

UC San Diego

UC San Diego Previously Published Works

Title

The geometry of quantum learning

Permalink

<https://escholarship.org/uc/item/47w3t4b5>

Journal

Quantum Information Processing, 9(3)

ISSN

1573-1332

Authors

Hunziker, Markus

Meyer, David A.

Park, Jihun

et al.

Publication Date

2010-06-01

DOI

10.1007/s11128-009-0129-6

Peer reviewed

The geometry of quantum learning

Markus Hunziker · David A. Meyer ·
Jihun Park · James Pommersheim ·
Mitch Rothstein

Published online: 23 September 2009

© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract Concept learning provides a natural framework in which to place the problems solved by the quantum algorithms of Bernstein-Vazirani and Grover. By combining the tools used in these algorithms—quantum fast transforms and amplitude amplification—with a novel (in this context) tool—a solution method for geometrical optimization problems—we derive a general technique for quantum concept learning. We name this technique “Amplified Impatient Learning” and apply it to construct quantum algorithms solving two new problems: BATTLESHIP and MAJORITY, more efficiently than is possible classically.

M. Hunziker · D. A. Meyer (✉) · J. Pommersheim
Project in Geometry and Physics, Department of Mathematics, University of California/San Diego,
La Jolla, CA 92093-0112, USA
e-mail: dmeyer@math.ucsd.edu

M. Hunziker · J. Park · M. Rothstein
Department of Mathematics, University of Georgia, Athens, GA 30602-7403, USA

M. Rothstein
e-mail: rothstei@math.uga.edu

M. Hunziker
Department of Mathematics, Baylor University, Waco, TX 76798-7328, USA
e-mail: Markus_Hunziker@baylor.edu

J. Park
Department of Mathematics, Pohang University of Science and Technology,
Pohang, Kyungbuk 790-784, Korea
e-mail: wlog@postech.ac.kr

J. Pommersheim
Department of Mathematics, Reed College, Portland, OR 97202-8199, USA
e-mail: jamie@reed.edu

Keywords Quantum algorithms · Procrustes problem

PACS 03.67.Ac

Mathematical Subject Classification (2010) 81P68 · 68Q12 · 68Q32 · 15A60

1 Introduction

Over the past decade increasing numbers of scientists have built quantum computation into an imposing edifice. The paucity of quantum algorithms, however, betrays a certain emptiness at its center. Only a handful of problems are known to be solvable more efficiently quantum mechanically than classically, and even fewer general quantum algorithmic techniques are known. The latter include quantum fast transforms [1–7] and amplitude amplification [8–13]. In this paper we explain how to combine these techniques with a new (in this context) one—a solution method for geometrical optimization problems—into quantum algorithms that solve new classes of problems.

These new problems can be thought of as generalizations of the structured and unstructured search problems solved by the quantum algorithms of Bernstein and Vazirani [14, 15] and Grover [8, 9]. Our thinking, however, is largely informed by a branch of classical artificial intelligence—machine learning (For the computer science perspective) [16], and (for a recent mathematical perspective) [17], or more specifically, computational learning theory [18].

In this subject, a *concept* is a map $c : X \rightarrow \mathbb{Z}_2$, defined on some discrete set X ; the support of the function, $c^{-1}(1) \subset X$, is the *extension* of the concept. For example, let X be the set of all balloons, and define $c(x) = 1$ if and only if $x \in X$ is red; this concept is “red balloon”. *Concept learning* is the process by which a *student* (the learner) identifies (or approximates) a target concept from a *concept class* \mathcal{C} of possible concepts. Learning can be *passive*—in situations where examples $x \in X$ are presented to the student by some external mechanism, or *active*—in situations where the student can query a *teacher* for information about the target concept. In the latter case, Angluin has defined a *minimally adequate teacher* to consist of a *pair* of oracles: a *membership oracle* that responds to a query $x \in X$ with $\bar{c}(x)$, where $\bar{c} \in \mathcal{C}$ is the target concept; and an *equivalence oracle* that responds to a query $c \in \mathcal{C}$ with $\delta_{c\bar{c}}$ [19].

The number of queries made by a learning algorithm is the *query complexity* of the algorithm; the number of queries to the membership oracle is its *sample complexity*. These are distinct from the *computational complexity* of the algorithm, which is defined in the usual way [20]. A *family* of concept classes \mathcal{C}_i , for $0 < i \in \mathbb{Z}$, is an infinite sequence of concept classes defined on a corresponding sequence of sets X_i . A learning algorithm for such a family is a sequence of learning algorithms, one for each \mathcal{C}_i . Since each algorithm in the sequence has a sample complexity, we can discuss the asymptotic sample complexity of the family. As we describe in detail in §2, both Bernstein and Vazirani’s and Grover’s algorithms can be interpreted as quantum algorithms for concept learning from a membership oracle, each with a sample complexity

that is asymptotically smaller than the sample complexity of the best possible classical learning algorithm for the same problem.

Bernstein and Vazirani’s algorithm is particularly striking because it requires only a single query to the membership oracle to learn any concept in the problem class. Only very special concept learning problems have quantum sample complexity 1 in this sense. In §2 we explain that these are learning problems in what should be described as “Hadamard” concept classes. Other learning problems, like the one solved by Grover’s algorithm, have quantum sample complexity greater than 1, but one can ask how well a student can learn with a single query. In §3 we pose this problem of “impatient learning” precisely, and show that it is answered by the solution to a certain geometric optimization problem.

With additional queries we should expect superior results. In §4 we show that the quantum computing technique of amplitude amplification [8–13] corresponds to querying also the other half of a minimally adequate teacher, the equivalence oracle. Using an equivalence oracle we can define a general quantum learning algorithm, but without the use of some structure in—or symmetry of—the concept class, it is precisely Grover’s algorithm, with the queries interpreted as being to the equivalence oracle, rather than to the membership oracle. In §5 we review group algebras, in order to describe particular symmetries of concept classes. These symmetries—via quantum fast transforms—allow equivalence queries to be combined with optimal impatient learning algorithms to achieve performance superior to use of equivalence queries alone. In §6 and §7 we analyze the resulting quantum algorithms for concept classes with \mathbb{Z}_N and \mathbb{Z}_2^n symmetry, respectively. We obtain efficient quantum algorithms for two novel problems: BATTLESHIP and MAJORITY.

We conclude in §8 with a discussion of the optimality of our quantum algorithms, and their relevance to a pair of conjectured upper bounds for the sample complexity of quantum learning algorithms.

2 Formalization of quantum learning algorithms

Bernstein and Vazirani’s search problem is the task of identifying $a \in \mathbb{Z}_2^n$, given a ‘sophisticated’ oracle that returns $a \cdot x \bmod 2$ when queried about $x \in \mathbb{Z}_2^n$ [21]. From our point of view, it can also be interpreted as an instance of active learning with access to a membership oracle. There is a family of concept classes \mathcal{BV}^n for $0 < n \in \mathbb{Z}$, with

$$\mathcal{BV}^n = \{p_a : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2 \mid p_a(x) = a \cdot x \bmod 2 \text{ for } a \in \mathbb{Z}_2^n\},$$

consisting of the concepts “bit string with odd inner product with a ” for $a \in \mathbb{Z}_2^n$. Since the concept class \mathcal{BV}^n is parameterized by $a \in \mathbb{Z}_2^n$, identifying a is equivalent to learning a target concept p_a by querying a membership oracle. Classically this learning problem has sample complexity $\Omega(n)$.

