A Generic System for Image Interpretation Using Flexible Templates

A. Hill, T. F. Cootes and C. J. Taylor

Department of Medical Biophysics, University of Manchester,
Oxford Road, Manchester M13 9PT, England

Abstract

We describe a generic approach to image interpretation, based on combining a general method of building flexible template models with Genetic Algorithm (GA) search. The method can be applied to a given image interpretation problem simply by training a Point Distribution Model (PDM), using a set of examples of the image structure to be located. A local optimisation technique, developed for use with PDMs, has been incorporated into the GA search with the aim of improving the speed of convergence and optimality of solution. We present results, from three practical applications, demonstrating that the new method offers significant improvements when compared to previously reported approaches to flexible template matching. The benefits include the ability to deal with different domains of application using a standard method, the ability to deal with complex multi–part models and improved search performance.

1 Introduction

Flexible templates have been employed widely as a means for model–based image interpretation. In most cases, however, the flexible template employed is *hand–crafted* and the search strategy used to locate instances of the template within a given image is problem specific [11,12,13,14]. In a previous publication [8] we described an approach to flexible template matching using a generic search strategy – Genetic Algorithms (GAs) [4,6,9]. The method generates good interpretation hypotheses robustly and at moderate computational cost. We presented results obtained using a system for automatically delineating the left ventricle of the heart in echocardiograms, employing a flexible template model of the left ventricle which was hand–crafted. We present here a development of that work which employs a generic technique for flexible template construction based on Point Distribution Models (PDMs) [2]. Several improvements to the original GA method result from the use of the PDM approach :

- The method is now entirely generic; it can be applied to new problems by training a PDM using examples of the image structures to be located.
- The ability of a PDM to capture both the variability in shape of an object and the spatial relationships between a number of different objects enables complicated biological structures to be modelled and subsequently located using a GA search.
- By incorporating into the GA a local optimisation technique developed for use with PDMs [3], significant improvements in the performance of the GA

search can be achieved both in terms of speed of convergence and optimality of solution; the ability of a GA search to extract multiple candidate interpretations from an image is also enhanced by the use of the local optimiser.

Our results demonstrate how complex biological structures can be automatically located using the combination of PDMs and a GA search. The delineation of the left ventricle in echocardiograms previously reported [8] has been enhanced by employing a more complex, PDM model of the heart. We also present results of locating the first and second ventricles of the brain in Magnetic Resonance images. The improved performance of the GA search when incorporating the PDM local optimiser is also demonstrated. The effect of the local optimiser when using a GA to extract multiple interpretation hypotheses is discussed and results are presented for locating simultaneously many resistors on a printed circuit board.

2 Genetic Algorithms

GAs employ mechanisms analogous to those involved in natural selection to conduct a search through a given parameter space for the global optimum of some objective function. The main features of the approach are as follows :

- A point in the search space is encoded as a *chromosome*.
- A *population* of N chromosomes/search points is maintained.
- New points are generated by probabilistically combining existing solutions.
- Optimal solutions are *evolved* by iteratively producing new *generations* of chromosomes using a *selective breeding* strategy based on the relative values of the objective function for the different members of the population.

A solution, $z = (z_1, z_2, ... z_n)$, is encoded as a string of *genes* to form a *chromosome* representing an *individual*. In many applications the gene values are [0,1] and the chromosomes are simply bit strings. An objective function, f, is supplied which can decode the chromosome and assign a *fitness value* to the individual a chromosome represents. In our case the z_i are model parameters which define the shape and pose of possible image objects. The objective function measures the extent to which the potential interpretation represented by a particular chromosome is supported by image evidence.

Given a population of chromosomes the genetic operators *crossover* and *mutation* can be applied in order to propagate *variation* within the population. Crossover takes two *parent* chromosomes, cuts them at some random gene/bit position and recombines the opposing sections to create two *children* e.g. crossing the chromosomes 010–11010 and 100–00101 at position 3–4 gives 010–00101 and 100–11010. Mutation is a background operator which selects a gene at random on a given individual and mutates the value for that gene (for bit strings the bit is complemented).

