

Modeling inter- and intra-patient anatomical variation using a bilinear model

Yongwon Jeong and Richard J. Radke

Rensselaer Polytechnic Institute

Department of Electrical, Computer, and Systems Engineering

110 8th St., Troy, NY 12180

jeongy@rpi.edu, rjradke@ecse.rpi.edu

Abstract

We introduce a novel method for modeling the natural anatomical variation of an organ along inter-patient and intra-patient axes using a bilinear model. Bilinear models are attractive for this purpose since one type of variation can be explored while the other is held constant. We apply our model to a total of 204 prostate shapes contoured from CT imagery of 12 different patients, and show that the learned bilinear models can fit both training and testing shapes accurately. We also demonstrate the superior performance of the bilinear model over a linear model with the same number of parameters in adapting to prostate shapes from a new patient acquired immediately prior to consecutive fractions of radiation therapy.

1. Introduction

Due to the introduction of imaging technologies such as CT and MRI and advances in computer vision techniques, image analysis has come to play an important role in the medical field [4]. Shape information about an anatomical object is often helpful for diagnostic and treatment purposes. For example, the shape variation of the organ over time can indicate the presence or absence of disease. Our particular interest is in shape variation as it applies to radiation therapy treatment planning. Better shape models can help localize a target object more precisely so the outcome of treatment can be improved. In this paper, we suggest a novel application of a bilinear model to the 3D shape modeling of the prostate for the purpose of guidance in radiotherapy treatment. While bilinear models have been introduced in other areas of computer vision, to our knowledge they have rarely been used in biomedical imaging applications. The advantage of the bilinear shape model is that it can decouple the expected variations both across the patient population and within a specific patient (*e.g.*, due to changes in bodily state). The paper is organized as follows. In Section

2, we briefly review several shape modeling techniques for anatomical objects. The details of our bilinear shape modeling technique are explained in Section 3. In Section 4, we describe experimental results that demonstrate the effective performance of the bilinear model on the task of adapting to an unseen patient. We summarize and discuss future work in Section 5.

2. Prior work, landmark selection, and linear models

There have been many shape modeling methods described for medical applications. A good review can be found in [10]. Here, we mention some of the methods that are relevant to our work. Huang and Amini [7] applied deformable models based on B-splines to calculate the volume of the brain and heart. Pizer *et al.* [11] proposed the M-rep medial representation for the segmentation of medical images. They extracted medial atoms from the medial plane of a 3D object and parameterized each atom. They used a hierarchical approach to control the shape of an object at various levels (*e.g.*, from global translation to local curvature). The popular Active Shape Model (ASM) proposed by Cootes *et al.* [2] has been used to segment medical images including the heart chamber, prostate, and ventricles of the brain. The ASM describes shape changes in a statistical framework by capturing how important points on an object vary based on training data. We will use the ASM for our basic comparisons because it can describe shape variation in a compact form, and is related to the bilinear shape model, as we explain below.

Both linear and bilinear shape modeling methods are based on Point Distribution Models (PDMs) [3]. In a PDM, the shape of an object is expressed as a set of points distributed along its boundary. These points, or “landmarks”, can be selected in various ways. Our method of selecting landmark points suitable for medical tomographic imagery is summarized in Figure 1, and is based on the method described in [6]. The training data is composed of parallel

CT slices. In each slice, an expert has indicated several points around the boundary of an organ. The difficulty is that the numbers of slices per scan, the numbers of points per slice and the relative positions of the points differ from dataset to dataset, making it difficult to establish correspondence for a shape model. An interpolating surface that goes through all of the original points in each dataset is constructed by fitting a smooth surface to the data points. The interpolating surface can be found using an implicit function method such as [14]. The interpolating surface is then resampled into a new set of parallel slices at uniform intervals and each slice is sampled into a pre-defined number of points at fixed locations. We select the location of points on a slice in the following way: from the geometric centroid of each slice, a fixed number of rays (e.g., 20) are extended until they meet the boundary of the slice. Then, the points where the rays and the boundary meet are selected as the set of points to be used as the training vector for the shape model. Let us denote the training data composed of T instances as $\{\mathbf{x}(t), t = 1, \dots, T\}$, where each instance is composed of N points at specified locations in 3D: $\mathbf{x}(t) = [x_1, y_1, z_1, \dots, x_N, y_N, z_N]^T$ (i.e. each training instance is a column vector).

