

Generalization bounds in hybrid quantum-classical machine learning models

Tongyan Wu^{1,*}, Amine Bentellis², Alona Sakhnenko¹, and Jeanette Miriam Lorenz^{3,†}

¹*Fraunhofer Institute for Cognitive Systems IKS and Technical University of Munich, Munich, Germany*

²*Fraunhofer Institute for Cognitive Systems IKS, Munich, Germany*

³*Fraunhofer Institute for Cognitive Systems IKS and Ludwig-Maximilian University, Munich, Germany*



(Received 26 May 2025; accepted 24 December 2025; published 17 February 2026)

Hybrid quantum-classical models aim to harness the strengths of both quantum computing and classical machine learning, but their practical potential remains poorly understood. In this work, we develop a unified mathematical framework for analyzing generalization in hybrid models, offering insight into how these systems learn from data. We establish a generalization bound of the form $\tilde{O}(\frac{\alpha^k}{\sqrt{N}}(k^{\frac{3}{2}}\sqrt{mn} + \sqrt{T \log T}))$ for N training data points, T trainable quantum gates, n -dimensional quantum circuit output, and k bounded linear layers $\|F_i\|_F \leq \alpha$, where $i = 1, \dots, k$ and $F \in \mathbb{R}^{m \times n}$ interspersed with activation functions. This generalization bound decomposes into quantum and classical contributions, providing a theoretical framework to separate their influence and clarifying their interaction. Alongside the bound, we highlight conceptual limitations of applying classical statistical learning theory in the hybrid setting and suggest promising directions for future theoretical work.

DOI: [10.1103/6ckt-hlh8](https://doi.org/10.1103/6ckt-hlh8)

I. INTRODUCTION

Quantum-classical hybrid machine learning models have emerged as a promising approach in an era where fully fault-tolerant, large-scale quantum computers remain out of reach. Currently available devices impose significant restrictions on width and depth of quantum circuits, making hybridization of quantum and classical computation an ideal candidate for exploiting early capabilities of quantum computing. In the current framework, classical optimization methods act as a complement for shallow-depth circuits. As the quantum technological stack matures into the fault-tolerant regime, hybrid algorithms will evolve as well, allowing a careful interleaving of classical and quantum models, leveraging the strength of both of these resources. Early proof-of-concept hybrid algorithms show promising results [1–8]; however, a solid theoretical understanding of the power of these architectures is lacking. Gaining such insights could shed light on the differences in learning capacity between quantum and classical algorithms.

There are several theoretically interesting aspects of a learning model's behavior, e.g., resource demands [1], sampling complexity [8], and learnability [9]. One of the most sought-after metrics is generalization, which allows one to predict the model's performance beyond the available dataset. Classical approaches in statistical learning theory establish

uniform upper bounds on the generalization error based on both the training set size and the inherent complexity of the model. In this context, complexity can be characterized through measures such as covering numbers and Rademacher complexity [10], which capture the richness of a function class associated with the hypothesis class¹ of the model. Several works have provided rigorous bounds on the generalization error for classical machine learning models [11–13] as well as for fully quantum learning models [1–3]. However, despite the growing popularity of hybrid models, the conditions under which these hybrid architectures can generalize accurately remain largely unexplored.

Inspired by the remarkable success of neural networks (NNs) in classical machine learning, the quantum machine learning community has pursued a similar line of research. In particular, layerwise organized quantum circuits with parametrizable gates that mimic classical feedforward architecture are known as quantum NNs (QNNs). A diverse zoo of models has emerged inspired by classical ideas, such as hybrid Boltzmann machines [5], hybrid autoencoders [6] and hybrid convolutional NNs [7]. QNNs fall under the category of quantum machine learning models (QMLMs) [14,15], which in general consist of a parametrized quantum circuit that is being trained on classical or quantum data.

In this paper, we derive and prove a generalization bound for hybrid (quantum-classical) machine learning models (hybrid QMLMs) that consist of general QMLMs combined with classical NNs. In the hybrid quantum-classical setting, a model's hypothesis class becomes richer due to integrating quantum with classical machine learning components. In our proof, we first separately derive complexity measures for the

*Contact author: tongyan.wu@tum.de

†Contact author: jeanette.miriam.lorenz@iks.fraunhofer.de

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

¹A concept that represents a set of all possible functions that are representable by a chosen algorithm.

hypothesis classes corresponding to the quantum and classical components. Due to the richness of the hybrid model's hypothesis class, we use covering numbers to quantify the complexity of the hypothesis class geometrically, rather than a direct worst-case analysis using Rademacher complexity. After establishing covering number bounds, we combine them to characterize the overall complexity of the hybrid hypothesis class. Using Dudley's entropy integral, we obtain generalization bounds that quantify a model's learning ability from finite data. A key implication of our results is that the complexity of a hybrid model that has minimal classical layers will perform similarly to a QMLM, providing theoretical support for the practical viability and design of quantum-assisted learning algorithms. As such, designing hybrid models with shallow classical components but expressive quantum circuits can achieve strong generalization while reducing implementation overhead, offering a principled framework for balancing quantum and classical resources in real-world QMLM architectures.

This paper is organized as follows: In Sec. II, we justify the chosen QMLM architecture and provide an overview of previous work on generalization bounds in both quantum and classical contexts. In Sec. III, we outline the fundamental concepts relevant to this work. In Sec. IV, we present a derived generalization bound for a hybrid quantum-classical setup and discuss its implications in Sec. V. Finally, Sec. VI contains a step-by-step derivation of the aforementioned bound.

II. RELATED WORK

On the quantum side, Caro *et al.* [1] proved generalization bounds for QMLMs and derived a bound for (fully) quantum convolutional neural networks (QCNNs) based on covering numbers with T trainable gates to scale at worst as $\tilde{O}(\sqrt{\frac{T \log T}{N}})$. Furthermore, they show that when rerunning the same gates at most M times, the generalization performance scales as $\tilde{O}(\sqrt{\frac{T \log MT}{N}})$, only worsening logarithmically. These findings suggest that these models can generalize well even with a small training dataset. In our work, we extend these results to capture the behavior of a hybrid architecture.

On the hybrid side, while the theoretical foundation is currently lacking (providing motivation for this work) numerous studies have demonstrated promising empirical results. For instance, Matic *et al.* [7] presented empirical evidence supporting the viability of a hybrid quantum-classical convolutional neural network (QCCNN) in a medical use case, demonstrating its potential practical utility in real-world scenarios. A subsequent study [16] studied theoretically derived generalization metrics for QMLMs, such as those from Refs. [2,17], applied to the architecture from Ref. [7] in an empirical context; however, these efforts revealed a significant lack of correlation between theoretical metrics and empirical accuracy. This implies that the general metrics did not capture the behavior of a hybrid model sufficiently well. In contrast, in this work, we take a fully theoretical approach and develop generalization bounds tailored to a hybrid setup.

On the classical side, extensive research has been conducted on generalization bounds for NNs. Lin and Zhang

[18] offer a comprehensive overview of various approaches to these bounds and provide a simplification and generalization of them under specific assumptions. Notably, the bounds $\mathcal{O}(\frac{2^k \eta^k}{\sqrt{N}})$ [19] and $\tilde{O}(\frac{\alpha^{k-1} k^{\frac{3}{2}} \eta \sqrt{d}}{\sqrt{N}})$ [10,20], where $\|F\|_\sigma \leq \eta$, emphasize that generalization depends on how weight magnitudes scale with depth and the activation functions are assumed to be 1-Lipschitz. The bounds illustrate how generalization bounds can be established via norm-based complexity measures, particularly the product of per-layer norms and depth. These bounds are reflected in our analysis of the classical component of hybrid QMLMs. An important point to highlight is that our work prioritizes the investigation into interaction between the classical and quantum components rather than the sharpness of the bounds. Future research can focus on refining these bounds, as our study establishes a robust foundation that facilitates this process.

III. BACKGROUND

A. Generalization

Generalization is the ability of a (trained) model to perform well on unseen data. We measure the generalization error by comparing the model's performance on training data with its expected performance on unseen data. Formally, given a training dataset $S = \{(x_i, y_i)\}, i = 1, \dots, N$, of size N and the model's hypothesis class H , the empirical risk associated with a hypothesis $h \in H$ is given by

$$\hat{R}_S(h) = \frac{1}{N} \sum_{i=1}^N \ell(h(x_i), y_i), \quad (1)$$

where ℓ is a predefined loss function that quantifies the discrepancy between the model prediction and the true label. This empirical risk, also known as the training error, serves as an estimate of the expected risk, which is defined as

$$R(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[\ell(h(x), y)], \quad (2)$$

where the expectation is taken over the unknown data distribution $\mathcal{D}$. The discrepancy between these two quantities, known as the generalization error,

$$\text{gen}(h) = R(h) - \hat{R}_S(h), \quad (3)$$

measures how well the training performance translates to unseen data. Bounding this generalization error is crucial in statistical learning theory and is the focus of the subsequent theoretical results, where maintaining good generalization to new data while fitting the training data well is a trade-off.

