



Original Paper

Quantitative Convergence of Trained Quantum Neural Networks to a Gaussian Process

Anderson Melchor Hernandez, Filippo Girardi,
Davide Pastorello and Giacomo De Palma

Abstract. We study quantum neural networks where the generated function is the expectation value of the sum of single-qubit observables across all qubits. In (Girardi et al., CMP (2025)), it is proven that the probability distributions of such generated functions converge in distribution to a Gaussian process in the limit of infinite width for both untrained networks with randomly initialized parameters and trained networks. In this paper, we provide a quantitative proof of this convergence in terms of the Wasserstein distance of order 1. First, we establish an upper bound on the distance between the probability distribution of the function generated by any untrained network with finite width and the Gaussian process with the same covariance. This proof utilizes Stein's method to estimate the Wasserstein distance of order 1. Next, we analyze the training dynamics of the network via gradient flow, proving an upper bound on the distance between the probability distribution of the function generated by the trained network and the corresponding Gaussian process. This proof is based on a quantitative upper bound on the maximum variation of a parameter during training. This bound implies that for sufficiently large widths, training occurs in the lazy regime, i.e., each parameter changes only by a small amount. While the convergence result of Girardi et al., CMP (2025) holds at a fixed training time, our upper bounds are uniform in time and hold even as $t \rightarrow \infty$.

Contents

1. Introduction
 - 1.1. Our Results

2. Preliminaries
 - 2.1. Training Data
 - 2.2. Quantum Neural Networks
 - 2.3. Light Cones
 - 2.4. Assumptions on Architecture and Initialization
 - 2.5. The Neural Tangent Kernel
 - 2.6. The Assumptions for the Neural Tangent Kernel
 - 2.7. The Expected Behavior of $N(m)$ in Circuits with Logarithmic Depth
 - 2.8. The Wasserstein Distance of Order 1
 3. Some Relevant Facts About Gaussian Processes
 4. Multidimensional Stein Methods
 - 4.1. The Analytic Solution for the Linearized Model
 5. Main Results
 6. Convergence at Initialization
 7. Convergence of the Trained Network
 - 7.1. Quantitative Lazy Training for a Finite Size Circuit
 - 7.2. The Wasserstein Distance Between f and f^{lin}
 - 7.3. The Closeness Between f^{lin} and a Gaussian Process
 8. Conclusions
- Acknowledgements
- Appendix A. Regularity of the Solutions of the Stein Equation
- Appendix B. Convergence Toward a Gaussian Process
- References

1. Introduction

In recent years, scientific communities have become increasingly involved in the use of Artificial Intelligence (AI) to analyze large databases [1, 2]. AI currently encompasses a vast number of subfields, ranging from learning theory to the mathematical foundations of its development. The core of AI and machine learning lies in the recognition and identification of complex patterns within vast amounts of data, enabling these systems to uncover hidden relationships, make informed predictions, and generate insights that would be difficult, if not impossible, for humans to discern on their own [3]. Among the new emerging disciplines, quantum machine learning (QML) is an interdisciplinary field that merges the principles of quantum computing with classical machine learning techniques [4–6]. One of the core ideas of QML is to harness quantum algorithms and the unique properties of quantum mechanics such as superposition, entanglement, and quantum parallelism to enhance the performance of deep neural models [7]. Quantum neural networks constitute the quantum version of deep neural models. These new models are based on quantum circuits and generate functions given by the expectation values of a quantum observable measured on the output of a quantum circuit made by parametric one-qubit

and two-qubit gates [8, 9]. The parameters of the circuit encode both the input data and the parameters of the model itself. These parameters are typically optimized by gradient descent, which involves iterative adjustment to minimize a cost function and improve the performance of the quantum circuit in the processing and analysis of data [10].

In this paper, we study quantum neural networks applied to supervised learning for binary classification. Let \mathcal{X} be the set of possible inputs (e.g., pictures encoded as points in \mathbb{R}^d), which we assume to be finite. Let Θ be the vector of the parameters, and let $x \mapsto f(\Theta, x)$ be the function generated by the quantum neural network. Let $\{(x^{(i)}, y^{(i)}) : i = 1, \dots, n\}$ be the set of the training examples made by the training inputs $x^{(i)} \in \mathcal{X}$ (e.g., pictures of dogs or cats) and the corresponding training labels $y^{(i)} \in \{\pm 1\}$ (e.g., $y^{(i)} = 1$ if $x^{(i)}$ represents a dog and $y^{(i)} = -1$ if $x^{(i)}$ represents a cat). The goal of supervised learning is to adjust the parameters Θ so that $f(\Theta, x)$ reproduces as closely as possible the training examples. This is usually achieved by minimizing a loss function such as the empirical quadratic loss

$$\mathcal{L}(\Theta) = \sum_{i=1}^n \left(f(\Theta, x^{(i)}) - y^{(i)} \right)^2 \quad (1.1)$$

via gradient descent. For simplicity, in this paper we will consider the continuous-time gradient flow rather than gradient descent.

Several works have focused on the analysis of quantum neural networks, as it is believed that they can combine the computational power of quantum computers with the capabilities of deep learning algorithms [11]. In recent work [12], the authors had shown that an exponential quantum speed-up can be obtained via the use of a quantum-enhanced feature space, where each data point is mapped in a nonlinear way to a quantum state, and then classified by a linear classifier in a high-dimensional Hilbert space [13]. Nevertheless, a significant disadvantage lies in the need to determine the appropriate parameters to configure the quantum circuit beforehand, and it is not yet clear whether these parameters can be effectively obtained using a variational technique [14]. In fact, one of the primary challenges associated with training quantum neural networks is the so-called barren-plateau phenomenon, related to different causes like entanglement in input states and the locality of observables, which poses a major obstacle by causing gradients to vanish during the optimization process, making it difficult to train the network effectively [15–17]. To address this challenge, several methods have been proposed recently. For instance, in [18], the authors consider the problem from the perspective of quantum tensor network optimization; in [19], the authors presented a general theory based on the Lie algebra of the circuit’s generators to describe the rise of barren plateaus; in [20], the authors solved the problem of barren plateaus in the specific case of the Hamiltonian variational ansatz, a kind of quantum circuit used in quantum many-body problems; in [21], an efficient state ansatz is proposed to mitigate barren plateaus in the context of variational quantum eigensolvers;

in [22], a lower bound on the variance of the circuit gradients is derived for quantum circuits composed by local 2-designs.

In [23], the authors considered an untrained variational quantum circuit and demonstrated that the associated cost function exhibits an exponentially vanishing gradient, highlighting the difficulty of overcoming this issue. In contrast, the authors in [24] explored more general cost functions inspired by the quantum generalization of optimal mass transport theory [25] developed in [26], offering a different approach to mitigate the challenges posed by the barren-plateau phenomenon. An important problem in the classical theory of deep learning is the analysis of the behavior of a neural network in the limit of infinite width. Significant progress has been made in addressing the question of whether training can perfectly fit the training examples while simultaneously avoiding overfitting. A fundamental breakthrough has been the proof that, in the limit of infinite width, the probability distribution of the function generated by a deep neural network trained on a supervised learning problem converges to a Gaussian process [27–29]. These recent developments have inspired a renewed interest in this important result within the field of quantum machine learning. Specifically, they have led to investigations into whether quantum neural networks exhibit similar properties. In this context, several studies have emerged. For example, [30] proves that in the infinite width limit, trained quantum neural networks with constant depth operate in the lazy regime (i.e., the maximum amount by which the training can change a parameter tends to zero) and are capable of perfectly fitting the training examples. For a contrasting scenario, see [31].

In the recent work [8], the authors rigorously generalized the above classical breakthrough to the context of quantum neural networks. They considered quantum neural networks trained on supervised learning tasks, where the objective function is defined as the expected value of the sum of single-qubit observables across all qubits. The authors proved, for the first time, the trainability in the limit of infinite width in any regime where the depth is allowed to grow with the number of qubits (denoted by m), as long as barren plateaus do not arise. More precisely, Ref. [8] first proves that the probability distribution of the function generated by a randomly initialized quantum neural network in the limit of infinite width converges in distribution to a Gaussian process when the parameters on which each measured qubit depends influence a small number of other measured qubits [8, Theorem 3.7]. Ref. [8] then proves that for quantum neural networks trained in continuous time with gradient flow, the training occurs in the lazy regime and is able to perfectly fit the training set. The key element of the proof is showing that the dependence of the generated function on the parameters can be approximated by its linearized version near the initialization values of the parameters. Consequently, the linearized evolution equation has an analytic solution whose probability distribution is Gaussian with analytically computable mean and covariance. As a result, the probability distribution of the function generated by the trained network is proved to converge in distribution to the aforementioned Gaussian process ([8, Theorem 4.15]).

1.1. Our Results

In this paper, we prove a quantitative version of the results of [8]. Our first goal is to establish explicit upper bounds on the distance between the probability distribution of the function generated by the quantum network at initialization and the corresponding limit Gaussian process. We employ the Wasserstein distance of order 1, which we denote with d_W , and we provide an explicit upper bound that tends to zero as the number of qubits diverges. Without giving all the details, we prove the following (see Theorem 5.1 for the formal statement):

Theorem 1.1 (Convergence at initialization, informal statement). *We denote with \bar{X} the vector made by the elements of \mathcal{X} , and with $f(\Theta, \bar{X})$ the vector made by the associated outputs. Let $\bar{\mathcal{K}}_0$ be the covariance matrix of $f(\Theta, \bar{X})$ when the parameters Θ are randomly initialized. Then, there exist positive numbers $\{\alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}\}$ given by (5.1) and depending on the number of qubits of the network m , on the architecture of the network, on the number of possible inputs \bar{N} and on the covariance matrix $\bar{\mathcal{K}}_0$ such that*

$$d_W(f(\Theta, \bar{X}), \mathcal{N}(0, \bar{\mathcal{K}}_0)) \leq \alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}, \quad (1.2)$$

where $\mathcal{N}(0, \bar{\mathcal{K}}_0)$ denotes the centered Gaussian distribution with covariance matrix $\bar{\mathcal{K}}_0$.

To obtain the explicit expression (5.1) for $\alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}$, we employ Stein's method, which is a powerful method to estimate the Wasserstein distance of order 1 of the sum of weakly dependent random variables [32–36]. We will prove that $\alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}$ depends on the size of the light cones of the quantum circuit and on a normalizing constant that quantifies the presence or absence of the barren-plateau phenomenon. The dependence of $\alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}$ on \bar{N} scales as \bar{N}^2 , causing $\alpha_m^{\bar{N}, \bar{\mathcal{K}}_0}$ to diverge as the number of inputs increases: the nature of the network introduces dependencies between different qubits, which cause the bound (1.2) to be worse than the bound obtained for the sum of independent random variables in the central limit theorem.

Then, we quantify the distance between the probability distribution of the function generated by the trained network and its associated Gaussian process. We prove a strong quantitative version of the results of [8] that can be applied to any quantum neural network with finite width and therefore does not require to build a sequence of networks of increasing width. As before, without giving all the details, we prove the following (see Theorem 5.2 for the formal statement):

Theorem 1.2 (Convergence of the trained network, informal statement). *With the same notation of Theorem 1.1, let us denote by $f(\Theta_t, \bar{X})$ the vector of the outputs of a quantum neural network trained via gradient flow for a time $t > 0$, where Θ_t is the vector of the trained parameters and \bar{X} is the vector of the possible inputs. Let X be the vector of the training inputs. Then, there exists a Gaussian probability distribution $\mathcal{N}(\mu_t(\bar{X}), \mathcal{K}_t(\bar{X}, \bar{X}))$ with mean $\mu_t := \mu_t(\bar{X})$*

and covariance matrix $\overline{\mathcal{K}}_t := \mathcal{K}_t(\overline{X}, \overline{X})$ (that will be determined analytically in (4.17) and (4.16)) such that

$$d_W^{(s)}(f(\Theta_t, \overline{X}), \mathcal{N}(\mu_t, \overline{\mathcal{K}}_t)) \leq \gamma_m^{\overline{N}, n, \overline{\mathcal{K}}_0}, \quad (1.3)$$

where $\{\gamma_m^{\overline{N}, n, \overline{\mathcal{K}}_0}\}$ are positive numbers given by (5.2) depending on the number of possible inputs \overline{N} , on the number of qubits m , on the architecture of the network, on the covariance matrix at initialization $\overline{\mathcal{K}}_0$ and on the number of training examples n , and $d_W^{(s)}$ denotes the truncated Wasserstein distance of order 1 defined in (2.57).

Let us point out that $\gamma_m^{\overline{N}, n, \overline{\mathcal{K}}_0}$ depends on the previous constant $\alpha_m^{\overline{N}, \overline{\mathcal{K}}_0}$, and thus, its dependence on \overline{N} is of order \overline{N}^2 . We stress that (1.3) is valid for any quantum neural network of finite width, unlike the results of [8] which require to build a sequence of quantum neural networks with diverging width. Furthermore, [8] proved the convergence to a Gaussian process at fixed training time, and such result cannot be directly extended to the limit $t \rightarrow \infty$, unless such limit is taken after the limit of infinite width. On the contrary, the constants $\{\gamma_m^{\overline{N}, n, \overline{\mathcal{K}}_0}\}$ do not depend on the training time, and therefore, the bound (1.3) is uniform in time and holds for finite width even in the limit $t \rightarrow \infty$. Lastly, let us notice that (1.2) and (1.3) imply the convergence results of [8].

As in [8], the proof of Theorem 1.2 is based on the following upper bound to the maximum variation of a parameter during training (see Theorem 7.1 for the formal statement):

Theorem 1.3 (Lazy training, informal statement). *With the same notation of Theorem 1.1 and Theorem 1.2, let us denote by $x \mapsto f^{\text{lin}}(\Theta, x)$ the first-order Taylor approximation of $x \mapsto f(\Theta, x)$ with respect to the parameters Θ expanded around their initialization values. Let $x \mapsto f(\Theta_t^{\text{lin}}, x)$ be the model obtained by randomly initializing Θ and training $x \mapsto f^{\text{lin}}(\Theta, x)$ via gradient flow for time t . Then, for any $0 < \delta < 1$ there exist positive numbers $\{\eta_{m, n, \delta}\}$ given by (7.10) depending on m, n, δ , and further constants, such that with probability at least $1 - \delta$, one gets that*

$$\sup_{\substack{x \in \mathcal{X} \\ t \geq 0}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)| \leq \eta_{m, n, \delta}, \quad (1.4)$$

where \mathcal{X} is the set of all the possible inputs.

We stress that, for any fixed n and δ , $\eta_{m, n, \delta} \rightarrow 0$ for $m \rightarrow \infty$ if barren plateaus do not arise. Therefore, (1.4) proves that if m is large enough, each parameter remains with high probability close to its initialization value and that the training happens in the lazy regime. In the formal statement of Theorem 1.3, we provide two main improvements with respect to [8, Theorems 4.10, 4.12]. First, [8, Theorems 4.10, 4.12] require to build a sequence of networks with increasing width and hold only if the width is larger than some value that has not been computed. Instead, we want to describe the training dynamics of

a network with finite width; therefore, we have replaced the hypotheses about the asymptotic behavior of the network with the quantitative requirement at finite size (7.6). Second, while the bounds of [8, Theorems 4.10, 4.12] contain constants that have not been computed and that depend on the failure probability of the training at finite width, we have provided explicit bounds.

The article is organized as follows. In Sect. 2, we introduce some preliminaries, notations, and our hypotheses. In Sect. 3, we recall some known facts about Gaussian processes. In Sect. 4, we recall some facts about the Stein method and its relationship to the Wasserstein distance of order 1. In Sect. 5, we state our main results. In Sect. 6, we prove Theorem Theorem1.1. In Sect. 7, we prove Theorem 1.3, which is then used to prove Theorem 1.2. In the same section, we prove Theorem 1.2. We conclude in Sect. 8. Lastly, for the sake of completeness, we include appendices to discuss the regularity of the solutions of the Stein’s equation (Appendix A) and some theorems proved in [8] regarding the convergence of the function generated by a quantum circuit to a Gaussian process (Appendix B).

2. Preliminaries

Let us start by introducing the notation of the present work (Tables 1 and 2).

TABLE 1. Notation concerning the general properties of the circuit and the Wasserstein distance

Symbol	Description	Introduced in
m	Number of qubits in the parameterized quantum circuit	Section 2.2
L	Number of layers in the parameterized quantum circuit	Definition 2.1
Θ	Vector of the parameters of the quantum circuit	Equation 2.3
\mathcal{P}	Denotes the parameter space, so that $\Theta \in \mathcal{P}$. Here $\mathcal{P} = [0, \pi]^{Lm}$	Section 2.1
$ \Theta $	Number of parameters $ \Theta := \dim \mathcal{P} = Lm$	Section 2.2
$U(\Theta, x)$	Parameterized quantum circuit (unitary operator)	<i>ibid.</i>
$f(\Theta, x)$	Function generated by the quantum neural network	Section 2.4
$N(m)$	Normalization factor of the model	<i>ibid.</i>
\mathcal{M}_i	(Extended) future light cone of the parameter i	Definition 2.13
\mathcal{N}_k	(Extended) past light cone of the observable k	Definition 2.7
$ \mathcal{M} $	Maximal cardinality of a future light cone in the circuit	Equation 2.14
$ \mathcal{N} $	Maximal cardinality of a past light cone in the circuit	<i>ibid.</i>
\mathcal{P}_i	Set of indices of the observables depending on the observable i	Equation 2.22
$\tilde{\mathcal{P}}_i$	Set representing the union of the sets \mathcal{P}_j for those j in \mathcal{P}_i	Equation 2.26
d_W	Wasserstein distance of order 1	Section 2.8
$d_W^{(s)}$	Truncated Wasserstein distance of order 1	<i>ibid.</i>
ω	Continuity modulus of the Hessian of the solution of Stein’s equation	Equation A.3

TABLE 2. Notation concerning the training of the circuit

Symbol	Description	Introduced in
\mathcal{X}	The feature space	Section 2.1
\bar{N}	Cardinality of the feature space	<i>ibid.</i>
\bar{X}	Vector containing all the possible inputs in \mathcal{X}	<i>ibid.</i>
x	A generic input belonging to \mathcal{X}	<i>ibid.</i>
\mathcal{Y}	The output space	<i>ibid.</i>
y	A generic output belonging to \mathcal{Y}	<i>ibid.</i>
\mathcal{D}	Training set, whose elements are denoted by $(x^{(i)}, y^{(i)})$ for $i = 1, \dots, n$	<i>ibid.</i>
n	Number of training samples (i.e., cardinality of \mathcal{D})	<i>ibid.</i>
X	Vector containing the inputs of the training set	Section 2.5
Y	Vector containing the outputs of the training set corresponding to X	<i>ibid.</i>
$f^{\text{lin}}(\Theta, x)$	Linearized model	Section 4.1
$\hat{K}_{\Theta}(x, x')$	Empirical neural tangent kernel	Definition 2.5
$K(x, x')$	Analytic neural tangent kernel	Definition 2.6
\mathcal{K}_0	Covariance matrix at initialization of size $\bar{N} \times \bar{N}$	Theorem 5.1
$\lambda_{\min}^{\mathcal{K}}$	Smallest eigenvalue of $K(X, X^T)$	Assumption A 3
t	Continuous or discrete training time	Section 2.5
$F(t)$	Vector containing the model function evaluated in the inputs of the dataset, i.e., $F(t) = f(\Theta_t, X)$	Equation 2.37
\mathcal{K}_t	Covariance matrix of $f(\Theta_t, \bar{X})$	Theorem 5.2
$F^{\text{lin}}(t)$	Vector containing the linearized model function evaluated in the inputs of the dataset, i.e., $F^{\text{lin}}(t) = f^{\text{lin}}(\Theta_t^{\text{lin}}, X)$	Equation 4.12
η	Learning rate, which enters the gradient flow equation and is a function of m	Equation 2.32
$\mathcal{L}(\Theta)$	Cost function for the original model according to the training set	Equation 2.2
Θ_t	Parameter vector evolving via gradient flow according to \mathcal{L}	Equation 2.32
$\mathcal{L}^{\text{lin}}(\Theta)$	Cost function for the linearized model	Section 4.1
Θ_t^{lin}	Parameter vector evolving via gradient flow according to \mathcal{L}^{lin}	<i>ibid.</i>

2.1. Training Data

Let $\mathcal{X} \subset \mathbb{R}^d$ be the feature space, i.e., the set of all the possible inputs, which we assume to have finite cardinality $N = |\mathcal{X}|$. We will often use the notation \bar{X} to represent the vector having as entries the elements of the input space \mathcal{X} . Let \mathbb{R} be the output space. Let

$$\mathcal{D} := \left\{ (x^{(i)}, y^{(i)}) : i = 1, \dots, n \right\} \subset \mathcal{X} \times \mathcal{Y} \quad (2.1)$$

be the training set. We set $n = |\mathcal{D}|$ to be the cardinality of \mathcal{D} . We let \mathcal{P} be the parameter space, and let $\Theta \in \mathcal{P}$ be the vector of the parameters. Let $f : \mathcal{P} \times \mathcal{X} \rightarrow \mathbb{R}$ be a generic parametric function. As a cost function, we consider the mean squared error on the training set \mathcal{D} of cardinality n

$$\mathcal{L}(\Theta) := \sum_{i=1}^n \left(f(\Theta, x^{(i)}) - y^{(i)} \right)^2. \quad (2.2)$$

2.2. Quantum Neural Networks

Let \mathbb{C}^2 be the Hilbert space of a single qubit. In what follows, we denote by $m \in \mathbb{N}$ the number of qubits of the quantum neural network. Hence, the Hilbert space of the system is $\mathcal{H} = (\mathbb{C}^2)^{\otimes m}$, and its dimension denoted as $\dim \mathcal{H} = 2^m$. Following the notations of [8], we recall what a ‘‘layer’’ is.

