Electricity Price Curve Modeling by Manifold Learning

Jie Chen, *Student Member, IEEE,* Shi-Jie Deng, *Senior Member, IEEE,* and Xiaoming Huo, *Senior Member, IEEE*

***Abstract*—This paper proposes a novel non-parametric approach for the analysis and prediction of electricity price curves by applying the manifold learning methodology. Cluster analysis based on the embedded low-dimensional manifold of the original price data is employed to identify characteristics of the price curve shape. The proposed price curve model performs well in forecasting both short-term price such as the day-ahead prices and longer term price such as the week-ahead prices. The forecast accuracy is demonstrated by numerical results using historical price data taken from the Eastern U.S. electric power markets.**

***Index Terms*—Electricity spot price, locational marginal price, electricity forward curve, forecasting, manifold learning**

I. Introduction

In the competitive electricity wholesale markets, market participants, including power generators and merchants alike, strive to maximize their profits through prudent trading and effective risk management against adverse price movements. A key to the success of market participants is how well they can model the electricity price dynamics and realistically capture its characteristics. One strand of research on modeling electricity price processes focuses on the aspect of derivative pricing and asset valuation which investigates the electricity spot and forward price models in a risk-neutral world (e.g., [1], [2], and [3]). Another research strand concerns the modeling of electricity prices in the physical world which offers price forecasts for assisting with physical trading and operational decision-making. An accurate short term price forecast over a time horizon of hours helps market participants to devise their bidding strategies in the auction based pool-type markets and to allocate generation capacity optimally among different markets. The medium term forecast with a time horizon spanning days to months is useful for balance sheet calculations and risk management applications [4].

In the second research strand of power price modeling, there is an abundant literature on forecasting spot or short-term electricity prices, especially the day-ahead prices ([5], [6], [7], [8], [9], [10], [11], [12]). Typically, the electricity prices are treated as hour-to-hour univariate time series and then modeled by either parametric models including ARIMA processes and their variants ([6], [7] and [8]), regime-switching or hidden Markov processes ([9] and [12]), Levy processes [11], and hybrid price models combining statistical modeling with fundamental supply-demand modeling ([5]); or non-parametric models such as the artificial neural networks ([13], [14], [15]). While spot price modeling is important, successful trading and risk management operations in electricity markets also requires knowledge on an electricity price curve consisting of prices for electricity delivered at a sequence of future times instead of only at the spot. For instance, in order to maximize the market value of generation assets, power generators would need to base their physical trading decisions over how much power to sell in the next day and in the long-term contract markets on both the short-term price forecast for electricity delivered in the next 24 hours and the electricity forward price with maturity ranging from weeks to years. The non-storable nature of electricity makes the electrons delivered at different time points essentially different commodities. Current market price (i.e., spot price) of electricity may have little correlation with that of electricity delivered in a few months from now. Thus, it is imperative to be able to model the electricity price curve as a whole. There is not much literature on modeling electricity price curves. Paper [16] proposes a parametric forward price curve model for the Nordic market, which does not model the movements of the expected future level of a forward curve. A recent paper [17] employs a weighted average of nearest neighbors approach to model and forecast the day-ahead price curve. These works offer little insight on understanding the main drivers of the price curve dynamics. Our paper contributes to this strand of research by proposing a novel non-parametric approach for modeling electricity price curves. Analysis on the intrinsic dimensionality of an electricity price curve is offered, which sheds light on identifying major factors governing the price curve dynamics. The forecast accuracy of our model compares favorably against that of the ARIMA type models in one-day ahead prediction, and is much better in prediction over longer horizons such as one week or one month.

In general, the task of analytically modeling the dynamics of such a price curve is daunting because the curve can be a very high-dimensional subject. Each price point on the curve essentially represents one dimension of uncertainty. To reduce the dimensionality of modeling a price curve and identify the major random factors influencing the curve dynamics, Principle Component Analysis (PCA) is proposed and has been widely applied in the real-world data analysis for industrial practices. As PCA is mainly suited for extracting the linear factors of a data set, it does not appear to perform well in fitting electricity price curves with a linear factor model in

Jie Chen, Shi-Jie Deng (corresponding author) and Xiaoming Huo are with the H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0205. Emails: jchen, deng, xiaoming@isye.gatech.edu.

This work has been partially supported by National Science Foundation grants 0604736 and 0134210.

a low-dimensional space. However, the following intuition suggests that there shall exist a low-dimensional structure capturing the majority of randomness in the electricity price curve dynamics. Take the day-ahead electricity price curve as an example. While electricity delivered in the next 24 hours are different commodities, the corresponding prices all result from equilibrating the fundamental supply and demand for electricity. The common set of demand and supply conditions in all 24 hours hints on a possible non-linear representation of the 24-dimensional price curve in a space of lower dimension. A natural extension to the PCA approach is to consider the manifold learning methods, which are designed to analyze intrinsic non-linear structures and features of high-dimensional price data, and then make short- and medium-term price predictions. Specifically, we propose to extract a low-dimensional structure (i.e., manifold) from the electricity price curves through a manifold learning method. After getting the low-dimensional manifold representation of a price curve, price forecasts are obtained by first predicting on each coordinate of the manifold and then utilizing a reconstruction method to map the forecasts back to the original price space. The conceptual flowchart of our modeling approach is illustrated by Fig. 1. Our main contributions lie in the dashed boxes 1 and 2 in which a proper dimension reduction method is identified and applied to analyze the intrinsic low-dimensional manifold of electricity price curves. The resulting analysis reveals the major drivers of the price curve dynamics and facilitates accurate price forecast. This work also enables the application of standard times series models such as Holt-Winter in the forecast step from box 1 to box 2.

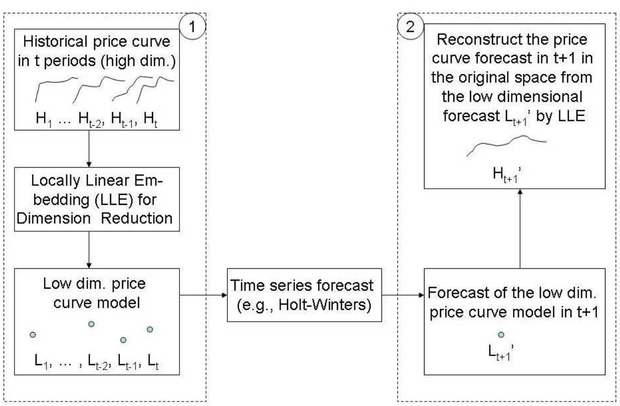

Fig. 1. The conceptual flowchart of the model.

Manifold learning is a new and promising nonparametric dimensionality reduction approach. Many high dimensional data sets that are encountered in real-world applications can be modeled as sets of points lying close to a low dimensional manifold. Given a set of data points $x_1, x_2, \cdots, x_N$, we can assume that they are sampled from the manifold with noise, i.e.,

$$x_i = f(y_i) + \varepsilon_i, i = 1, \cdots, N \tag{1.1}$$

where $x_i \in \mathbb{R}^D$, $y_i \in \mathbb{R}^d$, $d \ll D$ and ε_i's are noise. The manifold based methodology offers a way to find the embedded low dimensional structure $y_i's$ from the high dimensional data x_i's.

