Quantum Circuits for Quantum Convolutions: A Quantum Convolutional Autoencoder

Javier Orduz[1], Pablo Rivas[2], and Erich Baker[3]

School of Engineering and Computer Science
Department of Computer Science
Baylor University, Texas, USA
[1]Javier_OrduzDucuara@baylor.edu [2]Pablo_Rivas@baylor.edu
[3]Erich_Baker@baylor.edu

Abstract. Quantum machine learning deals with leveraging quantum theory with classic machine learning algorithms. Current research efforts study the advantages of using quantum mechanics or quantum information theory to accelerate learning time or convergence. Other efforts study data transformations in the quantum information space to evaluate robustness and performance boosts. This paper focuses on processing input data using randomized quantum circuits that act as quantum convolutions producing new representations that can be used in a convolutional network. Experimental results suggest that the performance is comparable to classic convolutional neural networks, and in some instances, using quantum convolutions can accelerate convergence.

Keywords: Quantum Computing, Convolutional Autoencoder, Quantum Machine Learning

1 Introduction

Neural network architectures based on convolutional operations are a trendy machine learning tool [22,1,9]. These kinds of models offer certain versatility when the learned filters are transferred into or out of the network [14]. Beyond the apparent filtering capabilities, convolutional neural networks (CNNs) have other interesting properties and architectures that have produced many exciting applications, and variants [27,28,1]. This research is based on a classic autoencoder CNN architecture with convolutional and pooling layers.

An autoencoder (AE) is considered an unsupervised learning model that reconstructs the input signal using a neural network [4]. AEs are notably known for some of their successful versions, including the Variational Autoencoder (VAE) [24], and the denoising AE [18,16]. Convolutional AEs, in particular, have been proven to be very robust in learning representations of the data, usually compressed, while at the same time retaining much of the information [7]. In this research, we focus on an AE for image-related tasks, learning representations leveraging **quantum computing**, and replacing the first layer of the model.

The field of quantum computing has grown recently, providing new means to calculate, represent, and read information [3]. A recent study shows quantum machine learning can be considered a kernel transformation in classic machine learning [19]. Therefore, here we further explore the idea of using a quantum circuit to process images and then feed them to a convolutional autoencoder, which we call a *quanvolutional autoencoder* [15]. This work is also motivated by others who explored a quanvolutional network for classification purposes [8].

The rest of the paper is organized as follows. Section 2 describes state of the art in quantum machine learning. Section 3 describes the proposed quanvolutional autoencoder architecture. Section 4 describes the results, and a discussion is presented in Section 5, along with our conclusions.

2 Background

Quantum computation (QC) has permeated different computation areas; in particular, machine learning is a tangible subject to explore with quantum algorithms. Scientific community explores different tools and potential applications tackled with quantum algorithms and concepts related to Quantum Mechanics [6]. On the other hand, Machine Learning (ML) techniques and algorithms, such as Supervised and Unsupervised, have been implemented in some python platforms. This intersection of ML and QC is called Quantum Machine Learning [3,20,25].

One of the most interesting applications in machine learning models, in particular, autoencoders [5]. This is a kind of artificial neural network belonging to the unsupervised learning, and it is useful to reduce the dimensions. The purpose of this document is to show some quantum advantages, several applications, in the Quanvolutional Neural Networks, which is model explored by [8].

Autoencoders are generally used in tasks associated with information compression, and summarization [21,26], and generally speaking in dimensionality reduction [23,16]. Fig. 1 depicts the basic representation of an autoencoder as a function with parameters that need to be learned efficiently. AEs are unsupervised models that aim to take an input, find intermediate linear and non-linear transformations, and reproducing or reconstructing the input at the output layer. Autoencoders usually minimize some reconstruction loss function such as the mean squared error or categorical cross-entropy reconstruction losses [1]. The model presented in this work is a convolutional autoencoder whose architecture we discuss next.

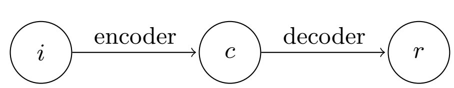

Fig. 1. Basic autoencoder, where i refers to the input, c is the code, and r is the output.

