Fast Linear Elastic Matching without Boundary Conditions

Samson J. Timoner[1], W. Eric L. Grimson[1], Ron Kikinis[2], and William M. Wells III[1,2]

[1] MIT AI Laboratory, Cambridge MA 02139, USA
samson@ai.mit.edu
[2] Brigham and Women's Hospital, Harvard Medical School, Boston MA, USA

Abstract. A new method is presented for quickly solving non-rigid registration problems that use linear elastic formulations, with no boundary conditions. The algorithm successfully solves ill-conditioned systems of equations formerly thought to be impractical. The running time of the method scales linearly as the number of nodes in the system. We demonstrate a matcher by warping one segmented amygdala onto the amygdalas of 29 other patients with good matching results. The technique required only a few minutes for each match on a desktop computer.

1 Introduction

The non-rigid matching of three dimensional shapes is an enabling method for many medical image analysis applications. Such matching is used to find the correspondence between an anatomical atlas and a volumetric dataset, so that information from the atlas can be transferred to a patient. Non-rigid warping of shapes is also used between two labeled volumes to derive a deformation field that can be used to match corresponding grayscale images [3, 7]. In addition, in order to correlate shape and disease, organs can be non-rigidly aligned enabling machine learning algorithms to classify them based on deformation fields [4].

Because regions of medical image data often contain little useful information, the wide variety of non-rigid warpers must usually constrain the displacement field in some way. Elastic models are a popular choice since they are easy to understand and simulate, and their smoothness properties may be as likely as other constraints. [3, 7, 13, 8].

The past several years have seen considerable activity in the usage of intensity-based mechanisms for solving medical image registration problems [10, 14, 11, 5, 9, 12, 6]. These approaches have the advantage of not requiring segmentation of the image data, and they have proven robust in a number of applications [15].

In light of these issues, it seems clear that grayscale-based matching using elastic deformation models is an attractive approach to solving such problems. There are, however, remaining issues of the computation time and the ill-conditioned nature of the simulation. Perhaps, because of these issues, there have been a number of approaches based on surface models or surface/volume

hybrids. While attractive from a computational point of view, these may be significant compromises because of the different mechanics of shell structures and volumetric structures. For example, it may be simpler to capture the property of thin rods being easier to bend than thick rods with a volumetric mechanical model rather than a surface model.

Based on the above observations, it has been our our primary goal in this project to implement a grayscale-based image matcher based on an elastic deformation model combined with an image agreement term. In this paper we present a fast algorithm that does exactly that. We illustrate our method by aligning 30 different segmented amygdala-hippocampus complexes to each other.

1.1 Related Work

Christensen et al. [2] introduce models based on viscous fluids that have have demonstrated impressive registration among different human brains. These models are far more flexible than models based on elasticity and in some applications may yield unphysical deformations spread over large regions[13].

Ferrant et. al implemented a matcher by first finding correspondences using an active surface algorithm and then using the displacements of the active surface model as boundary conditions for an elastic model of the interior warping[7]. This method is generally limited by errors in the correspondences found in the active surface algorithm.

Davatzikos and Prince examined a method which is closer to a full elastic solver[3]. They establish correspondences between surfaces at a series of points. They then use an elastic force and an image driving force to fill in the rest of the deformation field. Wang and Staib [13] found that constraining the deformation field at a series of points led to a "jiggling effect" near the surface. They therefore modified this method by not forcing the deformation field to be fixed at the surface points, but rather to penalize the deformation field for deviating from the set points.

Our methods are similar to those of Wang and Staib [13] in the use an elastic energy objective function with a statistical image agreement term. However, they note that with these two terms alone, the resulting equations are ill-conditioned and difficult to solve. They add landmark points to improve the conditioning of the equations to make them easier to solve. Our main contribution in the present paper is to develop a linear-time approach that overcomes the ill-conditioning of the problem.

2 Methods

The matching process begins by filling the volume of one shape with tetrahedra. The mesh then aligns itself to the other label volume may by minimizing a combination of an elastic energy term and an image agreement term. The resulting minimization is non-linear and leads to ill-conditioned equations. We describe here a solver that handles these difficulties.

