

Quantum Nearest-Neighbor Algorithms for Machine Learning

Nathan Wiebe[†], Ashish Kapoor^{*}, and Krysta M. Svore[†]

[†]Quantum Architectures and Computation Group, Microsoft Research, Redmond, WA (USA)

^{*}Adaptive Systems and Interaction Group, Microsoft Research, Redmond, WA (USA)

We present several quantum algorithms for performing nearest-neighbor learning. At the core of our algorithms are fast and coherent quantum methods for computing distance metrics such as the inner product and Euclidean distance. We prove upper bounds on the number of queries to the input data required to compute these metrics. In the worst case, our quantum algorithms lead to polynomial reductions in query complexity relative to the corresponding classical algorithm. In certain cases, we show exponential or even super-exponential reductions over the classical analog. We study the performance of our quantum nearest-neighbor algorithms on several real-world binary classification tasks and find that the classification accuracy is competitive with classical methods.

The discipline of machine learning has exploded in recent years due to the increasing demand for automated modeling of large amounts of data. Every day people interact with hundreds of systems developed using machine learning techniques, including internet search engines, speech recognition applications, GPS-based navigation tools, and automated robots. Consider the task faced by the U.S. postal service of routing over 160 billion pieces of mail. The sheer magnitude of this problem necessitates the use of software to automatically recognize the handwritten digits and letters that form the address of a recipient. The nearest-neighbor algorithm is commonly used to solve tasks such as handwriting recognition due to its simplicity and high performance accuracy [1].

Nearest-neighbor classification is a supervised machine learning technique that relies on labeled *training* data to make a classification decision. Consider the binary classification task of determining if a given unlabeled *test* digit is *even* or *odd*. The *training* data consists of handwritten digits expressed as multidimensional feature vectors, each with a human-assigned label of either even or odd. The entries of each feature vector $\mathbf{v} \in \mathbb{R}^N$, where N is the number of features used to characterize the digit, may be the pixel values that comprise the image. Figure 1 shows an example of 25 digits, each of which is represented by a 256-dimensional feature vector of pixel values.

We can divide the training set into two sets, or clusters, of vectors in \mathbb{R}^N , $\{A\}$ and $\{B\}$, such that $\{A\}$ contains only odd examples, $\{B\}$ contains only even examples, and $|\{A\}| + |\{B\}| = M_A + M_B = M$. The goal is to classify, or label, a given unlabeled *test* point $\mathbf{u} \in \mathbb{R}^N$ as $\mathbf{u} \in A$ or $\mathbf{u} \in B$. Here we take N and M to be large; an algorithm is efficient if it requires time $O(\text{polylog}(NM))$.

The nearest-neighbor algorithm first computes the distance between the test vector \mathbf{u} and each training vector \mathbf{v} and then assigns \mathbf{u} to the cluster that contains the closest vector to \mathbf{u} . Specifically, it assigns \mathbf{u} to $\{A\}$ if

$$\min_{\mathbf{a} \in \{A\}} |\mathbf{u} - \mathbf{a}| \leq \min_{\mathbf{b} \in \{B\}} |\mathbf{u} - \mathbf{b}|,$$

for an appropriate distance metric $|\mathbf{u} - \mathbf{v}|$. Extensions of nearest-neighbor include k -nearest-neighbors, where a




FIG. 1: An example of 25 handwritten digits. Each digit is stored as a 256-pixel greyscale image and represented as a unit vector with $N = 256$ features.

function, such as majority, of the labels of the k nearest training vectors is used to determine the label of the test point [2]. The classical nearest-neighbor algorithm requires time $O(NM)$.

Quantum computation shows promise as a powerful resource for accelerating certain classical machine learning algorithms [3–6]. However, a major challenge facing the development of practical quantum machine learning algorithms is the need for an oracle to return, for example, the distances between elements in the test and training sets. Lloyd, Mohseni, and Rebentrost [5] recently proposed an efficient quantum algorithm to address this problem. Namely, their algorithm computes a representative vector for each set, referred to as a *centroid*, by averaging the vectors in A and B , respectively. The test point \mathbf{u} is assigned to $\{A\}$ if the distance to the centroid of $\{A\}$, written as $\text{mean}(\{A\})$, is smallest,

$$|\mathbf{u} - \text{mean}(\{A\})| \leq |\mathbf{u} - \text{mean}(\{B\})|,$$

and $\{B\}$ otherwise.

Note that this algorithm, which we refer to as *nearest-centroid classification* is a form of nearest-neighbor classification, where the nearest *centroid* is used to determine the label, as opposed to the nearest training point. If the number of clusters equals the number of points, then it reduces to nearest-neighbor classification. In practice, nearest-centroid classification can perform poorly because $\{A\}$ and $\{B\}$ are often embedded in a complicated manifold where the mean values of the sets are not within the manifold [7]. In contrast, nearest-neighbor classification tends to work well in practice and often out-

performs centroid-based classification but can be prohibitively expensive on classical computers [8].

Therefore, we present a quantum nearest-neighbor algorithm that assigns a point \mathbf{u} to either cluster $\{A\}$ or $\{B\}$ such that both the probability of a faulty assignment and the number of quantum oracle queries is minimized. We consider two choices of distance metrics within our algorithm:

1. the inner product, $|\mathbf{u} - \mathbf{v}| = |\mathbf{u}||\mathbf{v}| - \mathbf{u} \cdot \mathbf{v}$,
2. the Euclidean distance, $|\mathbf{u} - \mathbf{v}| = \sqrt{\sum_{i=1}^N (u_i - v_i)^2}$.

Our quantum algorithm overcomes the main drawback of the nearest-centroid approach in [5]: low assignment accuracy in many real-world problems.

Throughout, the test point is set to $\mathbf{v}_0 := \mathbf{u}$ and the training set consists of \mathbf{v}_j , for $j = 1, \dots, M$. We assume the following: (1) The input vectors are d -sparse, i.e., contain no more than d non-zero entries. (2) Quantum oracles are provided in the form

$$\begin{aligned} \mathcal{O} |j\rangle |i\rangle |0\rangle &:= |j\rangle |i\rangle |v_{ji}\rangle, \\ \mathcal{F} |j\rangle |\ell\rangle &:= |j\rangle |f(j, \ell)\rangle, \end{aligned} \quad (1)$$

where v_{ji} is the i^{th} element of the j^{th} vector and $f(j, \ell)$ gives the location of the ℓ^{th} non-zero entry in \mathbf{v}_j . (3) The user knows an upper bound r_{\max} on the maximum value of any feature in the dataset. (4) Each vector is normalized to 1, for convenience (this is not necessary). (5) The run time of the algorithm is dominated by the number of queries made to oracles \mathcal{O} and \mathcal{F} .

We show for our algorithm that:

1. The number of queries depends on dr_{\max}^2 rather than on the feature dimension N or the sparsity d alone. Thus, for practical applications the query complexity is typically independent of the number of features (i.e., $r_{\max} \propto 1/\sqrt{d}$).
2. The number of queries scales as $O(\sqrt{M} \log(M))$ rather than M . Furthermore, the query complexity of our Euclidean-based method can be independent of M .
3. Our algorithm can tolerate relatively large errors in distance calculations when applied to real-world classification problems.

Implementation of oracle-based algorithms will require instantiation of the oracles, which are an abstraction of the many ways for algorithms to interact with data. If the task is to classify chemicals, the oracle query could represent a call to an efficient quantum simulation algorithm that yields physical features of the chemicals [9, 10]. In other cases, the oracle query could represent accesses to a large quantum database that contains classical bit strings. One way to construct such a database is to use a quantum random access memory (qRAM) [11], however alternate implementations are possible. In this work, we assume oracles are provided and show how to minimize the number of queries to the oracle.

Inner Product Classification. We first describe a quantum nearest-neighbor algorithm that uses the inner product as the distance metric. We show, somewhat surprisingly, that the required number of oracle queries does not explicitly depend on the number of features N . Rather, it depends implicitly on N through dr_{\max}^2 and ϵ .

Theorem 1. *Let \mathbf{v}_0 and $\{\mathbf{v}_j : j = 1, \dots, M\}$ be d -sparse unit vectors such that $\max_{j,i} |v_{ji}| \leq r_{\max}$, then the task of finding $\max_j |\langle \mathbf{u} | \mathbf{v}_j \rangle|^2$ within error at most ϵ and with success probability at least $1 - \delta_0$ requires an expected number of queries that is bounded above by*

$$1080\sqrt{M} \left[\frac{4\pi(\pi+1)d^2r_{\max}^4}{\epsilon} \right] \left[\frac{\ln\left(\frac{81M(\ln(M)+\gamma)}{\delta_0}\right)}{2(8/\pi^2 - 1/2)^2} \right],$$

where $\gamma \approx 0.5772$ is Euler's constant.

Two important scaling factors in the theorem should be emphasized. First, the scaling of the query complexity with M is near-quadratically better than its classical analog. Second, if $r_{\max} \propto 1/\sqrt{d}$ then the scaling is independent of both d and N . We expect this condition to occur when all input vectors have at least $\Theta(d)$ sparsity.

Note that in cases where the sign of the inner product is necessary for assignment, we can generalize the above method by performing

$$|\mathbf{v}_\ell\rangle \mapsto \frac{|0\rangle |0^{\otimes \log_2 N}\rangle + |1\rangle |\mathbf{v}_\ell\rangle}{\sqrt{2}}, \quad (2)$$

and then using these states in **Theorem 1**. This allows direct estimation of the cosine distance and in turn the inner product. Non-unit vectors can be trivially handled given access to quantum oracles that return the norms of the vectors [6].

We prove **Theorem 1** by the following steps (see supplementary material for more detail). Assume that we want to compute the inner product between two states \mathbf{v}_j and $\mathbf{v}_0 := \mathbf{u}$ and let $v_{ji} = r_{ji}e^{i\phi_{ji}}$, where r_{ji} is a positive number. This can be achieved using a coherent version of the swap test [12] on the states

$$\begin{aligned} d^{-1/2} \sum_i |i\rangle \left(\sqrt{1 - \frac{r_{ji}^2}{r_{\max}^2}} e^{-i\phi_{ji}} |0\rangle + \frac{v_{ji}}{r_{\max}} |1\rangle \right) |1\rangle, \\ d^{-1/2} \sum_i |i\rangle |1\rangle \left(\sqrt{1 - \frac{r_{0i}^2}{r_{\max}^2}} e^{-i\phi_{0i}} |0\rangle + \frac{v_{0i}}{r_{\max}} |1\rangle \right). \end{aligned} \quad (3)$$

which, as we show in the appendix, can be prepared using six oracle calls and two single-qubit rotations. If the swap test is applied to these states and a probability of obtaining outcome '0', denoted $P(0)$, is found then

$$|\langle \mathbf{u} | \mathbf{v}_j \rangle|^2 = (2P(0) - 1)d^2r_{\max}^4.$$

Statistical sampling requires $O(M/\epsilon^2)$ queries to achieve the desired error tolerance, which can be expensive if small values of ϵ are required.

We reduce the scaling with ϵ to $O(1/\epsilon)$ by removing the measurement in the swap test and applying amplitude estimation (AE) [13] to estimate $P(0)$ within error ϵ , denoted $\tilde{P}(0)$. This can be done because the state preparation procedure and the measurement-free swap test are invertible. If a register of dimension R is used in AE then the inference error obeys

$$|P(0) - \tilde{P}(0)| \leq \frac{\pi}{R} + \frac{\pi^2}{R^2}.$$

Choosing R to be large enough so that the error in AE is at most $\epsilon/2$ yields

$$R \geq \left\lceil \frac{4\pi(\pi+1)d^2 r_{\max}^4}{\epsilon} \right\rceil. \quad (4)$$

The scaling with M can also be quadratically reduced by using the maximum/minimum finding algorithm of Dürr and Høyer [14], which combines Grover's algorithm with exponential search to find the largest or smallest element in a list. In order to apply the algorithm, we need to make the AE step reversible. We call this form of AE *coherent amplitude estimation*.