3 Autoencoder Architecture

We proposed to use the autoencoder architecture shown in Fig. 2. At the

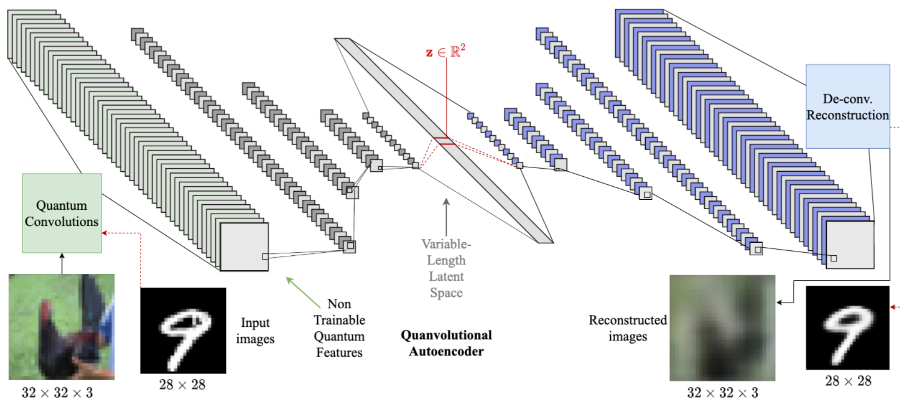

Fig. 2. *Top:* Quanvolutional autoencoder architecture, where the quantum filtering (quantum convolutions) are described as quantum circuits, as shown in Fig. 3. The quantum layer is non-trainable by the network. *Bottom:* Classic convolutional autoencoder with all its layers trainable. Note that the dense layer in the center of the AEs can also vary, producing a two-dimensional latent space or a larger space.

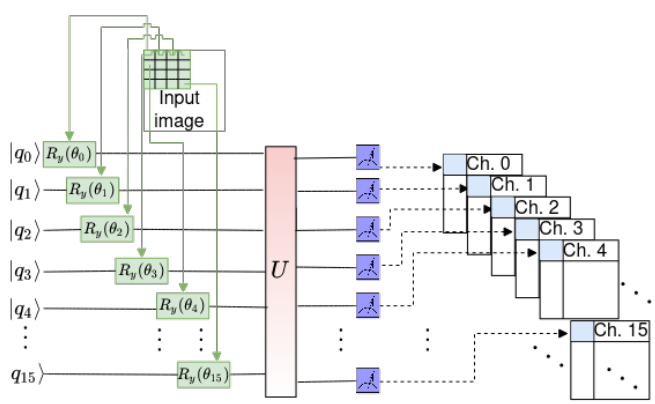

Fig. 3. The quantum circuit for the autoencoder considering 16 qubits for a 4×4 quantum convolutional filter.

top and bottom of the figure, we have two similar configurations which differ in the first convolutional layer. The traditional autoencoder at the bottom of

Fig. 2 shows an autoencoder that has the traditional convolutional, pooling, de-convolutional, and upsampling layers. However, the top of the figure shows an alternative configuration that uses quantum-based filters. These filters are implemented as quantum circuits that can produce an interesting and alternative way of representing images in lower-dimensional filters. Note, however, that in this particular configuration, we are using a non-trainable set of quantum-based filters. To this kind of architecture, we refer to as **Quanvolutional Autoencoder**. The name is inspired by the work of Henderson et al. [8] on quanvolutional neural networks. Although a quantum version of the autoencoder has already been proposed by Romero et al. in 2017 [17], such model and the one presented here are substantially different and come from different motivations; our approach, in contrast, uses randomized quantum circuits as image convolutions.

Fig. 3 shows $|q_i\rangle$ as the qubits on a 4×4 region, each element of this region is associated to the same numbers of qubits initialized in ground state after we apply rotations, which are parametrized to angle θ and scaled by π, to embed the information in the qubits. This rotation is given by,

$$R_y(\theta) = e^{-i\theta\frac{\sigma_y}{2}} = \begin{pmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}, \tag{1}$$

where eq. (1) represents the single qubit Y rotation, σ_y is the Pauli matrix, and U is a random circuit, this part was implemented using PennyLane [2].

