

Improving Segmentation of the Left Ventricle using a Two-Component Statistical Model

Sebastian Zambal, Jiří Hladůvka, and Katja Bühler

VRVis Research Center for Virtual Reality and Visualization,
Donau-City-Strasse 1, 1220 Vienna, Austria
<http://www.vrvis.at>

Abstract. Quality of segmentations obtained by 3D Active Appearance Models (AAMs) crucially depends on underlying training data. MRI heart data, however, often come noisy, incomplete, with respiratory-induced motion, and do not fulfill necessary requirements for building an AAM. Moreover, AAMs are known to fail when attempting to model local variations. Inspired by the recent work on split models [1] we propose an alternative to the methods based on pure 3D AAM segmentation. We interconnect a set of 2D AAMs by a 3D shape model. We show that our approach is able to cope with imperfect data and improves segmentations by 11% on average compared to 3D AAMs.

1 Introduction

Correct segmentation of the left ventricle (LV) in cardiac MRI short-axis images is very important for further diagnosis of the heart's function. Manual segmentation is very time-consuming and therefore efforts are taken to automate this process to the largest possible extent.

Although variants of Active Appearance Models (AAMs) [2] have been applied to this problem [3,4,5], a robust and completely automatic segmentation still poses a challenge.

Several problems appear in the context of shape- and appearance-based modeling of the left ventricle:

- The appearance of papillary muscles and trabeculae varies irregularly from patient to patient. Especially in slices close to the heart's apex these fuzzy anatomical structures are hard to model statistically. The necessary requirement for building an AAM – that the underlying training data has a gaussian distribution – is not satisfied in the apical/apex region.
- Often, only a part of the left ventricle is captured. Training data consisting of poorly corresponding volumes decreases the quality of the resulting model.
- Respiratory-induced motion introduces shifts of neighboring slices and so leads to irregular local variances in the data.
- Principal Component Analysis (PCA) involved in the build leads to a model where variances apply to the complete model. This makes it hard to handle local variations, e.g. different brightness of individual slices.

In this paper we propose a new and robust approach to segmentation of the left ventricle. By linking a set of 2D AAMs with a simple 3D shape model we perform model matching on a global and local context iteratively. With this approach we handle the above list of problems.

This paper is structured as follows. In section 2 we review related work on statistical models of shape and appearance in context of LV segmentation. In section 3 we introduce a two-component model of the left ventricle. In section 4 we outline how the two-component model is matched to unseen data. Results are presented in section 5 and a conclusion is drawn in section 6.

2 Statistical Modeling of the Left Ventricle

Our approach benefits from the combination of both Active Shape Models (ASMs) [6] and Active Appearance Models (AAMs) [2]. ASMs are statistical models of shape that can be used for identification and segmentation of known objects in unknown images. AAMs are a direct extension of ASMs and incorporate not only information about shape but also model the underlying texture. Variations of both methods have successfully been applied to segmentation in medical imaging [4,7,8].

Statistical models of shape and appearance are created from a finite number of training data sets. In other words, the model is only able to represent and thus segment new data sets that are similar to the ones in the training set. Increasing the size of the training set is a rather limited way to improve the generalization ability of the model. By adding local deformability the generalization ability can be increased explicitly. The statistical model is used to find a rough initial segmentation which is refined by restricted local deformations [8].

Several authors have discussed local and global deformability of a model. Independent Component Analysis (ICA) [9,10] leads to modes of variation which have a rather local effect compared to the global modes of variation derived from the conventionally used PCA. To deal with local variations of a model it can explicitly be split into sub-models. The split is achieved either with the help of human expertise or by means of optimization based on the Minimal Description Length (MDL) [1].

A great problem that arises in segmentation of short-axis MRI images of the left ventricle is respiratory-induced motion. Such artifacts appear as slight displacements of spatially neighboring slices. Handling of these displacements is crucial to achieve both a correct model and correct segmentations. A recently proposed approach [11,12] first detects the left ventricle in one or multiple slices. In a second step image alignment is used to compensate the shifts of neighboring slices. Finally a 3D AAM search is performed to achieve the final segmentation.

