Fuzzy regularisation of deformation fields in image registration

Antonio Tristán-Vega
Laboratory of Image Processing.
University of Valladolid. Spain
atriveg@lpi.tel.uva.es

Santiago Aja-Fernández
Laboratory of Image Processing.
University of Valladolid. Spain
santi@lpi.tel.uva.es

Abstract

We introduce the use of a variable smoothing kernel, whose width is driven by a fuzzy controller, to regularise a deformation field in the context of image registration. Our experiments show that such a technique outperforms the classical fixed-width regularisation, being capable of removing irregularities in the deformation field while maintaining an adequate adaptive behaviour for localised deformations, thus preserving fine details.

Keywords: Image registration, regularisation, fuzzy.

1 Introduction

Elastic registration is an important and challenging issue in medical imaging [1], which implies the alignment of a moving image or scene with a fixed image or model by means of a free-form deformation. Its applications on real world problems comprise the compensation of brain-shift on neurosurgery interventions, free-hand US volumes compounding, and practically any situation where imaging of soft tissues is needed. The more general form in which a free-form deformation may be represented is a displacement field defined for every voxel in the image [2, 3]. Nevertheless, this technique requires an adequate regularisation in order to achieve visually acceptable results and to eliminate discontinuities in the registered image, due to the uncertainty int the estimation of deformations. However, an excessive regularisation hinders a proper match of the model in zones of localised deformations, so a trade-off is necessary.

There exist quite diverse techniques to regularise the deformation field. The most standard one is the convolution with Gaussian kernels [2, 3], whose variance should be large enough to avoid singularities, but at the same time small enough to preserve details. In [4] the authors proposed to model the images as a deformable material, whose elastic properties are taken into account to regularise the deformation by means of the Navier-Stokes equations. This approach is reformulated in [5] to model the images as viscous fluids, applying the Navier-Stokes equations to the velocity field instead of the deformation field. The problem with this approaches is that they consider a physical model that does not apply to the problem at hand, since the images are neither elastic bodies nor viscous fluids. Besides, the computational complexity grows with this method, since a partial derivatives equation must be solved. Although this difficulty may be overcome with the use of truncated impulse responses [6], this in fact results in a convolution with a smoothing kernel, in quite a similar way to the Gaussian smoothing. Finally, in [7] it is proposed to project the deformation in a more restrictive transformation space, so the deformation is smooth due to the restrictions of such a space. Evidently, the power of the algorithm to recover irregular deformations is then reduced. In this paper, the use of a variable width smoothing kernel is introduced, so that the regularisa-

tion is more aggressive in homogeneous or uniformly deformed regions, meanwhile preserving the local details where needed. In such a way, the smoothing is no longer a convolution with a Gaussian kernel or a truncated impulse response, but a locally adaptive nonlinear operator. This way, the expected qualitative behaviour is translated into semantic descriptions which are further used to design a fuzzy controller, whose output manages the width of the kernel based on local properties of both the images to register and the deformation field itself.

The remainder of the paper is organised as follows: in Section 2 we briefly introduce the registration algorithm to be used, based on Thirion's demons. In section 3 we explain in detail the fuzzy system designed for adaptive regularisation. Section 4 show some significant results that demonstrate the potential of the technique here presented. Finally, in Section 5 we give some final remarks and conclude.

