

Tracking of Tagged MR Images by Bayesian Analysis of a Network of Quads

Delman Lee, John T. Kent, and Kanti V. Mardia

Department of Statistics, University of Leeds, Leeds, LS2 9JT, U.K.

Abstract. Automatic tracking of tagged MR image sequences is done frame-by-frame. For each frame, a quadrilateral (quad) detector is run over the image to give a set of “potential quads”. A likelihood function is specified for the detection of potential quads from an image. Quads are picked from the set of potential quads to form a “quilt”. Quads are present where a grid structure is apparent in the image. A prior is specified to govern how the quads should be joined up to form the quilt. The prior for the quilt (i) encourages quads to be close to their positions predicted from the last frame, (ii) encourages neighbouring quads to be close to each other, (iii) discourages intersecting quads, (iv) avoids “tears” in the quilt, and (v) encourages connectedness of quads. With the likelihood and prior densities, a Bayesian analysis is carried out using the Markov Chain Monte Carlo method on the posterior density to give an estimate of the posterior mode.

1 Introduction

A non-invasive method for monitoring the deformation of the heart wall is by tagged Magnetic Resonance Imaging (MRI) [1]. A typical tagged MRI frame 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 (in planes perpendicular to the imaging plane), 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. This paper addresses the issue of tracking the tagged grid through time.

Many past tracking techniques for tagged MRI images [2–5] 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 Principles of the Method

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”. Quads are picked from the set of potential quads to form a “quilt”. The quads in the quilt are connected by a square graph structure and can be either alive or dead. For a particular frame, live quads account for tagged materials, for example heart wall, that are visible, and dead quads




Fig. 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.

account for tagged materials whose structure is no longer apparent, for example blood. A prior $f(x)$ is specified to govern how the quads should be joined up to form the quilt x (§ 2.2). A likelihood function $f(y, \mu|x)$ is specified for the detection of potential quads y (and some unobserved data μ) from a quilt x (§ 2.3).

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 $f(x, \mu|y) \propto f(y, \mu|x)f(x)$. Inference about the posterior density is done by sampling using the Markov Chain Monte Carlo (MCMC) method (§ 3) [6]. The MCMC method is chosen in favour of a deterministic method to avoid the solution being trapped in minor subsidiary modes of a “bumpy” posterior.

2.1 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. For each frame, the image is pre-processed by a quad detector to give a set of P “potential quads” $\tilde{q} = \{\tilde{q}_p\}_{p=1}^P$, $\tilde{q}_p = (\tilde{c}_p, \tilde{v}_p)$ and their “certainties” $e = \{e_p\}_{p=1}^P$. The certainty $0 < e_p < 1$ is a measure of how sure we are of the existence of the quad \tilde{q}_p . Fig. 1 shows a sample output of the quad detector. The set of potential quads and their certainties $y = (\tilde{q}, e)$ is our observed data.

Let \mathcal{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. The extent of the square lattice is specified by the user in the first frame. In subsequent frames, \mathcal{G} is the output from processing of the previous frame.

Once \mathcal{G} is specified at a time frame, the components of our model x are (i) the maximum number of quads, n , that can make up a quilt, (ii) the positions and vertices of the quads, $q = \{q_i\}_{i \in \mathcal{G}}$, and (iii) the statuses, dead ($s_i = 0$) or alive ($s_i = 1$), of the quads, $s = \{s_i\}_{i \in \mathcal{G}}$. Two quads q_i and q_j are neighbours if the nodes i and j share an edge in the graph \mathcal{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. Let \mathcal{L} denote the set of indices of the live quads in x , and $n_{\text{alive}} = \sum_{i \in \mathcal{L}} s_i$ the number of live quads. Thus the model $x = (q, s)$ (taking the underlying graph structure \mathcal{G} for granted) defines a quilt in a certain region.

2.2 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, we assume an Ising model truncated to have support on \mathcal{S} , where $\mathcal{S} \subset \{0, 1\}^n$ 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

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

with the following notations.

Based on information from previous frames, we have a set of n predicted quad positions for the current frame $\hat{c} = \{\hat{c}_i\}_{i \in \mathcal{G}}$ and a set of $n \ 2 \times 2$ 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 $\mathbb{I}[\cdot]$ 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.

2.3 Incomplete Data and Likelihood

Consider a hypothetical quad detector which takes a quilt x as input and generates a list of P potential quads \tilde{q} , their certainties e , and a mapping $\mu(\cdot)$. The function $\mu(\cdot)$, which

maps \mathcal{G} to $\{0, 1, \dots, P\}$, describes which potential quad in \tilde{q} corresponds to which model quad in x . Further $\mu(i)=0$ means that quad i does not correspond to any potential quads in \tilde{q} . Let $z(\mu)$ be the number of zeros in $\{\mu(i)\}_{i \in \mathcal{G}}$. We assume that all live quads in x are detected exactly once, i.e. the mapping from the set $\mathcal{L}=\{i : \mu(i) > 0\}$ to $\{1, \dots, P\}$ is injective. The actual quad detector takes an image as input and generates the potential quads \tilde{q} and their certainties e . The mapping μ is unobserved and represents missing data, while $y=(\tilde{q}, e)$ 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 \tilde{q} , certainties e , and mapping $\mu(\cdot)$? We model our likelihood function $f(y, \mu|x)$ as a product of three factors: $f(y, \mu|x) = f(\tilde{q}, e, \mu|x) = f(\tilde{q}|\mu, x)f(e|\mu, x)f(\mu|x)$, where \tilde{q} and e are assumed to be independent given μ and x .

