New Methods for the Automated Analysis of Massive, Parallel, and Distributed Data Streams

Volker Markl

Professor of Computer Science, TU Berlin

Chief Scientist, DFKI Berlin

Director, Berlin Institute for the Foundations of Learning and Data (BIFOLD)

❶ Data as a Factor of Production

❷ Selected Research Contributions

❸ Current Research and Vision

Big Data and Machine Learning are the key drivers of innovation in AI and Data Science.

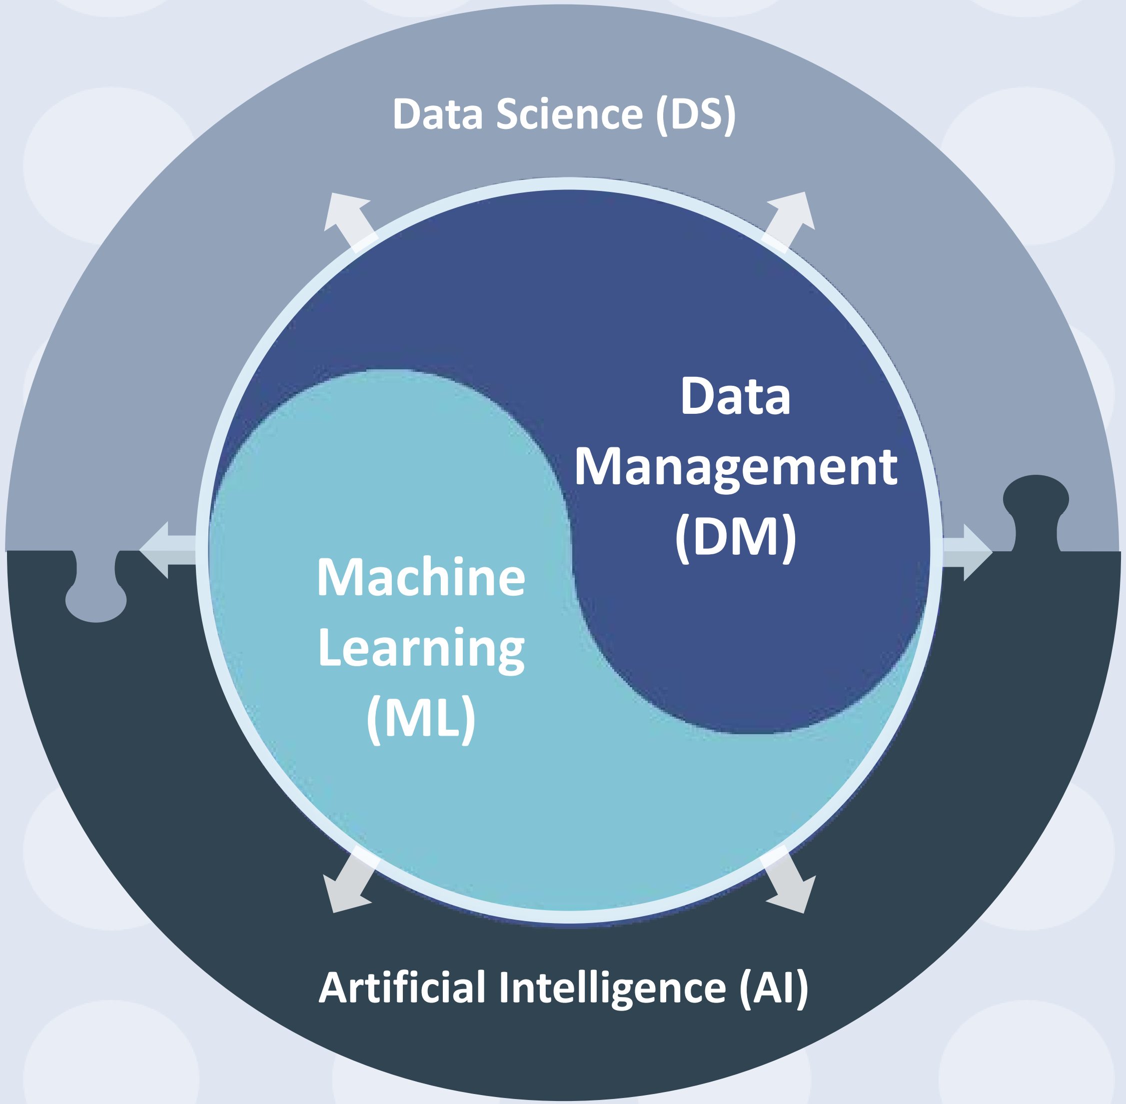

Data Management jointly with ML are disruptive in the Sciences, Humanities, and Industry.

Sciences and Humanities

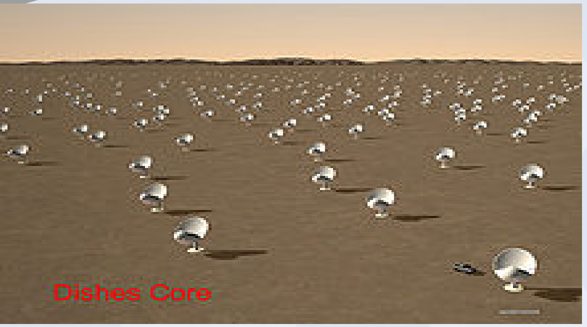

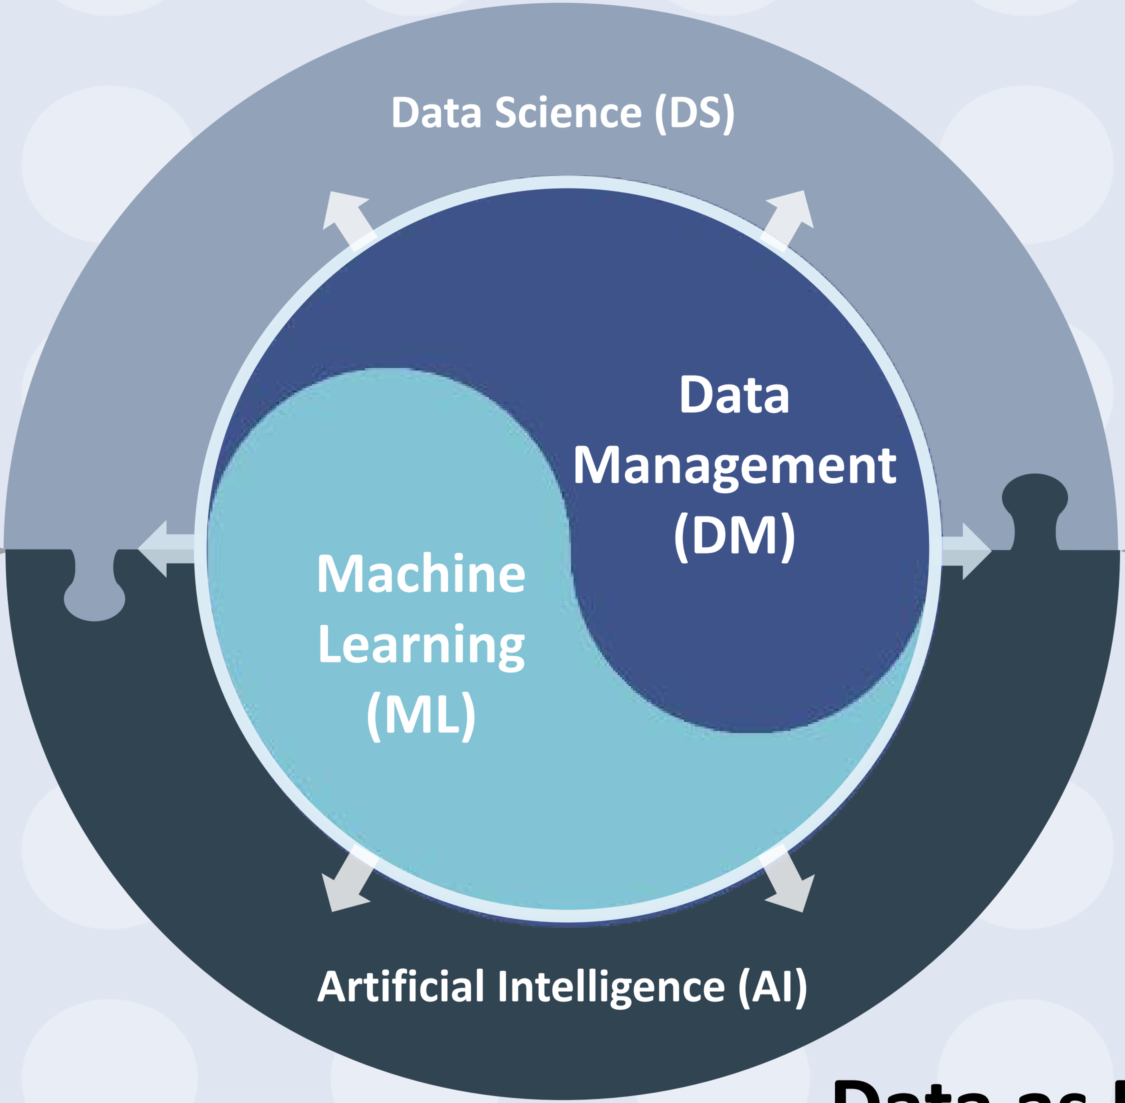

Industry

4th Paradigm

Data as Factor of Production

Remote Sensing

Nominated ESA–EGU EOS Excellence Award

Analysis of massive satellite data archives (Sentinel-2)

Labeled training data archive for AI models

3 TB / 5 days multimodal data

Fast classification and categorization

© scihub.copernicus.eu

- Enables:
  - analysis of trends (e.g., deforestation)
  - predictions about regions (e.g., draught)
- Added to popular "big data" catalogs: (e.g., Google Earth, Radiant MLHub, TensorFlow)

[1] BigEarthNet: A Large-Scale Benchmark Archive. IGARSS 2019.
[2] Multi-Label Remote Sensing. IEEE Access 2020.
[3] http://bigearth.net

(Begüm Demir, TU Berlin)

Industrie/Industry 4.0

Exploratory realtime analysis of sensor data streams

Prediction of paper quality during production

97 heterogeneous sensors

low latency

complex model building + information extraction and integration

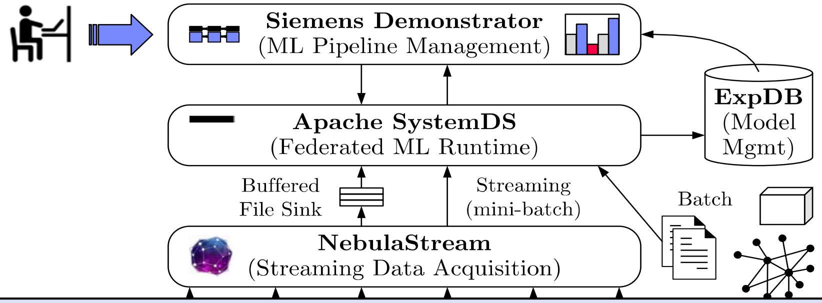

- Enables:
  - faster reaction to paper quality issues
  - cost reduction
- Data science on real time data streams

[1] The NebulaStream Platform. CIDR 2020.
[2] SystemDS. CIDR 2020.

(TU Graz, Siemens)

Data science process is complex and time-consuming

Goal of data management research:
scaling the process and its operationalization with respect to human and technical latency

Draws on results from:
computer architecture, statistics, machine learning, distributed systems, compilers, programming languages, etc.

