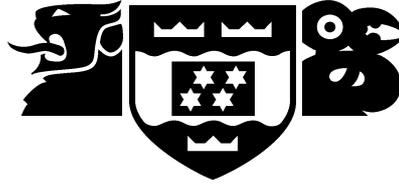


VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui



School of Mathematical and Computing Sciences
Computer Science

Genetic Programming with Gradient
Descent Search for Multiclass Object
Classification

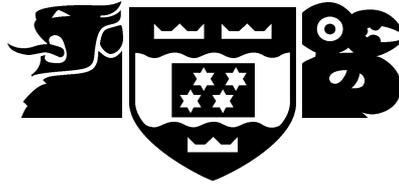
Mengjie Zhang, Will Smart

Technical Report CS-TR-04/4
Feb 2004

School of Mathematical and Computing Sciences
Victoria University
PO Box 600, Wellington
New Zealand

Tel: +64 4 463 5341
Fax: +64 4 463 5045
Email: Tech.Reports@mcs.vuw.ac.nz
<http://www.mcs.vuw.ac.nz/research>

VICTORIA UNIVERSITY OF WELLINGTON
Te Whare Wananga o te Upoko o te Ika a Maui



School of Mathematical and Computing Sciences
Computer Science

PO Box 600
Wellington
New Zealand

Tel: +64 4 463 5341, Fax: +64 4 463 5045
Email: Tech.Reports@mcs.vuw.ac.nz
<http://www.mcs.vuw.ac.nz/research>

Genetic Programming with Gradient
Descent Search for Multiclass Object
Classification

Mengjie Zhang, Will Smart

Technical Report CS-TR-04/4
Feb 2004

Abstract

This paper describes an approach to the use of gradient descent search in genetic programming (GP) for object classification problems. In this approach, pixel statistics are used to form the feature terminals and a random generator produces numeric parameter terminals. The four arithmetic operators and a conditional operator form the function set and the classification accuracy is used as the fitness function. Gradient descent search is introduced to the GP mechanism and is embedded into the genetic beam search, which allows the evolutionary learning process to globally follow the beam search and locally follow the gradient descent search. Two different methods, an online gradient descent scheme and an offline gradient descent scheme, are developed and compared with the basic GP method on three image data sets with object classification problems of increasing difficulty. The results show that both the online and the offline gradient descent GP methods outperform the basic GP method in both classification accuracy and training time and that the online scheme achieved better performance than the offline scheme. This suggests that the GP method with gradient descent search is more effective and more efficient than without and that the online gradient descent algorithm is best suited to object classification problems.

Keywords Genetic programming, genetic algorithms, online Learning, offline learning, target recognition, object recognition.

Author Information

Mengjie Zhang is an academic staff member in computer science and Will Smart is a MSc (part 2) student in computer science. Both authors are in the School of Mathematical and Computing Sciences, Victoria University of Wellington, New Zealand.

Genetic Programming with Gradient Descent Search for Multiclass Object Classification

Mengjie Zhang and Will Smart

School of Mathematical and Computer Science
Victoria University of Wellington,
P. O. Box 600, Wellington, New Zealand
Email: {mengjie, smartwill}@mcs.vuw.ac.nz

Abstract. This paper describes an approach to the use of gradient descent search in genetic programming (GP) for object classification problems. In this approach, pixel statistics are used to form the feature terminals and a random generator produces numeric parameter terminals. The four arithmetic operators and a conditional operator form the function set and the classification accuracy is used as the fitness function. Gradient descent search is introduced to the GP mechanism and is embedded into the genetic beam search, which allows the evolutionary learning process to globally follow the beam search and locally follow the gradient descent search. Two different methods, an online gradient descent scheme and an offline gradient descent scheme, are developed and compared with the basic GP method on three image data sets with object classification problems of increasing difficulty. The results show that both the online and the offline gradient descent GP methods outperform the basic GP method in both classification accuracy and training time and that the online scheme achieved better performance than the offline scheme. This suggests that the GP method with gradient descent search is more effective and more efficient than without and that the online gradient descent algorithm is best suited to object classification problems. Although developed for object classification problems, this approach is expected to be able to be applied to general classification and prediction tasks.

