A comparative analysis of B-spline deformation models in 3D shape matching

T.R.Langerak
University Medical Center Utrecht
Heidelberglaan 100
3584 CX Utrecht, The Netherlands
robin@isi.uu.nl

Y. Song
Delft University of Technology
Landbergstraat 15
2628 CE Delft, The Netherlands
Y.Song@tudelft.nl

ABSTRACT
Lattice-based B-spline driven deformation has become a useful method in shape matching applications. Here, the challenge is to find a configuration of the B-spline deformation model that effectuates a deformation that spatially aligns one shape (the source) to another shape (the target). Literature study indicates that few B-Spline deformation based algorithms were implemented that target polygonal meshes. In contrast, in the field of medical image registration, B-spline deformation has been extensively applied in matching shapes that use a voxel-based shape representation. For exploring the opportunities of applying these voxel-based methods to the shape matching of polygonal meshes, in this paper we propose to match polygon meshes by transforming them to voxel models and apply established techniques from the field of medical imaging. Two voxel-based methods are selected and implemented: Global Optimization methods, which globally optimize the B-spline model, and Markov Random Field methods, which locally optimize the B-spline model. These methods are compared to parameterized B-spline-based shape matching methods previously proposed by the authors. These methods directly match polygon meshes. Results indicate that the proposed methodology outperforms the parameterized approach in terms of accuracy and computation time and therefore is a promising alternative to existing methods.

Keywords
Shape matching, B-spline deformation, freeform deformation, Markov Random Fields, active shape models.

1. INTRODUCTION
In 1986, Sederberg and Parry introduced the concept of lattice-based B-spline driven deformation that allowed the freeform deformation of (in their case) solid geometric models [Sed86]. B-spline driven freeform deformation makes use of the correspondence between a target shape and an overlying grid of B-Spline control points. A change in the configuration of the B-Spline control points results, through a trivariate Bernstein polynomial function, in a deformation of the underlying shape (see figure 1). B-spline deformation has turned out to be an intuitive method of shape modeling, as the nature of the B-Spline paradigm allows both global and local deformation. Long after its introduction, B-spline driven deformation has been and is being applied in shape modeling applications (e.g. [Igar99], [Karp06], [Song06], [Pern05]). However, there is also another important application of B-spline driven freeform deformation: that of shape matching. In shape matching the challenge is to find a transformation that spatially aligns two shapes. This transformation can be given as a transformation matrix, but it can also be enacted by a matrix of B-Spline control points.

B-Spline driven shape matching techniques have been applied to reverse engineering [Lang07][Ver01] and shape retrieval. In addition, B-spline driven shape matching has found an application in the area of medical image registration [Glo08], [Rue99], [Mattes03]. Here, the main advantage of the B-spline deformation model is that it is successful in mimicking natural tissue deformation and expressing anatomical variance.

When looking closely at the existing literature on B-spline driven shape matching, one cannot fail to notice that shape matching has been given much more attention in the field of medical image registration than it has in the registration of polygonal meshes or solid models.

