BAYESIAN REGISTRATION OF MODELS USING FEM EIGENMODES

M.H-M. Syn and R.W. Prager

CUED/F-INFENG/TR 213

September 1995

Cambridge University Engineering Department
Trumpington Street
Cambridge CB2 1PZ
England

E-mail: `mhs@eng.cam.ac.uk, rwp@eng.cam.ac.uk`

Abstract

Highest Confidence First (HCF) estimation is applied to deterministic scaled-ordered non-rigid registration of organ models. A local Posterior energy measure is computed from Bayesian combination of local Prior and Likelihood energy measures, over a Markov Random Field (MRF) defined over the Finite Element neighbourhood of every element node. Prior energy is derived from the Gompertz metric of biological growth, and Likelihood energy is derived from the biologically meaningful similarity between local FEM eigenmode displacement components. The Centroid Size metric is generalised to give the characteristic scale of an organ model, which allows for normalisation of model size and eigenmode magnitude. Linear axes along which modal moments act are used as an estimate of intrinsic model pose, so that initial rigid-body registration can be achieved.

Contents

Chapter 1

Introduction

In this report we describe a deterministic algorithm for computing a map of homologous sites between two organ models. This algorithm uses Bayesian combination of evidence to integrate a local metric of biological growth with a metric of the likelihood of biological homology.

The likelihood features are FEM eigenmodes, which have a number of useful properties when used as scale-ordered shape features (Syn & Prager 1995*b*). Linear axes for both models are used to approximate eigenmodes of fundamental spatial frequency, allowing pose estimation for an initial rigid-body registration. FEM eigenmodes are then used in ascending order of spatial frequency in order to refine the homology map.

Such a non-rigid registration procedure has a number of applications in a 3D ultrasound imaging system

- register models with different spatial sampling eg. comparing an organ imaged using 3D ultrasound and MR;
- resampling a warped mesh resulting from a 3D segmentation;
- automatic registration of training models can be use for building a statistical description of shape variation;
- automatic registration of segmented organs to examine which points are vary least and are therefore best suited for use as landmarks;
- a speculative application is the registration of a model to unordered boundary elements derived from 3D ultrasound images using texture, intensity and inter-frame velocity edges (cf. Sclaroff & Pentland (1995)).

1.1 Prospectus

Size metric

We generalise the Centroid Size metric (Bookstein 1991) to apply over the triangular surface faces of an organ model so that models can be size-normalised before registration.

FEM eigenmodes

We review the properties of FEM eigenmodes which are not only modes of mass flux in biological growth, but also features of biological homology. Eigenmodes are only determined up to a scale-factor, and magnitude can be normalised using the size metric derived above, whereas direction normalisation is handled using a relatively *ad hoc* scheme.

Pose estimation for initial registration

The first three eigenmodes are of fundamental spatial frequency, and the modal moment vector (which is usually non-zero due to sampling and numerical roundoff) is a convenient linear axis which describes a major plane of symmetry. We use modal moment vectors to construct a trio of *principal modal axes* which define the model's intrinsic pose, and hence can be used to estimate a rigid-body registration.

Bayesian registration

We apply Highest Confidence First (HCF) estimation to local Posterior energies at each model site, in computing a homology map which defines the non-rigid registration of two models. The local Posterior energy is a Bayesian combination of a Prior energy derived from the Gompertz metric of biological growth, and a Likelihood energy derived from the similarity between local eigenmode components. HCF estimation is performed using eigenmodes in increasing order of spatial resolution, so that a graduated improvement in localisation of the homology map is achieved.

Chapter 2

Size metric

2.1 Covariances between size and shape

Bookstein (1991, Ch. 5.5) presents the Centroid Size metric as the only size variable invariant to independent, identical and isotropic noise in landmark specification. There are a number of equivalent definitions of Centroid Size, one of which is as the summation of all inter-landmark squared-distances.

We generalise this to operate over a 3D model consisting of triangular surface facets. The model interior is ignored since it adds nothing to our shape measure - all internal displacements are determined from surface landmarks and displacements in our applications.

2.2 Coordinates in a triangle

Figure 2.1 shows a triangle with node positions $(\mathbf{a}, \mathbf{b}, \mathbf{c})$.

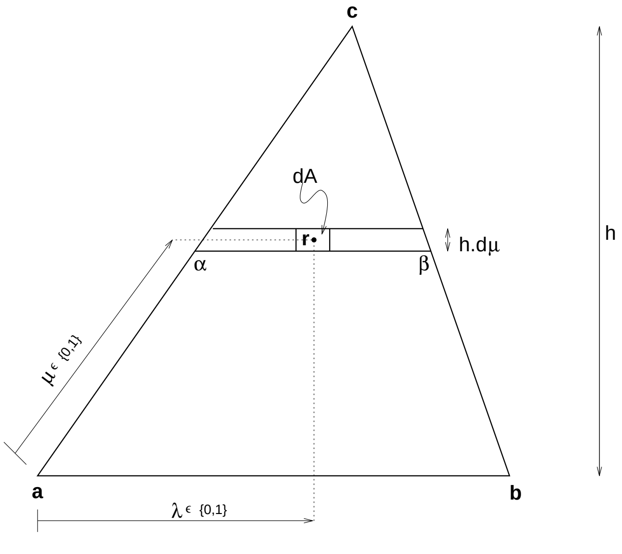

Figure 2.1: Coordinate parameters of triangle in 3D space

Consider a small strip $(\boldsymbol{\alpha},\boldsymbol{\beta})$ parallel to the baseline $(\mathbf{a},\mathbf{b})$. Let $\mathbf{r}$ denote a position in this strip.

$$\begin{aligned}
\boldsymbol{\alpha} &= \mathbf{a}+\mu(\mathbf{c}-\mathbf{a}) \quad (2.1)\\
\boldsymbol{\beta} &= \mathbf{b}+\lambda(\mathbf{c}-\mathbf{b}) \quad (2.2)\\
\mathbf{r} &= \boldsymbol{\alpha}+\lambda(\boldsymbol{\beta}-\boldsymbol{\alpha}) \quad (2.3)\\
dA &= hd\mu.\|\boldsymbol{\beta}-\boldsymbol{\alpha}\|d\lambda \quad (2.4)\\
\boldsymbol{\beta}-\boldsymbol{\alpha} &= (\mathbf{b}-\mathbf{a})(1-\mu) \quad (2.5)
\end{aligned}$$

We can find the height of the triangle from the vector $(\mathbf{c}-\mathbf{a})$ by subtracting the component along the triangle's baseline.

$$h=\left\|(\mathbf{c}-\mathbf{a})-\left(\frac{\mathbf{b}-\mathbf{a}}{\|\mathbf{b}-\mathbf{a}\|}\cdot(\mathbf{c}-\mathbf{a})\right)\frac{\mathbf{b}-\mathbf{a}}{\|\mathbf{b}-\mathbf{a}\|}\right\| \tag{2.6}$$

2.3 Centroid Size integrated over two triangles

Let us extend Bookstein's metric to operate between two triangles T_1 and T_2, respectively bounded by the nodal coordinates $(\mathbf{a},\mathbf{b},\mathbf{c})$ and $(\mathbf{d},\mathbf{e},\mathbf{f})$.