To build the standard active shape model as in [3], principal component analysis is applied as follows. First, the mean shape is obtained by averaging the training vectors:

$$\bar{\mathbf{x}} = \frac{1}{T} \sum_{t=1}^T \mathbf{x}(t). \quad (1)$$

The covariance matrix of the training data is computed as

$$\Sigma = \frac{1}{T-1} \sum_{t=1}^T (\mathbf{x}(t) - \bar{\mathbf{x}})(\mathbf{x}(t) - \bar{\mathbf{x}})^T. \quad (2)$$

Then, eigenvalue decomposition is carried out for the covariance matrix to obtain mode vectors

$$\Phi = [\phi_1 \phi_2 \dots \phi_T]. \quad (3)$$

where ϕ_i 's are the eigenvectors ordered according to decreasing magnitude of their corresponding eigenvalues λ_i . Retaining only L significant eigenvectors has the advantages of 1) reducing the error induced by inconsistent contouring, and 2) reducing the dimensionality of the model. For example, when only a portion α ($0 \leq \alpha \leq 1$) of the total variance is to be captured in the retained mode vectors, L can be determined as the minimum l such that

$$\frac{\sum_{i=1}^l \lambda_i}{\sum_{i=1}^T \lambda_i} \geq \alpha. \quad (4)$$

The mode vectors express the orthogonal components of variation from the mean shape in the training data. Therefore, the resulting shape model can generate shapes in terms




Figure 1. Illustrating contour resampling. (a) Original contours arranged in 17 axial slices describe a prostate shape. (b) A smooth surface interpolating the original contours is constructed. (c) The interpolating surface is resampled into parallel slices and each slice is resampled at corresponding locations (e.g., 20 slices, 20 points per slice).

of a weighted sum of mode vectors plus the mean shape:

$$\hat{\mathbf{x}} = \bar{\mathbf{x}} + \Phi \mathbf{p}. \quad (5)$$

where \mathbf{p} is the vector of model control parameters.

3. Shape modeling using bilinear models

The linear model described in the previous section can be thought of as a special case of a bilinear model [8], instances of which have been used in several computer vision applications in recent years. For example, Marimont and Wandell [9] used a bilinear model to decouple sensor responses for collections of varying surfaces and illuminants. Tenenbaum and Freeman [13] applied bilinear models to several tasks that have two natural independent variations, such as images of the same set of faces seen in different poses or lighting conditions, alphabetical letters in different fonts, and English vowels spoken by different speakers.

Bilinear and, more generally, multilinear models have been popularly used for face modeling in particular. Basicle and Blake [1] used a bilinear model to decouple pose and expression from facial images for facial animation. Vasilescu and Terzopoulos [15] used a multilinear model to capture facial pose, illumination, and expression separately. In the same multilinear framework, Vlasic *et al.* [16] built a generative model for variations in face, expression, and viseme (mouth articulations related to speech).

Finally, while we do not explore them here, there have also been approaches to separating two factors using nonlinear models. For example, Elgammal and Lee [5] embedded the observations from dynamic objects such as gait images into nonlinear manifolds to separate two factors (e.g., person and body configuration) in a generative model. Soatto and Yezzi [12] separated motion and deformation to track deformable objects that are also moving (such as images of a jellyfish) in a level set framework.

