

A Domain-Independent Window Approach to Multiclass Object Detection Using Genetic Programming

Mengjie Zhang

School of Mathematical and Computing Sciences, Victoria University of Wellington, P.O. Box 600, Wellington, New Zealand
Email: mengjie@mcs.vuw.ac.nz

Victor B. Ciesielski

School of Computer Science and Information Technology, RMIT University, GPO Box 2476v Melbourne, 3001 Victoria, Australia
Email: vc@cs.rmit.edu.au

Peter Andrae

School of Mathematical and Computing Sciences, Victoria University of Wellington, P.O. Box 600, Wellington, New Zealand
Email: pondy@mcs.vuw.ac.nz

Received 30 June 2002 and in revised form 7 March 2003

This paper describes a domain-independent approach to the use of genetic programming for object detection problems in which the locations of small objects of multiple classes in large images must be found. The evolved program is scanned over the large images to locate the objects of interest. The paper develops three terminal sets based on domain-independent pixel statistics and considers two different function sets. The fitness function is based on the detection rate and the false alarm rate. We have tested the method on three object detection problems of increasing difficulty. This work not only extends genetic programming to multiclass-object detection problems, but also shows how to use a single evolved genetic program for both object classification and localisation. The object classification map developed in this approach can be used as a general classification strategy in genetic programming for multiple-class classification problems.

Keywords and phrases: machine learning, neural networks, genetic algorithms, object recognition, target detection, computer vision.

1. INTRODUCTION

As more and more images are captured in electronic form, the need for programs which can find objects of interest in a database of images is increasing. For example, it may be necessary to find all tumors in a database of x-ray images, all cyclones in a database of satellite images, or a particular face in a database of photographs. The common characteristic of such problems can be phrased as “given $subimage_1, subimage_2, \dots, subimage_n$ which are examples of the objects of interest, find all images which contain this object and its location(s).” Figure 10 shows examples of problems of this kind. In the problem illustrated by Figure 10b, we want to find centers of all of the Australian 5-cent and 20-cent coins and determine whether the head or the tail side is up. Examples of other problems of this kind include target detection problems [1, 2, 3], where the task is to find, say, all tanks, trucks, or helicopters in an image. Unlike most of the cur-

rent work in the object recognition area, where the task is to detect only objects of one class [1, 4, 5], our objective is to detect objects from a number of classes.

Domain independence means that the same method will work unchanged on any problem, or at least on some range of problems. This is very difficult to achieve at the current state of the art in computer vision because most systems require careful analysis of the objects of interest and a determination of which features are likely to be useful for the detection task. Programs for extracting these features must then be coded or found in some feature library. Each new vision system must be handcrafted in this way. Our approach is to work from the raw pixels directly or to use easily computed pixel statistics such as the mean and variance of the pixels in a subimage and to evolve the programs needed for object detection.

Several approaches have been applied to automatic object detection and recognition problems. Typically, they use

multiple independent stages, such as preprocessing, edge detection, segmentation, feature extraction, and object classification [6, 7], which often results in some efficiency and effectiveness problems. The final results rely too much upon the results of earlier stages. If some objects are lost in one of the early stages, it is very difficult or impossible to recover them in the later stage. To avoid these disadvantages, this paper introduces a single-stage approach.

There have been a number of reports on the use of genetic programming (GP) in object detection and classification [8, 9]. Winkeler and Manjunath [10] describe a GP system for object detection in which the evolved functions operate directly on the pixel values. Teller and Veloso [11] describe a GP system and a face recognition application in which the evolved programs have a local indexed memory. All of these approaches are based on detecting one class of objects or two-class classification problems, that is, objects versus everything else. GP naturally lends itself to binary problems as a program output of less than 0 can be interpreted as one class and greater than or equal to 0 as the other class. It is not obvious how to use GP for more than two classes. The approach in this paper will focus on object detection problems in which a number of objects in more than two classes of interest need to be localised and classified.

1.1. Outline of the approach to object detection

A brief outline of the method is as follows.

- (1) Assemble a database of images in which the locations and classes of all of the objects of interest are manually determined. Split these images into a training set and a test set.
- (2) Determine an appropriate size ($n \times n$) of a square which will cover all single objects of interest to form the input field.
- (3) Invoke an evolutionary process with images in the training set to generate a program which can determine the class of an object in its input field.
- (4) Apply the generated program as a moving window template to the images in the test set and obtain the locations of all the objects of interest in each class. Calculate the detection rate (DR) and the false alarm rate (FAR) on the test set as the measure of performance.

1.2. Goals

The overall goal of this paper is to investigate a learning/adaptive, single-stage, and domain-independent approach to multiple-class object detection problems without any preprocessing, segmentation, and specific feature extraction. This approach is based on a GP technique. Rather than using specific image features, pixel statistics are used as inputs to the evolved programs. Specifically, the following questions will be explored on a sequence of detection problems of increasing difficulty to determine the strengths and limitations of the method.

- (i) What image features involving pixels and pixel statistics would make useful terminals?

- (ii) Will the 4 standard arithmetic operators be sufficient for the function set?
- (iii) How can the fitness function be constructed, given that there are multiple classes of interest?
- (iv) How will performance vary with increasing difficulty of image detection problems?
- (v) Will the performance be better than a neural network (NN) approach [12] on the same problems?

1.3. Structure

The remainder of this paper gives a brief literature survey, then describes the main components of this approach including the terminal set, the function set, and the fitness function. After describing the three image databases used here, we present the experimental results and compare them with an NN method. Finally, we analyse the results and the evolved programs and present our conclusions.

2. LITERATURE REVIEW

2.1. Object detection

The term *object detection* here refers to the detection of small objects in large images. This includes both *object classification* and *object localisation*. *Object classification* refers to the task of discriminating between images of different kinds of objects, where each image contains only one of the objects of interest. *Object localisation* refers to the task of identifying the positions of all objects of interest in a large image. The object detection problem is similar to the commonly used terms *automatic target recognition* and *automatic object recognition*.

We classify the existing object detection systems into three dimensions based on whether the approach is segmentation free or not, domain independent or specific, and on the number of object classes of interest in an image.

2.1.1 Segmentation-based versus single stage

According to the number of independent stages used in the detection procedure, we divide the detection methods into two categories.

(i) *Segmentation-based approach*, which uses multiple independent stages for object detection. Most research on object detection involves 4 stages: *preprocessing*, *segmentation*, *feature extraction*, and *classification* [13, 14, 15], as shown in Figure 1. The preprocessing stage aims to remove noise or enhance edges. In the segmentation stage, a number of coherent regions and “suspicious” regions which might contain objects are usually located and separated from the entire images. The feature extraction stage extracts domain-specific features from the segmented regions. Finally, the classification stage uses these features to distinguish the classes of the objects of interest. The algorithms or methods for these stages are generally domain specific. Learning paradigms, such as NNs and genetic algorithms/programming, have usually been applied to the classification stage. In general, each independent stage needs a program to fulfill that specific task and, accordingly, multiple programs are needed for object detection problems. Success at each stage is critical




FIGURE 1: A typical procedure for object detection.

to achieving good final detection performance. Detection of trucks and tanks in visible, multispectral infrared, and synthetic aperture radar images [2], and recognition of tanks in cluttered images [6] are two examples.

(ii) *Single-stage approach*, which uses only a single stage to detect the objects of interest in large images. There is only a single program produced for the whole object detection procedure. The major property of this approach is that it is segmentation free. Detecting tanks in infrared images [3] and detecting small targets in cluttered images [16] based on a single NN are examples of this approach.

While most recent work on object detection problems concentrates on the segmentation-based approach, this paper focuses on the single-stage approach.

2.1.2 Domain-specific approach versus domain-independent approach

In terms of the generalisation of the detection systems, there are two major approaches.

(i) *Domain-specific object detection*, which uses specific image features as inputs to the detector or classifier. These features, which are usually highly domain dependent, are extracted from entire images or segmented images. In a lentil grading and quality assessment system [17], for example, features such as brightness, colour, size, and perimeter are extracted and used as inputs to an NN classifier. This approach generally involves a time-consuming investigation of good features for a specific problem and a handcrafting of the corresponding feature extraction programs.

(ii) *Domain-independent object detection*, which usually uses the raw pixels directly (no features) as inputs to the detector or classifier. In this case, feature selection, extraction, and the handcrafting of corresponding programs can be completely removed. This approach usually needs learning and adaptive techniques to learn features for the detection task. Directly using raw image pixel data as input to NNs for detecting vehicles (tanks, trucks, cars, etc.) in infrared images [1] is such an example. However, long learning/evolution times are usually required due to the large number of pixels. Furthermore, the approach generally requires a large number of training examples [18]. A special case is to use a small number of domain-independent, pixel level features (referred to as *pixel statistics*) such as the *mean* and *variance* of some portions of an image [19].

2.1.3 Multiple class versus single class

Regarding the number of object classes of interest in an image, there are two main types of detection problems.

(i) *One-class object detection problem*, where there are multiple objects in each image, however they belong to a sin-

gle class. One special case in this category is that there is only one object of interest in each source image. In nature, these problems contain a binary classification problem: *object* versus *nonobject*, also called *object* versus *background*. Examples are detecting small targets in thermal infrared images [16] and detecting a particular face in photograph images [20].

(ii) *Multiple-class object detection problem*, where there are multiple object classes of interest, each of which has multiple objects in each image. Detection of handwritten digits in zip code images [21] is an example of this kind.

It is possible to view a multiclass problem as series of binary problems. A problem with objects 3 classes of interest can be implemented as class1 against everything else, class2 against everything else, and class 3 against everything else. However, these are not independent detectors as some methods of dealing with situations when two detectors report an object at the same location must be provided.