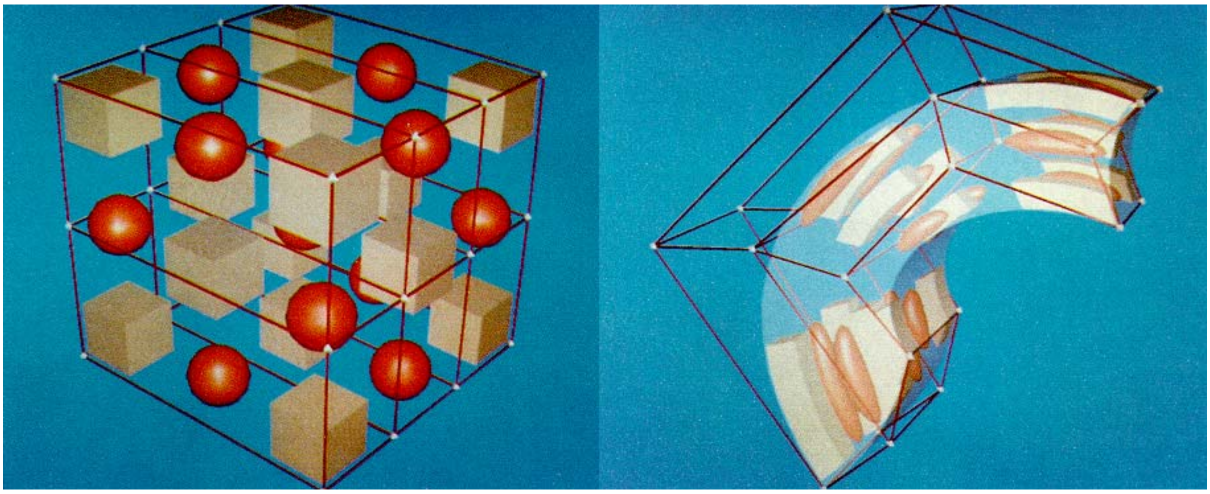

Figure 1: *Example of B-Spline driven freeform deformation (taken from [Sed86])*

Partly, this can be explained by the fact that for voxel models, B-spline driven shape matching is less computationally costly. Because the validity of voxels is easier to maintain (through interpolation) than the validity of polygonal models, computation costs are considerably lower, especially for large and/or complex shapes. In this paper we therefore propose a new technique for matching shapes represented by polygonal meshes (see figure 2).

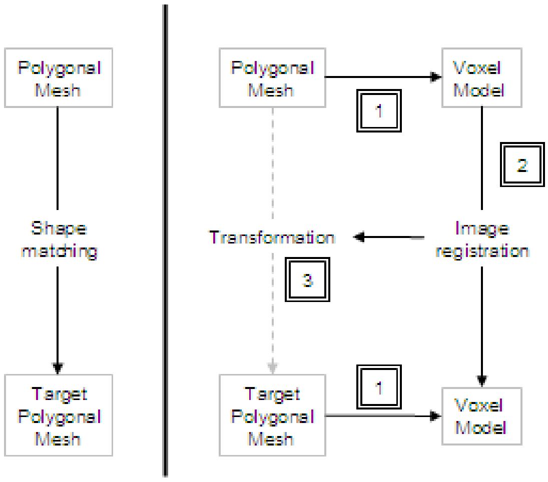

Figure 2: *The direct (left) and the proposed shape matching method (right)*

In a first step, both meshes are converted to a voxel model. Then, using proven techniques from the field of medical image registration, these voxel models are registered. This results in a transformation that can also be applied to the original polygonal mesh to align it with the target mesh, after which only some finetuning is needed. The purpose of this approach is to reduce the computation cost of the shape matching of polygonal meshes, and hopefully increase the accuracy.

In this paper we will investigate two medical image registration methods: global optimization and Markov Random Field minimization. Both will be applied to polygonal models after they have been converted to a voxel model, and the result will be compared to a method that makes use of parametric models for matching polygonal models without converting them to voxel models (i.e. as depicted on the left side of figure 2). This method is an adaptation of a method previously proposed by the authors ([Song05],[Lang07], [Lang08]).

In section 2 we will first define the B-Spline deformation model. In section 3, we will review the three employed shape matching techniques. In section 4 we will compare the performance of all three techniques when applied to four polygonal test models and investigate if the techniques that are taken from the field of medical image registration outperform parametric models. Finally, in section 5 we will draw conclusions and speculate on future research.

2. B-SPLINE DEFORMATION MODELS

The basic principle of the B-spline deformation model is the duality between the parametric coordinates of the model and the spatial coordinates of the shape that is being deformed. The basic elements of a B-spline deformation model are its control points. Although theoretically there is no reason for a specific starting point of the B-spline control points, usually they are assumed to be initially positioned in a uniform lattice, as in figure 1, left. Alternatively, they can be positioned with regard to a previously identified shape [Song05], or in lattices of arbitrary topology [Mac96].

2.1 Definition of the Deformation Model

In three dimensions, the B-spline deformation model $\mathcal{B}$ can be defined as $\mathcal{B}=\{P,d\}$ with P being a set of $l \times m \times n$ control points $P_{ijk} = \left(x_{ijk}, y_{ijk}, z_{ijk}\right)$. For i, j and k it holds that $0 \le i < l$ $0 \le j < m$ and $0 \le k < n$; d is the degree of the B-spline model. Without loss of generality we assume here that B-Splines are uniform and rational, for the reason that most of the methods that are dealt with in this paper have only been described for these types of B-Splines. A NURBS deformation model would also incorporate knot vectors and control point weights and although this makes the B-spline deformation more complex, in an application to shape matching the only relevant difference lies in the dimensionality of the shape similarity function.

