A Bayesian Approach to Tracking Tagged Cardiac MR Images

Delman Lee, John T. Kent and Kanti V. Mardia
Department of Statistics, University of Leeds, LS2 9JT

Abstract

In cardiac magnetic resonance imaging, the heart can be "tagged" at the start of the heart cycle, partitioning the heart into a set of quadrilaterals (quads). The objective is to track the heart through time using the tags to measure the deformation of the heart. A Bayesian method is proposed to carry out the tracking. First, the heart image at each time t is pre-processed to give a set of "potential quads". A prior distribution on the quads at time t is specified to ensure regularity with respect to the structure of the quads and to ensure similarity to their positions at time $t-1$. A likelihood is specified for the generation of the potential quads given the true quads. A posterior estimate of the quad structure at time t is obtained by a Markov Chain Monte Carlo algorithm.

1 Introduction

A non-invasive method for monitoring the deformation of the heart wall is by tagged Magnetic Resonance (MR) imaging (Axel and Dougherty, 1989). A typical tagged MR image of the left ventricle (long-axis) is shown on the left of Fig. 1. The dark lines forming the grid are "tagged" to the material points, and as such follow the material points through time. In subsequent time frames, the contrast between the tagged and non-tagged materials decreases, and, furthermore, lines tagged to fluid material (e.g. blood) disappear. Tracking the grid through time provides deformation information that is useful for diagnosis of certain heart diseases.

Many past tracking techniques for tagged MR images (Guttman et al., 1994; Kumar and Goldgof, 1994; Kraitchman et al., 1995; Radeva et al., 1996) follow the dark tagged lines through time with deformable splines. In this paper we investigate the dual problem of tracking the network of (approximate) quadrilaterals formed by the grid lines, rather than the lines themselves. *Quadrilateral* will hereafter be shortened to *quad*.

2 The Bayesian Framework

Tracking of tagged MR image sequences is done frame-by-frame. For each frame, a quad detector is run over the image to give a set of "potential quads". Associated with each potential quad

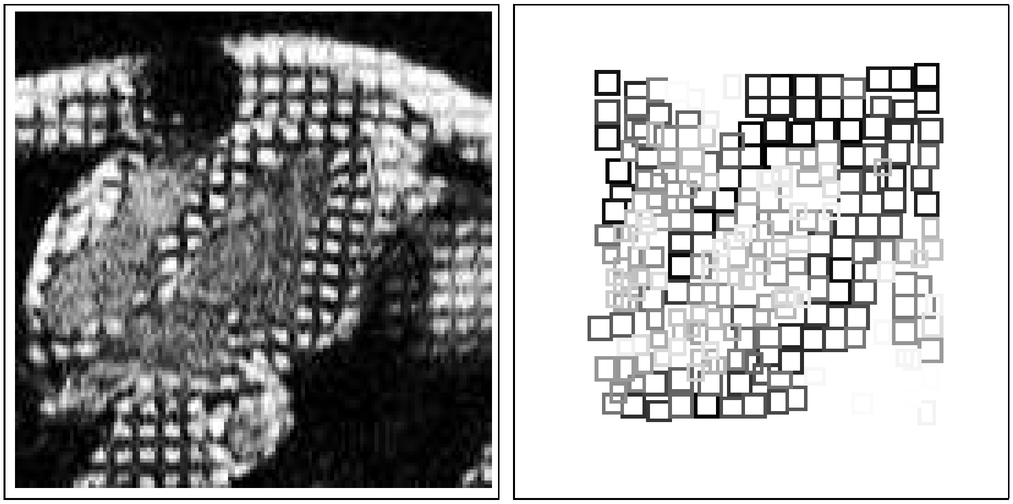

Figure 1: The right picture shows the potential quads generated by the quad detector on the left image (the third frame of a sequence). The darker the potential quad, the higher the certainty.

is a "certainty" which measures how sure we are of the existence of the quad. Quads are picked from the set of potential quads to form a "quilt", denoted by x.

