

Transfer learning in hybrid classical-quantum neural networks

Andrea Mari, Thomas R. Bromley, Josh Izaac, Maria Schuld, and Nathan Killoran

Xanadu, 777 Bay Street, Toronto, Ontario, Canada.

We extend the concept of transfer learning, widely applied in modern machine learning algorithms, to the emerging context of hybrid neural networks composed of classical and quantum elements. We propose different implementations of hybrid transfer learning, but we focus mainly on the paradigm in which a pre-trained classical network is modified and augmented by a final variational quantum circuit. This approach is particularly attractive in the current era of intermediate-scale quantum technology since it allows to optimally pre-process high dimensional data (e.g., images) with any state-of-the-art classical network and to embed a select set of highly informative features into a quantum processor. We present several proof-of-concept examples of the convenient application of quantum transfer learning for image recognition and quantum state classification. We use the cross-platform software library PennyLane to experimentally test a high-resolution image classifier with two different quantum computers, respectively provided by IBM and Rigetti.

1 Introduction

Transfer learning is a typical example of an artificial intelligence technique that has been originally inspired by biological intelligence. It originates from the simple observation that the knowledge acquired in a specific context can be transferred to a different area. For example, when we learn a second language we do not start from scratch, but we make use of our previous linguistic knowledge. Sometimes transfer learning is the only way to approach complex cognitive tasks, e.g., before learning quantum mechanics it is advisable to first study linear algebra. This general idea has been successfully applied also to design artificial neural networks [31, 36, 47]. It has been shown [38, 51] that in many situations, instead of training a full network from scratch, it is more efficient to start from a pre-trained deep network and then optimize only some of the final layers for a particular task and dataset of interest (see Fig. 1).

The aim of this work is to investigate the potential of the transfer learning paradigm in the context of quantum machine learning [10, 15, 40]. It is not in

Andrea Mari: andreamari84@gmail.com

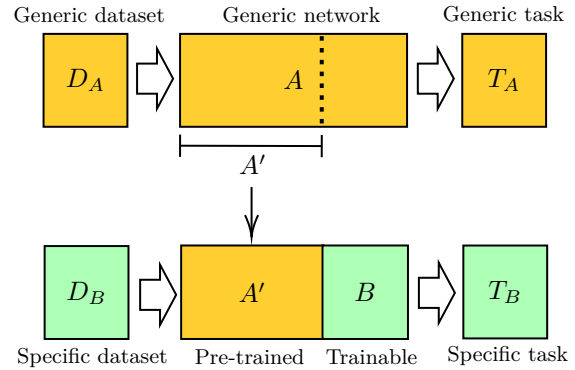


Figure 1: General representation of the *transfer learning* method, where each of the neural networks A and B can be either classical or quantum. Network A is pre-trained on a dataset D_A and for a task T_A . A reduced network A' , obtained by removing some of the final layers of A , is used as a fixed feature extractor. The second network B , usually much smaller than A' , is optimized on the specific dataset D_B and for the specific task T_B .

the scope of this study to make claims about the performance of quantum models with respect to classical models, but just to demonstrate that quantum transfer learning is feasible. We focus on hybrid models [17, 30, 39], i.e., the scenario in which quantum variational circuits [24, 30, 33, 34, 41, 43] and classical neural networks can be jointly trained to accomplish hard computational tasks. In this setting, in addition to the standard *classical-to-classical* (CC) transfer learning strategy in which some pre-acquired knowledge is transferred between classical networks, three new variants of transfer learning naturally emerge: *classical to quantum* (CQ), *quantum to classical* (QC) and *quantum to quantum* (QQ).

In the current era of Noisy Intermediate-Scale Quantum (NISQ) devices [37], CQ transfer learning is particularly appealing since it opens the possibility to classically pre-process large input samples (e.g., high resolution images) with any state-of-the-art deep neural network and to successively manipulate few but highly informative features with a variational quantum circuit. This scheme is quite convenient since it makes use of the power of quantum computers, combined with the successful and well-tested methods of classical machine learning. On the other hand, QC and QQ transfer learning might also be very interesting approaches especially once large quantum computers will be available. In this case, fixed quantum

circuits might be pre-trained as generic quantum feature extractors, mimicking well known classical models which are often used as pre-trained blocks: e.g., AlexNet [28], ResNet [20], Inception [46], VGGNet [44], etc. (for image processing), or ULMFiT [22], Transformer [48], BERT [14], etc. (for natural language processing). In summary, such classical state-of-the-art deep networks can either be *used* in CC and CQ transfer learning or *replaced* by quantum circuits in the QC and QQ variants of the same technique.

Up to now, the transfer learning approach has been largely unexplored in the quantum domain with the exception of a few interesting applications, for example, in modeling many-body quantum systems [12, 23, 52], in the connection of a classical autoencoder to a quantum Boltzmann machine [35] and in the initialization of variational quantum networks [49]. With the present work we aim at developing a more general and systematic theory, specifically tailored to the emerging paradigms of variational quantum circuits and hybrid neural networks.

For all the models theoretically proposed in this work, proof-of-principle examples of practical implementations are presented and numerically simulated. Moreover we also experimentally tested one of our models on physical quantum processors—*ibmqx4* by IBM and *Aspen-4-Q-A* by Rigetti—demonstrating the successful classification of high resolution images with a hybrid classical-quantum system.

2 Hybrid classical-quantum networks

Before presenting the main ideas of this work, we begin by reviewing basic concepts of hybrid networks and introduce some notation.

2.1 Classical neural networks

A very successful model in classical machine learning is that of deep feed-forward neural networks [18]. The elementary block of a deep network is called a *layer* and maps input vectors of n_0 real elements to output vectors of n_1 real elements. Its typical structure consists of an affine operation followed by a non-linear function φ applied element-wise,

$$L_{n_0 \rightarrow n_1} : \mathbf{x} \rightarrow \mathbf{y} = \varphi(W\mathbf{x} + \mathbf{b}). \quad (1)$$

Here, the subscript $n_0 \rightarrow n_1$ indicates the number of input and output variables, \mathbf{x} and \mathbf{y} are the input and output vectors, W is an $n_1 \times n_0$ matrix and \mathbf{b} is a constant vector of n_1 elements. The elements of W and \mathbf{b} are arbitrary real parameters (respectively known as *weights* and *baises*) which are supposed be *trained*, i.e., optimized for a particular task. The nonlinear function φ is quite arbitrary but common choices are the *hyperbolic tangent* or the *rectified linear unit* defined as $\text{ReLU}(x) = \max(0, x)$.

