Exploring Developmental Robotics via Predicted Progress

Leslie McTavish, Julia Ferraioli and George Dahl – Bryn Mawr College, Swarthmore College
Faculty Advisor Douglas Blank

Abstract: Developmental robotics involves designing learning systems for robots such that they are capable of displaying sophisticated behaviors that have not been directly programmed into them (Meeden and Blank, 2006). Traditionally, robots can only do what we program them to do. They are not capable of acquiring any new information and therefore are not able to adapt to new situations. The developmental approach to learning attempts to create a models that mimic human learning development. Using the Intelligent Adaptive Curiosity model, the robot learns how to effectively explore its environment and associate changes in its motor values with its sensor readings.

Intelligent Adaptive Curiosity (IAC):

Experiment: This experiment is a reenactment of the Playground Experiment conducted by Pierre-Yves Oudeyer, Frederic Kaplan, Verena V. Hafner and Andrew Whyte at the Sony Computer Science Lab, Paris. The goal is to see if a Robot can learn about itself and its environment using its sensory and motor values to predict the outcome of a motor sequence. The motivation system being used is called Intelligent Adaptive Curiosity (Oudeyer & Kaplan 2004) and it has been extended using a neural network and K-means to train and split the memory regions.

Methodology: The robot 'brain' is given a set of sensory and motor values which it uses to try to predict what the next sensory values will be. Each time the robot makes a move the system generates 3 sets of random candidates for the next motion. Each candidate contains values that indicate whether the robot with bite or bash, and the head and leg positions. Most of the time the brain will choose the candidate for which the maximal progress is expected and occasionally it will pick a random candidate. Each step is added to a Memory Expert which is a neural network. When the size of the Experts reaches a threshold, the region is split using K-means.