We achieve this by introducing a coherent majority voting scheme on a superposition over k -copies of the output of AE. AE outputs a state of the form $a|y\rangle + \sqrt{1-|a|^2}|y^\perp\rangle$, where y is a bit-string that encodes $P(0)$ and $|y^\perp\rangle$ is orthogonal to $|y\rangle$. The median of k bitstrings x_k is computed coherently by $\mathcal{M} : |x_1\rangle \cdots |x_k\rangle |0\rangle \mapsto |x_1\rangle \cdots |x_k\rangle |\bar{x}\rangle$, where \bar{x} is the median of $[x_1, \dots, x_k]$. For this application, AE guarantees that $|a|^2 \geq 8/\pi^2 > 1/2$ and Hoeffding's inequality shows that $\bar{y} = y$ with overwhelming probability if k is sufficiently large. In particular, it is straightforward to show using the binomial theorem that we can write

$$\begin{aligned} \mathcal{M}(a|y\rangle + \sqrt{1-|a|^2}|y^\perp\rangle)^{\otimes k} |0\rangle \\ = A|\Psi\rangle |y\rangle + \sqrt{1-|A|^2}|\Phi\rangle |y^\perp\rangle, \end{aligned} \quad (5)$$

where $|A|^2 > 1 - \Delta$ for $k \geq \ln(\frac{1}{\Delta}) / (2(8/\pi^2 - \frac{1}{2}))$ and states $|\Psi\rangle$ and $|\Phi\rangle |y^\perp\rangle$ are computationally irrelevant. We then use coherent majority voting to construct a $\sqrt{2\Delta}$ -approximate oracle that maps $|j\rangle |0\rangle \mapsto |j\rangle |\bar{y}\rangle$. This approximate oracle is then used in the Dürr Høyer minimum finding algorithm.

We then make the pessimistic assumption that if the use of an approximate oracle leads to an erroneous outcome from the minimum finding algorithm even once then the whole algorithm fails. Fortunately, since the number of repetitions of AE scales as $k \in O(\log(1/\Delta))$, this probability can be made vanishingly small at low cost. Our final cost estimate then follows by multiplying, k , R , the costs of state preparation, and the number of iterations used in the Dürr Høyer algorithm.

Euclidean Distance Classification. The classification problem can also be solved using the Euclidean distance. We now describe a quantum nearest-neighbor algorithm

based on the Euclidean distance between \mathbf{u} and the cluster *centroids*, i.e., the mean values of the vectors within each cluster. This can be viewed as a step in a k -means clustering algorithm [5, 15]. We refer to this algorithm as the *nearest-centroid* algorithm.

Our nearest-centroid algorithm differs substantially from that of [5] in that (1) we normalize the computed distances, and (2) we consider a generalization to cases where each cluster is subdivided into M' clusters that each contain $M_1, \dots, M_{M'}$ vectors respectively. If $M' = M$ then the algorithm reduces to nearest-neighbor classification with the Euclidean distance metric.

These differences help address two central problems of centroid-based classification. First, imagine that cluster $\{A\}$ is dense but cluster $\{B\}$ is sparse. Then even if $|\mathbf{u} - \text{mean}(\{A\})| \leq |\mathbf{u} - \text{mean}(\{B\})|$, it may be much more likely that $\mathbf{u} \in \{B\}$ because the probability of a large deviation from the centroid is much greater for $\{B\}$ than $\{A\}$. Normalizing the distance by the width of the cluster can help address this issue [16]. We also show in the supplemental material that this assignment reduces to the likelihood ratio test under certain assumptions. Second, if $\{A\}$ and $\{B\}$ are non-convex then the centroid of $\{A\}$ *may actually be in* $\{B\}$. Segmenting the data into M' smaller clusters can help address this issue.

The following theorem gives the query complexity for our quantum nearest-centroid algorithm. The normalization of the distance by the cluster width can easily be omitted from the algorithm if desired.

Theorem 2. *Let \mathbf{v}_0 and $\{\mathbf{v}_j^{(m)} : j = 1, \dots, M_m, m = 1, \dots, M'\}$ be d -sparse unit vectors such that the components satisfy $\max_{m,j,i} |v_{ji}^{(m)}| \leq r_{\max}$ and $\sigma_m = \frac{1}{M_m} \sum_{p=1}^{M_m} \|\mathbf{v}_p^{(m)} - \mathbf{v}_p^{(m)}\|_2^2 + \frac{1}{M_m} \sum_j \|\mathbf{v}_j^{(m)}\|_2^2$, if $M_m > 1$ and $\sigma_m = 1$ otherwise. The task of finding*

$$\min_m \left(\frac{\|\mathbf{v}_0 - \frac{1}{M_m} \sum_{j=1}^{M_m} \mathbf{v}_j^{(m)}\|_2^2}{\sigma_m} \right),$$

with error in the numerator and denominator bounded above by ϵ and with success probability at least $1 - \delta_0$, requires an expected number of queries that is bounded above by

$$900\sqrt{M'} \left\lceil \frac{8\pi(\pi+1)dr_{\max}^2}{\epsilon} \right\rceil \left\lceil \frac{\log\left(\frac{81M'(\log(M')+\gamma)}{\delta_0}\right)}{2((8/\pi^2)^2 - 1/2)^2} \right\rceil.$$

If $M' \in O(\text{polylog}(MN))$ and $dr_{\max}^2 \in O(1)$ then the learning problem is efficient, which motivates the use of centroid-based classification for supervised learning problems where $\{A\}$ and $\{B\}$ can be partitioned into a union of a small number of disjoint training sets that are both unimodal and convex. Even if $M' \in \Theta(M)$ is required, then the query complexity of this method is at most comparable to the inner-product-based approach.

The proof of [Theorem 2](#) follows similarly to that of [Theorem 1](#). We use coherent amplitude estimation (AE) to find the numerator and the distance between \mathbf{u} and the centroid as well as the intra-cluster variance. We then use a reversible circuit to divide the distance by the variance and use the Dürr Høyer algorithm [14] to find the minimum relative distance over all M' clusters. The biggest conceptual difference is that in this case we do not use the swap test; instead we use a method from [6, 17].

The result of [17] shows that for any unitary V and transformation mapping $|j\rangle|0\rangle \mapsto |j\rangle|\mathbf{v}_j\rangle$, a measurement can be performed that has success probability

$$P(0) \propto \left| \sum_{j=0}^{M_m} |V_{j0}\rangle^2 \mathbf{v}_j \right|^2, \quad (6)$$

for any M_m . If we choose $\mathbf{v}_0 = -\mathbf{u}$, then by (6) and

$$|V_{j0}\rangle = \begin{cases} \frac{1}{\sqrt{2}}, & j = 0 \\ \frac{1}{\sqrt{2M_m}}, & \text{otherwise} \end{cases},$$

the probability of success gives the square of the Euclidean distance between \mathbf{u} and the cluster centroid. Note that non-unit vectors can be accommodated by doubling the number of vectors and setting $\mathbf{v}_j = \mathbf{x}_j + \mathbf{y}_j$ where \mathbf{x}_j and \mathbf{y}_j are unit vectors. We show in the supplemental material that

$$\left| \mathbf{u} - \frac{1}{M_m} \sum_{j \geq 1} \mathbf{v}_j \right|^2 = 4dr_{\max}^2 P(0).$$

The operator V can be implemented efficiently using techniques from quantum simulation (see supplemental material), and the \mathbf{v}_j are prepared using (3). The process of estimating the distance is therefore efficient.

The remainder of the procedure is identical to that of the inner-product-based classification. The most notable technical difference is that the phase estimation procedure must succeed in both the distance and the intra-cluster variance calculations. This results in the success probability in phase estimation dropping from at least $8/\pi^2$ to at least $(8/\pi^2)^2 \approx 2/3$. Thus, quantum nearest-centroid classification (based on Euclidean distance) requires more iterations than quantum nearest-neighbor classification (based on inner product distance).

Numerical Experiments. We evaluate the performance of our algorithms on several real-world tasks. The first task is based on handwritten digits from the MNIST digits database [18]. Given a training set of M handwritten digits (see [Figure 1](#)) and their labels (even or odd), assign a label (even or odd) to an unlabeled test digit. Each digit is represented by a 256-dimensional feature unit vector of pixel values [18]. These pixel values are rescaled to have zero mean and unit variance over the data set before normalizing. In all plots, error bars indicate standard deviation of the accuracy.




FIG. 2: Classification accuracy for digit data vs ϵ for cases where half the dataset is used for training.

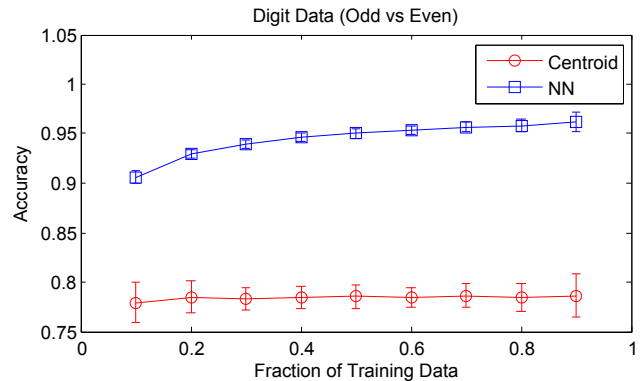


FIG. 3: Classification accuracy for digit data for fixed noise $\epsilon = 10^{-5}$ as a function of the training data size.

First, we compare the accuracy of our nearest-neighbor algorithm (NN) based on the inner product to our nearest-centroid algorithm (Centroid) based on Euclidean distance, as a function of noise ϵ in the distance computation. For Centroid, we set $M' = 1$. [Figure 2](#) plots ϵ versus the accuracy of both NN (blue squares) and Centroid (red circles), for noise drawn from $\mathcal{N}(0, \epsilon)$ independently for each distance computation, where $\epsilon \in [10^{-5}, 10^5]$. The accuracy is averaged across classification of 100 random test examples using $M = 2000$ training points. In the low-noise regime, NN significantly outperforms Centroid by roughly 20%. At $\epsilon \approx 0.1 \approx 1/\sqrt{N}$, both algorithms exhibit significant loss in accuracy and eventually degrade to 50% accuracy. Both NN and Centroid tolerate relatively large errors without sacrificing classification accuracy.

The tolerance against noise up to size $O(1/\sqrt{N})$ is well-justified for high-dimensional systems by concentration of measure arguments [19], so we anticipate that $\epsilon \in \Theta(1/\sqrt{N})$ should be appropriate for problems that lack underlying symmetry in the assignment set (such as even/odd classification).

Second, we study the effect of training data size on the performance accuracy of the two algorithms for a fixed noise rate $\epsilon = 0.1$. [Figure 3](#) plots the training data

size versus the performance accuracy of both NN (blue squares) and Centroid (red circles). We vary the training set size by taking random fractions of $M = 4000$ points, at fractions of 0.1, 0.2, \dots , 0.9. For all training set sizes, NN significantly outperforms Centroid. In addition, NN exhibits increasing performance accuracy as M increases, from 84% to 90%. In contrast, Centroid’s accuracy hovers around 73%, even as M increases.

While NN outperforms Centroid on the digit classification task, we find that for other tasks, outlined in the supplemental material, Centroid outperforms NN. However, in tasks where Centroid performs well, we find that both methods exhibit low classification accuracy. This could indicate the need for more training data, in which case NN may begin to outperform Centroid as the amount of training data M grows.

For the digit classification task, we estimate that accuracy α can be obtained using NN with a number of oracle queries that scales (for constant success probability) as $O(\sqrt{M}) = O((1 - \alpha)^{-5/4} \log((1 - \alpha)^{-1}))$. In contrast, the number of queries in the classical nearest-neighbor algorithm scales as $O((1 - \alpha)^{-5/2})$. The centroid-based algorithm achieves at best $\alpha \approx 0.78$. In addition, we find that $dr_{\max}^2 \approx 2.8$ for this problem, indicating that the cost of state preparation will likely become negligible as $\alpha \rightarrow 1$ and in turn as $M \rightarrow \infty$.

Conclusions. We have presented two quantum algorithms for performing nearest-neighbor classification that promise significant speed-ups over their classical counterparts. Our algorithms enable classification over datasets with both a high-dimensional feature space as well as a large number of training examples. Computation of distance metrics is extremely common in machine learning algorithms; we have developed two fast methods for computing distance metrics on a quantum computer that can be implemented coherently. Finally, we have shown that our algorithms are robust to noise and perform well when applied to typical real-world tasks.

