A Probabilistic Fitness Measure for Deformable Template Models

J. Haslam, C.J. Taylor, T.F. Cootes

Dept. of Medical Biophysics,
Stopford Building,
University of Manchester,
Manchester,
M13 9PT.
Phone: 061 275–5130

Abstract

Methods for automatic image interpretation based on the use of deformable template models have proved very successful. Whatever deformable template scheme is used, one of the basic requirements is a method for assessing the likelihood that a particular model instance is the correct interpretation of a given image. We describe a Bayesian 'fitness' measure which combines the likelihood of the model shape with the evidential support in a principled way. Image search is carried out by minimising the fitness measure using multi–scale quasi–Newtonian optimisation. We have previously compared the performance of different fitness measures. Here we give results for the new method and show that, by making optimal use of the image evidence, it achieves more accurate interpretation than the best of the methods we have previously tested.

Introduction

Flexible template models have been used successfully for many applications of automatic image interpretation[1,2,3,4]. The template embodies *a priori* knowledge of the typical shape and shape variability of the modelled object(s). A number of modelling methods have been employed to represent structures of interest. Lipson et al [1] model vertebral cross–sections using a simple ellipsoidal template, Bajcsy and Kovavic [2] use a 3–dimensional voxel representation of an anatomical brain atlas, whilst Staib and Duncan [3] parameterize the left ventricle of the heart in echocardiograms using an elliptic Fourier decomposition its closed boundary. Generally, image search for flexible template models is performed by moving and deforming the template so as to minimise some energy function, or fitness measure, which has been defined to correspond to the best fit with the image. The fitness measure typically favours matching the model to strong edge evidence in the image. However, little attention has been paid to making proper use of the evidence. For example, it is often implicit that all of the supporting evidence extracted from the image is independent and of equal reliability; neither assumption is, in general, realistic. In this paper, we attempt to tackle this problem in a more principled way by defining a Bayesian fitness measure for use with flexible template models.

Our fitness measure combines both grey–level evidence and *a priori* knowledge of shape probability to produce a likelihood estimate that a given model instance is the correct interpretation. It can therefore be used with any shape modelling method for which an estimate of shape probability is available. In order to obtain the fitness measure, we build statistical models of the grey–level evidence extracted from the image, and the object shape parameters. The models are built using 'training data' gathered from a number of example images. The shape models are Point Distribution Models (PDMs), based on the work of Cootes et al. [4].

BMVC 1994 doi:10.5244/C.8.3

During image search, the fitness measure is minimised using a standard quasi-Newtonian optimisation algorithm. We use a multi-resolution search strategy, performing the optimisation on each level of an image pyramid in turn. We present results for trans-axial Magnetic Resonance images of the lower abdomen (male). These images are particularly difficult to interpret because they are 'cluttered' and the contrast between structures of interest is poor. We have run multiple-trial experiments comparing the accuracy of interpretation obtained using the new fitness measure with that achieved using alternative approaches and show a significant improvement in performance.

The Bayesian Fitness Measure

The likelihood of a flexible template instance being the correct fit to the *image* evidence may be expressed using Bayes' Theorem :

$$P(\mathcal{F}|E) = \frac{P(E|\mathcal{F})P(\mathcal{F})}{P(\mathcal{F})P(E|\mathcal{F}) + P(\mathcal{B})P(E|\mathcal{B})} \tag{1}$$

where $P(\mathcal{F}|E)$ is the probability of the *foreground* (or correct) state $\mathcal{F}$ given image evidence E, $P(E|\mathcal{F})$ is the probability of observing evidence E in the foreground state $\mathcal{F}$, $P(\mathcal{F})$ is the prior probability of observing the foreground state $\mathcal{F}$, $P(E|\mathcal{B})$ is the probability of observing evidence E in the *background* state (ie any state other than the foreground state), and $P(\mathcal{B})$ is the prior probability of observing the background state $\mathcal{B}$. The foreground model summarises the distribution of grey-level evidence observed when the model represents the correct solution; the background model summarises the distribution of grey-level evidence observed when the model is *not* at the correct solution.

