An MRF-Based Discrete Optimization Framework for Combined DCE-MRI Motion Correction and Pharmacokinetic Parameter Estimation

Monica Enescu[1], Mattias P. Heinrich[3], Esme Hill[2], Ricky Sharma[2], Michael A. Chappell[1], and Julia A. Schnabel[1]

[1] Institute of Biomedical Engineering, University of Oxford, UK
[2] Department of Oncology, University of Oxford, UK
[3] Institute of Medical Informatics, University of Lüebeck, Germany

Abstract. Dynamic contrast-enhanced MRI (DCE-MRI) images are increasingly used for assessing cancer treatment outcome. These time sequences are typically affected by motion, which causes significant errors in tracer kinetic model analysis. Current intra-sequence registration methods for contrast enhanced data either assume restricted transformations (e.g. translation) or employ continuous optimization, which is prone to local optima. In this work, we propose a new approach to DCE-MRI intra-sequence registration and pharmacokinetic modelling, which is formulated in an MRF optimization framework. The complete $4D$ graph corresponding to a DCE-MRI sequence is reduced to a concatenation of minimum spanning trees, which can be optimized more efficiently. To address the changes due to contrast, a data cost function which incorporates pharmacokinetic modelling information is formulated. The advantages of this method are demonstrated on 8 DCE-MRI image sequences of patients with advanced rectal tumours, presenting mild to severe motion.

1 Introduction

DCE-MRI has become an important tool for assessing early phase clinical trials of cancer therapy, as it can measure in vivo tumour vasculature changes that occur due to treatment. The underlying tissue physiology is typically derived from the DCE-MRI image signal by fitting a pharmacokinetic (PK) model to the contrast enhancement-time curve on a voxel-by-voxel basis. Typically, tissue perfusion, permeability and the volume occupied by tumour cells are obtained in terms of PK model parameters. As a DCE-MRI acquisition takes several minutes, with volumes being acquired every $5 - 10$ seconds, the resulting time sequence is inherently affected by patient and physiological motion. This motion may introduce significant errors to the per-voxel PK model fitting, as anatomical features of interest might move to different voxel locations in subsequent volumes. To correct for this motion, image registration is needed. DCE-MRI registration is a particularly challenging problem, as observed changes throughout the time series can be either due to motion or due to contrast enhancement.

M.J. Cardoso et al. (Eds.): BAMBI 2014, LNCS 8677, pp. 73–84, 2014.

Moreover, contrast arrival can give rise to image features that were not present in the baseline image. In the literature, approaches for time series motion correction broadly fall into two categories: 1. strategies that try to alleviate the effects of contrast enhancement i.e. by using a multi-modal similarity metric [1], or by restricting the applied transformation [2]; 2. strategies that implicitly derive [3] or explicitly assume [4,5,6] a model of contrast enhancement which is used in the registration algorithm. Among the most prominent approaches, we note the work of Buonaccorsi et al. [5], who are among the first to explore an explicit kinetic model-based registration. In that work, PK parameter estimation and registration to the model predicted sequence are performed iteratively. Their work is mainly limited by allowing only 3D translation transforms. In a more recent approach, Bhushan et al. [6] address this issue by proposing a simultaneous non-rigid motion correction and PK parameter estimation method. However, as their approach uses a Gauss-Newton optimization, this method is sensitive to initialization and is likely to be trapped in local optima. A data-driven approach, presented by Melbourne et al. [3], proposes modelling the time series data using principal component analysis. The underlying assumption is that the first few principal components will contain information about contrast enhancement trends and the remaining principal components contain noise related to motion. This assumption is valid for small peristaltic motion, but does not hold in the case of larger and periodic motion (e.g. breathing).

