






Benchmarking adversarially robust quantum machine learning at scaleMaxwell T. West ^{1,*} Sarah M. Erfani ² Christopher Leckie ² Martin Sevier ¹
Lloyd C. L. Hollenberg^{1,3} and Muhammad Usman ^{1,3,4,†}¹*School of Physics, The University of Melbourne, Parkville, Victoria 3010, Australia*²*School of Computing and Information Systems, Melbourne School of Engineering, The University of Melbourne, Parkville, Victoria 3010, Australia*³*Center for Quantum Computation and Communication Technologies, The University of Melbourne, Parkville, Victoria 3010, Australia*⁴*Data61, CSIRO, Research Way Clayton, Victoria 3168, Australia*

(Received 21 November 2022; accepted 26 May 2023; published 23 June 2023)

Machine learning (ML) methods such as artificial neural networks are rapidly becoming ubiquitous in modern science, technology, and industry. Despite their accuracy and sophistication, neural networks can be easily fooled by carefully designed malicious inputs known as adversarial attacks. While such vulnerabilities remain a serious challenge for classical neural networks, the extent of their existence is not fully understood in the quantum ML setting. In this paper, we benchmark the robustness of quantum ML networks, such as quantum variational classifiers (QVC), at scale by performing rigorous training for both simple and complex image datasets and through a variety of high-end adversarial attacks. Our results show that QVCs offer a notably enhanced robustness against classical adversarial attacks by learning features, which are not detected by the classical neural networks, indicating a possible quantum advantage for ML tasks. Contrarily, and remarkably, the converse is not true, with attacks on quantum networks also capable of deceiving classical neural networks. By combining quantum and classical network outcomes, we propose an adversarial attack detection technology. Traditionally quantum advantage in ML systems has been sought through increased accuracy or algorithmic speed-up, but our study has revealed the potential for a kind of quantum advantage through superior robustness of ML models, whose practical realization will address serious security concerns and reliability issues of ML algorithms employed in a myriad of applications including autonomous vehicles, cybersecurity, and surveillance robotic systems.

DOI: [10.1103/PhysRevResearch.5.023186](https://doi.org/10.1103/PhysRevResearch.5.023186)**I. INTRODUCTION**

The past decade has seen an extraordinary multidisciplinary uptake of machine learning (ML) methods, driven largely by the success of deep neural networks [1], in a diverse set of scientific applications including, for example, image classification [2], natural language processing [3], protein structure prediction [4], and quantum circuit optimization [5]. At the same time, the rise of autonomous vehicles, drones, and robots has seen ML technology increasingly enter industrial and military use, with attention naturally turning to the reliability of such ML tools in the face of malicious actors who may seek to exploit them [6–10]. Most notably, a serious blow to the prospect of reliably using neural networks in security conscious environments has been delivered in the discovery of adversarial attacks, and the resulting field of adversarial ML [6–10]. The key finding of adversarial ML is that even well trained, high performing neural networks will generally

possess serious vulnerabilities to inputs, which have been carefully crafted in order to deceive them, despite possibly being all but indistinguishable [7] from genuine inputs, which the network can classify easily. This discovery renders the use of artificial neural networks in environments where the input source cannot be trusted a dangerous proposition, and is a key roadblock on the path to widespread deployment of artificial intelligence in general scenarios [10]. Significant efforts have been made to strengthen neural networks against adversarial attacks [11–16], including, for example, the judicious injection of small amounts of randomness in order to disturb any sensitively constructed adversarial inputs [15], or the inclusion of adversarial examples at training time in order to build robustness [11,16], but a universal and reliable defense mechanism still remains elusive. While for now the long-term prospects of classical neural networks in the face of adversarial attacks remains unclear, recently attention has been turning to how emerging quantum ML solutions will fare against adversarial attacks.

Quantum ML is a new paradigm for the design of ML solutions [17] with the possibility of exploiting the capabilities of quantum computing for superior performance in ML applications, either by utilizing quantum subroutines to enhance the performance of classical ML classifiers [18,19], or through the development of classifiers, which are themselves inherently quantum [20–26]. Remarkable advances

*westm2@student.unimelb.edu.au

†musman@unimelb.edu.au

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

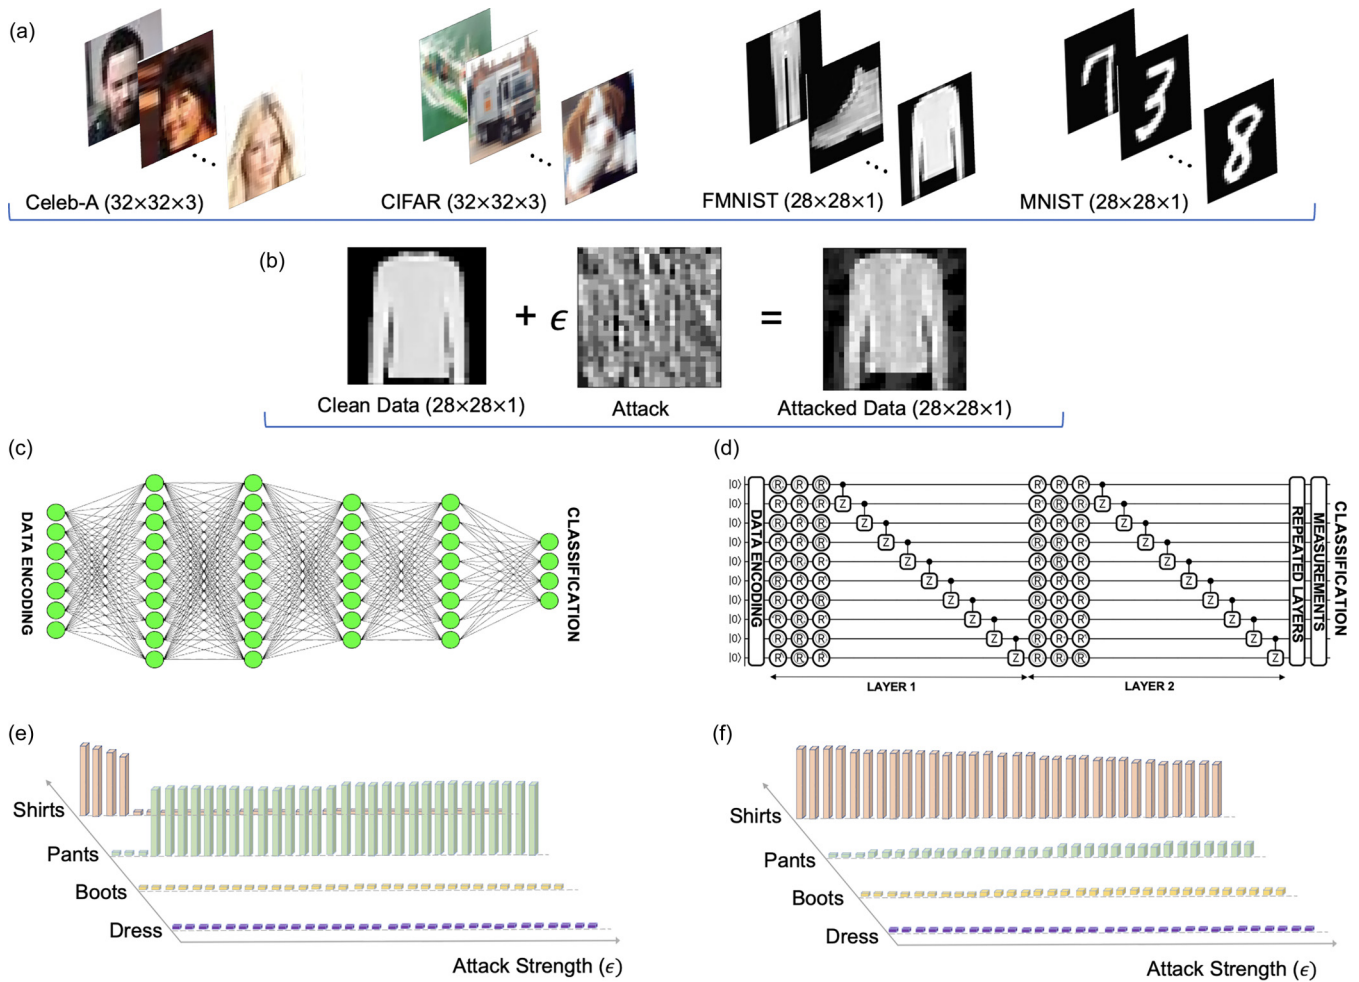


FIG. 1. Adversarial machine learning benchmark framework. A flowchart diagram of the quantum/classical adversarial ML framework developed to benchmark their robustness against a range of sophisticated adversarial attacks. (a) A variety of image datasets (MNIST, FMNIST, CIFAR, and Celeb-A) are selected to train and test classical and quantum ML networks. (b) An example image from the FMNIST dataset along with an adversarial attack and the attacked image is shown. For each image dataset, adversarial attacks (PGD, FGSM, and AutoAttack) are generated with variable strengths controlled by ϵ (with respect to the l_∞ norm). Importantly, we generate both quantum and classical attacks to test their transferability across networks, i.e., quantum attacks on classical networks and vice versa. (c) Schematic diagram of a typical classical neural network is illustrated. In our paper, two classical neural networks (ConvNet and ResNet) were trained to directly compare their performance against the quantum ML networks. The exact configurations of the classical networks are described in the Appendix. (d) Schematic diagram of a quantum variational classifier (QVC) network trained and tested in this paper. The QVC network consists of a data loading layer where the input image data is amplitude encoded into a quantum state. The variational part of the QVC consists of a repeated pattern of layers—only two such layers are shown. We trained and tested a number of QVC networks with varying number of repeated layers, which are labeled based on the number of repeated layers, e.g., a QVC200 network contains 200 layers. The final stage of the network is a measurement layer, which determines the classification label. [(e),(f)] Here we depict the performance of a quantum and a classical network against an adversarial attack, where a small, carefully chosen perturbation is added to an image [as shown in (b)], which is then passed to both a neural network and a QVC for classification. The probabilities that are assigned to various labels by the networks are shown as a function of the strength of the attack. At a certain critical attack strength the correct label (“Shirts”, shown in orange) is no longer considered the most probable by the classical network [see (e)], and the attack has succeeded in fooling the model. Contrarily, the same attacked image when passed to a QVC is still correctly predicted as “Shirts” [see (f)], even when the attack strength is increased well beyond what is tolerated by the classical network.

in both quantum hardware and software development [27] have led to a great interest in developing and benchmarking a variety of quantum ML methods [28]. This naturally leads to an important question: Can quantum ML algorithms be designed to achieve superior defense against adversarial attacks compared to their classical counterparts? This is the key open question, which we aim to address in this paper. There have been a few recent studies, which have explored the extent to

which quantum classifiers themselves suffer from adversarial examples [24,29–36], giving birth to a new field of quantum adversarial ML (QAML) [36]. However, the QAML literature has to date been limited to only small-scale proof-of-concept studies such as based on downscaled MNIST data [24,29] or other simple binary classification problems [34]. Furthermore, the field has so far primarily focused on the vulnerability assessment for white-box adversarial attacks, wherein the

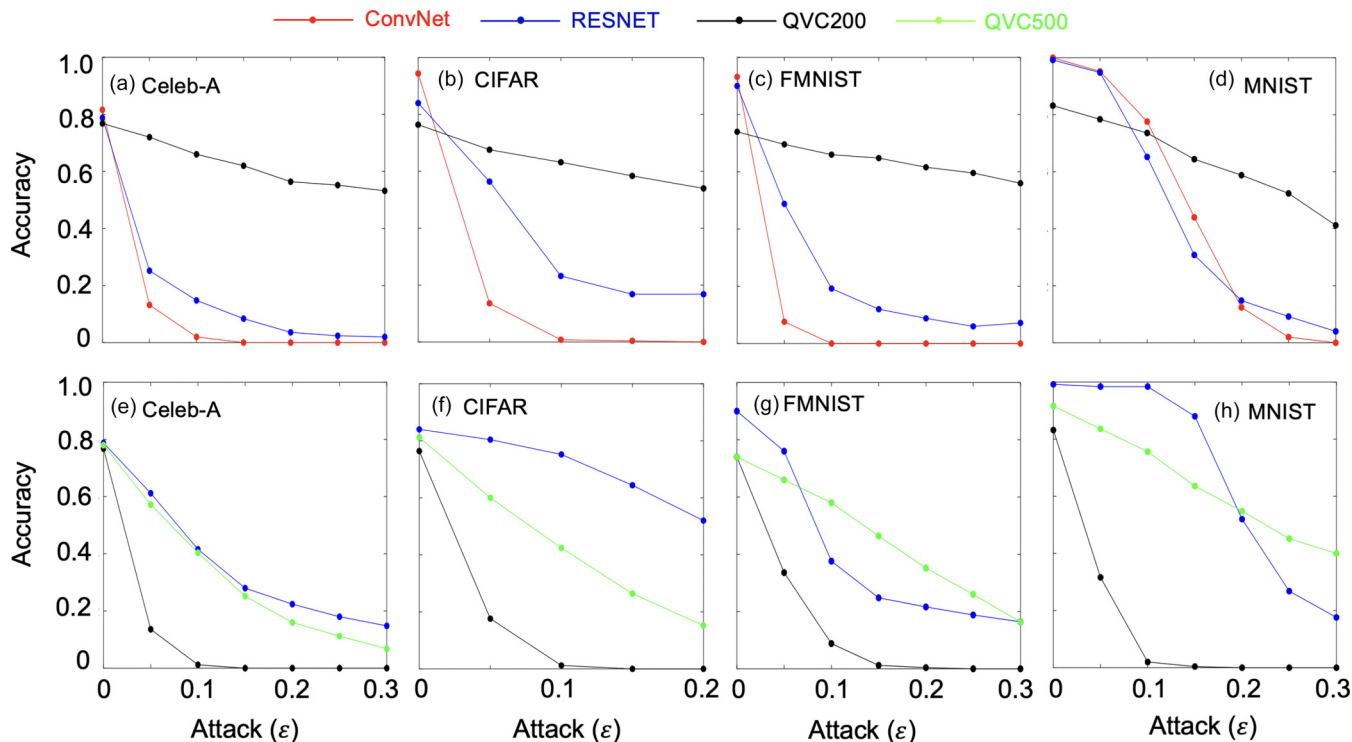


FIG. 2. Transferability and defence. The accuracy achieved by classical and quantum networks on sets of 250 adversarially attacked test images from each of the considered datasets in the cases of white-box PGD attacks on the convolutional network (top row), and 200 layer quantum variational classifier (QVC200, bottom row) as a function of attack strength ϵ (measured with the l_∞ norm). In both cases we see the accuracy of the network under attack decrease sharply. The tendency of the accuracy of the independent networks to also decrease is a manifestation of the transferability of adversarial examples—they are typically capable of fooling even networks, which they were not explicitly designed to attack. We see an exception to this in the top row, with the quantum classifier largely resisting the attacks generated with respect to the convolutional neural network. The behavior of QVCs of different depths under adversarial attacks is shown in Fig. 20.

adversary is assumed to have complete knowledge of the target network [29,30,33,34]. In the case of a classical neural network, this knowledge would include the architecture and numerical values of all weights in the network, while for a QVC it would include the architecture and the values of all trainable parameters. In contrast to white-box attacks, a black-box attack is one which is generated without precise knowledge of the ML network structure, a more realistic threat model for attacking production ML systems entrusted with security-sensitive tasks. Typically, black-box attacks are carried out by first implementing a white-box attack on a network of one’s own, and then submitting the generated adversarial example to the target network. The success of this procedure is known as the transferability of adversarial examples in the ML literature [7,37–39].

By developing an adversarial ML benchmark framework (see Fig. 1 for a flowchart description), our paper evaluates the robustness of quantum networks under true vulnerability tests through benchmarking their performance at scale, i.e., based on full-sized simple and complex image datasets without any downscaling and by creating both white and black-box attacks. The presented detailed and systematic study of defense and transferability of quantum and classical ML networks in the presence of adversarial attacks not only provides key insights, but also reveals a potential for future quantum advantage in ML applications, which could be unlocked in the next few years with the anticipated scale-up

of quantum processors coupled with the development of sophisticated error correction schemes.

In this paper, we carried out a systematic set of quantum and classical simulations across a range of image datasets [40–43] and by creating a variety of adversarial attacks. Our results analyze and compare both the defense of classical(quantum) networks against quantum(classical) adversarial attacks, and the transferability of adversarial examples within classical and quantum ML methods in black-box settings. The focus of our paper is on comparing the performance of quantum variational classifiers (QVCs), classical convolutional neural networks (CNNs), and the well-known classical neural network architecture ResNet18 [2]. Our results reveal a surprising discovery that, while the adversarial examples constructed by carrying out white-box attacks on the QVCs tend to transfer well to the classical networks, the converse is not true, with the QVCs displaying a remarkable resilience to the classical adversarial attacks (see Fig. 1) in a black-box setting. Based on the analysis of perturbations generated from classical and quantum attacks, we interpret that the observed difference between classical and quantum defense mechanisms arises because the QVCs learn a different, but highly meaningful, set of features to the classical networks, which rely on informative but nonrobust features of the data. We also investigate the performance of adversarial training of quantum networks under both white-box and black-box settings, which provides an important insight that while highly successful




FIG. 3. Content of adversarial perturbations. The adversarial perturbations generated by $\epsilon = 0.1$ PGD attacks on the convolutional neural network, ResNet and the 500 layer quantum variational classifier (QVC500) are shown for several examples from the FMNIST and Celeb-A databases. In the case of FMNIST the task is to distinguish between classes of images representing various types of clothing, e.g., T-shirts, pants, shoes, while in the case of Celeb-A the classification problem is to determine whether or not the featured person has black hair. The perturbations generated by the classical networks are incomprehensible to humans, exploiting the highly abstract features discovered by these networks. The attack on the QVC, on the other hand, yields perturbations whose semantic content is clear; they constitute concrete steps towards actually changing the label of the clean image (i.e., filling in the gap between the legs of the pants, adding sleeves to the T-shirt, removing sleeves from the jumper, changing the hair color of the imaged person).

for classical networks [11,16], the benefits of adversarial training are quite insignificant in further improving the performance of already resilient quantum ML networks against classical attacks. Interestingly, against the quantum attacks, the computationally expensive adversarial training improves the accuracy of quantum ML networks in a white-box setting, but offers diminishing return in black-box scenarios. Finally, we propose an adversarial attack detection technology, which relies on classical and quantum ML frameworks both working together to rapidly identify the presence of an adversarial attack. Overall, our results have established that a future deployment of quantum ML solutions in security conscious practical applications could offer a kind of advantage in terms of robustness of ML frameworks against adversarial attacks, which will be in addition to commonly sought enhancements in speed and/or accuracy.

II. RESULTS AND DISCUSSION

The testing and benchmarking of QAML models involve selection of datasets, generation of adversarial attacks, implementation of classical and quantum ML models, and a systematic investigation of attack transferability and defense. In this paper, we investigate QAML across a diverse set of well-known image datasets, including both grey-scale (MNIST [40] and FMNIST [41]), and RGB color (CIFAR [42] and Celeb-A [43]) images. While we use all ten classes of the MNIST and FMNIST datasets, although not a limitation of our paper, we restrict to binary classification (ships vs trucks) in the case of CIFAR in order to reduce the computational bur-

den of training the large quantum classifiers. The classification challenge considered for the Celeb-A dataset is to determine whether the pictured person has black hair. Example images from each of the datasets are shown in Fig. 1(a).

After the selection of datasets, we implemented three different types of adversarial attacks: PGD [44], FGSM [45], and AutoAttack [46]. These are some of the strongest attacks commonly used in the classical ML literature to test and benchmark the adversarial vulnerability of classical neural networks. The details of how each of these attacks attempts to find adversarial perturbations are given in Appendix B. On the quantum side, our focus is on standard quantum variational classifiers (QVCs), while on the classical side we consider convolutional neural networks (henceforth labeled as ConvNet for simplicity) and the well-known neural network architecture ResNet18 (henceforth labeled as ResNet for simplicity) [2]. The architectures of ConvNet, ResNet, and the QVCs are schematically shown in Fig. 6 below.

The first component of the QVCs is the data loading stage, in which the images are encoded into a quantum state. There are many popular strategies for encoding data into quantum computers, including angle encoding [47], amplitude encoding [47], matrix product state based block encoding [48], and the flexible representation of quantum images (FRQI) [48], each with their own advantages and disadvantages. In this paper, owing to the high dimensionality of the image data, we employ the method of amplitude encoding, which is very efficient in terms of the number of qubits required. While amplitude encoding is highly space efficient, it is not specifically tailored for image data, and for example does

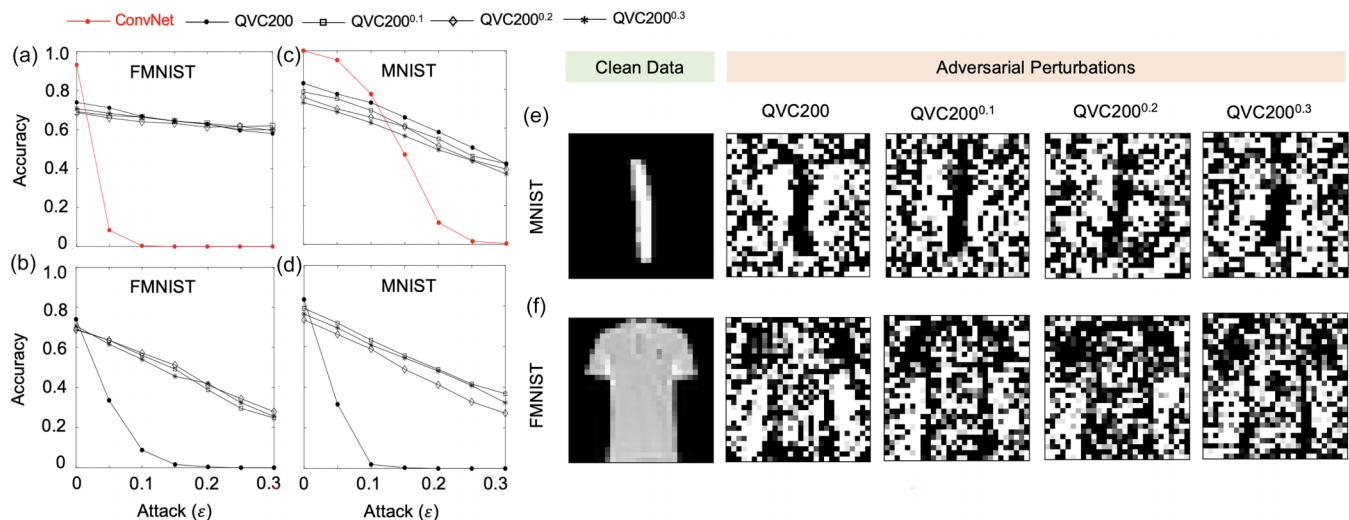


FIG. 4. Quantum ML adversarial training. (a) The accuracies achieved by quantum ML networks are plotted as a function of attack strength on a set of 250 adversarially attacked test images from the FMNIST dataset in the cases of white-box attacks on the ConvNet. The attack is applied to QVC200 and adversarially trained QVC200 networks [QVC200^{0.1}, QVC200^{0.2}, and QVC200^{0.3} where adversarial training is performed with PGD attacks of (l_∞) strength 0.1, 0.2, and 0.3 respectively]. The adversarial training makes a negligible difference to the accuracy of the quantum network against classical attacks, which are already highly resilient. This resilience is strong enough that the networks are almost entirely immune to increases in the strength of the attacks from $\epsilon = 0.2$ to $\epsilon = 0.3$, with their accuracies remaining roughly constant. In fact, there is a small nonmonotonic component in the behavior of the adversarially trained QVCs, with increases of 0.04% in test accuracy after an increase to the attack strength obtained by QVC200^{0.1} and QVC200^{0.2}. With 250 test examples, this corresponds to correctly classifying one more of the images and is due to the perturbations, despite being “adversarial” from the classical network’s point of view, looking essentially random to the QVCs. (b) The accuracies achieved by quantum ML networks are plotted as a function of attack strength on a set of 250 adversarially attacked test images from the FMNIST dataset in the cases of white-box attacks on the QVC200. The attack is applied to QVC200 and adversarially trained QVC200 networks [QVC200^{0.1}, QVC200^{0.2}, and QVC200^{0.3} where adversarial training is performed with PGD attacks of (l_∞) strength 0.1, 0.2, and 0.3 respectively]. The adversarial training significantly improves the accuracy of the quantum network. [(c),(d)] The accuracies achieved by ConvNet, ConvNet^{0.1}, ConvNet^{0.2}, ConvNet^{0.3}, QVC200^{0.1}, QVC200^{0.2}, and QVC200^{0.3} on the FMNIST dataset in the case of white-box attacks on the QVC200. The classical networks exhibit high accuracy when adversarially trained. As before, adversarial training makes negligible difference to the accuracy of the QVC200. [(e),(f)] The adversarial perturbations generated by PGD attacks on QVC200, QVC200^{0.1}, QVC200^{0.2}, and QVC200^{0.3} are shown for a sample image from the MNIST and FMNIST datasets. The perturbations corresponding to QVC200 show the presence of clear features, which actually change the label of the clean image (i.e., adding sleeves to the T-shirt). The perturbations from adversarially trained QVCs are very similar to the QVC, which underwent standard training.

not explicitly take into account spatial information such as the relative closeness of pixels in the image. The exact role of encoding schemes and in particular the investigation of image-specific encoding schemes such as FRQI in the context of the adversarial robustness of QML would be an interesting question to investigate in a future study. Specifically, given an input data vector \mathbf{x} , amplitude encoding constructs the state

$$|\psi(\mathbf{x})\rangle = \sum_i x_i |i\rangle.$$

We are therefore able to encode the 28×28 grey-scale images of the MNIST and FMNIST datasets into 10 qubits, and the $3 \times 32 \times 32$ RGB images of CIFAR and Celeb-A into 12 qubits (as $28 \times 28 < 2^{10}$ and $3 \times 32 \times 32 < 2^{12}$).

After encoding, the resulting quantum state is processed by a variable number of layers of parametrised (trainable) single-qubit rotations and entangling CZ gates [refer to Figs. 1(c) and 6(c)]. The QVC networks are labeled based on the number of layers in the architecture, e.g., QVC200 consists of 200 layers. Both classical and quantum networks were rigorously trained to achieve high accuracies (see details in Figs. 7 and 8). We note that the learning accuracies of QVCs trained

in our paper are quite close to the outcomes from the classical networks even for complex RGB datasets (CIFAR and Celeb-A) and despite that the classical networks utilize significantly more resources than the QVCs. Furthermore, our primary aim in this paper is to evaluate the robustness of quantum ML in the presence of adversarial attacks, and not so much on its performance for classification tasks, hence QVC parameters were not fine tuned for this purpose. Further details pertaining to the QVC architectures, CNNs, ResNet, and their training procedures are provided in the Appendix.

A. Adversarial attack transferability and defence

A complete testing and benchmarking process of classical and quantum ML networks involves two steps: (1) Adversarial examples created by various white-box attacks on a classical network, e.g., ConvNet and their application to a different classical network, e.g., ResNet, which analyses the transferability of attacks, as the attack generation and testing is done on two entirely different architectures. In the quantum case, this could be performed by generating adversarial examples

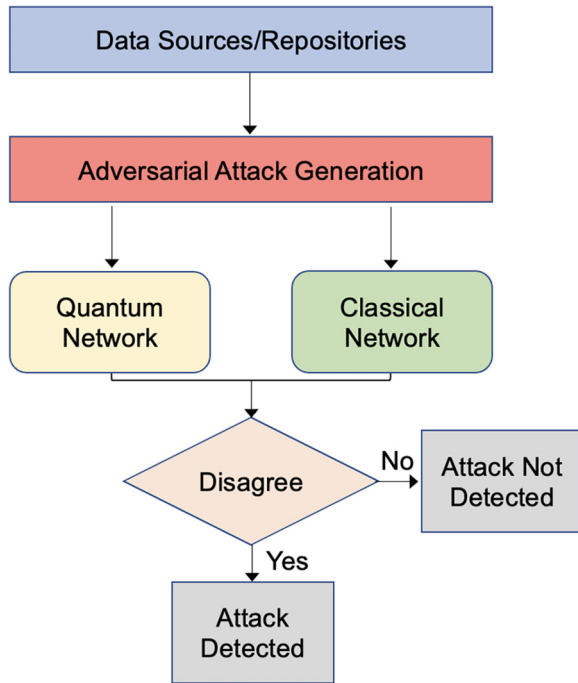


FIG. 5. Attack detection technology. A flowchart diagram illustrating the proposed attack detection strategy based on a hybrid quantum/classical approach. An attack is detected when the predictions of the classical and quantum networks disagree. A future technology developed based on such a combination of classical and quantum networks could provide the critical capability of rapid detection of adversarial attacks on ML models.

on, e.g., QVC200 network and assessing its transferability to a QVC500 network. (2) Adversarial examples created by various white-box attacks on a classical network such as ConvNet and their application to a quantum network (such as QVC200), which evaluates the defense of a quantum network against that classical attack. This also involved the testing of the defense of the classical networks (ConvNet and ResNet) against the attacks generated from quantum networks such as QVC200. Figure 9 provides a simple illustration that defines defense mechanisms and transferability of attacks across classical and quantum networks.

Figure 2 plots our results benchmarking the transferability and defense of both classical and quantum networks for all four datasets in the presence of PGD attacks. The corresponding results for FGSM and AutoAttack are plotted in Figs. 10 and 11 respectively, which exhibit very similar trends. Let us first discuss the transferability of adversarial attacks (from one classical network to another classical network, and from one quantum network to another quantum network). Along the top row [(a)–(d)] of Fig. 2, we witness the well-documented transferability of adversarial examples between independent classical networks: Adversarial examples created by various white-box attacks on ConvNet transfer well to ResNet, despite it having an entirely different architecture. Being under white-box attack, the accuracy of the ConvNet itself falls quickly (the red lines). The accuracy of an independent classical network, ResNet, also falls quite rapidly (the blue lines),

demonstrating the transferability of adversarial attacks in the classical setting, and as seen in Figs. 2(a)–2(d).

Along the bottom row [(e)–(h)] of Fig. 2 we show, similarly, a successful transfer of adversarial examples generated by a white-box attack on QVC200 to an independently trained QVC500 network, although the transferability across quantum networks is relatively weak compared to the classical case. This is perhaps due to the fact that both QVC200 and QVC500 share the same architectural design, with the only difference being the circuit depths, or number of layers. In the future, it would be interesting to test transferability of quantum networks with different architectural designs, e.g., how well the attacks generated from a QVC network transfer to a quantum convolutional network. Nevertheless, Figs. 2(e)–2(h) show that the accuracy of QVC200 itself falls rapidly in the face of the adversarial examples, which are being generated specifically with respect to it (the black lines), as does that of QVC500 (the green lines), demonstrating the successful transferability between two QVC networks. While the transferability of adversarial examples between QVCs has to our knowledge not been studied in the literature, we consider this result unsurprising: QVC200 demonstrates a significant vulnerability to white-box adversarial attack, consistent with theoretical expectations [30] and previous empirical studies [29,34], and the successful transfer of these attacks to QVC500 demonstrates that the quantum networks are independently utilizing similar sets of nonrobust features, just as their classical counterparts do [37], although to a weaker extent.

Next, we turn our attention to the defense of classical networks against a quantum attack and vice versa. As exhibited by the results plotted in Figs. 2(a)–2(d), in the presence of PGD attacks generated from the classical ConvNet for various datasets, the QVC200 network demonstrates far superior robustness, retaining reasonable accuracies even in the face of very strong attacks, i.e., $\epsilon \geq 0.2$. Conversely, the results plotted in Figs. 2(e)–2(h) show a failure of the classical networks to maintain their accuracy in the face of adversarial examples generated by attacks on the QVC200 networks. Similar conclusions can be drawn from Figs. 10 and 11, where classical and quantum defenses are tested in the presence of adversarial examples generated by FGSM and AutoAttack. One might have expected, *a priori*, that the classical and quantum networks would learn different features of the data, and that therefore neither attack would transfer well across the classical/quantum divide. The success of the quantum adversarial examples in deceiving the classical networks, however, necessitates a more careful explanation, which will be discussed in the next section based on an analysis of the underpinning adversarial perturbations in the attacked images in both the classical and quantum realms.

B. Adversarial perturbations

To gain further insight into the success of the quantum networks in resisting the classical adversarial attacks, and conversely the failure of the classical networks to do the same, we plot the adversarial perturbations generated by PGD attacks on ConvNet, ResNet, and the 500 layer QVC on examples from the FMNIST and Celeb-A datasets in Fig. 3. Similar

examples from the MNIST and CIFAR datasets are shown in Fig. 12, along with perturbations from the FGSM and AutoAttack in Figs. 13 and 14. These perturbations highlight a very similar trend. The perturbations generated in the attacks on the classical networks display complicated high-frequency patterns, which are not readily human understandable. The fact that these “worst case” perturbations (from the networks’ perspective) take such a form is indicative of the fact that the classical neural networks have learnt to classify the images by identifying extremely complicated patterns in the data rather than following the data distributions, or in other words using a more human-like recognition of the large scale features in the images. While these features may be highly informative, the fact that classical networks tend to independently discover similar sets of them contributes to the vulnerability of the networks to transferred attacks, as the features can then be simultaneously targeted by an adversarial attack on any one of the classifiers individually. The inability of the classical attacks to transfer to the QVCs then indicates that the quantum models are utilizing a different set of features.

Indeed, the perturbations generated by attacking QVC200 demonstrate a markedly different story: in each case we can clearly identify meaningful information in the perturbation. For example, in the first column of Fig. 3, the perturbation is filling in the gap between the legs of the pants, which would indeed alter the true label of the image where the perturbation strong enough. Similarly, in the second column we see the attack adding sleeves to the T-shirt, and in the third column removing them. On the right-hand side, where the classification task is to determine whether or not the pictured person has black hair, we can again see meaningful perturbations generated by the QVC, either lightening or darkening the hair of the pictured person so as to change the label of the data sample image. The fact that, in the case of the QVC, the “worst case perturbations” do genuinely correspond to efficiently changing the true label of the sample, rather than incomprehensible noise as is the case for classical networks, implies that the QVC is learning more meaningful patterns in the data, and so is not easily fooled except by meaningful perturbations. This discovery is reminiscent of a similar phenomenon observed in the case of (classical) adversarially trained classifiers, where the adversarial perturbations of classifiers, which have been explicitly designed to be robust display more meaningful structure [37,38] than those which have not. We will discuss adversarial training in more detail in the next section. Remarkably, the ability of the QVCs to display similar behavior despite having undergone only standard (nonadversarial) training is highly interesting and is an indication of a kind of possible quantum advantage in ML tasks.

The qualitative differences between the classical and quantum perturbation landscapes explain both the success of the QVC adversarial examples in transferring to the classical networks, and the failure of the adversarial examples generated by attacking ConvNet to transfer to the QVCs. The classical networks rely on nonrobust, but highly informative features, which allow them to achieve high accuracy on clean data, but are highly susceptible to adversarial attacks. As the QVCs do not seem to rely on such features to the same extent, an adversarial attack, which targets them is limited in its ability

to fool the QVCs, but the attacks on QVC200 transfer well to the classical networks by virtue of their meaningful content.

These findings suggest extending the distinction drawn in the classical adversarial ML literature [37,38] between robust and nonrobust features to include a third category, classically intractable features. While these classically intractable features may themselves be robust or nonrobust, we expect them to be robust against classically generated attacks in practice, susceptible at most to perturbations generated by attacking a quantum model. The QVCs studied in this paper have discovered features consistent with this classically intractable but nonrobust category, being susceptible to attacks transferred from quantum models, but not classical ones. The investigation of quantum models, which are also robust against attacks transferred from quantum classifiers is an interesting direction for future work.

C. Adversarial training

The field of adversarial ML has seen a long battle between new attacks, proposed defense mechanisms, reformulated counterattacks, updated defense mechanisms, and so on [12,14,39,49–55]. Throughout this evolving landscape of attacks and defences, an enduring strategy on the defensive side has been that of adversarial training, in which adversarial examples are calculated and included in the training set of the ML model [11,12,16,39]. Despite its simplicity, and the lack of a rigorous guarantee of its success, adversarial training has been found to be effective in practice, and is considered to be one of the strongest methods for building adversarial robustness in classical ML [11,16]. In this paper, we assess the capability of adversarial training for quantum ML networks in the context of the FMNIST and MNIST datasets, with the training-time adversarial examples generated by the PGD, FGSM, and AutoAttack with $\epsilon = 0.1, 0.2, \text{ and } 0.3$. While a recent study considered adversarial training of a shallow QVC in the context of only binary classification on the MNIST dataset [29], we note that in this study a deep quantum ML network has undergone adversarial training with complex multiclass data, i.e., all ten classes of the FMNIST and MNIST datasets in Figs. 4(a)–4(d). Due to the extensive computational requirements, we leave the adversarial training of the 12-qubit classifiers employed for the CIFAR and Celeb-A datasets to a future study.

We plot in Figs. 4(a)–4(d) the results generated by creating adversarial attacks on ConvNet and QVC200, and applying to adversarially trained QVC200 networks. In Figs. 4(a) and 4(c), ConvNet is attacked in a white-box setting and the attack is transferred to traditionally trained QVC200 and adversarially trained QVCs (black-box scenario). We use superscripts to indicate the strength of the attacks used in the adversarial training, i.e., the labels QVC200^x indicate adversarial training with perturbations of l_∞ norm x , where $x \in \{0.1, 0.2, 0.3\}$. Our results show that adversarial training of QVC200 for both FMNIST and MNIST datasets makes only a minor difference in the accuracy of quantum networks. This is in stark contrast to classical ML where adversarial training has been shown to work quite well. Interestingly, it has been reported in the literature that adversarial training of classical ML also significantly reduces their clean learning accuracies (at $\epsilon = 0$), but

for QVCs no considerable reduction is observed in our paper, in particular for the FMNIST data. As the adversarial training is notoriously computationally expensive, our results indicate that in terms of QVC defense against classical attacks, expensive adversarial training is not justified. However, we acknowledge that these insights are based on the conducted simulations and further studies may be needed to establish this conclusion in a fully general set-up.

Next, we generate quantum attacks on QVC200 in a white-box setting and evaluate the accuracy of adversarially trained QVC200 networks in the presence of those quantum attacks. Figures 4(b) and 4(d) plot our results for FMNIST and MNIST datasets. In the quantum attack case, we show that the adversarial training improves the accuracy of quantum networks, although it does not completely restore it. We also note that the extent of adversarial training (0.1 or 0.2 or 0.3) only has a little impact, which is again in contrast to the classical ML literature where the level of adversarial training makes a big performance difference. The impact of QVC200 adversarial training in the presence of FGSM and AutoAttack is qualitatively similar to the PGD attacks as indicated from the comparison of results between Figs. 4 and 16. Finally, we also tested the resiliency of adversarially trained QVC200s in a black-box setting and the results are plotted in Fig. 17. Here, the attacks are generated in a white-box setting on QVC500 and transferred to QVC200 and adversarially trained QVC200s, which indicate that adversarial training only negligibly improves the accuracy.

To gain further insights into the adversarial training and its impact on the accuracy of the quantum ML networks, we show the perturbations generated by adversarial attacks on networks both with and without adversarial training in Figs. 4(e) and 4(f). Further examples of such images are shown in Figs. 18, 16(e), and 16(f) show examples for FGSM and AutoAttack. The overall conclusions from all these examples is the same that the perturbations generated by attacks on the adversarially trained QVC models display meaningful semantic content and comparable to those generated by attacking the (regularly trained) QVCs. The relatively weak impact of adversarial training on QVC200 accuracy is perhaps due to the fact that, even when undergoing standard training, the QVC networks learnt features quite similar to those they learnt when undergoing adversarial trainings. Also, the perturbation contents from quantum attacks are fundamentally different from the contents of classically generated perturbations (Fig. 3). Therefore, their susceptibility (or lack thereof) to a given classical attack is similar. This is in contrast to the classical networks, which will default to learning nonrobust features unless they are explicitly and rigorously trained in an adversarial setting (see for example Fig. 15), and therefore adversarial training makes a significant difference.

D. Prospects for future hardware implementation

The QVC results discussed in this paper are based on quantum simulations in a noiseless environment. Here we briefly discuss if/when the QVC networks benchmarked in our paper can be implemented on quantum processors. The quantum circuits corresponding to the QVC200 networks employed for the MNIST and FMNIST datasets are made up of 200

layers, with each layer consisting of 39 quantum gates (single qubit rotation gates and two-qubit CZ gates), as shown by the schematic diagram in Fig. 6(c). The amplitude encoding of the input data into an initial quantum state also requires a few thousand quantum gates. Overall, we estimate that a QVC200 network such as trained in our paper would consist of 10^3 – 10^4 quantum gates. The current generation of quantum devices are not capable of implementing quantum circuits with such deep circuit depths due to limitations imposed by relatively high level of noise or errors. However, the impact of noise can be analyzed to some extent by running quantum simulations with noise models. In Fig. 19, we have plotted the results obtained from QVC200 on the MNIST dataset in the presence of four different noise models: depolarization noise, amplitude damping noise, and bit-flip errors [56]. These noise models are described in the Appendix. Based on the noisy simulations, our results show that the accuracy is independent of noise strength if the noise model is depolarizing noise [31] and is only slightly reduced when the bit-flip error rate is increased to 0.5%. On the other hand, the presence of the amplitude damping noise severely impacts the accuracy of the developed quantum solution. Similarly, the use of the realistic hardware error models available within the Qiskit simulation environment, which are designed to simulate the noise found in real IBM devices, lead to low fidelity machine learning predictions, which were overwhelmed with errors and reduced to randomly guessing. This is anticipated as the quantum circuit is very deep and therefore the cumulative impact of the noise accumulated over several layers is quite significant. Despite the limitations imposed by noise, experimental work on a QAML implementation has already begun, with a first experimental demonstration just performed in the literature on superconducting quantum hardware [34], where MNIST and magnetic resonance imaging datasets were used for binary classification using 10-qubits and 26-layers, showing accuracies of above 90%. Although this study is at the proof-of-concept level, it is already indicating that near-term quantum devices should be able to handle quantum ML problems within the next few years. In our recent survey article [36], we have discussed the prospects for a surface code [57] based fault-tolerant implementation of QAML networks and based on our qualitative estimates, we anticipate that a fault-tolerant QAML implementation may be possible on 4000-qubit devices, which are expected to become available in 2025 [56]. The theoretical analysis presented in our paper has established a clear pathway for quantum advantage in QAML and its future fault-tolerant implementation will be a key milestone in the field of quantum computing, unlocking a quantum advantage for a range of real-world ML applications.

E. Adversarial attack detection

The results plotted in Figs. 2, 10, and 11 show that quantum networks defend remarkably well in the presence of classical adversarial attacks. Based on these simulations, we formulate an adversarial attack detection strategy, which will be highly useful in practical settings where a rapid detection of adversarial attacks is crucial for reliable ML based solutions. The flowchart diagram in Fig. 5 illustrates our attack detection technology, which is based on the hybrid operation of classical

and quantum networks working in conjunction to detect the presence of an adversarial attack. In this strategy, an attack is detected if the outcomes from a classical (ConvNet or ResNet) and a quantum network (QVC) disagree. We note that the proposed attack detection technology is only conceptually described here as it is limited by the clean accuracy of the quantum ML network and its practical application would require further development and optimization of QVC networks to attain clean learning accuracies (at $\epsilon = 0$) close to 1.0. We tested the working of the technology by generating 500 clean and 500 attacked images with $\epsilon = 0.1$ PGD attacks. This process was repeated 10 times. For the FMNIST case, comparing the results of ConvNet and QVC500 resulted in a false positive rate of $22.3 \pm 1.5\%$ and a true positive rate of $79.0 \pm 1.8\%$, whereas comparing ResNet and QVC500 led to a false positive rate of $22.9 \pm 1.7\%$ and a true positive rate of $67.2 \pm 1.6\%$. Here, a false positive means an attack is erroneously detected on a clean image and a true positive means an attack is detected on a perturbed image. These values are consistent with the relatively low clean accuracy of QVC500 with a significant improvement in performance expected if the clean accuracies of QVC500 can be improved in the future.

III. SUMMARY AND OUTLOOK

Our study has performed a systematic and detailed assessment of the QAML robustness in the presence of various attacks and based on an analysis for a variety of datasets. The presented results have revealed a number of important insights into the working of QAML and bridge a critical knowledge gap by answering key open questions such as: How well will QAML fare against strong classical adversarial attacks, to what extent will attacks transfer across the quantum and classical boundaries, and whether promising classical defense strategies such as adversarial training will work for QML networks. Our study will pave the way for future development of the QAML field and may lead to experimental demonstration of quantum advantage for ML tasks.

While our study provides many useful insights, which are important to fully understand the working of QAML, there still remains more work to be done for its practical implementation targeting real-world applications. Below, we highlight a few important areas for future development:

(i) We have found that the adversarial training of our QVCs produces relatively small improvements in accuracy, compared to those seen classically in the literature, where it significantly improves the accuracy of classical networks. It will be important to investigate this further, in particular by going beyond QVC architectures and by constructing and benchmarking more complex quantum networks such as quantum convolutional networks. This is of great significance because quantum ML networks have been found to be vulnerable against quantum attacks. If adversarial training could work as well for quantum networks as it works for classical networks, it may allow secure universal robustness of quantum networks against both classical and quantum attacks. Another important line of study would be to further optimize quantum ML networks to improve their clean learning accuracies.

TABLE I. The resource requirements for the different networks considered. In all cases the quantum variational classifiers require drastically less trainable parameters than their classical counterparts. Due to computational restrictions QVC1000 was not run on the 12-qubit datasets.

Network	MNIST and FMNIST			CIFAR-10 and Celeb-A		
	n_{qubits}	n_{qubits}	n_{layers}	n_{qubits}	$n_{\text{parameters}}$	n_{layers}
ConvNet		$\sim 10^6$	5		$\sim 10^6$	5
ResNet18		$\sim 10^7$	10		$\sim 10^7$	18
QVC200	10	6000	200	12	7200	200
QVC500	10	15000	500	12	18000	500
QVC1000	10	30000	1000			

(ii) While our paper has focused on image data, which has been strongly represented throughout the adversarial ML literature, the extent to which similar results may be found on distinct data types (e.g., audio, text, the environment of reinforcement learning agents, and fully quantum data, such as that arising in problems of the phase classification of quantum states [58]) also remains an interesting question for future investigations, which will enable a wider adaptation of QAML solutions.

(iii) The construction of deeper QVCs is seen to have limited effect on both the initial accuracy (Fig. 7) and the robustness to the adversarial attacks (see Fig. 20). This may be due to the generic form of the QVCs considered here not being particularly well suited to image classification. Analogously to how the original fully connected artificial neural networks were supplanted by convolutional neural networks for image classification and feature detection, these general purpose QVCs may be replaced by specialized quantum generalizations of convolutional neural networks [59] in order to construct models whose performance on image data exhibits more favorable scaling with model size. Moreover, beyond a certain depth, the QVCs will become maximally expressible and further increases to the depth will not allow for the finding of better solutions. This occurs at a point where the circuits contain $\mathcal{O}(2^{n_{\text{qubits}}})$ trainable parameters and the rank of the quantum Fisher information (QFI) matrix saturates [60]. Although we found that calculating the rank of the QFI matrix was computationally intractable for the circuits, which we used, the results of Fig. 7 suggest that we are operating in this regime.

(iv) The quantum circuits considered in this paper are quite large, consisting of 10 to 12 qubits, 200–1000 layers of single qubit rotation gates and entangling CZ gates, for a total of $\mathcal{O}(10^4)$ quantum gates. The accurate evaluation of such a circuit is beyond the capabilities of the noisy, nonerror corrected quantum computers available today. However, it may be possible using sophisticated error correction methods such as surface-code algorithms to implement the proposed QAML solutions on quantum processors in the near future. We qualitatively estimate that such implementations would require about 5000 or more qubits with error rates requirements below 0.5% or so. While quantum processors with such configurations may become available in the next few years [56], it would still require significant

development to transpile QAML circuits to error-corrected versions directly implementable on a quantum processor. Further optimization of the data loading step, which in our current paper is very expensive, should ease resource requirements for an implementation on a quantum processor, further pushing forward the practical realisability of QAML approaches.

In summary, vulnerability to adversarial examples has recently emerged as a serious issue confronting classical ML algorithms, raising ongoing concerns about their security and reliability when classifying data from untrusted sources. In the case of quantum ML, too, a similar susceptibility exists, and adversarial examples must be reckoned with if the field is to achieve its expected revolutionary potential. We address an important gap in the QAML literature by thoroughly studying the transferability of adversarial examples between classical and quantum neural networks in the context of common image datasets, discovering a surprising one-way resiliency between quantum and classical networks. The failure of the classical adversarial examples to transfer to the QVCs is characterized as a result of the classical and quantum networks learning different features of the various image datasets, and the quantum ML networks being

largely impervious to the specialized attacks, which targeted the precise features employed by the classical networks. Although demonstrated on various standard image datasets in this paper, the consequence of the differences between the architectures of the classical and quantum classifiers, which allow quantum ML to learn from classically intractable and robust features will only become more drastic in future large-scale quantum classifiers. In this case we would suspect even more strongly that classically generated adversarial examples will fail to transfer to these powerful quantum classifiers, which are relying on features that are invisible to the classical networks, and therefore cannot be targeted in an attack. Such a scenario will offer a new form of advantage in QML, orthogonal to the commonly anticipated gains in terms of speed or accuracy—quantum ML classifiers, which may not be necessarily more accurate than their classical counterparts, but exhibit superior robustness to adversarial attacks.

ACKNOWLEDGMENTS

The authors acknowledge useful discussions with Charles D. Hill, Shu-Lok Tsang, and Jia S. Low at the University

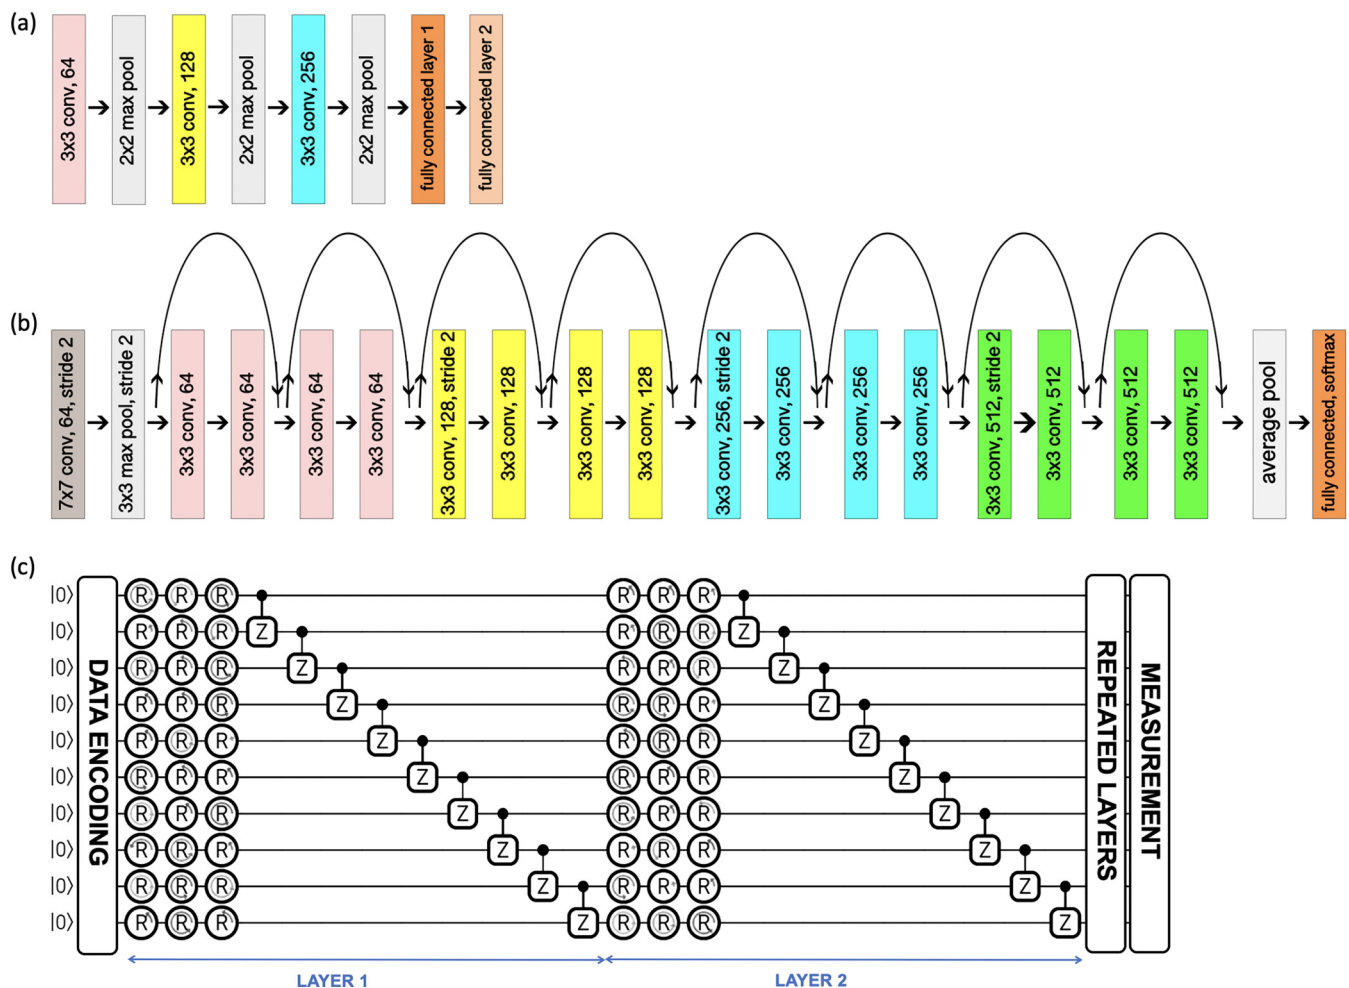


FIG. 6. Schematic depictions of the considered architectures are shown for (a) our simple convolutional neural network, (b) ResNet18 (as in Ref. [2]), (c) our QVCs. The QVC architecture is discussed in detail in Appendix A.

of Melbourne. M.T.W. acknowledges the support of the Australian Government Research Training Program Scholarship. S.M.E. is in part supported by Australian Research Council (ARC) Discovery Early Career Researcher Award (DECRA) No. DE220100680. The authors acknowledge the support from IBM Quantum Hub at the University of Melbourne and Australian Army Research through Quantum Technology Challenge (QTC22). The computational resources were provided by the National Computing Infrastructure (NCI) and Pawsey Supercomputing Center through National Computational Merit Allocation Scheme (NCMAS). The quantum circuit diagram shown in Fig. 1 was generated by using the Quantum User Interface (QUI) developed at the University of Melbourne.

The authors declare no competing financial or non-financial interests.

M.U. conceived the initial idea and supervised the project. M.T.W. developed classical and quantum ML frameworks and performed the simulations, with input from M.U. All authors contributed to the analysis of the data. M.U. and M.T.W. wrote the manuscript with input from all authors.

APPENDIX A: CLASSIFIER DETAILS

Here we provide various further details regarding the implementation of the classifiers and adversarial attacks em-

ployed in this paper. The total number of trainable parameters for each network may be found in Table I.

Classical network implementations. Our convolutional neural network begins with three layers of 3×3 filters, containing 64, 128, and 256 feature maps respectively, with 2×2 maxpooling and the ReLu activation function. These convolutional layers are followed by two fully connected layers, which also utilize the ReLu activation function. The architecture of ResNet18 is as described in [2]. The networks were implemented in Pytorch [61] and trained with the Adam optimiser and the cross-entropy loss function [62]. Further details of the network architectures are shown in Appendix A.

Quantum network implementations. Our QVCs follow a standard three step process for processing the input image data before outputting a predicted label. In the first step the images are encoded into a quantum state via the method of amplitude encoding [47]. In the second step the encoded images are processed by passing through a parameterised quantum circuit consisting of a variable number of layers, with each layer consisting of a parameterised rotation to each qubit followed by nearest neighbor CZ gates. In this paper we employ deep circuits with 200, 500, or 1000 layers. Finally, in the third step we measure the z expectation value of the first m qubits, where m is the number of classes in the particular classification problem being considered. The

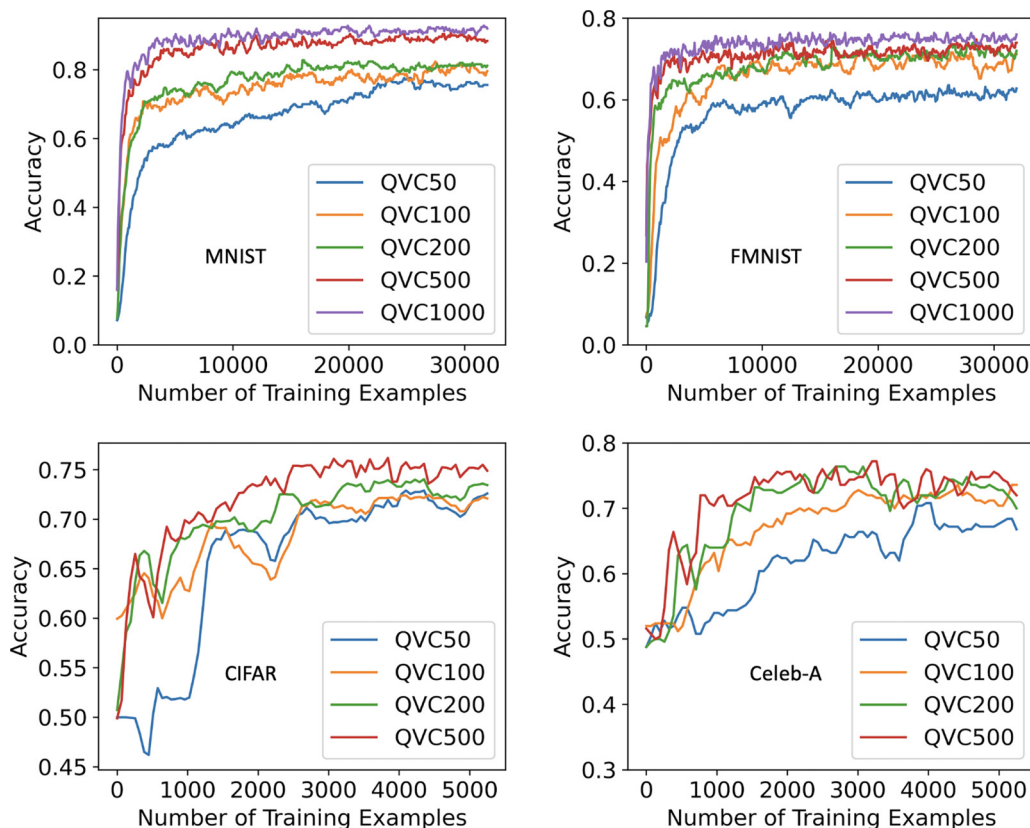


FIG. 7. The accuracy achieved by the quantum variational classifiers on test sets of 250 images from each dataset throughout the training process. We see consistent but modest gains in accuracy as a function of model size in the cases of MNIST and FMNIST (10 class classification problems), but less consistency in the cases of CIFAR-2 and Celeb-A (binary classification problems). Architectures more suited to image classification than QVCs (e.g., quantum convolutional neural networks [59]) may display better scaling behavior.

Learning Accuracies				
Datasets	CONV Net	RESNET	QVC200	QVC500
MNIST	1	0.992	0.832	0.916
FMNIST	0.932	0.9	0.74	0.74
CIFAR	0.944	0.84	0.764	0.812
Celeb-A	0.816	0.788	0.768	0.78

FIG. 8. Test accuracies achieved by classical and quantum networks for various datasets.

prediction of the QVC is defined to be the index of the highest of these values. The QVCs are trained via a full statevector simulation (i.e., using an infinite number of shots). Further details may be found in Appendix B. The QVCs were implemented in PennyLane [63] and trained with the Adam optimiser and the cross-entropy loss function [62]. On a real quantum computer, numerically evaluating gradients by making very small changes to the parameters of the circuit may be difficult to accomplish with high fidelity, and alternate techniques such as the param-shift rule may need to be considered [64].

Adversarial training. We employ adversarial training with the PGD attack with three gradient descent steps. Each batch of training examples consists of 50% clean and 50% adversarial examples.

Quantum noise. We test our QVCs in the presence of depolarization noise, amplitude damping, and bit flip noise.

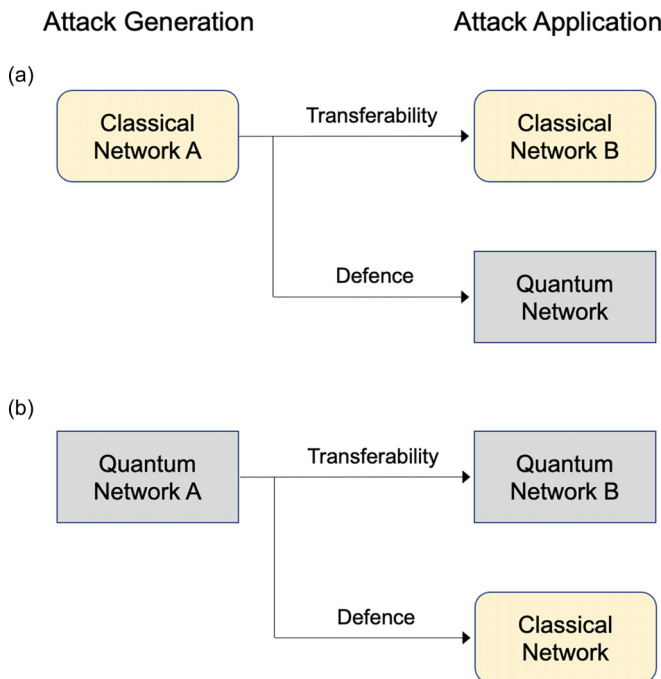


FIG. 9. A simple illustration to show our definition of transferability and defence. The attack is generated by either a classical network or a quantum network (Attack Generation). The effectiveness of the attack generated by a classical (quantum) network when applied to a different classical (quantum) network is defined as Transferability. The robustness of a quantum (classical) network against a classical (quantum) attack is defined as Defence.

The results are shown in Fig. 19. Depolarization noise: In this model, the state of the system is replaced with the maximally mixed state with probability p , which indicates the strength of the noise. Amplitude damping: This model is a quantum channel, which gives a simple model for relaxation of an excited state to the ground state. In the qubit case, with some probability there is a spontaneous transition $|0\rangle \rightarrow |1\rangle$. The lifetime is governed by a parameter γ , which determines the strength of the noise. Bit-flip: In this model, an X gate is applied with probability p , which determines the strength of the noise.

APPENDIX B: ADVERSARIAL MACHINE LEARNING

We begin with a brief review of the aspects of adversarial machine learning relevant to our paper, drawing on both the classical [7] and quantum [29] supervised machine learning literature. The key discovery of adversarial machine learning is that standard ML frameworks are highly susceptible to being deceived by subtle, malicious tampering with their input data [6–9]. The results of such tampering, adversarial examples, can be readily produced by taking a data sample (which for a high performing ML method will likely be correctly classified) and attempting to find a tiny perturbation, which when added to the data causes a misclassification. By insisting that the perturbation is small one guarantees that the true label of the constructed adversarial example is the same as the label of the original datapoint, despite the model classifying them differently if the attack is successful.

Concretely, suppose that we are attempting to train a classifier to label data from a set \mathcal{X} , with corresponding labels from another set \mathcal{Y} . The goal of supervised machine learning is then to learn a parameterised function $C_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ in order to minimize the empirical loss obtained on a given training set consisting of pairs (x_i, y_i) of labeled examples,

$$\theta = \operatorname{argmin}_{\theta \in \Theta} \frac{1}{N} \sum_i \mathcal{L}(C_\theta(x_i), y_i) \tag{B1}$$

where Θ is the set of possible parameter values, and \mathcal{L} is a chosen loss function (e.g., the cross-entropy loss [62]). Alternately, in the setting of adversarial machine learning, one is given a trained classifier C_θ and an input sample x , and seeks to construct an adversarial perturbation δ_{adv} by maximizing the loss function,

$$\delta_{\text{adv}} = \operatorname{argmax}_{\delta \in \Delta} \mathcal{L}(C_\theta(x + \delta), y) \tag{B2}$$

where Δ is the set of allowable (i.e., acceptably small) perturbations. While the optimization problem of Eq. (B2) may be highly nonconvex and analytically intractable, many strong strategies (“types of attacks”) have emerged for tackling it in practice. In this paper we consider three standard attacks from the classical ML literature, namely projected gradient descent (PGD) [44], the fast gradient sign method (FGSM) [45], and AutoAttack [46].

One of the first white-box algorithms proposed to attack deep networks for image classification, FGSM obtains the adversarial perturbation δ by computing the contribution of each element of the input vector x to the loss function, restricted by the ℓ_∞ norm, by taking

$$\delta = \epsilon \operatorname{sign}(\nabla_x \mathcal{L}(C(x; \theta^*), y)) \tag{B3}$$

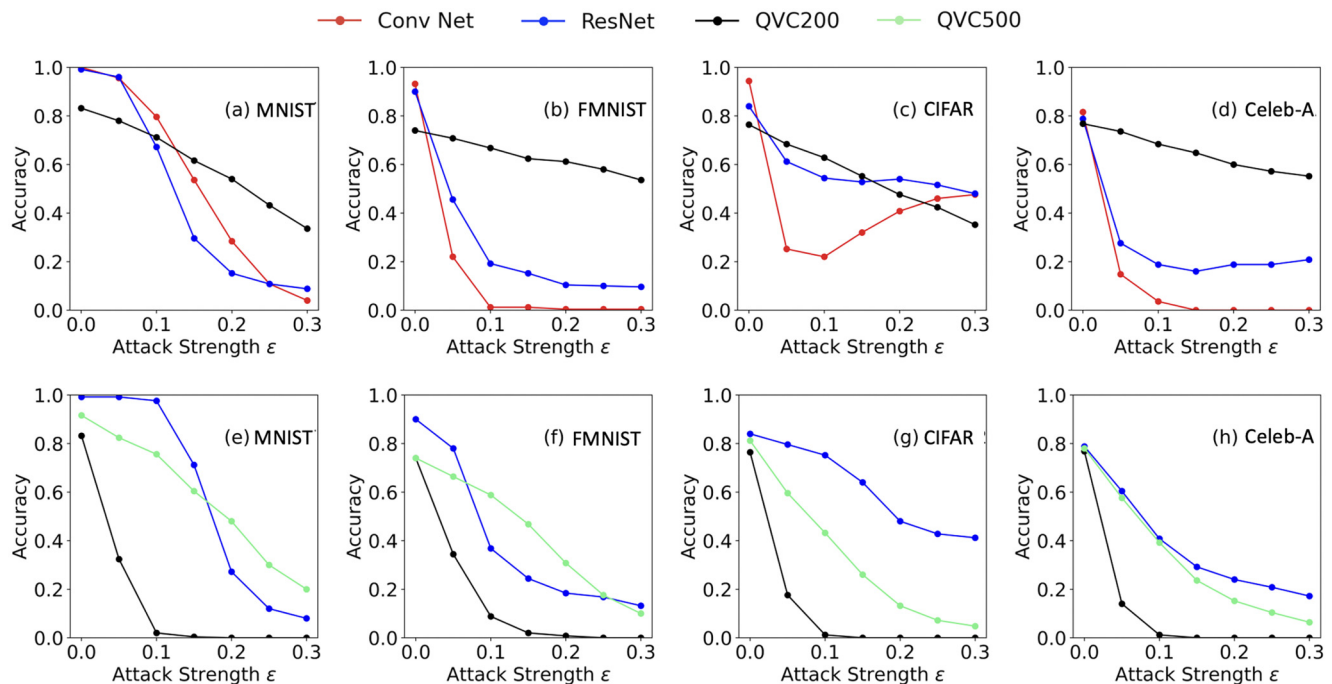


FIG. 10. The accuracy achieved by classical and quantum networks on sets of 250 adversarially attacked test images from each of the considered datasets in the cases of white-box FGSM attacks on the convolutional network (top row), and 200 layer quantum variational classifier (QVC200, bottom row) as a function of attack strength (measured with the l_∞ norm). In both cases we see the accuracy of the network under attack decrease sharply. The tendency of the accuracy of the independent networks to also decrease is a manifestation of the transferability of adversarial examples—they are typically capable of fooling even networks, which they were not explicitly designed to attack. We see an exception to this in the top row, with the quantum classifier usually resisting the attacks generated with respect to the convolutional neural network.

Despite its simplicity, FGSM can consistently and efficiently generate perturbations that cause untargeted misclassifications with a high confidence in deep networks, which is hypothesized to be due to the linearity of typical neural networks [45]. Another popular attack, PGD is essentially an iterated extension of FGSM. Starting uniformly randomly in the space of allowable perturbations Δ , it computes an adversarial perturbation iteratively by taking

$$\delta'_0 = \text{Unif}(\Delta), \delta'_{t+1} = \mathcal{P}(\delta'_t + \alpha \nabla_{\delta} \mathcal{L}(C(\mathbf{x} + \delta'_t; \theta^*), y)) \tag{B4}$$

given some step size α . After each gradient step the perturbation is projected back onto the allowable region Δ by the projector \mathcal{P} (such as clipping individual values between ϵ under the l_∞ norm). PGD is an important attack as it is conjectured that it is a universal first-order adversary in the sense that no other polynomially bounded attacks using first-order gradients will find a significantly better perturbation [44]. Finally, AutoAttack is an ensemble method in which multiple attacks are employed, with an input that resists an attack being fed into the next attack in the list. Our implementation follows Ref. [46].

An important classification of adversarial attacks is into white-box and black-box attacks. In a white-box attack, the adversary is assumed to have full access to the network under attack, including familiarity with its architecture and the values of its weights and biases. As a result, the adversary

possesses the ability to differentiate with respect to an input to the network, and therefore carry out optimization based attacks such as the PGD, FGSM, and AutoAttack considered here. While such attacks generally have a devastating effect on the accuracy of a network, the assumption of such intimate familiarity with the target network may not always be satisfied. A more realistic case is that where the adversary has access to the target network only through submitting queries, and is therefore forced to attempt a black-box attack. The feasibility of black-box attacks stems largely from the discovery of the remarkable transferability of adversarial examples between various ML frameworks: adversarial examples constructed to attack a specific target tend to transfer to entirely independent networks, deceiving them as well [7,37–39]. This discovery alleviates the need for direct access to a network in order to reliably attack it; the adversary may construct a network of their own, perform a white-box attack on it, and then simply submit the generated adversarial examples to the target network. The transferability of adversarial examples between classical and quantum networks is the main topic of study in this paper.

