Workshop on Non-Linear Model Based Image Analysis, NMBIA-98, University of Strathclyde, 1-3 July, 1998.

Image Segmentation & the use of Genetic Algorithms for Optimising Parameters and Algorithm choice

P.G.Ducksbury
Defence Evaluation & Research Agency
Malvern, UK

Abstract

One of the difficulties that has been apparent in applying image processing and understanding algorithms is that of the optimal choice of parameters and the algorithms themselves. Firstly we must select an algorithm and secondly the actual parameters that are required by that algorithm. It is also the case that using a chosen algorithm on a different image class yields results of a totally different quality, we have considered three image classes, namely infra-red linescan, Russian satellite and SPOT imagery. We have explored the use of genetic algorithms for the purpose of parameter and algorithm selection and will show how the approach can successfully obtain results which in the past have tended to be obtained somewhat heuristically. Once a reliable region has been obtained then we can represent its shape using a curvature scale space description.The main application of this work will be in the area of image databases.

1 Introduction

The automatic choice of optimal parameters and algorithms within image processing is a difficult area. Previous research [1, 2] has concentrated on using Bayesian methods (Pearl Bayes Networks) for combining information from feature detectors to obtain more reliable and consistent results. An alternative to actually combining the results of individual detectors is to explore the possibility of actually optimising the choice of those detectors and the associated parameters. We have explored the use of genetic algorithms for this purpose. Our research shows how the approach can successfully obtain results which in the past have tended to be obtained somewhat heuristically. For example, it is often the case that with inherited software and applications that it is known that 'algorithm W is good on infrared (IR) imagery with parameter set X', but 'algorithm Y is better on SPOT imagery with parameter set Z'. We show how a genetic algorithm is used to optimise both the parameter set for a given feature detector and the feature detector itself. One of the difficulties in a problem like this can be the size of the search space. Here however the population used for the genetic algorithm is reduced considerably by considering only a constrained space, for example a given parameter is likely only have a certain valid range of values (a low hysteresis threshold should be less than the high one).

2 Feature Detection

The detection of textured regions is achieved using three approaches, Belief Network [1], Fractal and Wavelet [3]. With the belief network the idea is to essentially combine a set of statistical evidence measures taken at multiple scales into a confidence (or belief) of some event occurring. Prior knowledge is required and this forms a set of multi-dimensional conditional probabilities which relate the inputs (causal information) to the outputs of a given node within the network. A set of rules are used which are based on the assumption that casual information that is tightly clustered should naturally generate a higher belief than that which is more dispersed. This is a reasonable assumption since if we consider segmentation of gray level imagery, then neighboring pixels from different regions which have intensity values lying at either end of the gray level spectrum have low probability and for those that belong to homogenous regions, high probability.

This now provides us with a set of feature detectors and their parameters to optimise. The next required stage prior to the GA is to produce some measure of confidence (cost function) as to how good the feature detectors are, this can be the most difficult stage as this is often a subjective task.

2.1 Confidence Estimation - Regions

In the case of region detection algorithms the boundary is taken and the skewness is computed for a line placed perpendicular to each boundary point. The skewness characterising the degree of asymmetry of a distribution around its mean. If a boundary was placed in the middle of a field as opposed to the edge of a region then we would naturally expect a different distribution and skewness measure. The sum of differences between the generated confidence outline and the similar for digital urban area outline data is taken as the confidence measure for a particular instance of an algorithm. Other methods have been proposed for confidence measures [4], alternatively a region either side of a segment of the boundary could have been analysed.

3 Genetic Algorithm

3.1 Search Strategies

Of the search and optimisation strategies that are available such as calculus based methods, enumerative search and random search we consider the genetic algorithm approach to be the more suitable albeit for this application.

Calculus based methods [5] search for local optima either indirectly by solving a set of equations which result from setting the gradient of an objective function to zero or directly by using the function and moving in a direction towards a local gradient referred to as either hill-climbing or steepest descent. The enumerative search methods essentially consider every point in the search space and evaluate the objective function at those points. They have the disadvantage of the actual size of the search space. For the same reason we also discounted the use of random search approaches, although these techniques do have their uses.