Many nonparametric methods are created for nonlinear manifold learning, including multidimensional scaling (MDS) [18], [19], locally linear embedding (LLE) [20], [21], ISOMAP [22], Laplacian eigenmaps [23], Hessian eigenmaps [24], local tangent space alignment (LTSA) [25], and diffusion maps [26]. Survey [27] gives a review on the above methods.

In this paper, locally linear embedding (LLE) is adopted for manifold learning and reconstruction. The study of the intrinsic dimension and embedded manifold indicates that there does exist a low dimension manifold with an intrinsic dimension that is between 3 and 4 for day-ahead electricity price curves in the New York Power Pool (known as NYPP).

Cluster analysis is integrated with manifold learning to identify the characteristics of the price data. For comparison, clustering is done in two cases: (1) with all the coordinates in the original space, and (2) part of the coordinates in the low dimensional space. Cluster analysis based on the low dimensional structure provides clearer view of the seasonal effect in the electricity prices as well as the shape changing of the price curves, relative to the clustering on the original price data.

The rest of the paper is organized as follows. Section II describes a manifold based method—locally linear embedding (LLE)—and the corresponding reconstruction method. In Section III, the LLE and its reconstruction are applied to model and analyze the day-ahead electricity price curves in NYPP. In Section IV, the application of cluster analysis with the low dimension coordinates of the price curves is illustrated. Section V presents the electricity price curve prediction based on manifold learning. Section VI concludes.

II. MANIFOLD LEARNING ALGORITHM

Among various manifold based methods, we find that locally linear embedding (LLE) works well in modeling electricity price data. Our purpose is two folded: to analyze the features of electricity price curves and predict the price curve at a future time. The reconstruction from the low dimensional coordinates to the high dimensional space is a necessary step for forecasting. In the literature, there are only a few reconstruction methods. Through extensive computational experiments, we conclude that LLE is more efficient in reconstruction relative to other methods for our purpose. Moreover, LLE is easy to implement and it is fast. In this section, we introduce the algorithms of LLE and its reconstruction.

A. Locally Linear Embedding(LLE)

Given a set of data points $x_1, x_2, \cdots, x_N$, with high dimension D. We are looking for the embedded low dimensional features $y_1, y_2, \cdots, y_N \in \mathbb{R}^d$. The LLE method is a nonparametric method that works as follows [20], [21], [27]:

1) Identify the k nearest neighbors based on Euclidean distance for each data point $x_i, 1 < i < N$. Let N_i denote the indices of the k nearest neighbors of the vector x_i.

2) Find the optimal local convex combinations of the k nearest neighbors to represent each original vector. That is, the following objective function (2.2) is minimized and the weights $w'_{ij}s$ of the convex combinations are calculated.

$$E(w) = \sum_i |x_i - \sum_{j \in N_i} w_{ij} x_j|^2. \quad (2.2)$$

3) Find the representations in the low dimensional space, such that the above local convex representations are best preserved. This is done by choosing the d dimensional coordinates $y_i, 1 < i < N$ that minimize the following objective function:

$$\Phi(y) = \sum_i |y_i - \sum_{j \in N_i} w_{ij} y_j|^2. \quad (2.3)$$

It can be shown that with some additional conditions, which make the problem well defined, solving the above minimization problem(2.3) is equivalent to solving an eigenvector problem with a sparse $N \times N$ matrix.

LLE does not impose any probabilistic model on the data; However, it implicitly assumes the convexity of the manifold. It can be seen later that this assumption is satisfied by the electricity price data.

B. LLE Reconstruction

Given a new data point in the embedded low dimensional space, the reconstruction method is used to find its counterpart in the high dimensional space based on the training data set, which is the price data in our setting. Reconstruction accuracy is critical for the application of manifold learning in prediction. There are some but limited number of methods for the reconstruction from the embedded space to the data space. For specific linear manifold, the reconstruction can be easily made with PCA. For nonlinear manifold, LLE reconstruction, which is in the same manner as LLE, is introduced in [20]. LTSA reconstruction, a interpolation-like reconstruction, and nonparametric regression reconstruction are introduced in [25]. Among all these reconstruction methods, LLE reconstruction in general has the best performance on the electricity data. This is a important reason for us to choose LLE in the present paper.

LLE reconstruction is derived in a similar manner as LLE. There are three steps. Suppose embedded manifold coordinates $y_1, y_2, \cdots, y_N$, have been obtained through LLE in the previous subsection. Denote the new data point as y_0 in the low dimensional space. LLE reconstruction is applied to find the approximation $\hat{x}_0$ of the original vector x_0 in the high dimensional space based on $x_1 \cdots, x_N$ and $y_1, \cdots, y_N$. The three steps are:

1) Identify the k nearest neighbors of y_0 among $y_1, \cdots, y_N$. Let N_0 denote the set of the indices of the k nearest neighbors of the vector y_0.
2) The weights of the convex combination are calculated by solving

$$E(w) = |y_0 - \sum_{j \in N_0} w_j y_j|^2. \quad (2.4)$$

3) Point $\hat{x}_0$ is reconstructed by $\hat{x}_0 = \sum_{j \in N_0} w_j x_j$.

Remark: Solving optimization problems (2.3) and (2.4) is equivalent to solving a linear system of equations. When there are more neighbors than high dimensions or low dimensions, i.e., $k > D$ or $k > d$. The coefficient matrix associated with the system of linear equations is singular, which means that the solution is not unique. In [20], this issue is resolved by adding the identity matrix multiplied with a small constant to the coefficient matrix. We follow this approach.

Suppose $x_0^{(j)}, 1 \leq j \leq D$, is the jth component of vector x_0. The reconstruction error (RE) of x_0 is defined as

$$RE(x_0) = \sum_{j=1}^{D} \frac{|x_0^{(j)} - \hat{x}_0^{(j)}|}{x_0^{(j)}} \quad (2.5)$$

The reconstruction error of the entire training data set (TRE)[1] is defined as

$$TRE = \sum_{i=1}^{N} \sum_{j=1}^{D} \frac{|x_i^{(j)} - \hat{x}_i^{(j)}|}{x_i^{(j)}} \quad (2.6)$$

by regarding each y_i as a new data point y_0.

III. Manifold Learning of Electricity Price Data

The day-ahead locational based marginal prices (LBMPs) of electricity in the Capital Zone operated by the New York Independent System Operator (NYISO) are collected from Nov 18, 1999 to Feb 12, 2006 for our analysis. The data are available online ($www.nyiso.com/public/market\_data/pricing\_data.jsp$). The price curve consisting of hourly prices on each day is considered as an observation. As there are 24 hourly electricity prices in one day in NYPP, the dimension D of the high dimensional space is 24. Because the embedded manifold structure might vary with time, the data set for manifold learning should not be too large. On the other hand, it should also contain enough data points to guarantee the accuracy of the manifold learning. Therefore, the size of training data set for manifold learning is set to be two years. In this section, two-year data from Feb 6, 2003 to Feb 5, 2005 are used as an illustration of manifold based methodology. Fig. 2 plots the hour-to-hour day-ahead LBMPs in dollars per MegaWatt-hour ($/MWh) during this period.