3 A Two-Component Model of the Left Ventricle

Inspired by the idea of modeling global and local features separately [1,8] we propose to combine a set of local 2D AAMs with a global 3D shape model. The purpose of the 2D AAMs is to precisely match the well articulate slices of the




Fig. 1. Typical examples of (a) basal, (b) mid (the arrows indicate papillary muscles), (c) apical, and (d) apex slices (top) and the first two principal components of texture after PCA (bottom).

heart’s base. The purpose of the 3D shape model is twofold: to propagate the position and size of the basal slices to apical ones and to keep the global shape characteristics plausible. In the following we describe the two components of our model.

3.1 Component 1: A Set of 2D Active Appearance Models

In a first step we consider the given training data, i.e. texture and annotation, slice-wise. From base to apex we identify the following four classes of slices:

Basal slices (fig. 1(a)) reside close to the base of the heart and do not contain papillary muscles.

Mid slices (fig. 1(b)) clearly show papillary muscles (arrows).

Apical slices (fig. 1(c)) contain trabeculae and papillary muscles, which are represented by irregular texture.

Apex slices (fig. 1(d)) show the very apex of the heart and contain highly irregular texture.

We performed PCA on the shape-normalized texture vectors of all slices in the training set. Scatter-plots of the first two principal components are shown in figure 1 for the according classes of slices. We observe that basal and mid slices show a relatively compact gaussian distribution and are thus well suited for linear statistical modeling using PCA. This is not the case for apical and apex slices, whose texture is much more irregular.

With a manual classification of slices we build a set of four individual 2D AAMs. This proceeding has the advantage that even incomplete MRI data sets can be included in the training set and used for building the first component of the model.

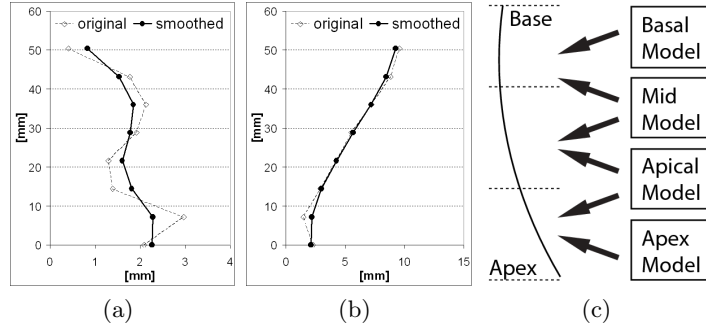


Fig. 2. The original and smoothed centerlines projected to XZ (a) and YZ (b) plane. (c) The assignment of sub-models to sections of the global shape model.




Fig. 3. The first (a), second (b), and third (c) modes of variation of the global shape model. The plots show the mean shape together with the largest possible deformations.

3.2 Component 2: A Global Shape Model

The intention of the global model is to represent the ventricle’s shape in an overall simplified way. Its task is to assure a valid relative placement and scaling of the four 2D AAMs.

The core of the global model is a centerline connecting the 2D AAMs in 3D. For the landmark points of each slice of the annotated training set a centroid is calculated. The assembly of these centroids forms the centerline. Figures 2(a) and 2(b) show such a centerline for one of the annotated data sets. It can be seen that respiratory-induced shifts of neighboring slices appear in the data. We use a gaussian filter to smooth the centerlines and to reduce respiratory motion artifacts in the training set.

Additionally to the centroids we incorporate the radii of slices into the global model. For each slice a radius as the mean distance of shape points to the centroid is calculated. This results in three variables that approximately describe the ventricle’s shape on one slice: x/y-coordinates and radius.

