HybridVQC: A Quantum-Inspired Neural Architecture for Autism Spectrum Disorder Classification

Shrivatsa Vaman Gudi[1*], Parimal Nitin Ganorkar[1], and Sasikala R[1]

[1] School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, India

Abstract. The classification of Autism Spectrum Disorder (ASD) using neuroimaging techniques is still not an easy task, as the features have a high dimension, inter-site variability, and a small number of labeled samples. In this paper, HybridVQC, a quantum-inspired hybrid neural architecture, incorporates mathematical concepts of quantum circuits into a fully classical, learning architecture that can be executed on a GPU. ABIDE Structural MRI slices are fed through a pretrained backbone of ResNet-18 to obtain deep representations, and the deep representations are further shrunk to 16 principal components in Principal Component Analysis (PCA). The low features are subsequently applied to a special QuantumLikeLayer, which uses trigonometric encoding and dense mixing of features in order to simulate quantum rotation and entanglement effects on ordinary CUDA. The results of experiments with 1,693 structural MRI slices reveal that the maximum validation and test accuracy is 80.63 and 75.0, respectively, versus 56% of a classical SVM baseline using the same parameters. The findings point to quantum-inspired non-linear transformations have the potential to enhance feature separability and training stability in neuroimaging classification, and do not use quantum simulators or physical quantum systems.

KEYWORDS: Quantum-Inspired Neural Network, HybridVQC, Neuroimaging Classification, ABIDE Dataset, Trigonometric Encoding, QuantumLikeLayer, Medical Image Analysis, Classical-Quantum Hybrid Learning

[*]Corresponding author: shrivatsagudi@gmail.com

INTRODUCTION:

Autism Spectrum Disorder (ASD) is a complicated neurodevelopmental disorder that impacts the communicational ability, sensory information processing, and behavior of people. The early and proper diagnosis of it is of paramount importance to increase the long-term effects. But the present practices largely depend on behavioural judgments. This has resulted in the application of neuroimaging particularly functional and structural Magnetic Resonance Imaging (fMRI and sMRI) as a more objective instrument to determine brain-based autism markers. Nonetheless, MRI data is neither simple nor noisy nor the same across various subjects, making it a major problem since automated classification mechanisms attempt to detect meaningful patterns in such data.

The general machine learning and deep learning approaches have been extensively researched in neuroimaging classification. It was demonstrated that such models as Support Vector Machines (SVMs), Multilayer Perceptrons (MLPs), and convolutional neural networks (CNNs) as ResNet-18 are capable of taking into account the spatial dependency, as well as high-level features. In this paper, we have taken ResNet-18 as a feature extractor in order to extract powerful representations of MRI slices. Principal Component Analysis (PCA) was then used to reduce the dimensions. Although these classical models prove to be effective, they tend to overlook intricate connections among features and higher-order neural data correlations. Simultaneously, emerging methods such as Quantum Machine Learning (QML) have demonstrated the possibility of modeling these complicated connections more effectively. Nevertheless, actual quantum hardware, in particular, modern Noisy Intermediate-Scale Quantum (NISQ) computers is constrained by qubits, decoherence, and noise. This renders them inapplicable to big datasets such as ABIDE.

To bridge this gap, we introduce a Quantum-Inspired Hybrid Neural Network (HybridVQC) to be a hybrid between the mathematics of quantum calculation and a fully classical, and GPU-accelerated, model. This design has an original QuantumLikeLayer that is to simulate quantum behaviours such as superposition, rotation, and entanglement by performing trigonometric transformations and dense mixing of features. This approach, in contrast to the real quantum models, involves quantum-inspired transformations on CUDA-enabled hardware. It is scalable, reproducible, and at the same time has the rich representation of quantum systems. Our model was trained on ABIDE dataset, with preprocessing of ResNet-18 and PCA. The findings indicated encouraging accuracy and stability, which indicates the possibilities of quantum-inspired learning in the application to real-world neuroimaging.

LITERATURE REVIEW:

Diagnosis of autism spectrum disorder (ASD) by neuroimaging is a field of study under investigation. The goal is to develop objective mechanisms of categorizing by biomarkers. It started with early initiatives, such as the Autism Brain Imaging Data Exchange (ABIDE) [1], which established a big pool of MRI and fMRI data. This enabled researchers to perform consistent studies on the brain connectivity in autism. It has since developed numerous machine learning and deep learning techniques to extract relevant neural signatures in this information.

