

Quantum diffusion map for nonlinear dimensionality reduction

Apimuk Sornsang , Ninnat Dangniam, Pantita Palittapongarnpim *, and Thiparat Chotibut [†]

Chula Intelligent and Complex Systems, Department of Physics, Faculty of Science, Chulalongkorn University, Bangkok 10330, Thailand



(Received 25 June 2021; accepted 24 September 2021; published 10 November 2021)

Inspired by random walks on graphs, the diffusion map (DM) is a class of unsupervised machine learning that offers automatic identification of low-dimensional data structure hidden in a high-dimensional data set. In recent years, among its many applications, the DM has been successfully applied to discover relevant order parameters in many-body systems, enabling automatic classification of quantum phases of matter. However, a classical DM algorithm is computationally prohibitive for a large data set, and any reduction of the time complexity would be desirable. With a quantum computational speedup in mind, we propose a quantum algorithm for the DM, termed the *quantum diffusion map* (qDM). Our qDM takes as an input N classical data vectors, performs an eigendecomposition of the Markov transition matrix in time $O(\log^3 N)$, and classically constructs the diffusion map via the readout (tomography) of the eigenvectors, giving a total expected runtime proportional to $N^2 \text{polylog } N$. Finally, quantum subroutines in the qDM for constructing a Markov transition matrix and for analyzing its spectral properties can also be useful for other random-walk-based algorithms.

DOI: [10.1103/PhysRevA.104.052410](https://doi.org/10.1103/PhysRevA.104.052410)

I. INTRODUCTION

Discovering statistical structure in high-dimensional data is essential for data-driven science and engineering. Advances in unsupervised machine learning offer a plethora of alternatives to automatically search for low-dimensional structure of data lying in a high-dimensional space. Many of these approaches involve dimensionality reduction. A classic approach is the principal component analysis (PCA), which projects high-dimensional data onto a low-dimensional linear space spanned by a set of orthonormal bases, whose directions capture significant data variations. A more compelling approach enables automatic searches for a low-dimensional data manifold embedded in a high-dimensional space. Well-known manifold learning algorithms include Isomap [1], Laplacian eigenmaps [2], uniform manifold approximation and projection (UMAP) [3], nonlinear PCA [4], t -distributed stochastic neighbor embedding (t-SNE) [5], and the diffusion map (DM) [6–8], which is the focus of this paper. Inspired by the appealing features of random walks on graphs, the DM and its variants have received increasing attention for data visualization in bioinformatics [9,10]. The DM has also recently been successfully applied to provide automatic classification of topological phases of matter, and offers automatic identification of quantum phase transitions in many-body systems [11–16].

Most dimensionality reduction methods require the computation of singular value decomposition (SVD) of a matrix constructed from a collection of high-dimensional data points. For a matrix of size, say, $N \times d$, where N is the number of data points and d is the dimensionality of each data vector

(assumed to be smaller than N), the computational cost of SVD typically grows with the number of data points as $O(N^3)$. Thus, classical dimensionality reduction can be computationally prohibitive for a large data sample. However, under moderate assumptions of accessibility to certain features of full-scale quantum computers, matrix exponentiation-based quantum algorithms have been proposed to perform SVD more efficiently [17]. In particular, assuming efficient encoding of classical data to quantum information and accessibility to appropriate quantum random access memory (qRAM), the quantum singular value decomposition (qSVD) algorithm's runtime for many nonsparse low-rank Hermitian matrices is $O(\text{polylog } N)$, which is exponentially faster than the classical counterpart and can be extended to nonsquare matrices. A classic quantum algorithm for dimensionality reduction with quantum computational speedup is quantum principal component analysis (qPCA), which exploits matrix exponentiation tricks as in qSVD [18]. More recently, quantum algorithms with quantum advantage for nonlinear dimensionality reduction with nonlinear kernel [19] and cluster identification based on spectral graph theory have also been proposed [20–22].

With a quantum computational speedup for dimensionality reduction in mind, we propose a quantum algorithm for unsupervised manifold learning called a quantum diffusion map (qDM). Under mild assumptions of appropriate oracles, the qDM has an expected runtime of roughly $\kappa_D^{0.625} N^2 \text{polylog } N$, where N is the number of data points and κ_D is the condition number of the degree matrix (see Theorem 1 and Lemma 1 for the precise statement), as opposed to $O(N^3)$ in the classical DM. Although κ_D depends strongly on the data structure and can take the value N in the worst case, such a worst case is highly atypical for a well-structured data set, in which $\kappa_D = O(1)$. Without the final readout step, our qDM algorithm prepares all necessary components for constructing the diffusion map in $O(\kappa_D^{0.625} \log^3 N)$ time.

*Corresponding author: panpalitta@gmail.com

[†]Corresponding author: thiparatc@gmail.com,
thiparat.c@chula.ac.th

Although the backbone of the DM is a graph-based dimensionality reduction method, the procedure is different from other spectral graph methods. Namely, rather than working with the data-induced graph Laplacian as in Laplacian eigenmaps or in spectral clustering, the DM involves a Markov transition matrix that defines random walks on a data-induced graph. Therefore, recipes for qDM dimensionality reduction are different from the recently proposed quantum spectral clustering [20,21], mainly in the efficient construction of the degree matrix, which was mentioned as not efficiently accessible [21]. Our algorithm also exploits the coherent state encoding scheme, initially proposed in Ref. [23], to provide an exact quantum encoding of a classical Gaussian kernel. This is in contrast to the quantum kernel principle component analysis (qKPCA) [19] for dimensionality reduction, which only encodes the nonlinear kernel approximately.

The paper is organized as follows. Section II provides necessary ingredients to understand the classical diffusion map and its runtime. Section III introduces quantum subroutines for the qDM, which is then discussed in detail in Secs. IV and V. The qDM runtime complexity appears in Sec. VI. We conclude with the discussion and outlook in Sec. VII. Additional details on the algorithmic complexity and the subroutines for the qDM are provided in Appendixes A and B.

II. CLASSICAL DIFFUSION MAP FOR DIMENSIONALITY REDUCTION

Given a set of N data vectors $X \equiv \{\mathbf{x}^{(i)}\}_{i=0}^{N-1}$, in which each data vector has d features (i.e., $\mathbf{x}^{(i)} \in \mathbb{R}^d$), unsupervised machine learning seeks to identify the structure of data distribution in a high-dimensional ($d \gg 1$) space. When the data distribution has a highly nonlinear structure, a classic linear approach such as PCA fails. Rather than using mere Euclidean distance as a similarity measure between data points, the manifold learning approach assigns the connectivity among data points in their neighborhood as a proxy for proximity between points; e.g., points on a toroidal helix embedded in three dimensions with equal Euclidean distance can have different geodesic distances on a manifold (different similarity; see Fig. 1, top). This approach allows one to perform a nonlinear dimensionality reduction by assigning an appropriate function that maps a high-dimensional data point into a relevant lower-dimensional Euclidean space embedding, while encapsulating the notion of proximity in high dimensions based on neighborhood connectivity.

Neighborhood connectivity leads to the development of a graph-based manifold learning approach, where one can assign a vertex i to a data vector $\mathbf{x}^{(i)}$ and assign an edge between a pair of data points that are considered to be neighbors. Isomap [1] and Laplacian eigenmaps [2] are among the first graph-based manifold learning algorithms, where relevant low-dimensional data embedding can be extracted from eigendecomposition of a data-induced graph Laplacian matrix.

A. Diffusion map in a nutshell

In the diffusion map, the similarity matrix between a pair of data vectors, or equivalently the weighted edge between a

pair of graph vertices, is often taken to be a Gaussian kernel:

$$W_{ij} \equiv K_\sigma(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \exp\left(-\frac{\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|_2^2}{2\sigma}\right), \quad (1)$$

where the adjustable parameter σ , called the *bandwidth*, sets the scale of the neighborhood. Two points whose squared distance exceeds this bandwidth contribute exponentially little to the weighted edge, suggesting that these two points are far away from being a neighbor in its original feature space. Given a graph with weighted edges W_{ij} , the DM assigns a discrete-time random walk on a data-induced graph, where the Markov transition probability from vertex i to j is given by the normalized weighted edge,

$$P_{ij} = \frac{W_{ij}}{\sum_{j=0}^{N-1} W_{ij}}. \quad (2)$$

One can compactly compute the Markov transition matrix P from the (weighted) *adjacency matrix* W , and the *degree matrix* $D \equiv \text{diag}\{d_i\}_{i=0}^{N-1}$ where $d_i \equiv \sum_{j=0}^{N-1} W_{ij}$ by

$$P = D^{-1}W. \quad (3)$$

The notion of proximity based on graph connectivity can be characterized by how fast random walkers on a graph starting at different data points visit each other. One expects that two points that are connected by multiple paths should be near, whereas two points that are sparsely connected should lie far from each other. In fact, it can be shown that [8] the proper distance function, called the *diffusion distance*, between two points $\mathbf{x}^{(i)}$ and $\mathbf{x}^{(j)}$ after a t -step random walk is given by

$$\text{Dist}_t^2(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \sum_{k=0}^{N-1} \frac{[(P^t)_{ik} - (P^t)_{jk}]^2}{u_0(\mathbf{x}^{(k)})}, \quad (4)$$

where u_0 is the *left* eigenvector with eigenvalue $\lambda_0 = 1$ of P and $u_0(\mathbf{x}^{(k)})$ denotes the k th coordinate of u_0 .