Similarly, the likelihood of a flexible template instance being an example of the model shape is given by :

$$P(\mathcal{F}|S) = \frac{P(S|\mathcal{F})P(\mathcal{F})}{P(S)} \tag{2}$$

where $P(\mathcal{F}|S)$ is the probability of the foreground state $\mathcal{F}$ given shape S, $P(S|\mathcal{F})$ is the probability of observing shape S in the foreground state $\mathcal{F}$, $P(\mathcal{F})$ is the prior probability of observing the foreground state $\mathcal{F}$, and $P(S)$ is the prior probability of observing shape S. We assume that $P(S)$ and $P(\mathcal{F})$ are constants ie all shape states are equally probable and there is always the same chance of finding a foreground model. Then,

$$P(\mathcal{F}|E,S) = P(\mathcal{F}|E)P(\mathcal{F}|S) \propto \frac{P(S|\mathcal{F})P(E|\mathcal{F})}{P(E|\mathcal{F}) + \frac{P(\mathcal{B})}{P(\mathcal{F})}P(E|\mathcal{B})} \tag{3}$$

We can replace the probabilities in equation (3) by probability *densities*, giving :

$$P(\mathcal{F}|E,S) \propto \frac{p(S|\mathcal{F})p(E|\mathcal{F})}{p(E|\mathcal{F}) + \frac{P(\mathcal{B})}{P(\mathcal{F})}p(E|\mathcal{B})} \tag{4}$$

In order to estimate $P(\mathcal{F}|E,S)$ we need to model each of the terms in equation (4). $p(E|\mathcal{F})$ and $p(E|\mathcal{B})$ may be determined by modelling the grey-level evidence observed for examples of the *foreground* and *background* states respectively. Both models can be obtained by performing a statistical analysis of grey-level data from a training set of images, each annotated with the correct interpretation. $p(S|\mathcal{F})$ is equal to the prior probability of observing shape S for a correct interpretation $\mathcal{F}$. We obtain a value proportional to $p(S|\mathcal{F})$ by building a parameterised model summarising the probability distribution for each of the shape parameters. Again, this is achieved by performing a statistical analysis of a number of training examples. Finally, we assume that the ratio of the Bayesian priors, $p(\mathcal{B}) : p(\mathcal{F})$ is constant, and equal to 1. We make this assumption on the grounds that it will be approximately true when the model instance is near the

correct solution. This means that towards the convergence of image search, we will be able to extract a final model probability which is reasonably accurate.

When investigating the behaviour of the fitness measure and in image search, it is convenient to deal with log probabilities because they allow a large dynamic range to be accommodated. Thus, we use a fitness measure given by :

$$f = -\ln\left(p(\mathcal{F}|E,S)\right) + C \tag{5}$$

$$= -\ln(p(E|\mathcal{F})) + \ln(-p(E|\mathcal{F}) - p(E|\mathcal{B})) - \ln(p(S|\mathcal{F})) + C \tag{6}$$

where C is a constant which we take to be zero.

Modelling Shape

We show how to obtain a value for $p(S|\mathcal{F})$ for a Point Distribution Model (PDM). PDMs model the shapes of variable objects by representing each object as a set of labelled points. Each point corresponds to a particular location on the object. The model consists of the mean positions of these points over a training set of examples and the main modes of variation describing how the points move about their mean positions. A shape instance may be calculated as :

$$\mathbf{s} = \bar{\mathbf{s}} + \mathbf{P}_s\mathbf{b}_s \tag{7}$$

where $\mathbf{s} = (x_0, y_0, x_1, y_1, \ldots, x_k, y_k, \ldots, x_{m-1}, y_{m-1})^T$, (x_k, y_k) is the position of point k, $\bar{\mathbf{s}}$ represents the mean shape, $\mathbf{P}_s$ is a $2m \times t_s$ matrix of the t_s unit eigenvectors of the covariance matrix of $\mathbf{S}$, and $\mathbf{b}_s = (b_{s1}\ b_{s2}\ \ldots\ b_{st_s})^T$ is a set of shape parameters

We estimate $p(S|\mathcal{F})$ for a PDM as follows. Firstly, we assume a normal distribution in each of the shape parameter values $\mathbf{b}_s$. For a shape instance $\mathbf{s}$, this gives a probability density:

$$p(\mathbf{s}|\mathcal{F}) \propto e^{-M_s/2} \quad ; \quad M_s = \sum_{j=1}^{t_s} \frac{b_{sj}^2}{\lambda_{sj}} \tag{8}$$

where M_s is the Mahalanobis distance, and λ_{sj} is the eigenvalue corresponding to the jth eigenvector.

