Reinforcement Learning for Control

Open Invited Track at IFAC World Congress 2017

Lucian Buşoniu * **Robert Babuška** ** **Jagannathan Sarangapani** ***

* *Technical University of Cluj-Napoca, Romania (lucian@busoniu.net)*
** *Delft University of Technology, the Netherlands (r.babuska@tudelft.nl)*
*** *Missouri University of Science and Technology, USA (sarangap@mst.edu)*

Abstract: Reinforcement learning (RL) offers a principled way to control nonlinear stochastic systems with partly or even fully unknown dynamics. Recent advances in areas such as deep learning and adaptive dynamic programming (ADP) have led to significant inroads in applications from robotics, automotive systems, smart grids, game playing, traffic control, etc. This open track provides a forum of interaction and an outlet for novel contributions at the cutting edge of RL for control, addressing open issues such as safety and performance guarantees, computational complexity, large-scale systems, multiagent or partially observable problems, etc. We are particularly interested in the – so far largely unexplored – interactions between artificial-intelligence and control-theoretic approaches to RL; deep learning and ADP are two typical such areas that have stayed mostly separate on the two respective sides. We seek contributions on methods and analysis of RL for control, as well as on its applications in engineering, artificial intelligence, operations research, economics, medicine, and other relevant fields. We equally welcome surveys by established researchers in the field.

Keywords: Reinforcement learning control, adaptive dynamic programming, deep learning, performance and safety guarantees, Markov decision processes.

Technical Committee: **TC3.2 - Computational Intelligence in Control**.

Most systems in practical control applications are partly unknown, often to such an extent that fully model-based design cannot achieve satisfactory results. Reinforcement learning (RL) offers a principled way to solve such problems, in cases where a cumulative performance index must be optimized. RL methods are enjoying great popularity due to their ability to deal with general and complex systems, which in addition to having partly or fully unknown dynamics, may also be highly nonlinear and stochastic. Recently, deep learning methods have found a fertile intersection with RL, with well-known success stories like AlphaGo. In control, a name often used for RL methods is adaptive dynamic programming (ADP), where the focus is placed on exploiting the known structure of the model and ensuring stability guarantees. RL and ADP techniques have made great inroads in application areas like robotics, automotive systems, smart grids, game playing, resource management, and traffic control.

However, many issues remain open, including strong safety and performance guarantees, computational complexity, scalability to large-dimension state and action variables, applications to multiagent, distributed or partially observable problems, etc. The Open Track on "Reinforcement Learning for Control" provides an outlet for novel contributions in these and other open areas, as well as a forum for interaction between researchers and practitioners in the field. We are particularly interested in the – so far largely unexplored – interactions between artificial-intelligence and control-theoretic approaches to RL; deep learning and ADP are two typical such areas that have stayed mostly separate on the two respective sides. Synergy between artificial intelligence and control theory in RL can lead to major breakthroughs such as computationally efficient algorithms with strong, simultaneous performance and safety (stability) guarantees.

We equally welcome contributions from control theory, computational intelligence, computer science, operations research, neuroscience, and other novel perspectives on RL. We host original papers on methods, analysis, and applications of RL for control, as well as surveys from established researchers in the field. We are interested in applications from engineering, artificial intelligence, operations research, economics, medicine, and other fields. Topics of interest include, but are not limited to:

- RL and ADP-based control
- Performance, stability, and complexity analysis
- Deep reinforcement learning
- Function approximation and feature discovery
- Multiagent and distributed RL
- Model-based and model-learning techniques
- Event-driven RL and ADP
- Policy gradient and actor-critic methods
- Partially observable Markov decision processes
- Tree search, planning, and receding-horizon methods
- Bayesian RL and exploration
- Applications of RL and ADP
- Novel perspectives, e.g. biologically inspired methods

For additional information and updates, see the website of the open track at `http://busoniu.net/rltrack`.

Proposal submitted to 20th IFAC World Congress. Received September 15, 2016.