

Is Social Learning More Than Parameter Tuning?

Jacqueline Heinerman
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
j.v.heinerman@vu.nl

Jörg Stork
TH Köln
Köln, Germany

Margarita Alejandra Rebolledo
Coy
TH Köln

Julien Hubert
Vrije Universiteit Amsterdam

A.E. Eiben
Vrije Universiteit Amsterdam

Thomas Bartz-Beielstein
TH Köln

Evert Haasdijk
Vrije Universiteit Amsterdam

CCS CONCEPTS

•Computing methodologies → Evolutionary robotics; Multi-agent systems; •Theory of computation → Evolutionary algorithms;

KEYWORDS

Evolutionary robotics; social learning; neural networks; parameter tuning

ACM Reference format:

Jacqueline Heinerman, Jörg Stork, Margarita Alejandra Rebolledo Coy, Julien Hubert, A.E. Eiben, Thomas Bartz-Beielstein, and Evert Haasdijk. 2017. Is Social Learning More Than Parameter Tuning?. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017*, 2 pages. DOI: <http://dx.doi.org/10.1145/3067695.3076059>

1 INTRODUCTION

Social learning enables multiple robots to share learned experiences while completing a task. The literature offers examples where robots trained with social learning reach a higher performance compared to their individual learning counterparts [e.g., 2, 4]. No explanation has been advanced for that observation. In this research, we present experimental results suggesting that a lack of tuning of the parameters in social learning experiments could be the cause. In other words: the better the parameter settings are tuned, the less social learning can improve the system performance.

To test our hypothesis, we generated 50 parameter settings using Design of Experiments (DoE) and tested them in an individual learning configuration with a single robot (i.e., with social learning disabled). The definitive screening DoE was created with the help of JMP software (SAS Institute Inc, JMP, Version 11.1.0). The experiments are conducted in simulation using JBotEvolver [1]. The experiment requires the robots to learn a foraging task. The environment is a square arena. Five pucks are randomly placed in the arena at the start of a run. The robots must collect the pucks and bring them to the nest located in the centre of the arena. Once

a puck is brought to the nest, it is immediately replaced at a random location in the environment. The performance of each robot, i.e. its fitness, is equal to the number of pucks it collected during a trial lasting 1000 time steps. The robots use embedded instances of the NEAT evolutionary algorithm for on-line learning [3]. The robot's controller is an artificial neural network. The neural network has 11 input and two output nodes. The input nodes consist of 8 proximity sensors, a nest sensor, a puck sensor, and a puck carrying sensor; the output nodes provide the right and left motor speed.

From the 50 parameter settings of the DoE, we selected the 10 settings with best performance and 10 with median performance. We compared the performance of these two groups of settings, where social learning is disabled, with two social learning configurations of 2 and 4 robots. Social learning is implemented as follows: first, the robots sequentially evaluate the controllers in the current generation. Then, the robots exchange information. Each robot randomly selects another robot from which it receives its current best controller, i.e. the controller with the highest fitness. The robot compares the received controller's fitness to that of its own worst controller. The new controller replaces the worst controller if it is better. The NEAT algorithm uses the updated list of controllers and fitness values to create the next generation. When social learning is applied, these robots have the same parameter settings as the individual learning mechanisms except for the population size. The population size for the 2 and 4 robot setup is the population size from the 1 robot setup divided by the number of robots (e.g., when the 1 robot setup has a population size of 100, the social learning experiments used a population size of 50 and 25 for the 2 robot and 4 robot setup respectively).

The robots operate in their own arena but they communicate across arenas. Consequently, the performance of the robot is only due to its own actions and not influenced by other robots in the same arena. Removing this inter robot collision allows for a better comparison between the individual and the social learning experiments. For each experiment, 20 replicate runs are performed with different random seeds.