Definition 2.1. A layer is a unitary operation $U(\Theta, x) \in \mathfrak{L}(\mathcal{H})$ resulting from:

1. the application on each qubit of a different parametrized single-qubit gate $W_i(\Theta) \in \mathfrak{L}(\mathbb{C}^2)$; each parametrized gate depends on a single parameter θ_i , which is different for each gate,

followed by

2. a set of one-qubit and two-qubit gates acting on disjoint qubits, that is, each qubit can be acted at most one gate; each gate may depend only on the input x ; the resulting unitary operation will be called $V \in \mathfrak{L}(\mathcal{H})$.

In what follows, a quantum circuit is a combination of parameterized layers $U_\ell(\Theta, x)$, $\ell \in \mathbb{N}$. Next, we let $L \in \mathbb{N}$ be the number of layers in a quantum circuit, which may depend on the number of qubits m .

In the next, we consider $\theta_1, \dots, \theta_{Lm}$ be the parameters of a quantum circuit, so that Θ will be the vector

$$\Theta := \begin{pmatrix} \theta_1 \\ \theta_2 \\ \vdots \\ \theta_{Lm} \end{pmatrix}, \quad (2.3)$$

of dimension $\dim \Theta := |\Theta| = Lm$. We now recall a convenient notation for the indices of the parameters used in [8].

Definition 2.2. Each parameter index $i \in \{1, \dots, Lm\}$ can be expressed in the form $i = m(\ell - 1) + k$ for some $\ell \in \{1, \dots, L\}$, and $k \in \{1, \dots, m\}$. Here, k refers to the qubit involved in the single-qubit gate parametrized by θ_i , while

ℓ refers to the layer in which such gate acts. The following compact notation, which we call layer-qubit representation of the parameter index i , simplifies the above form:

$$i = [\ell m] \equiv m(\ell - 1) + k. \quad (2.4)$$

Therefore, a layer $U_\ell(\Theta, x)$ can be written as

$$\begin{aligned} U_\ell(\Theta, x) &:= V_\ell(x) (W_{[\ell 1]} \otimes \cdots \otimes W_{[\ell m]}) (\Theta) \\ &= V_\ell(x) W_\ell(\Theta), \end{aligned} \quad (2.5)$$

where we have set $W_\ell(\Theta) := (W_{[\ell 1]} \otimes \cdots \otimes W_{[\ell m]}) (\Theta)$. The result of the circuit on an initial state $|\psi_0\rangle$ is described by the unitary operation

$$U(\Theta, x) := U_L(\Theta, x) \cdots U_1(\Theta, x); \quad |\psi_{\text{out}}\rangle := U(\Theta, x)|\psi_0\rangle. \quad (2.6)$$

2.3. Light Cones

The architecture of the network generates a causal structure where the probability distribution of the outcome of the measurement of each output qubit can depend only on some of the parameters, and each parameter can influence only some output qubits. Such causal structure is formalized by the notion of light cones:

Definition 2.3. (*Light cones*) For any $i \in \{1, \dots, |\Theta|\}$, we define the future light cone \mathcal{L}_i^f of the parameter θ_i as the subset

$$\mathcal{L}_i^f := \{k \in \{1, \dots, m\} : f_k(\Theta, x) \text{ depends on } \theta_i\}. \quad (2.7)$$

Analogously, we define the past light cone \mathcal{L}_k^p of the qubit k as the subset

$$\mathcal{L}_k^p := \{i \in \{1, \dots, |\Theta|\} : f_k(\Theta, x) \text{ depends on } \theta_i\}. \quad (2.8)$$

Both sets \mathcal{L}_i^f , and \mathcal{L}_k^p are useful for tracking the dependence of observables on the parameters. In general, it is difficult to provide an explicit representation of them. For this reason, we now introduce another family of sets that can help us to explicitly track the dependence on the parameters. For any quantum circuit U , we define the following sets. For each layer ℓ , and qubit k , we set

$$\begin{aligned} \mathcal{J}_{\ell, k} &:= \{k' \in \{1, \dots, m\} : \text{the qubit } k \text{ interacts with the qubit } k' \text{ in the layer } \ell\} \\ &\cup \{k\}. \end{aligned} \quad (2.9)$$

We now set,

$$\mathcal{J}_k^\ell := \begin{cases} \mathcal{J}_{L, k} & \text{if } \ell = L, \\ \bigcup_{k' \in \mathcal{J}_k^{\ell+1}} \mathcal{J}_{\ell, k'} & \text{if } \ell < L. \end{cases} \quad (2.10)$$

In particular, \mathcal{J}_k^1 is the set of qubits in the past light cone of the observable k , i.e., the qubits involved in the computation of its expectation value.

Furthermore, we set

$$\mathcal{N}_k^\ell := \bigcup_{k' \in \mathcal{J}_k^\ell} \{[k'k']\}. \quad (2.11)$$

Definition 2.4 (*Extended light cones*). Let us fix a quantum circuit U . Given any qubit index $k \in \{1, \dots, m\}$, we define the extended past light cone \mathcal{N}_k as the subset of the parameter indices $\{1, \dots, |\Theta|\}$ given by

$$\mathcal{N}_k := \bigcup_{\ell=1}^L \mathcal{N}_k^\ell \quad (2.12)$$

Similarly, we define the extended future light cone of a parameter index $i \in \{1, \dots, |\Theta|\}$, as

$$\mathcal{M}_i := \{k \in \{1, \dots, m\} : i \in \mathcal{N}_k\}. \quad (2.13)$$

In the next, we set

$$|\mathcal{M}| := \max_i |\mathcal{M}_i|; \quad |\mathcal{N}| := \max_k |\mathcal{N}_k| \quad (2.14)$$

the maximal cardinalities of the extended light cones.

Remark 2.1. We notice that $\mathcal{L}_k^p \subset \mathcal{N}_k$, and $\mathcal{L}_i^f \subset \mathcal{M}_i$, and thus from now on, we can only consider the extended light cones.

2.4. Assumptions on Architecture and Initialization

We assume the following:

A 1. We consider an observable \mathcal{O} given by the sum of single-qubit observables \mathcal{O}_k :

$$\begin{aligned} \mathcal{O} &= \sum_{k=1}^m \mathcal{O}_k = \mathcal{O}_1 \otimes \mathbb{1}_2 \otimes \dots \otimes \mathbb{1}_m + \mathbb{1}_1 \otimes \mathcal{O}_2 \otimes \dots \otimes \mathbb{1}_m \\ &\quad + \mathbb{1}_1 \otimes \mathbb{1}_2 \otimes \dots \otimes \mathcal{O}_m, \end{aligned} \quad (2.15)$$

where each \mathcal{O}_k is traceless and has the spectrum contained in the interval $[-1, 1]$. We further assume that the parametric one-qubit gates of the circuit $W_i(\theta_i)$ can be written as time evolutions generated by hermitian hamiltonians \mathcal{G}_i with spectrum in $\{-1, 1\}$, i.e.,

$$W_i(\theta_i) = e^{-i\mathcal{G}_i\theta_i}, \quad \mathcal{G}_i = \mathcal{G}_i^\dagger = \mathcal{G}_i^{-1}. \quad (2.16)$$

We notice that $\theta_i \mapsto W_i(\theta_i)$ is periodic with period π up to an irrelevant multiplicative constant, so we fix the parameter space to be $\mathcal{P} = [0, \pi]^{Lm}$, and thus $|\Theta| = \dim \mathcal{P} = Lm$. The function generated by the network is then

$$\begin{aligned} f(\Theta, x) &:= \frac{1}{N(m)} \langle 0^m | U^\dagger(\Theta, x) \mathcal{O} U(\Theta, x) | 0^m \rangle \\ &= \frac{1}{N(m)} \sum_{k=1}^m f_k(\Theta, x) \end{aligned} \quad (2.17)$$

where

$$f_k(\Theta, x) := \langle 0^m | U^\dagger(\Theta, x) \mathcal{O}_k U(\Theta, x) | 0^m \rangle. \quad (2.18)$$

A 2. We assume that each parameter θ_i is initialized independently by sampling it from the uniform distribution on $[0, \pi]$ and that the final layer is chosen so that $\mathbb{E}_{\Theta} f_k(\Theta, x) = 0$ for any $k = 1, \dots, m$ and any $x \in \mathcal{X}$. For any $x, x' \in \mathcal{X}$, let

$$\mathcal{K}_0(x, x') = \mathbb{E}(f(\Theta, x) f(\Theta, x')). \quad (2.19)$$

We assume that $\mathcal{K}_0(x, x) > 0$ for any $x \in \mathcal{X}$ and that $N(m)$ is chosen such that

$$\max_{x \in \mathcal{X}} \mathcal{K}_0(x, x) = 1. \quad (2.20)$$

Finally, we assume that $N(m) \geq 1$ (our bounds would anyway become trivial if $N(m) \leq 1$).

In what follows, let us set

$$V := \{1, \dots, m\}. \quad (2.21)$$

For each $k \in V$, we define

$$\mathcal{P}_k := \{k' \in V : f_{k'}(\Theta, x) \text{ is not independent from } f_k(\Theta, x)\}. \quad (2.22)$$

This set is crucial since it takes track of the number of random variables $f_{k'}(\Theta, x)$ that have correlation with $f_k(\Theta, x)$. Let $\mathcal{G} = (V, E)$ be the graph with vertices V , and edges E defined as follows. We say

$$(k, k') \in E \text{ if and only if } k' \in \mathcal{P}_k. \quad (2.23)$$

Furthermore, we define the *maximal degree* D of \mathcal{G} as the maximum number of edges containing any fixed vertex as

$$D := \max_{k \in V} \deg k = \max_{k \in V} |\{k' \in V : (k, k') \in E\}| = \max_{k \in V} |\mathcal{P}_k|. \quad (2.24)$$

Let us notice that according to the definition of \mathcal{P}_i , we have that $j \in \mathcal{P}_i$ if and only if $i \in \mathcal{P}_j$. Let dist be the distance on \mathcal{G} given by the length of the shortest path, such that for any $i \in V$ we have

$$\mathcal{P}_i = \{j \in V : \text{dist}(i, j) \leq 1\}. \quad (2.25)$$

In what follows, we set

$$\tilde{\mathcal{P}}_i := \bigcup_{j \in \mathcal{P}_i} \mathcal{P}_j = \{j \in V : \text{dist}(i, j) \leq 2\}, \quad (2.26)$$

and

$$\tilde{D} := \max_{1 \leq i \leq m} \left| \{(i, j) : \tilde{\mathcal{P}}_i \cap \tilde{\mathcal{P}}_j \neq \emptyset\} \right| = \max_{1 \leq i \leq m} |\{(i, j) : \text{dist}(i, j) \leq 4\}|. \quad (2.27)$$

Lemma 2.1 [8, Lemma 2.9]. *For any $k \in \{1, \dots, m\}$, let \mathcal{P}_k be defined as in (2.22). Then*

$$|\mathcal{P}_k| \leq |\mathcal{M}| |\mathcal{N}|. \quad (2.28)$$

In particular,

$$D \leq |\mathcal{M}| |\mathcal{N}|. \quad (2.29)$$

Lemma 2.2. *We have*

$$\tilde{D} \leq |\mathcal{M}|^4 |\mathcal{N}|^4. \quad (2.30)$$

Proof. Note that by definition of D , we have that

$$\tilde{D} \leq D^4,$$

and so that by Lemma 2.1, we are done. \square

Remark 2.2. For any $i = 1, \dots, m$, let

$$Z_i := \sum_{k \in \mathcal{P}_i} \frac{f_k(\Theta, \bar{X})}{N(m)}. \quad (2.31)$$

Then if $\mathcal{P}_k \cap \tilde{\mathcal{P}}_j = \emptyset$, we have that Z_j is independent of Z_k . This fact will be significant when proving Theorem 5.1.

2.5. The Neural Tangent Kernel

Before presenting our main result, let us review some relevant facts about the quantum neural tangent kernel as presented in [8]. We are interested in the analysis of the minimization of the cost function (2.2) via gradient flow:

$$\frac{d\Theta_t}{dt} = -\eta \nabla_{\Theta} \mathcal{L}(\Theta_t) \quad (2.32)$$

where $\eta > 0$ is the learning rate that can be reabsorbed by rescaling the training time, and the initial value of Θ is given by the random sampling of A 2. Notice that

$$\frac{d}{dt} \mathcal{L}(\Theta_t) = \frac{d\Theta_t}{dt} \cdot \nabla_{\Theta} \mathcal{L}(\Theta_t) = -\eta \|\nabla_{\Theta} \mathcal{L}(\Theta_t)\|_2^2 \leq 0. \quad (2.33)$$

We stress that in general, the loss function $\mathcal{L}(\Theta)$ is not convex; hence, gradient flow is not guaranteed to converge to a global minimum. Given a training set $\mathcal{D} = \{(x^{(i)}, y^{(i)})\}_{i=1, \dots, n} \subseteq \mathcal{X} \times \mathcal{Y}$, we will call $n = |\mathcal{D}|$ the number of examples and we will represent it in a vectorized form as follows

$$X = \begin{pmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(n)} \end{pmatrix}, \quad Y = \begin{pmatrix} y^{(1)} \\ y^{(2)} \\ \vdots \\ y^{(n)} \end{pmatrix}. \quad (2.34)$$

Given any function $g : \mathbb{R} \rightarrow \mathbb{R}$, we will often use the following notation:

$$g(X) := \begin{pmatrix} g(x^{(1)}) \\ g(x^{(2)}) \\ \vdots \\ g(x^{(n)}) \end{pmatrix}, \quad g(X^T) := (g(x^{(1)}) \ g(x^{(2)}) \ \cdots \ g(x^{(n)})) \quad (2.35)$$

Similarly, for any bivariate function $K : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ we will write $K(X, X^T)$ to indicate the $n \times n$ matrix with entries $(K(X, X^T))_{ij} := K(x^{(i)}, x^{(j)})$ for $1 \leq i, j \leq n$. In what follows, we denote by

$$X = (x^{(1)}, \dots, x^{(n)})^T \quad (2.36)$$

the vector of the training inputs and by \bar{X} the vector of all the possible inputs in \mathcal{X} . We set

$$F(t) := \begin{pmatrix} f(\Theta_t, x^{(1)}) \\ f(\Theta_t, x^{(2)}) \\ \vdots \\ f(\Theta_t, x^{(n)}) \end{pmatrix} = f(\Theta_t, X). \quad (2.37)$$

From the gradient flow equation (2.32) and the chain rule, the evolution equations for the parameters and the model function can be written as

$$\begin{cases} \frac{d\Theta_t}{dt} = -\eta \nabla_{\Theta} f(\Theta_t, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t), \\ \frac{d}{dt} f(\Theta_t, x) = -\eta (\nabla_{\Theta} f(\Theta_t, x))^T \nabla_{\Theta} f(\Theta_t, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t), \end{cases} \quad (2.38)$$

where $\nabla_{\Theta} f(\Theta_t, X^T)$ denotes the gradient of $f(\Theta_t, X^T)$ with respect to Θ while $\nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t)$ indicates the gradient of the cost function \mathcal{L} with respect to $f(\Theta_t, X)$. Recall that T is the transposition operator.

Definition 2.5. We define the empirical neural tangent kernel, in short empirical neural tangent kernel, as

$$\hat{K}_{\Theta}(x, x') := (\nabla_{\Theta} f(\Theta_t, x))^T \nabla_{\Theta} f(\Theta_t, x'). \quad (2.39)$$

(2.38) can be written as

$$\begin{cases} \frac{d\Theta_t}{dt} = -\eta \nabla_{\Theta} f(\Theta_t, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t), \\ \frac{d}{dt} f(\Theta_t, X) = -\eta \hat{K}_{\Theta_t}(x, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t). \end{cases} \quad (2.40)$$

We mention that the quantum neural tangent kernel has been studied in several works inspired by the success of the neural tangent kernels (NTKs) in describing the dynamics of classical neural networks and the convergence of the training processes [37–42]. In [37], a quantum NTK for deep quantum circuits is introduced considering a first-order expansion with respect to the parameters since, during training, for a deep enough circuit the parameters do not move much from initialization. In [38], the authors propose the use of efficient quantum algorithms to construct kernel functions based on properties that are difficult to compute by classical computations. In [39], the authors consider a class of hybrid neural networks composed of a quantum data encoder followed by a classical network where the quantum part is randomly initialized according to unitary 2-designs, and the classical part is also randomly initialized according to Gaussian distributions, then they apply the NTK theory to construct an effective quantum kernel, which is in general nontrivial to design. In [40], neural tangent kernels for variational circuits are defined, then the authors derive dynamical equations for their associated loss function in optimization and learning tasks, they analytically solve the dynamics in the lazy training regime. In [41], the authors study the connections between the expressibility and value concentration of quantum tangent kernel models. In particular, for global loss functions, they rigorously prove that high expressibility of quantum encodings can lead to exponential concentration of

quantum tangent kernel values to zero. In [42], there is an analysis on the effects of hyperparameter choice on a quantum kernel model performance and the generalization gap between classical and quantum kernels.

2.6. The Assumptions for the Neural Tangent Kernel

In what follows, we introduce some useful assumptions to treat the behavior of the neural tangent kernel. Before, let us recall the definition of the analytic neural tangent kernel.

Definition 2.6. We define the analytic neural tangent kernel as the expectation of the empirical neural tangent kernel:

$$K(x, x') := \mathbb{E}(\hat{K}_\Theta(x, x')). \quad (2.41)$$

A 3. We suppose that the neural tangent kernel restricted to the training inputs $K = K(X, X^T)$ is invertible. We denote with λ_{\max}^K and λ_{\min}^K its maximum and minimum eigenvalue, respectively.

A 4. Let us furthermore assume that

$$\lambda_{\min}^K \geq \frac{96n\sqrt{Lm}|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \sqrt{\log(2n^2\sqrt{N(m)})}, \quad (2.42)$$

$$\frac{Lm|\mathcal{M}|^2|\mathcal{N}|^2}{N^3(m)} \leq \frac{1}{256 \log 2} \frac{1}{N^2}. \quad (2.43)$$

Eq. (2.42) will be needed to enforce that the empirical neural tangent kernel is strictly positive with high probability, while eq. (2.43) will be needed to enforce that the empirical neural tangent kernel is close to the analytic neural tangent kernel with high probability.

2.7. The Expected Behavior of $N(m)$ in Circuits with Logarithmic Depth

In the quantitative bounds that we are going to provide in the next sections, the dependence on the number of qubits will appear as

$$\left(\frac{L^\alpha m^\beta |\mathcal{M}|^\gamma |\mathcal{N}|^\delta (\log N(m))^\sigma}{N(m)} \right)^\nu \quad (2.44)$$

for some $\alpha, \beta, \gamma, \delta, \sigma, \nu > 0$ according to the precise statement. As discussed in [8], in some cases it is possible to estimate $N(m)$, $|\mathcal{M}|$ and $|\mathcal{N}|$ so that the asymptotical behavior of the bounds for wide circuits can be studied. In the example provided by E. Abedi & al. in [30] we have $L = O(1)$ and

$$N(m) = \sqrt{m}, \quad |\mathcal{M}| = O(L), \quad |\mathcal{N}| = O(L^2). \quad (2.45)$$

As we will see, in general $\beta < \frac{1}{2}$, this kind of circuit always satisfies all our theorems. However, since $\dim \mathcal{H}_{\text{loc}} = 2^{O(L)} = O(1)$, such circuit is classically simulable, so we do not expect to achieve quantum advantages with such architecture.

A suitable class of circuits for investigating quantum advantages is the one studied by J. C. Napp in [43], where the following bound is provided

$$N(m) \geq \frac{\sqrt{m}}{2^{CL}} \quad \text{for some fixed } C. \quad (2.46)$$

Furthermore, we will prove that (see (6.18))

$$N(m) \leq \sqrt{m|\mathcal{M}||\mathcal{N}|}. \quad (2.47)$$

In a generic setting the number of qubits in the past light cone of any observable O_k can grow as

$$|\mathcal{J}_k^1| = O(2^L). \quad (2.48)$$

Under the hypothesis of geometrical locality (i.e., each qubit can interact only with the nearest neighbor qubits), a d -dimensional lattice of qubits has

$$|\mathcal{J}_k^1| = O(L^d). \quad (2.49)$$

Let us assume to choose $L = \epsilon \log_2 m$. Without assumptions on the geometrical locality, $|\mathcal{J}_k^1| = O(m^\epsilon)$, which is an upper bound, so any growth $|\mathcal{J}_k^1| = \Theta(m^{\epsilon'})$ with $\epsilon' \leq \epsilon$ can be achieved by an appropriate choice of the interactions. In this case, the dependence on the number of qubits will provide a prefactor (2.44) asymptotically vanishing provided that ϵ, ϵ' are small enough. Indeed, up to multiplicative constants, (2.44) will behave as

$$\lim_{m \rightarrow \infty} \left(\frac{(\log_2 m)^\alpha m^{\beta + (\gamma + \delta)\epsilon'}}{m^{1/2 - C\epsilon}} (\log m)^\sigma \right)^\nu = 0 \quad (2.50)$$

since in all our bounds we have $\beta < 1/2$. However, the local Hilbert spaces have a quasi-exponential dimension $2^{m^{\epsilon'}}$. In the geometrically local setting, we can choose $|\mathcal{J}_k^1| = \Theta((\epsilon \log_2 m)^d)$; if $d \geq 2$, then the local Hilbert spaces have a super-polynomial dimension, since $2^{(\epsilon \log_2 m)^d} = m^{\epsilon^d \log_2^{d-1} m}$. Besides, (2.44) is asymptotically vanishing for ϵ small enough.

