

Dynamic quantum clustering: A method for visual exploration of structures in dataMarvin Weinstein¹ and David Horn²¹*Stanford Linear Accelerator Center, Stanford, California 94025, USA*²*School of Physics and Astronomy, Tel Aviv University, Tel Aviv 69978, Israel*

(Received 17 August 2009; published 17 December 2009)

A given set of data points in some feature space may be associated with a Schrödinger equation whose potential is determined by the data. This is known to lead to good clustering solutions. Here we extend this approach into a full-fledged dynamical scheme using a time-dependent Schrödinger equation. Moreover, we approximate this Hamiltonian formalism by a truncated calculation within a set of Gaussian wave functions (coherent states) centered around the original points. This allows for analytic evaluation of the time evolution of all such states opening up the possibility of exploration of relationships among data points through observation of varying dynamical distances among points and convergence of points into clusters. This formalism may be further supplemented by preprocessing such as dimensional reduction through singular-value decomposition or feature filtering.

DOI: [10.1103/PhysRevE.80.066117](https://doi.org/10.1103/PhysRevE.80.066117)

PACS number(s): 89.75.Fb, 89.90.+n, 95.75.Pq, 89.75.Kd

I. INTRODUCTION

Clustering of data is, in general, an ill-defined problem. Nonetheless, it is a very important one in many scientific and technological fields of study. Given a set of data points one looks for possible structures by sorting out which points are close to each other and, therefore, in some sense belong together. This is a preliminary stage taken before investigating what properties are common to these subsets of the data.

It has recently become quite popular to investigate such questions in a dynamic framework, thus, allowing for more fluid associations of data points rather than rigid clusters. Diffusion geometry [1–3] is such a method based on a discrete analog of the heat equation

$$-\frac{\partial \Phi}{\partial t} = H\Phi, \quad (1)$$

where H is some operator with positive eigenvalues guaranteeing that the temporal evolution of $\Phi(\vec{x}, t)$ is that of diffusion. Thus, starting out with $\Phi(\vec{x}, 0)$, e.g., a Gaussian concentrated around some data point one would expect $\Phi(\vec{x}, t)$ to spread over all space that is occupied by the data points.

In this paper we advocate the use of a Schrödinger Hamiltonian H that is intimately connected to the data structure as defined by the quantum clustering method [4] and summarized below. We extend it into a time-dependent Schrödinger equation:

$$i\frac{\partial \psi(\vec{x}, t)}{\partial t} = H\psi(\vec{x}, t). \quad (2)$$

The ensuing dynamic quantum clustering (DQC) formalism allows us, by varying a few parameters, to study in detail the temporal evolution of wave functions representing the original data points. In turn, this dynamical behavior allows us to explore the structure of the quantum potential function defined by the quantum clustering method.

DQC begins by associating each data point with a state in Hilbert space. The temporal development of the centroids of these states may be viewed in the original data space as moving images of the original points. Their distances to one

another change with time, thus, representing associations they form with each other. Convergence of many points onto a common center at some instant of time is an obvious manifestation of clustering. Many transitional relationships may occur revealing substructures in clusters or even more complex associations. For this reason we propose this approach as a general method for visually and interactively searching for and exploring structures in sets of data.

II. QUANTUM CLUSTERING

The quantum clustering approach begins, as does the well-known Parzen-window estimator [5], by associating to each of n data points \vec{x}_i in a Euclidean space of d dimensions a Gaussian wave function $\psi_i(\vec{x}) = e^{-(\vec{x} - \vec{x}_i)^2/2\sigma^2}$ and then constructing the sum of all these Gaussians,

$$\psi(\vec{x}) = \sum_i e^{-(\vec{x} - \vec{x}_i)^2/2\sigma^2}. \quad (3)$$

Conventional scale-space clustering [6] views this function as a probability distribution (up to an overall factor) that could have generated the observed points and regards, therefore, its maxima as determining locations of cluster centers. Often these maxima are not very prominent and, in order to uncover more of them, one has to reduce σ down to low values where the number and location of the maxima depend sensitively upon the choice of σ .

Quantum clustering took a different approach requiring ψ to be the ground state of the Hamiltonian

$$H\psi \equiv \left[-\frac{\sigma^2}{2}\nabla^2 + V(\mathbf{x}) \right] \psi = E_0\psi. \quad (4)$$

By positing this requirement [4], the potential function $V(\mathbf{x})$ has become inextricably bound to the system of data points, since $V(\vec{x})$ is determined, up to a constant, by a simple algebraic inversion of Eq. (4) [7]. Moreover, one may expect V to have minima in regions where ψ has maxima and, furthermore, that these minima will be more pronounced than the corresponding maxima found in the Parzen estimator. In fact,

it frequently turns out that a concentration of data points will lead to a local minimum in V , even if ψ does not display a local maximum. Thus, by replacing the problem of finding maxima of the Parzen estimator by the problem of locating the minima of the associated potential, $V(\vec{x})$, we simplify the process of identifying clusters. The effectiveness of this approach has been demonstrated in the work by Horn and Gottlieb [4,8]. It should be noted that the enhancement of features obtained by applying Eq. (4) comes from the interplay of two effects: attraction of the wave function to the minima of V and spreading of the wave function due to the second derivative (kinetic term). This may be viewed as an alternative model to the conventional probabilistic approach incorporating attraction to cluster centers and creation of noise, both inferred from—or realized by—the given experimental data.

DQC drops the probabilistic interpretation of ψ and replaces it by that of a probability amplitude as customary in quantum mechanics. DQC is set up to associate data points with cluster centers in a natural fashion. Whereas in QC this association was done by finding their loci on the slopes of V , here we follow the quantum-mechanical temporal evolution of states associated with these points. Specifically, we will view each data point as the expectation value of the position operator in a Gaussian wave function $\psi_i(\vec{x}) = e^{-(\vec{x}-\vec{x}_i)^2/2\sigma^2}$; the temporal development of this state traces the association of the data point it represents with the minima of $V(\vec{x})$ and, thus, with the other data points.

III. DQC

As we already noted, the conversion of the static QC method to a full dynamical one begins by focusing attention on the Gaussian wave function, $\psi_i(\vec{x}) = Ce^{-(\vec{x}-\vec{x}_i)^2/2\sigma^2}$, associated with the i th data point, where C is the appropriate normalization factor. Thus, by construction, the expectation value of the operator \vec{x} in this state is simply the coordinates of the original data point; i.e.,

$$\vec{x}_i = \langle \psi_i | \vec{x} | \psi_i \rangle = \int d\vec{x} \psi_i^*(\vec{x}) \vec{x} \psi_i(\vec{x}). \quad (5)$$

The dynamical part of the DQC algorithm is that, having constructed the potential function $V(\vec{x})$, we study the time evolution of each state $\psi_i(\vec{x})$ as determined by the time-dependent Schrödinger equation; i.e.,

$$i \frac{\partial \psi_i(\vec{x}, t)}{\partial t} = H \psi_i(\vec{x}, t) = \left[-\frac{\nabla^2}{2m} + V(\vec{x}) \right] \psi_i(\vec{x}, t), \quad (6)$$

where m is an arbitrarily chosen mass for a particle moving in d dimensions. If we set $m=1/\sigma^2$ then, by construction, $\psi_i(\vec{x})$ of Eq. (3) is the lowest energy eigenstate of the Hamiltonian.

