Latest Trends at DKRZ

Installations and Machine Learning Activities

Prof. Dr. Thomas Ludwig
German Climate Computing Center (DKRZ)
University of Hamburg, Department for Computer Science (UHH/FBI)

Installations at DKRZ

35 Years of Computer and Storage Systems (1987-2022)

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

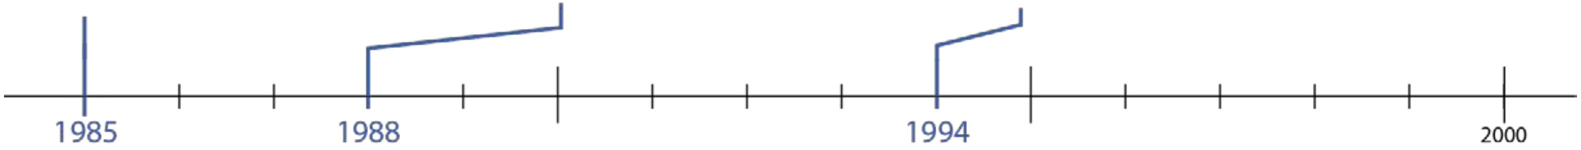

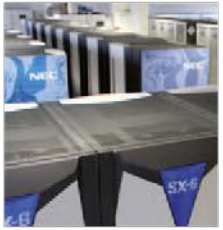

2002: NEC SX-6 „Hurrikan“

- 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 processor cores**
- **3.6 Petaflops**
- **240 Terabyte memory**
- **54 Petabyte disk space**
- **up to 500 Petabyte tape archive**
  120 PetaByte are used