A prior $f(x)$ is specified to govern how the quads should be joined up to form the quilt x. The prior enforces our belief of what kinds of quilts are plausible (§ 2.1). Let y denote the potential quads and their certainties collectively. A likelihood function $f(y, \mu|x)$ is specified for the detection of the potential quads and their certainties (and some unobserved data μ) from a quilt x (§ 2.2). A Bayesian strategy is used to infer about the quilt x and the unobserved data μ from the observed data y. Information about the model and the unobserved data are contained in the posterior distribution given by $f(x, \mu|y) \propto f(y, \mu|x) f(x)$. Samples from the posterior density are simulated by the Markov Chain Monte Carlo (MCMC) method (§ 3) (Besag et al., 1995). The posterior mode is then estimated from the samples.

Notations

Let a quad be represented by $q_i = (c_i, v_i)$, where c_i is the centre of the quad, and v_i is the set of the four vertices of the quad. The centre of the quad c_i is not strictly necessary, but it is kept in the representation for convenience. The set of P potential quads and their certainties are denoted by $\check{q} = \{\check{q}_p\}_{p=1}^{P}$, $\check{q}_p = (\check{c}_p, \check{v}_p)$ and $e = \{e_p\}_{p=1}^{P}$ $(0 < e_p < 1)$, respectively. Fig. 1 shows a sample output of the quad detector. The set of potential quads and their certainties $y = (\check{q}, e)$ is our observed data.

Let $\mathscr{G}$ be an n-node undirected graph associated with a subset of a 2D square lattice. Nodes in the graph are indexed by $i = (i_1, i_2) \in \mathbb{Z}^2$, and are connected with 8-adjacency. With $\mathscr{G}$ specified, the components of our model quilt x are (i) the maximum number of quads, n, that can make up a quilt, (ii) the positions and vertices of the quads, $q = \{q_i\}_{i \in \mathscr{G}}$, and (iii) the statuses, dead $(s_i = 0)$ or alive $(s_i = 1)$, of the quads, $s = \{s_i\}_{i \in \mathscr{G}}$. Two quads q_i and q_j are neighbours if the nodes i and j share an edge in the graph $\mathscr{G}$, and we write $i \sim j$ in such a case.

The set of live quads in x forms the "fabric" of the quilt, while the set of dead quads corresponds to "holes" in the quilt. The live quads account for visible tagged materials, for example heart wall, and dead quads account for tagged materials whose structure is no longer apparent, for example blood.

2.1 Quilt Prior

We model the prior for the quilt $x=(q,s)$ by $f(x)=f(q)f(s)$, where q and s are assumed to be independent. For the prior on the status s, we assume an Ising model truncated to have support on $\mathscr{S}$, where $\mathscr{S} \subset \{0,1\}^n$ is the set of configurations such that the live quads form one connected component. The truncated Ising model encourages neighbouring nodes to take similar values, discouraging a checkerboard pattern for s. We model the prior on the quads by

$$
\begin{aligned}
f(q) \propto \; & \exp\left(-\frac{1}{2}\sum_{i\in\mathscr{G}} \Delta_i^T R_i \, \Delta_i\right) \times \exp\left(-\frac{1}{2}\sum_{i\sim j} \Delta_j^T S_{ij} \, \Delta_i\right) \times \\
& \prod_{\substack{i\sim j,k \\ j\neq k}} \mathrm{I}\left[\cos(\psi_{jik} - \theta_{jik}) > \cos\phi\right] \times \exp\left(-\gamma \sum_{i,j\in\mathscr{G}} q_i \cap q_j\right)
\end{aligned}
\tag{1}
$$

with the following notations.

Based on information from previous frames, we have a set of n predicted quad positions for the current frame $\hat{c}=\{\hat{c}_i\}_{i\in\mathscr{G}}$ and a set of n 2×2 covariance matrices Σ_i of the prediction error. Let $\Delta_i = c_i - \hat{c}_i$ be the deviation from the predicted position of the centre of the quad i. The first factor in (1) enforces our belief that the positions of the quads are not too far away from their predicted values, while the second factor in (1) encourages departures from predicted positions to be similar for neighbouring quads. The matrices R_i and S_{ij} are chosen such that the covariance matrix of the resultant quadratic form, when the first and second factors are combined, is positive definite.

Let ψ_{jik} denote the angle formed by the three points (c_j, c_i, c_k) with c_i at the apex, and θ_{jik} be the predicted angle formed by the centres of the three quads (q_j, q_i, q_k) with q_i at the apex. Thus, with $\mathrm{I}[\cdot]$ denoting the indicator function, ϕ in the third factor of (1) is the allowed deviation from the expected angle θ_{jik}. For an 8-adjacency graph, we chose $\phi=\pi/4$. The third factor in (1) aims to preserve the regularity of the quilt, avoiding tears.