If $\mathbf{r}$ is a point within T_1 and $\mathbf{s}$ is a point within T_2, then the Centroid Size metric is the double integral

$$\begin{aligned}
\int_{T_1}\int_{T_2}\|\mathbf{r}-\mathbf{s}\|^2 dA_1.dA_2 &= \int_0^1\int_0^1\int_0^1\int_0^1 (h_1\|\beta_1-\alpha_1\|)(h_2\|\beta_2-\alpha_2\|)\|\mathbf{r}-\mathbf{s}\|^2.d\lambda_1 d\mu_1.d\lambda_2 d\mu_2\\
&= h_1h_2\|\mathbf{a}-\mathbf{b}\|\|\mathbf{d}-\mathbf{e}\|\left(\frac{3k_1+3k_2-4k_3}{4\times 18}\right) \quad (2.7)
\end{aligned}$$

where

$$\begin{aligned}
k_1 &= \|\mathbf{a}\|^2+\|\mathbf{b}\|^2+\|\mathbf{c}\|^2+\|\mathbf{d}\|^2+\|\mathbf{e}\|^2+\|\mathbf{f}\|^2 \quad (2.8)\\
k_2 &= \mathbf{a}.\mathbf{b}+\mathbf{b}.\mathbf{c}+\mathbf{c}.\mathbf{a}+\mathbf{d}.\mathbf{e}+\mathbf{e}.\mathbf{f}+\mathbf{f}.\mathbf{d} \quad (2.9)\\
k_3 &= \mathbf{a}.\mathbf{d}+\mathbf{a}.\mathbf{e}+\mathbf{a}.\mathbf{f}+\mathbf{b}.\mathbf{d}+\mathbf{b}.\mathbf{e}+\mathbf{b}.\mathbf{f}+\mathbf{c}.\mathbf{d}+\mathbf{c}.\mathbf{e}+\mathbf{c}.\mathbf{f} \quad (2.10)
\end{aligned}$$

Summed over all N triangle pairs (T_i,T_j) on the model surface, the Centroid Size metric becomes

$$cs=\frac{1}{N}\left(\sum_{i=1}^{N}\sum_{j\geq i}\int_{T_i}\int_{T_j}\|\mathbf{r}-\mathbf{s}\|^2 dA_1.dA_2\right)^{\frac{1}{2}} \tag{2.11}$$

Equation 2.7 was algebraically integrated by computer and has some expected properties, being

1. symmetric under permutations of $(\mathbf{a},\mathbf{b},\mathbf{c})$ or $(\mathbf{d},\mathbf{e},\mathbf{f})$ and their components.

2. symmetric under swapping $(\mathbf{a},\mathbf{b},\mathbf{c})$ with $(\mathbf{d},\mathbf{e},\mathbf{f})$.

2.4 Scale-normalised liver models

Figure 2.2 shows twelve volumetric liver models of varying size and sampling which have been scale-normalised using the Centroid Size metric. These models were sampled and segmented from MR scans of fetal cadaveric livers (Pasapula 1995).

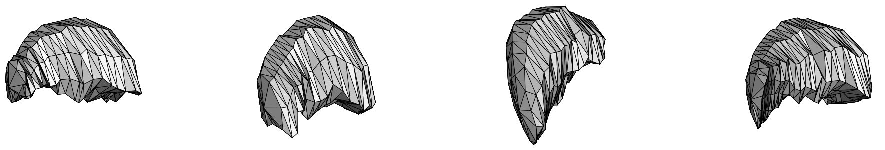

(a) Liver 1 (b) Liver 2 (c) Liver 3 (d) Liver 4

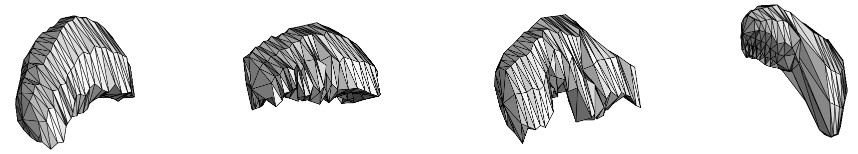

(e) Liver 5 (f) Liver 6 (g) Liver 7 (h) Liver 8

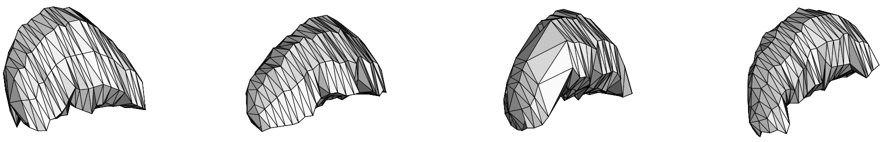

(i) Liver 9 (j) Liver 10 (k) Liver 11 (l) Liver 12

Figure 2.2: Scale-normalised volumetric liver models

Chapter 3

FEM eigenmodes

3.1 FEM eigenmodes as shape features

Sclaroff & Pentland (1995) describes the use of 2D eigenmodes in *modal matching* of images in a video database, leading to applications in object recognition and classification. A connectivity model is built using Gaussian interpolation of unordered edge elements from an edge finder. An *affinity table* is compiled of the strongest correspondences between two objects which should have only a small rotation offset, from which the overall homology map is derived. We have applied a rotation-invariant extension of this algorithm (Syn & Prager 1994) with some success to an artificial dataset of 3D ellipsoidal models, but it fails with more complex 3D models such as those shown in Figure 2.2.

3.2 Eigenmode symmetries

Eigenmodes are ordered in descending order of spatial scale, and when comparing between two similar models are also ranked in descending order of similarity, since the two models have more shape differences at smaller spatial scales.

Elastic eigenmodes exploit local structural weaknesses of models, which at large spatial frequencies expose axes of global symmetry. Again these axes of symmetry are scale-ordered since the modal energy of these structural weaknesses are dependant on spatial scale.

In fact the first three eigenmodes for a model in 3D will always be of fundamental spatial frequency, so that there is only one spatial phase change over the surface. These fundamental modes act about the *principal modal axes* which are used in Chapter 4 to determine model pose.

3.2.1 Gallbladder models

Figure 3.1 shows two gallbladder models in the same pose. The first three non-rigid eigenmodes are the same, and there is a *single* linear bending axis about which each mode bends or twists, since each mode is of fundamental spatial frequency. Figures 4.1–4.3 show the bending axes more clearly.

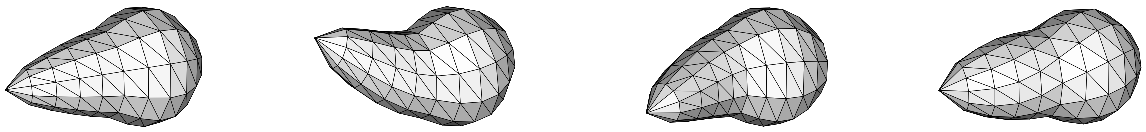

(a) Gallbladder (b) First mode (c) Second mode (d) Third mode

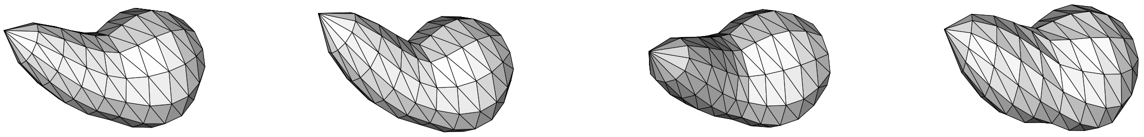

(e) Bent gallbladder (f) First mode (g) Second mode (h) Third mode

Figure 3.1: First three non-rigid eigenmodes of two volumetric gallbladder models (in the same pose)

3.2.2 Liver models

Figure 3.2 shows the first three non-rigid eigenmodes of the Liver1 and Liver5 models. Arrows indicate the gross modal action, from which the bending axis for each mode can be deduced.

