



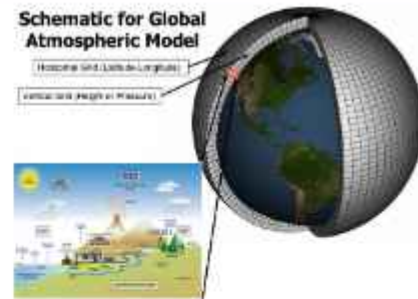
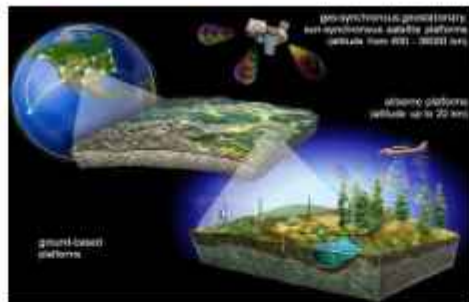
WELCOME!

Environmental Intelligence @ Exeter

11th September 2025

AI for Environment

Monitoring,
modelling,
predicting



New tools &
technologies

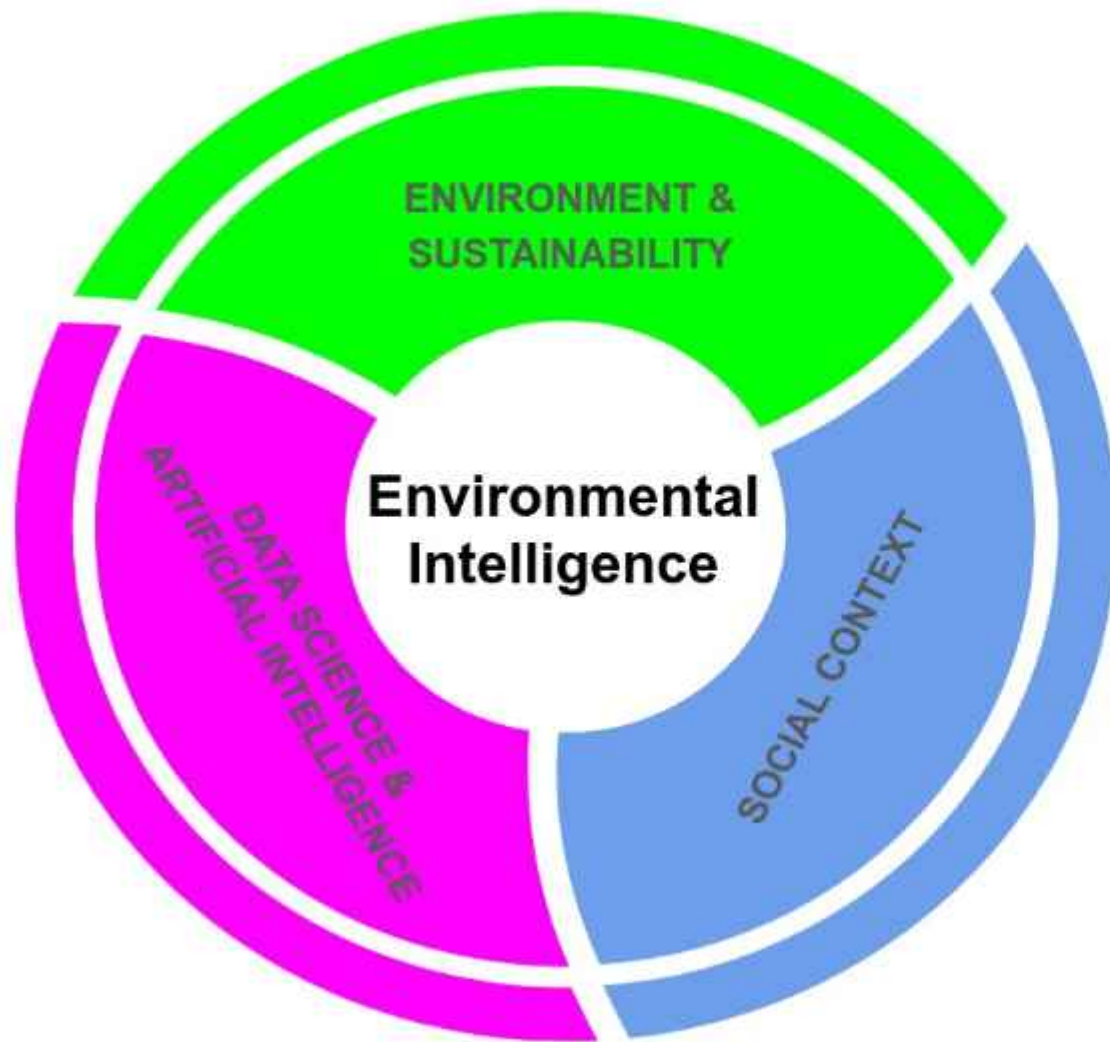


Managing
complex
systems





University
of Exeter



“Using data
science and AI to
solve
environmental
challenges”

**Interdisciplinary
Applied
Exeter USP?**

El@Exeter



Environmental
Intelligence CDT



Centre for Environmental Intelligence

Today: Time to celebrate & look ahead :-)

0900	ARRIVAL, POSTERS & REFRESHMENTS		
1000	SESSION 1: WELCOME & KEYNOTE Venue: Henderson LT Tim Lenton: Positive tipping to avoid climate tipping		
1100	SESSION 2A: EI for Biodiversity & Nature Venue: Henderson LT	SESSION 2B: EI for Earth & Climate Venue: Conf Room 1+2	SESSION 2C: EI for People & Health Venue: Sem Room A+B
1200	LUNCH BREAK & POSTERS		
1300	SESSION 3: PANEL DISCUSSIONS Venue: Henderson LT Social Dimensions of Environmental Intelligence Environmental Intelligence in Practice		
1430	SESSION 4A: Machine Learning for Earth Obs Venue: Henderson LT	SESSION 4B: Climate Comms & Data Sci (C3DS) Venue: Conf Room 1+2	
1530	REFRESHMENTS & POSTERS		
1600	SESSION 5A: Earth Rover Spotlight Venue: Henderson LT	SESSION 5B: AI for Water Systems Venue: Conf Room 1+2	
1700	SESSION 6: EI @ Exeter: Past, present, future Venue: Henderson LT		
1730	SOCIAL EVENT - FREE BEER! Sponsored by Earth Rover Project!		

~60 talks & posters!

Parallel sessions:
Make a choice!

Detailed programme
available online

<https://tinyurl.com/2pnfuktc>



Housekeeping

Catering in XFi Atrium

- Lunch 1200-1300
- Tea/coffee/cake 1530-1600
- Social 1730 (thanks to Earth Rover Program!)
- Gluten-free available

View posters in breaks

Photos may be taken

No fire alarms expected

Be kind & respectful – enjoy the day!





Centre for Environmental Intelligence

Environmental solutions with AI and data science

Research Services · 161 followers · 11-50 employees



Andrew works here

Message

Following



Home

About

Posts

Jobs

People



Centre for
Environmental
Intelligence

161 followers

All

Images

Videos

Articles

Documents

Sort by: Top



Centre for Environmental Intelligence

161 followers

2d ·

We're excited to announce the official launch of the Centre for Environmental Intelligence (CEI) at the [University of Exeter](#)! ...more

LinkedIn



CEI LinkedIn



Sarah Spaul &
EI Network team

Thanks!



All the speakers,
panellists & chairs



Environmental
Intelligence CDT





University
of Exeter

Global Systems
Institute

Anticipating tipping points

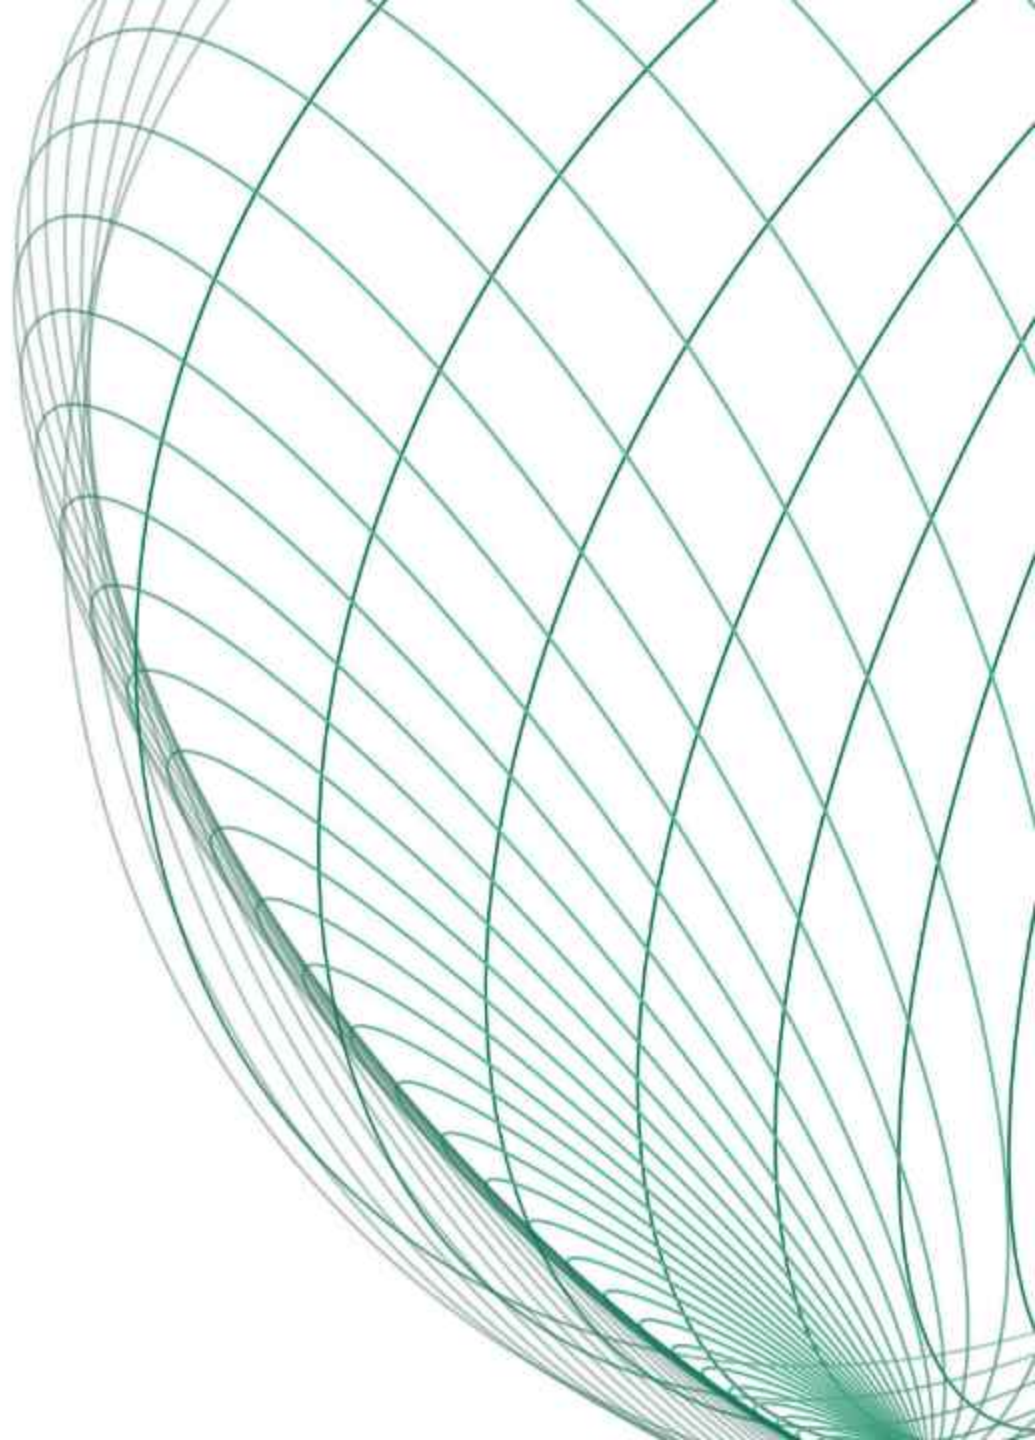
Tim Lenton

t.m.lenton@exeter.ac.uk

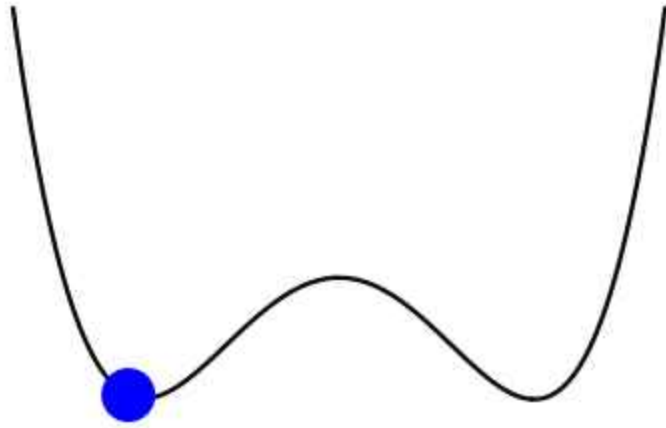
exeter.ac.uk/gsi



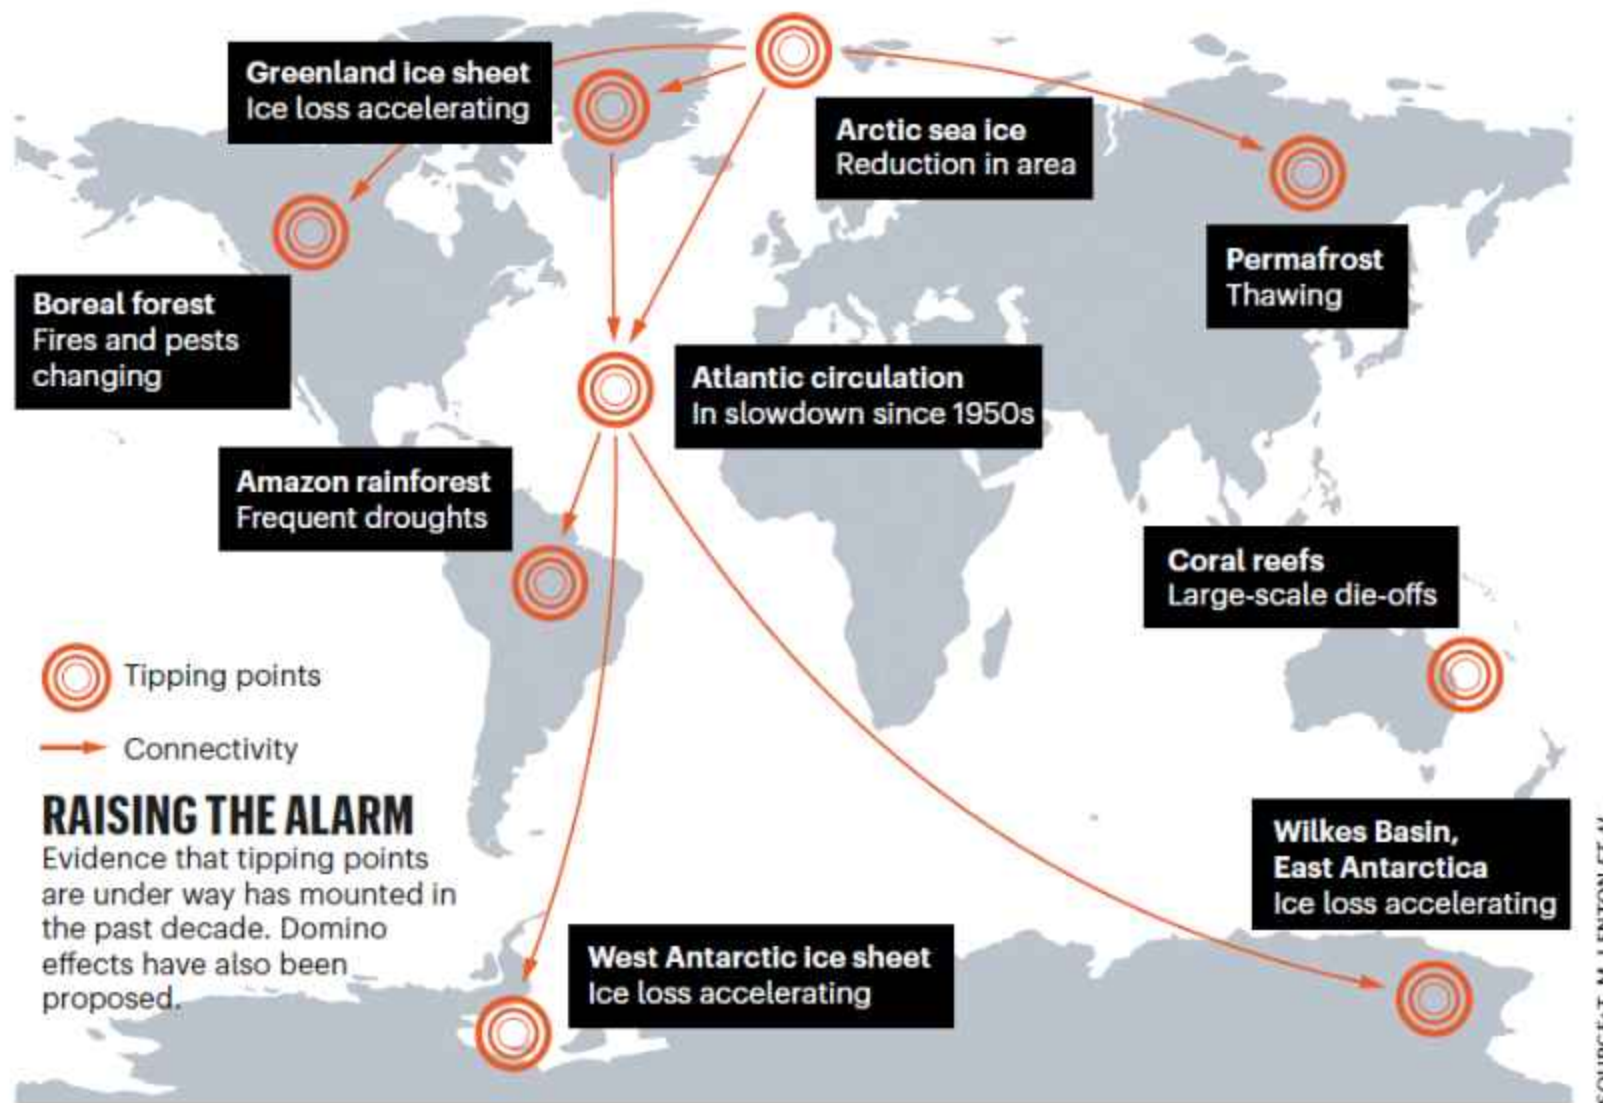
EARTH COMMISSION
GLOBAL COMMONS ALLIANCE



Passing a tipping point

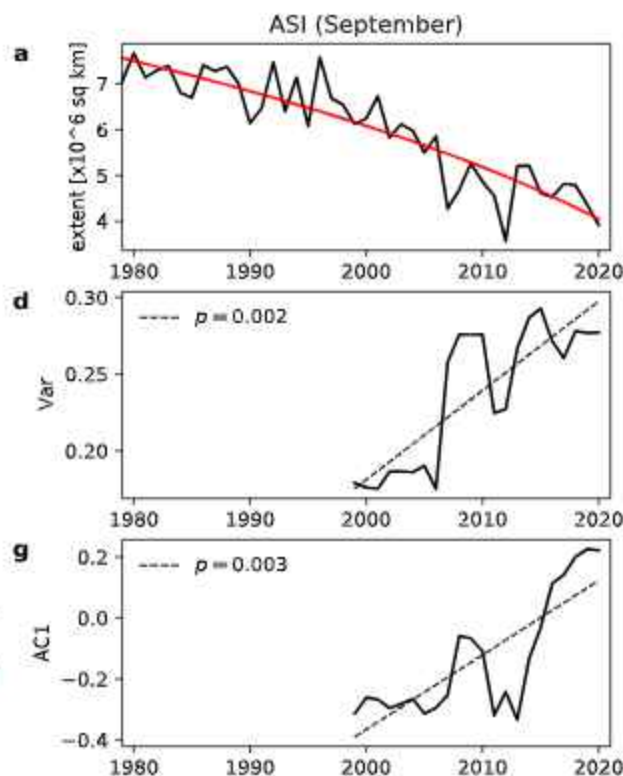


CC BY 4.0 - Chris A. Boulton (UoE, UK)

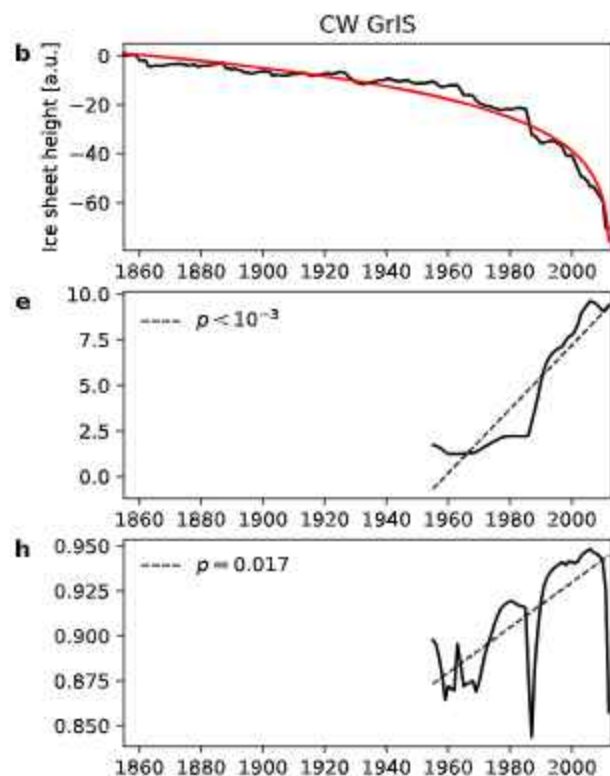


Early warning signals in observational data

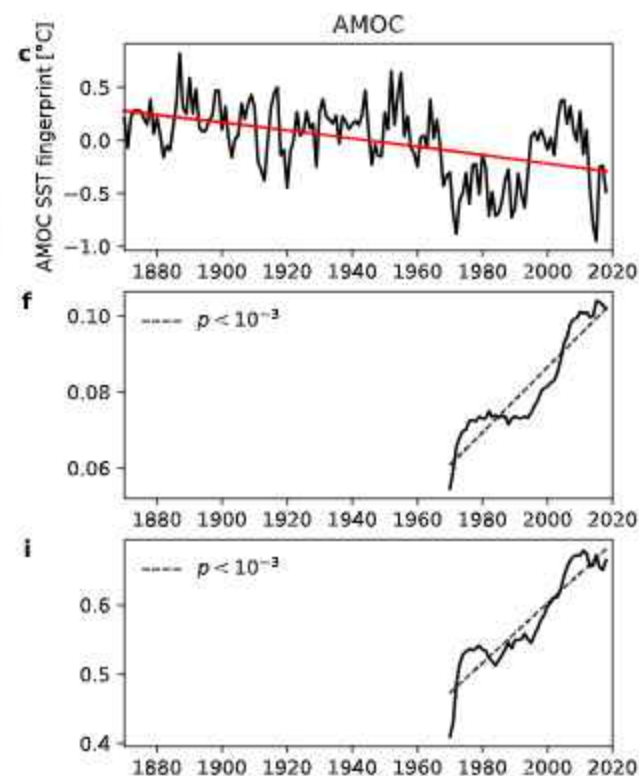
Arctic sea-ice
September extent



Central western
Greenland Ice Sheet



Atlantic Meridional
Overturning Circulation

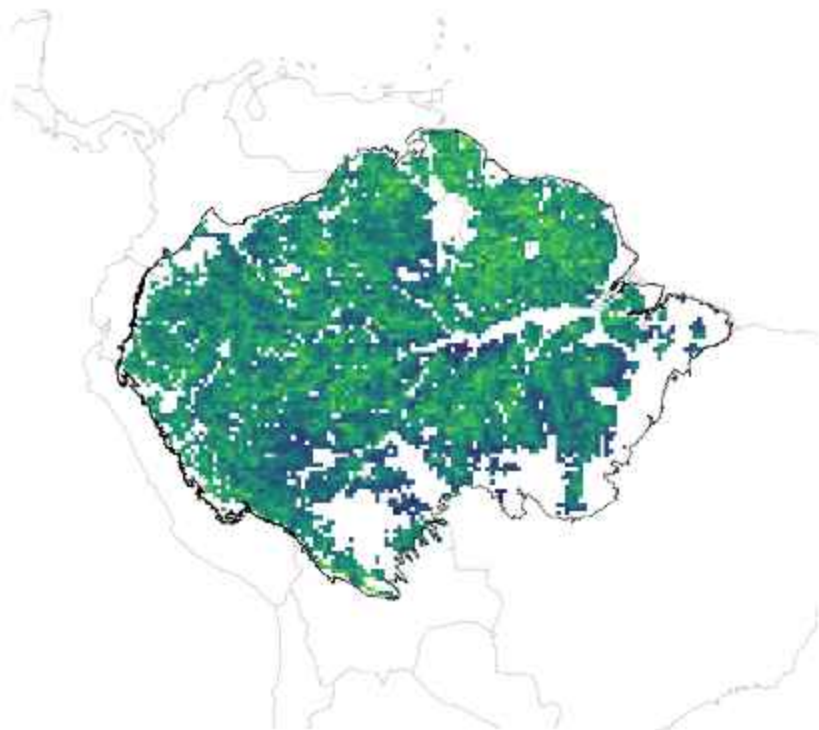


Variance

AR(1) auto-
correlation

Remotely sensing Amazon rainforest resilience

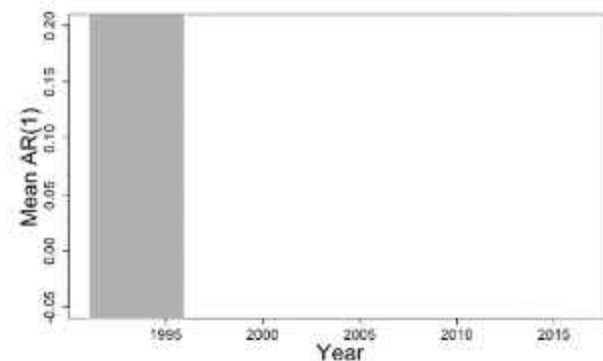
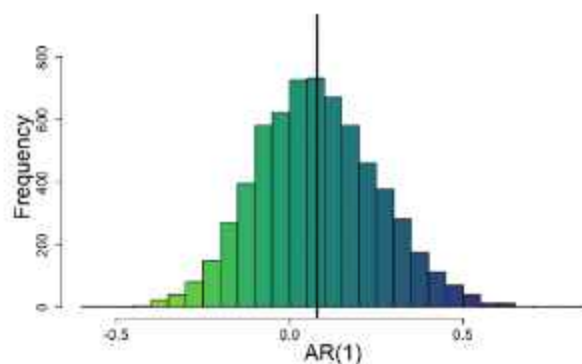
AR(1) of VOD



Key:

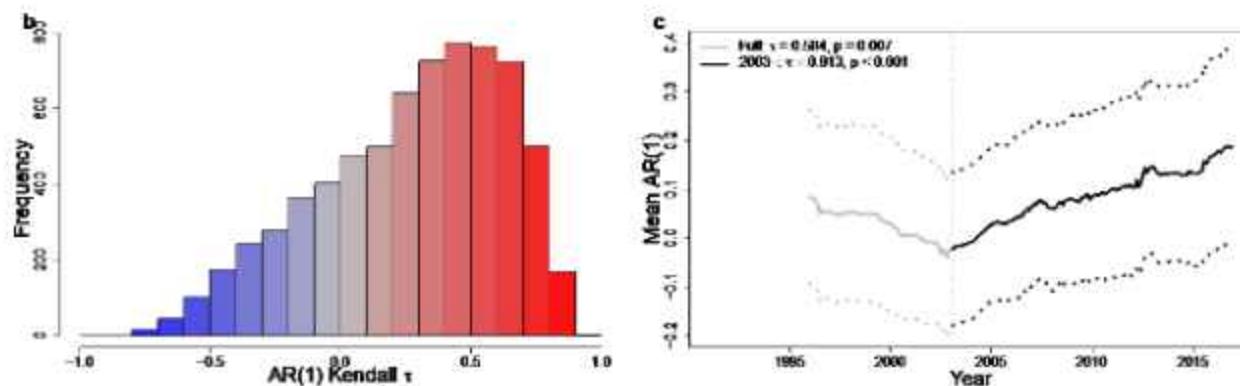
Paler = more resilient

Darker = less resilient

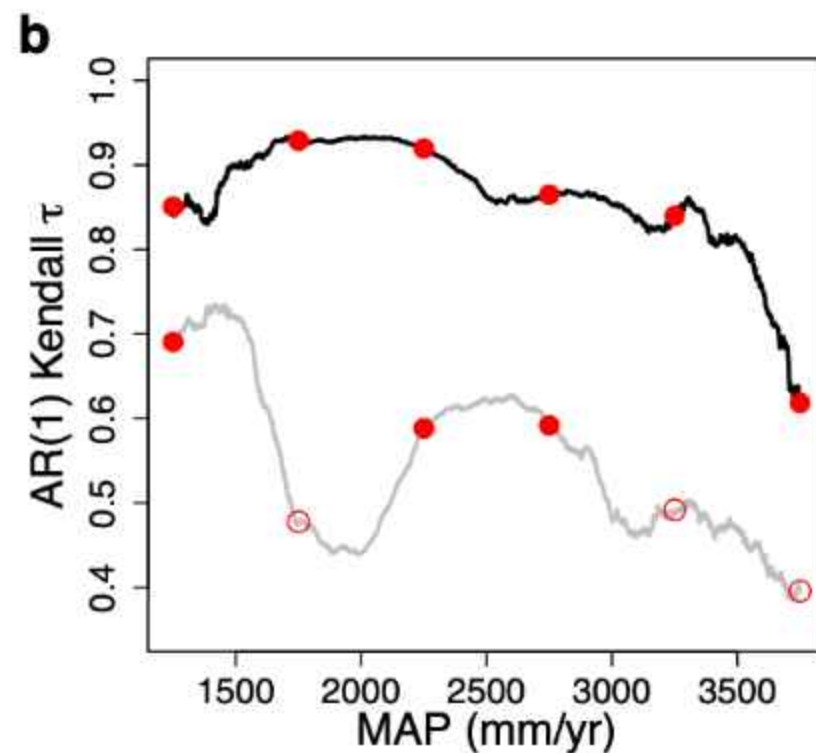
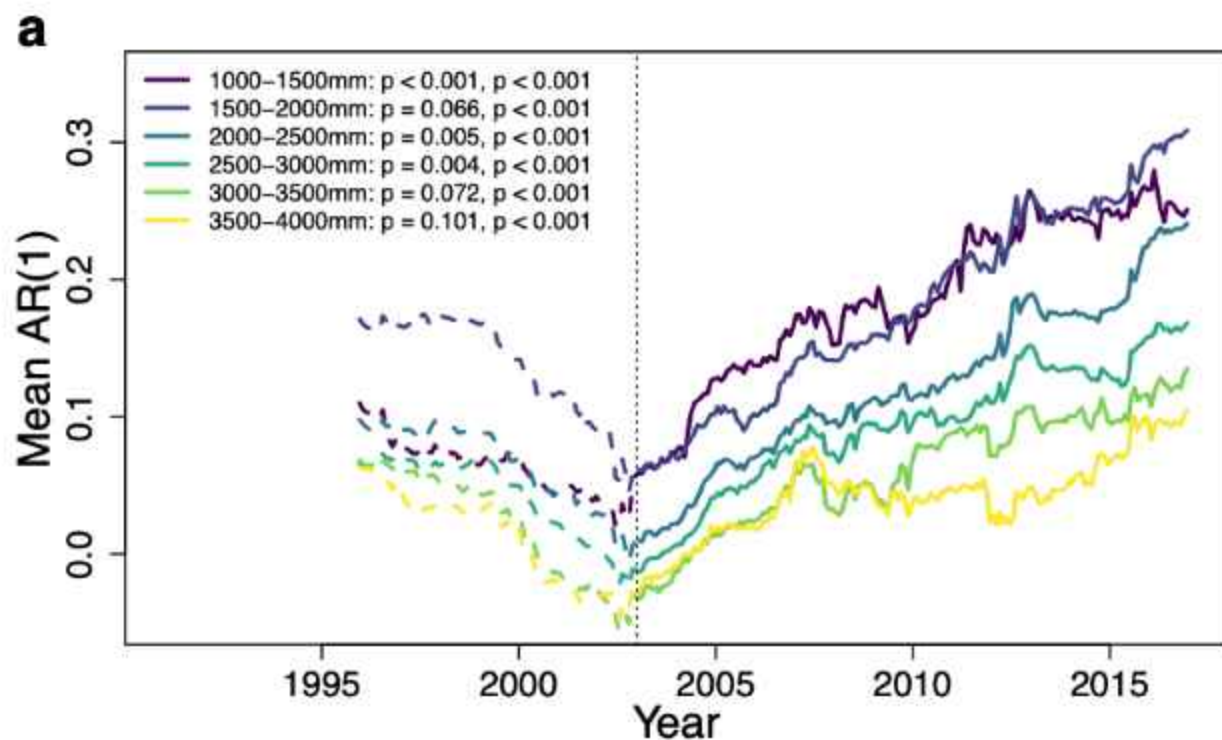


