ITERATIVE MULTI-SCALE REGISTRATION WITHOUT LANDMARKS

Philippe Thévenaz, Urs E. Ruttimann and Michael Unser
National Institutes of Health
Bethesda MD 20892–5766, USA

ABSTRACT

We present an automatic sub-pixel registration algorithm that minimizes the mean square difference of intensities between a reference and a test data set (volumes or images). It uses spline processing, is based on a coarse-to-fine pyramid strategy, and performs minimization according to a variation of the iterative Marquardt-Levenberg scheme. The geometric deformation model is a general affine transformation that one may optionally restrict to a rigid-body (isometric scale, rotation and translation), procrustean (rotation and translation) or translational case; it also includes an optional parameter for the linear adaptation of intensity. We present several PET and fMRI experiments and show that this algorithm provides excellent results. We conclude that the multi-resolution refinement strategy is faster and more robust than a comparable single-scale one.

1. INTRODUCTION

Image registration plays a major role in at least two types of application: data fusion, where noise reduction is often desired, and data comparison, for the detection of significant differences.

In the first case, one takes advantage of the availability of multiple instances of supposedly identical data. The registration of these instances with one another allows the extraction of common features, for example by averaging, or by more refined processes [1]. One can find successful examples of this approach in the correlation-averaging of virus particles in high-resolution electron microscopy [2]. It has also been applied in the consolidation of sparse data, for example when mapping range images of the sea floor [3].

In the second case, a new problem appears, due to the very fact that the registration process tries to align data that are possibly dissimilar. This last consideration sometimes leads to robust registration methods using an internal criterion that is not sensitive to outliers [4, 5, 6]. After registration, the task usually proceeds by the detection of dissimilar regions, given a statistically significant level of confidence in the decision with respect to the type I and type II errors [7].

Be it fully automated [8] or not [9], landmark registration is often proposed as a versatile solution. It allows for a geometric transformation whose properties can be as general as desired [10], or in the contrary be tuned to the problem by additional constraints [11, 12]. However, the solution can be no better than the initial selection of landmarks allows, which is the crucial component of this approach; moreover, it fails when sub-pixel accuracy is desired, since the landmarks are usually selected directly on the discretization grid.

Registration without landmark is the preferred solution whenever applicable, since it often takes into account the whole of the available data by implementing some kind of correlation process. The resulting benefit is that the robustness of this class of techniques is much higher than that of landmark-based registration [4]. One may consider any kind of geometric transformation, be it local and computationally expensive, like a deformation field [13], or inexpensive, but restricted to a global translation [14].

In this paper, we address the problem of registration for data comparison purposes, without landmarks, and with a global affine transformation paradigm. This transformation is more general than the one used in [15]. We present a fully automatic iterative multi-scale registration with sub-pixel accuracy, extending the work in [16] by considering a full tri-dimensional (3-D) approach. In the terminology of [17], we use the raw intensity as our feature space, thus effectively exploiting all the information in the image; our search space is the continuous set of non-singular 3-D affine transforms, which are the most general linear transformations available, and where by continuity we imply sub-pixel accuracy; our dissimilarity measure is Euclidean, which is maximum-likelihood assuming additive white Gaussian noise; finally, our search strategy is multi-scale, for fast convergence, and iterative, based on a variation of the Marquardt-Levenberg (ML) algorithm for non-linear least-square optimizations [18].

2. SYSTEM

2.1. Criterion

Any automatic registration method requires the optimization of an objective criterion, whose role is to measure the similarity of the test data with respect to the reference. As criterion, we select ε^2 the average square difference. Letting $f_{\mathrm{R}}(\mathbf{x})$ and $f_{\mathrm{T}}(\mathbf{x})$ be respectively the reference and test data, our criterion reads

$$\varepsilon^2 = \frac{1}{|V|} \int_{\mathbf{x} \in V} \left(f_{\mathrm{R}}(\mathbf{x}) - Q_{\mathbf{p}}\{f_{\mathrm{T}}(\mathbf{x})\} \right)^2 \mathrm{d}\mathbf{x} \tag{1}$$

where $Q_{\mathbf{p}}\{f\}$ is the transformation of interest described by $\mathbf{p}$, and where V represents the volume of interest. This criterion offers the advantage that it is well understood and lends itself well to minimization with respect to $\mathbf{p}$. Its drawback is that, in the presence of severe outliers, its minimum may become less pronounced, and may not even be located where one would expect it to be.

