

Fast Deformable Registration of 3D-Ultrasound Data Using a Variational Approach

Darko Zikic¹, Wolfgang Wein¹, Ali Khamene², Dirk-André Clevert³, and Nassir Navab¹

¹ Chair for Computer Aided Medical Procedures (CAMP)
Technische Universität München, Germany
{zikic,wein,navab}@cs.tum.edu

² Imaging and Visualization Dept., Siemens Corporate Research
Princeton, NJ, USA
ali.khamene@siemens.com

³ Department of Clinical Radiology,
University Hospitals Munich-Grosshadern, Germany
dirk.clevert@med.uni-muenchen.de

Abstract. We present an intensity based deformable registration algorithm for 3D ultrasound data. The proposed method uses a variational approach and combines the characteristics of a multilevel algorithm and the properties of ultrasound data in order to provide a fast and accurate deformable registration method. In contrast to previously proposed approaches, we use no feature points and no interpolation technique, but compute a dense displacement field directly. We demonstrate that this approach, although it includes solving large PDE systems, reduces the computation time if implemented using efficient numerical techniques. The performance of the algorithm is tested on multiple 3D US images of the liver. Validation is performed by simulations, similarity comparisons between original and deformed images, visual inspection of the displacement fields and visual assessment of the deformed images by physicians.

1 Introduction

Being a non-invasive and cost effective modality, 3D ultrasound (US) has certain advantages over other modalities such as CT, MR or PET imaging. However, as confirmed by physicians, particular clinical applications of 3D US such as follow-up studies can be improved if the deformations between subsequent examinations can be undone.

Furthermore, since ultrasound is a real-time technique, it can be used in image-guided interventions. In this context, precise deformable mapping of the pre- and intra-operative images would facilitate the navigation process.

To this end, fast and accurate deformable registration methods for 3D US are required. This is a more challenging task compared to modalities such as CT or MRI since speckle, occlusion and other artifacts are inherent to the imaging technique. Hence, the registration method has to be able to cope with these

modality-specific problems. Recently, a number of papers addressed these difficulties by proposing deformable registration algorithms.

Foroughi and Abolmaesumi [1] present a feature based approach, using attribute vectors as a similarity measure. These vectors contain image intensities smoothed by a Gaussian filter and computed on multiple resolution scales. A set of leading points is computed in the reference and the template image according to the distinction of the attribute vectors. Within this set, corresponding points are found by search in a preset range which yields a sparse displacement field. In the final step, the dense displacement field is computed by propagating the displacement values at the leading points using Gaussian kernels. Being among the most recent work on deformable 3D ultrasound registration, this paper reports runtimes of 430 seconds for images of the size $185 \times 113 \times 199$.

In [2] Xiao et al. present a scheme, using block matching and local statistics for estimating the local deformation of the data. A Bayesian regularization method is used for constraining the deformation. In a final step the dense deformation field is computed using B-splines. Here the algorithm runtimes were about 360 seconds for $149 \times 154 \times 121$ images. However, the weaker performance of the hardware (as of 2001) has to be taken into account.

Krücker et al. [3] developed a two-stage, sub-volume based method. In the first stage, a global affine registration is computed. In the second stage the sub-volumes are registered locally, while the connectivity of the whole volume is achieved by an interpolation using thin-plate splines. Different similarity measures are used, and the reported runtimes were below 300 seconds (again it should be considered that hardware available in 2001 was used).

All these methods have in common that the displacement is computed only at certain points of the volume by finding correspondences between selected features and that they use an interpolation technique to determine the global displacement field. Feature based approaches however have certain intrinsic drawbacks. First, not all of the available information is used for registration. Furthermore, the search for good features is a complex issue. And finally, the search range for finding correspondences must be set, which results in either lower confidence levels for the results or higher computation times.

In this work we overcome these difficulties by directly computing the dense displacement field. To this end, we employ the known variational approach for deformable registration. To the best of our knowledge, this is the first use of such an approach for 3D ultrasound data. Furthermore, we show that although this approach involves solving a large system of partial differential equations (PDE), an efficient numerical implementation can reduce the computation time compared to the previously proposed methods.