Modelling the Foreground Evidence

We obtain a value for $p(E|\mathcal{F})$ by modelling the grey-level evidence observed for examples of the *foreground* state. We gather training data for the model as follows :

- For each model point in each member of the training set, we extract a profile from the image, of length n_p pixels, centred on the model point, and perpendicular to the model boundary.
- For each example in the training set, we concatenate the profiles to give a single grey level vector $\mathbf{F} = (\mathbf{F}_1, \mathbf{F}_2, \ldots \mathbf{F}_{n_m})$ as shown in Figure 1.

  Each vector is thus of length $n = n_p \times n_m$ pixels, where n_m is the number of points in the shape model.

Since each model point corresponds to the same structure in every training image, we expect the extracted profiles to be similar in every example. We build the grey-level model from the training data similarly to Cootes et al. [5]. We calculate a mean profile, $\bar{\mathbf{F}}$, and construct an n^2 covariance matrix $\mathbf{S}_\mathbf{F}$. This gives a second order statistical description of the expected grey-level evidence. The variation about the mean is described by $\mathbf{P}_\mathbf{F}$, the eigenvectors of $\mathbf{S}_\mathbf{F}$ corresponding to the $t_\mathbf{F}$ ($< n$) largest eigenvalues. We can thus express an instance of the foreground profile model as :

$$\mathbf{F} = \bar{\mathbf{F}} + \mathbf{P}_\mathbf{F}\mathbf{b}_\mathbf{F} \tag{9}$$

where $\mathbf{b}_\mathbf{F}$ is a set of t_F parameters describing the profile model.

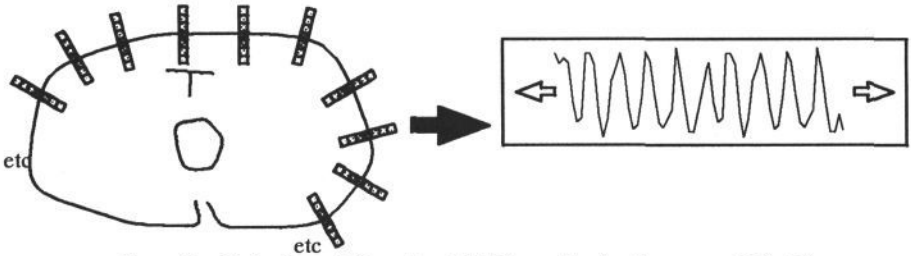

Figure 1 : Gathering of Grey–Level Evidence for the Foreground Model – Grey–Level Profiles Gathered Perpendicular to the Model Boundary at each Model Point are Concatenated Together

The difference between our approach and that of Cootes et al. lies in our use of *concatenated* training profiles. Whereas Cootes builds an individual profile model for each shape model point, we build a single, concatenated profile model. We do this because there is a difficulty in combining probabilities generated for the individual models to give a single foreground evidence probability. The grey–level profile data gathered from different model points may be correlated, particularly for nearby points. We cannot therefore simply assume independence of the individual model probabilities – doing so will give us an incorrect estimate of overall model probability. A solution to the problem is to use a single profile model, which implicitly takes care of any correlation. This allows us to obtain a more principled estimate of model probability.

If we assume that the original foreground profile data follows a normal distribution, the probability density function for observing a foreground grey–level model instance **F** is given by :

$$p(\mathbf{F}\,|\,\mathcal{F}) \propto e^{-M_{E_F}/2}\;;\quad M_{E_F} = \sum_{j=1}^{t_F} \frac{b_{Fj}^2}{\lambda_{Fj}} \tag{10}$$

where M_{E_F} is the Mahalanobis distance of the multi–variate normal probability density function, λ_{Fj} is the eigenvalue corresponding to the jth eigenvector, $\lambda_{Fj} > \lambda_{Fj+1}$, and t_F is the number of eigenvectors used in the model. M_{E_F} thus gives a fitness measure for the grey–level evidence *explained* by the foreground model.

However, unless we have a very large number of training examples, t_F will normally be less than n. This means that we are modelling the grey–level evidence using a basis set which does not span the observation space; there are in principle $(n–t_F)$ unmodelled dimensions. If we reconstruct a grey–level profile from a valid interpretation using equation (9), we find that the reconstructed profile is slightly different from the original. This difference or *unexplained evidence* is caused by the model truncation, and must be accounted for if we are to accurately calculate a probability for a foreground model instance. Cootes et al. [5] do this by estimating the contribution of unexplained evidence to the foreground model probability; we choose to build a simple single–parameter *unexplained evidence model*.