The important feature of quantum dynamics, which makes the evolution so useful in the clustering problem, is that according to Ehrenfest's theorem, the time-dependent expectation value

$$\langle \psi_i(t) | \vec{x} | \psi_i(t) \rangle = \int d\vec{x} \psi_i^*(\vec{x}, t) \vec{x} \psi_i(\vec{x}, t) \quad (7)$$

satisfies the equation

$$\frac{d^2 \langle \vec{x}(t) \rangle}{dt^2} = -\frac{1}{m} \int d\vec{x} \psi_i^*(\vec{x}, t) \vec{\nabla} V(\vec{x}) \psi_i(\vec{x}, t) \quad (8)$$

$$= -\frac{1}{m} \langle \psi_i(t) | \vec{\nabla} V(\vec{x}) | \psi_i(t) \rangle. \quad (9)$$

If $\psi_i(\vec{x})$ is a narrow Gaussian, this is equivalent to saying that the center of each wave function rolls toward the nearest minimum of the potential according to the classical Newton's law of motion. This means we can explore the relation of this data point to the minima of $V(\vec{x})$ by following the time-dependent trajectory $\langle \vec{x}_i(t) \rangle = \langle \psi_i(t) | \vec{x} | \psi_i(t) \rangle$. Clearly, given Ehrenfest's theorem, we expect to see any points located in, or near, the same local minimum of $V(\vec{x})$ to oscillate about that minimum coming together and moving apart. In our numerical solutions we generate animations which display this dynamics for a finite time. This allows us to visually trace the clustering of points associated with each one of the potential minima.

In their quantum clustering paper Horn and Gottlieb successfully used classical gradient descent to cluster data by moving points (on classical trajectories) to the nearest local minimum of $V(\vec{x})$. The idea being that points which end up at the same minimum are in the same cluster. At first glance it would seem that DQC replaces the conceptually simple problem of implementing gradient descent with the more difficult one of solving complicated partial differential equations. We will show the difficulty is only apparent. In fact, the solution of the Schrödinger equation can be simplified considerably and will also allow further insights than the gradient descent method. The DQC algorithm translates the problem of solving the Schrödinger equation into a matrix form which captures most of the details of the analytic problem, but which involves $N \times N$ matrices whose dimension, N , is less than or equal to the number of data points. This reduction is independent of the data dimension of the original problem. From a computational point of view there are many advantages to this approach. First, the formulas for constructing the mapping of the original problem to a matrix problem are all analytic and easy to evaluate, thus, computing the relevant reduction is fast. Second, the evolution process only involves matrix multiplications, so many data points can be evolved simultaneously and, on a multicore processor, in parallel. Third, the time involved in producing the animations showing how the points move in data space scales linearly with the number of dimensions to be displayed. Finally, by introducing an m that is different from $1/\sigma^2$ we allow ourselves the freedom of employing low σ , which introduces large numbers of minima into V , yet also having a low value for m which guarantees efficient tunneling, thus, connecting points that may be located in nearby, nearly degenerate potential minima. By using this more general Hamiltonian, we reduce the sensitivity of the calculation to the specific choice of σ .

One final point worth making before describing the method of calculation is that the use of Gaussian wave functions to represent data points allows us to develop a number of flexible strategies for handling very large data sets. This issue will be addressed below.

IV. CALCULATION METHOD

Before discussing how this method works in practice we will give a brief outline of the details of the general procedure. We begin by assuming that there are n -data points that we wish to cluster. To these data points we associate n states, $|\psi_i\rangle$. These states are n Gaussian wave functions such that the i th Gaussian is centered on the coordinates of the i th data point. These states form a basis for the vector space within which we calculate the evolution of our model.

Let us denote by N , the $n \times n$ matrix formed from the scalar products

$$N_{i,j} = \langle \psi_i | \psi_j \rangle, \quad (10)$$

and by \mathcal{H} , the $n \times n$ matrix

$$\mathcal{H}_{i,j} = \langle \psi_i | H | \psi_j \rangle, \quad (11)$$

and by $\vec{X}_{i,j}$ the matrix of expectation values

$$\vec{X}_{i,j} = \langle \psi_i | \vec{x} | \psi_j \rangle. \quad (12)$$

The calculation process can be described in five steps. First, begin by finding the eigenvectors of the symmetric matrix N which correspond to states having eigenvalues larger than some preassigned value; e.g., 10^{-5} . These vectors are linear combinations of the original Gaussians which form an orthonormal set. Second, compute \mathcal{H} in this orthonormal basis, \mathcal{H}^{tr} . Do the same for $\vec{X}_{i,j}$. Fourth, find the eigenvectors and eigenvalues of \mathcal{H}^{tr} , construct $|\psi_i(t)\rangle = e^{-i\mathcal{H}^{tr}t} |\psi_i\rangle$, that is the solution to the reduced time-dependent Schrödinger problem

$$i \frac{\partial}{\partial t} |\psi_i(t)\rangle = \mathcal{H}^{tr} |\psi_i(t)\rangle \quad (13)$$

such that $|\psi_i(t=0)\rangle = |\psi_i\rangle$. Finally, construct the desired trajectories

$$\langle \vec{x}_i(t) \rangle = \langle \psi_i | e^{i\mathcal{H}^{tr}t} \vec{X} e^{-i\mathcal{H}^{tr}t} | \psi_i \rangle \quad (14)$$

by evaluating this expression for a range of t and use them to create an animation. Stop the animation when clustering of points is obvious.

It is clear that restricting attention to the truncated Hamiltonian perforce loses some features of the original problem; however, its advantage is that we can derive analytic expressions for all operators involved (see Appendixes A and B). As a result, the numerical computations can be done very quickly. Experience has shown that as far as clustering is concerned this approximation causes no difficulties.

V. EXAMPLE: RIPLEY'S CRAB DATA

To test our method we apply it to a five-dimensional data set with 200 entries used in Ripley's text book [9]. This data

set records five measurements made on male and female crabs that belong to two different species. This data set has been used in the original paper on quantum clustering by Horn and Gottlieb [4]. It is being used here to allow readers to compare the two techniques. Applications to other data sets will be discussed below. Our main motivation is to provide a simple example which exhibits the details of the DQC method. In particular, we wish to show that the simplest computational scheme for implementing the general program captures the essential features of the problem and does as well as one can reasonably expect to do.

