A Superpolynomial Lower Bound for the Covering Numbers of PQC-induced Function Classes

Romualdo Santolla
Bachelor's Thesis
Supervisors: Vedran Dunjko (PhD.), Serge Fehr (PhD.)

July 2026

Abstract

A central question in quantum machine learning is to identify learning tasks for which quantum algorithms offer a computational advantage over their classical counterparts. Within the Probably Approximately Correct (PAC) learning framework, the complexity of a learning problem is captured by its concept class, and covering numbers provide a key tool for establishing efficient PAC learning guarantees.

Recent work by Barthe et al. [BRGD25] established a superpolynomial upper bound on the covering number of a family of PQC-induced concept classes, but conjectured that this bound might be an artefact of the proof and that a polynomially sized ε-cover, and hence a simple brute-force learning strategy, might nevertheless exist.

This thesis resolves this question negatively. We show that there exist PQCs of depth $O(\log n)$ and constant frequency width whose induced function classes have L^2-covering number at least $\Omega(n^{\log \log n})$. Our result shows that the previously obtained superpolynomial upper bound cannot, in general, be improved to a polynomial one. Consequently, any finite ε-net approximating these classes must contain a superpolynomial number of functions ruling out the existence of an efficient the brute-force learning strategy.

Contents

1 Introduction

1.1 Background

Quantum computing has the potential to solve certain computational problems more efficiently than classical computers. An important question is whether similar advantages can be achieved in machine learning. Quantum machine learning (QML) investigates the use of quantum computers for learning tasks, with the long-term goal of identifying problems for which quantum algorithms outperform their classical counterparts.

The Probably Approximately Correct (PAC) learning framework was introduced by Valiant [Val84]. Rather than analyzing a particular learning algorithm, the PAC framework studies whether a family of target functions can be learned efficiently from labeled examples. Efficiency is measured in terms of both computational complexity and sample complexity.

Central to PAC learning is the notion of a *concept class*, namely a collection of functions representing all possible target concepts of a learning problem. Various complexity measures have been proposed to quantify the richness of concept classes. Among these, covering numbers are of particular interest, since polynomial covering number bounds yield efficient PAC learning guarantees under suitable assumptions.

Parameterized quantum circuits (PQCs) provide a natural way of constructing concept classes for quantum machine learning. A PQC consists of a quantum circuit containing trainable parameters whose output expectation values define functions on the input space. Varying the trainable parameters therefore gives rise to a family of functions, allowing PQC-induced concept classes to be studied within the PAC framework.

1.2 Established lower bounds of PQC-induced Concept Classes

As discussed in the previous section, every parameterized quantum circuit (PQC) induces a concept class by varying its trainable parameters. The functions in such a class are obtained as expectation values of a fixed measurement operator applied to the output state of the circuit. Consequently, the learnability of PQCs can be studied within the PAC learning framework by analyzing the complexity of the corresponding concept classes.

Barthe et al. [BRGD25] studied concept classes induced by parameterized quantum circuits acting on the binary hypercube. For technical convenience, we instead work with the equivalent input space

$$\mathcal{X} = \{\pm 1\}^n \times \{0, 1\},$$

which is naturally in bijection with a binary hypercube of the same cardinality. This reparametrization does not affect the underlying learning problem or the covering-number analysis, but leads to a considerably simpler construction.

A central quantity in the analysis of Barthe et al. [BRGD25] is the $L^2(\mathcal{X})$-covering number of concept classes induced by expectation values of measurements. A formal description of this is given in Section 2.3. They established a superpolynomial upper bound on this covering number and showed that the corresponding learning problem nonetheless admits an efficient learning algorithm. This upper bound was improved to a $\Omega(n^{\log \log n})$ upperbound by Caro et al. [CHC+22], however Barthe et al. conjectured that this upper bound might be tightened further and that the superpolynomial bound is merely an artefact of the proof. This would admit a simpler, efficient learning algorithm through brute-force, that is, by iterating over polynomially many representative functions until a match has been found.