Recent deep learning research investigated various possibilities of detecting ASD with MRI and fMRI. Convolutional neural networks (CNNs) in their traditional form have been used to learn significant brain characteristics using imaging data alone. As an example, the functional connectivity matrices have been analyzed using one-dimensional CNN models to detect patterns of features over time and enhance the accuracy of early diagnosis [2]. Similarly, CNN architectures that are optimized including convolutional, pooling, dropout, etc. have been used to attain impressive accuracy in detecting ASD in children [3]. Besides the regular CNNs, graph-based networks such as ASD-GraphNet [4] use brain regions as nodes and the links between them as edges. This has given it improved interpretability and classification on ABIDE. These experiments show the effectiveness of deep feature extraction, although they also indicate the problems of imbalanced data, inter-subject variation, and constraints of traditional architectures to simulate more intricate neural interactions.

Along with the development of deep learning, quantum machine learning (QML) has also emerged as providing new methods to represent high-dimensional data using transformations of quantum mechanics. More attention has been given to hybrid structures that combine classical neural networks with quantum elements. It has been demonstrated that studies on Quantum Support Vector Machines (QSVM) and Variational Quantum Circuits (VQC) can better attempt to describe non-linear decision boundaries in medical datasets [5]-[7]. As an example, the detection of ASD and Alzheimer disease by the use of QSVM-based classifiers has been used, utilizing concepts such as quantum superposition and entanglement to enhance feature separation [6], [8]. But the majority of these studies rely on simulated quantum systems or small datasets due to the current constraints of noisy intermediate-scale quantum (NISQ) hardware, which makes their generalization to large neuroimaging studies such as ABIDE possible.

To address these hardware limits, an alternative set of quantum-inspired algorithms has developed. The objectives of these directions are to simulate the mathematical strength of quantum computing based solely on classical models. They commonly make use of trigonometric transformations, feature rotations that are complex, and entanglements like operations implemented on GPUs. Surveys and research have revealed that the area of research has developed at a very high rate in the recent past. They emphasize the flexibility of quantum-inspired metaheuristics and learning algorithms to deal with complex optimization and classification issues [9]. Together with conventional neural networks, quantum-inspired deep learning systems have already shown the capability of re-creating the representational richness of quantum systems with no reliance on quantum processors [10], [11].

The present study is anchored on this emerging literature. Past studies of deep learning in relation to the ABIDE have established firm guidelines on ASD classification. Quantum representations have meanwhile demonstrated the potential benefits of quantum representations as demonstrated by quantum and hybrid models. Nevertheless, there is a lack of literature investigating a GPU-based, quantum-inspired neural network capable of integrating quantum-like mapping into a practicable deep learning framework on large neuroimaging datasets. The proposed HybridVQC model bridges this gap by withholding trigonometric encoding and dense feature-mixing operations to simulate the behaviour of quantum entanglement and rotation in an entirely classical system. Based on the results of previous CNN, graph-learning, and hybrid quantum models, the paper moves the neuroimaging classification forward and demonstrates that quantum-inspired computation is practical and can be trained on standard, non-specialised GPU hardware.

METHODOLOGY:

1. Overview of the Proposed Framework

The suggested framework, HybridVQC, presents a new hybrid neural architecture that is quantum-inspired and that fills the gap present between Deep Learning (DL) and Quantum Machine Learning (QML). The system is created with the purpose of classifying Autism Spectrum Disorder (ASD) on the ABIDE dataset and applies a Classical-to-Quantum-Inspired paradigm of transfer learning.

The architectural process, as shown in Fig. 1, handles the high dimensionality of medical imaging data with a piping process that decomposes the process into four sequential phases:

1.Deep Feature Extraction: use a results of ResNet-18 backbone to transform raw MRI pixels to a high level semantic feature space.

2.Manifold Reduction: The use of Principal Component Analysis (PCA) to map the sparse feature vectors of 512 dimensions onto a dense 16-dimensional subspace that can be efficiently simulated on a quantum computer.

3.Quantum-Inspired Processing: Ingesting the compressed embeddings in a special Quantum-Inspired Neural Network (QINN). It is an essential component, the QuantumLikeLayer, which is a mathematical simulation of qubit superposition (through phase encoding) and entanglement (through dense feature mixing) on classical hardware.

4.Classification & Evaluation: Deriving diagnostic probabilities via an ultimate soft-margin classifier.

A pretrained ResNet-18 network converts each slice of the MRI to a small, data-rich embedding. The embeddings obtained as 512 dimensional are condensed into 16 major components and are fed into the proposed Quantum-Inspired Neural Network (QINN), which has the custom QuantumLikeLayer. This layer mathematically models quantum circuit behavior (rotational and entanglement) of quantum circuits. It performs trigonometric transformations that are all on the classical CUDA-enabled GPUs.