A classical deep neural network is the concatenation of many layers, in which the output of the first is the input of the second and so on:

$$C = L_{n_{d-1} \rightarrow n_d} \circ \dots \circ L_{n_1 \rightarrow n_2} \circ L_{n_0 \rightarrow n_1}, \quad (2)$$

where different layers have different weights. Characteristic hyper-parameters of a deep network are its *depth* d (number of layers) and the number of *features* (number of variables) for each layer, i.e., the sequence of integers $n_0, n_1 \dots n_{d-1}$.

2.2 Variational quantum circuits

One of the possible quantum generalizations of feed-forward neural networks can be given in terms of *variational quantum circuits* [17, 24, 30, 30, 33, 34, 39, 41, 43]. Following the analogy with the classical case, one can define a quantum layer as a unitary operation which can be physically realized by a low-depth variational circuit acting on the input state $|x\rangle$ of n_q quantum subsystems (e.g., qubits or continuous variable modes) and producing the output state $|y\rangle$:

$$\mathcal{L} : |x\rangle \rightarrow |y\rangle = U(\mathbf{w})|x\rangle, \quad (3)$$

where \mathbf{w} is an array of classical variational parameters. Examples of quantum layers could be: a sequence of single-qubit rotations followed by a fixed sequence of entangling gates [41, 43] or, for the case of optical modes, some active and passive Gaussian operations followed by single-mode non-Gaussian gates [24]. Notice that, differently from a classical layer, a quantum layer preserves the Hilbert-space dimension of the input states. This fact is due to the fundamental unitary nature of quantum mechanics and, as discussed at the end of this section, should be taken into account when designing quantum networks.

A variational quantum circuit of depth q is a concatenation of many quantum layers, corresponding to the product of many unitaries parametrized by different weights:

$$\mathcal{Q} = \mathcal{L}_q \circ \dots \circ \mathcal{L}_2 \circ \mathcal{L}_1. \quad (4)$$

In order to inject classical data in a quantum network we need to embed a real vector \mathbf{x} into a quantum state $|x\rangle$. This can also be done by a variational embedding layer depending on \mathbf{x} and applied to some reference state (e.g., the vacuum or ground state),

$$\mathcal{E} : \mathbf{x} \rightarrow |x\rangle = E(\mathbf{x})|0\rangle. \quad (5)$$

Typical examples are single-qubit rotations or single-mode displacements parametrized by \mathbf{x} . Notice that, differently from \mathcal{L} , the embedding layer \mathcal{E} is a map from a classical vector space to a quantum Hilbert space.

Conversely, the extraction of a classical output vector \mathbf{y} from the quantum circuit can be obtained by

measuring the expectation values of n_q local observables $\hat{\mathbf{y}} = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_{n_q}]$. We can define this process as a measurement layer, mapping a quantum state to a classical vector:

$$\mathcal{M} : |x\rangle \rightarrow \mathbf{y} = \langle x | \hat{\mathbf{y}} | x \rangle. \quad (6)$$

Globally, the full quantum network including the initial embedding layer and the final measurement can be written as

$$Q = \mathcal{M} \circ \mathcal{Q} \circ \mathcal{E}. \quad (7)$$

The full network is a map from a classical vector space to a classical vector space depending on classical weights. Therefore, even though it may contain a quantum computation hidden in the quantum circuit, if considered from a global point of view, Q is simply a black-box analogous to the classical deep network defined in Eq. (2).

However, especially when dealing with real NISQ devices, there are technical limitations and physical constraints which should be taken into account: while in the classical feed-forward network of Eq. (2) we have complete freedom in the choice of the number of features for each layer; in the quantum network of Eq. (7) all these numbers are often linked to the size of the physical system. For example, even if not strictly necessary, typical variational embedding layers encode each classical element of \mathbf{x} into a single subsystem and so, in many practical situations, one has:

$$\#\text{inputs} = \#\text{subsystems} = \#\text{outputs}. \quad (8)$$

This common constraint of a variational quantum network could be overcome by:

1. adding ancillary subsystems and discarding/measuring some of them in the middle of the circuit;
2. engineering more complex embedding and measuring layers;
3. adding pre-processing and post-processing classical layers.

In this work, mainly because of its technical simplicity, we choose the third option and we formalize it through the notion of *dressed* quantum circuits introduced in the next subsection.

2.3 Dressed quantum circuits

In order to apply transfer learning at the classical-quantum interface, we need to connect classical neural networks to quantum variational circuits. Since in general the size of the classical and quantum networks can be very different, it is convenient to use a more flexible model of quantum circuits.

Let us consider the variational circuit defined in Eq. (7) and based on n_q subsystems. With the aim of adding some basic pre-processing and post-processing of the input and output data we place a classical layer at the beginning and at the end of the quantum network, obtaining what we might call a *dressed* quantum circuit:

$$\tilde{Q} = L_{n_q \rightarrow n_{\text{out}}} \circ Q \circ L_{n_{\text{in}} \rightarrow n_q}. \quad (9)$$

where $L_{n \rightarrow n'}$ is given in Eq. (1) and Q is the associated *bare* quantum circuit defined in Eq. (7). Differently from a complex hybrid network in which the computation is shared between cooperating classical and quantum processors, in this case the main computation is performed by the quantum circuit Q , while the classical layers are mainly responsible for the data embedding and readout. A similar hybrid model was applied to a generative quantum Helmholtz machine in Ref. [7] and to a variational circuit in Ref. [5].

We can say that from a hardware point of view a dressed quantum circuit is almost equivalent to a bare one. On the other hand, it has two important advantages:

1. the two classical layers can be trained to optimally perform the embedding of the input data and the post-processing of the measurement results;
2. the number of input and output variables are independent from the number of subsystems, allowing for flexible connections to other classical or quantum networks.

Even if our main motivation for introducing the notion of dressed quantum circuits is a smoother implementation of transfer learning schemes, this is also a quite powerful machine learning model in itself and constitutes a non-trivial contribution of this work. In the Examples section, a dressed quantum circuit is successfully applied to the classification of a non-linear benchmark dataset (2D spirals).

3 Transfer learning

In this section we discuss the main topic of this work, i.e., the idea of transferring some pre-acquired “knowledge” between two networks, say from network A to network B , where each of them could be either classical or quantum.

As discussed in the previous section, if considered as a black box, the global structure of a quantum variational circuit is similar to that of a classical network (see Eqs. (7), (9) and (2)). For this reason, we are going to define the transfer learning scheme in terms of two generic networks A and B , independently from their classical or quantum physical nature.

