

Report on the Eighth European Workshop on Learning Robots (EWLR'99)

Dolores Cañamero¹, Pier Luca Lanzi², Marnix Nuttin³, Jeremy Wyatt⁴

²Artificial Intelligence Research Institute (IIIA)
Spanish Scientific Research Council (CSIC)
E-08193 Bellaterra, Barcelona, Spain

`lola@iiaa.csic.es`, `http://www.iiaa.csic.es/~lola`

²Artificial Intelligence and Robotics Project
Dipartimento di Elettronica e Informazione
Politecnico di Milano, Milano, Italy

`lanzi@elet.polimi.it`

³K.U.Leuven - Department of Mechanical Engineering
Division of Production Engineering, Machine design and Automation (PMA)
Celestijnenlaan 300B, B-3001 Heverlee, Belgium
`marnix.nuttin@mech.kuleuven.ac.be`

⁴School of Computer Science
Birmingham University
Edgbaston, Birmingham
`jlw@cs.bham.ac.uk`

The Eighth European Workshop on Learning Robots (EWRL99) was held September 18th 1999 at the *Ecolé Polytechnique Fédérale de Lausanne* in Switzerland. The workshop, which started immediately after the European Conference on Artificial Life (ECAL-99), brought together some of the most active researchers who work on Machine Learning techniques for autonomous robotics applications. The workshop had nine presentations and two invited talks, delivered by Henrik Lund of Aarhus University (Denmark) and Jun Tani of Sony Computer Science Lab of Tokyo (Japan). This report outlines the work presented at the workshop and also outlines a possible road map for robot learning research arising out this.

1 General Overview

The works presented during the workshop concerned different and varied applications of robotics, ranging from “*classical*” navigation-related tasks (e.g., homing, map building, light-seeking, and landmark selection), to “*higher-level*” cognitive tasks (e.g., planning), or to more “*uncommon*” robotics applications such as entertainment and education.

1.1 The Invited Presentations

There were two invited talks: one by Henrik Lund of Aarhus University, one by Jun Tani of Sony Computer Science Laboratory.

Henrik Lund presented some educational applications of Evolutionary Robotics (ER) developed with LEGO Mindstorms. Lund introduced a tool developed in the LEGO laboratory of Aarhus University, called *Toybots*, which allows children to play with Evolutionary Robotics and LEGO Mindstorm. Toybots is a tool for supporting the evolution of populations of robots (e.g. pets); the users (mainly children) can select the behaviors they like better while the robots are acting in the real world. Toybots is basically an interface to a user-guided genetic algorithm in which the user selection substitutes the classical formal fitness function. This approach allows for personalized evolution according to the child's tastes, and at the same time letting children learn about evolution in an intuitive way. The principle of using minimal simulation coupled with reinforcement learning-based re-training in the real world helps to bridge the gap between simulation and reality. Complex behavior is achieved using an evolvable behavior-based system, where both the basic behavior modules and the behavior arbitrators are evolved by the children. These ideas have been applied in the Interactive LEGO-Mindstorms Football, where children can select behavioral sequences to configure their players, and in the Adaptive LEGO Pet Robot. Some videos from the demonstration made during IJCAI99 were presented.

Jun Tani presented an overview of the research he has been carrying out with mobile robots in the last decade. The long term goal of Tani's research is to explain structural coupling in a qualitative manner by using the basic elements of dynamical systems. Tani showed how his research progressed "from simple to complex" in three main stages. At the early stage of his research, he investigated reactive behaviors within context using steady strategies for navigation and homing tasks. Subsequently, he moved to studying internal modeling and its situatedness using robots that become situated through travel and can perform one-step predictions for the future states through mental rehearsing. The underlying idea here is a relational view of embodiment as structural coupling. In the last two years, he has been exploring the open dynamics of this coupling, and in particular how to deal with incompleteness in learning through exploration. Some of the problems he has studied are forward modeling for planning with incomplete knowledge in real time, and novelty rewarding using the prediction error as rewarding, in order to have the model predict how much the robot can predict.

