World Scientific
www.worldscientific.com

ROBUST MULTISCALE DEFORMABLE REGISTRATION OF 3D ULTRASOUND IMAGES

IOANNIS PRATIKAKIS[§]

*Institute of Informatics and Telecommunications,
National Center for Scientific Research "Demokritos", 15310 Athens, Greece
ipratika@iit.demokritos.gr*

CHRISTIAN BARILLOT[*], PIERRE HELLIER[†] and ETIENNE MEMIN[‡]

[§]*IRISA*
[*]*CNRS*
[†,§]*INRIA*
[‡]*Universitaire Rennes I*
[*,†,‡,§]*Project ViSTA,
Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France*
[*]*cbarillo@irisa.fr*
[†]*phellier@irisa.fr*
[‡]*memin@irisa.fr*

In this paper, we embed the minimization scheme of an automatic 3D non-rigid registration method in a multiscale framework. The initial model formulation was expressed as a robust multiresolution and multigrid minimization scheme. At the finest level of the multiresolution pyramid, we introduce a focusing strategy from coarse-to-fine scales which leads to an improvement in the accuracy of the registration process. A focusing strategy has been tested for a linear and a non-linear scale-space. Results on real 3D ultrasound images are discussed.

Keywords: 3D ultrasound; non-rigid registration; multiscale.

1. Introduction

Many image-guided operation systems are developed for minimally invasive surgery especially in the field of neurosurgery. However, they may not accurately depict the position of the specific anatomical structures due to the intraoperative movements of the brain, so-called the brain shift. This shift is caused by a variety of surgical factors including Cerebrospinal Fluid (CSF) drainage, tumor resection, hyperventilation, edema and hemorrhage. This phenomenon may change the geometry of brain structures and it may become a serious problem since preoperative images, like CT and MR, cannot be used anymore as a reference. The registration of preoperative

MRI to intraoperative 3D ultrasound is the main motivation of the presented work. Intraoperative 3D ultrasound imaging can provide real-time images of the brain during surgery and enables us to determine the extent of the intraoperative structure shift between the stage prior to and during surgery. For this purpose, several US acquisitions are necessary and particularly, one has to be done before the opening of the dura which will provide us an imaging of the same topology of structures that can be identified in the preoperative MRI. For the sake of clarity, at this point we have to mention that in this work we have not dealt with the problem of MR and US registration. The construction and use of such an intraoperative ultrasound set-up has also been reported by other authors.[1,2,3]

The problem of registration in ultrasound images has been treated by different researchers. Morcy and Von Ramm[4] investigated the implementation of a correlation search scheme to estimate the 3D motion vectors and they have demonstrated the advantages over 2D correlation search using the Sum Absolute Difference (SAD) as a similarity measure. Strintzis and Kokkinidis[5] introduced a maximum likelihood block matching technique which correspond to an accurate statistical description of ultrasound images. In the work of Yeung *et al.*[6] a multi-level motion model-based approach to ultrasonic speckle tracking has been developed that addresses the inherent trade-offs associated with traditional single-level block matching methods. Furthermore, an adaptive mesh has been proposed for non-rigid tissue motion estimation from ultrasound image sequences.[7] A deformable blocking matching algorithm has been developed which takes into consideration both similarity measures and strain energy caused by mesh deformation. Pennec *et al.*[8] disseminated results regarding 3D ultrasound registration using the demon's algorithm and a straightforward minimization of the sum of square of intensity differences criterion. This registration approach along with an exploitation of the temporal continuity of the deformation led to a suggestion for a tracking algorithm.[9] Arbel *et al.*[10] followed a strategy based on cross-correlation measurements to register preoperative US to intraoperative US. The preoperative data are "pseudo-ultrasound" images that were based on preoperative segmented MRI and were constructed using the predicted appearance of neuroanatomical structures in ultrasound. As an extension to the standard Bayesian image analysis paradigm, King *et al.*[11] proposed a new technique that incorporates a multiscale approach. This new technique is demonstrated by applying it to the problem of compensating for soft tissue deformation of pre-segmented surfaces for imageguided surgery using 3D ultrasound. The solution is regularized using the knowledge of the mean and Gaussian curvatures of the surface estimate. Accurate estimates of the deformed surfaces were successfully computed using the algorithm, based on prior probabilities defined using a minimal amount of human intervention.