A	B	Transf. learn. scheme
Classical	Classical	CC ([31, 36, 38, 47, 51])
Classical	Quantum	CQ (Examples 2 and 3)
Quantum	Classical	QC (Example 4)
Quantum	Quantum	QQ (Example 5)

Table 1: Transfer learning schemes in hybrid classical-quantum networks. The corresponding proof-of-principle examples are indicated in the table and presented in Section 4.

Generic transfer learning scheme (see Fig. 1):

1. Take a network A that has been pre-trained on a dataset D_A and for a given task T_A .
2. Remove some of the final layers. In this way, the resulting truncated network A' can be used as a feature extractor.
3. Connect a new trainable network B at the end of the pre-trained network A' .
4. Keep the weights of A' constant, and train the final block B with a new dataset D_B and for a new task of interest T_B .

Following the common convention used in classical machine learning [31, 36, 38, 47, 51], all situations in which there is a change of dataset $D_B \neq D_A$ and/or a change of the final task $T_B \neq T_A$ can be identified as *transfer learning* methods. The general intuition behind this training approach is that, even if A has been optimized for a specific problem it can still act as a convenient feature extractor also for a different problem.

The concept of feature extraction is much more general than what is actually needed in this work (see, e.g., Ref. [8] for a good introduction to this topic). In the context of transfer learning, the truncated pre-trained network A' is interpreted as a feature extractor because, after removing some of the final layers of A (step 2), A' can be used to produce features which are not problem-specific. The motivation for truncating the final layers of A is that the final activations of a network are usually more tuned to the specific problem, while intermediate features are more generic and so more suitable for transfer learning. For example, in a convolutional neural network, a feature extractor can be obtained by removing the final dense layers while preserving all the pre-trained convolution layers.

In our hybrid setting, the fact that the networks A and B can be either classical or quantum gives rise to a rich variety of hybrid transfer learning models summarized in Table 1. For the reader familiar with quantum communication theory, this kind of classification might look similar to that of hybrid channels in which information can be exchanged between quantum and

classical systems. Here however there is a fundamental difference: what is actually transferred in this case is not raw information but some more structured and organized *learned representations*. We expect that the problem of transferring structured knowledge between systems governed by different physical laws (classical/quantum) could stimulate many interesting foundational and philosophical questions. The aim of the present work is however much more pragmatic and consists of studying practical applications of this idea.

3.1 Classical to quantum transfer learning

As discussed in the introduction, the CQ transfer learning approach is perhaps the most appealing one in the current technological era of NISQ devices. Indeed today we are in a situation in which intermediate-scale quantum computers are approaching the quantum supremacy milestone [6, 19] and, at the same time, we have at our disposal the very successful and well-tested tools of classical deep learning. The latter are universally recognized as the best-performing machine learning algorithms, especially for image and text processing.

In this classical field, transfer learning is already a very common approach, thanks to the large zoo of pre-trained deep networks which are publicly available [11]. CQ transfer learning consists of using exactly those classical pre-trained models as feature extractors and then post-processing such features on a quantum computer; for example by using them as input variables for the dressed quantum circuit model introduced in Eq. (9). This hybrid approach is very convenient for processing high-resolution images since, in this configuration, a quantum computer is applied only to a fairly limited number of abstract features, which is much more feasible compared to embedding millions of raw pixels in a quantum system. We would like to mention that also other alternative approaches for dealing with large images have been recently proposed [21, 29, 35, 42].

We applied our model for the task of image classification in several numerical examples and we also tested the algorithm with two real quantum computers provided by IBM and Rigetti. All the details about the technical implementation and the associated results are reported in the next Section, Examples 2 and 3.

3.2 Quantum to classical transfer learning

By switching the roles of the classical and quantum networks, one can also obtain the QC variant of transfer learning. In this case a pre-trained quantum system behaves as a kind of feature extractor, i.e., a device performing a (potentially classically intractable) computation resulting in an output vector of numerical values associated to the input. As a second step,

a classical network is used to further process the extracted features for the specific problem of interest. This scheme can be very useful in two important situations: (i) if the dataset consists of quantum states (e.g., in a state classification problem), (ii) if we have at our disposal a very good quantum computer which outperforms current classical feature extractors at some task.

For case (i), one can imagine a situation in which a single instance of a variational quantum circuit is first pre-trained and then used as a kind of *multipurpose measurement device*. Indeed one could make many different experimental analyses by simply letting input quantum systems pass through the same fixed circuit and applying different classical machine learning algorithms to the associated measured variables.

For case (ii) instead, one can envisage a multi-party scenario in which many classical clients B can independently send samples of their specific datasets to a common quantum server A which is pre-trained to extract generic features by performing a fixed quantum computation. Server A can send back the resulting features to the classical clients B , which can now locally train their specific machine learning models on pre-processed data.

Given the current status of quantum technology, case (ii) is likely beyond a near-term implementation. On the other hand, case (i) could already represent a realistic scenario with current technology.

In Example 4 of the next Section, we present a proof-of-concept example in which a pre-trained quantum network introduced in Ref. [24] is combined with a classical post-processing network for solving a quantum state classification problem.

3.3 Quantum to quantum transfer learning

The last possibility is the QQ transfer learning scheme, where the same technique is applied in a fully quantum mechanical fashion. In this case a quantum network A is pre-trained for a generic task and dataset. Successively, some of the final quantum layers are removed, and replaced by a trainable quantum network B which will be optimized for a specific problem. The main difference from the previous cases is that, since the process is fully quantum without intermediate measurements, features are implicitly transferred in the form of a quantum state, allowing for coherent superpositions.

The main motivation for applying a QQ transfer learning scheme is to reduce the total training time: instead of training a large variational quantum circuit, it is more efficient to initialize it with some pre-trained weights and then optimize only a couple of final layers. From a physical point of view, such optimization of the final layers could be interpreted as a change of the measurement basis which is tuned to the specific

problem of interest.

If compared with classical computers, current NISQ devices are not only noisy and small: they are also relatively slow. Training a quantum circuit might take a long time since it requires taking many measurement shots (i.e., performing a large number of actual quantum experiments) for each optimization step (e.g., for computing the gradient). Therefore any approach which can reduce the total training time, as for example the QQ transfer learning scheme, could be very helpful.

In Example 5 of the next Section, we trained a quantum state classifier by following a QQ transfer learning approach.

4 Examples

The Python code related to the examples presented in this section is available at [2].

