

UCLA

UCLA Previously Published Works

Title

Prediction of high-dimensional states subject to respiratory motion: a manifold learning approach

Permalink

<https://escholarship.org/uc/item/64n7p9cm>

Journal

Physics in Medicine and Biology, 61(13)

ISSN

0031-9155

Authors

Liu, Wenyang
Sawant, Amit
Ruan, Dan

Publication Date

2016-07-07

DOI

10.1088/0031-9155/61/13/4989

Peer reviewed



Published in final edited form as:

Phys Med Biol. 2016 July 7; 61(13): 4989–4999. doi:10.1088/0031-9155/61/13/4989.

Prediction of high-dimensional states subject to respiratory motion: a manifold learning approach

Wenyang Liu¹, Amit Sawant^{3,4}, and Dan Ruan^{1,2}

¹Department of Bioengineering, University of California, Los Angeles ²Department of Radiation Oncology, University of California, Los Angeles ³Department of Radiation Oncology, University of Texas Southwestern ⁴Department of Radiation Oncology, University of Maryland

Abstract

The development of high-dimensional imaging systems in image-guided radiotherapy provides important pathways to the ultimate goal of real-time full volumetric motion monitoring. Effective motion management during radiation treatment usually requires prediction to account for system latency and extra signal/image processing time. It is challenging to predict high-dimensional respiratory motion due to the complexity of the motion pattern combined with the curse of dimensionality. Linear dimension reduction methods such as PCA have been used to construct a linear subspace from the high-dimensional data, followed by efficient predictions on the lower-dimensional subspace. In this study, we extend such rationale to a more general manifold and propose a framework for high-dimensional motion prediction with manifold learning, which allows to learn more descriptive features compared to linear methods with comparable dimensions. Specifically, a kernel PCA is used to construct a proper low-dimensional feature manifold, where accurate and efficient prediction can be performed. A fixed-point iterative pre-image estimation method is used to recover the predicted value in the original state space. We evaluated and compared the proposed method with PCA-based approach on level-set surfaces reconstructed from point clouds captured by a 3D photogrammetry system. The prediction accuracy was evaluated in terms of root-mean-squared-error (RMSE). Our proposed method achieved consistent higher prediction accuracy (sub-millimeter) for both 200ms and 600ms lookahead lengths compared to the PCA-based approach, and the performance gain was statistically significant.

1. Introduction

Respiratory motion can cause significant discrepancy between planned and delivered dose distribution in radiotherapy. Conventional motion management methods often detect, estimate, and predict a motion characterization with low degrees of freedom. Even image-based techniques translate their intensity input to the positional monitoring of ROI's rigid motion. However, physiological motion is complex and the anatomical variation is beyond the description of a few fiducial points or their equivalent. A full volumetric description is conceptually possible given the developments of imaging systems of increased capability, ranging from volumetric MRI and CBCT [10], projection monitoring (kV, MV) [25], to

photogrammetry surface imaging [22, 11, 12]. Accurate and efficient extraction of such rich information requires sophisticated processing. In addition, effective motion management during radiation treatment often requires prediction to account for system latency and extra signal/image processing time. It is challenging to predict high-dimensional respiratory motion because the complexity of performing respiratory motion estimation and prediction grows exponentially with the dimensionality of the state space, an issue well known as “curse of dimensionality”. Fortunately, it is conceivable that a full characterization of respiratory motion lies in a smooth manifold of much lower dimension than the extrinsic volumetric imaging space. Linear dimension reduction methods such as PCA have been used to perform prediction in a subspace of reduced dimension [20]. As nonlinear extensions, manifold learning methods have also been utilized to extract respiratory gating navigators from MR/ultrasound images and to reconstruct dynamic 3D image volumes [26, 2]. There are two major benefits of manifold learning: compared with linear feature subspaces that are of similar dimensions, nonlinear approaches are able to provide more descriptive features; with the same descriptive power, nonlinear approaches require lower effective dimensions. In this study, we propose a framework for high-dimensional motion prediction with manifold learning. Specifically, we utilize kernel PCA to learn and construct a low-dimensional feature subspace, where efficient and accurate prediction can be performed. The predicted value in the feature subspace is subsequently mapped back to the original state space via an iterative preimage estimation algorithm. Fig. 1 illustrates the schematic of the proposed method.

