

A Comprehensive Review of Hybrid Quantum-Classical Approaches for Satellite-Based Deforestation Detection

Prathamesh Balasaheb Kadam, Shreyas Subhash Raut, Tejas Uttam Shitole

Department of Artificial Intelligence and Data Science
AISSMS Institute of Information Technology, Pune

Abstract: Monitoring deforestation accurately and on a large scale is a vital environmental issue, increasingly dependent on analyzing large amounts of satellite images. Although traditional deep learning models like Convolutional Neural Networks (CNNs) have shown effectiveness, they are nearing their limits in terms of performance and computational capacity as data volumes continue to grow. Quantum Machine Learning (QML) offers a revolutionary alternative by utilizing quantum computing principles to identify complex patterns in high-dimensional feature spaces that classical methods cannot handle. This paper thoroughly reviews circuit-based hybrid quantum-classical models, which are suitable for current Noisy Intermediate-Scale Quantum (NISQ) devices, and assesses their use in classifying land use and land cover from satellite data. We examine and contrast two leading architectures: the Quantum Convolutional Neural Network (QCNN), which incorporates a quantum processing layer into a classical network, and the Neural Quantum Kernel (NQK) approach, which employs a quantum circuit to create a robust feature kernel for a classical Support Vector Machine (SVM). The review emphasizes that the effectiveness of these models is greatly affected by the incorporation of quantum entanglement in the circuit design, and they can achieve accuracy levels comparable to or better than the most advanced classical models with significantly reduced complexity. By synthesizing the current state of these emerging technologies, this review highlights the substantial potential of hybrid quantum-classical systems to enhance real-time environmental monitoring, while also identifying the main challenges and future research directions in this promising field.

Keywords: Remote sensing, EuroSAT, quantum machine learning, hybrid models, neural quantum kernels, deforestation detection, dimensionality reduction, quantum kernels

I. INTRODUCTION

The precise observation of global forest cover is vital for monitoring biodiversity, resources, and the global carbon cycle. Earth Observation missions deliver massive data volumes exceeding 150 terabytes daily, transforming land use and land cover (LULC) classification into a Big Data challenge. Classical machine learning (ML) and deep learning (DL) models like Convolutional Neural Networks (CNNs) have been leading tools for analyzing this data, achieving success in LULC classification. However, as satellite data grows, these methods face limitations, leading to exploration of new technologies. Quantum Computing (QC) has emerged as a breakthrough technology. Unlike classical computers using bits, quantum computers use qubits in superposition of states. This feature enables quantum systems to operate in vast Hilbert spaces, providing powerful computation capabilities. This has created Quantum Machine Learning (QML), offering new approaches to uncover patterns. Given current Noisy Intermediate-Scale Quantum (NISQ) limitations, hybrid quantum-classical models combine classical networks with quantum circuits for processing. This paper reviews hybrid quantum-classical models for satellite imagery classification, focusing on deforestation mapping. We analyze two architectures: the Quantum Convolutional Neural Network (QCNN) incorporating quantum processing, and the Neural Quantum Kernel (NQK) method for feature kernel generation. We examine their architectures, advantages, and potential to overcome classical limitations, addressing future research directions.



II. BACKGROUND

To understand the application of hybrid quantum-classical models to remote sensing, it is essential to first grasp the fundamental principles of quantum computing and the rationale behind the hybrid architectural paradigm. This section provides the necessary background on these core concepts.

A. Fundamentals of Quantum Computing for ML

At the heart of quantum computing is the qubit, or quantum bit, the fundamental unit of quantum information. Unlike a classical bit, which can only be in a state of 0 or 1, a qubit can exist in a superposition of both states simultaneously. This state, denoted $| \rangle$, is described as a linear combination of the basis states $| 0 \rangle$ and $| 1 \rangle$:

Here, α and β are complex numbers known as probability amplitudes, which are constrained by the rule $|\alpha|^2 + |\beta|^2 = 1$, where $|\alpha|^2$ and $|\beta|^2$ represent the respective probabilities of measuring the qubit in state $| 0 \rangle$ or $| 1 \rangle$. This probabilistic nature is a key departure from classical deterministic systems. Geometrically, the state of a single qubit can be visualized as a point on the surface of a 3D sphere known as the Bloch sphere.