1 Introduction

Since the early 1990s, there has been a number of reports on applying genetic programming techniques to object recognition problems [1–7]. Typically, these GP systems used either high level or low level image features as the terminal set, arithmetic and conditional operators as the function set, and classification accuracy, error rate or similar measures as the fitness function. During the evolutionary process, selection, crossover and mutation operators were applied to the genetic beam search to find good solutions. While most of these GP systems achieved reasonable even good results, they usually spent a long time for training/learning good programs for a particular task. In addition, because of the long training times, the evolutionary process was often stopped when a maximum number of generations was reached, rather than an ideal solution was found.

Gradient descent is a long term established search/learning technique and commonly used to train multilayer feed forward neural networks [8]. This algorithm can guarantee to find a local minima for a particular task. While the local minima is not the best solution, it often meets the request of that task.

In this paper, we apply the gradient descent search to genetic programming, so that a hybrid beam-gradient descent search scheme can be formed. During the evolutionary process, the beam search is still the basic global search mechanism, but the gradient descent search is locally applied to individual programs in the population inside a particular generation.

The goal of this research is to develop such a search scheme in genetic programming for multiclass object classification problems, and to investigate whether this approach can perform better than the basic GP approach in terms of training efficiency and classification performance.

The rest of the paper is organised as follows. Section 2 describes our basic GP approach to object classification, including terminals, functions, fitness function and genetic operations. Section 3 specifies the gradient descent search algorithm in our approach. Section 4 gives the image data sets used in the experiments. Section 5 presents experimental results and section 6 draws the conclusions and gives future work.

2 GP Applied to Object Classification

In the basic GP approach, we used the tree-structure to represent genetic programs [9]. The ramped half-and-half method was used for generating programs in the initial population and for the mutation operator [10]. The proportional selection mechanism and the reproduction [11], crossover and mutation operators [12] were used in the learning and evolutionary process.

In the remainder of this section, we address the other aspects of our basic GP system: (1) Determination of the terminal set; (2) Determination of the function set; (3) Construction of the fitness measure; and (4) Selection of the input parameters and determination of the termination strategy.

2.1 Terminals

In this approach, we used two kinds of terminals: *feature terminals* and *numeric parameter terminals*.

Feature Terminals. Feature terminals form the inputs from the environment and usually correspond to image features in object recognition and image analysis. To achieve domain independent object classification, our system used *pixel statistics*, domain independent low level image features, as the feature terminals.

The pixel statistics considered in this approach are the means and variances of certain regions in object cutout images. Two such regions were used, the entire object cutout image and the central square region. This makes four feature terminals. Since the ranges of these four feature terminal values are quite different,

we linearly scaled them into the range $[-1, 1]$ based on all object image examples to be classified.

The values of the feature terminals are the values of the pixel statistics for certain object examples. Similarly to neural network inputs which are not changed during network training, these values would remain unchanged in the evolutionary process, although different objects usually have different feature values.

Numeric Parameter Terminals. Numeric parameter terminals are floating point numbers randomly generated using a uniform distribution at the beginning of evolution. To be consistent with the feature terminals, we also set the range of these parameter terminals to $[-1, 1]$.

Unlike the feature terminals, the values of this kind of terminals are the same for all object images. Similarly to feature terminals, they would remain unchanged during the evolutionary recombination (crossover and reproduction) in the basic GP approach. However, they will be changed and updated with a continuous parameter entity when the gradient descent algorithm (see section 3) is applied, which are similar to the weights and biases in neural networks.

2.2 Functions

In the function set, the four standard arithmetic and a conditional operation were used to form the function set:

$$FuncSet = \{+, -, *, /, if\} \quad (1)$$

The $+$, $-$, and $*$ operators have their usual meanings — addition, subtraction and multiplication, while $/$ represents “protected” division which is the usual division operator except that a divide by zero gives a result of zero. Each of these functions takes two arguments. The *if* function takes three arguments. The first argument, which can be any expression, constitutes the condition. If the first argument is negative, the *if* function returns its second argument; otherwise, it returns its third argument.

2.3 Classification Strategy

The output of a genetic program in the standard GP system is a floating point number. Generally genetic programs can perform two class object classification tasks quite well where the division between positive and negative numbers of a genetic program output corresponds to the separation of the two object classes. However, for multiple class object classification problems described here, where more than two classes of objects are involved, the standard genetic programming classification strategy mentioned above cannot be applied.