The fact that there are circuits such that the quantities of the form (2.44) are asymptotically vanishing—as in (2.50)—is fundamental to understand why both conditions of Assumption A 4 are reasonable. Since in (2.43) a quantity of the form (2.44) appears, we immediately see that it can be satisfied by the circuits described above when the number of qubits is large enough. Similarly, for any fixed $\lambda_0 > 0$, if we restrict the class of circuits considered above to the ones having $\lambda_{\min}^K \geq \lambda_0$, then (2.42) holds provided that m is large enough.

We should mention that the exponential decrease of $N(m)$ on the number of layers is a manifestation of the phenomenon of *barren plateaus* [43]. We will not further investigate the particular architectures which satisfy our theorems, but our statements will be motivated by the existence of examples such as the ones we have just presented.

2.8. The Wasserstein Distance of Order 1

In this subsection, we introduce the distance that we employ to quantify the closeness between the probability distribution of the functions generated by quantum neural networks and the associated Gaussian processes (for a precise definition of what a Gaussian process is, see Section 3): the Wasserstein distance of order 1, also called Monge–Kantorovich distance or earth mover’s distance [44–46]. The Wasserstein distance of order 1 has been introduced in

the context of optimal mass transport as the minimum cost required to transport a mass distribution into another [47, 48]. It admits applications in several areas of mathematics [25] and in the general scenario of transporting resources in the cheapest way [49, 50]. As a distance among probability distributions, the Wasserstein distance finds natural applications in statistics [51] and machine learning [52–55] and it is also be extended to the quantum realm and considered in the context of quantum machine learning [24, 56].

Let (Ω, \mathbb{P}) be a probability space. Let us denote by $\mathcal{C} := \mathcal{C}(\mathbb{R}^d)$ a collection of Borel-measurable real-valued functions on \mathbb{R}^d . We say that the class \mathcal{C} is separating if the following property holds: any two \mathbb{R}^d -valued random variables F, G satisfying $h(F), h(G) \in L^1(\Omega, \mathbb{P})$ and $\mathbb{E}(h(F)) = \mathbb{E}(h(G))$ for any $h \in \mathcal{C}$, have necessarily the same law. The distance between the laws of F and G induced by \mathcal{C} is

$$d_{\mathcal{C}}(F, G) := \sup \{ |\mathbb{E}(h(F)) - \mathbb{E}(h(G))| : h \in \mathcal{C} \}. \quad (2.51)$$

The Wasserstein distance of order 1 between the laws of F and G , denoted by $d_W(F, G)$, is obtained from (2.51) by taking \mathcal{C} to be the set of all functions $h : \mathbb{R}^d \rightarrow \mathbb{R}$ such that $\|h\|_{\text{Lip}} \leq 1$, where

$$\|h\|_{\text{Lip}} := \sup_{\substack{x \neq y \\ x, y \in \mathbb{R}^d}} \frac{|h(x) - h(y)|}{\|x - y\|_2}, \quad (2.52)$$

and $\|\cdot\|_2$ denotes the usual Euclidean norm in \mathbb{R}^d . The Wasserstein distance of order 1 admits the following equivalent formulation:

$$d_W(F, G) := \inf \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|_2 d\gamma(x, y) : \gamma \in \Gamma(\text{law}(F), \text{law}(G)) \right\}, \quad (2.53)$$

where $\Gamma(\text{law}(F), \text{law}(G))$ denotes the set of all probability measures supported on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals $\text{law}(F)$, and $\text{law}(G)$. In (2.53), we defined the Wasserstein distance of order 1 between the law of two random variables, F, G . In general, this definition applies for any probability measures μ, ν on \mathbb{R}^d . Furthermore, the distance (2.53) can also be expressed in terms of the so-called Monge–Kantorovich formulation [57]. Let $T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be a Borel map, and denote by $T_{\#}\mu$ the push forward of μ through T defined by

$$T_{\#}\mu(E) := \mu(T^{-1}(E)) \quad \forall E \subset \mathbb{R}^d, \text{ Borel}. \quad (2.54)$$

The Monge version of the transport problem is the following: Let μ, ν be probability measures on \mathbb{R}^d absolutely continuous with respect to the d -dimensional Lebesgue measure. Minimize

$$T \mapsto \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - T(x)\|_2 d\mu(x), \quad (2.55)$$

among all the transport maps T from μ to ν , that is, all maps T such that $T_{\#}\mu = \nu$. It is well-known that (2.53) is equivalently written in terms of the minimization problem (2.55). That is, there exists an optimal transport map

$T : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$d_W(\mu, \nu) = \int_{\mathbb{R}^d} \|x - T(x)\|_2 d\mu(x). \tag{2.56}$$

In what follows, given $s > 0$, we will also need the following truncated version of the Wasserstein distance of order 1, denoted as $d_W^{(s)}(F, G)$:

$$d_W^{(s)}(F, G) := \inf \left\{ \int_{\mathbb{R}^d \times \mathbb{R}^d} d^{(s)}(x, y) d\gamma(x, y) : \gamma \in \Gamma(\text{law}(F), \text{law}(G)) \right\}, \tag{2.57}$$

where

$$d^{(s)}(x, y) := \|x - y\|_2 \wedge s, \quad \text{for } x, y \in \mathbb{R}^d. \tag{2.58}$$

Let us denote by $\mathcal{C}_{\text{Lip}(1),s}$ the class of 1-Lipschitz functions with respect to $d^{(s)}$, and by $\mathcal{C}_{\text{Lip}(1)}$ the class of 1-Lipschitz functions with respect to $\|\cdot\|_2$. Notice that $\mathcal{C}_{\text{Lip}(1),s} \subset \mathcal{C}_{\text{Lip}(1)}$, and thus by the dual characterization of the 1-Wasserstein distance, we get that

$$d_W^{(s)}(F, G) \leq d_W(F, G). \tag{2.59}$$

3. Some Relevant Facts About Gaussian Processes

In this part, we collect some of the properties needed in the study of the convergence of a sequence of random variables $\{X_n\}_{n \in \mathbb{N}}$ to a Gaussian processes. Let us start by recalling what a stochastic process is.

Definition 3.1. A stochastic process is a collection of random variables $(X_\alpha)_{\alpha \in A}$ defined on a same probability space, where A is a set of indices.

It is customary to write $X \sim Y$ when the random variables X , and Y have the same probability distribution. We say that a random variable X is Gaussian, or distributed according to a Gaussian probability distribution $N(\mu, \sigma^2)$ with $\mu \in \mathbb{R}$, $\sigma \neq 0$, if its probability density is given by

$$p(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}, \quad x \in \mathbb{R}.$$

In the vectorial case, given $\mu \in \mathbb{R}^d$, and $\mathcal{K} \in \mathbb{R}^{d \times d}$ is positive definite matrix, we write $X \sim N(\mu, \mathcal{K})$ if X is a random vector with probability density function

$$p(x) = \frac{1}{(2\pi)^{\frac{N}{2}}} \frac{1}{\sqrt{\det(\mathcal{K})}} e^{-\frac{1}{2}(x-\mu)^T \mathcal{K}^{-1}(x-\mu)}, \quad x \in \mathbb{R}^d.$$

In what follows, we recall what a Gaussian process is.

Definition 3.2. A stochastic process $\{X_\alpha\}_{\alpha \in A}$ is called Gaussian process (GP) if for every subset $F \subset A$ there exist a vector $\mu_F \in \mathbb{R}^{|F|}$, and a positive semidefinite covariance matrix \mathcal{K}_F such that

$$\{X_\alpha\}_{\alpha \in F} \sim N(\mu_F, \mathcal{K}_F). \tag{3.1}$$

In particular, if we define the mean function $\mu : A \rightarrow \mathbb{R}$, and the covariance function $\mathcal{K} : A \times A \rightarrow \mathbb{R}$ as

$$\mu(\alpha) := \mathbb{E}(X_\alpha) \tag{3.2}$$

$$\mathcal{K}(\alpha, \alpha') := \mathbb{E}((X_\alpha - \mu(\alpha))(X_{\alpha'} - \mu(\alpha'))), \tag{3.3}$$

then, we can consider μ_F , and \mathcal{K}_F as the vector, and the matrix given by the evaluations of these functions, that is, we may set

$$\mu_F = \mu(F), \quad \mathcal{K}_F = \mathcal{K}(F, F^T). \tag{3.4}$$

An important fact about the Gaussian processes is that it can be characterized by the mean and the covariance function. Then, we can refer to a Gaussian process $\{X_\alpha\}_{\alpha \in A}$ by writing

$$\{X_\alpha\}_{\alpha \in A} \sim \text{GP}(\mu, \mathcal{K}). \tag{3.5}$$

In what follows, we aim to recall some about the convergence in distribution of random variables.

Definition 3.3. We say that a sequence of real random variables $\{X_n\}_{n \in \mathbb{N}}$ converges in distribution to a real-valued random variable X if for any bounded, and continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$, one has that

$$\lim_{n \rightarrow +\infty} \mathbb{E}(f(X_n)) = \mathbb{E}(f(X)).$$

In this case, we denote it as $X_n \xrightarrow{d} X$.

4. Multidimensional Stein Methods

In this section, we recall some concepts related to Stein’s method, which we will employ to study the closeness between the law of a function generated by a quantum circuit and that of a Gaussian random variable. Stein’s method is a probabilistic technique that enables the assessment of the distance between two probability measures via second-order differential operators. Specifically, we aim to bound the Wasserstein distance of order one using bounds on solutions to the so-called Stein equations, which we briefly review in this section.

In what follows, we set $d \geq 1$, and denote by $\mathbb{M}_{d \times d}$ the set of all real $d \times d$ -matrices. Given a positive semidefinite matrix $\Sigma \in \mathbb{M}_{d \times d}$, we denote by $\mathcal{N}_d(0, \Sigma)$ the law of a \mathbb{R}^d -valued Gaussian vector with mean zero and covariance matrix Σ . When Σ is the identity matrix, we write $\mathbb{1}_d$. In what follows, we closely follow the notation of [32, Chapter 4]. We denote the Hilbert–Schmidt inner product and the Hilbert–Schmidt norm by $\langle \cdot, \cdot \rangle_{\text{HS}}$ and $\|\cdot\|_{\text{HS}}$, respectively:

$$\langle A, B \rangle_{\text{HS}} := \text{tr}(AB^T), \quad \|A\|_{\text{HS}} := \sqrt{\langle A, A \rangle_{\text{HS}}} \quad \forall A, B \in \mathbb{M}_{d \times d}, \tag{4.1}$$

where $\text{tr}(\cdot)$ denotes the usual trace operator. In the next, we denote by $\|A\|_{\text{op}}$ the operator norm of A

$$\|A\|_{\text{op}} := \sup \{ \|Ax\|_2 : x \in \mathbb{R}^d \text{ such that } \|x\|_2 = 1 \}, \tag{4.2}$$

where $\|\cdot\|_2$ is the Euclidean norm of \mathbb{R}^d . The following lemma, commonly referred to as Stein's Lemma, will be used to derive bounds on the Wasserstein distance of order one between probability measures:

Lemma 4.1 [32, Lemma 4.1.3]. *Let Σ be a positive semidefinite $d \times d$ matrix. Let $N = (N_1, \dots, N_d)$ be a random vector with values in \mathbb{R}^d . Then N has Gaussian $\mathcal{N}_d(0, \Sigma)$ distribution if and only if*

$$\mathbb{E}(\langle N, \nabla f(N) \rangle_{\mathbb{R}^d}) = \mathbb{E}(\langle \Sigma, \text{Hess } f(N) \rangle_{\text{HS}}) \quad (4.3)$$

for every C^2 function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ having bounded first and second derivatives. Here, we have denoted by $\text{Hess } f$ the Hessian matrix of f .

In the statement of the previous lemma, there is no need to restrict Σ to be positive definite. However, in the next definition, we make such an assumption.

Definition 4.1. Let Σ be a positive definite $d \times d$ matrix. Let $N \sim \mathcal{N}_d(0, \Sigma)$. Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be such that $\mathbb{E}|h(N)| < +\infty$. The Stein equation associated to h and N is the partial differential equation

$$\langle \Sigma, \text{Hess } f(x) \rangle_{\text{HS}} - \langle x, \nabla f(x) \rangle_{\mathbb{R}^d} = h(x) - \mathbb{E}(h(N)). \quad (4.4)$$

A solution to Eq. (4.4) is a C^2 -function f satisfying (4.4) for every $x \in \mathbb{R}^d$.

The following proposition relates Lipschitz functions with Eq. (4.4).

Proposition 4.1 [32, Proposition 4.3.2]. *Let $N \sim \mathcal{N}_d(0, \Sigma)$. Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function with Lipschitz constant $K > 0$. Then the function $f_h : \mathbb{R}^d \rightarrow \mathbb{R}$ given by*

$$f_h(x) := \int_0^\infty \mathbb{E} \left(h(N) - h \left(\exp(-t)x + \sqrt{1 - \exp(-2t)}N \right) \right) dt \quad (4.5)$$

is well-defined, it is a C^2 -function and satisfies (4.4) for all $x \in \mathbb{R}^d$. Furthermore,

$$\sup_{x \in \mathbb{R}^d} \|\text{Hess } f_h(x)\|_{\text{HS}} \leq K \sqrt{d} \|\Sigma^{-1}\|_{\text{op}} \|\Sigma\|_{\text{op}}^{\frac{1}{2}}. \quad (4.6)$$

The following theorem provides an explicit relationship between the Wasserstein distance of order one and the bound given in (4.6).

Theorem 4.1 [32, Theorem 4.4.1]. *Let Σ be a positive definite $d \times d$ matrix, and let $N \sim \mathcal{N}_d(0, \Sigma)$. For any square integrable random vector F with values in \mathbb{R}^d , we have*

$$d_W(F, N) \leq \sup_{f \in \mathcal{F}_d(\Sigma)} |\mathbb{E}(\langle \Sigma, \text{Hess } f(F) \rangle_{\text{HS}}) - \mathbb{E}(\langle F, \nabla f(F) \rangle_{\mathbb{R}^d})| \quad (4.7)$$

where

$$\mathcal{F}_d(\Sigma) := \left\{ f : \mathbb{R}^d \rightarrow \mathbb{R} \text{ is } C^2, \text{ and } \sup_{x \in \mathbb{R}^d} \|\text{Hess } f(x)\|_{\text{HS}} \leq \sqrt{d} \|\Sigma^{-1}\|_{\text{op}} \|\Sigma\|_{\text{op}}^{\frac{1}{2}} \right\}. \quad (4.8)$$

This theorem provides an explicit form for bounding the Wasserstein distance of order one. However, this comes at the cost of a less favorable bound with respect to the dimensionality of the ambient space \mathbb{R}^d , and the fact that the right-hand side depends on the inverse of the covariance matrix Σ , which can be challenging to determine explicitly in certain cases [58].

4.1. The Analytic Solution for the Linearized Model

By using a linearized version of (2.40), it is possible to study the behavior of this evolution equation directly. In particular, the first-order approximation of the function $f(\Theta_t, \cdot)$ enables the analysis of the asymptotic properties of the function itself [8, 30]. In what follows, we consider

$$\begin{cases} \frac{d\Theta_t^{\text{lin}}}{dt} = -\eta \nabla_{\Theta} f(\Theta_0, X^T) \nabla_{f^{\text{lin}}(\Theta_t, X)} \mathcal{L}^{\text{lin}}(\Theta_t^{\text{lin}}), \\ \frac{d}{dt} f^{\text{lin}}(\Theta_t^{\text{lin}}, X) = -\eta \hat{K}_{\Theta_0}(x, X^T) \nabla_{f^{\text{lin}}(\Theta_t^{\text{lin}}, X)} \mathcal{L}^{\text{lin}}(\Theta_t^{\text{lin}}), \end{cases} \quad (4.9)$$

where the linearized version $f^{\text{lin}}(\Theta_t^{\text{lin}}, \cdot)$ of the system (2.40) is given by

$$f^{\text{lin}}(\Theta_t^{\text{lin}}, x) = f(\Theta_0, x) + \nabla_{\Theta} f(\Theta_0, x)^T (\Theta_t^{\text{lin}} - \Theta_0). \quad (4.10)$$

Furthermore, \mathcal{L}^{lin} is the loss function \mathcal{L} computed on \mathcal{D} using the linearized model function (4.10) instead of the original model function $f(\Theta_t, \cdot)$. Using the notation (2.37), we may write the equation for $f(\Theta_t, \cdot)$, and for $f^{\text{lin}}(\Theta_t^{\text{lin}}, \cdot)$ as

$$\frac{d}{dt} F(t) = -\eta \hat{K}_{\Theta_t}(F(t) - Y), \quad (4.11)$$

$$\frac{d}{dt} F^{\text{lin}}(t) = -\eta \hat{K}_{\Theta_0}(F^{\text{lin}}(t) - Y) \quad (4.12)$$

where $F^{\text{lin}}(t) := f^{\text{lin}}(\Theta_t^{\text{lin}}, X)$, $\hat{K}_{\Theta_t} = \hat{K}_{\Theta_t}(X, X^T)$, for $t \geq 0$. For one moment, suppose that \hat{K}_{Θ_0} is invertible. The solution of (4.12) gives the evolution of the linearized function evaluated at the inputs of the dataset. Since $F^{\text{lin}}(0) = F(0)$, we have that

$$F^{\text{lin}}(t) = e^{-\eta t \hat{K}_{\Theta_0}} (F(0) - Y) + Y. \quad (4.13)$$

For a new input, one gets that

$$\frac{d}{dt} f^{\text{lin}}(\Theta_t^{\text{lin}}, x) = -\eta \hat{K}_{\Theta_0}(x, X^T) e^{-\eta \hat{K}_{\Theta_0} t} (F(0) - Y). \quad (4.14)$$

Therefore,

$$f^{\text{lin}}(\Theta_t^{\text{lin}}, x) = f(\Theta_0, x) - \hat{K}_{\Theta_0}(x, X^T) \hat{K}_{\Theta_0}^{-1} \left(\mathbb{1} - e^{-\eta \hat{K}_{\Theta_0} t} \right) (F(0) - Y). \quad (4.15)$$

Based on this observation, [8] proved that the probability distribution of $\{f^{\text{lin}}(\Theta_t, x)\}_{x \in \mathcal{X}}$ converges in distribution to a Gaussian process in the limit of infinite width. This motivates the following notations. We denote by $f^{(\infty)}(\cdot)$

the Gaussian process to which $f(\Theta_0, \cdot)$ converges in distribution. Let us denote by \bar{X} the vectors of all the inputs of \mathcal{X} . For $t > 0$, we set

$$\begin{aligned} \mathcal{K}_t(x, x') &:= \mathcal{K}_0(x, x') - K(x, X^T)K^{-1}(\mathbb{1} - e^{-t\eta K})\mathcal{K}_0(X, x') \\ &\quad - K(x', X^T)K^{-1}(\mathbb{1} - e^{-t\eta K})\mathcal{K}_0(X, x) \\ &\quad + K(x, X^T)K^{-1}(\mathbb{1} - e^{-t\eta K})\mathcal{K}_0(X, X^T) \\ &\quad (\mathbb{1} - e^{-t\eta K})K^{-1}K(X, x'), \end{aligned} \tag{4.16}$$

$$\mu_t(x) := K(x, X^T)K^{-1}(\mathbb{1} - e^{-t\eta K})Y. \tag{4.17}$$

More precisely, [8, Corollary 4.9] proved that $\{f^{\text{lin}}(\Theta_t, x)\}_{x \in \mathcal{X}}$ converges to the Gaussian process with mean and covariance defined by (4.17) and (4.16), respectively.

5. Main Results

In this section we state our main results.

Theorem 5.1 (Convergence at initialization, formal statement). *Suppose that A 1-A 2 hold true. Let $\bar{\mathcal{K}}_0 := \mathcal{K}_0(\bar{X}, \bar{X}^T)$ be the covariance matrix of the quantum neural network at initialization. Then, we have*

$$\begin{aligned} d_W(f(\Theta, \bar{X}), \mathcal{N}(0, \bar{\mathcal{K}}_0)) &\leq \bar{N}^2 \left(2 \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \\ &\quad \frac{m|\mathcal{M}|^{7/2}|\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)), \end{aligned} \tag{5.1}$$

where $\mathcal{N}(0, \bar{\mathcal{K}}_0)$ denotes a centered Gaussian process with covariance $\bar{\mathcal{K}}_0$.

