Sparse update for loopy belief propagation: fast dense registration for large state spaces

Pengdong Xiao Nick Barnes Paulette Lieby Tiberio Caetano
College of Engineering & Computer Science
The Australian National University
Canberra, Australia 2601
and
National ICT Australia
pengdong.xiao,nick.barnes,paulette.lieby,tiberio.caetano@anu.edu.au

Abstract

A dense point-based registration is an ideal starting point for detailed comparison between two neuroanatomical objects. This paper presents a new algorithm for global dense point-based registration between anatomical objects without assumptions about their shape. We represent mesh models of the surfaces of two similar 3D anatomical objects using a Markov Random Field and seek correspondence pairs between points in each shape. However, for densely sampled objects the set of possible point by point correspondences is very large. We solve the global non-rigid matching problem between the two objects in an efficient manner by applying loopy belief propagation. Typically loopy belief propagation is of order m^3 for each iteration, where m is the number of nodes in a mesh. By avoiding computation of probabilities of configurations that cannot occur in practice, we reduce this to order m^2. We demonstrate the method and its performance by registering hippocampi from a population of individuals aged 60-69. We find a corresponding rigid registration, and compare the results to a state-of-the-art technique and show comparable accuracy. Our method provides a global registration without prior information about alignment, and handles arbitrary shapes of spherical topology.

1. Introduction

3D shape matching is a key problem in computer vision and medical image analysis. Elegant diagnostic results are arising from analysis of correlations between pathology and neuroanatomical shape [7]. However, building an accurate model of shape across a population [20], or performing detailed comparisons between anatomical objects depends on high quality registration.

This paper describes a new global dense point-based registration method that is computationally efficient. The approach takes a mesh model of spherical topology and represents it as a Markov Random Field (MRF). We then apply loopy belief propagation (LBP) to find matches between two such meshes. To make the algorithm efficient, we have developed a new sparse update for belief propagation that exploits the state space topology inherent to spatial matching. Our method provides a global registration without prior information about alignment, and handles arbitrary shapes of spherical topology.

Dense registration is common practice in computer vision, such as bundle adjustment in multi-view geometry [8]. It is also a well established technique for scanned 3D data of manufactured objects [10]. However, in those cases a near exact match between points is possible as the structures are identical. In contrast, for registration of neuroanatomical objects across individuals, or of the same individual over a time period, there will be structural differences.

There are many methods for shape representation and matching, see [15] for a review. Standard approaches to matching include Fourier analysis (e.g. [19] for 2D and [4] for 3D), moments [11], curvature scale space [5, 12], level sets [16], wavelet transform [14], and shape contexts [3]. Most approaches avoid the combinatorics of dense correspondences by matching a lower dimensional representation (e.g., moments [11], first order (FO) ellipsoid (FOE) [9]). The approach by Kelemen *et al.* [9] aligns the axes of the ellipsoids corresponding to the FO spherical harmonic coefficients of the meshes to obtain the registration.

Alignment using the FO ellipsoid is effective in practice. However it assumes that: (i) the shapes are pre-aligned to avoid the rotational and symmetry ambiguities of ellipsoids; and, (ii) they effectively can be approximated by an ellipsoid. For example FOE alignment would not work with

spherical or cuboid objects. The registration method we propose does not rely on an assumption of pre-alignment. We demonstrate registration of shapes without unique main axes for alignment.

We may solve the registration problem between two meshed surfaces as a correspondence matching problem between the nodes of the two surfaces. Based on the work published in [17], we may frame this as energy minimisation using MRF techniques. Here we embed the meshed surfaces into a Markov random field framework, in which geometric information (in this paper, curvature) is captured by the nodes, and topological information is preserved via their Markov network.

In [17], dense registration over large point sets is performed. However, minimisation of the energy function is NP hard. Thus for a moderate sized mesh, an exhaustive search rapidly becomes intractable. Gibbs sampling as used in [17] does not scale to shapes whose meshes have more than 2000 nodes. In this paper, we investigate applying LBP to the probabilistic inference problem of minimising the energy function.

Although belief propagation is guaranteed to converge to a global maximum for tree-like structures [13], the meshes we consider in this paper have loops, so global convergence is not guaranteed. However, in practice, belief propagation is known to find good solutions for MRFs with loops; loss of guaranteed convergence is the price to be paid for tractability (see [2]).