In this approach, we used a different strategy — a variant version of the *program classification map* [11]. This variation situates class regions sequentially on the floating point number line. The object image will be classified to the

class of the region that the program output with the object image input falls into. Class region boundaries start at some negative number, and end at the same positive number. Boundaries between the starting point and the end point are allocated with an identical interval of 1.0. For example, a five class problem would have the following classification map:

$$\mathbf{class} = \begin{cases} \text{Class 1, } r < -1.5 \\ \text{Class 2, } -1.5 \leq r < -0.5 \\ \text{Class 3, } -0.5 \leq r < 0.5 \\ \text{Class 4, } 0.5 \leq r < 1.5 \\ \text{Class 5, } 1.5 \leq r \end{cases} \quad (2)$$

where r is the program output.

2.4 Fitness Function

We used classification accuracy on the training set of object images as the fitness function. The classification accuracy of a genetic program classifier refers to the number of object images that are correctly classified by the genetic program classifier as a proportion of the total number of object images in the training set. According to this design, the best fitness is 100%, meaning that all object images have been correctly recognised.

2.5 Parameters and Termination Criteria

The parameter values used in this approach are shown in table 1.

Table 1. Parameters used for GP training for the three datasets.

Parameter Kinds	Parameter Names	Shape	coin1	coin2
Search Parameters	population-size	300	300	500
	initial-max-depth	3	3	3
	max-depth	5	6	8
	max-generations	51	51	51
	object-size	16×16	70×70	70×70
Genetic Parameters	reproduction-rate	20%	20%	20%
	cross-rate	50%	50%	50%
	mutation-rate	30%	30%	30%
	cross-term	15%	15%	15%
	cross-func	85%	85%	85%

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

- The classification problem has been solved on the training set, that is, all objects of interest in the training set have been correctly classified without any missing objects or false alarms for any class.

- The accuracy on the validation set starts falling down.
- The number of generations reaches the pre-defined number, *max-generations*.

3 Gradient Descent Applied to Genetic Programming

In this section, we describe how to apply the gradient descent search to genetic programming. In this approach, both the online scheme and the offline scheme were introduced.

3.1 Overview of the Algorithm

We assume the successful degree to which the task has been performed can be measured on some real scale and that the performance is scalar and continuous. We also assume that the *cost function* C is used as the measure.

A continuous cost surface can be formed for a given task based on a set of parameters. The lower the point in the cost surface, the better the performance of the system. To improve the system performance, the gradient descent search is applied to take steps “downhill” on C from the current parameter θ .

The gradient of C is found as the vector of partial derivatives with respect to the parameter values. This gradient vector points along the surface, in the direction of maximum-slope at the point used in the derivation. Changing the parameters proportionally to this vector (negatively, as it points to “uphill”) will move the system down the surface C . The distance moved should therefore be the length of the vector times a factor α .

$$\Delta\theta_i = -\alpha \cdot \frac{\partial C}{\partial \theta_i} \quad (3)$$

where θ_i is the i 'th parameter, α is a factor.

It is important to note that the gradient descent algorithm does not replace *any* of the normal genetic operators that produce populations from generation to generation. Instead, the gradient descent algorithm augments the existing GP system, by locally applying gradient descent search to each program in the current population in a particular generation.

In our GP system, the parameters to be applied to the gradient descent algorithm are the numeric parameter terminals in each individual program. The feature terminals cannot be changed, as they are computed based on actual object examples. The gradient of the cost surface with respect to the values of the numeric parameter terminals can be found through the gradient descent algorithm, which is similar to the back propagation algorithm used in training neural networks.

3.2 The Cost Function

In this approach, we used half of the squared difference between the program actual output and the desired output for a particular object input as the cost function, as shown in equation 4.

$$C_\theta = \frac{(y_\theta - Y)^2}{2} \quad (4)$$

where C_θ is the value of the cost surface at the current parameters θ , y_θ is the actual output of program with the current object as input. Y is the corresponding desired output and is calculated by equation 5.

$$Y = \mathbf{class} - \frac{\mathbf{numclass} + 1}{2} \quad (5)$$

where **class** is the class label of the object and **numclass** is the total number of classes. For example, for a five class problem as described in equation 2, the desired outputs are $-2, -1, 0, 1,$ and 2 for object classes 1, 2, 3, 4, and 5, respectively. In other words, we take the centres of the class regions as desired outputs except the beginning and the end classes.

