

Hybrid Quantum Variational Autoencoders for Representation Learning

Pablo Rivas , Senior, IEEE
 School of Eng. and Computer Science
 Dept. of Computer Science, Baylor University
 Email: Pablo_Rivas@Baylor.edu

Liang Zhao 
 Marketing Department
 St. Ambrose University
 Email: ZhaoLiang@sau.edu

Javier Orduz 
 School of Eng. and Computer Science
 Dept. of Computer Science, Baylor University
 Email: Javier_OrduzDucuara@Baylor.edu

Abstract—Representation learning is a standard area that has seen many improvements based on machine learning advances. Quantum machine learning advances are now spreading across different application areas such as representation learning. This paper introduces a novel hybrid quantum machine learning approach to representation learning by using a quantum variational circuit that is trainable with traditional gradient descent techniques. We use marketing data to showcase the learning potential of our model.

Index Terms—quantum machine learning, marketing, quantum variational circuits

I. INTRODUCTION

Creative neural network architectures are emerging consistently, aiming to solve interesting problems [1]–[3]. These kinds of models offer certain versatility when one focuses on learning representations [4]. Interestingly, unsupervised approaches tend to be preferred to remove any label bias that might be introduced and that may not be wanted. This research is based on a classic autoencoder architecture that is combined with a novel quantum variational approach.

An autoencoder (AE) is considered an unsupervised learning model that reconstructs the input signal using a neural network [5]. AEs are notably known for some of their successful versions, including the Variational Autoencoder (VAE) [6], and the denoising AE [7], [8]. Dense 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 [9].

There is a growing interest among marketing researchers to leverage machine learning in recent years. Compared with traditional statistical and econometric methods, machine learning methods can process large scaled data, unstructured data, have flexible model structures and yield better predictions. Autoencoder is beginning to be used to generate meaningful descriptions from complex data in marketing contexts such as consumer social networks or consumer-product networks [10]. The methodological challenge in analyzing large-scale network is the high dimensionality. A recent study on a user-brand network from user engagement data on Facebook used a deep autoencoder to perform embedding, and showed that a market structure of brand that is more fluid and overlapping than what standard industry classification would suggest [11]. In addition, a variational autoencoder has been developed to

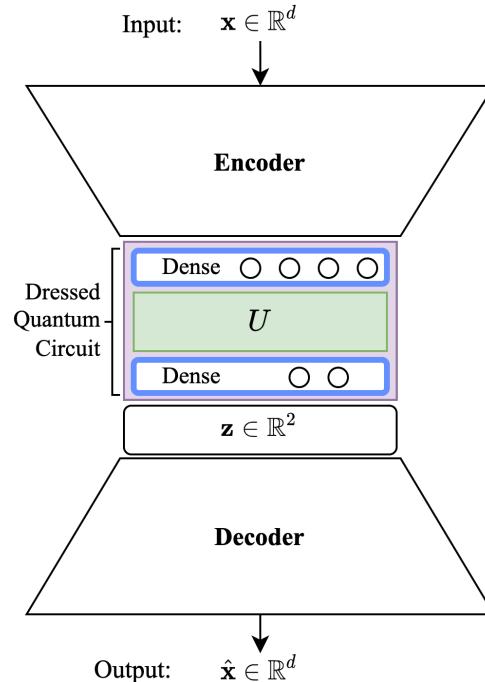


Fig. 1. Dressed quantum circuit making a hybrid autoencoder.

combine several data sources that contain both text and images to help generate new logo designs for marketers [12].

In this research, we focus on an AE for learning representations leveraging **quantum computing**, and introducing a trainable quantum variational circuit into the model. Fig. 1 depicts the circuit placement around dense layers, thereby dressing the quantum circuit to be compatible with the rest of the autoencoder architecture.

The rest of the paper is organized as follows. The associated background material is reviewed in Section II. In section III, we present our methodology, including a description of the dataset we used, a layout of the neural architecture we developed, and an overview and evaluation of the experiments we ran, including the results. Conclusions are presented in Section IV.

II. BACKGROUND

A. Related work in quantum machine learning

Much of the work related to quantum machine learning has been popularized in recent years. Some of the most notable efforts involve variational approaches [13]–[15]. Researchers have shown that these models are effective in complex tasks that grant further studies and open new doors for applied quantum machine learning research.