The main contribution of this thesis is to resolve this conjecture negatively. In fact, we show that the previously attained upper bound by C. Caro et al. is sharp. We identify a family of logarithmic-depth parameterized quantum circuits, namely the signed probe family, whose induced concept class has superpolynomial $L^2(\mathcal{X})$-covering number. Consequently, the previously established upper bound cannot, in general, be improved to a polynomial one.

Theorem 1 (Lower bound of the Covering Number of PQCs)**.** *There exist a fixed observable O and a family of logarithmic-depth parameterized quantum circuits with constant frequency width such that the induced concept class satisfies*

$$\mathcal{N}\Big(\mathcal{C}^{\mathrm{PQC}}_{n,\log,O}, L^2(\mathcal{X}), \varepsilon\Big) = \Omega\big(n^{c \log \log n}\big)$$

for every sufficiently small fixed $\varepsilon > 0$ and some constant $c > 0$.

We prove this by constructing *the signed probe family* and showing that its covering number has the quasi-polynomial lower bound as in Theorem 1.

The remainder of this thesis is organized as follows. Chapter 2 introduces the necessary notation and preliminaries on covering numbers and parameterized quantum circuits. Chapter 3 constructs the signed probe family. Chapter 4 establishes a superpolynomial lower bound on the covering number of a subset of its function family. Finally, Chapter 5 combines the results to prove the main theorem.

2 Preliminaries

This chapter introduces the notation and definitions used throughout the remainder of the thesis. Unless stated otherwise, we follow the formalism of Barthe et al. [BRGD25].

2.1 Function Spaces and Covering Numbers

Throughout this thesis we work with the input space

$$\mathcal{X} = \{\pm 1\}^n \times \{0, 1\},$$

equipped with the uniform probability measure. As discussed in Section 1, this is merely a convenient reparametrization of the binary hypercube considered in [BRGD25].

Definition 2 (L^2-norm)**.** *Let $f : \mathcal{X} \to \mathbb{R}$. The L^2-norm of f is defined by*

$$\|f\|_{L^2(\mathcal{X})} := \left(\frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} |f(x)|^2 \right)^{1/2} = \left(\mathbf{E}_{x \sim \mathrm{Unif}(\mathcal{X})} \left[|f(x)|^2 \right] \right)^{1/2}.$$

For $f, g : \mathcal{X} \to \mathbb{R}$, the corresponding distance is

$$\|f - g\|_{L^2(\mathcal{X})}.$$

Since the underlying domain is fixed throughout the thesis, we write $\|\cdot\|_{L^2}$ or simply L^2 whenever no confusion can arise.

Definition 3 (Covering Number)**.** *Let $(X, \|\cdot\|)$ be a normed vectorspace and let $F \subseteq X$. For $\varepsilon > 0$, the ε-covering number of F, denoted by $\mathcal{N}(F, \|\cdot\|, \varepsilon)$, is the smallest integer N such that there exist points $x_1, \ldots, x_N \in X$ with*

$$F \subseteq \bigcup_{i=1}^{N} B_d(x_i, \varepsilon),$$

where

$$B_d(x, \varepsilon) = \{y \in X : \|x - y\| < \varepsilon\}.$$

Equivalently, $\mathcal{N}(F, \|\cdot\|, \varepsilon)$ is the minimum cardinality of an ε-net for F with respect to $\|\cdot\|$.

2.2 Parameterized Quantum Circuits

We use Dirac (bra-ket) notation. For a vector $|\psi\rangle \in \mathbb{C}^2$, we write

$$\langle\psi| := |\psi\rangle^\dagger$$

for its conjugate transpose.

We write

$$A \doteq B$$

to denote equality up to a global phase; that is, there exists $\lambda \in \mathbb{C}$ with $|\lambda| = 1$ such that $A = \lambda B$.