A deformation T can be denoted as a function $T(I,B,\Delta)$, where Δ is a set of $l \times m \times n$ vectors that contains, for each control point, a three-dimensional spatial transformation vector $\Delta_{ijk} = (\Delta x_{ijk}, \Delta y_{ijk}, \Delta z_{ijk})$. The coordinates of the B-spline control points can be expressed both in a (x,y,z) coordinate system and in a (u,v,w) coordinate system. The (x,y,z) coordinates indicate the spatial location of the control point. These coordinates can be directly manipulated manually or automatically, and are therefore the main focus of both shape modeling and shape matching techniques. The (u,v,w) coordinates indicate the position of the control point in the parametric space of the B-spline model and it is an important notion of B-spline deformation that this position is invariant under freeform deformation.

Given the (u,v,w) coordinates of an arbitrary point in an image I, the corresponding (x,y,z) coordinates can easily be found using the trivariate Bernstein polynomial:

$$I(x,y,z) = \sum_{i=0}^{l}\sum_{j=0}^{m}\sum_{k=0}^{n} \mu P_{ijk} \text{, where}$$

$$\mu = \binom{l}{i}\binom{m}{j}\binom{n}{k}(1-u)^{l-i}(1-v)^{m-j}(1-w)^{n-k}u^i v^j w^k$$

Then, the deformation of an arbitrary point in the image I can be described as:

$$T(I(x,y,z),\mathcal{B},\Delta) = \sum_{i=0}^{l}\sum_{j=0}^{m}\sum_{k=0}^{n} \mu\left(P_{ijk} + \Delta_{ijk}\right) - I(x,y,z)$$

Unfortunately, (u,v,w) coordinates are not always readily available and must first be found. Given the (x,y,z) coordinates of a point, this can be done by iterative subdivision of the B-spline control point lattice, for example by using one of the many variants of Newton's iteration.

2.2 The Shape Matching Problem

In this paper we will focus on an application of the lattice-based B-spline deformation model to shape matching. Using the definitions given before, this problem can be defined as follows:

Given a source image I_s, a target image I_t, a shape similarity function $\delta(\ )$ and a B-Spline deformation model B, find the set of translation vectors Δ such that $\delta(T(I_s,\mathcal{B},\Delta), I_t)$ is minimal.

In other words, to align two images, coordinates for each control point in the B-spline deformation model must be found such that the resulting deformation minimizes the distance between source and target image. Since there are many strategies to approach this problem, without loss of generality the following assumptions are made in this paper:

1. *Only the source image is deformed*

   In some applications, it may be logical to deform both the fixed and the moving image and end up somewhere in the middle [Yang08], to average the deformation fields of both, or to enforce that both registrations are consistent [Chri01]. Such strategies may be beneficial in cases where one or both of the images contain noise. In general, a deformation of the target image can be approximated as a deformation of the source image through a process of interpolation, but this process may suffer from many off-topic difficulties and for simplicity's sake we therefore assume that deformation only occurs in one direction.

2. *B-spline models are non-hierarchical*

   A common strategy for B-spline driven shape matching uses a multi-resolution approach, in which the resolution of the B-spline control point mesh increases in a step-wise fashion, so that in each step the level of detail of the procedure is increased [Szel97]. By performing global deformations first, and gradually increasing the resolution of the deformation, local deformation becomes less dependent on the global image similarity. However, in each resolution, the formulation of the problem is the same and because our findings can easily be generalized to multiple resolutions, we assume that the shape matching procedure only takes place within a single resolution.

3. *Source and target image have been affinely registered prior to B-spline deformation.*

   In the normal practice of image registration, a rigid or affine transformation is performed before applying a B-spline deformation, to make sure that scale and position do not have an effect on the B-spline deformation and to reduce the risk of getting stuck in a local minimum of the search space. Similar to multi-resolution registration, this improves the local effects of a B-spline deformation. Because a rigid or affine transformation can in principle be expressed as a strictly constrained B-spline deformation, such a registration does not compromise the validity of our deformation model.

