Proper versus Improper Quantum PAC Learning

Ashwin Nayak [*]
University of Waterloo

Pulkit Sinha [†]
University of Waterloo

June 5, 2024

Abstract

A basic question in the PAC model of learning is whether *proper* learning is harder than *improper* learning. In the classical case, the answer to this question, with respect to sample complexity, is known to depend on the concept class. While there are concept classes for which the two modes of learning have the same complexity, there are examples of concept classes with VC dimension d that have sample complexity $\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right)$ for proper learning with error ϵ, while the complexity for improper learning is $\mathrm{O}\left(\frac{d}{\epsilon}\right)$. One such example arises from the Coupon Collector problem.

Motivated by the efficiency of proper versus improper learning with *quantum* samples, Arunachalam, Belovs, Childs, Kothari, Rosmanis, and de Wolf [2] studied an analogue, the Quantum Coupon Collector problem. Curiously, they discovered that for learning size k subsets of $[n]$ the problem has sample complexity $\Theta(k\log\min\{k, n-k+1\})$, in contrast with the complexity of $\Theta(k\log k)$ for Coupon Collector. This effectively negates the possibility of a separation between the two modes of learning via the quantum problem, and Arunachalam *et al.* posed the possibility of such a separation as an open question.

In this work, we first present an algorithm for the Quantum Coupon Collector problem with sample complexity that matches the sharper lower bound of $(1-\mathrm{o}_k(1))k\ln\min\{k, n-k+1\}$ shown recently by Bab Hadiashar, Nayak, and Sinha [8], for the entire range of the parameter k. Next, we devise a variant of the problem, the Quantum *Padded* Coupon Collector. We prove that its sample complexity matches that of the classical Coupon Collector problem for both modes of learning, thereby exhibiting the same asymptotic separation between proper and improper quantum learning as mentioned above. The techniques we develop in the process can be directly applied to any form of padded quantum data. We hope that padding can more generally lift other forms of classical learning behaviour to the quantum setting.

1 Introduction

A fundamental task in machine learning is to predict certain properties of objects, given access to labeled data for other objects drawn from the same distribution. Problems like these form the basis of supervised machine learning. Setting aside the algorithmic aspects of the problem, a natural question to ask is: how many labeled data are required to achieve a certain degree of accuracy in the prediction? To study this formally, Valiant [14] developed the *probably approximately correct* (PAC) model.

[*]Department of Combinatorics and Optimization, and Institute for Quantum Computing, University of Waterloo, 200 University Ave. W., Waterloo, ON, N2L 3G1, Canada. Email: academic@ashwinnayak.info .

[†]School of Computer Science, and Institute for Quantum Computing, University of Waterloo, 200 University Ave. W., Waterloo, ON, N2L 3G1, Canada. Email: psinha@uwaterloo.ca .

In the PAC model, binary properties are thought of as Boolean functions, mapping each object in the population to a 0 or 1 corresponding to whether they have the property or not. Depending on the learning task at hand, we may have some prior knowledge about the Boolean functions that may be allowed. Each Boolean function is called a *concept*, and the set of allowed functions, a *concept class*. Given labeled data according to an unknown concept in the concept class and an unknown distribution on the population, the goal is to reconstruct a concept (with probability at least $1 - \delta$) which on the same distribution matches with the unknown concept with probability at least $(1 - \epsilon)$. The number of data needed is formalised as the sample complexity, which is simply the number of labeled samples. This is called the (ϵ, δ)-PAC sample complexity of the learning task.

Formalised this way, the hardness of the learning task only depends on the concept class. As shown in a series of works [5, 9], this hardness is almost fully characterized by the *VC dimension* d of the concept class, and the sample complexity turns out to be

$$\Theta\left(\frac{d}{\epsilon} + \frac{\log(1/\delta)}{\epsilon}\right) \quad . \tag{1.1}$$

The PAC learning model was extended to the quantum setting by Bshouty and Jackson [6], where random samples of labeled data were replaced with *quantum* samples. In quantum samples, the data points are in a superposition, weighted with the square root of the corresponding probability. Learning using these samples is readily seen to be at least as easy as with classical samples, as we can apply classical learning techniques on the distribution obtained after measuring in the standard basis. Hence the upper bound on sample complexity in Eq. (1.1) applies. Arunachalam and de Wolf [4] showed that the lower bound in (1.1) also holds in the quantum setting. (See also Ref. [8] for an arguably simpler, information-theoretic proof.)

Interestingly, optimal PAC learners are not guaranteed to output a concept within the concept class. The optimal learner, even though equipped with information about the possible actual classifications, can choose to not classify based on any of those. This kind of learning, where the output concept is not guaranteed to belong to the concept class, is called *improper learning*. If on the other hand, we enforce the condition that the PAC learner only output concepts from within the class, then the learning task is called *proper learning*.

Does the requirement of proper learning make the task harder? In the classical case, the answer to this question, with respect to sample complexity, is known to depend on the concept class. For the concept class containing all possible concepts, the answer is a clear no, as proper and improper learning are equivalent. On the other hand, for each d and ϵ, there are examples of concept classes that have sample complexity

$$\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon} + \frac{\log(1/\delta)}{\epsilon}\right) \tag{1.2}$$

for proper learning. These classes exhibit an asymptotic separation between the sample complexity of proper and improper PAC learning. An apparently folklore example of such concept classes arises from the *Coupon Collector* problem [10].

In the Coupon Collector problem, a classic problem from probability theory, we are given integers k, n with $1 < k < n$, and have access to independent, uniformly distributed elements from a fixed, unknown size-k subset S of $[n]$. The task is to draw enough samples so as to "collect" (i.e., observe) all k "coupons" (elements) in S. Viewed as a learning problem, the task is to determine the unknown set S. It is well-known that $k \ln k + \Theta(k)$ samples are necessary and sufficient to observe all k elements of S with constant probability of success (see, e.g., [12, Theorem 5.13]).

Taking $n \coloneqq \lceil 1/\epsilon \rceil$ and $k \coloneqq n - 1$ in the Coupon Collector problem, we get a concept class with VC-dimension 1, and hence the class is $(1/n, 1/4)$-PAC learnable with $\Theta(n)$ samples, albeit *improperly*. On

the other hand, *proper* PAC learning the class with the same parameters requires $\Theta(n \ln n)$ samples, as this entails correctly producing the size-$(n-1)$ subset with probability at least $3/4$. This gives us a separation of $\Omega((1/\epsilon)\log(1/\epsilon))$ versus $\mathrm{O}(1/\epsilon)$ for the sample complexity of proper versus improper learning. A generalization of this construction gives us a concept class with VC-dimension d which has sample complexity as stated in Eq. (1.2) [10].

Motivated by the efficiency of proper versus improper learning with *quantum* samples, Arunachalam, Belovs, Childs, Kothari, Rosmanis, and de Wolf [2] studied a quantum analogue of the Coupon Collector problem. In the *Quantum Coupon Collector* problem, the learner has access to quantum samples, i.e., uniform superpositions $|\psi_S\rangle$ over the elements of an unknown size-k subset S of $[n]$. The task is to learn S with high probability. Arunachalam *et al.* discovered a curious phenomenon—the quantum sample complexity of the problem matches that of classical coupon collection for k up to roughly $n/2$, but then *decreases* until k reaches n. More precisely, with $m \coloneqq n-k$, they proved that the sample complexity of the Quantum Coupon Collector problem with constant probability of error < 1 is $\Theta(k \log \min\{k, m+1\})$. Bab Hadiashar, Nayak, and Sinha [8] proved a sharper lower bound of $(1-\mathrm{o}_k(1))k \ln \min\{k, m+1\}$. More precisely, they established the following result.

Theorem 1.1 ([8], Theorem IV.17 and proof of Corollary IV.18)**.** *Let $\delta \in (0, 1/40]$ and $c_0 \coloneqq \frac{1}{2}\ln\left(\frac{1-\delta}{32\delta}\right)$. Any algorithm for the Quantum Coupon Collector problem with parameters n, k, and error probability at most δ has sample complexity*

- *at least $k \ln m + c_0 n$ when $1 \le m \le \delta n$ and $m \ln m \le c_0 n/20$, and*
- *at least $k \ln k - k \ln \ln k - \mathrm{O}(k)$ otherwise.*

The above bound has the exact optimal leading order term for k sufficiently smaller than n (listed as the second case in the theorem). Bab Hadiashar *et al.* conjectured that this sharp optimality of the bound extends to all k up to n.

First, we present an algorithm that confirms this conjecture.

Theorem 1.2. *There is an algorithm for the Quantum Coupon Collector problem with parameters n, k, where $n \ge 1$ and $1 < k < n$, that has probability of error at most $\delta \in (0, 1]$ and uses*

- *at most $k \ln m + k \ln \frac{\mathrm{e}}{\delta}$ samples when $3m \ln(\mathrm{e}\, m) \le n$, and*
- *at most $k \ln k + k \ln \frac{1}{\delta}$ samples otherwise.*

We may verify that the number of samples matches the lower bound in Theorem 1.1 exactly in the leading order term for all k, for constant error δ. Moreover, the second order term is also tight up to a factor of 2, as a function of k and δ, in the first case listed in Theorem 1.1.

When k is sufficiently small, i.e., in the "large" m regime, there is a straightforward algorithm. We measure the quantum samples and to obtain classical ones, and follow the classical learning scheme. The non-trivial case is when m is "small". In the small m regime, the previous best algorithm due to Arunachalam *et al.* attempts to learn the complement $\overline{S}$ rather than S. The algorithm repeatedly measures the quantum samples $|\psi_S\rangle$ according to the measurement $(|\psi_{[n]}\rangle\langle\psi_{[n]}|, \mathbb{I} - |\psi_{[n]}\rangle\langle\psi_{[n]}|)$, where $|\psi_{[n]}\rangle$ is the uniform superposition over all elements of $[n]$. The algorithm further measures the residual state in the computational basis when the second outcome is observed. As we may readily verify, the elements in $\overline{S}$ are considerably more likely to be observed. Thus it suffices to estimate the frequency with which elements occur to determine which ones belong to $\overline{S}$.

We begin with the basic idea underlying the above algorithm, and build on it. Instead of estimating frequencies, we take a more optimistic approach, in that we gamble on the element observed in the second measurement as being one of the more likely elements. I.e., we gamble on the element belonging to $\overline{S}$. Of course, the gamble may fail, and we may incorrectly add an element to our guess for $\overline{S}$. We devise a refined set of measurements to either identify and remove such "rogue" elements from our guess, or to discover other elements that are likely to be in $\overline{S}$. Perhaps surprisingly, the behaviour of the resulting random process resembles that of the classical Coupon Collector process, and we show that it converges sufficiently rapidly to the correct set.

