Learning with Quantum Examples: Multiclass, Online, and Smoothed Settings

by

Preetham Mohan

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Applied and Interdisciplinary Mathematics)
in the University of Michigan
2026

Doctoral Committee:

Professor Ambuj Tewari, Co-Chair
Professor Shravan Veerapaneni, Co-Chair
Dr. Srinivasan Arunachalam, IBM
Professor Martin Strauss

Preetham Mohan

preetham@umich.edu

ORCID iD: 0000-0003-3401-2985

Dedication

To my parents. None of this would have been possible without your fierce support and the many sacrifices you made, allowing me to be limitless in pursuing opportunities at every step.

Acknowledgments

At a time when I had little clarity about the direction I wanted my research to take and was quite apprehensive about approaching a potential advisor, I had somehow mustered the courage to reach out to Prof. Ambuj Tewari. I remain immensely grateful that he not only agreed to take me on but also helped shape an area of focus that played to my strengths. Over the several years that have gone into shaping this thesis, through rewarding highs and more-than-occasional disappointing lulls, I have remained indebted to him for his patience, countless enlightening discussions, and for demonstrating how much can be accomplished through sheer resolve, persistence, confidence, and energy.

I would also like to thank my co-advisor, Prof. Shravan Veerapaneni, for his steadfast support of this thesis and for lending a near-term practitioner's perspective to the work, and Prof. Martin Strauss for his guidance as a member of my thesis committee, particularly during its final stages. I am also sincerely grateful to Dr. Srinivasan Arunachalam for graciously serving on my thesis committee and for several insightful email exchanges that informed and motivated the research presented in this thesis. Much of the work presented here draws direct inspiration from his pioneering contributions to the field.

I would further like to thank Sagnik Chatterjee for several insightful conversations during the early years of this thesis, which broadened my exposure to many research frontiers within quantum. I am also grateful to my peers in the CASI group. Through your reading-group presentations and our discussions in Slack and during advisor meetings, I was exposed to a wealth of knowledge and ideas, from which I benefited greatly, even if I could perhaps imbibe only a small portion. It was a privilege to experience being in such a vibrant research environment. Thanks, in particular, to Vinod and Unique for insightful messages about the classical learning theory counterparts to the work presented in this thesis.

A primary reason I was able to pursue a trajectory in academic research at all is my undergraduate advisor, Prof. Thomas P.-Y. Yu. Thank you for your immense mentorship, guidance, and care.

To my fellow graduate students, thank you for reminding me, especially during difficult moments, that kindness, generosity, and authenticity remain abundant in the world. Simply being surrounded by such remarkable people often provided the perspective and encouragement I needed to keep moving forward.

Regardless of how much comfort I seemed to find in solitude, I don't think I would have had the energy to complete this thesis without the many friendships, interactions, and experiences that

provided me with enough social stimulus to remain well-regulated. So, if we have ever shared a positive interaction, please know that it likely helped me move through my day a little more smoothly.

I cherish the countless hours spent playing (and anticipating playing) soccer in Ann Arbor. It gave me community, routine, and a reason to step away from my desk when I most needed to. I am proud to have been part of the long lineage of our departmental (now interdepartmental) *Eigenvictors* team and grateful to everyone who shared the field with me. Some of my fondest memories of graduate school were made there.

To Tejaswi and Sridhar, an extraordinary number of the experiences that defined these years were shared with you both, and I am deeply grateful for your friendship, support, and care. Thank you for being the people I could always rely on. To Shravan and Shabarish, thank you for making things a bit more lively in my final year.

I don't know if I would have ever felt as though I truly belonged somewhere, anywhere, had it not been for my closest friends from my undergraduate years: Anu, Merlin, Omesh, Shayna, Soham, Sravya, Vinay, Yuvraj, and Zeel. It is wildly impressive that we have continued to check on each other, and celebrate each other's successes with much gusto.

To Kayla, Teresa, and the departmental staff, thank you for your patience, kindness, and for making the countless logistical aspects of graduate school feel effortless. I am also grateful for the numerous departmental events that helped foster a sense of community and connection throughout these years.

Finally, I owe the deepest gratitude to my parents. Any accomplishment reflected in these pages rests upon the foundation of your love, sacrifices, encouragement, and unwavering belief in me. I am also grateful to my relatives, whose care, generosity, and constant support have added to that source of strength throughout my life.

This thesis is as much a testament to the support I have received as it is to any effort of my own.

TABLE OF CONTENTS

LIST OF FIGURES

FIGURE

LIST OF TABLES

TABLE

ABSTRACT

As quantum computing progresses toward fault-tolerant architectures, the question of which computational tasks admit provable quantum advantages and which do not has become increasingly central. Learning theory, and in particular learning from quantum examples, provides one of the few settings in which unconditional quantum-classical separations can be established. In distribution-free (i.e., worst-case) PAC learning, existing results show that quantum examples provide no asymptotic advantage in sample complexity. In contrast, under the uniform distribution, unbounded quantum-classical separations are known for learning Fourier-sparse Boolean functions. Together, these results reveal a striking dichotomy. However, this understanding has largely been developed in the context of learning Boolean functions in the batch setting, leaving open how these phenomena extend more broadly. This thesis develops the theory of learning with quantum examples beyond the batch Boolean setting along three directions: multiclass learning, online learning, and smoothed learning.

In the multiclass PAC setting, we establish upper and lower bounds on quantum sample complexity in both the realizable and agnostic regimes, finding that quantum examples continue to yield no distribution-independent separation from classical examples, with learning rates governed by the Natarajan dimension up to logarithmic factors in the label-space size. We next study online learning, where no standard framework for learning with quantum examples existed prior to this work. We provide such a model by lifting the classical online framework to one in which the adversary provides distributions over labeled examples, and then by encoding these distributions as quantum examples. We establish expected regret guarantees for binary and multiclass classification in both the realizable and agnostic settings. The central finding is that unrestricted adversarial power permits highly concentrated distributions that dequantize the learning problem.

Motivated by this dequantization phenomenon, we develop a smoothed learning framework that constrains distributions to be smooth, interpolating between the concentrated-distribution regime, in which no quantum advantage exists, and the uniform-distribution regime, in which unbounded separations are known. For the class of Fourier-sparse Boolean functions, we establish the robustness of quantum Fourier sampling throughout a nontrivial near-uniform regime in the batch setting and observe that these guarantees extend naturally to the online setting. We further analyze several approaches to recovering the target function once its Fourier support has been

identified and, based on these observations, formulate conjectures supporting the persistence of the unbounded quantum advantage throughout this regime.

Together, these results paint a coherent picture of learning with quantum examples beyond the batch Boolean setting, showing that quantum-classical separations depend on the interplay between hypothesis class structure, distributional assumptions, and the degree of adversarial control permitted in the learning process.

CHAPTER 1

Introduction

1.1 Why Learning Theory Matters for Quantum Computing

Understanding the relative power of quantum versus classical computation is one of the major goals of complexity theory. Following the seminal work of Shor [Sho99], it is widely believed that quantum computation is exponentially more powerful than classical computation for certain tasks; however, since we are unable to prove super-polynomial classical lower bounds for strong models of computation, there are relatively few settings in which such a separation can be *unconditionally* established. Query complexity and sample complexity are important examples of such settings where exponential quantum speedups can be proved unconditionally [AA15, AC16, ACL+21, GS25], and it is precisely this feature that makes the learning-theoretic lens so valuable for quantum computing research.

Quantum computing also enjoys a peculiar privilege of timeline. Its foundational mathematics, rooted in linear algebra, probability theory, and complexity theory, is so rigorously developed that we can prove theorems about computational tasks that we may not be able to execute experimentally for decades. Many of the predicted advantages would require large-scale fault-tolerant devices that remain beyond current technological reach. This stands in stark contrast to the trajectory of classical machine learning, where a large part of the motivation for theoretical study arose *after* its empirical success. That is, algorithms demonstrably "worked" in practice long before we could adequately explain why. We are therefore fortunate to find ourselves at a moment when the explosion of machine learning theory coincides with an urgent search for quantum advantage. The unconditional separations afforded by sample complexity sit at precisely that intersection, making it an especially fertile ground for the two fields to inform and sharpen each other. Indeed, understanding where quantum examples provably help is only half the picture. Identifying settings where they do *not* is equally illuminating, since both sides sharpen our understanding of the structural origins of quantum speedups.

1.2 The Emerging Path to Fault-Tolerant Quantum Computing

The theoretical investigation of quantum learning is not merely an abstract exercise; it is propelled by a hardware ecosystem advancing at a pace that would have seemed overly ambitious even five years ago. While no single architecture has yet emerged as a clear winner, there is a credible and converging path toward scalable, fault-tolerant quantum computing across multiple modalities.

On the superconducting front, Google's *Willow* chip (105 qubits, 2024) crossed the "below threshold" point in quantum error correction, demonstrating that logical error rates decrease exponentially as physical qubit counts increase [Goo25]. IBM's roadmap targets *Starling*, its planned fault-tolerant machine, to achieve 200 logical qubits capable of executing 100 million quantum gates by 2029 [IBM26]. Neutral-atom platforms, pursued by QuEra, Atom Computing, and Pasqal, offer long coherence times and high connectivity at the cost of slower gate speeds; strikingly, Cain et al. [CXK+26] recently showed that Shor's algorithm can be executed at cryptographically relevant scales with as few as 10,000 reconfigurable atomic qubits, a dramatic reduction from prior estimates in the millions. Trapped-ion systems (IonQ, Quantinuum) offer very high gate fidelities [HSL+25], while Microsoft's topological qubit program promises intrinsically lower error-correction overhead if physical noise can be suppressed at the hardware level [AAA+25].

With this converging evidence of a credible path to fault tolerance, the central theoretical question sharpens: *for which tasks does a quantum computer offer a provable advantage, and for which tasks is such an advantage provably out of reach?* Mapping this landscape, that is, determining where quantum resources are genuinely necessary and where they are not, is not merely an academic exercise; it is what will direct the use of fault-tolerant hardware as it matures, and consequently serves as the central organizing theme of this thesis.

The Quantum–Classical Feedback Loop. Even if scalable fault-tolerant quantum hardware were never realized, this theoretical exploration would still be worth pursuing. Once a purported quantum speedup is proposed, researchers are compelled to verify that no efficient classical algorithm exists for the same task, and this pressure has repeatedly produced classical improvements. Tang's *dequantization* of the quantum recommendation systems algorithm [Tan19] is the canonical example, producing a novel and efficient classical algorithm as a direct byproduct of the quantum investigation. This cycle of quantum proposal and classical response has been productive enough to spawn an entire subfield: quantum-inspired machine learning, which develops classical algorithms that leverage the principles, if not the hardware, of quantum mechanics [Tan22, HHM+23]. This makes the theoretical study of quantum learning valuable regardless of when, or indeed *whether*, the hardware fully arrives.

1.3 Quantum Learning Theory: Where We Fit

The survey of Anshu and Arunachalam [AA24] on the complexity of learning quantum states organizes the field into a natural 2×2 framework along two axes: **what** is being learned – *all* quantum states versus *some subclass*, and **how well** it needs to be learned – under *strong* requirements (e.g., recovery up to small trace distance) versus *weaker* ones (e.g., PAC-style generalization, statistical query learning). The quadrant most relevant to this thesis is *learning some subclass of states under weaker requirements*, and specifically, the subclass corresponding to classical functions encoded as quantum examples. A quantum example for a target function $h^\star$ under distribution D is a coherent superposition of labeled examples of the form $\sum_x \sqrt{D(x)}\,|x, h^\star(x)\rangle$, which replaces a classical labeled sample $(x, h^\star(x))$ drawn from D, allowing the learner to exploit quantum interference and measurement to extract information about the target concept.

The study of learning classical functions from quantum examples was initiated by Bshouty and Jackson [BJ95], who formulated this notion of a quantum example and gave an efficient quantum algorithm for learning DNFs under the uniform distribution. The subsequent literature has developed in several directions: establishing tight characterizations of quantum sample complexity for PAC and agnostic learning [AS05, Zha10, AdW18]; developing explicit quantum algorithms in both the realizable and agnostic settings [JTY02, AS07, KRS19, Car20, CSS15, GKZ19, BA22]; and broadening the model to accommodate quantum membership queries [SG04, Mon12], quantum statistical queries [ABC$^+$20], and quantum distribution learning [SSHE21].

This body of work, however, has focused almost exclusively on learning *Boolean* functions in the *batch* setting, leaving several natural extensions open prior to this thesis:

(i) **Multiclass Learning.** In contrast to Boolean concept classes, whose combinatorial complexity is governed by the VC dimension, multiclass learning exhibits a substantially richer structural landscape. Learnability in this setting is characterized by notions such as the Natarajan dimension [Nat89, CEH$^+$25] and the Daniely–Shalev-Schwartz (DS) dimension [DSS14, BCD$^+$22], both natural generalizations of the VC dimension to multiclass hypothesis classes. Despite the central role of multiclass learning in modern learning theory, the quantum sample complexity of multiclass learning had not been systematically studied prior to our work.

(ii) **Online learning.** In the online model, examples arrive one at a time and the learner must predict before observing the label; performance is measured by the *mistake bound* over an adversarial sequence. While the online learnability of quantum *states* has been studied [ACH$^+$19, AQS22], an online learning framework for learning classical function classes with quantum examples had not previously been formulated. This is particularly well-motivated given the rich classical online learning literature [Lit88, BDPSS09, DSBDSS15, SSBD14],

where combinatorial characterizations of learnability and mistake rates are already well understood.

(iii) **Smoothed learning.** Haghtalab et al. [HRS20] initiated the smoothed analysis of classical online learning, demonstrating that slight relaxations of worst-case adversarial assumptions can recover learnability for classes that are otherwise intractable in the fully adversarial setting. Concretely, the adversary is required to sample from distributions that are not overly concentrated, thereby interpolating between worst-case online and statistical i.i.d.-like models. Extending this perspective to quantum learning settings is, as a consequence of (ii), entirely novel to this thesis.

1.4 Contributions and Thesis Outline

Building on the gaps identified above, this thesis develops the theory of learning with quantum examples along the axes of multiclass learning, online learning, and smoothed learning.

Chapter 2: Batch Learning with Multiclass Quantum Examples. We introduce the framework for quantum batch (PAC) learning in both the realizable and agnostic settings, together with the necessary preliminaries. We then motivate the study of multiclass classification in this framework and establish upper and lower bounds on the quantum sample complexity of multiclass learning in both the realizable and agnostic PAC settings.

Chapter 3: Online Learning with Adversarially Chosen Distributions and Quantum Examples. We revisit the canonical classical online learning model and progressively generalize it to a corresponding quantum online learning framework. As an intermediate step, we introduce a classical *adversary-provides-a-distribution* model, formulated as an adaptive-adversary variant of the framework of [DT22]. This model is of independent interest as a classical online learning framework in its own right, and serves additionally as a conceptual bridge to the quantum setting. We then introduce a quantum online learning model for learning classical function classes from quantum examples and establish expected regret bounds for binary and multiclass classification in both the realizable and agnostic settings.

Chapter 4: Smoothed Learning with Quantum Examples. Motivated by the observations in Chapters 2 and 3 that highly concentrated example distributions obstruct quantum learning advantages, we introduce a framework for smoothed learning from quantum examples under σ-smooth distributions. We study both the batch and online settings, focusing on the class of k-Fourier-sparse Boolean functions. Our analysis separates the learning process into two stages: first,

quantumly identifying the Fourier support of the target function, and second, classically recovering the target function from this support. For the first stage, we establish that quantum Fourier sampling remains robust throughout a nontrivial near-uniform regime in the batch setting, and we observe that these guarantees extend naturally to the online setting even against adaptive, σ-smooth adversaries. For the second stage, we analyze three classical recovery approaches that succeed under the uniform distribution and investigate their prospects in the smoothed setting. These observations lead us to formulate conjectures that, if resolved affirmatively, would imply that the unbounded quantum advantage previously established under the uniform distribution [ACL+21] persists throughout the identified near-uniform regime for both batch and online learning.

Chapter 5: Conclusion. We conclude by synthesizing the main conceptual themes and technical contributions of the thesis, discussing the limitations and robustness of quantum learning advantages across the settings studied, and highlighting several open problems and directions for future work.

CHAPTER 2

Batch Learning with Multiclass Quantum Examples

Bshouty and Jackson [BJ95] introduced a quantum extension of Valiant's [Val84] Probably Approximately Correct (PAC) learning framework by formalizing what it means to learn from quantum examples. In Section 2.1, we develop the technical framework necessary to describe this model and discuss existing sample complexity results for quantum PAC learning of Boolean hypothesis classes. The implications of these results motivate our study of quantum PAC learning for multiclass hypothesis classes (Section 2.2).

This chapter is based largely on material from [MT25]. Section 2.1 adapts the preliminary material developed in Section 2 of that paper, while the principal technical development in Section 2.2 is adapted from Section 3. Minor revisions have been made to the exposition, notation, and organization for inclusion in this thesis.

2.1 Preliminaries

In this section, we introduce the key concepts and notation used throughout the chapter, and more broadly throughout the thesis. We begin with a brief review of the fundamentals of quantum computing, and then outline the batch learning frameworks in both the classical and quantum settings, setting the stage for our main results in Section 2.2.

2.1.1 Notation

In the bra-ket (Dirac) notation, a ket, $|x\rangle$, denotes a column vector in a complex vector space with an inner product (i.e., a Hilbert space). It is used primarily in the context of describing the state of a quantum system (e.g., see Definition 2.1.1). A bra $\langle\cdot|$ is the dual of the ket, in that $\langle x| = |x\rangle^\dagger$, where the $\dagger$ operator denotes the conjugate transpose. Typically, the bra notation is used for operators $\langle M|$ (e.g. measurement operators) acting on a ket. This notation lends itself naturally to the notion of inner product $\langle x|x\rangle = \|x\|^2$, and matrix-vector multiplication $\langle M|x\rangle$. Furthermore, note that $|x, y\rangle$ denotes the tensor product $|x\rangle \otimes |y\rangle$, where $\otimes$ denotes the standard tensor product

of two vector spaces. The comma may be omitted, and we have numerous equivalent notations for the tensor product: e.g., $|0^2\rangle = |00\rangle = |0,0\rangle = |0\rangle\,|0\rangle = |0\rangle \otimes |0\rangle = |0\rangle^{\otimes 2}$.

2.1.2 Quantum Basics

Analogous to how a classical bit (*bit*) is a unit of classical information, a quantum bit (*qubit*) is a unit of quantum information. The difference between the two is best illustrated by considering how each is realized. A bit is realized via well separated values of a physical property of a system (e.g., voltage across an element in an electric circuit). If the value is higher than a certain threshold, the bit assumes the value 1. Otherwise, it assumes the value 0. Thus, a bit carries the information equivalent of its namesake, a binary digit. A qubit, on the other hand, is realized as a two-level quantum system; e.g., as the spin (up, down) of an electron, the polarization (horizontal, vertical) of a photon, or the discrete energy levels (ground, excited) of an ion. Consequently, it is governed by the postulates of quantum mechanics [NC10], as detailed in the following paragraphs.

The *state space* of a qubit is a 2-dimensional complex vector space, denoted as $\mathbb{C}^2$. The definition of a qubit as a *state vector* within this space is presented in the following definition.

Definition 2.1.1 (Qubit)**.** *A single (isolated) qubit is described by a state vector $|\psi\rangle$, which is a unit vector in the state space $\mathbb{C}^2$. Mathematically,*

$$|\psi\rangle = \alpha_0\,|0\rangle + \alpha_1\,|1\rangle\,, \quad \alpha_0, \alpha_1 \in \mathbb{C}, \quad |\alpha_0|^2 + |\alpha_1|^2 = 1,$$

where $|0\rangle = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ *and* $|1\rangle = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$ *are basis vectors for the state space.*

So, although we have two basis states (much as we did for the classical bit), the qubit is allowed to be in a (complex) superposition of the two, whereas a classical bit must deterministically be in one of the basis states. Additionally, as our learning examples (refer to (2.1), (2.2)) will involve multiple qubits, it is important to note that the state space of the composite system, comprising many qubits, is the tensor product of the state spaces of its components (i.e., the individual qubits). The joint state of the composite system formed by n qubits, each in state $|\psi_i\rangle\,,\ i \in \{1, \ldots, n\}$, is given by, $|\psi_1\rangle \otimes |\psi_2\rangle \otimes \cdots \otimes |\psi_n\rangle \in \mathbb{C}^{2^n}$. Now, let us express the state vector for multi-qubit system in terms of the standard basis elements.

Definition 2.1.2 (Multi-qubit system)**.** *The state vector, $|\Psi\rangle \in \mathbb{C}^{2^n}$, describing the system of n qubits can be expressed in terms of the standard basis elements $\{|b\rangle = |b_1\rangle \otimes \cdots \otimes |b_n\rangle \,|\, b = (b_1, \ldots, b_n) \in \{0,1\}^n\}$ as follows:*

$$|\Psi\rangle = \sum_{b \in \{0,1\}^n} \alpha_b\,|b\rangle\,,$$

where $\alpha_b = \langle \Psi | b \rangle \in \mathbb{C}$, $\sum_b |\alpha_b|^2 = 1$.

In contrast, the joint state of n classical bits is described by their Cartesian product. This essential distinction between Cartesian and tensor products is precisely the phenomenon of quantum entanglement, namely the existence of (pure) states of a composite system that are not product states of its parts. Quantum entanglement, alongside superposition, lies at the heart of intrinsic advantages of quantum computing.

Any manipulation of a quantum system is confined to unitary evolution. In the context of computation, this implies that *all quantum gates are unitary operators*, restricting their application to reversible computations. An avenue for irreversible computation, and the only way to obtain classical outputs in the quantum realm, is the notion of a *measurement*.

Definition 2.1.3 (Measurement)**.** *Quantum measurements are described by a collection* $\{M_m\}$ *of measurement operators acting on the state space of the system. The index* m *denotes the possible classical outcomes of the measurement. If the quantum system is in the state* $|\psi\rangle$ *before measurement, then the probability that result* m *occurs is given by* $p(m) = \langle \psi | M_m^\dagger M_m | \psi \rangle$*, and the state of the system after the measurement, if* $m^\star$ *is observed, "collapses to"* $M_{m^\star} |\psi\rangle / \sqrt{p(m^\star)}$*. To ensure the conservation of total probability,* $\sum_m M_m^\dagger M_m = I$ *is satisfied.*

Here are a couple of examples to illustrate the above definition:

- A measurement in the standard basis is implemented by measurement operators $M_0 = |0\rangle \langle 0|$ and $M_1 = |1\rangle \langle 1|$.
- A measurement in the standard basis of the state $|\Psi\rangle = \sum_{b \in \{0,1\}^n} \alpha_b |b\rangle$ yields the classical outcome b with probability $|\alpha_b|^2$.

2.1.3 PAC Learning Framework

In the classical PAC (Probably Approximately Correct) learning model [Val84], a learner is provided oracle access to samples (x, y), where x is sampled from some unknown distribution D on $\mathcal{X}$ and $y = h^\star(x)$, for some *target* hypothesis $h^\star : \mathcal{X} \to \mathcal{Y}$. We assume that $h^\star \in \mathcal{H}$, where $\mathcal{H}$ is a predefined hypothesis class, i.e., the learner has prior knowledge of $\mathcal{H}$. The goal of the learning problem is to find[1] $h : \mathcal{X} \to \mathcal{Y}$ such that the generalization error, given by the loss function $\mathcal{L}(h, D, h^\star) = \mathbb{P}_{x \sim D}(h(x) \neq h^\star(x))$, is minimized.

Definition 2.1.4 (PAC learner)**.** *An algorithm* $\mathcal{A}$ *is an* (ϵ, δ)*-PAC learner for a hypothesis class* $\mathcal{H}$ *if, for any unknown distribution* D *and for all* $h^\star \in \mathcal{H}$*,* $\mathcal{A}$ *takes in* m *pairs of labeled instances, i.e.,* $\{(x_i, h^\star(x_i))\}_{i=1}^m$*, each drawn i.i.d. from* D*, and outputs a hypothesis* h *such that*

[1] Note that h need not necessarily belong to $\mathcal{H}$. If it does, the learner is called *proper*. If not, the learner is *improper*.

$\mathbb{P}[\mathcal{L}(h, D, h^\star) \leq \epsilon] \geq 1 - \delta$*, where the outer probability is over the sequence of examples and the learner's internal randomness.*

Indeed, an (ϵ, δ)-PAC learner outputs a hypothesis that is, with high probability ($\geq 1 - \delta$), approximately correct ($\mathcal{L} \leq \epsilon$). A hypothesis class $\mathcal{H}$ is *PAC-learnable* if there exists an algorithm $\mathcal{A}$ that is an (ϵ, δ)-PAC learner for $\mathcal{H}$. In this case, we define the *sample complexity* of $\mathcal{H}$ as the minimum number of copies of quantum examples m for which there exists an (ϵ, δ)-PAC learner for $\mathcal{H}$, in the worst case over all target functions $h^\star \in \mathcal{H}$ and all distributions D.

When $\mathcal{Y} = \{0, 1\}$, we are in the setting of binary classification. To express the sample complexity of learning Boolean function classes later on, we define below a key combinatorial parameter known as the VC dimension.