Deformable registration based on the variational minimization of a functional is a well-known technique. A number of different similarity measures and possible regularization terms was proposed in literature, e.g. [4, 5]. Many useful references are also to be found in the Computer Vision community, where the 2D Optical Flow problem leads to the same equations [6, 4]. Especially the efficient

solution of the resulting PDE by using a multigrid solver is described in [7]. An introduction to multigrid methods can be found in [8].

2 Method

In the following we briefly present a deformable registration method based on the variational approach. The method yields a dense displacement field as the solution to the deformable registration problem.

In order to allow for larger displacements and to speed up the algorithm, we employ a multi-resolution pyramid. The pyramid is computed by reducing the resolution by a factor of 2 along each dimension from level to level by applying a discrete approximation to the Gaussian filter. On every level of the pyramid, the variational registration is performed using the result of the next coarser level as initial guess. Besides the general and well known advantages of this approach like accounting for large displacements, it is in particular beneficial when applied to ultrasound images.

Because of the speckle inherent to ultrasound images, these reveal large structures better when they are smoothed. Hence performing a large number of iterations on the coarser levels seems reasonable. This intuition is supported by the work presented in [1], where the entries of the attribute vectors were image intensities smoothed by a Gaussian filter and computed on different resolution scales. This directly corresponds to solving the problem on the coarser resolutions. Especially the fact that the intensities from the finest level are included only after Gaussian filtering suggests that in our approach no iterations on the finest level have to be performed. Furthermore, the results in [6] demonstrate that the variational registration should always be performed on smoothed images. In the validation section we confirm the benefit resulting from using coarser resolution levels.

Since in this work we consider deformable registration, we assume that the input images are already rigidly registered. Otherwise, they would be deformed in order to account for rigid motion which would lead to incorrect deformations. The rigid registration of the images, computing the pure rotation and translation parameters was performed using cross-correlation as similarity measure.

We now briefly present the variational method used on every level of the Gaussian pyramid.

2.1 Variational Deformable Registration

We set $\Omega = [0, 1]^3 \subset \mathbb{R}^3$ as the spatial domain considered, and define the reference $R : \Omega \rightarrow [0, 1] \subset \mathbb{R}$ and template image $T : \Omega \rightarrow [0, 1] \subset \mathbb{R}$.

The registration problem is to compute the deformation function $\varphi : \Omega \rightarrow \Omega$, $\varphi(x) = x + u(x)$, s.t. the reference $R(x)$ and the deformed template $T(\varphi(x))$ are possibly similar with respect to a certain similarity measure. Here $u : \Omega \rightarrow \Omega$ is the displacement field.

Since the registration problem is ill-posed we do not solve it directly but we minimize a regularized energy functional

$$J[u] = D[u] + \alpha S[u] \quad (1)$$

with $D[u]$ a dissimilarity measure and a regularizer (smoothing operator) $S[u]$ which penalizes unlikely deformations [5]. The influence of the regularizer is determined by the scalar α . The solution of $u = \arg \min_{\hat{u}} J[\hat{u}]$ is the displacement field that represents the solution of the registration problem. The regularization term enforces that u is a diffeomorphism. This is an important property for the displacement field, since this class of functions maintains the topology of the images.

We use the sum of square differences (SSD) as dissimilarity measure [5]:

$$D[R, T, u] = \int_{\Omega} (R(x) - T(x + u(x)))^2 dx \quad (2)$$

The regularizer is chosen such that strong variations in the displacement field are penalized, i.e.

$$S[u] = \int_{\Omega} |\nabla u_x|^2 + |\nabla u_y|^2 + |\nabla u_z|^2 dx \quad (3)$$

This is the well-known diffusion regularization term, compare e.g. [4, 5]. We solve the minimization problem (1) by applying the calculus of variations. This yields a partial differential equation (Euler-Lagrange equation), the solution to which is a necessary condition for a stationary point of the functional. With the above selection of the dissimilarity and regularization term we get the following non-linear partial differential equation.

$$\alpha \Delta u = f(u) \quad \text{with} \quad f(u) = (R - T(\varphi)) \nabla T(\varphi) \quad (4)$$

This equation is the well known Poisson equation. For the solution of the system we assume the displacement function to be zero at the boundary, i.e. we impose homogeneous Dirichlet boundary conditions. For solving this non-linear PDE we apply a modified fix-point iteration, where only the incremental updates are regularized, compare [5].