We define the unexplained evidence as a vector $\mathbf{u_F}$ of length n, given by :

$$\mathbf{u_F} = \mathbf{F}^{obs} - \mathbf{F}^{\mathrm{mod}} \tag{11}$$

where $\mathbf{F}^{obs}$ is the grey–level intensity vector for the concatenated profile, and $\mathbf{F}^{\mathrm{mod}}$ is the grey–level intensity vector as it is best expressed in terms of the foreground model.

We make a number of measurements of $\mathbf{u_F}$; in order to obtain a representative sample, we measure $\mathbf{u_F}$ for examples not originally used to train the foreground model. The procedure for obtaining n_t measurements of $\mathbf{u_F}$ is as follows. For each example in the

original training set, we train a foreground model using the other $n_t - 1$ examples. We then measure $\mathbf{u_F}$ for the example using the 'leave one out' model.

Since we expect the unexplained evidence to be largely uncorrelated, we assume that we can only extract one unexplained evidence model parameter reliably. This implies that the variance in every dimension of the unexplained evidence model space is the same. We can use the n_t unexplained evidence examples to estimate this variance λ_{U_F} :

$$\lambda_{U_F} = \frac{1}{n_t(n - t_F)} \sum_{ij} u_{ij}^2 \tag{12}$$

where u_{ij} is the grey level intensity of the *ith* pixel of the *jth* example.

Taking the unexplained evidence into account, the Mahalanobis distance for the full n–dimensional observation space M_F is given by :

$$M_F = \sum_{j=1}^{t_F} \frac{b_{Fj}^2}{\lambda_{Fj}} + \sum_{j=t_F+1}^{n} \frac{b_{Fj}^2}{\lambda_{Fj}} \tag{13}$$

$$= M_{E_F} + M_{U_F} \tag{14}$$

where M_{U_F} is the contribution to the total Mahalanobis distance from the unexplained evidence.

Since we have used only t_F modes in our model, we have no knowledge of b_{Fj} or λ_{Fj} for $j > t_F$. However, it can be shown (see Cootes et al [5]) that

$$\sum_{j=t_F+1}^{n} b_{Fj}^2 = R_F^2 \tag{15}$$

where R_F^2 is given by :

$$R_F^2 = (\mathbf{F}^{obs} - \overline{\mathbf{F}})^T(\mathbf{F}^{obs} - \overline{\mathbf{F}}) - \mathbf{b}_\mathbf{F}^T\mathbf{b}_\mathbf{F} \tag{16}$$

Since we assume the variance for all dimensions of the unexplained evidence model to be equal to λ_{U_F}, M_{U_F} may be written as :

$$M_{U_F} = \frac{R_F^2}{\lambda_{U_F}} \tag{17}$$

If we assume a multi–variate normal distribution, the probability density of observing foreground evidence $\mathbf{F}$ may therefore be expressed as :

$$p(\mathbf{F} \mid \mathcal{F}) \propto e^{-(M_{U_F} + M_{E_F})/2} \tag{18}$$

Modelling the Background Evidence

We obtain a value for $P(E \mid \mathcal{B})$ by modelling the grey–level evidence observed for examples of the *background* state. In order to build the model we must define what we mean by 'background'. The proper definition of background is '*every possible model configuration in the image except the foreground state*'. However, it is impractical to gather training data for such a model. Instead, we define a subset of all possible background states : '*those states of the model close to, but not at, the correct model state*'. This definition allows us to discriminate between the foreground and background states for model instances near the correct answer. We therefore train the background model as follows :

- For each training image, we generate a number of random model poses, with values of x, y, scale and rotation near, but not at, the correct model pose for the image. In each case the correct shape parameters are used, though they could also have been varied.

- An area of pose-space close to the correct (foreground) model configuration is excluded from the background model. The amount of excluded pose-space is determined by the inherent inaccuracy in position of the PDM shape model points due to the limited number of shape modes used to construct the shape model.
- For each member of the randomly generated training set of model configurations, we extract profiles from the corresponding image, and train a concatenated profile model as for the foreground.