Although LBP will typically converge, it is still computationally intensive: it is of order kmn^2, where k is the number of iterations, m is the size of the MRF, and n is the number of states at each node. For this paper, the two meshes generally have a similar number of nodes, as human hippocampi are of similar size. Thus, since the number of possible assignments for each node –the state space– is the number of nodes of the other mesh, so that $m \sim n$, the order is approximately km^3. In this paper we exploit the fact that due to the topology of the meshes involved we may take advantage of a locality principle in matching. We reduce the computational complexity to order $km^2\overline{n}$, where $\overline{n}$ is the mean number of neighbours of a node. In our experiments, for hippocampal meshes of more than 2000 nodes, the time for an iteration is reduced by a factor of 150 with respect to standard LBP. Multiple iterations are required for LBP convergence. For large databases of objects to be registered this becomes a critical time saving.

The remainder of this paper is organised as follows. Section 2 formulates the correspondence matching problem of anatomical shapes in the MRF framework and describes the energy function for correspondence. Section 3 gives the belief propagation formulation of the matching problem while Section 4 describes the sparse update for belief propagation. Section 5 describes one method for extracting a rigid registration from point correspondences; this method is implemented in the experiments which are presented in Section 6. The experiments show the registration of 23 hippocampi selected from a database of individuals aged 60-69.

2. MRF representation for dense shape matching

In this section, we re-present the framework of the method in [17] which solves the curvature-based matching problem. Note that in [17] the Gibbs sampler is used to solve the problem; here, in Section 3, we detail our approach using loopy belief propagation.

Let M and M' be two meshed surfaces, and X be an MRF with $m = |M|$ nodes each with a state space of size $n = |M'|$. Following [17], we perform a one-to-one mapping $M \rightarrow X$ which preserves Gaussian curvature at nodes and mesh topological structure at edges in X. The correspondence matching problem between M and M' is then an energy function minimisation problem in X, and a point correspondence between M and M' is a realization of the MRF X in M'. By finding the most likely variable assignment of the MRF we then obtain the best match under the energy function specified by the MRF.

We define this energy function H as

$$H(x) = H_1(x) + H_2(x), \qquad x \in X, \tag{1}$$

where $H_1(x)$ describes the energy function for the singletons of X and is based on curvature difference and $H_2(x)$, the energy function for the 2-cliques, is based on pairwise distances.

3. Loopy belief propagation for dense shape matching

For ease of exposition, we introduce the following notation: given $i \in \{1, \ldots, m\}$, we associate node i in M to the i th variable in X, denoted by X_i, and to x_i, an arbitrary realisation of X_i in M'. As seen in Section 2 the energy function $H(x)$, $x \in X$, comprises two components. The first component $H_1(x)$ is defined on singletons and is based on curvature difference. Let $K(u)$ denote the Gaussian curvature of u, u a node of M or M', and let C_1 and C_2 denote the set of singletons and 2-cliques of X respectively. Then

$$H_1(x) = \sum_{X_i \in C_1} h_1(x_i), \quad \text{with } h_1(x_i) = |K(X_i) - K(x_i)|^2. \tag{2}$$

The second component of $H(x)$ incorporates spatial restrictions based on the pairwise distance defined on 2-element cliques:

$$H_2(x) = \sum_{(X_i, X_j) \in C_2} h_2(x_i, x_j), \quad \text{with } h_2(x_i, x_j) = d(x_i, x_j), \tag{3}$$

where d is the distance between x_i and x_j in M'. That is, d is the distance between the nodes that are assigned as correspondences to the 2-clique (X_i, Y_j). These pairwise distances preserve the topological information of the shape.

For a pairwise MRF the probability measure for the joint distribution with respect to a configuration x in X is $P(x) = \frac{1}{Z}\exp(-H(x))$ with Z a normalisation constant. This gives

$$P(x) = \frac{1}{Z}\prod_i \exp\left(-|K(X_i) - K(x_i)|^2\right) \prod_{(i,j)} \exp\left(-d(x_i, x_j)\right). \tag{4}$$

Setting $\phi_i(x_i) = \exp\left(-|K(X_i) - K(x_i)|^2\right)$, $\psi_{ij}(x_i, x_j) = \exp\left(-d(x_i, x_j)\right)$ yields

$$P(x) = \frac{1}{Z}\prod_i \phi_i(x_i) \prod_{(i,j)} \psi_{ij}(x_i, x_j). \tag{5}$$