A parameterized quantum circuit (PQC) is a family of unitary operators

$$U_\alpha(x) : \mathbb{C}^{2n} \to \mathbb{C}^{2n}$$

parameterized by a classical input $x \in \mathcal{X}$ and a trainable parameter vector $\alpha \in [0,1]^d$. We restrict our attention to parameterized quantum circuits acting on the two-dimensional Hilbert space $\mathbb{C}^2$, commonly referred to as a single qubit.

Definition 4 (Observable)**.** *An observable is a Hermitian operator*

$$O = O^\dagger$$

acting on the Hilbert space.

We will make use of the Pauli-observables in particular:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix},$$

and their corresponding rotation gates

$$R_X(\phi) = e^{-i\phi X/2}, \qquad R_Z(\phi) = e^{-i\phi Z/2}.$$

2.3 PQC-induced concept class

Definition 5 (PQC-induced function)**.** *Given a parameterized quantum circuit $U_\alpha(x)$, an observable O, and a fixed initial state $|\psi_0\rangle$, define*

$$f_\alpha(x) = \langle\psi_0|U_\alpha(x)^\dagger O U_\alpha(x)|\psi_0\rangle.$$

This function represents the expected outcome of measuring the observable O after applying the circuit to the initial state $|\psi_0\rangle$.

Definition 6 (PQC-induced concept class)**.** *The concept class induced by the parameterized quantum circuit $U_\alpha(x)$ and fixed observable O is*

$$\mathcal{C}^{\mathrm{PQC}}_{n,f,O} = \{f_\alpha : \mathcal{X} \to \mathbb{R} \mid \alpha \in [0,1]^{f(n)}\}.$$

When it is clear what observable we are talking about, we shall omit O from the subscript.

3 A One-Qubit Construction of the Signed-Probe Family

3.1 Elementary Identities

The construction in Section 3.3 relies on the following elementary identities.

Lemma 7. *For every $\phi \in \mathbb{R}$,*

$$XR_Z(\phi)X = R_Z(-\phi).$$

Furthermore,

$$R_X(\pi) \doteq X,$$

where $\doteq$ denotes equality up to a global phase.

Proof. The first identity follows directly from the relation $XZX = -Z$:

$$XR_Z(\phi)X = Xe^{-i\phi Z/2}X = e^{-i\phi(XZX)/2} = e^{i\phi Z/2} = R_Z(-\phi).$$

The second equality following from

$$Xe^{Z}X = X(I_2 + Z + Z^2/2! + ...)X = I_2 + XZX + XZ^2X/2! + ... = e^{XZX}$$

For the second identity,

$$R_X(\pi) = e^{-i\pi X/2} = -iX,$$

which differs from X only by the global phase $-i$. □

3.2 The Signed-Probe Family

We introduce the auxiliary function families that is core in proving Theorem 1.

Definition 8 (Signed-Probe Family)**.** *Let $p = \log n$ and $A > 0$. For each $\alpha \in [0, A]^p$, define*

$$f_\alpha : \{\pm 1\}^p \times \{0, 1\} \to \mathbb{R}$$

by

$$f_\alpha(s, b) = \cos\left(\frac{\pi b}{2} + \pi s \cdot \alpha\right).$$

Equivalently,

$$f_\alpha(s, 0) = \cos(\pi s \cdot \alpha), \qquad f_\alpha(s, 1) = -\sin(\pi s \cdot \alpha).$$

We define

$$\mathcal{F}_{n,\log,A} = \{f_\alpha : \alpha \in [0, A]^p\}.$$

and

$$\mathcal{F}_{p,\log} := \mathcal{F}_{p,\log,1}$$

Unless stated otherwise, we let $A = \frac{1}{4}$.