2   The demons algorithm

The demons algorithm was first introduced in [2] by Thirion. It aims to align the isocontours of the images to register by iteratively estimating the displacement field between them by means of the optical flow equation:

$$\mathbf{f}(x,y) = \frac{(m-s)\nabla s}{\|\nabla s\|^2 + (m-s)^2} \qquad (1)$$

where m and s are respectively the grey values of the model and the scene at a given location (x,y), and $\mathbf{f}$ is the correction for the displacement field $\mathbf{d}$ at the current iteration. If we pay attention to Fig. 1, the interpretation of eq. (1) is immediate: the isocontours we have depicted are supposed to correspond to equal grey levels in S and M, and we are supposing as well that inside these isocontours the grey level values are higher than those outside them, so that the gradient of S points inside. In P_1, the difference $m-s$ is then positive (because s is outside the contour, and therefore $s < m$), so P_1 is pushed in the direction of ∇s, which is, insides. On the contrary, in P_2 the
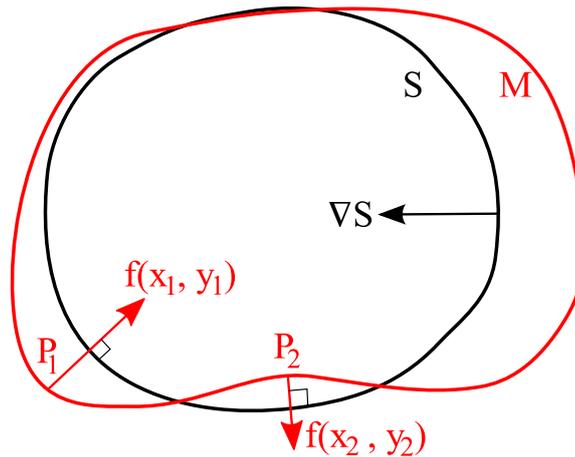

Figure 1: Two corresponding isocontours in the fixed (S) and moving (M) images. We assume that the grey level is greater inside the isocontour, so the gradient points in this direction. The isocontour M is pushed so that it tends to be aligned with that of S.

difference $m-s$ is negative, and P_2 is pushed outsides. Moreover, the larger the error (the value of $|m-s|$) is, the greater correction to the displacement field is applied. When this process is iteratively repeated, the isocontours of both the images are aligned, and therefore the images are coregistered.

The main problem with this formulation is that it requires that the grey levels of corresponding pixels are the same, which makes it too sensitive to noise and not able to deal with changes in contrast. To overcome this difficulty, a generalisation of the algorithm was introduced in [3]; the key is that eq. (1) may be seen as a Levenberg-Marquardt optimisation of the mean squared error (MSE) between the images, and so this similarity measure is substituted in [3] by a local normalised correlation coefficient (NCC). This implies that the grey level of corresponding pixels has not to be the same in both the images, but instead it is enough that a linear relation between them, at least in a local sense, holds. This is especially useful in the case of MRI imaging, since the illumination conditions may severely vary across the image. Whether we use the original formulation of eq. (1) or the improved estimation introduced in [3], a regularisation step is needed for the deformation to be phys-
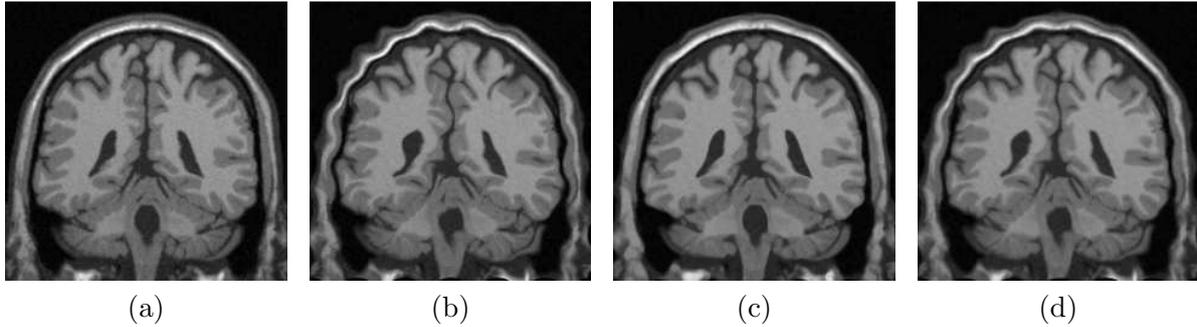
| (a) | (b) | (c) | (d) |