Example 1 - A 2D classifier based on a dressed quantum circuit

This first example demonstrates the *dressed* quantum circuit model introduced in Eq. (9). We consider a




Figure 2: Classification based on a classical network (left) and a dressed quantum circuit (right) of the same dataset consisting of two set of points (blue and red) organized in two concentric spirals. Training points are pale-colored while test points are sharp-colored. The model decision function is evaluated for the whole 2D plane, determining the blue and red regions which should ideally match the color of the data points. The test accuracy of each model is reported in the bottom-right of the corresponding plot.

typical benchmark dataset consisting of two classes of points (blue and red) organized in two concentric spirals as shown in Fig. 2. Each point is characterized by two real coordinates and we assume to have at our disposal a quantum processor of 4 qubits. Since we have two real coordinates as input and two real variables as output (one-hot encoding the blue and red classes), we use the following model of a dressed quantum circuit:

$$\tilde{Q} = L_{4 \rightarrow 2} \circ Q \circ L_{2 \rightarrow 4}, \quad (10)$$

where $L_{2 \rightarrow 4}$ represents a classical layer having the structure of Eq. (1) with $\varphi = \tanh$, Q is a (bare)

variational quantum circuit, and $L_{4 \rightarrow 2}$ is a linear classical layer without activation i.e., with $\varphi(\mathbf{y}) = \mathbf{y}$. The structure of the variational circuit is $Q = \mathcal{M} \circ \mathcal{Q} \circ \mathcal{E}$ as in Eq. (7). The chosen embedding map prepares each qubit in a balanced superposition of $|0\rangle$ and $|1\rangle$ and then performs a rotation around the y axis of the Bloch sphere parametrized by a classical vector \mathbf{x} :

$$\mathcal{E}(\mathbf{x}) = \bigotimes_{k=1}^4 \left(R_y(x_k \pi/2) H \right) |0000\rangle, \quad (11)$$

where H is the single-qubit Hadamard gate. The trainable circuit is composed of 5 variational layers $\mathcal{Q} = \mathcal{L}_5 \circ \mathcal{L}_4 \circ \mathcal{L}_3 \circ \mathcal{L}_2 \circ \mathcal{L}_1$, where

$$\mathcal{L}(\mathbf{w}) : |x\rangle \rightarrow |y\rangle = K \bigotimes_{k=1}^4 R_y(w_k) |x\rangle, \quad (12)$$

and K is an entangling unitary operation made of three controlled NOT gates:

$$K = (CNOT \otimes \mathbb{I}_{3,4})(\mathbb{I}_{1,2} \otimes CNOT)(\mathbb{I}_1 \otimes CNOT \otimes \mathbb{I}_4). \quad (13)$$

Finally, the measurement layer is simply given by the expectation value of the $Z = \text{diag}(1, -1)$ Pauli matrix, locally estimated for each qubit:

$$\mathcal{M}(|y\rangle) = \mathbf{y} = \begin{bmatrix} \langle y | Z \otimes \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I} | y \rangle \\ \langle y | \mathbb{I} \otimes Z \otimes \mathbb{I} \otimes \mathbb{I} | y \rangle \\ \langle y | \mathbb{I} \otimes \mathbb{I} \otimes Z \otimes \mathbb{I} | y \rangle \\ \langle y | \mathbb{I} \otimes \mathbb{I} \otimes \mathbb{I} \otimes Z | y \rangle \end{bmatrix} \quad (14)$$

Given an input point of coordinates $\mathbf{x} = (x_1, x_2)$, the classification is done according to $\text{argmax}(\mathbf{y})$, where $\mathbf{y} = (y_1, y_2)$ is the output of the dressed quantum circuit (10).

For training and testing the model, the dataset has been divided into 2000 training points (pale-colored in Fig. 2) and 200 test points (sharp-colored in Fig. 2). As typical in classification problems, the cross entropy (implicitly preceded by a LogSoftMax layer) was used as a loss function. The LogSoftMax layer acts as $y_j \rightarrow \hat{p}_j = e^{y_j} / (e^{y_1} + e^{y_2})$, such that $\{\hat{p}_j\}$ is a valid probability distribution for $j = 1, 2$. The predicted distribution $\{\hat{p}_j\}$ should approximate the true one $\{p_j\}$, which in this example is given by: $\{p_1 = 1, p_2 = 0\}$, if \mathbf{x} belongs to the red class; $\{p_1 = 0, p_2 = 1\}$, if \mathbf{x} belongs to the blue class. The cross entropy loss function is defined as:

$$J = -\mathbb{E} \sum_j p_j \log(\hat{p}_j) = -\mathbb{E} \sum_j p_j [y_j - \log(\sum_k e^{y_k})], \quad (15)$$

where the average \mathbb{E} is with respect to the train/test datasets. The right-hand-side of Eq. (15) was minimized via the Adam optimizer [26]. A total number of 1000 training iterations were performed, each of them

with a batch size of 10 input samples. The numerical simulation was done through the PennyLane software platform [9].

The results are reported in Fig. 2, where the dressed quantum network is also compared with an entirely classical counterpart in which the quantum circuit is replaced by a classical layer, i.e., $C = L_{4 \rightarrow 2} \circ L_{4 \rightarrow 4} \circ L_{2 \rightarrow 4}$. The corresponding accuracy, i.e., the fraction of test points correctly classified, is 0.97 for the dressed quantum circuit and 0.85 for the classical network.

The presented results suggest that a dressed quantum circuit is a very flexible quantum machine learning model which is capable of classifying highly non-linear datasets. We would like to remark that the classical counter-part has been presented just as a qualitative benchmark: even if in this particular example the quantum model outperforms the classical one, any general and rigorous comparison would require a much more complex and detailed analysis which is beyond the aim of this work. It is also worth to mention that the training time was significantly longer for the (simulated) quantum model (≈ 2 sec) with respect to the classical model (≈ 0.01 sec).

Example 2 - CQ transfer learning for image classification (ants / bees)

In this second example we apply the classical-to-quantum transfer learning scheme for solving an image classification problem. We first numerically trained and tested the model, using PennyLane with the PyTorch [32] interface. Successively, we have also run it on two real quantum devices provided by *IBM* and *Rigetti*. To our knowledge, this is the first time that high resolution images have been classified with a hybrid classical-quantum system.

Our example is a quantum model inspired by the official PyTorch tutorial on classical transfer learning [1]. The model can be defined in terms of the general CQ scheme proposed in Section 3 and represented in Fig. 1, with the following specific settings:

$D_A = \text{ImageNet}$: a public image dataset with 1000 classes [13].

$A = \text{ResNet18}$: a pre-trained residual neural network introduced by Microsoft in 2016 [20].

$T_A = \text{Classification (1000 labels)}$.

$A' = \text{ResNet18 without the final linear layer, obtaining a pre-trained extractor of 512 features}$.

$D_B = \text{Images of two classes: ants and bees (Hymenoptera subset of ImageNet), separated into a training set of 245 images and a testing set of 153 images}$.

