Interactive Model-based Image Registration

Thomas Schiwietz[1,2] and Joachim Georgii[1] and Rüdiger Westermann[1]

[1] Technische Universität München
[2] Siemens Corporate Research
Email: {schiwiet,georgii,westermann}@in.tum.de

Abstract

We present an interactive technique for the registration of captured images of elastic and rigid body parts in which the user is given flexible control over material specific deformation properties. Our method can effectively handle arbitrary stiffness distributions and it achieves an accurate matching without sacrificing the physical correctness of the simulated deformations. The algorithm consists of three steps, which are performed iteratively until an optimal spatial mapping is determined: First, the optical flow is used to predict an initial image transformation. Second, a priori knowledge of the deformation model is used to refine the predicted field. A physics-based filter operation generates a transformation that is consistent with the model of linear elasticity. Third, the process is repeated using the displaced template image. To achieve accurate image deformations we employ implicit multigrid solvers using finite differences (optical flow) and finite elements (linear elasticity). The robustness and accuracy of our method is validated using synthetic and real clinical data composed of heterogeneous materials exhibiting different stiffness characteristics.

1 Introduction

Especially in medical applications there is an ever growing interest in techniques that can accurately relate information in multiple data sets resulting from different measurements. Image registration techniques are designed for that purpose as they try to find the correspondence between images of the same anatomical structure taken under different conditions (e.g. different relative camera-patient position, different methods of acquisition, etc.). Such techniques compute a spatial mapping of structures into a common coordinate system or aim at compensating individual structural characteristics and differences due to movements over time.

One of the main challenges in image registration is to accurately model and simulate the deformation behavior of the structures contained in the data. Scanned body parts are usually composed of highly heterogeneous tissue types and in particular an elastic modulus with a dynamic range of several orders of magnitude is not unusual to be found in a single scan. Registration techniques capable of dealing with such data require to adjust the stiffness of the transformation to permit realistic displacements over all parts in the data. Physics-based deformation models then have to be considered to realistically simulate the structural changes.

In this work, we present an algorithm for image registration that addresses the aforementioned requirements by using a predictor-corrector approach to iteratively compute an optimal spatial mapping. The predicted deformation field can be computed from image pairs using any standard evaluation method. The corrector step implements a finite-element solver for the Lagrangian equation of motion simulating linear elasticity. It takes as input the predicted vector field and modifies it according to the underlying physical model. As all parts of the algorithm run interactively, intuitive control mechanisms to flexibly adjust the deformation properties of specific regions in the data are integrated.

In particular, we utilize a GPU-based random walker [14] for assigning specific stiffness values to parts of an image. While it is generally possible to assign these properties using a painting tool on a per-pixel basis this is a tedious and time-consuming task not feasible in clinical practice. Instead, we use the segmentation algorithm to assign stiffness on a per-object basis. The segmentation algorithm extracts parts from the image automatically by single mouse clicks or strokes on top of the object, and it then fills the object's interior with a user-defined stiffness.

VMV 2007 H. P. A. Lensch, B. Rosenhahn, H.-P. Seidel, P. Slusallek, J. Weickert (Editors)

1.1 Related Work

With respect to the assumed image transformation, image registration techniques can be classified into parametric and non-parametric approaches. Comprehensive surveys of these approaches along with a number of application areas can be found in [20, 28, 21]. Parametric approaches impose specific restrictions on the transformation, e.g. requiring the transformation to be rigid, polynomial or affine. Non-parametric, i.e. non-linear, approaches, on the other hand, are far more flexible in the type of transformation they compute. Such transformations rely on additional constraints like the regularization of the displacement field, which can be enforced by explicit smoothing of the deformation field [1, 22, 25]. Over the last years special emphasis has been put on the development of variational PDE approaches to obtain approximate solutions to the non-linear registration problem [7]. In this case an additional smoothing term is considered in the global energy function to be minimized. In particular, physics-based regularizers based on linear elastic and viscous fluid models have been shown to be effective in penalizing unrealistic local and global displacements [6, 27, 26, 4]. While finite difference methods are frequently applied to discretize the linear elasticity equations [21], finite element methods are so far rarely applied to the best of our knowledge. This is due to the fact that they are computationally more complex, and thus result in poor performance rates [15]. However, as has be shown by Georgii and Westermann [12], multigrid algorithms can efficiently be applied to speed up the solution process of finite element techniques.