Returning to the question of proper versus improper learning, consider the case $k = n - 1$. The sample complexity of the Quantum Coupon Collector problem is $\Theta(n)$. Thus, unlike in the classical case, the quantum version of the problem *does not* separate the sample complexity of proper from that of improper quantum PAC learning. Arunachalam *et al.* concluded their work by raising the question of the asymptotic equality of the complexity of learning in the two modes [2, Section 5].

The second result we present resolves this question.

Theorem 1.3. *For each $d \geq 5$ and $\epsilon < 1$, there are concept classes for which $\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon} + \frac{1}{\epsilon}\log\frac{1}{\delta}\right)$ quantum samples are required for proper (ϵ, δ)-PAC learning, while $\mathrm{O}\left(\frac{d}{\epsilon} + \frac{1}{\epsilon}\log\frac{1}{\delta}\right)$ quantum samples suffice for improper (ϵ, δ)-PAC learning.*

To establish this separation, we construct concept classes that are "padded" versions of those used to show the separation via the classical Coupon Collector problem. The *Padded Coupon Collector* problem is obtained by modifying the standard Coupon Collector problem by appending arbitrary values from a large set to each element of the set $[n]$. Its quantum version, the *Quantum Padded Coupon Collector* (QPCC) is defined analogously. The goal in both the cases is still to learn an unknown subset of $[n]$ as in the standard coupon collector problem, or more generally, to learn a large part of the subset. (See Section 2 for the precise definition of the problems.) In the classical case, the addition of padding has no real consequence. As the padding for any element does not carry any information about the set, an optimal learner always ignores it. In the quantum case, however, the situation becomes interesting.

In contrast with the Quantum Coupon Collector problem, we show that the behaviour of the quantum samples for this new problem, with respect to learning, is very close to that in the classical case. The techniques we develop approximately preserve not only the classical sample complexity of learning the exact set, but also the for the case when we allow a small number l of *mismatches*—elements that are in the hypothesis, but not in the underlying set.

Theorem 1.4. *Consider the Quantum Padded Coupon Collector problem for size k subsets of $[n]$ with at most l mismatches and pads ranging in $\mathbb{N}_p$. Any learner for this problem that has success probability at least $1 - \delta$ uses at least*

$$t_0 \quad := \quad k \ln \frac{k+1}{10l+1} + k \ln(1 - 4\delta)$$

samples, provided $n \geq k + 5l$, $l \geq 1$, $p \geq k^{t_0}/\delta^2$, and $\delta \in (0, 1/4)$.

Note that $t_0 \in \Omega(k \log \frac{k}{l})$ for constant δ. Theorem 1.4 helps us establish not only the $\Omega\left(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\right)$ proper PAC learning lower bound for $d \in \Theta(1)$ known for classical learning, but also the more general lower bound of $\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right)$ stated in Theorem 1.3.

To understand the role of padding, it is instructive to examine how sample-efficient learning algorithms for Quantum Coupon Collector work. The algorithm in Ref. [2] and the new algorithm we present in Section 3.1 both heavily exploit the property that when k is "close" to n the quantum samples corresponding to

different subsets are very close to the uniform superposition of all elements. This helps us convert quantum samples for the sets to those of their *complements*, with high probability. Padding enables us to circumvent such algorithms. Since the padding of each element in $[n]$ is an unknown arbitrary value, the quantum samples corresponding to the different subsets with the padding are no longer close to a common pure state, immediately thwarting the known algorithms. Intuitively, due to the independence of the padding from the sets, if the number of possible ways of padding is very large, one may expect that any attempt by a learner to incorporate the register holding the padding in its operations will disturb the coherence of the quantum sample. On the other hand, not using the padding at all is equivalent to tracing it out, again reducing quantum samples to classical samples.

The techniques underlying the proof of Theorem 1.4, however, formalize this intuition only indirectly. They have their origins in the spectral analysis that lies at the heart the optimal lower bounds for PAC and agnostic quantum learning due to Bab Hadiashar, Nayak, and Sinha [8]. We consider the learning problem for sets and the padding chosen uniformly. Taking advantage of the symmetry introduced by the padding, we show that the quantum state corresponding to the random quantum samples may be viewed as a mixture of states that do not contain sufficient information about the set if the number of samples is not sufficiently large. In fact, we relate the information in the states to that obtained from the same number of classical samples. Known properties of the Coupon Collector process then lead to the theorem.

The techniques we develop for QPCC can be applied to padded quantum data in other contexts as well. So we hope that padding can more generally lift other forms of classical learning behaviour to the quantum setting. In particular, these ideas may help answer some questions in Quantum Learning Theory raised in the recent survey by Anshu and Arunachalam [1].

On a more philosophical note, at first glance, the idea of padding might seem artificial. However, in the setting of classical learning, in many cases padding gives a more accurate picture of the kind of data that are available, and of the corresponding learning process. Often, for simplicity and/or computational efficiency, while trying to learn aspects of a population via sampling, we ignore parts of the data from each sample. As a simple example, consider estimating the average height of individuals in a population. For this task, most if not all sensible estimators ignore features such as weight and eye-color, even if such information is present in the samples. This extra unused information behaves like padding.

Organization of the paper. In Section 2 we describe the notation for properties of sequences, as well as for notions in quantum information. This is followed by basics of Quantum Learning Theory and the PAC learning model in both the classical and quantum settings. We then describe the relevant variants of the Coupon Collector problem: the Quantum Coupon Collector, the Padded Coupon Collector, and the Quantum Padded Coupon Collector. At the end we state a tail bound relevant to the properties of the Classical Coupon Collector problem.

Section 3 contains the relevant arguments for proving Theorem 1.2; Section 3.1 contains the details of the algorithm while Section 3.2 establishes the proof of correctness.

Section 4 is devoted to the Quantum Padded Coupon Collector problem, in order to prove the lower bound in Theorem 1.4. In Section 4.1, we reduce the problem to finding a lower bound on the average error over a suitably chosen adversarial distribution, finding a diagonalized form for the ensemble average. In Section 4.2, we exhibit the emerging classical behaviour of this ensemble as the size of the padding increases, proving Theorem 1.4. Finally, we explain how this result gives a separation between proper versus improper learning, i.e., we prove Theorem 1.3 in Section 4.3.

In the appendix, in Appendix A, we bound the deviation from classical behaviour of the ensemble average in Section 4.2. Appendix B presents known results about the Coupon Collector process for com-

pleteness; Section 4.2 calls for these results.

Acknowledgements. This research was supported in part by NSERC Canada. P.S. is also supported by a Mike and Ophelia Lazaridis Fellowship.

2 Preliminaries

General notation. For simplicity we omit ceilings or floors when specifying some integral quantities such as the number of samples used by learning algorithms.

For a positive integer d, we denote the set $\{0, 1, \ldots, d-1\}$ as $\mathbb{N}_d$ and the set $\{1, \ldots, d\}$ as $[d]$. For a sequence $\boldsymbol{b} \in \mathbb{N}_p^n$, define $\mathrm{supp}(\boldsymbol{b}) \coloneqq \{i \in [n] : \boldsymbol{b}_i \neq 0\}$, the set of coordinates of $\boldsymbol{b}$ with non-zero elements. For a sequence $\boldsymbol{i} \in [n]^t$, define $\mathrm{range}(\boldsymbol{i}) \coloneqq \{\boldsymbol{i}_j : j \in [t]\}$, the set of elements of $[n]$ occurring in the sequence $\boldsymbol{i}$. These definitions arise naturally by treating sequences as functions. For a function f with domain $[n]$, and $\boldsymbol{i} \in [n]^t$, we denote the coordinate-wise application of f to $\boldsymbol{i}$ as $f(\boldsymbol{i})$. The scalar product $\boldsymbol{x} \cdot \boldsymbol{y}$ for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{N}_p^t$ is defined as

$$\boldsymbol{x} \cdot \boldsymbol{y} \quad \coloneqq \quad \sum_{i=1}^{t} \boldsymbol{x}_i \boldsymbol{y}_i \pmod{p} \ .$$

Quantum Information notation. For a thorough introduction to the basics of quantum information, we refer the reader to the book by Watrous [16]. We briefly review the notation that we use in this article.

We consider only finite dimensional Hilbert spaces in this work, and denote them either by capital script letters like $\mathcal{A}$ and $\mathcal{B}$, or directly as $\mathbb{C}^m$ for a positive integer m. A *register* is a physical quantum system, and we denote it by a capital letter, like A or B. A quantum register A is associated with a Hilbert space $\mathcal{A}$, and the state of the register is specified by a unit-trace positive semi-definite operator on $\mathcal{A}$. The state is called a *quantum state*, or also as a *density operator*. We denote quantum states by lower case Greek letters, like ρ and σ. We use notation such as ρ^A to indicate that register A is in state ρ, and may omit the superscript when the register is clear from the context. We use *ket* and *bra* notation to denote unit vectors and their adjoints in a Hilbert space, respectively. A quantum state is *pure* if it has rank one, i.e., $\rho = |\phi\rangle\langle\phi|$ where $|\phi\rangle$ is a unit vector. For convenience we sometimes also use the ket notation for unormalised vectors, with an explicit mention. For a linear operator M on some Hilbert space, $\|M\|_{\mathrm{tr}} \coloneqq \mathrm{Tr}\sqrt{M^*M}$ denotes the trace norm of M (also called the ℓ_1 or Schatten 1 norm).

For a set $X \subseteq [n]$, we define $|\psi_X\rangle \in \mathbb{C}^n$ as the uniform superposition over elements of the set X, and Π_X as the orthogonal projection onto the subspace spanned by the elements of X in the same space. I.e.,

$$|\psi_X\rangle \quad \coloneqq \quad \frac{1}{\sqrt{|X|}} \sum_{i \in X} |i\rangle \ , \qquad \text{and}$$

$$\Pi_X \quad \coloneqq \quad \sum_{i \in X} |i\rangle\langle i| \ .$$

Quantum learning theory. We refer the readers to the book [13] for an introduction to machine learning theory and the survey [3] for an introduction to quantum learning theory. Here, we briefly review the notation related to the PAC model that we study in this article.

For some finite, non-empty domain X, we refer to a Boolean function $c : X \to \{0,1\}$ as a *concept*. We may also think of a concept as a bit-string in $\{0,1\}^n$ for $n := |X|$ which lists the value c assigns to each element of X. A *concept class* is a subset $\mathcal{C} \subseteq \{0,1\}^X$ of Boolean functions. For a concept c, we refer to $c(x)$ as the *label* of $x \in X$, and the tuple $(x, c(x))$ as a *labeled example*. We say a concept class $\mathcal{C}$ is *non-trivial* if it contains two distinct concepts c_1, c_2 such that for some $x_1, x_2 \in X$, we have $c_1(x_1) = c_2(x_1)$ and $c_1(x_2) \neq c_2(x_2)$.

