

Architecture of ML Systems (AMLS) 01 Introduction and Overview

Prof. Dr. Matthias Boehm

Technische Universität Berlin Berlin Institute for the Foundations of Learning and Data Big Data Engineering (DAMS Lab)





2 Matthias Boehm | FG DAMS | AMLS SoSe 2025 – 01 Introduction and Overview

Announcements / Org

#1 Hybrid & Video Recording

- Hybrid lectures (in-person, zoom) with optional attendance <u>https://tu-berlin.zoom.us/j/9529634787?pwd=R1ZsN1M3SC9BOU1OcFdmem9zT202UT09</u>
- Zoom video recordings, links from website <u>https://mboehm7.github.io/teaching/ss25_amls/index.htm</u>

#2 Course Registrations

- TU Berlin ISIS registrations as of Apr 17
- Bachelor/Master/PhD ratio? CS/others ratio?

zoom







3 Matthias Boehm | FG DAMS | AMLS SoSe 2025 – 01 Introduction and Overview

#3 Faculty IV - Team Awareness and Antidiscrimination https://www.tu.berlin/eecs/awan

- Goal
 - Low-barrier approachability for spectrum of awareness and antidiscrimination issues
- Team
 - Irene Hube-Achter (MTSV)
 - Matthias Boehm (professors)
 - Nadine Karsten (scientific personnel)
 - Tom Hersperger (students)

Mission Statement

- Account for heterogeneity and complexity of modern societies at TU Berlin
- All employees and students are committed to
 - #1 Treat all persons with fairness and respect
 - #2 Ensure a safe environment for all
 - #3 Comply with our duty of care towards others
 - #4 Actively support the implementation of the above guidelines and contribute

Contact: private email, eecs-TB-awareness@win.tu-berlin.de, or <u>AwAn@dams.tu-berlin.de</u>







Agenda

berlin

- FG Big Data Engineering (DAMS Lab)
- Motivation and Goals
- Course Organization and Logistics
- Course Outline, and Projects
- Apache SystemDS and DAPHNE





FG Big Data Engineering (DAMS Lab)

https://www.tu.berlin/dams



About Me

- Since 09/2022 TU Berlin, Germany
 - University professor for Big Data Engineering (DAMS)

• 2018-2022 TU Graz, Austria

- BMK endowed chair for data management + research area manager
- Data management for data science (DAMS), SystemDS & DAPHNE

2012-2018 IBM Research – Almaden, CA, USA

- Declarative large-scale machine learning
- Optimizer and runtime of Apache SystemML
- 2007-2011 PhD TU Dresden, Germany
 - Cost-based optimization of integration flows
 - Time series forecasting / in-memory indexing & query processing













FG Big Data Engineering (DAMS Lab) – Team

Secretary (09/2023) Sarah Hashmi



- Postdoc (04/2019) Arnab Phani
- **PhD Student** (01/2020) Sebastian Baunsgaard Graduating 2024/25
- PhD Student (10/2023) **Christina Dionysio**



- PhD Student (10/2023) **Philipp Ortner** → 🕅 FONDA
- PhD Student (03/2023) **David Justen** \rightarrow BIFOLD GS
- PhD Student (01/2024) Ramon Schöndorf CHARITÉ Agility Project
- PhD Student (02/2023) **Carlos E. Muniz Cuza**
 - \checkmark \rightarrow Agility Project
- PhD Student (10/2024) Grigorii Turchenko

BIFOLD GS w/ Steffen Zeuch



PhD Student (04/2021) Saeed Fathollahzadeh **Concordia** external



berlin

PhD Student (01/2025) **Martin Seyferth**

Physikalisch Technische Burdesanstalt external

PhD Student (02/2025) Alexander Wolotschai

Mercedes-Benz external

- PhD Student (04/2025) Martin Moesmann external AALBORG UNIVERSITET
- 3x Student Assistants
- 25x Bachelor/Master Students
- 3x Interns and Visitors







FG Big Data Engineering (DAMS Lab) – Teaching

Successfully Established TUB Teaching Portfolio (modules, slides)









Motivation and Goals



Example ML Applications (Past/Present)



Transportation / Space

- Lemon car detection and reacquisition (classification, seq. mining)
- Airport passenger flows from WiFi data (time series forecasting)
- Data analysis for driving assistance (blind spot detection, emergency braking, lane centering)
- Automotive vehicle development (ML-assisted simulations, hyper-parameter tuning, ejector optimization)
- Earth observation and local climate zone classification and monitoring, reproducibility, compression