In our case (medical imaging), outliers do exist, for example when we compare two functional magnetic resonance images (fMRI) of a brain in two different states of activation. However, we do not expect these outliers to be dominant; in fact, the very need for registration arises from the fact that the differences between the two brain instances are extremely faint, and cannot be detected without their careful alignment.

2.2. Transformation

As transformation of interest $Q_{\mathbf{p}}\{f\}$, we consider the general affine transformation described by a 3×3 matrix $\mathbf{A}$ and by a translation vector $\mathbf{b}$

$$Q_{\mathbf{p}}\{f(\mathbf{x})\} = Q_{\mathbf{A},\mathbf{b}}\{f(\mathbf{x})\} = f(\mathbf{A}\,\mathbf{x} + \mathbf{b}) \tag{2}$$

This transformation includes translation, rotation around any center, skewing, shearing and scaling. We have also developed more constrained versions (e.g., rigid body, procrustean, translation only), and, as in [4], we have added to all of them a possibility for gray-scale adaptation; but we do not report on any of these extra versions here.

2.3. Multi-scale processing

We will see below that our approach is iterative. We presently take advantage of this fact by introducing a change of resolution between some of the iterations. Using a coarse-to-fine strategy enables us to perform a first parameter estimation at the coarsest level of a resolution pyramid; after this first estimation is deemed satisfactory, we switch to a finer level, using the previous parameters as initial conditions for carrying the optimization on the new level. At each new level the convergence is fast, since these propagated initial conditions are close to the actual optimum.

In fact, for practical purposes, the quasi-Newtonian optimization method that we use can converge in just one step if the previous estimate is sufficiently good, which is often the case for all but the coarsest level where most of the iterations are spent, but where fortunately

the data set is small. This saving reflects itself in the computation time, which is greatly reduced.

Another very important advantage of using a resolution pyramid is that its embedded smoothing properties tend to regularize the problem by letting the surface $\varepsilon^2(\mathbf{p})$ be also smoother at coarser scales: the image size reduction effectively removes most of the noise. With most images, it is very likely that the location of a near-absolute optimum in the coarsest level will not be missed, whereas it would possibly be, if only the finest level would be looked at. This is important since, contrary to [19], we do not make use of an optimization procedure that would be able to escape from local minima. The pyramid is computing according to [20, 21] with cubic splines.

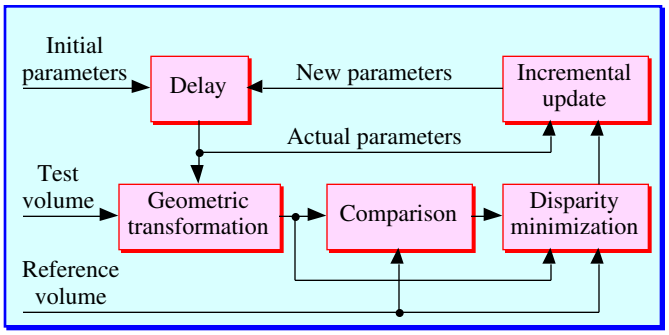

Figure 1: Outline of the registration algorithm

2.4. Optimization

Our choice for solving $\partial\varepsilon^2(\mathbf{p})/\partial\mathbf{p}=0$ is ML, an iterative gradient-based algorithm for non-linear least-squares optimization problems [18]. The system is depicted in Figure 1, where we see that at each step $f_T(\mathbf{x})$ has to undergo $Q_\mathbf{p}\{f\}$ before it is compared to $f_R(\mathbf{x})$. The reslicing needed in the geometric transformation is performed by resampling a cubic spline model fitted to the volume. The optimization requires the explicit knowledge of the partial derivatives $\partial f_T/\partial\mathbf{p}$, which are computed by using a cubic spline fit again, along with the techniques described in [20, 21]. In section 3 we will propose an approach to further accelerate the standard ML algorithm when the transformation satisfies certain commutativity and distributivity rules.