A crucial combinatorial quantity in learning Boolean functions is the *VC dimension* of a concept class, introduced by Vapnik and Chervonenkis [15]. We say a set $S := \{s_1, \ldots, s_d\} \subseteq X$ is *shattered* by a concept class $\mathcal{C}$ if for every $a \in \{0,1\}^d$, there exists a concept $c \in \mathcal{C}$ such that $(c(s_1), \ldots, c(s_d)) = a$. The *VC dimension* of $\mathcal{C}$, denoted as $\text{VC-dim}(\mathcal{C})$, is the size of the largest set shattered by $\mathcal{C}$.

PAC model. Consider a concept class $\mathcal{C} \subseteq \{0,1\}^X$. The PAC (*probably approximately correct*) model for learning concepts was introduced—in the classical setting—by Valiant [14], and was extended to the quantum setting by Bshouty and Jackson [6]. In the quantum PAC model, a learning algorithm is given a *quantum PAC example oracle* $\text{QPEX}(c, D)$ for an unknown concept $c \in \mathcal{C}$ and an unknown distribution D over X. The oracle $\text{QPEX}(c, D)$ does not have any inputs. When invoked, it outputs a superposition of labeled examples of c with amplitudes given by the distribution D, namely, the pure state

$$\sum_{x \in \{0,1\}^n} \sqrt{D(x)}\, |x, c(x)\rangle \ ,$$

which we call a *quantum sample*, or simply a *sample*. Note that measuring a quantum sample in the computational basis gives us a labeled example distributed according to D, i.e., a classical sample. We say a Boolean function h, commonly called a *hypothesis*, is an *ϵ-approximation* of c (or has error ϵ) with respect to distribution D, if

$$\Pr_{x \sim D}[h(x) \neq c(x)] \quad \leq \quad \epsilon \ . \tag{2.1}$$

Given access to the oracle $\text{QPEX}(c, D)$, the goal of a quantum PAC learner is to find a hypothesis h that is an ϵ-approximation of c with sufficiently high success probability.

Definition 2.1. *For $\epsilon, \delta \in [0,1]$, we say that an algorithm $\mathcal{A}$ is an* (ϵ, δ)-PAC quantum learner *for the concept class $\mathcal{C}$ if for every $c \in \mathcal{C}$ and distribution D, given access to $\text{QPEX}(c, D)$, with probability at least $1 - \delta$, the algorithm $\mathcal{A}$ outputs a hypothesis $h \in \{0,1\}^X$ which is an ϵ-approximation of c. We say that $\mathcal{A}$ is a* proper *learner if it always outputs a hypothesis $h \in \mathcal{C}$, and say that it is* improper *otherwise.*

The *sample complexity* of a quantum learner $\mathcal{A}$ is the maximum number of times $\mathcal{A}$ invokes the oracle $\text{QPEX}(c, D)$ for any concept $c \in \mathcal{C}$ and any distribution D over X. The *(ϵ, δ)-PAC quantum sample complexity* of a concept class $\mathcal{C}$ is the minimum sample complexity of a (ϵ, δ)-PAC quantum learner for $\mathcal{C}$. Since we can readily derive classical samples from quantum ones, quantum learning algorithms are at least as efficient as classical ones in terms of sample complexity, as well as other measures such as time and space complexity.

Quantum Coupon Collector. Let n be an integer ≥ 3. For a positive integer $k \in (1, n)$, let S be a k-element subset of $[n]$, and let $|\psi_S\rangle$ denote the uniform superposition over the elements of S:

$$|\psi_S\rangle \quad := \quad \frac{1}{\sqrt{k}} \sum_{i \in S} |i\rangle \ .$$

This is a quantum analogue of a uniformly random sample from S, and we call this a *quantum sample* of S. For ease of notation, we define $m \coloneqq n - k$.

In the Quantum Coupon Collector problem, we are given n, k and quantum samples of an arbitrary but fixed, unknown k-element subset S, and we would like to learn the subset using as few samples as possible. By "learning the subset", we mean that we would like to determine, with probability at least $1 - \delta$ for some parameter $\delta \in [0, 1)$, all the k elements of the set S. We are interested in the quantum sample complexity of the problem, i.e., the least number of samples required by a quantum algorithm to learn the set with probability of error at most δ. Observe that by permuting the elements of $[n]$ by a uniformly random permutation, we can show that the optimal worst-case error equals the optimal average-case error under the uniform distribution over the sets.

We may view the (Quantum) Coupon Collector problem as a problem in the PAC learning model as follows. The concept class $\mathcal{C}_{k,n}$ corresponding to the problem with parameters k, n is the set of characteristic vectors of all size k subsets of $[n]$, i.e., $\mathcal{C}_{k,n} \coloneqq \{c \in \{0,1\}^n : |c| = k\}$. Labeled examples for a concept $c \in \mathcal{C}_{k,n}$ are always drawn from the uniform distribution over $i \in [n]$ such that $c(i) = 1$. Hence the label in such examples is superfluous, and the quantum sample considered above is equivalent to a sample in the quantum PAC model. Learning the unknown subset corresponds to *properly* learning the concept with approximation error less than $1/k$.

Padded Coupon Collector. The *Padded Coupon Collector* problem is a variant of the Coupon Collector problem, defined as follows. Let n, k, p be positive integers with $1 < k < n$, and l a non-negative integer. For an unknown set $S \subseteq [n]$ of size k and unknown function $f : [n] \to \mathbb{N}_p$, we are given independent random samples $(i, f(i))$ for i chosen uniformly at random from S. Given $\delta \in [0, 1)$, the goal is to output, with probability at least $1 - \delta$, a size k subset S' such that $|S' \setminus S| \leq l$, using as few samples as possible.

We call $f(i)$ as the *padding* of i, p as the *padding length*, and l as the number of *mismatches*. In the classical case, this problem is as hard as the analogous Coupon Collector problem (in which we allow l mismatches), as we can convert instances of either problem to those of the other, with the same parameters (i.e., k, n, l) by adding or removing the padding. This holds for any value of p.

The quantum analogue, the *Quantum Padded Coupon Collector* problem, with parameters n, k, p, l as before is defined as follows. For an unknown set $S \subseteq [n]$ of size k and unknown function $f : [n] \to \mathbb{N}_p$, we are given copies of *quantum samples* $|\psi_{S,f}\rangle$, where

$$|\psi_{S,f}\rangle \quad \coloneqq \quad \frac{1}{\sqrt{k}} \sum_{i \in S} |i\rangle|f(i)\rangle \ .$$

Given $\delta \in [0, 1)$, the goal is again to output, with probability at least $1 - \delta$, a size k subset S' with at most l mismatches, i.e., with $|S' \setminus S| \leq l$, using as few samples as possible.

As before, quantum samples reduce to random samples as in the classical problem, when measured in the computational basis. Unlike the classical case, it is not possible to convert individual quantum samples to those of the unpadded variant, i.e., the Quantum Coupon Collector problem, when $p > 1$. Applying the unitary transform $|i\rangle|f(i)\rangle \to |i\rangle|0\rangle$ requires knowledge of the entire function f.

Tail bounds. To analyse the classical Coupon Collector process in Section B, we need the Hoeffding Bound for Hypergeometric series [11]. If out of n balls, αn balls are red and $(1 - \alpha)n$ balls are blue for some $\alpha \in (0, 1)$, and X is the random variable corresponding to the number of red balls picked when

picking t out of n balls without replacement, then we have the following concentration bounds.

$$\Pr[X - \alpha t \geq \lambda] \leq \exp\left(-\frac{2\lambda^2}{t}\right) \quad , \quad \Pr[X - \alpha t \leq -\lambda] \leq \exp\left(-\frac{2\lambda^2}{t}\right). \tag{2.2}$$

3 Quantum Coupon Collector

In this section, we present a more efficient learning algorithm for the Quantum Coupon Collector problem. The number of samples used by the algorithm matches the recently established sharp lower bound [8] exactly in the leading order term, for constant probability of error.

3.1 The algorithm

When $3m\ln(\mathrm{e}\,m) > n$, we reduce the Quantum Coupon Collector problem to its classical version. I.e., we measure $k\ln k + \Theta(k)$ quantum samples $|\psi_S\rangle$ in the standard basis and output the set of all coupons observed. So, we need only focus on the "small" m regime, i.e., when $3m\ln(\mathrm{e}\,m) \leq n$.

As in Ref. [2], we try to learn the complement $\overline{S}$ of the set S when m is small, but we do this differently. We maintain a guess G for the set $\overline{S}$, and use the quantum samples to improve our guess via suitable measurements. The measurements of a quantum sample may reveal that an element x has been misclassified as being in $\overline{S}$, in which case we remove it from G. Alternatively, we may discover a new element x that is likely to be in $\overline{S}$. In this case, we add x to G. Sometimes, the measurements do not reveal either type of element, and we move to the next sample. We show that after sufficiently many such iterations, the guess G equals $\overline{S}$ with constant probability close to 1.

More formally, the algorithm keeps track of the following subsets of $[n]$:

1. G: This is the current guess for the complement $\overline{S}$.

2. U: This is a set that the algorithm currently believes contains S.

Initially, G is empty. Ideally the algorithm keeps adding elements from the complement $\overline{S}$ to this set G until all the desired elements have been collected. The set U is initially set to $[n]$, and is updated to $[n] \setminus G$, whenever G is updated. We refer to elements $i \notin \overline{S}$ that have been incorrectly added to G as *rogue* coupons. The set of rogue coupons at any time is precisely $G \cap S$.

Recall the notation $|\psi_X\rangle$ and Π_X for the uniform superposition over elements of the set $X \subseteq [n]$ and the orthogonal projection onto the subspace spanned by the elements of X, respectively. Let $M_0 \coloneqq \Pi_G$ and $M_1 \coloneqq \mathbb{I} - M_0$. The algorithm first measures a quantum sample $|\psi_S\rangle$ according to (M_0, M_1). If the outcome is 0, the guess G necessarily has rogue coupons. We measure the residual state, viz. $|\psi_{G\cap S}\rangle$, to identify a rogue coupon and remove it from G. If the outcome is 1, the algorithm further measures the residual state $|\psi_{U\cap S}\rangle$ according to (E_0, E_1), where $E_0 \coloneqq |\psi_U\rangle\langle\psi_U|$ and $E_1 \coloneqq \mathbb{I} - E_0$. If the outcome is 1, as we show later, the elements of $\overline{S}\cap U$ occur with higher (positive) amplitude in the resulting residual state. We therefore measure the residual state in the computational basis, and add the outcome to G. Algorithm 1 summarises all the steps of the algorithm.

