

Preface

This book is about the Thirteenth Workshop on Genetic Programming Theory and Practice, a workshop held this year from May 14 to 16, 2015, at the University of Michigan under the auspices of the Center for the Study of Complex Systems. The workshop is a forum for theorists and users of genetic programming to come together and share ideas, insights, and observations. It is designed to be speculative in nature by encouraging participants to discuss ideas or results that are not necessarily ready for peer-reviewed publication.

To facilitate these goals, the time allotted for presentations is longer than is typical at most conferences, and there is also more time devoted for discussion. For example, presenters usually have 40 min to present their ideas and take questions, and then, before each break, there is open discussion on the ideas presented in a session. Additionally, at the end of each day, there is a review of the entire day and the ideas and themes that have emerged during the sessions. Looking back at the schedule, in a typical day, there was 240 min of presentation and 55 min of discussion or fully 19 % of the time spent in open discussion.

In addition to the regular sessions, each day starts with a keynote speaker who gets a full hour of presentation and 10 min of Q&A. By design, the keynotes are generally not about genetic programming but come from a related field or an application area that may be fertile ground for GP. This year, the first keynote speaker was Dave Ackley from the University of New Mexico who delivered and addressed the topic titled “A Requiem for Determinism.” This provocative presentation argued that from the beginning of modern computing, people such as John von Neumann argued that hardware could not be relied on to work perfectly in all cases—just because of the nature of electronics in that they will fail some number of times. These days, the growth of complexity of software has added to this problem. Modern software depends on the user’s ability to reboot the system when things get out of sync or when hardware fail. Ackley argues that the correct response (as foreseen by von Neumann) is to make systems that continue to function even when the system nominally fails. Dave went on to suggest that given that GP takes its cues from nature, we should consider incorporating methods that survive “mistakes” in execution.

The second keynote speaker was Larry Burns, who had been an executive at General Motors and is now a consultant with Google on their autonomous vehicle project. Larry's talk was about the development of autonomous vehicles and the likely arc of adoption of autonomous vehicles, but he went on to discuss the fact that technology cannot be thought of in isolation and in particular that it exists in a cultural context and is co-dependent on the infrastructure. As engineers, we tend to think only of the technology we are developing, but Larry made a strong case for thinking about work in a larger context.

The third keynote was Julian Togelius on "Games Playing Themselves: Challenges and Opportunities for AI Research in Digital Games." Games have been at the center of AI development since the beginning of modern computers. Turing mused on chess-playing computers. Samuel's checker playing system could be argued to be the beginning of neural nets, at least on an engineering level. Deep Thought attracted worldwide attention when it beat Garry Kasparov, the then-reigning world chess champion. Julian posed a number of interesting questions relating to AI, particularly about the human traits of curiosity and what it means to "like" something. He turned the usual dynamic of interaction around by asking the questions whether games could be "curious" about people and later asked whether computers could "like" games or even "like" making good games. It was an interesting reversal on the usual questions about AI work and was an interesting discussion in the context of GP.

While the keynotes at the workshop were provocative and interesting, the chapters in this book are the core of GPTP. The first chapter by Kommenda et al. is titled "Evolving Simple Symbolic Regression Models by Multi-objective Genetic Programming." This interesting chapter revisits the question of evaluating the complexity of GP expressions as part of the fitness measure for evolution. Most previous efforts focused either on the structural complexity of the expression or an expensive calculation of subtrees and their components. This chapter proposes a lightweight semantic metric which lends itself to efficient multi-modal fitness calculations without using input data.

The second chapter, by Elyasaf et al., titled "Learning Heuristics for Mining RNA Sequence-Structure Motifs" explores the difficult problem correlating RNA sequences to biological functionality. This is a critical problem to finding and understanding biological mechanisms derived from specific RNA sequences. The authors use GP to create hyper-heuristics that find cliques within the graphs of RNA. Though the chapter only describes the approach and does not show concrete results, it is a clever approach to a complex problem, and we look forward to seeing results in a future GPTP.

