Vision, Modeling, and Visualization (2012)
M. Goesele, T. Grosch, B. Preim, H. Theisel, and K. Toennies (Eds.)

Registration of temporal ultrasonic image sequences using Markov random fields

Sebastian Schäfer and Klaus Toennies

Department of Simulation and Graphics, University of Magdeburg, Germany

Abstract

Ultrasound perfusion imaging is a rapid and inexpensive technique which enables observation of a dynamic process with high temporal resolution. The image acquisition is disturbed by various motion influences due to the acquisition procedure and patient motion. To extract valid information about perfusion for quantification and diagnostic purposes this influence must be compensated. In this work an approach to account for non-linear motion using a markov random field (MRF) based optimization scheme for registration is presented. Optimal transformation parameters are found all at once in a single optimization framework. Spatial and temporal constraints ensure continuity of a displacement field which is used for image transformation. Simulated datasets with known transformation fields are used to evaluate the presented method and demonstrate the potential of the system. Experiments with patient datasets show that superior results could be achieved compared to a pairwise image registration approach. Furthermore, it is shown that the method is suited to include prior knowledge about the data as the MRF system is able to model dependencies between the parameters of the optimization process.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement—Registration

1. Introduction

Acquisition of dynamic medical images is used to measure functional processes for early detection and diagnosis of diseases and pathologies. As an important part of this, perfusion imaging describes and quantifies the passage of fluids through blood vessels, the lymphatic system, organs or tissue. Signal acquisition is performed consecutively, depicting multiple instances of the same region of interest (ROI) over time, resulting in an additional function dimension t. The spatial domain of the ROI can either be two-dimensional or three-dimensional leading to 2D+t or 3D+t data.

2D ultrasonography (US) enables immediate and inexpensive examinations with high temporal resolution. It is well suited for imaging abdominal and thoracic organs. There are no contraindications and the patient is not exposed to radiation. US is also used for perfusion imaging employing contrast agents (CA) [CCA*08]. CA consists of gas-filled micro bubbles that have a high degree of echogenicity as they increase the US backscatter [PG11]. By acquiring 2D contrast enhanced US (CEUS) multi-frame sequences, propagation and contrast uptake after the injection of the CA can be measured to assess perfusion kinetics. This is used to delineate the vascular structure in normal and pathological tissue to detect tumors or metastases in various organs or to assess disease activity in Crohn's disease [KJG*04, NØH*09]. The perfusion is obtained by extracting and analyzing the perfusion kinetics of the blood in the tissue of interest from the acquired multi-frame data. The acquisition procedure of CEUS usually produces two parallel image sequences: standard brightness mode (b-mode) and the measured CA enhancement (Fig. 1a, 1b). CEUS imaging is able to produce a high temporal resolution ($\approx$ 10 frames per second) [CCA*08].

During hand-held CEUS examinations the sonographer normally holds the probe still in a particular position to image a slice of interest during CA administration. However, significant motion is still present in the data. On the one hand patient movement (e.g. caused by breathing) affects the probe position and orientation. On the other hand motion is induced intrinsically by perfusion, digestion and breathing.

While these motion types can be interpreted by well trained physicians [RTP*05], in computer-assisted analysis the different image frames of a time-dependent acquisition

DOI: 10.2312/PE/VMV/VMV12/151-158

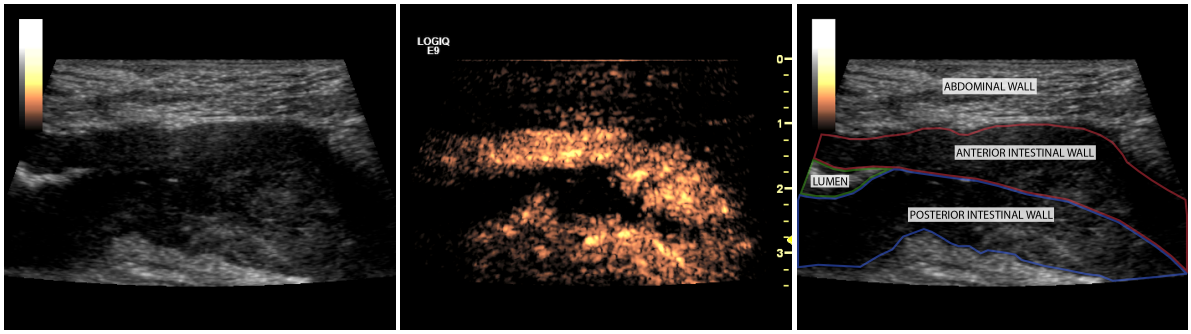

(a) (b) (c)

Figure 1: A representative frame of a CEUS acquisition of a patient with a stenosis of the small bowel due to Crohn's disease, showing b-mode data (a) and contrast data (b), the most important regions labeled in (c).

need to be aligned in order to extract valid perfusion parameters over time. This motion compensation in medical image analysis can be achieved by registration [HHH01].

