

Registration of PET and MR hand volumes using Bayesian Networks

Derek Magee¹, Steven Tanner¹, Michael Waller², Dennis McGonagle², and
Alan P. Jeavons¹

¹ School of Computing/Academic Unit of Medical Physics, University of Leeds, UK

² Leeds Teaching Hospitals NHS Trust, Leeds, UK
drm@comp.leeds.ac.uk

Abstract. A method for the non-rigid, multi-modal, registration of volumetric scans of human hands is presented. PET and MR scans are aligned by optimising the configuration of a tube based model using a set of Bayesian networks. Efficient optimisation is performed by posing the problem as a multi-scale, local, discrete (quantised) search, and using dynamic programming. The method is to be used within a project to study the use of high-resolution HIDAC PET imagery in investigating bone growth and erosion in arthritis.

1 Introduction

In this paper we present a novel method for the non-rigid registration of high-resolution HIDAC Positron Emission Tomography (PET) and Magnetic Resonance³ (MR) scan volumes of human hands. To our knowledge we are the first to tackle this particular multi-modal registration problem. This work is part of a wider project to investigate the use of high-resolution list-mode QuadHIDACTM PET imagery ($\sim 0.5mm^3$ voxel size, Fluorine-18 tracer) for the study of the location of bone growth and erosion in the hands of patients suffering from arthritis. Our method involves fitting a pair of models, based on a set of cylindrical tubes, to the two data sets to be registered and calculating a rigid and a non-rigid (piecewise rigid) transform. The models are fitted by the optimisation of a set of Bayesian networks with respect to annotated volumes. The method is made computationally tractable by posing it as a multi-scale local search problem. Quantisation of the search spaces allows the efficient use of dynamic programming to obtain a globally optimal solution in these spaces.

Lower resolution PET imaging ($\sim 5mm^3$ voxel size) is widely used in brain imaging to extract functional information. However, PET imaging provides little anatomical information. Therefore, it is routinely used in conjunction with MR imaging, which can provide this anatomical information. As PET and MR scans are rarely co-located, data sets must be registered. To quote Myers [1]; “[this has] become a matter of routine in the analysis of brain PET studies”. The most popular and successful methods of non-fiducial (physical marker) based registration are based around the maximisation of mutual information, or voxel similarity, over a rigid transform (e.g. [2–4], see [5] for an overview and evaluation of a number of such techniques). Clearly such rigid transforms are unsuitable in our application domain (figure 5). PET has also been applied to cardiac imaging.

³ T2-weighted, fat suppressed, spin-echo coronal images (voxel size $\sim 0.5 \times 0.5 \times 2mm$).

Here PET provides metabolism information, which must again be augmented with anatomical information from MR imaging. Akin to our work, Makela *et al.* [6] use a model based technique to perform multi-modal registration. This is based on the fact that the thorax and lung surfaces are clearly visible in both imaging modalities. A deformable template model is fitted to each data set, and the results used to calculate a rigid transformation. Farahani *et al.* [7] describe a prototype system for the spatio-temporally co-located acquisition of PET and MR for brain imaging. There remain some technical hurdles to be overcome with this approach. If this were widely used (unlikely any time soon due to cost and technical constraints) this will eliminate the need for volume registration. However, multi-modal volume registration is likely to be required for some time to come, especially for high-resolution PET scans (where few uni-modal scanners exist, let alone multi-modal scanners). In our application PET imaging provides information about bone growth and/or erosion. As with previous applications of PET, MR imaging is required to provide anatomical information such as blood vessel and tendon location.

Our approach to the registration problem is partly inspired by the 3D geometric models used in visual tracking. Perhaps the first example of such a model is the WALKER model of Hogg [8] where the body and limbs of a human are modelled as a collection of cylinders. This model is fitted to 2D visual data using edge information. Rehg and Kanade [9] use a similar approach to track the human hand. A simplified 28 D.O.F. cylinder based model is fitted to stereo data using a local optimisation method (Levenburg-Marquart). Stenger *et. al* [10] use a similar model based on a number of 3D quadratics, rather than cylinders, for added realism. Visual tracking has the advantage over medical image analysis applications of having state estimate(s) from previous timesteps to work from. As such, local optimisation over such a large configuration space is possible (although a good initialisation method is required). Felzenswalb and Huttenlocher [11] describe how the effective dimensionality of such optimisation problems may be reduced for certain classes of model using quantisation and dynamic programming. This work is our starting point, and it is described more fully in the following section.