The quantum filtering layer is not available for training; this layer can be pre-computed prior to the training of the network to reduce the computation time. Otherwise, the quantum convolutional filtering process can be executed during mini-batch training on demand. The experimental design of this architecture is tested and compared to a classic approach where the entire network is trainable, as we discuss next.

4 Experiments and Results

The experiments were carried out using a standard GPU-based system running Python 3 with PennyLane and TensorFlow libraries. We divided the experiment into two different ones with different datasets, MNIST and CIFAR-10. The objective of each experiment is to test the autoencoder's ability on a standard simple image dataset (MNIST) and a more complex, color-based one (CIFAR-10). We want to study the reconstruction ability of each model, inspect its latent space and its rate of convergence for both a quantum or a classic convolutional layer.

4.1 MNIST

The MNIST dataset [12] consists of grayscale images of size 28×28. The number of images used for training the AEs is $60,000$ and $10,000$ for testing. For this problem, we used the architecture described in [15]. The results of the reconstruction using

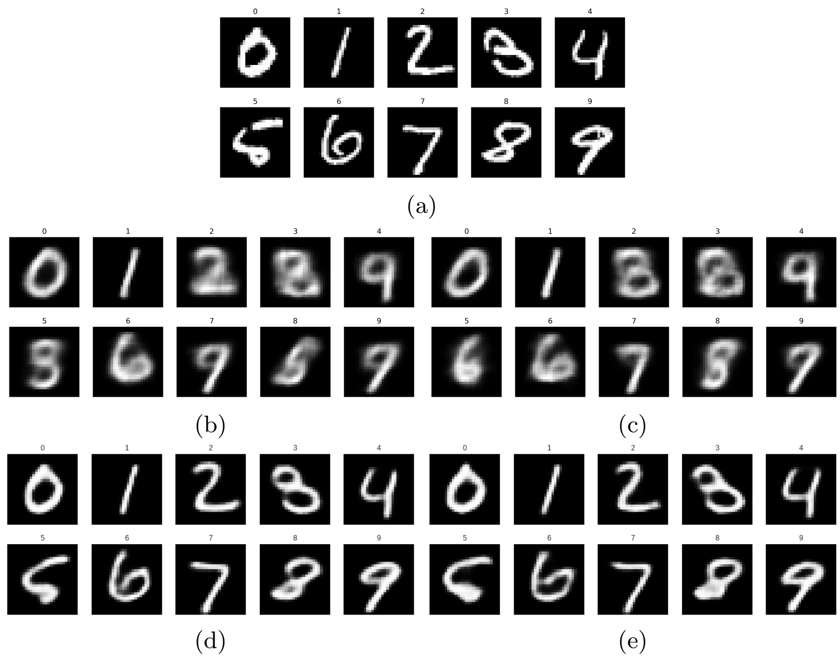

Fig. 4. MNIST data and reconstruction results from different latent spaces. (a) Original. (b) Classic Convolutional AE, $\mathbf{z} \in \mathbb{R}^2$. (c) Quanvolutional AE, $\mathbf{z} \in \mathbb{R}^2$. (d) Classic Convolutional AE, $\mathbf{z} \in \mathbb{R}^{64}$. (e) Quanvolutional AE, $\mathbf{z} \in \mathbb{R}^{64}$.

both the classic and quantum-based methodologies are shown in Fig. 4. In (a) we have a sample of the original digits in the test set, then in (b) we have the classic convolutional AE approach in reconstruction while the quantum-based one is shown in (c). Clearly, the reconstruction abilities of the network are comparable. This initial test goes down to a latent space in $\mathbf{z} \in \mathbb{R}^2$; therefore, the reconstruction results appear to have considerable room for improvement. However, a second experiment was conducted where the autoencoders were set to find a latent space in $\mathbf{z} \in \mathbb{R}^{64}$, which yields the reconstruction results in Fig. 4 (d) and (e) for the classic and quantum approach, respectively. Evidently, these results in this larger latent space are much better in terms of reconstruction and denoising abilities. A visual comparison between both methods suggests that they have comparable reconstruction performance.