The data are stored in a matrix M which has 200 rows and five columns. Following an often-used dimensional reduction method, we preprocess our data with a singular-value decomposition (SVD)

$$M = USV^\dagger, \quad (15)$$

where U is a unitary 200×200 matrix and S is the 200×5 matrix of singular values, the latter occurring on the diagonal of its upper 5×5 entries. The submatrix of U consisting of the first five columns, the so-called five principal components (PCs), can be thought of as assigning to each sample a unique point in a five-dimensional vector space. We may study the problem in the full five-dimensional space or within any subspace by selecting appropriate principal components. In [4] QC was applied to this problem in a two-dimensional subspace consisting of PC2 and PC3. In what follows we will discuss the application of DQC to the three-dimensional data composed of the first three PCs (although there would be no change in the results if we used all five dimensions). In order to give the reader some feeling for how the quantum potential associated with the data looks in a simple case, we have included Fig. 1 where we exhibit the original data points (different colors of points correspond to known classes), placed upon the associated two-dimensional quantum potential.

As is obvious from the plot the minima of the potential function do a very good job of capturing the different clusters. Moreover, letting data points roll downhill to the nearest minimum will produce a reasonable separation of the clusters.

Clearly, when we restrict attention to the first three PCs, the rows of the matrix obtained by restricting U to its first three columns are not guaranteed to be normalized to unity. Hence, we employ the conventional approach of projecting all points onto the unit sphere [10].

In what follows we study the temporal behavior of the curves $\langle \vec{x}_i(t) \rangle$ for all i . Henceforth, we will refer to this as the "motion of points." Figure 2 shows the distribution of the original data points plotted on the unit sphere in three dimensions. This is the configuration before we begin the dynamic quantum evolution. To visually display the quality of the separation we have colored the data according to its known four classes; however, this information is not incorporated into our unsupervised method. To begin with, we see that the two species of crabs (one species in the upper left- and right-hand corners, and the other in the lower left- and right-hand corners) are fairly well separated; however, separating the sexes in each species is problematic. The middle plot in Fig.




FIG. 1. (Color online) This is a plot of the quantum potential function for the two-dimensional problem of the Ripley’s crab data where the coordinates of the data points are chosen to be given by the second and third principal components. The four known classes of data points are shown (online only) in different colors proceeding counterclockwise from the upper right-hand corner (blue, green, yellow, and red) and they are placed upon the potential surface at their original locations.

2 shows the distribution of the points after a single stage of quantum evolution stopped at a time when points first cross one another and some convergence into clusters has occurred. It is immediately obvious that the quantum evolution has enhanced the clustering and made it trivial to separate clusters by eye. Once separation is accomplished, extracting the clusters can be performed by eye from the plots or by any conventional technique, e.g., *k* means.

An alternative way of displaying convergence is shown in Fig. 3, where we plot the Euclidean distance from the first point in the data set to each of the other points. The clusters lie in bands which have approximately the same distance from the first point. It is difficult to get very tight clusters since the points, while moving toward cluster centers, oscillate

late around them, and arrive at the minima at slightly different times. Given this intuition, it is clear that one way to tighten up the pattern is to stop DQC evolution at a point where the clusters become distinct, and then restart it with the new configuration, but with the points redefined at rest. We refer to this as iterating the DQC evolution. The right-hand plots in Figs. 2 and 3 show what happens when we do this. The second stage of evolution clearly tightens up the clusters significantly as was expected.

By the end of the second iteration, there can be no question that it is a simple matter to extract the clusters. As is quite evident, clustering does not agree completely with the expert classification; i.e., points with different colors may be grouped together. This is, however, the best one can do by color-blind treatment of the information provided in the data matrix. As we already noted the full five-dimensional study of the crab data set can proceed in the same manner, although it does not lead to new insights.

VI. DYNAMIC DISTANCES

The fact that data points of different classes happen to lie close to each other in the data matrix can be due to various factors: errors in data measurements, errors in the expert assignment to classes, true proximity of data points in spite of differences of origin (an extreme example would be similarities of phenotypes in spite of differences in genotypes), or—the simplest possibility—the absence of some discriminative features in the feature space that spans the data measurements. But there is another important conceptual message to be learned here—clustering and/or classification may not capture all the interesting lessons that may be derived from the data. A similar message is included in the diffusion geometry approach [1,2] that advocates measuring diffusion distances among points rather than Euclidean ones. Diffusion distances are influenced by the existence of all other points. In our DQC analysis this may be replaced in a straightforward manner by defining dynamic distances among points

$$d_{i,j}(t) = \|\langle \vec{x}_i(t) \rangle - \langle \vec{x}_j(t) \rangle\| \tag{16}$$

with the norm being Euclidean or any other suitable choice.

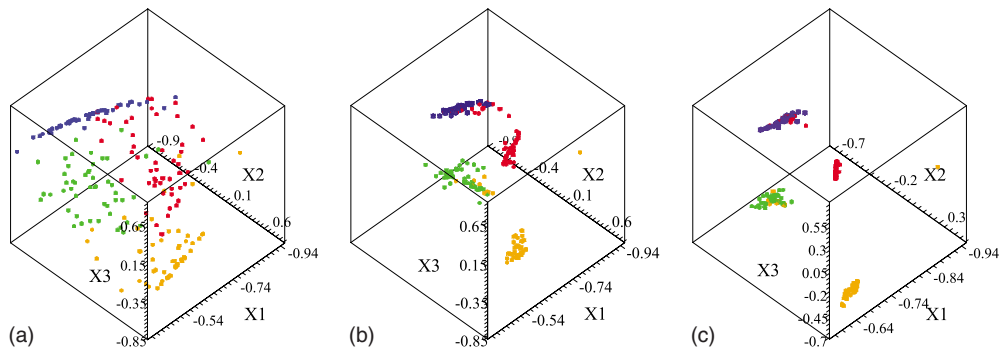


FIG. 2. (Color online) The left-hand plot shows three-dimensional distribution of the original data points before quantum evolution. The middle plot shows the same distribution after quantum evolution. The right-hand plot shows the results of an additional iteration of DQC. The values of parameters used to construct the Hamiltonian and evolution operator are $\sigma=0.07$ and $m=0.2$. Colors indicate the expert classification of data into four classes unknown to the clustering algorithm. Note that small modifications of the parameters lead to the same results. The first species of crabs (male and female) are designated by the blue and red points in the upper left- and right-hand corners, respectively, and the second by the green and orange points in the lower left- and right-hand corners.