An example of the profile grey-levels for a background instance can, by analogy with (9), be reconstructed from :

$$\mathbf{B} = \overline{\mathbf{B}} + \mathbf{P}_{\mathbf{B}}\mathbf{b}_{\mathbf{B}} \tag{19}$$

where $\mathbf{b}_{\mathbf{B}}$ is a set of t_B parameters .

As with the foreground model, we also model the training set background evidence *not* explained by the background model, gathering examples of the *unexplained background evidence* on a 'leave one out' basis as before, and calculating a parameter λ_{U_B} which we take to be the variance for all dimensions of the unexplained evidence model.

By analogy with equation (18), the probability density of observing grey-level background model instance **B** is given by :

$$p(\mathbf{B}|\mathfrak{B}) \propto e^{-(M_{U_B}+M_{E_B})/2} \tag{20}$$

where M_{E_B} is the Mahalanobis distance for the explained evidence background model:

$$M_{E_B} = \sum_{j=1}^{t_B} \frac{b_{Bj}^2}{\lambda_{Bj}} \tag{21}$$

λ_{Bj} is the eigenvalue corresponding to the jth eigenvector, and M_{U_B} is the Mahalanobis distance for the unexplained evidence background model :

$$M_{U_B} = \frac{R_B^2}{\lambda_{U_B}} \tag{22}$$

(all other symbols have analogous meanings to those used for the foreground model)

Results with Real Images

Combining equations (6),(18) and (20) we obtain a grey-level evidence fitness measure, f_g, given by :

$$f_g = M_{E_F} + M_{U_F} + \ln(M_{E_F} + M_{U_F} + M_{E_B} + M_{U_B}) \tag{23}$$

Finally, including the shape evidence gives us an overall fitness measure, f :

$$f = f_g + M_s \tag{24}$$

We have performed experiments on a set of real images to investigate the behaviour of f_g. We used a PDM representing some of the structure visible in a trans-axial Magnetic Resonance image of the lower abdomen (male), approximately mid-way through the prostate gland (see Figure 3 a)). The PDM, trained on 40 examples, uses 79 model points to represent the structure of the prostate, a T-shaped structure called the symphasis pubis and the skin. 7-pixel profiles for building both foreground and background models were extracted at each model point in each training set image.

Figure 2 gives results for the components of f_g vs. x displacement of the model in the image, keeping all other model pose and shape parameters at their correct values. Figure 2a) plots the Mahalanobis distance for the explained evidence models (eqns (10),(21)); 2b) plots the Mahalanobis distance for the unexplained evidence models

(eqns (17),(22)), and 2c) shows f_g itself (eqn (23)). We see that f_g is dominated by our estimate of the unexplained evidence. This also appears to give a smoother, more robust estimate of model fitness than the explained evidence contribution. The important features to note about the shape of f_g are : i) the foreground contribution to the function has a sharp minimum at the correct model instance ii) conversely, the background contribution to the function has a much less pronounced maximum at the same position iii) the foreground model is calculated to be more probable than the background model for a small region around the correct answer.