The search for an optimal solution starts with a randomly generated population of chromosomes; an iterative procedure is used to conduct the search. For each iteration a process of *selection* from the current *generation* of chromosomes is followed by application of the genetic operators and re-evaluation of the resulting

chromosomes. *Selection* allocates a number of trials to each individual according to its *relative fitness value* $f_i/\bar{f}$, $\bar{f} = 1/N \{f_1 + f_2 + .. + f_N\}$. The *fitter* an individual the more trials it will be allocated and vice versa. Average individuals are allocated a single trial.

Trials are conducted by applying the genetic operators (in particular crossover) to selected individuals, thus producing a new generation of chromosomes. The algorithm progresses by allocating, at each iteration, ever more trials to the high performance areas of the search space under the assumption that these areas are associated with short sub–sections of chromosomes which can be recombined using the random cut–and–mix of crossover to generate even better solutions.

The major feature of GAs which makes them attractive for object location is the use of a population of solutions, allowing competition between alternative interpretations. If there are several possible candidates, within an image, for the object we wish to locate, either because there are several instances of the object itself or instances of similar objects, the manner in which a GA search is conducted allows the various plausible interpretations to compete with one another, the strongest solution having the greatest probability of success. We have shown previously [8] how this facet of a GA search can be exploited to extract multiple plausible interpretations from an image by allowing separate *species* to adapt to various *niches* within the search space.

3 Point Distribution Models

We have employed the method described by Cootes et al [2] for constructing flexible templates. The technique captures the statistical variation in the distribution of sets of points to produce Point Distribution Models (PDMs). We describe the construction of PDMs for a chamber of the heart in echocardiograms and structures in the brain in Magnetic Resonance images. We also describe briefly *Active Shape Models* : instances of PDMs which deform to improve their fit to image data.

3.1 Constructing Point Distribution Models

The key steps in constructing a PDM from a set of examples of an object to be modelled are as follows :

- Generate a set of object descriptions upon which the model is to be trained. An object description is simply a labelled set of points $(x_0,y_0;\ x_1,y_1;\ ...)$. Each labelled point represents a particular position on the object (for example, the corner of a boundary). Corresponding points on different objects represent equivalent locations on each object.
- Align the sets of points and perform a Principle Components Analysis to the locations of the points. This involves finding the mean position of each point and the co–variance matrix of the position variables $(x_0,y_0;\ x_1,y_1;\ ...)$. The principle eigenvectors of the matrix give the main modes of variation of the training set.

This procedure results in a model with a small set of parameters $\mathbf{b} = (b_1, b_2, .., b_m)$ which act as weights for the major m eigenvectors of the co-variance matrix. These weights can be manipulated to create new instances from the class of objects modelled.

Because the technique is concerned only with the statistical variations of sets of points it can be employed not only to model objects which are, for example, closed boundaries but can also be used to capture the spatial relationships between objects. In this paper we have constructed PDMs of two complex biological structures : the left ventricle, septum and mitral valve of the heart (as imaged in echocardiography – see figure 3) and the first and second ventricles of the brain (as imaged in Magnetic Resonance Imaging – see figure 4). Both of these models exhibit complex structure (several parts) as well as variability in shape (see figure 1).
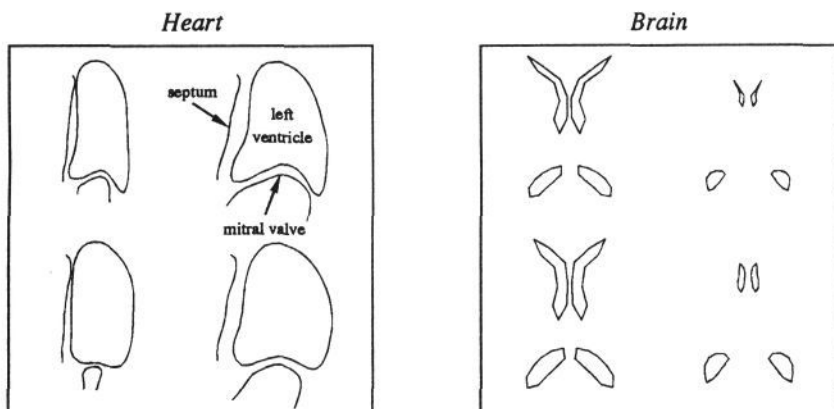

Figure 1 : Examples from Heart and Brain Ventricle Models.