Projecting the test set on the learned latent space is one interesting way of examining what representations are being learned. To this end, we present Fig. 5, which shows the comparison between (a) the classic approach and (b) the quantum-based approach. The differences here can be appreciated by noticing cluster separation, within class distributions, and cluster overlaps. The figure suggests that the learned spaces are rich and discriminative as they appear, and if they are further re-trained separately for classification, it might yield very good results, even if the latent space is only two-dimensions. However, for a

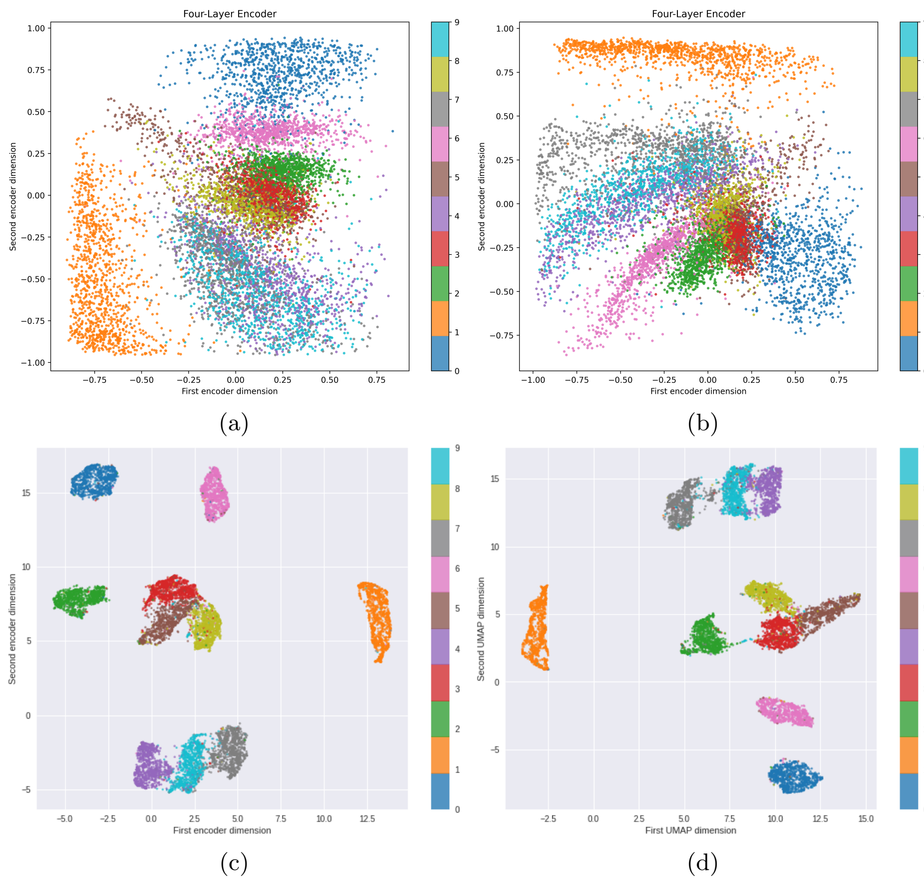

Fig. 5. Two-dimensional representation of the latent space. (a) and (c) display classic AE while (b) and (d) depict the quantum-based.

64-dimensional space there are clear improvements in regards to cluster features, as seen in Fig. 5 (c) and (d) for classic and quantum-based, respectively. Note that (c) and (d) are created with UMAP [13], which finds a two-dimensional representation of the $\mathbb{R}^{64}$ space.

Finally, another objective way to compare performances is to observe the average behavior of the learning across epochs. Fig. 6 shows the results of the two experiments with MNIST. In (a) we can compare the results using the two-dimensional latent space reconstructions. It shows that both methodologies are comparable and converge to a minimum without being statistically different, i.e., the standard deviation of both overlap most of the time. These results are produced with ten (10) randomized runs of the same experiment choosing random initial weights every time. However, in (b) we can observe that the gap in performance between the two methodologies is statistically significant. What we can observe here is that both methodologies converge to different minima

