



Received 22 December 2022; revised 16 June 2023; accepted 16 June 2023; date of publication 20 June 2023; date of current version 15 September 2023.

Digital Object Identifier 10.1109/TQE.2023.3287736

The Quantum Path Kernel: A Generalized Neural Tangent Kernel for Deep Quantum Machine Learning

MASSIMILIANO INCUDINI¹, MICHELE GROSSI²,
ANTONIO MANDARINO³, SOFIA VALLECORSÀ²,
ALESSANDRA DI PIERRO¹, AND DAVID WINDRIDGE⁴

¹Department of Computer Science, University of Verona, 37134 Verona, Italy

²European Organization for Nuclear Research (CERN), 1211 Geneva, Switzerland

³International Centre for Theory of Quantum Technologies (ICTQT), University of Gdansk, 80-309 Gdańsk, Poland

⁴Department of Computer Science, Middlesex University, The Burroughs, NW44BT London, U.K.

Corresponding author: Massimiliano Incudini (e-mail: massimiliano.incudini@univr.it).

The work of A. Mandarino was supported in part by the Foundation for Polish Science (FNP), IRAP project ICTQT, contract 2018/MAB/5, and in part by EU Smart Growth Operational Programme and the University of Verona throughout Mobility under Grant PIA2022_CATB_DIPIERRO.

ABSTRACT Building a quantum analog of classical deep neural networks represents a fundamental challenge in quantum computing. A key issue is how to address the inherent nonlinearity of classical deep learning, a problem in the quantum domain due to the fact that the composition of an arbitrary number of quantum gates, consisting of a series of sequential unitary transformations, is intrinsically linear. This problem has been variously approached in literature, principally via the introduction of measurements between layers of unitary transformations. In this article, we introduce the quantum path kernel (QPK), a formulation of quantum machine learning capable of replicating those aspects of deep machine learning typically associated with superior generalization performance in the classical domain, specifically, *hierarchical feature learning*. Our approach generalizes the notion of quantum neural tangent kernel, which has been used to study the dynamics of classical and quantum machine learning models. The QPK exploits the parameter trajectory, i.e., the curve delineated by model parameters as they evolve during training, enabling the representation of differential layerwise convergence behaviors, or the formation of hierarchical parametric dependencies, in terms of their manifestation in the gradient space of the predictor function. We evaluate our approach with respect to variants of the classification of Gaussian XOR mixtures: an artificial but emblematic problem that intrinsically requires multilevel learning in order to achieve optimal class separation.

INDEX TERMS Machine learning, neural tangent kernel (NTK), quantum kernel, quantum machine learning, quantum neural networks (QNNs), support vector machine (SVM).

I. INTRODUCTION

Bridging classical deep neural networks and quantum computing represents a key research challenge in the field of *quantum machine learning* [1], [2]. The potential for improvement offered by quantum computing in the machine learning domain may be characterized in terms of its impact on algorithmic efficiency, generalization error, or else its capacity for treating quantum data [3].

A notable recent result in the field has been the introduction of the concept of *variational quantum algorithms* and the related neural network analog referred to as the quantum neural network (QNN) [4]. This, in essence, consists of a

feature map encoding data into a quantum Hilbert space upon which certain parameterized unitary rotations are applied prior to final measurement in order to obtain a classification or regression output. The system as a whole is then optimized by classical methods. Such models provably lead to a computational advantage over classical models on certain artificial tasks [5], and in respect to the analysis of specific physical systems [6]. It has been quantitatively shown that QNNs can be trained faster than their classical analogues [4]. However, QNNs remain problematic in various respects. One limitation arises from the so-called barren plateau problem [7], in which the variance of the gradient vanishes exponentially

with the system size as the parameterized transformation becomes increasingly expressive [8]. A number of approaches, including layerwise training of QNNs [9], have been proposed to mitigate the issue.

A second problematic aspect of QNNs, and the one that constitutes our principal focus here, is the linearity of the dynamics of quantum systems. Concatenations of linear unitary transformations remain unitary, and thus, “stacked” quantum transformations, in effect, collapse to a single linear transformation, appearing to rule out de facto the hierarchical feature learning of classical deep neural networks, which relies on nonlinearities to separate feature layers. This property makes the QNN essentially a kernel machine [10]. In terms of the predictor function, however, the QNN is composed of multiplications of rotation operators parameterized by both the feature and model weights. The nonlinearity of projections of rotation operators can be exploited to replicate a very constrained form of nonlinearity for feature learning [11], [12]. Another strategy is to introduce nonlinearity via the measurement operation, i.e., a *dissipative QNN* [13]. Both approaches involve the projection of the quantum state into a subspace of the original Hilbert space.

Much of the recent study of the dynamics of deep neural networks in the classical realm has focused on the neural tangent kernel (NTK) [14] which represents the network in terms of the corresponding training gradients in the model parameter space. The NTK hence approximates the behavior of predictors via a linear model. It is often therefore applied to study neural networks in their asymptotic, infinite-width, limit. In this regime, the network exhibits *lazy training* [15], i.e., parameter gradients remain at their initial values during the entirety of the training. The NTK thus accurately characterizes the dynamics of such infinite-width neural networks, but is otherwise only an approximation [16]. The difference in test error between the predictor and its linearized version depends on the problem structure [17], with hierarchical feature learning capability being crucial to obtaining superior performance [18]. However, the kernel nature of the NTK means that it shares with quantum computing a ready interpretation within a Hilbert space, and is thus of considerable interest within quantum machine learning. The first explicit application of NTK to QNNs, the quantum neural tangent kernel (QNTK) was given in [19].

In this article, we propose a method for overcoming the de facto lack of hierarchical feature learning capability in QNNs. We propose the application of path kernels [20] to QNNs, which we call the quantum path kernel (QPK). Such an approach generalizes the QNTK so that the resulting kernel is representative of the ensemble of NTKs calculated over the full parameter path trajectory, i.e., the function describing the evolution of model parameters over time, including implicitly any parametric evolutions corresponding to hierarchical feature learning. We show experimentally an increased expressivity of the resulting model relative to linearized equivalents, evaluating our method on the Gaussian XOR mixture classification problem. For this problem,

finite-width neural networks have both theoretically and empirically shown to be close-to-optimal performance, whereas linear NTK models fail [21], suggesting that it cannot be effectively resolved without implicating multi-level learning behavior. Furthermore, we discuss possible improvements for the proposed approach, which can be obtained by considering only the contribution of the parameter gradient path that gives rise to the most decorrelated feature representation. These specifically corresponds to the contributions associated with the maximally nonlinear point of the parameter path, corresponding to the largest (positive or negative) eigenvalues of the Hessian of the predictor function [22]. We further enhance the decorrelation between feature representations via a stochastic, noisy, or nongradient-descent-based training algorithm in which the averaging operation between decorrelated representations allows us to interpret the model as an ensemble technique.

The rest of this article is organized as follows. In Section II, we briefly review the necessary conceptual background. In Section III, we present the QPK and discuss the hierarchical feature learning of the induced model. In Section IV, we demonstrate how this leads to superior performance in solving the Gaussian XOR mixture classification problem. Finally, Section V concludes this article.

A. CONTRIBUTIONS

- 1) We propose the QPK as a mechanism for building hybrid classical/quantum machine learning models which are able to emulate the hierarchical feature learning structure of deep neural networks without violating the underlying linearity of the quantum dynamics. We conjecture such a kernel can lead to improved performance of the kernel machine on tasks involving *feature learning*.
- 2) We provide numerical evidence of the superior performance of the QPK compared to the QNTK on the Gaussian XOR mixture problem, which is Bayes optimally soluble only through implicating layerwise nonlinear separability.
- 3) We consider the importance of the extraction of non-correlated feature representations corresponding to maximally varying portions of the parameter gradient path.

B. RELATED WORKS

The introduction of the NTK by [14] has marked a significant step in the theory of machine learning, shedding new light on discussions regarding the relative performance of linear and nonlinear models. For example, Ghorbani et al. [17] suggested that tasks in which kernel methods (including NTK) perform worse than neural networks are those in which the kernel suffers from the curse of dimensionality, whereas neural networks, in learning some useful lower dimensional representation, do not. One example of such a problem is the Gaussian XOR mixture classification task [21]. Furthermore, linearized models have been shown to perform slightly

worse than wide (i.e., large, but noninfinite) neural networks on a CIFAR-10 benchmark [23], with the gap between the approaches increasing for finite width networks [24].

In relation to quantum computation, researchers have spent substantial effort on the limitations imposed by the linear dynamics of quantum systems. Authors in [25] review early approaches to the formulation of nonlinear quantum machine learning models: some have focused on developing a *quantum perceptron* equivalent or *quantum neuron*, i.e., a candidate building block for the quantum analog of neural networks; Schuld et al. [26] used phase estimation to implement the functioning of a step function; the authors in [27] and [28] proposed to exploit the repeat until success (RUS) policy to mimic the behavior of tangent and sigmoid activation functions, while Gili et al. [29] used RUS to construct a Born machine; Tacchino et al. [30] emulated the nonlinearity of perceptrons using measurements. In relation to QNNs, Sharma et al. [13] proposed dissipative QNNs in which the nonlinearity is obtained via intertwining measurements between unitary gates; the authors in [31] and [32] proposed the use of a larger Hilbert space to implement the nonlinear transformation, while Daksin [33] exploited the exponential form of unitary gate to achieve periodic activation functions. Finally, nonlinear models of quantum mechanics have been conjectured by [34], although these violate some computational complexity assumptions [35].

Regarding QNTK literature, authors in [36] have quantified the benefic effect of introducing handcrafted symmetries in the QNN using tools derived from the QNTK. Authors in [37] and [38] investigated lazy training in the quantum machine learning context.

II. BACKGROUND

This section briefly introduces the key concepts and notations in relation to deep learning and quantum machine learning through which we develop our results. We denote by $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \subseteq \mathcal{X} \times \mathcal{Y}$ a labeled dataset of pairs that are i.i.d. sampled from an unknown probability distribution. We indicate the data vector space with $\mathcal{X} = \mathbb{R}^d$, and the target space with either $\mathcal{Y} = \mathbb{R}$ or $\mathcal{Y} \subseteq \mathbb{Z}$, $|\mathcal{Y}| < \infty$ for regression or classification tasks, respectively. We indicate uniform sampling from a uniform discrete distribution with $\sim \{v_i\}_{i=1}^n$ and sampling from a normal distribution of mean μ and variance σ^2 with $\sim \mathcal{N}(\mu, \sigma)$.