Operations on qubits are performed by quantum gates, which are mathematically represented by unitary matrices and are the building blocks of quantum circuits. Key single-qubit gates include the Hadamard (H) gate, which creates an equal superposition from a basis state, and Rotation gates (e.g., $R_y(\theta)$), which rotate the qubit's state vector on the Bloch sphere by a specified angle θ . These rotation gates are critical for Parameterized Quantum Circuits (PQCs), as they provide a mechanism to encode classical data into the quantum state. Beyond single-qubit operations, multi-qubit gates enable one of quantum computing's most powerful features: entanglement. The Controlled-NOT (CNOT) gate, for instance, is a two-qubit gate that flips a target qubit if and only if a control qubit is in the state $| 1 \rangle$. When applied to qubits in superposition, this gate creates an entangled state, where the qubits become intrinsically linked. In such a state, the measurement outcome of one qubit instantly influences the state of the other, regardless of the distance separating them—a non-classical correlation that dramatically expands the computational space of the system.

B. Hybrid Quantum-Classical Architectures

Current quantum processors operate in what is known as the Noisy Intermediate-Scale Quantum (NISQ) era. This era is defined by quantum devices that have a limited number of qubits and are susceptible to environmental noise and gate errors, which corrupt the delicate quantum states before a computation is complete. Due to these technological limitations, building and running a large, fault-tolerant quantum algorithm for a real-world problem remains a significant challenge.

To harness the power of quantum mechanics in this near term, the field has converged on hybrid quantum-classical models. This paradigm leverages the respective strengths of both computational worlds. A classical computer, typically equipped with GPUs, handles tasks it excels at: loading and pre-processing large datasets, executing deep neural networks for efficient feature extraction, and optimizing model parameters. The quantum processor, or Quantum Processing Unit (QPU), is then used as a specialized co-processor to perform a specific, complex task where it may hold a computational advantage.

In a typical hybrid QML model, classical data is encoded into a PQC through a process known as data embedding. For instance, the feature vector extracted by a classical CNN can be used as the set of angles $\{\theta\}$ for the rotation gates within the PQC. The quantum circuit processes this information, and its final state is measured to produce a classical output vector. This output can then be fed into subsequent classical layers for final processing or used to compute a loss function. This entire hybrid system is often trained end-to-end, with gradients being calculated for both the classical and quantum parameters to iteratively optimize the model's performance.

C. Quantum Machine Learning Overview

Quantum circuits can encode classical vectors into quantum states using angle, amplitude, or structured entangling feature maps. Parameterized circuits enable trainable transformations but risk barren plateaus when depth scales [5]. Quantum kernel estimation computes inner products between embedded quantum states, feeding a classical kernel



method (e.g., SVM) [4]. Neural Quantum Kernels extend this by first training a PQC on classification signals, then freezing it to generate a data-driven kernel [8].

III. EXISTING SYSTEM

Building on these foundational concepts, researchers have developed several distinct hybrid architectures for image classification. This section reviews two of the most prominent approaches: the Quantum Convolutional Neural Network (QCNN) and the Neural Quantum Kernel (NQK) method.

A. Quantum Convolutional Neural Networks (QCNNs)

A QCNN is a hybrid architecture that enriches a classical CNN by introducing a quantum layer into its structure. A prime example of this model is the modified LeNet-5 architecture proposed for LULC classification, which features a "sandwich-like" design where a PQC is placed between classical fully connected layers. The workflow begins with classical convolutional layers that perform initial feature extraction from the input image. The resulting feature map is flattened and passed through a classical fully connected layer, which downsamples it into a lower-dimensional vector suitable for being encoded onto the quantum circuit.

