

# Autonomic Computing: A New Challenge for Machine Learning Combined Tutorial and Workshop

Friday September 22, 2006

Irina Rish, *IBM T.J. Watson Research Center*  
Gerry Tesauro, *IBM T.J. Watson research Center*  
Rajarshi Das, *IBM T.J. Watson research Center*  
Jeff Kephart, *IBM T.J. Watson research Center*

In this combined tutorial and workshop, we will explore the many new challenges for machine learning that arise in the context of the emerging field of Autonomic Computing. Large-scale distributed computing systems are evolving into multi-vendor tangles of heterogeneous components, each of which may have hundreds of configuration and tuning parameters, and they are rapidly becoming too difficult for humans to configure, tune and maintain. The goal of Autonomic Computing is to avert this looming complexity crisis by giving computing systems and their components the ability to manage themselves in accordance with high-level objectives from administrators.

Machine learning appears to be a promising approach to the above challenges which could replace laborious, time consuming, and often suboptimal "hand-crafted" rules and models that are typically used for management tasks in state-of-art systems, by adaptive online learning and decision-making methods, such as, for example, reinforcement learning. Another hope is that adaptive online decision-making via ML approaches can cope better with the complex non-stationary dynamic system behavior than the explicit control and queuing models used in the most advanced systems management today. Moreover, manual selection of relevant metrics and measurements that is essential for effective management could be replaced by dimensionality reduction and active learning approaches that would help to "squeeze out" at minimal cost the most-relevant systems applications, it does not necessarily mean that all the data being collected are informative with respect to the question of interest or are relevant for decision needs to be made. Indeed, "we are drowning in data but starving for knowledge".

The good news about systems management domain is that it is naturally suited for active learning and exploration vs exploitation in reinforcement learning, since various measurements and tests can be constructed and performed on demand - the feature not always present in more "natural" domains such as biology or medicine where tests are more costly and constrained. In the tutorial, we plan to discuss several systems management applications, including event mining, network diagnosis, end-user transaction recognition, performance prediction, server selection in content-distribution systems, and online resource allocation, as well as examples of machine-learning approaches that were successfully applied to the above problems, e.g., probabilistic inference, classification, reinforcement learning, active learning, and collaborative prediction.

In the workshop following the tutorial we hope to bring together academic and industrial researchers to share their experience in applying machine learning to complex distributed systems management, identify new challenges to machine-learning community, match mature technologies to current problems, and chart the trajectory of inter-disciplinary research techniques that can be applied in Autonomic Computing.

## Workshop/Tutorial Schedule

### Morning session: 10:40–12:10

10:40am    **Tutorial: Part 1**, *Irina Rish and Gerry Tesauro*

12:10        *lunch*

### Afternoon session: 13:30–18:00

13:30        **Tutorial: Part 2**, *Irina Rish and Gerry Tesauro*

14:30        **Reliable, Adaptive Distributed Systems: RADical New Challenges For Machine Learning**,  
*Armando Fox (invited talk)*

15:10        *coffee break*

15:30        **Inferring Network Structure from Co-Occurrences**, *Michael Rabbat (invited talk)*

16:10        **Statistical Software Debugging**, *Alice Zheng (invited talk)*

16:50        *10 min break*

17:00        **Resource Access Pattern Mining for Dynamic Energy Management**, *Dinesh Rajan*

17:20        **A New Distributed Data Mining System on Grid**, *Huaiguo Fu*

17:40        *Discussion*