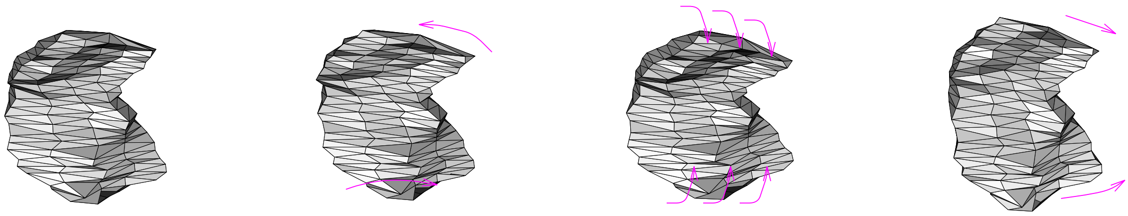

(a) Liver1 (b) First mode (c) Second mode (d) Third mode

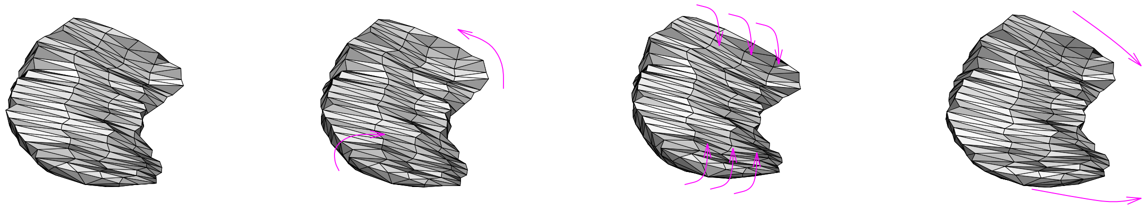

(e) Liver5 (f) First mode (g) Second mode (h) Third mode

Figure 3.2: First three eigenmodes of liver models (Liver1 and Liver5)

3.3 Eigenmode similarity

Syn & Prager (1995*a*) develop a model of biological growth which, given a number of simplifying assumptions, shows that eigenmodes describe the directions of mass transfer during growth. We have already argued above that for two similar models, their corresponding eigenmodes are ranked in order of similarity because the models have accentuated differences at smaller spatial scales (ie. fine details are more likely to differ).

When comparing two biological models, we may treat them as noisy instantiations of some prototype. Such noise is manifested in all stages of biological morphogenesis and growth[1], in genetic shape and control parameters as well as in the imaging and shape-

---

[1] Syn & Prager (1995*a*) review current work on the robustness of biological morphogenesis in the presence

modelling process. When the models are very similar to each other (and hence to the prototype), then the modes of growth of the prototype are also very similar to the modes of the two models for the reason mentioned above. Additionally the shape difference between model and prototype lies within the linearisation limit of more eigenmodes[2].

Since each position on the prototype grows into a point on both models, there is a biological correspondence or *homology* between all points on the eigenmodes of two biological models. In establishing corresponding points between two models, we can therefore compare the eigenmode components at these points to give the likelihood that these points are homologous.

Figure 3.3 illustrates the situation when both models are similar to each other and to the prototype.

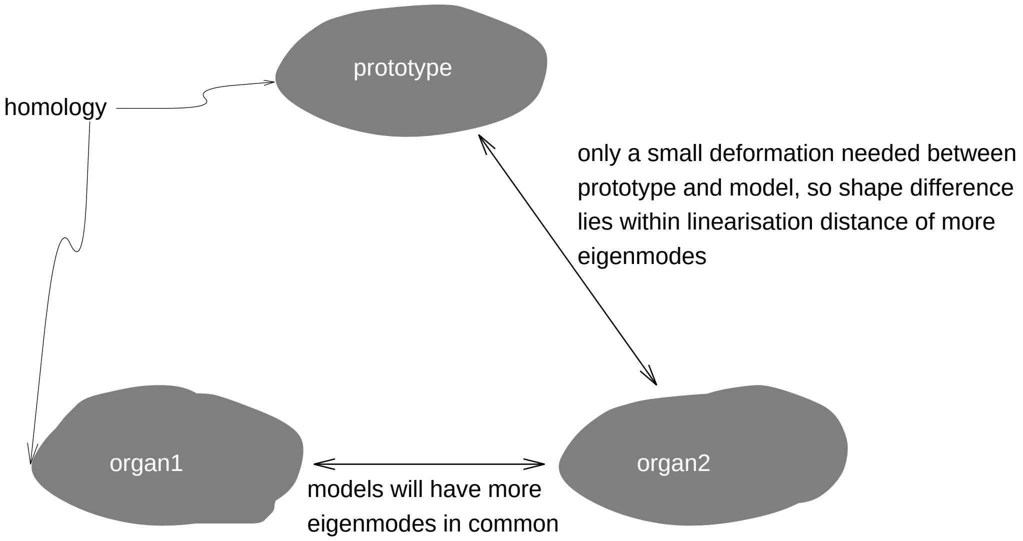

Figure 3.3: Organ models are noisy instantiations of a prototype

A related issue is that of choosing biological landmarks. The more variable the position of a point on an organ, when compared to other organs, the more statistical power it has in describing the current instantiation of organ shape. Sites with large eigenmode displacement components are ones which vary most between organs, according to the model of biological growth presented in (Syn & Prager 1995*a*), and will therefore be most significant in determining the likelihood of homology between models. This accords with our use of FEM eigenmodes as *shape features*.

3.4 Mode normalisation

In order for eigenmodes to be used as shape features for comparison and registration between two models, they must first be normalised in magnitude and direction, since they are

---

of noise

[2] Eigenmodes of smaller spatial scale will have smaller linearisation limits.

arbitrarily scaled when derived as eigenvectors. The models being compared are sampled with arbitrary density and distribution, so that the standard method of scaling the 2-norm is inadequate. Normalising the mean displacement component at surface sample points is also inadequate because eigenmodes components have near-zero net displacement (Syn & Prager 1995*b*, Ch. 3).

We use a normalisation scheme which scales the largest single component at any sample point to a fraction (normally 10%) of the characteristic scale of the model, which represents the linearisation limit of the eigenmode. An additional constraint is that the net volume change associated with the desired sign of each non-rigid eigenmode is the largest positive one.

Eigenmodes with large spatial frequency are used as shape features, and they are discretely sampled vectors of a displacement *field* which has only a very small number (1–3) of phase changes and local displacement maxima over the model surface. Even with noise in spatial sampling and mode computation, these maxima will be spatially large and vary smoothly, which is why the normalisation used above is fairly robust.

Using this normalisation scheme, we avoid the circular problem of normalising eigenmode direction and magnitude by first having to determine initial pose and correspondence.

3.4.1 Volume change in symmetric eigenmodes

The sign of volume change is a macroscopic characteristic of eigenmode direction, which we have found to be robust enough to be used for direction normalisation. It is possible when comparing two noisy models for a sparsely sampled eigenmode to give volume changes of opposite signs. We have found this to be the case for higher eigenmodes of some of the liver models, and this is because of the sparse sampling of the model which falls below the minimum Nyquist frequency needed to specify both model shape and eigenmode features.

Note that non-rigid eigenmodes always cause a change in model volume. Displacement fields which do not are by definition *rigid-body displacements.* Note also that symmetric eigenmodes acting on symmetric models will give the same volume change for positive and negative eigenmode displacement.

3.4.2 Normalising higher eigenmodes

The standard FEM constraint is to ensure that the largest single component has a positive x-direction (Hitchens 1992, Ch. 8), but this is pointless when the initial pose is undetermined. However after initial pose has been estimated using the principal modal axes, this constraint can be used to normalise the direction of higher eigenmodes.