In this work, we propose a new framework for combined DCE-MRI intra-sequence registration using discrete optimization and PK parameter estimation (DireP). The problem is formulated using a Markov random field (MRF) which involves the optimization on a $4D$ graph for each DCE-MRI sequence. This method addresses the sensitivity to initialization of continuous approaches and is less prone to local optima by offering increased flexibility over the space of possible displacements $\mathcal{L}$. To reduce the computational costs which are typically associated with discrete optimization of a full 4D graph, the nodes are connected as follows: The minimum spanning tree (MST) which best replicates the underlying anatomy of the pre-contrast image is calculated [7]. This structure is assumed to be identical in all the subsequent volumes, as they have the same anatomy as the baseline. These structures are connected through time at every node, to preserve the temporal continuity of each voxel (Fig. 1).

The paper is structured as follows. Section 2 describes the methodological contributions of this paper: First, a DCE-MRI tailored similarity metric which incorporates PK modelling information is formulated. Next, in order to reduce computational costs, we construct a reduced $4D$ graph corresponding to the DCE-MRI sequence, and optimize it using a message passing approach. To our knowledge, this is the first work to perform combined DCE-MRI motion correction and PK analysis, where the registration is performed in a discrete optimization framework. Section 3 describes the results of the proposed algorithm (DireP), on both synthetic and real DCE-MRI images of rectal cancer. DireP is compared to a recent DCE-MRI non-rigid registration algorithm (referred to

Method1) [6] which employs continuous optimization. We conclude this paper with Section 4, and present future work plans in Section 5.

2 Methods

Discrete optimization is typically formulated as an MRF labelling problem [8]. For deformable image registration purposes, a graph is defined in which the nodes $p \in \Omega$ represent voxels or groups of voxels, and the edges connect voxels with similar anatomical features and spatial proximity. For every node p, there is a set of labels $l_p \in \mathcal{L}$ which represent possible discrete displacements of the source image volume with respect to the target image volume. Finding the optimum displacement at each voxel equates to finding the labelling that minimizes the MRF energy function:

$$E(l) = \sum_{p \in \Omega} C_D(l_p) + \gamma \sum_{(p,q) \in \mathcal{N}} C_R(l_p, l_q) \tag{1}$$

The unary term represents the data cost C_D, which measures the similarity of a voxel in the target image to the corresponding voxel in the source image displaced with l_p. The pairwise term represents the regularization cost $C_R(l_p, l_q)$ and is used to smooth the displacements of directly connected voxels $(p, q) \in \mathcal{N}$. Here, $\mathcal{N}$ represents the neighborhood of a voxel, as given by the MST. γ weights the amount of regularization.

Methods to solve the labelling problem can be roughly divided into two categories: graph-cuts and message passing approaches. Popular graph-cuts algorithms include α-expansion [9]. Depending on the complexity of the graph to be optimized, message passing can range from dynamic programming, over loopy belief propagation (LBP) [10], to tree-reweighted message passing (TRW-S) [11]. For an overview of discrete methods for deformable registration, we refer the reader to the work of Sotiras et al. [12].

In this work, each image volume of the DCE-MRI sequence is registered to the pre-contrast volume by optimizing a reduced $4D$ graph corresponding to the underlying anatomy. To address the intensity differences caused by contrast inflow, a DCE-MRI tailored data cost function incorporating PK information is proposed. The optimization is performed using belief propagation. An overview of the entire algorithm for finding the optimal displacement and updating the PK model can be found in Algorithm 1.

2.1 Data Cost Calculation Using Pharmacokinetic Model Prediction

As mentioned above, the DCE-MRI intensities of different volumes in the image sequence are not comparable using standard similarity metrics such as the sum of squared differences (SSD). To address this issue, we propose to compare the similarity between the intensities at the volume to be registered I_{t_i}, and the PK model predicted intensity at that volume $PK(I_{t_0}, K^{trans}, v_e, t_i)$:

$$C_D(l_p) = SSD(I_{t_i}(p + l_p), PK(I_{t_0}(p), K^{trans}(p), v_e(p), t_i)) \tag{2}$$

Here, t_i is the current time point of the DCE-MRI sequence, t_0 is the first time point, and K^{trans} and v_e are parameters of the PK model. To predict the appearance of the baseline I_{t_0} at t_i, an initial estimate of the PK parameters K^{trans}, v_e is required. This initial estimate is obtained by least squares (lsq) fitting of the model to the data.