2. Methods

2.1. Manifold learning – kernel PCA

The basic idea of manifold learning is to first map the high-dimensional data to a feature space and a locally linear subspace is subsequently constructed from that feature space through some nonlinear mapping. We utilize kernel PCA in this study to learn and construct the low-dimensional feature subspace, but the rationale generalizes to other manifold learning methods. Given a set of training data $\Omega = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, for each $\mathbf{x}_j \in \mathcal{R}^d$ †, we define a kernel function $\kappa: \Omega \times \Omega \rightarrow \mathcal{R}$, such that there exists a mapping to a Hilbert space $\varphi: \Omega \rightarrow \mathcal{H}$:

$$\kappa(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle. \quad (1)$$

Similar to PCA, the eigen decomposition is performed on the centered kernel matrix:

$$H K H = U \Lambda U^T, \quad (2)$$

†Unless otherwise stated, we use bold letters to represent vectors, which are by default column. We use capital letters to represent matrices.

where K is the kernel matrix with each entry $K_{ij} = \kappa(\mathbf{x}_i, \mathbf{x}_j)$, $H = I - \frac{1}{n} \mathbf{1}\mathbf{1}^T$ is the centering matrix, $U = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_n]$ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$. The k th orthonormal eigenvector of V of the covariance matrix in the feature space can be represented as:

$$V_k = \sum_{i=1}^n \frac{\alpha_{k,i}}{\sqrt{\lambda_k}} \tilde{\phi}(\mathbf{x}_i), \quad (3)$$

where $\alpha_{k,i}$ is the i th element of the k th eigenvector \mathbf{a}_k , and $\tilde{\phi}(\mathbf{x}_i)$ is the centered representation in the feature space $\tilde{\phi}(\mathbf{x}_i) = \phi(\mathbf{x}_i) - \bar{\phi}$, with $\bar{\phi} = \frac{1}{n} \sum_{i=1}^n \phi(\mathbf{x}_i)$.

2.2. Prediction in the feature subspace

Let $\mathbf{x}(t) \in \mathcal{R}^d$ denote the high-dimensional state at time t in the original space, upon determining the dimensionality m of the feature subspace, the low-dimensional representation $\tilde{\mathbf{x}}_k(t)$ onto the k th eigenvector is calculated by taking the inner product:

$$\tilde{\mathbf{x}}_k(t) = \langle \tilde{\phi}(\mathbf{x}), V_k \rangle = \frac{1}{\sqrt{\lambda_k}} \sum_{i=1}^n \alpha_{k,i} \tilde{\kappa}(\mathbf{x}, \mathbf{x}_i), \quad (4)$$

where V_k represents the eigenvector corresponding to the k th largest eigenvalue, and $\tilde{\kappa}$ represents the normalized kernel defined as:

$$\tilde{\kappa}(\mathbf{x}, \mathbf{x}_i) = \kappa(\mathbf{x}, \mathbf{x}_i) - \frac{1}{n} \mathbf{1}^T \kappa_{\mathbf{x}} - \frac{1}{n} \mathbf{1}^T \kappa_{\mathbf{x}_i} + \frac{1}{n^2} \mathbf{1}^T \kappa \mathbf{1}. \quad (5)$$

We perform prediction in the low-dimensional (m -D) feature subspace for $\tilde{\mathbf{x}}$ instead of the original high-dimensional (d -D) state space for \mathbf{x} . In the reduced space, the aim is to predict the signal state τ time ahead response $\tilde{\mathbf{y}}_t = \tilde{\mathbf{x}}(t + \tau)$, given a length of mp covariate variable $\check{\mathbf{s}}_t = [\mathbf{x}(t - (p - 1))^T, \mathbf{x}(t - (p - 2))^T, \dots, \mathbf{x}(t)^T]^T \in \mathcal{R}^{mp}$ of p states. Various algorithms are available for low-dimensional prediction such as: linear models, neural networks [16], and kernel density estimation based method [19]. As the specific prediction module is not the focus of this paper, we utilize the vector autoregressive (VAR) model for demonstration purposes:

$$f(\check{\mathbf{s}}_t) = A\check{\mathbf{s}}_t + \mathbf{b}, \quad (6)$$

where $A = [A_1, A_2, \dots, A_p]$ is the coefficient matrix with $A_i \in \mathcal{R}^{m \times m}$, and $\mathbf{b} \in \mathcal{R}^m$ is the bias coefficient vector. Given a series of T covariate-response pairs $\{(\check{\mathbf{s}}_i, \mathbf{y}_i) \mid i \in [1, T]\}$, the coefficient matrix can be obtained by minimizing the sum of square residuals:

$$\min_B \|Y - BS\|^2, \quad (7)$$

where $B = [\mathbf{b}, A_1, A_2, \dots, A_p] \in \mathcal{R}^{m \times (mp+1)}$, $Y = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T] \in \mathcal{R}^{m \times T}$, and

$$S = \begin{bmatrix} 1 & 1 & \dots & 1 \\ \check{s}_1 & \check{s}_2 & \dots & \check{s}_T \end{bmatrix} \in \mathcal{R}^{(mp+1) \times T}. \quad \text{Eq. 7 has a closed-form solution: } B = \hat{Y} S^T (SS^T)^{-1}.$$

Given a testing covariate $\check{\mathbf{s}}$, we can estimate its response as $\check{\mathbf{y}} = \hat{B} \begin{bmatrix} 1 \\ \check{\mathbf{s}} \end{bmatrix}$.