In Bernstein and Vazirani’s quantum algorithm for this problem, as well as in all the quantum concept learning algorithms we consider in this paper, the “data structure” consists of a query “register” and a response “register”—the Hilbert space of states is $\mathbb{C}^{|\mathbb{Z}_2^n|} \otimes \mathbb{C}^2$. A membership oracle for target concept \bar{c} acts via the unitary

transformation $U_{\bar{c}}$ defined by linear extension from its action on the *computational basis*, $\{|x, b\rangle \mid x \in X, b \in \mathbb{Z}_2\}$, namely $U_{\bar{c}}|x, b\rangle = |x, b + \bar{c}(x)\rangle$, where “+” denotes addition modulo 2. Let $|-\rangle = (|0\rangle - |1\rangle) / \sqrt{2}$. Then $U_{\bar{c}}|x\rangle|-\rangle = (-1)^{\bar{c}(x)}|x\rangle|-\rangle$. We will use this “phase kickback” trick [22] throughout, so we need only concentrate on the query register and, abusing notation slightly, write $U_{\bar{c}}|x\rangle = (-1)^{\bar{c}(x)}|x\rangle$.

With this notation, Bernstein and Vazirani’s algorithm is summarized by the equation:

$$H^{\otimes n} U_{p_a} H^{\otimes n} |0\rangle = |a\rangle, \tag{2.1}$$

where $H = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} / \sqrt{2}$ and $0 \in \mathbb{Z}_2^n$. That is, from an initial state $|0\rangle$, we apply the *Hadamard transform*, $H^{\otimes n}$; query the membership oracle; and apply the Hadamard transform again. The result is the state $|a\rangle$, so a measurement in the computational basis identifies the target concept p_a with probability 1. The quantum sample complexity of this algorithm is 1, a substantial improvement over the classical sample complexity.

To understand why this algorithm works, notice that after the first Hadamard transform in (2.1), the state of the query register is in an equal superposition of all possible queries:

$$H^{\otimes n} |0\rangle = \frac{1}{\sqrt{2^n}} \sum_{x \in \mathbb{Z}_2^n} |x\rangle.$$

Such an equal superposition is the state before the initial query in each of the quantum algorithms we discuss in this paper. Acting on this state by U_{p_a} produces one of $|\mathbb{Z}_2^n| = 2^n$ possible vectors, according to the value of a . Let $A_{\mathcal{B}\mathcal{V}}$ be the matrix that has these vectors as columns. In general we make the following definition.

Definition For any concept class \mathcal{C} defined over a set X , define the *membership query matrix* $A_{\mathcal{C}}$ to be the $|X| \times |\mathcal{C}|$ matrix with c th column

$$U_c \frac{1}{\sqrt{|X|}} \sum_{x \in X} |x\rangle,$$

for $c \in \mathcal{C}$. In this paper we only consider concept classes for which there is a bijection between X and \mathcal{C} ; we call these *matched* concept classes. For matched concept classes, the membership query matrix is square.

For the Bernstein and Vazirani problem, the membership query matrix has entries

$$(A_{\mathcal{B}\mathcal{V}})_{xa} = \frac{(-1)^{x \cdot a}}{\sqrt{2^n}},$$

which we recognize as the entries of $H^{\otimes n}$. Thus the final Hadamard transform in (2.1) acts as

$$H^{\otimes n} (A_{\mathcal{B}\mathcal{V}})_a = (H^{\otimes n} A_{\mathcal{B}\mathcal{V}})_a = (H^{\otimes n} H^{\otimes n})_a = (I)_a = |a\rangle, \tag{2.2}$$

since $H = H^{-1}$. That is, it inverts the query matrix. Clearly, then, the sample complexity of quantum learning in any concept class with a unitary membership query matrix is 1. Since such a membership query matrix is just a Hadamard matrix in the traditional sense (an orthogonal matrix with entries ± 1) [23, 24], normalized by $\sqrt{|X|}$, we refer to such concept classes as *Hadamard* concept classes.¹

Not all learning problems, of course, are this easy. Grover’s search problem can also be interpreted as an instance of active concept learning with access to a membership oracle. In this case there is a family $\{\mathcal{G}_N\}$ of concept classes, for $0 < N \in \mathbb{Z}$, with

$$\mathcal{G}_N = \{\delta_a : \mathbb{Z}_N \rightarrow \mathbb{Z}_2 \mid \delta_a(x) = \delta_{ax} \text{ for } a \in \mathbb{Z}_N\},$$

which consists of the concepts “is the number a ” for $a \in \mathbb{Z}_N$. The task is to identify a given an oracle that returns δ_{ax} when queried about x . Since the concept class \mathcal{G}_N is parameterized by $a \in \mathbb{Z}_N$, identifying a is equivalent to identifying a target concept δ_a . Classically this learning problem has sample complexity $\Omega(N)$.

Quantum mechanically, this oracle acts by a unitary matrix U_{δ_a} , so the membership query matrix for this problem has entries

$$(A_G)_{xa} = \frac{(-1)^{\delta_{xa}}}{\sqrt{N}} = \frac{1}{\sqrt{N}} \left(NF^\dagger |0\rangle\langle 0| F - 2I \right)_{xa}, \tag{2.3}$$

where F is the N -dimensional discrete Fourier transform. Clearly A_G is not unitary, so \mathcal{G}_N is not a Hadamard concept class, and a single query does not suffice to learn a target concept. In fact, Bennett, Bernstein, Brassard, and Vazirani [26] showed that (in our language) the sample complexity of Grover’s learning problem is $\Omega(\sqrt{N})$. Nevertheless, one might ask how well it is possible to do with a single query. That is, if we can make any unitary transformation (independent of a) after a single query, how do we maximize the probability that a measurement in the *computational basis* $\{|x\rangle \mid x \in \mathbb{Z}_N\}$ returns a ? We give a general solution to this problem of *impatient learning* in the next section, and then apply it to Grover’s problem in §5.

3 Impatient learning

The column vectors of a membership query matrix—the possible states of the query register after a single equal superposition membership query—form a special case of a general situation we can consider, namely a quantum system whose state is one of a set of $0 < N \in \mathbb{Z}$ unit vectors $\{|v_i\rangle \mid i \in \mathbb{Z}_N\}$ in an N -dimensional Hilbert space, \mathcal{H} . The task is to select a measurement to perform that will maximize the probability of correctly guessing which state the system was in before the measurement was made. This is a special case of the problem originally considered by Helstrom [27] and Kholevo [28], *quantum hypothesis testing*, namely identifying one from among a set of pure quantum states, no matter their provenance.

¹ This nomenclature is motivated by van Dam’s paper on a quantum algorithm for the quadratic residue problem [25].

Recall that a von Neumann measurement [29] is defined by an orthogonal direct sum decomposition of the Hilbert space. The measurement is *complete* if the summands are one-dimensional. Belavkin [30] and Kennedy [31] have shown that when the $\{|v_i\rangle\}$ are linearly independent the optimal quantum measurement is, in fact, a complete von Neumann measurement. Such a measurement determines an orthonormal basis $\{|e_i\rangle \mid i \in \mathbb{Z}_N\}$, up to phases. The probability that the system will be in state $|e_i\rangle$ after this measurement, given that it was in state $|v\rangle$ before the measurement, is $|\langle v|e_i\rangle|^2$. If we assume that the system has been prepared in one of the states $\{|v_i\rangle\}$, chosen uniformly at random, then the quantity we want to maximize is

$$\sum_{i=1}^N |\langle v_i|e_i\rangle|^2. \tag{3.1}$$

Necessary and sufficient criteria for solutions to this optimization problem, in the more general case of arbitrary prior probabilities for the $\{|v_i\rangle\}$, can be found in the early quantum hypothesis testing literature [28, 30, 32]. In the following we provide a brief, geometrical derivation of such a criterion.