2 Model Fitting as Bayesian Network Optimisation using Dynamic Programming

Felzenswalb and Huttenlocher [11] pose the problem of fitting a model to sensor data (in the visual tracking domain) as the optimisation of the parameters of a directed graph structured network (a type of Bayesian network). Probabilistic dependencies exist between the data and individual model graph nodes, and between parent and child nodes joined by the directed vertices. In general, finding a globally optimal solution over such a network is an N-P complete problem (and in practice costly approximate methods are often used). However, if the model graph is tree-structured (i.e. there are no loops), and the space of solutions is quantised, a globally optimal solution may be obtained in linear time (w.r.t. the number of nodes in the graph) using dynamic programming (see [11] for full details). The solution in fact has complexity $O(q^2n)$ (where n is the number of nodes, and q is the number of quantisations of each node parameterisation). Furthermore, the method is made up of two parts; individual node evaluation (complexity $O(qn)$), and belief propagation over the network (complexity $O(q^2n)$). If

the the majority of the complexity of the network dependencies is in the former, the solution is approximately linear in both the number of nodes, and the number of quantisations of the individual node parameterisations. The phalanges (rigid finger sections) of a human hand in our PET/MR registration problem may be modelled as tree-structured networks, as illustrated in figure 1.

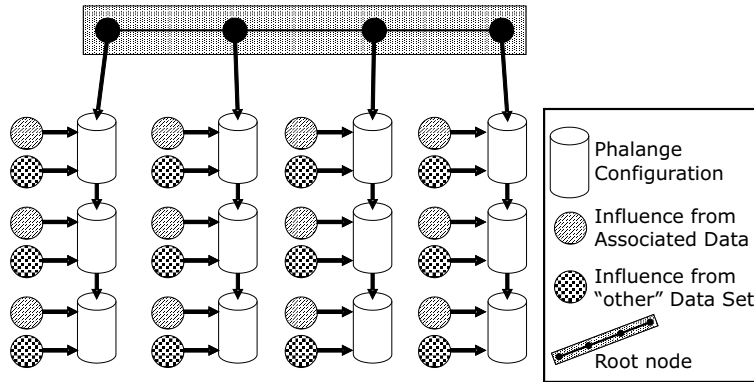


Fig. 1. Hand (four fingers) PET/MR Bayesian Networks Formulation

In our scenario there are two models to be fitted to two data sets (PET and MR). The data sets are linked in that they relate to the same individual. As such, phalange lengths should be similar. The phalanges are parameterised as a tube with 7 parameters (3D start/end and radius). We choose not to constrain the motion of the fingers to a plane (as in [9]), as this is not a useful approximation for our data. However, even with a small number of quantisations per dimension this produces an excessively large number of configurations of each phalange (individual node evaluation has complexity $O(q_d^{n_d})$, where q_d is the number of quantisations in each dimension, and n_d is the number of dimensions). The following sections describe an approach that uses two networks per data set, and local quantisation to overcome this computational hurdle.

3 Efficient Model Networks and their Optimisation

The raw data used in this paper is volumetric and comes from PET and MR scans. Typically volume dimensions are $256 \times 256 \times 16$ for the MR data and $256 \times 256 \times 440$ for the PET data. With this volume of data it is desirable to pre-process the data before evaluating possible model configurations for computational reasons. Another reason for data pre-processing is to identify volumetric areas relating to physical features to assist the model matching/registration process. As there is essentially no common information in the PET and MR scans, other than that an area is within the hand (or not), this is what is used. One common method of identifying regions is volume segmentation. In this process each voxel is labelled as belonging (or not belonging) to a physical structure. For our data this is problematic, as there is much uncertainty over many voxels (especially in the PET data). This task is therefore hard for a skilled human expert, let alone an automated system. This approach also performs little data reduction. We propose a simpler alternative approach to segmentation; point annotation. The principle behind this is that a relatively sparse set of points that

are definitely “hand” or “not hand” are identified. This is currently performed in 2D by hand (taking <15 minutes per data set), and is illustrated in figure 2. We




Fig. 2. Examples of Hand Annotation of Data in 2D: a) MR ‘hand’, b) PET ‘not hand’, 3D Visualisation of ‘hand’ PET & MR point clouds

believe annotation to be a much simpler process to automate (using region based classifiers) than segmentation for this data, as uncertain points may simply be excluded from the labelling process. It is also a much faster process to carry out by hand. Automation of this process is planned for the near future. The 2D point annotation of multiple volumetric slices is used to form 3D “point clouds” using the slice number to form the third dimension. This is illustrated in figure 2.c. These point clouds allow the definition of probabilistic metrics over the model configuration parameters (see later), and allow quite effective visualisation of the quality of final model configurations.

