

The Roles of Machine Learning in Robust Autonomous Systems

David Kortenkamp

NASA Johnson Space Center – ER2
Houston, TX 77058
kortenkamp@jsc.nasa.gov

Extended Abstract

Robust autonomous systems will need to be adaptable to changes in the environment and changes in the underlying physical system. This is especially critical for long-duration missions. For example, autonomous robots that explore other planets for years will need to adapt to degradation in their capabilities and to unforeseen environmental factors. Another NASA domain requiring robust, adaptable autonomy is control of closed-loop systems that will provide life support to crews on long-duration missions. We have been investigating autonomous control of advanced life support systems for many years (Schreckenghost *et al.* 1998; Kortenkamp, Keirn-Schreckenghost, & Bonasso 2000). Recently we have begun investigating learning with respect to advanced life support systems (Kortenkamp, Bonasso, & Subramanian 2001). In this abstract I briefly discuss some of the roles machine learning can play with respect to control of advanced life support systems or any complex, real-time system.

Detecting signatures

Certain sensor “signatures” require specific responses from the autonomous control system. These signatures are often hand-coded by the programmer. For example, the programmer might state that if the temperature is above 100 and the pressure is above 1000 then vent the tank. This is a trivial example and real-world examples will be more complex and involved. This means that hand-coded signatures may not accurately capture the event, especially if the environment or the physical plant are changing. A variety of machine learning techniques could be used to look at the history of the system and adjust the signatures automatically. A specific research challenge in this area is:

- Can we learn to parse real-world data streams into recognizable system modes or events?

Optimizing control/resource usage

Long-duration missions, both robotic and crewed, face severe resource constraints. Robots are often con-

strained in the amount of energy (from batteries or solar panels) that they have. Crewed missions are often constrained by life support consumables such as oxygen, water and food. Given adequate simulations, machine learning techniques such as reinforcement learning or genetic algorithms can search through combinations of control actions to discover a control policy that optimizes usage of a particular resource (or combination of resources). The resulting policy can be implemented by the autonomous control system to increase mission duration. We have conducted experiments using genetic algorithms and reinforcement learning to optimize resource utilization in a simulated advanced life support system (see (Kortenkamp, Bonasso, & Subramanian 2001) for details).

Refining models

Many autonomous controllers contain a model of the system they are trying to control, either explicitly or implicitly. Typically, these models are hand-coded and do not change. However, in long-duration missions the underlying physical system may change dramatically due to damage or degradation. For example, a robot’s solar panels may accumulate dust that reduces their power output. Or, a robot’s wheels may slip more than anticipated. Or a filter in a life support system may clog more frequently than expected. In each of these cases, machine learning techniques could be used to modify the internal model of the system based on real-world data. This is especially necessary if other machine learning techniques are using these models for optimization (as discussed in the previous section). A specific research challenge in this area is:

- Can feedback from actual operation of the system be used to automatically refine our simulations and models?

Learning/optimizing sequences

The behavior of most autonomous systems can be characterized as sequence of actions that lead to some desired result. These sequences can be generated by planners, be hand-coded by programmers, or “emerge” from the interactions of independent behaviors. In any case,

sequences are at the heart of a robust autonomous system. Machine learning techniques can be used to optimize sequences, especially in the case of sequences that are generated by planners or are hand-coded. Optimization of sequences can occur either by experimenting with a simulation or by looking at data from previous executions of sequences in the real world. Some research challenges in this area are:

- How can we provide safety guarantees when transferring a controller learned from simulations into the real system?
- How can the system learn not only sequences but also the contexts in which the sequences apply? Or perhaps the system starts with a standard set of sequences and learns when and where each sequence applies.
- What are the tradeoffs between having the system learn sequences from scratch as opposed to “tweaking” existing, working sequences?

Integration with autonomous control architectures

A large complex system will have existing control procedures and possibly even an overarching control architecture. Very few current control architectures incorporate learning or adaptation. Integration of learning into an existing (or proposed) autonomous control system raises a number of important research issues, such as:

- How good does the initial set of control behaviors or rules need to be for learning to be effective? Can we start *tabula rasa* or do we need very effective initial strategies?
- What are the differences between learning control information, learning procedural information, learning qualitative modeling information and learning planning information?
- What are the criteria whereby the autonomous control system turns learning on or off?
- How does the autonomous control system decide when to use new learned actions?

Control system design methodology

The solution space for a closed loop control policy in any complex interacting system is enormous. We believe that machine learning techniques are a useful way to “probe” the solution space and give control system designers an idea as to its topology and the important control variables. In this way, machine learning algorithms become not just a tool for adjusting the on-line control system, but also a tool for helping programmers design an *a priori* control policy. There are some important research questions that need to be answered with respect to using machine learning techniques as “probes” of the solution space, including:

- For what class of dynamical systems is open loop analysis (e.g., genetic algorithms) helpful for guiding state space design for closed loop policy search?
- How do we design reinforcement learning algorithms that can automatically design the right state space features?
- Which machine learning techniques (e.g., reinforcement learning, memory-based learning, genetic algorithms, etc.) are the most useful for probing the solution space?
- How can the discoveries of machine learning algorithms be presented to control system designers so that they can easily understand the topology of the search space and the important control variables and interactions?
- How detailed must the simulation be for results from a machine learning algorithm that probes the solution space to be applicable to the real physical system?

Conclusions

The use of machine learning to increase the robustness of autonomous control of long-duration missions poses some specific challenges to the machine learning community. In particular,

- Dealing with non-stationary dynamics
- Global optimization of large numbers of parameters
- Issues of safety and performance guarantees (verification and validation of machine learning techniques)
- Inspectibility (so crew members know why certain actions have been chosen)
- Long-term (or life-long) learning
- Coping with drastic changes in the underlying environment or physical system
- Integration of machine learning techniques with traditional AI techniques like planners, sequencers, model-based reasoners

Acknowledgements

The ideas in this paper grew out of conversations with the Advanced Life Support (ALS) Machine Learning Group, which consists of R. Peter Bonasso (NASA JSC/Metrica), Justin Boyan (NASA Ames), Dan Clancy (NASA Ames), Greg Dorais (NASA Ames), Land Fleming (NASA JSC/Lockheed), Matt MacMahon (NASA JSC/SKE), Jane Malin (NASA JSC), Jeff Schneider (Carnegie Mellon University), Debra Schreckenghost (NASA JSC/Metrica), Jay Schroeder (NASA JSC), Alan Schultz (Naval Research Laboratory), Richard Simpson (NASA JSC/Metrica) and Devika Subramanian (Rice University). This work is supported by a NASA CETDP Thinking Space Systems grant entitled “Evaluating the Application of Machine Learning to the Control of Advanced Life Support Systems.”

References

- Kortenkamp, D.; Bonasso, R. P.; and Subramanian, D. 2001. Distributed autonomous control of space habitats. In *Proceedings of the IEEE Aerospace Conference*.
- Kortenkamp, D.; Keirn-Schreckenghost, D.; and Bonasso, R. P. 2000. Real-time autonomous control of space habitats. In *AAAI 2000 Spring Symposium on Real-Time Autonomous Control*.
- Schreckenghost, D.; Ryan, D.; Thronesbery, C.; Bonasso, R. P.; and Poirot, D. 1998. Intelligent control of life support systems for space habitats. In *Proceedings of the Conference on Innovative Applications of Artificial Intelligence*.