We begin by dividing a 3D shape into a mesh of small volumetric elements. Most packages that divide objects into elements are designed for regularly shaped mechanical parts, not irregular medical organs[6]. Our two main concerns are to produce a mesh that accurately represents a surface with as few skewed tetrahedra as possible. Skewed tetrahedra lead to errors in solving elasticity equations and can halt the progress of a typical solver[1].We have constructed our own mesher which accomplishes these two goals.

We use a probabilistic framework to match two shapes by maximizing the log probability of the deformation field given the data. Using Bayes rule,

$$\arg\max_{\boldsymbol{r}} \log P(\boldsymbol{r}|\text{data}) = \arg\max_{\boldsymbol{r}} \log P(\text{data}|\boldsymbol{r}) + \log P(\boldsymbol{r}) - \log P(\text{data})$$

where $\boldsymbol{r}$ is the deformation field. The probability of the data, $P(\text{data})$, is independent of the warping, so we ignore this term.

To estimate $P(\boldsymbol{r})$, we turn to statistical physics, where the probability that a system is in a configuration is proportional to $e^{-E/(K_B T)}$. We use the linear elastic energy as described in the introduction. The elastic energy of the mesh is given by the integral of the stress strain product over the volume. Through standard means [16], the integral can be linearized about a current configuration. One can then approximate the energy of the system as $\frac{1}{2}\boldsymbol{r}^T K \boldsymbol{r}$ where K is an elasticity matrix. The matrix K is proportional to Young's modulus Y and also depends on Poisson's ratio ν. Note that the matrix K changes each time the configuration of nodes changes.

Our aim is to create a full grayscale nonrigid alignment method, therefore, we use a volumetric data agreement term rather than a voxel-based agreement term. In the present case of label data, we set the center of the voxel values to 0 and 1 according to the label, and interpolate for values elsewhere. We use the average intensity, $\hat{I}_T(\boldsymbol{r})$, inside each tetrahedron as the observed data. Further, we assume that the probability of $\hat{I}_T(\boldsymbol{r})$ for a given tetrahedron is independent of $\hat{I}_T(\boldsymbol{r})$ for all other tetrahedra. Therefore, $\log P(\text{data}|\boldsymbol{r}) = \sum_{\text{Tetrahedra}} \log P(\hat{I}_T(\boldsymbol{r})|l_T)$. where l_T is the label of a tetrahedron and the warping $\boldsymbol{r}$ determines the position of each tetrahedron. We use a simple linear probability so that $P(\hat{I}_T(\boldsymbol{r})|l_T) = \hat{I}_T(\boldsymbol{r})P(1|l_T) + (1 - \hat{I}_T(\boldsymbol{r}))P(0|l_T)$.

$$\arg\max_{\boldsymbol{r}} \frac{-\boldsymbol{r}^T K \boldsymbol{r}}{2K_B T} + \sum_{\text{Tetrahedra}} \log P(\hat{I}_T|l_T). \tag{1}$$

One standard way to maximize a function is to use a Newton based solver to find the zeros in the gradient of the maximization function. Such a solver would find that given the guess $\boldsymbol{r_k}$ in the k^{th} iteration, the change in the deformation field, $\boldsymbol{\delta r}$, in the $k+1^{th}$ iteration solves

$$[\frac{K}{K_B T} - H_P]\boldsymbol{\delta r} = \frac{-K\boldsymbol{r_k}}{K_B T} + \sum_{\text{Tetrahedra}} \frac{d(\log P(\hat{I}_T|l_T))}{d\boldsymbol{r}} \tag{2}$$

where H_P is the Hessian of the image agreement term. We modify the standard Newton solver to stably and consistently find a good maximum. Since we model intensities as being linear in the neighborhood of each tetrahedron, H_P is singular. Worse, it has positive eigenvalues making the maximization unstable. We therefore do not use the Hessian of the image agreement term.