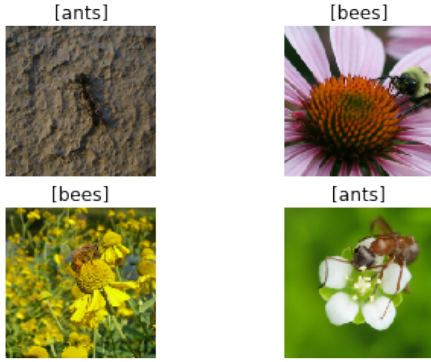


Figure 3: Random batch of 4 images sampled from the test dataset D_B and classified by our classical-quantum model (numerically simulated). Predictions are reported in square brackets above each image.

$B = \tilde{Q} = L_{4 \rightarrow 2} \circ Q \circ L_{512 \rightarrow 4}$: i.e., a 4-qubit dressed quantum circuit (9) with 512 input features and 2 real outputs.

$T_B = \text{Classification (2 labels)}$.

The bare variational circuit is essentially the same as the one used in the previous example (see Eqs. (11,12,13,14)), with the only difference that in this case the quantum depth is set to 6. The cross entropy is used as a loss function and minimized via the Adam optimizer [26]. We trained the variational parameters of the model for 30 epochs over the training dataset, with a batch size of 4 and an initial learning rate of $\eta = 0.0004$, which was successively reduced by a factor of 0.1 every 10 epochs. After each epoch, the model was validated with respect to the test dataset, obtaining a maximum accuracy of 0.967. A visual representation of a random batch of images sampled from the test dataset and the corresponding predictions is given in Fig. 3.

We also tested the model (with the same pre-trained parameters), on two different real quantum computers: the *ibmqx4* processor by IBM and the *Aspen-4-4Q-A* processor by Rigetti (see Fig. 4). The corresponding classification accuracies, evaluated on the same test dataset, are reported in Table 2. For the simulation, the exact expectation values were numerically evaluated. On the other hand, for both experiments performed on real quantum processors, we averaged 1024 measurement shots for each expectation value.

We did not apply any particular strategy to reduce the execution time of the hybrid model since it was not relevant for our aims. However, for the sake of completeness, we report that the time necessary to test the model over the full dataset was approximately: ≈ 26 sec when using *ibmqx4* with a free account, ≈ 7 sec when using *Aspen-4-4Q-A*, and ≈ 2 sec when using the default PennyLane simulator

QPU	Accuracy
Simulator	0.967
ibmqx4	0.95
Aspen-4-4Q-A	0.80

Table 2: Image classification accuracy associated to different quantum processing units (QPU): exact simulator, *ibmqx4* (IBM) and *Aspen-4-4Q-A* (Rigetti).




Figure 4: Two high-resolution images sampled from the dataset D_B and experimentally classified with two different quantum processors: *ibmqx4* by IBM (first line) and *Aspen-4-4Q-A* by Rigetti (second line). In both cases the same pre-trained classical network (ResNet18 by Microsoft [20]) was used to pre-process the input image, extracting 512 highly informative features. The rest of the computation was performed by a trainable variational quantum circuit “dressed” by two classical encoding and decoding layers, as described in Eq. (9).

on a laptop. As a qualitative reference, we also report that time corresponding to the classical network ResNet18 (without any quantum layer) was ≈ 1 sec.

Our results demonstrate the promising potential of the CQ transfer learning scheme applied to current NISQ devices, especially in the context of high-resolution image processing.

Example 3 - CQ transfer learning for image classification (CIFAR)

We now apply the same CQ transfer learning scheme of the previous example but with a different dataset D_B . Instead of classifying images of *ants* and *bees*, we use the standard CIFAR-10 dataset [27] restricted to the classes of *cats* and *dogs*. Successively we also repeat again the training and testing phases with the CIFAR-10 dataset restricted to the classes of *planes* and *cars* (see Fig. 5).

We remark that, in both cases, the feature extractor ResNet18 is pre-trained on ImageNet. Despite CIFAR-10 and ImageNet being quite different datasets (they also have very different resolutions), the CQ transfer learning method achieves nonetheless relatively good results.




Figure 5: Random batches of 4 images sampled from the CIFAR-10 test dataset restricted to the classes of *cats* and *dogs* (a), and to the classes of *planes* and *cars* (b). In both cases the binary classification problem is solved by our hybrid classical-quantum model (numerically simulated). Predictions are reported in square brackets above each image.

	ants/ bees	dogs/ cats	planes/ cars
Quantum depth	6	5	4
Number of epochs	30	3	3
Batch size	4	8	8
Learning rate	0.0004	0.001	0.0007
Accuracy	0.976	0.8270	0.9605

Table 3: Hyper-parameters used for the classification of different image datasets D_B . The last line reports the corresponding accuracy achieved by our model (numerically simulated).

The hyper-parameters used for all the different datasets (including the previous example) are summarized in Table 3. The corresponding test accuracies are also reported in Table 3, while some of the predictions for random samples of the restricted CIFAR datasets are visualized in Fig. 5.

Example 4 - QC transfer learning for quantum state classification

Quantum to classical (QC) transfer learning consists of using a pre-trained quantum circuit as a feature extractor and in post-processing its output variables with a classical neural network. In this case only the final classical part will be trained to the specific problem of interest.

The starting point of our example is the pre-trained continuous-variable quantum network presented in Ref. [24], Section IV.D, Experiment C. The original aim of this network was to encode 7 different 4×4 images, representing the (L,O,T,I,S,J,Z) tetrominos (popularized by the video game *Tetris* [3]), in the Fock basis of two-mode quantum states. The expected input of the quantum network is one of the following



Figure 6: Images of the 7 different tetrominos encoded in the photon number probability distribution of two optical modes, after projecting on the subspace of up to 3 photons. These images are extracted from Fig. 10 of Ref. [24].

combinations of two-mode coherent states:

$$\begin{aligned}
 |\varphi_1\rangle &= |\alpha\rangle|\alpha\rangle \\
 |\varphi_2\rangle &= |-\alpha\rangle|\alpha\rangle \\
 |\varphi_3\rangle &= |\alpha\rangle|-\alpha\rangle \\
 |\varphi_4\rangle &= |-\alpha\rangle|-\alpha\rangle \\
 |\varphi_5\rangle &= |i\alpha\rangle|i\alpha\rangle \\
 |\varphi_6\rangle &= |-\alpha\rangle|i\alpha\rangle \\
 |\varphi_7\rangle &= |i\alpha\rangle|-\alpha\rangle
 \end{aligned} \tag{16}$$

where the parameter $\alpha = 1.4$ is a fixed constant. In Ref. [24] the network was successfully trained to generate an optimal unitary operation $|\tilde{\varphi}_j\rangle = U|\varphi_j\rangle$, such that the probability of finding i photons in the first mode and j photons in the second mode is proportional to the amplitude of the image pixel (i, j) . More precisely, the network was trained to reproduce the tetromino images after projecting the quantum state on the subspace of up to 3 photons (see Fig. 6).