For the PDM of the heart we employed echocardiogram time sequences from 33 individuals. From each of these sequences two images were selected which showed the left ventricle in its most contracted and extended states, giving a training set of 66 images. These images were labelled by an expert. For the ventricles of the brain we employed 3D Magnetic Resonance data sets from ten individuals. Because the ventricles of the brain are 3D structures we were able to employ several "slices" from each data set to give variability in shape due to slice position as well as variability in shape between individuals. On average, nine images were used from each sequence giving a total training set of 88 images.

3.2 Active Shape Models

Cootes and Taylor [3] have developed a local optimisation technique which can deform a PDM to fit image data. The technique is similar to that employed in the so called "Snake" approach presented by Kass et al [10] in that image evidence is used to suggest deformations of the model in order to improve the correspondence between model and data. An important property of the method presented by Cootes and Taylor, however, is that the shape constraints learned during the

training of the PDM are never violated during model deformation i.e. the shape of the model is always legal during the iterative deformation process.

Each iteration of the procedure is as follows :

- Place the current instance of the model onto the image and interrogate the image data locally to deduce the deformation of each point required to improve the correspondence between model and data.
- Compute changes in the translation (t_x, t_y) scale (s) and orientation (θ) of the model from these local deformations. Apply these changes to $(t_x, t_y), s, \theta$ and evaluate any residual deformations which remain.
- The residual deformations are now employed to suggest changes to the model parameters which control shape, **b**. The current model parameters **b** are updated to reflect these changes in shape.

By applying this procedure iteratively the position, scale, orientation and shape of the model which best fit the data in the locality of the original estimate can be determined.

4 Combining Active Shape Models and Genetic Algorithms

4.1 Possible Frameworks

There are two possible ways we might combine Active Shape Models (ASMs) and GAs:

- Consider the techniques as separate but complementary. Here the GA search would be conducted as normal and the ASM applied to the solutions suggested by the GA. This uses the ASM as a *refinement* procedure; if the GA suggests a solution in a non–optimal area of the search space, the ASM can do no better than locate the local optimum for that area of the search space.
- Incorporate the ASM directly into the GA search; in the GA literature it has been suggested that incorporating heuristic information and local optimisation techniques within a GA search can improve performance significantly (see the discussion of the Travelling Salesman Problem in [4]). The basis of this approach is that the GA can locate the *hills* in the search space while the local optimiser embedded within the GA can climb to the top of these hills.

We are particularly interested in the second of these approaches which allows a more thorough investigation of the search space, especially when a *speciated* version of the GA is employed.

4.2 Genetic Algorithms and Active Shape Mutation

In a single iteration of the ASM procedure, a set of local deformations are computed from the image data and applied to the model to generate suggested changes to $(t_x, t_y), s, \theta$, and **b** in order to improve the correspondence between model and data. When using a chromosomal representation of the parameters these changes can be realised by first decoding the chromosome to generate the current values

of $(t_x, t_y), s, \theta, \mathbf{b}$, then updating the parameter values and re–encoding the parameters to produce an updated chromosome. In effect, the genes on the chromosome are mutated to reflect the suggested changes in the parameters. This mutation will, in general, be beneficial to the individual concerned because the template which the chromosome encodes will represent a better fit to the image data; consequently the objective function value associated with the template will be improved. There are two parameters which control the incorporation of the ASM mutation into a GA :

- The rate at which the mutation is applied (M_{asm}); each individual in the current population is mutated with probability M_{asm}.
- The number of iterations of the ASM technique, M_i, to be applied for each mutation.

It will often be the case that the local deformations applied to the model can be generated as a by–product of evaluating the objective function which we are attempting to minimise/maximise i.e. when estimating how well the model fits the data, estimates of improved positions of model points suggest themselves naturally. This means that we can apply a single iteration of the ASM for "free" every time the objective function is evaluated – the only additional computing requirements being those to calculate the model transformation from the given local deformations. Taking this into consideration, together with the fact that the ASM mutation is generally beneficial, we suggest the values $M_{asm} = 1$ and $M_i = 1$ i.e. a single iteration of the local optimiser is applied every time the objective function is evaluated.