2.3. Pre-image estimation

With the predicted response in the low-dimensional space $\check{\mathbf{y}} = [\check{y}_1, \check{y}_2, \dots, \check{y}_m] \in \mathcal{R}^m$, its projection onto the feature subspace can be represented as:

$$\sum_{k=1}^m \check{y}_k V_k + \bar{\phi} = \sum_{k=1}^m \check{y}_k \sum_{i=1}^n \frac{\alpha_{k,i}}{\sqrt{\lambda_k}} (\phi(\mathbf{x}_i) - \frac{1}{n} \sum_{i=1}^n \phi(\mathbf{x}_i)) + \frac{1}{n} \sum_{i=1}^n \phi(\mathbf{x}_i) = \sum_{i=1}^n \tilde{\gamma}_i \phi(\mathbf{x}_i), \quad (8)$$

where $\gamma_i = \sum_{k=1}^m \check{y}_k \frac{\alpha_{k,i}}{\sqrt{\lambda_k}}$, and $\tilde{\gamma}_i = \gamma_i + \frac{1}{n} (1 - \sum_{j=1}^n \gamma_j)$. For visualization/denoising purpose, we would also like to recover the counter part of this projection in the original high-dimensional state space. However, exact recovery typically does not exist [15], and one usually seeks an approximate solution, which is known as the pre-image estimation problem. One straightforward approach is to directly recover the pre-image \mathbf{x}^* by minimizing the squared distance between $\phi(\mathbf{x}^*)$ and the projection $\sum_{i=1}^n \tilde{\gamma}_i \phi(\mathbf{x}_i)$ in the feature space [15]:

$$\operatorname{argmin}_{\mathbf{x}^*} \left\| \phi(\mathbf{x}^*) - \sum_{i=1}^n \tilde{\gamma}_i \phi(\mathbf{x}_i) \right\|^2 = \operatorname{argmin}_{\mathbf{x}^*} \left\| \phi(\mathbf{x}^*) \right\|^2 - 2 \sum_{i=1}^n \tilde{\gamma}_i \kappa(\mathbf{x}^*, \mathbf{x}_i) + \eta, \quad (9)$$

where η includes terms independent of \mathbf{x}^* . Eq. 9 is equivalent to:

$$\operatorname{argmin}_{\mathbf{x}^*} \sum_{i=1}^n \tilde{\gamma}_i \kappa(\mathbf{x}^*, \mathbf{x}_i). \quad (10)$$

For kernels that have analytical forms, such as Gaussian or polynomial, Eq. 10 can be solved by fixed point iterations [15]. With a Gaussian kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2/c)$, we solve \mathbf{x}^* by the following fixed-point iterations:

$$\hat{\mathbf{x}}_{t+1}^* = \frac{\sum_{i=1}^n \tilde{\gamma}_i \exp(-\|\hat{\mathbf{x}}_t^* - \mathbf{x}_i\|^2/c) \mathbf{x}_i}{\sum_{i=1}^n \tilde{\gamma}_i \exp(-\|\hat{\mathbf{x}}_t^* - \mathbf{x}_i\|^2/c)}. \quad (11)$$

The selection of parameter c in the kernel function is important to properly construct the manifold. Specifically, c needs to be large enough to capture the local structure of the manifold, while being small enough to differentiate between neighborhoods. In practice, setting c comparable to the magnitude of distance between different point pairs serves as a good choice. In this study, we set $c = 2500$ in all our experiment, and we find our algorithm robust to different values of c , as long as its value is within a proper range of training pair distances.

Algorithm 1 illustrates details of the proposed prediction framework.

