FORSCHUNGSZENTRUM JÜLICH

JSC
Data Activities

26 October 2015 | Daniel Mallmann, Morris Riedel

RWTH Aachen Workshop Data Science: Theory and Application

Overview

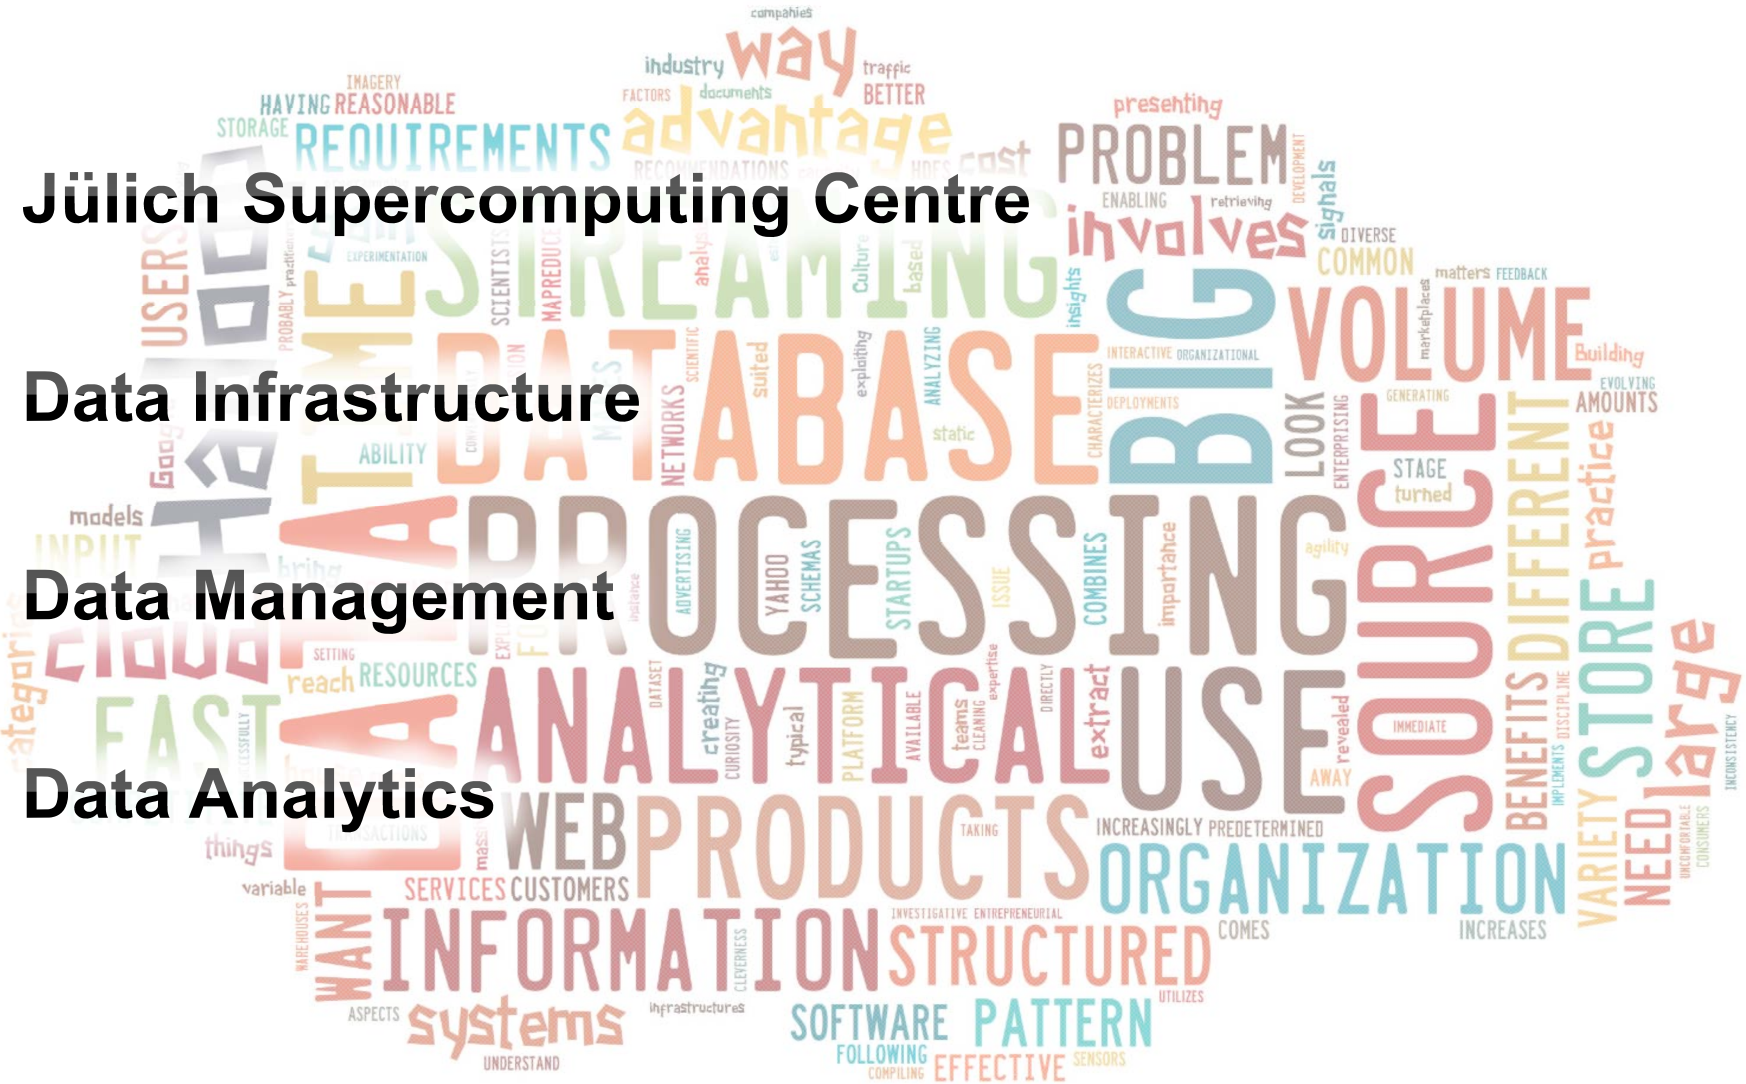

Jülich Supercomputing Centre

Data Infrastructure

Data Management

Data Analytics

Jülich Supercomputing Centre

Supercomputer operation for

- Centre
- Region
- Germany
- Europe (PRACE, EU projects)

Application support

- Unique support & research environment at JSC
- Peer review support and coordination

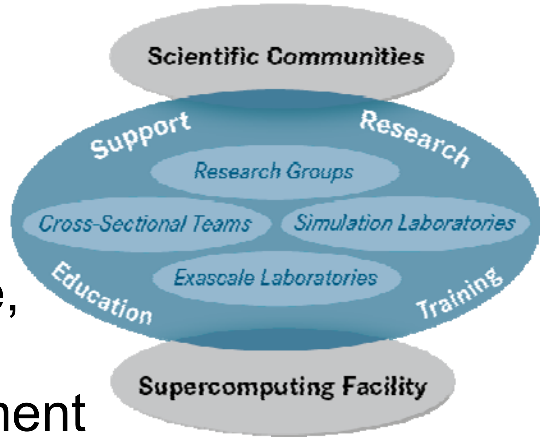

R&D work

- Methods and algorithms, computational science, performance analysis and tools
- Scientific Big Data Analytics and data management
- Computer architectures, Co-Design
  Exascale Laboratories: EIC, ECL, NVIDIA