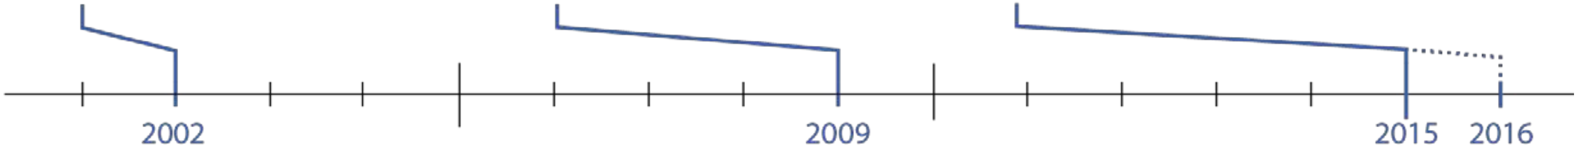

Powerful Computers for Competitive Research

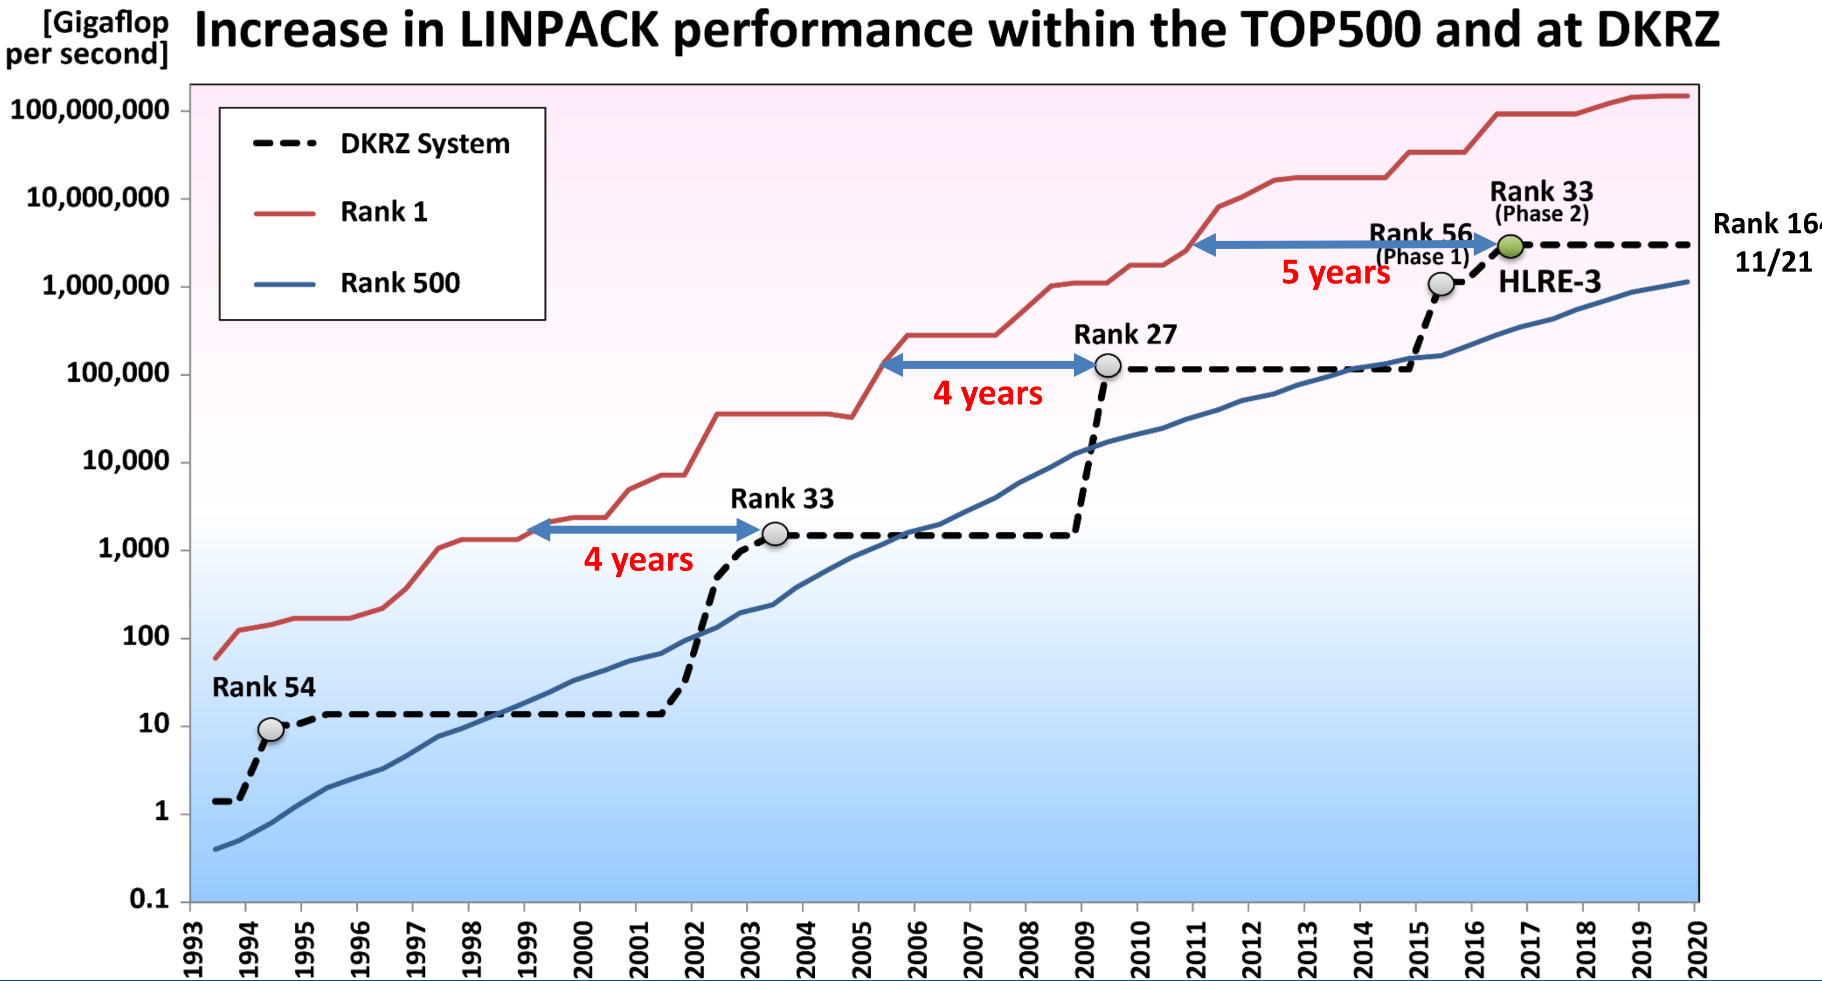

HLRE-3 – Mistral (2015-2022)

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

HLRE-3: Proposal 2011 vs. Real System 2015/16

Maße	2011	2015/16	Faktor
compute performance (no accelerators)	3 PFLOPS	3.6 PFLOPS	**(20) 22**
# compute nodes	3,000	3,300	(12) 13
# processor cores	120,000+	100,000+	(14) 12
main memory	360 TB	260 TB	(18) 13
disk capacity	120 PB	54 PB	**(20) 9**
data rate memory to disk	600 GB/s	450 GB/s	(20) 15
capacity tape library (2015, 2020)	650 PB	~400 PB	(10) 6
data rate disk to tape	30 GB/s	20 GB/s	(10) 7
power consumption	1.6 MW	1.4 MW	0.9
investment costs (compute, disk, tape)	€ 36M	€ 36M	

