is between single-qubit and multi-qubit gates. Single-qubit gates (some examples in Tab. 2.1) act on individual qubits and are represented by  $2 \times 2$  unitary matrices. They manipulate the state by rotating it around the Bloch sphere. For example, the Hadamard gate ( $H$ ) creates superposition by mapping basis states to equal-weight combinations. Rotation gates  $R_x$ ,  $R_y$ , and  $R_z$  perform arbitrary-angle rotations around the respective axes, with  $\theta$  specifying the rotation angle.

Table 2.1: Example of single-qubit gate operations:  $X$ ,  $Y$ ,  $Z$ ,  $H$ , and rotations  $R_x(\theta)$ ,  $R_y(\theta)$ ,  $R_z(\theta)$

| Gate          | Matrix                                                                                                                                   | Description                                  |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------|
| $X$           | $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$                                                                                           | Bit flip (rotation about $x$ axis)           |
| $Y$           | $\begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}$                                                                                          | Bit and phase flip (rotation about $y$ axis) |
| $Z$           | $\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$                                                                                          | Phase flip (rotation about $z$ axis)         |
| $H$           | $\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$                                                                       | Creates equal superposition                  |
| $R_x(\theta)$ | $\begin{bmatrix} \cos(\frac{\theta}{2}) & -i \sin(\frac{\theta}{2}) \\ -i \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}$ | Rotation around $x$ axis by angle $\theta$   |
| $R_y(\theta)$ | $\begin{bmatrix} \cos(\frac{\theta}{2}) & -\sin(\frac{\theta}{2}) \\ \sin(\frac{\theta}{2}) & \cos(\frac{\theta}{2}) \end{bmatrix}$      | Rotation around $y$ axis by angle $\theta$   |
| $R_z(\theta)$ | $\begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}$                                                                  | Rotation around $z$ axis by angle $\theta$   |

For example, consider a qubit initially in state  $|0\rangle$ . Applying the Hadamard gate transforms it to  $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ , creating an equal superposition. A subsequent measurement yields  $|0\rangle$  or  $|1\rangle$  with 50% probability each. The  $R_z(\theta)$  gate rotates the qubit's phase; while this does not change measurement probabilities directly in the computational basis, it modifies the interference patterns in subsequent operations. These gates serve as building blocks for quantum algorithms.

Multi-qubit gates enable interactions between qubits and are essential for generating entanglement. The most common one is the controlled-NOT (CNOT) gate:

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad (2.3)$$

CNOT is a two-qubit operation where the second (target) qubit is flipped if and only if the first (control) qubit is in state  $|1\rangle$ :

$$\text{CNOT} |00\rangle = |00\rangle, \quad \text{CNOT} |10\rangle = |11\rangle \quad (2.4)$$

These gates can be combined in sequence to build a quantum circuit where arbitrary unitary transformations are applied on multi-qubit systems. A finite set of gates, such as  $\{R_x(\theta), R_z(\theta), \text{CNOT}\}$ , forms a universal gate set capable of approximating any unitary operation to arbitrary precision. Gate operations must also be designed with hardware constraints in mind, especially in the NISQ era. Constraints include number of qubits, gate error rates, and coherence time. Here, qubit coherence refers to the ability to maintain a well-defined relative (complex) phase between  $|0\rangle$  and  $|1\rangle$ , so the coherence time is the duration over which the relative phase remains stable, before noise causes errors.

Quantum information is also governed by no-cloning [14, 15] and no-deleting theorems [16], which set foundational constraints on data copying and erasure. These constraints impact the design of QML algorithms and limit redundancy mechanisms common in classical ML pipelines.

### 2.2.2 From Classical to Quantum Machine Learning

To understand the transition from classical to quantum ML, it is useful to compare how information is represented and processed. Classical ML algorithms typically operate on data encoded as vectors in real-valued spaces ( $\mathbb{R}^d$ ), transforming them through non-linear functions (e.g., neural network layers) to learn decision boundaries.

In contrast, QML encodes classical data into quantum states and processes it through quantum operations in a Hilbert space whose dimension grows as  $2^n$  with the number of qubits  $n$ . This can enable rich feature mappings and compact parameterizations, although practical advantages depend on the chosen encoding, circuit structure, and constraints related to data loading and measurement. The differences can be summarized in three key aspects:

- *Data representation:* A quantum state over  $n$  qubits is described by amplitudes in a  $2^n$ -dimensional Hilbert space. This provides a high-dimensional representation, but the accessible information is constrained by state preparation and measurement.
- *Feature spaces and kernels:* Quantum circuits can implement feature maps and kernel evaluations for certain data embeddings, potentially inducing highly non-linear decision boundaries in the original input space.
- *Optimization and training:* Many QML models are trained in hybrid quantum–classical loops (e.g., variational circuits optimized with classical routines). While interference shapes the measurement statistics, training can suffer from issues such as shot noise and barren plateaus.

### 2.2.3 Data Encoding

A fundamental prerequisite for QML is the mapping of classical EO data into quantum states, a process known as *quantum data encoding*. This operation defines how information is represented in the Hilbert space and directly affects the model’s expressive power, hardware efficiency, and resilience to noise.

Let  $\mathbf{x} \in \mathbb{R}^n$  represent a classical input vector (e.g., a pixel spectrum or spatial feature descriptor). A quantum encoding routine prepares a quantum state  $|\psi(\mathbf{x})\rangle$  such that the classical data is embedded into the amplitudes or parameters of the state. Formally:

$$|\psi(\mathbf{x})\rangle = U(\mathbf{x})|0\rangle^{\otimes n} \quad (2.5)$$

where  $U(\mathbf{x})$  is a unitary transformation that encodes  $\mathbf{x}$  into the quantum register. The main encoding schemes can be categorized as follow:



- *Angle Encoding*: Each feature  $x_i$  modulates a rotation gate  $R_x(x_i)$  /  $R_y(x_i)$  /  $R_z(x_i)$  applied to the  $i$ -th qubit. This is hardware-efficient and suitable for near-term devices:

$$|\psi(\mathbf{x})\rangle = \bigotimes_{i=1}^n R_y(x_i) |0\rangle \quad (2.6)$$

For instance, Sebastianelli et al. [17] used this method to encode optical data.

- *Amplitude Encoding*: The vector  $\mathbf{x}$  is normalized and used to directly specify the amplitudes of the quantum state:

$$|\psi(\mathbf{x})\rangle = \sum_{i=0}^{2^n-1} x_i |i\rangle, \quad \text{with } \sum_i |x_i|^2 = 1 \quad (2.7)$$

This allows exponential compression but requires complex state preparation, and it is typically less used when coming to hardware implementation.

- *Hybrid and Physics-Informed Embeddings*: EO-specific encoding schemes have been proposed to incorporate domain knowledge. For instance, Otgonbaatar et al. [18] introduced a method for encoding Stokes parameters from polarimetric SAR imagery into quantum gates, aligning the quantum representation with physical observables.

The choice of encoding impacts circuit depth, fidelity, and interpretability. For example, angle encoding is shallow and robust to noise, making it well-suited for NISQ devices, while amplitude encoding offers higher expressiveness at the cost of deeper circuits and more complex state preparation [19].

## 2.2.4 Quantum Machine Learning Paradigms

QML includes a variety of algorithmic paradigms tailored to leverage QC for learning. Here, a few notable examples are described.

*Variational Quantum Algorithms (VQAs)* utilize PQCs where parameters are iteratively optimized to minimize a classical objective function. This approach is well-suited to NISQ devices due to short circuit depth and hybrid computation. In EO, VQAs can be used for solving inverse problems such as deducing physical parameters from measurements:

$$\min_{\theta} \langle \psi(\theta) | \hat{H} | \psi(\theta) \rangle \quad (2.8)$$

where  $\hat{H}$  is an observable encoding the loss function, and  $|\psi(\theta)\rangle$  is the quantum state.

*Quantum Support Vector Machines (QSVMs)* extend classical Support Vector Machines (SVMs) by defining quantum kernels that evaluate inner products in high-dimensional feature spaces.

These kernels are computed by preparing quantum states corresponding to data points and evaluating their inner products through quantum circuits. While interference can enhance certain kernel implementations, the fundamental mechanism is the evaluation of inner products in a quantum Hilbert space, which can capture complex nonlinear relationships and offer computational gains when classical kernels become costly, such as in multispectral pixel classification [5].

*Quantum Principal Component Analysis (QPCA)* enables efficient extraction of dominant eigenmodes from quantum-encoded density matrices. This is particularly relevant in dimensionality reduction of EO datasets, e.g. identifying major variance modes in sea surface temperature data using quantum phase estimation [5].

*Quantum Neural Networks (QNNs)* emulate neural architectures using quantum gates with tunable parameters. Circuit-based QNNs exploit superposition and entanglement to model complex functions and




Figure 2.2: High level scheme representing the process of training a quantum machine learning model using a classical optimizer.

decision boundaries. For example, they can be trained to detect abnormal vegetation indices or rainfall patterns from spectral time series.

### 2.2.5 Training Quantum Models

Training QML models requires efficient and hardware-aware optimization techniques (high level scheme presented in Fig. 2.2).

A widely used method for gradient-based learning is the *parameter-shift rule*, which provides exact gradients for expectation values under certain conditions:

$$\frac{\partial \langle \hat{O} \rangle}{\partial \theta} = \frac{1}{2} [\langle \hat{O} \rangle_{\theta + \frac{\pi}{2}} - \langle \hat{O} \rangle_{\theta - \frac{\pi}{2}}] \quad (2.9)$$

This method allows tuning parameters using classical optimizers and supports compatibility with the gradient descent and the Adam algorithm, widely used in deep learning [20].

Nevertheless, QML optimization landscapes are known to suffer from *barren plateaus*, where gradients vanish exponentially as the number of qubits or circuit depth increases. These regions make training difficult and require circuit designs that preserve local structure. Mitigation strategies include careful initialization (layerwise, orthogonal), localized cost functions, and use of shallow, task-specific ansätze [21].

## 2.3 Literature review on QML4EO

### 2.3.1 Review Methodology

To ensure a comprehensive and unbiased review, it was employed a structured methodology to collect relevant works on QML applied to EO. The approach was designed to maximize coverage while maintaining high standards of relevance and quality. The key steps of our methodology are outlined as follows.

Systematic searches were conducted across multiple prominent scientific databases, including Web of Science, Scopus, IEEE Xplore, and arXiv. To enhance the completeness of this search, they were formulated queries incorporating a combination of key terms related to QML and EO. Boolean operators and




Figure 2.3: Geographical distribution of the contributions identified in our study, based on the affiliation of the first author. The map highlights a strong European presence, with Germany, Italy, and Poland leading in the number of contributions. The right-hand box provides a detailed numerical breakdown by country, illustrating the global interest in this emerging research area.

wildcards were used to refine search results and optimize the balance between precision and recall. An example of a search query used in our process is as follows:

*("Quantum Machine Learning" OR "Quantum Neural Networks") AND ("Earth Observation" OR "Remote Sensing" OR "Satellite Data").*

The search terms were iteratively refined based on the initial retrievals, and a set of inclusion and exclusion criteria was applied. Specifically, the following were included:

- Peer-reviewed journal articles that provide substantial contributions to QML in EO.
- Conference papers that introduce novel methodologies, frameworks, or case studies with sufficient technical depth.
- Preprints available on arXiv, recognizing that many cutting-edge studies in this rapidly evolving domain are yet to be formally published.

Conversely, the following were excluded:

- Poster presentations and unpublished manuscripts hosted solely on personal webpages or non-reputable repositories.
- Articles lacking sufficient technical depth, such as those providing only high-level discussions without quantitative results.
- Works that, despite mentioning QML and EO, did not present a clear methodological application or technical contributions within the domain.

To enhance transparency and reproducibility, a comprehensive list of all works included in this review was compiled. This list is publicly accessible through a dedicated webpage, awesome-qc4eo<sup>3</sup>, where

<sup>3</sup>Awesome QC4EO GitHub page: <https://alessandrosebastianelli.github.io/awesome-QC4EO/>

readers can explore the complete bibliography along with associated metadata. The webpage is hosted on GitHub, facilitating community engagement and fostering contributions from researchers and practitioners interested in expanding this evolving repository. By providing an open and interactive platform, the goal is to encourage collaboration and the continuous enrichment of resources in the field of QML4EO.

The outcome of this research led to the identification of 76 relevant publications. Considering the relatively early stage of this research area, this represents a significant milestone in the field. To gain further insight into the global landscape, the geographical distribution of contributions based on the affiliation of the first author was analyzed, as shown in Figure 2.3.

The analysis reveals a strong European presence, with Germany leading by a substantial margin (23 publications), followed by Italy (12) and Poland (9). Contributions from outside Europe include the United States (8), India (5), and Taiwan (4), indicating growing interest in the topic across multiple continents. Other countries with notable contributions include France (3), China and Switzerland (2), and several others with single entries, such as Canada, the Republic of Korea, Spain, and Greece. The full breakdown is shown in the right box in Figure 2.3. It is important to specify that they were collected and categorized papers until June 15, 2025, once publicly available.

### 2.3.2 Collected Contributions

Figure 2.4 provides a comprehensive overview of QC libraries and platforms, categorized by their compatibility with quantum hardware, implementation status (simulated or deployed), and degree of quantum-classical hybridization. The visualization maps major frameworks such as D-Wave Ocean, PennyLane, Qiskit, and TorchQuantum, indicating whether they are hybrid, fully quantum, or quantum-inspired, and whether they are tied to specific hardware backends like IBM, D-Wave, or Rigetti. Notably, the distribution reveals a strong bias toward platforms associated with unknown or unspecified quantum hardware, suggesting that most contributions are currently focused more on algorithmic development and simulation rather than direct hardware deployment. This highlights both the exploratory nature of the field and the current limitations in hardware accessibility, guiding researchers toward tools aligned with their computational goals.

Building on this landscape-level categorization, it is possible now to delve into the collected list of contributions according to the underlying QML model. This finer-grained analysis allows to assess how different quantum machine learning approaches are employed across various platforms and research efforts.

In order to introduce the identified methodologies, better detailed in the next sections, Fig. 2.5 illustrates a summary graphic showing their configuration.

### Quantum Neural Networks

Quantum Neural Networks (QNNs) are quantum algorithms that leverage QC to learn and extract features from input data for analysis.

**Hybrid Quantum Neural Network** This type of model combines quantum and classical computing components to perform machine learning tasks. Depending on the role and placement of the quantum component in the architecture, these hybrid models can be further broadly classified into two categories, here called Early Hybrid Schemes and Late Hybrid Schemes.

- *Early Hybrid Schemes:* These models mainly exploit QC to process input data through iterative encoding and analysis using QML techniques. A representative example is the quanvolutional neural network [22], which applies a PQC to small regions of the input to extract local features




Figure 2.4: Classification of quantum computing libraries and platforms based on their quantum-classical nature, implementation status (simulated or deployed), and associated hardware (e.g., IBM, D-Wave, Rigetti). Each marker represents a specific software-hardware combination; the number displayed inside the marker indicates the number of reviewed works adopting that configuration, while the marker area is scaled proportionally to the same count to provide an immediate visual indication of concentration density. The figure also highlights a strong bias toward platforms with unknown or unspecified quantum hardware, suggesting a current emphasis on algorithm development rather than hardware-specific implementations. Jitter has been applied to the  $(x, y)$  scatter coordinates to improve readability and reduce overlap.

for subsequent analysis. By systematically sliding the PQC across the input, this approach enables large-scale EO image processing. The resulting quantum-processed features are then fed into classical models for downstream tasks, offering a promising strategy for integrating QML into EO data analysis during the NISQ era.

Based on these principles, various models have been proposed and have demonstrated strong potential for EO data analysis. For example, Henderson et al. [23] and Ghosh et al. [24] developed specialized PQCs to extract enhanced features from multispectral imagery and radar sounder data, respectively. In the context of Synthetic Aperture Radar (SAR) data, Miller et al. [25] employed quanvolution to generate feature maps for classification tasks, while Mauro et al. [26] explored its potential for speckle noise reduction. Furthermore, Sebastianelli et al. [27] introduced a




Figure 2.5: Identified quantum machine learning methodologies in the field of earth observation. From left to right: 1) Hybrid quantum neural networks, divided into A) early hybrid scheme, where QC is added internally to an AI module; B) late hybrid scheme, where QC is applied to AI-reduce set of feature; C) fully quantum scheme, where no classical AI is involved, which represents the aimed solution in the future; 2) Quantum Kernels; 3) Quantum Annealing; 4) Quantum Encoding Methods.

quantum-based framework utilizing random circuits to process large volumes of diverse EO datasets, and proposed a parallelized quantum operation to improve computational efficiency. Building on the findings reported in [27], Mauro et al. [28] enhanced a framework for estimating turbidity levels with data from the  $\Phi$ Sat-2 mission, by applying quantum filters, as detailed in [29]. Departing from the conventional approach, Chalumuri et al. [30] developed a row-wise quantum operation to extract local features for EO data classification. Otgonbaatar et al. [18] applied a PQC for pixel-wise classification of PolSAR imagery, and Lin et al. [31] integrated pixel-level features derived from QML into a Graph Neural Network (GNN) architecture for the hyperspectral image change detection task. In contrast, some studies [32–36] have explored full-input quantum encoding for feature extraction. However, these approaches face significant practical limitations due to the constraints of current NISQ hardware. Current implementations typically require input downsampling to fit within available quantum resources, which can lead to the loss of critical spatial-spectral information. This fundamental trade-off poses a significant challenge, particularly in EO applications where the preservation of fine-grained details could be crucial for effective analysis. Furthermore, this principle has also been applied to explore the potential of QML for fusing multiple EO modalities in data analysis. For instance, Majji et al. [37] and Miller et al. [38] employed different PQCs to integrate SAR and optical imagery, with the fused features subsequently processed by classical deep learning models for classification.

Park et al. [39] propose a novel hybrid quantum federated learning framework tailored for satellite-ground communication. Leveraging slimmable quantum federated learning (SQFL) and Slimmable quantum neural network (SQNN)s with angle and pole configurations, their approach enhances communication and computing efficiency. Techniques like superposition coding and successive decoding further increase communication opportunities. Their extensive experiments demonstrate that the satellite-ground SQFL framework outperforms classical federated learning and conventional QFL methods in both computational and communicational aspects, marking a promising advance in distributed QML4EO applications.

- **Late Hybrid Schemes:** These hybrid models employ classical ML techniques to extract global feature

representations, which are then processed by QML components to enhance higher-level feature refinement for final predictions. As a result, the selection and design of classical feature reduction methods are critical to the overall performance of such hybrid architectures, particularly under the constraints of NISQ hardware.

To date, various classical algorithms have been integrated with PQC for EO tasks. Principal Component Analysis (PCA) is one typical algorithm that linearly transforms input data into a set of principal components. Gawron et al. [40] demonstrated the effectiveness of this framework for binary classification of multispectral images. Similar strategies have been extended to hyperspectral imagery: Priyanka et al. [41] implemented a QNN, and Shaik et al. [42] employed a Quantum Support Vector Machine (QSVM) for pixel-wise classification, both using PCA-based feature reduction.

In addition, DL, aiming to learn deep representations of data, has been extensively studied and has led to significant breakthroughs over the past decade, as reviewed in [43]. Utilizing such techniques in QML and combining them with QML has triggered great attention in research. For instance, Liliopoulos et al. [44] concentrated on Multilayer Perceptron (MLP)-based models that integrate quantum neurons in the hidden classical layers and achieved enhanced performance. Rainjonneau et al. [45] combined MLPs with a PQC as a quantum layer for satellite mission planning. Fan et al. [46] exploited MLP to assist EO data encoding for efficient QML analysis.

Besides traditional feedforward neural networks, Convolutional Neural Networks (CNNs) have proven particularly effective for EO imagery processing due to their ability to exploit spatial locality and perform hierarchical feature learning for efficient and robust feature extraction. Accordingly, many hybrid models, such as [47–49], incorporate CNN-based architectures to extract a set of informative features from EO data, and these features are then further processed using PQCs to perform the final classification with fully connected layers. In another study [50], Malarvanan et al. divided the features extracted by a CNN into four segments, each processed by a separate PQC. The outputs of these PQCs were then concatenated and passed through a fully connected layer for classification. Besides directly applying CNNs for feature reduction before feeding the data into PQCs, Zollner et al. [51] leveraged a deep convolutional network (VGG16) prior to the QML component for classification, and Chang et al. [52] utilized a convolutional autoencoder in their proposed model to reduce the input data dimensionality. The compressed feature representations were subsequently processed by a PQC for binary classification. In a study by Zollner et al. [53], several dimensionality reduction techniques were compared, and the results demonstrated the superiority of the autoencoder models in generating compact representations of EO imagery in comparison to traditional methods such as PCA. Otgonbaatar et al. [54] combined a convolutional autoencoder with a VGG16 to extract global features from multispectral EO images, which were then processed by a PQC for further classification. Such autoencoder-based hybrid models have been applied not only to classification tasks but also to EO data generation-related applications. For example, [55, 56] proposed to use QML for latent space transformations in hyperspectral image restoration. Ghosh et al. [57] introduced a hybrid model with a UNet-like architecture using a PQC to process the bottleneck features for semantic segmentation of radar sounder data.

Moreover, GNN architectures and attention mechanism have also gained popularity in EO applications, and their integration with QML has also been explored. For example, Mauro et al. [58] proposed combining a GNN with QML to predict the Ocean Niño Index. Papa et al. [59] examined the validity of integrating QML into ViTs for EO applications.

To better tackle the challenges in EO data analysis using DL models, various comprehensive training strategies have been explored and widely adopted in practice. These strategies have also been

integrated into hybrid QML models. One representative example is the use of pretrained models, which aim to enhance performance and generalization under limited annotated data via transfer learning. A study by Otgonbaatar et al. [60] demonstrated the effectiveness of this framework when applied to a small but high-dimensional annotated EO dataset. Chang et al. [61] used a pretrained autoencoder to map high-dimensional data into a latent space and proposed a hybrid classical-quantum GAN for an image generation task. In a related study, De Falco et al. [62] also exploited QML within the latent space and developed a quantum latent diffusion model integrated with a convolutional autoencoder for EO imagery generation. Besides that, ensemble learning has also gained attention in the development of hybrid models. For instance, Bhavsar et al. [63] employed multiple classical ML algorithms for robust feature reduction, followed by a PQC for prediction.

Van et al. [64] compare three distinct quantum neural network architectures for multispectral image classification using the EuroSAT dataset. Their study includes a fully quantum CNN, as well as hybrid and orthogonal variants, implemented via TensorFlow Quantum. The results highlight that both hybrid and fully quantum approaches offer improvements in classification performance, suggesting a growing relevance of QNNs for RS tasks as QC matures.

Glatting et al. [65] propose a cognitive SAR architecture for ship detection based on quantum reinforcement learning. The system utilizes two satellites in coordination: one with wide-swath, low-resolution scanning and another with high-resolution beam refocusing. Quantum variational circuits are employed within the reinforcement learning framework to dynamically adapt sensing strategies, highlighting the potential of QML for intelligent decision-making in EO scenarios.

A recent study by Otgonbaatar et al. [66] emphasizes that a true quantum advantage can only be expected if quantum models are able to generalize better than classical ones and avoid limitations such as weight symmetry during training. The authors highlight hyperspectral image classification as a particularly promising application domain due to its smaller input sizes and limited availability of training data, although current hardware constraints remain a bottleneck.

These studies collectively highlight the paradigm for feasibly processing EO data using QC in the NISQ era. However, whether QNNs can consistently outperform their classical counterparts remains an open question and requires further investigation.

**Fully Quantum Neural Network** This type of model employs QC exclusively for all stages of the learning process, including data encoding, feature extraction and prediction. Unlike hybrid approaches, it does not incorporate any classical components for processing and analyzing data.

Due to the limitations of quantum resources in the NISQ area, there are relatively few studies employing such models for analyzing EO data, which typically contain rich spatial and spectral information. Nevertheless, notable efforts have been made to explore the potential of QML in the EO domain under current constraints of quantum resources. For instance, Painchart et al. [67] introduced a QNN with a pyramid structure to reduce quantum resources for analyzing and classifying temporal sequence EO data. Yu et al. [68] introduced an evolutionary algorithm-driven method for quantum circuit design, aimed at minimizing quantum resource usage within an image similarity learning framework for effective representation learning. Pastori et al. [69] explore the use of QNNs to model cloud cover parameterizations in climate models, finding out that QNNs can capture non-linear relationships relevant to cloud dynamics and demonstrating potential for parameter reduction and improved generalization in constrained settings.



## Quantum Kernels

Recent research has increasingly explored quantum kernel methods as a promising approach to EO tasks, particularly in settings where classical kernel-based models like SVMs already perform well. In this context, quantum kernels refer to similarity measures computed using quantum circuits, where classical data is embedded into a quantum Hilbert space and inner products are estimated via quantum measurements. These kernels can, in theory, capture highly nonlinear structures in data using fewer resources than their classical counterparts, depending on the nature of the embedding.

Mirowski et al. [8] provide a broad overview of QML applications in RS, identifying both the potential of quantum kernels and the limitations posed by current hardware. A series of focused studies by the same group demonstrates practical implementations, notably in cloud detection using multispectral Sentinel-2 imagery [70, 71] and semantic segmentation using hyperspectral sensors [72]. Their results show that quantum kernel SVMs can match or slightly outperform classical counterparts on small-scale datasets, offering benefits in generalization and feature mapping. Shaik et al. [73] test simulated quantum kernels on a feature-reduced hyperspectral dataset, obtaining a 5% improvement in accuracy over RBF kernels for 12 features. Shaik et al. [42] extend the utility of quantum kernel methods into the semi-supervised learning regime, proposing a quantum-based pseudo-labelling approach for hyperspectral image classification that improves performance when labeled data is scarce. Sarkar et al. [74] apply quantum-enhanced SVMs to multiclass LULC classification, reporting encouraging results that point toward hybrid workflows as a practical direction. Glatting et al. [65] introduce quantum kernels for InSAR phase unwrapping, a novel application where phase continuity and optimization complexity make quantum approaches conceptually attractive. Meanwhile, Rodriguez-Grasa et al. [75] explore neural quantum kernels for satellite image classification, proposing a hybrid architecture that combines classical neural networks with quantum feature maps. Theoretical and practical considerations are deepened by Mirowski et al. [76], who analyze the impact of shot count on kernel estimation, highlighting the trade-offs between quantum fidelity and computational noise. Across these studies, a consistent theme emerges: quantum kernels are not yet outperforming classical baselines across the board. Nevertheless, they offer competitive accuracy, greater flexibility in non-linear mapping, and a testbed for exploring quantum advantage in real EO problems [77, 78].

## Quantum Data Encoding

Both QNNs and quantum kernels require classical data to be converted to the quantum domain, in a step called quantum data encoding. The choice of encoding affects not only the expressiveness and efficiency of the quantum model but also its compatibility with hardware constraints and the nature of the input data. Some works have put strong emphasis on this step, finding specific approaches for EO data.

Zollner et al. [53] explore different satellite image representation strategies for quantum classifiers, comparing encoding schemes such as amplitude encoding, angle encoding, and hybrid variants in terms of dimensionality, noise resilience, and classification accuracy. Their findings emphasize the trade-offs between circuit depth, interpretability, and performance, particularly in the context of image data with strong spatial structure. In a more specialized application, Otgonbaatar and Datcu [18] propose a natural quantum embedding of Stokes parameters derived from polarimetric SAR imagery. By directly mapping the physical quantities of polarimetric data into quantum gates, their method aligns with the physics of both the sensor and the quantum system, preserving important structural and statistical properties. Together, these works highlight the emerging focus on encoding as a design challenge, especially when adapting EO data, which often has high dimensionality and complex modality, to the constraints and potential of quantum models.

## Quantum Annealing

Quantum Annealing (QA) is a metaheuristic optimization technique that finds low-energy configurations in an energy landscape. The intrinsic suitability of QA for combinatorial optimization makes it appealing for EO problems, which often involve high-dimensional spaces, non-convex optimization, and discrete decision variables. Quantum annealers developed by D-Wave Systems are already being used to address real-world EO tasks.

Classification of RS data remains one of the most explored areas for QA applications, where optimization tasks appear. Boyda et al. [79] demonstrated an early proof-of-concept by detecting tree cover in aerial images using QA. Though limited by hardware constraints at the time, their study paved the way for later work on QA-enabled environmental monitoring. Otgonbaatar and Datcu [80] proposed a QA-based feature subset selection method for hyperspectral image classification. Their approach enhances classification accuracy by focusing on the most informative spectral bands, which also reduces the Quadratic Unconstrained Binary Optimization (QUBO) size and improves quantum resource efficiency.

Delilbasic et al. [81] developed a quantum-trained SVM method tailored to binary EO datasets, leveraging QA to solve the QUBO form of SVMs. Their results marked a significant step toward quantum-enhanced EO data classification. In later work, Delilbasic et al. [82] extended the quantum-trained SVM method to multiclass classification. This approach formulates the entire multiclass classification problem as a single optimization task, avoiding one-vs-all decomposition. Their method demonstrates competitive accuracy compared to classical approaches while reducing computational overhead. Then, Zardini et al. [83] proposed a framework for adding locality to quantum-trained binary and multiclass SVMs, demonstrating how restricting the training set to the  $k$  nearest neighbors can improve scalability and robustness under current hardware constraints. Delilbasic et al. [84] summarized advantages and limitations of SVM models trained with quantum annealing.

Regression tasks in EO, such as estimating biophysical variables, present unique challenges due to the continuous nature of outputs. Pasetto et al. [85] proposed a QA-based support vector regression model for estimating variables like leaf area index. Their implementation translates regression objectives into QUBO formulations, enabling deployment on existing quantum hardware. Pasetto et al. [86] also investigated kernel approximation techniques on a QA for EO regression tasks. By approximating non-linear kernels efficiently within the QUBO framework, the study highlights new avenues for leveraging QA in continuous prediction scenarios. Efficient data handling is critical for EO, especially with high-dimensional inputs. Otgonbaatar and Datcu [87] also tackled the coresets selection problem, aiming to build smaller yet informative datasets for downstream tasks. Their method leverages classical algorithms to select representative subsets of images, balancing fidelity with computational feasibility. They then train a QSVM using these subsets.

## Quantum-inspired Methods

For completeness, a few works based on quantum-inspired methods are also included, i.e., classical algorithms or architectures that draw conceptual or mathematical inspiration from quantum mechanics but are designed for conventional hardware. These approaches often aim to capture the advantages of quantum formalisms without the overhead and noise of real quantum devices. One typical model is the tensor network, an important concept in many-body quantum systems that provides an efficient representation of quantum states [88]. Matrix Product State (MPS) is one of the fundamental examples, widely used to represent entangled states in low-dimensional quantum systems. Inspired by this, Otgonbaatar et al. [89] applied tensor network techniques to reduce the number of parameters in a DL model and improve the spectral resolution of Hyperspectral Images (HSIs) by decomposing them into

sparse factor matrices.

In another example, motivated by the quantum theory, Zhang et al. [90] proposed a spectral-spatial network for HSI feature extraction. The network includes a phase-prediction module that assigns a phase for each input feature based on discriminative information, and a measurement-like fusion module that generates a low-dimensional feature representation for further analysis.

## 2.4 Application Domains and Research Directions in QML4EO

QML has emerged as a promising avenue for tackling the computational and data complexity challenges inherent in modern EO, particularly in processing high-dimensional, multi-source RS datasets. This section synthesizes the current landscape of QML applications in EO, highlighting prevalent data types, methodological trends, task categories, and institutional actors shaping the field.

### 2.4.1 Emerging Patterns in EO Sensors and Tasks

The conducted analysis reveals a concentration of QML applications around specific sensor modalities and EO task types, as illustrated in Figure 2.6.




Figure 2.6: Distribution of QML applications across EO sensor modalities (left) and task types (right).

More than 50% of the studies leverage multispectral imagery, followed by hyperspectral at 18.6% and SAR data at 13.6%. A smaller subset of works explores multi-sensor fusion or incorporates in-situ measurements.

Classification tasks dominate the field, accounting for nearly 62% of QML applications, these are further detailed in Section 2.4.2. Segmentation (16.7%) and object detection (8.3%) follow, commonly applied to urban mapping and ship detection. Regression, forecasting, and filtering remain underexplored but offer fertile ground for future developments, particularly in climate modeling and environmental monitoring.

### 2.4.2 Common Quantum Approaches and Technological Trends

Most QML models for EO rely on hybrid quantum-classical architectures tailored to the constraints of NISQ hardware. At the core of these systems are Variational Quantum Circuit (VQC)s, which enable low-depth, parameter-efficient learning via classical optimization.

Key trends include:

- VQCs paired with classical CNNs for land cover classification [48];
- Quantum kernel methods for hyperspectral pixel classification and cloud detection [70, 75];
- Quantvolutional layers integrated into classical deep networks for enhanced feature extraction from multispectral and SAR data [17, 52];
- Quantum-inspired tensor networks for dimensionality reduction in hyperspectral data cubes [89];

While classification is the dominant focus, a small but growing number of studies extend QML methods to structured prediction, regression, and data fusion tasks. A summary of representative QML4EO applications, including the data modalities and quantum methods used, is presented in Table 2.2.

Table 2.2: Representative QML applications in Earth Observation (EO).

| Application Area                        | EO Data Type        | Quantum Method                            | Reference                 |
|-----------------------------------------|---------------------|-------------------------------------------|---------------------------|
| Land cover classification               | Multispectral       | VQC + CNN                                 | Sebastianelli et al. [48] |
| Deforestation detection (time series)   | SAR (Sentinel-1)    | Quantum CNN + RNN                         | Painchart et al. [67]     |
| Cloud detection                         | Multispectral       | Quantum kernel                            | Miroszewski et al. [70]   |
| EO classification (EuroSAT, Sentinel-1) | Multispectral + SAR | Quantvolutional neural network (QuantvNN) | Sebastianelli et al. [17] |
| Semantic segmentation                   | Multispectral       | Quantum annealing                         | Delilbasic et al. [84]    |
| Multisource fusion for classification   | SAR + Optical       | Quantum-enhanced preprocessing            | Majji et al. [37]         |
| Feature learning                        | Multispectral       | Quantum contrastive similarity learning   | Yu et al. [68]            |
| Hyperspectral compression               | Hyperspectral       | Tensor networks                           | Otgonbaatar et al. [89]   |

#### Land Cover and Crop Type Classification

Land cover and crop type classification represent the most active areas of QML application in EO, often using multispectral or hyperspectral inputs.

VQCs have shown strong potential in these tasks, offering comparable performance to classical CNNs with significantly fewer parameters. Sebastianelli et al. [48], for example, matched CNN accuracy on the EuroSAT dataset using a compact hybrid quantum-classical model.

In parallel, quantum kernel methods have gained traction for their ability to embed input data into high-dimensional Hilbert spaces. Rodriguez et al. [75] introduced a neural quantum kernel classifier that outperformed classical SVM baselines in multispectral classification, demonstrating the value of quantum-aware feature encodings.

More recently, Sebastianelli et al. [17] proposed the Quantv4EO model, which incorporates quantvolutional layers into the preprocessing pipeline of hybrid quantum-classical architectures. Applied to multisource datasets such as EuroSAT and Sentinel-1, this model demonstrated improved classification accuracy (up to 5% gain) and 80% faster convergence during training compared to classical baselines, along with notable parameter efficiency.

#### Semantic Segmentation

Quantum-enhanced models are also emerging in structured output tasks such as segmentation and detection, where spatial context and fine-grained boundaries are essential.

Delilbasic et al. [84] investigated the use of quantum annealing for semantic segmentation using multispectral EO imagery. Their study highlighted both the potential and limitations of current quantum annealing hardware, emphasizing the importance of problem encoding and postprocessing for effective label assignment.

### **Geophysical Parameter Estimation**

QML has also been applied to continuous geophysical variable retrieval from satellite data, an inherently nonlinear inverse problem.

Painchart et al. [67] introduced a QNN to detect deforestation from temporal SAR sequences. Although framed as a classification task, their model demonstrated the feasibility of applying quantum convolutional and recurrent layers to extract dynamic features from time-series EO data.

### **Representation Learning for Feature Similarity**

Recent studies have explored QML's potential for learning efficient and transferable data representations.

Yu et al. [68] proposed a quantum contrastive learning approach based on unsupervised similarity triplets. Applied to multispectral imagery, the model achieved significant reduction in quantum resource usage while maintaining strong performance in downstream tasks such as land surface classification and deforestation detection.

### **Multimodal Data Fusion and Dimensionality Reduction**

As EO data becomes increasingly multimodal, QML offers novel strategies for data fusion and compression.

Majji et al. [37] proposed a quantum-inspired framework for enhancing SAR and optical image fusion. By applying a set of quantum image processing techniques, including quantum Fourier transform and Hadamard-based fusion operators, the authors demonstrated improved classification accuracy in land-use tasks, with performance gains over SAR-only inputs and standard fusion methods. Though the processing was simulated, the study illustrates the potential of quantum-inspired operations to improve EO data interpretability.

In parallel, Otgonbaatar et al. [89] employed tensor network decomposition techniques, originally developed in quantum many-body physics, for hyperspectral data compression. The approach achieved competitive classification accuracy while reducing input dimensionality, offering an efficient pathway for spectral feature extraction in high-dimensional EO datasets.

### **Generative Models and Data Augmentation**

Quantum generative models are still nascent but show promise for EO data synthesis and augmentation. Zhang et al. [90] developed a Quantum GAN (QGAN) to generate realistic hyperspectral image patches. Used for data augmentation, the synthetic samples improved downstream classification accuracy by 3–4%, underscoring their utility in low-sample regimes.

## **2.4.3 Emerging Use Cases and Research Directions**

Beyond pattern recognition, recent studies point to broader QML applications in EO, including:

- **Time-series modeling** with quantum recurrent architectures for change detection and climate dynamics;

- **Probabilistic inversion** using quantum-enhanced Bayesian inference for biophysical parameter retrieval;
- **Onboard processing** via lightweight quantum circuits for low-latency satellite edge computing.

These emerging directions highlight QML’s potential to advance EO well beyond classification and segmentation, into dynamic, probabilistic, and real-time domains.

#### 2.4.4 Geographic and Institutional Ecosystem

The field of QML4EO is supported by a growing and geographically diverse research community. After the earliest contributions that helped define the scope of this research area, several academic and institutional actors have played a key role in advancing methods, applications, and benchmarking frameworks. Notable contributors include:

- University of Sannio and University of Pavia, among the earliest contributors to QML4EO, with a focus on hybrid quantum neural networks applied to multispectral and SAR EO data;
- Jülich Supercomputing Centre, with contributions leveraging quantum and high-performance computing infrastructures for QML experimentation and benchmarking;
- The European Space Agency, in particular through the  $\Phi$ -lab, supporting exploratory research and applied studies at the intersection of QC and EO;
- Technical University of Munich (TUM) and the German Aerospace Center (DLR), contributing to foundational studies and benchmarking efforts in quantum-enhanced EO;
- CERN, providing advanced computational expertise and contributing to the exploration of quantum algorithms and infrastructures relevant to large-scale scientific data analysis;
- Asian academic institutions, primarily focusing on quantum-inspired approaches to hyperspectral image analysis;
- NASA-affiliated researchers and U.S. universities, exploring quantum annealing and optimization techniques for mission planning, scheduling, and resource allocation in EO contexts.

A significant step toward structural consolidation has been the establishment of the QUEST TC within the IEEE GRSS. QUEST plays a pivotal role in shaping the research agenda by promoting open-source software tools, organizing community benchmarks, and fostering interdisciplinary dialogue between EO scientists and QC researchers.

## 2.5 Discussion

The literature on QML4EO has grown exponentially since 2017, yet its centre of gravity remains narrow and preliminary. Figure 2.4 shows that about 73% of the papers analysed in this dissertation executed their quantum circuits only on classical simulators, and a further 15% combined simulation with a single “proof-of-concept” hardware run involving  $< 10$  qubits. Only 4% relied exclusively on physical devices and none of these exceeded 33 qubits. This heavy tilt towards simulated back-ends explains why reported accuracies sometimes match or slightly exceed classical CNN or SVM baselines: the experiments are effectively noise-free, unconstrained by qubit topology, and run with unlimited shot counts.

A second trend is the over-concentration on pixel-wise or patch-wise land-cover classification (60% of studies) and SAR/hyperspectral variants thereof. Tasks that arguably showcase the distinctive strengths of quantum kernels, e.g. data fusion across modalities, probabilistic inversion, or large-scale spatio-temporal forecasting, account for  $< 10\%$  of the corpus. This thematic imbalance limits the external validity of today’s “quantum advantage” claims: if the community keeps benchmarking on simplified, low-resolution or heavily down-sampled datasets, it will remain unclear whether any reported speed-ups survive contact with real EO workloads.

### 2.5.1 Methodological Weak Spots

Classical baselines are rarely optimised to parity. Several “quantum-wins” disappear when the comparator model is given the same parameter budget or when standard regularisation is applied. Worse, inference latency and energy for the cryogenic stack are never included in cost metrics, creating a distorted “qubit-seconds per prediction” narrative.

Only 48% of the paper surveyed publish code or circuit specifications. Different cloud providers still lack a common intermediate representation, making cross-platform replication difficult. Without transparent kernels and random seeds, it is impossible to separate algorithmic merit from hyper-parameter luck.

