A Developmental Robotics Manifesto

Douglas Blank, Bryn Mawr College (dblank@cs.brynmawr.edu)
James Marshall, Sarah Lawrence College (jmarshall@sarahlawrence.edu)
Lisa Meeden, Swarthmore College (meeden@cs.swarthmore.edu)

We are largely in agreement with Tani's approach to developmental robotics as elucidated in this dialog and his recent book. The basic assumptions inherent in his approach, such as that agents are embodied in the world and that neural systems are capable of complex learning, are now established wisdom. Although this has been a relatively recent shift in AI and Cognitive Science, we consider these underlying assumptions to be a given and thus do not address them further. Here we expand on Tani's questions and offer a broader set of principles for guiding developmental robotics research.

The learning system should be capable of taking advantage of the world's structure and continuity.

For any specific learning algorithm, there will always be some types of problems that cannot be easily learned. Thus one should choose a learning method that best matches the domain. Our world is a highly structured and largely continuous environment through time and space. Because of these regularities, we should use those learning systems that can exploit the gradients between similar situations. Gradient descent procedures, such as backpropagation, are well suited for domains of this type.

In addition, an embodied agent moving through the real world typically needs to execute a sequence of actions in order to achieve its current goal or plan. Recurrent neural networks are able to construct representations of sequential memory using gradient descent procedures. Thus, Tani's approach of using BPTT applied to a recurrent neural system clearly meets our guidelines. However, we do not want to commit to any particular architecture or technique, even though we use very similar implementations in our own research.

Intrinsic motivation and the ability to make predictions and abstractions should be innate.

Developmental robotic agents are confronted by a dynamically changing and immensely complex world, which is only partially predictable. An agent's sensory systems provide a ceaseless flood of multimodal information about the surrounding environment. Initially, a robotic agent has no understanding of the relationship between its sensors and the world, nor how its actions affect its sensors. Intrinsic motivation imbues the robot with curiosity and a desire to learn, which guides the robot in seeking out a low-level understanding of itself. At this stage, moment-to-moment sensory predictions about the outcomes of actions drive the robot to make low-level abstractions. Even though the system will never be able to perfectly predict the world, attempting to predict it will generate an error signal that gives the robot useful information. This information can be exploited in a variety of ways, for example to measure learning progress, to trigger attention, or to recognize sources of variability, which could include other agents.

Intelligence emerges through bottom-up and top-down interactions.

Once the robot has developed a bottom-up understanding of how its sensors and actions interact, and has created low-level abstractions based on this understanding, higher-order predictions and chunked abstractions can emerge. To be useful, *meaning* must be extracted from the sensory stream, in a continuous process that filters out enormous amounts of noisy, extraneous, redundant, or irrelevant information, depending on the situation at hand, and a coherent, abstract interpretation of the situation must be constructed.

This abstract interpretation can be bootstrapped from the knowledge gained at the lower levels. First, a higher level would learn to predict the sequence of states that occur at a lower level, leading to the development of its own higher-level abstractions. We note that these higher-level abstractions can be much more sophisticated than merely predicting one's own sensor readings. For example, a system could compare possible future action sequences in a type of counterfactual exploration. A higher level could then manipulate the sequence of states at the lower lever in order to achieve a chosen goal.

This process proceeds simultaneously on many levels of abstraction, and gradually, through development, becomes ever more efficient over the lifetime of the agent, as its knowledge of the world increases and it learns to better exploit that knowledge in pursuit of its goals.

We believe that continual, sustained, and interacting pressures will be necessary to create a system of lifelong learning. Over time, such an emergent, self-ratcheting system will have the potential to achieve robust levels of intelligent behavior in dynamic, unpredictable environments.

Conclusions

Returning to Tani's primary questions: we see predictive coding as essential; BPTT as a promising approach, though not essential; staged development as emergent; and foresee that by following the philosophy outlined above, robots could one day be conscious. We believe that Tani's work is a valuable contribution to better understanding the potential of developmental robotics, and, if combined with self-motivation and self-ratcheting pressures, is firmly in line with our manifesto.