Definition 2.1.5 (VC dimension)**.** *Given a hypothesis class $\mathcal{H} = \{h : \mathcal{X} \to \{0, 1\}\}$, a set $S = \{s_1, \ldots, s_t\} \subseteq \mathcal{X}$ is said to be shattered by $\mathcal{H}$ if, for every labeling $\ell \in \{0, 1\}^t$, there exists an $h \in \mathcal{H}$ such that $(h(s_1), h(s_2), \ldots, h(s_t)) = \ell$. The VC dimension of $\mathcal{H}$, VCdim($\mathcal{H}$), is the size of the largest set S that is shattered by $\mathcal{H}$.*

2.1.4 Agnostic Learning Framework

In the PAC learning framework, we worked with the *realizability assumption*, namely that $h^\star \in \mathcal{H}$. If we omit this rather strong assumption, we are able to generalize the PAC learning framework to the agnostic learning framework [KSS92]. Here, a learner is provided with oracle access to samples (x, y), sampled from some unknown distribution D on $\mathcal{X} \times \mathcal{Y}$. The learner has knowledge of a predefined hypothesis class $\mathcal{H}$. The objective of the learning problem is to find[1] $h : \mathcal{X} \to \mathcal{Y}$ such that the *regret*

$$\mathcal{R}(h, D) = \mathbb{P}_{(x,y)\sim D}(h(x) \neq y) - \inf_{h_c \in \mathcal{H}} \mathbb{P}_{(x,y)\sim D}(h_c(x) \neq y),$$

is minimized. One can notice that if the labels happen to satisfy some $h^\star \in \mathcal{H}$, $\mathcal{R} \equiv \mathcal{L}$.

Definition 2.1.6 (Agnostic learner)**.** *An algorithm $\mathcal{A}$ is an (ϵ, δ)-agnostic learner for a hypothesis class $\mathcal{H}$ if, for any unknown distribution D, $\mathcal{A}$ takes in m pairs of labeled instances, i.e., $(x_i, y_i)_{i=1}^m$, each drawn i.i.d. from D, and outputs a hypothesis h such that $\mathbb{P}[\mathcal{R}(h, D) \leq \epsilon] \geq 1 - \delta$, where the outer probability is over the sequence of examples and the learner's internal randomness.*

2.1.5 Quantum PAC and Agnostic Learning Frameworks

In the quantum setting, the primary difference from the classical setting lies in how the examples are provided. In particular, in the PAC learning setup, a quantum example [BJ95] takes the form

$$\sum_{x \in \{0,1\}^n} \sqrt{D(x)}\, |x, h^\star(x)\rangle\,, \tag{2.1}$$

for some $h^\star \in \mathcal{H}$, where $D : \{0,1\}^n \to [0,1]$ is a distribution over the instance space[2], as before. This might appear slightly strange, as a single example seemingly contains information about *all* possible classical examples. However, if we view it via the lens of measurement (see Definition 2.1.3), then it is clear that measuring a quantum example will provide the learner with a *single* classical example $(x, h^\star(x))$ with probability $D(x)$, exactly how it was in the classical PAC learning setup. While we have argued that the quantum example is a natural generalization of the classical example, the question still remains as to whether any sample complexity advantages in the quantum realm arise from the intrinsic description of a quantum example or from the quantum algorithm used or from both.

In the agnostic learning setting, a quantum example takes the form

$$\sum_{(x,y) \in \{0,1\}^{n+1}} \sqrt{D(x,y)}\, |x, y\rangle\,, \tag{2.2}$$

where, now, $D : \{0,1\}^{n+1} \to [0,1]$. These examples, like in the quantum PAC setting above, are typically *prepared* by acting on the all-zero state $|0^n, 0\rangle$ via an appropriate quantum circuit.

Given quantum examples (instead of classical examples), Definitions 2.1.4 and 2.1.6 otherwise stay exactly the same in the quantum setting. Under the quantum (batch) learning frameworks outlined here in Section 2.1.5, we investigate the sample complexity of batch learning a hypothesis class $\mathcal{H}$ in the following sections (Sections 2.1.6 and 2.2). Specifically, we address the question of how many copies of quantum examples, as given in (2.1) (resp. (2.2)), are required to (ϵ, δ)-quantum PAC (resp. quantum agnostic) learn $\mathcal{H}$.

2.1.6 Quantum Batch Binary Classification

In the binary classification setting, this question has been conclusively answered, and we reproduce the corresponding theorem below.

[2]Here, we have taken $\mathcal{X} = \{0,1\}^n$ for convenience and ease of analysis. However, any finite $\mathcal{X}$ could be mapped to this one, if needed.

Theorem 2.1.7 (Sample complexity bounds[3] for quantum batch binary classification; Theorems 23 and 25 in [AdW18])**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$. For every $\delta \in (0, 1/2)$ and $\epsilon \in (0, 1/20)$, the sample complexity of an (ϵ, δ)-quantum PAC learner (and, respectively, an (ϵ, δ)-quantum agnostic learner) for the hypothesis class $\mathcal{H}$ is given by:*

$$m^{PAC} = \Theta\left(\frac{VCdim(\mathcal{H}) + \log(\frac{1}{\delta})}{\epsilon}\right), \quad and \quad m^{agnostic} = \Theta\left(\frac{VCdim(\mathcal{H}) + \log(\frac{1}{\delta})}{\epsilon^2}\right).$$

Arunachalam and de Wolf [AdW18] established the *optimal* lower bounds via quantum state identification, employing ideas from Fourier analysis to assess the performance of the Pretty Good Measurement. They also derived *near-optimal* information-theoretic lower bounds, which are off by a factor[4] of $\frac{1}{\log(\text{VCdim}(\mathcal{H})/\epsilon)}$ in the term involving $\text{VCdim}(\mathcal{H})$ in both the PAC and agnostic settings. In particular, in the context of this thesis, we note that the *hard* distribution used to establish these lower bounds takes the form of a *near* point-mass distribution supported on the VC-shattered set.

Returning, there are two key takeaways from Theorem 2.1.7, namely:

1. No new combinatorial dimension is needed to characterize quantum batch learnability of Boolean functions, namely that the VC dimension continues to do so.
2. There is *at most a constant* sample complexity advantage for quantum batch learning of Boolean functions, in both the realizable and agnostic settings, as compared to the corresponding classical sample complexities.

Having resolved the quantum sample complexity in the (batch) binary classification setting, Arunachalam and de Wolf [AdW18] left open the natural and fundamental question of extending this characterization to the multiclass setting, i.e., when $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ with $|\mathcal{Y}| = k > 2$. This question is particularly significant because the binary case suggests a complete picture governed solely by the VC dimension, while the multiclass regime introduces qualitatively new combinatorial structure. In the following section, we resolve this question by providing a tight characterization of quantum sample complexity for multiclass classification.

2.2 Quantum Batch Multiclass Classification

In this section, we address the aforementioned question: what is the quantum sample complexity for learning concepts whose range is $[k]$ rather than $\{0,1\}$, for some $k > 2$? To state the corresponding sample complexity results in the multiclass setting, we first introduce a combinatorial

[3]The upper bounds are obtained trivially via a measure-and-learn-classically quantum learner, whereas matching lower bounds are provided by [AdW18].

[4]In [HNS24], the information-theoretic approach has since been refined to provide the *optimal* lower bounds.

parameter, the Natarajan dimension ($\mathrm{Ndim}(\cdot)$), which generalizes the VC dimension to multiclass hypothesis classes.

Definition 2.2.1 (Natarajan dimension)**.** *Given a hypothesis class $\mathcal{H} = \{h : \mathcal{X} \to [k]\}$, a set $S = \{s_1, \ldots, s_t\} \subseteq \mathcal{X}$ is said to be N-shattered by $\mathcal{H}$ if there exist two "witness" functions $f_0, f_1 : S \to [k]$ such that:*

- *For every $x \in S$, $f_0(x) \neq f_1(x)$.*
- *For every $R \subseteq S$, there exists a function $h \in \mathcal{H}$ such that*

$$\forall x \in R, h(x) = f_0(x) \text{ and } \forall x \in S \setminus R, h(x) = f_1(x).$$

The Natarajan dimension of $\mathcal{H}$, $\mathrm{Ndim}(\mathcal{H})$, is the size of the largest set S that is N-shattered by $\mathcal{H}$.

2.2.1 Lower Bounds on the Sample Complexity

Theorem 2.2.2 (Sample complexity lower bounds for quantum batch multiclass classification)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$. For every $\delta \in (0, 1/2)$ and $\epsilon \in (0, 1/20)$, the sample complexity of an (ϵ, δ)-quantum PAC learner (and, respectively, an (ϵ, δ)-quantum agnostic learner) for the hypothesis class $\mathcal{H}$ is bounded below as follows:*

$$m^{PAC} = \Omega\left(\frac{Ndim(\mathcal{H}) + \log(\frac{1}{\delta})}{\epsilon}\right), \text{ and } m^{agnostic} = \Omega\left(\frac{Ndim(\mathcal{H}) + \log(\frac{1}{\delta})}{\epsilon^2}\right).$$

At its core, the proof involves reducing the problem to the quantum binary case – establishing that a learning algorithm for $\mathcal{H}$ implies a learning algorithm for $\mathcal{H}_d$, where $\mathrm{VCdim}(\mathcal{H}_d) = \mathrm{Ndim}(\mathcal{H}) = d$. This, in turn, enables us to deduce a sample complexity lower bound for learning $\mathcal{H}$ based on the corresponding lower bound for learning $\mathcal{H}_d$. A key step in the reduction involves the following transformation of a quantum binary example[5] into a quantum multiclass example,

$$\sum_{x\in[d]} \sqrt{D(x)}\,|x, y\rangle \to \sum_{x\in[d]} \sqrt{D(x)}\,|x, f_y(x)\rangle\,, \text{ where } y \in \{0, 1\}, \text{ and } f_0, f_1 : [d] \to [k]. \quad (2.3)$$

While in the corresponding classical reduction proof, converting $(x, y) \to (x, f_y(x))$ is entirely trivial with the knowledge of x, y, f_0, f_1, performing the transformation in (2.3) using only unitary operations (in a reversible manner) in the quantum realm involves delicate reasoning using an explicit quantum circuit. In particular, it is noteworthy as its existence hinges on the reversibility of the transformation $y \leftrightarrow f_y$, which is guaranteed precisely due to the definition of N-shattering.

[5]To maintain consistency with Section 2.1.5, the input space $[d]$ can be identified with $\{0, 1\}^{\lceil \log_2 d \rceil}$.

As preliminaries for the proof, we first introduce the quantum X, CNOT, TOFFOLI gates, and quantum oracles for computing classical functions. The notation $\oplus$ refers to the classical XOR operation (i.e., addition modulo 2).

Definition 2.2.3 (X gate)**.** *X (or the Pauli-X) gate is the quantum equivalent of the classical NOT gate. It operates on one qubit, mapping $|0\rangle \to |1\rangle$ and $|1\rangle \to |0\rangle$ (i.e. it "flips" the qubit).*

Definition 2.2.4 (CNOT gate)**.** *CNOT is a quantum gate that operates on two qubits, one control and one target. If the control qubit is in the state $|1\rangle$, it flips (i.e., applies an X gate to) the target qubit.*

Definition 2.2.5 (TOFFOLI gate)**.** *TOFFOLI is a quantum gate that operates on three qubits, two control and one target. If the control qubits are both in the state $|1\rangle$, it flips (i.e., applies an X gate to) the target qubit.*

Definition 2.2.6 (Quantum oracle U_f)**.** *For classical functions $f : \{0,1\}^m \to \{0,1\}^n$, there exists*[6] *a quantum oracle U_f that performs the unitary evolution*

$$U_f |x, y\rangle = |x, y \oplus f(x)\rangle ,$$

for $x \in \{0,1\}^m$ and $y \in \{0,1\}^n$.

Proof. (of Theorem 2.2.2) Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class of Natarajan dimension d and let $\mathcal{H}_d = \{0,1\}^{[d]}$. Let $\mathcal{A}$ be a *quantum* PAC (corresp. *quantum* agnostic) learning algorithm for $\mathcal{H}$. We proceed to show that it is possible to construct a quantum PAC (corresp. quantum agnostic) learning algorithm, $\bar{\mathcal{A}}$, for $\mathcal{H}_d$. Therefore, by reduction, we would obtain $m_{\bar{\mathcal{A}},\mathcal{H}_d} \leq m_{\mathcal{A},\mathcal{H}}$, and thus $m^{\text{PAC}}_{\mathcal{H}_d} \leq m^{\text{PAC}}_{\mathcal{H}}$ (corresp. $m^{\text{agnostic}}_{\mathcal{H}_d} \leq m^{\text{agnostic}}_{\mathcal{H}}$). Since, by construction, $\text{VCdim}(\mathcal{H}_d) = d = \text{Ndim}(\mathcal{H})$, the reduction allows us to obtain the sample complexity lower bounds being proven here (for the multiclass case), from the corresponding lower bounds for quantum batch binary classification (Theorem 2.1.7). Now, for the key step of the proof, given a *quantum* learner $\mathcal{A}$ for $\mathcal{H}$, it is possible to construct a *quantum* learner $\bar{\mathcal{A}}$, for $\mathcal{H}_d$, as follows. We show this for the *quantum PAC* case, and comment here that this reduction in the *quantum agnostic* case will proceed identically.

The learner $\bar{\mathcal{A}}$ receives m-copies of the quantum example $\sum_{x\in[d]} \sqrt{D(x)}\,|x, y\rangle$, where $(x, y) \in [d] \times \{0,1\}$ and $D : [d] \to [0,1]$ is an *arbitrary* distribution on $[d]$. Now, let $S = \{s_1, \ldots, s_d\} \subseteq \mathcal{X}$ be a set and f_0, f_1 be the functions that witness the N-shattering of S by $\mathcal{H}$. The learner $\bar{\mathcal{A}}$ will now attempt to convert[7] each of its m-copies of $\sum_{x\in[d]} \sqrt{D(x)}\,|x, y\rangle$ to $\sum_{x\in[d]} \sqrt{D(x)}\,|s_x, f_y(s_x)\rangle$. However, as s_x is simply an indexing of the elements of the set S, without loss of generality, we let

[6] In fact, the quantum oracle U_f can be implemented in a rather straightforward way, by using the truth table of f and generalizations of the CNOT gate that use several qubits as controls.

[7] The learner $\bar{\mathcal{A}}$ will then present these transformed examples to $\mathcal{A}$, the quantum PAC learner for $\mathcal{H}$.

$\bar{\mathcal{A}}$ convert each of its m-copies of $\sum_{x\in[d]} \sqrt{D(x)}\,|x,y\rangle$ to $\sum_{x\in[d]} \sqrt{D(x)}\,|x,f_y(x)\rangle$ instead. We claim that the transformation

$$|\psi\rangle_{\bar{\mathcal{A}}} = \sum_{x\in[d]} \sqrt{D(x)}\,|x,y\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x,f_y(x)\rangle = |\psi\rangle_{\mathcal{A}}\,, \tag{2.4}$$

is attainable. Indeed, the quantum circuit shown in Figure 2.1 (and described subsequently) performs the following augmented transformation,

$$\overline{|\psi\rangle_{\bar{\mathcal{A}}}} = \sum_{x\in[d]} \sqrt{D(x)}\,|x,y,0^{3\lceil\log_2 k\rceil}\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x,0,0^{2\lceil\log_2 k\rceil},f_y(x)\rangle = \overline{|\psi\rangle_{\mathcal{A}}}. \tag{2.5}$$

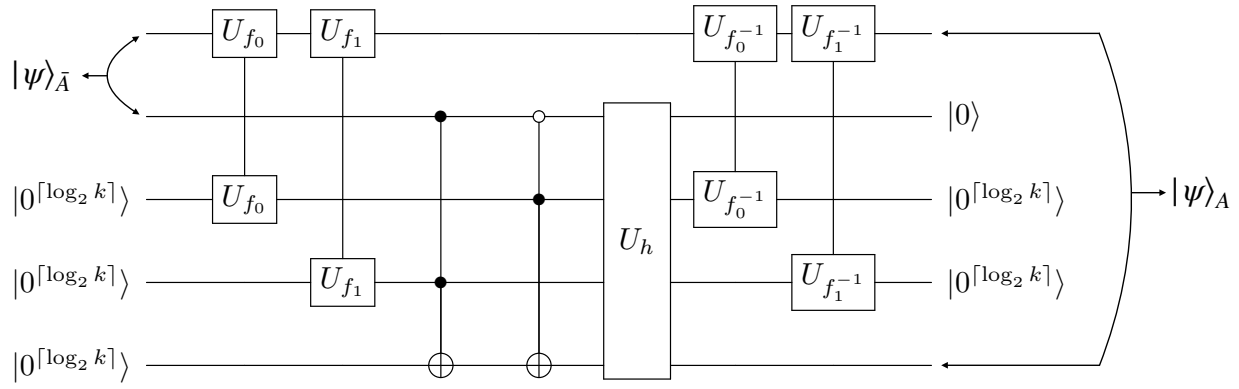

Figure 2.1: A quantum circuit that performs the transformation in (2.5). The first row denotes the $\lceil\log_2 d\rceil$ qubits involved in encoding x. The second row denotes the single qubit involved in encoding y. The third and fourth row each denote the $\lceil\log_2 k\rceil$ ancillary qubits designed to hold the intermediate computation of $f_0(x)$ and $f_1(x)$ respectively. The fifth row denotes the $\lceil\log_2 k\rceil$ qubits designed to hold the output $f_y(x)$.

We form the augmented state $\overline{|\psi\rangle_{\bar{\mathcal{A}}}}$ by appending $3\lceil\log_2 k\rceil$ qubits in the state $|0\rangle$ to the input $|\psi\rangle_{\bar{\mathcal{A}}}$. We intend to use one set of $\lceil\log_2 k\rceil$ qubits to encode each of $f_0(x)$, $f_1(x)$ (both ancillary) and $f_y(x)$ (solution). First, we pass the qubits encoding x and one set of $\lceil\log_2 k\rceil$ qubits in the state $|0\rangle$ to the quantum oracle U_{f_0}. From this, we transform

$$\overline{|\psi\rangle_{\bar{\mathcal{A}}}} = \sum_{x\in[d]} \sqrt{D(x)}\,|x,y,0^{3\lceil\log_2 k\rceil}\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x,y,f_0(x),0^{2\lceil\log_2 k\rceil}\rangle\,.$$

Next, we pass the qubits encoding x and another set of $\lceil\log_2 k\rceil$ qubits in the state $|0\rangle$ to the quantum oracle U_{f_1}. From this, we transform

$$\sum_{x\in[d]} \sqrt{D(x)}\,|x,y,f_0(x),0^{2\lceil\log_2 k\rceil}\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x,y,f_0(x),f_1(x),0^{\lceil\log_2 k\rceil}\rangle\,.$$

Now, we apply $\lceil \log_2 k \rceil$ TOFFOLI gates to each set of the qubit encoding y, a qubit involved in encoding $f_1(x)$ and a qubit $|0\rangle$ in the set of $\lceil \log_2 k \rceil$ remaining qubits that have not yet been operated on (that are designed hold the final result $f_y(x)$). From this, we transform

$$\sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), 0^{\lceil \log_2 k \rceil}\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), y.f_1(x)\rangle\,.$$

Next, we apply the X gate[8] to the qubit encoding y, and then apply $\lceil \log_2 k \rceil$ TOFFOLI gates to each set of the qubit (now encoding) $1 - y$, a qubit involved in encoding $f_0(x)$ and a qubit that is holding the final result. Finally, we apply the X gate again to the qubit encoding $1 - y$ to revert the effect of the first X gate. All together, we have transformed

$$\begin{aligned}
\sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), y.f_1(x)\rangle &\mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, 1-y, f_0(x), f_1(x), y.f_1(x)\rangle \\
&\mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, 1-y, f_0(x), f_1(x), y.f_1(x) \oplus (1-y).f_0(x)\rangle \\
&\mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), y.f_1(x) \oplus (1-y).f_0(x)\rangle \\
&= \sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), f_y(x)\rangle\,,
\end{aligned}$$

where, in the last line, we recognize that $y.f_1(x) \oplus (1-y).f_0(x) = f_y(x)$.

Now that we have computed our solution, $f_y(x)$, to ensure we are not introducing extraneous outputs to satisfy reversibility of unitary computation, we must *uncompute* all ancillary qubits to their original form. Furthermore, we also do not want the input $|y\rangle$ in our output[9], and would like to "remove" it, by transforming it to the $|0\rangle$ state. We proceed with the *uncomputation* as follows. First, we claim, with the following subproof, that there exists a (genuine) boolean *function* h that takes as input $f_0(x), f_1(x)$, and $f_y(x)$ and outputs y.

> For a short (sub)proof-by-contradiction, consider a particular input $f_0(x), f_1(x)$, and $f_y(x) = y.f_1(x) \oplus (1-y).f_0(x)$ (for some x) that maps to both $y = 0$ and $y = 1$. This would imply that $f_0(x) = f_1(x)$ for that particular x. However, as f_0 and f_1 are witnesses of the N-shattering of $\mathcal{H}$, they must disagree on all inputs x, providing us with the contradiction.

As h is a genuine (boolean) function, we can apply the oracle U_h to the qubits encoding $f_0(x), f_1(x)$,

[8]The concise description for controlling on $|y\rangle = |0\rangle$ (instead of explicitly on $|1-y\rangle = |1\rangle$) is depicted in Figure 2.1.

[9]We do not want to allow $\mathcal{A}$ to be able to cheat by providing it with this additional knowledge.

and $f_y(x)$ and the qubit encoding y. This gives us the transformation[10]

$$\sum_{x\in[d]} \sqrt{D(x)}\,|x, y, f_0(x), f_1(x), f_y(x)\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, y\oplus y, f_0(x), f_1(x), f_y(x)\rangle$$
$$= \sum_{x\in[d]} \sqrt{D(x)}\,|x, 0, f_0(x), f_1(x), f_y(x)\rangle\,.$$

Lastly, we apply $U_{f_0^{-1}}$ to the qubits encoding x and the ancillary qubits encoding $f_0(x)$. And then, we apply $U_{f_1^{-1}}$ to the qubits encoding x and the ancillary qubits encoding $f_1(x)$. All together, we have transformed,

$$\sum_{x\in[d]} \sqrt{D(x)}\,|x, 0, f_0(x), f_1(x), f_y(x)\rangle \mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, 0, 0^{\lceil \log_2 k\rceil}, f_1(x), f_y(x)\rangle$$
$$\mapsto \sum_{x\in[d]} \sqrt{D(x)}\,|x, 0, 0^{2\lceil \log_2 k\rceil}, f_y(x)\rangle = \overline{|\psi\rangle}_{\mathcal{A}}.$$

Now that we have performed the transformation in (2.5) quantumly, ignoring[11] ancillary (and the "removed" y) qubits in the output, we note that the output of the transformation in (2.4) is attained.

The learner $\bar{\mathcal{A}}$ now feeds the m-copies of $\sum_{x\in[d]} \sqrt{D(x)}\,|x, f_y(x)\rangle$ as input to the *quantum* learner $\mathcal{A}$ and obtains (from $\mathcal{A}$) a *classical* (black box) function[12] $g : S \to \mathcal{Y}$ (recall $S = \{s_1, \ldots, s_d\} \subseteq \mathcal{X}$ is the subset that is N-shattered by $\mathcal{H}$). Finally, $\bar{\mathcal{A}}$ outputs the hypothesis (as a black box function) $f : [d] \to \{0, 1\}$, given by $f(i) = 1$ if and only if $g(s_i) = f_1(s_i)$ (which by construction learns $\mathcal{H}_d$).

■

2.2.2 Upper Bounds on the Sample Complexity

In general, classical sample complexity upper bounds trivially translate to the corresponding quantum ones, as the quantum learner always has the option of simply performing a measurement on each quantum example, and perform the classical learning algorithm on the resulting m classical examples. We include the theorem statement (Theorem 2.2.7) and proof below for completeness.

Theorem 2.2.7 (Sample complexity upper bounds for quantum batch multiclass classification)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$. The sample complexity of an (ϵ, δ)-quantum PAC learner (and, respectively, an (ϵ, δ)-quantum agnostic learner) for the hypothesis class $\mathcal{H}$ is bounded above as*

[10] In other words, we have been able to "remove" y. It is important to note that this was possible precisely because y was recoverable from f_y (rendering $y \leftrightarrow f_y$ reversible), ensuring its ability to be encoded in a unitary operation.

[11] They are each deterministically $|0\rangle$.