A. PRIMER ON QUANTUM MACHINE LEARNING MODELS

Here, we fix the notation for our quantum machine learning models. The state of a quantum system of m -qubits is described by a density matrix $\rho \in \mathcal{H} \equiv \mathbb{C}^{2^m \times 2^m}$. The initial state of a quantum computation is denoted by $\rho_0 = |0\rangle\langle 0|$, and the (possibly parametric) unitary transformations by U, V , and W . Any parametric unitary can be written as

$$U(\boldsymbol{\theta}) = \exp \left\{ -i \sum_{k=1}^m f_j(\boldsymbol{\theta}) \sigma_{\alpha_1, \dots, \alpha_k}^{(q_1, \dots, q_k)} \right\} \quad (1)$$

where $\alpha_i \in \{X, Y, Z, \mathbf{0}\}$ for $i = 1, \dots, k$, and $\sigma_{\alpha_1, \dots, \alpha_k}$ is a tensor product of one or more corresponding Pauli matrixes applied to qubits q_1, \dots, q_k . The same transformation may be interpreted as a rotation and be equivalently denoted by $R_{\alpha_1, \dots, \alpha_k}^{(i_1, \dots, i_k)}(\boldsymbol{\theta})$, where $\boldsymbol{\theta} \in \mathbb{R}^P$ are rotational angles. A QNN is a function of the form¹

$$f(\mathbf{x}; \boldsymbol{\theta}) = \text{Tr}[\rho_{\mathbf{x}, \boldsymbol{\theta}} O] = \text{Tr}[V^\dagger(\boldsymbol{\theta}) U_\phi^\dagger(\mathbf{x}) \rho_0 U_\phi(\mathbf{x}) V(\boldsymbol{\theta}) O] \quad (2)$$

where O indicates any measurement operator. Both the matrixes U and V are decomposed in single and two-qubits parametric rotations interspersed with nonparametric gates (e.g., CNOT).

B. KERNEL METHODS AND KERNEL MACHINES

A kernel on \mathcal{X} is a binary, symmetric, and positive definite function $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. The kernel function generalizes the notion of an inner product by mapping elements of \mathcal{X} into a Hilbert space \mathcal{H} , which typically has a richer, more task-specific, structure compared to the original space equipped with linear inner products [40]. The mapping is achieved via an *implicit feature map* $\phi : \mathcal{X} \rightarrow \mathcal{H}$. This formalism enables the exploitation of potentially infinite-dimensional Hilbert spaces (such as those implied by the Gaussian feature map or radial basis function, in which \mathbf{x} is mapped to a multivariate Gaussian of mean \mathbf{x} and fixed covariance), residing in the space of square-integrable multivariate functions.

From the perspective of feature transformation, the kernel function κ can be expressed as follows:

$$\kappa(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle_{\mathcal{H}}. \quad (3)$$

In the majority of tasks, it is not necessary to have knowledge of or utilize the explicit form of the feature map ϕ ; it is sufficient that function κ obeys the Mercer's condition.

A *kernel machine* is a function that can be expressed as a linear combination of kernel evaluations over the training objects $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$

$$f(\cdot) = \sum_{i=1}^n \alpha_i \kappa(\mathbf{x}_i, \cdot) = \sum_{i=1}^n \alpha_i \phi_i(\cdot) \quad (4)$$

where $\alpha_i \in \mathbb{R}$ (the first equality being a statement of the celebrated *representer theorem*). Different kernels capture different aspects of the data, resulting in the corresponding kernel machines having differing characteristics (and ultimately differing performances in the context of a classification problem).

The associated learning problem is then one of finding the kernel machine $f = \arg \min R[f]$ that minimizes the empirical risk function

$$R[f] = \sum_{i=1}^n \ell(f(\mathbf{x}_i), y_i) + \lambda \|f\|^2 \quad (5)$$

¹The most general form of QNN proposed is the *data reuploading* QNN, which allows the interspersing of data encoding and trainable transformations. Such a form, however, does not add any computational power to the standard QNN approach [39].

with ℓ being a convex loss function, and $\lambda \geq 0$ regularization term. The learning problem is thus a convex, n -dimensional optimization problem independent of the dimensionality of the original space in which the \mathbf{x}_i were embedded prior to feature transformation.

C. NOTIONS OF NONLINEARITY IN CLASSICAL AND QUANTUM LEARNING MODELS

With respect to both kernel machines and layerwise deep learning, the concepts of *linear model*, *nonlinear model*, and *feature learning* that we utilize here are as formalized in [41].

A *linear model* is thus a function of the form

$$f(\mathbf{x}; \boldsymbol{\theta}) = \sum_{j=1}^p \theta_j \phi_j(\mathbf{x}) \quad (6)$$

where $\{\phi_j : \mathcal{X} \rightarrow \mathbb{R}\}_{j=0}^p$ are the *feature functions*, whose values correspond with the model features. We might consider an additional feature $\phi_0 \equiv 1$ that incorporates the bias. The formula in (6) is linear with respect to the space of the parameters \mathbb{R}^p ; in fact, we can interpret the function as an inner product in that space, i.e.,

$$f(\mathbf{x}; \boldsymbol{\theta}) = \langle \boldsymbol{\theta}, \boldsymbol{\phi}(\mathbf{x}) \rangle_{\mathbb{R}^p} \quad (7)$$

with $\boldsymbol{\theta} = (\theta_1, \dots, \theta_p)$ and $\boldsymbol{\phi}(\mathbf{x}) = (\phi_1(\mathbf{x}), \dots, \phi_p(\mathbf{x}))$, while it behaves nonlinearly with respect to the original feature space \mathcal{X} , due to the feature functions. In this sense, kernel machines are linear models.

A *nonlinear model* is a function of the form

$$\begin{aligned} f(\mathbf{x}; \boldsymbol{\theta}) = & \sum_{j=1}^p \theta_j \phi_j(\mathbf{x}) + \frac{\epsilon}{2} \sum_{j,k=1}^p \theta_j \theta_k \psi_{j,k}(\mathbf{x}) \\ & + \frac{\epsilon}{3!} \sum_{j,k,\ell=1}^p \theta_j \theta_k \theta_\ell \psi_{j,k,\ell}(\mathbf{x}) + \dots \end{aligned} \quad (8)$$

The higher order terms of the expansion are characterized by their own set of features, e.g., $\{\psi_{j,k} : \mathcal{X} \rightarrow \mathbb{R}\}_{j,k=1}^p$ for the second-order term. The elements of such sets are unique up to a permutation of their variables, thus the terms $1/2!, 1/3!, \dots$ compensate the multiple counting of such elements in (8). The term $\epsilon \ll 1$ adjusts the contribution of the nonlinear terms. If the model is truncated to the second term it is denoted a *quadratic model* [41], [42]. In such a case, the loss function is quartic, and thus, we cannot analytically determine the optimal parameters as for linear regression. The parameter dynamics of such a model is given by

$$\begin{aligned} f(\mathbf{x}, \boldsymbol{\theta} + d\boldsymbol{\theta}) = & f(\mathbf{x}, \boldsymbol{\theta}) \\ & + \sum_{j=1}^p d\theta_j \left[\phi_j(\mathbf{x}) + \epsilon \sum_{k=1}^p \theta_j \psi_{j,k}(\mathbf{x}) \right] \end{aligned} \quad (9)$$

$$+ \frac{\epsilon}{2!} \sum_{j,k=1}^p d\theta_j d\theta_k \psi_{j,k}(\mathbf{x}) \quad (10)$$

$$\begin{aligned} = & f(\mathbf{x}, \boldsymbol{\theta}) + \sum_{j=1}^p d\theta_j \phi_j^E(\mathbf{x}; \boldsymbol{\theta}) \\ & + \frac{\epsilon}{2} \sum_{j,k=1}^p d\theta_j d\theta_k \psi_{j,k}(\mathbf{x}) \end{aligned} \quad (11)$$

where ϕ^E are *effective feature functions*, i.e., features that depend on, and evolve with, the model parameters, which are learnt during the optimization phase. This behavior can be generalized to consider terms of even higher orders: the presence of order n terms make the feature functions of order $n - 1$, effective, which may further influence the lower order terms. Models having effective feature functions have *feature learning* capabilities. A *deep learning model* is both capable of feature learning and composed of several nonlinear modules arranged in a hierarchical fashion [43]; such that differing layers can follow differing (albeit hierarchically conditioned) gradient paths.

Turning to QNNs, the quantum model

$$f(\mathbf{x}; \boldsymbol{\theta}) = \text{Tr}[\rho_{\mathbf{x},\boldsymbol{\theta}} O] \quad (12)$$

with $\rho_{\mathbf{x},\boldsymbol{\theta}} = V^\dagger(\boldsymbol{\theta}) U_\phi^\dagger(\mathbf{x}) |0\rangle\langle 0| U_\phi(\mathbf{x}) V(\boldsymbol{\theta})$ and O a Hermitian observable, is a linear model in the space of density matrixes of the quantum system \mathcal{H} : the trace operation $\text{Tr}[A^\dagger B]$ is an inner product for the space of matrixes $\mathbb{C}^{k \times k}$. Such a property implies that the construction of a layerwise architecture for v , i.e., $v(\boldsymbol{\theta}) = \prod_i V_i(\boldsymbol{\theta})$ effectively collapses to a single operation: this may add more degrees of freedom to the linear transformation² but cannot make the model nonlinear in \mathcal{H} in the manner of a classical deep learning model.

However, in terms of the predictor function $f(\mathbf{x}; \boldsymbol{\theta})$, the quantum model does not necessarily fit the form set out (6) since the parameters of the QNN model, in particular the angle of rotation operation (in the form of imaginary exponential function), are subject to the trace operation. Thus, for example, consider a single-qubit quantum model acting on a single input $\mathbf{x} \in \mathbb{R}^1$, depending on a single parameter $\boldsymbol{\theta} \in \mathbb{R}^1$, with feature map $U_\phi(\mathbf{x}) = \exp(-i x \sigma_x)$, variational form $V(\boldsymbol{\theta}) = \exp(-i \theta \sigma_x)$ and measurement operator $O = \sigma_z$, in which case $f(\mathbf{x}; \boldsymbol{\theta})$ has the form

$$\begin{aligned} f(\mathbf{x}; \boldsymbol{\theta}) = & \text{Tr} \\ & \left[\begin{pmatrix} \cos^2(\theta + x) & -i \sin(\theta + x) \cos(\theta + x) \\ \frac{1}{2} i \sin(2(\theta + x)) & \sin^2(\theta + x) \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \right] \\ = & \cos(2(\theta + x)) \end{aligned} \quad (13)$$

which is nonlinear in its weights. Clearly, if we were to consider a model other than a QNN then the predictor function would change, for example, as in [30]; however, it does not alter our argument here.

