

# Integrating Domain-Knowledge into Deep Learning

Ruslan Salakhutdinov

Apple

rsalakhu@cs.cmu.edu

## ABSTRACT

In this talk, I will first discuss deep learning models that can find semantically meaningful representations of words, learn to read documents and answer questions about their content. I will show how we can encode external linguistic knowledge as an explicit memory in recurrent neural networks, and use it to model co-reference relations in text. I will further introduce methods that can augment neural representation of text with structured data from Knowledge Bases for question answering, and show how we can use structured prior knowledge from Knowledge Graphs for image classification. Finally, I will introduce the notion of structured memory as being a crucial part of an intelligent agent's ability to plan and reason in partially observable environments. I will present a modular hierarchical reinforcement learning agent that can learn to store arbitrary information about the environment over long time lags, perform efficient exploration and long-term planning, while generalizing across domains and tasks.

## BIOGRAPHY

Ruslan Salakhutdinov is a UPMC Professor of Computer Science in the Department of Machine Learning at CMU and is a director of AI research at Apple. He received his PhD in computer science from the University of Toronto in 2009. After spending two post-doctoral years at the Massachusetts Institute of Technology Artificial Intelligence Lab, he joined the University of Toronto as an Assistant Professor in the Departments of Statistics and Computer Science. In 2016 he joined CMU. Ruslan's primary interests lie in deep learning, machine learning, and large-scale optimization. He is an action editor of the Journal of Machine Learning Research and served on the senior programme committee of several top-tier learning conferences including NeurIPS and ICML.



He is an Alfred P. Sloan Research Fellow, Microsoft Research Faculty Fellow, Canada Research Chair in Statistical Machine Learning, a recipient of the Early Researcher Award, Google Faculty Award, Nvidia's Pioneers of AI award, and is a Senior Fellow of the Canadian Institute for Advanced Research.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

KDD '19, August 4–8, 2019, Anchorage, AK, USA.

© 2019 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-6201-6/19/08.

DOI: <https://doi.org/10.1145/3292500.3340416>