B. Complexity measures

Complexity measures quantify how well functions from a hypothesis class are able to fit various training data, as overfitting can be understood as a class having many functions that are distinct from the target function but can still fit the training data well. The Rademacher complexity and covering numbers are complexity measures based on probability and geometric characterizations of hypothesis classes and have well-established connections to translate between them. Our approach builds on existing uniform-bound-based measures

in the quantum setting, where these have proven effective in bounding QMLMs [1].

1. Rademacher complexity

Definition 1 (Empirical Rademacher complexity). Consider arbitrary spaces $\mathcal{X}, \mathcal{Y}$, and define $\mathcal{Z} := \mathcal{X} \times \mathcal{Y}$. Given a training set $S = \{z_1, \dots, z_n\}$ and a real-valued hypothesis class H on $\mathcal{Z}$, the empirical Rademacher complexity is defined as the expected supremum of the correlation between the functions in H and the random Rademacher variables σ_i applied to the dataset:

$$\mathcal{R}_S(H) := \mathbb{E}_\sigma \left[\sup_{h \in H} \frac{1}{n} \sum_{i=1}^n \sigma_i h(z_i) \right], \quad (4)$$

where the Rademacher variables $\sigma_i \sim U_{\{-1,1\}}$ are random variables that take values ± 1 with equal probability.

From the definition, we see that empirical Rademacher complexity varies with the training set, reflecting the actual learning difficulty of each dataset, explaining its name. This data dependence can yield tighter generalization bounds than distribution-independent worst-case bounds, yet it can still be used to derive bounds that hold for all distributions. As such, Rademacher complexity often serves as a key step in proving general results in learning theory. The following theorem is fundamental in this context.

Theorem 1 (Rademacher generalization bounds [22,23,25]). Consider arbitrary spaces $\mathcal{X}, \mathcal{Y}$. We define $\mathcal{Z} := \mathcal{X} \times \mathcal{Y}$ and a real-valued hypothesis class H on $\mathcal{Z}$. For any $\delta > 0$ and any probability measure $\mathbb{P}$ on $\mathcal{Z}$ we have with probability at least $(1 - \delta)$ for the training data $S \in \mathcal{Z}^n$ obtained by n -times repeated sampling with respect to $\mathbb{P}$,

$$\text{gen}(h) \leq 2\mathcal{R}_S(\ell \circ H) + c\sqrt{\frac{\ln \frac{1}{\delta}}{n}}, \quad \forall h \in H, \quad (5)$$

where $c > 0$ is some constant and $\ell \circ H := \{(x, y) \mapsto \ell(h, y, x) \mid h \in H\}$ is the loss-transformed hypothesis class.

This theorem allows us to bound the generalization error using Rademacher complexities.

2. Covering Number

We begin by defining the ε cover and the covering number.

Definition 2 (ε cover and covering number [23]). Let $(V, \|\cdot\|)$ be a normed space, and let $U \subseteq V$. Then U is an ε cover of V if $\forall v \in V$ there exists $u \in U$ such that $\|u - v\| \leq \varepsilon$. The covering number of the normed space $(V, \|\cdot\|)$ with any $\varepsilon > 0$ is the size of the smallest ε cover:

$$\mathcal{N}(V, \|\cdot\|, \varepsilon) := \min\{|U| : U \text{ is an } \varepsilon \text{ cover of } V\}. \quad (6)$$

We call the logarithm-transformed covering number $\ln(\mathcal{N}(V, \|\cdot\|, \varepsilon))$ the metric entropy of V .

The larger the covering number, the more complex the hypothesis class, because it indicates that more functions are needed to approximate the entire class within a small error tolerance. Intuitively, if a hypothesis class H has a large covering number for small ε , it means the class is complex, as there are many distinct functions that cannot be closely approximated by one another. On the other hand, if the covering number is small, the class is simpler, with many hypotheses being

similar to each other within the margin of error ε . There is a common argument chain to arrive at generalization bounds via Theorem 1 and an integral (developed by Dudley [24]) once a covering number has been determined for a hypothesis class.

Theorem 2 (Dudley's entropy integral [25]). Let H be a hypothesis class on $\mathcal{X}$ equipped with a norm $\|\cdot\|$ and let $\gamma_0 := \sup_{h \in H} \|h\|$. We can then bound the (true) Rademacher complexity by Dudley's entropy integral: the usage of

$$\hat{\mathcal{R}}_n(H) \leq \inf_{\alpha \in [0, \frac{\gamma_0}{2})} 4\alpha + \frac{12}{\sqrt{n}} \int_\alpha^{\gamma_0} \sqrt{\ln \mathcal{N}(H, \|\cdot\|, \varepsilon)} d\varepsilon,$$

where $\hat{\mathcal{R}}_n(H)$ is the true Rademacher complexity.

The integral can be understood as summing the contributions of each infinitesimal resolution to the Rademacher complexity of the hypothesis class. At each level of resolution, the function class will become better at fitting the Rademacher random variables.

IV. RESULTS

With this work, we lay the foundation for a more theoretically grounded understanding of the capability of hybrid quantum-classical models. We derive and prove a generalization bound $\tilde{\mathcal{O}}(\frac{\alpha^k}{\sqrt{N}}(k^{\frac{1}{2}}\sqrt{mn} + \sqrt{T \log T}))$ for N training data points, T trainable quantum gates, n -dimensional quantum circuit output, and k bounded fully connected layers $\|F_i\|_F \leq \alpha$ where $i = 1, \dots, k$, $F \in \mathbb{R}^{m \times n}$. The bound can be decomposed into classical and quantum components, which enables investigation of their interaction. Just like the bounds presented by Lin and Zhang [18] on the classical side, we assumed 1-Lipschitz activation functions when presenting our bound; however, the derivation in the Appendix includes the general case for L -Lipschitz activation functions.

Our analysis leverages concepts from statistical learning theory and quantum information, allowing us to theoretically predict the empirical performance of hybrid models. The derived generalization bound indicates how the model's complexity, dictated by the number of quantum gates and the characteristics of the classical layers, influences its ability to generalize from the training data to unseen examples. A key implication of our result is that introducing classical layers into the architecture does not introduce significant generalization disadvantage compared to a fully quantum model with equivalent quantum capacity. In fact, the bound suggests that a well-designed hybrid model with a moderate number of classical layers can match the generalization capabilities of a purely quantum model while offering practical advantages such as reduced quantum circuit depth and better noise tolerance. On the other side, classical NNs typically exhibit generalization bounds that scale unfavorably with the number of layers, often exponentially. In contrast, the hybrid model's quantum component contributes a polylogarithmic factor, potentially offering more favorable scaling under certain conditions. This provides a theoretical justification for hybridization as a strategy for improving scalability without sacrificing learning performance.

V. DISCUSSION

While the theoretical bounds investigated in this work provide valuable insights, they have been shown to struggle in the overparametrized regime. Recent works, such as Ref. [13] and its quantum variant [26], argue that traditional generalization bounds fall short of explaining the empirical success despite high complexities of modern machine learning models. To bridge the gap between theoretical bounds and empirical observations, future research should explore new paradigms of generalization that synergize current understanding of generalization based on model complexities with emerging evidence from edge cases. This approach will foster a more comprehensive understanding of generalization in diverse scenarios. A recent approach [27] has demonstrated strong empirical support for theoretically established bounds in an overparametrized regime. The authors leveraged the relationship between kernel methods and deep learning via Neural Tangent Kernels, validating the idea presented in Ref. [28] that comprehending generalization in deep learning necessitates understanding generalization in kernel methods. This suggests that future research into generalization bounds for NNs can benefit from a synergistic investigation into their more foundational cousin—the kernel methods. Similar connections between models have been identified in the quantum case [29,30], suggesting that this line of research might shed light on generalization bounds of QNNs as well.

The classical part of the bound presented in our work suffers from the same problems mentioned above. While these bounds align with earlier covering-number results, such as those in Ref. [21], they are looser than the more recent bounds found in Refs. [19,20]. However, our primary contribution is the separation of the classical and quantum parts by their covering numbers. The submultiplicative property is applicable to any chosen covering number for each component under appropriate norms. Our approach prioritizes a methodological perspective over the tightness offered by modern results.