2.5. Convergence

We consider three concurrent criteria for deciding when to stop the iteration process. The first one is self-evident: we stop when a near-perfect match is met, that is, when $\varepsilon^2 \le \mathrm{T}$, where T is some prescribed threshold. The second has to do with the observed relative gain $\Delta\varepsilon^2/\varepsilon^2$ at each successful iteration step: we declare that the convergence is reached whenever this gain is below another user-selected threshold. The third one concerns itself with $|\Delta p_i/p_i|$ the relative change of parameter values at every iteration step: when at least one of the parameters is alive, we keep on going; but when all are motionless we stop. If any parameter p_i happens to be zero, or near zero, we substitute $|\Delta p_i/p_i|$ by a similar, although possibly less effective criterion extracted from some of the parameters of the ML algorithm, guaranteed not to suffer from normalization problems, and guaranteed to converge towards zero when the number of iterations grows high enough.

3. ACCELERATED MARQUARDT-LEVENBERG

3.1. Requirements

Let us decompose the operator $Q_\mathbf{p}\{f\}$ into a translation and a matrix operation given respectively by

$$Q_\mathbf{b}\{f(\mathbf{x})\} = f(\mathbf{x}+\mathbf{b}) \qquad A_\mathbf{A}\{f(\mathbf{x})\} = f(\mathbf{A}\mathbf{x}) \tag{3}$$

The composition rules for these operators are

$$\begin{aligned}
T_\mathbf{b}\{T_\mathbf{a}\{f(\mathbf{x})\}\} &= f(\mathbf{x}+\mathbf{a}+\mathbf{b}) = T_{\mathbf{a}+\mathbf{b}}\{f(\mathbf{x})\}\\
A_\mathbf{B}\{A_\mathbf{A}\{f(\mathbf{x})\}\} &= f(\mathbf{B}\mathbf{A}\mathbf{x}) = A_{\mathbf{B}\mathbf{A}}\{f(\mathbf{x})\}\\
T_\mathbf{b}\{A_\mathbf{A}\{f(\mathbf{x})\}\} &= f(\mathbf{A}\mathbf{x}+\mathbf{b}) = A_\mathbf{A}\{T_{\mathbf{A}^{-1}\mathbf{b}}\{f(\mathbf{x})\}\}\\
A_\mathbf{A}\{T_\mathbf{b}\{f(\mathbf{x})\}\} &= f(\mathbf{A}(\mathbf{x}+\mathbf{b})) = T_{\mathbf{A}\mathbf{b}}\{A_\mathbf{A}\{f(\mathbf{x})\}\}
\end{aligned} \tag{4}$$

These operators are distributive with respect to the addition

$$\begin{aligned}
T_\mathbf{b}\{f(\mathbf{x})+g(\mathbf{x})\} &= T_\mathbf{b}\{f(\mathbf{x})\}+T_\mathbf{b}\{g(\mathbf{x})\}\\
A_\mathbf{A}\{f(\mathbf{x})+g(\mathbf{x})\} &= A_\mathbf{A}\{f(\mathbf{x})\}+A_\mathbf{A}\{g(\mathbf{x})\}
\end{aligned} \tag{5}$$

Looking at the squared norm of the signal, we have

$$\|T_\mathbf{b}\{f(\mathbf{x})\}\|^2 = \|f(\mathbf{x})\|^2 \qquad \|A_\mathbf{A}\{f(\mathbf{x})\}\|^2 = \|f(\mathbf{x})\|^2/|\mathbf{A}| \tag{6}$$

3.2. Criterion revisited

Using the previous expressions, we can now rewrite a criterion exactly equivalent to (1) as

$$\varepsilon^2 = \left\| A_{\mathbf{I}+\Delta\mathbf{A}}\{T_{\Delta\mathbf{b}}\{f_T(\mathbf{x})\}\} - A_{\mathbf{A}^{-1}}\{T_{-\mathbf{b}}\{f_R(\mathbf{x})\}\}\right\|^2 \tag{7}$$

$$\varepsilon^2 = \frac{1}{\sqrt{|\mathbf{I}+\Delta\mathbf{A}|}}\cdot\left\| f_T(\mathbf{x}) - A_{(\mathbf{I}+\Delta\mathbf{A})^{-1}\mathbf{A}^{-1}}\{T_{-\mathbf{b}-\mathbf{A}(\mathbf{I}+\Delta\mathbf{A})\Delta\mathbf{b}}\{f_R(\mathbf{x})\}\}\right\|^2 \tag{8}$$

$$\varepsilon^2 = \sqrt{|\mathbf{A}|}\cdot\left\| T_{\mathbf{b}+\mathbf{A}(\mathbf{I}+\Delta\mathbf{A})\Delta\mathbf{b}}\{A_{\mathbf{A}(\mathbf{I}+\Delta\mathbf{A})}\{f_T(\mathbf{x})\}\} - f_R(\mathbf{x})\right\|^2 \tag{9}$$