Based on equations 4 and 5, we will obtain the derivative of the whole genetic program which contains one or more numeric parameter terminals, as shown in equation 6.

$$\frac{\partial C}{\partial y} = \frac{\partial \left(\frac{(y-Y)^2}{2} \right)}{\partial y} = y - Y \quad (6)$$

The partial derivatives for the numeric parameter terminals can be obtained based on the *chain rule*, which will be described in the next sub section.

3.3 Chain Rules in Genetic Programs

In general, if f , g and h are functions where f depends on g and g depends on h , then the chain rule can be represented as several parts based on some partial derivatives:

$$\frac{\partial f}{\partial h} = \frac{\partial f}{\partial g} \times \frac{\partial g}{\partial h} \quad (7)$$

Now we use the program tree shown in figure 1 as an example to describe the chain rule in our algorithm. Assume O_j is the evaluated result of node j , then $y = O_1$ is the final output of the program.

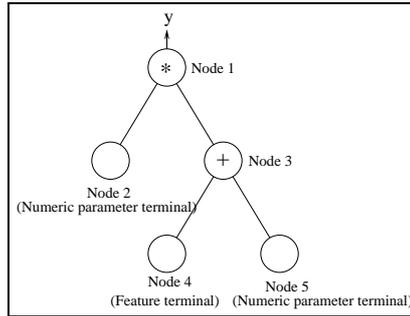


Fig. 1. An example program.

Since nodes 2 and 5 are numeric parameter terminals, the gradient vector should contain values from the partial derivatives from nodes 2 and 5. The partial derivatives of the cost function on node 2 and node 5 are:

$$\begin{aligned}\frac{\partial C}{\partial O_2} &= \frac{\partial C}{\partial y} \cdot \frac{\partial y}{\partial O_2} = \frac{\partial C}{\partial y} \cdot \frac{\partial(O_2 O_3)}{\partial O_2} = (y - Y) \cdot O_3 \\ \frac{\partial C}{\partial O_5} &= \frac{\partial C}{\partial y} \cdot \frac{\partial y}{\partial O_5} = \frac{\partial C}{\partial y} \cdot \frac{\partial(O_2 O_3)}{\partial O_5} = \frac{\partial C}{\partial y} \cdot \frac{\partial(O_2 O_3)}{\partial O_3} \cdot \frac{\partial O_3}{\partial O_5} \\ &= \frac{\partial C}{\partial y} \cdot \frac{\partial(O_2 O_3)}{\partial O_3} \cdot \frac{\partial(O_4 + O_5)}{\partial O_5} = (y - Y) \cdot O_2 \cdot 1 = (y - Y) \cdot O_2\end{aligned}$$

As y, Y, O_2 , and O_3 can be obtained from evaluation of the whole program or a part of the program or calculated by equation 5, these gradients can be calculated accordingly. In other words, using the chain rule the gradients can be broken down to evaluated values and derived mathematical operators. The chain rule can be applied many times to accommodate programs of any depth.

The derivatives of the various functions used in this approach are listed in table 2.

Table 2. Function derivatives.

Function f	meanings	$\frac{\partial f}{\partial a_1}$	$\frac{\partial f}{\partial a_2}$	$\frac{\partial f}{\partial a_3}$
$(+ a_1 a_2)$	$a_1 + a_2$	1	1	n/a
$(- a_1 a_2)$	$a_1 - a_2$	1	-1	n/a
$(* a_1 a_2)$	$a_1 \times a_2$	a_2	a_1	n/a
$(/ a_1 a_2)$	$a_1 \div a_2$	a_2^{-1}	$-a_1 \times a_2^{-2}$	n/a
(if $a_1 a_2 a_3$)	if $a_1 < 0$ then a_2 else a_3	0 0	1 if $a_1 < 0$ 0 if $a_1 \geq 0$	0 if $a_1 < 0$ 1 if $a_1 \geq 0$

3.4 Calculation of the Factor α

In this approach, the factor α in equation 3 is proportional to the inversed sum of the square gradients on all numeric parameter terminals along the cost surface, as shown in equation 8.

$$\alpha = \eta \cdot \frac{1}{\sum_i^N \left(\frac{\partial y}{\partial O_i}\right)^2} \quad (8)$$

where N is the number of numeric parameter terminals in the program, i indexes the numeric parameter terminals, y is the output of the program, O_i is the output of the i th numeric parameter terminal, η is a learning rate defined by the user.