With "high" probability, in step 21, we keep adding elements of $\overline{S}$ to G. Due to our choice of U as the complement of G (and therefore a set disjoint from G), the element added in this step is distinct from all the elements currently in G. Steps 11 to 15 are designed to correct our guess G, since it is possible that in step 21, we mistakenly add elements from S to G. By measuring $|\psi_S\rangle$ with $|\psi_G\rangle\langle\psi_G|$, we check if any rogue coupons were previously collected in G. If rogue coupons are detected, we remove them from G. If

Algorithm 1: Quantum-Coupon-Collector

Input : positive integers n, k such that $1 < k < n$, error parameter $\delta \in (0,1)$, and quantum samples $|\psi_S\rangle$ for an unknown size k subset $S \subset [n]$

Output : subset $T \subset [n]$

1 $m \leftarrow n - k$;
2 **if** $3m\ln(\mathrm{e}\, m) > n$ **then** $\ell \leftarrow k\ln k + k\ln\frac{1}{\delta}$;
3 **else** $\ell \leftarrow k\ln m + k\ln\frac{\mathrm{e}}{\delta}$;
4 Obtain ℓ copies of the quantum sample, viz., $|\psi_S\rangle^{\otimes\ell}$;
5 **if** $3m\ln(\mathrm{e}\, m) > n$ **then**
6 Measure each of the ℓ quantum samples $|\psi_S\rangle$ in the computational basis ;
7 $T \leftarrow$ the set of all the outcomes observed ;
8 **else** /* $3m\ln(\mathrm{e}\, m) \le n$ */
9 $G \leftarrow \emptyset, U \leftarrow [n]$;
10 **for** $t = 1$ *to* ℓ **do**
11 $M_0 \leftarrow \Pi_G$, $M_1 \leftarrow \mathbb{I} - M_0$;
12 Measure the t-th quantum sample $|\psi_S\rangle$ according to (M_0, M_1) to get an outcome $a \in \{0,1\}$ and corresponding residual state $|\xi_a\rangle$;
13 **if** *outcome* $a = 0$ **then**
14 Measure the residual state $|\xi_0\rangle$ in the computational basis to get some outcome x ;
15 $G \leftarrow G \setminus \{x\}, U \leftarrow U \cup \{x\}$;
16 **else** /* outcome $a = 1$ */
17 $E_0 \leftarrow |\psi_U\rangle\langle\psi_U|$, $E_1 \leftarrow \mathbb{I} - E_0$;
18 Measure the residual state $|\xi_1\rangle$ according to (E_0, E_1) to get an outcome $b \in \{0,1\}$ and corresponding residual state $|\phi_b\rangle$;
19 **if** *outcome* $b = 1$ **then**
20 Measure the residual state $|\phi_1\rangle$ in the computational basis to get some outcome x ;
21 $G \leftarrow G \cup \{x\}, U \leftarrow U \setminus \{x\}$;
22 $T \leftarrow G$;
23 **return** T

no rogue coupons were collected, or if all the ones that were collected have been removed, then steps 14 and 15 are not executed. We claim that at the end of this procedure, G is, with all but a small probability of failure, the set $\overline{S}$. We prove the correctness of the algorithm in the next section.

3.2 The analysis

The correctness of Algorithm 1 when $3m\ln(\mathrm{e}\, m) > n$ follows from well-known properties of the classical Coupon Collector process. A straightforward calculation shows that the expected number of coupons that remain to be collected after t samples have been obtained is at most $k(1 - \frac{1}{k})^t$, which is bounded above by δ when $t = k\ln k + k\ln\frac{1}{\delta}$. This implies that the probability that the algorithm *does not* collect all the k coupons in S with $k\ln k + k\ln\frac{1}{\delta}$ samples is at most δ.

So we need only analyse Algorithm 1 when $3m\ln(\mathrm{e}\, m) \le n$. We do so by studying how far the guess G is from $\overline{S}$, i.e., the size of the symmetric difference of the two sets. We show that in expectation, this distance

is strictly decreasing, at a sufficiently high rate. To implement this approach, we keep track of two more sets.

1. R: the set of rogue coupons. This is the set $S \cap G$.
2. C: the set of coupons from $\overline{S}$ that remain to be collected. This is the set $\overline{S} \setminus G$

We denote the cardinality of R and C by J_t and L_t, respectively, after iteration t. So $J_0 = 0$ and $L_0 = m$. Note that $(J_t : t \geq 0)$ and $(L_t : t \geq 0)$ are random walks over non-negative integers, with $J_t \in [0, k]$ and $L_t \in [0, m]$.

The number of rogue coupons J_t may decrease in step 15, or increase by 1 in step 21, or stay the same, in one iteration. The number of coupons L_t in C may stay the same, or decrease by 1 in step 21, in one iteration. Finally, at most one of J_t or L_t changes in one iteration. The algorithm succeeds if and only if there are neither any rogue coupons in G nor any coupons from $\overline{S}$ that remain to be collected when it terminates, i.e., $J_\ell = L_\ell = 0$, where ℓ is the total number of iterations. This is equivalent to the sum $J_\ell + L_\ell$ being 0. So we track the random walk $K_t \coloneqq J_t + L_t$ instead.

Conditioned on $J_t = j$ and $L_t = l$ for some $j \geq 0$ and $l \in [0, m]$, we compute some states and the probability of some events occurring in the $(t+1)$-th iteration. The probability of obtaining outcome 0 in step 12, and thus reducing J_t by 1, is

$$\|\Pi_G |\psi_S\rangle\|^2 \quad = \quad \sum_{i \in G \cap S} \frac{1}{k} \quad = \quad \frac{j}{k} \ .$$

The probability of executing step 16 is therefore $1 - \frac{j}{k}$. Conditioned on executing step 16, i.e., getting outcome 1 in step 12, the residual state is $|\psi_{S \cap U}\rangle$. So the conditional probability of getting outcome 1 in step 18 is

$$1 - \langle \psi_U | \psi_{S \cap U} \rangle^2 \quad = \quad 1 - \frac{|S \cap U|}{|U|} \quad = \quad 1 - \frac{k-j}{k-j+l} \quad = \quad \frac{l}{k-j+l} \ ,$$

as $S \cap U = S \setminus R$, and $U = (S \cap U) \cup C$. The unnormalised resultant state on getting outcome 1 is

$$\begin{aligned}
(\mathbb{I} - |\psi_U\rangle\langle\psi_U|)|\psi_{S\cap U}\rangle \quad &= \quad |\psi_{S\cap U}\rangle - \langle\psi_U|\psi_{S\cap U}\rangle|\psi_U\rangle \\
&= \quad |\psi_{S\cap U}\rangle - \sqrt{\frac{k-j}{k-j+l}}\,|\psi_U\rangle \\
&= \quad |\psi_{S\cap U}\rangle - \frac{k-j}{k-j+l}\,|\psi_{S\cap U}\rangle - \frac{\sqrt{(k-j)l}}{k-j+l}\,|\psi_C\rangle \\
&= \quad \frac{l}{k-j+l}\,|\psi_{S\cap U}\rangle - \frac{\sqrt{(k-j)l}}{k-j+l}\,|\psi_C\rangle \ .
\end{aligned}$$

So the normalised state $|\phi_1\rangle$ is

$$|\phi_1\rangle \quad = \quad \sqrt{\frac{l}{k-j+l}}\,|\psi_{S\cap U}\rangle - \sqrt{\frac{k-j}{k-j+l}}\,|\psi_C\rangle \ .$$

From the above expression for $|\phi_1\rangle$, it is immediate that the probability of adding a new rogue coupon to G in step 21, conditioned on getting this state, is $\frac{l}{k-j+l}$. Similarly, the conditional probability for adding a

fresh coupon from $\overline{S}$ to G is $\frac{k-j}{k-j+l}$. The corresponding unconditional probabilities in iteration $t+1$ (for J_t increasing or L_t decreasing by 1) are

$$\frac{k-j}{k} \times \frac{l}{k-j+l} \times \frac{l}{k-j+l} \qquad \text{and} \qquad \frac{k-j}{k} \times \frac{l}{k-j+l} \times \frac{k-j}{k-j+l} \ ,$$

respectively.

Using these results, we can compute the transition probabilities for the random walk (K_t) in the $(t+1)$-th iteration. We have $K_{t+1} = K_t + 1$ when the number of rogue coupons J_t increases by 1. This happens with probability

$$\frac{(k-j)l^2}{k(k-j+l)^2} \ .$$

We have $K_{t+1} = K_t - 1$ when either the number of rogue coupons J_t decreases, or when we collect a new coupon from $\overline{S}$, i.e., L_t decreases. This happens with probability

$$\frac{j}{k} + \frac{(k-j)^2 l}{k(k-j+l)^2} \ .$$

Thus, defining $r \coloneqq j + l$, we get

$$\begin{aligned}
\mathbb{E}\left[K_{t+1} \middle| J_t = j,\, L_t = l\right] &= r + \frac{(k-j)l^2}{k(k-j+l)^2} - \frac{j}{k} - \frac{(k-j)^2 l}{k(k-j+l)^2} \\
&= r - \frac{(k-j)(k-j-l)l}{k(k-j+l)^2} - \frac{r-l}{k} \\
&= r\left(1 - \frac{1}{k}\right) - \frac{l}{k}\left(\frac{(k-j)(k-j-l)}{(k-j+l)^2} - 1\right) \\
&= r\left(1 - \frac{1}{k}\right) + \frac{l^2}{k}\left(\frac{3(k-j)+l}{(k-j+l)^2}\right) \\
&\leq r\left(1 - \frac{1}{k}\right) + \frac{3l^2}{(k-j+l)k} \\
&= r\left(1 - \frac{1}{k}\right) + \frac{3rl^2}{k(k-j+l)(j+l)} \ .
\end{aligned}$$

We bound the second term from above as follows. For $l = 0$, it is bounded by 0. Otherwise, as a function of $j \in [0, k]$, the expression $(k-j+l)(j+l)$ is minimised at $j \in \{0, k\}$, and the minimum is $(k+l)l$. Moreover, as a function of l the expression $l/(k+l)$ is maximised at $l = m$. So

$$\frac{l^2}{(k-j+l)(j+l)} \quad \leq \quad \frac{l^2}{l(k+l)} \quad \leq \quad \frac{m}{n} \ .$$

Using this, we get

$$\begin{aligned}
\mathbb{E}\left[K_{t+1} \middle| J_t = j,\, L_t = l\right] &\leq r\left(1 - \frac{1}{k}\right) + \frac{r}{k} \times \frac{3m}{n} \\
&= r\left(1 - \frac{1}{k}\left(1 - \frac{3m}{n}\right)\right) \ .
\end{aligned}$$

Since $K_t = r = j + l$ when $J_t = j, L_t = l$, we have

$$\begin{aligned}
\mathbb{E}[K_{t+1}] &\leq \mathbb{E}[K_t]\left(1 - \frac{1}{k}\left(1 - \frac{3m}{n}\right)\right) \\
&\leq \mathbb{E}[K_0]\left(1 - \frac{1}{k}\left(1 - \frac{3m}{n}\right)\right)^{t+1} \\
&= m\left(1 - \frac{1}{k}\left(1 - \frac{3m}{n}\right)\right)^{t+1} \quad . \qquad (3.1)
\end{aligned}$$

Note that the behaviour of the random variable K_t is very similar to that of the number of coupons that remain to be collected in the classical Coupon Collector process.