Loss of resilience of the Amazon rainforest

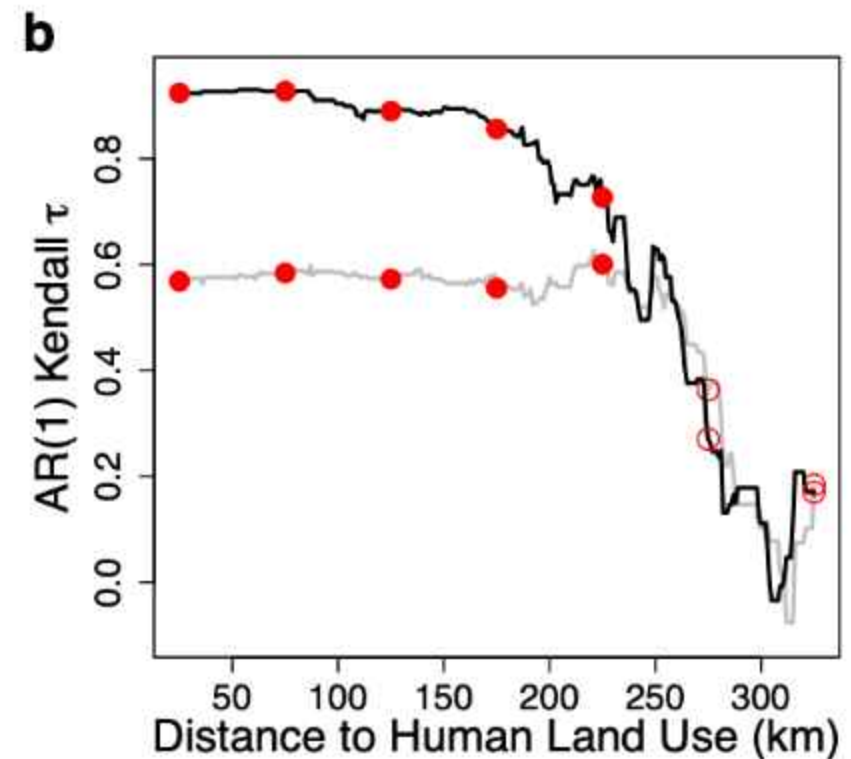
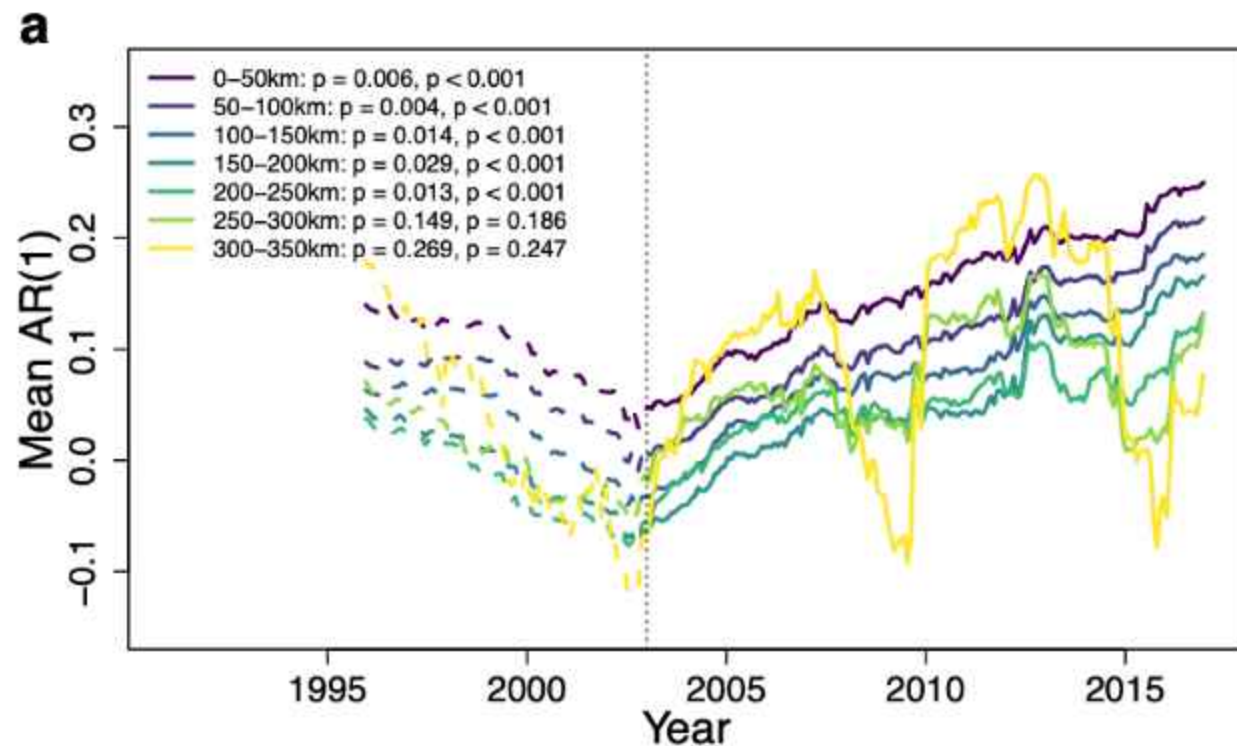
Trend in AR(1) of VOD since the early 2000s



Dependence on mean annual precipitation

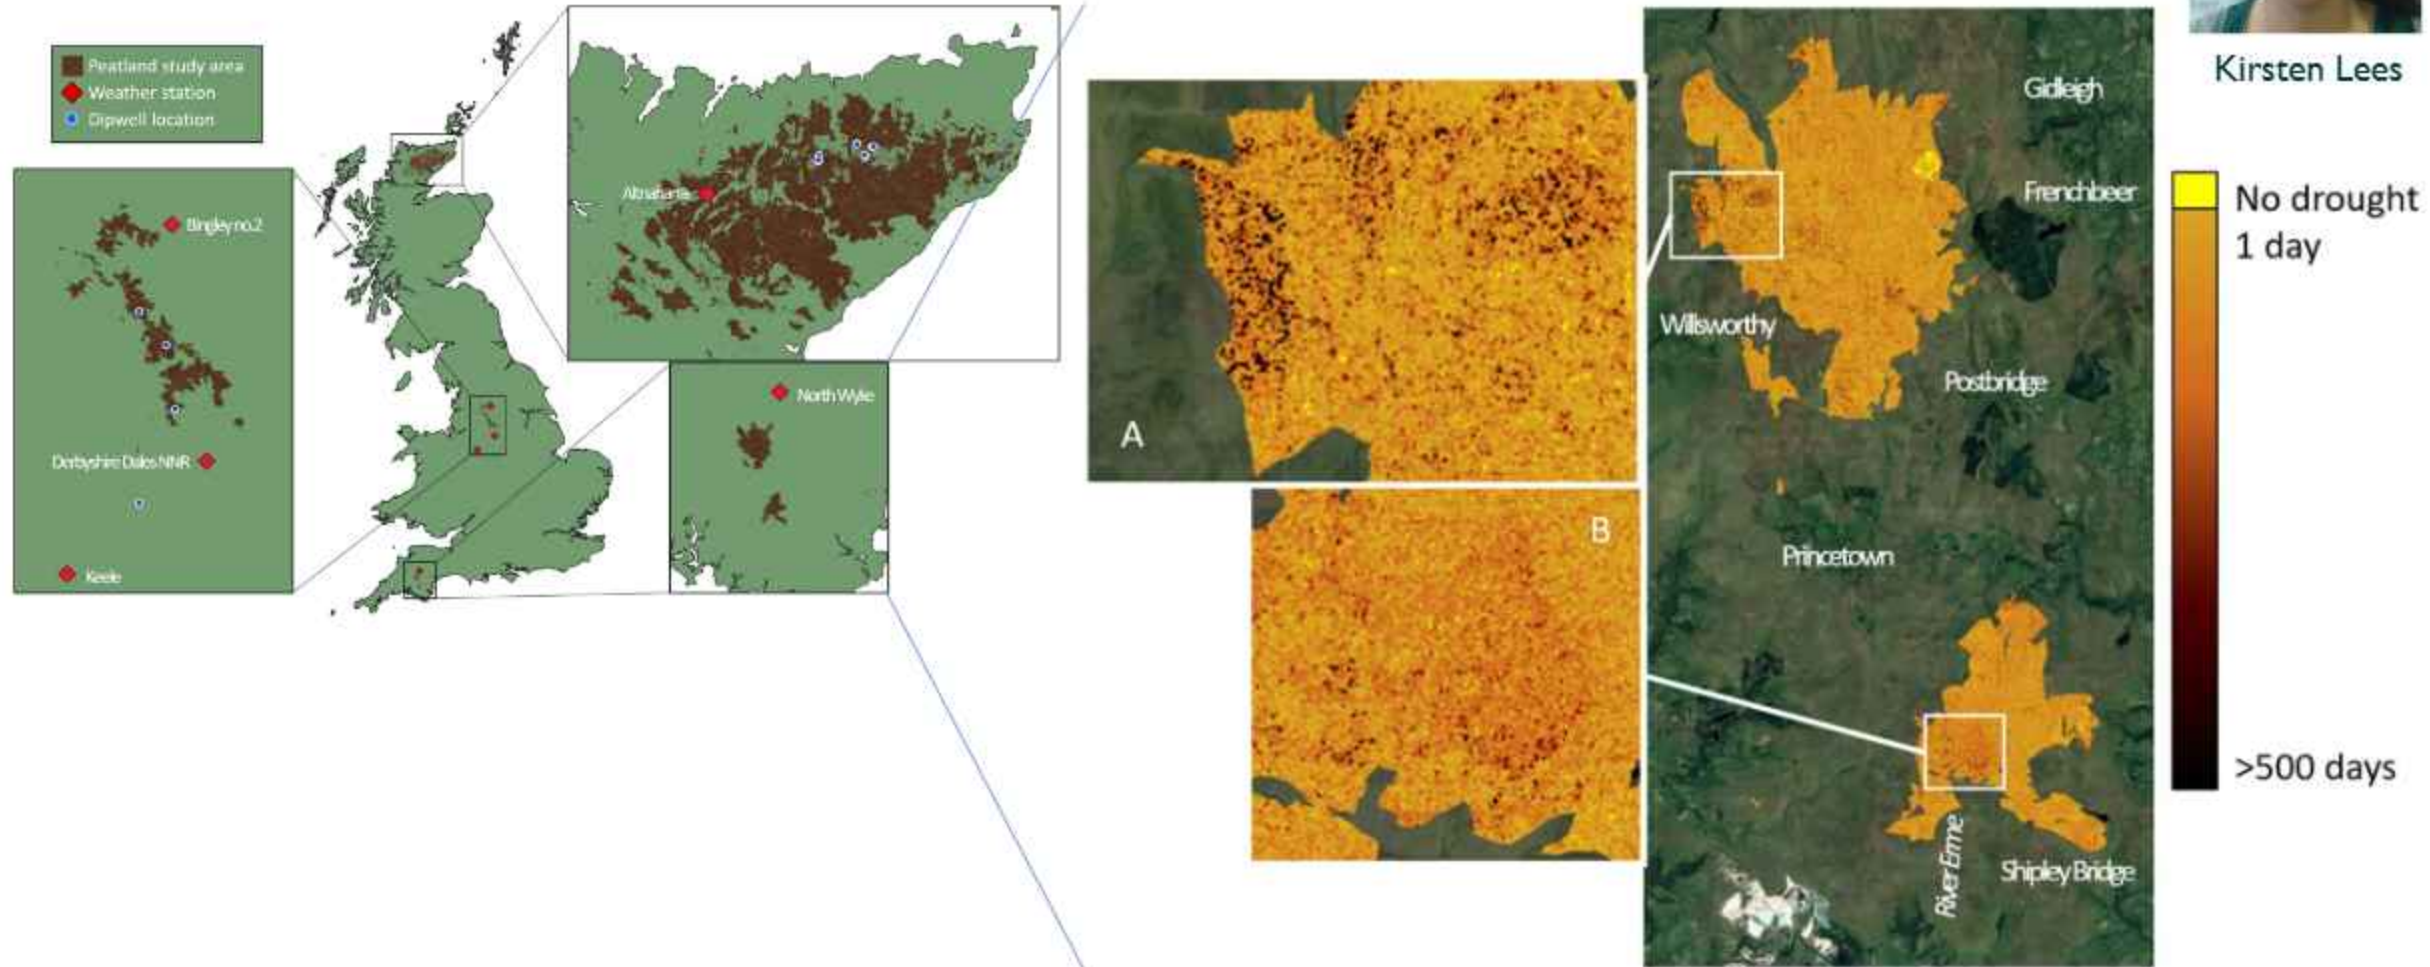


Dependence on distance to human activities



Sensing peatland resilience to drought

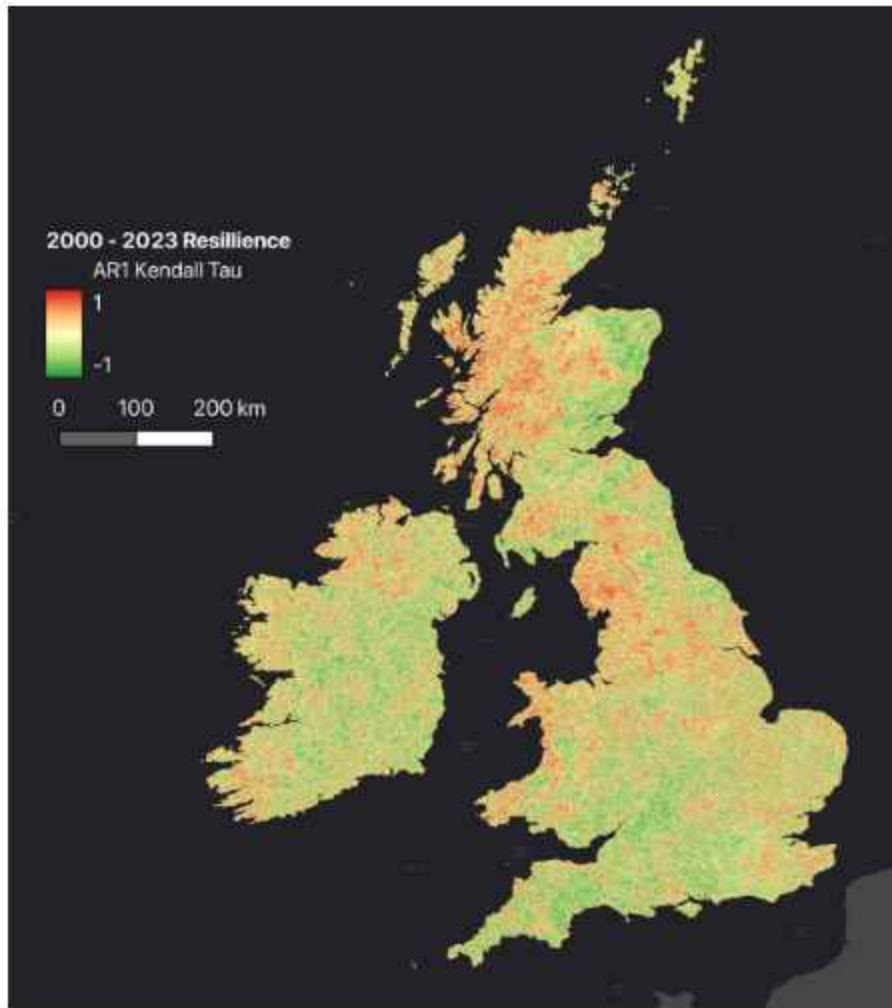
Water table recovery rate using Sentinel-1 SAR backscatter data



Kirsten Lees

Resilience sensing nature in the UK

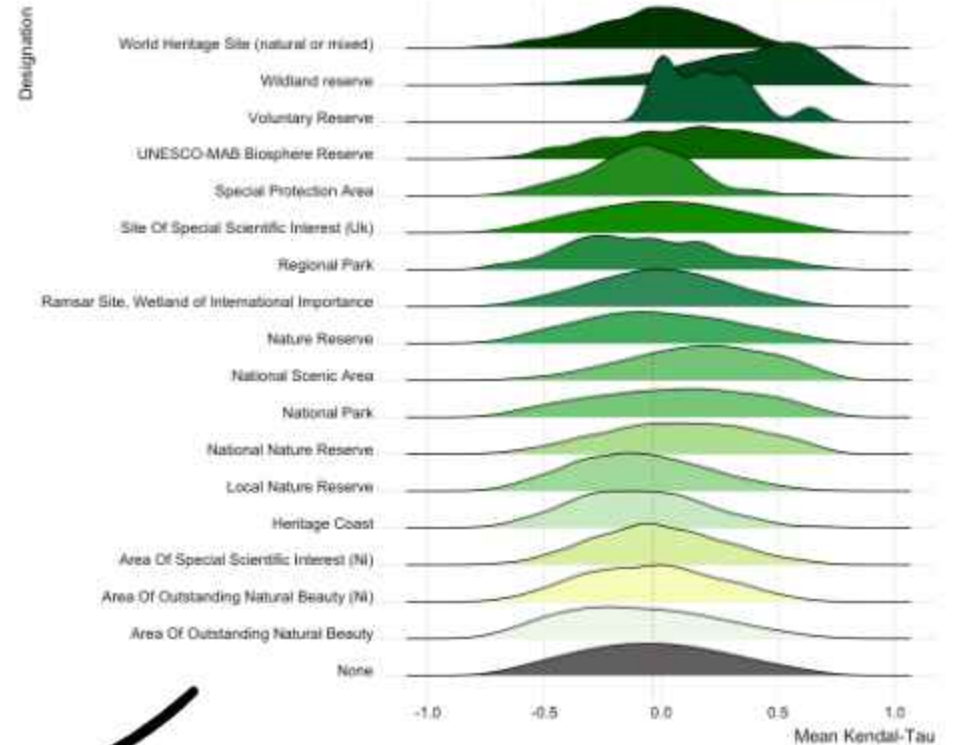
Daneen Cowling
EI PhD student



Resilience trends are spatially variable across the UK landscape and the network of protected sites

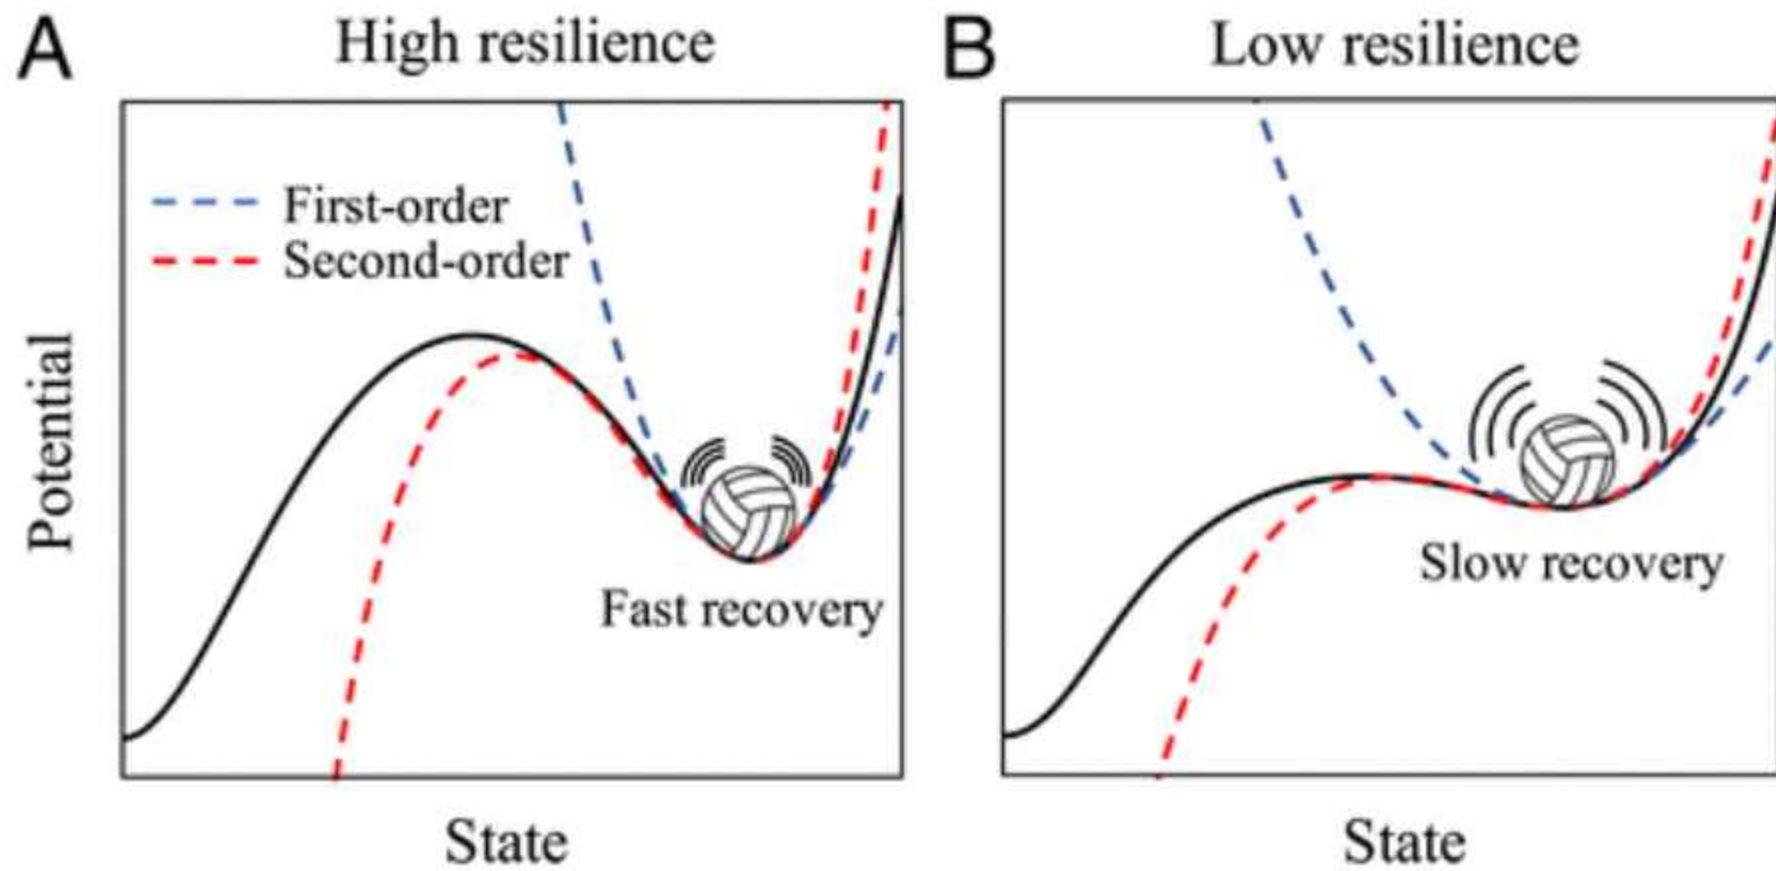


Using spectral signatures of different habitat behaviour (e.g., species rich grasses, heather moorlands, woodlands) we can track recovery time to perturbations (e.g., drought) to understand why some areas are more resilient than others

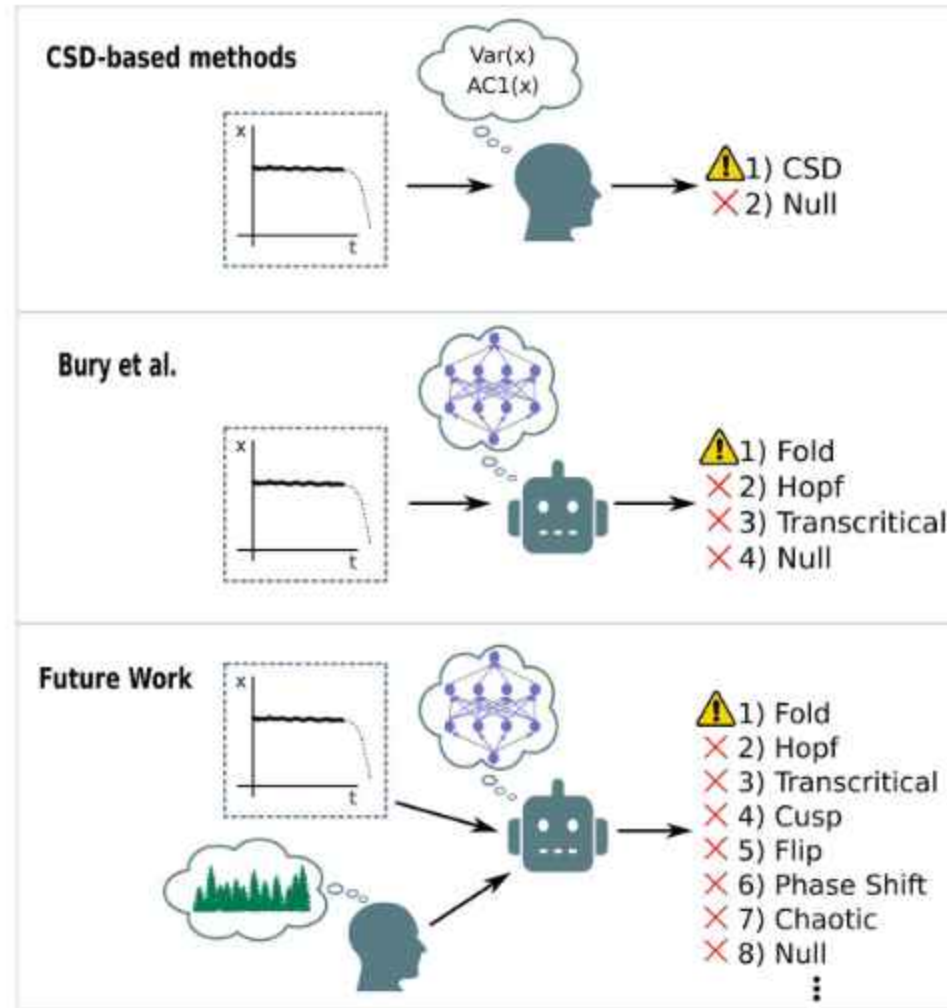


Averaged trends (Kendall Tau) of AR1 for NDVI for protected sites in the UK, relative to non-protected 2000-2023
Red line shows mean AR1 KT for non-protected

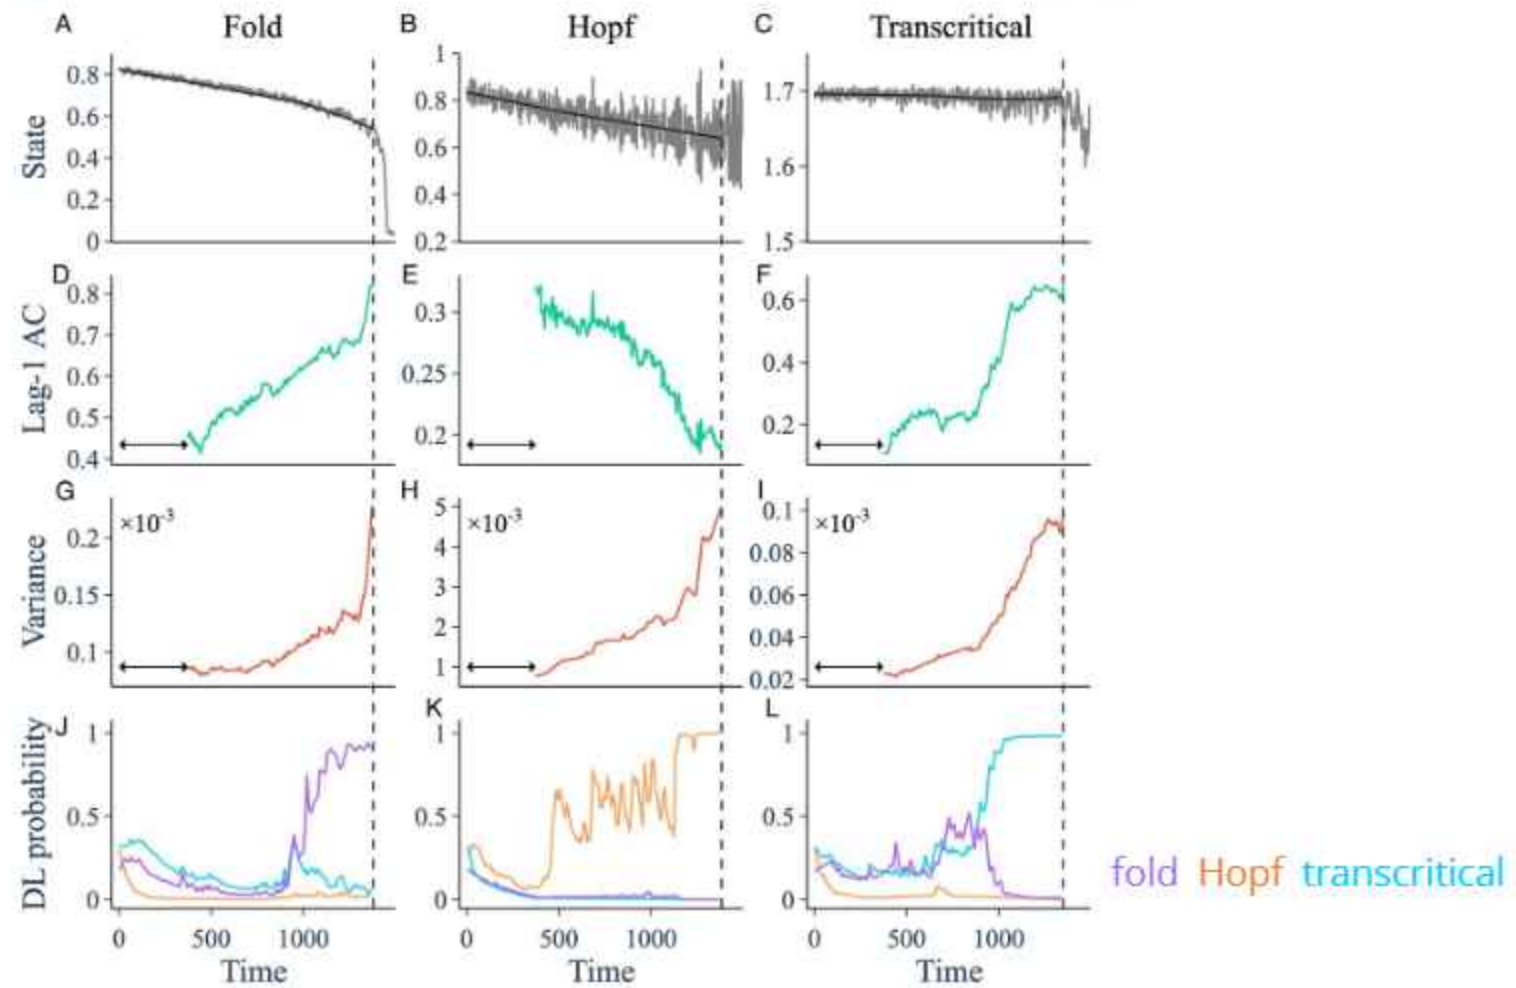
Can we get better early warning?