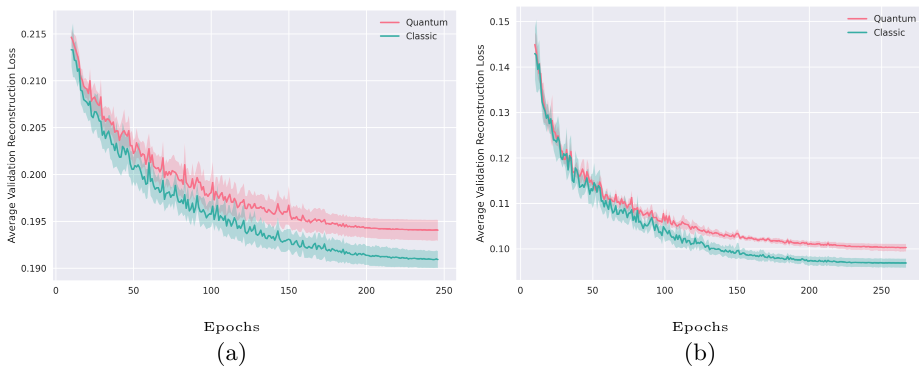

Fig. 6. Average loss during gradient descent over a reconstruction loss on MNIST data. (a) is for $\mathbf{z} \in \mathbb{R}^2$ (b) is for $\mathbf{z} \in \mathbb{R}^{64}$.

with nearly a 5×10^{-3} difference in magnitude of the loss. The loss function used in all of our experiments is the *binary cross entropy* loss with the *adam* optimizer [10].

4.2 CIFAR-10

The dataset known as CIFAR-10 [11] consists of 50,000 color images of size $32 \times 32 \times 3$ pixels. These images correspond to 10 different categories and are labeled; however, the labels are not used during training. A sample of these pictures is shown in Fig. 7 (a). The convolutional network for this dataset has the architecture shown in Table 1. From the table, we can see that the first four elements of the table are marked with †, indicating that those layers are the ones replaced with the quantum-based based methodology to replace traditional convolutions. These quantum convolutions start with the same single input image of $28 \times 28 \times 1$ and end up in multiple convolved images of size $7 \times 7 \times 16$. This makes the comparison between the classic convolutional AE approach and the quantum-based one, with the exception, of course, of the fact that in the classic approach, all convolutional layers are trainable. In contrast, the quantum-circuit convolutions, shown in Fig. 3, are fixed. These facts related to the architecture are depicted in Fig. 2.

Similar to the experiments performed over MNIST in the previous section, we executed the same tests of performance, obtaining the reconstruction results shown in Fig. 7. In (b) we can see the reconstruction results of the classic autoencoder, and in (c) we see the corresponding quantum-based implementation. The results are comparable in terms of a visual inspection. However, when we display the average behavior of the optimization process, we can encounter interesting results, see Fig. 8. From the figure, we observe that for an embedding space in $\mathbb{R}^2$, Fig. 8 (a), the loss function minimization is comparable for both models and appears unstable at earlier iterations. However, for the larger embedding space in $\mathbb{R}^{128}$, Fig. 8 (b), the quantum-based approach exhibits strong stability in earlier iterations

Table 1. Configuration of the classic convolutional neural network for CIFAR-10. The layers marked with † are the ones replaced by the quantum circuit convolutions. Note that for the dense layer marked with $*$ we also experimented with two neurons instead of 128 while the rest of the architecture remains the same.

Operation or Layer	Filters	Filter Size	Stride	Padding	Size of Output
†Input image	-	-	-	-	$32 \times 32 \times 3$
†Convolution †ReLU	48 -	4×4 -	1×1 -	same -	$32 \times 32 \times 46$ $32 \times 32 \times 48$
†Max pooling	-	4×4	4×4	same	$8 \times 8 \times 48$
Convolution ReLU	24 -	3×3 -	1×1 -	same -	$8 \times 8 \times 24$ $8 \times 8 \times 24$
Convolution ReLU	12 -	2×2 -	1×1 -	same -	$8 \times 8 \times 12$ $8 \times 8 \times 12$
Max pooling	-	2×2	2×2	same	$4 \times 4 \times 12$
Flatenning Dropout 20% Dense (linear)	- - -	- - -	- - -	- - -	192 192 192
*Dense (tanh)	-	-	-	-	128
Dense (linear) Reshape	- -	- -	- -	- -	192 $4 \times 4 \times 12$
Convolution ReLU	12 -	2×2 -	1×1 -	same -	$4 \times 4 \times 12$ $4 \times 4 \times 12$
Up Sampling	-	2×2	2×2	same	$8 \times 8 \times 12$
Convolution ReLU	24 -	3×3 -	1×1 -	same -	$8 \times 8 \times 24$ $8 \times 8 \times 24$
Convolution ReLU	48 -	1×1 -	1×1 -	same -	$8 \times 8 \times 48$ $8 \times 8 \times 48$
Up Sampling	-	4×4	2×2	same	$32 \times 32 \times 48$
Convolution Sigmoid	3 -	4×4 -	1×1 -	none -	$32 \times 32 \times 3$ $32 \times 32 \times 3$