5 Results

The results we present below were obtained using a GA with "standard" parameter values [7] (unless otherwise stated): rate of crossover = 0.6, population size = 50, rate of random mutation = 0.005. The crossover operator employed was the constrained, two–point version suggested by Booker in [4]. The Remainder Stochastic Independent Sampling (RSIS) algorithm suggested by Baker [1] was used for selection. All model parameters were encoded as unsigned gray–code binary integers as suggested by Fitzpatrick et al [5]; 8 bits were used to represent each parameter. The same objective function was employed in all cases. The function was constructed in such a manner that it was minimised when strong edges of similar magnitude were located within the image close to the boundary of a given instance of a PDM (see [8] for more detail).

5.1 Improved Performance of GAs using ASM Mutations

In [8] we showed how GAs could be employed together with a flexible template model of the left ventricle of the heart to locate left ventricular boundaries in echocardiograms (see figure 3). We have compared the performance of GAs with and without ASM mutations using this exemplar. We applied a GA search to 5 echocar-

diagrams 10 times, starting from different (random) initial populations. Two measures were recorded at each iteration of the GA procedure and average values of these measures computed for all 50 applications. The measures used were :

- the value of the objective function averaged over the current population.
- the best objective function value in the current population.

On average the performance of the GA incorporating the ASM was significantly better than that of the GA without the ASM. Convergence was more rapid and the best solution found was also improved. The results are shown in figure 2.
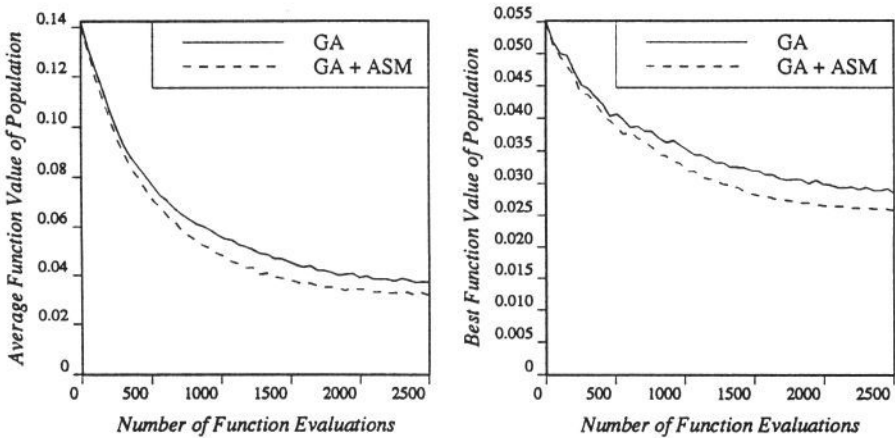