A. Log Transform and Outlier Preprocessing

We take the logarithmic (log) transform of the electricity prices. There are several advantages to deal with the log prices. First, the electricity prices are well known for the non-constant variance, and log transformation could make the prices less volatile. Second, the embedded manifold is more uniformly distributed in the low dimensional space. Third, the reconstruction error of the entire training data set (TRE) can be reduced by taking log transform.

Outlier preprocessing is also a necessary step for manifold learning. Outliers in the paper are defined as the extreme electricity price spikes that are quite different from the ones

[1] When TRE is calculated, y_i itself is not included in its k nearest neighbors.

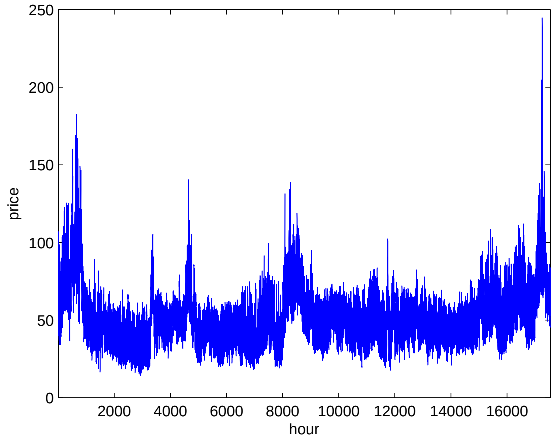

Fig. 2. Hour-to-hour day-ahead LBMPs in Capital Zone from Feb 6, 2003 to Feb 5, 2005 in NYISO.

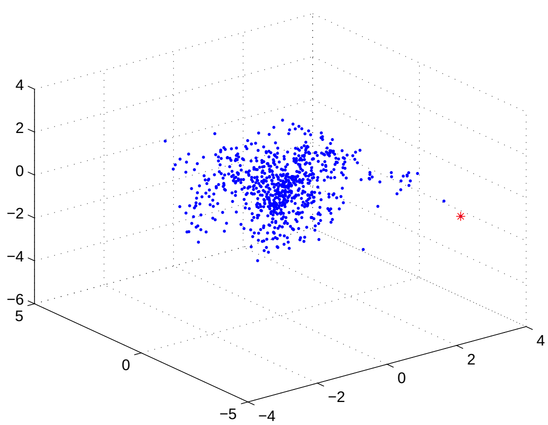

Fig. 3. Embedded 3-D manifold without any outlier preprocessing(but with log transform and LLP smoothing). "$*$" indicates the outlying day—Jan 24, 2005

in the neighborhood. In Fig. 2, there is an extreme spike on the right of the plot that belongs to Jan 24, 2005. In the low dimensional manifold, the days with outliers can also be detected by the points that stand far away from the other points. Fig. 3 shows that the point corresponding to Jan 24, 2005 lies out of the main cloud of the points on the embedded three-dimension manifold. Thus, we regard Jan 24, 2005 as an outlying day. To deal with an outlier, we replace the prices in the outlying day with the average of prices in the days that are right before and right after. We preprocess the outliers because the embedded low dimensional manifold is supposed to extract the trendy feature of the entire data set. Extreme price spikes are considered as individual and local features. Outlier preprocessing is able to improve the efficiency of manifold learning. Fig. 4 shows that the low dimensional manifold after removing the outliers is more uniform. Moreover, we find that the outliers, as very special phenomena in the past, do not severely affect the prediction of the near-term regular prices in the forecasting application. This observation illustrates the robustness of our model.

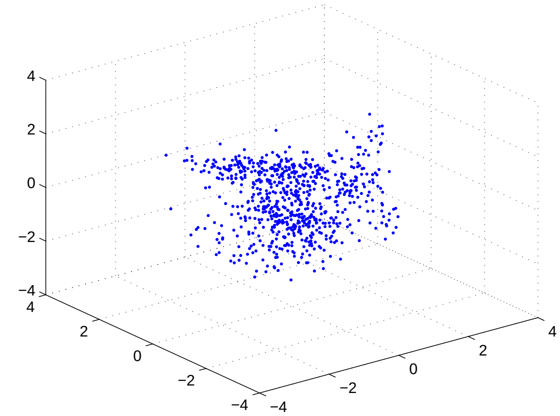

Fig. 4. Embedded 3-D manifold with log transform, outlier preprocessing and LLP smoothing.

B. LLP Smoothing

In our model, the noise in (1.1) can contaminate the learning of the embedded manifold and the estimation of the intrinsic dimension. Therefore, locally linear projection (LLP) [28], [29], [27] is recommended to smooth the manifold and reduce the noise. In the first step, neighbor observations are identified. In the second step, singular value decomposition (SVD) or principal component analysis (PCA) is used to estimate the local linear subspace. Finally, the observation is projected into this subspace. The detailed description of the algorithm is given as follows:

ALGORITHM: LLP

For each observation $y_i, i = 1, 2, \cdots, N$,

1) Find the k-nearest neighbors of y_i. The neighbors are denoted by $\widetilde{y}_1, \widetilde{y}_2, \cdots, \widetilde{y}_k$.
2) Use PCA or SVD to identify the linear subspace that contains most of the information in the vectors $\widetilde{y}_1, \widetilde{y}_2, \cdots, \widetilde{y}_k$. Suppose the linear subspace is A_i. Let k_0 denote the assumed dimension of the embedded manifold. Then subspace A_i can be viewed as a linear subspace spanned by the vectors associated with the largest k_0 singular values.
3) Project y_i into the linear subspace A_i and let $\hat{y}_i, i = 1, \cdots, N$, denote the projections.

After denoising, the efficiency of manifold learning is improved, and the reconstruction error of the entire training data set (TRE) is reduced in general. For $d = 3$ and the illustrated train data set, TRE is 4.32% after LLP smoothing, compared to 6.02% without LLP smoothing. The choice of the two parameters in LLP, the dimension of the linear space and the number of the nearest neighbors will be discussed in detail in following subsections.

C. Estimation of Intrinsic Dimension

Intrinsic dimension (i.e., the dimension of the low dimensional space d) is an important parameter of manifold learning.

Papers [30] and [31] provide several ways of estimating the intrinsic dimension. In [30], the maximum likelihood estimator of the intrinsic dimension is established, and Matlab code is available online. In [31], the intrinsic dimension is estimated based on a nearest neighbor algorithm. Without LLP smoothing, the two methods show that the intrinsic dimension is between 4 and 5[2]. Thus, it is reasonable to set the dimension of the linear space as 4 in LLP smoothing. After LLP smoothing, the intrinsic dimension is reduced to a value between 3 and 4. The numerical experiments seem to indicate that LLP smoothing can not only denoise, but also improve the efficiency of estimating the intrinsic dimension. In this section, only manifold learning with $d = 3$ is illustrated; manifold learning with $d = 4$ will be in the same manner.