It is customary that the two variations of a bilinear model are called “style” and “content”. Their roles are mathematically interchangeable, but there is usually a natural semantic

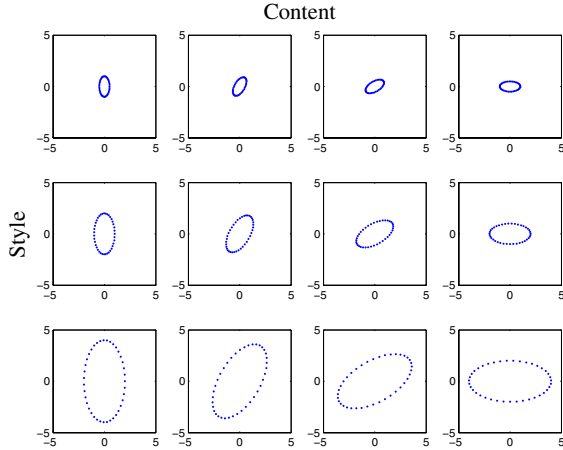


Figure 2. Examples of observations of ellipses of varying size and orientation. “Style” is defined as the size of the ellipse and “content” is defined as the orientation of the ellipse.

interpretation. An example of style and content is illustrated in Figure 2 for an observation vector composed of 40 points distributed around an ellipse. In this case, the style and content are defined as the size and the orientation of the ellipse, respectively.

In building a bilinear model, basis functions along with style and content parameters are found such that the resulting model minimizes the total squared error between the model and the given observation vectors in the training set. Before training, each observation must be classified into a style category s and a content category c (again, there is usually a natural semantic basis for this categorization). Suppose there are S total style categories and C total content categories. It is not necessary that each style/content combination has an instance in the training set, or that there are the same number of instances in each category.

The mathematical goal is to find K basis function matrices, denoted $\{\mathbf{W}_k \in \mathbb{R}^{I \times J}, k = 1, \dots, K\}$, and vectors of style and content parameters, $\{\mathbf{a}_s \in \mathbb{R}^I, s = 1, \dots, S\}$ and $\{\mathbf{b}_c \in \mathbb{R}^J, c = 1, \dots, C\}$ respectively, that minimize the cost function

$$\sum_{t=1}^T \sum_{s=1}^S \sum_{c=1}^C \sum_{k=1}^K \chi_{sc}(t) (\mathbf{x}_k(t) - \mathbf{a}_s^T \mathbf{W}_k \mathbf{b}_c)^2, \quad (6)$$

where T is the number of K -dimensional observation vectors in the training set and $\chi_{sc}(t)$ is an indicator function which is 1 if $\mathbf{x}(t)$ is in style s and content c and 0 otherwise. The dimensionality of the style control vector (I) and content control vector (J) are chosen so that $1 \leq I \leq S$ and $1 \leq J \leq C$, analogous to the way the number of dominant modes is chosen in a linear PCA model. The set of matrices \mathbf{W}_k plays a similar role to the orthogonal mode matrix Φ in a linear PCA model, and describes the interaction between the style and content factors.

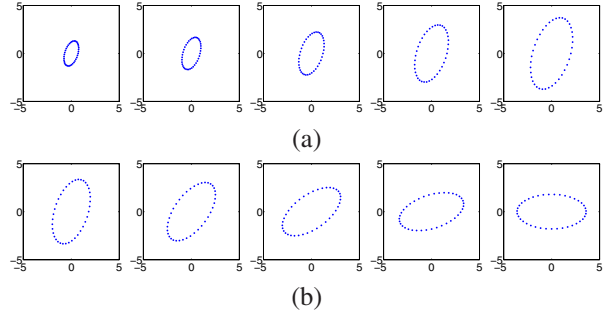


Figure 3. New observations generated using a bilinear model built from the observations in Figure 2. (a) Varying the style control parameter with content control parameter fixed. (b) Varying the content parameter with style control parameter fixed.