Before obtaining a relevant lower-dimensional Euclidean space representation of the original data point $\mathbf{x}^{(i)}$, we first note the following important identity. Define the *diffusion map* with m bases as

$$\phi_{t,m}(\mathbf{x}^{(i)}) = \begin{pmatrix} \lambda_1^t v_1(\mathbf{x}^{(i)}) \\ \vdots \\ \lambda_m^t v_m(\mathbf{x}^{(i)}) \end{pmatrix}, \quad (5)$$

where v_i is the *right* eigenvector of P with eigenvalue λ_i , $v_i(\mathbf{x}^{(j)})$ denotes the j th coordinate of v_i , and the eigenvalues are ordered such that $\lambda_0 = 1 > \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{N-1} \geq 0$ assuming no degeneracy in the largest eigenvalue (which is always possible given σ is sufficiently large). Note that $\lambda_0 = 1$ is always the eigenvalue of a Markov transition matrix P with a normalized eigenvector $(1 \ \dots \ 1)^T/N$, corresponding to the stationary uniform distribution over all vertices. Then, the diffusion distance (4) can be *exactly* computed from the Euclidean representation in Eq. (5) with $m = N - 1$ bases [8]:

$$\text{Dist}_t^2(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) = \|\phi_{t,N-1}(\mathbf{x}^{(i)}) - \phi_{t,N-1}(\mathbf{x}^{(j)})\|_2^2. \quad (6)$$

The above equality states that the distance in the diffusion space (based on graph connectivity) is identical to the Euclidean embedding distance (induced by the diffusion map). Thus, the notion of proximity between data points from

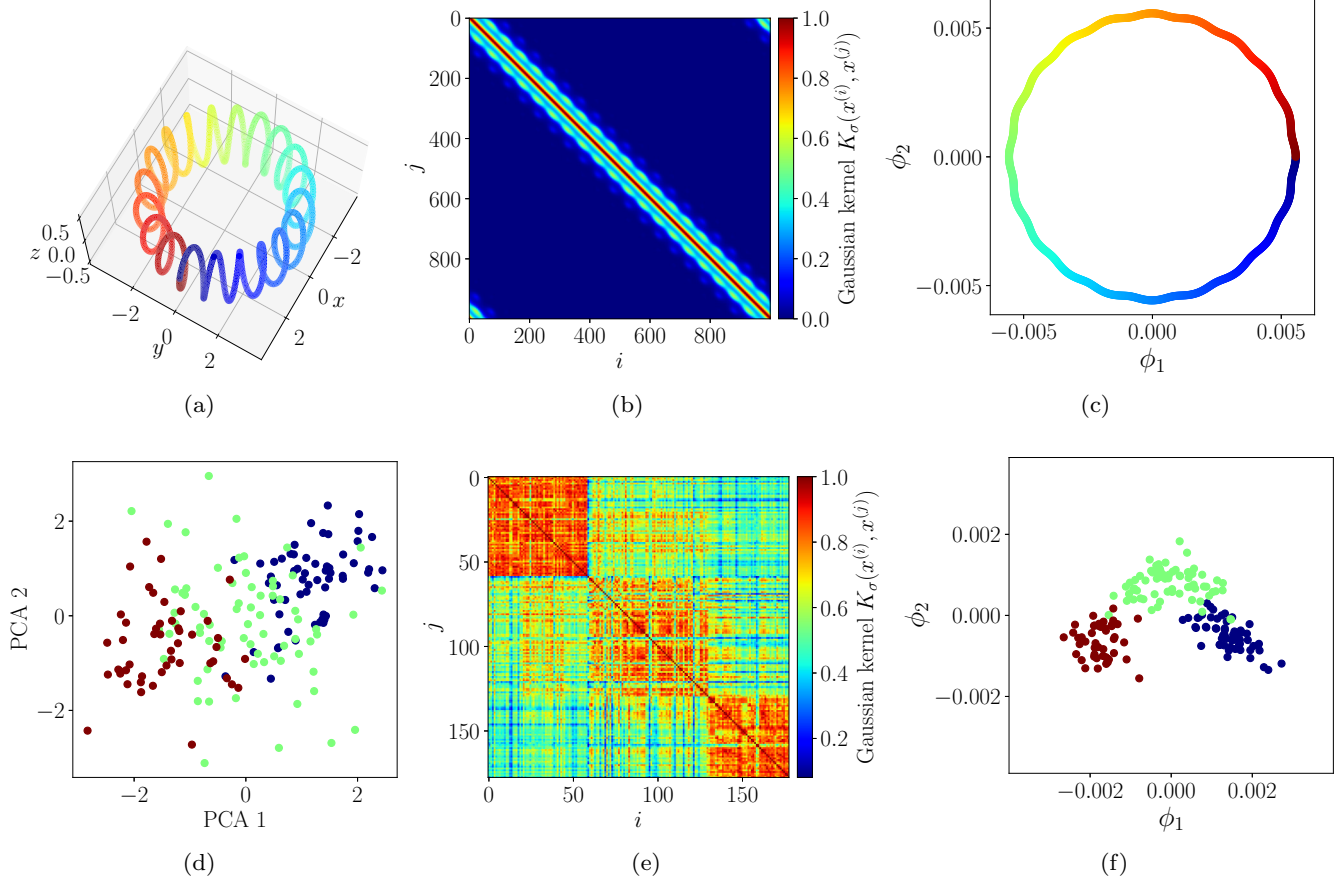


FIG. 1. The application of a classical DM for nonlinear dimensionality reduction (a)–(c) to identify the proximity structure of data points distributed as a toroidal helix in three dimensions, and (d)–(f) to reveal the proximity structure of 13 chemical components’ concentrations of 178 wines ($d = 13$, $N = 178$) derived from three different cultivars grown in the same region in Italy, taken from the WINE data set in Ref. [24]. (a) In a toroidal helix data set, color variation signifies a difference in the geodesic distance between points on a one-dimensional manifold. (b) The Gaussian kernel with $\sigma = 1$ of Eq. (1) shows the neighborhood connectivity has a local, periodic structure. (c) The dimensionality reduction into the first two components ($m = 2$) of the DM in Eq. (5) reveals a one-dimensional structure with an appropriate notion of geodesic distance (see color variations). In the WINE data set, wines grown from three different cultivars are labeled in three different colors (red, green, blue). (d) Projecting the 13 data features (chemical concentrations) of all wines into the first two principal components using PCA does not suffice to distinguish between the three types. However, (e) with the Gaussian kernel with $\sigma = 50$, the DM with two bases ($m = 2$) can reveal quite a clear distinction between types (f). In addition, the proximity structure also suggests that the green type has its chemical components closer to those of the other two types. In both data sets, we use $t = 1$ in the DM. Thus, the random walk on data-induced graphs can reveal salient low-dimensional data structure within just a one-step walk, provided σ is chosen appropriately.

its graph connectivity can be simply computed from the Euclidean embedding by the diffusion map (5).

What about dimensionality reduction? Since P is a Markov transition matrix, its eigenvalues are $\lambda_0 = 1 \geq \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{N-1} \geq 0$. Bases in Eq. (5) with low-lying eigenvalues are then exponentially suppressed as t step increases. In the long-time limit, dimensionality reduction thus naturally arises in the DM. One may take the number of bases m in Eq. (5) corresponding to the number of top eigenvalues of P and still obtain meaningful low-dimensional Euclidean representation of points in the diffusion space. In practice, for the purpose of visualization, taking $m = 2 \ll d$ is a drastic dimensionality reduction, yet a DM with such a small number of bases can yield insights into the approximate proximity of data in the original high-dimensional space (see Fig. 1). Using DM (5) to extract low-dimensional data representation that approximately preserves the notion of distance from neigh-

borhood connectivity in the original high-dimensional space is also the first step towards data clustering algorithms. Namely, one may employ standard clustering algorithms, such as a k -means algorithm, on the low-dimensional outputs of the DM without suffering from the curse-of-dimensionality problem.

