Service Development for Machine Learning for Earth System Science

France-Japan-Germany Trilateral Workshop

Convergence of HPC and Data Science for Future Extreme Scale Intelligent Applications

Prof. Dr. Thomas Ludwig
Deutsches Klimarechenzentrum (DKRZ)
Universität Hamburg, Fachbereich Informatik (UHH/FBI)

German Climate Computing Centre

Mission

DKRZ – Partner for Climate Science.

Maximum Compute Performance.
Sophisticated Data Management.
Competent Service.

Vision

DKRZ reliably unlocks the potential of the accelerating technological progress for climate research

German Climate Computing Centre

Non-profit limited company since 1987

- Share-holders Max-Planck-Gesellschaft, City of Hamburg, Alfred-Wegener-Institute, Helmholtz Center Geesthacht
- 80+ staff at DKRZ
- 10+ staff at university research group (focus on parallel I/O)

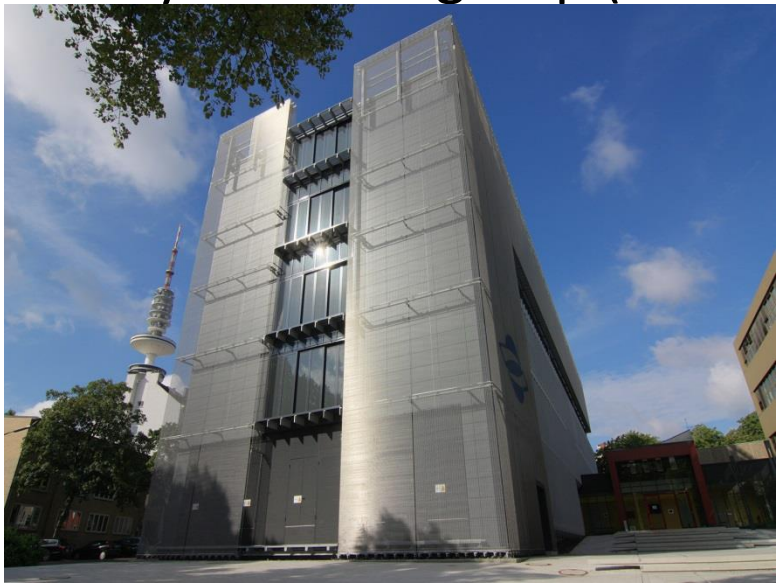

Governance

- Financial contribution
  - Currently
    - Infrastructure by Helmholtz Association, Max-Planck-Gesellschaft, and City of Hamburg (€45M/5a)
    - Operational costs by share-holders (€10M/a)
- Resource distribution
  - 50% to share-holders according to their share
  - 50% to German research community
    - Assigned by Scientific Steering Committee

30 Years of Computer and Storage Systems (1987-2019)

1985: Control Data Cyber-205

- 1 processor
- 0.2 Gigaflops
- 0.03 Gigabyte memory

1988: Cray 2S

- 4 processors
- 2 Gigaflops
- 1 Gigabyte memory

1994: Cray C-916 „Sea“

- 16 processors
- 16 Gigaflops
- 2 Gigabyte memory
- 128 Gigabyte disc space
- 10 Terabyte tape archive

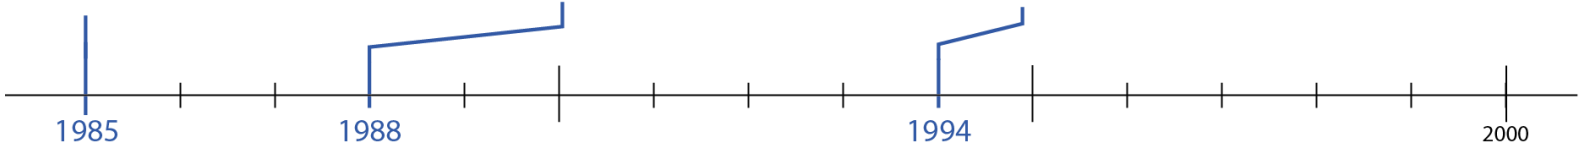

2002: NEC SX-6 „Hurrikan“

- 192 processors
- 1.5 Teraflops
- 1.5 Terabyte memory
- 60 Terabyte disc space
- 3.4 Petabyte tape archive

2009: IBM Power6 „Blizzard“

- 8500 processors
- 158 Teraflops
- 20 Terabyte memory
- 6 Petabyte disc space
- 60 Petabyte tape archive

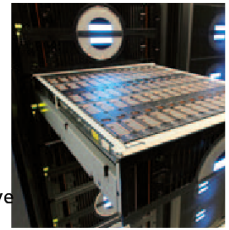

2015/16: bullx B700 DLC „Mistral“

- **>100,000 processors**
- **3.6 Petaflops**
- **240 Terabyte memory**
- **54 Petabyte disk space**
- **up to 500 Petabyte tape archive**