The typical approach to belief propagation (see [18]) relies on the concept of message passing: a message m_{ij} is passed from node i to node j indicating the belief node i has about the state node j should be in. In BP the belief b_i at a node i is proportional to $\phi_i(x_i)$ and the product of all incoming messages:

$$b_i(x_i) = k\phi_i(x_i) \prod_{j \in N(i)} m_{ji}(x_i), \tag{6}$$

where k is a normalisation constant and $N(i)$ denotes the set of neighbours of i. The messages between nodes are updated as follows:

$$m_{ij}(x_j) = \sum_{x_i} \left[\phi_i(x_i)\psi_{ij}(x_i, x_j) \prod_{k \in N(i)\setminus j} m_{ki}(x_i)\right]. \tag{7}$$

Equation (7) is the basis of the sum product algorithm. We note that on the right-hand side, we take the product of all the messages incoming into node i, except the one coming from node j. Further, the summation ranges over all the states x_i the node i may find itself in. Note that message normalisation must be implemented carefully to avoid underflow.

4. Sparse update for LBP for shape matching

When matching two mesh models representing the shape of objects, there is a topological structure which ensures that certain combinations of point matches cannot occur. For example, given a match between a node X_i in X and a node x_i in the other mesh model M', an immediate neighbour Y_i of X_i cannot match a point y_i that is very far from x_i in M'. In other words, one cannot match a point on one hippocampus to the head of the other hippocampus while matching its immediate neighbour to the tail of the other. To embody this we define the pairwise potential function ψ_{ij} as a step function:

$$\psi_{ij}(x_i, x_j) = \begin{cases} 1, & \text{if } x_j \in N(x_i) \\ 0, & \text{otherwise,} \end{cases} \tag{8}$$

where $N(x_i)$ is the set of neighbours of x_i, $x_i \in M'$. This is an implicit property of spatial matching. This case will occur in a general class of problems where the structure of the state space prohibits certain configurations of matches.

Given Equation (8), when a message is passed from node i to node j, this message will have a positive value if x_i is in the neighbourhood of x_j, and zero otherwise. Thus, the sum product of Equation (7) may now be stated as a max product:

$$m_{ij}(x_j) = \max_{x_i \in N(x_j)} \left[\phi_i(x_i)\psi_{ij}(x_i, x_j) \prod_{k \in N(i)\setminus j} m_{ki}(x_i)\right]. \tag{9}$$

In (9) we have replaced the sum operator by the max operator to privilege the state among $N(x_j)$ which has highest probability.

In Equation (8), any match outside the immediate neighbourhood will have zero probability. Such a cut-off is beneficial to computation speed as each expansion to the neighbourhood increases the computation time by a constant factor. We conducted trials relaxing the neighbourhood constraints (including nodes with shortest distance of up to four) and found negligible benefits in registration quality. This suggests that the restriction of belief over neighbourhoods only marginally impacts on the final converged state of the MRF.

For a given node neighbourhood, max product will give a high score if at least one node in the neighbourhood has high probability, whereas sum product only gives a high score if the average over neighbours is high. Therefore max product may assist with convergence, and may facilitate better performance over small neighbourhoods. This was verified when conducting trials using max product vs sum product; the former showed reliable and faster convergence throughout.

5. Rigid registration

Three non-collinear point correspondences are sufficient to give a rigid body registration. To find such a registration one must select from the entire set of possible correspondences a point set that is consistent with a good registration. We achieve this by applying Random Sample Consensus (RANSAC) [6], whereby we choose a random set of three non-collinear point correspondences to form a rigid body transformation, and then check the Euclidean distance between the transformed locations of all remaining point correspondences. When choosing sets of three non-collinear

points we ensure that those points are spatially apart in order to avoid poor hypotheses arising from instability of close point sets.

We perform many draws of three points, and choose the transformation with the largest inlier set. The inlier set will consist of those points whose distance from their corresponding points will be under some threshold. The rigid body transformation is then formed by taking the L2 norm best fit over the inlier set.

Note that methods other than RANSAC may be applicable here and may give better results. Further, the inlier set may be chosen according to some other criteria than distance between point correspondences. Indeed we may compute the Jaccard index between the shapes to evaluate the quality of the tentative registration. However, computation of the Jaccard index is computationally expensive and improved results come at the cost of complexity.