Chapter 4

Pose estimation for initial registration

We have shown that FEM eigenmodes are shape features with the desirable properties of scale-ordering, similarity-ordering and global shape support. Sclaroff & Pentland (1995) regard eigenmode features as generalised axes of symmetry, and in this chapter we utilise this property to perform pose estimation for initial rigid-body registration of structures which are to be finely registered in Chapter 5

4.1 Eigenmode moments

We have seen in Section 3.2 that the net moment of the eigenmode field about the centroid must be zero. If we compute the moment using only the boundary components of the eigenmode field, which have the largest magnitude and are perceptually most significant, then there will be a net moment vector for models such as those shown in Figure 3.1

4.1.1 Gallbladder model

Figures 4.1–4.3 show superposed views of each eigenmode, together with surface normals which are used to highlight local displacement components. The approximate position is also indicated of the modal moment axis about which the components of each fundamental mode acts.

4.1.2 Liver model

Figures 4.4–4.6 shows the surface normals and bending axes for the Liver9 model.

4.2 Principal modal axes

The modal moment axes indicated above are linearised approximations to the line of zero phase change in each eigenmode of fundamental spatial frequency. A mutually-orthogonal set of three axes can be computed as an estimate of model pose.

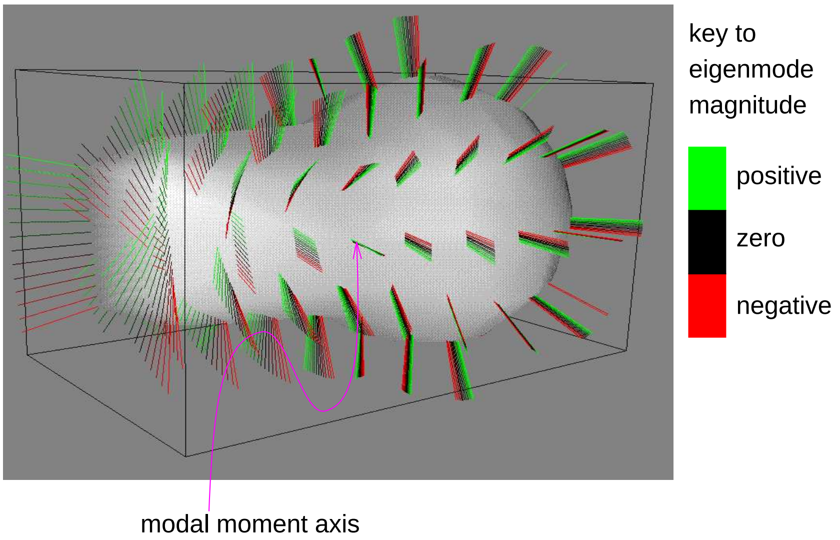

Figure 4.1: Modal moment axis for first eigenmode of gallbladder model

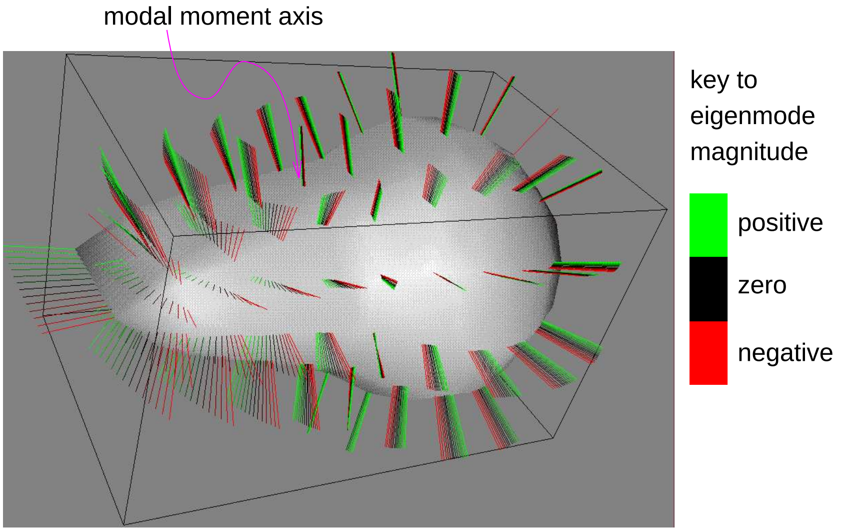

Figure 4.2: Modal moment axis for second eigenmode of gallbladder model

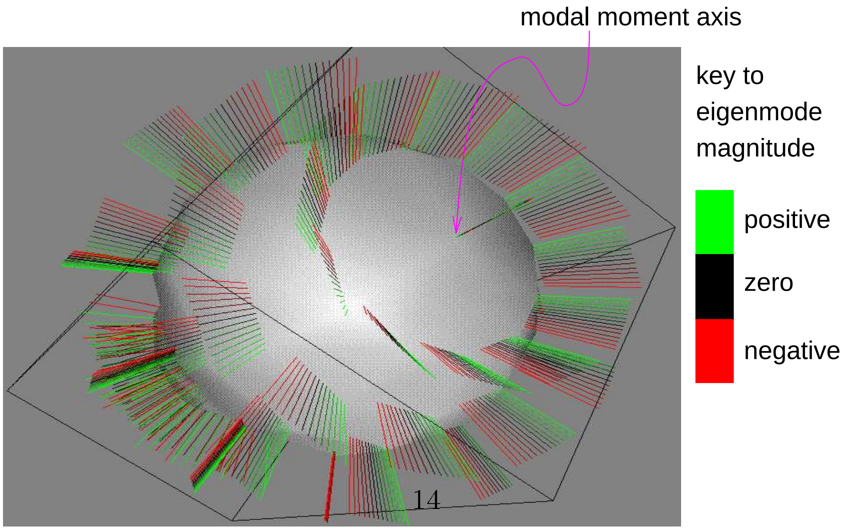

Figure 4.3: Modal moment axis for third eigenmode of gallbladder model

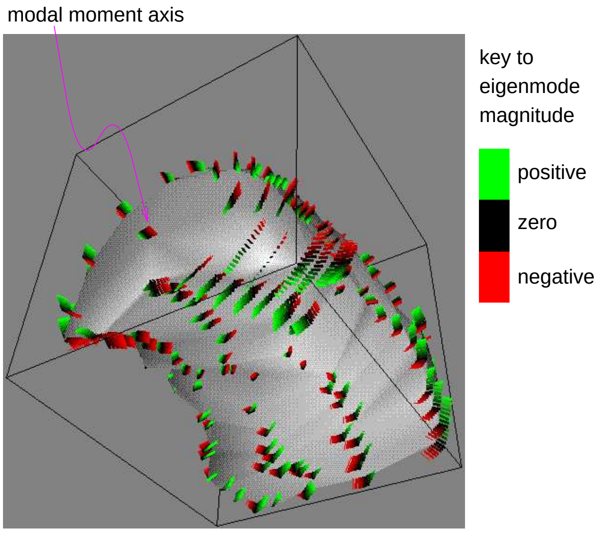

Figure 4.4: Modal moment axis for first eigenmode of Liver9 model

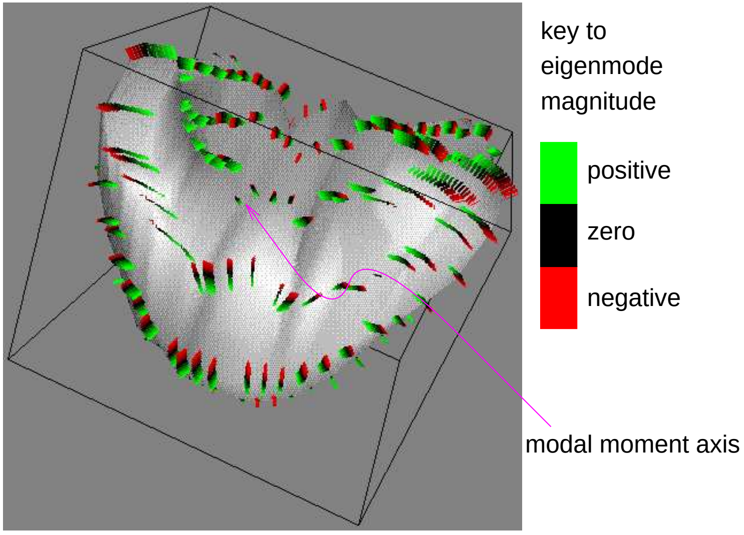

Figure 4.5: Modal moment axis for second eigenmode of Liver9 model

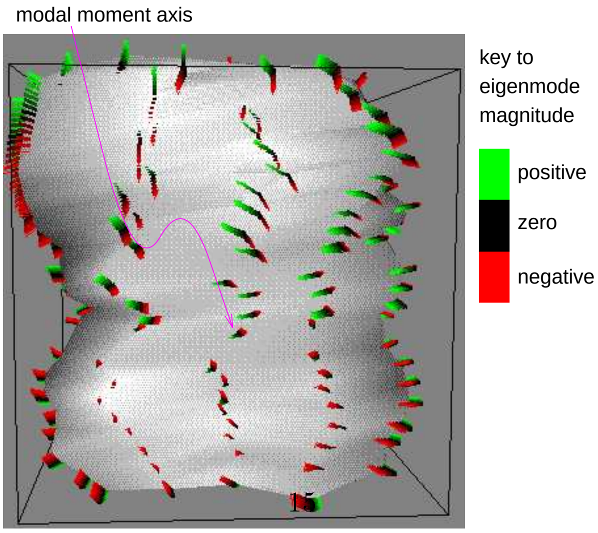

Figure 4.6: Modal moment axis for third eigenmode of Liver9 model

Vector moments are taken about the model centroid, of each eigenmode component about its point of action on each element node. The modal moment axis is defined as the axis which maximises the projection of the moment vector distribution, and can be computed as the primary eigenvector of the covariance matrix of moment vectors. Two modal moment axes are computed from the first two fundamental eigenmodes, and are mutually orthogonalised to give a right handed trio which forms the set of *principal modal axes*.

It can be seen that the ordering of the fundamental eigenmodes reflects the mirror and rotational symmetries of the gallbladder and liver models. The first two gallbladder modes give modal moment axes which are normals defining the two dominant planes of mirror symmetry which run along the major axis. The third gallbladder mode shows the plane of minor symmetry which is perpendicular to the long axis.

The first liver mode is perpendicular to the long axis and is the dominant plane of symmetry. The two other modes show up planes of symmetry which are less well defined.

4.2.1 Accuracy of modal axes

The directional accuracy of modal moment axes is dependant to some extent upon the sampling distribution of the Finite Element nodes, regardless of the accuracy of the eigenvector computation routine. In the worse case represented by the liver models in Figure 2.2, there is an error of approximately 30°, which is sufficiently accurate as an initial rigid-body registration.