Non-rigid registration can be considered as a motion estimation problem which can be solved by minimizing an objective function. This function is the energy which usually consists of two terms. The first term represents the interaction between the unknown variables and the data while the second one explores some kind of

prior information. Within this context, Mémin and Pérez[12] proposed a motion estimator which makes use of the optical flow constraint along with an associated smoothness regularizing prior. Both terms have been constructed with an outlier rejection mechanism, originated from robust statistics. For the minimization of their function they used a multiresolution and multigrid scheme. The multiresolution part is dedicated to grasp large displacements while the multigrid approach is invoked for accelerating the estimation. The extension of this work to treat 3D data has been done by Hellier *et al.*[13]

In this paper, we embed the above mentioned minimization scheme in a multiscale framework aiming to improve the estimates by making them less sensitive to the noise of acquisition. In the same spirit, Weber and Malik[14] proposed a model for multiscale motion estimation. They convolved an image sequence with a set of linear, separable spatiotemporal filter kernels and applied a robust version of the total least squares on the filtered responses in a two step method. In the first step, every scale group individually form an estimate for the velocity. Checking of the validity of the estimate is performed and all the remaining valid estimates are combined into a second total least squares formulation. This approach cannot guarantee the computation of a dense deformation field due to production of holes (lack of estimates) in the case that different spatial scales overlap regions of different motions due to a motion boundary. Niessen *et al.*[15] reported a reconciliation of optical flow and scale-space theory. They computed both zeroth and first order optic flow at multiple spatial and temporal scales and they applied a scale selection criterion which attributes in each pixel the optic flow at the chosen scale. In this work, the proposed model was centralized on the data-derived information without suggesting the incorporation of any explicit physical knowledge. Alvarez *et al.*[16] proposed a method that is much related to the method that we introduce. They have presented an interpretation of a classic optical flow method by Nagel and Enkelmann[17] as a tensor-driven anisotropic diffusion approach. They avoided convergence to irrelevant local minima by embedding their method into a linear scale-space framework. However, they cannot cope with possible modifications of structures' topology.

Non-rigid registration of 3D ultrasound images poses a significant challenge due to the following shortcomings: (i) Low SNR of ultrasound images which are characterized by Rayleigh-governed speckle noise and corrupted by Gaussian-distributed electronic noise. (ii) Motion ambiguities which arise when there is insuffcient representation of spatial information. This holds in regions of image saturation or specular reflection and in homogeneous regions of weak acoustic scatters. (iii) Speckle decorrelation. Since speckle patterns result from the constructive and destructive interference of ultrasonic echoes from numerous subresolvable elements, non-uniform movement of these scatters in the tissue volume can cause temporal decorrelation of the speckle patterns. The algorithm which is presented in this paper is designed to overcome the above shortcomings and they lead to an accurate registration.

The paper is organized as follows. In Sec. 2 we present in detail the multiresolution and multigrid optimization scheme. Section 3 describes the multiscale

framework that the optimization scheme is embedded. Section 4 is dedicated to the experimental results and conclusions are drawn in Sec. 5.

2. Primary Registration Method

2.1. *Formulation of the registration problem*

In this section, we present the primary registration method that has been proposed by Hellier *et al.*[13] In this work, the registration problem is considered as a motion estimation problem. The optical flow hypothesis, introduced by Horn and Schunck,[18] leads then to the minimization of the following cost function:

$$U(\boldsymbol{w}; f) = \sum_{s \in S} [\boldsymbol{\nabla} f(s,t) \cdot \boldsymbol{w}_s + f_t(s,t)]^2 + \alpha \sum_{\langle s,r \rangle \in \mathcal{C}} \|\boldsymbol{w}_s - \boldsymbol{w}_r\|^2 \,, \tag{1}$$

where s is a voxel of the volume, t is the temporal index of the volumes, f is the luminance function, $\boldsymbol{w}$ is the expected 3D displacement field, S is the voxel lattice, $\mathcal{C}$ is the set of neighboring pairs and α controls the balance between the two energy terms. The first term is the first order Taylor-expansion of the luminance conservation equation and it represents the interaction between the field and the data, whereas the second term expresses the smoothness constraint.