Another popular approach is to perform kernel learning using a quantum approach [16]–[18]. In this case the kernel-based projection of data \mathbf{x} produces a feasible linear mapping to the desired target y as follows:

$$y(\mathbf{x}) = \text{sign} \left(\sum_{j=1}^M \alpha_j k(\mathbf{x}_j, \mathbf{x}) + b \right) \quad (1)$$

for hyper parameters b, α that need to be provided or learned. This enables the creation of some types of support vector machines whose kernels are calculated such that the data \mathbf{x} is processed in the quantum realm. That is $|\mathbf{x}_j\rangle = 1/|\mathbf{x}_j| \sum_{k=1}^N (\mathbf{x}_j)_k |k\rangle$. The work of Schuld *et al.* [17], expands the theory behind this idea and show that all kernel methods can be quantum machine learning methods [19].

Recently, in 2020, Mari *et al.* [20], worked on variational models that are hybrid in format. Particularly, the authors focused on transfer learning, i.e., the idea of bringing a pre-trained model (or a piece of it) to be part of another model. In the case of [20] the larger model is a computer vision model, e.g., ResNet [21], which is part of a variational quantum circuit that performs classification. The work we present here follows a similar idea, but we focus in the autoencoder architecture, rather than a classification model, and we focus on learning representations in comparison between a classic and a variational quantum fine-tuned model.

B. Variational quantum circuits

Variational quantum circuits have been recently studied in combination with different models, including neural networks, support vector machines, and other linear classifiers [17], [22]. The authors in [20] define a quantum layer as a unitary operation, U , implemented as a variational circuit on an input state $|\hat{\mathbf{x}}\rangle$, that produces the output state $|y\rangle$ as follows:

$$|\hat{\mathbf{x}}\rangle \rightarrow |y\rangle = U(\mathbf{w})|\hat{\mathbf{x}}\rangle, \quad (2)$$

where \mathbf{w} denotes the parameters of the variational circuit. Fig. 2 depicts this idea.

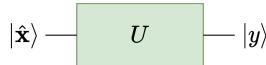


Fig. 2. Quantum layer of a circuit seen as a unitary operation.

A quantum circuit can be composed of entangling gates, qubit rotations, and others [23], [24], which together compose a quantum layer that can preserve the Hilbert-space dimension

of the input states $\hat{\mathbf{x}}$ [20]. Our work follows the paradigm of creating a variational circuit with quantum layers, as we discuss next.

III. METHODOLOGY

In this section, we present our methodology, including a description of the dataset we used, a layout of the neural architecture we developed, and an overview and evaluation of the experiments we ran, including the results.

A. Classic Neural Autoencoder

A classic autoencoder can be broken down into two major components that serve specific purposes during an unsupervised learning process. Fig. 3 shows an autoencoder that is implemented using fully connected (dense) layers and other traditional layers that act as regularizers and reduce overfitting, e.g., dropout, batch normalization. It receives as input some vector, $\mathbf{x} \in \mathbb{R}^d$, and then it goes into a number layers, which are meant to compress the input data down to \hat{d} dimensions. This first set of layers is known as the encoder. The second set of neurons is meant to reconstruct the input data back to its original dimensionality and values $\hat{\mathbf{x}} \in \mathbb{R}^d$ using a similar number of layers; this group of layers is known as the decoder.

In this case, the classic autoencoder shown in Fig. 3 acts as a compression network, in the sense that after training the model to achieve good reconstruction, if we disconnect the decoder, we end up with a neural network that encodes the input data into \hat{d} dimensions (we will use two for visualization purposes). This presents a unique advantage over supervised models: in a supervised model, we teach a network to look for a pattern that will permit an association with a given target label; however, in unsupervised learning (or in this autoencoder, for example), the network does not look for a specific pattern but rather learns to use the input space in any way that preserves the most representative and most important information of the input data, so as to allow good reconstruction in the decoder.

We train our classic autoencoder to minimize an L_1 loss, i.e., mean absolute error loss, using an RMSProp optimizer with an initial learning rate of 0.001, reducing a 1% every 10 epochs, for 500 epochs.