Let $q_i \cap q_j$ denote the area of intersection of the interior of the two quads q_i and q_j. With $\gamma>0$, the fourth factor in (1) penalizes against intersecting quads, and allows quads to be packed more closely in regions of small quads. This is the only factor in (1) which depends on the vertices of the quads.

Design Issues

Note that the prior $f(q)$ in (1) has the Markov property with neighbourhoods up to second order, and thus allowing efficient implementation for the MCMC simulation in § 3. However, there is a non-local operation for some types of transitions in the MCMC algorithm. This is due to the restriction of Ising model to the set of configurations where the live quads form a single connected component. For certain types of transitions in the MCMC algorithm, we need to find the articulation points[1] of the graph formed by the live quads. This is found by a depth-first search,

[1] An articulation point of a singly-connected graph is a node, whose removal results in two separate components.

which is a non-local operation of $O(N, E)$, where N and E are the number of nodes and edges in the graph.

The motivation for requiring the live quads in the quilt to remain a single connected piece is to avoid many disconnected pieces of quilt, each misaligned with respect to one another. We have considered the notion of *local connectivity* to solve this problem. Consider the graph associated with a local l-neighbourhood of a node (all nodes reachable by l steps or less). A node is a *local articulation point* if its removal disconnects the l-neighbourhood graph into two separate connected components. The local connectivity idea works well when there are no large holes in the quilt. For our tagged MR images of the left ventricle, where there is a large region of without tags, the local connectivity approach is impractical.

Note also that the sums and products in (1) are over both live and dead quads. A more suitable prior $f(q|s)$ might be one where the sums and products are only over the live quads (the prior becomes conditional on s). However, the normalization $Z(s)$ for such a prior $f(q|s)$ comes into play in the MCMC algorithm. To be precise, one needs the ratio of the normalizations $Z(s)/Z(s')$ in the acceptance probability, where s and s' are respectively the current and proposed statuses. This ratio is intractable and one may have to resort to approximations like the pseudo-likelihood to make progress.

2.2 Incomplete Data and Likelihood

We will only provide an outline of the likelihood terms here, see (Lee et al., 1997b) for details. Consider a hypothetical quad detector which takes a quilt x as input and generates a list of P potential quads $\check{q}$, their certainties e, and a mapping $\mu(\cdot)$. The function $\mu(\cdot)$, which maps $\mathscr{G}$ to $\{0, 1, \cdots, P\}$, describes which potential quad in $\check{q}$ corresponds to which model quad in x. Further $\mu(i)=0$ means that quad i does not correspond to any potential quads in $\check{q}$.

We assume that all live quads in x are detected exactly once, i.e. the mapping from the set $\{i : \mu(i) > 0\}$ to $\{1, \cdots, P\}$ is injective. The actual quad detector takes an image as input and generates the potential quads $\check{q}$ and their certainties e. The mapping μ is unobserved and represents missing data, while $y=(\check{q}, e)$ are the observed, incomplete data.

To specify the likelihood for (y, μ) given x, we need to answer the question: given a quilt x, what is the probability of the hypothetical quad detector generating potential quads $\check{q}$, certainties e, and mapping $\mu(\cdot)$? We model our likelihood function $f(y, \mu|x)$ as a product of three factors: $f(y, \mu|x) = f(\check{q}, e, \mu|x) = f(\check{q}|\mu, x)f(e|\mu, x)f(\mu|x)$, where $\check{q}$ and e are assumed to be independent given μ and x. The three factors are modelled respectively by:

1. Given the quilt x and the mapping μ, the quad detector is assumed to detect all live quads in x precisely, and gives out $P - n_{\text{alive}}$ artefactual potential quads, where n_{alive} is the number of live quads. The artefactual potential quads are assumed to be uniformly distributed on the space of possible outputs of the quad detector.

2. Given the quilt x and the mapping μ, the potential quad that corresponds to a live quad should have a certainty close to unity, and one that does not correspond to any live quads should have a certainty close to zero. We model these two cases with beta densities. As-

suming independence between the certainties of different potential quads, $f(e|\mu, x)$ is thus a product of P independent beta densities.