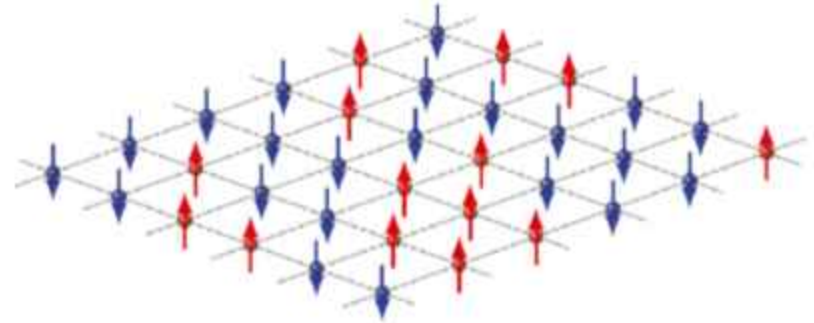
Train a deep learning algorithm



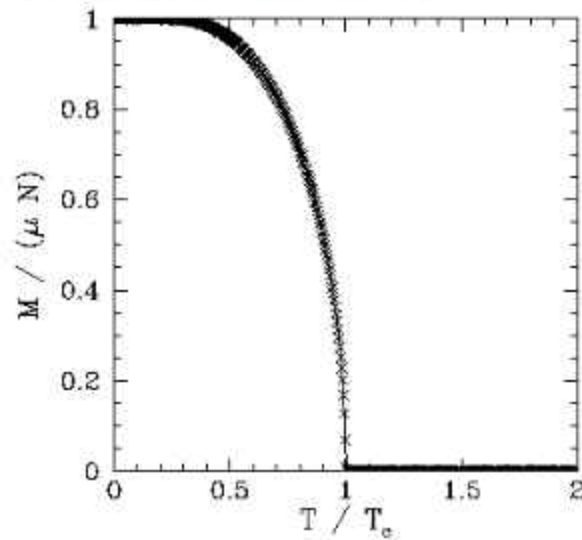
Deep learning performance versus AR(1), variance



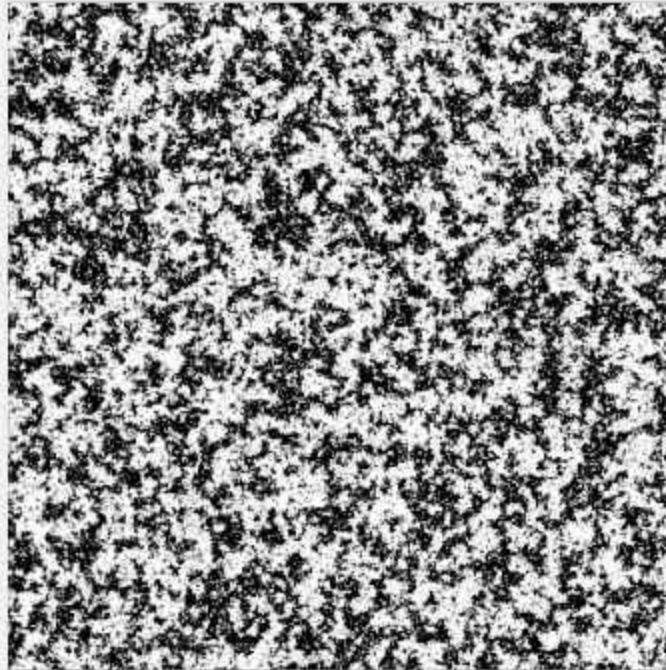
The Ising(-Lenz) model



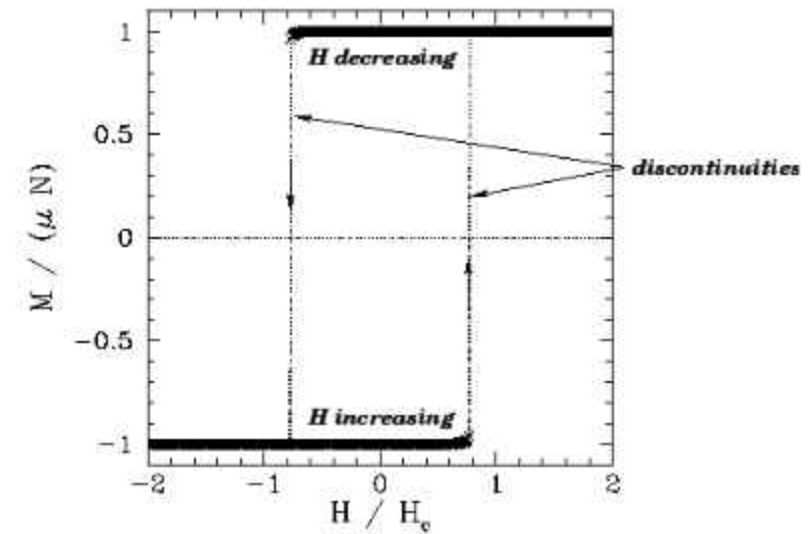
Second-order phase transition
(varying temperature)



Near the critical temperature



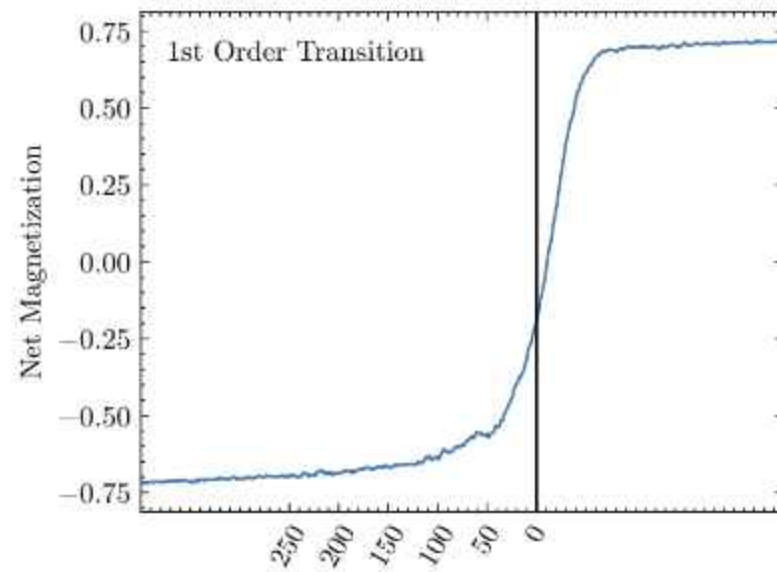
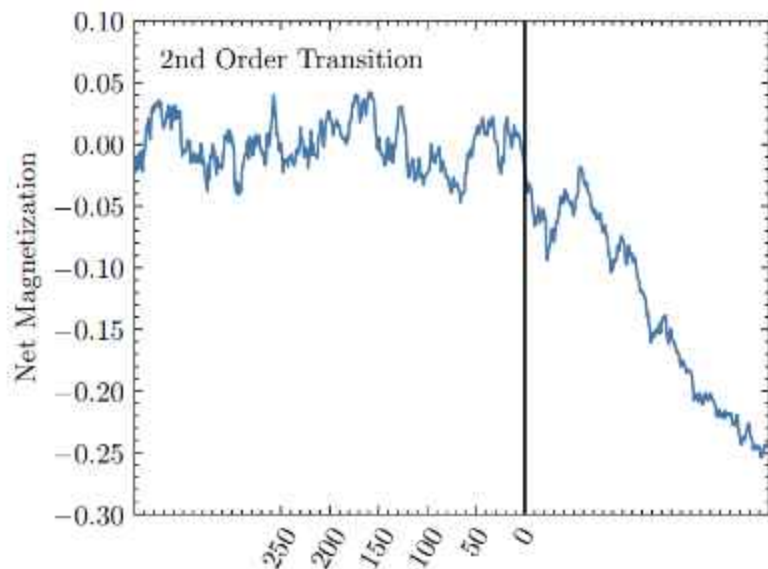
First-order phase transition
(varying external magnetic field)



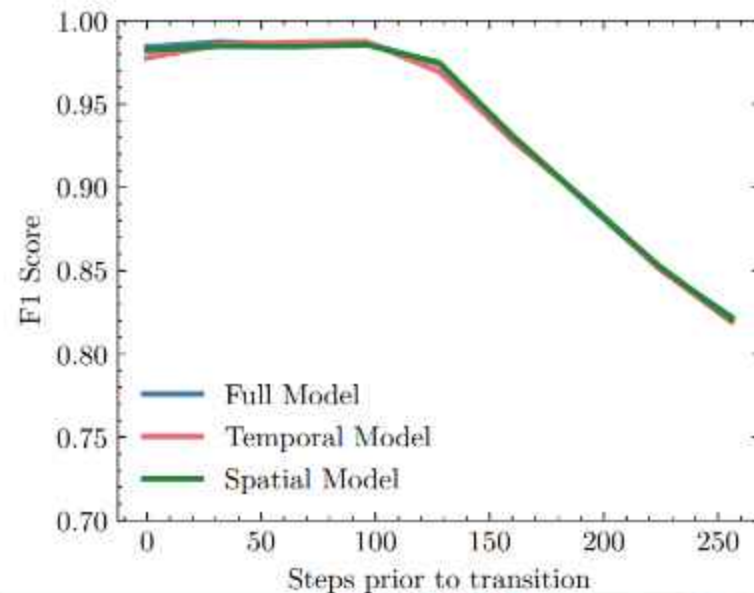
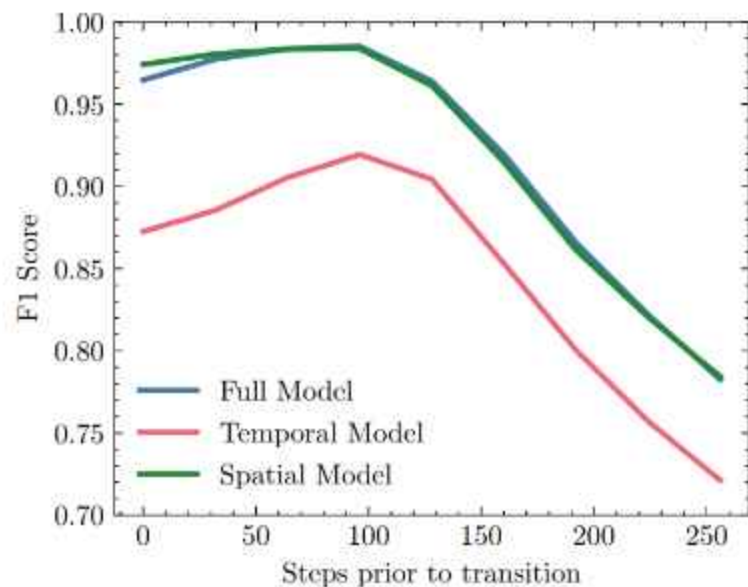


Daniel
Dylewsky

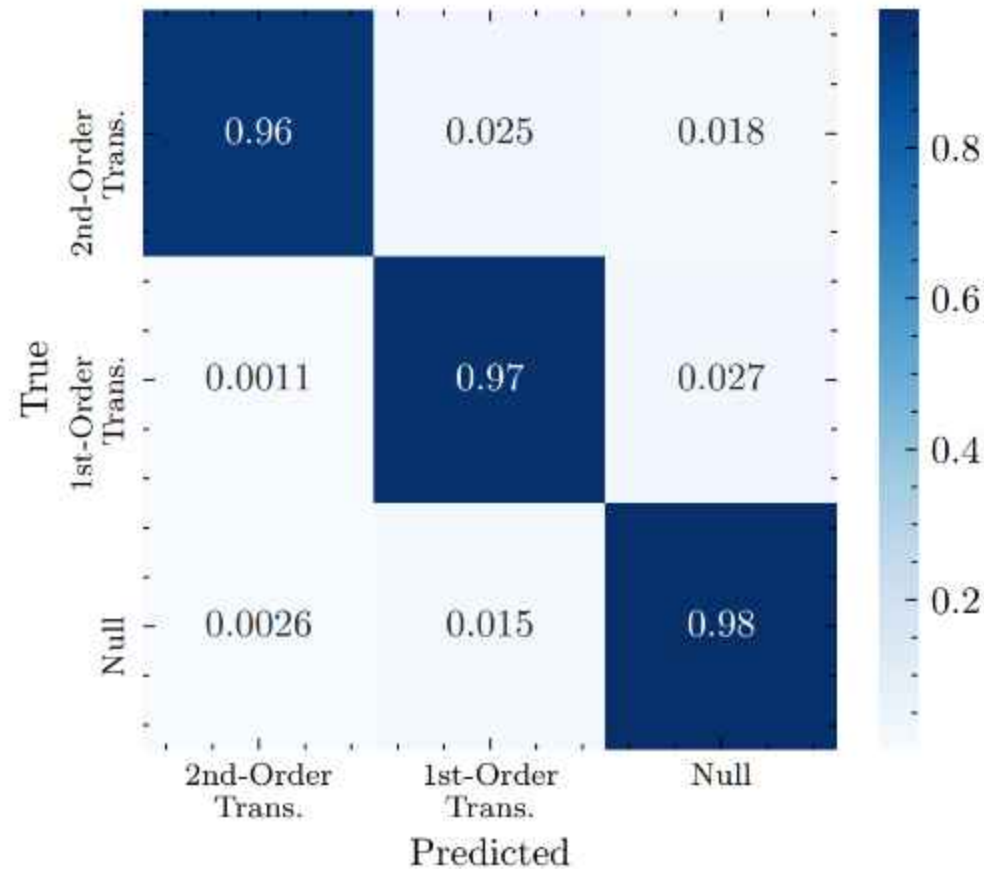
Example run



Model accuracy



Distinguish 1st & 2nd-order phase transitions



TIM LENTON



*how to fix
the climate crisis*

POSITIVE TIPPING POINTS

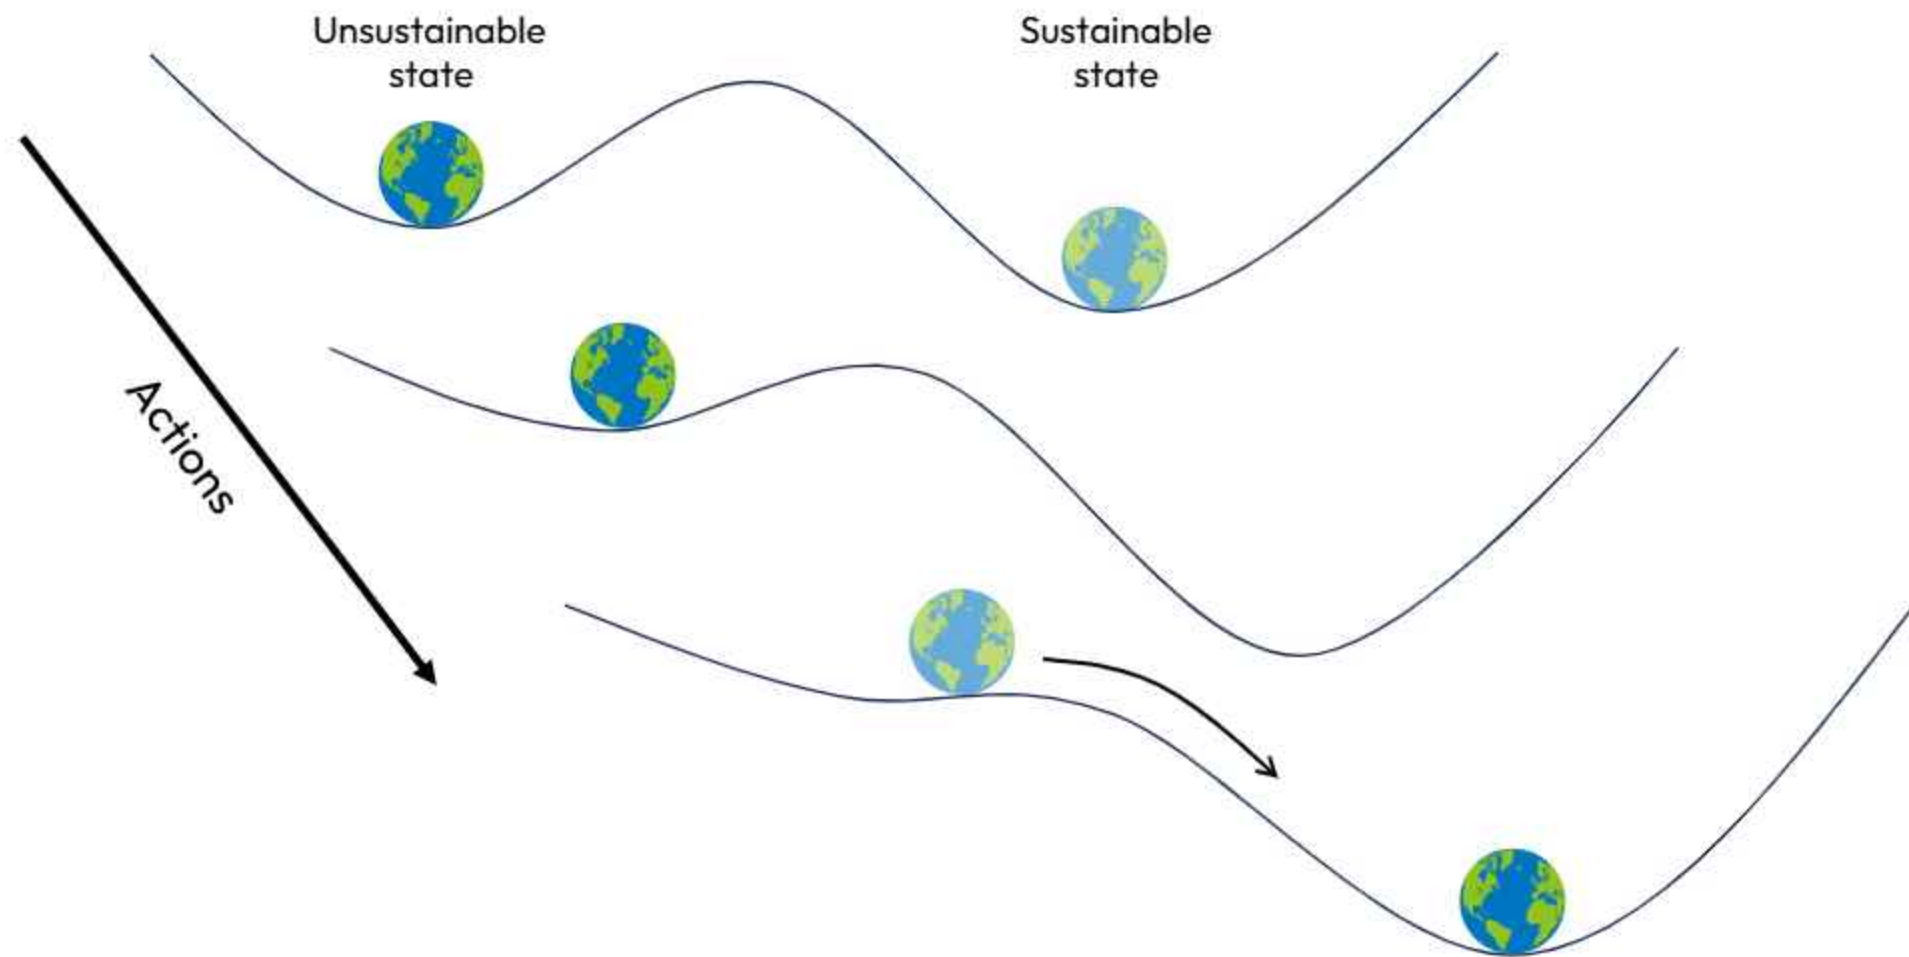
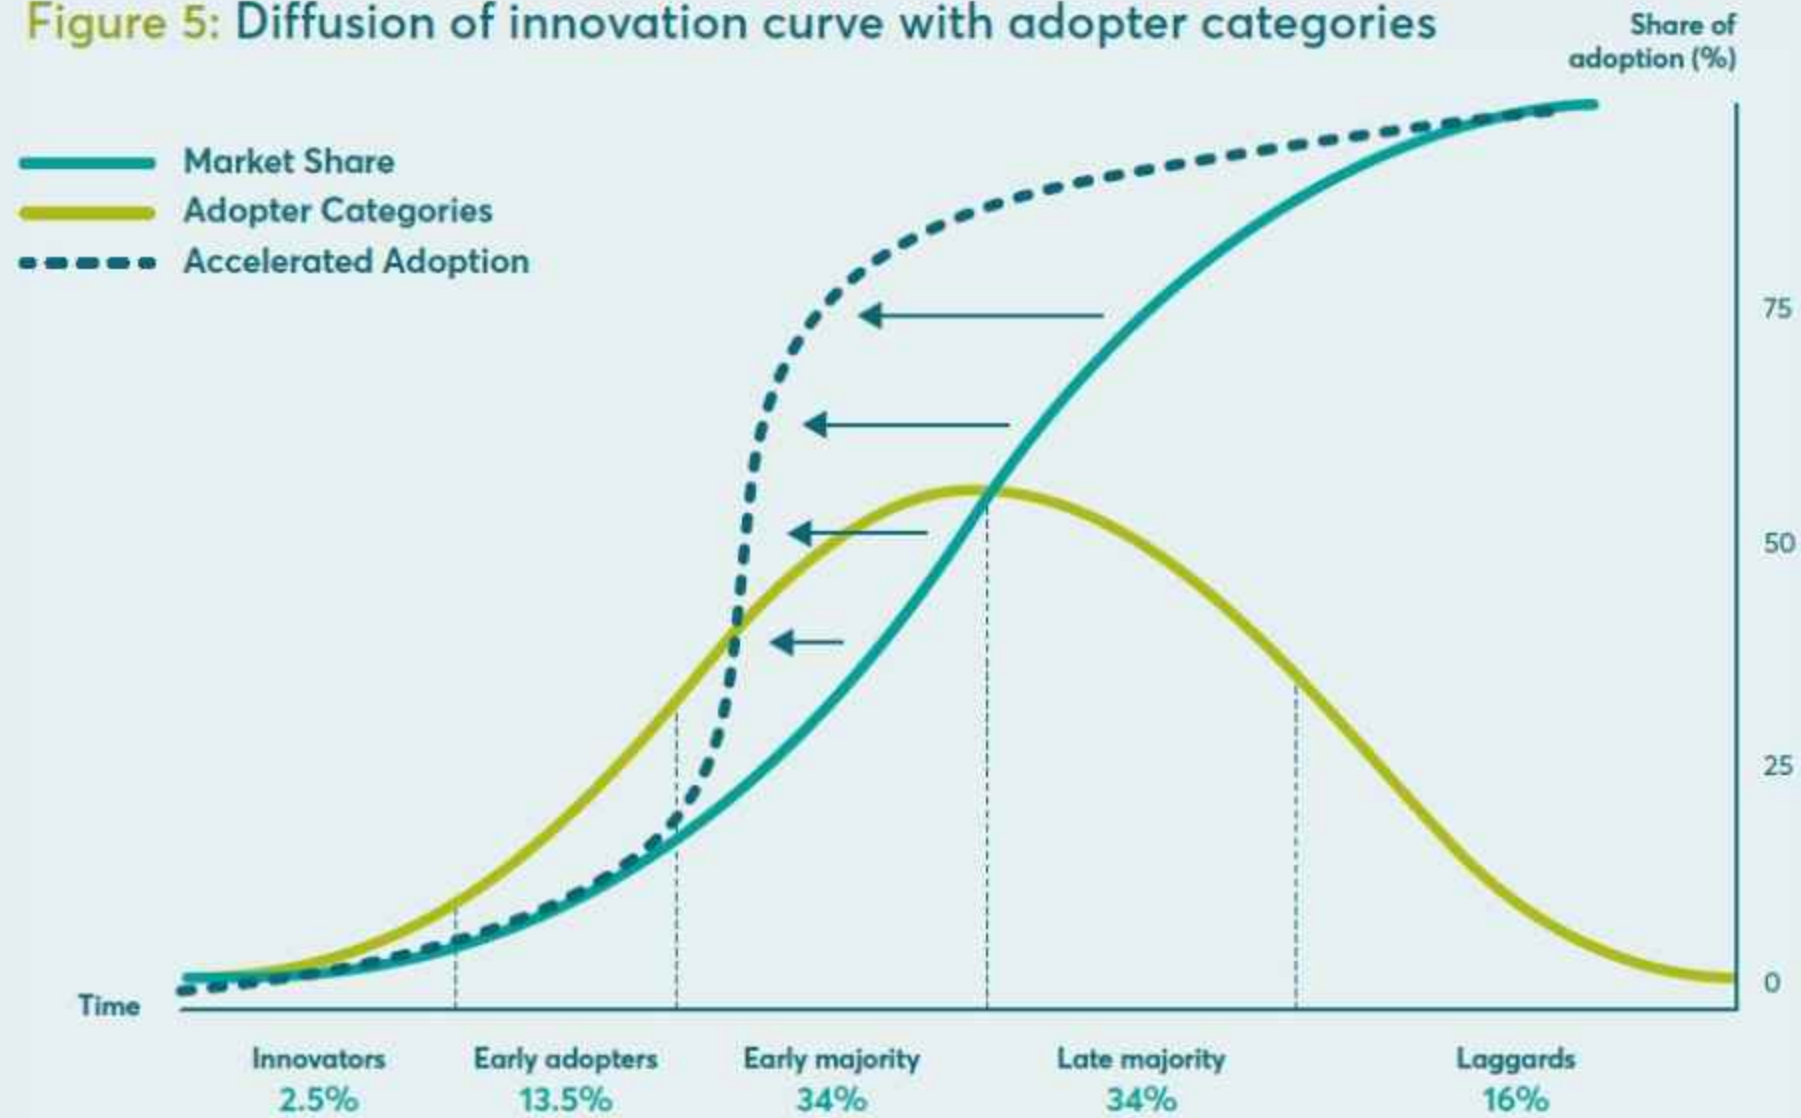
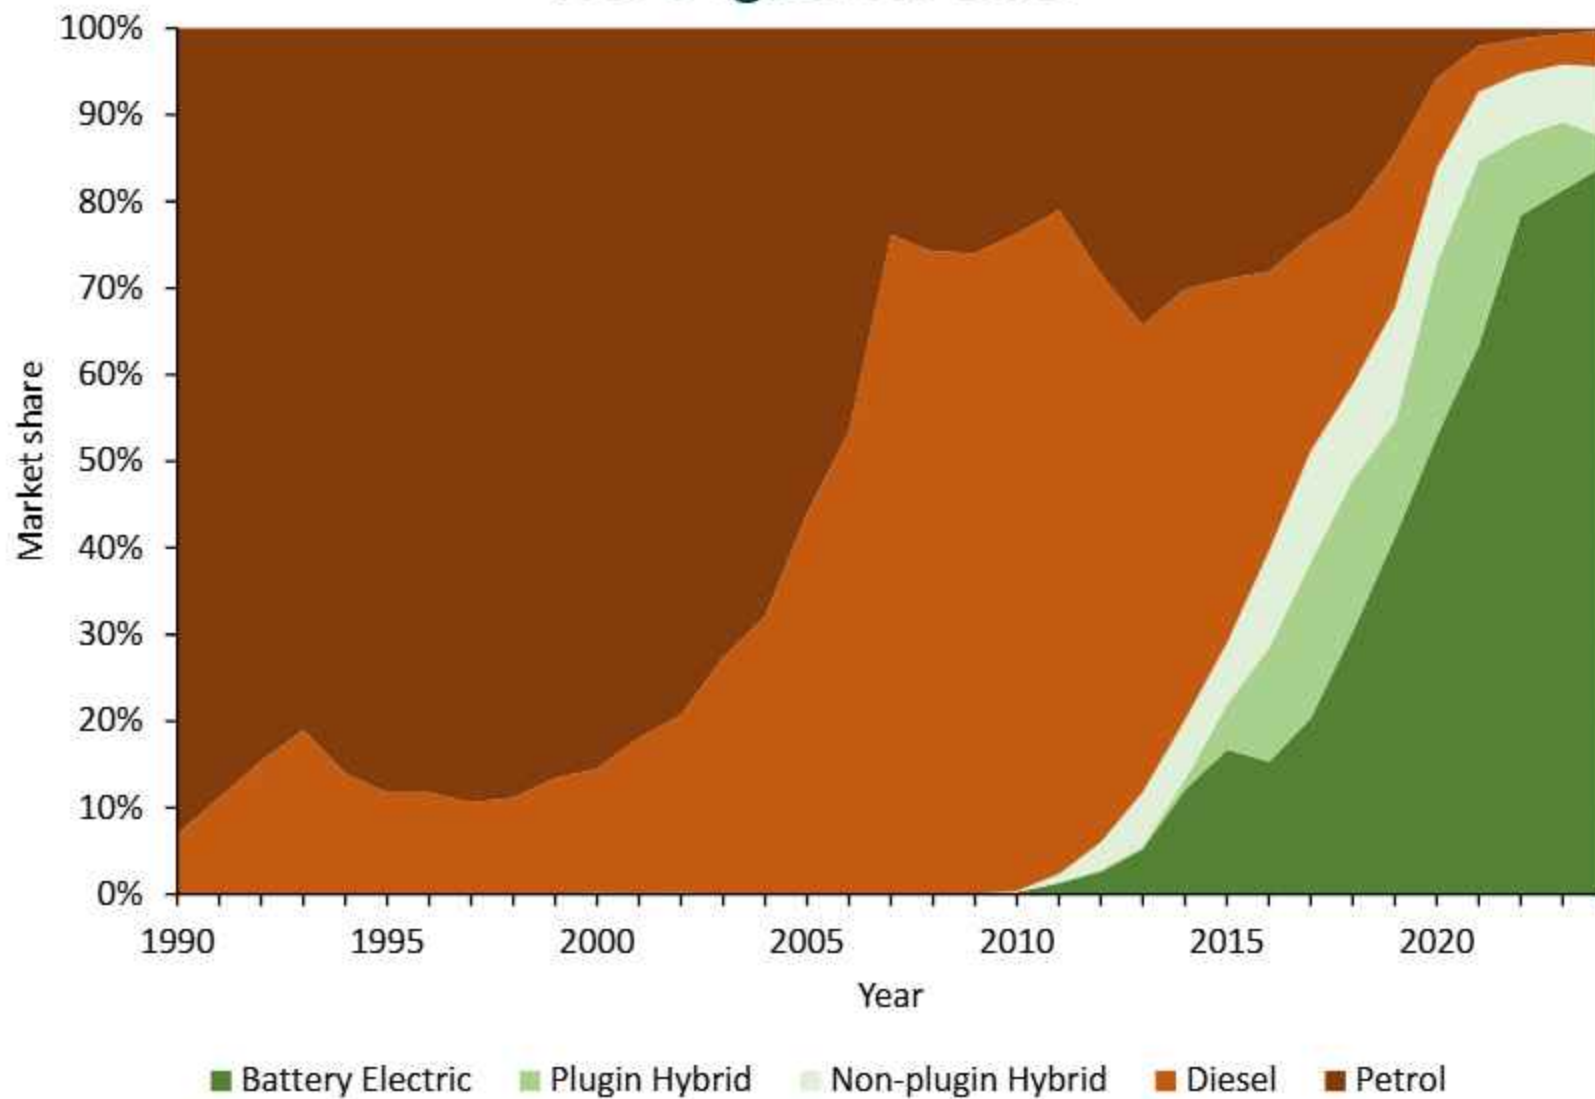