PK parameter estimation was performed using the standard Tofts model [13] with the Orton [14] bi-exponential arterial input function with population averaged parameters. The Tofts model offers an estimate of tissue perfusion through K^{trans}, the transfer constant between blood plasma and the extravascular extracellular space (EES), which is defined as the product between k_{ep}, the rate constant between the EES and plasma, and v_e, the fractional volume of the EES.

2.2 Optimization on the Reduced 4D Graph

For the DCE-MRI motion correction problem, the graph to be optimized can be reduced to temporal concatenation of identical spatial trees (Fig. 1). The pre-contrast image anatomy is taken as a reference, against which all the subsequent volumes are aligned. We assume that they obey the same anatomy and differ only by contrast enhancement or motion. Under this assumption, we estimate the unique MST of the pre-contrast volume using Prim's algorithm [7] akin to the work of Heinrich et al. [15]. This structure is replicated for each subsequent volume, and the temporally corresponding nodes are additionally connected across time, as they represent the same anatomical landmark over time. For each individual tree, the optimization is possible using belief propaga-

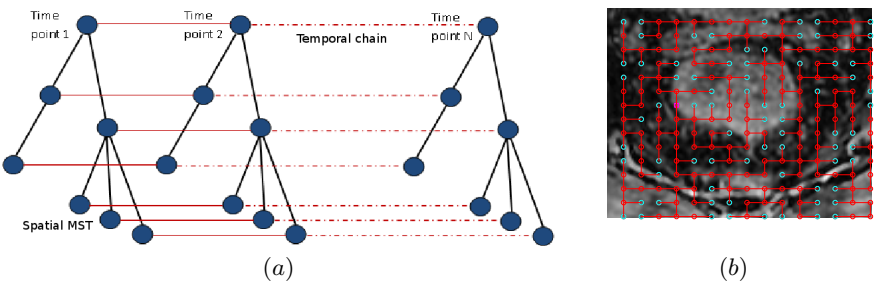

Fig. 1. (a) An illustration of the graphical structure used in this algorithm. While the spatial connectivity is captured by minimum spanning trees (MST), the temporal continuity is captured by temporal chains. (b) An example of MST for a 2D slice through the mesorectal area.

tion on the MST [16]. At each node p, a message vector m_p containing the cost of the best displacement l_p can be found given the displacement of its parent q, l_q, and the messages from its children c, m_c:

$$m_p^{tree}(l_q) = \min_{l_p}(C_D(l_p) + \gamma_{sp} C_R(l_p, l_q) + \sum_c m_c(l_p)) \tag{3}$$

For leaf nodes, Eq. 3 can be evaluated directly, as there are no incoming messages from children. The messages are then forward-passed from the leaves to the root node, and then backward-passed, from the root towards the leaves. For each individual temporal chain, the optimization can be performed using belief propagation on the chain [17]. The message of the current node p, m_p, can be found given the displacement of the previous node q, l_q, and the message coming from the subsequent node r, $m_r(l_p)$:

$$m_p^{chain}(l_q) = \min_{l_p}(C_D(l_p) + \gamma_{temp} C_R(l_p, l_q) + m_r(l_p)) \tag{4}$$

To optimize the reduced 4D structure, for each node we perform independent optimization on each temporal chain and on each spatial tree, and average the resulting marginals. This procedure is repeated for a number of iterations akin LBP [10], where the messages towards a node in the previous iteration are added to the marginal of that node in the current iteration. The best displacement can then be found by calculating argmin based on the marginal distribution. Although this approach is not guaranteed to converge to a global optimum, it is physically motivated, and provides a good trade-off between optimality and efficiency. The averaging of marginals from multiple graph models has previously been used in stereo processing using an approach called semi-global matching [18], yielding excellent results. The entire DireP algorithm is presented in Algorithm 1. In each optimization step (Algorithm 1, 2.2, 2.3), the full distribution of pseudo-marginals for the space of displacements is estimated.