For the purposes of our example, we now assume that the previous 7 input states (16) are subject to random Gaussian displacements in phase space:

$$|\varphi_j\rangle \xrightarrow{\text{random displacement}} \hat{D}(\delta\alpha_1, \delta\alpha_2)|\varphi_j\rangle, \tag{17}$$

where $\hat{D}(\alpha_1, \alpha_2)$ is a two-mode displacement operator [50], the values of the complex displacements $\delta\alpha_1$ and $\delta\alpha_2$ are sampled from a symmetric Gaussian distribution with zero mean and quadrature variance $N \geq 0$, and $j \in \{1, 2, 3, 4, 5, 6, 7\}$ is the label associated to the input states (16). The noise is similar to a Gaussian additive channel [50]; however, for simplifying the numerical simulation, here we assume that the unknown displacements remain constant during the estimation of expectation values. Physically, this situation might represent a slow phase-space drift of the input light mode.

We also assume that, differently from the original image encoding problem studied in Ref. [24], our new task is to classify the noisy input states. In other words, the network should take the states defined in (17) as inputs, and should ideally produce the correct label $j \in \{1, 2, 3, 4, 5, 6, 7\}$ as output. In order to tackle this problem, we apply a QC transfer learning approach: we pre-process our random input states with the quantum network of Ref. [24] and we consider the corresponding images as features which we

are going to post-process with a classical layer to predict the state label j . In simple terms, the QC transfer learning method allows us to convert a quantum state classification problem into an image classification problem.

Also in this case we can summarize the transfer learning scheme according to the notation introduced in Section 3 and represented in Fig. 1:

$D_A = 7$ two-mode coherent states defined in Eq. (16).

$A =$ Photonic neural network introduced in Ref. [24], consisting of an encoding layer, 25 variational layers, and a final (Fock) measurement layer.

$T_A =$ Fock basis encoding of tetrominos images (see Fig. 6).

$A' =$ Pre-trained network A , truncated up to a quantum depth of $q = 15$ variational layers.

$D_B =$ Same states of the original dataset D_A but subject to random phase-space displacements as described in Eq. (17).

$B = L_{16 \rightarrow 7}$: i.e., a classical linear layer having the structure of Eq. (1), without activation ($\varphi(\mathbf{y}) = \mathbf{y}$).

Also in this case we used the Adam optimizer [26] to minimize a cross-entropy loss function associated to our classification problem. For each optimization step we sampled independent random displacements with variance $N = 0.6$ that we applied to a batch of 7 states defined in Eq. (16). We optimized the model over 1000 training batches with a learning rate of $\eta = 0.01$, obtaining a classification accuracy of 0.803. The numerical simulation was performed with the Strawberry Fields software platform [25], combined with the TensorFlow [4] optimization back-end.

A summary of the hyper-parameters and of the corresponding accuracy is given in Table 4.

	QC Classifier
Quantum depth q	15
Classical depth	1
Noise variance N	0.6
Training batches	1000
Batch size	7
Learning rate	0.01
Fock-space cutoff	11
Accuracy	0.803

Table 4: Hyper-parameters used for our quantum state classifier based on the QC transfer learning scheme. The last line reports the corresponding accuracy achieved by the model, simulated on Strawberry Fields with a fixed cutoff in the Fock basis.




Figure 7: Features of a batch of 7 noisy states extracted by the quantum network A' and represented as 4×4 gray-scale images. The associated classes and the predictions made by our quantum-classical model are reported above each image.




Figure 8: Accuracy of the hybrid QC classifier with respect to the quantum depth of the pre-trained network A' , evaluated for three different classical networks B of depth 1, 2, and 3 respectively. The existence of an intermediate optimal value for the quantum depth is a characteristic signature typical of the transfer learning method.

Finally, the predictions for a sample of 7 noisy states are graphically visualized in Fig. 7, where the features extracted by the pre-trained quantum network A' are represented as 4×4 gray scale images. The features of Fig. 7 are quite different from the original tetrominos images shown in Fig. 6. This due to the truncation of network A and to the presence of input noise. However, as long as the images of Fig. 7 are distinguishable, this is not a relevant issue since the final classical layer is still able to correctly classify the input states with high accuracy.

We conclude this example with an analysis of the model performance with respect to the values of the quantum and classical depths. Since the original pre-trained network A has 25 quantum layers, for the truncated network A' we can choose a quantum depth q within the interval 0-25. For the classical network B we consider the cases of 1, 2 and 3 layers, corresponding to the models $L_{16 \rightarrow 7}$, $L_{16 \rightarrow 7} \circ L_{16 \rightarrow 16}$ and $L_{16 \rightarrow 7} \circ L_{16 \rightarrow 16} \circ L_{16 \rightarrow 16}$, respectively.

The results are shown in Fig. 8. By direct inspection we can see that increasing the classical depth is helpful but it saturates the accuracy already after two layers. On the other hand, it is evident that the quantum depth has an optimal value around $q = 15$ while for larger values the accuracy is reduced. This

QQ Classifier	
Depth of A'	1
Depth of B	3
Training batches	500
Batch size	8
Learning rate	0.01
Fock-space cutoff	15
Accuracy	0.869

Table 5: Hyper-parameters used for our quantum state classifier based on the QQ transfer learning scheme. The last line reports the corresponding accuracy achieved by the model (numerically simulated).

is a paradigmatic phenomenon well known in classical transfer learning: better features are usually extracted after removing some of the final layers of A . Notice that because of the quantum nature of the system, the quantum state produced by the truncated variational circuit could be entangled and/or not aligned with the measurement basis. So the numerical evidence that the truncation of a quantum network does not always reduce the quality of the measured features, but it can actually be a convenient strategy for transfer learning, is a notable result.

Example 5 - QQ transfer learning for quantum state classification (Gaussian / non-Gaussian)

Finally, our last example is a proof-of-principle demonstration of QQ transfer learning. In this case we train an optical network A to classify a particular dataset D_A of Gaussian and non-Gaussian quantum states. Successively, we use it as a pre-trained block for a dataset D_B consisting of Gaussian and non-Gaussian states which are different from those of D_A . The pre-trained block is followed by some quantum variational layers that will be trained to classify the quantum states of D_B .