Education and Training

DEEP

Jülich Supercomputing Centre
Domain-specific User Support and Research

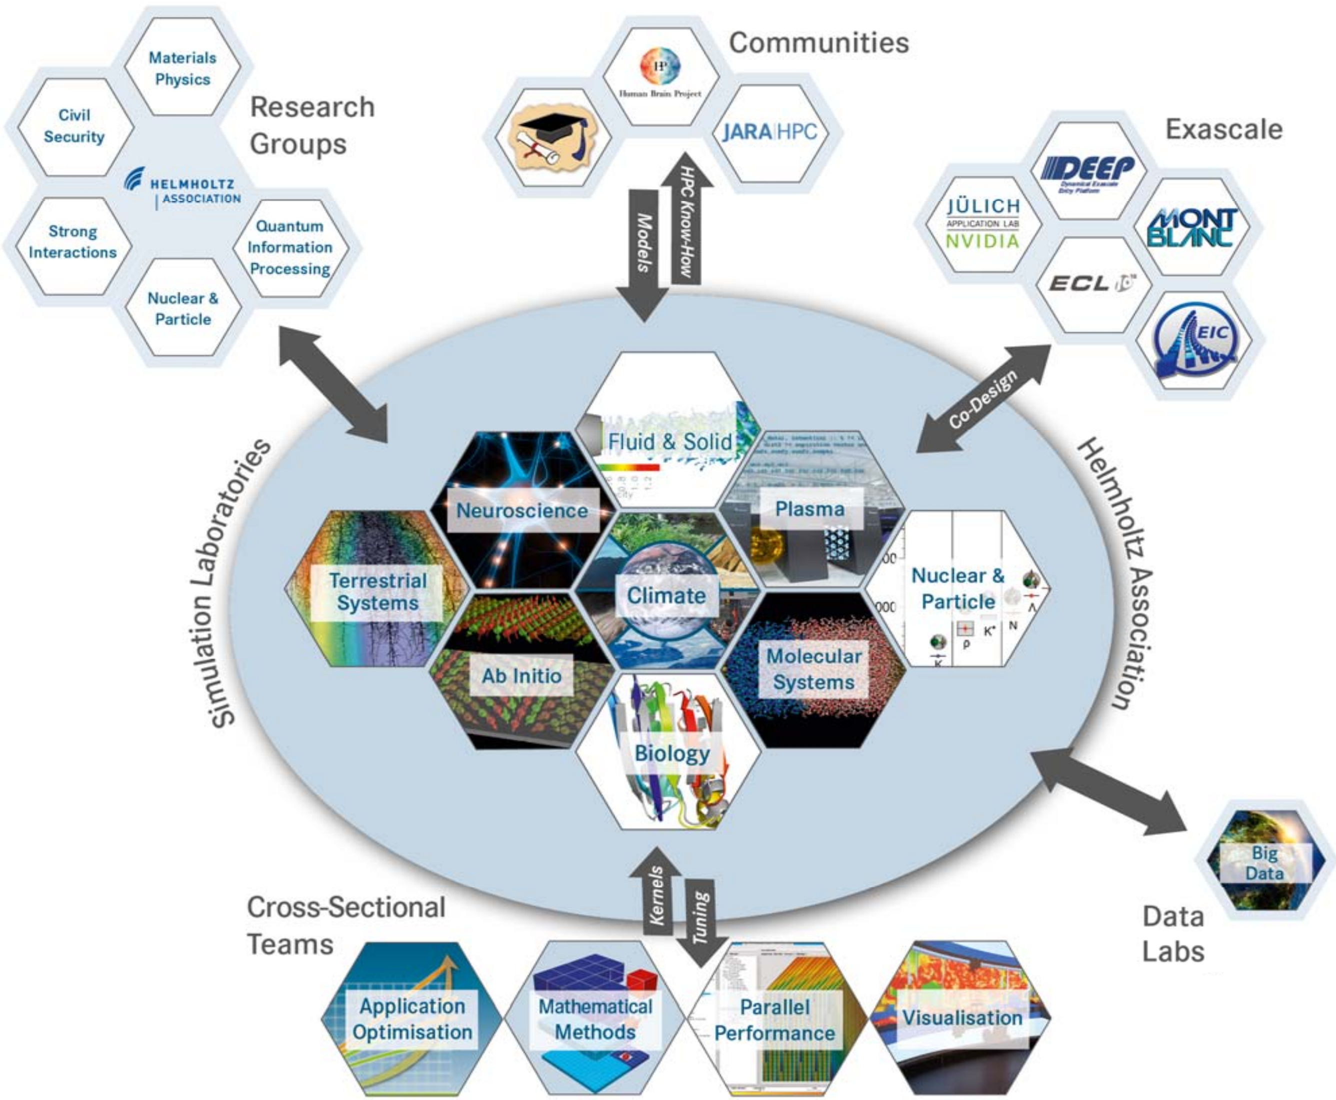

Jülich Supercomputing Centre
Supercomputer Systems

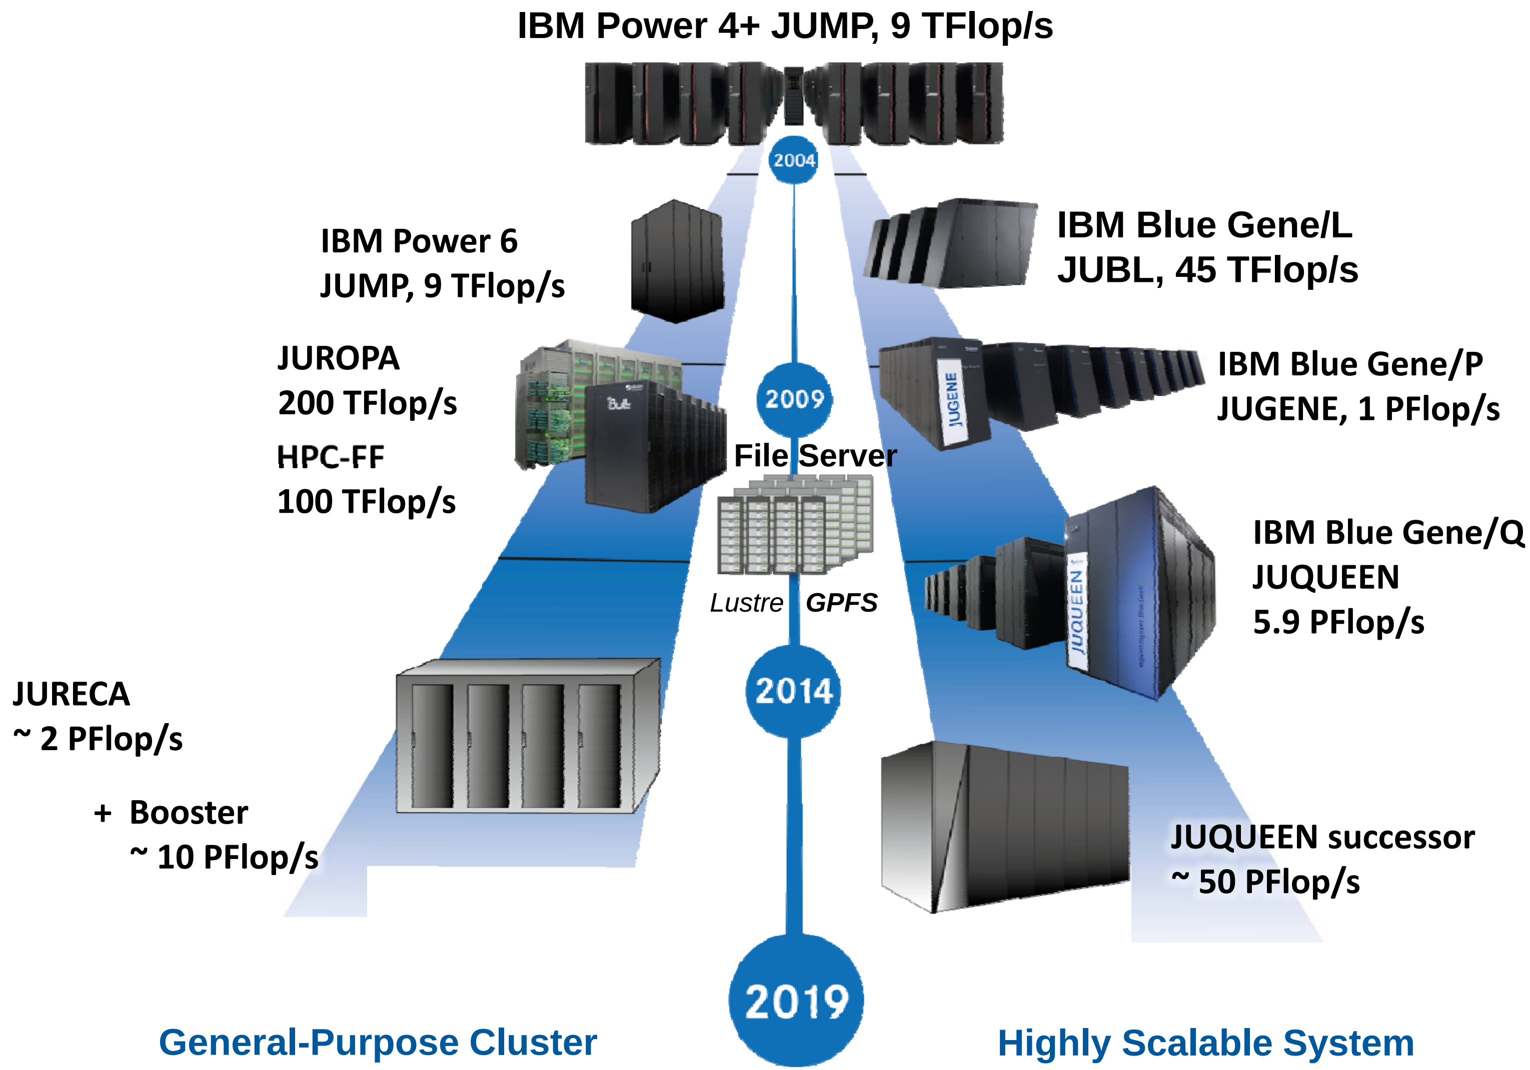

Data Infrastructure
JUST: Juelich Storage Cluster

- IBM-GPFS
  (General Parallel File System)
- 20.3 PB online storage
  (16.2 PB net)
- 22 Racks, 84 Server
- 8,300 disks, MTBF 3 disks per week
- RAID6, RAID5, RAID1
- Fileserver for
  - HPC systems: JUQUEEN, JURECA
  - Clusters: JUDGE, JUVIS (visualisation)
    DEEP (Dynamical Exascale Entry Platform)
  - Infrastructure collaborations

