

## Editorial

# Introduction to developmental robotics

LISA A. MEEDEN\*† and DOUGLAS S. BLANK‡§

†Computer Science, Swarthmore College, Swarthmore, PA 19081, USA

‡Computer Science, Bryn Mawr College, Bryn Mawr, PA 19010, USA

Developmental robotics is a broad, new discipline that lies at the intersections of psychology, biology, artificial intelligence (AI) and robotics. This new field was inspired by the fact that most complex and intelligent biological organisms (as opposed to artificial ones) undergo an extended period of development before reaching their adult form and adult abilities. This new rubric captures the essential features of many related, previous research agendas, including embodied cognition, evolutionary robotics and machine learning. Although developmental robotics combines many of these previous efforts, it also has fundamental differences that separate it in a number of interesting ways.

To appreciate these differences, it is useful to reflect on the history of robotics and AI. Since the inception of AI in the 1950s, its practitioners have been striving to create intelligent machines. There have been some notable successes in restricted domains, such as game playing. However, the vision of creating general-purpose, human-like intelligence has not yet been achieved. To date, there have been three primary approaches to trying to create intelligent robots: direct programming; supervised machine learning; and evolutionary adaptation (Weng *et al.* 2001).

In direct programming, a human engineer analyses the problem domain, determines a solution and then implements the solution in a program. Here the intelligence resides solely in the human engineer, the robot is merely acting out the pre-programmed commands. Robots created by direct programming tend to be brittle and fail in new situations not anticipated by the human engineer.

In supervised learning, a human engineer creates a series of training situations describing how the robot should respond to particular sensory inputs. The robot learns to mimic the training data and typically makes useful generalizations that apply to novel situations that were never seen during training. This is an improvement over direct programming in that the robot, rather than the human engineer, determines how to solve the problem, and the robot can go beyond what it was initially exposed to, leading to more robust behaviour. However,

---

\*Email: meeden@cs.swarthmore.edu

§Email: dblank@cs.brynmawr.edu

robots created by supervised learning can only be trained on a small set of tasks and have not been able to learn general-purpose strategies that would apply across a wide variety of tasks.

In evolutionary adaptation, a human engineer creates a fitness function for measuring a robot's success. Then an initial population of potential robot controllers is randomly created, and each one is scored based on the fitness function. High-scoring robot controllers are allowed to reproduce with one another to create a new generation of slightly better controllers. This process is repeated over many generations, and eventually good robot controllers emerge. This is an improvement over the supervised learning approach in that the fitness function can be much more abstract than the training data needed for supervised machine learning. Thus, robots created by evolutionary adaptation can produce more general-purpose strategies that are not necessarily tied to a specific task. Although the evolutionary approach has shown promise, the level of generality that has been achieved so far is limited.

For our discussion, we shall eliminate direct programming from consideration; it is still the most effective method for solving a particular task quickly, but seems unlikely ever to lead to open-ended, general-purpose behaviour. The other two methodologies employ incremental, constructivist techniques. That is, supervised learning and evolutionary adaptation both work by taking small steps—starting with a set of current abilities, and then building on them by extending them a little bit. This property is captured by the aphorism, 'you can only learn what you almost already know'. However, these models usually can only learn to take a few small steps from the starting point. The obvious solution to this problem is the developmental paradigm. Under a developmental process, a system can continually advance what it knows by placing itself in situations where it almost knows something new, and then learning it. Applied repeatedly, this process can lead to models capable of behaviour far different from their initial abilities.

How actually to implement such a developmental system is an open question. Having an implementation will have implications for at least the four areas mentioned, psychology, biology, AI and robotics. Thus, the field of developmental robotics has a wide range of practitioners using the tools of their respective disciplines. Although these fields bring their own perspectives, a common set of attributes have begun to emerge in developmental robotics that distinguish this approach from previous methods.

One key difference between the developmental approach to robotics and the three previous AI approaches is that the human engineer is much more removed from the system. A robot's body (including its actuators and sensors) is fundamentally different from the human body. Therefore, the kinds of categories and concepts a robot may develop through its interactions with the environment are likely to be quite different from our own. In developmental robotics, the human engineer creates a developmental architecture, one that can autonomously determine what to learn and how to learn it, and then allows the robot to construct its own representations of its body and its environment.

A number of papers in this special issue illustrate this fundamental difference in the developmental approach by demonstrating robots that create their own sensor and action models from scratch. Stronger and Stone, in their paper *Towards autonomous sensor and actuator model induction on a mobile robot*, describe a system that simultaneously creates sensor and action models by bootstrapping one from the other. Olsson, Nehaniv and Polani, in their paper *From unknown sensors and actuators to actions grounded in sensorimotor perceptions*, start with a basic assumption of continuity and, using information theory, are able to construct sensor and action models that allow the robot to move from random actions to visually guided movement.