Minimizing (7) with respect to $\Delta\mathbf{p}=(\Delta\mathbf{A},\Delta\mathbf{b})$ is equivalent to the more traditional minimization of (9) with respect to $\mathbf{p}$ when written with $\Delta\mathbf{p}=\mathbf{0}$. The advantage of the minimization over $\Delta\mathbf{p}$ instead of $\mathbf{p}$ is that the Hessian matrix required in ML has now to be computed once only, at the parameter value $\Delta\mathbf{p}=\mathbf{0}$; the same is true for the derivatives $(\partial f_T/\partial\Delta\mathbf{p})|_{\Delta\mathbf{p}=\mathbf{0}}$.

It also simplifies the dependencies between the Δp_i parameters, the only cost being a slightly more complicated parameter update procedure, that is, $\mathbf{p}\to\mathbf{p}+\Delta\mathbf{p}$ has to be replaced by the set of rules given in (8). In essence, we now update the transformation that is applied to f_R instead of f_T.

3.3. Level change

Since our approach is pyramidal, at each level change we not only have to compute new derivatives and Hessian, but we also need to update the parameters between levels. Consider an initial geometric transformation $T_\mathbf{b}\{A_\mathbf{A}\{f(\mathbf{x})\}\}$ given by

$$\begin{cases} y_1 = b_1 + a_{11}x_1 + a_{12}x_2 + a_{13}x_3\\ y_2 = b_2 + a_{21}x_1 + a_{22}x_2 + a_{23}x_3\\ y_3 = b_3 + a_{31}x_1 + a_{32}x_2 + a_{33}x_3 \end{cases} \tag{10}$$

Apply some scale factors $(\lambda_1,\lambda_2,\lambda_3)$ and compute a new, equivalent transformation $T_{\mathbf{b}'}\{A_{\mathbf{A}'}\{s(\mathbf{x}')\}\}$ as in

$$\begin{cases} y_1' = b_1' + a_{11}'x_1' + a_{12}'x_2' + a_{13}'x_3' \\ y_2' = b_2' + a_{21}'x_1' + a_{22}'x_2' + a_{23}'x_3' \\ y_3' = b_3' + a_{31}'x_1' + a_{32}'x_2' + a_{33}'x_3' \end{cases} \quad \begin{cases} x_1' = \lambda_1 x_1 \\ x_2' = \lambda_2 x_2 \\ x_3' = \lambda_3 x_3 \end{cases} \quad \begin{cases} y_1' = \lambda_1 y_1 \\ y_2' = \lambda_2 y_2 \\ y_3' = \lambda_3 y_3 \end{cases} \tag{11}$$

These four systems of three equations are satisfied by

$$\begin{cases} b_1' = \lambda_1 a_1 \\ b_2' = \lambda_2 a_1 \\ b_3' = \lambda_3 a_1 \end{cases} \quad \begin{cases} a_{11}' = a_{11} & a_{12}' = \frac{\lambda_1}{\lambda_2}a_{12} & a_{13}' = \frac{\lambda_1}{\lambda_3}a_{13} \\ a_{21}' = \frac{\lambda_2}{\lambda_1}a_{21} & a_{22}' = a_{22} & a_{23}' = \frac{\lambda_2}{\lambda_3}a_{23} \\ a_{31}' = \frac{\lambda_3}{\lambda_1}a_{31} & a_{32}' = \frac{\lambda_3}{\lambda_2}a_{32} & a_{33}' = a_{33} \end{cases} \tag{12}$$

This last equation allows us to carry a description of the geometric transformation from one level of the pyramid to the next. It also needs to be applied at the change of level, when the origin of the coordinate system is displaced.