### 2.5.2 Hardware Reality Check

In late 2023, IBM announced the 1121-qubit “Condor” chip, and in 2025 a roadmap toward the 1386-qubit multi-chip “Kookaburra”. These numbers sound impressive until one realises that (i) usable qubit counts after error-mitigation are one to two orders of magnitude lower, and (ii) single- and two-qubit gate fidelities remain in the 99%–99.5% band, meaning error probabilities compound rapidly with circuit depth. Coherence times on today’s superconducting platforms are still in the order of  $100\mu\text{s} - 1\text{ms}$ , constraining practical circuits to a few hundred layers at best. Consequently, most QML4EO prototypes cap circuit depth at  $< 20$  parameterised gates, a regime where classical tensor-network simulators can still track the state vector efficiently.

## 2.6 Conclusions and Future Prospects

This chapter reviews the foundations, the current state of the art, and the emerging research ecosystem of QML4EO. Starting from key quantum information principles and the main learning paradigms (e.g., VQA, quantum kernels, and QNN), it is discussed how QML can address some of the computational and methodological bottlenecks of modern EO, while also acknowledging the constraints imposed by present-day NISQ hardware. The systematic literature review highlighted that the field is rapidly expanding, with a strong prevalence of hybrid quantum-classical solutions and a clear concentration on classification tasks using multispectral and hyperspectral data, followed by early attempts on segmentation, regression, filtering, and multimodal fusion.

Overall, three high-level insights emerge. First, most evidence to date suggests that QML is currently best positioned as a *complementary* component in EO pipelines rather than a full replacement of mature classical DL models: hybrid schemes, compact PQC, and carefully designed encodings represent the most practical path under today’s resource limits. Second, the choice of data encoding and circuit design is not a technical detail but a primary driver of performance, robustness, and scalability; in EO, physically meaningful or task-aware embeddings can become a key enabler for generalization under limited training data. Third, while claims of quantum advantage remain premature for most EO settings, the surveyed works collectively indicate that competitive accuracy, parameter efficiency, and promising

generalization behavior can be achieved in small-to-medium regimes, motivating continued investment in methodological rigor and reproducible evaluation.

Looking ahead, transforming QML4EO into a mature and widely adoptable research direction requires progress along four converging axes. (i) **Hardware readiness:** improved qubit counts, connectivity, coherence times, and error profiles are necessary to move beyond heavily downsampled inputs and toy-scale demonstrations. (ii) **Algorithmic scalability:** circuit architectures must explicitly address trainability issues (e.g., barren plateaus), shot-noise trade-offs, and resource-aware optimization, including error mitigation strategies and noise-resilient ansätze. (iii) **Benchmarking and standardization:** the community would benefit from shared datasets, reporting guidelines, and reference baselines that separate genuine quantum contributions from classical preprocessing gains; standardized protocols are essential to compare results across platforms and backends in a fair manner. (iv) **EO-driven problem selection:** near-term impact is most likely in specific EO applications domain where data are high-dimensional but labeled samples are scarce, where kernels and compact models are already competitive, or where discrete optimization naturally arises (e.g., mission planning, resource allocation, and certain inversion formulations).

In this context, community-building initiatives and coordination efforts (e.g., open repositories, shared tools, and the IEEE GRSS QUEST TC) play a crucial role in consolidating terminology, accelerating reproducibility, and bridging EO and QC expertise. Ultimately, the most credible short-term objective is not a generic quantum advantage claim, but the development of reliable, benchmarked, and hardware-aware quantum components that can be integrated into operational EO workflows. As quantum technologies mature, these components may evolve from hybrid add-ons into genuinely quantum-native models, opening new opportunities for efficient learning and inference in large-scale EO applications.



## Bibliography

- [1] A. Sebastianelli, **F. Mauro**, A. Delilbasic, P. Di Stasio, F. Fan, G. Meoni, G. Cavallaro, X. X. Zhu, P. Gamba, and S. L. Ullo, “Quantum Machine Learning for Earth Observation: A Review and Future Prospects,” *Submitted to IEEE Geoscience and Remote Sensing Magazine*, 2025.
- [2] D. Tuia, K. Schindler, B. Demir, X. X. Zhu, M. Kochupillai, S. Džeroski, J. N. van Rijn, H. H. Hoos, F. Del Frate, M. Datcu *et al.*, “Artificial intelligence to advance earth observation: A review of models, recent trends, and pathways forward,” *IEEE Geoscience and Remote Sensing Magazine*, 2024.
- [3] M. Wang, W. Fu, X. He, S. Hao, and X. Wu, “A survey on large-scale machine learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 6, pp. 2574–2594, 2022.
- [4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, Dec. 2010.
- [5] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Physics*, vol. 56, no. 2, pp. 172–185, 2015.
- [6] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, Sep 2017. [Online]. Available: <https://doi.org/10.1038/nature23474>
- [7] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, Jun 2021. [Online]. Available: <https://doi.org/10.1038/s43588-021-00084-1>
- [8] A. Miroszewski, J. Nalepa, B. L. Saux, and J. Mielczarek, “Quantum machine learning for remote sensing: Exploring potential and challenges,” in *Proceedings of the 2023 conference on Big Data from Space (BiDS'23): 6-9 November 2023, Austrian Center, Vienna, 2023*.
- [9] A. Nammouchi, A. Kassler, and A. Theocharis, “Quantum machine learning in climate change and sustainability: A short review,” in *Proceedings of the AAAI Symposium Series*, vol. 2, no. 1, 2023, pp. 107–114.
- [10] Otgonbaatar, Soronzonbold and Nurmi, Olli and Johansson, Mikael and Makela, Jarmo and Gawron, Piotr and Puchala, Zbigniew and Mielczarek, Jakub and Miroszewski, Artur and Dumitru, Corneliu Octavian and others, “Quantum computing for climate change detection, climate modeling, and climate digital twin,” *Authorea Preprints*, 2023.
- [11] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
- [12] E. Rieffel and W. Polak, “An Introduction to Quantum Computing for Non-Physicists,” *ACM Comput. Surv.*, vol. 32, no. 3, p. 300–335, Sep. 2000. [Online]. Available: <https://doi.org/10.1145/367701.367709>
- [13] D. P. DiVincenzo, “Quantum gates and circuits,” *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1969, pp. 261–276, 1998.
- [14] J. L. Park, “The concept of transition in quantum mechanics,” *Foundations of physics*, vol. 1, no. 1, pp. 23–33, 1970.

- 
- [15] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned," *Nature*, vol. 299, no. 5886, pp. 802–803, 1982.
- [16] A. Kumar Pati and S. L. Braunstein, "Impossibility of deleting an unknown quantum state," *Nature*, vol. 404, no. 6774, pp. 164–165, 2000.
- [17] A. Sebastianelli, **F. Mauro**, G. Ciabatti, D. Spiller, B. Le Saux, P. Gamba, and S. Ullo, "Quanv4eo: empowering earth observation by means of quanvolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
- [18] S. Otgonbaatar and M. Datcu, "Natural embedding of the stokes parameters of polarimetric synthetic aperture radar images in a gate-based quantum computer," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–8, 2021.
- [19] J. Preskill, "Quantum computation (physics 219/computer science 219)," *California Institute of Technology. Fall. Pasadena, California*, 2021.
- [20] E. N. Evans, D. Byrne, and M. G. Cook, "A quick introduction to quantum machine learning for non-practitioners," *arXiv preprint arXiv:2402.14694*, 2024.
- [21] Y. Du, X. Wang, N. Guo, Z. Yu, Y. Qian, K. Zhang, M.-H. Hsieh, P. Rebentrost, and D. Tao, "Quantum machine learning: A hands-on tutorial for machine learning practitioners and researchers," *arXiv preprint arXiv:2502.01146*, 2025.
- [22] M. Henderson, S. Shakya, S. Pradhan, and T. Cook, "Quanvolutional neural networks: powering image recognition with quantum circuits," *Quantum Machine Intelligence*, vol. 2, no. 1, p. 2, 2020.
- [23] M. Henderson, J. Gallina, and M. Brett, "Methods for accelerating geospatial data processing using quantum computers," *Quantum Machine Intelligence*, vol. 3, no. 1, p. 4, Jun. 2021. [Online]. Available: <https://link.springer.com/10.1007/s42484-020-00034-6>
- [24] R. Ghosh, A. Delilbasic, G. Cavallaro, and F. Bovolo, "A cnn architecture tailored for quantum feature map-based radar sounder signal segmentation," in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 442–445.
- [25] L. Miller, G. Uehara, A. Sharma, and A. Spanias, "Quantum machine learning for optical and sar classification," in *2023 24th International Conference on Digital Signal Processing (DSP)*. IEEE, 2023, pp. 1–5.
- [26] **F. Mauro**, A. Sebastianelli, M. P. Del Rosso, P. Gamba, and S. L. Ullo, "Qspecklefilter: A quantum machine learning approach for sar speckle filtering," in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, 2024, pp. 450–454.
- [27] A. Sebastianelli, **F. Mauro**, G. Ciabatti, D. Spiller, B. Le Saux, P. Gamba, and S. Ullo, "Quanv4eo: empowering earth observation by means of quanvolutional neural networks," *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
- [28] F. Razzano, P. Di Stasio, **F. Mauro**, G. Meoni, M. Esposito, G. Schirinzi, and S. L. Ullo, "AI Techniques for Near Real-Time Monitoring of Contaminants in Coastal Waters on Board Future  $\Phi$ Sat-2 Mission," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.
- [29] **F. Mauro**, F. Razzano, P. Di Stasio, A. Sebastianelli, G. Meoni, G. Schirinzi, P. Gamba, and S. L. Ullo, "Quantum-Enhanced Water Quality Monitoring: Exploiting  $\Phi$ Sat-2 Data with Quanvolution," *IEEE Geoscience and Remote Sensing Letters*, 2025.
-

- 
- [30] A. Chalumuri, R. Kune, S. Kannan, and B. S. Manoj, “Quantum-enhanced deep neural network architecture for image scene classification,” *Quantum Information Processing*, vol. 20, no. 11, p. 381, Nov. 2021.
- [31] C.-H. Lin, T.-H. Lin, and J. Chanussot, “Quantum Information-Empowered Graph Neural Network for Hyperspectral Change Detection,” *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [32] Y. Xu, H. Huang, and R. State, “Cropland quantum learning: A hybrid quantum-classical neural network for cropland classification,” in *2024 IEEE 3rd International Conference on Computing and Machine Intelligence (ICMI)*. IEEE, 2024, pp. 1–7.
- [33] F. Fan, Y. Shi, and X. X. Zhu, “Earth observation data classification with quantum-classical convolutional neural network,” in *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2022, pp. 191–194.
- [34] —, “Urban land cover classification from sentinel-2 images with quantum-classical network,” in *2023 Joint Urban Remote Sensing Event (JURSE)*. IEEE, 2023, pp. 1–4.
- [35] F. Fan, Y. Shi, T. Guggemos, and X. X. Zhu, “Hybrid quantum-classical convolutional neural network model for image classification,” *IEEE transactions on neural networks and learning systems*, 2023.
- [36] F. Fan, Y. Shi, and X. X. Zhu, “Land Cover Classification From Sentinel-2 Images With Quantum-Classical Convolutional Neural Networks,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 12 477–12 489, 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10602728/>
- [37] S. R. Majji, A. Chalumuri, R. Kune, and B. Manoj, “Quantum processing in fusion of sar and optical images for deep learning: A data-centric approach,” *IEEE Access*, vol. 10, pp. 73 743–73 757, 2022.
- [38] L. Miller, G. Uehara, and A. Spanias, “Quantum image fusion methods for remote sensing,” in *2024 IEEE Aerospace Conference*. IEEE, 2024, pp. 1–9.
- [39] S. Park, S. Jung, and J. Kim, “Dynamic Quantum Federated Learning for Satellite-Ground Integrated Systems Using Slimmable Quantum Neural Networks,” *IEEE Access*, vol. 12, pp. 58 239–58 247, 2024.
- [40] P. Gawron and S. Lewiński, “Multi-spectral image classification with quantum neural network,” in *IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2020, pp. 3513–3516.
- [41] G. Priyanka and M. Venkatesan, “Hyperspectral image classification using quantum machine learning,” in *2024 1st International Conference on Advanced Computing and Emerging Technologies (ACET)*. IEEE, 2024, pp. 1–7.
- [42] R. U. Shaik, A. Unni, and W. Zeng, “Quantum based pseudo-labelling for hyperspectral imagery: A simple and efficient semi-supervised learning method for machine learning classifiers,” *Remote Sensing*, vol. 14, no. 22, p. 5774, 2022.
- [43] X. X. Zhu, D. Tuia, L. Mou, G.-S. Xia, L. Zhang, F. Xu, and F. Fraundorfer, “Deep learning in remote sensing: A comprehensive review and list of resources,” *IEEE geoscience and remote sensing magazine*, vol. 5, no. 4, pp. 8–36, 2017.
-

- 
- [44] I. Liliopoulos, G. D. Varsamis, K. Milchanowski, R. Martin-Cuevas, K. Safouri, P. Dimitrakis, and I. G. Karafyllidis, “Hybrid classical-quantum multilayer neural networks for monitoring agricultural activities using remote sensing data,” *Quantum Machine Intelligence*, vol. 7, no. 1, p. 4, 2025.
  - [45] S. Rainjonneau, I. Tokarev, S. Iudin, S. Rayaprolu, K. Pinto, D. Lemtiuzhnikova, M. Koblan, E. Barashov, M. Kordzanganeh, M. Pflitsch *et al.*, “Quantum algorithms applied to satellite mission planning for Earth observation,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 16, pp. 7062–7075, 2023.
  - [46] F. Fan, Y. Shi, T. Guggemos, and X. X. Zhu, “Hybrid Quantum Deep Learning With Superpixel Encoding for Earth Observation Data Classification,” *IEEE Transactions on Neural Networks and Learning Systems*, 2025.
  - [47] D. A. Zaidenberg, A. Sebastianelli, D. Spiller, B. Le Saux, and S. L. Ullo, “Advantages and bottlenecks of quantum machine learning for remote sensing,” in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 2021, pp. 5680–5683.
  - [48] A. Sebastianelli, D. A. Zaidenberg, D. Spiller, B. Le Saux, and S. L. Ullo, “On circuit-based hybrid quantum neural networks for remote sensing imagery classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 565–580, 2021.
  - [49] A. Sebastianelli, M. P. Del Rosso, S. L. Ullo, and P. Gamba, “On quantum hyperparameters selection in hybrid classifiers for earth observation data,” *IEEE Geoscience and Remote Sensing Letters*, 2023.
  - [50] A. S. Malarvanan, “Hybrid quantum neural network advantage for radar-based drone detection and classification in low signal-to-noise ratio,” *arXiv preprint arXiv:2403.02080*, 2024.
  - [51] J. M. Zollner, “Quantum classifiers for remote sensing,” in *Proceedings of the 30th International Conference on Advances in Geographic Information Systems*, 2022, pp. 1–2.
  - [52] S. Y. Chang, B. Le Saux, S. Vallecorsa, and M. Grossi, “Quantum convolutional circuits for earth observation image classification,” in *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2022, pp. 4907–4910.
  - [53] J. M. Zollner, P. Walther, and M. Werner, “Satellite image representations for quantum classifiers,” *Datenbank-Spektrum*, vol. 24, no. 1, pp. 33–41, 2024.
  - [54] S. Otgonbaatar and M. Datcu, “Classification of remote sensing images with parameterized quantum gates,” *IEEE Geoscience and Remote Sensing Letters*, vol. 19, pp. 1–5, 2021.
  - [55] C.-H. Lin and Y.-Y. Chen, “Hyperqueen: Hyperspectral quantum deep network for image restoration,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–20, 2023.
  - [56] S.-M. Hsu, T.-H. Lin, and C.-H. Lin, “Hyperqueen-mf: Hyperspectral quantum deep network with multi-scale feature fusion for quantum image super-resolution,” in *2024 IEEE 13rd Sensor Array and Multichannel Signal Processing Workshop (SAM)*. IEEE, 2024, pp. 1–5.
  - [57] R. Ghosh, A. Delilbasic, G. Cavallaro, and F. Bovolo, “A hybrid quantum-classical cnn architecture for semantic segmentation of radar sounder data,” in *2024 IEEE Mediterranean and Middle-East Geoscience and Remote Sensing Symposium (M2GARSS)*. IEEE, 2024, pp. 366–370.
  - [58] **F. Mauro**, A. Sebastianelli, B. L. Saux, P. Gamba, and S. L. Ullo, “A hybrid mlp-quantum approach in graph convolutional neural networks for oceanic nino index (oni) prediction,” *arXiv preprint arXiv:2401.16049*, 2024.
-

- 
- [59] L. Papa, A. Sebastianelli, G. Meoni, and I. Amerini, “On the impact of key design aspects in simulated Hybrid Quantum Neural Networks for Earth Observation,” *arXiv preprint arXiv:2410.08677*, 2024.
- [60] S. Otgonbaatar, G. Schwarz, M. Datcu, and D. Kranzlmüller, “Quantum transfer learning for real-world, small, and high-dimensional remotely sensed datasets,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2023.
- [61] S. Y. Chang, S. Thanasilp, B. L. Saux, S. Vallecorsa, and M. Grossi, “Latent style-based quantum gan for high-quality image generation,” *arXiv preprint arXiv:2406.02668*, 2024.
- [62] F. De Falco, A. Ceschini, A. Sebastianelli, B. Le Saux, and M. Panella, “Quantum latent diffusion models,” *Quantum Machine Intelligence*, vol. 6, no. 2, p. 85, 2024.
- [63] R. Bhavsar, N. K. Jadav, U. Bodkhe, R. Gupta, S. Tanwar, G. Sharma, P. N. Bokoro, and R. Sharma, “Classification of potentially hazardous asteroids using supervised quantum machine learning,” *IEEE Access*, vol. 11, pp. 75 829–75 848, 2023.
- [64] M. van Waveren, M. Savinaud, G. Pasero, V. Defonte, P.-M. Brunet, O. Faucoz, P. Gawron, B. Gardas, Z. Puchała, and Ł. Paweła, “Comparison of quantum neural network algorithms for earth observation data classification,” in *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2023, pp. 780–783.
- [65] K. Glatting, J. Meyer, S. Huber, and G. Krieger, “Quantum kernel methods for insar phase unwrapping,” in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 437–441.
- [66] S. Otgonbaatar and D. Kranzlmüller, “Exploiting the quantum advantage for satellite image processing: Review and assessment,” *IEEE Transactions on Quantum Engineering*, 2023.
- [67] H. Painchart, M. Van Waveren, B. Mora, and G. Pasero, “Quantum algorithm for the analysis of temporal sequences of satellite images,” in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 804–807.
- [68] L.-H. Yu, X.-Y. Li, G. Chen, Q.-S. Zhu, H. Li, and G.-W. Yang, “Qusl: Quantum unsupervised image similarity learning with enhanced performance,” *Expert Systems with Applications*, vol. 258, p. 125112, 2024.
- [69] L. Pastori, A. Grundner, V. Eyring, and M. Schwabe, “Quantum neural networks for cloud cover parameterizations in climate models,” *arXiv preprint arXiv:2502.10131*, 2025.
- [70] A. Miroszewski, J. Mielczarek, F. Szczepanek, G. Czelusta, B. Grabowski, B. Le Saux, and J. Nalepa, “Cloud detection in multispectral satellite images using support vector machines with quantum kernels,” in *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2023, pp. 796–799.
- [71] A. Miroszewski, J. Mielczarek, G. Czelusta, F. Szczepanek, B. Grabowski, B. Le Saux, and J. Nalepa, “Detecting clouds in multispectral satellite images using quantum-kernel support vector machines,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS)*, 2023.
-

- 
- [72] A. Miroszewski, B. Le Saux, N. Longép  , and J. Nalepa, “Utility of quantum kernel machines in remote sensing applications,” in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 799–803.
- [73] R. U. Shaik and S. Periasamy, “Accuracy and processing speed trade-offs in classical and quantum svm classifier exploiting prisma hyperspectral imagery,” *International Journal of Remote Sensing*, vol. 43, no. 15-16, pp. 6176–6194, 2022.
- [74] D. Sarkar, E. Dimitrov, P. S. Vieites, M. Fernandez-Urrutia, V. Kannan, and P. PS, “Multiclass land use/land cover (lulc) classification using quantum enhanced support vector machines,” in *IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2024, pp. 446–449.
- [75] P. Rodriguez-Grasa, R. Farzan Rodr  guez, G. Novelli, Y. Ban, and M. Sanz, “Satellite image classification with neural quantum kernels,” *Machine Learning: Science and Technology*, 2024.
- [76] A. Miroszewski, M. F. Asiani, J. Mielczarek, B. L. Saux, and J. Nalepa, “In search of quantum advantage: Estimating the number of shots in quantum kernel methods.”
- [77] M. K. Gupta, M. Romaszewski, and P. Gawron, “Potential of quantum machine learning for processing multispectral earth observation data,” *Authorea Preprints*, 2023.
- [78] M. K. Gupta, M. Beseda, and P. Gawron, “How quantum computing-friendly multispectral data can be?” in *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2022, pp. 4153–4156.
- [79] E. Boyda, S. Basu, S. Ganguly, A. Michaelis, S. Mukhopadhyay, and R. R. Nemani, “Deploying a quantum annealing processor to detect tree cover in aerial imagery of california,” *PLOS ONE*, vol. 12, no. 2, pp. 1–22, 02 2017.
- [80] S. Otgonbaatar and M. Datcu, “A quantum annealer for subset feature selection and the classification of hyperspectral images,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 14, pp. 7057–7065, 2021.
- [81] A. Delilbasic, G. Cavallaro, M. Willsch, F. Melgani, M. Riedel, and K. Michielsen, “Quantum Support Vector Machine Algorithms for Remote Sensing Data Classification,” in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. Brussels, Belgium: IEEE, Jul. 2021, pp. 2608–2611.
- [82] A. Delilbasic, B. Le Saux, M. Riedel, K. Michielsen, and G. Cavallaro, “A Single-step Multiclass SVM Based on Quantum Annealing for Remote Sensing Data Classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing (JSTARS)*, 2023.
- [83] E. Zardini, A. Delilbasic, E. Blanzieri, G. Cavallaro, and D. Pastorello, “Local Binary and Multiclass SVMs Trained on a Quantum Annealer,” *IEEE Transactions on Quantum Engineering*, vol. 5, pp. 1–12, 2024. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/10706813>
- [84] A. Delilbasic, B. Le Saux, M. Riedel, K. Michielsen, and G. Cavallaro, “Quantum Annealing for Semantic Segmentation in Remote Sensing: Potential and Limitations,” in *2024 IEEE Mediterranean and Middle-East Geoscience and Remote Sensing Symposium (M2GARSS)*. Oran, Algeria: IEEE, Apr. 2024, pp. 376–380. [Online]. Available: <https://ieeexplore.ieee.org/document/10537465/>
-

- [85] E. Pasetto, A. Delilbasic, G. Cavallaro, M. Willsch, F. Melgani, M. Riedel, and K. Michielsen, “Quantum support vector regression for biophysical variable estimation in remote sensing,” in *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2022, pp. 4903–4906.
- [86] E. Pasetto, M. Riedel, K. Michielsen, and G. Cavallaro, “Kernel approximation on a quantum annealer for remote sensing regression tasks,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2024.
- [87] S. Otgonbaatar and M. Datcu, “Assembly of a coreset of earth observation images on a small quantum computer,” *Electronics*, vol. 10, no. 20, p. 2482, 2021.
- [88] R. Orús, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states,” *Annals of physics*, vol. 349, pp. 117–158, 2014.
- [89] S. Otgonbaatar and D. Kranzlmüller, “Quantum-inspired tensor network for earth science,” in *IGARSS 2023-2023 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2023, pp. 788–791.
- [90] J. Zhang, Y. Zhang, and Y. Zhou, “Quantum-inspired spectral-spatial pyramid network for hyperspectral image classification,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 9925–9934.





## Chapter 3

# Quanv4EO: Empowering Earth Observation by means of Quantonvolutional Neural Networks

A significant amount of remotely sensed data is generated daily by many EO spaceborne and airborne sensors over different countries of our planet. Different applications use those data, such as natural hazard monitoring, global climate change, urban planning, and more. Many challenges are brought by the use of these big data in the context of RS applications. In recent years, the employment of ML and DL-based algorithms has allowed a more efficient use of these data, but the issues in managing, processing, and efficiently exploiting them have even increased as classical computers have reached their limits. This chapter highlights how my research represents a significant shift towards leveraging QC techniques in processing large volumes of RS data. The proposed Quanv4EO framework introduces a quantum-inspired preprocessing strategy that leverages simulated quantum circuits to extract high-dimensional and information-rich features from EO imagery, enabling compact, efficient, and scalable learning pipelines. Its effectiveness was first demonstrated on standard image classification datasets (MNIST and FashionMNIST), and then on the **EuroSAT (Sentinel-2)** for LULC classification [1, 2]; **Sentinel-1** for building segmentation [3] and speckle filtering [4]; **ΦSat-2** for turbidity regression [5]. The proposed results underscore the potential of quantum-enhanced approaches in RS data analysis, paving the way for more efficient and effective solutions for wide geographical area EO data exploitation.

### 3.1 Introduction

Many challenges emerge with the use of big data in the context of RS applications, and even if the employment of ML and DL-based algorithms has allowed more efficient use of these data, the issues in managing, processing, and efficiently exploiting them have increased since classical computers have reached their limits [6], [7].

Quantum Convolutional or Quantonvolutional Neural Network (QuanvNN)s represent an exciting discovery and development in the field of DL, where the principles of QC are integrated into CNNs. QuanvNNs aim to enhance data processing and representation in the computer vision and EO domain by introducing quantum-inspired operations.

At the core of QuanvNNs are the quantonvolutional operations, which replace or augment the traditional convolutions used in classical CNNs. Quantonvolutions leverage quantum-inspired filters that capture complex quantum states, allowing for more expressive feature representations. These filters can be implemented using various quantum-inspired techniques such as quantum circuits, quantum gates, or other

quantum-inspired operations. By using these quantum filters, QuanvNNs can learn to extract features that exploit the inherent quantum nature of the data, leading to improved discrimination and generalization capabilities in various computer vision tasks [8].

This chapter introduces Quanv4EO, a quanvolution method for preprocessing multi-dimensional EO data. The work presented represents a significant milestone in the field of EO by introducing the application of QuanvNNs to the field of RS and paving the way for overcoming current limitations of classical computing and further enhancing the use of QC for EO tasks. One hurdle in utilizing QC for EO tasks is related to the requirement for large quantum circuits to process the sheer size of EO images. To harness the full potential of QC in this domain, the proposed approach provides a promising solution to handle large RS images with remarkable effectiveness.

With the proposed method, only a few qubits are needed to encode the information, rather than hundreds of them, while performance is maintained at the same level (or effectively improved) compared to comparable SOTA approaches.

QuanvNNs, as already mentioned, incorporate quantum-inspired operations into the architecture of CNNs, which have demonstrated in the last years to draw new paths for EO image processing and classification. The proposed solution opens up new avenues for leveraging the capabilities of QC in many EO tasks, such as image classification, object detection, and semantic segmentation.

Differently from previous works by Zaidenberg et al. [9, 10], where a hybrid QNN with a classical convolutional branch and a quantum layer as classifier was proposed, in Quanv4EO the quantum layers have been moved to the beginning, while the classification has been left to classical strategies. Moreover, kernel functions in high-dimensional Hilbert spaces have been accessed.

Overall, the quantum advantage of QuanvNNs is demonstrated with applications specifically to RS and EO domains. Moreover, a lazy training-like procedure is proposed: a method of learning wherein generalization from the training data is postponed until a query is directed to the system. All the above characteristics make the proposed model completely new and able to overcome some of the highlighted challenges. As discussed ahead, the main objective of this chapter is to demonstrate that the proposed approach could potentially serve as a resolution to challenges encountered in fully quantum approaches. The chapter is organized as follows. The SOTA is analyzed in Section 5.3 by introducing related works on QuanvNNs. The quanvolution approach and the related kernel functions are introduced in Section 5.4. In Subsection 3.4.2, the lazy training-like procedure proposed for the proposed model is explained. In order to reduce the simulation time, a parallelized version of the quantum convolution operation is developed, as explained in Subsection 3.4.3. In Subsection 3.4.4, the processing time is evaluated by varying a list of certain parameters, such as the size of the input image, the stride, the number of qubits, the kernel size, and the number of output channels. In Subsection 3.4.5, the evaluation of the feature maps, produced by a quanvolutional layer by varying the number of qubits and kernel size, is presented. In Section 5.6, the quality and usability of features extracted by these layers is evaluated by validating the proposed method through image classification and segmentation, filtering and regression tasks using both common datasets for image classification in the data science field and EO datasets. Comparisons carried out with the classical counterparts demonstrate how the feature maps extracted by using quantum kernels increase the performance of the proposed method. Extensive analysis is also performed, highlighting how the main objective of this research is not necessarily to assert superior performance but rather to demonstrate that the proposed approach could reach these performances with a significantly reduced number of parameters. Discussions and conclusions, reported in Section 3.6, close the chapter.

## 3.2 Related works on Quanvolution

CNNs have rapidly gained popularity in numerous ML applications, especially within the domain of image recognition. Much of the significant advantage offered by these networks stems from their capacity to extract features from data in a hierarchical fashion. These features are obtained through various transformational layers, with the convolutional layer being a prominent component that lends the model its name.

QuanvNNs were first introduced in [11], where a transformational layer known as a "quantum convolution" or "quanvolutional" layer was proposed. Quanvolutional layers process input data by locally transforming it through a series of random quantum circuits, similar in principle to the transformations performed by randomly applied convolutional filter layers. Assuming that these quantum transformations yield meaningful features suitable for classification, the overall algorithm may prove highly valuable in the context of near-term quantum computing. This is because it necessitates the use of compact quantum circuits with minimal or no error correction.

Specifically, in the quanvolutional approach, akin to a traditional convolutional layer, the quanvolutional layer generates feature maps by applying localized transformations to input tensors. However, instead of conducting element-wise matrix-matrix multiplications, the quanvolutional layer initiates the process by encoding an image patch into a quantum state. Subsequently, this state undergoes a sequence of operations involving two-qubit gates, like the CNOT gate, and parameterized one-qubit gates, such as rotations around the three axes of the Bloch sphere [12]. Given the constraints on the quantum hardware's size, processing an entire image is impractical. A workaround for this limitation is the use of quanvolution with smaller patches, as employed in convolution, using the available qubit count to encode the information. Importantly, the encoding of compact filter sizes does not necessitate Quantum Random Access Memory (QRAM)<sup>1</sup> technology, which remains under development.

QuanvNNs have a remarkable ability to discern and understand localized patterns within data, which is crucial for various tasks such as image recognition or signal processing. This capability stems from their inherent design, which allows them to effectively capture and represent intricate details present in different regions of an image. Moreover, QuanvNNs maintain translational invariance across the entire image, meaning that they can recognize patterns regardless of their position within the image. This characteristic is essential as it ensures that the network's learning remains consistent irrespective of shifts or transformations in the input data.

Furthermore, the integration of quanvolutional layers into modern layered ML models marks a significant advancement in leveraging quantum technology for practical applications. These quanvolutional layers seamlessly blend into existing architectures, facilitating the adoption of quantum techniques without requiring a complete overhaul of established frameworks.

Building upon these advancements, the research reported in [12] delves into the exploration of quantum image encoding techniques and their impact on the performance of a hybrid quanvolutional algorithm inspired by convolutional methods. The results indicate that the choice of image encoding is important, that there is no one-size-fits-all solution, and the final selection depends on the specific application.

Since the majority of current quantum circuits operate on one-dimensional objects, when dealing with two-dimensional data, like images, to encode them in quantum states, they must be transformed into

---

<sup>1</sup>QRAM is necessary for algorithms that require efficient access to large amounts of quantum data, such as Grover's algorithm, which benefits from quick access to a quantum database for searching a specific element. The main limitation of QRAMs is storage capacity, which is influenced by the complexity and scalability of quantum systems. Not having QRAM can be seen as positive in certain contexts because it avoids the technical challenges and limitations associated with their implementation. QRAMs are susceptible to decoherence and noise, which can make scaling them up difficult. Additionally, the need for external control for each QRAM query can introduce a significant opportunity cost, where the control hardware could be more efficiently used to run extremely parallel classical algorithms [13].

a one-dimensional array. In [14], the authors present various techniques for mapping 2D data to a 1D format. These techniques are explored within the domain of QuanvNNs. Every patch of an input image holds pixel data that is employed for quanvolutional processing in QuanvNNs, akin to the traditional CNN's convolution. The pixel arrangements within each patch may be altered into various scanning sequences before being inputted into the quantum encoder for the extraction of features. Experimental results reveal that the suggested mapping methods deliver substantial improvements in comparison to the conventional raster scan approach.

A QuanvNN extends a CNN by incorporating quanvolutional layers with quantum circuits for filters, showing potential benefits. In the use case presented in [15], a convolutional layer was replaced with a quanvolutional layer, resulting in similar classification accuracy but reduced loss. Models with quanvolutional layers resisted adversarial examples and were resilient to Circuit-Weighted Adversarial attacks. In a Quanvolution algorithm, every pixel's data is encoded within individual qubits. Once feature extraction is performed, a feature corresponding to the measurement outcome of each qubit in the PQC is derived. It is important to note that the quantity of features obtained through this process is equivalent to the number of qubits in the PQC, which coincides with the size of the kernel. In [16], a scaled QuanvNN was introduced to enhance the number of filters. Since an increased number of qubits in a single quantum circuit often leads to barren plateaus, the objective of Scalar QuanvNN is to mitigate barren plateaus by augmenting the number of filters, which in turn increases the number of quantum circuits rather than the qubits within the circuit. The scalar QuanvNN employs multiple filters, thereby providing flexibility in adjusting the quantity of extracted features.

It is important to highlight that Quanvolutional layers leverage quantum circuits in the creation of feature maps. The research proposed in [17] utilizes random quantum circuits as filters. A comparison between QuanvNN and CNN models can help determine whether the addition of a quanvolutional layer enhances model accuracy and performance. Both models were trained using various training dataset sizes. Beyond dataset size, the model was also tested with different numbers of Quanvolutional layers (1, 4, 8, 16, 32) using 600 training images and 150 test images. The model's performance improves as the number of layers increases, but this also leads to increased computational time.


Several recent studies have explored the use of trainable quantum circuits for quanvolution, which allow quantum filters to adapt their parameters during training. Notably, Chen et al. [18] proposed QuanvNNs with trainable circuits for analyzing high-energy physics data, showing that optimized quantum circuits can improve the extraction of relevant features. Similarly, Matic et al. [19] applied a quantum-classical hybrid model with trainable quantum layers to radiological image classification, demonstrating that trainable quantum layers could achieve better performance than fixed random circuits.

### 3.3 Datasets

The experiments conducted in this study utilized several datasets, analyzing both data science and EO benchmark data. Below, it is provided a detailed description of each dataset and its role in the analysis:

**MNIST and FashionMNIST** The MNIST dataset [20] is a widely used benchmark for image classification, consisting of 70,000 (60,000 for training and 10,000 for testing) grayscale images of handwritten digits (0-9) with a resolution of  $28 \times 28$  pixels. It is commonly employed to evaluate the performance of classification algorithms due to its simplicity and the availability of numerous pre-existing results for comparison. Similarly, the FashionMNIST dataset [21] includes 70,000 (60,000 for training and 10,000 for testing) grayscale images of clothing items, also with a resolution of  $28 \times 28$  pixels, making it a more challenging alternative to MNIST with the same structure. Both datasets were used to validate

**Sentinel-1 Data for SAR Filtering** In addition to optical data, the method was tested on SAR data from Sentinel-1 for evaluating the capabilities of the Quanv4EO framework in reducing speckle noise. SAR images are known for their resilience against weather conditions, making them valuable for continuous EO. However, they are affected by speckle noise, which can hinder image interpretation. The Sentinel-1 data (2637 with speckle and 2637 without speckle) used in this study were downloaded from the Google Earth Engine (GEE) and processed to evaluate the speckle-filtering capabilities of the proposed quantum model, following the same pipeline presented in [22]. Specifically, 2000 samples were used for training, 600 for validation, and 37 for testing as specified in [22]. This dataset was crucial for demonstrating the model’s potential in a practical application of quantum-enhanced feature extraction and noise reduction.



Data were retrieved from the *S1 SAR GRD collection* on GEE, containing both VV and VH polarizations. As Ground Truth (GT), they were utilized building polygons from the Open Buildings 2.5D Temporal dataset, which includes over 1.8 billion building footprints across Africa, Latin America, the Caribbean, South Asia, and Southeast Asia. The dataset offers data on building presence, counts, and heights at a 4m resolution, updated annually from 2016 to 2023.

To prepare the training and test datasets, it was performed a patching operation on both the input image and the corresponding GT. The original image was divided into  $256 \times 256$  pixel patches with a factor of stride equal to 128, generating partially overlapping *chips* to augment the dataset. The resulting patches were split into training and testing datasets (80-20), with the training set used for model training and the test set for evaluation. This dataset was essential for validating the Quanv4EO framework on a downstream task that is entirely different from both classification and regression.

**EuroSAT** The EuroSAT dataset [23] consists of Sentinel-2 satellite images covering 10 different land cover classes, including agricultural, residential, and water bodies, among others. The dataset includes 27,000 labeled images, and each of them has a resolution of  $64 \times 64$  pixels and provides 13 spectral bands. This dataset is specifically designed for LULC classification, making it a suitable choice for evaluating the capabilities of quantum-enhanced ML models like Quanv4EO in handling multi-dimensional EO data. The EuroSAT dataset was employed to assess the effectiveness of quanvolutional layers in extracting relevant features from complex EO images and improving classification accuracy. Each dataset was divided into a training set and a test set (80/20 ratio, respectively). Finally, it was ensured that the split was applied class-wise.

**$\Phi$ Sat-2 data for water quality monitoring** To develop a quantum-enhanced AI system for water quality monitoring, it was built a dataset integrating  $\Phi$ Sat-2 spectral data with Copernicus Marine Service products, ensuring strong alignment between satellite observations and reference bio-geochemical measurements [5, 24].

To enable fast and efficient water quality monitoring, simulated  $\Phi$ Sat-2 data, generated through the ESA OrbitalAI Challenge Simulator, were employed.  $\Phi$ Sat-2 provides high-resolution spectral data tailored for environmental applications, such as the water quality monitoring.  $\Phi$ Sat-2 is a 6U CubeSat equipped with a Visible to Near-Infrared (VIS/NIR) multispectral instrument, capable of capturing eight spectral bands (7 multispectral + 1 panchromatic) at medium-to-high resolution.

The  $\Phi$ Sat-2 Simulator processes Sentinel-2 L1C data to generate optical multispectral observations that closely resemble those of  $\Phi$ Sat-2. This choice ensures a higher alignment with  $\Phi$ Sat-2's sensor characteristics, leveraging its improved spatial resolution and spectral properties for enhanced environmental monitoring and analysis. The transformation involves: 1) spectral transformation, aligning Sentinel-2 bands with  $\Phi$ Sat-2's higher spatial resolution; 2) spatial resampling, adjusting pixel sizes to match  $\Phi$ Sat-2's sensor characteristics; 3) simulation of sensor-specific noise models, including variations in Signal-to-Noise Ratio (SNR) and optical distortions and 4) reflectance computation, producing AI-ready multispectral image patches.

To establish a high-quality GT dataset, it was relied on High-Resolution Ocean Colour (HR-OC) products from the Copernicus Marine Service, specifically the Mediterranean Sea, Bio-Geo-Chemical, L3, daily observation dataset. These Level 3 daily satellite-derived products provide essential bio-geochemical indicators relevant to coastal water quality assessment. The HR-OC dataset, derived from Sentinel-2/MSI sensor data, includes RS Reflectance, Turbidity, Solid Particulate Matter Concentration, Spectral Particulate Backscattering, Chlorophyll-a Concentration. For this study, turbidity is the primary focus, as it represents a key indicator of water quality and its variations over time. These products are processed at 100 m spatial resolution, covering a 20 km-wide coastal zone in the Mediterranean. The extracted latitude and longitude coordinates from Copernicus HR-OC data were used to define the precise locations from which  $\Phi$ Sat-2 patches were extracted.

The dataset construction starts with the selection of an Area Of Interest (AOI) within the Mediterranean Sea, identified using the Copernicus Marine Service Data Store based on the availability of HR-OC bio-




Figure 3.2: Area of Interest (AOI) selection and turbidity distribution. The left panel shows the satellite image of the selected AOI, with the red bounding box marking the analysis area. The right panel displays the corresponding turbidity map from Copernicus HR-OC data, where colors indicate turbidity levels (low values in blue, high values in red).

geochemical data. The primary selection criterion was the presence of turbidity measurements. From the Copernicus Marine Data Store, the extracted geospatial coordinates (latitude, longitude) corresponding to HR-OC turbidity observations, were extracted. For each recorded measurement, the dataset includes timestamp of acquisition, latitude and longitude of the observed point, and turbidity values. These extracted coordinates served as reference points for spatially aligning  $\Phi$ Sat-2 spectral data with Copernicus HR-OC measurements. To maintain geospatial consistency, the  $\Phi$ Sat-2 Simulator was configured to generate multispectral images of  $4240 \times 4240$  pixels, covering the selected AOI and ensuring alignment with GT marine data (Figure 3.2).