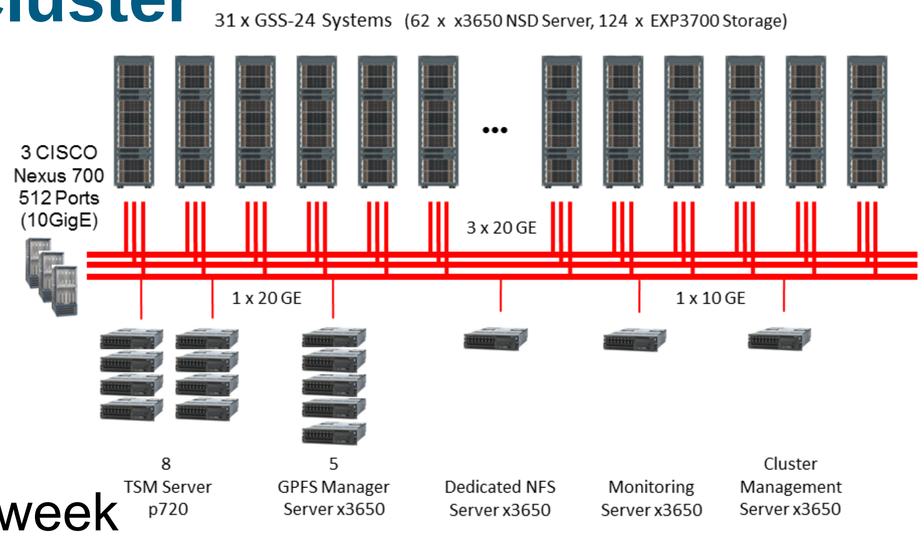

Data Infrastructure
Tape Libraries

- Automated cartridge systems
- 99 PB
- Used for
  - Backup
  - Long term archive
  - Migration of active (online) data to less expensive storage media
- 2 libraries (in 2 buildings)
  48 tape drives
- 16,600 tapes

Data Infrastructure
Distributed Storage

dCache

- 2.5 PB disk and tape
- User communities
  - ILDG for QCD on HPC systems
  - LOFAR for long term storage

iRODS

- 1 PB disk and tape
- Usage
  - EUDAT Safe Replication Service
  - EUDAT Data Staging Service

Data Management
EUDAT

Objective

- Build a cost-efficient high-quality Collaborative Data Infrastructure

Duration

- First phase: 1st October 2011 – 28th February 2015
- Second phase: 1st March 2015 – 28th February 2018

Budget

- First phase: 16.3 M€ (9.3 M€ European Commission)
- Second phase: 18.8 M€, 1900 Person Months

Consortium

- Second phase: 35 partners from 14 countries
- 15 thematic data centres, 14 generic data centers,
  3 technology providers, 2 libraries, 1 communication expert

EUDAT

JÜLICH
FORSCHUNGSZENTRUM

Data Management
EUDAT – Partner

CSC

UNINETT Sigma

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Science & Technology Facilities Council

CERFACS

UCL

rzg

CINECA

Trust-IT Services Ltd

Jisc

grnet

CINES

KIT

CERN

DKRZ

LIBER

umweltbundesamt

DANS

EBERHARD KARLS UNIVERSITÄT TÜBINGEN

e-Science Data Factory

EMBL

epcc

SNIC

INGV

Royal Netherlands Meteorological Institute
Ministry of Infrastructure and the Environment

