Developing and Evaluating Incremental Evolution using High Quality Performance Measures for Genetic Programming

A thesis presented in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Computer Science
at Massey University, Albany, Auckland,
New Zealand

Matthew Garry William Walker

2007

To Mum and Dad

Abstract

This thesis is divided into two parts. The first part considers and develops some of the statistics used in genetic programming (GP) while the second uses those statistics to study and develop a form of incremental evolution and an early termination heuristic for GP.

The first part looks in detail at success proportion, Koza's minimum computational effort, and a measure we rename "success effort". We describe and develop methods to produce confidence intervals for these measures as well as confidence intervals for the difference and ratio of these measures.

The second part studies Jackson's fitness-based incremental evolution. If the number of fitness evaluations are considered (rather than the number of generations) then we find some potential benefit through reduction in the effort required to find a solution. We then automate the incremental evolution method and show a statistically significant improvement compared to GP with automatically defined functions (ADFs).

The success effort measure is shown to have the critical advantage over Koza's measure as it has the ability to include a decreasing cost of failure. We capitalise on this advantage by demonstrating an early termination heuristic that again offers a statistically significant advantage.

Preface

Acknowledgements

First, I'd like to thank Mum and Dad. Without their support and—equally importantly—encouragement, I wouldn't be here at the end of this experience. Thank you for your enthusiasm, for putting up with draft chapters being read to you, and for listening to ideas when mere mortals would have fallen asleep.

Chris Messom, my supervisor, has been great. My PhD-studying friends have had some horror stories to tell regarding their supervisors, but Chris was fantastic. He organised funding (thank you to the Government for their "Top Doctoral" programme), was ridiculously patient, always supportive and thorough and motivational when things weren't working. Another of Chris's wonders is his ability to respond. The university recommends forming an agreement with supervisors to return work within a specific time. They recommended a turn-around time of two weeks. I think Chris might have averaged two hours. Thanks Chris.

Howard Edwards has been a huge resource for the statistics in this thesis. He provided excellent assistance and discussion. He also provided great motivation as it was fantastic to work with someone who was excited by the ideas. Although all the work presented in this thesis is my own, it is Howard and Chris and I who are the "we" in this thesis.

Hopefully we can convince Bruce Mills to return from the United Arab Emirates. Bruce provided the most entertaining conversations (that could easily eliminate an afternoon) introducing me to evolutionary ideas that I still think about. But possibly most importantly to me was that it was his suggestion that gave me a concrete starting point. The ideas transitioned unrecognisably but finally they formed this thesis. His influence was an excellent catalyst just when I needed it.

Christian Gagné produced Open BEAGLE, GP software that I have contributed to and have used almost exclusively throughout the thesis. He provided excellent practical discussions surrounding GP and was very helpful when I requested large datasets of GP experiments.

Steffen Christensen was also of considerable assistance when he provided his absolutely enormous artificial ant datasets.

I'd also like to say a very big thank you to my little sister and my friends. You provided excellent conversations and emails, or lovely distracting phone calls and good vibes, or excellent dinners, or just a break away from this work. You also applied what was highly necessary pressure to get this work done. Thanks.

Finally, I'd like to thank the anonymous reviewers of papers I submitted, especially the reviewer who voted as best paper the work that turned into chapter 3.

Publications

Parts of this thesis have already been published. The initial work surrounding the Wilson-Dependent method in chapter 3 was published at *EuroGP 2007* [114]. The discussion of its reliability (section 3.2) was published at *GECCO* a few months later [115], as was a summary of the confidence interval method for the success effort statistic (chapter 4) [117]. The more detailed analysis of success effort that appears in chapter 4 along with parts of the literature review in chapter 2 and the comparisons that appear in chapter 5 were later published at *CEC* [116].

Table of Contents

List of Tables

List of Figures