Naively calculating the pair-wise regularization cost in Eqs. 3, 4 would require $|\mathcal{L}|^2$ computations for every voxel. To reduce this cost, we employed the min-convolution technique [10], which reduces the complexity to $|\mathcal{L}|$. To further reduce computational costs, a multi-resolution approach was employed. For each resolution level, the image is divided into non-overlapping cubic groups of voxels that are represented by a single node in the graph. The regularization term is calculated only for each group of voxels. At finer levels, the previous solution is upsampled and used as a prior for the optimization algorithm. The final (dense) solution is also obtained by upsampling using a first order spline interpolation.

Additionally, if we treat the deformation field as a velocity field, it can be transformed into a diffeomorphic mapping [15] by using the scaling and squaring method [19]. Diffeomorphism is a desired property for the deformation field, as it prevents transformations that are not physically feasible, i.e. folding. This is particularly important for soft tissue images, as it agrees with the tissue incompressibility assumption.

3 Results

3.1 Algorithm Evaluation on Synthetic Data

The discrete registration and pharmacokinetic estimation (DireP) algorithm proposed in this paper was first tested on synthetic data, where the ground truth intra-sequence motion and PK parameters are known. To simulate a realistic

Algorithm 1. DireP: Discrete motion correction and pharmacokinetic estimation

```
1. PK parameter estimation on uncorrected time series (lsq fitting)
while n_iterPK do
  2. Groupwise registration of all volumes to the pre-contrast volume:
  2.1 Initialize marginals and messages:
  foreach node do marginal_c[ node ]=C_D(l_p); marginal_t[ node ]=C_D(l_p);
  message[ node ]=0;
  while n_iterMRF do
    2.1 Re-initialize marginals with values from previous iteration
    2.2 foreach timepoint pass messages on the corresponding spatial MST
    Forward pass
    for node=leaves to root-1 do
      cost = marginal_t[node];
      message[ node ]=min-sum(cost); (see Eq. 3)
      marginal_t[ parent ] = marginal_t[ parent ]+message[ node ];
    end for
    Backward pass
    for node=root-1 to leaves do
      cost=marginal_t[ parent ] - message[ node ]+ message[ parent ];
      message[ node ]=min-sum(cost);
    end for
    foreach node marginal_t[ node ]=marginal_t[ node ]+message[ node ];
    2.3 foreach node of the spatial MST, pass messages on the temporal chain
    Forward pass
    for t=tdim-1 to 1 do
      cost = marginal_c[ node_t ];
      message[ node_t ]=min-sum(cost); (see Eq. 4)
      marginal_c[ node_{t-1} ] = marginal_c[ node_{t-1} ]+message[ node_t ];
    end for
    Backward pass
    for t=1 to tdim-1 do
      cost=marginal_c[node_{t-1} ] - message[ node_t ]+ message[ node_{t-1} ];
      message[ node ]=min-sum(cost);
    end for
    foreach node marginal_c[ node ]=marginal_c[ node ]+message[ node ];

    2.4 Average the two marginals and use the result in the next iteration
  end while
  3. Re-estimate PK parameters on corrected time series (lsq fitting),
  use in the next iteration
end while
```