6. Results

We have performed 3D rigid body registration experiments using human hippocampus data. We chose 23 hippocampi of the non-demented subjects aged 60-69 from a large database Open Access Series of Imaging Studies (OASIS), which is publicly available at http://www.oasis-brains.org/. OASIS is a project aimed at making MRI data sets of the brain freely available to the scientific community. In this age band, some atrophy will be present, making this an interesting test of rigid body registration. We register each of the 23 hippocampi to another arbitrarily selected hippocampus from the same age band of the database. For each pair we visually verified the match in a 3D rendering (e.g., Figure 1). We also registered the same pairs using FOE alignment as in [9], where ambiguities arising from the eight symmetries implicit in this method were solved by L2 norm alignment to a given template (referred to subsequently as FOE-template).

We found that all registrations were globally correct, i.e., there were no rotations (note that no flips about planes may occur in this method). Overall, from visual inspection we found the results were of comparable quality to FOE-template registration. In some cases MRF gave better results than FOE-template (e.g., Figure 1(a) and (b), Jaccard index MRF: 0.7198, FOE-template: 0.6946); in other cases the reverse was true (e.g., Figure 1(c) and (d), MRF: 0.6292, FOE-template: 0.7019).

MRF gave a visually correct registration for all of the 23 hippocampi. Over the 23 registered pairs, the average Jaccard index was 0.6852 for MRF and 0.6927 for FOE-template.

The actual point-wise correspondences between two well registered hippocampi are shown in Figure 2(a), and the inliers found among these are shown in Figure 2(b). For clarity we only show matches from the tail of the hippocampus. It can be seen that although there are many poor matches, there is an inlier set of correct correspondences that the rigid registration technique finds to give a globally correct registration. From the RANSAC process, we found that inlier sets generally approximated 15% percent of the nodes.

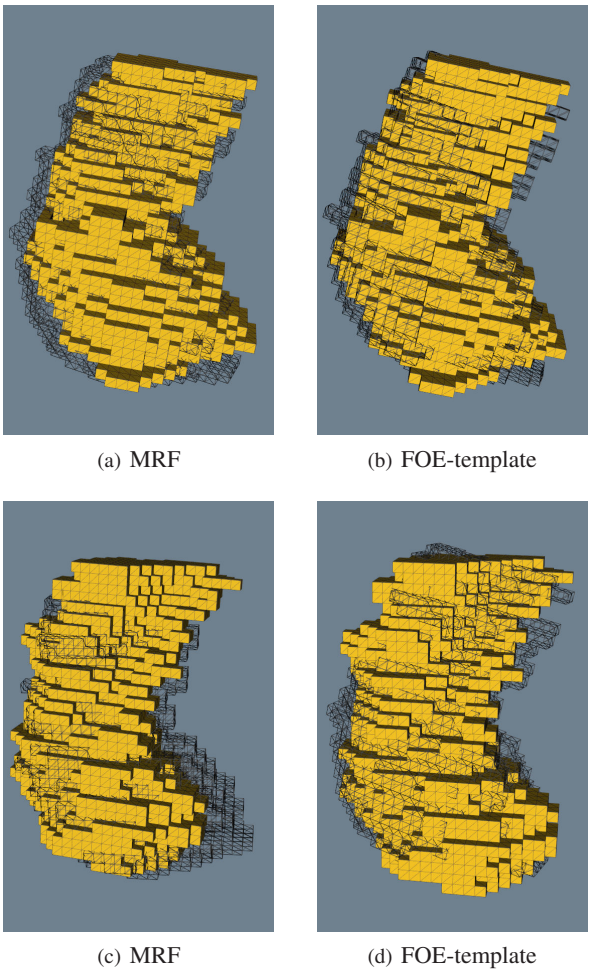

(a) MRF (b) FOE-template

(c) MRF (d) FOE-template

Figure 1. *MRF vs FOE-template registration: (a), (b): good registration of MRF compared with FOE-template approach; (c), (d): an example where FOE-template based approach performs better.*

We have also used a clamping technique to increase the size of the inlier set of correct correspondences. The clamping technique combines the belief propagation message updating algorithm with the 3D rigid body registration process. The clamping technique clamps the inlier set of correct correspondences obtained by 3D rigid body registration in the previous iteration before performing the next iteration of LBP message update. After a few iterations, typically 2 to 4, between LBP and 3D rigid body registration, the ratio of the size of the inlier set to the number of nodes can be increased to more than 90%. Figure 2(c) and (d) show shape matching results using the clamping technique. The