From the simulated  $\Phi$ Sat-2 imagery, they were extracted  $20 \times 20$  pixel patches, each precisely corresponding to the geospatial locations of HR-OC turbidity measurements retrieved from the Copernicus database. To ensure relevance to water quality analysis, patches were selected exclusively over marine areas. This process generated a dataset of 8256 patches, ensuring that each patch accurately reflects the spectral properties of  $\Phi$ Sat-2. The extracted patches are directly linked to the turbidity values provided by HR-OC, establishing a clear correspondence between the satellite data and the reference measurements. This alignment ensures that the dataset effectively captures spatial variations in turbidity, enabling a detailed analysis of water quality dynamics.

The careful selection of the Mediterranean AOI and the precise geospatial alignment between  $\Phi$ Sat-2 spectral observations and HR-OC turbidity measurements provide a solid basis for developing and testing advanced models, such as the Quanv4EO framework, on a completely new type of data ( $\Phi$ sat-2).

**Data Preprocessing** For each dataset, preprocessing steps were applied to ensure compatibility with the quantum processing pipeline. This included normalization of pixel values and resizing of images where necessary. Overall, these datasets provided a comprehensive benchmark for evaluating the performance of the Quanv4EO framework across diverse data types, ranging from simple handwritten digit recognition to complex multi-spectral satellite imagery, thus highlighting the versatility and potential of quantum-enhanced ML in RS applications.






Figure 3.3: Quanvolution Scheme, where a quantum kernel of  $2 \times 2$  qubits is shown (top). Features extracted with the quantum kernel on Sentinel-1 and Sentinel-2 images (bottom).

### 3.4 Methodology

A possible approach and description for quantum convolution (Quanvolution) directly derive from its classical counterpart. The establishing procedure is schematized in Fig. 3.3, where, for simplicity, we show the quantum convolution with a kernel of size  $2 \times 2$  and a quantum circuit of 4 qubits.

Let's assume that the quantum circuit used for the quanvolution operator is the VQC in Figure 3.3. With this configuration, a small region of the input image defined by the *kernel\_size* is encoded in a quantum circuit of at least  $kernel\_size^2$  qubits to the quantum state of Hilbert  $\mathbb{H}^{2^n}$  through the operator  $U_\phi(x)$ . The ansatz  $U_W(\vartheta)$ , with random parameters  $\vartheta$  and gates ( $R_x$ ,  $R_y$ ,  $R_z$ , and CNOT) [11], is applied to the quantum state  $|\phi(x)\rangle$  [25]. Measurement operations are performed to get the output, the expected values with respect to the observable  $\hat{O}$  are calculated, and the results are obtained as:

$$f(x, \vartheta) = \langle \phi(x) | U_W^\dagger(\vartheta) \hat{O} U_W(\vartheta) | \phi(x) \rangle \quad (3.1)$$

Under the conditions of the “lazy training regime,” as explained later in Section 3.4.2, the parameters  $\vartheta$  are frozen, but ultimately they can be optimized using a common gradient-descent technique:

$$\nabla_{\vartheta}(f(x, \vartheta)) = \frac{1}{2} \left[ f(x, \vartheta + \frac{\pi}{2}) - f(x, \vartheta - \frac{\pi}{2}) \right] \quad (3.2)$$

The mathematical formulation of quanvolution follows that of convolution in signal theory. Let us assume that  $X \in \mathbb{R}^{W \times H}$  is a single-band image of spatial dimension  $W$  and  $H$ , width and height respectively, and  $k \in \mathbb{R}^{M \times N}$  is a kernel with dimensions  $M$  and  $N$ , where typically  $M = N$  and  $M \ll W$  and  $N \ll H$ . The convolution of  $X$  and  $k$  is given by Equation (3.3).



$$\begin{aligned}
 y[m, n] &= X[m, n] * k[m, n] = \\
 &= \sum_{i=-\infty}^{+\infty} \sum_{j=-\infty}^{+\infty} X[i, j] \cdot k[m-i, n-j]
 \end{aligned} \tag{3.3}$$

Considering that both  $X$  and  $k$  are finite signals, the equation can be rewritten as in Equation (3.4):

$$y[m, n] = \sum_{i=-m}^{m+W} \sum_{j=-n}^{n+H} X[i, j] \cdot k[m-i, n-j] \tag{3.4}$$

For the case of quanvolution, Equation (3.4) can be rewritten as in Equation (3.5):

$$\begin{aligned}
 y[m, n] &= \sum_{i=-m}^{m+W} \sum_{j=-n}^{n+H} f(X[i, j], k[m-i, n-j]) = \\
 &= \sum_{i=-m}^{m+W} \sum_{j=-n}^{n+H} \langle \phi(X[i, j]) | U_W^\dagger(k[m-i, n-j]) \\
 &\quad \hat{O} U_W(k[m-i, n-j]) | \phi(X[i, j]) \rangle
 \end{aligned} \tag{3.5}$$

The introduced formulation assumes that the input image has only one channel and the convolution is 2D. Yet, this assumption can be easily extended to the case of multi-channel images with 2D convolution (each channel is processed by a 2D kernel and the results are combined properly), as well as multi-channel images with 3D convolution (the kernel is 3D, so the channels are processed together). Moreover, also in the case of quanvolution, the stride has the same meaning with respect to its classical counterpart, unlike the depth of the feature map, which for quanvolution is directly linked to the number of qubits since it cannot surpass the number of defined qubits. A similar rule applies to the kernel size; indeed, for a kernel size  $(n \times n)$ , at least  $n^2$  qubits are required. Using more qubits than those required by the kernel size results only in a larger amount of feature maps, as detailed in the next sections.

### 3.4.1 Encoding of Input Image Patches

To represent classical data  $\mathbf{x} \in \mathbb{R}^n$  within an  $n$ -qubit quantum circuit, a quantum feature map  $\phi : \mathbb{R}^n \rightarrow H^{2^n}$  is utilized, where  $H^{2^n}$  denotes a Hilbert space of dimension  $2^n$ . This process involves applying a unitary operator  $U_\phi(\mathbf{x})$  to the initial state  $|0\rangle^{\otimes n}$ , resulting in:

$$U_\phi(\mathbf{x}) |0\rangle^{\otimes n} = |\phi(\mathbf{x})\rangle = |\psi\rangle. \tag{3.6}$$

In the proposed Quanv4EO framework, each patch of the input image is encoded into a quantum state using **angle encoding**. Specifically, the pixel values of each patch are normalized and then mapped via rotational gates. For an input patch of size  $n \times n$ , the normalized pixel intensities are used to modulate the phase of each qubit in an  $n$ -qubit quantum circuit. This angle encoding approach allows for the transformation of classical information (pixel intensities) into the quantum domain, making it suitable for subsequent processing through quantum gates in the quanvolution operation.

It is worth mentioning that other encoding techniques exist, for instance **amplitude encoding**, which can also be used to represent classical data in quantum states. This method may offer different advantages depending on the specific characteristics of the input data and the structure of the quantum circuits [26].

### 3.4.2 Lazy Training Regime

By definition, “Lazy Learning” or “Lazy Training” includes those ML algorithms where the generalization of the training data is delayed until a query to the system is made, as opposed to “Eager Learning” or “Feature Training” algorithms, where generalization is attempted before receiving a query and some features are learned [27], [28]. This approach proves to be particularly useful when the training data are continuously updated with new samples and in the case of large, continuously changing training sets with few attributes that are commonly queried. A typical example of a lazy classifier is the K-Nearest Neighbors (KNN) algorithm. Specifically, the target function is approximated locally: this allows solving multiple problems and dealing with changes in the problem domain. In the case of DL, when dealing with the lazy-training regime, the dynamics are almost linear, and the Neural Tangent Kernel (NTK) almost does not change after initialization. In fact, in the NTK limit, the dynamics become linear with respect to the weight changes, hence they can be described by a frozen kernel  $\Theta$ . In the NTK limit, the learning dynamics are described as:

$$\Theta(w, x_1, x_2) = \nabla_w f(w, x_1) \cdot \nabla_w f(w, x_2) \quad (3.7)$$

which is the NTK definition, where  $x_1, x_2$  are the inputs, and  $\nabla_w$  is the gradient with respect to the parameters  $w$ . Defining with  $h$  the number of hidden neurons in the neural network, the  $\Theta$  kernel evolves in time, but, as  $h \rightarrow \infty$ ,  $\Theta(w, x_1, x_2) \rightarrow \Theta_\infty(x_1, x_2)$ . Hence, the kernel is frozen, and the dynamics converge in a time independent of  $h$  to a global loss minimum.

The Lazy Training Regime has been discussed and tested also in the case of QML [11], [29]. In particular, in [29], a version of *Quantum* NTK is defined and its dynamics in the frozen limit, where variational angles do not change much - i.e., in the lazy-training regime - are analytically solved.

In [11], a more practical application of lazy training for the quantum part of the QuanvNN architecture is discussed. Here, a quantum transformation layer is added to a classical CNN architecture, enabling the embedding of a set of quantvolutional filters within the network’s architecture. Notably, these quantvolutional filters utilize random quantum circuits. Their unique characteristic lies in their ability to efficiently access kernel functions in high-dimensional Hilbert spaces. This lazy training regime is adopted for several reasons. Firstly, it is derived from the insights presented in [11]. Additionally, the reported experiments have yielded promising results even without training the quantum layers. Furthermore, the optimization of quantum convolution for EO use cases remains a time-consuming task, although ongoing research efforts, as discussed in [30], are directed towards addressing this challenge.

### 3.4.3 Parallelization and Optimization

One of the main challenges for quantum convolution consists of the processing time. Contrary to its classical counterpart (e.g., convolutional layers implemented in TensorFlow, PyTorch, etc.), which have been developed and optimized over years, there is not an optimized version of the quantum convolution. To mitigate this problem, in the proposed framework it is developed a parallelized version of the quantum convolution operation.

To show the validity of the method and the advantages introduced in the proposed implementation through parallelization, several experiments were conducted on different hardware setups and under varying testing conditions. The two experimental setups used are as follows:

- **Setup-A:** Ubuntu, CPU: 11th Gen Intel(R) Core(TM) i7-11800H @ 2.30GHz, CORES: 8, LOGICAL PROCESSORS: 16, RAM: 32.0 GB, GPU: NVIDIA GeForce RTX 3080 with 16GB of dedicated RAM.
- **Setup-B:** MacOS, CPU: Intel Core i5 @ 2.4 GHz, CORES: 4, LOGICAL PROCESSORS: 8, RAM: 8.0 GB, GPU: Intel Iris Plus Graphics 655 with 1536 MB of dedicated RAM.

To evaluate the processing time of the quantum convolution, it was used an image with dimensions (64, 64, 3). Specifically, for the 3-channel images, a 2D quanvolution approach was applied. The choice between 2D and 3D quanvolution significantly impacts both computational efficiency and feature extraction quality. In 2D quanvolution, each channel of multi-channel data is processed separately, preserving the spatial relationships across channels while maintaining lower computational costs. This approach is particularly suitable for datasets where inter-channel correlations are moderate. Conversely, 3D quanvolution processes the entire multi-channel input simultaneously, capturing more complex inter-channel dependencies. However, this method requires a higher computational load and a greater number of qubits, which may not always be practical with current quantum hardware.

While this analysis focuses on 2D quanvolution, the proposed implementation is also capable of handling 3D quanvolution. This feature can be particularly beneficial for higher-dimensional datasets, such as hyperspectral data from PRISMA. However, exploring this option is beyond the scope of this study and is suggested as a direction for future research.

In Table 3.1, they are reported the processing times by comparing the two experimental setups and quantum convolution with and without parallelization. As expected, Table 3.1 clearly shows that each configuration with parallelized quantum convolution results in reduced processing times. The best results are obtained with Setup-A when using 16 logical processors.

Table 3.1: Processing time for Quanvolution

| Modality             | Processing Time (s) |         |
|----------------------|---------------------|---------|
|                      | Setup-A             | Setup-B |
| No Para.             | ~ 5.72              | ~ 9.3   |
| Para. 4-logi. proc.  | ~ 1.54              | ~ 4.4   |
| Para. 8-logi. proc.  | ~ 1.10              | ~ 3.2   |
| Para. 16-logi. proc. | ~ 0.95              | -       |

Parallelization was achieved through the Python library `joblib`, and the implementation can be found in GitHub - main branch. Furthermore, it was achieved a significantly lower processing time (as better described in Section 3.4.4) by developing an additional version using the `Jax` library, which can be found in GitHub - Jax branch.

The processing time analysis presented in this study utilized the most optimized version of the quanvolutional operator available at the time. Specifically, a parallelized implementation of the quantum convolution was used, which significantly reduced computation times by leveraging multiple logical processors. As shown in Table 3.1, parallelization led to a substantial decrease in the time required to process EO images, enabling the practical application of the Quanv4EO framework to real-world data. Following the release of a new version of the PennyLane library, additional optimization opportunities have emerged, particularly for quantum circuits with more than 20 qubits. These optimizations include the use of GPUs for accelerating quantum simulations, which could further reduce processing times for more complex circuits. Specifically, the use of GPUs could enable faster training and inference times, especially for larger-scale EO datasets that require the manipulation of higher-dimensional quantum states.

It is important to note that the current analysis reflects the performance of the quanvolutional operator using CPU-based parallelization, which already provides significant speed improvements compared to a non-parallelized implementation. However, as hardware and software tools continue to evolve, incorporating GPU-based acceleration as a future step could yield further enhancements, particularly in scenarios involving larger quantum circuits.

### 3.4.4 Processing Time Analysis

In this section, the processing time is evaluated by varying parameters such as the size of the input image, the stride, the number of qubits, the kernel size, and the number of output channels. It is important to underline that the processing time analysis was conducted after parallelization.




Figure 3.4: Quanvolution Image Processing Time by varying size of the input image, stride, number of qubits, kernel size, and number of output channels. The top left plot shows how timing is influenced by the number of qubits, image size, and stride. The top right shows how timing is influenced by the number of qubits, image size, and kernel size. The bottom plot shows how timing is influenced by the number of qubits, image size, and number of output feature maps.

As Figure 3.4 shows, processing time is influenced by image size, as it slowly increases with the size of the input image. The number of qubits also affects the time required by the quanvolutional layer to process an image. This, in combination with the image size and stride, is more evident at the bottom of Fig. 3.4. The stride significantly influences processing time; with an increase in stride size (from deep blue to yellow), processing time drastically decreases. This result was expected since a larger stride results in fewer applications of the quantum circuit and thus fewer operations necessary to process an image.

Kernel size also impacts processing time, although less than stride. In fact, the top right plot of Fig. 3.4 shows that with an increase in kernel size (from deep blue to yellow), processing time worsens. This result was also expected since kernel size is strictly related to the number of qubits. The same applies to the number of output feature maps: the bottom plot shows that increasing the number of output feature maps results in higher processing time, which again relates to the number of qubits. For simplicity, breaking apart graphs in Fig. 3.4, they were selected experiments with 8 and 16 qubits, a kernel size of  $2 \times 2$ , and 4 output feature maps. This allows for a clearer view of processing time (depending on image size), as shown in Fig. 3.5.




Figure 3.5: Quanvolution processing time with a configuration using 8/16 qubits, a  $2 \times 2$  kernel size, and 4 output feature maps.

### 3.4.5 Feature Maps Consistency

In this section, we evaluate the feature maps produced by a quanvolutional layer by varying the number of qubits and kernel size.

**Random Circuit and Reproducibility of Randomized Experiments** Given the structure of quanvolution with a random layer, as proposed by Henderson *et al.* [11], there is a connection between the kernel size and the number of feature maps that can be produced. Indeed, as shown in Fig. 3.6 and Fig. 3.7, when it is requested  $\#features > (kernel\_size)^2$ , it is only possible to get  $(kernel\_size)^2$  informative feature maps. This limitation arises because a random layer does not fully connect all the qubits, meaning that only the qubits used to ingest the image pixels are well-activated to produce the feature maps.

The Random Layers feature in PennyLane [31], which utilizes randomly parameterized single-qubit rotations and entangling gates, introduces variability in the quantum state evolution. However, due to the randomness in the connections and parameters, only a limited subset of the qubits effectively contributes to the generated features. This limitation arises because random layers may not induce a rich enough entanglement structure to fully explore the high-dimensional Hilbert space, especially when the number of requested features exceeds the squared kernel size. Consequently, as demonstrated in the figures, the number of feature maps that provide distinct information tends to saturate at  $(kernel\_size)^2$ , aligning with the number of active qubits.

In the experiments involving random quantum circuits, multiple runs were conducted with different random seeds. This approach ensures that the results are not dependent on a single random initialization, thus providing a more reliable assessment of the model's performance.

**Custom Circuit** A solution to the aforementioned problem is the implementation of a circuit with a custom component on top of the random layer. For example, by mutually connecting all the qubits with CNOT gates and then applying a random layer, it is possible to obtain 16 feature maps even with a kernel size of  $3 \times 3$  or  $2 \times 2$ , as shown in Fig. 3.8.




Figure 3.6: Quanvolutional feature maps generated using a random quantum circuit with 16 qubits, for different kernel sizes ( $4 \times 4$ ,  $3 \times 3$ , and  $2 \times 2$ ).




Figure 3.7: Quanvolutional feature maps obtained from a random quantum circuit with 9 qubits, shown for kernel sizes  $3 \times 3$  and  $2 \times 2$ .



Figure 3.8: Quanvolutional feature maps produced by a custom quantum circuit with 16 qubits, combining a highly entangling layer with a random layer, for kernel sizes  $3 \times 3$  and  $2 \times 2$ .




Figure 3.9: Comparison between quanvolutional and classical convolutional feature maps (same spatial scale). Despite the absence of training, quanvolutional features exhibit richer and more structured representations than classical convolutional features.

**Differences with Convolutional Layers** A comparison between quantum convolutional and classical convolutional feature maps is presented in Fig. 3.9, showcasing qualitative differences (both are shown on the same scale). Notably, despite lacking training, quantum feature maps appear to contain more information than their classical counterparts. This enhancement is especially evident in classical feature maps with minimal information, indicating the effectiveness of the quantum approach. Additionally, the use of random circuits parallels classical weight initialization, underscoring the distinction in initialization methods between classical and quantum contexts (weights in classical vs. gates in quantum).

### 3.5 Experimental Results

In QML, the training process often involves optimizing quantum circuits or variational parameters representing the model. Training a quantum model necessitates manipulation of quantum states, measurements, and optimization algorithms, tasks that can be computationally expensive and resource-intensive. The proposed quanvolutional layers, essentially fixed quantum kernels processing EO data, aim to efficiently access kernel functions in high-dimensional Hilbert space. It's crucial to note that while QML can exploit quantum system properties for efficient computations, direct manipulation of a "huge Hilbert space" in terms of individual quantum states is not typical. QC operates within a Hilbert space, representing a quantum system's state space, with its size exponentially growing with qubit count, a factor promising for specific computations. QML algorithms often focus on quantum variational algorithms (e.g., Variational Quantum Eigensolver, Quantum Approximate Optimization Algorithm) or QNNs, utilizing quantum circuits for tasks like optimization or classification, leveraging quantum mechanics principles. Despite their potential, QML algorithms still face limitations and are in early development stages. Directly accessing and manipulating a "huge Hilbert space" typically necessitates a fully realized, fault-tolerant quantum computer, a long-term goal in QC research [32].

To assess feature quality and usability extracted by these layers, they were conducted several experiments, firstly on data science datasets and then on EO datasets. Regarding the former, it was ran MNIST and Fashion MNSIT classification. For the latter, it was ran LULC classification experiments within a simulated *Pennylane* environment, utilizing various methodologies: 1) Balanced Iterative Reducing and Clustering using Hierarchies (BIRCH) clustering, an unsupervised method, 2) Automated Machine Learning (AutoML), a supervised ML method, and 3) Automated Deep Learning (AutoDL), a supervised DL method. Moreover, in other case studies, it was demonstrated the efficiency of QuanvNNs in estimating turbidity levels from  $\Phi$ Sat-2 data, extracting speckle noise and segmenting buildings from SAR data. Finally, it was also explored the possibility to train the quanvolutional layer.

Furthermore, the proposed architecture was tested on satellite images with a size of ((64, 64, 3)) for EuroSAT in the frozen configuration, on  $\Phi$ Sat-2 patches for turbidity regression with a size of ((20, 20, 7)), on Sentinel-1 SAR patches for building segmentation with a size of ((256, 256, 2)) and ((64, 64, 2)) for speckle filtering. In the trainable configuration, the model was also evaluated on resized EuroSAT images of ((16, 16, 3)), Finally, MNIST and Fashion-MNIST datasets with size ((28, 28, 1)) were used as additional benchmarks.

#### 3.5.1 Framework and Model Configurations

In this study, various AutoML and AutoDL frameworks, alongside clustering algorithms, were employed to develop efficient and scalable models for classification and unsupervised learning tasks. Additionally, the proposed approach known as quanvolution was integrated into the analysis, leveraging quantum properties to enhance traditional convolutional operations. Each method, was configured to optimize both accuracy and computational cost, considering the characteristics and size of the dataset. The following paragraphs outline the specific configurations and models explored.

**Features Clustering** BIRCH Clustering is a hierarchical clustering algorithm designed for large-scale datasets. It aims to efficiently and effectively cluster data by incrementally building a clustering feature tree. BIRCH algorithm clusters data points by maintaining a compact summary of the dataset, called the Cluster Feature tree, which contains information about the clustering structure. The algorithm utilizes a combination of distance metrics and density-based clustering to determine cluster centers and assign data points to clusters. BIRCH is particularly suited for applications with large datasets and provides

a scalable approach to clustering [33]. Clustering was conducted using the BIRCH algorithm from the *scikit-learn* library.

A hyperparameter search was conducted to optimize the BIRCH clustering, focusing on the *threshold* (0.84) and *branching factor* (9) parameters, which control cluster compactness and tree branching, respectively. The best configuration was selected based on the lowest error in matching predicted cluster counts to actual label distributions.

**AutoML for Classification** AutoML is a process that automates the development of ML models. It simplifies tasks such as data pre-processing, feature engineering, algorithm selection, hyperparameter tuning, and model evaluation. AutoML leverages AI and ML techniques to automatically search for the best model and configuration. Overall, AutoML accelerates model development, enables non-experts to use ML, and improves efficiency and effectiveness in building models [34]. Moreover, AutoML helps to reduce/remove the model-selection bias that can be introduced by researchers.

AutoML was implemented using *AutoSklearn* to automate model selection and hyperparameter optimization for classification tasks. The configuration set a total execution time of 1200 seconds and a per-model time limit of 120 seconds (i.e., *time left for this task=1200*, *per run time limit=120*). For simplicity and interpretability, the final output was constrained to a single model (*ensemble size=1*).

AutoSklearn was configured to explore a variety of standard classifiers, such as *Decision Tree*, *Random Forest*, and *SVM*, alongside feature preprocessors, including *PCA*, *Independent Component Analysis (ICA)*, and *Feature Agglomeration*. This setup effectively balanced predictive performance with computational efficiency.

**AutoDL for Classification** Similarly to AutoML, which automates the development of ML models, AutoDL automates the development of DL models. It simplifies tasks such as architecture design, hyperparameter tuning, and model evaluation. AutoDL utilizes algorithms and techniques to automatically search for the best DL model, saving time and effort. By automating the process, AutoDL enables users to focus on problem formulation and interpretation of results. It accelerates the DL model development and it makes it more accessible [35], and also in this case, it helps to reduce/remove the model-selection bias that can be introduced by researchers.

For AutoDL, *AutoKeras* was employed to perform model search and hyperparameter tuning for classification tasks. AutoKeras utilizes a Neural Architecture Search (NAS) framework, systematically exploring model architectures and configurations to identify the best-performing network for a given dataset.

In this study, the configuration was limited to the *vanilla* model set, with no custom network overwrites. The search parameters were set to a maximum of 100 trial (*max trials=100*), allowing each model to train for up to 100 epochs (*epochs=100*). This approach enabled a streamlined yet effective exploration of DL architectures, balancing model quality with computational constraints.

### 3.5.2 Results on Data Science datasets

In order to validate the proposed method some experiments were initially ran using non EO dataset, particularly the MNIST [20] and the Fashion MNIST dataset [21].

**MNIST dataset** The results obtained by the Quanv4EO framework on the MNIST dataset demonstrate high accuracy, reaching 0.9984, with a remarkably small model size of only 42k parameters plus 16 qubits (frozen). This positions the proposed method in competition with the most advanced models in image classification on MNIST dataset, such as Branching CNN + HVC, EnsNet, and Efficient-



CaspNet. Despite its small size (significantly reduced compared to other models in the SOTA), the proposed method maintains high precision, offering an efficient and promising alternative for image data classification.

Table 3.2: Comparisons with SOTA on MNIST data classification. (Classical Model rank)

| Model                    | Accuracy | Model Size          |
|--------------------------|----------|---------------------|
| Branching CNN + HVC [36] | 0.9987   | 1.5 M               |
| EnsNet [37]              | 0.9984   | -                   |
| Efficient-CaspNet [38]   | 0.9984   | 160 k               |
| SOPCNN [39]              | 0.9983   | 1.4 M               |
| RMDL [40]                | 0.9982   | -                   |
| Quanv4EO                 | 0.9984   | 42 k + 16q (frozen) |

**FashionMNIST dataset** Similarly, on FashionMNIST dataset, the Quanv4EO framework proved to be highly competitive, achieving an accuracy of 96.81%. Furthermore, the model remains incredibly small, with the same size as the model used for MNIST classification (42k + 16 qubits, frozen). This result is really encouraging, since the proposed method offers an excellent balance between model size and performance.

Table 3.3: Comparisons with SOTA on Fashion MNIST data classification. (Classical Model rank)

| Model                       | Accuracy | Model Size          |
|-----------------------------|----------|---------------------|
| Fine-Tuning DARTS [41]      | 0.9691   | 3.9 M               |
| Shake-Shake [42]            | 0.9641   | -                   |
| PreAct-ResNet18 + FMix [43] | 0.9636   | -                   |
| Random Erasing [44]         | 0.9635   | -                   |
| E2E-3M [45]                 | 0.9592   | -                   |
| Quanv4EO                    | 0.9681   | 42 k + 16q (frozen) |

In both cases, the results confirm the effectiveness and efficiency of the Quanv4EO framework in image data classification on standard datasets like MNIST and FashionMNIST. The implemented technique provides a solid foundation for application across a wide range of contexts, ensuring high-level performance with reduced computational requirements.

### 3.5.3 Results on EO datasets

**Features Clustering** This section presents the clustering results when using the unsupervised BIRCH Clustering to the feature maps extracted using Quanvolution. Moreover, it was repeated the same experiments, under the same settings, using instead feature maps extracted with classical convolution. Before presenting the results, it is worth to highlight that BIRCH Clustering is not meant to work with images. The adaptation to images worked smoothly. Nevertheless, the scope was to test Quanvolutional features maps using as many learning strategies as possible, thus leading us to the choice of using this kind of unsupervised strategy also in this case.

In order to work with a data format as close as possible to the one more compatible with BIRCH Clustering, 2 quanvolutional layers were applied with a kernel size of 4 and stride 4 on the EuroSAT dataset

[23] obtaining for each image in the dataset a feature map of  $4 \times 12$ . The PCA was then applied to compress the dimension of the features from 192 ( $4 \times 12$ ) to 10. The BIRCH Clustering was then applied to this new set of variables.

In the end, different plots were used to show that quanvolution could better differentiate the features, as shown at the top of Figure 3.10. Indeed, they are more distant with respect to the CNN ones shown at the bottom of Figure 3.10. Even if the clustering worked perfectly on both methods, the quanvolution works required a simpler model, indeed drawing a straight line is sufficient to split the features with an acceptable error, while for the convolutional one a polynomial function is needed. In a first phase (simple scenario) clustering was applied on only 2 classes of our dataset (Figure 3.10) and then (more complex scenarios) the analysis was expanded to the whole dataset (Figure 3.11). Since EuroSAT is a labeled dataset it was possible to compute the quality of clustering by counting the items clustered in the right class. Findings suggested that, in the simple scenario both models achieved a score of 100%, but the quantum model is better in making the features more differentiable; in the more complex scenario, quanvolution produced 8% miss-classified labels, while the simple convolution produced 11% miss-classified labels.

Since the dataset is labeled, although the method is unsupervised, it is still possible to calculate classification (clustering) accuracy. Results are reported in Tab. 3.4, where they are compared the results of the same BIRCH clustering model trained on Quanvolution feature maps and on classic features maps, respectively. Results show that the latter performed better in most of the classes, apart from River and Permanent Crop class where the quanvolutional features maps are beaten by the classic ones, while for the class Sea Lake, both models have shown the same performance. Generally speaking, the feature maps extracted using a quantum kernel introduce a 2% of gain.

Table 3.4: Results of BIRCH Clustering for QuanvNN and CNN feature maps.

| Class         | CNN         | QCNN        |
|---------------|-------------|-------------|
| Annual Crop   | 0.87        | <b>0.99</b> |
| Forest        | 0.85        | <b>0.89</b> |
| H. Vegetation | 0.99        | <b>0.99</b> |
| Highway       | 0.97        | <b>0.99</b> |
| Industrial    | 0.99        | <b>0.99</b> |
| Pasture       | 0.85        | <b>0.92</b> |
| P. Crop       | <b>0.99</b> | 0.93        |
| Residential   | 0.90        | <b>0.94</b> |
| River         | <b>0.97</b> | 0.95        |
| Sea Lake      | <b>0.99</b> | <b>0.99</b> |
| Mean Acc.     | 0.94        | <b>0.96</b> |

**AutoML Classifier with quantum features** This section presents the classification results when using the AutoML strategy to the feature maps extracted using Quanvolution. Moreover, the same experiments were repeated, under the same settings, using feature maps extracted with classical convolution.

Before presenting the results, it is worth highlighting that AutoML is designed for tabular data. Although it was successfully adapted for images, this adaptation might affect the results. In any case, since the goal was to test Quanvolution feature maps with a variety of learning strategies, the choice was to use also the AutoML.

Results are reported in Tab. 3.5, where results of the same AutoML model trained on Quanvolution




Figure 3.10: Features clustering with Quanvolution (top) and with CNN (bottom) - Binary Case (considered as example Annual Crop and River classes of EuroSAT).




Figure 3.11: Features clustering with Quanvolution (top) and with CNN (bottom) - Multiclass Case

feature maps are compared with ones on Classic features maps respectively. Results show that the former performed better in most of the classes, apart from Highway and Permanent Crop class where both models performed the same. Generally speaking, the feature maps extracted using a quantum kernel introduced 2% of gain.

**AutoDLL classifier with quantum features** As for the previous section, this one presents the results when AutoDL is applied on quanvolution and classic convolution feature maps.

Table 3.5: Results of AutoML applied to QCNN and CNN feature maps.

| Class         | CNN         | QCNN        |
|---------------|-------------|-------------|
| Annual Crop   | 0.90        | <b>0.91</b> |
| Forest        | 0.96        | <b>0.98</b> |
| H. Vegetation | 0.84        | <b>0.85</b> |
| Highway       | <b>0.85</b> | <b>0.85</b> |
| Industrial    | 0.93        | <b>0.95</b> |
| Pasture       | 0.89        | <b>0.91</b> |
| P. Crop       | <b>0.85</b> | <b>0.85</b> |
| Residential   | 0.90        | <b>0.93</b> |
| River         | 0.87        | <b>0.89</b> |
| Sea Lake      | 0.95        | <b>0.96</b> |
| Mean Acc.     | 0.89        | <b>0.91</b> |

Results are reported in Tab. 3.6, where the results of the same AutoDL model trained on Quanvolution feature maps are compared with ones on Classic features maps respectively. Results showed that the former performed better for all the classes. The feature maps extracted using a quantum kernel have introduced a 5% of gain.

Table 3.6: Results of AutoDL applied to QCNN and CNN feature maps.

| Class         | CNN  | QCNN        |
|---------------|------|-------------|
| Annual Crop   | 0.88 | <b>0.93</b> |
| Forest        | 0.94 | <b>0.98</b> |
| H. Vegetation | 0.80 | <b>0.87</b> |
| Highway       | 0.81 | <b>0.88</b> |
| Industrial    | 0.92 | <b>0.95</b> |
| Pasture       | 0.83 | <b>0.91</b> |
| P. Crop       | 0.77 | <b>0.88</b> |
| Residential   | 0.94 | <b>0.97</b> |
| River         | 0.85 | <b>0.90</b> |
| Sea Lake      | 0.96 | <b>0.98</b> |
| Mean Acc.     | 0.88 | <b>0.93</b> |

In summary, by means of quanvolution it was possible to reach better performance, with an improvement of 2-5% with respect to standard CNN. It can be observed that the BIRCH clustering is working the best, and this is amenable to the fact that hyperparameter selection for this algorithm is much easier than AutoML and AutoDL. Indeed, it is worth highlighting that for both cases, AutoML and AutoDL, the search was limited to a certain amount of time. Moreover, all the possibilities of these two tools were not explored. Given that, it is possible, with more time and extended exploration, to reach better performances.

The important finding relies on the fact that in all cases it was reached an improvement, and specifically for the case of AutoDL, which reached better performances for all the classes. This is magnified by the fact that quanvolutional layers are “frozen”, hinting that results can be even better. Moreover, only two layers of quanvolution were used, so the global model is not that complex. This finding

is in line with [10], indeed this paper demonstrates that hybrid-quantum solution can easily achieve better/comparable performance with respect to classical solutions, but with a light-weight structure.

Moreover it was also observed that quantum models have a faster convergence - less number of epochs to reach the convergence - with respect to standard models, of 80% on average. This finding is also in line with other works in the same field [10, 46–48].

Table 3.7 poses a challenge in explanation. In this study, quantum processing is strategically positioned at the network’s outset, delegating the classification task to clustering (unsupervised ML), AutoML, and AutoDL. The intention is not necessarily to assert superior performance but rather to sufficiently demonstrate that the proposed approach could potentially serve as a resolution to challenges encountered in fully quantum approaches. In future works, this can be reached by integrating a quantum classifier [10] following the quanvolution layers, thereby avoiding significant issues tied to machine size when applied to EO data.

Table 3.7: Comparisons with SOTA on EuroSAT dataset classification

| Model                                            | Overall Accuracy | Model Size            |
|--------------------------------------------------|------------------|-----------------------|
| Helber et Al. [23] ResNet-50                     | 0.98             | 25 M                  |
| Helber et Al. [23] GoogleNet                     | 0.98             | 7 M                   |
| Li et Al. [49] ResNet-18                         | 0.98             | 11 M                  |
| Sumbul et Al. [50]                               | 0.70             | 23 k                  |
| Sebastianelli et Al. [10] (Ry Circuit)           | 0.79             | $42k + 4q$            |
| Sebastianelli et Al. [10] (Bellman Circuit)      | 0.84             | $42k + 4q$            |
| Sebastianelli et Al. [10] (Real Amplitudes)      | 0.92             | $42k + 4q$            |
| Sebastianelli et Al. [10] (Coarse-to-fine grain) | 0.97             | $4 \times (42k + 4q)$ |
| Quanv4Eo + Clustering                            | $0.96 \pm 0.03$  | 16q (frozen)          |
| Quanv4Eo + AutoML                                | $0.91 \pm 0.04$  | 16q (frozen)          |
| Quanv4Eo + AutoDL                                | $0.93 \pm 0.03$  | $42k + 16q$ (frozen)  |

**Turbidity Regression on  $\Phi$ Sat-2 Data** Coastal water quality monitoring is crucial for environmental sustainability and public health. In [5] a very cutting-edge methodology is introduced, using  $\Phi$ Sat-2 multispectral data and QuanvNNs to explore quantum-enhanced ML for water contaminant assessment. To evaluate the effectiveness of the proposed quantum approach, simulations were performed across multiple quantum circuit configurations and compared against a classical baseline model (AI4EDoET [24]). The results are summarized in Table 3.8. According to the constraint on kernel size and number of qubits, for 4, 9, and 16 qubits, kernel sizes of  $2 \times 2$ ,  $3 \times 3$ , and  $4 \times 4$  were used, respectively.

One of the most notable advantages of this approach is the significant reduction in the number of parameters while maintaining high performance. The classical model AI4EDoET consists of 818,349 parameters, whereas the proposed quantum circuits require substantially fewer parameters, ranging from 4,641 to 9,857 for quantum circuits with up to 16 qubits. This represents a drastic reduction in computational complexity, highlighting the efficiency of quantum processing in parameter optimization. Moreover, the Pearson correlation coefficient, calculated between each channel and the GT and then averaged across all channels, remains consistently high across all quantum configurations: its values remain above 0.83, with most configurations achieving values close to 0.90. The Root Mean Squared

Error (RMSE) values remain also competitive, with quantum circuits yielding results comparable to or even better than the classical counterpart, particularly for lower parameter configurations.

Table 3.8: Simulation results for different quantum circuit configurations and classical baseline (AI4EDoET)

| # Qubits                                  | # Parameters | Pearson Mean $\pm$ STD                 | RMSE          | MAE           |
|-------------------------------------------|--------------|----------------------------------------|---------------|---------------|
| <b>Basic Entangler Circuit</b>            |              |                                        |               |               |
| 4                                         | 4641         | $0.8880 \pm 0.01752$                   | 0.2450        | 0.1886        |
| 4                                         | 9089         | $0.8880 \pm 0.01752$                   | 0.3338        | 0.2611        |
| 9                                         | 4801         | $0.8986 \pm 0.00064$                   | 0.3797        | 0.2962        |
| 9                                         | 9409         | $0.8986 \pm 0.00064$                   | 0.2639        | 0.2085        |
| 16                                        | 5025         | $0.8952 \pm 0.00097$                   | 0.3838        | 0.3229        |
| 16                                        | 9857         | $0.8952 \pm 0.00097$                   | 0.3362        | 0.2730        |
| <b>Strongly Entangling Circuit</b>        |              |                                        |               |               |
| 4                                         | 4641         | $0.8981 \pm 0.00022$                   | 0.2521        | 0.1969        |
| 4                                         | 9089         | $0.8981 \pm 0.00022$                   | 0.2652        | 0.2050        |
| 9                                         | 4801         | $0.8986 \pm 0.00064$                   | 0.2286        | 0.1789        |
| 9                                         | 9409         | $0.8986 \pm 0.00064$                   | 0.3230        | 0.2568        |
| 16                                        | 5025         | $0.8308 \pm 0.16260$                   | 0.4501        | 0.3802        |
| 16                                        | 9857         | $0.8308 \pm 0.16260$                   | 0.5469        | 0.5037        |
| <b>Random Circuit</b>                     |              |                                        |               |               |
| 4                                         | 4641         | $0.8966 \pm 0.00052$                   | 0.2190        | 0.1745        |
| <b>4</b>                                  | <b>9089</b>  | <b><math>0.8966 \pm 0.00052</math></b> | <b>0.1980</b> | <b>0.1591</b> |
| 9                                         | 4801         | $0.8957 \pm 0.00053$                   | 0.2262        | 0.1806        |
| 9                                         | 9409         | $0.8957 \pm 0.00053$                   | 0.6608        | 0.5591        |
| 16                                        | 5025         | $0.8952 \pm 0.00097$                   | 0.4591        | 0.3333        |
| 16                                        | 9857         | $0.8952 \pm 0.00097$                   | 0.7184        | 0.5522        |
| <b>Classical Baseline (AI4EDoET [24])</b> |              |                                        |               |               |
| -                                         | 818349       | $0.8409 \pm 0.1119$                    | 0.2135        | 0.1606        |

This work shows that quantum-enhanced preprocessing can make water quality monitoring models both lighter and more accurate: using quanvolution with  $\Phi$ Sat-2 data, the model needs 98% fewer parameters while improving both Pearson correlation and RMSE over a classical baseline. The new  $\Phi$ Sat-2 + Copernicus turbidity dataset shows that different quantum circuits give similar correlations, with Random circuits slightly better on error metrics, supporting the idea that quantum-assisted feature extraction is an efficient and scalable option for EO regression tasks.

**QSpeckleFilter** Observing Earth through SAR provides valuable insights regardless of weather conditions. However, SAR imagery often suffers from speckle noise, hindering accurate interpretation. As an application of QuanvNNs, in [4] a QSpeckleFilter is proposed, which outperforms its classical counterpart (based on CNNs), demonstrating promising advancements in EO applications.

QSpeckleFilter builds upon the previous work by Sebastianelli et al. [22], utilizing quantum convolution and a modified speckle denoiser applied to Ground Range Detected (GRD) Sentinel-1. The proposed method significantly enhances performance metrics, indicating the effectiveness of quanvolution in SAR image processing. By processing the dataset with Quanvolutions, the dataset dimensions were expanded, enhancing noise extraction capabilities. Following preprocessing using quanvolution, the architecture proposed in [22] has been fine-tuned to incorporate quanvolution-preprocessed data. The

proposed method utilizes quanvolved feature maps to estimate the speckle noise within each image. Subsequently, the identified speckle is subtracted from the Sentinel-1 image using a *skip connection* and a *subtraction layer*, resulting in a residual model. The skip connection serves as a bypass for the CNN, while the subtraction layer performs mathematical subtractions between the input data propagated through the skip connection and the CNN output. The methodology demonstrates improved PSNR and SSIM values compared to previous method, as shown in Table 3.5.3.