To recap, a QNN is a linear model in the Hilbert space of the density matrixes due to the linearity of the evolution of

²Depending on the generators involved and up to a maximum of $4^n - 1$ (where n is the number of qubits).

closed quantum systems. However, its predictor is nonlinear in the parameter vector θ since its structure results in a composition of trigonometric functions. This potentially allows a limited degree of representational learning capability if aggregated layerwise (limited in the sense of applying only to a highly constrained set of activation functions). However, due to the Lie algebraic equivalence of any given sequence of quantum transformations to some single unitary operation in the absence of the trace operation, we are still not able to characterize truly deep models in the quantum domain.

D. CHARACTERIZATION OF MODEL DYNAMICS THROUGH THE NEURAL TANGENT KERNEL

The output $f(\mathbf{x}; \theta)$ of a machine learning model trained via (possibly stochastic) gradient descent can be approximated as a first-order Taylor expansion $f(\mathbf{x}; \theta) \approx f(\mathbf{x}; \theta_0) + \nabla_{\theta} f(\mathbf{x}; \theta_0)(\theta - \theta_0)$. Such an approximation allows the representation of machine learners as linear (kernel) models via the NTK ([14])

$$\kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}'; \theta) = \nabla_{\theta} f(\mathbf{x}; \theta) \cdot \nabla_{\theta} f(\mathbf{x}'; \theta). \quad (14)$$

This is widely used to characterize the dynamics of infinite-width neural networks [15], in which the NTK is independent of the random initialization and constant in time. On a coarse level of detail, we can assert that model training that takes place in the *lazy-training regime*, i.e., when the evolution of $\theta(t)$ during the training of the model $f(\mathbf{x}, \theta)$ closely follows the tangent path, can be decently approximated by the NTK. A more detailed analysis in [16] has revealed that the NTK is constant if and only if the model is linear (in its parameters). Such a result allows us to quantify the nonlinearity of a model through its Hessian norm of the predictor function: if $\|H_f\| \ll \|\nabla_w f\|$ then the model is nearly linear. This has been used in [11] to analyze the behavior of the QNNs in the lazy training regime.

III. QUANTUM PATH KERNEL FRAMEWORK

Extant quantum methods hence do not fully capture the deviations from gradient path linearity exhibited by empirically optimal finite-width learners in the classical domain during hierarchical feature learning (where, of necessity, lower level features reach learning convergence prior to the higher level features that utilize them). Thus, in order to encompass this learning behavior within the (implicitly kernel-based) quantum machine learning regime, we introduce into the quantum domain a key idea of Domingo [20], namely, *path kernelization*

$$\kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}'; \gamma) = \int_{\gamma} \nabla_{\theta} f(\mathbf{x}; \theta) \cdot \nabla_{\theta} f(\mathbf{x}'; \theta) d\theta \quad (15)$$

where k_{pk} is the *path kernel*, i.e., the line integral of κ_{ntk} over the multidimensional curve representing the evolution of the parameters $\theta = \gamma(t)$, $t \in [0, 1]$ during training, with $\bar{\theta} = \bar{\gamma}(1)$. Appendix A1 gives a proof that path kernel is effectively Mercer, and sets out the pseudocode for its construction.

In general, chain rule dependencies arising from the specifics of the architecture of the classical deep networks will imply that hierarchical dependencies develop among the parameters during learning. The result holds even for stochastic gradient descent optimization, in which case (15) is a stochastic integral. It is thus key to our argument to consider the parameter path γ and its morphological evolution. For linear models, assuming a vanilla gradient descent training over a convex loss function \mathcal{L} , the parameter path is described by a linear vector $\{(1-t)\theta_0 + t\theta_f \mid t \in [0, 1]\}$, where $\theta_0 \in \mathbb{R}^p$ are the parameters at their initialization, and $\theta_f \in \mathbb{R}^p$ are the parameters at their convergence on the (ideally global) minima of \mathcal{L} . In this case, it is immediately possible to check that the derivative of the linear model $\nabla_{\theta} f$ is independent of θ , and thus, that the NTK is constant. For nonlinear models, the loss function \mathcal{L} may become nonconvex and γ is not constrained to have a linear trajectory. In this latter case, both the $\nabla_{\theta} f$ and the NTK will vary over time during learning.

The path kernel was originally introduced in [20] in terms of a direct functional equivalence between nonlinear models trained with gradient descent and path kernel machines (cf Appendix A2). We do not, however, here, focus on the potential of path kernels in approximating nonlinear models. Instead, we seek to exploit the intrinsic potential of path kernels to capture implicit feature hierarchy formation in a manner capable of implementation on a hybrid deep machine learning model within a QNN setting. We thus seek to establish whether incorporation of a quantum version of path construction into kernel machines, such as the support vector machine (SVM), can yield performance enhancements in a feature learning setting: Section IV will seek to quantify this directly.

We depict the construction of the path kernel in Fig. 1. The parameter trajectory for a nonlinear model is described by a complex, nonstraight curve. Each point of the parameter path $\theta_t = \gamma(t)$ may be used to define a new kernel representation for the training data, namely, $\kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}'; \theta_t)$. We can then define a sequence of kernels stacked in a hierarchical way (whose structure, in passing, resembles the layers of a deep neural network, though this observation is peripheral to the argument being made here). Thus, each new “layer” is a source of representation learning: the new representation (i.e., kernel matrix) is the result of an optimization process that further adapts the previous representation to the given data discrimination problem (which resembles, though is again not equivalent to, classifier boosting).

It thus becomes possible, via explicit substitution for the corresponding QNTK previously defined, to construct a QPK as follows:

$$\begin{aligned} \kappa_{\text{qpk}}(\mathbf{x}, \mathbf{x}'; \gamma) = & \int_{\gamma} \nabla_{\theta} \langle 0 | V^{\dagger}(\mathbf{x}) U^{\dagger}(\theta) O U(\theta) V(\mathbf{x}) | 0 \rangle^T \\ & \cdot \nabla_{\theta} \langle 0 | V^{\dagger}(\mathbf{x}') U^{\dagger}(\theta) O U(\theta) V(\mathbf{x}') | 0 \rangle d\theta \end{aligned} \quad (16)$$




FIGURE 1. Computation of the path kernel. (Bottom left) A typical parameter trajectory γ is depicted, representing parametric evolution during the training phase. (Top left) As θ evolves, it gives rise to differing NTK matrixes, corresponding to distinct representations of the data. Such a sequence of matrixes thus give rise to a hierarchical stack of representations in the feature learning regime. (Middle) As the training approaches convergence, subsequent matrixes become similar to each other, and thus their corresponding representations are correlated. (Right) The path kernel constitutes the average over these representations.

$$\begin{aligned} &\approx \frac{1}{T} \sum_{i=0}^{T-1} \nabla_{\theta} \langle 0|V^{\dagger}(\mathbf{x})U^{\dagger}(\theta_i)OU(\theta_i)V(\mathbf{x})|0\rangle^T \\ &\quad \cdot \nabla_{\theta} \langle 0|V^{\dagger}(\mathbf{x}')U^{\dagger}(\theta_i)OU(\theta_i)V(\mathbf{x}')|0\rangle \end{aligned} \quad (17)$$

where T is the number of points sampled from the curve γ , and $t = i/T \in [0, 1]$. Equation (16) defines the QNTK as its classical analog. Equation (17) is the discretized version of the preceding equations, corresponding to actual implementation in a gradient descent-trained model. The equation representing the corresponding kernel machine is given by

$$f(\cdot) = \sum_{i=1}^n \alpha_i \kappa_{\text{qpk}}(\mathbf{x}_i, \cdot) \quad (18)$$

where the coefficients α_i are determined during the training process (such as is the case for SVMs and kernel ridge regressors), and \mathbf{x}_i are the elements of the training set. Importantly, these coefficients do not depend on the value of \mathbf{x} .

The resulting QPK is consequently both a quantized version of Domingo’s path kernel as well as a generalization of the QNTK, one that is implicitly capable of embodying the complex parametric interactions (such as transient parametric coevolutions) that occur during learning in order to arrive at the final trained model, including those implicated in hierarchal feature learning.

A. QUANTUM PATH KERNEL AS A GENERALIZATION OF QUANTUM NEURAL TANGENT KERNEL

The primary motivation for introducing the QPK is to enhance the performance of kernel machines in the quantum domain. As such, it represents a departure from the typical focus of the QNTK theory, which primarily aims to predict specific properties of the underlying quantum model. Nevertheless, it is possible to interpret the QPK as a generalization of the QNTK with corresponding insights into the dynamics of the training process. In particular, it may be seen that the

QNTK is constant only when it is fully independent of θ , in which case

$$\begin{aligned} \kappa_{\text{qpk}}(\mathbf{x}, \mathbf{x}'; \gamma) &= \int_{\gamma} \kappa_{\text{qntk}}(\mathbf{x}, \mathbf{x}'; \theta) d\theta \\ &= \kappa_{\text{qntk}}(\mathbf{x}, \mathbf{x}'; \mathbf{0}) \int_{\gamma} d\theta = \kappa_{\text{qntk}}(\mathbf{x}, \mathbf{x}'; \mathbf{0}). \end{aligned} \quad (19)$$

That is, the QPK becomes identical to the QNTK. However, as set out in Section II-C, the particular structure of QNNs will, of itself, give rise to a nonlinear predictor. Thus, in principle, the QNTK would not be expected to be constant in output terms in the finite width regime [11]. However, a close-to-constant behavior can be expected for quantum machine learning models whose training is lazy (i.e., lazy training induced via overparameterization of the QNN, such that a large number of parameters result in a simplified loss landscape [44], [45], leading to rapid convergence to a global minima).