The elasticity matrix K is also singular. For example a rigid translation of the entire mesh causes no elastic energy increase, and therefore a rigid translation is in the null space of the elasticity matrix. One standard method for solving linear systems with singular matrices is to increment the diagonal of the matrix. We therefore add a small constant ϵ multiplied by the identity, I, to the elasticity matrix to improve the conditioning of the matrix. This term essentially penalizes any displacement and is similar to an inertial energy or mass. We use the following iteration:

$$[\frac{K}{K_B T} + \epsilon I]\boldsymbol{\delta r} = \frac{-K\boldsymbol{r_k}}{K_B T} + \sum_{\text{Tetrahedra}} \frac{d(\log P(\hat{I}_T|l_T))}{d\boldsymbol{r}}. \tag{3}$$

This iteration will cease only when the driving forces, the right hand side is zero.

While the "mass" term improves the conditioning of the elasticity matrix, the elasticity matrix is still poorly-conditioned. Solving a poorly-conditioned linear system is not a problem for standard Krylov subspace solvers. Such methods attempt to iteratively minimize the squared error in the system of equations while minimizing the norm of the solution vector. However, because there may be no exact solution, an iterative solver will require many iterations before it converges. We therefore do not try to solve the system of equations to high precision. A standard technique in a Newton solve is to find an approximate solution to the matrix equation (3), using only a few iterations of the iterative matrix solve. Using this technique, the matrix solve time can be improved by more that a factor of 100. And, because there may be no exact solution, little is lost in solving the system approximately.

We make one additional modification to the non-linear maximization to speed the convergence to a final solution. We start the maximization by only considering tetrahedra that were in the first labeled volume. After 90% of the those tetrahedra are overlapping the second labeled shape, we consider tetrahedra outside the labeled volume as well. By having both labeled and unlabeled tetrahedra applying forces, the final convergence rate is increased. We stop the iteration when the norm of the force vector has decreased below a threshold and the objective function has leveled off.

3 Experiment

We report the results of our matcher applied to a the segmented, left hippocampus-amygdala from a data set of 30 patients including 15 normal and 15 schizophrenic. The data set was hand segmented.

We meshed a left amygdala chosen randomly from one of the data sets. We initially aligned the mesh with the other 29 data sets using the second order moments of the mesh and the labeled data. We then ran our matcher 29 times.

3.1 Matching Results

Of the 29 matches, we divide the results into three groups based on a qualitative shape description. The first group includes 13 amygdalas that have relatively smooth surfaces and similar features. The second group includes 14 amygdalas that have small protrusions and sharp discontinuities in their surfaces. The final group includes 2 amygdalas with large protrusions and very different features from the remaining amygdalas.

In the first group, the shapes of the amygdalas are comparable. They have similar bulbous ends, a long curved middle region and a "head" on the other end (see Figure 1). For this group, after the mesh is warped, more than 90% of the voxels inside the mesh are labeled correctly. Figure 1 shows an example match from this set. The initial alignment in the left most image between the mesh (dark) and a label map (light) shows a slight rotation between the two amygdalas. There are width differences in the middle and base of the amygdalas. The "head" portion of the mesh is significantly smaller than the head portion of the label volume. The match shown in Figure 1 has a mesh surface that cuts through the surface of the label map, sometimes being behind that surface and sometimes in front. A cross section through the mesh, (not shown) shows that the surface of the mesh is essentially a smoothed version of the surface of the label map.

The convergence of the algorithm on all matches in this group resembled the convergence shown in Figure 1. There is an initial alignment phase which take roughly 6 iterations. At the end of this phase the great majority of the alignment between the mesh and label map is completed. Generally, between iterations 20 and 30, the forces due to the tetrahedra in the background of the mesh are activated. This causes label tetrahedra previously inside the labeled region to move towards the boundary of the label volume. After roughly 40 iterations, the final shape of the mesh is determined. However, a total of 70-80 iterations are required to maximize the objective function. In these remaining iterations, very small motions in the mesh are observed. All matchings in this group are complete by the 80th iteration and reached a maximum in the objective function between -1440 and -1640.

