Leveraging Triplet Autoencoders with Quantum Tensor Networks for Multiclass Image Classification*

Vito Nicola Losavio[1,*,†], Maria Grazia Miccoli[1,†] and Donato Malerba[1,†]

[1]*Department of Computer Science, University of Bari Aldo Moro, Italy*

Abstract

This paper presents a hybrid quantum-classical pipeline for multiclass classification, integrating a triplet autoencoder (T-AE) with quantum circuits inspired by tensor network structures. The T-AE compresses high-dimensional input data into a semantically structured latent space, thereby enhancing class separability through triplet loss regularisation. These latent embeddings are then mapped onto quantum states using a modified angle encoding scheme, after which they are processed by parameterised quantum circuits inspired by Matrix Product States (MPS) and Tree Tensor Networks (TTN). This architecture captures both local and global correlations while maintaining a reduced number of parameters. Experimental results on the MNIST dataset demonstrate that the proposed pipeline achieves superior accuracy to PCA-based baselines, particularly in scenarios with limited data. These results highlight the potential of combining structured latent representations with tensor-inspired quantum models for efficiently addressing complex classification tasks.

Keywords

Quantum tensor network, variational quantum algorithm, triplet loss

1. Introduction

Tensor networks (TNs) have emerged as a modern numerical method for representing and processing high-dimensional data in computational tasks for a growing number of applications. Their main strength lies in the ability to reduce both data complexity and computational cost through the decomposition of global objects into networks of locally connected components. Originally developed in the field of quantum many-body physics, structures such as matrix product states (MPS) and tree tensor networks (TTN), have proven effective at capturing complex correlations with a reduced number of parameters, which gives a clear demonstration of their ability in providing efficient computation. This makes them a scalable alternative to dense representations in classical machine learning systems[1, 2]. TNs have more recently been explored in machine learning as tools for building expressive and compact models that perform structured feature selection through their topology and bond dimensions [3].

It has been proved that tensor networks can represent losslessly the 2^n quantum states of a n-qubit circuit [4]. They are able to accommodate quantum properties, such as entanglement, thanks to bond dimensions and TN topologies [5]. Additionally, contraction operations on small-rank tensors allows to simulate the property the quantum superposition [6]. This has inspired the design of efficient quantum models and opened to new approaches in quantum machine learning (QML). One of the main results has been the development of variational models that incorporate physical constraints while maintaining learning capabilities [7]. This synergy is particularly promising in the current Noisy Intermediate-Scale Quantum (NISQ) era, where circuit depth and qubit connectivity are limited [8].

The literature shows that tensor networks have been particularly successful in supervised learning [1], especially in classification problems. Previous studies have demonstrated the advantages on binary classification [3], while more intricate tasks are still to be explored. Probably, the first one of this list

*AIQ*QIA 2025 3nd International Workshop on AI for Quantum and Quantum for AI*

*You can use this document as the template for preparing your publication. We recommend using the latest version of the ceurart style.

*Corresponding author.

†These authors contributed equally.

v.losavio5@studenti.uniba.it (V. N. Losavio); m.miccoli45@studenti.uniba.it (M. G. Miccoli); donato.malerba@uniba.it (D. Malerba)

is represented by multi-class classification, where one cannot design or employ directly approaches originally developed for binary case, since the generalization of methods is not trivial.

2. Contribution

Multi-class classification represents the problem of categorization which is more frequent finding in the real-world applications and represents the increasing of the challenge for binary classificaiton in QML. In [9] , the authors apply tensor network-based circuit architectures to digit recognition using a subset of the MNIST dataset and quantum phase recognition. Results for digit recognition were compared against classical benchmarks. The circuits tested were based on tree tensor networks (TTN) and multiscale entanglement renormalization ansätze (MERA). Two approaches for implementing the multi-class setup were explored: amplitude decoding, which extracts predictions from computational basis measurements, and qubit decoding with binary labels, which computes predictions using the expectation value of single-qubit observables. The study focused on circuits utilizing eight qubits.

Another approach in this field is proposed by Murota et al.[10] propose the forest tensor network classifier, a framework comprising multiple TTN classifiers that form an intermediate representation used subsequently for classification.