APPENDIX C: QUANTUM MACHINE LEARNING

The quantum machine learning (QML) models which we will consider throughout this paper belong to the category of quantum variational classifiers (QVCs). As in the classical case, these models will be parameterised functions $\mathcal{A}_\theta: \mathcal{X} \mapsto \mathcal{Y}$, where we consider an input set \mathcal{X} with associated labels




FIG. 11. The accuracy achieved by classical and quantum networks on sets of 250 adversarially attacked test images from each of the considered datasets in the cases of white-box AutoAttack on the convolutional network (top row), and 200 layer quantum variational classifier (QVC200, bottom row) as a function of attack strength (measured with the l_∞ norm). In both cases we see the accuracy of the network under attack decrease sharply. The tendency of the accuracy of the independent networks to also decrease is a manifestation of the transferability of adversarial examples—they are typically capable of fooling even networks which they were not explicitly designed to attack. We see an exception to this in the top row, with the quantum classifier largely resisting the attacks generated with respect to the convolutional neural network.

from a further set \mathcal{Y} and denote a QVC parameterised by parameters θ as \mathcal{A}_θ . Our QVCs follow a standard three step process for processing the input image data before outputting

a predicted label. In the first step the images are encoded into a quantum state. Due to the strong limitations on the numbers of qubits currently available to quantum computers and




FIG. 12. The adversarial perturbations generated by $\epsilon = 0.1$ PGD attacks on the convolutional neural network, ResNet18 and the 500 layer quantum variational classifier (QVC500) are shown for several examples from the MNIST and CIFAR databases.




FIG. 13. The adversarial perturbations generated by $\epsilon = 0.1$ FGSM attacks on the convolutional neural network, ResNet18 and the 500 layer quantum variational classifier (QVC500) are shown for several examples from the FMNIST and Celeb-A databases.

simulators, and our need to encode high dimensional image data, we employ the method of amplitude encoding [47]. Having represented the image by a vector x containing its pixel values, amplitude encoding is the mapping

$$x \mapsto \sum_{i=0}^{2^n} x_i |i\rangle \quad (C1)$$

where the set $\{|i\rangle\}_{i=0}^{2^n-1}$ forms the computational basis of the Hilbert space. As amplitude encoding makes use of the entire Hilbert space, which has a dimension exponentially large in the number of qubits, it can encode a vector $x \in \mathbb{R}^m$ into $\lceil \log_2(m) \rceil$ qubits. This extreme compression allows us to encode images into a manageable number of qubits; in the case of the MNIST and FMNIST datasets, which consist of 28×28 greyscale images, only 10 qubits are needed, while for




FIG. 14. The adversarial perturbations generated by $\epsilon = 0.1$ FGSM attacks on the convolutional neural network, ResNet18 and the 500 layer quantum variational classifier (QVC500) are shown for several examples from the MNIST and CIFAR databases.

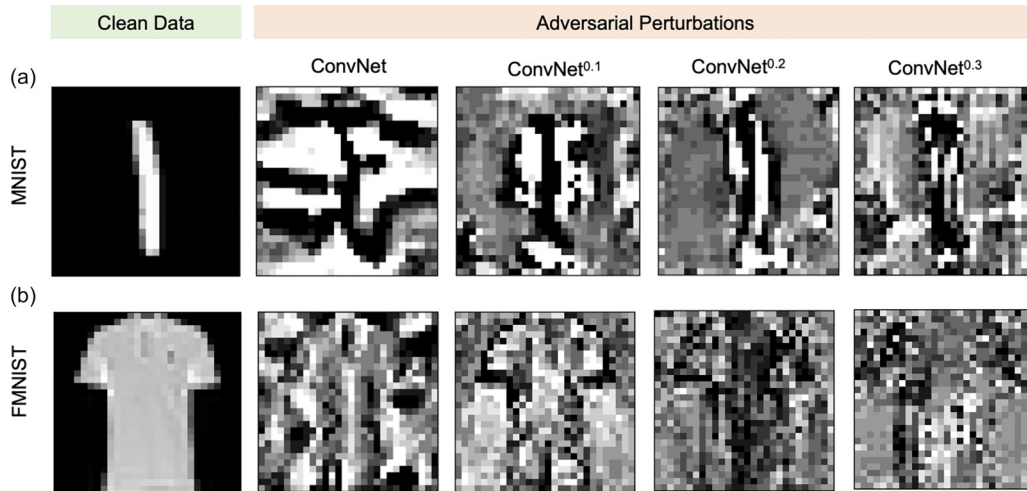


FIG. 15. [(a),(b)] The adversarial perturbations generated by PGD attacks on ConvNet, ConvNet^{0.1}, ConvNet^{0.2}, and ConvNet^{0.3} are shown for sample images from the MNIST and FMNIST datasets. While the perturbation for ConvNet exhibits highly abstract features, which are incomprehensible to humans, the perturbations from the adversarial training show hints of systematic features, reminiscence of quantum networks.

the 3×32×32 RGB images from the CIFAR-10 and CelebA datasets, 12 qubits are required. In both cases, zeros are appended to the input vectors in order to make their length a power of two. Having loaded the images into the quantum model in the first stage, in the second stage they are processed by passing through a parameterised quantum circuit, the parameters of which are updated during training to minimize the average cross entropy loss on the training dataset as in Eq. (B1). We choose parameterised circuits consisting of a variable number of layers, with each layer consisting of a parameterised rotation to each qubit followed by nearest neighbor CZ gates. In this paper we employ deep circuits with 200, 500, or 1000 layers. Finally, in the third stage a set of measurements are performed in order to determine the predicted output label. In our case we measure the z expectation value of the first m qubits, where m is the number of classes in the particular classification problem being considered (10 for

MNIST and FMNIST, two for CIFAR-10 and CelebA). Given an input \mathbf{x} , then, the final output of the quantum classifier is

$$\mathcal{A}_\theta(\mathbf{x}) = \operatorname{argmax}_{k \leq m} \operatorname{Tr}[U_\theta(\mathbf{x})\sigma_k^z] \quad (C2)$$

where we denote the unitary representing the action of the parameterised quantum circuit on the input \mathbf{x} by $U_\theta(\mathbf{x})$. The procedure is depicted in Fig. 1.

The accuracies obtained by the various QVCs on the considered datasets are shown in Fig. 7.

APPENDIX D: FURTHER DETAILS

This Appendix contains Figs. 6–20 supporting various aspects of this paper, including plots of adversarial accuracies and perturbations omitted from the main text, diagrams of

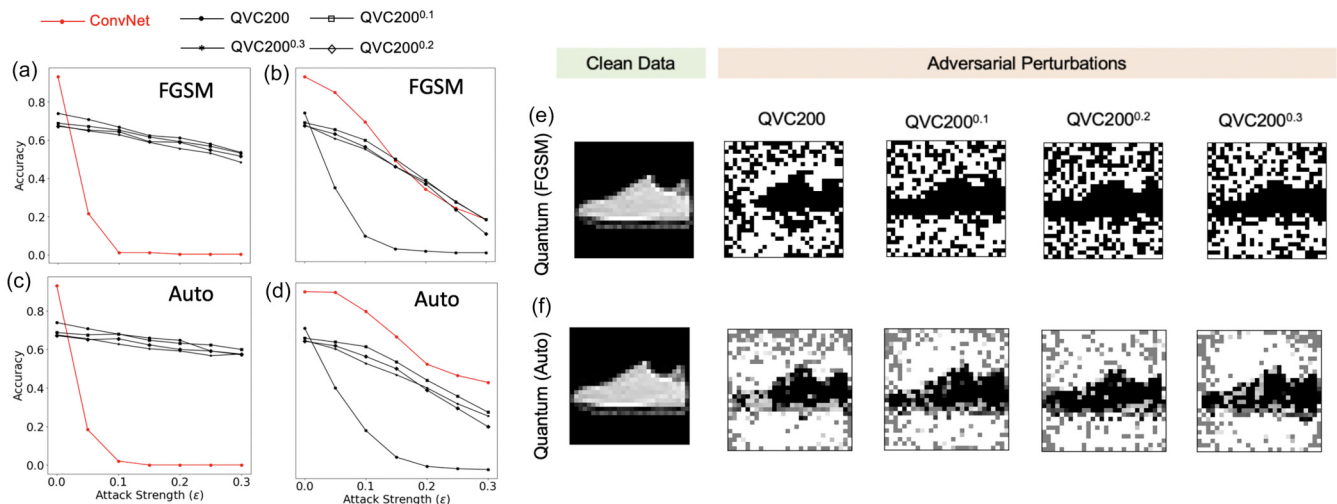


FIG. 16. As in Fig. 4, but using the FGSM and AutoAttack for FMNIST dataset.




FIG. 17. The accuracies achieved by quantum ML networks are plotted as a function of attack strength on a set of 250 adversarially attacked test images from the FMNIST dataset in the cases of white-box attacks on the QVC500 network. The attack is applied to QVC200 and adversarially trained QVC200 networks [QVC200^{0.1}, QVC200^{0.2}, and QVC200^{0.3} where adversarial training is performed with PGD attacks of (l_∞) strength 0.1, 0.2, and 0.3 respectively]. The adversarial training makes a negligible difference to the accuracy of the quantum network.




FIG. 18. The adversarial perturbations generated by $\epsilon = 0.1$ PGD attacks on QVC200 networks, with and without adversarial training, for a few examples from the FMNIST dataset. The QVCs display large-scale features whose meaning is often understandable, unlike ConvNet with standard training.

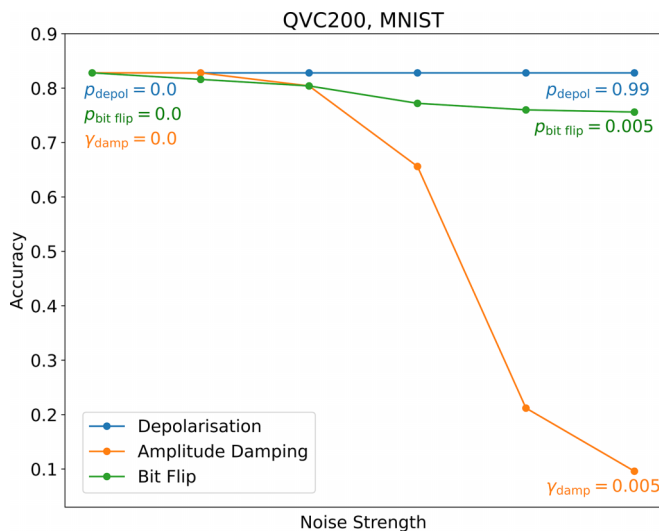


FIG. 19. The test accuracies of QVC200 on the MNIST dataset under various levels of QC noise in the simulation. The description of QC noise models is provided in the Methods section. We sample $p_{\text{bit flip}}, \gamma_{\text{damp}} \in \{0, 0.01, 0.02, 0.03, 0.04, 0.05\}$, and $p_{\text{depol}} \in \{0, 0.2, 0.4, 0.6, 0.8, 0.99\}$. The robustness to depolarization noise is perfect for all $p_{\text{depol}} < 1$, as proven for this class of QML model in Ref. [31].

the networks employed, the validation accuracy of the QVCs throughout the training process, and the performance of the QVCs in the presence of noise.

