

Model-Based Image Interpretation Using Genetic Algorithms

A. Hill and C. J. Taylor

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

Abstract

We describe the application of genetic algorithms in model-based image interpretation. The delineation of left ventricular boundaries in apical 4-chamber echocardiograms is used as an illustrative exemplar. The suitability of genetic algorithms for the model/objective-function/search procedure is presented.

1 Introduction

In model-based vision we generally wish to find the model-to-image transformation which explains some observed image. For example, for a 2D image of a known, rigid 3D object, the goal is to find the parameters (position and orientation) of the model which best explain the projection observed in the image. In general, additional parameters are required to account for variability in the object itself. A common approach to finding the parameter values is to extract image primitives and solve the combinatorial problem of establishing a correspondence between these and model primitives [1]. This assumes that a robust mechanism exists for extracting well-defined primitives, an assumption which often does not hold, particularly in fields such as medical image interpretation.

An alternative is to search through the space of possible parameter values, projecting model instances back into the image until one which is consistent with the observed image is found. Such an approach has often been applied starting from one or more approximate solutions but the initial approximations can be difficult to find when primitive extraction is problematic. In the absence of an initial approximation or cue, the alternative of a “blind” search through the parameter space seems unattractive because the search space can be huge (2^{57} possible states in the example we present below), however, the problem is well-matched to a class of optimisation methods known as Genetic Algorithms (GAs) [2,3,4] which can robustly find good solutions in large search spaces using very few trials. Given an objective function, f , which measures the evidential support for any particular projection into the image of the model, a GA search can find a set of parameters which provide a good explanation (or interpretation) of the image. Our results demonstrate the feasibility of this approach.

2 The Exemplar

We have evaluated the method using ultrasound images of the heart (apical 4-chamber echocardiograms). A typical example is shown in figure 1.a. The problem we address is that of locating the boundary of the left ventricle (LV) as shown in figure 1.b. The aim is to provide quantitative information concerning LV function, by analysing the motion of the LV boundary over a time sequence of images.




Figure 1.a : An apical 4-chamber echocardiogram




Figure 1.b : Associated LV boundary

The features of echocardiograms which make this a challenging exemplar are :

- Drop-out of the echo from some sections of the left ventricular wall.
- Considerable biological variation in LV size, shape and position.
- Occlusion of the LV boundary by other structures in the heart.
- Noise and artefacts.

There have been many attempts to automate the delineation of the LV [5,6] but nearly all rely on some operator intervention. Even so, such systems perform unreliably. Our objective is to achieve robust performance for images of clinical (poor) quality using a completely automatic approach.

3 Genetic Algorithms

GAs employ mechanisms analogous to those involved in natural selection to conduct a search through a given parameter space for the maximum/minimum 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, rather than a single point.
- New points in the search space 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.

Under the right circumstances such methods have been shown to converge to good solutions remarkably rapidly and have the advantage that the rate of convergence varies in accordance with the complexity of the search space, thus achieving robust performance over a range of conditions.


3.1 Representation and Evaluation of Solutions

A solution $\underline{x} = (x_1, x_2, \dots, x_n)$ is encoded as a string of *genes* to form a *chromosome* representing an *individual*. Each gene can take one of several values or *alleles*. Holland [2] showed that long chromosomes with few alleles per gene are preferable to shorter chromosomes with many alleles per gene. This implies that the optimal case is binary alleles. Consequently, in many applications the chromosomes are simply bit strings.

An objective/utility function, f , is supplied which can decode the chromosome and assign a *fitness value* to the individual the chromosome represents. In many applications the fitness value is simply the objective function evaluated at the point in the search space represented by the chromosome, $f(\underline{x})$.

3.2 The Genetic Operators

Given a population of chromosomes the genetic operators *crossover* and *mutation* can be applied. Crossover takes two *parent* chromosomes, cuts them at some random gene/bit position and recombines the opposing sections to create two *children*:



Mutation is a background operator which selects a gene at random on a given individual and mutates the allele for that gene (for bit strings the bit is complemented). Mutation is used to reintroduce alleles which may have been lost from the population for purely stochastic reasons.

3.3 Iterative Solution

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. *Selection* allocates a number of trials

to each individual according to its *relative fitness value* f_i/\bar{f} , $\bar{f} = \frac{1}{N} \sum_{i=1}^N f_i$. The *fitter*

an individual the more trials it will be allocated and vice versa. Average individuals are allocated only one trial. A trial is conducted by applying the genetic operators (in particular crossover) to selected individuals to produce a new generation of chromosomes.