Research presented in this work will be targeted on the formulation and implementation to approach this specific problem with CEUS data. Restrictions and constraints of the data have to be included as a priori knowledge when solving the registration problem to compensate for deficiencies within the data such as noise or small out-of-plane motion.

To compensate homogeneous motion influence within a predefined ROI linear translations are applied to each image in the sequence. As most of the image data contains large parts of soft tissue locally varying motion influence has to be expected as well. Most applications use free-form deformations combined with spatial constraints to restrict the motion to be locally smooth [LCKD*10]. The high temporal resolution of the data also suggests to temporally constrain transformation. We introduce a markov random field (MRF) formulation which is an excellent method to model these dependencies. We show that the approach is more robust against noise as a consequence of additional information introduced by the spatial and temporal smoothness constraints.

2. Related Work

Intra-modal registration of time-dependent data has been used in different applications. To calculate and analyze the deformation of the human heart Ledesma-Carbayo et al. presented a combined spatio-temporal registration procedure [LCKD*05, LCKD*10]. Similarity of the deformed 2D US frames was measured by the mean squared distance of all frames in the temporal sequence to a specified reference frame. Transformation parameters for B-Splines are found for all frames simultaneously restricting the parameters to be continuously smooth over time. This stabilizes the approach against outliers using prior knowledge in form of anticipated motion over time through a smoothness constraint. Another approach uses contour tracking on echocardiographic ultrasound images learning motion characteristics from training data [JNMPB99]. However, this implies that strong contour information is available and the motion influence is interpretable (e.g. reoccurring regularly) and can be formalized. Temporally constrained registration has also been used to be able to quantify dilation of the brachial artery in ultrasound image sequences by a Kalman filter [FLL03]. An important assumption for this is a gradual motion process over time to be modeled by the linear Kalman filter.

Optimization with MRF has been adapted recently to solve registration problems with spatial prior terms [GPK*07, KLY08, SKH08, MS12]. A MRF is an undirected graph with nodes representing the unknowns of the system. Nodes can take labels which express a certain configuration. In case of the registration problem this configuration represents the transformation parameters for compensation of motion effects. The configuration with the highest probability leads to a low energy at the nodes (denoted as singleton energy). Constraints can be included through edges. Edges combine exactly two nodes evaluating the probability of the connected nodes configurations (pairwise/doubleton energy). The overall MRF energy must be minimized to obtain the global best configuration of the system [Bes86, GPK*07]. As a specific number of possible configurations is required, the search space is discretized and encoded in a label set.

A congeneric problem of registration is optical flow calculation e.g in video data. Glocker et al. introduced a spatio-temporal framework where temporal continuity is achieved through edges connecting nodes over time in a MRF [GPK*08]. Hereby, local uncertainties are compensated through constraining the possible solution space. It was used to solve the morphing of two images over a specific number of intermediate steps using cubic B-Splines.

The advantage of MRF-based methods is that prior in-

formation of the scene, such as smoothness of transformation parameters, can be specified on a non-object basis. This makes it suitable for a large class of problems [Li94, KZ04]. Mahapatra et al. combine the registration task with the knowledge coming from segmentation to determine the elasticity of the transformation parameters for each location individually [MS12]. Additionally, the discrete search space allows the efficient computation of all transformation parameters in a single system without using derivative-driven objective functions [GPK*08]. MRF-based objective functions can be solved efficiently using graph cut-based approaches [BLRB01, BK04, KTP08].

Perfusion imaging of US data is a temporal sequence of images which may be disturbed by motion influence stemming from different sources. The CA enhancement as time-varying signal reveals perfusion which can only be quantified if correct spatial correspondence of temporal samples is achieved. The literature does not yet offer adapted methods for the scenario of automatic motion compensation of CEUS perfusion images, taking account for the low signal-to-noise ratio and the huge number of image frames to be registered.

3. Method

To approach the specific problem of motion compensation in a time dependent 2D CEUS image sequence discrete optimization using a MRF is employed. The spatial and temporal resolution of the measured CEUS data induce a dependency of transformation parameters. Corresponding parameters belonging to a local neighborhood establish spatial and temporal continuity. This can be included in the MRF-based system through edges between parameter nodes constraining the result of the final transformation parameters.

To avoid registration of 2D frames which do not depict the same slice (due to out-of-plane motion), temporal regions are defined containing frames with in-plane motion only and registration is applied to theses regions individually [SAN*11]. Additionally, the user has to specify a ROI (Fig. 5a, 5b, 5c) where the motion compensation is applied to. This ROI is determined in a representative frame of the sequence and can be applied to the rest of the image frames. It should not include discontinuities, so that the above assumptions for the spatial and temporal priors will hold.

As a preprocessing step the sequence is linearly registered (using a single translation vector for each image) with temporally constrained translation to remove global motion influence beforehand. This step reduces the non-linear registration task to be targeted on local motion only. As a result, the search space for the non-linear transformation parameters can be constrained to a few pixel translation for each location in the images. If the dataset does not exhibit global motion beforehand, the similarity term will only provide weak energies and the temporal constraint ensures continuity of the translation resulting in zero translation.

