

Trans-Dimensional MCMC Methods for Fully Automatic Motion Analysis in Tagged MRI

Ihor Smal, Noemí Carranza-Herrezuelo, Stefan Klein,
Wiro Niessen, Erik Meijering

Biomedical Imaging Group Rotterdam,
Departments of Medical Informatics and Radiology
Erasmus MC – University Medical Center Rotterdam, The Netherlands

Abstract. Tagged magnetic resonance imaging (tMRI) is a well-known noninvasive method allowing quantitative analysis of regional heart dynamics. Its clinical use has so far been limited, in part due to the lack of robustness and accuracy of existing tag tracking algorithms in dealing with low (and intrinsically time-varying) image quality. In this paper, we propose a novel probabilistic method for tag tracking, implemented by means of Bayesian particle filtering and a trans-dimensional Markov chain Monte Carlo (MCMC) approach, which efficiently combines information about the imaging process and tag appearance with prior knowledge about the heart dynamics obtained by means of non-rigid image registration. Experiments using synthetic image data (with ground truth) and real data (with expert manual annotation) from preclinical (small animal) and clinical (human) studies confirm that the proposed method yields higher consistency, accuracy, and intrinsic tag reliability assessment in comparison with other frequently used tag tracking methods.

Key words: Tracking, motion analysis, tagged MRI, cardiac imaging, particle filtering, Markov chain Monte Carlo

1 Introduction

In magnetic resonance imaging (MRI), tissue tagging by spatial modulation of magnetization (SPAMM) [2] has shown great potential for studying myocardial motion and treatment effects after myocardial infarction. Contrary to conventional MRI, tagged MRI (tMRI) allows for quantitative assessment of the myocardium, for example by analyzing regional strain. With tMRI, two orthogonal sets of magnetic saturation planes, each orthogonal to the image plane, can be created in short time. The deformation of the resulting tag pattern over time reflects the deformation of the underlying cardiac tissue, and is of clinical significance for assessing dynamic properties of the heart. Existing methods, based on active contours [1, 5], non-rigid image registration [4, 7], and optical flow [3], are not robust or accurate enough to deal with the strongly varying image quality of typical experimental data. The inevitable fading of the tag pattern and acquisition noise (Fig. 1) also affect commercially available tracking methods, based on harmonic phase (HARP) MRI [12].




Fig. 1. Examples of images from studies of rat hearts, with frames 6, 12, and 18 (out of 24 per heart cycle) showing the fading of the tagging.

In this paper we propose a novel Bayesian method for motion analysis in tMRI, implemented by means of particle filtering (PF), and a trans-dimensional Markov chain Monte Carlo (TDMCMC) approach. Using a specially designed tag likelihood measure and heart motion model (based on non-rigid image registration), the method allows for more robust and accurate tag position estimation in a sequence of noisy images. Moreover, TDMCMC makes statistical inference about both the *existence* and *position* of hundreds of appearing and disappearing tags possible and tractable. We demonstrate how tag existence can be used for automatic detection of left ventricular (LV) contours (or the myocardium). The performance of the proposed method is evaluated using realistic synthetic image data of several types, and real data from both preclinical (small animal) and clinical (human) experiments.