Data science process is complex and time-consuming

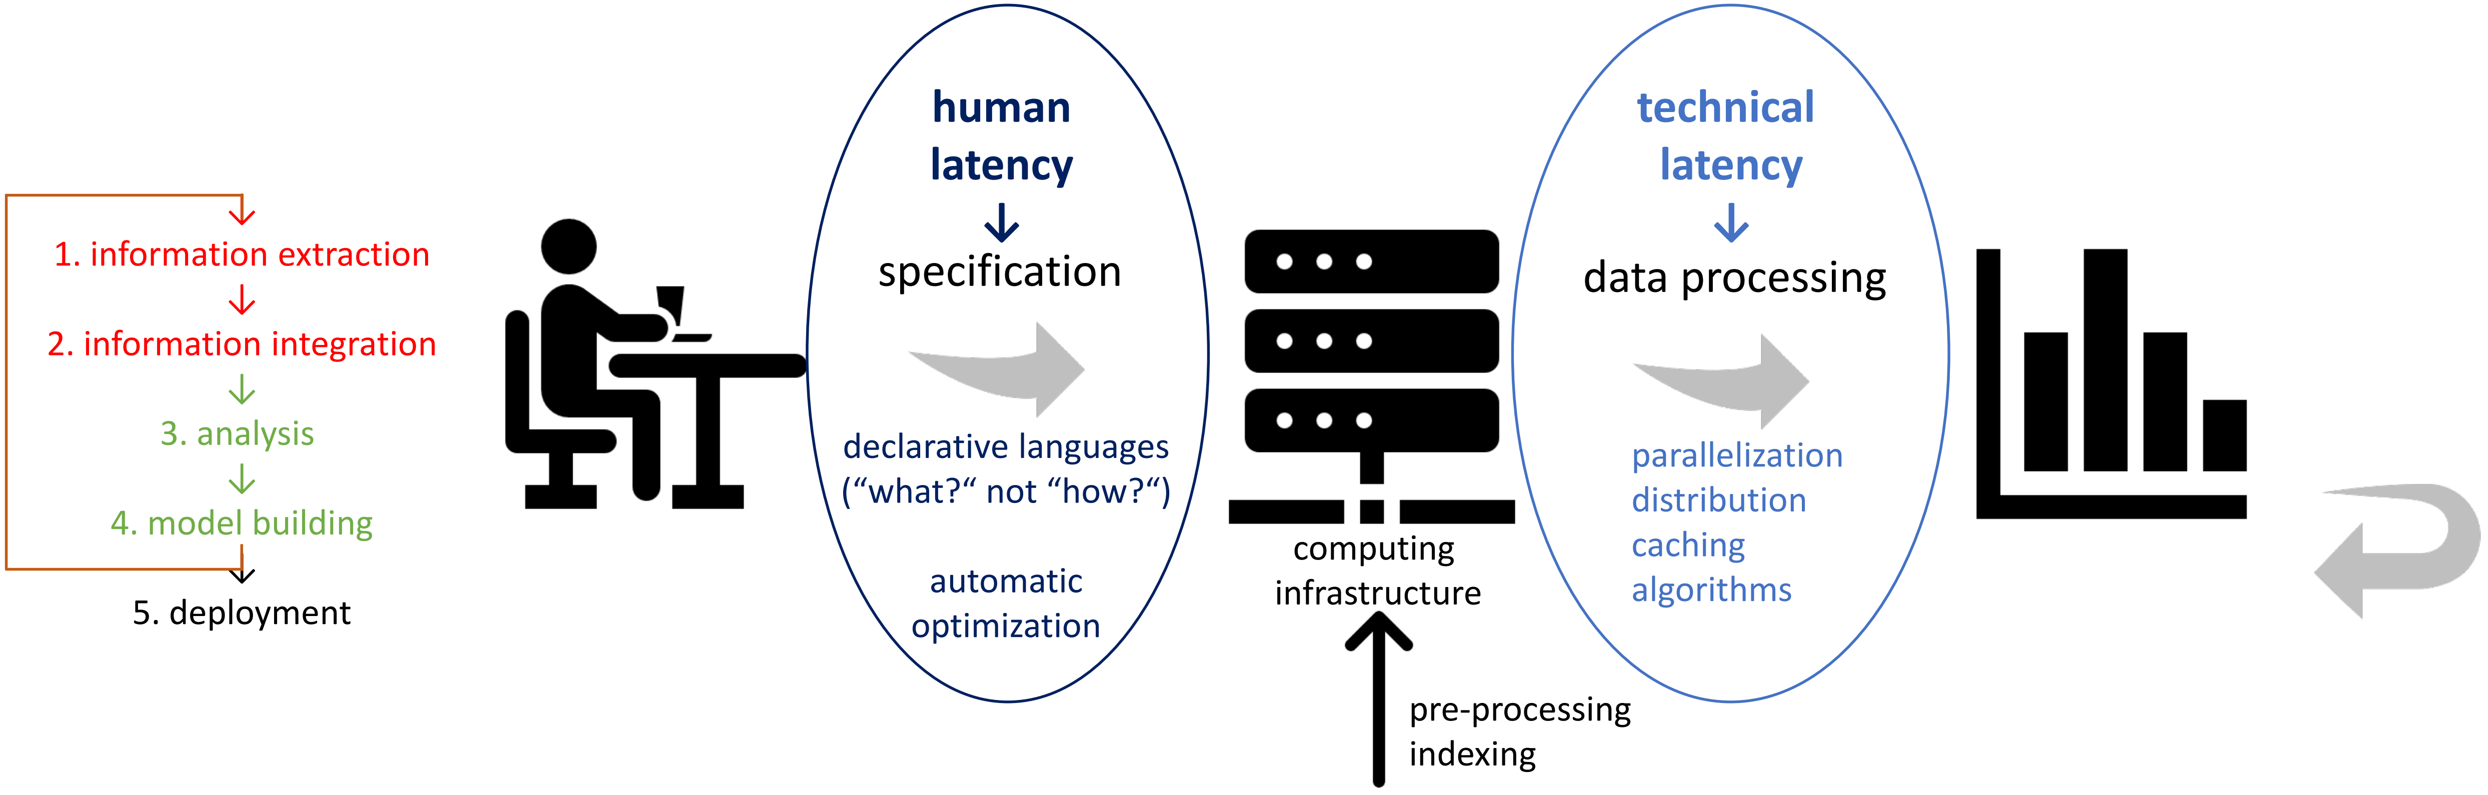

Goal of data management research:
scaling the process and its operationalization with respect to human and technical latency

Draws on results from:
computer architecture, statistics, machine learning, distributed systems, compilers, programming languages, etc.

Example: "3-Means-Clustering"

