A Kernel Ridge Regression Model for Respiratory Motion Estimation in Radiotherapy

Tobias Geimer[1,2], Adriana Birlutiu[3], Mathias Unberath[1,2], Oliver Taubmann[1,2], Christoph Bert[2,4], Andreas Maier[1,2]

[1]Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg
[2]Graduate School in Advanced Optical Technologies (SAOT), Erlangen
[3]Computer Science Department, "1 December 1918" University of Alba Iulia, Romania
[4]Department of Radiation Oncology, Universitätsklinikum Erlangen, Friedrich-Alexander-Universität Erlangen-Nürnberg
tobias.geimer@fau.de

Abstract. This paper discusses a kernel ridge regression (KRR) model for motion estimation in radiotherapy. Using KRR, dense internal motion fields are estimated from high-dimensional surrogates without the need for prior dimensionality reduction. We compare the proposed model to a related approach with dimensionality reduction in the form of principal component analysis and principle component regression. Evaluation was performed in a simulation study based on nine 4D CT patient data sets achieving a mean estimation error of 0.84 ± 0.21 mm for our approach.

1 Introduction

Respiratory motion is of concern for several medical procedures in the thoracic and abdominal areas. In external beam radiation therapy, intra-fractional motion may lead to underdosing the clinical target volume (CTV), if not addressed properly [1]. The patient is irradiated according to a treatment plan based on CT imaging, defining an optimized dose distribution. However, respiratory motion typically leads to displacement of the CTV, resulting in insufficient dose in the target and thus the potential survival of malignant cells. One option is to introduce additional margins covering the extent of the CTV's motion at the cost of higher dose to healthy tissue. More preferably, real-time motion estimation can be used to either restrict exposure time to certain parts of the respiratory cycle (gating) or adjust the beam according to the target volume (tracking).

For real-time motion estimation, a patient-specific motion model can be trained pre-procedurally connecting a highly correlated external surrogate signal to the corresponding internal deformation [2]. The ground-truth deformation field is usually obtained from 4D imaging by registration to a reference phase. In recent literature, various methods for ground-truth-to-surrogate correspondence estimation such as (multi-)linear regression [3] are employed. Then, intra-procedural acquisition of just the surrogate signal allows for inference of the internal motion field [2].

Wilms et al. [3] investigated multi-variate regression approaches based on range imaging among other surrogates. While these approaches operate directly on the acquired data, others also based on multi-linear regression incorporate an additional generalization step in the form of Principal Component Analysis (PCA) to describe both the internal and external variation of a patient's breathing cycle [4,5]. As an additional benefit, the reduced dimensionality in feature space also reduces the complexity of the regression problem. Another way of dealing with high-dimensional domains is to incorporate kernels into the regression. Li and Xing [6] investigated kernel-based respiratory motion estimation but only used a 1-D surrogate. The major benefit of such a kernel approach is its ability to represent non-linear mappings between internal and external motion.

In this work, we present a correlation model based on Kernel Ridge Regression (KRR) to estimate dense internal motion fields from two marker-less high-dimensional surrogates: range imaging [7] and X-ray fluoroscopy. We evaluate our approach in a simulation study on 4D CT data of nine cancer patients and compare it to Principal Component Regression (PCR).

2 Materials and Methods

First, we will introduce our notation. Then, the two correlation approaches will be covered: a) related work in the form of PCR and b) KRR without prior dimensionality reduction. An overview of our data and evaluation methods will conclude the section.

2.1 Data Matrices

The respiratory motion model can be trained pre-procedurally from 4D imaging such as the planning CT for radiotherapy. Performing demons-based non-rigid registration [8], n internal deformation fields $\{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_n\}, \boldsymbol{t}_i \in \mathbb{R}^{d_t}$ are obtained, that are stored column-wise in the data matrix $\boldsymbol{T} \in \mathbb{R}^{d_t \times n}$. Similarly, $\boldsymbol{S} \in \mathbb{R}^{d_s \times n}$ denotes the n corresponding surrogate observations $\boldsymbol{s}_i \in \mathbb{R}^{d_s}$. These can either be the patient's thorax surface motion fields or fluoroscopic images at the same breathing phase. For training purposes, they are extracted from the 4D CT as well (see Sec. 2.4).

2.2 Principal Component Regression

Principal Component Analysis. PCA is a popular linear dimensionality reduction technique [9]. It can be used to decompose a given data set into mutually orthogonal modes of variation, called principal components. With the first few components often being sufficient to represent more than 90% of the variance present in the data set, the number of basis vectors is less than that of the original domain. Using PCA, the data sets $\boldsymbol{S} \in \mathbb{R}^{d_s \times n}$ and $\boldsymbol{T} \in \mathbb{R}^{d_t \times n}$ can, therefore, be represented by a set of features $\boldsymbol{F}_S \in \mathbb{R}^{p_s \times n}$ and $\boldsymbol{F}_T \in \mathbb{R}^{p_t \times n}$ of reduced dimensionality.