dataset, synthetic images were generated as follows: 8 real DCE-MRI sequences were selected, and PK model fitting was performed on each of them. Using the resulting parameter maps and the pre-contrast image volumes, 8 model predicted sequences of size $120 \times 120 \times 52 \times 29$ each were generated. These images constituted the ground truth motion free dataset. To simulate motion, a random displacement field of size $6 \times 6 \times 4$ was generated for each image volume independently. This displacement was upsampled to the image volume size, and smoothed with a Gaussian kernel. We also applied Gaussian smoothing on the temporal dimension, as the motion of each voxel is expected to have some degree of temporal smoothness due to periodic motion patterns i.e. breathing. The parameters for the synthetic motion were chosen to obtain a diffeomorphic deformation field in the interval $\pm$7mm along each direction. The experiments were run with $n_{iterPK} = 2$, $n_{iterMRF} = 5$ (see Algorithm 1), $\gamma_{sp} = 0.01$ and $\gamma_{temp} = 0.1$, which were empirically chosen. For both the synthetic and the real data, the label space was chosen as $L = \{0, \pm q, \pm 2q, .., \pm\frac{u}{2}q\}^3$. 3 resolution levels were used in the registration. Depending on the resolution level, u was $8, 6, 4$, with a quantization q of $2, 1, 0.5$mm. We used groups of voxels of sizes 8^3, 6^3 and 4^3, with corresponding label spaces of size 9^3, 7^3 and 5^3.

Table 1. Registration results on synthetic data. The average target registration error (TRE) and residual fitting errors in PK parameters are reported, together with the corresponding standard deviations.

Measure	Before	Method1	DireP $3D$	DireP $4D$
TRE (mm)	1.40 ± 0.72	1.02 ± 0.65	$\mathbf{0.59 \pm 0.38}$	0.62 ± 0.39
Error in k_{ep}	0.39 ± 0.03	0.15 ± 0.02	$\mathbf{0.13 \pm 0.01}$	$\mathbf{0.13 \pm 0.01}$
Error in v_e	0.20 ± 0.01	0.18 ± 0.02	$\mathbf{0.03 \pm 0.002}$	$\mathbf{0.03 \pm 0.002}$

Quantitative results on the synthetic dataset are presented in Table 1. The target registration error (TRE) is defined as the average difference between the ground truth deformation field and the deformation field estimated by the algorithms. DireP was compared to a recent non-rigid registration and PK estimation approach using continuous optimization (Method1) [6]. DireP $3D$ is the algorithm version without temporal regularization and DireP $4D$ is the variant including temporal regularization. The results show that while both DireP $3D$ and DireP $4D$ outperform Method1 and recover a good part of the synthetic motion, DireP $3D$ has a slightly better performance in terms of TRE. This is due to trading a solution that is globally optimal for each volume (DireP $3D$) for a sub-optimal solution on the entire $4D$ graph (DireP $4D$). After applying each method, the PK parameters on the corrected datasets were generated by least-squares fitting. The errors in k_{ep} and v_e are defined as the average absolute difference between the ground truth PK parameters and the parameters obtained on the corrected datasets. These values were calculated on a circular region delineated around the tumour in the pre-contrast image volume (which

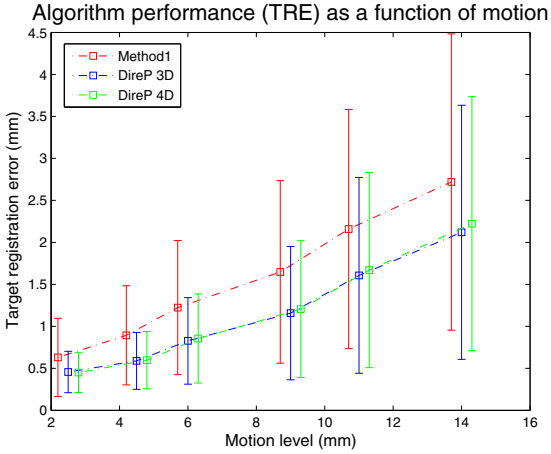

Fig. 2. The TREs (mm) for the three different algorithms are shown in function of the motion level. The performance of Method1 is represented with red error bars, DireP $3D$ is represented in blue, and DireP $3D$ is represented in green. The coloured error bars are slightly displaced for better visualization, but they correspond to the same level of motion.

was taken from a real dataset as explained above). The PK parameter errors decrease considerably after DireP $3D$ and DireP $4D$ registration.