B. DECORRELATION IN FEATURE REPRESENTATION

The QPK clearly exhibits dependency on the training initialization: different initial parameter values, optimization algorithms, or learning rates may lead to differing QPK kernel matrixes. In particular, the utilization of “vanilla” gradient-descent optimization algorithms, with a fixed number of training epochs, may introduce subtle biases in the QPK. For example, if training were to converge rapidly, any contribution between the instance of convergence and the end of the training will be effectively identical and oversampled: this contribution will hence outweigh the others, biasing the “stack” of aggregated kernel matrixes toward its final layer, as per Fig. 1.

To avoid this, more sophisticated optimization algorithms can be considered. For example, the ADAM optimizer adaptively increases the learning rate in locally convex portions




FIGURE 2. Gaussian XOR Mixture classification experiment workflow. (a) Generate dataset. (b) Train QNN. (c) Compare NTK and PK.

of the loss landscape, leading to fewer similar contributions within the path kernel. Furthermore, it is possible to perturb parameter paths via stochastic, noisy, or nongradient-descent-based optimization techniques in order to decorrelate subsequent contributions to the QPK. Having different, highly decorrelated contributions would allow us to interpret the QPK as an ensemble technique analogous to bootstrap aggregation (bagging) often used for tuning the bias/variance trade off in classical machine learning. (Multiple Kernel Learning [46] might also be used to optimally weigh individual contributions over the kernel at the expense of interpretability in path terms.)

Appendix A3 discusses implementation details for the QPK and its tested variants. We, therefore, now turn to an examination of the test regime.

IV. EXPERIMENTAL EVALUATION OF THE QUANTUM PATH KERNEL IN CLASSIFYING GAUSSIAN XOR MIXTURES

Machine-learning nonlinearities, such as those underpinning feature learning in empirical DNNs can thus be feasibly implemented in a quantum setting via the QPK. It remains to demonstrate that this can yield superior generalization performance on plausible quantum devices. Our evaluation, therefore, considers the reference case of the Gaussian XOR mixture classification problem [47], [48], [49].

In particular, the Gaussian XOR mixture classification problem is an important benchmark for highlighting layer-wise learning capabilities of a model (or the lack of them), in that it intrinsically requires a two-layer solution in order to achieve Bayes optimal class separation. Theoretical evidence has shown that kernel methods, in particular, those with random features, struggle to accurately classify XOR data vector mixtures [21]. In Appendix B, we further analyze the problem, reproducing the results of [21], and proposing an interpretation of the success of feature learning models in tackling the Gaussian XOR mixture problem.

Our experimental workflow is pictured in Fig. 2. First, we generate the dataset for the above described problem. Second, we train several QNNs to best fit the generated data. Third, we use the training information to create the QNTK and QPK matrixes; the latter are used to train

a kernel machine (specifically the SVM) to obtain final classifications. Then, our analysis begins with convergence study of the QNNs with an increasing number of layers, to highlight the effect of architectural parametrization in QNNs. Finally, we compare the performances of the QNTK and QPK approaches in terms of testing and training accuracy. The simulation details are shown in Appendix C.

A. EXPERIMENTAL SETUP

The ground truth Gaussian XOR mixture dataset is specified by d the dimensionality of the features, $d' \leq d$ the number of nonzero features representing the multidimensional Gaussian XOR mixture, $\bar{\epsilon}$ the variance of the Gaussian noise, and n the number of data points; it is composed as follows:

$$\mathcal{D}_{d,d',\bar{\epsilon},n} = \left\{ \left([x_1 + \epsilon_1, \dots, x_{d'} + \epsilon_{d'}, 0, \dots, 0]^T, y_i \right) \in \mathbb{R}^d \times \{\pm 1\} \right\}_{j=1}^n \quad (20)$$

where $x_i \sim \{\pm 1\}$, $\epsilon_i \sim \mathcal{N}(0, \bar{\epsilon})$ for $i = 1, \dots, d'$, and $y_i = \prod_{i=1}^{d'} x_i$. Such a dataset is optimally classified via the oracle function

$$f_{\text{oracle}}(\mathbf{x}) = \prod_{i=1}^{d'} x_i. \quad (21)$$

We generate multiple datasets $\mathcal{D}_{d,d',\epsilon,n}$ having feature dimensionality ranging in $d = 2, 3, \dots, 10$, noise ranging in $\epsilon = 0.1, 0.2, \dots, 1.0$, number of nonzero features fixed to $d' = 2$, and number of elements fixed at $n = 32$. Then, each dataset has been randomly partitioned into a training set $\mathcal{D}_{\text{train}}$ and a testing set $\mathcal{D}_{\text{test}}$.

Each dataset is processed by a distinct QNN, each sharing the same structure described by

$$\begin{aligned} f(\mathbf{x}; \theta) &= \text{Tr}[\rho_{\mathbf{x},\theta} O] \\ &= \text{Tr}[V^\dagger(\theta) U_\phi^\dagger(\mathbf{x}) \rho_0 U_\phi(\mathbf{x}) V(\theta) O] \end{aligned} \quad (22)$$

with data encoding

$$U_\phi(\mathbf{x}) = \prod_{j=1}^d \exp \left\{ -i x_j \sigma_y^{(j)} \right\} \quad (23)$$




FIGURE 3. Quantum circuit schematic of the classification model used for $d = 3$ qubits and $L = 2$ layers.

such that the trainable ansatz is described as

$$V(\theta) = \prod_{j=1}^L \exp \left\{ -i \theta_{2i+1} \sigma_x^{(j)} \right\} \times \exp \left\{ -i \theta_{2i} \sigma_z^{(j)} \otimes \sigma_z^{(j+1 \bmod d)} \right\} \quad (24)$$

with the L hyperparameter representing the number of layers of the model. Finally, the observable is $O = \sigma_z^{(0)}$.

This data encoding is chosen for its simplicity: the encoding of one feature for each qubit results in a constant-depth circuit. The choice of the trainable ansatz, though, is particularly important: the underlying functional transformation has the potential to be affected by barren plateau issues if it is too expressive [8], for example when the parametric transformation is able to approximate any arbitrary unitary matrix. The expressibility of a quantum transformation can be examined using Lie-algebraic tools, as shown in [50]. Among the class of unitaries that are nonmaximally expressive, we have selected a specific form that has empirically demonstrated favorable trainability, as detailed in [45, Fig. 7a]. The choice of the observable is also guided by the necessity of avoiding the barren plateau issue. According to [51], global observables are likely to exhibit vanishing gradients; we thus apply the simplest possible classifier observable acting on a single qubit. The circuit is pictured in Fig. 3. In our experiment, the observed qubit is the uppermost; although any other qubit choice would result in a similar predictor due to the symmetric structure of the circuit.

Each dataset is processed with the above described QNN employing a number of layers ranging from $L = 1$ to 20. According to [52], the QNNs should be initialized at $\theta = \mathbf{0}$ to avoid further trainability issues. However, we do not need to consider such initialization strategy for the variational unitary since the previous expedients were sufficient to allow successful training. Thus, the parameters θ_j are sampled from a standard normal distribution. Each QNN is trained using the stochastic gradient-descent algorithm ADAM for 1000 epochs using an initial (adaptive) learning rate $\eta = 0.1$. The loss function is either binary cross entropy or mse and, for the sake of simplification, the batch size is equal to the total cardinality of the training set. These loss functions are commonly used for classification and regression tasks, which differs from the loss functions used in the works related to the QNTK theory.

In the experimental setup described above, we study, both epochwise and depthwise, the effect induced by different initialization parameters on the convergence of the loss function during training.

B. RESULTS

We evaluate the depthwise convergence characteristics of the respective $f(\mathbf{x}; \theta)$ models in terms of the corresponding accuracies of the QPK and QNTK under SVM final classification. Of particular interest is evaluating the closeness of models to the *lazy training* regime, indicative of the model being near to linear. Lazy training, in classical machine learning, typically occurs for very wide neural networks with the loss decreasing to zero exponentially rapidly, while network parameters stay close to their initialization values throughout training. In the current context, this would correspond to the QPK collapsing to the QNTK, and we would anticipate convergent classification performances for the two approaches.

We, therefore, evaluate training loss for each of the QNN models over the respective training epochs with an increasing number of QNN layers $L = 1, \dots, 20$. This will be used to determine proximity to the lazy training regime (i.e., identifying if the QNN converges exponentially fast to zero loss). We additionally plot the norm difference between the parameters during training compared to their initialization values. These will be used to determine the extend to which parameters vary from their initialization, indicating the training richness of models in the *nonlazy* training regime.

We are also interested in determining the robustness of the classifiers to stochastic noise influences during training and their corresponding resilience to overfitting (or the extent to which *benign overparameterization* [44] effects exists), measured in terms of generalization performance. Therefore, the above evaluations are repeated for datasets additively noise-perturbed in an increasing signal-to-noise ratio.

Finally, we are interested in comparing the generalization performances of our approach to that of the QNTK. For this, we evaluate test accuracy score for the QPK and QNTK, against the oracle. Superior performance of the QPK, in solving the Gaussian XOR mixture problem, will be taken to be indicative of superior ability to replicate the layerwise feature-learning capability of classical multilayer networks.

1) DEPTHWISE CONVERGENCE CHARACTERISTICS

Fig. 4 indicates the respective convergence behavior of the evaluated quantum machine learning models with respect to




FIGURE 4. Behavior of the quantum machine learning models $f(x; \theta)$ over the training phase. (a) Training dataset for the parameter selection $d = 4, \epsilon = 0.1$. (b) Evolving loss for each of the 20 evaluated depthwise models ($L = 1, \dots, 20$) during training. (c) Deviation of the parameter vector from its initialization. (d)–(f) Corresponding information when $d = 4, \epsilon = 0.4$. (g)–(i) for $d = 4, \epsilon = 1.0$.

the increasing number of layers. Column 1 has illustrative samples from the training distribution with rowwise decrements in the signal-to-noise ratio, column 2 gives the corresponding loss curves during training, and column 3 indicates the corresponding change in the magnitude of the parameter vector offset from initialization

$$\frac{\|\theta(n) - \theta(0)\|}{\|\theta(0)\|} \quad (25)$$

where $\theta(0)$ is the value of the parameters at their initialization, and $\theta(n)$ is their value at the n th epoch.