More importantly, the architecture uses hierarchical abstraction strategy in managing the complexity of neuroimaging information. The classical ResNet-18 backbone is in a good position to learn local spatial hierarchies of edges, texture and shape in the cortical structures but is prone to fail to learn non-local interactions over small datasets because of the large parameter space of fully connected layers. By transferring these features extracted to the Quantum-Inspired Neural Network (QINN) the system becomes no longer about perceptual processing (seeing the picture) but relational inference (understanding the patterns). The entanglement-like nature of the QINN is particularly tailored to finding non-linear associations between the compressed principal components and more effectively serves as a very efficient, dense associative memory.

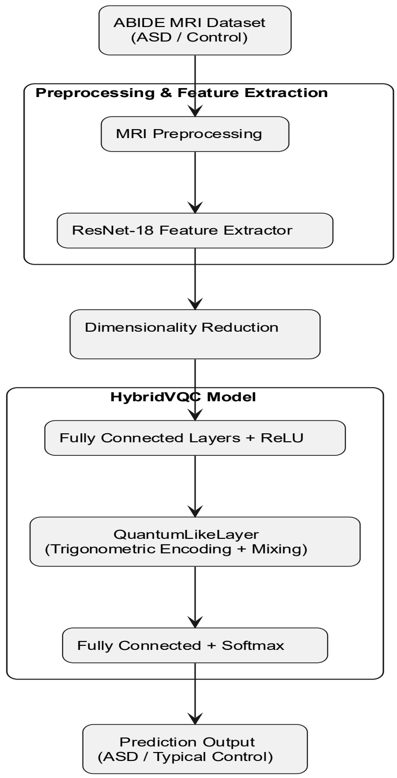

Figure 1: **Architecture of the proposed *HybridVQC* framework.**

2. Data Preprocessing and Feature Extraction

ABIDE dataset is a multi-site multi-site scanner-based MRI, which by nature causes variability in scanner types, voxel resolution and signal intensities. In order to deal with this heterogeneity and enable the data to go to the hybrid quantum architecture, there was a strict three-stage preprocessing pipeline.

2.1. Slice Selection and Normalization:

This paper uses a slice-based method (so-called 2D) instead of volumetric (3D) processing, which is characterized by high computational cost. For every subject volume$V \in R^{D\times H\times W}$, we algorithmically identified the geometric center slice Z_{mid} = D // 2. To capture sufficient anatomical context, we extracted a triplet of axial slices S ={Z_{mid}-1, Z_{mid}, Z_{mid}+1}. This specific spatial window ensures the capture of the Corpus Callosum and Lateral Ventricles, structural biomarkers which have been statistically shown to exhibit volumetric deviations in ASD patients.

After extraction, intensity inhomogeneity correction was done. The pixel intensities of MRI are arbitrary relative units so we did normalization of Minmax per slice to bring intensity value within the range of [0, 1]. Given a 2D slice I the normalized pixel value is given by:

$$I_{norm} = \frac{I - I_{min}}{I_{max} - I_{min} + \epsilon} \quad (1)$$

where I_{min} and I_{max} are the minimum and maximum intensity values of the slice, and $\epsilon = 10^{-8}$ is a small constant added for numerical stability. This step prevents gradient explosion during the initial training phase.

2.2. Deep Feature Extraction (Transfer Learning):

Since the ABIDE data is a small dataset in comparison to natural image datasets, inductive training with a deep Convolutional Neural Network (CNN) is highly susceptible to overfitting. As a solution to this, we used Transfer Learning with the ResNet-18 architecture that had been pretrained on ImageNet dataset (~1.2 million images).

Architecture Modification: The last Fully Connected (FC) classification layer of ResNet-18 has been taken out revealing the global average pooling layer.

Input Adaptation: The input of one channel of the single channel grayscale MRI slices was copied to three channels to fit the RGB input specification (3 x 224 x 224) of the pretrained network.

Forward Pass: Each preprocessed slice was passed through the frozen backbone to yield a high-level feature vector $f \in R^{512}$ This vector, denoted as X_{512}, encapsulates abstract spatial hierarchies (edges, textures, shapes) learned from massive datasets, providing a robust starting point for the subsequent quantum-inspired classification task.

3. Dimensionality Reduction Using PCA

The dimension of the feature vectors available as a result of the ResNet-18 backbone is high (512 features). Direct processing of this would incur too much computational overhead and could easily cause overfitting to smaller medical datasets. In order to deal with this we use Principal Component Analysis (PCA).