We find that quantum algorithms for machine learning can provide algorithmic improvements over classical machine learning techniques. Our algorithms are a step toward blending fast quantum methods with proven machine learning techniques. In future work, these methods could be used to develop quantum k -nearest-neighbor algorithms or more generally other supervised, unsupervised or semi-supervised learning algorithms.

ACKNOWLEDGMENTS

We thank Matt Hastings for valuable comments and feedback.

-
- [1] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *Information Theory, IEEE Transactions on*, 13(1):21–27, 1967.
 - [2] Keinosuke Fukunaga and Patrenahalli M. Narendra. A branch and bound algorithm for computing k -nearest neighbors. *Computers, IEEE Transactions on*, 100(7):750–753, 1975.
 - [3] Esmat Aïmeur, Gilles Brassard, and Sébastien Gambs. Machine learning in a quantum world. In *Advances in Artificial Intelligence*, pages 431–442. Springer, 2006.
 - [4] Daoyi Dong, Chunlin Chen, Hanxiong Li, and Tzyh-Jong Tarn. Quantum reinforcement learning. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 38(5):1207–1220, 2008.
 - [5] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.
 - [6] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big feature and big data classification. *arXiv preprint arXiv:1307.0471*, 2013.
 - [7] Ilya Levner. Feature selection and nearest centroid classification for protein mass spectrometry. *BMC bioinformatics*, 6(1):68, 2005.
 - [8] Nuanwan Soonthornphisaj, Kanokwan Chaikulseriwat, and Piyanan Tang-On. Anti-spam filtering: a centroid-based classification approach. In *Signal Processing, 2002 6th International Conference on*, volume 2, pages 1096–1099. IEEE, 2002.
 - [9] James D Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure hamiltonians using quantum computers. *Molecular Physics*, 109(5):735–750, 2011.
 - [10] Dave Wecker, Bela Bauer, Bryan K Clark, Matthew B Hastings, and Matthias Troyer. Can quantum chemistry be performed on a small quantum computer? *arXiv preprint arXiv:1312.1695*, 2013.
 - [11] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Quantum random access memory. *Physical review letters*, 100(16):160501, 2008.
 - [12] Harry Buhrman, Richard Cleve, John Watrous, and Ronald De Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.
 - [13] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. *arXiv preprint quant-ph/0005055*, 2000.
 - [14] Christoph Durr and Peter Hoyer. A quantum algorithm for finding the minimum. *arXiv preprint quant-ph/9607014*, 1996.
 - [15] Stuart Lloyd. Least squares quantization in pcm. *Information Theory, IEEE Transactions on*, 28(2):129–137, 1982.
 - [16] Robert Tibshirani, Trevor Hastie, Balasubramanian Narasimhan, and Gilbert Chu. Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of the National Academy of Sciences*, 99(10):6567–6572, 2002.
 - [17] Andrew M Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Information & Computation*, 12(11-12):901–924, 2012.
 - [18] Yann LeCun and Corinna Cortes. The mnist database of handwritten digits, 1998.

- [19] Michel Ledoux. *The concentration of measure phenomenon*, volume 89. AMS Bookstore, 2005.
- [20] Anil K Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010.
- [21] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [22] Olvi L Mangasarian, W Nick Street, and William H Wolberg. Breast cancer diagnosis and prognosis via linear programming. *Operations Research*, 43(4):570–577, 1995.
- [23] John Ross Quinlan, PJ Compton, KA Horn, and L Lazarus. Inductive knowledge acquisition: a case study. In *Proceedings of the Second Australian Conference on Applications of expert systems*, pages 137–156. Addison-Wesley Longman Publishing Co., Inc., 1987.
- [24] Yoshua Bengio. Learning deep architectures for AI. *Foundations and Trends in Machine Learning*, 2(1):1–127, 2009. Also published as a book. Now Publishers, 2009.
- [25] Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 20–29. ACM, 2003.
- [26] Andrew M Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 59–68. ACM, 2003.
- [27] Nathan Wiebe, Dominic W Berry, Peter Høyer, and Barry C Sanders. Simulating quantum dynamics on a quantum computer. *Journal of Physics A: Mathematical and Theoretical*, 44(44):445308, 2011.
- [28] Ashwin Nayak and Felix Wu. The quantum query complexity of approximating the median and related statistics. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 384–393. ACM, 1999.
- [29] Sylvester Eriksson-Bique, Mary Solbrig, Michael Stefanelli, Sarah Warkentin, Ralph Abbey, and Ilse CF Ipsen. Importance sampling for a monte carlo matrix multiplication algorithm, with application to information retrieval. *SIAM Journal on Scientific Computing*, 33(4):1689–1706, 2011.
- [30] Dénes Petz and Júlia Réffy. On asymptotics of large haar distributed unitary matrices. *Periodica Mathematica Hungarica*, 49(1):103–117, 2004.

Appendix A: Additional Numerical Experiments

We evaluate the performance of our nearest-neighbor (NN) and nearest-centroid (Centroid) algorithms on several additional machine learning tasks. A list of datasets, their respective training set sizes, and feature dimensions are listed in Table I. Each task is mapped to a binary classification problem (two classes). The data sets do not, in general, contain an equal number of training vectors per class. We denote the number of training vectors in classes A and B to be M_A and M_B , respectively.

The noise induced by inaccurate estimation of the distances is modeled by introducing Gaussian random noise with zero mean and variance ϵ^2 and then clipping the result to the interval $[0, \infty)$. Other distributions, such as uniformly distributed noise, gave qualitatively similar results. The features used in each data set can take on dramatically different value types. For example, the diabetes data set contains features such as patient age and blood pressure. In all tasks, we scale each feature to have zero mean and unit variance over the given data set.

We do not scale the vectors to unit length because the length of the vector is important for classification since points in one class are likely to be nearly co-linear with those in another class in low-dimensional spaces. Non-unit vectors can be easily accommodated by our algorithms by multiplying by the norms of the vectors in the inner-product based approach or by increasing the number of vectors used in the centroid approach. This also means that $|\mathbf{v}_j|$ will typically be on the order of \sqrt{N} , which suggests that for the data sets that we consider $|\mathbf{v}_j| \in [1, 10]$ is not unreasonable. Hence we will refer to the regime where $\epsilon \leq 1$ as the low-noise regime and $\epsilon \in (1, 10]$ as the high-noise regime.

We first evaluate our algorithms on a standard machine learning benchmark commonly referred to as the “half moon” dataset, which consists of two synthetically generated crescent-shaped clusters of points, as shown in Figure 4. The dataset challenges classification algorithms since the convex hulls of the two “moons” overlap and the mean value for each cluster (denoted by a star) sits in a region not covered by points. This data set will be hard to classify with centroid-based methods (using one cluster) because 14.3% of the data is closer to the centroid of the opposite set than to its own centroid. This means that the accuracy of centroid-based assignment will be at most 85.7%. In contrast, we expect nearest-neighbor classification to work well because the typical (Euclidean) distance between points is roughly 0.03, whereas the two classes are separated by a distance of approximately 0.5. This means that NN should succeed with near 100% probability, except in cases where the training set size is very small.

In Figure 5, we plot the accuracies of our nearest-neighbor algorithm (NN; blue squares) and our nearest-centroid algorithm (Centroid; red circles) as functions of noise ϵ in the distance computation. NN significantly outperforms Centroid in the low-noise regime, exhibiting an accuracy near 100% versus Centroid’s 86% accuracy. As the noise level increases, the accuracy of both algorithms decays; however, in the low noise regimes, NN outperforms Centroid with statistical significance. At high noise levels, both algorithms decay to 50% accuracy as expected.

Figure 6 shows accuracy as a function of training data size. Here the training data size is taken to be a fraction, f , of the 2000 vectors in the set and the remaining fraction, $1 - f$, of the 2000 vectors was used to test the accuracy of the assignments. Again, NN is almost always successful in classifying vectors; whereas Centroid achieves accuracies




FIG. 4: Half-moon data set, vectors are unnormalized. The two clusters of red and blue vectors correspond to the two classes used in the assignment set and the red and blue stars give the centroids of the corresponding cluster.




FIG. 5: Accuracy as a function of noise ϵ in distance computation for half-moon data. 50% of the data was used to train the classifier and the remaining 50% was used to test it.




FIG. 6: Accuracy as a function of training data size for half-moon data. Noise of $\epsilon = 10^{-5}$ was used here.

between 84 – 88%. Neither algorithm exhibits significant improvements in learning as the training set size is increased. This behavior indicates the difficulty of this classification task for Centroid.

There are of course other methods that can be employed in order to boost the success probability of centroid-based classification. The simplest is to cluster the data using a k -means clustering algorithm to subdivide each of the half moons into two or more clusters. This semi-supervised approach often works well, but can be expensive for certain representations of the data [20].

The next tasks that we consider consist of determining whether a given disease was present or not based on patient data. The diseases considered include breast cancer, heart disease, thyroid conditions, and diabetes. All data is taken from the UCL Machine Learning Repository [21]. Details on the features and data size are given in Table I.

	N , Number of Features	M , Number of Points	M_A	M_B	Year
Half Moon	2	2000	1000	1000	–
Breast Cancer [21, 22]	9	683	239	444	1992
Heart Disease (Statlog Data Set) [21]	13	270	120	150	1993
Thyroid [21, 23]	5	215	150	65	1987
Diabetes (Pima) [21]	7	532	177	355	1990

TABLE I: Evaluation datasets. Sizes of each data set for the conditions examined in Figure 7.