Curvature-based constraints for non-linear registration problem have been discussed in [11, 16]. While effective in computing deformations that comply with the incorporated regularization models, only a few approaches have attempt to also take heterogeneous materials exhibiting varying stiffness into account [9, 19, 18].

2 Method

We propose a model-based approach for the computation of a non-linear transformation of a template image T onto a reference image R. A priori knowledge about the physical deformation model is exploited to make the transformation consistent with this model. Our algorithm proceeds in multiple steps, each of which is performed interactively to accommodate steering of model-based parameters by the user. The major parts of our algorithm are illustrated in Figure 1.

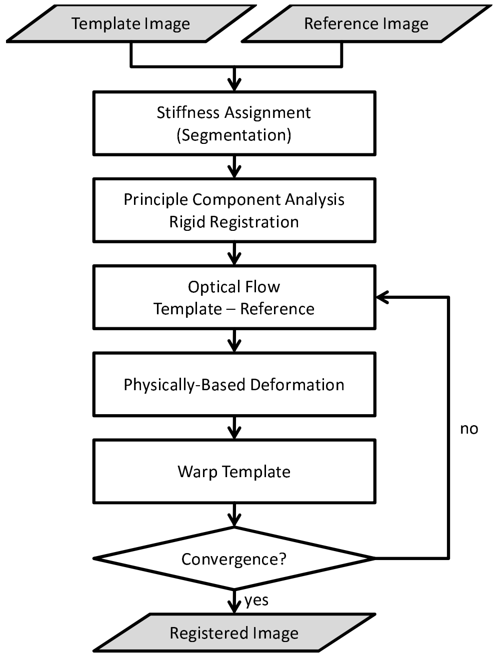

Figure 1: Overview of major parts of the registration system. A template and a reference image are registered using our algorithm: After stiffness assignment the images are rigidly registered by PCA. Then, our iterative loop estimates a vector field using the optical flow method. This field is corrected by physically based image deformation. The loop stops if the convergence criterium is met.

We utilize the Random Walker segmentation algorithm [14] to automatically extract objects from the image, for which we then specify the respective stiffness value. For our purposes, non-binary segmentation algorithms like the Random Walker provide attractive properties since the segmentation results yield a smooth transition between the segmented object and the background. A smooth stiffness distribution is required by the finite element method, since material stiffness is only approximated C^0-continuous, and thus large jumps can introduce instabilities. The smooth transition is highly desirable for stiffness assignment because the segmentation result can be mapped directly to stiffness using a transfer function. In this way, abrupt stiffness changes in the simulation grid can be avoided at no extra cost, enabling realistic deformations at boundary transitions.

In a pre-process a rigid PCA (principle component analysis) registration is performed to obtain an initial mapping between T and R. Therefore, the covariance matrix of both images is computed by weighting pixel-positions with the intensity values. An eigenvalue decomposition of that matrix reveals the relative rotation and scaling of the two images. The relative translation is the distance between the mean positions of both images.

2.1 Gridding

From the segmented and pre-registered template a regular simulation grid consisting of quadrangular elements is constructed by placing exactly one grid vertex at each pixel center. Pixel colors and associated stiffness values are respectively assigned as vertex and element attributes. Given such a grid in the reference configuration $x \in \Omega$, a deformed grid is modelled using a displacement function $u(x), u : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ yielding the deformed configuration $x + u(x)$. The transformed template image is then generated by rendering the deformed grid onto a regular pixel grid. This image is used in the next iteration to calculate the optical flow.

