REVIEW ARTICLE



Generalization error bound for quantum machine learning in NISQ era—a survey

Bikram Khanal¹ · Pablo Rivas¹ · Arun Sanjel¹ · Korn Sooksatra¹ · Ernesto Quevedo¹ · Alejandro Rodriguez¹

Received: 25 November 2023 / Accepted: 16 October 2024 / Published online: 4 December 2024 © The Author(s), under exclusive licence to Springer Nature Switzerland AG 2024

Abstract

Despite the mounting anticipation for the quantum revolution, the success of quantum machine learning (QML) in the noisy intermediate-scale quantum (NISQ) era hinges on a largely unexplored factor: the generalization error bound, a cornerstone of robust and reliable machine learning models. Current QML research, while exploring novel algorithms and applications extensively, is predominantly situated in the context of noise-free, ideal quantum computers. However, quantum circuit (QC) operations in NISQ-era devices are susceptible to various noise sources and errors. In this article, we conducted a systematic mapping study (SMS) to explore the state-of-the-art generalization error bound for QML in NISQ-era devices and analyze the latest practices in the field. Our study systematically summarizes the existing computational platforms with quantum hardware, datasets, optimization techniques, and the proposed error bounds detailed in the literature. It also highlights the limitations and challenges in QML in the NISQ era and discusses future research directions to advance the field. Using a detailed Boolean operators query in five reliable indexers, we collected 544 papers and filtered them to a small set of 37 relevant articles. This filtration was done following the best practice of SMS with well-defined research questions and inclusion and exclusion criteria.

Keywords Quantum machine learning · Generalization error bound · NISQ devices · Quantum circuits

1 Introduction

In the field of machine learning, the fundamental theory of learning from data (Abu-Mostafa et al. 2012), it is defined that any model aims to learn an unknown target function

Arun Sanjel and Korn Sooksatra contributed equally to this work

 Pablo Rivas pablo_rivas@baylor.com
 Bikram Khanal bikram_khanal1@Baylor.edu
 Arun Sanjel arun_sanjel1@baylor.edu
 Korn Sooksatra korn_sooksatra1@baylor.edu
 Ernesto Quevedo ernesto_quevedo1@baylor.edu
 Alejandro Rodriguez alejandro_rodriguez4@baylor.edu
 School of Engineering & Computer Science, Baylor

University, One Bear Place, Waco 76798, TX, USA

using a learning algorithm \mathcal{A} . The algorithm produces a hypothesis g from a hypothesis set \mathcal{H} that approximates f. In classical machine learning, the generalization error or generalization gap is the difference between a model's performance on the training data and unseen data (Emami et al. 2020; Jakubovitz et al. 2019; Nadeau and Bengio 1999). The true error $\mathbb{E}(g) = \mathbb{P}[g(x_i) \neq f(x_i)]$ is the expected difference between the hypothesis output $g(x_i)$ and the actual output $f(x_i), i = 1, ..., N$, and the empirical error $\hat{\mathbb{E}}(g) = \frac{1}{N} \sum_{i=1}^{N} [g(x_i) \neq f(x_i)]$ is the average difference between true label y_i and predicted label \hat{y}_i over the dataset. In this context, g(x) is the model's current best approximation of f and $\mathbb{P}[\cdot]$ is the probability of an event. A bound on the generalization error provides an upper limit to the model's error rate on unseen data. It is typically derived using statistical learning theory and depends on factors such as the complexity of the model (hypothesis class), the number of examples in the training data, and the randomness in the data generation process. Hoeffding's inequality (Hoeffding 1994) gives a common form of the generalization bound in

 $f: \mathcal{X} \to \mathcal{Y}$ from a dataset $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_N)\}$

classical machine learning. Hoeffding's inequality provides an error bound for f based on \mathcal{D} (Duchi n.d.) by giving the deviation probability of $\hat{\mathbb{E}}(g)$ from $\mathbb{E}(g)$ as function of a positive tolerance ϵ and total number of training samples N and is given by

$$\mathbb{P}[|\mathbb{E}(g) - \hat{\mathbb{E}}(g)| \ge \epsilon] \le 2e^{-2N\epsilon^2} \tag{1}$$

This equation states that the probability that the absolute difference between the in-sample and out-of-sample errors being greater than ϵ is less than or equal to $2e^{-2N\epsilon^2}$. However, the effectiveness of this bound may be reduced or invalidated when the underlying random variable is affected by the noise, as is the case with NISQ devices (De Palma et al. 2023; Du et al. 2021; Hakkaku et al. 2022). While some noise is beneficial for classical machine learning (Nee-lakantan et al. 2015), excessive noise in quantum systems can limit the performance of QML algorithms (Bharti et al. 2021; Wang et al. 2021). To overcome this challenge, NISQ-era hardware-related noise and QC operations-induced noise must be considered when developing reliable QML models.

QML is an emerging field with great promise for revolutionizing learning from data (Arunachalam and de Wolf 2017; Biamonte et al. 2017; Carleo et al. 2019; Wittek 2014). It focuses on improving machine learning using quantum systems through a mathematical framework (Arunachalam and de Wolf 2017). Several learning models (Arunachalam and de Wolf 2017; Martín-Guerrero and Lamata 2022) have been proposed in QML, including probably approximately correct (PAC) that explores how quantum resources, such as superposition and entanglement, can improve the sample complexity or computational efficiency of learning classical concepts (Rocchetto et al. 2019). However, our ability to harness QML capabilities is particularly impacted by the practical limitations of NISQ devices, which are presently the most advanced quantum computers available (Bharti et al. 2021; Preskill 2018). Variational quantum computing (Cerezo et al. 2021) uses imperfect NISQ-era devices for computation. The variational quantum circuit (VQC) model is a QML model that can be described as a quantum circuit model (Schuld and Petruccione 2021; Wittek 2014) and is defined as

$$f_{\theta}(x) = \langle \psi(x,\theta) | \mathcal{M} | \psi(x,\theta) \rangle.$$
(2)

with $|\psi(x,\theta)\rangle$ being a quantum state prepared by a parameterized quantum circuit (PQC) $U(x,\theta)$ with trainable parameters θ (Benedetti et al. 2019) and \mathcal{M} a measurement operator. On a system with *n* number of qubits, the circuit $U(x,\theta)$ initiates with a sequence of quantum gates from an initial state predominantly $|0\rangle^{\otimes n}$ (Biamonte et al. 2017). The relationship between the PQC and the unitary is crucial and can be defined as follows:

$$|\psi(x,\theta)\rangle = U(x,\theta)|0\rangle^{\otimes n} \tag{3}$$

Though theoretically promising, studies have shown that deep quantum circuits are especially vulnerable to noise, accumulating gate errors and experiencing significant decoherence (Alam and Ghosh 2022; Wang et al. 2022). In addition, the benefits of quantum kernels (Heyraud et al. 2022; Kübler et al. 2021; Schuld 2021; Thanasilp et al. 2024; Wang et al. 2021) are reduced in the presence of large system noise and a higher number of measurements (Huang et al. 2021; Preskill 2018; Thanasilp et al. 2024; Wang et al. 2021). Furthermore, the limited data available for training quantum models becomes even more challenging to work with due to the noise inherent in NISQ devices, leading to potential misdirection in learning and increasing generalization error (Schuld and Killoran 2019).

The implication is clear: algorithms implemented on NISQ devices are susceptible to considerable noise and may not work as expected (Bharti et al. 2021; Wang et al. 2021), causing a divergence between theoretical predictions and empirical results. To fully exploit the potential of QML models in the NISQ era, it is crucial to develop a deep understanding of these constraints and establish a robust error bound to account for the impact of the existing system noise and hardware limitations. Consequently, a dedicated study to define and understand the generalization error bound (GEB) for QML in the NISQ era is not merely a theoretical interest but a critical necessity for the practical realization of the potential of quantum computation.

In this paper, we investigate the GEB in supervised QML and its validity in the NISQ era. Additionally, we seek to explore the types of algorithms used in QML research, the platform of choice for these algorithm implementations, optimization techniques, datasets, and whether most of the work is theoretical or experimental. Further, we investigate the effectiveness of these QML models by presenting their performance metrics on classical benchmark datasets such as MNIST, Fashion MNIST, and IRIS, along with the number of classes analyzed and the context of the experiments, whether conducted in noisy or ideal settings, as described in Section 3.1.1. We conducted a systematic literature review (SLR) designed as an SMS to achieve these objectives. Research has demonstrated that SMS is a valuable tool for organizing and categorizing existing discoveries while also identifying limitations and gaps for improvements (Kitchenham et al. 2011). This paper is structured as follows: Section 2 provides an overview of the methodology for SLR, including

its protocols. Section 3 presents the result of this SLR, and Section 4 provides the discussion. Finally, in Section 5, we conclude this review with a summary of our findings and their implications.

2 Methodology

SLR in this study followed the Kitchenham et al. (2011) process, which consists of three phases: planning, conducting, and reporting. This approach helped us to systematically identify, analyze, evaluate, and interpret the literature to answer the proposed research questions. In the planning phase, we selected the database for the literature search and defined research questions along with the inclusion and exclusion criteria. We also designed the boolean algebra for the search query. The conducting phase involved identifying relevant articles based on the defined protocols. Finally, the results were interpreted and reported in a structured format in the reporting stage.

For this review, we used two applications: "Publish or Perish"¹ and "Zotero".² "Publish or Perish" was used for searching the literature across multiple platforms, while "Zotero" was used for organizing the literature, checking for duplicates, and generating a bibliography. Additionally, we used Microsoft Excel to keep track of our progress during various stages of the review.