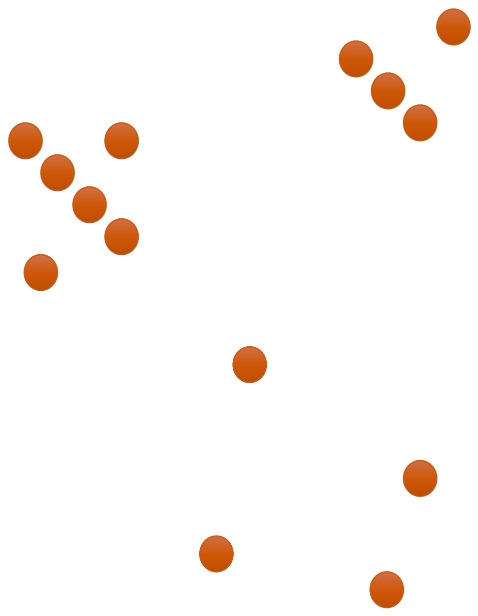

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

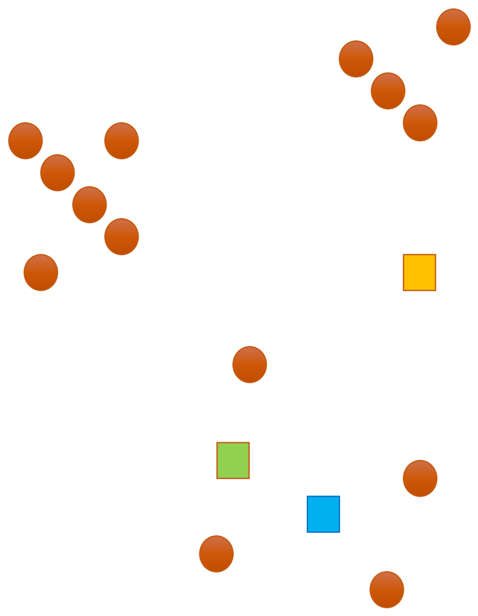

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

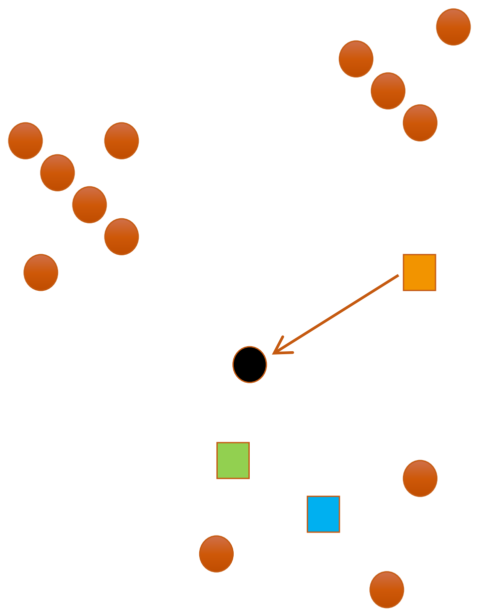

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

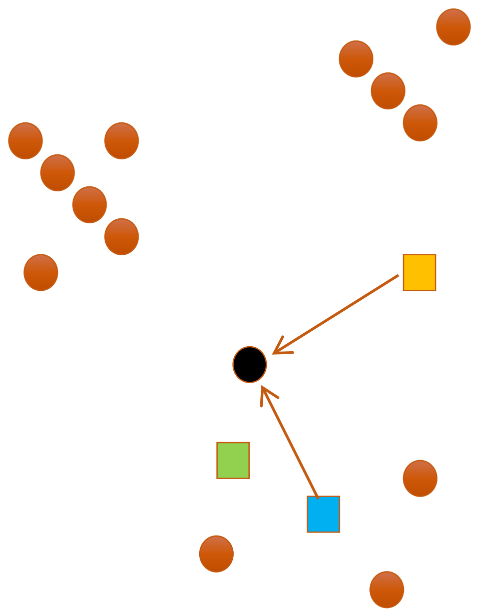

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

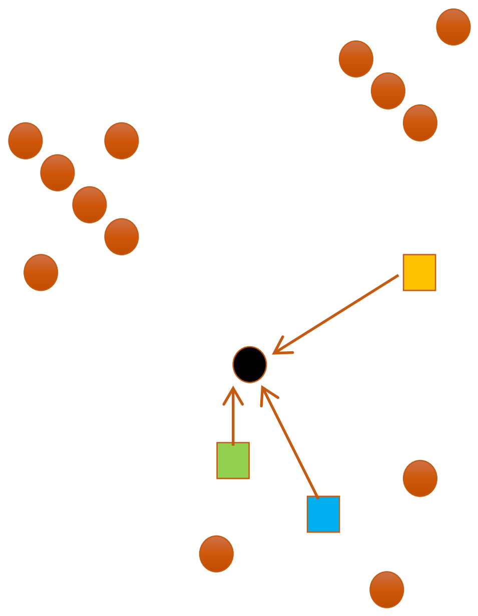

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

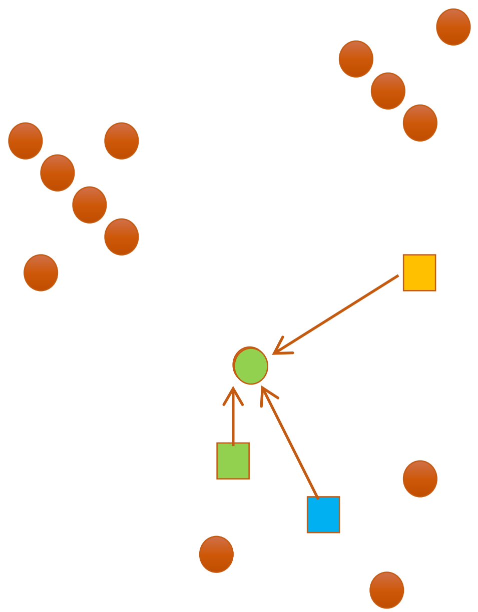

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

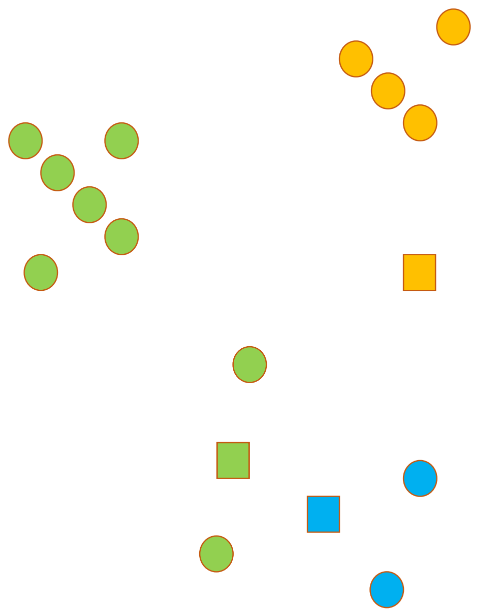

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

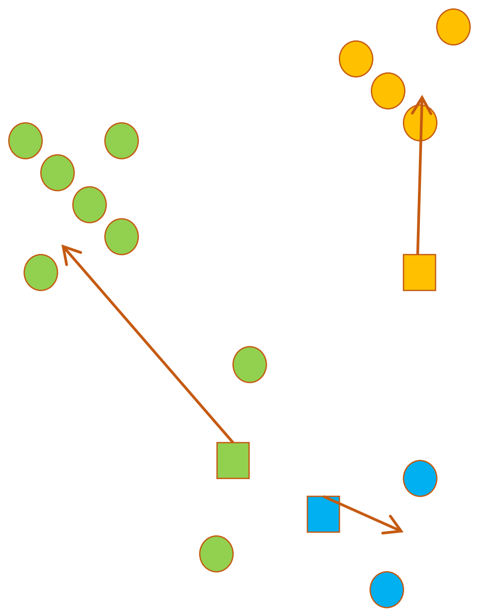

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

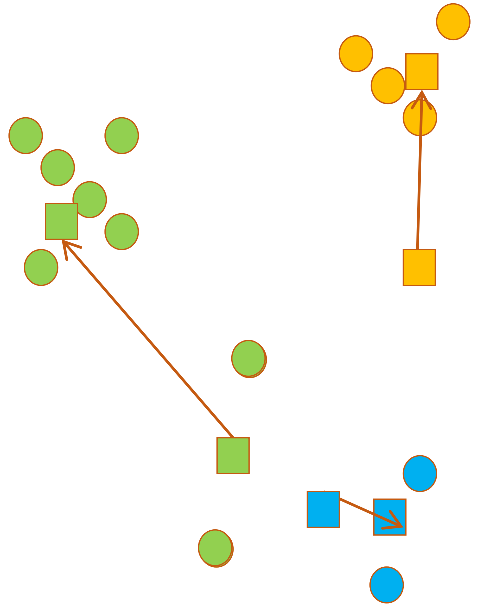

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

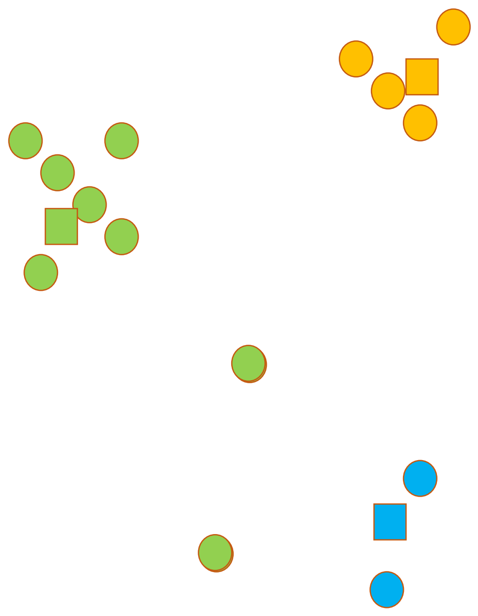

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

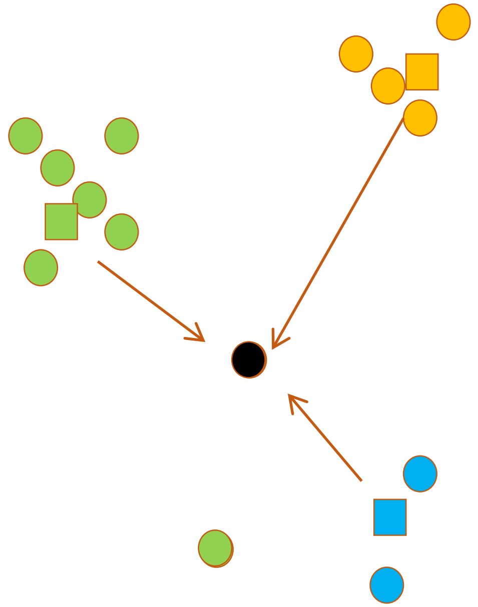

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

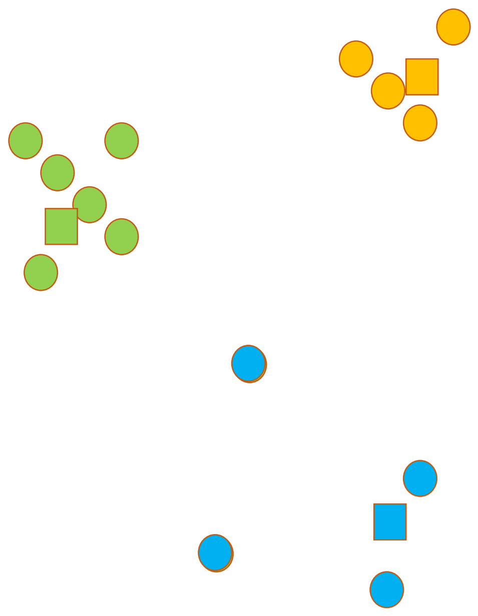

Choose 3 random cluster centers

Iterate until convergence:

Compute distance of each point to each center

Assign each point to the closest cluster

Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

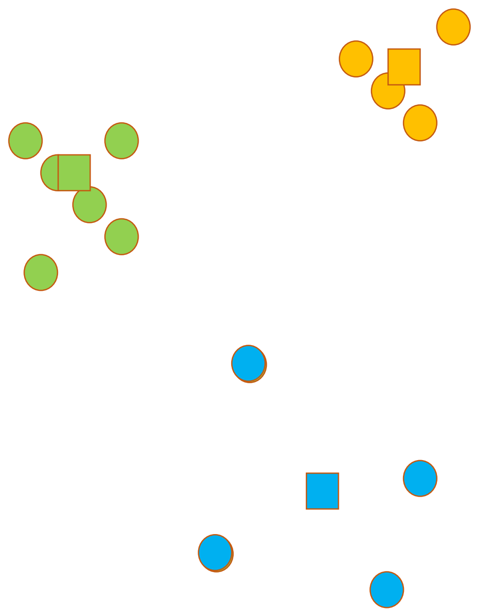

Choose 3 random cluster centers

Iterate until convergence:

Compute distance of each point to each center

Assign each point to the closest cluster

Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance. The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

Example: "3-Means-Clustering"

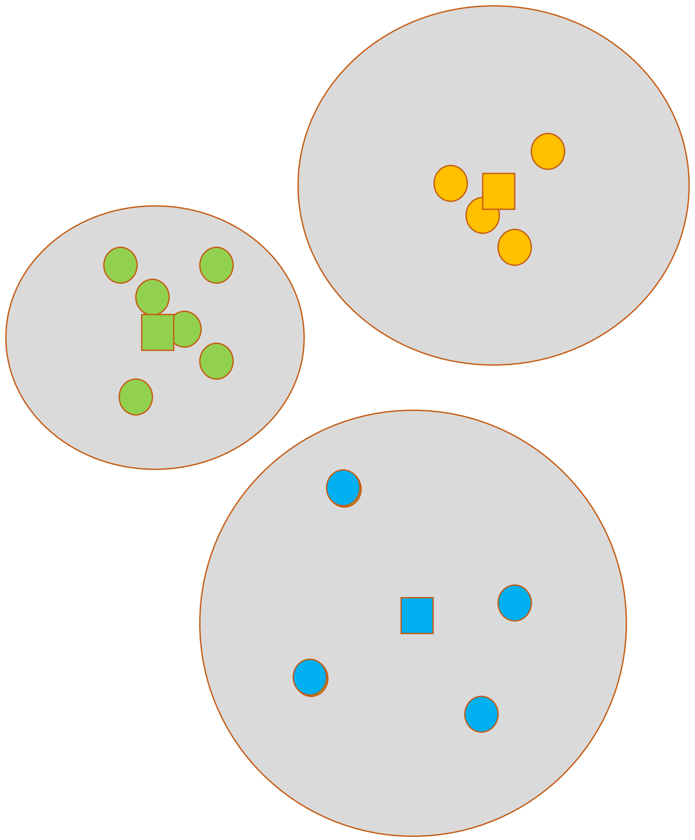

Choose 3 random cluster centers

Iterate until convergence:

- Compute distance of each point to each center
- Assign each point to the closest cluster
- Move centers

"3-Means-Clustering" is a simple data analysis method that divides a dataset into k groups (clusters) with respect to their relative distance.
The example illustrates an iterative algorithm to determine three groups (clusters) for a set of points according to a Euclidean distance.

“What“ not “How“

Example “k-Means-Clustering“

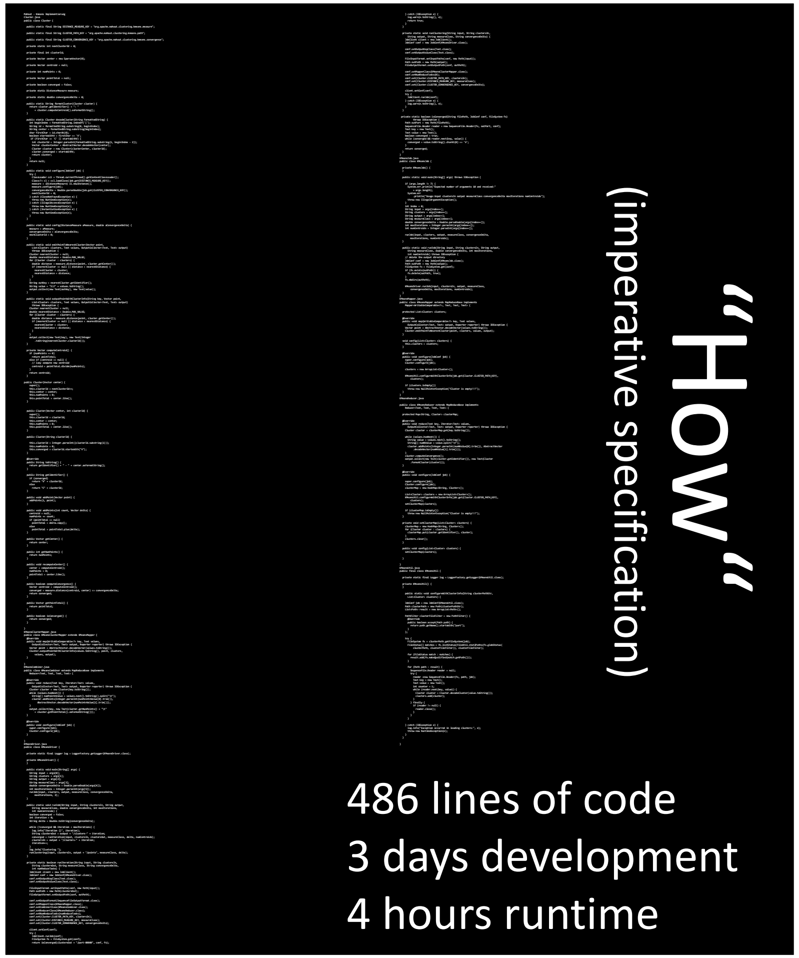

Hand-optimized code
(data-, workload- and system-dependent)