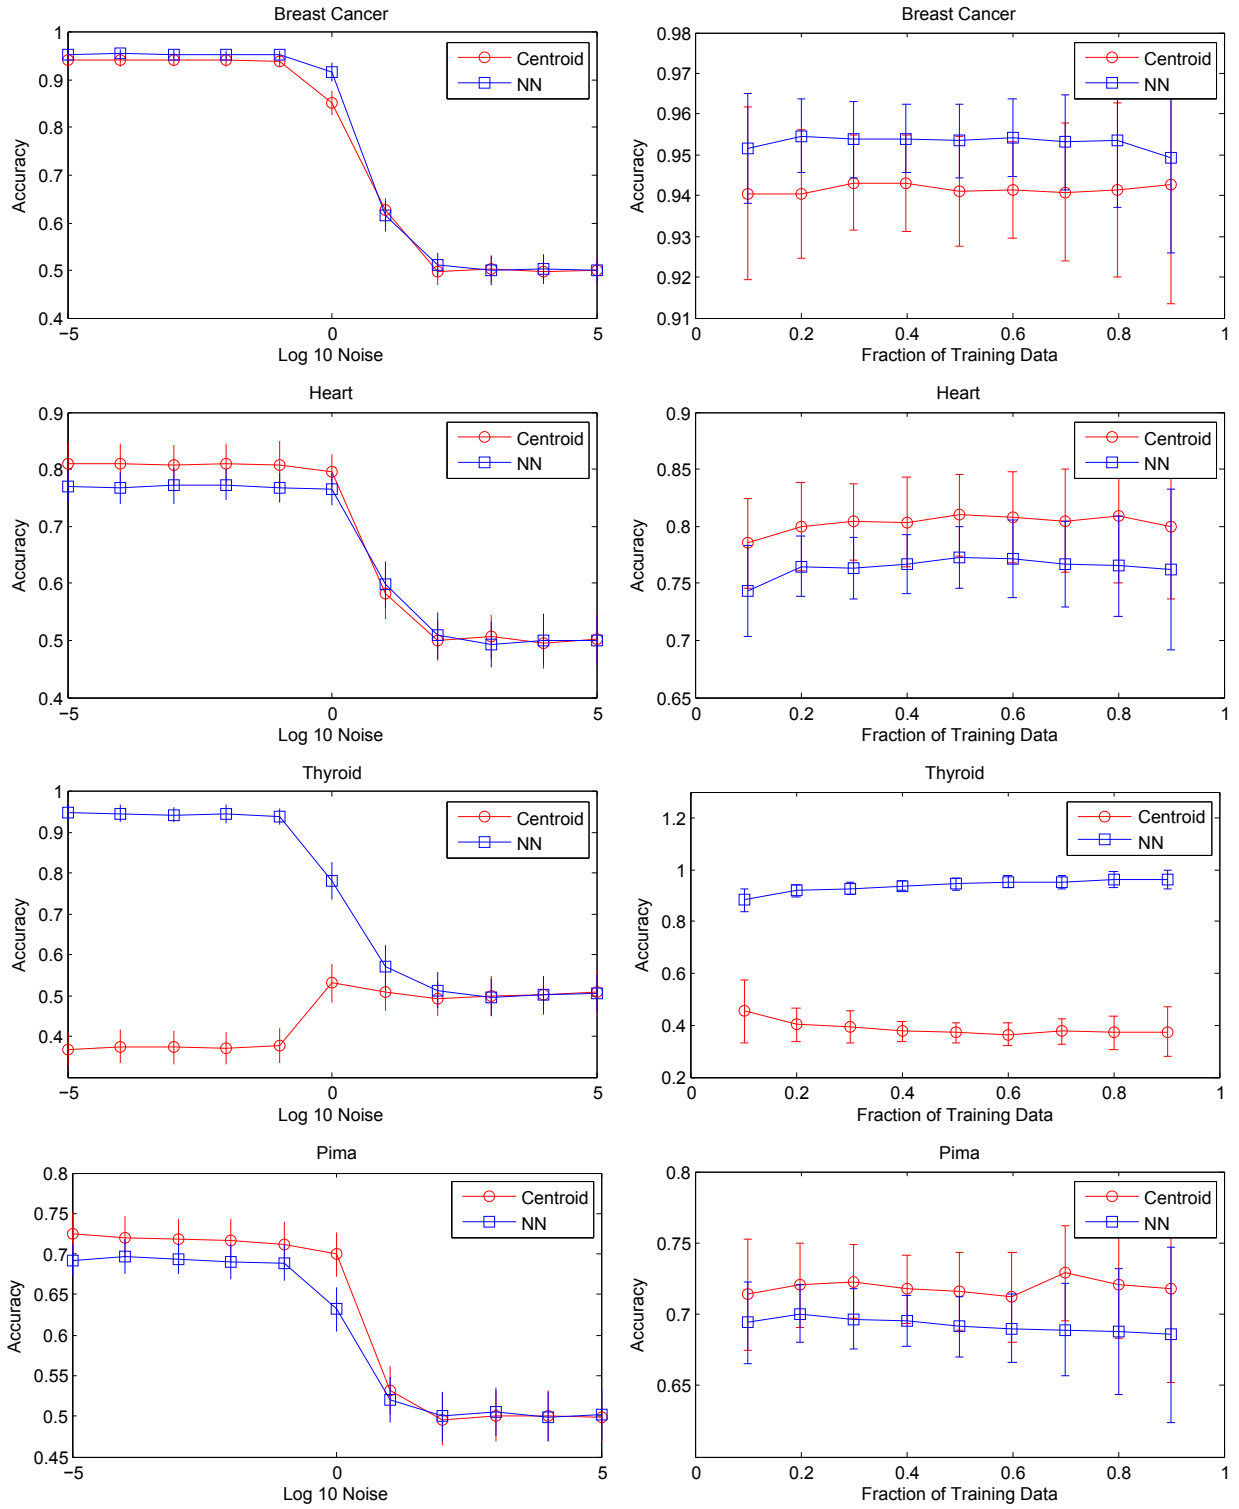


FIG. 7: (Left Column) Accuracy as a function of noise ϵ in the distance computation; (Right Column) Accuracy as a function of training set size for breast cancer (first row), heart disease (second row), thyroid (third row), and diabetes (fourth row) data. 50% of the data is used for training and the remainder for testing for all data in the left column. $\epsilon = 10^{-5}$ is taken for all data in the right column.

In some cases, we modified the data slightly. The breast cancer data, thyroid data, and the Pima diabetes study all contained instances of missing data. In each case we removed any vector that had a missing value. We also removed




FIG. 8: Accuracy as a function of noise ϵ in the distance computation and the fraction of the total data that is used for training for thyroid data set where the normalization step in the distances has been omitted. 50% of the data is used for training and the remainder for testing in the left plot. $\epsilon = 10^{-5}$ is taken for all data in the right plot.

boolean features from the thyroid and Pima diabetes data sets.

The left column of Figure 7 shows the accuracy of NN (blue squares) and Centroid (red circles) as a function of noise ϵ in the distance computations. The first row shows the accuracies on the breast cancer data. Both algorithms exhibit similarly high accuracies above 94% in the low-noise regime, with NN outperforming Centroid with significance only at $\epsilon = 1$. In the extreme noise regime, NN performs just slightly better than random as expected.

In the second and last rows, the accuracies for heart disease and diabetes data are shown. In these tasks, we find that in the low-noise regime, Centroid slightly outperforms NN, without statistical significance (except when $\epsilon = 1$). In the presence of high amounts of noise, both methods exhibit some learning; however, in all cases, learning is limited to around 55%.

In the third row, accuracy for the thyroid data is shown. NN exhibits significantly better accuracy of 90% as compared to less than 40% for Centroid. In this case, the centroid-based algorithm performed worse than random guessing. Poor accuracy is caused, in part, by our decision to divide the distance by the standard deviation in the distances as seen in Figure 8. We found that the variance of the hypothyroid cases (X_B) was high enough that the mean of the training vectors that tested negative for thyroid conditions (X_A) was within one standard deviation of it. In particular, $\sqrt{\mathbb{E}_{\mathbf{v} \in X_A} (|\mathbf{v} - \text{mean}(X_B)|_2^2) / \sigma_B} \approx 0.49$ and $\sqrt{\mathbb{E}_{\mathbf{v} \in X_B} (|\mathbf{v} - \text{mean}(X_A)|_2^2) / \sigma_A} \approx 4.4$. Thus this test will incorrectly assign vectors from X_A with high probability and correctly assign vectors from X_B with high probability. We therefore expect the accuracy to be roughly 30% since the probability of drawing a vector from X_B is roughly 65/215. This is close to the observed accuracy of $37\% \pm 4\%$.

The data in Figure 8, which forgoes normalizing the computed distances in Centroid, is devoid of these problems. For low noise, Centroid succeeds roughly 86% of the time and falls within statistical error of the NN data at $\epsilon \approx 1$. Also, we observe that the assignment accuracy increases for both methods as more training data is used. This is in stark contrast to the data in Figure 7; however, this does not imply that the centroid-based method is actually performing well. If we were to assign the data to class A every time, regardless of the distance, we would succeed with probability 70%. If Centroid is used, then the accuracy only increases by roughly 15%. Also, since the two clusters strongly overlap, distance to the centroid is not a trustworthy statistic on which to base classification. For these reasons, the use of Centroid to diagnose thyroid conditions, either with or without normalization, is inferior to using other methods.

The right column of Figure 7 shows the accuracy of our two algorithms as a function of training set size. In the breast cancer task (first row), we see that both NN and Centroid exhibit little variation in accuracy as the amount of training data increases. Similarly, in the heart disease and diabetes tasks (second and last rows), an increase in training data size does not imply significant increases in accuracy. However, in the thyroid task, we see some differences in learning between NN and Centroid as the training data size increases. NN's accuracy improves, from 85% to 96%, while Centroid's accuracy decreases slightly.

It is hard to determine in general why Centroid sometimes outperforms NN, but outliers in the data are frequently one reason. Outliers can cause problems for NN because it becomes increasingly likely as more training data is included that an outlier point from X_B is close to any given element of X_A . Thus increasing the training size can

actually be harmful for certain nearest-neighbor classification problems. Centroid is less sensitive to these problems because averaging over a data set reduces the significance of outliers. Such problems can be addressed in the case of NN by using k -nearest neighbor classifiers instead of nearest-neighbor classification [2]. Our quantum algorithms can be trivially modified to output the classes of each of the k closest vectors. Alternatively, such problems can also be addressed by using alternative machine learning strategies such as deep learning [24].

In summary, our numerical results indicate that classification accuracy, and in turn the best choice of algorithm, is highly dependent on the particular task and dataset. While nearest-neighbor classification appears to be the preferred algorithm on most of the tasks presented here, in practice, a highly non-linear combination of classification algorithms is more commonly used [24]. However, such classical approaches can be computationally expensive, in particular when classification over a large dataset is required. Our quantum algorithms for classification offer the advantage of fast classification in conjunction with high performance accuracy, and may enable accurate classification of datasets that otherwise classically would not be possible.

Appendix B: Proofs of Theorems 1 and 2

We present the proofs of [Theorem 1](#) and [Theorem 2](#) by way of a number of propositions that can be independently verified. We begin with preliminary results that show that the state preparations used in our algorithms are efficient. We then review known results on the performance of the quantum minimum finding algorithm and amplitude estimation. We present our coherent majority voting scheme and variant of the swap test and provide intermediate results needed to apply the Dürr Høyer algorithm and amplitude estimation coherently. We then use these results to prove [Theorem 1](#). Finally, we turn our attention to proving [Theorem 2](#) which uses many of the same techniques used to prove [Theorem 1](#), but in addition requires the introduction of new methods for computing the distances to the cluster centroids and the intra-cluster variance.

1. Preliminary Results

We begin by introducing a method to implement the operator V which is needed for our nearest-centroid classification algorithm.

Lemma 3. *A unitary V can be efficiently synthesized within error $O(\epsilon)$ on a quantum computer equipped with H (Hadamard), T ($\pi/8$) and CNOT gates such that*

$$|V_{j0}\rangle = \begin{cases} \frac{1}{\sqrt{2}}, & j = 0 \\ \frac{1}{\sqrt{2M}}, & \text{otherwise.} \end{cases}$$

Proof. Since H is unitary and Hermitian it is a straightforward exercise in Taylor's theorem to show that for any $t \in \mathbb{R}$

$$e^{-iH^{\otimes m}t} = \mathbb{1} \cos(t) - iH^{\otimes m} \sin(t). \quad (\text{B1})$$

Thus if we choose V to be e^{-iHt} for some fixed value of t then

$$V|0\rangle = \left(\cos(t) - i \frac{\sin(t)}{\sqrt{M+1}} \right) |0\rangle - i \frac{\sin(t)}{\sqrt{M+1}} \sum_{j>0} |j\rangle. \quad (\text{B2})$$

The value of t is found by setting $P(0)/P(j > 0) = 1/M$, which yields

$$t = \sin^{-1} \left(\sqrt{\frac{M+1}{2M}} \right). \quad (\text{B3})$$

Finally, H can be made sparse efficiently by

$$(HT^2)H(T^6H) = \begin{bmatrix} 0 & e^{i\pi/4} \\ e^{-i\pi/4} & 0 \end{bmatrix}, \quad (\text{B4})$$

and hence $H^{\otimes n}$ can be transformed into a one-sparse matrix by applying this basis transformation to each qubit. One-sparse matrices can be efficiently simulated [25–27], completing the proof of the lemma. \square





FIG. 9: Circuit for performing CMP illustrated for a single qubit inputs $|i_p\rangle$ and $|M_p\rangle$ after repeating this circuit n times, the lower most register will contain $|1\rangle$ if $i > M$.

Lemma 4. *The state $\frac{1}{\sqrt{M}} \sum_{j=1}^M |j\rangle$ can be prepared efficiently and deterministically using a quantum computer.*

Proof. Let $m = \lceil \log_2(M+1) \rceil$, and let $|M\rangle$ be the computational basis state that stores M as a binary string. The proof follows from the fact that the following circuit



prepares the desired state given measurement outcomes of 0 and 0, which occurs with probability $\frac{M}{2^m}$. Here the operation CMP obeys

$$\text{CMP } |i\rangle |M\rangle |0\rangle = \begin{cases} i \leq M, & |i\rangle |M\rangle |0\rangle \\ i > M, & |i\rangle |M\rangle |1\rangle. \end{cases} \quad (\text{B5})$$

Here CMP can be implemented using the circuit in [Figure 9](#). Hence the state can be prepared efficiently and with high probability if measurements are used. Also, note that the quantum control on the value of M_p can be replaced with classical control in cases where a quantum superposition over different values of M is not needed.

Since the success probability is known, the success probability can be boosted to certainty through amplitude amplification which requires $\Theta\left(\sqrt{\frac{2^m}{M}}\right) = \Theta(1)$ applications of CMP [13]. This means that the measurement can be removed in the state preparation step without sacrificing the efficiency of the algorithm. \square

Another important result is the method of Dürr and Høyer which is given as the following lemma [14].