Shortcomings of this formulation are well-known:

(a) The optical flow constraint (OFC) is not valid in cases of large displacements because of the linearization.
(b) The OFC might not be valid everywhere, because of the noise, the intensity non-uniformity, and occlusions.
(c) The "real" field probably contains discontinuities that might not be preserved.

To cope with the (b) and (c) limitations, the quadratic cost has been replaced by robust functions. To face the problem (a), a multiresolution and multigrid strategy has been designed.

2.2. *Robust estimators*

Cost function (1) does not make any difference between relevant data and inconsistent data, and it is sensitive to noise. Therefore, robust M-estimators have been introduced in the formulation.[19] An M-estimator is a function ρ that is increasing on $\mathbb{R}^+$, such that (i) $\phi(u) \triangleq \rho(\sqrt{u})$ is strictly concave on $\mathbb{R}^+$ and (ii) $\lim_{x \to \infty} \rho'(x) < \infty$. The main benefit of robust M-estimators is the semi-quadratic formulation that can be deduced from (i):

$$\exists \psi \in C^1([0,M], \mathbb{R}) : \forall u, \rho(u) = \min_{z \in [0,M]} (zu^2 + \psi(z)) \,. \tag{2}$$

Two robust estimators have therefore been introduced: the first one on the data term (ρ_1) and the second one on the regularization term (ρ_2). According to Eq. (2), the minimization of the cost function (1) is equivalent to the minimization of the augmented function, noted $\overset{\star}{U}$:

$$\overset{\star}{U}(\boldsymbol{w}, \delta, \beta; f) = \sum_{s \in S} \{\delta_s (\boldsymbol{\nabla} f(s,t) \cdot \boldsymbol{w}_s + f_t(s,t))^2 + \psi_1(\delta_s)\}$$

$$+ \alpha \sum_{\langle s,r \rangle \in C} \{\beta_{sr}(\|\boldsymbol{w}_s - \boldsymbol{w}_r\|)^2 + \psi_2(\beta_{sr})\}, \tag{3}$$

where δ_s and β_{sr} are auxiliary variables acting as "weights". This cost function has the advantage to be quadratic with respect to $\boldsymbol{w}$. Furthermore, when the adequation of data with the model is not correct, its contribution gets lower as the associated weight δ_s decreases ($\delta_s = \phi'_1([\boldsymbol{\nabla} f(s,t) \cdot \boldsymbol{w}_s + f_t]^2)$), and function ϕ' decreases as well, making this formulation more robust.

2.3. *Multiresolution and multigrid minimization*

In order to cope with large displacements, a classical incremental multiresolution procedure has been developed. A pyramid of volumes $\{f^k\}$ is constructed by successive Gaussian smoothing and subsampling. At the coarsest level, the linearization of the conservation equation can be hopefully used. For the next resolution levels, only an increment $d\boldsymbol{w}^k$ is estimated to refine estimate $\hat{\boldsymbol{w}}^k$, obtained from the previous level (Eq. (4)).

$$\overset{\star}{U}^k(d\boldsymbol{w}^k, \delta^k, \beta^k; f^k, \hat{\boldsymbol{w}}^k) = \sum_{s \in S^k} \{\delta_s^k (\boldsymbol{\nabla} f^k(s + \hat{\boldsymbol{w}}_s^k, t_2) \cdot d\boldsymbol{w}_s^k + f^k(s + \hat{\boldsymbol{w}}_s^k, t_2)$$

$$- f^k(s, t_1))^2 + \psi_1(\delta_s^k)\} + \alpha \sum_{\langle s,r \rangle \in \mathcal{C}^k} \{\beta_{sr}^k(\|(\hat{\boldsymbol{w}}_s^k + d\boldsymbol{w}_s^k)$$

$$- (\hat{\boldsymbol{w}}_r^k + d\boldsymbol{w}_r^k)\|)^2 + \psi_2(\beta_{sr}^k)\}. \tag{4}$$