D. Manifold Learning by LLE

We can set the intrinsic dimension d as 3 or 4 and apply manifold based method LLE to the denoised data $\hat{y}_i, i = 1, \cdots, N$, that are obtained after LLP smoothing. For LLP smoothing, LLE, and LLE reconstruction, the number of nearest neighbors is required. A common number k is chosen. We set k to be 23 for all the numerical study. If k is too small, the reconstruction error of the entire training data set (TRE) increases. The reconstruction error, however, is not sensitive to k, and the similar reconstruction error is produced with k around 23. Fig. 4 provides the plot of the embedded three-dimension manifold. As the low dimensional manifold is nearly convex and uniformly distributed, LLE is an appropriate manifold based method.

Fig. 5 provides the plots of coordinates of the embedded three-D manifold. There are some interesting (and perhaps coincidental) interpretations for each coordinate in the low dimensional space. The first coordinate represents the mean of the prices in a day, and this can be seen via the high correlation coefficient (0.9964) with the mean. The second coordinate is highly correlated with the standard deviation of the prices in a day with a correlation coefficient 0.7073. This also means that second coordinate contains some other information besides standard deviation. For each statistic of mean, standard deviation, range[3], skewness and kurtosis of the prices in a day, Table. I presents the coordinates with the highest correlation coefficient with the statistics among the three dimensions, and the corresponding correlation coefficients. The table shows that the second coordinate is also correlated with range and skewness, but not significantly. The third coordinate shows both weekly and yearly seasonality. Weekly seasonality is well known for electricity price, and it can be easily detected by spectral analysis. Yearly seasonality may be caused by shape changing of the price curves over the year, which will be seen more clearly in the application of clustering in the next section.

[2]The estimation of the intrinsic dimension varies slightly with the choices of the number of the nearest neighbors. Therefore, a range of the estimates of the intrinsic dimension is given.

[3]Range is the difference of the maximum price and the minimum price in a day.

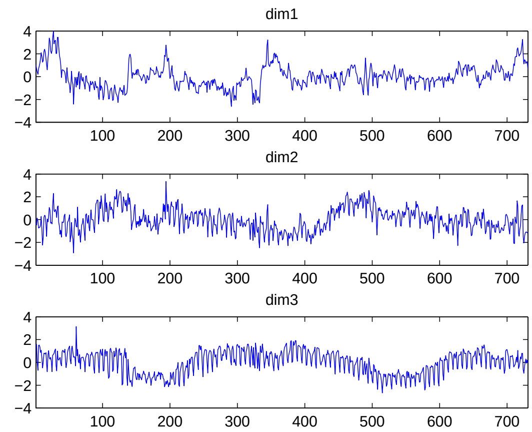

Fig. 5. Coordinates of the embedded 3-D manifold.

TABLE I

COORDINATES THAT HAVE THE MAXIMUM CORRELATION COEFFICIENTS WITH THE STATISTICS (OUT OF MEAN, STANDARD DEVIATION, RANGE, SKEWNESS AND KURTOSIS) AMONG 3 DIMENSIONS IN THE REDUCED-DIMENSION SPACE.

	Mean	Standard Deviation	Range	Skewness	Kurtosis
Coordinate	1st	2nd	2nd	2nd	3rd
Correlation Coefficient	0.9964	0.7073	0.5141	0.5646	0.2612

IV. CLUSTER ANALYSIS

We apply cluster analysis to the price curves. Cluster analysis [32] (also known as data segmentation) groups or segments a collection of objects into subsets (i.e., clusters), such that those within each cluster are more closely related to each other than those assigned to different clusters.

The K-means clustering algorithm is one of the mostly used iterative clustering methods. Assume that there are K clusters, the algorithm begins with a guess of the K cluster centers. Then, the algorithm iterates between the following two steps until convergence. The first step is to identify the closest cluster center for each data point based on the Euclidean distance. The second step is to replace each cluster center with the coordinate-wise average of all the data points that are closest to it.

For the electricity price data, we apply K-means clustering to the low dimensional coordinates that are obtained after applying manifold learning. In K-means algorithm, the number of clusters is supposed to be known. In [33], an information theoretic approach is provided to find the number of clusters in a data set. The R code of the algorithm is available online. For all the low dimensional coordinates illustrated, the number of clusters is estimated to be 6. After clustering, the cluster mean of price curves is plotted in Fig. 6. The clusters are distinguished at price levels and price curve shapes.

A significant advantage of applying cluster analysis to the low dimensional coordinates rather than the original price data is that the features of the price curves can be separated,

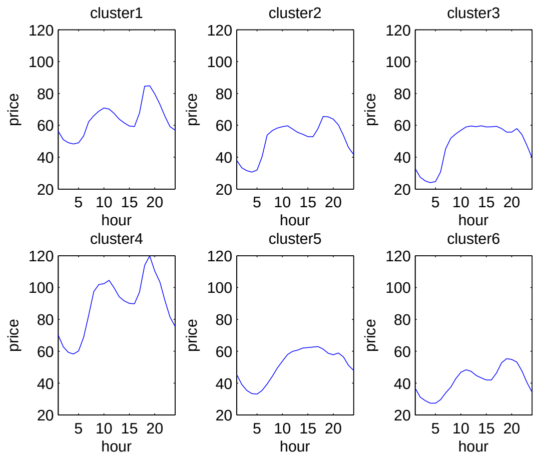

Fig. 6. The mean of the price curves in each cluster. The clustering is based on all the low dimensional coordinates.

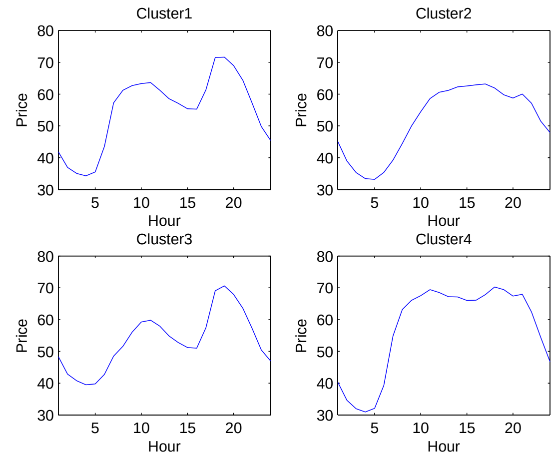

Fig. 7. The mean of the price curves in each cluster. The clustering is based on the 2nd and 3rd coordinates in the low-D space.

and then clustered respectively. Each coordinate in the low dimensional space extracts some main features of the price curves. As mentioned, the first coordinate always represents the mean of the prices in a day, and it is orthogonal to other coordinates, as LLE outputs a null-space [27]. We remove the first coordinate and cluster the remaining coordinates in the low dimensional space. The clustering focuses on the other features, e.g., the shape of the price curves, rather than the level of the price curves. According to the method in [33], the number of clusters is estimated to be equal to 4. Fig. 7 provides the mean of the price curves in each cluster when the second and third coordinates are used. The shape of price curves varies over the clusters. Fig. 8 provides the distribution of the clusters in the same scenario. It is obvious that there is a strong yearly seasonality in the shape of price curves. It seems that the shape of price curves is unimodal in the summer while it is bimodal in the winter.

For the low-dimensional manifold with $d = 4$, the number of clusters is 7 with all the coordinates clustered. When all the coordinates except the first one are clustered, the number of clusters is reduced to 5. The clusters have the similar properties as those in the case of $d = 3$.

