Bio-Inspired Computation: Coursework 3
Experimenting with Genetic Programming

Genetic programming (GP) can be applied to a wide variety of problems, and often finds good solutions in comparison with other computational methods. However, in practice good performance requires a certain amount of parameter tuning to be done. A typical GP system has many parameters, and many values that can be chosen for these. These choices affect how quickly and, in the case of multiple runs, how often a good solution will be found. They can also affect other properties of the solutions, such as their size.

The aim of this coursework is for you to gain some understanding of how parameter choices affect the performance of GP in practice. You will be using ECJ to do this. ECJ provides a mature implementation of tree-based GP. Using parameter files, it allows you to experiment with a range of different parameter settings. ECJ also comes with a selection of different benchmarking problems, which you can use to measure the effect of different parameter settings.

What you are asked to do:

1. Familiarise yourself with ECJ and its GP facilities.
2. Investigate how different parameter settings affect the performance of GP.
3. Write a report, of up to 5 pages, describing what you have learnt.

Again, in more detail:

1. ECJ

A copy of ECJ is included in the assessment zip file. Download this archive, decompress it, and place the `ecj` directory somewhere in your filing system.

General usage instructions, including a manual and tutorials, can be found on the ECJ website: http://cs.gmu.edu/~eclab/projects/ecj/docs/

You may find it useful to look through the tutorials, since these contain information about using ECJ and its parameter files. An overview of parameters and parameter files can also be found at: http://cs.gmu.edu/~eclab/projects/ecj/docs/parameters.html

ECJ is a command line programme. Instructions for installing and running it can be found in the `ecj_readme.txt` file in the `ecj/cw3` directory.

2. Experimentation

The aim of this coursework is for you to gain an understanding of how parameter settings affect the performance of GP in practice. This will require you to carry out a series of GP runs using different parameter settings on more than one problem. You will then present the results of these experiments in a suitable manner, and discuss your findings.

ECJ allows you to investigate many parameters and many benchmark problems. To simplify things, the `ecj/cw3` directory contains three ready-made parameter files, each containing a subset of the parameters available in ECJ. Each of these files is for a different problem: two kinds of

symbolic regression, and the Santa Fe ant trail problem (these problems were briefly covered in the lecture slides).

You should experiment with the parameter settings in these files to see how they affect the behaviour of the GP system. You can do this by editing the parameter files, or you can append new settings to the command line. For example, to run ECJ on the quartic regression problem using a population size of 200, run this command from the `ecj` directory:

```
java -cp . ec.Evolve -file cw3/quartic_regression.params
 -p pop.subpop.0.size=200
```

After some initial experimentation, you should choose the **three** parameters that you think have the most notable affect on this system and then carry out a more thorough investigation of their effects across the three problems. It is up to you exactly which parameter values you use for your experiments, but the range and quantity of values should be sensible and justified, for example by looking at which values people use in books and published papers. (Note, if you choose to look at the effect of changes to the function set, the entire function set should be considered to be a single parameter, and different combinations of functions its values).

Bear in mind that GP is a non-deterministic algorithm. This means that it will generate different results in different runs for the same parameter settings. Hence, you should look at performance measures across a number of repeated runs using the same parameter settings. For instance, you should consider measures such as mean fitness, standard deviation, and success rate when determining how a parameter affects performance.

3. Report

Your report (up to 5 pages in length) should summarise the experiments you have carried out. It should:

- explain the purpose of each parameter and briefly discuss the affect that it has on the behaviour of the GP system, based on your initial observations and any information from other sources (such as books and papers)
- explain the motivation behind the three parameters you selected for in-depth study, and the precise values you looked at for each of these parameters
- clearly present the results of the experiments you carried out your in-depth study
- discuss these results, highlighting any insights that you have discovered
- use tables, plots, charts and citations as appropriate.

HAND-IN:

1) hand-in your report at the School Office by 3pm on Thursday 3rd December 2015
2) also email an electronic copy to M.Lones@hw.ac.uk with subject line “BIC CW3”

Marking scheme and assessment criteria

This is worth 20% of the module. Grading will use the assessment criteria given on the following page, and will take into account the quality and quantity of work done. You are encouraged to carry out wider reading, and extra marks will be awarded if you show evidence of this.

Criteria	Weight	A (70-100%)	B (60-69%)	C (50-59%)	D (40-49%)	E/F (<40%)
Discussion of subject matter and results *(introduction, discussion of results, conclusions etc.)*	35%	Thoughtful discussion with insightful conclusions supported by data and evidence of wider learning	Meaningful discussion and conclusions which are supported by data and/or wider learning	Some flaws in the analysis of results and conclusions, limited indication of wider learning	Significant flaws in analysis or conclusions, no useful indication of wider learning	No real indication of understanding, no wider learning
Design of experiments *(choice of parameters and values, multiple repeats etc.)*	25%	Well designed and justified experiments acknowledging wider learning.	Sensible experiments, with reasonable justification	Some weakness in experimental design and/or limited justification	Significant weakness in experimental design, no useful justification	Inappropriate experiments and justification
Presentation of results *(tables, plots, choice of metrics etc.)*	25%	Clear, insightful presentation of experimental results.	Use of appropriate presentation that contributes to understanding of experimental results	Some weakness in how results are presented, limiting understanding	Significant weakness in how results are presented, impairing understanding	Inappropriate form of presenting results, leading to no useful understanding
Report style *(structure, language, formatting, citation style etc.)*	15%	Well structured report, consistent and readable format, citations according to standards	Readable and generally well structured report with usable citations	Some issues with writing and/or structure, citations not standardised	Difficult to read and/or follow, incomplete and/or messy citations	Uninteligible languable, structure and/or citations