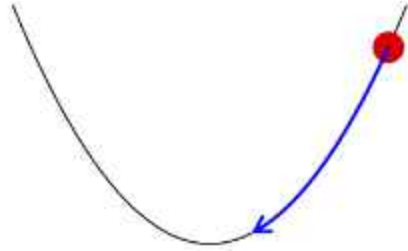
Figure 5: Diffusion of innovation curve with adopter categories



Norwegian car sales



Resilience = recovery rate
following a perturbation

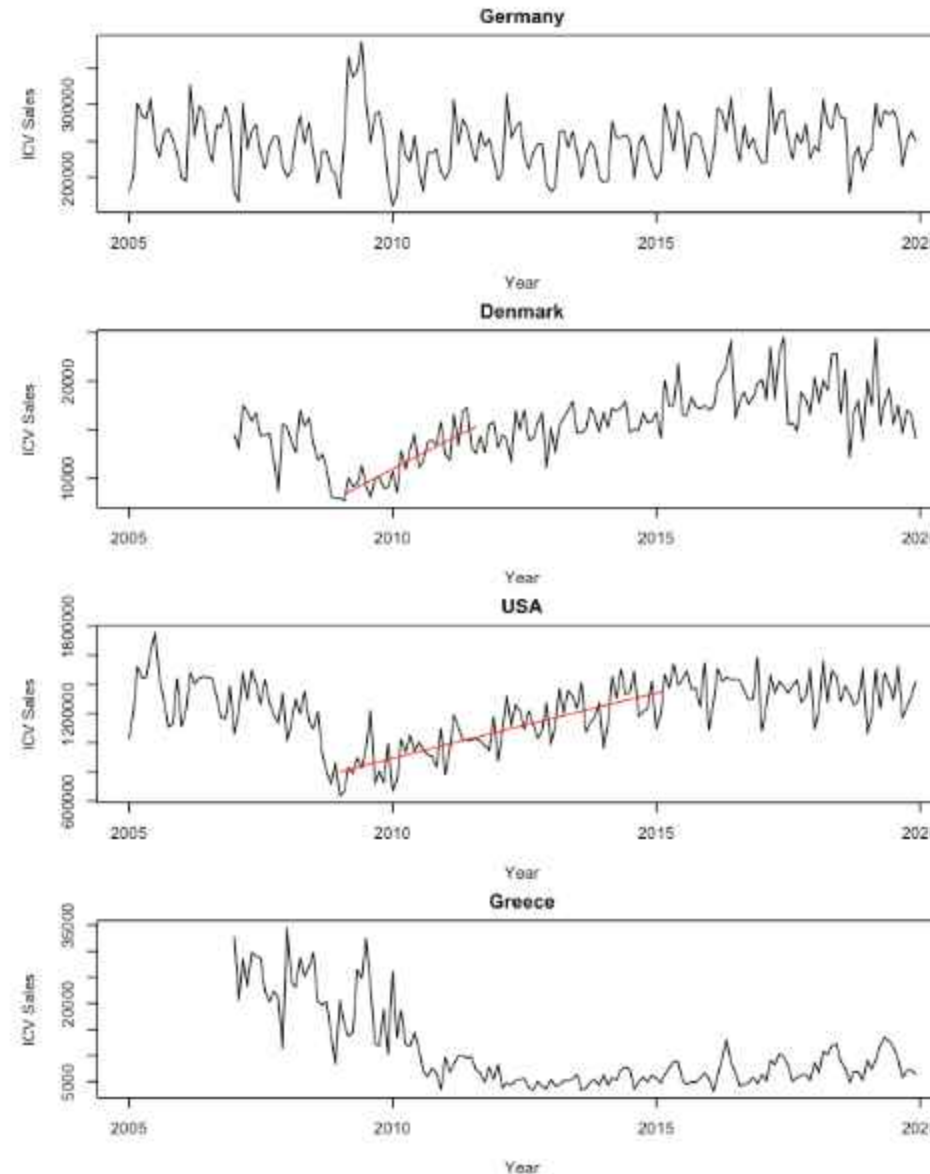


Fast recovery => High resilience



Slow recovery => Low resilience

Early opportunity signal
to trigger a tipping point

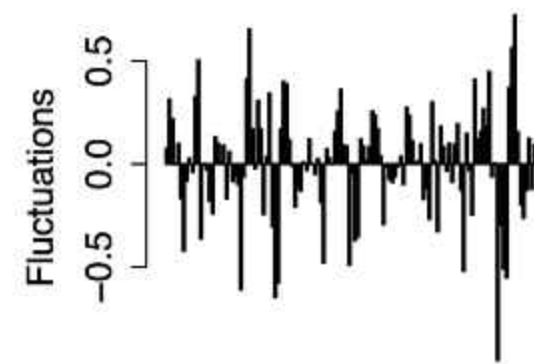
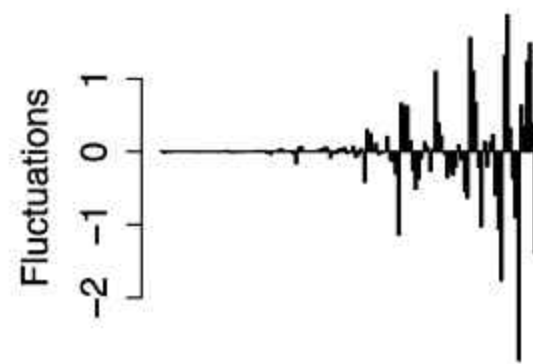
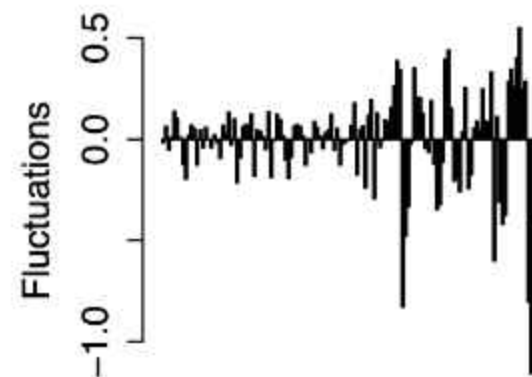
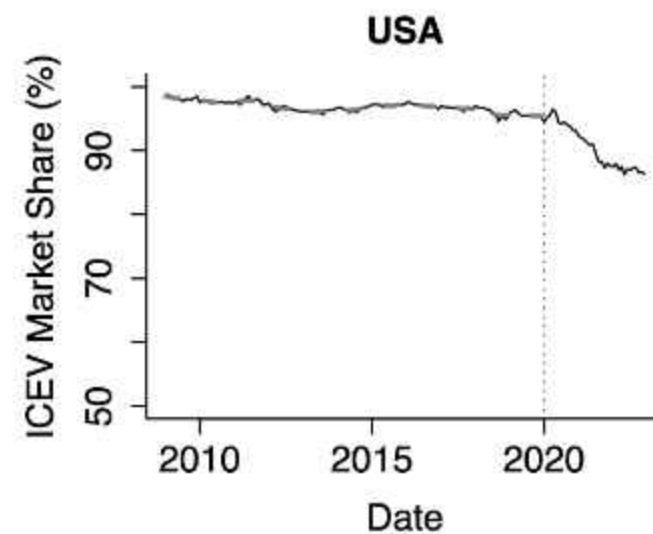
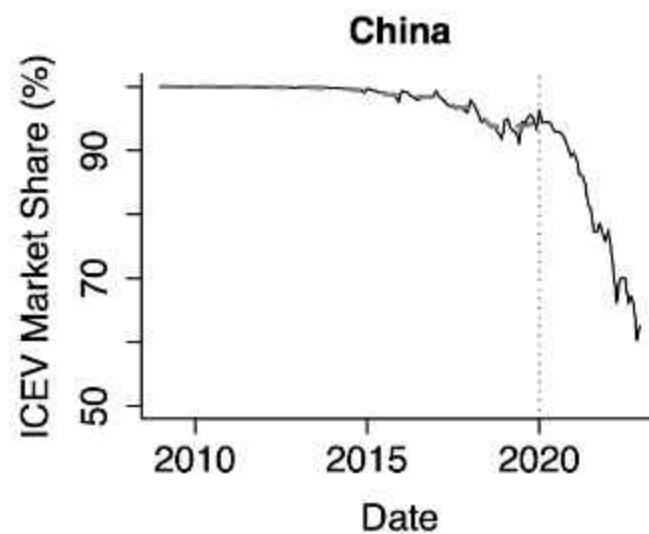
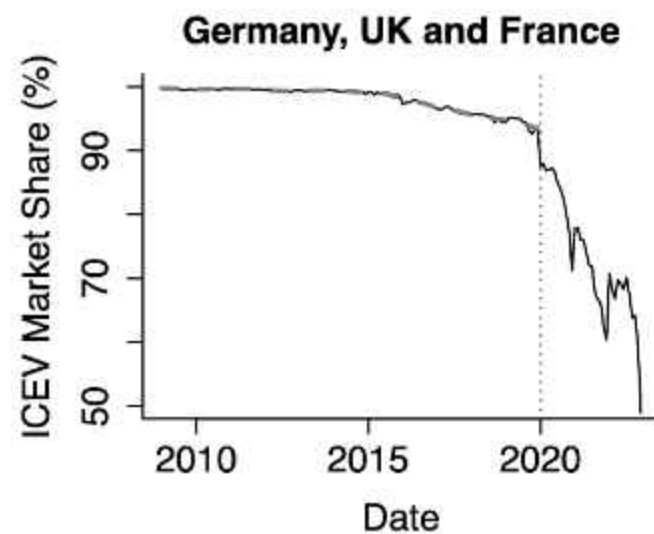


No decline in sales (due
to Government policy)

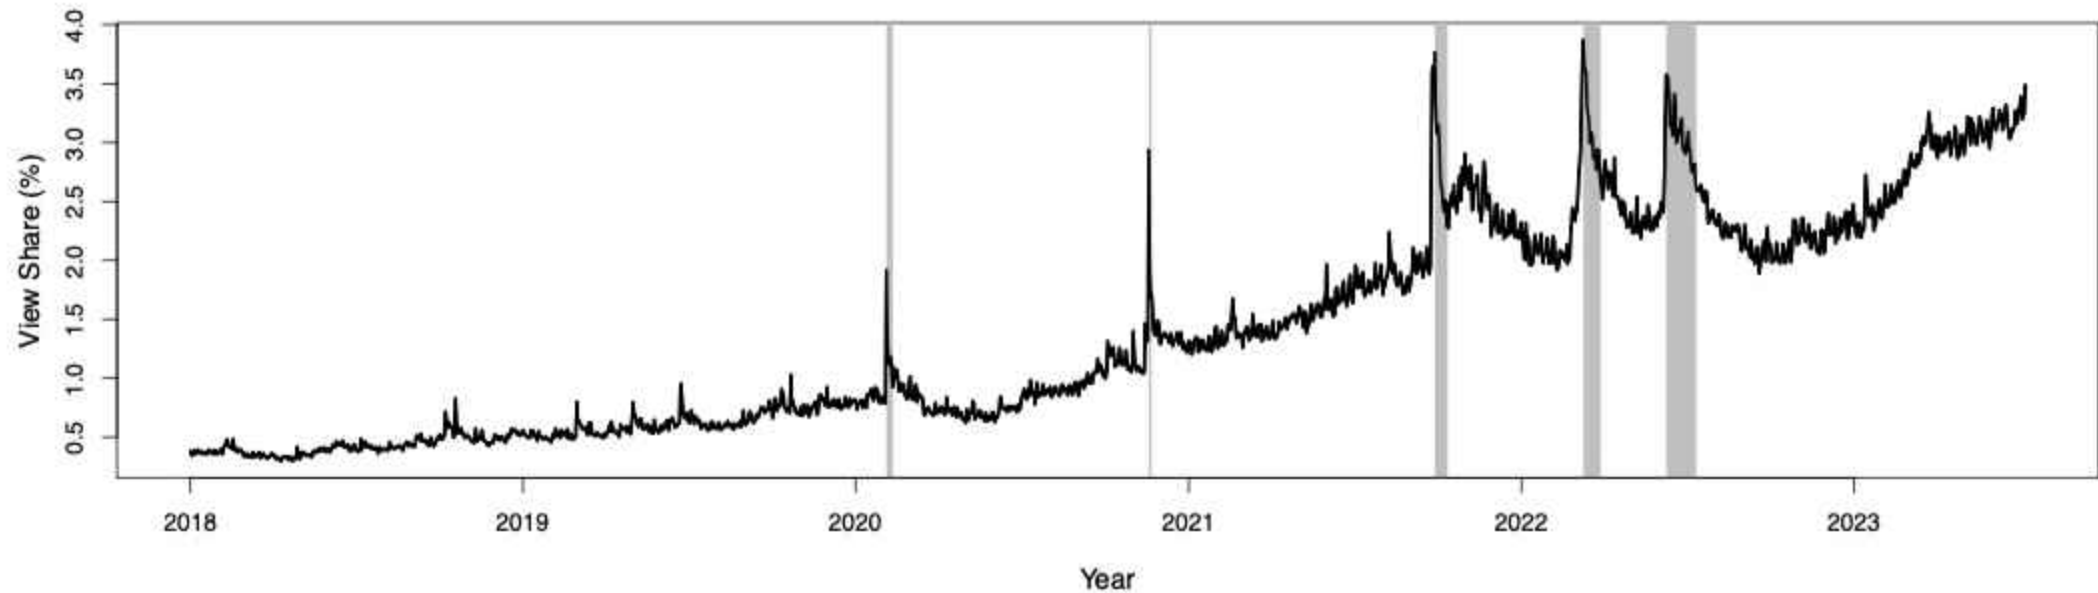
Decline in sales,
followed by recovery

Decline in sales,
slower recovery

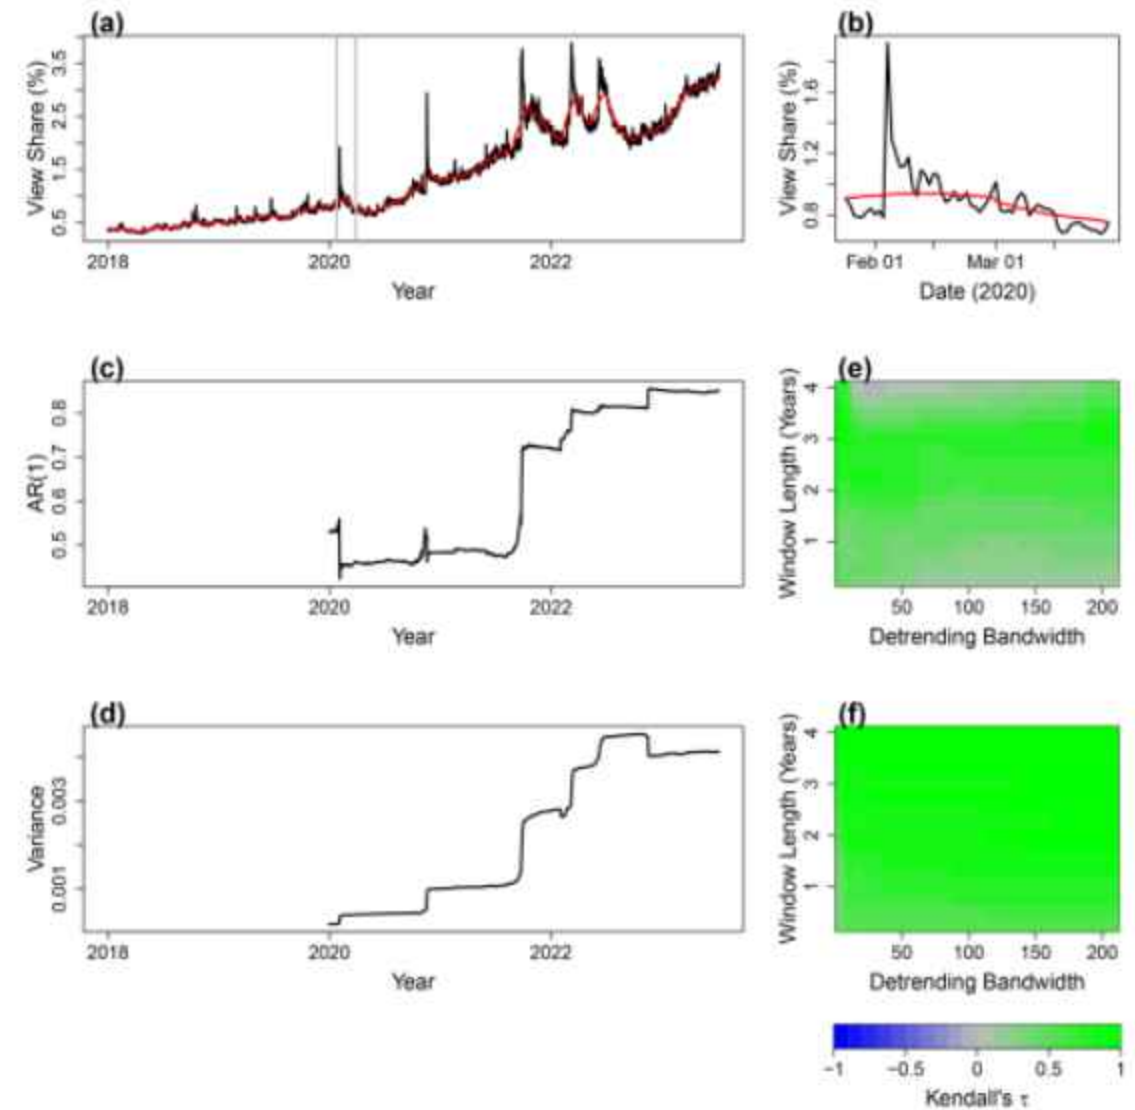
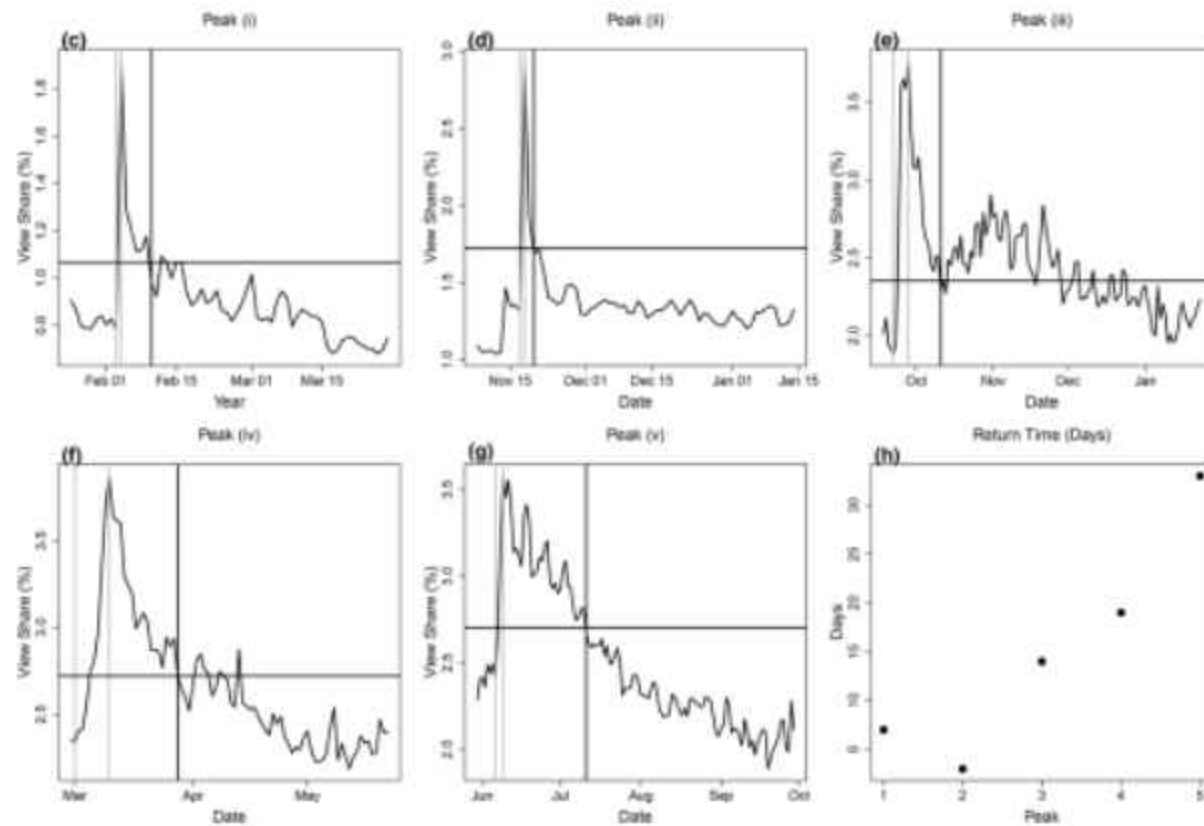
Transition to a lower
sales regime



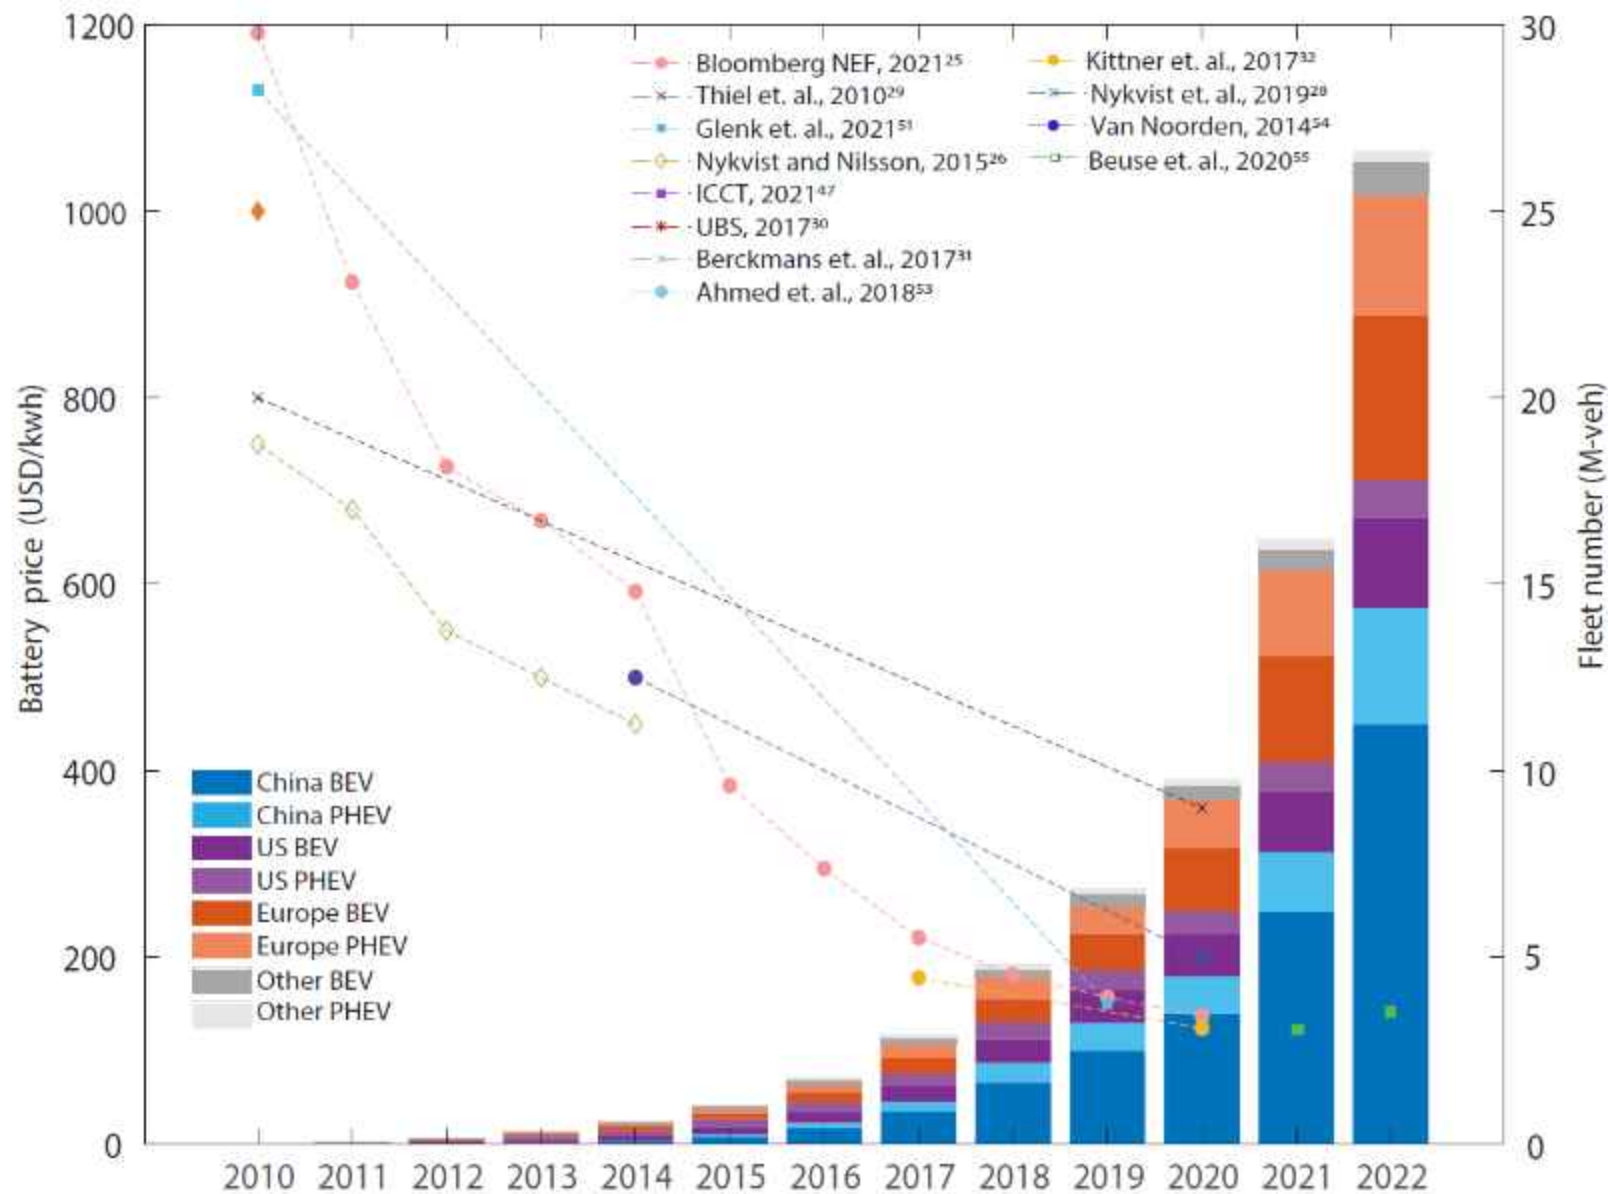
Electric vehicle advert view share (Autotrader, UK)



