



Science and  
Technology  
Facilities Council

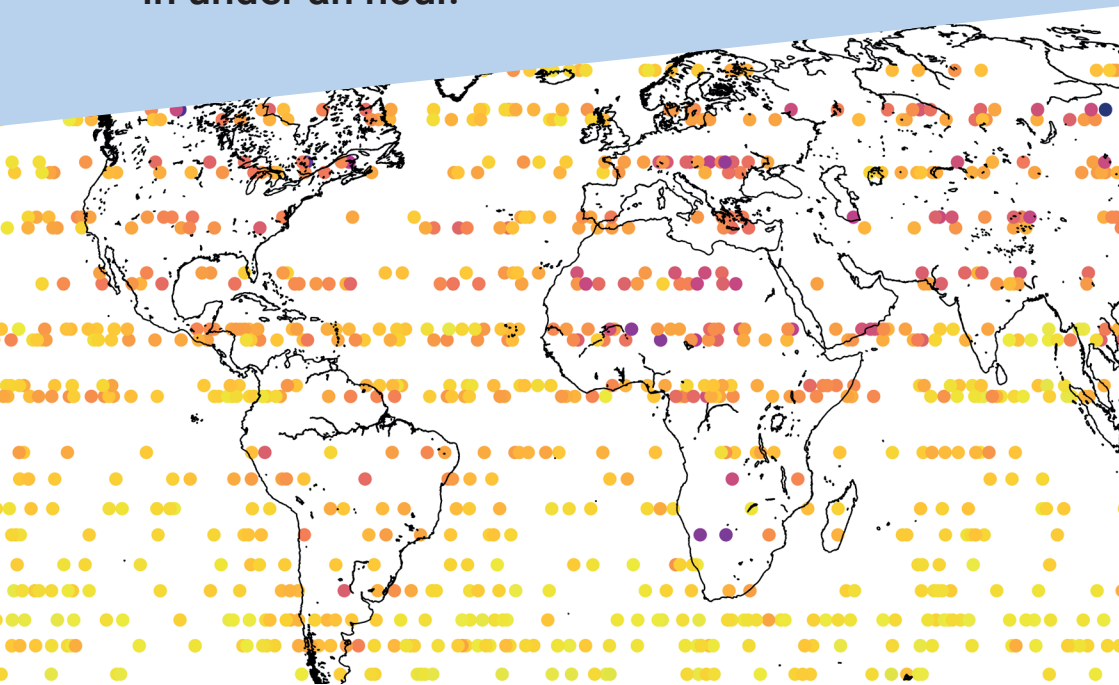
# Pearl

*A small but powerful deep learning system.*



# About Pearl

Small but powerful, Pearl is a deep learning system with the capacity to greatly speed up the training process of complex artificial intelligence models. It comprises two tightly coupled, state-of-the-art NVIDIA DGX2 nodes, and a large array of high-speed storage. The system is expected to accelerate the research from multiple scientific domains. For instance, a complex model that would require about 500 hours on a single CPU can be performed on Pearl in under an hour.





# Impact studies

Cryo-electron microscopy allows researchers to obtain high-resolution atomic structures of biological molecules such as proteins or viruses. These structures are sought after by scientists who can use this detailed knowledge to identify, for example, potential new drug targets. Manual analysis of cryo-electron microphotographs is time-consuming and dull so novel machine-learning techniques, powered by Pearl, aim to streamline the process.

Weather observational stations across the world predict cloud coverage in their immediate area by taking and analysing photographs, which has an important impact on weather and environmental predictions. Unfortunately, efficiency is far from optimal. Pearl is powering an AI model which combines multiple satellites, observational stations and datasets to improve efficiency of cloud quantification.

runs on two  
**DGX2**  
systems

**4PFLOPS**  
performance capacity

**600TB**  
storage

**180,000**  
GPU cores



The  
Alan Turing  
Institute

Pearl is funded by the Alan Turing Institute and hosted and managed at STFC by the Scientific Machine Learning group within the Scientific Computing Department.