In general, multiple-class object detection problems are more difficult than one-class detection problems. This paper is focused on detecting multiple objects from a number of classes in a set of images, which is particularly difficult. Most research in object detection which has been done so far belongs to the one-class object detection problem.

2.2. Performance evaluation

In this paper, we use the DR and FAR to measure the performance of multiclass object detection problems. The DR refers to the number of small objects correctly reported by a detection system as a percentage of the total number of actual objects in the image(s). The FAR, also called false alarms per object or *false alarms/object* [16], refers to the number of nonobjects incorrectly reported as objects by a detection system as a percentage of the total number of actual objects in the image(s). Note that the DR is between 0 and 100%, while the FAR may be greater than 100% for difficult object detection problems.

The main goal of object detection is to obtain a high DR and a low FAR. There is, however, a trade-off between them for a detection system. Trying to improve the DR often results in an increase in the FAR, and vice versa. Detecting objects in images with very cluttered backgrounds is an extremely difficult problem where FARs of 200–2000% (i.e., the detection system suggests that there are 20 times as many objects as there really are) are common [5, 16].

Most research which has been done in this area so far only presents the results of the classification stage (only the final stage in Figure 1) and assumes that all other stages have been properly done. However, the results presented in this paper are the performance for the whole detection problem (both the localisation and the classification).

2.3. Related work—GP for object detection

Since the early 1990s, there has been only a small amount of work on applying GP techniques to object classification, object detection, and other vision problems. This, in part, reflects the fact that GP is a relatively young discipline compared with, say, NNs.

2.3.1 Object classification

Tackett [9, 22] uses GP to assign detected image features to a *target* or *nontarget* category. Seven primitive image features and twenty statistical features are extracted and used as the terminal set. The 4 standard arithmetic operators and a logic function are used as the function set. The fitness function is based on the classification result. The approach was tested on US Army NVEOD Terrain Board imagery, where vehicles, such as tanks, need to be classified. The GP method outperformed both an NN classifier and a binary tree classifier on the same data, producing lower rates of false positives for the same DRs.

Andre [23] uses GP to evolve functions that traverse an image, calling upon coevolved detectors in the form of hit-miss matrices to guide the search. These hit-miss matrices are evolved with a two-dimensional genetic algorithm. These evolved functions are used to discriminate between two letters or to recognise single digits.

Koza in [24, Chapter 15] uses a “turtle” to walk over a bitmap landscape. This bitmap is to be classified either as a letter “L,” a letter “I,” or neither of them. The turtle has access to the values of the pixels in the bitmap by moving over them and calling a detector primitive. The turtle uses a decision tree process, in conjunction with negative primitives, to walk over the bitmap and decide which category a particular landscape falls into. Using automatically defined functions as local detectors and a constrained syntactic structure, some perfect scoring classification programs were found. Further experiments showed that detectors can be made for different sizes and positions of letters, although each detector has to be specialised to a given combination of these factors.

Teller and Veloso [11] use a GP method based on the PADO language to perform face recognition tasks on a database of face images in which the evolved programs have a local indexed memory. The approach was tested on a discrimination task between 5 classes of images [25] and achieved up to 60% correct classification for images without noise.

Robinson and McIlroy [26] apply GP techniques to the problem of eye location in grey-level face images. The input data from the images is restricted to a 3000-pixel block around the location of the eyes in the face image. This approach produced promising results over a very small training set, up to 100% true positive detection with no false positives, on a three-image training set. Over larger sets, the GP approach performed less well however, and could not match the performance of NN techniques.

Winkeler and Manjunath [10] produce genetic programs to locate faces in images. Face samples are cut out and scaled, then preprocessed for feature extraction. The statis-

tics gleaned from these segments are used as terminals in GP which evolves an expression returning how likely a pixel is to be part of a face image. Separate experiments process the grey-scale image directly, using low-level image processing primitives and scale-space filters.

2.3.2 Object detection

All of the reported GP-based object detection approaches belong to the *one-class object detection* category. In these detection problems, there is only one object class of interest in the large images.

Howard et al. [19] present a GP approach to automatic detection of ships in low-resolution synthetic aperture radar imagery. A number of random integer/real constants and pixel statistics are used as terminals. The 4 arithmetic operators and min and max operators constitute the function set. The fitness is based on the number of the true positive and false positive objects detected by the evolved program. A two-stage evolution strategy was used in this approach. In the first stage, GP evolved a detector that could correctly distinguish the target (ship) pixels from the nontarget (ocean) pixels. The best detector was then applied to the entire image and produced a number of false alarms. In the second stage, a brand new run of GP was tasked to discriminate between the clear targets and the false alarms as identified in the first stage and another detector was generated. This two-stage process resulted in two detectors that were then fused using the min function. These two detectors return a real number, which if greater than zero denotes a ship pixel, and if zero or less denotes an ocean pixel. The approach was tested on images chosen from commercial SAR imagery, a set of 50 m and 100 m resolution images of the English Channel taken by the European Remote Sensing satellite. One of the 100 m resolution images was used for training, two for validation, and two for testing. The training was quite successful with perfect DR and no false alarms, while there was only one false positive in each of the two test images and the two validation images which contained 22, 22, 48, and 41 true objects.

Isaka [27] uses GP to locate mouth corners in small (50×40) images taken from images of faces. Processing each pixel independently using an approach based on relative intensities of surrounding pixels, the GP approach was shown to perform comparably to a template matching approach on the same data.

A list of object detection related work based on GP is shown in Table 1.

3. GP ADAPTED TO MULTICLASS OBJECT DETECTION

3.1. The GP system

In this section, we describe our approach to a GP system for multiple-class object detection problems. Figure 2 shows an overview of this approach, which has a learning process and a testing procedure. In the learning/evolutionary process, the evolved genetic programs use a square input field which is large enough to contain each of the objects of interest. The programs are applied in a moving window fashion to the

TABLE 1: Object detection-related work based on GP.

Problems	Applications	Authors	Year	Source
Object classification	Tank detection (classification)	Tackett	1993	[9]
		Tackett	1994	[22]
	Letter recognition	Andre	1994	[23]
		Koza	1994	[24]
	Face recognition	Teller and Veloso	1995	[11]
	Small target classification	Stanhope and Daida	1998	[28]
		Winkeler and Manjunath	1997	[10]
	Shape recognition	Teller and Veloso	1995	[25]
Object detection	Eye recognition	Robinson and McIlroy	1995	[26]
	Ship detection	Howard et al.	1999	[19]
	Mouth detection	Isaka	1997	[27]
	Small target detection	Benson	2000	[29]
	Vehicle detection	Howard et al.	2002	[30]
Other vision problems	Edge detection	Lucier et al.	1998	[31]
	San Mateo trail problem	Koza	1992	[32]
		Koza	1993	[33]
	Image analysis	Howard et al.	2001	[34]
		Poli	1996	[35]
	Model interpretation	Lindblad et al.	2002	[36]
	Stereoscopic vision	Graae et al.	2000	[37]
	Image compression	Nordin and Banzhaf	1996	[38]

entire images in the training set to detect the objects of interest. In the test procedure, the best evolved genetic program obtained in the learning process is then applied to the entire images in the test set to measure object detection performance.

The learning/evolutionary process in our GP approach is summarised as follows.

- (1) Initialise the population.
- (2) Repeat until a termination criterion is satisfied.
 - (2.1) Evaluate the individual programs in the current population. Assign a fitness to each program.
 - (2.2) Until the new population is fully created, repeat the following:
 - (i) select programs in the current generation;
 - (ii) perform genetic operators on the selected programs;
 - (iii) insert the result of the genetic operations into the new generation.
- (3) Present the best individual in the population as the output—the learned/evolved genetic program.

In this system, we used a tree-like program structure to represent genetic programs. The ramped half-and-half method was used for generating the programs in the initial population and for the mutation operator. The proportional selection mechanism and the reproduction, crossover, and mutation operators were used in the learning process.

In the remainder of this section, we address the other aspects of the learning/evolutionary system: (1) determination of the terminal set, (2) determination of the function set, (3) development of a classification strategy, (4) construction of the fitness measure, and (5) selection of the input parameters and determination of the termination strategy.

3.2. The terminal sets

For object detection problems, terminals generally correspond to image features. In our approach, we designed three different terminal sets: local rectilinear features, circular features, and “pixel features.” In all these cases, the features are statistical properties of regions of the image, and we refer to them as pixel statistics.

3.2.1 Terminal set I—rectilinear features

In the first terminal set, twenty pixel statistics, F_1 to F_{20} in Table 2, are extracted from the input field as shown in Figure 3. The input field must be sufficiently large to contain the biggest object and some background, yet small enough to include only a single object. In this way, the evolved program, as a detector, could automate the “human eye system” of identifying pixels/object centres which stand out from their local surroundings.

In Figure 3, the grey-filled circle denotes an object of interest and the square $A_1B_1C_1D_1$ represents the input field.

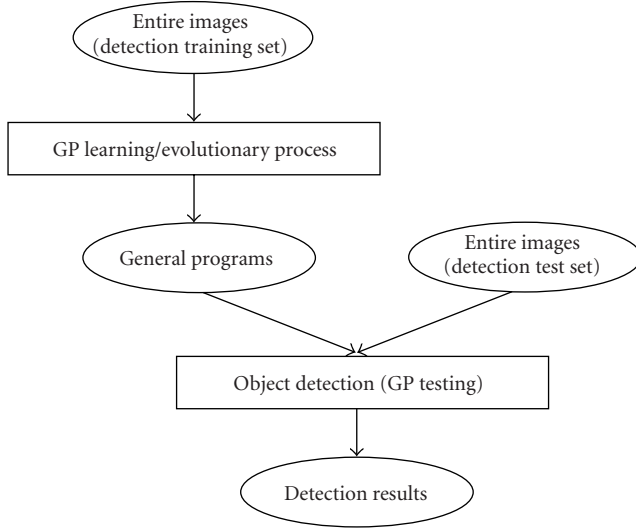


FIGURE 2: An overview of the GP approach for multiple-class object detection.