3. Given the quilt x, we place a uniform weight on all permissible mappings μ.

3 Posterior Sampling

In order to gain information about the posterior, we used the MCMC method with a Metropolis-Hastings type algorithm that allows jumps between different dimensional spaces (Geyer and Møller, 1994; Green, 1995). The dimension changing aspect of our model comes from sites being alive or dead. An outline of the algorithm is provided below.

At each MCMC iteration, the algorithm proposes one of four types of transitions: birth, death, live-displacement and dead-displacement. The live- and dead-displacement transitions displace the positions of a live and dead quad, respectively. The birth and death transitions toggle the statuses of a dead and live quad, and propose a new position for them. The birth and death transitions are constrained so that the live quads in the quilt remain a single connected component. The four types of transitions are proposed with equal probabilities. For each type of transition, a node i is chosen randomly, and new values for (q_i, s_i) and $\mu(i)$ are proposed with a certain probability. The proposed values are then accepted with a certain acceptance probability. The proposal distributions and acceptance probabilities for the four types of transitions are chosen to satisfy detailed balance so that the Markov chain has the desired stationary distribution.

A greedy version of the Markov chain (similar to the Iterative Conditional Mode (ICM) of (Besag, 1986)) is run where only births with potential quads nearest to the predicted position of the nodes are proposed, and the proposed birth is accepted only if the posterior density is increased. The quilt from such a greedy Markov chain is the initial state of the full MCMC algorithm. Samples from the Markov chain are collected after a burn-in period.

The posterior is summarized by its "restricted" marginal modes, arrived at by the following procedure. During sampling, for each live quad in the quilt, a histogram of which potential quad it has picked is maintained. At the end of the sampling, we look at quads that have been alive for over a certain percentage of the time; we call these quads the "resilient quads". We scan through the resilient quads in an order based on the entropies of their histograms. Resilient quads with lowest entropies are scanned first. For each resilient quad, the potential quad that MCMC has picked most often is assigned to it, provided that this modal potential quad has not been assigned to another resilient quad. After the assignment process, the set of quads which forms the largest single connected component is our estimate of the posterior mode.

Averages of the following quantities are inappropriate summary statistics. For example, since the indexing of our potential quads is arbitrary, the average of the histogram is meaningless. For each node, one can find the average position of all the potential quads it has picked during sampling. However, such an average is not a useful quantity in our case where the distribution is multi-modal and an average position may not correspond to any real quads in the image. The question of how best to summarize the sampled quilts remains open.