2.2 Implementation Details

For the actual solution, we discretize the PDE (4) by the finite difference method using central differences [5]. After this step, the resulting system of linear algebraic equations $Au = f$ remains to be solved in every step of the fix-point iteration. This system is sparse, yet the size of the problem is extremely large. The matrix A is $N \times N$ with $N = 3n$ and n the number of the voxels in the volume. For a realistic problem size of $256 \times 256 \times 256$ we end up with a system of the size $N \approx 5 \cdot 10^7$. Hence an efficient solver is needed. For such large systems one of the most efficient and general solvers is the multigrid method, which

has the theoretical complexity of $O(N)$ [8]. Furthermore, the multigrid solvers are very well understood for the Poisson problem and were already applied to registration problems before [5].

For solving the linear algebraic equations $Au = f$ we apply the Full Multigrid cycle (FMG) [8]. Experiments show that the number of pre- and post-smoothing Gauss-Seidel iterations for FMG can be chosen as low as 1.

For the fast implementation of the algorithm, it is required that the volumes have the size $S_1 \times S_2 \times S_3$ with $S_i = 2^{n_i} + 1$, hence the images are padded if necessary. Furthermore, we use the Euler reference frame for the displacement field. This enables simple and fast warping of the images within the algorithm. This is important since it is the most time-consuming step of the algorithm besides the solving of the equation system. We also assume that the diffeomorphism condition is sufficiently fulfilled, which allows for computing the inverse of the actual displacement field. This results in computing the registration from the reference to the template image. The force term then changes to $-(R - T(u)) \nabla R$ which has the advantage of having to compute the gradient only once for the whole algorithm instead of once in every fix-point iteration.

The alpha values are chosen differently for each resolution level, such that the performed regularization is weaker on the finer levels. However, the same values can be used for all data sets. Empirical tests let us select $\alpha = 0.005$ for the finest level and increase it by factors 2 – 4 for the coarser levels.

The stopping criterion for the number of fix-point iterations per level can be set manually, or a threshold ϵ for the relative decrease of dissimilarity can be defined. If the threshold is set in a conservative way (ca. 2% for coarsest level and 0.5% for finest), the number of iterations resembles the manually selected values.

3 Results and Validation

The volumes have been acquired from a healthy liver using a Siemens Sonoline Antares ultrasound system. We used a C5F1 curved-array mechanically swept 3D/4D ultrasound transducer at 3.5Mhz. Several volumes with large field-of-view were recorded at deep and medium inhale, and different orientations. For the evaluation we used 5 pairs of data sets. The digital scanline data stored by ultrasound device was converted to rectilinear volumes of size 257^3 . The voxel size for the data sets was approximately 0.87, 0.72 and 0.59 mm/voxel along the respective dimensions.

The implementation was done in C++, and the experiments were performed on an AMD Athlon PC with 2.2GHz and 2 GB RAM.

The runtimes of the algorithm range from 10 to 40 sec. (av. 27 sec.), depending on the similarity of the images and thus on the number of iterations at respective resolutions levels. The computation time for one iteration is 15 sec. for the 257^3 resolution, 2.0 sec. for 129^3 , 0.25 sec. for 65^3 and 0.02 sec. for 33^3 .

3.1 Simulation Studies

For this test we use a known displacement field u to deform a given real US image T in order to artificially compute the reference image R . The images R and T are then used as input for the registration algorithm. This way, we are able to use the field u as ground truth, to which the result u' of the registration procedure can be compared.