Furthermore, at each level of resolution, a multigrid minimization based on successive partitions of the initial volume is achieved (see Fig. 1). For each cube of a given grid level ℓ (partition of cubes), a 12-parametric increment field is estimated. The result over the grid level is a rough estimate of the desired solution, and it is used to initialize the next grid level. This hierarchical minimization strategy improves the quality and the convergence rate.

The partition at the coarsest grid level is initialized with a binary segmentation mask of the structure of interest (template). The octree partition which is thus defined is anatomically relevant. When we change the grid level, each cube is divided adaptively. The criterion of subdivision may be either the measure of the way that model fits the data, or a prior knowledge such as the presence of an important anatomical structure where the estimation must be accurate. Consequently, we can distinguish between the regions of interest where the estimation must be precise and the other regions where computation efforts are useless.

Fig. 1.   Example of multiresolution/multigrid minimization. For each resolution level (on the left), a multigrid strategy (on the right) is performed. For clarity reasons, this is a 2D illustration of our 3D algorithm.

3. Embedded Multiscale Framework

The multigrid scheme which has already been described is bound to a good initialization of the flow. To improve the quality of the initial estimates we propose to incorporate the scale of image measurements by exploring the scale-space of the data-derived information. This improvement is expected to come due to the manipulation of data in different scales with reduced speckle noise that is inherent to the echographic data. Specifically, since we deal with the optical flow constraint we experiment with two scale-spaces which are characterized by the luminance conserving principle. These are the linear scale-space[20] and the one which is constructed by the regularized version[21] of the Perona–Malik algorithm.[22] Let f_τ^0 be the luminance of a voxel at the finest spatial resolution which has been diffused at the scale quantization level τ. Then, a linear scale-space is denoted as:

$$f_\tau = f_o * G_\sigma \,, \tag{5}$$

where $*$ denotes convolution, f_o is the original image and G_σ is the Gaussian kernel of standard deviation σ.

If no scale is preferred, the natural way to travel through a linear multiscale can be realized via a sampling which should follow a linear and dimensionless scale parameter $\delta\lambda$ which is related to σ by :

$$\sigma_\tau = e^{\lambda_0 + \tau\delta\lambda} \,, \tag{6}$$

where τ denotes the quantization levels.

The *regularized* Perona–Malik scale-space in its discretized form is denoted as:

$$f_\tau = f_{\tau-1} + \lambda \sum_i c_i (G_\sigma * \Delta_i f) \,, \tag{7}$$

where $i \in \{N, S, E, W, F, B\}$ and N, S, E, W, F, B denote Northern, Southern, Eastern, Western, Forward and Backward neighbor, respectively.

$$c_i = g\|G_\sigma * \Delta_i f\| \,, \tag{8}$$

where c_i is a decreasing function of the image gradient that has been determined at a scale σ to compensate for noise and to assure well-posedness of the diffusion equation.

$\Delta_i f = f_i - f_\star$ where $f_\star$ denotes the central pixel in a 6-neighbor, 3-dimensional mask:

$$g\|\Delta_i f\| = e^{-(\frac{\Delta_i f^2}{k})}\,, \tag{9}$$

$$g\|\Delta_i f\| = \frac{1}{(1 + \frac{\Delta_i f^2}{k})}\,, \tag{10}$$

where k is a contrast parameter and can be interpreted as a threshold, which determines whether a gradient is significant or not.

