



**THE UNIVERSITY OF GEORGIA
DEPARTMENT OF STATISTICS**

Colloquium Series

Csaba Szepesvari

Department of Computing Science, University of Alberta

Canada CIFAR AI Chair, Amii, DeepMind

3:45 PM, Thursday, April 29, 2021

Zoom link: <https://zoom.us/j/7389376577>

Planning, Learning, and generalization in Reinforcement Learning

Markov decision processes (MDPs) provide a powerful framework to capture some of the most important aspects of decision making under uncertainty. However, MDPs lack structure and as such planning and learning in MDPs with multidimensional state spaces, common in applications, is strongly intractable. Reinforcement learning methods introduce either value-function approximation, or policy approximation, or both to address this issue. The key idea is that using a powerful function approximation method, such as neural networks, algorithms can extrapolate far away from the data that they are trained on, which ideally leads to increased efficiency and effectiveness. In this talk, building on recent results, I will discuss whether (and when) it is reasonable to expect this approach to work, describing situations where intractability still remains, while also describing results when it is avoided. I will conclude with discussing open questions.

Bio: Csaba Szepesvari is a Canada CIFAR AI Chair, the team-lead for the “Foundations” team at DeepMind and a Professor of Computing Science at the University of Alberta. He earned his PhD in 1999 from Jozsef Attila University, in Szeged, Hungary. In addition to regularly publishing at top tier journals and conferences, he has (co-)authored three books. Currently, he serves as the action editor of the Journal of Machine Learning Research and as an associate editor of the Mathematics of Operations Research journal, in addition to serving regularly on program committees of various machine learning and AI conferences. Dr. Szepesvari's main interest is developing principled, learning-based approaches to artificial intelligence (AI). He is the co-inventor of UCT, an influential Monte-Carlo tree search algorithm, a variant of which was used in the AlphaGo program which, in a landmark game, defeated the top Go professional Lee Sedol in 2016, ten years of the invention of UCT. In 2020, Dr. Szepesvari co-founded the weekly “Reinforcement Learning Theory virtual seminar series”, which showcases top theoretical work in the area of reinforcement learning with speakers and which is open to attendees from all over the world.

For more information, please contact us at: stat@uga.edu