

Dargana

Aspia Space: M Smith, L Fleming, J Geach,
R Roberts, F Kalaitzi, J Banister

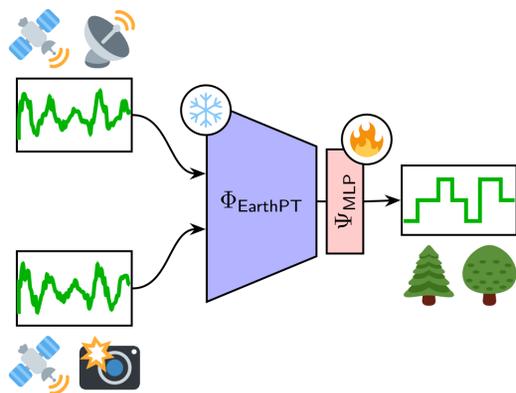
The EarthPT LOM

EarthPT is an EO foundation model (a 'Large Observation Model') trained on 30 billion 8×8 pixel patches of Earth observation data from ClearSky and Sentinel-1 C-band imagery, spanning eight years over 640 000 hectares of the South West of England. In previous work (Smith+2023), we demonstrated that **EarthPT has the ability to encode rich representations of the dynamic properties of the Earth's surface.**

Here, we demonstrate that it is possible to efficiently leverage these representations by fine-tuning the EarthPT foundation model on specialised data.

Fine-tuning EarthPT

We present Dargana, a fine-tuned variant of the EarthPT LOM that achieves specialisation using **<3% of its pre-training data volume and 5% of its pre-training compute:**



This efficiency comes from using standard fine-tuning techniques that have gained popularity in natural language processing; freezing the base model (Φ_{EarthPT}) weights and attaching a multilayer perceptron (Ψ_{MLP}) to the central layer, updating only the MLP head during training.

Case study: Cornwall, UK

Dargana is fine-tuned to generate regularly updated classification of tree canopy cover at a 10m resolution, distinguishing conifer and broadleaved tree types. Using Cornwall, UK as our test case, **we achieve a pixel-level ROC-AUC of 0.98 on unseen data.**

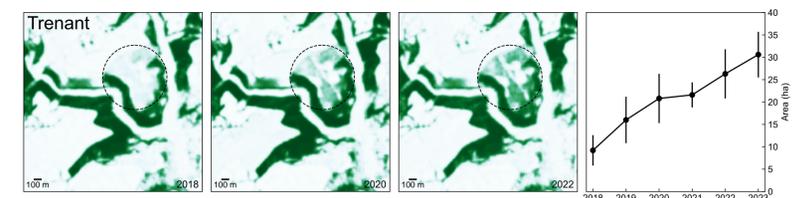
Fine-tuning the EarthPT foundation model for accurate tree canopy classification from space.



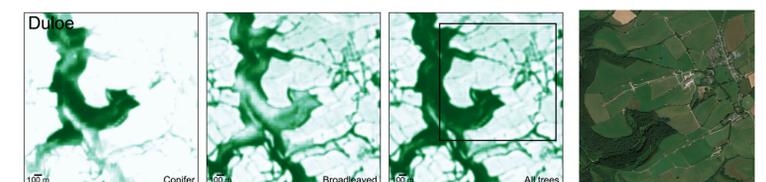
Results

Our headline figure (centre) presents a full mosaic of Cornwall as predicted by Dargana on the 31st December 2023. For illustration purposes, we combine all 'tree-like' objects into a single class, coloured green.

Due to the architecture of EarthPT, we can continuously update class probability estimates as new observations become available. This enables **regular near-time monitoring of change**, such as new planting establishment. Below we show a sequence of Dargana predictions over the Trenant region of Cornwall. We can see an increasing probability of 'tree-like' objects as newly planted saplings grow and establish themselves.



Dargana extrapolates below the limit of its fine-tuning data—derived from the National Forest Inventory (NFI)—to identify trees at granularities well below the NFI's 0.5 hectares minimum mapping unit. Below we show the Duloe region of Cornwall, where we can see **the model predicts the presence of broadleaved trees tracing hedgerows, coppice, trees within built-up environment and even large lone trees.** Gaps and low cover within areas of contiguous woodland are also seen.



A scalable future

The scalability of our approach is particularly promising. Having demonstrated success at county scale ($O(100 \text{ km})$), there are no fundamental barriers to extending coverage nationally or internationally as the underlying architecture is inherently scalable to larger geographical regions. We are also not restricted to land use classification; we can alter the MLP head for any task where we have sufficient data for fine-tuning. Future work will focus on expanding both geographical scope and the range of downstream tasks, working toward a comprehensive suite of LOM-driven environmental monitoring tools.