3.5 Online Gradient Descent Algorithm

Based on equation 3, the change of the value of the i th numeric parameter terminal would be:

$$\Delta O_i = -\eta \cdot \frac{1}{\sum_i^N \left(\frac{\partial y}{\partial O_i}\right)^2} \cdot \frac{\partial C}{\partial O_i} \quad (9)$$

Accordingly, the new value of the numeric parameter terminal would be:

$$(O_i)_{new} = O_i + \Delta O_i = O_i - \eta \cdot \frac{1}{\sum_i^N (\frac{\partial y}{\partial O_i})^2} \cdot \frac{\partial C}{\partial O_i} \quad (10)$$

The online Gradient descent algorithm in GP is summarised as follows. For each program on each object, do the following:

- Evaluate program, save the outputs of all nodes in the program.
- Calculate the (partial) derivatives of the cost function at numeric parameter terminals using the chain rule and table 2.
- Calculate the change of each numeric parameter terminal using equation 9.
- Update the numeric parameter terminals according to equation 10.

Based on this definition, if the program only has linear operators performed on subtrees with numeric parameter terminals, the changes to the numeric parameter terminals on the program are independent, and the learning rate η is 1.0, then the program with the changes driven by this online gradient descent algorithm would produce ideal result for this object input, that is, the error would be zero (the prove is omitted here due to the page limitation).

3.6 Offline Gradient Descent Algorithm

The offline learning algorithm is similar to the online learning algorithm except that the change of each numeric parameter terminal in a genetic program is based on the sum of all object examples in the training set rather than on each single object. Accordingly, the numeric parameter terminals will be only updated after a whole cycle of all examples in the training set.

4 Image Data Sets

We used three image data sets in the experiments. Example images are shown in figure 2.

The first set of images (figure 2 (a)) was generated to give well defined objects against a noisy background. The pixels of the objects were produced using a Gaussian generator with different means and variances for each class. Four classes of 713 small objects were cut out from these images to form the classification data set *shape*. The four classes are: black circles, light grey squares, white circles, and grey noisy background. This set was considered to include an easy object classification problem.

The second data set (*coin1*, figure 2 (b)) was intended to be harder than data set 1 and consist of scanned 5 cent and 10 cent New Zealand coins. There are five classes of 576 object cutouts: 5 cent heads, 5 cent tails, 10 cent heads and 10 cent tails, and a relatively uniform background. Compared with the *shape* data

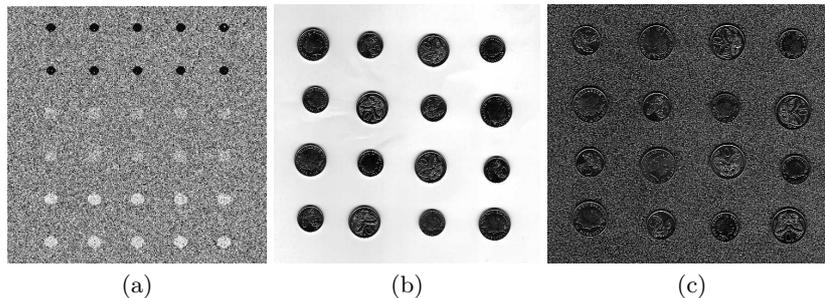


Fig. 2. Example images from Shape (a), Coin1 (b) and Coin2 (c)

set, the objects in this set are more difficult to distinguish and this set has more classes.

The third data set (*coin2*, figure 2 (c)) also consists of five classes of object cutouts, but the background is highly clustered, which makes the classification problems much harder. Human eyes also have difficulty to distinguish heads from tails for either 5 cent or 10 cent coins. Thus, these datasets provide three object classification problems of increasing difficulty.

The object cutout images in each of these datasets were equally split into three separate data sets: one third for the training set used directly for learning the genetic program classifiers, one third for the validation set for controlling over-fitting, and one third for the test set for measuring the performance of the learned program classifiers.

5 Results and Discussion

This section presents the results of our GP approach with both the online and the offline gradient descent algorithms on the four datasets and compare them with the results of the basic GP approach without the gradient descent search. For all cases, 10 runs were carried out and the average results on the test set are presented.