1.2 The Contributed Presentations

The applications presented in contributed papers were at least as much existing as invited ones. Here, we shortly summarize the contributed presentations grouping them by applications while the next application will focus on the techniques used to implement learning within the robotic

Navigation. Maze navigation was the application domain chosen by Inuzuka, Onda, and Itoh to test their ILP-based algorithm to learn control rules through iterative collection of examples. Two other talks were concerned with navigation problems. Davesne and Barret have developed Constraint-based Memory Units (CbMU), a new kind of learning agents to incrementally build complex behavior from a set of basic tasks and perceptive constraints on the behavior. The Khepera simulator was used to test these agents on a goal-seeking navigation task in maze-like environments with different configurations of walls and obstacles. Visually-guided landmark learning and navigation by a mobile (Nomad200) robot in natural (not pre-engineered) indoor environments was the application used by Bianco and Cassinis to test their biologically-inspired learning algorithm.

Map building. This topic was addressed by several speakers. Quoy and Gaussier presented joint work with other colleagues on a neural model for trajectory planning by mobile robots based on the elaboration of a “cognitive” planning map. In a simulated animat environment containing a nest, obstacles, landmarks, and resources (food and water), a robot facing an action selection problem builds a planning map to learn paths that allow it to better satisfy its internal motivations—hunger, thirst, and rest. The elaboration of occupancy maps by a robot navigating an indoors, laboratory environment, was the problem addressed by Rodríguez and colleagues. These authors proposed two new count-based methods that proved computationally very efficient when compared with well known probabilistic and histogrammic methods.

Robot pushing task. Ishiguro and colleagues discussed their current investigation which involves adaptive controllers for autonomous mobile robots based on the idea of dynamic rearrangement of neural networks. Taking inspiration from neuromodulators, they evolve neural networks which are able not only to adjust their weights, but also their own structure by blocking and/or activating synapses and neurons. Simulations applied to a peg-pushing task for a Khepera robot resulted in greater adaptability and better results than conventional evolutionary robotics approaches where only synaptic weights are evolved.

Walking. Mark Pendrith used the problem of training a six-legged robot to walk as the testbed for his novel algorithm, *ccBeta*, which he introduced to improve the performance of Reinforcement Learning algorithms in real world applications.

Feature Extraction. Research on adaptive controllers is also being conducted by Hosoda (in joint work with Asada), in this case with the purpose of having robots with multiple degrees of freedom, multiple sensors, and which act in a dynamic environment find by themselves the number of DOF needed to accomplish a given task. They propose an architecture consisting of several subcontrollers, where each subcontroller is responsible for some of the DOF and must find out its own redundancy so that the redundant DOF can be spared to other controllers. They presented an application of this architecture to a visual servoing control task.

2 The Learning Techniques

In the works presented at the workshop we identified three main learning techniques.

Neural Networks. Peter Eggenberger and colleagues al. followed a dynamically rearranging neural network approach. They were inspired by the existence of neuromodulators and a 1991 article of Meyrand. Several kinds of neuromodulators spread through the network and their effect depends on the specific receptors which can be different for different neurons. This approach allows to adapt not only the weights of a neural network, but also its structure. Their approach allows also to correlate neural activity from distant cells.

The inspiration of this approach comes from biology, where neuroscientific results suggest that biological networks may not only adjust the weights but also the neural structure. E.g. the stomatogastric network of a lobster consists of three networks (oesophageal, pyloric and gastric), which usually show independent behaviour. However, when eating, these networks can rearrange to form a new one: the swallowing network.

In the simulation part, an evolutionary approach is followed using genetic algorithms. The following mechanisms are evolved: Diffusion of neuromodulators (when and which type of

neuromodulators are diffused from each neuron), and the reaction to neuromodulators (how do the receptors on each synapse interpret the received neuromodulators). E.g. depending on the received neuromodulators, the mechanism for updating the synaptic weights is either Hebbian learning, anti-Hebbian learning, non-learning or blocking. After the simulation part, the evolved controller is used in real-world. In their approach, the authors attempt to bridge the gap between simulation and real-world.