Proof. See Sect. 6. □

Theorem 5.2 (Convergence of the trained network, formal statement). *Suppose that A 1, A 2, A 3, and A 4 hold true. Let $\bar{\mathcal{K}}_0 := \mathcal{K}_0(\bar{X}, \bar{X}^T)$ be the covariance matrix of the quantum neural network at initialization. Then, for any $s > 0$ and any $t > 0$, we have*

$$\begin{aligned} d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\ \leq C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \bar{\mathcal{K}}_0, s) \frac{L^2 m^{9/4} |\mathcal{M}|^{11/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)), \end{aligned} \tag{5.2}$$

where

$$\begin{aligned}
 C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \bar{\mathcal{K}}_0, s) &:= 96s\bar{N}^2 + 8\sqrt{\bar{N}} \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} (\bar{N}^{3/2} + 4n^{5/2}) \\
 &\quad + 132n^2\sqrt{\bar{N}} (2\|Y\|_2^2 + 4n) \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
 &\quad + 2(\sqrt{n} + \|Y\|_2) \left(\frac{1}{\lambda_{\min}^K} + \frac{6}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \right),
 \end{aligned} \tag{5.3}$$

and $\mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)$ denotes Gaussian process with mean vector $\mu_t(\bar{X})$ and covariance matrix $\bar{\mathcal{K}}_t := \mathcal{K}_t(\bar{X}, \bar{X}^T)$, where μ_t and \mathcal{K}_t are defined by (4.17) and (4.16), respectively.

Proof. See Sect. 7. □

Remark 5.1. The convergence rates in Theorem 5.1 and Theorem 5.2 strongly depends on the number of possible inputs \bar{N} . We stress that, if we are interested in only a limited set of inputs, we can improve the rates by restrict \mathcal{X} to such set. For example, we can choose \mathcal{X} to be the set of the training inputs together with a single test input and get $\bar{N} = n + 1$. Furthermore, let us notice that the right hand side of (5.2) does not depend on the training time, and therefore, this bound is uniform in time and holds for finite width even in the limit of $t \rightarrow +\infty$.

6. Convergence at Initialization

In this Section, we prove Theorem 5.1.

Proof. Let $h : \mathbb{R}^N \rightarrow \mathbb{R}$ be a Lipschitz function with Lipschitz constant equal to one. Let \bar{X} be the vector of all the inputs in \mathcal{X} . Let

$$\begin{aligned}
 W_i &:= \sum_{k \notin \mathcal{P}_i} \frac{f_k(\Theta, \bar{X})}{N(m)}, \quad Z_i := \sum_{k \in \mathcal{P}_i} \frac{f_k(\Theta, \bar{X})}{N(m)}. \quad \text{for } i = 1, \dots, m; \\
 X_i &:= \frac{f_i(\Theta, \bar{X})}{N(m)} \quad \text{for } i = 1, \dots, m.
 \end{aligned} \tag{6.1}$$

For the sake of a clean notation, let $W := f(\Theta, \bar{X})$, and we set $Hf := \text{Hessian} f$ to denote the Hessian of a generic smooth function $f : \mathbb{R}^d \rightarrow \mathbb{R}$. Notice that

$$W = W_i + Z_i. \tag{6.2}$$

Let us denote by g one solution of (4.4) associated to h . Further, by Proposition 4.1 one gets that

$$\mathbb{E} (h(W) - h(\mathcal{N}(0, \bar{\mathcal{K}}_0))) = \mathbb{E} (\langle \bar{\mathcal{K}}_0, Hg(W) \rangle_{\text{HS}} - W \cdot \nabla g(W)). \tag{6.3}$$

Let us now choose a further random variable θ independent from all the variables X_i and uniformly distributed in $[0, 1]$. By the fundamental theorem of calculus, we write for $x, u, v \in \mathbb{R}^N$

$$\langle \nabla g(x + v) - \nabla g(x), u \rangle = \mathbb{E} \left(\langle Hg(x + \theta v), uv^T \rangle_{\text{HS}} \right). \quad (6.4)$$

Since W_i, X_i are independent, then we obtain

$$\begin{aligned} \mathbb{E} (\langle X_i, \nabla g(W_i + Z_i) \rangle) &= \mathbb{E} (\langle X_i, \nabla g(W_i + Z_i) - \nabla g(W_i) \rangle) \\ &= \mathbb{E} (\langle Hg(W_i + \theta Z_i), X_i Z_i^T \rangle_{\text{HS}}). \end{aligned} \quad (6.5)$$

Then, we insert this equality into (6.3) to obtain that

$$\begin{aligned} &\mathbb{E} (h(W) - h(\mathcal{N}(0, \overline{\mathcal{K}}_0))) \\ &= \mathbb{E} \left(\langle \overline{\mathcal{K}}_0, Hg(W) \rangle_{\text{HS}} - \sum_{i=1}^m \langle X_i Z_i^T, Hg(W_i + \theta Z_i) \rangle_{\text{HS}} \right). \end{aligned} \quad (6.6)$$

Furthermore, we write

$$\begin{aligned} &\mathbb{E} (h(W) - h(\mathcal{N}(0, \overline{\mathcal{K}}_0))) \\ &= \mathbb{E} \left(\left\langle \overline{\mathcal{K}}_0 - \sum_{i=1}^m X_i Z_i^T, Hg(W) \right\rangle_{\text{HS}} \right) \\ &\quad + \mathbb{E} \left(\sum_{i=1}^m \langle X_i Z_i^T, Hg(W_i + Z_i) - Hg(W_i + \theta Z_i) \rangle_{\text{HS}} \right). \end{aligned} \quad (6.7)$$

Note then that

$$\begin{aligned} &\left| \mathbb{E} \sum_{i=1}^m \langle X_i Z_i^T, Hg(W_i + Z_i) - Hg(W_i + \theta Z_i) \rangle_{\text{HS}} \right| \\ &\leq \mathbb{E} \sum_{i=1}^m \|X_i\|_2 \|Z_i^T\|_2 \|Hg(W_i + Z_i) - Hg(W_i + \theta Z_i)\|_{\text{op}}. \end{aligned} \quad (6.8)$$

On the other hand, notice that by Proposition A.1 with $d = \overline{N}$, we get for $x, y \in \mathbb{R}^{\overline{N}}$,

$$\|Hg(x) - Hg(y)\|_{\text{op}} \leq \omega(\overline{N}, \|x - y\|_2), \quad (6.9)$$

where ω denotes the modulus of continuity of Hg . Hence, by (6.9), and since the maximum number of qubits involved in the definition of Z_i are D as defined

in (2.31) and (2.24), we have that

$$\begin{aligned}
 & \left| \mathbb{E} \sum_{i=1}^m \langle X_i Z_i^T, Hg(W_i + Z_i) - Hg(W_i + \theta Z_i) \rangle_{\text{HS}} \right| \\
 & \leq \mathbb{E} \sum_{i=1}^m \|X_i\|_2 \|Z_i^T\|_2 \omega(\bar{N}, (1-\theta) \|Z_i\|_2) \\
 & \leq \mathbb{E} \sum_{i=1}^m \|X_i\|_2 \|Z_i^T\|_2 \omega(\bar{N}, \|Z_i\|_2) \\
 & \leq \frac{Dm\bar{N}}{(N(m))^2} \omega\left(\bar{N}, \frac{D}{N(m)} \sqrt{\bar{N}}\right).
 \end{aligned} \tag{6.10}$$

Let us now proceed to bound the first term on the right hand side of (6.7). Let us define

$$C_i := \mathbb{E}(X_i Z_i^T), \tag{6.11}$$

and observe that $\sum_{i=1}^m C_i = \mathbb{E} \sum_{i=1}^m X_i W^T = \bar{\mathcal{K}}_0$ where we have used that X_i and W_i are independent, for all $i = 1, \dots, m$. Since $g \in \mathcal{F}_d(\bar{\mathcal{K}}_0)$, we have that Hg satisfies (4.6), and thus

$$\begin{aligned}
 & \left| \mathbb{E} \left\langle \bar{\mathcal{K}}_0 - \sum_{i=1}^m X_i Z_i^T, \text{Hessian } g(W) \right\rangle_{\text{HS}} \right| \\
 & \leq \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \mathbb{E} \left\| \bar{\mathcal{K}}_0 - \sum_{i=1}^m X_i Z_i^T \right\|_{\text{HS}} \\
 & \leq \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \sqrt{\mathbb{E} \left\| \bar{\mathcal{K}}_0 - \sum_{i=1}^m X_i Z_i^T \right\|_{\text{HS}}^2} \\
 & = \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \sqrt{\mathbb{E} \left\| \sum_{i=1}^m (C_i - X_i Z_i^T) \right\|_{\text{HS}}^2} \\
 & = \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \sqrt{\mathbb{E} \left\| \sum_{i=1}^m M_i \right\|_{\text{HS}}^2},
 \end{aligned} \tag{6.12}$$

where we have defined $M_i := C_i - X_i Z_i^T$, for $i = 1, \dots, m$. Notice that from Remark 2.2, by the definition of Z_i , and Z_j , these are independent whenever $\mathcal{P}_i \cap \tilde{\mathcal{P}}_j = \emptyset$. Notice that the same holds true for M_i and M_j for any $i, j = 1, \dots, m$. Furthermore, since the maximum number of qubits involved in the

definition of Z_i is D as defined in (2.24), then

$$\|C_i\|_{\text{HS}}, \quad \|X_i Z_i^T\|_{\text{HS}} \leq \frac{D\bar{N}}{(N(m))^2}, \quad (6.13)$$

so that

$$\|M_i\|_{\text{HS}} \leq \frac{2D\bar{N}}{(N(m))^2}. \quad (6.14)$$

Then, by inserting in (6.7), inequalities (6.10), (6.12), and (6.14), we conclude that

$$\begin{aligned} & \mathbb{E}|h(W) - h(\mathcal{N}(0, \bar{\mathcal{K}}_0))| \\ & \leq \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \sqrt{\mathbb{E} \left\| \sum_{i=1}^m M_i \right\|_{\text{HS}}^2} + \frac{Dm\bar{N}}{(N(m))^2} \omega \left(\bar{N}, \frac{D}{N(m)} \sqrt{\bar{N}} \right) \\ & \leq \sqrt{\bar{N}} \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} \sqrt{\frac{\tilde{D}m(2\bar{N}D)^2}{(N(m))^4} + \frac{Dm\bar{N}}{(N(m))^2} \omega \left(\bar{N}, \frac{D}{N(m)} \sqrt{\bar{N}} \right)} \\ & =: C_m^{\bar{N}, \bar{\mathcal{K}}_0}, \end{aligned} \quad (6.15)$$

where in the last inequality, we have used that if i, j are observable such that M_i, M_j are independent, then their components do not give any contribution to the sum on the right hand side of (6.15) (recall that \tilde{D} is the constant defined in (2.27), and denotes the maximum over all possible related observables i, j). By taking the supremum over all the functions with Lipschitz constant one, we obtain the left-hand side of (5.1). That is,

$$d_W(f(\Theta, \bar{X}), \mathcal{N}(0, \bar{\mathcal{K}}_0)) \leq C_m^{\bar{N}, \bar{\mathcal{K}}_0}. \quad (6.16)$$

It remains to bound this constant $C_m^{\bar{N}, \bar{\mathcal{K}}_0}$. First, we notice that

$$\begin{aligned} \mathcal{K}_0(x, x) &= \frac{1}{N^2(m)} \sum_{k, k'=1}^m \mathbb{E}[f_k(\Theta, x) f_{k'}(\Theta, x)] \\ &= \frac{1}{N^2(m)} \sum_{k=1}^m \left(\sum_{k' \in \mathcal{P}_k} \mathbb{E}[f_k(\Theta, x) f_{k'}(\Theta, x)] + \sum_{k' \notin \mathcal{P}_k} \mathbb{E}[f_k(\Theta, x)] \mathbb{E}[f_{k'}(\Theta, x)] \right) \\ &= \frac{1}{N^2(m)} \sum_{k=1}^m \sum_{k' \in \mathcal{P}_k} \mathbb{E}[f_k(\Theta, x) f_{k'}(\Theta, x)] \\ &\leq \frac{1}{N^2(m)} \sum_{k=1}^m \sum_{k' \in \mathcal{P}_k} \mathbb{E}[|f_k(\Theta, x)| |f_{k'}(\Theta, x)|] \\ &\leq \frac{1}{N^2(m)} \sum_{k=1}^m \sum_{k' \in \mathcal{P}_k} 1 \leq \frac{1}{N^2(m)} m \max_k |\mathcal{P}_k| \leq \frac{1}{N^2(m)} m |\mathcal{M}| |\mathcal{N}|, \end{aligned} \quad (6.17)$$

where the last inequality follows from Lemma 2.1. Therefore,

$$1 = \max_{x \in \mathcal{X}} \mathcal{K}_0(x, x) \leq \frac{1}{N^2(m)} m |\mathcal{M}| |\mathcal{N}|. \quad (6.18)$$

We recall that

$$\omega(d, x) := \begin{cases} x(C(1, d) - 2 \log x), & \text{if } x \leq 1, \\ C(1, d), & \text{if } x > 1. \end{cases} \quad \text{with } C(1, d) := 2^{\frac{3}{2}} \frac{1 + 2d}{d} \frac{\Gamma(\frac{1+d}{2})}{\Gamma(\frac{d}{2})}. \quad (6.19)$$

Recalling Wendel's inequality [59], valid for any $x > 0$ and $s \in (0, 1)$,

$$\frac{\Gamma(x+s)}{\Gamma(x)} \leq x^s, \quad (6.20)$$

we set $x = d/2$ and $s = 1/2$ in order to upper bound

$$C(1, d) \leq 2^{\frac{3}{2}} \frac{1 + 2d}{d} \sqrt{\frac{d}{2}} \leq 2 + 4\sqrt{d}. \quad (6.21)$$

In particular, for $d \geq 1$ we can upper bound $C(1, d) \leq 6\sqrt{d}$. Furthermore,

$$\begin{aligned} \omega\left(d, \frac{D}{N(m)} \sqrt{N}\right) &= \begin{cases} \frac{D\sqrt{N}}{N(m)} \left(C(1, d) - 2 \log D\sqrt{N} + \log N(m)\right) & \text{if } D\sqrt{N} \leq N(m), \\ C(1, d) & \text{if } D\sqrt{N} > N(m), \end{cases} \\ &\leq \begin{cases} \frac{D\sqrt{N}}{N(m)} (C(1, d) + \log N(m)) & \text{if } D\sqrt{N} \leq N(m), \\ \frac{D\sqrt{N}}{N(m)} C(1, d) & \text{if } D\sqrt{N} > N(m), \end{cases} \\ &\leq \frac{D\sqrt{N}}{N(m)} C(1, d) (1 + \log N(m)) \\ &\leq 6\sqrt{d} \frac{D\sqrt{N}}{N(m)} (1 + \log N(m)), \end{aligned} \quad (6.22)$$

where we have noticed that $d, D, N(m) \geq 1$. This, combined with Lemma 2.1 ($D \leq |\mathcal{M}||\mathcal{N}|$) and Lemma 2.2 ($\tilde{D} \leq (|\mathcal{M}||\mathcal{N}|)^4$), implies

$$\begin{aligned} C_m^{\overline{N}, \overline{\mathcal{K}}_0} &\leq 2\overline{N} \sqrt{\overline{N}} \left\| \overline{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \overline{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} \frac{\sqrt{m} |\mathcal{M}|^3 |\mathcal{N}|^3}{N^2(m)} + 6\overline{N}^2 \frac{m |\mathcal{M}|^2 |\mathcal{N}|^2}{N^3(m)} (1 + \log N(m)) \\ &\leq \overline{N}^2 \left(2 \left\| \overline{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \overline{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \frac{m |\mathcal{M}|^{7/2} |\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)). \end{aligned} \quad (6.23)$$

where in the second inequality we have used (6.18). \square

7. Convergence of the Trained Network

7.1. Quantitative Lazy Training for a Finite Size Circuit

In this section, we provide a renewed proof of some results in [8] that takes into account the quantitative contribution of the finite-size effects and where we do not need to build a sequence of networks with diverging width. Therefore, the results will apply to any network satisfying the expressivity condition (7.6). We will often need the following lemma.

Lemma 7.1. (Lipschitzness of the gradient [8, Lemma 4.20]) *The following inequalities hold:*

$$|\partial_{\theta_i} f(\Theta, x)| \leq 2 \frac{|\mathcal{M}_i|}{N(m)}, \quad (7.1)$$

$$\|\partial_{\theta_i} f(\Theta, X)\|_2 \leq 2\sqrt{n} \frac{|\mathcal{M}|}{N(m)}, \quad (7.2)$$

$$\|\nabla_{\Theta} f(\Theta, x)\|_1 \leq 2L \frac{m}{N(m)} |\mathcal{M}|, \quad (7.3)$$

$$\|\nabla_{\Theta} f(\Theta, x) - \nabla_{\Theta} f(\Theta', x)\|_{\infty} \leq 4 \frac{|\mathcal{M}|^2 |\mathcal{N}|}{N(m)} \|\Theta - \Theta'\|_{\infty}. \quad (7.4)$$

Theorem 7.1. (Lazy training, formal statement) *Suppose that A 1, A 2, and A 3 hold true, and let*

$$R(\delta) := \|Y\|_2 + \sqrt{\frac{2n}{\delta}}. \quad (7.5)$$

Let us assume that, fixed a constant $0 < \delta < 1$, it holds that

$$\lambda_{\min}^K \geq \frac{96n\sqrt{Lm}|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \sqrt{\log \frac{2n^2}{\delta}}. \quad (7.6)$$

Then, there exist a positive number $\tilde{\lambda}_{\min}(\delta)$ satisfying

$$\tilde{\lambda}_{\min}(\delta) \geq \frac{1}{3} \lambda_{\min}^K, \quad (7.7)$$

whose explicit expression is provided in (7.20) below, such that, when applying gradient flow with learning rate η , the following inequalities hold with probability at least $1 - \delta$ over random initialization:

$$\mathcal{L}(\Theta_t) \leq \frac{R^2(\delta)}{2} e^{-2\eta\tilde{\lambda}_{\min}(\delta)t} \quad \forall t \geq 0, \quad (7.8)$$

$$\|\Theta_t - \Theta_0\|_{\infty} \leq \frac{2\sqrt{n}R(\delta)|\mathcal{M}|}{N(m)} \frac{1}{\tilde{\lambda}_{\min}(\delta)} \left(1 - e^{-\eta\tilde{\lambda}_{\min}(\delta)t}\right) \quad \forall t \geq 0, \quad (7.9)$$

$$\sup_{\substack{x \in \mathcal{X} \\ t \geq 0}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)|$$

$$\leq 132n^2 R^2(\delta) \left(1 + \left(\tilde{\lambda}_{\min}(\delta)\right)^{-3}\right) \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} \log N(m). \quad (7.10)$$

Proof. We recall Chebyshev's inequality for a random vector V with $\mathbb{E}[V] = v$ and $\mathbb{E}[(X - v)^2] = \sigma^2$:

$$\mathbb{P}(\|V - v\|_2 \geq k\|\sigma\|_2) \leq \frac{1}{k^2}. \quad (7.11)$$

Calling $F(0) = f(\Theta_0, X)$, we have $\mathbb{E}[F(0)] = 0$ and $\mathbb{E}[F^2(0)] = \text{diag}(\mathcal{K}_0(X, X^T))$. Let us define

$$R(\delta) := \|Y\|_2 + \sqrt{\frac{2}{\delta}} \left\| \left(\text{diag}(\mathcal{K}_0(X, X^T)) \right)^{1/2} \right\|_2 \quad (7.12)$$

Hence,

$$\begin{aligned} \mathbb{P}(\|F(0) - Y\|_2 \geq R(\delta)) &\leq \mathbb{P}(\|F(0)\|_2 \geq R(\delta) - \|Y\|_2) \\ &\leq \mathbb{P}\left(\|F(0)\|_2 \geq \sqrt{\frac{2}{\delta}} \left\| \left(\text{diag}(\mathcal{K}_0(X, X^T))^{1/2} \right) \right\|_2\right) \leq \frac{\delta}{2}. \end{aligned} \quad (7.13)$$

Therefore,

$$\mathbb{P}(\|F(0) - Y\|_2 < R(\delta)) \geq 1 - \frac{\delta}{2}. \quad (7.14)$$

Using that

$$\sum_{i=1}^n \mathcal{K}_0(x^{(i)}, x^{(i)}) = n, \quad (7.15)$$

we have

$$\begin{aligned} R(\delta) &:= \|Y\|_2 + \sqrt{\frac{2}{\delta}} \left\| \left(\text{diag}(\mathcal{K}_0(X, X^T))^{1/2} \right) \right\|_2 \\ &= \|Y\|_2 + \sqrt{\frac{2}{\delta}} \left(\sum_{i=1}^n \mathcal{K}_0(x^{(i)}, x^{(i)}) \right)^{1/2} \\ &\leq \|Y\|_2 + \sqrt{\frac{2n}{\delta}}. \end{aligned} \quad (7.16)$$

where we have used that $\mathcal{K}_0(x^{(i)}, x^{(i)}) \leq 1$. In particular, $R(\delta) \geq 1$, which means $R(\delta) \leq R^2(\delta)$, which will be used later.