As to the binary classification with tensor networks, [11] provides several contributions through an hybrid method. Firstly, they propose a kernel encoding of quantum data, XX entanglement gate to share the correlation between adjacent qubits, and use stochastic gradient descent to avoid the analysis gradient too costly.

There is growing list of quantum-inspired/based versions of CNN methods for multi-class classification. [12] proposed a quantum CNN architecture for multiclass image classification, using a new convolutional architecture using 3-qubit gates, a new pooling architecture that leverages the reversibility of quantum gates and tested a set of configurations for the densely connected layer. In [13], after convolutional layers, a few fully connected layers are followed by a quantum variational circuit, whose outputs are passed to a softmax loss function and optimized by adjusting the quantum parameters. A modified quantum perceptron is also introduced, enabling high classification accuracy with a compact quantum circuit.

It is evident that the approaches that have been employed with great success are those that utilise Quantum Convolutional Neural Networks (QCNN) for multiclass image classification by analysing either an increasing number of classes [14, 15] or different domains [16], achieving good performance.

A different approach is proposed by Shi et al.[17], that involves the use of a convolutional neural network (CNN) to extract meaningful information, in combination with a quantum neural network (QNN) to handle the multi-class classification problem.

In this field, we also have the paper proposed by Senokosov et al.[18]. This paper sets out a two-hybrid architecture that incorporates a quantum layer inside the net, using the quantum properties to enhance the performance of the model with fewer parameters.

In this work, we address the multiclass classification task by developing a hybrid-quantum machine learning pipeline that coherently integrates classical and quantum techniques. Specifically, we explore the effective employment of quantum tensor network circuits as variational models for distinguishing between multiple classes.

To prepare the high-dimensional input data for quantum processing, we use a T-AE, so trained with a triplet loss function. This approach enables us to reduce the dimensionality of the original feature space while imposing a structure that draws samples belonging to the same class closer together and pushes samples from different classes further apart. This makes the downstream classification problem more tractable and aligns the data with the constraints of quantum hardware by reducing the number of qubits required for encoding.

We then feed this compact, semantically structured latent representation into quantum circuits that leverage the inherent inductive biases of tensor network topologies. Our aim is to demonstrate that it is possible to achieve competitive multiclass classification performance with fewer parameters than more

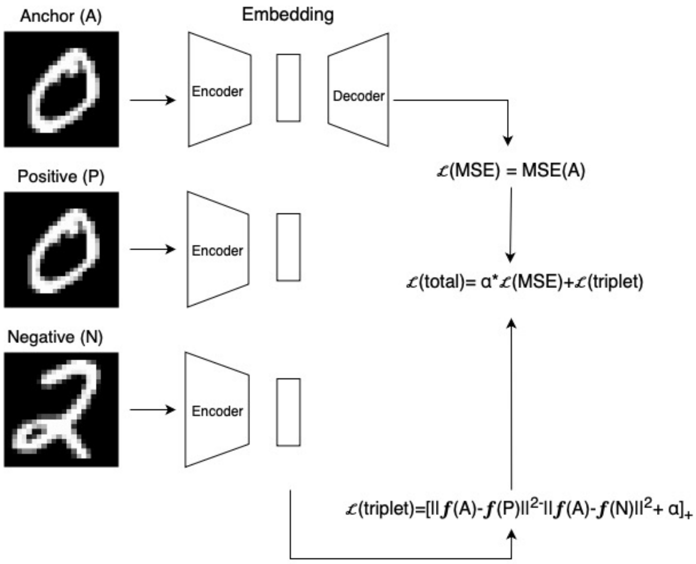

Figure 1: Autoencoder trained using triplet loss.

generic quantum variational models. The combination of advanced feature learning and architecturally motivated quantum circuits has been demonstrated to facilitate scalable, resource-efficient quantum machine learning solutions that are capable of addressing complex multiclass scenarios.