The relationship between the algorithms' performance and the level of motion in the synthetic data was also investigated, and the results are shown in Fig. 2. It can be seen that DireP $3D$ and DireP $4D$ outperform Method1 regardless of the motion level, and the difference is more pronounced for bigger motion.

3.2 Pharmacokinetic Modelling and Motion Correction on DCE-MRI Images of Rectal Cancer

The presented algorithm was also applied to a dataset of DCE-MRI images from 8 patients with advanced rectal tumours. T1-weighted dynamic images of the pelvis were acquired using the spoiled gradient echo LAVA protocol. Contrast agent (ProHance) was injected at a rate of 3ml/sec, 0.1 mmol/kg body weight. An image volume was acquired every 9.5 seconds for approximately 5 minutes, yielding a sequence of $512 \times 512 \times 52 \times 29$ with a resolution of $0.7813 \times 0.7813 \times 2$mm. The algorithms were applied to a $120 \times 120 \times 52 \times 29$ ROI which contained the tumour volume. The tumours were manually delineated by our clinical collaborators. The experiments were run using $n_{iterPK} = 3$, $n_{iterMRF} = 5$, $\gamma_{sp} = 0.01$ and $\gamma_{temp} = 0.1$, which were empirically chosen. For the registration, 4 resolution levels were used. Depending on the resolution level, u was $8, 6, 4, 4$, with a quantization q of $2, 1, 0.5, 0.5$mm. We used groups of voxels of sizes 8^3, 6^3, 4^3 and 2^3, with corresponding label spaces of size 9^3, 7^3, 5^3 and 5^3. The results for the entire clinical dataset are shown in Fig. 3. The

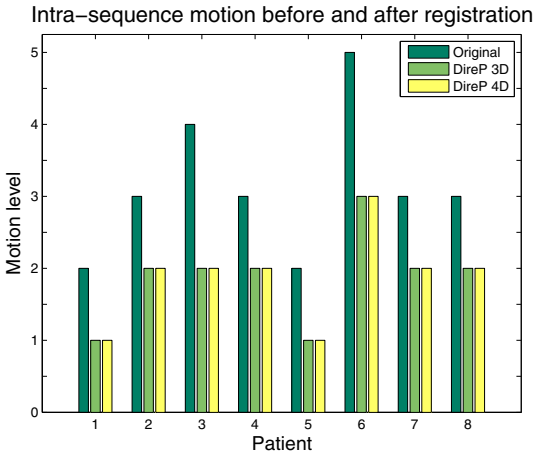

Fig. 3. The DCE-MRI images were blindly graded before and after registration by a clinical expert. A score of '1' represents no motion, '2' is mild or minimal motion, '3' is moderate motion, '4' is significant motion and '5' is severe motion. Both DireP $3D$ and DireP $4D$ reduce the amount of motion in all the images of the rectal cancer dataset.