Early Opportunity Signals

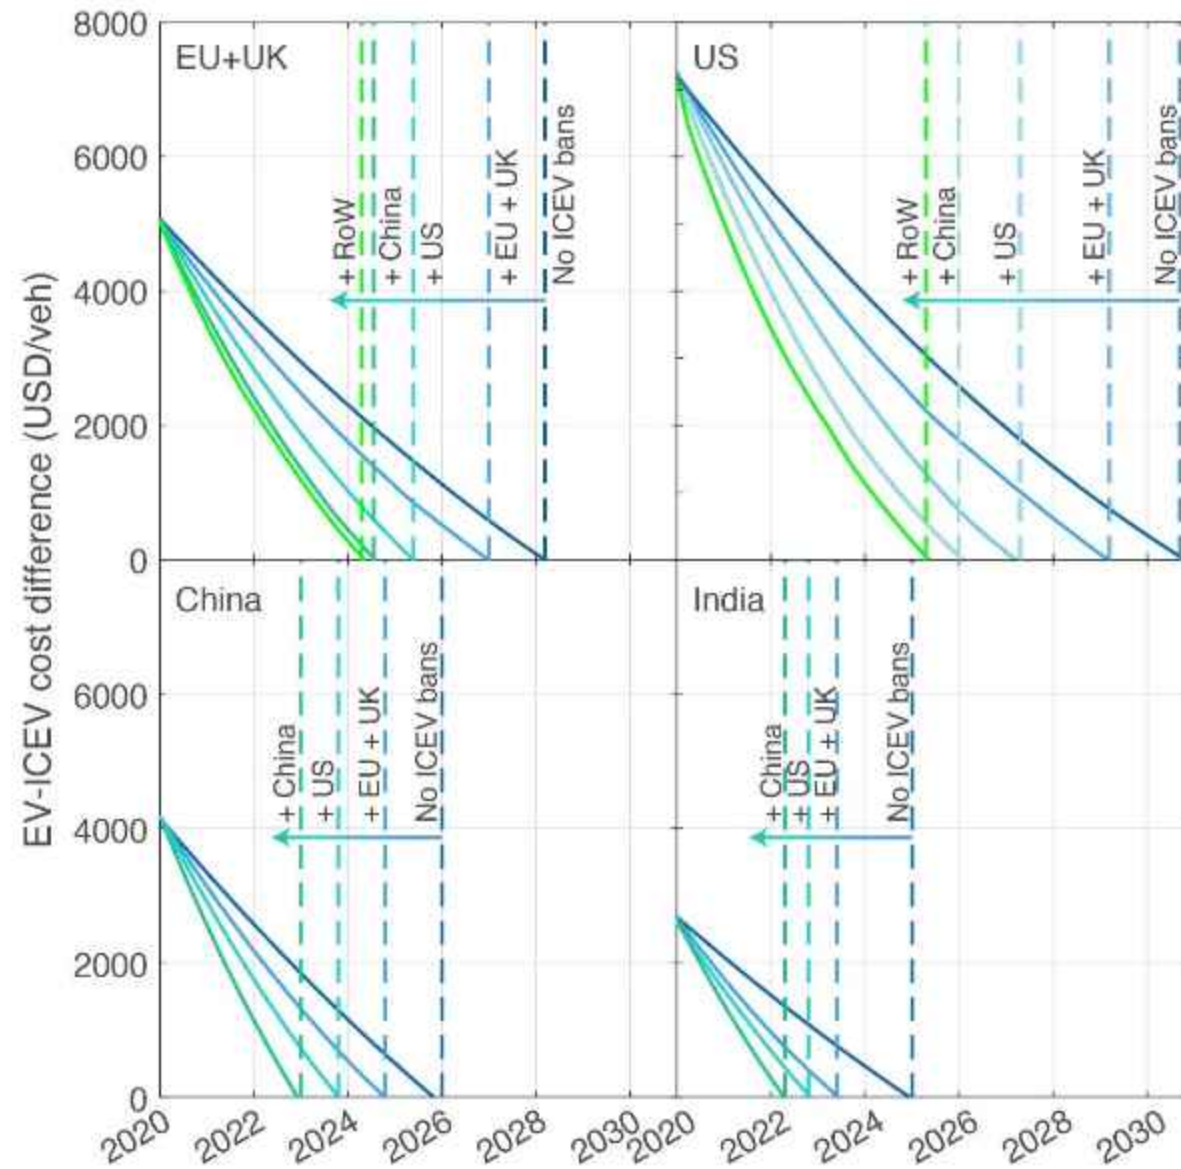


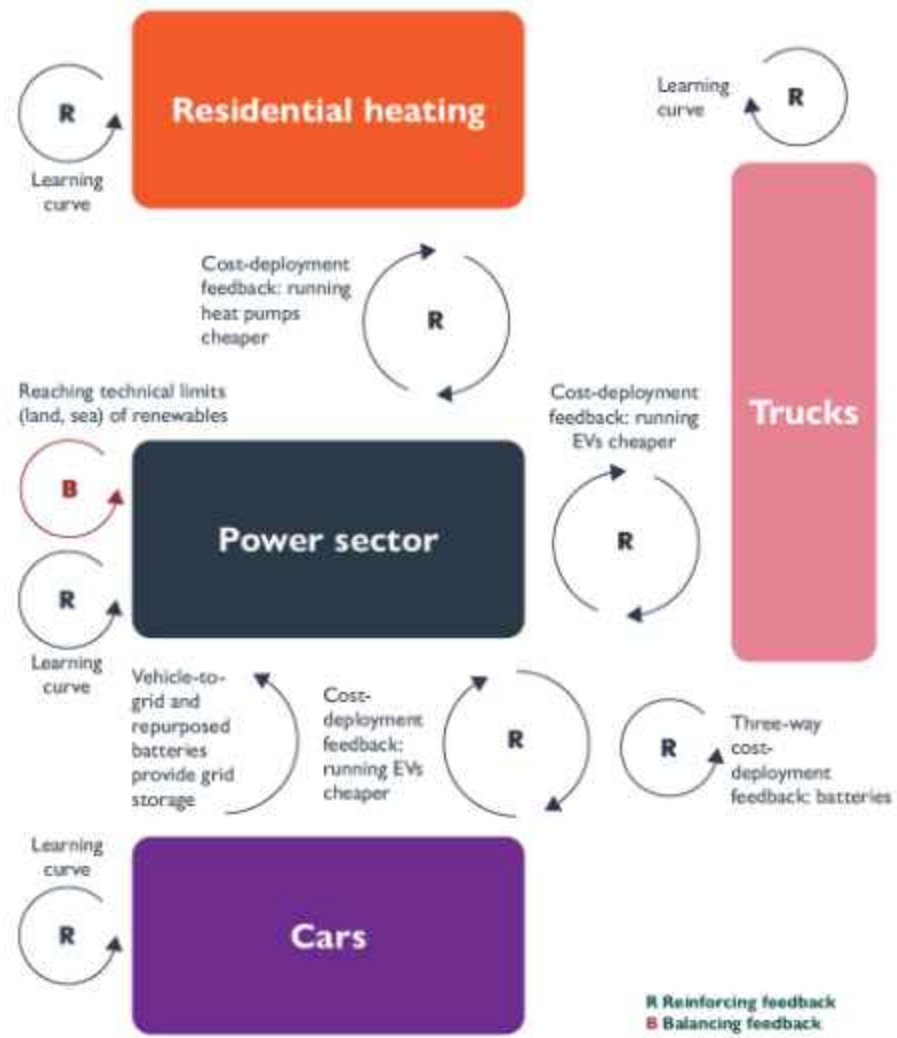
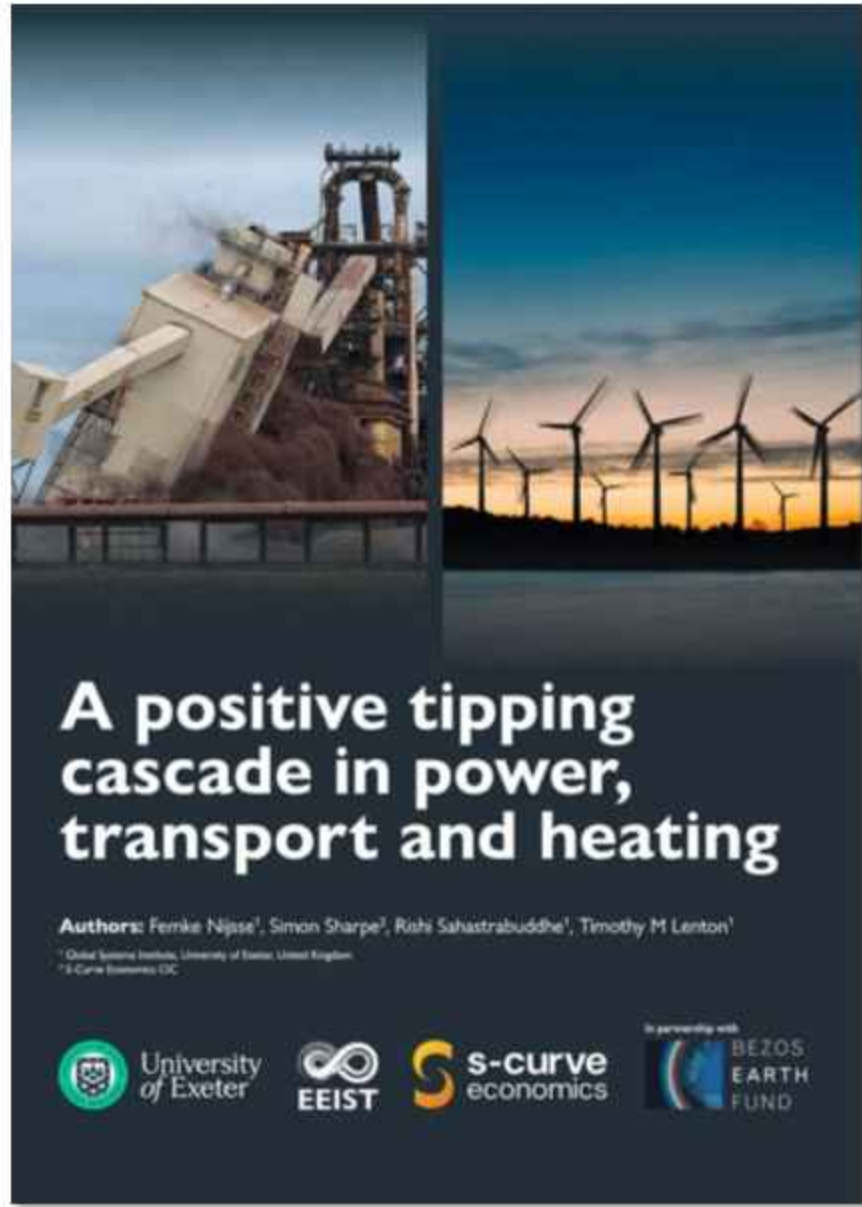
Battery price (\$/kWh)

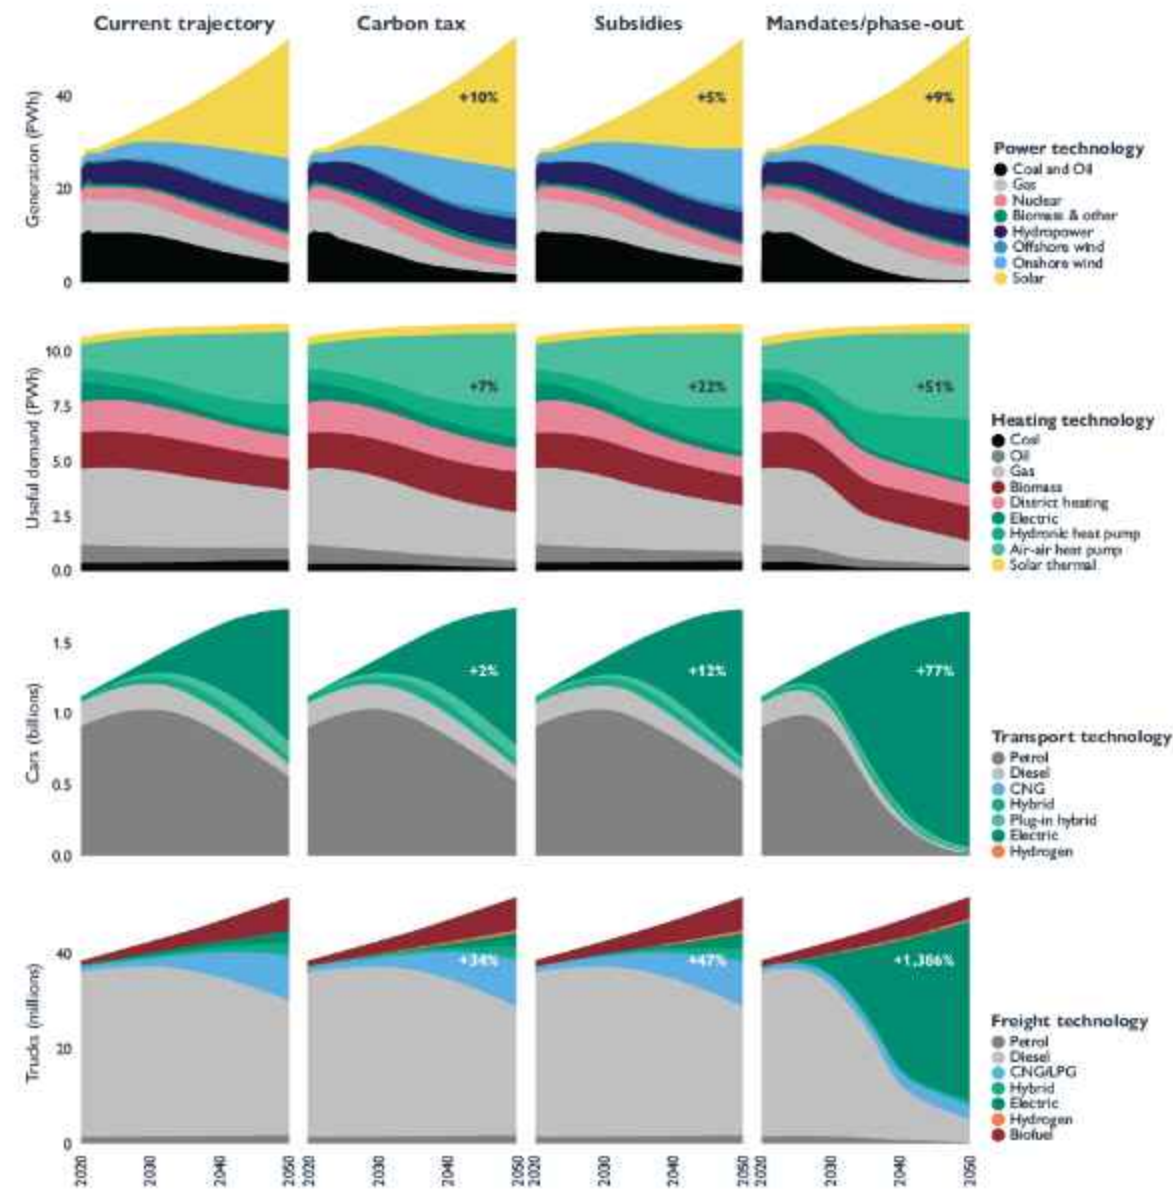


Global EV fleet (millions)

Banning ICEVs brings forward EV cost parity







Summary

- Robust early warning signals of loss of resilience are already seen in four coupled climate tipping elements
- Early opportunity signals of loss of resilience of incumbent fossil-fuelled technologies have been found in the car market
- Deep learning can improve early warning skill and may have the potential to forecast how far away a tipping points is

Critical slowing down

$$\frac{dx}{dt}(t) = -U'(x(t)) + \eta(t)$$

- state variable x , nonlinear Potential function U , additive white noise η with standard deviation σ .
- Assume a stable fixed point (i.e., an equilibrium state) x^* around which the dynamics takes place.
- Close to this fixed point, the potential U can be approximated by a quadratic function with minimum at x^* , i.e. for x near x^* , we have for some parameter $\lambda < 0$ that $U(x) \approx -\frac{\lambda}{2} x^2$.
- Approximating the potential by a quadratic function corresponds to a linearization of the equation of motion around x^* , i.e., for the fluctuations $\Delta x = x - x^*$ we obtain

$$\frac{d}{dt} \Delta x(t) = \lambda \Delta x(t) + \eta(t)$$

- This defines an Ornstein-Uhlenbeck (OU) process with restoring rate λ .
- Discretizing this process into time steps Δt yields an autoregressive process for which the variance and the autocorrelation function $\alpha(n)$ at n time steps are analytically known:

$$\text{Var}(\Delta x) = \frac{\sigma^2}{1 - e^{2\lambda\Delta t}} \approx -\frac{\sigma^2}{2\lambda}, \quad \alpha(n) = e^{n\lambda\Delta t}$$

Spatio-Temporal Categorical Output Model Emulation

Bertrand Nortier

September 10, 2025

Spatial Categorical Emulation

Categorical output model

- ▶ Let's assume that we want to train an emulator for an original model whose outputs are categorical, for example, a vector of words or a vector of colours

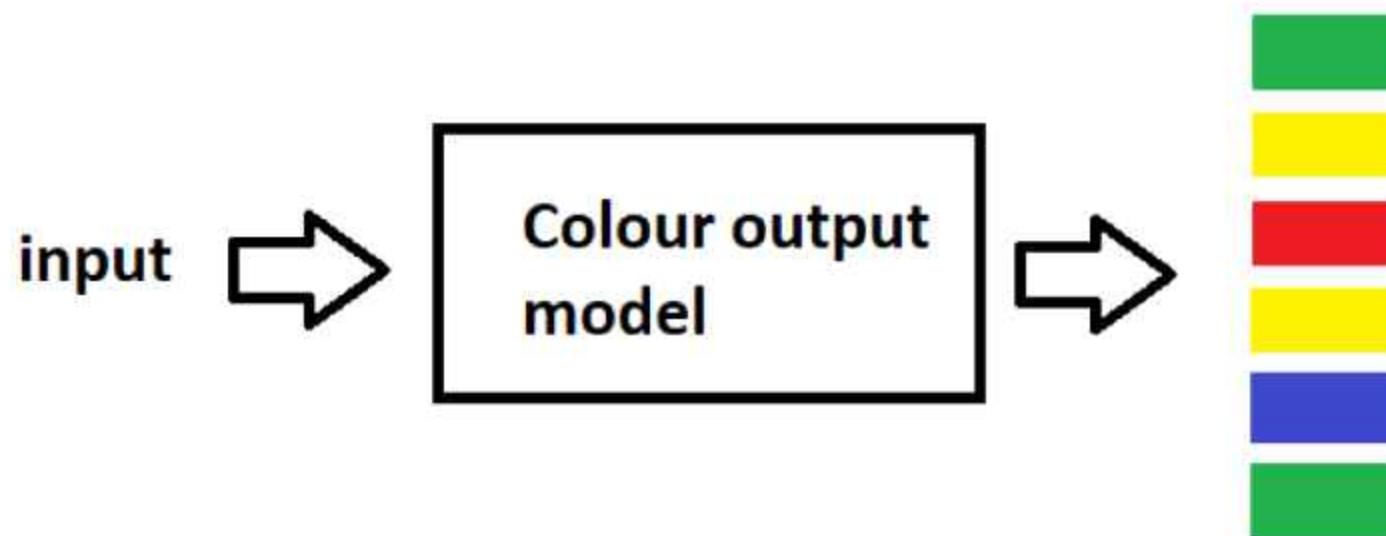


Figure: Colour output model

Latent continuous space

- ▶ Instead of trying to “code” the categories with arbitrary numbers, we can represent it the following way.
- ▶ First, we transform the categories into binary vectors:

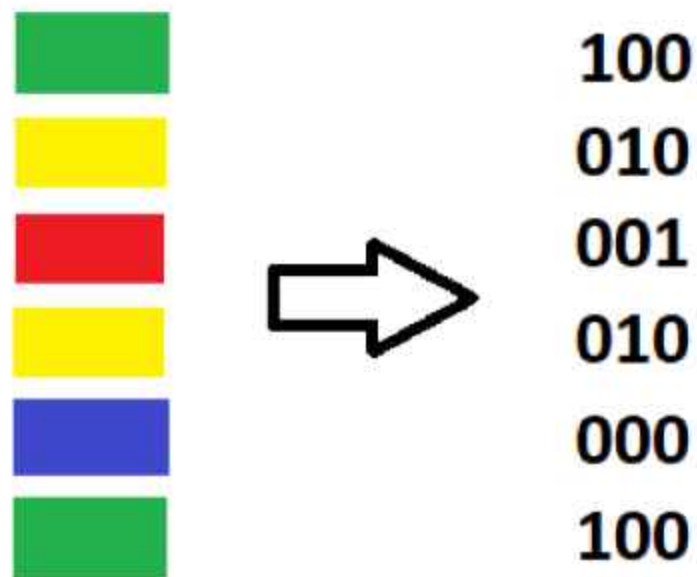


Figure: Colour output model

Latent Reduced Rank representation

- ▶ The model is then parametrized by continuous latent parameters $\gamma_{ijk} \in \mathbb{R}$ such that:

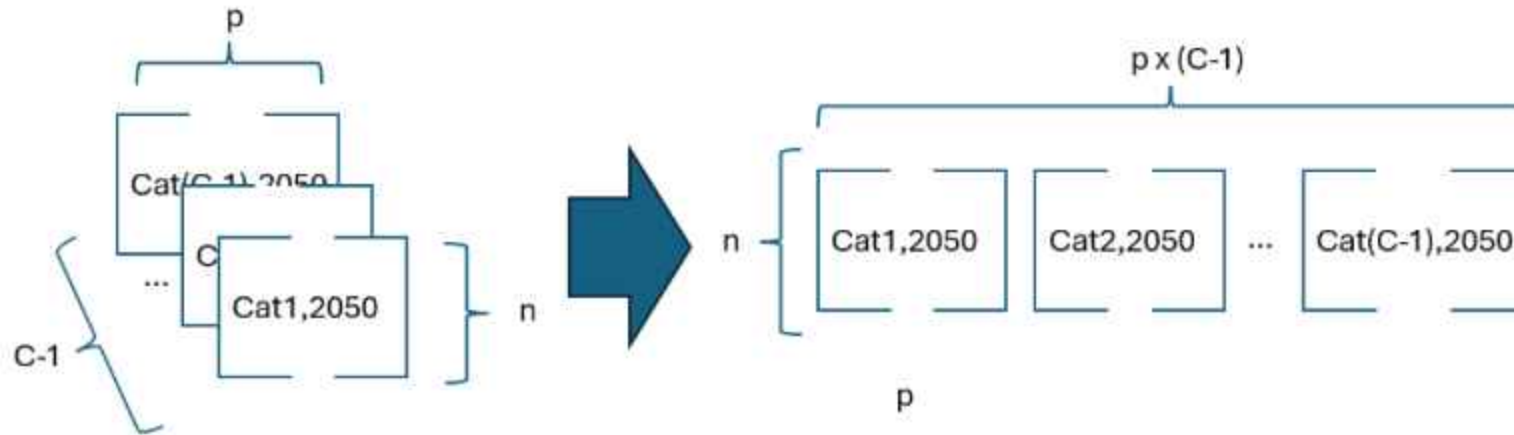
$$P(y_{ijk} = 1) = \frac{\exp(\gamma_{ijk})}{1 + \sum_{c=1}^{C-1} \exp(\gamma_{ijc})}$$

- ▶ vector of size p of C categories with n design points $\Rightarrow n \times p \times (C - 1)$ parameters.
- ▶ We may want to find a more parcimonious representation.
- ▶ Let's call $\mathbf{\Gamma}$ the tensor of size $n \times p \times (C - 1)$. We can reformat it into an $np \times (C - 1)$ matrix.
- ▶ We then want to find a representation of the form $([3], [6])$:

$$\mathbf{\Gamma}_{n \times p(C-1)} = \mathbf{1}_n \boldsymbol{\mu}_{p(C-1)}^T + \mathbf{W}_{n \times J} (\mathbf{K}_{p(C-1) \times J})^T$$

Tensor reformat

- ▶ An important step is to reformat the tensor Γ into a matrix
- ▶ When doing this **1) we are back to a case similar to the binary case 2) We can now apply standard SVD to obtain a reduced rank representation 3) we do not remove the correlations as data for design point $i, i = 1, \dots, n$ is still grouped together. 4) the information of which slide of the tensor the category is not important are categories are not ordered**



Algorithm

- ▶ We iterate 2 steps:

- 1) **likelihood optimization**

$$\mathbf{U}^{(m+1)} = \mathbf{\Gamma}^{(m)} + 2 \left[\mathbf{Y}_{1:(C-1)} - \mathbf{\Pi}_{1:(C-1)}^{(m)} \right]$$

- 2) **SVD/reduced rank projection:**

$$\left(\mathbf{U}^{(m+1)} - \mathbf{1}_n \left(\boldsymbol{\mu}^{(m+1)} \right)^T \right) = \mathbf{A}^{(m+1)} \mathbf{D}^{(m+1)} \left(\mathbf{B}^{(m+1)} \right)^T$$

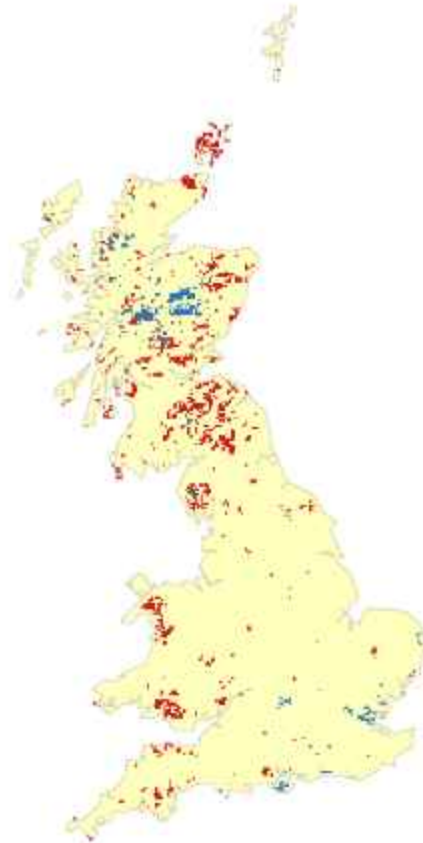
We can then keep the J first components:

$$\mathbf{\Gamma}_{n \times p(C-1)}^{(m+1)} = \mathbf{1}_n \left(\boldsymbol{\mu}_{p(C-1)}^{(m+1)} \right)^T + \left(\mathbf{A}_{1:n,1:J}^{(m+1)} \mathbf{D}_{1:J,1:J}^{(m+1)} \right) \left(\mathbf{B}_{p(C-1) \times J}^{(m+1)} \right)^T$$

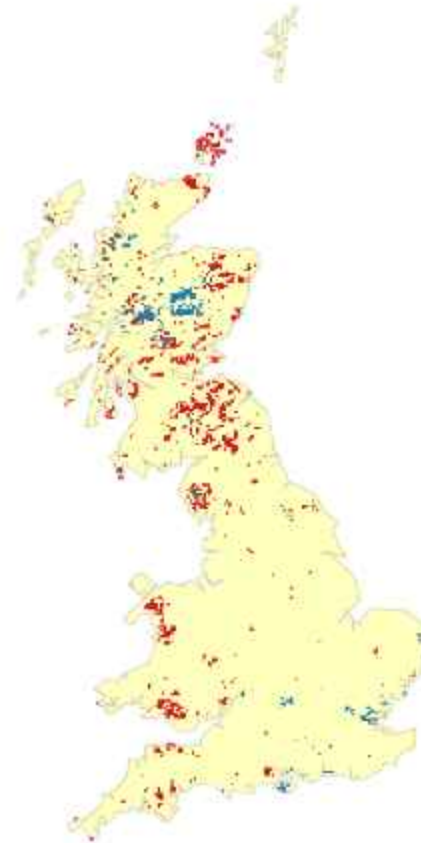
We then choose: $\mathbf{W}^{(m+1)} = \left(\mathbf{A}_{1:n,1:J}^{(m+1)} \mathbf{D}_{1:J,1:J}^{(m+1)} \right)$ as the matrix of basis multipliers and $\mathbf{K}^{(m+1)} = \mathbf{B}_{pC \times J}^{(m+1)}$ as the basis. The columns of \mathbf{W} can then be emulated with GPs/DGPs.

Application to Greenhouse gas removal model

true map



low rank reconstruction



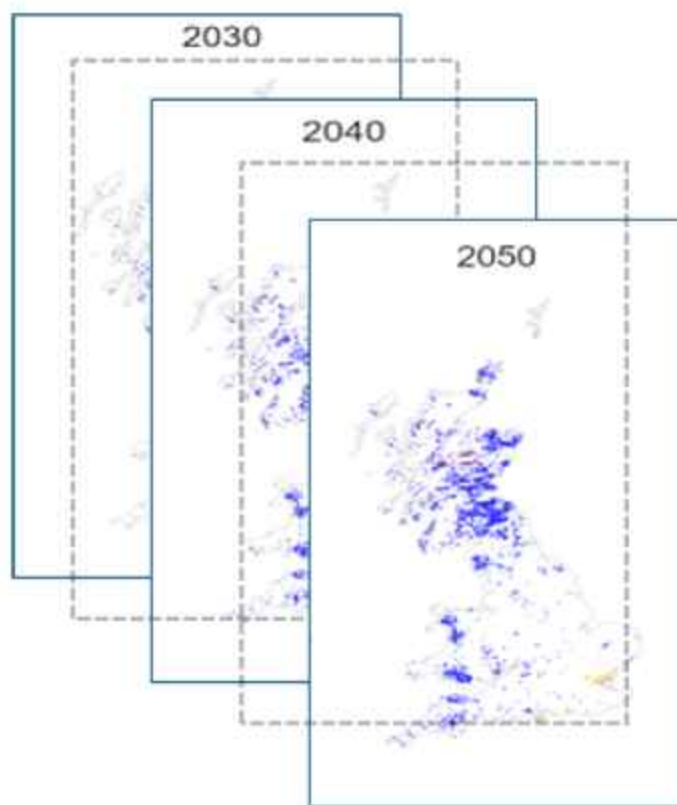
Type of trees

- no planting
- felled conifers
- felled deciduous
- unfelled conifers
- unfelled deciduous

Time Series of Categorical Maps

Time Series of Maps I

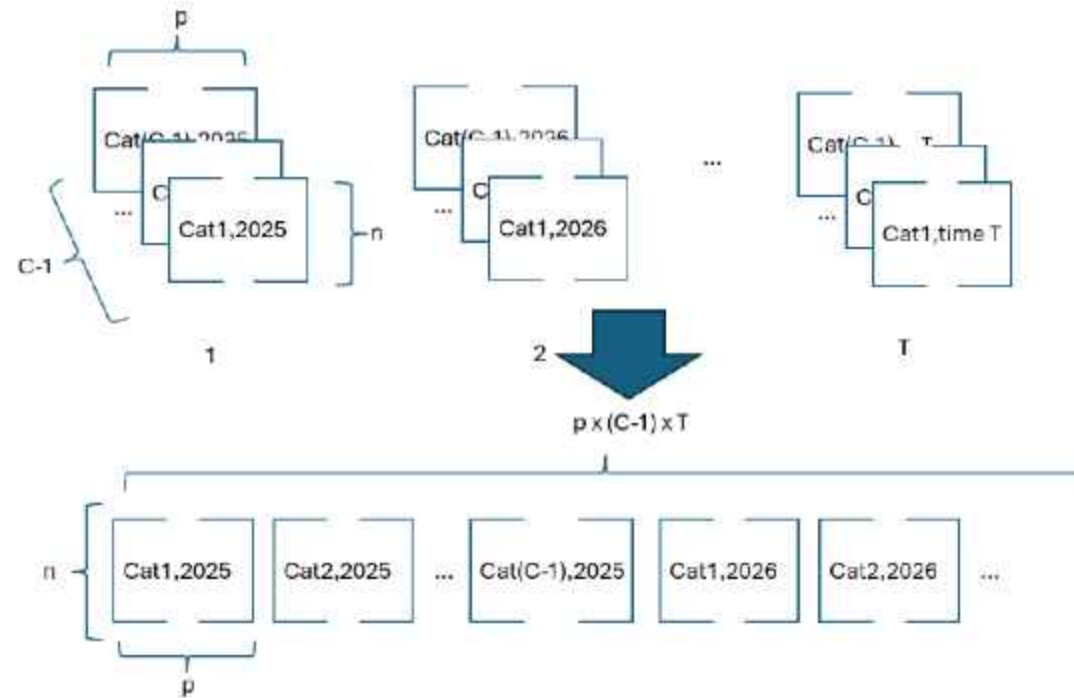
- The model does not solely output Categorical maps but entire time series. What method could we use to emulate this?