Another model, described in two papers within this special issue, strives to learn common-sense knowledge starting from simple sensorimotor experiences. Kuipers, Beeson, Modayil and Provost, in their paper *Bootstrap learning of foundational representations*, describe the

overall structure of this model, known as the Spatial Semantic Hierarchy (SSH). SSH initially creates sensor and action models, then uses these models to discover high-level actions as well as high-level perceptual features, and finally uses these features to construct cognitive maps. Provost, Kuipers and Mikkulainen, in their paper *Developing navigation behaviour through self-organizing distinctive state abstraction*, describe in more detail how the high-level actions and features of the SSH provide a much more efficient abstraction for learning to navigate within an environment.

These four papers illustrate that sensor and action models need not be human-engineered, but can instead be grounded in (and emerge from) the robot's own sensorimotor experience. Kuipers, Beeson, Modayil and Provost argue that the ability of a robot to learn high-level concepts based on its own low-level experience is the foundation for having its own original intentionality and, eventually, consciousness.

A second key difference between the developmental approach and previous AI approaches is that, typically, no tasks are pre-defined and no external goals are provided. Instead, the robot is self-motivated to choose its own actions. Many researchers in developmental robotics have been exploring curiosity as an essential self-motivating influence.

Schmidhuber, in his paper *Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts*, provides a categorization of methods that have been used to implement curiosity in developmental systems. He argues that intelligent agents generate actions so as to create new and interesting patterns in their own input stream. Furthermore, he suggests that art and creativity can be explained as by-products of this intrinsic curiosity.

Oudeyer and Kaplan, in their paper *Discovering communication*, describe a self-motivation device that is based on optimizing learning progress. This method allows the robot to focus on what they call *progress niches*, situations that are neither too predictable nor too unpredictable. This scheme allows the robot 'to learn what it almost already knows', and thus provides a series of stepping stones towards more complex behaviour.

If these developmental systems have no pre-defined tasks or goals, how can one judge if the system is successful? What measurements can we use to gauge progress in the developing system? Because the human engineer may be far removed from the developmental system, such questions may, in the long run, be irrelevant. However, we currently need evidence to indicate whether we are on the right path in designing initial structures. One solution to these questions may be to look for patterns in the resulting behaviour of the developing system. Of course, all of the papers in this special issue explore models that change over time. Oudeyer and Kaplan argue that their Intelligent Adaptive Curiosity model exhibits qualitative behavioural changes indicating stages in a developmental sequence.

The next paper in this special issue is an example of a sub-area within developmental robotics, known as epigenetic robotics, that has the goal of creating robotic models of psychological development (Berthouze and Prince 2003). In his paper *Decomposing infant's object representations: A dual-route processing account*, Matthew Schlesinger explores a model that mimics human infant behaviour and is supported by evidence of the structure of the human visual system. The model implements complementary processes for identifying salient objects and locations (the 'what' and 'where' channels). This eye-movement model is instantiated with a real camera and demonstrates that prediction errors can drive top-down, deliberate analysis.

A model need not adopt all of these extreme developmental properties to exhibit advantages of the developmental paradigm. One key advantage is that a developing system can more readily adapt to the environmental conditions that it encounters during its lifetime. Developmental mechanisms found in biological organisms can be incorporated into the design of robots with the hope of creating better, more robust robots. In their paper *Learning acceptable windows of contingency*, Kevin Gold and Brian Scassellati explore the problem of how a robot can learn to recognize what aspects of its sensory experience result from self-generated actions.

They also examine the problem of detecting social responses that are contingent on the robot's actions. This is accomplished by learning a range of possible times over which the effect of an action can occur. This adaptive process allows the robot to reason more effectively about causal relationships between itself and others.

Developmental robotics can be seen as the next stage in a new, interdisciplinary developmental science. Like all developmental systems, it builds on previous knowledge, while seeking out novel, interesting areas to explore. It is fuelled by bottom-up data, and guided by top-down theory. We hope that this special issue contributes to both the fuelling and guiding of this emerging field.

## Reference

- L. Berthouze and C.G. Prince, "Introduction" in *Proceedings of the Third International Workshop on Epigenetic Robotics: Modelling Cognitive Development in Robotic Systems*, 2003.
- J. Weng, J. McClelland, A. Pentland, O. Sporns, I. Stockman, M. Sur and E. Thelen, "Autonomous mental development by robots and animals", *Science*, 291, pp. 599–600, January 2001.