Available online at www.sciencedirect.com

ScienceDirect

Cognitive Systems Research 36–37 (2015) 15–29

Cognitive Systems RESEARCH

www.elsevier.com/locate/cogsys

A single computational model for many learning phenomena

Action editor: Vasant Honavar

Giancarlo Petrosino *, Domenico Parisi

Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

Received 30 October 2014; accepted 2 June 2015
Available online 11 June 2015

Abstract

Simplicity is a basic principle of science and this implies that, if we want to explain the behaviour of animals by constructing robots that behave like real animals, one and the same robot should reproduce as many behaviours and as many behavioural phenomena as possible. In this paper we describe robots that both evolve and learn in their "natural" environment and, in addition, learn in the equivalent of an experimental laboratory and reproduce a variety of results of experiments on learning in animals. We introduce a new model of learning in which the weights of the connections that link the units of the robots' neural network are genetically inherited and do not change during the robots' life but what changes during life and makes the robots learn new behaviours is the synaptic receptivity of a special set of network units which we call learning units. The robots evolve in a variety of different environments and they learn in a variety of different ways including imprinting and learning by imitating the behaviour of others. Then we test the robots in the controlled conditions of an artificial laboratory and we reproduce a number of experimental results on both operant learning and classical conditioning, including learning and extinction curves, the role of the temporal interval between conditioned and unconditioned stimuli, and the influence of motivation on learning.

Keywords: Robots; Learning; Evolution; Neural network

1. One robot, many phenomena

Robots are a new way of expressing theories of behaviour. If a robot reproduces the behaviour of one particular animal, the theory which has been used to construct the robot explains the behaviour of the animal. The problem with current robots is that the same behaviour can be reproduced by different robots and it is difficult to decide which robot incorporates the correct theory. To solve this problem we suggest one should apply the principle "one robot, many phenomena". If one and the same robot reproduces many different behaviours and many different aspects of these behaviours, we can be more confident that the robot is not a "toy" and it really helps us to understand and explain the behaviour of animals.

In this paper we apply the principle "one robot, many phenomena" to the study of how animals acquire their behaviour. Robots as practical applications can be programmed by us to exhibit some desired behaviour. But robots as theories of behaviour cannot be programmed by us because real animals are not programmed by anyone and they acquire their behaviour as a result of evolution, development, and learning. Evolution is changes in the genes of a population of individuals due to the selective reproduction of the best individuals and the addition of random mutations to the genes that offspring inherit from parents. Development is changes that take place during the life of an individual and that are encoded in the individual's inherited genes. Learning is changes due to the particular

* Corresponding author.
E-mail address: giancarlo.petrosino@istc.cnr.it (G. Petrosino).

http://dx.doi.org/10.1016/j.cogsys.2015.06.001

experiences that the individual has during its life. The principle "one robot, many phenomena" requires that robots acquire their behaviours as a result of all these different processes.

Most current robots acquire their behaviour through *either* evolution *or* learning but not both. (In this paper we ignore development. For robots that develop during their life, see Cangelosi & Schlesinger, 2014.) Some robots evolve their behaviour because the connection weights and other parameters of the neural network that controls their behaviour are encoded in the genes they inherit from their parents and the genetic pool of the population changes from one generation to the next. The behaviour of these robots evolves in a succession of generations but it does not change during an individual's life. Other robots learn during their life because the connection weights of their neural network change as a result of the particular experiences that a robot has during its life but these robots do not have inherited genes and they start their learning from a random set of connection weights for their neural network. The principle "one robot, many phenomena" requires that robots *both* evolve *and* learn.

Some research has been dedicated to robots that both evolve and learn but these robots do not evolve and learn in what we might call their "natural" environment (Baxter, 1992; Chalmers, 1990; Dasdan & Oflazer, 1993). In the natural environment the movements with which animals respond to the sensory inputs to a large extent determine the successive sensory inputs that arrive to their sensory organs, and this sensory-motor loop is a crucial aspect of behaviour: animals are not passive receivers of sensory input (Parisi, 1994). In the experimental laboratory the inputs that arrive to the individual's sensory organs are decided by the experimenter and there is no sensory-motor loop. The best empirical results collected by psychologists are obtained in the experimental laboratory and this is inevitable because collecting empirical data in the real ecology of animals is both difficult and expensive. But this problem can be solved with robots. Today's robots either evolve in what we may call their natural environment or they learn in the equivalent of an experimental laboratory. This is a limitation that robots as scientific theories will have to overcome because real animals *both* evolve *and* learn in their natural environment and, therefore, robotics should be an ecological robotics (Arkin, Cervantes-Perez, & Weitzenfeld, 1998; Duchon, Kaelbling, & Warren, 1998; Parisi, Cecconi, & Nolfi, 1990). Individual robots that learn in the equivalent of a psychologist's experimental laboratory tend to exhibit different behaviours because the connection weights that have been assigned to their neural network at the beginning of learning are random, not because they have different genes or they have had different experiences. Real animals are inter-individually different because they both inherit different genes and have different experiences during their life. The interaction between evolution and learning has been studied by using disembodied neural networks or robots whose behaviour is controlled by a neural network, and these studies have shown that neural networks evolve the capacity to learn if the environment is unstable (Todd & Miller, 1991) and robots obtain better results if the initial connection weights of their neural network are evolved and not random (Nolfi & Parisi, 1996). Learning in the "natural" environment has been studied by using different learning algorithms such as actor-critic (Barto, 1995; Schembri, Mirolli, & Baldassarre, 2007), integrated CTRRN (Tuci, Quinn, & Harvey, 2002), and reinforcement learning (Niv, Joel, Meilijson, & Ruppin, 2002), and Floreano, Dürr, and Matiussi (2008) have proposed an interesting interaction between evolution and learning in which a neural network evolves both its architecture and its capacity to learn.

In this paper we apply the principle "one robot, many phenomena" by having our robots both evolve and learn in their "natural" environment and, after they have evolved and learned in their "natural" environment, we bring them into an experimental laboratory and we replicate some of the results obtained by psychologists in experiments on learning. Referring to computational models of learning, Balkenius and Morén have written that "it is very hard to construct a model that covers a large area of experimental conditions" (Balkenius & Morén, 1998). This is an important point. Experiments on learning in real animals generate many different types of data. Learning curves and extinction curves are only one example. Other examples are time constraints on learning such as the temporal interval between pressing a lever and the appearance of food in operant learning or between the sound of the bell and taste in the mouth in classical conditioning. Furthermore, learning depends on motivation. A mouse will learn to press a lever to obtain food but this will only happen if the mouse is hungry and, therefore, is motivated to eat. Another phenomenon which is linked to motivation is devaluation. After the mouse has learned to press a lever, the mouse will not press the lever unless it is hungry. Our robots try to reproduce *all* these types of data.

Another important way in which our robots try to adhere to the principle "one robot, many phenomena" is that we use the same model of learning for different types of learning: sensitization, habituation, reinforcement learning, classical conditioning, imprinting, imitation. We describe this model of learning in the next section.

2. The learning model

In traditional neural networks the activation level of a post-synaptic unit (post-S unit) is a function of the activation arriving from its pre-synaptic units (pre-S units), where the activation arriving from each pre-S unit depends on two factors: (a) the current level of activation of the pre-S unit, and (b) the weight of the connection linking the pre-S unit to the post-S unit (Fig. 1). Learning occurs because the weights of the connections change as a result of the neural network's experiences.