It was empirically determined that the number of principal components is 16 because with this number the cumulative variance in the extracted feature space was roughly 95%. Adding more components after 16 led to less significant improvements in accuracy but significantly increased the computational complexity of the quantum-inspired layer. Therefore, 16 elements offered the best trade-off between representational ability, generalization accuracy and computational efficiency.

4. Quantum-Inspired Hybrid Neural Network (HybridVQC_CUDA)

4.1. Architectural Composition

The HybridVQC network consists of three functional blocks (see Fig. 1):

- Input Block (fc_{in}) – two fully connected layers projecting the 16-D input into a latent subspace, followed by ReLU activations.
- QuantumLikeLayer – a custom transformation layer that mimics quantum circuit operations through trigonometric encoding and dense feature mixing.
- Output Block (fc_{out}) – a classifier head producing two logits corresponding to ASD and Control categories.

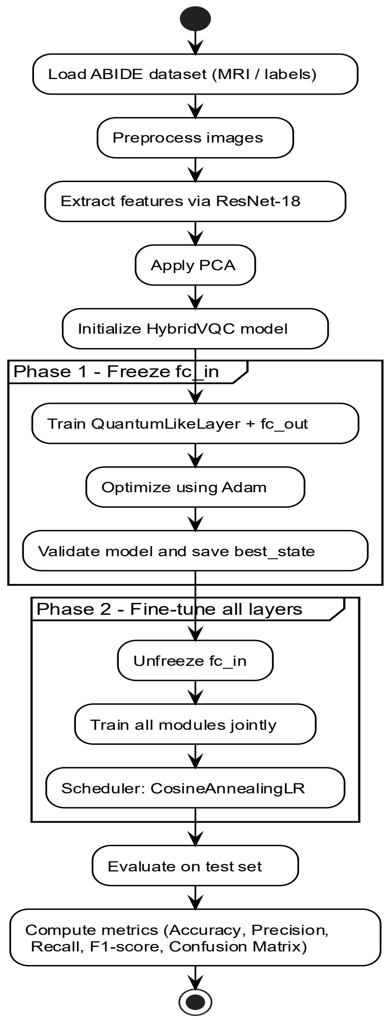

Figure 2: **Training and processing flow of the *HybridVQC* model**

4.2. Quantum-Inspired Encoding

At the heart of the model lies the QuantumLikeLayer, which performs three concurrent operations resembling quantum gate rotations.
For each feature vector $x \in \mathbb{R}^F$, with learnable parameters $\theta_{\ell,q,1}, \theta_{\ell,q,2}, \theta_{\ell,q,3}$ at layer ℓ and qubit q, the encoding is defined as:

$$h = \sin(x + \theta_1) + \cos(x + \theta_2) + \sin(x \times \theta_3) \qquad (2)$$

Equation (2) introduces phase-like rotations using sine and cosine functions, where each term corresponds to a transformation analogous to quantum rotation gates (R_x, R_y, R_z).

- The sine term models constructive interference,

- the cosine term captures complementary phase information, and
- the multiplicative term $x \times \theta_3$ provides amplitude modulation, enabling non-linear feature expansion.

This combination allows the network to learn richer internal representations than traditional activation functions.

4.3. Entanglement-based Feature Mixing

To imitate quantum entanglement, where qubits share correlated states, a dense mixing operation is applied across the simulated qubit dimensions:

$$h' = \frac{1}{Q}\ h\ 1_{Q\times Q} \qquad (3)$$

Here, $Q = 6$ represents the number of simulated qubits, and $1_{Q\times Q}$ is an all-ones matrix. Equation (3) enforces global dependency among features by averaging and redistributing information across all qubit channels.
This operation means that every output dimension of the layer is contingent on all the input dimensions, which recreates the entire coupling of entangled quantum states, but which is computationally efficient on GPUs.

The encoded representation h' is then forwarded to the classifier head fc_{out}, which performs linear projection and applies a softmax function to output final class probabilities.