We conducted a systematic literature search across five academic databases to identify relevant papers for our study. The databases included Google Scholar, Scopus, ACM Digital Library, IEEE Xplore, and Semantic Scholar. We collected a total of 678 articles from various platforms, which, after accounting for overlaps and duplicates, resulted in 534 unique publications. We added 10 articles that did not appear in the initial articles but were relevant to this study. These articles were hand-picked during the snowballing process. Among databases, Google Scholar had the highest number of relevant papers, totaling 183, after removing duplicates, representing 33.7% of the entire collection. ACM Digital Library and Semantic Scholar followed, with 108(20%) and 106(19.5%) papers, respectively. Scopus sourced 79 articles, making up 14.54%, while IEEE Xplore contributed 58 articles, which was 10.68% of the total. About 1.6% of articles were manually added. Additionally, we used query strings "Quantum" OR "quantum" AND "Quantum Machine Learning" AND "error bound" AND "noisy" AND "NISQ" or "Quantum Machine Learning" AND "error bound" to search



Fig. 1 The Prisma diagram provides the literature counts at various stages of the SLR process. From the collected paper count of 688, the filtration steps excluded 144 papers, and the eligibility steps excluded 507 papers, resulting in 37 papers for the analysis

for the relevant literature from 2010 to 2023. At the end of the study, we identified 37 articles that met the inclusion criteria. We summarized the identification and selection process using the PRISMA diagram in Fig. 1.

Figures 2, 3, and 4 reveal key trends in the field's research practices. Figure 2 shows an increase in publications number as the year progresses, suggesting a growing research interest. However, only a few of these publications meet the high standards for inclusion in this review, indicating rigorous selection criteria. Figure 3's analysis of the databases shows Google Scholar as the most significant source, but Scopus and Google Scholar lead in quality, as they contribute the most to the final paper selection. This suggests these databases might have more relevant and higher-quality papers. Lastly, Fig. 4 indicates a common trend of 2-5authors per paper, pointing to collaboration in the field, and highlights journal articles as the favored publication type, suggesting a preference for formal, peer-reviewed research channels over conference papers. This analysis highlights the importance of research quantity and quality, the majority of collaborative efforts, and a tendency toward standard publication types.

¹ https://harzing.com/resources/publish-or-perish/

² https://www.zotero.org/



Fig.2 a Bar graph presenting the number of publications per year from 2010 to 2023 among initially extracted papers. It is evident that the number of publications has increased over the years, mainly in the last 5 years. **b** Bar graph for the number of papers extracted per year among the final papers

2.1 Planning

The study primarily focuses on current practice in QML, possible developments, and the difficulties that accompany it, especially regarding the generalization error bound. We formed a team of six members to ensure the study maintains rigor and avoids biases. These members had specialized backgrounds in machine learning, quantum computing, and quantum machine learning. This collective expertise ensured comprehensive coverage and allowed us to view the articles from multiple directions.

The QML field has gathered much attention in the last decade or so. For this reason, we limited our search from 2010





Fig. 3 a Figure showing the number of papers extracted from each database source. Google Scholar was the most significant source for this study. **b** Bar graph for the number of papers extracted from each publisher among the final papers. We can see that the majority of the

papers were published in ArXiv. This graph represents the final papers after the filtration process. Naturally, it is expected that the platform with free access to full-text articles would have the most papers



(a)

Fig. 4 a Bar graph for a count of authors per paper. For most papers, we can see a collective research effort trend of 2-5 authors per paper. **b** Bar graph for the number of papers per publication type. The major-

to 2023.³ We selected the dataset mentioned above because these databases provided access to a large number of articles with full-length searches or customized searches of an article and were available for free via the licenses held by the University. Once the relevant research papers were collected, the next step involved meticulously extracting the information. We focused on details like the authors, publication year, venue, research approach, tools and techniques employed, a platform used for quantum computation, outcomes, challenges, datasets and data encoding strategy, optimization technique, and further research suggestions. This structured extraction method ensured that the data we gathered was comprehensive and easy to interpret and analyze.

2.2 Research questions

In this article, we focused on answering the following questions:

- 1. RQ1: What is the current state-of-the-art generalization error bound for quantum machine learning applied to noisy intermediate-scale quantum (NISQ) devices?
- 2. RQ2: What are the current standard practices in QML in the context of NISQ?
- 3. RQ2a. Is the majority of the research focused on theoretical, empirical, or other approaches?
- 4. RQ2b. How is success measured in QML research: complexity, accuracy, training time, or other metrics?
- 5. RQ2c. What types of datasets are commonly used in QML research: real, synthetic, or others?



ity of the papers were journal articles. This suggests a preference for formal, peer-reviewed research channels over conference papers

6. RQ3. What computing platforms/devices are used in the experiments?

2.3 Inclusion and exclusion criteria

We defined the following inclusion and exclusion criteria for filtering the extracted article. The filtration process was divided into three phases: phase 1, phase 2, and phase 3 (which shall be discussed later). Any article satisfying at least two inclusion criteria in phase 1 moved to phase 2. Also, any article satisfying at least one exclusion criterion was directly excluded in phase 1.

Inclusion criteria:

- 1. Topic: Papers focused on QML in the context of the NISQ era.
- 2. Error bounds: Papers discussing Hoeffding's error bound or related error bounds or noise for QML in NISQ.
- 3. Research approach: Papers presenting theoretical, empirical, or other research approaches relevant to the research questions.

Exclusion criteria:

- 1. Language: Papers not published in English.
- 2. Topic: Papers not focused on QML or NISQ devices.
- 3. Error bounds: Papers that do not discuss generalization error bound, noise in NISQ, or related error bounds.
- Relevance: Papers that do not provide sufficient information or context to address the research questions.
- 5. Publication type: Non-peer-reviewed articles, such as opinion pieces, editorials, or preprints without substantial evidence or contribution to the field or Book.

³ Some of the recently published articles were included as part of the snowballing phase.

6. Duplicates: Papers that have already been included in the review.

2.4 Search process and query string

The initial phase of this review involved gathering relevant publications as described previously. The search process consisted of formulating search queries tailored to each database and successively filtering out papers based on inclusion and exclusion criteria until final selections were made. We defined the following keywords to aid the search process: *quantum machine learning, generalization error; NISQ devices, quantum circuits.* With these keywords defined, we constructed search queries for various platforms as follows:

- 1. Google Scholar: (*all fields*: Quantum OR quantum AND Quantum Machine Learning AND error bound AND data AND noisy AND NISQ. Date: 2010 – 2023)
- 2. Scopus: (*title, abstract, and keywords*: Quantum Machine Learning AND noisy AND NISQ AND PUBYEAR > 2010 AND PUBYEAR <= 2023)
- 3. ACM Digital Library: (*title*: quantum) AND (*Abstract*: quantum machine learning) AND (*E-Publication Date*: (01/01/2010 TO 03/31/2023))
- IEEE Xplore: (*all metadata*: quantum machine learning) AND (*Abstract*: quantum machine learning) AND (*Abstract*: bound). Date:2010 2023.
- 5. Semantic Scholar: (*all fields*: Quantum Machine Learning AND bound AND NISQ, Date: 2010 – 2023)

From these searches, we identified 688 papers. Upon removing duplicates, we had 534 unique articles. Table 1 details the search counts for each platform based on the formulated queries and fields. Figure (2, 3, 4) visually illustrates the publication rates, source-specific paper counts, and distribution of paper types.

Table 1	Search	results	from	various	sources	
---------	--------	---------	------	---------	---------	--

Source	Field	Result count (ini- tial/filtered)
Google Scholar	All fields	183
ACM Digital Library	Title, abstract	108
Semantic Scholar	All fields	106
Scopus	Title, abstract, keywords	79
IEEE Xplore	Abstract	58
Snow balling		10
Duplicates		118
Article published be	fore 2010	26
Total		688/544

The filtration and information extraction were divided into the following phases:

Phase 1: At this stage, each paper was evaluated by two members based solely on its title, abstract, and keywords following the inclusion/exclusion criteria. Papers receiving at least one approval proceeded to phase 2. During this process, we removed 388 irrelevant papers and proceeded with 146 for phase 2.

Phase 2: In this phase, we addressed any inconsistencies that had occurred in phase 1. We flagged that the decision for an article is inconsistent if both the reviewers had different decisions for a paper during the phase 1 review. Of the 146 papers, 46 had inconsistent decisions. Different team members who had not previously worked on inconsistency papers re-evaluated each of these papers using the inclusion/exclusion criteria. This phase concluded with 23 of the 46 disputed papers advancing to phase 3, totaling 123 papers.

Phase 3: Each team member was assigned a subset of papers and was tasked with thoroughly reading assigned papers and determining their suitability for the review. They were thereupon asked to decide if the paper should be included in the study. By the end of this phase, 27 papers were deemed pertinent to our study.

Phase 4: Those who undertook comprehensive paper assessments in phase 3 were then responsible for extracting key data based on a predefined coding schema.

Snow balling: This process included adding the papers that, for some reason, did not show up in the search result or were filtered during the above phase but are relevant to the research. We added 10 papers manually during the review process.

With the final papers selected, we extracted the relevant information from each paper. Table 2 presents the final list of papers. In the next section, we present the results of our analysis of these papers.

3 Result

In this section, we present the results of our SLR. We begin by discussing the datasets and optimization techniques used in QML research. We then present the performance metrics of the QML models and the platforms used for the experiments. We also discuss the research approach, generalization and other relatable error bounds, and the experimental and theoretical nature of the research.