TABLE 2: Twenty pixel statistics. (SD: standard deviation.)

Pixel statistics		Regions and lines of interest
Mean	SD	
F_1	F_2	big square $A_1B_1C_1D_1$
F_3	F_4	small central square $A_2B_2C_2D_2$
F_5	F_6	upper left square $A_1E_1OG_1$
F_7	F_8	upper right square $E_1B_1H_1O$
F_9	F_{10}	lower left square $G_1OF_1D_1$
F_{11}	F_{12}	lower right square $OH_1C_1F_1$
F_{13}	F_{14}	central row of the big square G_1H_1
F_{15}	F_{16}	central column of the big square E_1F_1
F_{17}	F_{18}	central row of the small square G_2H_2
F_{19}	F_{20}	central column of the small square E_2F_2

The five smaller squares represent local regions from which pixel statistics will be computed. The 4 central lines (rows and columns) are also used for a similar purpose.¹ The mean and standard deviation of the pixels comprising each of these regions are used as two separate features. There are 6 regions giving 12 features, F_1 to F_{12} . We also use pixels along the main axes (4 lines) of the input field, giving features F_{13} to F_{20} .

In addition to these pixel statistics, we use a terminal which generates a random constant in the range $[0, 255]$. This corresponds to the range of pixel intensities in grey-level images.

These pixel statistics have the following characteristics.

- (i) They are symmetrical.

- (ii) Local regional features (from small squares and lines) are included. This assists the finding of object centres in the sweeping procedure—if the evolved program is considered as a moving window template, the match between the template and the subimage forming the input field will be better when the moving template is close to the centre of an object.
- (iii) They are domain-independent and easy to extract. These features belong to the pixel level and can be part of a domain-independent preexisting feature library of terminals from which the GP evolutionary process is expected to automatically learn and select only those relevant to a particular domain. This is quite different from the traditional image processing and computer vision approaches where the problem-specific features are often needed.
- (iv) The number of these features is fixed. In this approach, the number of features is always twenty no matter what size the input field is. This is particularly useful for the generalisation of the system implementation.

3.2.2 Terminal set II—circular features

The second terminal set is based on a number of circular features, as shown in Figure 4. The features were computed based on a series of concentric circles centred in the input field. This terminal set focused on boundaries rather than regions. The gap between the radii of two neighbouring circles is one pixel. For instance, if the input field is 19×19 pixels, then the number of central circles will be $19/2 + 1 = 10$ (the central pixel is considered as a circle with a zero radius); accordingly, there would be 20 features. Compared with the rectilinear terminal set, the number of these circular features in this terminal set depends on the size of the input field.

3.2.3 Terminal set III—pixels

The goal of this terminal set is to investigate the use of raw pixels as terminals in GP. To decrease the computation cost, we considered a 2×2 square, or 4 pixels, as a single pixel. The average value of the 4 pixels in the square was used as the value of this pixel, as shown in Figure 5.

3.3. The function sets

We used two different function sets in the experiments: 4 arithmetic operations only, and a combination of arithmetic and transcendental functions.

3.3.1 Function set I

In the first function set, the 4 standard arithmetic operations were used to form the nonterminal nodes:

$$\text{FuncSet1} = \{+, -, *, /\}. \quad (1)$$

The $+$, $-$, and $*$ operators have their usual meanings—addition, subtraction, and multiplication, while $/$ represents “protected” division which is the usual division operator

¹These lines can be considered special local regions. If the input field size n is an even number, each of these “lines” is a rectangle consisting of two rows or two columns of pixels.




FIGURE 3: The input field and the image regions and lines for feature selection in constructing terminals.

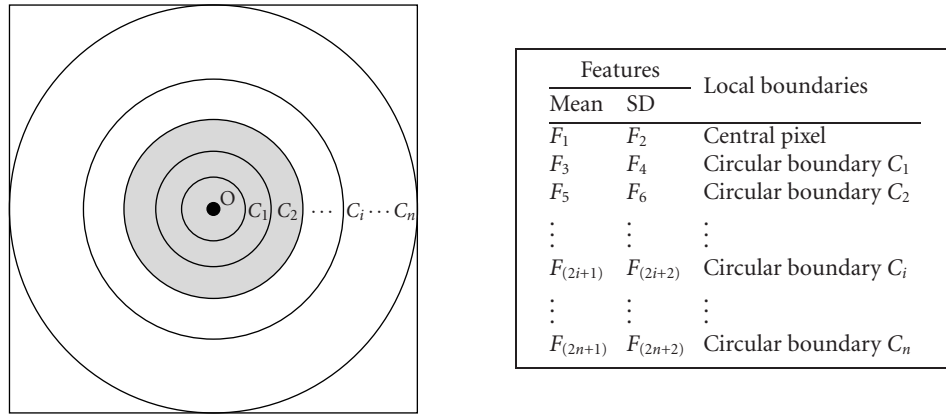


FIGURE 4: The input field and the image boundaries for feature extraction in constructing terminals.




FIGURE 5: Pixel terminals.

except that a divide by zero gives a result of zero. Each of these functions takes two arguments. This function set was designed to investigate whether the 4 standard arithmetic functions are sufficient for the multiple-class object detection problems.

A generated program consisting of the 4 functions and a number of rectilinear terminals is shown in Figure 6. The LISP form of this program is shown in Figure 7.

This program performed particularly well for the coin images.

3.3.2 Function set II

We also designed a second function set. We hypothesized that convergence might be quicker if the function values were close to the range $(-1, 1)$ and more functions might lead to better results if the 4 arithmetic functions were not sufficient. We introduced some transcendental functions, that is, the absolute function dabs, the trigonometric sine function sin, the logarithmic function log, and the exponent (to base e) function exp, to form the second function set:

$$\text{FuncSet2} = \{+, -, *, /, \text{dabs}, \sin, \log, \exp\}. \quad (2)$$

3.4. Object classification strategy

The output of a genetic program in a standard GP system is a floating point number. Genetic programs can be

$$\begin{aligned}
& \frac{F_{16}}{F_{14}} + F_5 + \frac{F_{14} \cdot F_{20}}{F_{11}} + F_{12} - F_{14} - (F_9 \cdot F_{11} \cdot F_1 \cdot F_{10} - F_9 \cdot F_{17}) \cdot \frac{F_5}{F_{18}} \\
& - \left[F_{17} + (F_{11} + F_{12}) \cdot F_{20} + \left(F_2 + 145.765 - \frac{F_6}{F_{11}} \right) \cdot (133.082 - F_{17}) \cdot \frac{F_{11}}{F_{14} \cdot F_{20}} \right] \\
& + \left[(F_6 - F_5 - F_3 \cdot F_6) \cdot \frac{F_1 + 145.765 + F_{16} \cdot F_{10}}{F_{18}} - F_{12} \right] \\
& \cdot [F_{17} + (F_{17} + F_{12}) \cdot F_{20} + F_{14} \cdot F_{12} \cdot (F_1 + F_{12} - F_{17})]
\end{aligned}$$

FIGURE 6: A generated program for the coin detection problem.

$$\begin{aligned}
& (+ (- (+ (+ (/ F_{16} F_{14}) F_5) (+ (/ (/ F_{11} (* F_{14} F_{20})) F_{11}) (- F_{12} \\
& F_{14}))) (- (* (- (* (* F_9 F_{11}) F_1) F_{10}) (* F_9 F_{17})) (/ F_5 F_{18})) (- \\
& (+ (+ F_{17} (* (+ F_{11} F_{12}) F_{20})) (* (- (+ F_2 145.765) (/ F_6 F_{11})) (- \\
& 133.082 F_{17}))) (/ F_{11} (* F_{14} F_{20})))) (* (- (* (- (- F_6 F_5) (* F_3 \\
& F_6)) (/ (+ (+ F_1 145.765) (* F_{16} F_{10})) F_{18})) F_{12}) (+ (+ F_{17} (* (+ F_{17} \\
& F_{12}) F_{20})) (* (+ F_{14} F_{12}) (- (+ F_1 F_{12}) F_{17}))))))
\end{aligned}$$

FIGURE 7: LISP format of the generated program in Figure 6.

used to perform one-class object detection tasks by utilising the division between negative and nonnegative numbers of a genetic program output. For example, negative numbers can correspond to the background and nonnegative numbers to the objects in the (single) class of interest. This is similar to binary classification problems in standard GP where the division between negative and nonnegative numbers acts as a natural boundary for a distinction between the two classes. Thus, genetic programs generated by the standard GP evolutionary process primarily have the ability to represent and process binary classification or one-class object detection tasks. However, for the multiple-class object detection problems described here, where more than two classes of objects of interest are involved, the standard GP classification strategy mentioned above cannot be applied.

In this approach, we develop a different strategy which uses a *program classification map*, as shown in Figure 8, for the multiple-class object detection problems. Based on the output of an evolved genetic program, this map can identify which class of the object located in the current input field belongs to. In this map, m refers to the number of object classes of interest, v is the output value of the evolved program, and T is a constant defined by the user, which plays a role of a threshold.

3.5. The fitness function

Since the goal of object detection is to achieve both a high DR and a low FAR, we should consider a multiobjective fitness function in our GP system for multiple-class object detection problems. In this approach, the fitness function is based on

a combination of the DR and the FAR on the images in the training set during the learning process. Figure 9 shows the object detection procedure and how the fitness of an evolved genetic program is obtained.