Apart from understanding the generalization abilities of quantum models, it is crucial to identify the optimal context (use case) for utilizing these models. Recent research [31] has shown that QCNNs can be efficiently simulated on classical devices through a process known as dequantization. This is particularly efficient on easy datasets that are usually used to demonstrate the utility of proof-of-concept implementations. These results can be extended to a broader QNN architecture. In pursuit of identifying a “killer application” for QMLMs, it is essential to evaluate both the architecture and its generalization capabilities, alongside the complexity of the specific use case at hand. By integrating the tools developed in the present work with insight from Ref. [31], we come one step closer to understanding the generalization benefits of hybrid QMLMs.

Our study has some limitations, which should be addressed in future work. First, we focused on a purely theoretical derivation of the bound, and empirical validation remains to be conducted. However, the authors of Ref. [1], which is foundational for our study, have already performed empirical verification of their bound, suggesting that our results should exhibit similar behavior. Second, we consider a hybrid architecture comprising a quantum layer followed by classical layers. While the derived generalization bound provides




FIG. 1. Visualization of the proof structure: We present metric-entropy and covering-number bounds for both quantum and classical components, demonstrating that the hybrid metric-entropy bound can be expressed through both. We then utilize this framework to derive our generalization bound via Dudley’s entropy integral and Rademacher bounds.

insight into the learning behavior of the hybrid quantum-classical model, it does not directly answer how to determine the optimal quantum-to-classical ratio of layers. The bound primarily consists of the insight provided in Ref. [1] where the generalization bound on the QMLM scales less than exponentially in the training data and the generalization bounds of NNs. However, since both these bounds are loose [13,26], they fail to capture the interplay between large circuits and deep classical layers.

VI. METHODS

In this section, we present the theoretical methods that lead to the generalization bounds for QMLMs by combining bounds from quantum (see Sec. VI B) and classical (see Sec. VI A) parts as illustrated in Fig. 1. For a hybrid QMLM and a sufficiently large training dataset, our bounds guarantee accurate generalization that is close to the performance of a fully QMLM with high probability on unseen data.

A. Classical part: Covering numbers in neural networks

For the classical part, we introduce the metric-entropy bound following the proofs in Ref. [21], which derived generalization bounds for NNs using covering numbers.

Intuitively, a covering number represents the smallest number of covers needed, based on a chosen radius or resolution, to cover a given space. In the context of function spaces, this concept characterizes how uniformly any two functions can vary from each other. When considering bounded linear functions with a bound α , we can visualize the space as being enclosed within a sphere of radius α . Lemma A1 extends this intuition by deriving the covering number, or, more precisely, the metric entropy, associated with this space. Realizing that a classical NN consists of linear layers and activation functions, we can write the hypothesis class of a NN with k layers and extend Lemma 4 to derive a NN’s metric entropy. The idea is to cover the changes in each layer through coverings of each linear layer while dealing with the interspersed activation function. However, separating out the L -Lipschitz activation functions will pick up a factor L for each layer together with the operator bound. The change to the uniform output metric reflects the datum-dependent formulation of Theorem 2, since

we are formulating the generalization bound over possible data input x . The proof can be found in the Appendix.

Lemma 1 (*Metric entropy for classical k -layer neural networks*). Let $\mathcal{X} = \{x \in \mathbb{R}^n : \|x\|_{\ell_2} \leq R\}$ be the input domain. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be an L -Lipschitz activation function applied coordinatewise, and consider the class of k -layer networks $\mathcal{C} = \{f(x) = F_k \sigma(F_{k-1} \sigma(\dots \sigma(F_1 x))), \|F_i\|_F \leq \alpha, F_i \in \mathbb{R}^{m \times n}\}$. Then, for any $\varepsilon > 0$, the covering number of $\mathcal{C}$ under the data-dependent ℓ_2 metric

$$\|f - g\|_{X, \ell_2} := \left(\frac{1}{n} \sum_{j=1}^n \|f(x_j) - g(x_j)\|_2^2 \right)^{1/2}$$

satisfies

$$\ln \mathcal{N}(\mathcal{C}, \|\cdot\|_{X, \ell_2}, \varepsilon) \leq kmn \ln \left(\frac{3kRL^{k-1} \alpha^k}{\varepsilon} \right)$$

for $\varepsilon \leq kRL^{k-1} \alpha^k$, and is zero for larger ε .

With the metric entropy for the k -layer NN, we can derive the generalization bound of the classical part. Inserting Lemma 1 into Theorem 2 we obtain a bound on the empirical Rademacher complexity of the hypothesis class. This induces the generalization bound using Theorem 1. Hence,

Theorem 3 (*Generalization bound for k -layer neural networks*). Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be an L -Lipschitz activation function applied coordinatewise, and consider the class of k -layer networks $\mathcal{C} = \{f(x) = F_k \sigma(F_{k-1} \sigma(\dots \sigma(F_1 x))), \|F_i\|_F \leq \alpha, F_i \in \mathbb{R}^{m \times n}\}$.

Suppose the loss ℓ is coordinatewise L -Lipschitz and bounded by M . Then with probability at least $1 - \delta$ over the choice of independent and identically distributed (i.i.d.) training data S of size N according to $\mathbb{P}$,

$$|R(f) - \hat{R}_S(f)| \in \mathcal{O} \left(L^k \alpha^k k^{\frac{3}{2}} \sqrt{\frac{mn}{N}} + \sqrt{\frac{\log(1/\delta)}{N}} \right). \quad (7)$$

B. Quantum part: Generalization bound in QMLMs

We introduce the generalization bound for a QMLM derived from the covering number of 2-qubit quantum channels and unitary gates, which is formalized in Lemma 2. This lemma provides an upper bound for the covering number of the set of 2-qubit unitaries by using the fact that the space of 2-qubit unitaries is bounded within a unit ball under the operator norm.

Lemma 2 (*Covering-number bounds for 2-qubit unitaries* [1]). Let $\|\cdot\|$ be a unitarily invariant norm on complex 4×4 matrices. The covering number of the set of 2-qubit unitaries $\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ with respect to the norm $\|\cdot\|$ can be bounded as

$$\mathcal{N}(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \varepsilon) \leq \left(\frac{6\|I_{\mathbb{C}^4}\|}{\varepsilon} \right)^{32}, \quad (8)$$

$0 < \varepsilon \leq \|I_{\mathbb{C}^4}\|$.

Remark. The metric entropy of 2-qubit unitaries can then easily be seen as

$$\ln \mathcal{N}(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \varepsilon) \leq 32 \ln \left(\frac{6\|I_{\mathbb{C}^4}\|}{\varepsilon} \right).$$

Theorem 4 (*Generalization bound for QMLM* [1]). For a QMLM $\mathcal{E}_{\theta}^{\text{QMLM}}$ using T parametrized local quantum channels,

we have with high probability over training data of size N that

$$|R(\mathcal{E}_{\theta}^{\text{QMLM}}) - \hat{R}(\mathcal{E}_{\theta}^{\text{QMLM}})| \in \mathcal{O} \left(\sqrt{\frac{T \log(T)}{N}} \right). \quad (9)$$

Remark. This implies that the required size of the data, N , scales as

$$\sqrt{\frac{T \log(T)}{N}} < \epsilon \Rightarrow N > \frac{T \log(T)}{\epsilon^2}.$$

A more refined version of this statement is the following theorem that considers a QMLM where many of the parametrized (local) gates are applied repeatedly. Assume each gate is repeated at most M times.

Theorem 5 (*Generalization bound for repeated local gates* [1]). Let $\mathcal{E}_{\theta}^{\text{QMLM}}$ be a QMLM with an architecture consisting of T independently parametrized 2-qubit Completely Positive and Trace-Preserving (CPTP) maps and at most repeated usage of these channels M times. Then, with probability at least $1 - \delta$ over the choice of i.i.d. training data S of size N according to $\mathbb{P}$,

$$R(\mathcal{E}_{\theta}^{\text{QMLM}}) - \hat{R}_S(\mathcal{E}_{\theta}^{\text{QMLM}}) \in \mathcal{O} \left(\sqrt{\frac{T \log(TM)}{N}} + \sqrt{\frac{\log(1/\delta)}{N}} \right). \quad (10)$$

C. Generalization bound in the hybrid QMLM

Now that we have derived the covering numbers for each part of the hybrid QMLM, we investigate how these layers interact to establish the generalization bound.

