assumption of tabular representations of the value function, for which generalization is impossible. Moreover, several researchers have argued that there are serious hazards intrinsic to the approximation of value functions by reinforcement learning methods. It is an important and still unclear question as to just how serious these hazards might be. In this workshop we will survey the substantial recent and ongoing work pertaining to this question. We will seek to answer it or, failing that, to identify the remaining outstanding issues.

If existing theory of reinforcement learning no longer applies when an approximate value function is learned, what are we to do? The workshop will explore the problem and the following possible responses:

1. Ignore the problem and proceed empirically, perhaps discovering as we create applications that some seem to work and some don't. A danger with this approach is that successes will be reported loudly and failures whispered quietly, and we may become a community of tweakers.
2. Analyze what properties function approximators need in order to work well with curent reinforcement learning algorithms: are there function approximators specially suited to reinforcement learning?
3. Invent new reinforcement learning methods that are specifically designed to work well with function approximators.
4. Invent new theory for value function approximation. What bounds can be placed on the errors? Can particular function approximators give us better bounds or stability guarantees? Can online training get us better bounds or stability guarantees?
5. What can we learn from the literature of other fields:
   - optimal control theory and LQG regulation
   - heuristic function generation techniques in computer game playing
   - operations research and differential games

WORKSHOP FORMAT

Several focussed sesssions with short (≈15 minute) talks followed by moderated discussion.

WHO SHOULD ATTEND

All researchers with empirical or theoretical experience with value function approximation in reinforcement learning or dynamic programming. As an expression of interest, please submit as soon as possible a short (one paragraph to one page) statement of your research interests in value function approximation, or a copy of a paper you have written in this area. We will use this material to help select and organize the sessions.

If you would like to make a presentation at the workshop, please also submit an extended abstract (2-5 pages) by May 1st, 1995.

SUBMISSION INFORMATION

Electronic submissions (ASCII or postscript) are preferred. E-mail submissions should be sent to:

`awm@cs.cmu.edu`

Post submission should be sent to:

Andrew W. Moore
Robotics Institute, Smith Hall 221
Carnegie Mellon University
Pittsburgh, PA 15213

IMPORTANT DATES

Submit statements of interest	ASAP
Submit extended abstract	May 1
Acceptance notifcation	May 19
Submit final camera-ready papers	June 9
Workshop	July 9

■ **WORKSHOP ANNOUNCEMENT** ■

VALUE FUNCTION APPROXIMATION IN REINFORCEMENT LEARNING

Sunday July 9th, 1995

workshop held in conjunction with
The Twelfth International Conference on Machine Learning
Tahoe City, California
July 9-12, 1995

Organizers:

Andrew Moore (chair), Carnegie Mellon
Rich Sutton, Stow Research
Justin Boyan, Carnegie Mellon

Program Committee

John Tsitsiklis, MIT LIDS
Satinder Singh, MIT Brain & Cognitive Sciences
Leemon Baird, Wright Patterson AFB

Web home page: `http://www.cs.cmu.edu:8001/Web/Groups/reinforcement/ml95/`

Overview

This workshop will explore the issues that arise in reinforcement learning when the value function cannot be learned exactly, but must be approximated. It has long been recognized that approximation is essential on large, real-world problems because the state space is too large to permit table-lookup approaches. In addition, we need to generalize from past experiences to future ones, which inevitably involves making approximations. In principle, all methods for learning from examples are relevant here, but in practice only a few have been tried, and fewer still have been effective. The objective of this workshop is to bring together all the strands of reinforcement learning research that bear directly on the issue of value function approximation in reinforcement learning. We hope to survey what works and what doesn't, and achieve a better understanding of what makes value function approximation special as a learning from examples problem.

Issues

The key computational idea underlying reinforcement learning is the iterative approximation of the value function---the mapping from states (or state-action pairs) to an estimate of the long-term future reward obtainable from that state. For large problems, the approximation of the value function must involve generalization from examples to reduce memory requirements and training time. Moreover, generalization may help beat the curse of dimensionality for problems in which the complexity of the value function increases sub-exponentially with the number of state variables. Generalizing function approximators have been used effectively in reinforcement learning as far back as Samuel's checker player, which used a linear approximator, and Michie and Chambers' BOXES system, which used state aggregation. Tesauro's TD-Gammon, which used a backpropagation network, provides a tantalizing recent demonstration of just how effective this can be.

However, almost all of the theory of reinforcement learning is dependent on the