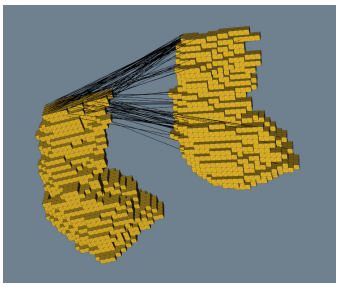

(a) MRF matches

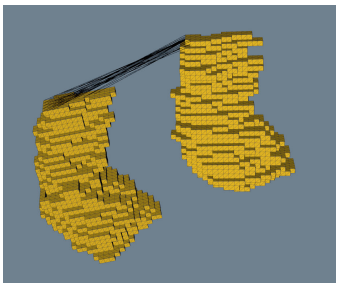

(b) Inliers

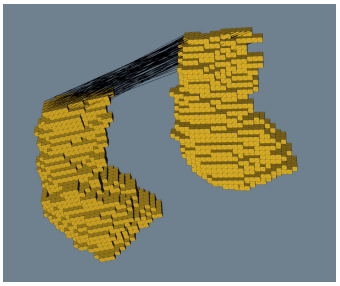

(c) MRF matches after clamping

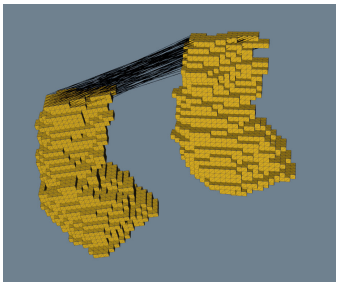

(d) Inliers after clamping

Figure 2. *(a) Matching results of two hippocampi of a good registration, showing only matches from an area of the tail of the hippocampus. There is a clear set of correct and consistent matches, amongst a set of incorrect and inconsistent matches. (b) Subset of inliers of MRF matches found by RANSAC. (c) Matching results of the same hippocampi of a good registration after using the clamping technique. The correct matching rate is increased significantly. (d) Subset of inliers of MRF matches found by RANSAC after using clamping technique.*

correct matching rate, the rate of the inlier numbers during the 3D rigid body registration, is 95% after four alternative iterations. These 95% of points constitute a point-by-point non-rigid registration. The inconsistent matches are identified in this process.

In order to demonstrate that our method works for registration of non-ellipsoidal shapes, we perform the registration of a non-ellipsoidal structure with the same structure whose curvature at each node is disturbed. The non-ellipsoidal structure is publicly available [1]. We disturbed the curvature by adding random and uniformly distributed noise with intervals from [-10%,10%] to [-90%,90%]. Figure 3 illustrates the registration for the non-ellipsoidal structure.

7. Conclusion

We have formulated a 3D dense point-based registration approach based on MRFs. We have presented a sparse update algorithm for loopy belief propagation that renders the dense registration computationally efficient. The algorithm reduces the computational complexity from order m^3 for each iteration for standard loopy belief to order m^2 by exploiting the topology of the state space.

By combining MRF based registration with RANSAC to select the correct matches, we were able, in all the cases in our dataset, to find a correct global registration up to a quality that is comparable with standard methods in use today. Our MRF-based technique may be used in an automated pipeline for registration when the shapes under consideration are not of ellipsoidal form and when their general orientation is not known *a priori*. This later problem has occurred in our study of ageing hippocampi where the shapes had to have their global alignment visually checked for shape analysis, as an L2 norm alignment to a template may not necessarily resolve the ambiguities inherent in a FOE fitting.

References

[1] http://www.cse.wustl.edu/~cmg/meshes.php. 5

[2] S. M. Aji and R. J. McEliece. The generalized distributive law. *IEEE Transactions on Information Theory*, 46(2):325–343, 2000. 2

[3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(24):509–522, 2002. 1

[4] C. Brechbuhler, G. Gerig, and O. Kulbler. Parameterization of closed surfaces for 3-d shape description. *Computer Vision and Image Understanding*, 61(2):154–170, 1995. 1

[5] G. Dudek and J. K. Tsotsos. Shape representation and recognition from multiscale curvature. *Computer Vision and Image Understanding*, 68(2):170–189, November 1997. 1