We can phrase this problem as a question about matrices: If we choose an isomorphism of Hilbert spaces, $\mathcal{H} \simeq \mathbb{C}^N$, then the list $(|v_1\rangle, \dots, |v_N\rangle)$ is identified with a square matrix $A \in M_N(\mathbb{C})$. Making an arbitrary complete measurement is equivalent to making an arbitrary unitary transformation, followed by a fixed complete measurement in, for example, the computational basis. Thus we should consider the matrices SA , for $S \in U(N)$, where $U(N)$ denotes the unitary group. We write $A \sim B$ if $B = SA$ for some $S \in U(N)$. Maximizing the quantity (3.1) is equivalent to maximizing the quantity

$$\|d(B)\|^2 \tag{3.2}$$

over the $U(N)$ -orbit of A , $\{B \mid B \sim A\}$, where $d : M_N(\mathbb{C}) \rightarrow M_N(\mathbb{C})$ is projection onto diagonal matrices and $\|\cdot\|$ is the L^2 (or Frobenius) norm. In the following, when we speak of critical points of the function (3.2), it will be implicit that the $U(N)$ -orbit of A is the domain. We have the following characterization of the critical points, which was stated and proved (differently) by Helstrom, in the more general setting of an arbitrary probability distribution over the state vectors [32, Chap. IV, eq. (1.30)]:

Proposition 3.1 *The matrix B is a critical point of $\|d(B)\|^2$ if and only if $Bd(B)^\dagger$ is Hermitian, i.e.,*

$$Bd(B)^\dagger = d(B)B^\dagger,$$

where † denotes the adjoint.

Proof Let $\mathfrak{u}(N)$ denote the Lie algebra of $U(N)$, i.e., the set of skew-Hermitian matrices. The criticality condition is that for all $\zeta \in \mathfrak{u}(N)$,

$$\frac{d}{dt} \Big|_{t=0} \|d((1 + t\zeta)B)\|^2 = 0,$$

which is true when $\operatorname{Re}(\operatorname{tr}(d(\zeta B)d(B)^\dagger)) = 0$. For all matrices X and Y , one has

$$\operatorname{tr}(d(X)d(Y)) = \operatorname{tr}(Xd(Y)).$$

Therefore

$$\operatorname{Re}\left(\operatorname{tr}\left(d(\zeta B)d(B)^\dagger\right)\right) = \operatorname{Re}\left(\operatorname{tr}\left(\zeta Bd(B)^\dagger\right)\right),$$

so the condition for B to be critical is that $Bd(B)^\dagger$ be orthogonal to all skew-Hermitian matrices, with respect to the inner product $\operatorname{Re}(\operatorname{tr}((\cdot)(\cdot)^\dagger))$. This proves the proposition, since the orthogonal complement to the space of skew-Hermitian matrices is the space of Hermitian matrices. \square

Since $\|d(B)\|^2$ is invariant under left multiplication by unitary diagonal matrices, we can restrict our attention to those critical points of (3.2) that have nonnegative real entries on the diagonal. Now the criticality condition reads

$$Bd(B) = d(B)B^\dagger. \tag{3.3}$$

We would like, however, an explicit solution to (3.3). Consider the Gram matrix of A , $G = A^\dagger A$, with components $G_{ij} = \langle v_i | v_j \rangle$. G is a positive semi-definite Hermitian matrix. Let \sqrt{G} denote the positive semi-definite Hermitian square root of G . By the polar decomposition of A , there is always a unitary matrix S such that $\sqrt{G} = SA$, so it is natural to ask whether $B = \sqrt{G}$ solves (3.3). Proposition 3.1 shows that this is generally not the case. More precisely, we have the following corollary:

Corollary 3.2 *\sqrt{G} is a critical point of (3.2) if and only if \sqrt{G} commutes with its own diagonal.*

If the off-diagonal part of \sqrt{G} is sufficiently general then the conclusion of Corollary 3.2 will force the diagonal to be constant. Although this is a strong condition in general, it is a very natural simplification [28,30,32,33]. We shall see that it occurs in many structured learning problems. Moreover, having a constant diagonal is precisely the condition needed to go beyond impatient learning—which we will do in the next section. So it is a case worth considering.

Proposition 3.3 *Let G be a positive semi-definite Hermitian matrix. Let \sqrt{G} denote the positive semi-definite Hermitian square root of G . Assume the diagonal of \sqrt{G} is constant. Let S denote the set of matrices B such that $B \sim \sqrt{G}$ and B has constant diagonal. Then the maximum of $\|d(B)\|^2$ over $B \in S$ occurs at \sqrt{G} .*

Proof If B has constant diagonal, then $\|d(B)\|^2 = |\operatorname{tr}(B)|^2/N$. So it suffices to prove that \sqrt{G} gives the maximum value of $|\operatorname{tr}(B)|^2$ over all $B \in S$. As in the proof of Proposition 3.1, the critical points occur when $\operatorname{Re}(\operatorname{tr}(B)\operatorname{tr}(\zeta B)) = 0$ for all $\zeta \in \mathfrak{u}(N)$. Writing $\operatorname{tr}(B)\operatorname{tr}(\zeta B) = \operatorname{tr}(\zeta B \operatorname{tr}(B))$, we see that the critical points are given by the condition that $B \operatorname{tr}(B)$ is Hermitian, also as in the proof of Proposition 3.1. Let

$B_h = (B + B^\dagger)/2$ and $B_s = (B - B^\dagger)/2$. We want the skew-Hermitian part of $B \operatorname{tr}(B)$ to vanish, thus

$$B_h \operatorname{tr}(B_s) + B_s \operatorname{tr}(B_h) = 0. \quad (3.4)$$

The trace of (3.4) shows that either $\operatorname{tr}(B_h) = 0$ or $\operatorname{tr}(B_s) = 0$. If both traces vanish then we get the minimum possible value, $|\operatorname{tr}(B)|^2 = 0$. If this is also the maximum then \sqrt{G} is forced to vanish since it is positive semi-definite, so the statement is true in this case. If only one of the traces vanishes, the maximum occurs at a point where $B_s = 0$ or $B_h = 0$. Since multiplication by i is a symmetry of $|\operatorname{tr}(B)|^2$, we may assume $B_s = 0$. Then B is some square root of G . The maximum of $|\operatorname{tr}(B)|^2$ will occur when one chooses the same sign for the square root of each eigenvalue, e.g., when $B = \sqrt{G}$. \square

Remark As we have noted above, both Proposition 3.1 and Proposition 3.3 have long been known in the context of quantum hypothesis testing. Nevertheless, we have included our proofs of these results in order to emphasize the connection with a similar optimization problem: These new proofs are inspired by the proof of the result that the minimum of the L^2 distance $\|B - I\|$ over the set $B \sim \sqrt{G}$ is given by $B = \sqrt{G}$, irrespective of any assumption about the diagonal. In particular, for an arbitrary invertible matrix A , with polar decomposition $A = S^{-1}P$, where S is unitary and P is positive definite Hermitian, the closest point to I in the $U(N)$ -orbit of A is P —this is the solution to the *Procrustes Problem* [34]. Most recently, Eldar and Forney have noted that (in our notation) when the diagonal of \sqrt{G} is constant, $B = \sqrt{G}$ not only minimizes the L^2 distance $\|B - I\|$ over the \sqrt{G} orbit, but also maximizes $\|d(B)\|^2$ over the orbit, and is thus the solution to the optimization problem (3.1) that is the relevant one for quantum measurement [35].

Thus we have the following quantum algorithm for a concept learning problem with invertible membership query matrix A_C :

Impatient learning

1. Prepare the query register in the equal superposition state, $F^\dagger|0\rangle$, where F is the $|X|$ -dimensional discrete Fourier transform. (Any unitary map taking $|0\rangle$ to the equal superposition state works; in the case where $X = \mathbb{Z}_2^n$, $H^{\otimes n}$ can be applied.)
2. Query the membership oracle, obtaining as the state the \bar{c} th column of A_C , $U_{\bar{c}}F^\dagger|0\rangle$.
3. Apply the unitary transformation $S_C = (A_C^\dagger A_C)^{1/2} A_C^{-1}$.
4. Measure the resulting state $S_C U_{\bar{c}} F^\dagger|0\rangle$ in the computational basis.