ALGORITHM 1 Prediction with manifold learning

Step 1: Perform eigen decomposition on kernel matrix and define projection matrix

- $HKH = U\Lambda U^T$, with $U = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_m]$ and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_m)$
- Define projection matrix $P = [V_1, V_2, \dots, V_m]$, for each $V_k = \sum_{i=1}^n \frac{\alpha_{k,i}}{\sqrt{\lambda_k}} \tilde{\phi}(\mathbf{x}_i)$

Step 2: Project training and testing data onto the constructed sub-manifold

- $\tilde{\mathbf{x}}_i = P^T \phi(\mathbf{x}_i)$

Step 3: Calculate VAR coefficient matrix given T covariate-response pairs $\{(\mathbf{s}_i, \mathbf{y}_i) | i \in [1, T]\}$, for each $\check{\mathbf{s}}_i = [\tilde{\mathbf{x}}_{i-(p-1)}^T, \dots, \tilde{\mathbf{x}}_i^T]^T$ and $\check{\mathbf{y}}_i = \mathbf{y}_{i+\tau}$

- Solve VAR coefficient matrix $B = \hat{Y}S^T(SS^T)^{-1}$

Step 4: Estimate testing response given testing covariates $\check{\mathbf{s}}$

- $\check{\mathbf{y}} = \hat{B} \begin{bmatrix} 1 \\ \check{\mathbf{s}} \end{bmatrix}$

Step 5: Estimate pre-image through fixed-point iterations

- $$\hat{\mathbf{x}}_{t+1}^* = \frac{\sum_{i=1}^n \tilde{\gamma}_i \exp(-\|\hat{\mathbf{x}}_t^* - \mathbf{x}_i\|^2/c) \mathbf{x}_i}{\sum_{i=1}^n \tilde{\gamma}_i \exp(-\|\hat{\mathbf{x}}_t^* - \mathbf{x}_i\|^2/c)},$$

Where $\gamma_i = \sum_{k=1}^m \check{y}_k \frac{\alpha_{i,k}}{\sqrt{\lambda_k}}$, and $\tilde{\gamma}_i = \gamma_i + \frac{1}{n} (1 - \sum_{j=1}^n \gamma_j)$
-

3. Experiments and results

3.1. Data description and experimental setup

We evaluated the proposed method on 200 level-set surfaces reconstructed from 200 point clouds. The point clouds were acquired at 15Hz by the VisionRT system from one patient during radiotherapy and spanned about 14 seconds. Fig. 2 presents the example level set

surface reconstructed from one acquired point cloud using the variational method [11]. We used the first 100 surfaces as the training set for manifold learning and training the VAR prediction algorithm. The rest of the 100 surfaces were used as the testing set.

We evaluated the proposed method at two different lookahead lengths: 200ms and 600ms, accounting for different system response time under various radiotherapy configurations [8, 17]. The parameter c of the kernel function was set to $c = 2500$. We compared the prediction accuracy of our method with PCA-based prediction method by varying the dimension m (from 1 to 3) of their corresponding feature subspaces. The covariate variable of both methods was of mp dimensions, with $p = 20$ states in our experiment. For the convenience of performance evaluation, each level set surface is represented as a height function $z(i, j)$ on a resampled rectangular grid of size 200×150 . The prediction accuracy was evaluated w.r.t. root-mean-squared-error (RMSE) and Variance by comparing the predicted surface with the ground truth surface. Specifically, on a size of $I = 200, J = 150$ rectangular grid, we

calculated RMSE and variance as: $\text{RMSE} = \sqrt{\frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J e_{i,j}^2}$ and

$\text{Variance} = \frac{1}{IJ} \sum_{i=1}^I \sum_{j=1}^J \left(e_{i,j} - \frac{\sum_{i=1}^I \sum_{j=1}^J e_{i,j}}{IJ} \right)^2$, where $e_{i,j}$ is the point-wise error at grid point (i, j) : $e_{i,j} = z_{\text{pred}}(i, j) - z_{\text{true}}(i, j)$, with z_{pred} and z_{true} being the height functions of the predicted and the ground truth level-set surfaces, respectively.

3.2. Manifold learning and prediction results

Fig. 3 illustrates the first-dimensional embeddings of the training data and the corresponding prediction results: the blue solid curve represents the first-dimensional embeddings of 100 training surfaces learned from kernel PCA, where a clear periodic pattern can be observed; the green solid curve represents the first-dimensional embeddings of 100 training surfaces learned from PCA, which is noisier (less representative) than that from the kernel PCA; the red dashed curve represents the first-dimensional projection of 100 testing surfaces in the feature space; the black dash-dot curve represents the estimated prediction with 200ms lookahead length by VAR.

3.3. Prediction error analysis

Fig. 4 illustrates and compares one example estimated surface with the ground-truth surface in the chest region. The RMSE and variance histograms from both our method and PCA-based method on 100 testing surface are compared in Fig. 5, where both the RMSE and variance from the proposed method show a clear concentration on the lower value side. The error statistics of both methods are reported in Table 1, where the proposed method achieves consistent lower RMSE for both 200ms and 600ms lookahead lengths compared to PCA-based prediction results. As resulting RMSE from both methods failed the Kolmogorov-Smirnov normality tests, non-parametrical paired Mann-Whitney U-tests were conducted. The p-values in Table. 2 reveals the statistical significance of the performance gain from our method.