KTH

JÜLICH

GFZ

Max-Planck-Institut für Meteorologie

POZNAN

SURF SARA

CLARIN

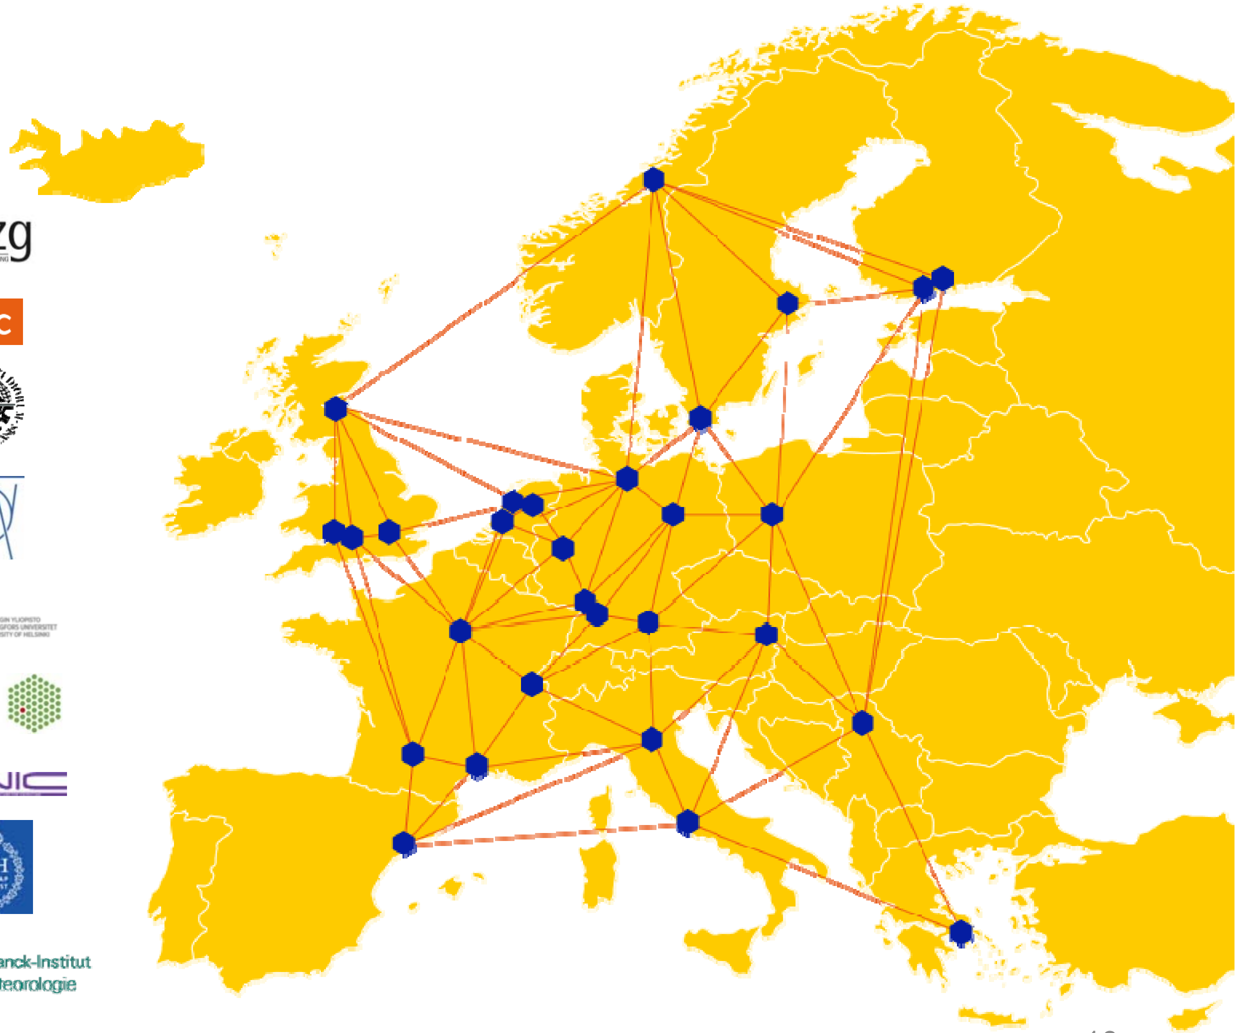

Data Management
EUDAT Collaborative Data Infrastructure

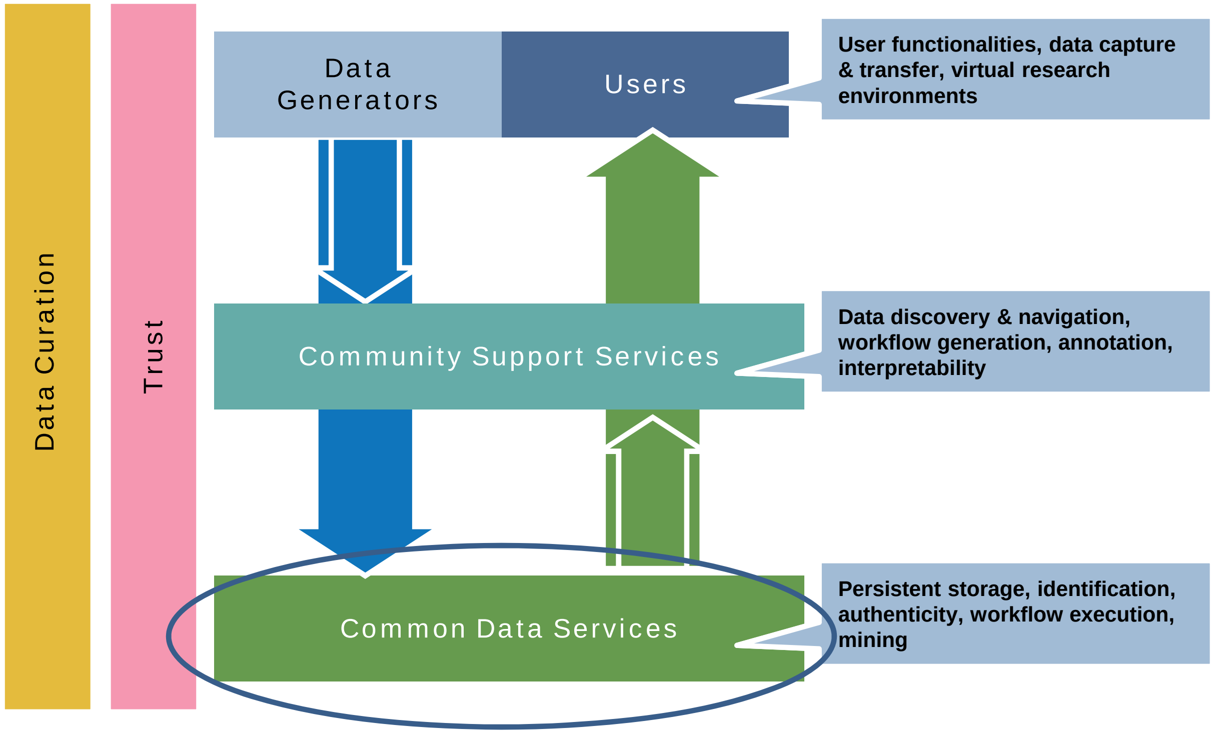

Data Management
EUDAT Services

B2DROP – Sync and Exchange Research Data

B2SHARE – Store and Share Research Data

B2SAFE – Replicate Research Data Safely

B2STAGE – Get Data to Computation

B2FIND – Find Research Data

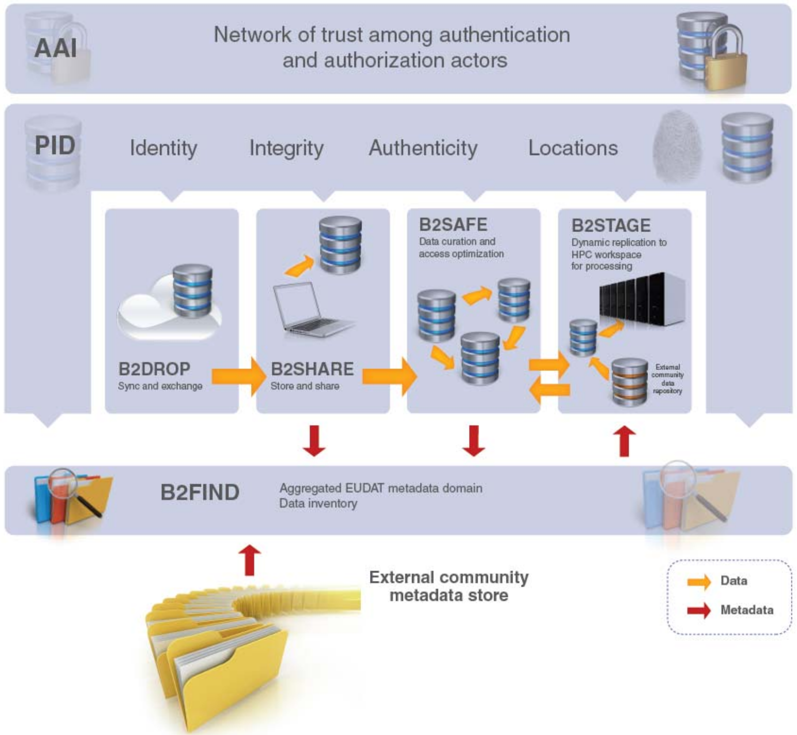

http://www.eudat.eu/services

UNICORE

JÜLICH FORSCHUNGSZENTRUM

Data Management UNICORE

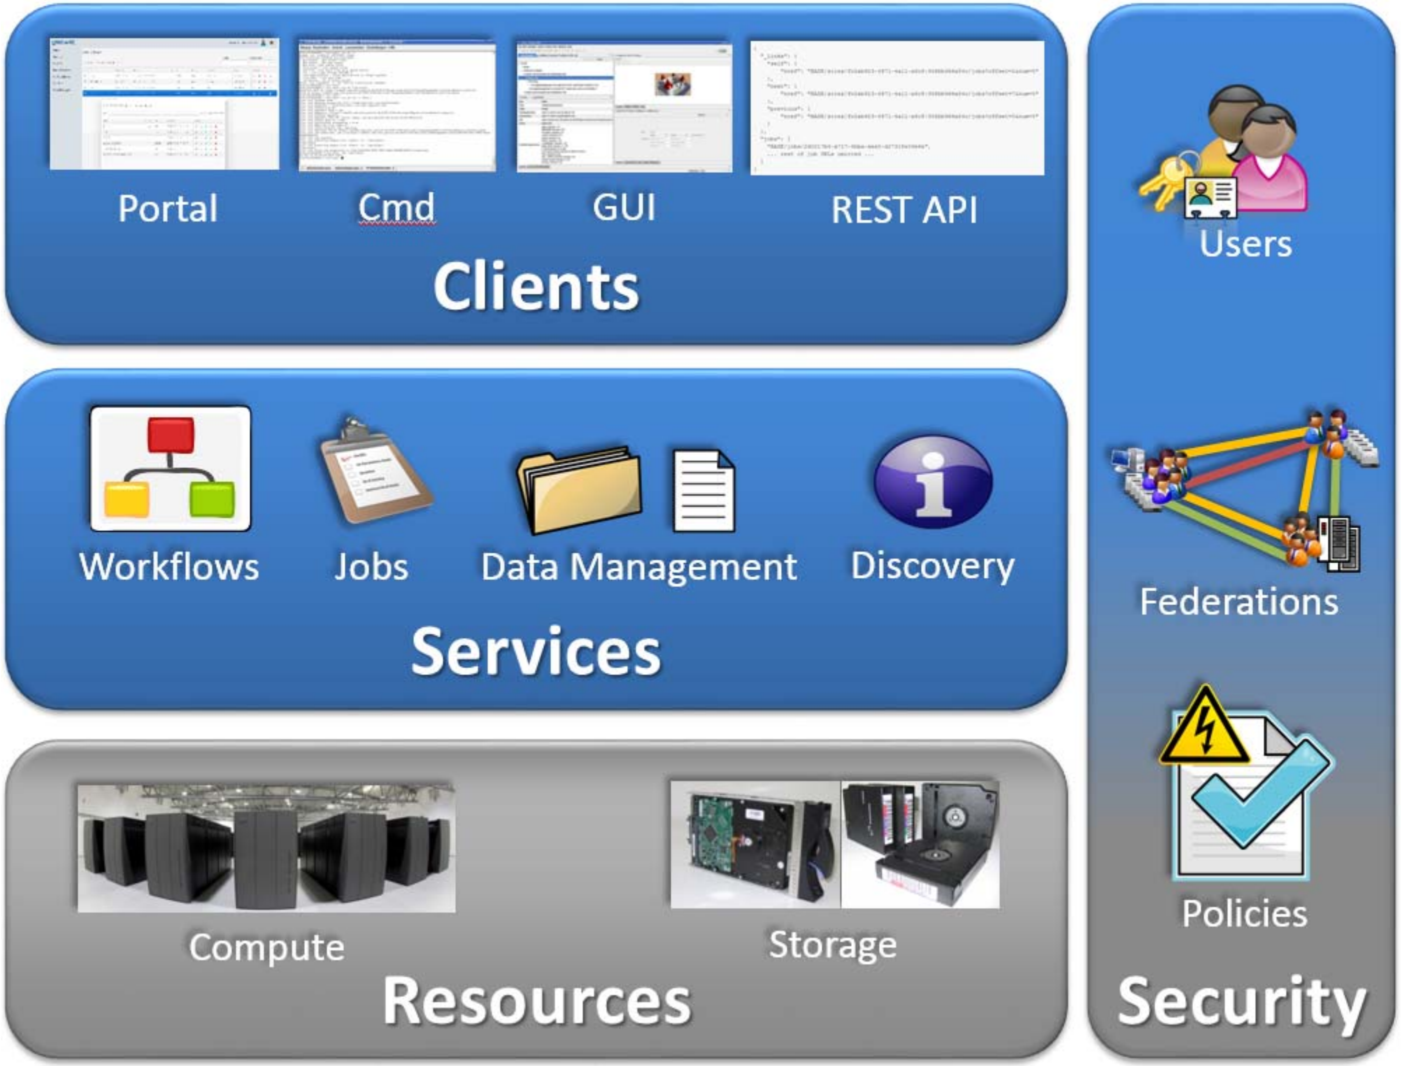

http://unicore.eu/

Data Management UFTP

High-performance file transfer

- Firewall-friendly: only a single open port is required
- Secure through use of out-of-band authentication
- Options: rsync, data compression and encryption
- Fully integrated with UNICORE and useable standalone