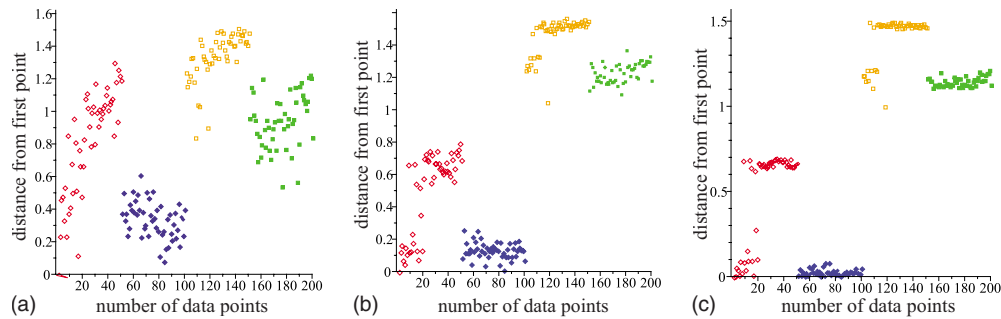


FIG. 3. (Color online) A plot of Euclidean distance of each point i from the first data point. Again, the left-hand plot shows the distances for the initial distribution of points. The middle plot shows the same distances after quantum evolution. The right-hand plot shows results after another iteration of DQC. The numbering of the data points is ordered according to the expert classification of these points into four classes containing 50 instances each. The colors of the points are, from left to right, red—open diamond, blue—filled diamonds, yellow—open squares, and green—filled squares.

Clearly $d_{i,j}(0)$ is the geometric distance as given by the original data matrix or by its reduced form that is being investigated. As DQC evolves with time $d_{i,j}(t)$ changes, and when some semiperfect clustering is obtained, it will be close to zero for points that belong to the same cluster. Figure 3 shows this change in time for all $d_{i,1}(t)$ in the crab-data example studied above. It is quite obvious that, in addition to the few cases in which clustering disagrees with classification, there are many intermediate steps where different data points are close to each other in spite of eventually evolving into different clusters and belonging to different classes. Thus, a close scrutiny of the dynamic distances matrix $d_{i,j}(t)$ may lead to interesting observations regarding the relationships among individual pairs of points in the original data, a relationship that is brought out by DQC as result of the existing information about all other data points. It may be used to further investigate the reason for such proximities, along any one of the lines mentioned above, and thus may lead to insights regarding the problem at hand.

VII. ANALYSIS OF LARGE DATA SETS

There are many scientific and commercial fields, such as cosmology, epidemiology, and risk management, where the data sets of interest contain many points, often also in large numbers of dimensions. We have already discussed how to overcome the problem of large dimensions. Dealing with large number of points requires yet a new angle. In problems of this sort it is clear from the outset that, especially on a PC, diagonalizing matrices which are larger than 2000×2000 is computationally intensive. It is obvious that using brute force methods to evolve sets of data having tens of thousands of points simply would not work. The solution to the problem of dealing with sets of data containing tens of thousands of entries each with N features lies in the fact that the SVD decomposition maps the data into an N -dimensional cube and the fact that the data points are represented by states in Hilbert space rather than N -tuples of real numbers. Since there is literature on ways to do SVD decomposition for large sets of data, we will not address this point here. What we do wish to discuss is how to exploit the representation of data points as states in Hilbert space in order to evolve large sets of data.

The trick is to observe that since Gaussian wave functions whose centers lie within a given cube have nonvanishing overlaps, as one chooses more and more wave functions one eventually arrives at a situation where the states become what we will refer to as *essentially linearly dependent*. In other words, we arrive at a stage at which any new wave function added to the set can, to some predetermined accuracy, be expressed as a linear combination of the wave functions we already have. Of course, since quantum mechanical time evolution is a linear process, this means that this additional state can be evolved by expressing it as a linear combination of the previously selected states and evolving them. Since computing the overlap of two Gaussians is done analytically (Appendix B) determining which points determine the set of *maximally essentially linearly independent states for the problem* is easy. Typically, even for data sets with 35 000 points, this is of the order of 1000 points. This works because, as we have already noted, we do not need high accuracy for DQC evolution. The quality of the clustering degrades very slowly with loss in accuracy. Thus, we can compute the time evolution operator in terms of a well chosen subset of the data and then apply it to the whole set of points [11].

To demonstrate this versatility of DQS we have analyzed a set of 35 213 points in 20 dimensions [12]. This data set definitely shows some nontrivial variations in density that can be made apparent by a visual inspection of the data plotted in different dimensional combinations. However, DQC is needed to obtain good visual deciphering of the different structures. Selecting a subset of 1200 points, whose Gaussians we consider to be a set of essentially linearly independent states for the problem, we construct \mathcal{H}^M . By expanding the remaining states in terms of these 1200 states we can easily evaluate the DQC evolution of all 35 213 points. The results are displayed in Fig. 4. It seems very clear how the structures develop with DQC. Using the last DQC stage it is possible to identify the main structures and assign each substructure a different color. One can then examine the colored version of a plot of the individual data points, discern the structures that belong together, and follow the DQC development tracing out dynamics distances between the different points and structures in all dimensions.

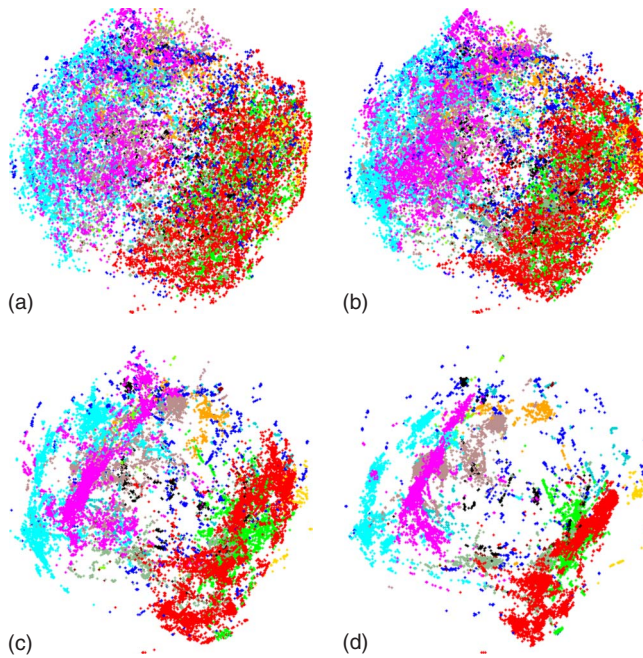


FIG. 4. (Color online) A plot of the first three principal components for a large data set comprising 35 213 points in 20 dimensions before and after DQC evolution. The potential was determined from the full data set and evolved using a subset of 1200 points whose Gaussians serve as an essentially linear set of independent states. Three stages of DQC development are shown. In this plot the colors have no significance; the important thing is that they make the clusters more obvious. The coloring was decided upon by selecting the most obvious clusters from the evolved data and assigning colors to them. The dark blue points (these are the darkest points) correspond to points that we did not bother to assign to clusters. The purpose of coloring is to be able to look at the points in the original data, discern those that belong to common structures, and follow their dynamic distances under DQC evolution.