Lemma 5 (Dürr Høyer). *The expected number of Grover iterations needed to learn $\min\{y_i : i = 1, \dots, M\}$ is bounded above by*

$$\frac{45}{2} \sqrt{M}.$$

Nayak and Wu show that this algorithm is also near-optimal by providing a matching lower bound of $\Omega(\sqrt{M})$ for minimum finding (proven using the polynomial method) [28].

Similarly, we also use the amplitude estimation result of Brassard et al. [13] to estimate the amplitude squared of a marked component of a quantum state, which we denote as a . The algorithm works by applying the phase estimation algorithm to an operator Q , which performs an iteration of Grover's algorithm where we wish to estimate the amplitude of the marked state. We provide a circuit for amplitude estimation in [Figure 10](#). The following theorem shows that amplitude estimation can learn the resultant probabilities quadratically faster than statistical sampling.

Theorem 6 (Brassard, Høyer, Mosca and Tapp). *For any positive integers k and L , the amplitude estimation algorithm of [13] outputs \tilde{a} ($0 \leq \tilde{a} \leq 1$) such that*

$$|\tilde{a} - a| \leq 2\pi k \frac{\sqrt{a(1-a)}}{L} + \left(\frac{\pi k}{L}\right)^2$$

with probability at least $8/\pi^2$ when $k = 1$ and with probability greater than $1 - 1/(2(k-1))$ for $k \geq 2$. It uses exactly L iterations of Grover's algorithm. If $a = 0$ then $\tilde{a} = 0$ with certainty, and if $a = 1$ and M is even, then $\tilde{a} = 1$ with certainty.




FIG. 10: Quantum circuit for amplitude estimation where F_L is the L -dimensional Fourier transform and the controlled Q^j operator applies j Grover iterations to the target state if the top most register is $|j\rangle$

2. Proof of Theorem 1

In order to find the minimum value of a set of different quantum amplitudes using [Lemma 5](#) we need to be able to perform iterations of Grover's algorithm using the result of [Theorem 6](#). This cannot be done directly (with high probability) because the traditional approach to amplitude estimation is not reversible. We provide below a reversible algorithm that uses a coherent form of majority voting to obtain a reversible analog for algorithms like amplitude estimation.

Lemma 7. *Let \mathcal{A} be a unitary operation that maps $|0^{\otimes n}\rangle \mapsto \sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle$ for $1/2 < |a_0| \leq |a| \leq 1$ using Q queries then there exists a deterministic algorithm such that for any $\Delta > 0$ there exists an integer k and a state $|\Psi\rangle$ can be produced that obeys $\| |\Psi\rangle - |0^{\otimes n k}\rangle |y\rangle \|_2 \leq \sqrt{2\Delta}$ using a number of queries bounded above by*

$$2Q \left\lceil \frac{\ln(1/\Delta)}{2(|a_0| - \frac{1}{2})^2} \right\rceil.$$

Proof. The basic idea behind the algorithm is to prepare k copies of the state $\sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle$, and then coherently compute the median via a reversible circuit and uncompute the k resource states used to find the median of the values of y . First, let \mathcal{M} be a circuit that performs

$$\mathcal{M} : |y_1\rangle \cdots |y_k\rangle |0\rangle \mapsto |y_1\rangle \cdots |y_k\rangle |\bar{y}\rangle, \quad (\text{B6})$$

where we use \bar{y}_k to denote the median. This transformation can be performed by implementing a sort algorithm using $O(kn \log(k))$ operations and hence is efficient.

The initial state for this part of the protocol is of the form

$$(\sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle)^{\otimes k}.$$

We can therefore partition the k -fold tensor product as a sum of two disjoint sets: the sum of states with median y and another sum of states with median not equal to y . We denote these two sums as $|\Psi\rangle$ and $|\Phi\rangle$ respectively, which is equivalent to expressing

$$(\sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle)^{\otimes k} := A|\Psi\rangle + \sqrt{1-|A|^2}|\Phi\rangle, \quad (\text{B7})$$

for some value of A . A direct consequence of [\(B7\)](#) is that there exists (a possibly entangled state) $|\Phi; y^\perp\rangle$ such that

$$\mathcal{M}(\sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle)^{\otimes k} = A|\Psi\rangle |y\rangle + \sqrt{1-|A|^2}|\Phi; y^\perp\rangle, \quad (\text{B8})$$

where $A|\Psi\rangle + \sqrt{1-|A|^2}|\Phi\rangle := (\sqrt{a}|y\rangle + \sqrt{1-a}|y^\perp\rangle)^{\otimes k}$ and $|\Psi\rangle$ represents the subspace where the median is y and $|\Phi\rangle$ is its complement. Our goal is now to show that $|A|^2 > 1 - \Delta$ for k sufficiently large.

To see this, let us imagine measuring the first register in the computational basis. The probability of obtaining p results, from the k resulting bit strings, that are not y is given by the binomial theorem:

$$P(p) = \binom{k}{p} |a|^p |1-a|^{k-p}. \quad (\text{B9})$$

Now we can compute the probability that a measurement of the last register will not yield y by observing the fact that in any sequence of measurements that contains more than $k/2$ y -outcomes, the median must be $k/2$. Therefore

the probability that the computed value of the median is not y is at most the probability that the measured results contain no more than $k/2$ y outcomes. This is given by (B9) to be

$$P(y^\perp) \leq \sum_{p=0}^{\lfloor k/2 \rfloor} \binom{k}{p} |a|^p |1 - |a||^{k-p}. \quad (\text{B10})$$

Using Hoeffding's inequality on (B10) and $|a| \geq |a_0| > 1/2$ we find that

$$P(y^\perp) \leq \exp\left(\frac{-2(k|a| - \frac{k}{2})^2}{k}\right) = \exp\left(-2k\left(|a_0| - \frac{1}{2}\right)^2\right). \quad (\text{B11})$$

Eq. (B11) therefore implies that $P(y^\perp) \leq \Delta$ if

$$k \geq \frac{\ln(\frac{1}{\Delta})}{2(|a_0| - \frac{1}{2})^2}. \quad (\text{B12})$$

Next, by applying $\mathcal{A}^{\dagger \otimes k}$ to the first register, we obtain

$$\begin{aligned} \mathcal{A}^{\dagger \otimes k} \left(A |\Psi\rangle |y\rangle + \sqrt{1 - |A|^2} |\Phi; y^\perp\rangle \right) &= \mathcal{A}^{\dagger \otimes k} \left(A |\Psi\rangle |y\rangle + \sqrt{1 - |A|^2} |\Phi\rangle |y\rangle \right) + \mathcal{A}^{\dagger \otimes k} \left(\sqrt{1 - |A|^2} (|\Phi; y^\perp\rangle - |\Phi\rangle |y\rangle) \right) \\ &= |0^{\otimes nk}\rangle |y\rangle + \mathcal{A}^{\dagger \otimes k} \left(\sqrt{1 - |A|^2} (|\Phi; y^\perp\rangle - |\Phi\rangle |y\rangle) \right). \end{aligned} \quad (\text{B13})$$

Note that $\langle y | y^\perp \rangle = 0$ and hence $|\Phi; y^\perp\rangle$ is orthogonal to $|\Phi\rangle |y^\perp\rangle$. If $|\cdot|$ is taken to be the 2-norm then (B13) gives that

$$\left| \mathcal{A}^{\dagger \otimes k} \left(A |\Psi\rangle |y\rangle + \sqrt{1 - |A|^2} |\Phi; y^\perp\rangle \right) - |0^{\otimes nk}\rangle |y\rangle \right| \leq \sqrt{2(1 - |A|^2)} \leq \sqrt{2\Delta}, \quad (\text{B14})$$

since $P(y^\perp) := 1 - |A|^2 \leq \Delta$ for k chosen as per (B12). The result then follows after noting that k must be chosen to be an integer and that the total number of queries made to prepare the state is $2Qk$. \square

Lemma 7 shows that coherent majority voting can be used to remove the measurements used in algorithms such as amplitude estimation at the price of introducing a small amount of error in the resultant state. We can use such a protocol in the Dürr Høyer algorithm to find the minimum value of all possible outputs of the algorithm, as shown in the following corollary.

Corollary 8. *Assume that for any $j = 1, \dots, M$, a unitary transformation $|j\rangle |0^{\otimes n}\rangle \mapsto \sqrt{a} |y_j\rangle + \sqrt{1 - |a|} |y_j^\perp\rangle$ for $1/2 < |a_0| \leq |a| \leq 1$ can be performed using Q queries then the expected number of queries made to find $\min_j y_j$ with failure probability at most δ_0 is bounded above by*

$$90\sqrt{M}Q \left\lceil \frac{\ln\left(\frac{81M(\ln(M)+\gamma)}{\delta_0}\right)}{2(|a_0| - \frac{1}{2})^2} \right\rceil.$$

Proof. **Lemma 5** states that at most $45\sqrt{M}/2$ applications of Grover's search are required, which requires $45\sqrt{M}$ queries to an (approximate) oracle that prepares each y_j since two queries are required per Grover iteration. **Lemma 7** therefore says that the cost of performing this portion of the algorithm is

$$N_{\text{queries}} \leq 90\sqrt{M}Q \left\lceil \frac{\ln(1/\Delta)}{2(|a_0| - \frac{1}{2})^2} \right\rceil. \quad (\text{B15})$$

Next, we need to find a value of Δ that will make the failure probability for this approximate oracle at most δ_0 . Now let us assume the worst case scenario that if the measurement of y_j fails to output the desired value even once, then the entire algorithm fails. We then upper bound the probability of failure by summing the probability of failure in each of the steps in the search. Assuming that the algorithm is searching for an element of rank at least r then the number of calls to the oracle yielding y_j is at most [13]

$$9\sqrt{\frac{M}{r-1}}.$$

This means that the amplitude of the erroneous component of the state (using subadditivity of quantum errors) is at most

$$9\sqrt{\frac{\Delta M}{r-1}}.$$

The worst case scenario is that the algorithm must search through all M entries (this is extremely unlikely if M is large because the average complexity is $O(\sqrt{M})$). This means that the probability of at least one failed observation occurring is at most

$$\sum_{r=2}^M \frac{81\Delta M}{r-1} = 81MH_{M-1} \leq 81M(\ln(M) + \gamma). \quad (\text{B16})$$

Here H_{M-1} is the $M-1$ th harmonic number and γ is Euler's constant. Therefore if we want the total probability of error to be at most δ_0 then it suffices to choose

$$\Delta = \frac{\delta_0}{81M(\ln(M) + \gamma)}. \quad (\text{B17})$$

Then combining (B15) and (B17) gives us that the average query complexity obeys

$$N_{\text{queries}} \leq 90\sqrt{M}Q \left[\frac{\ln\left(\frac{81M(\ln(M)+\gamma)}{\delta_0}\right)}{2(|a_0| - 1/2)^2} \right]. \quad (\text{B18})$$

□

Note that we want to maximize the value of $\sin^2(\pi y_j/R)$ that is yielded by the amplitude estimation algorithm for each j . This maximization is equivalent to minimizing $|R/2 - y_j|$. Given that y_j is returned coherently by our de-randomized amplitude estimation circuit, a measurement-free circuit can be used that computes $|R/2 - y_j|$ for any input y_j . This requires no further oracle calls. Hence [Corollary 8](#) applies to our circumstances with no further modification.

The results of [Theorem 1](#) and [Theorem 2](#) follow directly from [Corollary 8](#) by substitution of appropriate values of Q and $|a_0|$. The remaining work focuses on devising an appropriate state preparation algorithm that can be used in [Theorem 6](#).

Lemma 9. *Let \mathbf{v}_j be d -sparse and assume that the quantum computer has access to O and F then a unitary transformation exists that can be implemented efficiently using 3 oracle calls and, for all j , maps*

$$|j\rangle |0\rangle \mapsto \frac{1}{\sqrt{d}} |j\rangle \sum_{i=1}^d |f(j, i)\rangle \left(\sqrt{1 - \frac{r_{jf(j,i)}^2}{r_{j\max}^2}} e^{-i\phi_{jf(j,i)}} |0\rangle + \frac{r_{jf(j,i)} e^{i\phi_{jf(j,i)}}}{r_{j\max}} e^{i\phi_{jf(j,i)}} |1\rangle \right).$$

Proof. We begin by preparing the state

$$\frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |i\rangle |0\rangle |0\rangle, \quad (\text{B19})$$

which can be prepared reversibly and efficiently by applying [Lemma 4](#).