Finance

- Insurance claim cost per customer (model selection, regression)
- Financial analysts survey correlation (bivariate stats w/ new tests)

Health Care

- Breast / lunch / stomach cancer from histopathology images (classification)
- Glucose trends and warnings (clustering, classification)
- Emergency room diagnosis / patient similarity (classification, clustering)
- Patient survival analysis and prediction (Cox regression, Kaplan-Meier)



A Car Reacquisition Scenario







Example ML Applications



Production/Manufacturing

- Paper and fertilizer production (regression/classification, anomalies)
- Semiconductor manufacturing (ion beam tuning), and material degradation modeling (survival analysis)
- Mixed waste stream sorting and recycling (composition, alignment, quality)

Other Domains

- Machine data: errors and correlation (bivariate stats, seq. mining)
- Smart grid: energy demand/RES supply, weather models (forecasting)

Information Extraction

- NLP contracts
 rights/obligations (classification, error analysis)
- PDF table recognition and extraction, OCR (NMF clustering, custom)
- Learning explainable linguistic expressions (learned FOL rules, classification)
- Algorithm Research (+ various state-of-the art algorithms)
 - User/product recommendations via various forms of NMF; word-embeddings via orthogonalized skip-gram
 - Localized, supervised metric learning (dim reduction and classification)



What is an ML System?





What is an ML System?, cont.



ML System

- Narrow focus: SW system that executes ML applications
- Broad focus: Entire system (HW, compiler/runtime, ML application)
- → Trade-off runtime/resources vs accuracy
- Early days: no standardizations (except some exchange formats), lots of different languages and system architectures, but many shared concepts

Course Objectives

- Architecture and internals of modern (large-scale) ML systems
 - Microscopic view of ML system internals
 - Macroscopic view of ML pipelines and data science lifecycle
- #1 Understanding of characteristics → better evaluation / usage
- #2 Understanding of effective techniques → build/extend ML systems





Course Organization and Logistics



Basic Course Organization



Staff

- Lecturer: Prof. Dr. Matthias Boehm (replacement: Dr. Patrick Damme)
- Teaching Assistants: M.Sc. Sebastian Baunsgaard (projects/exercises)

Language

- Lectures and slides: English
- Communication and examination: English/German

Course Format

- TU Berlin: VL+UE 3+2 SWS, 6 ECTS (3x 1 ECTS + 2x 1.5 ECTS), master / TU Graz: VU 3 SWS, 5 ECTS
- Weekly lectures (start 4.15pm sharp, including Q&A), attendance optional
- Mandatory programming project or exercises (~3 ECTS)
- Recommended papers for additional reading on your own
- Prerequisites (preferred)
 - Basic courses: data management/databases, distributed systems, applied ML
 - Completed bachelor (not mandatory)





Course Logistics

Website

- https://mboehm7.github.io/teaching/ss25_amls/index.htm
- All course material (lecture slides) and dates
- Video Recording / Live Streaming Lectures (zoom)

Communication

- Informal language (first name is fine)
- Please, immediate feedback (unclear content, missing background)
- ISIS forum for offline questions, after lecture, and via email/PR discussions
- Office hours: Tue 4pm-5.30pm (Sebastian in B 119) or after lecture

Exam

- Completed project / exercise (checked by me/staff, no plagiarism incl *-GPT)
- Final exam (written exam or delegated oral exams, if >70 students take the exam)
- Grading (project/exercises completion, 100% exam) → new: +5 exam points if exercises >=90 points









Course Logistics, cont.

berlin

Course Applicability TU Berlin

- Master programs computer engineering, computer science, electrical engineering, information systems management
 - Area Data and Software Engineering
 - Area Cognitive Systems
 - Area Distributed Systems

Course Applicability TU Graz (remote – offered in exceptions)

- Master programs computer science (CS), as well as software engineering and management (SEM)
 - Catalog Data Science (compulsory course in major, and elective)
 - Catalog Machine Learning (elective course)
 - Catalog Interactive and Visual Information Systems (elective course)
 - Catalog Software Technology (elective course)
- PhD CS doctoral school list of courses