2 Proposed Tracking Framework

2.1 Particle Filtering and Importance Sampling

Bayesian estimation infers the posterior distribution $p(\mathbf{s}_t|\mathbf{z}_{1:t})$ of the unobserved state \mathbf{s}_t of a system (the tag pattern in our case), which changes over time, taking into account noisy measurements (images in our case) $\mathbf{z}_{1:t} \triangleq \{\mathbf{z}_1, \dots, \mathbf{z}_t\}$ up to time t . The conceptual solution is a two-step recursion $p(\mathbf{s}_t|\mathbf{z}_{1:t-1}) = \int D(\mathbf{s}_t|\mathbf{s}_{t-1})p(\mathbf{s}_{t-1}|\mathbf{z}_{1:t-1})d\mathbf{s}_{t-1}$ and $p(\mathbf{s}_t|\mathbf{z}_{1:t}) \propto L(\mathbf{z}_t|\mathbf{s}_t)p(\mathbf{s}_t|\mathbf{z}_{1:t-1})$, where two application specific models (the state transition $D(\mathbf{s}_t|\mathbf{s}_{t-1})$ and the likelihood $L(\mathbf{z}_t|\mathbf{s}_t)$) need to be specified. In practice, $p(\mathbf{s}_t|\mathbf{z}_{1:t})$ is obtained using a particle filtering (PF) approximation [6], which represents the posterior by a set of N_s random samples (“particles”), and associated normalized weights $\{\mathbf{s}_t^{(i)}, w_t^{(i)}\}_{i=1}^{N_s}$. With standard PF, these samples and weights are then propagated through time to give an approximation of the filtering distribution at subsequent time steps as $\mathbf{s}_t^{(i)} \sim D(\mathbf{s}_t|\mathbf{s}_{t-1}^{(i)})$ and $w_t^{(i)} \propto L(\mathbf{z}_t|\mathbf{s}_t^{(i)})$, $i = \{1, \dots, N_s\}$, $t = \{1, 2, \dots\}$. At each time step, the optimal state is estimated from $p(\mathbf{s}_t|\mathbf{z}_{1:t})$, for example using the MMSE estimator as $\hat{\mathbf{s}}_t = \sum_{i=1}^{N_s} w_t^{(i)} \mathbf{s}_t^{(i)}$.

More advanced and efficient PF variants (such as sequential importance sampling (SIS) [6]) choose the weights using an importance function describing which areas of the state space contain most information about the posterior. Here, we

propose to use as importance function a time-dependent transform, \mathbf{T}_t , obtained using mutual-information based non-rigid image registration (NRR) [4], to predict the position of tags. The transform maps any point (x, y) in the reference image $I_{t=0}$ to its corresponding point (x', y') in the image I_t at time t , and serves to *approximately* estimate the global motion of the myocardium, after which particle sampling can be constrained to the relevant parts of the state space, much closer to the optimal solution. To utilize \mathbf{T}_t as importance function, we transform the original image sequence $\{I_t\}_{t=0}^{T-1}$ to form new measurements (images) $\mathbf{z}_t(x, y) = I_t(\mathbf{T}_t(x, y))$, which all resemble the initial frame $I_{t=0}$. The proposed PF-based approach is subsequently applied to \mathbf{z}_t and refines the initial coarse tag position estimation (obtained by NRR).

2.2 Refining the Particle Filtering Framework

At each time step t , the configuration of M tags (taken as the intersections of the lines in the tagging pattern) is represented by a set of state vectors $\mathbf{S}_t = \{\mathbf{s}_{m,t}\}_{m=1}^M$. Each tag m is described by $\mathbf{s}_{m,t} = (x_{m,t}, y_{m,t}, \theta_{m,t}, e_{m,t})$, where $(x_{m,t}, y_{m,t})$ and $\theta_{m,t}$ define the spatial position and local orientation, respectively. The binary “existence” variable $e_{m,t}$ indicates whether tag m lies within ($e_{m,t} = 1$) or outside ($e_{m,t} = 0$) the myocardium. In the initial configuration at $t = 0$, the coordinates $(x_{m,0}, y_{m,0})$ of the state vectors correspond to the intersections of perfectly orthogonal tag lines, and the local orientation $\theta_{m,0}$ is known from the imaging protocol (see also Fig. 2). For each tag m with $e_{m,t} = 1$, a neighborhood system $N_{m,t} = \{m' : m \sim m', e_{m',t} = 1\}$ is defined, where “ \sim ” indicates that tag m' belongs to one of the four possible nearest neighbors of tag m . At each time point, the set of “live” tags $E_t = \{m : e_{m,t} = 1\}$ defines the lattice $\mathbf{s}_t = \{\mathbf{s}_{m,t} : m \in E_t\}$, which corresponds to the myocardium.