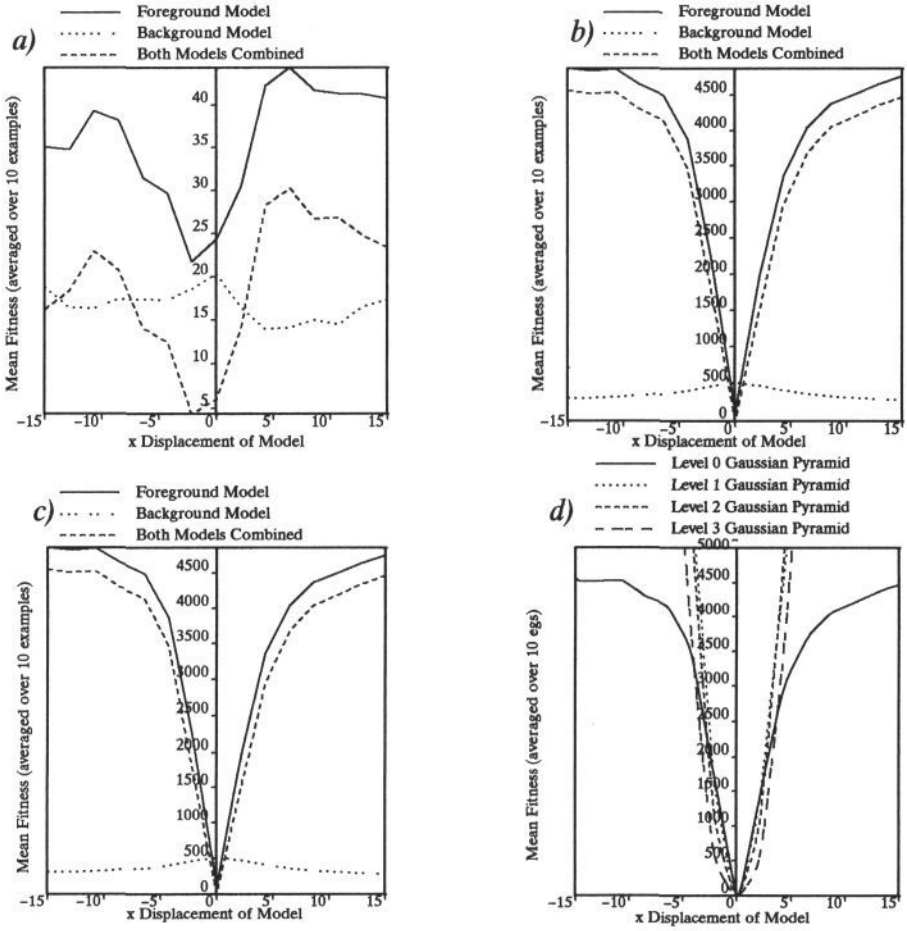

Figure 2: a) Mahalanobis Distance for Explained Evidence Model vs. x, b) Mahalanobis Distance for Unexplained Evidence Model vs. x, c) Combined Grey–Level Fitness Measure f_g vs. x d) f_g vs. x Plotted for 4 Levels of a Gaussian Image Pyramid

Varying Image Resolution

We also plot f_g as a function of x displacement for grey–level models generated at each level of a Gaussian image–pyramid. Figure 2 d) shows results for a 4–level pyramid. The specificity of the function at the correct solution increases as we increase the resolution of the model. We therefore have a method of controlling how strongly peaked our fitness measure is at the correct solution. We believe this can be used to improve the speed and robustness of our search method in two ways. Firstly, a highly smoothed lower–resolution image will be less noisy, giving a search–space less subject to local minima, and increasing the likelihood of the optimiser finding a global minimum. Se-

condly, as the image smoothing is increased, the fitness function becomes less like a δ–function in nature; this increases the speed of the optimiser.

Image Search

The current search method used for image search with PDMs is known as an Active Shape Model (ASM). This method finds the best fit of a PDM to an example image [6], given an initial estimate of model pose and shape. The model is gradually deformed to fit the observed data, and the global shape constraints of the model are enforced by ensuring that the model shape parameters remain within sensible limits. It has also been shown how models may be built of the expected grey–levels at each point in the PDM [5,7]. During ASM search, the suggested displacement for each model point is determined by finding the best local fit of its grey–level model to the image evidence.