When viewing animated eigenmodes, the human eye does a very good job of estimating the orientation of modal moment axes, and this is because the displacement field is integrated continuous over the model surface. This can also be done computationally by using FEM interpolation of the displacement field, at the cost of more computation.

Chapter 5

Bayesian registration

We present here a deterministic algorithm for determining a non-rigid registration between two biological organs using Bayes rule to combine a likelihood measure based on eigenmode similarity and a prior constraint based on a model of biological growth.

5.1 Growth model is the prior function

Syn & Prager (1995*a*) present a new biologically motivated metric of organ shape difference. The Gompertz function is suggested as an energetic metric of biological growth, based on experimental evidence from the literature as well as a thermodynamic model of growth control. It is shown that when comparing two fully grown organs which are treated as noisy instantiations of some prototype, the linearised metric is simply the mass difference between the organs.

This metric can be computed locally during the registration process, by an exactly constrained fit of principal eigenmodes of mass transfer to the number of committed homologies. This is explained in more detail below in Section 5.8.2 as the metric of local *Prior energy.*

5.2 Eigenmodes are the likelihood features

Syn & Prager (1995*b*) explain the properties of FEM eigenmodes when used as shape features, and Section 3.3 shows that eigenmode similarity between two points is a measure of the likelihood of their biological homology.

We use the monotonic and efficient 2-norm metric to quantify eigenmode similarity, given that eigenmode components have been normalised as described in Section 3.4. This gives us a metric of local *Likelihood energy.*

5.3 Markov neighbourhoods in Finite Element models

The Finite Element Method utilises shape functions to interpolate stresses and strains within each element (Syn & Prager 1995*b*), so that the displacement of a point within an element is solely a function of the displacements of the element nodes.

Consider the special case when this point coincides with an element node. This node is shared between adjacent elements, and its displacement is therefore only influenced by the nodal displacements of all these adjacent elements. The neighbourhood formed by such a group of adjacent elements is Markovian, and we can see that each nodal component of an eigenmode exists within a *Markov neighbourhood*.

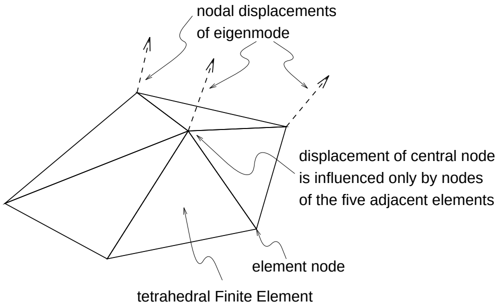

Figure 5.1: Five Finite Elements in a Markov neighbourhood

5.4 MRF-Gibbs equivalence

Bayesian combination of the Prior probability P_{Pr} and Likelihood probability P_{Lk}, leads to the Posterior probability P_{Ps} that a site on one model is homologous to a site on another model. Therefore the Posterior probability associated with the central site is computed only from eigenmode components within the Markov neighbourhood, and Bayes Theorem gives

$$P_{Ps} = \frac{P_{Pr} P_{Lk}}{Z} \tag{5.1}$$

where Z is a normalisation constant.

Likelihood and Prior energies over each Markov Random Field (MRF) can be described at a neighbourhood site f using a MRF-Gibbs equivalence (Chou & Brown 1990)

$$\begin{aligned} P_{Pr}(f) &= A_1 e^{-U_{Pr}(f)/T_1} & (5.2) \\ P_{Lk}(f) &= A_2 e^{-U_{Lk}(f)/T_2} & (5.3) \end{aligned}$$

Equations 5.2 and 5.3 have the probabilistic constraint that

$$\sum_{f \in MRF} P(f) = 1 \tag{5.4}$$

In order that the Posterior energy can be derived directly by adding Prior and Likelihood energies, we use this one constraint to normalise the energies of both Prior and Likelihood. Another constraint is simply that

$$A_1 = A_2 = 1 \tag{5.5}$$

since this only maps to a constant offset to U_{Pr} and U_{Lk} in each MRF.

