Project: **1176**

Project title: **AIM**

Principal investigator: **Tobias Weigel**

Report period: **2021-05-01 to 2022-04-30**

Text: Maximum of two pages incuding figures. Reports for joint projects may be longer.

Vouchers using DKRZ resources

Project 1176 (AIM) uses DKRZ resources to facilitate advanced ML applications in Earth and Environmental Sciences for scientific cases from users across the Helmholtz association. These are formulated and managed as support project vouchers in the framework of Helmholtz AI. Out of all the voucher projects active in the report period, the following ones used substantial granted storage and/or compute resources on Mistral.

Center	Title	Start/End
GFZ	Extension of ML to new data version of CyGNNS (v3)	20/11 - 21/06
Hereon	Detection and short-term forecast potential for rogue waves	21/02 - 21/08
KIT	FastChem atmospheric chemistry: 1y simulation data and box model verification assessment	21/02 - 21/08
GEOMAR	Preparation and conduct of a ML training course for mardata	21/09 - 21/10
KIT	Modelling atmospheric chemistry tracer transport with ML	21/09 - 21/12
AWI	SWIFT-AI: Coupling to ICON-ART	22/02 - 22/08
Hereon	Python-Fortran bridges for ICON	22/02 - 22/08

Expired resources:

A substantial fraction of requested CPU compute resources expired; consumption of compute time was overestimated due to lack of knowledge of how much compute time ML projects require. In consequence, already in our last project application, we only requested a very small amount. Most of the compute is done on GPU nodes and this is expected to be similar for Levante as well, though compute time required for some data preprocessing tasks done on CPUs is also significant for some voucher projects.

Overall, ML projects cannot work efficiently without GPU access, but, on the other hand, also do not require substantial amount of GPU compute time as most of them still run on single nodes. However, as opposed to requiring substantial node-hours, ML applications are rather bound by the size of GPU memory, age of the GPU architecture, memory throughput and latency of disk-main memory-GPU pipelines.

Storage consumption (/work) has been largely as planned, reaching quota limits occasionally, which was dealt with by removing or archiving older data.

Tape archive usage:

Data files of older vouchers or completed stages of ongoing vouchers were archived; in particular, working data of the GFZ CyGNSS and KIT atmospheric chemistry work was put to tape.