2.2 Deformation Estimate

An initial deformation estimate is computed using classical optical flow as introduced by [17]. The main idea behind this algorithm is to minimize a global cost function representing the rate of change of image brightness from one image to the other. As the optical flow estimates the apparent motion of brightness patterns in two images it is clear that this particular kind of deformation estimate cannot be applied in general for multi-modal registration. However, the optical flow works fine for single modalities, where all images are generated with the same scanning conditions. In the multi-modal case, deformation estimates based on common similarity measures like mutual information have to be favored [21].

Since techniques based on the optical flow are considered to be rather slow due to the numerical complexity of the employed solvers, in a number of research projects considerable effort has been put in the development of advanced numerical techniques like multigrid schemes. In this work, we make use of the implementation by Bruhn et al. [4] to significantly speedup the prediction of an initial deformation field.

2.3 Model-based Correction

Rather than penalizing the optical flow with a model-based regularizer as proposed in [23], the optical flow field is used as an external force field driving the deformation of structures contained in the template image. The deformation is based on a linear elasticity model, where the dynamic behavior is governed by the Lagrangian equation of motion. An implicit multigrid solver presented in [13] is extended to efficiently compute the resulting displacements.

In contrast to previous registration approaches using finite differences, it is worth noting that our method employs a finite-element discretization. In this way, improved physical accuracy is achieved, but we have to pay for that by an increasing numerical complexity. As our timings show, however, the multigrid solver we have developed performs favorable to respective finite difference techniques.

By using a model-based approach as described it is clear, that the simulated displacement field enforced by the elasticity equation differs from the optical flow field. Physically speaking, the optical flow field is corrected towards a displacement field that complies with the underlying model.

2.3.1 Linear Elasticity Model

If the object to be simulated obeys to the model of linear elasticity the dynamic behavior is governed by the Lagrangian equation of motion

$$M\ddot{u} + C\dot{u} + Ku = f \tag{1}$$

where M, C, and K are respectively known as the mass, damping and stiffness matrices. u consists of the linearized displacement vectors of all vertices and f are the linearized force vectors applied to each vertex. It is worth noting that in linear elasticity we only consider material that has a linear relationship between how hard it is squeezed or torn (stress) and how much it deforms (strain). This relationship is expressed by the material law, which is accounted for by the stiffness matrix K.

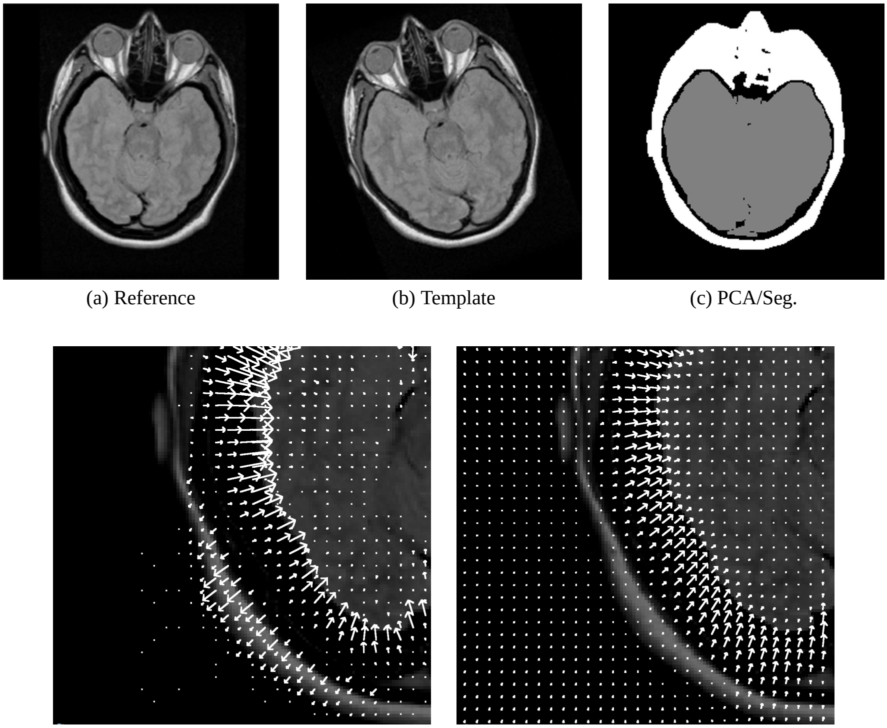

(a) Reference (b) Template (c) PCA/Seg.