Figure 2: An example of the behaviour of the demons algorithm for a T1-MRI coronal view of a human brain: (a) the original image; (b) the deformed image, following eq. (2); (c) the recovered image using Gaussian convolution-based regularisation, using $\sigma = 0.6$, and (d) using $\sigma = 1.8$. All images are 232×232 pixels in size.

ically acceptable since the estimation of the displacement field is an ill-posed problem, due to the enormous degrees of freedom of the solution [2]. The entire registration algorithm may be summarised in the following steps, comprised in every single iteration[1]:

1. Initialise the deformation field to null displacements at every location (x, y).

2. For the given scene S and the currently deformed model M, use eq. (1) (or its NCC-based analogous) to compute the update of the displacement field ($\mathbf{f}$).

3. Update the displacement field $\mathbf{d}$ by adding the correction computed in the previous step, $\mathbf{f}$.

4. Regularise the so obtained deformation field to compute the current displacements at each image location (x, y).

5. If the predefined maximum number of iterations has not been reached, return to step 2 and continue.

This procedure is common to all of the registration algorithms of this kind, and therefore we do not give more details for the shake of brevity, except for the step 4, which is the focus of our paper. To stress the importance

of an adequate regularisation, let us fix our attention in Fig. 2. We have used a coronal view of a synthetic MRI volume as the fixed image or scene, and scaled it in order to obtain an image 232×232 pixels in size. This phantom is available in the public BrainWeb database[2] [8]. As the moving image to register, and like in [3], we have taken a deformed version of this same 2-D view, obtained from a synthetic deformation field that obeys the following equation:

$$d_i(x, y) = \delta \cos \left(\frac{6 \cdot 2\pi x}{X} \right) \cos \left(\frac{6 \cdot 2\pi y}{Y} \right) \quad (2)$$

where $d_i(x, y)$, with $i = \{x, y\}$, are the displacements at pixel (x, y) in directions x and y, respectively, $X = Y = 232$, and $\delta = 3$. At the sight of Fig. 2, it remains evident that an excessive regularisation, as in case (d), results in the algorithm being incapable of recovering the irregular deformation given by eq. (2); the deformation field in this case varies too fast, so the cut-off frequency of the low-pass filter given by the Gaussian filter with $\sigma = 1.8$ is inadequate. On the contrary, with $\sigma = 0.6$, the corresponding cut-off frequency is one third of that with $\sigma = 1.8$, and the local behaviour of the deformation field is well accounted. It is clear that in large homogeneous regions with a constant grey level , such as the background of the image, or the grey matter of the brain, the regularisation may be more

---

[1]In practise, this iterations are usually enclosed in a multi-resolution scheme, and at the end of each of this levels the obtained displacement field is interpolated to double its resolution. In our paper, we have used 5 resolution levels in all cases.

[2]http://www.bic.mni.mcgill.ca/brainweb/

aggressive than in pixels near the edges because relevant details of the scene will not be lost; moreover, such a regularisation may be preferable, since in homogeneous zones there are no meaningful structures and therefore a reliable estimation of the displacement could not be possible. This uncertainty in the estimation may be partially palliated by assuming that the deformation is smooth enough, and therefore we can use a wider Gaussian kernel to reduce the larger amount of noise. In a similar way, it would be desirable to use a coarser regularisation where the deformations are locally smooth, and vice-versa. This two issues are exploited in the next section to design a regularisation technique based on Gaussian kernels of variable width, depending on local characteristics of both the image and the deformation field.

3 Regularisation of the deformation field

Two desirable properties for our regularisation scheme arise from the example of Fig. 2:

1. If the local variability near a pixel of the scene is small, the width of the Gaussian kernel should be larger due to the lack of reliability in the estimation.

2. If the variability of the displacement field at a given point is far larger compared to the variability of its neighbouring pixels, the Gaussian kernel should be larger, since this corresponds to a weak estimation of the displacement and not to a local property of the field, and vice-versa.