A hybrid QMLM integrates both classical and quantum computational parts. Mathematically, this can be viewed as a hypothesis class that consists of functions mapping classical data to quantum states, then back to a classical result via measurement. Formally, we denote a QMLM as

$$\mathcal{E}_{\theta, x}^{\text{QMLM}}(\cdot),$$

where θ are continuous parameters of the quantum gates, i.e., angles in rotation channels, and x is the classical input datum.

We assume in the hybrid setting that the classical datum x is encoded into a quantum state with density matrix $\rho(x)$. As the specific encoding technique has no impact on our proof, we will sometimes omit its dependence on x and just write ρ . The classical NN receives the expected measurement outcome $\text{tr}(M \mathcal{E}_{\theta, x}^{\text{QMLM}}(\rho(x)))$, where M is a measurement operator. We also assume that the classical NN follows empirical risk minimization (ERM) and the optimization over quantum parameters is performed using a classical optimizer, typically gradient-based methods. The model, therefore, consists of a composite hypothesis class

$$H = \{h(x) = F \text{tr}(M \mathcal{E}_{\theta, x}^{\text{QMLM}}) | F \in \mathcal{C}\}, \quad (11)$$

where the quantum component $\mathcal{E}_{\theta, x}^{\text{QMLM}}$ represents a parametrized quantum channel applied to the encoded classical input x , followed by a measurement operator M . The classical component $F \in \mathcal{C}$ is a classical NN with k layers of linear operators with bounded norm and L -Lipschitz activation functions as defined in Lemma 1.

Theorem 6 (Generalization bound in quantum-classical hybrid models). Let h_θ be a hybrid QMLM with an architecture consisting of T independently parametrized 2-qubit CPTP maps followed by bounded measurement operators $\|M\| \leq \beta$ and a k -layer classical NN consisting of activation functions and bounded linear layers $\|F\|_F \leq \alpha$. Then, with probability at least $1 - \delta$ over the choice of i.i.d. training data $X = \{x_i\}_{i=1}^N$ of size N according to $\mathbb{P}$,

$$R(h_\theta) - \hat{R}_X(h_\theta) \in \tilde{O}\left(\frac{\alpha^k}{\sqrt{N}}(k^{\frac{3}{2}}\sqrt{mn} + \sqrt{T \log T})\right), \quad (12)$$

where $\tilde{O}(n)$ denotes Big-O notation with polylogarithmic terms in n hidden.

The proof consists of combining the two results from Lemmas 3 and 2 to obtain a covering number for the hybrid model, after which we establish the generalization bound. Essentially, we demonstrate that the covering numbers exhibit a submultiplicative property across each layer by addressing how the classical and quantum parts interact as separate layers. This then leads to a standard argument of bounding the Rademacher complexity via Dudley's entropy integral which then delivers the generalization bound.

We combine previous work on both classical and quantum machine learning models to claim a similar result for hybrid models with multiple classical layers. For a hybrid QMLM and a sufficiently large training dataset, a slightly adjusted proof gives a bound that guarantees accurate generalization close to the performance of a fully QMLM with high probability on unseen data.

Remark. For a single layer this bound reduces to the quantum part asymptotically,

$$\tilde{O}\left(\frac{\alpha}{\sqrt{N}}\sqrt{mn} + \sqrt{T \log T}\right) = \tilde{O}\left(\sqrt{\frac{T \log(T)}{N}}\right),$$

and implies that the generalization error of a hybrid model with a single classical layer is not much worse than a fully quantum model.

VII. CONCLUSION

This work contributes a theoretical foundation for understanding generalization in hybrid quantum-classical machine learning models. We provide a characterization of learning capacity in hybrid models by decomposing into distinct quantum and classical contributions. Our result shows that introducing bounded classical layers on top of a trainable quantum model does not degrade generalization performance. Instead, hybrid architectures can retain the expressivity and learning guarantees of their fully quantum counterparts while offering practical benefits, such as reduced quantum circuit depth and improved robustness to noise. While our bound advances the theoretical understanding of hybrid models, it leaves open the important question of how to determine the optimal balance between quantum and classical components. Addressing this challenge will require new theoretical tools or empirical studies that go beyond the current statistical learning theory tools. Nonetheless, our work provides a stepping stone for principled design and evaluation of hybrid models and opens

several promising directions for future research in quantum machine learning theory.

ACKNOWLEDGMENTS

This research is supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy, with funds from the Hightech Agenda Bayern.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX

In this section, we present the proof for the claimed generalization bounds for hybrid QMLMs by combining bounds from the quantum (Sec. VIB) and the classical (Sec. VIA) components. The proof is based on showing the submultiplicativity property of the covering numbers when combining both parts.

The first part introduces two lemmata that lead to the classical result from Ref. [21] on generalization bounds for NNs by deriving a metric-entropy bound on the NN. The presented Lemma A1 on bounded linear functionals has been adapted from Ref. [23]. By characterizing the interaction of each layer in the NN, since the layers are bounded linear functionals interspersed with activation functions, we can derive a metric-entropy bound for the entire NN.

Lemma A1 (ℓ_2 metric entropy for bounded linear functionals). Consider the set $\mathcal{V} = \{v \in \mathbb{R}^n \mid \|v\|_{\ell_2} \leq \alpha\}$ of bounded linear functionals (by identifying the vectors v with their duals and interpreting them as functionals on $\mathbb{R}^n$). Then, for any $\varepsilon > 0$, we have

$$\ln \mathcal{N}(\mathcal{V}, \|\cdot\|_{\ell_2}, \varepsilon) \leq n \ln \left(\frac{3\alpha}{\varepsilon}\right). \quad (A1)$$

Proof. Notice that $\mathcal{V} = \{v \in \mathbb{R}^n \mid \|v\|_{\ell_2} \leq \alpha\}$, which is an ℓ_2 ball of radius α . The result follows from plugging into the metric entropy for norm balls. ■

Lemma A2 (Metric entropy for fully connected layer). Let $\mathcal{F} = \{F \in \mathbb{R}^{m \times n} \mid \|F\|_F \leq \alpha\}$, where m is the output dimension and n the input dimension. This set represents bounded matrices presenting a fully connected layer in a neural network. Its metric-entropy bound is given by

$$\ln \mathcal{N}(\mathcal{F}, \|\cdot\|_F, \varepsilon) \leq nm \ln \left(\frac{3\alpha}{\varepsilon}\right). \quad (A2)$$

Proof. Recall that $\mathbb{R}^{m \times n}$ is isomorphic to $\mathbb{R}^{mn}$ and notice that $\|F\|_F = \|\phi(F)\|_{\ell_2}$, where ϕ is the isomorphism between the two spaces. Thus, we apply Lemma A1 and derive the result. ■

Lemma A3 (Metric entropy for classical k -layer neural networks). Let $\mathcal{X} = \{x \in \mathbb{R}^n \mid \|x\|_{\ell_2} \leq R\}$ be the input domain. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be an L -Lipschitz activation function applied coordinatewise, and consider the class of k -layer networks $\mathcal{C} = \{f(x) = F_k \sigma(F_{k-1} \sigma(\dots \sigma(F_1 x))), \|F_i\|_F \leq \alpha, F_i \in \mathbb{R}^{m \times n}\}$. Then, for any $\varepsilon > 0$, the covering number of

$\mathcal{C}$ under the data-dependent ℓ_2 metric

$$\|f - g\|_{X, \ell_2} := \left(\frac{1}{n} \sum_{j=1}^n \|f(x_j) - g(x_j)\|_2^2 \right)^{1/2}$$

satisfies

$$\ln \mathcal{N}(\mathcal{C}, \|\cdot\|_{X, \ell_2}, \varepsilon) \leq kmn \ln \left(\frac{3kRL^{k-1}\alpha^k}{\varepsilon} \right)$$

for $\varepsilon \leq kRL^{k-1}\alpha^k$ and equals zero for larger ε .