3.1 Model Initialisation and Local Search

Our approach to reducing the computational complexity of the Bayesian network optimisation is to pose the optimisation as a local search problem. Initial model configurations are specified by hand, and the set of possible configurations investigated are local offsets from this configuration. This allows a rather smaller number of quantisations of each dimension than quantising the complete 7D space of solutions for each phalange. Hand specification of the initial configuration is done by clicking on the approximate start/end points of each phalange in a 2D slice view (as in the annotation in the previous section). The end of a parent

phalange and the start a connected child phalange are deemed to be coincident, thus only 16 clicks are required for four fingers (the thumb is ignored only as it is absent in the majority of the data sets used). The radius of each phalange is estimated as $K \times$ phalange length (where K is typically 0.25, which is usually a slight underestimate). Figures 3.a and 3.d show examples of the variable quality of this initial configuration. In particular, the length (and thus radius) of the finger tip phalanges in the MR data is incorrect as they lie partially outside the scan volume. The location of PET data phalanges, especially at the base of the fingers, is up to 10mm from the correct location. Estimated radii in all cases are very approximate, even when the length is near correct.




Fig. 3. Example of model Initialisation (MR Data): a) Initial guess, b) Corrected Tube lengths, c) Output of ‘Greedy’ algorithm, (PET Data): a) Initial guess, b) Corrected Tube lengths, c) Output of ‘Greedy’ algorithm

Two methods are applied to the initial configurations before Bayesian Network optimisation is performed. Firstly a check is made on whether the ends of a phalange lie very close to the edge of the scan volume (i.e. within 5% of the complete range of any dimension from the edge). We term this “phalange validity”. If it is the case that one phalange is valid and the other is not the length of the invalid phalange is set to be equal to the valid one. This is done using the valid end of the invalid phalange and its direction vector, as in equation 1.

$$End_{new} = Start_{orig} + Dir_{orig} \times Len_{other} \quad (1)$$

The results of this process are illustrated in figures 3.b and 3.e (which are data sets imaged from the same individual, and only the MR data has invalid phalanges). Once the lengths are corrected, a “greedy” algorithm is used to obtain a better initial estimate of the radius. This works by assigning annotated hand points to phalanges that contain them⁴. The radii of each tube is increased by increasing factors (typically 1:1.5 in 0.01 steps). The minimum radii that contains a larger number of unassigned annotated hand points than the initial estimate, without containing points assigned to another phalange, is selected. If increasing the radius by one step increases the number of “not hand” annotated points contained within the phalange tube, without increasing the number of “hand” annotated points, expansion of that phalange is halted. This process is repeated until convergence. The results of this process are illustrated in figures 3.c and 3.f.

⁴ If a point is inside two tubes it is assigned to the tube it is furthest inside.

3.2 Model Fitting using Multi-scale Optimisation of Two Networks

If each dimension of the 7D configuration space of each phalange tube is divided into 9 evenly spaced quantisations⁵ (centred around the initial configuration) the total number of quantisations (q) of each tube is $9^7 = 4,782,969$. As the complexity of the belief propagation is $O(q^2n)$ evaluation of such a network takes of the order of an hour or more on current standard hardware (PIV 3GHz). As we wish to perform this optimisation at multiple scales this would make the registration process rather time consuming. As a computationally efficient alternative, we divide the model into two networks; one relating to the start/end point configuration of the phalange tubes (6D) and one relating to the Radius (1D). Optimisation of a pair of the former is performed in around 30 seconds, and the later is optimised in interactive time. Tables 1-4 describe the various probabilistic factors that are used to form the model networks. These fall into two categories; i) Simple probabilities, calculated as a data proportion (or similar), and ii) Normalised Gaussian probabilities, calculated from some distance d , as in equation 2. Normalised Gaussians are used as $P = 1$ for $d = 0$ (i.e. no influence is had).