Table 3.9: Proposed model’s average scores on the testing dataset: (a) GT, (b) Input with speckle, (c) Sebastianelli et Al. [22] and (d) Proposed **QSpeckleFilter**.

| Model                         | PSNR $\uparrow$ | SSIM $\uparrow$ |
|-------------------------------|-----------------|-----------------|
| (a) Ground Truth              | $+\infty$       | 1.0             |
| (b) Speckled                  | 15.70           | 0.58            |
| (c) Sebastianelli et Al. [22] | 19.21           | 0.75            |
| (d) <b>QSpeckleFilter</b>     | 21.72           | 0.81            |

**Building Segmentation** Building segmentation in urban areas is essential in fields such as urban planning, disaster response, and population mapping. Yet accurately segmenting buildings in dense urban regions presents challenges due to the large size and high resolution of satellite images. This study presented in [3] investigates the use of a Quanvolutional pre-processing to enhance the capability of the Attention U-Net model in the building segmentation.

The results, summarized in Table 3.10, indicate that integrating quanvolutional operators with the Attention U-Net architecture significantly reduces the number of trainable parameters while maintaining competitive performance in terms of Overall Accuracy (OA).

The classic Attention U-Net achieves the highest OA of 0.9491, but it requires 34.8 million parameters, which increases computational demands. In contrast, the quantum-assisted models achieve OA values between 0.9343 and 0.9384 with only 2.1 million parameters, a reduction of over 93%. An example of the quanvolved feature maps and the corresponding segmentation prediction, compared with the GT on the test set, is shown in Fig. 3.12.




Figure 3.12: Comparison of quanvolved feature maps, prediction, and corresponding GT for a random sample from the test set.

Three different circuit configurations were evaluated for the Quanvolution-enhanced models: *Strongly Entangled Circuit*, *Random Circuit* and *Basic Entangled Circuit*. This last circuit reached the highest OA

among the quantum models at 0.9384, approaching the classic Attention U-Net’s performance with significantly fewer trainable parameters.

Table 3.10: Performance comparison between the standard Attention U-Net and the Quantum-assisted Attention U-Net using different quantum configurations.

| Model Type            | Circuit            | OA            | #Parameters  | #Qubits | Kernel |
|-----------------------|--------------------|---------------|--------------|---------|--------|
| Att. U-Net            | -                  | 0.9491        | 34.8 million | -       | -      |
| Quan4EO + Att. U-Net  | Strongly Entangled | 0.9373        | 2.1 million  | 9       | 3x3    |
| Quanv4EO + Att. U-Net | Random             | 0.9343        | 2.1 million  | 9       | 3x3    |
| Quanv4EO + Att. U-Net | Basic Entangled    | <b>0.9384</b> | 2.1 million  | 9       | 3x3    |

This study demonstrates that integrating Quanvolution into an Attention U-Net pipeline enables competitive building-segmentation performance on Sentinel-1 SAR data while reducing trainable parameters by over 90%, highlighting Quanvolution as an efficient and scalable alternative to fully classical models for large-scale urban mapping.

**Trainable quanvolution** Building on the Quanv4EO framework, the study presented in [2] introduces a trainable quanvolutional layer within a hybrid neural network, enabling dynamic parameter optimization during training. Encouraging insights arise from the comparison between the frozen and trainable configurations of the Quanv4EO layer.

Table 3.11: Performance comparison of models (5 epochs). The kernel size value is the same for all the models and it is set to  $3 \times 3$ . The number of qubits for the quantum models is set to 9.

| Model                     | # Classic Parameters | # Quantum Parameters | Accuracy     |
|---------------------------|----------------------|----------------------|--------------|
| CNN                       | $\sim 96k$           | -                    | 0.943        |
| Reduced CNN               | $\sim 6k$            | -                    | 0.935        |
| Quanv4EO [1]              | $\sim 10k$           | 27 (frozen)          | 0.942        |
| <b>Trainable Quanv4EO</b> | $\sim 5k$            | 27 (trainable)       | <b>0.947</b> |

Although the number of quantum parameters is constant across both settings, enabling their joint optimization with classical weights (trainable configuration) yields a deeper reduction in the number of classical parameters, as reported in Tab. 3.11. However, since training a quantum circuit is computationally demanding when simulated on classical hardware, this experiment was performed in a reduced setting with two classes and input images resized to  $16 \times 16$  pixels. Despite this simplification, the results highlight the promising role of trainable quantum layers for future implementations on real quantum hardware.

### 3.6 Discussion and Conclusions

The conducted investigation into the realm of QML has revealed promising advancements in the domain of EO data processing. Through the utilization of quanvolutional layers, this line of research has introduced the **Quanv4EO framework**, a quanvolution method for (pre)processing multi-dimensional EO data. Quanv4EO represents a novel approach to feature extraction, leveraging the power of quantum circuits to enhance classification, segmentation, regression and filtering tasks. Indeed, the training



process, involving the optimization of quantum circuits or variational parameters, presents computational challenges. Nevertheless, the introduced methodology offers efficient access to high-dimensional Hilbert space, facilitating the manipulation of quantum states, measurements, and optimization algorithms.

The experiments conducted on both standard datasets like MNIST and FashionMNIST, as well as EO datasets such as EuroSAT, Sentinel-1 and  $\Phi$ Sat-2 data, demonstrate the potential of the proposed Quanv4EO framework. The results indicate high accuracy in multiple and different down-stream tasks, surpassing or matching SOTA models while maintaining significantly smaller model sizes. However, it is important to acknowledge that these findings are primarily derived from simulations and a limited scale of experiments, which restricts the scope of generalization beyond a proof of concept.

Looking ahead, the findings suggest avenues for future research, particularly in exploring the integration of quantum classifiers following the quanvolutional layers, potentially overcoming challenges associated with fully quantum approaches. Additionally, future work should extend the ongoing exploration of trainable quanvolutional filters, as their optimization may lead to further performance improvements, especially considering the already encouraging obtained results. Future studies could also explore the application of 3D quanvolution for data types with higher spectral dimensions, leveraging the flexibility of our implementation to extend beyond 2D preprocessing. From a theoretical point of view, 3D quanvolution has the potential to capture richer inter-channel relationships compared to 2D quanvolution. However, it also comes with increased computational requirements, making it more demanding in terms of quantum resources. Evaluating whether this trade-off results in significantly improved model performance remains an open question, which could be addressed in future studies. Furthermore, while angle encoding has proven effective in this study, future research could explore alternative encoding methods, such as amplitude encoding, to further optimize the representation of image patches in the quantum state. Finally, while the study contributes to advancing quantum-enhanced ML for EO data analysis, it remains essential to validate these concepts through larger-scale experiments and hardware-based implementations before asserting broader applicability.

## Bibliography

- [1] A. Sebastianelli, **F. Mauro**, G. Ciabatti, D. Spiller, B. Le Saux, P. Gamba, and S. Ullo, “Quanv4eo: empowering earth observation by means of quanvolutional neural networks,” *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
- [2] **F. Mauro**, F. De Falco, A. Ceschini, G. Meoni, A. Sebastianelli, M. Panella, P. Gamba, and S. L. Ullo, “Advancing Earth Observation with Trainable Quanvolutional Neural Networks for classification tasks,” *IGARSS 2025 - 2025 IEEE International Geoscience and Remote Sensing Symposium*, 2025.
- [3] L. Russo, **F. Mauro**, B. Memar, A. Sebastianelli, S. L. Ullo, and P. Gamba, “A quantum-assisted attention u-net for building segmentation over tunis using sentinel-1 data,” in *2025 Joint Urban Remote Sensing Event (JURSE)*. IEEE, 2025, pp. 1–4.
- [4] **F. Mauro**, A. Sebastianelli, M. P. Del Rosso, P. Gamba, and S. L. Ullo, “Qspecklefilter: A quantum machine learning approach for sar speckle filtering,” in *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, 2024, pp. 450–454.
- [5] **F. Mauro**, F. Razzano, P. Di Stasio, A. Sebastianelli, G. Meoni, G. Schirinzi, P. Gamba, and S. L. Ullo, “Quantum-enhanced water quality monitoring: Exploiting  $\phi$ sat-2 data with quanvolution,” *IEEE Geoscience and Remote Sensing Letters*, 2025.
- [6] M. Chi, A. Plaza, J. Benediktsson, Z. Sun, J. Shen, and Y. Zhu, “Big data for remote sensing: Challenges and opportunities,” *Proceedings of the IEEE*, vol. PP, pp. 1–13, 09 2016.
- [7] N. Barde, D. THAKUR, P. Bardapurkar, and S. Dalvi, “Consequences and limitations of conventional computers and their solutions through quantum computers,” *Leonardo Electronic Journal of Practices and Technologies*, vol. 10, pp. 161–171, 07 2011.
- [8] M. Henderson, S. Shakya, S. Pradhan, and T. Cook, “Quanvolutional Neural Networks: Powering Image Recognition with Quantum Circuits,” Apr. 2019, arXiv:1904.04767 [quant-ph]. [Online]. Available: <http://arxiv.org/abs/1904.04767>
- [9] D. A. Zaidenberg, A. Sebastianelli, D. Spiller, B. Le Saux, and S. L. Ullo, “Advantages and bottlenecks of quantum machine learning for remote sensing,” in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 2021, pp. 5680–5683.
- [10] A. Sebastianelli, D. A. Zaidenberg, D. Spiller, B. Le Saux, and S. L. Ullo, “On circuit-based hybrid quantum neural networks for remote sensing imagery classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 565–580, 2021.
- [11] M. Henderson, S. Shakya, S. Pradhan, and T. Cook, “Quanvolutional neural networks: powering image recognition with quantum circuits,” *Quantum Machine Intelligence*, vol. 2, no. 1, pp. 1–9, 2020.
- [12] D. Mattern, D. Martyniuk, H. Willems, F. Bergmann, and A. Paschke, “Variational quanvolutional neural networks with enhanced image encoding,” *arXiv preprint arXiv:2106.07327*, 2021.
- [13] L. Bugalho, E. Z. Cruzeiro, K. C. Chen, W. Dai, D. Englund, and Y. Omar, “Resource-efficient simulation of noisy quantum circuits and application to network-enabled QRAM optimization,” *npj Quantum Information*, vol. 9, no. 1, p. 105, 2023.

- 
- [14] C. Onuoha, J. Flaherty, T. Nguyen, and T. C. Thang, “Data scanning methods for quantum-classical interface in quanvolutional neural networks,” in *2022 IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia)*. IEEE, 2022, pp. 1–4.
  - [15] K. Sooksatra, P. Rivas, and J. Orduz, “Evaluating accuracy and adversarial robustness of quanvolutional neural networks,” in *2021 International Conference on Computational Science and Computational Intelligence (CSCI)*. IEEE, 2021, pp. 152–157.
  - [16] H. Baek, W. J. Yun, and J. Kim, “Scalable quantum convolutional neural networks,” *arXiv preprint arXiv:2209.12372*, 2022.
  - [17] A. P. Patil, S. Pandey, N. D. Kasat, S. Modi, S. Raj, and R. Kulkarni, “Implementation of handwritten character recognition using quanvolutional neural network,” in *2022 IEEE North Karnataka Subsection Flagship International Conference (NKCon)*. IEEE, 2022, pp. 1–4.
  - [18] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, “Quantum convolutional neural networks for high energy physics data analysis,” *Physical Review Research*, vol. 4, no. 1, p. 013231, 2022.
  - [19] A. Matic, M. Monnet, J. M. Lorenz, B. Schachtner, and T. Messerer, “Quantum-classical convolutional neural networks in radiological image classification,” in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2022, pp. 56–66.
  - [20] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
  - [21] H. Xiao, K. Rasul, and R. Vollgraf, “Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms,” *arXiv preprint arXiv:1708.07747*, 2017.
  - [22] A. Sebastianelli, M. P. Del Rosso, S. L. Ullo, and P. Gamba, “A speckle filter for sar sentinel-1 grd data based on residual convolutional neural networks,” 2022.
  - [23] P. Helber, B. Bischke, A. Dengel, and D. Borth, “Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
  - [24] F. Razzano, P. Di Stasio, F. Mauro, G. Meoni, M. Esposito, G. Schirinzi, and S. L. Ullo, “AI techniques for near Real-Time monitoring of contaminants in coastal waters on board future  $\Phi$ sat-2 Mission,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 16 755–16 766, 2024.
  - [25] F. De Falco, A. Ceschini, A. Sebastianelli, B. L. Saux, and M. Panella, “Towards efficient quantum hybrid diffusion models,” *arXiv preprint arXiv:2402.16147*, 2024.
  - [26] M. Weigold, J. Barzen, F. Leymann, and M. Salm, “Data encoding patterns for quantum computing,” in *Proceedings of the 27th Conference on Pattern Languages of Programs*, 2020, pp. 1–11.
  - [27] M. Geiger, S. Spigler, A. Jacot, and M. Wyart, “Disentangling feature and lazy training in deep neural networks,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2020, no. 11, p. 113301, nov 2020.
  - [28] L. Chizat, E. Oyallon, and F. Bach, “On lazy training in differentiable programming,” 2020.
-

- 
- [29] J. Liu, F. Tacchino, J. R. Glick, L. Jiang, and A. Mezzacapo, “Representation learning via quantum neural tangent kernels,” *PRX Quantum*, vol. 3, p. 030323, Aug 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PRXQuantum.3.030323>
- [30] A. Ceschini, A. Carbone, A. Sebastianelli, M. Panella, and B. Le Saux, “On Hybrid Quantvolutional Neural Networks Optimization,” *Submitted to Springer Quantum Machine Intelligence*, 2024.
- [31] H. Suryotrisongko and Y. Musashi, “Evaluating hybrid quantum-classical deep learning for cybersecurity botnet DGA detection,” *Procedia Computer Science*, vol. 197, pp. 223–229, 2022.
- [32] M. Schuld and F. Petruccione, *Machine learning with quantum computers*. Springer, 2021.
- [33] T. Zhang, R. Ramakrishnan, and M. Livny, “Birch: A new data clustering algorithm and its applications,” *Data mining and knowledge discovery*, vol. 1, pp. 141–182, 1997.
- [34] F. Hutter, L. Kotthoff, and J. Vanschoren, *Automated machine learning: methods, systems, challenges*. Springer Nature, 2019.
- [35] H. Jin, Q. Song, and X. Hu, “Auto-keras: An efficient neural architecture search system,” in *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, 2019, pp. 1946–1956.
- [36] A. Byerly, T. Kalganova, and I. Dear, “No routing needed between capsules,” *Neurocomputing*, vol. 463, pp. 545–553, 2021.
- [37] D. Hirata and N. Takahashi, “Ensemble learning in cnn augmented with fully connected subnetworks,” *IEICE TRANSACTIONS on Information and Systems*, vol. 106, no. 7, pp. 1258–1261, 2023.
- [38] V. Mazzia, F. Salvetti, and M. Chiaberge, “Efficient-capsnet: Capsule network with self-attention routing,” *Scientific reports*, vol. 11, no. 1, p. 14634, 2021.
- [39] Y. Assiri, “Stochastic optimization of plain convolutional neural networks with simple methods,” *arXiv preprint arXiv:2001.08856*, 2020.
- [40] K. Kowsari, M. Heidarysafa, D. E. Brown, K. J. Meimandi, and L. E. Barnes, “Rmdl: Random multimodel deep learning for classification,” in *Proceedings of the 2nd international conference on information system and data mining*, 2018, pp. 19–28.
- [41] M. S. Tanveer, M. U. K. Khan, and C.-M. Kyung, “Fine-tuning darts for image classification,” in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 4789–4796.
- [42] P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur, “Sharpness-aware minimization for efficiently improving generalization,” *arXiv preprint arXiv:2010.01412*, 2020.
- [43] E. Harris, A. Marcu, M. Painter, M. Niranjana, A. Prügell-Bennett, and J. Hare, “Fmix: Enhancing mixed sample data augmentation,” *arXiv preprint arXiv:2002.12047*, 2020.
- [44] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang, “Random erasing data augmentation,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, 2020, pp. 13 001–13 008.
- [45] N. H. Phong and B. Ribeiro, “Rethinking recurrent neural networks and other improvements for image classification,” *arXiv preprint arXiv:2007.15161*, 2020.
-

- [46] A. Ceschini, A. Rosato, and M. Panella, “Hybrid quantum-classical recurrent neural networks for time series prediction,” in *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2022, pp. 1–8.
- [47] A. Sebastianelli, M. P. Del Rosso, S. L. Ullo, and P. Gamba, “On quantum hyperparameters selection in hybrid classifiers for earth observation data,” 2023.
- [48] M. Incudini, M. Grossi, A. Ceschini, A. Mandarino, M. Panella, S. Vallecorsa, and D. Windridge, “Resource saving via ensemble techniques for quantum neural networks,” *arXiv preprint arXiv:2303.11283*, 2023.
- [49] J. Li, D. Lin, Y. Wang, G. Xu, Y. Zhang, C. Ding, and Y. Zhou, “Deep discriminative representation learning with attention map for scene classification,” *Remote Sensing*, vol. 12, no. 9, p. 1366, 2020.
- [50] G. Sumbul, M. Charfuelan, B. Demir, and V. Markl, “Bigearthnet: A large-scale benchmark archive for remote sensing image understanding,” in *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2019, pp. 5901–5904.



## Chapter 4

# Quantum Graph Neural Networks

Building on the previous discussion of QuannvNNs, and acknowledging their inherent limitations, which will be examined in detail in this chapter, now the focus is moved to Quantum Graph Neural Network (QGNN)s. This shift reflects the need for quantum models capable of handling rich relational structures and long-range dependencies that extend beyond grid-based representations [1]. QGNNs represent a novel fusion of QC and GNNs, aimed at overcoming the computational and scalability challenges inherent in classical models that are powerful tools for analyzing data with complex relational structures but suffer from limitations such as high computational complexity and over-smoothing in large-scale applications [2]. QC, leveraging principles like superposition and entanglement, offers a pathway to enhanced computational capabilities. This chapter critically reviews the SOTA in QGNNs, exploring various architectures. Their applications across diverse fields such as high-energy physics, molecular chemistry, finance and earth sciences are discussed, highlighting the potential for quantum advantage. Additionally, the significant challenges faced by QGNNs are addressed, including noise, decoherence, and scalability issues, proposing potential strategies to mitigate these problems. This chapter aims to provide a foundational understanding of QGNNs, fostering further research and development in this promising interdisciplinary field.

### 4.1 Introduction

While QuannvNNs excel at extracting localized spatial features, they exhibit inherent limitations in capturing global-scale dependencies or large-scale geophysical events, such as El Niño fluctuations. Motivated by this observation, this chapter extends the investigation to Hybrid Quantum Graph Convolutional Neural Network (HQGCNN)s, which integrate a Graph CONvolutional Neural Network (GCNN) with a quantum MLP. Through this hybrid formulation, it was possible to demonstrate the model's capability to learn long-range spatial relationships between EO measurements and Earth phenomena evolving over time.

GNNs (GNNs) have emerged as powerful tools for analyzing and processing graph-structured data, a representation that appears in a wide array of domains such as social networks, recommender systems, cybersecurity, sensor networks, and natural language processing. Since their introduction [3–10], GNNs have shown remarkable performance in graph-related tasks thanks to their ability to model complex structural relationships encoded through nodes and edges. At their core, GNNs operate by iteratively aggregating information from neighboring nodes to generate expressive node-level representations dictated by the graph topology. However, this message-passing paradigm presents notable challenges. First, its computational complexity grows rapidly with graph size, leading to substantial memory consumption and long training times for large-scale graphs [11]. Second, GNNs are prone to *over-smoothing* [12],

whereby repeated aggregation causes node representations to become nearly indistinguishable, degrading their discriminative power.

Research on GNNs has become central within modern ML, and has increasingly converged with emerging computational paradigms. Among these, QC has attracted growing interest as a potential means to address intrinsic GNN limitations. QML models have demonstrated provable advantages over classical counterparts in regression and classification tasks involving both classical and quantum data [13–16]. More broadly, the integration of QC into ML workflows is motivated by the exponential advantages quantum processors may offer in analyzing certain classes of data, particularly those arising from quantum processes themselves [17]. Consequently, QNNs have emerged as a promising direction within DL, leveraging principles such as superposition and entanglement to process large-scale heterogeneous data more efficiently than classical neural networks [18, 19].

Yet, adopting a quantum framework introduces its own challenges, including limited coherence times and constraints on qubit connectivity. Despite the current infeasibility of fault-tolerant quantum computers, the rapid progress in NISQ devices has made VQCs a practical tool for exploiting quantum resources [20–23]. QNNs implemented through PQCs are particularly suitable for NISQ architectures; however, they still face issues such as barren plateaus—regions of exponentially vanishing gradient variance—that severely hinder trainability [24–26].

In an effort to mitigate these drawbacks and improve the applicability of neural architectures to complex systems, recent research has converged on the development of QGNNs. These models integrate the representational flexibility of GNNs with quantum computational primitives, offering a framework designed to capture the rich hierarchical and relational information inherent in graph-structured data. Their promise lies in leveraging QC to process highly complex or large-scale graphs more efficiently than classical GNNs. The growing interest in QGNNs is reflected in the increasing number of publications over recent years, as illustrated in Fig. 4.1, with a marked rise in contributions since 2021.

The literature considered in this chapter includes peer-reviewed journal articles, conference papers, and preprints published up to **December 2024**, corresponding to the submission date of the review paper [2].




Figure 4.1: Number of publications on QGNNs over the years.

However, despite the theoretical advantages of QC, there remains considerable skepticism regarding



whether current QGNN implementations can truly achieve quantum advantage, particularly given the limitations of present NISQ hardware and variational algorithms. Many real-world problems involve large, dense graphs that exceed the practical capabilities of today’s quantum processors. The aim of this chapter is to clarify whether—and under which conditions, QGNNs may outperform classical GNNs. By examining the current state of the field, the goal is to distinguish realistic expectations from hype and provide an evidence-based perspective on the feasibility, benefits, and limitations of quantum-enhanced graph learning. This analysis aims to inform future research directions and guide practitioners in deciding when and how quantum techniques may be meaningfully integrated into graph-based workflows. The structure of the chapter is as follows. Sect. 4.2 provides an overview of classical GNNs and their main application domains. Sect. 4.3 presents an in-depth analysis of QGNN architectures, including a taxonomy of existing models, their ansatz, hardware requirements, and integration with classical components. A comparison with classical GNNs is also provided, highlighting theoretical distinctions and practical implications. Sect. 4.4 discusses QGNN applications and evaluates their potential for quantum advantage. Sect. 4.5 examines the challenges and limitations of deploying QGNNs on NISQ devices. Finally, Sect. 4.6 concludes the chapter with a summary of findings and future research directions.

## 4.2 Preliminaries on GNNs

### 4.2.1 Definitions and Notations

Graphs are mathematical structures that represent a set of objects (nodes) and the pairwise relations (edges) existing between them. Formally, a graph  $\mathcal{G}$  can be defined as  $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ , where  $\mathcal{N}$  and  $\mathcal{E}$  are the set of nodes and edges respectively. Graphs with directional edges are directed graphs, conversely, they are undirected.

A graph with  $n$  nodes can be represented in matrix form through the Adjacency matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$ , where  $A_{ij} = 1$  if  $n_i$  and  $n_j$  are connected, else 0. The Degree matrix  $\mathbf{D}$  is a diagonal matrix which represents the number of edges connected to each node  $n$ , where  $D_{ii} = \sum_{j \in V} A_{ij}$ . Laplacian Matrix  $\mathbf{L}$  encodes the whole structure of the graph in matrix form:  $\mathbf{L} = \mathbf{D} - \mathbf{A}$ . Imbalanced weights may undesirably affect the matrix spectrum, leading to the need of normalization:  $\tilde{\mathbf{A}} = \mathbf{A} + \mathbf{I}$ ,  $\tilde{\mathbf{D}}$  is built from  $\tilde{\mathbf{A}}$  and finally  $\mathbf{L}_{\text{norm}} = \mathbf{D}^{-1/2} \tilde{\mathbf{L}} \mathbf{D}^{-1/2} = \mathbf{I} - \mathbf{D}^{-1/2} \tilde{\mathbf{A}} \mathbf{D}^{-1/2}$ .

A feature vector of arbitrary dimensions  $f$  can be linked to each node (or also edge): the node feature matrix  $\mathbf{X} \in \mathbb{R}^{n \times f}$  represent all the features of the graph, where  $X_{ij}$  represents the  $j$ -th feature of the  $i$ -th node.

### 4.2.2 Architecture of GNNs

Traditional ML struggles with graph data because it’s designed for structured data with consistent size and structure, like images or audio. Graphs, however, are unstructured, varying in form, size, and connections. GNNs address this by specializing in tasks at different graph levels, such as node, edge, and overall graph classification or regression, as illustrated in Fig. 4.2.

The primary component of GNNs lies in the use of Graph Convolutional Layer (GCL)s, which, when stacked in sequence, form a complete GNN. A GCL is built upon three main principles. Firstly, it leverages the locality of graph structures, performing convolutions on each node and its local neighborhood. Secondly, it is permutation equivariant, meaning it does not depend on the arbitrary ordering of graph features when converted into matrix form, which is crucial for enabling differentiable processing. Lastly, a GCL is differentiable, using a trainable weight matrix  $\mathbf{W} \in \mathbb{R}^{f \times f'}$  that works in conjunction with a message-passing stream of information between neighboring nodes, as defined by the adjacency matrix




Figure 4.2: Illustration of GNN and QGNN framework. In the classical path, graph data is encoded, optimized, and trained to produce the output. In the quantum path, quantum encoding is applied, followed by circuit optimization (purely quantum or hybrid), and measurements yield the output.

A. Mathematically, a GCL can be represented by:

$$f(\mathbf{X}, \mathbf{A}) = \varphi(\mathbf{A}(\mathbf{X}\mathbf{W})), \quad (4.1)$$

where  $\varphi$  is a differentiable activation function.

By stacking GCLs, node or edge representations are updated iteratively, integrating information from neighboring nodes by flowing through the graph structure. Specifically, stacking  $n$  GCLs allows information to reach nodes that are  $n$ -hops away.

Several types of GNNs have been developed: Recurrent Graph Neural Network (RecGNN)s, Convolutional Graph Neural Network (ConvGNN)s, Graph Autoencoder (GAE)s, Graph Transformers (GTs), and Spatio Temporal Graph Neural Network (STGNN)s. RecGNNs, pioneered by [27, 28], emphasize information diffusion through iterative updates. ConvGNNs, introduced by [5], extend CNNs to graph data, supporting effective information propagation among neighboring nodes. GAEs [29] aim to learn low-dimensional graph representations in an unsupervised manner: they consist of an encoder, which maps the input graph to a latent space, and a decoder, which attempts to reconstruct the graph from this latent representation. More recently, STGNNs, introduced by [30], address data that evolves over space and time by combining GCLs with temporal processing units, like recurrent or attention mechanisms, to capture dynamic patterns. GTs [31] integrate graph learning with transformer architectures, demonstrating versatile performance across diverse tasks. These advancements in GNNs provide powerful tools for analyzing complex graph-structured data, unlocking applications across multiple domains.

### 4.3 QGNNs Models

In this chapter QGNNs are classified into three main categories: Fully Quantum GNNs, Hybrid Quantum-Graph Neural Network (HQGNN)s, and Quantum-Assisted Graph Neural Network (QAGNN)s. This taxonomy, summarized in Fig. 4.3, reflects their progressive development in integrating QC with graph-based tasks, balancing opportunities with inherent challenges.




Figure 4.3: Proposed taxonomy for QGNNs. QGNN approaches are categorized into two primary groups: Fully/Hybrid Quantum and Quantum Assisted. The Fully/Hybrid Quantum group includes subcategories such as Quantum Graph Neural Networks, Recurrent, Convolutional, Time Series Convolutional, and Equivariant models, with references to key works in the field. The Quantum Assisted group highlights methods leveraging quantum-classical hybrid architectures for specific applications. Each category is annotated with relevant citations to provide context and background.

Fully Quantum GNNs perform all processing stages within the quantum domain, maximizing the potential of QC by eliminating classical-quantum transitions. While theoretically promising, their practical scalability is currently constrained by hardware limitations, including noise, limited qubit counts, and shallow circuit depths.

To mitigate these constraints while looking for quantum advantages, HQGNNs have emerged as an alternative. These models tightly integrate classical and quantum layers directly within the GNN structure. By embedding quantum operations into core GNN functionalities, HQGNNs selectively employ quantum resources for tasks where they can offer unique benefits, while classical layers efficiently handle the rest. This approach balances quantum capabilities with classical strengths, providing a more feasible path for near-term applications.

In contrast, QAGNNs take a different approach. Rather than embedding quantum operations within the GNN structure, QAGNNs retain fully classical GNN layers and utilize quantum components externally as auxiliary tools. Quantum resources are applied either before the GNN for preprocessing input data, or after for refining outputs or feature representations. Unlike HQGNNs, where QC is integral to the GNN architecture, QAGNNs use quantum capabilities to enhance specific stages without altering the classical nature of the network.

In the following sections, these approaches will be deeper presented, examining their architectural designs, the quantum ansatzes employed, and their potential to advance graph-based QML. It is also highlighted the main differences between GNNs and QGNNs in the reported Table 4.1.

#### 4.3.1 Fully and Hybrid Quantum

The integration of QC with GNNs has led to two main approaches: Fully QGNNs and HQGNNs.

Fully QGNNs process graph data entirely within the quantum domain using PQCs to encode structures and relationships into Hilbert space, uncovering patterns challenging for classical methods. Introduced

Table 4.1: Comparison of GNNs and QGNNs.

| Property                 |      | Graph Neural Network (GNN)                                       | Quantum Graph Neural Network (QGNN)                                 |
|--------------------------|------|------------------------------------------------------------------|---------------------------------------------------------------------|
| Data Representation      |      | Represents graphs using vectors and weighted matrices            | Represents graphs through quantum states and operators              |
| Update Mechanism         |      | Relies on convolutions and node aggregation                      | Utilizes quantum evolutions and quantum gates                       |
| Computational Complexity | Com- | Performance depends on graph density and model depth             | Offers potential speedups using quantum computers                   |
| Scalability              |      | Faces oversmoothing challenges in very deep graphs               | Limited by the current state of quantum technology                  |
| Advantages               |      | Unconstrained by specialized hardware requirements               | Enables parallelism and efficient handling of high-complexity tasks |
| Limitations              |      | Sensitive to graph size with convergence issues in deep networks | Encounters challenges associated with quantum hardware reliability  |

by Verdon et al. [32] and inspired by QAOA techniques, these models face scalability issues due to noise, limited qubits, and shallow circuits [33, 34].

HQGNNs combine quantum and classical layers, selectively applying quantum resources to enhance feature mappings or graph relationships. For example, [35] demonstrated a hybrid approach using ansatzes such as MPS and Tree Tensor Network (TTN)s, improving expressivity and mitigating barren plateaus. Other studies confirmed HQGNNs' balance between scalability and performance, making them suitable for near-term applications [36, 37].

Both approaches leverage quantum advantages for graph tasks, differing mainly in their strategies: Fully QGNNs emphasize purely quantum implementations despite hardware constraints, whereas HQGNNs pragmatically integrate classical methods to enhance scalability and practical feasibility on NISQ devices.

### Fully and Hybrid Quantum

**Quantum Graph Neural Networks** A QGNN that employs a quantum circuit's ansatz to learn and make predictions about the underlying graph is considered a groundbreaking approach in the field of QML. The underlying idea is to leverage Hilbert space to embed data and thus introduce a richer feature map, with characteristics that are hard to access for classical methods. Quantum dynamics, in the meantime, allows to explore and capture complex patterns and relationships within the graph that might be infeasible for classical GNNs.

QGNNs were first introduced in 2019 by [32], where the authors presented the concept of Networked Quantum Systems. They consider a graph  $G = (V, E)$ , with  $V$  as the set of nodes and  $E$  as the set of edges. Each node  $v \in V$  is associated with a quantum subsystem defined by a Hilbert space  $\mathcal{H}_v$ , resulting in a global Hilbert space  $\mathcal{H}_V \equiv \bigotimes_{v \in V} \mathcal{H}_v$ . Nodes or subgraphs correspond to one or more qubits in this representation. The same holds for the edges  $e \in E$ , although the authors do not consider them in their work:  $H_E \equiv \bigotimes_{e \in E} H_e$ , with  $H_E \otimes H_V$  being the total Hilbert space for the graph.

These QGNNs are inherently inspired by the structure of the QAOA, leveraging quantum circuits to en-

code problem-specific Hamiltonians that represent the graph structure. The evolution of the quantum state follows techniques such as the Trotter-Suzuki decomposition to approximate the Hamiltonian dynamics and principles of adiabaticity, ensuring a gradual evolution toward the ground state. This design enables QGNNs to capture complex graph features and optimize solutions in ways that classical GNNs cannot.

As proposed in [32], the most general QGNN ansatz is a PQC on a network that consists of a sequence of  $Q$  different Hamiltonian evolutions, with the whole sequence repeated  $P$  times:

$$U_{\text{QGNN}}(\boldsymbol{\eta}, \boldsymbol{\theta}) = \prod_{p=1}^P \prod_{q=1}^Q e^{-i\eta_{pq} \hat{H}_q(\boldsymbol{\theta})}, \quad (4.2)$$

where the product is time-ordered, the  $\boldsymbol{\eta}$  and  $\boldsymbol{\theta}$  are variational (trainable) parameters, and the Hamiltonians  $\hat{H}_q(\boldsymbol{\theta})$  can generally be any parameterized Hamiltonians whose topology of interactions is that of the problem graph:

$$\begin{aligned} \hat{H}_q(\boldsymbol{\theta}) \equiv & \sum_{(j,k) \in E} \sum_{r \in \mathcal{J}_{jk}} W_{qrjk} \hat{O}_j^{(qr)} \otimes \hat{P}_k^{(qr)} + \\ & \sum_{v \in V} \sum_{r \in \mathcal{J}_v} B_{qrv} \hat{R}_j^{(qv)}. \end{aligned} \quad (4.3)$$

Here, the  $W_{qrjk}$  and  $B_{qrv}$  are real-valued coefficients which can generally be independent trainable parameters, forming a collection  $\boldsymbol{\theta} \equiv \cup_{q,j,k,r} \{W_{qrjk}\} \cup_{q,v,r} \{B_{qrv}\}$ . The operators  $\hat{R}_j^{(qv)}$ ,  $\hat{O}_j^{(qr)}$ , and  $\hat{P}_j^{(qr)}$  are Hermitian operators which act on the Hilbert space of the  $j$ -th node of the graph. The sets  $\mathcal{J}_{jk}$  and  $\mathcal{J}_v$  are index sets for the terms corresponding to the edges and nodes, respectively. To make compilation easier, one can enforce that the terms of a given Hamiltonian  $\hat{H}_q$  commute with one another, but different  $\hat{H}_q$  need not commute.

### Quantum Recurrent GNNs

Quantum Recurrent Graph Neural Networks (QRecGNNs) are advanced architectures designed to model temporal dependencies within graph-structured data using QC principles. By extending classical Recurrent Neural Network (RNN)s to the quantum domain, QRecGNNs leverage PQCs and Hamiltonian dynamics to process temporal and relational graph information.

The foundational concept of QRecGNNs was introduced by [32], who proposed a fully quantum architecture based on Hamiltonian evolution. In this approach, the QRecGNNs process all data within the quantum domain, tying parameters across time steps to achieve recurrence. The general ansatz for QRecGNNs is expressed as:

$$U_{\text{QRecGNN}}(\boldsymbol{\eta}, \boldsymbol{\theta}) = \prod_{p=1}^P \prod_{q=1}^Q e^{-i\eta_{pq} \hat{H}_q(\boldsymbol{\theta})}, \quad (4.4)$$

where  $\eta_q$  are temporal parameters,  $\hat{H}_q$  are parameterized Hamiltonians encoding the graph topology, and  $\boldsymbol{\theta}$  are trainable parameters. This fully quantum design enables QRecGNNs to simulate continuous quantum dynamics, capturing both temporal and structural dependencies in a unified framework. Inspired by the Trotter-Suzuki decomposition and QAOA, exemplifies a physics-aligned approach suitable for tasks requiring quantum coherence.

Subsequent works have proposed hybrid adaptations of s, combining quantum and classical elements to address the scalability challenges of fully quantum models. For example, [38] introduced a hybrid that leverages Pauli-rotation encoding and Quantum Fourier Transform (QFT) to encode complex graph

relationships. While quantum circuits handle data encoding and temporal evolution, classical optimization tools, such as those provided by PennyLane, are used to train the parameters. This hybrid structure maintains the advantages of quantum computation while relying on classical resources for optimization and scalability.

Another hybrid approach was proposed by [39], who embedded PQC's within GNN layers to process graph features as quantum states. Temporal dependencies were modeled using unitary transformations in the quantum domain, while sequence updates relied on classical RNN layers. This architecture exemplifies a partially quantum design, where quantum operations enhance specific aspects of the model without fully replacing classical components.

A key application of Ns is the training of Ising Hamiltonians, as demonstrated in [32, 39]. Fully quantum implementations leverage VQE to prepare initial quantum states and use fidelity-based loss functions to train the parameters. For example, the loss function for QRecGNNs evaluating the fidelity between evolved and target states is:

$$L(\delta, \theta) = 1 - \frac{1}{N} \sum_{j=1}^N |\langle \psi_t | U(\delta, \theta) | \psi_0 \rangle|^2, \quad (4.5)$$

where  $N$  is the batch size,  $\delta$  represents time parameters, and  $|\psi_0\rangle$  is the initial state.

QRecGNNs offer a versatile framework for modeling temporal dynamics on graph-structured data, bridging quantum principles with recurrent processing. Fully quantum implementations, such as that of [32], maximize quantum coherence but face hardware constraints. Hybrid models, including those of [38] and [39], balance quantum advantages with classical scalability, providing practical solutions for near-term quantum devices.

### Quantum Convolutional GNNs

Quantum Convolutional Graph Neural Networks (QConvGNNs) represent a rapidly evolving area of quantum-enhanced ML, specifically designed to perform convolution operations on graph-structured data. By leveraging quantum principles, these models aim to address the limitations of classical GNNs in capturing complex graph relationships. The approaches vary from fully quantum architectures to hybrid designs, balancing quantum operations with classical optimization to align with current hardware capabilities.

At their core, QConvGNNs use quantum systems to encode and process graph features, often representing nodes, edges, or adjacency relations as quantum states or Hamiltonians. The mathematical formalism underlying these models is the same as that presented for s by [32], where GCLs are implemented through PQC's based on Hamiltonian evolution aligned with graph topology. Specifically, the unitary operator for a convolutional layer follows the structure:

$$U_{\text{QGCNN}} = \prod_{p=1}^P \prod_{q=1}^Q e^{-i\eta_{pq} H_q(\theta)}, \quad (4.6)$$

where  $H_q(\theta)$  encodes the graph structure and interactions,  $\eta_{pq}$  determines the weight of the Hamiltonian evolution at each layer and sub-layer,  $P$  is the number of layers, and  $Q$  is the number of sub-layers within each layer.

However, unlike QRecGNNs, which impose parameter sharing across temporal iterations to capture sequential dependencies, QConvGNNs enforce permutation invariance over node indexes. This means that the Hamiltonian parameters are constrained to be globally shared across the graph rather than locally tuned per node, making the model more analogous to classical convolutional graph networks. This

design choice ensures that QConvGNNs effectively generalize across different graph structures without requiring node-specific parameterization, although at the cost of potentially reduced expressivity in highly heterogeneous graphs. While this approach provides a direct implementation of quantum graph convolutions, its scalability remains limited by current quantum hardware constraints.

Building on these foundations, researchers in [34] proposed another fully quantum model that encodes graph adjacency using Givens rotations. In this method, adjacency relations are transformed into unitary operations:

$$U(\theta) = \prod_i G(\theta_i), \quad (4.7)$$

where  $G(\theta_i)$  are Givens rotation gates applied to qubit pairs, parameterizing interaction strengths with  $\theta_i$ . While this approach emphasizes adjacency encoding, its scalability is similarly constrained by the growing number of qubits required for large graphs.

Hybrid QConvGNNs attempt to overcome these scalability issues by integrating quantum and classical components. Authors in [40] introduced the Quantum Spatial Graph Convolutional Neural Network (QSGCN), which uses quantum amplitude encoding for graph features and controlled-unitary operations for convolution and pooling. Classical optimization is employed to update parameters, while dimensionality reduction is achieved by collapsing quantum states. This hybrid design improves scalability but introduces additional complexity.

To further adapt to near-term quantum devices, [36] developed a NISQ-targeted using subgraph decomposition and modular adjacency encoding with Pauli rotations  $R_y(\theta) = e^{-i\theta\sigma_y/2}$ . This method reduces resource requirements but may lose some structural fidelity due to its modular approach.