The next chapter, by de Melo and Banzhaf, "Kaizen Programming for Feature Construction for Classification" adopts the Japanese practice of Kaizen (roughly, continuous improvement) to GP in the domain of classification problems. In this case, they use GP to generate new ideas in the Kaizen algorithm where in this case "ideas" mean classifier rules that are recursively improved, removed, or refined. It is an interesting idea that takes advantage of GP's ability to generate novel partial solutions and then refine them using the Kaizen approach.

In chapter “GP As If You Mean It: An Exercise for Mindful Practice” by William Tozier, Bill argues that pathologies of result in GP sometimes inform us as to the nature of the problem we are trying to solve and that our (learned) instinct of changing GP parameters or even mechanisms to produce a “better” result may be misguided. He goes from there to a practice of learning adapted for GP that can improve how we use GP by being mindful of how it behaves as we change single features in the problem. He borrows from Pickering’s *Mangle* to create consistent ways to use GP to learn from the problem rather than to adjust the GP until you get a result you expected.

In chapter “nPool: Massively Distributed Simultaneous Evolution and Cross Validation in EC-Star,” Hodjat and Shahrzad continue work on EC-Star, a GP system designed to be massively parallel using the Cloud. This chapter focuses on evolving classifiers by using local populations with k-fold cross-validation that is later tested across different segments of the samples. Additionally, they are developing these classifiers using time series data, which adds an additional challenge to the problem by requiring a lag as part of the operator set. It is a challenging project that has elements of standard cross-validation with island populations but where learning is not permitted between islands and testing is done entirely on different islands with different samples. This creates a danger of premature convergence/overfitting since populations only have one set of samples to learn on, but they control this as compensated for by extensive validation using the other islands. While this is clearly an interesting approach with some good results, the authors suggest that more work needs to be done before it’s ready for commercial use.

In chapter “Highly Accurate Symbolic Regression with Noisy Training Data”, Michael Korn continues his pursuit of improving an almost plug-and-play approach to solving symbolic regression problems that verge on the pathologic from a GP perspective. Here he introduces an improved algorithm and adds noise to the input data and is able to show that he can still produce excellent results for out-of-sample data. He also makes this system available for further testing by other researchers, inviting them to test it on different symbolic regression problems.

The seventh chapter, by Gustafson et al., is titled “Using Genetic Programming for Data Science: Lessons Learned.” The authors are well versed in industrial applications of computational systems and survey the strengths and weaknesses of GP in such applications. They identify a number of areas where GP offers significant value to Data Scientists but also observe some of the faults of GP in such a context. For those seeking to make GP a more accessible technology in the “real world,” this chapter should be carefully considered.

The eighth chapter is a highly speculative effort by Bill Worzel titled “The Evolution of Everything (EvE) and Genetic Programming.” This chapter sets out to explore more open-ended uses of GP. In particular, he focuses on the coming impact of the Internet of Things (sometimes called the Internet of Everything) on the computing world and speculates that with a constant stream of real-world data, GP could break the mold of generational limits and could constantly evolve solutions that change as the world changes. The effort proposes combining GP,

functional programming, particulate genes, and neural nets and (most speculatively) suggests that if the singularity is reachable, it probably will be evolved rather than autonomously springing into being.

The ninth chapter, titled “Lexicase Selection for Program Synthesis: A Diversity Analysis,” by Spector and Helmuth, is an exploration of the hypothesis that lexicase selection improves diversity in a population. Lexicase selection is compared with tournament selection and implicit fitness sharing. Lexicase showed improved error diversity, which suggests improved population diversity, thus supporting the hypothesis and the expected mechanism for lexicase selection.

In the next chapter, “Behavioral Program Synthesis: Insights and Prospects,” by Krawiec et al., the authors argued at the workshop that a single-valued fitness function “abuses” program evolution by forcing it to evolve a lump sum of what is often a complex set of samples. Instead, they propose using an interaction matrix as a more useful metric as it gives information on specific tests. They argue that not only is information being “left on the table” with single-valued metrics but that the overall behavioral characteristic of an evolved solution is lost and a great deal of nuance and understanding goes missing. They then go on to propose what they call behavioral synthesis which focuses on the behavior of evolved solutions as the dominant factor in evolution. This paper suggests that we need a more nuanced notion of fitness.