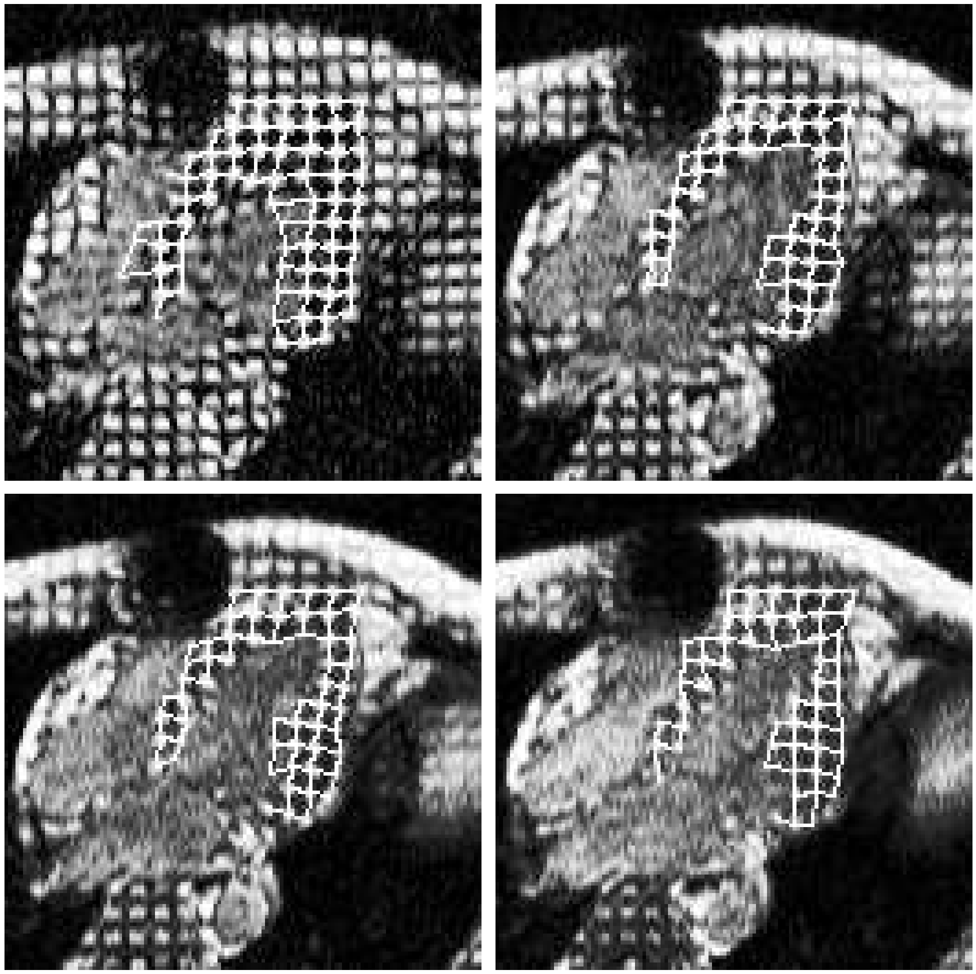

Figure 2: The graph structure of the network of quads for four selected frames of a nine-frames long-axis sequence of the left ventricle. For clarity, only the 4-adjacent edges of the graph are drawn.

4 Results & Discussion

The successful tracking in a region of interest of two image sequences are shown in Figs. 2 & 3, where the nodes of the graph correspond to the centres of the quads. Four selected frames of a long-axis sequence are shown. The results are for a MCMC simulation with 300 sweeps, of which 100 are for burn-in. The tracking for the nine frames took 5mins on an UltraSPARC-167MHz. Note the successful tracking in Fig. 2 where a noticeable proportion of the live quads dies through time.

The number of burn-in sweeps of the MCMC is arrived at by inspection. Many convergence diagnostics have been proposed recently (see e.g. (Cowles and Carlin, 1996)) and will be incorporated into the algorithm. Related to the rate of convergence is the mixing property of the Markov chain. Our current algorithm proposes single-node transitions. Transition that updates a block of nodes at a time should improve the mixing property of the chain.

Since there is no substantial inter-frame motion, the following approximations work satisfactorily: (i) the estimated positions of quads from the previous frame is used as the initial estimate