The fitness of a genetic program is obtained as follows.

- (1) Apply the program as a moving $n \times n$ window template (n is the size of the input field) to each of the training images and obtain the output value of the program at each possible window position. Label each window position with the “detected” object according to the object classification strategy described in Figure 8. Call this data structure a detection map. An object in a detection map is associated with a floating point program output.
- (2) Find the centres of *objects of interest only*. This is done as follows. Scan the detection map for an object of interest. When one is found, mark this point as the centre of the object and continue the scan $n/2$ pixels later in both horizontal and vertical directions.
- (3) Match these detected objects with the known locations of each of the desired true objects and their classes. A match is considered to occur if the detected object is within *tolerance* pixels of its known true location. A tolerance of 2 means that an object whose true location is (40, 40) would be counted as correctly located at (42, 38) but not at (43, 38). The *tolerance* is a constant parameter defined by the user.
- (4) Calculate the DR and the FAR of the evolved program.
- (5) Compute the fitness of the program as follows:

$$\text{fitness}(\text{FAR}, \text{DR}) = W_f \times \text{FAR} + W_d \times (1 - \text{DR}), \quad (3)$$

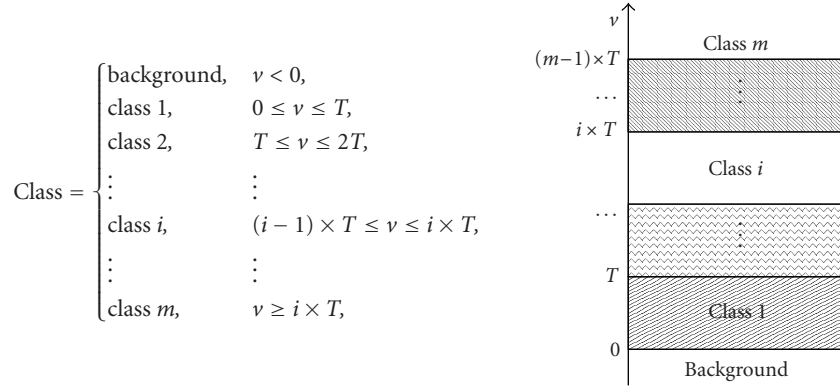


FIGURE 8: Mapping of program output to an object classification.

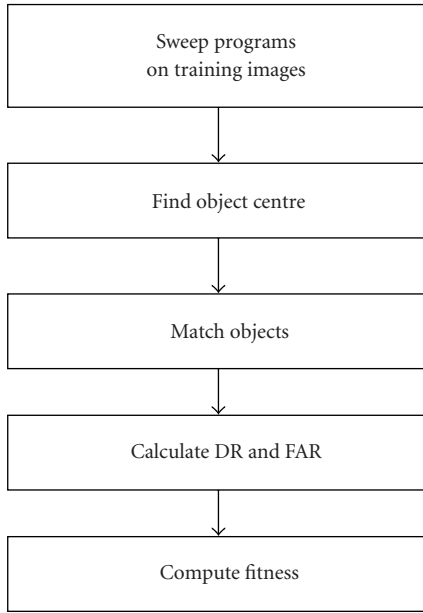


FIGURE 9: Object detection and fitness calculation.

where W_f and W_d are constant weights which reflect the relative importance of FAR versus DR.²

With this design, the smaller the fitness, the better the performance. Zero fitness is the ideal case, which corresponds to the situation in which all of the objects of interest in each class are correctly found by the evolved program without any false alarms.

3.6. Main parameters

Once a GP system has been created, one must choose a set of parameters for a run. Based on the roles they play in the learning/evolutionary process, we group these parameters

into three categories: search parameters, genetic parameters, and fitness parameters.

3.6.1 Search parameters

The search parameters used here include the number of individuals in the population (*population-size*), the maximum depth of the randomly generated programs in the initial population (*initial-max-depth*), the maximum depth permitted for programs resulting from crossover and mutation operations (*max-depth*), and the maximum generations the evolutionary process can run (*max-generations*). These parameters control the search space and when to stop the learning process. In theory, the larger these parameters, the more the chance of success. In practice, however, it is impossible to set them very large due to the limitations of the hardware and high cost of computation.

There is another search parameter, the size of the input field (*input-size*), which decides the size of the moving window in which a genetic program is computed in the program sweeping procedure.

3.6.2 Genetic parameters

The genetic parameters decide the number of genetic programs used/produced by different genetic operators in the mating pool to produce new programs in the next generation. These parameters include the percentage of the best individuals in the current population that are copied unchanged to the next generation (*reproduction-rate*), the percentage of individuals in the next generation that are to be produced by crossover (*cross-rate*), the percentage of individuals in the next generation that are to be produced by mutation (*mutation-rate* = 100% – *reproduction-rate* – *cross-rate*), the probability that, in a crossover operation, two terminals will be swapped (*cross-term*), and the probability that, in a crossover operation, random subtrees will be swapped (*cross-func* = 100% – *cross-term*).

3.6.3 Fitness parameters

The fitness parameters include a threshold parameter (T) in the object classification algorithm, a tolerance parameter

²Theoretically, W_f and W_d could be replaced by a single parameter since they have only one degree of freedom. However, the two cases of using a single and double parameters have different effects for stopping the evolutionary process. For convenience, we use two parameters.

TABLE 3: Parameters used for GP training for the three databases.

Parameter kinds	Parameter names	Easy images	Coin images	Retina images
Search parameters	Population-size	100	500	700
	Initial-max-depth	4	5	6
	Max-depth	8	12	20
	Max-generations	100	150	150
	Input-size	14×14	24×24	16×16
Genetic parameters	Reproduction-rate	10%	1%	2%
	Cross-rate	65%	74%	73%
	Mutation-rate	25%	25%	25%
	Cross-term	15%	15%	15%
	Cross-func	85%	85%	85%
Fitness parameters	T	100	100	100
	W_f	50	50	50
	W_d	1000	1000	3000
	Tolerance (pixels)	2	2	2

(*tolerance*) in object matching, and two constant weight parameters (W_f and W_d) reflecting the relative importance of the DR and the FAR in obtaining the fitness of a genetic program.

3.6.4 Parameter values

Good selection of these parameters is crucial to success. The parameter values can be very different for various object detection tasks. However, there does not seem to be a reliable way of *a priori* deciding these parameter values. To obtain good results, these parameter values were carefully chosen through an empirical search in experiments. Values used are shown in Table 3.

For detecting circles and squares in the easy images, for example, we set the population size to 100. On each iteration, 10 programs are created by reproduction, 65 programs by crossover, and 25 by mutation. Of the 65 crossover programs, 10 (15%) are generated by swapping terminals and 55 (85%) by swapping subtrees. The programs are randomly initialised with a maximum depth of 4 at the beginning and the depth can be increased to 8 during the evolutionary process. We also use 100, 50, 1000, and 2 as the constant parameters T , W_f , W_d , and *tolerance*, which are used for the program classification and the calculation of the fitness function. The maximum generation permitted for the evolutionary process is 100 for this detection problem. The size of the input field is the same as that used in the NN approach [12], that is, 14×14 .

3.7. Termination criteria

In this approach, the learning/evolutionary process is terminated when one of the following conditions is met.


- (i) The detection problem has been solved on the training set, that is, all objects in each class of interest in the training set have been correctly detected with no false

alarms. In this case, the fitness of the best individual program is zero.

- (ii) The number of generations reaches the predefined number, *max-generations*. *Max-generations* was determined empirically in a number of preliminary runs as a point before overtraining generally occurred. While it would have been possible to use a validation set to determine when to stop training, we have not done this. Comparison of training and test DRs and FARs indicated that overfitting was not significant.

4. THE IMAGE DATABASES

We used three different databases in the experiments. Example images and key characteristics are given in Figure 10. The databases were selected to provide detection problems of increasing difficulty. Database 1 (easy) was generated to give well-defined objects against a uniform background. The pixels of the objects were generated using a Gaussian generator with different means and variances for each class. There are three classes of small objects of interest in this database: black circles (*class1*), grey squares (*class2*), and white circles (*class3*). The Australian coin images (database 2) were intended to be somewhat harder and were taken with a CCD camera over a number of days with relatively similar illumination. In these images, the background varies slightly in different areas of the image and between images, and the objects to be detected are more complex, but still regular. There are 4 object classes of interest: the head side of 5-cent coins (class *head005*), the head side of 20-cent coins (class *head020*), the tail side of 5-cent coins (class *tail005*), and the tail side of 20-cent coins (class *tail020*). All the objects in each class have a similar size. They are located at arbitrary positions and with some rotations. The retina images (database 3) were taken by a professional photographer with special apparatus at a clinic and contain very irregular objects on a very



Number of images: 10
Object classes: 3
Image size 700×700

(a) Easy (circles and squares).



Number of images: 20
Object classes: 4
Image size 640×680

(b) Medium difficulty (coins).



Number of images: 15
Object classes: 2
Image size 1024×1024

(c) Very difficult (retinas).

FIGURE 10: Object detection problems of increasing difficulty.




FIGURE 11: An enlarged view of one piece of the retina images.

cluttered background. The objective is to find two classes of retinal pathologies—haemorrhages and microaneurisms. To give a clear view of representative samples of the target objects in the retina images, one sample piece of these images is presented in Figure 11. In this figure, haemorrhage and microaneurism examples are labeled using white surrounding squares.

5. EXPERIMENTAL RESULTS

We performed three groups of experiments, as shown in Table 4. The first group of experiments is based on the first two terminal sets (rectilinear features and circular features) and the first function set (the 4 standard arithmetic functions). The second group of experiments uses the third terminal set consisting of raw “pixel” and the first function set. The third group of experiments uses the first terminal set consisting of rectilinear features and the second function set consisting of additional transcendental functions.