Finally, [41] proposed a hybrid model integrating quantum walks with classical convolutional layers. Continuous-time quantum walks propagate information through the graph, with the average mixing matrix  $Q$  capturing vertex connectivity:

$$Q = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T e^{iHt} e^{-iHt} dt, \quad (4.8)$$

where  $H$  is the Hamiltonian, typically the adjacency matrix, and  $T$  denotes the total evolution time. This matrix enhances feature propagation, which is then processed by classical convolutional layers, allowing for practical integration with existing frameworks.

These diverse approaches highlight a spectrum of strategies for leveraging quantum principles in graph-structured learning. Fully quantum models excel in directly encoding graph structures but face hardware limitations, while hybrid models balance quantum and classical techniques to achieve greater scalability.

### Quantum Time Series Convolutional GNNs

Quantum Time Series Convolutional Graph Neural Network (TS-QGCNN)s were first introduced by [42] to leverage quantum mechanics for capturing complex temporal and spatial dependencies in time series analysis. In the temporal domain, this approach employs the Schrödinger equation to model the dynamics of time series data, representing the evolution of the data as a quantum state. This provides a closed-form solution that effectively captures periodic oscillations and fluctuations typical of time series patterns. The evolution of the state over time is described by a unitary operator  $U = e^{iHt}$ , where  $H$  is the Hamiltonian that governs the temporal dynamics of the system:

$$|\psi(t)\rangle = U|\psi(0)\rangle = e^{iHt}|\psi(0)\rangle, \quad (4.9)$$

where  $t$  is time. Through this modeling, the TS-QGCNN can handle periodicities and dynamic changes in the data, producing a stable temporal representation that aligns with recurring patterns and trends commonly found in time series.

In the *spatial domain*, the network applies a to process data with non-Euclidean structures, such as those found in complex networks or structured time series data. Each node in the network is represented as a quantum state, and connections between nodes are encoded with controlled quantum gates, which map adjacency relations directly into the quantum circuit. This quantum convolutional framework allows the model to efficiently capture dependencies across interconnected nodes, making it highly suitable for applications in graph-structured data analysis.

To capture both local and global patterns, the model employs parameterized quantum gates, including rotation gates  $R_z$  and  $R_y$ , which dynamically adjust interactions between nodes. The operation of these gates is expressed as:

$$U_j = e^{i\alpha} R_z(\theta) R_y(\phi) R_z(\gamma),$$

where  $\theta$ ,  $\phi$ , and  $\gamma$  are adjustable parameters that optimize the network's response to varying patterns within the data. These quantum operations introduce non-linear transformations essential for recognizing complex patterns, and ultimately provide predictive outputs after each convolutional layer.

By combining a Schrödinger-based temporal model with QConvGNNs for spatial relationships, the TS-QGCNN effectively captures intricate spatio-temporal dependencies, making it a powerful tool for dynamic and structured time series analysis. However, while representing an important class of QGNNs, further exploration of this domain is necessary to fully understand its efficiency and advantages.

### Equivariant QGNNs

An Equivariant Quantum Graph Neural Network (EQGNN) combines QC with geometric DL to process graph data while preserving symmetry properties (equivariance), such as node permutations, rotations, or translations. EQGNNs thus encode graph structure and features within quantum circuits, maintaining relational consistency under symmetry operations.

Mernyei et al. [43] introduced Equivariant Quantum Graph Circuits (EQGCs), where graph nodes and their features are encoded into quantum states through parameterized unitaries. These circuits apply node-permutation-equivariant unitaries dependent on the graph's adjacency matrix, followed by node-permutation-invariant measurements (e.g., averaging expectation values over all nodes).

In the SOTA two main implementations were proposed: Equivariant Hamiltonian Quantum Graph Circuits (EHQGCs) and Equivariantly Diagonalizable Unitary Quantum Graph Circuits (EDU-QGCs). EHQGCs rely on Hamiltonians with graph-like topology similar to [32], while EDU-QGCs, introduced by Ryu et al. [44], use special two-node diagonalizable unitaries (EDUs) composed of local and diagonal unitary operations. EDU-QGCs can approximate any real-valued function on bounded graphs and pass the 1-Weisfeiler-Lehman test, outperforming classical message-passing networks (MPNNs).

However, face scalability limitations due to large quantum circuit requirements. Ai et al. [33] addressed this by proposing a decomposition strategy, partitioning large graphs into smaller subgraphs handled separately by quantum circuits connected via CNOT gates, improving scalability.

In summary, EQGNNs provide an innovative approach combining equivariance-preserving quantum circuits with GNNs, overcoming certain classical model limitations. Nevertheless, scalability challenges persist, restricting EQGNNs mostly to small-scale graphs. Future research should focus on improving circuit scalability and optimizing architectures for larger, practical applications.

### 4.3.2 Quantum Assisted GNNs

QAGNNs typically consist of a combination of classical GNN layers and quantum circuits, where the quantum component acts as an auxiliary tool rather than being integrated directly into the GNN layers themselves. This quantum functionality is often applied either upstream, for preprocessing or encoding



input data, or downstream, for enhancing feature representations or refining predictions. In contrast, HQGNNs incorporate quantum operations directly within the GNN layers, leveraging quantum mechanics as an integral part of the message-passing or feature aggregation processes. This fundamental distinction highlights the difference in how quantum resources are employed in these two architectures. Several recent studies exemplify this. Mauro et al. [1] proposed a QAGNN combining a classical ConvGNN (from [45]) with a Quantum MLP, showing superior performance and significantly fewer training epochs compared to classical benchmarks.

Ray et al. [46] presented another architecture where a classical GNN generates embeddings subsequently processed by a downstream Variational Quantum Classifier. This hybrid setup demonstrated improved results, especially with end-to-end training, slightly outperforming purely classical GNN models.

Similarly, Vitz et al. [47] designed a QAGNN employing a classical GENConv layer for dimensionality reduction, followed by amplitude encoding into a quantum circuit. Despite competitive outcomes, this approach did not surpass classical counterparts, suggesting optimization is still needed to clearly demonstrate quantum advantages.

Overall, QAGNNs represent a pragmatic, near-term approach, potentially addressing current quantum hardware limitations (limited qubits, coherence times) by strategically applying quantum resources only where beneficial. Their strength lies in scalability, as most computational workloads remain classical, enabling the processing of larger graphs compared to fully quantum or HQGNN models. However, unlike fully quantum approaches, QAGNNs lack a robust theoretical foundation, making their long-term effectiveness and potential advantages still experimental and uncertain.

## 4.4 Applications and Potential for Quantum Advantage

QGNNs can be applied across a variety of domains, ranging from High-Energy Physics (HEP) and the analysis of complex systems to applications in molecular biology and chemistry. Furthermore, they have significant potential in areas directly related to quantum systems. This section explores specific models and implementations of QGNNs tailored to these applications, providing a comparison with the current SOTA in classical GNNs. Fig. 4.4 provides an overview of the distribution of QGNN applications across various domains, the methodological approaches employed, and the simulation platforms used, illustrating the dominant trends in current research.

In this Section, it is evaluated not only the potential advantages offered by QGNNs, but also their current limitations and challenges as they are introduced into increasingly realistic and complex scenarios. It is important to highlight that QGNNs do not always deliver a true quantum advantage, defined as scenarios where quantum algorithms demonstrably outperform classical algorithms by solving problems faster or more accurately than is feasible with classical computing alone [48]. Instead, QGNNs often provide quantum utility, which refers to the practical benefits and enhancements that QC can bring to existing computational tasks, even when these improvements do not fundamentally surpass classical methods, as defined in Table 4.2.

### 4.4.1 Earth Observation

In the field of EO, QGNNs have the potential to serve as a powerful tool for capturing global-scale relationships and time-dependent variations in climatic systems. A compelling example is the modeling of the Ocean Niño Index (ONI), which tracks the El Niño phenomenon, i.e. anomalous warming in the Pacific Ocean that has widespread climatic impacts [45, 92]. Predicting ONI fluctuations demands models capable of addressing spatially distant but temporally correlated events, a hallmark of complex




Figure 4.4: Distribution of QGNN categories across three main areas: (a) application domains of the analyzed works, including Quantum Hamiltonian Dynamics (QHD), finance, complex systems, Molecular Chemistry and Biology (MC&B), HEP, and other categories; (b) types of approaches used, distinguishing between hybrid, fully quantum, and quantum-assisted methods; (c) simulation platforms utilized, such as PennyLane, Qulacs, OriginQ, CirQ, Qiskit, TensorFlowQ, and unspecified cases.

Table 4.2: Summary of Applications and Potential Quantum Advantages in QGNNs.

| Application Domain              | Authors                    | Description                                                                                                                                                                | Potential Quantum Advantage                                                                                                                        |
|---------------------------------|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| High-Energy Physics             | [35, 49–65]                | Tasks: jet tagging, particle track reconstruction, particle identification. Handles sparse data and complex graph structures in high-collision experiments (e.g., HL-LHC). | Equivariant architectures improve robustness. Reduction parameters numbers and computational complexity. Quantum speedups observed in jet tagging. |
| Molecular Chemistry and Biology | [16, 44, 46, 47, 66–85]    | Tasks: molecular property prediction (e.g., HOMO-LUMO gap, intramolecular forces) or classification. Focus on geometric and quantum chemical properties.                   | Capturing symmetries and geometric properties efficiently. Reduction parameters numbers and achieving competitive performance.                     |
| Complex Systems                 | [1, 34, 35, 42, 45, 86–97] | Tasks: traffic prediction, climate modeling, fraud detection. Addresses spatio-temporal dependencies in large-scale systems.                                               | Improving accuracy and training efficiency (e.g., climate modeling). Achieving higher efficiency (e.g., fraud detection).                          |
| Quantum Mechanics               | [32, 39, 98–106]           | Tasks: Hamiltonian dynamics modeling, quantum state preparation (e.g., GHZ state), graph isomorphism, spectral clustering.                                                 | Efficient Hamiltonian reconstruction, reduced complexity in state preparation, and accurate spectral clustering with low qubit precision.          |

systems. The QAGNN, presented in [1], addresses these challenges by combining classical convolution techniques with quantum layers. This quantum assisted architecture not only improves predictive accuracy compared to SOTA classical models [45] but also reduces training time by an order of magnitude, completing tasks in just five epochs instead of fifty. In this case, the results clearly demonstrate a quantum advantage. However, given that this is a QAGNN, the quantum component remains extremely limited, consisting only of a VQC applied at the final stage. It would be valuable to further investigate how the integration of even a small quantum component into a classical model can yield such significant benefits in a task as complex as this one. These results underscore the potential of quantum-enhanced models to revolutionize environmental analytics by improving efficiency and scalability in ways unattainable with classical methods.

#### 4.4.2 High-Energy Physics

QGNNs are becoming essential in the field of HEP especially for tasks as jet tagging (identifying the likely parton-level particle from which a jet originated) and track reconstruction (tracing the trajectories of particles). These tasks require precise data extraction and processing from experiments in particle accelerator, such as those at the Large Hadron Collider (LHC) [49, 50]. GNNs initially emerged as the best suited architecture for handling complex graph structures inherent in particle jets and tracks, allowing effective analysis of interactions between particles [51, 52]. However, the upcoming High-Luminosity Large Hadron Collider (HL-LHC) upgrade [53] has driven a demand for fast-processing models that can handle both high collision rates and sparse data, further advancing the use of QGNNs in HEP [35].

The motivation behind adopting QGNNs lies in leveraging a potential quantum advantage that can not only improve accuracy but, more importantly, enable the handling of large-scale, sparse data with a possible speedup over traditional models. For jet tagging, two different architectures have been proposed.

The first [54] performs binary classification of jets originating from either quarks or gluons, using the Pythia 8 Quark and Gluon Jets for Energy Flow dataset [55]. It incorporates a model carefully designed to address this specific problem, constructing a fully parameterized Hamiltonian based on the adjacency matrix of the jet graph, which is then made permutation-equivariant, yielding an EQGNN model. Equivariance, in particular, makes the architecture more robust and effective, achieving approximately a 7% increase in Area Under the ROC Curve (AUC) compared to a classical Equivariant Graph Neural Network (EGNN) model with a similar parameter count.

The second architecture, proposed in [56], addresses jet discrimination in both binary (using the TOP dataset [57, 58]) and multiclass scenarios (with the jetNET dataset [59]), employing data reuploading to facilitate more complex quantum state interactions and learning through aggregation via a SUM operation, similar to classical Message Passing Graph Neural Network (MPGNN)s. Notably, the implementation reveals an immediate quantum advantage, as the proposed quantum architecture has a complexity of  $O(N)$  when trained on  $N$  particles, assuming the parametrized operators are sufficiently deep. This offers a polynomial speedup over classical models, which require  $O(N^2)$  instead. Moreover, the quantum model does exhibit more stable training compared to classical models, especially in the case of multiclass classification. However, in terms of error performance, it proves to be comparable to a classical model.

Another classification problem, similar to jet tagging, involves identifying particle types ( $\mu^+$ ,  $e^-$ ,  $\pi^+$ , and  $p$ ) from single-particle images generated through realistic simulations of detector response and particle interactions. This task is particularly important in the DUNE experiment [60], a long-baseline neutrino oscillation experiment, where the DUNE detector digitally records high-resolution images of particle interactions. The proposed hybrid quantum-classical architecture [61] for this task addresses the challenge of sparse data by flattening the original input images  $X$  and performing  $n$  successive matrix multiplications with the adjacency matrix  $A$ . The aggregated vector  $A^n X$  is then encoded into a quantum state, processed by two VQCs, and finally passed through a single-layer classical neural network. In this case, a quantum utility is observed, as the hybrid architecture matches or slightly outperforms comparable architectures (MLP, CNN, QuannvNN) while reducing parameters by over 50%. Additionally, another potential application is highlighted in [62], where two different QGNN models are proposed and tested on a relatively simple task: simulating the interactions of three particles confined within a box. The two models are inspired by the classical approach presented in [63], but differ in their design. One model is speculative and cannot be physically implemented on a quantum computer, while the other is feasible for practical implementation. This distinction arises primarily because the speculative model exploits quantum superposition states to propagate information within the architecture. Specifically, it accesses and stores the qubits' superposition amplitude states, using them classically. This approach requires an excessive number of measurements from the quantum computer, making it impractical. Both models are simulated and compared with their classical GNN counterpart in terms of loss and accuracy. On the one hand, all the considered models (the two QGNNs and the classical GNN) demonstrate high learning efficiency. Furthermore, the implementable QGNN achieves a quantum utility, improving the accuracy obtained by the classical model. On the other hand, the accuracy achieved indicates that the models still struggle to capture more complex behaviors, underscoring that while the introduction of QC into GNNs is promising, it still requires further refinement to achieve significantly better performance.

Lastly, QGNNs are increasingly applied to particle track reconstruction, driven by the need for faster processing. Current models handle present collision rates but scale more than quadratically, posing challenges for the high collision rates anticipated with the HL-LHC. In [35], a QGNN model is proposed that integrates quantum components and is designed to optimally account for the sparse nature and complex geometry of tracking detector data. This model begins with a single fully connected neural network layer, which expands the input data size to a hidden size of  $N_d$ . Graph features, where nodes

represent particle hits and edges represent track segment candidates, are then processed iteratively using two quantum circuits (called Edge and Node Networks) before a final classical layer. The hybrid model with attention passing is thus designed to explore a potential quantum advantage in reducing the hidden dimension  $N_d$  using quantum circuits. However, in terms of performance, it remains comparable to classical models. Further studies are certainly needed to fully explore all possible benefits, as while there are clear advantages, such as the ability to achieve a reduced hidden dimension, there are also significant challenges that need to be addressed. Among these is the difficulty of training larger VQCs models, as highlighted in [64, 65], partly due to the fact that the proposed QGNN architecture involves iterative usage of the two quantum circuits. This further necessitates keeping the models small, given the limitations imposed by simulation constraints.

In general, QGNNs are demonstrating a tangible advantage in their application to HEP. These advantages are evident in the reduction of parameters and model size compared to their classical counterparts [61], as well as in achieving a polynomial speedup [56]. These features align well with the demands of the field, which increasingly require the ability to process large volumes of data at higher speeds. It is important to note, however, that alongside these advantages, several challenges have been identified that need to be addressed for QGNNs to fully mature in this domain. Firstly, their performance in terms of error metrics remains comparable to that of classical models. Furthermore, scaling these models is increasingly difficult, particularly due to the challenges of working with large VQCs. The use of QGNNs in HEP is undoubtedly promising and, indeed, necessary. At the same time, there is a clear need for continued exploration and refinement of the existing models to overcome these limitations and unlock their full potential.

#### 4.4.3 Molecular Chemistry and Biology

In the domains of molecular chemistry and biology, the integration of GNNs has proven exceptionally impactful, particularly for tasks such as molecular classification and the prediction of specific molecular and material properties. Traditionally, laboratory methods like in-silico simulations were the primary means available to accelerate property prediction. However, as the size of the molecules or materials under study increases, the computational cost of these methods grows polynomially [66, 67, 107]. This limitation has prompted an increased reliance on data-driven approaches. Among various classical ML models, GNNs have emerged as the preferred solution due to their intrinsic capacity to process graph-structured data, ideal for representing the complex structures of molecules and materials [68–70].

Despite these advances, several challenges remain. Specifically, the scarcity of training data for certain rare compounds, alongside the necessity for precise feature extraction from extensive feature sets, has underscored the potential of QGNNs in these fields. QML has already demonstrated the advantage of requiring significantly less training data to build predictive models than traditional methods [16]. Moreover, the enhanced ability of quantum techniques to map high-dimensional feature spaces facilitates superior performance on complex data structures, which is particularly advantageous for specialized applications in chemistry and materials science [71].

One of the additional advantages of adopting QGNNs lies in the ability to design architectures that effectively capture the geometric nature of the problem. An example of this is the model proposed in [72], which surpasses comparable quantum and classical models in terms of interpretability, enabling the exploration and exploitation of symmetries inherent in the problem. Specifically, the objective here is to predict intramolecular forces and the total molecular energy in water molecules ( $H_2O$ ) using a model with 9 qubits, where each qubit represents one of the three spatial coordinates of the three water atoms. Moreover, the model also considers the interatomic distances within the molecule by encoding them, while the trainable layer part is built using parameterized rotations on each qubit and leveraging swap gates to enforce a degree of permutation invariance. Although this architecture does not yet

achieve state-of-the-art performance on the dataset considered and introduced in [73], it nonetheless demonstrates significant effectiveness in addressing its intended task.

Another example in which leveraging the geometry and symmetries of the problem is crucial to design an accurate architecture is presented in [44], where the task involves predicting the energy gap between the highest occupied and lowest unoccupied molecular orbitals (HOMO-LUMO gap) of small organic molecules using the QM9 dataset [74, 75]. The base architecture includes encoding the input features, specifically the node features, followed by the design of an EDU-QGCs module capable of processing link features, and the implementation of a readout function that is invariant to qubit permutations, ensuring robustness. Two potential enhancements are proposed for this base model. The first involves introducing a classical neural network aimed at determining the angles of rotation used for quantum encoding. This enhancement is particularly valuable because encoding a large number of features within a quantum model is challenging, yet having a rich feature set to encode improves the model’s performance. Thus, incorporating a classical auxiliary model serves as a practical compromise to reduce the complexity associated with high-dimensional quantum encoding while maintaining robust information encoding. The second proposed improvement involves adding a qubit representing a master node linked to the molecular graph. In this case, the accuracy achieved does not indicate a significant quantum advantage, as SOTA models such as ALIGNN [77] and SphereNet [76] achieve better accuracy. However, the utility lies in the fact that, when comparing performance relative to the number of trainable parameters, the quantum model outperforms classical models. Also in [46], where a quantum assisted approach is employed to tackle the 3-class binary classification problem in breast cancer subtyping using the BReAst Cancer Subtyping (BRACS) dataset [78], the importance of encoding and the dimensionality of the encoded vector for achieving accurate learning is highlighted. In this work, HACT-NET [78], a SOTA GNN optimized for hierarchical analysis of digital pathology tissue images, is combined with a VQC. The output of the final layer of the classical GNN is subsequently encoded into the VQC using various encoding methods, such as angle encoding and amplitude encoding. Depending on the encoding method employed and the number of qubits that make up the model, a classical layer is introduced between the final layer of the classical GNN and the VQC. This classical layer adjusts the dimensions of the GNN’s output to match the requirements of the VQC. A key finding from several ablation studies is that the quantum assisted model outperforms the classical model when the VQC is fed a higher-dimensional vector to encode. This underscores how the compression performed by the classical layer can result in a loss of information, while amplitude encoding of a higher-dimensional vector preserves the information required for accurate high-dimensional learning.

In [79], the challenge of encoding high-dimensional data, which is crucial for applications in molecular chemistry and biology, is addressed under the constraint of a limited qubit count. Here, a trainable mapping technique is introduced to preserve data point distances and mitigate information loss during mapping. Additionally, a decomposition strategy based on ego-graph decomposition [85] is proposed to maximize qubit efficiency: the entire graph is divided into ego-graphs, which are processed sequentially by quantum circuits. This approach opens promising avenues for leveraging quantum advantage in realistic settings. The egoQGNN architecture, which incorporates tensor product and unitary matrix representations to reduce the number of required model parameters, demonstrates competitive performance. Notably, when tested on graph classification with realistic datasets such as MUTAG [84], PTC [82], and PROTEINS [83], the architecture achieves strong results, using only 1.68% of the parameters of its classical counterparts.

Finally, another study explores QGNNs in the context of materials design [47], predicting formation energies of perovskites. The proposed quantum assisted architecture initially employs a GENConv model [81] to transform the input graph into an intermediate representation, which is then encoded and processed by a quantum circuit. The results were compared with classical models such as GENConv and XGBoost [80] and show that, in this case, the quantum model provides no added benefit in terms of

achieved metrics. In fact, it performs slightly worse, underscoring that classical models remain a robust and state-of-the-art approach, making them challenging to surpass.

Overall, while the use of GNNs has marked a significant advancement in applications within chemistry and molecular biology, enabling data-driven models that operate directly on molecular graph structures, the introduction of QGNNs aims to further enhance this progress. On the one hand, their use demonstrates clear utility, mainly because it allows for the development of models that leverage the geometry and symmetry of the problem, as shown in [44, 72]. However, in none of the cases studied so far has a true quantum advantage been achieved. This limitation is largely due to the need for large datasets and the requirement to consider a substantial number of features to achieve good performance in relevant applications. Current quantum technologies struggle to process and encode such extensive feature sets. To mitigate this limitation, many studies have incorporated classical preprocessing techniques. These include using auxiliary neural networks, as in [44], or employing classical decomposition strategies, as in [79]. Additionally, quantum assisted approaches have been adopted, maximizing the use of classical models while restricting the use of the quantum component [46, 47]. This highlights that quantum models still rely heavily on classical assistance and, on their own, are not yet capable of delivering superior performance, demonstrating the need for further research to address the outlined challenges and potentially surpass GNNs, which remain the SOTA approach for applications in this field.

#### 4.4.4 Complex Systems

Another promising area where QGNNs have seen interesting development is in the analysis of complex systems. Here, QGNNs are used to harness quantum mechanical properties in order to analyze these systems more effectively. Starting from their application in traffic analysis and prediction, where entanglement can be utilized for dynamic modeling of interactions between adjacent roads, the fields of application are vast. This extends even to financial systems, where superposition can be used to simultaneously examine multiple data points, accelerating tasks such as pattern recognition and data classification, crucial elements in detecting fraudulent activities.

In particular, one domain where QGNNs have shown significant promise is intelligent transportation, particularly in traffic congestion prediction. The rapid advancement of commercial 5G networks and the proliferation of the intelligent Internet of Things (IoT) [86, 87] have made it possible to leverage vast quantities of real-time data to monitor and manage traffic conditions more effectively [88]. By employing precise and adaptive modeling techniques, traffic systems, recognized as inherently complex due to their dynamic nature, can be better understood and controlled. These systems are characterized by non-Euclidean network topologies and intricate interdependencies, where the flow on adjacent roads influences and interacts with each other. Classical GNNs, while adept at processing the non-Euclidean structures of traffic networks, struggle to fully capture the dynamic interplay of spatial and temporal factors [89, 90].

This challenge has inspired the development of architectures such as the TS-QGCNN [42]. The strength of this architecture lies in its dual ability to capture and process both the temporal and spatial dynamics of traffic data. On the temporal side, the Schrödinger equation is employed to model the evolution of time series, offering a closed-form solution that effectively captures the periodic oscillations and fluctuations inherent in such datasets. Simultaneously, it processes the non-Euclidean structures of traffic data, providing a sophisticated representation of spatial relationships. However, while this architecture proves to be innovative because it can handle both the temporal and spatial aspects and is also robust to external perturbations, it does not demonstrate an advantage compared to other classical prediction models, which show better performance on the task in question.

In another example [91], the same architecture used in [35] for particle track reconstruction problems shows potential in industrial world applications, such as in the sector of reconstructing airplane routes.

Here, the basic idea remains the same, but the particles from [35] now correspond to airplanes. The objective is thus to optimize airport capacities and improve predictions of potential collisions by the Traffic Collision Avoidance System (TCAS). In fact, TCAS has become crucial for potential collisions after two aircraft collided, and in this context, the use of QGNNs highlights how quantum technology is now finding applications in the industry as well. However, in this case, there is a lack of experimental results that validate the use of this model for non-HEP applications, as well as an adequate comparison with classical models to highlight the potential advantages and limitations of the approach.

In financial systems, where detecting fraudulent activities requires uncovering hidden patterns in vast, interconnected transaction networks, QGNNs are proving transformative. Fraud often emerges through subtle, non-obvious relationships that classical GNNs struggle to detect, particularly as data scales exponentially [93]. The introduction of QGNNs enables the utilization of key quantum mechanical features: on one hand, quantum entanglement allows these models to capture the intricate dependencies within fraud networks, uncovering previously hidden connections [94], while superposition can provide benefits by accelerating certain tasks as data classification and pattern recognition. In the work by [95], a QGNN is proposed and compared with classical GNN model, specifically with GraphSAGE [96], conducting tests on the available credit card fraud detection dataset [97]. The proposed quantum model, that takes inspiration from the work of [34], demonstrated its superiority, achieving a remarkable 94.5% accuracy compared to 92.4% with its classical counterpart. Despite its simplicity, comprising only six qubits and approximately 200 parameters, the quantum models clearly represents a significant utility, highlighting the usefulness of QGNNs even in fraud detection applications.

Overall, the analysis of complex systems represents a significant area of application for QGNNs. This is because, as mentioned, certain key properties of quantum mechanics find immediate relevance: for instance, entanglement proves essential for modeling interconnections within data, while superposition can offer advantages in fields where large datasets must be processed efficiently and with speed. However, despite some innovative approaches, such as the one presented in [42], which enables both temporal and spatial data analysis, it remains challenging to claim a true quantum advantage in this context. While certain models, like those in [1] and [95], achieve good performance, the quantum component often relies heavily on classical support. As shown, in [1], where the quantum part is confined to the final layer of a predominantly classical model, the quantum contribution still requires classical assistance to handle the complexity of such applications effectively. In the most recent case examined ([95]), it is somewhat easier to discuss quantum utility, as there is an undeniable benefit; however, it remains modest, with only a 2% improvement in accuracy. Therefore, while QGNNs hold great promise for analyzing complex systems, further research is necessary to achieve a true quantum advantage in this field.

#### 4.4.5 Quantum Mechanics

A unique application of QGNNs, which has no direct counterpart in classical GNNs, is their ability to learn and estimate specific features and properties of quantum systems, that can subsequently be leveraged in various applications [98–100]. QGNNs leverage quantum mechanical principles, enabling them to model, predict, and analyze quantum phenomena with remarkable accuracy. This distinct capability makes QGNNs invaluable in fields like quantum information processing, quantum dynamics, and quantum sensing. For instance, the work presented in [32] demonstrates how QRecGNNs can learn the complex dynamics of an Ising spin system. By targeting an Ising Hamiltonian with a transverse field on a specific graph, QRecGNNs efficiently reconstructs the correct Hamiltonian parameters. This is achieved by evolving an initial quantum state  $|\psi_0\rangle$  under the ansatz and comparing the resulting state with the target state  $|\psi_T\rangle$  providing access to the output of quantum dynamics at multiple time intervals. Also in [39], it is highlighted that the Ising model Hamiltonian can be trained well via QRecGNNs,



demonstrating great implications since the Ising model emulates the QUBO problem, which is effective in modeling non-deterministic polynomial problems [101, 102].

Beyond modeling quantum Hamiltonian dynamics, QGNNs excel in tasks such as state preparation. A notable example, shown in [32], is the use of  $s$  to prepare the Greenberger–Horne–Zeiling (GHZ) state [103], a crucial factor for achieving quantum advantage in specific applications. In quantum sensing, which is a leading area of quantum technology, the GHZ state enables the precise estimation of parameters hidden in weak qubit phase rotation signals. Such signals often arise in interactions between artificial atoms and weak electric fields. Specifically, it is demonstrated that QGNNs offer an advantage by reducing the complexity of the problem while effectively achieving the desired state, as confirmed by practical tests.

Other applications, also discussed in [32], involve classical problems such as spectral clustering [104] and verifying whether two graphs are isomorphic, a task that has become a benchmark for evaluating the representation power of classical GNNs [105]. Both applications leverage the spectral properties of a matrix associated with the graph, aligning particularly well with the inherent capabilities of QGNNs. In particular, concerning spectral clustering, which entails finding low-frequency eigenvalues of the graph Laplacian and clustering node values to identify graph clusters, a QSGCN can be utilized. Notably, it is observed that using low-qubit precision produces meaningful results, suggesting that spectral clustering could become a promising application for near-term quantum devices. Finally, the process of verifying whether two graphs are isomorphic uses a QSGCN to sample the eigenvalues of the coupling Hamiltonian on both graphs through standard basis measurements of the qubits and computing the eigenvalue at each sample of the wavefunction. The Kolmogorov-Smirnov test [106] is then used to determine if the graphs are isomorphic. It is demonstrated that highly accurate performance can be achieved even with low sample sizes. This suggests that the QSGCN is fully capable of identifying graph isomorphism, as required for graph convolutional network benchmarks.

Overall, QGNNs represent a powerful tool for estimating specific features and properties of quantum systems, with potential applications spanning various domains. These include solving QUBO problems, quantum sensing, and quantum communication through the generation of specific states, such as the GHZ state, as well as addressing classical problems like spectral clustering and graph isomorphism verification. In all these scenarios, the structure and quantum mechanical properties underlying QGNNs align naturally with the requirements of the tasks, effectively demonstrating their utility. Specifically, it is possible to talk about quantum advantage in cases where problem complexity is reduced, such as in the generation of the GHZ state, and quantum utility in broader applications where QGNNs offer practical solutions. For instance, as demonstrated in the case of spectral clustering, excellent results are achieved even with low qubit precision, and in graph isomorphism, highly accurate performance can be obtained even with a limited number of samples. Among the various potential applications, including HEP, chemistry, and complex systems, the ability of QGNNs to effectively model and characterize quantum phenomena remains particularly well-suited and inherently aligned with their nature.

## 4.5 Challenges and Limitations

Despite significant advancements in QGNNs, numerous challenges remain that must be overcome to realize the potential for quantum speedups as these architectures are scaled, as described in Table 4.3. It is essential to thoroughly characterize the limitations of QGNNs to develop effective strategies for designing more robust algorithms, proving performance guarantees, and ultimately enhancing quantum hardware. Addressing these challenges will be of utmost importance in transitioning from theoretical potential to practical benefits in QC.

Table 4.3: Summary of Challenges, Descriptions, and Potential Solutions in QGNNs.

| Problem                        | Description                                                                                                                                                                                                                                            | Potential Solution                                                                                                                                                                                                                                                                  |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Noise and Decoherence          | Quantum hardware is inherently susceptible to noise and decoherence, leading to performance degradation in QGNNs. Local and correlated noise (e.g., dephasing, bit-flip, crosstalk) and limited error correction mechanisms exacerbate this challenge. | Quantum error mitigation techniques (e.g., zero-noise extrapolation, probabilistic error cancellation) and QGNN-specific noise-resilient circuit designs that leverage graph structure. Developing adaptive QGNN architectures to balance circuit depth and input graph properties. |
| Scalability and Adaptability   | The linear growth in qubit demand with graph size makes large-scale graph analysis impractical. QGNNs also face challenges in adapting to different graph topologies dynamically, requiring re-design of quantum circuits.                             | Use topology-aware quantum circuit synthesis, hybrid quantum-classical architectures, and circuit optimization techniques (e.g., gate pruning, reinforcement learning). Explore adaptable QGNN designs for dynamic graph structures.                                                |
| Lack of Performance Guarantees | QGNNs lack robust theoretical performance guarantees, unlike classical ML and QAOA-based algorithms. Theoretical analysis is hindered by diverse and heuristic QGNN architectures.                                                                     | Develop frameworks for theoretical guarantees (e.g., analyzing the role of entanglement and quantum correlations). Benchmark QGNNs against classical methods on various graph tasks.                                                                                                |
| Barren Plateaus                | Gradients in QGNNs vanish as the number of qubits or layers increases, hindering training. Noise and cost function design further exacerbate this phenomenon.                                                                                          | Use local cost functions, gradient-free optimizers, and layer-wise training to mitigate barren plateaus. Design architectures like EQGCs, which are less prone to barren plateaus.                                                                                                  |
| Initialization Strategies      | Random initialization often leads to barren plateaus. Limited literature attention on effective initialization strategies hinders performance improvements.                                                                                            | Use heuristic-based or warm-start initialization methods, Gaussian initialization, and progressive layer-wise training to ensure favorable starting points for optimization.                                                                                                        |

#### 4.5.1 Noise and Decoherence

Quantum hardware is inherently susceptible to noise and decoherence, which can significantly degrade the performance of quantum algorithms, including QGNNs. Noise, arising from environmental interactions, imperfect gate operations, and readout errors, can be classified into local (uncorrelated) and correlated types. Decoherence, the loss of quantum information due to these interactions, poses a particular threat to the fragile quantum states used in QGNNs.

Recent studies have investigated the effects of local, uncorrelated noise on VQAs [108–111]. For example, authors in [112] emphasize that the inherent noise in near-term quantum devices introduces systematic bias in objective function evaluations. This bias necessitates a reevaluation of convergence analysis for classical optimization procedures in quantum contexts. The quantum system’s susceptibility to various local noise channels, such as dephasing, bit-flip, and depolarizing, further complicates this analysis. Correlated errors, including crosstalk [113], non-Markovian  $1/f$  noise [114], and interactions with environmental fluctuators [115], have been shown to be prevalent in NISQ devices, further complicating the implementation of QGNNs.

While quantum error correction (QEC) codes can mitigate hardware noise, they require significant qubit and gate operation overhead [116]. Quantum error mitigation techniques such as zero-noise extrapolation [117], probabilistic error cancellation [118] and their combination [119] have shown promise on NISQ devices, but face scalability and large-scale implementation challenges.

A significant gap in current research is the absence of comprehensive theoretical and empirical studies that characterize the performance degradation of QGNN as a function of noise strength and graph size. As quantum noise increases, it could be expected QGNN performance to degrade. However, the precise nature of this degradation remains unclear. The complexity of the input graph likely influences how severely noise affects QGNN performance: larger graphs may be more susceptible to noise-induced

errors, as they require longer circuits. Moreover, deeper circuits with more layers potentially offer better expressibility and graph representation capabilities, but they also provide more opportunities for errors to accumulate due to prolonged exposure to noise.

If someone had a quantitative understanding of these relationships, it would enable researchers to make informed design decisions about QGNN architecture for real-world applications, could facilitate the determination of optimal circuit depth for specific noise levels and graph sizes, and lead to the development of adaptive QGNN designs that dynamically adjust their structure based on input graph properties and quantum hardware noise characteristics. One reason for this research gap may be the current limitations in simulating and executing large QGNNs on actual quantum hardware, as illustrated in the following Sect. 4.5.2.

A primary strategy for noise mitigation is reducing gate count in compiled quantum circuits, as different hardware architectures impact circuit depth and connectivity. In this sense, developing QGNN-specific graph-to-circuit mapping and noise-resilient techniques that exploit graph data structure could be a promising research direction. This approach could leverage graph connectivity and redundancy for more efficient error correction strategies, potentially leading to feasible QGNN implementations on NISQ devices.

### 4.5.2 Scalability and Adaptability

Scalability is a critical challenge for QGNNs. As the size of the graph and the number of qubits increase, the complexity of the quantum circuits also grows, leading to greater demands on quantum hardware. Current QGNN approaches face the challenge of qubit demand growing linearly with graph size, making large-scale graph analysis on NISQ devices impractical.

One approach to addressing scalability is using topology aware quantum circuit synthesis for qubit-efficient encoding and mappings of graph data into quantum circuits [120, 121]. Unitary synthesis techniques offer optimal gate counts while respecting hardware-specific qubit topologies. Moreover, circuit optimization strategies, including gate pruning, reinforcement learning and parallelization, can help manage circuit depth and execution times [122–125].

Hybrid quantum-classical approaches also offer a practical solution for scalability. By offloading parts of the computation to classical layers, it is possible to handle larger graphs that would be otherwise unfeasible to process on quantum hardware alone. This approach has been successfully applied in [1, 35, 61].

### 4.5.3 Lack of Performance Guarantees

Unlike classical ML and QAOA-based algorithms, where performance can often be analyzed and bounded theoretically, QGNNs currently lack robust performance guarantees. This gap arises from the inherent challenges in theoretical analysis and the necessity for empirical validation to evaluate performance effectively. The diverse, fragmented and heuristic nature of QGNN architectures (see Sect. 4.3) further complicate this issue.

Recent works on VQAs have begun to partially address this gap. For example, several studies have provided bounds on the expressiveness of QNNs, suggesting conditions under which quantum models might outperform classical ones [14, 126, 127]. However, these results are preliminary and often depend on idealized assumptions about ansatz design, quantum hardware and noise levels.

A promising research direction is the development of theoretical frameworks that can provide more concrete performance guarantees for QGNNs. This includes understanding the role of entanglement and quantum correlations in enhancing the expressive power of QGNNs. Additionally, benchmarking

QGNNs against classical counterparts on a variety of graph-based tasks could provide empirical insights into their practical advantages and limitations.

#### 4.5.4 Barren Plateaus

Training variational QNNs, including QGNNs, is often significantly challenged by the barren plateau phenomenon [64]. Unlike classical neural networks, where gradients vanish as the number of layers increases, QNNs encounter this issue more broadly as the number of qubits grows. Despite [47] used a gradient-free optimizer for their QGNN, barren plateaus may hinder the training of quantum models designed for high-dimensional data, even when using gradient-free optimization algorithms [128]. The barren plateau problem is further exacerbated by quantum noise [129] and the design of the cost function [26]. Notably, there has been experimental evaluation of local cost functions compared to global cost functions in addressing the barren plateau issue within QGNNs [44]. One potential reason for the observed improvement is that optimization algorithms tend to encounter barren plateaus more frequently with global readout functions, as suggested by [26]. However, the difference in performance is minimal, and [44] provides limited insight into barren plateaus for QGNNs.

Moreover, the architecture of the quantum circuit plays a crucial role in the presence of barren plateaus. For instance, Equivariant Quantum Graph Convolutional Neural Network (EQGCNN)s proposed in [43] have been observed to scale effectively with model depth and are less prone to barren plateaus than other QNNs architecture. Nonetheless, these findings are based on circuits with only 6 to 10 qubits, while barren plateaus are more likely to emerge in deeper circuits, necessitating further empirical validation. Overall, while the barren plateau problem in QGNNs is widely recognized, it remains insufficiently addressed in the literature. Several training strategies and initialization techniques have been proposed to mitigate this issue [130, 131] (see Sect. 4.5.5), but further research and development are required to fully understand and resolve the challenges posed by barren plateaus in variational quantum graph-based models.

#### 4.5.5 Initialization Strategies

Quantum circuits' parameters initialization in QGNNs is a critical step that can significantly impact their performance. Random initialization can lead to barren plateaus [64], therefore effective initialization strategies are needed to provide good starting points for the parameters of quantum circuits.

Despite its importance, the topic of initialization is not extensively covered in the literature, and many papers give it minimal attention. For example, [44] compared with classical GNNs, using three different training runs with varied initial parameters for the classical models to find optimal weights. However, for the quantum and hybrid models, only a single training run was conducted, with all initial quantum weights set to 1. Such an evaluation approach is puzzling, given that the classical models' training results showed a significant dependence on initial weight values. The decision to use a single, uniform initialization for quantum models might be influenced by performance considerations or limitations specific to quantum systems, but it calls for further investigation and clarity. Additional evidence of the importance of initialization strategies is found in [43], where a model experienced a failure due to bad initialization, resulting in a poor 50% accuracy. Despite using only a small number of qubits (6 to 10), unfortunate starting points can lead to significant learning issues.

Heuristic-based initialization methods, such as initializing parameters based on classical solutions or using domain-specific knowledge, have been shown to improve convergence in VQAs [130]. Another approach is to use warm-start techniques, where the parameters are initialized close to known good solutions from related problems or smaller instances of the current problem [132]. Techniques such as Gaussian initialization [131], which uses Gaussian distributions to set initial parameters, and layer-

wise training [133], where parameters are initialized progressively from shallow to deeper layers, could provide promising alternatives. These approaches might help mitigate the barren plateau problem by ensuring more favorable starting conditions for the optimization process.

