Adaptive Behavior

http://adb.sagepub.com

Reply to Dario Floreano's "Engineering Adaptive Behavior"

Marco Dorigo and Marco Colombetti

Adaptive Behavior 1997; 5; 417

DOI: 10.1177/105971239700500310

Reply to Dario Floreano's "Engineering Adaptive Behavior"

Marco Dorigo*
Université Libre de Bruxelles

Marco Colombetti†
Politecnico di Milano

Dario Floreano's review of *Robot Shaping: An Experiment in Behavior Engineering* is very stimulating and deserves a long answer. However, we do not want to abuse our privilege of having the last word, which the editor of this special issue has accorded to us. Therefore, we shall limit our reply to a few points: (1) the relatively high learning time required by our system, (2) the urge for more sophisticated sensors, (3) the possibility of exploiting a human trainer, (4) the issue of modularity, and (5) the role and stage of development of the BAT methodology.

s ALECSYS a Slow Learner?

We believe that a comparison based only on learning speed between our experimental results and those obtained by Nehmzov and McGonigle (1994) is not fair for at least three reasons. First, Nehmzov and McGonigle's learning system is a pattern associator that can learn only linearly separable functions, whereas ALECSYS is a more general system that does not have this limit: A possible lower speed of ALECSYS is counterbalanced by its greater generality.

Second, ALECSYS learns by reinforcement learning, whereas Nehmzov and McGonigle's pattern associator learns by supervised learning. Again, ALECSYS is more general and also more indicated for robotics applications, wherein labeled training pairs are difficult, if not impossible, to provide. In fact, although Nehmzov and McGonigle devised a clever way to automatically generate training pairs, their approach is feasible only if the number of possible actions for each input pattern is very small, as is the case in their experiments: They have 3 possible actions, as opposed to the 16 we use in most of our experiments; obviously, the greater the number of actions, the smoother is the resulting movement of the robot.

Third, it is very difficult to compare the relative learning speed of different robots, particularly when no precise quantitative measure is used. In fact, in Nehmzov and

Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle (IRIDIA), Université Libre de Bruxelles, Avenue Franklin Roosevelt 50, 1050 Bruxelles, Belgium; E-mail: mdorigo@ulb.ac.be; URL: http://iridia.ulb.ac.be/dorigo/dorigo.html

Artificial Intelligence and Robotics Project (PM-AI&R Project), Dipartimento di Elettronica e Informazione, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy; E-mail: colombet@elet.polimi.it; URL: http://www.elet.polimi.it/people/colombet/

Adaptive Behavior Vol. 5, No. 3/4, 417–420

McGonigle (1994, section 3.1) we read: "The teaching process takes under one minute." Unfortunately, they do not say how they differentiate between a robot that has learned a behavior and one that has not yet. Consider a robot, such as AutonoMouse II, which has four possible motor moves for each motor (move one step backward, stay still, move one step forward, move two steps forward). As we explain at the end of section 5.2.4.1 of the book, "The robot starts . . . to approach the light source much before it has reached a high frequency of correct moves." This is easily explained if you recall that the correct behavior is more frequent than the wrong one as soon as performance is higher than 50 percent. As a consequence, a human observer would say that the robot has learned the approaching behavior well before the robot performance is 100 percent correct. With AutonoMouse II, for example, this happens very quickly, in less than the 1 minute reported by Nehmzov and McGonigle. In contrast, learning could take much longer if we say that a behavior is learned when the robot is performing correctly all, or most, of the moves, and that it is not yet learned in all other cases.

In conclusion, on the one hand it is not completely clear that ALECSYS is less efficient than the pattern associator proposed by Nehmzov and McGonigle; on the other hand, ALECSYS's higher degree of generality could justify its slightly lower performance, if such exists.

2 Do We Need More Sophisticated Sensors?

When Floreano says, "In nature, performance accuracy often relies on smart combinations and processing of information coming from a set of redundant, simple, and noisy sensors," he probably is thinking of simple organisms such as insects; indeed, we would classify birds' vision and many mammals' olfactory systems as very sophisticated equipment. However, such sensors are extremely well adapted to carrying out a specific kind of behavior in a specific type of environment; in other words, they are not just simple and noisy, but they are so in exactly the right way. Trying to adopt a similar solution in the realm of robotics might not be a sensible engineering choice for at least two reasons: (1) Cheap commercial sensors might be simple and noisy *in the wrong way*—that is, they might lack the capacity of adapting to the kind of behaviors and environments we have in mind for our robots; and (2) the effort to develop the appropriate integration of redundant information might exceed the cost of more sophisticated sensors.

In our experimental activity, we often found ourselves longing for richer sensory devices. Certainly, much can be done by giving our agents some limited image-processing capability, an objective that we have begun to pursue. However, we agree that trying to exploit cheap commercial sensors at their best is itself a very interesting research topic.

:xploiting Human Trainers

We agree with Floreano that "it would be desirable for the system to be open to reinforcements externally given, when necessary, by a human trainer." From the point of view of reinforcement learning, the problem is to find a way of integrating immediate reinforcement coming from the reinforcement program with sparse, delayed reinforcements coming from a different source. Our experience is that current reinforcement learning methods are not robust with respect to the discrepancies that inevitably arise in similar situations (see, for example, Caironi & Dorigo, 1994). Again, we face here an interesting topic for future research.

'he Issue of Modularity

The role of modularity in classical design is reasonably well understood. Also the idea that "biological organisms are ... modular at several levels" can be accepted, even though many biologists probably would disagree sharply if we tried to provide a detailed description of an organism's modules. However, before we try to find out "how to chop, link, and coordinate modules automatically and autonomously," we would like to understand whether a *lack* of modularity, at least at some level, is essential for the success of natural systems.

The agents we have developed thus far are all modular at the level of behavior organization. However, the use of evolutionary techniques, at least in principle, opens the way to forms of nonmodular design. We argue in the final chapter of our book that one of the main concerns of artificial intelligence is finding ways to cope with problems that do not admit of modular solutions. If this view is correct, lack of modularity and what we usually call *intelligence* might be connected in some important way that we still do not understand. We find this idea fascinating and worth some research effort.

'he Behavior Analysis and Training (BAT) Methodology

Our reviewer is surprised to find the BAT methodology described at the end of the book. However, at least in the form described in *Robot Shaping*, the methodology did take its present form only very recently. BAT was not formulated a priori: It resulted from our direct experience and from the analysis of our own mistakes. By postponing its presentation, we certainly saved some time but, more importantly, we did not want to give the reader the false impression that a systematic approach can be devised before a lot of exploratory work has been carried out.

Finally, let us remark that Floreano is quite right in criticizing BAT's lack of concern for the economical and technological facets of behavior engineering. To add a further criticism, our methodology is clearly insufficient from the point of view of quality metrics. Our only excuse is that we did not have enough energy, time, and professional competence to cover these essential aspects. Given the importance of the subject, we hope to contribute to its development in the near future. (A first step in this direction can be found in Colombetti & Dorigo, 1997.)

References

Caironi, P. V. C., & Dorigo, M. (1994). *Training and delayed reinforcements in Q-learning agents* (Tech. Rep. No. IRIDIA/94-14). Belgium: Université Libre de Bruxelles.

Colombetti, M., & Dorigo, M. (1997). Evolutionary computation in behavior engineering. In S. Yao (Ed.), *Evolutionary computation: Theory and applications*. Singapore: World Scientific.

Nehmzov, U., & McGonigle, B. (1994). Achieving rapid adaptation in robots by means of external tuition. In J. Meyer, H. L. Roitblat & S. W. Wilson (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.