(d) Prediction (e) Correction

Figure 2: The template image (b) is registered to the reference image (a) by first applying a PCA rigid registration (c). Then, stiffness values are assigned to segmented parts (c). The iterative loop predicts a vector field (d) using the optical flow method and corrects it by our physical deformation model (e) until convergence is reached.

Equipped with any suitable discretization finite element methods typically built these matrices by assembling all element matrices to yield a sparse system of linear equations. For the details on the discretization process as well as the numerical schemes used to solve the resulting system let us refer the reader to [2, 3].

2.3.2 Finite Element Method

To improve simulation accuracy we have integrated quadrangular elements with bilinear nodal basis functions into our approach. Quadrangular elements consist of 4 supporting nodes $\underline{v}_k$, thus interpolating the deformation in the interior as

$$u(x) = \sum_{k=1}^{4} N_k(x)\underline{u}_k$$

where

$$N_k(x) = c_0^k + c_1^k x_1 + c_2^k x_2 + c_3^k x_1 x_2$$

and $\underline{u}_k$ is the displacement of the k-th node. The coefficients c_i^k can be easily derived from $N_k(\underline{v}_k) = 1$ and $N_k(\underline{v}_i) = 0$ if $k \neq i$. The shape functions $N_k(x)$ and its derivatives are needed to build the final element matrices, which are then assembled into the global stiffness matrix,

By using quadrangular elements the overall number of elements is reduced. This allows for a slightly faster assembly process of the global system matrix. Furthermore, the semi-quadratic interpolation scheme increases the simulation accuracy and thus improves physical correctness.

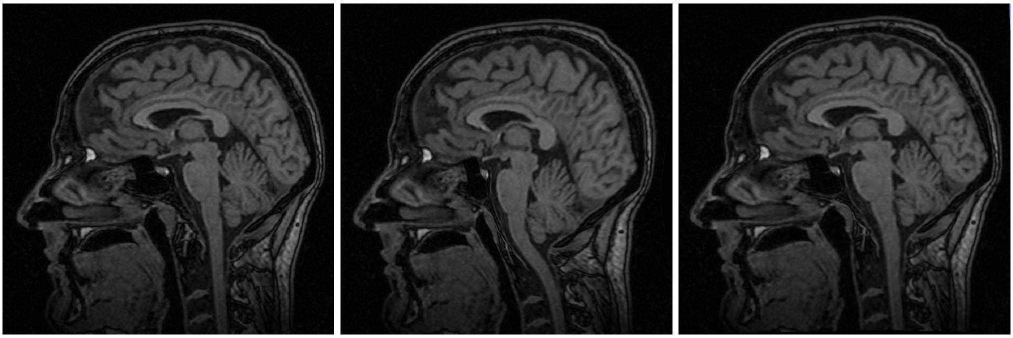

Figure 3: The first image show the reference configuration. The second image was generated by a physically deformation using the finite element method. The third image shows the result of our registration algorithm.

3 System Assembly

In the following, we will show how we interconnect the different parts of our algorithm into one system. An illustration of our algorithm is depicted in Figure 2.

Once the user has selected two images to register, stiffness values are assigned to all pixels of the template image using a painting tool, a segmentation algorithm (Random Walker), or other automatic assignment methods. Once the stiffness assignment is done, the images are pre-registered using PCA.

All the upcoming steps are then performed automatically, including mesh generation, deformation estimate, model-based correction and image warping. These steps are performed iteratively until the template image matches the reference image with respect to any suitable metric, i.e. the mean square deviation of the current per-vertex deformations to the previous ones.