It is evident that none of the models reach the interpolation threshold [53], i.e., the point at which the training data are fitted perfectly with zero training error. To fit the training dataset we would need at least 32 parameters (two nonzero coordinates per point per 16 points). However, we are not able to reach the interpolation threshold even in the deepest configuration with a total of 40 parameters. This behavior is expected by the choice of a parametrically constrained U in effect acting as a form of regularization. As in the case of

classical DNN, an increasing number of parameters results in a decrease in the loss [Fig. 4(b), (e), and (h)], and in an increase in the proximity between the parameter vectors and their initialization [Fig. 4(c), (f), and (i)].

We can conclude that none of the QNN models exhibit evidence of lazy training. In particular, while models having a higher number of parameters do indeed converge more rapidly, parameters are nonetheless varying substantially from their initialization. This behavior is even more noticeable in the smaller models, with a norm difference oscillating substantially prior to the convergence. Such nontrivial training is suggestive of the QPK differing largely from the QNTK in its training characteristics.

2) TEST AND TRAIN ACCURACY OF THE QUANTUM PATH KERNEL VERSES THE QUANTUM NTK

Fig. 5 indicates the corresponding test accuracies, measuring how well the respective models generalize to unseen data. While the QPK and QNTK models both perform similarly at low signal-to-noise ratios, it is particularly striking to observe




FIGURE 5. Respective test accuracy scores for the quantum path kernel model, the Quantum NTK, and the oracle. Error bars represent the standard deviation over three (otherwise identical) experiments having parametric specifications (a) $d = 4, \epsilon = 0.1$. (b) $d = 4, \epsilon = 0.4$. (c) $d = 4, \epsilon = 1.0$.




FIGURE 6. Respective training accuracies of the quantum path kernel model, the Quantum NTK, and the oracle. Error bars represent the standard deviation over three (otherwise identical) experiments having specifications (a) $d = 4, \epsilon = 0.1$. (b) $d = 4, \epsilon = 0.4$. (c) $d = 4, \epsilon = 1.0$.

the outperformance of the QPK over the quantum NTK with increasing hierarchical depth at the highest signal-to-noise setting.

Fig. 6 indicates the training accuracy with depth at the point of convergence. It may be observed that the QPK exhibits lower loss than the quantum NTK across the full signal-to-noise range, with the effect becoming more marked at higher noise levels (ultimately overfitting relative to the noise-free oracle in panel c), consistent with the expectation that QPK has a lower bias than the quantum NTK.

In sum, results confirm the anticipated improvement in performance for the QPK over the QNTK in the Gaussian XOR mixture setting.

V. CONCLUSION

In this article, we have introduced the QPK as a mechanism for incorporating key complex classical multilayer network learning behaviors, in particular hierarchical feature learning, within QNNs via an appropriately expressive kernelization of the training process. We evaluate our approach on the

Gaussian XOR mixture classification problem, a straightforward benchmark of multilayer learning capacity that requires a minimum two-layer solution in order to approach Bayes optimally. The experimental results indicate superior generalization performance relative to the QNTK, an advantage which is especially pronounced in high-depth, low signal-to-noise settings.

We have shown theoretically that the QPK converges to the QNTK only in the lazy training regime, i.e., when the training loss decreases to zero exponentially fast whilst model parameters stay close to their initializations across training. Such behavior is classically seen in infinite-wide neural networks, whose behavior is then close to that of a linear model. Our experiments, by contrast, indicate that QNNs do not operate in the linear regime.

We have discussed, though do not evaluate in the current article, the potential for using stochastic, noisy, or nongradient descent-based optimization techniques to artificially perturb parameter paths within the QPK in order to implicate more decorrelated feature representations. We furthermore, propose in future to extend the QPK approach via

weighting of individual kernel representations in a more heuristic way, for example via multiple kernel learning. We have also referred in passing to the interpretation of the QPK as an ensemble method due to the averaging operation over its kernel matrixes. This will be explored more fully in future investigations.

APPENDIX A THEORETICAL AND IMPLEMENTATIONAL DETAILS OF THE PATH KERNEL IN THE CLASSICAL MACHINE LEARNING DOMAIN

The path kernel was introduced in [20] as a means of replicating arbitrary gradient descent-based machine learning models in the form of kernel machines, under some weak assumptions. The path kernel is consequently of inherent interest in the theory of classical machine learning in that it grants a further layer of interpretability to models, including those, such as the neural networks, that often lacks this [54]. In contrast, kernel machines permit a clear interpretation of prediction functions in terms of linear combinations of data in the training set as a consequence of the Representer Theorem. In particular, [20, Theorem 1] indicates that the machine learning model $h(\mathbf{x}; \mathbf{w}) : \mathbb{R}^D \times \mathbb{R}^P \rightarrow \mathbb{R}$ (with D the dimensionality of the data and P the number of model parameters) can be rewritten equivalently as a function f utilizing the path kernel:

$$h(\mathbf{x}; \bar{\mathbf{w}}) = f(\mathbf{x}, \bar{\gamma}) = \sum_{i=1}^m \alpha_i(\mathbf{x}) \kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i; \bar{\gamma}) + \alpha_0(\mathbf{x}) \quad (26)$$

where

$$\kappa_{\text{pk}} : \mathbb{R}^D \times \mathbb{R}^D \times ([0, T] \rightarrow \mathbb{R}^P) \rightarrow \mathbb{R} \quad (27)$$

$$\kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i; \gamma) = \int_0^T \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \gamma(t)) dt \quad (28)$$

is the path kernel, a parametric kernel function (this parameterization has been rendered explicit in current formulation). In this case, $\bar{\gamma} : [0, 1] \rightarrow \mathbb{R}^P$ is the parameter path, as detailed in Section III, with a terminal parameter value $\bar{\gamma}(1) = \bar{\mathbf{w}}$ and for particular values of α_i defined in Appendix A2. The NTK can also be expressed as a parametric kernel

$$\kappa_{\text{ntk}} : \mathbb{R}^D \times \mathbb{R}^D \times \mathbb{R}^P \rightarrow \mathbb{R} \quad (29)$$

$$\kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \mathbf{w}) = \nabla_{\mathbf{w}} h(\mathbf{x}; \mathbf{w}) \cdot \nabla_{\mathbf{w}} h(\mathbf{x}_i; \mathbf{w}). \quad (30)$$

Equation (26) holds under the proviso that h is differentiable in \mathbf{w} , and trained via gradient descent (GD) for the given training dataset $\{(\mathbf{x}_i, y_i^*)\}_{i=1}^m \subseteq \mathbb{R}^D \times \mathbb{R}$ using the convex differentiable loss function $L(\mathbf{w}) = \sum_{i=1}^M \ell(h(\mathbf{x}_i, \mathbf{w}), y_i^*)$.

Equation (26) differs from a linear model due to the explicit dependency of the data \mathbf{x} in the weights α_i , and it remains a matter of discussion as to whether the path kernel in fact represents a more generalized model class than that of kernel machines (although it is clearly equivalent for infinitely small learning rates [55]). This debate need not

concern us for the present purposes, where the intent is to obtain a class of models capable of representing the network gradient trajectory in a manner expressible on current quantum computers.

As the path kernel is not widely deployed in practical machine learning, we detail here some of its properties. In Appendix A1 we prove the path kernel is a Mercer kernel. In Appendix A2 we briefly comment on the proof of [20, Th. 1]. In Appendix A3 we demonstrate a numerical implementation of the path kernel.

A. PATH KERNEL IS A MERCER KERNEL

Given any $\bar{\gamma}$, the function $\bar{\kappa}_{\text{pk}}(\mathbf{x}, \mathbf{x}') = \kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}'; \bar{\gamma})$ is a positive definite or Mercer kernel on \mathbb{R}^D . A Mercer kernel satisfies

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \kappa(\mathbf{x}_i, \mathbf{x}_j) \geq 0 \quad (31)$$

for all sequence of elements $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^D$ and constants $c_1, \dots, c_n \in \mathbb{R}$.

It is straightforward to demonstrate that such a condition is valid of the path kernel. First, $\bar{\kappa}_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i) = \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \mathbf{w})$ is a positive definite function for any \mathbf{w} in consequence of the positive definiteness of the Gram matrix of inner products in the Hilbert space of the kernel. Second, since both the positive combination and the infinitesimal limit of combinations of positive definite kernels still satisfy the Mercer condition, then the preceding is immediately valid for the path kernel in both its discrete and continuous formulations.

B. COMMENT ON THEOREM 1 IN DOMINGO'S WORK

In this section, we comment on [20, Th. 1] in order to highlight some of its limitations. The dynamics of any predictor under training via gradient descent may be described by a first-order nonhomogeneous differential equation

$$\frac{dh(\mathbf{x}; \mathbf{w})}{dt} = - \sum_{j=1}^P \frac{\partial h}{\partial w_j} \cdot \frac{\partial L}{\partial w_j} \quad (32)$$

where $h(\mathbf{x}; \mathbf{w}) : \mathbb{R}^D \times \mathbb{R}^P$ and L is the convex differentiable loss function. We can describe these predictor dynamics over training in terms of the tangent kernel

$$\frac{dh(\mathbf{x}; \mathbf{w}(t))}{dt} = \sum_{j=1}^d \frac{\partial h(\mathbf{x}; \mathbf{w})}{\partial w_j} \cdot \frac{dw_j}{dt} \quad (33)$$

$$= \sum_{j=1}^d \frac{\partial h(\mathbf{x}; \mathbf{w})}{\partial w_j} \cdot \left(- \frac{\partial L(\mathbf{w}(t))}{\partial w_j} \right) \quad (34)$$

$$= \sum_{j=1}^d \frac{\partial h(\mathbf{x}; \mathbf{w})}{\partial w_j} \cdot \left(- \sum_{i=1}^m \frac{\partial \ell(y_i^*, h(\mathbf{x}_i; \mathbf{w}))}{\partial w_j} \right) \quad (35)$$

$$= \sum_{j=1}^d \frac{\partial h(\mathbf{x}; \mathbf{w})}{\partial w_j} \cdot \left(- \sum_{i=1}^m \frac{\partial \ell(y_i^*, y_i)}{\partial y_i} \frac{\partial h(x_i; w)}{\partial w_j} \right) \quad (36)$$

$$= - \sum_{i=1}^m \frac{\partial \ell}{\partial y_i} \sum_{j=1}^d \frac{\partial h(\mathbf{x}; \mathbf{w})}{\partial w_j} \frac{\partial h(\mathbf{x}_i; \mathbf{w})}{\partial w_j} \quad (37)$$

$$= - \sum_{i=1}^m \frac{\partial \ell}{\partial y_i} \nabla_w h(\mathbf{x}; \mathbf{w}) \cdot \nabla_w h(\mathbf{x}_i; \mathbf{w}) \quad (38)$$

$$= - \sum_{i=1}^m \frac{\partial \ell}{\partial y_i} \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \mathbf{w}). \quad (39)$$