The regularisation step cannot be seen this way as a Gaussian convolution, since our system is no longer shift-invariant. Therefore, we substitute the standard sum $(\mathbf{d} * g_\sigma)(x, y) = \sum_{u,v} \mathbf{d}(u,v) g_\sigma(x-u, y-v)$ with:

$$\tilde{\mathbf{d}}(x, y) = \sum_{u,v} \mathbf{d}(u,v) g_{\sigma(x,y)}(x-u, y-v) \quad (3)$$

where the term $\sigma(x, y)$ accounts for the adaptivity of the Gaussian kernel g, whose width has to be fixed for each location (x, y) by
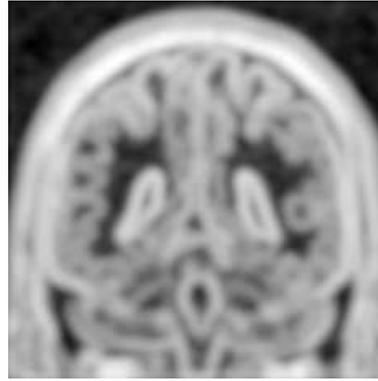

Figure 3: The local variability of the scene, for the example of Fig. 2.

means of a fuzzy system, whose design is described in what follows.

3.1 Local variability of the scene

This parameter is characterised with a simple gradient-based technique: after smoothing the image with a homogeneous Gaussian kernel with $\sigma = 1$, we compute both the components of the gradient as centred differences. The magnitude of the gradient is then smoothed with a Gaussian kernel with $\sigma = 3$, and the resulting image is log-compressed and converted to the interval $[0, 1]$. All these stages, as well as the involved parameters, have been validated using the same example as in Fig. 2, whose result is shown in Fig. 3. It is worth noticeable that this parameter is calculated on the scene (fixed image), so its computation is required only once (or once per resolution level). Without the log-compression and range scaling, the first input to the fuzzy system is:

$$V(x, y) = \|\nabla (s * g_1)(x, y)\| * g_3(x, y) \quad (4)$$

3.2 Local smoothness of deformations

In this case, we compute a local average of the displacement field as a convolution with a Gaussian kernel with fixed $\sigma = 1.4$, which is useful to compute the local deviation of the displacement with respect to its neighbours:
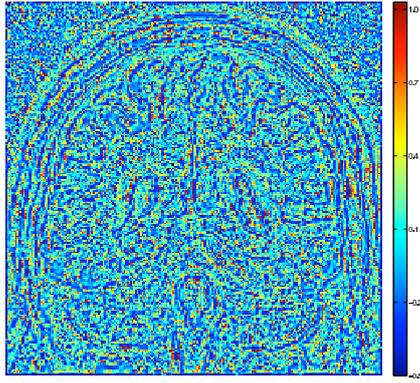
Figure 4: The local irregularity of the displacement field for the example in Fig. 2.

$$\Delta(x,y) = \|\mathbf{d}(x,y) - (\mathbf{d} * g_{1.4})(x,y)\| \quad (5)$$

If this value is large, it may be because the displacement field is irregular near the location (x,y), and in this case the parameter Δ should be equally large for the neighbouring pixels as well; on the other hand, if this large deviation is due to an outlier, neighbouring pixels will not show such an irregular behaviour, so we compare the deviation in (x,y) with the average deviation on its vicinity:

$$I(x,y) = \frac{\Delta(x,y)}{(\Delta * g_\sigma)(x,y)} \quad (6)$$

This value is clipped to the range $[-0.5, 1]$, which has been empirically tested to be adequate in our example, and then log-compressed to obtain the second input to the system. An example of the result of this procedure may be seen in Fig. 4.

