Opponent's review of the dissertation thesis

Author: RNDr. Peter Krčah
Title: Evolution and Learning of Virtual Robots
Opponent: Mgr. Martin Pilát, Ph.D, KTIML MFF UK, Martin.Pilat@mff.cuni.cz

In his work, Peter Krčah extends the results from his master thesis and uses the HierarchicalNeat algorithm to evolve virtual creatures. The main contribution is the addition of lifetime learning to the search process which allows for faster convergence. Another contribution is the use of novelty search in the design of virtual creatures.

The thesis provides a nice introduction to evolutionary robotics and to HierarchiicalNEAT in Chapter 2. Especially, the HierarchicalNEAT algorithm is described in a great detail.

The main contribution of the thesis is in Chapters 3 and 4. The third chapter includes the description of the algorithm for the design of virtual creatures based on HierarchicalNEAT and lifelong learning. The lifelong learning consists of the changing of sizes of the parts of the virtual being. To speed the learning up, only a short (2 seconds) simulation is used instead of the long one, which is used to compute the fitness in the main evolutionary algorithm. The learning is also performed by means of hill climbing instead of an evolutionary algorithm and only 8 iterations of the algorithm are performed. The author demonstrated that the addition of lifelong learning improves the quality of the individuals found by the algorithm in a given time. The author provides a throughout explanation of the algorithm and detailed description of the tuning of the parameters used. On the other hand, the reason for the selection of hill-climbing as the learning algorithm is discussed only briefly. It is only mentioned that it works better than random search. The results of the experiments are deeply statistically evaluated and nicely commented. I have a few questions to this part:

1. Did you try using any other methods than hill-climbing and random search in the learning phase?

2. I believe that interesting results could be obtained if instead of the learning phase the HierarchicalNEAT algorithm used the faster evaluation instead of the full evaluation in some of the generations. For example, if it used the fast evaluation in 8 generations and then used the full evaluation in one generation, the overall computational time would be the same as in the version with learning. Have you considered such a setting? It appears (to me) more natural than the addition of learning and it could also provide a nice baseline that would show the learning is in fact useful.

In the fourth chapter, the author describes the use of novelty search in the design of virtual creatures. Novelty search is a technique which discards the objectives and aims only at finding novel behavior. In this case, the novelty is defined as the final position of the individual. The author demonstrates that the creatures evolved by novelty search can outperform those evolved by a fitness-based search in cases where the fitness in deceptive. The author also tries to combine these two methods by switching from novelty search to the fitness-based search after a given number of generations. In my opinion, the switching from novelty search to the fitness-based search is a

nice idea, but I would suggest using the best individuals (according to the fitness) in the initial generation of the fitness-based search instead of the last generation of the novelty search. Again, I have a few questions to this part:

1. Have you considered using the best individuals found by the novelty search instead of the last generation? It would seem more reasonable.

2. On a related note, how many individuals from the last generation of the novelty search can avoid the obstacle?

Overall, the work is written in a good English, and all the parts contain all important information. The experiments in the thesis are performed with a focus on detail and are statistically evaluated.

The weakest point of the thesis seems to be the fact that the its actual contribution seems rather small compared to other works I have seen so far. The introduction of lifelong learning is a rather simple extension of the HierarchicalNEAT algorithm. The application of the novelty search in this case also seems rather straightforward and the algorithm is not tuned for the problem at hand in any way. On the other hand, all the experiments presented in the thesis are performed in a very careful way, various possibilities are tested and rigorously evaluated. Moreover, the HierarchicalNEAT algorithm is also created by the author (in his master thesis).

Therefore, I believe the author proved in his thesis he is capable of high quality scientific work. I recommend the work for defense and I believe the author should be awarded the Ph.D. degree.

In Prague August 5, 2016. Martin Pilát, opponent