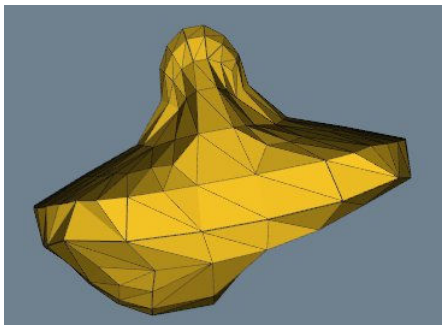

(a) Non-ellipsoidal structure

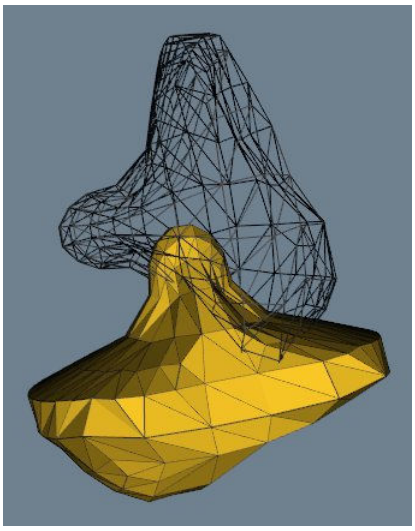

(b) Before registration

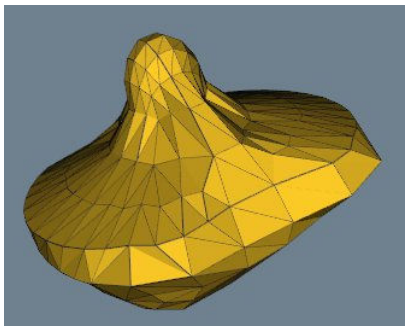

(c) Global minimum

Figure 3. *Registration of a non-ellipsoidal structure: (a) a non-ellipsoidal structure, (b) the non-ellipsoidal structure and the same structure whose curvature at each node is disturbed before registration, (c) global minimum.*

[6] M. A. Fischler and R. C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 3

[7] P. Golland, W. E. Grimson, M. E. Shenton, and R. Kikinis. Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis*, 9(1):69–85, 2005. 1

[8] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition, 2004. 1

[9] A. Kelemen, G. Szekely, and G. Gerig. Elastic model-based segmentation of 3d neuroradiological data sets. *IEEE Trans. Medical Imaging*, 18(10):828–839, Oct. 1999. 1, 4

[10] P. Liang and J. S. Todhunter. Representation and recognition of surface shapes in range images: A differential geometry approach. *Computer Vision, Graphics, and Image Processing*, 52(10):78–109, 1990. 1

[11] J.-F. Mangin, F. Poupon, E. Duchesnay, D. Riviere, A. Chaia, D. L. Collins, A. C. Evans, and J. Regis. Brain morphometry using 3d moment invariants. *Medical Image Analysis*, 8(3):187–196, 2004. 1

[12] F. Mohanna and F. Mokhtarian. An efficient active contour model through curvature scale space filtering. *Multimedia Tools and Applications*, 21:225–242, 2003. 1

[13] J. Pearl. Probabilistic reasoning in intelligent systems: Network of plausible inference, 1988. San Mateo, California: Margan Kaufmann Publishers, INC. 2

[14] I. E. Rube, M. Ahmed, and M. Kamel. Wavelet approximation-based affine invariant shape representation functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(2):323–327, February 2006. 1

[15] R. C. Veltkamp and M. Hagedoorn. State-of-the-art in shape matching. Technical Report UU-CS-1999-27, Department of Information and Computing Sciences, Utrecht University, 1999. 1

[16] B. Venuri, J. Ye, Y. Chen, and C. Leonard. Image registration via level-set motion: Applications to atlas-based segmentation. *Medical Image Analysis*, 7(1):1–20, 2003. 1

[17] P. Xiao, N. Barnes, T. Caetano, and P. Lieby. An mrf and gaussian curvature based shape representation for shape matching. In J. O. ad E Angelopoulou, G. Kamberov, and Q. Dinh, editors, *IEEE CVPR Workshop Beyond Multiview Geometry: Robust Estimation and Organization of Shapes from Multiple Cues*, Minneapolis, USA, June 2007. 2

[18] J. Yedidia, W. Freeman, and Y. Weiss. Understanding belief propagation and its generalizations. In *Exploring Artificial Intelligence in the New Millennium*, pages 239–269, 2003. 3

[19] C. Zahn and R. Roskies. Fourier descriptors for plane closed curves. *IEEE Transactions on Computers*, 21(3):269–281, March 1972. 1

[20] L. Zhou, R. Hartley, P. Lieby, N. Barnes, K. Anstey, N. Cherbuin, and P. Sachdev. A study of hippocampal shape difference between genders by efficient hypothesis test and discriminative deformation. In *MICCAI 2007*, 2007. 1