UNIVERSITÄT SALZBURG

Gastvortrag
Fachbereich Computerwissenschaften
31. Jänner 2019 um 15:00 Uhr

Der Fachbereich Computerwissenschaften der Paris-Lodron-Universität Salzburg

lädt am

Donnerstag, 31. Jänner 2019 um 15:00 Uhr

im Hörsaal T03 des Fachbereichs Computerwissenschaften,
Jakob-Haringer-Straße 2,

zum
Gastvortrag
von

Univ.-Prof. Dipl.-Wirt.-Inf. Dr.-Ing. Matthias Böhm
TU Graz

ein.

Univ.-Prof. Dr. Nikolaus Augsten
Host

On Optimizing Operator Fusion Plans for Large-Scale Machine Learning in SystemML

Large-scale machine learning (ML) underpins many applications that profoundly transform our lives, but ML systems to execute these workloads are still in their infancy. In a first part of this talk, we give an overview of Apache SystemML as a representative ML system for declarative, large-scale ML. SystemML provides an R-like syntax and automatically compiles these high-level linear algebra programs into hybrid runtime plans of single-node, in-memory operations, and distributed operations on Spark. In a second part, we then present a selected research result on optimizing operator fusion plans. The opportunities for fused operators - in terms of fused chains of basis operators - are ubiquitous, and include fewer intermediates, scan sharing, and sparsity exploitation across operators. However, existing fusion heuristics struggle to find good plans for complex operator DAGs or hybrid plans. Therefore, we introduce an exact yet practical cost-based optimization framework for fusion plans, including techniques for candidate exploration, candidate selection, and code generation of local and distributed operations over dense, sparse, and compressed data. Finally, we share some lessons learned and ongoing work on properly supporting the entire end-to-end data science lifecycle.

Zur Person Matthias Böhm:

Matthias Boehm is a professor for data management in data science at Graz University of Technology, Austria, where he holds a BMVIT-endowed chair for data management. Prior to joining TU Graz in 2018, he was a research staff member at IBM Research - Almaden, USA, with a major focus on optimization and runtime techniques for declarative, large-scale machine learning. Since 2015, Matthias also serves as a PMC member for Apache SystemML. He received his Ph.D. from Dresden University of Technology, Germany in 2011 with a dissertation on cost-based optimization of integration flows. His previous research also includes systems support for time series forecasting as well as in-memory indexing and query processing. Matthias is a recipient of the 2016 VLDB Best Paper Award, and a 2016 SIGMOD Research Highlight Award.