3.2 Constraints

The problem we are addressing is essentially a constrained optimisation one. Whenever a constraint is violated the solution would be infeasible and should therefore have no fitness value. In some highly constrained domains obtaining feasible points may prove to be a difficult task. Penalty type methods try to obtain information from infeasible solutions by the degrading of the fitness in relation to the relative degradation in the constraint violation. This approach is somewhat difficult in our domain as the objective function is actually an image processing or image understanding program (or even a suite of programs) and we may not be able to allow an infeasible member in the first place as the decoded set of parameters may have no sensible meaning to the algorithm. An interesting item for future work may be to incorporate sensitivity analysis in an attempt to gain information from infeasible points.

3.3 Parameter Set Encoding

For encoding the parameter set for these problems we use a simple intuitive and direct mapping of the parameters into substrings within a string (the analogy here being the genes within the chromosome). The parameter set for both of the problems is therefore encoded into a bit string with distinct substrings representing particular algorithms and parameters. The bits essentially represent the parameterisation of the problem, for example currently parameter one represents the choice between the algorithms. Different algorithms may result in differing string lengths, thus zero padding is included to the size of the maximum string.

For the region detection algorithms the parameters represent sampling window size, thresholds, filtersizes and the number of scales. Each of these parameters having a predefined dynamic range.

Given this encoding an initial population can be built by randomly generating bit strings. The constraints that are applied at this stage cause the removal of members (chromosomes) which have parameters (genes) outside the constrained search space as well as the removal of duplicate members. This results in a possibly highly constrained but feasible population. Given this initial population we can now perform the operations of crossover and mutation as described below.

3.4 Crossover

In performing crossover we are essentially playing a game of mix and match with members of the population that are highly correlated with previous success.

When it comes to the operations of crossover and mutation a variety of different approach's exist [5]. Our approach consists of a combination of the elitist strategy with a tournament selection, all members being chosen using a roulette wheel based selection strategy.

The elitist approach copies the 'best' solutions from one generation to the next thereby guaranteeing that the GA never loses its best solutions. A variety of thresholds being experimented for this with 0.05 being the chosen. The tournament selection chooses one member at random from the population as a parent and then chooses two others, these then essentially compete for the right to breed with the first chosen parent based on which of the two has the highest confidence measure. The roulette wheel selection is used whenever a member of the population is required to be chosen at random. Individual slots within the wheel are weighted in proportion to the members confidence measure (or fitness value). This being implemented along the lines of computing the cumulative histogram of the fitness values. A random point then indicates the location at which we take the first member for which the cumulative sum is greater.

The actual crossover of the selected parents is done by a simple exchange of substrings between the two, the location within the parents being randomly chosen. If the new members are either infeasible or duplicates, then they are rejected on the grounds that they would waste valuable population space and reduce the diversity. This crossover process continues to iterate until a new population has been generated. The initial population size is chosen according to a rule of thumb which chooses a population that is at least an order of magnitude greater than the number of parameters to optimise.

3.5 Mutation

The mutation operation again consists of choosing members at random using the roulette wheel selection and then mutating just a single bit within the members string chosen at random. The relative amount of mutation was set at 0.05 although different ratios where tried, larger amount having a slight detrimental effect on the overall population quality. Again new infeasible and duplicate members are rejected.

4 Shape description

Once a region is obtained we adopt the curvature scale space approach [6, 2] for forming a very concise description of the region which could be used for indexing and searching of an image database. This technique involves iteratively convolving the shape with gaussian filters of increasing σ until no more zero crossings of curvature are located. A plot of the contour position against scale space is then searched for the major peaks, its is these that are used as the shape descriptors for indexing and retrieval from a database.

5 Results

Figure 1 shows result for the region optimisation using three region finding algorithms. This problem was actually treated as a minimisation of the sum of differences of the outline with that of an urban area digital outline dataset. The top left image being the selection of the Belief Network algorithm (this algorithm was implemented from a cueing perspective and was not intended to produce a detailed outline). The top right image being the worst result obtained, in this case from the fractal algorithm with a poor choice of parameters in particular the thresholding parameter that is applied to the feature image in order to generate the outline. The bottom image is the curvature scale space image obtained from the top left outline by iteratively convolving with gaussians of increasing σ and tracking zero crossings.

We have shown that in principle a genetic algorithm is a very suitable search mechanism for optimising both the choice of parameters and algorithms for image processing operations. In fact any image processing/Understanding algorithm is suitable provided some confidence judgement can be made as to its effectiveness ([7] contains more details such as roads, regions, ship detection and registration). An alternative to its use here may be in situations like an algorithm that is composed of several sub-algorithms these themselves could be optimised more systematically. Another area that might be of interest for future work would be to examine the posibility of incorporating sensitivity analysis in an attempt to gain some information from infeasible points.