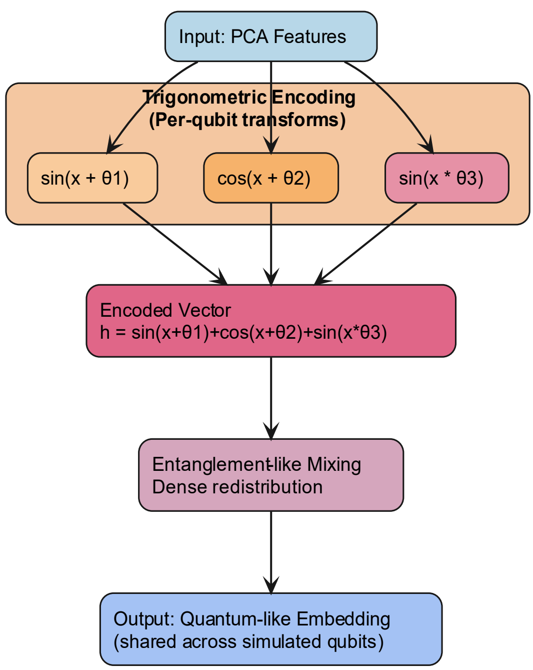

Figure 3: **Internal operations of the *QuantumLikeLayer* (encoding and).**

5. Training Strategy

Training of HybridVQC follows a two-phase optimization schedule, depicted in Fig. 2.

5.1. Phase 1 – Feature Freezing:
The parameters of the input block fc_{in} are frozen, while the QuantumLikeLayer and fc_{out} are trained with a learning rate of 3×10^{-3}. This enables the quantum-inspired layer to balance itself and acquire its internal changes on its own.

5.2. Phase 2 – The network is fine-tuned end-to-end and all layers are unfrozen using a smaller learning rate of 1×10^{-3}. The Adam algorithm opts with Cross-Entropy Loss and label smoothing (0.05) are used to counter over-confidence. CosineAnnealingLR scheduler is a dynamically decreasing learning rate schedule that runs for 60 epochs.

Optimization uses the Adam algorithm with Cross-Entropy Loss and label smoothing (0.05) to prevent over-confidence.

A CosineAnnealingLR scheduler dynamically reduces the learning rate over 60 epochs. The batch size is 64, and the gradients are capped by 1.0, to ensure stability.

6. Evaluation Metrics

The trained HybridVQC model was evaluated with the help of a set of metrics that were chosen considering the conditions of automated medical screening. The misclassification cost is asymmetric in a diagnostic setting, so we were more concerned with measures of clinical reliability than with statistical correctness.

Sensitivity-Driven Assessment (Recall): Recall was our major concern. False Negative (where an ASD subject is diagnosed as neurotypical) is the worst error in the context of Autism Spectrum Disorder (ASD), where early intervention therapies are crucial in the early stages. As a result, the model analysis focused on the positive case detection rate, whereby the quantum-inspired layer was to be able to identify the subtle structural biomarkers in the MRI images needed to identify possible cases.

Precision and Diagnostic Confidence: We also used Precision which was used to regulate the False Positive. Although high level of sensitivity is essential, low level of precision would yield too many false alarms hence unnecessary anxiety amongst patients and strain on the clinicians resources. The F1-Score was the harmonic balance between these two antagonistic goals that demonstrated that the model did not acquire high accuracy by merely predicting the majority class.

Generalization Verification: In order to ensure that the model was learning useful features as opposed to memorizing patient-specific noise (overfitting), we applied a strict control on the difference between Training and Validation measures. The last run was measured on a hold-out Test Set- data of subjects with whom the model had not been presented in the optimization stage.

Validation Strategy: A strict hold-out strategy was used to test the generalization ability of the model. The data was divided into training, validation and testing data to ensure that the reported measures are the performance of the model when applied to unfamiliar patients, and not the memorization of the patterns of features during the training process.

EXPERIMENTAL SETUP

1. Dataset Description

The suggested model was tested against Autism Brain Imaging Data Exchange (ABIDE) dataset [1]. This data is a open-source neuroimaging repository, which contains resting-state functional MRI (rs-fMRI) and structural MRI data recorded at different research locations across the world. It has subjects with Autism Spectrum Disorder (ASD) and Typical Control (TC) in a broader age range that is 7-64 years. The data is also associated with corresponding data like IQ, gender and diagnostic scores.

In this paper, a subset of structural MRI data was utilized in order to keep the imaging type and preprocessing procedures consistent. On quality filtering and on slice extraction, 1,693 MRI slices were incorporated. This sum is comprised of 860 slices with Typical Controls and 833 slices with ASD. Three central axial slices (center +/-1) were used by each subject. This made structural features, which were adjacent to the corpus callosum and thalamic areas well represented in the data.

This slice-based implementation achieves a dimensional reduction (over 90% of the raw data) of full 3D volumetric processing of data, which can then be trained on standard GPU hardware without losing diagnostically useful information.

In order to deal with class imbalance and provide a fair model training we resampled minority samples with the help of class-balanced loader. A final training set of 1,600 balanced slices was obtained in this process. The rest of the 20 percent (approximately 340 samples) was reserved to be validated and tested. This even distribution enabled the model to learn equally between the two categories of diagnostic and this prevented bias, whereas enhanced generalization performance.