Categorical map time series

First Possibility: Stack the map for each year as we have stacked the Categories

- Issue: We end up with very large dimension. Example:
 $n = 300$, $p = 50,000$, $C = 5$, $T = 25$, then the matrix is $300 \times (50000 * 4 * 25) = 300 \times 5,000,000$.

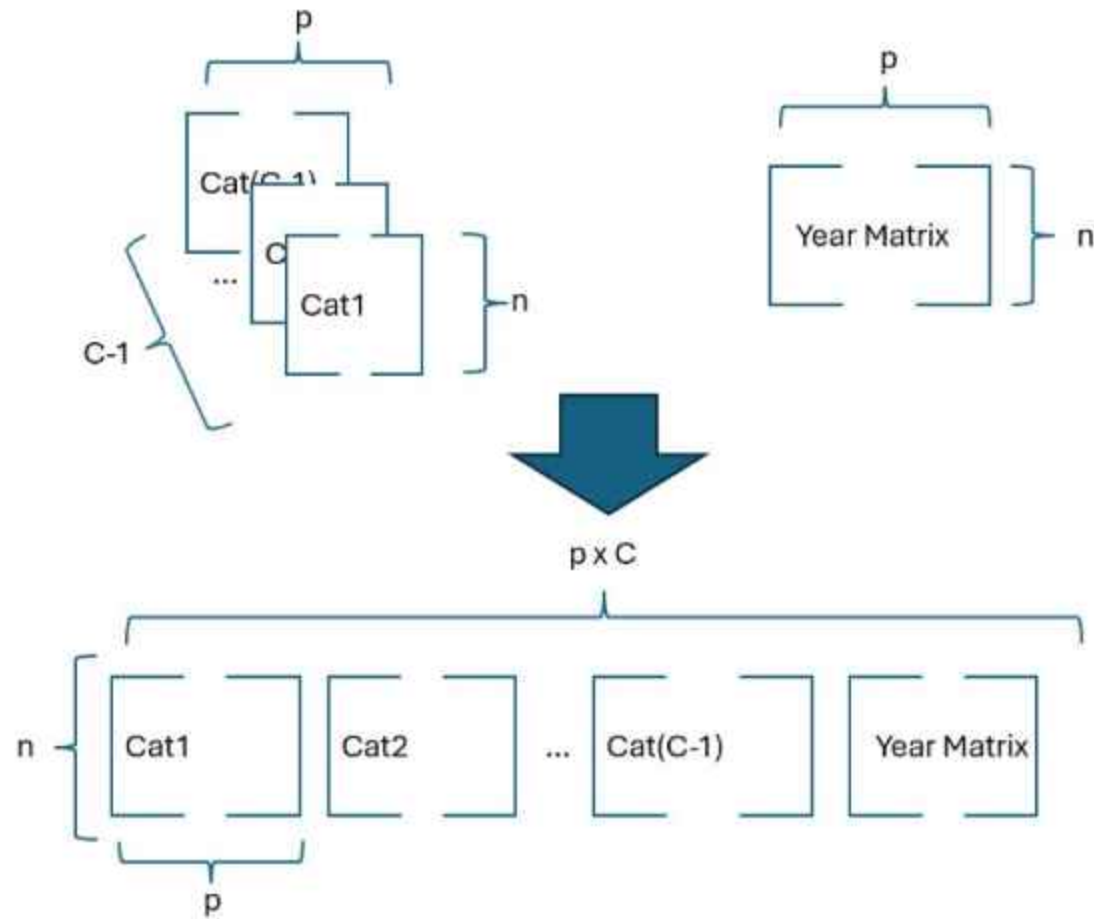


Possibility 2: reformat the time and stack Category with year I

- ▶ We note that in the models we are emulating, the decision to plant occurs at a single time point.
- ▶ We propose another data format that is more parsimonious

id	Cat	year		id	Cat Bin	year
1	1	2024		1	1000	2024
3	1	2041		3	1000	2041
141	1	2023		141	1000	2023
142	1	2023		142	1000	2023
143	4	2023		143	0001	2023
153	1	2023	↔	153	1000	2023
212	2	2039		212	0100	2039
229	2	2026		229	0100	2026
237	2	2033		237	0100	2023
249	4	2029		249	0001	2029
...				...		

Possibility 2: reformat the time and stack Category with year II



Possibility 2: reformat the time and stack Category with year III

- We now have 1 more step

1) **Categorical part of the** likelihood optimization

$$\mathbf{U}^{(m+1)} = \mathbf{\Gamma}^{(m)} + 2 \left[\mathbf{Y}_{1:(C-1)} - \mathbf{\Pi}_{1:(C-1)}^{(m)} \right].$$

- 2) Year Matrix Stacking: $\tilde{\mathbf{U}}_{n \times pC}^{(m+1)} = [\tilde{\mathbf{U}}_{n \times p(C-1)}^{(m+1)} | \mathbf{Y}_{n \times p}]$

- 3) **SVD/reduced rank projection:**

$$\left(\tilde{\mathbf{U}}^{(m+1)} - \mathbf{1}_n \left(\boldsymbol{\mu}^{(m+1)} \right)^T \right) = \mathbf{A}^{(m+1)} \mathbf{D}^{(m+1)} \left(\mathbf{B}^{(m+1)} \right)^T$$

We can then keep the J first components:

$$\tilde{\mathbf{I}}_{n \times pC}^{(m+1)} = \mathbf{1}_n \left(\boldsymbol{\mu}_{pC}^{(m+1)} \right)^T + \left(\mathbf{A}_{1:n, 1:J}^{(m+1)} \mathbf{D}_{1:J, 1:J}^{(m+1)} \right) \left(\mathbf{B}_{pC \times J}^{(m+1)} \right)^T$$

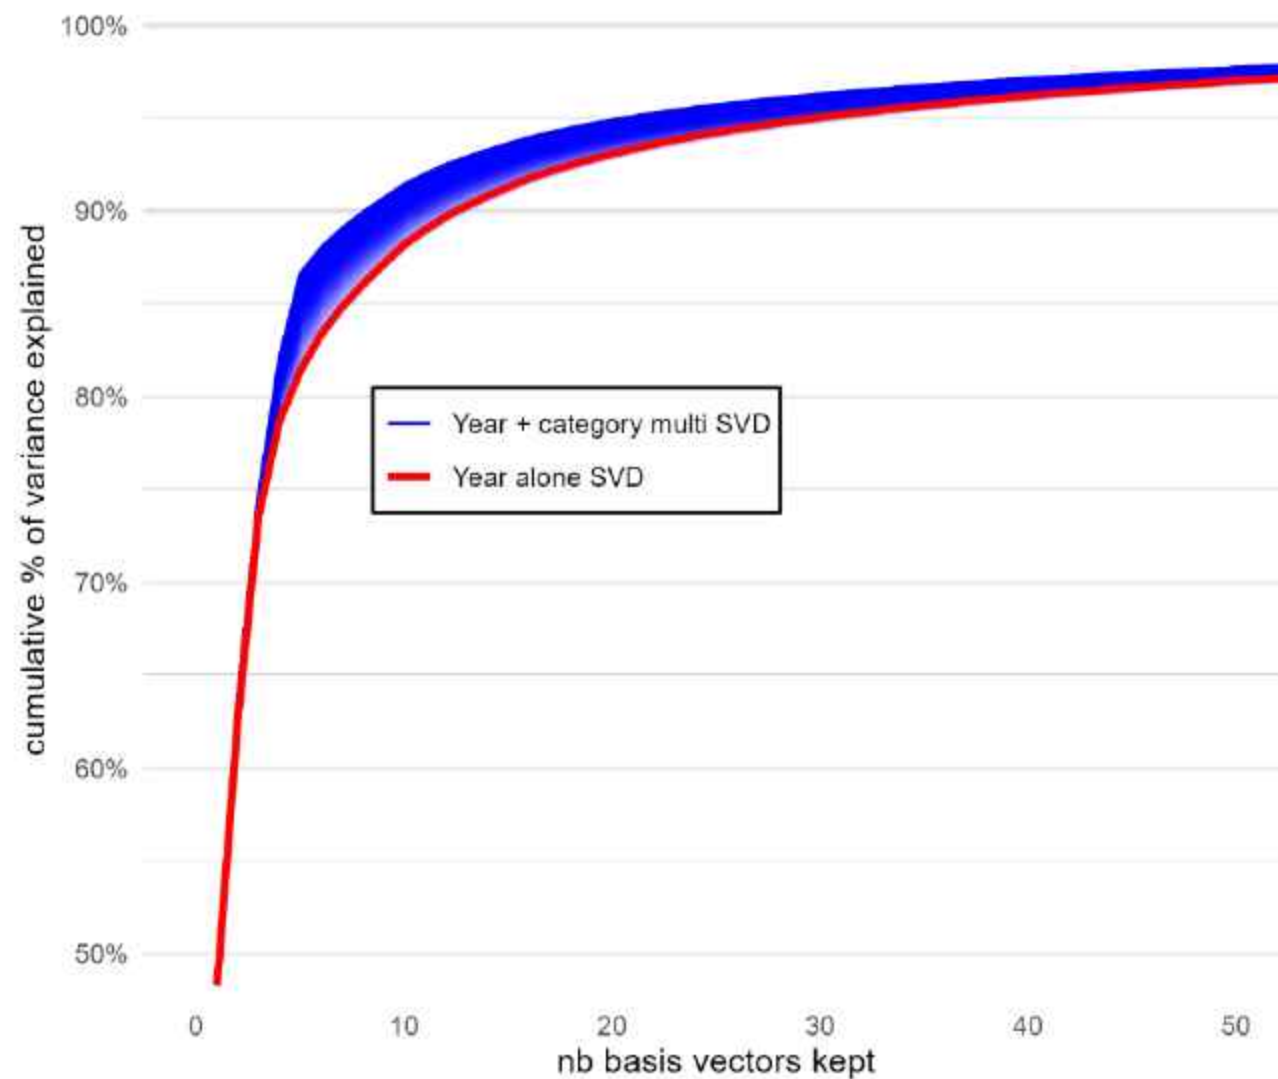
Likelihood and multitask learning

- ▶ this stacking of the categorical and year matrix could be viewed as **multi-task learning**.
- ▶ The likelihood or nll are then a composite of the multinomial likelihood and the likelihood for the year. Assuming we use the negative log-likelihood:

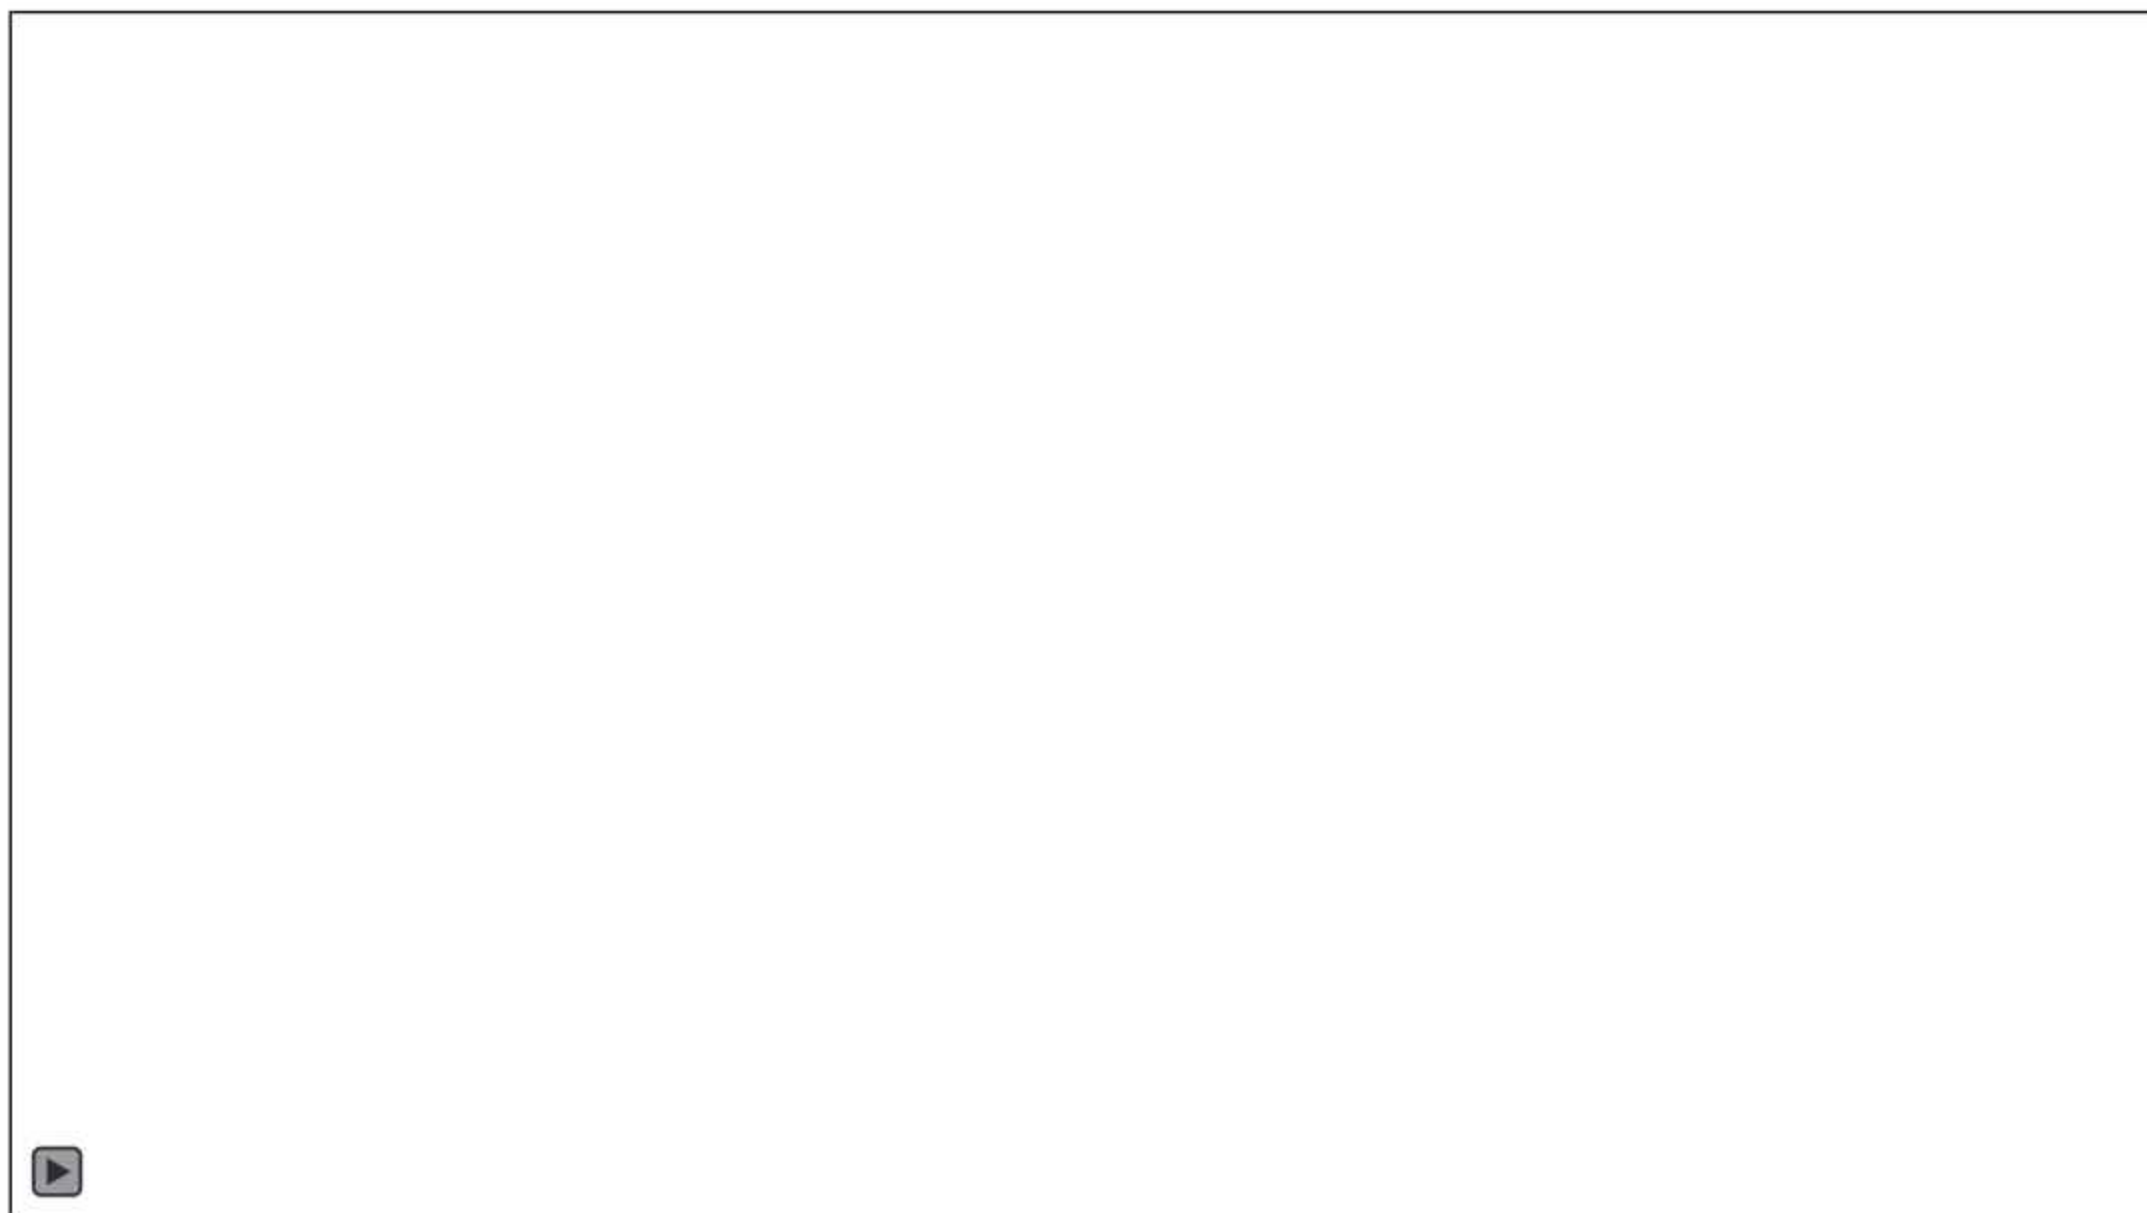
$$\ell = \lambda \ell_{cat} + (1 - \lambda) \ell_{lin}$$

- ▶ $\ell_{cat} = -\frac{1}{np} \sum_{i=1}^n \sum_{j=1}^p \sum_{k=1}^C z_{ijk} \log(\hat{\pi}_{ijk}),$
 $\ell_{lin} = \frac{1}{np} \sum_{i=1}^n \sum_{j=1}^p (y_{ijk} - \hat{y}_{ijk})^2$

How does stacking the Year matrix with Gamma Matrix affect the SVD of year?



Example of integration of the method in R shiny app



References I

- [1] S. Ahmadi-Asl, S. Abukhovich, M. G. Asante-Mensah, A. Cichocki, A. H. Phan, T. Tanaka, and I. Oseledets.
Randomized algorithms for computation of tucker decomposition and higher order svd (hosvd).
IEEE Access, 9:28684–28706, 2021.
- [2] B. W. Bader and T. G. Kolda.
Tensor decompositions and their application.
<https://www.osti.gov/biblio/1147546>, 2007.
- [3] W. Chang, M. Haran, P. Applegate, and D. Pollard.
Calibrating an ice sheet model using high-dimensional binary spatial data.
Journal of the American Statistical Association, 111(513):57–72, 2016.
- [4] G. H. Golub and C. F. Van Loan.
Matrix computations.
JHU press, 2013.
- [5] G. Gopalan and C. K. Wikle.
A higher-order singular value decomposition tensor emulator for spatiotemporal simulators.
Journal of Agricultural, Biological and Environmental Statistics, 27(1):22–45, 2022.

References II

- [6] P. J. Groenen and J. Josse.
Multinomial multiple correspondence analysis.
arXiv preprint arXiv:1603.03174, 2016.
- [7] M. E. Kilmer, L. Horesh, H. Avron, and E. Newman.
Tensor-tensor algebra for optimal representation and compression of multiway data.
Proceedings of the National Academy of Sciences, 118(28):e2015851118, 2021.
- [8] C.-T. Kuo.
Higher order svd: theory and algorithms.
en. In:(), page 8, 2013.

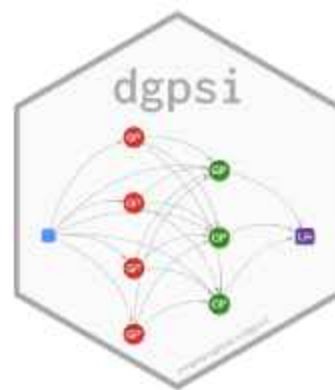


Where should all the trees go?

Danny Williamson, Deyu Ming, Timothee Bacri, Paul Wright, Amy Binner, Jonathan Mosedale, Ivis Kerama, Muhammad Hasan, Bertrand Nortier, Mingda Yuan and the ADD-TREES team



University
of Exeter



Nitrate price per kg:



Phosphate price per ha:



Carbon price per tonne:



Total budget (£bn):