When the value of $I(x,y)$ is large, it means that the pixel shows a highly irregular behaviour with respect to its neighbours, so it is probably an outlier; when it is not the case, it may be because the deformation field is locally smooth, if $\Delta(x,y)$ is small, or because we have a fine detail on the deformation, when $\Delta(x,y)$ is larger.

3.3 Inputs to the fuzzy controller

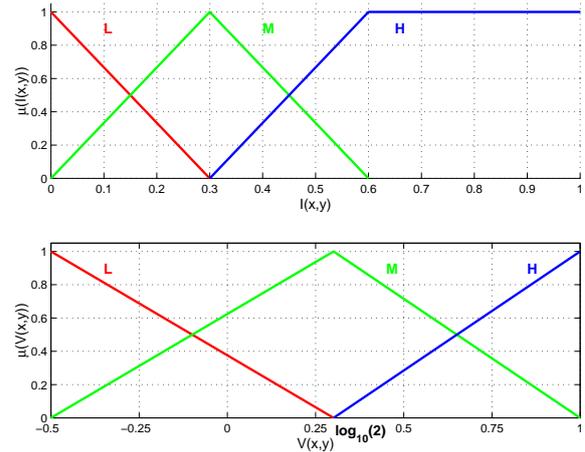The local variability of the scene and the local irregularity of the deformation are trans-

Figure 5: The membership functions of the fuzzy variables $\mathcal{I}$ and $\mathcal{V}$.

lated into two fuzzy variables $\mathcal{I}$ and $\mathcal{V}$ whose universes are the ranges of variation of the parameters we have previously introduced, $\bar{I}$ and $\bar{V}$, so adopting the notation $A_l = \log_{10} A$:

$$\bar{I}(x,y) = \frac{I_l(x,y) - \min_{x,y} I_l(x,y)}{\max_{x,y} I_l(x,y) - \min_{x,y} I_l(x,y)}$$
$$\bar{V}(x,y) = \max\{\min\{V_l(x,y), 1\}, -0.5\} \quad (7)$$

and therefore the universe of $\mathcal{I}$ is $U_{\mathcal{I}} = [0,1] \subset \mathcal{R}$ and the universe of $\mathcal{V}$ is $U_{\mathcal{V}} = [-0.5, 1] \subset \mathcal{R}$. The membership functions are depicted in Fig. 5. We have defined three fuzzy sets in both cases, corresponding to low (L), medium (M) or high (H) values of each linguistic variable. The shape of the fuzzy sets, as well as the rest of the system, have been designed and optimised for the example in Fig. 2.

3.4 Output of the fuzzy controller

The output of the system is a fuzzy variable $\mathcal{O}$ that controls the width of the Gaussian kernel, which may be low (L), medium (M), or high (H); the universe of this variable is $U_{\mathcal{O}} = [0,1] \subset \mathcal{R}$, with 0 corresponding to the minimum σ and 1 to the maximum. Since we use a standard additive model (SAM), there is no need to define the shape of the output sets [9], but it is enough instead to determine their centroids and areas; we use 0, 0.5, and 1 as centroids for the respective L, M, and H sets, and equal areas for all of them.

In practise, the defuzzyfied output is further discretised into seven integer values, which is, $i = \lceil 7 \cdot o \rceil$, with o the defuzzyfied value of the output variable $\mathcal{O}$ and $\lceil x \rceil$ the minor integer greater than x. Then we use i as an index in the vector $\mathcal{C} = \{0.6, 0.8, 1.0, 1.2, 1.4, 1.6, 1.8\}$: the displacement field is regularised with seven Gaussian kernels g_{σ_i}, $\sigma_i \in \mathcal{C}$, and at each (x, y) we choose the i-th output as the deformation field for the location (x, y).