3.3 The Construction

We now show that the Signed-Probe Family is indeed realized by a one-qubit parameterized quantum circuit.
Let $p = \log n$ and let

$$\alpha = (\alpha_1, \ldots, \alpha_p) \in [0, A]^p.$$

The classical input consists of

$$q = (q_1, \ldots, q_p) \in \{0,1\}^p, \qquad b \in \{0,1\}.$$

Define the one-qubit circuit

$$U_{\alpha,q,b} = R_Z\left(\frac{\pi b}{2}\right) \prod_{j=1}^{p} \left(R_X(\pi q_j) R_Z(\pi \alpha_j)\right),$$

where

$$\prod_{j=1}^{p} A_j := A_p A_{p-1} \cdots A_1.$$

The real-valued output of the circuit is defined by

$$F_\alpha(q,b) = \langle + | U_{\alpha,q,b}^\dagger X U_{\alpha,q,b} | + \rangle.$$

We first put the circuit into a normal form. Since

$$R_X(\pi q_j) \doteq X^{q_j},$$

and since

$$X R_Z(\phi) X = R_Z(-\phi),$$

we may commute all X-factors to the right. Each time an X-factor passes a Z-rotation, it changes the sign of the rotation angle.
Define

$$Q(q) = q_1 + \cdots + q_p \pmod 2$$

and

$$s_j(q) = (-1)^{q_j + q_{j+1} + \cdots + q_p}, \qquad j = 1, \ldots, p.$$

Then, up to a global phase,

$$U_{\alpha,q,b} \doteq R_Z\left(\frac{\pi b}{2} + \pi \sum_{j=1}^{p} s_j(q) \alpha_j\right) X^{Q(q)}.$$

Since $X|+\rangle = |+\rangle$, the factor $X^{Q(q)}$ does not affect the input state $|+\rangle$. Hence the state before measurement is, up to a global phase,

$$R_Z\left(\frac{\pi b}{2} + \pi \sum_{j=1}^{p} s_j(q) \alpha_j\right) |+\rangle.$$

For every real ϕ,

$$\langle +|R_Z(\phi)^\dagger X R_Z(\phi)|+\rangle = \cos\phi.$$

Therefore

$$F_\alpha(q,b) = \cos\left(\frac{\pi b}{2} + \pi\sum_{j=1}^{p} s_j(q)\alpha_j\right).$$

Finally, the map

$$q \mapsto s(q) = (s_1(q), \ldots, s_p(q))$$

is a bijection from $\{0,1\}^p$ onto $\{\pm 1\}^p$. Thus, after reparametrizing the input by the sign vector $s = s(q)$, the circuit realizes the family

$$f_\alpha(s,b) = \cos\left(\frac{\pi b}{2} + \pi s\cdot\alpha\right), \qquad s \in \{\pm 1\}^p,\ b \in \{0,1\}.$$

Equivalently,

$$f_\alpha(s,0) = \cos(\pi s\cdot\alpha), \qquad f_\alpha(s,1) = -\sin(\pi s\cdot\alpha).$$

Hence the circuit output class is precisely as defined in Definition 8. In particular $\mathcal{F}_{p,\log}$ is precisely a PQC-induced concept class as defined in Definition 6.

4 A Cover Lower Bound for the Signed-Probe Family

We begin by deriving an explicit expression for the L^2-distance between two functions in $\mathcal{F}_{n,\log,A}$.

Lemma 9 (Distance between Probe-Functions)**.** *For every $\alpha, \beta \in [0,A]^p$ define $\delta := \alpha - \beta$. We have*

$$\|f_\alpha - f_\beta\|_{L^2}^2 = 1 - \prod_{k=1}^{p} \cos(\pi\delta_k).$$

Proof. Let $(s,b) \sim uniform$ We have that

$$\mathbf{E}_b[(f_\alpha(s,b) - f_\beta(s,b))^2] = \frac{1}{2}(f_\alpha(s,0) - f_\beta(s,0))^2 + \frac{1}{2}(f_\alpha(s,1) - f_\beta(s,1))^2 = \frac{1}{2}|e^{i\pi s\cdot\alpha} - e^{i\pi s\cdot\beta}|^2$$