3.4. Prediction on simulated irregular breathing pattern

We further evaluated the performance of the proposed method on irregular breathing patterns. As the key factor that affects the performance in our case is the prediction module utilized in the reduced feature space, we evaluated the performance of the VAR model on a set of simulated respiratory patterns by warping the first-dimensional trajectory learned from kernel PCA. Specifically, given a discrete observation x_j obtained at t_j , we warped the temporal axis by $\tilde{t} = t^\eta$, where we set $\eta = 0.7$. The function $g(t) = f(\tilde{t})$ provides a continuous warping, where discrete samples of the waveform g can be obtained with proper interpolation. The prediction result on the simulated patterns is illustrated in Fig. 6, where we observed certain performance degradation due to the lack of adaptivity of the VAR model. Quantitatively, the mean absolute prediction error in the feature space increased from 0.13 to 0.14, compared to the performance on regular breathing patterns. Paired Mann-Whitney U-test yielded a p-value of 0.15, which indicated lack of evidence for statistical significant difference in performance between those two settings. We are in the process of integrating the proposed approach with more sophisticated alternative prediction modules in the post dimension-reduction space. We chose not to evaluate the surface reconstruction error from the predicted irregular pattern, as the ground truth surfaces are not available and can only be interpolated among original high-dimensional surfaces, which is subject to large errors and may mislead the conclusion.

4. Discussion and conclusion

We have proposed a prediction framework for high-dimensional states subject to respiratory motion. The learning approach is particularly suitable to manage respiratory motion in image-guided radiotherapy, with quasi-periodic breathing pattern. The high-dimensional data is mapped to a low-dimensional feature subspace by kernel PCA, followed by efficient VAR prediction in the constructed low-dimensional space. An iterative fixed-point pre-image estimation method is utilized to map the predicted projection in the feature subspace to its original high-dimensional state space. We evaluated the proposed method on predicting level-set surfaces reconstructed from point clouds captured by a photogrammetry system during radiotherapy. Our method outperformed PCA-based prediction and achieved sub-millimeter RMSE accuracy for both 200ms and 600ms lookahead lengths. Paired Mann-Whitney U-tests further demonstrated the statistical significance of the superiority of our method.

Despite the non-convex nature of the iterative fixed-point pre-image estimation method and its reported numerical instability [9, 6], we observed robust performance of this method under various initialization conditions in our application. This could be partially explained by the relatively clear patterns of the respiratory surface images, as opposed to more complicated/obscure patterns from other applications such as hand-written numbers and human faces. There exists alternative and often more complex pre-image estimation methods, e.g., multi-dimensional scaling based method [9] and conformal map approach [6], which yield closed-form solutions and have been shown to achieve more accurate/clean pre-image estimation results [7]. However, such comparison is application-dependent and we are actively evaluating their performance in our application.

In the current work, we chose kernel PCA to learn the nonlinear embeddings from the high-dimensional data. In general, other manifold learning techniques, e.g., Locally Linear Embedding (LLE) [18], Isomap [23], Laplacian Eigenmaps (LE) [3] are also applicable and can be incorporated into our proposed prediction framework. However, unlike kernel PCA, the nonlinear mappings of LLE, Isomap and LE are only defined over the training set, which do not extrapolate well to out-of-samples [1]. In other words, to incorporate new sample into the manifold, one may need to re-compute the whole manifold. Out-of-sample extension for LLE, Isomap and LE is not only challenging but also subject to large errors if recursive schemes are used. On the other hand, the reconstruction from kernel PCA is straightforward and stable, using linear combinations. Given the complexity and lack of agreed kernel extension methods for the other nonlinear methods [4], we decided to omit their implementation to avoid misleading conclusions in this study.

When applying the proposed method to other high-dimensional data such as volumetric MR or CBCT images, it generally needs more training samples, as the intrinsic dimensions of their state spaces are larger than that from surface images. This issue may be further composed with the fact that volumetric images are harder and slower to acquire, so the number of observations are expected to be low compared to the photogrammetry acquisitions in this study. On the other hand, since the training samples are static, there are various ways to “augment” the training set. For example, one may use model-based 4D reconstruction method to general infinite number of training samples [24]. Biomechanical models can also be used to “interpolate” high-dimensional states [21].