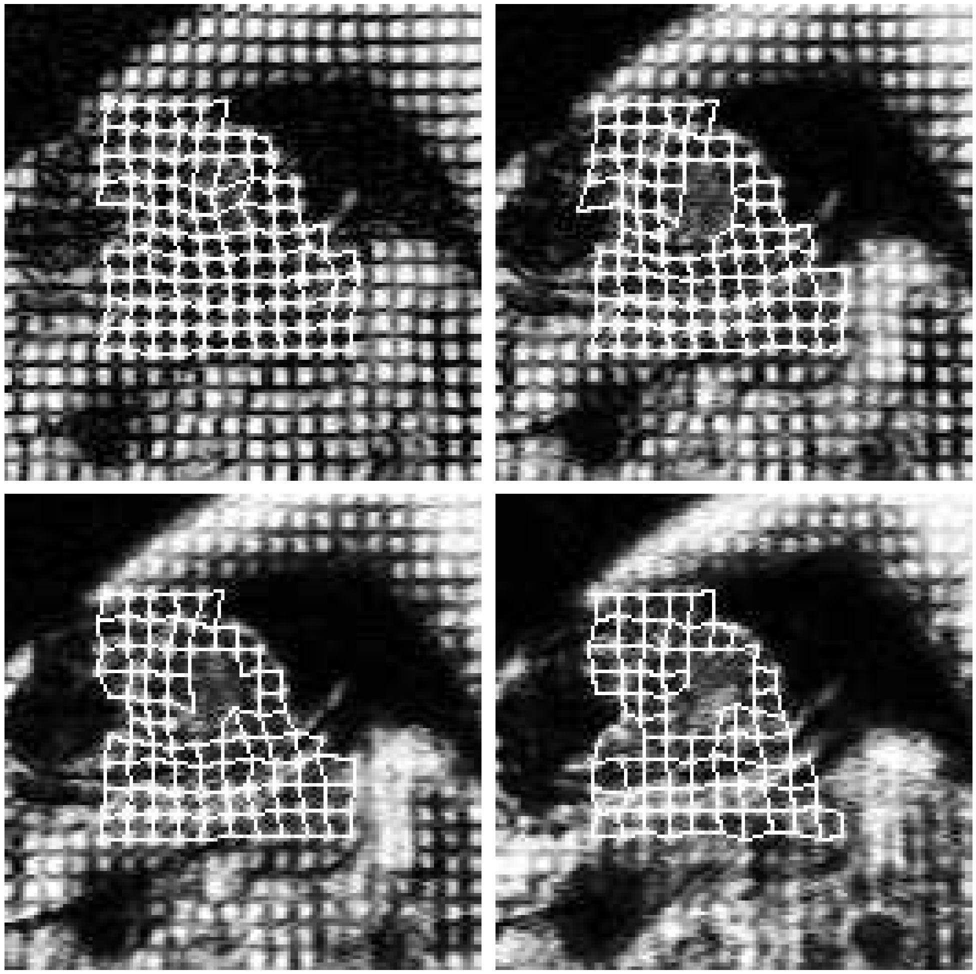

Figure 3: The graph structure of the network of quads for four selected frames of a nine-frames short-axis sequence of the left ventricle. For clarity, only the 4-adjacent edges of the graph are drawn.

in the current frame, and (ii) a quad detector that handles only mild deformation from a square. If the inter-frame motion becomes severe, more sophisticated motion prediction and quad detection will be necessary. The paper has focussed on the problem of tracking the tagged grid. The next step is to deduce the full 3D motion of the left ventricle, see e.g. (Park et al., 1996). With current MR tagging techniques, motion perpendicular to the acquisition plane (off-plane motion) are not accounted for. Image sequences acquired at orthogonal planes are required to resolve the 3D motion. We have tackled the through-plane problem in (Lee et al., 1997a).

Acknowledgments

The work is supported under the EPSRC Stochastic Modelling in Science and Technology programme. We would like to thank Liz Berry, Bill Crum, John Ridgway and U. Sivananthan for their clinical input, and EPSRC visiting Research Fellows Andre Guéziec, Valen Johnson and Demetri Terzopoulos for their useful comments.

References

Axel, L. and Dougherty, L. (1989). Heart wall motion: Improved method of spatial modulation of magnetization for MR imaging. *Radiology*, 172:349–350.

Besag, J. (1986). On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society, B*, 48(3):259–302.

Besag, J., Green, P., Higdon, D., and Mengersen, K. (1995). Bayesian computation and stochastic systems. *Statistical Science*, 10(1):3–66.

Cowles, M. and Carlin, B. (1996). Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434):883–904.

Geyer, C. and Møller, J. (1994). Simulation procedures and likelihood inference for spatial point processes. *Scandinavian Journal of Statistics*, 21:359–373.

Green, P. (1995). Reversible jump Markov chain Monte-Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732.

Guttman, M., Prince, J., and McVeigh, E. (1994). Tag and contour detection in tagged MR images of the left ventricle. *IEEE Transactions on Medical Imaging*, 13(1):74–88.

Kraitchman, D., Young, A., Chang, C., and Axel, L. (1995). Semi-automatic tracking of myocardial motion in MR-tagged images. *IEEE Transactions on Medical Imaging*, 14(3):422–433.

Kumar, S. and Goldgof, D. (1994). Automatic tracking of SPAMM grid and the estimation of deformation parameters from cardiac MR images. *IEEE Transactions on Medical Imaging*, 13(1):122–132.

Lee, D., Kent, J., and Mardia, K. (1997a). The through-plane problem for tagged MR images. In *Medical Image Understanding and Analysis*, Oxford.

Lee, D., Kent, J., and Mardia, K. (1997b). Tracking of tagged MR images. Technical Report STAT 97/08, Department of Statistics, University of Leeds, Leeds, LS2 9JT, U.K. A shorter version appears in the proceedings of XVth International Conference on Information Processing in Medical Imaging, 1997.

Park, J., Metaxas, D., and Axel, L. (1996). Analysis of left ventricular wall motion based on volumetric deformable models and MRI-SPAMM. *Medical Image Analysis*, 1(1):53–71.

Radeva, P., Amini, A., Huang, J., and Martí, E. (1996). Deformable B-solids and implicit snakes for localization and tracking of SPAMM MRI-data. In *Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 192–201.