2.3 Regularization

Although B-spline driven deformation is a powerful deformation model that is particularly well-suited to performing local deformations, there is also a risk: because the variables of the shape matching problem are the B-spline control points rather than the underlying image, it may happen that a configuration of the B-spline control point network results in a deformation under which the underlying image violates some validity constraints. The most common violation of image validity is folding, which occurs when two or more control points swap their relative position. Rather than checking the validity of the underlying image during the shape matching procedure, often a regularization mechanism or penalty is included in the optimization routine, such that the problem of shape matching can be reformulated as:

Given a source image I_s*, a target image* I_t*, a shape similarity function* $\delta(\ )$*, a B-Spline deformation model* $\mathcal{B}$ *and a regularization function* ρ*, find the set of translation vectors* Δ *such that* $\delta(T(I_s), I_t) + \rho(\Delta)$ *is minimal*

The difficulty here lies in constructing the regularization function and throughout the literature several options have been proposed.

The most common approach is to relate the regularization function to the elastic energy of the B-spline deformation [Rue99][Kybic03] [Sorz05], such that the translation of a single B-spline control point is penalized depending on the magnitude of the translation vector. In [Hart00] it has been shown that folding can be detected using the Jacobian determinant of the deformation field and in [Rohl03] it has been shown that deviations of this Jacobian can be used to obtain an incompressibility constraint that can be used to guarantee a volume-preserving deformation. Staring et al. [Star06] propose a local rigidity constraint that can be used to penalize large deformation in some areas, but less so in other areas.

3. DEFORMATION METHODS

As the deformation that is enacted by the B-spline model depends on the position of its control points, a logical approach to the optimization problem described in section 2 is to treat the coordinates of these control points as variables of the minimization problem. Although this seems straightforward, there are several approaches to this problem, which will be discussed in this section. This section will review two approaches that are used in the area of medical image registration, and will also discuss a parametric approach that the authors of this paper have successfully applied in the area of reverse engineering. In section 4, we will test all three methods in an application to polygonal meshes to see if they lead to comparable results.

3.1 Global Optimization

The most straightforward way of finding a deformation that minimizes the distance between source and target image is a global optimization approach. A global optimization approach treats all *n* B-spline control point coordinates as if they are coordinates of a global vector Δ, which can be defined as $\Delta = \{\delta x_0, \ldots, \delta y_1, \ldots, \delta z_n\}$, where $n = i \cdot j \cdot k$

This vector can be minimized using any known multi-dimensional function optimization technique. Most approaches use a gradient descent (GD) approach [Rue99][Matt03], in which an iterative approach to the minimization of Δ is taken. In each iteration step, a number of points are sampled in the target image and projected backwards onto the source image. The distance between the sample points and their backwards projection is used to calculate the distance between the source and target image and in the next step the direction of further optimization is chosen to be the direction of steepest descent of the distance function. This method is straightforward but computationally costly because it requires a computation of the gradient of the distance function.

Although gradient descent is the most popular global optimization approach, other (gradient-based) methods are also possible. Klein et al. [Klein07] compare different methods, including gradient descent, quasi-Newton, nonlinear conjugate gradient, simultaneous perturbation, evolution strategy, Robbins-Monro and Kiefer-Wolfowitz. Although small differences between the methods can be observed, it is unclear to what extent these differences are application-independent.

3.2 Markov Random Fields

Another approach to the shape matching problem, in which there is no need to compute the gradient, is to formulate it as a discrete Markov Random Field (MRF) objective function [Glo08]. Basically this boils down to treating the deformation vectors for each point as independent, such that the movement of a single control point is evaluated regardless of the movement of its neighbors. Note that under certain constraints, this approach is similar to Powell's direction set approach to global optimization.

Clearly, such a search covers a much larger part of the search space, involving all the risks and benefits that go along with it. In addition, a MRF approach requires a much more thoroughly defined (and, ideally, an adaptive) regularization term: the independent movement of control points increases the risk of folding and violation of other validity constraints. A primary-dual linear programming approach can be used to quickly compute the optimal MRF [Komo07].