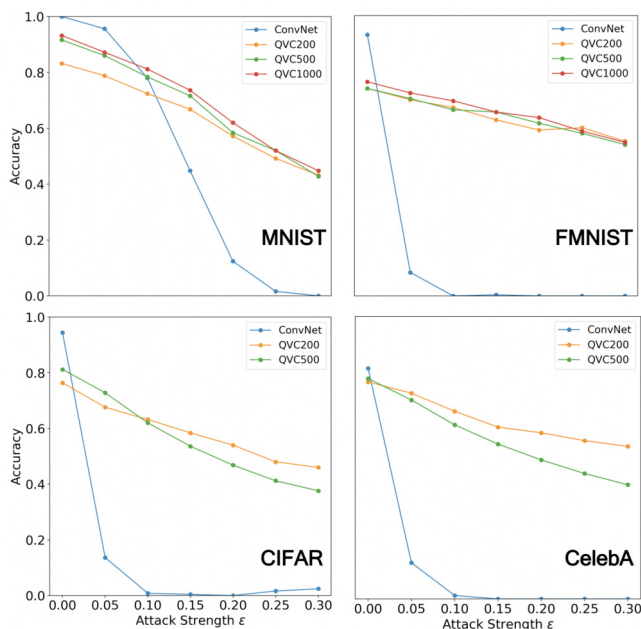


FIG. 20. The accuracies of QVCs of various depths on images generated by a PGD attack on ConvNet. We find that increasing the depth of the quantum networks has a limited effect on both clean accuracy and adversarial robustness.

- [1] Y. LeCun, Y. Bengio, and G. Hinton, Deep learning, *Nature (London)* **521**, 436 (2015).
- [2] 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* (IEEE, New York, 2016), pp. 770–778.
- [3] D. W. Otter, J. R. Medina, and J. K. Kalita, A survey of the usages of deep learning for natural language processing, *IEEE Trans. Neural Netw. Learn. Syst.* **32**, 604 (2020).
- [4] J. Jumper *et al.*, Highly accurate protein structure prediction with alphafold, *Nature (London)* **596**, 583 (2021).
- [5] T. Fösel, M. Y. Niu, F. Marquardt, and L. Li, Quantum circuit optimization with deep reinforcement learning, [arXiv:2103.07585](https://arxiv.org/abs/2103.07585).
- [6] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, Evasion attacks against machine learning at test time, in *Machine Learning and Knowledge Discovery in Databases*, edited by H. Blockeel *et al.* (Springer, Berlin, 2013), pp. 387–402.
- [7] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, R. Fergus, Intriguing properties of neural networks, [arXiv:1312.6199](https://arxiv.org/abs/1312.6199).
- [8] L. Huang, A. D. Joseph, B. Nelson, B. I. Rubinstein, and J. D. Tygar, Adversarial machine learning, in *Proceedings of the 4th ACM Workshop on Security and Artificial Intelligence*, AISec '11 (ACM Press, New York, NY, USA, 2011), pp. 43–58.
- [9] A. Kurakin, I. Goodfellow, and S. Bengio, Adversarial machine learning at scale, [arXiv:1611.01236](https://arxiv.org/abs/1611.01236).
- [10] S. Goldwasser, M. P. Kim, V. Vaikuntanathan, and O. Zamir, Planting undetectable backdoors in machine learning models, in *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2022), pp. 931–942.
- [11] E. Wong, L. Rice, and J. Z. Kolter, Fast is better than free: Revisiting adversarial training, [arXiv:2001.03994](https://arxiv.org/abs/2001.03994).
- [12] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, Towards deep learning models resistant to adversarial attacks, [arXiv:1706.06083](https://arxiv.org/abs/1706.06083).
- [13] I. Goodfellow, P. McDaniel, and N. Papernot, Making machine learning robust against adversarial inputs, *Commun. ACM* **61**, 56 (2018).
- [14] D. J. Miller, Z. Xiang, and G. Kesidis, Adversarial learning targeting deep neural network classification: A comprehensive review of defenses against attacks, *Proc. IEEE* **108**, 402–433 (2020).
- [15] J. Cohen, E. Rosenfeld, and Z. Kolter, Certified adversarial robustness via randomized smoothing, in *International Conference on Machine Learning* (PMLR, 2019), pp. 1310–1320.
- [16] T. Bai, J. Luo, J. Zhao, B. Wen, and Q. Wang, Recent advances in adversarial training for adversarial robustness, [arXiv:2102.01356](https://arxiv.org/abs/2102.01356).
- [17] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [18] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [19] A. M. Childs, R. Kothari, and R. D. Somma, Quantum algorithm for systems of linear equations with exponentially improved dependence on precision, *SIAM J. Comput.* **46**, 1920 (2017).
- [20] K. Beer, D. Bondarenko, T. Farrelly, T. J. Osborne, R. Salzmann, D. Scheiermann, and R. Wolf, Training deep quantum neural networks, *Nat. Commun.* **11**, 808 (2020).
- [21] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [22] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, *Quantum Sci. Technol.* **2**, 045001 (2017).
- [23] P.-L. Dallaire-Demers, and N. Killoran, Quantum generative adversarial networks, *Phys. Rev. A* **98**, 012324 (2018).
- [24] A. Kehoe, P. Wittek, Y. Xue, and A. Pozas-Kerstjens, Defence against adversarial attacks using classical and quantum-enhanced Boltzmann machines, *Mach. Learn. Sci. Technol.* **2**, 045006 (2021).
- [25] J. Heredge, C. Hill, L. Hollenberg, and M. Sevior, Quantum support vector machines for continuum suppression in B meson decays, *Comput. Softw. Big Sci.* **5**, 1 (2021).
- [26] M. Srikumar, C. D. Hill, and L. C. Hollenberg, A kernel-based quantum random forest for improved classification, [arXiv:2210.02355](https://arxiv.org/abs/2210.02355).
- [27] P. Jurcevic *et al.*, Demonstration of quantum volume 64 on a superconducting quantum computing system, *Quantum Sci. Technol.* **6**, 025020 (2021).
- [28] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Challenges and opportunities in quantum machine learning, *Nat. Comput. Sci.* **2**, 567 (2022).
- [29] S. Lu, L.-M. Duan, and D.-L. Deng, Quantum adversarial machine learning, *Phys. Rev. Res.* **2**, 033212 (2020).
- [30] N. Liu and P. Wittek, Vulnerability of quantum classification to adversarial perturbations, *Phys. Rev. A* **101**, 062331 (2020).
- [31] Y. Du, M.-H. Hsieh, T. Liu, D. Tao, and N. Liu, Quantum noise protects quantum classifiers against adversaries, *Phys. Rev. Res.* **3**, 023153 (2021).
- [32] J. Guan, W. Fang, and M. Ying, Robustness verification of quantum classifiers, in *International Conference on Computer Aided Verification* (Springer International Publishing, Cham, 2021), pp. 151–174.
- [33] M. Weber, N. Liu, B. Li, C. Zhang, and Z. Zhao, Optimal provable robustness of quantum classification via quantum hypothesis testing, *npj Quantum Inf.* **7**, 1 (2021).
- [34] W. Ren, W. Li, S. Xu, K. Wang, W. Jiang, F. Jin, X. Zhu, J. Chen, Z. Song, P. Zhang *et al.*, Experimental quantum adversarial learning with programmable superconducting qubits, *Nat. Comput. Sci.* **2**, 711 (2022).
- [35] H. Liao, I. Convy, W. J. Huggins, and K. B. Whaley, Robust in practice: Adversarial attacks on quantum machine learning, *Phys. Rev. A* **103**, 042427 (2021).
- [36] M. T. West, S.-L. Tsang, J. S. Low, C. D. Hill, C. Leckie, L. C. L. Hollenberg, S. M. Erfani, and M. Usman, Towards quantum enhanced adversarial robustness in machine learning, *Nat. Mach. Intell.* (2023), <https://doi.org/10.1038/s42256-023-00661-1>.
- [37] A. Ilyas *et al.*, Adversarial examples are not bugs, they are features, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Red Hook, NY, 2019).

- [38] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry, Robustness may be at odds with accuracy, [arXiv:1805.12152](https://arxiv.org/abs/1805.12152).
- [39] I. J. Goodfellow, J. Shlens, and C. Szegedy, Explaining and harnessing adversarial examples, [arXiv:1412.6572](https://arxiv.org/abs/1412.6572).
- [40] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [41] H. Xiao, K. Rasul, and R. Vollgraf, Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms, [arXiv:1708.07747](https://arxiv.org/abs/1708.07747).
- [42] A. Krizhevsky, Learning multiple layers of features from Tiny Images, Master's thesis, Department of Computer Science, University of Toronto, 2009.
- [43] S. Yang, P. Luo, C.-C. Loy, and X. Tang, From facial parts responses to face detection: A deep learning approach, in *Proceedings of the IEEE International Conference on Computer Vision* (IEEE, New York, 2015), pp. 3676–3684.
- [44] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, Towards deep learning models resistant to adversarial attacks, [arXiv:1706.06083](https://arxiv.org/abs/1706.06083).
- [45] I. J. Goodfellow, J. Shlens, and C. Szegedy, Explaining and harnessing adversarial examples, [arXiv:1412.6572](https://arxiv.org/abs/1412.6572).
- [46] F. Croce and M. Hein, Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks, in *International Conference on Machine Learning* (PMLR, 2020), pp. 2206–2216.
- [47] R. LaRose, and B. Coyle, Robust data encodings for quantum classifiers, *Phys. Rev. A* **102**, 032420 (2020).
- [48] R. Dilip, Y. J. Liu, A. Smith, and F. Pollmann, Data compression for quantum machine learning, *Phys. Rev. Res.* **4**, 043007 (2022).
- [49] A. Athalye, N. Carlini, and D. Wagner, Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples, in *International Conference on Machine Learning* (PMLR, 2018), pp. 274–283.
- [50] N. Carlini and D. Wagner, Adversarial examples are not easily detected: Bypassing ten detection methods, in *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security* (ACM Press, New York, 2017), pp. 3–14.
- [51] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, Countering adversarial images using input transformations, [arXiv:1711.00117](https://arxiv.org/abs/1711.00117).
- [52] J. Buckman, A. Roy, C. Raffel, and I. Goodfellow, Thermometer encoding: One hot way to resist adversarial examples (2018), <https://openreview.net/pdf?id=S18Su--CW>.
- [53] R. Feinman, R. R. Curtin, S. Shintre, and A. B. Gardner, Detecting adversarial samples from artifacts, [arXiv:1703.00410](https://arxiv.org/abs/1703.00410).
- [54] H. Salman, G. Yang, J. Li, P. Zhang, H. Zhang, I. Razenshteyn, and S. Bubeck, Provably robust deep learning via adversarially trained smoothed classifiers, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Red Hook, NY, 2019).
- [55] H. Zhang *et al.*, Theoretically principled trade-off between robustness and accuracy, in *International Conference on Machine Learning* (PMLR, 2019), pp. 7472–7482.
- [56] IBM Quantum road-map: <https://research.ibm.com/blog/ibm-quantum-roadmap-2025>.
- [57] S. Gicev, L. C. Hollenberg, and M. Usman, A scalable and fast artificial neural network syndrome decoder for surface codes, [arXiv:2110.05854](https://arxiv.org/abs/2110.05854).
- [58] M. Caro *et al.*, Generalization in quantum machine learning from few training data, *Nat. Commun.* **13**, 4919 (2022).
- [59] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [60] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, Theory of overparametrization in quantum neural networks, [arXiv:2109.11676](https://arxiv.org/abs/2109.11676).
- [61] A. Paszke *et al.*, Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32* (Curran Associates, Red Hook, NY, 2019), pp. 8024–8035.
- [62] Z. Zhang and M. Sabuncu, Generalized cross entropy loss for training deep neural networks with noisy labels, in *Advances in Neural Information Processing Systems 31* (Curran Associates, Red Hook, NY, 2018).
- [63] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. Akash Narayanan, A. Asadi, J. Miguel Arrazola *et al.*, PennyLane: Automatic differentiation of hybrid quantum-classical computations, [arXiv:1811.04968](https://arxiv.org/abs/1811.04968).
- [64] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019).