3. Method

As mentioned earlier, our approach starts with transforming the original, high-dimensional input data into a more compact representation using an autoencoder. The autoencoder is specifically trained to minimise a combined loss function that incorporates a *triplet loss* component designed to enforce a geometry in which samples of the same class are pulled closer together and samples of different classes are pushed apart, as well as a reconstruction term based on *mean squared error* (MSE). The total loss is defined as follows:

$$\mathscr{L} = \mathscr{L}_{triplet} + \alpha \times \mathscr{L}_{MSE} \tag{1}$$

where α balances the reconstruction penalty relative to the triplet term.

The triplet loss function was originally proposed by Schroff et al.[19]. This loss function is designed to operate on three samples simultaneously: the current training example (anchor), another sample from the same class (positive), and a sample from a different class (negative). The aim is to ensure that the distance between the anchor and the positive is smaller than the distance between the anchor and the negative by at least a predefined margin. Formally, it is defined as:

$$\mathscr{L}_{triplet} = \sum_{i}^{N} \left[\| f(x_i^a) - f(x_i^p) \|_2^2 - \| f(x_i^a) - f(x_i^n) \|_2^2 + \alpha \right]_+ \tag{2}$$

where x_i^a is the *anchor*, x_i^p is the *positive* and x_i^n is the negative samples, and $[\cdot]_+$ mean that only the positive value are considered, ensuring the hinge loss behaviour.

For the reconstruction component, we use the standard mean squared error (MSE) loss function, which is widely used in autoencoder applications to measure the similarity between the network's output $\hat{x}_i$ and the original input x_i. This term ensures that the autoencoder maintains the ability to faithfully reproduce the input data, in addition to structuring the latent space through the triplet objective. Formally, the MSE is defined as:

$$\mathscr{L}_{MSE} = \frac{1}{n} \sum_{i=1}^{n} (x_i - \hat{x}_i)^2 \tag{3}$$

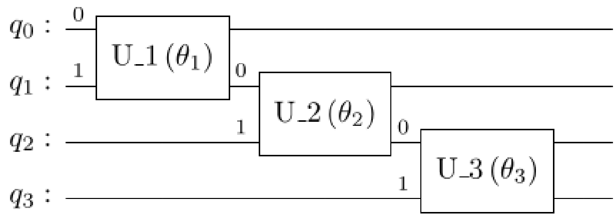

Figure 2: Example of a Matrix Product State (MPS) quantum circuit, where local parameterized unitaries $U_k(\theta_k)$ act sequentially on overlapping subsets of qubits, propagating correlations along the chain.

The combination of this reconstruction loss with the triplet loss serves to guide the autoencoder in its learning of latent representations that are discriminative across different classes, thereby ensuring the preservation of the essential input data features, as visually summarised in Figure 1, which depicts the overall architecture and training objective of our T-AE.

Once the data has been transformed into a compact latent space by means of the T-AE, we turn our attention to modeling the relationships present within this reduced representation. To achieve this, we draw inspiration from well-established methods in classical tensor methods, such as Matrix Product States (MPS) and Tree Tensor Networks (TTN). These frameworks are known for their ability to decompose complex systems into structured forms that efficiently capture correlations across different components. Translating these ideas into the quantum domain, we design circuits that reflect the same underlying principles.

Instead of applying generic, fully connected quantum models to this latent space, we deliberately exploit the taxonomic properties of tensor networks, which naturally organize computations into sequential or hierarchical patterns. We implement these patterns as parameterized quantum circuits, arranging local unitary operations according to specific tensor network topologies like MPS and TTN. This architectural choice allows us to propagate information and build correlations throughout the system in a controlled and interpretable manner, aligning with the structural priors imposed by the original tensor network formulations. In doing so, we aim to harness the efficiency and expressiveness of tensor networks directly within our hybrid-quantum learning pipeline, tailoring it to handle the intricate patterns necessary for robust multiclass classification.

Formally, we can denote each local unitary block in the circuit as $U_k(\theta_k)$, where θ_k represents the associated set of variational parameters. The specific tensor network topology dictates the arrangement of these local unitaries and the establishment and propagation of correlations, such as entanglement, across the system. Within this framework, the global quantum state prepared by the circuit can be expressed as follows:

$$|\psi(\Theta)\rangle = \sum_{k\in\mathscr{E}} U_k(\theta_k)\,|0\rangle^{\otimes n} \tag{4}$$

where $\mathscr{E}$ denotes the set of connections (or edges) dictated by the tensor network topology. Each $U_k(\theta_k)$ acts locally on one or more qubits and contributes to building up the global quantum state. The symbol $\sum$ here does not represent a simple algebraic product, but rather a structured composition of unitary operations, explicitly ordered and interconnected according to the topology of $\mathscr{E}$. It serves as a compact way to express how the network orchestrates the local transformations into a coherent global circuit.