The main focus and contribution of this work is to construct a nonlinear sub-manifold where high-dimensional states would reside, which is an extension of PCA-based methods that have been used to derive low-dimensional subspaces [14]. It is different from conventional respiratory motion models that directly parameterize the temporal respiratory dynamics with few degrees of freedom, such as using B-spline functions or polynomials [13, 5]. The constructed manifold can be considered as a constraint that supplements those motion models.

As the specific prediction module is not the focus of this paper, a simple and crude VAR prediction module without adaptivity is used for demonstration purposes. We are in process of incorporating more sophisticated prediction modules into our proposed prediction framework.

Acknowledgments

This work is supported in part by NIH grant R01 CA169102-02. We also acknowledge the support from UCLA dissertation year fellowship and Vision RT.

References

1. Arias, P.; Randall, G.; Sapiro, G. Computer Vision and Pattern Recognition, 2007 CVPR'07 IEEE Conference on. IEEE; 2007. Connecting the out-of-sample and pre-image problems in kernel methods; p. 1-8.
2. Baumgartner CF, Kolbitsch C, Balfour DR, Marsden PK, McClelland JR, Rueckert D, King AP. High-resolution dynamic mr imaging of the thorax for respiratory motion correction of pet using

- groupwise manifold alignment. *Medical image analysis*. 2014; 18(7):939–952. [PubMed: 24972374]
3. Belkin M, Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Computation*. 2003; 15(6):1373–1396.
 4. Bengio Y, Paiement JF, Vincent P, Delalleau O, Le Roux N, Ouimet M. Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. *Advances in Neural Information Processing Systems*. 2004; 16:177–184.
 5. Blackall J, Ahmad S, Miquel M, McClelland J, Landau D, Hawkes D. MRI-based measurements of respiratory motion variability and assessment of imaging strategies for radiotherapy planning. *Physics in Medicine and Biology*. 2006; 51(17):4147. [PubMed: 16912374]
 6. Honeine P, Richard C. A closed-form solution for the pre-image problem in kernel-based machines. *Journal of Signal Processing Systems*. 2011; 65(3):289–299.
 7. Honeine P, Richard C. Preimage problem in kernel-based machine learning. *Signal Processing Magazine, IEEE*. 2011; 28(2):77–88.
 8. Keall PJ, Cattell H, Pokhrel D, Dieterich S, Wong KH, Murphy MJ, Vedam SS, Wijesooriya K, Mohan R. Geometric accuracy of a real-time target tracking system with dynamic multileaf collimator tracking system. *International Journal of Radiation Oncology Biology Physics*. 2006; 65(5):1579–1584.
 9. Kwok JTY, Tsang IWH. The pre-image problem in kernel methods. *Neural Networks, IEEE Transactions on*. 2004; 15(6):1517–1525.
 10. Létourneau D, Wong JW, Oldham M, Gulam M, Watt L, Jaffray DA, Siewerdsen JH, Martinez AA. Cone-beam-CT guided radiation therapy: technical implementation. *Radiotherapy and Oncology*. 2005; 75(3):279–286. [PubMed: 15890424]
 11. Liu W, Cheung Y, Sabouri P, Arai TJ, Sawant A, Ruan D. A continuous surface reconstruction method on point cloud captured from a 3D surface photogrammetry system. *Medical Physics*. 2015; 42(11):6564–6571. [PubMed: 26520747]
 12. Liu W, Cheung Y, Sawant A, Ruan D. A robust real-time surface reconstruction method on point clouds captured from a 3d surface photogrammetry system. *Medical Physics*. 2016; 43(5):2353–2360. [PubMed: 27147347]
 13. McClelland JR, Blackall JM, Tarte S, Chandler AC, Hughes S, Ahmad S, Landau DB, Hawkes DJ. A continuous 4d motion model from multiple respiratory cycles for use in lung radiotherapy. *Medical Physics*. 2006; 33(9):3348–3358. [PubMed: 17022231]
 14. McClelland JR, Hawkes DJ, Schaeffter T, King AP. Respiratory motion models: a review. *Medical image analysis*. 2013; 17(1):19–42. [PubMed: 23123330]
 15. Mika, S.; Schölkopf, B.; Smola, AJ.; Müller, KR.; Scholz, M.; Rätsch, G. *NIPS*. Vol. 4. Citeseer; 1998. Kernel PCA and de-noising in feature spaces; p. 7
 16. Murphy MJ, Dieterich S. Comparative performance of linear and nonlinear neural networks to predict irregular breathing. *Physics in Medicine and Biology*. 2006; 51(22):5903. [PubMed: 17068372]
 17. Poulsen PR, Cho B, Ruan D, Sawant A, Keall PJ. Dynamic multileaf collimator tracking of respiratory target motion based on a single kilovoltage imager during arc radiotherapy. *International Journal of Radiation Oncology Biology Physics*. 2010; 77(2):600–607.
 18. Roweis ST, Saul LK. Nonlinear dimensionality reduction by locally linear embedding. *Science*. 2000; 290(5500):2323–2326. [PubMed: 11125150]
 19. Ruan D. Kernel density estimation-based real-time prediction for respiratory motion. *Physics in Medicine and Biology*. 2010; 55(5):1311. [PubMed: 20134084]
 20. Ruan D, Keall P. Online prediction of respiratory motion: multidimensional processing with low-dimensional feature learning. *Physics in Medicine and Biology*. 2010; 55(11):3011. [PubMed: 20442460]
 21. Saadé J, Didier AL, Villard PF, Buttin R, Moreau JM, Beuve M, Shariat B. A preliminary study for a biomechanical model of the respiratory system. *Engineering and Computational Sciences for Medical Imaging in Oncology-ECSMIO*. 2010; 2010:7.

