Preface

Complex artificial systems have become integral and critical components of modern society. The unprecedented rate at which computers, networks, and other advanced technologies are being developed ensures that our dependence on such systems will continue to increase. Examples of such systems include computer and communication networks, transportation networks, banking and finance systems, electric power grid, oil and gas pipelines, manufacturing systems, and systems for national defense. These are usually multi-scale, multi-component, distributed, dynamic systems. While advances in science and engineering have enabled us to design and build complex systems, comprehensive understanding of how to control and optimize them is clearly lacking.

There is an enormous literature describing specific methods for controlling specific complex systems or instruments based on various simplifying assumptions and requiring a range of performance compromises. While much has been said about complex systems, these systems are usually too complex for the conventional mathematical methodologies that have proven to be successful in designing complex instruments. The mere existence of complex systems does not necessarily mean that they are operating under the most desirable conditions with enough robustness to withstand the kinds of disturbances that inevitably arise. This was made clear, for example, by the major power outage across dozens of cities in the Eastern United States and Canada in August of 2003.

Dynamic programming is a well known, general-purpose method to deal with complex systems, to find optimal control strategies for nonlinear and stochastic dynamic systems. It is based on the Bellman equation which suffers from a severe "curse of dimensionality" (for some problems, there can even be *three* curses of dimensionality). This has limited its applications to very small problems. The same may be said of the classical "min-max" algorithms for zero-sum games, which are closely related. Over the past two decades, substantial progress has been made through efforts in multiple disciplines such as adaptive/optimal/robust control, machine learning, neural networks, economics, and operations research. For the most part, these efforts have not been cohesively linked, with multiple parallel efforts sometimes being pursued without knowledge of what others have done. A major goal of the 2002 NSF workshop was to bring these parallel communities together to discuss progress and to share ideas. Through this process, we are hoping to better define a community with common interests and to help develop a common vocabulary to facilitate communication.

Despite the diversity in the tools and languages used, a common focus of these researchers has been to develop methods capable of finding high-quality approximate solutions to problems whose exact solutions via classical dynamic programming are not attainable in practice due to high computational complexity and lack of accurate knowledge of system dynamics. At the workshop, the phrase *approximate dynamic programming* (ADP) was identified to represent this stream of activities.

A number of important results were reported at the workshop, suggesting that these new approaches based on approximating dynamic programming can indeed scale up to the needs of large-scale problems that are important for our society. However, to translate these results into systems for the management and control of real-world complex systems will require substantial multi-disciplinary research directed toward integrating higher-level modules, extending multi-agent, hierarchical, and hybrid systems concepts. There is a lot left to be done!

This book is a summary of the results presented at the workshop, and is organized with several objectives in mind. First, it introduces the common theme of ADP to a large, interdisciplinary research community to raise awareness and to inspire more research results. Second, it provides readers with detailed coverage of some existing ADP approaches, both analytically and empirically, which may serve as a baseline to develop further results. Third, it demonstrates the successes that ADP methods have already achieved in furthering our ability to manage and optimize complex systems.

The organization of the book is as follows. It starts with a strategic overview and future directions of the important field of ADP. The remainder contains three parts. Part One aims at providing readers a clear introduction of some existing ADP frameworks and details on how to implement such systems. Part Two presents important and advanced research results that are currently under development and that may lead to important discoveries in the future. Part Three is dedicated to applications of various ADP techniques. These applications demonstrate how ADP can be applied to large and realistic problems arising from many different fields, and they provide insights for guiding future applications.

Additional information about the 2002 NSF workshop can be found at
http://www.eas.asu.edu/~nsfadp

JENNIE SI, TEMPE, AZ

ANDREW G. BARTO, AMHERST, MA

WARREN B. POWELL, PRINCETON, NJ

DONALD C. WUNSCH, ROLLA, MO