[12] In fact, our identification $s_x \leftrightarrow x$ earlier means that $g : [d] \to \mathcal{Y}$ instead, and $f : [d] \to \{0, 1\}$ would then be given by $f(i) = 1$ if and only if $g(i) = f_1(i)$.

follows:

$$m^{PAC} = \mathcal{O}\left(\frac{Ndim(\mathcal{H})\log(k)\log(\frac{1}{\epsilon}) + \log(\frac{1}{\delta})}{\epsilon}\right); \quad m^{agnostic} = \mathcal{O}\left(\frac{Ndim(\mathcal{H})\log(k) + \log(\frac{1}{\delta})}{\epsilon^2}\right).$$

Proof. The quantum PAC (resp. quantum agnostic) learner performs a measurement on each of the m examples $\sum_x \sqrt{D(x)}\,|x, y\rangle$ (corresp. $\sum_{x,y} \sqrt{D(x,y)}\,|x, y\rangle$) in the standard computational basis. This provides m classical examples, i.e. gives us the training set $\{(x_i, y_i)\}_{i=1}^m$, where a given (x_i, y_i) appears with probability $D(x_i)$ (corresp. $D(x_i, y_i)$). Now, the quantum PAC (resp. quantum agnostic) learners calls upon an (ϵ, δ)-PAC (resp. agnostic) *classical* learner to learn on the training set $\{(x_i, y_i)\}_{i=1}^m$, and outputs the resulting *classically learned* hypothesis. Thus, the classical sample complexity sufficiency requirements [DSBDSS15, BDCBHL95] continue to hold, and our proof is complete. ■

Before we proceed, we take a moment to summarize the sample complexity of classical and quantum batch learning that we have discussed thus far.

	Boolean		**Multiclass**	
	Realizable	**Agnostic**	**Realizable**	**Agnostic**
Classical	$\Theta\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon}\right)$ [BEHW89, Han16]	$\Theta\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ [KSS92, Tal94]	$\Omega\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon}\right)$ [Nat89] $\mathcal{O}\left(\frac{d\log(k)\log(\frac{1}{\epsilon})+\log(\frac{1}{\delta})}{\epsilon}\right)$ [DSBDSS15]	$\Omega\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ [BDCBHL95] $\mathcal{O}\left(\frac{d\log(k)+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ [BDCBHL95]
Quantum	$\Theta\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon}\right)$ [AdW18]	$\Theta\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ [AdW18]	$\Omega\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon}\right)$ (Thm. 2.2.2) $\mathcal{O}\left(\frac{d\log(k)\log(\frac{1}{\epsilon})+\log(\frac{1}{\delta})}{\epsilon}\right)$ (Thm. 2.2.7)	$\Omega\left(\frac{d+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ (Thm. 2.2.2) $\mathcal{O}\left(\frac{d\log(k)+\log(\frac{1}{\delta})}{\epsilon^2}\right)$ (Thm. 2.2.7)

Table 2.1: An overview of sample complexity results for **batch** learning in classical and quantum paradigms. Our novel contributions are presented in boxes shaded gray. d denotes the VC dimension for Boolean cases ($d := \text{VCdim}(\mathcal{H})$), and the Natarajan dimension for Multiclass cases ($d := \text{Ndim}(\mathcal{H})$).

In [DSBDSS15], the following classical multiclass agnostic sample complexity upper bound

$$m^{\text{PAC}} = \mathcal{O}\left(\frac{\text{Ndim}(\mathcal{H})(\log(k) + \log(\frac{1}{\epsilon}) + \log(\text{Ndim}(\mathcal{H}))) + \log(\frac{1}{\delta})}{\epsilon}\right),$$

was shown to hold which has a tighter dependence on ϵ, but a looser dependence on $\text{Ndim}(\mathcal{H})$, as compared to the one we have worked with (see Table 2.1). However, the proof of Theorem 2.2.7 would naturally extend this bound to the quantum setting as well. Indeed, the *measure-and-*

learn-classically argument used there allows *any* classical sample complexity upper bounds to be immediately carried over to the quantum setting. We discuss a notable consequence of this observation in the following section (Section 2.2.3).

2.2.3 Towards a Full Characterization of Quantum Multiclass Learnability for Unbounded Label Spaces

In both the realizable and agnostic settings, our $\mathrm{Ndim}(\cdot)$-based sample complexity upper bounds carry a $\log(k)$ factor that is absent from the corresponding lower bounds. Consequently, these bounds do not yield a full characterization of quantum multiclass learnability in the case of an unbounded label space (i.e., as $k \to \infty$). This gap has since been resolved in the *classical* setting by Brukhim et al. [BCD+22], who showed that the Daniely–Shalev-Schwartz (DS) dimension[13] – rather than the Natarajan dimension alone – provides a full characterization of multiclass sample complexity, including when the label space is unbounded. In particular, they proved that finiteness of the DS dimension is both necessary and sufficient for multiclass PAC learnability, and exhibited a multiclass hypothesis class with finite Natarajan dimension that is nevertheless not PAC-learnable (Theorem 2 of [BCD+22]). This demonstrates that the k-dependence inherent in $\mathrm{Ndim}(\cdot)$-based bounds cannot, in general, be removed.

At the same time, the Natarajan dimension remains fundamentally relevant in the agnostic setting. Indeed, subsequent work [CEH+25] showed that the sharpest known agnostic multiclass sample complexity bounds depend simultaneously on both the DS and Natarajan dimensions, indicating that the two notions capture genuinely different aspects of multiclass complexity.

While this thesis does not focus on the DS dimension directly, we note that the classical upper bounds involving the DS dimension can be transferred to the quantum setting via the *measure-and-learn-classically* reduction discussed in Theorem 2.2.7. The classical DS-dimension lower bound argument (Theorem 2 of [DSS14]) proceeds through a transductive learning framework, where the learner is given access to an unlabeled sample and must infer the corresponding labels. In the quantum example model, however, there is no clear analogue of separating the "instance space" from the labels, since quantum examples are presented as joint superpositions over both registers (see Examples 2.1 and 2.2), and accessing or measuring one component generally disturbs the other. Consequently, the transductive techniques underlying the classical DS-dimension lower bounds do not appear to admit a straightforward quantum analogue. This obstruction is primarily methodological, in that, it rules out a straightforward adaptation of the known classical proof, but does not preclude the possibility of obtaining quantum DS-dimension lower bounds through different techniques, such as direct information-theoretic arguments. Resolving this gap would be a

[13] Informally, the DS dimension refines the combinatorial structure captured by the Natarajan dimension by controlling how flexibly hypotheses can realize alternative labelings across a set of inputs.

significant step toward a characterization of quantum multiclass learnability in the $k \to \infty$ regime. We therefore identify the following as an open question for future work:

- What is the *tight* quantum sample complexity bound for batch multiclass learning, in both the realizable and agnostic settings, when the label space is unbounded (i.e., when $k \to \infty$)?

Chapter Summary and Related Work. Quantum learning of classical functions in the batch setting has has predominantly focused on learning Boolean functions (see [AdW17] for a survey). In our work, the multiclass extension in the batch setting (Section 2.2) is inspired by an open question of [AdW18]. The lower bound (Section 2.2.1) in this extension is made possible due to the presence of an explicit quantum circuit (see Figure 2.1) that allows for a *black box*-style reduction of the quantum multiclass case to the quantum binary classification case, enabling the existing quantum batch binary lower bound (and, techniques therewith, see Section 2.1.6 for a detailed treatment) in [AdW18] to yield the desired multiclass bound. This reduction is analogous to the establishment of the lower bound in the classical batch multiclass setting (as presented in Theorem 5 of [DSBDSS15]). While the batch classification model in this thesis assumes the learner has access to copies of quantum states $|\psi_c\rangle$ (i.e., an *example oracle*), with no (i.e., at most a constant) advantage over classical batch models, Grover's algorithm – along with its quadratic quantum sample complexity advantage – has been successfully adapted in [SSG24] to batch *binary* classification in a stronger learning model where the learner has access to a quantum circuit Q_c that generates a quantum sample $|\psi_c\rangle$ (i.e., a *unitary oracle*[14]).

Furthermore, in this chapter, we do not explicitly address whether learning is proper or improper in the various settings. However, since most of our proofs involve reductions to the corresponding classical settings, the nature of the algorithms that provide tight bounds in those settings remains unchanged in our cases. The separation between proper and improper quantum learning has garnered significant interest in the community [ABC$^+$20, NS24]. In particular, this separation has been studied in the context of quantum extensions of the coupon collector problem, which provides a classical example of such a separation. Notably, [NS24] established a provable separation between proper and improper learning in the quantum setting using a "padded" variant of the quantum coupon collector problem. On the classical side, [BHMZ20] provides a full characterization of when proper learning achieves optimal sample complexity, based on a combinatorial parameter known as the dual Helly number. To the best of our knowledge, an analogous characterization for quantum proper learning remains an open question.

[14]The sample complexity here refers to the number of calls to the unitary oracle, providing access to both Q_c and its Hermitian conjugate $Q_c^\dagger$

CHAPTER 3

Online Learning with Adversarially Chosen Distributions and Quantum Examples

So far, we have been working with learning in the batch setting, where we are provided with all the examples at once[1]. For several practical applications, it is either impossible to obtain all the examples at once (e.g., recommendation systems), or we simply wish to evolve our learning over time. In these cases, *online learning* [Lit88] – where we iteratively improve our hypothesis using examples we receive over time, and using our current hypothesis to predict for the upcoming example – is the appropriate framework to be placing ourselves in.

Unlike the batch setting, quantum online learning of classical functions, to the best of our knowledge, has no predefined model. One possible explanation is that we need, as an intermediary, a new classical online learning model (see Section 3.2) where, at each round, the adversary provides a *distribution* over the example (input-label) space instead of a single example. With this new classical model, and the definition of a quantum example, a model for online learning in the quantum setting arises as a natural extension (see Figure 3.1).

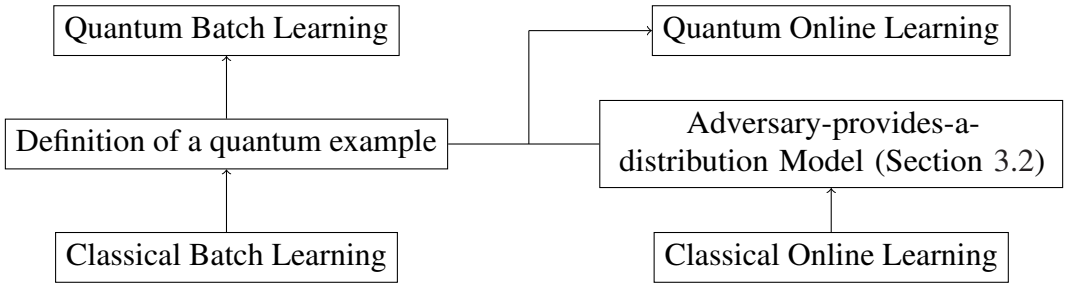

Figure 3.1: Mapping of the tools necessary for generalizations of learning paradigms from classical to quantum.

To bring out this narrative, we proceed chronologically. In the following section (Section 3.1), we start with an overview of known models in classical online learning, and discuss the obstacles to a direct quantum generalization of these models, motivating why the intermediary model is needed.

[1]This is typical for most settings where we are trying to learn a hypothesis via inductive reasoning (e.g. learning a function to fit data, etc.)

With ease of exposition in mind, we begin with a focus on Boolean function classes in the realizable setting.

This chapter is based largely on material from [MT25], particularly Sections 4 and 5 of that paper. Minor revisions have been made to the exposition, notation, and organization for inclusion in this thesis.

3.1 Canonical Classical Online Model: Adversarially Chosen Inputs

Let $\mathcal{C} := \{c : \mathcal{X} \to \{0,1\}\}$, and $\mathcal{H} \subseteq \mathcal{C}$ (i.e., $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$). A protocol for online learning is a T-round procedure described as follows: at the t-th round,

1. Adversary provides input point in the domain: $x_t \in \mathcal{X}$.
2. Learner uses a hypothesis[2] $h_t \in \mathcal{C}$, and makes the prediction $\hat{y}_t = h_t(x_t) \in \{0,1\}$.
3. Adversary provides the input point's label, $y_t = h^\star(x_t)$, where $h^\star \in \mathcal{H}$.
4. Learner suffers a loss of 1 (a 'mistake'), if $\hat{y}_t \neq y_t$, i.e. $\mathcal{L}_{\mathbb{I}}(h_t, x_t, h^\star) = \mathbf{1}[h_t(x_t) \neq h^\star(x_t)]$.

Therefore, the learner's total loss is given by,

$$\mathcal{L}_{\mathbb{I}}(\mathbf{h}, \mathbf{x}, h^\star) = \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, x_t, h^\star) = \sum_{t=1}^{T} \mathbf{1}[h_t(x_t) \neq h^\star(x_t)], \tag{3.1}$$

where we use $\mathbf{h} := (h_1, \ldots, h_T)$ to denote the sequence of hypotheses that the learner uses, and $\mathbf{x} := (x_1, \ldots, x_T)$ to denote the sequence of instances that the adversary provides. The subscript $\mathbb{I}$ (in $\mathcal{L}_{\mathbb{I}}$) indicates that this is the input-based indicator (0-1) loss function[3].

The learner chooses an algorithm $\mathcal{A}$ that will generate the sequence $\mathbf{h}_{\mathcal{A}}$ following the protocol above. The learner's goal is to minimize $\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}}, \mathcal{H}) = \sup_{\mathbf{x}, h^\star \in \mathcal{H}} \mathbb{E}[\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}}, \mathbf{x}, h^\star)]$, i.e., make as few mistakes, on average, as possible regardless of the adversary's (potentially worst-case) choices of sequence of instances $\mathbf{x} := (x_1, \ldots, x_T)$ and labeling function $h^\star \in \mathcal{H}$. For the subsequent bound on $\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}}, \mathcal{H})$, we first define the combinatorial parameter, Littlestone dimension, $\mathrm{Ldim}(\mathcal{H})$.

Definition 3.1.1 (Littlestone dimension)**.** *Let T be a rooted tree whose internal nodes are labeled by elements from $\mathcal{X}$. Each internal node's left edge and right edge are labeled 0 and 1, respectively. The tree T is L-shattered by $\mathcal{H}$ if, for every path from root to leaf which traverses the nodes $x_1, \ldots, x_d$, there exists a hypothesis $h \in \mathcal{H}$ such that, for all i, $h(x_i)$ is the label of the edge (x_i, x_{i+1}). We define the Littlestone dimension, Ldim($\mathcal{H}$), to be the maximal depth of a complete binary tree that is L-shattered by $\mathcal{H}$.*

[2]Note that the learner may choose a hypothesis $h_t \in \mathcal{C} \setminus \mathcal{H}$, i.e., we do not require the learner to be *proper*.

[3]This is distinct from the probabilistic loss function $\mathcal{L}_P$ that we will encounter later in Section 3.2.

The classical online learning model in this section (Section 3.1) has been thoroughly studied, and the following theorem characterizes it in terms of the Littlestone dimension.

Theorem 3.1.2 (Bounds on $\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}}, \mathcal{H})$ for the canonical classical online model; Corollary 21.8 in [SSBD14] and Theorem 24 in [DSBDSS15])**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$ be a hypothesis class. The Standard Optimal Algorithm (SOA)*[4] *is a deterministic algorithm that achieves a worst-case total loss of Ldim$(\mathcal{H})$, i.e. $\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{SOA}, \mathcal{H}) =$ Ldim$(\mathcal{H})$. Furthermore, for any algorithm $\mathcal{A}$, the total expected loss on the worst-case sequence is at least*[5] $\frac{1}{2} \cdot$ *Ldim$(\mathcal{H})$, i.e. $\mathcal{L}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}}, \mathcal{H}) \geq \frac{1}{2} \cdot$ Ldim$(\mathcal{H})$.*

3.1.1 Obstacles to a Direct Quantum Generalization

Given the popularity and widespread applications of the classical online model, we explore the feasibility of developing a quantum version of the above online model to ultimately inquire whether such a quantum adaptation would be any more powerful from the perspective of the learner and/or the adversary. In essence, as there exists a well-defined "landscape" for classical and quantum *batch* learning, we seek to delineate the analogous landscape in the *online* learning context.

To this, if one attempts to naïvely generalize the above classical model to the quantum setting, an obvious issue arises: the quantum examples of the form (2.1) do not split the input-label pair. In particular, an adversary cannot temporally separate its provision of the input point and its label. A first step towards a model that can be generalized to the quantum setting, then, is to reorder the steps at the t-th round to $2, 1 \,\&\, 3, 4$ (i.e., where the learner provides a prediction $\hat{y}_t$ after which the adversary presents *both* the input and its label (x_t, y_t)).

While this reordering gives an entirely equivalent model that is, once again, characterized by Littlestone dimension (Theorem 3.1.2), it is not sufficient for a natural quantum generalization. The issue now is that a *classical* adversary only ever presents one (classical) example at each round. How do we go about generalizing a single classical example to a *quantum* adversary's (quantum) example that, in general, sits in superposition? It appears futile to attempt to do so. The missing piece, evidently, is the lack of a notion of a distribution over examples in the classical online model(s) examined so far.

3.2 An Intermediate Classical Online Model: Adversarially Chosen Distributions

Now that we have identified the unfilled gap to transition to the quantum setting, we first state the appropriate *classical* generalization of the canonical classical model (in Section 3.1) by asking

[4] At round t, given input x_t, the SOA predicts $\hat{y}_t \in \{0,1\}$ that maximizes the Littlestone dimension of the version space consistent with $\hat{y}_t$.

[5] The adversary traverses the shattered tree and provides, at every round, the label that the (randomized) algorithm $\mathcal{A}$ is less likely to predict.

the adversary to, at each t, choose a distribution over a set of input-label pairs, from which an explicit input-label pair is then drawn. The protocol for the T-round procedure will be as follows: at the t-th round,

1. Learner provides a hypothesis $h_t \in \mathcal{C}$.
2. Adversary chooses a distribution $D_t : \mathcal{X} \to [0, 1]$ on the instance space, draws $x \sim D_t$, and reveals $(x, h^\star(x))$ to the learner, where $h^\star \in \mathcal{H}$.
3. Learner suffers, but does not "see", a loss of $\mathcal{L}_P(h_t, D_t, h^\star) := \mathbb{P}_{x \sim D_t}(h_t(x) \neq h^\star(x))$.

Here, the learner's total loss is given by,

$$\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) = \sum_{t=1}^{T} \mathcal{L}_P(h_t, D_t, h^\star) = \sum_{t=1}^{T} \mathbb{P}_{x \sim D_t}(h_t(x) \neq h^\star(x)), \tag{3.2}$$

where we additionally use $\mathbf{D} := (D_1, \ldots, D_T)$ to denote the sequence of distributions that the adversary chooses. Analogously, the learner's objective is to choose $\mathbf{h}$ to minimize $\mathcal{L}_P(\mathbf{h}, \mathcal{H}) = \sup_{\mathbf{D},\, h^\star \in \mathcal{H}} \mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)]$.

We identify this model as the adaptive adversary variant of the online learning model recently considered in Dawid and Tewari [DT22]. Note that, if we restrict the adversary, allowing it to choose only point masses, we recover the *reordered* model in Section 3.1.

From a learning standpoint, the adversary-provides-a-distribution model differs fundamentally from the canonical model in that the learner, here, does not have full information about its own loss at any given round. Since the learner does not know D_t, it cannot compute $\mathcal{L}_P(h_t, D_t, h^\star)$ for any t. In other words, the learner seeks to minimize a quantity that it cannot even compute. This *partial information* setting here, at least at first, appears to be more challenging for the learner as it not only grapples with the inability to compute its loss but also contends with the larger space available to the adversary for its choices ($D_t \in [0, 1]^{\mathcal{X}}$ vs. $x_t \in \{0, 1\}^{\mathcal{X}}$).

However, as we will soon illustrate, this perceived challenge proves not to be the case. The key factor influencing this distinction lies in the learner's ability to calculate $\mathcal{L}_{\mathbb{I}}(\mathbf{h}, \mathbf{x}, h^\star)$ for the observed sequence of examples $\mathbf{x}$, providing an unbiased estimator for its total loss $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)$. We demonstrate that it is indeed (necessary and) sufficient for a learner in the adversary-provides-a-distribution model to execute SOA on the observed sequence of examples $\mathbf{x}$ to achieve a bound analogous to that in the canonical model (cf. Theorem 3.1.2). Before delving into the results, we formally define what a learner, an adversary, and learnability entails for the adversary-provides-a-distribution model we have just discussed. We also introduce some definitions from probability theory that will be used throughout the subsequent proofs.

3.2.1 Definitions and Preliminaries

3.2.1.1 Learner and Learnability

Definition 3.2.1 (Classical online learner)**.** *An algorithm $\mathcal{A}$ is a classical online learner for a hypothesis class $\mathcal{H} \subseteq \mathcal{C}$ if having received a sequence of examples over the first t rounds, $(x_1, h^\star(x_1)), \ldots, (x_t, h^\star(x_t))$ where $x_i \sim D_i$ with D_i arbitrary (unknown), $\mathcal{A}$ outputs a hypothesis $h_{t+1} \in \mathcal{C}$ at round*[6] *$t+1$.*

Definition 3.2.2 (Classical adversary)**.** *Having received a sequence of hypothesis $\mathbf{h}|_t = (h_1, \ldots, h_t)$ from the learner, and a sequence of examples $\mathbf{x}|_t = (x_t, \ldots, x_t)$ drawn previously from its own prior choices of distributions $\mathbf{D}|_t = (D_1, \ldots, D_t)$ over the first t rounds, at round $t+1$, a classical (online) adversary chooses a distribution $D_{t+1} : \mathcal{X} \to [0,1]$ on the instance space, draws $x_{t+1} \in D_{t+1}$ and reveals $(x_{t+1}, h^\star(x_{t+1}))$ to the learner, where $h^\star \in \mathcal{H}$ is consistent with all preceding labeled examples.*

Definition 3.2.3 (Classical online learnability)**.** *A hypothesis class $\mathcal{H}$ is classical online learnable if there exists a classical online learning algorithm $\mathcal{A}$ such that:*
$\mathcal{L}_P(\mathbf{h}_\mathcal{A}, \mathcal{H}) = \sup_{\mathbf{D},\, h^\star \in \mathcal{H}} \mathbb{E}[\mathcal{L}_P(\mathbf{h}_\mathcal{A}, \mathbf{D}, h^\star)] = o(T)$.

3.2.1.2 Probability Theory Basics

Let $\mathbf{X} := (X_t)_{t \geq 1}$ be a sequence of random variables defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where each X_t represents the outcome of an experiment at time t. Associated with this sequence is a filtration $\mathbb{F} := (\mathcal{F}_t)_{t \geq 1}$, an increasing family of σ-algebras such that $\mathcal{F}_t$ captures the information available up to time t. Formally, $\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \cdots \subseteq \mathbb{F}$, where $\mathcal{F}_t$ contains events determined by $\mathbf{X}|_t = (X_1, X_2, \ldots, X_t)$.

Definition 3.2.4 (Martingale)**.** *A sequence of random variables $\mathbf{M} := (M_t)_{t \geq 1}$ is said to be a martingale with respect to the filtration $\mathbb{F} := (\mathcal{F}_t)_{t \geq 1}$ if it satisfies the following conditions:*

1. *Adaptation: M_t is $\mathcal{F}_t$-measurable for all $t \geq 1$ (i.e., we say $\mathbf{M}$ is adapted to $\mathbb{F}$),*
2. *Integrability: $\mathbb{E}[|M_t|] < \infty$ for all $t \geq 1$,*
3. *Martingale Property: $\mathbb{E}[M_{t+1}|\mathcal{F}_t] = M_t$ for all $t \geq 1$.*

Definition 3.2.5 (Martingale Difference Sequence)**.** *Given a martingale $\mathbf{M} := \{M_t\}_{t \geq 1}$, the corresponding martingale difference sequence is defined as $\mathbf{D} := \{D_{t+1}\}_{t \geq 1}$, where $D_{t+1} = M_{t+1} - M_t$. By the martingale property, the sequence $\mathbf{D}$ satisfies $\mathbb{E}[D_{t+1}|\mathcal{F}_t] = 0$ for all $t \geq 1$, indicating that each increment D_{t+1} has conditional mean zero given the information up to time t.*

[6] Prior to receiving any examples, $\mathcal{A}$ outputs some arbitrary hypothesis $h_1 \in \mathcal{C}$ at round 1.

In general, any sequence of random variables $\mathbf{X} := \{X_t\}_{t\geq 1}$ is a martingale difference sequence with respect to the filtration $\mathbb{F} := (\mathcal{F}_t)_{t\geq 1}$ if it satisfies the following conditions:

1. *Adaptation: X_t is $\mathcal{F}_t$-measurable for all $t \geq 1$ (i.e. $\mathbf{X}$ is adapted to $\mathbb{F}$),*
2. *Integrability: $\mathbb{E}[|X_t|] < \infty$ for all $t \geq 1$,*
3. *Conditional Mean-Zero Property: $\mathbb{E}[X_t|\mathcal{F}_{t-1}] = 0$ for all $t \geq 1$*[7].

With these definitions and preliminaries in place, our next objective is to characterize learnability in the adversary-provides-a-distribution framework. When we later introduce the quantum online learning model (Section 3.3), these classical insights will serve as a foundation, enabling us to draw direct comparisons and understand the strong links connecting the *classical* adversary-provides-a-distribution model to the *quantum* online learning setup.