The core of the QCNN is its quantum layer. The classical vector is embedded into a PQC, often with 4 to 8 qubits, by using its values as the parameters for a series of rotation gates. A key finding is that the circuit's design, particularly its use of entanglement, is crucial for the model's performance. In comparative studies, circuits that leveraged entanglement (such as the Bellman Circuit and Real Amplitudes Circuit) achieved significantly higher classification scores than simpler circuits with no entanglement. This suggests that the correlations enabled by entanglement allow the model to capture more complex patterns within the remote sensing data. After the quantum processing, the qubits are measured, yielding a classical output vector. This vector is then up-sampled by another classical fully connected layer to match the number of output classes before a final softmax activation.

To further enhance performance, this architecture can be deployed in a coarse-to-fine classification framework. In this strategy, a first QCNN is trained to categorize images into broad "macro-classes" (e.g., Vegetation, Urban, Water Bodies). Subsequently, specialized QCNNs, each trained on a specific subset of data, perform fine-grained classification within those macro-classes. This hierarchical approach proved highly effective, achieving a final accuracy of 97% on the 10-class EuroSAT dataset—a result on par with much deeper and more complex classical models like ResNet-50, but with a vastly simpler 6-layer architecture.

a) B. Neural Quantum Kernels (NQK) : An alternative and highly promising hybrid approach is the Neural Quantum Kernel (NQK) method. This architecture fundamentally changes the role of the quantum circuit. Instead of acting as a trainable classification layer within a network, the quantum circuit is used to engineer a powerful kernel function. In machine learning, a kernel is a function that measures the similarity between two data points in a high-dimensional feature space. The NQK method uses a quantum circuit to compute this similarity in the exponentially large Hilbert space, effectively creating a "quantum kernel".

The NQK workflow is distinct from the QCNN's end-to-end training. It typically involves a multi-stage process:

1) Classical Feature Extraction and Reduction: An input image is first passed through a powerful classical model, such as a pre-trained CNN or an autoencoder, to extract a meaningful, low-dimensional feature vector. This initial step is critical as it addresses the data embedding bottleneck, ensuring the vector is small enough to be efficiently encoded onto a NISQ-era quantum device.

2) Quantum Kernel Computation: A PQC is trained on labeled data to act as a feature map. Once trained, this circuit is "frozen" and used to estimate the similarity between pairs of input vectors by calculating their inner product in the quantum feature space. This process is repeated for all pairs of data points in the training set to construct a complete kernel matrix.

3) Classical Classification: Finally, this quantum-generated kernel matrix is fed into a powerful classical kernel-based classifier, most commonly a Support Vector Machine (SVM), which then learns the final decision boundary for classification.

This decoupled architecture offers several advantages for near-term applications. By relying on small and shallow quantum circuits, the NQK method is inherently more robust to the noise present in NISQ hardware. Furthermore, its modularity allows for flexibility, as different classical front-ends or SVM configurations can be tested without needing to retrain the quantum feature map. This approach has proven successful on real remote sensing tasks, demonstrating comparable or improved performance over classical methods while being highly efficient for current quantum systems

b) C. Dataset and Training Setup: The model was trained and tested on the EuroSAT dataset, which consists of 27,000 labeled Sentinel-2 images divided into 10 classes. Training used the Adam optimizer with a learning rate of 0.0002 and 50 epochs. Quantum simulations were performed using IBM Qiskit and Google Colab GPU backends.

c) D. Performance and Results:

- The Real Amplitudes QCNN achieved an overall accuracy of 92%, outperforming classical CNN baselines with equivalent complexity.
- The Coarse-to-Fine hybrid QCNN further improved accuracy to 97%, achieving results comparable to state-of-the-art ResNet-50 models but with far fewer layers (6 vs. 50).
- Reported F1-Scores ranged between 94% and 99% across fine-grained subclasses.
- The No Entanglement circuit performed the weakest (failing to classify the Highway class), demonstrating the significance of entanglement in learning complex spatial relationships.

d) E. Architectural Benefits:

1) Reduced Complexity:

The hybrid QCNN uses significantly fewer parameters than deep CNNs while maintaining high accuracy.