TABLE 4: Three groups of experiments.

Experiments	Terminal sets	Function sets
I	TermSet1 (rectilinear)	FuncSet1
	TermSet2 (circular)	FuncSet1
II	TermSet3 (pixels)	FuncSet1
III	TermSet1 (rectilinear)	FuncSet2

In these experiments, 4 out of 10 images in the easy image database are used for training and 6 for testing. For the coin images, 10 out of 20 are used for training and 10 for testing. For the retina images, 10 are used for training and 5 for testing. The total number of objects is 300 for the easy image database, 400 for the Australian coin images, and 328 for the retina images. The results presented in this section were achieved by applying the evolved genetic programs to the images in the test sets.

5.1. Experiment I

This group constitutes the major part of the investigation. The main goal here is to investigate whether this GP approach can be applied to multiple-class object detection problems of increasing difficulty. The parameters used in these experiments are shown in Table 3 (Section 3.6.4). The average performance of the best 10 genetic programs (evolved from 10 runs) for the easy and the coin databases, and the average performance of the best 5 genetic programs (out of 5 runs, due to the high computational cost) for the retina images are presented.

The results are compared with those obtained using an NN approach for object detection on the same databases

[12, 39]. The NN method used was the same as the GP method shown in Section 1.1, except that the evolutionary process was replaced by a network training process in step (3) and the generated genetic program was replaced by a trained network. In this group of experiments, the networks also used the same set of pixel statistics as TermSet1 (rectilinear) as inputs. Considerable effort was expended in determining the best network architectures and training parameters. The results presented here are the best results achieved by the NNs and we believe that the comparison with the GP approach is a fair one.

5.1.1 Easy images

Table 5 shows the best results of the GP approach with the two different terminal sets (GP1 with TermSet1, GP2 with TermSet2) and the NN method for the easy images. For *class1* (black circles) and *class3* (grey circles), all the three methods achieved a 100% DR with no false alarms. For *class2* (grey squares), the two GP methods also achieved 100% DR with zero false alarms. However, the NN method had an FAR of 91.2% at a DR of 100%.

5.1.2 Coin images

Experiments with coin images gave similar results to the easy images. These are shown in Table 6. Detecting the heads and tails of 5 cents (class *head005*, *tail005*) appears to be relatively straight forward. All the three methods achieved a 100% DR without any false alarms. Detecting heads and tails of 20-cent coins (class *head020*, *tail020*) is more difficult. While the NN method resulted in many false alarms, the two GP methods had much better results. In particular, the GP1 method achieved the ideal results, that is, all the objects of interest were correctly detected without any false alarms for all the 4 object classes.

5.1.3 Retina images

The results for the retina images are summarised in Table 7. Compared with the results for the other image databases, these results are not satisfactory.³ However, the FAR is greatly improved over the NN method.

The results over the three databases show similar patterns: the GP-based method always gave a lower FAR than the NN approach for the same detection rate. While GP2 also gave the ideal results for the easy images, it produced a higher FAR on both the coin and the retina images than the GP1 method. This suggests that the local rectilinear features are more effective for these detection problems than the circular features.

5.1.4 Training times

We performed these experiments on a 4-processor ULTRA-SPARC4. The training times for the three databases are very

TABLE 5: Comparison of the object detection results for the easy images: the GP approaches versus the NN approach. (Input field size = 14×14 ; repetitions = 10.)

Easy images	Object classes		
	class1	class2	class3
Best detection rate (%)	100	100	100
False alarm rate (%)	NN	0	91.2
	GP1	0	0
	GP2	0	0

TABLE 6: Comparison of the object detection results for the coin images. The GP approaches versus the NN approach. (Input field size = 24×24 , repetitions = 10.)

Coin images	Object classes			
	head005	tail005	head020	tail020
Best detection rate (%)	100	100	100	100
False alarm rate (%)	NN	0	0	182
	GP1	0	0	0
	GP2	0	0	38.4

TABLE 7: Comparison of the object detection results for the retina images. The GP approaches versus the NN approach. (Input field size = 16×16 , repetitions = 5.)

Retina images	Object classes	
	Haem	Micro
Best detection rate (%)	73.91	100
False alarm rate (%)	NN	2859
	GP1	1357
	GP2	1857

different due to various degrees of difficulty of the detection problems. The average training times used in the GP evolutionary process (GP1) for the easy, the coin, and the retina images are 2 minutes, 36 hours, and 93 hours, respectively.⁴ This is much longer than the NN method, which took 2 minutes, 35 minutes, and 2 hours on average. However, the GP method gave much better detection results on all the three databases. This suggests that the GP method is particularly applicable to tasks where accuracy is the most important factor and training time is seen as relatively unimportant.

³With the current techniques applied in this area, detecting objects in images with a highly cluttered background is an extremely difficult problem [5, 16]. In fact, these results are quite competitive to other methods for very difficult detection problems. As a young discipline, it is quite promising for GP to achieve such results.

⁴Even if the training time for difficult problems is very long, the time spent on applying the learned genetic program to the test set is usually very short, say, from several seconds to about one minute.

TABLE 8: Results with the second function set.

	Easy images			Coin images				Retina images	
	Class1	Class2	Class3	Head005	Tail005	Head020	Tail020	Haem	Micro
Best detection rate (%)	100	100	100	100	100	100	100	73.91	100
False alarm rate (%)	0	0	0	0	0	0	0	1214	463

5.2. Experiment II

Instead of using rectilinear and circular features (pixel statistics) as in experiment I, experiment II directly uses the pixel values as terminals (the third terminal set). For the input field sizes of 14×14 , 24×24 , and 16×16 , for the easy, the coin, and the retina images, the number of terminals are 49 (7×7), 144 (12×12), and 64 (8×8), respectively. For the easy images, the learning took about 70 hours on a 4-processor ULTRA-SPARC4 machine to reach perfect detection performance on the training set and 78 generations were taken. The population size used was 1000, the maximum depth of the program was 30, the maximum initial depth 10, the maximum number of generations 100. For the coin images and the retina images, the situation was worse. Since a large number of terminals were used, the maximum depth of the program trees was increased to 50 for the coin images and 60 for the retina images. The population size for both databases used was 3000 with a maximum number of generations of 100. The evolutionary process took three weeks to complete 50 generations for the coin images and five weeks to complete 50 generations for the retina images. The best detection results were overall 22% FAR at a 100% DR for the coin images, and about 850% FAR at a DR of 100% for microaneurisms in the retina images.

While these results are worse than those obtained by the GP1 and GP2 using the rectilinear and circular features, they are still better than the NN approach. If we use a larger population (e.g., 10000 or 50000), a larger program size (e.g., 100), and a larger number of generations (e.g., 300), the results could be better according to our experience. While this is not possible to investigate with the current hardware we use, it shows a promising future direction with the improvement and development of more powerful hardware, for example, parallel or genetic hardware.

5.3. Experiment III

Instead of using the four standard arithmetic functions, this experiment focused on using the extended function set (FuncSet2), as shown in Section 3.3.2. The parameters shown in Table 3 (Section 3.6.4) were used in this experiment. The best detection results for the three databases are shown in Table 8.

As can be seen from Table 8, this function set also gave ideal results for the easy and the coin images and a better result for the retina images. The best DR for detecting *micro* is 100% with a corresponding FAR of 463%. The best DR for *haem* is still 73.91% but the FAR is reduced to 1214%. In

addition, convergence was slightly faster for training the coin and retina images. This suggests that dabs, sin, log, and exp are particularly useful for more difficult problems.

6. DISCUSSION

6.1. Analysis of results on the retina images

The GP-based approach achieved the ideal results on the easy images and the coin images, but resulted in some false alarms on the retina images, particularly for the detection of objects in class *haem* in which the FAR was very high and more than a quarter of the real objects in this class were not detected by the evolved genetic program.

We identified two possible reasons for the results on the retina images being worse than the results on the easy and the coin images. The first reason concerns the complexity of the background. In the easy and coin images, the background is relatively uniform, whereas in the retina images it is highly cluttered. In particular, the background of the retina images contains many objects, such as veins and other anatomical features, that are not members of the two classes of interest (microaneurisms and haemorrhages). These objects of noninterest must be classified as “background,” in just the same way as the genuine background. The more complex the boundary between classes in the input space, the more complex an evolved program has to be to distinguish the classes. It may be that the more complex background class in the retina images requires a more complex evolved program than the GP system was able to discover. It may even be that the set of terminals and functions is not adequate/sufficient to represent an evolved program to distinguish the objects of interest from such a rich background.