B. Time complexity of classical diffusion map

Numerical recipes to obtain nonlinear dimensionality reduction or manifold learning in the classical DM consist of four steps:

- (1) Construct the data-induced similarity matrix (weighted adjacency matrix) W from Eq. (1).
- (2) Construct the Markov transition matrix P from Eq. (3); i.e., $P = D^{-1}W$.
- (3) Compute the eigendecomposition of P .
- (4) Construct the diffusion map (5).

We now discuss the time complexity for each step as a function of the number of data points, N . Assuming the Gaussian kernel function (1), the weighted adjacency matrix W is symmetric. The number of the elements that the algorithm needs to calculate is then $\sum_{i=1}^N i = (N^2 + N)/2$. Therefore, the time complexity of calculating the kernel matrix in step 1 is $O(N^2)$.

For step 2, the calculation of P involves normalizing each row of W as in Eq. (2). Numerical computation of the normalization factor, which is simply the summation, takes the time $O(N)$. Then normalizing each row costs another $O(N)$. Applying these procedures over the $N \times N$ matrix thus yields the time cost $O(N^2) + O(N^2)$. Hence, the time complexity for computing P is $O(N^2)$.

In step 3, the time cost stems from finding $\{v_i\}$ and $\{\lambda_i\}$ of P , using SVD, which can be applied to any nonsquare matrices. This algorithm comprises two steps [25]. The first step is to use Householder transformations to reduce P to a bidiagonal form. Then the QR algorithm is applied to find singular values. The time complexity of these two steps combined is $O(N^3)$.

The last step is computing the diffusion map ϕ according to Eq. (5), which involves the multiplication of λ_i^t with its corresponding right eigenvector v_i . Although we might not need to perform the multiplication for all $N - 1$ right eigenvectors as some λ_i^t maybe negligible in the long-time limit, we consider the worst-case scenario where all $N - 1$ eigenvectors are required. With the length of each vector being N , the time cost is $O(N^2 - N)$ for computing ϕ and thus the time complexity for this step is $O(N^2)$.

Combining all the steps above, which are implemented in sequence, the classical DM has a time complexity of $O(N^2) + O(N^2) + O(N^3) + O(N^2)$ leading to an overall time complexity that scales as $O(N^3)$. Note here that the time complexity is primarily dominated by the eigendecomposition algorithm, which is subject to change if a different algorithm is chosen in the implementation. In fact, this dominating complexity of $O(N^3)$ from eigendecomposition urges for the development of quantum algorithms to achieve a quantum computational speedup, which we discuss next.

III. QUANTUM DIFFUSION MAP ROADMAP

In this section, we first give an outline of the qDM algorithm summarized in Fig. 2. Then we proceed with classical to quantum data encoding, and the construction of the kernel matrix in Secs. III A and III B, respectively. Constructing the Markov transition matrix and extracting the relevant eigendecomposition are not straightforward, and will be discussed in Secs. IV and V, respectively.

From now on, we will focus on the case where the weighted adjacency matrix is the Gaussian kernel matrix, so that $W = K$. To reduce the dominating complexity in the construction of the diffusion map, we first turn to the idea of quantum-mechanically encoding the (normalized) kernel matrix as a density matrix \hat{K} without explicitly evaluating each matrix element. Using the fact that the Gaussian kernel (1) is the inner product of canonical coherent states, this can be done in $O(\log N)$ time assuming qRAM by exploiting “quantum

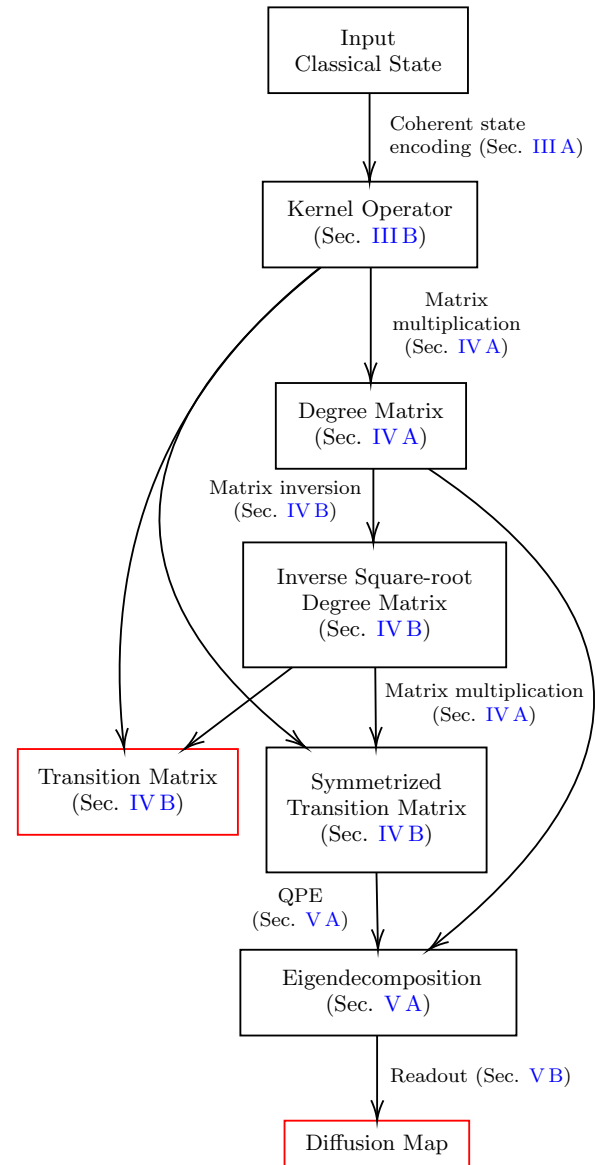


FIG. 2. The roadmap for the quantum diffusion map algorithm, which also enables an efficient construction of the transition matrix for generating random walks on graphs as a byproduct. Note that we do not perform the eigendecomposition of the transition matrix directly since quantum phase estimation (QPE) only eigendecomposes symmetric matrices. Nevertheless, the eigenpairs of the original transition matrix can be reconstructed from those of the symmetrized transition matrix.

parallelism,” the ability to query data in superposition. We further discuss the assumption of qRAM in Appendix B 1.

To construct the transition matrix, our algorithm uses as subroutines density matrix exponentiation [18], the quantum matrix algebra toolbox (QMAT) [26], and quantum matrix inversion [27,28]. While it is straightforward to classically compute the degree matrix D given elements of K , here we do not have direct access to elements of the density matrix \hat{K} without performing full tomography. Instead, we compute D via the identity

$$D = (K\mathbb{1}) \odot I, \quad (7)$$

where \odot is element-wise multiplication, also known as a Hadamard product, $\mathbb{1}$ is the all-ones matrix, and I is the identity matrix. Such matrix arithmetic can be done in the QMAT framework given the exponential of the relevant matrices either by density matrix exponentiation for \hat{K} , or Hamiltonian simulation to exponentiate sparse matrices [29]. We then use quantum phase estimation (QPE) to construct the inverse square root of D (akin to the Harrow-Hassidim-Lloyd (HHL) algorithm [27]). Then the (right) eigenvectors and eigenvalues of the transition matrix $P \equiv D^{-1}K$ can be obtained from the eigendecomposition (via another QPE) of the symmetrized transition matrix $S \equiv D^{-1/2}KD^{-1/2}$. The diffusion map is then constructed classically. The time complexity is given in Sec. VI.

A. Quantum state encoding into coherent states

By the fact that the Gaussian kernel arises naturally as the inner product of coherent states, Ref. [23] proposed encoding classical data into multimode coherent states,

$$\mathbf{x}^{(i)} \mapsto |\mathbf{x}^{(i)}\rangle = \bigotimes_{p=0}^{d-1} |x_p^{(i)}\rangle, \quad (8)$$

where each single-mode coherent state $|x_p^{(i)}\rangle$ represents the p th feature of datum $\mathbf{x}^{(i)}$. In continuous-variable quantum systems, the canonical coherent states are states with minimal and equal uncertainties in both quadratures. They are displaced ground states of the harmonic oscillator, and can be realized, for example, as the state of classical electromagnetic radiation. A single-mode coherent state is defined as

$$|x_p^{(i)}\rangle = \exp\left(-\frac{(x_p^{(i)})^2}{2}\right) \sum_{n=0}^{\infty} \frac{(x_p^{(i)})^n}{\sqrt{n!}} |n\rangle, \quad (9)$$

where $|n\rangle$ is the n th eigenstate of the harmonic oscillator.