3.1 Dataset

Table 3 shows a wide diversity in both the types of datasets and optimization techniques used in QML research, offering

Table 2 List of final papers

Table 2 Elst of final papers				
Title	Reference	Published year	Source	Publisher
Robust classification with adiabatic quantum optimization	Denchev et al. (2012)	2012	Semantic Scholar	ArXiv
Implementable quantum classifier for nonlinear data	Du et al. (2018)	2018	Google Scholar	ArXiv
Towards quantum machine learning with tensor networks	Huggins et al. (2019)	2019	Google Scholar	ArXiv
Error-mitigated data-driven circuit learning on noisy quantum hard- ware	Hamilton and Pooser (2020)	2020	Scopus	Springer
Quantum classifier with tailored quantum kernel	Blank et al. (2020)	2020	Google Scholar	Nature
Quantum error mitigation with arti- ficial neural network	Kim et al. (2020)	2020	IEEE	IEEE
The Born supremacy: quantum advantage and training of an Ising Born machine	Coyle et al. (2020)	2020	Scopus	NPJ Quantum Inform
Layerwise learning for quantum neural networks	Skolik et al. (2021)	2021	Scopus	Springer
On the expressibility and overfitting of quantum circuit learning	Chen et al. (2021)	2021	ACM	Quantum (ACM)
A rigorous and robust quantum speed-up in supervised machine learning	Liu et al. (2021)	2021	Scopus	Nature Physics
Robust quantum classifier with min- imal overhead	Park et al. (2021)	2021	Scopus	IEEE
I-QER: an intelligent approach towards quantum error reduction	Basu et al. (2022)	2021	ACM	ACM
Can noise on qubits be learned in quantum neural network? A case study on QuantumFlow	Liang et al. (2021)	2021	Scopus	IEEE
Power of data in quantum machine learning	Huang et al. (2021)	2021	Scopus	Nature communications
Generalization in quantum machine learning: a quantum information perspective	Banchi et al. (2021)	2021	Semantic Scholar	APS
Quantum one-class classification with a distance-based classifier	De Oliveira et al. (2021)	2021	IEEE	IEEE
Noise-induced barren plateaus in variational quantum algorithms	Wang et al. (2021)	2021	Google Scholar	Nature
Encoding-dependent generalization bounds for parametrized quantum circuits	Caro et al. (2021)	2021	Snow balling	Quantum
The power of quantum neural net- works	Abbas et al. (2021)	2021	Snow balling	Nature Communication
The inductive bias of quantum ker- nels	Kübler et al. (2021)	2021	Snow balling	NeurIPS
Towards understanding the power of quantum kernels in the NISQ era	Wang et al. (2021)	2021	Scopus	Quantum
Problem-dependent power of quan- tum neural networks on multi-class classification	Du et al. (2022b)	2022	Snow balling	ArXiv
Theoretical error performance anal- ysis for variational quantum circuit- based functional regression	Qi et al. (2023)	2022	Google Scholar	Nature

Table 2 continued

Title	Reference	Published year	Source	Publisher
QOC: quantum on-chip training with parameter shift and gradient pruning	Wang et al. (2022)	2022	Semantic Scholar	ACM
Quantum perceptron revisited: computational-statistical tradeoffs	Roget et al. (2022)	2022	Scopus	AUAI
Implementation and empirical eval- uation of a quantum machine learn- ing pipeline for local classification	Zardini et al. (2022)	2022	Google Scholar	ArXiv
Noisy quantum kernel machines	Heyraud et al. (2022)	2022	Google Scholar	APS
A kernel-based quantum random forest for improved classification	Srikumar et al. (2022)	2022	Google Scholar	ArXiv
Bandwidth enables generalization in quantum kernel models	Canatar et al. (2022)	2022	Snow balling	ArXiv
The dilemma of quantum neural networks	Qian et al. (2022)	2022	IEEE	IEEE
Generalization with quantum geometry for learning unitaries	Haug and Kim (2023)	2023	Snow balling	ArXiv
Ensemble-learning variational shallow-circuit quantum classifiers	Li et al. (2023)	2023	Google Scholar	ArXiv
Generalization in quantum machine learning from few training data	Caro et al. (2022)	2023	Google Scholar	Nature
Understanding quantum machine learning also requires rethinking generalization	Gil-Fuster et al. (2023)	2023	Snow balling	ArXiv
Quantum machine learning beyond kernel methods	Jerbi et al. (2023)	2023	Snow balling	Nature Communications
Out-of-distribution generalization for learning quantum dynamics	Caro et al. (2023)	2023	Snow balling	Nature Communications
Dynamical simulation via quan- tum machine learning with provable generalization	Gibbs et al. (2024)	2024	Snow balling	APS

a glimpse into the current practice of the field. The dataset types vary from synthetic cases, often tailored to quantumspecific issues, to established ones like MNIST (LeCun et al. 2010), Fashion MNIST (Xiao et al. 2017), Iris (Anderson 1936; Fisher 1936), and UCI⁴ datasets. The use of synthetic data in many studies suggests that QML is often operating in a proof-of-concept stage, possibly due to the NISQ constraints. We observed a recurring use of the MNIST dataset and its variants, which, while well-established in classical machine learning, raises questions in the quantum context. It is well known that quantum advantage is not universal but problem-specific (Ball et al. 2020; Herrmann et al. 2023). Therefore, the frequency of MNIST and other classical datasets might unintentionally misguide the field into a comparability trap with classical machine learning. This might hint at an ongoing struggle to balance between specificity and generalizability in QML models. The use of synthetic datasets, while necessary for proof-of-concept,

should be complemented with real-world datasets to ensure the practicality of the models. Studies should ideally focus on problems that are inherently difficult for classical algorithms but are solvable more efficiently on a quantum setup. This observation leads to the question: Are we perhaps focusing on familiar grounds at the expense of uncovering quantum advantage?

On the other hand, the optimization of QML models is a complex task. The use of classical techniques such as stochastic gradient descent (SGD) or backpropagation, and their quantum counterparts, is a topic of ongoing discussion between classical and quantum computation (Lavrijsen et al. 2020; Khairy et al. 2020). However, these classical techniques may not be the most suitable for optimizing quantum circuits, particularly in the presence of quantum noise (Khairy et al. 2020). The optimization landscape of these models is highly non-convex, leading to low convergence rates for SGD. Moreover, the use of Nystrom approximation or Hamming distance-based optimization

⁴ https://archive.ics.uci.edu/

 Table 3
 Comparative overview of datasets and optimization techniques in quantum machine learning research

Reference	Type of dataset	Optimization technique
Basu et al. (2022)	Synthetic: Tabular with quantum gate as features	Grid search
Blank et al. (2020)	Synthetic: Quantum states parametrized by angles	Optimize1qGates (Qiskit)
Du et al. (2018)	Synthetic: Linear and nonlinear dataset generated following (Havlíček et al. 2019)	Parameter-shift rule
Chen et al. (2021)	Real: Iris dataset and synthetic: non-linear	Backpropagation (Simulator) following (Watabe et al. 2019)
Srikumar et al. (2022)	Real: Fashion MNIST, breast cancer and heart dis- eases	Nystrom approximation
Heyraud et al. (2022)	Real: MNIST	Least-square loss function minimization
Zardini et al. (2022)	Real: UCI and Iris	Hamming distance optimization
Qi et al. (2023)	Real: MNIST	SGD with Adam optimizer
Li et al. (2023)	Real: MNIST and synthetic: phase recognition	Automatic differentiation with Adam
Caro et al. (2022)	Synthetic: Ground states phase	SPSA with Matrix Product State
Roget et al. (2022)	Real: Iris and synthetic: following (Mohri et al. 2018)	Gradient descent
Huang et al. (2021)	Real: Fashion-MNIST and synthetic: Engineered dataset	Gradient descent
Liu et al. (2021)	Synthetic: 2D points with hyperplane distance	Convex quadratic optimization
Skolik et al. (2021)	Real: MNIST	Parameter-shift with binary cross-entropy
Liang et al. (2021)	Real: MNIST	Qubit Mapping
Wang et al. (2021)	Real: Fashion-MNIST and synthetic: Engineered dataset	Grid search (regularization parameter)
Coyle et al. (2020)	Synthetic: Engineered dataset (QCIBM)	SGD (parameter-shift rule)
Hamilton and Pooser (2020)	Synthetic: 4-qubit circuit targets	Gradient-based (Adam, parameter-shift)
Huggins et al. (2019)	Real: MNIST	SPSA, finite difference gradient
Wang et al. (2022)	Real: MNIST, Vowel-4	SGD with Adam
Banchi et al. (2021)	Synthetic: 2-Moon	Variational quantum info bottleneck
Denchev et al. (2012)	Real: UCI and synthetic: Long-Servedio, Mease- Wyner	Adiabatic quantum optimization
Kim et al. (2020)	Synthetic: Rand. quant. circuits	Gradient descent (RMSE)
De Oliveira et al. (2021)	Real: Iris	SGD
Qian et al. (2022)	Real: Wine, MNIST	SGD, SQNGD
Wang et al. (2021)	Synthetic: Randomly generated graphs follow- ing (ERDdS and R&wi 1959)	Quantum approximate optimization algorithm
Abbas et al. (2021)	Real: Iris (First two classes)	Cross-entropy loss with Adam
Kübler et al. (2021)	Synthetic: drawn from a uniform distribution on $[0, 2\pi]^d$	Mean square error with Kernel-target Alignment
Canatar et al. (2022)	Real: FMNIST, KMNIST, PLAsTiCC and syn- thetic following (Shaydulin and Wild 2022)	Convex quadratic optimization
Jerbi et al. (2023)	Real: FMNIST	Gradient descent with Adam
Caro et al. (2023)	Synthetic: Random product states	Gradient Free Nelder-Mead
Gibbs et al. (2024)	Synthetic: Har-random product states	Gradient descent
Gil-Fuster et al. (2023)	Synthetic: Generalized cluster Hamiltonian of n qubits	Covariance matrix adaptation evolution strategy
Haug and Kim (2023)	Synthetic: Random product states	Gradient descent
Du et al. (2022b)	Parity and FMNIST	Gradient descent with Adagrad optimizer

methods may be an attempt to circumvent quantum hardware limitations. Furthermore, the intrinsic difficulty of these optimization problems is highlighted by the NP-hard nature of training variational quantum algorithms (Bittel and Kliesch 2021). Even shallow variational quantum models, devoid of barren plateaus (Marrero et al. 2021; Arrasmith et al. 2021; Wang et al. 2021; McClean et al. 2018; Holmes et al. 2022; Zhao and Gao 2021), have a superpolynomially small fraction of local minima within any constant energy from the global minimum (Anschuetz and Kiani 2022). Just as in classical machine learning (Bermeitinger et al. 2019; Skorski et al. 2021), these models become untrainable without an appropriate initial estimate of the optimal parameters. In addition, the exponential suppression of cost function differences in a barren plateau hampers the progress of gradient-free optimizers without exponential precision (Arrasmith et al. 2021). Learning can also be hindered without multiple copies of a state or if there is an excess of entanglement within the circuit (Abbas et al. 2024; Marrero et al. 2021). These challenges underscore the complexity of the optimization landscape in QML and the need for further research. While innovative, these methods should be critically assessed to ensure they do not compromise the potential advantages of a fully quantum approach, as they could result in quantum solutions that are neither faster nor more accurate than their classical counterparts (Dunjko and Briegel 2018).