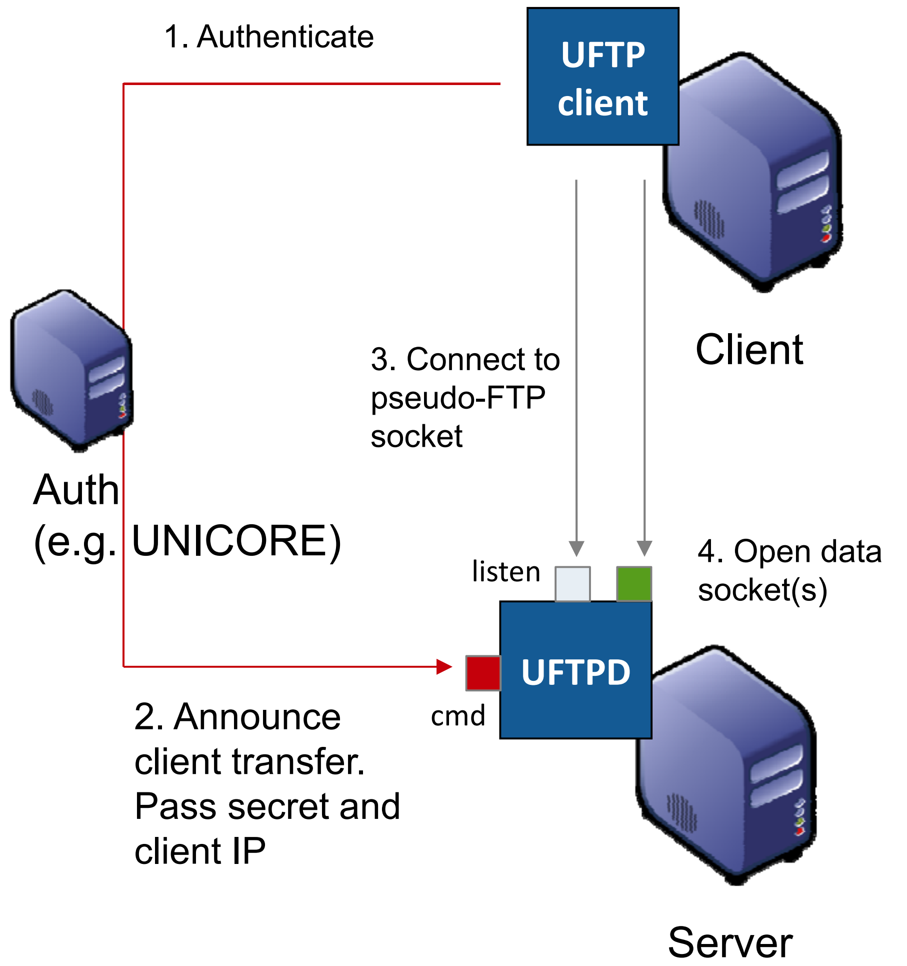

http://sourceforge.net/projects/unicore/files/Clients/UFTP-Client/

Data Analysis
Domain-specific User Support and Research

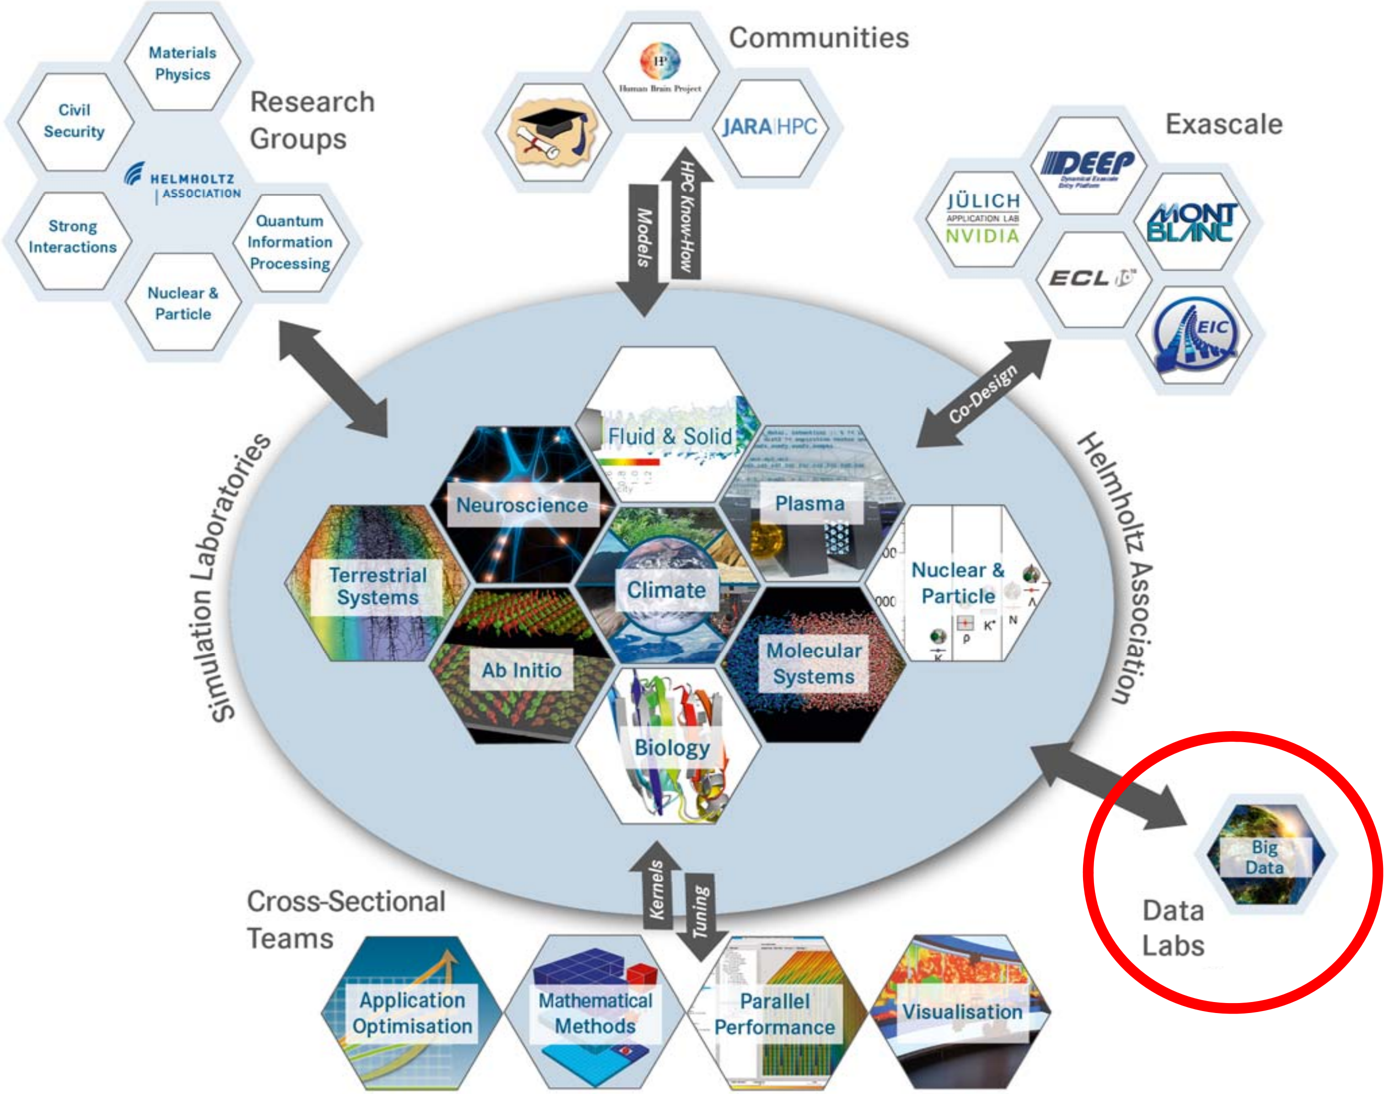

Data Analysis

Classification

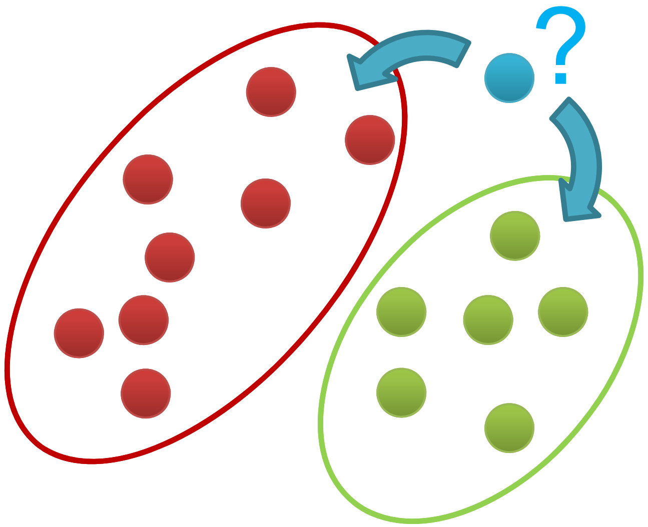

Groups of data exist
New data classified
to existing groups

Clustering

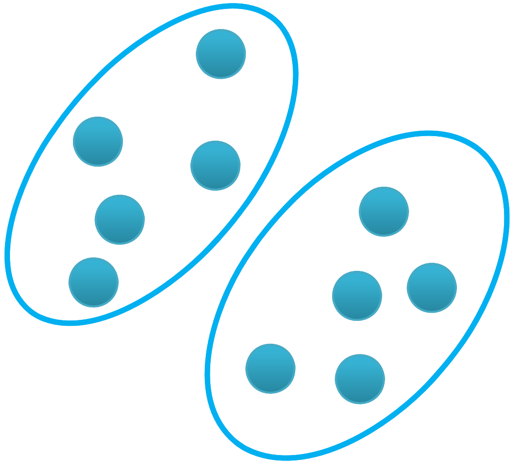

No groups of data exist
Create groups from
data close to each other

Regression

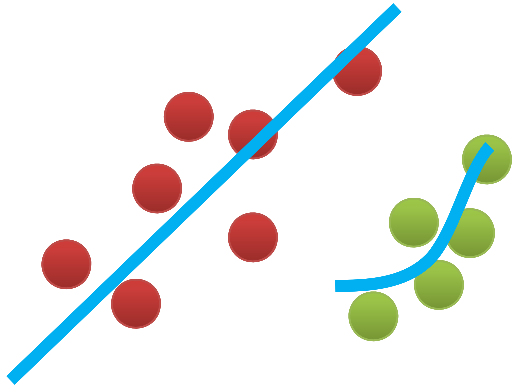

Identify a line with
a certain slope
describing the data

Data Analysis – Classification
Remote Sensing Images

Challenges: high number of classes, less samples, mixed pixels

Class		Number of samples		Class		Number of samples	
Number	Name	Training	Test	Number	Name	Training	Test
1	Buildings	1720	15475	27	Pasture	1039	9347
2	Corn	1778	16005	28	pond	10	92
3	Corn?	16	142	29	Soybeans	939	8452
4	Corn-EW	51	463	30	Soybeans?	89	805
5	Corn-NS	236	2120	31	Soybeans-NS	111	999
6	Corn-CleanTill	1240	11164	32	Soybeans-CleanTill	507	4567
7	Corn-CleanTill-EW	2649	23837	33	Soybeans-CleanTill?	273	2453
8	Corn-CleanTill-NS	3968	35710	34	Soybeans-CleanTill-EW	1180	10622
9	Corn-CleanTill-NS-Irrigated	80	720	35	Soybeans-CleanTill-NS	1039	9348
10	Corn-CleanTilled-NS?	173	1555	36	Soybeans-CleanTill-Drilled	224	2018
11	Corn-MinTill	105	944	37	Soybeans-CleanTill-Weedy	54	489
12	Corn-MinTill-EW	563	5066	38	Soybeans-Drilled	1512	13606
13	Corn-MinTill-NS	886	7976	39	Soybeans-MinTill	267	2400
14	Corn-NoTill	438	3943	40	Soybeans-MinTill-EW	183	1649
15	Corn-NoTill-EW	121	1085	41	Soybeans-MinTill-Drilled	810	7288
16	Corn-NoTill-NS	569	5116	42	Soybeans-MinTill-NS	495	4458
17	Fescue	11	103	43	Soybeans-NoTill	216	1941
18	Grass	115	1032	44	Soybeans-NoTill-EW	253	2280
19	Grass/Trees	233	2098	45	Soybeans-NoTill-NS	93	836
20	Hay	113	1015	46	Soybeans-NoTill-Drilled	873	7858
21	Hay?	219	1966	47	Swampy Area	58	525
22	Hay-Alfalfa	226	2032	48	River	311	2799
23	Lake	22	202	49	Trees?	58	522
24	NotCropped	194	1746	50	Wheat	498	4481
25	Oats	174	1568	51	Woods	6356	57206
26	Oats?	34	301	52	Woods?	14	130

remote sensing cube & ground reference

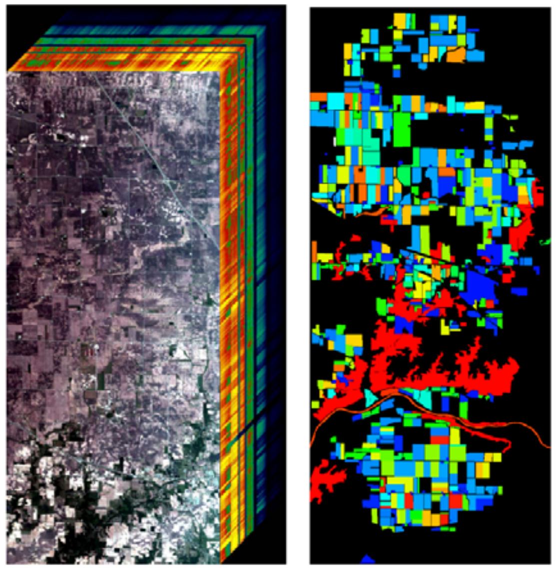

(1) Scenario ‘unprocessed data‘