For non-linear diffusion schemes there is no global scale parameter because they adapt the diffusion locally. However, we may synchronize their scale parameter with the one of linear diffusion. This holds due to the fact that the scalar diffusivity c_i in Eq. (8) is constructed such that $\|c\| \leq 1$. Therefore, an upper bound is derived for the nonlinear schemes which permits us to recall the relation between the evolution parameter and the standard deviation of the Gaussian $\tau_n = (1/2)\sigma_n^2$ for the creation of the *regularized* Perona–Malik scale quantization space.[23]

The construction of any of the above scale-spaces leads to a stack of volumes $\{f_\tau^0\}$ which is the source of the data measurements for every successive quantization scale during a coarse-to-fine parameter estimation. This can be explained by Eq. (11).

$$\overset{\star}{U}_\tau^0(d\boldsymbol{w}^0, \delta^0, \beta^0; f_\tau^0, \hat{\boldsymbol{w}}^0) = \sum_{s \in S^0} \{\delta_s^0(\boldsymbol{\nabla} f_\tau^0(s + \hat{\boldsymbol{w}}_s^0, t_2) \cdot d\boldsymbol{w}_s^0 + f_\tau^0(s + \hat{\boldsymbol{w}}_s^0, t_2)$$

$$- f_\tau^0(s, t_1))^2 + \psi_1(\delta_s^0)\} + \alpha \sum_{\langle s,r \rangle \in \mathcal{C}^0} \{\beta_{sr}^0(\|(\hat{\boldsymbol{w}}_s^0 + d\boldsymbol{w}_s^0)$$

$$- (\hat{\boldsymbol{w}}_r^0 + d\boldsymbol{w}_r^0)\|)^2 + \psi_2(\beta_{sr}^0)\}\,, \tag{11}$$

f_τ^0 denotes the data measurement at the finest pyramid resolution and the τ scale quantization level.

Our goal is the estimation of parameter $\hat{\boldsymbol{w}}^0$ which is refined at each quantization scale by only an increment of $d\boldsymbol{w}_s^0$. Minimization remains in the same multigrid fashion.

4. Experimental Results

4.1. *Artificially deformed US3D volumetric data*

In this work, our efforts were motivated by the application of tissue deformation tracking which can result in brain shift correction. In view of this, we have conducted a number of experiments using an original 3D ultrasound image ($256 \times 256 \times 128$)

of the brain of an 8-month old baby (Fig. 2) and its deformed counterpart (Fig. 3). In the ideal case, the accuracy of our algorithm in registering volumes should be tested in a situation that the actual motion should be known. Due to the diffculty in producing known non-rigid motion fields in biological tissues we have chosen


Fig. 2.   Preoperative 3D ultrasound.

Fig. 3.   Simulated intraoperative (deformed) 3D ultrasound.

to simulate this phenomenon. We have created an artificially deformed volume by using a Thin Plate Spline deformation.[24] Although this approach produces a global smooth deformation, we were very careful in the distribution of the point landmarks over the whole volume to cope with local deformations. The deformed volume and the produced velocity field can be seen in Figs. 3 and 4, respectively.

Fig. 4.   The artificial deformation field.

Fig. 5.   Difference between the original and the deformed volume.

Fig. 6.   Difference between the original and the reconstructed volume.

In our experimental work we strived towards an overall comparison between the primary non-rigid registration model of Sec. 2 and the model with an embedded scale-space framework of Sec. 3. Both qualitative and quantitative evaluation have been conducted. As a qualitative measure we have chosen to use the difference image between the original volume and the reconstructed one. All of the registration models produced difference images without significant differences, implying a visually correct registration (Fig. 6). For the sake of comparison we provide the difference image between the original volume and the deformed one in Fig. 5. The difference image in Fig. 6 has appeared after the application of the algorithm which uses the embedded *regularized* P&M scale-space.

For a quantitative evaluation we have considered the following measures: (i) Mean Squared Error (MSE), where the error is expressed as the difference between the original and the reconstructed volume; (ii) the average angular error between the correct $\vec{v_c}$ and the estimated $\vec{v_e}$ velocity: $\psi = \arccos(\vec{v_c} \cdot \vec{v_e})$ along with (iii) its standard deviation.