2) Quantum Efficiency:

Shallow and small quantum circuits are ideal for NISQ devices, making the model feasible for near-term implementation.

3) Enhanced Generalization:

Quantum entanglement improves feature separability and reduces overfitting.

4) Scalability and Modularity:

The Coarse-to-Fine strategy enables independent training of sub-classifiers, supporting scalable environmental applications such as deforestation monitoring.

IV. RESEARCH GAP AND MOTIVATION

Despite the strong performance of the hybrid QCNN presented in [4], several limitations restrict its applicability to deforestation mapping and real-world environmental monitoring. The model was validated only on the EuroSAT benchmark, which contains general-purpose land-cover categories such as vegetation, water bodies, and urban areas. While useful for demonstrating classification capability, this dataset does not focus on temporal forest-cover change, tree density variation, or biomass degradation, which are essential for accurate deforestation detection. Consequently, the existing architecture lacks sensitivity to subtle spectral differences that indicate progressive forest loss.

Furthermore, the quantum component in the base model primarily serves as a fixed transformation layer within a convolutional framework. It does not exploit quantum kernels or quantum similarity measures, which could enhance separability between complex vegetation patterns. The absence of Neural Quantum Kernels (NQKs) limits its capacity to learn nonlinear feature relationships in high-dimensional quantum feature spaces. Additionally, the QCNN relies on simulated quantum circuits with a small number of qubits (typically four), which restricts scalability and generalization to larger, real-world remote sensing datasets.

From a system design perspective, the coarse-to-fine classification structure improves accuracy but increases model redundancy and computational cost due to multiple specialized classifiers. For real-time deforestation monitoring, a more lightweight and explainable hybrid framework is required—one that balances accuracy, interpretability, and computational efficiency.

To address these gaps, this research proposes an NQK-based hybrid quantum-classical architecture for satellite-driven deforestation mapping. The motivation lies in utilizing classical preprocessing for dimensionality reduction and a quantum kernel module to compute pairwise similarities in the quantum Hilbert space. This allows for efficient Support



Vector Machine (SVM) classification with improved generalization, even in limited-data scenarios. By combining quantum feature encoding, kernel estimation, and explainable output metrics (confidence scores and class probabilities), the proposed system aims to deliver a practical, high-accuracy, and NISQ-compatible approach for environmental monitoring and policy-driven deforestation assessment.

A. Hybrid Quantum-Classical Architectures

Current quantum processors operate in what is known as the Noisy Intermediate-Scale Quantum (NISQ) era. This era is defined by quantum devices that have a limited number of qubits and are susceptible to environmental noise and gate errors, which corrupt the delicate quantum states before a computation is complete. Due to these technological limitations, building and running a large, fault-tolerant quantum algorithm for a real-world problem remains a significant challenge.

To harness the power of quantum mechanics in this near term, the field has converged on hybrid quantum-classical models. This paradigm leverages the respective strengths of both computational worlds. A classical computer, typically equipped with GPUs, handles tasks it excels at: loading and pre-processing large datasets, executing deep neural networks for efficient feature extraction, and optimizing model parameters. The quantum processor, or Quantum Processing Unit (QPU), is then used as a specialized co-processor to perform a specific, complex task where it may hold a computational advantage.

In a typical hybrid QML model, classical data is encoded into a PQC through a process known as data embedding. For instance, the feature vector extracted by a classical CNN can be used as the set of angles () for the rotation gates within the PQC. The quantum circuit processes this information, and its final state is measured to produce a classical output vector. This output can then be fed into subsequent classical layers for final processing or used to compute a loss function. This entire hybrid system is often trained end-to-end, with gradients being calculated for both the classical and quantum parameters to iteratively optimize the model's performance.