Free subject course in any other study program





Outline and Projects

Created SoSe 2019, partially based on



[Matthias Boehm, Arun Kumar, Jun Yang: Data Management in Machine Learning Systems. Synthesis Lectures on Data Management, Morgan & Claypool Publishers 2019] Major updates in SoSe 2020 – SoSe 2025



Part A: Overview and ML System Internals

- 01 Introduction and Overview [Apr 17]
- 02 Languages, Architectures, and System Landscape [Apr 24]
- 03 Size Inference, Rewrites, and Operator Selection [Fri, May 02]
- 04 Operator Fusion and Runtime Adaptation [Fri, May 09]
- 05 Data- and Task-Parallel Execution [May 15]
- O6 Parameter Servers [May 22]
- 07 LLM Training and Inference [Jun 05]
- 08 Hybrid Execution and HW Accelerators [Jun 12]
- 09 Caching, Partitioning, Indexing, and Compression [Jun 19]



Labor Day May 01 80y EoWW2 May 08

Ascension of Jesus May 29



Part B: ML Lifecycle Systems

- 10 Data Acquisition, Cleaning, and Preparation [Jun 26]
- 11 Model Selection and Management [Jul 03]
- 12 Model Debugging, Fairness, and Explainability [Jul 10]
- 13 Model Serving Systems and Techniques [Jul 17]
 - Incl Q&A and Exam Preparation

Planned Written Exams

- Thu Jul 24, 4pm (in ???)
- Thu Jul 31, 4pm (in ???)
- Thu Aug 14, 4pm (in ???)





Programming Projects

Open-Source Projects

- Programming project in context of open-source projects
 - Apache SystemDS: <u>https://github.com/apache/systemds</u>
 - DAPHNE: <u>https://github.com/daphne-eu/daphne</u>
 - Other OSS projects possible, but harder to merge PRs
- Commitment to open source and open communication (PRs, mailing list)
- **Remark:** Don't be afraid to ask questions / develop code in public

Objectives

- Non-trivial feature in an ML system (3 ECTS → 75-90 hours)
- OSS processes: Break down into subtasks, code/tests/docs, one or more pull-requests per project, code review, incorporate review comments, etc

Team

Individuals or up to three-person teams (w/ separated responsibilities)



Lists of projects available on website

Project/Exercise Selection by May 01 EOD: https://tinyurl.com/yeyrjyxw





Alternative Exercise (published: Apr 12, submission: Jul 15)

berlin

- Task: ECG Time Series Classification
 - Task Description: <u>https://mboehm7.github.io/</u> teaching/ss25_amls/AMLS_2025_Exercise.pdf
 - Data Exploration (15): data acquisition, data characteristics, train/validation splits
 - Modeling (40): train ML models for classifying the ECG time series; hyper-parameter tuning
 - Data Augmentation and Feature Engineering (30)
 - Data Reduction (15): sampling, compression, others
- Objectives
 - End-to-end development of an ML pipeline for ECG data
 - Handle data exploration, modeling and tuning, data augmentation, and compression

Team

Individuals or up to three-person teams (w/ separated responsibilities)



0 .. Normal heart rhythm1 .. Tachyarrhythmia2 .. Other 3 .. Noisy





Apache SystemDS and DAPHNE



https://github.com/apache/systemds



https://github.com/daphne-eu/daphne





Apache SystemDS: A Declarative ML System for the End-to-End Data Science Lifecycle

https://github.com/apache/systemds





Landscape of ML Systems

Existing ML Systems

- #1 Numerical computing frameworks
- #2 ML Algorithm libraries (local, large-scale)
- #3 Linear algebra ML systems (large-scale)
- #4 Deep neural network (DNN) frameworks
- #5 Model management, and deployment

Exploratory Data-Science Lifecycle

- Open-ended problems w/ underspecified objectives
- Hypotheses, data integration, run analytics
- Unknown value \rightarrow lack of system infrastructure
 - → Redundancy of manual efforts and computation

Data Preparation Problem

- 80% Argument: 80-90% time for finding, integrating, cleaning data
- Diversity of tools → boundary crossing, lack of optimization



"Take these datasets and show value or competitive advantage"



berlin



The Data Science Lifecycle (aka KDD Process, aka CRISP-DM)







Apache SystemDS [https://github.com/apache/systemds]



berlin

Open Source SystemML Educate One Million

Establish Spark Technology Center

Language Abstractions and APIs









Basic HOP and LOP DAG Compilation





1.6GB

r'(CP)

