A Neural Network Pole Balancer that Learns and Operates on a Real Robot in Real Time

Dean Hougen hougen@cs.umn.edu
John Fischer jfischer@cs.umn.edu
Deva Johnam johnam@cs.umn.edu

Artificial Intelligence, Robotics, and Vision Laboratory
Department of Computer and Information Sciences
University of Minnesota
4-192 Electrical Engineering and Computer Science Building
200 Union St. SE, Minneapolis, MN 55455

Abstract

A neural network approach to the classic inverted pendulum task is presented. This task is the task of keeping a rigid pole, hinged to a cart and free to fall in a plane, in a roughly vertical orientation by moving the cart horizontally in the plane while keeping the cart within some maximum distance of its starting position. This task constitutes a difficult control problem if the parameters of the cart-pole system are not known precisely or are variable. It also forms the basis of an even more complex control-learning problem if the controller must learn the proper actions for successfully balancing the pole given only the current state of the system and a failure signal when the pole angle from the vertical becomes too great or the cart exceeds one of the boundaries placed on its position.

The approach presented is demonstrated to be effective for the real-time control of a small, self-contained mini-robot, specially outfitted for the task. Origins and details of the learning scheme, specifics of the mini-robot hardware, and results of actual learning trials are presented.