3.2.2 Binary Classification in the Realizable Setting

Theorem 3.2.6 (Upper bound on the expected loss for the classical adversary-provides-a-distribution model)**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$ be a hypothesis class, and $h^\star \in \mathcal{H}$. For every adversary, there exists a classical online learner for $\mathcal{H}$ that satisfies*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathcal{O}(\mathit{Ldim}(\mathcal{H})).$$

Proof. Let $\mathbf{D}$ and $h^\star \in \mathcal{H}$ be arbitrarily chosen. We proceed by first obtaining a high-probability bound for $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)$, and then converting it to an in-expectation one. To obtain the high-probability bound, we begin by establishing that the difference between the total (probabilistic) loss $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)$, and the indicator (mistake) loss $\mathcal{L}_{\mathbb{I}}(\mathbf{h}, \mathbf{x}, h^\star)$ (see Section 3.1, (3.1)) on the revealed stream of examples $\mathbf{x} = (x_t)_{t=1}^T$ (with each $x_t \sim D_t$) is the sum of a martingale difference sequence.

Let $M_t := \underbrace{\mathcal{L}_P(h_t, D_t, h^\star)}_{P_t} - \underbrace{\mathcal{L}_{\mathbb{I}}(h_t, x_t, h^\star)}_{I_t}$, where $x_t \sim D_t$. With the filtration $\mathbb{F} := (\mathcal{F}_t)_{t=1}^T$, where $\mathcal{F}_t$ corresponds to the information revealed[8] up to (and, including) round t, namely[9] $\mathbf{h}|_t$, $\mathbf{D}|_t$, $\mathbf{x}|_t$ and $\mathbf{y}|_t := (h^\star(x_i))_{i=1}^t$, we note that $\mathbf{M} := (M_t)_{t=1}^T$ is adapted to $\mathbb{F}$ and $\forall t$,

$$\mathbb{E}[M_t] = \mathbb{E}[P_t] - \mathbb{E}[I_t] < \infty$$
$$\mathbb{E}[M_t|\mathcal{F}_{t-1}] = \mathbb{E}[P_t|\mathcal{F}_{t-1}] - \mathbb{E}[I_t|\mathcal{F}_{t-1}] = P_t - P_t = 0,$$

[7] By convention, we treat the "conditioning" on $\mathcal{F}_0$ as an unconditional expectation, i.e., $\mathbb{E}[X_1] = 0$.

[8] Note that this is not alluding to the information revealed to the learner. Instead, we can think of this information as having been revealed to an arbiter until the end of round t, where during each round the arbiter performs the draw $x_t \sim D_t$ on the adversary's communicated choice of D_t and provides $(x_t, h^\star(x_t))$ to the learner.

[9] Recall, $\mathbf{v}|_t = (v_1, \ldots, v_t)$, i.e. $\mathbf{v}$ restricted to the first t rounds.

where the first line is due to the boundedness ($0 \leq P_t, I_t \leq 1,\ \forall t$) of P_t and I_t. And, the second line is due to $\mathbb{E}[I_t|\mathcal{F}_{t-1}] = 1 \cdot \mathbb{P}_{x_t \sim D_t}(h_t(x_t) \neq h^\star(x_t)) + 0 \cdot \mathbb{P}_{x_t \sim D_t}(h_t(x_t) = h^\star(x_t)) = P_t$ and $\mathbb{E}[P_t|\mathcal{F}_{t-1}] = P_t$ (see Remark 1). Therefore, $\mathbf{M}$ is a martingale difference sequence.

We first bound the *predictable quadratic variation* $\langle M_T \rangle$ of $\mathbf{M}$, as follows:

$$\begin{aligned}
\langle M_T \rangle &:= \sum_{t=1}^{T} \mathbb{E}[M_t^2|\mathcal{F}_{t-1}] = \sum_{t=1}^{T} \mathbb{E}[(P_t - I_t)^2|\mathcal{F}_{t-1}] \\
&= \sum_{t=1}^{T} \Big(\mathbb{E}[P_t^2|\mathcal{F}_{t-1}] - 2\mathbb{E}[P_t I_t|\mathcal{F}_{t-1}] + \mathbb{E}[I_t^2|\mathcal{F}_{t-1}] \Big) \\
&= \sum_{t=1}^{T} (P_t^2 - 2P_t^2 + P_t) = \sum_{t=1}^{T} (P_t - P_t^2) \\
&\leq \sum_{t=1}^{T} P_t, \qquad (3.3)
\end{aligned}$$

where the second line is due to linearity of expectation, the third line uses $\mathbb{E}[P_t I_t|\mathcal{F}_{t-1}] = P_t \mathbb{E}[I_t|\mathcal{F}_{t-1}] = P_t \cdot P_t = P_t^2$, and the fourth line is due to $P_t^2 \geq 0,\ \forall t$. Now, by Theorem 1 of [BLL$^+$11], with probability $1 - \delta$ (for any $\delta > 0$), we have

$$\sum_{t=1}^{T} M_t \leq \log\Big(\frac{1}{\delta}\Big) + (e-2)\langle M_T \rangle \iff \sum_{t=1}^{T} P_t \leq \sum_{t=1}^{T} I_t + \log\Big(\frac{1}{\delta}\Big) + (e-2)\langle M_T \rangle.$$

Note here that $\sum_{t=1}^{T} P_t = \mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)$. Therefore, appealing to Theorem 3.1.2 and the inequality on $\langle M_T \rangle$ in (3.3), we have that (with probability $1 - \delta$),

$$\begin{aligned}
&\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) \leq \text{Ldim}(\mathcal{H}) + \log\Big(\frac{1}{\delta}\Big) + (e-2)\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) \\
\implies &\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) \leq \frac{1}{1-(e-2)}\text{Ldim}(\mathcal{H}) + \frac{1}{1-(e-2)} \log\Big(\frac{1}{\delta}\Big),
\end{aligned}$$

or equivalently, for any $\delta > 0$,

$$\mathbb{P}\Big[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \frac{1}{1-(e-2)} \cdot \text{Ldim}(\mathcal{H}) + \frac{\delta}{1-(e-2)}\Big] \leq e^{-\delta}. \qquad (3.4)$$

Now, we compute the in-expectation bound (i.e. a bound on $\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)]$) guaranteed by the

above tail bound (3.4). Let $c := \dfrac{1}{1-(e-2)}$.

$$\begin{aligned}
\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] &= \int_0^\infty \mathbb{P}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \ell]\, d\ell \\
&= \int_0^{c\cdot \text{Ldim}(\mathcal{H})} \mathbb{P}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \ell]\, d\ell + \int_{c\cdot \text{Ldim}(\mathcal{H})}^\infty \mathbb{P}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \ell]\, d\ell \\
&= \int_0^{c\cdot \text{Ldim}(\mathcal{H})} \mathbb{P}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \ell]\, d\ell \\
&\qquad\qquad + \; c\int_0^\infty \mathbb{P}\Big[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > c\cdot \text{Ldim}(\mathcal{H}) + c\delta\Big]\, d\delta \\
&\leq \int_0^{c\cdot \text{Ldim}(\mathcal{H})} 1\, d\ell \; + \; c\int_0^\infty e^{-\delta}\, d\delta \\
&= c\cdot \text{Ldim}(\mathcal{H}) + c \\
&= \mathcal{O}(\text{Ldim}(\mathcal{H})),
\end{aligned}$$

where the first line holds as $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) \geq 0$, the third line uses the change of variable $\ell = c\cdot \text{Ldim}(\mathcal{H}) + c\delta$, the fourth line uses a naïve bound of $\mathbb{P}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \ell] \leq 1$ (for $0 \leq \ell \leq c\cdot \text{Ldim}(\mathcal{H})$) on the first integral, and our high-probability bound in (3.4) on the second integral. ■

Remark 1 (Internal Randomness)**.** *In the proof of Theorem 3.2.6 (and in subsequent proofs throughout the thesis), the filtration $\mathbb{F}$ does not include any information on the internal randomness of the learner or the (adaptive) adversary at any round. In other words, the bound in Theorem 3.2.6 applies to a T-round procedure that is behaving deterministically (in as much as it applies to the learner and the adversary). Crucially, however, it applies to all such procedures. Therefore, the bounds presented in Theorem 3.2.6 (as well as others throughout the thesis) imply identical bounds on a T-round procedure where internal randomness is allowed for both the learner and the adversary.*

To summarize, Theorem 3.2.6 tells us that $\text{Ldim}(\mathcal{H})$ continues to be a sufficient condition for learnability under the new (adversary-provides-a-distribution) model in Section 3.2. In other words, a learner that performs SOA on the observed sequence of examples $\mathbf{x}$ only suffers a constant overhead under the adversary-provides-a-distribution model as compared to SOA under the canonical online model (Section 3.1). Next, we show (Theorem 3.2.7) that $\text{Ldim}(\mathcal{H})$ is also a necessary condition for learnability, and thus fully characterizes the learnability of the adversary-provides-a-distribution model.

Theorem 3.2.7 (Lower bound on the expected loss for the classical adversary-provides-a-distribution model)**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$ be a hypothesis class, and $h^\star \in \mathcal{H}$. For every classical online learner of*

$\mathcal{H}$, there exists an adversary such that

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \Omega(\textit{Ldim}(\mathcal{H})).$$

Proof. Consider an adversary which chooses each D_t to be a point mass on the instance space, i.e. the adversary simply chooses an instance x_t at each t. Since each $x_t \sim D_t$ is deterministic, we have

$$\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) = \mathcal{L}_{\mathbb{I}}(\mathbf{h}, \mathbf{x}, h^\star) = \Omega(\text{Ldim}(\mathcal{H})),$$

where the second equality is due to (the lower bound part of) Theorem 3.1.2. By taking expectations, we conclude our proof. ■

Now that we have characterized the learnability of the adversary-provides-a-distribution model in Section 3.2 (via Theorems 3.2.6 and 3.2.7), we proceed to introduce and present results for related classical online models which arise from successively relaxing, first, the *realizability* assumption and then, the *boolean function class* assumption. These serve to define our scope and lay the groundwork for a comprehensive understanding before introducing the anticipated quantum generalization.

3.2.3 Binary Classification in the Agnostic Setting

In the agnostic framework, we dispense with the realizability assumption that $h^\star \in \mathcal{H}$ (i.e., the labels need not be consistent with any hypothesis in the hypothesis class). In fact, the labels need not arise from a labeling function at all, i.e., the examples may be inconsistent[10]. The agnostic generalization of the adversary-provides-a-distribution model in Section 3.2 is given by the following protocol for the T-round procedure: at the t-th round,

1. Learner provides a hypothesis $h_t \in \mathcal{C}$.
2. Adversary chooses a distribution $D_t : \{0,1\}^{n+1} \to [0,1]$ on the instance space, draws and reveals $z_t = (x_t, y_t) \sim D_t$ to the learner.
3. Learner suffers, but does not "see", a loss of $\mathcal{L}_P(h_t, D_t) := \mathbb{P}_{(x,y)\sim D_t}(h_t(x) \neq y)$.

As there may be no hypothesis that provides the true label on every instance over the T rounds, we resort to comparing the learner to the best fixed hypothesis in $\mathcal{H}$ in hindsight. In other words, the learner's total *regret* is given by,

$$\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) = \sum_{t=1}^{T} \mathcal{L}_P(h_t, D_t) - \inf_{h\in\mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_P(h, D_t). \tag{3.5}$$

[10]That is, it is entirely possible to encounter both $(x, 0)$ and $(x, 1)$ in the sequence of examples.

Definitions 3.2.1, 3.2.2, and 3.2.3 are adapted analogously to give us the notions of a classical online learner, adversary and learnability in the agnostic setting. Additionally, we introduce further notation and definitions to facilitate the subsequent theorem.

An $\mathcal{X}$**-valued tree** $\underline{X}$ **of depth** T is a rooted complete binary tree with nodes labeled by elements of $\mathcal{X}$. We identify the tree $\underline{X}$ with the sequence $x^{(1)}, \ldots, x^{(T)}$ of labeling functions $x^{(i)} : \{\pm 1\}^{i-1} \to \mathcal{X}$ which provide the labels for each node. Here, $x^{(1)}$ labels the root of the tree, and $x^{(i)}$ for $i > 1$ labels the node obtained by following the path of length $i-1$ from the root, with $+1$ indicating 'right' and -1 indicating 'left'. A path of length T is given by the sequence $\epsilon = (\epsilon_1, \ldots, \epsilon_T) \in \{\pm 1\}^T$. We denote the label at round t along this path as $x^{(t)}(\epsilon)$, understanding that $x^{(t)}$ depends only on the prefix $(\epsilon_1, \ldots, \epsilon_{t-1})$ of ϵ. With this notion of a tree, we define the *sequential Rademacher complexity* of a hypothesis class, $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ [RST15].

Definition 3.2.8 (Sequential Rademacher complexity of $\mathcal{H}$)**.** *The sequential Rademacher complexity of a function class $\mathcal{H} \subseteq \boldsymbol{R}^{\mathcal{X}}$ on an $\mathcal{X}$-valued tree $\underline{X}$ is defined as*

$$\mathfrak{R}_T(\mathcal{H}, \underline{X}) = \mathbb{E}\Big[\sup_{h \in \mathcal{H}} \frac{1}{T} \sum_{t=1}^{T} \epsilon_t h(x^{(t)}(\epsilon))\Big],$$

and

$$\mathfrak{R}_T(\mathcal{H}) = \sup_{\underline{X}} \mathfrak{R}_T(\mathcal{H}, \underline{X}),$$

where the outer supremum is taken over all $\mathcal{X}$-valued trees of depth T; $\epsilon = (\epsilon_1, \ldots, \epsilon_T) \in \{\pm 1\}^T$ is a sequence of i.i.d. Rademacher random variables.

Definition 3.2.9 (The loss class $\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}$)**.** *The loss class, $\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}$, is a boolean hypothesis class given by*

$$\mathcal{L}_{\mathbb{I}} \circ \mathcal{H} = \{l_h : (x, y) \mapsto \mathbf{1}[h(x) \neq y] \,|\, h \in \mathcal{H}\}.$$

Note that $\mathcal{L}_{\mathbb{I}} \circ \mathcal{H} \subseteq \{0,1\}^{\mathcal{X} \times \mathcal{Y}}$, and that its sequential Rademacher complexity is defined analogously to Definition 3.2.8. With these definitions in place, we can proceed to present the theorems for the bounds on expected regret for the adversary-provides-a-distribution model in the agnostic setting.

Theorem 3.2.10 (Upper bound on the expected regret for the classical adversary-provides-a-distribution model in the agnostic setting)**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$ be a hypothesis class. For every adversary, there exists a classical online learner for $\mathcal{H}$ that satisfies*

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \mathcal{O}(\sqrt{\mathit{Ldim}(\mathcal{H}) \cdot T}).$$

Proof. Let $\mathbf{D}$ be an arbitrary sequence of distributions, and let $\mathbf{z} = (z_1, \ldots, z_t)$ be a sequence of instances such that $z_i \sim D_i$. Then, defining[11]

$$\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) := \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, z_t) - \inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h, z_t),$$

we have

$$\begin{aligned}\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) &= \mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) - \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) \\ &= \underbrace{\sum_{t=1}^{T} \mathcal{L}_P(h_t, D_t) - \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, z_t)}_{\Delta_1} + \underbrace{\inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h, z_t) - \inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_P(h, D_t)}_{\Delta_2} + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}).\end{aligned}$$

We proceed by bounding the expected value of Δ_1, Δ_2 and $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})$ separately.

Working with Δ_1, we let $M_t := \mathcal{L}_P(h_t, D_t) - \mathcal{L}_{\mathbb{I}}(h_t, z_t)$. With the filtration $\mathbb{F} := (\mathcal{F}_t)_{t=1}^{T}$, where $\mathcal{F}_t$ corresponds to the information revealed up to (and, including) round t, namely $\mathbf{h}|_t$, $\mathbf{D}|_t$ and $\mathbf{z}|_t$, we note that $\mathbf{M} := (M_t)_{t=1}^{T}$ is adapted to $\mathbb{F}$ and $\forall t$,

$$\begin{aligned}\mathbb{E}[M_t] &= \mathbb{E}[\mathcal{L}_P(h_t, D_t)] - \mathbb{E}[\mathcal{L}_{\mathbb{I}}(h_t, z_t)] < \infty \\ \mathbb{E}[M_t | \mathcal{F}_{t-1}] &= \mathbb{E}[\mathcal{L}_P(h_t, D_t) | \mathcal{F}_{t-1}] - \mathbb{E}[\mathcal{L}_{\mathbb{I}}(h_t, z_t) | \mathcal{F}_{t-1}] = \mathcal{L}_P(h_t, D_t) - \mathcal{L}_P(h_t, D_t) = 0,\end{aligned}$$

where the first line is due to the boundedness ($0 \leq \mathcal{L}_P(h_t, D_t), \mathcal{L}_{\mathbb{I}}(h_t, z_t) \leq 1, \ \forall t$) of $\mathcal{L}_P(h_t, D_t)$ and $\mathcal{L}_{\mathbb{I}}(h_t, z_t)$. And, the second line is due to $\mathbb{E}[\mathcal{L}_{\mathbb{I}}(h_t, z_t) | \mathcal{F}_{t-1}] = 1 \cdot \mathbb{P}_{z_t \sim D_t}(h_t(x_t) \neq y_t) + 0 \cdot \mathbb{P}_{z_t \sim D_t}(h_t(x_t) = y_t) = \mathcal{L}_P(h_t, D_t)$. Therefore, $\mathbf{M}$ is a martingale difference sequence. Now, as $\Delta_1 = \sum_{t=1}^{T} M_t$, by Azuma-Hoeffding's inequality, since $|M_t| < 1$ for all t, we have that

$$\mathbb{P}[|\Delta_1| \geq \delta] \leq 2 \exp\Big(- \frac{\delta^2}{2T} \Big), \quad \text{for all} \ \ \delta \in \mathbf{R}^+, T \in \mathbb{Z}^+. \tag{3.6}$$

This allows us to compute the following bound on $\mathbb{E}[\Delta_1]$:

$$\begin{aligned}\mathbb{E}[\Delta_1] \leq \mathbb{E}[|\Delta_1|] &\leq \int_0^\infty \mathbb{P}[|\Delta_1| \geq \delta] \, dB \\ &\leq \int_0^\infty 2 \exp\Big(- \frac{\delta^2}{2T} \Big) \, dB = \sqrt{2\pi T} = \mathcal{O}(\sqrt{T}),\end{aligned}$$

where the first line holds as $\Delta_1 < |\Delta_1|$ (for the first inequality) and $|\Delta_1| \geq 0$ (for the second inequality), and the second line uses the bound in (3.6).

[11] $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})$ is precisely the regret of an algorithm in the agnostic generalization of the canonical (adversary-provides-an-input) classical online model in Section 3.1.

Next, working with Δ_2, we obtain the following chain of (in)equalities:

$$\begin{aligned}
\Delta_2 &= \inf_{h\in\mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h, z_t) - \inf_{h\in\mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_P(h, D_t) \\
&= -\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} -\mathcal{L}_{\mathbb{I}}(h, z_t) + \sup_{h\in\mathcal{H}} \sum_{t=1}^{T} -\mathcal{L}_P(h, D_t) \\
&= -\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} (-\mathcal{L}_{\mathbb{I}}(h, z_t) + \mathcal{L}_P(h, D_t) - \mathcal{L}_P(h, D_t)) + \sup_{h\in\mathcal{H}} \sum_{t=1}^{T} -\mathcal{L}_P(h, D_t) \\
&\leq \sup_{h\in\mathcal{H}} \sum_{t=1}^{T} (\mathcal{L}_{\mathbb{I}}(h, z_t) - \mathcal{L}_P(h, D_t)),
\end{aligned}$$

where the final inequality is a result of the subadditivity of the supremum; to elaborate, we derive $\sup_h f \leq \sup_h (f-g) + \sup_h g \iff \sup_h f - \sup_h g \leq \sup_h (f-g)$, to which we substitute $f = -\mathcal{L}_P(h, D_t)$ and $g = -\mathcal{L}_{\mathbb{I}}(h, z_t) + \mathcal{L}_P(h, D_t) - \mathcal{L}_P(h, D_t)$. So far, taking expectations both sides, we have $\mathbb{E}[\Delta_2] \leq \mathbb{E}[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} (\mathcal{L}_{\mathbb{I}}(h, z_t) - \mathcal{L}_P(h, D_t))]$. Next, we bound the in-expectation quantity on the right-hand side to obtain,

$$\begin{aligned}
\mathbb{E}\Big[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} (\mathcal{L}_{\mathbb{I}}(h, z_t) - \mathcal{L}_P(h, D_t))\Big] &= \mathbb{E}\left[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} \Big(\mathbf{1}[h(x_t) \neq y_t] - \mathbb{P}_{(x_t,y_t)\sim D_t}[h(x_t) \neq y_t]\Big)\right] \\
&= \mathbb{E}\left[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} \Big(\mathbf{1}[h(x_t) \neq y_t] - \mathbb{E}_{(x_t,y_t)\sim D_t}\Big[\mathbf{1}[h(x_t) \neq y_t]|\mathcal{F}_{t-1}\Big]\Big)\right] \\
&\leq 2T \cdot \mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}) \\
&\leq 2T \cdot \mathfrak{R}_T(\mathcal{H}) \\
&= \mathcal{O}(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T}),
\end{aligned}$$

where the second line uses a conditioning on the filtration $\mathcal{F}_{t-1}$, which corresponds to the information revealed up to (and, including) round $t-1$ (namely $\mathbf{h}|_{t-1} := (h_1, \ldots, h_{t-1})$, $\mathbf{D}|_{t-1} := (D_1, \ldots, D_{t-1})$ and $\mathbf{z}|_{t-1} := ((x_1, y_1), \ldots, (x_{t-1}, y_{t-1}))$), the third line is due to Theorem 2 in [RST15][12], the fourth line is due to Theorem 16 of [DT22] (with $\mathcal{Y} = \{0, 1\}$) and the last line is from the proof of Theorem 12.1 in [ABED+21].

[12]Theorem 2 in [RST15], as stated, provides a bound on $\mathbb{E}\Big[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} \Big(\mathbb{E}_{(x_t,y_t)\sim D_t}\Big[\mathbf{1}[h(x_t) \neq y_t]|\mathcal{F}_{t-1}\Big] - \mathbf{1}[h(x_t) \neq y_t]\Big)\Big]$. However, the proof of Lemma 9 in [RST15] notes its validity even with absolute values around the sum, which subsequently ensures that Theorem 2 in [RST15] also holds in the same generality. This justifies its use here in the following sense: $\mathbb{E}\Big[\sup_{h\in\mathcal{H}} \sum_{t=1}^{T} \Big(\mathbf{1}[h(x_t) \neq y_t] - \mathbb{E}_{(x_t,y_t)\sim D_t}\Big[\mathbf{1}[h(x_t) \neq y_t]|\mathcal{F}_{t-1}\Big]\Big)\Big] \leq \mathbb{E}\Big[\sup_{h\in\mathcal{H}} \Big|\sum_{t=1}^{T} \Big(\mathbb{E}_{(x_t,y_t)\sim D_t}\Big[\mathbf{1}[h(x_t) \neq y_t]|\mathcal{F}_{t-1}\Big] - \mathbf{1}[h(x_t) \neq y_t]\Big)\Big|\Big] \leq 2T \cdot \mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H})$.

Finally, also from Theorem 12.1 of [ABED$^+$21], we have $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) = \mathcal{O}(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T})$. Putting everything together, using the linearity of expectation, we have,

$$\begin{aligned}
\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] &= \mathbb{E}[\Delta_1 + \Delta_2 + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] \\
&= \mathbb{E}[\Delta_1] + \mathbb{E}[\Delta_2] + \mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] \\
&= \mathcal{O}(\sqrt{T}) + \mathcal{O}(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T}) + \mathcal{O}(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T}) \\
&= \mathcal{O}(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T}).
\end{aligned}$$

■

In summary, our theorem reveals that the optimal classical learner for the canonical (adversary-provides-an-input) agnostic model, when provided with the observed sequence of instances $\mathbf{z}$ in the new protocol, experiences, at most, a constant overhead when assessed under the new (adversary-provides-a-distribution) framework. Upon closer examination, it was critical for the bound on $\mathbb{E}[\Delta_2]$ in our proof not to be worse than the bound on $\mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})]$. It is noteworthy that the bound on $\mathbb{E}[\Delta_2]$ is guaranteed by the rate of online uniform convergence (in the frameworks of the sequential Rademacher complexity [RST15] and the adversarial (uniform) laws of large numbers [ABED$^+$21]), whereas the bound on $\mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})]$ is guaranteed by the rate of canonical (agnostic) online learnability. The equivalence between these two rates, for the boolean function class case, played a pivotal role in establishing our result.

Next, we establish (Theorem 3.2.11) a matching lower bound for expected regret within the agnostic adversary-provides-a-distribution framework. This fully characterizes agnostic learnability under the adversary-provides-a-distribution framework. As with all lower bound proofs within this framework, we efficiently conclude the proof statement by considering an adversary that exclusively plays point masses.