1 PetaByte are used

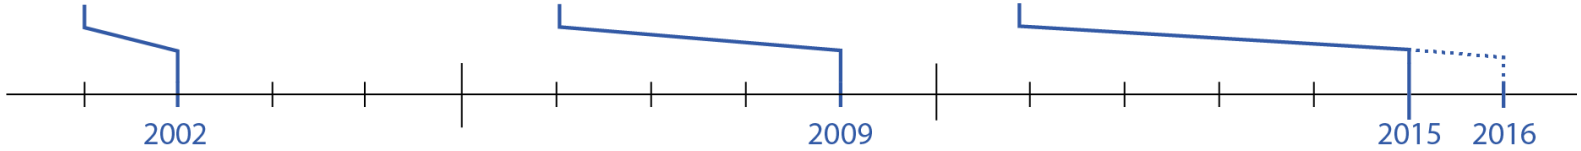

Powerful Computers for Competitive Research

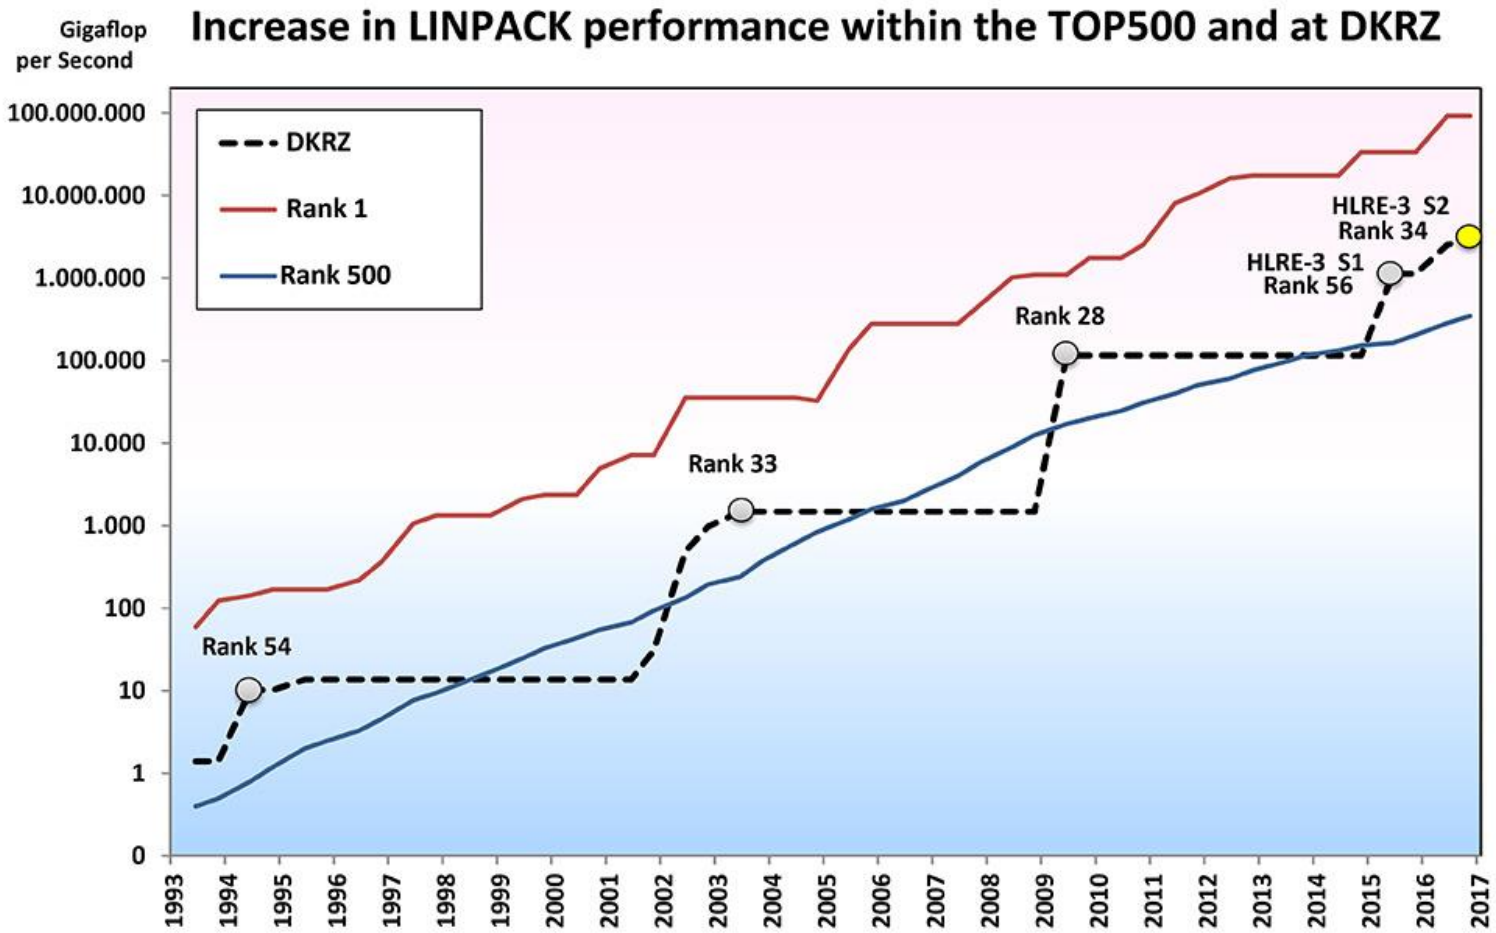

HLRE-3 – Mistral (2015-2020)

bullx DLC 720, 3,500+ nodes, 100,000+ cores, Haswell/Broadwell, 3.6 PFLOPS
240 TB main memory, 54 PB disk storage, 450 GB/s mem-disk rate, FDR network
21 nodes for visualization
hot liquid cooling with high efficiency

no GPUs for acceleration or ML

High Volume Data Archive

- 65,000 slots for tapes in Hamburg (10,000 remote)
- 120+ PB of climate data, increase 40 PB in 2018
- 500 PB capacity until 2020