Dropping the normalisation offset Z from Equation 5.1 gives the Posterior energy U_{Ps} (scaled by some T_3)

$$U_{Ps} = (U_{Pr}/T_1) + (U_{Lk}/T_2) \tag{5.6}$$

5.5 Normalising Prior and Likelihood energies

Use the change of variable $T' = -\frac{1}{T}$ over the clique, then

$$\begin{aligned} \sum_{i=1}^{N} e^{U_i T'} &= 1 \\ &= (e^{U_1})^{T'} + (e^{U_2})^{T'} + \ldots + (e^{U_N})^{T'} \end{aligned} \tag{5.7}$$

We can find T' by computing the root of the monotonic function

$$y = (e^{U_1})^{T'} + (e^{U_2})^{T'} + \ldots + (e^{U_N})^{T'} - 1 \tag{5.8}$$

using a simple Newton-Raphson iteration

$$\begin{aligned} T'_{next} &= T' - \frac{y}{\frac{dy}{dT'}} \\ &= T' - \frac{(e^{U_1})^{T'} + (e^{U_2})^{T'} + \ldots + (e^{U_N})^{T'} - 1}{U_1(e^{U_1})^{T'} + U_2(e^{U_2})^{T'} + \ldots + U_N(e^{U_N})^{T'}} \end{aligned} \tag{5.9}$$

which gives us T_1 and T_2 in Equation 5.6.

5.6 Scale-ordered registration

We have now established an energetic measure which describes the likelihood of two sites on different models being homologous, conditioned by the constraint of biological plausibility. This local information has to be integrated over the entire surface so that we arrive at a global homology map which minimises the Posterior energy over all sites.

The first non-rigid eigenmode is of fundamental spatial frequency so that there is only one phase change over the model. A non-rigid registration using this eigenmode should give only one of two locally optimal fits[1]. Similarly registration using subsequent eigenmodes gives a spatial localisation determined by the eigenmode's spatial wavelength.

Chapter 4 shows how a pose estimate can be arrived at to give an approximate rigid-body registration, *which should lie within the domain of the "correct" local optimum* of the first eigenmode's non-rigid registration. The homology map is then estimated to the spatial resolution of the first eigenmode, and refined using as many eigenmodes as are similar between the two models. The number of similar eigenmodes two models have in common therefore determines the resolution to which a biological homology can be determined.

---

[1] Two local optima occur because the two phases of a fundamental eigenmode can be matched in two ways to a corresponding fundamental eigenmode.

5.7 Highest Confidence First estimation

The minimisation problem outlined above has the following characteristics

1. initialisation within the domain of a local optimum
2. graduated presentation of increasingly localised optima

so the algorithm used needs to

1. initialise the homology map using the first eigenmode
2. refine the homology map using subsequent eigenmodes, while taking into account the registration given by previous eigenmodes

We have implemented the Highest Confidence First (HCF) algorithm to meet these criteria. HCF is an efficient algorithm which embodies the principles of graceful degradation and least committment, with small and predictable run-time (Chou & Brown 1990). It performs an approximate Maximum A-Posteriori (MAP) optimisation using local homology estimates, with global scheduling to rank committment decisions in order of importance.

5.7.1 Augmented search space

Chou & Brown (1990) define an *augmented search space* as the set of possible homologies or *committments* for a site, augmented with the additional option of being *uncommitted.* After a site has committed to a homology, it can still change its committment but cannot nullify it.

The *augmented Posterior energy* is a modified measure which only considers committed sites in the MRF. An uncommitted site always takes into account the states of active neighbours when making a committment however.

5.7.2 Local stability measures

The order in which sites have their committments updated, is determined by a *priority queue* with the least stable site at the top of the heap.

The stability measure for a committed site is the difference in Posterior energy between the current committment f_1 and a committment (different from the current one) f_2 which gives the smallest Posterior energy. This represents the largest reduction in local Posterior energy which can be achieved by a change in committment.

$$G = \min_{f_2 \in MRF, f_2 \neq f_1} U_{Ps}(f_2) - U_{Ps}(f_1) \tag{5.10}$$

The stability measure for an uncommitted site is the negative difference in Posterior energy between the sites which give the two smallest Posterior energy values. This indicates the depth of the energetic "well" which the best committment sits in, and is always negative.

$$G = \min_{f_2 \in MRF, f_2 \neq f_1} -\Big(U_{Ps}(f_2) - U_{Ps}(f_1)\Big) \tag{5.11}$$

$$\text{s.t. } U_{Ps}(f_1) = \min_{f \in MRF} U_{Ps}(f) \tag{5.12}$$

5.8 Implementation

All sites are initialised as being *uncommitted* to begin with, so that the augmented local Posterior energies are essentially determined by local Likelihood. Therefore sites with strong Likelihood evidence of one committment over others will be visited early in the construction of the homology map. Sites without strong Likelihood evidence will take the neighbourhood configuration into account when making committments.

A serial implementation of HCF still has the weakness that the first committment made strongly influences subsequent committments but may be inaccurate. Chou, Cooper, Swain, Brown & Wixson (1993) present Local HCF as a parallel implementation which addresses this potential instability by integrating initial committments over the network of sites.

5.8.1 Spatial refinement

HCF terminates when the top of the priority queue, which represents the least stable site, is positive. The homology map obtained is used to update the rigid-body registration, which improves the accuracy of the local Prior energy measures).

The map initialises another HCF run using the next eigenmode, which gives a principled method of propagating Posterior estimates from all previous runs as Prior information in the current run.

5.8.2 Computing Prior energy

Local Prior energies in a MRF are usually defined with respect to cliques, which are totally connected subgraphs in the MRF. Cliques allow the Prior energy to be built up piecewise from different configurations of spatial dependency. Examples of potential functions assembled from heuristically assigned clique functions are found in (Chou & Brown 1990).

The Gompertz metric of biological growth gives a Prior energy measure computed from local mass change, which can described locally as a linear sum of eigenmode components. Each eigenmode component has an associated local mass change, which varies non-linearly with modal amplitude. Given any committment configuration in a MRF, an exactly constrained least-squares fit is performed to find eigenmode amplitudes from homology positions, whence the local Prior energy.

The qualitative behaviour of the Gompertz metric encourages homology vectors in each element's neighbourhood to not only minimise displacement, but also form a coherent cluster of committments. This coherence propagates through overlapping MRFs from the top of the stability queue downwards to less stable sites.

In cases where the mesh is both noisily and sparsely sampled, the determination of eigenmode amplitudes can be made overconstrained by using fitting only a few eigenmodes to a larger number of committment constraints. This reduces the ability of local Prior energy to discriminate between similar committment patterns, and will dampen the local variability of the homology map.

Chapter 6

Results

6.1 Rigid-body registration using principal modal axes

Figure 6.1 shows the rigid-body registration of two gall-bladder models using principal modal axes.

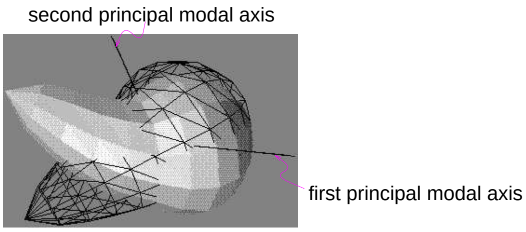

Figure 6.1: Rigid-body registration of gall-bladder models using principal modal axes

Figure 6.2 shows the refinement in pose estimation derived from the first and second modal homology maps.

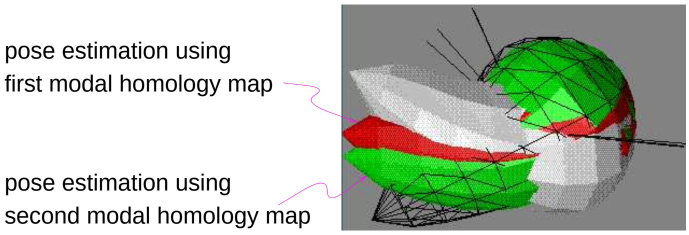

Figure 6.2: Refinement of rigid-body registration using homology maps

6.2 Relative importance of Likelihood features

Gee, Le Briquer & Barillot (1995) note that the exact form of the Prior function is unimportant, so long as it penalises "unlikely" homologies and encourages "likely" ones. The quality of the homology map is determined very much by the quality of the Likelihood features, which is demonstrated below. Note that the size of the wireframe model has been scaled up by 10% to show the homology vectors more clearly.