in contrast to the classic approach. Both models eventually converge to a similar solution after a number of iterations.

5 Discussion and Conclusions

The experiments performed here show that the proposed work is comparable to the classic approach in terms of reconstruction ability and convergence. However, in the case of a general-purpose image dataset such as CIFAR-10, the quantum-based approach offered early learning stability in learning higher-dimensional latent spaces, in contrast to the classic approach that converges slowly at first.

This paper presented an AE application with a particular quantum implementation in the algorithm in the form of quantum convolutions. We call this implementation a Quanvolutional autoencoder. Our deployment shows the potential applications of quantum algorithms in learning image representations.

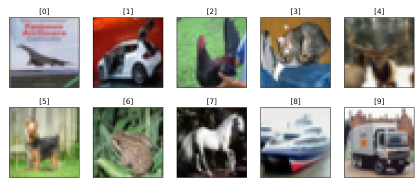

(a)

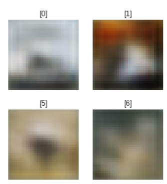

(b)

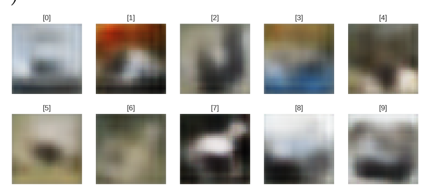

(c)

Fig. 7. CIFAR-10 dataset samples and reconstruction. (a) Original samples. (b) Images reconstructed with the classic approach. (c) Images reconstructed with the proposed quantum approach.

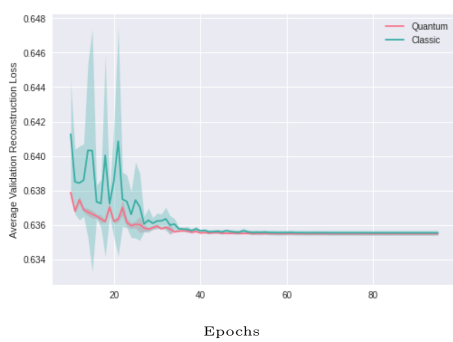

(a)

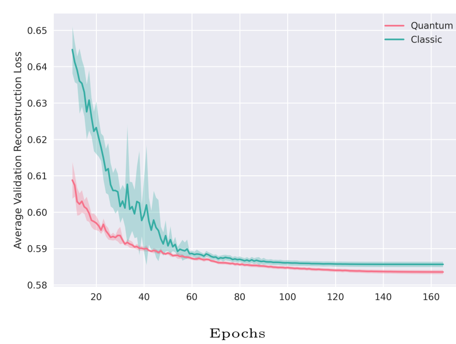

(b)

Fig. 8. Loss function minimization across different number of epochs. (a) is for a two-dimensional embedding space and (b) is for a 128-dimensional space.

Acknowledgements

The authors thank the Department of Computer Science at Baylor University for their support. This research was also funded, in part by the Baylor AI lab.

