

# **Data Management in Machine Learning Systems**

# Synthesis Lectures on Data Management

## Editor

**H.V. Jagadish**, *University of Michigan*

## Founding Editor

**M. Tamer Özsu**, *University of Waterloo*

*Synthesis Lectures on Data Management* is edited by H.V. Jagadish of the University of Michigan. The series publishes 80–150 page publications on topics pertaining to data management. Topics include query languages, database system architectures, transaction management, data warehousing, XML and databases, data stream systems, wide scale data distribution, multimedia data management, data mining, and related subjects.

## Data Management in Machine Learning Systems

Matthias Boehm, Arun Kumar, and Jun Yang  
2019

## Non-Volatile Memory Database Management Systems

Joy Arulraj and Andrew Pavlo  
2019

## Data Exploration Using Example-Based Methods

Matteo Lissandrini, Davide Mottin, Themis Palpanas, and Yannis Velegrakis  
2018

## Data Profiling

Ziawasch Abedjan, Lukasz Golab, Felix Naumann, and Thorsten Papenbrock  
2018

## Querying Graphs

Angela Bonifati, George Fletcher, Hannes Voigt, and Nikolay Yakovets  
2018

## Query Processing over Incomplete Databases

Yunjun Gao and Xiaoye Miao  
2018

### Natural Language Data Management and Interfaces

Yunyao Li and Davood Rafiei

2018

### Human Interaction with Graphs: A Visual Querying Perspective

Sourav S. Bhowmick, Byron Choi, and Chengkai Li

2018

### On Uncertain Graphs

Arijit Khan, Yuan Ye, and Lei Chen

2018

### Answering Queries Using Views

Foto Afrati and Rada Chirkova

2017

### Databases on Modern Hardware: How to Stop Underutilization and Love Multicores

Anatasia Ailamaki, Erieta Liarou, Pınar Tözün, Danica Porobic, and Iraklis Psaroudakis

2017

### Instant Recovery with Write-Ahead Logging: Page Repair, System Restart, Media Restore, and System Failover, Second Edition

Goetz Graefe, Wey Guy, and Caetano Sauer

2016

### Generating Plans from Proofs: The Interpolation-based Approach to Query Reformulation

Michael Benedikt, Julien Leblay, Balder ten Cate, and Efthymia Tsamoura

2016

### Veracity of Data: From Truth Discovery Computation Algorithms to Models of Misinformation Dynamics

Laure Berti-Équille and Javier Borge-Holthoefer

2015

### Datalog and Logic Databases

Sergio Greco and Cristina Molinaro

2015

### Big Data Integration

Xin Luna Dong and Divesh Srivastava

2015

### Instant Recovery with Write-Ahead Logging: Page Repair, System Restart, and Media Restore

Goetz Graefe, Wey Guy, and Caetano Sauer

2014

### Similarity Joins in Relational Database Systems

Nikolaus Augsten and Michael H. Böhlen  
2013

### Information and Influence Propagation in Social Networks

Wei Chen, Laks V.S. Lakshmanan, and Carlos Castillo  
2013

### Data Cleaning: A Practical Perspective

Venkatesh Ganti and Anish Das Sarma  
2013

### Data Processing on FPGAs

Jens Teubner and Louis Woods  
2013

### Perspectives on Business Intelligence

Raymond T. Ng, Patricia C. Arocena, Denilson Barbosa, Giuseppe Carenini, Luiz Gomes, Jr.,  
Stephan Jou, Rock Anthony Leung, Evangelos Milios, Renée J. Miller, John Mylopoulos, Rachel  
A. Pottinger, Frank Tompa, and Eric Yu  
2013

### Semantics Empowered Web 3.0: Managing Enterprise, Social, Sensor, and Cloud-based Data and Services for Advanced Applications