6 Acknowledgements

Some of this work has resulted from collaboration with the Universities of Oxford and Surrey. In particular we wish to thank Zi-Yan Xie at Oxford for software which implements the Fractal and Wavelet research and Farzin Mokhtarin at Surrey for a copy of the Curvature scale space software

References

[1] P. G. Ducksbury, Parallel Texture Region Segmentation using a Pearl Bayes Network, British Machine Vision Conference, BMVC-93, University of Surrey, Sep 1993, pp 187-196.

[2] P. G. Ducksbury, M. J. Varga, Region Based Image Content Descriptors and Representation, 6^{th} IEE Int. Conf. on Image Processing & its Applications, Trinity College, Dublin, July 1997.

[3] Zhi-Yan Xie, Multi-scale Analysis and Texture Segmentation, PhD Thesis, Dept of Eng Science, University of Oxford, 1994.

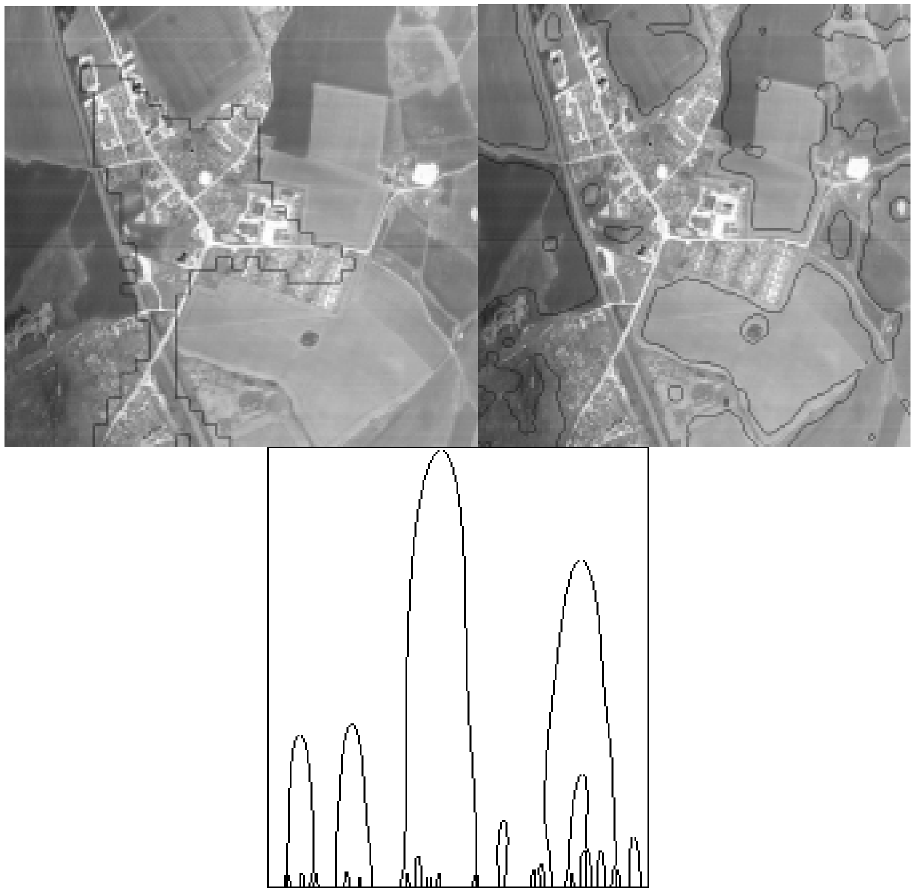

Figure 1: *Urban region optimisation, Top left - best result, Top right - worst result, Bottom - Curvature scale space image for top left region.*

[4] P. L. Palmer, H. Dabis, J. Kittler, A performance measure for boundary detection algorithms, Computer Vision and Image Understanding, Vol 63, No 3, pp 476-494, 1996.

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

[6] F. Mokhtarian, Silhouette-Based isolated object recognition through curvature scale space, IEEE Trans PAMI, vol 17, no 5, May 1995.

[7] P.G.Ducksbury, M.J. Varga, P.K. Kent, S. Foulkes, D.M. Booth, Genetic algorithms for automatic algorithm and parameter selection in ATR applications', SPIE Aerosense-98, Conf 3371, Orlando, 13-17th April, 1998.