“What“, not “How“:
Example “k-Means-Clustering“

Declarative data flow program with automatic optimization, parallelization, and hardware adaption

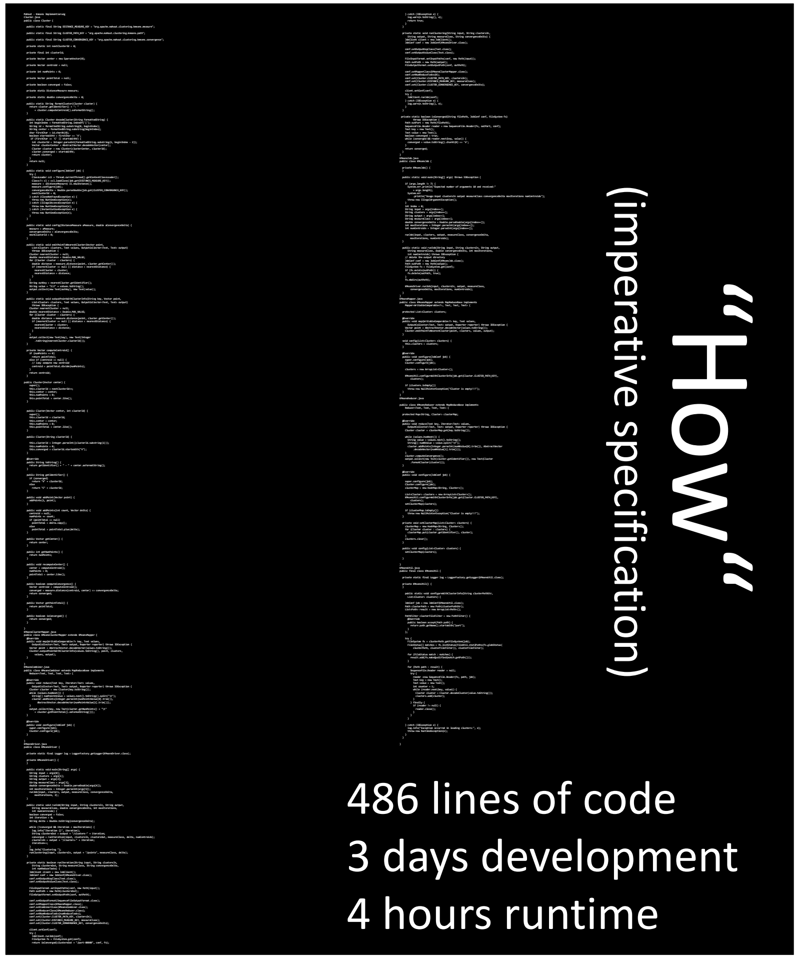

Hand-optimized code (data-, workload- and system-dependent)

❶ Data as a Factor of Production

❷ Selected Research Contributions: Reduction of Human and Technical Latency

① Automatic Optimization of Complex Analysis Algorithms for Big Data

② Analysis of Heterogeneous and Distributed Data Streams in the Internet of Things

❸ Current Research and Vision

Apache Flink: Data Programmability and Scalable Data Stream Analytics

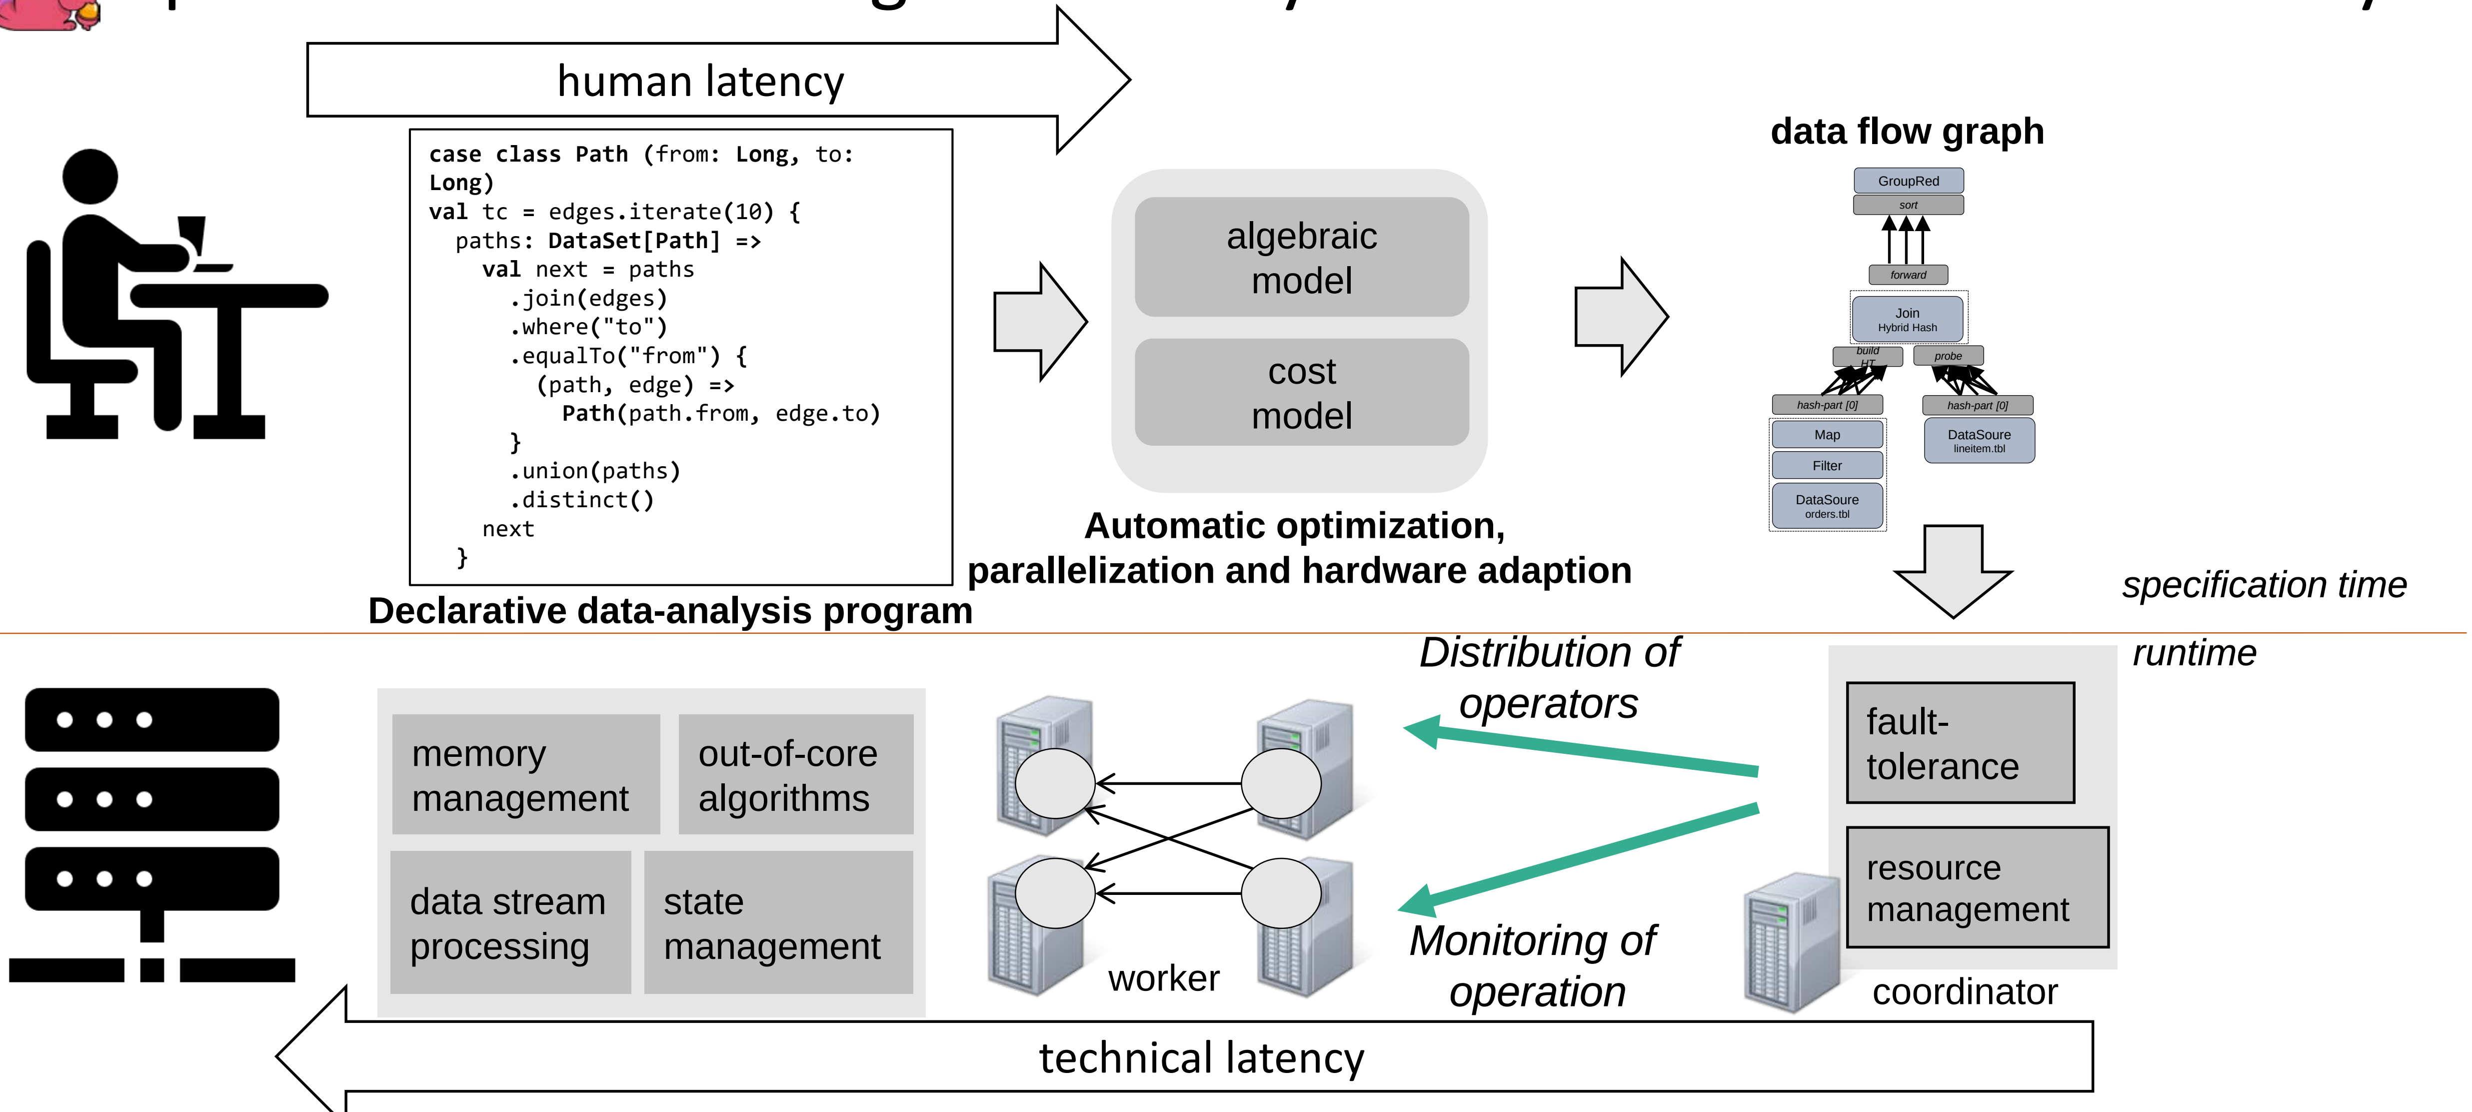

D. Battré, S. Ewen, F. Hueske, O. Kao, V. Markl, D. Warneke: *Nephele/PACTs: a programming model and execution framework for web-scale analytical processing. SoCC 2010: 119-130*

P. Carbone, A. Katsifodimos, S. Ewen, V. Markl, S. Haridi, K. Tzoumas: *Apache Flink™: Stream and Batch Processing in a Single Engine. IEEE Data Eng. Bull. 38(4): 28-38 (2015)*

Why Optimization?

