



Laboratorio di Acquisizione
della Conoscenza e
Apprendimento delle Macchine



Centro
Interdipartimentale di
Logica e
Applicazioni



DIPARTIMENTO
DI INFORMATICA

Dottorato di Ricerca in Informatica e Matematica

Prof. Dr. Kristian Kersting

Computer Science Department Data Mining Group

TU Dortmund University, Germany

Martedì 27 settembre, ore 10,30, Sala Consiglio, VII Piano, Dipartimento di Informatica

Lifted Machine Learning

Our minds make inferences that appear to go far beyond standard data science approaches. Whereas people can learn richer representations and use them for a wider range of data science tasks, data science algorithms have been mainly employed in a stand-alone context, constructing a single function from a table of training examples. In this seminar, I shall touch upon an approach to data science that can capture these human learning aspects by combining graphs, databases, and relational logic in general with statistical learning and optimization. Here, high-level (logical) features such as individuals, relations, functions, and connectives provide declarative clarity and succinct characterizations of the data science problem. While attractive from a modeling viewpoint, this declarative data science programming also often assuredly complicates the underlying model, making solving it potentially very slow. Hence, I shall also touch upon ways to reduce the solver costs. One promising direction to speed up is to cache local structures in the computational models. I shall illustrate this for probabilistic inference, linear programs, and convex quadratic programs, all working horses of data science.

Based on joint works with Martin Mladenov, Amir Globerson, Martin Grohe, Christopher Re, Sriraam Natarjan and many more.

Kristian Kersting is an Associate Professor for Computer Science at the TU Dortmund University, Germany. He received his PhD from the University of Freiburg, Germany, in 2006. After a PostDoc at MIT, he moved to the Fraunhofer IAIS and the University of Bonn using a Fraunhofer ATTRACT Fellowship. His main research interests are data mining, machine learning, and statistical relational AI, with applications to medicine, plant phenotyping, traffic, and collective attention. Kristian has published over 150 technical papers, and his work has been recognized by several awards, including the ECCAI Dissertation Award for the best AI dissertation in Europe.

He gave several tutorials at top venues and serves regularly on the PC (often at the senior level) of the top machine learning, data mining, and AI venues. Kristian co-founded the international workshop series on Statistical Relational AI and co-chaired ECML PKDD 2013, the premier European venue for Machine Learning and Data Mining, as well as the Best Paper Award Committee of ACM KDD 2015, the premier international venue for Data Mining.

In 2017, he will co-chair UAI, one of the premier international conferences on research related to knowledge representation, learning, and reasoning in the presence of uncertainty. Currently, he is an action editor of DAMI, MLJ, AIJ, and JAIR, the editor of JAIR's Da special track on Deep Learning, Knowledge Representation, and Reasoning, as well as on the editorial boards of Information and of New Generation Computing.