The size of each MRI slice was resized to 224 x 224 pixels and normalized within the range [0, 1] and placed in a two-class folder structure; Autistic and typical Control. This configuration allowed loading of data in a reproducible way in training. The data was arranged according to the initial subject labelling convention of ABIDE [1] so that traceability could be made between imaging samples and metadata.

2. Environment and Implementation Details

Experiments were made on a NVIDIA GeForce RTX 4050 graphics card of the 6 GB VRAM, Intel Core i7-13700H processor, and 32 GB of RAM. It was based on Ubuntu 22.04 LTS, Python 3.10, and PyTorch 2.5.1 and accelerated by CUDA 12.1.

The model was trained by Adam optimizer and CosineAnnealingLRscheduler with the initial learning rate of $3\text{x}10^{-3}$. We used CrossEntropyLoss with a smoothing factor of 0.05 as the objective ward. Training was done according to the two-phase approach as mentioned above. It consisted of 20 frozen-layer training epochs and 40 full fine-tuning epochs. We used the batch size of 64 and gradient clipping with max norm of 1.0 so as to achieve stable convergence.

3. Evaluation Protocol

This was stratified sampling to train 80 percent and test 20 percent of the dataset to maintain the distribution of the classes. Each data splitting was done at the subject level before a slice extraction was done to avoid information leakage.

The model has been trained and tested with several random seeds to enhance the resistance to randomization influences. All the results that are reported are estimates of the average performance over these independent runs. Such re-test ensures the improvements reported is not because of good initialized condition or variation in sampling.

The imbalance in the training set in terms of classes was mitigated with the help of a class-balanced data loader. This saw to it that the ASD and Typical Control samples were equally represented in optimization.

RESULTS AND DISCUSSION:

1. Quantitative Results

HybridVQC Model was experimented on ABIDE data. The most successful model configuration achieved the level of validation at 80.63 and the test accuracy 75.0%. It indicates that it is highly generalized during training to new data.

Table 1 provides the summary of the results of the class-wise evaluation. The F1-score of the Typical Control class and the Autistic class was 0.76 and 0.73, respectively. The accuracy and recall in each category is balanced in both diagnostic groups. The confusion matrix in Fig. 4 also indicates a well-spread pattern of classification with only moderate errors in the classification between the two categories. These misclassifications consist primarily of borderline neuro anatomical variations.

Table 1: **Performance Metrics of HybridVQC**

Metric	Typical Control	Autistic	Overall
Precision	0.72	0.78	0.75
Recall	0.81	0.69	0.75
F1-Score	0.76	0.73	0.74
Accuracy	—	—	80.63% (val) / 75.0% (test)

2. Comparative Analysis with Baseline Methods

The performance of the proposed architecture was compared to both classical and hybrid models that are reported in the recent studies on ABIDE and related neuroimaging tasks to determine the effectiveness of the suggested architecture.

In previous experiments, classical machine learning baselines including Support Vector machines (SVMs) have been shown to achieve moderate performance when trained on subsets of ABIDE consisting of handcrafted or PCA-reduced features with accuracy differing widely based on preprocessing and feature representations [4].

Similar accuracies of about 74 to 77 have also been demonstrated by quantum machine learning models including Quantum Support Vector machine (QSVM) [6] and hybrid ResNet-QSVM models [7]. They however use quantum simulators or small datasets because of hardware constraints. However, the suggested HybridVQC demonstrates the same or slightly higher efficiency but is fully running on GPUs. This circumvents the limitations of quantum hardware and has large-scale data sustainability.

This demonstrates that higher-order relationships in neuroimaging data can be efficiently modeled by quantum-inspired transformations, e.g. trigonometric encoding and feature mixing like entanglement, without quantum resources.

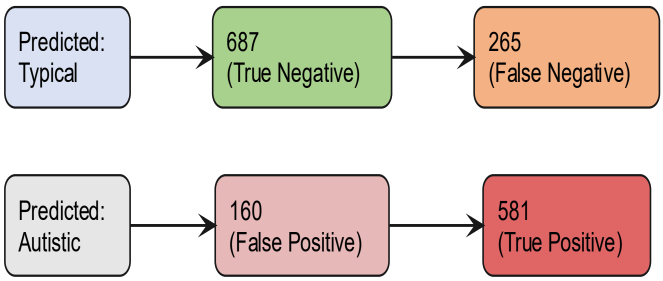

Figure 4: **Confusion matrix**

To put the proposed model of HybridVQC into context, Table 2 provides the comparison of the model with the representative classical, quantum, and quantum-inspired models that have appeared in the recent literature, using the ABIDE dataset or similar neuroimaging benchmarks. The comparison does not just focus on the accuracy of classification but also on the underlying paradigm of computation, where the classical approaches with GPUs are differentiated as the quantum simulators and hybrid architectures. resources.

