Model-Image Registration of Parametric Shape Models

Gavin Baker[1] and Nick Barnes[2]

[1] Department of Computer Science & Software Engineering
The University of Melbourne, VIC 3010 Australia
gavinb@cs.mu.oz.au
[2] National ICT Australia
Locked Bag 8001, Canberra, ACT 2601 Australia

Abstract. Mathematical models are often used to describe natural phenomena, organisms and anatomy. This paper presents a Model-Image Registration framework that can be used to evaluate models of anatomical shape. An optimisation process is applied to the model parameters to fit the model to an organ of interest in a volumetric image. The fit is obtained by maximising a metric that measures how closely the model surface matches the image edge features, according to gradient magnitude and orientation with respect to the surface normal. The system was tested using a spiral shell model, with generated data as ground truth, and also with CT scans of the temporal bone. The parameters converge to within a close tolerance after a few hundred iterations on the test data, and show promising results on registering with the clinical data.

1 Introduction

Anatomical shape models can give us insight into normal and abnormal shape variation. Such models can be used for a variety of purposes: segmentation, by capturing *a priori* information and acting as an initialisation; diagnosis, by providing quantitative measures that can be compared to known normal forms; simulation and training, by enabling clinicians to generate a variety of plausible forms; and morphometry, by providing a direct means to measure organ shape.

Deformable models such as Level Sets [1], Snakes [2] and Fourier boundary models [3], have proven effective for the segmentation of anatomical objects. However the result of such a segmentation process does not readily yield clinically relevant metrics. For example, the coefficients of a Fourier boundary model are not clinically meaningful, and arise from the mathematical representation rather than that organism.

The cochlea is the shell-like organ of hearing in the middle ear. Its shape is frequently described as resembling a shell or a logarithmic spiral [4]. The cochlea's spiral shape and its lack of distinct landmarks makes it difficult to employ standard morphometric techniques for measurement and shape analysis. A shape model of the cochlea, incorporating statistical variation, would have many clinical applications. For example, to improve cochlear implant design, aid in diagnosis, and generate models for virtual surgery.

While detailed studies have been carried out on the morphology of the bony labyrinth (incorporating the cochlea) [5], the metrics are typically linear measurements or relative orientations — for example, the width of the basal turn of the cochlea, or the angle of the lateral semicircular canal in the saggital plane, all of which are most relevant to phylogenic studies. Most of the cochlea models described are spacecurves that either don't take into account cross-sectional shape [6, 4], or use an approximation.

Other metrics are far more clinically relevant for the purposes of planning for cochlear implant surgery, such as the average diameter of the basal turn or the curvature of the centreline of the otic capsule. For a clinical model to be truly useful, it should incorporate or yield metrics that are directly relevant in the clinical domain [7].

Viable shape models can be generated from parametric functions in geometric space (as in the initial experiments described herein). Alternatively, a parametric model derived from Generalised Procrustes Analysis [8] could be used, where the parameters correspond to the eigenmodes that represent the greatest shape variation.

We therefore need a means to evaluate different parametric models, to compare their performance, accuracy and generalisability on a series of datasets. This paper presents a such system, developed with the Insight Toolkit [9].

2 Method

We have developed a Model-Image Registration framework, as illustrated in Figure 1. The goal is to register a model M to an object within a volumetric image I. The optimiser varies the model parameters $\boldsymbol{p}$ to maximize the fitness function F.

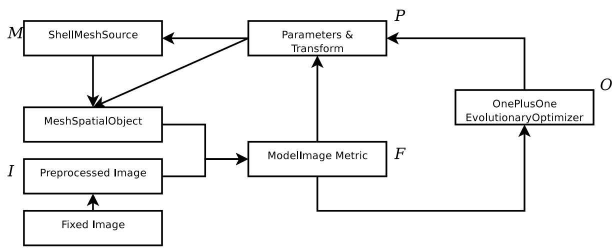

Fig. 1. Model-Image Registration framework

2.1 Parametric Model

The first model investigated is the 3D parametric form of a spiral seashell as described by Nordstrand [10]. This model was chosen as it has low number of parameters, and produces a surface that resembles the cochlea. The Cartesian form is:

$$x(\theta, \phi) = a\left[1 - \frac{\phi}{2\pi}\right]\cos(N\phi)\,(1 + \cos\theta) + c\cos(N\phi) \tag{1}$$

$$y(\theta, \phi) = a\left[1 - \frac{\phi}{2\pi}\right]\sin(N\phi)\,(1 + \cos\theta) + c\sin(N\phi) \tag{2}$$

$$z(\theta, \phi) = b\frac{\phi}{2\pi} + a\sin\theta\left[1 - \frac{\phi}{2\pi}\right] \tag{3}$$

where a controls the diameter, b controls the vertical growth rate, c determines the radial expansion, and N is the number of turns. These constitute the model parameters; ϕ and θ are swept through the range $[0, 2\pi]$.