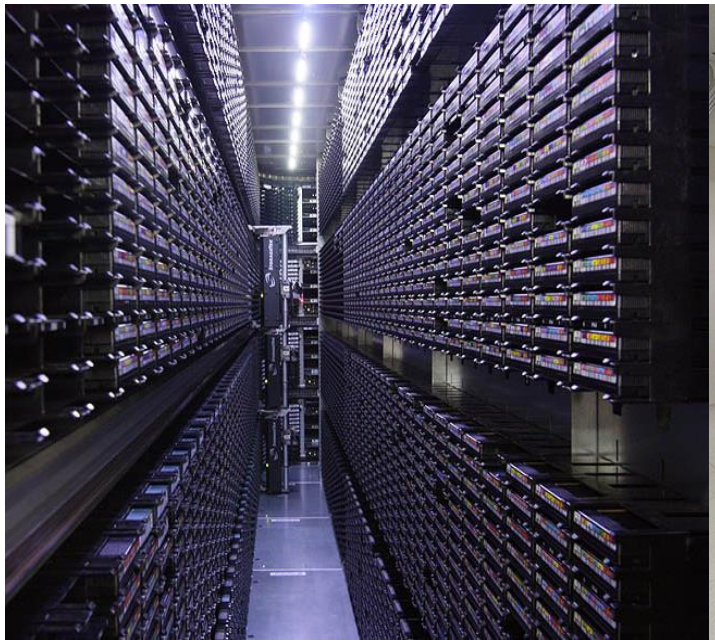

Service Profile

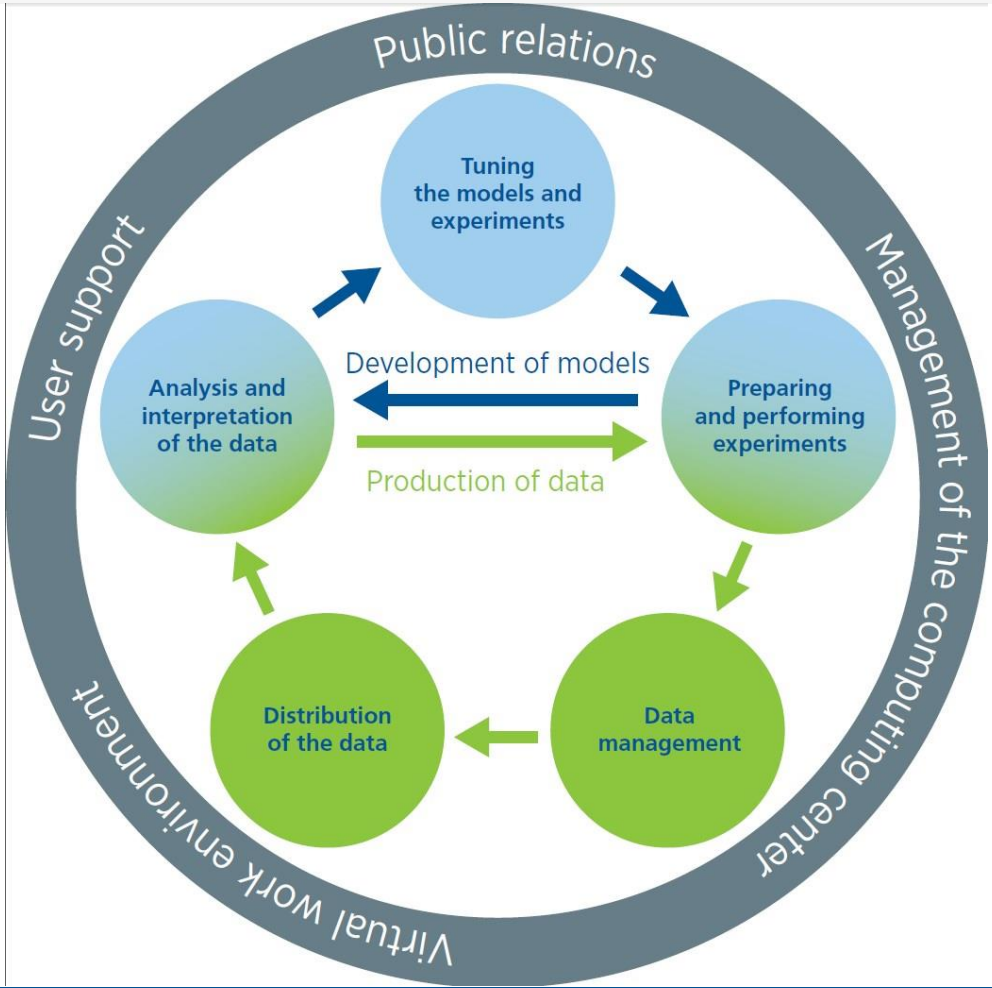

Service Profile – Model Development

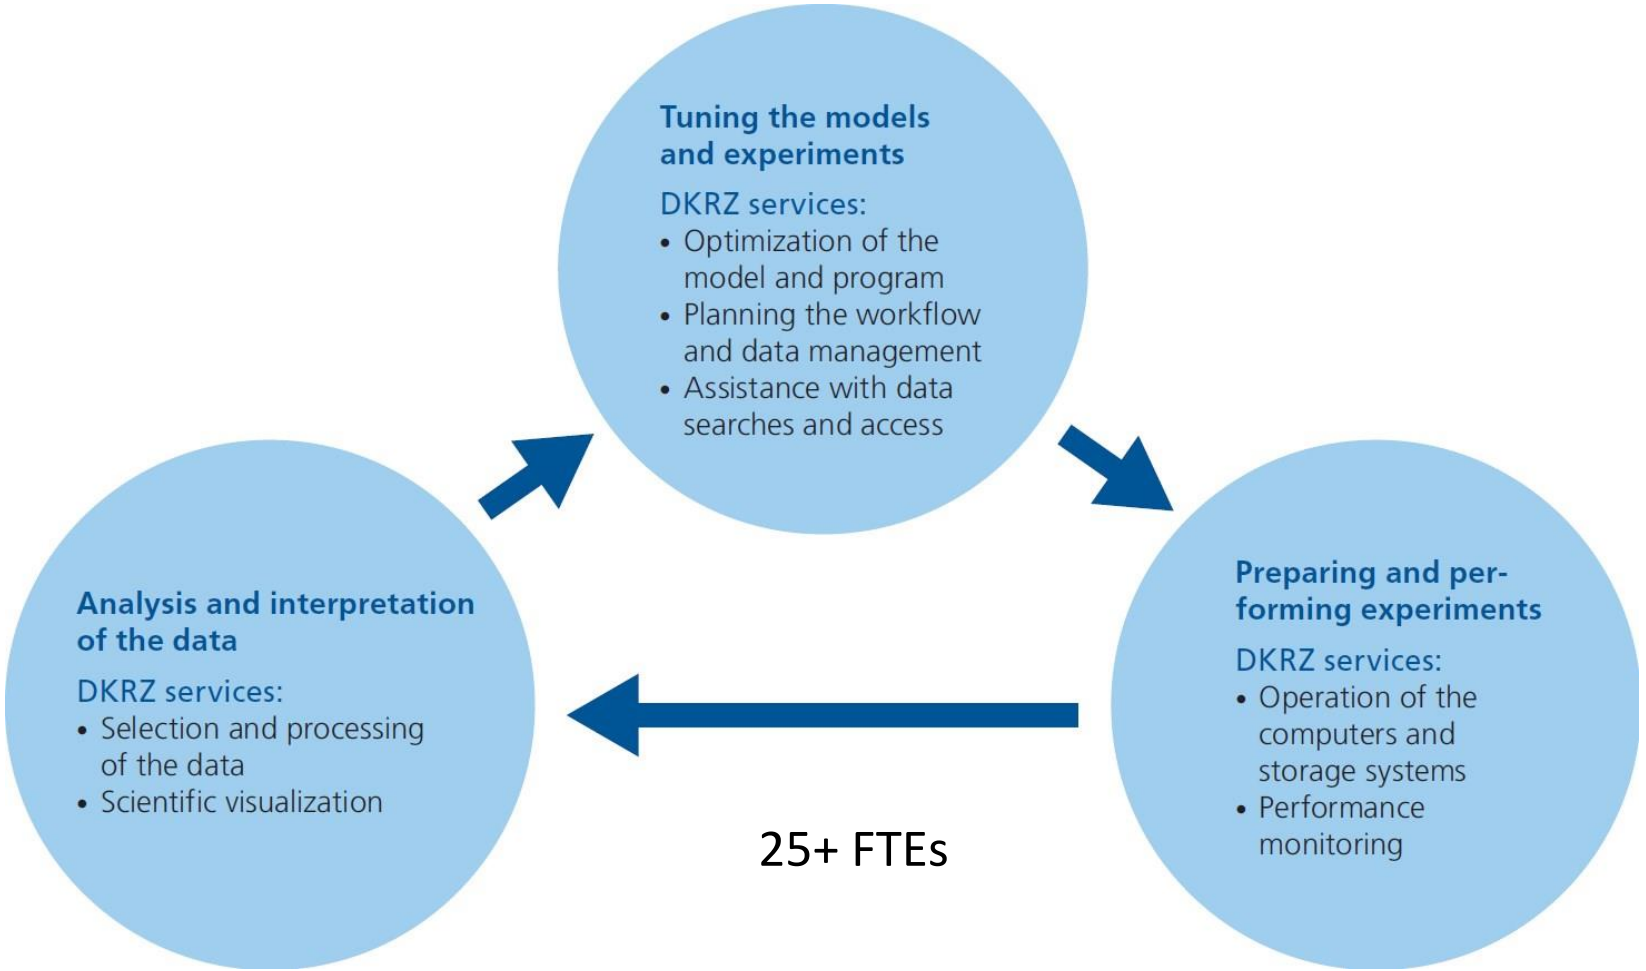

Service Profile – Data Production

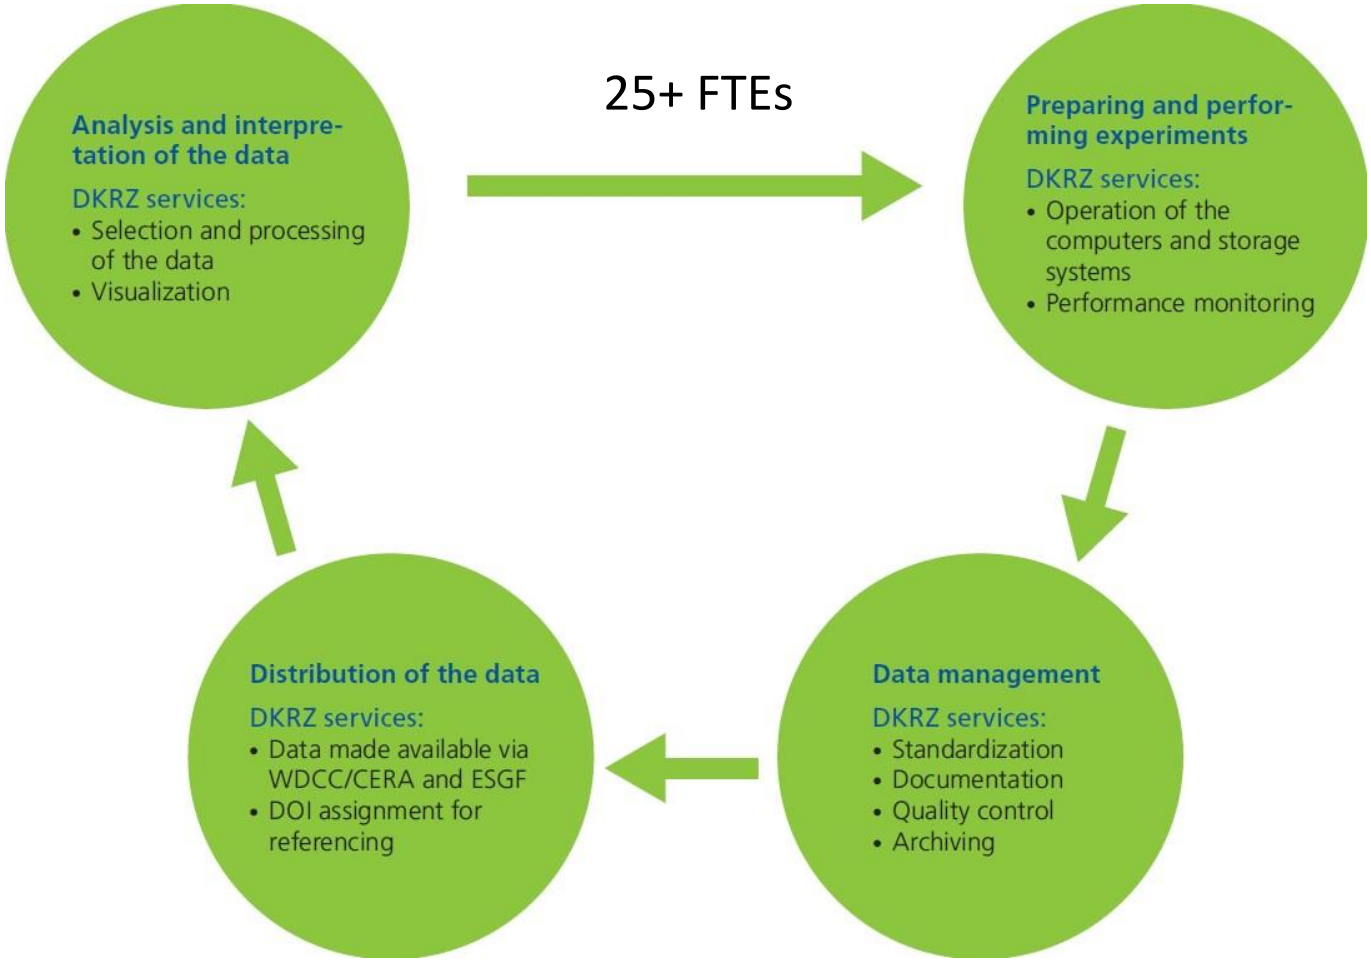

Challenges for HPC Center

Complexities

- Increasing for computer architecture
  - Systems with heterogeneous architectures and increasing number of components (processors, disk, …)
- Increasing for numerical modeling
  - Higher spatial and temporal resolution, further processes, generation of huge data volumes
- Increasing for data usage