The algorithm progresses by allocating, at each iteration, ever more trials to the high performance (better than average) areas of the search space under the assumption that these areas are associated with short sub-sections of chromosomes (hyperplanes/schemata/building blocks) which can be recombined using the random cut-and-mix of crossover to generate even better solutions.

4 The Model-Based Approach

The approach we adopt is to create a model of the LV which captures the variability in LV shape and allows missing parts of the boundary within the image to be inferred. The model is instantiated by choosing values for a set of 6 shape parameters ($\alpha, \beta, \delta_1, \delta_2, \delta_3, \delta_4$) (see figure 2) and 4 transformation parameters (r, θ, s, ϕ). An objective function, derived by considering edge evidence within the image along profiles constructed perpendicular to the candidate boundary, is employed to evaluate the degree to which image evidence supports any particular instantiation of the model.

4.1 The Model

LV shape (see figure 2) is defined as follows : The parameter α controls the width/height ratio. The co-ordinate of a control point representing the apex is thus given by $(-\alpha, 1)$ and the control point representing the attachment of the mitral valve to the septum (see figure 1.a) by $(-\alpha, -1)$. The angle of the mitral valve to the septum is governed by the parameter β which defines a vertical displacement from the position $(\alpha, -1)$ i.e. a control point is positioned at $(\alpha, -1 + \beta)$. The movement of the mitral valve is governed by the parameter δ_1 which describes the perpendicular displacement of a control point lying midway between the control points $(-\alpha, -1)$ and $(\alpha, -1 + \beta)$. In a similar manner δ_2 governs the position of the control point at the centre of the septum. The two remaining parameters δ_3, δ_4 dictate the position of the control points which modify the shape of the outside wall of the LV. $\delta_3\alpha$ defines a horizontal displacement from the control point $(-\alpha, 1)$ and δ_4 defines a vertical displacement from the control point $(\alpha, -1 + \beta)$.

For any instantiation of the shape model a continuous ventricular boundary can be obtained by first generating the control points in the manner described above and then employing a parametric cubic spline which interpolates these control points. Any number of points on the continuous boundary can be selected via the cubic spline.

The translation, scale and orientation of the model co-ordinate system with respect to the image are defined by $(r, \theta), s, \phi$ respectively. The translation is given in polar co-ordinates because the echocardiogram is a sector of a circle. All 10 parameters $\alpha, \beta, \delta_{1..4}, r, \theta, s, \phi$ are quantised appropriately - see table 1 (note that r and s are dimensionless $[0,1]$ variables i.e. for a square image of dimension L the actual values of r and s employed are rL and sL respectively).




Figure 2 : The model for the left ventricle.

Table 1 : Quantisation of model parameters.

Symbol	Min	Max	Bits	Symbol	Min	Max	Bits
r	0.2	0.475	7	β	-0.1	0.1	5
θ	$\pi/3$	$2\pi/3$	7	δ_1	-0.5	0	5
s	0.13	0.3	6	δ_2	0	0.35	5
ϕ	$-\pi/6$	$\pi/6$	6	δ_3	0.25	0.75	5
α	0.35	0.6	6	δ_4	0.1	1.0	5

4.2 The Objective Function

To perform the experiments described below we have used a simple and rather ad-hoc objective function. A number of points, P , on the continuous boundary of the candidate ventricle are selected for processing, $P = [P_\rho s (1 + a)]$ where $[\]$ indicates rounding to the nearest integer and P_ρ governs the density of points. Using $P_\rho = 200$ gives typically 30–100 points on the boundary. For each of these points a grey level profile perpendicular to the boundary at that point is extracted. For each profile the position (p_i , $p_{\min} \leq p_i \leq p_{\max}$) and strength (g_i) of the largest intensity step along the profile are recorded. Each profile is of length 20 pixels ($p_{\min} = -16$, $p_{\max} = 4$). The objective function is then given by :


$$f = \frac{1}{\sqrt{\bar{g}}} \frac{1}{P} \sum_{i=1}^P \left\{ \left| \frac{p_i}{p_{\min}} \right| + \left| \frac{g_i}{\bar{g}} - 1 \right| \right\} \quad \text{where} \quad \bar{g} = \frac{1}{P} \sum_{i=1}^P g_i$$

When we use the objective function to evaluate instantiations of the ventricle model we seek to minimise f , thus favouring solutions with strong edges (\bar{g} large) of equal magnitude ($|g_i/\bar{g} - 1| \rightarrow 0$) located close to the boundary position predicted by the model ($p_i \rightarrow 0$).