For the generation of a displacement field u , we use two different alternatives u_G and u_R . For the first option u_G we use a Gaussian $G_{\sigma,s}$ for every dimension of the field. Here, $G_{\sigma,s}$ has the standard deviation σ , and is scaled such that the maximum value equals s . We find that this gives deformations similar to the ones observed in the data. More challenging, the second option for the field u_R is to use the result of a previous registration process.

The results of this experiment are summarized in following table, where μ and σ denote the respective mean and standard deviation. In the first four columns we give the information describing the ground truth u , while the next three columns describe the difference between the ground truth and the computed displacement field $u - u'$.

type of u	$max(u)$ [mm]	μ_u [mm]	σ_u [mm]	$max(u - u')$	$\mu_{u-u'}$	$\sigma_{u-u'}$
$u_G: G_{0,2,5}$	9.035	0.096	2.174	2.341	0.052	0.170
$u_G: G_{0,2,20}$	36.138	0.386	0.543	11.243	0.161	0.461
u_R ,	20.091	2.344	3.180	7.926	0.497	0.574

For the evaluation of the results, the mean error and the standard deviation of the error are the most important values. We observe a clear decrease in the mean error. The small standard deviation about the error mean describes the quality of the solution. It states that errors much larger than the mean and especially the maximum error rarely occur.

Furthermore, this test allows us to measure the contribution of the single resolution levels to the computed field u' . The results support the decision to reduce the number of iterations on the finer levels since the contribution to the ground truth field u is negligible and this at a high computational cost. Since the presentation of all possible combinations would be a five-dimensional problem, we list only some of the measurements, which illustrate the contribution of iterations on respective levels. In the left-most column, the table lists the number of iterations per level with the coarsest level first and the finest last. The tests were performed using ground truth field $G_{0,2,20}$ from the table above.

Row No.	Iterations per Level	$max(u - u')$	$\mu_{u-u'}$	$\sigma_{u-u'}$
1	(20 30 30 10 5)	11.243	0.161	0.461
2	(20 30 30 10 0)	11.277	0.164	0.466
3	(20 30 30 0 0)	12.343	0.179	0.529
4	(20 30 0 0 0)	22.427	0.338	1.164
5	(0 0 40 0 0)	14.210	0.203	0.657
6	(20 30 40 10 0)	10.261	0.161	0.438

The rows 1-4 show some results where we continuously decrease the number of iterations on the finer levels. In row 5 we observe that comparable results can be achieved by performing more iterations only on the middle resolution level. Row 6 finally demonstrates that a combination of lower levels can be superior in terms of accuracy to using levels on the finest resolution.

3.2 Similarity Measurements

For the 5 data sets an increase in the similarity between the reference and the template before and after registration was measured. Here, we do not only measure the gain in the mean sum of square differences, but also compute Normalized Cross-Correlation (NCC) and Mutual Information (MI). Simply evaluating the SSD might be biased, since this measure is used for the registration itself. Comparison to more general metrics like NCC and MI [4, 3] confirms the validity of using the SSD as a similarity measure for the registration of ultrasound images. In order to measure the gain in the similarity s we give the ratio of the differences of the similarity before and after the registration to the original similarity: $r = (s_{after} - s_{before}) / abs(s_{before})$. In order to perform the comparison, the dissimilarity SSD is transformed to a similarity measure. For the SSD, the values for r ranged between 32.82% and 48.47% with an average of 40.78%. The results for NCC were 5.99%-10.94% with average 8.85% and for MI 21.21%-39.14% with average 30.49%. Additionally, we observe a strong correlation between the gain in the SSD and in NCC and MI.

3.3 Visual Assessment

Since the deformable registration is an inherently ill-posed problem, the inspection and visual assessment of the deformed images by a trained physician is of extreme importance for the validation.

First, the gain in the similarity and thus the benefit of the method for the medical use was verified. Besides using plain difference images, we deploy different colors (red and blue) for the two images and display the differences by fading between the images. This has the advantage that not only the similarity of the images but also the deformation can be perceived to a certain extent.

Second, the structure of the deformed data set was inspected to assure that the method did preserve the anatomy of the data. This test was also positive. Furthermore it was confirmed that the speckle pattern in the original data sets was not destroyed by the deformation. This is important since the speckle carries important diagnostic information about the imaged tissue.