22. Shah AP, Dvorak T, Curry MS, Buchholz DJ, Meeks SL. Clinical evaluation of interfractional variations for whole breast radiotherapy using 3-dimensional surface imaging. *Practical Radiation Oncology*. 2013; 3(1):16–25. [PubMed: 24674259]
23. Tenenbaum JB, De Silva V, Langford JC. A global geometric framework for nonlinear dimensionality reduction. *Science*. 2000; 290(5500):2319–2323. [PubMed: 11125149]
24. Thomas D, Lamb J, White B, Jani S, Gaudio S, Lee P, Ruan D, McNitt-Gray M, Low D. A novel fast helical 4d-ct acquisition technique to generate low-noise sorting artifact-free images at user-selected breathing phases. *International Journal of Radiation Oncology Biology Physics*. 2014; 89(1):191–198.
25. Timmerman, RD.; Xing, L. *Image-guided and adaptive radiation therapy*. Lippincott Williams & Wilkins; 2012.
26. Wachinger C, Yigitsoy M, Rijkhorst EJ, Navab N. Manifold learning for image-based breathing gating in ultrasound and mri. *Medical Image Analysis*. 2012; 16(4):806–818. [PubMed: 22226466]




Figure 1.

Schematic of the proposed method: (1) mapping the original high-dimensional states to a reduced feature space via implicit nonlinear mapping ϕ using kernel operator, (2) performing prediction in the low-dimensional feature subspace, (3) mapping the prediction value back to the original high-dimensional state space via pre-image estimation.




Figure 2. Example level set surface reconstructed from one point cloud acquired by VisionRT, represented in height function.




Figure 3. Manifold learning and prediction results on 200 level-set surfaces: the blue solid curve represents the first-dimensional embeddings of 100 training surfaces learned by kernel PCA; the green solid curve represents the first-dimensional embeddings of 100 training surfaces learned by PCA; the red dashed curve represents the first-dimensional embeddings of 100 testing surfaces; the black dash-dot curve represents the estimated prediction in the feature subspace by VAR with 200ms lookahead.

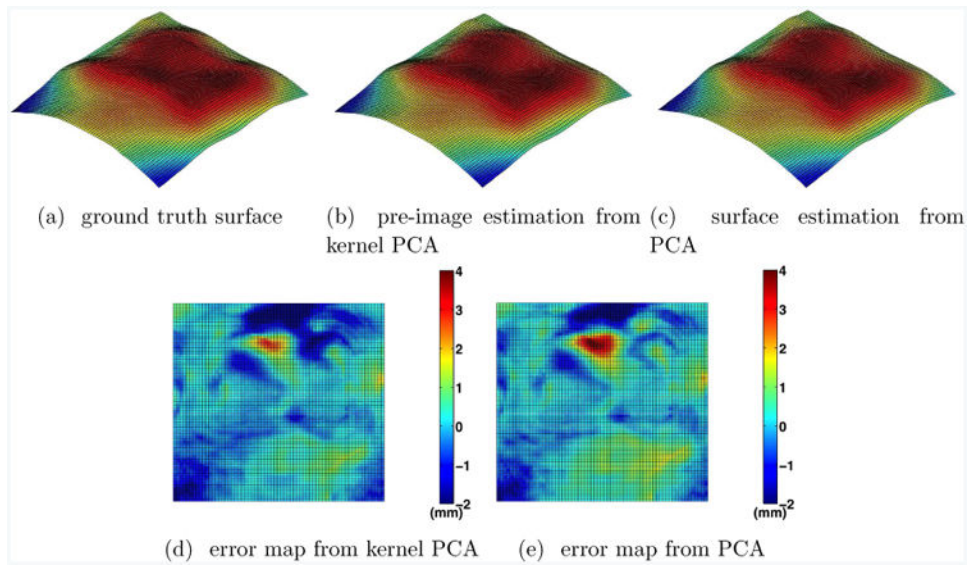


Figure 4. Comparison of estimated surfaces from kernel PCA and PCA against the ground truth: (a) the ground truth surface, (b) pre-image estimation from kernel PCA, (c) estimated surface from PCA, (d) error map depicting difference between (a) and (b), (e) error map depicting difference between (a) and (c).