M. Quoy, P. Gaussier et al. present a neural approach for mobile robot trajectory planning. Their approach can be situated as follows. An animat moves around in an environment. There are energy levels (food, water and sleep) and motivations (hunger, thirst and rest). The network consists of two levels. There is a recognition level and a goal level. The neurons in the recognition level represent locations in the environment and respond more if the animat is closer to that particular location. In the second level a graph is built linking reachable places. There is no external explicit description of the environment.

Evolutionary Computation During the workshop, three presentations concerned the use of evolutionary computation techniques for learning robot controllers. In particular two of this presentations discussed an “hot” topic of evolutionary robotics, i.e., the delicate equilibrium between evolution in simulated and real-world environments. Evolutionary techniques usually necessitate a large number of experiments before adequate behaviors are evolved. Accordingly, Evolutionary Robotics usually involves an initial step in which controllers are trained in simulation and then they are *refined* in the environment. Unfortunately, the performance of controllers evolved in simulation usually drop when the controller are used in the real environment. To alleviate this problem, Peter Eggenberger and colleagues al. (see also the above discussion) evolve neural networks which can adapt both their synaptic weights and *also* their structure through a mechanism of inhibition and activation of synapsis. They tested their solution with a set of experiments on a peg pushing task with a Kephera robots showing promising results. Lund argued that it is *possible* to evolve controllers directly in the real world by providing a tool for letting users assist the evolution of robot controllers. He supported his claim showing how children could interactively guide the evolution of controllers for LEGO Mindstorm robots. Finally, with a more theoretical approach, Ficici and colleagues presented a work on embodied evolution, a new evolutionary robotics methodology where a population of physical robots evolve by reproducing with one another in the task environment in a distributed, asynchronous, and autonomous manner. They showed an application of this methodology in a population of small mobile robots performing a light-seeking task, where results outperformed a designed solution previously proposed by these authors.

Reinforcement Learning. Mark Pendrith gave a talk which had a more theoretical flavor than the other ones. In particular, Pendrith showed how some of the fundamental hypotheses of theoretical reinforcement learning are inapplicable in practice. In particular he discussed the role of the learning rate in applications with real world applications which usually involve noisy sensory inputs. He argued that the agent’s learning rate should adapt through the interaction with the environment instead of being fixed (as usually done in practice) or converging asymptotically to zero (as required by the theory). Pendrith introduced an algorithm, *ccBeta*, to adapt the learning rate by using a local estimator of the variance of agent’s expected payoff. The new *ccBeta* algorithm was applied in simulated environment and to the problem of training a six-legged robot to walk.

3 A Roadmap for Robot Learning Research

What indications do the papers presented at this workshop – together with those presented at the companion event ECAL'99, and reports of other recent advances – give for the direction of robot learning research?

Broadly speaking there currently seem to be five popular classes of techniques or mathematical languages for talking about robot learning: neural approaches, evolutionary approaches, reinforcement learning approaches, probabilistic approaches and control theoretic approaches. These overlap to some considerable degree, and one of the trends in this workshop was toward integrating techniques in order to solve more challenging tasks. These classes of techniques are likely to remain highly significant in the near future.

There have been significant advances in several of the underlying technologies, e.g. support vector machines and graphical models in neural networks [2]; improved links between reinforcement learning and stochastic control theory [1]; advances in planning and learning methods for stochastic environments [3, 5]; and improved theoretical models of simple genetic algorithms [9]. Some of these advances have been driven by robot learning researchers; many others have yet to be exploited, and this is an important area for robot learning research in the near future.

Robot learning algorithms are now being used in some real-world applications [7, 4]; although these are still limited to domains such as entertainment and education. Real-world applications, whatever the domain, are important in driving the integration of techniques from separate fields. Competitions, such as the AAI robot competitions and Robocup are important in this respect and also in that they enable researchers to compare the performance of a wide range of techniques (albeit on very specific tasks).