Proof. Write the network recursively as

$$\begin{aligned} f_0(x) &:= x, \quad a_i(x) := \sigma(f_i(x)), \\ f_i(x) &:= F_i a_{i-1}(x) \quad (i = 1, \dots, k). \end{aligned} \quad (\text{A3})$$

By induction using $\|F_i\|_F \leq \alpha$ and Lipschitzness of σ ,

$$\begin{aligned} \|a_{i-1}(x)\|_{\ell_2} &\leq (L\alpha)^{i-1} \|x\|_{\ell_2} \leq (L\alpha)^{i-1} R, \\ \forall x \in \mathcal{X}, \quad i &= 1, \dots, k. \end{aligned} \quad (\text{A4})$$

Let $F = f_k(x)$ and $K = h_k(x)$ be two k -layer NNs. We want to bound $\|f_k - h_k\|_{X, \ell_2}$. Replacing layers one at a time to obtain the standard telescoping decomposition and using $\|(F_i - K_i)v\|_{\ell_2} \leq \|F_i - K_i\|_F \|v\|_{\ell_2}$ and Eq. (A4),

$$\begin{aligned} \|f_k(x) - h_k(x)\|_{\ell_2} &\leq \sum_{i=1}^k (L\alpha)^{k-i} \|F_i - K_i\|_F \|a_{i-1}(x)\|_{\ell_2} \\ &\leq \sum_{i=1}^k (L\alpha)^{k-i} \|F_i - K_i\|_F (L\alpha)^{i-1} R \\ &= R(L\alpha)^{k-1} \sum_{i=1}^k \|F_i - K_i\|_F, \quad \forall x \in \mathcal{X}. \end{aligned}$$

Since this bound is uniform in x , it implies the same inequality in the empirical metric:

$$\|f_k(x) - h_k(x)\|_{X, \ell_2} \leq R(L\alpha)^{k-1} \sum_{i=1}^k \|F_i - K_i\|_F. \quad (\text{A5})$$

Choose a per-layer Frobenius radius

$$\delta := \frac{\varepsilon}{kR(L\alpha)^{k-1}}.$$

If each layer F_i is approximated by K_i within δ in Frobenius norm, then Eq. (A5) gives $\|f_k(x) - h_k(x)\|_{X, \ell_2} \leq \varepsilon$. Thus, an ε cover of $\mathcal{C}$ is obtained by taking the product of δ covers of each layer's parameter ball $\{F_i : \|F_i\|_F \leq \alpha\}$.

By Lemma A2, taking products over k layers and logarithms,

$$\begin{aligned} \ln \mathcal{N}(\mathcal{C}, \|\cdot\|_{X, \ell_2}, \varepsilon) &\leq \sum_{i=1}^k mn \ln \left(\frac{3\alpha}{\delta} \right) \\ &= kmn \ln \left(\frac{3\alpha kR(L\alpha)^{k-1}}{\varepsilon} \right) \\ &= kmn \ln \left(\frac{3kRL^{k-1}\alpha^k}{\varepsilon} \right). \end{aligned} \quad (\text{A6})$$

Theorem A1 (Generalization bound in quantum-classical hybrid models). Let H be a hypothesis class for a hybrid machine learning model with T quantum gates and a k -layer NN, and $X = \{x_i\}_{i=1}^N$ be a sample of size N . Assume that the covering number of H with respect to $\|\cdot\|_{X, \ell_2}$ satisfies, for any $\varepsilon > 0$,

$$\begin{aligned} \mathcal{N}(H, \|\cdot\|_{X, \ell_2}, \varepsilon) &\leq \mathcal{N}\left(\mathcal{F}, \|\cdot\|_{X, \ell_2}, \frac{\varepsilon}{2\beta\sqrt{n}}\right) \\ \mathcal{N}\left(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \frac{\varepsilon}{4TL^{k-1}\alpha^k\beta\sqrt{n}}\right)^T, \end{aligned} \quad (\text{A7})$$

where $\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ denotes the set of 2-qubit unitary operators, and $\mathcal{F}$ denotes a k -layer NN with linear layers as defined in Lemma 1. M is a fixed measurement operator with operator norm $\|M\| \leq \beta$, T is the number of parametrized unitaries, and n is the number of output registers of the quantum circuit.

Furthermore, assume that the loss function $\ell : H \times \mathcal{X} \times \mathcal{Y} \rightarrow [0, M]$ is coordinatewise L -Lipschitz continuous in its first argument with respect to the norm $\|\cdot\|$.

Then, with probability at least $1 - \delta$ over the random draw of a sample X with size N , for all hypotheses $h \in H$, the expected loss satisfies

$$\text{gen}(h) \in \tilde{O}\left(\frac{L^k \alpha^k}{\sqrt{N}} (k^{\frac{3}{2}} \sqrt{mn} + \sqrt{T \log T})\right), \quad (\text{A8})$$

where $\tilde{O}(n)$ denotes big-O notation with polylogarithmic terms in n hidden.

Proof. The proof will be split over the next sections. We first derive the covering number of the hypothesis class H of the hybrid QMLM. This is done by realizing that any hypothesis can be covered by coverings of the classical and the quantum component. After deriving the covering number, we will calculate the Rademacher complexity of the hypothesis class via Dudley's entropy integral. Finally, we bound the Rademacher complexity of the loss-transformed hypothesis class by scaling the Rademacher complexity of the hypothesis class by a Lipschitz constant that is averaged over all coordinates L . This delivers the generalization bound. ■

Submultiplicativity of covering numbers for a hybrid quantum-classical model

The hybrid model under consideration comprises two main components: a quantum component and a classical component. The essential step to finding a covering number for the hybrid model is correctly identifying how a covering can be constructed from coverings of the classical and quantum components. Lemmas 2 and A1 each provide covering numbers and we will show that multiplying both provides an upper bound on the hybrid model's covering number. This submultiplicativity property is the following lemma.

Lemma A4. Let $\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ denote the set of 2-qubit unitary operators, let n be the number of output registers in the quantum circuit, and let $\mathcal{F}$ denote a k -layer NN with Frobenius norm $\|F\|_F \leq \alpha$ for each linear layer F_i and L -Lipschitz activation functions. Let M be a fixed measurement operator with operator norm $\|M\| \leq \beta$. Define the hypothesis class H as introduced in Eq. (11) with T the number of parametrized

unitaries. Let $\mathcal{N}(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \varepsilon)$ and $\mathcal{N}(\mathcal{F}, \|\cdot\|_F, \varepsilon)$ be the covering numbers for the 2-qubit unitary operators and the k -layer NN derived previously and let $X = \{x_i\}_{i=1}^N$ be a sample of size N .

Then, for any $\varepsilon > 0$, the covering number of H with respect to $\|\cdot\|_{X, \ell_2}$ satisfies

$$\mathcal{N}(H, \|\cdot\|_{X, \ell_2}, \varepsilon) \leq \mathcal{N}\left(\mathcal{F}, \|\cdot\|_{X, \ell_2}, \frac{\varepsilon}{2\beta\sqrt{n}}\right) \cdot \mathcal{N}\left(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \frac{\varepsilon}{4TL^{k-1}\alpha^k\beta\sqrt{n}}\right)^T.$$

Proof. We will build coverings for (i) each unitary layer in operator norm $\|\cdot\|$, and (ii) the classical NN F in the sample-dependent seminorm $\|\cdot\|_{X, \ell_2}$. The overall cover of H will be the corresponding product covering.

We aim to construct an ε covering for the hypothesis class H by combining ε_F coverings for the classical component $\mathcal{F}$ and ε_U coverings for the quantum component $\mathcal{U}$. The key idea is to ensure that perturbations in F and U individually lead to a controlled perturbation in the composite hypothesis h . Let $\mathcal{C}_{\mathcal{F}}(\varepsilon_F)$ be an ε_F covering of the fully connected layers $\mathcal{F}$ with respect to $\|\cdot\|_F$, and $\mathcal{C}_{\mathcal{U}}(\varepsilon_U)$ be an ε_U covering of the unitaries $\mathcal{U}$ with respect to the spectral norm (induced 2-norm) $\|\cdot\|_{\ell_2}$. We determine ε_F and ε_U accordingly, such that a hypothesis in H is covered by its components with respect to $\|\cdot\|_{\ell_2}$.

We will write a hypothesis h as

$$h(x) = F \cdot z(x), \quad z(x) = \text{tr} [M \mathcal{E}_{\theta, x}^{\text{QMLM}}(\rho(x))]_{i=1}^n \in \mathbb{R}^n,$$

where $\rho(x)$ is the density operator of the encoded data x and M is a measurement operator.² This separates the classical part $F \in \mathcal{C}$ and the quantum part $z(x)$. As a reminder,

$$\mathcal{E}_{\theta}^{\text{QMLM}} = U_T \cdots U(\cdot) U^\dagger \cdots U_T^\dagger,$$

where the $U_i \in \mathcal{U}$ are 2-qubit unitaries depending on θ_i .

Consider two hypotheses $h, h' \in H$. As introduced above, we can write them as

$$\begin{aligned} h(x) &= Fz(x), \\ h'(x) &= Kz'(x), \end{aligned} \quad (\text{A9})$$

where $F, K \in \mathcal{F}$ are the classical parts and z, z' are the quantum parts.