(2) Scenario ‘preprocessed data‘

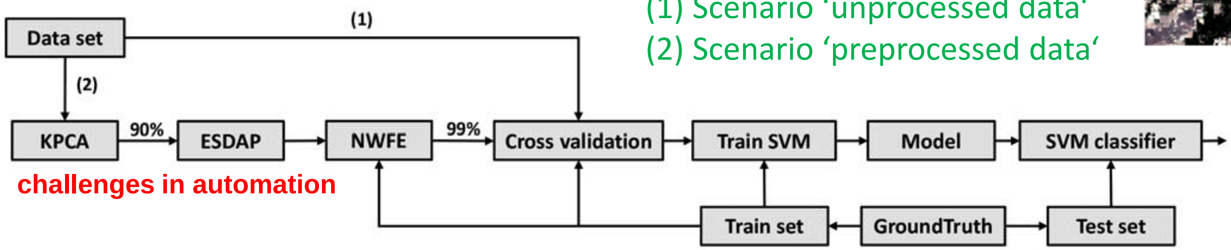

challenges in automation

[1]

Data Analysis – Classification Methods

Perceptron Learning Algorithm – simple linear classification

- Enables binary classification with ‘a line‘ between classes of seperable data

Support Vector Machines (SVMs) – non-linear (‘kernel‘) classification

- Enables non-linear classification with maximum margin (best ‘out-of-the-box‘)

Reasoning: achieves often better results than other methods in tackled application domain

Decision Trees & Ensemble Methods – tree-based classification

- Grows trees for class decisions, ensemble methods average n trees

Artificial Neural Networks (ANNs) – brain-inspired classification

- Combine multiple linear perceptrons to a strong network for non-linear tasks

Naive Bayes Classifier – probabilistic classification

- Use of the Bayes theorem with strong/naive independence between features

Data Analysis – Classification
SVM Algorithm

- Introduced 1995 by C.Cortes & V. Vapnik et al. *[2]*
- Creates a ‘maximal margin classifier‘ to get future points (‘more often‘) right and take advantage of kernel methods
- Uses quadratic programming & Lagrangian method with N x N

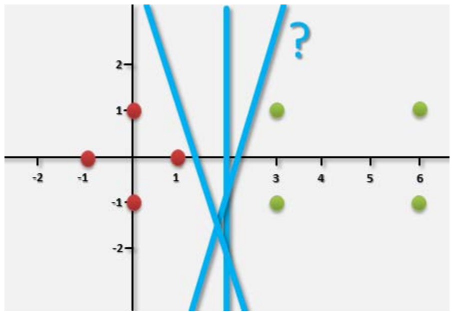

linear

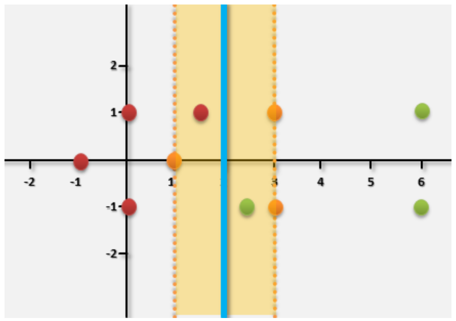

maximal margin clasifier

Data Analysis – Classification
SVMs – Review of Available Technologies

Technology	Platform Approach	Analysis
Apache Mahout	Java; Hadoop	No parallelization strategy for SVMs
Apache Spark/MLlib	Java; Spark	Parallel linear SVMs (no multi-class)
Twister/ParallelSVM	Java; Twister; Hadoop 1.0	Parallel SVMs, open source; developer version 0.9 beta
scikit-learn	Python	No parallelization strategy for SVMs
piSVM 1.2 & piSVM 1.3	C; MPI	Parallel SVMs; stable; not fully scalable
GPU LibSVM	CUDA	Parallel SVMs; hard to programs, early versions
pSVM	C; MPI	Parallel SVMs; unstable; beta version

[3]

Data Analysis – Classification
Parallel & Scalable piSVM Open Source Tool

Original parallel piSVM tool 1.2 [4]

- Open-source and based on libSVM library, C, 2011
- Message Passing Interface (MPI)
- New version 1.3 published 2014-10
- Lack of ‘big data‘ support (memory, layout, etc.)