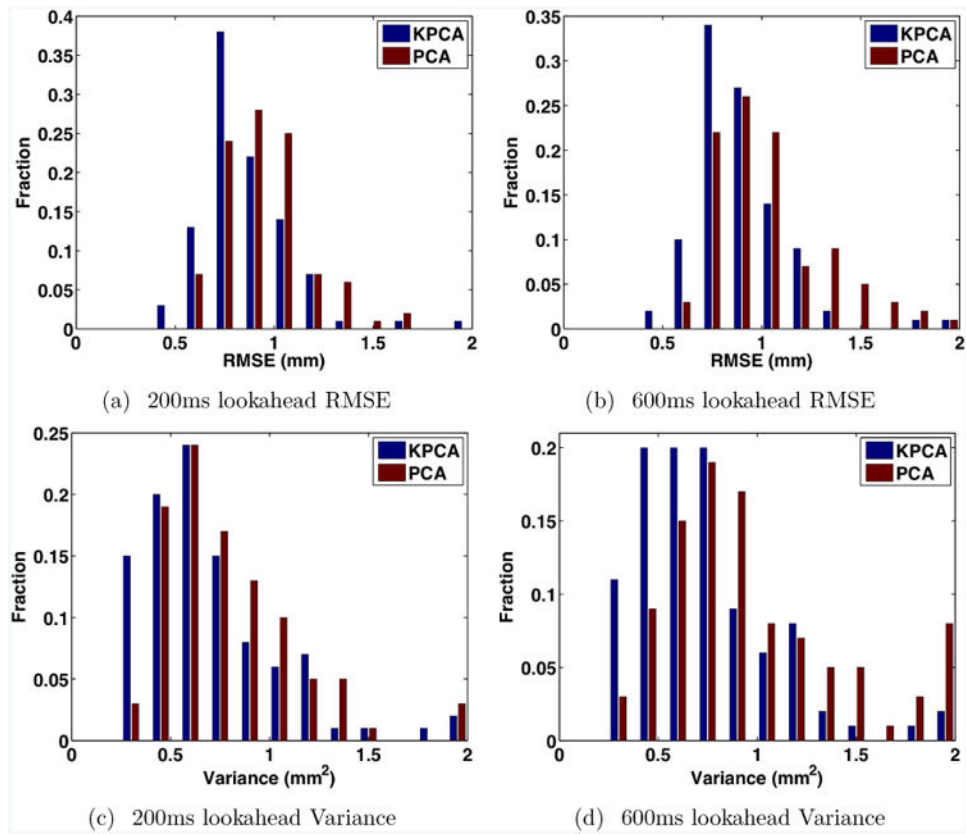


Figure 5. Histograms of RMSE and Variance from 100 prediction results using the first-dimensional embeddings.




Figure 6.

Prediction results on simulated irregular trajectory: the blue solid curve represents the training trajectory after warping; the red dashed curve represents the testing trajectory after warping; the black dash-dot curve represents the estimated prediction by VAR with 200ms lookahead.

Error statistics of the estimated pre-images from 100 testing surfaces based on RMSE (in mm).

Table 1

dimension of feature subspace	200ms lookahead			600ms lookahead		
	1	2	3	1	2	3
PCA						
Mean	0.95	1.06	1.45	1.04	1.33	1.89
S.D.	0.22	0.31	0.58	0.28	0.42	0.73
Median	0.92	1.03	1.32	0.97	1.26	1.83
KPCA						
Mean	0.86	0.93	1.08	0.89	1.00	1.44
S.D.	0.24	0.31	0.40	0.24	0.32	0.49
Median	0.81	0.86	0.97	0.85	0.92	1.37

P-values from paired Mann-Whitney U-test on RMSE from KCPA and PCA prediction results.

Table 2

dimension of feature subspace	200ms lookahead			600ms lookahead		
	1	2	3	1	2	3
P-value	4.41e-6	2.26e-7	2.22e-8	2.40e-8	8.29e-13	1.74e-6