☒ Show full parcel map *i*

☐ Annual Area Limit (10 kha) *i*

☐ Annual Budget Limit *i*

☐ Spatial Prioritisation *i*

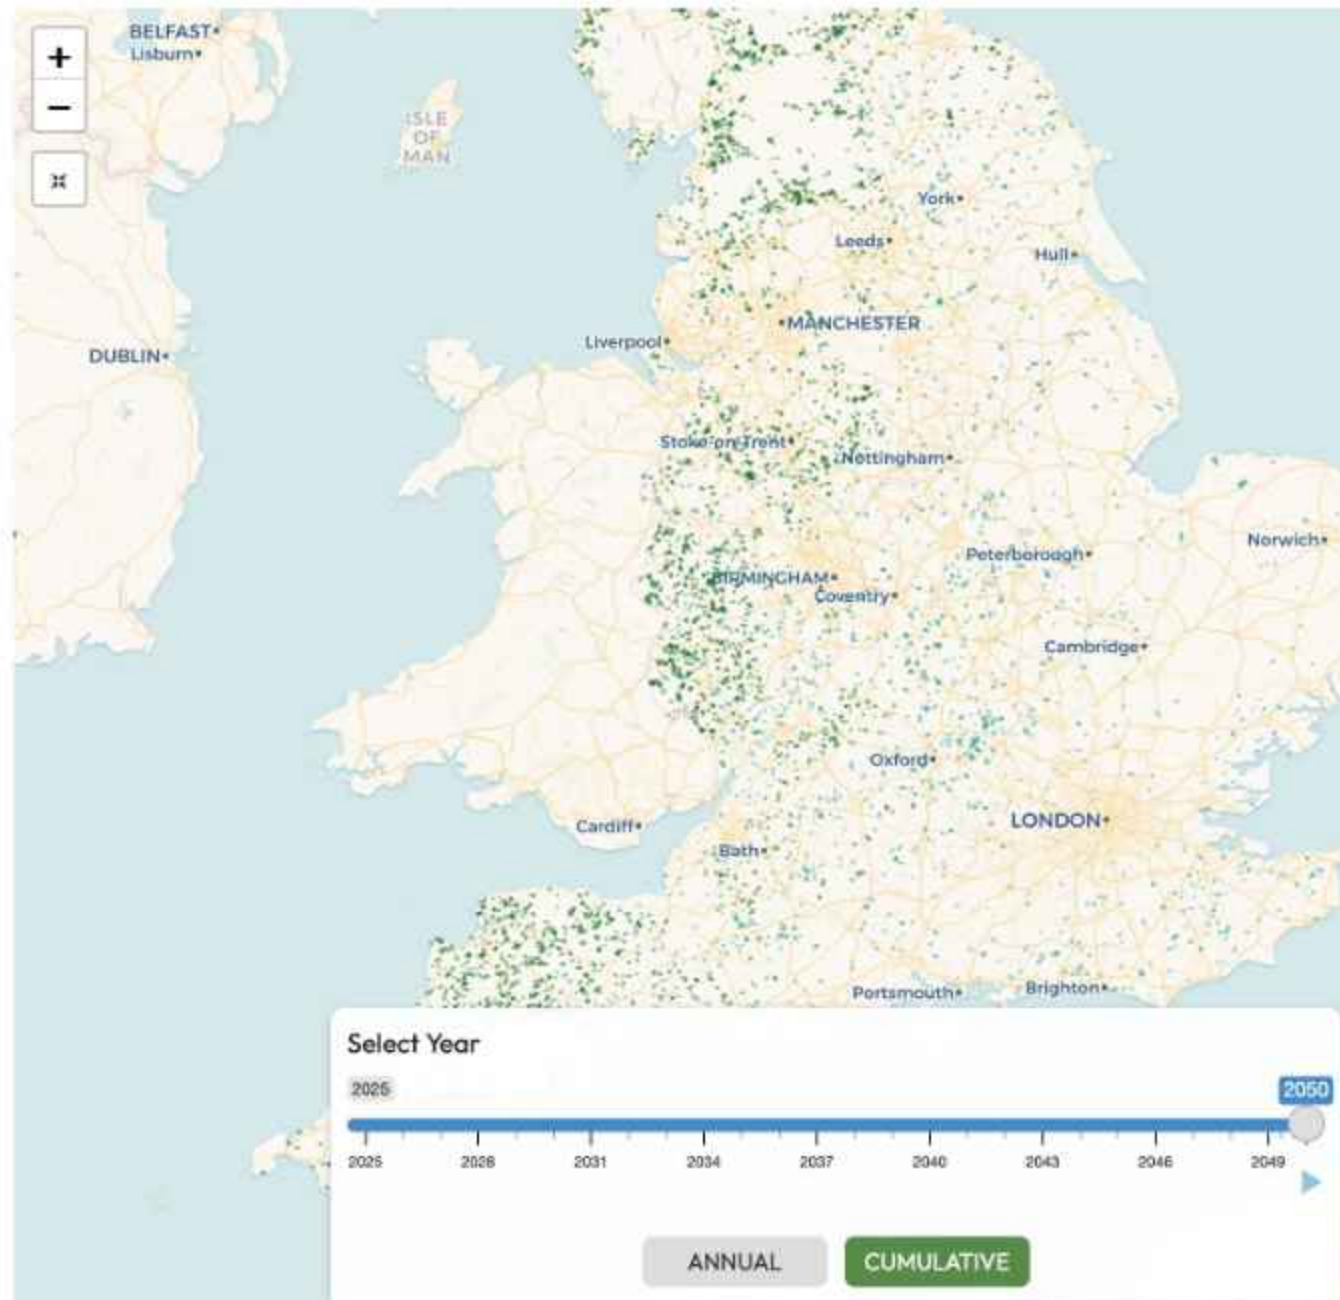
SUBMIT

RESET

SAVE

COMPARE

Saved Strategies *^*



Outcomes

Carbon (MtCO₂) 1.37

Trees (Mha) 2.23

Nitrate (Mkg) 0.73

Phosphate (Mkg) 17.65

Biodiversity (SR) 1052.64

Grassland (Mha) -167.58

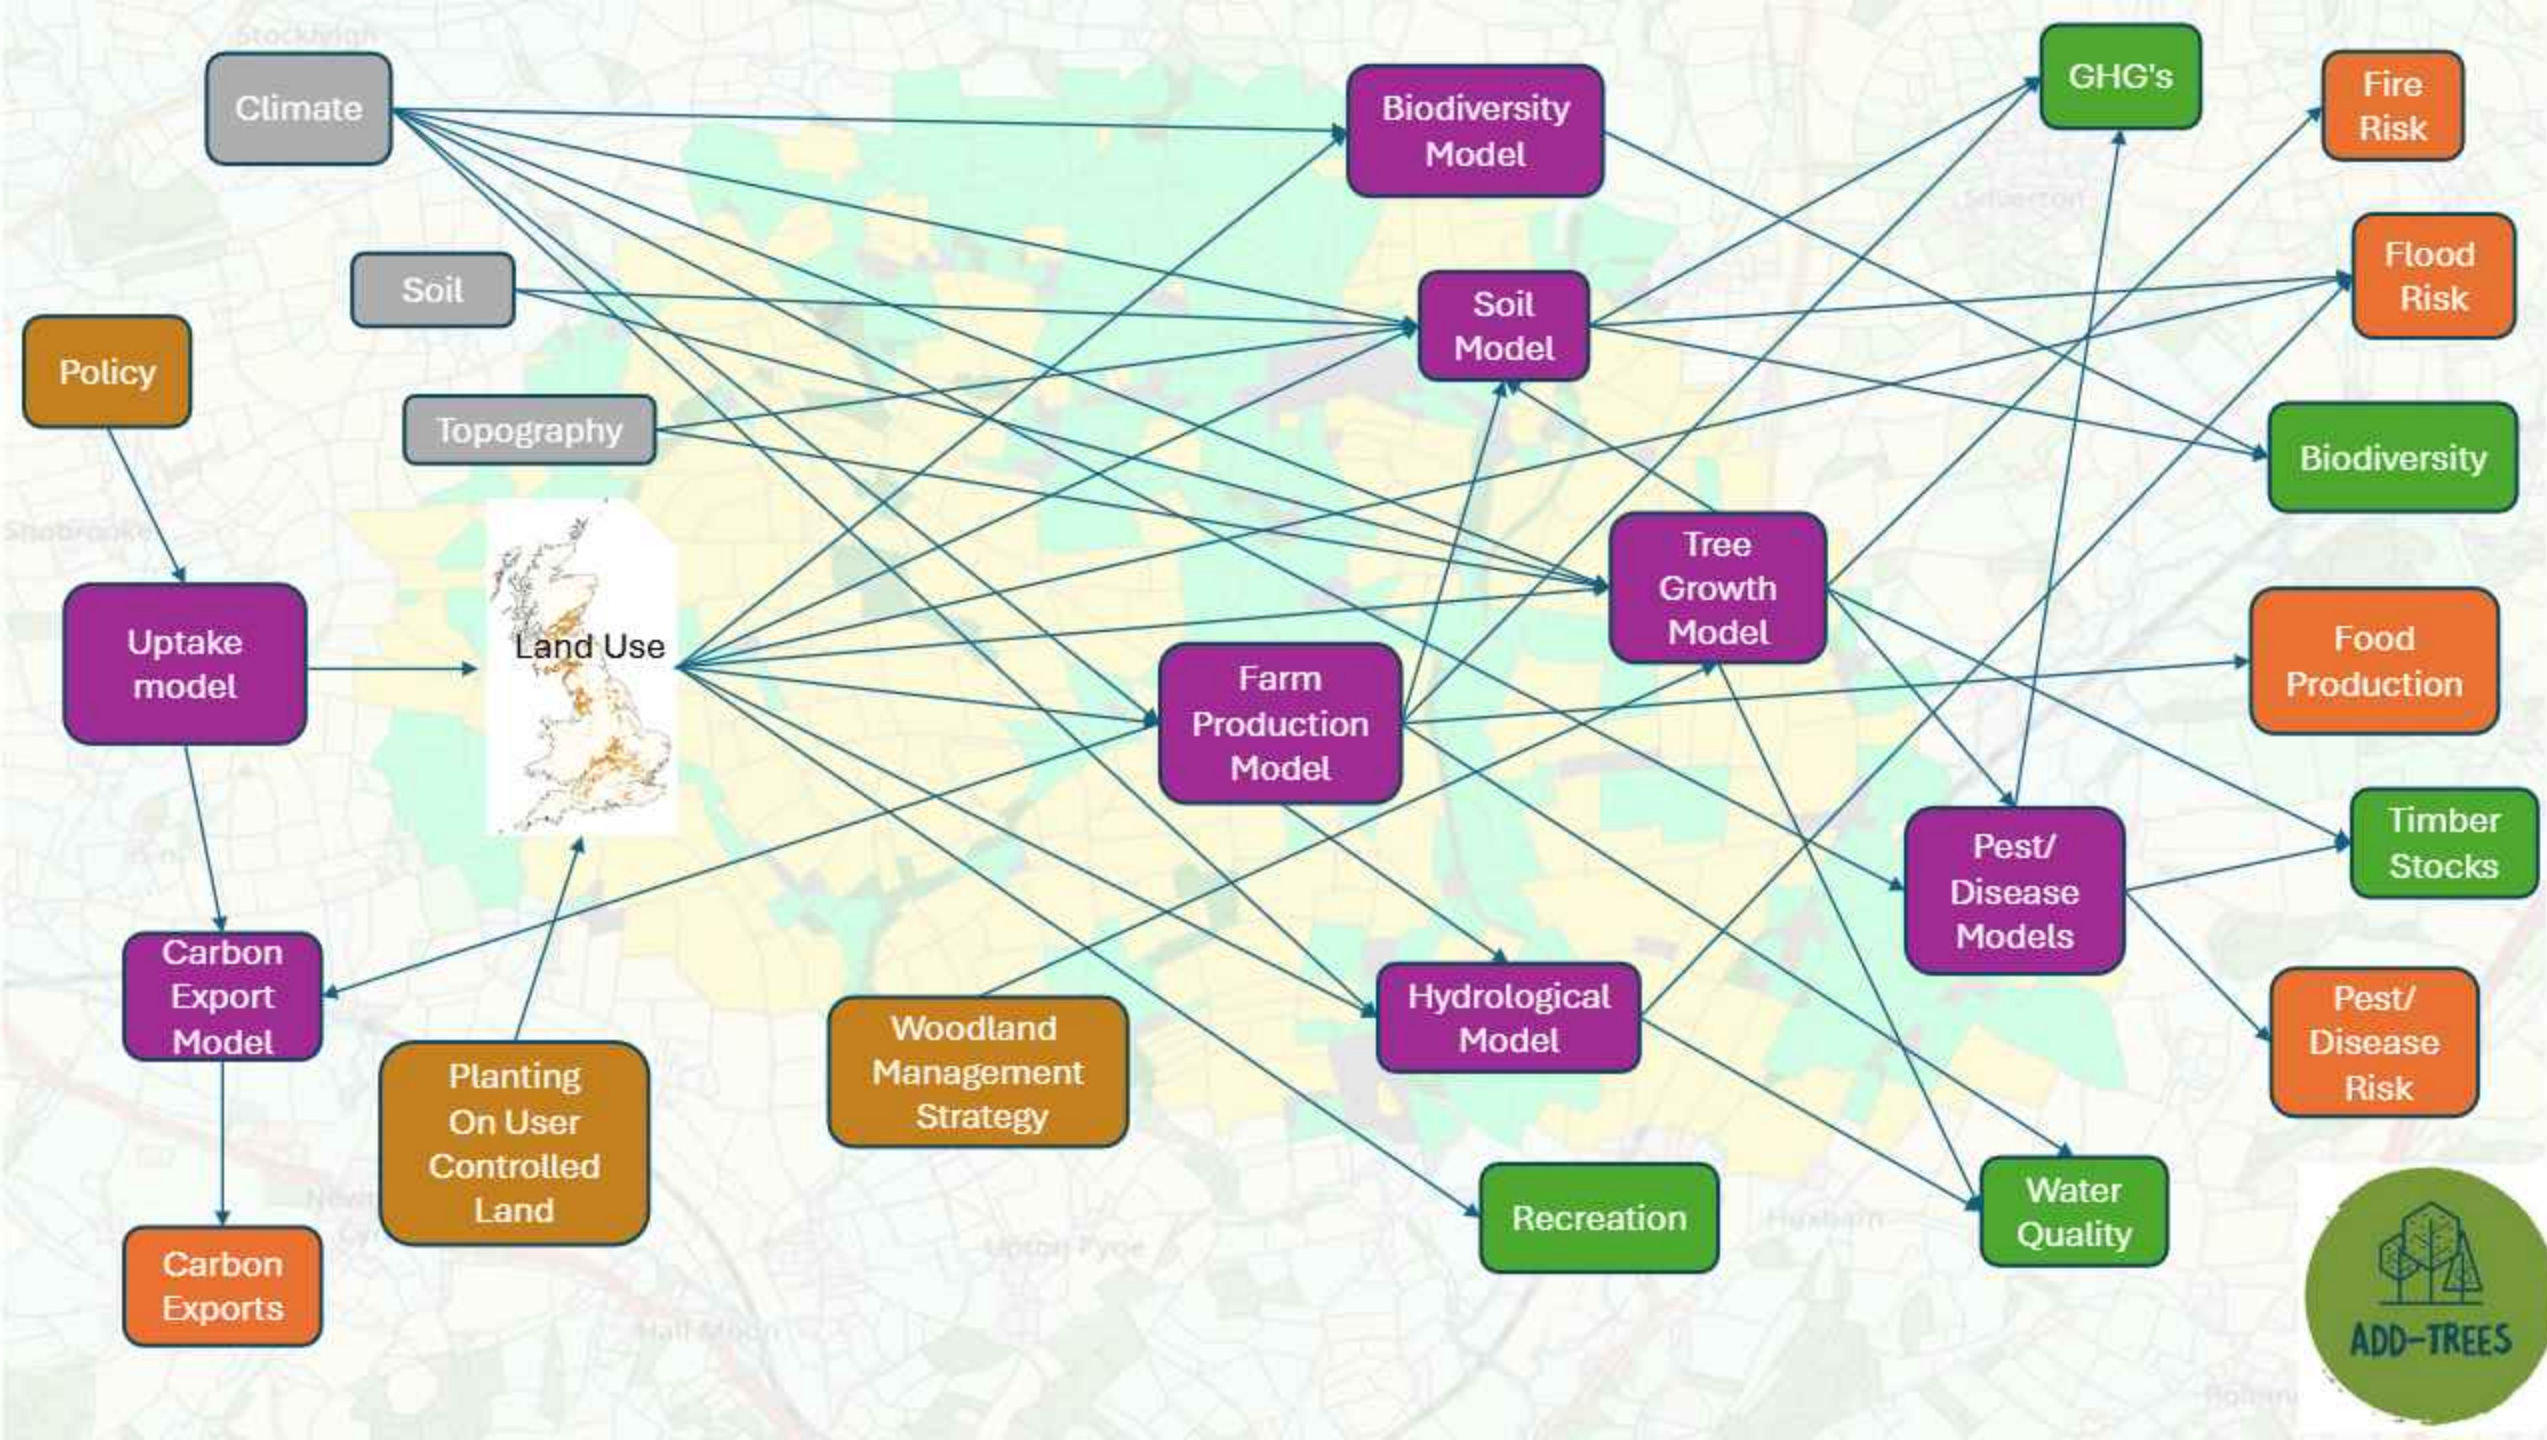
Total spend (£bn) 4.38

Show Time Series

Planting Type

Legend for Planting Type:

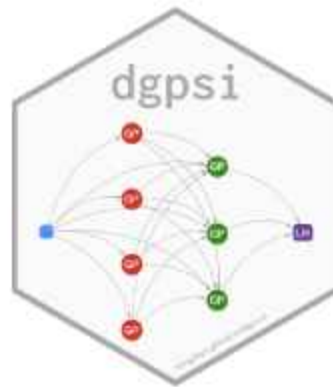
- Unfelled Conifers
- Unfelled Deciduous
- Felled Conifers
- Felled Deciduous
- Species-Rich Grassland







Linked GP emulators (a network model)

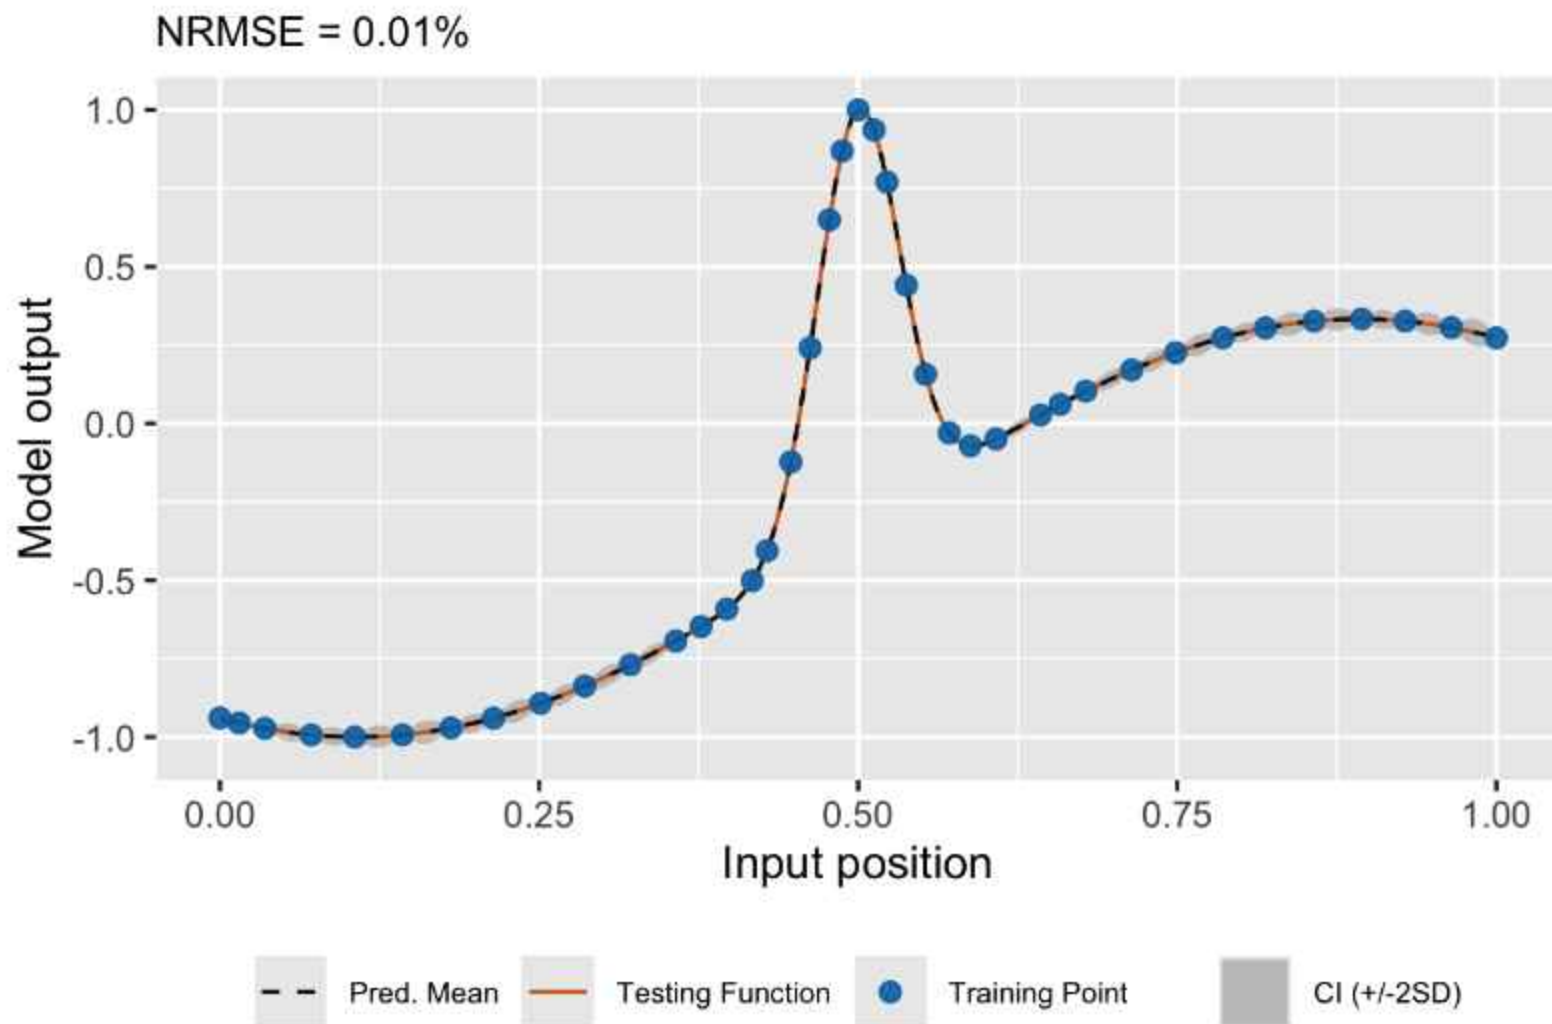


Linked GP (a network model)

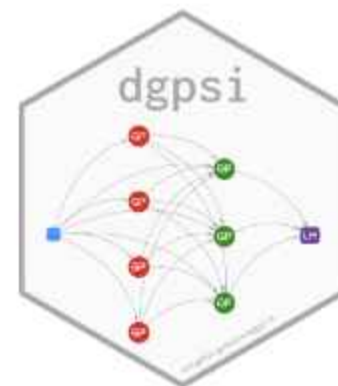


f2 emulator

ators



1 Process Emulator

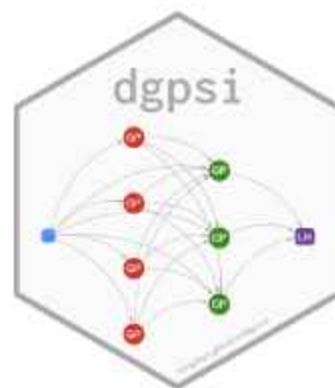


NRMSE = Normalized Root Mean Squared Error

Deep GP (a network model for a single node)



=

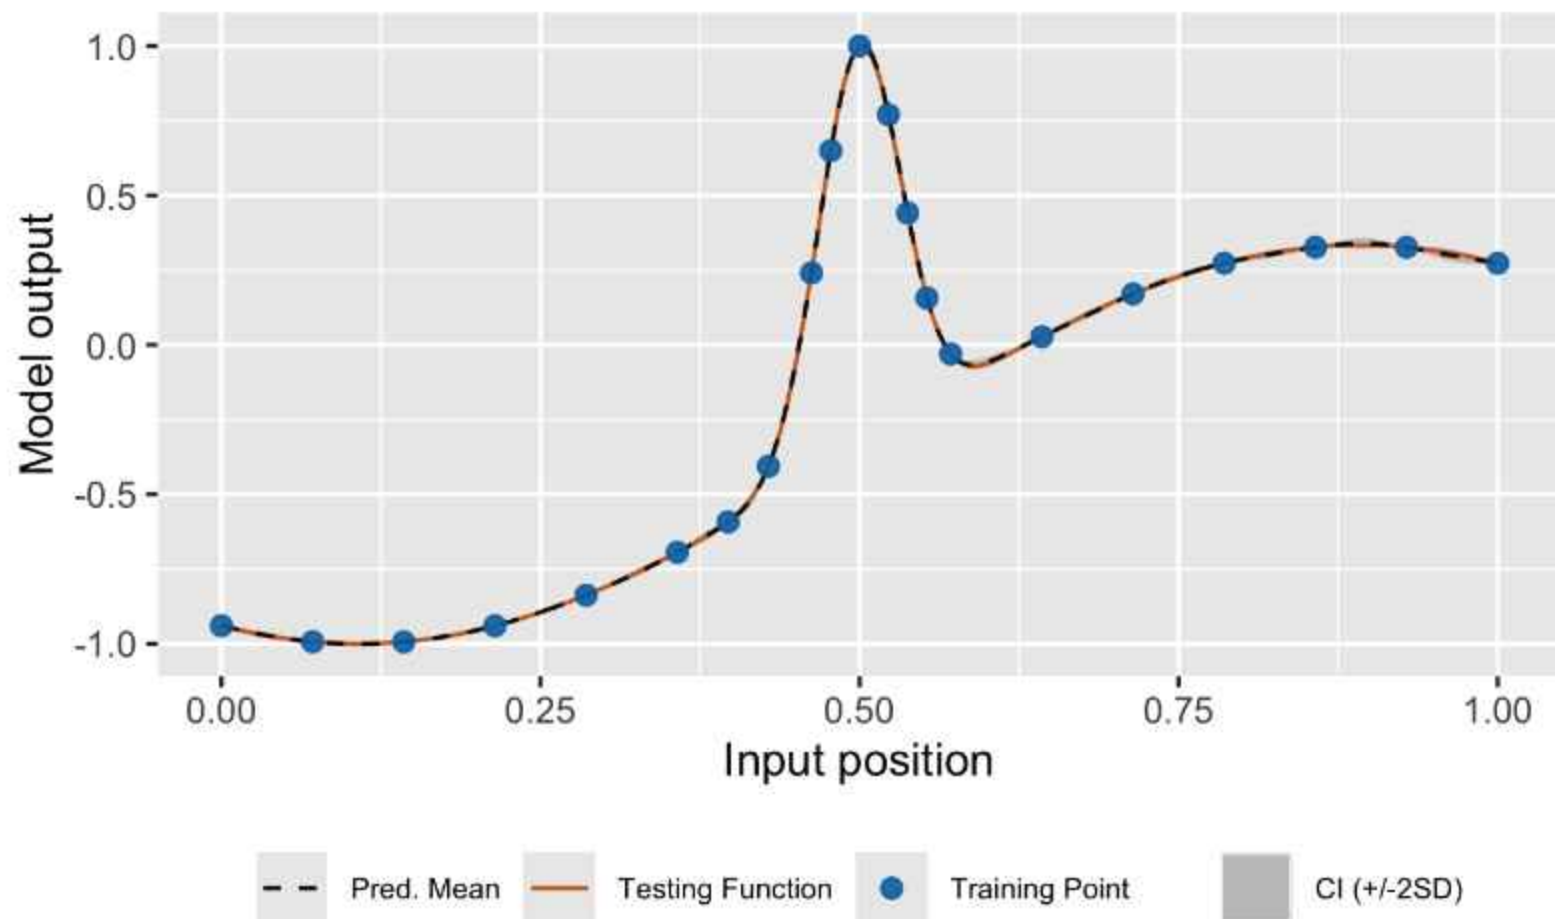


Deep GP (a network model for a single node)

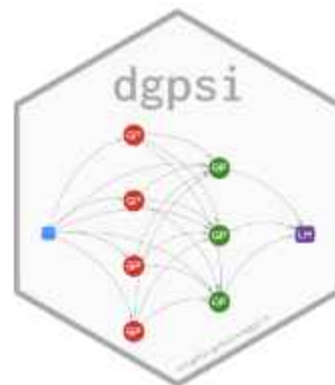


f2 emulator

O1: NRMSE = 0.12%



O_i = Output *i* of the DGP emulator
NRMSE = Normalized Root Mean Squared Error

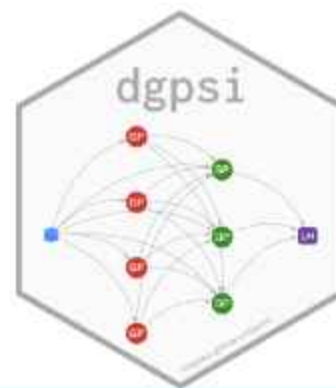




Linked DGP (a network model for model networks)



- We could replace each node in the LGP with a DGP (LDGP).

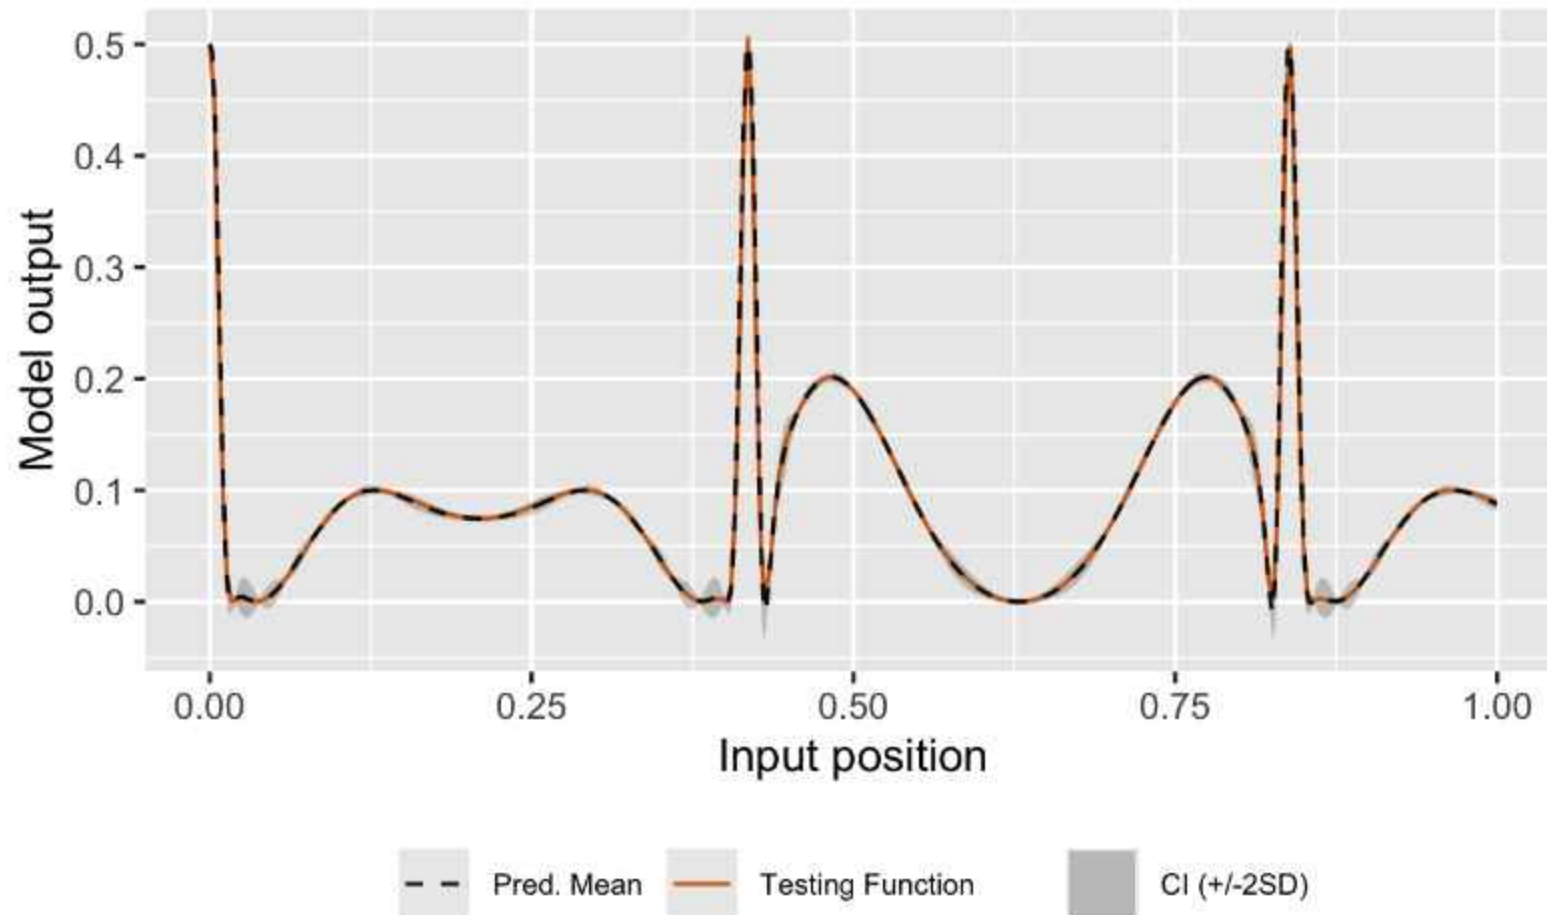


Linked DGP (a network model for model networks)



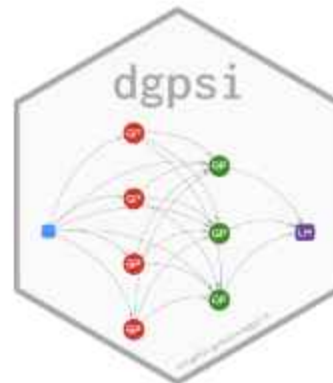
Linked Emulator

E1O1: NRMSE = 1.09%



E_{iOj} = Output j of Emulator i in the final layer of the linked emulator

NRMSE = Normalized Root Mean Squared Error



Following the Paws Behind the Patterns



Oscar Rodriguez de Rivera Ortega

Lecturer in Statistics, Faculty of Environment, Science and Economy

Introduction

- Movement is a fundamental component of animal ecology, shaping how individuals interact with their environment.



- Gaining insight into the underlying mechanisms that generate spatiotemporal patterns of movement is thus critical for the effective conservation and management of wildlife populations

Advances in biologging technologies

particularly GPS-based telemetry, now allow researchers to collect animal tracking data at fine spatial and temporal resolutions



However, they also introduce analytical **challenges**

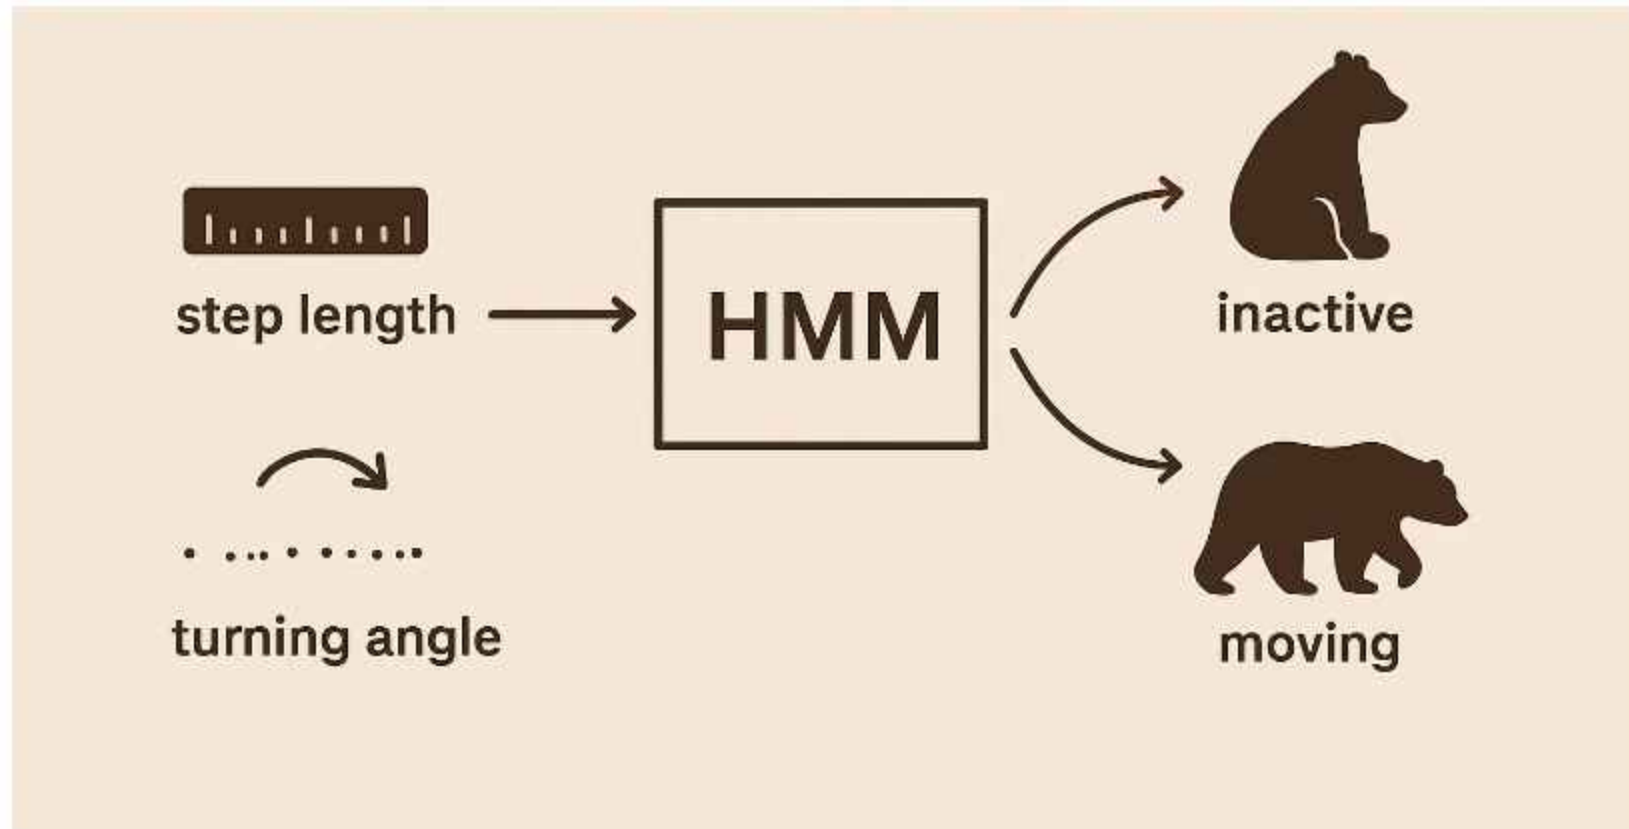


A crucial first step in many movement ecology studies is to segment animal tracks into basic functional units of movement that carry biological meaning

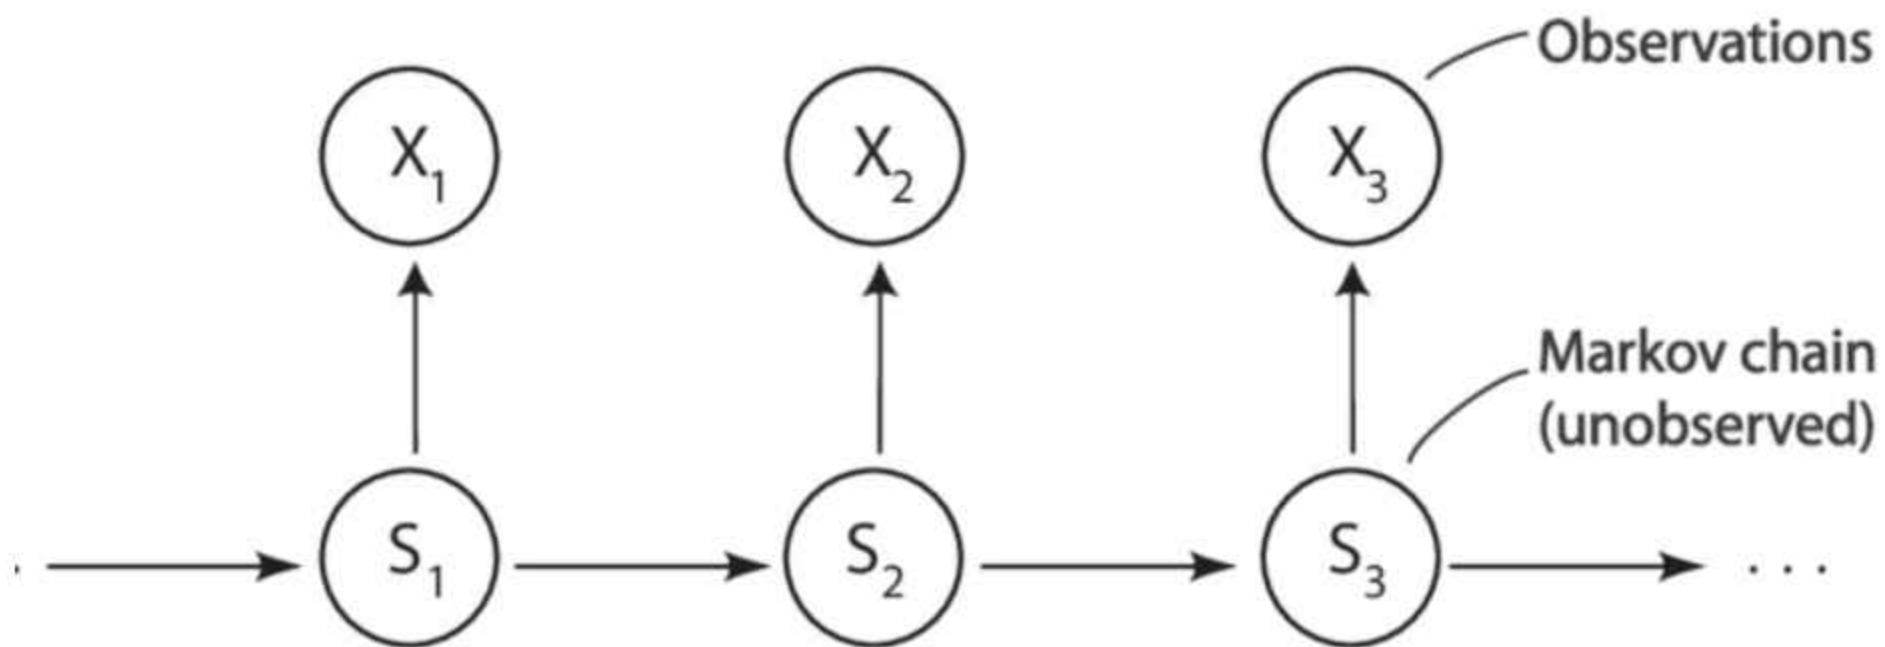
By identifying distinct movement modes from these tracks, researchers can begin to explore the biological processes that drive movement and shape patterns of space use.