Before presenting our model we need to define a continuous-variable single-mode variational layer, the analog of Eq. (12). We follow the general structure proposed in Ref. [24]:

$$\mathcal{L} = \Phi \circ \mathcal{D} \circ \mathcal{R} \circ \mathcal{S} \circ \mathcal{R}, \quad (18)$$

where \mathcal{R} is a phase space rotation, \mathcal{S} is a squeezing operation, \mathcal{D} is a displacement and Φ is a cubic phase gate. All operations depend on variational parameters and, for sufficiently many layer applications, the model can generate any single-mode unitary operation. Moreover, by simply removing the last non-Gaussian gate Φ from (18), we obtain a Gaussian layer which can generate all Gaussian unitary operations.

We can express the QQ transfer learning model of this example according to the notation introduced in Section 3 and represented in Fig. 1:





Figure 9: Evolution of the loss function (cross entropy) with respect to the number of training iterations. The top plot represents a network of total depth 3 optimized with a QQ transfer learning scheme (orange) compared with a network trained from scratch (blue). In the bottom plot instead the total depth is fixed to 4 layers.

D_A = Two classes, 0 and 1, of quantum states generated by two different variational random circuits. States of class 0 are generated by a random single-mode Gaussian layer applied to the vacuum. States of class 1 are generated by a random non-Gaussian layer applied to the vacuum.

A = Single-mode variational quantum layer followed by an on/off threshold detector.

T_A = Classification (labels: 0 and 1).

A' = Network A without the measurement layer.

D_B = Two classes, 0 and 1, of quantum states. States of class 0 are generated by a random single-mode Gaussian layer applied to the coherent state $|\alpha\rangle$ with $\alpha = 1$. States of class 1 are generated by a random Gaussian layer applied to the Fock state $|1\rangle$.

B = Single-mode variational quantum circuit of depth q , followed by a on/off threshold detector.

A summary of the hyper-parameters used for defining and training this QQ model is reported in Table

5, together with the associated accuracy. In Fig. 9 we plot the loss function (cross entropy) of our quantum variational classifier with respect to the number of training iterations. We compare the results obtained with and without the pre-trained layer A' (i.e., with and without transfer learning), for a fixed total depth of 3 or 4 layers. It is clear that the QQ transfer learning approach offers a strong advantage in terms of training efficiency.

For a sufficiently long training time however, the network optimized from scratch achieves the same or better results with respect to the network with a fixed initial layer A' . This effect is well known also in the classical setting and it is not surprising: the network trained from scratch is in principle more powerful by construction, because it has more variational parameters. However, there are many practical situations in which the training resources are limited (especially when dealing with real NISQ devices) or in which the dataset D_B is experimentally much more expensive with respect to D_A . In all these kind of practically constrained situations, QQ transfer learning could represent a very convenient strategy.

5 Conclusions

We have outlined a framework of transfer learning which is applicable to hybrid computational models where variational quantum circuits can be connected to classical neural networks. With respect to the well-studied classical scenario, in hybrid systems several new and promising opportunities naturally emerge as, for example, the possibility of transferring some pre-acquired knowledge at the classical-quantum interface (CQ and QC transfer learning) or between two quantum networks (QQ transfer learning). As an additional contribution, we have also introduced the notion of “dressed quantum circuits”, i.e., variational quantum circuits augmented with two trainable classical layers which improve and simplify the data encoding and decoding phases.

Each theoretical idea proposed in this work is supported with a proof-of-concept example, numerically demonstrating the validity of our models for practical applications such as image recognition or quantum state classification. Particular focus has been dedicated to the CQ transfer learning scheme because of its promising potential with currently available quantum computers. In particular we have used the CQ transfer learning method to successfully classify high resolution images with two real quantum processors (by IBM and Rigetti).

From our theoretical and experimental analysis, we can conclude that transfer learning is a promising approach which can be particularly convenient in the context of near-term quantum devices.

Acknowledgments

We thank Christian Weedbrook for helpful discussions. The authors would like to thank Rigetti for access to their resources, Forest [45], QCS and Aspen-4-4Q-A backend. We also acknowledge the use of the IBM Q Experience, Qiskit [16] and IBM Q 5 Tenerife v1.0.0 (ibmqx4) backend.

References

- [1] Sasank Chilamkurthy, PyTorch transfer learning tutorial. https://pytorch.org/tutorials/beginner/transfer_learning_tutorial.html. Accessed: 2019-08-08.
- [2] <https://github.com/XanaduAI/quantum-transfer-learning>. Accessed: 2020-29-06.
- [3] Tetris, Wikipedia, 2019. <https://en.wikipedia.org/wiki/Tetris>. Accessed: 2019-08-08.
- [4] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [5] Soumik Adhikary, Siddharth Dangwal, and Debanjan Bhowmik. Supervised learning with a quantum classifier using multi-level systems. *Quantum Information Processing*, 19(3): 89, 2020. DOI: 10.1007/s11128-020-2587-9.
- [6] Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. DOI: 10.1038/s41586-019-1666-5.
- [7] Marcello Benedetti, John Realpe-Gómez, and Alejandro Perdomo-Ortiz. Quantum-assisted Helmholtz machines: A quantum–classical deep learning framework for industrial datasets in near-term devices. *Quantum Science and Technology*, 3(3):034007, 2018. DOI: 10.1088/2058-9565/aabd98.
- [8] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. DOI: 10.1109/tpami.2013.50.
- [9] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, and Nathan Killoran. PennyLane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*, 2018.
- [10] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549 (7671):195, 2017. DOI: 10.1038/nature23474.