A Euclidean 3D transform is used to apply translation and rotation to the model. A full affine transform is not required, as scaling is already effectively provided by the model parameters. The parameter vector $\boldsymbol{p}$ is thus:

$$\boldsymbol{p}^T = \left[\, a\ b\ c\ t_x\ t_y\ t_z\ r_x\ r_y\ r_z \,\right] \tag{4}$$

The operator manually provides an approximate initialisation of all the model parameters $\boldsymbol{p}$, visually matching the model to the image data. A mesh is then generated that represents the surface of the model according to the initial model parameters. The mesh is a quadrilateral mesh of cylindrical topology, and the triangular tessellation includes the centroid and surface normal at each cell.

The volumetric data is preprocessed in order to maximise the metric at object boundaries. A Perona-Malik anisotropic diffusion filter [] is applied, which smooths the internal image data while preserving edge contrast. This is followed by a Sigmoid filter to normalise the image scale. Finally the image gradient $\nabla\mathcal{I}$ is computed with a recursive Gaussian image filter.

The registration metric is designed to maximise the fit of the surface to the object boundary. It is similar to the external deformation force for the simplex mesh described by Delignette[11]. Each cell of the mesh is evaluated to measure how well it fits the local image features. The first component of the fitness is taken from the magnitude of the image gradient at the cell centroid. The second component measures how well the surface orientation aligns with the image gradient. Thus:

$$\mathcal{F} = \alpha\|\nabla\mathcal{I}\| + \beta\langle\nabla\mathcal{I}, \hat{\boldsymbol{n}}\rangle^2 \tag{5}$$

where α and β are independent weights, ∇I is the image gradient and $\hat{\boldsymbol{n}}$ is the surface normal on the mesh.

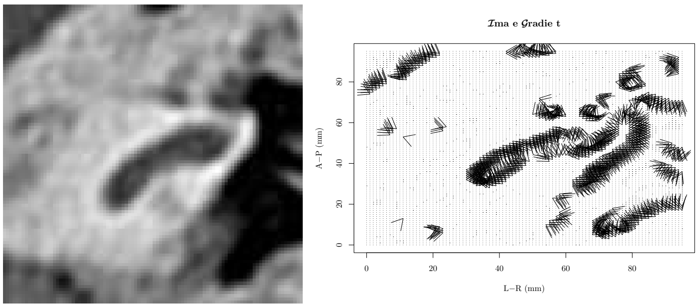

(a) (b)

Fig. 2. (a) A sample CT slice showing a turn of the cochlea, (b) The image gradient

The One Plus One Evolutionary Optimizer [12] was used to separately optimise the rotation, translation and model parameters in turn. It was found experimentally that this 3-stage approach was necessary to avoid the interaction effects of simultaneously varying two types of parameters.

This system was implemented using the Insight Toolkit (ITK) [9]. The Model-Image Registration framework in ITK was extended to permit optimising a general set of parameters, rather than a particular type of transform. This abstraction is a generalisation of the existing approach. The `MeshSource` was extended to create a class that takes an arbitrary set of parameters, permitting any type of model to be plugged into the system. The triangular mesh cell was extended to incorporate the centroid and normal, and a mesh visitor was developed to recalculate this information.

3 Results

To validate the approach, a seashell model with known parameters was rasterised into a low-resolution (96x96x96) volume as ground-truth. This image was then loaded, and the model parameters and transform set by visual inspection. The results of the registration process are shown in Figure 3. The parameters were optimised to within 5% for A and C, and around 8% for B after 350 iterations.

For the clinical experiments, a CT scan of a human temporal bone was used, cropped to the region about the cochlea (volume size 71x84x51). The model's initial transform was manually set to align the position and orientation as close as possible to the otic capsule, while remaining slightly smaller. After 500 iterations, the result was visually inspected to evaluate the fit to the outer wall. The shape matched well in the first turn, but did not match the curvature of the basal turn, which opens out more than the rest of the spiral. These results are shown in Figure 4 in the 3 standard projections.

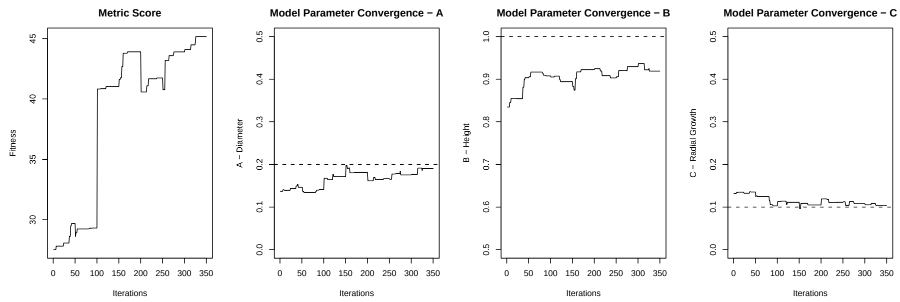

Fig. 3. Validation using a generated image, showing the overall metric score and convergence of the 3 model parameters. Dashed line indicates ground-truth value.