Once basis function matrices and control parameters that minimize the total squared error are calculated, the generative bilinear model for a new style vector \mathbf{a} and content vector \mathbf{b} is simply

$$\hat{\mathbf{x}}_k = \mathbf{a}^T \mathbf{W}_k \mathbf{b}. \quad (7)$$

The procedure to estimate the parameters of the model uses either iterative-SVD (Singular Value Decomposition) or direct minimization, depending on how the training observations are distributed in each style and content class. For a more detailed explanation for the building and fitting of the bilinear model, see [13]. New observations generated from a bilinear model with $I = S = 3$ and $J = C = 4$ for the ellipse example are shown in Figure 3. Observations in a new style (*i.e.*, size) and content (*i.e.*, orientation) are shown in the figure. In each case, examples unobserved (but consistent with) the training set can be generated.

In this paper, the underlying motivation for using a bilinear model is that there are two general types of variation inherent in images of the same organ: intra-patient variation caused by changes in a fixed patient’s bodily state, and inter-patient variation caused by differences between patients across the population. The idea is that we can effectively decouple these two factors in a dataset containing many different scans from each of many different patients - something that is not possible with the linear model. Thus, we define “style” as patient ID and “content” as the bodily state within a patient. The parameter vector \mathbf{a} controls the inter-patient variation and the parameter vector \mathbf{b} controls the intra-patient variation. Consequently, the model can adapt to a new patient by estimating the inter-patient (style) control parameter \mathbf{a} using a few scans and the adapted model should behave consistently with the training data by varying the intra-patient (content) control parameter \mathbf{b} . The mathematical approach is the following:

1. Assemble a training database of organ shapes from many patients observed at multiple times.




Figure 4. Axial CT images with manually-outlined prostates from Patient 1 (top row) and Patient 2 (bottom row) on two different days.

Set number	Total shapes	Slices per shape
1	17	11–13
2	17	10–12
3	19	9–18
4	19	9–15
5	13	8–9
6	14	6–14
7	18	11–21
8	18	10–17
9	18	9–18
10	17	9–15
11	17	8–14
12	17	7–15

Table 1. Summary of the dataset used for the experiment.

- Sort each scan into a style category and a content category (see Section 4.1 below).
- Learn the basis matrices $\mathbf{W}_k \in \mathbb{R}^{I \times J}$ by minimizing (6).
- For M scans of a previously unseen patient, $\{\mathbf{x}(m), m = 1, \dots, M\}$, estimate the fixed style vector \mathbf{a} and the content vectors $\{\mathbf{b}_1, \dots, \mathbf{b}_M\}$ by minimizing

$$\sum_{m=1}^M \sum_{k=1}^K \|\mathbf{x}_k(m) - \mathbf{a}^T \mathbf{W}_k \mathbf{b}_m\|^2. \quad (8)$$

Recall that the \mathbf{W}_k are fixed after the training process.

- Generate a style-specific model for the new patient by fixing \mathbf{W} and \mathbf{a} and varying \mathbf{b} as in (7).

4. Datasets and experiments

4.1. Training

Our training dataset is composed of prostate shapes from 99 3D CT images of the male pelvis from 6 patients. The




Figure 5. Original sets of points describing prostates from Patient 1 (top row) and Patient 2 (bottom row) on two different days.

images from each patient were acquired immediately prior to consecutive fractions of radiation treatment. An expert planner manually contoured the prostate in every slice of each dataset; several examples of contoured images are shown in Figure 4. The corresponding prostate contour points arranged in axial slices (prior to resampling) are shown in Figure 5. The original sets of points are resampled to yield 400 points (20 slices per shape, 20 points per slice) following the procedure described in Section 2, resulting in the shapes illustrated in Figure 6. The resampled points are converted into column vectors and used to build and test the linear and bilinear models. Our testing dataset is composed of 105 prostate shapes from a different set of six patients. The number of shapes and original number of slices per shape in both datasets are shown in Table 1. Sets 1–6 were used for model building and sets 7–12 were used for model fitting.