Figure 2 presents an example of the proposed formalisation, illustrating how local parameterized unitaries are sequentially applied according to the MPS topology.

Having formally introduced the tensor network framework and described the T-AE that produces compact latent representations, we will now detail the feature map that embeds these vectors into quantum states. This completes the set of building blocks required to fully specify the quantum circuit architecture proposed in this work.

Specifically, we use a variant of the standard angle encoding scheme, preceding each data-driven rotation R_Y with a *Hadamard gate*. This modification initialises each qubit in a superposition spanning the entire Hilbert space. This choice ensures that subsequent parameterised operations have access to a richer quantum state manifold, which could enhance the circuit's ability to model complex correlations

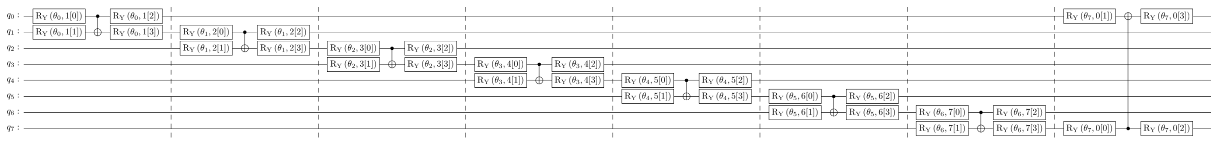

Figure 3: The MPS fully connected block proposed in this work. This section of the circuit sequentially entangles pairs of qubits, directly interrelating all components of the latent embedding. This ensures that each part of the input contributes to the global quantum state, capturing widespread correlations.

within the input data.

Melo et al.[20] and Monnet et al.[21] have already examined a similar type of feature map in other quantum machine learning contexts.

We now bring together all the elements previously introduced into a unified quantum circuit architecture. This design integrates the latent representations, the feature map, and the tensor network-inspired structures into a single workflow tailored for multiclass classification.

The circuit is organized into three main stages, each serving a distinct purpose:

1. *Feature map*: As previously mentioned, it begins by embedding the latent vector into quantum states using our Hadamard angle encoding. This initialises each qubit into a rich superposition, preparing the system to capture intricate correlations through interference and entanglement.
2. *Fully connected MPS section*: The next stage of the circuit employs a sequence of entangling operations arranged in a fully connected manner, inspired by Matrix Product States. The propagation of correlations in a conventional MPS is primarily conducted sequentially along a chain, whereas our design propagates correlations back to the initial qubit.
   This approach is driven by the nature of the input provided. Each qubit encodes an embedding component extracted from the same original image. This ensures that all qubits contain relevant information. Closing the chain in this way ensures that correlations are distributed across the entire register, allowing the circuit to capture global dependencies spanning the full set of embeddings.
   The model is present in figure 3.
3. *TTN hierarchy with measurements*: The final segment of the circuit is inspired by Tensor Tree Networks (TTNs) and implements a tree-like structure, progressively merging and propagating information through successive layers of local unitaries. This hierarchical composition concentrates correlations into fewer qubits as we move up the structure, resulting in a compact set of qubits encapsulating the most globally aggregated features.
   At the end of this process, we perform explicit measurements on these final qubits. The number of qubits measured is determined by the multiclass nature of the classification task. For example, when classifying four categories, we read the states of the last two qubits and map them onto two classical registers. Similarly, for problems involving six, eight or ten classes, we extend the measurement to the last four qubits and use four classical registers to capture a richer set of outcomes. This strategy enables us to adapt the quantum-classical interface to the complexity of the multiclass problem by allocating more measurement capacity as the number of categories increases.
   We are fully aware that measuring only a subset of the qubits results in the partial collapse of the overall quantum state and the inevitable loss of some quantum information. However, in this work, our primary aim is to preserve and exploit the hierarchical informational structure imposed by the tensor network by focusing our measurements precisely on the qubits at the top of the TTN hierarchy. These qubits carry the culmination of correlations and entanglement propagated through all preceding layers, effectively distilling the most globally representative features of the input data.
   By measuring these final qubits explicitly, we extract maximally aggregated information according to the multi-scale architecture of the tensor network. For instance, when measuring the final four qubits of the circuit, up to 16 possible output states are inherently obtained. In our multiclass