The iterative registration loop starts by computing the optical flow field between the current warped template image and the reference image. Next, optical flow vectors are applied as external forces to the finite element vertices. Bilinear interpolation of the four closest optical flow vectors is used to obtain the force for one finite element vertex in the deformed grid. Note that the external forces are accumulate in each step while the previous accumulated force is damped by a small percentage. The average force vector of all finite element vertices is subtracted from all vertices in order to keep the mesh in place. Since we cannot compute the exact force field required for convergence in one single step, we choose an iterative approach using a dynamic simulation. Thus, in every time step we only do a small step into the direction predicted by the optical flow force field; this loop is iterated until convergence is finally achieved.

To correct the predicted optical flow field via the linear elasticity module we simulate the deformation over one time step until the multigrid solver converges.

Resulting displacements are send to the graphics card where the simulation mesh is updated accordingly. This mesh is then rendered onto a regular pixel grid using hardware supported interpolation to generate the warped template image used in the next iteration. As the update step is entirely performed on graphics hardware, and because sending displacement values from the CPU to the graphics card does not introduce any time constraints, its contribution to the overall runtime is negligible. Note, that both the optical flow and the physical correction are computed on the CPU.

4 Results

4.1 Performance Measurement

For our experiments we used a computer with an Intel Core 2 6600 2,4 GHz CPU, equipped with 2 GB RAM and an NVIDIA GeForce 8800 GTX graphics card. The following table shows the performance (one iteration) of the most important steps of our algorithm on various grid sizes.

As can be seen in Table 1, one single iteration loop runs at interactive speed. Therefore, param-

Grid size	128^2	256^2	512^2
OF	18.3 ms	79 ms	342 ms
Defo	28.8 ms	116 ms	478 ms
Update	5.4 ms	17 ms	52 ms
Total	52.5 ms	212 ms	872 ms

Table 1: Timing statistics for one iteration of the proposed algorithm. In the first row, the time required by the optical flow computations (OF) are listed. Then, timings for the elastic deformation engine (Defo) and for the update of the vertices, forces and template image (Update) are shown. Due to the multigrid approach timings scale linearly with grid size.

eters like force field scaling, stiffness values and stiffness distribution can be changed interactively within the iteration loop. Our examples shown in the next section demonstrate that five to thirty iterative steps are required to achieve the final result. Note, that it is in general possible to use smaller grids for deformation and optical flow calculations. In all examples, the mesh resolution was set to 128^2.

4.2 Examples

On the last page we show three results that have been generated using our approach. Figure 5 shows a synthetic data set demonstrating varying stiffness distribution. Next, we created a semi-synthetic data set from a MR scan of a brain by manually shrinking the grey matter in the reference image (see Figure 6). Finally, we evaluate our method on a real data set shown in Figure 7.

Our results demonstrated the accuracy of realistic material deformation at interactive rates and the effective application to image registration problems. The registration process converges if the accumulated force field does not change anymore. This is the case if all vectors produced by the optical flow estimate are below a certain threshold, e.g. 0.25 pixel width. Then, the internal forces of the deformed material are in balance with the external forces of the predictor.

4.3 Validation

Schnabel et. al. have developed a general framework for validation of non-rigid registration algorithms [24]. Based on this framework, we have validated our approach. We generated various image pairs with a physically accurate finite element method and registered them by means of our algorithm. Figure 3 shows, that our algorithm achieves highly accurate results even for large deformations. The mean intensity difference between theses images is below 10^{-2}.

4.4 Discussion

In comparison to Modersitzki [21], our method distinguishes in several parts. First, we use a finite element discretization rather than a finite difference method. This allows a much better approximation of the partial derivatives and thus improves both stability and accuracy. Second, the external forces applied to the deformation engine are computed using an optical flow model rather than a gradient based approach. Therefore, the elastic deformation is driven more precisely to its final state, and thus a significant smaller number of iteration steps is required until convergence. This effect is illustrated in Figure 4. As a sidenote let us mention that both the optical flow method [8] and the finite element method [24] are used for validation of non-rigid registration techniques. By incorporating these techniques into our registration algorithm, high accuracy can be achieved.