Likelihood For each tag m , the likelihood $L_s(\mathbf{z}_t | \mathbf{s}_{m,t}, N_{m,t})$ of the state $\mathbf{s}_{m,t}$ is dependent on the tag neighborhood, and given by one of 16 models (Fig. 2). The complete likelihood is factorized as $L(\mathbf{z}_t | \mathbf{s}_t) = \prod_{m \in E_t} L_s(\mathbf{z}_t | \mathbf{s}_{m,t}, N_{m,t})$. To measure the likelihood of the state $\mathbf{s}_{m,t}$, the model is positioned at $(x_{m,t}, y_{m,t})$ with orientation $\theta_{m,t}$, the mean and variance of the image intensities in the “black” and “gray” regions (Fig. 2) are computed (μ_0 and σ_0^2 versus μ_1 and σ_1^2), and the discriminative power of the t -value (related to the Student’s t -test) is used to measure whether the intensity samples from the likelihood model in the two regions are from the same distribution. Thus, the likelihood is defined as $L_s(\mathbf{z}_t | \mathbf{s}_{m,t}, N_{m,t}) = \max\{0, (\mu_1 - \mu_0)(\sigma_1^2 n_1^{-1} + \sigma_0^2 n_0^{-1})^{-1/2}\}$, where n_0 and n_1 denote the number of counted samples in the corresponding regions of the model. Due to the fact that a larger number of samples is used to estimate the means and variances, the proposed likelihood model is both accurate and very robust, even for extremely low-SNR image data.

Transition Prior and Tag Interactions The state evolution is modeled as $D(\mathbf{s}_t | \mathbf{s}_{t-1}) = D_s(\mathbf{s}_t | \mathbf{s}_{t-1}) D_N(\mathbf{s}_t)$, where $D_N(\mathbf{s}_t)$ models tag interactions, and the




Fig. 2. From left to right: example of initial configuration with “live” (black dots) tags and the corresponding schematic LV contours interpolation (gray curves), positioning of the likelihood models depending on the neighborhood ($\theta_{exp} = 45^\circ$), and the possible neighborhood configurations and corresponding models with indicated number of possible unique rotations (d and h are known from the acquisition protocol).

multi-tag transition prior is factorized as $D_s(\mathbf{s}_t|\mathbf{s}_{t-1}) = \prod_{m \in E_t} D_s(\mathbf{s}_{m,t}|\mathbf{s}_{m,t-1})$, which is valid due to the fact that the global (correlated) tag motion was removed by the transform \mathbf{T}_t . To assess the remaining (uncorrelated) deviations of the tags in \mathbf{z}_t from the perfect square grid, the transition prior for $(x_{m,t}, y_{m,t}, \theta_{m,t})$ of a tag is modeled as a constrained random walk. State transitions are further constrained by explicitly modeling the interaction between the tags using a Markov random field (MRF) [10] with respect to the neighborhood systems, using a spring-like model $D_N(\mathbf{s}_t) = \prod_{m \in E_t} \prod_{m' \in N_{m,t}} \Phi(\mathbf{s}_{m,t}, \mathbf{s}_{m',t})$, where $\Phi(\mathbf{s}_{m,t}, \mathbf{s}_{m',t})$ is maximal for tags separated by distance d , and goes to zero if the distance is larger or smaller than d (see also [10]), which penalizes neighboring tag locations that are either too close or too far from each other.

Initialization The initial tag positions for configuration S_0 are obtained by a completely automatic initialization procedure [14], which matches the position of an artificial pattern of orthogonal lines (defined by θ_{exp} , d , and h) with the observed tag pattern in the image $I_{t=0}$. Each tag is tested if it falls into the myocardial region or the background, using the likelihood L_s and the one-tailed t -test: $e_{m,0} = 1$ if the p -value < 0.01 , and $e_{m,0} = 0$ otherwise.

2.3 MCMC-Based Particle Filtering

Tracking of multiple objects using PF in a joint state space (where the state vector \mathbf{s}_t is a collection of individual object states) suffers from exponential complexity in the number of tracked objects, rendering it infeasible for more than a few objects [6, 10]. To deal with a large and variable number of tags (variable dimensionality of the state vector \mathbf{s}_t), we employ TDMCMC-based PF [10]. MCMC methods define a Markov chain over the space of configurations, such that the stationary distribution of the chain is equal to the sought posterior $p(\mathbf{s}_t|\mathbf{z}_{1:t})$. The simulation of the chain (the sets of samples $\{\mathbf{s}_t^{(i)}\}_{i=1}^{N_s}$ representing the posterior), using for example the Metropolis-Hastings (MH) algorithm [8], includes proposing a new state \mathbf{s}_t^* (as $\mathbf{s}^{(i)}$), sampled from a proposal density $q(\mathbf{s}_t^*|\mathbf{s}_t^{(i-1)})$, and accepting or rejecting it according to an acceptance ratio $\alpha = \min(1, p(\mathbf{s}_t^*|\mathbf{z}_{1:t})q(\mathbf{s}_t^{(i-1)}|\mathbf{s}_t^*)/(p(\mathbf{s}_t^{(i-1)}|\mathbf{z}_{1:t})q(\mathbf{s}_t^*|\mathbf{s}_t^{(i-1)})))$.