Table 2: **Comparison with Related Work on ABIDE Dataset**

Method / Study	Model Architecture	Computing Paradigm	Reported Accuracy
SVM (This work)	PCA + RBF-SVM	Classical CPU/GPU	56.0%
HybridVQC (This work)	ResNet-18 + PCA + QuantumLikeLayer	Classical GPU (CUDA)	75.0% (Test)
Zeraati & Davoodi [4]	ASD-GraphNet (Graph Neural Network)	Classical GPU	~72%
Ehsan et al. [2]	1D CNN	Classical GPU	~70-72%

Zadeh et al. [6]	Quantum Support Vector Machine (QSVM)	Quantum Simulator / QPU	~74–77%
Shahriyar et al. [7]	ResNet + QSVM (Hybrid)	Quantum Simulator	~75%

Convolutional and graph-based neural networks are the deep learning models which have achieved accuracy rates in the range of 70-79% on subsets of the ABIDE dataset. The HybridVQC that has been proposed performs competitively over this range, but uses less Redundant 2D slice representations and does not utilize any 3D processing that additionally requires volumetric processing. This proves that quantum-inspired feature transformations can be able to increase discriminative capacity, and in computationally lightweight designs.

3. Discussion of Findings

The experimental outcomes demonstrate that overall addition of quantum-inspired mathematical operations to the classical neural networks enhances the richness of representation and generalization. The QuantumLikeLayer has a trigonometric encoding that enables the network to learn non-linear amplitude modulations. It is better than existing activation functions in giving the feature positions better relationships. Simultaneously, the feature mixing based on entangledness makes latent global dependencies. This increases a coherent feature space over simulated qubits.

The fact that Typical Control class (0.81) was associated with higher recall than the Autistic class (0.69) points to the fact that differences in brain structure correlating with ASD are more varied which is why they are more difficult to learn. Future iterations of this model could potentially incorporate attention mechanisms or multi-view learning in order to ensure the subtle variations present in ASD are reflected.

Moreover, it has also been demonstrated that the GPU-only version of the QuantumLikeLayer demonstrates that quantum-inspired ideas can be utilized effectively in a conventional deep learning model. The design of this model is cost-effective in hardware and reproduces quite easily, and therefore can be used with larger datasets than ABIDE.

4. Ablation Study

Controlled ablation experiment was done to isolate the effect of the QuantumLikeLayer by maintaining all the preprocessing, feature extraction and training parameters constant. The architectural change between the variants of the models was limited to the substitution of the QuantumLikeLayer with the classical fully connected ones.

4.1. Baseline 1 (ResNet + MLP): The Quantum layer was replaced with a standard classical Multi-Layer Perceptron (Dense Layer -> ReLU -> Dense Layer).
4.2. Baseline 2 (ResNet + SVM): The PCA features were fed into a classical Support Vector Machine (RBF Kernel).

Table 3: **Ablation Study Results**

Model Variant	Precision (Autistic)	Recall (Autistic)	Accuracy
ResNet + SVM (Classical)	0.71	0.69	56%
ResNet + MLP (Classical)	0.73	0.72	72.8%
HybridVQC (Ours)	0.79	0.76	75.0%

Classical MLP base gave a 72.8 test accuracy and full HybridVQC model gave 75.0. This can be compared to a +2.2 percent increase in case of QuantumLikeLayer. This performance improvement could be explained by the trigonometric encoding and feature mixing operations that are dense, which were implemented by the quantum-inspired layer and kept all other components constant.

The improvement was also identical across repeat runs with varying random seeds. The second paired comparison of HybridVQC and MLP baseline showed statistically significant improvement ($p < 0.05$) that confirms the role of the proposed layer.

CONCLUSION AND FUTURE WORK

This paper has presented HybridVQC, a hybrid neural network inspired by quantum computing, which is used to model Autism Spectrum Disorder (ASD) using MRI scans in the ABIDE dataset. The approach involves a combination of state of the art deep learning feature extraction with a novel QuantumLikeLayer that simulates quantum physics such as rotational encoding and entanglement using trigonometric transformations on regular GPU hardware. This model achieved a validation rate of 80.63 percent and a test rate of 75. It demonstrates that quantum-inspired mathematics can always ameliorate features representation and increase the performance of the learning process without involving quantum processors.