For the sake of a lean notation, we define

$$g(\delta) := \frac{16n\sqrt{Lm}|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \sqrt{\log \frac{2n^2}{\delta}}, \quad (7.17)$$

$$\begin{aligned} \rho(m) &:= \frac{1}{2} \frac{N^2(m)}{16nLm|\mathcal{M}|^2|\mathcal{N}|} (\lambda_{\min}^K - g(\delta)) \\ &\quad \left(1 - \sqrt{1 - \frac{128n\sqrt{n}R(\delta)}{(\lambda_{\min}^K - g(\delta))^2} \frac{Lm|\mathcal{M}|^3|\mathcal{N}|}{N^3(m)}} \right) \\ &\leq \frac{4\sqrt{n}R(\delta)}{\lambda_{\min}^K - g(\delta)} \frac{|\mathcal{M}|}{N(m)}, \end{aligned} \quad (7.18)$$

$$h(\delta) := 16n \frac{Lm}{N^2(m)} \rho(m), \quad (7.19)$$

$$\tilde{\lambda}_{\min}(\delta) := \lambda_{\min}^K - g(\delta) - h(\delta), \quad (7.20)$$

where the bound follows from $1 - \sqrt{1-x} \leq x$. Let us informally clarify the meaning of these definitions; in particular, the following claims, which will be formally proved later, might help to understand the idea behind the apparently complicated form of (7.17)-(7.20). We call $\text{Spec}(A)$ the spectrum of a symmetric matrix A .

- The minimal eigenvalue λ_{\min}^K of the analytic kernel $K(X, X^T)$ is not in general the same one of the empirical kernel at initialization \hat{K}_{Θ_0} and during the training \hat{K}_{Θ_t} . We will need to lower bound $\min_{0 \leq t \leq t_1} \min \text{Spec}(\hat{K}_{\Theta_t})$ and we are going to call $\tilde{\lambda}_{\min}(\delta)$ our estimate. t_1 is a time that will be discussed later. Of course the bound will worsen with the size of the trajectory of Θ_t ($0 \leq t \leq t_1$) in the parameter space: if $\rho(m)$ is the maximal possible radius that such trajectory can achieve with respect to the starting point, $\tilde{\lambda}_{\min}(\delta)$ will decrease with $\rho(m)$.
- $\tilde{\lambda}_{\min}(\delta)$ is made of two corrections to λ_{\min}^K : the first one— $g(\delta)$ —emerges from the fact that, even if the empirical neural tangent kernel at initialization is concentrated on the analytic neural tangent kernel, they are different matrices; the second one— $h(\delta)$ —is due to the evolution from \hat{K}_{Θ_0} to \hat{K}_{Θ_t} . The latter grows with the radius of the trajectory.
- In the proof we will study the trajectory from $t = 0$ to $t = t_1$, where the final time is the first instant when the parameter vector Θ_{t_1} is far from Θ_0 more than $\rho(m)$ in the parameter space, where it will be crucial to require $\rho(m)$ to satisfy

$$\rho(m) = \frac{2\sqrt{n}R(\delta)|\mathcal{M}|}{N(m)} \frac{1}{\tilde{\lambda}_{\min}(\delta)} \quad (7.21)$$

in order to conclude the proof of (7.8) and (7.9). Since $\rho(m)$ linearly appears in $\tilde{\lambda}_{\min}(\delta)$, in order to compute $\rho(m)$ we need to solve the second-order equation provided by (7.21). The solution is (7.18).

Therefore, by our definitions, we have that $\rho(m)$ satisfies (7.21) (which can be explicitly verified by using the definitions provided above). In particular, by (7.18),

$$\tilde{\lambda}_{\min}(\delta) = \frac{2\sqrt{n}R(\delta)|\mathcal{M}|}{N(m)} \frac{1}{\rho(m)} \geq \frac{1}{2}(\lambda_{\min}^K - g(\delta)), \quad (7.22)$$

and, by hypothesis (7.6), which can be restated as

$$g(\delta) \leq \frac{1}{6}\lambda_{\min}^K, \quad (7.23)$$

we have that

$$\tilde{\lambda}_{\min}(\delta) \geq \frac{1}{3}\lambda_{\min}^K > 0. \quad (7.24)$$

Let $B_r(\Theta_0) = \{\Theta : \|\Theta_0 - \Theta\|_\infty < r\}$ be the ball of center Θ_0 and radius r .

Let us now recall the following quantitative result about the closeness between \hat{K} and K proved in [8]:

Theorem 7.2. [8, Theorem 4.7] *Suppose that we initialize a parameterized quantum circuit randomly, that is, the parameters θ_i are taken as independent random variables. Then, for any $x, x' \in \mathcal{X}$ we have*

$$\mathbb{P}\left(|\hat{K}_{\Theta}(x, x') - K(x, x')| \geq \varepsilon\right) \leq \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^4|\mathcal{N}|^2} \varepsilon^2\right). \quad (7.25)$$

Let $\|\cdot\|_{\text{HS}}$ be the Hilbert–Schmidt norm and let us apply Theorem 7.2 to $F_{ij} := K(x^{(i)}, x^{(j)}) - \hat{K}_{\Theta_0}(x^{(i)}, x^{(j)})$, where $1 \leq i, j \leq n$, i.e. $F = K - \hat{K}_{\Theta_0}$:

$$\begin{aligned} \mathbb{P}[\|F\|_{\text{HS}} \geq \epsilon] &= \mathbb{P}\left[\sum_{i,j=1}^n (F_{ij})^2 \geq \epsilon^2\right] \leq \mathbb{P}\left[\max_{ij} |F_{ij}| \geq \epsilon/n\right] \\ &\leq \sum_{i,j=1}^n \mathbb{P}[|F_{ij}| \geq \epsilon/n] \leq n^2 \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^4|\mathcal{N}|^2} \frac{\epsilon^2}{n^2}\right), \end{aligned} \quad (7.26)$$

whence, by (7.17),

$$\mathbb{P}[\|F\|_{\text{HS}} < g(\delta)] \leq 1 - \frac{\delta}{2}. \quad (7.27)$$

When $\|F\|_{\text{HS}} < g(\delta)$, the maximum eigenvalue of $|F|$ is $\lambda_F < g(\delta)$, so

$$F \preceq |F| \preceq \lambda_F \mathbb{1} \prec g(\delta) \mathbb{1}.$$

The previous equation implies that

$$K - \hat{K}_{\Theta_0} \prec g(\delta) \mathbb{1} \quad \Rightarrow \quad \hat{K}_{\Theta_0} \succ K - g(\delta) \mathbb{1} \succeq (\lambda_{\min}^K - g(\delta)) \mathbb{1}.$$

So,

$$\hat{K}_{\Theta_0}(X, X^T) \succ (\lambda_{\min}^K - g(\delta)) \mathbb{1} \quad (7.28)$$

with probability at least $1 - \frac{\delta}{2}$. Let

$$t_1 = \inf\{t : \|\Theta_t - \Theta_0\|_{\infty} \geq \rho(m)\}. \quad (7.29)$$

For $t \leq t_1$ we have, by [8, Lemma 4.22],¹ we have

$$|\hat{K}_{\Theta_t}(x, x') - \hat{K}_{\Theta_0}(x, x')| \leq 16 \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \|\Theta_t - \Theta_0\|_{\infty} \leq 16 \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m). \quad (7.30)$$

Therefore,

$$\|\hat{K}_{\Theta_t} - \hat{K}_{\Theta_0}\|_{\text{HS}} \leq 16n \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m) = h(\delta), \quad (7.31)$$

whence

$$\hat{K}_{\Theta_t}(X, X^T) \succ (\lambda_{\min}^K - g(\delta) - h(\delta)) \mathbb{1} = \tilde{\lambda}_{\min}(\delta) \mathbb{1} \quad \forall t \leq t_1, \quad (7.32)$$

with probability at least $1 - \delta$ (by the union bound applied to the events described above). Recalling that

$$\frac{d}{dt} f(\Theta_t, x) = -\eta \hat{K}_{\Theta_t}(x, X^T) \cdot (F(t) - Y),$$

for $t \leq t_1$ we have, with probability at least $1 - \delta$ (see eq. (7.14)),

$$\frac{d}{dt} \|F(t) - Y\|_2^2 = -2\eta (F(t) - Y)^T \hat{K}_{\Theta_t}(F(t) - Y) \leq -2\eta \tilde{\lambda}_{\min}(\delta) \|F(t) - Y\|_2^2,$$

¹since in the current work we do not introduce the normalization $N_K(m)$ for the neural tangent kernel, it is sufficient to set $N_K(m) = 1$, and $\Sigma_1 = Lm$ in the quoted result.

$$\begin{aligned} \Rightarrow \mathcal{L}(\Theta_t) &= \frac{1}{2} \|F(t) - Y\|_2^2 \leq \frac{1}{2} e^{-2\eta\tilde{\lambda}_{\min}(\delta)t} \|F(0) - Y\|_2^2 \\ &\leq e^{-2\eta\tilde{\lambda}_{\min}(\delta)t} \frac{R^2(\delta)}{2}. \end{aligned} \quad (7.33)$$

Recalling also that

$$\dot{\Theta}_t = -\eta \nabla_{\Theta} f(\Theta_t, X) \cdot (F(t) - Y)$$

and using Lemma 7.1,

$$\begin{aligned} \frac{d}{dt} |\theta_i(t) - \theta_i(0)| &\leq \left| \frac{d}{dt} \theta_i(t) \right| = \left| \eta \partial_{\theta_i} f(\Theta, X) \cdot (F(t) - Y) \right| \\ &\leq \eta \|\partial_{\theta_i} f(\Theta, X)\|_2 \|F(t) - Y\|_2 \\ &\leq 2\eta\sqrt{n} \frac{|\mathcal{M}|}{N(m)} R(\delta) e^{-\eta\tilde{\lambda}_{\min}(\delta)t}, \end{aligned} \quad (7.34)$$

$$\Rightarrow |\theta_i(t) - \theta_i(0)| \leq \frac{2\sqrt{n}R(\delta)|\mathcal{M}|}{N(m)} \frac{1}{\tilde{\lambda}_{\min}(\delta)} \left(1 - e^{-\eta\tilde{\lambda}_{\min}(\delta)t}\right) \quad \forall t \leq t_1, \quad (7.35)$$

$$\Rightarrow \|\Theta_t - \Theta\|_{\infty} \leq \rho(m) \left(1 - e^{-\eta\tilde{\lambda}_{\min}(\delta)t}\right) \quad \forall t \leq t_1, \quad (7.36)$$

with probability at least $1 - \delta$, where the last implication follows from (7.21) combined with the definition (7.24) of $\tilde{\lambda}_{\min}(\delta)$. If $t_1 < \infty$, then

$$\|\Theta_{t_1} - \Theta\|_{\infty} \leq \rho(m) \left(1 - e^{-\eta\tilde{\lambda}_{\min}(\delta)t_1}\right) < \rho(m) \quad \forall t \leq t_1, \quad (7.37)$$

with probability at least $1 - \delta$, but this contradicts the definition (7.29) of t_1 , so we must have $t_1 = \infty$. This implies that (7.33) and (7.36) hold for any $t > 0$; so, we have proved (7.8) and (7.9). Combining the latter with [8, Theorem 4.23], we can estimate

$$\sup_t |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t, x)| \leq \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \|\Theta_t - \Theta_0\|_2^2 \leq \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \rho^2(m), \quad (7.38)$$

with probability at least $1 - \delta$. Now, we want to prove that

$$\|F(t) - F^{\text{lin}}(t)\|_2 \leq \frac{16nR(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m) \left(1 - e^{-\eta\tilde{\lambda}_{\min}(\delta)t}\right), \quad (7.39)$$

with probability at least $1 - \delta$. We follow the strategy of [30], using our results to improve the final bound. Let us define

$$\Delta(t) = \|F(t) - F^{\text{lin}}(t)\|_2$$

and compute

$$\begin{aligned} &\frac{1}{2} \frac{d}{dt} \Delta^2(t) \\ &= \sum_{i=1}^n \left(f(\Theta_t, x^{(i)}) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x^{(i)}) \right) \left(\frac{d}{dt} f(\Theta_t, x^{(i)}) - \frac{d}{dt} f^{\text{lin}}(\Theta_t^{\text{lin}}, x^{(i)}) \right) \end{aligned}$$

$$\begin{aligned}
 &= -\eta \sum_{i=1}^n \left(f(\Theta_t, x^{(i)}) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x^{(i)}) \right) \\
 &\quad \times \left(\hat{K}_{\Theta_t}(x^{(i)}, X^T)(f(\Theta_t, X) - Y) - \hat{K}_{\Theta}(x^{(i)}, X^T)(f(\Theta_t^{\text{lin}}, X) - Y) \right) \\
 &= -\eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta_t}(F(t) - Y) + \eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta}(F^{\text{lin}}(t) - Y) \\
 &= -\eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta_t}(F(t) - Y) \\
 &\quad - \eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta}(F(t) - F^{\text{lin}}(t)) \\
 &\quad + \eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta}(F(t) - Y)
 \end{aligned} \tag{7.40}$$

Noticing that \hat{K}_{Θ} is positive semidefinite,

$$-\eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta}(F(t) - F^{\text{lin}}(t)) \leq 0, \tag{7.41}$$

so (7.40) becomes

$$\begin{aligned}
 \frac{1}{2} \frac{d}{dt} \Delta^2(t) &\leq -\eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta_t}(F(t) - Y) \\
 &\quad + \eta(F(t) - F^{\text{lin}}(t))^T \hat{K}_{\Theta}(F(t) - Y) \\
 &= -\eta(F(t) - F^{\text{lin}}(t))^T (\hat{K}_{\Theta_t} - \hat{K}_{\Theta})(F(t) - Y)
 \end{aligned} \tag{7.42}$$

whence

$$\begin{aligned}
 \left| \Delta(t) \frac{d}{dt} \Delta(t) \right| &\leq \eta \|F(t) - F^{\text{lin}}(t)\|_2 \|\hat{K}_{\Theta_t} - \hat{K}_{\Theta}\|_{\text{op}} \|F(t) - Y\|_2 \\
 &= \eta \Delta(t) \|\hat{K}_{\Theta_t} - \hat{K}_{\Theta}\|_{\text{op}} \|F(t) - Y\|_2.
 \end{aligned} \tag{7.43}$$

Therefore,

$$\left| \frac{d}{dt} \Delta(t) \right| \leq \eta \|\hat{K}_{\Theta_t} - \hat{K}_{\Theta}\|_{\mathcal{L}} \|F(t) - Y\|_2. \tag{7.44}$$

By (7.33)

$$\|F(t) - Y\|_2 \leq R(\delta) e^{-\eta \tilde{\lambda}_{\min}(\delta)t}, \tag{7.45}$$

and by (7.31)

$$\|\hat{K}_{\Theta_t} - \hat{K}_{\Theta_0}\|_{\text{op}} \leq \|\hat{K}_{\Theta_t} - \hat{K}_{\Theta_0}\|_{\text{HS}} \leq 16n \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m), \tag{7.46}$$

we have

$$|\partial_t \Delta(t)| \leq 16nR(\delta)\eta \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m) e^{-\eta \tilde{\lambda}_{\min}(\delta)t} \tag{7.47}$$

whence

$$\Delta(t) \leq \frac{16nR(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \rho(m) \left(1 - e^{-\eta \tilde{\lambda}_{\min}(\delta)t} \right),$$

and this proves the claim (7.39). Now we have all the ingredients to finish the proof. We adapt to our case the strategy of [60]. Let us compute

$$\|\dot{\Theta}_t - \dot{\Theta}_t^{\text{lin}}\|_{\infty} = \eta \|\nabla_{\Theta} f(\Theta_t, X^T)(F(t) - Y) - \nabla_{\Theta} f(\Theta_0, X^T)(F^{\text{lin}}(t) - Y)\|_{\infty}$$

$$\begin{aligned}
&\leq \eta \left\| (\nabla_{\Theta} f(\Theta_t, X^T) - \nabla_{\Theta} f(\Theta_0, X^T)) (F(t) - Y) \right\|_{\infty} \\
&\quad + \eta \left\| \nabla_{\Theta} f(\Theta_0, X^T) (F^{\text{lin}}(t) - F(t)) \right\|_{\infty} \\
&\leq \eta \sup_i \|\partial_{\theta_i} f(\Theta_t, X) - \partial_{\theta_i} f(\Theta_0, X)\|_2 \|F(t) - Y\|_2 \\
&\quad + \eta \sup_i \|\partial_{\theta_i} f(\Theta_0, X)\|_2 \|F^{\text{lin}}(t) - F(t)\|_2
\end{aligned} \tag{7.48}$$

Let us bound the previous expression term by term. Combining the Lipschitzness result of Lemma 7.1 with the lazy training bound (7.36), and using the convergence to the examples (7.33), we control the first term:

$$\begin{aligned}
&\eta \sup_i \|\partial_{\theta_i} f(\Theta_t, X) - \partial_{\theta_i} f(\Theta_0, X)\|_2 \|F(t) - Y\|_2 \\
&\leq 4\eta\sqrt{n}R(\delta) \frac{|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \rho(m) e^{-\eta\tilde{\lambda}_{\min}(\delta)t} \\
&=: A(t).
\end{aligned} \tag{7.49}$$

Regarding the second term, we need two different estimates to be used for “small” and “large” t , as we will show soon. The first estimate is based on the Lipschitzness of the gradient Lemma 7.1 and on (7.39):

$$\begin{aligned}
\eta \sup_i \|\partial_{\theta_i} f(\Theta_0, X)\|_2 \|F^{\text{lin}}(t) - F(t)\|_2 &\leq \frac{32\eta n\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^3|\mathcal{N}|}{N^3(m)} \rho(m) \\
&=: B(t).
\end{aligned} \tag{7.50}$$

The second estimate exploits again Lemma 7.1 and the convergence to the examples for the original model (7.39); we also need to quantify the convergence to the examples for the linearized model; this immediately follows from the analytic solution for the evolution of linearized model (4.10) and from (7.13):

$$\begin{aligned}
&\eta \sup_i \|\partial_{\theta_i} f(\Theta_0, X)\|_2 \|F^{\text{lin}}(t) - F(t)\|_2 \\
&\leq 2\eta\sqrt{n} \frac{|\mathcal{M}|}{N(m)} (\|F^{\text{lin}}(t) - Y\|_2 + \|F(t) - Y\|_2) \\
&\leq 2\eta\sqrt{n} \frac{|\mathcal{M}|}{N(m)} \left(e^{-\eta\lambda_{\min}^K t} \|F(0) - Y\|_2 + R(\delta) e^{-\eta\tilde{\lambda}_{\min}(\delta)t} \right) \\
&\leq 2\eta\sqrt{n} \frac{|\mathcal{M}|}{N(m)} R(\delta) \left(e^{-\eta_0\lambda_{\min}^K t} + e^{-\eta\tilde{\lambda}_{\min}(\delta)t} \right) \\
&\leq 4\eta\sqrt{n}R(\delta) \frac{|\mathcal{M}|}{N(m)} e^{-\eta\tilde{\lambda}_{\min}(\delta)t} =: C(t).
\end{aligned} \tag{7.51}$$

Hence,

$$\|\dot{\Theta}_t - \dot{\Theta}_t^{\text{lin}}\|_{\infty} \leq A(t) + B(t) \quad \text{and} \quad \|\dot{\Theta}_t - \dot{\Theta}_t^{\text{lin}}\|_{\infty} \leq A(t) + C(t). \tag{7.52}$$

Defining

$$t^* = \frac{1}{\eta\tilde{\lambda}_{\min}(\delta)} \log N(m), \tag{7.53}$$

we integrate

$$\begin{aligned}
 \|\Theta_t - \Theta_t^{\text{lin}}\|_\infty &\leq \int_0^\infty A(t)dt + \int_0^{t^*} B(t)dt + \int_{t^*}^\infty C(t)dt \\
 &= 4\sqrt{n}R(\delta) \frac{|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \frac{\rho(m)}{\tilde{\lambda}_{\min}(\delta)} \\
 &\quad + \frac{32n\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^3|\mathcal{N}|}{N^3(m)} \frac{\rho(m)}{\tilde{\lambda}_{\min}(\delta)} \log N(m) \\
 &\quad + \frac{4\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{|\mathcal{M}|}{N(m)} e^{-\log N(m)} \\
 &= \frac{4\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{|\mathcal{M}|}{N(m)} \left\{ \left(|\mathcal{M}||\mathcal{N}| + \frac{8n}{\tilde{\lambda}_{\min}(\delta)} \right. \right. \\
 &\quad \left. \left. \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \log N(m) \right) \rho(m) + \frac{1}{N(m)} \right\}. \tag{7.54}
 \end{aligned}$$

So, using [8, Theorem 4.23] together with (7.36) and (7.54), we conclude by computing

$$\begin{aligned}
 |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)| &= |f(\Theta_t, x) - f(\Theta_0, x) - \nabla_{\Theta} f(\Theta_0, x)^T (\Theta_t^{\text{lin}} - \Theta_0)| \\
 &\leq |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t, x)| + |\nabla_{\Theta} f(\Theta_0, x)^T (\Theta_t^{\text{lin}} - \Theta_t)| \\
 &\leq \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \|\Theta_t - \Theta_0\|_\infty^2 + 2 \frac{Lm|\mathcal{M}|}{N(m)} \|\Theta_t^{\text{lin}} - \Theta_t\|_\infty \\
 &\leq \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N(m)} \rho^2(m) + \frac{8\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^2}{N^2(m)} \\
 &\quad \left\{ \left(|\mathcal{M}||\mathcal{N}| + \frac{8n}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^2|\mathcal{N}|}{N^2(m)} \log N(m) \right) \rho(m) + \frac{1}{N(m)} \right\} \\
 &\leq \frac{20nR^2(\delta)}{\lambda_{\min}^2(\delta)} \frac{Lm|\mathcal{M}|^4|\mathcal{N}|}{N^3(m)} + \frac{128n^2R^2(\delta)}{(\tilde{\lambda}_{\min}(\delta))^3} \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|}{N^5(m)} \log N(m) \\
 &\quad + \frac{8\sqrt{n}R(\delta)}{\tilde{\lambda}_{\min}(\delta)} \frac{Lm|\mathcal{M}|^2}{N^3(m)} =: \xi_m(\delta). \tag{7.55}
 \end{aligned}$$

We can use inequality (6.18) to bound

$$\begin{aligned}
 \xi_m(\delta) &\leq \frac{n^2R^2(\delta)}{\tilde{\lambda}_{\min}(\delta)} \left(8 + \frac{20}{\tilde{\lambda}_{\min}(\delta)} + \frac{128}{(\tilde{\lambda}_{\min}(\delta))^2} \right) \\
 &\quad \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|^2}{N^5(m)} (1 + \log N(m))
 \end{aligned}$$

$$\begin{aligned}
&\leq \frac{4n^2 R^2(\delta)}{\tilde{\lambda}_{\min}(\delta)} \left(2 + \frac{5}{\left(\tilde{\lambda}_{\min}(\delta)\right)^2} + \frac{32}{\left(\tilde{\lambda}_{\min}(\delta)\right)^3} \right) \\
&\quad \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) \\
&\leq 132n^2 R^2(\delta) \left(1 + \frac{1}{\left(\tilde{\lambda}_{\min}(\delta)\right)^3} \right) \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) ,
\end{aligned} \tag{7.56}$$

where we have recalled that $R(\delta) \leq R^2(\delta)$ and, in the last inequality, we have used that, for $x \geq 0$,

$$2x + 5x^2 + 32x^3 \leq 2x + 5x^2 + 32x^3 + (x-1)^2 + (x-4)^2(x+2) = 33(1+x^3). \tag{7.57}$$

This concludes the proof. \square

7.2. The Wasserstein Distance Between f and f^{lin}

In this part, we prove an upper bound to the Wasserstein distance of order 1 between the objective function $f(\Theta_t, x)$ and its linearized version $f^{\text{lin}}(\Theta_t^{\text{lin}}, x)$.