Using the identity $|z_1 - z_2|^2 = |z_1|^2 + |z_2|^2 - 2\,\mathrm{Re}(z_1\bar{z_2})$ for complex numbers, we get

$$\begin{aligned}
\|f_\alpha - f_\beta\|_{L^2}^2 &= \mathbf{E}_{(s,b)}[|f_\alpha - f_\beta|^2] \\
&= \frac{1}{2}\mathbf{E}_s[|e^{i\pi s\cdot\alpha} - e^{i\pi s\cdot\beta}|^2] \\
&= 1 - \mathbf{E}_s[\mathrm{Re}(e^{i\pi s\cdot\delta})] \\
&= 1 - \mathbf{E}_s[\cos(\pi s\cdot\delta)]
\end{aligned}$$

We can simplify further, as we have

$$\begin{aligned}
\mathbf{E}_s[\cos(\pi s \cdot \delta)] &= \mathrm{Re}(\mathbf{E}_s[e^{i\pi s \cdot \delta}]) \\
&= \mathrm{Re}(\prod_{k=1}^{p} \mathbf{E}_{s_k}[e^{i\pi s_k \delta_k}]) \\
&= \mathrm{Re}(\prod_{k=1}^{p} [\frac{1}{2}(\cos(\pi\delta_k) + i\sin(\pi\delta_k) \\
+ &\qquad \frac{1}{2}(\cos(\pi\delta_k) - i\sin(\pi\delta_k))] \\
&= \prod_{k=1}^{p} \cos(\pi\delta_k)
\end{aligned}$$

consequently

$$\|f_\alpha - f_\beta\|_{L^2}^2 = 1 - \prod_{k=1}^{p} \cos(\pi\delta_k)$$

□

Lemma 10. *Suppose $A \leq \frac{1}{4}$. Then there exists a constant $m > 0$ such that*

$$\cos(x) \leq e^{-mx^2}$$

for every $x \in [-\pi A, \pi A]$. Consequently,

$$\|f_\alpha - f_\beta\|_{L^2}^2 \geq 1 - e^{-m\pi^2\|\alpha-\beta\|_2^2}$$

for all $\alpha, \beta \in [0, A]^p$.

Proof. Define

$$g(x) = \begin{cases} \dfrac{-\log(\cos x)}{x^2}, & 0 < |x| \leq \pi A, \\ \dfrac{1}{2}, & x = 0. \end{cases}$$

Since $A \leq 1/4$, we have $\cos(x) > 0$ on $[-\pi A, \pi A]$. Furthermore, g is continuous on this compact interval and strictly positive, so it attains a positive minimum $m > 0$. Hence

$$-\log(\cos x) \geq mx^2,$$

or equivalently,

$$\cos(x) \leq e^{-mx^2}.$$

Applying this coordinatewise gives

$$\prod_{k=1}^{p} \cos(\pi\delta_k) \leq \exp\left(-m\pi^2 \sum_{k=1}^{p} \delta_k^2\right),$$

which proves the claim. □

Theorem 11 (Covering Number Lower Bound for the Signed-Probe Family)**.** *For every sufficiently small fixed $\varepsilon > 0$,*

$$\mathcal{N}(\mathcal{F}_{n,\log,A}, L^2, \varepsilon) = \Omega\big(n^{C \log\log n}\big)\,.$$

for some constant $C > 0$.