Building a linear shape model with this data is straightforward as described in Section 2. However, to build the bilinear model, we need to categorize shapes in the training data into different style and content classes. Style is simply defined as patient ID, so there are 6 styles in our training data (and we choose the style vector to be 6-dimensional). However, content categories are more difficult to define, because there is no obvious way to parameterize the inherent variation of the prostate independent of personal variation.¹

In the experiment reported here, we defined the content classes based on the relative volume of each patient’s prostate as follows. For all the datasets from a given patient, the volume of the prostate is measured and normalized so that the maximum volume has value 1 and the minimum

¹This task would be easier, for example, in the lung, in which content classes could be naturally defined as phases in the breath cycle.




Figure 6. Resampled sets of points describing prostates from Patient 1 (top row) and Patient 2 (bottom row) on two different days.




Figure 7. Mean training shapes assigned to each style and content class (there is one blank entry for a class that has no assigned training shapes).

volume has value 0, *i.e.*, the normalized volume is calculated as

$$v_{normalized}^i = \frac{v^i - v_{min}}{v_{max} - v_{min}}. \quad (9)$$

The samples from each patient are then binned with respect to the normalized volume. We varied the number of content bins from 1 to 6. Figure 7 shows the mean shapes of each style and content class for $S = 6$ and $C = 5$. The training shapes labeled with corresponding style and content are used to build the bilinear model as described in Section 3. The bilinear model with $I = S = 6$ and $J = C = 5$ using the training set can generate new shapes as shown in Figure 8. The effect of varying the content for a fixed style is easy to interpret, but it is more difficult to characterize in words the style variation for fixed content.

We must turn to more quantitative measures of how faithfully each model can approximate a test shape. The

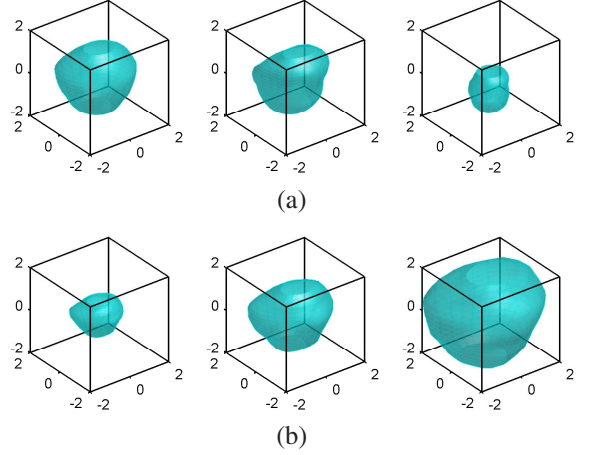


Figure 8. New observations generated using the prostate bilinear model. (a) Varying the style control parameter with content control parameter fixed. (b) Varying the content parameter with style control parameter fixed.

average Euclidean distance is used as a distance measure between a test shape and its projected model. That is, for a set of N points in 3D describing a shape, $\mathbf{x} = \{x_1, y_1, z_1, \dots, x_N, y_N, z_N\}$, the approximation error is defined as

$$E = \frac{1}{N} \sum_{n=1}^N \left\| \begin{bmatrix} x_n \\ y_n \\ z_n \end{bmatrix}_{test} - \begin{bmatrix} x_n \\ y_n \\ z_n \end{bmatrix}_{approximated} \right\|. \quad (10)$$

For the linear model specified by $(\bar{\mathbf{x}}, \Phi)$, the minimum squared error approximation of a test shape \mathbf{x} is carried out by projecting it into the model space via

$$\hat{\mathbf{x}} = \bar{\mathbf{x}} + \Phi \mathbf{p}_{projected}, \quad (11)$$

where

$$\mathbf{p}_{projected} = \Phi^T (\mathbf{x}_{test} - \bar{\mathbf{x}}). \quad (12)$$

To approximate a test shape \mathbf{x} using the bilinear model specified by \mathbf{W} , optimum model parameters are found that minimize the modeling error, *i.e.*

$$\hat{\mathbf{x}}_k = \mathbf{a}^T \mathbf{W}_k \mathbf{b} \quad (13)$$

where

$$\{\mathbf{a}, \mathbf{b}\} = \arg \min_{\{\mathbf{a}, \mathbf{b}\}} \|\mathbf{x}_k - \mathbf{a}^T \mathbf{W}_k \mathbf{b}\|^2. \quad (14)$$

We measure the accuracy in three situations: accuracy over all training data, accuracy over all testing data, and accuracy in an adaptation experiment that reflects how the models might actually be used in a fractionated radiotherapy context.