Fig. 8 and Fig. 9 plot the mean of the price curves in each cluster and the distribution of the clusters respectively, when the original electricity price curves are clustered. The number of clusters is estimated to be 4 by using the method in [33]. The clusters differentiate in various features, e.g., the price level and the price curve shape level. We also notice that the distribution of the clusters for the original price data is very unbalanced. There are only 28 data points in the second cluster (corresponding to the cluster label '3' in Fig. 8), which is 3.8% of the two year data. For the clustering with low dimensional coordinates, there is no such phenomena. Furthermore, seasonal effect can not be clearly observed in the distribution of the clusters that are from the original price data.

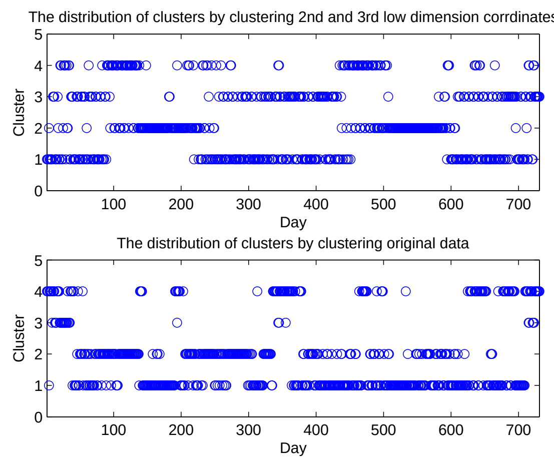

Fig. 8. Distribution of clusters.

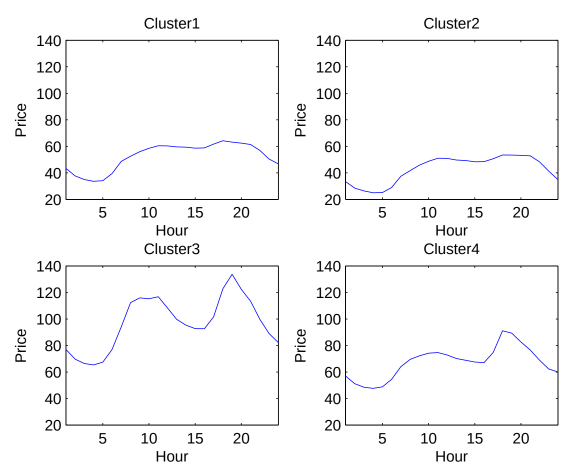

Fig. 9. The mean of the price curves in each cluster. The clustering is based on original price data.

V. PREDICTION

Prediction of future electricity price curves is an important issue in the electricity price market, accurate prediction enables market participants to increase their profit by trading energy and hedge the potential risk successfully. However, it is hard to make accurate prediction due to the high dimensionality of the price curve. Although there are various methods for multivariate time series modeling and forecasting, e.g., vector ARIMA, they are always applied to low dimensional multivariate data. With high dimensional data, e.g., electricity price curves with 24 dimensions, the model often finally turns out to be too complicated to be useful. Another intuitive way of prediction is to treat the electricity prices at the same hour over the days as a univariate time series, and then apply the univariate time series forecasting to 24 series. This approach ignores both the correlation among the prices at different hours in a day and the integrity of the price curve.

A large amount of existing forecasting methods focus on one-day ahead price prediction, i.e., the horizon of prediction is one day (24 hours). The longer horizon prediction from days to weeks has not drawn much attention, however it also plays an important role in biding strategy and risk management. Our method of prediction based on the embedded low dimensional manifold provides a good tool in forecasting both one-day ahead price curve and price curves in a longer horizon, e.g., one week and one month. In ARIMA and its related methods, the time series model needs to be identified by people, and it may vary for different modelers. Our method, however, is nonparametric and robust, and it does not require any model assumption. Furthermore, the size of the training data for ARIMA and its related model is often restricted within half a year, while two-year data is applied in the proposed method for prediction.

A. Prediction Method

In our method, we make the prediction in the low dimensional space, and then re-map them back to the original high dimensional space. There are three step:

1) Learn the low dimensional manifold of electricity price curves using the LLE method.
2) Predict each low dimensional coordinate via univariate time series forecasting, e.g., Holt-Winters Algorithm, Structural models, etc. For one-day ahead price curve prediction, only one data point needs to be forecasted in each dimension; for one-week ahead price curve prediction, seven data points need to be forecasted in each dimension.
3) Reconstruct the prediction in low dimensional space to the original high dimensional space using the LLE reconstruction.

The first and third step have been described in the previous sections. In the second step, we make the prediction via univariate time series forecasting rather than multivariate time series forecasting, because the low dimensional coordinates are orthogonal to each other. There are a variety of methods of univariate time series forecasting. In this paper, we choose Holt-Winters algorithm and the structural model, as they do not require model identification and are easy to implement (there are existing functions in statistical software R). For details regarding the algorithms of Holt-Winters and the structural model, we refer to [34].

B. The Definition of Weekly Average Prediction Error

To assess the predictive accuracy of our methodology, weekly average prediction error is computed in this paper.

1) For one-week ahead prediction, weekly average prediction error (WPE$_w$) is defined in the same way as in [8] [17]. Let $\hat{x}_i$ and x_i denote the one-day ahead prediction for the ith day in a week and the corresponding true price curve, respectively. $\|\cdot\|_1$ is the L_1 norm of a vector, which is the sum of the absolute values of all the components in the vector. We define

$$\bar{x} = \frac{1}{7\times 24}\sum_{i=1}^{7}\|x_i\|_1$$

and

$$\text{WPE}_d = \frac{1}{7\times 24}\sum_{i=1}^{7}\frac{\|x_{(i)} - \hat{x}_{(i)}\|_1}{\bar{x}}$$

2) Since the weekly average prediction error of one-week ahead prediction has not been defined, we use the following definition. Let i denote the position of the first predicted day in one-week prediction horizon, i.e., the training data set is two-year observations before day i. Let $\hat{x}_{(i,j)}$ denote the jth predicted daily prices with day i as the starting day of this predicted week, where $1 \le j \le 7$. When i increases by one, the prediction window (with length of seven days) is moved forward by one. Let $x_{(i,j)}$ be the corresponding observed electricity prices. Similar to WPE$_d$, we define

$$\bar{x}_{(i)} = \frac{1}{7\times 24}\sum_{j=1}^{7}\|x_{(i,j)}\|_1$$

and

$$\text{WPE}_w = \frac{1}{7\times 7\times 24}\sum_{i=1}^{7}\sum_{j=1}^{7}\frac{\|x_{(i,j)} - \hat{x}_{(i,j)}\|_1}{\bar{x}_{(i)}}$$

3) For one-month ahead prediction, we use the following definition which is similar to WPE$_w$. Notation $\hat{x}_{(i,j)}$ and $x_{(i,j)}$ are defined in the same way as in WPE$_w$ except that the prediction horizon is four weeks (approximately one month) rather than one week, i.e. $1 \le j \le 28$. Similar to WPE$_w$, we define

$$\bar{x}_{(i)} = \frac{1}{28\times 24}\sum_{j=1}^{28}\|x_{(i,j)}\|_1$$

and

$$\text{WPE}_m = \frac{1}{7\times 28\times 24}\sum_{i=1}^{7}\sum_{j=1}^{28}\frac{\|x_{(i,j)} - \hat{x}_{(i,j)}\|_1}{\bar{x}_{(i)}}$$

We define σ_d, σ_w and σ_m as the standard deviations of WPE$_d$, WPE$_w$ and WPE$_m$, respectively.