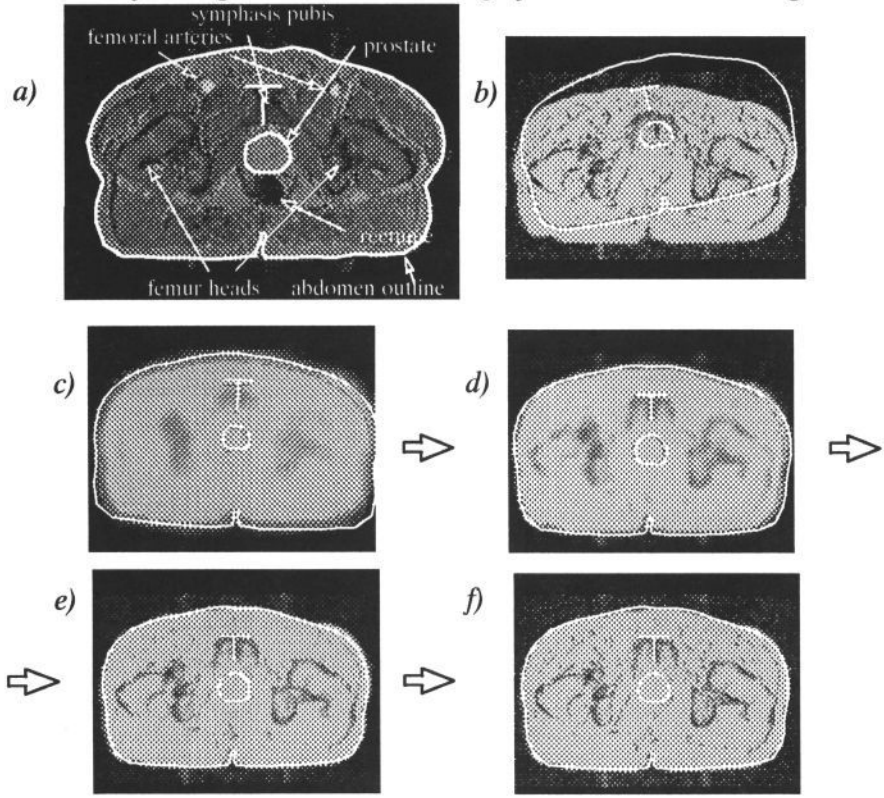

Figure 3 : a) Example Labelled Abdominal MR Image Showing Modelled Structures, b) Model Start Position Shown in Gaussian Pyramid Level 0 (ie original image), c)After Search in Gaussian Pyramid Image Level 3 (ie $^1/_{64}$ original area), d) After Search in Gaussian Pyramid Level 2 (ie $^1/_{16}$ original area), e) After Search in Gaussian Pyramid Level 1 (ie $^1/_4$ original area), f) After Search in Gaussian Pyramid Level 0 (ie original area)

This has been shown to improve the performance of the search method over simple edge–based techniques [8]. Finally, a multi–resolution version of the ASM with grey–level models has been implemented [9]; a grey–level model for each model point is trained at each level of a gaussian image pyramid. The search starts at the coarsest image resolution, and progresses by moving to finer and finer resolutions until it converges on a final solution. For comparison with our fitness measure, we may define an

overall fitness measure for an instance of the ASM as :

$$f_{ASM} = \sum_{i=1}^{nm} M_i \tag{25}$$

where M_i is the total Mahalanobis distance for the individual profile model describing the *foreground* grey-level data observed at model point i given by:

$$M_i \approx \sum_{j=1}^{t_i} \frac{b_{ij}^2}{\lambda_{ij}} + \frac{2}{\lambda_{t_i}} R_i^2 \tag{26}$$

and

$$R_i^2 = (\mathbf{F}_i^{obs} - \overline{\mathbf{F}}_i)^T (\mathbf{F}_i^{obs} - \overline{\mathbf{F}}_i) - \mathbf{b}_i^T \mathbf{b}_i \tag{27}$$

(other symbols have analogous meanings to those used previously; see [5] for further details).

We also wish to use an iterative method of finding a correct interpretation of a given image by minimising the fitness function. We cannot however use the ASM algorithm because the concatenated profile models we have used are not appropriate to that formulation. We use a quasi-Newtonian optimisation algorithm, (NAG routine E04JAF [10]) which makes no assumptions about the problem domain. Although the algorithm is about 50 times slower than the iterative approach of the Active Shape Model, the separation of fitness measure and search strategy makes the overall approach more generic, allowing it to be applied to any form of flexible template model. No attempt has been made to automate the search through scale-space. The optimisation is simply performed at each image pyramid resolution in turn, starting with the coarsest resolution.

We tested our method using the 2D trans-axial Magnetic Resonance images of the abdomen. Figure 3b)-f) shows the results of an example optimisation for an unseen example.

We performed a series of experiments comparing the Bayesian fitness measure and search strategy to that of a multi-resolution ASM which uses local grey-level models and an automatic scale-space search strategy. We tested each method on 20 training set images, for various initial model positions. Initially, the model was instantiated with its correct pose parameters, and with shape parameters set to their mean training set values. We then randomly displaced the model pose parameters from their correct values by the the following amounts :

x : by up to the amount given by the x ordinate of Figure 4
y : by up to the amount given by the x ordinate of Figure 4
scale : up to ±5%, (except for $\Delta x = 0$, when Δscale also $= 0$)
rotation : up to $\pm 0.5^c$, (except for $\Delta x = 0$, when Δscale also $= 0$)

Figure 4 shows the results obtained. The following observations may be made.