Proof. Let

$$\varepsilon' = \sqrt{\frac{-\log(1-\varepsilon^2)}{m\pi^2}}.$$

Suppose $\{f_{\alpha^1}, \ldots, f_{\alpha^N}\} \subseteq \mathcal{F}_{n,\log,A}$ is an ε-cover of $\mathcal{F}_{n,\log,A}$. Then for every $\alpha \in [0,A]^p$, there exists $i \in \{1, \ldots, N\}$ such that

$$\|f_\alpha - f_{\alpha^i}\|_{L^2} \le \varepsilon.$$

By Lemma 10,

$$1 - e^{-m\pi^2\|\alpha-\alpha^i\|_2^2} \le \varepsilon^2,$$

which implies

$$\|\alpha - \alpha^i\|_2 \le \varepsilon'.$$

Hence $\{\alpha^1, \ldots, \alpha^N\}$ forms an ε'-cover of the cube $[0,A]^p$. Therefore,

$$\mathcal{N}(\mathcal{F}_{n,\log,A}, L^2, \varepsilon) \ge \mathcal{N}([0,A]^p, \|\cdot\|_2, \varepsilon').$$

By the volumetric estimate (Proposition 4.2.12 in [Ver18]) and the asymptotic formula for the volume of the Euclidean unit ball (obtained from Stirling's formula),

$$\mathcal{N}([0,A]^p, \|\cdot\|_2, \varepsilon') = \Omega\left(\left(\frac{A\sqrt{p}}{\varepsilon'}\right)^p\right).$$

Since A and ε are fixed, ε' is a constant

$$\mathcal{N}(\mathcal{F}_{n,\log,A}, L^2, \varepsilon) = \Omega((c\sqrt{p})^p)\,,$$

for some constant $c > 0$. Finally, since $p = \log n$, we have

$$\mathcal{N}(\mathcal{F}_{n,\log,A}, L^2, \varepsilon) = \Omega\Big((\sqrt{\log n})^{\log n}\Big) = \Omega\big(n^{C\log\log n}\big)\,,$$

for some constant $C > 0$. □

5 Proof of the Main Theorem

Proof of Theorem 1. Since

$$\mathcal{F}_{n,\log,A} \subseteq \mathcal{F}_{n,\log}$$

We have $\mathcal{N}(\mathcal{F}_{n,\log,A}, \|\cdot\|_{L^2}, \varepsilon) \le \mathcal{N}(\mathcal{F}_{n,\log}, \|\cdot\|_{L^2}, \varepsilon)$ so Theorem 11 directly implies $\mathcal{N}(F_d, \|\cdot\|_{L^2}, \varepsilon) = \Omega(n^{c\log\log n})$

□

6 Conclusion

We studied the covering numbers of concept classes induced by parameterized quantum circuits. We constructed a family of logarithmic-depth PQCs with constant frequency width whose induced concept class has superpolynomial L^2-covering number. Consequently, the previously established superpolynomial upper bound on the covering number cannot, in general, be improved to a polynomial one, thereby resolving the conjecture of Barthe et al. [BRGD25] negatively.

Although this result rules out the existence of a polynomially sized ε-cover, and hence a brute-force learning strategy based on enumerating such a cover, the learning problem remains efficiently solvable using the sparse recovery algorithm of Barthe et al. This highlights that large covering numbers do not necessarily preclude efficient learning and suggests that covering-number techniques alone are insufficient to characterize the learnability of PQC-induced concept classes.

An interesting direction for future work is to further investigate the relationship between the structural properties of PQC-induced concept classes, their metric complexity, and their learnability. While this thesis shows that polynomial sized frequency space alone does not imply polynomial covering numbers, the precise relationship between the structure of the frequency spectrum and covering numbers remains open. At the same time, it would be valuable to understand which structural properties of PQC-induced concept classes permit efficient learning despite their superpolynomial covering numbers.

References

[BRGD25] Alice Barthe, Mahtab Yaghubi Rad, Michele Grossi, and Vedran Dunjko. Quantum advantage in learning quantum dynamics via fourier coefficient extraction. *arXiv preprint*, 2025. Available at `https://arxiv.org/abs/2506.17089`.

[CHC+22] Matthias C. Caro, Hsin-Yuan Huang, M. Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J. Coles. Generalization in quantum machine learning from few training data. *Nature Communications*, 13(1):4919, 2022.

[Val84] Leslie G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 2018.