In TDMCMC [9], the approach is modified to define a Markov chain over a variable-dimensional state space. The algorithm first selects a move type v from a finite set of allowed moves designed to explore a variable-dimensional state space. We propose three moves, which increase the dimensionality of the state (tag birth), decrease it (tag death), or leave it unchanged (tag update). Choosing the move type is done by sampling from a prior distribution $p(v)$. TDMCMC PF generates new samples from the proposal distribution $q(\mathbf{s}_t^*|\mathbf{s}_t^{(i-1)})$, where the move-specific proposal $q_v(m)$ selects (uniformly from the “live” or “dead” tags) a tag that should be updated, added, or removed from \mathbf{s}_t , and the tag-specific proposal $q_v(\mathbf{s}_t^*|\mathbf{s}_t, m)$ (in our case $N_s^{-1} \sum_i D_s(\mathbf{s}_t|\mathbf{s}_{t-1}^{(i)})$) samples its new state. In this case, the corresponding acceptance ratios can be readily computed [9, 10].

Having defined the necessary moves, models and acceptance ratios, we run the TDMCMC PF to obtain an approximation of the filtering distribution $p(\mathbf{s}_t|\mathbf{z}_{1:t})$, where the Markov chain constructed by the TDMCMC represents the believe distribution of the current tag states given the observations. Point estimates are obtained by computing the MMSE position of each tag. The filtering distribution is also used to compute an estimate of tag existence in time at any point within the myocardium by computing the MMSE estimate $\hat{e}_{m,t} = N_s^{-1} \sum_i e_{m,t}^{(i)}$. This information is further used to segment the myocardium by averaging the existence variables for each tag $\hat{e}_{m,t}$ over time and thresholding at value 0.5. The resulting boundary tags are used for spline interpolation of the LV contours. The final result is a dense motion field estimation together with the indication of the tag existence within the myocardial region during motion. The results of tag tracking using the TDMCMC PF are used to refine the B-spline representation of the dense displacement field (initially obtained by the NRR step) using the B-spline refinement procedure proposed in [11].

3 Experimental Results

3.1 Evaluation on Synthetic Image Data

The performance of tag tracking depending on image quality (SNR) was assessed using synthetic image sequences, for which the ground truth was available. To realistically simulate LV motion through the cardiac cycle, we used a numerical phantom [3], which modeled simultaneous radial expansion/contraction and rotation of “tagged” myocardium (Fig. 3). The image sequences (30 frames of size 256×176 pixels) were created for SNRs of $\{18.06, 12.04, 8.51, 6.02\}$ dB, by corrupting the noise-free images with Rician noise (Fig. 3), which is typical in MRI. Additionally, the typically observed fading of the tag pattern during a heart cycle was modeled, resulting in sequences with time-varying SNR (changing from 18.06 dB in the initial frame to 6.02 dB in the last frame) closely resembling the SPAMM-tagged images of a gel phantom used in [7].

The accuracy of tag tracking was evaluated using the root mean square error (RMSE) in tag positions, averaged for each method over five independent runs on different realizations of the synthetic data. The NRR step was done using




Fig. 3. Example of synthetic data with time-varying SNR (modeling tag fading) in frame 1 (a) and frame 10 (b), and zoomed image regions demonstrating the modeled tag and myocardium appearance for different SNR levels, corresponding to 18.06 dB (c), 12.04 dB (d), 8.51 dB (e), 6.02 dB (f).