  - Visualization and world-wide usage of high data volumes

Problems

- Limited finances for investments in research infrastructures and for their operation
- Limited number of human resources with interdisciplinary scientific education and competences
- **Slowdown of hardware improvement – End of Moore´s Law**

Potential System in 2020

	real 2015/16	2020	Faktor
compute performance (no accelerators)	3.6 PFLOPS	14 PFLOPS	~~10~~ 4
# compute nodes	3,300	GPUs for acceleration and ML (share to be decided)	
# processor cores	100,000+		
main memory	260 TB		
disk capacity	54 PB	120 PB	~~5~~ 2,2
data rate memory to disk			
capacity tape library (2015, 2020)			
data rate disk to tape			
power consumption	1.4 MW	2.0 MW	~~1.0~~ 1.5
Investment costs	€ 36M	€ 40M	1,1

The End of "Traditional" HPC

Traditional HPC

- Mathematics is differential equations
- Processors are conventional homogeneous CPUs
- Profiting from Moore´s Law
  - now: economically reasonable chip improvements very difficult

Example DKRZ (6 years operation, increasing investments)

- NEC -> IBM in 2009: 100x performance
- IBM -> Bull/Atos in 2015: 24x performance
- Bull/Atos -> N.N. in 2020: 4x performance

no more low hanging fruits!