Notice that $B_C = S_C A_C$ satisfies (3.3). As a consequence of the preceding discussion we have:

Theorem 3.4 *For concept \bar{c} , Impatient Learning succeeds with probability $|(B_C)_{\bar{c}\bar{c}}|^2$ and, when these diagonal elements are all equal, is optimal (in the sense of maximizing (3.1)) among single query quantum algorithms that begin with an equal superposition over membership queries.²*

As we saw in (2.2), for $A_{\mathcal{B}\mathcal{V}}, B_{\mathcal{B}\mathcal{V}} = H^{\otimes n} A_{\mathcal{B}\mathcal{V}} = I$ maximizes (3.1), so the Bernstein–Vazirani algorithm is Impatient Learning, and succeeds with probability 1 for every target concept. Furthermore, as we will see in §5, for Grover’s problem, (3.1) is maximized by $B_G = (2F^\dagger|0\rangle\langle 0|F - I) A_G$. Using (2.3) it is then easy to compute that the diagonal entries of B_G are all $(3 - 4/N)/\sqrt{N}$, so Impatient Learning succeeds with asymptotic probability $9/N$ as $N \rightarrow \infty$. Theorem 3.4 says that this is the best we can do using only a single membership query. Although it is certainly an improvement over the success probability $1/N$ of random guessing, Impatient Learning is far from satisfactory for this problem.

4 Beyond impatient learning

In fact, for most concept learning problems a single membership query simply does not provide enough information to learn the target concept with probability close to 1. A specific target concept \bar{c} defines a subspace of the Hilbert space $\mathbb{C}^{|X|}$, namely $\text{span}\{|\bar{c}\rangle\}$ (recall that there is a bijection between X and \mathcal{C}), however, so we can apply one of the few general quantum algorithm techniques—amplitude amplification [8–13]. This technique, invented by Brassard and Høyer [10] as a generalization of Grover’s algorithm [8, 9], can be described in terms of concepts:

AMPLITUDE AMPLIFICATION ([13], Theorem 2). *Let χ be a concept over X ; let \mathcal{H}_1 denote the subspace of $\mathbb{C}^{|X|}$ spanned by the vectors labeled by the elements in the extension of χ , $\chi^{-1}(1)$; and let Π denote the projection $\mathbb{C}^{|X|} \rightarrow \mathcal{H}_1$. For any unitary transformation W of $\mathbb{C}^{|X|}$, let $p(W) = |\Pi W|0\rangle|^2$ be the probability with which the state $W|0\rangle$ is measured to be in the subspace \mathcal{H}_1 . As long as $p(W) > 0$, we can set $\sin^2 \theta = p(W)$ for $0 < \theta \leq \pi/2$. In this case, repeatedly applying the unitary transformation $WU_{\delta_0}W^\dagger U_\chi$ amplifies the probability of measuring the state to be in the subspace \mathcal{H}_1 . More precisely,*

$$p\left((WU_{\delta_0}W^\dagger U_\chi)^m W\right) \geq \max\{1 - p(W), p(W)\},$$

where $m = \lfloor \frac{\pi}{4\theta} - \frac{1}{2} \rfloor$, the nearest integer to $\frac{\pi}{4\theta} - \frac{1}{2}$.

After step 3 of Impatient Learning, the state is $S_C U_{\bar{c}} F^\dagger |0\rangle$, where S_C was chosen to maximize $\sum_{\bar{c}} |\langle \bar{c} | S_C U_{\bar{c}} F^\dagger |0\rangle|^2$. Thus, letting $\chi = \delta_{\bar{c}}$, $\mathcal{H}_1 = \text{span}\{|\bar{c}\rangle\}$ and $W_C = S_C U_{\bar{c}} F^\dagger$, applying Amplitude Amplification gives a new quantum algorithm:

² Even when the diagonal of $B_C = (A_C^\dagger A_C)^{1/2}$ is not constant, a result of Barnum and Knill [36] implies that the success probability is within a factor of 2 of the optimum.

Amplified impatient learning

1. Prepare the query register in the equal superposition state, $F^\dagger|0\rangle$, where F is the $|X|$ -dimensional discrete Fourier transform. (Any unitary map taking $|0\rangle$ to the equal superposition state works; in the case where $X = \mathbb{Z}_2^n$, $H^{\otimes n}$ can be applied.)
2. Query the membership oracle, obtaining as the state the \bar{c} th column of A_C , $U_{\bar{c}}F^\dagger|0\rangle$.
3. Apply an impatient learning transform S_C , producing $S_C U_{\bar{c}}F^\dagger|0\rangle = W_C|0\rangle$.
4. Apply $W_C U_{\delta_0} W_C^\dagger U_{\delta_{\bar{c}}}$ m times, where $m = \lfloor \frac{\pi}{4\theta} - \frac{1}{2} \rfloor$, and $\sin \theta = |\langle \bar{c} | W_C | 0 \rangle| = |(B_C)_{\bar{c}\bar{c}}|$, with $0 < \theta \leq \frac{\pi}{2}$.
5. Measure the resulting state in the computational basis.

As a consequence of Theorem 3.4 and Amplitude Amplification we have:

Theorem 4.1 *For problems with B_C having constant diagonal element s , Amplified Impatient Learning succeeds with probability at least $\max\{1 - s^2, s^2\}$. Since each of W_C and W_C^\dagger includes calls to the membership oracle via $U_{\bar{c}}$, Amplified Impatient Learning has sample complexity $2m + 1$, i.e., $O(1/s)$.*

Notice, however, that the algorithm uses more than membership queries. The operation $U_{\delta_{\bar{c}}}$ in step 4 is the action of an *equivalence* oracle responding to a queried *concept* (rather than a concept argument). Roughly speaking, the Impatient Learning part of this algorithm maximizes the amplitude for the target concept \bar{c} after a single membership query; then an equivalence oracle is queried about the correctness of this concept. Thus Amplified Impatient Learning uses *both* oracles comprising the minimally adequate teacher defined in §1, making $2m + 1$ membership queries and m equivalence queries.

At the risk of confusing the membership and equivalence oracles, we can apply Amplified Impatient Learning to Grover’s problem. As we noted at the end of §3 and as we will compute in §5, for A_G the post-membership query transform is $S_G = 2F^\dagger|0\rangle\langle 0|F - I = F^\dagger U_{\delta_0} F$. So for this problem,

$$W_G = S_G U_{\bar{c}} F^\dagger = F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} F^\dagger,$$

where this use of $U_{\delta_{\bar{c}}}$ is a query to the membership oracle. The iterated transformation is

$$\begin{aligned} W_G U_{\delta_0} W_G^\dagger U_{\delta_{\bar{c}}} &= (F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} F^\dagger) U_{\delta_0} (F U_{\delta_{\bar{c}}} F^\dagger U_{\delta_0} F) U_{\delta_{\bar{c}}} \\ &= F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} \cdot F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} \cdot F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} \\ &= (F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}})^3, \end{aligned}$$

where the $U_{\delta_{\bar{c}}}$ in the first expression is the operation of the equivalence oracle but the distinction between the two kinds of oracles is ignored in the last expression. The complete algorithm is then

$$(F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}})^{3m} F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}} F^\dagger |0\rangle = (F^\dagger U_{\delta_0} F U_{\delta_{\bar{c}}})^{3m+1} F^\dagger |0\rangle,$$

where $m = \lfloor \frac{\pi}{4\theta} - \frac{1}{2} \rfloor$, and $\theta = \arcsin |(B_G)_{\bar{c}\bar{c}}| = \arcsin \left((3 - 4/N)/\sqrt{N} \right)$. Thus $m \sim \frac{\pi}{4} \sqrt{N}/3$ so the iterated transformation is applied $\frac{\pi}{4} \sqrt{N}$ times, asymptotically. This is, in fact, exactly Grover’s algorithm [8,9], although one usually sees it factored differently (and with F and F^\dagger replaced by $H^{\otimes n}$).

5 Concept classes with group symmetry

The sets \mathbb{Z}_2^n and \mathbb{Z}_N , over which the Bernstein-Vazirani and Grover concept classes are defined, are abelian groups under componentwise addition modulo 2 and addition modulo N , respectively. In each case the Hilbert space \mathbb{C}^G becomes a ring, with multiplication law defined by linear extension from

$$|x\rangle * |y\rangle = |x + y\rangle \quad \text{for } x, y \in G,$$

where G is \mathbb{Z}_2^n or \mathbb{Z}_N .