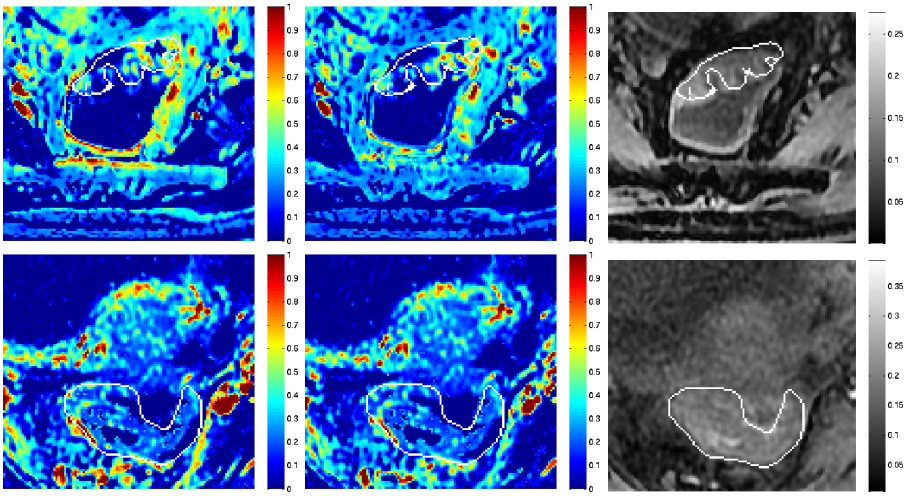

Fig. 4. K_{trans} maps for one slice of the image volume. Images before (a) and after DireP $3D$ (b) motion correction are shown, together with the corresponding anatomical image (c). The top row corresponds to a patient exhibiting severe motion, and the bottom row corresponds to a patient exhibiting moderate motion.

motion level in each time series was blindly graded before and after registration by a clinical expert. It can be seen that the observed motion level decreases in all the patients after applying our discrete registration framework in $3D$ and $4D$,

respectively. We also note that the differences between DireP $3D$ and DireP $4D$ were not detectable on a visual evaluation. Figure 4 presents the effect of motion correction on K_{trans}. Images before and after DireP $3D$ are shown, for a patient exhibiting severe motion, and for a patient exhibiting moderate motion. It can be seen that in both cases the K_{trans} maps become sharper after registration, with a clearer separation between individual voxels, and motion artefacts at the tumour boundary are reduced. This effect is particularly visible for the severe motion case, where the non-registered sequence yields a heavily blurred K_{trans} map.

On a 2.93GHz CPU, using C++ code, runtimes are: Method1 14.29min, DireP $3D$ 3.15min, DireP $4D$ 7.02min. For both the $3D$ and the $4D$ algorithm, the Jacobian determinant of the deformation field is positive. The Jacobian determinants were calculated for the real dataset, as well as for the synthetic images.

4 Discussion and Conclusion

We have proposed a new algorithm for DCE-MRI time series motion correction and PK estimation (DireP), which is formulated on an MRF and uses discrete optimization. The PK estimation is performed iteratively with the deformable registration. Two variants of the algorithm, one without temporal regularization (DireP $3D$) and one with temporal regularization (DireP $4D$) were tested on a synthetic dataset and a dataset comprising 8 DCE-MRI sequences of rectal cancer patients. Both DireP $3D$ and DireP $4D$ reduce the amount of motion in all the images of the rectal cancer dataset, especially in challenging sequences exhibiting severe motion. When tested on synthetic data, both the variants show an improvement over a state-of-the art algorithm using continuous optimization [6]. Although DireP $3D$ slightly outperforms DireP $4D$ in terms of TRE ($0.59mm$ vs. $0.62mm$), the key advantage of DireP $4D$ lies in its ability to impose a degree of temporal smoothness, which in itself is desirable to avoid unrealistic fitting of the PK model. At the same time, motions such as peristalsis with a high frequency cannot be captured by a transformation that enforces temporal smoothness. We expect the positive effect of the 4D regularization to be much more visible for applications where lower frequency motion, i.e. breathing, is predominant. Examples include liver DCE-MRI, or rectal data where peristalsis is controlled by drug administration.

The pre-contrast image was chosen as a reference for registration as it more closely represents the true appearance of the anatomy, where features are not distorted by contrast arrival. This also avoids registration error propagation. For the PK modelling part, the standard Tofts model is assumed, as it is widely used in clinical practice. Our algorithm is expected to give comparable if not better results using a more complex model.

5 Future Work

In this work, the $4D$ optimization problem was solved by alternating between spatial MST and temporal chain minimization, and averaging the resulting

marginals. This procedure was repeated a couple of times akin LBP [10]. Although this approach has the advantage of reduced complexity, as it involves optimizing loop-free graphs, convergence might be slow and is not guaranteed. In future work, algorithms for joint minimization of the spatial and temporal problem will be investigated. A possible approach is employing a $4D$ generalization of the TRW-S algorithm [11].