5.1 Shape Dataset

Table 3. Results of the shape Datasets.

Dataset	η	Generations			Accuracy (%)		
		Basic	Online	Offline	Basic	Online	Offline
Shape	0.0	9.56	n/a	n/a	99.48	n/a	n/a
	0.2	n/a	1.00	3.20	n/a	100.00	99.77
	0.4	n/a	1.00	2.20	n/a	100.00	99.86
	0.7	n/a	1.00	2.80	n/a	100.00	99.81
	1.0	n/a	1.00	3.30	n/a	100.00	99.63
	1.4	n/a	1.00	5.50	n/a	99.95	99.77

Table 3 shows the results on the Shape data set using different learning rates. The first line shows that for the Shape data set, the basic GP approach without using the gradient descent algorithm (η is 0.0) achieved an average accuracy of 99.48% over 10 runs on the test set and the average number of generations of the 10 runs spent on the training process was 9.56.

For dataset *shape*, the online gradient descent algorithm achieved the best performance of all the three methods in terms of both classification accuracy and training time. In particular, using the online gradient descent algorithm at learning rates of 0.2, 0.4, 0.7 and 1.0, ideal performances were achieved and only one generation¹ was required for the evolutionary learning process. This was considerably faster than the basic GP approach. While the offline algorithm also achieved better accuracy and a faster training than the basic GP approach for these learning rates, it was not as good as the online algorithm for this data set.

5.2 Coin Datasets

As shown in table 4, the results for the two coin datasets show a similar pattern to the *shape* dataset. In both datasets, for all the learning rates investigated here, the accuracy performances of the GP approach with the two gradient descent algorithms were always superior to those without using the gradient descent search. Again, the online algorithm was better than the offline algorithm on these data sets. These results suggests that a combination of the basic genetic beam search with the gradient descent search could always find better genetic program classifiers than the genetic beam search only for these object classification tasks.

Table 4. Results of the two *coin* Datasets.

Dataset	η	Generations			Accuracy (%)		
		Basic	Online	Offline	Basic	Online	Offline
Coin1	0.0	51.00	n/a	n/a	82.12	n/a	n/a
	0.2	n/a	20.40	35.70	n/a	98.94	94.19
	0.4	n/a	25.40	43.40	n/a	98.94	93.06
	0.7	n/a	23.30	46.60	n/a	98.81	91.81
	1.0	n/a	36.70	47.30	n/a	97.69	92.44
	1.4	n/a	34.00	37.50	n/a	97.94	94.38
Coin2	0.0	51.00	n/a	n/a	73.83	n/a	n/a
	0.2	n/a	51.00	51.00	n/a	83.50	72.17
	0.4	n/a	51.00	51.00	n/a	82.83	77.67
	0.7	n/a	50.20	51.00	n/a	85.17	78.83
	1.0	n/a	51.00	51.00	n/a	86.50	84.17
	1.4	n/a	51.00	51.00	n/a	80.83	79.00

¹ One generation in GP with the two gradient descent algorithms usually requires a slightly longer time than that in the basic GP approach. However, this was considerably improved by only applying the gradient descent algorithm to the top 5% of the programs in the population. Thus, the actual times for a single generation in the three methods are still very similar.

In terms of the number of generations used in the training process, the GP approach with the two gradient descent algorithms required fewer generations to achieve good results for dataset *coin1*. For example, the approach with the online algorithm at a learning rate of 0.2 only used 20.4 generations and achieved almost ideal performance (98.94%), while the basic GP approach used 51 generations but only obtained 82.12% accuracy. For dataset *coin2*, all the GP approaches used almost all the 51 generations (one of the stopping criteria). However, after examining the internal behaviour of the evolutionary process, we found that the method with the online gradient descent algorithm converged at generations 10-20 in all the cases while the basic method without this algorithm got converged at generations 45-51 in almost all the cases. The offline algorithm was a bit slower than the online algorithm, but it still converged much faster than the basic GP method. These results suggest that the GP approach with the gradient descent search converged faster than without.

For dataset *coin2*, the accuracy performance could not improved with even more generations. This is probably because only four pixel statistics were used as features terminals, which were not sufficient for this difficult object classification task.