The inner product between two data vectors is the Gaussian kernel

$$\begin{aligned} \langle \mathbf{x}^{(i)} | \mathbf{x}^{(j)} \rangle &= \prod_{p=0}^{d-1} \langle x_p^{(i)} | x_p^{(j)} \rangle = \prod_{p=0}^{d-1} \exp\left(-\frac{(x_p^{(i)} - x_p^{(j)})^2}{2}\right) \\ &= \exp\left(-\frac{\|\mathbf{x}^{(i)} - \mathbf{x}^{(j)}\|_2^2}{2}\right) = K(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}). \end{aligned} \quad (10)$$

Here, we take $\sigma = 1$ for simplicity (though σ can be incorporated by a scaling factor during state encoding in classical preprocessing). Recall that we focus on the kernel matrix as the weight matrix, and thus we denote K for describing the weight matrix W .

B. Kernel matrix

Now we exploit superposition access to the encoded data to implicitly evaluate the kernel matrix. This can be done efficiently, assuming qRAM, for instance [30,31]. Calling the oracle with the state $\frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle$ entangles the data vectors with their labels:

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle_{\text{label}} \otimes |\mathbf{x}^{(i)}\rangle_{\text{position}}. \quad (11)$$

The reduced density matrix of the label space is given by the partial trace over the position space:

$$\begin{aligned} \text{Tr}_{\text{position}}(|\psi\rangle\langle\psi|) &= \frac{1}{N} \sum_{i,j=0}^{N-1} \langle \mathbf{x}^{(i)} | \mathbf{x}^{(j)} \rangle |i\rangle\langle j| \\ &= \frac{1}{N} \sum_{i,j=0}^{N-1} K(\mathbf{x}^{(i)}, \mathbf{x}^{(j)}) |i\rangle\langle j| = \frac{K}{N} \equiv \hat{K}, \end{aligned} \quad (12)$$

where $\text{Tr}(\hat{K}) = \sum_{i=0}^{N-1} K(\mathbf{x}^{(i)}, \mathbf{x}^{(i)}) = N$. Copies of \hat{K} will be used to prepare the degree matrix and the transition matrix.

IV. QUANTUM SUBROUTINE FOR CONSTRUCTING THE TRANSITION MATRIX

In this section, we show that the degree matrix and its inverse, and consequently the transition matrix, can be obtained efficiently in a coherent fashion. In particular, if one can efficiently perform matrix multiplication and a Hadamard product \odot involving \hat{K} , then one can efficiently obtain the degree matrix according to Eq. (7). Such matrix arithmetic can be done in the QMAT framework [26] by first exponentiating the relevant (possibly non-Hermitian) matrices A in the form $e^{iX_i(A)}$, where

$$X_i(A) = R_i \otimes A + R_i^\dagger \otimes A^\dagger, \quad (13)$$

and

$$R_1 = |0\rangle\langle 1|, \quad R_2 = |1\rangle\langle 2|, \quad R_3 = |0\rangle\langle 2| \quad (14)$$

are qutrit operators. Once one has the ability to apply, say, $e^{iX_1(A_1)}$ and $e^{iX_2(A_2)}$ (shuffling between R_i , $i = 1, 2, 3$, can be done by a simple permutation of the qutrit basis states), then using the commutation relation between $X_1(A_1)$ and $X_2(A_2)$ one can approximate $e^{iX_3(A_3)}$ where A_3 is the result of the desired binary operation between A_1 and A_2 .

A. Exponential of degree matrix

To compute the degree matrix, we require QMAT subroutines for (i) matrix multiplication, (ii) the tensor product, and (iii) the Hadamard product. The latter two subroutines reduce to matrix multiplication as $A_1 \otimes A_2 = (A_1 \otimes I)(I \otimes A_2)$, and

$$\Xi(A_1 \otimes A_2) \Xi^\dagger = (A_1 \odot A_2) \otimes |0\rangle\langle 0|, \quad (15)$$

where

$$\Xi = \sum_{i=0}^{N-1} |i\rangle\langle i| \otimes |0\rangle\langle i|, \quad (16)$$

and $|0\rangle = |0 \cdots 0\rangle$ [26]. The multiplication subroutine and its complexity are summarized in Appendix A.

Define the (rank-one) density matrix

$$\hat{\mathbb{1}} \equiv \frac{1}{N} \sum_{i,j=0}^{N-1} |i\rangle\langle j| = \frac{\mathbb{1}}{N}. \quad (17)$$

If $e^{iX_i(\hat{K})\tau}$, $e^{iX_i(\hat{\mathbb{1}})\tau}$, $e^{iX_i(\Xi)\tau}$, and $e^{iX_i(\Xi^\dagger)\tau}$ can be constructed efficiently for some time τ , then one can construct for a




FIG. 3. Quantum circuit for square-root degree matrix inversion. Note that QFT denotes quantum Fourier transform and \mathcal{U}_D is defined by Eq. (18).

larger time t (to be determined later for the purpose of phase estimation) the controlled- $e^{iX_3(D)t}$ unitary

$$e^{iX_3(\Xi(\hat{K}\hat{1}\otimes I)\Xi^\dagger)t} = e^{iX_3(D)t} \otimes |\mathbf{0}\rangle\langle\mathbf{0}| + I \otimes (I - |\mathbf{0}\rangle\langle\mathbf{0}|) \equiv \mathcal{U}_D. \quad (18)$$

The last register will be discarded in the inversion step.

Now we discuss an efficient construction of the unitary matrices required in the above procedure. Ξ is sparse and thus can be exponentiated efficiently [32]. For the kernel matrix \hat{K} and $\hat{1}$ which are not sparse, we can use the technique of density matrix exponentiation with time cost of $O(\tau^2/\epsilon)$ with accuracy ϵ [18]. [Note that we cannot exponentiate $X_i(A)$ directly using this method since $X_i(A)$ is not a density matrix. However, since $R_i + R_i^\dagger = |+\rangle\langle+| - |-\rangle\langle-|$, where $|\pm\rangle$ are the ± 1 eigenvectors of $R_i + R_i^\dagger$, we can separately exponentiate $|\pm\rangle\langle\pm| \otimes A$ with the time parameters $\pm\tau$, respectively, and concatenate them to obtain the desired unitary operator.] Hence, $e^{iX_3(\hat{K}\hat{1})\tau}$ and automatically $e^{iX_3((\hat{K}\hat{1})\otimes I)\tau}$ (since $e^{iA} \otimes I = e^{iA\otimes I}$) can be created efficiently via the multiplication subroutine.

B. Inverse square root of degree matrix and transition matrix

Having $e^{iX_3(D)t}$ already in the exponential form, it is natural to construct the inverse square root of D using QPE akin to the HHL algorithm for matrix inversion [27], with the following differences:

- (1) Our degree matrix is always well conditioned, having all eigenvalues in the range $1 \leq d_i < N$.
- (2) To construct the (symmetrized) transition matrix, we would like the resulting inverse square-root matrix in the form of a density matrix to be fed into the QMAT framework once more via density matrix exponentiation. Following Ref. [28], this implies that the function of the eigenvalues we want to compute is $1/\sqrt[4]{d_i}$.

The quantum circuit for the QPE (Fig. 3) is initialized in the state

$$|\psi_0\rangle = |\mathbf{0}\rangle_{\text{val}} \otimes |+\rangle_R \otimes |1\rangle_{\text{vec}} \otimes |\mathbf{0}\rangle_{\text{junk}}, \quad (19)$$

where $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |2\rangle)$ and $|1\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle$. The first register “val” will eventually contain the estimated eigenvalues after the QPE, which are the components of the degree matrix. The next two registers, “R” and “vec,” will contain the eigenvectors of $X_3(D)$, and the last one, “junk,” is added for collapsing other terms except $e^{iX_3(D)t}$ in Eq. (18). Denote by $\kappa_D = d_{\max}/d_{\min}$ the condition number of D , i.e., the ratio of the largest eigenvalue to the smallest eigenvalue. Setting $t = O(\kappa_D^{0.25}/\epsilon_D)$ with accuracy ϵ_D in the QPE results in the final state $\frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |\tilde{d}_i\rangle_{\text{val}} \otimes |+\rangle_R \otimes |i\rangle_{\text{vec}} \otimes |\mathbf{0}\rangle_{\text{junk}}$, where the tilde denotes the binary representation of the estimated degree.

After the conditional rotation on the ancilla qubit, “anc,” the state is

$$\sum_{i=0}^{N-1} \left(\sqrt{1 - \frac{\check{C}}{d_i}} |0\rangle_{\text{anc}} + \sqrt{\frac{\check{C}}{d_i}} |1\rangle_{\text{anc}} \right) \otimes |\tilde{d}_i\rangle_{\text{val}} \otimes |+\rangle_R \otimes |i\rangle_{\text{vec}} \otimes |\mathbf{0}\rangle_{\text{junk}}, \quad (20)$$

where \check{C} is of order $O(d_{\min})$. We then perform postselection on the outcome “1” on the ancilla qubit and take the partial trace over “val,” “R,” and “junk” registers to obtain

$$\hat{D}^{-1/2} = \check{C} \sum_{i=0}^{N-1} \sqrt{\frac{1}{d_i}} |i\rangle\langle i|, \quad (21)$$

where $\check{C} = O(\{\sum_{i=0}^{N-1} d_i^{-1/2}\}^{-1})$ with success probability $\Omega(1/\kappa_D^{0.5})$ [28]. The success probability can be amplified to near certainty by repeating $O(\kappa_D^{0.25})$ rounds of amplitude amplification on the ancilla register [27]. Via another multiplication subroutine, we now have access to the symmetrized transition matrix $S = D^{-1/2}KD^{-1/2}$ in an efficient manner.