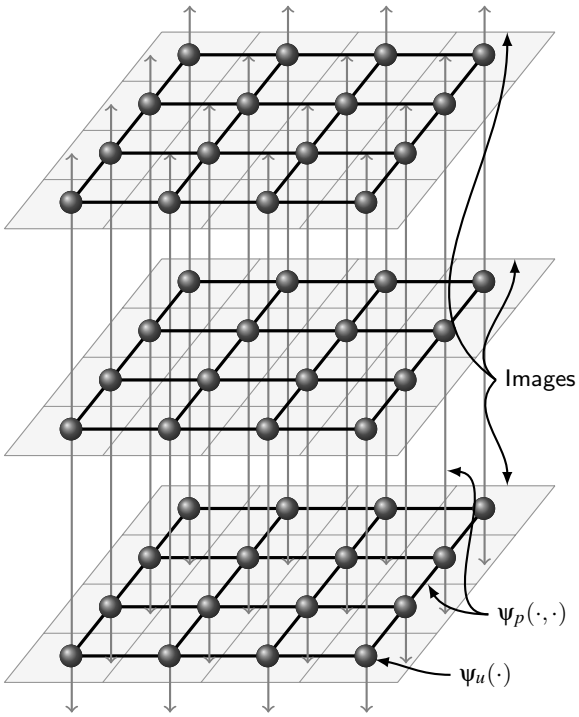

Figure 2: The MRF model represented by a graph with nodes being assigned a unary potential function ψ_u, spatial edges within the image frames and temporal edges between image frames (both assigned a pairwise potential function ψ_p).

The MRF model in general consists of nodes $v_i \in \mathcal{V}$ and edges $e \in \mathcal{E}$ of an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. An edge always connects exactly two nodes $e = (v_i, v_j) \in \mathcal{E}$ with $v_i, v_j \in \mathcal{V}$. The goal is to find the most probable configuration of random variables $\mathcal{X}$ within the system. This is determined by the probability of the random variables $x_i \in \mathcal{X}$ at respective nodes v_i at all cliques in the graph. The different types of cliques are defined by the neighborhood system N. The Hammersley-Clifford theorem [Bes86] stipulates the random variables $\mathcal{X}$ to be a MRF with respect to a neighborhood N if and only if the probability distribution of $P(\mathcal{X})$ is a Gibbs distribution:

$$P(\mathcal{X}) = Z^{-1} \times e^{-U(\mathcal{X})}. \quad (1)$$

Z is a normalizing constant and $U(\mathcal{X})$ is the energy defined by the sum of all different clique potentials depending on N. In image analysis and computer vision the problem is often regarded as an energy minimization (according to the terminology of similar problems) using a neighborhood of pairwise interaction only (clique number of 2 given by Von Neumann neighborhood). Nodes are assigned an unary potential function $\psi_u(\cdot)$ and edges are assigned a pairwise

potential function $\psi_p(\cdot,\cdot)$ (Fig. 2). The global energy of the MRF is defined by the sum of both potentials

$$E_{global} = \sum_{v_i \in \mathcal{V}} \psi_u(v_i) + \sum_{e=\langle v_i, v_j \rangle \in \mathcal{E}} \psi_p(v_i, v_j) \qquad (2)$$

and must be minimized in order to find the best solution according to the constraints made in the system by varying the labels of the nodes in the MRF.

3.1. MRF for registration

For image registration tasks the potential functions are used to evaluate different transformation parameters for a local neighborhood of the respective node. Hence, each node v_i in the MRF is assigned a label $l_k \in \mathcal{L}$ representing a transformation parameter. In our case, this is a translation in 2D defined by a vector $l_k = (t_x, t_y)$.

In the case of image registration the potential functions are used to evaluate different transformation parameters with respect to the overall problem. Unary potential functions $\psi_u(\cdot)$ represent the contribution of the current parameter to the fitting quality in terms of a similarity definition. Edges can be inserted between two nodes exhibiting a certain dependency or constraint (see Fig. 2). Pairwise potential functions $\psi_p(\cdot,\cdot)$ are used to model the dependencies. They evaluate the probability in terms of an energy value of two neighboring labels (transformation parameters).