$$P = e^{-\frac{d^2}{2\sigma^2}} \quad (2)$$

Variable	Description	Form	Notes
P_{cgl_start} , P_{cgl_end}	Distance of start/end from 'curve gradient line' start/end	Norm. Gaussian	S.D. = $0.5 \times \text{length}$ [1 if 'invalid']
P_{assign}	Proportion of 'assigned' data enclosed	Simple prob.	If < 0.95 set to 0.01
$P_{n_unassign}$	Proportion of nearby (within search range) 'unassigned' data enclosed	Simple prob.	Clipped at 0.9 and scaled [0,1]
P_{len_p}	Difference of length from mean of PET & MR original lengths of this tube	Norm. Gaussian	S.D. = $0.5 \times \text{length}$, =PET/MR orig if the other 'invalid', 1 if both 'invalid'
P_{len_sim}	Difference in length from same tube in other data set (current guess)	Norm. Gaussian	S.D. = $0.25 \times \text{length}$

N.B. Where length is specified in the notes, this is the initial estimate for that data set
Table 1. Probabilistic influence factors on individual tube configurations (position/length)

The Curve gradient line (CGL), used to calculate P_{cgl_start} in table 1), is calculated for each phalange from the mean of the points initially associated with that phalange. For each finger (trio of phalanges), a quadratic is fitted to the points in the 2D plane defined by these points. The gradient (direction) of this quadratic at each mean point defines the CGL (a straight line which passes through the mean). The extent (start/end) of these lines are found by projecting each of the associated points onto the line and calculating the maximum distance in each direction from the mean point. Validity of these lines is calculated in exactly the same way as described in section 3.1. Independence of each factor is assumed, and individual phalange tube configuration influences are combined as a product (equation 3).

$$P_{tube_n} = P_{cgl_start_n} \times P_{cgl_end_n} \times P_{assign_n} \times P_{n_unassign_n} \times P_{len_p_n} \times P_{len_sim_n} \quad (3)$$

Independence of topological factors is also assumed, and configuration influences are combined as a product (equation 4).

⁵ This is approaching the minimum sensible before solutions are lost between steps

Variable	Description	Form	Notes
P_{near}	Distance between parent end and child start	Norm. Gaussian	S.D. = 0.025×parent length
$P_{straight}$	Dot product of unit direction vectors of parent/child	Simple prob.	Enforces finger straightness

Table 2. Probabilistic influence topological factors between parent/child configurations (position/length)

$$P_{topol_{n,m}} = P_{near_{n,m}} \times P_{straight_{n,m}} \quad (4)$$

Variable	Description	Form	Notes
P_{assign}	Proportion of ‘assigned’ data enclosed	Simple prob.	If <0.95 set to 0.01
$P_{unassign}$	Proportion of ‘unassigned’ data enclosed	Simple prob.	Clipped at 0.9 and scaled [0,1]
P_{not}	1 - (No. ‘not hand’ points / No. ‘associated’ hand points)	Simple prob.	Lower clip at 0.01
P_{enc_other}	If no. assigned to another tube enclosed > no. assigned to this tube × 0.05, P_{enc_other} =0.01, else 1	Simple Prob.	
$P_{rad_{match}}$	Difference between radius and initial other model radius	Norm. Gaussian	S.D. = 0.25 × tube radius

Table 3. Probabilistic influence factors on individual tube configurations (radius)

Variable	Description	Form	Notes
$P_{dradius}$	Difference in radius between parent and child	Norm. Gaussian	S.D. = 0.25 × parent radius

Table 4. Probabilistic influence topological factors between parent/child configurations (radius)

Radius model network influences are combined as a product in exactly the same way as the start/end model. Belief propagation through each finger (trio of phalange tubes) of the network also assumes independence of factors and is performed separately for each finger (equation 5).

$$P_{finger} = P_{root} \times P_{tube_1} \times P_{topol_{1,2}} \times P_{tube_2} \times P_{topol_{2,3}} \times P_{tube_3} \quad (5)$$

The globally optimal (maximum) value of P_{finger} for each finger (within the quantised space) is found using dynamic programming. Efficiency savings may be made in this process by bearing in mind all probabilities in equation 5 are < 1 by design. Thus a partial solution that has lower probability than the best solution so far need not be evaluated further in the forward part of the algorithm (in practice saving many evaluations of $P_{topol_{n,m}}$).