AI/ML Exceeds Moore´s Law

Anthony Sarkis: “Why AI progress is faster than Moore’s Law — the age of the algorithm” (Sep 2018)

https://medium.com/@anthony_sarkis/the-age-of-the-algorithm-why-ai-progress-is-faster-than-moores-law-2fb7d5ae7943

- deep learning algorithms improvement
  - region proposal network
  - image classification
- specialized hardware

Climate Science with Machine Learning Gains Momentum!

Exascale Deep Learning for Climate Analytics

Kurth, Treichler, Romero, Mudigonda, Luehr, Phillips, Mahesh,
Matheson, Deslippe, Fatica, Prabhat, Houston

2018

We extract pixel-level masks of extreme weather patterns using variants of Tiramisu and DeepLabv3+ neural networks. We describe improvements to the software frameworks, input pipeline, and the network training algorithms necessary to efficiently scale deep learning on the Piz Daint and Summit systems. The Tiramisu network scales to **5300 P100 GPUs** with a sustained throughput of **21.0 PF/s** and parallel efficiency of 79.0%. DeepLabv3+ scales up to **27360 V100 GPUs** with a sustained throughput of **325.8 PF/s** and a parallel **efficiency of 90.7%** in **single precision**. By taking advantage of the **FP16 Tensor Cores**, a half-precision version of the DeepLabv3+ network achieves a peak and sustained throughput of 1.13 EF/s and 999.0 PF/s respectively.

A Comment on Weather, Climate, and ML

- Weather

  Weather is the **state of the atmosphere**

- Climate

  Climate is the **statistics of weather** over a usually 30 years interval

- Computational weather prediction
  - Make it quick
- Computational climate projection
  - Make it exact

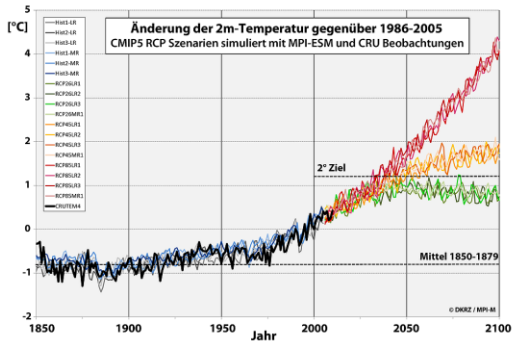

Build up ML Competences and Services

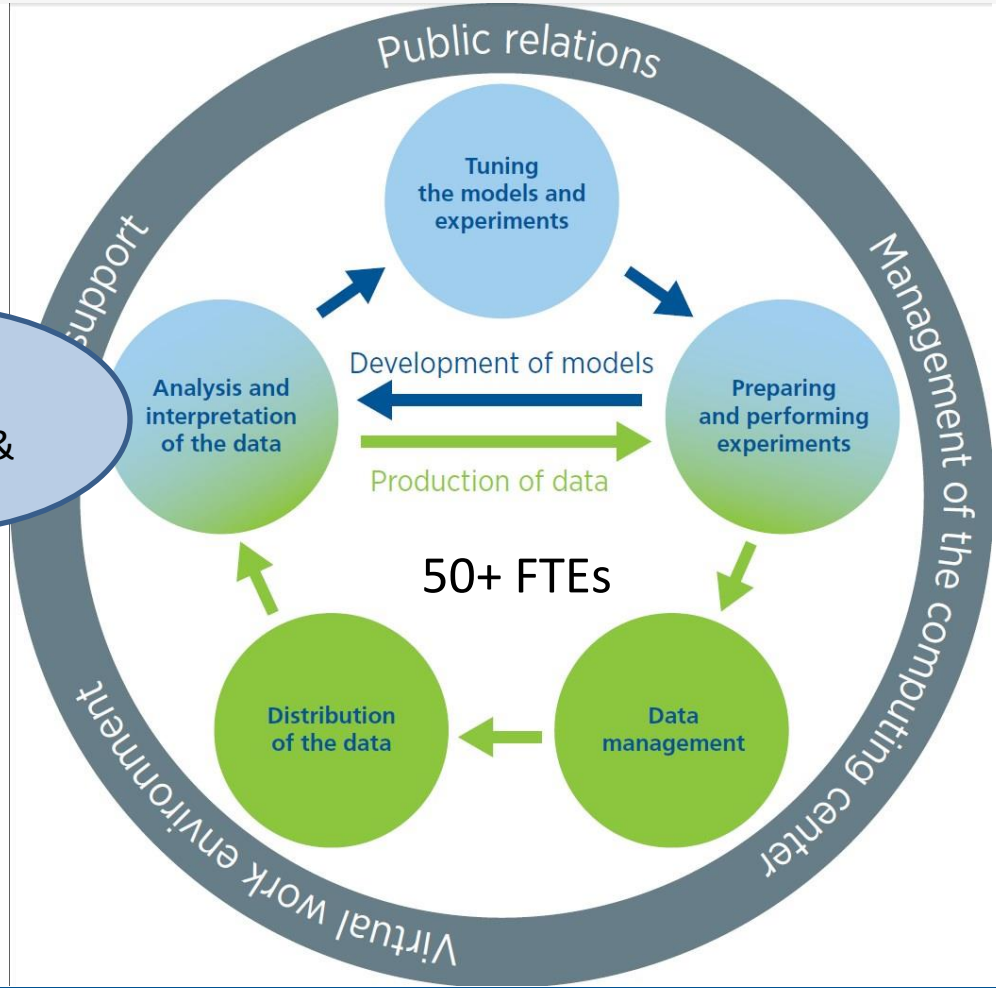

Machine Learning Methods & Services

8 FTEs

Project: HAICU
AI research group

HELMHOLTZ AI | ARTIFICIAL INTELLIGENCE COOPERATION UNIT

OUR MODEL OUR RESEARCH YOU + HELMHOLTZ AI

https://www.haicu.de/

HELMHOLTZ AI COOPERATION UNIT

Vision

Our ambition at the Helmholtz AI Cooperation Unit (**HAICU**) is to reach an internationally visible leadership position in applied Artificial Intelligence (AI)/Machine Learning (ML) by combining unique research questions, data sets and expertise with newly developed AI/ML-based tools and democratized access to them in an open and dynamic community.

