

Diffeomorphic Demons using Normalised Mutual Information, Evaluation on Multi-Modal Brain MR Images

M. Modat^a, T. Vercauteren^b, G.R. Ridgway^a, D.J. Hawkes^a, N.C. Fox^a and S. Ourselin^a

^aUniversity College London, UK. ^bMauna Kea Technologies, Paris, France.

ABSTRACT

The demons algorithm is a fast non-parametric non-rigid registration method. In recent years great efforts have been made to improve the approach; the state of the art version yields symmetric inverse-consistent large-deformation diffeomorphisms. However, only limited work has explored inter-modal similarity metrics, with no practical evaluation on multi-modality data. We present a diffeomorphic demons implementation using the analytical gradient of Normalised Mutual Information (NMI) in a conjugate gradient optimiser. We report the first qualitative and quantitative assessment of the demons for inter-modal registration. Experiments to spatially normalise real MR images, and to recover simulated deformation fields, demonstrate (i) similar accuracy from NMI-demons and classical demons when the latter may be used, and (ii) similar accuracy for NMI-demons on T1w-T1w and T1w-T2w registration, demonstrating its potential in multi-modal scenarios.

1. INTRODUCTION

The demons algorithm¹ is a well-established technique for non-rigid registration. Its key advantage over competing algorithms like fluid² and Free-Form Deformation (FFD)³ is computational efficiency. Recent work has improved the demons' registration quality and its applicability. The diffeomorphic approach⁴ makes the algorithm suitable for statistical morphometric analyses that require topology-preservation. The symmetric and inverse-consistent formulation⁵ enhances reliability for longitudinal data, and permits consistent forward and inverse transformations (e.g. allowing both voxel-wise and vertex-wise data to be transformed). These modifications have been accomplished without unduly increasing the demons' computation time, thus maintaining its speed advantage. However, the demons framework has not yet been shown to be suitable for multi-modal applications; while some papers have proposed Mutual Information as a similarity measure,⁶⁻⁸ they have shown no successful results on practical inter-modal examples.

We present a novel implementation of the diffeomorphic demons using Normalised Mutual Information (NMI) with its analytical gradient in a non-linear conjugate gradient optimiser. Our evaluation employs simulated data from BrainWeb,⁹ using its perfectly aligned T1- and T2-weighted images to perform several related comparisons. Firstly, we compare the performance of atlas construction when registering T1w images to the T1w BrainWeb template using the conventional mean-squared error (MSE) demons algorithm and our new NMI-demons. We then compare the latter to the novel use of the NMI-demons on registering the same T1w images to the T2w template. In a third experiment, we first use an FFD algorithm to register the BrainWeb T1w image to a number of different subjects' T1w scans, we then apply the estimated transformations to the T1w BrainWeb image, simulating a set of images with known ground-truth correspondence to BrainWeb. This allows quantitative measurement of the error in registering these images back to either the T1w or T2w BrainWeb originals, evaluating both intra- and inter-modal accuracy.

2. METHOD

2.1 Diffeomorphic demons

For fixed image F and moving image M , typical non-linear registration algorithms seek a spatial transformation g that optimizes a weighted average of an image similarity term $\text{Sim}(F, M \circ g)$ and a regularization term $\text{Reg}(g)$. Coupling of these terms usually leads to computationally intensive optimization. In contrast, Thirion proposed to follow an unregularized optical-flow step by a decoupled regularization using Gaussian smoothing.¹ By

Send correspondence to Marc Modat, E-mail: m.modat@ucl.ac.uk

introducing a hidden variable c that represents observed point correspondences, the demons algorithm can be seen as an optimization of a global energy,¹⁰ where a distance term $\text{dist}(g, c)$ ensures g remains close to c :

$$E(c, g) = \frac{1}{\sigma_i^2} \text{Sim}(F, M \circ c) + \frac{1}{\sigma_x^2} \text{dist}(g, c)^2 + \frac{1}{\sigma_T^2} \text{Reg}(g). \quad (1)$$

The auxiliary variable justifies the demons’ two step procedure as a well-posed alternate optimization over c and g . The correspondence step optimizes $\frac{1}{\sigma_i^2} \text{Sim}(F, M \circ c) + \frac{1}{\sigma_x^2} \text{dist}(g, c)^2$, with respect to c , for fixed g . This is typically done by a gradient descent step starting from $c = g$. The regularization step optimizes $\frac{1}{\sigma_x^2} \text{dist}(g, c)^2 + \frac{1}{\sigma_T^2} \text{Reg}(g)$, with respect to g , for fixed c , typically by simple Gaussian smoothing. σ_i , σ_x and σ_T are weights that are not explicitly used in the practical implementation.

