

Autonomous Agents: Characterization and Requirements

Jose C. Brustoloni
November 1991
CMU-CS-91-204

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

A fresh assessment of what autonomous agents are or may aspire to be is presented in the paper. We dismiss the traditional dichotomy between reactive and symbolic architectures. A classification in terms of the amount of knowledge embedded in the system is made, instead, running from *regulatory* to *planning* to *adaptive* agents. Regulatory agents are implementable as *control systems* or *automata*, and represent the ideal case in which the agent has all the knowledge it needs. We find that each type of agent is best suited for a particular kind of behavior. The different kinds of behavior form a *hierarchy* in terms of frequency of occurrence and response time, and do not gracefully reduce to each other. This suggests implementing agents as hierarchies of simpler agents, each specialized in the kind of behavior it implements. In this scheme, agents at one level act through agents at the next lower level. We introduce the concept of *drive* to address the issue of where an agent's goals come from. A discussion is made of common pitfalls and suggestions for the design of planning agents. Finally, a *parallel* architecture that embodies many of the ideas presented is sketched and discussed. This architecture solves the two major problems present in Georgeff's earlier architecture, PRS: strictly sequential execution, and inability to learn.

This research was sponsored by the Office of Naval Research under Contract N00014-86-K-0678.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.