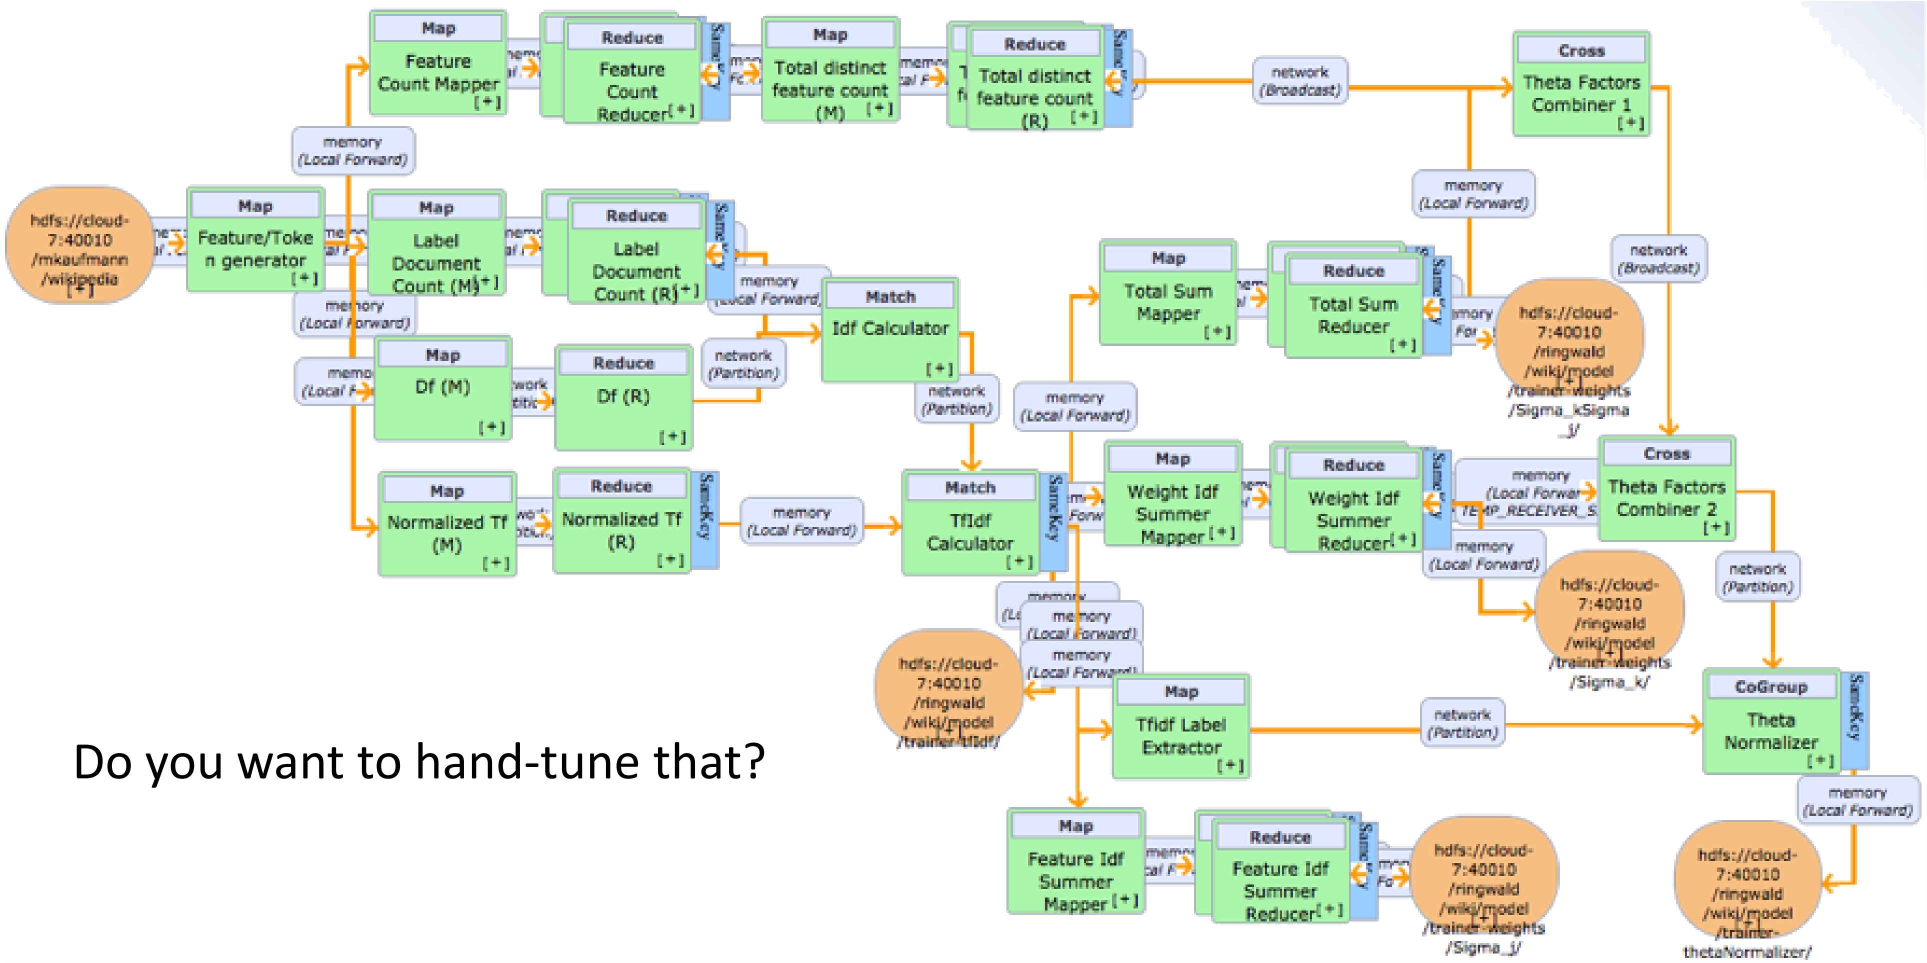

Do you want to hand-tune that?

F. Hueske, M. Peters, A. Krettek, M. Ringwald, K. Tzoumas, V. Markl, J.C. Freytag: Peeking into the optimization of data flow programs with MapReduce-style UDFs. ICDE 2013: 1292-1295