3.3 Parametric Models

Both of the previous methods assume that no information is available when translating the B-spline control points. In other words: these methods look for *any* deformation that is optimal given a certain shape similarity measure: it is not guaranteed, however, that this deformation is valid or that it is the least complex way of deforming the source image. Imagine, for example, the alignment of two 2D pictures of faces: in this case from the point of view of the shape similarity function it is valid to deform the source image in such a way that the eyes are swapped while the rest of the model is kept constant. However, in practice this deformation is invalid and unnecessary and therefore an undesirable result of the shape matching procedure.

In this particular example, it makes much more sense to incorporate knowledge of the shapes in question into the shape matching procedure. Models on the way a shape can be deformed are called Active Shape Models (ASM) or Active Appearance Models (AAM) [Cootes95]. In the application of reverse engineering, a similar concept is the freeform feature model [Song05],[Lang07]. In terms of the definitions given earlier in this paper, these models operate by enforcing a set of displacement vectors $\Delta^{\pi} = \left\{\delta^{\pi} x_0,\ldots,\delta^{\pi} y_1,\ldots,\delta^{\pi} z_n\right\}$, such that the shape matching procedure is reduced to finding the set of scalar values $\Pi = \left\{\pi_0,\ldots,\pi_n\right\}$, such that $\Delta = \left\{\pi_0\delta^{\pi} x_0,\ldots,\pi_1\delta^{\pi} y_1,\ldots,\pi_n\delta^{\pi} z_n\right\}$. If Δ^{π} is unknown, then it can be estimated and iteratively maximized using an expectation-maximization approach (as suggested in [Lang08]), but such an estimation is only meaningful if there is indeed a (perceived) parametric difference between source and target image. If this is not the case, the shape matching routine is likely to become over-parameterized. Π can be optimized using global optimization methods or MRF models, but the latter would be difficult to implement because a regularization penalty cannot be directly defined on Π. Song et al. [Song05] use a quasi-Newton approach to match a parameterized shape, or feature, to a target shape. In [Lang07] an evolutionary approach to global optimization is proposed, which is computationally costly but turns out to be successful in traversing the search space in meaningful directions.

3.4 A Comparative Analysis

The main difference between global optimization and Markov Random Fields lies in two aspects: computational complexity and allowed deformation. Global optimization allows less deformation during the shape matching procedure, but is more stable as a result and converges to a solution in a much smoother way. The fact that the global deformation is taken into account during the optimization means that it is slower, especially so when the gradient is used. In an application to the registration of voxel models, the need for a gradient descent approach is justified by the fact that shape similarity measures are often dependent on other aspects than spatial correspondence.

Markov Random Fields methods have the advantage that they explore much more of the search space, and although this also increases the risk of getting stuck in a local minimum, there are plenty of strategies available in the literature to reduce this risk. That deformation is not looked at from a global perspective allows for fast computation, for example using linear programming. MRF optimization has been reported to be up to four times as fast as global optimization.

4. RESULTS

To test the merit of the method proposed in section 2 we have applied both shape matching using parametric models and shape matching using the two reviewed approaches to the registration of voxel models to four polygonal meshes. These meshes have been selected to pose a variety of problems that are commonly encountered in shape matching problems, specifically varying polygon density and complex topology. In addition, the test cases had several characteristics that can be expected to be a challenge in converting them to voxel models, most importantly the aforementioned varying polygon density and connectivity. The four test models are shown in figure 3. Test models a) and b) are industrial parts that have been obtained from the practice of computer-aided design and that pose various challenges: model a) exhibits a large variety in the level of detail while model b) has a complex topology. Models c) and d) were taken from the internet. Model c) has a straightforward topology but a large number of polygons, due to the texture of the skin of the figure. Model d) poses a challenge because of its topology. In addition, some parts of the model are connected only through thin strips, which may form a challenge when converting the model to a voxel model. In table 1, data for all four models is given.