Lemma 7.2. *Suppose that A 1, A 2, A 3, and A 4 hold true. Let \bar{X} be the vector of all the inputs of \mathcal{X} . Then, for any $s > 0$ we have*

$$\begin{aligned}
&d_{\text{W}}^{(s)}(f(\Theta_t, \bar{X}), f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})) \\
&\leq 132n^2 \sqrt{\bar{N}} \left(\|Y\|_2 + \sqrt{2n\sqrt{N(m)}} \right)^2 \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
&\quad \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) + \frac{s}{\sqrt{N(m)}}.
\end{aligned} \tag{7.58}$$

Proof. Note that

$$\begin{aligned}
d_{\text{W}}^{(s)}(f(\Theta_t, \bar{X}), f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})) &\leq \mathbb{E} \min \left\{ \|f(\Theta_t, \bar{X}) - f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})\|_2, s \right\} \\
&\leq \mathbb{E} \min \left\{ \sqrt{\bar{N}} \sup_{x \in \mathcal{X}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)|, s \right\}
\end{aligned}$$

where \bar{N} indicates the number of inputs in \bar{X} . Further, let us take

$$\delta_m := \frac{1}{\sqrt{N(m)}}; \tag{7.59}$$

since (7.6) holds with δ_m , we then apply Theorem 7.1, and so that when applying the gradient flow with learning rate η , the following inequality holds

with probability at least $1 - \delta_m$ over random initialization:

$$\sup_{\substack{x \in \mathcal{X} \\ t \geq 0}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)| \leq 132n^2R^2(\delta_m) \left(1 + \frac{1}{\left(\tilde{\lambda}_{\min}(\delta_m)\right)^3} \right) \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) \quad (7.60)$$

where $\tilde{\lambda}_{\min}(\delta_m)$ is defined according to (7.20), and satisfies (7.7) due to (2.42) in A 4. Let us call \mathcal{E} the event in which (7.60) holds, and let \mathcal{E}^c be its complementary. Clearly $\mathbb{P}(\mathcal{E}) \leq 1$ and $\mathbb{P}(\mathcal{E}^c) \leq \delta_m$, and thus

$$\begin{aligned} & \mathbb{E} \min \left\{ \sqrt{N} \sup_{x \in \mathcal{X}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)|, s \right\} \\ & \leq \mathbb{E} \left(\sqrt{N} \sup_{\substack{x \in \mathcal{X} \\ t \geq 0}} |f(\Theta_t, x) - f^{\text{lin}}(\Theta_t^{\text{lin}}, x)| \middle| \mathcal{E} \right) + \mathbb{E}(s | \mathcal{E}^c) \delta_m. \end{aligned} \quad (7.61)$$

Hence, by applying (7.60) and (7.61), one gets that

$$\begin{aligned} & d_{\text{W}}^{(s)}(f(\Theta_t, \bar{X}), f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})) \\ & \leq 132\sqrt{N}n^2R^2(\delta_m) \left(1 + \frac{1}{\left(\tilde{\lambda}_{\min}(\delta_m)\right)^3} \right) \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) \\ & \quad + \frac{s}{\sqrt{N(m)}} \\ & \leq 132\sqrt{N}n^2R^2(\delta_m) \left(1 + \frac{27}{\left(\lambda_{\min}^K\right)^3} \right) \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|^2}{N^5(m)} \\ & \quad \times (1 + \log N(m)) + \frac{s}{\sqrt{N(m)}}, \end{aligned} \quad (7.62)$$

where the last inequality follows by (7.7). Now by the choice of δ_m and $R(\delta_m)$ defined according to (7.16), we are done. \square

7.3. The Closeness Between f^{lin} and a Gaussian Process

Let us define the following processes:

$$R_t(\bar{X}) := -\hat{K}_{\Theta_0}(\bar{X}, X^T)\hat{K}_{\Theta_0}^{-1} \left(\mathbb{1} - e^{-\eta t \hat{K}_{\Theta_0}} \right) (F(0) - Y), \quad (7.63)$$

$$R_t^\infty(\bar{X}) := -K(\bar{X}, X^T)K^{-1} \left(\mathbb{1} - e^{-\eta t K} \right) (f^{(\infty)}(X) - Y). \quad (7.64)$$

In the next, we show that for any $t > 0$, $f^{\text{lin}}(\Theta_t, \bar{X})$ is close to $Z_t(\bar{X})$ that is defined as

$$Z_t(\bar{X}) := f^{(\infty)}(\bar{X}) - K(\bar{X}, X^T)K^{-1} \left(\mathbb{1} - e^{-\eta t K} \right) (f^{(\infty)}(X) - Y). \quad (7.65)$$

Notice that Z_t follows a normal probability law $\mathcal{N}(\mu_t(\bar{X}), \mathcal{K}_t(\bar{X}, \bar{X}^T))$, where $\mu_t(\bar{X})$ represents its mean vector, and $\bar{\mathcal{K}}_t := \mathcal{K}_t(\bar{X}, \bar{X}^T)$ its covariance operator. These quantities are the vectorial form of (4.17), and (4.16), respectively.

Lemma 7.3. *Let us fix $0 < s < +\infty$. Let \bar{X} be the vector of all the inputs in \mathcal{X} , and let X be the vector of the n training inputs. For any $t > 0$, we set $\bar{\mathcal{K}}_t := \mathcal{K}_t(\bar{X}, \bar{X}^T)$ defined according to (4.16). Suppose that A 1, A 2, and A 3 hold true. Then*

$$\begin{aligned} d_W^{(s)}(f^{\text{lin}}(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) &\leq C_m^{\bar{N}, \bar{\mathcal{K}}_0} + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} C_m^{\bar{N}, \bar{\mathcal{K}}_0} \\ &+ 2 \left(\frac{1}{\sqrt{N(m)} \lambda_{\min}^K} + \frac{6}{\sqrt{N(m)} (\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^2(m)} \right) \mathbb{E} \|F(0) - Y\|_2 \\ &+ \exp \left(-\frac{1}{256} \frac{N^3(m)}{Lm|\mathcal{M}|^2|N|^2} \frac{1}{\bar{N}^2} \right) s \end{aligned} \quad (7.66)$$

where $C_m^{\bar{N}, \bar{\mathcal{K}}_0}$ is defined according to (6.15), and $C_m^{\bar{N}, \bar{\mathcal{K}}_0}$ defined as

$$C_m^{\bar{N}, \bar{\mathcal{K}}_0} := \sqrt{n} \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} \sqrt{\frac{\tilde{D}(2nD)^2}{(N(m))^4} + \frac{Dmn}{(N(m))^2} \omega \left(n, \frac{D}{N(m)} \right)}, \quad (7.67)$$

where D , \tilde{D} , and ω are defined in (2.24), (2.27), and (A.3), respectively.

Proof. Consider $0 < s < +\infty$, and let h be a 1-Lipschitz function. Let us define

$$\begin{aligned} \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) &= f(\Theta_0, \bar{X}) + \tilde{R}_t(\bar{X}), \\ \tilde{R}_t(\bar{X}) &:= -K(\bar{X}, X^T)K^{-1}(\mathbb{1} - e^{-\eta tK})(F(0) - Y). \end{aligned} \quad (7.68)$$

Notice that by the triangle inequality

$$\begin{aligned} d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), Z_t) \\ \leq d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})) + d_W^{(s)}(\tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), Z_t) \end{aligned} \quad (7.69)$$

and so

$$\begin{aligned} d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), Z_t) &\leq \mathbb{E} \min \left\{ \left\| f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) \right\|_2, s \right\} \\ &+ \mathbb{E} \min \left\{ \left\| \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - Z_t \right\|_2, s \right\}. \end{aligned} \quad (7.70)$$

Notice that by the definition of f and \tilde{f}^{lin} , we have

$$\mathbb{E} \min \left\{ \left\| f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) \right\|_2, s \right\} \leq \mathbb{E} \min \left\{ \left\| R_t - \tilde{R}_t \right\|_2, s \right\}. \quad (7.71)$$

In what follows, since η is a constant, we can absorb it into the variable t , and therefore, we will not consider it further. Hence, for each matrix Σ depending

on \bar{X}, X such that the inverse matrix $(\Sigma(X, X^T))^{-1}$ is well defined, we define the matricial function,

$$\begin{aligned} M(\Sigma) &= \Sigma(\bar{X}, X^T)\Sigma^{-1}(\mathbb{1} - e^{-t\Sigma}), \text{ where } \Sigma = \Sigma(X, X^T), \text{ and} \\ \Sigma^{-1} &= (\Sigma(X, X^T))^{-1}. \end{aligned} \quad (7.72)$$

For the sake of simplicity, given a covariance matrix Σ of a set of real-valued random variables indexed by the elements of \mathcal{X} , we write $\Sigma_* := \Sigma(\bar{X}, X^T)$. Then

$$M(K) = K_*K^{-1}(\mathbb{1} - e^{-tK}), \quad M(\hat{K}) = \hat{K}_*\hat{K}^{-1}(\mathbb{1} - e^{-t\hat{K}}), \quad \text{where } \hat{K} = K_{\Theta_0}. \quad (7.73)$$

Notice that

$$\begin{aligned} dM &= d\Sigma_*\Sigma^{-1}(\mathbb{1} - e^{-t\Sigma}) - \Sigma_*\Sigma^{-1}d\Sigma\Sigma^{-1}(\mathbb{1} - e^{-t\Sigma}) \\ &\quad + \Sigma_*\Sigma^{-1} \int_0^1 e^{-s\Sigma t} t d\Sigma e^{-(1-s)\Sigma t} ds \end{aligned} \quad (7.74)$$

Let us denote by λ_{\min}^Σ the smallest eigenvalue of Σ , and suppose that $\Sigma \geq \lambda_{\min}^\Sigma \mathbb{1}$. Notice that

$$\begin{aligned} \|dM\|_{\text{op}} &\leq \|d\Sigma_*\|_{\text{op}} \frac{1}{\lambda_{\min}^\Sigma} + \|\Sigma_*\|_{\text{op}} \|d\Sigma\|_{\text{op}} \frac{1}{(\lambda_{\min}^\Sigma)^2} \\ &\quad + t \|\Sigma_*\|_{\text{op}} \frac{1}{\lambda_{\min}^\Sigma} \|d\Sigma\|_{\text{op}} e^{-t\lambda_{\min}^\Sigma} \\ &\leq \|d\Sigma_*\|_{\text{op}} \frac{1}{\lambda_{\min}^\Sigma} + \|\Sigma_*\|_{\text{op}} \|d\Sigma\|_{\text{op}} \frac{1}{(\lambda_{\min}^\Sigma)^2} \\ &\quad + \frac{1}{(\lambda_{\min}^\Sigma)^2 e} \|\Sigma_*\|_{\text{op}} \|d\Sigma\|_{\text{op}}. \end{aligned} \quad (7.75)$$

On the other hand, by the definition of R_t , and \tilde{R}_t

$$\mathbb{E} \min \left\{ \left\| R_t - \tilde{R}_t \right\|_2, s \right\} = \mathbb{E} \min \left\{ \left\| \left(M(\hat{K}) - M(K) \right) (F(0) - Y) \right\|_2, s \right\}, \quad (7.76)$$

where $\hat{K} = \hat{K}_{\Theta_0}$ is the empirical tangent kernel, and K the analytic tangent kernel. In what follows, fix $0 < \varepsilon$, and we call $\mathcal{E} := \left\{ \omega \in \Omega : \left\| K(\bar{X}, \bar{X}^T) - \hat{K}(\bar{X}, \bar{X}^T) \right\|_{\text{op}} < \varepsilon \right\}$. We write

$$\begin{aligned} &\mathbb{E} \min \left\{ \left\| \left(M(\hat{K}) - M(K) \right) (F(0) - Y) \right\|_2, s \right\} \\ &= \mathbb{P}(\mathcal{E}) \mathbb{E} \left(\min \left\{ \left\| \left(M(\hat{K}) - M(K) \right) (F(0) - Y) \right\|_2, s \right\} \middle| \mathcal{E} \right) \\ &\quad + \mathbb{P}(\mathcal{E}^c) \mathbb{E} \left(\min \left\{ \left\| \left(M(\hat{K}) - M(K) \right) (F(0) - Y) \right\|_2, s \right\} \middle| \mathcal{E}^c \right), \end{aligned} \quad (7.77)$$

and thus

$$\begin{aligned}
& \mathbb{E} \min \left\{ \left\| R_t - \tilde{R}_t \right\|_2, s \right\} \\
& \leq \mathbb{P}(\mathcal{E}) \mathbb{E} \left(\min \left\{ \left\| (M(\hat{K}) - M(K)) (F(0) - Y) \right\|_2, s \right\} \middle| \mathcal{E} \right) + \mathbb{P}(\mathcal{E}^c) s \\
& \leq \mathbb{P}(\mathcal{E}) \mathbb{E} \left(\min \left\{ \left\| (M(\hat{K}) - M(K)) \right\|_{\text{op}} \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right) + \mathbb{P}(\mathcal{E}^c) s.
\end{aligned} \tag{7.78}$$

Let us take $\alpha \in [0, 1]$. By the fundamental theorem of calculus, we write

$$M(K) - M(\hat{K}) = \int_0^1 \frac{dM(\Sigma_\alpha)}{d\alpha} d\alpha, \quad \Sigma_\alpha := \alpha K + (1 - \alpha) \hat{K}. \tag{7.79}$$

By (7.75) we have that

$$\begin{aligned}
& \left\| M(K) - M(\hat{K}) \right\|_{\text{op}} \\
& \leq \left\| \frac{d(\Sigma_\alpha)_*}{d\alpha} \right\|_{\text{op}} \frac{1}{\lambda_{\min}^K} + \|K_*\|_{\text{op}} \left\| \frac{d\Sigma_\alpha}{d\alpha} \right\|_{\text{op}} \frac{1}{(\lambda_{\min}^K)^2} \\
& \quad + \frac{1}{(\lambda_{\min}^K)^2 e} \|K_*\|_{\text{op}} \left\| \frac{d\Sigma_\alpha}{d\alpha} \right\|_{\text{op}} \\
& = \frac{1}{\lambda_{\min}^K} \left\| K(\bar{X}, X^T) - K_{\Theta_0}(\bar{X}, X^T) \right\|_{\text{op}} + \frac{1}{(\lambda_{\min}^K)^2} \|K_*\|_{\text{op}} \left\| K - \hat{K} \right\|_{\text{op}} \\
& \quad + \frac{1}{(\lambda_{\min}^K)^2 e} \|K_*\|_{\text{op}} \left\| K - \hat{K} \right\|_{\text{op}}.
\end{aligned} \tag{7.80}$$

Hence, we obtain that

$$\begin{aligned}
& \mathbb{E} \left(\min \left\{ \left\| M(\hat{K}) - M(K) \right\|_{\text{op}} \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right) \\
& \leq \frac{\varepsilon}{\lambda_{\min}^K} \mathbb{E} \left(\min \left\{ \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right)
\end{aligned} \tag{7.81}$$

$$\begin{aligned}
& + \frac{\varepsilon}{(\lambda_{\min}^K)^2} \mathbb{E} \left(\min \left\{ \|K_*\|_{\text{op}} \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right) \\
& + \frac{\varepsilon}{(\lambda_{\min}^K)^2 e} \mathbb{E} \left(\min \left\{ \|K_*\|_{\text{op}} \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right).
\end{aligned} \tag{7.82}$$

In order to bound $\|K_*\|_{\text{op}}$ and $\|\hat{K}_*\|_{\text{op}}$, let us first notice that, for a generic real $n_1 \times n_2$ matrix A such that $|A_{ij}| \leq a$ for any $1 \leq i \leq n_1$ and $1 \leq j \leq n_2$,

by convexity of $x \mapsto x^2$ we have

$$\begin{aligned} \|A\|_{\text{op}} &= \sup_{\|v\| \leq 1} \sqrt{\sum_{i=1}^{n_1} \left(\sum_{j=1}^{n_2} A_{ij} v_j \right)^2} \leq \sup_{\|v\| \leq 1} \sqrt{n_2 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} A_{ij}^2 v_j^2} \\ &\leq a \sup_{\|v\| \leq 1} \sqrt{n_1 n_2 \sum_{j=1}^{n_2} v_j^2} \leq a \sqrt{n_1 n_2} \end{aligned} \quad (7.83)$$

where $v \in \mathbb{R}^{n_2}$. A uniform bound on the matrix elements of $\hat{K}_{\mathcal{D}}$ can be obtained as follows. By Lemma 7.1, for any $1 \leq i \leq Lm$ and $x \in \mathcal{X}$ we have

$$|\partial_{\theta_i} f(\Theta_0, x)| \leq 2 \frac{|\mathcal{M}|}{N(m)}; \quad (7.84)$$

therefore,

$$\begin{aligned} |\hat{K}_{\Theta_0}(x, x')| &= \left| \sum_{j=1}^{Lm} \partial_{\theta_j} f(\Theta_0, x) \partial_{\theta_j} f(\Theta_0, x') \right| \leq \sum_{j=1}^{Lm} |\partial_{\theta_j} f(\Theta_0, x)| |\partial_{\theta_j} f(\Theta_0, x')| \\ &\leq 4 \frac{Lm |\mathcal{M}|^2}{N^2(m)} \end{aligned} \quad (7.85)$$

whence

$$\|\hat{K}_*\|_{\text{op}} \leq 4\sqrt{\bar{N}n} \frac{Lm |\mathcal{M}|^2}{N^2(m)} \quad \text{and} \quad \|K_*\|_{\text{op}} \leq \mathbb{E}[\|\hat{K}_{\mathcal{D}}\|_{\text{op}}] \leq 4\sqrt{\bar{N}n} \frac{Lm |\mathcal{M}|^2}{N^2(m)}. \quad (7.86)$$

Hence, we obtain that

$$\begin{aligned} &\mathbb{E} \left(\min \left\{ \|M(\hat{K}) - M(K)\|_{\text{op}}, \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right) \\ &\leq \frac{\varepsilon}{\lambda_{\min}^K} \mathbb{E} \left(\min \{ \|F(0) - Y\|_2, s \} \middle| \mathcal{E} \right) \\ &\quad + \frac{3}{2} \frac{\varepsilon}{(\lambda_{\min}^K)^2} \mathbb{E} \left(\min \left\{ 4\sqrt{\bar{N}n} \frac{Lm |\mathcal{M}|^2}{N^2(m)} \|F(0) - Y\|_2, s \right\} \middle| \mathcal{E} \right) \\ &\leq \left(\frac{\varepsilon}{\lambda_{\min}^K} + \frac{6\varepsilon}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm |\mathcal{M}|^2}{N^2(m)} \right) \mathbb{E} (\|F(0) - Y\|_2 | \mathcal{E}). \end{aligned} \quad (7.87)$$