The classical demons algorithm works on the complete space of dense non-parametric spatial transformations and uses a second-order Newton-type optimisation (see §2.3) of an MSE criterion, $\text{Sim}(F, M \circ c) = \|F - M \circ c\|^2$. This approach was extended⁴ to generate smooth topology-preserving maps by working on a space of diffeomorphisms. Efficient implementation is achieved by using diffeomorphisms that are the exponential map of stationary velocity fields, and hence can be computed rapidly using a scaling and squaring approach. In practice, the diffeomorphic demons involves both the above-mentioned Gaussian smoothing (of the displacement field, analogous to diffusion regularisation) and a second Gaussian smoothing of the velocity field — a visco-diffusion model¹¹ analogous to the fluid algorithm.² There are two tunable parameters for the Gaussian widths, σ_{diffn} and σ_{fluid} .

2.2 Normalised mutual information and its gradient

The Normalised Mutual Information (NMI) similarity measure¹² aims to quantify the amount of information one image has about another. It is suitable for multi-modal applications as it does not rely on a simple functional relationship between intensities. $\text{NMI} = (H(F) + H(M \circ g)) / H(F, M \circ g)$, where $H(F)$ and $H(M \circ g)$ are the marginal entropies of the fixed and warped images, and $H(F, M \circ g)$ is their joint entropy. To optimize NMI we use its analytical gradient $\partial \text{NMI} / \partial \mu$ for each degree of freedom μ , which we have previously derived.¹³

2.3 Optimisation

As previously mentioned, the classical demons optimises the MSE using a Newton-type method (multiplying the gradient by the inverse-Hessian). Because MSE is a local similarity measure, and the demons has a ‘non-parametric’ voxel-wise transformation, the Hessian consists only of 3×3 blocks on the diagonal. This can be seen by considering the gradient with respect to movement of one voxel; the rates of change of the x-, y- and z-components at that voxel may depend on each other but not on the motion of other voxels.

This approach performs well, but it cannot be applied in the multi-modal case. NMI is a global metric where all voxels with the same intensity are linked through the joint histogram, so the Hessian is much less sparse than for MSE. The Hessian’s size (square in the number of voxels — e.g. $10^6 \times 10^6$), therefore makes inversion almost impossible, and we have found simple approximations of the NMI Hessian are inadequate. We therefore propose conjugate gradient (CG) optimisation, requiring only the first derivative of NMI. The interleaving of descent steps and Gaussian smoothing means it is not theoretically clear that CG will work well, but in practice, we have found that it significantly improves on gradient descent.

A disadvantage of CG compared to Newton-type optimisation is that an appropriate step-size must be found. Furthermore, the scale of the NMI gradient is not compatible with realistic displacement; in one example on 3D brain MRI, the update field of MSE-demons showed displacements of ~ 1 voxel, while the maximal NMI-demons step was just 10^{-6} . We therefore rescale the NMI gradient to have maximal displacement norm of 1 voxel.

3. EXPERIMENTS

3.1 Comparison of the MSE- and NMI-demons

The difference between the MSE criterion and the NMI is assessed using T1w brain images. Forty diagnosed Alzheimer’s disease (AD) patient scans and 18 age-matched control scans are registered affinely and then non-rigidly to the (fixed) BrainWeb T1w image. Scans were acquired on a 1.5 T GE Signa scanner with an inversion-prepared fast SPGR sequence (TE=6.4 ms, TI=650 ms, TR=3s, bandwidth=16 kHz, $256 \times 256 \times 128$ matrix, field