The results also show that different performances were obtained if different learning rates were used. It did not appear to have a reliable way to determine a good learning rate. As in neural networks, this learning parameter needs to be empirically searched through tuning certain experiments to obtain good results. However, if this can lead to considerably better results, this price is worth to pay. Our experiments suggest that a learning rate between 0.2 to 1.0 is a good starting point for object classification problems.

6 Conclusions

The goal of this paper is to develop an approach to integrating gradient descent search to genetic programming (GP) and to investigate whether the hybrid approach is better than the basic GP approach for object classification problems. This goal was achieved by introducing and embedding the gradient descent search into the genetic beam search to form the online and offline gradient descent algorithms, allowing the evolutionary process to globally follow the beam search and locally follow the gradient descent search to find good solutions.

On all the datasets investigated here, the two GP methods with the gradient descent search always achieved better results than the method without in both classification accuracy and training generations. In particular, the online gradient descent algorithm suited the best to these classification problems. As expected, the performances on the three image datasets deteriorated as the degree of difficulty of the object classification problem was increased.

Although developed for object classification problems, this new method is expected to be able to be applied to general classification and prediction tasks.

For future work, we will investigate whether the performance on the difficult coin data sets can be improved if more features are added to the feature terminal

set. We will also investigate the power and reliability of the gradient descent GP methods on even more difficult image classification problems such as face recognition problems and satellite image detection problems, and compare the performance with other long-term established methods such as decision trees, neural networks, and support vector machines.

References

1. David Andre. Automatically defined features: The simultaneous evolution of 2-dimensional feature detectors and an algorithm for using them. In Kenneth E. Kinneer, editor, *Advances in Genetic Programming*, pages 477–494. MIT Press, 1994.
2. Daniel Howard, Simon C. Roberts, and Richard Brankin. Target detection in SAR imagery by genetic programming. *Advances in Engineering Software*, 30:303–311, 1999.
3. Thomas Loveard and Victor Ciesielski. Representing classification problems in genetic programming. In *Proceedings of the Congress on Evolutionary Computation*, volume 2, pages 1070–1077, COEX, World Trade Center, 159 Samseong-dong, Gangnam-gu, Seoul, Korea, 27-30 May 2001. IEEE Press.
4. Andy Song, Vic Ciesielski, and Hugh Williams. Texture classifiers generated by genetic programming. In *Proceedings of the 2002 Congress on Evolutionary Computation CEC2002*, pages 243–248. IEEE Press, 2002.
5. Walter Alden Tackett. Genetic programming for feature discovery and image discrimination. In *Proceedings of the 5th International Conference on Genetic Algorithms, ICGA-93*, pages 303–309, University of Illinois at Urbana-Champaign, 17-21 July 1993. Morgan Kaufmann.
6. Jay F. Winkeler and B. S. Manjunath. Genetic programming for object detection. In *Genetic Programming 1997: Proceedings of the Second Annual Conference*, pages 330–335, Stanford University, CA, USA, 13-16 July 1997. Morgan Kaufmann.
7. Mengjie Zhang and Victor Ciesielski. Genetic programming for multiple class object detection. In *Proceedings of the 12th Australian Joint Conference on Artificial Intelligence (AI'99)*, pages 180–192, Sydney, Australia, December 1999. Springer-Verlag Berlin Heidelberg. Lecture Notes in Artificial Intelligence (LNAI Volume 1747).
8. D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, and the PDP research group, editors, *Parallel distributed Processing, Explorations in the Microstructure of Cognition, Volume 1: Foundations*, chapter 8. The MIT Press, Cambridge, Massachusetts, London, England, 1986.
9. John R. Koza. *Genetic programming : on the programming of computers by means of natural selection*. Cambridge, Mass. : MIT Press, London, England, 1992.
10. Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, and Frank D. Francone. *Genetic Programming: An Introduction on the Automatic Evolution of computer programs and its Applications*. San Francisco, Calif. : Morgan Kaufmann Publishers; Heidelberg : Dpunkt-verlag, 1998. Subject: Genetic programming (Computer science); ISBN: 1-55860-510-X.
11. Mengjie Zhang, Victor Ciesielski, and Peter Andreae. A domain independent window-approach to multiclass object detection using genetic programming. *EU-RIASP Journal on Signal Processing, Special Issue on Genetic and Evolutionary Computation for Signal Processing and Image Analysis*, 2003(8):841–859, 2003.

12. John R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*. Cambridge, Mass. : MIT Press, London, England, 1994.