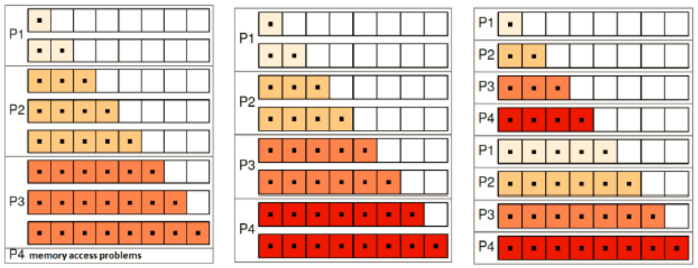

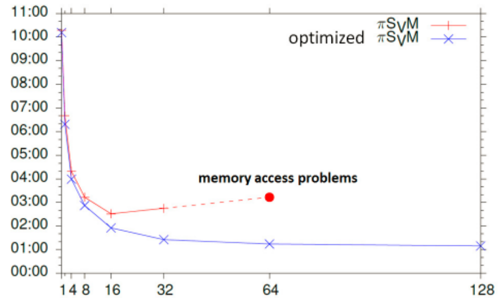

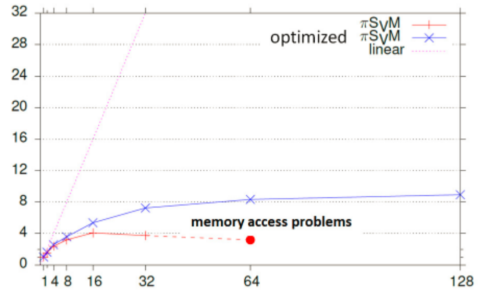

Tuned scalable parallel piSVM tool 1.2.1

- Highly scalable version maintained by Juelich
- Based on original piSVM 1.2 tool
- Open-source (repository to be created)
- Optimizations: load balancing, MPI collectives

Data Analysis

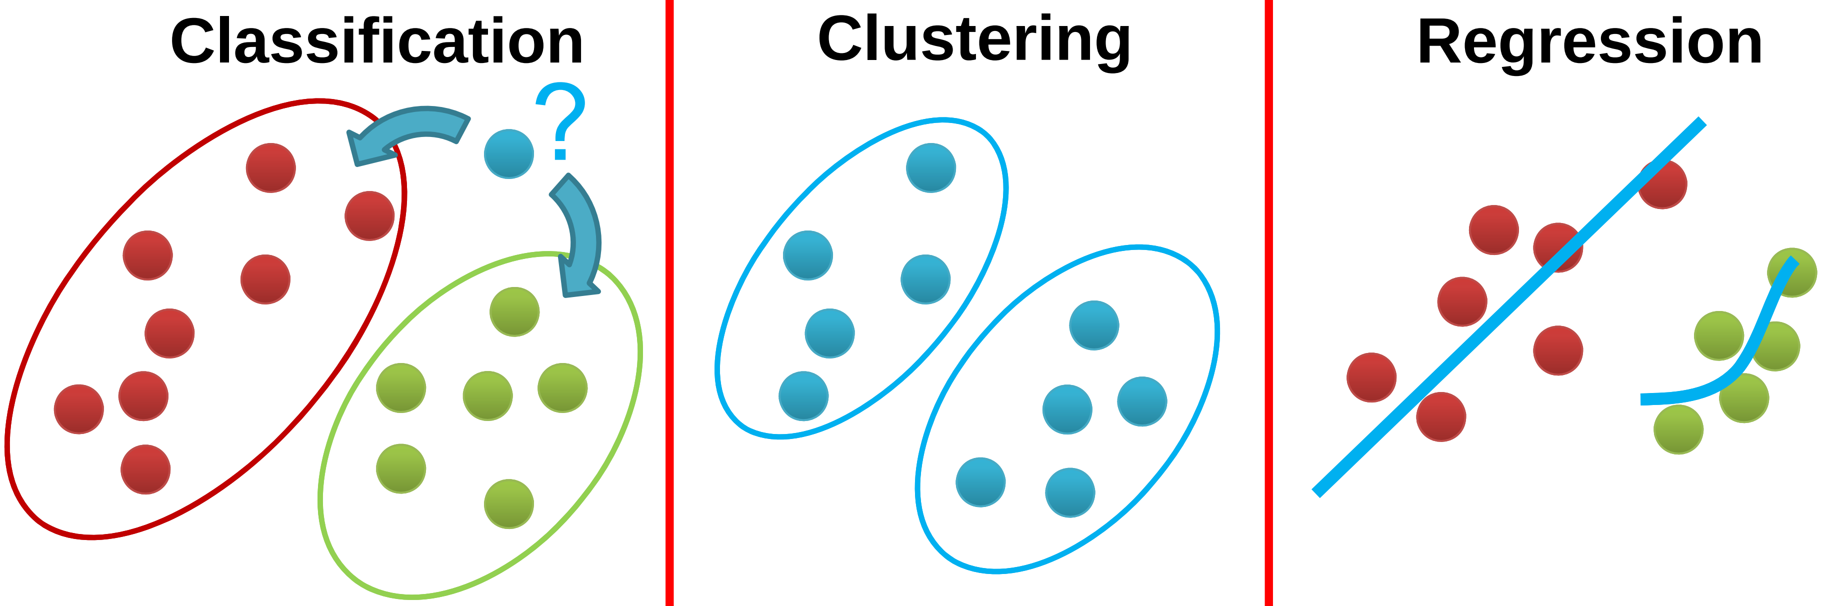

Groups of data exist
New data classified
to existing groups

No groups of data exist
Create groups from
data close to each other

Identify a line with
a certain slope
describing the data

Data Analysis – Clustering Large Point Clouds

‘Big Data‘: 3D/4D laser scans

- Captured by robots or drones
- Millions to billion entries
- Inner cities (e.g. Bremen inner city)
- Whole countries (e.g. Netherlands)

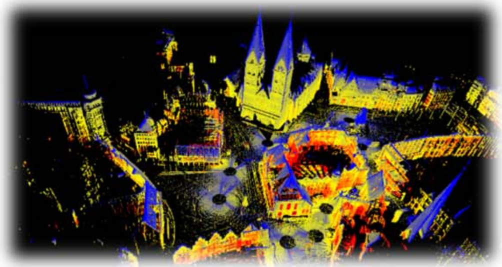

Selected Scientific Cases

- Filter noise to better represent real data
- Grouping of objects (e.g. buildings)

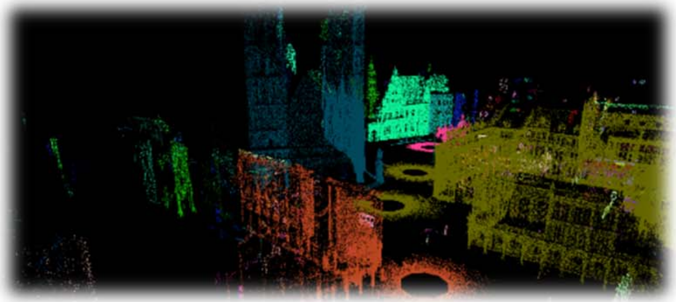

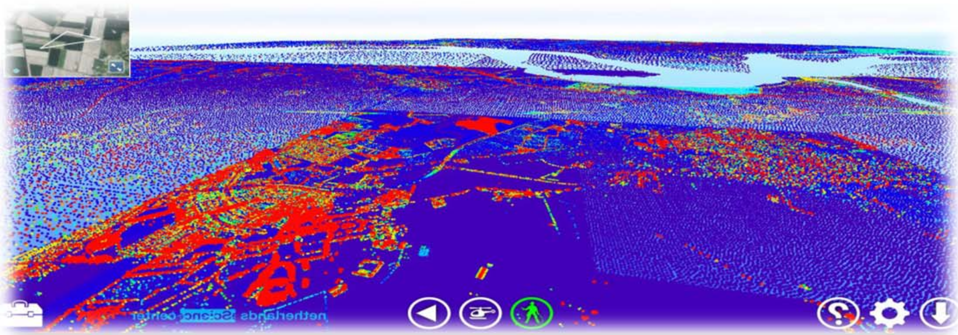

Data Analysis – Clustering Methods

K-Means Clustering – Centroid based clustering
- Partitions a data set into K distinct clusters (centroids can be artificial)

K-Medoids Clustering – Centroid based clustering (variation)
- Partitions a data set into K distinct clusters (centroids are actual points)

Sequential Agglomerative hierarchic nonoverlapping (SAHN)
- Hiearchical Clustering (create tree-like data structure → 'dendrogram')

Clustering Using Representatives (CURE)
- Select representative points / cluster; as far from one another as possible

Density-based spatial clustering of applications + noise (DBSCAN)
- Assumes clusters of similar density or areas of higher density in dataset

Reasoning: density similiarity measure helpful in our driving applications

Data Analysis – Clustering DBSCAN

DBSCAN Algorithm [5]

- Introduced 1996 by Martin Ester et al.
- Groups number of similar points into clusters of data
- Similarity is defined by a distance measure (e.g. *euclidean distance*)

Distinct Algorithm Features

- Clusters a variable number of clusters
- Forms arbitrarily shaped clusters
- Identifies outliers/noise

Understanding Parameters for MPI/OpenMP tool [6]

- Looks for similar points within a given search radius → Parameter *epsilon*
- A cluster consists of given minimum number of points → Parameter *minPoints*