Let $\ell \coloneqq k \ln m + ck$ for a positive parameter c to be specified later. As $3m \ln(\mathrm{e}\, m) \leq n$, by Eq. (3.1),

$$\begin{aligned}
\mathbb{E}[K_\ell] &\leq m\left(1 - \frac{1}{k}\left(1 - \frac{3m}{n}\right)\right)^{k \ln m + ck} \\
&\leq \exp\left(\frac{3m \ln m - c(n - 3m)}{n}\right) \qquad \text{(using } 1 + z \leq \mathrm{e}^z\text{)} \\
&\leq \exp(1 - c) \ .
\end{aligned}$$

Thus, the probability that K_ℓ is not 0 is at most $\exp(1 - c)$. Taking $c \coloneqq \ln \frac{\mathrm{e}}{\delta}$, we get a probability of failure of at most δ. This proves Theorem 1.2.

4 Quantum Padded Coupon Collector

In this section, we establish strong bounds on the sample complexity of the Quantum Padded Coupon Collector problem, thereby proving Theorem 1.4. Observe that at the heart of the problem we have a question about learning a subset from an ensemble of quantum states. In Section 4.1, we analyse this ensemble to bring it into diagonal form—a form that is more amenable to analysis. We then relate the problem to the *classical* Coupon Collector process via an approximation of the ensemble in Section 4.2) to conclude the desired bound. Finally, in Section 4.3, we show how this bound gives us a strong lower bound for proper PAC learning.

4.1 Simplified ensemble

Assume we have an algorithm for the problem with parameters n, k, l, p that uses t samples. We consider the average-case success probability of the algorithm when the underlying set-function pair S, f are chosen uniformly at random from their respective domains. We show that this probability is suitably small, if t is bounded as in Theorem 1.4. For ease of reference, we call the version of the Quantum Padded Coupon Collector problem with uniformly random S, f, with access to a total of t samples, as $\mathcal{T}_1$.

The goal of the task $\mathcal{T}_1$ is to approximate the unknown set S regardless of the padding function f. So we consider the quantum state σ_S corresponding to t quantum samples $|\psi_{S,f}\rangle^{\otimes t}$ for uniformly random f. Treating f as an element of $\mathbb{N}_p^n$, we have

$$\sigma_S = \frac{1}{p^n} \sum_{f \in \mathbb{N}_p^n} (|\psi_{S,f}\rangle\langle\psi_{S,f}|)^{\otimes t} = \frac{1}{p^n k^t} \sum_{\boldsymbol{i},\boldsymbol{j} \in S^t} |\boldsymbol{i}\rangle\langle\boldsymbol{j}| \otimes \sum_{f \in \mathbb{N}_p^n} |f(\boldsymbol{i})\rangle\langle f(\boldsymbol{j})| \ , \qquad (4.1)$$

where, for ease of notation, we have rearranged the registers in the final expression to collect the padding together.

Viewing $\mathbb{N}_p^n$ as the additive group $\mathbb{Z}_p^n$, we see that σ_S is invariant under the action induced by $f \mapsto f+g$, for $g \in \mathbb{N}_p^n$. Thus σ_S is block-diagonal in the same basis as this action, for every S. With this in mind, consider $\rho_S \coloneqq (\mathbb{I} \otimes F^{\otimes t})\, \sigma_S (\mathbb{I} \otimes (F^*)^{\otimes t})$ where F is the Fourier transform over $\mathbb{Z}_p$:

$$F \quad \coloneqq \quad \frac{1}{\sqrt{p}} \sum_{j \in \mathbb{N}_p} \omega^{ij} |j\rangle\langle i| \ ,$$

and $\omega \coloneqq \exp\left(\frac{2\pi \mathrm{i}}{p}\right)$ is a primitive p-th root of unity. We have

$$\begin{aligned} \rho_S \quad &= \quad \frac{1}{p^n k^t} \sum_{\boldsymbol{i},\boldsymbol{j} \in S^t} |\boldsymbol{i}\rangle\langle \boldsymbol{j}| \otimes \sum_{f \in \mathbb{N}_p^n} F^{\otimes t} |f(\boldsymbol{i})\rangle\langle f(\boldsymbol{j})| (F^*)^{\otimes t} \\ &= \quad \frac{1}{p^{n+t} k^t} \sum_{\boldsymbol{i},\boldsymbol{j} \in S^t} |\boldsymbol{i}\rangle\langle \boldsymbol{j}| \otimes \sum_{\boldsymbol{x},\boldsymbol{y} \in \mathbb{N}_p^t} \left(\sum_{f \in \mathbb{N}_p^n} \omega^{f(\boldsymbol{i})\cdot \boldsymbol{x} - f(\boldsymbol{j})\cdot \boldsymbol{y}} \right) |\boldsymbol{x}\rangle\langle \boldsymbol{y}| \ . \end{aligned} \tag{4.2}$$

Now, we compute the value in brackets in Eq. (4.2), i.e.,

$$\sum_{f \in \mathbb{N}_p^n} \omega^{f(\boldsymbol{i})\cdot \boldsymbol{x} - f(\boldsymbol{j})\cdot \boldsymbol{y}} \ . \tag{4.3}$$

The scalar product $f(\boldsymbol{i}) \cdot \boldsymbol{x}$ can be interpreted as a weighted sum of the coordinates of the vector $f(\boldsymbol{i})$, with weights given by $\boldsymbol{x}$. Consider what happens if we increase one component of f by 1. More specifically, suppose we increase the value of $f(q)$ by 1 for $q \in [n]$, and keep the rest of f the same to get a new function f'. Then, $f'(x) - f(x) = \mathbb{1}[x = q]$, and we have

$$(f(\boldsymbol{i}) \cdot \boldsymbol{x} - f(\boldsymbol{j}) \cdot \boldsymbol{y}) - (f'(\boldsymbol{i}) \cdot \boldsymbol{x} - f'(\boldsymbol{j}) \cdot \boldsymbol{y}) \quad = \quad \sum_{r \in [t] \,:\, \boldsymbol{i}_r = q} \boldsymbol{x}_r - \sum_{r \in [t] \,:\, \boldsymbol{j}_r = q} \boldsymbol{y}_r \ .$$

If this quantity is $0 \pmod p$, then changing $f(q)$ does not change $f(\boldsymbol{i}) \cdot \boldsymbol{x} - f(\boldsymbol{j}) \cdot \boldsymbol{y} \pmod p$. Otherwise, as $f(q)$ cycles through $\mathbb{N}_p$, the expression $f(\boldsymbol{i}) \cdot \boldsymbol{x} - f(\boldsymbol{j}) \cdot \boldsymbol{y}$ cycles through all multiples of the quantity $\pmod p$. Since $\sum_{i=0}^{p-1} \omega^{ij} = 0$ for any non-zero $j \in \mathbb{N}_p$, we have that if

$$\sum_{r \in [t] \,:\, \boldsymbol{i}_r = q} \boldsymbol{x}_r - \sum_{r \in [t] \,:\, \boldsymbol{j}_r = q} \boldsymbol{y}_r \quad = \quad 0 \pmod p$$

for all q, then the sum in Eq. (4.3) is p^n, and otherwise it equals 0. This leads us to defining the *modular signature* $\mathrm{ms} : [n]^t \times \mathbb{N}_p^t \to \mathbb{N}_p^n$ of a pair $(\boldsymbol{i}, \boldsymbol{x})$; for each $q \in [n]$, the q-th coordinate of the modular signature is given by

$$\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x})_q \quad \coloneqq \quad \sum_{r \in [t] \,:\, \boldsymbol{i}_r = q} \boldsymbol{x}_r \pmod p \ . \tag{4.4}$$

With this notion, we can evaluate the expression in Eq. (4.3) as

$$\sum_{f \in \mathbb{N}_p^n} \omega^{f(\boldsymbol{i})\cdot \boldsymbol{x} - f(\boldsymbol{j})\cdot \boldsymbol{y}} \quad = \quad \begin{cases} p^n & \text{if } \mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \mathrm{ms}(\boldsymbol{j}, \boldsymbol{y}) \\ 0 & \text{otherwise.} \end{cases}$$

Combining this with Eq. (4.2), we get

$$\begin{aligned}\rho_S &= \frac{1}{p^t k^t} \sum_{\substack{\boldsymbol{i},\boldsymbol{j}\in S^t\,;\,\boldsymbol{x},\boldsymbol{y}\in\mathbb{N}_p^t\,:\\ \mathrm{ms}(\boldsymbol{i},\boldsymbol{x})=\mathrm{ms}(\boldsymbol{j},\boldsymbol{x})}} |\boldsymbol{i}\rangle\langle\boldsymbol{j}| \otimes |\boldsymbol{x}\rangle\langle\boldsymbol{y}| \\ &= \frac{1}{p^t k^t} \sum_{\boldsymbol{b}\in\mathbb{N}_p^n} |\phi_{S,\boldsymbol{b}}\rangle\langle\phi_{S,\boldsymbol{b}}| \ , \end{aligned} \tag{4.5}$$

where $|\phi_{S,\boldsymbol{b}}\rangle$ is the unnormalised state defined as

$$|\phi_{S,\boldsymbol{b}}\rangle \ := \sum_{\boldsymbol{i}\in S^t\,,\,\boldsymbol{x}\in\mathbb{N}_p^t\,:\,\mathrm{ms}(\boldsymbol{i},\boldsymbol{x})=\boldsymbol{b}} |\boldsymbol{i}\rangle|\boldsymbol{x}\rangle \ .$$

Note that the states $|\phi_{S,\boldsymbol{b}}\rangle$ are orthogonal for different $\boldsymbol{b}$, and

$$\sum_{\boldsymbol{b}} \|\phi_{S,\boldsymbol{b}}\|^2 = k^t p^t \ . \tag{4.6}$$

Note also that if $|\phi_{S,\boldsymbol{b}}\rangle \neq 0$, i.e., the $\boldsymbol{b}$ occurs as the modular signature of some pair $\boldsymbol{i} \in S^t, \boldsymbol{x} \in \mathbb{N}_p^t$, then for any $i \in [n]$, if we have $\boldsymbol{b}_i \neq 0$, then $i \in S$. In other words, if $|\phi_{S,\boldsymbol{b}}\rangle \neq 0$, we have $\mathrm{supp}(\boldsymbol{b}) \subseteq S$.

Since the ensemble of states (ρ_S) is related to (σ_S) by the Fourier transform, the task $\mathcal{T}_1$ is equivalent to approximating S given ρ_S, for a uniformly random subset S.

4.2 Relation to the classical case

Next, we show that the states ρ_S may be approximated closely by states ρ'_S which we define below, and the latter are easier to analyse. In fact, we show that estimating the set S given ρ'_S is closely related to the *classical* Coupon Collector process.