Amit Sheth and Krishnaprasad Thirunarayan  
2012

### Data Management in the Cloud: Challenges and Opportunities

Divyakant Agrawal, Sudipto Das, and Amr El Abbadi  
2012

### Query Processing over Uncertain Databases

Lei Chen and Xiang Lian  
2012

### Foundations of Data Quality Management

Wenfei Fan and Floris Geerts  
2012

### Incomplete Data and Data Dependencies in Relational Databases

Sergio Greco, Cristian Molinaro, and Francesca Spezzano  
2012

### Business Processes: A Database Perspective

Daniel Deutch and Tova Milo  
2012

### Data Protection from Insider Threats

Elisa Bertino  
2012

### Deep Web Query Interface Understanding and Integration

Eduard C. Dragut, Weiyi Meng, and Clement T. Yu  
2012

### P2P Techniques for Decentralized Applications

Esther Pacitti, Reza Akbarinia, and Manal El-Dick  
2012

### Query Answer Authentication

HweeHwa Pang and Kian-Lee Tan  
2012

### Declarative Networking

Boon Thau Loo and Wenchao Zhou  
2012

### Full-Text (Substring) Indexes in External Memory

Marina Barsky, Ulrike Stege, and Alex Thomo  
2011

### Spatial Data Management

Nikos Mamoulis  
2011

### Database Repairing and Consistent Query Answering

Leopoldo Bertossi  
2011

### Managing Event Information: Modeling, Retrieval, and Applications

Amarnath Gupta and Ramesh Jain  
2011

### Fundamentals of Physical Design and Query Compilation

David Toman and Grant Weddell  
2011

### Methods for Mining and Summarizing Text Conversations

Giuseppe Carenini, Gabriel Murray, and Raymond Ng  
2011

### Probabilistic Databases

Dan Suciu, Dan Olteanu, Christopher Ré, and Christoph Koch  
2011

### Peer-to-Peer Data Management

Karl Aberer  
2011

### Probabilistic Ranking Techniques in Relational Databases

Ihab F. Ilyas and Mohamed A. Soliman  
2011

### Uncertain Schema Matching

Avigdor Gal  
2011

### Fundamentals of Object Databases: Object-Oriented and Object-Relational Design

Suzanne W. Dietrich and Susan D. Urban  
2010

### Advanced Metasearch Engine Technology

Weiyi Meng and Clement T. Yu  
2010

### Web Page Recommendation Models: Theory and Algorithms

Sule Gündüz-Ögüdücü  
2010

### Multidimensional Databases and Data Warehousing

Christian S. Jensen, Torben Bach Pedersen, and Christian Thomsen  
2010

### Database Replication

Bettina Kemme, Ricardo Jimenez-Peris, and Marta Patino-Martinez  
2010

### Relational and XML Data Exchange

Marcelo Arenas, Pablo Barcelo, Leonid Libkin, and Filip Murlak  
2010

### User-Centered Data Management

Tiziana Catarci, Alan Dix, Stephen Kimani, and Giuseppe Santucci  
2010

### Data Stream Management

Lukasz Golab and M. Tamer Özsu  
2010

### Access Control in Data Management Systems

Elena Ferrari  
2010

### An Introduction to Duplicate Detection

Felix Naumann and Melanie Herschel

2010

### Privacy-Preserving Data Publishing: An Overview

Raymond Chi-Wing Wong and Ada Wai-Chee Fu

2010

### Keyword Search in Databases

Jeffrey Xu Yu, Lu Qin, and Lijun Chang

2009

© Springer Nature Switzerland AG 2022

Reprint of original edition © Morgan & Claypool 2019

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

Data Management in Machine Learning Systems

Matthias Boehm, Arun Kumar, and Jun Yang

ISBN: 978-1-031-00741-5      paperback

ISBN: 978-3-031-01869-5      ebook

ISBN: 978-1-031-00056-6      hardcover

DOI 10.1007/978-3-031-01869-5

A Publication in the Springer series

*SYNTHESIS LECTURES ON DATA MANAGEMENT*

Lecture #57

Series Editor: H.V. Jagadish, *University of Michigan*

Founding Editor: M. Tamer Özsu, *University of Waterloo*

Series ISSN

Print 2153-5418    Electronic 2153-5426



# Data Management in Machine Learning Systems

Matthias Boehm

Graz University of Technology

Arun Kumar

University of California, San Diego

Jun Yang

Duke University

*SYNTHESIS LECTURES ON DATA MANAGEMENT #57*

