Preface

We begin our journey by motivating our studies in two different ways, one more pragmatic and one more philosophical.

Machine learning, and in particular its recent manifestation in deep learning in the last two decades, has been transformative for computer science and information technology. The promise, perils, and possibility of generative artificial intelligence have seeped from Silicon Valley to the public discourse, and the ultimate contours of its potential are the subject of intense speculation. Granted all of the recent developments in contemporary machine learning, many of the core ideas derive from *statistical learning theory*, which had its heyday in the 1990's and early 2000's. This is a rigorous mathematical subject which conceives of learning in a probabilistic and often Bayesian manner, drawing on probability theory and empirical process theory, while utilizing information-theoretic concepts from Shannon's foundational work. Since contemporary machine learning is mostly an *empirical* subject pertaining to extraordinarily sophisticated statistical models which defy comprehensive characterization, the particularities of the theorems developed in statistical learning theory are not often used; however, the intuitions these rigorous results provide are essential for designing new neural network architectures, loss functions, training algorithms, and datasets. As such, the afterlife of statistical learning theory is that its quantitative knowledge in mathematically simple settings has been lifted to qualitative but indispensable wisdom about highly complex systems.

Then one motivation for our studies is to develop a quantum version of statistical learning theory (or more succinctly, quantum learning theory), suitable for future application by quantum computers. Our studies will focus on quantum learning for *quantum data* as opposed to classical data, for reasons that will be explained. (Indeed, the latter setting is very interesting but has a somewhat different character.) The subject will necessarily be mathematically rigorous to cement our understanding of quantum data and quantum learning algorithms, as well as to develop robust methods with provable performance guarantees suitable for scientific applications. We emphasize that at this moment in time, quantum learning theory is not chiefly an empirical subject such as contemporary machine learning; this underscores the necessity of mathematical rigor and the importance of the foundational development of basic quantum learning algorithms and methods that future theoretical or empirical inquiries may build on. We will focus on developments in quantum learning theory mostly from 2019 onward, which saw the development of fruitful foundations and applications of the subject.

There is also a second, more philosophical motivation for our study of quantum learning theory. *Epistemology* is the philosophical study of what we can know about the world, and how we come to know it. One of the earlier treatments of the subject goes back over 2000 years to Plato, although among scientists Descartes' maxim

"I think therefore I am" (*cogito ergo sum*) may be more familiar. Specifically, Descartes was concerned with what he could know with certainty about the world, and upon wrestling with various uncertainties he concludes that he knows at the very least that he himself exists, since for him to even render the thought requires his own existence.

A persistent thread in epistemology since the beginning is that there may be aspects of our reality that we can never come to know. A particularly incisive analysis along these lines was developed by Immanuel Kant in the late 18th century, in which he detailed how the physicality of our corporeal beings and the constitution of our minds place a priori fundamental limitations on what we can know about the world, leaving certain truths necessarily out of our reach. While this premise is widely accepted by philosophers, it is often frowned upon by scientists; after all, we are children of the Enlightenment for which scientific knowledge is infinitely extensible and far-reaching. If you feel such an urge to frown on epistemology, consider this more modern example: we live in a universe which is expanding and accelerating. Eventually, the expansion will be so fast that light from the early universe will become so redshifted as to be undetectable. As such, if there is life that develops somewhere in the universe at such a time, they will never be able to empirically determine that there was a Big Bang. Thus a truth about the universe is, to them, out of reach.

In the early 20th century, David Hilbert famously declared that the *mathematical* world was fundamentally knowable, and that every precise mathematical statement was either true or false. This epistemic totalism was shockingly undermined by Kurt Gödel in 1931, when he showed that there must always exist mathematical statements which can neither be proved true nor false. This death blow to Hilbert's (and Bertrand Russell's) conception of mathematical knowledge was concretized by Alan Turing in his foundational work on computer science, the pragmatic heir to mathematical logic. Turing famously showed in 1936 that there is no algorithm (which is guaranteed to terminate in finite time) that can conclusively decide if any given algorithm will halt or not. Thus Turing's theory of undecidability cleaves out facts about the world which are fundamentally unknowable to us, furnishing totally precise examples of epistemic roadblocks.

Decades later starting in the early 1970's, the subject of *computational complexity* began to emerge. Instead of being concerned with whether the solution to a computational problem was knowable or unknowable, the subject focused on the *difficulty* of computational problems. For example, one can show that sorting a list of n items (on a classical computer) requires *at most* $\sim n \log n$ computational steps, but also *at least* $\sim n \log n$ computational steps, thus pinpointing the absolute difficulty of the problem. Some computational problems have polynomial difficulty whereas others have exponential difficulty, and are stratified according to *computational complexity classes*. In this way, computational complexity theory comprises a quantitative form of epistemology, circumscribing how difficult it is to obtain computational knowledge.

Having set the scene, we turn to a deep question: how do we come to learn about the world through scientific inquiry? A key facet is that we interrogate the natural world through experiment, and algorithmically process our collected data to reveal hitherto unknown properties of nature. More formally, we can conceptualize a system in nature – such as a superconductor, a vat of chemicals, a biological organism, etc. – as a source of *data* which is not fully characterized (or else we would not need to run the experiment); then our experiment comprises a series of interactions with the world to sample data, subsequent computational processing of the data, and possibly additional interactions with the world predicated on the processing of previous data. The ultimate outcome is that we learn a property of the world, such as the charge of the electron, the symmetry of a crystal, etc. In this manner, we see that scientific experiments can be beautifully and precisely abstracted into the framework of learning theory. Therefore, a quantitative study of learning theory can reveal what we can fundamentally come to know about the world, and how difficult it is to do so.

Since the laws of nature are quantum-mechanical, any theory of learning the natural world must take quantum mechanics into account. In particular, the natural systems we seek to understand may be quantum-mechanical; the data we extract can be quantum-mechanical; and our means of processing that data can be quantum-mechanical. Thus we necessitate a quantum theory of learning. Such a theory reveals that there are facets of the natural world which are inaccessible to us unless we can harness quantum computers to couple to natural systems and perform quantum information processing. More bluntly, quantum learning gives us access to properties of the natural world which are otherwise unknowable by classical means. And yet the same theory shows us which properties of the natural world are forever out of reach, even with the aid of vast quantum computational power.

Quantum learning theory circumscribes what is knowable and unknowable about the natural world, providing a quantitative epistemology of the grasp of scientific inquiry. With so much at stake, let us begin.