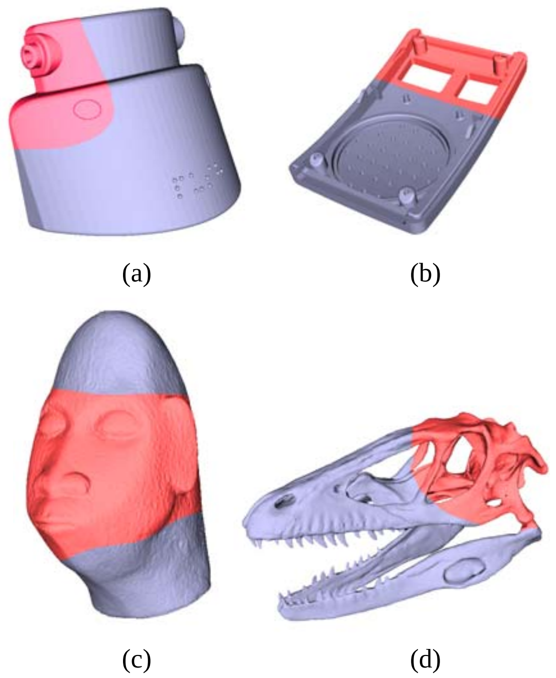

Figure 3: *The four test shapes*

Of each model, a region was selected to serve as the target polygonal mesh (see figure 3). This region was chosen to be smaller than the source image, because this is a requirement of the reviewed image registration methods due to the backwards projection of the sample points used to compute the shape similarity. Table 1 also displays the data for the selected regions, as well as the dimensions of the voxel models that were generated.

Table 1: *Basic data on experiment models*

	Model (a)	*Model (b)*	*Model (c)*	*Model (d)*
Model:				
#polygons	46930	57914	71924	53688
#voxels	53x57x38	112x61x23	62x97x63	72x59x123
Selected Region				
#polygons	6441	2967	48703	6083
#voxels	31x11x33	42x61x23	62x68x63	64x52x48

The original polygonal meshes were chosen as the source images and are denoted as I_s^0 to I_s^3. The selected regions were used as the target images and denoted as I_t^0 to I_t^3.

The proposed method proceeded in the following steps:

1. Noise was added to the target images as follows: a B-spline control point lattice with a resolution of 20x20x20 was constructed around each target image, such that the entire part was contained in the lattice. All B-spline deformation vectors were initialized with a zero displacement vector, to which white Gaussian noise in the domain $\left[-1,1\right]$ was added. The resulting deformation was applied to the polygonal meshes and the displacement vectors were recorded as the set of vectors Δ_{noise}.
2. The source images were registered to the target images using a parametric model that was initialized by the normalized vector set Δ_{noise}, such that it indicated the direction but not the magnitude of the deformation that was performed in step 1. Although in practice Δ_{noise} is not known, in this case we felt that using it was justified because no obvious parameterization is available. In section 5 we will argue why this does not affect our findings.
3. Both source and target images were voxelized using the VoxelModeller algorithm implemented in the Visualisation Toolkit (VTK) Library. The voxelization took less than a minute for models a, b, and d, and a little over two minutes for model c.
4. All voxelized source images were registered to the corresponding voxelized target image using both an MRF approach and a global optimization approach, using a 20x20x20 B-spline control point grid.
5. The transformation that resulted from step 4 was applied to the original polygonal meshes.

For the registrations, the following methods were used:

- For the gradient descent registration of the voxel models, we made use of the open source software library Image Registration Toolkit (ITK).
- For the registration of voxel models using Markov Random Fields, we used the DROP package, as described in [Glo08]
- To register the polygonal meshes we used a reimplementation of the method given in [Lang07].

For the experiments with the voxel models, mutual information was used as a shape similarity measure. For the polygonal meshes, the Mean Directed Hausdorff Distance was used. Values for the regularization penalty were experimentally determined. Unfortunately this had to be done for each test model individually and therefore the values for the regularization penalty could not be compared.

The experiments were done using a computer with four 1.6 GHz processors and 8 GB of RAM memory. Table 2 shows the results of the experiments, where accuracy is measured as the distance in mm. of the source and target image after shape matching.