B. Quantum Hybrid Autoencoder

The quantum hybrid autoencoder is based on the variational circuit shown in Fig. 4. The figure has many gates and operators in the form of quantum layers. In the next paragraphs we discuss briefly the meaning of these.

1) *Hadamard operators layer*: The Hadamard operator on a qubit is denoted as:

$$H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \quad (3)$$

Its primary purpose is to create superposition. This is one of the most relevant concepts coming from quantum mechanics, implemented in quantum computing. Superposition and other concepts taken from quantum mechanics are what make quantum information work to our advantage.

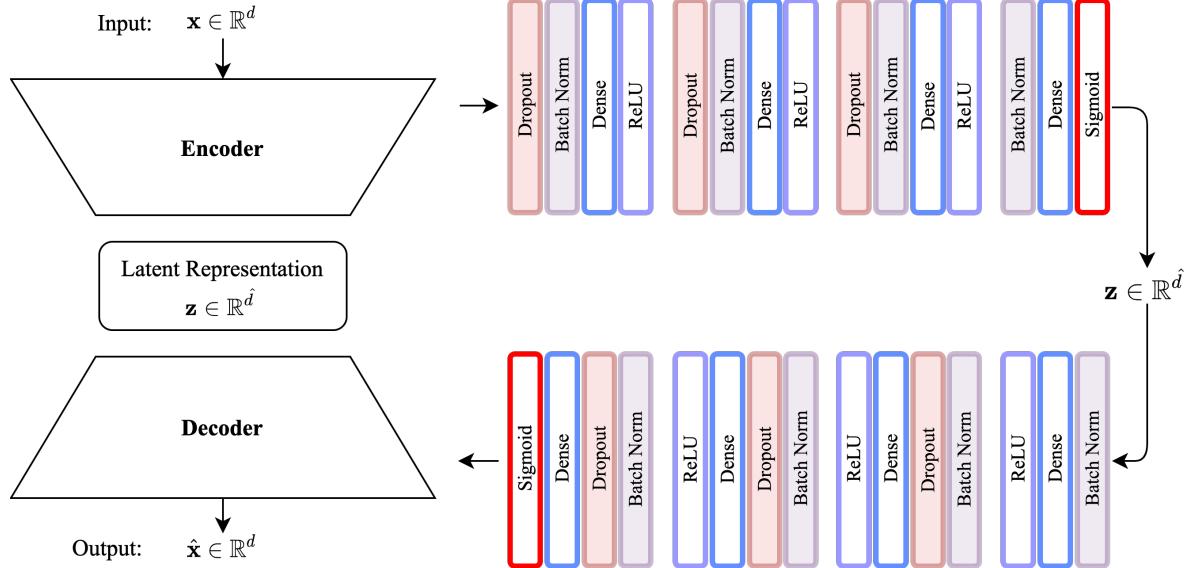


Fig. 3. Two representations of our autoencoder. Left: compact and abstracted model representation. Right: full and layer-descriptive model.

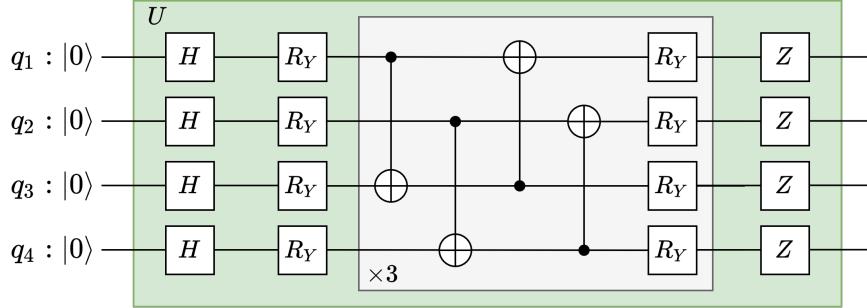


Fig. 4. Quantum variational circuit. The angle of the rotations on Y are fully trainable. The Pauli Z operator provides the means for the measurement of the expected outputs. The entire circuit is denoted as U and has three layers in our study.

2) *Single qubit Y rotation layer*: A rotation of a qubit makes a qubit change the spin based on the rotation angle, ϕ , as follows:

$$R_Y(\phi) = e^{-i\phi\sigma_Y/2} = \begin{bmatrix} \cos(\phi/2) & -\sin(\phi/2) \\ \sin(\phi/2) & \cos(\phi/2) \end{bmatrix}. \quad (4)$$

The rotation angle ϕ in our research is a trainable parameter.