## ABSTRACT

Large-scale data analytics using machine learning (ML) underpins many modern data-driven applications. ML systems provide means of specifying and executing these ML workloads in an efficient and scalable manner. Data management is at the heart of many ML systems due to data-driven application characteristics, data-centric workload characteristics, and system architectures inspired by classical data management techniques.

In this book, we follow this data-centric view of ML systems and aim to provide a comprehensive overview of data management in ML systems for the end-to-end data science or ML lifecycle. We review multiple interconnected lines of work: (1) ML support in database (DB) systems, (2) DB-inspired ML systems, and (3) ML lifecycle systems. Covered topics include: in-database analytics via query generation and user-defined functions, factorized and statistical-relational learning; optimizing compilers for ML workloads; execution strategies and hardware accelerators; data access methods such as compression, partitioning and indexing; resource elasticity and cloud markets; as well as systems for data preparation for ML, model selection, model management, model debugging, and model serving. Given the rapidly evolving field, we strive for a balance between an up-to-date survey of ML systems, an overview of the underlying concepts and techniques, as well as pointers to open research questions. Hence, this book might serve as a starting point for both systems researchers and developers.

## KEYWORDS

ML systems, data management, data science, ML lifecycle, ML training, ML serving, in-database analytics, linear algebra, optimizing compilers, distributed machine learning, hardware accelerators, data access methods, resource elasticity, data cleaning, model management, benchmarking

# Contents

<b>Preface</b> .....	<b>xiii</b>
<b>Acknowledgments</b> .....	<b>xv</b>
<b>1 Introduction</b> .....	<b>1</b>
1.1 Overview of ML Lifecycle and ML Users .....	1
1.2 Motivation .....	3
1.3 Outline and Scope .....	4
<b>2 ML Through Database Queries and UDFs</b> .....	<b>7</b>
2.1 Linear Algebra .....	8
2.2 Iterative Algorithms .....	12
2.3 Sampling-Based Methods .....	14
2.4 Discussion .....	16
2.5 Summary .....	18
<b>3 Multi-Table ML and Deep Systems Integration</b> .....	<b>21</b>
3.1 Learning over Joins .....	22
3.2 Statistical Relational Learning and Non-IID Models .....	26
3.3 Deeper Integration and Specialized DBMSs .....	29
3.4 Summary .....	32
<b>4 Rewrites and Optimization</b> .....	<b>33</b>
4.1 Optimization Scope .....	33
4.2 Logical Rewrites and Planning .....	36
4.3 Physical Rewrites and Operators .....	42
4.4 Automatic Operator Fusion .....	47
4.5 Runtime Adaptation .....	50
4.6 Summary .....	52

<b>5</b>	<b>Execution Strategies</b>	<b>53</b>
5.1	Data-Parallel Execution	54
5.2	Task-Parallel Execution	58
5.3	Parameter Servers (Model-Parallel Execution)	62
5.4	Hybrid Execution Strategies	66
5.5	Accelerators (GPUs, FPGAs, ASICs)	68
5.6	Summary	71
<b>6</b>	<b>Data Access Methods</b>	<b>73</b>
6.1	Caching and Buffer Pool Management	73
6.2	Compression	76
6.3	NUMA-Aware Partitioning and Replication	79
6.4	Index Structures	80
6.5	Summary	83
<b>7</b>	<b>Resource Heterogeneity and Elasticity</b>	<b>85</b>
7.1	Provisioning, Configuration, and Scheduling	86
7.2	Handling Failures	90
7.3	Working with Markets of Transient Resources	93
7.4	Summary	98
<b>8</b>	<b>Systems for ML Lifecycle Tasks</b>	<b>101</b>
8.1	Data Sourcing and Cleaning for ML	101
8.2	Feature Engineering and Deep Learning	105
8.3	Model Selection and Model Management	110
8.4	Interaction, Visualization, Debugging, and Inspection	115
8.5	Model Deployment and Serving	117
8.6	Benchmarking ML Systems	119
8.7	Summary	120
<b>9</b>	<b>Conclusions</b>	<b>123</b>
	<b>Bibliography</b>	<b>127</b>
	<b>Authors' Biographies</b>	<b>157</b>