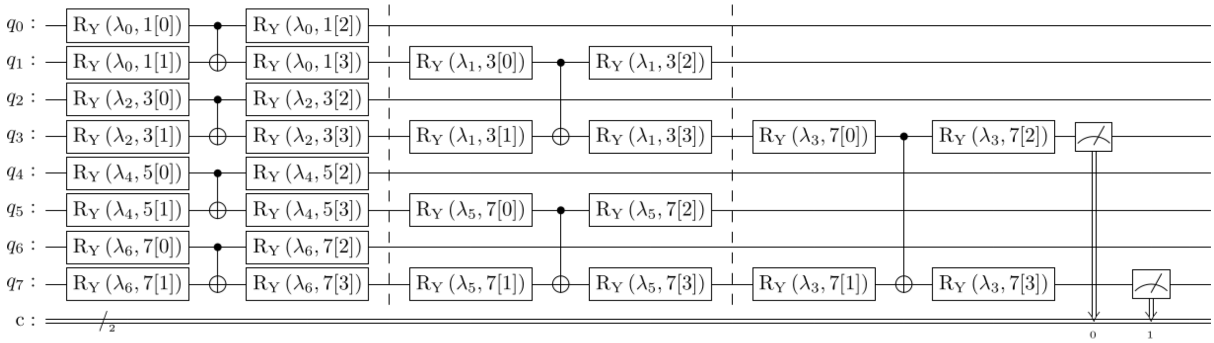

Figure 4: Proposed TTN-inspired quantum circuit used for multiclass classification. The hierarchical structure progressively aggregates information from the fully connected MPS stage, routing correlations upwards through successive layers of local unitaries. Measurements are performed on the final qubits at the top of the hierarchy, encapsulating the most globally representative features.

experiments, we typically target a smaller set of categories, such as *four*, *six*, *eight* or *ten* classes. To map the raw quantum measurement outcomes into the predefined class labels, we incorporate an additional interpret function. This strategy enables us to adapt the expressive capacity of the quantum circuit to the specific classification task at hand while maintaining the critical flow of hierarchical information enabled by the TTN structure.
This architecture is represented by the figure shown in Figure 4.

The interpret function that we use is designed to take each bitstring obtained from quantum measurements and perform a simple modulo operation with respect to the number of target classes. This ensures that each outcome is consistently assigned to one of the predefined class labels, regardless of how many distinct bitstring outcomes can be produced by measuring the final qubits. This strategy offers a straightforward, computationally efficient method of linking the quantum output space to the classical multiclass decision space, thus seamlessly integrating the final stage of our quantum classification pipeline.

4. Experiments

For our experimental evaluation, we implemented the proposed quantum tensor network models using the Qiskit [22] framework along with the Qiskit Machine Learning extension [23], which provides convenient abstractions for building variational quantum algorithm (VQA) and integrating them into hybrid learning pipelines. As a benchmark dataset, we selected MNIST[24], a widely used collection of handwritten digit images that offers a standard testing ground for multiclass classification algorithms. This combination of robust quantum computing libraries and a well-understood dataset allows us to rigorously assess the performance and behavior of the proposed architectures under realistic conditions. In the quantum circuits of the tensor network designed for this study, the local unitary blocks U_k introduced in our formalisation are instantiated as *RealAmplitudes circuits* on two-qubit with a single repetition. This ensures that each core block can generate entanglement among the qubits on which it acts, providing the essential mechanism for propagating correlations throughout the network. The effect of these RealAmplitudes circuits in establishing entanglement and mixing the input features can also be seen in Figures 3 and 4. For our quantum experiments, we utilized the Qiskit Aer simulator backend, which facilitates high-fidelity emulation of quantum circuits on classical hardware. This choice enables scalable and reproducible testing of our variational quantum models, free from the noise and hardware limitations associated with current NISQ devices.

During the training phase, we constructed our datasets by selecting varying numbers of samples per class from the original MNIST training set. In particular, we explored training regimes with 50, 100, 150, 200, 250, and 300 examples per class to investigate how the model's performance evolves as the number of available examples increases. For each multiclass scenario, we included only the digits corresponding

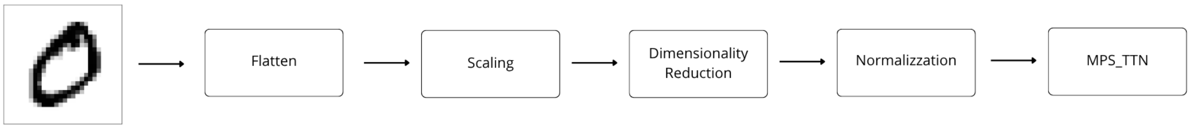

Figure 5: Proposed pipeline: starting with the image flattening, then scaling and dimensionality reduction with AE, followed by normalisation in quantum states to give the model this representation.

to the number of target classes under consideration; for instance, when tackling a four-class problem, we used samples drawn from the first (the 0 digit) four digits (the 3 digit). To thoroughly evaluate generalisation, we retained all available test samples per class from the MNIST test set, which typically contains around $1,000$ examples for each digit. This experimental design aligns with the concept of few-shot learning [25], as the training set remains relatively small compared to the extensive test set, thereby highlighting the model's ability to learn robust decision boundaries even with limited data.

First, the T-AE was trained on these normalized training examples, where pixel values were scaled by dividing by 255 to lie within the range $[0, 1]$. We set $\alpha = 1$, which means that the reconstruction loss was given the same importance as the triplet loss during training, thus balancing the objective of accurately reconstructing the input data with the objective of structuring the latent space to reflect class similarities. Once the latent space representations had been learned, the embeddings were extracted and further normalised to be within the interval $[0, \pi]$. This step was specifically designed to align with the angle encoding requirements used in the quantum circuit. We then trained the VQA on these normalised latent vectors, using the *COBYLA* optimiser for $1,000$ iterations.

During the testing phase, we evaluated the trained quantum models using the full set of test samples available for each selected class in the MNIST dataset. This approach ensures a rigorous assessment of the generalisation capabilities of our quantum tensor network architectures, which is particularly important given the relatively limited number of training samples employed. The latent embeddings for these test inputs were generated by passing the normalised images through the previously trained T-AE and scaling them to the range $[0, \pi]$, as was done during training. Predictions were obtained by feeding these latent representations into the trained variational quantum circuits. The final class labels were then assigned using the interpret function described earlier.

The proposed pipeline is illustrated in Figure 5, which shows the process from data pre-processing to quantum classification. Within this framework, we conducted a comparative analysis to evaluate the effectiveness of our latent space based autoencoder. Specifically, we replaced the autoencoder with the standard dimensionality reduction technique of *Principal Component Analysis* (PCA) and repeated classification experiments using the same quantum tensor network architecture.

The aim of this comparison is to assess whether training a compact latent representation using a T-AE can outperform conventional linear methods such as PCA, which are commonly used in hybrid quantum-classical machine learning pipelines. By keeping the quantum architecture the same and only varying the feature extraction mechanism, we can isolate the autoencoder's contribution and evaluate its effect on model performance in settings with limited data.

Table 1 compares the classification accuracy achieved with standard PCA and the proposed T-AE as dimensionality reduction strategies. As can be seen, the T-AE consistently outperforms PCA across all configurations, particularly in complex scenarios or when data is limited.

In the four-class task, the T-AE achieves its best performance with 200 or 300 training samples per class, reaching 97.19% and 98.00%, respectively. While PCA yields comparable accuracy with 100 samples (86.05%), its performance is inconsistent and deteriorates rapidly with 150 samples (69.64%), indicating a lack of robustness. In contrast, the T-AE maintains high accuracy across different sample sizes, confirming the stability of the learned representations.

The difference becomes even more striking in the six-class task. Although PCA struggles to surpass 65.5%, even with 200 samples, T-AE achieves 93.22% under the same conditions and delivers consistent

Samples per class	4 Classes		6 Classes	
	PCA	Triplet AE	PCA	Triplet AE
50	81.09	78.42	56.66	77.62
100	86.05	94.49	60.11	67.73
150	69.64	92.13	53.86	90.17
200	82.22	**97.19**	65.54	**93.22**
250	80.78	96.92	59.71	81.46
300	80.30	**98.00**	60.85	86.40

Table 1
Classification accuracy (%) for 4 and 6 classes using PCA and T-AE

performance above 86% from 150 samples onward. Even in the most data-scarce scenario of 50 samples, the autoencoder achieves a substantial improvement on PCA (77.62% versus 56.66%).

These results suggest that the proposed T-AE learns more expressive latent features and provides a more stable and scalable representation as the number of classes and samples increases. This supports its suitability for integration with quantum tensor networks, particularly for low-data and high-complexity classification tasks.

5. Conclusion

In this study, we demonstrated that combining a compact, semantically rich latent representation produced by a T-AE with a quantum circuit architecture inspired by tensor network structures can effectively address complex multiclass classification tasks. Leveraging the expressiveness of hierarchical quantum models and the discriminative power of the learned latent space enabled us to achieve high classification performance, even in low-data regimes.

The proposed hybrid circuit has demonstrated robust generalisation capabilities across various classification settings. Our results confirm that integrating structured quantum circuits with meaningful input representations enables the efficient use of quantum resources by reducing the number of trainable parameters, while achieving performance that is competitive with, and sometimes superior to, that of standard baselines and classical dimensionality reduction approaches such as PCA.

These findings reinforce the idea that quantum tensor network models are theoretically elegant and practically viable tools for quantum machine learning, particularly when dealing with complex decision boundaries or limited training data. The combination of expressive quantum architectures and task-specific encoding strategies suggests promising directions for future developments in scalable, efficient quantum learning systems.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-4o in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] E. M. Stoudenmire, D. J. Schwab, Supervised learning with tensor networks, in: NIPS, 2016, pp. 4799–4807.

[2] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, E. M. Stoudenmire, Towards quantum machine learning with tensor networks, Quantum Science and technology 4 (2019) 024001.

[3] D. Guala, S. Zhang, E. Cruz, C. A. Riofrío, J. Klepsch, J. M. Arrazola, Practical overview of image classification with tensor-network quantum circuits, Scientific Reports 13 (2023) 4427.