## 4.6 Conclusion

QGNNs remain a largely unexplored yet promising area of research, as highlighted by the aforementioned studies and preliminary experimental results. These findings suggest that QGNNs can represent a viable and, in some cases, superior alternative to classical GNNs, outperforming them in terms of performance, parameter efficiency, and problem complexity reduction.

Quantum utility has been repeatedly demonstrated in a variety of tasks, showing that different QGNN ansatzes can yield promising outcomes in fields such as quantum dynamics learning, quantum sensor network optimization, unsupervised graph clustering, supervised graph isomorphism classification, image classification, particle track reconstruction, and other graph-oriented problems.

However, this chapter also emphasizes several critical limitations and open challenges that are considered to be key directions for future development. As extensively discussed, current QGNN implementations still rely heavily on classical components. For instance, in HQGNNs classical layers are essential, particularly in the input encoding phase, to manage large-scale data that cannot yet be efficiently handled by quantum circuits. Similarly, in QAGNNs, the quantum layer often serves only as an auxiliary module embedded within a predominantly classical architecture.

While this hybrid approach has undoubtedly facilitated the application of QGNNs to increasingly complex and realistic tasks, it also constrains the possibility of achieving full quantum advantage. Among the most pressing limitations are the effects of noise and decoherence inherent in current quantum hardware. Additionally, the application of QGNNs to real-world problems involving large-scale graphs leads to a linear increase in qubit requirements, which, in turn, exacerbates issues such as barren plateaus, making the training of such architectures challenging, if not infeasible.

Looking ahead, one promising direction for overcoming these limitations and moving toward more fully quantum architectures is to draw inspiration from the QAOA. The latter is particularly compelling in that it encodes the structure of the problem (graph) directly into the quantum circuit, eliminating the need for classical preprocessing layers. Numerous QAOA variants have been proposed in the literature [134], and it is reasonable to explore whether these enhancements could be adapted for use in QGNN frameworks.

Moreover, several interesting properties of QAOA merit investigation within the context of QGNNs, such as parameter symmetries and the relation between graph's properties and the performance of QGNNs. Additional future work could involve generalizing QGNN architectures to include quantum degrees of freedom on graph edges, as well as extending the models to support more complex data structures like hypergraphs and simplicial complexes. Furthermore, in order to mitigate the effects of noise and decoherence, the topology of the quantum hardware, specifically, the qubit connectivity, should be carefully considered during the design of QGNN circuits.

Ultimately, the analysis and results presented in this chapter should not be interpreted as definitive statements on the current efficacy of QGNNs, nor should they discourage further exploration. Instead, this work aims to provide a comprehensive and critical foundation that encourages continued research into the vast potential of QGNNs and their practical applications.

## Bibliography

- [1] **F. Mauro**, A. Sebastianelli, B. L. Saux, P. Gamba, and S. L. Ullo, “A Hybrid MLP-Quantum approach in Graph Convolutional Neural Networks for Oceanic Nino Index (ONI) prediction,” *arXiv preprint arXiv:2401.16049*, 2024.
- [2] A. Ceschini, **F. Mauro**, F. De Falco, A. Sebastianelli, A. Verdone, A. Rosato, B. L. Saux, M. Panella, P. Gamba, and S. L. Ullo, “From Graphs to Qubits: A Critical Review of Quantum Graph Neural Networks,” *Submitted to Neural Computing and Applications (arXiv preprint arXiv:2408.06524)*, 2025.
- [3] F. Scarselli *et al.*, “The graph neural network model,” *IEEE Trans. on neural networks*, vol. 20, no. 1, 2008.
- [4] M. Gori *et al.*, “A new model for learning in graph domains,” in *Proc. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2, Montreal, Canada, 2005, pp. 729–734.
- [5] T. N. Kipf and M. Welling, “Semi-supervised classification with graph convolutional networks,” *arXiv preprint arXiv:1609.02907*, 2016.
- [6] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun, “Spectral Networks and Locally Connected Networks on Graphs,” in *Proc. of the 2nd International Conference on Learning Representations (ICLR)*, Banff, Canada, 2014, pp. 1–14.
- [7] W. L. Hamilton, R. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *Conference on Neural Information Processing Systems (NeurIPS)*, Long Beach, California, 2017, pp. 1025–1035.
- [8] D. Duvenaud, D. Maclaurin, J. Aguilera-Iparraguirre, R. Gómez-Bombarelli, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, “Convolutional networks on graphs for learning molecular fingerprints.” in *Conference on Neural Information Processing Systems (NeurIPS)*, Montreal, Canada, 2015, pp. 2224–2232.
- [9] F. Gama, A. G. Marques, G. Leus, and A. Ribeiro, “Convolutional neural network architectures for signals supported on graphs,” *IEEE Trans. on Signal Processing*, vol. 67, no. 4, 2019.
- [10] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, “Neural message passing for quantum chemistry,” in *International conference on machine learning*, Sydney, Australia, 2017, pp. 1263–1272.
- [11] S. Abadal, A. Jain, R. Guirado, J. López-Alonso, and E. Alarcón, “Computing graph neural networks: A survey from algorithms to accelerators,” *ACM Computing Surveys (CSUR)*, vol. 54, no. 9, pp. 1–38, 2021.
- [12] X. Zhang, Y. Xu, W. He, W. Guo, and L. Cui, “A comprehensive review of the oversmoothing in graph neural networks,” in *CCF Conference on Computer Supported Cooperative Work and Social Computing*. Springer, 2023, pp. 451–465.
- [13] A. Sebastianelli, D. A. Zaidenberg, D. Spiller, B. L. Saux, and S. L. Ullo, “On circuit-based hybrid quantum neural networks for remote sensing imagery classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 15, pp. 565–580, 2022.

- 
- [14] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, 2021.
  - [15] Y. Liu, S. Arunachalam, and K. Temme, “A rigorous and robust quantum speed-up in supervised machine learning,” *Nature Physics*, vol. 17, no. 9, pp. 1013–1017, 2021.
  - [16] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, “Generalization in quantum machine learning from few training data,” *Nature communications*, vol. 13, no. 1, p. 4919, 2022.
  - [17] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, “Quantum advantage in learning from experiments,” *Science*, vol. 376, no. 6598, pp. 1182–1186, 2022.
  - [18] F. Tacchino, C. Macchiavello, D. Gerace, and D. Bajoni, “An artificial neuron implemented on an actual quantum processor,” *npj Quantum Information*, vol. 5, no. 1, pp. 1–8, 2019.
  - [19] J. Tian, X. Sun, Y. Du, S. Zhao, Q. Liu, K. Zhang, W. Yi, W. Huang, C. Wang, X. Wu *et al.*, “Recent advances for quantum neural networks in generative learning,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 10, pp. 12 321–12 340, 2023.
  - [20] J. Preskill, “Quantum computing in the nisc era and beyond,” *Quantum*, vol. 2, p. 79, 2018.
  - [21] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, “Quantum circuit learning,” *Physical Review A*, vol. 98, no. 3, p. 032309, 2018.
  - [22] F. Tacchino, P. K. Barkoutsos, C. Macchiavello, D. Gerace, I. Tavernelli, and D. Bajoni, “Variational learning for quantum artificial neural networks,” in *2020 IEEE international conference on quantum computing and engineering (QCE)*. IEEE, 2020, pp. 130–136.
  - [23] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
  - [24] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, “Supervised learning with quantum-enhanced feature spaces,” *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
  - [25] Z. Holmes, N. Coble, A. T. Sornborger, and Y. Subaşı, “On nonlinear transformations in quantum computation,” *arXiv preprint arXiv:2112.12307*, 2021.
  - [26] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, “Cost function dependent barren plateaus in shallow parametrized quantum circuits,” *Nature communications*, vol. 12, no. 1, pp. 1–12, 2021.
  - [27] A. Sperduti and A. Starita, “Supervised neural networks for the classification of structures,” *IEEE Transactions on Neural Networks*, vol. 8, no. 3, pp. 714–735, 1997.
  - [28] A. Micheli, D. Sona, and A. Sperduti, “Contextual processing of structured data by recursive cascade correlation,” *IEEE Transactions on Neural Networks*, vol. 15, no. 6, pp. 1396–1410, 2004.
  - [29] T. N. Kipf and M. Welling, “Variational graph auto-encoders,” *arXiv preprint arXiv:1611.07308*, 2016.
-

- 
- [30] B. Yu, H. Yin, and Z. Zhu, “Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting,” *arXiv preprint arXiv:1709.04875*, 2017.
  - [31] J. Gong, Z. Teng, Q. Teng, H. Zhang, L. Du, S. Chen, M. Z. A. Bhuiyan, J. Li, M. Liu, and H. Ma, “Hierarchical graph transformer-based deep learning model for large-scale multi-label text classification,” *IEEE Access*, vol. 8, pp. 30 885–30 896, 2020.
  - [32] G. Verdon, T. McCourt, E. Luzhnica, V. Singh, S. Leichenauer, and J. Hidary, “Quantum graph neural networks,” *arXiv preprint arXiv:1909.12264*, 2019.
  - [33] X. Ai, Z. Zhang, L. Sun, J. Yan, and E. Hancock, “Decompositional quantum graph neural network,” *arXiv preprint arXiv:2201.05158*, 2022.
  - [34] Z. Hu, J. Li, Z. Pan, S. Zhou, L. Yang, C. Ding, O. Khan, T. Geng, and W. Jiang, “On the design of quantum graph convolutional neural network in the nisq-era and beyond,” in *2022 IEEE 40th International Conference on Computer Design (ICCD)*. IEEE, 2022, pp. 290–297.
  - [35] C. Tüysüz, C. Rieger, K. Novotny, B. Demirköz, D. Dobos, K. Potamianos, S. Vallecorsa, J.-R. Vlimant, and R. Forster, “Hybrid quantum classical graph neural networks for particle track reconstruction,” *Quantum Machine Intelligence*, vol. 3, pp. 1–20, 2021.
  - [36] S. M. A. Shah, H. Ge, S. A. Haider, M. Irshad, S. M. Noman, J. A. Meo, A. Ahmad, and T. Younas, “A quantum spatial graph convolutional network for text classification,” *Comput. Syst. Sci. Eng.*, vol. 36, no. 2, pp. 369–382, 2021.
  - [37] J. Zheng, Q. Gao, and Y. Lü, “Quantum graph convolutional neural networks,” in *2021 40th Chinese Control Conference (CCC)*. IEEE, 2021, pp. 6335–6340.
  - [38] P. Medisetty, L. K. Kumar Pallapothu, P. Chand Evuru, K. B. Prakash, V. M. S. Vulavalapudi, and G. S. Varma, “The quantum graph recurrent neural network,” in *2023 Second International Conference On Smart Technologies For Smart Nation (SmartTechCon)*, 2023, pp. 667–672.
  - [39] J. Choi, S. Oh, and J. Kim, “A tutorial on quantum graph recurrent neural network (qgrnn),” in *2021 International Conference on Information Networking (ICOIN)*. IEEE, 2021, pp. 46–49.
  - [40] J. Zheng, Q. Gao, M. Ogorzałek, J. Lü, and Y. Deng, “A Quantum Spatial Graph Convolutional Neural Network Model on Quantum Circuits,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–15, 2024.
  - [41] L. Bai, Y. Jiao, L. Cui, L. Rossi, Y. Wang, S. Y. Philip, and E. R. Hancock, “Learning graph convolutional networks based on quantum vertex information propagation,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 2, pp. 1747–1760, 2021.
  - [42] Z. Qu, X. Liu, and M. Zheng, “Temporal-spatial quantum graph convolutional neural network based on schrödinger approach for traffic congestion prediction,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 8, pp. 8677–8686, 2023.
  - [43] P. Mernyei, K. Meichanetzidis, and İ. İ. Ceylan, “Equivariant quantum graph circuits: Constructions for universal approximation over graphs,” *Quantum Machine Intelligence*, vol. 5, no. 1, p. 6, 2023.
  - [44] J.-Y. Ryu, E. Elala, and J.-K. K. Rhee, “Quantum graph neural network models for materials search,” *Materials*, vol. 16, no. 12, p. 4300, 2023.
-



- 
- [45] S. R. Cachay, E. Erickson, A. F. C. Bucker, E. Pokropek, W. Potosnak, S. Bire, S. Osei, and B. Lütjens, “The World as a Graph: Improving El Niño Forecasts with Graph Neural Networks,” *arXiv preprint arXiv:2104.05089*, 2021.
- [46] A. Ray, D. Madan, S. Patil, M. A. Rapsomaniki, and P. Pati, “Hybrid quantum-classical graph neural networks for tumor classification in digital pathology,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.11353>
- [47] M. Vitz, H. Mohammadbagherpoor, S. Sandeep, A. Vlastic, R. Padbury, and A. Pham, “Hybrid quantum graph neural network for molecular property prediction,” *ArXiv*, vol. abs/2405.05205, 2024. [Online]. Available: <https://api.semanticscholar.org/CorpusID:269626205>
- [48] A. W. Harrow and A. Montanaro, “Quantum computational supremacy,” *Nature*, vol. 549, no. 7671, p. 203–209, Sep. 2017. [Online]. Available: <http://dx.doi.org/10.1038/nature23458>
- [49] P. Baldi, P. Sadowski, and D. Whiteson, “Searching for Exotic Particles in High-Energy Physics with Deep Learning,” *Nature Commun.*, vol. 5, p. 4308, 2014.
- [50] D. Guest, K. Cranmer, and D. Whiteson, “Deep Learning and its Application to LHC Physics,” *Ann. Rev. Nucl. Part. Sci.*, vol. 68, pp. 161–181, 2018.
- [51] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban, and D. Whiteson, “Jet Flavor Classification in High-Energy Physics with Deep Neural Networks,” *Phys. Rev. D*, vol. 94, no. 11, p. 112002, 2016.
- [52] S. Zhang, H. Tong, J. Xu, and R. Maciejewski, “Graph convolutional networks: Algorithms, applications and open challenges,” in *International Conference on Computational Social Networks*, 2018. [Online]. Available: <https://api.semanticscholar.org/CorpusID:54526291>
- [53] G. Apollinari, O. Brüning, T. Nakamoto, and L. Rossi, “High luminosity large hadron collider hl-lhc,” 2015. [Online]. Available: <https://cds.cern.ch/record/2120673>
- [54] R. T. Forestano, M. Comajoan Cara, G. R. Dahale, Z. Dong, S. Gleyzer, D. Justice, K. Kong, T. Magorsch, K. T. Matchev, K. Matcheva *et al.*, “A comparison between invariant and equivariant classical and quantum graph neural networks,” *Axioms*, vol. 13, no. 3, p. 160, 2024.
- [55] P. T. Komiske, E. M. Metodiev, and J. Thaler, “Energy flow networks: deep sets for particle jets,” *Journal of High Energy Physics*, vol. 2019, no. 1, Jan. 2019. [Online]. Available: [http://dx.doi.org/10.1007/JHEP01\(2019\)121](http://dx.doi.org/10.1007/JHEP01(2019)121)
- [56] Y.-A. Chen and K.-F. Chen, “Jet discrimination with quantum complete graph neural network,” 2024. [Online]. Available: <https://arxiv.org/abs/2403.04990>
- [57] G. Kasieczka, T. Plehn, J. Thompson, and M. Russel, “Top quark tagging reference dataset,” Mar. 2019. [Online]. Available: <https://doi.org/10.5281/zenodo.2603256>
- [58] G. Kasieczka, T. Plehn, A. Butter, K. Cranmer, D. Debnath, B. M. Dillon, M. Fairbairn, D. A. Faroughy, W. Fedorko, C. Gay, L. Gouskos, J. F. Kamenik, P. T. Komiske, S. Leiss, A. Lister, S. Macaluso, E. M. Metodiev, L. Moore, B. Nachman, K. Nordström, J. Pearkes, H. Qu, Y. Rath, M. Rieger, D. Shih, J. M. Thompson, and S. Varma, “The Machine Learning landscape of top taggers,” *SciPost Phys.*, vol. 7, p. 014, 2019. [Online]. Available: <https://scipost.org/10.21468/SciPostPhys.7.1.014>
-

- 
- [59] R. Kansal, J. Duarte, H. Su, B. Orzari, T. Tomei, M. Pierini, M. Touranakou, J.-R. Vlimant, and D. Gunopulos, “Jetnet,” Aug. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6975118>
- [60] B. A. et al., “Deep underground neutrino experiment (dune), far detector technical design report, volume i: Introduction to dune,” 2020. [Online]. Available: <https://arxiv.org/abs/2002.02967>
- [61] S. Y.-C. Chen, T.-C. Wei, C. Zhang, H. Yu, and S. Yoo, “Hybrid quantum-classical graph convolutional network,” *arXiv preprint arXiv:2101.06189*, 2021.
- [62] B. Collis, S. Patel, D. Koch, M. Cutugno, L. Wessing, and P. M. Alsing, “Physics simulation via quantum graph neural network,” 2023. [Online]. Available: <https://arxiv.org/abs/2301.04702>
- [63] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. W. Battaglia, “Learning to simulate complex physics with graph networks,” *ArXiv*, vol. abs/2002.09405, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:211252550>
- [64] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature communications*, vol. 9, no. 1, p. 4812, 2018.
- [65] V. Leyton-Ortega, A. Perdomo-Ortiz, and O. Perdomo, “Robust implementation of generative modeling with parametrized quantum circuits,” 2019. [Online]. Available: <https://arxiv.org/abs/1901.08047>
- [66] L. E. Ratcliff, S. Mohr, G. Huhs, T. Deutsch, M. Masella, and L. Genovese, “Challenges in large scale quantum mechanical calculations,” *WIREs Computational Molecular Science*, vol. 7, no. 1, p. e1290, 2017. [Online]. Available: <https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/wcms.1290>
- [67] R. A. Friesner, “Ab initio quantum chemistry: methodology and applications.” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102 19, pp. 6648–53, 2005. [Online]. Available: <https://api.semanticscholar.org/CorpusID:25199073>
- [68] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, “A comprehensive survey on graph neural networks,” *IEEE transactions on neural networks and learning systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [69] T. Xie and J. C. Grossman, “Crystal graph convolutional neural networks for an accurate and interpretable prediction of material properties,” *Physical Review Letters*, vol. 120, no. 14, Apr. 2018. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevLett.120.145301>
- [70] C. Chen, W. Ye, Y. Zuo, C. Zheng, and S. P. Ong, “Graph networks as a universal machine learning framework for molecules and crystals,” *Chemistry of Materials*, vol. 31, no. 9, p. 3564–3572, Apr. 2019. [Online]. Available: <http://dx.doi.org/10.1021/acs.chemmater.9b01294>
- [71] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, “Power of data in quantum machine learning,” *Nature Communications*, vol. 12, no. 1, May 2021. [Online]. Available: <http://dx.doi.org/10.1038/s41467-021-22539-9>
- [72] S. Piperno, A. Ceschini, S. Y. Chang, M. Grossi, S. Vallecorsa, and M. Panella, “A study on quantum graph neural networks applied to molecular physics,” *arXiv preprint arXiv:2408.03427*, 2024.
-

- [73] O. Kiss, F. Tacchino, S. Vallecorsa, and I. Tavernelli, “Quantum neural networks force fields generation,” *Machine Learning: Science and Technology*, vol. 3, no. 3, p. 035004, jul 2022. [Online]. Available: <https://dx.doi.org/10.1088/2632-2153/ac7d3c>
- [74] L. Ruddigkeit, R. van Deursen, L. Blum, and J. Reymond, “Enumeration of 166 billion organic small molecules in the chemical universe database gdb-17,” *Journal of Chemical Information and Modeling*, vol. 52, no. 11, pp. 2864–2875, 2012.
- [75] R. Ramakrishnan, P. O. Dral, P. O. Dral, M. Rupp, and O. A. von Lilienfeld, “Quantum chemistry structures and properties of 134 kilo molecules,” *Scientific Data*, vol. 1, 2014. [Online]. Available: <https://api.semanticscholar.org/CorpusID:15367821>
- [76] Y. Liu, L. Wang, M. Liu, Y. Lin, X. Zhang, B. Oztekin, and S. Ji, “Spherical message passing for 3d molecular graphs,” in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=givsRXsOt9r>
- [77] K. Choudhary and B. DeCost, “Atomistic line graph neural network for improved materials property predictions,” *npj Computational Materials*, vol. 7, no. 1, Nov. 2021. [Online]. Available: <http://dx.doi.org/10.1038/s41524-021-00650-1>
- [78] P. Pati, G. Jaume, A. Foncubierta, F. Feroce, A. M. Anniciello, G. Scognamiglio, N. Brancati, M. Fiche, E. Dubruc, D. Riccio, M. D. Bonito, G. D. Pietro, G. Botti, J.-P. Thiran, M. Frucci, O. Goksel, and M. Gabrani, “Hierarchical graph representations in digital pathology,” 2021. [Online]. Available: <https://arxiv.org/abs/2102.11057>
- [79] X. Ai, Z. Zhang, L. Sun, J. Yan, and E. Hancock, “Towards quantum graph neural networks: An ego-graph learning approach,” 2024. [Online]. Available: <https://arxiv.org/abs/2201.05158>
- [80] T. Chen and C. Guestrin, “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD ’16. New York, NY, USA: Association for Computing Machinery, 2016, p. 785–794. [Online]. Available: <https://doi.org/10.1145/2939672.2939785>
- [81] G. Li, C. Xiong, A. Thabet, and B. Ghanem, “Deepergcn: All you need to train deeper gcns,” 2020. [Online]. Available: <https://arxiv.org/abs/2006.07739>
- [82] C. Helma, R. King, S. Kramer, and A. Srinivasan, “The predictive toxicology challenge 2000-2001,” *Bioinformatics*, vol. 17, no. 1, pp. 107–108, 2001.
- [83] R. Lab, “Protein function prediction via graph kernels,” *Oral Radiology*, vol. 6, no. 2, pp. 29–35, 1990.
- [84] A. K. Debnath, R. L. L. D. Compadre, G. Debnath, A. J. Shusterman, and C. Hansch, “Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity,” *Journal of Medicinal Chemistry*, vol. 34, no. 2, pp. 786–797, 1991.
- [85] D. Sandfelder, P. Vijayan, and W. L. Hamilton, “Ego-GNNs: Exploiting ego structures in graph neural networks,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Jun. 2021, p. 8523–8527. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP39728.2021.9414015>

- 
- [86] F. Zhang, G. Han, L. Liu, M. Martínez-García, and Y. Peng, “Joint optimization of cooperative edge caching and radio resource allocation in 5g-enabled massive iot networks,” *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 14 156–14 170, 2021.
- [87] G. Han, J. Tu, L. Liu, M. Martínez-García, and Y. Peng, “Anomaly detection based on multidimensional data processing for protecting vital devices in 6g-enabled massive iiot,” *IEEE Internet of Things Journal*, vol. 8, no. 7, pp. 5219–5229, 2021.
- [88] M. Veres and M. Moussa, “Deep learning for intelligent transportation systems: A survey of emerging trends,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 8, pp. 3152–3168, 2020.
- [89] B. Yan, G. Wang, J. Yu, X. Jin, and H. Zhang, “Spatial-temporal chebyshev graph neural network for traffic flow prediction in iot-based its,” *IEEE Internet of Things Journal*, vol. 9, no. 12, pp. 9266–9279, 2022.
- [90] L. Zhao, Y. Song, C. Zhang, Y. Liu, P. Wang, T. Lin, M. Deng, and H. Li, “T-gcn: A temporal graph convolutional network for traffic prediction,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 9, pp. 3848–3858, 2020.
- [91] K. Novotny, C. Tüysüz, C. Rieger, D. Dobos, K. Potamianos, B. Demirkoz, F. Carminati, S. Vallecorsa, J. R. Vlimant, and F. Fracas, “Quantum track reconstruction algorithms for non-hep applications,” *Proceedings of 40th International Conference on High Energy physics — PoS(ICHEP2020)*, 2021. [Online]. Available: <https://api.semanticscholar.org/CorpusID:236769635>
- [92] E. S. Sarachik and M. A. Cane, *The El Niño-Southern Oscillation Phenomenon*. Cambridge: Cambridge University Press, 2010.
- [93] S. Yu, J. Ren, S. Li, M. Naseriparsa, and F. Xia, “Graph learning for fake review detection,” *Frontiers in Artificial Intelligence*, vol. 5, 2022. [Online]. Available: <https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.922589>
- [94] M. Grossi, N. Ibrahim, V. Radescu, R. Loredó, K. Voigt, C. von Altrock, and A. Rudnik, “Mixed quantum-classical method for fraud detection with quantum feature selection,” *IEEE Transactions on Quantum Engineering*, vol. 3, p. 1–12, 2022. [Online]. Available: <http://dx.doi.org/10.1109/TQE.2022.3213474>
- [95] N. Innan, A. Sawaika, A. Dhor, S. Dutta, S. Thota, H. Gokal, N. Patel, M. A.-Z. Khan, I. Theodonis, and M. Bennai, “Financial fraud detection using quantum graph neural networks,” *Quantum Machine Intelligence*, vol. 6, no. 1, pp. 1–18, 2024.
- [96] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *Advances in Neural Information Processing Systems*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds., vol. 30. Red Hook, NY: Curran Associates, Inc., 2017.
- [97] ULB MLG, “Credit card fraud detection,” <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>, 2013.
- [98] H. J. Kimble, “The quantum internet,” *Nature*, vol. 453, no. 7198, pp. 1023–1030, 2008.
-

- 
- [99] K. Qian, Z. Eldredge, W. Ge, G. Pagano, C. Monroe, J. V. Porto, and A. V. Gorshkov, “Heisenberg-scaling measurement protocol for analytic functions with quantum sensor networks,” *Physical Review A*, vol. 100, no. 4, p. 042304, 2019.
  - [100] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, “Hamiltonian learning and certification using quantum resources,” *Physical review letters*, vol. 112, no. 19, p. 190501, 2014.
  - [101] C. Papalitsas, T. Andronikos, K. Giannakis, G. Theocharopoulou, and S. Fanarioti, “A qubo model for the traveling salesman problem with time windows,” *Algorithms*, vol. 12, no. 11, p. 224, 2019.
  - [102] M. Z. Alom, B. Van Essen, A. T. Moody, D. P. Widemann, and T. M. Taha, “Quadratic unconstrained binary optimization (qubo) on neuromorphic computing system,” in *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017, pp. 3922–3929.
  - [103] D. M. Greenberger, M. A. Horne, and A. Zeilinger, “Going beyond Bell’s theorem,” in *Bell’s theorem, quantum theory and conceptions of the universe*. Dordrecht, The Netherlands: Springer, 1989, pp. 69–72.
  - [104] A. Ng, M. Jordan, and Y. Weiss, “On spectral clustering: Analysis and an algorithm,” *Advances in neural information processing systems*, vol. 14, 2001.
  - [105] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How powerful are graph neural networks?” 2019. [Online]. Available: <https://arxiv.org/abs/1810.00826>
  - [106] H. W. Lilliefors, “On the kolmogorov-smirnov test for normality with mean and variance unknown,” *Journal of the American Statistical Association*, vol. 62, no. 318, pp. 399–402, 1967. [Online]. Available: <http://www.jstor.org/stable/2283970>
  - [107] S. Piperno, A. Ceschini, S. Y. Chang, M. Grossi, S. Vallecorsa, and M. Panella, “A study on quantum graph neural networks applied to molecular physics,” 2024. [Online]. Available: <https://arxiv.org/abs/2408.03427>
  - [108] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New J. Phys.*, vol. 18, no. 2, p. 023023, Feb. 2016.
  - [109] J. R. McClean, M. E. Kimchi-Schwartz, J. Carter, and W. A. De Jong, “Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states,” *Physical Review A*, vol. 95, no. 4, p. 042308, 2017.
  - [110] L. Gentini, A. Cuccoli, S. Pirandola, P. Verrucchi, and L. Banchi, “Noise-resilient variational hybrid quantum-classical optimization,” *Physical Review A*, vol. 102, no. 5, p. 052414, 2020.
  - [111] E. Fontana, N. Fitzpatrick, D. M. Ramo, R. Duncan, and I. Rungger, “Evaluating the noise resilience of variational quantum algorithms,” *Physical Review A*, vol. 104, no. 2, p. 022403, 2021.
  - [112] V. Kungurtsev, G. Korpas, J. Marecek, and E. Y. Zhu, “Iteration complexity of variational quantum algorithms,” *arXiv preprint arXiv:2209.10615*, 2022.
  - [113] I. Heinz and G. Burkard, “Crosstalk analysis for single-qubit and two-qubit gates in spin qubit arrays,” *Physical Review B*, vol. 104, no. 4, p. 045420, 2021.
  - [114] G. Burkard, “Non-markovian qubit dynamics in the presence of  $1/f$  noise,” *Physical Review B—Condensed Matter and Materials Physics*, vol. 79, no. 12, p. 125317, 2009.
-

- 
- [115] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, “Correlating decoherence in transmon qubits: Low frequency noise by single fluctuators,” *Physical review letters*, vol. 123, no. 19, p. 190502, 2019.
- [116] J. Roffe, “Quantum error correction: an introductory guide,” *Contemporary Physics*, vol. 60, no. 3, pp. 226–245, 2019.
- [117] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits,” *Physical review letters*, vol. 119, no. 18, p. 180509, 2017.
- [118] E. Van Den Berg, Z. K. Mineev, A. Kandala, and K. Temme, “Probabilistic error cancellation with sparse pauli-lindblad models on noisy quantum processors,” *Nature physics*, vol. 19, no. 8, pp. 1116–1121, 2023.
- [119] A. Mari, N. Shammah, and W. J. Zeng, “Extending quantum probabilistic error cancellation by noise scaling,” *Physical Review A*, vol. 104, no. 5, p. 052607, 2021.
- [120] M. G. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, and C. Iancu, “Towards optimal topology aware quantum circuit synthesis,” in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)*. IEEE, 2020, pp. 223–234.
- [121] M. Weiden, J. Kalloor, J. Kubiawicz, E. Younis, and C. Iancu, “Wide quantum circuit optimization with topology aware synthesis,” in *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)*. IEEE, 2022, pp. 1–11.
- [122] Y. Nam, N. J. Ross, Y. Su, A. M. Childs, and D. Maslov, “Automated optimization of large quantum circuits with continuous parameters,” *npj Quantum Information*, vol. 4, no. 1, p. 23, 2018.
- [123] Y. Ge, W. Wenjie, C. Yuheng, P. Kaisen, L. Xudong, Z. Zixiang, W. Yuhan, W. Ruocheng, and Y. Junchi, “Quantum circuit synthesis and compilation optimization: Overview and prospects,” *arXiv preprint arXiv:2407.00736*, 2024.
- [124] S. Sim, J. Romero, J. F. Gonthier, and A. A. Kunitsa, “Adaptive pruning-based optimization of parameterized quantum circuits,” *Quantum Science and Technology*, vol. 6, no. 2, p. 025019, 2021.
- [125] M. Ostaszewski, L. M. Trenkwalder, W. Masarczyk, E. Scerri, and V. Dunjko, “Reinforcement learning for optimization of variational quantum circuit architectures,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 18 182–18 194, 2021.
- [126] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, “Expressive power of parametrized quantum circuits,” *Physical Review Research*, vol. 2, no. 3, p. 033125, 2020.
- [127] Y. Yang, Z. Zhang, A. Wang, X. Xu, X. Wang, and Y. Li, “Maximizing quantum-computing expressive power through randomized circuits,” *Physical Review Research*, vol. 6, no. 2, p. 023098, 2024.
- [128] A. Arrasmith, M. Cerezo, P. Czarnik, L. Cincio, and P. J. Coles, “Effect of barren plateaus on gradient-free optimization,” *Quantum*, vol. 5, p. 558, 2021.
- [129] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, “Noise-induced barren plateaus in variational quantum algorithms,” *Nature communications*, vol. 12, no. 1, p. 6961, 2021.
-

- [130] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, “An initialization strategy for addressing barren plateaus in parametrized quantum circuits,” *Quantum*, vol. 3, p. 214, 2019.
- [131] K. Zhang, L. Liu, M.-H. Hsieh, and D. Tao, “Escaping from the barren plateau via gaussian initializations in deep variational quantum circuits,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 18 612–18 627, 2022.
- [132] F. Truger, J. Barzen, M. Bechtold, M. Beisel, F. Leymann, A. Mandl, and V. Yussupov, “Warm-starting and quantum computing: A systematic mapping study,” *ACM Computing Surveys*, vol. 56, no. 9, pp. 1–31, 2024.
- [133] A. Skolik, J. R. McClean, M. Mohseni, P. Van Der Smagt, and M. Leib, “Layerwise learning for quantum neural networks,” *Quantum Machine Intelligence*, vol. 3, pp. 1–11, 2021.
- [134] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, “A review on quantum approximate optimization algorithm and its variants,” *arXiv preprint arXiv:2306.09198*, 2023.





## Chapter 5

# Quantum Diffusion Models

The rapid adoption of Diffusion Model (DM)s in the EO domain has unlocked unprecedented generative capabilities for synthetic satellite images and their corresponding labels. However, classical DMs still face significant computational limitations and difficulties in capturing intricate spatial and spectral correlations present in EO data. Recent researches on QML, introducing first attempts of Quantum Generative Adversarial Network (QGAN)s and QDMs, offer a fundamentally different approach to overcome these challenges. Motivated by this progress, and leveraging the expressive power of quantum states and the high dimensionality of Hilbert space, this chapter introduces two novel approaches: a Quantum Latent Diffusion Model (QLDM) [1] and a class-conditioned QDM [2]. The proposed models are specifically designed for generating synthetic EO imagery, addressing the fundamental complications posed by the high cost and human effort required to produce annotated datasets.

### 5.1 Introduction

The use of DMs for generating synthetic satellite images and their corresponding labels is becoming an essential tool in EO, especially considering the high cost and human effort required to produce annotated datasets.

While traditional data augmentation techniques, such as geometric transformations and color adjustments have shown limited capability in capturing the underlying data distributions, mainly due to their reliance on fixed mathematical formulation, recent advancements in DL have introduced powerful generative approaches. For instance, AI methods such as Generative Adversarial Network (GAN)s and DMs [3–5] manage to synthesize highly realistic EO imagery.

More in detail, these models have been applied to a wide range of RS tasks such as cloud removal [6], hyperspectral and multispectral data generation [7, 8], and super-resolution enhancement [9]. However, challenges remain in computational efficiency, spectral consistency, and domain generalization [10, 11]. Precisely, despite GANs provide fast inference, they often suffer from instability and mode collapse. On the other hand, DMs can provide more stable and high fidelity training, but require hundreds to thousands of iterations at inference time.

Following this research path, recent work explored hybrid architectures that combine previous convolutional and vision transformers techniques into GANs and DMS strategies to balance realism, controllability, and computational cost [5]. Additionally, in the field of EO research, text-to-image generative models, which leverage semantic features, have enabled conditional image generation [12]. However, generative modeling for EO remains only partially explored, with a limited number of studies addressing its full potential.

In parallel, the research field of QML has gained traction with the development of quantum algorithms

for supervised and unsupervised learning tasks [13–16]. Moreover, quantum generative models, including QGANs [17] and QDMs [18, 19], are offering a fundamentally different approach to data generation by leveraging the expressive power of quantum states and the high dimensionality of Hilbert space [20, 21].

In light of this, [22] recently proposed a hybrid quantum GAN architecture (LaSt-QGAN) applied to the SAT4 EO dataset, marking one of the first applications of quantum generative modeling to EO. Furthermore, [19] designed a QLDM that leverages the established idea of classical latent diffusion models with a quantum denoiser circuit. However, QDMs performances are hindered by the information compression due to the autoencoder process. Specifically, in order to implement a fully quantum denoiser circuit, the images must be compressed to extremely low-dimensional representations, such that they can be encoded via angle encoding on a limited number of qubits. Nevertheless, as reported in [19], the results are encouraging and support the feasibility of leveraging QC within generative models in few-shot learning scenarios. Building on this idea, in [1] was proposed a QLDM to generate EO data. Thus, while the quantum latent model outperforms its classical counterpart, the strong compression required to operate in the latent space makes it less suitable for EO data, where preserving spatial and radiometric detail is essential. These limitations motivate a hybrid architecture: quantum modules are used only where they are most effective, without constraining the entire diffusion process. This relaxes the requirement of a tiny architecture and allows working directly with full-resolution data.

Motivated by the above considerations, in [2] is proposed the **Quanvolutional Conditioned U-Net (QCU-Net)** model, a hybrid quantum-classical architecture which applies quantum operations to feature extraction stages via a novel quanvolutional approach within a conditioned diffusion framework. Precisely, a class-conditioning mechanism is leveraged to generate class-dependent synthetic images. This approach, when compared with classical data augmentation techniques, enables the creation of more diverse and semantically rich samples, potentially better suited for enhancing the performance of general-purpose downstream EO tasks. Moreover, its value becomes particularly evident in scenarios where labeled data are scarce.

The main novelties introduced by the proposed **QCU-Net** model are hereafter reported:

- a novel hybrid class-conditioned QDM is proposed to EO data for RGB imagery, where the U-Net is enriched with (i) quantum layers at the bottleneck, where spatial resolution is minimal, and (ii) an additional quanvolutional layer placed earlier in the encoder;
- a class-conditioned generation process is introduced to guide the model using labeled information, and a specific protocol is used to assess the quality of the class-conditional generation;
- multiple quantum circuit ansatz and configurations are benchmarked to improve generation expressiveness and validate the results via ablation studies.

The proposed **QCU-Net** model achieves substantial improvements over classical diffusion-based models, as well as other quantum architectures. Notably, QCU-Net achieves a 64% reduction in FID and a 76% reduction in KID compared to classical baselines. It also exhibits enhanced semantic accuracy, particularly in difficult classes such as urban–vegetation boundaries, which are typically prone to misclassification due to high inter-class variability. Finally, the proposed model is shown to produce samples that better approximate the true distribution of satellite imagery compared to a classical counterpart.

Consequently, by introducing quantum generative modeling to EO, this work aims to lay the groundwork for future quantum-enhanced applications in RS. These include tasks such as multi-sensor simulation and quantum-enhanced EO data processing, as well as the generation of synthetic data to support general-purpose downstream applications in contexts where labeled data is limited or difficult to obtain.

The rest of the chapter is organized as follows. Section 5.2 introduces the fundamentals of QC and DMs, offering the theoretical foundation needed to grasp the proposed hybrid quantum-classical generative technique. Section 5.3 discusses the SOTA, reviewing both classical generative models, including GANs and DMs in EO, and recent developments in quantum generative models. Section 5.4 describes the proposed model, presenting in detail the QLDM and the QCU-Net architecture designed for EO imagery synthesis, as well as the other quantum model integrated within the diffusion pipeline used for comparison. It is also outlined the class-conditioning mechanism that guides generation. Section 5.5 summarizes the experimental setup, including software, training procedures, and evaluation metrics, along with the specifics of the EuroSAT dataset utilized in the conducted experiments. Section 5.6 presents a comprehensive set of experimental results, highlighting key improvements in image realism, semantic fidelity, and computational efficiency, as well as the insights derived from a detailed ablation study. Finally, Section 5.7 concludes the chapter, summarizing the contributions and proposing future research directions for quantum-enhanced EO applications.

## 5.2 Background

This section introduces the fundamental concepts behind the proposed research area, focuses on DMs and their relevance to the proposed methodology.

### 5.2.1 Fundamentals of Diffusion Models

In this section some fundamental concepts about DMs are first introduced, with a specific focus on the parameters that may impact generation performance.

In particular, DMs are characterized by two main phases: the forward (diffusion) process and the reverse process.

During the forward process, the noise schedule  $\beta_t$ , as defined in [23], governs a Markov chain that can be described as:

$$(\mathbf{x}_{0:T}) = q(\mathbf{x}_0) \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I})$$

Data samples  $\mathbf{x}_0 \sim q$  from the true distribution are gradually corrupted through  $T$  steps, resulting in  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

The reverse process instead has the aim of inverting this diffusion by sampling from  $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$ . However, this conditional distribution is generally intractable due to its dependence on the true data distribution, and thus a data-driven approach is adopted.

Let  $p_\theta$  be a neural network-based approximation of  $q(\mathbf{x}_{t-1} | \mathbf{x}_t)$ , modeled as:

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t))$$

In simplified settings,  $\boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)$  is kept fixed equal to  $\beta_t \mathbf{I}$ , and the network is trained to predict the noise term  $\epsilon_\theta(\mathbf{x}_t, t)$ , since the mean  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  can be expressed as a function of the predicted noise as follows:

$$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}} \left( \mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right).$$

As a result, the training process is driven by the following loss function:

$$L_t^{\text{simple}} = \mathbb{E}_{\mathbf{x}_0 \sim q, t, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\epsilon_\theta(\mathbf{x}_t, t) - \epsilon\|^2]$$

with the configuration of the diffusion rate  $\beta$  fundamental for its final performance. Different solutions are adopted. In particular, in [23, 24] authors set a *linear*  $\beta$  variance ranging from  $\beta_1 = 10^{-4}$  to  $\beta_T = 0.02$  with  $T = 1000$  steps. In contrast, in [25], authors propose to improve diffusion models with a reparametrization of the generation process variance, i.e., replacing the linear schedule with a squared *cosine* to prevent abrupt changes of noise levels. This choice leads to a slower forward process with  $T = 4000$  steps while increasing reconstructed image details.