Mission

We are a research-driven hub for applied AI that

- fosters cross-field creativity by stimulating collaborative research projects
- identifies and leverages similarities between applications to advance generalised AI/ML methods
- integrates field-specific excellence and AI/ML prowess
- improves the quality, scalability and timely availability of emerging methods and tools
- empowers and trains the current and next generation of scientists

to enable the efficient and agile development and implementation of AI/ML assets across the whole **Helmholtz Association**.

As one of the platforms (**Helmholtz Information & Data Science Academy** (HIDA), **Helmholtz Federated IT Services** (HIFIS), **Helmholtz Imaging Platform** (HIP), **Helmholtz Metadata Collaboration** (HMC) and HAICU) initiated by the **Helmholtz Incubator Information and Data Science** cooperation and openness is hardwired into our approach. Be curious.

The Helmholtz Association's research fields tackle the major challenges and pressing issues facing society and develop sustainable solutions for tomorrow and beyond. As an enabler, HELMHOLTZ AI will transfer methods of applied AI in all of these six research fields:

https://www.haicu.de/themenmenue/research/research-fields/index.html




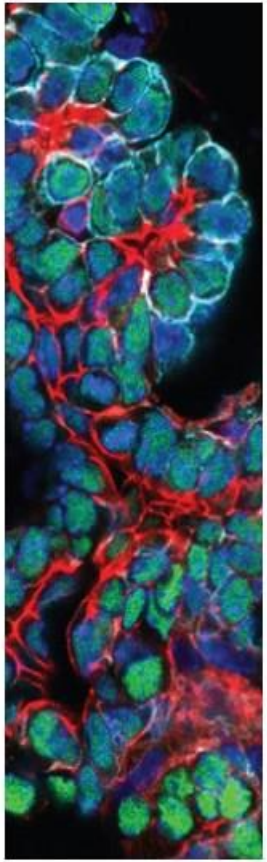


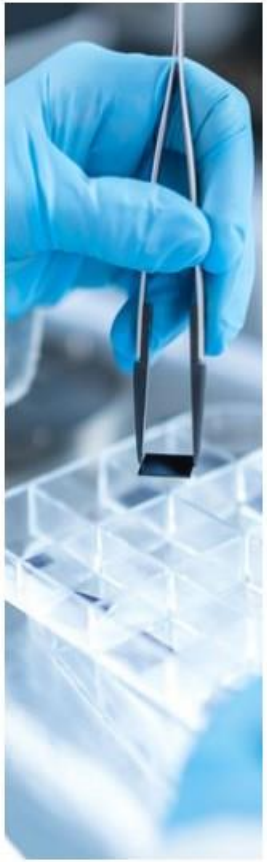

AIM - Artificial Intelligence Innovation in Earth System Analytics and Modelling

Units

- Young investigators group (HZG)
  - Broad expertise in earth system modelling and data science methods
  - Will conduct application-oriented research
- High level support team (DKRZ)
  - Create expertise on AI/ML in earth system modelling and analytics
  - Support earth system science community in applying AI/ML methods to HPC
  - Contribute to education together with AIM-YIG, HAICU locals and central and HIDA
  - Facilitate AI/ML research in earth system science context
    - Software environment and reproducibility, data reduction

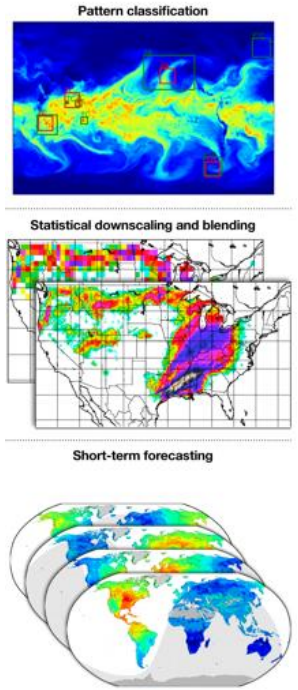

Use Cases Earth and Environment: ML Applications to Climate Research

Improved reconstruction of past ocean temperature based on 'climate-proxy' records (e.g. tree-rings)

- If tested with synthetic data from climate simulations traditional linear methods (linreg) underestimate past variations.
- Gaussian Process Regression (gpreg) yields reconstructions closer to the truth (orig).