VIII. INTERPLAY OF FEATURE SELECTION WITH DQC

Data exploration involves not only the instances, or data points, but also the features (coordinates) with which the instances are defined. By performing SVD, and selecting a subset of coordinates, we define superpositions of the original features within which we search for clustering of the instances. In problems with very many features, it is advantageous to also perform some feature filtering employing a judicious selection of subsets of the original features. Clearly, the effectiveness of preprocessing data using some method for selecting important features is well appreciated. What we wish to show in this discussion is how easily one distinguishes the effects of feature filtering in our visual approach and how easy it is, in problems where one has an expert classification, to see if the unsupervised method used to select important features is working well. Furthermore, we wish to show the power of combining iterations of an SVD-based feature filtering algorithm in conjunction with iterations of DQC. To do this we will show what happens when one applies these ideas to the data set of Golub *et al.* [13].

The Golub *et al.* data set contains gene chip measurements on cells from 72 leukemia patients with two different

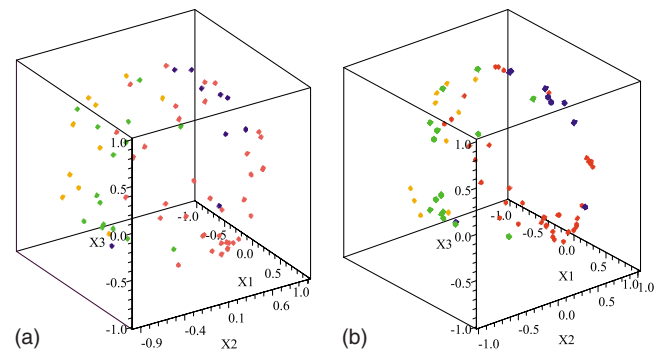


FIG. 5. (Color online) The left-hand picture is the raw data from the Affymetrix Chip plotted for principal components 2, 3, and 4. Clearly, without the coloring it would be hard to identify clusters. The darkest (blue) points are somewhat close to one another but the remaining points are spread out enough so classification is difficult. The right-hand picture is the same data after DQC evolution using $\sigma=0.2$ and a mass $m=0.01$. The different classes are shown as blue (darkest), red (next darkest gray), green (next darkest), and orange (lightest gray).

types of leukemia, acute lymphocytic leukemia (ALL) and acute myeloid leukemia (AML). The expert identification of the classes in this data set is based upon dividing the ALL set into two subsets corresponding to T-cell and B-cell Leukemia. The AML set is divided into patients who underwent treatment and those who did not. In total the Affymetrix GeneChip used in this experiment measured the expression of 7129 genes. The feature filtering method we employ is based on SVD entropy and is a simple modification of a method introduced by Varshavsky *et al.* [14] and applied to the same data.

The method begins by computing the SVD-based entropy [15] of a data set M [matrix of n instances by m features of Eq. (15)] based on the eigenvalues s_j of its diagonal matrix S . Defining normalized relative variance values $v_j = s_j^2 / \sum_k s_k^2$, the data set entropy is defined through

$$E = - \frac{1}{\log r} \sum_{j=1}^r v_j \log(v_j), \quad (17)$$

where r is the rank of the data matrix typically much smaller than m . Given the data set entropy of the matrix M , define the contribution of the i th feature to the entropy using a leave-one-out comparison; i.e., for each feature we construct the quantity

$$CE_i = E(M_{(n \times m)}) - E(M_{[n \times (m-1)]}), \quad (18)$$

where the second entropy is computed for the matrix with the i th feature removed. Our filtering technique will be to remove all features for which $CE_i \leq 0$.

Figure 5 displays the raw data in the three-dimensional space defined by PCs 2 to 4 and the effect that DQC has on these data. In Fig. 6 we see the result of applying feature filtering to the original data, represented in the same three dimensions, followed by DQC evolution. Applying a single

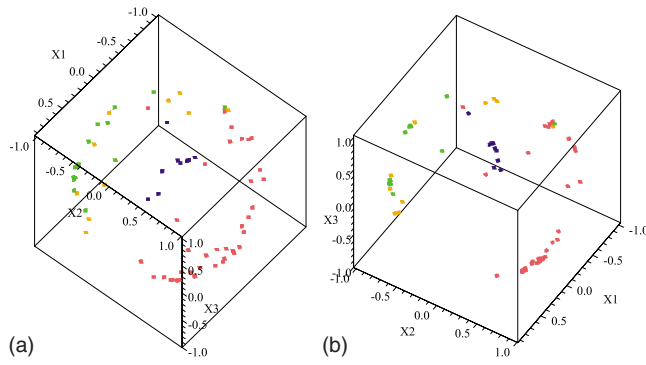


FIG. 6. (Color online) The left-hand plot is the Golub data after one stage of SVD-entropy-based filtering but before DQC evolution. The right-hand plot is the same data after DQC evolution.

stage of filtering has a dramatic effect upon clustering even before DQC evolution. The latter helps sharpening the cluster separation.

Figure 7 shows the results of three iterations of SVD entropy, before and after DQC evolution. These plots, especially the after DQC pictures, show dramatic clustering, especially for the blue points. With each stage of filtering we see that the blue points cluster better and better, in that the single red outlier separates from the cluster and the cluster separates more and more from the other points. The blue points are what we will refer to as an *obviously robust cluster* which has been identified in early stages of filtering. If one continues iterating past the fifth stage, however, the clear separation of the blue (darkest gray) points from the others begins to diminish. Thus, we see that the SVD-entropy-based filtering, in trying to enhance the clumping of the red points, starts throwing away those features which make the blue cluster distinct. Since this effect is quite pronounced we would say that features that are important to distinguishing the blue cluster from the others begin to be removed at the sixth and higher iterations of filtering. This is, of course, just what we are looking for, a way of identifying those features which are important to the existing biological clustering. Out of the original 7129 features, we have reduced ourselves to 2766 features by the fifth iteration. In going from step five to step six this gets further reduced to 2488 features, so we

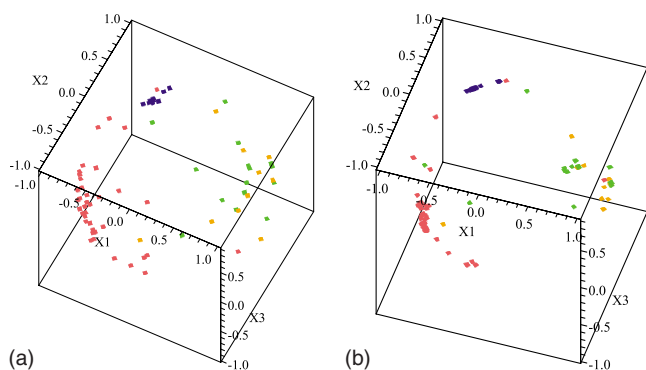


FIG. 7. (Color online) The left-hand plot is the data after three stages of SVD-entropy-based filtering but before DQC evolution. The right-hand plot is the same data after DQC evolution.