Common choices for similarity calculation are Normalized Correlation and Mean Squared Distance. As both acquired sequences are recorded simultaneously, we use the images from b-mode to calculate the transformation and then apply the result to the contrast sequence. In b-mode data the intensity can be assumed to be at the same level for comparison, thus the simpler Mean Squared Distance measure is used. To obtain reliable information about observed motion a local area of a 7×7 pixel neighborhood M_{v_i} around the pixel of current node v_i is used to calculate similarity:

$$\psi_u(v_i) = \frac{1}{|M_{v_i}|} \sum_{v_m \in M_{v_i}} \left(H(v_m) - H(v_i)\right)^2 \qquad (3)$$

$$H(v) = I(P(v) + L(v)) \qquad (3a)$$

$$P(v_i) = (x, y) \quad x, y \in \Omega \qquad (3b)$$

$$L(v_i) = l_k = (t_x, t_y) \quad l_k \in \mathcal{L}, \qquad (3c)$$

where $I(\cdot)$ is the image function contributing the appropriate intensity values to given indices. $P(\cdot)$ represents the corresponding pixel indices in the image space Ω and $L(\cdot)$ represents the current label of given nodes v_i.

The formation of the unary potential function is done through similarity calculation of each frame to a predefined fixed image frame. This assures best fitting to a fixed basis

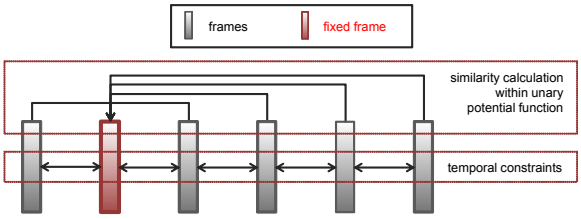

Figure 3: Scheme of the unary energy calculation to a predefined fixed frame.

under the restriction of temporal and spatial smoothness to the neighbor images both controlled by the pairwise term. (Fig. 3). As fixed frame of a sequence the frame with maximum average similarity to all other frames of the sequence is chosen [SNGT12]. To evaluate the pairwise potential energies between labels the Euclidean distance of corresponding translation vectors is calculated. Moreover, the pairwise energy ψ_p of the MRF energy term (see Eq. 2) consists of energies for temporal edges and spatial edges:

$$\psi_p = \psi_{pt} + \alpha \cdot \psi_{ps} \qquad (4)$$

$$\psi_{pt}(v_1, v_2) = \psi_{ps}(v_1, v_2) = d(P(v_1), P(v_2)), \qquad (5)$$

with $d(\cdot,\cdot)$ being the Euclidean distance of two pixel coordinates. The weighting parameter α has been set empirically to two. Varying this parameter has shown to be very robust towards the quality of the results.

The resulting MRF can be efficiently solved using the graph-cut based α-expansion algorithm [BK04]. It reaches the optimal solution for a two-label problem. The deviation to the optimal solution is bounded for multiple labels [KT07].

For the linear registration preprocessing step the MRF reduces to a markov chain (MC), as there is only one transformation parameter for each image of the sequence. Spatial continuity is given by default. Thus, nodes in the MC represent the images and temporal smoothness is achieved by the edges in the MC connecting neighbor images. This preprocessing registration reduces the overall motion in the sequence and, by association, the search space for the nonlinear registration step, where we assume that 5 pixel translations in each direction are sufficient to correct patient datasets for motion. This results in 121 possible configurations (labels) per node in the MRF.

4. Results and Evaluation

Results are produced using two different simulated datasets at different noise levels (Fig. 4, 5) to evaluate the performance of the different constraints. Additionally, three patient datasets containing between 200 and 800 time frames are

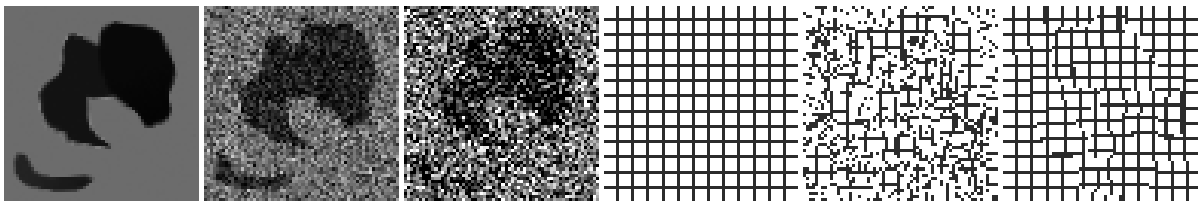

(a) (b) (c) (d) (e) (f)

Figure 4: Simulated dataset 1 with no noise (a), with a SNR level of 3 (b), with a SNR level of 1.5 (c) and a transformed grid image (d) with no constraints enabled (e) and spatial and temporal constraints used (f) for registration of images with SNR=3. It can be seen, that the constraint algorithm produces a smoother deformation field which is less susceptible to outliers (f).