Given the quilt x , we place a uniform weight on all permissible mappings μ . Thus $f(\mu|x)$ is a discrete distribution with probability $(P - n_{\text{alive}})!/P!$ if $z(\mu) = P - n_{\text{alive}}$, and 0 otherwise. 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 $e^2/3$ and $(1-e)^2/3$ for $0 < e < 1$, respectively. Assuming independence between the certainties of different potential quads, $f(e|\mu, x)$ is thus a product of P independent beta densities.

Finally, 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. The artefactual potential quads are assumed to be uniformly distributed on the space of possible outputs of the quad detector. We denote the ‘‘volume’’ of such 8-dimensional space by V . We thus have $f(\tilde{q}|\mu, x) = V^{-z(\mu)} \prod_{i \in \mathcal{L}} \mathbb{I}[\tilde{q}_{\mu(i)} = q_i]$. Note that an artefactual potential quad \tilde{q}_i has a probability density on \mathbb{R}^8 if $s_i=0$, and has a point mass if $s_i=1$.

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 [7, 8]. The dimension changing aspect of our model comes from sites being alive or dead. Due to page limitations, only an outline of the algorithm is provided here.

At each MCMC iteration, the algorithm proposes one of four types of transitions: live-displacement, dead-displacement, birth and death. 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 (feasible) dead and live quad, and propose a new position for them, respectively. 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 from the feasible nodes, and new values for the component i of the quilt (q_i, s_i) and for the component i of the mapping $\mu(i)$ 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 carefully chosen to satisfy detailed balance so that the Markov chain has the desired stationary distribution.

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. For a fixed node $i \in \mathcal{G}$, consider the histogram of the mapping μ , $h_{\mu(i)}(p)$, whose entry is the frequency that the p th potential quad is picked by quad i in the quilt during simulation. Picking a potential quad of 0 means that the quad is dead. For all the nodes that have been alive for more than a certain percentage of iterations of the chain (10% in our simulations), rescale the histogram $h_{\mu(i)}(p)$ to sum up to 1 for $p \neq 0$. We scan through the nodes according to the entropies of their rescaled histograms. Nodes with lower entropies are scanned first. For each scanned node i , the potential quad \check{q}_p is picked, where p is the mode of the rescaled histogram. Node i is then born if \check{q}_p is not already associated with a live node. After scanning through all the nodes, the set of live quads which forms the largest single connected component is our estimate of the quilt for the current frame.

4 Results & Discussion

The successful tracking in a region of interest of an image sequence is shown in Fig. 2, 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 20mins and 5mins on an i486-133MHz and UltraSPARC-167MHz, respectively. 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. [9]) 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 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 accounting for “through-plane motion” (see e.g. [10]).

References

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

Acknowledgment: 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 useful discussions and for providing us with the data set.




Fig. 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.

2. M. Guttman, J. Prince, and E. McVeigh, "Tag and contour detection in tagged MR images of the left ventricle," *IEEE Trans. Med. Imaging*, **13**:74–88, 1994.
3. S. Kumar and D. Goldgof, "Automatic tracking of SPAMM grid and the estimation of deformation parameters from cardiac MR images," *IEEE Trans. Med. Imaging*, **13**:122–132, 1994.
4. D. Kraitchman, A. Young, C. Chang, and L. Axel, "Semi-automatic tracking of myocardial motion in MR-tagged images," *IEEE Trans. Med. Imaging*, **14**:422–433, 1995.
5. P. Radeva, A. Amini, J. Huang, and E. Martí, "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*, 192–201, June 1996.
6. J. Besag, P. Green, D. Higdon, and K. Mengersen, "Bayesian computation and stochastic systems," *Statistical Science*, **10**:3–66, 1995.
7. C. Geyer and J. Møller, "Simulation procedures and likelihood inference for spatial point processes," *Scandinavian J. of Statistics*, **21**:359–373, 1994.
8. P. Green, "Reversible jump Markov chain Monte-Carlo computation and Bayesian model determination," *Biometrika*, **82**:711–732, 1995.
9. M. Cowles and B. Carlin, "Markov chain Monte Carlo convergence diagnostics: A comparative review," *J. of the American Statistical Association*, **91**:883–904, 1996.
10. J. Park, D. Metaxas, and L. Axel, "Analysis of left ventricular wall motion based on volumetric deformable models and MRI-SPAMM," *Medical Image Analysis*, **1**:53–71, 1996.