Theorem 3.2.11 (Lower bound on the expected regret for the classical adversary-provides-a-distribution model in the agnostic setting)**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$ be a hypothesis class. For every classical online learner of $\mathcal{H}$, there exists an adversary such that*

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\textit{Ldim}(\mathcal{H}) \cdot T}).$$

Proof. As in Theorem 3.2.7, we again consider an adversary which chooses each D_t to be a point mass on the instance space, i.e. the adversary simply chooses an instance z_t at each t. Since each $z_t \sim D_t$ is deterministic, we have

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] = \Omega(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T}),$$

where the second equality is due to (the lower bound part of) Theorem 21.10 in [SSBD14]. ■

3.2.4 Multiclass Classification in the Realizable Setting

In Sections 3.1, 3.2.2 and 3.2.3, we concerned ourselves with learning Boolean hypothesis classes, i.e. $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$. Now, we consider the adversary-provides-a-distribution model (introduced in Section 3.2) extended to the setting of multiclass learning, i.e. $\mathcal{H} \subseteq \mathcal{C} := \{c : \mathcal{X} \to \mathcal{Y}\}$, with $|\mathcal{Y}| > 2$. As stated earlier, the objective here is to lay the groundwork, with a clear understanding of these models in the classical paradigm, before delving into their quantum generalizations.

The online learning protocol in the realizable setting is identical to that specified in Section 3.2, with the added specification of a multiclass hypothesis (and, concept) class. To express our results in this setting, we first define the combinatorial parameter, multiclass Littlestone dimension (mcLdim($\mathcal{H}$)), which is a generalization of the Littlestone dimension to the multiclass setting.

Definition 3.2.12 (Multiclass Littlestone dimension)**.** *Let T be a rooted tree whose internal nodes are labeled by elements from $\mathcal{X}$ and whose edges are labeled by elements from $\mathcal{Y}$, such that the edges from a single parent to its child nodes are each labeled with a different label*[13]*. The tree T is mcL-shattered by $\mathcal{H}$ if, for every path from root to leaf which traverses the nodes $x_1, \ldots, x_d$, there exists a hypothesis $h \in \mathcal{H}$ such that, for all i, $h(x_i)$ is the label of the edge (x_i, x_{i+1}). We define the multiclass Littlestone dimension, mcLdim($\mathcal{H}$), to be the maximal depth of a complete **binary** tree that is mcL-shattered by $\mathcal{H}$.*

Theorem 3.2.13 (Upper bound on the expected loss for the classical adversary-provides-a-distribution model in the multiclass setting)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every adversary, there exists a classical online learner for $\mathcal{H}$ that satisfies*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathcal{O}(\textit{mcLdim}(\mathcal{H})).$$

Proof. We follow the steps in the proof of Theorem 3.2.6, which continue to hold in the multiclass setting, with a minor difference: the upper bound on $\sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, x_t, h^\star)$ is now mcLdim($\mathcal{H}$) instead of Ldim($\mathcal{H}$). As a result, we arrive at the following high-probability bound: with probability $1-\delta$ (for any $\delta > 0$), $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) \leq \frac{1}{1-(e-2)}\text{mcLdim}(\mathcal{H}) + \frac{1}{1-(e-2)}\log(\frac{1}{\delta})$, or equivalently, the following tail bound: for any $\delta > 0$,

$$\mathbb{P}\Big[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) > \frac{1}{1-(e-2)} \cdot \text{mcLdim}(\mathcal{H}) + \frac{\delta}{1-(e-2)}\Big] \leq e^{-\delta},$$

from which we obtain the desired in-expectation result: $\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathcal{O}(\text{mcLdim}(\mathcal{H}))$. ■

[13] In the binary case (where the only "different labels" are 0 and 1), it is not hard to see that the definition reduces to that of the Littlestone dimension (Definition 3.1.1).

Theorem 3.2.14 (Lower bound on the expected loss for the classical adversary-provides-a-distribution model in the multiclass setting)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every classical online learner of $\mathcal{H}$, there exists an adversary such that*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \Omega(\textit{mcLdim}(\mathcal{H})).$$

Proof. As in the online lower bound proofs thus far, we consider an adversary which chooses each D_t to be a point mass on the instance space, i.e. the adversary simply chooses an instance x_t at each t. Since each $x_t \sim D_t$ is deterministic, we have

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathbb{E}[\mathcal{L}_{\mathbb{I}}(\mathbf{h}, \mathbf{x}, h^\star)] = \Omega(\text{mcLdim}(\mathcal{H})),$$

where the second equality is due to (the lower bound part of) Theorem 24 in [DSBDSS15]. ■

Theorems 3.2.13 and 3.2.14 together imply that $\text{mcLdim}(\mathcal{H})$ continues to characterize multiclass learnability in the realizable setting under the adversary-provides-a-distribution framework. Given that the rate $\Theta(\text{mcLdim}(\mathcal{H}))$ is independent of k, $\text{mcLdim}(\mathcal{H})$ characterizes realizable learnability even in the unbounded label space case, mirroring the scenario in the adversary-provides-an-input framework (Theorem 24, [DSBDSS15]).

3.2.5 Multiclass Classification in the Agnostic Setting

Here, the online learning protocol is identical to that specified in Section 3.2.3, with the added specification of a multiclass hypothesis (and, concept) class. We introduce specific definitions and lemmas to facilitate the subsequent theorem, which provides an upper bound on the expected regret of an online learner in this multiclass agnostic adversary-provides-a-distribution setting. For notation related to trees, please refer to Section 3.2.3 (paragraph preceding Definition 3.2.8).

Definition 3.2.15 (0-cover of a hypothesis class $\mathcal{H}$ on a tree $\underline{X}$)**.** *A set $\mathcal{V}$ of $\mathcal{Y}$-valued trees is a 0-cover of $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ on an $\mathcal{X}$-valued tree $\underline{X}$ of depth T if*

$$\forall h \in \mathcal{H}, \forall \epsilon \in \{\pm 1\}^T, \exists \underline{V} \in \mathcal{V},\ \textit{s.t.},\ v^{(t)}(\epsilon) = h(x^{(t)}(\epsilon)),$$

for all $t \in [T]$.

Definition 3.2.16 (Covering number of a hypothesis class $\mathcal{H}$ on a tree $\underline{X}$)**.** *The covering number of a hypothesis class $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ on an $\mathcal{X}$-valued tree $\underline{X}$, $\mathcal{N}(\mathcal{H}, \underline{X})$, is defined as follows:*

$$\mathcal{N}(\mathcal{H}, \underline{X}) := \min\{|\mathcal{V}| : \mathcal{V} \textit{ is a 0-cover of } \mathcal{H} \textit{ on } \underline{X}\},$$

i.e. the size of the smallest set $\mathcal{V}$ (of trees) that 0-covers $\underline{X}$.

Lemma 3.2.17 ($\mathcal{N}(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}, \underline{Z}) \leq \mathcal{N}(\mathcal{H}, \underline{X})$)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$. Let $\underline{Z}$ be a $(\mathcal{X} \times \mathcal{Y})$-valued tree,*[14] *and let $\underline{X}$ be the $\mathcal{X}$-valued tree obtained by extracting the $\mathcal{X}$-component from each node of $\underline{Z}$. Then,*

$$\mathcal{N}(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}, \underline{Z}) \leq \mathcal{N}(\mathcal{H}, \underline{X}).$$

Proof. Let the set $\mathcal{V}$ be the smallest set of trees that form a 0-cover of $\mathcal{H}$ on $\underline{X}$. For each tree $\underline{V} \in \mathcal{V}$, we construct (and add to $\mathcal{W}$) a tree $\underline{W}$, given by

$$w^{(t)}(\epsilon) = \begin{cases} 0, & \text{if } \ v^{(t)}(\epsilon) = y^{(t)}(\epsilon) \\ 1, & \text{otherwise} \end{cases},$$

where $y^{(t)}(\epsilon)$ is the y-label of the node of $\underline{Z}$ encountered after having traversed the path $(\epsilon_1, \ldots, \epsilon_{t-1})$. We claim that $\mathcal{W}$ provides a 0-cover of $\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}$ on $\underline{Z}$.

As the set $\mathcal{V}$ of trees forms a 0-cover of $\mathcal{H}$ on $\underline{X}$, we know that

$$\forall h \in \mathcal{H}, \forall \epsilon \in \{\pm 1\}^T, \exists \underline{V} \in \mathcal{V}, \text{ s.t., } v^{(t)}(\epsilon) = h(x^{(t)}(\epsilon)),$$

for all $t \in [T]$. Therefore, by construction,

$$\forall h \in \mathcal{H}, \forall \epsilon \in \{\pm 1\}^T, \exists \underline{W} \in \mathcal{W}, \text{ s.t., } w^{(t)}(\epsilon) = \mathbf{1}[h(x^{(t)}(\epsilon)) \neq y^{(t)}(\epsilon)],$$

for all $t \in [T]$. So, we obtain

$$\mathcal{N}(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}, \underline{Z}) \leq |\mathcal{W}| = |\mathcal{V}| = \mathcal{N}(\mathcal{H}, \underline{X}),$$

completing our proof. ■

Lemma 3.2.18 ($\mathcal{N}(\mathcal{H}, \underline{X}) \leq (Tk)^{\text{mcLdim}(\mathcal{H})}$)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ with $|\mathcal{Y}| = k > 2$, and $\underline{X}$ be an $\mathcal{X}$-valued tree. Then,*

$$\mathcal{N}(\mathcal{H}, \underline{X}) \leq (Tk)^{\textit{mcLdim}(\mathcal{H})}.$$

Proof. Our main idea is that the set of experts, as defined in the proof of Theorem 25 in [DSBDSS15] (reproduced below), can be used to construct a 0-cover of $\mathcal{H}$ on $\underline{X}$.

> Given time horizon T, let $A_T = \{A \subset [T] \mid |A| \leq \text{mcLdim}(\mathcal{H})\}$. For every $A \in A_T$ and $\phi : A \to \mathcal{Y}$, we define an expert $E_{A,\phi}$. The expert $E_{A,\phi}$ imitates the SOA algorithm

[14] The nodes of $\underline{Z}$ are of the form $z = (x, y)$, where $x \in \mathcal{X}$ and $y \in \mathcal{Y}$.

when it errs exactly on the examples $\{x_t \mid t \in A\}$ and the true labels of these examples are determined by ϕ. Formally, the expert $E_{A,\phi}$ proceeds as follows:

> Set $V_1 = \mathcal{H}$.
> For $t = 1, 2, \ldots, T$:
> Receive x_t.
> If $t \in A$, set $\hat{y}_t = \phi(t)$.
> If $t \notin A$, set $\hat{y}_t = \operatorname{argmax}_{y \in \mathcal{Y}} \text{mcLdim}(\{h \in V_t : h(x_t) = y\})$.
> Predict $\hat{y}_t$ and update $V_{t+1} = \{h \in V_t : h(x_t) = \hat{y}_t\}$.

Throughout this proof, we use the notation $E_{A,\phi}(x_1, \ldots, x_t)$ to denote the label $\hat{y}_t$ predicted at round t by the expert $E_{A,\phi}$ after processing the sequence of instances $(x_1, \ldots, x_t)$. Similarly, given any true labeling function $h^\star \in \mathcal{H}$, we use the notation $\text{SOA}^{h^\star}(x_1, \ldots, x_t)$ to denote the label predicted at round t by the *standard* SOA algorithm, after processing the sequence of instances $(x_1, \ldots, x_t)$ and updating its version space based on the labels $h^\star(x_1), \ldots, h^\star(x_{t-1})$.

Now, the set of experts $\mathcal{E} = \{E_{A,\phi}\}$ has size $|\mathcal{E}| = \sum_{j=0}^{\text{mcLdim}(\mathcal{H})} \binom{T}{j} k^j \leq (Tk)^{\text{mcLdim}(\mathcal{H})}$ with the following *expert guarantee*:

> For any sequence $(x_1, \ldots, x_T)$ of instances, and any $h \in \mathcal{H}$, there exists an expert $E^\star \in \mathcal{E}$ such that $E^\star(x_1, \ldots, x_t) = h(x_t)$, $\forall\, t \in [T]$.

Indeed, given an arbitrary sequence of instances $(x_1, \ldots, x_T)$ and an arbitrary $h \in \mathcal{H}$, the expert $E^\star = E_{A^\star,\phi^\star}$ that satisfies the guarantee corresponds to $A^\star := \{t \in [T] : \text{SOA}^h(x_1, \ldots, x_t) \neq h(x_t)\}$ and $\phi^\star : A^\star \to \mathcal{Y}$ defined by $\phi^\star(t) := h(x_t)$. Since the SOA algorithm makes at most $\text{mcLdim}(\mathcal{H})$ mistakes, we have $|A^\star| \leq \text{mcLdim}(\mathcal{H})$, and therefore $E^\star \in \mathcal{E}$ by construction.

Next, for each expert $E \in \mathcal{E}$, we add a tree $\underline{V_E}$ to $\mathcal{V}$, given by

$$\underline{V_E}^{(t)}(\epsilon) = E(x^{(1)}(\epsilon), \ldots, x^{(t)}(\epsilon)), \quad \forall\, t \in [T],$$

where $\epsilon = (\epsilon_1, \ldots, \epsilon_T) \in \{\pm 1\}^T$ is a sequence of i.i.d. Rademacher random variables. We now verify that the set of trees $\mathcal{V}$ forms a 0-cover of $\mathcal{H}$ on $\underline{X}$.

Fix an arbitrary $h \in \mathcal{H}$ and an arbitrary $\epsilon = (\epsilon_1, \ldots, \epsilon_T) \in \{\pm 1\}^T$. For the sequence of examples on the path $(x^{(1)}(\epsilon), \ldots, x^{(T)}(\epsilon))$, by the expert guarantee and our construction of $\mathcal{V}$ above, we have that there exists $E^\star \in \mathcal{E}$ such that $\underline{V_{E^\star}}^{(t)}(\epsilon) = E^\star(x^{(1)}(\epsilon), \ldots, x^{(t)}(\epsilon)) = h(x^{(t)}(\epsilon))$, $\forall t \in [T]$. Therefore, by Definition 3.2.15, we see that $\mathcal{V}$ forms a 0-cover of $\mathcal{H}$ on $\underline{X}$. Hence,

$$\mathcal{N}(\mathcal{H}, \underline{X}) = \min\{|\mathcal{V}| : \mathcal{V} \text{ is a 0-cover of } \mathcal{H} \text{ on } \underline{X}\} \leq |\mathcal{V}| = |\mathcal{E}| \leq (Tk)^{\text{mcLdim}(\mathcal{H})},$$

as desired. ■

Theorem 3.2.19 (Upper bound on the expected regret for the classical adversary-provides-a-distribution model in the multiclass agnostic setting)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every adversary, there exists a classical online learner for $\mathcal{H}$ that satisfies*

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}))] = \mathcal{O}(\sqrt{\textit{mcLdim}(\mathcal{H}) \cdot T \log(Tk)}).$$

Proof. We follow the steps in the proof of Theorem 3.2.10, which, in a general sense, are applicable in the multiclass setting. However, a key bound used in the proof of Theorem 3.2.10 ($\mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}) \leq \mathfrak{R}_T(\mathcal{H})$) is not known to continue to hold in the multiclass setting, forcing us to handle the rest of the proof differently. Some interesting insights follow from this deviation, which is elaborated in the proof below, as well as the discussion that follows.

We recall the preliminaries. Let $\mathbf{D}$ be an arbitrary sequence of distributions, and let $\mathbf{z} = (z_1, \ldots, z_t)$ be a sequence of instances such that $z_i \sim D_i$. Then, defining[15] $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) = \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, z_t) - \inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h, z_t)$, we have

$$\begin{aligned}
\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) &= \mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) - \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) \\
&= \underbrace{\sum_{t=1}^{T} \mathcal{L}_P(h_t, D_t) - \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h_t, z_t)}_{\Delta_1} + \underbrace{\inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_{\mathbb{I}}(h, z_t) - \inf_{h \in \mathcal{H}} \sum_{t=1}^{T} \mathcal{L}_P(h, D_t)}_{\Delta_2} + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}).
\end{aligned}$$

We proceed by bounding the expected value of Δ_1, Δ_2 and $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})$ separately. First, our bound $\mathbb{E}[\Delta_1] \leq \mathcal{O}(\sqrt{T})$ using Azuma-Hoeffding's inequality in Theorem 3.2.10 is independent of the form of $\mathcal{H}$ and, in particular, continues to hold for a multiclass $\mathcal{H}$.

Next, with regard to $\mathbb{E}[\Delta_2]$, we recover the chain of inequalities in Theorem 3.2.10 leading up to

$$\mathbb{E}[\Delta_2] \leq 2T \cdot \mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}). \tag{3.7}$$

However, the result in Theorem 16 of [DT22] ($\mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}) \leq \mathfrak{R}_T(\mathcal{H})$) only applies when $\mathcal{H}$ is a boolean hypothesis class. Therefore, we proceed with an explicit bound on $\mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H})$ using a covering number argument. Let $\underline{Z}$ be an $(\mathcal{X} \times \mathcal{Y})$-valued tree, and let $\underline{X}$ be the $\mathcal{X}$-valued tree obtained by extracting the $\mathcal{X}$-component from each node of $\underline{Z}$. We provide a chain of inequalities

[15] $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})$ is the regret of an algorithm in the *multiclass agnostic* generalization of the canonical (adversary-provides-an-input) classical online model in Section 3.1.

starting from (3.7):

$$\begin{aligned}
\mathbb{E}[\Delta_2] &\leq 2T \cdot \mathfrak{R}_T(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}) \\
&\leq \frac{24T\sqrt{\log \mathcal{N}(\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}, \underline{Z})}}{\sqrt{T}} \\
&\leq \frac{24T\sqrt{\log \mathcal{N}(\mathcal{H}, \underline{X})}}{\sqrt{T}} \\
&\leq \frac{24T\sqrt{\text{mcLdim}(\mathcal{H}) \cdot \log(Tk)}}{\sqrt{T}} \\
&= \mathcal{O}(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log(Tk)}),
\end{aligned}$$

where the second line is from Theorem 3 and Definition 5 of [RST15], the third line is from Lemma 3.2.17, and the fourth line is from Lemma 3.2.18. Finally, from Theorem 4 of [HMR$^+$23], we have $\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H}) = \mathcal{O}\Big(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log\left(\frac{T}{\text{mcLdim}(\mathcal{H})}\right)}\Big)$, when $T \geq 2 \cdot \text{mcLdim}(\mathcal{H})$. Putting everything together, using the linearity of expectation, we have,

$$\begin{aligned}
\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] &= \mathbb{E}[\Delta_1 + \Delta_2 + \mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] \\
&= \mathbb{E}[\Delta_1] + \mathbb{E}[\Delta_2] + \mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] \\
&= \mathcal{O}(\sqrt{T}) + \mathcal{O}(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log(Tk)}) \\
&\qquad + \mathcal{O}\left(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log\left(\frac{T}{\text{mcLdim}(\mathcal{H})}\right)}\right) \\
&= \mathcal{O}(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log(Tk)}),
\end{aligned}$$

where the third line holds as point-wise bounds imply bounds in-expectation[16], and the last line holds as $k > 2$ and $\text{mcLdim}(\mathcal{H}) \geq 1 \implies k \geq \frac{1}{\text{mcLdim}(\mathcal{H})}$. ■

In summary, our theorem provides a k-dependent upper bound on $\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})]$, while the corresponding lower bound (presented next, Theorem 3.2.20), is k-independent. Although the optimal classical learner for the canonical multiclass agnostic model bridges this gap, as indicated by the k-independent bound used on $\mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})]$, it remains unclear whether we can establish a k-*independent* upper bound on $\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})]$. Our current proof strategy[17] faces challenges in achieving this goal due to the following observation. The bound on $\mathbb{E}[\Delta_2]$ is guaranteed by the

[16]Since the point-wise bound holds for $T \geq 2 \cdot \text{mcLdim}(\mathcal{H})$, it is clear that $\mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] \leq 2 \cdot \text{mcLdim}(\mathcal{H}) + \mathcal{O}\Big(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T \log\left(\frac{T}{\text{mcLdim}(\mathcal{H})}\right)}\Big)$, where the latter term dominates when $T \geq 4 \cdot \text{mcLdim}(\mathcal{H})$.

[17]Currently, we employ the optimal classical learner for the canonical multiclass agnostic model by presenting it with the observed sequence of instances $\mathbf{z}$ in the adversary-provides-a-distribution protocol and evaluate its performance under the adversary-provides-a-distribution framework.

rate of online uniform convergence of the loss ($\mathcal{L}_{\mathbb{I}} \circ \mathcal{H}$) class (in the frameworks of sequential Rademacher complexity [RST15] and adversarial (uniform) laws of large numbers [ABED$^+$21]). Meanwhile, the bound on $\mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})]$ is guaranteed by the rate of canonical (agnostic) online learnability of $\mathcal{H}$. However, in the multiclass function class case, as demonstrated by Theorem 7 and Example 1 in [HMR$^+$23], these two rates are *not* equivalent.

Theorem 3.2.20 (Lower bound on the expected regret for the classical adversary-provides-a-distribution model in the multiclass agnostic setting)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class. For every classical online learner of $\mathcal{H}$, there exists an adversary such that*

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\textit{mcLdim}(\mathcal{H}) \cdot T}).$$

Proof. As in the online lower bound proofs thus far, we consider an adversary which chooses each D_t to be a point mass on the instance space, i.e., the adversary simply chooses an instance z_t at each t. Since each $z_t \sim D_t$ is deterministic, we have

$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \mathbb{E}[\mathcal{R}_{\mathbb{I}}(\mathbf{h}, \mathbf{z}, \mathcal{H})] = \Omega(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T}),$$

where the second equality is due to Theorem 26 in [DSBDSS15]. ■

Due to Theorems 3.2.19 and 3.2.20, we have characterized multiclass agnostic learnability in the adversary-provides-a-distribution setting for the *bounded* label space ($k < \infty$) case.

Related Work. The online *adversary-provides-a-distribution* model presented in this section (Section 3.2) has analogues in other works. For example, [HRS20] and [BDGR22] consider a setting where an adaptive adversary provides a distribution. Their setup is designed for smoothed learning, where at each time step, the adversary provides a sample drawn from their chosen distribution. However, their framework assumes a full-information setting, where the learner observes all losses, whereas in our model, the learner only receives partial feedback. Another related connection can be found in the online game introduced in [LN24] to analyze batch generalization error. While their setting, like ours, involves partial information on losses (and, therefore, regret), the probabilistic aspect of their model arises from the learner selecting a distribution over a hypothesis class, rather than the adversary providing one.

3.3 Online Learning with Quantum Examples

Equipped with our model in Section 3.2, we are finally ready to introduce our quantum online learning model. However, prior to the model description, we clarify our scope. In its nascent

existence, quantum online learning has primarily focused on the online learning of *quantum states* [ACH+19, QAS21, AA24]. In contrast, our focus in this thesis is on the online learning of *classical functions* via quantum examples. Our scope is motivated by the abundance of classical online learning literature [Lit88, BDPSS09, DSBDSS15, SSBD14], as well as our results in Sections 3.2, that presents us with an at-the-ready comparison.

3.3.1 Model Description

Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ be a hypothesis class. Identifying the T-round protocol in Section 3.2 (corresp. Section 3.2.3) with the definition of a quantum example in (2.1) (corresp. (2.2)), we obtain the following "natural" model for quantum online learning. The T-round protocol proceeds as follows: at the t-th round,

1. Learner provides a hypothesis $h_t : \mathcal{X} \to \mathcal{Y}$.
2. Adversary reveals an example $|\psi_t\rangle$ where
   (a) $|\psi_t\rangle = \sum_{x\in\mathcal{X}} \sqrt{D_t(x)}\,|x, h^\star(x)\rangle$ for some $D_t : \mathcal{X} \to [0,1]$ and $h^\star \in \mathcal{H}$ (realizable),
   (b) $|\psi_t\rangle = \sum_{x\in\mathcal{X},\, y\in\mathcal{Y}} \sqrt{D_t(x,y)}\,|x, y\rangle$ for some $D_t : \mathcal{X}\times\mathcal{Y} \to [0,1]$ (agnostic[18]).
3. Learner incurs loss[19]
   (a) $\mathcal{L}_P(h_t, D_t, h^\star) := \mathbb{P}_{x\sim D_t}(h_t(x) \neq h^\star(x))$ (realizable),
   (b) $\mathcal{L}_P(h_t, D_t) := \mathbb{P}_{(x,y)\sim D_t}(h_t(x) \neq y)$ (agnostic).

As in the classical adversary-provides-a-distribution models, the learner's total loss in the realizable case continues to be given by $\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star) = \sum_{t=1}^T \mathcal{L}_P(h_t, D_t, h^\star)$ (ref. (3.2)), while in the agnostic case, the learner's total regret continues to be expressed as $\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H}) = \sum_{t=1}^T \mathcal{L}_P(h_t, D_t) - \inf_{h\in\mathcal{H}} \sum_{t=1}^T \mathcal{L}_P(h, D_t)$ (ref. (3.5)).

For this model, we formally define what a quantum online learner, a quantum adversary, and quantum online learnability entails. While these definitions are analogous to Definitions 3.2.1, 3.2.2, and 3.2.3, we present them here for the sake of completeness.