6.2.1 Registration using only Prior energy

Figure 6.3(a) shows the registration achieved using only a simple Euclidean metric in comparing homologous sites. The Euclidean metric performs relatively well in this situation, even without neighbourhood information, because of the lack of noise, good initial pose estimation, and the coherent relative spacing of the two models' element nodes.

Figure 6.3(b) shows the registration achieved using only the Gompertz metric which operates over an MRF when comparing homologous sites. In general the Gompertz metric is more reliable because it has neighbourhood support in deciding the biological plausibility of a suggested homology.

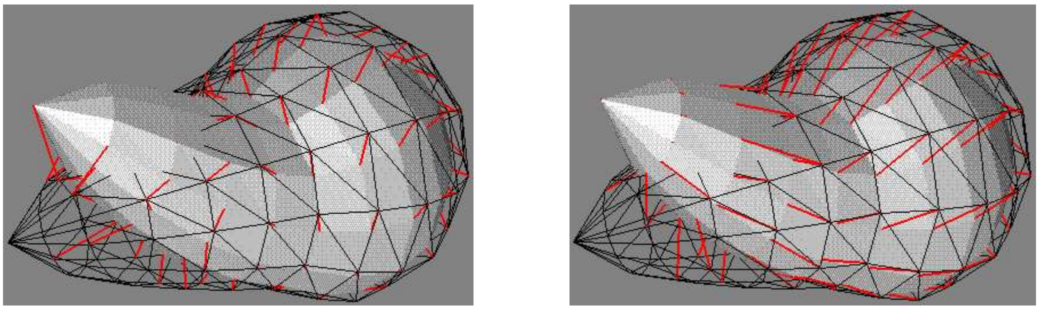

(a) Euclidean Prior

(b) Gompertz Prior

Figure 6.3: Registration using only Prior energy

6.2.2 Registration using only Likelihood

Figure 6.4 shows the homologies obtained by using only Likelihood features in HCF estimation. There is no Prior constraint on a site to commit to another nearby site. Figure 6.4(a) shows that all sites commit to the correct side of the other model, but with bad spatial localisation since only the first eigenmode is used. Figure 6.4(b) shows that with the first two (mutually-orthogonal) fundamental eigenmodes, spatial localisation is improved.

The two gall-bladder models have 18 non-rigid modes in common, where the 18th mode has a spatial wavelength of approximately 1/4 the length of the model. Figure 6.4(c) shows that with gradual spatial localisation using all 18 eigenmodes, all sites including the awkward ones near the gallbladder tip have committed with good spatial resolution. The hidden part of the homology map to the rear of the models, also has improved localisation compared to Figure 6.4(b). The improved pose estimation by the homology maps have also given a better underlying rigid-body registration.

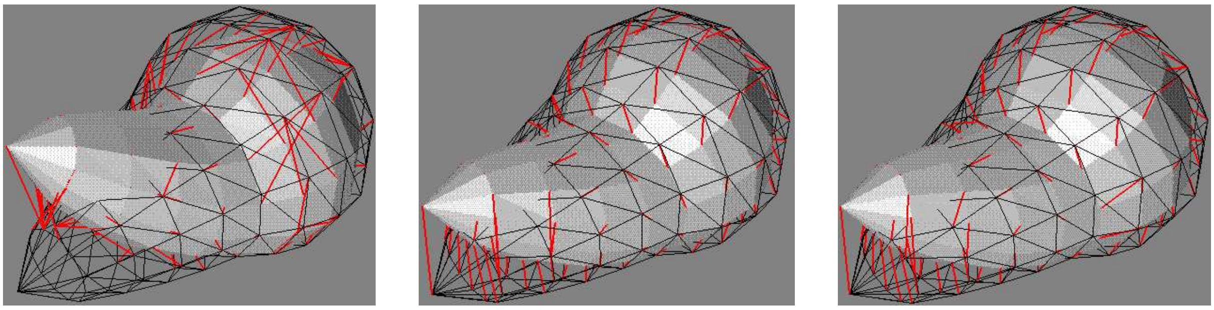

(a) 1st eigenmode (b) 2nd eigenmode (c) 18th eigenmode

Figure 6.4: Registration using only modal Likelihood features

6.3 Registration using three eigenmodes

Figure 6.5 shows the successively refined HCF estimate of 3D registration between two gall-bladder models. Figure 3.1 showed that the first three eigenmodes are of fundamental spatial wavelength, so that the registration achieved in Figure 6.5(b) should be localised to approximately half the length of the gall-bladder model. In fact the presence of the Gompertz prior and the regular spacing of the two meshes, means that the homology map is almost perfectly recovered[1].

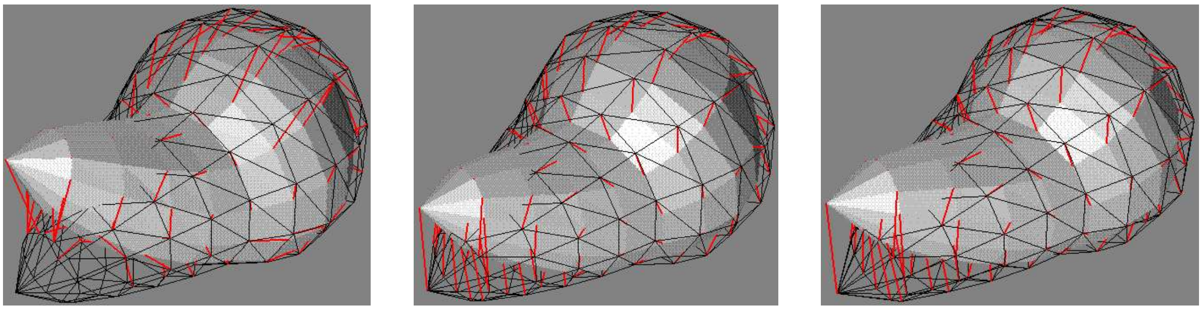

(a)
HCF estimate using 1st eigenmode

(b)
HCF estimate using 2 eigenmodes

(c)
HCF estimate using 18 eigenmodes

Figure 6.5: HCF estimate and refinement of gall-bladder registration using first three eigenmodes

The effect of the Gompertz prior can be seen in Figure 6.5(a), which has a more coherent

---

[1] Standard FEM modelling practice suggests 8 elements per feature wavelength, so the maximum registration resolution of the gallbladder models according to this criterion is approximately the length of the gallbladder.

homology map than Figure 6.4(a), in which only the Likelihood feature was used.

Chapter 7

Discussion

7.1 Initial rigid-body registration

Section 3.4 described an eigenmode normalisation scheme whereby eigenmode directions and magnitudes are uniquely determined for use as Likelihood features. The initial pose estimate is also uniquely determined by this normalisation, which relies on volume-change as a reliable macroscopic property of the eigenmode field. When the model's spatial sampling is sparse and noisy however, a more robust scheme is required.

Without a unique normalisation of the first two modal moment axes, there are four possible initial pose estimates, and the best one is determined using a nearest-neighbour metric which sums the distance between surface element mid-points.

The reason we separate the processes of initial pose estimation, and subsequent non-rigid registration is that in almost all cases of MR or ultrasound scanning, the initial pose estimate is in fact available from information about patient orientation.

7.2 Registration without initial pose estimate

The Necker cube illusion strongly suggests that the Human Visual System (HVS) performs local optimisation in fitting a model to the image, and actively peturbs each local optimum to converge onto alternative optima. This is consistent with a scheme in which a number of hypothesised registrations are computed and compared, leading to a more robust registration algorithm.