Let us now compute $\mathbb{P}(\mathcal{E}^c)$. Proceeding as in (7.83) we get

$$\|K(\bar{X}, \bar{X}^T) - \hat{K}(\bar{X}, \bar{X}^T)\|_{\text{op}} \leq \bar{N} \sup_{x, x' \in \mathcal{X}} |K(x, x') - \hat{K}(x, x')|. \quad (7.88)$$

Then, by Theorem 7.2 and (7.88), one gets

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &\leq \mathbb{P}\left(\bar{N} \sup_{x, x' \in \mathcal{X}} |K(x, x') - \hat{K}(x, x')| \geq \varepsilon\right) \\ &= \mathbb{P}\left(\sup_{x, x' \in \mathcal{X}} |K(x, x') - \hat{K}(x, x')| \geq \frac{\varepsilon}{\bar{N}}\right) \\ &\leq \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{\varepsilon^2}{\bar{N}^2}\right). \end{aligned} \tag{7.89}$$

Therefore, by (7.71) together with (7.87) and (7.89), one has

$$\begin{aligned} &\mathbb{E} \min \left\{ \left\| f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) \right\|_2, s \right\} \\ &\leq \left(\frac{\varepsilon}{\lambda_{\min}^K} + \frac{6\varepsilon}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^2(m)} \right) \mathbb{E} (\|F(0) - Y\|_2 | \mathcal{E}) \\ &\quad + \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{\varepsilon^2}{\bar{N}^2}\right) s. \end{aligned} \tag{7.90}$$

Let us now estimate the remaining term $\mathbb{E} \min \left\{ \left\| \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - Z_t \right\|_2, s \right\}$. Notice that

$$\begin{aligned} \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - Z_t &= f(\Theta_0, \bar{X}) - f^{(\infty)}(\bar{X}) \\ &\quad - \left(K(\bar{X}, X^T) K^{-1} (\mathbb{1} - e^{-\eta t K}) (f(\Theta_0, X) - f^{(\infty)}(X)) \right), \end{aligned} \tag{7.91}$$

and thus

$$\begin{aligned} &\mathbb{E} \min \left\{ \left\| \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - Z_t \right\|_2, s \right\} \\ &\leq \mathbb{E} \left\| f(\Theta_0, \bar{X}) - f^{(\infty)}(\bar{X}) \right\|_2 + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} \mathbb{E} \left\| f(\Theta_0, X) - f^{(\infty)}(X) \right\|_2 \\ &= d_W(f(\Theta_0, \bar{X}), \mathcal{N}(0, \bar{\mathcal{K}}_0)) + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} d_W(f(\Theta_0, X), \mathcal{N}(0, \bar{\mathcal{K}}_0)) \end{aligned} \tag{7.92}$$

where the last identity follows by the existence of an optimal plan. Then by (6.16), one gets that

$$\mathbb{E} \min \left\{ \left\| \tilde{f}^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}) - Z_t \right\|_2, s \right\} \leq C_m^{\bar{N}, \bar{\mathcal{K}}_0} + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} C_m^{n, \bar{\mathcal{K}}_0}. \tag{7.93}$$

Therefore, combining (7.90) and (7.93) with (7.70), we obtain

$$\begin{aligned}
 & d_W^{(s)}(f^{\text{lin}}(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
 & \leq C_m^{\bar{N}, \bar{\mathcal{K}}_0} + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} C_m^{\bar{N}, \bar{\mathcal{K}}_0} \\
 & \quad + \left(\frac{\varepsilon}{\lambda_{\min}^K} + \frac{6\varepsilon}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^2(m)} \right) \mathbb{E}(\|F(0) - Y\|_2 | \mathcal{E}) \\
 & \quad + \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{\varepsilon^2}{\bar{N}^2}\right) s. \tag{7.94}
 \end{aligned}$$

Let us choose $\varepsilon = \frac{1}{\sqrt{N(m)}}$. Notice that

$$\mathbb{E}(\|F(0) - Y\|_2 | \mathcal{E}) \leq \frac{1}{\mathbb{P}(\mathcal{E})} \mathbb{E}\|F(0) - Y\|_2 \leq 2\mathbb{E}\|F(0) - Y\|_2, \tag{7.95}$$

where

$$\mathbb{P}(\mathcal{E}) = 1 - \mathbb{P}(\mathcal{E}^c) \geq 1 - \exp\left(-\frac{1}{256} \frac{N^3(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{1}{\bar{N}^2}\right) \geq \frac{1}{2} \tag{7.96}$$

which follows from A 4, and our conclusion follows. \square

We now are in position to prove Theorem 5.2.

Proof of Theorem 5.2. Let us fix $s > 0$.

Notice that thanks to Lemma 7.2

$$\begin{aligned}
 & d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
 & \leq d_W^{(s)}(f(\Theta_t, \bar{X}), f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X})) + d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
 & \leq 132n^2\sqrt{\bar{N}} \left(\|Y\|_2 + \sqrt{2n\sqrt{N(m)}} \right)^2 \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
 & \quad + \frac{L^2m^2|\mathcal{M}|^5|\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) + \frac{s}{\sqrt{N(m)}} \\
 & \quad + d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)). \tag{7.97}
 \end{aligned}$$

On the other hand, by Theorem 7.3, we obtain that

$$\begin{aligned}
 & d_W^{(s)}(f^{\text{lin}}(\Theta_t^{\text{lin}}, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
 & \leq C_m^{\bar{N}, \bar{\mathcal{K}}_0} + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} C_m^{\bar{N}, \bar{\mathcal{K}}_0} \\
 & \quad + 2 \left(\frac{1}{\sqrt{N(m)}\lambda_{\min}^K} + \frac{6}{\sqrt{N(m)}(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^2(m)} \right) \mathbb{E}[\|F(0) - Y\|_2] \\
 & \quad + \exp\left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{1}{N(m)\bar{N}}\right) s \tag{7.98}
 \end{aligned}$$

Hence, combining (7.97), and (7.98) we get

$$\begin{aligned}
& d_{\mathbb{W}}^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
& \leq \frac{s}{\sqrt{N(m)}} + C_m^{\bar{N}, \bar{\mathcal{K}}_0} + \frac{\|K(\bar{X}, X^T)\|_{\text{op}}}{\lambda_{\min}^K} C_m^{n, \bar{\mathcal{K}}_0} \\
& \quad + 132n^2 \sqrt{\bar{N}} \left(\|Y\|_2 + \sqrt{2n\sqrt{N(m)}} \right)^2 \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
& \quad \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) \\
& \quad + 2 \left(\frac{1}{\sqrt{N(m)} \lambda_{\min}^K} + \frac{6}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^{5/2}(m)} \right) \mathbb{E}[\|F(0) - Y\|_2] \\
& \quad + \exp \left(-\frac{1}{256} \frac{N^4(m)}{Lm|\mathcal{M}|^2 |\mathcal{N}|^2} \frac{1}{N(m)\bar{N}} \right) s. \tag{7.99}
\end{aligned}$$

We recall that

$$C_m^{n, \bar{\mathcal{K}}_0} := \sqrt{n} \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} \sqrt{\frac{\tilde{D}(2nD)^2}{(N(m))^4}} + \frac{Dmn}{(N(m))^2} \omega \left(n, \frac{D}{N(m)} \right), \tag{7.100}$$

where we recall that

$$\omega(d, x) := \begin{cases} x(C(1, d) - 2 \log x) & \text{if } x \leq 1, \\ C(1, d) & \text{if } x > 1. \end{cases} \quad \text{with } C(1, d) := 2^{\frac{3}{2}} \frac{1 + 2d \Gamma(\frac{1+d}{2})}{d \Gamma(\frac{d}{2})}. \tag{7.101}$$

We notice that by proceeding as in (6.23), we also obtain that

$$C_m^{n, \bar{\mathcal{K}}_0} \leq n^2 \left(2 \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \frac{m|\mathcal{M}|^{7/2} |\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)). \tag{7.102}$$

Furthermore, in order to simplify (7.99), we recall (7.86)

$$\|K(\bar{X}, X^T)\|_{\text{op}} \leq 4\sqrt{\bar{N}n} \frac{Lm|\mathcal{M}|^2}{N^2(m)}, \tag{7.103}$$

we upper bound

$$\left(\|Y\|_2 + \sqrt{2n\sqrt{N(m)}} \right)^2 \leq 2\|Y\|_2^2 + 4n\sqrt{N(m)} \leq (2\|Y\|_2^2 + 4n) \sqrt{N(m)} \tag{7.104}$$

and, by virtue of A 2, we compute

$$\begin{aligned}
 \mathbb{E}[\|F(0) - Y\|_2]^2 &\leq \mathbb{E}[\|F(0) - Y\|_2^2] \\
 &= \sum_{i=1}^n \mathbb{E} \left[\left(f(\Theta_0, x^{(i)}) - y^{(i)} \right)^2 \right] \\
 &= \sum_{i=1}^n \left(\mathcal{K}_0(x^{(i)}, x^{(i)}) + (y^{(i)})^2 \right) \\
 &\leq n + \|Y\|_2^2,
 \end{aligned} \tag{7.105}$$

whence

$$\mathbb{E}[\|F(0) - Y\|_2] \leq \sqrt{n + \|Y\|_2^2} \leq \sqrt{n} + \|Y\|_2. \tag{7.106}$$

Now, we know that

$$xe^{-x} \leq \frac{1}{e} \tag{7.107}$$

whence

$$e^{-x} \leq \frac{1}{e} \frac{1}{x} \tag{7.108}$$

holds. Therefore,

$$\exp \left(-\frac{1}{256} \frac{N^3(m)}{Lm|\mathcal{M}|^2|\mathcal{N}|^2} \frac{1}{N^2} \right) \leq \frac{256}{e} \frac{N^2}{N^3(m)} \frac{Lm|\mathcal{M}|^2|\mathcal{N}|^2}{N^3(m)}. \tag{7.109}$$

The last ingredient is again (6.18):

$$N(m) \leq \sqrt{m|\mathcal{M}||\mathcal{N}|}. \tag{7.110}$$

Then, it's laborious but easy to show that

$$\begin{aligned}
 &d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t(\bar{X}), \bar{\mathcal{K}}_t)) \\
 &\leq \frac{s}{\sqrt{N(m)}} + N^2 \left(2 \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \frac{m|\mathcal{M}|^{7/2}|\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)) \\
 &\quad + 4\sqrt{N}n \frac{Lm|\mathcal{M}|^2}{N^2(m)} n^2 \left(2 \|\bar{\mathcal{K}}_0^{-1}\|_{\text{op}} \|\bar{\mathcal{K}}_0\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \\
 &\quad \frac{m|\mathcal{M}|^{7/2}|\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)) \\
 &\quad + 132n^2 \sqrt{N} (2\|Y\|_2^2 + 4n) \sqrt{N(m)} \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
 &\quad \frac{L^2 m^2 |\mathcal{M}|^5 |\mathcal{N}|^2}{N^5(m)} (1 + \log N(m)) \\
 &\quad + 2 \left(\frac{1}{\sqrt{N(m)} \lambda_{\min}^K} + \frac{6}{(\lambda_{\min}^K)^2} \sqrt{N}n \frac{Lm|\mathcal{M}|^2}{N^{5/2}(m)} \right) (\sqrt{n} + \|Y\|_2) \\
 &\quad + \frac{256}{e} s N^2 \frac{Lm|\mathcal{M}|^2|\mathcal{N}|^2}{N^3(m)}
 \end{aligned}$$

$$\begin{aligned}
&\leq s \frac{m^{9/4} |\mathcal{M}|^{9/4} |\mathcal{N}|^{9/4}}{N^5(m)} + \bar{N}^2 \\
&\times \left(2 \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \frac{m^2 |\mathcal{M}|^{9/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)) \\
&+ 4n^{5/2} \sqrt{\bar{N}} \left(2 \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} + 6 \right) \frac{Lm^2 |\mathcal{M}|^{11/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)) \\
&+ 132n^2 \sqrt{\bar{N}} (2 \|Y\|_2^2 + 4n) \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
&\frac{L^2 m^{9/4} |\mathcal{M}|^{21/4} |\mathcal{N}|^{9/4}}{N^5(m)} (1 + \log N(m)) \\
&+ 2 \left(\frac{1}{\lambda_{\min}^K} \frac{m^{9/4} |\mathcal{M}|^{9/4} |\mathcal{N}|^{9/4}}{N^5(m)} + \frac{6}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \frac{Lm^{9/4} |\mathcal{M}|^{13/4} |\mathcal{N}|^{5/4}}{N^5(m)} \right) \\
&(\sqrt{n} + \|Y\|_2) + \frac{256}{e} s \bar{N}^2 \frac{Lm^2 |\mathcal{M}|^3 |\mathcal{N}|^3}{N^5(m)} \\
&\leq C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \bar{\mathcal{K}}_0, s) \frac{L^2 m^{9/4} |\mathcal{M}|^{11/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)) \quad (7.111)
\end{aligned}$$

where

$$\begin{aligned}
C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \bar{\mathcal{K}}_0, s) &:= 96s\bar{N}^2 + 8\sqrt{\bar{N}} \left\| \bar{\mathcal{K}}_0^{-1} \right\|_{\text{op}} \left\| \bar{\mathcal{K}}_0 \right\|_{\text{op}}^{\frac{1}{2}} (\bar{N}^{3/2} + 4n^{5/2}) \\
&+ 132n^2 \sqrt{\bar{N}} (2 \|Y\|_2^2 + 4n) \left(1 + \frac{27}{(\lambda_{\min}^K)^3} \right) \\
&+ 2 (\sqrt{n} + \|Y\|_2) \left(\frac{1}{\lambda_{\min}^K} + \frac{6}{(\lambda_{\min}^K)^2} \sqrt{\bar{N}n} \right). \quad (7.112)
\end{aligned}$$

□

8. Conclusions

We have studied quantum neural networks whose generated function is given by the expectation value of the sum of single-qubit observables on the state generated by a parametric quantum circuit (A 1). Our main result is a quantitative proof of the convergence to a Gaussian process of the probability distribution of the functions generated by such networks in the limit of infinite width for both untrained networks with randomly initialized parameters and trained networks.

First, we have considered the probability distribution of the function generated by a quantum neural network with untrained randomly initialized

parameters. In Theorem 5.1, we have provided an upper bound to the Wasserstein distance of order 1 between such probability distribution and the Gaussian distribution with the same covariance.

Then, we have considered the probability distribution of the function generated by the network during the training dynamics via gradient flow. In Theorem 5.2, we have proved an upper bound to a truncated version of the Wasserstein distance of order 1 between such probability distribution and a suitable Gaussian process. The proof of Theorem 5.2 is based on Theorem 7.1, which states that, for large enough width, the generated function can be approximated by truncating to the first order its Taylor series with respect to the parameters of the model around their initialization value, implying that each parameter remains with high probability close to its initialization value and that the training happens in the *lazy regime*.

Both upper bounds Theorem 5.1 and Theorem 7.1 tend to zero in the limit of infinite width, provided that the network does not suffer from barren plateaus. Therefore, our results provide a quantitative proof of the convergence in distribution to a Gaussian process proved in [8].

While [8] requires to build a sequence of networks with diverging width, our results can be applied to any quantum neural network with finite width. Moreover, [8] proves the convergence in distribution to a Gaussian process only at fixed training time, and such result holds in the limit $t \rightarrow \infty$ only if such limit is taken after the limit of infinite width. Instead, all the upper bounds proved in this work do not depend on the training time and hold in the limit $t \rightarrow \infty$ for finite width.

Our results open the way to several possible research directions:

- It would be interesting to extend our results to the case of an infinite input space, such as a subset of \mathbb{R}^d . In particular, we would like to generalize the functional bounds for the Wasserstein distance for classical neural networks, see for instance [61, Theorem 1.1], and [62, Theorem 3.12].
- It would also be interesting to determine lower bounds to the distance between the probability distribution of the function generated by the trained network and the corresponding Gaussian process, or at least to determine the asymptotic scaling of such corrections in the limit of infinite width, providing a quantum generalization of the classical results in this direction [63–66].

Acknowledgements

GDP has been supported by the HPC Italian National Centre for HPC, Big Data and Quantum Computing – Proposal code CN00000013 – CUP J33C220 01170001 and by the Italian Extended Partnership PE01 – FAIR Future Artificial Intelligence Research – Proposal code PE00000013 – CUP J33C2200283000 6 under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU. Funded by the European Union – NextGenerationEU under the National Recovery and Resilience Plan (PNRR) – Mission

4 Education and research – Component 2 From research to business – Investment 1.1 Notice Prin 2022 – DD N. 104 del 2/2/2022, from title “understanding the LEarning process of QUantum Neural networks (LeQun)”, proposal code 2022WHZ5XH – CUP J53D23003890006. DP has been supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU. GDP and DP are members of the “Gruppo Nazionale per la Fisica Matematica (GNFM)” of the “Istituto Nazionale di Alta Matematica “Francesco Severi” (INdAM)”. AMH has been supported by project PRIN 2022 “understanding the LEarning process of QUantum Neural networks (LeQun)”, proposal code 2022WHZ5XH – CUP J53D23003890006. The author AMH is a member of the “Gruppo Nazionale per l’Analisi Matematica, la Probabilità e le loro Applicazioni (GNAMPA)” of the “Istituto Nazionale di Alta Matematica “Francesco Severi” (INdAM)”.

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement.

Declarations

Conflict of interest The authors confirm that there are no conflict of interest.

Ethical Approval The authors confirm that our research paper does not involve human participants and/or animals as a dataset.

Open Access. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

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

Appendix A. Regularity of the Solutions of the Stein Equation

In what follows, we recall some recent result about the regularity of the solutions of the Stein equation (4.4) obtained in [36]. Let us fix $d \geq 1$, and suppose that $\Sigma = \mathbb{1}_{d \times d}$ is identity matrix in \mathbb{R}^d . Given $h : \mathbb{R}^d \rightarrow \mathbb{R}$ a 1-Lipschitz function, we consider f_h as a solution of the Stein equation (4.4).

Proposition A.1. *Let $h : \mathbb{R}^d \rightarrow \mathbb{R}$ be a 1-Lipschitz function. Then the solution f_h of the Stein equation (4.4) satisfies,*

$$\|\text{Hessian} f_h(x) - \text{Hessian} f_h(y)\|_2$$

$$\leq \begin{cases} \|x - y\|_2 (C(1, d) - 2 \log(\|x - y\|_2)), & \text{if } \|x - y\|_2 \leq 1; \\ C(1, d) & \text{if } \|x - y\|_2 > 1. \end{cases} \quad (\text{A.1})$$

Here, we have denoted by $C(1, d)$ the constant given by

$$C(1, d) := 2^{\frac{3}{2}} \frac{1 + 2d \Gamma\left(\frac{1+d}{2}\right)}{d \Gamma\left(\frac{d}{2}\right)}. \quad (\text{A.2})$$

Proof. This result was proved in [36, Proposition 2.3] in a more general case. \square

In the next, let us define $\omega(d, \cdot) : [0, \infty) \rightarrow [0, \infty)$ the following function:

$$\omega(d, x) := \begin{cases} x (C(1, d) - 2 \log x), & \text{if } x \leq 1, \\ C(1, d), & \text{if } x > 1. \end{cases} \quad (\text{A.3})$$

Notice that $\omega(d, 0) = 0$, $\lim_{x \rightarrow 0^+} \omega(d; x) = 0$, and it is non-decreasing, then it is a modulus of continuity.

Appendix B. Convergence Toward a Gaussian Process

In this Appendix, we briefly summarize some of the key results established in [8]. There, the authors consider a sequence of quantum neural networks with diverging width satisfying A 1-A 4 and the following further assumptions:

A 5. Suppose that the covariance of the generated function at initialization converges uniformly:

$$\lim_{m \rightarrow +\infty} \sup_{x, x' \in \mathcal{X}} |\mathbb{E}(f(\Theta, x)f(\Theta, x')) - \mathcal{K}_0^\infty(x, x')| = 0, \quad (\text{B.1})$$

where $\mathcal{K}_0^\infty : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel such that

$$\mathcal{K}_0^\infty(x, x) > 0 \text{ for all } x \in \mathcal{X}. \quad (\text{B.2})$$

A 6. Suppose that $N(m)$ grows sufficiently fast such that

$$\lim_{m \rightarrow +\infty} \frac{m |\mathcal{M}|^2 |\mathcal{N}|^2}{N^3(m)} = 0. \quad (\text{B.3})$$

A 7. Alternatively, suppose that

$$\lim_{m \rightarrow +\infty} \frac{L^2 m^2 |\mathcal{M}|^6 |\mathcal{N}|^3}{N^5(m)} \log N(m) = 0. \quad (\text{B.4})$$

A 8. Suppose that there exists a positive constant $N_K(m)$ depending on the number of qubits m and a kernel $K^\infty : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that the analytic neural tangent kernel rescaled by $N_K(m)$ converges uniformly to K^∞ :

$$\lim_{m \rightarrow +\infty} \sup_{x, x' \in \mathcal{X}} \left| \frac{K(x, x')}{N_K(m)} - K^\infty(x, x') \right| = 0. \quad (\text{B.5})$$

Further, suppose that the minimum eigenvalue of $K^\infty = K^\infty(X, X^T)$ is strictly positive, and denote it by λ_{\min}^∞ .

The following is one of the main achievements of [8].

