Neurocontrol and Neurobiology: New Developments and Connections

Paul J. Werbos, Room 1151
National Science Foundation*
Washington D.C. 20550

and

Andras J. Pellionisz
NASA Ames Research Center* 261-3
Moffett Field, CA 94035

ABSTRACT

The past two years have seen major progress in neurocontrol, particularly in the use of complex reinforcement learning schemes which appear ever more relevant to understanding neurobiology. For example, at McDonnell-Douglas, controllers which combine adaptive critic networks (i.e., approximate dynamic programming) with the use of backpropagation in real time have solved difficult control problems -- resistant to classical methods and simpler ANNs -- crucial to the feasibility of building an airplane (the National Aerospace Plane, NASP) able to reach earth orbit. Such developments led to a joint NSF-McDonnell workshop in October 1990, and to a new book [1] which provides extensive implementation details. As these details emerged -- particularly in relation to planning, chunking and real-time adaptation of time-lagged recurrent networks -- parallels to neurobiology have grown stronger, and have begun to lead to empirical possibilities of importance to neuroscience. This has led to thoughts of NSF-NASA-NIH-NIMH-(?) collaboration, in facilitating what could become a Newtonian revolution in neuroscience, with cognitive implications as well. This paper will elaborate on each of these points in turn. Because of page limits, it will summarize important conclusions, and leave it to the citations to provide more details and the reasons behind the conclusions.

RECENT PROGRESS IN NEUROCONTROL IN GENERAL

Artificial Neural Networks (ANNs) have performed four types of tasks in control applications. First, they have been used in subordinate roles -- for sensor fusion, pattern recognition, etc. -- in systems where the control signals were not generated by an ANN. Such applications have been very common and useful, but they do not meet the definition of neurocontrol. Second, ANNs have been used to "clone" human experts (or to clone automatic controllers too slow to use in real time). People have often trained ANNs to reproduce the actions of a human, taking as input the current sensor data available to the human; however, because good human controllers (like good automatic controllers) are highly sensitive to dynamics, it works better to treat these applications as an exercise in dynamic modeling or system identification or emulation of the human expert. Third, ANNs have been used to make a plant follow a desired trajectory, or stay at a desired set-point, or follow a desired reference model. This can be done in a direct way ("direct inverse control") or in an indirect way. The direct way fits the biologists' notion of learning the mapping from spatial coordinates (in which the desired path is encoded) to motor coordinates (actions required to meet that path), but it works better if dynamics are accounted for. In the indirect method, one defines an error function or disutility function which combines some measure of tracking error, jerkiness of movement, etc., and one then uses an optimization method to minimize that error over time. Narendra says that the latter works better and he has a stability proof for it[1]. Finally, there are two classes of neurocontrol designs to do optimization over time -- the backpropagation of utility (U), and the adaptive critic family of designs.

All four types of application have seen great advances in the past few years. Many tricks

*This paper gives personal views, not the views of NSF or NASA. As government work, however, it is in the public domain. Many thanks are due to C.Lau and A.Penz for making this possible.