3.5 Rule base of the fuzzy controller

Seven rules are enough to implement the two qualitative requirements we have reasoned in the first part of Section 3 (see Table 1). With a low local variability ($\mathcal{I} = L$), there are no significant structures, so the estimations are not robust and a strong regularisation is needed in all cases ($\mathcal{O} = H$). On the contrary, when $\mathcal{I} = H$ the estimations are more robust, so the regularisation may be very localised ($\mathcal{O} = L$), although it depends on the irregularity of the deformation: if $\mathcal{V} = M, H$, the displacement is most likely to be an outlier, and thus we should perform a stronger regularisation ($\mathcal{O} = M$). With an intermediate local variability ($\mathcal{I} = M$), the situation is almost the same, but since the structures in this case are not so reliable, a greater regularisation is needed ($\mathcal{O} = H$ instead of $\mathcal{O} = M$).

4 Results

To test the algorithm, we have to use a different image from that used for the design step. We choose an axial view of the same T1-MRI BrainWeb phantom [8], with a size of 210×252 pixels, as shown in Fig. 6. This is the fixed image or scene; to obtain the model or moving image, we use a synthetic deformation similar to that of eq. (2), with $X = 210$, $Y = 252$, and a varying δ ranging from 0.4

Table 1: Rule base for the fuzzy controller.

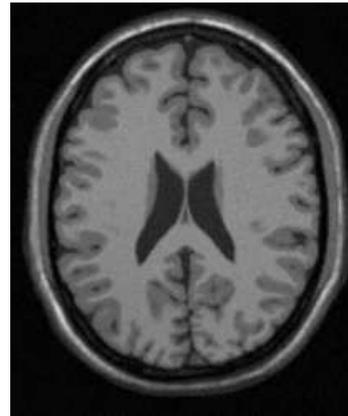| $\mathcal{I}$	$\mathcal{V}$	$\mathcal{O}$	$\mathcal{I}$	$\mathcal{V}$	$\mathcal{O}$	$\mathcal{I}$	$\mathcal{V}$	$\mathcal{O}$
L	*	H	M	L	L	H	L	L
			M	M	M	H	M	M
			M	H	H	H	H	M

Figure 6: We use a 210×252 axial view of a phantom T1-MRI volume to validate the regularisation technique.

to 5.0 which is used to parametrise the extent of the deformation. For each value of δ, we perform eight registration experiments, using (1) a fixed-width Gaussian regularisation with each of the values $\sigma_i \in \mathcal{C}$ (seven experiments), and (2) adaptive regularisation based on our fuzzy controller. As the performance measure, we use the MSE. We do so instead of using the euclidean error in the estimation of the displacements, because in homogeneous zones the estimation is not robust, and this could bias the results; moreover, in these zones the accuracy of the estimation is not so relevant, since for the majority of clinical applications the requirement is that corresponding tissues are aligned, and it is still met even if the displacement is not precisely determined. We present the results in Fig. 7.
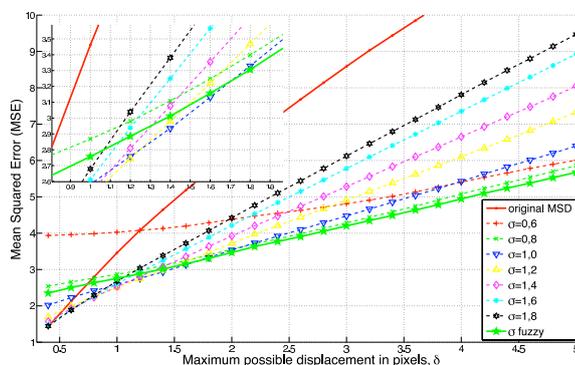

Figure 7: MSE vs. δ for each regularisation technique. We present the original MSE as well for comparison purposes.