Impact of Automatic Optimization

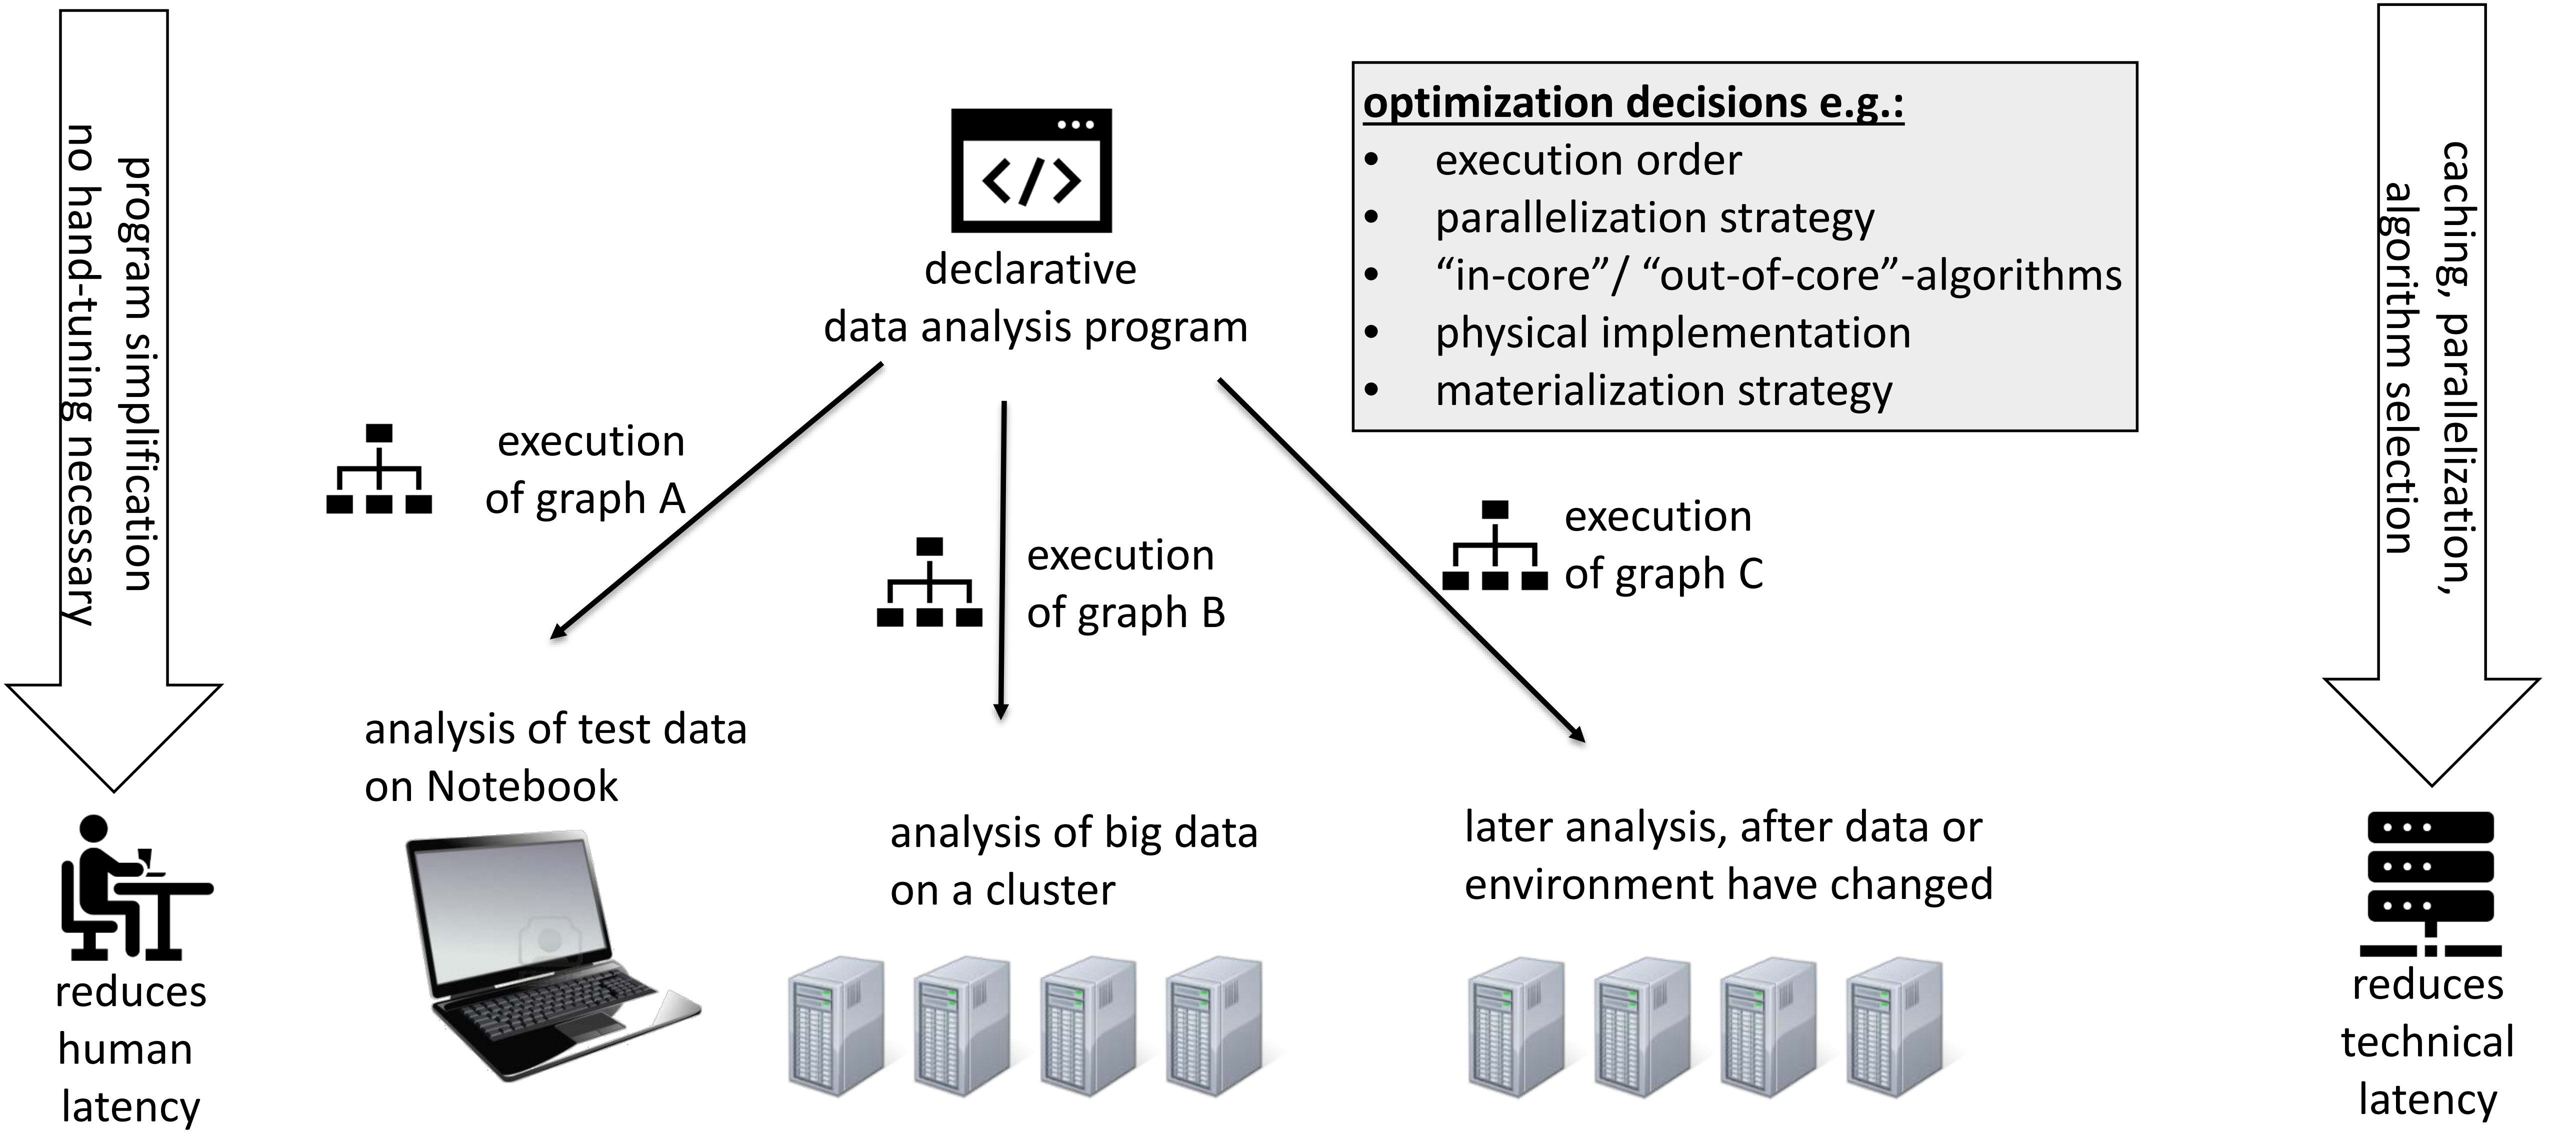

Effect of Optimization:
Lower latency and higher throughput in particular for streaming applications

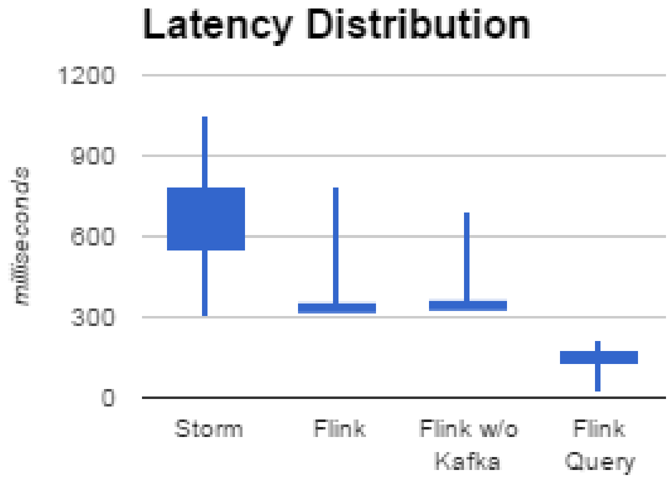

2x to 4x lower latency

50x higher throughput

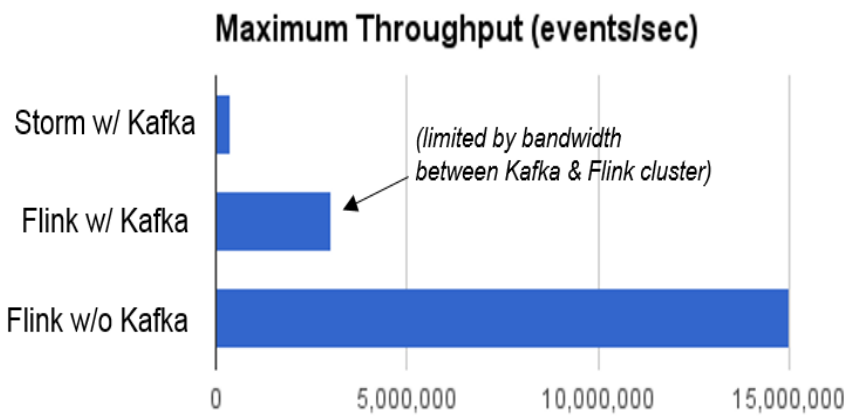

[1] https://www.ververica.com/blog/extending-the-yahoo-streaming-benchmark

compression of data streams for visualization

implicit parallelism through deep language embedding

latency reduction by result sharing

Apache Flink

https://www.meetup.com/topics/apache-flink/
https://flink.apache.org/poweredby.html
https://github.com/apache/flink

Flink Community

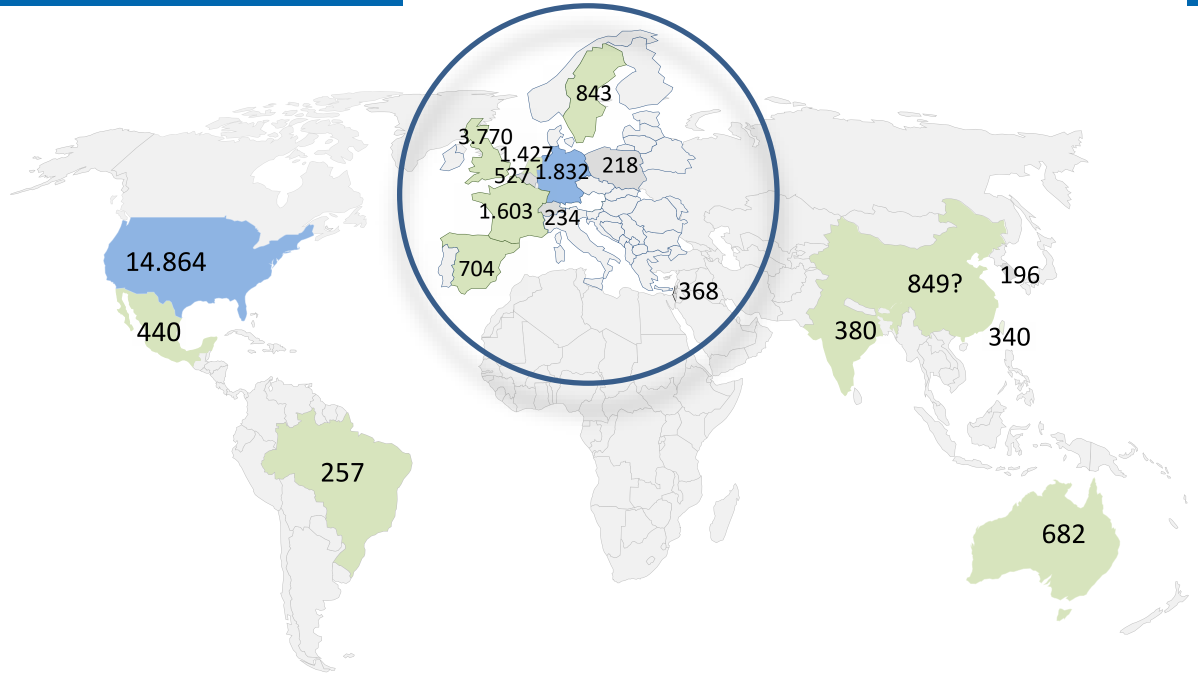

Flink Contributors

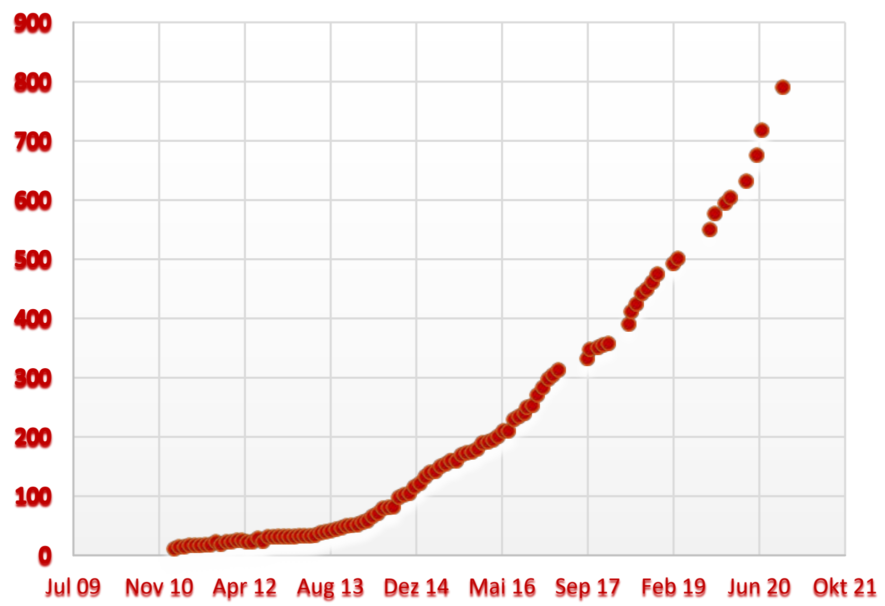

29,500+	Meetup Members Worldwide	18	Countries that Regularly Hold Meetups
870+	Open Source Contributors/Developers	49+	Companies using Apache Flink
49	Meetup Groups Worldwide		Startup data Artisans (now Ververica)

Last updated: May 2021

Some Highly Engaged Users

Largest job has > 20 operators, runs on > 5000 vCores in 1000-node cluster, processes millions of events per second

Complex jobs of > 30 operators running 24/7, processing 30 billion events daily, maintaining state of 100s of GB with exactly-once guarantees

30 Flink applications in production for more than one year; 10 billion events (2TB) processed daily

Courtesy of Kostas Tzoumas

Apache Flink Users

aeris
alephd We help publishers
Alibaba Group
amadeus
Atos
BetterCloud
Tencent 腾讯
altice
bouygues TELECOM
Capital One
dataArtisans
DRIVETRIBE
EURA NOVA
ERICSSON
FIRELAYERS
aws
HUAWEI
IMR INTERNET MEMORY RESEARCH
ING
KEEDIO
king
LINE
mbr targeting
MediaMath
MUX
KLAVIYO
NETFLIX
N MUSIC
OKKAM
otto group
Pragsis Bidoop Big Data Analytics
RADICALBIT Deep, Different, Disruptive.
InnoGames
ResearchGate
lyft
THIRD WATCH
Google
UBER
zalando
LIMEROAD New. Every day.
ROVIO
oppo
bol.com
Styx
yelp
COMCAST
ebay
Telefónica
Touk
criteo
D TRB

https://cwiki.apache.org/confluence/display/FLINK/Powered+by+Flink

➊ Data as a Factor of Production

➋ Selected Research Contributions: Reduction of Human and Technical Latency:

① Automatic Optimization of Complex Analysis Algorithms for Big Data

② Analysis of Heterogeneous and Distributed Data Streams in the Internet of Things

➌ Current Research and Vision

NebulaStream: Data Processing for the IoT

Reliable and efficent data stream management in a massively distributed, heterogeneous environment

Millions of data sources

Distribution

Heterogeneous hardware and data formats

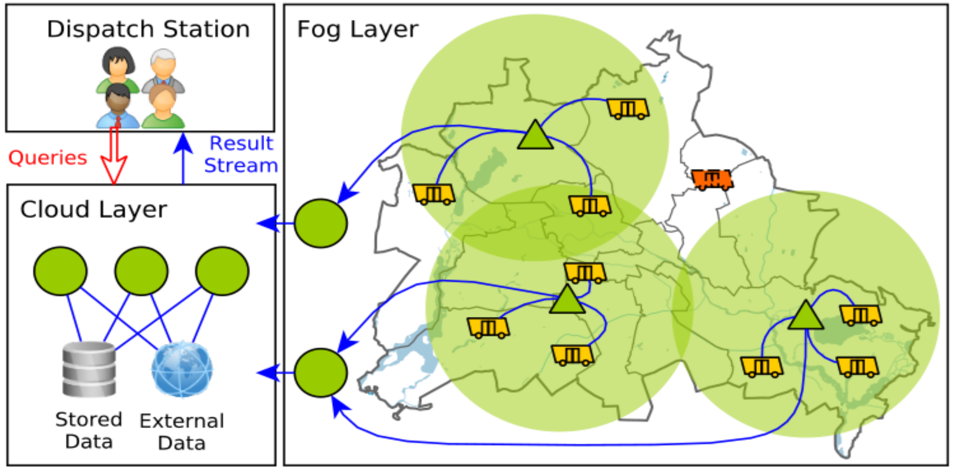

ACM SIGMOD Best Paper Award 2020

Latency reduction by heterogeneous hardware

"In-network"-pre-processing and "on-demand"-data processing

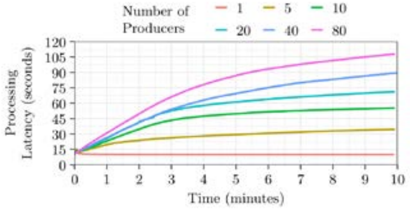

Hardware-targeted code generation

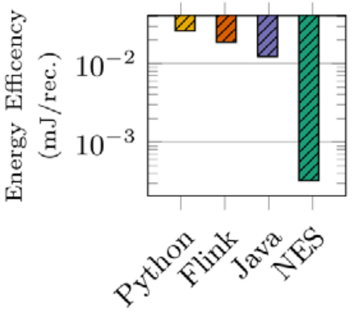

Intermediate representation to optimize heterogeneous data analysis progams

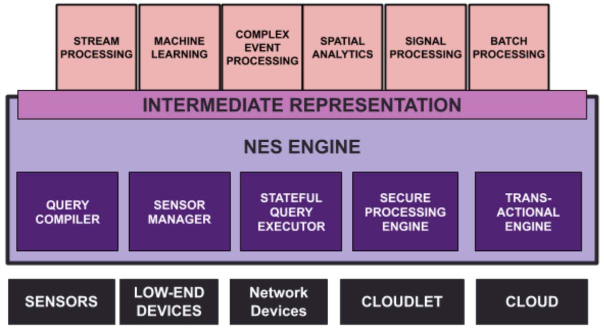

Decentralized error detection and intervention

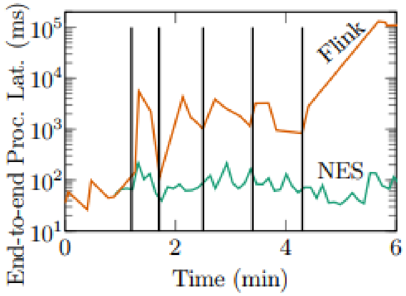

Dynamic optimization and lightweight adaptation of the distributed processing topology

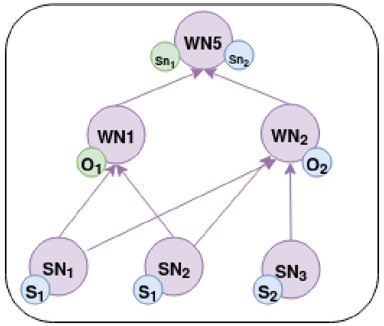

❶ Data as a Factor of Production

❷ Selected Research Contributions: Reduction of Human and Technical Latency

❸ Current Research and Vision:

Methods and Technologies for a <u>single</u> data infrastructure for AI-applications

Current Situation: Considerable backlog in AI innovations due to …

Fragmented AI project landscape

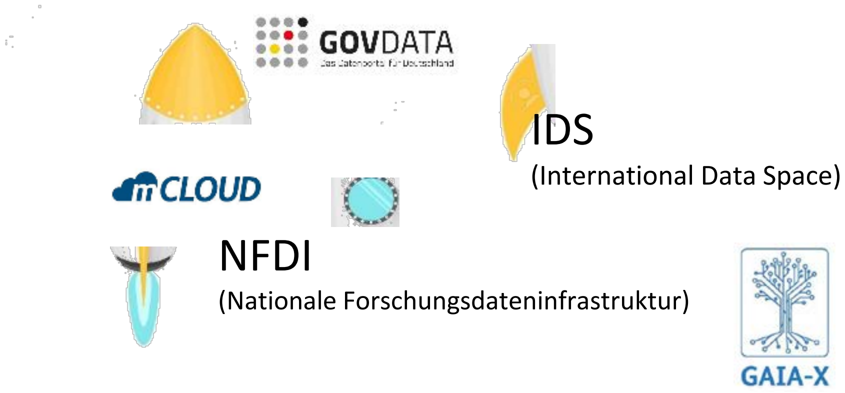

- Focus often on data exchange
- Interesting concepts exist – lack of implementation
- Multitude of efforts – lack of critical mass

Many little platforms
-> Redundant efforts
-> Competition for the best minds
-> Fragmented resources and expertise