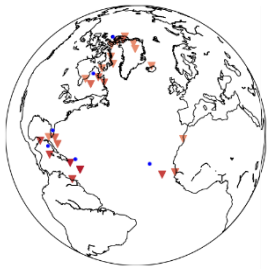

Proxy network (predictors)

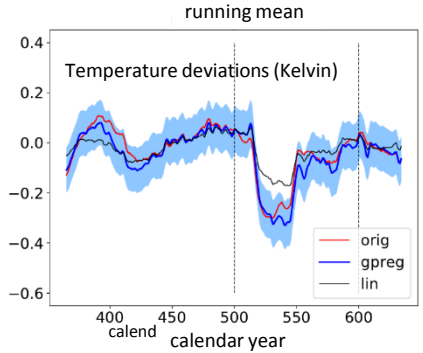

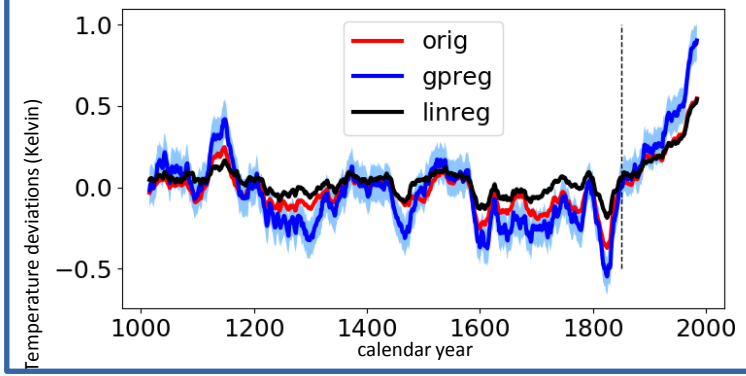

Improve short and long-term prediction of coastal flooding?

- Dynamical storm-surge models, based on wind, water temperature profiles and topography, underestimate surge height.
- Classification or pattern recognition methods based on Random Forest or Neural Networks may better predict extreme surges.
- They can be applied to the output of future climate simulations to estimate future changes.

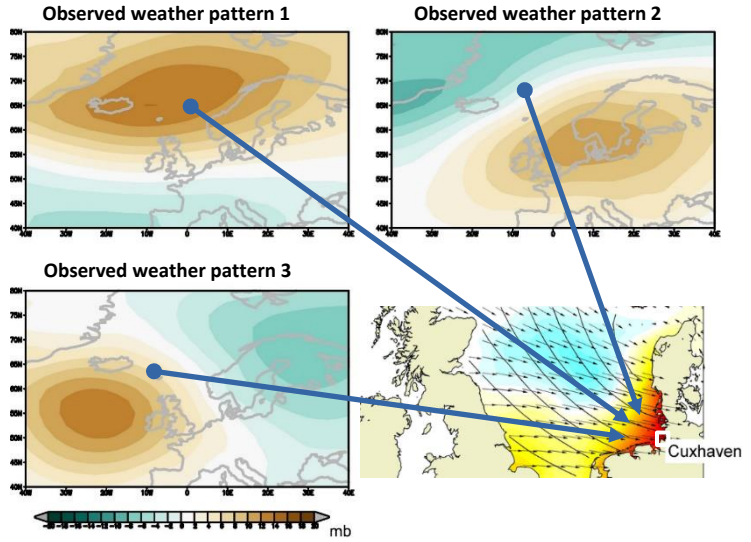

➡ ***Which observed weather pattern leads to storm-surge and how intense?***

➡ ***How will storm surge relevant weather patterns change in the future?***

AI/ML Research Group

Climate Informatics and Technologies

- Interface between AI/ML and climate science
- AI/ML for DKRZ HPC Infrastructure
- Knowledge transfer and method Research
- Utilization of cutting-egde AI/ML technologies for climate scientists
- Focus on research challenges like e.g. climate prediction and missing climate information of the past

Climate Models

Machine Learning

Example: Missing Historical Climate Information in Observations

Nvidia Technology: Image inpainting on irregular holes using deep convolutional neural networks

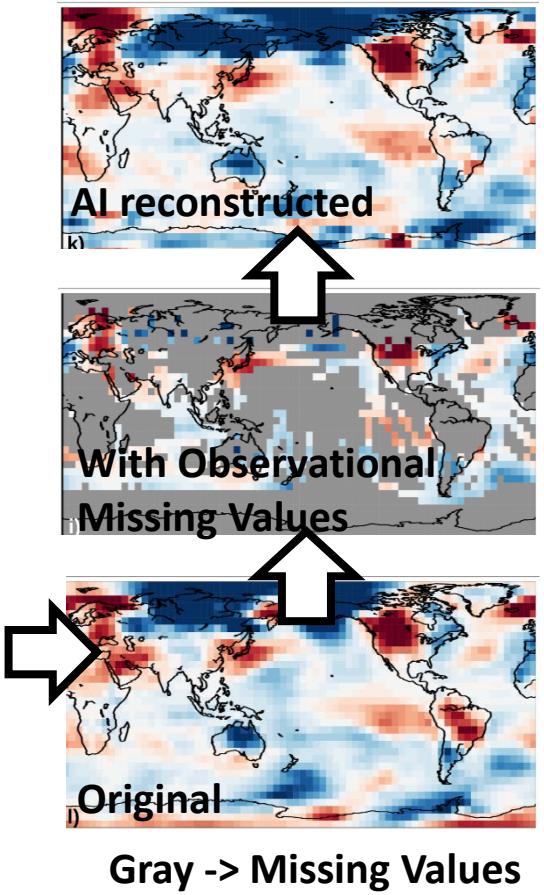

Retrospect and Prospect

- Machine learning and specialized hardware have a 60+ years success record
- Traditional HPC slows down because of HW issues
- ML is now being adopted by classical HPC communities

- With computational climate science we still explore the application fields
- DKRZ will extend its methods and services portfolio