Proposal 2016 for a System in 2020 (HLRE-3: 2015)

	real 2015/16	2020	Faktor
compute performance (1/5 with GPUs)	3.6 PFLOPS	35 PFLOPS	**10**
# compute nodes	3,300	3,100	1
# processor cores	100,000+	220,000	2
main memory	260 TB	640 TB	2,5
disk capacity	54 PB	250 PB	**5**
power consumption	1.5 MW	1.4 MW	**1.0**
total costs (incl. HSM, computer room refurbishment)	€ 41M	€ 45M	1,1

HLRE-4 System in 2022

	real 2015/16	2022	Faktor
compute performance (1/5 with GPUs)	3.6 PFLOPS	16+ PFLOPS	**4.5**
# compute nodes	3,300	2,900	0,9
# processor cores	100,000+	370,000+	3.7
main memory	260 TB	815 TB	3
disk capacity	54 PB	130 PB	**2.4**
power consumption	1.5 MW	3.0 MW	**2**
total costs (incl. HSM, computer room refurbishment)	€ 41M	€ 45M	1,1

HLRE-4 “Levante”

Data Sheet

Installation: 2021-2022
Producer: Atos
Model: Atos BullSequana XH2000
No. of cores: 370.000
Network: HDR-Infiniband, NVIDIA Mellanox InfiniBand HDR 100G/200G
Disk system: 130 Petabyte (Lustre) von DDN
HSM system: Cristie / StrongBox Data Solutions / Huawei

CPU-Partition

- 2.832 compute nodes
  - 2.520 nodes with 2 processors AMD 7763 (256 GB memory)
  - 294 nodes with 2 processors AMD 7763 (512 GB memory)
  - 18 nodes with 2 processors AMD 7763 (1024 GB memory)
- Peak performance: 14 PetaFLOPS
- Main memory: 815 TB

GPU-Partition

- 60 GPU nodes with
  - 2 processors AMD 7713 (512 GB memory)
  - 4 Nvidia A100 GPUs (56 nodes with 80 GB, 4 nodes with 40 GB local memory)
- Peak performance: 2.8 PetaFLOPS
- Main memory: 30 TB

New Cooling System

The Deceleration of Acceleration

- Systems at DKRZ (same investment in € plus inflation)
  - 2009: NEC -> IBM with 100x improvement
  - 2015: IBM -> Bull/Atos with 22x improvement
  - 2022: Bull/Atos -> Bull/Atos with 4.5x improvement
    - 2x power consumption
- Partner centers in the last years
  - LRZ (Garching/München): 4x
  - HLRS (Stuttgart): 3.5x

Use Case from the Proposal in 2016

On-line nesting approach in chemistry-climate models

Field of application:		**Model used:**		**Project:**
Climate-Chemistry		MECO(n): T42L90MA, T106L191MA ICON (chem. enabled)		
	Compute time [node hours / sim years]	**Job size [nodes]**	**Storage size [GByte / year]**	**Model turnaround [sim years / day]**
today	5100	-		
tomorrow	~ 300.000	-		
growth factor	**60**	-	**14**	**6**
Main reasons for increased computational demand:	– **Cascades of nesting levels down to 10 km or below** – **Increased EMAC resolution** – **Refined chemical mechanism**			

Use Cases from the Proposal in 2016

Not a single one of the 11 use cases can be satisfied

they range from 7x to 1000x

with their compute performance requirements

The End of "Traditional" HPC

Traditional HPC

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

- No more low hanging fruits!
- No more fruits?

Preview to System in 2028

- System
  - Presumably installed only after 7 years
  - Performance increase 2x?
  - Disk space increase?
  - Power consumption increase? 2x?
  - Weight and foot print?
- Scientists

  What new and exciting research to conduct?
- Vendors

  How to make business with HPC systems?

