Reliable, Adaptive Distributed Systems: RADical New Challenges For Machine Learning

Armando Fox
University of California, Berkeley

The video hosting site YouTube.com became the 15th most popular Internet site within 18 months of its founding. Complex services such as Google operate multiple datacenters comprising hundreds of thousands of individual machines. Yet consumers have come to expect near-100availability. No wonder up to 85typical corporate IT budget is dedicated to just keeping things working, according to the former president of Oracle Online, hence the recent interest in applying existing SML techniques to systems analysis.

I will argue that systems is a promising application area both because it offers new research challenges to SML, and also because the systems community's extensive domain-specific knowledge offers promising collaborative approaches to these challenges. System experts' knowledge can inform issues such as parameter drift and structural drift of system models. "Ground truth" in systems is notoriously elusive and simulation is not considered sufficiently reliable to establish it, yet good systems intuition combined with SML techniques can sometimes provide useful proxies for correctly-labeled datasets. And while the end goal of autonomic computing is to obviate the need for human operators, the systems community still recognizes operator knowledge as an important information source that should be incorporated into system models.

There is a tremendous opportunity for systems building and SML to "meet in the middle" and tackle these challenges to enable the next century of IT innovation. I will discuss some of these problems in the context of the new Berkeley RAD Lab, whose ambitious five- year mission statement to enable a single individual to develop, analyze, deploy and operate the next-generation IT service at scales comparable to Google's today.

Speaker's Bio:
Armando Fox (fox@cs.berkeley.edu) recently joined UC Berkeley as a co-founder of the Berkeley RAD Lab. Prior to that he was an Assistant Professor of Computer Science at Stanford. His recent collaboration with David Patterson on Recovery-Oriented Computing earned him the distinction of being included in the "Scientific American 50" of 2003; he is also the recipient of an NSF CAREER award and teaching awards from Stanford University, the Society of Women Engineers, and Tau Beta Pi. In previous lives he helped design the Intel Pentium Pro microprocessor and founded a small company to commercialize his UC Berkeley dissertation research on mobile computing. He received his other degrees in electrical engineering and computer science from MIT and the University of Illinois.