Although we will not deal with the bona fide transition matrix P in the rest of this work, it is important to note that $P = D^{-1}K = D^{-1/2}D^{-1/2}K$ can be efficiently constructed via the multiplication subroutine, given that we already have an access to $D^{-1/2}$ and K . In particular, the exponential $e^{iX_3(P)\tau}$ can be constructed in roughly $O(\kappa_D^{0.625} \tau^{1.5} \log^2 N)$ steps [the same as $T_S(\tau)$ in Theorem 1]. Our algorithm thus can result in P as a by-product, which may be useful for other random-walk-based algorithms.

V. EIGENDECOMPOSITION

Recall that the necessary ingredients for constructing the diffusion map are the (right) eigenvectors and the eigenvalues of the transition matrix $P = D^{-1}K$. We thus need to solve the eigenvalue problem

$$D^{-1}K|v_i\rangle = \lambda_i|v_i\rangle. \quad (22)$$

QPE would immediately provide such an eigendecomposition if P is symmetric, which is typically not the case. Note that if one applies QPE to the Hermitized $\begin{pmatrix} 0 & P \\ P^\dagger & 0 \end{pmatrix}$, one would end up with the singular vectors of P , rather than the desired eigenvectors of nonsymmetric P . Therefore, we are in fact dealing with the generalized or nonsymmetric eigenvalue problem.




FIG. 4. Quantum circuit for transition matrix P eigendecomposition. \mathcal{U}_D is defined by Eq. (18).

One resolution is to transform Eq. (22) to a standard symmetric eigenvalue problem

$$D^{-1/2}KD^{-1/2}|s_i\rangle = \lambda_i|s_i\rangle, \quad (23)$$

where $|s_i\rangle \equiv D^{1/2}|v_i\rangle$. Then, the desired eigenvectors $|v_i\rangle$ can be recovered given an efficient rotation conditional on $D^{-1/2}$. Now QPE can be applied, the output of which will be an estimate of the eigenvalues and the eigenvectors of

$$X_3(S) = (R_3 + R_3^\dagger) \otimes S, \quad (24)$$

which are $\{\pm\lambda_i\}_{i=0}^{N-1}$ and $|\pm\rangle \otimes |s_i\rangle$. The rest of this section elaborates on the procedure for an eigendecomposition of S , and on how to extract the eigenpairs of P in order to construct the diffusion map classically.

A. Transition matrix eigendecomposition

The complete quantum circuit for the eigendecomposition of P is shown in Fig. 4. An input state is prepared in registers “val,” “R,” and “vec” as

$$|\mathbf{0}\rangle_{\text{val}} \otimes |+\rangle_R \otimes |\psi_0\rangle_{\text{vec}} = |\mathbf{0}\rangle_{\text{val}} \otimes |+\rangle_R \otimes \sum_{i=0}^{N-1} \beta_i |s_i\rangle_{\text{vec}}, \quad (25)$$

where $|\psi_0\rangle$ is an arbitrary initial state such that $\beta_i \equiv \langle s_i | \psi_0 \rangle$. The output state of the QPE after evolving for time $t = 2\pi/\epsilon_S$, where ϵ_S is the accuracy of the estimated phases, reads

$$|\psi_1\rangle_{\text{val},R,\text{vec}} \equiv \sum_{i=0}^{N-1} \beta_i |\tilde{\lambda}_i\rangle_{\text{val}} \otimes |+\rangle_R \otimes |s_i\rangle_{\text{vec}}, \quad (26)$$

where $\tilde{\lambda}_i$ denotes the binary representation of the estimated eigenvalues in the range $[0,1)$.

The resulting eigenvectors $|s_i\rangle$ are not yet the components of the diffusion map; they need to be conditionally rotated to the right eigenvectors $|v_i\rangle = D^{-1/2}|s_i\rangle$ of P via another QPE

using $e^{iX_3(D)t}$ as the time-evolution operator. (Note that this step is more similar to the original HHL algorithm than our first application of QPE to construct $D^{-1/2}$; cf. Ref. [28]) For this reason, we add to the initial state register “deg,” which will store the estimated eigenvalues \tilde{d}_i of the degree matrix, and register “junk” to eliminate the unwanted terms in Eq. (18):

$$|\mathbf{0}\rangle_{\text{deg}} \otimes |\psi_1\rangle_{\text{val},R,\text{vec}} \otimes |\mathbf{0}\rangle_{\text{junk}}. \quad (27)$$

The output of the QPE after evolving for time $t = O(\kappa_D^{0.5}/\epsilon_D)$ is

$$\sum_{ij=0}^{N-1} |\tilde{d}_j\rangle_{\text{deg}} \otimes \beta_i |\tilde{\lambda}_i\rangle_{\text{val}} \otimes |+\rangle_R \otimes (s_i)_j |j\rangle_{\text{vec}} \otimes |\mathbf{0}\rangle_{\text{junk}}, \quad (28)$$

where $(s_i)_j$ is the j th component of $|s_i\rangle$. After the rotation conditioning on \tilde{d}_i , and $O(\kappa_D^{0.5})$ rounds of amplitude amplification on an ancilla qubit as in Eq. (20) (not shown here), the eigenvector register is now, up to normalization, in the desired eigenstate

$$\sum_{j=0}^{N-1} \frac{(s_i)_j}{\sqrt{d_j}} |j\rangle_{\text{vec}} = D^{-1/2}|s_i\rangle_{\text{vec}} = |v_i\rangle_{\text{vec}}. \quad (29)$$

B. Extracting desired eigenvalues and eigenvectors

Measuring the eigenvalue register “val” gives one of the eigenvalues and collapses the register “vec” to the eigenvector corresponding to that eigenvalue with probability $|\beta_i|^2$ [assuming no degeneracy as in the classical case (Sec. II A)]. The eigenvector can then be extracted using a number of copies that scales almost linearly in N [33]. Repeating the procedure until we recover all $N - 1$ distinct eigenvalues and eigenvectors allows us to classically construct the diffusion map (5).

The process of recovering all eigenpairs is not deterministic; in particular, one would not be able to recover an eigenpair for which β_i is close to zero. What this means in practice is that one may need to change the initialization $|\psi_0\rangle$ of the QPE every so often, perhaps adaptively. We give below the expected runtime of this process in terms of the β_i 's of a fixed fiducial state based on the classical coupon collector's problem [34]. Interestingly, in the best case where the β_i 's are all equal, the quantum coupon collector's algorithm in Ref. [35] guarantees that we will see all eigenpairs at least once in time $O(N)$.

Algorithm 1. Quantum diffusion map.

Input : Classical data X

Output: Eigenvalues $\{\lambda_i\}$ and right eigenvectors $\{v_i\}$ of the Markov transition matrix P

1. Embed X into coherent state qRAM
 2. Perform partial trace to obtain \hat{K}
 3. Create state $\hat{\mathbb{1}}$ and construct $e^{iX_1(\hat{K})t}$ and $e^{iX_2(\hat{\mathbb{1}})t}$
 4. Construct $e^{iX_3(D)t}$ where $t = O(\kappa_D^{0.25}/\epsilon_D)$
 5. Invert square-root of D by
 - (a) Using $e^{iX_3(D)t}$ in QPE initialized using state (19)
 - (b) Rotating the ancilla qubit controlled by the eigenvalue register “val”
 - (c) Amplifying an ancilla qubit “1” and perform partial trace on the “val”, “R”, and “junk” registers to get $\hat{D}^{-1/2}$
 6. Construct $e^{iX_3(\hat{D}^{-1/2})t}$ and $e^{iX_3(S)t}$ where $t = O(1/\epsilon_S)$
 7. Use $e^{iX_3(S)t}$ in QPE initialized using state (25)
 8. Operate $D^{-1/2}$ to obtain the eigen-decomposition of P (29)
 9. Measure the eigenvalue register “val” to get an approximated eigenvalue λ_i
 10. Perform tomography to extract the eigenvector v_i corresponding to the λ_i measured in 9
-

VI. TIME COMPLEXITY OF QUANTUM DIFFUSION MAP

The complexity of the qDM is summarized in Theorem 1 and Lemma 1, the proof of which is provided in Appendix B. In Theorem 1, we provide the time complexity for performing eigendecomposition of the transition matrix as we proposed in Secs. III–V. In Lemma 1, we provide the expected time complexity for reading out all eigenpairs using the classical coupon collector's algorithm.