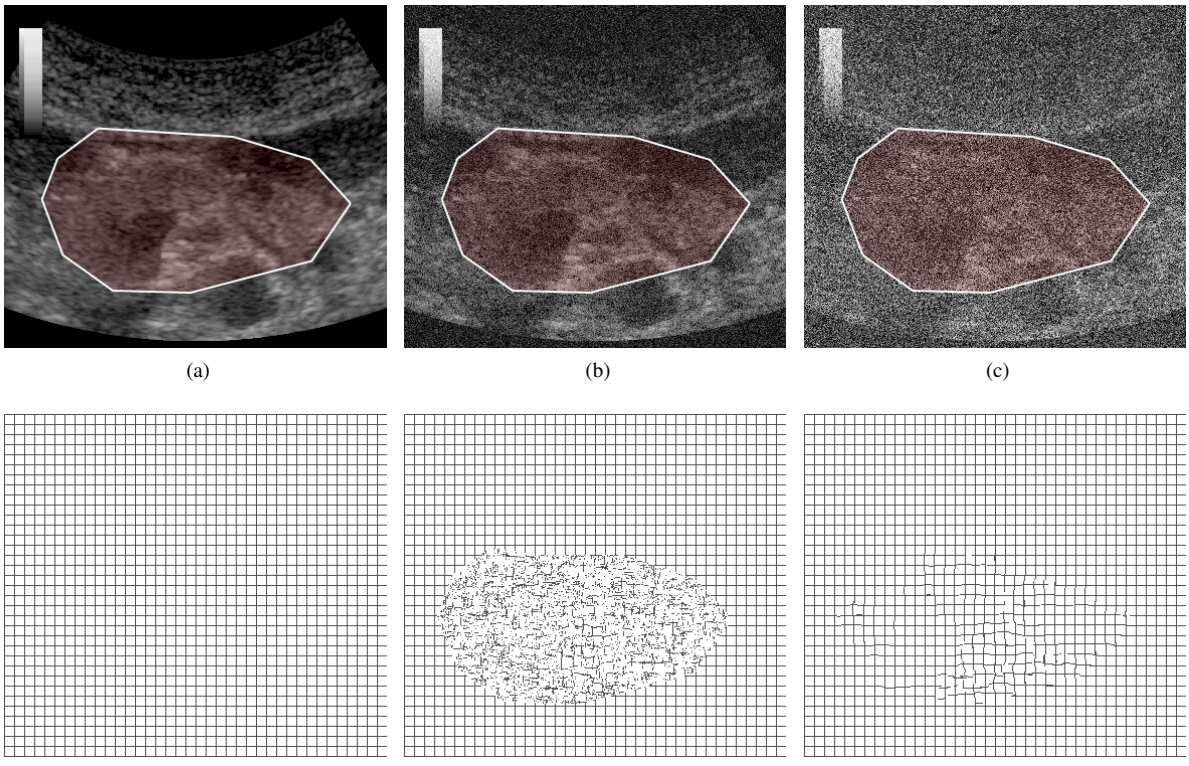

(a) (b) (c)

(d) (e) (f)

Figure 5: Simulated dataset 2 with no noise (a), with a SNR level of 3 (b), with a SNR level of 1.5 (c). The polygonal area shows the region of interest the motion compensation is targeted on. A grid image (d) is transformed with parameters determined with no constraints enabled (e) and spatial and temporal constraints used (f) for registration of images with SNR=3. Both calculated deformation fields (e) and (f) produce comparable similarity measures. However, spatially and temporally constraining the approach leads to smoother and more plausible transformations.

used to produce first results and document the potential of the presented approach.

The simulated datasets contain five time frames and have been transformed with a known deformation field. The first simulated dataset (Fig. 4) shows different objects based on organic shapes, the second one (Fig. 5) is a single frame extracted from a b-mode sequence which has been transformed with a known deformation field. Additive gaussian noise at different signal-to-noise-ratios (SNR=3 and SNR=1.5, ratio of signal mean and the standard deviation of the noise) has been added to both datasets (Fig. 4b, 4c, 5b, 5c). After recovering the transformation parameters with our method, the

average distance to the ground truth parameters can be calculated to measure the performance of the system. The simulated deformation fields contain parameters of zero translation. As they have to be recovered correctly by the algorithm as well they are included in the distance measure. However, it is more likely to recover zero translation when optimization is constrained spatially and temporally. Therefore, the deviation from ground truth is also calculated for all non-zero translation parameters in the deformation field.

The evaluation of the motion compensation of simulated datasets shows that the more prior knowledge is introduced the better the true deformation field is recaptured. This is confirmed by both, the average distance of all transformation parameters in the deformation field and non-zero parameters in the field (Fig. 6). For both simulated datasets the method using no constraints, temporal constraints, spatial constraints and both constraints achieves 0.75 pixel, 0.47 pixel, 0.45 pixel and 0.30 pixel average distance to all parameters of the ground truth deformation, respectively. Regarding non-zero parameters average distances of 1.47 pixel, 1.29 pixel, 1.05 pixel and 0.99 pixel are obtained, respectively. For the simulated images without added noise the method was able to almost reproduce the correct parameters, independent from the used constraints (Fig. 6). This shows that if the similarity term (unary potential) provides very reliable information, neither spatial nor temporal constraints are necessary. At lower SNR levels the improvement induced by the constraints is higher as the deficiency of the data can be compensated by coupling information from the neighborhood (Fig. 6a-6d). Although the major contribution originates from the spatial smoothness terms, the temporal smoothness term enhances the overall accuracy in almost all cases. For both simulated datasets the improvement aggregates to 0.11 pixel for all parameters and 0.01 pixel considering non-zero parameters in the deformation field.