B. Quantum Machine Learning Overview

Quantum circuits can encode classical vectors into quantum states using angle, amplitude, or structured entangling feature maps. Parameterized circuits enable trainable transformations but risk barren plateaus when depth scales [5]. Quantum kernel estimation computes inner products between embedded quantum states, feeding a classical kernel method (e.g., SVM) [4]. Neural Quantum Kernels extend this by first training a PQC on classification signals, then freezing it to generate a data-driven kernel [8].

V. EXISTING SYSTEM

Building on these foundational concepts, researchers have developed several distinct hybrid architectures for image classification. This section reviews two of the most prominent approaches: the Quantum Convolutional Neural Network (QCNN) and the Neural Quantum Kernel (NQK) method.

A. Quantum Convolutional Neural Networks (QCNNs)

A QCNN is a hybrid architecture that enriches a classical CNN by introducing a quantum layer into its structure. A prime example of this model is the modified LeNet-5 architecture proposed for LULC classification, which features a "sandwich-like" design where a PQC is placed between classical fully connected layers. The workflow begins with classical convolutional layers that perform initial feature extraction from the input image. The resulting feature map is flattened and passed through a classical fully connected layer, which downsamples it into a lower-dimensional vector suitable for being encoded onto the quantum circuit.

The core of the QCNN is its quantum layer. The classical vector is embedded into a PQC, often with 4 to 8 qubits, by using its values as the parameters for a series of rotation gates. A key finding is that the circuit's design, particularly its use of entanglement, is crucial for the model's performance. In comparative studies, circuits that leveraged entanglement (such as the Bellman Circuit and Real Amplitudes Circuit) achieved significantly higher classification scores than simpler circuits with no entanglement. This suggests that the correlations enabled by entanglement allow



the model to capture more complex patterns within the remote sensing data. After the quantum processing, the qubits are measured, yielding a classical output vector. This vector is then up-sampled by another classical fully connected layer to match the number of output classes before a final softmax activation.

To further enhance performance, this architecture can be deployed in a coarse-to-fine classification framework. In this strategy, a first QCNN is trained to categorize images into broad "macro-classes" (e.g., Vegetation, Urban, Water Bodies). Subsequently, specialized QCNNs, each trained on a specific subset of data, perform fine-grained classification within those macro-classes. This hierarchical approach proved highly effective, achieving a final accuracy of 97% on the 10-class EuroSAT dataset—a result on par with much deeper and more complex classical models like ResNet-50, but with a vastly simpler 6-layer architecture.

a) B. Neural Quantum Kernels (NQK) : An alternative and highly promising hybrid approach is the Neural Quantum Kernel (NQK) method. This architecture fundamentally changes the role of the quantum circuit. Instead of acting as a trainable classification layer within a network, the quantum circuit is used to engineer a powerful kernel function. In machine learning, a kernel is a function that measures the similarity between two data points in a high-dimensional feature space. The NQK method uses a quantum circuit to compute this similarity in the exponentially large Hilbert space, effectively creating a "quantum kernel".

The NQK workflow is distinct from the QCNN's end-to-end training. It typically involves a multi-stage process:

1) Classical Feature Extraction and Reduction: An input image is first passed through a powerful classical model, such as a pre-trained CNN or an autoencoder, to extract a meaningful, low-dimensional feature vector. This initial step is critical as it addresses the data embedding bottleneck, ensuring the vector is small enough to be efficiently encoded onto a NISQ-era quantum device.

2) Quantum Kernel Computation: A PQC is trained on labeled data to act as a feature map. Once trained, this circuit is "frozen" and used to estimate the similarity between pairs of input vectors by calculating their inner product in the quantum feature space. This process is repeated for all pairs of data points in the training set to construct a complete kernel matrix.

3) Classical Classification: Finally, this quantum-generated kernel matrix is fed into a powerful classical kernel-based classifier, most commonly a Support Vector Machine (SVM), which then learns the final decision boundary for classification.