# parameters	Training err (mm)	Testing err (mm)
10	1.43	1.52
11	1.30	1.40
12	1.23	1.32

Table 2. Modeling accuracy: projection error.

4.2. Training and testing accuracy

First, we measured the training accuracy by evaluating the bilinear model over the same 99 shapes used to build it. The results are shown in the second column of Table 2, giving the average Euclidean distance between corresponding points across all the shapes in the training sets. The training prostates can be all enclosed in a cube with side length of 5 cm. The dimension of the bilinear model is equal to 6 (the dimension of style control parameter) + J (the dimension of the content control parameter, which was set to be equal to the number of contents, C). The results show that for $J = 4, 5, 6$, the bilinear model fits the data quite well (*i.e.* to within 1.5 mm).

Next, the 105 prostate shapes from the 6 patients that were not included in building the bilinear model building were approximated. This test determines how well the model generalizes to unseen data. The results are shown in the third column of Table 2. The results show the same trend as in the first test, and the errors are only slightly higher.

4.3. Adaptation accuracy

In this experiment, we test the adaptability of the bilinear model to a new patient. We imagine that the model is applied in the following scenario. A new patient enters the clinic for radiation treatment (usually spaced over 20+ days, or “fractions”). The first M scans from the new patient are used to estimate the style parameter \mathbf{a} and content parameters $(\mathbf{b}_1, \dots, \mathbf{b}_M)$ based on the bilinear model matrices \mathbf{W}_k previously estimated offline. Subsequently, the patient-specific model with fixed \mathbf{a} and \mathbf{W}_k is used to approximate the shape at the remaining fractions. In our experiments, we learned the bilinear model matrices \mathbf{W}_k from the 6 training patients, and then did 6 adaptation experiments, using the first M scans from each testing patient to estimate the style parameters. The accuracy is measured as the average approximation error over the remaining shapes for each patient.

The dimension of the bilinear model is equal to the number of content control parameters, since the style control parameter is fixed after adaptation. Generally, we expect there to be an elbow point in approximation error as the number of adaptation datasets increases. That is, the modeling accuracy will increase as more shapes are used for the estimation of the style parameter, until there is relatively little to be gained by seeing new examples. For example, for different

numbers of content bins, Figure 9 shows how the projection error changes as the number of adaptation datasets increases for testing patient 12. In the figure, the elbow points occur at around $M = 4$ for models of dimensionality $J = 4, 5, 6$. This means that reasonable adaptation can be made using the first 4 new scans for this patient.

For comparison, we also did 6 adaptation experiments using a linear PCA model (*i.e.* using the 99 training datasets plus the first M shapes from each testing dataset) using the same number of control parameters as in the bilinear model. We compared the adaptability of both the linear and bilinear models for patient 12 in Figure 10, using a 4-dimensional parameter vector for each. The performance of the bilinear model improves as the the number of adaptation datasets increases, while the performance of the linear model does not improve much. This confirms the potential effectiveness of the bilinear model for adapting to a new patient. We also performed the adaptation experiment using $M = 1 - 8$ for the bilinear model with $J = 4$ (4 content bins) using testing sets 7–12. The point at which the minimum projection error occurs with respect to the number of sets used for adaptation is summarized as a histogram in Figure 11. It can be seen that adaptation after 4 to 6 scans seems to be reasonable based on this dataset.