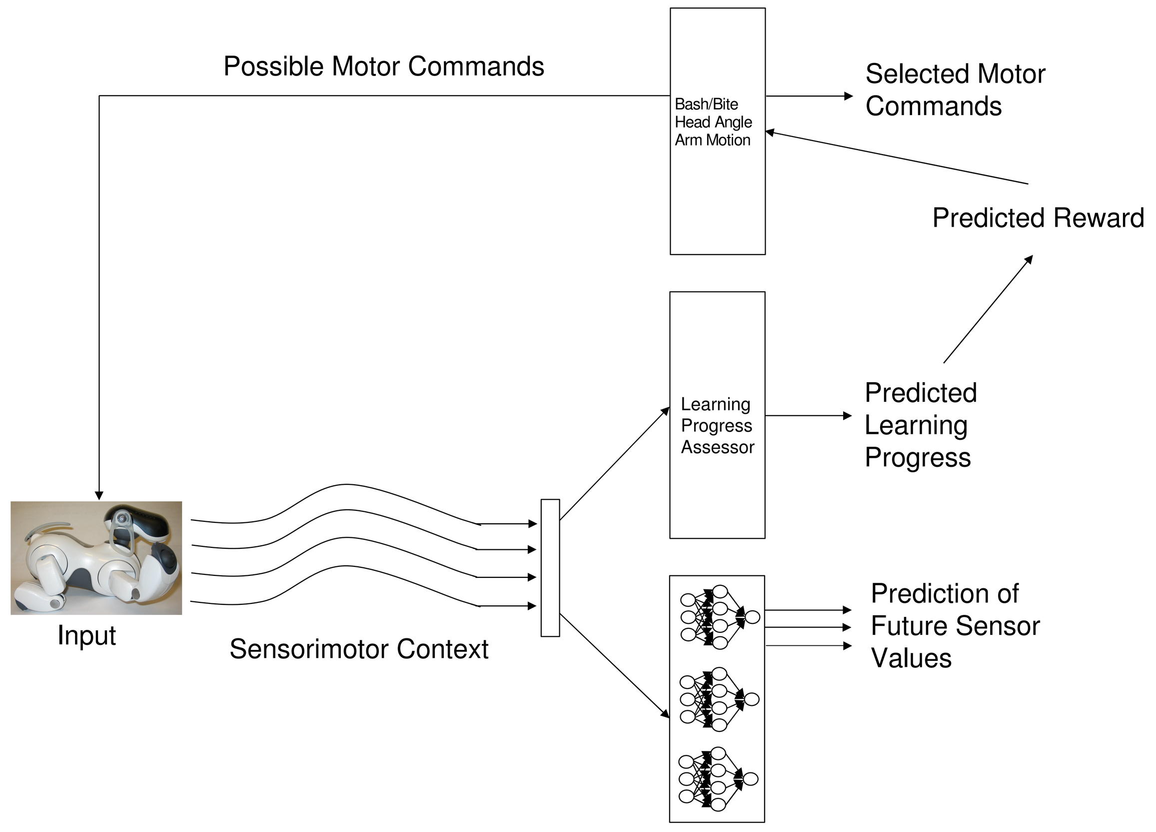

Environment: The robot is placed in a baby's playground area in which there is an elephant's ear which it can bite, and a suspended object that it can bash. It is also capable of making movements with its head each time it makes a biting or bashing motion and there are red target tags set up so that the robot can see a tag if it does a successful bite or bash.

Results: After 21,900 cycles, the IAC resulted in 339 regions, each containing neural network expert and a maximum of 50 exemplars. As expected, the results show that the number of successful bites is higher than successful bashes. This is because there are fewer random choices involved in a successful bite movement. The red lines indicate the random probability of success for each movement. It should be noted that the robot has not been told that it should bite or bash, it is only trying to make the action that best predicts maximal progress.