Instead: Invest in Brainware

- Better codes, better software engineering
- Adaptations to heterogeneous hardware
- Interdisciplinary teams of domain scientists, computer scientists, mathematicians, …
- Adopt new methods
  - Data intensive science
  - Machine learning (artificial intelligence)

Machine Learning Activities

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 (GPGPU,…)

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

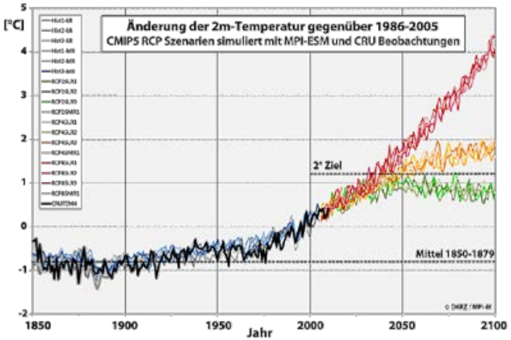

Build up ML Competences and Services

Machine Learning Methods & Services

8 FTEs

Project: HAICU
AI research group

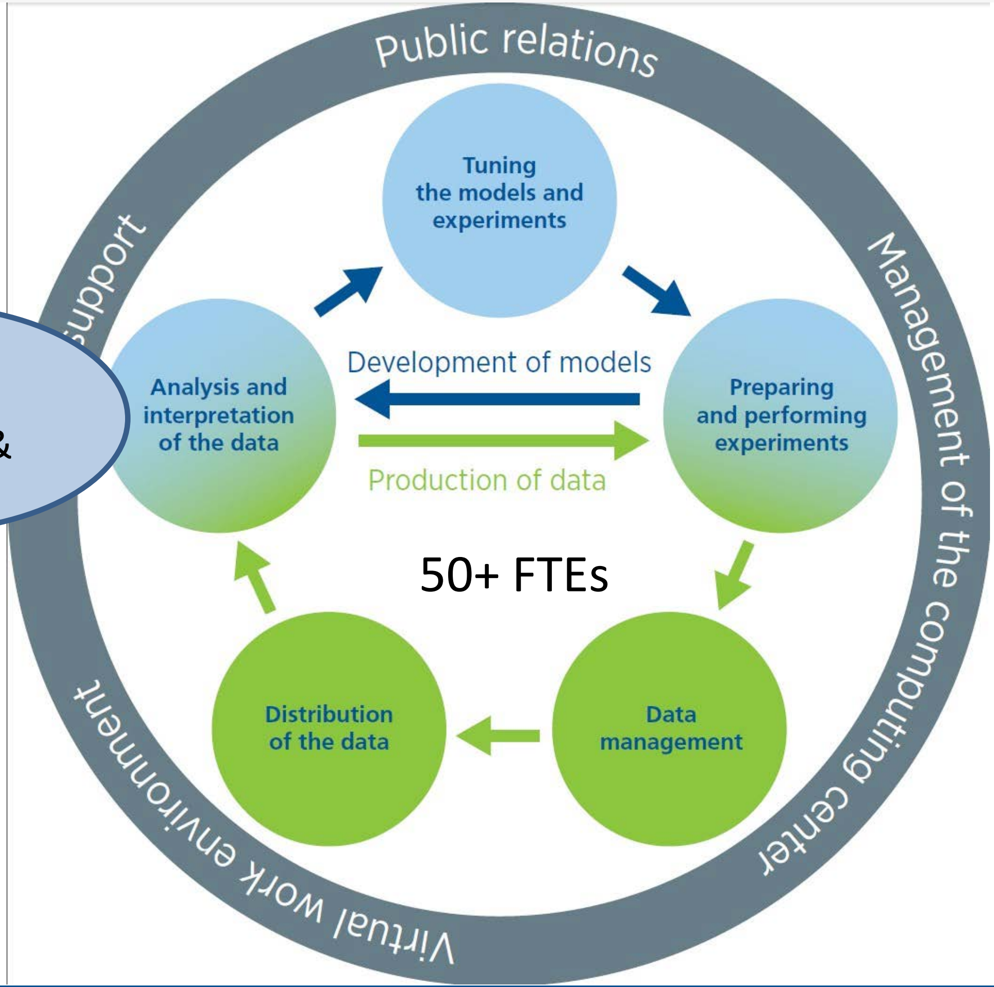

AIM - Artificial Intelligence Innovation in Earth System Analytics and Modelling

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

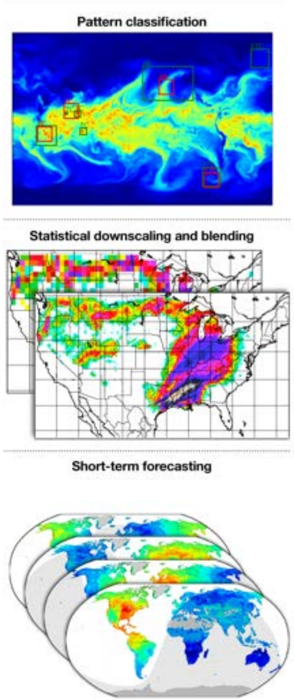

Use Cases Earth and Environment: ML Applications to Climate Research

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).