Last, we display the computed deformation fields in order to verify that no apparently wrong displacements are computed. Standard visualizations of the displacement by displaying sampled vector fields or the deformed grid fail in 3D. The display is either too complex for the validation or the sampling is very low. Thus we display one slice of the deformed image and visualize the corresponding displacement vectors. This way, the origin of the voxels in the plane is visualized. This can be done with a much higher sampling than if the




Fig. 1. Results of the deformable registration: the first column contains the reference, the second the template, and the third the deformed template image. The first row shows the results in the YZ and the second in the XZ plane. Corresponding slices are shown. It must be considered that apparently large deformations in the displayed slice are actually caused by smaller deformations along the axis orthogonal to the slice.

whole 3D field were visualized at once. This approach has the further advantage that the displacement is visualized simultaneously with the image data. So not only the coherence of the field but also its relationship to the anatomy can be verified.

Using the above tests, the significance of the proposed method as well as its validity were confirmed a physician.

4 Conclusion and Future Work

In this paper we presented a deformable registration method for ultrasound data. To the best of our knowledge this is the first time a variational approach is used for 3D US data. It allows a direct computation of a dense displacement field even for large displacements and overcomes the difficulties inherent to feature-based methods. Besides, we demonstrate that using a Gaussian pyramid and efficient PDE multigrid solvers results in the currently fastest implementation. We evaluate the algorithm with a number of different methods and also present useful techniques for visual assessment of 3D deformable registration methods.

As future work we will evaluate the presented method on a large number of different data sets. The cooperation with medical partners will be extended




Fig. 2. Left: visualization of the 3D displacement field, any slice from the US volume can be selected. The two right images depict the difference images of the reference and template before and after registration.

in order to integrate the method into existing medical workflows or to define new possible applications. We plan to test the performance of other similarity measures and regularization operators within the variational framework on ultrasound data [4, 5] and for inter-modality registration. Finally, we plan to tackle the specific US problem of dependence on scanning position. This can be different images depending on the scanning position, due to e.g. shadows from air or bones. We plan to integrate the results from [9] in order to solve that problem.

References

1. Foroughi, P., Abolmaesumi, P.: Elastic registration of 3d ultrasound images. In Duncan, J.S., Gerig, G., eds.: MICCAI. Volume 3749 of Lecture Notes in Computer Science., Springer (2005) 83–90
2. Xiao, G., Brady, M., Noble, J.A., Burcher, M., English, R.: Non-rigid registration in 3D free-hand ultrasound imaging of the breast. In Lemke, H.U., Vannier, M.W., Inamura, K., Farman, A.G., Doi, K., eds.: CARS. Volume 1230 of International Congress Series., Elsevier (2001) 1199–1200
3. Krucker, J.F., LeCarpentier, G.L., Fowlkes, J.B., Carson, P.L.: Rapid elastic image registration for 3-D ultrasound. *IEEE Trans. Medical Imaging* **21** (2002) 1384–1394
4. Hermosillo, G.: Variational Methods for Multimodal Image Matching. PhD thesis, Université de Nice - Sophia Antipolis (2002)
5. Modersitzki, J.: Numerical methods for image registration. Oxford University Press (2004)
6. Álvarez, L., Weickert, J., Sánchez, J.: Reliable estimation of dense optical flow fields with large displacements. *International Journal of Computer Vision* **39** (2000) 41–56
7. Bruhn, A., Weickert, J., Feddern, C., Kohlberger, T., Schnörr, C.: Variational optical flow computation in real time. *IEEE Transactions on Image Processing* **14** (2005) 608–615
8. Briggs, W.L., Henson, V.E., McCormick, S.F.: A multigrid tutorial. 2. ed. edn. SIAM, Society for Industrial and Applied Mathematics (2000)
9. Wein, W., Roeper, B., Navab, N.: Automatic registration and fusion of ultrasound with ct for radiotherapy. In: MICCAI 2005 Proceedings. Lecture Notes in Computer Science, Springer (2005)