TU Delft

CWI
Centrum Wiskunde & Informatica

Mekelweg 4, 2628 CD Delft
The Netherlands
T +31 (0)15 27 89803
W www.tudelft.nl
E P.A.N.Bosman@tudelft.nl

Science Park 123, P.O. Box 94079, 1090 GB Amsterdam
The Netherlands
T +31 (0)20 592 9333
W www.cwi.nl
E Peter.Bosman@cwi.nl

Date: 04/04/2022

Subject: Evaluation of PhD thesis of Erik Derner

Dear office for doctoral study and research at Czech Technical University in Prague,

I have read the Ph.D. thesis of Mr. Erik Derner with much interest, especially since much of the research revolves around the use of genetic programming (GP), which is a technique studied within some of my research lines that I oversee at CWI and that I teach in my course on evolutionary algorithms at Delft University of Technology.

The thesis contributes to the field of evolutionary computation (EC), a well-established form of nature-inspired heuristic optimization algorithms. The thesis has a particular focus on (the use of) GP, a particular form of EC in which executable programs or expressions are optimized. Specifically, the use of GP for symbolic regression is considered, which is a form of machine learning whereby (typically small) expressions are optimized to fit a given dataset, using a limited set of primitives, such as addition, multiplication, logarithm, and constants.

The contribution to the field of EC from the foundational or fundamental perspective is somewhat limited, as much of the research focuses on novel uses of GP within robotics applications. Nevertheless, such applications and use-cases of the technology are also of prime importance to demonstrate the real-world applicability of EC techniques.

Within the application domain of robotics, this thesis achieves exactly this. Motivated from problems hailing from the domain of robotics, in this thesis, particular demand-driven uses of and additions to GP are described. In particular, an approach to discover dynamic process models from data is proposed, as well as a way to sub select samples from a data set to improve the learning process and its outcomes using only a few samples, and a multi-objective approach to incorporating problem-specific knowledge. The last chapter focuses on a different type of technology that leverages an analysis of observed features to detect changes in dynamic environments (e.g., mobile robots). All results pertain to well-designed and studied experiments and, in case of Chapter 3, even include a comparison with (modern) deep learning approaches.

I would however still like to make a few critical notes.

The form of GP that is used mostly throughout the thesis, is so-called single-node GP (SNGP). However, this form of GP is not the most common, or classical, form of GP, which would be tree-based GP. While in Chapter 3 MGGP is used, which uses some parts of classical GP, no comparison is made with classical GP anywhere in the thesis. While the omission of such a baseline to compare with is unfortunate, the lack of motivation for the choice of SNGP is a bigger issue in my opinion. Why is SNGP the best choice for the research in this thesis? What are the advantages over other types of GP? Would the conclusions change if another type of GP was used? Are parts of the conclusions and the recommendations that follow from this thesis dependent on the use of SNGP? Such questions are left unanswered, unfortunately.

It would also have been good to have read a bit more discussion about the weak points of GP in general. In particular, GP is not necessarily easy to scale to large input dimensions and can take a long time to find a good solution, even longer than deep learning. Related to this, it would have been good to have read a bit more reflection about the (complexity of) the scalability of GP.

Another issue is that it is well known that the performance of methods in EC are strongly dependent on the size of the population. However, in this thesis, this is left unacknowledged, and no motivation is given why certain population sizes are used. While the use of a single population size is not entirely uncommon in GP literature, a motivation for the choice is still important, be it from similar applications or from own (initial) experimentation. Results may be different if the population size (and accordingly, the maximum evaluations budget) are increased. Better (closer to optimal, within the representation limit) answers may then be found, which may influence the conclusions drawn about the applicability of GP to the situation. Conversely, if the population size is too small, the potential of GP may be underestimated. Hence, better motivation for the choice of population size and the impact of changing it, would have been a welcome addition.

Similarly, while I welcome the comparisons made with deep learning approaches in Chapter 3, I do feel that the exploration and motivation for the types of neural networks used, is rather shallow. It is stated what networks were used, but nothing is said about how the hyperparameters were chosen, which is equally important to setting the right population size for GP. This has a negative impact on the strength of the results.

Despite these critical notes, I cannot reach any other conclusion than that this PhD thesis is worthy of an academic defence. The main body of the thesis is based on published work in the form of 3 journal articles and 1 conference paper, which is a nice achievement for a PhD thesis in computer science. Moreover, even though I feel there are a few shortcomings, the writing of the thesis is very good and very clear. It stays on target and delivers on its promises.

The author of the thesis proved to have an ability to perform research and to achieve scientific results. I do recommend the thesis for presentation with the aim of receiving a Ph.D. degree.

Sincerely,

Prof.dr. Peter A.N. Bosman
Senior Researcher in the Life Sciences and Health Research Group
Centrum Wiskunde & Informatica (CWI)

Professor in the Algorithmics Research Group
Faculty of Electrical Engineering, Mathematics & Computer Science
Delft University of Technology