P_{root} relates to the probabilistic influence from the root node (figure 1). This is calculated as a normalised Gaussian (equation 2) based on the distance (in the 2D plane defined by the root node, and tube directions) of the start of the base phalange from its “root node projection” (σ is $0.5 \times$ the tube radius). The root node configuration is defined as a straight line in 3D space. This line is calculated by calculating a least squares rigid registration (translation and rotation) between the corresponding start and end points of the PET and MR models (pre-optimisation), using the method of Horn [12]. A straight line is fitted to the 8 registered base phalange start points, again using a least squares error minimisation approach. The pairs of start points (one for each data set) are projected onto this line, and the mean taken as the “root node projection”

for that finger. These points are projected back into the original space for the transformed data set (choice of which set to transform makes little difference to the root node projections calculated). It should be noted that the root node configuration is pre-calculated from the initial configuration rather than optimised within the dynamic programming stage as it is the same for both PET and MR networks (which are currently optimised separately). This gives consistency of the starting points of the phalange tubes, which would be expected for data taken from the same patient. The networks are optimised at increasingly finer quantisations / smaller search ranges using the previous optimal as a starting point, as illustrated in figure 4.




Fig. 4. Multi-scale Network optimisation flowchart

Parent/child end/start points are aligned after each application of the Start-End network optimiser. The iteration detailed in figure 4 is performed 4 times (with the start/end search range halved at each iteration). The total model fitting process (including initialisation) takes just a few minutes to fit models to both data sets

3.3 Non-rigid Transform Calculation

Once the tube based model has been fitted to corresponding PET and MR datasets, it is a reasonably simple process to calculate a non-rigid (piecewise rigid) transform between the two data sets. This transform may be used to warp one of the data sets (we warp the PET data) into correspondence with the other. First, a global rigid (translation and rotation) transform is calculated from the corresponding start and end points of the phalange tubes in the two models. The closed form least-squares error minimisation method of Horn is used [12]. Calculating local rigid transformations for each (globally aligned) tube pair is performed using the same method. However, at least three points are required for this method, and only two are available (the start & end of the tube). A third point is generated by using principal components analysis (PCA) to calculate the eigenvector of the globally aligned annotated hand points (both sets) with the smallest eigenvalue. This ‘minor axis vector’ (V_{ma}) defines a plane with the tube direction (T_{dir}) in which the third point lies. This third point is calculated as in equation 6.

$$P_3 = T_{start} + T_{length} \times D_3 \quad (6)$$

Where D_3 is the vector in the plane defined by V_{ma} and T_{dir} perpendicular to T_{dir} , where the dot product of D_3 and V_{ma} is positive. Transformation of any point or voxel within a single tube is simply a matter of performing the global transform, followed by the appropriate tube transform (or the inverse operations in the opposite order for the reverse transform). Points outside the tubes may simply be transformed using the global transform only. In fact we interpolate the transform near the edges and when a point is in more than one tube. Details are omitted for brevity. Figure 5 shows some results of these transforms applied to annotated point and voxel data.




Fig. 5. Registration results; a)&b) Global (Rigid) and Local (Non-rigid) registration applied to annotated hand points, c) Warped PET slice, d) Corresponding MR slice

4 Evaluation

Our method was evaluated by application to a number of data sets. The principal behind the evaluation is that the closest annotated point in the ‘other’ data set to a registered/transformed ‘hand’ point should be a hand point. Our chosen evaluation metric is to count the proportion of these ‘correct points’ for the transforms in either direction. It should be noted that the absolute values presented are fairly meaningless as PET and MR data sets don’t image exactly the same part of the hand (parts may be missing in one set or another). However, relative values for a data set demonstrate improvement in registration quality. Results are presented in table 5. Results show a statistically significant improvement