C. Prediction of electricity data

In our numerical experiments, we forecast the one-day ahead electricity price curve for each day of the second week from February 2005 to January 2006. The training data set of both manifold learning and prediction is the prices in 731 days (two years) prior to the day of the week when prices are to be predicted. Simultaneously, for each training data set, the prediction of one-week ahead electricity price curves is made. In the previous manifold learning example, the embedded manifold was trained with the data from Feb 6, 2003 to Feb 5, 2005, which is two years prior to the first day of the second week in February 2005. The one-day ahead price prediction for Feb 6, and the one-week ahead price prediction for the week from Feb 6 to Feb 12 are made in the experiment.

For each training data set, the same parameter set illustrated in the previous section is used. The number of nearest neighbors and the dimension of the linear space in LLP smoothing are set to be 23 and 4 respectively. Only one outlying day—Jan 24, 2005—is identified. For the second step of prediction, Holt-Winters algorithm is executed with starting period equal to 7 days and 14 days respectively. This choice is due to the weekly effect of the electricity prices. In the output tables, HW7 and HW14 stand for Holt-Winter algorithm with starting period equal to 7 days and 14 days respectively, and STR stands for the structural model.

Table II and III provide the weekly average prediction errors of one-day ahead price prediction for the 12 weeks and their standard deviations. The naive prediction of a certain week is given by the actual prices of the previous week. Our prediction method outperforms the naive method for both $d = 3$ and $d = 4$. Compared with the the ARIMA model, our prediction method has competitive prediction accuracy. The prediction errors of the proposed method are slightly smaller than those of ARIMA for $d = 4$, however slightly larger for $d = 3$. Therefore, the experiment seems to indicates that four is a more appropriate intrinsic dimension to make the prediction for the electricity prices. The three prediction methods (HW7, HW4 and STR) are applied in low dimensional space and generate similar performance; the structural model is slightly more accurate.

In Table IV and V, the weekly average prediction errors of one-week ahead price prediction for the 12 weeks and their standard deviations are presented. Our proposed method outperforms both the ARIMA model and the naive method with either $d = 3$ or $d = 4$. The ARIMA model even acts worse than the naive method for one-week ahead prediction. For the electricity price data, the ARIMA model is a very complicated model with multiple seasonality. Although the model often fits the data very well, it might cause the overfitting problem in the prediction, especially for longer horizon. For one-week ahead prediction, the ARIMA model needs to predict 168 data points, while the proposed method only needs to predict 7 data points for each coordinate. The three prediction methods in low dimensional space (HW7, HW4 and STR) generate similar performances, with structural model being slightly more accurate again.

The proposed method can be applied to forecast prices in

TABLE II

COMPARISON OF $\mathrm{WPE}_d(\%)$ OF ONE-DAY AHEAD PREDICTION FOR 12 WEEKS.

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	ARIMA	Naive
1	7.94	8.21	7.86	7.15	7.02	7.37	8.14	15.84
2	7.05	6.64	6.48	6.24	5.82	5.45	5.58	10.06
3	6.68	7.11	7.25	6.55	7.10	6.58	6.11	12.39
4	6.75	6.34	6.50	6.25	5.85	6.06	7.28	5.83
5	10.26	9.77	9.94	9.67	9.24	9.72	9.67	31.78
6	7.81	8.58	7.61	7.83	8.53	7.61	7.48	17.49
7	5.34	5.47	5.49	5.15	5.16	5.42	5.98	13.02
8	7.13	8.29	7.63	6.88	8.07	7.48	7.19	14.67
9	6.05	6.07	6.27	6.08	6.02	6.38	6.37	9.68
10	7.85	7.73	7.54	6.68	6.62	6.10	5.87	18.74
11	8.80	9.12	8.99	8.46	8.64	8.79	8.52	27.86
12	9.25	9.26	8.71	8.62	8.70	8.25	10.50	15.42
Mean	7.58	7.72	7.52	7.13	7.23	7.10	7.39	16.07

TABLE III

COMPARISON OF $\sigma_d(\%)$ OF ONE-DAY AHEAD PREDICTION FOR 12 WEEKS.

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	ARIMA	Naive
1	2.75	3.00	2.90	3.27	3.66	3.36	5.06	4.07
2	2.24	2.73	2.45	1.81	2.12	1.85	1.92	6.33
3	2.69	3.10	2.33	2.31	2.85	2.23	1.87	3.43
4	3.20	2.83	2.97	3.27	2.83	2.98	3.45	1.08
5	3.10	3.68	3.96	2.65	3.12	3.91	4.08	8.07
6	3.90	5.23	4.15	3.99	5.25	4.19	3.86	7.68
7	1.61	1.52	2.17	1.63	1.51	2.30	2.89	4.17
8	3.30	3.33	2.78	3.36	3.39	2.82	4.37	8.84
9	2.06	1.76	2.05	2.06	1.87	2.04	2.12	4.85
10	2.07	1.87	2.13	2.44	2.28	2.54	2.23	6.77
11	3.27	3.41	3.07	3.84	4.00	3.45	3.20	12.08
12	4.46	4.32	3.46	4.68	4.46	3.60	5.98	4.22
Mean	2.89	3.07	2.87	2.94	3.11	2.94	3.42	5.96

TABLE IV

COMPARISON OF $\mathrm{WPE}_w(\%)$ OF ONE-WEEK AHEAD PREDICTION FOR 12 WEEKS

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	ARIMA	Naive
1	8.91	9.09	8.28	8.31	8.33	7.65	14.57	12.19
2	11.28	11.01	10.67	10.85	10.59	10.19	13.28	12.33
3	11.30	11.05	10.92	11.03	10.88	10.53	10.68	14.59
4	7.96	7.54	7.77	7.55	7.15	7.36	14.21	6.71
5	15.91	15.58	14.35	15.58	15.23	13.88	21.64	26.68
6	11.06	10.21	10.53	10.85	10.04	10.41	14.63	14.44
7	6.94	7.22	6.59	6.82	7.09	6.42	9.49	10.28
8	7.68	8.11	7.47	7.46	7.95	7.31	10.36	13.17
9	9.68	9.51	10.03	9.78	9.60	10.11	11.84	11.57
10	10.01	10.06	9.70	9.45	9.44	8.99	11.24	15.18
11	13.52	13.64	13.73	12.94	13.09	13.30	21.78	23.94
12	14.35	14.46	13.68	13.95	14.08	13.28	26.01	11.52
Mean	10.72	10.62	10.31	10.38	10.29	9.95	14.98	14.38