This decoupled architecture offers several advantages for near-term applications. By relying on small and shallow quantum circuits, the NQK method is inherently more robust to the noise present in NISQ hardware. Furthermore, its modularity allows for flexibility, as different classical front-ends or SVM configurations can be tested without needing to retrain the quantum feature map. This approach has proven successful on real remote sensing tasks, demonstrating comparable or improved performance over classical methods while being highly efficient for current quantum systems

b) C. Dataset and Training Setup: The model was trained and tested on the EuroSAT dataset, which consists of 27,000 labeled Sentinel-2 images divided into 10 classes. Training used the Adam optimizer with a learning rate of 0.0002 and 50 epochs. Quantum simulations were performed using IBM Qiskit and Google Colab GPU backends.

c) D. Performance and Results:

- The Real Amplitudes QCNN achieved an overall accuracy of 92%, outperforming classical CNN baselines with equivalent complexity.
- The Coarse-to-Fine hybrid QCNN further improved accuracy to 97%, achieving results comparable to state-of-the-art ResNet-50 models but with far fewer layers (6 vs. 50).
- Reported F1-Scores ranged between 94% and 99% across fine-grained subclasses.
- The No Entanglement circuit performed the weakest (failing to classify the Highway class), demonstrating the significance of entanglement in learning complex spatial relationships.

d) E. Architectural Benefits:

1) Reduced Complexity:

The hybrid QCNN uses significantly fewer parameters than deep CNNs while maintaining high accuracy.

2) Quantum Efficiency:

Shallow and small quantum circuits are ideal for NISQ devices, making the model feasible for near-term implementation.

3) Enhanced Generalization:

Quantum entanglement improves feature separability and reduces overfitting.

4) Scalability and Modularity:

The Coarse-to-Fine strategy enables independent training of sub-classifiers, supporting scalable environmental applications such as deforestation monitoring.

VI. RESEARCH GAP AND MOTIVATION

Despite the strong performance of the hybrid QCNN presented in [4], several limitations restrict its applicability to deforestation mapping and real-world environmental monitoring. The model was validated only on the EuroSAT benchmark, which contains general-purpose land-cover categories such as vegetation, water bodies, and urban areas. While useful for demonstrating classification capability, this dataset does not focus on temporal forest-cover change, tree density variation, or biomass degradation, which are essential for accurate deforestation detection. Consequently, the existing architecture lacks sensitivity to subtle spectral differences that indicate progressive forest loss.

Furthermore, the quantum component in the base model primarily serves as a fixed transformation layer within a convolutional framework. It does not exploit quantum kernels or quantum similarity measures, which could enhance separability between complex vegetation patterns. The absence of Neural Quantum Kernels (NQKs) limits its capacity to learn nonlinear feature relationships in high-dimensional quantum feature spaces. Additionally, the QCNN relies on simulated quantum circuits with a small number of qubits (typically four), which restricts scalability and generalization to larger, real-world remote sensing datasets.

From a system design perspective, the coarse-to-fine classification structure improves accuracy but increases model redundancy and computational cost due to multiple specialized classifiers. For real-time deforestation monitoring, a more lightweight and explainable hybrid framework is required—one that balances accuracy, interpretability, and computational efficiency.

To address these gaps, this research proposes an NQK-based hybrid quantum-classical architecture for satellite-driven deforestation mapping. The motivation lies in utilizing classical preprocessing for dimensionality reduction and a quantum kernel module to compute pairwise similarities in the quantum Hilbert space. This allows for efficient Support Vector Machine (SVM) classification with improved generalization, even in limited-data scenarios. By combining quantum feature encoding, kernel estimation, and explainable output metrics (confidence scores and class probabilities), the proposed system aims to deliver a practical, high-accuracy, and NISQ-compatible approach for environmental monitoring and policy-driven deforestation assessment.

VII. COMPARATIVE STUDY

Recently, machine learning and deep learning models have been utilized for deforestation monitoring. A comparative analysis of already available models, ResNet, CNN, QCNN, and NQK methods, provides useful information:

a) Complexity: CNN: Moderately complex computationally; can be utilized for typical image classification use cases but can become computationally expensive on larger resolution satellite images.