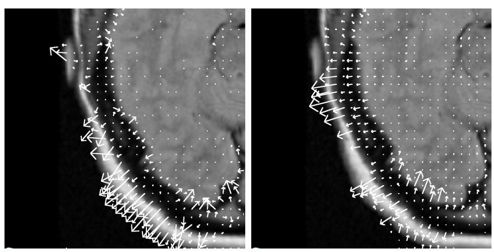

Figure 4: Comparison of image gradient (left) and optical flow (right) as deformation estimator. As can be seen, the optical flow yields generally much more accurate vector fields.

Another benefit of our system is that both the prediction and the correction stage are based on standard methods that can be used as black boxes. As a consequence, the predictor as well as the corrector can be easily adjusted to different methods that are either faster or specificly adapted to the underlying image modalities. Especially, the deformation engine used has several advantages: It can efficiently handle different kinds of material parameters (heterogeneous, soft and stiff material) and strain formulations (linear Cauchy strain, corotated Cauchy strain). Due to its implicit nature, stability can be guaranteed during the whole convergence process.

Especially for large deformations one might wish to address the problem that the linear approximation of the strain tensor lacks rotational invariance, thus resulting in unrealistic deformations once elements get rotated out of their reference configuration. Fortunately, as has been shown in [13], co-rotated elements [10] could be easily integrated into the registration process by adapting the deformation engine to the co-rotated strain measurement.

Note, that the method is especially applicable to detect regions of the image where the optical flow model and the elastic model contradict each other. For example, this might be the case when registrating images before and after a bone fracture if the bone parts are not healed at exactly the same cut surface. In those cases, the registration process can be either manually controlled, or a-priori knowledge can be incorporated into the whole process. In this example, one would allow the bone to deform such that both images can match exactly.

Furthermore, due to this system design, our method can be easily adapted to 3D registration problems. Prediction of a 3D vector field can be achieved via 3D optical flow or similar methods and physical deformation engines based on volumetric elements are already fast [13].

5 Conclusion

In this work we have described a physics-based technique for image registration that enables flexible control over the kind of deformations to perform. In extending previous methods we have proposed a sophisticated predictor of an initial deformation field, and we have shown that highly accurate and stable finite-element methods can be integrated into interactive scenarios. In this way, image registrations can be achieved at high performance while material specific deformation properties as they arise in reality are simulated.

In the future, we will investigate more advanced optical flow algorithms, e.g. the one proposed by Bruhn et al. [5]. Furthermore, we will try to achieve faster convergence of our registration algorithm; one possibility is to optimize the way the optical flow field is applied as external forces to the finite element vertices. Additionally, we plan to improve the algorithm by integrating mixed boundary conditions (forces and displacements) into the dynamic elasticity problem efficiently. This allows us to enforce displacements at specific points, e.g. bones.

The extension of our registration algorithm to 3D is straightforward enabling accurate registrations of volumetric bodies. To deal with multi-modal registration problems, we can replace our deformation estimator by an appropriate method like mutual information. Note, that due to the design of our system the deformation engine is in no way affected by this change.

References

[1] Luis Alvarez, Joachim Weickert, and Javier Sanchez. Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision*, 39(1):41–56, 2000.

[2] Klaus-Jurgen Bathe. *Finite Element Procedures*. Prentice Hall, 2002.

[3] Morten Bro-Nielsen and Stephane Cotin. Real-time volumetric deformable models for surgery simulation using finite elements and condensation. *Computer Graphics Forum*, 15(3):57–66, 1996.

[4] Andrés Bruhn, Joachim Weickert, Christian Feddern, Timo Kohlberger, and Christoph Schnörr. Variational optical flow computation in real-time. *IEEE Transactions on Image Processing*, 14:608–615, 2005.

[5] Andrés Bruhn, Joachim Weickert, Timo Kohlberger, and Christoph Schnörr. A multigrid platform for real-time motion computation with discontinuity-preserving variational methods. *Int. J. Comput. Vision*, 70(3):257–277, 2006.

[6] Gary E. Christensen, Richard D. Rabbitt, and Michael I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996.

[7] Ulrich Clarenz, Marc Droske, and Martin Rumpf. Towards fast non-rigid registration. *DFG 1114*, December 2002.