The eleventh chapter, “Using Graph Databases to Explore the Dynamics of Genetic Programming Runs,” McPhee et al. continues the search for understanding diversity in GP populations, a long-standing focus for research in the GP community. However, in this case, the authors are more interested in looking for “critical moments in the dynamics of a run.” To do this, they use a graph database to manage the data and then query the database to search for these crucial inflection points. They focus on the question of whether lexicase selection is truly better than tournament selection and why this might be. Though a work still in progress, this chapter suggests that this method of analyzing GP populations is a valuable addition to the GP toolset and re-raises some of the issues explored in chapter “GP As If You Meant It: An Exercise for Mindful Practice” by Tozier about looking at the process and not just the outcome and chapter “Behavioral Program Synthesis: Insights and Prospects” about the study of behavioral synthesis suggesting that this is an area where we will see more study in the near future.

The twelfth chapter is titled “Product Choice with Symbolic Regression and Classification,” by Truscott and Korns. This is one of the first, if not the first use of GP in market research. Huge amounts of money are spent surveying customers, and this data is used to predict brand popularity. The authors describe a survey of cell phones and the analysis produced using the ARC symbolic regression system adapted to classification. The results show well compared to existing methods and suggest that more work in this field may be productive.

The thirteenth chapter by Silva et al., is titled “Multiclass Classification Through Multidimensional Clustering” and revisits the difficult problem of multiclass classifications using GP. This builds from their earlier work which mapped values into higher-dimensional space during the training phase and then collected samples into

the closest cluster in the higher-order space. This chapter extends this idea by adding a pool of groups of possible GP trees and combining them selectively (via evolution) to create an ensemble of high-dimensional mapping functions. In some ways, this suggests a more transparent version of SVM, and the results presented suggest that this extension produces improved results with less overfitting.

The final chapter was written by Stijven et al. and is titled “Prime-Time Symbolic Regression Takes Its Place in the Real World.” With over 25 years of experience in applying symbolic regression to real-world problems, the authors make a strong case for GP to take its place in the frontlines of business. They give examples of how symbolic regression can be applied to business forecasting, commercial process optimization, and policy decision making in addition to their previous demonstration of applications in commercial R&D. Because many business applications are proprietary, they give an example of their methodology, which critically includes careful attention to the design of experiment (DOE) in a model of infectious disease epidemics that can inform policy decisions. All told, it is hard to find a group of people who have done more to advance the acceptance of GP in the real world.

Acknowledgments

We would like to thank all of the participants for again making GP Theory and Practice a successful workshop. As always, it produced a lot of high energy and interesting and topical discussions, debates, and speculations. The keynote speakers added a lot of food for thought and raised some interesting questions about GP’s place in the world. We would also like to thank our financial supporters for making the continued existence of GP Theory and Practice possible. These include:

- The Center for the Study of Complex Systems (CSCS)
- John Koza, Third Millennium Venture Capital Limited
- Michael Korns and Gilda Cabral
- Jason Moore, Computational Genetics Laboratory at Dartmouth College
- Mark Kotanchek and Evolved Analytics
- Babak Hodjat at Sentient
- Steve Everist and Everist Life Sciences
- Heuristic and Evolutionary Algorithms Laboratory, Upper Austria University of Applied Science
- Kordon Consulting

A number of people made key contributions to running the workshop and assisting the attendees while they were in Ann Arbor. Foremost among them was Linda Wood and Susan Carpenter, who made this GPTP workshop run smoothly with their diligent efforts before, during, and after the workshop itself. After the workshop, many people provided invaluable assistance in producing this book. Special thanks go to Kala Groscurth who did a wonderful job working with the authors, editors, and publishers and providing editorial and other assistance to

get the book completed. Jennifer Malat and Melissa Fearon provided invaluable editorial efforts, from the initial plans for the book through its final publication. Thanks also to Springer for helping with various technical publishing issues.

Ann Arbor, MI, USA
Ann Arbor, MI, USA

Midland, MI, USA
November 2015

Rick Riolo
W.P. Worzel
Arthur Kordon
Mark Kotanchek

Genetic Programming Theory and Practice XIII

Riolo, R.; Worzel, B.; Kotanchek, M.; Kordon, A. (Eds.)

2016, XX, 262 p. 69 illus., 31 illus. in color., Hardcover

ISBN: 978-3-319-34221-4