Multi-linear Regression. Let $\boldsymbol{F}_S \in \mathbb{R}^{p_s \times n}$ and $\boldsymbol{F}_T \in \mathbb{R}^{p_t \times n}$, where n is the number of observations and p_s, d_t denote feature dimensionality chosen for the PCA. Correlation between the target and the surrogate domain can be formulated as a multi-linear regression problem [4]. Combined with Tikhonov regularization, this is also known as (standard) Ridge Regression with the corresponding objective function

$$\underset{\boldsymbol{W}}{\operatorname{argmin}} \left(\frac{1}{2} ||\boldsymbol{W}\boldsymbol{F}_S - \boldsymbol{F}_T||_{\mathcal{F}}^2 + \alpha \frac{1}{2} ||\boldsymbol{W}||_{\mathcal{F}}^2 \right), \tag{1}$$

where $||\cdot||_{\mathcal{F}}$ is the Frobenius norm. The closed-form solution involves the Moore-Penrose pseudo-inverse

$$\boldsymbol{W} = \underbrace{\boldsymbol{F}_T \boldsymbol{F}_S^\top}_{p_t \times p_s} (\underbrace{\boldsymbol{F}_S \boldsymbol{F}_S^\top}_{p_s \times p_s} + \alpha \mathbb{I}_{p_s})^{-1} \in \mathbb{R}^{p_t \times p_s}, \tag{2}$$

which is computed using singular value decomposition (SVD). For low-dimensional feature spaces ($p_s, p_t \leq n$) $\boldsymbol{W}$ can be calculated explicitly.

2.3 Kernel Ridge Regression

The KRR method [9] is a regularized least squares method for classification and regression. For high-dimensional data $\boldsymbol{T}$ and $\boldsymbol{S}$ an explicit computation of $\boldsymbol{W}$ as presented in Eqn. 2 without prior dimensionality reduction is computationally expensive (as $d_s, d_t \gg p_s, p_t$). Fortunately, Eqn. 2 can be rewritten to:

$$\begin{aligned} \boldsymbol{W} &= \underbrace{\boldsymbol{T}\boldsymbol{S}^\top}_{d_t \times d_s} (\underbrace{\boldsymbol{S}\boldsymbol{S}^\top}_{d_s \times d_s} + \alpha \mathbb{I}_{d_s})^{-1} \\ &= \underbrace{\boldsymbol{T}}_{d_t \times n} (\underbrace{\boldsymbol{S}^\top \boldsymbol{S}}_{n \times n} + \alpha \mathbb{I}_n)^{-1} \underbrace{\boldsymbol{S}^\top}_{n \times d_s}. \end{aligned} \tag{3}$$

Making use of the kernel trick, the observations $\boldsymbol{s}_i$ are implicitly mapped to an even higher-dimensional Reproducing Kernel Hilbert space [10]:

$$\boldsymbol{\Phi} = \left[\phi(\boldsymbol{s}_1), \ldots, \phi(\boldsymbol{s}_n) \right]. \tag{4}$$

When predicting a target $\boldsymbol{t}_{\text{pred}}$ from a new observation $\boldsymbol{s}_{\text{new}}$, explicit access to $\boldsymbol{\Phi}$ is never actually needed:

$$\begin{aligned} \boldsymbol{t}_{\text{pred}} &= \boldsymbol{T} \left(\boldsymbol{\Phi}^\top \boldsymbol{\Phi} + \alpha \, \mathbb{I}_n \right)^{-1} \boldsymbol{\Phi}^\top \phi(\boldsymbol{s}_{\text{new}}) \\ &= \boldsymbol{T} \left(\boldsymbol{K} + \alpha \, \mathbb{I}_n \right)^{-1} \boldsymbol{\kappa} \left(\boldsymbol{s}_{\text{new}} \right). \end{aligned} \tag{5}$$

With $\boldsymbol{K}_{ij} = \boldsymbol{\phi}\left(\boldsymbol{s}_i\right)^\top \boldsymbol{\phi}\left(\boldsymbol{s}_j\right)$ and $\boldsymbol{\kappa}\left(\boldsymbol{s}_{\text{new}}\right)_i = \boldsymbol{\phi}\left(\boldsymbol{s}_i\right)^\top \boldsymbol{\phi}\left(\boldsymbol{s}_{\text{new}}\right)$, the prediction can be described entirely in terms of inner products in the higher-dimensional space. Not only does this approach work on the original data sets as well as the principal component representations, it also opens up ways to introduce non-linear mappings into the regression.