## 5.3 Related Works

This section reviews the evolution of generative methods for EO image-synthesis, from early GAN-based approaches (in Section 5.3.1) to classical DMs (in Section 5.3.2). Additionally, recent quantum generative techniques are also investigated in Section 5.3.3.

### 5.3.1 Generative Adversarial Networks for EO Image Synthesis

Generative models, particularly GANs, have become central to EO image synthesis tasks such as cloud removal, super-resolution, hyperspectral generation, and data augmentation. Early augmentation methods based on geometric or radiometric transformations [26] lacked the ability to capture the underlying data distribution, limiting variability and realism.

GANs introduced adversarial learning to model complex EO data distributions and have been successfully applied across multiple EO domains. Conditional GANs and CycleGANs have enabled both paired and unpaired image translation, including for tasks like cloud-to-clear image mapping [6, 27, 28], SAR-to-optical conversion [29], and temporal gap filling [30]. Attention-based architectures, such as transformer-enhanced GANs, have improved image sharpness and spatial consistency in thin-cloud removal [31] and super-resolution [32]. More recently, He et al. [33] proposed an Optimal-Transport GAN that couples dual attention modules with a semi-discrete OT loss to mitigate large intra-class and small inter-class variance in EO scenes, achieving SOTA fidelity on LoveDA and GID-15 benchmarks.

In hyperspectral and multispectral settings, GANs have been employed for super-resolution and pan-sharpening [34], as well as for class-conditioned data generation to augment small or imbalanced datasets [35]. These approaches demonstrate GANs' strength in producing visually realistic and geometrically detailed outputs, often improving downstream classification or detection performance.

However, GAN-based generation in EO still suffers from challenges such as training instability, mode collapse, and spectral inconsistency [4, 36]. Moreover, generated outputs may introduce artifacts features, making interpretability and physical validity critical concerns in operational contexts. Consequently, in order to overcome such limitations, the proposed strategy leverages the use of DMs for a more stable and coherent EO image generation.

### 5.3.2 Diffusion Models in Remote Sensing

In the context of EO, DMs have emerged as leading generative approaches, demonstrating a valuable contribution across a wide range of tasks, including image generation [37], super-resolution, cloud removal [11, 38], landscape classification, as well as change and climate prediction [39].

In particular, image generation plays a crucial role in EO applications, especially in scenarios involving hyperspectral imaging, where data scarcity is primarily due to the high acquisition cost [39]. Similarly,

in domains such as digital elevation models, data loss often results from limitations in LiDAR data capture systems [40].

SOTA generative models like DMs help overcome these limitations, enabling improved performance not only in data reconstruction but also in subsequent downstream tasks [41]. Recently, [42] introduced a four-stage diffusion-based augmentation pipeline, combining meta-prompting, InstructBLIP captioning, and LoRA fine-tuning of Stable Diffusion within an iterative synthesis loop that boosts CLIP top-1 accuracy on EuroSAT dataset. Among the most advanced models in this field, DiffusionSat [43] and HyperLDM [44] enhance spectral and spatial consistency through pixel-space and latent-space diffusion architectures, respectively. COP-GEN-Beta [45] extends diffusion modeling to the multimodal EO domain by learning joint distributions across SAR, optical, and elevation modalities.

Nonetheless, training and deploying DMs presents several challenges, primarily due to the need for highly parameterized neural networks, which are computationally intensive and difficult to train [12]. Some EO-specific solutions have been proposed in [46], such as an efficient diffusion transformer via self-attention distillation [47], which significantly reduces the number of parameters. Another approach involves using latent-space diffusion models [44], which mitigate computational costs by avoiding pixel-space processing.

### 5.3.3 Quantum Generative Models in Machine Learning

Quantum generative models leverage quantum mechanics to encode and generate complex data distributions. QGANs [13, 16], Quantum Variational Autoencoders, and Quantum Born Machines [48] were among the first proposed architectures, showing promise in low-data regimes.

Recently, QDMs have been introduced as quantum analogs of Denoising Diffusion Probabilistic Model (DDPM)s. The Quantum Generative Diffusion Model (QGDM) [49] defines a forward noisy process and a learnable reverse process via quantum circuits. QuDDPM [50] further extends this with parameterized quantum noise injection and denoising for data generation. Hybrid quantum-classical approaches [18] combine classical neural networks with quantum layers to enhance expressive power and reduce the number of parameters. Other proposals, such as quantum-noise-driven diffusion [51], treat quantum decoherence as a generative feature.

Although GANs and DMs are increasingly applied to EO data and general-purpose QGANs/QDMs are gaining traction, the application of QDMs to EO tasks remains largely unexplored with only one prior work introducing a QDM, specifically a QLDM that couples the classical latent-DM paradigm with a quantum denoiser circuit [19], highlighting a clear opportunity that the proposed research addressed by introducing the first class-conditioned QDM for EO image generation.

The novel proposed QCU-Net model introduces a framework well suited to EO high dimensionality, spectral complexity, and label scarcity data, and with respect to [1], there is no reduction in quality due to compression of images, and moreover, the class-conditioned new model is able to generate labeled images for the different classes. The results offer a powerful tool for the EO domain and community of RS researchers.

## 5.4 Methodology

This section presents the methodological framework adopted in this chapter, which encompasses both the QLDM and the proposed hybrid architecture, the QCU-Net. These models are explored as two complementary approaches for integrating quantum generative models, specifically QDMs, into the EO domain.

The QLDM provides a first demonstration of the potential benefits of introducing quantum operations within the diffusion process. However, despite outperforming its classical counterpart, the strong image compression required to operate in the latent space makes this approach less suitable for EO data, where preserving spatial and radiometric detail is essential. These limitations motivate the introduction of a hybrid architecture, in which quantum modules are applied only where they are most effective, without constraining the entire diffusion pipeline. This design relaxes the requirement for extremely compact latent representations and enables the model to operate directly on full-resolution data.

In addition to these two main models, a simpler conditioned baseline, the Quantum Vertex Conditioned U-Net (QVCU-Net), is also implemented to provide a controlled comparison. This leads to the following structure of the section:

1. **QLDM** (Section 5.4.1): a convolutional auto-encoder compresses each RGB image into a latent vector, and three VQCs are used to predict the denoising component at each timestep.
2. **QVCU-Net** (Section 5.4.2): a conditioned baseline model that integrates quantum layers exclusively at the bottleneck, providing a controlled quantum-enhanced reference for comparison.
3. **QCU-Net** (Section 5.4.3): the main hybrid architecture proposed in this work, which enriches a U-Net-like structure with (i) quantum layers at the bottleneck, where spatial resolution is minimal, and (ii) an additional quanvolutional layer placed earlier in the encoder [52, 53]. This allows the model to incorporate quantum feature transformations while preserving full-resolution spatial information.

Furthermore, Section 5.4.4 provides a detailed description of the quantum variational circuits and the ansätze employed across all architectures, while Section 5.4.5 describes the protocol used for assessing the quality of class-conditional generation.

### 5.4.1 Quantum Latent Diffusion Model




Figure 5.1: Architecture of the proposed quantum latent diffusion model.

The classical diffusion process involves two main stages: the forward phase, where noise is gradually introduced through a Markov chain, and the reverse phase, where the reverse Markov chain is reconstructed by a parametric model. Although this work maintains this process, the key difference is that the parametric model is implemented using VQCs. In order to effectively utilize VQCs, the model operates within the latent space, which is reached by employing a classical convolutional autoencoder. This autoencoder has been separately trained on the dataset of interest and can also introduce non-linearities into the process, which the proposed QLDM can leverage. In contrast to [19], the autoencoder processes 32x32x3 images, allowing compatibility with higher resolution images.

Downstream of the encoder, the latent vector of dimension 10 is utilized for the diffusion process and the subsequent image generation. As shown in Fig. 5.1, the architecture employs three distinct VQCs, each designed with a specific role. The first VQC processes the downsampled image from the encoder, degraded by noise added during the forward diffusion process. The second VQC addresses temporal information, essential for incorporating diffusion times into the model. The third VQC combines the outputs of the first two by processing their sum. The expectation value measured by the third VQC is used to directly estimate the noise  $\epsilon_\theta(\mathbf{x}_t, t)$  at time step  $t$ .

Each VQC is implemented using angle encoding, a feasible approach given the reduced size of the latent vector after using the encoder. In order to assess the potential advantages of quantum computation, the performances of the proposed QLDM were compared against:

- a custom classical Latent Diffusion Model (LDM) with the same structure as the QLDM, except that three MLPs replace the three VQCs;
- a QGAN inspired by LaSt-QGAN [22], featuring a quantum generator implemented with a VQC analogous to those in the proposed QLDM, paired with a classical discriminator;
- a classical GAN operating in the latent space.

All models were configured with a comparable number of parameters to ensure a fair comparison and a more accurate analysis.

While the latent-based architecture enables the integration of a fully-quantum denoiser into the diffusion process, its strong compression requirements make it less suitable for high-resolution EO imagery. To overcome these limitations and evaluate quantum contributions in a less constrained setting, a hybrid model operating directly in the image domain is also considered. This motivates the introduction of the following architectures.

#### 5.4.2 QVCU-Net

The U-Net, proposed for the first time in [54] as a neural network for image segmentation, has been established as a standard architecture used for the reverse diffusion process in a DM. Characterized by its distinctive U-shaped structure, it is a highly parameter-rich neural network consisting of an encoder and a symmetric decoder responsible for reconstructing the image from the feature space.

Starting from a SOTA U-Net architecture, composed of ResNet blocks and Multihead Attention mechanism with four attention heads as suggested in [55], quantum layers were integrated into this classical architecture, following the approach proposed in [18].

The rationale behind incorporating quantum layers into a classical network is motivated by the strengths of both paradigms. Precisely, while classical model can handle high-dimensional data as EO data, which is infeasible for the current NISQ devices, quantum layers offer a potential advantage to overcome some classical issues, e.g. the computational intensity and training difficulties of traditional models.




Figure 5.2: Class-conditioned *Quantum Diffusion Model*. **Top:** forward chain  $q(x_t | x_{t-1})$  (green) corrupts  $x_0$  into standard Gaussian noise  $x_T$ ; the learned reverse  $p_\theta(x_{t-1} | x_t, y)$  (blue) denoises step-wise under label  $y$ . **Bottom:** UNet denoiser with classical Conv/ResNet/Attention blocks plus two quantum residual variants: *QuanResNet* (red) replaces the *first* conv by a quanvolution, while *QResNet* (purple) at the bottleneck replaces *both* convs with a VQC acting on  $2 \times 2 \times 3$  features, applied to a fraction  $\rho$  of channels.

Specifically, in the proposed QVCU-Net architecture the quantum component is integrated into the ResNet block located at the vertex as shown in Fig. 5.2. Given that the image has the lowest spatial dimension at the vertex of the network, few qubits can be used. This motivates the use of angle data embedding. At the same time, having such low spatial dimension allows the implementation of a specific ansatz, described in Section 5.4.4, designed to process groups of adjacent channels. Thereby, the network can exploit both intra-channel and inter-channel correlations to the fullest extent.

Nevertheless, it is worth noting that the feature representation at the bottleneck consists of 60 channels (given the considered configuration). To better assess the contribution of the quantum layers, VQCs are gradually introduced over some vertex channels.

Ultimately, compared to the architecture proposed in [18], two main innovations have been introduced: *I*) the architecture has been adapted to handle EO images, specifically for RGB bands. This adaptation required an increase in the number of blocks within the U-Net, thereby making the model more complex but also more expressive and performative *II*) conditional generation has been incorporated. In addition to the image and temporal information, the model now also receives as input the class label  $y$ . This conditioning mechanism is particularly important in the context of the targeted task, i.e. to generate labeled EO images.

### 5.4.3 QCU-Net

Starting from the QVCU-Net architecture presented in the previous section, an additional quantum layer was introduced in the initial part of the encoder, leading to the development of the QCU-Net, illustrated in Fig. 5.2. In fact, the encoder is the portion of the U-Net responsible for extracting hierarchical features and capturing contextual information. Introducing quantum operations in such a critical and sensitive region of the network could therefore significantly enhance overall performance.

Unlike for the vertex, a key challenge in this encoder branch is that the feature maps have higher dimensions, depending on the ResNet block being used. As a result, encoding the entire feature map into a single VQC becomes infeasible, so an alternative approach is required.

The adopted solution is the quanvolutional approach, inspired by [52, 53]. Here, a quantum layer is



applied to  $k \times k$  patches of the input, akin to a classical convolutional operation, until the entire feature map has been processed, as depicted in Fig. 5.3. The key advantage of this method lies not only in the ability to effectively integrate quantum computation into critical stages of the encoder, but also in the construction of a feature map that transforms spatially-local subsections of the input tensor.

By default the quanvolution operation has been introduced in the 3 – rd upper layer of the QCU-Net. An ablation study has been conducted and reported in Section 5.6.2 to justify such a placement of the quanvolutional layer within the encoder branch.

#### 5.4.4 Variational Quantum Circuits and Ansatz Design

The VQCs used in both architectures consist of three main components. The first part is dedicated to data encoding. It is adopted angle encoding, where classical inputs  $\mathbf{x} \in \mathbb{R}^n$  are encoded into quantum states using rotations parameterized by the input values. Specifically, rotations are employed along the  $x$ -axis, as follows:

$$|\phi(\mathbf{x})\rangle = R_x(\mathbf{x})|0\rangle^{\otimes n} \quad (5.1)$$

The second part corresponds to the trainable portion of each variational circuit, namely the ansatz. For a comprehensive evaluation of the QLDM, three distinct ansatzes were employed. The first consisted of layered rotations along  $R_x$ ,  $R_z$ , and  $R_x$ , following the approach outlined in [19]. The second was the universal ansatz, as proposed in [56], while the third was the matchgate ansatz, introduced in [57]. Measurements were performed using the  $Z$  observable.

As regards the hybrid model, three different options were evaluated: the Hierarchical Quantum Convolutional (HQConv), the Flat Quantum Convolutional (FQConv), and the Only Rotations ansatz respectively. Fig. 5.3 shows the HQConv ansatz, which proves to be the most effective solution, as also explained in Section 5.6. The first two ansatz, based on [58], are particularly well-suited to the application context, where the quantum layer is integrated either at the vertex of a U-Net or within the encoder. These ansatz are specifically designed to process three image channels simultaneously. For example, at the vertex, where the image has spatial dimensions  $2 \times 2$ , a total of 12 qubits is used ( $2 \times 2$  spatial dimensions  $\times$  3 channels). The same principle applies to the QuanvUnet, where the variational quantum circuit acts as a quanvolutional filter, i.e., processing the image in  $2 \times 2$  patches and using three channels at a time.

Operating on three channels concurrently allows the model to extract features not only within each individual channel, but also across channels, potentially enhancing feature richness.

In the following list, these 3 ansatzs are detailed:

1. HQConv (see Fig. 5.3): this ansatz is structured in two stages. The first stage processes intra-channel correlations: each block (labeled Block A in the figure) acts only on qubits belonging to the same channel, and can be expressed as:

$$|q_t^o, q_{t+r}^o\rangle = [CR_x(\theta_{x,t}) \circ CR_z(\theta_{z,t})] |q_t^i, q_{t+r}^i\rangle \quad (5.2)$$

Where, superscripts  $i$  and  $o$  indicate the quantum state before and after the application of the block, respectively. The subscript  $t$  denotes the control qubit, while  $t + r$  indicates the corresponding target qubit. The second stage of the ansatz, marked as Block B in Fig. 5.3, is designed to capture inter-channel correlations. It operates as:

$$|q_t^o, q_{t+s}^o\rangle = [CR_x(\theta_x) \circ CR_z(\theta_z)] |q_t^i, q_{t+s}^i\rangle \quad (5.3)$$

Similarly, superscripts  $i$  and  $o$  represent the state before and after the application of this block respectively. This structure enables information exchange across different channels, which is particularly advantageous for multi-channel image processing tasks.

2. FQConv: this ansatz is also composed of two stages, but unlike HQConv, it immediately incorporates both intra-channel and inter-channel information at each layer. Both stages use the same general block structure, described by:

$$|q_t^o, q_{t+r}^o\rangle = [CR(\theta_t)] |q_t^i, q_{t+r}^i\rangle \quad (5.4)$$

Where, superscripts  $i$  and  $o$  denote the quantum states before (input) and after (output) the application of the controlled rotation. Precisely, in the first stage, the controlled rotation is performed along the  $z$ -axis, i.e.,  $CR_z$ , while in the second stage, it is along the  $x$ -axis, i.e.  $CR_x$ .

3. Only-Rotations: this ansatz consists solely of single-qubit rotations along the three axes ( $x$ ,  $y$ , and  $z$ ), with no entanglement operations. It was included exclusively for comparative analysis purposes, as discussed in Section 5.6.2, to evaluate how the more structured ansatz may offer performance improvements, particularly in multi-channel scenarios where entanglement and feature interactions across channels are expected to be beneficial.

The third part of each variational circuit is represented by the measurement stage, where quantum states are mapped back to classical data. In this case, the Pauli-Z observable was used, which has eigenvalues of  $+1$  and  $-1$ .

#### 5.4.5 Conditioning Evaluation Protocol

Among the objectives of this work it was necessary to assess whether the proposed QCU-Net and also the implemented QVCU-Net architectures were not only capable of producing visually and quantitatively good samples (for evaluation metrics and results, refer to Sections 5.6), but also able to generate images that correctly belong to the desired classes.

To this end, after generating  $n$  samples, each associated with a label  $y \in \{0, 9\}$ , which can be selected at inference time, a SOTA ResNet50 classifier [59–62] was trained on EuroSAT.

Next, the classifier’s weights were frozen, and the  $n$  generated samples were fed through the network in inference mode to obtain the predicted labels  $\tilde{y} \in \{0, 9\}$ . Assuming the classifier’s error rate is minimal (approximately 1%), the quality of the conditioning was assessed by measuring how often the predicted labels  $\tilde{y}$  matched the intended labels  $y$  used during sample generation. This evaluation is formalized by the following expression:

$$\text{Accuracy}_{\text{cond}} = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{y_i = \tilde{y}_i\}}, \quad (5.5)$$

where  $\mathbb{1}_{\{y_i = \tilde{y}_i\}}$  is the indicator function that returns 1 if  $y_i = \tilde{y}_i$ , and 0 otherwise.

## 5.5 Experimental Setup

To ensure the reproducibility and clarity of the experiments, this section outlines the complete setup used throughout the investigation. Precisely, in Sections 5.5.1 software environment, libraries, architectures, and training configurations are respectively reported. Subsequently, Section 5.5.2 presents the hyperparameter tuning strategy and provide an overview of the different model sizes evaluated.




Figure 5.3: Hybrid quantum–classical ResNet. *Top*: comparison between the classical ResNet block and the QuanResNet block, in which one convolutional layer is replaced by a Quanvolutional (quantum) layer. *Bottom*: schematic of the HQConvAnsatz variational quantum circuit: angle encoding of feature maps, parameterized rotations, entangling gates, and measurements to generate quantum-enhanced features.

Finally, in Sections 5.5.3 and 5.5.4 evaluation metrics used to assess model performance and the EuroSAT dataset, which serves as the primary benchmark for the conducted experiments, are respectively discussed.

### 5.5.1 Software, libraries and training configurations

All code is written in *Python* 3.10. Core computation relies on *JAX* (0.4.x) for accelerated tensor algebra and automatic differentiation, with *Flax* (0.8.x) providing the high-level neural-network API. Data ingestion is handled by *TensorFlow Datasets*. Quantum layers are implemented through *PennyLane* and integrated into custom residual blocks; they are simulated in a noiseless environment due to current computational constraints and are not executed on real quantum hardware. The use of *Flax* in the proposed framework enables a flexible and efficient execution of hybrid neural networks through compilation. In particular, it allows the selection of the *JAX* back-end for the quantum layers implemented in *PennyLane*. This back-end provides access to just-in-time compilation and automatic vectorization

of quantum circuits, resulting in a significant reduction in runtime (up to  $10\times$  faster compared to implementations of *PennyLane* within other neural network frameworks such as *PyTorch*) [63]. Practically, this enables the integration of a significant quantum contribution into a hybrid neural network model without compromising computational efficiency.

Regarding the evaluation metrics used for image assessment, the *Torchvision* library [64] was employed for image related metrics and *scikit-learn* [65] for classification metrics. All experiments were executed on the *DEEP Machine* cluster at the Jülich Supercomputing Centre.

### 5.5.2 Hyperparameters tuning and models' size

**Hyperparameter space** To ensure fair comparisons and establish a strong classical baseline for the hybrid model, a hyperparameter tuning of the classical model was performed using the Adam optimizer. Indeed, as mentioned earlier, suitable baselines were available for the latent model, whereas no comparable quantum-enhanced baselines existed for the hybrid model. For this reason, the classical counterpart was optimized to enable a fair comparison. The ablation focused on key optimizer parameters: learning rate (lr), first and second momentum coefficients ( $\beta_1$ ,  $\beta_2$ ), and numerical stability term ( $\epsilon$ ). Specifically, the combinations of 3 lr values ( $1e-2$ ,  $1e-3$  and  $1e-4$ ) were tested, 3 values for  $\beta_1$  (1.85, 1.90 and 1.95), 3 values for  $\beta_2$  (0.85, 0.90 and 0.95) and 3 values from  $\epsilon$  ( $1e-7$ ,  $1e-8$  and  $1e-9$ ). Results of the hyperparameters tuning are reported in Section 5.6, paragraph 5.6.3.

**Models' size** Table 5.1 illustrates the models' sizes in terms of trainable parameters and the percentage difference compared to the classical model. It is evident that quantum-based models consistently exhibit a reduced number of trainable parameters, with QVCU-Net displaying the fewest. However, the proposed model, QCU-Net, still allows for a reduction in the number of parameters compared to the classical one. A similar trend is observed in the latent-space setting, as shown in Tab. 5.2 where the QLDM variants also require fewer parameters once quantum layers replace their classical counterparts, further confirming the efficiency of quantum-enhanced formulations.

Table 5.1: List of models's size in terms of trainable parameters and the delta in percentage with respect to the classical model (hybrid configurations).

| Model     | Trainable Parameters | $\Delta_{classical}\%$ |
|-----------|----------------------|------------------------|
| Classical | 1.280.843            | -                      |
| QVCU-Net  | 1.157.723            | 9.6                    |
| QCU-Net   | 1.260.939            | 1.6                    |

Table 5.2: List of models's size in terms of trainable parameters and the delta in percentage with respect to the classical model (latent configurations).

| Model            | Total Parameters | $\Delta_{classical} (\%)$ |
|------------------|------------------|---------------------------|
| Classical LDM    | 330              | –                         |
| QLDM (Rx–Rz–Rx)  | 360              | 9.09                      |
| QLDM (universal) | 342              | 3.64                      |
| QLDM (matchgate) | 336              | 1.82                      |
| Classical GAN    | 341              | 3.33                      |
| QGAN             | 351              | 6.36                      |

### 5.5.3 Evaluation Metrics

The quality of the generated images was assessed using three metrics: *FID*, *KID*, and *IS*, which collectively evaluate the quality and diversity of generated images [66].

The FID quantifies the similarity between feature distributions extracted from real and generated images using activations of a pre-trained Inception network. Mathematically, the FID is expressed as:

$$\text{FID}(\mathbf{X}, \mathbf{Y}) = \|\mu_X - \mu_Y\|_2^2 + \text{Tr}(\Sigma_X + \Sigma_Y - 2(\Sigma_X \Sigma_Y)^{1/2}), \quad (5.6)$$

where  $\mu_X$ ,  $\mu_Y$  and  $\Sigma_X$ ,  $\Sigma_Y$  represent the mean and covariance matrices of the feature distributions for real ( $\mathbf{X}$ ) and generated ( $\mathbf{Y}$ ) images respectively.

The KID evaluates distribution similarity using the squared Maximum Mean Discrepancy (MMD) between features from real and generated images, employing a polynomial kernel defined as:

$$\begin{aligned} \text{KID}(\mathbf{X}, \mathbf{Y}) = & \mathbb{E}_{x, x' \sim \mathbf{X}}[k(x, x')] + \mathbb{E}_{y, y' \sim \mathbf{Y}}[k(y, y')] - 2\mathbb{E}_{x \sim \mathbf{X}, y \sim \mathbf{Y}}[k(x, y)] \end{aligned} \quad (5.7)$$

with kernel function  $k(x, y) = \left(\frac{1}{d}x^\top y + 1\right)^3$ , where  $d$  is the feature dimensionality.

The IS metric quantifies both quality and diversity of the generated images by evaluating their class distribution predicted by the Inception network:

$$\text{IS} = \exp(\mathbb{E}_{x \sim \mathbf{Y}}[D_{KL}(p(y|x})||p(y))]), \quad (5.8)$$

where  $D_{KL}$  is the Kullback–Leibler divergence,  $p(y|x)$  the conditional class distribution given an image  $x$ , and  $p(y)$  the marginal distribution over generated images.

Lower FID and KID values indicate closer alignment with real image distributions, whereas a higher IS value denotes superior quality and greater diversity in the generated samples. In the considered experimental setup, the metrics were computed using the TorchMetrics library [64]. Specifically, the KID was computed using subsets of 100 image pairs (100 real and 100 generated samples) to estimate the mean and variance, whereas the IS was averaged over 10 equal splits of the generated-image set. Finally, the class-conditioning capability was assessed by generating scikit-learn’s classification report, which calculates per-class precision, recall, F1-score, and support, along with overall accuracy and the macro and weighted averages of those metrics [67]; in addition, accuracy was computed separately for each individual class.

### 5.5.4 EuroSAT dataset

The EuroSAT dataset [68] consists of 27,000 Sentinel-2 satellite images covering 10 LULC classes, including agricultural areas, residential zones, and water bodies. Each image has a spatial resolution of  $64 \times 64$  pixels and includes 13 spectral bands. Although the dataset was originally developed for supervised classification tasks, in this work it was employed to generate synthetic satellite images using quantum diffusion models for data augmentation.

Since both quantum diffusion architectures considered in this study operate on three-channel inputs, the experiments were restricted to the RGB bands of EuroSAT. This choice ensures compatibility with evaluation metrics commonly used in generative modelling and reduces the computational complexity of the pipeline. However, the two quantum models use the dataset in partially different configurations.

**QLDM.** To accelerate training and accommodate the computational cost of simulating the quantum latent module, the RGB images were resized from their original resolution to  $32 \times 32$  pixels. Moreover, to conduct controlled experiments on generative performance, the work focused on a subset of

three classes: *Forest*, *Herbaceous Vegetation*, and *SeaLake*, each containing 2,700 samples. This yields a balanced dataset of 8,100 images.

**Hybrid QDM.** For the hybrid architecture, the EuroSAT RGB images were used at their original resolution of  $64 \times 64$  pixels. This configuration allowed to retain richer spatial information and to assess model scalability in a setting closer to the original EuroSAT distribution. The dataset serves as a benchmark for testing the capability of the hybrid quantum–classical model to generate high-quality synthetic images suitable for downstream data augmentation.

## 5.6 Results

The following section provides a comprehensive evaluation of the latent and hybrid quantum frameworks on the EuroSAT benchmark. The section is organized as follows: Section 5.6.1 reports the main results by comparing 1) the QLDM against its classical counterpart, the QGAN, and the latent-space GAN and 2) the baseline classical convolutional diffusion model and the QVCU-Net with the proposed QCU-Net, highlighting gains in sample realism, diversity, and label fidelity.

Finally, Section 5.6.2 presents a comprehensive ablation study analyzing the impact of quantum layer placement and circuit ansatz. Building upon these results, Section 5.6.3 offers a deeper investigation into specific factors influencing model behavior, including convergence dynamics and class-wise conditioning accuracy.

### 5.6.1 Main results

**Latent Model Generation Capabilities** To assess the performances of the proposed QLDM, the FID, KID and IS were evaluated. A comparison of these metrics is provided in Table 5.3, where lower FID/KID indicate better quality while higher IS reflects greater diversity. Specifically, in the KID computation, a subset size of 100 was used to calculate the mean and variance while, for the IS calculation, the dataset was divided into 10 splits to compute the mean and variance.

Table 5.3: Comparison of generative models by parameters, ansatz, and metrics (FID, KID, IS).

| Model            | #Classic parameters | #Quantum parameters | ansatz    | FID          | KID                                   | IS                                    |
|------------------|---------------------|---------------------|-----------|--------------|---------------------------------------|---------------------------------------|
| Classical LDM    | 330                 | 0                   | -         | 35.28        | $0.0368 \pm 0.0021$                   | $1.2659 \pm 0.0053$                   |
| QLDM (Rx-Rz-Rx)  | 0                   | 360                 | Rx-Rz-Rx  | 30.39        | $0.0311 \pm 0.0019$                   | $1.3228 \pm 0.0133$                   |
| QLDM (universal) | 0                   | 342                 | universal | <b>27.69</b> | <b><math>0.0258 \pm 0.0018</math></b> | $1.3152 \pm 0.0091$                   |
| QLDM (matchgate) | 0                   | 336                 | matchgate | 41.15        | $0.0484 \pm 0.0026$                   | <b><math>1.4440 \pm 0.0115</math></b> |
| Classical GAN    | 110+231             | 0                   | -         | 33.25        | $0.0311 \pm 0.0012$                   | $1.1847 \pm 0.01224$                  |
| QGAN             | 231                 | 120                 | Rx-Rz-Rx  | 33.54        | $0.0384 \pm 0.0013$                   | $1.1950 \pm 0.0071$                   |

The QLDM with universal ansatz achieved the best FID and KID performances, with scores of 27.69 and  $0.0258 \pm 0.0018$ , respectively. These results indicate a significant improvement in image quality compared to the classical LDM and the other QLDM variants. The universal ansatz appears to enable the most effective representation of the data, likely due to its flexibility in capturing complex correlations. The QLDM with the Rx-Rz-Rx ansatz also demonstrated strong performance, with an FID of 30.39 and a KID of  $0.0311 \pm 0.0019$ . However, its performance was slightly lower than the universal ansatz. Interestingly, the matchgate-based QLDM exhibited a notable trade-off: while its FID (41.15) and KID ( $0.0484 \pm 0.0026$ ) were worse than those of the other QLDM variants, it achieved the highest IS ( $1.4440 \pm 0.0115$ ), suggesting that the generated images might have higher perceptual diversity despite lower overall quality.

Classical models such as Classical LDM and GANs performed comparably, with the Classical LDM achieving an FID of 35.28 and a KID of  $0.0368 \pm 0.0021$ . The Classical GAN and QGAN, while similar in architecture and number of parameters, show distinct performance differences. The QGAN achieves a higher FID score overall. This suggests that the QGAN failed to demonstrate any significant advantage, possibly due to limitations in the quantum ansatz or parameterization.




Figure 5.4: FID scores over training epochs for different generative models.

Moreover, the evolution of FID scores across training epochs is shown in Fig. 5.4. Specifically, the FID for the universal ansatz decreases steadily, converging to the lowest value among all models. Notably, the Classical GAN is not included in Fig. 5.4 due to its significantly slower training, which would compress the other curves and make comparisons less clear. Meanwhile, the QGAN exhibits slower convergence compared to the other models.

Overall, the results highlight the potential of quantum-enhanced generative models, particularly the QLDM with the universal ansatz, achieving superior image generation quality for data augmentation. However, trade-offs in diversity and quality remain a critical area for further research.

**Hybrid Model Generation Capabilities** Table 5.4 compares the classical diffusion model and the QVCU-Net with the proposed QCU-Net model, using a single HQConv circuit, after 60,000 training iterations on the EuroSAT dataset.

The hybrid solution achieves substantial improvements across key metrics: FID drops from 7.22 to 2.57, while KID from  $3.4 \times 10^{-3}$  to  $0.8 \times 10^{-3}$ , representing relative reductions of 64% and 76%, respectively. These improvements indicate that QCU-Net model produces samples that are numerically closer to real images. The improvement of QCU-Net is significant even relative to the QVCU-Net model used as the comparison framework.

Visual inspection supports these findings: as shown in Fig. 5.5, QCU-Net model generates scenes that

Table 5.4: Comparisons of best classical and quantum models after 60,000 iterations on the EuroSAT dataset

| Model     | FID         | KID                                   | IS                                    |
|-----------|-------------|---------------------------------------|---------------------------------------|
| Classical | 7.22        | $0.0034 \pm 0.0008$                   | $1.8837 \pm 0.0306$                   |
| QCU-Net   | <b>2.57</b> | <b><math>0.0008 \pm 0.0003</math></b> | $1.7678 \pm 0.0241$                   |
| QVCU-Net  | 7.60        | $0.0037 \pm 0.0008$                   | <b><math>1.9011 \pm 0.0255</math></b> |




Figure 5.5: Qualitative comparison between GT and generated images on the EuroSAT dataset. Columns correspond (left to right) to the ten land-cover classes: Annual Crop, Forest, Herbaceous Vegetation, Highway, Industrial, Pasture, Permanent Crop, Residential, River, and Sea Lake. Rows show GT reference images (first row), outputs of the classical model (second row), and outputs of the quantum models (QVCU-Net and QCU-Net) (third and fourth rows).

are noticeably clearer and more visually aligned with GT samples. In contrast, outputs from the classical baseline tend to suffer from blurriness, reduced spatial coherence, and recurrent artifacts.

Figure 5.6 provides a comparative analysis of the FID convergence trends between the classical, QCU-Net and QVCU-Net models, throughout 60,000 optimisation steps. The figure shows a distinct trend: the QCU-Net rapidly surpasses the classical baseline and the QVCU-Net. By 20,000 iterations, QCU-Net achieves an FID close to 5, while the classical network remains near 10, and QVCU-Net positions just below 9.5. So it is evident that with 30% of total iterations the proposed model reaches already a satisfactory result.

QVCU-Net, improves more slowly and, as the ablation study will show, the position at which the quantum module is inserted plays a pivotal role in the quality of the generated results.

**Conditioning Fidelity:** To evaluate how accurately each model follows the specified class labels, approximately 10,000 synthetic images were generated (1000 per class) and classify them using a ResNet-50 model pre-trained on real EuroSAT data, which achieves a validation accuracy of 98.4%.

Table 5.5 shows a clear trend: overall QCU-Net is the better in the class-conditioning task. On average precision it surpasses the classical baseline of 25% on the precision and 40% on the recall. Consequently the F1-score follows the same pattern. Conversely, the QVCU-Net offers only a modest step beyond the classical baseline.






Figure 5.6: Convergence trends over training iterations for different models and evaluation metrics (FID, KID, IS, and Accuracy). Each plot shows the evolution of performance across iterations for the classical baseline, the QVCU-Net and the proposed QCU-Net. In addition to FID (shown in the main text), similar trends are evident for the other metrics. Notably, qCU-Net not only achieves superior scores but also demonstrates a significantly faster convergence compared to the other approaches.

Table 5.5: Comparison of performance metrics for each class across Classical Model, QVCU-Net, and QCU-Net. Best values per row are highlighted in bold.

| Class                | Classical   |      |      |      | QVCU-Net |             |      |      | QCU-Net     |             |             |             | Support |
|----------------------|-------------|------|------|------|----------|-------------|------|------|-------------|-------------|-------------|-------------|---------|
|                      | Prec.       | Rec. | F1   | Acc. | Prec.    | Rec.        | F1   | Acc. | Prec.       | Rec.        | F1          | Acc.        |         |
| AnnualCrop           | 0.61        | 0.83 | 0.71 | 0.83 | 0.66     | 0.84        | 0.74 | 0.83 | <b>0.75</b> | <b>0.92</b> | <b>0.83</b> | <b>0.92</b> | 1027    |
| Forest               | 0.27        | 0.50 | 0.35 | 0.50 | 0.28     | 0.54        | 0.37 | 0.53 | <b>0.88</b> | <b>0.56</b> | <b>0.68</b> | <b>0.55</b> | 1027    |
| HerbaceousVegetation | 0.42        | 0.62 | 0.50 | 0.61 | 0.49     | 0.59        | 0.53 | 0.58 | <b>0.67</b> | <b>0.90</b> | <b>0.77</b> | <b>0.90</b> | 1027    |
| Highway              | 0.81        | 0.59 | 0.68 | 0.58 | 0.85     | 0.62        | 0.72 | 0.61 | <b>0.91</b> | <b>0.74</b> | <b>0.81</b> | <b>0.73</b> | 1027    |
| Industrial           | 0.80        | 0.84 | 0.82 | 0.83 | 0.79     | 0.87        | 0.83 | 0.87 | <b>0.92</b> | <b>0.93</b> | <b>0.93</b> | <b>0.93</b> | 1027    |
| Pasture              | 0.72        | 0.34 | 0.46 | 0.33 | 0.72     | 0.36        | 0.48 | 0.35 | <b>0.81</b> | <b>0.83</b> | <b>0.82</b> | <b>0.83</b> | 1027    |
| PermanentCrop        | 0.74        | 0.29 | 0.41 | 0.28 | 0.68     | 0.37        | 0.48 | 0.36 | <b>0.85</b> | <b>0.66</b> | <b>0.74</b> | <b>0.65</b> | 1027    |
| Residential          | <b>0.98</b> | 0.53 | 0.69 | 0.53 | 0.96     | 0.63        | 0.76 | 0.63 | <b>0.98</b> | <b>0.97</b> | <b>0.97</b> | <b>0.97</b> | 1027    |
| River                | 0.88        | 0.41 | 0.56 | 0.40 | 0.86     | 0.44        | 0.58 | 0.43 | <b>0.90</b> | <b>0.68</b> | <b>0.77</b> | <b>0.68</b> | 1027    |
| SeaLake              | 0.58        | 0.98 | 0.73 | 0.98 | 0.61     | <b>0.99</b> | 0.75 | 0.98 | <b>0.69</b> | <b>0.99</b> | <b>0.81</b> | <b>0.99</b> | 948     |
| <b>Macro Avg</b>     | 0.68        | 0.59 | 0.59 | 0.59 | 0.69     | 0.62        | 0.62 | 0.64 | <b>0.84</b> | <b>0.82</b> | <b>0.81</b> | <b>0.83</b> | 948     |

**Preservation of radiometric statistics:** To evaluate the preservation of the input distribution, a comparative analysis between the distribution of the GT images and the generated ones was conducted, by using a common binning in  $[0, 1]$ , as illustrated in Fig. 5.7. Regarding EO image synthesis tasks, preserving the distribution is crucial because it ensures that the generated images accurately reflect the statistical properties of real-world data. This fidelity is vital for maintaining the integrity and usefulness of the models in practical applications.

QCU-Net demonstrates greater fidelity in producing an output whose distribution more closely aligns

with the original distribution. Specifically, on *AnnualCrop*, QCU-Net closely follows GT in the blue channel and is consistently closer than the classical diffusion baseline in green; in red, all methods show an accumulation near 1.0. Nevertheless, QCU-Net retains a summary statistics (mean and median) nearer to GT than the baseline. On *River*, the classical baseline shows a systematic right-shift across the RGB channels (higher mean and median), resulting in over-brightness compared to GT. By contrast, QCU-Net substantially reduces this bias: its red, green, and blue distributions overlap more closely with the GT. Overall, QCU-Net narrows the pixel-level distribution gap.




Figure 5.7: Per-channel intensity distributions (densities; common bins in  $[0, 1]$ ) for *AnnualCrop* (top) and *River* (bottom). Vertical dashed and dotted lines denote mean and median. We compare the distribution for the GT, Classic model and the QCU-Net model.

### 5.6.2 Ablation Study

**Position Analysis.** Table 5.6 presents the results of an ablation study evaluating the impact of quantum layer placement within the QCU-Net architecture. Two configurations were tested:  $p_1$ , where the quantum layer is applied early in the encoder (via a quanvolutional operation) i.e. 3- $rd$  layer, and  $p_2$ , where it is placed deeper in the network i.e. 4- $th$  layer.

Based on the results it can be noticed that the  $p_1$  setup yields the best performance across all key metrics: the model achieves the lowest FID = 2.57, the lowest KID =  $0.0008 \pm 0.0003$ , IS =  $1.7678 \pm 0.0241$ , and the best classification accuracy on generated samples (81.74%).

These findings suggest that quantum operations are most effective when applied during the early stages of feature extraction. Therefore, it is possible to consider that quantum feature maps have richer spatial and contextual information. In contrast, the applications of quantum operations to a deeper level of model, results in a decreased performance across all metrics. This suggests that applying quantum layers at this stage may limit the ability of quantum features to contribute meaningfully to representation learning.

Table 5.6: Ablation study on position analysis for QCU-Net after 60,000 iterations on the EuroSAT dataset

| Position | FID         | KID                        | IS                         | Accuracy (%) |
|----------|-------------|----------------------------|----------------------------|--------------|
| $p_1$    | <b>2.57</b> | <b>0.0008</b> $\pm 0.0003$ | <b>1.7678</b> $\pm 0.0241$ | <b>81.74</b> |
| $p_2$    | 3.32        | 0.0016 $\pm 0.0004$        | 1.7400 $\pm 0.0242$        | 80.86        |

**Ansatz Analysis** This ablation study focuses on comparing different quantum circuit designs to assess their impact on the quality and fidelity of generated output. Table 5.7 shows results for ansatz presented in Section 5.4.4. Although, all the tested configuration provide comparable results. best performance are reached with HQConv, producing the lowest FID (2.57), the best KID ( $0.0008 \pm 0.0003$ ) and IS ( $1.7780 \pm 0.0241$ ).