ResNet: More complicated due to residual blocks but it enables deep architectures, which helps mitigate the issues of vanishing gradients.

QCNN (Quantum Convolutional Neural Networks): Incorporates quantum layers; complexity is dependent on the number of qubits and the number of gate operations; could be lower for datasets with few images, but continues to grow non-linearly as the number of qubits increases.

NQK based models: New quantum kernel-based models can have high theoretical complexity but efficiently encode richer correlations in the data.

b) Performance: ResNet and CNN perform comparably on accuracy metrics on classic models.

QCNN models have demonstrated promising results on increasingly non-linear patterns with fewer parameters.



NQK based approaches can out-perform classical models on smaller datasets due to quantum encoding benefits, especially in fine-grained land-cover classification.

c) Scalability: Classical models (CNN/ResNet) scale better with existing GPU infrastructure.

QCNN and NQK models face challenges scaling with number of qubits to make them difficult to deploy in large-scale applications.

d) Hardware Feasibility: Classical models are compatible with hardware, deployable in GPUs and the cloud. In contrast, quantum models are limited by the hardware of NISQ (Noisy Intermediate-Scale Quantum), which includes limitations such as qubit decoherence, poor connectivity, and error rates.

VIII. CHALLENGES AND MITIGATION (CONCEPTUAL)

- Qubit Limitations: Cut dimensions aggressively; consider features selected by variance or mutual information.
- Noise / Decoherence: Use shallow circuits (little entangling depth); implement some form of error-correcting protocol or readout calibration.
- Barren Plateaus: Use shallow depth, initialize parameters with small angles, and assemble depth layerwise with small angles.
- Kernel Scaling: Work with (the Nystro'm approximation, or do some kind of active selection of a subset of the data).
- Overfitting: Use cross-validation, early stopping, set constraints on circuit expressivity (keep a limited number of entanglers).
- Explainability: Perturb the reduced features, and analyze sensitivity to kernel similarities.

IX. OPEN RESEARCH QUESTIONS

Key questions are:

- What conditions will result in learned quantum kernels outperforming classical RBF or polynomial kernels at the same latent dimension?
- What is the best topology for a circuit to balance feasible expressivity and trainability with NISQ noise?
- What can we automate by jointly selecting a classical encoder and quantum map?
- How can domain priors (to implement spectral index gates) be injected?
- How can we build scalable approximate kernels?

X. ANTICIPATED (NON-EMPIRICAL) OUTCOMES

Anticipated conclusions are: improved F1 score compared to a classical SVM on the same reduced features, and/or smaller parameter footprint in regard to deeper CNNs at similar accuracy, and/or the ability to reuse a learned quantum map as a module. There are no empirical conclusions provided.



XI. ARCHITECTURE DIAGRAM

Neural Quantum Kernel (NQK) Architecture

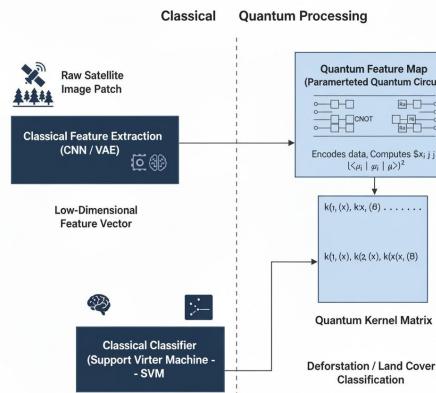


Fig. 1: System Architecture of the Proposed Model

XII. ETHICAL AND SUSTAINABILITY CONCERNS

Improving deforestation detection leads to improved environmental governance. Hybrid models may be less energy demanding compared to retraining deep models at scale. Confidence scoring and uncertainty need to be accounted for to avoid bad policy decisions. The use of the data must respect data licensing and ensure sensitive geographic data is not exploited.

