Learning to control advanced life support systems: final report

Devika Subramanian
Rice University

Decmber 2004

1 Goals and Achievements

Advanced life support systems have many interacting processes and limited resources. Controlling and optimizing advanced life support systems presents unique challenges. In particular, advanced life support systems are nonlinear coupled dynamical systems and it is difficult for humans to take all interactions into account to design an effective control strategy. In this project, we developed several reinforcement learning controllers that actively explore the space of possible control strategies, guided by rewards from a user specified long term objective function. We evaluated these controllers using a discrete event simulation of an advanced life support system. This simulation, called BioSim, designed by Nasa scientists David Kortenkamp and Scott Bell has multiple, interacting life support modules including crew, food production, air revitalization, water recovery, solid waste incineration and power. They are implemented in a consumer/producer relationship in which certain modules produce resources that are consumed by other modules. Stores hold resources between modules. Control of this simulation is via adjusting flows of resources between modules and into/out of stores. We developed adaptive algorithms that control the flow of resources in BioSim. Our learning algorithms discovered several ingenious strategies for maximizing mission length by controlling the air and water recycling systems as well as crop planting schedules. By exploiting non-linearities in the overall system dynamics, the learned controllers easily outperformed controllers written by human experts.

In sum, we accomplished three goals. We (1) developed foundations for learning models of coupled dynamical systems by active exploration of the state space, (2) developed and tested algorithms that learn to efficiently control air and water recycling processes as well as crop scheduling in Biosim, and (3) developed an understanding of the role machine learning in designing control systems for advanced life support.

2 Interaction

We interacted closely with Nasa scientists during the period of this grant visiting them on a regular basis to discuss aspects of the Biosim simulation as well as the design of the reinforcement learning algorithms. Our work on reinforcement learning was vital in uncovering crucial simulator bugs during January to June 2004; our reinforcement learner exploited loopholes in the simulator design to extend the mission length indefinitely! The exploration of the state space by the reinforcement learning allowed us to build models of system-level interactions in the Biosim.

3 Papers and Theses

- Learning to control coupled dynamical systems, D. Kortenkamp, P. Bonasso and D. Subramanian, Rice University Technical Report, 2003.
- Using reinforcement learning to control life support systems, T. J. Klein, D. Subramanian, D. Kortenkamp and S. Bell, ICES 2004.
- Using reinforcement learning to control advanced life support systems, T. J. Klein, Master thesis, Department of Electrical Engineering, Rice University, December 2004.

4 Training

Theresa Klein was trained on this grant; she obtained her Masters degree in Electrical Engineering as a result of the research she performed under the auspices of this grant. Theresa is currently pursuing a Ph.D. in Electrical Engineering at the University of Arizona. Her final career goal is to work for Nasa missions, designing robots and life support systems for interplanetary exploration. Theresa's thesis work will be reported in journal length submissions that are in progress.

5 Acknowledgement

We thank NASA for the support of this project. It allowed us to explore new directions in machine learning, and it helped train a new scientist in the design of autonomous control systems for complex dynamical systems.