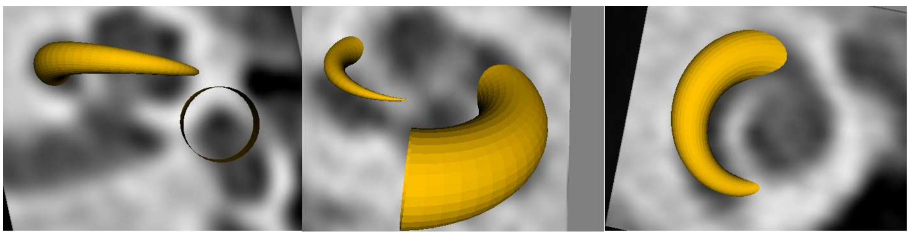

Fig. 4. Registering shell model to human Cochlea, showing axial, saggital and coronal projections.

4 Discussion & Conclusion

The geometry of a generated shape model is, by its very nature, highly constrained, thus limiting its ability to conform to a wide variety of samples. However, the results above show real promise for evaluating parametric models. The approach can be useful as a first approximation, or as a good initialisation to a fully deformable model.

The performance is sensitive to initialisation. However, with practice it is possible to effectively initialise the model within the object of interest. A repeated sequence of optimising the translation, rotation and then model parameters was found to be the most effective, rather than optimizing all the parameters at once.

The seashell model examined above is parameterised in terms that can be related to clinically relevant metrics. The results on the cochlea registration show promise as a first-order approximation, despite not having the flexibility to fit the basal turn[3].

We have demonstrated a framework that enables the evaluation of different types of parametric shape models, and examined its performance on generated and clinical data. We also demonstrated the application of this technique to

[3] Cohen's 2D logarithmic spiral [6] features a piecewise model specifically to take into account the change in curvature of the basal turn.

fitting a simple model to the CT scan of a cochlea. Model-Image Registration through direct optimisation of model parameters is viable, and enables the use of models with clinically relevant parameters.

The next phase of this work will take the resulting model from this process and compare it to an unconstrained deformable model. The deformation field between the two meshes will not only provide a measure of the accuracy, but also indicate the primary modes of deformations the model is incapable of matching, given its geometry. This information can in turn be used to improve the model.

Another logical extension of this work is to evaluate the performance of various optimizers. Indeed, a closer analysis of optimizer performance may reveal the need for an optimizaiton strategy more closely attuned to this particular type of registration.

The source code, test data, movies and scripts are available for download from the author's website at `http://www.cs.mu.oz.au/~ gavinb/`.

References

1. Malladi, R., Sethian, J., Vemuri, B.: Shape modeling with front propagation: a level set approach. IEEE Transactions on Pattern Analysis and Machine Intelligence **17** (1995) 158–175
2. Kass, M., Witkin, A., Terzopolous, D.: Snakes: Active contour models. International Journal of Computer Vision **1** (1988) 321–331
3. Staib, L., Duncan, J.: Boundary finding with parametrically deformable models. IEEE Transactions on Pattern Analysis and Machine Intelligence **14** (1992)
4. Yoo, S., Rubinstein, J., Vannier, M.: Three-dimensional geometric modeling of the cochlea using helico-spiral approximation. IEEE Transactions on Biomedical Engineering **47** (2000) 1392–1402
5. Spoor, F., Zonnenveld, F.: Comparative review of the human bony labyrinth. Yearbook of Physical Anthropology **41** (1998) 211–251
6. Cohen, L., Xu, J., Xu, S., Clark, G.: Improved and simplified methods for specifying positions of the electrode bands of a cochlear implant array. American Journal of Otology **17** (1996) 859–865
7. Baker, G., O'Leary, S., Barnes, N., Kazmierczak, E.: Cochlea modelling: Clinical challenges and tubular extraction. In: 17th Australian Joint Conference on Artificial Intelligence, Cairns, Australia (2004)
8. Dryden, I., Mardia, K.: 4. In: Statistical Shape Analysis. John Wiley & Sons (1998)
9. Ibáñez, L., Schroeder, W., Ng, L., Cates, J.: 1. In: The ITK Software Guide. Kitware Inc (2003)
10. Weisstein, E.W.: Seashell from mathworld – a wolfram web resource. http://mathworld.wolfram.com/Seashell.html (2005)
11. Delignette, H.: General object reconstruction based on simplex meshes. International Journal of Computer Vision **32** (1999) 111–142
12. Styner, M., Gerig, G., Brechbuehler, C., Szekely, G.: Parametric estimate of intensity inhomogeneities applied to MRI. IEEE Transactions On Medical Imaging **19** (2000) 153–165