[4] A. A. Melnikov, A. A. Termanova, S. V. Dolgov, F. Neukart, M. Perelshtein, Quantum state preparation using tensor networks, Quantum Science and Technology 8 (2023) 035027.
[5] E. Campos, A. Nasrallah, J. Biamonte, Abrupt transitions in variational quantum circuit training, Phys. Rev. A 103 (2021) 032607. URL: https://link.aps.org/doi/10.1103/PhysRevA.103.032607. doi:10.1103/PhysRevA.103.032607.
[6] D. Guala, E. Cruz-Rico, S. Zhang, J. M. Arrazola, Tensor-network quantum circuits, https://pennylane.ai/qml/demos/tutorial_tn_circuits, 2022. Date Accessed: 2025-07-04.
[7] R. Huang, X. Tan, Q. Xu, Variational quantum tensor networks classifiers, Neurocomputing 452 (2021) 89–98.
[8] J. Preskill, Quantum Computing in the NISQ era and beyond, Quantum 2 (2018) 79. URL: https://doi.org/10.22331/q-2018-08-06-79. doi:10.22331/q-2018-08-06-79.
[9] M. Lazzarin, D. E. Galli, E. Prati, Multi-class quantum classifiers with tensor network circuits for quantum phase recognition, Physics Letters A 434 (2022) 128056.
[10] K. Murota, T. Kobori, Trainable quantum neural network for multiclass image classification with the power of pre-trained tree tensor networks, CoRR abs/2504.14995 (2025). URL: https://doi.org/10.48550/arXiv.2504.14995. doi:10.48550/ARXIV.2504.14995. arXiv:2504.14995.
[11] R. Huang, X. Tan, Q. Xu, Variational quantum tensor networks classifiers, Neurocomputing 452 (2021) 89–98.
[12] S. Kashyap, S. S. Garani, Quantum convolutional neural network architecture for multi-class classification, in: 2023 International Joint Conference on Neural Networks (IJCNN), IEEE, 2023, pp. 1–8.
[13] D. Bokhan, A. S. Mastiukova, A. S. Boev, D. N. Trubnikov, A. K. Fedorov, Multiclass classification using quantum convolutional neural networks with hybrid quantum-classical learning, Frontiers in Physics 10 (2022) 1069985.
[14] M. Mordacci, D. Ferrari, M. Amoretti, Multi-class quantum convolutional neural networks, in: A. Pavone, C. Viola (Eds.), Proceedings of the 2024 Workshop on Quantum Search and Information Retrieval, QUASAR 2024, Pisa, Italy, 3 June 2024, ACM, 2024, pp. 9–16. URL: https://doi.org/10.1145/3660318.3660326. doi:10.1145/3660318.3660326.
[15] S. Shi, Z. Wang, J. Li, Y. Li, R. Shang, G. Zhong, Y. Gu, Quantum convolutional neural networks for multiclass image classification, Quantum Inf. Process. 23 (2024) 189. URL: https://doi.org/10.1007/s11128-024-04360-7. doi:10.1007/S11128-024-04360-7.
[16] H. K. Ahmed, B. Tantawi, G. I. Sayed, Multiclass image classification based on quantum-inspired convolutional neural network, in: S. Subbotin (Ed.), Proceedings of The Sixth International Workshop on Computer Modeling and Intelligent Systems (CMIS 2023), Zaporizhzhia, Ukraine, May 3, 2023, volume 3392 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 177–187. URL: https://ceur-ws.org/Vol-3392/paper15.pdf.
[17] M. Shi, H. Situ, C. Zhang, Hybrid quantum neural network structures for image multi-classification, Physica Scripta 99 (2024) 056012. URL: https://dx.doi.org/10.1088/1402-4896/ad3e3d. doi:10.1088/1402-4896/ad3e3d.
[18] A. Senokosov, A. Sedykh, A. Sagingalieva, B. Kyriacou, A. Melnikov, Quantum machine learning for image classification, Machine Learning: Science and Technology 5 (2024) 015040. URL: https://dx.doi.org/10.1088/2632-2153/ad2aef. doi:10.1088/2632-2153/ad2aef.
[19] F. Schroff, D. Kalenichenko, J. Philbin, Facenet: A unified embedding for face recognition and clustering, in: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, 2015, p. 815–823. URL: http://dx.doi.org/10.1109/CVPR.2015.7298682. doi:10.1109/cvpr.2015.7298682.
[20] A. Melo, N. Earnest-Noble, F. Tacchino, Pulse-efficient quantum machine learning, Quantum 7 (2023) 1130.
[21] M. Monnet, N. Chaabani, T.-A. Drăgan, B. Schachtner, J. M. Lorenz, Understanding the effects of data encoding on quantum-classical convolutional neural networks, in: 2024 IEEE International Conference on Quantum Computing and Engineering (QCE), volume 1, IEEE, 2024, pp. 1436–1446.
[22] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation,

L. S. Bishop, A. W. Cross, et al., Quantum computing with qiskit, arXiv preprint arXiv:2405.08810 (2024).

[23] M. E. Sahin, E. Altamura, O. Wallis, S. P. Wood, A. Dekusar, D. A. Millar, T. Imamichi, A. Matsuo, S. Mensa, Qiskit machine learning: an open-source library for quantum machine learning tasks at scale on quantum hardware and classical simulators, arXiv preprint arXiv:2505.17756 (2025).

[24] Y. LeCun, C. Cortes, C. Burges, Mnist handwritten digit database, ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist 2 (2010).

[25] Y. Wang, Q. Yao, J. T. Kwok, L. M. Ni, Generalizing from a few examples: A survey on few-shot learning, ACM computing surveys (csur) 53 (2020) 1–34.