The next step is to apply the oracle F to the result, this performs

$$\frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |i\rangle |0\rangle |0\rangle \mapsto \frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |f(j, i)\rangle |0\rangle |0\rangle. \quad (\text{B20})$$

Then querying O implements

$$\frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |f(j, i)\rangle |0\rangle |0\rangle \mapsto \frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |f(j, i)\rangle |v_{j,f(j,i)}\rangle |0\rangle. \quad (\text{B21})$$




FIG. 11: The swap test [12]. The probability of measuring the top qubit to be zero is $1/2 + |\langle\phi|\psi\rangle|^2/2$, which allows statistical testing to be used to efficiently discriminate the states.

By applying $R_y(2\sin^{-1}(r_{jf(j,i)}/r_{j\max}))$ on the final qubit in (B21), we obtain

$$\frac{1}{\sqrt{d}} \sum_{i=1}^d |j\rangle |f(j,i)\rangle |v_{j,f(j,i)}\rangle |0\rangle \mapsto \frac{1}{\sqrt{d}} |j\rangle \sum_{i=1}^d |f(j,i)\rangle |v_{j,f(j,i)}\rangle \left(\sqrt{1 - \frac{r_{jf(j,i)}^2}{r_{j\max}^2}} |0\rangle + \frac{r_{jf(j,i)}}{r_{j\max}} |1\rangle \right). \quad (\text{B22})$$

The result then follows by applying $R_z(2\phi_{jf(j,i)})$ to the last qubit in (B22) and using O^\dagger to clean the ancilla register containing $|v_{j,f(j,i)}\rangle$. Three queries are used in this process. \square

Next we use the swap test to provide a method to compute the inner product between two vectors. The test is implemented by the circuit in Figure 11 for arbitrary states $|\phi\rangle$ and $|\psi\rangle$. The resultant state before measurement is

$$\frac{1}{2} |0\rangle (|\phi\rangle |\psi\rangle + |\psi\rangle |\phi\rangle) + \frac{1}{2} |1\rangle (|\phi\rangle |\psi\rangle - |\psi\rangle |\phi\rangle),$$

and the probability of measuring the first qubit to be 1 is $1/2 - |\langle\phi|\psi\rangle|^2/2$. We do not ignore this measurement since we want to use the swap test within the Grover iterations used in Theorem 6.

Lemma 10. *For any fixed $\epsilon > 0$ and any pair of d -sparse unit vectors $\mathbf{u} \in \mathbb{C}^n$ and $\mathbf{v}_j \in \mathbb{C}^n$ a state of the form $\sqrt{|A|} |\Psi\rangle |y\rangle + \sqrt{1-|A|} |\Phi; y_\perp\rangle$ can be efficiently prepared where y encodes $|\langle\mathbf{u}|\mathbf{v}_j\rangle|^2$ within error ϵ and $|A| \geq 8/\pi^2$ using a number of queries that is bounded above by*

$$Q \leq 12 \left\lceil \frac{4\pi(\pi+1)d^2 r_{0\max}^2 r_{j\max}^2}{\epsilon} \right\rceil,$$

where $|\langle i|\mathbf{v}_j\rangle| \leq r_{j\max}$ for any $i \geq 0$ and d is a power of 2.

Proof. Lemma 9 provides a method for constructing the states

$$|\psi\rangle := \frac{1}{\sqrt{d}} \sum_{i=1}^d |f(j,i)\rangle |0^{\otimes n'}\rangle \left(\sqrt{1 - \frac{r_{jf(j,i)}^2}{r_{j\max}^2}} e^{-i\phi_{jf(j,i)}} |0\rangle + \frac{v_{jf(j,i)}}{r_{j\max}} |1\rangle \right) |1\rangle. \quad (\text{B23})$$

$$|\phi\rangle := \frac{1}{\sqrt{d}} \sum_{i=1}^d |f(0,i)\rangle |0^{\otimes n'}\rangle |1\rangle \left(\sqrt{1 - \frac{r_{0f(0,i)}^2}{r_{0\max}^2}} e^{-i\phi_{0f(0,i)}} |0\rangle + \frac{v_{0f(0,i)}}{r_{0\max}} |1\rangle \right). \quad (\text{B24})$$

We then see that

$$\langle\phi|\psi\rangle = \frac{1}{d} \sum_i \frac{v_{ji} v_{0i}^*}{r_{j\max} r_{0\max}} = \frac{\langle v_0 | v_j \rangle}{dr_{j\max} r_{0\max}}. \quad (\text{B25})$$

Note that we do not directly apply the swap test between the two states here because an undesirable contribution to the inner product of the form $\sqrt{1 - \frac{r_{jf(j,i)}^2}{r_{j\max}^2}} e^{-i\phi_{jf(j,i)}} \sqrt{1 - \frac{r_{0f(0,i)}^2}{r_{0\max}^2}} e^{i\phi_{0f(0,i)}}$ would arise. We remove the possibility of such terms appearing by adding an ancilla qubit in (B23) and (B24) that selects out only the component that gives us information about the inner product of \mathbf{v}_j and \mathbf{u} .

The probability of measuring 0 in the swap test is $|A| = \frac{1}{2}(1 + |\langle\phi|\psi\rangle|^2)$, which implies that

$$(2|A| - 1)d^2 r_{j\max}^2 r_{0\max}^2 = |\langle\mathbf{u}|\mathbf{v}_j\rangle|^2. \quad (\text{B26})$$

At this point we could learn A by sampling from the distribution given by the swap test, but it is more efficient to use amplitude estimation. Amplitude estimation, in effect, uses a controlled Grover search oracle. In this case the Grover oracle requires that we implement a reflection about the initial state and also reflect about the space orthogonal to the target state. We refer to this reflection as S_χ . Unlike Grover's problem, the reflection about the target state is trivial since the target state is obtained when the swap test yields 0. This means that no oracle calls are needed to implement the reflection about the final state. The reflection about the initial state is of the form $\mathcal{A}S_0\mathcal{A}^\dagger$, where \mathcal{A} is an algorithm that maps $|0^{2n'+2n+4}\rangle \rightarrow |\phi\rangle|\psi\rangle$, and S_0 is of the form

$$S_0|x\rangle = \begin{cases} |x\rangle & , x \neq 0 \\ -|x\rangle & , x = 0 \end{cases}. \quad (\text{B27})$$

This can be implemented using a multi-controlled Z gate, and hence is efficient. The prior steps show that \mathcal{A} can be implemented using 6 oracle calls: [Lemma 9](#) implies that three queries are needed for the preparation of $|\psi\rangle$ and three more are needed for the preparation of $|\phi\rangle$. This implies that a step of Grover's algorithm, given by $\mathcal{A}S_0\mathcal{A}^\dagger S_\chi$, can be implemented using 12 oracle queries.

Amplitude estimation requires applying an iteration of Grover's algorithm in a controlled fashion at most R times (where R is the Hilbert-space dimension of the register that contains the output of the amplitude estimation algorithm). The controlled version of $\mathcal{A}S_0\mathcal{A}^\dagger S_\chi$ requires no additional oracle calls because $\mathcal{A}S_0\mathcal{A}^\dagger$ can be made controllable by introducing a controlled version of S_0 and using a controlled R_y rotation in the implementation of \mathcal{A} (given by [Lemma 9](#)); furthermore, both S_0 and S_χ do not require oracle calls and hence can be made controllable at no additional cost.

The error in the resultant estimate of $P(0)$ after applying amplitude estimation is with probability at least $8/\pi^2$ at most [\[13\]](#)

$$|A - \tilde{A}| \leq \frac{\pi}{R} + \frac{\pi^2}{R^2}. \quad (\text{B28})$$

This means that $|P(0) - \tilde{P}(0)| \leq \delta$ if

$$R \geq \frac{\pi(\pi + 1)}{\delta}. \quad (\text{B29})$$

Now given that we want the total error, with probability at least $8/\pi^2$, to be $\epsilon/2$ (the factor of $1/2$ is due to the fact that the calculation of $\sin^{-1}(r_{ji}/r_{j\max})$ is inexact) then [\(B26\)](#) gives us that choosing δ to be

$$\delta = \frac{\epsilon}{4d^2r_{0\max}^2r_{j\max}^2}, \quad (\text{B30})$$

is sufficient. Eq. [\(B30\)](#) then gives us that it suffices to take a number of steps of Grover's algorithm that obeys

$$R \geq \left\lceil \frac{4\pi(\pi + 1)d^2r_{0\max}^2r_{j\max}^2}{\epsilon} \right\rceil. \quad (\text{B31})$$

Since each Grover iteration requires 12 applications of the oracles F or O the query complexity of the algorithm is

$$Q = 12R \leq 12 \left\lceil \frac{4\pi(\pi + 1)d^2r_{0\max}^2r_{j\max}^2}{\epsilon} \right\rceil, \quad (\text{B32})$$

as claimed.

Now lastly, we need to show that the error in the resultant probabilities from inexactly evaluating $\sin^{-1}(r_{ji}/r_{j\max})$ can be made less than $\epsilon/2$ at polynomial cost. As shown previously, if $n' \in O(\log(1/\epsilon))$ and if $|\tilde{\phi}\rangle$ and $|\tilde{\psi}\rangle$ are the approximate versions of the two inner products then

$$\left| \left| \langle \phi | \psi \rangle \right|^2 - \left| \langle \tilde{\phi} | \tilde{\psi} \rangle \right|^2 \right| \in O(\epsilon).$$

This means that since $\sin^{-1}(\cdot)$ is efficient (it can be computed easily using a Taylor series expansion or approximation by Chebyshev polynomials), the cost of making the numerical error sufficiently small is at most polynomial. \square

The proof of [Theorem 1](#) now trivially follows:


Proof of Theorem 1. Proof follows as an immediate consequence of [Corollary 8](#) and [Lemma 10](#) using $|a_0| \geq 8/\pi^2$ and $r_{j,\max} \leq r_{\max}$. \square

3. Proof of Theorem 2

The structure of the proof of [Theorem 2](#) is similar to that of [Theorem 1](#). The biggest difference is that the swap test is not used to compute the Euclidean distance in this case. We use the method of Childs and Wiebe [\[17\]](#) to perform this state preparation task. To see how their method works, let us assume that we have access to an oracle U such that for any j

$$U |j\rangle |0\rangle = |j\rangle |v_j\rangle. \quad (\text{B33})$$

Then for any unitary $V \in \mathbb{C}^{\log_2 M \times \log_2 M}$, the following circuit



has the property that the probability of the measurement yielding 0 is $\|\sum_{j=0}^M |V_{j,0}|^2 \mathbf{v}_j\|^2$. Thus the Euclidean distance can be computed by this approach by setting $|v_0\rangle = -|u\rangle$ and choosing the unitary V appropriately. We employ a variant of this in the following lemma, which explicitly shows how to construct the required states.

Lemma 11. *For any fixed $\epsilon > 0$, the quantities a state of the form $\sqrt{|A|} |\Psi\rangle |y\rangle |y'\rangle + \sqrt{1-|A|} |\Phi_{\text{bad}}\rangle$ can be efficiently and reversibly prepared such that y encodes $\| -|u\rangle + \frac{1}{M} \sum_j |v_j\rangle \|_2^2$ within error ϵ and y' encodes $\frac{1}{M} \sum_p \| -|v_p\rangle + \frac{1}{M} \sum_j |v_j\rangle \|_2^2$ within error ϵ such that $|A| \geq 64/\pi^4 \approx 2/3$ using a number of queries bounded above by*


$$Q \leq 10 \left\lceil \frac{8\pi(\pi+1)dr_{\max}^2}{\epsilon} \right\rceil,$$

where $r_{\max} \geq \max_j r_j$.