2.4 Data and Evaluation

Evaluation was conducted on nine 4D CT data sets of lung tumor patients treated at UK Erlangen. Each data set provided nine respiratory deformation fields of spacing $0.97 \times 0.97 \times 2.5\,\mathrm{mm}^3$, which were manually cropped to an internal region of interest (ROI) to represent the ground truth. For simulating the range imaging surrogate data, the fields were interpolated at a surface mesh directly extracted from the CT acquisition at the reference phase. This provided between 2252 and 4667 surface points for each patient. For fluoroscopy, Digitally Reconstructed Radiographs (DRR) with 1024×768 pixels and 0.39 mm pixel spacing were generated by forward-projecting the 4D CT volumes.

Leave-one-out evaluation was performed for each of the patients individually. Each phase was subsequently tested and, for this purpose, removed from the training sample together with its two neighbors. Estimation accuracy was assessed by the L2-norm of the residual vectors between estimated field and ground truth. Parameters for regularization and the Gaussian kernel were determined using grid search in the same leave-one-out manner. Similarly, internal model accuracy was determined, given perfect estimation of the weights from the surrogate. For PCR, the generalization ability of the internal model is the limiting factor. For KRR, each deformation field $\boldsymbol{t}_i$ was compared to the optimal guess $\boldsymbol{t}_{i,\mathrm{optim}} = \boldsymbol{T}\boldsymbol{a}_i$, where $\boldsymbol{a}_i = \mathrm{argmin}_{\tilde{\boldsymbol{a}}_i} \left(||\boldsymbol{T}\tilde{\boldsymbol{a}}_i - \boldsymbol{t}_i||_2^2 \right)$. This deformation field arises from interpolation weights that are optimal in a least-squares sense. As a result, we have a lower bound on the estimation error (dashed lines in Fig. 3).

3 Results

Fig. 1 shows the grid search result for PCR. Estimation error decreases for higher internal model dimensionality as more components also explain a higher amount of variance present in the data. The need for regularization was low for both surrogate types. The overall estimation error averaged over all patients and

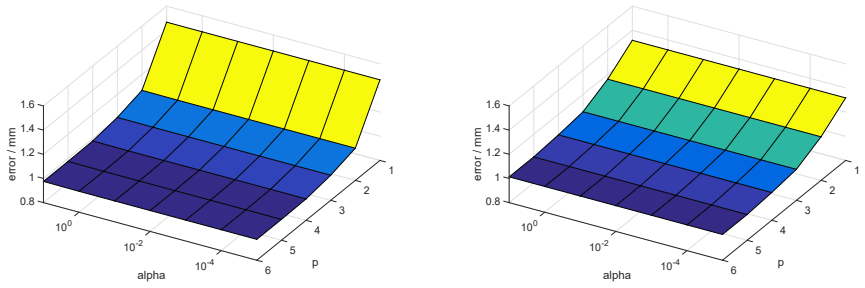

(a) Range imaging (b) Fluoroscopy

Fig. 1. Estimation error using surrogates from range imaging (a) and fluoroscopy (b) for $p = 1, \ldots, 6$ components and different regularization strengths.

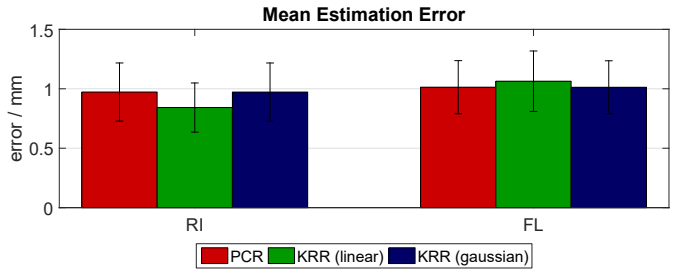

Fig. 2. Mean estimation error using PCR and KRR with a linear and a Gaussian kernel for the surface (RI) and the fluoroscopy (FL) surrogate.

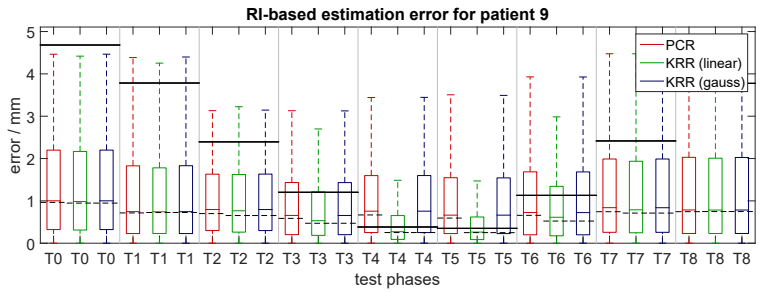

Fig. 3. Surface-based estimation error for each phase of patient 9. The black bars show the mean magnitude of the ground truth deformation field, while the dashed black lines indicate the lower bound for the mean error achievable with the respective model.

phases is plotted in Fig. 2. For a more detailed look, Fig. 3 shows the estimation error of a single patient for each breathing phase for the surface surrogate only, including the breathing magnitude and the inherent model error for each phase.