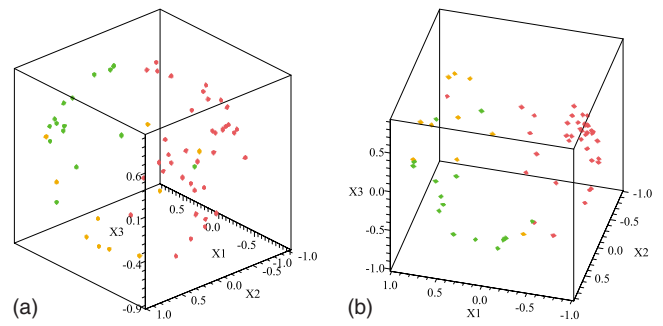


FIG. 8. (Color online) The left-hand plot is what the starting data look like if one first removes the blue points and does one stage of SVD-entropy-based filtering. The right-hand plot is what the data look like after three stages of filtering.

could begin searching among the 278 eliminated features to isolate those most responsible for the separation of the blue cluster from the others. Instead, we will take another track and, since it is so robust and easily identified, remove the blue cluster from the original data and repeat the same process without this cluster. The idea here is that now the SVD-entropy-based filtering will not be pulled by the blue cluster and so it will do a better job of sorting out the red, green, and orange clusters. As we will see, this is in fact the case.

In Fig. 8 we see a plot of what the starting configurations look like if one takes the original data, removes the identified blue cluster, and re-sorts the reduced data set according to the SVD-entropy-based filtering rules. The left-hand plot is what happens if one filters a single time, removing those features, i , whose one-left-out comparison, CE_i , is less than or equal to zero. The right-hand plot shows what happens if one repeats this procedure two more times, each time removing features for which $CE_i \leq 0$. There is no problem seeing that each iteration of the SVD-entropy-based filtering step improves the separation of the starting clusters. By the time we have done five SVD-entropy-based filtering steps the red, green, and orange clusters are distinct, if not obviously separated.

Finally, to complete our discussion, we show Fig. 9. This figure shows the results of doing five iterations of the SVD-entropy-based filtering and following that with three stages of DQC evolution. The dramatic clustering accomplished by DQC evolution makes it easy to extract clusters. Note, however, that in the second plot we see what we have seen throughout, that the red points first form two distinct sub-clusters which only merge after two more stages of DQC evolution. This constant repetition of the same phenomenon, which is only made more apparent by SVD-entropy-based filtering, is certainly a real feature of the data. It presumably says that what appears to be a sample of a single type of cell at the biological level is in reality two somewhat different types of cells when one looks at gene expression. A measure of the success of clustering is given by the Jaccard score [16] which, for this result is 0.762, higher than the value 0.707 obtained by [14].

IX. SUMMARY AND OUTLOOK

We have proposed a dynamical method for exploring proximity relationships among data points in large spaces.

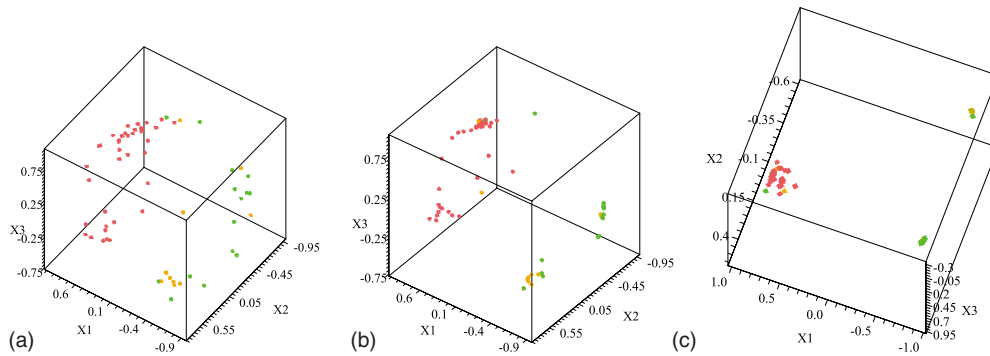


FIG. 9. (Color online) The left-hand plot is what the starting data look like if one first removes the blue points and does five stages of SVD-entropy-based filtering. The right-hand plot is what happens after one stage of DQC evolution. The bottom plot is the final result after iterating the DQC evolution step two more times. At this point the clusters are trivially extracted with the red points clustered in the left-hand side of the plot and the yellow in the lower right-hand corner and the greens in the upper right-hand corner.

Starting with the potential function of quantum clustering [4] we have shown how to embed it into a dynamical theory so as to provide a visual exploratory tool. Formulating the theoretical treatment using coherent (Gaussian) states allows us to derive analytic expressions for all necessary calculations of the temporal evolution. This allows us to treat quite complicated data and put them into a visual framework that can be easily manipulated by the user who wishes to search for structures in the data. We have tested the system on random data to make sure that it does not produce unwarranted clustering structures.

Throughout this paper we represent the DQC evolution of the Gaussians associated with the original data points by following the centers of the evolving wave functions. It should be noted that there is more information to be gained from the full wave function of a data point: it is expected to expand, at times, and cover a large fraction of the domain of the cluster with which this point is associated. It may also tunnel into neighboring clusters with which the point has small dynamic distances. We expect this notion to be particularly useful when the data may be better described in terms of “elongated clusters,” i.e., when cluster cores are not points but lines (e.g., a ring) or higher-dimensional manifolds. Note that our methodology is not confined to potentials that have only well-separated minima.

We have discussed the virtues of combining DQC with some preprocessing tools. The first was SVD, which was used to limit the range of the data values and to allow us to do some dimensional reduction. While dimensional reduction is a must for handling data in very large dimensions, and it helps to remove noise from the data, we wish to point out that DQC can handle a large number of features without much difficulty. The computational complexity of the problem is controlled by the number of data points, since this defines the size of the matrix to be exponentiated. The computational cost associated with keeping more features is only related to computing the matrices associated with multiplying a wave function by a given coordinate. This is a one time cost. The computational cost of computing the values of these operators only grows linearly with the number of features. Clearly it is possible to avoid these costs by keeping a large number of features when constructing the quantum po-

tential, $V(\vec{x})$, and plotting a much smaller number of features when constructing the animations. Experience has shown that after one stage of DQC evolution, clustering which occurs because of structures in $V(\vec{x})$ that are only seen in features that are not plotted in the animations becomes readily visible in those plots that we do construct. This aspect of DQC allows us to avoid some of the problems associated with using SVD to strongly reduce the number of dimensions. In addition to dimensional reduction based upon simply making an SVD decomposition of the data, we discussed one scheme for selecting individual features that are judged to be relevant to the data at hand. Since our problem is unsupervised, we employed a feature filtering method that depends on the contribution of the features to SVD entropy. The examples showed that the visual nature of DQC made it easy to judge the effectiveness of feature filtering especially after iterative applications of DQC evolution.