These results suggest that the HQConv ansatz can be better suited in capturing the high-dimensional dependencies characteristic of EO imagery, likely due to entanglement operations and structured gate composition. In contrast, the Only-Rotations achieves weaker generative metrics, although still performs adequately in terms of Accuracy.

Table 5.7: Ablation study on ansatz choice for QCU-Net after 60,000 iterations on the EuroSAT dataset

| Ansatz         | FID         | KID                 | IS                  | Accuracy (%) |
|----------------|-------------|---------------------|---------------------|--------------|
| HQConv         | <b>2.57</b> | $0.0008 \pm 0.0003$ | $1.7780 \pm 0.0241$ | 81.74        |
| FQConv         | 3.10        | $0.0013 \pm 0.0004$ | $1.7547 \pm 0.0228$ | 81.31        |
| Only-Rotations | 2.91        | $0.0010 \pm 0.0004$ | $1.7693 \pm 0.0227$ | <b>82.03</b> |

### 5.6.3 Additional studies

**Only Rotation vs Entangled Circuit** To investigate the effect of entanglement, a detailed per-class evaluation comparing two ansatz variants was performed, i.e. only rotation with no entanglement and the HQConv circuits.

Table 5.8, shows that the use of entanglement led to some improvement on the classification performance. Indeed, the F1-scores increases in challenging classes e.g. River ( $0.72 \rightarrow 0.75$ ). However, in other cases the no-entanglement version achieved slightly higher precision and recall, suggesting that the benefit of entanglement may be context-dependent.

Table 5.8: Per-class comparison of precision, recall, and F1-score between only rotation with no entanglement (N) and HQConv (E) circuits for QCU-Net.

| Class                | Prec (N/E)                | Rec (N/E)                 | F1 (N/E)                  | Support |
|----------------------|---------------------------|---------------------------|---------------------------|---------|
| AnnualCrop           | <b>0.73</b> / 0.70        | <b>0.91</b> / 0.89        | <b>0.81</b> / 0.79        | 1027    |
| Forest               | <b>0.89</b> / 0.87        | <b>0.67</b> / 0.54        | <b>0.77</b> / 0.67        | 1027    |
| HerbaceousVegetation | 0.71 / <b>0.72</b>        | <b>0.90</b> / 0.89        | <b>0.80</b> / <b>0.80</b> | 1027    |
| Highway              | 0.87 / <b>0.88</b>        | 0.79 / <b>0.80</b>        | 0.83 / <b>0.84</b>        | 1027    |
| Industrial           | 0.92 / <b>0.93</b>        | <b>0.96</b> / 0.93        | <b>0.94</b> / 0.93        | 1027    |
| Pasture              | 0.81 / <b>0.82</b>        | 0.80 / <b>0.83</b>        | 0.81 / <b>0.82</b>        | 1027    |
| PermanentCrop        | 0.85 / <b>0.90</b>        | <b>0.61</b> / 0.60        | 0.71 / <b>0.72</b>        | 1027    |
| Residential          | <b>0.97</b> / 0.96        | <b>0.98</b> / <b>0.98</b> | <b>0.97</b> / <b>0.97</b> | 1027    |
| River                | <b>0.92</b> / <b>0.92</b> | 0.60 / <b>0.63</b>        | 0.72 / <b>0.75</b>        | 948     |
| SeaLake              | <b>0.68</b> / 0.62        | <b>1.00</b> / 0.99        | <b>0.81</b> / 0.76        | 948     |

**Hyperparameters Tuning:** Table 5.9 presents the results of the hyperparameters tuning process. The experiments reveal that the learning rate has the most significant impact on model performance. A high learning rate ( $1e^{-2}$ ) leads to unstable training and poor generative quality (a very high FID = 206.82 and low IS =  $1.03 \pm 0.002$ ). Conversely, reducing the learning rate to  $1e^{-3}$  give improvements across all metrics (FID dropping to as low as 7.83 and IS rising to  $1.89 \pm 0.0268$ ).

Further fine-tuning of the  $\beta_1$  parameter show that the value 1.95 provides the best results (FID = 7.83), indicating improved stability and convergence behavior. The variation of  $\beta_2$  show a smaller but still noticeable impact, with values around 0.990 producing better performance.

Than, while keeping the best lr,  $\beta_1$ , and  $\beta_2$  fixed, it was noticed that raising the numerical-stability term from  $\epsilon = 10^{-8}$  to  $\epsilon = 10^{-7}$  reduced the FID from 7.83 to 7.22 and KID from  $3.9 \times 10^{-3}$  to  $3.4 \times 10^{-3}$ , while reducing it to  $10^{-9}$  pushed both metrics in the opposite direction. It was therefore selected  $\epsilon = 10^{-7}$ .

Table 5.9: Results with hyperparameters tuning for the classical baseline.

| lr   | $\beta_1$ | $\beta_2$ | $\epsilon$ | FID ↓         | KID ↓                | IS ↑                 |
|------|-----------|-----------|------------|---------------|----------------------|----------------------|
| 1e−2 | 1.90      | 0.990     | 1e−8       | 206.8200      | $0.2288 \pm 0.0016$  | $1.0269 \pm 0.0020$  |
| 1e−3 | 1.90      | 0.990     | 1e−8       | 8.2600        | $0.0046 \pm 0.0008$  | $1.8718 \pm 0.0233$  |
| 1e−4 | 1.90      | 0.990     | 1e−8       | 11.5200       | $0.0057 \pm 0.0010$  | $1.8138 \pm 0.0269$  |
| 1e−3 | 1.85      | 0.990     | 1e−8       | 8.0300        | $0.0040 \pm 0.0007$  | $1.8977 \pm 0.0231$  |
| 1e−3 | 1.95      | 0.990     | 1e−8       | 7.8300        | $0.0039 \pm 0.0007$  | $1.8922 \pm 0.0268$  |
| 1e−3 | 1.95      | 0.895     | 1e−8       | 8.4902        | $0.0048 \pm 0.0009$  | $1.9005 \pm 0.0249$  |
| 1e−3 | 1.95      | 0.995     | 1e−8       | 8.2423        | $0.0042 \pm 0.0008$  | $1.9033 \pm 0.0312$  |
| 1e−3 | 1.95      | 0.990     | 1e−7       | <b>7.2248</b> | <b>0.0034±0.0008</b> | $1.8837 \pm 0.0306$  |
| 1e−3 | 1.95      | 0.990     | 1e−9       | 8.2325        | $0.0042 \pm 0.0008$  | <b>1.9210±0.0324</b> |

## 5.7 Conclusion

In this chapter, the integration of QC with class-conditioned diffusion models for generating synthetic EO imagery was introduced. By combining the flexibility and modularity of classical architectures, such as U-Net, with quantum circuits, the proposed QCU-Net represents a significant quantum utility for advancement in generative modeling for EO tasks.

The results highlight that the introduction of quantum layers into the diffusion process considerably enhances both the realism and accuracy of the generated images, especially when these layers are inserted in the encoder branch of the U-Net. Specifically, the QCU-Net model demonstrated marked improvements over both classical and other simpler quantum models, which were included in this analysis to ensure a fair comparison and to highlight the potential utility of the proposed quantum framework. In terms of FID and KID, QCU-Net generates samples that are much closer to the true distribution of real satellite images. Notably, the QCU-Net model showed a 64% improvement in FID and a 76% reduction in KID compared to classical models, proving its superior performance. Furthermore, the proposed model exhibited improved semantic accuracy, particularly in challenging classes such as AnnualCrop, Industrial, and River, which are often difficult for conventional models due to high inter-class variability. In addition to the hybrid architecture, this work also investigated a QLDM, which integrates quantum circuits directly into the compressed representation of the data. The QLDM demonstrated clear improvements over its classical counterpart and over the latent GAN baselines, confirming that quantum layers can effectively enhance denoising and generative capabilities even under strong dimensionality reduction. However, the severe compression required to operate in the latent space, necessary to keep quantum simulation tractable, inevitably limits the preservation of spatial and radiometric detail, which is crucial in EO applications. These observations motivated the transition toward a hybrid, image-domain architecture, where quantum modules can act on richer feature representations without the constraints imposed by latent compression.

A unique contribution of this study is the class-conditioning mechanism, which enables the generation of labeled EO imagery directly applicable to real-world RS tasks, such as data augmentation or class-specific simulation. This is particularly valuable in scenarios where labeled data are scarce or costly to acquire. By combining class-conditioned QDMs with U-Net architectures, the proposed approach demonstrates how quantum circuits can significantly enhance generative modeling, especially in RS, where high-dimensional data and labeled images are often scarce.

The results of the conducted ablation studies support this, revealing that the integration of quantum layers in the middle layers of the model, within the feature extraction stages, leads to a notable performance boost. This finding underscores the importance of strategically placing quantum circuits within the model to achieve optimal results. Additionally, the HQConv quantum circuit ansatz, which incorporates entangling gates, proved effective in capturing the complex dependencies inherent in EO images, further emphasizing the role of quantum entanglement in enhancing generative models. However, in some cases the no-entanglement version achieved slightly better results, suggesting that the benefit of entanglement may be context-dependent and merits further investigation.

To offer future perspectives, it is important to highlight some challenges of the proposed approach. The QCU-Net, while outperforming the classical approach, still relies on classical components for certain stages, especially in managing high-dimensional feature maps in the encoder. Therefore, it necessitates careful optimization of both quantum and classical layers, leading to a trade-off between performance and computational feasibility. Moreover, the reliance on the EuroSAT RGB dataset, though widely used, limits the generalizability of the obtained results to other EO datasets that may differ in spectral bands or spatial resolutions. Additionally, the computational demands of quantum-classical hybrid models require substantial resources, posing a potential barrier to wider adoption, especially given the current limitations of NISQ devices.

Looking ahead, future research should explore additional quantum features, such as quantum noise-driven diffusion and quantum error correction, which could address the limitations of current NISQ devices and enhance model robustness. Furthermore, the scalability and efficiency of QDMs across various quantum hardware platforms remain critical areas of investigation, as the computational demands of quantum circuits are still a significant challenge. Finally, extending this work to multi-modal EO datasets that incorporate other data sources, such as SAR or LiDAR, could facilitate more comprehensive and multi-sensor EO applications.

In any case, the proposed framework has demonstrated to represent a powerful tool for the EO domain and the community of RS researchers, paving the way for quantum-enhanced applications in RS.

## Bibliography

- [1] F. De Falco, **F. Mauro**, A. Ceschini, A. Sebastianelli, P. Gamba, S. L. Ullo, and M. Panella, “Leveraging Quantum Latent Diffusion Models for Data Augmentation on the EuroSAT Dataset,” *IGARSS 2025 - 2025 IEEE International Geoscience and Remote Sensing Symposium*, 2025.
- [2] **F. Mauro**, F. De Falco, L. Papa, A. Ceschini, A. Sebastianelli, P. Gamba, M. Panella, and S. L. Ullo, “Enriching earth observation labeled data with quantum conditioned diffusion models,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 64, pp. 1–16, 2026.
- [3] A. Alzahem, W. Boulila, A. Koubaa, Z. Khan, and I. Alturki, “Improving satellite image classification accuracy using GAN-based data augmentation and vision transformers,” *Earth Science Informatics*, vol. 16, no. 4, pp. 4169–4186, 2023.
- [4] O. Adedeji, P. Owoade, O. Ajayi, and O. Arowolo, “Image augmentation for satellite images,” *arXiv preprint arXiv:2207.14580*, 2022.
- [5] S. Ghelichkhani, Y. A. M. A.-D. Salem, H. Salahov, F. A. I. Adam, A. J. Charbatji, and M. I. Abdulkareem, “Generative adversarial networks for artificial satellite image creation and manipulation,” in *Adversarial Multimedia Forensics*. Springer, 2023, pp. 125–143.
- [6] A. Sebastianelli, E. Puglisi, M. P. Del Rosso, J. Mifdal, A. Nowakowski, P. P. Mathieu, F. Pirri, and S. L. Ullo, “PLFM: Pixel-level merging of intermediate feature maps by disentangling and fusing spatial and temporal data for cloud removal,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2022.
- [7] D. A. Abuhani, I. Zualkernan, R. Aldamani, and M. Alshafai, “Generative Artificial Intelligence for Hyperspectral Sensor Data: A Review,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 6422–6439, 2025.
- [8] “Multispectral satellite image generation using StyleGAN3, author=Alibani, Michael and Acito, Nicola and Corsini, Giovanni,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 4379–4391, 2024.
- [9] S. Filali Boubrahimi, A. Neema, A. Nassar, P. Hosseinzadeh, and S. M. Hamdi, “Spatiotemporal data augmentation of MODIS-landsat water bodies using adversarial networks,” *Water Resources Research*, vol. 60, no. 3, p. e2023WR036342, 2024.
- [10] Y. Liu, J. Yue, S. Xia, P. Ghamisi, W. Xie, and L. Fang, “Diffusion models meet remote sensing: Principles, methods, and perspectives,” *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [11] D. Tang, X. Cao, X. Hou, Z. Jiang, J. Liu, and D. Meng, “Crs-diff: Controllable remote sensing image generation with diffusion model,” *IEEE Transactions on Geoscience and Remote Sensing*, 2024.
- [12] A. Sebaq and M. ElHelw, “Rsdiff: Remote sensing image generation from text using diffusion model,” *Neural Computing and Applications*, vol. 36, no. 36, pp. 23 103–23 111, 2024.
- [13] E. Farhi and H. Neven, “Classification with Quantum Neural Networks on Near Term Processors,” 2018.
- [14] M. Schuld, M. Fingerhuth, and F. Petruccione, “Implementing a distance-based classifier with a quantum interference circuit,” *EPL (Europhysics Letters)*, vol. 119, no. 6, p. 60002, sep 2017. [Online]. Available: <https://doi.org/10.1209%2F0295-5075%2F119%2F60002>

- 
- [15] E. Aïmeur, G. Brassard, and S. Gambs, “Quantum speed-up for unsupervised learning,” *Machine Learning*, vol. 90, no. 2, pp. 261–287, Feb. 2013. [Online]. Available: <https://inria.hal.science/hal-00736948>
- [16] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Science and Technology*, vol. 4, no. 4, p. 043001, nov 2019. [Online]. Available: <https://doi.org/10.1088%2F2058-9565%2Fab4eb5>
- [17] H.-L. Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li, F. Liang, J. Lin, Y. Xu, R. Yang, T. Liu, M.-H. Hsieh, H. Deng, H. Rong, C.-Z. Peng, C.-Y. Lu, Y.-A. Chen, D. Tao, X. Zhu, and J.-W. Pan, “Experimental quantum generative adversarial networks for image generation,” *Phys. Rev. Appl.*, vol. 16, p. 024051, Aug 2021. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.16.024051>
- [18] F. De Falco, A. Ceschini, A. Sebastianelli, B. Le Saux, and M. Panella, “Quantum Hybrid Diffusion Models for Image Synthesis,” *KI - Künstliche Intelligenz*, vol. 38, pp. 311–326, 2024.
- [19] —, “Quantum latent diffusion models,” *Quantum Machine Intelligence*, vol. 6, no. 2, p. 85, 2024.
- [20] S. Bravyi, D. Gosset, and R. König, “Quantum advantage with shallow circuits,” *Science*, vol. 362, no. 6412, pp. 308–311, oct 2018. [Online]. Available: <https://doi.org/10.1126%2Fscience.aar3106>
- [21] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, jun 2021. [Online]. Available: <https://doi.org/10.1038%2Fs43588-021-00084-1>
- [22] S. Y. Chang, S. Thanasilp, B. L. Saux, S. Vallecorsa, and M. Grossi, “Latent Style-based Quantum GAN for high-quality Image Generation,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.02668>
- [23] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International Conference on Machine Learning*. PMLR, 2015, pp. 2256–2265.
- [24] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 6840–6851, 2020.
- [25] A. Q. Nichol and P. Dhariwal, “Improved denoising diffusion probabilistic models,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 8162–8171.
- [26] A. Kamel and B. Issam, “Data Augmentation for Land Cover Classification Using Generative Adversarial Networks,” in *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*. IEEE, 2021, pp. 2309–2312.
- [27] P. Sarukkai, S. Xu, J. Kelly, Y. Lu, M. Konsolakis, and G. DiCaprio, “Cloud Removal from Satellite Images using Spatiotemporal Generator Networks,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2020, pp. 1792–1801.
- [28] P. Ebel, M. Schmitt, and X. X. Zhu, “Cloud removal in unpaired sentinel-2 imagery using cycle-consistent gan and sar-optical data fusion,” in *IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium*. IEEE, 2020, pp. 2065–2068.
-



- 
- [29] F. Xu, Y. Shi, P. Ebel, L. Yu, G.-S. Xia, W. Yang, and X. X. Zhu, “GLF-CR: SAR-enhanced cloud removal with global–local fusion,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 192, pp. 268–278, 2022.
- [30] M. Gonzalez-Calabuig, M.-Á. Fernández-Torres, and G. Camps-Valls, “Generative networks for spatio-temporal gap filling of Sentinel-2 reflectances,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 220, pp. 637–648, 2025.
- [31] J. Han, Y. Zhou, X. Gao, and Y. Zhao, “Thin Cloud Removal Generative Adversarial Network Based on Sparse Transformer in Remote Sensing Images,” *Remote Sensing*, vol. 16, no. 19, p. 3658, 2024.
- [32] D. Liu, L. Zhong, H. Wu, S. Li, and Y. Li, “Remote sensing image Super-resolution reconstruction by fusing multi-scale receptive fields and hybrid transformer,” *Scientific Reports*, vol. 15, no. 1, p. 2140, 2025.
- [33] J. He, X. Yuan, Y. Kou, and Y. Zhang, “A Novel Framework for Remote Sensing Image Synthesis with Optimal Transport,” *Sensors (Basel, Switzerland)*, vol. 25, no. 6, p. 1792, 2025.
- [34] J. Wang, X. Zhu, L. Jing, Y. Tang, H. Li, Z. Xiao, and H. Ding, “HyperGAN: A Hyperspectral Image Fusion Approach Based on Generative Adversarial Networks,” *Remote Sensing*, vol. 16, no. 23, p. 4389, 2024.
- [35] C. Sun, X. Zhang, H. Meng, X. Cao, and J. Zhang, “AC-WGAN-GP: Generating labeled samples for improving hyperspectral image classification with small-samples,” *Remote Sensing*, vol. 14, no. 19, p. 4910, 2022.
- [36] Y. Yuan, L. Lin, Q. Xin, Z.-G. Zhou, and Q. Liu, “An Empirical Study on Data Augmentation for Pixel-Wise Satellite Image Time Series Classification and Cross-Year Adaptation,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2025.
- [37] S. Sastry, S. Khanal, A. Dhakal, and N. Jacobs, “GeoSynth: Contextually-Aware High-Resolution Satellite Image Synthesis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 460–470.
- [38] F. Sanguigni, M. Czerkawski, L. Papa, I. Amerini, and B. L. Saux, “Diffusion models for earth observation use-cases: from cloud removal to urban change detection,” *arXiv preprint arXiv:2311.06222*, 2023.
- [39] L. Pang, X. Cao, D. Tang, S. Xu, X. Bai, F. Zhou, and D. Meng, “Hsigene: A foundation model for hyperspectral image generation,” *arXiv preprint arXiv:2409.12470*, 2024.
- [40] K. S.-H. Lo and J. Peters, “Diff-DEM: A Diffusion Probabilistic Approach to Digital Elevation Model Void Filling,” *IEEE Geoscience and Remote Sensing Letters*, 2024.
- [41] F. Zhou, H. Li, R. Hu, S. Wu, H. Feng, Z. Du, and L. Xu, “ControlCity: A Multimodal Diffusion Model Based Approach for Accurate Geospatial Data Generation and Urban Morphology Analysis,” *arXiv preprint arXiv:2409.17049*, 2024.
- [42] T. Sousa, B. Ries, and N. Guelfi, “Data augmentation in earth observation: A diffusion model approach,” *arXiv preprint arXiv:2406.06218*, 2024.
- [43] S. Khanna, P. Liu, L. Zhou, C. Meng, R. Rombach, M. Burke, D. B. Lobell, and S. Ermon, “DiffusionSat: A Generative Foundation Model for Satellite Imagery,” *arXiv preprint arXiv:2308.08369*, 2023.
-

- 
- [44] L. Liu, B. Chen, H. Chen, Z. Zou, and Z. Shi, “Diverse hyperspectral remote sensing image synthesis with diffusion models,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 61, pp. 1–16, 2023.
  - [45] J. Häberle, H. Hoffmann, X. Zhu, M. Schmitt, J. Eberle, and S. Hinz, “COP-GEN-Beta: Unified Generative Modelling of Copernicus Imagery Thumbnails,” *arXiv preprint arXiv:2504.08548*, 2025.
  - [46] Y. Cao, S. Geng, J. Zhao, and X. Wang, “LiteDiT: An Efficient Diffusion Transformer Model for Remote Sensing Image Synthesis Focus on Object,” in *2024 International Conference on Artificial Intelligence and Digital Technology (ICAIDT)*. IEEE, 2024, pp. 116–119.
  - [47] L. Papa, P. Russo, I. Amerini, and L. Zhou, “A survey on efficient vision transformers: Algorithms, techniques, and performance benchmarking,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, p. 7682–7700, Dec. 2024. [Online]. Available: <http://dx.doi.org/10.1109/TPAMI.2024.3392941>
  - [48] M. Hibat-Allah, M. Mauri, J. Carrasquilla, and A. Perdomo-Ortiz, “A framework for demonstrating practical quantum advantage: comparing quantum against classical generative models,” *Communications Physics*, vol. 7, no. 1, p. 68, 2024.
  - [49] C. Chen, Q. Zhao, M. Zhou, Z. He, Z. Sun, and H. Situ, “Quantum Generative Diffusion Model: A Fully Quantum-Mechanical Model for Generating Quantum State Ensemble,” *arXiv preprint arXiv:2401.07039*, 2024.
  - [50] B. Zhang, P. Xu, X. Chen, and Q. Zhuang, “Generative quantum machine learning via denoising diffusion probabilistic models,” *Physical Review Letters*, vol. 132, no. 10, p. 100602, 2024.
  - [51] M. Parigi, S. Martina, and F. Caruso, “Quantum-Noise-Driven Generative Diffusion Models,” *Advanced Quantum Technologies*, p. 2300401, 2024.
  - [52] M. Henderson, S. Shakya, S. Pradhan, and T. Cook, “Quanvolutional Neural Networks: Powering Image Recognition with Quantum Circuits,” 2019.
  - [53] A. Sebastianelli, **F. Mauro**, G. Ciabatti, D. Spiller, B. L. Saux, P. Gamba, and S. Ullo, “Quanv4EO: Empowering Earth Observation by means of Quanvolutional Neural Networks,” *IEEE Transactions on Geoscience and Remote Sensing*, pp. 1–1, 2025.
  - [54] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.
  - [55] A. Q. Nichol and P. Dhariwal, “Improved Denoising Diffusion Probabilistic Models,” in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 8162–8171. [Online]. Available: <https://proceedings.mlr.press/v139/nichol21a.html>
  - [56] J. M. Arrazola, O. Di Matteo, N. Quesada, S. Jahangiri, A. Delgado, and N. Killoran, “Universal quantum circuits for quantum chemistry,” *Quantum*, vol. 6, p. 742, Jun. 2022. [Online]. Available: <http://dx.doi.org/10.22331/q-2022-06-20-742>
-

- 
- [57] R. Jozsa and A. Miyake, “Matchgates and classical simulation of quantum circuits,” *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 464, no. 2100, p. 3089–3106, Jul. 2008. [Online]. Available: <http://dx.doi.org/10.1098/rspa.2008.0189>
- [58] Y. Jing, X. Li, Y. Yang, C. Wu, W. Fu, W. Hu, Y. Li, and H. Xu, “RGB image classification with quantum convolutional ansatz,” *Quantum Information Processing*, vol. 21, no. 3, Feb. 2022. [Online]. Available: <http://dx.doi.org/10.1007/s11128-022-03442-8>
- [59] M. Harini, S. Selvavarshini, P. Narmatha, V. Anitha, S. K. Selvi, and V. Manimaran, “Resnet-50 integrated with attention mechanism for remote sensing classification,” in *International conference on advances in distributed computing and machine learning*. Springer, 2024, pp. 255–265.
- [60] A. Bhatt and V. T. Bhatt, “Dcrff-Lhrf: an improvised methodology for efficient land-cover classification on eurosat dataset,” *Multimedia Tools and Applications*, vol. 83, no. 18, pp. 54 001–54 025, 2024.
- [61] S. Bobba, “Leveraging Pre-trained Deep Learning Models for Remote Sensing Image Classification: A Case Study with ResNet50 and EfficientNet,” *American Journal of Science, Engineering and Technology*, vol. 9, no. 3, pp. 150–162, 2024.
- [62] H. Dastour and Q. K. Hassan, “A comparison of deep transfer learning methods for land use and land cover classification,” *Sustainability*, vol. 15, no. 10, p. 7854, 2023.
- [63] S. Chan, P. Kulkarni, P. H. Yi, and V. S. Parekh, “Expanding the horizon: Enabling hybrid quantum transfer learning for long-tailed chest x-ray classification,” 2024. [Online]. Available: <https://arxiv.org/abs/2405.00156>
- [64] N. S. Detlefsen, J. Borovec, J. Schock, A. H. Jha, T. Koker, L. Di Liello, D. Stancl, C. Quan, M. Grechkin, and W. Falcon, “Torchmetrics-measuring reproducibility in pytorch,” *Journal of Open Source Software*, vol. 7, no. 70, p. 4101, 2022.
- [65] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [66] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium,” 2018. [Online]. Available: <https://arxiv.org/abs/1706.08500>
- [67] M. Sokolova and G. Lapalme, “A systematic analysis of performance measures for classification tasks,” *Information Processing & Management*, vol. 45, no. 4, pp. 427–437, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306457309000259>
- [68] P. Helber, B. Bischke, A. Dengel, and D. Borth, “Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
-



## Chapter 6

# Future Prospects and Conclusions

This dissertation investigated the role of QML4EO, with the overarching goal of understanding *whether, where, and how* quantum paradigms can complement SOTA ML and DL for RS under current NISQ constraints. By combining a structured review of the emerging landscape, a set of hybrid quantum–classical methodological contributions, and an empirical validation strategy across representative EO tasks, spanning classification, regression, segmentation, forecasting and generation, this work has aimed to move beyond speculative claims toward measurable evidence.

A central premise emerging across all chapters is that, in the NISQ era, the most credible pathway is not a fully-quantum end-to-end replacement of mature classical pipelines, but a *co-design paradigm*: classical backbones provide scalability, optimization stability, and mature inductive biases, while quantum modules are explored as specialized operators (such as feature maps extractors) that introduce novel representational capacity or favorable trade-offs in several scenarios.

This chapter closes the dissertation by connecting the research questions presented in Chapter 1 with the empirical and methodological evidence developed throughout the thesis, spanning Quanv4EO, QGNNs, and QDMs. The resulting synthesis is used to articulate the main outcomes and to clarify the limitations that currently shape QML4EO in the NISQ era. The chapter then outlines the most credible research directions and ends by discussing where *domain-relevant* quantum advantage may realistically emerge in EO.

### 6.1 Mapping the findings to the Research Questions

The dissertation was structured around four research questions. The following paragraphs summarize the answers supported by the empirical evidence reported in Chapters 3–5.

#### **RQ1: In which contexts do QML models provide a tangible advantage over traditional DL approaches?**

The evidence reported in this dissertation does not justify statements of a broad quantum advantage over SOTA DL for operational EO. Under current NISQ constraints, a more accurate interpretation is that of *quantum-assisted* ML/DL: the empirical gains, when present, emerge from hybrid pipelines in which a quantum block performs a specific transformation, while the classical backbone preserves scalability and training stability. In other words, the relevant question is not whether quantum models replace DL, but whether they can add measurable value when used as a targeted component.

Within this framing, the strongest and most recurrent outcome is a favourable trade-off between performance and model size. In the **Quanv4EO** experiments (Chapter 3), competitive building segmentation

performance on the Tunis area is obtained with more than 90% fewer trainable parameters than a classical Attention U-Net, indicating that local quantum feature operators can support compact discriminative models. A consistent pattern appears in the  $\Phi$ Sat-2 water-quality regression task, where high predictive quality is maintained (Pearson correlation  $\sim 0.90$ ) despite a drastically reduced parameter budget. These results suggest that, in the near term, the main benefit is not higher peak accuracy, but higher *expressivity per parameter*.

A related message arises from the generative models (Chapter 5). Hybrid diffusion models that introduce quantum residual components improve sample fidelity under controlled protocols, provided that the quantum part is placed where it can act on informative feature representations. This again points to *targeted hybridization* rather than to a quantum end-to-end alternative: the impact depends on *what* is hybridized and *where* the quantum block is inserted (encoding, depth, and interface with the classical network).

Overall, the dissertation supports a pragmatic conclusion. The most credible near-term gains in QML4EO are expected to be localized within larger learning pipelines, while the broader challenge becomes identifying which submodules benefit from hybridization and defining implementation choices that remain robust and reproducible under realistic constraints.

## **RQ2: How can hybrid quantum–classical architectures be effectively integrated into EO tasks?**

The dissertation successfully validated multiple integration patterns viable under NISQ constraints, moving beyond simple proof-of-concept designs toward structurally integrated hybrids. A primary effective strategy involves the deployment of *Front-end Quantum Feature Maps* (Quanv4EO), where fixed or lightly-trainable quantum circuits act as local kernel extractors. This approach facilitates the ingestion of high-dimensional EO data, such as Sentinel-1, Sentinel-2 and  $\Phi$ Sat-2 imagery, by leveraging the Hilbert space for feature encoding.

Building upon this, the research explored the utility of *Quantum-Assisted Graph Heads* (QGNNs) for modeling complex spatio-temporal dependencies. As demonstrated in the ONI prediction task (Chapter 4), inserting a VQC into a graph-based pipeline allows for the efficient capture of global teleconnections; notably, this architecture achieved convergence in significantly fewer iterations: 5 epochs compared to the 50 required by purely classical benchmarks.

Finally, a deeper level of integration was achieved through *Embedded Quantum Residuals* within generative frameworks (Chapter 5). By placing quantum layers directly within the encoder branch of a U-Net (QCU-Net), the quantum circuit was able to influence the underlying feature geometry while the classical backbone ensured the maintenance of spatial coherence. This specific configuration proved more robust than fully quantum latent-space approaches, demonstrating that quantum components can effectively act as specialized feature refiners within established DL architectures.

## **RQ3: Can quantum-enhanced feature extraction improve representational power for EO data?**

Yes, specifically by projecting data into high-dimensional Hilbert spaces that capture correlations often missed by classical kernels. The *Quanv4EO* framework demonstrated that quanvolutional filters can extract features from SAR and multispectral imagery that are robust and semantically rich enough to support downstream tasks like filtering, classification, regression and segmentation with minimal classical post-processing. The consistent performance of random quantum circuits in these tasks further suggests that the *structure* of the quantum feature map itself acts as a strong inductive bias for EO data.

## RQ4: To what extent can quantum probabilistic models enhance the generation of synthetic EO data?

The investigation into *QDMs* (Chapter 5) provided strong evidence that quantum components can improve generative fidelity. The proposed QCU-Net achieved a 64% reduction in FID and a 76% reduction in KID compared to classical baselines on the EuroSAT dataset. Furthermore, the ability to generate class-conditioned samples with high semantic accuracy (particularly for challenging classes like *River* and *Industrial*) indicates that quantum circuits can effectively model complex joint distributions, offering a viable path for data augmentation in label-scarce regimes.

## 6.2 Main Outcomes

### 6.2.1 From field overview to actionable gaps

The comprehensive review in Chapter 2 highlighted that while QML4EO is expanding, it remains fragmented and heavily reliant on simulators. A key outcome of this thesis was the shift from simple classification benchmarks to complex, structured tasks. By addressing classification, regression, segmentation, and time-series forecasting, this work has broadened the operational scope of QML, moving beyond the "toy problem" phase toward relevant geophysical applications.

### 6.2.2 Hybrid architectures are the practical path

A consistent conclusion is that hybrid designs are a *methodological necessity*. Fully quantum models (e.g., the pure QLDLM investigated in Chapter 5) suffer from severe information bottlenecks due to the need for extreme dimensionality reduction. In contrast, the QCU-Net demonstrates that embedding quantum modules as specialized operators within classical backbones allows to leverage quantum expressivity without sacrificing the spatial resolution required for EO.

### 6.2.3 Where Does This Leave “Quantum Advantage”?

The evidence collected in this work suggests that an unambiguous, end-to-end quantum speed-up for operational EO pipelines is not yet within reach; however, the data also reveal encouraging trends that deserve equal attention. First, hardware scaling is beginning to intersect with EO-relevant problem sizes. When the field launched in 2017, public devices exposed fewer than 20 high-fidelity qubits; today laboratory prototypes already pass the one-thousand-qubit mark, and cloud back-ends with 100–200 usable qubits are routinely available. Although error-rates and qubit connectivity still limit circuit depth, the ability to embed tens of independent EO features without extreme dimensionality reduction is now realistic. This leap removes one of the most frequently cited “show-stoppers” and enables explorations that were previously confined to idealised simulators.

Second, hybrid quantum–classical training loops are maturing quickly. Gradient-free optimisers, layer-wise learning-rate schedules and automatic shot-allocation have collectively contributed to the reduction of wall-clock training time. These engineering improvements do not constitute “quantum advantage” in a strict complexity-theoretic sense, but they do narrow the practical gap between research prototypes and production-grade EO analytics, particularly for tasks where classical meta-learning struggles with vanishing gradients or prohibitive memory footprints.

Third, community practice is becoming more rigorous. More than half of the papers published after 2024 deposit circuit descriptions, random seeds and pre-processing scripts in public repositories, making independent replication possible for the first time. In parallel, several open-access benchmark suites

(e.g., QEarthNet-B and Q4SAR-Lite) now specify fixed qubit budgets, shot limits and energy-usage reporting. These resources provide the uniform baselines that are essential for credible “advantage” claims and will help filter genuine algorithmic breakthroughs from artefacts of hyper-parameter tuning.

Taken together, these observations motivate a more nuanced position. QML4EO has transitioned from speculative curiosity to methodologically sound, steadily progressing research programme. Near-term “quantum advantage” will almost certainly be task-specific and resource-bounded, e.g., accelerated kernel evaluations in small-footprint edge constellations, or improved robustness against covariate shift in multi-modal fusion, rather than a sweeping replacement of classical DL. The path forward is therefore clear:

1. Target EO problems whose structure aligns with currently available qubit counts and noise profiles;
2. Report holistic metrics, including latency, energy and reproducibility, so that efficiency claims remain grounded; and
3. Invest in open benchmarks that couple real-world data complexity with transparent, hardware-aware constraints.

In that frame, the question “Where does this leave quantum advantage?” can be answered cautiously but optimistically: on the horizon, no longer on the drawing board. While these results do not yet constitute conclusive proof, the convergence of maturing hardware, sharper algorithms and disciplined benchmarking suggests that domain-relevant advantage, though likely incremental rather than monumental, may emerge sooner than many sceptics predicted.

## 6.3 Current Limitations

Despite the promising results, several limitations identified during this research define the boundaries of current applicability.

### 6.3.1 Encoding and information bottlenecks

EO data is high-dimensional. The process of encoding dense satellite imagery (e.g.,  $64 \times 64$  multi-spectral patches) into limited qubit registers remains the primary bottleneck. While *angle encoding* and *quantum convolutional patching* provided an efficient way of managing EO data, they are not scalable solutions for full-scene processing. Future progress will depend on the development of EO-aware encoding strategies, such as tensor-network embeddings or physics-informed encodings, that preserve spectral coherence without requiring prohibitive qubit counts.

### 6.3.2 Scalability and the simulation-to-hardware gap

As noted in Chapter 2 and verified throughout the experiments, the majority of current research relies on noise-free simulations. While this thesis utilized robust frameworks (PennyLane/JAX) to simulate up to 16 qubits, real hardware implementation faces noise, connectivity constraints, and shot noise. Consequently, the “advantage” reported in terms of convergence and accuracy must be weighed against the overhead of error mitigation and cryogenic control in a real-world deployment.



### 6.3.3 Benchmarking and Reproducibility

A recurring challenge was the lack of standardized baselines. To address this, the dissertation prioritized rigorous comparisons against optimized classical models (e.g., Attention U-Net, ResNet-50) using identical pre-processing and splits. This approach revealed that while quantum models are highly parameter-efficient, they do not always win on absolute accuracy, highlighting the importance of holistic metrics that include training time, model size, and inference energy.

Given the diversity of these integration patterns, it is becoming increasingly evident that establishing a set of general guidelines for validating QML4EO approaches is essential. The definition of standardized protocols for data encoding, circuit initialization, and benchmarking would facilitate more transparent and fair comparisons across different architectures. Such a framework would not only ensure the reproducibility of results but also provide a rigorous basis for assessing the actual quantum advantage over SOTA classical models in EO tasks.

## 6.4 Future Prospects and Research Directions

Building on the contributions of this dissertation, the following directions appear most impactful for the next phase of QML4EO.

A first important step is to fuse QC theory with EO data physics and rather than optimising for a single application, it is necessary to design *general-purpose quantum building blocks* that can transfer across EO domains. This thesis already supports this vision by introducing the *quanvolutional layer*, a quantum feature extractor that has since been applied to diverse EO tasks. An essential component within the quanvolutional framework, and, more generally, within any QML algorithm, lies in the data encoding and measurement stages. These stages determine how classical EO information is mapped to quantum states and subsequently read out, directly influencing the performance of the subsequent model. For these reasons, the first thread of future work is the design of suitable encoding strategies for EO data.

A second, essential thread concerns hardware suitability for EO tasks. Given the current scarcity of studies deploying real EO pipelines on quantum hardware, a future goal is to identify which platforms and technologies are most appropriate for such applications. This path will start with a systematic characterisation of hardware noise and stability by encoding and measuring EO data directly on different devices and progress towards the deployment of complete algorithms. This activity also represents a tangible progression in Technology Readiness Level (TRL) maturity, moving from purely simulated environments toward experimental verification on quantum devices.

A third thread could explore the possibility of bringing QC closer to on-board computation. Once shallow, few-qubit circuits demonstrate stable and meaningful signals, future works could outline a minimal path towards *quantum-on-board*. Pursuing this line could therefore bridge two complementary research domains, QC and AI at the edge, especially as emerging photonic quantum computers appears increasingly suitable for in-orbit deployment.

These research directions could bridge QC with operational EO tasks, following a realistic roadmap that connects the maturity of quantum hardware with the type of problems that can be addressed. As current devices gradually increase in qubit count, coherence time, and noise control, it becomes possible to move from small experimental tasks to more demanding applications. In the short term, use cases presented in this dissertation, such as *SAR image denoising and feature extraction*, *quantum-enhanced land classification*, and *water quality assessment* can already be tested on available NISQ platforms such as IBM Quantum, IonQ, or Rigetti. QDMs also offer potential for EO data generation. Moreover, with the next era of hardware, more complex use cases like *multi-sensor data fusion*, *mission planning for EO acquisition*, and *regional climate modelling* will become feasible. In the longer term, large-scale

systems could support *near-real-time EO data processing*, *global environmental simulations*, and *mission-level optimisation for Digital Twin Earth*.

Now, it is important to prioritise *available* and *emerging* NISQ platforms for near-term EO utility, while monitoring *on-the-horizon* capabilities as roadmap inputs. Notably, the recent *verifiable quantum advantage* demonstrated by Google Quantum AI's 105-qubit *Willow* processor (Google Quantum AI, 2025) further highlights the growing maturity of hardware progress and reinforces the rationale for focusing on near-term hybrid experiments. Recent developments are also redefining how quantum, High Performance Computing (HPC), and AI systems can operate together. For instance, NVIDIA has recently introduced NVQLink, a technology that directly connects quantum processors with GPU-based supercomputers, enabling real-time data exchange between qubits and classical accelerators. Such advances open new possibilities for hybrid quantum–HPC–AI workflows. This combined approach is particularly relevant for EO pipelines, allowing future research to prototype realistic hybrid configurations and assess their scalability. These directions define a possible future goal: to bring quantum utility from simulation to hardware validation, establishing the foundations for the next generation of EO processing architectures.

## 6.5 Final Conclusions

This dissertation contributes to the field of QML4EO by providing a structured view of the landscape and by proposing and testing hybrid quantum–classical models spanning discriminative, relational, and generative learning.

The results support a balanced but optimistic conclusion: the path forward is not a competition between quantum and classical learning, but their *co-design*. By leveraging classical scalability for structure and quantum transformations for feature expressivity, it is possible to build systems that exceed the sum of their parts. As quantum hardware matures and algorithmic designs evolve, the methodologies outlined in this thesis, from quanvolutional feature extraction to quantum-conditioned diffusion, provide a concrete roadmap to evolve QML4EO from an exploratory field into a rigorous, operational component of the EO ecosystem.