As can be seen in the figure, and especially in the zoom over the crossing point of the curves, for small δ (more regular deformations), a stronger regularisation is preferred, meanwhile as δ grows a more localised smoothing yields better results. The crossing point is near $\delta = 1.5$, and from this value he fuzzy regularisation is always better than any fixed Gaussian kernel, although the kernel with $\sigma = 0.8$ shows a similar behaviour. For larger deformations an smaller kernel is preferred (note the difference between $\sigma = 0.8$ and $\sigma = 1.8$), but even so the fuzzy regularisation, that combines kernels as wide as $\sigma = 1.8$ with smaller ones (like $\sigma = 0.6$, with a clearly worse isolated behaviour), is able to yield better results. For smoother deformations a more aggressive regularisation is better, but once again the fuzzy regularisation reaches a compromise (note the difference with $\sigma = 0.6$ in this case). Summarising, the fuzzy technique is able to give a trade-off over a wide range of possible deformations. Given the similar behaviour with the case of fixed $\sigma = 0.8$, a question arises about the need to use a variable Gaussian kernel; to justify it, we introduce Fig. 8, where the experiments in Fig. 7 are repeated, but in this case we use $\mathcal{C}' = \{0.6, 0.7, 0.8, 0.9, 1.0, 1.2, 1.4\}$ in the fuzzy controller. For large deformations we achieve more accurate results, with a more noticeable difference with fixed kernels, but at the expense of a worsening for smaller deformations. Nevertheless, there still exists a compromise (flatter response) over the entire
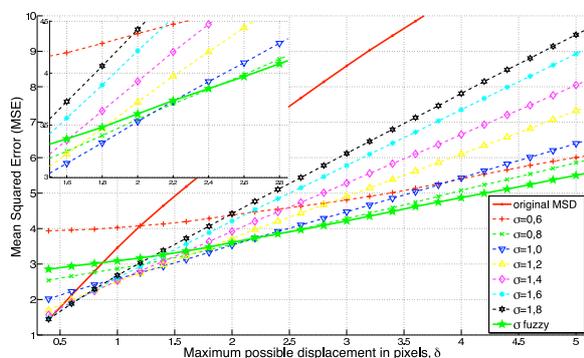
range of δ. It is worth notice that we have not changed the fuzzy system at all, but only the interpretation of its output.

As a final result, we give in Fig. 9 a glance of the performance of the fuzzy system at a given iteration, where we show the inputs, the activation of each rule, and the output that controls the width of the corresponding Gaussian kernel. As can be seen from the output, all variances σ_i^2 are used at some point, which partially guarantees the consistency of our approach, since it proves that the combined use of several Gaussian kernels (as shown in the final picture of Fig. 9) yields better results (as shown in Figs. 7 and 8) than the isolated use of each of them.

5 Conclusion

We have showed that the combined use of Gaussian kernels with different widths is an efficient way to adaptively regularise a deformation field in the context of image registration. We have considered two design issues: first, in homogeneous regions the estimations are not robust, and therefore more regularisation is needed. Besides, a great local variation of the deformation field may be due to a variety of factors, and we must distinguish between the presence of outliers and the actual existence of highly irregular local deformations. The implementation and validation of a fuzzy controller based on these two qualitative criteria demonstrates their usefulness. Regarding the fuzzy controller, it has been proved that the same design may be adapted, only by changing the range of the output, so that it yields a better performance either for large, irregular deformations, or for small, smoother ones. This makes of our fuzzy system a very powerful tool if *a priori* knowledge of the extent of the deformations may be retrieved for a given problem. The main drawback of this technique is its computational overload; we perform seven convolutions per iteration instead of a single one in order to exploit the matrix-based computation of MatLab, in whose language we have programmed the entire algorithm, but more conservative