Table 2: Results of the direct registration

(acc./time)	*GO*	*MRF*	*Par.*
Model (a)	4.8/6min	3.9/7min	3.4/24min
Model (b)	4.5/5min	4.1/6min	3.4/16min
Model (c)	4.8/6min	4.1/6min	3.3/48min
Model (d)	6.2/8min	5.3/8min	2.8/31min

Of the three methods, parameterized registration was the most accurate. This is not surprising, because this method alone made use of additional information regarding the optimal direction of a search for the most optimal deformation. To come to a more fair comparison between the methods, in a post-processing step we applied the parameterized shape matching method to the results of the voxel-based methods until A) the accuracy as given in the rightmost column of table 2 was achieved and B) until convergence (but forced to stop when the time given in table 2 was reached). A) gives an indication of the improvement in computation time that can be achieved and B) gives an impression of the improvement in accuracy that can be achieved. In table 3, the computation times and accuracy including this post-processing step are given. As an example, for model (c), Global Optimization and MRF minimization achieved an accuracy of 3.3 mm. in respectively 25 and 27 minutes, where this took 48 minutes using a parametric approach. In respectively 37 minutes and 44 minutes, an accuracy of 2.9, respectively 3.0 was achieved, compared to the 3.3 using a parametric approach.

Table 3: Results of the post-processing stage

	A		*B*	
	GO	*MRF*	*GO*	*MRF*
Model (a)	3.4/4min	3.4/4min	2.8/12min	2.6/17min
Model (b)	3.4/2min	3.4/3min	3.0/16min	3.1/11min
Model (c)	3.3/19min	3.3/21min	2.9/31min	3.0/38min
Model (d)	2.8/4min	2.8/5min	2.5/19min	2.3/24min

5. CONCLUSION AND DISCUSSION

The purpose of this paper was to demonstrate that voxelizing a polygonal mesh may be beneficial for the registration of the mesh. Voxel images can be registered much quicker than polygonal meshes, because the validity maintenance of a voxel model is much easier than that of a polygonal mesh.

Our results show that for the four presented test models, voxel-based shape matching is indeed much quicker than matching the polygonal meshes directly. In addition, in a post-processing, the results of the voxel-based methods can be used to improve the accuracy of the shape matching. As we said in section 4, making use of the noise vector does not affect our results: without this additional information, parametric shape matching would have performed even worse.

Although the results of these tests are interesting, one must keep in mind that four test models is not enough to reach a strong conclusion on the results of the reviewed methods. In addition, the parametric approach to shape matching has not been optimized to the extent that image registration techniques have been optimized. Optimization of the parametric approach may lead to better and faster results, but this also applies to the post-processing step of which the results are shown in table 3.

Nevertheless, the results indicate that, to obtain a better registration accuracy, it may pay off to first voxelise a polygonal mesh and then apply known registration methods. We intend to more thoroughly investigate this line of research in the future.

6. ACKNOWLEDGEMENTS

Test models a and b were kindly provided by NewProducts. Test models c and d were taken from the internet. Our work made extensive use of the ITK open source library for image registration methods and also of the DROP package, developed by Glocker et al. in München.

7. REFERENCES

[Chri01] Christensen, G.E., Johnson, H.J., Consistent Image registration, IEEE Transactions on Medical Imaging, Vol. 20, No. 7, pp. 568-582, 2001.

[Cootes95] Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J., Active shape models - their training and application. Computer Vision and Image Understanding, Vol. 61, pp. 38-59, 1995.

[Glo08] Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N., Dense Image registration through MRFs and efficient linear programming, Medical Image Analysis, Vol. 12, No. 6, pp. 731-741, 2008.

[Han08] Sass Hansen, M, Larsen, R., Glocker, B., Navab, N., Adaptive Parametrization of

Multivariate B-splines for Image Registration, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1-8, 2008.

[Hart00] Hartkens, T., Hill, D.L.G., Maurer, C.R., Martin, A.J.,Hall, W.A., Hawkes, D.J., Rueckert, D., Truwit, C.L., Quantifying the intraoperative brain deformation using interventional MR imaging, Proc. Int. Soc. Magn. Reson. Med., Vol. 8, 2000.

[Igar99] Igarashi, T., Matsuoka, S., Tanaka, H., *Teddy: A sketching interface for 3D freeform design*, Proceedings of SIGGRAPH, pp. 409-416, 1999.

[Karp06] Karpenko, O.A., Hughes, J.F., *SmoothSketch: 3D free-form shapes from complex sketches*, Transactions on Graphics, Vol. 25, No. 3, pp. 589-598, 2006.