Unclustered Data

Clustered Data

Data Analysis – Clustering
DBSCAN – Review of Available Technologies

Technology	Platform Approach	Analysis
HPDBSCAN (authors implementation)	C; MPI; OpenMP	Parallel, hybrid, DBSCAN
Apache Mahout	Java; Hadoop	K-means variants, spectral, no DBSCAN
Apache Spark/MLlib	Java; Spark	Only k-means clustering, No DBSCAN
scikit-learn	Python	No parallelization strategy for DBSCAN
Northwestern University PDSDBSCAN-D	C++; MPI; OpenMP	Parallel DBSCAN

[3]

Data Analysis – Clustering
Parallel & Scalable HP-DBSCAN Open Source Tool

Parallelization Strategy [7]

- Smart ‘Big Data‘ Preprocessing into Spatial Cells (‘indexed‘)
- OpenMP standalone
- MPI (+ optional OpenMP hybrid)

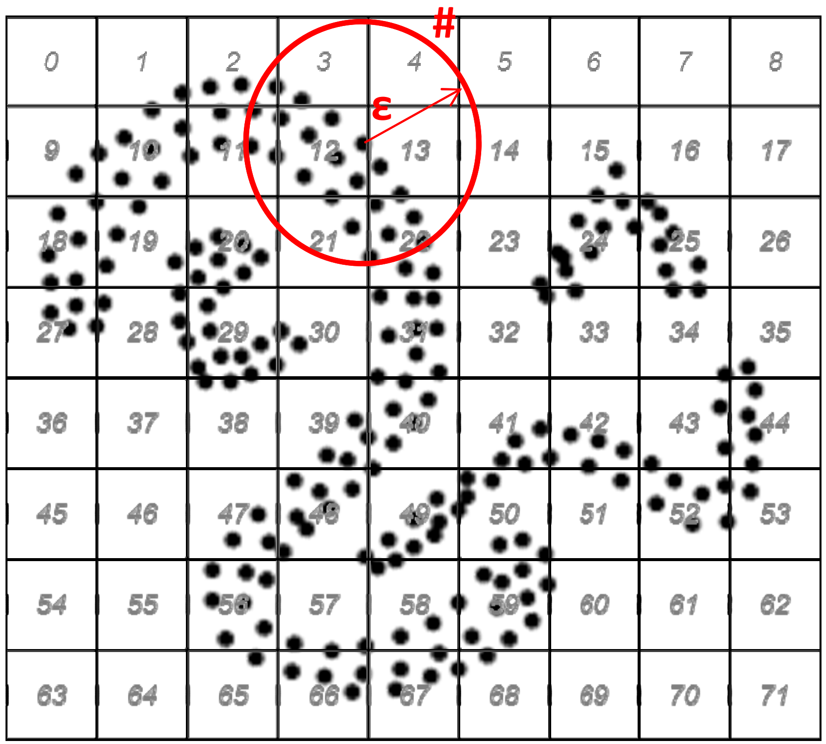

Preprocessing Step

- Spatial indexing and redistribution according to the point localities
- Data density based chunking of computations

Computational Optimizations

- Caching of point neighborhood searches
- Cluster merging based on comparisons instead of zone reclustering

Data Analysis – Clustering
Many Time Series Events

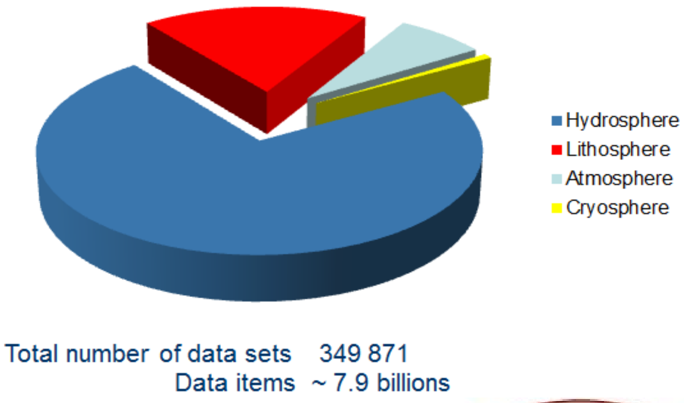

Earth Science Data Repository

- Time series measurements (e.g. salinity)
- Millions to billions of data items/locations
- Less capacity of experts to analyse data

Selected Scientific Case

- Data from Koljöfjords in Sweden (Skagerrak)
- Each measurement small data, but whole sets are ‘big data‘
- Automated water mixing event detection & quality control (e.g. biofouling)
- Verification through domain experts

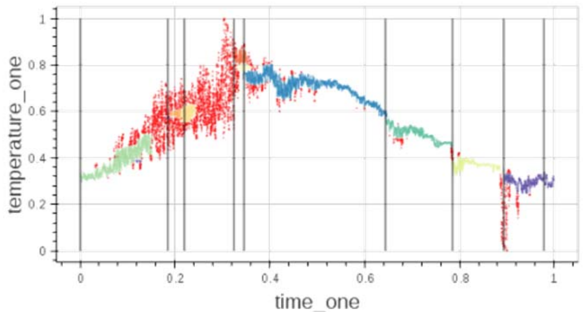

Mitglied der Helmholtz-Gemeinschaft

Data Analysis – Clustering Neuroscience Images

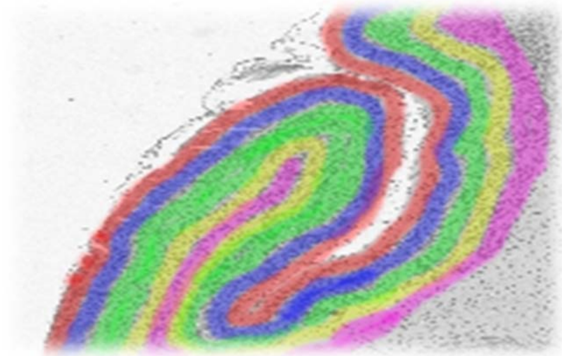

Large Brain Images

- High resolution scans of post mortem brains
- Rare ‘groundtruth available‘

Selected Scientific Case

- Cell nuclei detection and tissue clustering
- Detect various layers (colored)
- Layers seem to have different density distribution of cells
- Extract cell nuclei into 2D/3D point cloud
- Cluster different brain areas by cell density

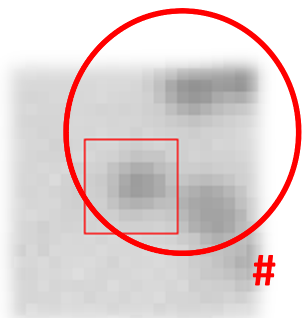

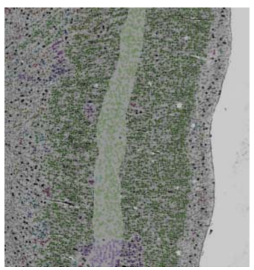

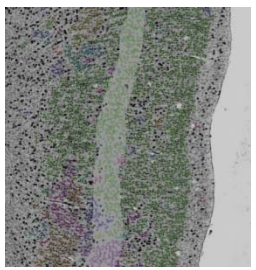

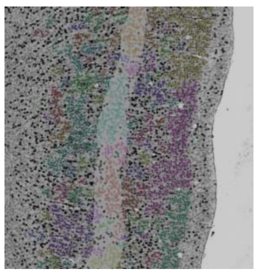

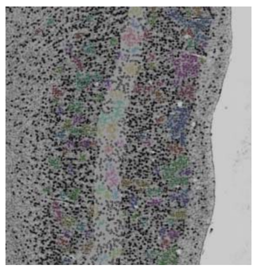

Summary

Data Management

- Services need to be customized for users, basic services could be generic
- JSC operates a central data system, that is part of data federations

EUDAT

- Transition from project to infrastructure
- Services are mature and production-ready

Data Analytics

- Focus on enabling and/or advancing data analytics in selected use cases through optimization and scaling of methods and tools

JÜLICH FORSCHUNGSZENTRUM

References

[1] G. Cavallaro, M. Riedel, J.A. Benediktsson et al., 'On Understanding Big Data Impacts in Remotely Sensed Image Classification using Support Vector Machine Methods', IEEE Journal of Selected Topics in Applied Earth Observation and Remote Sensing, 2015, DOI: 10.1109/JSTARS.2015.2458855

[2] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20(3), pp. 273–297, 1995

[3] M. Goetz, M. Riedel et al.,' On Parallel and Scalable Classification and Clustering Techniques for Earth Science Datasets' 6th Workshop on Data Mining in Earth System Science, Proceedings of the International Conference of Computational Science (ICCS), Reykjavik, Online: http://www.proceedings.com/26605.html

[4] Original piSVM tool, online: http://pisvm.sourceforge.net/

[5] M. Ester, et al. "A density-based algorithm for discoveringclusters in large spatial databases with noise." Kdd. Vol. 96. 1996.

[6] M. Goetz & C. Bodenstein, Clustering Highly Parallelizable DBSCAN Algorithm, JSC, Online: http://www.fz-juelich.de/ias/jsc/EN/Research/DistributedComputing/DataAnalytics/Clustering/Clustering_node.html

[7] M. Goetz, M. Riedel et al.,'HPDBSCAN – Highly Parallel DBSCAN', accepted for MLHPC Workshop at Supercomputing 2015, Online: http://www.wikicfp.com/cfp/servlet/event.showcfp?eventid=46948