The second group of images includes 14 matches where a maximum was found in the objective function, but roughly 60-80% of the voxels inside the deformed mesh were labeled correctly. These hand-segmented amygdalas show many small protrusions as well as sharp discontinuities in the surface. The deformed mesh cuts through the surface of the label map, smoothing the bumps and protrusions. These examples reached a maximum in the objective function of between -1620 and -1800, in approximately 80-100 iterations.

The last 2 amygdalas are shaped very differently from the remaining amygdalas. Figure 2 shows one of those two amygdalas. The "head" of the amygdala

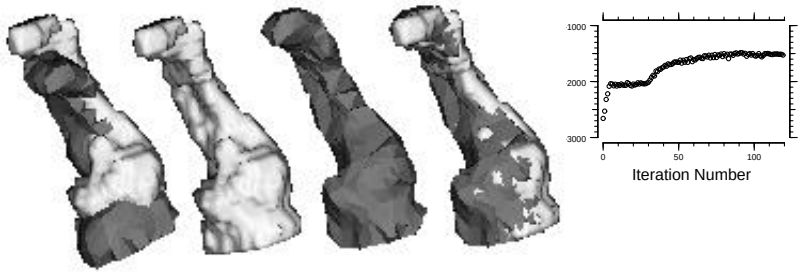

Fig. 1. The left image shows the initial alignment of the mesh (dark) with an unsmoothed triangulated surface of the labeled amygdala (light). The second image shows the triangulated surface alone. The third image shows the result of the non-rigid alignment on the mesh. The right most image shows the result of the warping overlayed on the triangulated surface. The plot shows the objective function versus iteration for a representative simulation in the first group.

does not seem to exist in this segmentation. The middle portion of the amygdala is in the shape of a half-cylinder shell, rather than a solid rod. There are numerous protrusions coming out of the amygdala. The matcher spreads the head-portion of the mesh out along the the semicircular. For these two cases, the algorithm converged to a maximum near -2100 in 80-100 iterations.

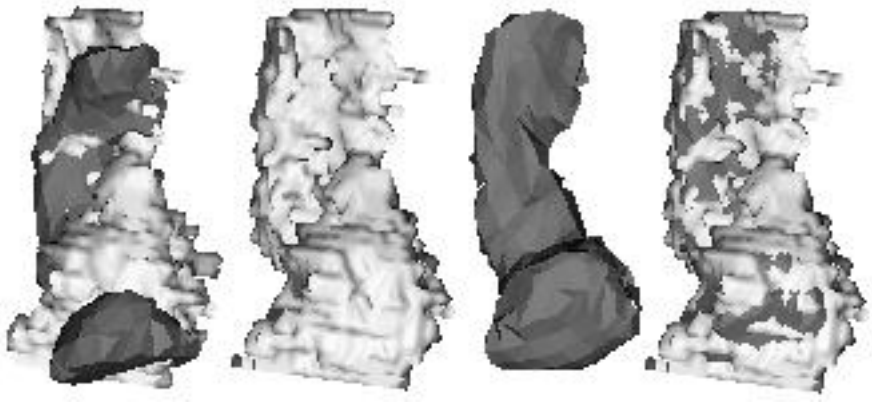

Fig. 2. An example of the matcher in group 3. The layout of this figure is identical to that of Figure 1

3.2 Matcher Performance

For the 1100 node meshes we examined in this paper, the matcher ran converged in 80 iterations in 3 minutes on a 450 Mhz Pentium III. On larger meshes, 80 iterations were still sufficient. For 2400 nodes, the time doubled to 6 minutes.

4 Discussion

Figure 1 illustrates that the head of the meshed amygdala is able slide into the head of the target-amygdala. This type of unconstrained movement is difficult to achieve and in one illustration of the effectiveness of our approach. Given the density of the mesh, the resulting match seems very reasonable.

The second and third group of amygdalas illustrate a fundamental problem with shape matching. While believable matches between smooth medical data is possible by a variety of methods, it is not obvious how a matcher should respond to small protrusions, segmentation errors, etc. Additionally, it is very difficult to find correspondences between data sets separately from finding a global optimization to a warping. In this sense, a method like ours that does not require identification of any correspondences is appropriate for this problem.