- Use of the quasi-Newtonian optimiser with the *original* Active Shape Model fitness measure significantly improves its search performance over that of the multi-resolution ASM.
- The Bayesian fitness measure combined with the quasi-Newtonian optimiser produces a still further improvement in search performance.
- The combination of quasi-Newtonian optimiser and unexplained foreground evidence fitness measure gives very similar results to those given by the overall fitness measure f. This is not surprising, as the unexplained evidence term dominates f, and the background model fitness term acts only to normalise the overall expression.

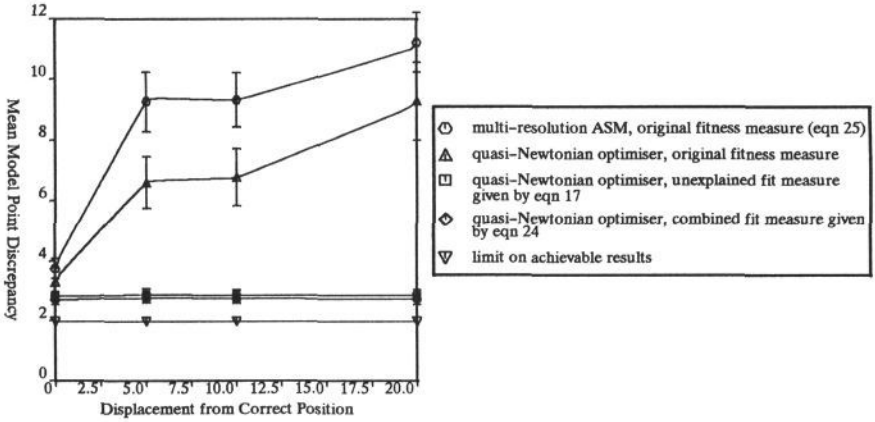

Figure 4 : Showing Search Results for Various Fitness Function / Optimisation Method Combinations for an Example PDM (error bars plotted as 1 standard error)

Conclusions

We have described a Bayesian probabilistic fitness measure which may be applied to flexible template models. The fitness measure, combined with a non-problem-dependent optimisation routine, has been shown to give a significant improvement in search performance for an example modelling technique (the Point Distribution Model) over the search method currently used (multi-resolution ASM). The measure attempts to combine shape and grey-level evidence in a truly probabilistic way, allowing the direct comparison of the fitness of different models. Further work may include automating the search through scale-space ie incorporating the model resolution as a parameter of the optimisation, and also improving the efficiency of our method.

References

[1] Lipson P., Yuille A.L., O'Keefe D., Cavanaugh J., Taffe J., Rosenthal D. Deformable Templates for Feature Extraction from Medical Images, *Proc. Euro. Conf. on Computer Vision,* (N. Faugeras Ed.), Springer Verlag, Berlin, 1990 pp 413–417

[2] Bajcsy R., Kovacic A. Multiresolution Elastic Matching, *Computer Vision, Graphics and Image Processing* 46, 1989, pp 1–21

[3] Staib L.H., Duncan J.S. Left Ventricular Analysis from Cardiac Images Using Deformable Models *Proc. Computers in Cardiology* 1989

[4] T.F.Cootes, C.J.Taylor, D.H.Cooper and J.Graham, Training Models of Shape from Sets of Examples. in *Proc. British Machine Vision Conference.* Springer-Verlag, 1992, pp.9–18.

[5] T.F.Cootes, C.J.Taylor, A.Lanitis,D.H.Cooper and J.Graham, Building and Using Flexible Models Incorporating Grey-Level Information, *Proc. International Conference on Computer Vision,* May 1993, pp.242–246.

[6] T.F.Cootes, C.J.Taylor, Active Shape Models - 'Smart Snakes'. in *Proc. British Machine Vision Conference.* Springer-Verlag, 1992, pp.266–275.

[7] T.F.Cootes, A. Hill, C.J.Taylor, J. Haslam, The Use of Active Shape Models for Locating Structures in Medical Images. *Proc. on Image Processing and Medical Imaging,* Arizona, June1993.

[8] T.F.Cootes, C.J.Taylor, Active Shape Model Search Using Local Grey-Level Models : A Quantitative Evaluation *accepted by British Machine Vision Conference,* April 1994.

[9] T.F.Cootes, C.J.Taylor, A. Lanitis, Multi-Resolution Search with Active Shape Models. *accepted by British Machine Vision Conference,* April 1994.

[10] P.E. Gill, W. Murray, Minimization Subject to Bounds on the Variables, *National Physical Laboratory Report NAC 72,* 1976.