Proof. First let us define an oracle W such that

$$W |j\rangle |p\rangle |i\rangle |0\rangle |0\rangle = \begin{cases} |0\rangle |p\rangle |i\rangle |v_{p,i}\rangle |v_{0,i}\rangle, & j = 0 \\ |j\rangle |p\rangle |i\rangle |v_{j,i}\rangle |v_{p,i}\rangle, & \text{otherwise} \end{cases} \quad (\text{B34})$$

A query to W can be implemented via



We see that the following transformation is efficient and can be performed using one query to F by applying V to the first register, and applying [Lemma 4](#) to the second and third registers

$$|0^{\otimes n}\rangle |0^{\otimes m}\rangle |0^{\otimes m}\rangle |0^{\otimes n'}\rangle |0\rangle \mapsto \frac{1}{\sqrt{Md}} \sum_{j=1}^M \sum_{p=1}^M \sum_{i=1}^d V_{j0} |j\rangle |f(j,i,p)\rangle |p\rangle |0^{\otimes n'}\rangle |0\rangle. \quad (\text{B35})$$

Here we take $f(j,i,p) = f(j,i)$ if $j \geq 1$ and $f(j,i,p) = f(p,i)$ if $j = 0$; furthermore, we use the convention that $|v_0\rangle = -\mathbf{u}$ and $|v_0^{(p)}\rangle = -\mathbf{v}_p$. We also take $m = \lceil \log_2 M \rceil$ and n' to be the number of bits needed to store the components of each \mathbf{v}_j .

The following state can be implemented efficiently within error at most $\epsilon/2$ using 3 oracle calls by applying [Lemma 9](#) with the modification that W is used in the place of O and then applying V^\dagger :

$$\frac{1}{\sqrt{Md}} \sum_{q,j=0}^M \sum_{i=1}^d \sum_{p=1}^M V_{qj}^\dagger V_{j0} |q\rangle |f(j,i,p)\rangle |p\rangle \left(\sqrt{1 - \left(\frac{r_{jf(j,i,p)}^{(p)}}{r_{\max}} \right)^2} e^{-i\phi_{jf(j,i,p)}^{(p)}} |0\rangle + \frac{r_{jf(j,i,p)}^{(p)}}{r_{\max}} e^{i\phi_{jf(j,i,p)}^{(p)}} |1\rangle \right) |\Theta(j,i,p)\rangle, \quad (\text{B36})$$

where $|\Theta(j, i, p)\rangle$ is a computational basis state that stores the ancilla qubits prepared by W . The qubit register containing $|\Theta(j, i, p)\rangle$ does not need to be cleaned since these qubits do not affect the trace.

We then use the definition of V in [Lemma 3](#) the probability of measuring the first register in the state $|0^{\otimes n}\rangle$ and the second-last register in the state $|1\rangle$ is:

$$\begin{aligned} & \text{Tr} \left(\frac{1}{dMr_{\max}^2} \sum_{i, i'} \sum_{j, j'} \sum_{p, p'} V_{j0} V_{0j}^\dagger V_{j'0}^* V_{0j'}^{\dagger*} V_{0j, i} v_{j, p}^* v_{j', i} \langle i | \langle i' | \otimes | p \rangle \langle p' | \right) \\ &= \left(\frac{1}{dMr_{\max}^2} \sum_i \sum_{j, j'} \sum_p V_{j0} V_{0j}^\dagger V_{j'0}^* V_{0j'}^{\dagger*} \langle i | v_j^{(p)} \rangle \langle v_{j'}^{(p)} | i \rangle \right) \\ &= \left(\frac{1}{Mdr_{\max}^2} \sum_p \left| \frac{-1}{2} |v_p\rangle + \frac{1}{2M} \sum_{j \geq 1} |v_j\rangle \right|^2 \right) \end{aligned} \quad (\text{B37})$$

We drop the state $|\Theta(j, i, p)\rangle$ above because it does not affect the trace. Thus the mean square distance between each \mathbf{v}_j and the centroid is

$$\frac{1}{M} \sum_p \left| -|v_p\rangle + \frac{1}{M} \sum_{j \geq 1} |v_j\rangle \right|^2 = 4dr_{\max}^2 P(0). \quad (\text{B38})$$

Three queries are needed to draw a sample from this distribution.

The distance can be computed similarly, except O can be queried directly instead of W and the “p”-register can be eliminated since we do not need to average over different distances. This saves one additional query, and hence it is straightforward to verify that the relationship between the distance squared and the probability of success is

$$\left| -|v_p\rangle + \frac{1}{M} \sum_{j \geq 1} |v_j\rangle \right|^2 = 4dr_{\max}^2 P(0). \quad (\text{B39})$$

Two queries are needed to draw a sample from this distribution.

Similar to the proof of [Lemma 10](#), we use amplitude estimation to estimate $P(0)$ in both cases. By following the same arguments used in [\(B28\)](#) to [\(B32\)](#) coherent AE can be used to prepare a state with probability of the form $\sqrt{|A|} |\Psi_{\text{good}}\rangle |y'\rangle + \sqrt{1 - |A|} |\Psi_{\text{bad}}; y'_\perp\rangle$ where $|A| \leq 8/\pi^2$ using a number of queries bounded above by

$$6 \left\lceil \frac{8\pi(\pi + 1)dr_{\max}^2}{\epsilon} \right\rceil. \quad (\text{B40})$$

Similarly, the cost of preparing a state of the form $\sqrt{|A|} |\Psi\rangle |y\rangle + \sqrt{1 - |A|} |\Phi; y_\perp\rangle$ where $|A| \leq 8/\pi^2$ and $|y\rangle$ encodes $|\mathbf{u} - \text{mean}(\{\mathbf{v}_j\})|$ is

$$4 \left\lceil \frac{8\pi(\pi + 1)dr_{\max}^2}{\epsilon} \right\rceil. \quad (\text{B41})$$

Therefore, by combining [\(B40\)](#) and [\(B42\)](#) we see that a state of the form $\sqrt{|A|} |\Psi\rangle |y\rangle |y'\rangle + \sqrt{1 - |A|} |\Phi_{\text{bad}}\rangle$ can be constructed where $|A| \leq (8/\pi^2)^2 \approx 2/3$ using a number of oracle calls bounded above by

$$10 \left\lceil \frac{8\pi(\pi + 1)dr_{\max}^2}{\epsilon} \right\rceil, \quad (\text{B42})$$

□

[Lemma 11](#) not only provides an essential step towards our proof of [Theorem 2](#) but it also provides an upper bound for the query complexity for nearest-centroid classification using $M' = 1$. It also gives an upper bound for the query complexity of the un-normalized centroid-based algorithm in [\[6\]](#). We give these bounds explicitly in the following corollary.

Corollary 12. Let \mathbf{v}_0 and $\{\mathbf{v}_j : j = 1, \dots, M\}$ be d -sparse unit vectors such that the components satisfy $\max_{j,i} |v_{ji}| \leq r_{\max}$. The task of finding

$$\left\| \mathbf{v}_0 - \frac{1}{M} \sum_{j=1}^M \mathbf{v}_j \right\|_2^2,$$

with error bounded above by ϵ and with success probability at least $1 - \delta_0$ requires an expected number of queries that is bounded above by

$$\left\lceil \frac{64 \ln(1/\delta_0)}{\pi^2(8/\pi^2 - 1/2)^2} \right\rceil \left\lceil \frac{8\pi(\pi + 1)dr_{\max}^2}{\epsilon} \right\rceil.$$

Proof. (B42) gives the cost of solving this problem within error ϵ and success probability $8/\pi^2$. The Chernoff bound then gives us that we can boost this success probability to $1 - \delta_0$ by repeating the experiment N_{samp} times where

$$N_{\text{samp}} = \frac{16 \ln(1/\delta_0)}{\pi^2(8/\pi^2 - 1/2)^2}. \quad (\text{B43})$$

The result then follows by multiplying (B42) by (B43) and taking the ceiling function of the pre-factor. \square

Theorem 2 is proved similarly to Corollary 12, with the exception that coherent amplitude amplification and the Durr Høyer minimization algorithm is used to coherently estimate the distance from \mathbf{u} to the centroid of a cluster. The proof follows trivially from the above results and is given below.

Proof of Theorem 2. Proof follows as a trivial consequence of taking $M = M'$ in Corollary 8, applying Lemma 11 and observing that dividing the calculated distance by σ_m is efficient irrespective of whether $M' > 1$ or $M' = 1$. Note that the upper bound is not tight in cases where $M' = 1$ because σ_m does not need to be computed in such cases; nonetheless, removing this cost only reduces the expected query complexity by a constant factor so we do not change the theorem statement for simplicity. Also, using a non-constant value for $M_1, \dots, M_{M'}$ does not change the problem since the state preparation method of Lemma 4 takes M as an input state that can be set to $M_1, \dots, M_{M'}$ coherently. \square

Appendix C: Justification for Normalizing Distance

An important question remains: when is the normalized distance between \mathbf{u} and the cluster centroid a useful statistic in a machine learning task? Of course, as mentioned earlier, this statistic is not always optimal. In cases where the training data points live on a complicated manifold, there will be many points that are close to the cluster centroid yet are not in the cluster. Even in such circumstances, the normalized distance leads to an upper bound on the probability that \mathbf{u} is in the cluster.

For concreteness, let us assume that

$$\begin{aligned} |\mathbf{u} - \text{mean}(\{A\})| &= \xi_A, \\ |\mathbf{u} - \text{mean}(\{B\})| &= \xi_B. \end{aligned} \quad (\text{C1})$$

If we then define the intra-cluster variances to be σ_A^2 and σ_B^2 for clusters $\{A\}$ and $\{B\}$, then Chebyshev's inequality states that regardless of the underlying distributions of the clusters that for any point x

$$\begin{aligned} \Pr(|x - \text{mean}(\{A\})| \geq \xi_A | x \in \{A\}) &\leq \frac{\sigma_A^2}{\xi_A^2}, \\ \Pr(|x - \text{mean}(\{B\})| \geq \xi_B | x \in \{B\}) &\leq \frac{\sigma_B^2}{\xi_B^2}. \end{aligned} \quad (\text{C2})$$

Eq. (C2) tells us that if the normalized distance is large then the probability that the point is in the corresponding cluster is small.

Unfortunately, (C2) does not necessarily provide us with enough information to merit use in a decision problem because there is no guarantee that the inequalities are tight. If Chebyshev's inequality is tight then basing a decision on the normalized distance is equivalent to the likelihood-ratio test, which is widely used in hypothesis testing.

Theorem 13. Assume there exist positive numbers a, b, α, β such that for all $\chi \geq \min\{\xi_A, \xi_B\}$

$$\begin{aligned} a \frac{\sigma_A^2}{\chi^2} &\leq \Pr(|x - \text{mean}(\{A\})| \geq \chi | x \in \{A\}) \leq \alpha \frac{\sigma_A^2}{\chi^2} \\ b \frac{\sigma_B^2}{\chi^2} &\leq \Pr(|x - \text{mean}(\{B\})| \geq \chi | x \in \{B\}) \leq \beta \frac{\sigma_B^2}{\chi^2}, \end{aligned}$$

and either $a \geq \beta$ or $\alpha \leq b$, then using the normalized distance to the cluster centroid to decide whether $\mathbf{u} \in \{A\}$ or $\mathbf{u} \in \{B\}$ is equivalent to using the likelihood ratio test.

Proof. The likelihood ratio test concludes that \mathbf{u} should be assigned to A if

$$\frac{\Pr(|\mathbf{u} - \text{mean}(\{A\})| \geq \chi | x \in \{A\})}{\Pr(|\mathbf{u} - \text{mean}(\{B\})| \geq \chi | x \in \{B\})} > 1. \quad (\text{C3})$$

Our assumptions show that (C3) is implied if

$$\frac{a \frac{\sigma_A^2}{\xi_A^2}}{\beta \frac{\sigma_B^2}{\xi_B^2}} > 1 \Rightarrow \frac{a}{\beta} > \left(\left(\frac{\sigma_B}{\xi_B} \right) \left(\frac{\sigma_A}{\xi_A} \right)^{-1} \right)^2. \quad (\text{C4})$$

If the normalized distance is used as the classification decision, then \mathbf{u} is assigned to $\{A\}$ if $\left(\left(\frac{\sigma_B}{\xi_B} \right) \left(\frac{\sigma_A}{\xi_A} \right)^{-1} \right) \leq 1$.

Therefore the two tests make the same assignment if $a \geq \beta$.