In the limit $\epsilon \rightarrow 0$ we obtain

$$\begin{aligned} h(\mathbf{x}) &= h(\mathbf{x}; \gamma(1)) \\ &= h(\mathbf{x}; \gamma(0)) - \int_0^1 \sum_{i=1}^m \frac{\partial \ell}{\partial y_i} \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \gamma(t)) dt. \end{aligned} \quad (40)$$

Such a function cannot be straightforwardly represented as a linear model. However, by multiplying and dividing by the path kernel itself we obtain the following equation, at the cost of introducing a dependency of \mathbf{x} in the model parameters:

$$\begin{aligned} h(\mathbf{x}; \gamma(1)) &= h(\mathbf{x}; \gamma(0)) \\ &+ \sum_{i=1}^m \left(- \frac{\int_0^1 \frac{\partial \ell}{\partial y_i} \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i; \gamma(t)) dt}{\kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i; \gamma)} \right) \kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i; \gamma) \\ &= h(\mathbf{x}; \gamma(0)) + \sum_{i=1}^m \alpha_i(\mathbf{x}) \kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i; \gamma). \end{aligned} \quad (41)$$

Various works have suggested that imposing stronger assumptions on training can remove the dependency of \mathbf{x} in the model parameters. For example, the authors in [55] achieved this by imposing a requirement that the loss derivative is of constant sign during training.

C. NUMERICAL CALCULATION OF THE PATH KERNEL

We can calculate the value of the path kernel by approximating the integral with a direct sum

$$\begin{aligned} \kappa_{\text{pk}}(\mathbf{x}, \mathbf{x}_i, \gamma) &= \int_0^1 \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i, \gamma(t)) dt \\ &\approx \frac{1}{T} \sum_{i=0}^{T-1} \kappa_{\text{ntk}}(\mathbf{x}, \mathbf{x}_i, \gamma(t)) \end{aligned} \quad (42)$$

where $t = i/T$.

The implementation details are reported in the following pseudocode listings. In Fig. 7, we indicate how to calculate the NTK of the predictor f once the parameter value w is fixed. In particular, the gradient can be calculated with the finite difference method or, if the predictor is implemented with a QNN, with the parameter-shift rule.

```

procedure CREATENEURALTANGENTKERNEL
  Input: predictor function  $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ , data set  $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=0}^{n-1}$ , parameter value  $w \in \mathbb{R}^p$ .
  Output: real symmetric matrix  $n \times n$  representing the neural tangent kernel of  $f$  over the given dataset.
   $M \leftarrow$  zero filled  $n \times p$  matrix
  for  $i \in 0, \dots, n-1$  do
     $M[i] \leftarrow \nabla f(\mathbf{x}_i, w)$ 
  return  $MM^T$ 
  
```

FIGURE 7. Pseudocode for the neural tangent kernel formulation.

```

procedure CREATEPATHKERNEL
  Input: predictor function  $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ , data set  $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=0}^{n-1}$ , parameter path  $\gamma \in \mathbb{R}^{p \times t}$  obtained during the gradient descent-based training phase.
  Output: real symmetric matrix  $n \times n$  representing the Path kernel of  $f$  over the given dataset.
   $M \leftarrow$  zero filled  $t$  elements array
  for  $j \in 0, \dots, t-1$  do
     $w \leftarrow \gamma[j]$ 
     $M \leftarrow M + \text{CREATENEURALTANGENTKERNEL}(f, \{\mathbf{x}_i\}_{i=0}^{n-1}, w)$ 
  return  $\frac{1}{t}M$ 
  
```

FIGURE 8. Pseudocode for the path kernel formulation.

The procedure for calculating the path kernel is shown in Fig. 8 and uses the NTK to calculate the individual contribution of each training epoch and thereafter calculates the average kernel matrix pointwise.

In Section III-B, we have discussed the potential significance of decorrelated features; here, we propose a numerical implementation of the *effective* path kernel. In contrast to the original path kernel, the effective path kernel seeks to avoid biasing due to multiple similar kernel contributions. This is especially important if the training has converged significantly earlier than the last training epoch: any contribution after convergence has the same NTK and will increase its relative weight as the number of epochs after convergence increases. Its formulation is given in Fig. 9. Both the path kernel and effective path kernel can be straightforwardly implemented in parallel over multiple CPUs (or multiple QPUs) for the evaluation of f .

APPENDIX B NUMERICAL EVIDENCE FOR THE INABILITY OF RANDOM FEATURE KERNEL TECHNIQUES IN SOLVING THE GAUSSIAN XOR MIXTURE CLASSIFICATION

In [21], the authors demonstrate that a two-layer-depth neural network with only a small number of neurons can easily outperform kernel methods on the Gaussian mixture classification problem, under the assumption that the number of training data points $n \rightarrow \infty$ is linearly proportional to the dimensionality of the data $d \rightarrow \infty$.

We modify Refinetti's experiment for the current purpose to show the same result in a more straightforward way.

We define the two-layer neural network as the function

$$\begin{aligned} f_{\text{nn}}(\mathbf{x}; W_1, W_2, W_3, b_1, b_2, b_3) \\ = W_3 \cdot \text{relu}(W_2 \cdot \text{relu}(W_1 \cdot \mathbf{x} + b_1) + b_2) + b_3 \end{aligned} \quad (43)$$

```

procedure CREATEEFFECTIVEPATHKERNEL
  Input: predictor function  $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ , data set  $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=0}^{n-1}$ ,
  number of training epochs  $t$ , parameter path  $\gamma \in \mathbb{R}^{p \times t}$  obtained during the
  gradient descent-based training phase, correlation threshold  $C \in [0, 1]$ .
  Output: real symmetric matrix  $n \times n$  representing the Effective Path kernel
  of  $f$  over the given dataset.
   $\ell \leftarrow$  CREATEEFFECTIVEPATHKERNELREC( $f, \{\mathbf{x}_i\}_{i=0}^{n-1}, \gamma, C, 0, t - 1$ )
   $n \leftarrow$  number of elements in the list  $\ell$ 
  return  $\frac{1}{n} \sum_{i=0}^{n-1} \ell_i$ 

procedure CREATEEFFECTIVEPATHKERNELREC
  Input: predictor function  $f : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}$ , data set  $\{\mathbf{x}_i \in \mathbb{R}^d\}_{i=0}^{n-1}$ ,
  parameter path  $\gamma \in \mathbb{R}^{p \times t}$  obtained during the gradient descent-based training
  phase, correlation threshold  $C \in [0, 1]$ , start instant  $t_s \in [0, \dots, t]$ , end instant
   $t_e \in [0, \dots, t]$ ,  $t_s < t_e$ .
  Output: array of  $t$  elements, each one coarsely quantifying the curvature
  of the parameter path at each training epoch.
  ▷ Start procedure
  if  $t_s \geq t_e$  then
    return [] ▷ Empty list
   $M_s \leftarrow$  CREATENEURALTANGENTKERNEL( $f, \{\mathbf{x}_i\}_{i=0}^{n-1}, \gamma[t_s]$ )
   $M_e \leftarrow$  CREATENEURALTANGENTKERNEL( $f, \{\mathbf{x}_i\}_{i=0}^{n-1}, \gamma[t_e]$ )
   $c \leftarrow$  correlation between  $M_s$  and  $M_e$  ▷ interpret the matrix as vectors to
  calculate the correlation, or change evaluation metric, e.g. Frobenius norm
  if  $|c| > C$  then
    return  $[M_s, M_e]$  ▷ Highly correlated representation
  else if  $t_s + 1 < t_e$  then
     $t_m \leftarrow \text{int}(t_s/2 + t_e/2)$ 
     $L \leftarrow$  CREATEEFFECTIVEPATHKERNELREC( $f, \{\mathbf{x}_i\}_{i=0}^{n-1}, \gamma, C, t_s, t_m$ )
     $R \leftarrow$  CREATEEFFECTIVEPATHKERNELREC( $f, \{\mathbf{x}_i\}_{i=0}^{n-1}, \gamma, C, t_m + 1, t_e$ )
    return  $[M_s, M_e] \cup L \cup R$  ▷ Concatenate lists

```

FIGURE 9. Pseudocode for the effective path kernel formulation.

parameterized by $W_1 \in \mathbb{R}^{h \times d}$, $W_2 \in \mathbb{R}^{h \times h}$, $W_3 \in \mathbb{R}^{1 \times h}$, b_1 , $b_2 \in \mathbb{R}^{h \times 1}$, $b_3 \in \mathbb{R}$, where h is the number of hidden neurons per layer (the number of hidden neurons here is fixed to $h = \lceil \sqrt{d} \rceil$). In our setting, we randomly initialize the weights W_1 , W_2 , and W_3 by sampling the matrix element i.i.d. from a Gaussian of zero mean and unitary variance. The model is then trained using the gradient-descent-based algorithm ADAM for a maximum 1000 epochs with learning rate 0.001 (the model is implemented in Python3 library scikit-learn, with the default configuration).

We define a random feature kernel machine as

$$\kappa_{\text{rf}}(\mathbf{x}, \mathbf{x}') = \langle \phi(\mathbf{x}), \phi(\mathbf{x}') \rangle, \quad \phi(\mathbf{x}) = \text{relu}(W \cdot \mathbf{x}) \quad (44)$$

with the activation weights parameterized by $W \in \mathbb{R}^{f \times d}$, $w_{i,j} \sim \mathcal{N}(0, 1)$, where f has been chosen such that the number of parameters of the random feature kernel is greater than or equal to the number of parameters in the neural network, thus

$$f = \frac{(dh + hh + h) + (h + h + 1)}{d}. \quad (45)$$

For $h = \lceil \sqrt{d} \rceil$ we can tightly upper bound f with $f < \lceil \sqrt{d} \rceil + 5$. This kernel function is then fed to a SVM for classification (as implemented in scikit-learn).