4.4. Segmentation of medical images using the bilinear model

Finally, we applied the bilinear shape model to the segmentation of CT images for prostate localization. The bilinear prostate model with $I = 6$ and $J = 4$ was built using the 6 training datasets and then was adapted to a new patient (Set 12) using the first 4 contoured image sets. Thus, the dimension of the adapted model is $J = 4$, the number of content control parameters. The remaining image sets were segmented using the 3D model-based algorithm we described in [6]. This algorithm evolves the segmenting surface, parameterized by the bilinear content vector, to minimize the difference between the histogram of pixel intensities inside the surface and a desired histogram learned from training data. The surface is initialized as the average of the first 4 shapes used for adaptation. We find the best content vector by minimizing the histogram-matching cost function, and the segmentation result is the shape defined by the optimal content vector. See [6] for more details about the appearance model, cost function, and optimization methods used in the segmentation process. An example result for one patient’s scan is shown in Figure 12, as 2D slices of the full 3D segmentation. The yellow line is the hand-labeled contour by a physician and the black line is the automatic segmentation result. Our preliminary tests indicate that the bilinear shape model with 4 control parameters can be effectively used to segment prostate CT images.




Figure 9. Projection error versus the number of shapes used for adaptation (* indicates the minimum projection error in each case).

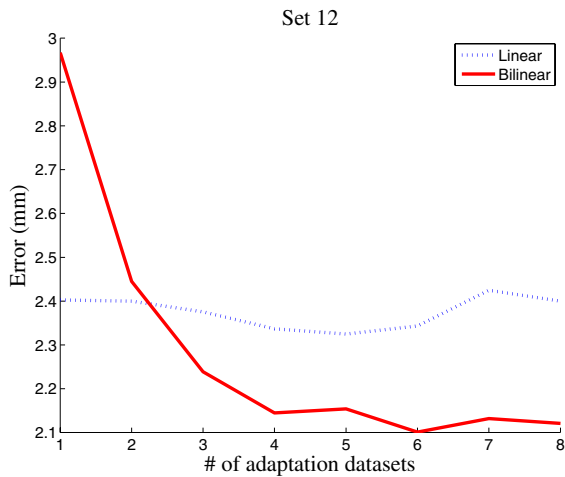


Figure 10. Adaptation accuracy of linear and style-adapted bilinear models, both with 4 control parameters.

5. Conclusions and future work

We demonstrated a novel application of bilinear models to the shape modeling of anatomical objects. The experimental results on prostate contours showed that the bilinear shape model performs well in fitting both training and testing shapes. In adaptation tests, the bilinear shape model performed better than a linear shape model using the same number of parameters. The results show the potential of bilinear models to compactly represent medical datasets that have two inherent types of variation. Our initial tests indicate that the adapted bilinear model can be used as an effective shape constraint in the segmentation of medical images.

When adapting the bilinear model to a new patient, it is not guaranteed that the control parameter estimation problem converges (or converges to a global minimum). However, the parameters always converged in our experiments

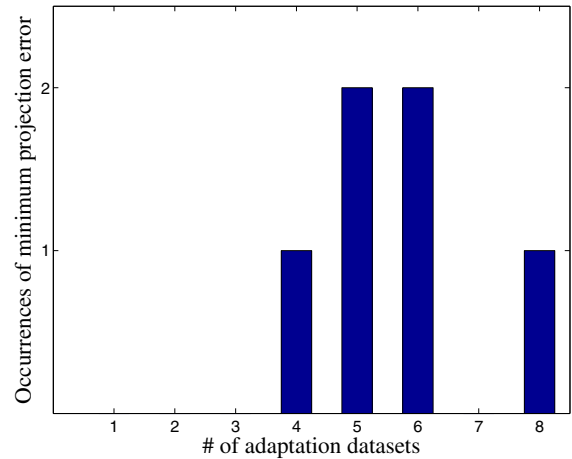


Figure 11. Histogram of minimum projection error occurrences w.r.t. the number of adaptation datasets for bilinear model with 4 control parameters.