Definition The *group algebra* of G is the Hilbert space \mathbb{C}^G (often written $\mathbb{C}[G]$), equipped with this ring structure.

The *regular* representation of the group algebra is the map

$$\mathbb{C}^G \ni |v\rangle \mapsto L_v \in M_{|G|}(\mathbb{C}),$$

where L_v is left multiplication by $|v\rangle$, a linear map on \mathbb{C}^G , hence a $|G| \times |G|$ complex matrix in the computational basis. We will identify the group algebra with its image in this representation.

For $G = \mathbb{Z}_2$, $\mathbb{C}^2 \ni |v\rangle = \alpha|0\rangle + \beta|1\rangle$ is a general element of the group algebra. From the definition,

$$\begin{aligned} |v\rangle * |0\rangle &= \alpha|0\rangle + \beta|1\rangle \\ |v\rangle * |1\rangle &= \alpha|1\rangle + \beta|0\rangle, \end{aligned}$$

so

$$L_v = \begin{pmatrix} \alpha & \beta \\ \beta & \alpha \end{pmatrix} = \alpha I + \beta X = \alpha X^0 + \beta X^1, \tag{5.1}$$

where $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = L_1$ is the usual Pauli matrix. More generally, we have:

Proposition 5.1 *The group algebra of \mathbb{Z}_2^n consists of $2^n \times 2^n$ dimensional matrices of the form*

$$L_v = \sum_{x \in \mathbb{Z}_2^n} v_x X^x,$$

for

$$|v\rangle = \sum_{x \in \mathbb{Z}_2^n} v_x |x\rangle \in \mathbb{C}^{\mathbb{Z}_2^n} = (\mathbb{C}^2)^{\otimes n}.$$

(In the expression for L_v , $x \in \mathbb{Z}_2^n$ is a multi-index, i.e., $X^x = X^{x_1 \dots x_n} = X^{x_1} \otimes \dots \otimes X^{x_n}$.) $|v\rangle$ is the first column of L_v ; L_v is symmetric and has constant diagonal. L_v is diagonalized by the Hadamard transform.

Proof By induction on n . That the Hadamard transform diagonalizes the elements of the \mathbb{Z}_2^n group algebra follows from the familiar Pauli matrix identity $Z = H \times H$, where $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$. □

The group \mathbb{Z}_N is generated by the element 1, and for $y \in \mathbb{Z}_N$,

$$L_1 : |y\rangle \mapsto |1 + y\rangle \implies L_1 = \begin{pmatrix} 0 & \dots & 1 \\ 1 & 0 & \\ & \ddots & \ddots & \vdots \\ & & \ddots & 0 \\ & & & 1 & 0 \end{pmatrix}.$$

Analogously to Proposition 5.1, for \mathbb{Z}_N we have:

Proposition 5.2 *The group algebra of \mathbb{Z}_N consists of $N \times N$ dimensional matrices of the form*

$$L_v = \sum_{x \in \mathbb{Z}_N} v_x L_1^x,$$

where

$$|v\rangle = \sum_{x \in \mathbb{Z}_N} v_x |x\rangle \in \mathbb{C}^{\mathbb{Z}_N} = \mathbb{C}^N.$$

$|v\rangle$ is the first column of L_v ; L_v need not be symmetric, but it has constant diagonal. L_v is diagonalized by the N -dimensional discrete Fourier transform.

Proof That the Fourier transform diagonalizes the elements of the \mathbb{Z}_N group algebra follows from the fact that $FL_1F^\dagger = \text{diag}(1, \omega, \omega^2, \dots, \omega^{N-1})$, where $\omega = e^{2\pi i/N}$. □

Propositions 5.1 and 5.2 allow us to characterize useful symmetries of concept classes:

Proposition 5.3 *Let C be a matched concept class over a finite abelian group G . Then A_C is in the group algebra of G if and only if it commutes with the action of G , i.e.,*

$$L_g A_C = A_C L_g, \quad \forall g \in G. \tag{5.2}$$

In components (5.2) becomes

$$(A_C)_{x+g, c+g} = (A_C)_{xc}, \quad \forall g \in G; \tag{5.3}$$

equivalently $c(x)$ is a function of $c - x$.

Proof Equations (5.2) and (5.3) are easily seen to be equivalent by direct computation. Suppose that the matrix A_C belongs to the group algebra of G . Then (5.2) follows from the fact that G is abelian, and hence L_g commutes with L_h for all $g, h \in G$. Conversely, if the matrix A_C satisfies (5.3), then it follows that A_C is determined by its first column (the column labeled by 0, the identity of G), as follows:

$$A_C = \sum_{k \in G} (A_C)_{k0} L_k.$$

Hence A_C belongs to the group algebra. □

In Grover’s problem, A_G satisfies (5.2) and (5.3) for $G = \mathbb{Z}_N$ (and for $G = \mathbb{Z}_2^n$, when $N = 2^n$). Thus this membership query matrix belongs to the group algebra of \mathbb{Z}_N (and of \mathbb{Z}_2^n , when $N = 2^n$) and is diagonalized by F (and by $H^{\otimes n}$, when $N = 2^n$). It is, furthermore, a real symmetric matrix. The following proposition explains how to compute the optimal transformation S required for Impatient Learning and Amplified Impatient Learning in this case.

Proposition 5.4 *Let A be a real, symmetric matrix in the group algebra of \mathbb{Z}_2^n (or \mathbb{Z}_N). Using the Spectral Theorem, define $|A|$ by*

$$|A|v = |\lambda|v,$$

for each eigenvector-eigenvalue pair (v, λ) of A . Then $|A|$ is also an element of the same group algebra. Moreover, the maximum value of $\|d(B)\|^2$ over matrices $B \sim A$ with constant diagonal occurs at $|A|$.

Proof Since A is real and symmetric, A is a square root of its Gram matrix. Conjugation by the appropriate transform ($H^{\otimes n}$ or F) diagonalizes A so $|A|$, having the same eigenvectors, is also an element of the same group algebra as A . Moreover, $|A|$ is the positive semi-definite square root of the Gram matrix. Thus the result follows from Proposition 3.3. □

Proposition 5.4, applied to a symmetric membership query matrix A_C satisfying the conditions of Proposition 5.3, implies that an optimal unitary transformation S_C in the Impatient Learning and Amplified Impatient Learning algorithms satisfies

$$|A_C| = S_C A_C. \tag{5.4}$$

When A_C is nonsingular, S_C is unique and (5.4) implies that

$$S_C = |A_C|A_C^{-1} =: \text{sign}(A_C),$$

where

$$\text{sign}(A)v = \text{sign}(\lambda)v = \frac{\lambda}{|\lambda|}v,$$

for all eigenvector-eigenvalue pairs (v, λ) of A .

To compute S_G for Grover’s concept class we diagonalize A_G :

$$F \frac{1}{\sqrt{N}} \left(N F^\dagger |0\rangle\langle 0| F - 2I \right) F^\dagger = \frac{1}{\sqrt{N}} (N|0\rangle\langle 0| - 2I).$$

This implies that

$$S_G = \text{sign}(A_G) = F^\dagger \text{diag}(1, -1, \dots, -1)F = F^\dagger (2|0\rangle\langle 0| - I) F = -F^\dagger U_{\delta_0} F,$$

which is the promised expression for S_G that we quoted in §3 and §4.

6 Learning problems with cyclic symmetry

Although recognizing Grover’s algorithm as an instance of Amplified Impatient Learning perhaps contributes to a better understanding of this basic quantum algorithm, we would like to apply the general formalism developed in the preceding sections to derive new quantum algorithms. So in this section we consider some new problems with cyclic symmetry.