Table 1 demonstrates the improvement in velocity estimation which has been achieved for the all the three above measures in the case of the embedded scale-space framework for both the linear and the regularized P&M case. The latter one has a slightly better behavior than the linear one.

Our basic argumentation for the advantageous use of a multiscale framework was that it can lead to an improvement in the quality of the initial estimates at the multigrid optimization scheme which subsequently will improve the quality of the final estimates. For the sake of clarity, Fig. 7 shows in terms of MSE the improvement that occurs during successive multigrid levels at the finest spatial

Table 1.   Quantitative comparison measures.

| Scale-space Framework	MSE	Mean Angular Error	Std Deviation
Without multiscale	10.2772	14.112656°	24.254787°
Embedded linear	9.73472	13.878700°	23.987515°
Embedded *regularized* P&M	9.6945	13.791579°	23.959972°

Fig. 7.   MSE improvement with respect to multigrid levels at the finest spatial resolution (artificially deformed volumetric data).

resolution for all the three examined cases. We may observe that in the case of the absence of a multiscale framework we get an initial estimate with an MSE that equals to 15.1756 while in the case of linear scale-space we get an initial estimate with an MSE that equals to 11.9146 and in the case of *regularized* P&M scale-space we get an initial estimate with an MSE that equals to 12.1082. The higher quality of the initial estimates was preserved till the final stage at the multigrid optimization scheme.

4.2. *US3D volumetric data with real deformations*

To validate the proposed method in the case of real deformations we conducted experiments with US3D volumetric datasets that have been produced out of a pig

brain at a postlethal status. Deformation has been invoked by inflating a balloon catheter that caused a volume increase by 0.25 cm^3. An instance of the acquired volumetric data ($100 \times 100 \times 80$) can be seen in Figs. 8 and 9. Similar to Sec. 4.1, we examine the quality of the registration, both qualitatively and quantitatively.


Fig. 8.   Original 3D ultrasound data.

Fig. 9.   Deformed 3D ultrasound data.

Qualitatively, a visual inspection shows that our registration is successful (Fig. 11). The real deformation field as it has been reconstructed by the proposed method is depicted in Fig. 10. Furthermore, the produced differences between the original volume and the reconstructed one are essentially diminished as it can be seen in


Fig. 10.   The deformation field.

Fig. 11.   The reconstructed 3D ultrasound data.

Fig. 13, compared to the differences produced between the original volume and its deformed counterpart (Fig. 12). For the sake of clarity, the values of differences have been augmented by 64 in both Figs. 12 and 13. In the case of a quantitative comparison, we consider the MSE measure (the error is expressed as the difference


Fig. 12.   Difference between the original and the deformed volume.

Fig. 13.   Difference between the original and the reconstructed volume.

Fig. 14. MSE improvement with respect to multigrid levels at the finest spatial resolution (volumetric data with real deformations).

between the original and the reconstructed volume) as the indicator of our registration validity. Furthermore, we demonstrate the improvement of the velocity estimation in terms of MSE during the complete multigrid level set at the spatial resolution for all the three examined cases. In agreement to our experiments with artificial deformations in Sec. 4.1 the initial higher quality of the estimates was preserved till the final stage of the multigrid optimization scheme (Fig. 14).

5. Conclusions and Perspectives

In this paper, we propose a *methodology* which embeds a multiscale framework in a multiresolution and multigrid optimization scheme that can lead to a successful non-rigid registration of 3D ultrasound images. It grasps its power from three fundamental features which operate as the remedy in the basic shortcomings of ultrasound images. Its multigrid nature responds to motion ambiguities in the case of insuffcient representation of spatial information, its estimate smoothness functional term can fight the speckle decorrelation which characterizes ultrasound while low SNR can be less disastrous for the estimates in the case of embedding a multiscale framework.

The embedded multiscale framework has demonstrated a superior performance compared to the primary registration method (Sec. 2) for both an isotropic and

an non-isotropic scale-space. Between the two scale-spaces, we have expected that the *regularized* P&M framework will be by far more robust and may even perform better. Our current experiments have shown that the *regularized* P&M performs slightly better but more experimental studies should be conducted to show agreement or not to the performance assumptions (on precision and robustness).