Theorem 1. Time complexity of transition matrix eigendecomposition. For a data set $X = \{\mathbf{x}^{(i)}\}_{i=0}^{N-1}$ such that each $\mathbf{x}^{(i)} \in \mathbb{R}^d$ and $d < N$, steps 1–8 of Algorithm 1 construct and eigendecompose the transition matrix with a runtime T_E of

$$O\left[T_S\left(\frac{1}{\epsilon_S}\right)\frac{\log N}{\epsilon_S} + T_D\left(\frac{\kappa_D^{0.5}}{\epsilon_D}\right)\frac{\kappa_D^{0.5} \log N}{\epsilon_D}\right], \quad (30)$$

where T_D denotes the time complexity to construct the degree matrix,

$$T_D(t) = O\left[\left(\frac{t^{1.5}}{\epsilon_m^{0.5}} + \frac{t^{1.25}}{\epsilon_e \epsilon_m^{0.75}}\right) \log N\right], \quad (31)$$

T_S denotes the time complexity to construct the symmetrized transition matrix,

$$O\left[\left(\frac{\kappa_D^{0.625}}{\epsilon_D^{2.5} \epsilon_m^{0.5}} + \frac{\kappa_D^{0.5625}}{\epsilon_D^{2.25} \epsilon_e \epsilon_m^{0.75}}\right) \left(\frac{t^{1.5}}{\epsilon_m^{0.5} \epsilon_e} + \frac{t}{\epsilon_e}\right) \log^2 N + \frac{t^{1.5}}{\epsilon_m^{0.5} \epsilon_e} \log N\right], \quad (32)$$

and ϵ_e and ϵ_m denote the errors of density matrix exponentiation and matrix multiplication, respectively.

Expressing Eq. (30), we simplify the time complexity of the transition matrix eigendecomposition to

$$O(\kappa_D^{0.625} \log^3 N + \kappa_D^{1.25} \log^2 N). \quad (33)$$

Lemma 1. The expected runtime to construct the diffusion map given a fixed initialization (25) of the QPE is upper bounded by

$$\frac{\max_i \{|\beta_i|^2\}}{\min_i \{|\beta_i|^2\}} N \log N \times O\left(\frac{N}{\epsilon_t} \log N\right) \times T_E, \quad (34)$$

where $i \in \{0, \dots, N-1\}$, T_E denotes the time complexity of eigendecomposing the transition matrix in Theorem 1, and ϵ_t denotes the error of quantum tomography.

In the special case when all the β_i 's are equal, we can recover the worst-case runtime by virtue of the recent quantum coupon collector's algorithm [35].

Lemma 2. Suppose that the QPE is initialized in the state such that all the β 's are equal in Eq. (25). Using the notations of Lemma 1, Algorithm 1 constructs the quantum diffusion map with a total runtime of

$$O\left(\frac{N^2}{\epsilon_t} \log N \times T_E\right). \quad (35)$$

When the errors and the condition number are independent of the data set size N , the expected runtime reduces to

$$N^2 \text{polylog } N, \quad (36)$$

compared to the classical diffusion map algorithm which requires $O(N^3)$ time in the worst case. In fact, the bottleneck factor of N^2 and the probabilistic nature of the runtime solely arise from the readout (tomography) alone (see Appendix B), which perhaps can be circumvented if one coherently outputs the diffusion map as quantum states, providing an end-to-end computation. Another possibility to bypass the probabilistic runtime is to extend the quantum coupon collector's algorithm in Lemma 2 to encompass a structured β_i and derive the worst-case runtime. However, we leave this open problem for future work. Regarding the condition number of the degree matrix, the condition number κ_D depends on the neighborhood connectivity of the data set, which will be tuned to be $O(1)$ and independent of N , provided the scale of the neighborhood controlled by σ in the kernel matrix is local, such as in Fig. 1 (middle).

VII. DISCUSSION AND CONCLUSION

In this work, we first review the classical diffusion map, an unsupervised learning algorithm for nonlinear dimensionality reduction and manifold learning. We then explain our efforts to construct a quantum algorithm for the diffusion map. Our quantum diffusion map (qDM) consists of five major steps: a coherent-state data encoding scheme, a natural construction of the kernel matrix from coherent states, a scheme to construct the Markov transition matrix from the kernel matrix, the eigendecomposition of the transition matrix, and extracting relevant quantum information to construct the diffusion map classically. The expected time complexity of the qDM is $N^2 \text{polylog } N$, compared to the worst-case runtime $O(N^3)$ of a classical DM. Importantly, from accessing of qRAM to performing an eigendecomposition of the Markov transition matrix, the total time complexity is only $O(\log^3 N)$. Such an exponential speedup for transition matrix construction and analysis could be useful in random-walk-based algorithms.

Note that other recently proposed quantum algorithms for manifold learning include qKPCA [19] and quantum spectral clustering [21]. However, the qDM has an appealing distance-preserving embedding property (6), possessed by neither qKPCA nor quantum spectral clustering. The inner working of the qDM is also starkly different from that of qKPCA and quantum spectral clustering. While qKPCA relies on approximating an arbitrary target kernel via the truncated Taylor’s expansion of linear kernels [19], the qDM proposes a method to *exactly* encode a Gaussian kernel from classical data, provided accessibility to qRAM. In addition, although both the qDM and quantum spectral clustering are based on spectral graph methods, the qDM manages to explicitly and efficiently construct the degree matrix inverse and the Markov transition operator, whose construction was mentioned not to be accessible efficiently [21]. Therefore, our work could provide a hint to achieve a quantum speedup in other algorithms that require a Markov transition operator and the analysis of its spectral properties.

Regarding the N^2 bottleneck, a quantum algorithm for the generalized coupon collector’s problem to speed up the process of finding all eigenpairs or a method to sort the eigenvalues and select only the top ones would reduce the complexity. Alternatively, one could devise a qDM algorithm that outputs the eigenvectors as quantum states and retain the exponential speedup. We leave open the question of how one could make use of such a state for nonlinear dimensionality reduction.

Recently, manipulation of matrices and their spectra by block encoding into submatrices of unitary matrices has become a topic of great interest due to its versatility in quantum algorithm design [36,37], which may provide an alternative route to the qDM algorithm.

ACKNOWLEDGMENTS

A.S. especially thanks Dimitris Angelakis from the Centre for Quantum Technologies (CQT), Singapore, for hospitality, and for useful advice on quantum algorithms. We also thank Jirawat Tangpanitanon and Supanut Thanasilp for useful dis-

cussions. This research has received funding support from the National Science, Research and Innovation Fund (NSRF) via the Program Management Unit for Human Resources & Institutional Development, Research and Innovation (Grant No. B05F630108), and Sci-Super VI fund from the faculty of science, Chulalongkorn University.

APPENDIX A: QMAT MATRIX MULTIPLICATION SUBROUTINE

As explained in Sec. IV A, all QMAT subroutines [26] utilized in our work reduce to that of matrix multiplication, which we now describe. Given matrices A_1 and A_2 , the idea of the multiplication subroutine is to construct a Hermitian embedding of the product $A_1 A_2$,

$$X_3(A_1 A_2) = \begin{pmatrix} 0 & 0 & A_1 A_2 \\ 0 & 0 & 0 \\ (A_1 A_2)^\dagger & 0 & 0 \end{pmatrix},$$

from the following commutator:

$$X_3(A_1 A_2) = U X_3(iA_1 A_2) U^\dagger = iU [X_1(A_1), X_2(A_2)] U^\dagger, \tag{A1}$$

where

$$U = \begin{pmatrix} \sqrt{-i}I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & \sqrt{i}I \end{pmatrix}.$$

Thus, the goal is to simulate the commutator of two Hermitian operators. This can be accomplished with bounded error via the second order of the Baker-Campbell-Hausdorff formula. Define

$$\tilde{I}(x, y) \equiv e^x e^y e^{-x} e^{-y} e^{-x} e^{-y} e^x e^y.$$

With the step size t/m , we have that [26]

$$\tilde{I}(xt/m, yt/m) = e^{2[x,y]t^2/m^2 + O(t^4/m^4)}. \tag{A2}$$

Thus, iterating the above $n' = m^2/(2t) = O(t^2/\epsilon_m)$ times,

$$\tilde{I}(xt/m, yt/m)^{n'/(2t)} = e^{[x,y]t + O(t^3/m^2)} \tag{A3}$$

correctly simulates $e^{[x,y]t}$ up to an error $\epsilon_m = O(t^3/m^2)$.