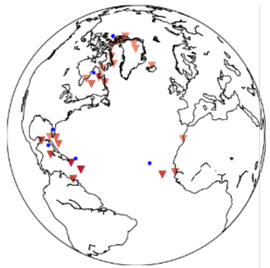

Proxy network (predictors)

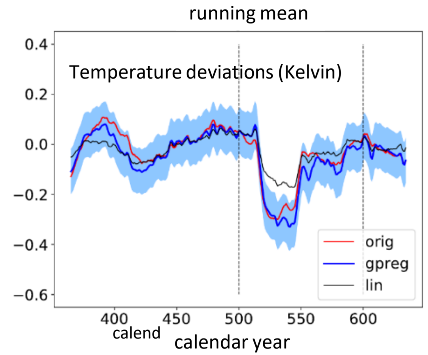

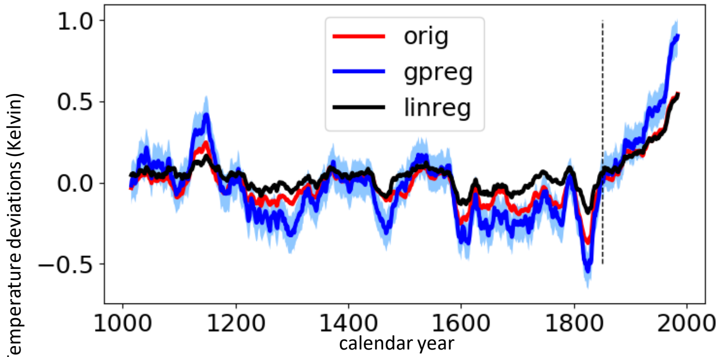

Improve short and long-term prediction of coastal flooding?

- 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.

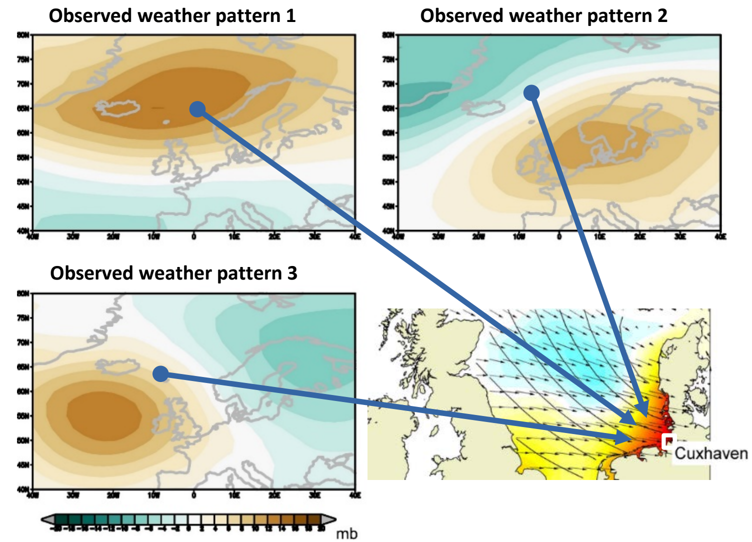

➡ ***Which observed weather pattern leads to storm-surge and how intense?***

➡ ***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

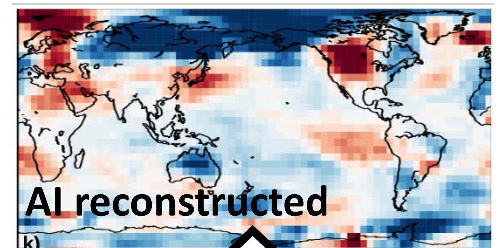

With Observational Missing Values

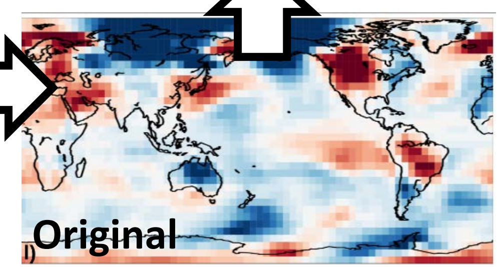

Gray -> Missing Values

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