Acknowledgments

References

1. R. Bucholtz, D. Yeh, J. Trobaugh, L. McDurmont, C. Sturm, C. Baumann, and J. Henderson, "The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device," in *CVRMed — MRCAS*, pp. 459–466 (1997).
2. R. Comeau, A. Fenster, and T. Peters, "Intraoperative US in interactive image-guided neurosurgery," *Radiographics* **18**, 1019–1027 (1998).
3. H. Erbe, A. Kriete, A. Jodicke, W. Deinsberger, and D. Boker, "3D-ultrasonography and image matching for detection of brain shift during intracranial surgery," in H. U. Lemke *et al.*, ed., *Computer Assisted Radiology* (*CAR*), pp. 225–230 (1996).
4. A. Morsy and O. VonRamm, "3D ultrasound tissue motion tracking using correlation search," *Ultrasonic Imaging* **20**, 151–159 (1998).
5. M. Strintzis and I. Kokkinidis, "Maximum likelihood motion estimation in ultrasound image sequences," *IEEE Signal Processing Letters* **4**(6), 156–157 (1997).
6. F. Yeung, S. Levinson, and K. Parker, "Multi-level and motion-model ultrasonic speckle tracking algorithms," *Ultrasound in Medicine and Biology* **24**(3), 427–441 (1998).
7. F. Yeung, S. Levinson, D. Fu, and K. Parker, "Feature-adaptive motion tracking of ultrasound image sequences using a deformable mesh," *IEEE Transactions on Medical Imaging* **17**(6), 945–956 (1998).
8. X. Pennec, P. Cachier, and N. Ayache, "Understanding the demon's algorithm: 3D non-rigid registration by gradient descent," in *Medical Image Computing and Computer Assisted Intervention*, *Lecture Notes in Computer Science* (Springer-Verlag, September 1999), pp. 597–605.
9. X. Pennec, P. Cachier, and N. Ayache, "Tracking brain deformations in time-sequences of 3D US images," *Pattern Recognition Letters — Special Issue on Ultrasonic Image Processing and Analysis* **24**(4–5), 801–813 (2003).
10. T. Arbel, X. Morandi, R. Comeau, and D. L. Collins, "Automatic non-linear MRI-ultrasound registration for the correction of intraoperative brain deformations," in W. J. Niessen and M. A. Viergever, ed., *Medical Image Computing and Computer Assisted Intervention*, Volume 2208 of *Lecture Notes in Computer Science* (Springer-Verlag, Utrecht, The Netherlands, 2001), pp. 913–921.

11. A. P. King, P. G. Batchelor, G. P. Penney, J. M. Blackall, D. L. G. Hill, and D. J. Hawkes, "Estimating sparse deformation fields using multiscale bayesian priors and 3D utrasound," in M. F. Insana and R. M. Leahy, ed., *Information Processing in Medical Imaging*, Number 2082 in Lecture Notes in Computer Science (Springer-Verlag, Davis, CA, 2001), pp. 155–161.

12. E. Mémin and P. Pérez, "Dense estimation and object-based segmentation of the optical flow with robust techniques," *IEEE Transactions on Image Processing* **7**(5), 703–719 (1998).

13. P. Hellier, C. Barillot, E. Mémin, and P. Pérez, "Hierarchical estimation of a dense deformation field for 3D robust registration," *IEEE Transactions on Medical Imaging* **20**(5), 388–402 (2001).

14. J. Weber and J. Malik, "Robust computation of optical flow in a multi-scale differential framework," *International Journal of Computer Vision* **14**, 67–81 (1995).

15. W. J. Niessen, J. S. Duncan, M. Nielsen, L. M. J. Florack, B. M. Romeny and M. A. Viergever, "A multiscale approach to image sequence analysis," *Computer Vision and Image Understanding* **65**(2), 259–268 (1997).