It is not clear that the way our method handles protrusions and irregularities, by cutting through them, is appropriate. We deliberately selected a coarse mesh that would not be able to align itself to these protrusions. However, we could have selected a finer mesh that would have deformed itself more closely to the irregular boundaries of many of the amygdalas. In this way the density of the mesh is a parameter that determines the results of the deformation. This is both an asset and a liability of our method.

The algorithm in this paper has shown its ability to solve ill conditioned problems very quickly and give reasonable results. It is fundamentally a linear algorithm, with time scaling like the number of nodes in the mesh. Figure 1 suggests that we may be able to stop the warping process with fewer iterations and thereby further increase the speed of our process.

5 Acknowledgements

S.J. Timoner is supported by the Fannie and John Hertz Foundation. This research is supported by the NSF ERC grant, Johns Hopkins Agreement #8810-274. The authors thank Dr. Martha Shenton for providing us with the segmented images for the amygdala matching experiments.

References

1. M. Berzins. Mesh quality: A function of geometry, error estimates or both? *Engineering With Computers*, 15(3):236–47, 1999.
2. G. Christensen, R. Rabbit, and M. Miller. Deformable templates using large deformation kinematics, 1996.

3. C. Davatzikos and J. Prince. Brain image registration based on curve mapping. *IEEE Workshop Biomedical Image Anal.*, pages 245–254, 1994.
4. C. Davatzikos, M. Vaillant, and S. Resnick et. al. A computerized approach for morphological analysis of the corpus callosum. *Journal Computer Assisted Tomography*, 20(1):88–97, 1996.
5. J. Dengler and M. Schmidt. The dynamic pyramid — a model for motion analysis with controlled continuity. *International Journal of Pattern Recognition and Artificial Intelligence*, 2:275–286, 1988.
6. Matthieu Ferrant, Simon Warfield, Charles R. G. Guttmann, Robert V. Mulkern, Ferenc A. Jolesz, and Ron Kikinis. 3d image matching using finite element based elastic deformation model. In *MICCAI*, pages 202–209, Cambridge,UK, October 1999.
7. Matthieu Ferrant, Simon K. Warfield, Arya Nabavi, Ferenc A. Jolesz, and Ron Kikinis. Registration of 3d intraoperative mr images of the brain using a finite element biomechanical model. In *MICCAI*, pages 19–28, 2000.
8. J. Gee, L. Briquer, C. Barillot, and D. Haynor. Probabilistic matching of brain images. *Information Processing in Medical Imaging*, pages 113–125, 1995.
9. Nobuhiko Hata, Arya Nabavi, Simon Warfield, W.M. Wells, Ron Kikinis, and F.A. Jolesz. A volumetric optical flow method for measurement of brain deformation from intraoperative magnetic resonance images. In *MICCAI*, pages 928–935, Cambridge,UK, October 1999.
10. M. Leventon, W. Grimson, and W. Wells. Multi-modal volume registration using joint intensity distributions. In *Proceedings of the first international conference on Medical Image Computing and Computer Assited Intervention*, 1998.
11. O.S. Skrinjar and J.S. Duncan. Real time 3d brain shift compensation. In *Information Processing in Medical Imaging*, volume 3, pages 42–55, Visegrad,Hungary, July 1999. Springer-Verlag.
12. J. Thirion. Non-rigid matching using demons. In *Proc. Conf. Computer Vision and Pattern Recognition*, pages 245–251, June 1996.
13. Y. Wang and L. H. Staib. Elastic model based non-rigid registration incorporating statistical shape information. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1162–1173, Cambridge, MA, October 1998.
14. W. Wells, P. Viola, S. Nakajima, H. Atsumi, and R. Kikinis. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–53, 1996.
15. J. West, JM. Fitzpatrick, M. Wang, B. Dawant, C. Maurer, R. Kessler, and R. Maciunas. Retrospective intermodality registration techniques: Surface-based versus volume-based. *IEEE Trans. Med. Imaging*, 18:144–150, 1999.
16. O.C. Zenkiewicz and R.L. Taylor. *The Finite Element Method*, volume 1. McGraw-Hill, Berkshire, England, fourth edition, 1989.