To build a statistical shape model, an equal number of corresponding landmarks has to be placed in every training example. We interpolate centroid positions and radii to get a total of 20 evenly-spaced slices. With the three features per slice this results in a total of 60 features included in the global shape model.

Before PCA can be applied, the individual training examples have to be aligned. Since we model not only the positions of centroids but also the radii we can not directly apply 3D Procrustes analysis. We rather align the center lines with respect to x/y-coordinates. The radii are thus not affected by the alignment. Figure 3 illustrates the first three modes of variation of the global shape model.

3.3 Combining Local Models and Global Model

In order to benefit from both components the local models have to be linked to the global model. As some data sets are delivered without apex or basal slices, we compensate this missing information by attaching two 2D AAMs to each slice. For example, the top most slice is assigned the basal model and the mid model. The better matching model is kept while the other one is ignored. Figure 2(c) shows the assignment of 2D AAM sub-models to the global shape model.

4 Matching the Two-Component Model

In the previous section we have outlined the idea of splitting a statistical model of the left ventricle into four parts and described how those parts are coupled over a global model. In this section we explain how the model iteratively is matched to unseen data.

4.1 Matching the Local Sub-models

We switch between two strategies. First a local matching of 2D AAMs is carried out for individual slices by standard AAM search [2]. In theory it should be possible to match all 2D AAMs in this way to get a valid global segmentation. While a valid match is obtained in basal and mid areas the 2D AAM search very often fails in the apical and apex slices. This is where the global model comes into play.

4.2 Updating the Global Model

Using the global component we propagate information from the better fitting slices to badly matching slices: The root mean square (RMS) texture error is calculated for all slices and the best 80% are used to update the global shape model. This update is done analogously to ASM search [6].

4.3 Iterative Global and Local Matching

After the global model has been updated the local AAMs are aligned to it. The search proceeds iteratively switching between 2D AAM search and update of the global model. In this way divergence due to bad initialization in critical slices is avoided and the overall segmentation gets improved. The loop of 2D and 3D updates is repeated until convergence.

The resulting segmentation still suffers from the lack of local variability. Further improvements are achieved performing additional 2D AAM search steps with position updates restricted by a constant maximum displacement.

Finally respiratory-induced shifts are compensated. The slices are shifted such that the centroids of the 2D AAMs align to the global model’s centroid coordinates. Compared to the work of Stegmann [11] our motion compensation thus benefits from prior knowledge that is encoded in the global shape model.

5 Results

We evaluated the described method with a set of 32 different short-axis studies. The quality of segmentation achieved with the two-component model was compared to that of 3D AAMs. Leave-one-out tests for all of the 32 data sets were performed. The average point-to-surface error (PSE) with respect to expert annotation was used to validate the segmentation quality. Figure 4 illustrates the measurements. Although the two-component model led to slightly worse results in 9 cases out of 32, the overall PSE got improved from an average of 2.20mm (3D AAM) to 1.96mm (the two-component model). Even though the average improvement by 0.24mm does not sound impressive, we emphasize that compared to the standard 3D AAM the two-component model performs better by 11%.

In figure 5 we provide a visual comparison for the largest improvement in terms of the average PSE which could be achieved for data set 10. It shows the result obtained with a standard 3D AAM (fig. 5(a)) and the improved segmentation achieved with the two-component model (fig. 5(b)).

Figure 6 demonstrates the motion compensation performed with the two-component model. A section of data set number 10 is depicted in figure 6(a).




Fig. 4. Point-to-surface error measured in leave-one-out tests for all training data sets.




Fig. 5. Results of matching data set 10 with 3D AAM (a) and two-component model (b).




Fig. 6. Motion compensation of data set 10 carried out by the two-component model. The original data (a) and the automatically corrected data (b) are shown.

Figure 6(b) shows the data set after matching and aligning the slices to the centerline of the global shape model.

Between 10 and 20 iterations are necessary for the two-component model matching algorithm to converge. The whole segmentation process typically takes a few seconds on standard PC hardware.