4. EXPERIMENTS

4.1. Ideal case

We begin our experiment series with a noise-free case, and we further simplify the problem by considering an image instead of a volume. First, we scale by λ, rotate by $\theta = 5°$ and translate by $(dx, dy) = (5,5)$ the Lena image using cubic spline resampling. We then use this image as our reference and try to align the original towards it. Since our registration procedure also uses cubic spline resampling, an exact match is a priori not impossible. Table 1 gives the considered parameters (including N the number of levels in the pyramid and t the relative computation time), their estimation by our registration algorithm, and the residual error ε^2 expressed in dB as a SNR between the aligned $f_T(\mathbf{x})$ and $f_R(\mathbf{x})$.

	i	*ii*	*iii*	*iv*	*v*	*vi*
N	4	4	4	3	2	1
t	0.30	0.23	0.22	0.22	0.31	1.00
dx'	5.0010	5.0002	4.9983	4.9983	4.9984	4.9999
dy'	5.0057	5.0002	5.0018	5.0018	5.0016	4.9992
θ'	5.0070	5.0003	5.0015	5.0015	5.0014	5.0035
λ'	0.7999	1.2500	1.0000	1.0000	1.0000	0.9999
λ	0.8	1.25	1	1	1	1
ε^2	42.05	43.29	43.26	43.26	43.28	43.03

Table 1: Ideal registration experiment

In this table, θ' has been estimated from the general affine transformation, and λ' from its determinant (λ is the correct value). We see that the fit is almost perfect, the relative error on the parameter estimation being most of the time lower than 1‰. We even had to resort giving many digit figures for showing the accuracy of the estimates.

4.2. Noisy case

We now add some white Gaussian noise to the previous test and reference images and we observe its effect on the registration. The amount of added noise is the same for $f_T(\mathbf{x})$ and $f_R(\mathbf{x})$, is independent for each of both images, and is reported in dB. Table 2 shows the results of these experiments, with $N = 4$, $(dx, dy) = (5,5)$, $\theta = 5°$, and $\lambda = 1$.

	i	*ii*	*iii*	*iv*	*v*	*vi*
SNR	25	20	15	10	5	0
dx'	4.9981	5.0008	4.9918	4.9883	4.9629	4.9058
dy'	5.0022	4.9988	5.0170	5.0314	5.0533	5.1996
θ'	5.0005	5.0024	5.0027	5.0112	5.0447	5.1102
λ'	0.9999	1.0000	0.9999	1.0000	0.9997	0.9995
ε^2	23.86	18.46	13.34	8.35	4.05	0.77

Table 2: Noisy registration experiment

Even when the signal is in a one-to-one ratio with noise (0 dB case), the parameter estimation is still very good (the relative error is lower than 5%). In this experiment, like in the previous one, we used the identity transformation as initial condition for the registration algorithm.

4.3. Controlled 3-D fMRI phantom

For this experiment, we used a 3-D phantom acquired through the functional Magnetic Resonance Imaging (fMRI) system described in [22]. We performed two separate sets of experiments. In the first one, we reconstructed the phantom from the Fourier domain with an added linear phase (synthetic translation); in the second one, we physically rotated the phantom by a known angle. Tables 3 shows the respective results of these two experiments, where Δx (in voxel units) stands for the synthetic displacement, $\Delta x'$ for its estimation through our registration procedure, and where θ (in degrees) stands for the measured physical rotation and θ' for its estimation.

	i	*ii*	*iii*	*iv*	*v*	*vi*
Δx	0.800	1.333	2.133	2.667	2.933	
$\Delta x'$	0.820	1.372	2.186	2.735	3.003	
θ	0.23	0.47	0.70	0.93	1.17	1.40
θ'	0.23	0.47	0.66	0.87	1.09	1.37

Table 3: Two fMRI experiments

We see that the agreement between the synthetic and the estimated displacements is fairly good; the same can be said of the measured and the estimated angle of rotation. Unfortunately, the precision of the mechanical device used for rotating the phantom is unknown; but the true difference $\theta - \theta'$ can be considered to lie within a tenth of a degree. We should add here that the limited range we explored was constrained by mechanical limitations. Given the identity transformation as initial guess, other experiments have shown that our registration algorithm still converges for angles as big as 30°.

4.4. Uncontrolled 3-D PET brain scans

In the previous experiments we showed cases where the deformation model (affine transformation) was consistent with the data. In the next example, this assumption will be no more valid: consider brain volumes acquired through a Positron Emission Tomography (PET) imaging device, and take several patients into account. The goal is now to register these brains with one another. Since they belong to different subjects, they exhibit not only size differences but also shape differences, which introduces outliers in the statistics of the difference volume. Figure 2 shows a set of unregistered slices of several patients, cut at the same nominal position in the original volumes. The top left is the reference slice, and the bottom right is the average of the selected slice for all 32 patients. Figure 3 shows the same configuration after registration in 3-D.

