The Biology and Technology of Intelligent Autonomous Agents

Luc Steels

Artificial Intelligence Laboratory, Vrije Universiteit Brussel

The NATO-sponsored Advanced Study Institute "The Biology and Technology of Intelligent Autonomous Agents" was an extraordinary event. For two weeks it brought together the leading proponents of the new behavior-oriented approach to Artificial Intelligence in Castel Ivano near Trento, Italy. The goal of the meeting was to establish a solid scientific and technological foundation for the field of intelligent autonomous agents with a bias towards the new methodologies and techniques that have recently been developed in Artificial Intelligence under the strong influence of biology. Major themes of the conference were: bottom-up AI research, artificial life, neural networks and techniques of emergent functionality. The meeting was such an extraordinary event because it not only featured very high quality lectures on autonomous agents and the various fields feeding it, but also robot laboratories which were set up by the MIT AI laboratory (with a lab led by Rodney Brooks) and the VU-B AI laboratory (with labs led by Tim Smithers and Luc Steels). This way the participants could also gain practical experience and discuss in detail what the difficulties and achievements were of different approaches. In fact, the meeting has been such a success that a follow up meeting is planned for September 1995 in Monte Verita (Switzerland). This meeting is organised by Roll Pfeifer (University of Zurich) and results of the meeting are planned to be published in one of the future issues of the journal. This special issue of Robotics and Autonomous Systems contains a representative sample of material presented at the Trento Advanced Study Institute. Papers have been retained that are of a more technical nature and directly relevant to the "Robotics and Autonomous Systems" community. Moreover the emphasis has been put on the problem of acquiring new behavior - as this is still the major bottleneck in building truly autonomous agents. A more complete collection of papers has been published as a book . This collection starts with a paper by Steels which is intended to illustrate the line of research advocated at the institute, namely an emphasis on biology. The paper attempts to characterise the notion of intelligent agents using concepts from biology. It can be seen as setting targets for what the field should accomplish. Then some overview papers representative of the main lectures at the institute have been included:

- Torras gives an overview of neural network techniques for robot adaptivity.
- Thrun and Mitchell survey the application of machine learning techniques to support life-long robot learning.
- Pfeifer critically examines the cognitive viewpoint and whether it is adequate for building intelligent autonomous agents.

Other overview lectures were given at the institute by David McFarland (biology of behavior), Agnessa Babloyantz (self-organisation), Rodney Douglas (biological neural networks), Peter Schuster (biological evolution), Rodney Brooks (robot anatomy), Tim Smithers (autonomy), and Walter Van de Velde (cognitive architectures). Next there is a series of research papers which are a direct consequence of the discussions, experiments, and presentations at the Trento gathering.

- The paper by Nehmzow (Animal and robot navigation) is an example of the kind of crossfertilisation that may take place between the biology and technology of intelligent agents. Nehmzow combines different mechanisms (landmarks, canonical paths, topological maps) in order to get robust navigation in mobile robots.
- The paper by Husband, Harvey and Cliff (Circle in the round: State space attractors for evolved sighted robots) uses mechanisms from evolution by natural selection to derive new robot behavior. The work is original particularly because it evolves visual morphologies as well as control networks.

- The paper by Smithers (On quantitative performance measures of robot behaviour) introduces a methodology for developing a quantitative approach towards the characterisation of performance.
- The paper by Weil] (Distributed reinforcement learning) introduces a novel application of reinforcement learning algorithms, namely in the direction of use by multiple agents.
- Finally the paper by Hasslacher and Tilden (Living machines) is a manifesto for merging biological principles and robotic engineering. The authors also give examples of various "biomorphs" which illustrate their approach.

Acknowledgement:

We thank Dr. L. Veiga da Cunha, director of the NATO ASI program for the support throughout the whole administrative process associated with the Nato funding.

This funding is hereby gratefully acknowledged.