According to Proposition 5.3, any learning problem with a transitive \mathbb{Z}_N action has the property that the oracle response $c(x)$ depends only on the difference $c - x \pmod N$. Thus we may write $c(x) = \phi(c - x)$ for some function $\phi : \mathbb{Z}_N \rightarrow \mathbb{Z}_2$, whence the membership query matrix A_C is

$$A_C = \frac{1}{\sqrt{N}} \sum_{k \in \mathbb{Z}_N} (-1)^{\phi(k)} L_1^k. \tag{6.1}$$

By Proposition 5.2, A_C is diagonalized by the Fourier transform; hence its eigenvalues are

$$\lambda_j = \frac{1}{\sqrt{N}} \sum_{k \in \mathbb{Z}_N} (-1)^{\phi(k)} \omega^{jk}, \tag{6.2}$$

for $j \in \mathbb{Z}_N$. By Proposition 5.4, the relevant quantity for Impatient Learning is the size of the diagonal elements of $|A_C|$. This matrix is in the \mathbb{Z}_N group algebra, and hence has constant diagonal. Furthermore, the diagonal element s is just the average of the

eigenvalues of $|A_C|$, i.e., the average of the absolute values of the eigenvalues of A_C . By Theorem 4.1, therefore, Amplified Impatient Learning requires $O(1/s)$ queries.

Consider a special class of cyclically symmetric problems, which we call BATTLESHIP, after the Milton Bradley game with the same name. Let $0 \leq r < N/2$. For any $a, x \in \mathbb{Z}_N$, set

$$b_a(x) = \begin{cases} 1 & \text{if } a - x \equiv -r, \dots, r \pmod N; \\ 0 & \text{otherwise.} \end{cases}$$

$d = 2r + 1$ is the length of the battleship, i.e., d counts the number of $x \in \mathbb{Z}_N$ that satisfy $b_a(x) = 1$ for any fixed a .

It turns out that the behavior of BATTLESHIP problems depends on the relative size of d with respect to N . Thus we consider two separate subfamilies of BATTLESHIP: For the problem SMALLSHIP(d), we fix the value of d and let N be arbitrary. For the problem BIGSHIP(α), we again let N be arbitrary, but fix the ratio $\alpha \in (0, 1/2)$ of d to N . That is, we take $d = \lfloor \alpha N \rfloor$.

Theorem 6.1 *For any fixed d , Amplified Impatient Learning solves SMALLSHIP(d) with $O(\sqrt{N})$ queries, which is optimal to within a constant factor. When applied to BIGSHIP(α), however, Amplified Impatient Learning requires $\Omega(\sqrt{N} / \log N)$ queries, which is far from optimal.*

Proof The eigenvalues of A_{BS} for the BATTLESHIP concept class with parameters N and r are

$$\lambda_j = \frac{1}{\sqrt{N}} \left(- \sum_{k=-r}^r \omega^{jk} + \sum_{k=r+1}^{N-r-1} \omega^{jk} \right).$$

In particular, this gives

$$\lambda_0 = \frac{1}{\sqrt{N}}(N - 2d),$$

while for $j > 0$,

$$\lambda_j = - \frac{2}{\sqrt{N}} \frac{\sin(\pi jd/N)}{\sin(\pi j/N)}. \tag{6.3}$$

First, consider the case of SMALLSHIP(d) for fixed d . Since the expression $\sin(\pi j/N)$ in the denominator of (6.3) is bounded above in absolute value by 1, it follows that

$$s\sqrt{N} \geq \frac{2}{N} \sum_{j=1}^{N-1} \left| \sin \frac{\pi jd}{N} \right|.$$

As N tends to infinity, the right hand side approaches the constant value $2 \int_0^1 |\sin d\pi x| dx$. We conclude that $s = \Omega(1/\sqrt{N})$ and hence Amplified Impatient Learning has query complexity $O(\sqrt{N})$.

To see that this is optimal, first note that the equivalence oracle can be simulated using exactly two calls to the membership oracle, since $c = b_a$ if and only if $c(a + r) = 1$ and $c(a + r + 1) = 0$. Hence the equivalence oracle is unnecessary. Now note that any solution to $\text{SMALLSHIP}(d)$ allows one to solve the restricted problem in which the value a is assumed to be divisible by d . This restricted problem is easily seen to be equivalent to an unstructured search on a set of size N/d , which is known to require $\Omega(\sqrt{N/d})$ (as shown in [26]). Since d is constant, we see that for $\text{SMALLSHIP}(d)$, Amplified Impatient Learning is optimal up to a constant factor.

Second, consider $\text{BIGSHIP}(\alpha)$ for fixed α . In this case we claim that $s = O((\log N)/\sqrt{N})$. To see this, note first that λ_0/N is $O(1/\sqrt{N})$. Bounding each of the sines in the numerator of (6.3) by 1, we find that for $j > 0$, λ_j/N is bounded above in absolute value:

$$\frac{1}{N} |\lambda_j| \leq \frac{1}{N\sqrt{N}} \csc \frac{\pi j}{N}.$$

It follows that λ_1/N and λ_{n-1}/N are $O(1/\sqrt{N})$, while the remaining sum

$$\frac{1}{N} \sum_{j=2}^{N-2} |\lambda_j| \leq \frac{1}{N\sqrt{N}} \sum_{j=2}^{N-2} \csc \frac{\pi j}{N} \leq \frac{1}{\sqrt{N}} \int_{1/N}^{1-1/N} \csc \pi x dx = O\left(\frac{\log N}{\sqrt{N}}\right).$$

Thus the number of steps required by Amplified Impatient Learning is $\Omega(\sqrt{N}/\log N)$. To see that this is not an optimal algorithm, consider using Grover’s algorithm to return some x for which $b_a(x) = 1$. This requires $O(\sqrt{N/d})$ quantum queries [37], and narrows the possible answer space to a set of size d . A classical binary search, requiring $\log d$ further (classical) queries can now be used to identify the answer a uniquely. This alternative algorithm solves $\text{BIGSHIP}(\alpha)$ with only $O(\sqrt{N/d} + \log d) = O(\sqrt{1/\alpha} + \log \alpha N) = O(\log N)$ queries, far fewer than the $\Omega(\sqrt{N}/\log N)$ required by Amplified Impatient Learning. \square

7 The Majority problem

The other group algebra introduced in §5 is that of \mathbb{Z}_2^n . In this section we study a novel problem, MAJORITY, that has this symmetry. Fix a positive integer n , and for each $a \in \mathbb{Z}_2^n$ define a function $m_a : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$ by

$$m_a(x) = \begin{cases} 1 & \text{if } \text{wt}(a - x) \leq n/2; \\ 0 & \text{otherwise.} \end{cases}$$

That is, $m_a(x) = 1$ when the bit strings a and x agree in at least as many bits as they disagree. The MAJORITY concept class \mathcal{MAJ}^n is defined to be the set of all functions

m_a , where a is any element of \mathbb{Z}_2^n . It is easy to see that any classical learning algorithm requires at least n queries. We can do better quantum mechanically:

Theorem 7.1 *Amplified Impatient Learning solves MAJORITY with $O(\sqrt{n})$ quantum queries, given access to both the membership oracle and the equivalence oracle.*

If n is an odd integer, the membership query matrix for \mathcal{MAJ}^n contains in its upper left-hand corner (rows and columns labeled $0, \dots, 0$ to $1, \dots, 1$ top-to-bottom and left-to-right, respectively) a $2^{n-1} \times 2^{n-1}$ submatrix proportional to the membership query matrix for \mathcal{MAJ}^{n-1} . Furthermore, if $a, b \in \mathbb{Z}_2^n$ are complementary bit strings then $m_a(x) = 1 - m_b(x)$ for all x , and hence the column in the membership query matrix corresponding to a equals the negative of the column corresponding to b . It follows that if one can learn a concept from \mathcal{MAJ}^{n-1} , then one can learn a concept from \mathcal{MAJ}^n with one additional membership query. Thus, in what follows, we will assume that n is an even integer.