Figure 2 : *Comparison of GA and GA+ASM Searches.*

5.2 Objects with Multiple Parts

As we have already stated, employing the PDM method as a generic flexible template construction technique extends greatly the applicability of image interpretation using GA search. Objects with many variable parts can be modelled using a single PDM and the search technique applied as before. We have found that the increased specificity of such models can help resolve potentially ambiguous interpretations of image data. For example, in the model of the left ventricle of the heart we have included both sides of the septum and mitral valve in the model (see figure 1) placing stricter constraints on possible interpretations of the image data than if only the boundary of the left ventricle itself had been modelled. Figure 3 shows the automatic location of the septum, mitral valve and left ventricle in an apical 4–chamber echocardiogram using the GA+ASM technique. Figure 4 shows the automatic delineation of the first and second ventricles in a Magnetic Resonance image of the brain. The PDM of the ventricles of the brain was trained on images from different individuals and also from images at various positions in a 3D data set; the 2D shapes and locations of the ventricles vary considerably from slice to slice. The image shown was selected at random from a 3D data set and no indication of the "slice index" was employed in the interpretation process.
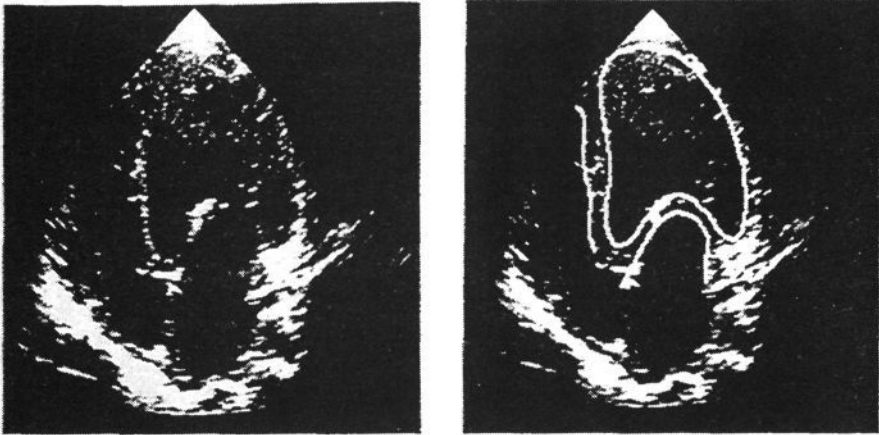
*Figure 3 : Automatic Delineation of the Left Ventricle,
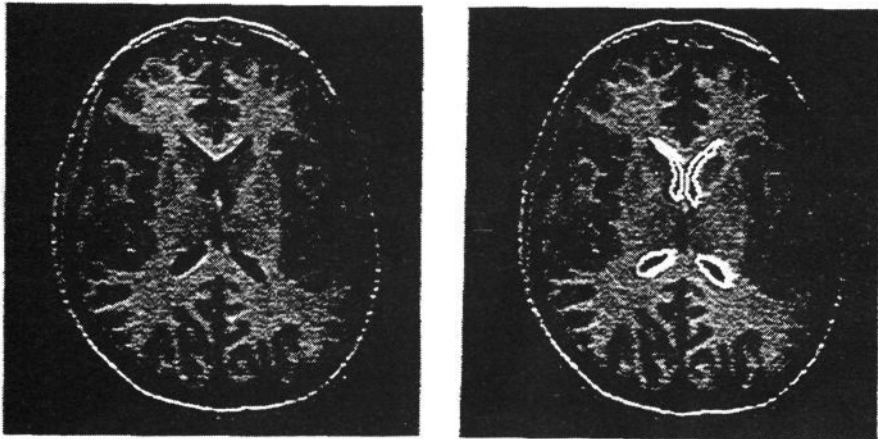Septum and Mitral Valve in the Heart.*

*Figure 4 : Automatic Delineation of the First and
Second Ventricles in the Brain.*

5.3 Extracting Multiple Plausible Interpretations using Speciation

In [8] we showed how a *speciated* version of the GA could be employed to extract multiple candidates for a given object within an image. In this case the GA is forced to spread its effort over different areas of the search space, rather than converging upon one particular area. This is accomplished by penalising individuals that reside in over–crowded areas of the search space in order to force migration to less crowded, yet still promising, areas. Speciation is also promoted by encouraging individuals to mate (crossover) with nearby rather than distant individuals. One problem with this approach is that adaptation of any particular species to the particular niche it occupies in the search space can be quite poor. This is because only a small gene pool is available for any given species due to the low numbers

of individuals belonging to each species (given a population of 100 individuals and just 3 or 4 possible interpretations, species might be represented by groups as small as 10 individuals). The danger here is that a species will become extinct in a promising area of the search space.

What is required is the ability to optimise locally with only a small number of individuals. This is exactly what the ASM mutation incorporated within the GA search can achieve. An example of the ability of the GA+ASM to maintain stable, small sub-populations is shown in figure 5 where a number of resistors have been located automatically on a printed circuit board using a PDM resistor model and a speciated GA+ASM search. The 7 resistors identified (5 "real" and 2 spurious) were located using a population size of 100 individuals and the 7 species shown were stable i.e. the GA maintained the sub-populations indefinitely without species becoming extinct.
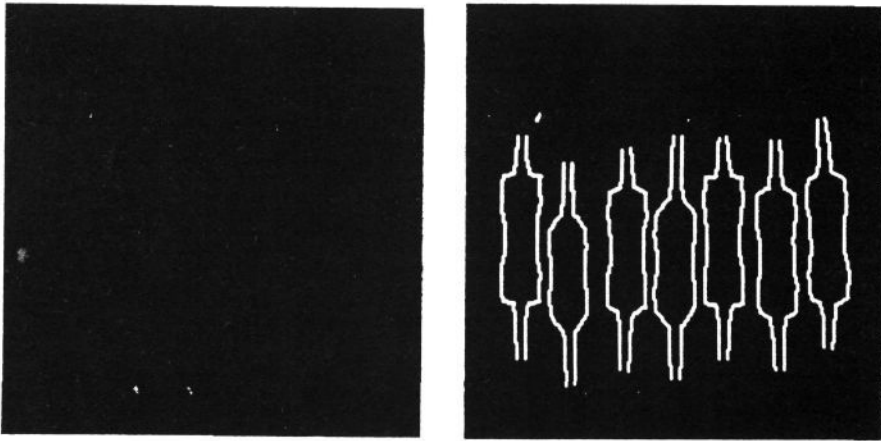