3.1.1 Algorithm performance

In this section, we present the performance of various models proposed among the selected papers. We focus on the performance of these models on the classical dataset. A quick glimpse of these datasets is presented in Table 3 and their performance in Table 4. We conveniently mapped the Fashion MNIST or FMINIST classes to 0 - 9 classes. When a study involves a subset of classes, it is indicated in the table under the number of classes column, e.g., (3 - (3,6,8)) represents a work focused on three specific classes-3, 6, and 8. We also note that accuracies are presented in percentages, and when individual noise rates were specified in the original works, we include them for context. It is apparent from Table 4 that the type and level of noise are pivotal in affecting the performance of the models. It is also evident that the performance of these models is sensitive to the presence of noise. For instance, the model by Liang et al. (2021) shows an impressive 98.04% accuracy on MNIST in a noiseless environment, but that number drops significantly under both flip error and phase error conditions. This sort of degradation is not unique and appears across multiple works, emphasizing the impact that different kinds of quantum noise can have on model performance. Similarly, Chen et al. (2021) model on the Iris dataset experiences an accuracy of 80.7% under gate noise conditions, suggesting that even a relatively low noise level can have a measurable impact. The same pattern is observed in work by Wang et al. (2021) on the Fashion MNIST dataset, where the accuracy drops from 96% in a noiseless setting to 91.2% under a depolarizing rate of 0.05. Additionally, models trained on actual quantum hardware generally have lower accuracies compared to those trained in noiseless or simulated noisy environments.

This suggests that how a model is designed could be integral in mitigating specific noise types. More importantly, these shifts in accuracy due to the presence of noise highlight the challenges of operating in the NISQ era, mainly when the noise rates are non-negligible. Additionally, the performance accuracy of models trained on real quantum hardware raises questions about hardware-specific optimizations and the challenges this presents for reproducibility. The performance also seems to vary when only subsets of classes are considered, as seen in work like Wang et al. (2021) and Huggins et al. (2019); this often raises concerns about the applicability of these models to real-world scenarios where class distributions are often imbalanced. It is also interesting to note that despite its simplicity, the Iris dataset tends to yield lower accuracies than more complex datasets like MNIST. Simpler datasets are often prone to be impacted by noise, and this is reflected in the performance of the models (Khanal and Rivas 2023). This suggests that the performance of QML models is not only sensitive to the type and level of noise but also to the complexity of the dataset.

3.2 Bounds

In this section, we discuss different bounds and complexities proposed in the selected literature. This focus is to discuss the theoretical guarantees in QML proposed across literature under our inclusion criteria for model performance. The authors have proposed numerous bounds in different categories to provide a robust framework for evaluating the performance of QML models. One of the most prominent categories is the generalization bound, which is a metric central to any machine learning task. We encourage readers to refer to the original works for a detailed derivation for each bound. In this paper, we discuss common properties observed across these bounds.

Generalization bound is an essential quantitative measure for assessing how well a model is expected to generalize to unseen data. In a QML setting, this bound often offers rich insights into the intricate interplay between quantum and classical computational resources beyond performance indicators. This contributes to our understanding of the capabilities and limitations of QML algorithms. A careful analysis of the generalization bounds provided in Abbas et al. (2021); Banchi et al. (2021); Caro et al. (2021, 2022, 2023); Gibbs et al. (2024); Gil-Fuster et al. (2023); Huang et al. (2021); Liu et al. (2021); Qi et al. (2023) reveals a universal dependency on the dataset size N. This dependency aligns with wellestablished understandings in classical machine learning that larger datasets result in better model generalization (Abu-Mostafa et al. 2012), reducing the model's uncertainty and error on unseen data. It is worth noting, however, that the influence of N on the bounds is not uniform across the board: the magnitude of its impact varies depending on other model parameters and method-specific assumptions. Furthermore, the generalization improvements appear to follow a sublinear trend, as most bounds show a $\mathcal{O}(\sqrt{N})$ behavior with respect to N. Additionally, these bounds frequently incorporate specific model parameters-such as the Hilbert space dimension d in Abbas et al. (2021); Caro et al. (2021); Chen et al. (2021); Huang et al. (2021); Jerbi et al. (2023); Qi et al. (2023), the number of trainable quantum circuit gates

T in Caro et al. (2022, 2023), and parameters w in Liu et al. (2021). The bounds from Abbas et al. (2021); Canatar et al. (2022); Caro et al. (2021); Gibbs et al. (2024); Heyraud et al. (2022); Oi et al. (2023) are additionally bounded in variables that are method specific. This implies two things: firstly, these bounds are often tailored to the specific algorithmic techniques or problem domains they are designed for, and secondly, the bounds suggest avenues for model fine-tuning, particularly by adjusting these specific parameters. Another noteworthy observation is that Quantum Kernel Theory is a recurring approach across many proposed generalization bounds (Blank et al. 2020; Canatar et al. 2022; Heyraud et al. 2022; Huang et al. 2021; Kübler et al. 2021; Liu et al. 2021; Wang et al. 2021). In QML landscape, kernel theory appears to be serving as a foundational technique for constructing algorithms with both robust performance and theoretically justifiable generalization guarantees (Schuld and Killoran

Table 4 Prediction accuracy of various works on the MNIST and IRIS datasets

References	Dataset	Number of classes	Accuracy	Setting
Heyraud et al. (2022)	MNIST	3 - (3,6,8)	94.50	Decoherence (spin dephasing)
Li et al. (2023)	MNIST	4 - (1,3,5,7)	95.00	Noiseless
Skolik et al. (2021)	MNIST	2 - (6,9)	90.00	Noiseless
			73.00	Shot noise
Liang et al. (2021)	MNIST	10	98.04	Noiseless
			88.24	Flip error (0.01)
			91.67	Phase error (0.01)
			77.78	Phase + flip error (0.01)
Huggins et al. (2019)	MNIST	2 ^a -(4,9)	88.00	Noiseless
			80.60	Amplitude(0.04) and dephasing(0.03) noise
Wang et al. (2022)	MNIST	4-(0,1,2,3)	63.70	Trained on-chip at ibmq_jakarta
		2-(3,6)	86.00	
Qian et al. (2022)	MNIST	10	94.00	Noiseless
			80.00	Gate noise
Srikumar et al. (2022)	FMNIST	10	93.30	Noiseless
Wang et al. (2021)	FMNIST	2 - (0,3)	96.00	Noiseless
			91.20	Depolarizing rate(0.05)
Wang et al. (2022)	FMNIST	4 - (0,1,2,3)	57.00	Trained on ibmq_manila
		2 - (3,6)	90.70	Trained on ibmq_santiago
Chen et al. (2021)	IRIS	3	80.70	Gate noise(0.01)
De Oliveira et al. (2021)	IRIS	2-(0,1)	98.89	Noiseless
Abbas et al. (2021)	IRIS	2-(0-1)	23.14 ^b	Trained on ibmq_montreal
	FMNIST	2	92.60	
Canatar et al. (2022)	KMNIST	2	91.50	Noiseless
	PLAsTiCC	2	78.90	
Du et al. (2022b)	FMNIST	9	50.00	Noiseless

^aSection IV on Huggins et al. (2019) presents results for 45 cases (10c2) combinations. Here, we only report the result that is severely impacted by the noise

^bIt is training loss of a model as reported in Abbas et al. (2021)

2019; Schuld 2021). However, it is crucial to emphasize that the effectiveness of kernel methods trainability guarantees, due to convex loss landscapes (Schuld 2021; Schuld and Killoran 2022), hinges on the efficient estimation of kernel values to a sufficient precision (Thanasilp et al. 2024). This is particularly challenging because, similar to the Barren plateau barrier in QNNs, hardware-induced noise in the near-term serves as a source of concentration for quantum kernel values to be exponentially concentrated towards some fixed value over different input data (Thanasilp et al. 2024).