Figure 12. 2D axial slices of the 3D segmentation result of one set of images from Patient 12 using the bilinear model (yellow: hand-labeled, black: segmentation result).

when the new patient showed similar variations as those in the training data. We believe a better shape model using the bilinear model is possible, since the assignment of content classes plays an important role in the bilinear model and could be improved. Specifically, we plan to investigate new ways of defining content and selecting content classes such that the bilinear shape model consistently outperforms linear shape models. Although we assumed two variations in our dataset, the number of true inherent variation axes may be more than two. In such case, a multilinear model [15] may express the variation better than bilinear model, although the difficulty of defining good content classes still remains. Finally, we plan to investigate the application of nonlinear multi-factor methods such as [5], although some such methods may be difficult to apply with very limited training data.

6. Acknowledgements

We wish to thank D. Michael Lovelock at Memorial Sloan-Kettering Cancer Center, New York, NY for providing us with clinical data. This work was supported in part by CenSSIS, the Center for Subsurface Sensing and Imaging Systems, under the Engineering Research Centers Program of the National Science Foundation (Award Number EEC-9986821).

References

- [1] B. Bascle and A. Blake. Separability of pose and expression in facial tracking and animation. *International Conference on Computer Vision*, pages 323–328, 1998. 2
- [2] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–366, 1994. 1
- [3] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models - Their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995. 1, 2
- [4] J. S. Duncan and N. Ayache. Medical image analysis: Progress over two decades and the challenges ahead. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:85–106, 2000. 1
- [5] A. Elgammal and C.-S. Lee. Separating style and content on a nonlinear manifold. *IEEE Conference on Computer Vision and Pattern Recognition*, 2004. 2, 7
- [6] D. Freedman, R. J. Radke, T. Zhang, Y. Jeong, D. M. Lovelock, and G. T. Y. Chen. Model-based segmentation of medical imagery by matching distributions. *IEEE Transactions on Medical Imaging*, 24:281–292, 2005. 1, 6
- [7] J. Huang and A. Amini. Anatomical object volumes from deformable B-spline surface models. *International Conference on Image Processing*, pages 732–736, 1998. 1
- [8] J. R. Magnus and H. Neudecker. *Matrix differential calculus with applications in statistics and econometrics*. Wiley, New York, 1988. 2
- [9] D. H. Marimont and B. A. Wandell. Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A: Optics, Image Science, and Vision*, 9:1905–1913, 1992. 2
- [10] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: A survey. *Medical Image Analysis*, 1(2):91–108, 1996. 1
- [11] S. M. Pizer, P. T. Fletcher, S. Joshi, A. Thall, J. Z. Chen, Y. Fridman, D. S. Fritsch, G. Gash, J. M. Glotzer, M. R. Jiroutek, C. Lu, K. E. Muller, G. Tracton, P. Yushkevich, and E. L. Chaney. Deformable m-reps for 3D medical image segmentation. *International Journal of Computer Vision*, 55, November 2003. 1
- [12] S. Soatto and A. J. Yezzi. Deformation: Deforming motion, shape average and the joint registration and segmentation of images. *Proceedings of the 7th European Conference on Computer Vision-Part III*, pages 32–57, 2002. 2
- [13] J. B. Tenenbaum and W. T. Freeman. Separating style and content with bilinear models. *Neural Computation*, 12(6):1247–1283, 2000. 2, 3
- [14] G. Turk, H. Q. Dinh, and J. F. O’Brien. Implicit surfaces that interpolate. *Proceedings of the 7th International Conference on Shape Modeling and Applications*, pages 62–71, 2001. 2
- [15] M. A. O. Vasilescu and D. Terzopoulos. Multilinear analysis of image ensembles: Tensorfaces. *Proceedings of the 7th European Conference on Computer Vision*, 2350:447–460, 2002. 2, 7
- [16] D. Vlasic, M. Brand, H. Pfister, and J. Popović. Face transfer with multilinear models. *ACM Transactions on Graphics*, 24:426–433, 2005. 2