Evaluation of patient datasets is performed on three different ROI per dataset in the b-mode and contrast data of the sequence. These regions are chosen to represent main areas of interest (e.g. Fig. 1c). First, the standard deviation within the regions is measured using the b-mode data. As a result of the registration it should decrease, as different tissue types are aligned and not mixed over a certain amount of frames. Second, the smoothness of the contrast enhancement signal in contrast data is measured. The smoothness is defined as the mean absolute difference (MAD) of the signal over time. This is an indicator of improved contrast data correspondence over the time sequence. Although the signal is still influenced by noise and speckle, MAD decreases with improved registration.

The measurements are performed before registration and after registration with the proposed method. Additionally, a classic registration approach performing pairwise registration of the sequences is tested as well for comparison [SNGT12]. Results of this experiment are shown in Tab. 1 and indicate that the presented method is able to produce superior results compared to the pairwise frame registration, which does not solve the problem concurrently by the help of information from neighboring frames. The MRF-based approach achieves an overall improvement for the three patient datasets of 18.1 % compared to 10.6 % of the pairwise image registration. The curve smoothness also leads to superior results, 3.2 % compared to 1.5 %. The visual registration quality could be enhanced by the proposed approach leading to a smoother appearance of transitions between the time frames.

The patient datasets contain between 300 and 800 frames. Depending on the number of frames regarded for temporal registration the calculation time for the process between 20 and 60 minutes on a standard quad core processor system. This is acceptable considering that the registered data is required for post-diagnostic analysis.

5. Conclusions and Future Work

In this work an approach for motion compensation of ultrasonic image sequences has been presented. The optimization scheme uses a MRF formulation allowing to include prior knowledge about the specifics given by the acquisition procedure to perform image registration. In our case, this is the spatial and temporal smoothness of the deformation field of the sequence of images. Optimization is conducted for registration of all image frames simultaneously and not between image pairs. Experiments show the influence of the different constraints of parameters at various noise levels and demonstrate the robustness against noise influence through incorporating prior information in form of spatial and temporal continuity of the transformation parameters. A performance improvement compared to pairwise image registration has been established as well.

The system is suited to incorporate more dependencies and constraints between its parameters. The variety of applications for MRF optimization is an example of this flexibility [Li94]. US image sequences have a low signal-to-noise ratio resulting in poor data quality and artifacts. To improve the search for correct transformation parameters, intensity distribution information can be used to generate more stable features for registration represented by labels [SNGT12]. We want to investigate if this property could be used to improve the results further as the changes of segment labels over time can be penalized in the MRF optimization scheme. It induces the configuration of the transformation parameters to generate better fitting quality in terms of the original data term (unary potential) and the energy minimization of the segment labels (pairwise potential). Another starting point for additional prior information is the use of the contrast sequence to improve fitting quality of transformation parameters.

To improve the calculation time of the approach, especially with regard to the planned extensions, we want to par-