In general, all proposed methods are suitable for motion estimation with a mean estimation error of around 1.0 ± 0.22 mm compared to a reference mean magnitude of 2.48 ± 0.81 mm without compensation. Data-based KRR for the surface surrogate achieved the best average estimation error of 0.84 ± 0.21 mm.

4 Discussion

Contrary to expectation, no significant benefits could be observed from applying non-linear KRR instead of PCR. This may be due to the low number of 6 phases available to train the models. Special attention should be given to the way the internal motion is reconstructed for PCR and KRR, respectively. In the case of KRR, the new phase is a linear interpolation between all observed phases. For PCR, the regression result yields the weights for the principal components from which the internal motion field is reconstructed as a linear combination of eigenvectors. Thus, PCA introduces an additional generalization step not present

in data-based KRR. However, from the dashed lines in Fig. 3 it can be seen that the principal components lack the variation to describe an entire cycle. Because of our 6 : 3 split for training and evaluation, the internal model only observed two thirds of a breathing cycle. This is particularly obvious for phases near end-exhale where the magnitude of breathing is low. Here, both PCR and Gaussian kernel produce errors higher than the reference magnitude. Only surface-based linear KRR was able to cope with low-magnitude phases near end-exhale. For the fluoroscopy surrogate, however, this trend could not be observed.

In summary, all mean estimation errors were close to their respective lower bound. Thus, the major bottleneck seems to be the reconstruction of the new phase from the internal model rather than the correlation problem itself.

For future work, an extended evaluation of both PCR and KRR is necessary. Using a patient's planning CT to train on an entire breathing cycle with evaluation on a potential follow-up CT will bring evaluation closer to the actual application case, where the treatment is performed days or even weeks after the planning images were acquired. This will also give a good indication whether the additional generalization step of PCA will proof beneficial in comparison to the phase interpolation of KRR.

References

1. Keall PJ, Mageras GS, Balter JM, Emery RS, Forster KM, Jiang SB, et al. The management of respiratory motion in radiation oncology report of AAPM Task Group 76. Med Phys. 2006;33(10):3874–3900.
2. McClelland JR, Hawkes DJ, Schaeffter T, King AP. Respiratory motion models: A review. Med Image Anal. 2013;17(1):19 – 42.
3. Wilms M, Werner R, Ehrhardt J, Schmidt-Richberg A, Schlemmer HP, Handels H. Multivariate Regression Approaches For Surrogate-based Diffeomorphic Estimation of Respiratory Motion in Radiation Therapy. Phys Med Biol. 2014;59(5):1147.
4. Taubmann O, Wasza J, Forman C, Fischer P, Wetzl J, Maier A, et al. Prediction of Respiration-Induced Internal 3-D Deformation Fields From Dense External 3-D Surface Motion. In: 28th Trans Comput Assist Radiol Surg (CARS). Heidelberg; 2014. p. 33–34.
5. Geimer T, Unberath M, Taubmann O, Bert C, Maier A. Combination of Markerless Surrogates for Motion Estimation in Radiation Therapy. In: 30th Trans Comput Assist Radiol Surg (CARS); 2016. p. 59–60.
6. Li R, Xing L. A Kernel Method for Real-Time Respiratory Tumor Motion Estimation Using External Surrogates. In: Machine Learning and Applications and Workshops (ICMLA), 10th International Conference. vol. 2; 2011. p. 206–209.
7. Wasza J, Fischer P, Leutheuser H, Oefner T, Bert C, Maier A, et al. Real-Time Respiratory Motion Analysis Using 4-D Shape Priors. IEEE Trans Biomed Eng. 2016;63(3):485–495.
8. Fischer B, Modersitzki J. A unified approach to fast image registration and a new curvature based registration technique. Linear Algebra Appl. 2004;380(0):107–124.
9. Friedman J, Hastie T, Tibshirani R. The Elements of Statistical Learning. vol. 1. Springer Series in Statistics, Berlin; 2001.
10. Berlinet A, Thomas-Agnan C. Reproducing Kernel Hilbert Spaces in Probability and Statistics. Springer Science & Business Media; 2011.