Figure 8: MSE vs. δ for each regularisation technique. In this case the fuzzy controller uses narrower kernels.
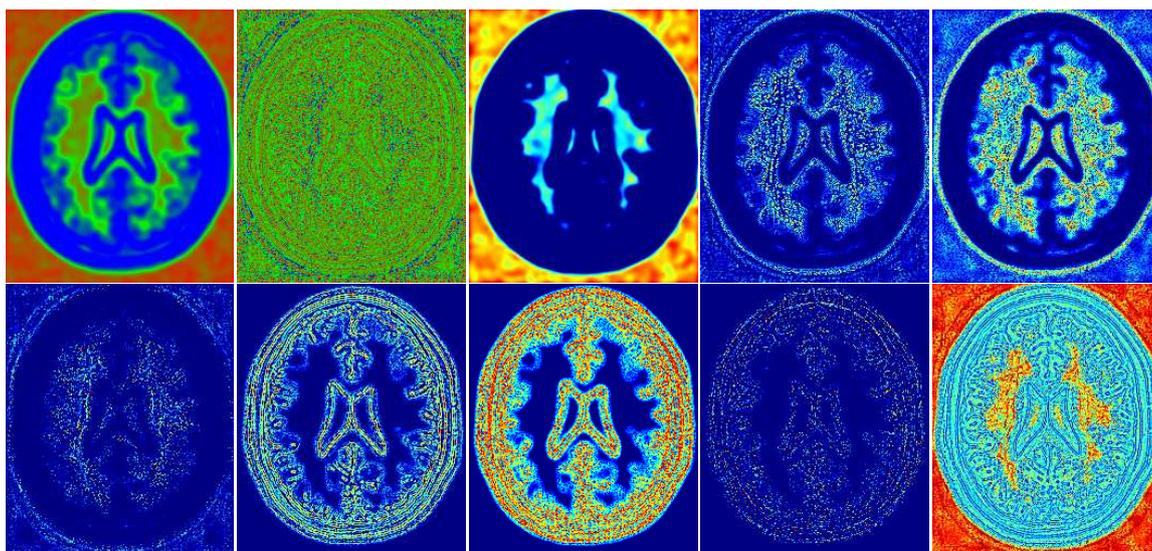
Figure 9: A glance of the fuzzy system at a given iteration; from left to right, the first two images are the two fuzzy inputs with the sets L, M, H coded as the R, G B, channels. The remaining 3 images of the upper row and the four images at the left in the second row are the seven rules described in Table 1, coded in pseudo-colour. The final image is the pseudo-colour representation of the output.

solutions could be found.

References

[1] B. Zitová and J. Flusser, "Image registration methods: A survey," *Im. Vision And Comput.*, vol. 21, pp. 977–1000, 2003.

[2] J.-P. Thirion, "Image matching as a diffusion process: an analogy with Maxwell's demons," *Med. Im. Anal.*, vol. 2, no. 3, pp. 243–260, Sept. 1998.

[3] P. Cachier and X. Pennec, "3D non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions," in *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, June 2000, pp. 182–189.

[4] G.E. Christensen, M.I. Miller, and M.W. Vannier, "A 3-D deformable magnetic resonance textbook based on elasticity," in *AAAI Spring Sympos.*, Palo Alto, CA, March 1994, pp. 153–156.

[5] G.E. Christensen, R.D. Rabbitt, and M.I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. on Im. Proc.*, vol. 5, no. 10, pp. 1435–1447, 1996.

[6] Morten Bro-Nielsen and Claus Gramkow, "Fast fluid registration of medical images," in *Procs. of VBC '96*, London, UK, 1996, pp. 267–276, Springer-Verlag.

[7] Mathieu De Craene, Aloys du Bois d'Aische, Benoit Macq, and Simon Warfield, "Incorporating metric flows and sparse jacobian transformations in ITK," *The Insight Journal*, Jan-Jun 2006.

[8] D.L. Collins, A.P. Zijdenbos, V. Kollokian, J.G. Sled, N.J. Kabani, C.J. Holmes, and A.C. Evans, "Design and construction of a realistic digital brain phantom," *IEEE Trans. on Med. Im.*, vol. 17, no. 3, pp. 463–468, 1998.

[9] Bart Kosko, *Fuzzy Engineering*, Prentice Hall, 1997.