of view $240 \times 240 \times 186$ mm). We used AD patients and control subjects to provide a large variety of deformations to recover (controls are quite similar to the target image, whereas the AD patients have much larger ventricles). Prior to any registration, a histogram-equalisation procedure non-linearly rescaled the voxel intensities between 0 and 64 for every image (to match the fixed image intensity histogram). Both MSE-demons and NMI-demons used 3 coarse-to-fine resolution levels. The number of iterations has been set to 50 for each level to ensure convergence in both algorithms. Different smoothing parameters have been selected for the two algorithms, based upon visual assessment of result and Jacobian images from test registrations. σ_{fluid} , has been set to 2 for both versions of the demons whereas σ_{diffn} is 2 for MSE-demons but 0.8 for NMI-demons.

3.2 Comparison of mono- and multi-modal NMI-demons

The 58 moving images described previously were all registered to the BrainWeb T2w image (aligned with the T1w fixed image used previously). The NMI-demons parameters were identical to the ones in the former experiment.

3.3 Assessment on multi-modal cases

Using 10 controls and 10 diagnosed AD scans from the 58 previously described, we generated known deformation fields. Firstly, the BrainWeb T1w image has been registered to each scan using the FFD. Secondly, the warped image was considered as fixed and, using the MSE- and NMI-demons, the deformation field was recovered from either the T1w or the T2w image. The T1w and T2w image being perfectly aligned, we had the possibility to directly compare the obtained result with a ground truth deformation field.

4. RESULTS AND DISCUSSION

4.1 Registration metrics

The MSE and NMI values were computed for each registered image and are summarized in the first part of Table 1, along with the overall extrema of the Jacobian determinant for each method and modality. As might be expected, the NMI-demons leads to higher NMI than the MSE-demons and *vice-versa*, i.e. each registration is ‘best’ by its own metric. The similarity of the ranges of the jacobian determinants suggests that the regularisation effects are similar, supporting our choice of different σ_{fluid} for MSE- and NMI-demons. Note that negative Jacobians can still arise in the diffeomorphic demons due to discretisation effects.⁴

		Affine (T1w-T1w)	MSE-demons (T1w-T1w)	NMI-demons (T1w-T1w)	NMI-demons (T1w-T2w)
Registration of 58 T1w images to BrainWeb T1w/T2w template	MSE value	72.6 (11.3)	27.3 (3.9)	30.8(5.2)	
	NMI value	1.094 (0.057)	1.149 (0.07)	1.161 (0.07)	1.159 (0.07)
	$ Jac $ extrema		[-0.93 9.03]	[0.04 8.30]	[0.05 7.31]
Recovery of 20 FFD simulated deformations	Absolute error		0.4677 (0.0235)	0.1025 (0.0318)	0.5473 (0.0255)

Table 1. Quantitative registration results, presented as ‘mean (standard deviation)’ or ‘[minimum maximum]’.

4.2 Atlasing results

The warped images from all subjects were averaged within-method and are displayed in Figure 1; all three are very similar, demonstrating that both intra- and inter-modal NMI-demons give similar results to MSE-demons. It appears that the MSE-demons result preserves slightly less detail in finer cortical folds. However, we only claim there is no evidence for marked inferiority of the NMI-demons. We cannot suggest that the NMI criterion outperforms the MSE one, because there are numerous differences in the two implementations: in σ_{diffn} ; in the use of CG or Newton-type optimisers; and in their moving-image interpolation (our NMI-demons code uses cubic spline, while the original MSE-demons uses linear). It seems inevitable that MSE-demons would be preferred when the data truly satisfy the assumption of additive Gaussian white noise, but it is conceivable that NMI-demons might be better even for intra-modal data if there are minor acquisition differences, e.g. across scanners.