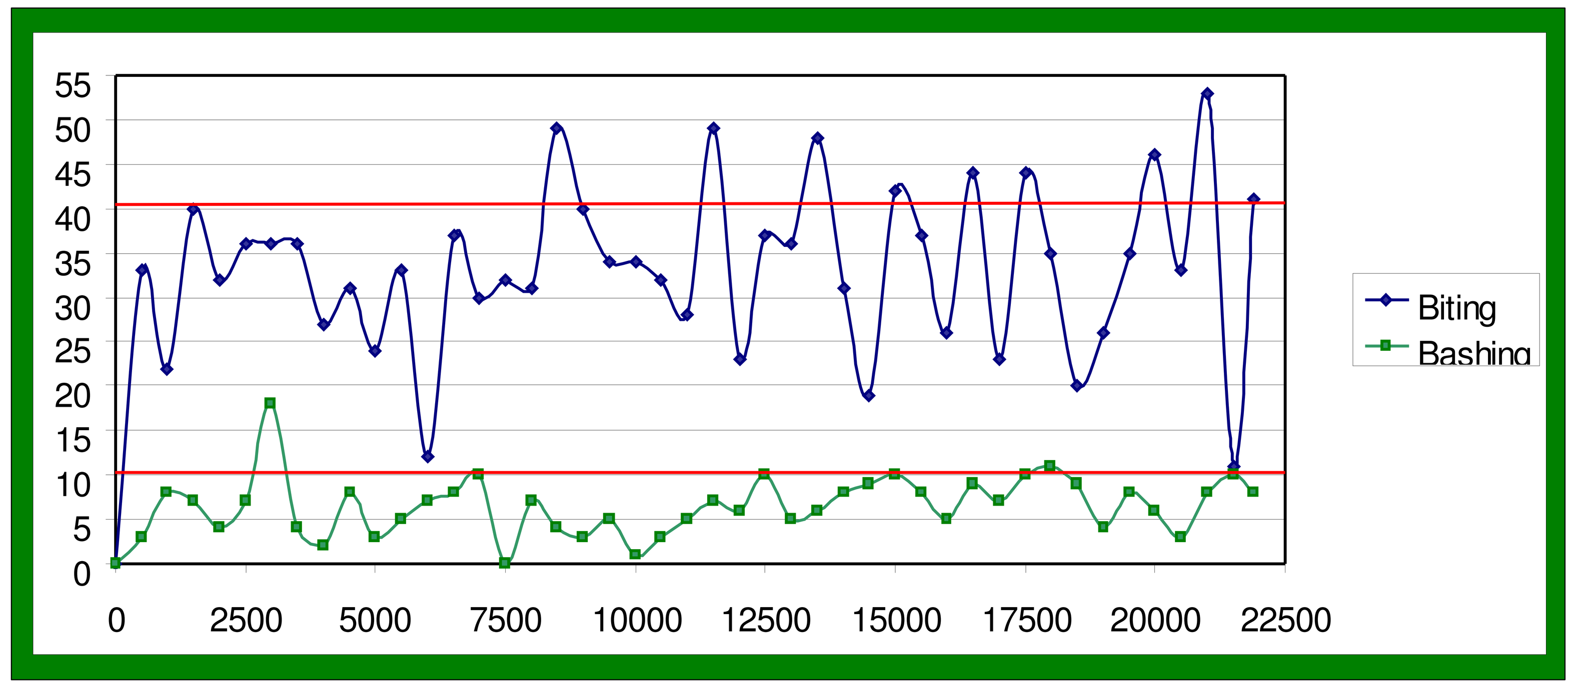

Number of Successful Actions per 500 cycles

Developmental Tools:

k-means: The k-means algorithm takes a data set and splits it up into *k* clusters based on the similarity of each data point to the data points around it. This clustering algorithm allows for generalizations to be made about possibly high dimensional data.

Neural Networks: Neural Networks, or Artificial Neural Networks (ANN), are non-linear statistical data modeling tools. They allow a computer to be able to generalize patterns from data, and to model the relationship between inputs and outputs. This model imitates the firing of neurons in the brain, so the output of the network depends on the values propagated from one neuron to another.

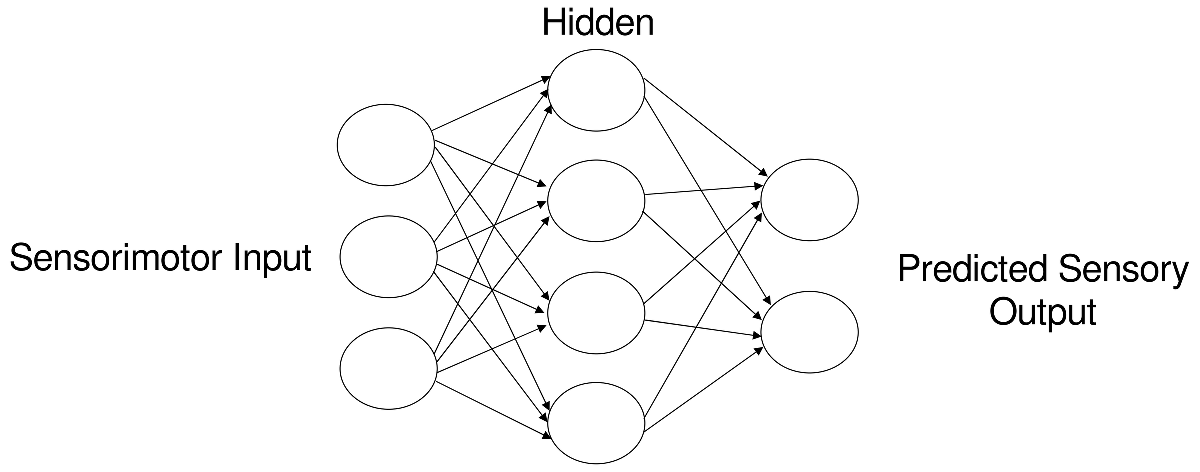

Future Work:

It would be interesting to try running this experiment with changes it the various parameters to try to increase the learning rate. Adjusting values such as the number of hidden layers, learning rate, and the size of the Memory Experts regions could have a dramatic effect on the results. We plan on exploring Cascade Correlation, an incremental algorithm (Fahlman and Lebiere, 1991).

References:

Fahman, S.E. and Lebiere, C. The Cascade-Correlation Learning Architecture. Carnegie Mellon Technical Report CMU-CS-90-100, 1991.

Meeden, L.A. and Blank, D. S.. Introduction to Developmental Robotics. In: Connection Science, Volume 18, Issue 2, June 2006.

Oudeyer, P., Kaplan, F., Hafner, V.V., and Whyte, A. The playground experiment: Task-independent development of a curious robot. In: Blank D., Meeden, L. (Eds.), Proceedings of the AAAI Spring Symposium on Developmental Robotics, 2005. Stanford, California, pp. 42-47.