We have already noted, that for sets of data containing entries with a very large number of features, DQC has the computational advantage that once one has formed the Hamiltonian of the system, the computational problem is carried out using a matrix which has no more rows and columns than the number of data points. Moreover, we have seen that the simplest reduction of the analytic problem of assigning data points to minima of the multidimensional potential function works remarkably well. Going beyond the truncation procedure explained in Appendix B, while easily doable, seems unnecessary for most problems, and this allows us to greatly speed up the computations. In our analysis we went on to discuss the case of data sets containing large numbers of points. It turns out that, using our Hilbert space representation of data points, we can naturally select a small set of points whose Gaussians span efficiently the entire Hilbert space. These Gaussians are then used as the basis for calculating the DQC evolution of all points. It is quite obvious from the example displayed in Fig. 4 how well these properties of DQC can be employed to discern structures in the large data set under consideration.

Finally, we wish to observe that the DQC methods described in this paper can be easily extended to general classification problems that are usually resolved by supervised machine learning methods. The point is that given a training

set, i.e., a data set that has been fully resolved by DQC once the appropriate stages of dimensional reduction and feature filtering has been applied, then one can use this set to classify new data. There are two different ways one can accomplish this task. In the first approach we use the fact that the training set has been successfully clustered to assign distinct colors to points that lie in the training set, so that they may be visually identified in all subsequent studies. Once this has been done, the classification of new data points can be accomplished in two steps. First, reduce the SVD matrix containing both the training set and the new data points (using the previously determined features) to an appropriate dimension and construct the QC potential for the full data set including the training set. Next, apply DQC to study the evolution of the full system using the new QC potential and see how the new points associate themselves with the points in the training set. Note, as always, both the intermediate dynamics and eventual coalescence of the full set into clusters can give useful information about the full data set. The fact that the old points have been colored according to the original classification scheme makes it possible to see if the SVD reduction of the full data set (training set plus new data) distorts the original classification. If this happens, i.e., if the original points fail to cluster properly, then one can go back and use the tools of feature filtering, etc. to analyze what has changed. This sort of visual identification of aspects of the data which distort clustering was already used in the case of the leukemia data set to see that the existence of a *strong* cluster can distort the clustering of the remaining data. Once this easily identified cluster was removed from the data set the clustering of the remaining data was significantly improved.

The second approach, which is necessary if the data set contains many entries and the training set is itself large, is to use only the training set to generate the quantum potential and the exponential of the Hamiltonian. Next, as we already discussed, use this operator to evolve the full data set including the training set. In this case the training set is guaranteed to cluster as before and we can categorize the new points according to how they associate with known clusters in the training data.

ACKNOWLEDGMENT

This work was supported by the U.S. DOE under Contract No. DE-AC02-76SF00515.

APPENDIX A: USEFUL OPERATOR IDENTITIES

Using conventional quantum-mechanical notation we represent the Gaussian wave function by

$$|\sigma\rangle = (\sqrt{\pi\sigma})^{-1/2} e^{-x^2/2\sigma^2}, \quad (\text{A1})$$

where we adopted Dirac's bra and ket notation [17] to denote $|\psi\rangle = \psi(x)$ and $\langle\psi| = \psi(x)^*$. Employing the operators x and $p = \frac{1}{i} \frac{d}{dx}$ obeying the commutation relations $[x, p] = i$, we define the annihilation operator

$$A_\sigma = i \frac{\sigma}{\sqrt{2}} p + \frac{1}{\sigma\sqrt{2}} x \quad (\text{A2})$$

obeying $A_\sigma|\sigma\rangle = 0$. Its Hermitian adjoint creation operator $A_\sigma^\dagger = -i \frac{\sigma}{\sqrt{2}} p + \frac{1}{\sigma\sqrt{2}} x$ obeys $[A_\sigma, A_\sigma^\dagger] = 1$.

We will need a few identities to derive the matrix elements we have to calculate. First we note the normal ordering identity (meaning rewriting by using the operator commutation relations so that A_σ 's appear to the right of all A_σ^\dagger 's):

$$e^{\alpha(A_\sigma^\dagger + A_\sigma)} = e^{\alpha^2/2} e^{\alpha A_\sigma^\dagger} e^{\alpha A_\sigma} \quad (\text{A3})$$

which may be proven by differentiation with respect to α . Next we note that

$$e^{g(\alpha)A_\sigma^\dagger} A_\sigma e^{-g(\alpha)A_\sigma^\dagger} = \sum_n \frac{g(\alpha)^n}{n!} [A_\sigma^\dagger [A_\sigma^\dagger [\dots [A_\sigma^\dagger, A_\sigma] \dots] \dots]_n = A_\sigma - g(\alpha), \quad (\text{A4})$$

which is easily derived by differentiating with respect to g and noting that only the first commutator is nonzero. A similar calculation proves the equally useful result:

$$e^{\alpha(A_\sigma^\dagger - A_\sigma)} = e^{-\alpha^2/2} e^{\alpha A_\sigma^\dagger} e^{-\alpha A_\sigma}. \quad (\text{A5})$$

Now, because the Parzen-window estimator is constructed using Gaussian wave functions centered about points other than $x=0$, it is convenient to have an operator expression which relates the Gaussian centered about $x=0$ to the Gaussian centered about $x=\bar{x}$.

Theorem. $|\sigma, \bar{x}\rangle = e^{-ip\bar{x}}|\sigma\rangle$ is a normalized Gaussian wave function centered at $x=\bar{x}$; i.e.,

$$|\sigma, \bar{x}\rangle = (\sqrt{\pi\sigma})^{-1/2} e^{-(x-\bar{x})^2/2\sigma^2}. \quad (\text{A6})$$

This state is known as a coherent state [18] obeying

$$A_\sigma|\sigma, \bar{x}\rangle = \bar{x}|\sigma, \bar{x}\rangle. \quad (\text{A7})$$

The generalization to Gaussians in any number of dimensions is straightforward, since they are just products of Gaussians defined in each one of the different dimensions.

APPENDIX B: MATRIX ELEMENTS

The states we start out with $|\sigma, \bar{x}_i\rangle$ have norm one and are, in general, linearly independent; however, they are not orthogonal to one another. In what follows we will need an explicit formula for the scalar product of any such Gaussian $|\sigma, \bar{x}_i\rangle$ with another $|\sigma, \bar{x}_j\rangle$. This is easily derived given the operator form for the shifted Gaussian derived in Appendix A. Thus, we find that

$$\langle\sigma, \bar{y}|\sigma, \bar{x}\rangle = \langle\sigma|e^{-ip(\bar{x}-\bar{y})}|\sigma\rangle = e^{-(\bar{x}-\bar{y})^2/4\sigma^2}, \quad (\text{B1})$$

which is needed for computing the matrix of scalar products $N_{ij} = \langle\sigma, \bar{x}_i|\sigma, \bar{x}_j\rangle$. Similarly, by employing $e^{ip\bar{y}}xe^{-ip\bar{y}} = x + \bar{y}$ we find that

$$\langle\sigma, \bar{y}|x|\sigma, \bar{x}\rangle = \frac{(\bar{x} + \bar{y})}{2} e^{-(\bar{x}-\bar{y})^2/4\sigma^2}. \quad (\text{B2})$$