The second possible reason concerns the variation in size of the objects. In the easy and coin images, all of the objects in a class have similar sizes, whereas in the retina images, the sizes of the objects in each class vary. This variation means that the evolved genetic program must cover a more complicated region of the input space. The sizes of the *micro* objects vary from 3×3 to 5×5 pixels and the sizes of the *haem* objects vary from 6×6 to 14×14 pixels. Given the size of the input field (16×16) and the choice of terminals, the variance in the size of the *haem* objects is particularly problematic since it ranges from just one quarter of the input field (hence entirely inside the central detection region) to almost the entire input field. The fact that the performance on the *haem* class is worse than the performance on the *micro* class (especially in experiment III) provides

Program 1	$\frac{\frac{F_3 \cdot F_{14} \cdot F_{15}}{F_6 \cdot F_{14}} - F_{19} - F_7 \cdot F_{10} \cdot F_{17} \cdot F_{16} \cdot F_{18} + \frac{F_5}{F_5} \cdot F_{14} + \frac{F_3}{F_5}}{\frac{F_3}{F_5} \cdot \frac{F_6}{F_{11}} + \frac{F_{19}}{F_5} + \frac{F_6}{F_{15}}}$
Program 2	$\frac{F_{18}}{F_3 + F_5 - F_3 \cdot F_5} - F_4 \cdot F_{16} \cdot \frac{F_{18} \cdot (F_5 + F_{18}) - (F_7 + F_4) + F_{10} - F_{19}}{F_4 \cdot F_{16}}$
Program 3	$\frac{(F_{16} + F_7) \cdot F_{15} \cdot F_4}{F_{19} - \frac{F_{13} \cdot F_5}{F_{18}} + F_{11}} \cdot \frac{F_9}{F_9 \cdot F_4}$

FIGURE 12: Three sample generated programs for simple object detection in the easy images.

```

(/ (+ (- (- (/ (* (* F3 F14) F15) (* F6 F14)) F19) (* (* (* (* F7 F10) F17)
F16) F18)) (+ (* (/ F5 F5) F14) (/ F3 F5))) (+ (* (/ F3 F5) (/ (/ F5 F6) (/
F11 F15)))) (/ (/ F19 F6) (/ F5 F15))))

```

FIGURE 13: LISP format of Program 1.

additional evidence that the size variation is a cause of the poor performance.

The first reason suggests that the current approach is limited on images containing cluttered backgrounds. One possible modification to address this limitation is to evolve multiple programs rather than a single program, either having a separate program for each class of interest, or having several programs to exclude different parts of the background. Another possible modification is to extend the terminal set and/or function set to enrich the expressive power of the evolved programs.

The second reason suggests that the current approach has limited applicability to scale invariant detection problems. This would not be surprising, given the current set of terminals and functions. In particular, although the pixel statistics used in the rectilinear and circular terminal sets are robust to small variations in scale, they are not robust to large variations. We will explore alternative pixel statistics that are more robust to scale variations, and also function sets that would allow disjunctive programs that could better represent classes that contained objects of several different size ranges.

6.2. Analysis of evolved programs

This section gives a brief analysis of the best generated programs for the three databases. The genetic programs evolved by the GP1 in experiment I are used as examples.

6.2.1 Easy images

Figure 12 shows three good sample evolved programs for the easy images. (These programs were the direct mathematical

conversion of the original LISP format programs evolved by the evolutionary process. The LISP format of the first program is, for example, shown in Figure 13. Note that we did not simplify them—simplification of evolved genetic programs is beyond the goal of this paper.) All of these programs achieved the ideal results: all of the circles and squares were correctly detected with no false alarms.

There are several things we can note about these programs. Firstly, the programs are not trivial, and are decidedly nonlinear. It is hard to interpret these programs even for the easy images. Secondly, the programs use many, but not all, of the terminals, but do not use any constants. There are no groups of the terminals that are unused—both the means and standard deviations of both the square regions and the lines are used in the programs, so it does not appear that any of the terminals could be safely removed. Thirdly, although the programs are not in their simplest form (e.g., the factor F_5/F_5 could be removed from the first program), there is not a large amount of redundancy, so that the GP search is finding reasonably efficient programs.

6.2.2 Coin images

In addition to the program shown in Figure 6, we present another generated program in Figure 14, which also performed perfectly for the coin images.

Compared with those for the easy images, these programs are more complex, which reflects the greater difficulty of the detection problem in the coin images. One difference is that these programs also contain constants. The set of possible programs is considerably expanded by allowing constants as well as the terminals, but the search for good values for the

$$\frac{F_{10} \cdot F_{12} - F_9 - F_2 + F_{12} \cdot \frac{F_{10} \cdot F_{12} - F_9}{87.251} - \left[F_2 - \left(\frac{F_{12} \cdot F_{12} - F_{17} - F_2}{F_1} + \frac{87.251}{F_5} \right) \right]}{F_{11} \cdot \left(\frac{F_9}{F_{16}} - \frac{F_{17} - F_2 + F_{12} \cdot F_{12} - F_{11}}{F_{16}} \right) - F_{15} + F_8} \cdot \frac{F_{15} \cdot \frac{F_{15}}{F_8}}{F_{17} - F_2} + \frac{\frac{F_9}{F_{13} - F_{15}} + \frac{87.251}{F_{19}}}{F_5} + F_{10} \cdot F_{12} - F_9$$

FIGURE 14: A sample generated program for regular object detection in the coin images.

constants is difficult. Our current GP is biased so that constants are only introduced rarely, but it is clear that the detection problem on the coin images is sufficiently difficult to require some of these constants.

6.2.3 Retina images

One evolved genetic program for the retina images is presented in Figure 15. (The program is presented in LISP format rather than standard format because of its complexity.)

This program is much more complex than any of the programs for the easy and the coin images. The program uses all 20 terminals and 8 constants. It does not seem possible to make any meaningful interpretation of this program. It may be that with high-level, domain-specific features and domain-specific functions, it would be possible for the GP system to construct simpler and more interpretable programs; however, this would be against one of the goals of this paper which is to investigate domain-independent approaches.

Even the best programs for the retina images gave quite a high number of false alarms, and it appears that the 20 terminals and 4 standard arithmetic functions are not sufficient for constructing programs for such difficult detection problems. Nonetheless, the program above still had much better performance than an NN with the same input features.

6.3. Analysis of classification strategy

As described in Figure 8, we used a program classification map as the classification strategy. In this map, a constant T was used to give “fixed”-size ranges for determining the classes of those objects from the output of the program. The parameter can be regarded as a *threshold* or a class boundary parameter. Using just a single value for T forces most of the classes to have an equal possible range in the program output, which might lead to a relatively long time of evolution. A natural question to raise is whether we can replace the single parameter T with a set of parameters, say, T_1, T_2, \dots, T_m , one for each class of interest.

To answer this question, we ran a set of experiments on the easy images with three parameters, T_1, T_2 , and, T_3 , for the thresholds in the program classification map. The

experiments showed that some sets of values of the parameters resulted in an ideal performance but other sets of values did not. Also, the learning/evolutionary process converged very fast with some sets of values but very slowly with others. However, the results of the experiments gave no guidelines for selecting a good set of values for these parameters. In some cases, using separate parameters for each threshold may lead to a better performance than using a single parameter, but appropriate values for the parameters need to be empirically determined. In practice, this is difficult because there is no *a priori* knowledge in most cases for setting these parameters.

We also tried an alternative classification strategy, which we called *multiple binary map*, to classify multiple classes of objects. In this method, we convert a multiple-class classification problem to a set of binary classification problems. Given a problem L with m classes $L = \{c_1, c_2, \dots, c_m\}$, the problem is decomposed into $L_1 = \{c_1, \text{other}\}, L_2 = \{c_2, \text{other}\}, \dots, L_m = \{c_m, \text{other}\}$, where c_i denotes the i th class of interest and *other* refers to the class of nonobjects of interest. In this way, a multiple-class object detection problem is decomposed into a set of one-class object detection tasks, and GP is applied to each of the subsets to obtain the detection result for a particular class of interest. We tested this method on the detection problems in the three image databases and the results were similar to those of the original experiments.

One disadvantage of this method is that several genetic programs have to be evolved. On the other hand, the genetic programs may be simpler, which may reduce the training time for each program. In fact, for the coin images problem, a considerably shorter total training time was required to create a set of one-class programs than to create a single multiple-class program. A more detailed discussion of this method is outside the goal of this paper, and is left to future work.

6.4. Analysis of crossover and mutation rates

Some GP researchers argue that mutation is useless and should not be used in GP [32], while some others insist that a high mutation rate would help the GP evolution converge [40, 41]. To investigate the effects of mutation in GP for multiclass object detection problems, we carried out ten

```

(* (* (- (/ F6 (+ (* (/ (* F2 (/ (* F6 (+ F1 (- F10 F15)))
                (- (- F18 F17) (- F19 87.05))))
            (+ 17.0792 (+ F9 F14)))
    (/ (+ F19 (* (+ (+ F11
                    (- (* (- (- F15 F18) (+ 40.58 F16))
                        (- (* F13 (+ (/ 57.64 F16) F13))
                            (- F9 F6)))
                    (/ (* F3 F1) F1)))
        (* (- (* (- (/ (+ (+ F18 (+ (/ (/ F14 F6)
                                (+ F6 F1))
                                    89.70))
                            (* F10 F12)) F2) F9)
            (+ (+ F16 14.75) F9)) F18)
        (/ (/ F13 F1) (* (+ F6 F12) F9))))
    (+ F16 F8)))
    (+ (- (- (+ (/ F10 (* F9 F6)) F13) F10) F18)
        (+ (* (- (+ F1 F2) (+ F17 F8)) F5)
            (* (* F20 F16) F10))))
    (* (+ (- (* (+ F11
                (+ (* F14 F3)
                    (/ F15 (/ (+ (* F2 14.5251)
                                (* (* (/ (* F18
                                    (/ (* F2 F13) F15))
                                        F1)
                                        (/ (/ F11 F13) (/ F7 F5)))
                                    (+ (+ F18 (* F2 F13))
                                        (/ F8 F12))))
                                F17)))) F11) F16)
        (* (- F1 (+ F3 F8)) F5))
    (/ (+ (- F7 F20) F18) F20))))
    (* (* (* (* F2 F13) F2)
        (/ (* F4 (/ (* F2 F13) F15)) (* F18 F12)))
        (* F14 F2)))
    (+ (+ (- (+ (- F19 F3) F2) F7) (- (+ F8 F17) F18))
        (/ (+ F15 60.10)
            (* (* F1 (/ (/ F12 (- (+ (/ (/ F12 F13) (/ F15 F5)) F17) F18))
                (/ F7 F5))) F8))))
    (* (/ (* F10 (/ (* F2 F13) F15)) F18)
        (* (* (* F2 F2) (/ (/ F18 (+ F1 F2)) F13)
            (/ (/ (- F15 96.16) (* F4 14.53)) F5))) F4)
    (/ (/ F12 F13) (/ F1 (+ (/ F10 F1) F4))))))

```

FIGURE 15: A sample generated program for very difficult detection problems in the retina images.

experiments for different rates of mutation versus crossover on the easy images, as shown in Figure 16. The reproduction rate was held constant at 10%, and the mutation rate varied from 0% to 40%. The graph shows the distribution of the number of generations to convergence by a box-and-whisker plot with the limits of the central box at the 30% and 70% percentiles. With both 0% and 40% mutation, the search sometimes did not converge within the limit of 250 generations. There was a clear effect of the mutation rate on the number of generations to convergence. The best mutation rate was 25%, where only 48 generations on average were required to find a good solution, with slower convergence at

both lower and higher mutation rates. Experiments on the coin and the retina images gave a similar trend. This suggests that, in GP for multiple-class object detection problems described in this paper, mutation plays an important role for keeping the diversity of the population, and that convergence could be sped up when an appropriate mutation rate was used. However, such a good mutation rate is generally task dependent, and 15%–30% is a good choice for similar tasks.