For learning problems with \mathbb{Z}_2^n symmetry we have the following analogues of (6.1) and (6.2): Since $c(x) = \phi(c - x)$ for some $\phi : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2$, the membership query matrix has the form

$$A_C = \frac{1}{\sqrt{2^n}} \sum_{b \in \mathbb{Z}_2^n} (-1)^{\phi(b)} X^b. \tag{7.1}$$

By Proposition 5.1, A_C is diagonalized by the Hadamard transform; hence its eigenvalues are

$$\lambda_c = \frac{1}{\sqrt{2^n}} \sum_{b \in \mathbb{Z}_2^n} (-1)^{\phi(b)} (-1)^{b \cdot c}, \tag{7.2}$$

for $c \in \mathbb{Z}_2^n$. With these preliminaries in place, we can prove Theorem 7.1:

Proof For the concept class \mathcal{MAJ}^n , $\phi(b) = \Theta(\frac{n}{2} - \text{wt}(b))$, where the Heaviside function $\Theta(z) = 1$ if $z \geq 0$; and vanishes otherwise. It is easy to see in (7.2) that for this problem the value of λ_c depends only on $k = \text{wt}(c)$. Thus, for $k \in \{0, \dots, n\}$, we may set $\lambda_{n,k} = \lambda_c$, where $c \in \mathbb{Z}_2^n$ is any bit string of weight k . To calculate $\lambda_{n,k}$, we consider the string $c = 0^{n-k}1^k$, which has weight k . For any $b \in \mathbb{Z}_2^n$, let s denote the number of 1s in the first $n - k$ bits of b , and let r denote the number of 1s in the remaining k bits of b . Then $b \cdot c = r$, and $\phi(b) = \Theta(\frac{n}{2} - (r + s))$. Since the number of bit strings b with given values for r and s is $\binom{n-k}{s} \binom{k}{r}$, we have

$$\lambda_{n,k} = \sum_{r,s} \binom{n-k}{s} \binom{k}{r} (-1)^{\Theta(\frac{n}{2} - (r+s))} (-1)^r.$$

Using standard combinatorial techniques, this sum simplifies to give

$$\lambda_{n,k} = \begin{cases} -\frac{(-1)^{k/2}}{\sqrt{2^n}} \binom{n}{n/2} \frac{1 \cdot 3 \cdots (k-1)}{(n-1) \cdot (n-3) \cdots (n-k+1)} & \text{for } k \text{ even;} \\ \lambda_{n,k-1} & \text{for } k \text{ odd.} \end{cases}$$

It follows that for any even number n , the eigenvalue of smallest absolute value is the middle eigenvalue $\lambda_{n,n/2}$, which is given by

$$|\lambda_{n,n/2}| = \begin{cases} \frac{1}{\sqrt{2^n}} \binom{n/2}{n/4} & \text{if } n \equiv 0 \pmod{4}; \\ \frac{2}{\sqrt{2^n}} \binom{(n-2)/2}{(n-2)/4} & \text{if } n \equiv 2 \pmod{4}. \end{cases}$$

Since, by Stirling's formula, each of these expressions is asymptotic to $1/\sqrt{n}$, we find that the average s of the absolute values of the eigenvalues of $A_{\mathcal{MAJ}^n}$ is $\Omega(1/\sqrt{n})$. Thus the quantum query complexity of MAJORITY is $O(1/s) = O(\sqrt{n})$, as claimed. \square

8 Conclusion

In this paper we have derived a general technique—Amplified Impatient Learning—for quantum concept learning from a minimally adequate teacher. We applied it to two novel problems: BATTLESHIP and MAJORITY, that like the problems of Bernstein-Vazirani and Grover, can be recognized as concept learning problems. Amplified Impatient Learning solves SMALLSHIP(d) with $O(\sqrt{N})$ queries, an improvement over the $\Omega(N)$ queries required classically. For BIGSHIP(α), Amplified Impatient Learning is not so good, but we gave an alternative quantum algorithm with sample complexity $O(\log N)$. Finally, Amplified Impatient Learning solves MAJORITY with $O(\sqrt{n})$ quantum queries, again an improvement over the $\Omega(n)$ required classically.

Quantum algorithms for concept learning were first considered by Bshouty and Jackson [38], who analyzed the traditional DNF learning problem [39]. Subsequently, Servedio and Gortler proved some general lower bounds on the quantum sample complexity of learning from any membership oracle [40]. Their results, together with algorithms derived in this paper, motivate us to make a pair of conjectures about general upper bounds on the quantum sample complexity of learning from a membership oracle:

Conjecture 1 *For any family of concept classes $\{\mathcal{C}_i\}$ with $|\mathcal{C}_i| \rightarrow \infty$, there exists a quantum learning algorithm with membership oracle query complexity $O(\sqrt{|\mathcal{C}_i|})$.*

Our quantum algorithm for SMALLSHIP(d) (which specializes to Grover's algorithm when $d = 1$) saturates this bound; the idea is that these minimally structured search problems are concept learning problems that are as difficult as any of the same size. As we noted in the proof of Theorem 6.1, the calls to the equivalence oracle in this problem can be replaced by calls to the membership oracle, so our results are consistent with Conjecture 1.

The difficulty of concept learning problems depends on more than the number of concepts $|\mathcal{C}|$ among which the target concept lies, however; it also depends on how similar distinct concepts are. Servedio and Gortler express their lower bounds in terms of a quantity $\gamma_{\mathcal{C}}$ that measures this similarity: For any $\mathcal{C}' \subset \mathcal{C}$, define $\mathcal{C}'_{x,b} = \{c \in \mathcal{C}' \mid c(x) = b\}$. Then

$$\gamma_{\mathcal{C}} = \min_{\mathcal{C}' \subset \mathcal{C}, |\mathcal{C}'| \geq 2} \max_{x \in X} \min_{b \in \mathbb{Z}_2} \frac{|\mathcal{C}'_{x,b}|}{|\mathcal{C}'|}.$$

$\gamma_{\mathcal{C}}$ is small if there is a large subset \mathcal{C}' from which the response to any query to the membership oracle rules out only a small fraction of the concepts. In this case we expect the concept class to be difficult to learn. (Since \mathcal{C}' can contain only 3 concepts, from which a query might eliminate only 1, $\gamma_{\mathcal{C}}$ cannot be greater than $1/3$.)

Conjecture 2 *For any family of concept classes $\{C_i\}$ with $|C_i| \rightarrow \infty$, there exists a quantum learning algorithm with membership oracle query complexity $O\left(\frac{\log |C_i|}{\sqrt{\gamma}}\right)$.*

The problems MAJORITY, SMALLSHIP(d), and BIGSHIP(α) studied in this paper provide examples of learning problems that satisfy the bounds given in the above conjectures. For the BATTLESHIP problem, one calculates that $\gamma_{BS} = \min\{d/N, 1/3\}$. Thus, for SMALLSHIP(d), whose quantum sample complexity is $O(\sqrt{N})$, Conjecture 1 is sharp, while Conjecture 2 provides the weaker bound of $O(\sqrt{N} \log N)$. For BIGSHIP(α), the situation is reversed: Conjecture 2 provides a sharp upper bound of $O(\log N)$, while Conjecture 1 gives the weaker upper bound of $O(\sqrt{N})$. For MAJORITY, whose quantum sample complexity is at most $\log N$, it is easy to see (using the fact that $\gamma \leq 1/3$) that Conjecture 2 holds.³

Acknowledgments We thank Zeph Landau, Jeff Remmel and Ronald de Wolf for useful discussions. This work has been partially supported by the National Science Foundation (NSF) under grant ECS-0202087, and by the National Security Agency (NSA) and Advanced Research and Development Activity (ARDA) under Army Research Office (ARO) grant number DAAD19-01-1-0520.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Shor, P.W.: Algorithms for quantum computation: discrete logarithms and factoring. In: Goldwasser, S. (ed.) Proceedings of the 35th Symposium on Foundations of Computer Science, Santa Fe, NM, 20–22 November 1994, pp. 124–134. IEEE, Los Alamitos, CA (1994)
2. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, quant-ph/9508027. SIAM J. Comput. **26**, 1484–1509 (1997)
3. Coppersmith, D.: An approximate Fourier transform useful in quantum factoring. IBM T. J. Watson Research Report RC 19642 (1994)
4. Høyer, P.: Efficient quantum transforms, quant-ph/9702028 (preprint 1997)
5. Klappenecker, A.: Wavelets and wavelet packets on quantum computers, quant-ph/9909014. In: Unser, M.A., Aldroubi, A., Laine, A.F. (eds.) Wavelet Applications in Signal and Image Processing VII, Denver, CO, 19–23 July 1999, SPIE Proceedings, vol. 3813, pp. 703–713. SPIE, Bellingham, WA (1999)
6. Klappenecker, A., Rötteler, M.: Discrete cosine transforms on quantum computers, quant-ph/0111038. In: Loncaric, S., Babic, H. (eds.) Proceedings of the 2nd International Symposium on Image and Signal Processing and Analysis, Pula, Croatia, 19–21 June 2001, pp. 464–468. IEEE, Los Alamitos, CA (2001)