On the other hand, measurement complexity, defined as the number and nature of quantum measurements required to extract classical information from a quantum system, plays a crucial role in determining the generalization capabilities of QML models, particularly in noisy quantum systems. Caro et al. (2021) establish a fundamental trade-off between the complexity of measurement observables and the amount of training data required for the effective generalization of ONNs. Their work demonstrates that while more complex measurements can enhance the expressivity of QNNs, they simultaneously demand larger training datasets to achieve robust generalization. This relationship is further explored in the context of quantum kernel methods, where Gentinetta et al. (2024); Liu et al. (2021) provide rigorous bounds on the number of measurement shots required to successfully train fidelity-based kernels (Havlíček et al. 2019). Furthermore, Wang et al. (2021) offer an optimistic perspective, showing that quantum kernel generalization can remain competitive with ideal scenarios when the number of measurements scales as $\mathcal{O}(N^3)$ provided the noise rate p remains low. However, Thanasilp et al. (2024) present a contrasting view, demonstrating that under conditions of exponential concentration in quantum kernel values, the required number of measurement shots for precise kernel estimation scales exponentially. This dichotomy underscores the critical nature of measurement complexity in QML model performance, especially in noisy environments. Additionally, measurement complexity bounds proposed in Blank et al. (2020); Park et al. (2021) offer insights into the algorithm's resource requirements, specifically, the number of quantum measurements needed to maintain a certain level of accuracy in kernel-based quantum classifier models. These bounds exhibit a dependence on p. The bound $\mathcal{O}\left(\frac{1}{(1-2p)^2}\right)$ from Blank et al. (2020), valid for p < 0.5, reveals a quadratic dependency on the noise rate p. This relationship implies that even small increases in noise can substantially elevate the measurement complexity, potentially limiting the algorithm's practicality in noisy environments. It is important to acknowledge that these bounds are primarily applicable within the specific context of kernel-based quantum models and may not necessarily extend to other QML frameworks.

Regardless, collectively, these works emphasize the delicate balance between measurement complexity, noise tolerance, and generalization performance in QML models.

Next, we discuss the query and runtime complexities, as these factors are instrumental in determining the practical feasibility of quantum approaches. Query complexity, which quantifies the number of interactions between an algorithm and an oracle or database, provides insight into the information-theoretic efficiency of quantum algorithms. Du et al. (2018) proposed a query complexity bound of $\mathcal{O}(poly(\log(d)\sqrt{N}))$ where d represents the feature space dimension, and N is the dataset size. The bound exhibits a potential quantum advantage, as it scales polynomially with $\log(d)$ and only with the \sqrt{N} , potentially outperforming classical algorithms for high-dimensional data. However, this advantage must be weighed against the challenges posed by measurement complexity in noisy quantum systems, as discussed earlier. Runtime complexity, on the other hand, directly reflects the required computational time. The proposed runtime complexity bound $\mathcal{O}(poly(\log d \log(d \log N)))\sqrt{\log N})$ by Du et al. (2018) demonstrates a more intricate scaling behavior, with polynomial dependencies on logarithmic terms of both d and N. While this scaling is generally favorable compared to many classical algorithms, especially for large datasets, it is important to note that the actual performance advantage can be mitigated by the overheads associated with quantum state preparation and measurement, particularly in near-term devices.

Furthermore, we observed that the authors also employ the Vapnik-Chervonenkis (VC) dimension to define generalization bounds. VC dimension plays a crucial role in understanding a model's capacity to generalize, providing an upper limit on the complexity of learnable functions. The VC dimension, defined as the largest number of points that a hypothesis class \mathcal{H} can shatter (i.e., perfectly separate regardless of their labeling), takes on a distinctive form in quantum systems. Chen et al. (2021) established a VC bound for quantum models: $2 \le d_{vc} \le (2\frac{n}{d}+1)^{2d}$, where d is the feature dimension and n is the number of qubits. This bound highlights a fundamental difference from classical models: the VC dimension in quantum settings depends explicitly on the number of qubits, suggesting that the expressive power of quantum models scales with the size of the quantum system. Moreover, in noisy quantum environments, the VC bound incorporates an additional dependency on the circuit depth L_c , as observed in works by Abbas et al. (2021); Caro et al. (2021, 2023, 2022). This three-way dependency on feature dimension, qubit count, and circuit depth represents a significant departure from classical machine learning, where generalization typically depends primarily on the feature

space dimension and sample size. The inclusion of circuit depth in the VC bound for noisy quantum systems emphasizes the intricate relationship between model complexity, noise, and generalization in QML. It suggests that deeper quantum circuits, while potentially more expressive, may face greater challenges in generalization, especially in the presence of noise. Furthermore, these bounds hint at both potential advantages and challenges for QML models. On one hand, the dependence on qubit count suggests that quantum models might offer enhanced expressive power that scales efficiently with system size. On the other hand, the sensitivity to circuit depth in noisy settings underscores the challenges of maintaining this expressivity in practical, noisy quantum devices.

3.3 Computing platforms

The choice of quantum computing platforms is a crucial aspect of QML research, as it directly impacts the experimental feasibility and the generalizability of the results. In Table 5, we provide the list of quantum computing platforms used for an experiment by the work listed in Table 2. The IBM Quantum Platform, such as Melbourne, Ourense, Rome, and Montreal, appears to be quite popular. The Melbourne processor, a 15 qubits system retired on 08/09/2021, tends to be the most utilized in the IBM quantum series, likely due to its higher qubit count and may be due to its early market entry. Specifically, Melbourne is used in works by Basu et al. (2022); Srikumar et al. (2022); Wang et al. (2021); Kim et al. (2020); De Oliveira et al. (2021), Ourense by Blank et al. (2020); Wang et al. (2021), Rome by Park et al. (2021), and Montreal by Abbas et al. (2021); Liang et al. (2021); Wang et al. (2021). While using these platforms is understandable, given their accessibility and the extensive support provided by IBM, it is important to note that the choice of platform can impact the results and potentially create research bias. For instance, the noise rates and error models of these platforms can vary, leading to different performance outcomes for the same model.

Another crucial observation is the use of Qiskit, a popular quantum computing software development kit (SDK) by IBM. Qiskit is used in works by Basu et al. (2022); Blank et al. (2020); Canatar et al. (2022); Chen et al. (2021); De Oliveira et al. (2021); Kim et al. (2020); Liang et al. (2021); Srikumar et al. (2022); Zardini et al. (2022). This is not surprising, given that Qiskit is one of the most widely used quantum computing frameworks. Software simulation often provides the first line of feasibility tests for quantum algorithms. However, these results can be optimistic compared to the NISQ devices since the simulators usually lack noise models under normal settings. Furthermore, the use of Pennylane only in Kübler et al. (2021); Qi et al. (2023); Qian et al. (2022) is

intriguing, mainly because Pennylane's focus on differentiable quantum computing makes it particularly well-suited for hybrid quantum-classical models. Additionally, the use of Julia in Li et al. (2023) and TensorFlow Quantum in Huang et al. (2021); Skolik et al. (2021) might hint at a move towards using more traditional machine learning frameworks. However, one could argue that these choices should not merely be about convenience or the ease of integration with classical models. They should also be critically evaluated for their capability to handle quantum-specific issues, such as error correction or the intricacies of quantum gate operations.

It is imperative to acknowledge, however, that this distribution does not necessarily mirror the broader QML community's platform preferences. Preliminary observations and informal surveys within the community suggest a significant and possibly growing interest in platforms like PennyLane, Tensorflow quantum, and QuTip, which may not be fully represented in our dataset. Hevia et al. (2022); Kordzanganeh et al. (2023); Serrano et al. (2022) are some of the works that provide a comprehensive study on quantum computing platforms used in literature.

In our findings, the platform selections are far from arbitrary, influenced by factors such as ease of use, capability, and perhaps even academic and commercial affiliations. Observing these diverse platforms used in the experiment raises the issue of reproducibility. How many of these papers provide adequate information for replicating their experiments on other platforms?⁵ The field risks becoming fragmented if results obtained on one platform cannot be compared or reproduced on another.

3.4 Research approach

Among the selected papers, we find that the research approaches are diverse, with a mix of theoretical and empirical work. Most of the research appears to use a theoretical approach with empirical validation. Such an approach is necessary, especially in the NISQ era, where empirical work can offer immediate insights into error rates, robustness, and other practical considerations that are crucial for theoretical generalizations. It is interesting to note that various articles address core concerns in machine learning from a quantum perspective, such as generalization bound (Abbas et al. 2021; Banchi et al. 2021; Canatar et al. 2022; Caro et al. 2022, 2023; Chen et al. 2021; Gibbs et al. 2024; Gil-Fuster et al. 2023; Huang et al. 2021; Wang et al. 2022; Heyraud et al. 2022; Jerbi et al. 2023; Kübler et al. 2021; Srikumar et al. 2022; Wang

⁵ Note that Pennylane has various plugins for an accessible dispatch of quantum functions to different quantum devices.

Table 5 Various quantum comput	ting platforms	used for experiments in literature	for QML experin	nents					
Reference	Qiskit	IBM Quantum Platform ^a	PyQuil	GENCI	Pennylane	Julia	Tensorflow	MGCF	QuTip
Basu et al. (2022)	>	>							
Du et al. (2018)			>						
Blank et al. (2020)	>	~							
Chen et al. (2021)	>								
Srikumar et al. (2022)	>	~							
Heyraud et al. (2022)				>					
Zardini et al. (2022)	>								
Qi et al. (2023)					>				
Li et al. (2023)						>			
Huang et al. (2021)							>		
Park et al. (2021)		>							
Skolik et al. (2021)							>		
Liang et al. (2021)	>	~							
Wang et al. (2021)	>	~							
Coyle et al. (2020)			>						
Hamilton and Pooser (2020)		~							
Huggins et al. (2019)								>	
Wang et al. (2022)		>							
Kim et al. (2020)	>	~							
De Oliveira et al. (2021)	>	~							
Qian et al. (2022)		>			>				
Wang et al. (2021)		~							
Abbas et al. (2021)		>							
Kübler et al. (2021)					>				
Canatar et al. (2022)	>								
Jerbi et al. (2023)							>		
Caro et al. (2023)		~							
Gibbs et al. (2024)		~							
Haug and Kim (2023)									>
Gil-Fuster et al. (2023)							>		
^a Refers to multiple IBM Quantum	1 devices. Deta	iils on these platforms are provided	1 in Section 3.3						

🖄 Springer

et al. 2021),⁶ and ensemble learning (Basu et al. 2022; Caro et al. 2023; Li et al. 2023; Srikumar et al. 2022). Unlike other works discussed in this review, Jerbi et al. (2023) provide the lower bound for qubit complexity for QNNs in an ideal setting. It is no surprise that these bounds are expressed in terms of the feature space dimension d. This suggests that the field is actively working towards addressing foundational learning problems, such as the true capabilities and limitations of QML models and their robustness and error tolerance in a NISQ environment. However, given that many of these approaches combine theory and experiment, one can hypothesize that the field is still working towards addressing foundational learning problems, such as the true capabilities and limitations of QML models and their robustness and error tolerance in a NISQ environment. We provide such empirical results on classical data in Table 4.