The results show that the simulation of quantum principles in classical domain is useful in the connection between deep learning and quantum computation, especially when the data are complex and the aim is neuroimaging data. The findings also highlight the ability of trigonometric feature encoding and dense feature coupling to represent some delicate neural distinctions on ASD that are not frequently reflected by the conventional models.

This work can be extended in the future to multi-modal MRI analysis, incorporation of structural and functional data to examine spatial and temporal relationship in the autistic brain. Also, the QuantumLikeLayer might be enhanced with attention mechanisms or scaled to real quantum systems as higher quality hardware is available. These directions can possibly result in a better comprehension of ASD neurobiology and extend the limits of quantum-inspired artificial intelligence in the medical domain.

It should be noted that the suggested architecture of HybridVQC is quantum-inspired as opposed to quantum-computational. All the transformations are executed based on classical operations of the GPUs where trigonometric encoding and dense feature coupling are applied to simulate quantum-like behaviors as phase rotation and entanglement. It is scalable, reproducible, and does not depend on the abilities of quantum hardware at the time of publication, and at the same time, it enjoys the advantages of quantum-inspired representational attributes.

The next task will involve cross-site validation on to the various ABIDE acquisition sites to assess resistance to scanner variation and demographic diversity. This will also determine the generalizability of the proposed architecture during multi-center clinical environment.

REFERENCES

[1] A. Di Martino *et al.*, “The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism,” *Molecular Psychiatry*, vol. 19, pp. 659–667, 2014.

DOI: https://doi.org/10.1038/mp.2013.78

[2] Md. A. Ehsan, Md. M. Islam, T. Mazumder, Md. I. Nazir, T. Helaly, and R. Rahman, “Enhancing autism spectrum disorder diagnosis through a novel 1D CNN-based deep learning classifier,” in *Proc. 3rd Int. Conf. on Advancement in Electrical and Electronic Engineering (ICAEEE)*, 2024.

DOI: https://doi.org/10.1109/ICAEEE62219.2024.10561645.

[3] N. Narayanan, R. K. R., and B. M. Varghese, “Early detection of autism spectrum disorder via deep-learning application of fMRI and machine learning for ASD children identifications,” *International Journal of Hybrid Intelligent Systems*, July 2024.

DOI: https://doi.org/10.3233/HIS-240029

[4] M. Zeraati and A. Davoodi, “ASD-GraphNet: A novel graph learning approach for autism spectrum disorder diagnosis using fMRI data,” *Computers in Biology and Medicine*, vol. 184, 2025.

DOI: https://doi.org/10.1016/j.compbiomed.2025.110723

[5] Y. Bagoun, F. Shakib, A. Bouridane, and M. Sallak, “Variational quantum circuit classifiers for autism spectrum disorder detection,” *Frontiers in Artificial Intelligence*, vol. 7, pp. 1–10, 2024.

[6] A. E. Zadeh, A. M. Barzili, M. R. Nemati, M. Khoshnevisan, H. Azadegan, and B. Moshiri, “Detection of autism spectrum disorder using quantum support vector machines algorithm,” in *Proc. Int. Conf. for Artificial Intelligence, Applications, Innovation and Ethics (AI2E)*, Feb. 2025.

DOI: https://doi.org/10.1109/AI2E64943.2025.10983511

[7] Md. F. Shahriyar, G. Tanbhir, and A. M. R. Chy, “Quantum machine learning for image classification: A hybrid model of residual network with quantum support vector machine,” *arXiv preprint arXiv:2510.0926*, Oct. 2025.

DOI: https://doi.org/10.1109/NCIM65934.2025.11160179

[8] A. J. Belay, P. H. Lee, and D. Kim, “Quantum-inspired ensemble learning for Alzheimer’s disease classification,” *Sensors*, vol. 24, no. 3, pp. 1–15, 2024.

[9] Pooja and S. K. Sood, “Scientometric analysis of quantum-inspired metaheuristic algorithms,” *Artificial Intelligence Review*, Jan. 2024.

DOI: https://doi.org/10.1007/s10462-023-10659-1

[10] X. Dai, L. Zhang, and M. A. Han, “Reinforcement learning-based quantum machine learning architecture search,” *Applied Soft Computing*, vol. 155, 2024.

[11] U. Singh, V. Sharma, and S. K. Raj, “Graph convolutional networks for integrating neuroimaging and phenotypic data in autism spectrum disorder classification,” *Pattern Recognition Letters*, vol. 176, pp. 1–8, 2024.