

Learning and Reasoning with Graph-Structured Data

ICML 2019 Workshop

Overview

Learning and reasoning with graph-structured representations is gaining increasing interest in both academia and industry, due to its fundamental advantages over more traditional unstructured methods in supporting interpretability, causality, and transferability / inductive generalization. Recently, there is a surge of new techniques in the context of deep learning, such as graph neural networks, for learning graph representations and performing reasoning and prediction, which have achieved impressive progress. However, it can still be a long way to go to obtain satisfactory results in long-range multi-step reasoning, scalable learning with very large graphs, flexible modeling of graphs in combination with other dimensions such as temporal variation and other modalities such as language and vision. New advances in theoretical foundations, models and algorithms, as well as empirical discoveries and applications are therefore all highly desirable. The aims of this workshop are to bring together researchers to dive deeply into some of the most promising methods which are under active exploration today, discuss how we can design new and better benchmarks, identify impactful application domains, encourage discussion and foster collaboration.

Important Dates

- **Submission deadline:** ~~Thursday, 18 April 2019~~ **Friday, 26 April 2019** (23:59 AoE)
- **Author notification:** Friday, 10 May 2019
- **Camera ready deadline:** Friday, 31 May 2019
- **Workshop:** 14 or 15 June 2019

Call for Papers

We invite the submission of papers on topics including, but not limited to:

- Deep learning methods on graphs/manifolds/relational data (e.g., graph neural networks)
- Deep generative models of graphs (e.g., for drug design)
- Unsupervised graph/manifold/relational embedding methods (e.g., hyperbolic embeddings)
- Optimization methods for graphs/manifolds/relational data
- Relational or object-level reasoning in machine perception
- Relational/structured inductive biases for reinforcement learning, modeling multi-agent behavior and communication
- Neural-symbolic integration
- Theoretical analysis of capacity/generalization of deep learning models for graphs/manifolds/relational data
- Benchmark datasets and evaluation metrics

Submissions should be no more than 4 pages, excluding references or supplementary materials. All submissions must be in pdf format as a single file (incl. supplementary materials) using the [workshop template](#) and submitted through CMT under [this link](#).

The review process is single-round and double-blind (submission files have to be anonymized). Previously published work and concurrent submissions are acceptable.

All accepted papers will be presented as posters during the workshop and listed on the website. Additionally, a small number of accepted papers will be selected to be presented as contributed or spotlight talks.

For any questions, please email graph-reason@googlegroups.com.

Speakers

- [Caroline Uhler](#), MIT
- [Song-Chun Zhu](#), UCLA
- [Pushmeet Kohli](#), DeepMind
- [Sanja Fidler](#), University of Toronto
- [William L. Hamilton](#), McGill University
- [Yaron Lipman](#), Weizmann Institute of Science
- [Tommi S. Jaakkola](#), MIT

Organizers

- [Ethan Fetaya](#), University of Toronto
- [Zhiting Hu](#), Carnegie Mellon University
- [Thomas Kipf](#), University of Amsterdam

- [Yujia Li](#), DeepMind
 - [Xiaodan Liang](#), Sun Yat-sen University
 - [Renjie Liao](#), University of Toronto
 - [Raquel Urtasun](#), University of Toronto
 - [Hao Wang](#), MIT
 - [Max Welling](#), University of Amsterdam
 - [Eric P. Xing](#), Carnegie Mellon University
 - [Richard Zemel](#), University of Toronto
-