Another interesting observation is the use of quantum kernel methods. Kernel methods seem to be an attempt to utilize classical machine learning techniques in quantum architectures to benefit from the mathematical rigor of kernel theory while aiming to harness quantum advantages. This could offer better generalization for QML models (Schuld and Petruccione 2021). Recent research has been motivated towards the quantum kernel-based theoretical and experimental advancement, with studies establishing a connection between supervised learning and quantum kernels (Schuld and Killoran 2019). For a class of machine learning problems, quantum kernel methods can solve them efficiently, which is hard for all classical methods (Liu et al. 2021). Furthermore, expressivity and generalization capacity of quantum kernels have been investigated in studies by Chen et al. (2021); Heyraud et al. (2022). Huang et al. (2021) found that data availability can modulate the computational hardness of learning tasks. However, Heyraud et al. (2022); Wang et al. (2021) reveal that the inner products or fidelity measures constituting quantum kernels can be particularly susceptible to noise, thereby affecting their overall performance and the reliability of the QML models built upon them.

Nonetheless, while quantum kernels have the potential to be advantageous in NISQ settings due to their ability to find better or equally good quantum models compared to variational circuit training (Schuld 2021), it is important to acknowledge that they might suffer from exponential concentration under certain conditions (Thanasilp et al. 2024). This phenomenon, while not necessarily affecting the trainability of quantum kernels, can lead to poor generalization, where the model's predictions on unseen data become independent of the input data, thereby undermining the expected advantages. Thanasilp et al. (2024) identified the expressivity of data embedding, global measurements, entanglement, and noise as a source of the concentration in quantum kernel models. This trade-off between the optimization advantages of quantum kernels and their potential generalization challenges is a careful consideration for kernel model design and implementation.

Furthermore, the efficiency of data encoding plays a crucial role, with compact encoding schemes potentially minimizing the number of gates required (Gan et al. 2023; Schuld et al. 2021). While both quantum kernel methods and QNNs utilize quantum circuits to encode data into quantum states, the role and function of these circuits differ between the two approaches. In quantum kernel methods, the quantum feature map is explicitly designed to project data into a highdimensional quantum space, where classical algorithms may then operate more effectively, depending on the problem. This can result in simpler classification circuits, as the complexity is handled by the classical algorithm post-quantum feature mapping. In contrast, QNNs embed the data encoding within a parameterized quantum circuit, where the entire model, including the quantum feature map, is optimized during training. As a result, the simplification observed in quantum kernel methods may not directly apply to QNNs. This potential for shallower circuits and fewer gates in quantum kernel methods aligns well with the limitations of NISQ devices, where circuit depth is constrained by noise (Wang et al. 2021). However, it is important to note that the best choice between quantum kernels and other QML methods depends on the specific dataset, data-encoding ansatz, learning problem, and the available hardware (Jerbi et al. 2023). On the other hand, QNNs can offer significant expressive power and flexibility in learning complex patterns (Abbas et al. 2021). However, this expressivity frequently comes with increased computation resource requirements, especially with deeper architectures (Qian et al. 2022; Skolik et al. 2021). This can pose challenges for NISQ devices. Furthermore, concerns exist regarding the learnability and trainability of QNNs in noisy settings (Du et al. 2021; Qian et al. 2022). In contrast, quantum kernel methods often benefit from guarantees of convex optimization and potential advantages in resource utilization (Park et al. 2021) but suffer from exponential concentration (Thanasilp et al. 2024). Yet, their data scaling limitations and reliance on specific kernel functions can restrict their applicability. Additionally, data requirements for QNNs often scale with circuit depth and problem complexity, potentially becoming substantial (Skolik et al. 2021).

Another research approach that appears in the final paper list is ensemble learning for error mitigation. Li et al. (2023); Srikumar et al. (2022) used this approach for error mitigation, suggesting that it might be a practical strategy to make quantum algorithms more robust. However, ensemble methods inherently require the collection of multiple models, which could be resource-intensive. We do not find any work explic-

⁶ Jerbi et al. (2023) provide a sample complexity which can be converted in to generalization bound in most of the cases.

itly addressing the trade-off between improved performance and increased resource utilization.

4 Discussion

In this section, we discuss the implications of our findings and the limitations of our methodology. We discuss the interplay between generalization bounds and others' complexity, the pitfalls in a dataset and optimization choices, platform standardization vs. research bias, the theoretical–experimental divide, fragmented research approaches, and the implications for the NISQ era. Our analysis suggests that there is a gap in the techniques and tendencies that could advance QML research in certain directions, both limiting and enabling.

The generalization error bound is interlinked with the measurement, sample, and qubits complexity. While generalization bounds give a theoretical metric to evaluate model robustness, qubits and sample complexity provide a practical measure of the resources required to achieve this robustness, their practical utility depends on the feasibility of measurement. It is striking that all these metrics are yet to be optimized together in the literature, an oversight that could hinder the transition from theory to practice. What remains unavailable is a framework that can combine these aspects, allowing for not just error prediction but also its empirical verification on actual quantum hardware.

The use of classical datasets like MNIST and IRIS raises the concern of what we term as familiarity bias. This leads to inevitable comparisons with classical algorithms, obscuring the specific advantages that quantum algorithms may bring. Concurrently, the dataset choices show a clear division between synthetic datasets and real-world datasets, with each having its own set of advantages and disadvantages. On one side, synthetic datasets may allow for a deeper understanding of quantum-specific phenomena but at the expense of broader applicability. On the other hand, classical datasets may not necessarily help demonstrate the quantum advantage.

The frequent use of IBM's quantum platforms suggests a potential trend toward platform standardization, but it also raises a concern about research bias. While standardization can facilitate result comparison and replication, the field should be wary of a one-size-fits-all approach. Different platforms have varying noise models, gate fidelities, and connectivity architectures that can significantly impact algorithm performance and, hence, the generalizability of the research findings.

The majority of theoretical work with empirical validation indicates a field that is cautiously optimistic. However, it is essential to question whether the empirical work truly validates the theory or simply provides a proof of concept under idealized/simulated conditions. The field needs to bridge the gap between theory and practice, ensuring that theoretical insights are validated in realistic settings.

Furthermore, the evident focus on kernel methods indicates a field searching for a stable theoretical foundation. Kernel methods offer mathematical rigor but could be susceptible to noise, whereas ensemble learning shows promise in error mitigation but may raise questions on resource optimization. These diverging approaches risk fragmenting the field unless they are part of a more extensive, unified strategy.

The QML community has yet to fully adapt to a unique set of constraints imposed by the NISQ era. While work on generalization bounds and other complexities indirectly acknowledges these limitations, direct strategies to navigate the NISQ landscape are conspicuously absent. The modifications or quantum patches to classical techniques (like SGD) can be seen as stop-gap measures but hardly a long-term solution.

Our analysis suggests a more unified and collaborative approach. This could involve sharing across generalization bounds, measurement complexity, dataset selection, optimization techniques, and platform-agnostic strategies. Caro et al. (2023) proposed to establish a general mathematical framework for quantum learning. However, it is essential to recognize that the diversity of machine learning models is a strength, not a weakness. Each model offers unique benefits and limitations, and there is no one-size-fits-all solution. The right model depends on the problem at hand.

The methodology of this survey is based on a systematic literature review, which is a well-established method for synthesizing and analyzing information from previously published literature. However, it is important to highlight that our search and filtration process is subject to certain constraints. The scope of our search was limited to certain academic databases and English-language publications, which may have inadvertently excluded relevant international research, details on Sections 2.1 and 2.3. Furthermore, the database search was conducted in early 2023 to collect the relevant papers from 2010 to 2023⁷ timeframe. The field of QML is rapidly evolving. Therefore, it is possible that some recent research may not have been included in our analysis. Furthermore, our inclusion and exclusion criteria were designed to maintain focus and relevance, but they may also introduce a degree of selection bias. By recognizing these limitations, we aim to ensure that the findings of our review are interpreted with an appropriate level of scrutiny and consideration for the broader research landscape.

 $^{^{7}}$ Some of the latest papers were hand-picked as the snowballing process.

5 Conclusion

This survey offered a comprehensive analysis of the current state of QML, focusing on generalization bounds, measurement complexities, datasets, optimization techniques, platforms, and research approaches. Our findings reveal a field's formative stages, struggling to balance theoretical robustness and practical applicability.

Generalization bounds are a pivotal element in evaluating QML models, revealing a dependency on dataset size, feature space dimension, and model-specific parameters. These bounds have also been closely tied to measurement sample and qubits complexities, other vital aspects of assessing the practicality of QML, especially in a NISQ environment. In practice, the descent number of research work (Canatar et al. 2022; Heyraud et al. 2022; Liang et al. 2021; Huggins et al. 2019; Wang et al. 2022, 2021) shows a trend toward using well-known classical datasets and optimization techniques. We presented the performance metrics of these works on IRIS, MNIST, and FMNIST datasets with gradient descent being the favorite choice for optimization. This focus raises questions about the field's ability to tackle quantum-specific challenges, calling for broader dataset and algorithmic diversity. It is worth noting that the use of classical datasets and optimization techniques is not necessarily a limitation but rather an intermediate step to understanding the possibilities and limitations of QML in the NISQ era. However, it is important to mention that there has been a shift to develop GEB in quantum setting with an experiment on quantum data (Caro et al. 2021, 2023; Gibbs et al. 2024; Gil-Fuster et al. 2023; Haug and Kim 2023). Despite the diverse directions of quantum learning theory, Caro et al. (2023) have propose a formalism for describing quantum learning tasks that involve training on classical-quantum data and then testing the learned hypothesis on new data. These new directions show that the QML community is moving towards more quantum-specific approaches.