To construct a covering for the hypothesis class H , we want to bound

$$\|h - h'\|_{X, \ell_2}. \quad (\text{A10})$$

We will bound this by separating the classical layers and the quantum layer and then bound them individually. The

separation uses the common insertion trick in triangle inequalities:

$$h(x) - h'(x) = (Fz(x) - Fz'(x)) + (F - K)z'(x). \quad (\text{A11})$$

Inserting this into Eq. (A10) gives us

$$\begin{aligned} \|h - h'\|_{X, \ell_2} &= \|(Fz - Fz') + (F - K)z'\|_{X, \ell_2} \\ &\leq \|(Fz - Fz')\|_{X, \ell_2} + \|(F - K)z'\|_{X, \ell_2}. \end{aligned} \quad (\text{A12})$$

Equation (A12) shows the varied parts that need to be covered. We (i) first bound the term $\|(Fz - Fz')\|_{X, \ell_2}$ corresponding to the varied quantum part and (ii) then bound the second term $\|(F - K)z'\|_{X, \ell_2}$ corresponding to the varied classical NN.

Starting with the quantum part, notice that, for any $z, z' \in \mathbb{R}^n$,

$$\begin{aligned} \|F(z) - F(z')\|_{\ell_2} &= \|F_k[\sigma(\cdots) - \sigma(\cdots)]\|_{\ell_2} \\ &\leq \|F_k\| \|\sigma(\cdots) - \sigma(\cdots)\|_{\ell_2} \\ &\leq \|F_k\|_F \|\sigma(\cdots) - \sigma(\cdots)\|_{\ell_2} \\ &\leq \alpha \|\sigma(\cdots) - \sigma(\cdots)\|_{\ell_2} \\ &\leq \alpha L \|F_{k-1}\sigma(\cdots) - F_{k-1}\sigma(\cdots)\|_{\ell_2} \\ &\vdots \\ &\leq L^{k-1} \alpha^k \|z - z'\|_{\ell_2}. \end{aligned}$$

Thus,

$$\|F(z) - F(z')\|_{X, \ell_2} \leq L^{k-1} \alpha^k \|z - z'\|_{X, \ell_2}.$$

We now examine the varied quantum parts $\|z - z'\|_{\ell_2}$. Notice that for a measurement operator M with $\|M\| \leq \beta$, and Φ, Ψ quantum channels, and a state ρ with $\|\rho\|_1 = 1$, we have

$$\begin{aligned} |z_j(x) - z'_j(x)| &= |\text{tr}(M(\Phi - \Psi)(\rho(x)))| \\ &\leq \|M\| \|(\Phi - \Psi)(\rho(x))\|_1 \quad (\text{H\"older's inequality (Schatten)}) \\ &\leq \|M\| \|\Phi - \Psi\|_{\diamond} \quad (\text{Def. Diamond norm}) \\ &\leq \beta \|\Phi - \Psi\|_{\diamond}. \quad (\|M\| \leq \beta). \end{aligned} \quad (\text{A13})$$

Consequently, if we define the n -dimensional measurement vector

$$z(x) := [\text{tr}(M \mathcal{E}_{\theta, x}^{\text{QMLM}}(\rho(x)))]_{j=1, \dots, n} \in \mathbb{R}^n$$

, then for any two channels $\mathcal{E}, \mathcal{E}'$,

$$\|z(x) - z'(x)\|_{\ell_2} \leq \beta \sqrt{n} \|\mathcal{E}_{\theta, x} - \mathcal{E}'_{\theta, x}\|_{\diamond}. \quad (\text{A14})$$

Using the same approach as in the proof of Lemma C.1 from Ref. [1], let $\mathcal{U} := \mathcal{E}_{\theta, x}$ and $\mathcal{V} := \mathcal{E}'_{\theta, x}$ be the unitary channels of the quantum parts z and z' . A telescoping expansion gives

$$\|\mathcal{U} - \mathcal{V}\|_{\diamond} \leq 2 \sum_{k=1}^T \|U_k - V_k\|. \quad (\text{A15})$$

²For complete mathematical accuracy, we should specify that the general form of the local 2-qubit unitary operators consists of tensor products $I \otimes \cdots \otimes U_k \otimes \cdots \otimes I$.

Combining Eq. (A14) and Eq. (A15) gives

$$\|z - z'\|_{X, \ell_2} \leq 2\beta\sqrt{n} \sum_{k=1}^T \|U_k - V_k\|. \quad (\text{A16})$$

Finally, for the quantum part we get

$$\|Fz - Fz'\|_{X, \ell_2} \leq 2L^{k-1}\alpha^k\beta\sqrt{n} \sum_{t=1}^T \|U_t - V_t\|. \quad (\text{i})$$

Moving on to the classical part, notice that

$$\|(F - K)z'\|_{X, \ell_2}$$

fulfills the conditions for Lemma 1 with $R = \beta\sqrt{n}$, since $\|z(x)\|_{\ell_2} \leq \beta\sqrt{n}$. This lets us know the covering number for any covering of radius ϵ for the k -layer NN F . Together, we see that

$$\begin{aligned} \|h - h'\|_{X, \ell_2} &\leq 2L^{k-1}\alpha^k\beta\sqrt{n} \sum_{k=1}^T \|U_k - V_k\| \\ &\quad + \|(F - K)z'\|_{X, \ell_2}. \end{aligned} \quad (\text{A17})$$

Set

$$\begin{aligned} \|U_k - V_k\| &\leq \frac{\epsilon}{2 \times 2TL^{k-1}\alpha^k\beta\sqrt{n}} = \epsilon_U, \\ \|(F - K)z'\|_{X, \ell_2} &\leq \frac{\epsilon}{2\beta\sqrt{n}} = \epsilon_F. \end{aligned} \quad (\text{A18})$$

Then, for any $h \in H$, there exists a covering $\mathcal{C}_H$ with $h' \in \mathcal{C}_H$ since

$$\|h - h'\|_{X, \ell_2} \leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \quad (\text{A19})$$

Hence, the constructed cover $\mathcal{C}_H$ satisfies

$$\sup_{h \in H} \inf_{h' \in \mathcal{C}_H} \|h - h'\|_{X, \ell_2} \leq \epsilon,$$

and since each hypothesis in H can be approximated within ϵ by combining an ϵ_F cover for $\mathcal{F}$ and $T \epsilon_U$ covers for $\mathcal{U}$, the total number of elements in $\mathcal{C}_H(\epsilon)$ is at most the product of the covering numbers of $\mathcal{F}$ and $\mathcal{U}$ T times, and we prove that

$$\begin{aligned} \mathcal{N}(H, \|\cdot\|_{X, \ell_2}, \epsilon) &\leq \mathcal{N}\left(\mathcal{F}, \|\cdot\|_{X, \ell_2}, \frac{\epsilon}{2\beta\sqrt{n}}\right) \\ &\quad \times \mathcal{N}\left(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \frac{\epsilon}{4TL^{k-1}\alpha^k\beta\sqrt{n}}\right)^T. \end{aligned} \quad (\text{A20})$$

Applying the logarithm and using the product rule gives the metric entropy H by

$$\begin{aligned} \ln \mathcal{N}(H, \|\cdot\|_{\ell_2}, \epsilon) &\leq \ln \mathcal{N}\left(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \frac{\epsilon}{4T\alpha\beta\sqrt{n}}\right)^T \\ &\quad + \ln \mathcal{N}\left(\mathcal{F}, \|\cdot\|_F, \frac{\epsilon}{2\beta\sqrt{n}}\right). \end{aligned} \quad (\text{A21})$$

a. Rademacher complexity

Lemma A5 (Radius of the hybrid hypothesis class). Let H be the hybrid hypothesis class from Lemma A4 producing n -dimensional outputs via a measurement operator M with $\|M\| \leq \beta$. Define $\gamma_0 := \sup_{h \in H} \|h\|_{X, \ell_2}$, where $\|h\|_{X, \ell_2}^2 := \frac{1}{N} \sum_{i=1}^N \|h(x_i)\|_{\ell_2}^2$. Then

$$\gamma_0 \leq \beta\sqrt{n}.$$

Proof. Each coordinate of $h(x)$ is a measurement expectation of the form $\text{tr}(M\rho)$ with $|\text{tr}(M\rho)| \leq \|M\| \|\rho\|_1 \leq \beta$, independently of the classical layers. Hence, $\|h(x)\|_{\ell_2} \leq \beta\sqrt{n}$ for every input x , and therefore $\|h\|_{X, \ell_2}^2 = \frac{1}{N} \sum_{i=1}^N \|h(x_i)\|_{\ell_2}^2 \leq \beta^2 n$. Taking the supremum over $h \in H$ gives $\gamma_0 \leq \beta\sqrt{n}$. ■