# Preface

Machine learning (ML) systems can be defined in a narrow or broad sense. First, there is a variety of ML applications like classification, regression, and clustering algorithms, as well as other tasks of the entire knowledge discovery or data science lifecycle. In a narrow sense, an ML system is simply the software system that executes these ML applications. Second, these software systems employ and extend techniques from multiple research communities. This includes: (1) compilation techniques from high performance computing (HPC), programming language compilers, and query compilation; (2) runtime techniques—for execution and data access—from database systems, operating systems, and distributed systems; as well as (3) hardware accelerators from computer architecture. Thus, in a broad sense, an ML system comprises the ML applications as well as the entire software and hardware system and its infrastructure.

The data management community and related fields have been working for over a decade on challenging ML workloads and dedicated ML systems. In modern ML systems, data management is central because: (1) ML applications benefit from high-quality and high-quantity input data; (2) data access is often a performance bottleneck whose repetitive workload characteristics can be optimized via classic data management techniques; and (3) common system architectures range from ML in database (DB) systems, through DB-inspired ML systems, to ML lifecycle systems. Especially in the last few years, many new ML systems have been introduced. These systems cover a range of very different architectures but also employ a number of common concepts and techniques. This rapidly evolving field, many discussions with team members from emerging systems, as well as our own experience with building such ML systems—like *RIOT*, *SystemML*, *Bismarck*, *Cumulon*, *Columbus*, *Hamlet*, and *Morpheus*—inspired us to give an overview of systems and techniques as well as research challenges and opportunities in the form of a SIGMOD 2017 tutorial.

The goal of this book is to expand on that effort by systematically structuring the up-to-date literature on ML systems and by providing an overview of the underlying techniques. In detail, we aim to survey the major categories of ML systems, as well as their key concepts and techniques at different levels of the system stack from programming model, through compilation and optimization, to runtime backends. The survey also includes orthogonal aspects such as resource elasticity and different tasks in the ML lifecycle. We hope this structure will also foster a discussion on common system architectures and components of ML systems. Such a consolidation of terminology and concepts could help systems researchers and developers alike to avoid inevitable redundancy when building end-to-end systems, identify baselines and research opportunities, and adopt recent advances of the field. Finally, the material of this book

might also serve as the basis for advanced courses on the architecture and implementation of ML systems at undergraduate or graduate level.

Matthias Boehm, Arun Kumar, and Jun Yang  
February 2019

# Acknowledgments

First of all, we want to thank H. V. Jagadish for encouraging and supporting us to extend our SIGMOD 2017 tutorial *Data Management in Machine Learning: Challenges, Techniques, and Systems* into this book project. Although we certainly underestimated the effort of writing this book, retrospectively, we are glad to have invested this effort. Furthermore, we are grateful for having worked on a variety of academic and industrial ML systems as part of larger teams. Hence, this book draws tremendously from ideas, contributions, and discussions of many people—too many to be named—from different communities. Obviously, the trend of open source ML systems and early sharing of system internals was an essential basis for the material in this book as well. Finally, we thank Bettina Kemme (and her students Joseph Vinish D’silva and Hanfeng Chen), Sudip Roy, and Zack Ives, for their valuable comments and suggestions on an early draft of this book.

Matthias Boehm, Arun Kumar, and Jun Yang  
February 2019