3) *CNOT qubit entangling layer*: The CNOT operation, defined as follows:

$$\text{CNOT} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad (5)$$

is aimed at linking qubits, combining them and propagating superposition across layers.

4) *Expectation layer over Pauli Z operators*: Finally, the output of the circuit is a measurement that is calculated over many observations returning the expected value. In our case

the measurements are applied after the Pauli Z operator defined as follows:

$$\sigma_z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (6)$$

The output of our circuit is considered to be produced after measuring the Pauli Z operator on every qubit.

5) *Quantum dressed circuit*: Once the circuit is defined, it needs to be dressed up to be combined with the classic autoencoder model. This process, shown in [20], involves a simple process that adds a single dense layer before and after the quantum circuit, as shown in Fig. 1. The number of neurons in the input layer is set to match the number of qubits, in this study is four. The last layer has two neurons, as set to match our visualization intent in two dimensions. This is because we are interested in inspecting the latent space in two dimensions. However, this can be arbitrarily set to any latent space dimension as desired.

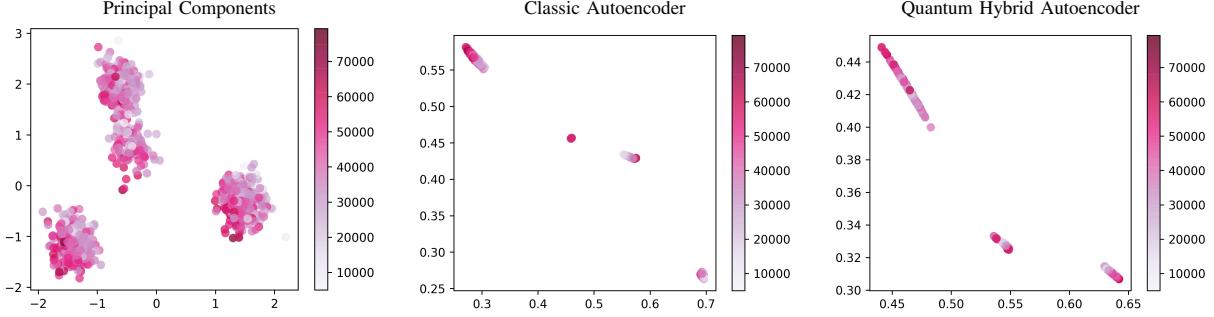


Fig. 5. Data visualization across different models, where the color indicates revenue returns. Left: PCA is used to show that there are three distinctive clusters. Center: AE is used in its classic form to show that there are four clusters of information. Right: The quantum hybrid approach also shows four distinctive clusters, although one of them is more prominent than the rest, see top left.

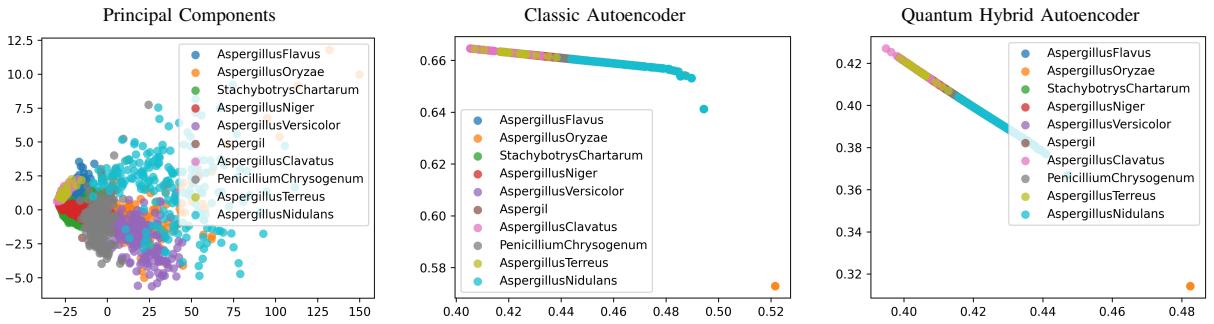


Fig. 6. Toxic mold dataset [25], [26]. The different molecular spectral signatures are learned in a discriminatory fashion that promotes simple linear separability using the proposed approach. In contrast, PCA offers a subspace that does not facilitate separability.