Define ρ'_S to be the following quantum state.

$$\rho'_S \ := \ \frac{1}{p^t k^t} \sum_{\boldsymbol{b}\in\mathbb{N}_p^n} \frac{\|\phi_{S,\boldsymbol{b}}\|^2}{\|\phi_{\boldsymbol{b}}\|^2} |\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}| \ ,$$

where

$$|\phi_{\boldsymbol{b}}\rangle \ := \sum_{\boldsymbol{i}\in\mathrm{supp}(\boldsymbol{b})^t,\,\boldsymbol{x}\in\mathbb{N}_p^t\,:\,\mathrm{ms}(\boldsymbol{i},\boldsymbol{x})=\boldsymbol{b}} |\boldsymbol{i}\rangle|\boldsymbol{x}\rangle \ .$$

I.e., we replace $|\phi_{S,\boldsymbol{b}}\rangle$ in ρ_S with the unormalised state $|\phi_{\boldsymbol{b}}\rangle$ and normalise appropriately to construct ρ'_S. Note that the states $|\phi_{\boldsymbol{b}}\rangle$ are also orthogonal to each other for different $\boldsymbol{b}$. Due to the properties of the states $|\phi_{S,\boldsymbol{b}}\rangle$ described in Section 4.1, we see that ρ'_S has support on $|\phi_{\boldsymbol{b}}\rangle$ only if $\mathrm{supp}(\boldsymbol{b}) \subseteq S$. Moreover, for modular signature $\boldsymbol{b}$ and any two subsets S, S' of size k both of which contain $\mathrm{supp}(\boldsymbol{b})$, we have $\|\phi_{S,\boldsymbol{b}}\|^2 = \|\phi_{S',\boldsymbol{b}}\|^2$, by symmetry. (We may permute $[n]$ to map S to S' while preserving $\mathrm{supp}(\boldsymbol{b})$ and the number of basis elements $|\boldsymbol{i},\boldsymbol{x}\rangle$ occurring in the two states.)

Consider the distance between ρ_S and ρ'_S.

$$\begin{aligned}\left\|\rho_S - \rho'_S\right\|_{\mathrm{tr}} &\leq \frac{1}{p^t k^t} \sum_{\boldsymbol{b}\in\mathbb{N}_p^n} \|\phi_{S,\boldsymbol{b}}\|^2 \left\| \frac{|\phi_{S,\boldsymbol{b}}\rangle\langle\phi_{S,\boldsymbol{b}}|}{\|\phi_{S,\boldsymbol{b}}\|^2} - \frac{|\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}|}{\|\phi_{\boldsymbol{b}}\|^2} \right\|_{\mathrm{tr}} \\ &\leq \max_{T\,:\,\mathrm{supp}(\boldsymbol{b})\subseteq T} \left\| \frac{|\phi_{T,\boldsymbol{b}}\rangle\langle\phi_{T,\boldsymbol{b}}|}{\|\phi_{T,\boldsymbol{b}}\|^2} - \frac{|\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}|}{\|\phi_{\boldsymbol{b}}\|^2} \right\|_{\mathrm{tr}} \ , \end{aligned} \tag{4.7}$$

by Eq. (4.6) and the relationship between $\mathrm{supp}(\boldsymbol{b})$ and the set S described above. In Lemma A.3 in Appendix A we show that this distance is bounded above by $\frac{1}{2}\sqrt{k^t/p}$. The high-level reason this bound holds is that most of the terms $|\boldsymbol{i}, \boldsymbol{x}\rangle$ occuring in $|\phi_{S,b}\rangle$ also occur in $|\phi_{\boldsymbol{b}}\rangle$.

Call the task of estimating S with at most l mismatches given ρ'_S instead of ρ_S as $\mathcal{T}_2$. (As before, the S is chosen uniformly at random.) Denote the maximum success probability over all algorithms for task $\mathcal{T}_i$ as $\mathrm{opt}(\mathcal{T}_i)$, for $i \in \{1, 2\}$. We have

$$\begin{aligned}
\mathrm{opt}(\mathcal{T}_1) &\leq \mathrm{opt}(\mathcal{T}_2) + \frac{1}{2} \max_{S\,:\,\mathrm{supp}(\boldsymbol{b})\subseteq S} \left\| \frac{|\phi_{S,\boldsymbol{b}}\rangle\langle\phi_{S,\boldsymbol{b}}|}{\|\phi_{S,\boldsymbol{b}}\|^2} - \frac{|\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}|}{\|\phi_{\boldsymbol{b}}\|^2} \right\|_{\mathrm{tr}} \\
&\leq \mathrm{opt}(\mathcal{T}_2) + \sqrt{\frac{k^t}{p}} \qquad \text{(By Lemma A.3).} \qquad (4.8)
\end{aligned}$$

So we focus our attention on bounding $\mathrm{opt}(\mathcal{T}_2)$.

Since ρ'_S is diagonalised via the modular signature, i.e., by the (unnormalised) states $|\phi_{\boldsymbol{b}}\rangle$, without loss in generality, we assume that an optimal learner $\mathcal{A}$ for task $\mathcal{T}_2$ first measures the modular signature. If $\mathcal{A}$ observes $\boldsymbol{b}$ as the modular signature, the resultant state is $|\phi_{\boldsymbol{b}}\rangle$, regardless of the unknown set S. Since $\mathcal{A}$ optimizes the average error, and the *a posteriori* probability of any size k subset S' containing $\mathrm{supp}(\boldsymbol{b})$ is the same, the optimal algorithm $\mathcal{A}$ outputs a uniformly random such subset S'. This implies that as long as the support of $\boldsymbol{b}$ has size at least $k - l$, the algorithm produces an estimate with at most l mismatches with S.

It remains to bound the probability of observing $\boldsymbol{b}$ with $|\mathrm{supp}(\boldsymbol{b})| \geq k - l$. The probability of observing a modular signature $\boldsymbol{b}$ with support size b for a given S and ρ'_S is

$$\begin{aligned}
\frac{1}{p^t k^t} \sum_{\boldsymbol{b}\,:\,|\mathrm{supp}(\boldsymbol{b})|=b} \|\phi_{S,\boldsymbol{b}}\|^2 &= \frac{1}{p^t k^t} \sum_{\boldsymbol{b}\,:\,|\mathrm{supp}(\boldsymbol{b})|=b} \Big|\{(\boldsymbol{i}, \boldsymbol{x}) : \boldsymbol{i} \in S^t, \boldsymbol{x} \in \mathbb{N}_p^t\,, \mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}\}\Big| \\
&= \frac{1}{p^t k^t} \Big|\{(\boldsymbol{i}, \boldsymbol{x}) : \boldsymbol{i} \in S^t, \boldsymbol{x} \in \mathbb{N}_p^t\,, |\mathrm{supp}(\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}))| = b\}\Big| \\
&= \Pr_{\boldsymbol{I}\sim S^t, \boldsymbol{X}\sim \mathbb{N}_p^t} \big[\, |\mathrm{supp}(\mathrm{ms}(\boldsymbol{I}, \boldsymbol{X}))| = b\big]\ ,
\end{aligned}$$

where $\boldsymbol{I}$ and $\boldsymbol{X}$ are drawn uniformly at random from S^t and $\mathbb{N}_p^t$, respectively. Thus, the probability that the modular signature $\boldsymbol{b}$ observed has $|\mathrm{supp}(\boldsymbol{b})| \geq b$ is bounded as

$$\Pr_{\boldsymbol{I}\sim S^t, \boldsymbol{X}\sim \mathbb{N}_p^t} \Big[\, |\mathrm{supp}(\mathrm{ms}(\boldsymbol{I}, \boldsymbol{X}))| \geq b\Big] \leq \Pr_{\boldsymbol{I}\sim S^t} \big[\, |\mathrm{range}(\boldsymbol{I})| \geq b\big]\ .$$

Note that the bound is exactly the probability of collecting at least b coupons from S in t steps. This leads us to an approximation variant of the classical Coupon Collector problem. Define $\mathcal{T}_3$ as the task of estimating the set S by a size k subset S' with $|S' \setminus S| \leq l$ mismatches, given t independent samples from an unknown, uniformly random subset $S \subseteq [n]$ of size k. The reasoning above gives us the following lemma.

Lemma 4.1. *The success probability* $\mathrm{opt}(\mathcal{T}_2)$ *is bounded above by the optimal probability of success for the task* $\mathcal{T}_3$.

The Coupon Collector process is well-studied, and a bound on the success probability for task $\mathcal{T}_3$ is likely known. For completeness, we derive a bound on the probability in Lemma B.3 in Appendix B.

We put all this together to prove Theorem 1.4, which we restate here for convenience.

Theorem 1.4. *Consider the Quantum Padded Coupon Collector problem for size k subsets of $[n]$ with at most l mismatches and pads ranging in $\mathbb{N}_p$. Any learner for this problem that has success probability at least $1-\delta$ uses at least*

$$t_0 \quad := \quad k \ln \frac{k+1}{10l+1} + k \ln(1-4\delta)$$

samples, provided $n \geq k+5l$, $l \geq 1$, $p \geq k^{t_0}/\delta^2$, and $\delta \in (0, 1/4)$.

Proof: Increasing n and decreasing the number of samples t only makes the problem harder, and thus it suffices to show that the success probability is strictly smaller than $1-\delta$ when $n := k+5l$ and $t := t_0 - 1$. Taking l and p as in the theorem statement, by Lemma B.3, the probability of estimating a uniformly random size k subset with at most l mismatches is $1-2\delta$. By Lemma 4.1 we get $\mathrm{opt}(\mathcal{T}_2) < 1-2\delta$. By our choice of p and Eq. (4.8) we have $\mathrm{opt}(\mathcal{T}_1) < 1-2\delta+\delta = 1-\delta$. Since $\mathcal{T}_1$ is the task of solving the Quantum Padded Coupon Collector problem for a uniformly random subset (and uniformly random padding), the theorem follows. ■

4.3 Reduction to proper PAC learning

Similar to the classical case, for each d and ϵ, we describe a concept class for which proper quantum PAC learning (with up to a small constant probability of error) has sample complexity $\Omega\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon} + \frac{1}{\epsilon}\log\frac{1}{\delta}\right)$. The latter summand is obtained via a straightforward argument as in Refs. [4, 8], so we focus on the first summand.

Fix positive integer parameters k, p, d. Let $n := k+d$. We define the concept class $\mathcal{C}_{n,k,p}$ over the domain $X := [n] \times \mathbb{N}_p$ as follows:

$$\mathcal{C}_{n,k,p} \quad := \quad \{g : \exists S \subseteq [n], |S| = k, g(i,x) = \mathbb{1}[i \in S] \; \forall i \in [n], x \in \mathbb{N}_p\} \ .$$

In the above, we refer to the concept corresponding to S as g_S. We may verify that for $k \geq d$, the VC dimension of the class $\mathcal{C}_{n,k,p}$ is exactly d. This is a "padded" version of the concept class $\mathcal{C}_{n,k}$ used to show a separation between proper and improper learning in the classical case. It is related to the Quantum Padded Coupon Collector problem in the following sense.