Definition 3.3.1 (Quantum online learner)**.** *An algorithm $\mathcal{A}$ is a quantum online learner for a hypothesis class $\mathcal{H}$ if having received a sequence of quantum examples $(|\psi_i\rangle)_{i=1}^t$ (of the form in 2. (a) or 2. (b) of Section 3.3.1) over the first t rounds, $\mathcal{A}$ outputs a hypothesis $h_{t+1} : \mathcal{X} \to \mathcal{Y}$ at round*[20] *$t+1$.*

[18] As in the classical case, the adversary need not be consistent: i.e., they could reveal, for e.g., both $|x, 0\rangle$ and $|x, 1\rangle$ during the T-round protocol.

[19] As an aside, for those who favor a mistake model, it is possible to define it by specifying a threshold ϵ. In this case, a mistake occurs in a round iff $\mathcal{L}_P > \epsilon$, i.e. $\mathcal{L}_{\mathbb{I}}^\epsilon = \mathbf{1}[\mathcal{L}_P > \epsilon]$. We do not investigate this mistake model.

[20] Prior to receiving any examples, $\mathcal{A}$ outputs some arbitrary hypothesis $h_1 : \mathcal{X} \to \mathcal{Y}$ at round 1.

Definition 3.3.2 (Quantum adversary)**.** *Having received a sequence of hypothesis $\mathbf{h}|_t = (h_1, \ldots, h_t)$ from the learner, and with knowledge of its own prior choices of quantum examples, $(|\psi_i\rangle)_{i=1}^t$, over the first t rounds, at round $t+1$, a quantum (online) adversary chooses a distribution D_{t+1} (on $\mathcal{X}$ (realizable) or on $\mathcal{X} \times \mathcal{Y}$ (agnostic)) and discloses the corresponding quantum example $|\psi_{t+1}\rangle$ (with consistent labeling throughout the protocol in the realizable case) to the learner.*

Definition 3.3.3 (Quantum online learnability)**.** *A hypothesis class $\mathcal{H}$ is quantum online learnable if there exists a quantum online learning algorithm $\mathcal{A}$ such that*

- $\mathcal{L}_P(\mathbf{h}_\mathcal{A}, \mathcal{H}) = \sup_{\mathbf{D}, h^\star \in \mathcal{H}} \mathbb{E}[\mathcal{L}_P(\mathbf{h}_\mathcal{A}, \mathbf{D}, h^\star)] = o(T)$ *(realizable),*
- $\mathcal{R}_P(\mathbf{h}_\mathcal{A}, \mathcal{H}) = \sup_{\mathbf{D}} \mathbb{E}[\mathcal{R}_P(\mathbf{h}_\mathcal{A}, \mathbf{D}, \mathcal{H})] = o(T)$ *(agnostic).*

3.3.2 Binary Classification

Under the quantum online learning model described above in Section 3.3, we bound the expected regret (in both the realizable and the agnostic cases) of a quantum online learner for a boolean hypothesis class.

Theorem 3.3.4 (Lower bounds on expected loss/regret for quantum online binary classification)**.** *Let $\mathcal{H} \subseteq \{0,1\}^\mathcal{X}$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every quantum online learner of $\mathcal{H}$, there exists a quantum adversary such that*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \Omega(\mathit{Ldim}(\mathcal{H})) \quad \textit{(realizable), and}$$
$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\mathit{Ldim}(\mathcal{H}) \cdot T}) \quad \textit{(agnostic).}$$

Proof. Let $\mathcal{A}_Q$ be an arbitrary, but fixed, *quantum* online learning algorithm for $\mathcal{H}$. We proceed using a reduction argument. To do this, we examine the scenario where a classical adversary chooses D_t to be a point mass for each t (i.e. the adversary simply chooses an instance x_t (realizable) or $z_t = (x_t, y_t)$ (agnostic) at each t), and analyze the loss/regret bound for the following classical learner $\mathcal{A}_C$ that accesses $\mathcal{A}_Q$ as a "black box". At the t-th round,

1. $\mathcal{A}_C$ provides hypothesis h_t^Q (received from $\mathcal{A}_Q$ in the previous round).
2. Adversary reveals (x_t, y_t) to $\mathcal{A}_C$ (in the realizable case, $y_t = h^\star(x_t)$ for some $h^\star \in \mathcal{H}$).
3. $\mathcal{A}_C$ state prepares $|\psi_t\rangle = |x_t, y_t\rangle$ and passes it as input to $\mathcal{A}_Q$.
4. $\mathcal{A}_Q$ outputs hypothesis h_{t+1}^Q to $\mathcal{A}_C$.

We provide a bound first for the realizable case. Since $\mathcal{A}_C$ plays h_t^Q at each t, it is clear, for our setup, that $\mathcal{L}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, h^\star) = \mathcal{L}_\mathbb{I}(\mathbf{h}_{\mathcal{A}_C}, \mathbf{x}, h^\star) \leq \mathcal{L}_\mathbb{I}(\mathbf{h}_{\mathcal{A}_Q}, \mathbf{x}, h^\star) = \mathcal{L}_P(\mathbf{h}_{\mathcal{A}_Q}, \mathbf{D}, h^\star)$. Taking expectations, and noting that $\mathbb{E}[\mathcal{L}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, h^\star)] = \Omega(\mathrm{Ldim}(\mathcal{H}))$ (from Theorem 3.2.7), we have

shown $\mathbb{E}[\mathcal{L}_P(\mathbf{h}_{\mathcal{A}_Q}, \mathbf{D}, h^\star)] = \Omega(\text{Ldim}(\mathcal{H}))$. Since, $\mathcal{A}_Q$ was chosen arbitrarily, we deduce that $\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \Omega(\text{Ldim}(\mathcal{H}))$.

The agnostic case follows an identical argument; we obtain the following chain of (in)equalities, $\mathcal{R}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, \mathcal{H}) = \mathcal{R}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}_C}, \mathbf{z}, \mathcal{H}) \leq \mathcal{R}_{\mathbb{I}}(\mathbf{h}_{\mathcal{A}_Q}, \mathbf{z}, \mathcal{H}) = \mathcal{R}_P(\mathbf{h}_{\mathcal{A}_Q}, \mathbf{D}, \mathcal{H})$. Taking expectations, and noting $\mathbb{E}[\mathcal{R}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T})$ (from Theorem 3.2.11) and that $\mathcal{A}_Q$ was chosen arbitrarily, we deduce $\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\text{Ldim}(\mathcal{H}) \cdot T})$. ■

Theorem 3.3.5 (Upper bounds on expected loss/regret for quantum online binary classification)**.** *Let $\mathcal{H} \subseteq \{0,1\}^{\mathcal{X}}$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every quantum adversary, there exists a quantum online learner for $\mathcal{H}$ that satisfies*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathcal{O}(\textit{Ldim}(\mathcal{H})) \quad \textit{(realizable), and}$$
$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \mathcal{O}(\sqrt{\textit{Ldim}(\mathcal{H}) \cdot T}) \quad \textit{(agnostic).}$$

Proof. For a naïve algorithm that, at each round t, measures $|\psi_t\rangle$ in the standard basis and employs a classical learner to learn from the observed classical outputs, the desired upper bounds are guaranteed by Theorems 3.2.6 and 3.2.10. ■

3.3.3 Multiclass Classification

Here, we present bounds on the expected regret (in both the realizable and the agnostic cases) of a quantum online learner for a multiclass hypothesis class.

Theorem 3.3.6 (Lower bounds on expected loss/regret for quantum online multiclass classification)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every quantum online learner of $\mathcal{H}$, there exists a quantum adversary such that*

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \Omega(\textit{mcLdim}(\mathcal{H})) \quad \textit{(realizable), and}$$
$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\textit{mcLdim}(\mathcal{H}) \cdot T}) \quad \textit{(agnostic).}$$

Proof. The proof is identical to that of Theorem 3.3.4, where now, for the corresponding classical learners, $\mathbb{E}[\mathcal{L}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, h^\star)] = \Omega(\text{mcLdim}(\mathcal{H}))$ (from Theorem 3.2.14), and $\mathbb{E}[\mathcal{R}_P(\mathbf{h}_{\mathcal{A}_C}, \mathbf{D}, \mathcal{H})] = \Omega(\sqrt{\text{mcLdim}(\mathcal{H}) \cdot T})$ (from Theorem 3.2.20). ■

Theorem 3.3.7 (Upper bounds on expected loss/regret for quantum online multiclass classification)**.** *Let $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$, with $|\mathcal{Y}| = k > 2$, be a hypothesis class, and $h^\star \in \mathcal{H}$. For every quantum adversary,*

there exists a quantum online learner for $\mathcal{H}$ that satisfies

$$\mathbb{E}[\mathcal{L}_P(\mathbf{h}, \mathbf{D}, h^\star)] = \mathcal{O}(\textit{mcLdim}(\mathcal{H})) \quad \textit{(realizable), and}$$
$$\mathbb{E}[\mathcal{R}_P(\mathbf{h}, \mathbf{D}, \mathcal{H})] = \mathcal{O}(\sqrt{\textit{mcLdim}(\mathcal{H}) \cdot T\log(Tk)}) \quad \textit{(agnostic).}$$

Proof. Once again, we consider the measure-and-learn-classically quantum learner for which the desired upper bounds are now guaranteed by Theorems 3.2.13 and 3.2.19. ■

Chapter Summary. Before proceeding to a discussion of the implications of our results, we briefly summarize the results of this chapter in the table below (Table 3.1).

	Boolean		**Multiclass**	
	Realizable	**Agnostic**	**Realizable**	**Agnostic**
Classical (Input-based)	$\Theta(d)$ [Lit88]	$\Theta(\sqrt{dT})$ [BDPSS09, ABED+21]	$\Theta(d)$ [DSBDSS15]	$\Omega(\sqrt{dT})$ [DSBDSS15] $\tilde{\mathcal{O}}(\sqrt{dT})$ [HMR+23]
Classical (Dist.-based)	$\Theta(d)$ (Thms. 3.2.6, 3.2.7)	$\Theta(\sqrt{dT})$ (Thms. 3.2.10, 3.2.11)	$\Theta(d)$ (Thms. 3.2.13, 3.2.14)	$\Omega(\sqrt{dT})$ (Thm. 3.2.20) $\mathcal{O}(\sqrt{dT\log(Tk)})$ (Thm. 3.2.19)
Quantum	$\Theta(d)$ (Thms. 3.3.4, 3.3.5)	$\Theta(\sqrt{dT})$ (Thms. 3.3.4, 3.3.5)	$\Theta(d)$ (Thms. 3.3.6, 3.3.7)	$\Omega(\sqrt{dT})$ (Thm. 3.3.6) $\mathcal{O}(\sqrt{dT\log(Tk)})$ (Thm. 3.3.7)

Table 3.1: An overview of expected regret bounds for **online** learning in the canonical classical (adversary-provides-an-input), classical *adversary-provides-a-distribution*, and quantum paradigms. Our novel contributions are presented in boxes shaded gray. d denotes the Littestone dimension for Boolean cases ($d := \text{Ldim}(\mathcal{H})$), and the multiclass Littlestone dimension for Multiclass cases ($d := \text{mcLdim}(\mathcal{H})$).

Notes on Table 3.1

- In all cases, we state known results that exhibit the tightest dependence on the combinatorial parameters that characterize learning in the respective settings.
- In the (canonical) classical online multiclass agnostic case, the $\tilde{\mathcal{O}}(\sqrt{\text{mcLdim}(\mathcal{H})T})$ bound hides a $\sqrt{\log\left(\frac{T}{\text{mcLdim}(\mathcal{H})}\right)}$ factor [HMR+23].
- The definition of loss/regret differs between the canonical classical adversary-provides-an-input model (Section 3.1) and both the classical adversary-provides-a-distribution model (Section 3.2) and the quantum online model (Section 3.3). Specifically, the former employs the indicator loss (mistake model), whereas the latter two involve probabilistic losses.

3.3.4 Takeaways

Before we end this section on online learning with quantum examples, we note that the proofs for the expected regret upper bounds were established by a quantum online learner that performs a measurement and subsequently learns classically. The fact that the upper bounds thus obtained are *identical* to the lower bounds, in all but one setting[21], shows that the performance of this measure-and-learn-classically learner is as good as the best "genuine" quantum online learner in these settings. We feel that this is consistent with the overall message of this thesis, viz., that there is limited power in quantum examples to speed up learning especially when the adversary is allowed to play arbitrary distributions (including very degenerate ones like point masses).

Recently, [HMR+23] improved the classical upper bound for the online multiclass agnostic case to $\tilde{\mathcal{O}}(\sqrt{\text{mcLdim}(\mathcal{H})T})$[22], which removes all k dependence (cf. the $\sqrt{\log k}$ factor that appears in our corresponding *quantum* upper bound in Theorem 3.3.7). Meanwhile, we believe our analysis in the proof of Theorem 3.2.19 (which establishes the classical bound for the measure-and-learn-classically quantum learner in Theorem 3.3.7) is tight, and so we suspect that any removal of the k-dependence in this setting would involve investigating into a "genuine" quantum online learning algorithm, which may involve a quantum-specific combinatorial parameter that characterizes learning. We identify this as an open question for future work.

- What is the *tight* expected regret bound for quantum online multiclass agnostic learning when the label space is unbounded (i.e., when the number of classes $k \to \infty$)?

In the batch setting, the sample complexity upper bounds were trivial to establish due to the quantum learner's ability to measure quantum examples and learn classically on the resulting output. In the online setting, the expected regret lower bounds, in turn, were trivial due to the adversary's ability to provide point masses D_t at each t, rendering each quantum example equivalent to a classical example. This prompts us to ask the following question.

- What happens when we impose restrictions on D_t to force it away from a point mass? Would the expected regret bounds for the canonical classical online model (Section 3.1), classical adversary-provides-a-distribution model (Section 3.2), and the quantum online model (Section 3.3) all diverge from one another?

[21] The exception is the online multiclass agnostic case, where the quantum upper and lower bounds differ by a factor of $\sqrt{\log(Tk)}$.

[22] Here, $\tilde{\mathcal{O}}(\cdot)$ hides $\sqrt{\log\left(\frac{T}{\text{mcLdim}(\mathcal{H})}\right)}$ factors.

CHAPTER 4

Smoothed Learning with Quantum Examples

The concluding question of Chapter 3 serves as the central motivation for this chapter. Throughout the two preceding chapters (Chapters 2 and 3), we studied learning in the *distribution-free* (i.e., worse-case) setting, where performance is measured against the worst-case input distribution. In each setting considered, we observed that learning from classical examples and learning from quantum examples yields guarantees of essentially the same order and form, whether measured in terms of sample complexity, loss, or regret.

A particularly notable feature of these results is that the hard distributions underlying the corresponding lower bounds are highly concentrated, either behaving as *near* point masses or, in some cases, being exact point masses. This phenomenon first appeared in the lower bounds for batch learning of Boolean functions (Theorem 2.1.7 in Chapter 2, reproduced from [AdW18]), and reappeared across the other learning settings studied thus far, including batch learning of multiclass functions (Section 2.2) and online learning of both Boolean and multiclass functions (Section 3.3). More broadly, it appears to be a recurring theme that sample complexity non-separations occur when the sampled variables concentrate on a small, a priori chosen, region of the domain.

On the other hand, when the underlying distribution is *uniform* (i.e., in arguably the quantum "best-case"), Arunachalam et al. [ACL$^+$21] establish an unbounded separation between the quantum $\mathcal{O}(k^{1.5}(\log k)^2)$ and classical $\Omega(kn)$ sample complexities for batch learning of k-Fourier-sparse n-bit Boolean functions. Taken together, these results highlight a sharp contrast:

- for *some* function classes, there exists an unbounded separation between quantum and classical sample complexities for batch learning under *uniform* examples;
- for *every* function class, in each learning setting considered thus far, there is no asymptotic separation between learning from quantum and classical examples under *some* sufficiently concentrated worst-case distribution over examples.

This naturally raises the question of what the landscape in between these two extremes looks like. In particular, are there function classes for which one can demonstrate an unbounded separation between quantum and classical learning under *near-uniform* example distributions? Conversely,

are there function classes for which a separation emerges when one deviates only *slightly* from the worst-case, near point-mass distributions?

One way to probe this landscape is through the lens of the *smoothed learning* framework introduced in [HRS20], which enforces a controlled level of anti-concentration in the example distribution. In the batch setting, this interpolates between highly concentrated and highly uniform example distributions, while later in the chapter we will see that an analogous notion in the online setting interpolates between fully adversarial and statistical i.i.d. learning. Here, we first develop a framework for learning from quantum examples under smooth distributions and use it to study the robustness of quantum sample-complexity separations.

4.1 Preliminaries

Recall that in the quantum PAC learning model (see Section 2.1.5), a learner is provided oracle access to quantum examples of the form

$$\sum_{x\in\{0,1\}^n} \sqrt{D(x)}\,|x, h^\star(x)\rangle\,, \tag{4.1}$$

for some target hypothesis $h^\star$ in a known hypothesis class $\mathcal{H}$, where the underlying distribution D over $\mathcal{X} = \{0,1\}^n$ was *unrestricted*.

4.1.1 Smoothed Distributions

In the smoothed setting, we relax this by requiring D to satisfy the following anti-concentration constraint.

Definition 4.1.1 (σ-smooth distribution)**.** *A distribution D on $\mathcal{X}$ is σ-smooth with respect to a reference distribution μ, for $\sigma \in (0,1]$, if $D \ll \mu$ (i.e., D is absolutely continuous*[1] *with respect to μ) and*

$$\operatorname{ess\,sup} \frac{dD}{d\mu} \le \frac{1}{\sigma}. \tag{4.2}$$

Our foray into smoothed learning is motivated by [ACL+21]'s results on exactly learning k-Fourier-sparse Boolean functions under the *uniform* distribution; accordingly, we specialize μ to the uniform distribution throughout. Definition 4.1.1 then reduces to the following explicit condition, with the uniform distribution itself recovered at $\sigma = 1$.

[1] Absolute continuity ensures the Radon–Nikodym derivative $\frac{dD}{d\mu}$ exists μ-almost everywhere, so the essential supremum in (4.2) is well-defined. Explicitly, $D \ll \mu$ means: $\forall A \subseteq \mathcal{X},\ \mu(A) = 0 \implies D(A) = 0$.

Definition 4.1.2 (σ-smooth distribution with respect to uniform)**.** *A distribution D on $\mathcal{X} = \{0,1\}^n$ is σ-smooth (with respect to the uniform distribution) for $\sigma \in (0,1]$ if*

$$\max_{x \in \mathcal{X}} D(x) \leq \frac{1}{\sigma 2^n}. \tag{4.3}$$

As before, the goal of the learning task is to find a hypothesis h such that the generalization error $\mathcal{L} = \mathbb{P}_{x \sim D}(h(x) \neq h^\star(x))$ is small, where the unknown distribution D is now constrained to be σ-smooth. We formalize this as follows.

Definition 4.1.3 (Smoothed Quantum PAC Learner)**.** *Let $\sigma \in (0,1]$. An algorithm $\mathcal{A}$ is an (ϵ, δ)-quantum PAC learner for a hypothesis class $\mathcal{H}$ under σ-smooth distributions if, for every σ-smooth distribution D on $\mathcal{X}$ and every target hypothesis $h^\star \in \mathcal{H}$, the learner $\mathcal{A}$, given m copies of the quantum example state in* (4.1)*, outputs a hypothesis h such that $\mathbb{P}[\mathcal{L} \leq \epsilon] \geq 1 - \delta$, where $\mathcal{L} = \mathbb{P}_{x \sim D}(h(x) \neq h^\star(x))$, and the outer probability is taken over the learner's internal randomness.*

A hypothesis class $\mathcal{H}$ is *smoothed quantum PAC-learnable* if there exists an (ϵ, δ)-quantum PAC learner for $\mathcal{H}$ under σ-smooth distributions. It is *smoothed quantum exactly learnable* if it admits a $(0, \delta)$-quantum PAC learner under σ-smooth distributions. The corresponding *sample complexity* is the minimum number m of quantum examples required in the worst case over all target hypotheses $h^\star \in \mathcal{H}$ and all σ-smooth distributions D on $\mathcal{X}$.

As mentioned at the start of this chapter, [ACL+21] established an unconditional and unbounded sample-complexity separation between quantum and classical learning for the task of exactly learning k-Fourier-sparse Boolean functions under the uniform distribution. Our goal was to investigate whether this separation persists when the example distribution is allowed to deviate slightly from uniformity. The framework of σ-smooth distributions now provides a natural way to formalize this interpolation between the uniform regime and increasingly concentrated distributions.

4.1.2 Basics of Boolean Fourier Analysis

Before we proceed, we first recall the basics of Fourier analysis on the Boolean cube (see [O'D14] for an in-depth background). For functions $f, g : \mathcal{X} \to \mathbb{R}$, where $\mathcal{X} = \{0,1\}^n$, we define the inner product

$$\langle f, g \rangle \;=\; \mathbb{E}_{x \sim \mathcal{X}}[f(x)g(x)],$$

where the expectation is with respect to the uniform distribution on $\mathcal{X}$. For each $S \in \{0,1\}^n$, the corresponding *character function* is given by

$$\chi_S(x) \;:=\; (-1)^{S \cdot x},$$

where $S \cdot x = \sum_{i=1}^{n} S_i x_i$ is the usual inner product over $\mathbb{F}_2$. The collection $\{\chi_S\}_{S \in \{0,1\}^n}$ forms an orthonormal basis for the space of real-valued functions on $\{0,1\}^n$, and so every $f : \{0,1\}^n \to \mathbb{R}$ has a unique *Fourier expansion*

$$f(x) \;=\; \sum_{S \in \{0,1\}^n} \widehat{f}(S)\, \chi_S(x),$$

where $\widehat{f}(S) = \langle f, \chi_S \rangle$ is the *Fourier coefficient* of f at S.

Parseval's identity states that $\sum_{S \in \{0,1\}^n} \widehat{f}(S)^2 = \mathbb{E}_{x \sim \mathcal{X}}[f(x)^2]$. For Boolean functions $f : \{0,1\}^n \to \{-1,1\}$, this gives $\sum_{S \in \{0,1\}^n} \widehat{f}(S)^2 = 1$, so $\{\widehat{f}(S)^2\}_{S \in \{0,1\}^n}$ forms a probability distribution. For $f : \{0,1\}^n \to \mathbb{R}$, the *Fourier support* is $\operatorname{supp}(\widehat{f}) = \{S : \widehat{f}(S) \neq 0\}$ and the *Fourier sparsity* is $|\operatorname{supp}(\widehat{f})|$. We say f is *k-Fourier-sparse* if $|\operatorname{supp}(\widehat{f})| \leq k$. Additionally, the *Fourier span* of f is the span of $\operatorname{supp}(\widehat{f})$ in $\mathbb{F}_2^n$, whereas the *Fourier dimension* of f is the dimension of its Fourier span.

4.1.3 Fourier Sampling

Any learning setting with an underlying Fourier structure is amenable to $\mathcal{O}(1)$ quantum Fourier sampling [BV93]. This procedure is the central technical tool underlying many quantum advantages in learning from quantum examples. In contrast, in the classical setting, Fourier sampling is known to require at least $\Omega(2^n)$ classical samples[2] [AC16]. We will now make precise how Fourier sampling can be implemented using *uniform quantum examples*. Before we proceed, we introduce the H (Hadamard) gate as a preliminary.

Definition 4.1.4 (H (Hadamard) gate)**.** *The H (Hadamard) gate is a quantum gate that operates on one qubit. It maps computational basis states to equal superpositions. For $b \in \{0,1\}$,*

$$H\,|b\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle + (-1)^b\,|1\rangle\right).$$

Lemma 4.1.5 (Lemma 4 in [ACL+21])**.** *Let $f : \{0,1\}^n \to \{-1,1\}$. There exists a procedure that uses one uniform quantum example and satisfies the following: with probability $1/2$ it outputs an S drawn from the distribution $\{\widehat{f}(S)^2\}_{S \in \{0,1\}^n}$, otherwise it rejects.*

Proof. Using a uniform quantum example, we first work toward preparing the phase-encoded state $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} f(x)\,|x\rangle$. Since $f(x) \in \{-1,1\}$, we encode the label into the computational basis bit $\frac{1-f(x)}{2} \in \{0,1\}$. Accordingly, the uniform quantum example corresponds to the state $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x, \frac{1-f(x)}{2}\rangle$.

[2] More precisely, [AC16] proves this lower bound even for the stronger model of classical query access, rather than merely access to uniform classical examples.

Applying a Hadamard gate (Definition 4.1.4) to the final qubit yields the state

$$\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle \frac{1}{\sqrt{2}} \left(|0\rangle + f(x)\,|1\rangle\right).$$

Measuring the final qubit yields outcome 0 with probability $\frac{1}{2}$, in which case the remaining state collapses to the uniform superposition $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} |x\rangle$, which contains no information about f, and the procedure rejects. Otherwise, outcome 1 is obtained, also with probability $\frac{1}{2}$, and the post-measurement state is $\frac{1}{\sqrt{2^n}} \sum_{x \in \{0,1\}^n} f(x)\,|x\rangle$. In this case, applying Hadamard gates to all n qubits yields

$$\sum_{S \in \{0,1\}^n} \widehat{f}(S)\,|S\rangle\,.$$

Measuring this state in the computational basis outputs S with probability $\widehat{f}(S)^2$. This procedure is known as *quantum Fourier sampling*. ■

To boost the success probability from $\frac{1}{2}$ to at least $1 - \delta$, it suffices to run the procedure t times independently, declaring success upon the first repetition that outputs an S. The probability that every repetition rejects is $(1/2)^t$, so choosing $t = \lceil \log_2(1/\delta) \rceil$ ensures that the failure probability is at most δ, using $\mathcal{O}(\log(1/\delta))$ independent uniform quantum examples in total. Since this overhead is independent of both n and k, the overall sample complexity remains $\widetilde{\mathcal{O}}(1)$.