- [11] Alfredo Canziani, Adam Paszke, and Eugenio Culurciello. An analysis of deep neural network models for practical applications. *arXiv preprint arXiv:1605.07678*, 2016.
- [12] Kelvin Ch'Ng, Juan Carrasquilla, Roger G Melko, and Ehsan Khatami. Machine learning phases of strongly correlated fermions. *Physical Review X*, 7(3):031038, 2017. DOI: [10.1103/physrevx.7.031038](https://doi.org/10.1103/physrevx.7.031038).
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. IEEE, 2009. DOI: [10.1109/CVPR.2009.5206848](https://doi.org/10.1109/CVPR.2009.5206848).
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 2019. DOI: [10.18653/v1/n19-1423](https://doi.org/10.18653/v1/n19-1423).
- [15] Vedran Dunjko, Jacob M Taylor, and Hans J Briegel. Quantum-enhanced machine learning. *Physical Review Letters*, 117(13):130501, 2016. DOI: [10.1103/physrevlett.117.130501](https://doi.org/10.1103/physrevlett.117.130501).
- [16] Héctor Abraham *et al.*. Qiskit: An open-source framework for quantum computing. doi: [10.5281/zenodo.2562110](https://doi.org/10.5281/zenodo.2562110), 2019.
- [17] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.
- [18] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [19] Aram W Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671):203, 2017. DOI: [10.1038/nature23458](https://doi.org/10.1038/nature23458).
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. DOI: [10.1109/cvpr.2016.90](https://doi.org/10.1109/cvpr.2016.90).
- [21] Maxwell Henderson, Samriddhi Shakya, Shashindra Pradhan, and Tristan Cook. Quantum neural networks: powering image recognition with quantum circuits. *Quantum Machine Intelligence*, 2(1), feb 2020. DOI: [10.1007/s42484-020-00012-y](https://doi.org/10.1007/s42484-020-00012-y).
- [22] Jeremy Howard and Sebastian Ruder. Universal language model fine-tuning for text classification. 2018. DOI: [10.18653/v1/p18-1031](https://doi.org/10.18653/v1/p18-1031).
- [23] Patrick Huembeli, Alexandre Dauphin, and Peter Wittek. Identifying quantum phase transitions with adversarial neural networks. *Physical Review B*, 97(13):134109, 2018. DOI: [10.1103/physrevb.97.134109](https://doi.org/10.1103/physrevb.97.134109).
- [24] Nathan Killoran, Thomas R. Bromley, Juan Miguel Arrazola, Maria Schuld, Nicolás Quesada, and Seth Lloyd. Continuous-variable quantum neural networks. *Physical Review Research*, 1(3), oct 2019. DOI: [10.1103/phys-revresearch.1.033063](https://doi.org/10.1103/phys-revresearch.1.033063).
- [25] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry Fields: A software platform for photonic quantum computing. *Quantum*, 3:129, 2019. DOI: [10.22331/q-2019-03-11-129](https://doi.org/10.22331/q-2019-03-11-129).
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012. DOI: [10.1145/3065386](https://doi.org/10.1145/3065386).
- [29] Ding Liu, Shi-Ju Ran, Peter Wittek, Cheng Peng, Raul Blázquez García, Gang Su, and Maciej Lewenstein. Machine learning by unitary tensor network of hierarchical tree structure. *New Journal of Physics*, 21(7):073059, 2019. DOI: [10.1088/1367-2630/ab31ef](https://doi.org/10.1088/1367-2630/ab31ef).
- [30] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):023023, 2016. DOI: [10.1088/1367-2630/18/2/023023](https://doi.org/10.1088/1367-2630/18/2/023023).
- [31] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2009. DOI: [10.1109/tkde.2009.191](https://doi.org/10.1109/tkde.2009.191).
- [32] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
- [33] Alejandro Perdomo-Ortiz, Marcello Benedetti, John Realpe-Gómez, and Rupak Biswas. Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers. *Quantum Science and Technology*, 3(3):030502, 2018. DOI: [10.1088/2058-9565/aab859](https://doi.org/10.1088/2058-9565/aab859).
- [34] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O'brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5: 4213, 2014. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [35] Sebastien Piat, Nairi Usher, Simone Severini, Mark Herbster, Tommaso Mansi, and Peter Mountney. Image classification with quantum pre-training and auto-encoders. *International Journal of Quantum*

- Information*, 16(08):1840009, 2018. DOI: [10.1142/s0219749918400099](https://doi.org/10.1142/s0219749918400099).
- [36] Lorien Y Pratt. Discriminability-based transfer between neural networks. In *Advances in Neural Information Processing Systems*, pages 204–211, 1993.
- [37] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [38] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y Ng. Self-taught learning: transfer learning from unlabeled data. In *Proceedings of the 24th International Conference on Machine Learning*, pages 759–766. ACM, 2007. DOI: [10.1145/1273496.1273592](https://doi.org/10.1145/1273496.1273592).
- [39] Maria Schuld and Nathan Killoran. Quantum machine learning in feature Hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019. DOI: [10.1103/physrevlett.122.040504](https://doi.org/10.1103/physrevlett.122.040504).
- [40] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015. DOI: [10.1080/00107514.2014.964942](https://doi.org/10.1080/00107514.2014.964942).
- [41] Maria Schuld, Alex Bocharov, Krysta M. Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3), mar 2020. DOI: [10.1103/physreva.101.032308](https://doi.org/10.1103/physreva.101.032308).
- [42] Kodai Shiba, Katsuyoshi Sakamoto, Koichi Yamaguchi, Dinesh Bahadur Malla, and Tomah Sogabe. Convolution filter embedded quantum gate autoencoder. *arXiv preprint arXiv:1906.01196*, 2019.
- [43] Sukin Sim, Peter D. Johnson, and Alán Aspuru-Guzik. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Advanced Quantum Technologies*, 2(12):1900070, 2019. DOI: [10.1002/qute.201900070](https://doi.org/10.1002/qute.201900070).
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [45] Robert S Smith, Michael J Curtis, and William J Zeng. A practical quantum instruction set architecture. *arXiv preprint arXiv:1608.03355*, 2016. DOI: [10.5281/zenodo.3677540](https://doi.org/10.5281/zenodo.3677540).
- [46] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. DOI: [10.1109/cvpr.2015.7298594](https://doi.org/10.1109/cvpr.2015.7298594).
- [47] Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of Research on Machine Learning Applications and Trends: Algorithms, Methods, and Techniques*, pages 242–264. IGI Global, 2010. DOI: [10.4018/978-1-60566-766-9.ch011](https://doi.org/10.4018/978-1-60566-766-9.ch011).
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [49] Guillaume Verdon, Michael Broughton, Jarrod R. McClean, Kevin J. Sung, Ryan Babush, Zhang Jiang, Hartmut Neven, and Masoud Mohseni. Learning to learn with quantum neural networks via classical neural networks. *arXiv preprint arXiv:1907.05415*, 2019.
- [50] Christian Weedbrook, Stefano Pirandola, Raúl García-Patrón, Nicolas J Cerf, Timothy C Ralph, Jeffrey H Shapiro, and Seth Lloyd. Gaussian quantum information. *Reviews of Modern Physics*, 84(2):621, 2012. DOI: [10.1103/RevModPhys.84.621](https://doi.org/10.1103/RevModPhys.84.621).
- [51] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems*, pages 3320–3328, 2014.
- [52] Remmy Zen, Long My, Ryan Tan, Frédéric Hébert, Mario Gattobigio, Christian Miniatura, Dario Poletti, and Stéphane Bressan. Transfer learning for scalability of neural-network quantum states. *Physical Review E*, 101(5), 2020. DOI: [10.1103/physreve.101.053301](https://doi.org/10.1103/physreve.101.053301).