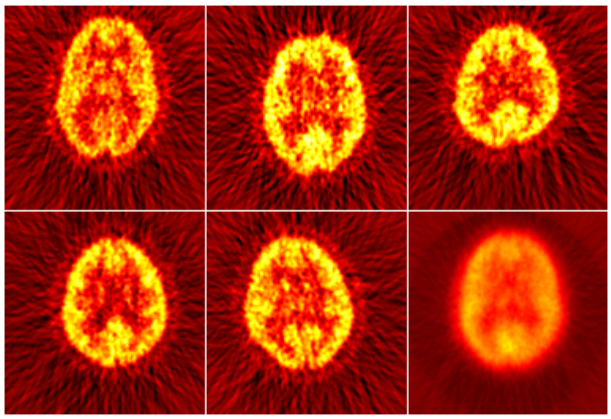

Figure 2: Unregistered slices and their average

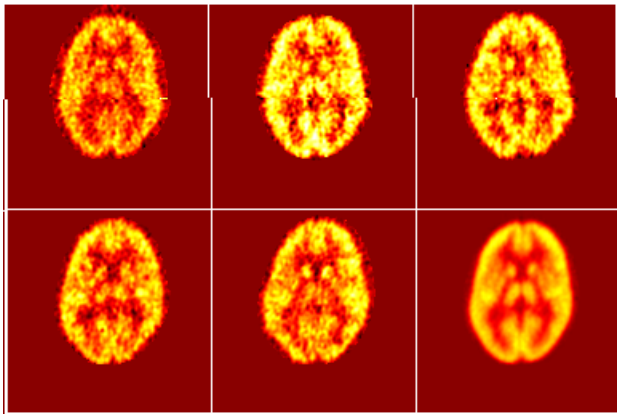

Figure 3: Registered slices and their average

In our registration algorithm, the specification of V, from equation (1), allows one to mask out irrelevant features. Since PET images are typically noisy, we took advantage of this masking and ignored most of the background contribution, which explains the disappearance of streaks in Figure 3.

Regarding the success of the registration, it is clear that most of the brain features are better resolved in the registered average volume than in the unregistered one (lower-right part of Figures 4 and 3 respectively). It tends to show that our least-squares registration criterion is robust enough to cope with outliers, i.e., deformations that cannot be modeled by an affine transformation. Another kind of outliers, not striking in the displayed figures but nevertheless present, has to do with the fact that the intensity contrast of the volumes differs from one patient to the next. Our registration algorithm can take this effect into account and determine a best linear remapping of the intensity that, when applied concurrently with the geometric affine transformation, minimizes the criterion shown in equation (1).

5. CONCLUSION

We have described a fully automatic registration algorithm that uses raw intensity as feature space and considers a Euclidean least-squares criterion for the simultaneous determination of a general 3-D affine transformation and of a linear change of intensity contrast. The search strategy takes advantage of a resolution pyramid and is based on a variation of the Marquardt-Levenberg algorithm for non-linear least-squares optimization. We perform all the processing steps with cubic spline models.

We have implemented this algorithm and we have presented several experiments involving 2-D images and 3-D volumes as well. We were able to show a good performance of our algorithm through a whole range of cases (synthetic problem, synthetic problem with added noise, controlled 3-D experiment with real data, real data with geometric outliers and intensity outliers).

An attractive feature of this registration algorithm is that it can easily incorporate a priori knowledge by the way of a (possibly multi-valued) mask. Another attractive feature is that the user can constrain the general affine transformation to be a rigid-body transformation, or even a procrustean transformation. Finally, the cost, in terms of computation time, is minimized through the use of a resolution pyramid.

REFERENCES

[1] M. Irani and S. Peleg, "Improving Resolution by Image Registration," *Computer Vision, Graphics, and Image Processing,* vol. 53, no. 3, pp. 231–239, 1991.

[2] J. Frank, A. Verschoor and M. Boubik, "Computer Averaging of Electron Micrographs of 40S Ribosomal Subunits," *Science*, vol. 214, pp. 1353–1355, 1981.

[3] B. Kamgar-Parsi, J.L. Jones and A. Rosenfeld, "Registration of Multiple Overlapping Range Images: Scenes without Distinctive Features," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 13, no. 9, pp. 857–871, 1991.