[8] James Cooper. Optical flow for validating medical image registration. In *IASTED International Conference on Signal and Image Processing*, pages 502–506, 2003.

[9] Valerie Duay, Pierre-François D'Haese, Rui Li, and Benoit M. Dawant. Non-rigid registration algorithm with spatially varying stiffness properties. In *ISBI*, pages 408–411, 2004.

[10] Olaf Etzmuß, Michael Keckeisen, and Wolfgang Straßer. A fast finite element solution for cloth modelling. In *Pacific Conference on Computer Graphics and Applications*, pages 244–251, 2003.

[11] Bernd Fischer and Jan Modersitzki. Curvature based image registration. *Journal of Mathematical Imaging and Vision (JMIV)*, 18(1):81–85, 2003.

[12] Joachim Georgii and Rüdiger Westermann. Interactive simulation and rendering of heterogeneous deformable bodies. In *Vision, Modeling and Visualization 2005*, pages 381–390, 2005.

[13] Joachim Georgii and Rüdiger Westermann. A multigrid framework for real-time simulation of deformable volumes. In *Proceedings of VRIPHYS Workshop*, pages 50–57, 2005.

[14] Leo Grady and Gareth Funka-Lea. Multi-label image segmentation for medical applications based on graph-theoretic electrical potentials. In *ECCV 2004 Workshops CVAMIA and MMBIA*, number LNCS3117 in Lecture Notes in Computer Science, pages 230–245. Springer, May 2004.

[15] Alexander Hagemann, Karl Rohr, H. Siegfried Stiehl, Uwe Spetzger, and Joachim M. Gilsbach. A biomechanical model of the human head for elasitc registration of MR-images. In *Bildverarbeitung fur die Medizin*, pages 44–48, 1999.

[16] Stefan Henn. A full curvature based algorithm for image registration. *Journal of Mathematical Imaging and Vision (JMIV)*, 24(2):195–208, 2006.

[17] Berthold Horn and Brian Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.

[18] Florian Jäger, Jingfeng Han, Joachim Hornegger, and Torsten Kuwert. A variational approach to spatially dependent non-rigid registration. In *Proceedings of the SPIE, Medical Image Processing*, volume 6144, pages 860–869, 2006.

[19] Sven Kabus, Astrid Franz, and Bernd Fischer. On elastic image registration with varying material parameters. In *Bildverarbeitung für die Medizin (BVM)*, pages 330–334. Springer Verlag, 2005.

[20] J.B.Antoine Maintz and Max A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.

[21] Jan Modersitzki. *Numerical Methods for Image Registration*. Oxford university press, New York, 2004.

[22] Esther Radmoser, Otmar Scherzer, and Joachim Weickert. Scale-space properties of regularization methods. In *Scale-Space Theories in Computer Vision*, pages 211–222, 1999.

[23] Paul Ruhnau and Christoph Schnörr. Optical stokes flow estimation: an imaging-based control approach. *Experiments in Fluids*, 42(1):61–78, 2007.

[24] J. A. Schnabel, C. Tanner, A. D. Castellano-Smith, A. Degenhard, M. O. Leach, D. R. Hose, D. L. G. Hill, and D. J. Hawkes. Validation of non-rigid image registration using finite element methods: Application to breast mr images. *IEEE Transactions on Medical Imaging*, 22:238–247, 2003.

[25] Marius Staring, Stefan Klein, and Josien P.W. Pluim. Nonrigid registration with adaptive, content-based filtering of the deformation field. In *Proceedings of SPIE Medical Imaging: Image Processing*, volume 5747, pages 212 – 221, February 2005.

[26] Yongmei Wang and Lawrence H. Staib. Integrated approaches to non rigid registration in medical images. In *Proceedings of WACV'98*, pages 102–108, 1998.

[27] Joachim Weickert and Christoph Schnörr. A theoretical framework for convex regularizers in pde-based computation of image motion. *International Journal of Computer Vision*, 45(3):245–264, 2001.

[28] Barbara Zitova and Jan Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, October 2003.