Likewise, the frequent use of specific platforms, particularly from IBM, hints at a potential research bias and raises questions about cross-platform reproducibility. The research approaches reflect a blend of theoretical and empirical work, with a focus on foundational problems in machine learning. Prominent techniques include kernel methods, ensemble learning, and randomized circuit learning. Yet, several unanswered questions remain, especially concerning the trade-offs between performance and resource utilization. Du et al. (2022b) have identified that the multi-class classification power of QNN classifiers is dominated by the training loss rather than generalization ability. Gil-Fuster et al. (2023) have challenged current understandings of generalization in QML. They presented a compelling result that showed that state-of-the-art QNNs can memorize random data, defying standard theories about how these models generalize. This questions the ability of current frameworks to guarantee how well a model performs on new data. Their work is key to understanding and making better QML algorithms.

The field of QML presents both opportunities and challenges. This survey focused on generalization bounds for supervised QML, but the field is actively exploring other branches, such as model expressibility and trainability. Du et al. (2022a) analyzed the expressivity of VQCs using qudits, finding an upper bound determined by the number of quantum gates and measurement observables used in the ansatz. Further research into qudit-based quantum computation and the expressibility and training of QML models can be promising future directions for the field.

As we navigate the intricacies of quantum computing, the task ahead is to address the gaps identified in this survey, striving for methodological diversity, practical applicability, and cross-platform standardization. We have identified several avenues for future work:

- A closer investigation of the trade-offs between generalization and computational costs.
- The need for more diverse datasets tailored for quantum phenomena.
- Scrutiny of the impact of platform biases on research outcomes.
- Evaluation of optimization techniques specific to quantum systems.
- Exploration of resource utilization in ensemble learning and other error mitigation strategies.
- A comprehensive theoretical development to lay the foundation for the unified understanding of quantum kernels (Gan et al. 2023) and quantum learning (Caro et al. 2023).

By confronting these challenges, the QML community can move closer to realizing the full potential of quantum computing in machine learning applications.

Acknowledgements The authors want to thank the Department of Computer Science at Baylor University for their support during the execution of the research. This work is partially supported by the National Science Foundation under Grant Nos. 2136961, 2210091, and 1905043. The views expressed herein are solely those of the author and do not necessarily reflect those of the National Science Foundation.

Author Contributions P.R. conceptualized the systematic study and the research questions. B.K. executed the entire study and prepared the manuscript. P.R. mentored B.K. during the process and reviewed the manuscript. K.S., A.S., and E.Q. helped during phases 1, 2, and 3 of the mapping study. A.R. helped during phases 1 and 2 of the mapping study. A.S. helped generate the figures and collaborated in the snowballing process. K.S. helped with proofreading the manuscript and the snowballing process.

Funding Part of this work was performed while P.R. was funded by the National Science Foundation under Grant Nos. 2136961, and 2210091, and 1905043. The views expressed herein are solely those of the author(s) and do not necessarily reflect those of the National Science Foundation.

Data Availability No datasets were generated or analyzed during the current study.

Declarations

Conflict of interest The authors declare no competing interests.

References

- Abbas A, King R, Huang HY, Huggins WJ, Movassagh R, Gilboa D, McClean J (2024) On quantum backpropagation, information reuse, and cheating measurement collapse. Adv Neural Inf Process Syst 36
- Abbas A, Sutter D, Zoufal C, Lucchi A, Figalli A, Woerner S (2021) The power of quantum neural networks. Nat Comput Sci 1(6):403–409
- Abu-Mostafa Y, Magdon-Ismail M, Lin HT. (2012) Learning from data: a short course. AML Book
- Alam M, Ghosh S (2022) Qnet: A scalable and noise-resilient quantum neural network architecture for noisy intermediate-scale quantum computers. Front Phys 9:702
- Anderson E (1936) The species problem in iris. Ann Missouri Bot Gard 23(3):457–509
- Anschuetz ER, Kiani BT (2022) Quantum variational algorithms are swamped with traps. Nat Commun 13(1):7760
- Arrasmith A, Cerezo M, Czarnik P, Cincio L, Coles PJ (2021) Effect of barren plateaus on gradient-free optimization. Quantum 5:558
- Arunachalam S, de Wolf R (2017) Guest column: a survey of quantum learning theory. ACM Sigact News 48(2):41–67
- Ball P et al (2020) Physicists in China challenge Google's quantum advantage'. Nature 588(7838):380
- Banchi L, Pereira J, Pirandola S (2021) Generalization in quantum machine learning: a quantum information standpoint. PRX Quantum 2(4):040321
- Basu S, Saha A, Chakrabarti A, Sur-Kolay S (2022) i-QER: an intelligent approach towards quantum error reduction. ACM Trans Quantum Comput 3(4):1–18
- Benedetti M, Lloyd E, Sack S, Fiorentini M (2019) Parameterized quantum circuits as machine learning models. Quantum Sci Technol 4(4):043001
- Bermeitinger B, Hrycej T, Handschuh S (2019) Singular value decomposition and neural networks. Artificial neural networks and machine learning-icann 2019: Deep learning: 28th international conference on artificial neural networks, munich, germany, September 17–19, 2019, proceedings, part ii 28 pp 153–164
- Bharti K, Cervera-Lierta A, Kyaw T, Haug T, Alperin-Lea S, Anand A, et al (2021) Noisy intermediate-scale quantum (NISQ) algorithms (2021). arXiv:2101.08448
- Biamonte J, Wittek P, Pancotti N, Rebentrost P, Wiebe N, Lloyd S (2017) Quantum machine learning. Nature 549(7671):195–202
- Bittel L, Kliesch M (2021) Training variational quantum algorithms is NP-hard. Phys Rev Lett 127(12):120502
- Blank C, Park DK, Rhee JKK, Petruccione F (2020) Quantum classifier with tailored quantum kernel. NPJ Quantum Inf 6(1):41
- Canatar A, Peters E, Pehlevan C, Wild SM, Shaydulin R (2022) Bandwidth enables generalization in quantum kernel models. arXiv:2206.06686

- Carleo G, Cirac I, Cranmer K, Daudet L, Schuld M, Tishby N, Zdeborová L (2019) Machine learning and the physical sciences. Rev Mod Phys 91(4):045002
- Caro M, Gur T, Rouzé C, Franca DS, Subramanian S (2023) Information-theoretic generalization bounds for learning from quantum data. arXiv:2311.05529
- Caro MC, Gil-Fuster E, Meyer JJ, Eisert J, Sweke R (2021) Encodingdependent generalization bounds for parametrized quantum circuits. Quantum 5:582
- Caro MC, Huang HY, Cerezo M, Sharma K, Sornborger A, Cincio L, Coles PJ (2022) Generalization in quantum machine learning from few training data. Nat Commun 13(1):4919
- Caro MC, Huang HY, Ezzell N, Gibbs J, Sornborger AT, Cincio L, Holmes Z (2023) Out-of-distribution generalization for learning quantum dynamics. Nat Commun 14(1):3751
- Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K et al (2021) Variational quantum algorithms. Nat Rev Phys 3(9):625– 644
- Chen CC, Watabe M, Shiba K, Sogabe M, Sakamoto K, Sogabe T (2021) On the expressibility and overfitting of quantum circuit learning. ACM Trans Quantum Comput 2(2):1–24
- Coyle B, Mills D, Danos V, Kashefi E (2020) The born supremacy: quantum advantage and training of an Ising born machine. NPJ Quantum Inf 6(1):60
- Denchev VS, Ding N, Vishwanathan S, Neven H (2012) Robust classification with adiabatic quantum optimization. arXiv:1205.1148
- De Oliveira NM, Lucas P, De Oliveira WR, Ludermir TB, Da Silva AJ (2021) Quantum one-class classification with a distance-based classifier. 2021 international joint conference on neural networks (ijcnn) pp 1–7
- De Palma G, Marvian M, Rouzé C, França DS (2023) Limitations of variational quantum algorithms: a quantum optimal transport approach. PRX Quantum 4(1):010309
- Du Y, Hsieh MH, Liu T, Tao D (2018). Implementable quantum classifier for nonlinear data. arXiv:1809.06056
- Du Y, Hsieh MH, Liu T, You S, Tao D (2021) Learnability of quantum neural networks. PRX Quantum 2(4):040337
- Du Y, Tu Z, Yuan X, Tao D (2022) Efficient measure for the expressivity of variational quantum algorithms. Phys Rev Lett 128(8):080506
- Du Y, Yang Y, Tao D, Hsieh MH (2022b) Problem-dependent power of quantum neural networks on multi-class classification. arXiv:2301.01597
- Duchi D (n.d.) CS229 supplemental lecture notes Hoeffding's inequality. https://cs229.stanford.edu/extra-notes/hoeffding.pdf
- Dunjko V, Briegel HJ (2018) Machine learning & artificial intelligence in the quantum domain: a review of recent progress. Rep Progress Phys 81(7):074001
- Emami M, Sahraee-Ardakan M, Pandit P, Rangan S, Fletcher A (2020) Generalization error of generalized linear models in high dimensions. International conference on machine learning pp 2892–2901
- ERDdS P, R&wi A (1959) On random graphs I. Publ Math Debrecen 6(290–297):18
- Fisher RA (1936) The use of multiple measurements in taxonomic problems. Annals of eugenics 7(2):179–188
- Gan BY, Leykam D, Thanasilp S (2023) A unified framework for traceinduced quantum kernels. arXiv:2311.13552
- Gentinetta G, Thomsen A, Sutter D, Woerner S (2024) The complexity of quantum support vector machines. Quantum 8:1225
- Gibbs J, Holmes Z, Caro MC, Ezzell N, Huang HY, Cincio L, Coles PJ (2024) Dynamical simulation via quantum machine learning with provable generalization. Phys Rev Res 6(1):013241
- Gil-Fuster E, Eisert J, Bravo-Prieto C (2023) Understanding quantum machine learning also requires rethinking generalization. arXiv:2306.13461