Theorem A2 (Rademacher complexity bound for the hybrid class). Under the setting of Lemma A4, with the neural-network metric-entropy lemma and the diamond-norm covering for the quantum part, we have, for all $\epsilon > 0$,

$$\begin{aligned} \ln \mathcal{N}(H, \|\cdot\|_{X, \ell_2}, \epsilon) &\leq \ln \mathcal{N}(\mathcal{F}, \|\cdot\|_{X, \ell_2}, \frac{\epsilon}{2\beta\sqrt{n}}) \\ &\quad + T \ln \mathcal{N}\left(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \frac{\epsilon}{4TL^{k-1}\alpha^k\beta\sqrt{n}}\right). \end{aligned} \quad (\text{A22})$$

Consequently, Dudley's entropy integral (with cutoff) yields

$$\hat{\mathcal{R}}_N(H) \in \tilde{\mathcal{O}}\left(\frac{L^{k-1}\alpha^k\beta\sqrt{n}}{\sqrt{N}}(kR\sqrt{kmn} + \sqrt{T \log T})\right).$$

Proof. We start by inserting the metric entropy derived in Lemma A4 into Dudley's Theorem (Theorem 2):

$$\hat{\mathcal{R}}_N(H) \leq \frac{12}{\sqrt{N}} \int_0^{\gamma_0} \sqrt{\ln \mathcal{N}(\mathcal{F}) + T \ln \mathcal{N}(\mathcal{U})} d\epsilon. \quad (\text{A23})$$

We omitted the metric and resolution size for readability. By Lemma A5, γ_0 represents the upper bound of the integration bounds, since it is the coarsest resolution for the hypothesis class H . Then, by applying the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$,

$$\hat{\mathcal{R}}_N(H) \leq \frac{12}{\sqrt{N}} \left(\int_0^{\gamma_0} \sqrt{\ln \mathcal{N}\left(\mathcal{F}, \|\cdot\|_{X, \ell_2}, \frac{\epsilon}{2\beta\sqrt{n}}\right)} d\epsilon \right. \quad (\text{A24})$$

$$\left. + \int_0^{\gamma_0} \sqrt{T \ln \mathcal{N}\left(\mathcal{U}\left(\mathbb{C}^2 \otimes \mathbb{C}^2, \|\cdot\|, \frac{\epsilon}{4TL^{k-1}\alpha^k\beta\sqrt{n}}\right)\right)} d\epsilon \right). \quad (\text{A25})$$

We analyze each summand separately: Starting from the classical contribution, insert Lemma A3. The integral becomes directly evaluable,

$$\frac{12}{\sqrt{N}} \int_0^{\gamma_0} \sqrt{kmn \ln \left(\frac{6kRL^{k-1}\alpha^k \beta \sqrt{n}}{\varepsilon} \right)} d\varepsilon = \frac{12\sqrt{kmn}}{\sqrt{N}} J, \quad (\text{A26})$$

where

$$J := \int_0^{\gamma_0} \sqrt{\ln \left(\frac{C}{\varepsilon} \right)} d\varepsilon, \quad C := 6kRL^{k-1}\alpha^k \beta \sqrt{n}.$$

Evaluate J by the change of variables $u = \ln(C/\varepsilon)$, i.e., $\varepsilon = Ce^{-u}$ and $d\varepsilon = -Ce^{-u} du$. The limits become $\varepsilon \downarrow 0 \Rightarrow u \rightarrow \infty$ and $\varepsilon = \gamma_0 \Rightarrow u = \ln(C/\gamma_0)$; hence,

$$J = C \int_{\ln(C/\gamma_0)}^{\infty} u^{1/2} e^{-u} du = C \Gamma\left(\frac{3}{2}, \log \frac{C}{\gamma_0}\right),$$

where $\Gamma(s, x)$ is the gamma function. Using the identity $\Gamma(\frac{3}{2}, x) = \frac{\sqrt{\pi}}{2} \operatorname{erf}(\sqrt{x}) + \sqrt{x} e^{-x}$ and $e^{-\ln(C/\gamma_0)} = \gamma_0/C$, we obtain

$$J = \frac{\sqrt{\pi}}{2} C \operatorname{erf}\left(\sqrt{\ln \frac{C}{\gamma_0}}\right) + \gamma_0 \sqrt{\ln \frac{C}{\gamma_0}},$$

where $\operatorname{erf}(x) := \frac{2}{\sqrt{\pi}} \int_0^x \exp(-t^2) dt$. Multiplying J by $\frac{12\sqrt{kmn}}{\sqrt{N}}$ gives us

$$\frac{12\sqrt{kmn}}{\sqrt{N}} \left(\frac{\sqrt{\pi}}{2} C \operatorname{erf}\left(\sqrt{\ln \frac{C}{\gamma_0}}\right) + \gamma_0 \sqrt{\ln \frac{C}{\gamma_0}} \right),$$

which is the classical contribution (A26) we started with.

Using the fact that $\operatorname{erf}(\cdot) \leq 1$ and hiding constants and polylogarithmic terms using the $\tilde{O}$ notation, we see that the Rademacher complexity of the classical part is

$$\tilde{O}\left(\sqrt{kmn} \frac{kRL^{k-1}\alpha^k \beta \sqrt{n}}{\sqrt{N}}\right).$$

Moving on to the quantum part, we see

$$\frac{12}{\sqrt{N}} \int_0^{\gamma_0} \sqrt{T \ln \mathcal{N}\left(\mathcal{U}\left(\mathbb{C}^2 \otimes \mathbb{C}^2, \|\cdot\|, \frac{\varepsilon}{4TL^{k-1}\alpha^k \beta \sqrt{n}}\right)\right)} d\varepsilon. \quad (\text{A27})$$

With the change of variables

$$\delta = \frac{\varepsilon}{4TL^{k-1}\alpha^k \beta \sqrt{n}}, \quad b = \frac{\gamma_0}{4TL^{k-1}\alpha^k \beta \sqrt{n}},$$

this becomes

$$\frac{48TL^{k-1}\alpha^k \beta \sqrt{n}}{\sqrt{N}} \int_0^b \sqrt{T \ln \mathcal{N}(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \delta)} d\delta. \quad (\text{A28})$$

By Lemma 2, for any unitarily invariant norm $\|\cdot\|$ and $0 < \delta \leq \|I_{\mathbb{C}^4}\|$,

$$\ln \mathcal{N}(\mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2), \|\cdot\|, \delta) \leq 32 \ln \left(\frac{6\|I_{\mathbb{C}^4}\|}{\delta} \right).$$

Plugging this into Eq. (A28) and using that for the spectral norm $\|I_{\mathbb{C}^4}\| = 1$ yields

$$\frac{48TL^{k-1}\alpha^k \beta \sqrt{n}}{\sqrt{N}} \sqrt{32T} \int_0^b \sqrt{\ln \left(\frac{6}{\delta} \right)} d\delta. \quad (\text{A29})$$

Evaluating the integral exactly gives

$$\int_0^b \sqrt{\ln \left(\frac{6}{\delta} \right)} d\delta = \frac{\sqrt{\pi}}{2} 6 \operatorname{erf}\left(\sqrt{\ln \frac{6}{b}}\right) + b \sqrt{\ln \frac{6}{b}}.$$

Hence, when inserting back $b = \frac{\gamma_0}{4TL^{k-1}\alpha^k \beta \sqrt{n}}$, the quantum Dudley term becomes

$$\frac{3\sqrt{512}}{\sqrt{N}} \gamma_0 \sqrt{T} \sqrt{\ln \left(\frac{24TL^{k-1}\alpha^k \beta \sqrt{n}}{\gamma_0} \right)}. \quad (\text{A30})$$

Using $\gamma_0 \leq L^{k-1}\alpha^k \beta \sqrt{n}$ we see that the quantum part scales as $\frac{L^{k-1}\alpha^k \beta \sqrt{n}}{\sqrt{N}} \sqrt{T \log T}$ up to explicit constants.

Combining both expressions, we arrive at the claimed

$$\tilde{O}\left(\frac{L^{k-1}\alpha^k \beta \sqrt{n}}{\sqrt{N}} (kR\sqrt{kmn} + \sqrt{T \log T})\right). \quad \blacksquare$$

b. Generalization bound

Having obtained an upper bound on the empirical Rademacher complexity $\hat{\mathcal{R}}_N(H)$ of the hypothesis class H , we aim to apply Theorem 1 to establish a corresponding generalization bound. However, it is important to observe that Theorem 1 requires the empirical Rademacher complexity of the loss-transformed hypothesis class $\mathcal{R}_S(\ell \circ H)$, rather than $\hat{\mathcal{R}}_S(H)$. To overcome this gap, we will introduce the following lemma which provides an upper bound on $\mathcal{R}(\ell \circ H)$ in terms of $\mathcal{R}(H)$ for loss functions that are Lipschitz continuous, such as the multiclass hinge loss. We will state it without proof.