References

1. Aloysius, N., Geetha, M.: A review on deep convolutional neural networks. In: 2017 International Conference on Communication and Signal Processing (ICCSP), pp. 0588–0592. IEEE (2017)
2. Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., Alam, M.S., Ahmed, S., Arrazola, J.M., Blank, C., Delgado, A., Jahangiri, S., McKiernan, K., Meyer, J.J., Niu, Z.,

Száva, A., Killoran, N.: Pennylane: Automatic differentiation of hybrid quantum-classical computations (2020)
3. Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., Lloyd, S.: Quantum machine learning. Nature **549**(7671), 195–202 (2017)
4. Dong, G., Liao, G., Liu, H., Kuang, G.: A review of the autoencoder and its variants: A comparative perspective from target recognition in synthetic-aperture radar images. IEEE Geoscience and Remote Sensing Magazine **6**(3), 44–68 (2018)
5. Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press (2016). http://www.deeplearningbook.org
6. Griffiths, David J and Schroeter, Darrell F: Introduction to quantum mechanics. Cambridge University Press (2018)
7. Guo, X., Liu, X., Zhu, E., Yin, J.: Deep clustering with convolutional autoencoders. In: International conference on neural information processing, pp. 373–382. Springer (2017)
8. Henderson, M., Shakya, S., Pradhan, S., Cook, T.: Quanvolutional neural networks: Powering image recognition with quantum circuits (2019)
9. Henning, R., Rivas-Perea, P., Shaw, B., Hamerly, G.: A convolutional neural network approach for classifying leukocoria. In: Image Analysis and Interpretation (SSIAI), 2014 IEEE Southwest Symposium on, pp. 9–12 (2014)
10. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
11. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
12. LeCun, Y., Cortes, C., Burges, C.: Mnist handwritten digit database. ATT Labs [Online]. Available: http://yann.lecun.com/exdb/mnist **2** (2010)
13. McInnes, L., Healy, J., Melville, J.: Umap: Uniform manifold approximation and projection for dimension reduction. arXiv preprint arXiv:1802.03426 (2018)
14. Rai, M., Rivas, P.: A review of convolutional neural networks and gabor filters in object recognition. In: 2020 International Conference on Computational Science and Computational Intelligence, pp. 1–8 (2020)
15. Rivas, P., Orduz, J., Baker, E.: A quanvolutional autoencoder. In: Proc. of the IEEE/CVF International Conf. on Computer Vision, p. Submitted (2021)
16. Rivas, P., Rivas, E., Velarde, O., Gonzalez, S.: Deep sparse autoencoders for american sign language recognition using depth images. In: 21st International Conference on Artificial Intelligence (ICAI 2019) (2019)
17. Romero, J., Olson, J.P., Aspuru-Guzik, A.: Quantum autoencoders for efficient compression of quantum data. Quantum Science and Technology **2**(4), 045,001 (2017)
18. Sagha, H., Cummins, N., Schuller, B.: Stacked denoising autoencoders for sentiment analysis: a review. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery **7**(5), e1212 (2017)
19. Schuld, M.: Quantum machine learning models are kernel methods. arXiv preprint arXiv:2101.11020 (2021)
20. Schuld, Maria and Sinayskiy, Ilya and Petruccione, Francesco: An introduction to quantum machine learning. Contemporary Physics **56**(2), 172–185 (2014)
21. Theis, L., Shi, W., Cunningham, A., Huszár, F.: Lossy image compression with compressive autoencoders. arXiv preprint arXiv:1703.00395 (2017)
22. Voulodimos, A., Doulamis, N., Doulamis, A., Protopapadakis, E.: Deep learning for computer vision: A brief review. Computational intelligence and neuroscience **2018** (2018)

23. Wang, Y., Yao, H., Zhao, S.: Auto-encoder based dimensionality reduction. Neurocomputing **184**, 232–242 (2016)
24. Wetzel, S.J.: Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders. Physical Review E **96**(2), 022,140 (2017)
25. Wittek, Peter: Quantum machine learning: what quantum computing means to data mining. Academic Press (2014)
26. Yousefi-Azar, M., Hamey, L.: Text summarization using unsupervised deep learning. Expert Systems with Applications **68**, 93–105 (2017)
27. Zhang, S., Tong, H., Xu, J., Maciejewski, R.: Graph convolutional networks: a comprehensive review. Computational Social Networks **6**(1), 1–23 (2019)
28. Zhiqiang, W., Jun, L.: A review of object detection based on convolutional neural network. In: 36th Chinese Control Conference, pp. 11,104–11,109. IEEE (2017)