

The Power of Relational Learning

Lise Getoor

Computer Science Department, University of California, Santa Cruz, US
getoor@soe.ucsc.edu

Abstract

We live in a richly interconnected world and, not surprisingly, we generate richly interconnected data. From smart cities to social media to financial networks to biological networks, data is relational. While database theory is built on strong relational foundations, the same is not true for machine learning. The majority of machine learning methods flatten data into a single table before performing any processing. Further, database theory is also built on a bedrock of declarative representations. The same is not true for machine learning, in particular deep learning, which results in black-box, uninterpretable and unexplainable models. In this talk, I will introduce the field of statistical relational learning, an alternative machine learning approach based on declarative relational representations paired with probabilistic models. I'll describe our work on probabilistic soft logic, a probabilistic programming language that is ideally suited to richly connected, noisy data. Our recent results show that by building on state-of-the-art optimization methods in a distributed implementation, we can solve very large relational learning problems orders of magnitude faster than existing approaches.

2012 ACM Subject Classification Computing methodologies → Machine learning; Information systems → Data management systems

Keywords and phrases Machine learning, Probabilistic soft logic, Relational model

Digital Object Identifier 10.4230/LIPICs.ICDT.2019.2

Category Invited Talk



© Lise Getoor;

licensed under Creative Commons License CC-BY

22nd International Conference on Database Theory (ICDT 2019).

Editors: Pablo Barcelo and Marco Calautti; Article No. 2; pp. 2:1–2:1

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany