Aggregate selection and direct competition in evolutionary robotics

Andrew L. Nelson[1] and Edward Grant[2]
[1]Androtics LLC, P.O. Box 44065, Tucson, AZ 85733, alnelson@ieee.org and [2]Center for Robotics and Intelligent Machines, North Carolina State University, Raleigh, NC 27695, egrant@ncsu.edu

Evolutionary robotics (ER) is a field of research that focuses on applying artificial evolution in order to generate learning in autonomous mobile robots. In a typical ER experiment, robot control systems are evolved to perform a particular task in a given environment. For example, robots might be evolved to locate objects (e.g. balls or small cubes) of one color in an environment containing walls and occluding obstacles of a different color. Robot controllers compete in the environment and are selected and propagated based on their ability (or fitness) to perform the desired task. A key component of this process is the manner in which the fitness of the evolving controllers is assessed.

In ER, fitness is measured by a *fitness function*. This function applies some criteria to determine which robots are better at performing the task at hand. Fitness functions can introduce *a priori* information or *features* related to how to perform the task. Populations of controllers evolved using such functions then reproduce these features and essentially evolve control systems that duplicate an *a priori* known algorithm. In contrast to this, it is sometimes possible to use an *aggregate* fitness function that incorporates no knowledge of how to achieve the desired task, but only measures final completion of the task. Such fitness functions are referred to as *aggregate* because they combine the benefit or deficit of all actions a given robot performs into a single success/failure term.

Our work has focused on combining two aspects of fitness assessment: aggregate fitness selection and direct competition among evolving agents (i.e. robot controllers). In some sense, all artificial evolution involves competition in that during the propagation phase, fitnesses of agents in the current population are compared, and the best (or fittest) are selected to create the next generation. However, the fitness of any one agent is absolute and does not usually directly alter the fitness of another. In our work we evolved robot control systems to perform a directly competitive searching task in which teams of robots played a simplified version of the game “capture the flag”. The fitness of each controller in the evolving population was defined only as a function of the number of games won in a tournament involving all members of the current generation. In this case performance of one controller would be directly affected by, and therefore *relative to*, that of the other agents in the population. We applied this form of *aggregate relative competitive fitness* evaluation to drive evolution of robot controllers in a coupled physical and simulation environment. Robot controllers in the form of artificial neural networks were evolved in simulation and then tested in real robots in a physical environment. Evolutionary runs on the order of 10^3 generations with a population size of 60 agents resulted in evolved controllers able to play competitively against a simple hand-coded rule-based controller.