The likelihood ratio test similarly assigns \mathbf{u} to $\{B\}$ if

$$\frac{\Pr(|\mathbf{u} - \text{mean}(\{A\})| \geq \chi | x \in \{A\})}{\Pr(|\mathbf{u} - \text{mean}(\{B\})| \geq \chi | x \in \{B\})} < 1, \quad (\text{C5})$$

which, similar to (C4) is implied by $\frac{a}{b} < \left(\left(\frac{\sigma_B}{\xi_B} \right) \left(\frac{\sigma_A}{\xi_A} \right)^{-1} \right)^2$ and is further equivalent to the distance-based assignment if $\alpha \leq b$. \square

Theorem 13 shows that the validity of distance-based assignment depends strongly on the tightness of Chebyshev's bound; however, it is not necessarily clear a priori whether lower bounds on $\Pr(|\mathbf{u} - \text{mean}(\{A\})| \geq \chi | x \in \{A\})$ and $\Pr(|\mathbf{u} - \text{mean}(\{B\})| \geq \chi | x \in \{B\})$ exist for values of a and b that are non-zero. Such bounds clearly exist if, for example, $\{A\}$ and $\{B\}$ are both drawn from Gaussian distributions. This follows because $\Pr(|\mathbf{u} - \text{mean}(\{B\})| \geq \chi | x \in \{B\})$ can be upper- and lower-bounded by a function of the distance and appropriate values of a and α can be extracted from the covariance matrix. Since both distributions are the same in this case, $a = b$, and $\alpha = \beta$, the normalized distance is well motivated if $\{A\}$ and $\{B\}$ are drawn from two Gaussian distributions whose centroids are sufficiently distant from each other.

Appendix D: Comparison to Monte-Carlo Approaches

It is often sufficient in practice to classify a test point based on a randomly chosen subset of training vectors (and perhaps also features) rather than the complete training set. These approaches are called Monte-Carlo algorithms and they are very powerful in cases where the training data is tightly clustered in high-dimensional spaces. The nearest-centroid classification algorithms are also useful in this regime, so it is natural to compare the cost of performing centroid-based classification using Monte-Carlo sampling to the cost of our nearest-centroid algorithm.

A Monte-Carlo approximation to the inner product of two d -sparse vectors a and b can be found via the following approach. First N_c samples of individual components of a and b are taken. If we assume that the locations where a and b are mutually non-zero are not known a priori then we can imagine that each vector is of dimension $D = \max(N, 2d)$. Let us denote the sequence of indexes to be i_t . Then each component of the D -dimensional vector should be drawn with uniform probability (i.e., $p(i_t = x) = 1/D$ for all x in the union of the set of vectors that support a and b). Then an unbiased estimator of the inner product is given by

$$X = \frac{D}{N_c} \sum_{t=1}^{N_c} a_{i_t} b_{i_t}. \quad (\text{D1})$$

In particular, it is shown in [29] that

$$\mathbb{E}[X] = a^T b, \quad \mathbb{V}[X] = \frac{1}{N_c} \left(D \sum_{i=1}^N a_{i_t}^2 b_i^2 - (a^T b)^2 \right) \in O \left(\frac{d^2 r_{\max}^4}{N_c} \right). \quad (\text{D2})$$

Chebyshev’s inequality therefore implies that for fixed vectors a and b that $N_c \in O(d^2 r_{\max}^4 \epsilon^{-2})$ is sufficient to guarantee that X is a correct estimate to within distance ϵ with high probability. Also, for random unit vectors, $D \sum_{i=1}^N a_{i_t}^2 b_i^2 - (a^T b)^2 = O(1)$ with high probability so typically the cost of the estimate will simply be $O(1/\epsilon^2)$. The cost of nearest neighbor classification is then $O(M d^2 r_{\max}^4 / \epsilon^2)$, which is (up to logarithmic factors) quadratically worse than our nearest-neighbor algorithm for cases where $d r_{\max}^2 \in O(1)$.

A similar calculation implies that we can estimate the components of the mean vector to within error ϵ/N_c (which guarantees that the overall error is at most ϵ). To estimate this, we need to have the variance of each component of the vector. Let $\{\mathbf{v}^{(m)}\}$ be a set of d -sparse unit vectors then

$$\mathbb{V}_m[\mathbf{v}_k^{(m)}] = \frac{1}{M} \sum_{m=1}^M (\mathbf{v}_k^{(m)} - \mathbb{E}_m[\mathbf{v}_k^{(m)}])^2 \leq 4r_{\max}^2. \quad (\text{D3})$$

Thus if we wish to estimate $\mathbb{E}_m[\mathbf{v}_j^{(m)}]$ within error ϵ/N_c (with high probability) then it suffices to take a number of samples for each vector component (i.e., each i_t) that obeys

$$N_s \in O \left(\frac{r_{\max}^2 N_c^2}{\epsilon^2} \right). \quad (\text{D4})$$

Since there are N_c different components, the total cost is $N_c N_s$ which implies that

$$\text{Cost} \in O \left(\frac{r_{\max}^2 N_c^3}{\epsilon^2} \right) \in O \left(\frac{d^6 r_{\max}^{14}}{\epsilon^8} \right). \quad (\text{D5})$$

Since $|a - b|_2 = a^T a + b^T b - 2a^T b$, it follows that the Euclidean distance to the centroid can be computed using a number of queries that scales as (D5).

The Euclidean distance between a test point and the centroid of a cluster can therefore be efficiently estimated, for fixed ϵ , using a classical sampling algorithm. The costs of doing so may be substantially worse than our quantum algorithm if non-constant ϵ is required, N is large and $d^{-3/7} \in o(r_{\max})$. We will see in the following that $\epsilon \in O(1/\sqrt{N})$ will typically be needed to make an accurate assignment in high-dimensional vector spaces, which implies that our quantum algorithms may typically offer substantial speedups over even classical Monte-Carlo algorithms in tasks where many features are present.

Appendix E: Sensitivity of Decision Problem

Although our quantum algorithms for computing the inner product and Euclidean distance provide better scaling with N and M (the dimension of the vectors and the number of vectors in the training set) than their classical analogs, the quantum algorithms introduce a $O(1/\epsilon)$ scaling with the noise tolerance (where ϵ is the error tolerance in the distance computation). If the typical distances in the assignment set shrink as $1/N^\gamma$ for positive integer γ , then it is possible that the savings provided by using a quantum computer could be negated because $\epsilon^{-1} \in \Omega(N^\gamma)$ in such cases.

We will now show that in “typical” cases where the vectors are uniformly distributed over the unit sphere that $\epsilon \in \Theta(1/\sqrt{N})$ will suffice with high probability. Since our nearest-neighbor algorithms scale as \sqrt{M}/ϵ (which in the Euclidean case corresponds to $M' = M$) this implies that our algorithms’ cost scales as $O(\sqrt{NM})$. This scales quadratically better than its classical analog or Monte-Carlo sampling (see appendix for more details on Monte-Carlo sampling).

This result is similar to concentration of measure arguments over the hypersphere, which show that almost all unit vectors are concentrated in a band of width $O(1/\sqrt{N})$ about any equator of the hypersphere [19]. The concentrated band of vectors is the origin of the “curse of dimensionality”, which implies that almost all random unit vectors are within a Euclidean distance of $\sqrt{2} - O(1/\sqrt{N})$ of any fixed unit vector. This means that the underlying distribution of distances in nearest-neighbor learning tends to flatten as $N \rightarrow \infty$, implying that very accurate estimates of the distances may be needed in high-dimensional spaces.

For Haar-random vectors, it can be shown that the probability distribution for the magnitude of each component of the vector is (to within negligible error in the limit of large N) independently distributed and has a probability density of [30]

$$p(|[\mathbf{v}_j]_k| = r) = 2(N-1)r(1-r^2)^{N-2}, \quad (\text{E1})$$

which after a substitution gives

$$p(|\sqrt{N}[\mathbf{v}_j]_k| = u) = \frac{2(N-1)u}{\sqrt{N}} \left(1 - \frac{u^2}{N}\right)^{N-2} \sim 2u\sqrt{N}e^{-u^2} = 2rNe^{-Nr^2}. \quad (\text{E2})$$

Chebyshev's inequality then can be used to show that with high probability $r \in \Theta(1/\sqrt{N})$ for each component, and (E2) shows that the distribution varies smoothly and hence it is highly probable that the differences between any two components of such random vectors are $\Theta(1/\sqrt{N})$.

Since the Haar measure is invariant under unitary transformation, we can take \mathbf{u} to be a basis vector without loss of generality. Then if we define \mathbf{w} and \mathbf{z} to be the two closest vectors we see that with high probability in the limit as $N \rightarrow \infty$

$$\begin{aligned} |\mathbf{u} - \mathbf{w}|^2 &= (1 - \mathbf{w}_1)^2 + \sum_{j=2}^N \mathbf{w}_j^2 \\ &= (1 - \mathbf{w}_1)^2 + (1 - |\mathbf{w}_1|^2) \\ &= 1 + (1 - \mathbf{w}_1)^2 + O(1/N) \\ &= 2 + \Theta(1/\sqrt{N}), \end{aligned} \quad (\text{E3})$$

where the last line follows from the observation that with high probability $|\mathbf{w}_1| \in \Theta(1/\sqrt{N})$. By repeating the same argument we see that

$$|\mathbf{u} - \mathbf{z}|^2 = 2 + \Theta(1/\sqrt{N}), \quad (\text{E4})$$

and hence from the fact that the distribution of distances is smooth and the components of \mathbf{w} and \mathbf{z} are independent we see that

$$|\mathbf{u} - \mathbf{z}|^2 - |\mathbf{u} - \mathbf{w}|^2 \in \Theta(1/\sqrt{N}). \quad (\text{E5})$$

This suggests that $\epsilon \in O(1/\sqrt{N})$ for the case where the members of the training set are Haar-random vectors. We demonstrate this point numerically below.




FIG. 12: Mean difference in Euclidean distance between two closest vectors to \mathbf{u} for vectors randomly chosen according to the Haar measure. The plot is computed for $M = 100$ using 100 trials per value of N . The blue crosses show the mean values of the distances and the dashed lines give a 50% confidence interval on the distance. The green line gives the the best powerlaw fit to the data.

We generate the data in [Figure 12](#) by generating a large set of random vectors chosen uniformly with respect to the Haar measure. We take $|u\rangle = |0\rangle$ in all these examples without loss of generality because the measure is rotationally invariant. We then compute for each $j = 1, \dots, M$ $|\mathbf{u} - \mathbf{w}_j|_2$ and sort the distances between \mathbf{u} and each of the randomly drawn vectors. Finally, we compute the distance gap, or the difference between the two smallest distances, and repeat this 100 times in order to get reliable statistics about these differences.

[Figure 12](#) shows that the difference between these two distances tends to be on the order of $1/\sqrt{N}$ as anticipated from concentration of measure arguments. It is easy to see from Taylor’s theorem that the differences in the square distances is also $O(1/\sqrt{N})$. Hence taking $\epsilon \in O(1/\sqrt{N})$ will suffice with high-probability since an error of this scale or smaller will not be sufficient to affect the decision about the identity of the nearest vector.

In contrast, the scaling with M is much less interesting because the volume expands exponentially with N . Hence it takes a very large number of points to densely cover a hypersphere. For this reason, we focus on the scaling with N rather than M . However, for problems with small N , large M , and no discernable boundaries between U and V , this issue could potentially be problematic.

We can now estimate the regime in which our quantum algorithms will have a definite advantage over a brute-force classical computation. We assume that $\epsilon = 1/\sqrt{N}$, $\delta_0 = 0.5$, $dr_{\max}^2 = 1$ and $M' = M$. We then numerically compute the points where the upper bounds on the query complexity in [Theorem 1](#) and [Theorem 2](#) equal the cost of a brute-force classical computation. We use these points to estimate the regime where our quantum algorithms be cost-advantageous over classical brute-force classification. As seen in [Figure 13](#), our quantum algorithms exhibit superior time complexities for a large range of M and N values. This trade-off point for the centroid method occurs when $M \approx 10^{16} N^{-1.07}$, and $M \approx 2 \times 10^{14} N^{-1.08}$ for the inner-product method.

It is important to note that the upper bounds on the query complexity are not expected to be tight, which means that we cannot say with confidence that our quantum algorithms will not be beneficial if the upper bounds are less than NM . Tighter bounds on the query complexity of the algorithm may be needed in order to give a better estimate of the performance of our algorithm in typical applications.




FIG. 13: Estimated regions where our quantum algorithms are cost-advantageous over a brute-force classical calculation. The shaded regions represent the parameter space where the upper bounds from [Theorem 1](#) and [Theorem 2](#) are greater than the brute-force classical cost of NM .