- Hakkaku S, Tashima Y, Mitarai K, Mizukami W, Fujii K (2022) Quantifying fermionic nonlinearity of quantum circuits. Phys Rev Res 4(4):043100
- Hamilton KE, Pooser RC (2020) Error-mitigated data-driven circuit learning on noisy quantum hardware. Quantum Mach Intell 2:1– 15
- Haug T, Kim M (2023) Generalization with quantum geometry for learning unitaries. arXiv:2303.13462
- Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2019) Supervised learning with quantumenhanced feature spaces. Nature 567(7747):209–212
- Herrmann N, Arya D, Doherty MW, Mingare A, Pillay JC, Preis F, Prestel S (2023) Quantum utility–definition and assessment of a practical quantum advantage. arXiv:2303.02138
- Hevia JL, Peterssen G, Piattini M (2022) QuantumPath: a quantum software development platform. Softw Pract Exper 52(6):1517– 1530
- Heyraud V, Li Z, Denis Z, Le Boité A, Ciuti C (2022) Noisy quantum kernel machines. Phys Rev A 106(5):052421
- Hoeffding W (1994) Probability inequalities for sums of bounded random variables. In: The collected works of Wassily Hoeffding, pp 409–426,
- Holmes Z, Sharma K, Cerezo M, Coles PJ (2022) Connecting ansatz expressibility to gradient magnitudes and barren plateaus. PRX Quantum 3(1):010313
- Huang HY, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, McClean JR (2021) Power of data in quantum machine learning. Nat Commun 12(1):2631
- Huggins W, Patil P, Mitchell B, Whaley KB, Stoudenmire EM (2019) Towards quantum machine learning with tensor networks. Quantum Sci Technol 4(2):024001
- Jakubovitz D, Giryes R, Rodrigues MR (2019) Generalization error in deep learning. In: Compressed sensing and its applications: Third international matheon conference 2017 pp 153–193
- Jerbi S, Fiderer LJ, Poulsen Nautrup H, Kübler JM, Briegel HJ, Dunjko V (2023) Quantum machine learning beyond kernel methods. Nat Commun 14(1):517
- Khairy S, Shaydulin R, Cincio L, Alexeev Y, Balaprakash P (2020) Learning to optimize variational quantum circuits to solve combinatorial problems. In: Proceedings of the aaai conference on artificial intelligence vol 34, pp 2367–2375
- Khanal B, Rivas P (2023) Evaluating the impact of noise on variational quantum circuits in NISQ era devices. In: Proc. of the international conference on emergent and quantum technologies (iceqt 2023) pp 1–7
- Kim C, Park KD, Rhee JK (2020) Quantum error mitigation with artificial neural network. IEEE Access 8:188853–188860
- Kitchenham BA, Budgen D, Brereton OP (2011) Using mapping studies as the basis for further research-a participant-observer case study. Inf Softw Technol 53(6):638–651
- Kordzanganeh M, Buchberger M, Kyriacou B, Povolotskii M, Fischer W, Kurkin A, Melnikov A (2023) Benchmarking simulated and physical quantum processing units using quantum and hybrid algorithms. Adv Quantum Technol 6(8):2300043
- Kübler J, Buchholz S, Schölkopf B (2021) The inductive bias of quantum kernels. Adv Neural Inf Process Syst 34:12661–12673
- Lavrijsen W, Tudor A, Müller J, Iancu C, De Jong W (2020) Classical optimizers for noisy intermediate-scale quantum devices. In: 2020 IEEE international conference on quantum computing and engineering (qce) pp 267–277
- LeCun Y, Cortes C, Burges C (2010) MNIST handwritten digit database. ATT Labs [Online] vol. 2. http://yann.lecun.com/exdb/mnist
- Li Q, Huang Y, Hou X, Li Y, Wang X, Bayat A (2023) Ensemble-learning variational shallow-circuit quantum classifiers. arXiv:2301.12707

- Liang Z, Wang Z, Yang J, Yang L, Shi Y, Jiang W (2021) Can noise on qubits be learned in quantum neural network? a case study on quantumflow. 2021 IEEE/ACM international conference on computer aided design (iccad) pp 1–7
- Liu Y, Arunachalam S, Temme K (2021) A rigorous and robust quantum speed-up in supervised machine learning. Nat Phys 17(9):1013–1017
- Marrero CO, Kieferová M, Wiebe N (2021) Entanglement-induced barren plateaus. PRX. Quantum 2(4):040316
- Martín-Guerrero JD, Lamata L (2022) Quantum machine learning: a tutorial. Neurocomputing 470:457–461
- McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H (2018) Barren plateaus in quantum neural network training landscapes. Nat Commun 9(1):4812
- Mohri M, Rostamizadeh A, Talwalkar A (2018) Foundations of machine learning. MIT press
- Nadeau C, Bengio Y (1999) Inference for the generalization error. Adv Neural Inf Process Syst 12
- Neelakantan A, Vilnis L, Le QV, Sutskever I, Kaiser L, Kurach K, Martens J (2015) Adding gradient noise improves learning for very deep networks. arXiv:1511.06807
- Park DK, Blank C, Petruccione F (2021) Robust quantum classifier with minimal overhead. In: 2021 international joint conference on neural networks (ijcnn) pp 1–7
- Preskill J (2018) Quantum computing in the NISQ era and beyond. Quantum 2:79
- Qi J, Yang CHH, Chen PY, Hsieh MH (2023) Theoretical error performance analysis for variational quantum circuit based functional regression. NPJ Quantum Inf 9(1):4
- Qian Y, Wang X, Du Y, Wu X, Tao D (2022) The dilemma of quantum neural networks. IEEE Trans Neural Netw Learn Syst
- Rocchetto A, Aaronson S, Severini S, Carvacho G, Poderini D, Agresti I, Sciarrino F (2019) Experimental learning of quantum states. Sci Adv 5(3):eaau1946
- Roget M, Di Molfetta G, Kadri H (2022) Quantum perceptron revisited: computational-statistical tradeoffs. In: Uncertainty in artificial intelligence pp 1697–1706
- Schuld M (2021) Supervised quantum machine learning models are kernel methods. arXiv:2101.11020
- Schuld M, Killoran N (2019) Quantum machine learning in feature Hilbert spaces. Phys Rev Lett 122(4):040504
- Schuld M, Killoran N (2022) Is quantum advantage the right goal for quantum machine learning? Prx Quantum 3(3):030101
- Schuld M, Petruccione F (2021) Machine learning with quantum computers. Springer
- Schuld M, Sweke R, Meyer JJ (2021) Effect of data encoding on the expressive power of variational quantum-machine-learning models. Phys Rev A 103(3):032430
- Serrano MA, Cruz-Lemus JA, Perez-Castillo R, Piattini M (2022) Quantum software components and platforms: overview and quality assessment. ACM Comput Surv 55(8):1–31
- Shaydulin R, Wild SM (2022) Importance of kernel bandwidth in quantum machine learning. Phys Rev A 106(4):042407
- Skolik A, McClean JR, Mohseni M, van der Smagt P, Leib M (2021) Layerwise learning for quantum neural networks. Quantum Mach Intell 3:1–11
- Skorski M, Temperoni A, Theobald M (2021) Revisiting weight initialization of deep neural networks. In: Asian conference on machine learning pp 1192–120
- Srikumar M, Hill CD, Hollenberg LC (2022) A kernel-based quantum random forest for improved classification. arXiv:2210.02355
- Thanasilp S, Wang S, Cerezo M, Holmes Z (2024) Exponential concentration in quantum kernel methods. Nat Commun 15(1):5200
- Wang H, Gu J, Ding Y, Li Z, Chong FT, Pan DZ, Han S (2022) Quantumnat: quantum noise-aware training with noise injec-

tion, quantization and normalization. In: Proceedings of the 59th ACM/IEEE design automation conference pp 1–6

- Wang H, Li Z, Gu J, Ding Y, Pan DZ, Han S (2022) QOC: quantum on-chip training with parameter shift and gradient pruning. In: Proceedings of the 59th ACM/IEEE design automation conference pp 655–660
- Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L, Coles PJ (2021) Noise-induced barren plateaus in variational quantum algorithms. Nat Commun 12(1):6961
- Wang X, Du Y, Luo Y, Tao D (2021) Towards understanding the power of quantum kernels in the NISQ era. Quantum 5:531
- Watabe M, Shiba K, Sogabe M, Sakamoto K, Sogabe T (2019) Quantum circuit parameters learning with gradient descent using backpropagation. arXiv:1910.14266
- Wittek P (2014) Quantum machine learning: what quantum computing means to data mining. Academic Press
- Xiao H, Rasul K, Vollgraf R (2017) Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. arXiv:1708.07747

- Zardini E, Blanzieri E, Pastorello D (2022) Implementation and empirical evaluation of a quantum machine learning pipeline for local classification. arXiv:2205.05333
- Zhao C, Gao XS (2021) Analyzing the barren plateau phenomenon in training quantum neural networks with the ZX-calculus. Quantum 5:466

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.