It is straightforward to generalize this derivation to obtain

$$\langle \sigma, \bar{y} | V(x) | \sigma, \bar{x} \rangle = e^{-(\bar{x} - \bar{y})^2/4\sigma^2} \langle \sigma | V \left(x + \frac{(\bar{x} + \bar{y})}{2} \right) | \sigma \rangle \quad (\text{B3})$$

for any function $V(x)$. Note that this expectation value can be evaluated by expanding V in a Taylor series about the point $(\bar{x} + \bar{y})/2$. The leading term is simply $e^{-(\bar{x} - \bar{y})^2/4\sigma^2} V(\frac{\bar{x} + \bar{y}}{2})$ and the remaining terms involving $\langle \sigma | x^n | \sigma \rangle$ can be evaluated from the identity

$$\langle \sigma | e^{\alpha x} | \sigma \rangle = \sum_{n=0}^{\infty} \frac{\alpha^n}{n!} \langle \sigma | x^n | \sigma \rangle = \sum_{p=0}^{\infty} \frac{\alpha^{2p} \sigma^{2p}}{4^p p!}. \quad (\text{B4})$$

To speed up computations we chose to approximate all expectation values of $V(x)$ by $V(\frac{\bar{x} + \bar{y}}{2})$, the first term in this series. One could obviously get a more accurate approximation to the original problem by including additional terms but explicit computation has shown that, for our purposes, this level of accuracy is sufficient.

The final formula we need to derive is that for

$$\langle \sigma, \bar{y} | p^2 | \sigma, \bar{x} \rangle = \langle \sigma | p^2 e^{-ip(\bar{x} - \bar{y})} | \sigma \rangle = \frac{(\bar{x} - \bar{y})^2}{2\sigma^2} e^{-(\bar{x} - \bar{y})^2/4\sigma^2}. \quad (\text{B5})$$

With these preliminaries behind us it only remains to describe the mechanics of the DQC evolution process, where we evaluate the Hamiltonian truncated to an $n \times n$ matrix in the nonorthonormal basis of shifted Gaussians:

$$\mathcal{H}_{i,j} = \langle \sigma, \bar{x}_i | H | \sigma, \bar{x}_j \rangle. \quad (\text{B6})$$

The time evolution of our original states is computed by applying the exponential of the truncated Hamiltonian to the state in question; i.e., $|\sigma, \bar{x}\rangle(t) = e^{-iHt} |\sigma, \bar{x}\rangle$. Computing the exponential of the truncated operator is quite simple except for one subtlety: we have defined \mathcal{H} by its matrix elements between a nonorthonormal set of states. Hence, to perform the exponentiation, we first find the eigenvectors and eigenvalues of the metric N_{ij} and use them to compute the matrix $N_{ij}^{-1/2}$ [19]. Then we construct the transformed \mathcal{H} by

$$\mathcal{H}_{i,j}^{tr} = \sum_{k,l} N_{i,k}^{-1/2} \mathcal{H}_{k,l} N_{l,j}^{-1/2}. \quad (\text{B7})$$

Now we can construct the exponential of this operator by simply finding its eigenvectors and eigenvalues. In order to compute the time evolution of one of the original states we simply write them in terms of the orthonormal basis.

The only step which remains is to explain how we compute the expectation values of the operator x as functions of time: we first construct, for each component, the operator

$$X_{i,j} = \langle \sigma, \bar{x}_i | x | \sigma, \bar{x}_j \rangle \quad (\text{B8})$$

and use $N_{i,j}^{-1/2}$ to put this into the same basis in which we exponentiate \mathcal{H} ; i.e., construct

$$X_{i,j} = \sum_{k,l} N_{i,k}^{-1/2} X_{k,l} N_{l,j}^{-1/2}. \quad (\text{B9})$$

-
- [1] R. R. Coifman, S. Lafon, A. B. Lee, M. Maggioni, B. Nadler, F. Warner, and S. W. Zucker, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 7426 (2005).
- [2] S. Lafon and A. B. Lee, *IEEE Trans. Pattern Anal. Mach. Intell.* **28**, 1393 (2006).
- [3] B. Nadler, S. Lafon, R. Coifman, and I. G. Kevrekidis, *Appl. Comput. Harmon. Anal.* **21**, 113 (2006).
- [4] D. Horn and A. Gottlieb, *Phys. Rev. Lett.* **88**, 018702 (2001).
- [5] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. (Wiley-Interscience, New York, 2001).
- [6] S. J. Roberts, *Pattern Recogn.* **30**, 261 (1997).
- [7] For a single data point $V(\bar{x})$ is just a quadratic harmonic potential centered on the original data point.
- [8] R. Varshavsky, M. Linal, and D. Horn, *COMPACT: A Comparative Package for Clustering Assessment*, Springer Lecture Notes in Computer Science Vol. 3759, edited by G. Chen, Y. Pan, M. Guo, and J. Lu (Springer, Berlin, 2005), pp. 159–167.
- [9] B. D. Ripley, *Pattern Recognition and Neural Networks* (Cambridge University Press, Cambridge, UK, 1996).
- [10] M. B. Wolfe, M. E. Schreiner, B. Rehder, D. Laham, P. W. Foltz, W. Kintsch, and T. K. Landauer, *Discourse Process.* **25**, 309 (1998).
- [11] This is particularly attractive for new multicore PCs and for computer clusters, since it is possible, even in MAPLE, to write multithreaded programs which farm the multiplication of large numbers of vectors out to different processors. This means that one can achieve a great improvement in the speed of the computation for very little additional work.
- [12] We are grateful to JoAnne Hewett and Tom Rizzo for providing us with this example. The points in the plot represent sample supersymmetric models for physics beyond the standard model that satisfy all experimental constraints in a parameter space of 20 dimensions.
- [13] T. R. Golub *et al.*, *Science* **286**, 531 (1999).
- [14] R. Varshavsky, A. Gottlieb, M. Linal, and D. Horn, *Bioinformatics* **22**, e507 (2006).
- [15] O. Alter, P. O. Brown, and D. Botstein, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 10101 (2000).
- [16] The Jaccard score is evaluated by considering all pairs of data points and asking if they cluster together and if they fit in the same class as judged by the expert. The Jaccard score is then defined by $J = \frac{tp}{tp+fp+fn}$, where tp, fp, fn stand for true positive, false positive, and false negative, correspondingly.
- [17] E. Merzbacher, *Quantum Mechanics*, 3rd ed. (Wiley, New York, 1998).
- [18] J. R. Klauder and B.-S. Skagerstam, *Coherent States* (World Scientific, Singapore, 1984).
- [19] If our original set of states is not linearly independent, then $N_{i,j}$ will have some zero eigenvalues. Clearly, we throw their corresponding eigenvectors away when computing $N_{i,j}^{-1/2}$. In practice we discard all vectors whose eigenvalue is smaller than 10^{-5} .