2 RESULTS

Our analysis shows that from the total of 21 investigated parameters from the NEAT learning algorithm, 14 parameters have a significant impact on the learning performance. The variable representing the chance to randomly reset a weight in the neural network has a

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '17 Companion, Berlin, Germany

© 2017 Copyright held by the owner/author(s). 978-1-4503-4939-0/17/07...\$15.00
DOI: <http://dx.doi.org/10.1145/3067695.3076059>




Figure 1: Mean performance with 95% confidence interval of the baseline experiments for all DoE parameter settings. The y-axis shows the performance, measured as the number of collected pucks. The x-axis shows the rank of the parameter setting. The results are compiled over 20 replicate runs.

big effect on the response and performance is best if it is turned off. Moreover, a large population size, a high mutation probability and a small value for the percentage of top individuals that can be parents have a significant beneficial influence on the fitness. This indicates that a large and diverse population of controllers seem to be advantageous for the learning rate.

Figure 1 shows the mean performance of the 1 robot performance at the final generation (generation 200). Performance, i.e., number of collected pucks, is plotted against the rank of the mean performance of the 50 parameter settings from the DoE. The data in figure 1 confirms that parameter settings significantly influence the performance of the controllers (Pearson's $r(50) = -0.9754128$, $p < 2.2e-1$).

Figure 2 shows the impact of social learning on performance for the best and median parameter settings. Ranks 1-10 refer to the best settings and ranks 21-30 to the median settings. Every setting is tested for a setup with 2 (red) and 4 (blue) robots in 20 replicate runs. The impact of social learning is measured as the ratio between the performance with social learning and the baseline performance. A ratio higher than 1 means an increase of performance due to social learning. There is a significant positive correlation between the rank number and the improvement ratio with social learning (for 2 robots Pearson's $r(20) = .56$, $p < 0.011$; for 4 robots Pearson's $r(20) = .63$, $p < 0.003$). Figure 2 shows that better parameter settings benefit less from social learning, indicated by performance ratio values around 1.

3 CONCLUSIONS

With this study we gained a better understanding of the benefits of social learning. Existing literature in social learning compares individual learning with social learning for only one parameter setting.




Figure 2: Performance ratio of 2 (blue) and 4 (red) robots compared to the baseline experiments. The y-axis shows the performance ratio, measured as the mean performance of 2 respectively 4 robots divided by the baseline performance. The x-axis shows the rank of the parameter setting (rank 1 is the rank resulting in the highest performance).

Results show an increased performance and increased learning speed. This study extended this comparison by using different parameter settings. We showed that the quality of the parameter settings influences how much social learning can improve the system performance: the better the parameter settings, the less social learning can contribute. Therefore, this study serves as a reminder that tuning the parameters can impact the conclusions drawn from an experiment. Nevertheless, tuning can be computationally expensive or not even possible when the optimal parameters may depend on the unknown environment that the robots operate in. Therefore, social learning can be a beneficial approach to increase performance and serve as an alternative to parameter tuning.

4 ACKNOWLEDGMENTS

This work is supported by the European Union's Horizon 2020 research and innovation programme under grand agreement No 640891 (DREAM project).

REFERENCES

- [1] Miguel Duarte, Fernando Silva, Tiago Rodrigues, Sancho Moura Oliveira, and Anders Lyhne Christensen. 2014. JBotEvolver: A versatile simulation platform for evolutionary robotics. In *Proceedings of the 14th International Conference on the Synthesis & Simulation of Living Systems*. MIT Press, Cambridge, MA, 210-211.
- [2] Jacqueline Heinerman, Massimiliano Rango, and A.E. Eiben. 2015. Evolution, Individual Learning, and Social Learning in a Swarm of Real Robots. In *Proceedings of the 2015 IEEE International Conference on Evolvable Systems (ICES)*. IEEE Press, New York, NY, USA, 1055-1062. DOI : <http://dx.doi.org/10.1109/SSCI.2015.152>
- [3] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99-127.
- [4] Wesley Tansey, Eliana Feasley, and Risto Miikkulainen. 2012. Accelerating evolution via egalitarian social learning. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, Terence Soule (Ed.). ACM, New York, NY, USA, 919-926.