We randomly generate the dataset $\mathcal{D}_{d,d',\epsilon,n}$, as detailed in Section IV-A. The experiment described below consists in comparing the performance of the neural network classifier with variations of the random feature kernel on the dataset $\mathcal{D}_{d,3,\epsilon,16d}$ for data point dimensionality $d = 4, 8, 12, 16, 20$ and noise $\epsilon = 0, 0.1, 0.2, \dots, 1.9, 2.0$. We keep the number of nonzero features $d' = 3$, meaning we are effectively classifying 3-D Gaussian XOR mixtures, with the number of

training vectors of the dataset fixed to be $16d$. The dataset is then randomly split 75% in the training dataset and 25% in the test set. For each dataset, we compare the performances of the oracle with the performances of the best of ten randomly initialized neural networks and the best of ten random feature kernels. For each dataset specification, we repeat this procedure ten times.

In Fig. 10, we set out the results of the above described experiments. It may be observed that neural networks outperform the kernel approach in each case, with the differential in accuracy increasing with the number of zero-valued features. Refinetti et al. [21] suggested that this difference in performance is accounted for by the fact that random feature kernels in high dimension behave as linear transformations [56].

Here, we have suggested a complementary interpretation of the results of such experiments. We have shown that the difference of performance between the two models is not uniquely determined by the failure of kernel methods per se. In fact, it is determined also by the feature learning capabilities of neural networks; inspecting the evolution of the W_1 parameters during the training of a neural network reveals that elements in W_1 related to the zero-features do indeed go to zero (Fig. 11). This results in having all of the hidden neurons (whose number is proportional to \sqrt{d} , and thus, increasing with the number of features) working adaptively to classify the three discriminatively informative components or features, thereby improving overall performance in contrast to the random feature kernel approach, for which adding feature (and parameters) drastically decreases performance (which is to say the path model outperforms the random feature kernel in this problem by being able to discharge junk features candidates, thus performing feature learning).

APPENDIX C DATA, CODE, AND SIMULATION DETAILS

Both the code to reproduce the indicated experiments and also the relevant data are freely available at <https://github.com/incud/QuantumPathKernel>. The code is released open-source.

The indicated experiments have been simulated on two devices as follows.

- 1) *one Dell Latitude 5510 having:* Intel Core i7-10610 U CPU with 4 physical cores, 16 GB RAM, without CUDA-enabled GPUs;
- 2) *one cluster node having:* Intel Xeon Silver 4216 CPU with 64 physical cores, 180 GB RAM, with $4 \times$ CUDA-enabled GPUs NVidia Tesla V100S 32 GB VRAM.

The software runs on Ubuntu 20.04 LTS and uses Python v3.9.11, PiP packet-manager v22.0.4 along with the other libraries listed in `requirements.txt` file in the root of the attached repository. Installation and simulation instructions are documented in the `README.md` file in the root




FIGURE 10. Comparison of the performance of random feature kernel and (two layer) neural networks over the 3-D Gaussian xor mixture problem with an increasing number of features set to zero. (a)–(f) have, respectively, 4, 8, 12, 16, 20, and 24 features per point, the first three being the only nonzero ones.




FIGURE 11. Values of the W_1 matrix for individual neural networks of the form of (43) during training on the Gaussian xor mixture datasets $\mathcal{D}_{24,3,0.8,384}$: (a)–(c) Coefficients at initialization, after 250 training epochs and after 750 training epochs of training with ADAM at a learning rate 0.001.

of the repository. Our code is based upon freely available, open-source frameworks only.

The framework used to define and simulate the quantum circuit is PennyLane [57]. The simulations have been accelerated using the JAX library [58]. (JAX might require installation from source code if used on operating systems different from Ubuntu). Alternatively, the source code can be set such that PennyLane does not require this library. (However, in this case, the circuit simulation might be substantially slower and would not benefit the full potential of multicore CPUs and GPUs.) These experiments have not been run on quantum hardware.

The input and output of each experiment are contained in different subfolders within the root directory. They contain the specifications needed to generate the training and testing

datasets, the datasets themselves, the trace of the parameters during the training for any model, and the QNTK and QPK gram matrixes for each model (which may be used to create a pretrained model), and also the resulting plots. The README.md explains in detail the commands needed to reproduce our results.

The simulations for all experiments have taken approximately 600 h across both machines used.

REFERENCES

- [1] P. Wittek, *Quantum Machine Learning: What Quantum Computing Means to Data Mining*. Cambridge, MA, USA: Academic, 2014, doi: [10.1016/C2013-0-19170-2](https://doi.org/10.1016/C2013-0-19170-2).
- [2] M. Schuld, I. Sinayskiy, and F. Petruccione, “An introduction to quantum machine learning,” *Contemporary Phys.*, vol. 56, no. 2, pp. 172–185, 2015, doi: [10.1080/00107514.2014.964942](https://doi.org/10.1080/00107514.2014.964942).

