



FASTEST PATH TO AI INNOVATION

ALTAIR SOLUTION AT LRZ ENABLES MACHINE LEARNING AT MCML


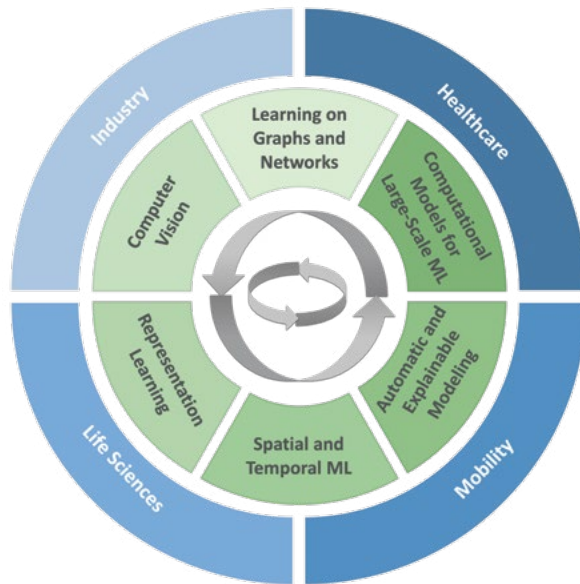
About the Customer

With artificial intelligence (AI) and machine learning (ML) increasingly on the rise in all industries, the global demand for scalable high-performance computing (HPC) capacity is constantly growing. As an institute of the Bavarian Academy of Sciences and Humanities, the Leibniz Supercomputing Centre (LRZ) is the high-performance computing center for Munich's top universities and research institutions in Germany and Europe. The LRZ provides robust and holistic IT infrastructure to its customers and the complete range of resources, services, consulting, and support enabling advanced research and discovery for all kinds of applications, including the highly demanding areas of AI and ML.



Thanks to Altair's expertise in GPU parallelization, we were able to provide the MCML with a perfectly customized and optimal solution, ensuring the performance and scalability it needs to further advance HPAI and AI research at the highest level in areas such as speech processing and image analysis."

Dr. Nicolay Hammer,
head of LRZ Big Data
and AI team



Their Challenge

LRZ was tasked with designing and installing a new infrastructure for and operate it on behalf of the **Munich Center for Machine Learning (MCML)**, one of the recently established centers of excellence for AI research and innovation. With 18 academic teams providing combined expertise in data science, computer science, and statistics, the MCML is focused on AI to further develop **ML technologies** and methods in different research fields such as natural language processing (NLP) and biomedical research. The challenge for LRZ was to design a customized computing system for MCML - with parallel GPUs allowing the MCML to scale up easily to adapt to their workload changes - and provide a **robust and accessible infrastructure** while meeting advanced storage, power supply, and cooling demands.

Our Solution

LRZ collaborates with Altair who provides and installs **NVIDIA's DGX systems** and infrastructure at the LRZ and at customer sites. To meet the special needs of best-in-class AI and ML research at the MCML, the Altair team had to design a cluster solution that provides strong nodes and big capacity to train a multitude of networks and is able to **develop highly scalable ML and AI applications**. Thanks to Altair's profound expertise in the parallelization of GPU applications the team was able to develop, prepare, and install a highly scalable system, consistent of eight DGX A100 systems combined to a powerful cluster called **DGX POD**, while ensuring continuous power supply and the necessary cooling of the cluster. Altair set up a system that provided the needed **connectivity, accessibility, and storage** and ensured that the data lake is open to any user for any simulation. These simulations and their respective result analyses constitute the starting point in AI research by using this data as the input for AI models to be trained. Ensuring smooth deployment according to NVIDIA standards and operation, Altair ensured that the new system allows for data handling between local and global network instances. Subsequent to the installation, Altair provided a comprehensive documentation of the new system.

Results

Altair's professional service team designed and installed the new solution, offering **excellent infrastructure** and service to one of LRZ's important customers. Enabling **fast and parallel computation**, the DGX POD installed at the LRZ ensures that the GPUs can fully leverage their performance allowing the MCML to **operationalize AI at scale**. The potential of the DGX POD shows in its scalability and performance with 8 NVIDIA DGX A100 systems delivering 40-petaFLOPS of AI performance. The system serves 18 teams drawn from two local universities that make up the MCML. This investment in the large-scale, multi-node AI infrastructure will enable academic teams at the MCML to efficiently process the demanding AI workloads and support the training of ML models, and gain insights faster than ever before. Researchers at LRZ also enjoy the ease of accessing containers on NGC Catalog, NVIDIA's hub for GPU-optimized software.

LEFT: MCML structure: MCML's fundamental research is bundled in six profile areas.
TOP: NVIDIA DGX A100 is the building block of the AI data center.
BOTTOM: The Leibniz Supercomputing Centre (LRZ) is the high-performance computer center for Munich's top universities and research institutions in Germany and Europe.