5 Implementation of Genetic Algorithm

5.1 Representation

As suggested by Fitzpatrick et al [7], the parameters of the model are encoded as unsigned gray-code binary integers. The number of bits encoding each parameter is given in table 1. For the translation vector (r, θ) it has proved useful to inter-leave the bits from each variable to form a single position vector variable with alleles 0..3, viz :



5.2 Search Parameters

We have employed the “standard” GA values (see Grefenstette [8]) of population size (N) = 50 and crossover rate (C) = 0.6 together with a mutation rate (M) of 0.005. In general, a fixed maximum number of objective function evaluations were allowed and this was set at 5000. The problem of unbiased trial allocation is discussed by Baker [9]. We have employed the Remainder Stochastic Independent Sampling (RSIS) algorithm as our selection strategy.

6 Results

We have acquired approximately 40 apical 4-chamber echocardiogram time sequences from several sources. All the images are of dimension 256^2 and pre-processed using a 3D grey-level morphological closing filter of dimensions $5 \times 5 \times 3$ (x, y and t respectively). The sequences exhibit all of the characteristics described in section 2. We have conducted two experimental studies; a comparison of different methods of optimisation (GAs and Simulated Annealing (SA) [3,10]) and a comparison of the performance of the model-based interpretation scheme described here with interpretations derived by an expert (manually drawn LV boundaries generated by a clinician).

For the first experiment a single image from each of five sequences was selected. These images were chosen on the basis that they were both difficult for an expert to interpret and also included characteristics which made the search space as difficult as possible (multiple optima, atypical LV shape). We applied both GA and SA to minimise the objective function described in section 4.2 for each of the five images and recorded the best objective function value located for each of 20 applications. Both methods of optimisation were allowed a maximum of 5000 function evaluations per application. Table 2 shows the two-tailed t-test results for the 20 sample distributions. It is clear from this table that, for the particular objective function, data and implementations described here, the performance of GA is significantly better than SA.

Table 2 : Comparison of different methods of optimisation.

Test Image	Genetic Algorithm		Simulated Annealing		t-test	
	mean f	σ of f	mean f	σ of f	t	signif.
1	0.1804	0.0165	0.1905	0.0129	2.2	0.037
2	0.1184	0.00551	0.1243	0.00447	3.7	0.001
3	0.1460	0.00685	0.1521	0.00646	2.9	0.006
4	0.1001	0.00477	0.1039	0.00447	2.6	0.013
5	0.1223	0.00891	0.1282	0.00805	2.2	0.034

For the second experiment we employed ten time sequences with boundaries drawn by an expert for the end systole and diastole frames i.e. two images per sequence giving a test suite of 20 images. It should be noted that several of these sequences were difficult to interpret, even for an expert. The images were interpreted using the model-based approach described here and the resulting LV boundaries compared with the expert-generated boundaries. The results are shown in table 3.

Table 3 : Comparison of model-based approach with expert.

x	Number of Images within $\pm x\%$			No. of images for Δ within x pixels [†]
	Area	Max Height	Max Width	
5	12	6	8	3 (x = 2)
10	14	11	13	14 (x = 4)
20	17	18	17	18 (x = 6)

In this table Δ is the mean distance between boundaries and is given by $\frac{1}{P} \sum_{i=1}^P \sqrt{(a_{i,x} - b_{j,x})^2 + (a_{i,y} - b_{j,y})^2}$ where $b_j = (b_{j,x}, b_{j,y})$ is the point on the expert boundary closest to the point on the GA generated boundary a_i . These results were produced by applying the GA search to each image once only and always from the same random set of starting points/chromosomes. A limit of 5000 function evaluations was again employed.

A "visual" comparison of the results revealed that the model-based approach failed to locate two out of the twenty left ventricles and obtained a less than acceptable interpretation for one other. The first two "failures" occur due to multiple plausible interpretations within the image and the third because of an abnormal ventricular shape (an aneurism). These incorrect interpretations are highlighted in the last row of table 3. Of the remaining 17 images, very plausible interpretations were obtained comparing favourably with those generated by an expert. We stress again the difficulty of the interpretation task for these images and the simplicity of the objective function we have employed.

[†]The mean absolute pixel height and width of the expert boundaries are 100 and 54 respectively.