Lemma 4.2. *Any proper (ϵ, δ)-PAC quantum learner $\mathcal{A}$ for $\mathcal{C}_{n,k,p}$ also solves the Quantum Padded Coupon Collector problem with the same parameters n, k, p, and with $l := \lfloor \epsilon k \rfloor$, with probability at least $1-\delta$, with the same number of samples.*

Proof: Consider the single quantum sample $|\psi_{S,f}\rangle$ for the Quantum Padded Coupon Collector problem, with an additional qubit in the state $|1\rangle$. The joint (pure) state is

$$|\psi_{S,f}\rangle|1\rangle \quad = \quad \sum_{i \in S} \frac{1}{\sqrt{S}} |i, f(i)\rangle|1\rangle \ .$$

Consider the distribution $D_{S,f}$ which is uniform over the the k elements $(i, f(i)) \in [n] \times \mathbb{N}_p$ given by $i \in S$, and the concept g_S. The state $|\psi_{S,f}\rangle|1\rangle$ is a quantum sample for the concept g_S with the distribution $D_{S,f}$. Hence, given samples $|\psi_{S,f}\rangle$, we can construct the same number of quantum samples for g_S corresponding to the distribution $D_{S,f}$, and then feed them to the quantum learner $\mathcal{A}$ for $\mathcal{C}_{n,k,p}$.

Since $\mathcal{A}$ is a proper learner, with probability at least $1-\delta$, it produces a concept $g_{S'} \in \mathcal{C}_{n,k,p}$ which is an ϵ approximation of g_S with respect to the distribution $D_{S,f}$. From this, we get

$$\begin{aligned}
& \Pr_{(I,X)\sim D_{S,f}}[g_{S'}(I,X) \neq g_S(I,X)] && \leq && \epsilon \\
\implies & \Pr_{I\sim S}[\mathbb{1}[I \in S'] \neq 1] && \leq && \epsilon \\
\implies & |S \setminus S'| && \leq && \epsilon \cdot k \\
\implies & |S' \setminus S| && \leq && \lfloor \epsilon \cdot k \rfloor \quad = \quad l \ .
\end{aligned}$$

Hence, with probability at least $1-\delta$, the output S' has at most l mismatches with S, and satisfies the requirements of the Quantum Padded Coupon Collector problem. ■

From this, and the bound in Theorem 1.4, the following is immediate.

Corollary 4.3. *For $d \geq 5$ and $k := \frac{d}{5\epsilon}$, the sample complexity of any quantum algorithm for proper (ϵ, δ)-PAC learning the concept class $\mathcal{C}_{n,k,p}$ is at least t_0, where p and t_0 are as in Theorem 1.4. In particular, for failure probability $\delta \leq \frac{1}{8}$ and padding length $p := 64k^{k \ln \frac{k+1}{10l+1}}$, the sample complexity is at least $\Omega\left(k \log \frac{k+1}{l+1}\right)$, i.e., $\Omega(\frac{d}{\epsilon} \log \frac{1}{\epsilon})$.*

Note that if we fix $\delta := \frac{1}{8}$, the lower bound holds for a fixed sufficiently large value of the padding length p, and hence for a fixed concept class $\mathcal{C}_{n,k,p}$ independent of δ. Combined with the general lower bound of $\Omega\left(\frac{1}{\epsilon} \log \frac{1}{\delta}\right)$ from [4, 8], and the general improper (classical) PAC learning algorithm due to Hanneke [9] which has sample complexity $\mathrm{O}\left(\frac{d}{\epsilon} + \frac{1}{\epsilon} \log \frac{1}{\delta}\right)$, we get Theorem 1.3.

References

[1] Anurag Anshu and Srinivasan Arunachalam. A survey on the complexity of learning quantum states. *Nature Reviews Physics*, 6(1):59–69, December 2023.

[2] Srinivasan Arunachalam, Aleksandrs Belovs, Andrew M. Childs, Robin Kothari, Ansis Rosmanis, and Ronald de Wolf. Quantum Coupon Collector. In Steven T. Flammia, editor, *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, volume 158 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:17, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

[3] Srinivasan Arunachalam and Ronald de Wolf. Guest column: A survey of Quantum Learning Theory. *SIGACT News*, 48(2):41–67, June 2017.

[4] Srinivasan Arunachalam and Ronald de Wolf. Optimal quantum sample complexity of learning algorithms. *Journal of Machine Learning Research*, 19(1):2879–2878, January 2018.

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

[6] Nader H. Bshouty and Jeffrey C. Jackson. Learning DNF over the uniform distribution using a quantum example oracle. *SIAM Journal on Computing*, 28(3):1136–1153, 1998.

[7] User "cardinal" (https://stats.stackexchange.com/users/2970/cardinal). What is a tight lower bound on the coupon collector time? Available at https://stats.stackexchange.com/q/7917.

[8] 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, March 2024.

[9] Steve Hanneke. The optimal sample complexity of PAC learning. *Journal of Machine Learning Research*, 17(38):1–15, January 2016.

[10] Steve Hanneke. Proper PAC learning VC dimension bounds. Available at https://cstheory.stackexchange.com/questions/40161/proper-pac-learning-vc-dimension-bounds, July 11, 2019.

[11] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

[12] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, second edition, 2017.

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

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

[15] Vladimir N. Vapnik and Alexey Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

[16] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, May 2018.

A Error due to approximation

In Section 4.2, we approximate an ensemble arising from the Quantum Padded Coupon Collector Problem with another ensemble that is easier to analyse. Here, we bound the error term that arises as a consequence; see Eq. (4.8).

We begin with two claims that are useful in bounding the error. The claims count the number of pairs of sequences $(\boldsymbol{i}, \boldsymbol{x})$ with a given modular signature under different conditions on $\boldsymbol{i}$. (See Eq. (4.4) for the definition of modular signature.)

Lemma A.1. *For any integers $n, t, p \geq 1$, and sequence $\boldsymbol{b} \in \mathbb{N}_p^n$,*

$$\begin{aligned} &\left| \left\{ (\boldsymbol{i}, \boldsymbol{x}) : \boldsymbol{i} \in \operatorname{supp}(\boldsymbol{b})^t, \boldsymbol{x} \in \mathbb{N}_p^t, \operatorname{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b} \right\} \right| \\ &\quad = \quad p^{t - |\operatorname{supp}(\boldsymbol{b})|} \times \left| \left\{ \boldsymbol{i} \in \operatorname{supp}(\boldsymbol{b})^t : \operatorname{range}(\boldsymbol{i}) = \operatorname{supp}(\boldsymbol{b}) \right\} \right| . \end{aligned}$$

Proof: Consider $(\boldsymbol{i}, \boldsymbol{x})$ such that $\boldsymbol{i} \in \operatorname{supp}(\boldsymbol{b})^t$ and $\operatorname{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}$. First, we have $\operatorname{range}(\boldsymbol{i}) \subseteq \operatorname{supp}(\boldsymbol{b})$. Second, if there is some element $q \in [n]$ that does not appear in the sequence $\boldsymbol{i}$, then $\boldsymbol{b}_q = 0$, i.e, $q \notin \operatorname{supp}(\boldsymbol{b})$. So we also have $\operatorname{range}(\boldsymbol{i}) \supseteq \operatorname{supp}(\boldsymbol{b})$, and hence $\operatorname{range}(\boldsymbol{i}) = \operatorname{supp}(\boldsymbol{b})$.

For a fixed $\boldsymbol{b}$ and $\boldsymbol{i} \in \mathrm{supp}(\boldsymbol{b})^t$ with $\mathrm{range}(\boldsymbol{i}) = \mathrm{supp}(\boldsymbol{b})$, the equation $\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}$ is a linear equation in the variables $\boldsymbol{x}$. Furthermore, for each $j \in [n]$, $\boldsymbol{b}_j$ is a sum of the variables in $X_j \coloneqq \{\boldsymbol{x}_r : \boldsymbol{i}_r = j\}$. The sets (X_j) form a partition of the t variables in $\boldsymbol{x}$. Let $t_j \coloneqq |X_j|$. Note that the equation with $\boldsymbol{b}_j$ is non-trivial if and only if j appears in $\boldsymbol{i}$, i.e., $j \in \mathrm{range}(\boldsymbol{i})$ and $t_j > 0$. If the equation is non-trivial, it has $p^{t_j - 1}$ solutions. Thus, the total number of solutions $\boldsymbol{x}$ to $\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}$ is

$$p^{\sum_{j=1}^{n} t_j - |\mathrm{range}(\boldsymbol{i})|} \quad = \quad p^{t - |\mathrm{range}(i)|} \quad = \quad p^{t - |\mathrm{supp}(\boldsymbol{b})|} \ .$$

Multiplying this with the number of $\boldsymbol{i}$'s with $\mathrm{range}(\boldsymbol{i}) = \mathrm{supp}(\boldsymbol{b})$, we get the claimed identity. ∎

Lemma A.2. *For any integers $n, t, p \geq 1$, set $S \subseteq [n]$, and sequence $\boldsymbol{b} \in \mathbb{N}_p^n$,*

$$\left| \left\{ (\boldsymbol{i}, \boldsymbol{x}) : \boldsymbol{i} \in S^t, \boldsymbol{x} \in \mathbb{N}_p^t, \mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b} \right\} \right| \quad = \quad \sum_{T \,:\, \mathrm{supp}(\boldsymbol{b}) \subseteq T \subseteq S} p^{t - |T|} \times \left| \left\{ \boldsymbol{i} \in T^t : \mathrm{range}(\boldsymbol{i}) = T \right\} \right| .$$

Proof: The proof is similar to that of the Lemma A.1. We count the total number of pairs $(\boldsymbol{i}, \boldsymbol{x})$ in the set on the LHS. Fix an $\boldsymbol{i} \in S^t$. Call $\mathrm{range}(\boldsymbol{i})$ as T. By the same reasoning as in the proof of Lemma A.1, since $\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}$ for some $\boldsymbol{x}$, we have $\mathrm{supp}(\boldsymbol{b}) \subseteq T$. Moreover, there are $p^{t-|T|}$ sequences $\boldsymbol{x}$ with $\mathrm{ms}(\boldsymbol{i}, \boldsymbol{x}) = \boldsymbol{b}$. Thus, summing over the possible values for T, we get the desired identity. ∎

Recall the unnormalised states $|\phi_{S,\boldsymbol{b}}\rangle$ and $|\phi_{\boldsymbol{b}}\rangle$ defined in Section 4, where $S \subset [n]$ and $\boldsymbol{b} \in \mathbb{N}_p^n$ is a modular signature of length t sequences. We show that the normalised states are close to each other in trace distance, provided the padding length p is sufficiently large.