4.1.4 Exact Learning Fourier-Sparse Boolean Functions

We now specialize attention to the hypothesis class of k-Fourier-sparse Boolean functions,

$$\mathcal{H}^{\mathrm{FS}} = \{h : \{0,1\}^n \to \{-1,1\} : |\operatorname{supp}(\widehat{h})| \leq k\}.$$

As discussed previously, [ACL+21] established an unconditional and unbounded separation between the classical and quantum sample complexities of exactly learning $\mathcal{H}^{\mathrm{FS}}$ under the uniform distribution. While their strongest upper bound achieves a quantum sample complexity of $\mathcal{O}(k^{1.5}(\log k)^2)$, the authors also present a simpler argument yielding an upper bound of $\mathcal{O}(k^2 \log k)$. Since this simpler bound already remains independent of n, it still suffices to demonstrate an unbounded quantum advantage over classical learning. Accordingly, we adopt this argument throughout the remainder of the chapter.

The following structural results concerning k-Fourier-sparse Boolean functions will be useful in the analysis that follows.

Lemma 4.1.6 (Theorem 12 of [GOS+11])**.** *Let $k \geq 2$. The Fourier coefficients of a k-Fourier-sparse Boolean function $f : \{0,1\}^n \to \{-1,1\}$ are integer multiples of $2^{1-\lfloor \log k \rfloor}$.*

Lemma 4.1.7 (Theorem 1.2 of [San19])**.** *The Fourier dimension of a k-Fourier-sparse Boolean function $f : \{0,1\}^n \to \{-1,1\}$ is $\mathcal{O}(\sqrt{k}\log k)$.*

We restate below the uniform-distribution exact-learning guarantee of [ACL+21].

Theorem 4.1.8 (Discussion under Theorem 6 and Section 3.1.2 of [ACL+21])**.** *Let $k \geq 2$. The hypothesis class $\mathcal{H}^{\mathrm{FS}}$ is exactly learnable under the uniform distribution by a quantum PAC learner with sample complexity $\mathcal{O}(k^2 \log k)$.*

Proof. Let $h^\star \in \mathcal{H}^{\mathrm{FS}}$ be arbitrary. By Lemma 4.1.6, every nonzero Fourier coefficient $\widehat{h^\star}(S)$ satisfies

$$|\widehat{h^\star}(S)| \geq 2^{1-\lfloor \log k \rfloor} \geq \frac{1}{k}.$$

Consequently, every $S \in \mathrm{supp}(\widehat{h^\star})$ satisfies $\widehat{h^\star}(S)^2 \geq 1/k^2$. Applying the Fourier sampling procedure of Lemma 4.1.5 to our uniform quantum example, any fixed $S \in \mathrm{supp}(\widehat{h^\star})$ is output with probability at least $1/k^2$ whenever the procedure succeeds. Hence, after T successful Fourier samples,

$$\mathbb{P}[S \text{ is never observed}] \leq \left(1 - \frac{1}{k^2}\right)^T \leq e^{-T/k^2}.$$

Since $|\mathrm{supp}(\widehat{h^\star})| \leq k$, a union bound gives

$$\mathbb{P}[\text{some element of } \mathrm{supp}(\widehat{h^\star}) \text{ is missed}] \leq k e^{-T/k^2}.$$

Taking $T = \mathcal{O}(k^2 \log k)$ allows us to obtain the entire Fourier support of $h^\star$ with constant probability. Now, knowing the entire Fourier support of $h^\star$ implies we have full knowledge of its Fourier span. From Lemma 4.1.7, we further know that the dimension of this span is bounded by $r = \mathcal{O}(\sqrt{k}\log k)$. Therefore, the classical procedure outlined in Section 3.1.2 of [ACL+21]) allows us to recover $h^\star$ with $\mathcal{O}(rk\log k) = \mathcal{O}(k^{1.5}(\log k)^2)$ uniform classical examples[3]. Therefore, the overall quantum sample complexity is bounded by the dominant term, $\mathcal{O}(k^2 \log k) + \mathcal{O}(k^{1.5}(\log k)^2) = \mathcal{O}(k^2 \log k)$, which completes our proof. ■

Remark 2 (1-Fourier-Sparse)**.** *Observe that Theorem 4.1.8 does not include the case $k = 1$, as the Fourier coefficient of a 1-Fourier-sparse function is 0-granular, which puts it out of the scope of Lemma 4.1.6. However, learning a 1-Fourier-sparse $h^\star$ in this context is trivial. With one successful application of the quantum Fourier sampling procedure outlined in Lemma 4.1.5, we learn the sole element, S_1, in the Fourier support of $h^\star$. Then, with one classical example $(x, h^\star(x))$, we fully recover $h^\star$ by learning the lone non-zero Fourier coefficient $\widehat{h^\star}(S_1) = \frac{h^\star(x)}{\chi_{S_1}(x)}$, which, in this case, would take the form of a sign.*

[3] Recall that these (uniform) classical examples can be trivially obtained by measuring the uniform quantum examples in the computational basis.

Discussion on the Proof of Theorem 4.1.8. Observe that the proof naturally decomposes into two distinct phases. In the first phase, we apply a *quantum* Fourier sampling procedure to learn the exact Fourier support of $h^\star$. Then, in the second phase, we revert to *classical* means to fully learn $h^\star$. In general, we can characterize the first phase as *quantumly* extracting *some* structural Fourier property of $h^\star$ (for instance, [ACL+21] first uses Fourier sampling to learn the Fourier *span* of $h^\star$), and the second phase as *classically* completing the learning process to recover $h^\star$. Before we proceed, we discuss two alternate ways in which the classical recovery of $h^\star$ could have proceeded in Theorem 4.1.8 once we had quantumly obtained the entire Fourier support of $h^\star$.

- **Empirical "estimation" of the Fourier coefficients of $h^\star$:** While the quantum Fourier sampling phase does also provide us with an estimate of the *squared* magnitudes of the Fourier coefficients, $\widehat{h^\star}(S_i)^2$, there is no immediate way to recover the signs of each $\widehat{h^\star}(S_i)$. So, we proceed via the standard learning-theoretic strategy (cf. Proposition 3.30 in [O'D14]) as follows. For each $S_i \in \mathrm{supp}(\widehat{h^\star})$, with access to uniform classical examples $(x, h^\star(x))$, we can compute $h^\star(x)\chi_{S_i}(x) \in \{-1, 1\}$ and, therefore, empirically estimate $\widehat{h^\star}(S_i) = \mathbb{E}_{x\sim\{0,1\}^n}[h^\star(x)\chi_{S_i}(x)]$. Indeed, due to the granularity of $\widehat{h^\star}(S_i)$ (Lemma 4.1.6), we merely need the empirical estimate to be accurate to within[4] $\pm\frac{1}{k}$ after which, by rounding to the nearest multiple of $2^{1-\lfloor \log k \rfloor}$, we would be able to pin down $\widehat{h^\star}(S_i)$ exactly. A standard application of Hoeffding's inequality, paired with a union bound over the k support elements, guarantees that $\mathcal{O}(k^2 \log k)$ classical examples are sufficient to recover all k non-zero Fourier coefficients with constant probability. While this approach is inferior to the one presented in the second phase of the proof of Theorem 4.1.8, it nevertheless guarantees the same $\mathcal{O}(k^2 \log k)$ overall sample complexity.
- **Solving a linear system for the Fourier coefficients of $h^\star$:** Without loss of generality, we let $\mathrm{supp}(\widehat{h^\star}) = \{S_1, \ldots, S_k\}$. The function we are trying to learn has the form $h^\star(x) = \sum_{j=1}^{k} \widehat{h^\star}(S_j)\chi_{S_j}(x)$. So, each uniform classical example $(x, h^\star(x))$ gives us a linear equation in the k unknown Fourier coefficients, $\{\widehat{h^\star}(S_1), \ldots, \widehat{h^\star}(S_k)\}$. After m samples, $\{(x_i, h^\star(x_i))\}_{i=1}^m$, we would obtain the following matrix system:

$$X\widehat{\mathbf{h^\star}} = \mathbf{h}^\star, \tag{4.4}$$

  where X is an $m \times k$ matrix with entries $X_{i,j} = \chi_{S_j}(x_i)$, $\widehat{\mathbf{h^\star}}$ is the *fixed* $k \times 1$ vector containing the unknown Fourier coefficients (i.e., $\widehat{\mathbf{h^\star}}_j = \widehat{h^\star}(S_j)$), and $\mathbf{h}^\star$ is an $m \times 1$ vector containing the labels of the m samples (i.e., $\mathbf{h^\star}_i = h^\star(x_i)$). We would be able to solve for $\widehat{\mathbf{h^\star}}$ as long as X has full column rank k. Therefore, the sample complexity is governed by how quickly sampling makes X achieve this desired full rank. Structurally speaking, each row $X_i \in \{-1, 1\}^{1\times k}$

[4] Note that $2^{1-\lfloor \log k \rfloor} \geq \frac{2}{k}$.

is an *independent*, *bounded* random vector. Furthermore, because we sample *uniformly* and the character functions are orthonormal, X_i also has the *isotropic* property, namely $\mathbb{E}_{x_i \sim \{0,1\}^n}[X_i^T X_i] = \mathbb{I}_{k \times k}$. Following the discussions of Theorem 4.6.1 in Sections 4.7 and 5.6 of [Ver18], $m = \mathcal{O}(k \log k)$ uniform classical examples would guarantee that X has full rank, which would allow us to solve the system in (4.4) to obtain all k unknown Fourier coefficients, and thereby recover $h^\star$. It is worth noting here that, as compared to our earlier approach of estimating the Fourier coefficients *separately*, learning them *simultaneously* here has unsurprisingly proved to be more efficient. Indeed, of the three classical methods for the "classical recovery" phase of the proof of Theorem 4.1.8, this is the most sample-efficient.

In fact, we may interpret the main overall sample complexity result of $\mathcal{O}(k^{1.5}(\log k)^2)$ in [ACL+21] as attempting to hit the "sweet spot" regarding sample efficiency in both phases, using fewer quantum samples to learn a weaker Fourier property (the Fourier span) of $h^\star$, while then using more classical samples to recover $h^\star$ exactly.

4.2 Smoothed Batch Learning of Fourier-Sparse Boolean Functions

The preceding theorem (Theorem 4.1.8) showed that $\mathcal{H}^{\mathrm{FS}}$ admits an unbounded quantum advantage under the uniform distribution. We now investigate the robustness of this separation when the example distribution is allowed to deviate slightly from uniformity. In particular, we study the quantum PAC learnability of $\mathcal{H}^{\mathrm{FS}}$ under σ-smooth (w.r.t. uniform) distributions, focusing on the regime where σ is close to 1. Our main result shows that, for sufficiently small deviations from uniformity, the *first phase* of the general learning procedure, i.e., using quantum examples to identify a collection $\mathcal{F}$ containing the entire Fourier support of $h^\star$, can still be achieved with a quantum sample complexity that is completely independent of n.

Theorem 4.2.1 (Fourier sampling guarantee for smoothed quantum batch learning of k-Fourier-sparse Boolean functions)**.** *Let $k \geq 2$ and let $0 \leq \xi \leq \frac{4^{-\lfloor \log k \rfloor}}{1+4^{-\lfloor \log k \rfloor}}$, so that $\sigma := 1 - \xi$ satisfies $\xi \leq \sigma \cdot 4^{-\lfloor \log k \rfloor}$. For any k-Fourier-sparse Boolean function $h^\star \in \mathcal{H}^{\mathrm{FS}}$, and under any distribution that is σ-smooth with respect to the uniform distribution (Definition 4.1.2), $\mathcal{O}\left(\frac{k^2 \log k}{\sigma}\right)$ quantum examples are sufficient to identify a collection $\mathcal{F}$ such that* $\mathrm{supp}(\widehat{h^\star}) \subseteq \mathcal{F}$.

Proof. Let $h^\star \in \mathcal{H}^{\mathrm{FS}}$, and σ-smooth (w.r.t. uniform) distribution D be arbitrarily fixed. We attempt to apply the Fourier sampling procedure outlined in Lemma 4.1.5 to our quantum example, $\sum_{x \in \{0,1\}^n} \sqrt{D(x)}\, |x, \frac{1-h^\star(x)}{2}\rangle$. If we succeed after applying a Hadamard gate to the final qubit and measuring it, the post-measurement state would be $\sum_{x \in \{0,1\}^n} \sqrt{D(x)} h^\star(x)\, |x\rangle$. If we now apply

Hadamard gates to all n qubits, we get

$$\sum_{x\in\{0,1\}^n} \sqrt{D(x)}h^\star(x)\Big(\frac{1}{\sqrt{2^n}}\sum_{S\in\{0,1\}^n}(-1)^{x\cdot S}\,|S\rangle\Big)$$
$$=\sum_{S\in\{0,1\}^n}\frac{1}{2^n}\sum_{x\in\{0,1\}^n}\underbrace{\sqrt{D(x)2^n}h^\star(x)}_{g(x)}(-1)^{x\cdot S}\,|S\rangle\,.$$

Defining $g(x)=\sqrt{D(x)2^n}h^\star(x)$, and recalling that

$$\widehat{g}(S)=\mathbb{E}_{x\sim\{0,1\}^n}[g(x)\chi_S(x)]=\frac{1}{2^n}\sum_{x\in\{0,1\}^n}g(x)(-1)^{x\cdot S},$$

the resulting state is precisely $\sum_{S\in\{0,1\}^n}\widehat{g}(S)\,|S\rangle$. So, measuring this state in the computational basis outputs S with probability $\widehat{g}(S)^2$. In other words, the procedure is now a mechanism to *Fourier sample from* g, from which we show it is nonetheless possible to collect the entire Fourier support of $h^\star$.

First, we define $a(x):=\sqrt{D(x)2^n}$ to write $g(x)=a(x)h^\star(x)$. Then,

$$g(x)=a(x)h^\star(x)=\sum_{T,U\in\{0,1\}^n}\widehat{a}(T)\widehat{h^\star}(U)\chi_T(x)\chi_U(x)=\sum_{T,U\in\{0,1\}^n}\widehat{a}(T)\widehat{h^\star}(U)\chi_{T\oplus U}(x),$$

where we used the identity $\chi_T(x)\chi_U(x)=\chi_{T\oplus U}(x)$. Regrouping terms according to $S:=T\oplus U$ (for each fixed T, the map $U\mapsto T\oplus U$ is a bijection on $\{0,1\}^n$), we obtain

$$\widehat{g}(S)=\sum_{T\in\{0,1\}^n}\widehat{a}(T)\widehat{h^\star}(S\oplus T). \tag{4.5}$$

Since $|\operatorname{supp}(\widehat{h^\star})|\le k$, there are at most k non-zero terms in the sum above. Without loss of generality, let $\operatorname{supp}(\widehat{h^\star})=\{S_1,\ldots,S_k\}$. We establish our result by providing a lower bound on the probability that a Fourier sample from g yields an element in $\operatorname{supp}(\widehat{h^\star})$, i.e., we proceed by lower bounding $\widehat{g}(S_i)^2$. For each $i\in[k]$, we can rewrite (4.5) as

$$\widehat{g}(S_i)=\sum_{j=1}^{k}\widehat{a}(S_i\oplus S_j)\widehat{h^\star}(S_j)=\widehat{a}(0^n)\widehat{h^\star}(S_i)+\sum_{j\neq i}\widehat{a}(S_i\oplus S_j)\widehat{h^\star}(S_j). \tag{4.6}$$

Now, by Parseval's identity applied to a,

$$\sum_{S\in\{0,1\}^n}\widehat{a}(S)^2=\mathbb{E}_{x\sim\{0,1\}^n}[a(x)^2]=\mathbb{E}_{x\sim\{0,1\}^n}[D(x)2^n]=\sum_{x\in\{0,1\}^n}D(x)=1.$$

On the other hand, since $a(x) \leq \frac{1}{\sqrt{\sigma}}$ for all $x \in \{0,1\}^n$ by σ-smoothness, we observe that

$$1 = \mathbb{E}_{x\sim\{0,1\}^n}[a(x)^2] \leq \max_{x\in\{0,1\}^n} a(x) \cdot \mathbb{E}_{x\sim\{0,1\}^n}[a(x)] \leq \frac{1}{\sqrt{\sigma}} \cdot \widehat{a}(0^n),$$

where we additionally use $\mathbb{E}_{x\sim\{0,1\}^n}[a(x)] = \mathbb{E}_{x\sim\{0,1\}^n}[a(x)\chi_{0^n}(x)] = \widehat{a}(0^n)$ in the final inequality. Altogether, we have shown that

$$\widehat{a}(0^n) \geq \sqrt{\sigma}. \tag{4.7}$$

Now, since $\widehat{a}(0^n)^2 \geq \sigma$, Parseval gives

$$\sum_{S\neq 0^n} \widehat{a}(S)^2 = 1 - \widehat{a}(0^n)^2 \leq 1 - \sigma = \xi.$$

In particular, since each S_j is distinct, and $S_i \oplus S_j \neq 0^n$ for $j \neq i$, the cross-term coefficients in (4.6) satisfy $\sum_{j\neq i} \widehat{a}(S_i \oplus S_j)^2 \leq \sum_{S\neq 0^n} \widehat{a}(S)^2 = \xi$. Furthermore, Parseval on Boolean $h^\star$ implies $\sum_{j\neq i} \widehat{h^\star}(S_j)^2 \leq \sum_{j=1}^{k} \widehat{h^\star}(S_j)^2 = 1$. Therefore, by the Cauchy-Schwarz inequality, we have the following inequality on the cross term in (4.6):

$$\sum_{j\neq i} |\widehat{a}(S_i \oplus S_j)||\widehat{h^\star}(S_j)| \leq \Big(\sum_{j\neq i} \widehat{a}(S_i \oplus S_j)^2\Big)^{\frac{1}{2}} \Big(\sum_{j\neq i} \widehat{h^\star}(S_j)^2\Big)^{\frac{1}{2}} \leq \sqrt{\xi} \cdot 1 = \sqrt{\xi}. \tag{4.8}$$

Next, applying the triangle inequality to (4.6), using (4.7) to bound the main term and (4.8) to bound the cross term, we obtain

$$|\widehat{g}(S_i)| \geq \widehat{a}(0^n)\,|\widehat{h^\star}(S_i)| - \sum_{j\neq i} |\widehat{a}(S_i \oplus S_j)|\,|\widehat{h^\star}(S_j)| \geq \sqrt{\sigma}\,|\widehat{h^\star}(S_i)| - \sqrt{\xi}. \tag{4.9}$$

Now, due to Lemma 4.1.6, every non-zero Fourier coefficient of $h^\star$ satisfies $|\widehat{h^\star}(S_i)| \geq 2^{1-\lfloor \log k \rfloor}$. Substituting into (4.9) and requiring $\sqrt{\xi} \leq \frac{\sqrt{\sigma}}{2} \cdot 2^{1-\lfloor \log k \rfloor}$, i.e., when $\xi \leq \sigma \cdot 4^{-\lfloor \log k \rfloor}$, we obtain

$$|\widehat{g}(S_i)| \geq \frac{\sqrt{\sigma}}{2} \cdot 2^{1-\lfloor \log k \rfloor}, \tag{4.10}$$

and therefore $\widehat{g}(S_i)^2 \geq \sigma \cdot 4^{-\lfloor \log k \rfloor} = \Omega(\sigma/k^2)$. This implies that each successful application of the Fourier sampling procedure outputs any specific support element $S_i \in \mathrm{supp}(\widehat{h^\star})$ with probability at least σ/k^2. Applying the coupon collector argument exactly as in Theorem 4.1.8, $T = \mathcal{O}\left(\frac{k^2 \log k}{\sigma}\right)$ Fourier samples are sufficient to observe all k elements of $\mathrm{supp}(\widehat{h^\star})$, and thereby identify a collection $\mathcal{F}$ containing the entire Fourier support, with constant probability. This completes our proof. ■

As a sanity check, setting $\sigma = 1$ in Theorem 4.2.1 exactly recovers the sample complexity guarantee from the *first phase* of the proof of Theorem 4.1.8. Furthermore, observe that the permitted regime for ξ in Theorem 4.2.1 is

$$0 \leq \xi \leq \frac{4^{-\lfloor \log k \rfloor}}{1 + 4^{-\lfloor \log k \rfloor}} = \mathcal{O}\left(\frac{1}{k^2}\right). \tag{4.11}$$

Thus, whenever the example distribution is $(1 - \xi)$-smooth for $\xi = \mathcal{O}(1/k^2)$, we preserve the n-independent quantum sample complexity of this initial phase.

Discussion on the classical recovery phase in the smoothed setting. In the proof of Theorem 4.1.8, and the subsequent discussion, we established three distinct classical avenues to recover $h^\star$ once its entire Fourier support was known. Transitioning to the smoothed setting introduces two critical complications. First, while our sampled collection $\mathcal{F}$ *contains* the entire Fourier support of $h^\star$, it may also contain extraneous elements. Second, we no longer have access to *uniform* classical examples. Below, we outline how these differences impact each recovery method and discuss their potential for completing the learning algorithm:

- **Learning from the Fourier span:** The uniform distribution recovery relied on bounding the dimension of the Fourier span and invoking a classical learning result due to [HR16]. To adapt this approach, we would need to restrict the concept class over the larger span dictated by $\mathcal{F}$. More fundamentally, we would have to carefully generalize the uniform distribution sample complexity guarantees of [HR16] to hold under σ-smooth distributions, which remains a non-trivial challenge.
- **Empirical estimation of the Fourier coefficients:** In the uniform case, Fourier coefficients were approximated via simple empirical averages. Under an unknown σ-smooth distribution D, creating an unbiased estimator would require importance weighting by the specific probabilities $D(x)$. Since D is unknown to the learner, this poses a significant, and perhaps insurmountable, bottleneck for direct empirical estimation.
- **Solving a linear system for the Fourier coefficients:** We can still construct a linear system analogous to (4.4), where the columns correspond to the elements of $\mathcal{F}$. While the rows, populated by examples $(x_i, h^\star(x_i))$ drawn from a σ-smooth distribution D, remain independent and bounded, they are crucially no longer isotropic. Successfully recovering $h^\star$ via this method requires leveraging non-isotropic random matrix theory to establish sufficient conditions for the system to attain full column rank, which includes analyzing whether the extraneous columns (the elements not in $\mathrm{supp}(\widehat{h^\star})$) introduce problematic linear dependencies. Nevertheless, this approach holds significant promise given the robust existing literature on random matrix relaxations.

We anticipate that the structural information provided by $\mathcal{F}$ is sufficient to enable efficient classical recovery, most likely via the linear system formulation. We formalize this expectation with the following conjecture regarding the overall sample complexity of a smoothed quantum exact learner.

Conjecture 4.2.2. *Let $k \geq 2$ and let $0 \leq \xi \leq \frac{4^{-\lfloor \log k \rfloor}}{1+4^{-\lfloor \log k \rfloor}}$, so that $\sigma := 1 - \xi$ satisfies $\xi \leq \sigma \cdot 4^{-\lfloor \log k \rfloor}$. The hypothesis class of k-Fourier-sparse Boolean functions $\mathcal{H}^{\mathrm{FS}}$ is learnable by a $(0, \frac{1}{3})$-quantum PAC learner under σ-smooth distributions (Definition 4.1.3) with sample complexity $\mathcal{O}(\mathrm{poly}(k, \frac{1}{\sigma}))$.*

Having elucidated the near-uniform regime in the batch setting, we now investigate smoothed learning in the online setting, when the adversary is similarly restricted to σ-smooth distributions.

4.3 Smoothed Online Learning with Quantum Examples

In Section 3.3, we introduced a model for online learning a hypothesis class $\mathcal{H} \subseteq \{-1, 1\}^{\mathcal{X}}$ from quantum examples. The T-round adversarial game proceeded as follows, where, at the t-th round:

1. Learner provides a hypothesis $h_t : \{0,1\}^n \to \{-1, 1\}$.
2. Adversary reveals an example $|\psi_t\rangle = \sum_{x \in \{0,1\}^n} \sqrt{D_t(x)}\, |x, h^\star(x)\rangle$ for some[5] $D_t : \mathcal{X} \to [0, 1]$ and a fixed $h^\star \in \mathcal{H}$.
3. Learner incurs loss $\mathcal{L}_t := \mathbb{P}_{x \sim D_t}(h_t(x) \neq h^\star(x))$.

When the adversary is unrestricted in its choice of distributions D_t, we saw in Section 3.3 that this game essentially reduces to the canonical classical online learning setting characterized by the Littlestone dimension. In particular, the adversary's ability to concentrate all probability mass on a single point allows it to effectively *dequantize* the learning problem.