For each of the four initial pose estimates derived from principal modal axes without direction normalisation, we initialise the homology map using only Prior energy since Likelihood features are not yet direction normalised. Then for every pair of corresponding eigenmodes, sign-indeterminate displacement components are compared at homologous sites on each model, which determines if the two eigenmodes are signed consistently.

Likelihood features are now used to refine the homology map, and the sign-determination process can be repeated again. The four final homology maps are compared, and the one which has the lowest total remaining Posterior energy is chosen.

7.3 Deterministic versus statistical algorithms

The deterministic local-search approach of HCF estimation is also in accordance with the Necker cube. HCF is preferable to stochastic algorithms such as simulated annealing which are computationally expensive although globally optimal. Stochastic algorithms do not give repeatable results so the quality of solutions needs to be assessed over a number of runs. HCF also has the additional advantage that there is an explicitly defined site-updating schedule.

Gee et al. (1995) make the point that Bayesian matching algorithms are highly dependant upon the quality of the Likelihood features used, and that the Prior has just to behave as a reasonable penalty function. Thus designing the registration algorithm to use powerful Likelihood features obviates the need for a globally-optimal stochastic algorithm.

7.4 Choosing an appropriate Prior

The Gompertz prior has a number of attractive properties:

1. it is both experimentally and thermodynamically motivated as a local metric of biological growth energy;
2. it can be used to compare adult organs from different patients;
3. it is computable for any clique configuration in a MRF;
4. it encourages a coherent (ie. consistent with our biological growth model) local pattern of committment, propagated by HCF over the entire homology map.

Most penalty functions applied in medical imaging are mechanically motivated eg. the membrane model (Amit, Grenander & Piccioni 1991), Laplacian model (Szeliski 1990), thin-plate or biharmonic model (Bookstein 1991), linear elasticity model (Syn & Prager 1995*b*) etc., and the viscous fluid prior used by Christensen, Rabbitt, Miller, Joshi, Grenander, Coogan & Van Essen (1995) has produced particularly impressive results in 2D and 3D mapping of images. In comparison with the Gompertz prior:

1. the viscous fluid prior was developed to allow large nonlinear strains, and has no biological basis;
2. there is no justification for using it to compare either models of a growing organ, or organs from different patients;
3. it is not linearised into eigenvectors since it is designed specifically to compensate for linearisation effects, so cannot be computed for arbitrary clique committment configurations;
4. it can therefore only constrain local homologies when it is gridded in a computable configuration.

Computations are run for hours on a massively-parallel supercomputer, and results have been shown which map (non-reversibly) a monkey's brain to a child's (Miller, Christensen, Amit & Grenander n.d.). In the context of medical imaging however, a Prior function should penalise such extreme examples as being biologically implausible instead of accomodating them.

7.5 Eigenmode similarity

We use the Lanczos algorithm with reorthogonalisation (Syn & Prager 1995*b*) to compute FEM eigenmodes. Our implementation which has been generalised to operate with eigenspaces so as to detect repeated eigenmodes. We use a slow but robust reorthogonlisation scheme which unfortunately means high-order eigenproblems can take a longer time to solve for a certain required accuracy.

When registering models which are more densely and accurately sampled, there are likely to be more similar eigenmodes which are also more accurately computed. Given enough eigenmodes registration can proceed using *projected eigenmode components* as Likelihood features (Syn & Prager 1994). These projected components are defined relative to lower eigenmodes, and are therefore rotation invariant so that there is no need for initial pose estimation.

7.5.1 Liver models

The liver models shown in Figure 2.2 were sampled with over 200 surface points and 10 internal points on the long axis. This results in long tetrahedral elements whose radial wavelength is already half the width of the liver, so that coupling between opposite sides of the liver model is quite coarse. This compromises the precision of the eigenmode displacement field, which is why the Liver models have at most 4 eigenmodes in common. These modes are all of fundamental spatial frequency, so that the derived homology map has poor spatial localisation.

A more sophisticated scheme for sampling internal points, such as the octree scheme described in Syn & Prager (1994), will improve both the quality and number of similar eigenmode features which can be used in registration. Even with sparse and very uneven sampling (eg. Liver11 has a few slices missing and one degenrate point near the thin end), the computed eigenmodes are still robust enough to give at least 2 sets of eigenmode features for use in registration.

7.6 Summary

We have presented a local Likelihood measure of biological homology, together with a Prior measure of biological growth energy. A Bayesian framework is used to compute a local Posterior measure for any committment configuration within MRFs defined in the neighbourhood of each element node. HCF estimation is used to initialise homology committments, and then in scale-ordered refinement of the homology map.

We will continue to update this report as experiments yield more results.

Anonymous ftp

Technical reports in this series can be obtained via anonymous ftp from our server

`svr-ftp.eng.cam.ac.uk`

World Wide Web

Reports can also be accessed through the Speech,Vision and Robotics Group's web server

`www.eng.cam.ac.uk`

Bibliography

Amit, Y., Grenander, U. & Piccioni, M. (1991), 'Structural image restoration through deformable templates', *Journal of the American Statistical Association* **86**(414), 376–387.

Bookstein, F. (1991), *Morphometric tools for landmark data : geometry and biology*, 1st edn, Cambridge University Press.

Chou, P. & Brown, C. (1990), 'The theory and practice of Bayesian image labeling', *International Journal of Computer Vision* 4, 185–210.

Chou, P., Cooper, P., Swain, M., Brown, C. & Wixson, L. (1993), Probabilistic network inference for cooperative high and low level vision, *in* 'Markov Random Fields : Theory and Application', Academic Press. FC.46.CS.

Christensen, G., Rabbitt, R., Miller, M., Joshi, S., Grenander, U., Coogan, T. & Van Essen, D. (1995), Topological properties of smooth anatomic maps, *in* Y. Bizais, C. Barillot & R. Di Paola, eds, 'Information Processing in Medical Imaging', Computational Imaging and Vision, Kluwer Academic Publishers, pp. 101–112.

Gee, J., Le Briquer, L. & Barillot, C. (1995), Probabilistic matching of brain images, *in* Y. Bizais, C. Barillot & R. Di Paola, eds, 'Information Processing in Medical Imaging', Computational Imaging and Vision, Kluwer Academic Publishers, pp. 113–125.

Hitchens, D., ed. (1992), *A Finite Element Dynamics Primer*, NAFEMS, Birniehill, East Kilbride, Glasgow.

Miller, M., Christensen, G., Amit, Y. & Grenander, U. (n.d.), 'Mathematical textbook of deformabel neuroanatomies', *Proceedings of the National Academy of Science.*

Pasapula, C. (1995), 3D shape analysis using point based distribution models on fetal cadaveric livers, internal report, Guy's and St. Thomas' Hospital.

Sclaroff, S. & Pentland, A. (1995), 'Modal matching for correspondence and recognition', *IEEE Trans. Pattern Analysis and Machine Intelligence* **17**(6), 545–561.

Syn, M. & Prager, R. (1994), Mesh models for three-dimensional ultrasound imaging, Technical Report CUED/F-INFENG/TR210, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Syn, M. & Prager, R. (1995*a*), A biological growth metric for 3D shape registration, Technical Report CUED/F-INFENG/TR225, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Syn, M. & Prager, R. (1995*b*), FEM eigenmodes as shape features, Technical Report CUED/F-INFENG/TR211, Cambridge University Engineering Department, Trumpington Street, CB2 1PZ.

Szeliski, R. (1990), *Bayesian modeling of unvertainty in low-level vision*, Kluwer Academic Publishers.