3.1 General Research Questions

There are an enormous number of important long term general research questions which need to be tackled. There are two many to be listed here, but they include:

- How can learning techniques be integrated the principled manner with prior knowledge (in all its forms) in real robot tasks? These could include knowledge learned by other robots, human skills, performances measures, partially specified process models, hand-coded controllers etc.
- How can disparate learning algorithms be integrated to provide complete robot learning systems?
- What bearing do robot learning algorithms have on the problems of sensory processing and sensory fusion?
- How can general learning algorithms be adapted to solve specific problems in robotic applications? e.g. Robot human interaction (learning by imitation), industrial processes employing robots, autonomous driving.
- What further insights can robotic implementations provide into the functioning of biological learning systems? This has been an important and active area of robot learning research to date. Such work also continues to provide inspiration for engineers.

- How can robot learning algorithms be used in multi-robot systems? What role do they have to play in the development of communication between robots, and the integration of information (in for example multiple robots used for map building).

3.2 Some Hot Topics in Robot Learning

There are a number of specific issues of particular interest. This list should in no way be seen as exhaustive. Many of the questions here can be seen as instantiations of the general questions above, in the context of specific techniques and/or task domains.

Probabilistic methods and reinforcement learning There have been a number of significant advances in the interface between these fields. In particular there has been recent work linking learning and planning through the theory of stochastic processes. There is now also a theory and a class of algorithms for learning hierarchical controllers. A good deal of work has been done implementing probabilistic methods in robots [8], and some regard this approach as a promising one for a framework for robot control.

Robot Navigation and Map Learning Map learning and map use is one of the most well studied problems in robotics. Some recent opinion suggests the navigation problem has essentially been solved. However, while sonar based approaches to navigation and map building, using probabilistic or neural techniques have been successful, the relationship between these methods and vision based navigation are still not completely clear. This is related to the general problem of sensory fusion.

Neural networks and evolutionary methods Researchers in evolutionary computing are currently interested in how the genotype unfolds into a phenotype through the processes of morphogenesis and learning. This dovetails neatly with work in evolutionary robotics, which seeks to improve the power of evolutionary algorithms, by evolving neural networks which are themselves adaptive. The state-of-the-art is arguably the sort of work presented by Eggenberger et al. in the workshop [10].

Learning and robot-human interaction Entertainment is currently one of the most important commercial applications of agent based technology, and particularly of robot technologies. Some researchers have argued for the use of humanoid robots as a means of facilitating interaction. Learning techniques (e.g. learning by imitation) have an important role to play in such systems.

Cognitive Robotics While numeric learning methods have dominated robot learning to date, a significant amount of work is now being done in the interface between robots and logic based representations. This work is sometimes known as cognitive robotics and there have been recent advances [6].

3.3 Summary

In summary while it is obviously difficult to set out a roadmap for robotics research with a great degree of confidence, we believe that the main thrust of work will focus around themes outlined above, namely:

- Testing the utility of new techniques in Machine Learning in robot domains.
- Integrating different learning techniques into complete robot learning systems.
- Integrating learning with prior knowledge.
- Using robot learning techniques to solve specific problems in real world applications.

References

- [1] Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996.
- [2] Nello Cristianini and John Shawe-Taylor. *An Introduction to Support Vector Machines : And Other Kernel-Based Learning Methods*. Cambridge University Press, 2000.
- [3] Michael Lederman Littamn. *Algorithms for Sequential Decision Making*. PhD thesis, Brown University, March 1996.
- [4] Henrik Hautop Lund and Luigi Pagliarini. Robocup jr. with lego mindstorms. In *Proceedings of Int. Conf. on Robotics and Automation (ICRA2000)*, 2000.
- [5] Ronald Parr. *Hierarchical Control and Learning for Markov Decision Processes*. PhD thesis, Computer Science, Berkeley, 1998.
- [6] Murray Shanahan and Mark Witkowski. Robot navigation and map building with the event calculus. In *Working Notes of the IJCAI 99 Workshop on Robot Action Planning*, 1999.
- [7] S. Thrun, M. Bennewitz, W. Burgard, A.B. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy, J.Schulte, and D. Schulz. Minerva: A second generation mobile tour-guide robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [8] Sebastian Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 2000.
- [9] Michael Vose. *The Simple Genetic Algorithm : Foundations and Theory*. Bradford Books, 1999.
- [10] Jeremy Wyatt and John Demiris, editors. *Advances in Robot Learning: Proceedings of the 8th European Workshop on Learning Robots*. Springer Verlag, forthcoming.