Putting all these together, the QMAT multiplication subroutine consumes $n' = m^2/(2t)$ copies of $e^{iX_1(A_1)t/m}$ and $e^{iX_2(A_2)t/m}$ (and their inverses), and outputs a unitary operator

$$U [I(iX_1(A_1)t/m, iX_2(A_2)t/m)]^{n'} U^\dagger, \tag{A4}$$

that approximates $e^{iX_3(A_1 A_2)}$ up to an error ϵ_m in the operator norm.

APPENDIX B: COMPLEXITY ANALYSIS

This Appendix describes the calculation for the time complexity that leads to the result in Theorem 1. We separate the algorithm into sequential steps as we have done with the classical DM (Sec. II B), namely, the embedding of classical data into the kernel \hat{K} , the calculation of the transition matrix P , and the calculation of the eigendecomposition.

1. Preparing the kernel

In creating \hat{K} , we consider only the time complexity coming from storing classical data into qRAM (Sec. III A) as the partial trace does not create additional time cost. Here we have assumed a common position among researchers in the field that encoding of classical data can be made efficient via the existence of an oracle [18,22,28,31,38,39], and analyze the algorithm based on this assumption. Note that practical feasibility and resource estimation of implementing qRAM is still an ongoing research topic [31,40–43]. Assuming that the data consist of N vectors each with d features, the data can be encoded into a qRAM of the size dN , which can be accessed using $O(\log(dN))$ operations (11) [38]. We assume here that we are dealing with big data, i.e., $d \leq N$. Therefore, the time cost becomes $O(\log N^2) = O(\log N)$.

2. Calculating the transition matrix

Next, we construct the exponential of P . We start with exponentiating the necessary matrices. Then we multiply these matrices in order to create $e^{iX_3(D)t}$. Next, we inverse the degree matrix D before constructing the transition matrix $P = D^{-1}K$.

The time cost of creating exponentials depends on whether the quantum simulation of matrices or quantum state is used. Creating exponentials from multiple copies of density matrices, i.e., $e^{iX_3(\hat{K})t}$ and $e^{iX_3(\hat{1})t}$, results in a time complexity of $O(t^2/\epsilon_e)$, where ϵ_e is the accuracy for the exponentiation [18]. Combined with the number of operations for storing the states, which is $O(\log N)$, the time complexity is $O(t^2/\epsilon_e \log N)$.

Exponentials of sparse matrices can be simulated using the method in Ref. [29] with query complexity of $O(t + \frac{\log(1/\epsilon_s)}{\log \log(1/\epsilon_s)})$, where ϵ_s is the accuracy of the method. For the matrices $e^{iX_3(\Xi)t}$ and $e^{iX_3(\Xi^\dagger)t}$, we also need to take into account the gate complexity $O(\log N)$. Thus the time complexity for preparing each of these exponentials is $O(t \log N)$.

To construct the degree matrix, we use the matrix multiplication subroutine in Appendix A. This subroutine uses four copies of each input exponential, which is multiplied and then applied n' times. In creating $e^{iX_3(\hat{K} \otimes \mathbb{1})t}$ (Sec. IV A), the time cost comes from multiplying the exponential of \hat{K} and of $\mathbb{1}$ whereas \otimes is assumed to be a free operation. Therefore, the time cost is $O(n'(\frac{t^2}{m^2 \epsilon_e} \log N))$. Given the condition that $n' = \frac{m^2}{2t} = O(\frac{t^2}{\epsilon_m})$ [26], we rewrite the time cost to $O(t \log N \epsilon_e^{-1})$.

In the next two steps, we multiply Ξ and Ξ^\dagger into $K \mathbb{1} \otimes I$. Again, we apply the multiplication subroutine iteratively with four copies of the result from the previous multiplication as one of the inputs. These steps admit an additional order of ϵ_m and t which leads the time cost of $e^{iX_3(D)t}$ to be

$$T_D(t) = O\left[\left(\frac{t^{1.5}}{\epsilon_m^{0.5}} + \frac{t^{1.25}}{\epsilon_e \epsilon_m^{0.75}}\right) \log N\right]. \quad (\text{B1})$$

Next, the degree matrix D is inverted (Sec. IV B) by first performing QPE that takes the prepared $e^{iX_3(D)t}$ as the unitary operator in QPE, then a controlled rotation. The time cost of the QPE is $O(T_D/\epsilon_D \log m)$ [44], where m is the number of eigenvalues which is $m = N$, and ϵ_D is the error of estimated phase. Assuming the controlled rotation is executed in time $O(1)$, we have accumulated a time cost of

$T_D + O(T_D/\epsilon_D \log N) + O(1)$. The time complexity is thus dominated by $O(T_D/\epsilon_D \log N)$.

The time T_D (B1) is a function of t , which we now find through error analysis of matrix inversion. As we are inverting the eigenvalues of D to a function of $1/\sqrt[4]{x}$ [28], the error after conditional rotation becomes $\epsilon_D = O(1/\sqrt[4]{d}t) < O(\kappa_D^{0.25}/t)$. Thus $t = O(\kappa_D^{0.25}/\epsilon_D)$.

As the success probability of measuring $|1\rangle$ in the ancilla qubit is $\Omega(1/\kappa_D^{0.5})$, the number of rounds necessary to obtain a successful result is $O(\kappa_D^{0.5})$. To accelerate the process, we use the amplitude amplification to amplify the success probability in $O(\kappa_D^{0.25})$ steps. Thus an overall complexity becomes

$$T_{D^{-1/2}} = O\left[\left(\frac{\kappa_D^{0.625}}{\epsilon_D^{2.5} \epsilon_m^{0.5}} + \frac{\kappa_D^{0.5625}}{\epsilon_D^{2.25} \epsilon_e \epsilon_m^{0.75}}\right) \log^2 N\right]. \quad (\text{B2})$$

The last step in constructing the symmetrized transition matrix is multiplying the inverse square-root degree matrix to \hat{K} , which we perform in the exponential (Sec. IV B). Therefore, we first need to exponentiate $D^{-1/2}$ and \hat{K} using a time of $O(t^2/(m^2 \epsilon_e))$ and $4n'$ copies each. The time complexity for constructing $e^{iX_3(S)t}$ is therefore

$$T_S(t) = O\left[T_{D^{-1/2}}\left(\frac{t^{1.5}}{\epsilon_m^{0.5} \epsilon_e} + \frac{t}{\epsilon_e}\right) + \frac{t^{1.5}}{\epsilon_m^{0.5} \epsilon_e} \log N\right].$$

3. Eigendecomposition

In this section, we find the eigenvalues and eigenvectors of S , using a second QPE algorithm (Sec. V). We take $t = O(1/\epsilon_S)$ to perform the QPE in time $O(\log N)$ with accuracy ϵ_S [44].

The third and final QPE is required to transform $|s_i\rangle$ to the right eigenvector $|v_i\rangle = D^{-1/2}|s_i\rangle$ of the transition matrix $P = D^{-1}K$. The complexity analysis proceeds analogously to the analysis in Appendix B 2: the error analysis of matrix inversion gives $t = O(\kappa_D^{0.5}/\epsilon_D)$. As the success probability of measuring $|1\rangle$ in the ancilla qubit is $\Omega(1/\kappa_D)$, the number of rounds necessary to obtain a successful result is $O(\kappa_D)$. To accelerate the process, amplitude amplification is used to amplify the success probability in $O(\kappa_D^{0.5})$ steps. Incorporating the time required to prepare copies of $e^{iX_3(D)t}$, the time complexity of the conditional rotation is hence $O(T_D(\kappa_D^{0.5}/\epsilon_D)\kappa_D^{0.5} \log N/\epsilon_D)$. Thus, the overall complexity of the eigendecomposition of the transition matrix becomes

$$O\left[T_S\left(\frac{1}{\epsilon_S}\right) \frac{\log N}{\epsilon_S} + T_D\left(\frac{\kappa_D^{0.5}}{\epsilon_D}\right) \frac{\kappa_D^{0.5} \log N}{\epsilon_D}\right], \quad (\text{B3})$$

which proves Theorem 1.

We repeat the QPE and measure the eigenvalue register in the computation basis until we find sufficiently many copies of all eigenpairs for tomography. The process is not deterministic and we can only talk about the expected number of repetitions required. The task of drawing items (with replacement) distributed according to a given probability distribution $\{p_i\}_{i=1}^N$ until one collects all N distinct items is known as the coupon collector's problem, and an upper bound on the expected number of draws is known to be $(\max_i\{p_i\} / \min_i\{p_i\})N \log N$ [34]. Once we repeat the QPE and measure the eigenvalue register in the computation basis until we obtain $O(N \log N/\epsilon_t)$ copies

of each eigenpair, the pure-state tomography algorithm of Ref. [33] recovers a classical vector that is ϵ_t close in ℓ_2 -norm to the eigenvector. The above facts together with Theorem 1 immediately imply Lemma 1.