Lemma A.3. *For $S \subset [n]$ and $\boldsymbol{b} \in \mathbb{N}_p^n$ such that $\mathrm{supp}(\boldsymbol{b}) \subseteq S$,*

$$\left\| \frac{|\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}|}{\|\phi_{\boldsymbol{b}}\|^2} - \frac{|\phi_{S,\boldsymbol{b}}\rangle\langle\phi_{S,\boldsymbol{b}}|}{\|\phi_{S,\boldsymbol{b}}\|^2} \right\|_1 \quad \leq \quad 2\sqrt{\frac{k^t}{p}} \ ,$$

where $k = |S|$, t is the number of samples, and p is the padding length.

Proof: For any two pure states $|\xi\rangle, |\zeta\rangle$ in the same space, we have

$$\| |\xi\rangle\langle\xi| - |\zeta\rangle\langle\zeta| \|_{\mathrm{tr}} \quad = \quad 2\sqrt{1 - |\langle\xi|\zeta\rangle|^2} \ .$$

So,

$$\left\| \frac{|\phi_{\boldsymbol{b}}\rangle\langle\phi_{\boldsymbol{b}}|}{\|\phi_{\boldsymbol{b}}\|^2} - \frac{|\phi_{S,\boldsymbol{b}}\rangle\langle\phi_{S,\boldsymbol{b}}|}{\|\phi_{S,\boldsymbol{b}}\|^2} \right\|_1 \quad = \quad 2\sqrt{1 - \left| \frac{\langle\phi_{\boldsymbol{b}}|\phi_{S,\boldsymbol{b}}\rangle}{\|\phi_{S,\boldsymbol{b}}\| \, \|\phi_{\boldsymbol{b}}\|} \right|^2} \ .$$

Since the amplitudes of the standard basis elements in the states are real, we have

$$\left| \frac{\langle\phi_{\boldsymbol{b}}|\phi_{S,\boldsymbol{b}}\rangle}{\|\phi_{S,\boldsymbol{b}}\| \, \|\phi_{\boldsymbol{b}}\|} \right|^2 \quad = \quad \frac{\langle\phi_{\boldsymbol{b}}|\phi_{S,\boldsymbol{b}}\rangle^2}{\|\phi_{S,\boldsymbol{b}}\|^2 \, \|\phi_{\boldsymbol{b}}\|^2} \quad = \quad \frac{\langle\phi_{\boldsymbol{b}}|\phi_{\boldsymbol{b}}\rangle^2}{\|\phi_{S,\boldsymbol{b}}\|^2 \, \|\phi_{\boldsymbol{b}}\|^2} \quad = \quad \frac{\|\phi_{\boldsymbol{b}}\|^2}{\|\phi_{S,\boldsymbol{b}}\|^2} \ .$$

Define $n_{t,T} \coloneqq \left|\left\{\boldsymbol{i} \in T^t : \mathrm{range}(\boldsymbol{i}) = T\right\}\right|$. By the definition of the states and Lemmas A.1 and A.2 we have

$$\begin{aligned}
\frac{\|\phi_{\boldsymbol{b}}\|^2}{\|\phi_{S,\boldsymbol{b}}\|^2} &= \frac{\left|\left\{(\boldsymbol{i},\boldsymbol{x}) : \boldsymbol{i} \in \mathrm{supp}(\boldsymbol{b})^t, \boldsymbol{x} \in \mathbb{N}_p^t, \mathrm{ms}(\boldsymbol{i},\boldsymbol{x}) = \boldsymbol{b}\right\}\right|}{\left|\left\{(\boldsymbol{i},\boldsymbol{x}) : \boldsymbol{i} \in S^t, \boldsymbol{x} \in \mathbb{N}_p^t, \mathrm{ms}(\boldsymbol{i},\boldsymbol{x}) = \boldsymbol{b}, \right\}\right|} \\
&= \frac{p^{t-|\,\mathrm{supp}(\boldsymbol{b})|}\, n_{t,\mathrm{supp}(\boldsymbol{b})}}{\sum_{T\,:\,\mathrm{supp}(\boldsymbol{b})\subseteq T\subseteq S} p^{t-|T|}\, n_{t,T}} \\
&= \frac{n_{t,\mathrm{supp}(\boldsymbol{b})}}{\sum_{T\,:\,\mathrm{supp}(\boldsymbol{b})\subseteq T\subseteq S} p^{|\,\mathrm{supp}(\boldsymbol{b})|-|T|}\, n_{t,T}} \\
&= \frac{n_{t,\mathrm{supp}(\boldsymbol{b})}}{n_{t,\mathrm{supp}(\boldsymbol{b})} + \sum_{T\,:\,\mathrm{supp}(\boldsymbol{b})\subsetneq T\subseteq S} p^{|\,\mathrm{supp}(\boldsymbol{b})|-|T|}\, n_{t,T}} \\
&\geq \frac{1}{1+\frac{k^t}{p}} \ .
\end{aligned}$$

Thus, the ℓ_1 distance is at most $2\sqrt{1-\frac{1}{1+\frac{k^t}{p}}} \leq 2\sqrt{\frac{k^t}{p}}$. This proves the lemma. ∎

B Properties of the Classical Coupon Collector process

In this section, we derive some properties of the Coupon Collector problem for completeness. These are used in the classical case to show a separation between proper and improper PAC learning, and also turn out to be relevant to the quantum case.

The parameters in the proof of Lemma B.1 below are based on the argument in Ref. [7].

Lemma B.1. *Let $\delta \in [0,1)$ and l be a non-negative integer. Given independent, uniform samples from a set $S \subset [n]$ of size k, at least $k\ln\frac{k+1}{l+1} + k\ln(1-\delta)$ samples are required to observe at least $k-l$ distinct elements of S with probability at least $1-\delta$.*

Proof: Suppose we draw samples from S until $k-l$ distinct elements are observed. Then, the total number of samples is the sum of $k-l$ independent geometric random variables. Let X_i be the number of samples drawn after the $(i-1)$-th distinct element is observed until the i-th distinct element is observed (with the last sample included). Then X_i has a geometric distribution with parameter $\frac{k-i+1}{k}$.

For any $s, \lambda \geq 0$ we have

$$\begin{aligned}
\Pr\left[\sum_{i=1}^{k-l} X_i \leq s\right] &= \Pr\left[\mathrm{e}^{-\lambda s} \leq \prod_{i=1}^{k-l} \mathrm{e}^{-\lambda X_i}\right] \\
&\leq \mathrm{e}^{\lambda s} \prod_{i=1}^{k-l} \mathbb{E}\left[\mathrm{e}^{-\lambda X_i}\right] && \text{(By the Markov Inequality)} \\
&= \mathrm{e}^{\lambda s} \prod_{i=1}^{k-l} \frac{\frac{k-i+1}{k}\,\mathrm{e}^{-\lambda}}{1-\frac{i-1}{k}\,\mathrm{e}^{-\lambda}} \\
&= \mathrm{e}^{\lambda s} \prod_{i=1}^{k-l} \frac{\frac{k-i+1}{k}}{\mathrm{e}^{\lambda}-\frac{i-1}{k}} \ .
\end{aligned}$$

Let $s \coloneqq k \ln \frac{k+1}{l+1} + ck$ for some $c \in \mathbb{R}$. Set $\lambda \coloneqq \frac{1}{k}$ and notice that $\exp\left(\frac{s}{k}\right) = \frac{k+1}{l+1}\,\mathrm{e}^{c}$, and $\exp\left(\frac{1}{k}\right) \geq 1 + \frac{1}{k}$. Using these, we get that

$$\begin{aligned}
\Pr\left[\sum_{i=1}^{k-l} X_i \leq s\right] &\leq \mathrm{e}^{c}\,\frac{k+1}{l+1} \prod_{i=1}^{k-l} \frac{\frac{k-i+1}{k}}{1+\frac{1}{k}-\frac{i-1}{k}} \\
&\leq \mathrm{e}^{c}\,\frac{k+1}{l+1} \prod_{i=1}^{k-l} \frac{k-i+1}{k-i+2} \\
&= \mathrm{e}^{c} \ .
\end{aligned}$$

If $c < \ln(1-\delta)$, the probability that $k-l$ distinct elements are observed with s samples is $< 1-\delta$. The lower bound stated in the lemma follows. ∎

The following lemma bounds the probability of correctly guessing most of the elements of an unknown set, when asked to guess a set of the same size.

Lemma B.2. *Let l, m be positive integers such that $l \geq 10m$. Suppose $T \subseteq [l+5m]$ is an unknown subset of size l. Then, if $T' \subset [l+5m]$ is a uniformly random subset of size l, the probability that it contains at least $l-m$ elements of T is at most $\frac{1}{2}$.*

Proof: For fixed m, the probability in question decreases as l increases, so we prove the result for $l \coloneqq 10m$.

For $i \in [l]$, let X_i be the random variable indicating whether the i-th element of S (say, in increasing order) is contained in T'. Note that $\sum X_i$ is the number of elements of T contained in T'. We have $\mathbb{E}[X_i] = \frac{l}{l+5m} = \frac{2}{3}$, and so $\mathbb{E}\left[\sum_i X_i\right] = \frac{l^2}{l+5m} = \frac{20}{3}m$. By the Hoeffding bound for the hypergeometric series (see Eq. (2.2)),

$$\begin{aligned}
\Pr\left[\sum X_i \geq l-m\right] &= \Pr\left[\sum X_i - \frac{20}{3}m \geq \frac{7}{3}m\right] \\
&\leq \exp\left(-\frac{98}{90}m\right) \ .
\end{aligned}$$

This is at most $\frac{1}{\mathrm{e}} \leq \frac{1}{2}$. ∎

Lemma B.2 helps bound the probability of properly learning the unknown set in the Coupon Collector problem.

Lemma B.3. *Let $\delta \in [0, 1)$. Given independent, uniformly random samples from an unknown set $S \subset [k + 5m]$ with $|S| = k$, at least $k \ln \frac{k+1}{10m+1} + k \ln(1 - 2\delta)$ samples are required in order to learn the set S with at most $m \geq 1$ mismatches, with average probability of failure at most δ when S is chosen uniformly at random.*

Proof: By the Lemma B.1, we know that if the number of samples drawn is strictly smaller than $k \ln \frac{k+1}{10m+1} + k \ln(1 - 2\delta)$, then there is at least 2δ probability of observing at most $k - 10m$ distinct elements from the set S.

Suppose $k - l$ distinct elements were observed for some $l \geq 10m$, since we are interested in decreasing the average probability of failure for uniformly random S, the best strategy is to guess the remaining l elements uniformly at random. By Lemma B.2, the probability that we produce a guess of size k with up to m mismatches is at most $\frac{1}{2}$. Hence, with probability at least $2\delta \times \frac{1}{2}$, we fail to learn S up to m mismatches. ■