TABLE V
COMPARISON OF $\sigma_w(\%)$ OF ONE-WEEK AHEAD PREDICTION FOR 12 WEEKS

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	ARIMA	Naive
1	1.45	1.82	1.21	1.32	1.66	1.14	7.06	2.35
2	3.72	3.66	3.67	3.80	3.73	3.86	7.06	1.18
3	3.31	3.61	3.54	3.12	3.48	3.41	3.50	2.47
4	3.27	2.99	3.36	3.45	3.13	3.50	6.28	0.88
5	6.07	6.29	6.29	5.91	6.10	6.46	10.30	4.75
6	3.14	3.42	3.52	3.10	3.44	3.43	5.51	1.50
7	1.00	0.99	1.28	1.02	0.95	1.40	3.77	1.48
8	3.07	2.86	2.27	3.19	2.96	2.33	4.88	2.28
9	2.01	1.67	2.00	2.06	1.73	2.01	5.50	1.34
10	2.17	2.30	2.42	2.17	2.42	2.30	3.92	3.05
11	4.47	3.96	4.88	4.88	4.41	5.17	13.77	4.65
12	3.47	3.63	2.43	3.25	3.37	2.15	11.60	2.66
Mean	3.10	3.10	3.07	3.11	3.12	3.10	6.93	2.38

a longer horizon than one week, e.g., two weeks or even one month. As there are only a few methods associated with one-month ahead price prediction, we apply three naive methods to compare with. The first naive method takes the last month prices in the training data set as its prediction. The second method repeats the last week prices four times, and the third one replicates the prices of last two weeks twice, respectively, as the prediction. Table VI and VII provide the weekly average prediction errors of the one-month ahead price prediction for 12 weeks and their standard deviations. The notations—naive1, naive3 and naive3—stand for the three naive methods. From the comparison, the proposed method outperforms all the naive methods in terms of the weekly average prediction errors. We notice that the stand deviations of the structural model are larger than those of naive models, and it is mainly due to an inaccurate prediction for one day in week five. Thus, Holt-Winters algorithm has the best performance among all the methods for one-month ahead price prediction.

In summary, the proposed method outperforms both the ARIMA model and the naive method for one-day ahead and one-week ahead price prediction when the intrinsic dimension is chosen appropriately. Compared with three naive methods, the proposed method is consistently better.

VI. CONCLUSION

We apply manifold-based dimensionality reduction to electricity price curve modeling. Locally Linear Embedding is demonstrated to be an efficient method for extracting the intrinsic low-dimension structure of electricity price curves. Using price data taken from the NYISO, we find that there exists a low-dimensional manifold representation of the day-ahead price curve in NYPP and specifically, the dimension of the manifold is between 3 and 4. The interpretation of each dimension in terms of basic statistics of day-ahead prices is also given. Numerical experiments show that:

1) Such a nonlinear dimensionality reduction approach enhances the clustering of daily price curves, and consequently leads to better understanding of the intrinsic trend in the electricity prices.
2) The prediction made in the reduced dimensional space is consistently better than that of other approaches such as naive methods, or the ARIMA models. In addition, our method facilitates medium-term prediction, which is difficult, even infeasible for other methods.

Our method does not require any model identification, unlike classical time series approaches such as ARIMA. The robustness and nonparametric feature of the proposed method make it appealing.

TABLE VI
COMPARISON OF $\text{WPE}_m(\%)$ OF ONE-MONTH AHEAD PREDICTION FOR 12 WEEKS

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	Naive1	Naive2	Naive3
1	9.2	9.7	9.6	8.9	9.3	9.4	32.1	12.1	25.7
2	12.9	12.4	12.5	12.5	12.1	12.2	17.2	13.7	14.4
3	17.8	17.5	16.8	17.6	17.3	16.5	15.0	16.0	14.2
4	12.3	12.1	12.2	12.0	11.8	11.9	13.6	11.5	12.0
5	17.1	17.1	25.9	16.7	16.7	25.5	24.4	23.9	27.5
6	16.0	15.6	16.7	15.9	15.6	16.7	17.0	13.6	13.9
7	10.8	10.8	11.4	10.7	10.7	11.4	13.0	12.8	15.2
8	14.7	14.4	13.7	14.5	14.2	13.7	19.5	15.9	18.9
9	18.1	18.3	19.9	18.2	18.4	20.0	19.7	17.1	19.7
10	14.7	14.6	15.3	14.0	13.8	14.6	31.9	18.4	28.5
11	19.3	19.2	19.7	19.0	18.9	19.5	30.0	26.2	27.6
12	13.4	13.4	12.5	13.1	13.1	12.0	31.8	14.3	15.0
mean	14.7	14.6	15.5	14.4	14.3	15.3	22.1	16.3	19.4

TABLE VII
COMPARISON OF $\sigma_m(\%)$ OF ONE-MONTH AHEAD PREDICTION FOR 12 WEEKS

Month	d=3 HW7	d=3 HW14	d=3 STR	d=4 HW7	d=4 HW14	d=4 STR	Naive1	Naive2	Naive3
1	1.24	1.70	2.65	1.13	1.55	2.63	0.17	2.79	6.82
2	3.58	3.36	4.01	3.64	3.33	4.11	0.47	0.46	1.01
3	3.02	1.84	3.20	3.02	2.06	3.49	0.46	2.25	1.08
4	1.30	1.23	1.34	1.22	1.18	1.22	0.40	1.38	0.65
5	4.63	4.32	26.01	4.59	4.28	25.41	0.36	5.96	0.62
6	3.82	4.47	3.57	3.95	4.50	3.58	1.15	2.90	0.89
7	0.32	0.35	0.99	0.35	0.36	1.03	0.62	1.24	0.58
8	3.13	3.72	3.03	3.22	3.82	2.82	0.92	0.50	0.46
9	4.13	3.97	6.79	4.20	4.01	6.80	0.80	2.77	1.95
10	1.26	1.30	1.93	1.09	1.21	1.62	1.51	3.47	3.77
11	2.26	2.08	3.12	2.40	2.21	3.29	0.54	4.33	1.82
12	6.21	6.46	3.62	6.04	6.25	3.44	4.16	1.83	1.76
mean	2.91	2.90	5.02	2.90	2.90	4.95	0.96	2.49	1.78

REFERENCES

[1] B. Johnson and G. Barz, *Energy Modelling and the Management of Uncertainty*. Risk Books, 1999, ch. Selecting Stochastic Processes for Modeling Electricity Prices, London.

[2] S. J. Deng, "Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes," *UCEI POWER Working Paper P-073*, 2000.

[3] J. J. Lucia and E. S. Schwartz, “Electricity prices and power derivatives: Evidence from the nordic power exchange,” *Review of Derivatives Research*, vol. 5, no. 1, pp. 5–50, 2002.

[4] A. Misiorek, S. Trueck, and R. Weron, “Point and interval forecasting of spot electricity prices: Linear vs. non-linear time series models,” *Studies in Nonlinear Dynamics and Econometrics*, vol. 10, no. 3, 2006, article 2.

[5] M. Davison, L. Anderson, B. Marcus, and K. Anderson, “Development of a hybrid model for electricity spot prices,” *IEEE Transactions on Power Systems*, vol. 17, no. 2, pp. 257–264, 2002.

[6] F. J. Nogales, J. Contreras, A. J. Conejo, and R. Espínola, “Forecast next-day electricity prices by time series models,” *IEEE Transactions on Power Systems*, vol. 17, no. 2, pp. 342–348, 2002.