[Klein07] Klein, S., Staring, M., Pluim, J.P.W., Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines, IEEE Transactions on Image Processing, Vol. 16, No. 12, pp. 2879-2890, 2007.

[Komo07] Komodakis, N., Tziritas, G., Approximate labeling via graph-cuts based on linear programming. IEEE Transactions on Pattern Analysis and Machine, Vol. 29, No. 8, pp. 1436 - 1453, 2007.

[Kybic03] Kybic, J., Unser, M., Fast parametric elastic image registration, IEEE Transactions on Image Processing, Vol 12, No.11, pp. 1427-1442, 2003.

[Lang07] Langerak, T.R., Vergeest, J.S.M., An evolutionary strategy for free form feature identification in 3D CAD models, Full Paper Proceedings of the WSCG conference, 2007, Plzen.

[Lang08] Langerak, T.R., Freeform feature recognition and manipulation to support shape design, Ph.D. Thesis, Delft University of Technology, 2008.

[Mac96] MacCracken, R., Joy, K.I., Freeform deformation with lattices of arbitrary topology, Proceedings of the 23rd annual conference on Computer graphics and interactive techniques, pp. 181-188, 1996.

[Mattes03] Mattes, D., Haynor, D.R., Veselle, H, Lewellen, T.K., Eubank, W., PET-CT Image registration in the chest using free-form deformations, IEEE Transactions on Medical Imaging, Vol. 22, No. 1, pp. 120-128, 2003.

[Pern05] Pernot, J.-P., Falcidieno, B., Giannini, F.; Léon, J.-C., Fully free-form deformation features for aesthetic shape design, Journal of Engineering Design, Vol. 16, No. 2, pp. 115-133, 2005.

[Rue99] Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D. L. G., Leach, M. O., Hawkes, D. J., Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images, IEEE Transactions on Medical Imaging, Vol. 18, No. 8, pp. 712-721, 1999.

[Szel97] R. Szeliski and J. Coughlan, Spline-based image registration, Int. J.Comput. Vis., vol. 22, pp. 199–218, 1997.

[Sed86] Sederberg, T.W, Parry, S.R., Free-form deformation of solid geometric models, Computer Graphics, Vol. 20, No. 4, pp. 151-160, 1986.

[Song05] Song, Y., Vergeest, J.S.M., Bronsvoort, W.F., Fitting and manipulating freeform shapes using templates, Journal of computing and information science in engineering, Vol. 5, No. 2, pp. 86-94, 2005.

[Sorz05] Sorzano, C.O.S., Thevenaz, P., Unser, M., Elastic registration of biological images using vector-spline regularization, IEEE Trans Biomed Eng, Vol. 52, No. 4, pp. 652-663, 2005.

[Star06] Staring, M., Klein, S., Pluim, J.P.W., A rigidity penalty term for nonrigid registration Medical Physics, Vol. 34, No. 11, pp. 4098-4108, 2007.

[Tanner00] Tanner, C., Schnabel, J.A., Chung, D., Clarkson, M.J., Rueckert, D., Hill, D.L.G., Hawkes, D.J., Volume and shape preservation of enhancing lesions when applying nonrigid registration to a time series of contrast enhancing MR breast images, Lecture Notes in Computer Science, Vol 1935, pp. 327–337, 2000.

[Rohl03] Rohlfing, T., Maurer, C.R., Bluemke, D.A., M., Jacobs, M.A., Volume-Preserving Nonrigid Registration of MR Breast Images Using Free-Form Deformation With an Incompressibility Constraint, IEEE Transactions on Medical Imaging, Vol. 22, No. 6, pp. 730-741, 2003.

[Ver01] Vergeest, J.S.M., Spanjaard, S., Horváth, I., Jelier, J.J.O., Fitting freeform shape patterns to scanned 3D objects, Journal of Computing and Information Science in Engineering, Vol. 1., No. 3, pp. 218-224, 2001.

[Yang08] Yang, D., Li, H., Low, D.A., Deasy, J.O., El Naqa, I., A fast inverse consistent deformable image registration method based on symmetric optical flow computation, Physics in Medicine and Biology, Vol. 53, pp. 6143-6165, 2008.