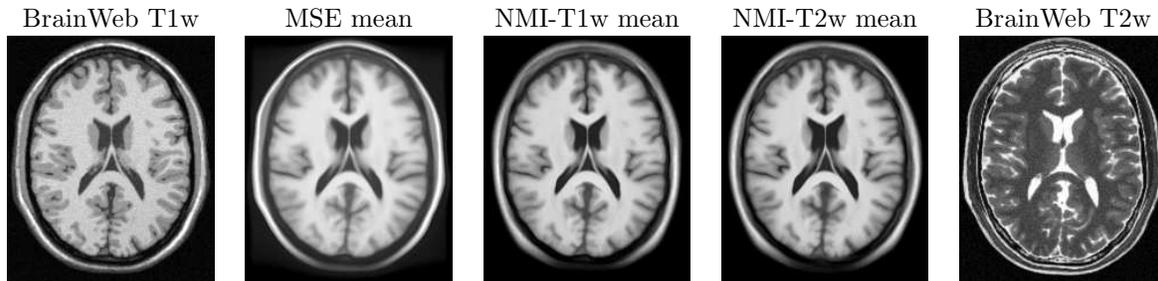


Figure 1. (b-d) Averages of 58 T1w brain MRI registered to (a) T1w, or (e) T2w, BrainWeb templates.

4.3 Registration errors

The bottom half of Table 1 presents results for the absolute errors at each voxel integrated over the brain area. It can be seen that the NMI-demons within modality performed better than the other two methods, which are themselves statistically significantly different, but in practical terms would be considered very close for most applications. The error between the ground truth and the recovered deformation field can be explained by the difference of deformation (and regularisation) model between the demons algorithm and the FFD.

4.4 Conclusion

Our novel NMI-demons performs similarly on intra- and inter-modal examples, and is comparable to the classical MSE-demons on examples where the latter is applicable. There is no significant difference in computation time. Future work could involve the exploration of theoretical and practical aspects of the optimisation scheme. Inter-modal capability is a valuable addition to the growing spectrum of fast demons-based methods.

REFERENCES

- [1] Thirion, J.-P., “Image matching as a diffusion process: An analogy with Maxwell’s demons,” *Med. Image Anal.* **2**(3), 243–260 (1998).
- [2] Christensen, G. E., Rabbitt, R. D., and Miller, M. I., “Deformable templates using large deformation kinematics,” *IEEE Trans. Image Process.* **5**, 1435–1447 (1996).
- [3] Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., and Hawkes, D., “Nonrigid registration using free-form deformations: Application to breast MR images,” *IEEE Trans. Med. Imag.* **18**, 712–721 (1999).
- [4] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N., “Diffeomorphic demons: Efficient non-parametric image registration,” *Neuroimage* **45**, S61–S72 (Mar. 2009).
- [5] Vercauteren, T., Pennec, X., Perchant, A., and Ayache, N., “Symmetric log-domain diffeomorphic registration: A demons-based approach,” in *[Proc. MICCAI’08]*, 754–761 (2008).
- [6] Guimond, A., Roche, A., Ayache, N., and Meunier, J., “Multimodal brain warping using the demons algorithm and adaptative intensity corrections,” *IEEE Trans. Med. Imag.* **20**, 58–69 (Jan. 2001).
- [7] Hermosillo, G., Chef d’Hotel, C., and Faugeras, O., “Variational methods for multimodal image matching,” *Int. J. Comput. Vis.* **50**, 329–343 (Dec. 2002).
- [8] Tao, G., He, R., Datta, S., and Narayana, P. A., “Symmetric inverse consistent nonlinear registration driven by mutual information,” *Comput. Methods Programs Biomed.* (2009). In press.
- [9] Cocosco, C., Kollokian, V., Kwan, R., and Evans, A., “Brainweb: Online interface to a 3D MRI simulated brain database,” *Neuroimage* **5**(4), 425 (1997).
- [10] Cachier, P., Bardinet, E., Dormont, D., Pennec, X., and Ayache, N., “Iconic feature based nonrigid registration: The PASHA algorithm,” *Comput. Vis. Image Underst.* **89**, 272–298 (Feb. 2003).
- [11] Modersitzki, J., *[Numerical methods for image registration]*, Oxford University Press (2004).
- [12] Studholme, C., Hill, D. L. G., and Hawkes, D. J., “An overlap invariant entropy measure of 3D medical image alignment,” *Pattern Recognit* **32**, 71–86 (1999).
- [13] Modat, M., Ridgway, G. R., Taylor, Z. A., Hawkes, D. J., and Fox, N. C., “A parallel-friendly normalized mutual information gradient for free-form registration,” *Proceedings of SPIE* (Jan 2009).