16. L. Alvarez, J. Weickert, and J. Sanchez, "A scale-space approach to nonlocal optical flow calculations," in *Scale-Space '99*, pp. 235–246 (1999).

17. H. H. Nagel and W. Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**, 565–593 (1986).

18. B. Horn and B. Schunck, "Determining optical flow," *Artificial Intelligence* **17**, 185–203 (1981).

19. M. Black and A. Rangarajan, "On the unification of line processes, outlier rejection and robust statistics with applications in early vision," *International Journal of Computer Vision* **19**(1), 57–91 (1996).

20. A. P. Witkin, "Scale-space filtering," in *International Joint Conference on Artificial intelligence* (Karlsruhe, W. Germany, 1983), pp. 1019–1023.

21. F. Catté, P. L. Lions, J. M. Morel, and T. Coll, "Image selective smoothing and edge detection by nonlinear diffusion," *SIAM Journal on Numerical Analysis* **29**, 182–193 (1992).

22. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **12**(7), 629–639 (1990).

23. W. J. Niessen, K. Vincken, J. Weickert, and M. A. Viergever, "Nonlinear multiscale representations for image segmentation," *Computer Vision and Image Understanding* **66**(2), 233–245 (1997).

24. F. L. Bookstein, "Principal warps: Thin-plate splines and the decomposition of defomations," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(6), 567–585 (1989).

Ioannis Pratikakis received his Diploma degree in Electrical Engineering from the Demokritus University of Thrace, Xanthi, Greece, as well as his PhD degree in Applied Sciences from the Vrije Universiteit Brussel, Brussels, Belgium, in 1992 and 1998 respectively. From March 1999 to March 2000, he was at the IRISA/VISTA group, Rennes, France as a postdoctoral fellow. He is currently working as a Research Associate at the Institute of Informatics and Telecommunications of the National Center for Scientific Research "Demokritos", Athens, Greece. His research interests include mathematical morphology, scale-space theory, image registration, color image processing and 3D reconstruction with applications in the field of medicine and robotics. He has been a member of the Hellenic Technical Chamber as well as the IEEE.

Christian Barillot is a senior researcher from the CNRS and he has more than 15 years of experience in the field of medical imaging (3D visualization, data fusion, image processing and registration). He received his PhD degree from the University of Rennes I on "Information Processing" in 1984. Thereafter, he was a Research Associate at the University of Rennes I. In 1986, he was recruited by the CNRS (National Center of Scientific Research) as a permanent researcher.

In 1987, 1988 and partially in 1991, he was a Research Associate at the Mayo Clinic, Rochester, MN, in the Biomedical Imaging Resources. Between 1988 and 1996 he worked in the INSERM Unit at the University of Rennes. In 1996, he joined the VISTA Project at IRISA. He is the (co)author of over 100 refereed scientific articles (abstracts excluded) on medical image processing. In addition, he was the collaborator and the principal investigator of over ten national and international grants. He was the co-chairman of the 14th edition of the IPMI (Information Processing in Medical Imaging) International Conference in 1995 and he will chair the next European edition of the International Conference on MICCAI in 2004 (Medical Image Computing and Computer Assisted Intervention).

Pierre Hellier graduated from the SUPAERO (an engineering school in space and aeronautics) with a postgraduate certificate in Applied Mathematics. He has obtained a PhD degree from INRIA in VISTA Project in December 2000. After a postdoc at the Utrecht University, he joined INRIA as a fulltime researcher in September 2001. His research interests include medical imaging, rigid and non-rigid registration, segmentation, calibration and acquisition of 3D ultrasound for neurosurgery.

Etienne Memin received his PhD degree in computer science from the University of Rennes I, France, in 1993. From 1994 to 1999, he was an Assistant Professor in computer science at the University of South Brittany (UBS Vannes).

He is currently an Assistant Professor in computer science at the University of Rennes I. His major research interests include motion analysis (motion estimation, motion segmentation and tracking), energy-based modeling and parallel processing. Since a couple of years ago, he has been particularly interested in motion analysis issues for image sequences depicting fluid phenomenon.