³ Added in proof: Recently, Conjecture 1 was proved by Ambainis et al. [41]. In addition, a version of Conjecture 2, weaker by only a logarithmic factor, was proved by Atici and Servedio [42]. Specifically, they prove that $O(\log |C_i| \log \log |C_i| / \sqrt{\gamma})$ is an upper bound on quantum query complexity.

7. Freedman, M.H.: Poly-locality in quantum computing, quant-ph/0001077. *Found. Comput. Math.* **2**, 145–154 (2002)
8. Grover L.K.: A fast quantum mechanical algorithm for database search. In: *Proceedings of the 28th Annual ACM Symposium on the Theory of Computing*, Philadelphia, PA, 22–24 May 1996, pp. 212–219. ACM, New York (1996)
9. Grover, L.K.: Quantum mechanics helps in searching for a needle in a haystack, quant-ph/9706033. *Phys. Rev. Lett.* **79**, 325–328 (1997)
10. Brassard, G., Høyer, P.: An exact quantum polynomial-time algorithm for Simon’s problem, quant-ph/9704027. In: *Proceedings of 5th Israeli Symposium on Theory of Computing and Systems*, Ramat-Gan, Israel 17–19 June 1997, pp. 12–23. IEEE, Los Alamitos, CA (1997)
11. Grover, L.K.: A framework for fast quantum mechanical algorithms, quant-ph/9711043. In: *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, Dallas, TX, 23–26 May 1998, pp. 53–62. ACM, New York (1998)
12. Brassard, G., Høyer, P., Tapp, A.: Quantum counting, quant-ph/9805082. In: *Proceedings of the 25th International Colloquium on Automata, Languages, and Programming*, Ålborg, Denmark, 13–17 July 1998. *Lecture Notes in Computer Science*, vol. 1443, pp. 820–831. Springer-Verlag, Berlin (1998)
13. Brassard, G., Høyer, P., Mosca, M., and Tapp, A.: Quantum amplitude amplification and estimation, quant-ph/0005055. In: Lomonaco, S.J. Jr., Brandt, H.E. (eds.) *Quantum Computation and Information*, *Contemporary Mathematics*, vol. 305, pp. 53–74. AMS, Providence, RI (2002)
14. Bernstein, E., Vazirani, U.: Quantum complexity theory. In: *Proceedings of the 25th ACM Symposium on Theory of Computing*, San Diego, CA, 16–18 May 1993, pp. 11–20. ACM Press, New York (1993)
15. Bernstein, E., Vazirani, U.: Quantum complexity theory. *SIAM J. Comput.* **26**, 1411–1473 (1997)
16. Mitchell, T.M.: *Machine Learning*. McGraw-Hill, San Francisco (1997)
17. Cucker, F.S., Smale, S.: On the mathematical foundations of learning. *Bull. Amer. Math. Soc.* **39**, 1–49 (2002)
18. Angluin, D.: Computational learning theory: survey and selected bibliography. In: *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, Victoria, British Columbia, Canada, 4–6 May 1992, pp. 351–369. ACM, New York (1992)
19. Angluin, D.: Queries and concept learning. *Mach. Learn.* **2**, 319–342 (1988)
20. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York (1979)
21. Meyer, D.A.: Sophisticated quantum search without entanglement, quant-ph/0007070. *Phys. Rev. Lett.* **85**, 2014–2017 (2000)
22. Cleve, R., Ekert, A., Macchiavello, C., Mosca, M.: Quantum algorithms revisited, quant-ph/9708016. *Proc. Roy. Soc. Lond. A.* **454**, 339–354 (1998)
23. Sylvester, J.J.: Thoughts on inverse orthogonal matrices, simultaneous sign-successions, and tessellated pavements in two or more colours, with applications to Newton’s rule, ornamental tile-work, and the theory of numbers. *Philos. Mag. Ser. IV.* **34**, 461–475 (1867)
24. Hadamard, M.J.: Résolution d’une question relative aux déterminants. *Bull. Sci. Math.* **17**, 240–246 (1893)
25. van Dam, W.: Quantum algorithms for weighing matrices and quadratic residues, quant-ph/0008059. *Algorithmica* OF1–OF16 (2002)
26. Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.: Strengths and weaknesses of quantum computing. *SIAM J. Comput.* **26**, 1510–1523 (1997)
27. Helstrom, C.W.: Detection theory and quantum mechanics. *Inf. Control.* **10**, 254–291 (1967)
28. Kholevo, A.S.: Quantum statistical decision theory. *J. Multivar. Anal.* **3**, 337–394 (1973)
29. von Neumann, J.: *Mathematische Grundlagen der Quantenmechanik* (Translated by Beyer, R.T. as *Mathematical Foundations of Quantum Mechanics*. Princeton University Press, Princeton 1955). Springer-Verlag, Berlin (1932)
30. Belavkin, V.P.: Optimal multiple quantum statistical hypothesis testing. *Stochastics.* **1**, 315–345 (1975)
31. Kennedy R.S.: On the optimal receiver for the M-ary pure state problem. *MIT Res. Lab. Electron. Quart. Prog. Rep.* **110** (15 July 1973) 142–146; see also [32], Appendix to Chap. IV
32. Helstrom, C.W.: *Quantum Detection and Estimation Theory*. Academic, New York (1976)
33. Yuen, H.P., Kennedy, R.S.: On optimal quantum receivers for digital signal detection. *Proc. IEEE* **58**, 1770–1773 (1970)
34. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press, Cambridge, p. 432, Theorem 7.4.9 (1985)

35. Eldar, Y.C., Forney, G.D. Jr.: On quantum detection and the square-root measurement, quant-ph/0001532. *IEEE Trans. Inf. Theory*. **47**, 858–872 (2001)
36. Barnum, H., Knill, E.: Reversing quantum dynamics with near-optimal quantum and classical fidelity. *J. Math. Phys.* **43**, 2097–2106 (2002)
37. Boyer, M., Brassard, G., Høyer, P., Tapp, A.: Tight bounds on quantum searching, quant-ph/9605034. *Fortsch. Phys.* **46**, 493–506 (1998)
38. Bshouty, N.H., Jackson, J.C.: Learning DNF over the uniform distribution using a quantum example oracle. *SIAM J. Comput.* **28**, 1136–1153 (1999)
39. Valiant, L.G.: A theory of the learnable. *Commun. ACM*. **27**, 1134–1142 (1984)
40. Servedio, R.A., Gortler S.J.: Quantum versus classical learnability. In: *Proceedings of the 16th Annual IEEE Conference on Computational Complexity*, Chicago, IL, 18–21 June 2001, pp. 138–148. IEEE, Los Alamitos, CA (2001)
41. Ambainis, A., Iwama, K., Kawachi, A., Masuda, H., Putra R.H., Yamashita, S.: Quantum identification of Boolean oracles. In: Diekert, V., Habib, M. (eds.) *Proceedings of STACS 2004, Lecture Notes in Computer Science*, vol. 2996, pp. 105–116. Springer, Berlin (2004)
42. Atici, A., Servedio, R.: Improved bounds on quantum learning algorithms. *Quant. Inf. Proc.* **4**, 355–386 (2005)