To prevent this degenerate behavior, we investigate a smoothed variant of the game where, at every round t, the adversary must choose a σ-smooth distribution D_t (Definition 4.1.1). This follows the analogous classical line of work on smoothed online learning [Hag18, HRS20, BDGR22, HRS24]. A related smoothed framework for online learning of quantum states was recently introduced in [MAGR25]; our setting instead concerns learning classical function classes from quantum examples.

The smoothing parameter σ elegantly interpolates between the adversarial and statistical regimes, providing a fascinating lens for quantum advantage. As $\sigma \to 0$, the adversary regains the power to deploy point-mass distributions, a regime where Theorems 3.3.4 and 3.3.5 preclude any quantum-classical separation in expected loss. Conversely, when $\sigma = 1$, the adversary is forced to play the

[5]Note that D_t may depend in any way on the adversary's past choices of distributions D_s, as well as the learner's past prediction hypotheses h_s, for all $s < t$.

reference distribution μ at every round. Fixing μ to be the uniform distribution and letting $\mathcal{H}$ be the class of k-Fourier-sparse functions $\mathcal{H}^{\mathrm{FS}}$ effectively recovers the batch setting of [ACL+21], which exhibits an unbounded quantum advantage (formalized below in Theorem 4.3.1). Motivated by the promise of this advantage persisting in the nontrivial near-uniform regime (Theorem 4.2.1 and Conjecture 4.2.2), we fix μ to be the uniform distribution for the remainder of this section.

4.3.1 Smoothed Online Learning of Fourier-Sparse Boolean Functions

Theorem 4.3.1 (Upper bound on the total expected loss for smoothed quantum online learning of $\mathcal{H}^{\mathrm{FS}}$ when $\sigma = 1$)**.** *Let $\mathcal{H}^{\mathrm{FS}}$ be the hypothesis class of k-Fourier-sparse Boolean functions. There exists a quantum online learner for $\mathcal{H}^{\mathrm{FS}}$ under 1-smooth "adversaries" that incurs total expected loss of at most $\mathcal{O}(k^2 \log k)$.*

Proof. As $\sigma = 1$, the adversary is forced to play the uniform distribution at every round. The learner can therefore sequentially simulate the exact batch learning algorithm of Theorem 4.1.8. At each round t, the learner predicts using an arbitrary consistent hypothesis $h_t \in \mathcal{H}^{\mathrm{FS}}$, incurs a loss of at most 1, and then consumes the received quantum example $|\psi_t\rangle$ to advance the batch learning procedure. By Theorem 4.1.8, after an expected[6] $\mathcal{O}(k^2 \log k)$ rounds, the learner is able to identify $h^\star$ exactly and predict perfectly thereafter. Consequently, the total expected loss is at most $\mathcal{O}(k^2 \log k)$. ■

We now extend our arguments to the near-uniform regime. We highlight two key observations regarding how the batch learning phases adapt when an adversary plays an adaptive sequence of σ-smooth distributions D_t.

- **Robustness of the Fourier sampling procedure.** By the analysis of Theorem 4.2.1, the probability of observing any Fourier support element $S_i \in \mathrm{supp}(\widehat{h^\star})$ is bounded below by $\frac{\sigma}{k^2}$. Crucially, this bound depends only on the smoothness parameter σ and *holds uniformly* over all σ-smooth distributions. Therefore, the adversary's freedom to adaptively vary D_t does not impede the Fourier sampling process. The learner can sequentially apply the Fourier sampling procedure, successfully identifying a collection $\mathcal{F} \supseteq \mathrm{supp}(\widehat{h^\star})$ in an expected $\mathcal{O}\left(\frac{k^2 \log k}{\sigma}\right)$ rounds.
- **Classical recovery under adaptive D_t.** To predict perfectly in subsequent rounds, the learner must exactly identify $h^\star$. While solving a linear system analogous to the batch setting provides a natural approach, other classical recovery techniques may also be viable. The unique challenge for any such approach in the online setting is that the classical examples

[6]While Theorem 4.1.8 guarantees the exact recovery of $h^\star$ with some constant probability, we can boost the success probability to $1 - \delta$ by standard repetition. Therefore, the expected sample complexity retains the same asymptotic order.

$(x_t, h^\star(x_t))$ are now drawn from *different*, adaptively changing, distributions D_t. Nevertheless, the overarching σ-smoothness constraint guarantees that these distributions remain globally bounded. We expect this restriction is sufficient to enable exact recovery within an n-independent number of rounds, whether by ensuring a linear system attains full column rank or by satisfying the requirements of an alternative robust classical algorithm.

Combining these observations, we anticipate that the structural guarantees of σ-smoothness permit efficient exact recovery of $h^\star$, bounding the total expected loss of a smoothed quantum online learner. We formalize this expectation with the following conjecture.

Conjecture 4.3.2. *Let $k \geq 2$ and let $0 \leq \xi \leq \frac{4^{-\lfloor \log k \rfloor}}{1+4^{-\lfloor \log k \rfloor}}$, so that $\sigma := 1 - \xi$ satisfies $\xi \leq \sigma \cdot 4^{-\lfloor \log k \rfloor}$. There exists a quantum online learner for $\mathcal{H}^{\mathrm{FS}}$ under σ-smooth adversaries that achieves a total expected loss bounded by $\mathcal{O}(\mathrm{poly}(k, \frac{1}{\sigma}))$.*

Chapter Summary. In this chapter, we introduced a framework for smoothed learning from quantum examples and studied both the batch and online learning settings under σ-smooth distributions. For the class of k-Fourier-sparse Boolean functions, we established preliminary results demonstrating the robustness of quantum Fourier sampling under σ-smooth distributions within a nontrivial near-uniform regime. In particular, we showed that quantum Fourier sampling allowed us to recover a collection $\mathcal{F} \supseteq \mathrm{supp}(\widehat{h^\star})$ with sample complexity depending only on k and σ (i.e., independent of n, as in the uniform case). Crucially, because our lower bound on successfully Fourier sampling an element in $\mathrm{supp}(\widehat{h^\star})$ held uniformly for any fixed σ-smooth class of distributions, we observed that the near-uniform robustness established in the batch setting naturally extended to the online framework, remaining unaffected even when an adversary could adaptively vary the σ-smooth distribution across rounds.

Beyond the core sampling mechanics, we analyzed three distinct classical avenues for sample-efficiently recovering the exact hypothesis $h^\star$ under the uniform distribution, once its Fourier support $\mathrm{supp}(\widehat{h^\star})$ has been quantumly identified. Then, we evaluated the viability of each of these approaches in the σ-smooth setting. Based on our observations and sampling guarantees, we formally conjectured that the unbounded quantum advantage for learning k-Fourier-sparse Boolean functions, a separation previously established only under the uniform distribution, persists throughout the identified near-uniform regime (4.11) for both batch and online learning.

Determining the smoothed learning rates for k-Fourier-sparse Boolean functions outside the regime of (4.11) remains an entirely open question. In particular, it is unclear whether the conjectured quantum advantage decays gradually beyond the near-uniform regime, or whether a sharp transition in the sample complexity and total expected loss bounds occurs as σ decreases.

CHAPTER 5

Conclusion

5.1 Summary of Contributions

This thesis developed the theory of learning with quantum examples along three interconnected directions: multiclass learning, online learning, and smoothed learning.

In the multiclass setting (Chapter 2), we extended quantum PAC learning beyond Boolean concept classes and established upper and lower bounds on the quantum sample complexity of batch multiclass learning in both the realizable and agnostic settings. Our results reinforce that, as in the Boolean setting, quantum examples do not yield quantum-classical sample complexity separations in distribution-independent (i.e., worst-case) PAC learning, with the resulting learning rates governed by the Natarajan dimension up to logarithmic factors in the label-space size.

In the online setting (Chapter 3), we introduced a framework for learning classical function classes from quantum examples under adversarially chosen distributions. As an intermediate step, we developed a classical *adversary-provides-a-distribution* model that naturally bridged classical and quantum online learning. Within these frameworks, we established expected regret guarantees for binary and multiclass classification in both the realizable and agnostic settings, with the key takeaway that *unrestricted* adversarial power permits highly concentrated distributions that can effectively dequantize the learning problem.

Motivated by these observations, Chapter 4 introduced a framework for smoothed learning with quantum examples based on σ-smooth distributions. Focusing on the class of k-Fourier-sparse Boolean functions, we established the robustness of quantum Fourier sampling throughout a nontrivial near-uniform regime in the batch setting, showing that a collection containing the Fourier support of the target function can be recovered with sample complexity depending only on k and σ. Crucially, because the lower bound governing the success probability of the Fourier sampling procedure held uniformly over any fixed σ-smooth class of distributions, we observed that this robustness extends naturally to the online framework, remaining unaffected even when an adversary adaptively varies the example distribution across rounds. We further investigated several classical approaches for recovering the target function from the information obtained via quantum Fourier

sampling and found that each requires additional nontrivial ingredients in the smoothed setting. These observations led us to formulate formal conjectures suggesting that the unbounded quantum advantage known under the uniform distribution persists throughout the identified near-uniform regime in both the batch and online settings.

5.2 Open Problems and Future Directions

The results of this thesis leave several natural directions open for future work, both in understanding the limitations of learning with quantum examples and in identifying settings where genuinely quantum advantages persist.

Multiclass Learning. A first direction concerns the characterization of quantum multiclass learnability when the label space is unbounded. In Chapter 2, our multiclass sample complexity bounds were governed by the Natarajan dimension up to logarithmic factors in the number of labels, k. While this matches the known finite-label classical theory, it does not yield a full characterization in the regime $k \to \infty$. In the classical setting, this gap was resolved through the Daniely–Shalev-Schwartz (DS) dimension [DSS14, BCD$^+$22], which characterizes multiclass PAC learnability even for unbounded label spaces. Although the corresponding classical upper bounds transfer directly to the quantum setting via measure-and-learn-classically reductions, the known DS-dimension lower-bound arguments rely on transductive learning techniques that do not appear to admit a straightforward quantum analogue (cf. Section 2.2.3 for more details), since quantum examples encode inputs and labels jointly in superposition. Determining the tight quantum sample complexity of multiclass learning for unbounded label spaces therefore remains open.

Online Learning. A second direction concerns the role of genuinely quantum strategies in online learning. In Chapter 3, nearly all of our optimal upper bounds were achieved by learners that simply measured the quantum example and subsequently learned classically. Together with the matching lower bounds, this suggests that unrestricted adversarial power can largely eliminate the benefits of quantum examples in online learning. The sole remaining gap appears in the agnostic multiclass setting, where the best known classical regret bound removes the residual dependence on the number of labels k [HMR$^+$23], whereas our corresponding quantum upper bound retains a $\sqrt{\log k}$ factor. Closing this gap may therefore require techniques beyond the current measure-and-learn-classically framework, potentially involving genuinely quantum online learning strategies.

Smoothed Learning. Finally, Chapter 4 leaves several open questions regarding smoothed learning with quantum examples. As a natural starting point, Conjectures 4.2.2 and 4.3.2 remain open. Resolving either conjecture would require understanding whether the exact target function can be

recovered efficiently from the information obtained via quantum Fourier sampling under σ-smooth distributions. Moreover, these conjectures rely on the robustness of the quantum Fourier sampling procedure established only within the near-uniform regime identified in (4.11). Beyond this regime, the exact learning complexity remains entirely open. In particular, it is unknown whether the conjectured quantum advantages decay gradually as σ decreases, or whether a sharp transition occurs in the sample complexity of batch learning and the total expected loss guarantees of online learning as σ approaches the unrestricted (resp. fully adversarial) setting.

More broadly, it would be interesting to understand which structural properties of a hypothesis class are necessary and sufficient for quantum advantages to survive under smoothness constraints. The k-Fourier-sparse setting benefited from an explicit Fourier structure that enabled a detailed analysis of the quantum Fourier sampling procedure, including uniform lower bounds on its success probability throughout the near-uniform regime identified in (4.11). It remains an open question whether analogous arguments can be developed for hypothesis classes lacking such explicit Fourier structure. Identifying a general smoothed quantum learnability criterion, analogous to the way the Littlestone dimension characterizes adversarial classical online learnability, would be a natural next step. Finally, an intriguing broader direction is to investigate whether these smoothed learning regimes admit formal connections to foundational quantum supremacy paradigms based on Fourier sampling [FU15, AC16, HE23].

5.3 Closing Remarks

The results of this thesis suggest that quantum learning advantages are neither universally present nor inherently fragile, but instead emerge from a subtle interplay between the hypothesis class, the example distribution, and the degree of adversarial control permitted in the learning process. While fully adversarial and highly concentrated distributions can effectively *dequantize* a learning problem, appropriately structured distributional assumptions can nevertheless preserve genuinely quantum learning behavior, particularly for hypothesis classes whose internal structure aligns naturally with the information accessible through quantum examples.

This thesis has mapped the first coordinates of that landscape. The terrain ahead is wide open.

Bibliography

[AA15] Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 307–316, 2015.

[AA24] Anurag Anshu and Srinivasan Arunachalam. A survey on the complexity of learning quantum states. *Nature Reviews Physics*, 6(1):59–69, 2024.

[AAA+25] David Aasen, Morteza Aghaee, Zulfi Alam, Mariusz Andrzejczuk, Andrey Antipov, Mikhail Astafev, Lukas Avilovas, Amin Barzegar, Bela Bauer, Jonathan Becker, et al. Roadmap to fault tolerant quantum computation using topological qubit arrays. *arXiv preprint arXiv:2502.12252*, 2025.

[ABC+20] Srinivasan Arunachalam, Aleksandrs Belovs, Andrew M. Childs, Robin Kothari, Ansis Rosmanis, and Ronald de Wolf. Quantum Coupon Collector. In *Proceedings of the 15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*, pages 10:1–10:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

[ABED+21] Noga Alon, Omri Ben-Eliezer, Yuval Dagan, Shay Moran, Moni Naor, and Eylon Yogev. Adversarial Laws of Large Numbers and Optimal Regret in Online Classification. In *Proceedings of the 53rd ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 447–455. ACM, 2021.

[AC16] Scott Aaronson and Lijie Chen. Complexity-theoretic foundations of quantum supremacy experiments. *arXiv preprint arXiv:1612.05903*, 2016.

[ACH+19] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. Online learning of quantum states. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124019, 2019.

[ACL+21] Srinivasan Arunachalam, Sourav Chakraborty, Troy Lee, Manaswi Paraashar, and Ronald de Wolf. Two new results about quantum exact learning. *Quantum*, 5:587, 2021.

[AdW17] Srinivasan Arunachalam and Ronald de Wolf. Guest Column: A Survey of Quantum Learning Theory. *ACM SIGACT News*, 48(2):41–67, 2017.

[AdW18] Srinivasan Arunachalam and Ronald de Wolf. Optimal Quantum Sample Complexity of Learning Algorithms. *Journal of Machine Learning Research*, 19(71):1–36, 2018.

[AQS22] Srinivasan Arunachalam, Yihui Quek, and John Smolin. Private and online quantum learning. In *Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS)*, 2022. arXiv:2103.00778.

[AS05] Alp Atıcı and Rocco A. Servedio. Improved Bounds on Quantum Learning Algorithms. *Quantum Information Processing*, 4(5):355–386, 2005.

[AS07] Alp Atıcı and Rocco A. Servedio. Quantum algorithms for learning and testing juntas. *Quantum Information Processing*, 6(5):323–348, 2007.

[BA22] Nader H. Bshouty and Hasan Anis. Efficient quantum agnostic improper learning of decision trees, 2022. arXiv:2210.00212.

[BCD+22] Nataly Brukhim, Daniel Carmon, Irit Dinur, Shay Moran, and Amir Yehudayoff. A Characterization of Multiclass Learnability. In *Proceedings of the 63rd Symposium on Foundations of Computer Science (FOCS)*, pages 943–955. IEEE, 2022.

[BDCBHL95] Shai Ben-David, Nicolo Cesa-Bianchi, David Haussler, and Philip M. Long. Characterizations of Learnability for Classes of $\{0, ..., n\}$-Valued Functions. *Journal of Computer and System Sciences*, 50(1):74–86, 1995.

[BDGR22] Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. Smoothed Online Learning is as Easy as Statistical Learning. In *Proceedings of the 35th Conference on Learning Theory (COLT)*, pages 1716–1786. PMLR, 2022.

[BDPSS09] Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic Online Learning. In *Proceedings of the 22nd Conference on Learning Theory (COLT)*, 2009.

[BEHW89] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM*, 36(4):929–965, 1989.

[BHMZ20] Olivier Bousquet, Steve Hanneke, Shay Moran, and Nikita Zhivotovskiy. Proper Learning, Helly Number, and an Optimal SVM Bound. In *Proceedings of the 33rd Conference on Learning Theory (COLT)*, pages 582–609. PMLR, 2020.

[BJ95] Nader H. Bshouty and Jeffrey C. Jackson. Learning DNF over the uniform distribution using a quantum example oracle. In *Proceedings of the 8th Conference on Computational Learning Theory (COLT)*, pages 118–127. ACM, 1995.

[BLL+11] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual Bandit Algorithms with Supervised Learning Guarantees. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 19–26. PMLR, 2011.

[BV93] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 11–20, 1993.

[Car20] Matthias C. Caro. Quantum learning Boolean linear functions w.r.t. product distributions. *Quantum Information Processing*, 19(6):1–33, 2020.

[CEH+25] Alon Cohen, Liad Erez, Steve Hanneke, Tomer Koren, Yishay Mansour, Shay Moran, and Qian Zhang. Sample complexity of agnostic multiclass classification: Natarajan dimension strikes back. *arXiv preprint arXiv:2511.12659*, 2025.

[CSS15] Andrew W. Cross, Graeme Smith, and John A. Smolin. Quantum learning robust to noise. *Physical Review A*, 92(1):012327, 2015. arXiv:1407.5088.

[CXK+26] Madelyn Cain, Qian Xu, Robbie King, Lewis RB Picard, Harry Levine, Manuel Endres, John Preskill, Hsin-Yuan Huang, and Dolev Bluvstein. Shor's algorithm is possible with as few as 10,000 reconfigurable atomic qubits. *arXiv preprint arXiv:2603.28627*, 2026.

[DSBDSS15] Amit Daniely, Sivan Sabato, Shai Ben-David, and Shai Shalev-Shwartz. Multiclass Learnability and the ERM Principle. *Journal of Machine Learning Research*, 16(72):2377–2404, 2015.

[DSS14] Amit Daniely and Shai Shalev-Shwartz. Optimal Learners for Multiclass Problems. In *Proceedings of the 27th Conference on Learning Theory (COLT)*, pages 287–316. PMLR, 2014.

[DT22] Philip Dawid and Ambuj Tewari. On Learnability under General Stochastic Processes. *Harvard Data Science Review* , 4(4), 2022.

[FU15] Bill Fefferman and Chris Umans. The power of quantum fourier sampling. *arXiv preprint arXiv:1507.05592*, 2015.

[GKZ19] Alex B. Grilo, Iordanis Kerenidis, and Timo Zijlstra. Learning-with-errors problem is easy with quantum samples. *Physical Review A*, 99(3):032314, 2019. arXiv:1702.08255.

[Goo25] Google Quantum AI and Collaborators. Quantum error correction below the surface code threshold. *Nature*, 638(8052):920–926, 2025.

[GOS+11] Parikshit Gopalan, Ryan O'Donnell, Rocco A Servedio, Amir Shpilka, and Karl Wimmer. Testing fourier dimensionality and sparsity. *SIAM Journal on Computing*, 40(4):1075–1100, 2011.

[GS25] Uma Girish and Rocco Servedio. Forrelation is extremally hard. *arXiv preprint arXiv:2508.02514*, 2025.

[Hag18] Nika Haghtalab. *Foundation of Machine Learning, by the People, for the People*. PhD thesis, Microsoft Research, 2018.

[Han16] Steve Hanneke. The Optimal Sample Complexity of PAC Learning. *Journal of Machine Learning Research*, 17(38):1–15, 2016.

[HE23] Dominik Hangleiter and Jens Eisert. Computational advantage of quantum random sampling. *Reviews of Modern Physics*, 95(3):035001, 2023.

[HHM+23] Larry Huynh, Jin Hong, Ajmal Mian, Hajime Suzuki, Yanqiu Wu, and Seyit Camtepe. Quantum-inspired machine learning: a survey. *arXiv preprint arXiv:2308.11269*, 2023.

[HMR+23] Steve Hanneke, Shay Moran, Vinod Raman, Unique Subedi, and Ambuj Tewari. Multiclass Online Learning and Uniform Convergence. In *Proceedings of 36th Conference on Learning Theory (COLT)*, pages 5682–5696. PMLR, 2023.

[HNS24] Shima Bab Hadiashar, Ashwin Nayak, and Pulkit Sinha. Optimal Lower Bounds for Quantum Learning via Information Theory. *IEEE Transactions on Information Theory*, 70(3):1876–1896, 2024.

[HR16] Ishay Haviv and Oded Regev. The list-decoding size of fourier-sparse boolean functions. *ACM Transactions on Computation Theory (TOCT)*, 8(3):1–14, 2016.

[HRS20] Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed Analysis of Online and Differentially Private Learning. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, pages 9203–9215. Curran Associates, Inc., 2020.

[HRS24] Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *Journal of the ACM*, 71(3):1–34, 2024.

[HSL+25] AC Hughes, R Srinivas, CM Löschnauer, HM Knaack, R Matt, CJ Ballance, M Malinowski, TP Harty, and RT Sutherland. Trapped-ion two-qubit gates with> 99.99% fidelity without ground-state cooling. *arXiv preprint arXiv:2510.17286*, 2025.

[IBM26] IBM Technology Atlas. Quantum Roadmap. `https://www.ibm.com/roadmaps/quantum/`, 2026. Accessed: 2026-05-16.

[JTY02] Jeffrey C. Jackson, Christino Tamon, and Tomoyuki Yamakami. Quantum DNF learnability revisited. In *Proceedings of the 8th Annual International Computing and Combinatorics Conference (COCOON)*, volume 2387 of *Lecture Notes in Computer Science*, pages 595–604. Springer, 2002. arXiv:quant-ph/0202066.

[KRS19] Varun Kanade, Andrea Rocchetto, and Simone Severini. Learning DNFs under product distributions via μ-biased quantum Fourier sampling. *Quantum Information and Computation*, 19(15–16):1261–1278, 2019. arXiv:1802.05690.

[KSS92] Michael J. Kearns, Robert E. Schapire, and Linda M. Sellie. Toward efficient agnostic learning. In *Proceedings of the 5th Workshop on Computational Learning Theory (COLT)*, pages 341–352. ACM, 1992.

[Lit88] Nick Littlestone. Learning Quickly When Irrelevant Attributes Abound: A New Linear-Threshold Algorithm. *Machine Learning*, 2(4):285–318, 1988.

[LN24] Gábor Lugosi and Gergely Neu. Online-to-PAC Conversions: Generalization Bounds via Regret Analysis, 2024.

[MAGR25] Maxime Meyer, Soumik Adhikary, Naixu Guo, and Patrick Rebentrost. Online learning of pure states is as hard as mixed states. *arXiv preprint arXiv:2502.00823*, 2025.

[Mon12] Ashley Montanaro. The quantum query complexity of learning multilinear polynomials. *Information Processing Letters*, 112(11):438–442, 2012.

[MT25] Preetham Mohan and Ambuj Tewari. Quantum learning theory beyond batch binary classification. *Quantum*, 9:1813, 2025.

[Nat89] Balas K. Natarajan. On Learning Sets and Functions. *Machine Learning*, 4(1):67–97, 1989.

[NC10] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010.

[NS24] Ashwin Nayak and Pulkit Sinha. Proper vs Improper Quantum PAC Learning, 2024.

[O'D14] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[QAS21] Yihui Quek, Srinivasan Arunachalam, and John A. Smolin. Private learning implies quantum stability. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NeurIPS)*, pages 20503–20515, 2021.

[RST15] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Sequential complexities and uniform martingale laws of large numbers. *Probability Theory and Related Fields*, 161(1):111–153, 2015.

[San19] Swagato Sanyal. Fourier sparsity and dimension. *Theory of Computing*, 15(1):1–13, 2019.

[SG04] Rocco A. Servedio and Steven J. Gortler. Equivalences and Separations Between Quantum and Classical Learnability. *SIAM Journal on Computing*, 33(5):1067–1092, 2004.

[Sho99] Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

[SSBD14] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.

[SSG24] Wilfred Salmon, Sergii Strelchuk, and Tom Gur. Provable Advantage in Quantum PAC Learning. In *Proceedings of the 37th Conference on Learning Theory (COLT)*, pages 4487–4510. PMLR, 2024.

[SSHE21] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert. On the quantum versus classical learnability of discrete distributions. *Quantum*, 5:417, 2021.

[Tal94] Michel Talagrand. Sharper Bounds for Gaussian and Empirical Processes. *The Annals of Probability*, 22(1):28–76, 1994.

[Tan19] Ewin Tang. A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st annual ACM SIGACT symposium on theory of computing*, pages 217–228, 2019.

[Tan22] Ewin Tang. Dequantizing algorithms to understand quantum advantage in machine learning. *Nature Reviews Physics*, 4(11):692–693, 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. Cambridge university press, 2018.

[Zha10] Chi Zhang. An improved lower bound on query complexity for quantum PAC learning. *Information Processing Letters*, 111(1):40–45, 2010.