Acknowledgements. Monica Enescu is funded from an EPSRC Doctoral Training Award, an Oxford Department of Engineering Science Scholarship and the Alberto Del Vicario Scholarship. Julia Schnabel and Esme Hill gratefully acknowledge funding from CRUK/EPSRC Centre for Cancer Imaging at Oxford. Esme Hill was also funded by the Oxfordshire Health Services (OHSRC). Ricky Sharma acknowledges funding from NIHR BRC Oxford.

References

1. Zoellner, F.G., Sancee, R., Rogelj, P., Ledesma-Carbayo, M.J., Rorvik, J., Santos, A., Lundervold, A.: Assessment of 3D DCE-MRI of the kidneys using non-rigid image registration and segmentation of voxel time courses. Comput. Med. Imag. Graphics 33(3) (2009)
2. Tanner, C., Schnabel, J.A., Chung, D., Clarkson, M.J., Rueckert, D., Hill, D.L.G., Hawkes, D.J.: Volume and shape preservation of enhancing lesions when applying non-rigid registration to a time series of contrast enhancing MR breast images. In: Delp, S.L., DiGoia, A.M., Jaramaz, B. (eds.) MICCAI 2000. LNCS, vol. 1935, pp. 327–337. Springer, Heidelberg (2000)
3. Melbourne, A., Atkinson, D., White, M.J., Collins, D., Leach, M., Hawkes, D.: Registration of dynamic contrast-enhanced MRI using a progressive principal component registration (PPCR). Phys. Med. Biology 52(17) (2007)
4. Hayton, P., Brady, M., Tarassenko, L., Moore, N.: Analysis of dynamic MR breast images using a model of contrast enhancement. Medical image analysis 1(3) (1997)
5. Buonaccorsi, G.A., O'Connor, J.P.B., Caunce, A., Roberts, C., Cheung, S., Watson, Y., Davies, K., Hope, L., Jackson, A., Jayson, G.C., Parker, G.J.M.: Tracer kinetic model-driven registration for dynamic contrast-enhanced MRI time-series data. Magnetic Resonance in Medicine 58(5) (2007)
6. Bhushan, M., Schnabel, J.A., Risser, L., Heinrich, M.P., Brady, J.M., Jenkinson, M.: Motion correction and parameter estimation in dceMRI sequences: Application to colorectal cancer. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part I. LNCS, vol. 6891, pp. 476–483. Springer, Heidelberg (2011)
7. Prim, R.: Shortest connection networks and some generalizations. Bell System Technical Journal (36) (1957)
8. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. Medical Image Analysis 12(6) (2008)
9. Boykov, V., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Tran. PAMI 23, 1222–1239 (2001)
10. Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient belief propagation for early vision. Intl. J. Computer Vision 70(1) (2006)

11. Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. IEEE Tran. PAMI 28(10) (2006)
12. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: A survey. IEEE Tran. Medical Imaging 32(7) (2013)
13. Tofts, P.S.: Modeling tracer kinetics in dynamic Gd-DTPA MR imaging. J. Magnetic Resonance Imaging 7(1) (1997)
14. Orton, M.R., d'Arcy, J.A., Walker-Samuel, S., Hawkes, D.J., Atkinson, D., Collins, D.J., Leach, M.O.: Computationally efficient vascular input function models for quantitative kinetic modelling using DCE-MRI. Phys. Med. Biology 53(5) (2008)
15. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: MRF-based deformable registration and ventilation estimation of lung CT. IEEE Tran. Medical Imaging 32(7) (2013)
16. Heinrich, M.P., Simpson, I., Jenkinson, M., Brady, M., Schnabel, J.A.: Uncertainty estimates for improved accuracy of registration-based segmentation propagation using discrete optimization. In: MICCAI SATA Workshop (2013)
17. Felzenszwalb, P.F., Zabih, R.: Dynamic programming and graph algorithms in computer vision. IEEE Tran. PAMI 33(4), 721–740 (2011)
18. Hirschmueller, H.: Stereo processing by semiglobal matching and mutual information. IEEE Tran. PAMI 30(2), 328–341 (2008)
19. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006)