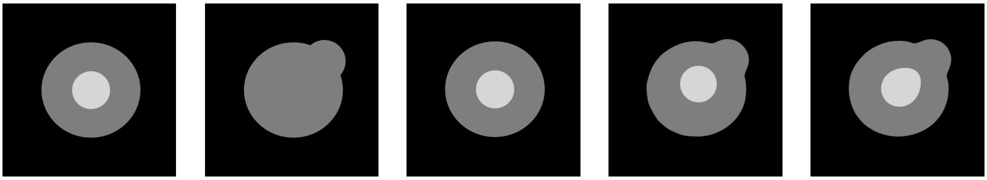

(a) Template (b) Reference (c) All hard (d) Mixed (e) All soft

Figure 5: This synthetic image simulates an object with two different tissue types. The template image (a) is registered to the reference image (b) using different tissue stiffness. Image (c) shows a uniformly hard stiffness distribution over the object so that no deformation is possible. Image (d) shows the result of assigning soft stiffness to the outer ring of the object. Now, the bump in the outer ring is matched while the hard white core remains undeformed. Image (e) shows results from a uniformly soft tissue throughout the object. Note, how the white core deforms.

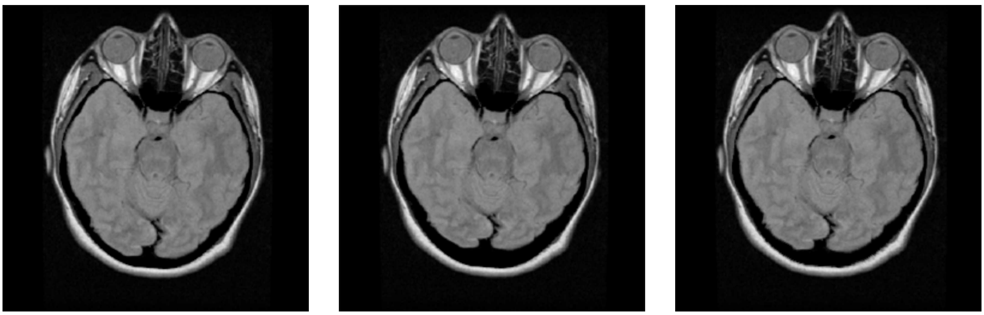

(a) Template (b) Reference (c) Registered

Figure 6: This Figure shows a semi-synthetic data set of a manually shrunken brain. The template image (a) is registered to the reference image (b). Image (c) shows the registered result using a 256×256 grid. The stiffness was set to 10^8 for the skull, 10^6 for the grey matter and 10^4 for dark area in between. Note, how the brain compacts, the soft tissue expands and the bones remain stiff. The images register in 5 iterations (approximately 1 second) with an average per pixel error of 10^{-2}.

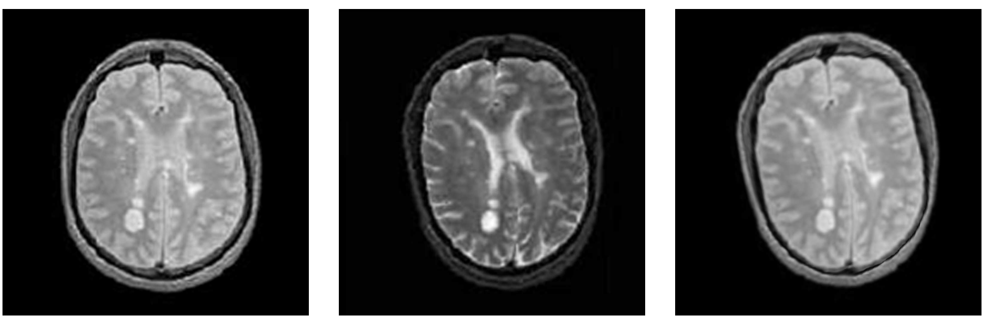

(a) Template (b) Reference (c) Registered

Figure 7: Note the different image contrasts of the template image (a) and the reference image (b). Our algorithm registers the two images accurately (c) in 0.5 seconds using a grid resolution of 128×128.