У

→ Hybrid Runtime Plans:

- Size propagation / memory estimates
- Integrated CP / Spark runtime
- Dynamic recompilation during runtime



(X_{2,1})

 $X_{m,1}$

(persisted in

MEM_DISK)

Static and Dynamic Rewrites

- Example Static Rewrites (size-independent)
 - Common Subexpression Elimination
 - Constant Folding / Branch Removal / Block Sequence Merge
 - Static Simplification Rewrites
 - Right/Left Indexing Vectorization
 - For Loop Vectorization
 - Spark checkpoint/repartition injection
- Example Dynamic Rewrites (size-dependent)
 - Dynamic Simplification Rewrites
 - Matrix Mult Chain Optimization



Sparsity Estimation & Sparse DP Enum [SIGMOD'19] berlin



 $sum(X+Y) \rightarrow sum(X)+sum(Y)$









Data Cleaning Pipelines [SIGMOD'24, under submission]



- Automatic Generation of Cleaning Pipelines
 - Library of robust, parameterized data cleaning primitives
 - Enumeration of DAGs of primitives & hyper-parameter optimization (evolutionary, HB)



University	Country		University	Country
TU Graz	Austria		TU Graz	Austria
TU Graz	Austria]	TU Graz	Austria
TU Graz	Germany]	TU Graz	Austria
IIT	India		IIT	India
IIT	IIT		IIT	India
IIT	Pakistan		IIT	India
IIT	India	1	IIT	India
SIBA	Pakistan		SIBA	Pakistan
SIBA	null		SIBA	Pakistan
SIBA	null]	SIBA	Pakistan

0.77	0.80	1	1	
0.96	0.12	1	1	
0.66	0.09	null	1	
0.23	0.04	17	1	
0.91	0.02	17	null	
0.21	0.38	17	1	
0.31	null	17	1	
0.75	0.21	20	1	
null	null	20	1	
0.19	0.61	20	1	
0.64	0.31	20	1	

C

D

в

A	В	С	D
0.77	0.80	1	1
0.96	0.12	1	1
0.66	0.09	17	1
0.23	0.04	17	1
0.91	0.02	17	1
0.21	0.38	17	1
0.31	0.29	17	1
0.75	0.21	20	1
0.41	0.24	20	1
0.19	0.61	20	1
0.64	0.31	20	1

Dirty Data

After imputeFD(0.5)



SliceLine for Model Debugging [SIGMOD'21b, BTW'25b]





- **Problem Formulation**
 - Intuitive slice scoring function
 - Exact top-k slice finding
 - $|S| \ge \sigma \land sc(S) > 0, \alpha \in (0,1]$
- Properties & Pruning
 - Monotonicity of slice sizes, errors
 - Upper bound sizes/errors/scores
 - \rightarrow pruning & termination
- Linear-Algebra-based Slice Finding
 - Recoded/binned matrix X, error vector e
 - Vectorized implementation in linear algebra (join & eval via sparse-sparse matmult)
 - Local and distributed task/data-parallel execution



slice size

 $= \alpha \left(\frac{|X|}{|S|} \cdot \frac{\sum_{i=1}^{|S|} es_i}{\sum_{i=1}^{|X|} e_i} - 1 \right) - (1 - \alpha) \left(\frac{|X|}{|S|} - 1 \right)$

slice error



0 2 0

== Level

20

0 2 0

111

0

0

Multi-level Lineage Tracing & Reuse



m>>n

t(X)

- Lineage as Key Enabling Technique
 - Trace lineage of ops (incl. non-determinism), dedup for loops/funcs
 - Model versioning, data reuse, incr. maintenance, autodiff, debugging

Full Reuse of Intermediates

- Before executing instruction, probe output lineage in cache Map<Lineage, MatrixBlock>
- Cost-based/heuristic caching and eviction decisions (compiler-assisted)
- Partial Reuse of Intermediates
 - Problem: Often partial result overlap
 - Reuse partial results via dedicated rewrites (compensation plans)
 - Example: stepIm
- Extensions: multi-backend (Local, GPU, Spark), unified memory management



Compressed Linear Algebra Extended

[PVLDB'16a, VLDBJ'18, SIGMOD'23a,

under submission













Improved general applicability (adaptive compression time, new compression schemes, new kernels, intermediates, workload-aware)

User Script:

X = read("data/X")

v = read("data/v")

w = 12svm(X, y, TRUE)

write(w,"data/wXy")

X = scale(X,TRUE,TRUE)

1e-9,1e-3,100)

Built-in

Functions:

■ Sparsity → Redundancy exploitation (data redundancy, structural redundancy)

Workload-aware Compression

- Workload summary
 - \rightarrow compression
- Compressed Representation
 - \rightarrow execution planning

BWARE: Compressed Feature Engineering

- Frame compression, compressed I/O
- Compressed feature transformations
- Morphing of compressed data

Federated Learning [SIGMOD'21c]













- Federated data (matrices/frames) as meta data objects
- Federated linear algebra, (and federated parameter server)





Federated Requests: READ, PUT, GET, EXEC_INST, EXEC_UDF, CLEAR

Design Simplicity: (1) reuse instructions (2) federation hierarchies





DAPHNE: An Open and Extensible System Infrastructure for Integrated Data Analysis Pipelines

https://github.com/daphne-eu/daphne







Motivation

→ DAPHNE Overall Objective: **Open and extensible system infrastructure**

dense



- Integrated Data Analysis Pipelines
 - Open data formats, query processing
 - Data preprocessing and cleaning
 - ML model training and scoring
 - HPC, custom codes, and simulations

Hardware Challenges

- DM+ML+HPC share compilation and runtime techniques / converging cluster hardware
- End of Dennard scaling:
 - $P = \alpha CFV^2$ (power density 1)
- End of Moore's law
- Amdahl's law: sp = 1/s
- → Increasing Specialization

Deployment Challenges





DAPHNE Use Cases



[Xiao Xiang Zhu et al: So2Sat LCZ42: A Benchmark Dataset for the Classification of Global Local Climate Zones. **GRSM 8(3) 2020**]

[So2Sat LC42: https://mediatum.ub.tum.de/1454690]



DLR Earth Observation

- ESA Sentinel-1/2 datasets → 4PB/year
- Training of local climate zone classifiers on So2Sat LCZ42 (15 experts, 400K instances, 10 labels each, 85% confidence, ~55GB H5)
- ML pipeline: preprocessing, ResNet20, climate models





- IFAT Semiconductor Ion Beam Tuning
- KAI Semiconductor Material Degradation
- AVL Vehicle Development Process (ejector geometries, KPIs)
- ML-assisted simulations, data cleaning, augmentation



AVL

DAPHNE System Architecture [CIDR'22]



DaphneLib	(API) Python API w/ lazy evaluation	
Daphn	eDSL (Domain-specific Language)	Extensible Extension Catalog
	DaphneIR (MLIR Dialect)	(new data types, kernels,
	Optimization Passes	Multi-level Compilation/scheduling algs)
MLIR-Based Compilation Chain	New Runtime Abstractions for Data, Devices, Operations	Runtime Sideways Entry, DSL-level
	Hierarchical Scheduling	Fine-grained constraints
Device Kernels (CPU, GPU, FPGA, Storage)	Vectorized Execution Engine (Fused Op Pipelines)Sync/Async I/O Buffer/Memory Management	Fusion and Parallelism(e.g., data/op placement)
Local (embe	dded) and Distributed Environments	Integration w/

(standalone, HPC, data lake, cloud, DB)

Integration w/ Resource Mgmt & Prog. Models

Vectorized (Tiled) Execution



(%9, %10) = fusedPipeline1(%X, %y, %colmu, %colsd) {



Default Parallelization Frame & Matrix Ops Locality-aware, Multi-device Scheduling Fused Operator Pipelines on Tiles/Scalars + Codegen

Summary & QA





- FG Big Data Engineering (DAMS Lab)
- Motivation and Goals
- Course Organization and Logistics
- Course Outline, and Projects
- Apache SystemDS and DAPHNE

Programming Projects in SystemDS/DAPHNE, or Alternative Exercise Project/Exercise Selection by May 01 EOD: <u>https://tinyurl.com/yeyrjyxw</u>

- Next Lectures (Part A)
 - **02** Languages, Architectures, and System Landscape [Apr 24]
 - 03 Size Inference, Rewrites, and Operator Selection [May 02]
 - **04 Operator Fusion and Runtime Adaptation** [May 09]
 - 05 Data- and Task-Parallel Execution [May 15]
 - 06 Parameter Servers [May 22]
 - 07 LLM Training and Inference [Jun 05]