[7] J. Contreras, R. Espínola, F. J. Nogales, and A. J. Conejo, “ARIMA models to predict next-day electricity prices,” *IEEE Transactions on Power Systems*, vol. 18, no. 3, pp. 1014–1020, 2003.

[8] A. J. Conejo, M. A. Plazas, R. Espínola, and A. B. Molina, “Day-ahead electricity price forecasting using the wavelet transform and ARIMA models,” *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 1035–1042, 2005.

[9] A. M. Gonzalez, A. M. S. Roque, and J. G. Gonzalez, “Modeling and forecasting electricity prices with input/output hidden markov models,” *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 13–24, 2005.

[10] C. Knittel and M. Roberts, “Empirical examination of deregulated electricity prices,” *Energy Economics*, vol. 27, no. 5, pp. 791–817, 2005.

[11] S. J. Deng and W. J. Jiang, “Levy process driven mean-reverting electricity price model: a marginal distribution analysis,” *Decision Support Systems*, vol. 40, no. 3-4, pp. 483–494, 2005.

[12] T. D. Mount, Y. Ning, and X. Cai, “Predicting price spikes in electricity markets using a regime-switching model with time-varying parameters,” *Energy Economics*, vol. 28, no. 1, pp. 62–80, 2006.

[13] B. Ramsay and A. J. Wang, “A neural network based estimator for electricity spot-pricing with particular reference to weekend and public holidays,” *Neurocomputing*, no. 47-57, 1998.

[14] B. R. Szkuta, L. A. Sanabria, and T. S. Dillon, “Electricity price short-term forecasting using artificial neural networks,” *IEEE Transactions on Power Systems*, vol. 14, no. 3, pp. 851–857, 1999.

[15] L. Zhang, P. B. Luh, and K. Kasiviswanathan, “Energy clearing price prediction and confidence interval estimation with cascaded neural networks,” *IEEE Transactions on Power Systems*, vol. 18, no. 1, pp. 99–105, 2003.

[16] N. Audet, P. Heiskanen, J. Keppo, and I. Vehvilainen, *Modelling Prices in Competitive Electricity Markets*. London: John Wiley & Sons, 2004, ch. Modeling Electricity Forward Curve Dynamics in the Nordic Market.

[17] A. T. Lora, J. M. R. Santos, A. G. Expósito, J. L. M. Ramos, and J. C. R. Santos, “Electricity market price forecasting based on weighted nearest neighbors techniques,” Working Paper, University of Sevilla, Spain, 2006.

[18] I. Borg and P. Groenen, *Modern Multidimensional Scaling: Theory and Applications*. New York: Springer-Verlag, 1997.

[19] J. B. Kruskal, “Multidimensal scaling by optimizing goodness of fit to a nonmetric hypothesis,” *Psychometrika*, vol. 29, pp. 1–27, 1964.

[20] L. K. Saul and S. T. Roweis, “Think globally, fit locally: unsupervised learning of low dimensional manifolds,” *Journal of Machine Learning Research*, vol. 4, pp. 119–155, 2003.

[21] L. K. Saul and S. T.Roweis, “Nonlinear dimensionality reduction by locally linear embedding,” *Science*, vol. 290, pp. 2323–2326, 2000.

[22] J. B. Tenenbaum, V. de Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science,*, vol. 290, pp. 2319–2323, 2000.

[23] M. Belkin and P. Niyogi, “Laplacian eigenmaps for dimensionality reduction and data representation,” *Neural Computation*, vol. 15, no. 6, pp. 1373–1396, June 2003.

[24] D. L. Donoho and C. Grimes, “Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data,” *Proceedings of the National Academy of Sciences*, vol. 100, pp. 5591–5596, 2003.

[25] Z. Zhang and H. Zha, “Principal manifolds and nonlinear dimension reduction via tangent space alignment,” *SIAM Journal of Scientific Computing*, vol. 26, no. 1, pp. 313–338, 2004.

[26] B. Nadler, S. Lafon, R. R. Coifman, and I. G. Kevrekidis, “Diffusion maps, spectral clustering and reaction coordinates of dynamical systems,” *Applied and Computational Harmonic Analysis: Special issue on Diffusion Maps and Wavelets*, vol. 21, pp. 113–127, July 2006.

[27] X. Huo, X. Ni, and A. K. Smith, *Mining of Enterprise Data*. Springer, 2005, new york Ch. A survey of manifold-based learning methods, invited book chapter, to appear, also available at http://www2.isye.gatech.edu/statistics/papers/06-10.pdf.

[28] X. Huo and J. Chen, “Local linear projection (LLP),” in *First IEEE Workshop on Genomic Signal Processing and Statistics (GENSIPS)*, Raleigh, NC, October 2002, http://www.gensips.gatech.edu/proceedings/.

[29] X. Huo, “A geodesic distance and local smoothing based clustering algorithm to utilize embedded geometric structures in high dimensional noisy data,” in *SIAM International Conference on Data Mining, Workshop on Clustering High Dimensional Data and its Applications*, San Francisco, CA, May 2003.

[30] E. Levina and P. J. Bickel, “Maximum likelihood estimation of intrinsic dimension,” in *Advances in Neural Information Processing Systems 17 (NIPS2004)*. MIT Press, 2005.

[31] P. Verveer and R. Duin, “An evaluation of intrinsic dimensionality estimators,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 1, pp. 81–86, 1995.

[32] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning*. Springer, 2001.

[33] C. Sugar and G. James, “Finding the number of clusters in a data set: An information theoretic approach,” *Journal of the American Statistical Association*, vol. 98, no. 750-763, 2003.

[34] P. J. Brockwell, *Introduction to Time Series and Forecasting*, 2nd ed. Springer, 2003.

PLACE PHOTO HERE

Jie Chen (S’05) received the B.S. degree in computational math from Nanjing University, China, in 2003. She is currently working towards the Ph.D. degree at the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta. Her research interests are applied statistics and data mining.

PLACE PHOTO HERE

Dr. Shi-Jie Deng (SM ’06) is Associate Professor of Industrial and Systems Engineering at Georgia Institute of Technology. Dr. Deng’s research interests include financial asset pricing and real options valuation, financial engineering applications in energy commodity markets, transmission pricing in electric power systems, stochastic modelling and simulation. He received the CAREER Award from the National Science Foundation in 2002. Dr. Deng has served as a consultant to several private and public organizations on issues of risk management and asset valuation in the restructured electricity industry. Dr. Deng holds a B.Sc degree in Applied Mathematics from Peking University in China, a M.Sc. in Mathematics from the University of Minnesota at Twin Cities, and M.S. and Ph.D in Industrial Engineering and Operations Research (IEOR) from the University of California at Berkeley.

PLACE PHOTO HERE

Dr. Xiaoming Huo (S’96–M’99–SM’04) received the B.S. degree in mathematics from the University of Science and Technology, China, in 1993 and the M.S. degree in electrical engineering and the Ph.D. degree in statistics from Stanford University, Stanford, CA, in 1997 and 1999, respectively. He is an Associate Professor with the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta. His research interests include dimension reduction and nonlinear methods. Dr. Huo received first prize in the Thirtieth International Mathematical Olympiad (IMO), which was held in Braunschweig, Germany.