Theorem B.1. [8, Theorem 3.7] *Suppose that A 1-A 2, A 5-A 6 hold true. Then as $m \rightarrow +\infty$, the probability distribution of the generated function at initialization converges in distribution to the Gaussian process with mean zero and covariance \mathcal{K}_0^∞ satisfying our assumption A 5.*

We notice that Theorem B.1 does not provide any convergence rate. Next, we prove that Theorem B.1 can be derived from Theorem 5.1.

Corollary B.1. *Let us suppose that A 1, A 2, and*

$$\lim_{m \rightarrow +\infty} \frac{m|\mathcal{M}|^{7/2}|\mathcal{N}|^{7/2}}{N^3(m)} (1 + \log N(m)) = 0. \quad (\text{B.6})$$

Then

$$\lim_{m \rightarrow +\infty} d_W(f(\Theta, \bar{X}), \mathcal{N}(0, \mathcal{K}_0^\infty)) = 0, \quad (\text{B.7})$$

so that $x \mapsto f(\Theta, x)$ converges in distribution to a Gaussian process with mean zero, and covariance operator \mathcal{K}_0^∞ .

Proof. Follows from Theorem 5.1 and (B.6). \square

Let us notice that to establish the convergence result Theorem B.1, in [8], the authors used the Lévy's continuity theorem combined with an appropriate upper bound on the cumulants of f . This approach was necessary because the random variables $\{f_k\}_{k=1}^m$ defined in (2.17) exhibit weak dependence, meaning the standard central limit theorem cannot be directly applied [67, 68]. Nevertheless, they were able to impose the weaker convergence assumption A 6 instead of (B.6). Next, let us recall another remarkable convergence result proved in [8] regarding the behavior of the trained objective function $f(\Theta_t, \cdot)$.

Theorem B.2. *Let us consider a sequence of quantum neural networks trained via gradient flow with learning rate $\frac{\eta}{N_K(m)}$*

$$\begin{cases} \frac{d\Theta_t}{dt} = -\frac{\eta}{N_K(m)} \nabla_{\Theta} f(\Theta_t, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t), \\ \frac{d}{dt} f(\Theta_t, x) = -\frac{\eta}{N_K(m)} (\nabla_{\Theta} f(\Theta_t, x))^T \nabla_{\Theta} f(\Theta_t, X^T) \nabla_{f(\Theta_t, X)} \mathcal{L}(\Theta_t), \end{cases} \quad (\text{B.8})$$

and satisfying A 1, A 2, A 5, 7, A 8. Then, for any $t \geq 0$, in the limit of infinitely many qubits $m \rightarrow +\infty$, $\{f(\Theta_t, x)\}_{x \in \mathcal{X}}$ converges in distribution to a Gaussian process $\{f_t^{(\infty)}(x)\}_{x \in \mathcal{X}}$ with mean and covariance

$$\begin{aligned} \mathcal{K}_t^\infty(x, x') &:= \mathcal{K}_0^\infty(x, x') - K^\infty(x, X^T)(K^\infty)^{-1} \left(\mathbb{1} - e^{-t\eta K^\infty} \right) \mathcal{K}_0^\infty(X, x') \\ &\quad - K^\infty(x', X^T)(K^\infty)^{-1} \left(\mathbb{1} - e^{-t\eta K^\infty} \right) \mathcal{K}_0^\infty(X, x) \\ &\quad + K^\infty(x, X^T)(K^\infty)^{-1} \left(\mathbb{1} - e^{-t\eta K^\infty} \right) \mathcal{K}_0^\infty(X, X^T) \\ &\quad \left(\mathbb{1} - e^{-t\eta K^\infty} \right) (K^\infty)^{-1} K^\infty(X, x'), \end{aligned} \quad (\text{B.9})$$

$$\mu_t^\infty(x) := K^\infty(x, X^T)(K^\infty)^{-1} \left(\mathbb{1} - e^{-t\eta K^\infty} \right) Y, \quad (\text{B.10})$$

where $K^\infty := K^\infty(X, X^T)$.

Next, we prove that the conclusion of Theorem B.2 can be obtained from Theorem 5.2.

Corollary B.2. *Let us consider a sequence of quantum neural networks trained via the gradient flow (B.8) and asymptotically satisfying A 1, A 2, A 5, A 8, and*

$$\lim_{m \rightarrow +\infty} \frac{L^2 m^{9/4} |\mathcal{M}|^{11/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)) = 0. \quad (\text{B.11})$$

Then, for any $s > 0$ and any $t > 0$, we have

$$\lim_{m \rightarrow +\infty} d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty)) = 0, \quad (\text{B.12})$$

so that the trained quantum neural function $x \mapsto f(\Theta_t, x)$ converges in distribution to a Gaussian process with mean μ_t^∞ , and covariance operator \mathcal{K}_t^∞ defined by (B.10), and (B.9), respectively.

Proof. Theorem 5.2 provides a bound on the truncated W_1 distance between the function generated by the trained network at finite width and the Gaussian process with mean and covariance

$$\begin{aligned} \tilde{\mathcal{K}}_t(x, x') &:= \mathcal{K}_0(x, x') - \tilde{K}(x, X^T) \tilde{K}^{-1} \left(\mathbb{1} - e^{-t\eta \tilde{K}} \right) \mathcal{K}_0(X, x') \\ &\quad - \tilde{K}(x', X^T) \tilde{K}^{-1} \left(\mathbb{1} - e^{-t\eta \tilde{K}} \right) \mathcal{K}_0(X, x) \\ &\quad + \tilde{K}(x, X^T) \tilde{K}^{-1} \left(\mathbb{1} - e^{-t\eta \tilde{K}} \right) \mathcal{K}_0(X, X^T) \\ &\quad \left(\mathbb{1} - e^{-t\eta \tilde{K}} \right) \tilde{K}^{-1} K(X, x'), \end{aligned} \quad (\text{B.13})$$

$$\tilde{\mu}_t(x) := \tilde{K}(x, X^T) \tilde{K}^{-1} \left(\mathbb{1} - e^{-t\eta \tilde{K}} \right) Y, \quad (\text{B.14})$$

where

$$\tilde{K}(x, x') := \frac{K(x, x')}{N_K(m)} \quad \forall x, x' \in \mathcal{X}. \quad (\text{B.15})$$

Notice that by the triangle inequality, we have that

$$\begin{aligned} d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty)) &\leq d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\tilde{\mu}_t(\bar{X}), \tilde{\mathcal{K}}_t)) \\ &\quad + d_W^{(s)}(\mathcal{N}(\tilde{\mu}_t(\bar{X}), \tilde{\mathcal{K}}_t), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty)). \end{aligned} \quad (\text{B.16})$$

Furthermore, since (B.11) implies A 4, by Theorem 5.2, and by [8, Remark 4.7], we simply let $N_K(m) \geq 1$, and then we still have that

$$d_W^{(s)}(f(\Theta_t, \bar{X}), \mathcal{N}(\tilde{\mu}_t(\bar{X}), \tilde{\mathcal{K}}_t)) \leq C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \mathcal{K}_0, s)$$

$$\frac{L^2 m^{9/4} |\mathcal{M}|^{11/2} |\mathcal{N}|^{9/2}}{N^5(m)} (1 + \log N(m)), \quad (\text{B.17})$$

where $C(\bar{N}, n, \|Y\|_2, \lambda_{\min}^K, \bar{\mathcal{K}}_0, s)$ is the positive constant given by (5.3). Then by (B.11) we obtain that (B.17) tends to zero as m goes to infinity. On the other hand, notice that

$$d_W^{(s)}(\mathcal{N}(\mu_t(\bar{X}), \mathcal{K}_t), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty)) \leq d_{W_2}(\mathcal{N}(\mu_t(\bar{X}), \mathcal{K}_t), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty)) \quad (\text{B.18})$$

where d_{W_2} is the Wasserstein distance of order 2. Moreover, by [69, Equation (4)]

$$\begin{aligned} & \left(d_{W_2} \left(\mathcal{N}(\tilde{\mu}_t(\bar{X}), \tilde{\mathcal{K}}_t), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty) \right) \right)^2 \\ &= \|\tilde{\mu}_t(\bar{X}) - \mu_t^\infty(\bar{X})\|_2^2 + \text{tr} \left(\tilde{\mathcal{K}}_t + \mathcal{K}_t^\infty - 2(\tilde{\mathcal{K}}_t \mathcal{K}_t^\infty)^{\frac{1}{2}} \right). \end{aligned} \quad (\text{B.19})$$

Then, by applying A 8, we conclude that

$$\lim_{m \rightarrow +\infty} d_{W_2} \left(\mathcal{N}(\tilde{\mu}_t(\bar{X}), \tilde{\mathcal{K}}_t), \mathcal{N}(\mu_t^\infty(\bar{X}), \mathcal{K}_t^\infty) \right) = 0. \quad (\text{B.20})$$

Therefore, by combining (B.6) with (B.16), (B.17), (B.20), we obtain our desired convergence. \square

References

- [1] Berlyand, L., Jabin, P.-E.: Mathematics of deep learning: an introduction. de Gruyter (2023)
- [2] Russell, S.J., Norvig, P.: Artificial intelligence: a modern approach. Pearson (2016)
- [3] Bishop, C.M.: Pattern recognition and machine learning. Springer Google Schola **2**, 1122–1128 (2006)
- [4] De Lima Marquezino, F., Portugal, R., Lavor, C.: A primer on quantum computing. Springer, Berlin (2019)
- [5] Schuld, M., Sinayskiy, I., Petruccione, F.: An introduction to quantum machine learning. *Contemp. Phys.* **56**(2), 172–185 (2015)
- [6] Pastorello, D.: Concise guide to quantum machine learning. Springer, Berlin (2023)
- [7] Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., Lloyd, S.: Quantum machine learning. *Nature* **549**(7671), 195–202 (2017)
- [8] Girardi, F., De Palma, G.: Trained quantum neural networks are gaussian processes. *Commun. Math. Phys.* **406**(4), (2025)
- [9] Schuld, Maria, Petruccione, Francesco: *Supervised learning with quantum computers*, volume 17. Springer, (2018)
- [10] Schuld, M., Sweke, R., Meyer, J.J.: Effect of data encoding on the expressive power of variational quantum-machine-learning models. *Phys. Rev. A* **103**(3), 032430 (2021)

- [11] Lloyd, S., Schuld, M., Ijaz, A., Izaac, J., Killoran, N.: Quantum embeddings for machine learning. arXiv preprint [arXiv:2001.03622](https://arxiv.org/abs/2001.03622) (2020)
- [12] Liu, Y., Arunachalam, S., Temme, K.: A rigorous and robust quantum speed-up in supervised machine learning. *Nat. Phys.* **17**(9), 1013–1017 (2021)
- [13] Havlíček, V., Córcoles, A.D., Temme, K., Harrow, A.W., Kandala, A., Chow, J.M., Gambetta, J.M.: Supervised learning with quantum-enhanced feature spaces. *Nature* **567**(7747), 209–212 (2019)
- [14] Cinelli, L.P., Marins, M.A., Antonio Barros Da Silva, E., Netto, S.L.: Variational methods for machine learning with applications to deep networks, vol 15. Springer, Berlin (2021)
- [15] McClean, J.R., Boixo, S., Smelyanskiy, V.N., Babbush, R., Neven, H.: Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**(1), 4812 (2018)
- [16] Shalev-Shwartz, S., Shamir, O., Shammah, S.: Failures of gradient-based deep learning. In: International Conference on Machine Learning, pp. 3067–3075. PMLR (2017)
- [17] Larocca, M., Thanasilp, S., Wang, S., Sharma, K., Biamonte, J., Coles, P.J., Cincio, L., McClean, J.R., Holmes, Z., Cerezo, M.: A review of barren plateaus in variational quantum computing. arXiv preprint [arXiv:2405.00781](https://arxiv.org/abs/2405.00781) (2024)
- [18] Martín, E.C., Plekhanov, K., Lubasch, M.: Barren plateaus in quantum tensor network optimization. *Quantum* **7**, 974 (2023)
- [19] Ragone, M., Bakalov, B.N., Sauvage, F., Kemper, A.F., Ortiz Marrero, C., Larocca, M., Cerezo, M.: A lie algebraic theory of barren plateaus for deep parameterized quantum circuits. *Nat. Commun.* **15**(1), 7172 (2024)
- [20] Park, C.-Y., Killoran, N.: Hamiltonian variational ansatz without barren plateaus. *Quantum* **8**, 1239 (2024)
- [21] Liu, X., Liu, G., Zhang, H.K., Huang, J., Wang, X.: Mitigating barren plateaus of variational quantum eigensolvers. *IEEE Trans. Quantum Eng.* (2024)
- [22] Zhang, H.-K., Liu, S., Zhang, S.-X.: Absence of barren plateaus in finite local-depth circuits with long-range entanglement. *Phys. Rev. Lett.* **132**(15), 150603 (2024)
- [23] Cerezo, M., Sone, A., Volkoff, T., Cincio, L., Coles, P.J.: Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.* **12**(1), 1791 (2021)
- [24] Kiani, B.T., Palma, G., Marvian, M., Liu, Z.-W., Lloyd, S.: Learning quantum data with the quantum earth mover’s distance. *Quantum Sci. Technol.* **7**(4), 045002 (2022)
- [25] Villani, C., et al.: Optimal transport: old and new, vol. 338. Springer, Berlin (2009)
- [26] De Palma, G., Trevisan, D.: Quantum optimal transport with quantum channels. In: *Annales Henri Poincaré*, vol. 22, pp. 3199–3234. Springer, Berlin (2021)
- [27] Hanin, B.: Which neural net architectures give rise to exploding and vanishing gradients? *Adv. Neural Inf. Process. Syst.* **31** (2018)
- [28] Lee, J., Bahri, Y., Novak, R., Schoenholz, S.S., Pennington, J., Sohl-Dickstein, J.: Deep neural networks as gaussian processes. arXiv preprint [arXiv:1711.00165](https://arxiv.org/abs/1711.00165) (2017)

- [29] Lee, J., Xiao, L., Schoenholz, S., Bahri, Y., Novak, R., Sohl-Dickstein, J., Pennington, J.: Wide neural networks of any depth evolve as linear models under gradient descent. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- [30] Abedi, E., Beigi, S., Taghavi, L.: Quantum lazy training. *Quantum* **7**, 989 (2023)
- [31] García-Martín, D., Larocca, M., Cerezo, M.: Deep quantum neural networks form gaussian processes. arXiv preprint [arXiv:2305.09957](https://arxiv.org/abs/2305.09957) (2023)
- [32] Nourdin, I., Peccati, G.: Normal approximations with Malliavin calculus: from Stein’s method to universality, vol. 192. Cambridge University Press, Cambridge (2012)
- [33] Nourdin, I., Peccati, G., Réveillac, A.: Multivariate normal approximation using stein’s method and malliavin calculus. *Annales de l’IHP Probabilités et statistiques* **46**, 45–58 (2010)
- [34] Stein, C.: Approximate computation of expectations. In *Lecture notes, IMS* (1986)
- [35] Chen, L.H.Y.: Stein meets malliavin in normal approximation. *Acta Mathematica Vietnamica* **40**, 205–230 (2015)
- [36] Gallouët, T., Mijoule, G., Swan, Y.: Regularity of solutions of the stein equation and rates in the multivariate central limit theorem. arXiv preprint [arXiv:1805.01720](https://arxiv.org/abs/1805.01720) (2018)
- [37] Shirai, N., Kubo, K., Mitarai, K., Fujii, K.: Quantum tangent kernel. *Phys. Rev. Res.* **6**(3), 033179 (2024)
- [38] Incudini, M., Martini, F., Pierro, A.D.: Toward useful quantum kernels. *Adv. Quantum Technol.* 2300298 (2024)
- [39] Nakaji, K., Tezuka, H., Yamamoto, N.: Quantum-classical hybrid neural networks in the neural tangent kernel regime. *Quantum Sci. Technol.* **9**(1), 015022 (2023)
- [40] Liu, J., Tacchino, F., Glick, J.R., Jiang, L., Mezzacapo, A.: Representation learning via quantum neural tangent kernels. *PRX Quantum* **3**(3), 030323 (2022)
- [41] Yu, L.-W., Li, W., Ye, Q., Lu, Z., Han, Z., Deng, D.-L.: Expressibility-induced concentration of quantum neural tangent kernels. *Rep. Progress Phys.* (2023)
- [42] Egginger, S., Sakhnenko, A., Lorenz, J.M.: A hyperparameter study for quantum kernel methods. *Quantum Mach. Intell.* **6**(2), 44 (2024)
- [43] Napp, J.: Quantifying the barren plateau phenomenon for a model of unstructured variational ansätze (2022)
- [44] Ambrosio, L., Gigli, N., Savaré, G.: Gradient flows: in metric spaces and in the space of probability measures. Springer, Berlin (2008)
- [45] Villani, C.: Optimal transport, volume 338 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer, Berlin, (2009)
- [46] Santambrogio, F.: Optimal transport for applied mathematicians. Birkäuser, NY **55**(58–63), 94 (2015)
- [47] Monge, G.: Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.* 666–704 (1781)
- [48] Kantorovich, L.V.: Mathematical methods of organizing and planning production. *Manag. Sci.* **6**(4), 366–422 (1960)

- [49] Evans, S.N., Matsen, F.A.: The phylogenetic kantorovich-rubinstein metric for environmental sequence samples. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **74**(3), 569–592 (2012)
- [50] Rachev, S.T., Stoyanov, S.V., Fabozzi, F.J.: A probability metrics approach to financial risk measures. Wiley (2011)
- [51] Panaretos, V.M., Zemel, Y.: An invitation to statistics in Wasserstein space. Springer, Berlin (2020)
- [52] Frogner, C., Zhang, C., Mobahi, H., Araya, M., Poggio, T.A.: Learning with a wasserstein loss. *Adv. Neural Inf. Process. Syst.* **28** (2015)
- [53] Gabriel, P., Marco, C.: Computational optimal transport: With applications to data science. *Found. Trends® Mach. Learn.* **11**(5-6):355–607 (2019)
- [54] Cheng, C., Zhou, B., Ma, G., Dongrui, W., Yuan, Y.: Wasserstein distance based deep adversarial transfer learning for intelligent fault diagnosis with unlabeled or insufficient labeled data. *Neurocomputing* **409**, 35–45 (2020)
- [55] Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International conference on machine learning, pp. 214–223. PMLR (2017)
- [56] Chakrabarti, S., Yiming, H., Li, T., Feizi, S., Wu, X.: Quantum wasserstein generative adversarial networks. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- [57] Feldman, M., McCann, R.J.: Uniqueness and transport density in monge’s mass transportation problem. *Calc. Var. Partial Differ. Equ.* **15**(1), 81–113 (2002)
- [58] Meckes, M.W.: Gaussian marginals of convex bodies with symmetries. arXiv preprint math/0606073 (2006)
- [59] Wendel, J.G.: Note on the gamma function. *Am. Math. Mon.* **55**(9), 563–564 (1948)
- [60] Chizat, L., Oyallon, E., Bach, F.: On lazy training in differentiable programming (2020)
- [61] Balasubramanian, K., Goldstein, L., Ross, N., Salim, A.: Gaussian random field approximation via stein’s method with applications to wide random neural networks (2024)
- [62] Favaro, S., Boris, H., Nourdin, I., Peccati, G.: Quantitative clts in deep neural networks, Domenico Marinucci (2024)
- [63] Hanin, B., Nica, M.: Finite depth and width corrections to the neural tangent kernel (2019)
- [64] Yaïda, S.: Non-gaussian processes and neural networks at finite widths (2020)
- [65] Roberts, D.A., Yaïda, S., Hanin, B.: The Principles of deep learning theory: an effective theory approach to understanding neural networks. Cambridge University Press, Cambridge (2022)
- [66] Hanin, B.: Random fully connected neural networks as perturbatively solvable hierarchies (2023)
- [67] Féray, V., Méliot, P., Nikeghbali, A., Féray, V., Méliot, P., Nikeghbali, A.: Dependency graphs and mod-gaussian convergence. *Mod- ϕ Convergence: normality Zones and Precise Deviations*, pp. 95–110 (2016)
- [68] Janson, S.: Normal convergence by higher semiinvariants with applications to sums of dependent random variables and random graphs. *Ann. Probab.* 305–312 (1988)

- [69] Dowson, D.C.: BV666017 Landau. The fréchet distance between multivariate normal distributions. *J. Multivar. Anal.* **12**(3), 450–455 (1982)

Quantitative Convergence of Trained Quantum

Anderson Melchor Hernandez, Davide Pastorello and Giacomo De Palma

Department of Mathematics

University of Bologna

Piazza di Porta San Donato 5

40126 Bologna

Italy

e-mail: anderson.melchor@unibo.it;

davide.pastorello3@unibo.it;

giacomo.depalma@unibo.it

Filippo Girardi

Scuola Normale Superiore

Piazza dei Cavalieri 7

56126 Pisa

Italy

e-mail: filippo.girardi@sns.it

and

Korteweg-de Vries Institute for Mathematics

University of Amsterdam

Science Park, 105-107

1098 XG Amsterdam

The Netherlands

and

QuSoft

Science Park 123

1098 XG Amsterdam

The Netherlands

Davide Pastorello

TIFPA-INFN

via Sommarive 14

38123 Povo Trento

Italy

Communicated by David Pérez-García a.

Received: March 19, 2025.

Accepted: September 29, 2025.