






MAGEO

- Large, multi-user GPU cluster for AI in EO
- Hardware selected for machine learning and deep learning applications
- Apply GPU compute for operational high resolution processing
- Supporting EO and AI researchers through an AI service under NEODAAS

Email: stgo@pml.ac.uk, info@neodaas.ac.uk

MAGEO - The Massive GPU Cluster For Earth Observation: Supporting Artificial Intelligence Research

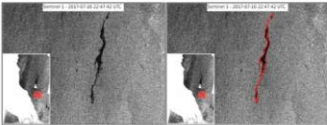
Stephen Goult, Peter Walker, Daniel Clewley, Steve Groom

Introduction
NERC recently recommissioned the NERC Earth Observation Data Acquisition and Analysis Service (NEODAAS) at Plymouth Marine Laboratory (PML) for a five-year period and awarded significant funding to create the Massive Graphic Processing Unit (GPU) cluster for Earth Observation (MAGEO). This facility will provide a high-performance compute capability to support Artificial Intelligence covering Deep Learning and Machine Learning (DL/ML) in the Earth Observation data space coupled with assistance from NEODAAS staff.

NEODAAS Artificial Intelligence Applications Service
Through NEODAAS, MAGEO will provide a service to UK researchers supporting Earth Observation scientists seeking to apply DL/ML tools to their own data, or artificial intelligence experts looking to apply their techniques to EO data. This service will come with the expertise in both fields hosted at PML, and opportunities for training in both the EO and ML/DL spaces.

Pilot projects
Several users have already expressed interest in using MAGEO for DL/ML projects, these include:

- Wildfire detection
- Oil spill detection
- Mangrove mapping
- Plastics detection
- Algal bloom and river plume segmentation
- Satellite and airborne data fusion



Sentinel 1 SAR data: detection and segmentation of oil slicks in the Strait of Malacca, Malaysia using ensemble ML methods developed by PML as part of the GASOS project

Proposed Technical Specification*:
Computation:

- 40 NVIDIA Tesla V100 Graphics Processing Unit cards
- >150,000 CUDA Cores for GPU processing
- >>30,000 Tensor cores for deep learning acceleration
- 500 CPU cores

Memory:

- 3 Terabyte RAM
- 1.5 Terabyte on-GPU RAM

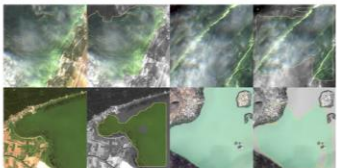
Storage:

- 40 Terabyte on-node solid state flash storage
- 500 Terabyte InfiniBand networked fast storage
- 6 Petabyte archive of NEODAAS satellite data

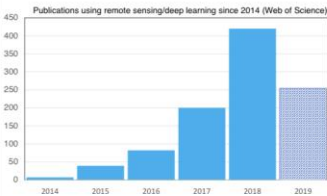
*indicative values

This specification will allow MAGEO to be shared between multiple users in parallel and support many configurations of resource allocation.

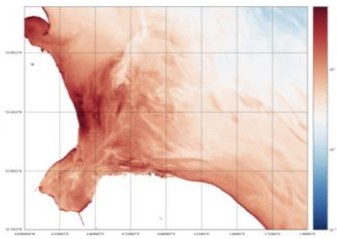
Beyond DL and ML applications
MAGEO will also support new approaches to processing of remote sensing data through access to the NVIDIA Compute Unified Device Architecture (CUDA) for GPUs designed to perform complex tasks in parallel across array based data. It is anticipated this will provide performance acceleration for processing of high resolution data, such as those produced with the Sentinel 2 MultiSpectral Instrument (MSI).



10m RGB data collected by Sentinel 2 MSI; algal blooms in the Baltic sea and Mexico have been segmented and identified from the original image by a DL Convolutional Neural Network. A yellow line and coloured region describes the extent identified. Developed as a collaborative project between PML and the University of Exeter.



Increasing Demand
The number of publications using deep learning with remote sensing data has consistently doubled year on year since 2014, with 2019 already over 250 papers. Demand for hardware like MAGEO will increase with future projects like these.



100m resampled turbidity data 5 day composite over the wash, East Anglia produced with data collected by Sentinel 2 MSI as part of the EMORES project

Contact us
For more information about usage of MAGEO and the AI applications service contact NEODAAS with the details below, or come and speak to the team at the NCEO NEODAAS stand
 Email: info@neodaas.ac.uk
 Website: www.neodaas.ac.uk