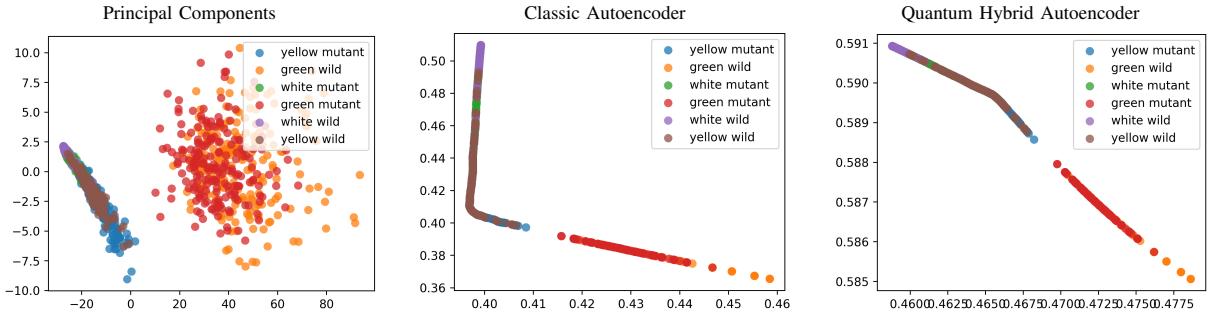


Fig. 7. Molecular origin dataset [27], [28]. PCA is able to capture two distinctive clusters of data although within those clusters separability can be difficult. On the other hand our proposed approach can provide representations that are suitable for other supervised machine learning tasks such as classification.

C. Marketing Dataset

The dataset used in our study comes from a publicly available Kaggle competition: <https://www.kaggle.com/irisfanfan/store-data>. The purpose of the dataset is to facilitate finding the media that has the greatest impact on sales among four different marketing efforts.

The authors of the case study initially wanted to facilitate the study of relationship between revenue and marketing investments. The dataset consists of six different attributes, with revenue being the target variable for study. Four of the attributes define the investment across the different media, i.e.,

local television, online, in store, and in person. It is expected to see some relationship to such four major groups in our visual analysis.

When deciding how to allocate promotional expenditures among local TV advertising, online advertising, in store displays, personal sales, and promotional events, marketing managers could improve the resource allocation efficiency by allocating more investment on clusters that are associated with higher revenue returns. Marketers could also manage promotional channels more effectively based on the unit of cluster, other than on the unit of specific channel.

D. Raman Spectroscopy Datasets

Our model was also tested on two datasets that were acquired using Raman spectroscopy [25], [27].

Han et al. [27], used shifted excitation Raman difference spectroscopy (SERDS), where a light source generates two spectra with different frequencies to eliminate the effects of fluorescence. The authors combine SERDS with the genetic breeding of mutant populations to gather information relevant to the pigment of molecules within cell walls.

Similarly, Strycker et al. [25], SERDS to measure the Raman spectra of 10 toxic mold species. Authors show that Raman signals originate from the melanin pigments bound within cell walls. In both of these datasets, the data is available online for research [26], [28].

E. Experiments, Evaluation, and Results

The model in Fig. 3 and Fig. 1 was implemented in Python using the PyTorch platform with Keras libraries and PennyLane. The training discussed below was performed on an NVIDIA P100 GPU system with 25 GB of RAM and 166 GB of storage.

We performed two major experiments: 1) one that visualizes the data using the classic autoencoder with the purpose of validating the learned representations; and 2) a quantum hybrid approach.

Both of our models optimized using gradient descent (RMSprop optimizer [29], with learning rate of 0.001) over the loss function named mean absolute error (MAE) loss:

$$L(y, \hat{y}) = \frac{1}{N} \sum_{i=0}^N |y - \hat{y}_i| \quad (7)$$

where N is the number of samples, $\hat{x} \in \mathbb{R}^d$ is the reconstructed data. The neurons in this output layer uses a sigmoid activation.

The results shown in Fig. 5 are including a classic benchmark known as principal component analysis which is used for dimensionality reduction with the standard settings.