7 Extensions – Niches and Species

If the GA could be employed to extract multiple optima, the functionality of the method would be greatly enhanced. The purpose of the search would no longer be to locate the single model instantiation most likely to be the LV in the current image, but would be to extract a handful of strong, ranked candidates for the LV. The problem of locating multiple optima when using GAs is discussed by Goldberg [4]. The approach adopted is to reduce the number of individuals in over-crowded areas of the search space by modifying their function values.

The major problem with this approach is how to decide how *close* individuals are to each other in order to determine crowding factors with which to modify the function values. We have implemented a crowding strategy which defines the closeness of individuals simply to be the 2D Euclidian distance separating the centres of the candidate ventricles within the image. The crowding factor (F) for any individual is simply the number of individuals that lie within a disc of a certain radius to the individual being considered. The modified function value is then given by $f(I - F/N)$ (when maximising) i.e. the greater the crowding, the worse the modified function values. A modified mating strategy has also been implemented in which close individuals are favoured as mates over distant individuals. This strategy has been found to be successful in maintaining separate sub-populations of individuals within the image which represent various candidate LVs.

As an example, the white dots in figure 3 represent the central co-ordinates of the individuals within the population after 40 generations; upon inspection it is clear that there are several separate *species* corresponding to the various cavities within the image.



Figure 3 : Species adapted to different chambers.

8 Conclusions

One of the major attributes of the model-based approach is the ability to correctly interpret incomplete and/or noisy image data by constraining all possible interpretations using knowledge represented by a model. In order for this process to be suc-

cessful it may be necessary to search a high-dimensional, non-linear, multi-modal and noisy search space. The chosen exemplar exhibits many of the problems frequently encountered within image interpretation (see section 2). The combination of a shape model which guarantees the feasibility of solutions, a naive objective function and a powerful search technique has been shown to yield very promising results.

The model-based approach described here would benefit greatly from a more generic approach to both model building (for generating candidate solutions) and evidential support mechanisms (the objective function). We are addressing both of these problems in our current work and progress in the former has already been made [11]. The sophistication of the objective function could be improved significantly by incorporating terms within the objective function which, rather than being ad-hoc as in section 4.2, are learned during the training phase of the model-building process.

9 Acknowledgements

Dr D. C. Beeton of Wythenshawe Hospital, Manchester for supplying some of the echocardiograms and Dr. K. Phadka of City General Hospital, Stoke-on-Trent for both supplying and labelling other sequences. This work was funded as part of IED project 1660. Dr. Hill is seconded from Active Memory Technology Ltd.

10 References

- 1 **Grimson W. E. L.** *Object Recognition by Computer : The Role of Geometric Constraints*, MIT Press, Cambridge, Massachusetts, 1990.
- 2 **Holland J. H.** *Adaptation in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor, 1975.
- 3 **Davis L.** *Genetic Algorithms and Simulated Annealing*, Pitman, London, 1987.
- 4 **Goldberg D. E.** *Genetic Algorithms in Search, Optimisation and Machine Learning*, Addison-Wesley, 1989.
- 5 **Chu C. H., Delp E. J., Buda A. J.** *Detecting Left Ventricular Endocardial and Epicardial Boundaries by Digital Two-Dimensional Echocardiography*, IEEE Trans. on Medical Imaging 1988; 7(2):81-90.
- 6 **Yamada H., Yamamoto K.** *Recognition of Echocardiograms by a Dynamic Programming Matching Method*, Pattern Recognition 1991; 24(2):147-155.
- 7 **Fitzpatrick J. M., Grefenstette J. J., Van Gucht D.** *Image Registration by Genetic Search*, Proc. IEEE Southeastcon, Louisville 1984; 460-464.
- 8 **Grefenstette J. J.** *Optimisation of Control Parameters for Genetic Algorithms*, IEEE Trans. on Systems, Man and Cybernetics 1986; 16(1):122-128.
- 9 **Baker J. E.** *Reducing Bias and Inefficiency in the Selection Algorithm*, Proc. of the 2nd Int. Conf. on Genetic Algorithms, Hillsdale, NJ, USA, 1987; 14-21.
- 10 **Kirkpatrick S., Gelatt C., Vecchi M.P.** *Optimisation by Simulated Annealing*, Science 1983; 220:671-680.
- 11 **Cootes T. F., Cooper D. H., Taylor C. J., Graham J.** *A Trainable Method of Parametric Shape Description*, Proc. British Machine Vision Conference, 1991.