SNR	6.02 dB	8.51 dB	12.04 dB	18.06 dB	Time-varying SNR
HARP	1.99±3.15	1.15±8.56	0.51±0.66	0.39±0.88	1.17±1.79
Standard PF	0.94±0.78	0.76±0.56	0.43±0.45	0.29±0.32	0.53±0.37
TDMCMC PF	0.64±0.37	0.39±0.22	0.24±0.16	0.13±0.10	0.34±0.25
Manual	1.24±0.71	0.88±0.55	0.71±0.50	0.61±0.41	0.64±0.38

Table 1. The results of tracking accuracy assessment (RMSE as a function of SNR), where the numbers represent RMSE \pm standard deviation, given in pixel units.

`elastix`, an open source image registration toolbox (<http://elastix.isi.uu.nl/>). The parameters for the TDMCMC simulation were fixed to $p(\text{death}) = p(\text{birth}) = 0.1$, $p(\text{update}) = 0.8$, $N_s = 100000$. All other parameters were optimized for best performance. The results of tag tracking using TDMCMC PF are shown in Table 1, in comparison with manual analysis, tracking using HARP [12] and a “standard” PF implementation (recently proposed by Smal et al. [13], which does not use NRR-based importance sampling). Our method clearly performs superiorly. The computational costs are given in Table 2.

3.2 Evaluation on Real Image Data

Experiments on real tMRI data were conducted on a clinical 3T MRI scanner (GE Medical Systems). Image sequences were collected from studies on healthy and diseased rats and pigs, and diseased human patients. Multiple short-axis view images of size 256×256 pixels were acquired with the following imaging parameters {human-, rat-, pig-related}: repetition time {6.5, 13, 4} msec, echo time {3.1, 4, 1.25} msec, flip angle {12, 7, 11} degrees, slice thickness {8, 1.6, 6} mm, spacing between slices {28, 1.6, 12} mm, pixel size $\{1.48 \times 1.48, 0.19 \times 0.19, 1.25 \times 1.25\}$ mm², frames per heart cycle {20, 24, 20}, number of slice positions {3, 7, 4}, tag spacing {7, 1.5, 6} mm, and $\theta_{exp} = 45^\circ$. Five tMRI sequences of each type (human, rat, pig) were analyzed using the proposed method. The algorithm parameters were fixed to the same values as in the experiments on synthetic data. As no ground truth was available this time, we measured the accuracy in comparison with manual tracking produced by experts.

The results of the comparison are given in Table 3 and typical tracking results in comparison with HARP are shown in Fig. 4. Visual inspection reveals that the proposed method correctly estimated the tag intersections within the myocardium through the whole image sequence. Contrary to our method, HARP

	Time
HARP	1.3
Standard PF	3.4
NRR+TDMCMC PF	11.7+10.2

Table 2. Average computational times in [s/frame] for considered image sequences.

	Human	Rat	Pig
HARP	1.65±1.73	1.36±0.93	1.43±1.35
Standard PF	1.48±1.43	1.35±0.87	1.27±1.17
TDMCMC PF	1.23±1.02	1.33±0.84	1.12±0.84

Table 3. The results of tracking accuracy assessment (RMSE) using real image data from human, pig and rat experiments.




Fig. 4. Tag trajectories obtained using HARP (a) and TDMCMC PF (b), with examples of Lagrangian circumferential strain maps obtained using TDMCMC PF for frames 12 (c) and 16 (d) of the human-heart image data.

produced many erroneous tracks, and reliable tag tracking was achieved only in the first 3–6 frames of the image sequence (where the tags had not faded yet) or in regions within the myocardium with sufficiently high SNR. The results in Table 3 do not include these erroneous tracks, because of their absence in the manually annotated data (it was difficult and tedious even for the experts to track the myocardial boundary points). Thus, the results in the table are biased towards tags that could be tracked relatively well, and the differences between the methods would likely be larger if all tags would be included.