Lemma A6 (Talagrand-ledoux's contraction lemma [32]). Let H be a hypothesis class and let $\ell : H \times \mathcal{Y} \rightarrow \mathbb{R}$ be an L -Lipschitz loss function with respect to its first argument. For any fixed sample $S = \{x_1, x_2, \dots, x_n\}$, the L -Lipschitz loss will scale as $\frac{L}{\sqrt{n}}$ when averaged over the sample. Thus, the empirical Rademacher complexity of the loss-transformed hypothesis class $\ell \circ H$ satisfies

$$\hat{\mathcal{R}}_S(\ell \circ H) \leq \frac{L}{\sqrt{n}} \hat{\mathcal{R}}_S(H). \quad (\text{A31})$$

Proof. By applying this lemma, we can estimate the Rademacher complexity of the loss-transformed hypothesis class by $L\hat{\mathcal{R}}_S(H)$.

Let ℓ be a mean-reduced, coordinatewise L -Lipschitz loss (so that the Talagrand-Ledoux contraction yields the factor $L/\sqrt{n}$). Combining Lemma A4, Theorem A2, and the Talagrand-Ledoux contraction, we obtain, for our dataset X ,

$$\begin{aligned} \hat{\mathcal{R}}_X(\ell \circ H) &\leq \frac{L}{\sqrt{n}} \hat{\mathcal{R}}_X(H) \in \tilde{O}\left(\frac{L^k \alpha^k \beta}{\sqrt{N}} (kR\sqrt{kmn} + \sqrt{T \log T})\right). \end{aligned}$$

To conclude the proof, we use Theorem 1: $\text{gen}(h) \leq 2\widehat{\mathcal{R}}_X(\ell \circ H) + c\sqrt{\ln(1/\delta)/N}$. The Talagrand-Ledoux contraction gives $\widehat{\mathcal{R}}_X(\ell \circ H) \leq \frac{L}{\sqrt{n}}\widehat{\mathcal{R}}_X(H)$. Substitute the bound from

Theorem A2: $\widehat{\mathcal{R}}_X(H) \in \tilde{O}(\frac{L^{k-1}\alpha^k\beta\sqrt{n}}{\sqrt{N}}(kR\sqrt{kmn} + \sqrt{T \log T}))$, which yields the stated result after canceling $\sqrt{n}$. ■

- [1] M. C. Caro, H.-Y. Huang, M. Cerezo, K. Sharma, A. Sornborger, L. Cincio, and P. J. Coles, Generalization in quantum machine learning from few training data, *Nat. Commun.* **13**, 4919 (2022).
- [2] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
- [3] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, Power of data in quantum machine learning, *Nat. Commun.* **12**, 2631 (2021).
- [4] C. Zoufal, A. Lucchi, and S. Woerner, Quantum generative adversarial networks for learning and loading random distributions, *npj Quantum Inf.* **5**, 103 (2019).
- [5] M. Kālis, A. Locāns, R. Šikovs, H. Naseri, and A. Ambainis, A hybrid quantum-classical approach for inference on restricted boltzmann machines, *Quantum Mach. Intell.* **5**, 44 (2023).
- [6] A. Sakhnenko, C. O'Meara, K. J. B. Ghosh, C. B. Mendl, G. Cortiana, and J. Bernabé-Moreno, Hybrid classical-quantum autoencoder for anomaly detection, *Quantum Mach. Intell.* **4**, 27 (2022).
- [7] A. Matic, M. Monnet, J. M. Lorenz, B. Schachtner, and T. Messerer, Quantum-classical convolutional neural networks in radiological image classification, in *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Broomfield, CO, USA, 2022), pp. 56–66.
- [8] K. Gili, M. Hibat-Allah, M. Mauri, C. Ballance, and A. Perdomo-Ortiz, Do quantum circuit born machines generalize? *Quantum Sci. Technol.* **8**, 035021 (2023).
- [9] R. Sweke, J.-P. Seifert, D. Hangleiter, and J. Eisert, On the quantum versus classical learnability of discrete distributions, *Quantum* **5**, 417 (2021).
- [10] P. L. Bartlett, D. J. Foster, and M. Telgarsky, Spectrally-normalized margin bounds for neural networks, in *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)* (Curran Associates, Inc., Long Beach, CA, USA, 2017), pp. 6241–6250.
- [11] F. He and D. Tao, Recent advances in deep learning theory, [arXiv:2012.10931](https://arxiv.org/abs/2012.10931).
- [12] A. Abbas, D. Sutter, A. Figalli, and S. Woerner, Effective dimension of machine learning models, [arXiv:2112.04807](https://arxiv.org/abs/2112.04807).
- [13] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, *Commun. ACM* **64**, 107 (2021).
- [14] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [15] M. Schuld, I. Sinayskiy, and F. Petruccione, An introduction to quantum machine learning, *Contemp. Phys.* **56**, 172 (2015).
- [16] M. Monnet, N. Chaabani, T.-A. Drăgan, B. Schachtner, and J. M. Lorenz, Understanding the effects of data encoding on quantum-classical convolutional neural networks, in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, Montral, QC, Canada, 2024), pp. 1436–1446.
- [17] M. Schuld, R. Sweke, and J. J. Meyer, Effect of data encoding on the expressive power of variational quantum-machine-learning models, *Phys. Rev. A* **103**, 032430 (2021).
- [18] S. Lin and J. Zhang, Generalization bounds for convolutional neural networks, [arXiv:1910.01487](https://arxiv.org/abs/1910.01487).
- [19] B. Neyshabur, R. Tomioka, and N. Srebro, Norm-based capacity control in neural networks, in *Proceedings of The 28th Conference on Learning Theory*, Proceedings of Machine Learning Research Vol. 40, edited by P. Grünwald, E. Hazan, and S. Kale (PMLR, Paris, 2015), pp. 1376–1401.
- [20] B. Neyshabur, S. Bhojanapalli, and N. Srebro, A PAC-bayesian approach to spectrally-normalized margin bounds for neural networks, in *International Conference on Learning Representations* (2018).
- [21] M. Anthony and P. L. Bartlett, *Neural Network Learning: Theoretical Foundations* (Cambridge University Press, Cambridge, U.K., 1999).
- [22] P. L. Bartlett and S. Mendelson, Rademacher and gaussian complexities: risk bounds and structural results, *J. Mach. Learn. Res.* **3**, 463 (2003).
- [23] M. M. Wolf, Mathematical foundations of supervised learning (2023), <https://mediatum.ub.tum.de/node?id=1723378>.
- [24] R. M. Dudley, *Uniform Central Limit Theorems* (Cambridge University Press, Cambridge, UK, 1999).
- [25] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, UK, 2014).
- [26] E. Gil-Fuster, J. Eisert, and C. Bravo-Prieto, Understanding quantum machine learning also requires rethinking generalization, *Nat. Commun.* **15**, 2277 (2024).
- [27] A. Canatar, B. Bordelon, and C. Pehlevan, Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks, *Nat. Commun.* **12**, 2914 (2021).
- [28] M. Belkin, S. Ma, and S. Mandal, To understand deep learning we need to understand kernel learning, in *International Conference on Machine Learning* (PMLR, 2018), pp. 541–549.
- [29] J. Liu, F. Tacchino, J. R. Glick, L. Jiang, and A. Mezzacapo, Representation learning via quantum neural tangent kernels, *PRX Quantum* **3**, 030323 (2022).
- [30] M. Incudini, M. Grossi, A. Mandarino, S. Vallecorsa, A. D. Pierro, and D. Windridge, The quantum path kernel: A generalized neural tangent kernel for deep quantum machine learning, *IEEE Trans. Quantum Eng.* **4**, 1 (2023).
- [31] P. Bermejo, P. Braccia, M. S. Rudolph, Z. Holmes, L. Cincio, and M. Cerezo, Quantum convolutional neural networks are (effectively) classically simulable, [arXiv:2408.12739](https://arxiv.org/abs/2408.12739).
- [32] M. Ledoux and M. Talagrand, *Probability in Banach Spaces: Isoperimetry and Processes*, Classics in Mathematics (Springer-Verlag, Berlin, 1991).