- [3] H.-Y. Huang et al., “Power of data in quantum machine learning,” *Nature Commun.*, vol. 12, no. 1, 2021, Art. no. 2631, doi: [10.1038/s41467-021-22539-9](https://doi.org/10.1038/s41467-021-22539-9).
- [4] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, “The power of quantum neural networks,” *Nature Comput. Sci.*, vol. 1, no. 6, pp. 403–409, 2021, doi: [10.1038/s43588-021-00084-1](https://doi.org/10.1038/s43588-021-00084-1).
- [5] Y. Liu, S. Arunachalam, and K. Temme, “A rigorous and robust quantum speed-up in supervised machine learning,” *Nature Phys.*, vol. 17, no. 9, pp. 1013–1017, 2021, doi: [10.1038/s41567-021-01287-z](https://doi.org/10.1038/s41567-021-01287-z).
- [6] H.-Y. Huang et al., “Quantum advantage in learning from experiments,” *Science*, vol. 376, no. 6598, pp. 1182–1186, 2022, doi: [10.1126/science.abn7293](https://doi.org/10.1126/science.abn7293).
- [7] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Commun.*, vol. 9, no. 1, 2018, Art. no. 4812, doi: [10.1038/s41467-018-07090-4](https://doi.org/10.1038/s41467-018-07090-4).
- [8] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, “Connecting ansatz expressibility to gradient magnitudes and barren plateaus,” *PRX Quantum*, vol. 3, no. 1, 2022, Art. no. 10313, doi: [10.1103/PRXQuantum.3.010313](https://doi.org/10.1103/PRXQuantum.3.010313).
- [9] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, “Layerwise learning for quantum neural networks,” *Quantum Mach. Intell.*, vol. 3, no. 1, Jan. 2021, Art. no. 5, doi: [10.1007/2Fs42484-020-00036-4](https://doi.org/10.1007/2Fs42484-020-00036-4).
- [10] M. Schuld, “Supervised quantum machine learning models are kernel methods,” 2021, *arXiv:2101.11020*, doi: [10.48550/arXiv.2101.11020](https://doi.org/10.48550/arXiv.2101.11020).
- [11] J. Liu, F. Tacchino, J. R. Glick, L. Jiang, and A. Mezzacapo, “Representation learning via quantum neural tangent kernels,” *PRX Quantum*, vol. 3, no. 3, 2022, Art. no. 030323, doi: [10.1103/PRXQuantum.3.030323](https://doi.org/10.1103/PRXQuantum.3.030323).
- [12] J. Liu, K. Najafi, K. Sharma, F. Tacchino, L. Jiang, and A. Mezzacapo, “Analytic theory for the dynamics of wide quantum neural networks,” *Phys. Rev. Lett.*, vol. 130, no. 15, 2023, Art. no. 150601, doi: [10.1103/PhysRevLett.130.150601](https://doi.org/10.1103/PhysRevLett.130.150601).
- [13] K. Sharma, M. Cerezo, L. Cincio, and P. J. Coles, “Trainability of dissipative perceptron-based quantum neural networks,” *Phys. Rev. Lett.*, vol. 128, no. 18, 2022, Art. no. 180505, doi: [10.1103/PhysRevLett.128.180505](https://doi.org/10.1103/PhysRevLett.128.180505).
- [14] A. Jacot, F. Gabriel, and C. Hongler, “Neural tangent kernel: Convergence and generalization in neural networks,” *Adv. Neural Inf. Process. Syst.*, vol. 31, pp. 8580–8589, 2018, doi: [10.1145/3406325.3465355](https://doi.org/10.1145/3406325.3465355).
- [15] L. Chizat, E. Oyallon, and F. Bach, “On lazy training in differentiable programming,” *Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 2937–2947, doi: [10.48550/arXiv.1812.07956](https://doi.org/10.48550/arXiv.1812.07956).
- [16] C. Liu, L. Zhu, and M. Belkin, “On the linearity of large non-linear models: When and why the tangent kernel is constant,” *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 15954–15964, 2020, doi: [10.48550/arXiv.2010.01092](https://doi.org/10.48550/arXiv.2010.01092).
- [17] B. Ghorbani, S. Mei, T. Misiakiewicz, and A. Montanari, “When do neural networks outperform kernel methods?,” *J. Statist. Mech.: Theory Exp.*, vol. 2021, no. 12, 2021, Art. no. 124009, doi: [10.1088/1742-5468/ac3a81](https://doi.org/10.1088/1742-5468/ac3a81).
- [18] M. Chen et al., “Towards understanding hierarchical learning: Benefits of neural representations,” *Adv. Neural Inf. Process. Syst.*, vol. 33, pp. 22134–22145, 2020, doi: [10.48550/arXiv.2006.13436](https://doi.org/10.48550/arXiv.2006.13436).
- [19] N. Shirai, K. Kubo, K. Mitarai, and K. Fujii, “Quantum tangent kernel,” 2021, *arXiv:2111.02951*, doi: [10.48550/arXiv.2111.02951](https://doi.org/10.48550/arXiv.2111.02951).
- [20] P. Domingos, “Every model learned by gradient descent is approximately a kernel machine,” 2020, *arXiv:2012.00152*, doi: [10.48550/arXiv.2012.00152](https://doi.org/10.48550/arXiv.2012.00152).
- [21] M. Refinetti, S. Goldt, F. Krzakala, and L. Zdeborová, “Classifying high-dimensional Gaussian mixtures: Where kernel methods fail and neural networks succeed,” in *Proc. 38th Int. Conf. Mach. Learn.*, 2021, pp. 8936–8947, doi: [10.48550/arXiv.2102.11742](https://doi.org/10.48550/arXiv.2102.11742).
- [22] B. Ghorbani, S. Krishnan, and Y. Xiao, “An investigation into neural net optimization via Hessian eigenvalue density,” in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 2232–2241, doi: [10.48550/arXiv.1901.10159](https://doi.org/10.48550/arXiv.1901.10159).
- [23] S. Arora, S. S. Du, W. Hu, Z. Li, R. R. Salakhutdinov, and R. Wang, “On exact computation with an infinitely wide neural net,” *Adv. Neural Inf. Process. Syst.*, vol. 32, pp. 8141–8150, 2019, doi: [10.48550/arXiv.1904.11955](https://doi.org/10.48550/arXiv.1904.11955).
- [24] Y. Bai, B. Krause, H. Wang, C. Xiong, and R. Socher, “Taylorized training: Towards better approximation of neural network training at finite width,” 2020, *arXiv:2002.04010*, doi: [10.48550/arXiv.2002.04010](https://doi.org/10.48550/arXiv.2002.04010).
- [25] M. Schuld, I. Sinayskiy, and F. Petruccione, “The quest for a quantum neural network,” *Quantum Inf. Process.*, vol. 13, no. 11, pp. 2567–2586, 2014, doi: [10.1007/s11228-014-0809-8](https://doi.org/10.1007/s11228-014-0809-8).
- [26] M. Schuld, I. Sinayskiy, and F. Petruccione, “Simulating a perceptron on a quantum computer,” *Phys. Lett. A*, vol. 379, no. 7, pp. 660–663, Mar. 2015, doi: [10.1016/j.physleta.2014.11.061](https://doi.org/10.1016/j.physleta.2014.11.061).
- [27] Y. Cao, G. G. Guerreschi, and A. Aspuru-Guzik, “Quantum neuron: An elementary building block for machine learning on quantum computers,” 2017, *arXiv:1711.11240*, doi: [10.48550/arXiv.1711.11240](https://doi.org/10.48550/arXiv.1711.11240).
- [28] W. Hu, “Towards a real quantum neuron,” *Natural Sci.*, vol. 10, no. 3, pp. 99–109, 2018, doi: [10.4236/ns.2018.103011](https://doi.org/10.4236/ns.2018.103011).
- [29] K. Gili, M. Sveistrys, and C. Ballance, “Introducing non-linearity into quantum generative models,” 2022, *arXiv:2205.14506*, doi: [10.1103/PhysRevA.107.012406](https://doi.org/10.1103/PhysRevA.107.012406).
- [30] F. Tacchino, C. Macchiavello, D. Gerace, and D. Bajoni, “An artificial neuron implemented on an actual quantum processor,” *NPJ Quantum Inf.*, vol. 5, no. 1, 2019, Art. no. 26, doi: [10.1038/s41534-019-0140-4](https://doi.org/10.1038/s41534-019-0140-4).
- [31] N. Guo, K. Mitarai, and K. Fujii, “Nonlinear transformation of complex amplitudes via quantum singular value transformation,” 2021, *arXiv:2107.10764*, doi: [10.48550/arXiv.2107.10764](https://doi.org/10.48550/arXiv.2107.10764).
- [32] Z. Holmes, N. Coble, A. T. Sornborger, and Y. Subaşlı, “On nonlinear transformations in quantum computation,” 2021, *arXiv:2112.12307*, doi: [10.48550/arXiv.2112.12307](https://doi.org/10.48550/arXiv.2112.12307).
- [33] A. Daskin, “A simple quantum neural net with a periodic activation function,” in *Proc. IEEE Int. Conf. Systems, Man, Cybern.*, 2018, pp. 2887–2891, doi: [10.48550/arXiv.1804.07633](https://doi.org/10.48550/arXiv.1804.07633).
- [34] S. Weinberg, “Precision tests of quantum mechanics,” *Phys. Rev. Lett.*, vol. 62, pp. 485–488, Jan. 1989, doi: [10.1103/PhysRevLett.62.485](https://doi.org/10.1103/PhysRevLett.62.485).
- [35] D. S. Abrams and S. Lloyd, “Nonlinear quantum mechanics implies polynomial-time solution for NP-complete and #P problems,” *Phys. Rev. Lett.*, vol. 81, pp. 3992–3995, Nov. 1998, doi: [10.1103/PhysRevLett.81.3992](https://doi.org/10.1103/PhysRevLett.81.3992).
- [36] X. Wang, J. Liu, T. Liu, Y. Luo, Y. Du, and D. Tao, “Symmetric pruning in quantum neural networks,” 2022, *arXiv:2208.14057*, doi: [10.48550/arXiv.2208.14057](https://doi.org/10.48550/arXiv.2208.14057).
- [37] J. Liu, Z. Lin, and L. Jiang, “Laziness, barren plateau, and noise in machine learning,” 2022, *arXiv:2206.09313*, doi: [10.48550/arXiv.2206.09313](https://doi.org/10.48550/arXiv.2206.09313).
- [38] E. Abedi, S. Beigi, and L. Taghavi, “Quantum lazy training,” *Quantum*, vol. 7, 2023, Art. no. 989, doi: [10.22331/q-2023-04-27-989](https://doi.org/10.22331/q-2023-04-27-989).
- [39] S. Jerbi, L. J. Fiderer, H. P. Nautrup, J. M. Kübler, H. J. Briegel, and V. Dunjko, “Quantum machine learning beyond kernel methods,” 2021, *arXiv:2110.13162*, doi: [10.1038/s41467-023-36159-y](https://doi.org/10.1038/s41467-023-36159-y).
- [40] H. Q. Minh, P. Niyogi, and Y. Yao, “Mercer’s theorem, feature maps, and smoothing,” in *Proc. Learn. Theory: 19th Annu. Conf. Learn. Theory*, 2006, pp. 154–168, doi: [10.1007/11776420_14](https://doi.org/10.1007/11776420_14).
- [41] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2022, doi: [10.1007/11776420_14](https://doi.org/10.1007/11776420_14).
- [42] D. Meltzer and J. Liu, “Catapult dynamics and phase transitions in quadratic nets,” 2023, *arXiv:2301.07737*, doi: [10.48550/arXiv.2301.07737](https://doi.org/10.48550/arXiv.2301.07737).
- [43] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015, doi: [10.1038/nature14539](https://doi.org/10.1038/nature14539).
- [44] E. Peters and M. Schuld, “Generalization despite overfitting in quantum machine learning models,” 2022, *arXiv:2209.05523*, doi: [10.48550/arXiv.2209.05523](https://doi.org/10.48550/arXiv.2209.05523).
- [45] M. Larocca, N. Ju, D. García-Martín, P. J. Coles, and M. Cerezo, “Theory of overparametrization in quantum neural networks,” 2021, *arXiv:2109.11676*, doi: [10.1038/s43588-023-00467-6](https://doi.org/10.1038/s43588-023-00467-6).
- [46] M. Gönen and E. Alpaydm, “Multiple kernel learning algorithms,” *J. Mach. Learn. Res.*, vol. 12, pp. 2211–2268, 2011, doi: [10.5555/1953048.2021071](https://doi.org/10.5555/1953048.2021071).
- [47] Z. Deng, A. Kammoun, and C. Thrampoulidis, “A model of double descent for high-dimensional binary linear classification,” 2019, *arXiv:1911.05822*, doi: [10.48550/arXiv.1911.05822](https://doi.org/10.48550/arXiv.1911.05822).
- [48] X. Mai and Z. Liao, “High dimensional classification via regularized and unregularized empirical risk minimization: Precise error and optimal loss,” 2019, *arXiv:1905.13742*, doi: [10.48550/arXiv.1905.13742](https://doi.org/10.48550/arXiv.1905.13742).
- [49] M. Lelarge and L. Miolane, “Asymptotic Bayes risk for Gaussian mixture in a semi-supervised setting,” in *Proc. IEEE 8th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, 2019, pp. 639–643, doi: [10.48550/arXiv.1907.03792](https://doi.org/10.48550/arXiv.1907.03792).

- [50] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, "Diagnosing barren plateaus with tools from quantum optimal control," *Quantum*, vol. 6, 2022, Art. no. 824, doi: [10.22331/q-2022-09-29-824](https://doi.org/10.22331/q-2022-09-29-824).
- [51] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, "Cost function dependent barren plateaus in shallow parametrized quantum circuits," *Nature Commun.*, vol. 12, no. 1, 2021, Art. no. 1791, doi: [10.1038/s41467-021-21728-w](https://doi.org/10.1038/s41467-021-21728-w).
- [52] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, "An initialization strategy for addressing barren plateaus in parametrized quantum circuits," *Quantum*, vol. 3, 2019, Art. no. 214, doi: [10.22331/q-2019-12-09-214](https://doi.org/10.22331/q-2019-12-09-214).
- [53] M. Belkin, D. Hsu, S. Ma, and S. Mandal, "Reconciling modern machine-learning practice and the classical bias–variance trade-off," *Proc. Nat. Acad. Sci.*, vol. 116, no. 32, pp. 15849–15854, 2019, doi: [10.1073/pnas.1903070116](https://doi.org/10.1073/pnas.1903070116).
- [54] A. Ghorbani, A. Abid, and J. Zou, "Interpretation of neural networks is fragile," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, no. 1, pp. 3681–3688, doi: [10.48550/arXiv.1710.10547](https://doi.org/10.48550/arXiv.1710.10547).
- [55] Y. Chen, W. Huang, L. Nguyen, and T.-W. Weng, "On the equivalence between neural network and support vector machine," *Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 23478–23490, 2021, doi: [10.48550/arXiv.2111.06063](https://doi.org/10.48550/arXiv.2111.06063).
- [56] N. El Karoui, "The spectrum of kernel random matrices," *Ann. Statist.*, vol. 38, no. 1, pp. 1–50, 2010, doi: [10.48550/arXiv.1001.0492](https://doi.org/10.48550/arXiv.1001.0492).
- [57] V. Bergholm et al., "PennyLane: Automatic differentiation of hybrid quantum-classical computations," 2018, *arXiv:1811.04968*, doi: [10.48550/arXiv.1811.04968](https://doi.org/10.48550/arXiv.1811.04968).
- [58] J. Bradbury et al., "JAX: Composable transformations of Python NumPy programs," 2018. [Online]. Available: <http://github.com/google/jax>

Open Access funding provided by 'Università degli Studi di Verona' within the CRUI CARE Agreement