6 Conclusion, Discussion, Future Work

We have introduced a new approach to 3D segmentation of the left ventricle from short-axis MRI images by interconnecting a set of 2D AAMs with a simple 3D shape model. The two-component model is more robust than a 3D AAM since local irregularities such as respiratory-induced motion and different intensities of gray values in individual slices can be handled.

Incomplete training data sets that do not contain the entire LV from base to apex pose a severe problem for 3D AAMs. If such data sets are present in the training set the correspondence between them is ill-posed. The outcoming model allows for shrinking within its modes of variation. The topmost or the

bottommost slices thus often remain unsegmented. This disadvantage is avoided with the two-component model which allows the use of incomplete data sets as training examples. Since slices are assigned to individual 2D AAM sub-models even a single annotated 2D short-axis slice can be added to the training set. As the two-component model attempts to match two neighboring 2D AAMs to each slice, segmentation of data sets that do not cover the entire LV is still possible.

We are convinced that the idea of combining different types of models over a global shape model is promising for other medical segmentation problems. Especially if local anatomical abnormalities (e.g., a tumor) appear, models consisting of multiple components will likely lead to more precise results.

References

1. Langs, G., Peloschek, P., Donner, R., Bischof, H.: A clique of active appearance models by minimum description length. In: British Machine Vision Conference (BMVC). Volume 2. (2005) 859–868
2. Cootes, T.F., Edwards, G.J., Taylor, C.J.: Active appearance models. In: European Conference on Computer Vision (ECCV). Volume 2. (1998) 484–498
3. Stegmann, M.B., Pedersen, D.: Bi-temporal 3D active appearance models with applications to unsupervised ejection fraction estimation. In: International Symposium on Medical Imaging. Volume 5747. (2005)
4. Mitchell, S.C., Lelieveldt, B.P.F., van der Geest, R.J., Bosch, J.G., Reiber, J.H.C., Sonka, M.: Multistage hybrid active appearance model matching: Segmentation of left and right ventricles in cardiac MR images. *IEEE Transactions on Medical Imaging* **20**(5) (2001) 415–423
5. Mitchell, S.C., Bosch, J.G., Lelieveldt, P.F., van der Geest, R.J., Reiber, J.H.C., Sonka, M.: 3D active appearance models: Segmentation of cardiac MR and ultrasound images. *IEEE Transactions on Medical Imaging* **21**(9) (2002) 1167–1178
6. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models – Their training and application. *Computer Vision and Image Understanding* **61**(1) (1995) 38–59
7. Beichel, R., Bischof, H., Leberl, F., Sonka, M.: Robust active appearance models and their application to medical image analysis. *IEEE Transactions on Medical Imaging* **24**(9) (2005)
8. Taylor, C.J., Cootes, T.F.: Combining elastic and statistical models of appearance variation. In: European Conference on Computer Vision (ECCV). (2000) 149–163
9. Üzümcü, M., Frangi, A., Reiber, J., Lelieveldt, B.: The use of independent component analysis in statistical shape models. In: SPIE Medical Imaging. Volume 5032. (2003) 375–383
10. Suinesiaputra, A., Frangi, A.F., Üzümcü, M., Reiber, J.H.C., Lelieveldt, B.P.F.: Extraction of myocardial contractility patterns from short-axes MR images using independent component analysis. In: ECCV Workshops CVAMIA and MMBIA. (2004) 75–86
11. Stegmann, M.B., Ólafsdóttir, H., Larsson, H.B.W.: Unsupervised motion-compensation of multi-slice cardiac perfusion MRI. *Medical Image Analysis* **9**(4) (2005) 394–410
12. Stegmann, M.B., Larsson, H.B.W.: Motion-compensation of cardiac perfusion MRI using a statistical texture ensemble. In: Functional Imaging and Modeling of the Heart, FIMH 2003. Volume 2674. (2003) 151–161