6.5. Analysis of reproduction

In early GP, the reproduction rule did a probabilistic selection of genetic programs from the current population based

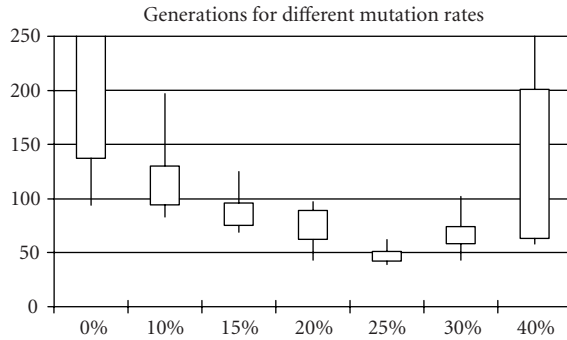


FIGURE 16: Convergence versus mutation rate.

on their fitness and allowed them to survive by copying them into the new population. The better the fitness, the more likely the individual program is to be selected [24, 42]. However, this mechanism does not guarantee that the best program will survive. An alternative reproduction rule is one that removes the probabilistic element, and simply reproduces the best n genetic programs from the current population. We ran experiments on the easy images with both reproduction rules and plotted the best fitness in each generation (see Figure 17). The dotted curve shows the best fitness with the probabilistic reproduction rule. Over the 100 generations, there are 4 clear intervals (at generation 7, 22, 45, and 67) where the fitness got worse rather than better, which delayed the convergence of learning. In contrast, the deterministic reproduction rule had a steady improvement in fitness. Furthermore, the deterministic reproduction rule converged on an ideal program after just 71 generations, while the probabilistic reproduction rule had still not converged on an ideal program after 100 generations. (In fact, the fitness did not improve at all during the final 30 generations!) Clearly, the new reproduction rule greatly improved the training speed and convergence.

7. CONCLUSIONS

The goal of this paper was to develop a domain-independent, learning/adaptive approach for detecting small objects of multiple classes in large images based on GP. This goal was achieved by the use of GP with a set of domain-independent pixel statistics as terminals, a number of standard operators as functions, and a linear combination of the DR and FAR as the fitness measure. A secondary goal was to compare the performance of this method with an NN method. Here the GP approach outperformed the NN approach in terms of detection accuracy.

The approach appears to be applicable to detection problems of varying difficulty as long as the objects are approximately the same size and the background is not too cluttered.

The paper differs from most work in object detection in two ways. Most work addresses the one-class problem, that is, *object* versus *nonobject*, or *object* versus *background*. This paper has shown a way of solving a multiple-class object detection problem without breaking it into a collection

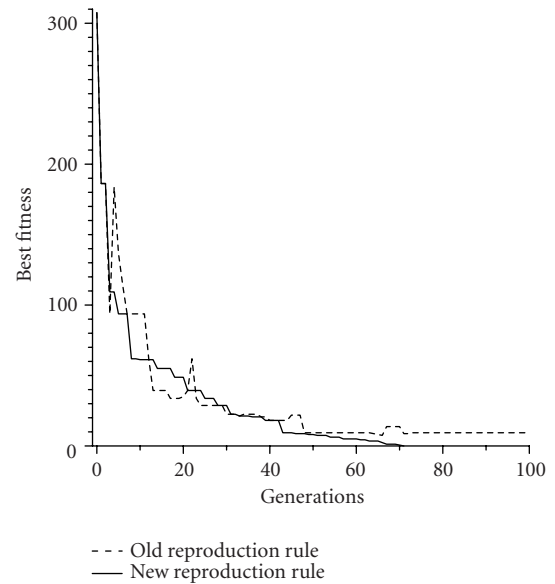


FIGURE 17: Training easy images based on the old and the new reproduction rules.

of one-class problems. Also, most current research uses different algorithms in multiple independent stages to solve the localisation problem and the classification problem; in contrast, this paper uses a single learned genetic program for both object classification and object localisation.

The experiments showed that mutation does play an important role in the three multiple-class object detection tasks. This is in contrast to Koza's early claim that GP does not need mutation. For GP applied to multiple-class object detection problems, the experiments suggest that a 15%–30% mutation rate would be a good choice.

The experiments also identified some limitations of the particular approach taken in the paper. The first limitation concerns the choice of input features and the function set. For the simple and medium-difficulty object detection problems, the 20 regional/rectilinear features and 4 standard arithmetic functions performed very well; however, they were not adequate for the most difficult object detection task. In particular, they were not adequate for detecting classes of objects with a range of sizes. Further work will be required to discover more effective domain-independent features and function sets, especially ones that provide some size invariance.

A second limitation is the high training time required. One aspect of this training time is the experimentation required to find good values of the various parameters for each different problem. The GP method appears to be applicable to multiple-class object detection tasks where accuracy is the most important factor and training time is seen as relatively unimportant, as is the case in most industrial applications. Further experimentation may reveal more effective ways of determining parameters which will reduce the training times.

Subject to these limitations, the paper has demonstrated that GP can be used effectively for the multiple-class

detection problem and provides more evidence that GP has a great potential for application to a variety of difficult problems in the real world.

ACKNOWLEDGMENTS

We would like to thank Dr. James Thom at RMIT University and Dr. Zhi-Qiang Liu at the University of Melbourne for a number of useful discussions. Thanks also to Peter Wilson whose basic GP package was used in this project and to Chris Kamusinski who provided and labelled the retina images.