4 Conclusions

Motion analysis in cardiac tMRI is a challenging problem in practice due to the poor quality of image data and complex motion scenarios. In this paper we have proposed and evaluated a novel Bayesian approach to tag tracking in tMRI, which combines prior knowledge about the heart dynamics (obtained using non-rigid image registration) with modeling of the myocardium appearance in noisy tMRI data. Straightforward generalization of the Bayesian formulation to the problem of multi-tag tracking is computationally prohibitive due to the increase in dimensionality of the state space. Therefore we have proposed a novel particle filtering (PF) method based on a trans-dimensional Markov chain Monte Carlo (TDMCMC) approach that efficiently deals with high-dimensional spaces and can track varying numbers of tags over time. The new method, which was evaluated using both synthetic and real tMRI data, demonstrated higher tracking consistency, accuracy and robustness in comparison with the commercially available HARP and a standard PF approach. Apart from yielding more accurate tracking results, the proposed TDMCMC PF is also capable of detecting the presence/absence of tags within the myocardium in a probabilistic fashion, which can be used to signify whether the tracking results in the different parts

of the myocardium are reliable or not. The analysis results (dense displacement and strain maps, obtained after refining the initial B-spline representation of the dense displacement field (obtained by the NRR step) with the tag tracking using the TDMCMC PF) can be used to develop new classification techniques for automatic diagnosis of healthy and diseased patients. This will be a subject of future work. As yet, the method is already being employed in ongoing longitudinal experiments in our institute, with the primary goal to quantify left ventricular remodeling after myocardial infarction in small animals.

References

1. Amini, A.A., Chen, Y., Curwen, R.W., Mani, V., Sun, J.: Coupled B-snake grids and constrained thin-plate splines for analysis of 2-D tissue deformations from tagged MRI. *IEEE Trans. Med. Imaging* 17(3), 344–356 (1998)
2. Axel, L., Dougherty, L.: MR imaging of motion with spatial modulation of magnetization. *Radiology* 171, 841–845 (1989)
3. Carranza-Herrezuelo, N., Bajo, A., Sroubek, F., Santamarta, C., Cristobal, G., Santos, A., Ledesma-Carbayo, M.J.: Motion estimation of tagged cardiac magnetic resonance images using variational techniques. *Comput. Med. Imaging Graph.* 34, 514–522 (2010)
4. Chandrashekhara, R., Mohiaddin, R.H., Rueckert, D.: Analysis of 3-D myocardial motion in tagged MR images using nonrigid image registration. *IEEE Trans. Med. Imaging* 23, 1245–1250 (2004)
5. Deng, X., Denney, Jr., T.S.: Three-dimensional myocardial strain reconstruction from tagged MRI using a cylindrical B-spline model. *IEEE Trans. Med. Imaging* 23(7), 861–867 (2004)
6. Doucet, A., de Freitas, N., Gordon, N.: *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, Berlin (2001)
7. Dougherty, L., Asmuth, J.C., Blom, A.S., Axel, L., Kumar, R.: Validation of an optical flow method for tag displacement estimation. *IEEE Trans. Med. Imaging* 18, 359–363 (1999)
8. Gilks, W.R., Richardson, S., Spiegelhalter, D.J.: *Markov Chain Monte Carlo in Practice*. Chapman and Hall, Boca Raton (1996)
9. Green, P.: Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82, 711–732 (1995)
10. Khan, Z., Balch, T., Dellaert, F.: MCMC-based particle filtering for tracking a variable number of interacting targets. *IEEE Trans. Pattern Anal. Mach. Intell.* 27, 1805–1819 (2005)
11. Lee, S., Wolberg, G., Shin, S.Y.: Scattered data interpolation with multilevel B-splines. *IEEE Trans. Vis. Comp. Graph.* 3, 228–244 (1997)
12. Osman, N.F., Kerwin, W.S., McVeigh, E.R., Prince, J.L.: Cardiac motion tracking using CINE harmonic phase (HARP) magnetic resonance imaging. *Magn. Reson. Med.* 42, 1048–1060 (1999)
13. Smal, I., Niessen, W.J., Meijering, E.: Particle filtering methods for motion analysis in tagged MRI. In: 7th IEEE International Symposium on Biomedical Imaging. pp. 488–491. IEEE Press, Piscataway (2010)
14. Urayama, S., Matsuda, T., Sugimoto, N., Mizuta, N., Yamada, N., Uyama, C.: Detailed motion analysis of the left ventricular myocardium using an MR tagging method with a dense grid. *Magn. Reson. Med.* 44, 73–82 (2000)