Figure 5 : Multiple Plausible Interpretations Employing a Speciated GA+ASM (the component on the left was excluded from the search area).

6 Conclusions

Combining a generic method of flexible template construction with the GA search technique has been shown to produce a very powerful and generally applicable method of model-based image interpretation. The ability of the PDM to capture both the variation in shape of single objects and the spatial relationships between different objects enables models of complex structures to be derived. We have shown that the GA search technique can be successfully employed to locate instances of these complex models within image data. Furthermore, the incorporation of a local optimisation technique for PDM models within the GA framework as an *Active Shape Mutation* improves considerably the performance of the GA search both in terms of speed of convergence and optimality of solution. The active shape mutation also improves the *speciated* version of the GA, in which many solutions are extracted simultaneously, by enabling small sub-populations to survive and improving the adaptation of each species to its particular environmental niche.

7 Acknowledgements

8 References

[1] J. E. Baker, Reducing Bias and Inefficiency in the Selection Algorithm, Proc. of the 2nd Int. Conf. on Genetic Algorithms, Hillsdale, NJ, USA, 1987; 14–21.

[2] T. F. Cootes, D. H. Cooper, C. J. Taylor, and J. Graham, Training Models of Shape from Sets of Examples, Proc. British Machine Vision Conference, Leeds, 1992.

[3] Cootes T. F., and Taylor C. J., Active Shape Models – Smart Snakes, Proc. British Machine Vision Conference, Leeds, 1992.

[4] L. Davis, Genetic Algorithms and Simulated Annealing, Pitman, London, 1987.

[5] J. M. Fitzpatrick, J. J. Grefenstette, and D. Van Gucht, Image Registration by Genetic Search, Proc. IEEE Southeastcon, Louisville 1984; 460–464.

[6] D. E. Goldberg, Genetic Algorithms in Search, Optimisation and Machine Learning, Addison–Wesley, 1989.

[7] J. J. Grefenstette, Optimisation of Control Parameters for Genetic Algorithms, IEEE Trans. on Systems, Man and Cybernetics 1986; 16(1):122–128.

[8] A. Hill, and C. J. Taylor, Model–Based Interpretation using Genetic Algorithms, Proc. British Machine Vision Conference, Glasgow, 1991, Springer Verlag; 266–274.

[9] J. H. Holland, Adaptation in Natural and Artificial Systems, University of Michegan Press, Ann Arbor, 1975.

[10] M. Kass, A. Witkin, and D. Terzopoulos, Snakes: Active Contour Models, 1st International Conference on Computer Vision, pub. IEEE Computer Society Press (1987) pp 259–268.

[11] P. Lilly, J. Jenkins, and P. Bourdillon, Automatic Contour Definition on Left Ventriculograms by Image Evidence and a Multiple Template–Based Model, IEEE Trans. on Medical Imaging (1989); 8(2):173–185.

[12] P. Lipson, A. L. Yuille, D. O'Keeffe, J. Cavanaugh, J. Taaffe, and D. Rosenthal, Deformable Templates for Feature Extraction from Medical Images, Proc. 1st European Conference on Computer Vision, Lecture Notes in Computer Science, Springer–Verlag, 1990; 413–417.

[13] L. H. Staib, and J. S. Duncan, Parametrically Deformable Contour Models, IEEE Computer Society Conference on Computer Vision and Pattern Recognition, San Diego, 1989.

[14] A. L. Yuille, D. S. Cohen, and P. Hallinan, Feature Extraction from Faces using Deformable Templates, Proc. Computer Vision, San Diego, 1989; 104–109.