REFERENCES

- [1] P. D. Gader, J. R. Miramonti, Y. Won, and P. Coffield, "Segmentation free shared weight networks for automatic vehicle detection," *Neural Networks*, vol. 8, no. 9, pp. 1457–1473, 1995.
- [2] A. M. Waxman, M. C. Seibert, A. Gove, et al., "Neural processing of targets in visible, multispectral IR and SAR imagery," *Neural Networks*, vol. 8, no. 7-8, pp. 1029–1051, 1995.
- [3] Y. Won, P. D. Gader, and P. C. Coffield, "Morphological shared-weight networks with applications to automatic target recognition," *IEEE Transactions on Neural Networks*, vol. 8, no. 5, pp. 1195–1203, 1997.
- [4] H. L. Roitblat, W. W. L. Au, P. E. Nachtigall, R. Shizumura, and G. Moons, "Sonar recognition of targets embedded in sediment," *Neural Networks*, vol. 8, no. 7-8, pp. 1263–1273, 1995.
- [5] M. W. Roth, "Survey of neural network technology for automatic target recognition," *IEEE Transactions on Neural Networks*, vol. 1, no. 1, pp. 28–43, 1990.
- [6] D. P. Casasent and L. M. Neiberg, "Classifier and shift-invariant automatic target recognition neural networks," *Neural Networks*, vol. 8, no. 7-8, pp. 1117–1129, 1995.
- [7] S. K. Rogers, J. M. Colombi, C. E. Martin, et al., "Neural networks for automatic target recognition," *Neural Networks*, vol. 8, no. 7-8, pp. 1153–1184, 1995.
- [8] J. R. Sherrah, R. E. Bogner, and A. Bouzerdoum, "The evolutionary pre-processor: automatic feature extraction for supervised classification using genetic programming," in *Proc. 2nd Annual Conference on Genetic Programming (GP-97)*, J. R. Koza, K. Deb, M. Dorigo, et al., Eds., pp. 304–312, Morgan Kaufmann, Stanford, Calif, USA, July 1997.
- [9] W. A. Tackett, "Genetic programming for feature discovery and image discrimination," in *Proc. 5th International Conference on Genetic Algorithms, ICGA-93*, S. Forrest, Ed., pp. 303–309, Morgan Kaufmann, Urbana-Champaign, Ill, USA, July 1993.
- [10] J. F. Winkeler and B. S. Manjunath, "Genetic programming for object detection," in *Proc. 2nd Annual Conference on Genetic Programming (GP-97)*, J. R. Koza, K. Deb, M. Dorigo, et al., Eds., pp. 330–335, Morgan Kaufmann, Stanford, Calif, USA, July 1997.
- [11] A. Teller and M. Veloso, "A controlled experiment: evolution for learning difficult image classification," in *Proc. 7th Portuguese Conference On Artificial Intelligence*, C. Pinto-Ferreira and N. J. Mamede, Eds., vol. 990 of *Lecture Notes in Computer Science*, pp. 165–176, Springer-Verlag, Funchal, Madeira Island, Portugal, October 1995.
- [12] M. Zhang and V. Ciesielski, "Centred weight initialization in neural networks for object detection," in *Computer Science '99: Proc. 22nd Australasian Computer Science Conference*, J. Edwards, Ed., pp. 39–50, Springer-Verlag, Auckland, New Zealand, January 1999.
- [13] T. Caelli and W. F. Bischof, *Machine Learning and Image Interpretation*, Plenum Press, New York, NY, USA, 1997.
- [14] O. Faugeras, *Three-Dimensional Computer Vision—A Geometric Viewpoint*, MIT Press, Cambridge, Mass, USA, 1993.
- [15] E. Gose, R. Johnsonbaugh, and S. Jost, *Pattern Recognition and Image Analysis*, Prentice-Hall, Upper Saddle River, NJ, USA, 1996.
- [16] M. V. Shirvaikar and M. M. Trivedi, "A neural network filter to detect small targets in high clutter backgrounds," *IEEE Transactions on Neural Networks*, vol. 6, no. 1, pp. 252–257, 1995.
- [17] P. Winter, S. Sokhansanj, H. C. Wood, and W. Crerar, "Quality assessment and grading of lentils using machine vision," in *Canadian Society of Agricultural Engineering Annual Meeting at the Agricultural Institute of Canada Annual Conference*, Lethbridge, AB, Canada, July 1996, CSAE paper No. 96-310.
- [18] E. Baum and D. Haussler, "What size net gives valid generalization?," *Neural Computation*, vol. 1, no. 1, pp. 151–160, 1989.
- [19] D. Howard, S. C. Roberts, and R. Brankin, "Target detection in SAR imagery by genetic programming," *Advances in Engineering Software*, vol. 30, no. 5, pp. 303–311, 1999.
- [20] S.-H. Lin, S.-Y. Kung, and L.-J. Lin, "Face recognition/detection by probabilistic decision-based neural network," *IEEE Transactions on Neural Networks*, vol. 8, no. 1, pp. 114–132, 1997.
- [21] Y. LeCun, B. Boser, J. S. Denker, et al., "Backpropagation applied to handwritten zip code recognition," *Neural Computation*, vol. 1, no. 4, pp. 541–551, 1989.
- [22] W. A. Tackett, *Recombination, selection, and the genetic construction of computer programs*, Ph.D. thesis, Faculty of the Graduate School, University of Southern California, Canoga Park, Calif, USA, April 1994.
- [23] D. Andre, "Automatically defined features: the simultaneous evolution of 2-dimensional feature detectors and an algorithm for using them," in *Advances in Genetic Programming*, K. E. Kinneer, Jr., Ed., pp. 477–494, MIT Press, Cambridge, Mass, USA, 1994.
- [24] J. R. Koza, *Genetic Programming II: Automatic Discovery of Reusable Programs*, MIT Press, Cambridge, Mass, USA, 1994.
- [25] A. Teller and M. Veloso, "PADO: learning tree structured algorithms for orchestration into an object recognition system," Tech. Rep. CMU-CS-95-101, Department of Computer Science, Carnegie Mellon University, Pittsburgh, Pa, USA, 1995.
- [26] G. Robinson and P. McIlroy, "Exploring some commercial applications of genetic programming," in *Proc. AISB Workshop on Evolutionary Computing*, T. C. Fogarty, Ed., vol. 993 of *Lecture Notes in Computer Science (LNCS)*, pp. 234–264, Springer-Verlag, Sheffield, UK, April 1995.
- [27] S. Isaka, "An empirical study of facial image feature extraction by genetic programming," in *Late Breaking Papers at the 1997 Genetic Programming Conference*, J. R. Koza, Ed., pp. 93–99, Stanford Bookstore, Stanford, Calif, USA, July 1997.
- [28] S. A. Stanhope and J. M. Daida, "Genetic programming for automatic target classification and recognition in synthetic aperture radar imagery," in *Evolutionary Programming VII: Proc. 7th Annual Conference on Evolutionary Programming*, V. W. Porto, N. Saravanan, D. Waagen, and A. E. Eiben, Eds., vol. 1447 of *Lecture Notes in Computer Science (LNCS)*, pp. 735–744, Springer-Verlag, San Diego, Calif, USA, March 1998.
- [29] K. Benson, "Evolving finite state machines with embedded genetic programming for automatic target detection within SAR imagery," in *Proc. 2000 Congress on Evolutionary Computation CEC00*, pp. 1543–1549, IEEE Press, La Jolla, Calif, USA, July 2000.

- [30] D. Howard, S. C. Roberts, and C. Ryan, "The boru data crawler for object detection tasks in machine vision," in *Proc. EvoWorkshops 2002, Applications of Evolutionary Computing*, S. Cagnoni, J. Gottlieb, E. Hart, M. Middendorf, and G. Raidl, Eds., vol. 2279 of *Lecture Notes in Computer Science (LNCS)*, pp. 220–230, Springer-Verlag, Kinsale, Ireland, April 2002.
- [31] B. J. Lucier, S. Mamillapalli, and J. Palsberg, "Program optimization for faster genetic programming," in *Proc. 3rd Annual Conference on Genetic Programming (GP-98)*, J. R. Koza, W. Banzhaf, K. Chellapilla, et al., Eds., pp. 202–207, Morgan Kaufmann, Madison, Wis, USA, July 1998.
- [32] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press, Cambridge, Mass, USA, 1992.
- [33] J. R. Koza, "Simultaneous discovery of reusable detectors and subroutines using genetic programming," in *Proc. 5th International Conference on Genetic Algorithms, (ICGA '93)*, S. Forrest, Ed., pp. 295–302, Morgan Kaufmann, Urbana-Champaign, Ill, USA, 1993.
- [34] D. Howard, S. C. Roberts, and C. Ryan, "Evolution of an object detection ant for image analysis," in *Genetic and Evolutionary Computation Conference Late Breaking Papers*, E. D. Goodman, Ed., pp. 168–175, San Francisco, Calif, USA, July 2001.
- [35] R. Poli, "Genetic programming for image analysis," in *Proc. 1st Annual Conference on Genetic Programming (GP-96)*, J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, Eds., pp. 363–368, MIT Press, Stanford, Calif, USA, July 1996.
- [36] F. Lindblad, P. Nordin, and K. Wolff, "Evolving 3d model interpretation of images using graphics hardware," in *Proc. 2002 Congress on Evolutionary Computation CEC2002*, pp. 225–230, Honolulu, Hawaii, USA, May 2002.
- [37] C. T. M. Graae, P. Nordin, and M. Nordahl, "Stereoscopic vision for a humanoid robot using genetic programming," in *Proc. EvoWorkshops 2000, Real-World Applications of Evolutionary Computing*, S. Cagnoni, R. Poli, G. D. Smith, et al., Eds., vol. 1803 of *Lecture Notes in Computer Science (LNCS)*, pp. 12–21, Springer-Verlag, Edinburgh, Scotland, UK, April 2000.
- [38] P. Nordin and W. Banzhaf, "Programmatic compression of images and sound," in *Proc. 1st Annual Conference on Genetic Programming (GP-96)*, J. R. Koza, D. E. Goldberg, D. B. Fogel, and R. L. Riolo, Eds., pp. 345–350, MIT Press, Stanford, Calif, USA, July 1996.
- [39] N. Rai, "Pixel statistics in neural networks for domain independent object detection," Minor thesis, Department of Computer Science, Faculty of Applied Science, RMIT University, 2001.
- [40] M. Fuchs, "Crossover versus mutation: an empirical and theoretical case study," in *Proc. 3rd Annual Conference on Genetic Programming (GP-98)*, J. R. Koza, W. Banzhaf, K. Chellapilla, et al., Eds., pp. 78–85, Morgan Kaufmann, Madison, Wis, USA, July 1998.
- [41] K. Harries and P. Smith, "Exploring alternative operators and search strategies in genetic programming," in *Proc. 2nd Annual Conference on Genetic Programming (GP-97)*, J. R. Koza, K. Deb, M. Dorigo, et al., Eds., pp. 147–155, Morgan Kaufmann, Stanford, Calif, USA, July 1997.
- [42] P. Wilson, "Development of genetic programming strategies for use in the robocup domain," Tech. Rep., Department of Computer Science, RMIT, 1998, Honours thesis.

Mengjie Zhang received a B.E. (mechanical engineering) and an M.E. (computer applications) in 1989 and 1992 from the Department of Mechanical and Electrical Engineering, Agricultural University of Hebei, China, and a Ph.D. in computer science from RMIT University, Melbourne, Australia, in 2000. During 1992–1995, he worked at the Artificial Intelligence Research Centre, Agricultural University of Hebei, China. In 2000, he moved to Victoria University of Wellington, New Zealand. His research is focused on data mining, machine learning, and computer vision, particularly genetic programming, neural networks, and object detection. He is also interested in web information extraction, and knowledge-based systems.



Victor B. Ciesielski received his B.S. and M.S. degrees in 1972 and 1975, respectively, from the University of Melbourne, Australia and his Ph.D. degree in 1980 from Rutgers University, USA. He is currently Associate Professor at the School of Computer Science and Information Technology, RMIT University, where he heads the Evolutionary Computation and Machine Learning Group. Dr. Ciesielski's research interests include evolutionary computation, computer vision, data mining, machine learning for robot soccer, and, in particular, genetic programming approaches to object detection and classification.



Peter Andreae received a B.E. (honours) in electrical engineering from the University of Canterbury, New Zealand, in 1977 and a Ph.D. in artificial intelligence from MIT in 1985. Since 1985, he has been teaching computer science at Victoria University of Wellington, New Zealand. His research interests are centered in the area of making agents that can learn behaviour from experience, but he has also worked on a wide range of topics, ranging from reconstructing vasculature from x-rays, clustering algorithms, analysis of micro-array data, programming by demonstration, and software reuse.