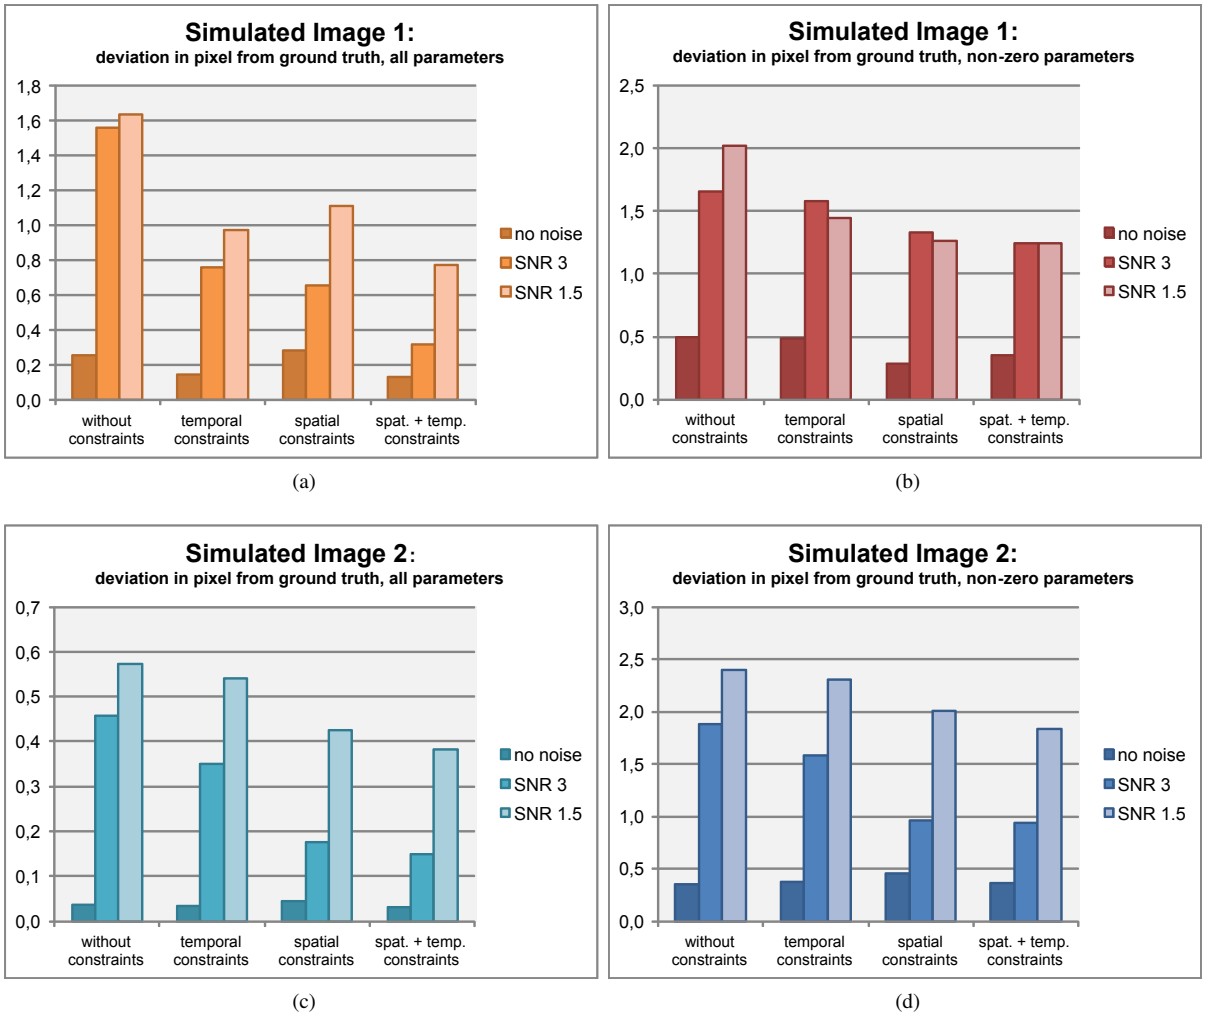

Figure 6: Performance of the presented approach using two simulated datasets and the spatial and temporal constraints disabled and enabled. The graphs show the average distance in pixel to the ground truth of all translation vectors of the deformation field (a, c). Additionally, the average distance in pixel to ground truth of all translation vectors $\neq (0, 0)$ are depicted in (b) and (d). In general, the distance to ground truth diminishes with constraints. The effect is bigger for datasets with strong noise influence.

allelize the calculation of energies for the MRF optimization by using either higher multi-core CPU environments or a GPGPU implementation.

6. Acknowledgements

The authors thank Kim Nylund and Odd Helge Gilja from the University of Bergen and Haukeland University Hospital in Bergen, Norway for providing the medical data.

References

[Bes86] BESAG J.: On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society, Series B (Methodological) 48*, 3 (1986), 259–302. 2, 3

[BK04] BOYKOV Y., KOLMOGOROV V.: An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 9 (Sept. 2004), 1124–37. 3, 4

[BLRB01] BOYKOV Y., LEE V., RUSINEK H., BANSAL R.: Segmentation of dynamic ND data sets via graph cuts using markov models. In *Medical Image Computing and Computer assisted Intervention (MICCAI)* (2001), Springer, pp. 1058–1066. 3

[CCA*08] CLAUDON M., COSGROVE D., ALBRECHT T., BOLONDI L., BOSIO M., CALLIADA F., CORREAS J., DARGE K., DIETRICH C., D'ONOFRIO M., ET AL.: Guidelines and good clinical practice recommendations for contrast enhanced ultrasound (CEUS). *Ultraschall in der Medizin 29*, 1 (2008), 28–44. 1

Table 1: Standard deviation within regions of interest in b-mode data and perfusion curve smoothness within contrast data is measured for three patient datasets before registration, after classic pairwise registration [SNGT12] and registration with the presented MRF-based approach. Improvement compared to values before registration are indicated in percent.

dataset	no. 1	no. 2	no. 3	overall avg.
standard deviation in b-mode data				
before registration	20.78	21.92	27.66	
classic pairwise registration	19.66 (6.1 %)	18.09 (19.2 %)	25.39 (6.7 %)	10.6 %
new MRF-based registration	18.85 (8.9 %)	16.11 (29.0 %)	23.35 (16.4 %)	18.1 %
perfusion curve smoothness in contrast data				
before registration	3.35	2.71	3.56	
classic pairwise registration	3.33 (0.7 %)	2.68 (0.8 %)	3.50 (3.0 %)	1.5 %
new MRF-based registration	3.29 (1.5 %)	2.56 (5.9 %)	3.48 (2.3 %)	3.2 %