From Fig. 5 on the left, we can observe three well-defined isotropic clusters; however, one of those is actually two clusters with high overlap.

From Fig. 5 on the center, we can see the results of applying the classic autoencoder to the same dataset producing four well-defined clusters. However, the proposed hybrid quantum approach leads to very similar results than the classic except for one cluster that appears to be lingering along the top left diagonal. Our initial exploration of the data suggests that such cluster of data belongs to data that produces the widest variety of uncertainty in the results related to revenue.

However, for our proposed approach as well as the classic one, cluster are concentrated, compressed within themselves, and have larger distances in relation to other clusters. This leads to excellent potential for learning other types of problems.

The Raman shift datasets depicted in Fig. 6 and Fig. 7 also exhibit similar characteristics of good representation learning. In Fig. 6 the latent space found using our hybrid approach can project the dataset and concentrate it along a line. The line can show how the different categories of toxic mold

molecules may have spectral similarities or simply classify the data automatically. For the second Raman shift-based molecular dataset, shown in Fig. 7, the difference between a classic approach to dimensionality reduction like PCA and our method is more palpable. While PCA offers a widespread data projection subspace, the data within itself is hardly separable between all categories. However, from visual inspection, our proposed method appears to facilitate easier class separation along with relatively linear space.

IV. CONCLUSIONS

We have presented a quantum hybrid approach to perform unsupervised dimensionality reduction of data. We use a quantum variational circuit whose parameters are trainable with traditional gradient descent techniques. The quantum circuit is dressed into a neural architecture based on a classic dense autoencoder.

Our preliminary results suggest that the model are capable of finding latent spaces, and learning representations that have high discriminative potential. Data from marketing shows that our model is comparable to a classic approach and better than PCA in finding clusters of data. Huge amount of social media and mobile usage data are being generated in an unstructured format, such as text, images, audio, video, network, or consumer online tracking data [30]. The hybrid quantum variational autoencoder can be beneficial to generate low-dimensional representations of the big data to facilitate feature extraction analysis. The hybrid quantum variational autoencoder could also be used in descriptive interpretation to generate meaningful descriptions from complex unstructured data [10]. Many important marketing research issues would be aided by using hybrid quantum variational autoencoder, and there are ample opportunities for new interdisciplinary research in the future.

Future work includes more datasets to showcase the potential of these kind of models.

ACKNOWLEDGMENT

The machine learning model presented here is based upon work supported in part by the National Science Foundation under Grant CHE-1905043.

REFERENCES

- [1] A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, "Deep learning for computer vision: A brief review," *Computational intelligence and neuroscience*, vol. 2018, 2018.
- [2] N. Aloysius and M. Geetha, "A review on deep convolutional neural networks," in *2017 International Conference on Communication and Signal Processing (ICCSP)*. IEEE, 2017, pp. 0588–0592.
- [3] R. Henning, P. Rivas-Perea, B. Shaw, and G. Hamerly, "A convolutional neural network approach for classifying leukocoria," in *Image Analysis and Interpretation (SSIAI), 2014 IEEE Southwest Symposium on*, April 2014, pp. 9–12.
- [4] M. Rai and P. Rivas, "A review of convolutional neural networks and gabor filters in object recognition," in *2020 International Conference on Computational Science and Computational Intelligence*, 2020, pp. 1–8.
- [5] G. Dong, G. Liao, H. Liu, and G. Kuang, "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*, vol. 6, no. 3, pp. 44–68, 2018.

[6] S. J. Wetzel, "Unsupervised learning of phase transitions: From principal component analysis to variational autoencoders," *Physical Review E*, vol. 96, no. 2, p. 022140, 2017.

[7] H. Sagha, N. Cummins, and B. Schuller, "Stacked denoising autoencoders for sentiment analysis: a review," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 7, no. 5, p. e1212, 2017.

[8] P. Rivas, E. Rivas, O. Velarde, and S. Gonzalez, "Deep sparse autoencoders for american sign language recognition using depth images," in *21st International Conference on Artificial Intelligence (ICAI 2019)*, 2019.

[9] X. Guo, X. Liu, E. Zhu, and J. Yin, "Deep clustering with convolutional autoencoders," in *International conference on neural information processing*. Springer, 2017, pp. 373–382.