Finally, in the case of uniformly random coupons, i.e., $p_i = p_j$ for all $i, j \in \{1, \dots, N\}$, there exists a gate-efficient quantum algorithm that solves the coupon collector's problem in time $O(N)$ [35], giving Lemma 2.

-
- [1] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, A global geometric framework for nonlinear dimensionality reduction, *Science* **290**, 2319 (2000).
- [2] M. Belkin and P. Niyogi, Laplacian eigenmaps and spectral techniques for embedding and clustering, in *Advances in Neural Information Processing Systems 14* (MIT Press, Cambridge, MA, 2001), pp. 585–591.
- [3] L. McInnes, J. Healy, N. Saul, and L. Großberger, UMAP: Uniform manifold approximation and projection, *J. Open Source Softw.* **3**, 861 (2018).
- [4] M. Scholz, F. Kaplan, C. L. Guy, J. Kopka, and J. Selbig, Nonlinear PCA: A missing data approach, *Bioinformatics* **21**, 3887 (2005).
- [5] L. van der Maaten and G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* **9**, 2579 (2008).
- [6] R. R. Coifman and S. Lafon, Diffusion maps, *Appl. Comput. Harmon. Anal.* **21**, 5 (2006).
- [7] S. Lafon, Diffusion maps and geometric harmonics, Ph.D. thesis, Yale University, 2004.
- [8] B. Nadler, S. Lafon, I. Kevrekidis, and R. Coifman, Diffusion maps, spectral clustering and eigenfunctions of Fokker-Planck operators, in *Proceedings of the 18th International Conference on Neural Information Processing Systems*, edited by Y. Weiss, B. Schölkopf, and J. Platt (MIT Press, Cambridge, MA, 2006), Vol. 18.
- [9] K. R. Moon, D. van Dijk, Z. Wang, S. Gigante, D. B. Burkhardt, W. S. Chen, K. Yim, A. van den Elzen, M. J. Hirn, R. R. Coifman, N. B. Ivanova, G. Wolf, and S. Krishnaswamy, Visualizing structure and transitions in high-dimensional biological data, *Nat. Biotechnol.* **37**, 1482 (2019).
- [10] C. L. Ecale Zhou, S. Malfatti, J. Kimbrel, C. Philipson, K. McNair, T. Hamilton, R. Edwards, and B. Souza, multiPhATE: Bioinformatics pipeline for functional annotation of phage isolates, *Bioinformatics* **35**, 4402 (2019).
- [11] J. F. Rodriguez-Nieva and M. S. Scheurer, Identifying topological order through unsupervised machine learning, *Nat. Phys.* **15**, 790 (2019).
- [12] Y. Long, J. Ren, and H. Chen, Unsupervised Manifold Clustering of Topological Phononics, *Phys. Rev. Lett.* **124**, 185501 (2020).
- [13] A. Lidiak and Z. Gong, Unsupervised Machine Learning of Quantum Phase Transitions Using Diffusion Maps, *Phys. Rev. Lett.* **125**, 225701 (2020).
- [14] A. Kerr, G. Jose, C. Riggert, and K. Mullen, Automatic learning of topological phase boundaries, *Phys. Rev. E* **103**, 023310 (2021).
- [15] Y. Che, C. Gneiting, T. Liu, and F. Nori, Topological quantum phase transitions retrieved through unsupervised machine learning, *Phys. Rev. B* **102**, 134213 (2020).
- [16] J. Wang, W. Zhang, T. Hua, and T.-C. Wei, Unsupervised learning of topological phase transitions using the Calinski-Harabaz index, *Phys. Rev. Research* **3**, 013074 (2021).
- [17] P. Reberstrost, A. Steffens, I. Marvian, and S. Lloyd, Quantum singular-value decomposition of nonsparse low-rank matrices, *Phys. Rev. A* **97**, 012327 (2018).
- [18] S. Lloyd, M. Mohseni, and P. Reberstrost, Quantum principal component analysis, *Nat. Phys.* **10**, 631 (2014).
- [19] Y. Li, R.-G. Zhou, R. Xu, W. Hu, and P. Fan, Quantum algorithm for the nonlinear dimensionality reduction with arbitrary kernel, *Quantum Sci. Technol.* **6**, 014001 (2020).
- [20] S. Apers and R. de Wolf, Quantum speedup for graph sparsification, cut approximation and Laplacian solving, in *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE, New York, 2020), pp. 637–648.
- [21] I. Kerenidis and J. Landman, Quantum spectral clustering, *Phys. Rev. A* **103**, 042415 (2021).
- [22] I. Kerenidis, J. Landman, A. Luongo, and A. Prakash, q-means: A quantum algorithm for unsupervised machine learning, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, 2019).
- [23] R. Chatterjee and T. Yu, Generalized coherent states, reproducing kernels, and quantum support vector machines, *Quantum Inf. Comput.* **17**, 1292 (2017).
- [24] The Wine data set is obtained from D. Hua and C. Graff, UCI machine learning repository (2017), <https://archive.ics.uci.edu/ml/datasets/Wine>.
- [25] L. N. Trefethen and D. Bau III, *Numerical Linear Algebra* (SIAM, Philadelphia, 1997), Vol. 50.
- [26] L. Zhao, Z. Zhao, P. Reberstrost, and J. Fitzsimons, Compiling basic linear algebra subroutines for quantum computers, *Quantum Mach. Intell.* **3**, 21 (2021).
- [27] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum Algorithm for Linear Systems of Equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [28] I. Cong and L. Duan, Quantum discriminant analysis for dimensionality reduction and classification, *New J. Phys.* **18**, 073011 (2016).
- [29] G. H. Low and I. L. Chuang, Optimal Hamiltonian Simulation by Quantum Signal Processing, *Phys. Rev. Lett.* **118**, 010501 (2017).
- [30] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Random Access Memory, *Phys. Rev. Lett.* **100**, 160501 (2008).
- [31] J. Biamonte, P. Wittek, N. Pancotti, P. Reberstrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [32] An $N \times N$ matrix is s sparse if there are at most s nonzero entries per row. An $N \times N$ matrix is *sparse* if it is at most $\text{polylog}(N)$ sparse.
- [33] I. Kerenidis and A. Prakash, A quantum interior point method for LPs and SDPs, *ACM Trans. Quantum Comput.* **1**, 1 (2020).
- [34] N. B. Shank and H. Yang, Coupon collector problem for non-uniform coupons and random quotas, *Electron. J. Combinatorics* **20**, P33 (2013).

- [35] S. Arunachalam, A. Belovs, A. M. Childs, R. Kothari, A. Rosmanis, and R. de Wolf, Quantum coupon collector, in *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, Leibniz International Proceedings in Informatics (LIPIcs), edited by S. T. Flammia (Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2020), Vol. 158, pp. 10:1–10:17.
- [36] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics, in *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing, STOC 2019* (Association for Computing Machinery, New York, 2019), pp. 193–204.
- [37] J. M. Martyn, Z. M. Rossi, A. K. Tan, and I. L. Chuang, A grand unification of quantum algorithms, [arXiv:2105.02859](https://arxiv.org/abs/2105.02859).
- [38] P. Rebentrost, M. Mohseni, and S. Lloyd, Quantum Support Vector Machine for Big Data Classification, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [39] H. Yamasaki, S. Subramanian, S. Sonoda, and M. Koashi, Learning with optimized random features: Exponential speedup by quantum machine learning without sparsity and low-rank assumptions, in *Advances in Neural Information Processing Systems 33*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (Curran Associates, Red Hook, NY, 2020), pp. 13674–13687.
- [40] N. Jiang, Y.-F. Pu, W. Chang, C. Li, S. Zhang, and L.-M. Duan, Experimental realization of 105-qubit random access quantum memory, *npj Quantum Inf.* **5**, 28 (2019).
- [41] C. T. Hann, C.-L. Zou, Y. Zhang, Y. Chu, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, Hardware-Efficient Quantum Random Access Memory with Hybrid Quantum Acoustic Systems, *Phys. Rev. Lett.* **123**, 250501 (2019).
- [42] O. Di Matteo, V. Gheorghiu, and M. Mosca, Fault-tolerant resource estimation of quantum random-access memories, *IEEE Trans. Quantum Eng.* **1**, 1 (2020).
- [43] C. T. Hann, G. Lee, S. M. Girvin, and L. Jiang, Resilience of quantum random access memory to generic noise, *PRX Quantum* **2**, 020311 (2021).
- [44] D. Dervovic, M. Herbster, P. Mountney, S. Severini, N. Usher, and L. Wossnig, Quantum linear systems algorithms: A primer, [arXiv:1802.08227](https://arxiv.org/abs/1802.08227).