[FLL03] FRANGI A. F., LACLAUSTRA M., LAMATA P.: A registration-based approach to quantify flow-mediated dilation (FMD) of the brachial artery in ultrasound image sequences. *IEEE Transactions on Medical Imaging 22*, 11 (Nov. 2003), 1458–69. 2

[GPK*07] GLOCKER B., PARAGIOS N., KOMODAKIS N., TZIRITAS G., NAVAB N.: Inter and intra-modal deformable registration: continuous deformations meet efficient optimal linear programming. *Information Processing in Medical Imaging 20*, July 2007 (Jan. 2007), 408–20. 2

[GPK*08] GLOCKER B., PARAGIOS N., KOMODAKIS N., TZIRITAS G., NAVAB N.: Optical flow estimation with uncertainties through dynamic MRFs. In *Computer Vision and Pattern Recognition (CVPR)* (2008), pp. 1–8. 2, 3

[HHH01] HAJNAL J. V., HAWKES D. J., HILL D. L. G.: *Medical image registration*. 2001. 2

[JNMPB99] JACOB G., NOBLE J. A., MULET-PARADA M., BLAKE A.: Evaluating a robust contour tracker on echocardiographic sequences. *Medical image analysis 3*, 1 (Mar. 1999), 63–75. 2

[KJG*04] KLEIN D., JENETT M., GASSEL H., SANDSTEDE J., HAHN D.: Quantitative dynamic contrast-enhanced sonography of hepatic tumors. *European Radiology 14*, 6 (2004), 1082–1091. 1

[KLY08] KWON D., LEE K., YUN I.: Nonrigid image registration using dynamic higher-order MRF model. In *European Conference on Computer Vision, ECCV* (2008), pp. 373–386. 2

[KT07] KOMODAKIS N., TZIRITAS G.: Approximate labeling via graph cuts based on linear programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 8 (Aug. 2007), 1436–53. 4

[KTP08] KOMODAKIS N., TZIRITAS G., PARAGIOS N.: Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies. *Computer Vision and Image Understanding 112*, 1 (Oct 2008), 14–29. 3

[KZ04] KOLMOGOROV V., ZABIH R.: What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence 26*, 2 (Feb. 2004), 147–59. 3

[LCKD*05] LEDESMA-CARBAYO M., KYBIC J., DESCO M., SANTOS A., SUHLING M., HUNZIKER P., UNSER M.: Spatio-temporal nonrigid registration for ultrasound cardiac motion estimation. *IEEE Transactions on Medical Imaging 24*, 9 (2005), 1113–26. 2

[LCKD*10] LEDESMA-CARBAYO M., KYBIC J., DESCO M., SANTOS A., UNSER M.: Cardiac motion analysis from ultrasound sequences using non-rigid registration. In *Medical Image Computing and Computer-assisted Intervention (MICCAI)* (2010), vol. 32, pp. 889–896. 2

[Li94] LI S.: Markov random field models in computer vision. In *Computer Vision ECCV* (1994), vol. 801, Springer, pp. 361–370. 3, 6

[MS12] MAHAPATRA D., SUN Y.: Integrating segmentation information for improved MRF-based elastic image registration. *IEEE Transactions on Image Processing 21*, 1 (Jan 2012), 170–83. 2, 3

[NØH*09] NYLUND K., ØDEGAARD S., HAUSKEN T., FOLVIK G., LIED G. A., VIOLA I., HAUSER H., GILJA O. H.: Sonography of the small intestine. *World Journal of Gastroenterology 15*, 11 (2009), 1319. 1

[PG11] POSTEMA M., GILJA O. H.: Contrast-enhanced and targeted ultrasound. *World Journal of Gastroenterology 17*, 1 (2011), 28. 1

[RTP*05] RENAULT G., TRANQUART F., PERLBARG V., BLEUZEN A., HERMENT A., FROUIN F.: A posteriori respiratory gating in contrast ultrasound for assessment of hepatic perfusion. *Physics in Medicine and Biology 50*, 19 (2005), 4465–80. 1

[SAN*11] SCHÄFER S., ANGELELLI P., NYLUND K., GILJA O. H., TÖNNIES K.: Registration of ultrasonography sequences based on temporal regions. In *Proc. of 7th Intl. Symp. on Image and Signal Processing and Analysis (ISPA 2011)* (Dubrovnik, Croatia, 2011), pp. 749–759. 3

[SKH08] SHEKHOVTSOV A., KOVTUN I., HLAVÁČ V.: Efficient MRF deformation model for non-rigid image matching. *Computer Vision and Image Understanding 112*, 1 (Oct. 2008), 91–99. 2

[SNGT12] SCHÄFER S., NYLUND K., GILJA O. H., TÖNNIES K.: Motion compensation of ultrasonic perfusion images. In *Medical Imaging 2012: Ultrasonic Imaging, Tomography, and Therapy* (2012), vol. 8320, SPIE. 4, 6, 8