[10] L. Ma and B. Sun, "Machine learning and ai in marketing—connecting computing power to human insights," *International Journal of Research in Marketing*, vol. 37, no. 3, pp. 481–504, 2020.

[11] Y. Yang, K. Zhang, and P. Kannan, "Identifying market structure: A deep network representation learning of social engagement," working paper, Tech. Rep., 2019.

[12] R. Dew, A. Ansari, and O. Toubia, "Letting logos speak: Leveraging multiview representation learning for data-driven logo design," *Available at SSRN*, vol. 3406857, 2019.

[13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, "Variational quantum algorithms," *Nature Reviews Physics*, pp. 1–20, 2021.

[14] A. Khoshman, W. Vinci, B. Denis, E. Andriyash, H. Sadeghi, and M. H. Amin, "Quantum variational autoencoder," *Quantum Science and Technology*, vol. 4, no. 1, p. 014001, 2018.

[15] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, "Theory of variational quantum simulation," *Quantum*, vol. 3, p. 191, 2019.

[16] C. Blank, D. K. Park, J.-K. K. Rhee, and F. Petruccione, "Quantum classifier with tailored quantum kernel," *npj Quantum Information*, vol. 6, no. 1, pp. 1–7, 2020.

[17] M. Schuld and N. Killoran, "Quantum machine learning in feature hilbert spaces," *Physical review letters*, vol. 122, no. 4, p. 040504, 2019.

[18] P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantum support vector machine for big data classification," *Physical review letters*, vol. 113, no. 13, p. 130503, 2014.

[19] M. Schuld, "Quantum machine learning models are kernel methods," *arXiv e-prints*, pp. arXiv–2101, 2021.

[20] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, "Transfer learning in hybrid classical-quantum neural networks," *Quantum*, vol. 4, p. 340, 2020.

[21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[22] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, "The theory of variational hybrid quantum-classical algorithms," *New Journal of Physics*, vol. 18, no. 2, p. 023023, 2016.

[23] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, "Circuit-centric quantum classifiers," *Physical Review A*, vol. 101, no. 3, p. 032308, 2020.

[24] N. Killoran, T. R. Bromley, J. M. Arrazola, M. Schuld, N. Quesada, and S. Lloyd, "Continuous-variable quantum neural networks," *Physical Review Research*, vol. 1, no. 3, p. 033063, 2019.

[25] B. D. Strycker, Z. Han, Z. Duan, B. Commer, K. Wang, B. D. Shaw, A. V. Sokolov, and M. O. Scully, "Identification of toxic mold species through raman spectroscopy of fungal conidia," *PloS one*, vol. 15, no. 11, p. e0242361, 2020.

[26] B. Strycker, Z. Han, Z. Duan, B. Commer, K. Wang, B. Shaw, A. Sokolov, and M. Scully, "Identification of toxic mold species through raman spectroscopy of fungal conidia," Oct 2020. [Online]. Available: osf.io/x6t5w

[27] Z. Han, B. D. Strycker, B. Commer, K. Wang, B. D. Shaw, M. O. Scully, and A. V. Sokolov, "Molecular origin of the raman signal from aspergillus nidulans conidia and observation of fluorescence vibrational structure at room temperature," *Scientific reports*, vol. 10, no. 1, pp. 1–8, 2020.

[28] Z. Han, B. D. Strycker, B. Commer, K. Wang, B. D. Shaw, M. O. Scully, and A. V. Sokolov, "Molecular origin of the raman signal from aspergillus nidulans conidia and observation of fluorescence vibrational structure at room temperature," Aug 2021. [Online]. Available: osf.io/bxd37

[29] M. C. Mukkamala and M. Hein, "Variants of rmsprop and adagrad with logarithmic regret bounds," in *International Conference on Machine Learning*. PMLR, 2017, pp. 2545–2553.

[30] L. Hagen, K. Uetake, N. Yang, B. Bollinger, A. J. Chaney, D. Dzyabura, J. Etkin, A. Goldfarb, L. Liu, K. Sudhir *et al.*, "How can machine learning aid behavioral marketing research?" *Marketing Letters*, vol. 31, no. 4, pp. 361–370, 2020.