Agora – One Data Management Infrastructure for AI-Applications

cross-sector
open and protected areas
collaborations
billing models

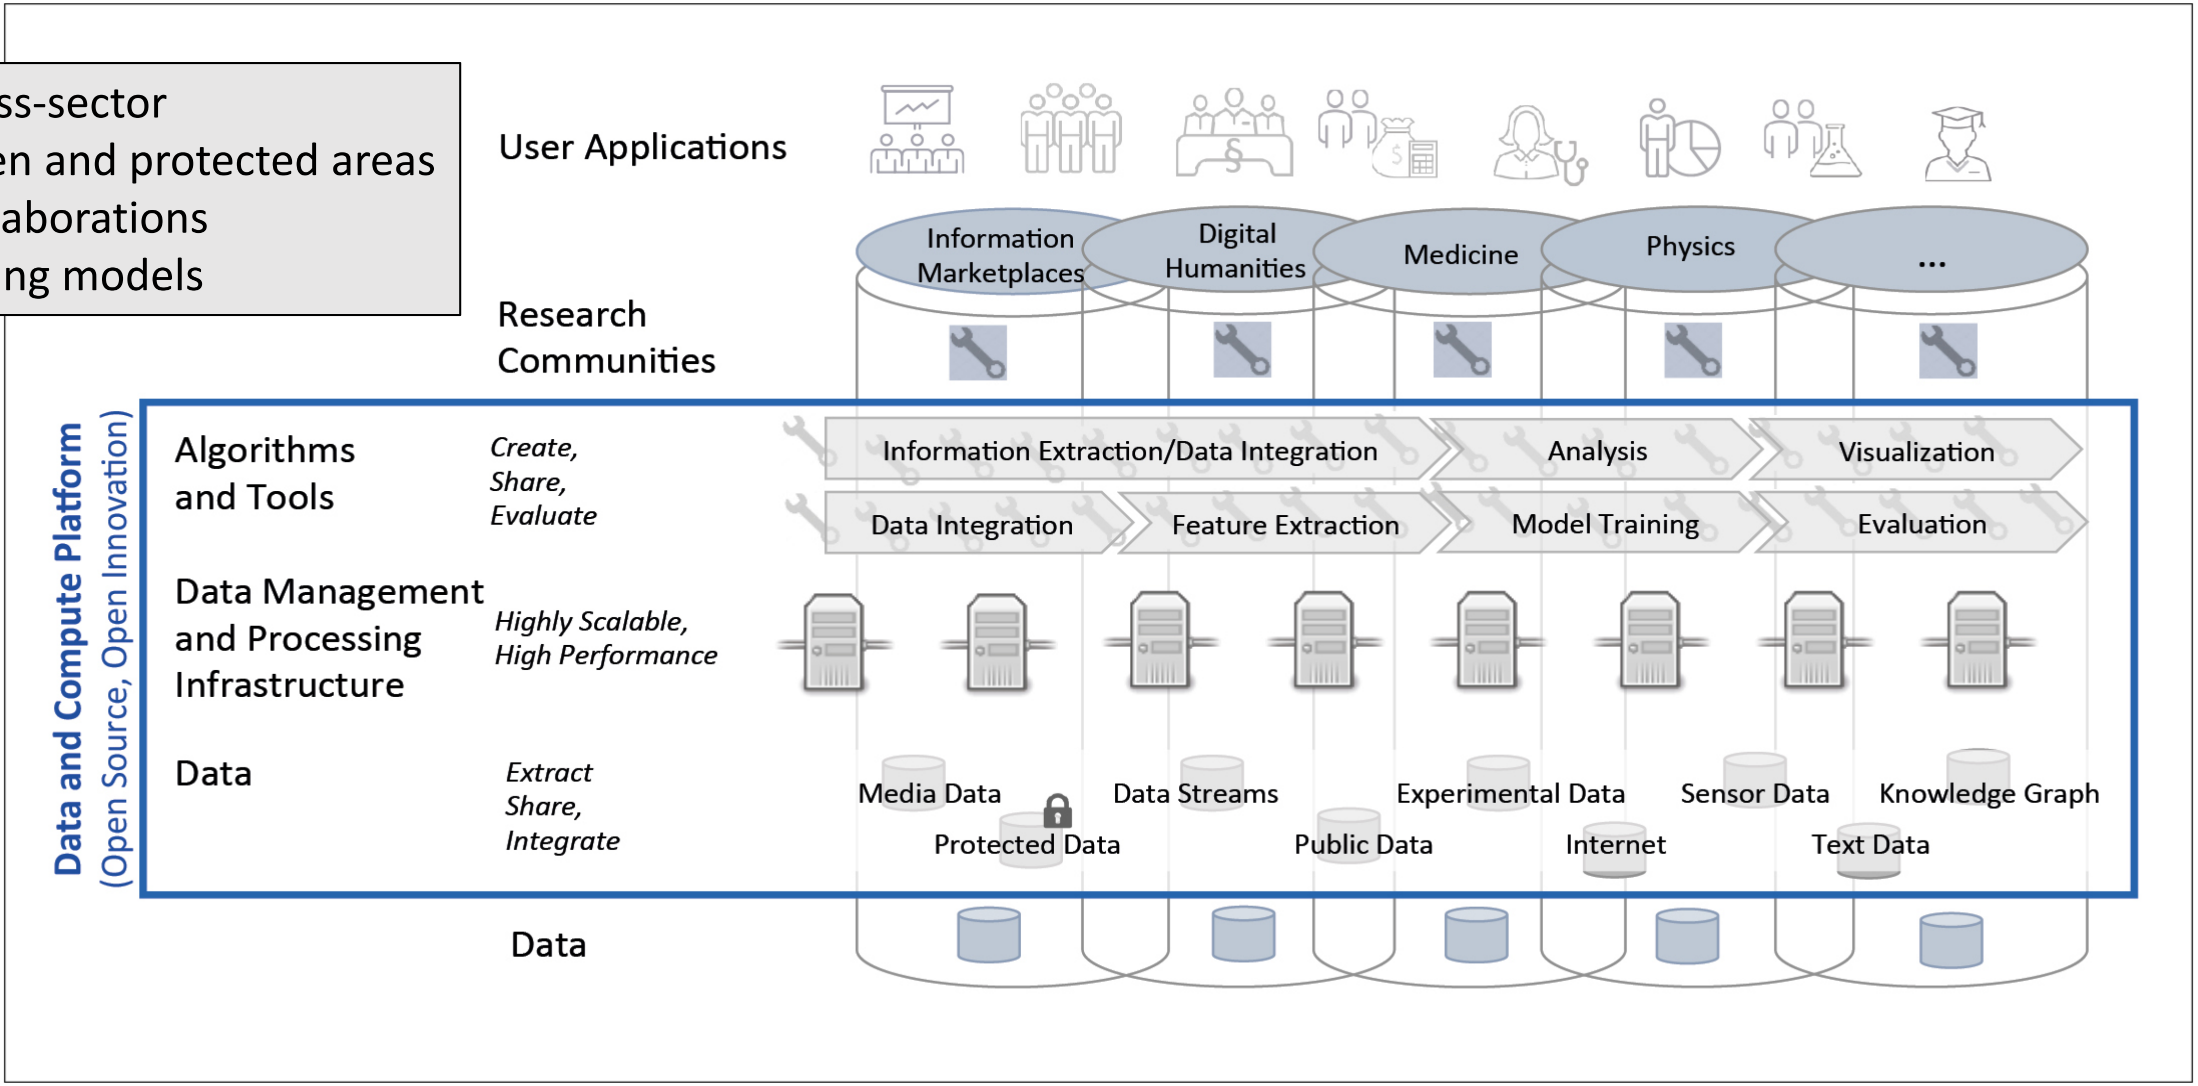

AGORA: One Data Management Infrastructure for AI-Applications
-- current research activities --

Multi-Platform Data Processing

Mediator-based systems:
- expensive data transfer
- additional system
- error-prone heuristics

Goal: Mediatorless processing
- orchestration of systems

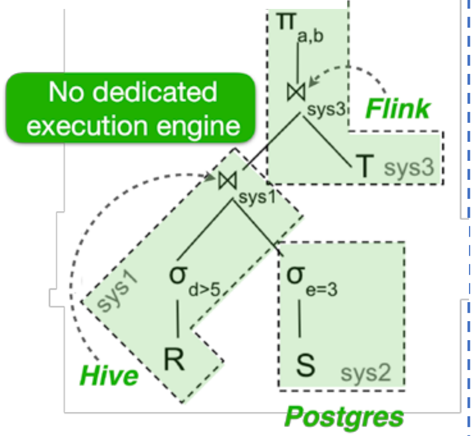

Compliant Distributed Processing

Data-flow constraints:
- transfer not permitted
- specification and enforcement of constraints

Goal: Automated derivation of compliant executions

Trustworthy Data Processing

Open environment:
- non-trusted servers
- inefficient processing

Goal: Efficient trustworthy processing through modern hardware.

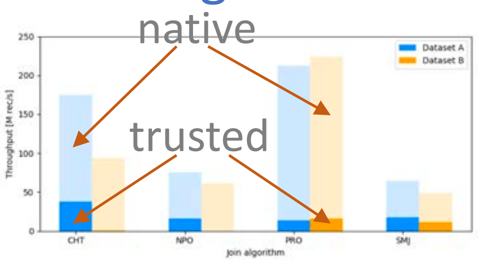

Performance gain through hardware-supported trusted processing

Debugging of Data Analysis Programs

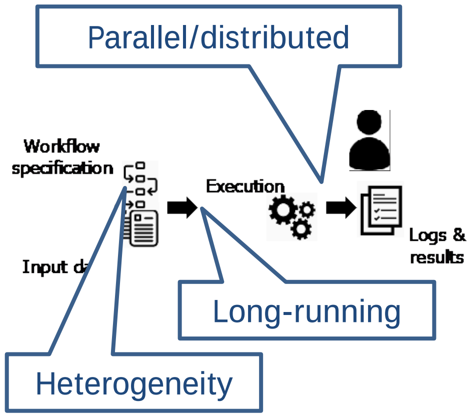

Online-error identification for big data streams:
- distributed processing
- continuous execution

Goal: Error correction by analyzing the processing history

Healthcare

Management and analysis of massive health data

Efficient integrated analysis of multimodal medical datasets (ECG, MRI, -"omics"-data)

Integration of image and time-series data

Processing of sensitive data

- "Life-Science-API" for Flink
- Demonstrator @ Forum Digitale Technologien
- Showcase at CeBIT

[1] Dictionary Learning. Information Technology 2020
[2] EMT network-based feature selection . PLoS ONE, 14 (1) 2019

(Tim Conrad, FU)

Digital Humanities

Efficient navigation through historical books

System for managing texts, illustrations and tables in historical books

Vast number of data sources

Text segmentation and semantic analysis

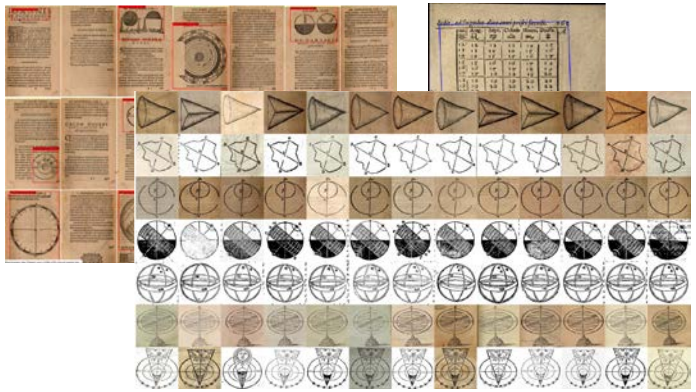

- New Paradigm for the historical sciences ("*longue-durée* "-studies with resolution of microstudies)
- Evolution of the role of archives and libraries

[1] Building and Interpreting Deep Similarity Models.
IEEE Trans. PAML, 2020

(Matteo Valeriani, MPG)

Environmental Data Analysis And "Citizen Science"

Deciphering air pollution data through data integration, explanation of outliers and visualization

1st Place
2019
GI FG DB Data Science Challenge

Analysis of the effects of traffic, driving restrictions, events, and weather on air quality

Heterogeneous data formats

Erroneous sensors

Number of data sources

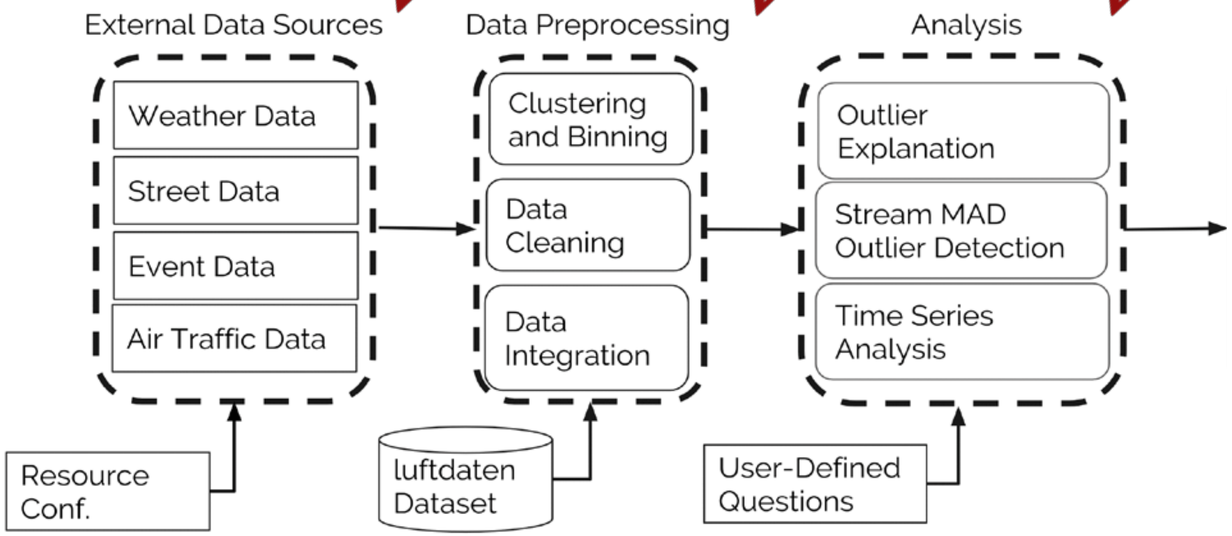

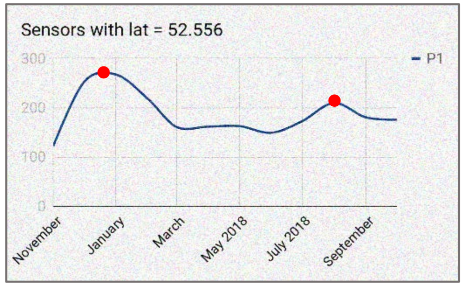

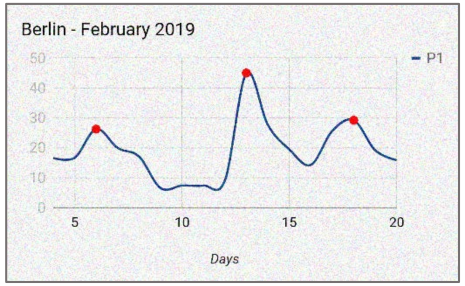

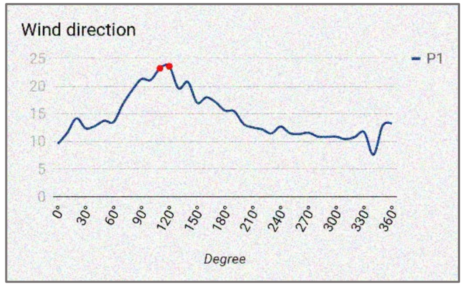

Selected results:

- Sensors near runways show peaks during vacation times.
- Events significantly impact air quality even without any traffic.
- Sensors near factories, stations, or highways show peaks depending on wind directions.

- Explanation of Air Pollution Using External Data Sources. M. Esmailoghli, S. Redyuk, R. Martinez, A. Ziehn, Z. Abedjan, T. Rabl, V. Markl, BTW 2019
- Particulate Matter Matters - The Data Science Challenge @ BTW'19. Meyer et al., Datenbank-Spektrum 2019

Conclusion

❶ Data are a new **factor of production** for sciences, humanities, and industry.
- Critical success factor is the reduction of **human latency** through intuitively usable systems.
- Reduction of the **technical latency** enables real-time analysis.

❷ Data management research is **interdisciplinary** within and outside computer science.

❸ Data management research has created important **technological foundations, new systems,** and **novel applications**:
- **Methods** for improving efficiency through automatic optimization, parallelization and hardware adaption.
- **Systems** for efficient, distributed, compliant analysis of large datasets and streams.
- **Applications** in the **Sciences and Humanities**, **Industry** and **Society** ("Citizen Science").

❹ **Current research challenges** in data management include:
- Methods and systems for the management and analysis of **distributed, heterogeneous data streams** ("IoT", "Industrie 4.0").
- Technologies for **data management infrastructures for open and protected, collaborative AI innovations** ("NFDI", "GAIA-X").

❺ The scientific community in Berlin is tackling these challenges in a comprehensive **ecosystem**:
- **Foundational research** (BIFOLD) and **applied research** (e.g., DFKI, Fraunhofer).
- **Applications** in mathematics/sciences, healthcare/biotech, humanities, as well as industry and startups.

Data – a New Factor of Production