XIII. CONCLUDING THOUGHTS

This overview outlines the motivation behind hybrid quantum-classical methods for deforestation monitoring within the constraints of NISQ technology, that is a Neural Quantum Kernel fit framework that offers the coupling of classical compressed data with a learned and reusable quantum feature map. The next steps are to measure the actual overhead costs to see if quantum-enhanced kernelization is defensibly superior to a classical baseline in a controlled dimensional setting.

ACKNOWLEDGEMENT

The authors recognize the mentorship received from Dr. Prajwal S. Gaikwad.

REFERENCES

- [1] 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 J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 12, no. 7, pp. 2217–2226, 2019. doi: 10.1109/JSTARS.2019.2918242.
- [2] 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 J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 15, pp. 565–580, 2022. doi: 10.1109/JSTARS.2021.3134785.
- [3] M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers*. Springer, 2018. doi: 10.1007/978-3-319-96424-9.
- [4] V. Havlíček, A. D. Corcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantumenthanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019. doi: 10.1038/s41586-019-0980-2.
- [5] M. Cerezo, A. Arrasmith, R. Babbush, et al., "Variational Quantum Algorithms," *Nature Reviews Physics*, vol. 3, pp. 625–644, 2021. doi: 10.1038/s42254-021-00348-9.



- [6] E. Farhi and H. Neven, "Classification with Quantum Neural Networks on Near Term Processors," arXiv:1802.06002, 2018.
- [7] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, "Parameterized quantum circuits as machine learning models," *Quantum Science and Technology*, vol. 4, p. 043001, 2019. doi: 10.1088/2058-9565/ab4eb5.
- [8] M. Schuld, "Quantum Machine Learning Models are Kernel Methods," arXiv:2101.11020, 2021.
- [9] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, "Quantum Embeddings for Machine Learning," arXiv:2001.03622, 2020.
- [10] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, "Connecting Ansatz Expressibility to Gradient Magnitudes and Barren Plateaus," *PRX Quantum*, vol. 3, p. 010313, 2022. doi: 10.1103/PRXQuantum.3.010313.
- [11] A. Mari, T. R. Bromley, J. Izaac, and N. Killoran, "Estimating the Gradient and Higher-Order Derivatives on Quantum Hardware," *Phys. Rev. A*, vol. 105, p. 012603, 2022. doi: 10.1103/PhysRevA.105.012603.
- [12] H. Abraham, et al., "Qiskit: An Open-source Framework for Quantum Computing," 2019. Available: <https://qiskit.org/>
- [13] V. Bergholm, J. Izaac, et al., "PennyLane: Automatic Differentiation of Hybrid Quantum-Classical Computations," arXiv:1811.04968, 2018.
- [14] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," arXiv:1412.6980, 2015.
- [15] N. Simidjievski, C. Bodnar, et al., "Variational Autoencoders for Cancer Data Integration," *Bioinformatics*, vol. 35, no. i, pp. i51–i60, 2019. doi: 10.1093/bioinformatics/btz341.
- [16] S. Wold, K. Esbensen, and P. Geladi, "Principal Component Analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 2, pp. 37–52, 1987. doi: 10.1016/0169-7439(87)80084-9.
- [17] T. M. Cover and P. E. Hart, "Nearest Neighbor Pattern Classification," *IEEE Trans. Inf. Theory*, vol. 13, no. 1, pp. 21–27, 1967. doi: 10.1109/TIT.1967.1053964.
- [18] M. C. Hansen, et al., "High-Resolution Global Maps of 21st-Century Forest Cover Change," *Science*, vol. 342, no. 6160, pp. 850–853, 2013. doi: 10.1126/science.1244693.
- [19] A. Dosovitskiy, et al., "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale," *ICLR*, 2021.
- [20] C. K. I. Williams and M. Seeger, "Using the Nyström Method to Speed Up Kernel Machines," *Advances in Neural Information Processing Systems*, vol. 13, pp. 682–688, 2001