[4] M. Herbin, A. Venot, J.-Y. Devaux, E. Walter, J.-F. Lebruchec, L. Dubertret and J.-C. Roucayrol, "Automated Registration of Dissimilar Images: Application to Medical Imagery," *Computer Vision, Graphics, and Image Processing*, vol. 47, no. 1, pp. 77–88, 1989.

[5] J.Y. Chiang and B.J. Sullivan, "Coincident Bit Counting—A New Criterion for Image Registration," *IEEE Trans. Medical Imaging*, vol. 12, no. 1, pp. 30–38, 1993.

[6] A. Venot, L. Pronzato and E. Walter, "Comments about the Coincident Bit Counting (CBC) Criterion for Image Registration," *IEEE Trans. Medical Imaging*, vol. 13, no. 3, pp. 565–566, 1994.

[7] U.E. Ruttimann, M. Unser and D. Rio, "Statistical Analysis of Image Differences by Wavelet Decompositions," in Proc. *16th Annual Int. Conf. of the IEEE Engineering in Medicine and Biology Society*, Baltimore, Maryland, U.S.A., November 3-6, 1994, pp. 28a–29a.

[8] Q. Zheng and R. Chellappa, "Motion Detection in Image Sequences Acquired from a Moving Platform," in Proc. *Int. Conf. Acoustics, Speech, and Signal Processing*, Minneapolis, Minnesota, U.S.A., April 27-30, 1993, pp. V-201–V-204.

[9] G.Q. Maguire, Jr., M.E. Noz, H. Rusinek, J. Jaeger, E.L. Kramer, J.J. Sanger and G. Smith, "Graphics Applied to Medical Image Registration," *IEEE Computer Graphics and Applications*, pp. 20–28, 1991.

[10] J. Flusser, "An Adaptative Method for Image Registration," *Pattern Recognition*, vol. 25, no. 1, pp. 45–54, 1992.

[11] B.C.S. Tom, S.N. Efstratiadis and A.K. Katsaggelos, "Motion Estimation of Skeletonized Angiographic Images Using Elastic Registration," *IEEE Trans. Medical Imaging*, vol. 13, no. 3, pp. 450–460, 1994.

[12] T. Wakahara, "An Iterative Image Registration Technique Using Local Affine Transformation," *Systems and Computers in Japan*, vol. 21, no. 12, pp. 78–89, 1990.

[13] Y. Amit, "A Nonlinear Variational Problem for Image Matching," *SIAM Journal on Scientific Computing*, vol. 15, no. 1, pp. 207–224, 1994.

[14] J. Noack and D. Sutton, "An Algorithm for the Fast Registration of Image Sequences Obtained with a Scanning Laser Ophtalmoscope," *Physics in Medicine & Biology*, vol. 39, no. 5, pp. 907–915, 1994.

[15] R.P. Woods, S.R. Cherry and J.C. Mazziotta, "Rapid Automated Algorithm for Aligning and Reslicing PET Images," *Journal of Computer Assisted Tomography*, vol. 16, no. 4, pp. 620–633, 1992.

[16] M. Unser and A. Aldroubi, "A Multiresolution Image Registration Procedure Using Spline Pyramids," in Proc. *SPIE*, San Diego, California, U.S.A., July 15-16, 1993, pp. 160–170.

[17] L.G. Brown, "A Survey of Image Registration Techniques," *ACM Computing Surveys*, vol. 24, no. 4, pp. 325–376, 1992.

[18] D.W. Marquardt, "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *Journal of the Society for Industrial and Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.

[19] V.R. Mandava, J.M. Fitzpatrick and D.R. Pickens, III, "Adaptive Search Space Scaling in Digital Image Registration," *IEEE Trans. Medical Imaging*, vol. 8, no. 3, pp. 251–262, 1989.

[20] M. Unser, A. Aldroubi and M. Eden, "B-Spline Signal Processing: Part I—Theory," *IEEE Trans. Signal Processing*, vol. 41, no. 2, pp. 821–832, 1993.

[21] M. Unser, A. Aldroubi and M. Eden, "B-Spline Signal Processing: Part II—Efficient Design and Applications," *IEEE Trans. Signal Processing*, vol. 41, no. 2, pp. 834–848, 1993.

[22] P. van Gelderen, N.F. Ramsey, G. Liu, J.H. Duyn, J.A. Frank, D.R. Weinberger and C.T.W. Moonen, "Three Dimensional Functional MRI of Human Brain on a Clinical 1.5T Scanner," *Proceedings of the National Academy of Sciences*, vol. in Press, 1995.