Data Set	PET→MR Nearest Correct Prop.	MR→PET Nearest Correct Prop.
A	[0.844/0.866] → [0.920/0.970] (0.026)	[0.774/0.821] → [0.816/0.919] (0.325)
B	[<i>0.834</i> /0.887] → [<i>0.829</i> /0.910] (0.018)	[0.766/0.835] → [0.769/0.909] (0.302)
C	[0.525/0.714] → [0.586/0.906] (0.132)	[0.552/0.762] → [0.651/0.931] (0.374)
D	[0.930/0.931] → [0.951/0.953] (0)	[0.712/0.732] → [0.757/0.841] (0)
E	[0.518/0.520] → [0.704/0.764] (0.054)	[0.689/0.680] → [0.908/0.941] (0.306)
F	[0.880/0.892] → [0.956/0.977] (0.553)	[0.758/0.788] → [0.890/0.914] (0.119)
Mean inc.	[0.069/0.111]	[0.090/0.139]

Results show initial [Global Trans./Local Trans.] → final [Global Trans./Local Trans.] (Raw, for reference)

italics imply registration fit decrease (for data set GJ PET→MR Global registration only)

Table 5. Nearest Neighbour Correct Evaluation Results

(at 5% confidence) in the local and global registrations (especially local) over all data sets after application of our method. Also, local registration always outperforms global registration. Our chosen metric says nothing about the magnitude of registration errors, although by inspection maximum error appears well under 5mm. Ground truth would be required to verify this. We plan experiments with imaging phantoms and pseudo-synthetic data to do this verification.

5 Discussion, Future Work and Acknowledgements

We have posed a problem of non-rigid multi-modal volume registration as one of model fitting by optimisation of a set of Bayesian networks. Quantisation, and local search, allow the efficient use of dynamic programming to find a globally optimal solution (within a locally quantised space). Applying this approach iteratively at multiple, decreasing, scales gives robust model fits within a few minutes. We have applied this approach to the registration of PET and MR volumes of hands, for the study of bone growth/erosion in arthritis. The networks we have proposed include no learned (or measured) prior terms. Such terms could easily be included if sufficiently accurate (and useful) models were available. However, it is debatable whether such models would add anything for the application presented as patients with arthritis often have rather unusual hand poses. In the near future we intend to automate the process of hand/not hand point annotation of volumetric data sets. We intend to further evaluate the exact accuracy of registrations using imaging phantoms with easily identifiable localisation points. Evaluation using pseudo-synthetic data is also planned.

D. McGonagle has been funded by the UK Medical Research Council to carry out the high-resolution imaging studies of arthritis used in this paper.

References

1. Myers, R.: The application of PET-MR image registration in the brain. *The British Journal of Radiology* **75** (2002) 31–35
2. Woods, R., Grafton, S., Holmes, C., Cherry, S., Mazziotta, J.: Automated image registration I. *Journal of Computer Assisted Tomography* **22**(1) (1998) 139–152
3. Wells, W., Viola, P., Atsumi, H., Nakajima, S., Kikinis, R.: Multi-modal volume registration by maximisation of mutual information. *Medical Image Analysis* **1**(1) (1996) 35–51
4. Studholme, C., Hill, D., Hawkes, D.: Automated 3D registration of magnetic resonance and positron emission tomography brain images by multi-resolution optimization of voxel similarity measures. *Medical Physics* **24**(1) (1997) 25–35
5. West, J., Fitzpatrick, M., Wang, M., et al.: Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography* **21**(4) (1997) 554–566
6. Makela, T., Pham, Q., Clarysse, P., Neonen, J., Lotjonen, J., Sipila, O., Hanninen, H., Lauerma, K., Knuuti, J., Katila, T., Magnin, I.: A 3D model-based registration approach for the PET, MR and MCG cardiac data fusion. *Medical Image Analysis* **7**(3) (2003) 377–389
7. Farahani, K., Slates, R., Shao, Y., Silverman, R., Cherry, S.: Contemporaneous positron emission tomography and MR imaging at 1.5T. *Journal of Magnetic Resonance Imaging* **9** (1999) 497–500
8. Hogg, D.: Model-based vision: A program to see a walking person. *Image and Vision Computing* **1** (1983) 5–20
9. Rehg, J., Kanade, T.: Visual tracking of high DOF articulated structures: An application to human hand tracking. In: *Proc. European Conference on Computer Vision*. (1994) 35–46
10. Stenger, B., Mendonca, P., Cipolla, R.: Model-based hand tracking using an unscented kalman filter. In: *Proc. British Machine Vision Conference*. (2001) 53–72
11. Felzenswalb, P., Huttenlocher, D.: Efficient matching of pictorial structures. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*. (2000)
12. Horn, B.: Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America* **4**(4) (1987) 629–642