HMM (Hidden Markov Models)

- HMMs infer underlying movement behaviours—commonly referred to as "hidden" states—based on observable movement metrics.



HMM



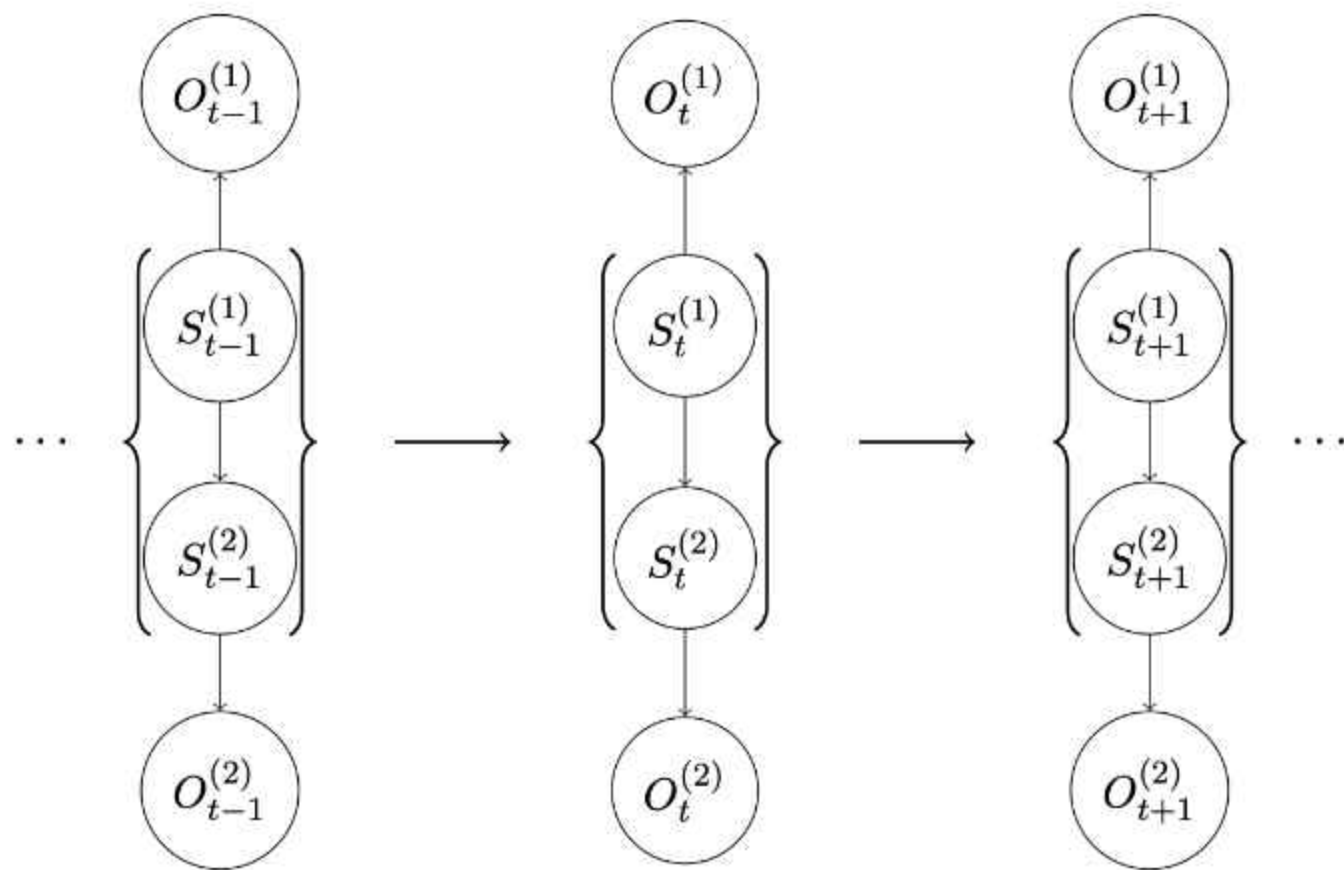
HMM

- More complex models can also incorporate additional state.
- HMMs can estimate the most likely sequence of states along the movement path, effectively providing a segmentation and classification of the track.
- Because step lengths and turning angles can be readily derived from tracking data, they are commonly used to characterize and classify activity modes in movement ecology.
- Animals also display movement patterns at broader temporal and spatial scales, here referred to as movement phases—such as residency, excursions, dispersals or migrations.

HMM

- we propose a simplified and accessible modeling framework that captures cross-scale movement dynamics using movement and space use variables with the same temporal resolution.
- Our framework is based on coupled hidden Markov models (CHMM) which jointly analyze two interacting processes operating at two different scales of movement. HHMMs have previously been applied to movement data
 - one fine-scale process that captures activity modes (inactive vs. moving) using two data streams—step lengths and turning angles—
 - and a coarse-scale process that represents movement phases (resident vs. non-resident) using residence time.

CHMM



Dependence structure of our formulation of an asymmetric coupled hidden Markov model. Observable sequences are denoted by O and latent (hidden) state sequences by S . The movement phase process is represented by (1) and the activity mode process is represented by (2). t indicates the time step.

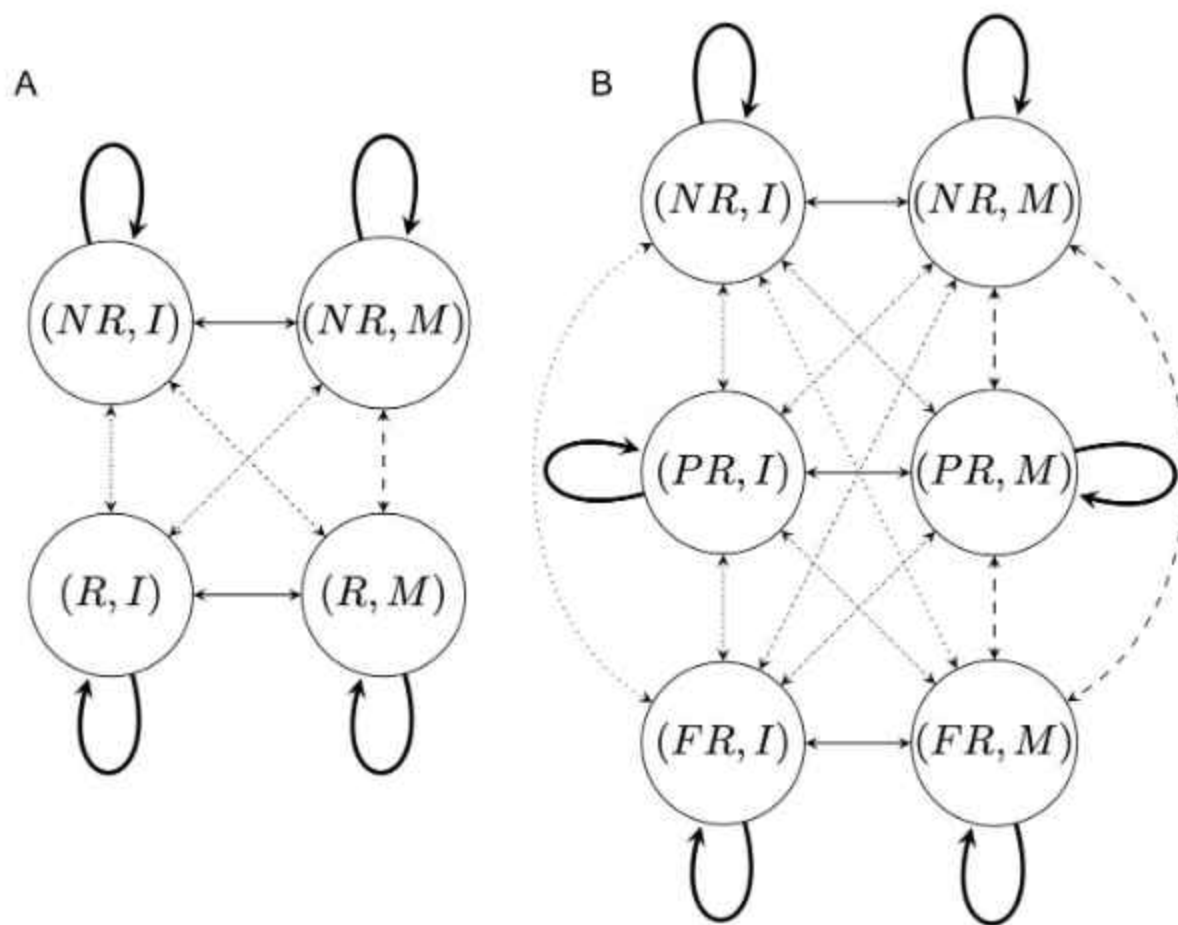
HMM

We first considered an CHMM specification that distinguished between two primary movement phases. To this end, we developed a 4-state model representing combinations of the two latent processes—activity mode (inactive *I* or moving *M*) and movement phase (resident *R* or non-resident *NR*)—resulting in four discrete states.

$$\Gamma(4) = \begin{pmatrix} a_{NR\ NR} & a_{NR\ R} \\ a_{R\ NR} & a_{R\ R} \end{pmatrix} \otimes \begin{pmatrix} b_{I\ I} & b_{I\ M} \\ b_{M\ I} & b_{M\ M} \end{pmatrix} = \begin{pmatrix} c_{NR\ I \rightarrow NR\ I} & c_{NR\ I \rightarrow NR\ M} & c_{NR\ I \rightarrow R\ I} & c_{NR\ I \rightarrow R\ M} \\ c_{NR\ M \rightarrow NR\ I} & c_{NR\ M \rightarrow NR\ M} & c_{NR\ M \rightarrow R\ I} & c_{NR\ M \rightarrow R\ M} \\ c_{R\ I \rightarrow NR\ I} & c_{R\ I \rightarrow NR\ M} & c_{R\ I \rightarrow R\ I} & c_{R\ I \rightarrow R\ M} \\ c_{R\ M \rightarrow NR\ I} & c_{R\ M \rightarrow NR\ M} & c_{R\ M \rightarrow R\ I} & c_{R\ M \rightarrow R\ M} \end{pmatrix}$$

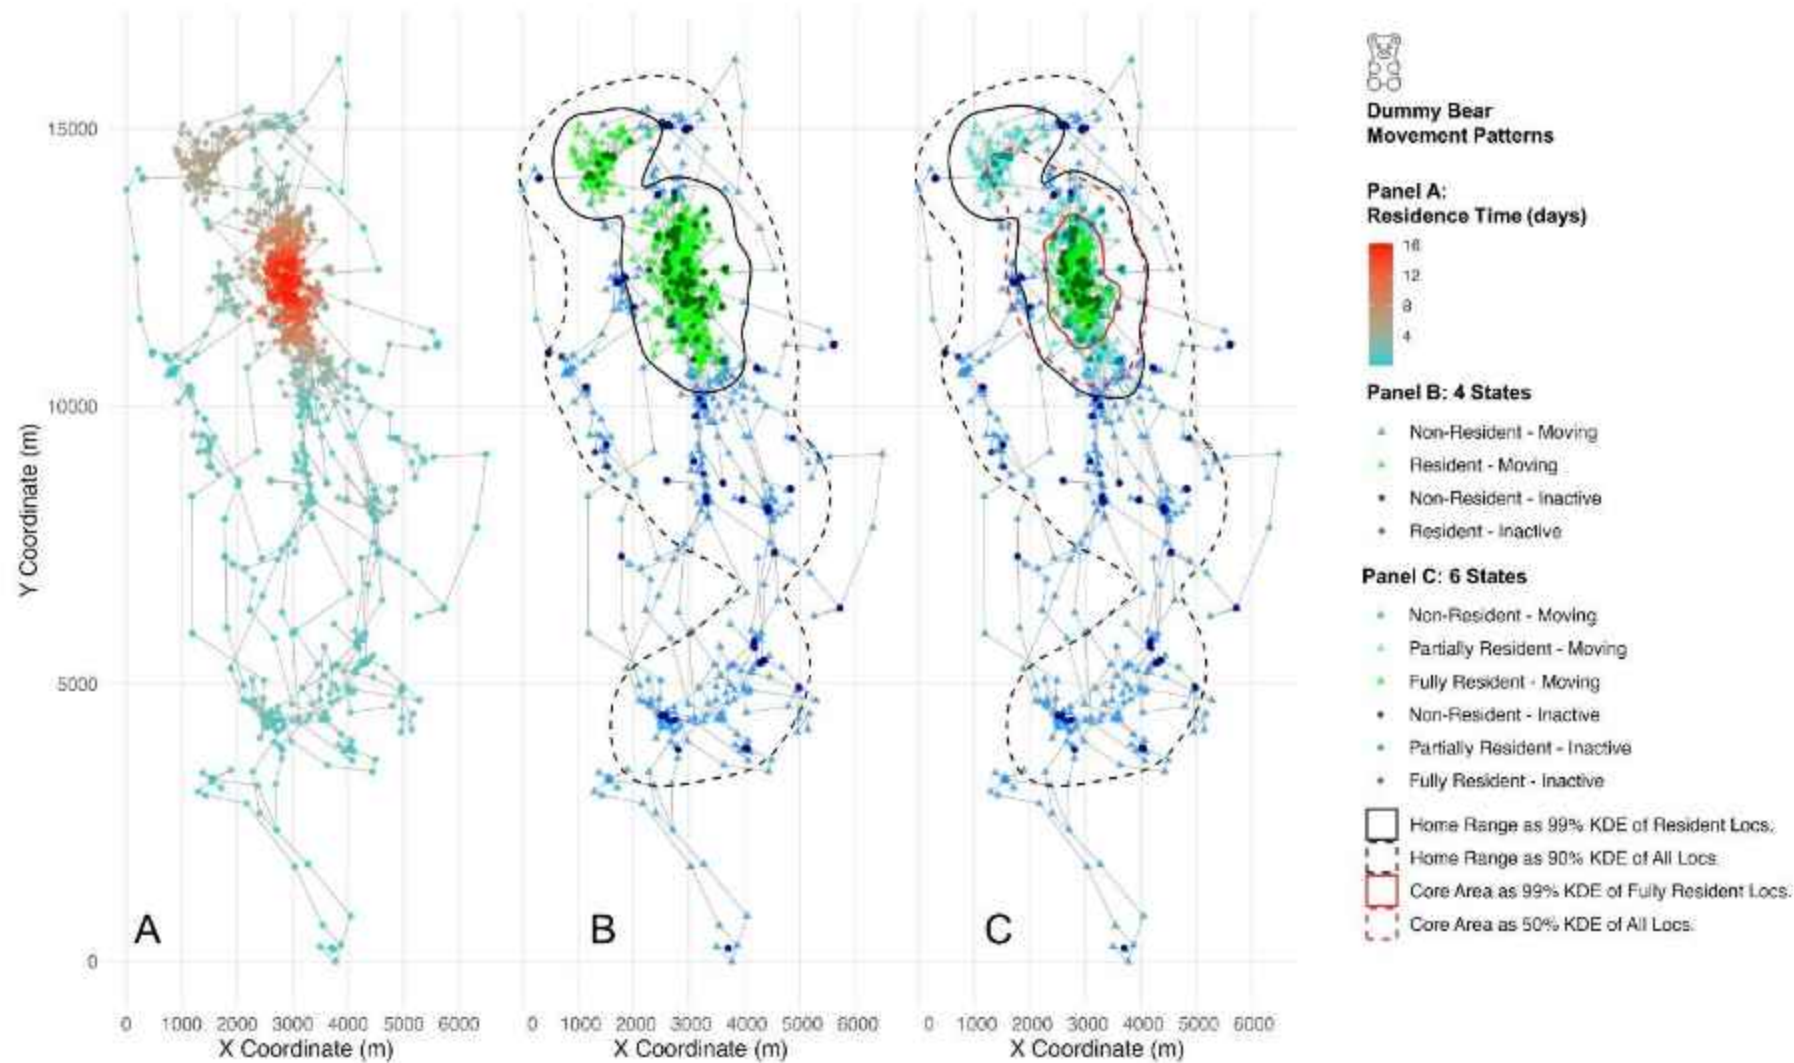
where *a* refers to the movement phase, *b* to the activity mode and *c* to the resulting transitions between states.

HMM

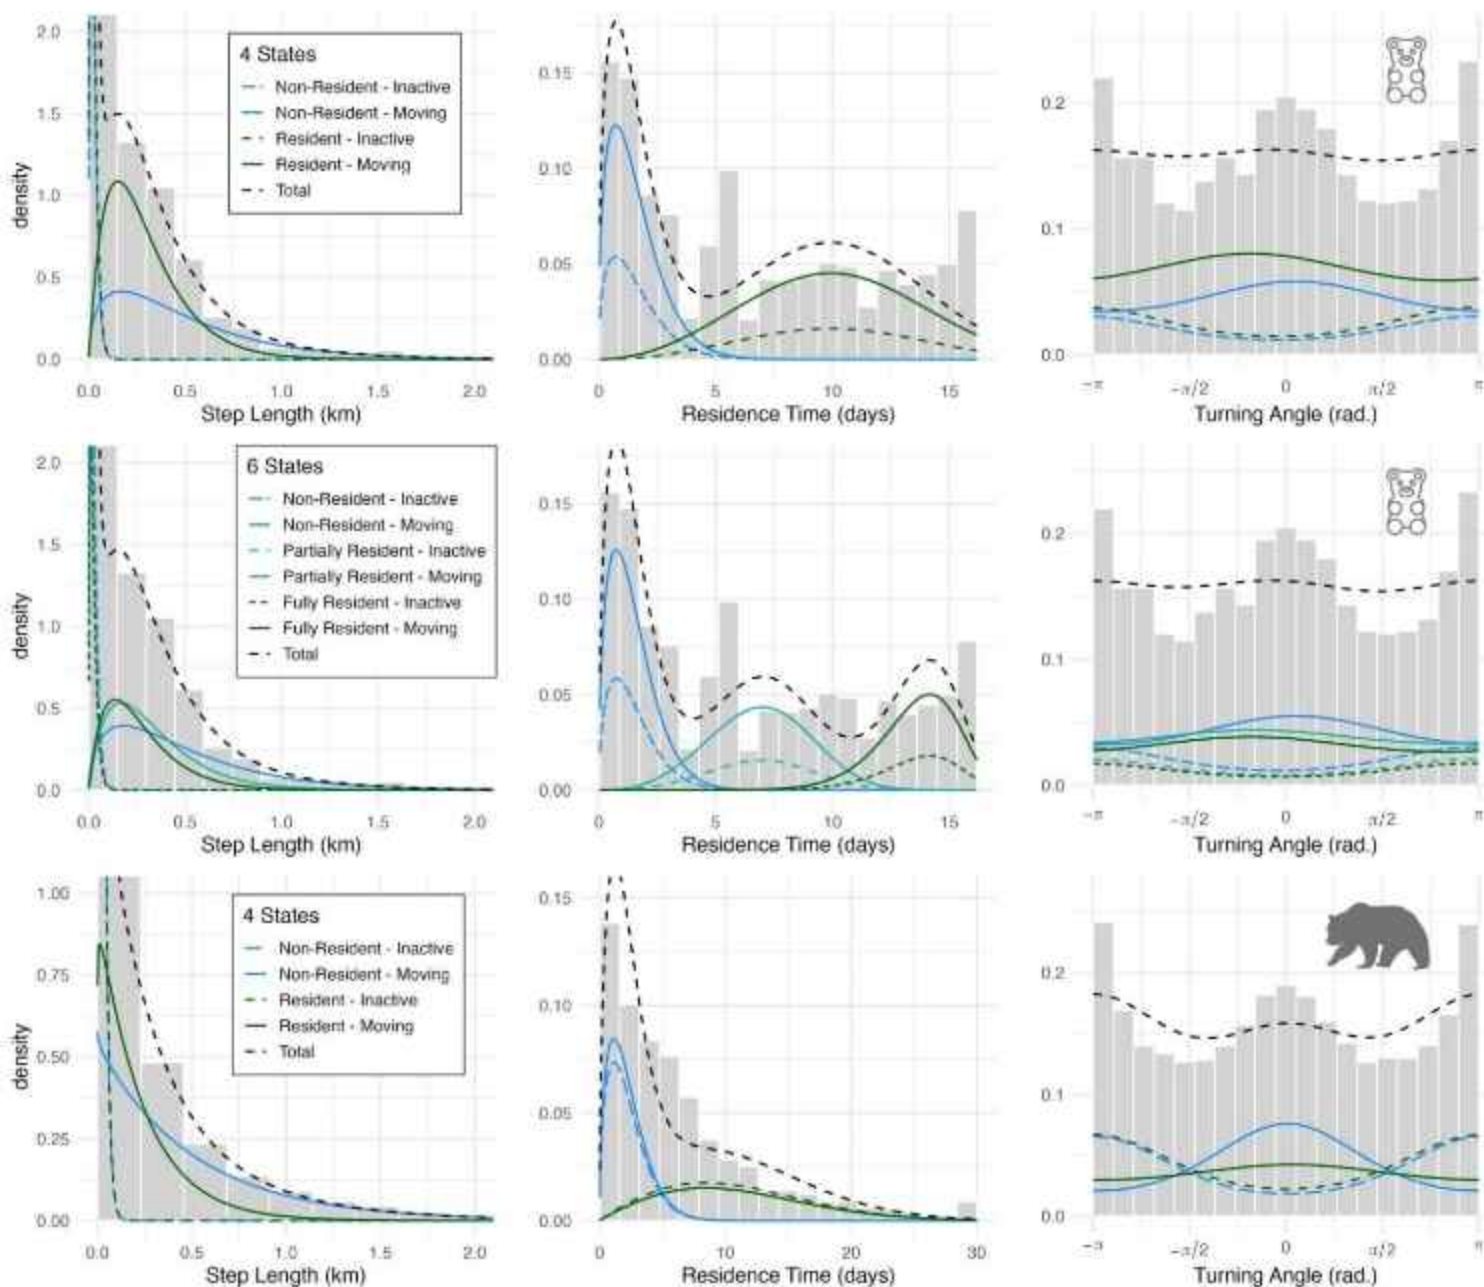


Schematic representation of the transition matrices for the (A) 4-state and (B) 6-state CHMMs.

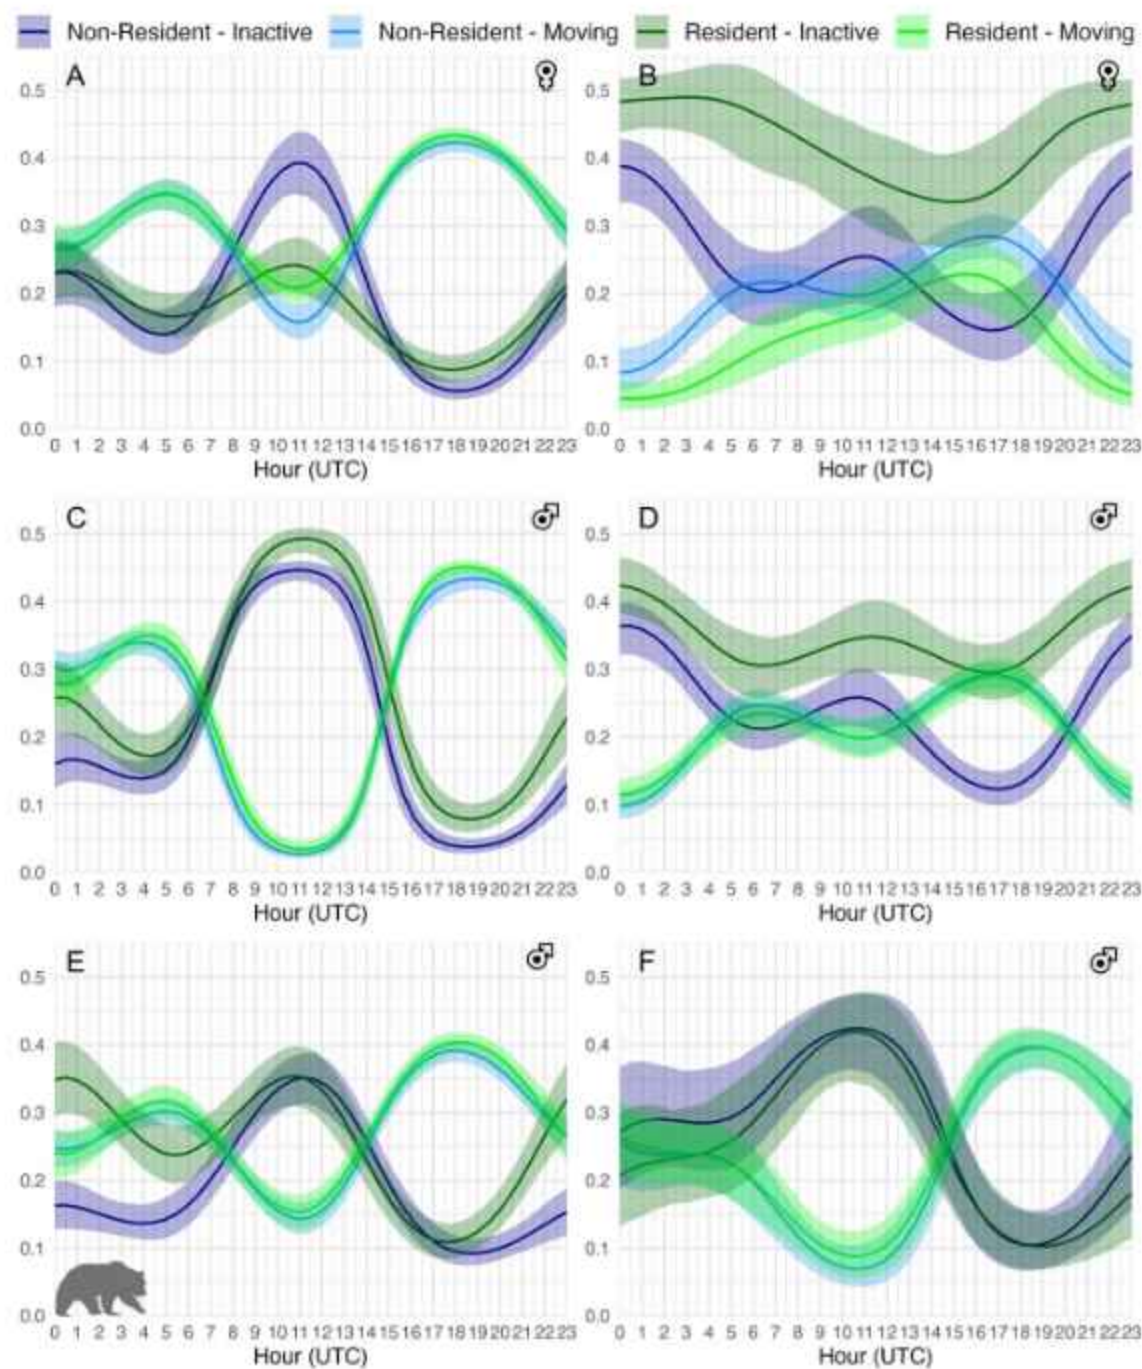
- In the 4-state model, states represent combinations of movement phase and activity mode: non-resident-inactive ($NR-I$), non-resident-moving ($NR-M$), resident-inactive ($R-I$), and resident-moving ($R-M$).
- The 6-state model includes an intermediate movement phase and consists of: non-resident-inactive ($NR-I$), non-resident-moving ($NR-M$), partially resident-inactive ($PR-I$), partially resident-moving ($PR-M$), fully resident-inactive ($FR-I$), and fully resident-moving ($FR-M$).



Movement patterns of the dummy bear: (A) GPS trajectory colored by residence time; (B) State-decoded time series from the 4-state CHMM, overlaid with home range contours estimated via kernel density estimators (KDE); (C) State-decoded time series from the 6-state CHMM, with KDE-derived contours delineating both home range and core areas.



Estimated state-dependent distributions for the three fitted CHMMs. Each row shows the distributions for step length (left), residence time (centre), and turning angle (right). Top row: 4-state model fitted to the dummy bear trajectory. Middle row: 6-state model fitted to the dummy bear. Bottom row: 4-state model fitted to Cantabrian brown bear trajectories.



Stationary state probabilities (y-axis) as a function of time of day (x-axis; hour in UTC) for six individual Cantabrian brown bears (2 females in panel A-B and 4 males in C-F) representing a range of circadian activity patterns. Corresponding plots for the remaining individuals are provided in Supporting Fig. 1.

Thanks for your attention



Oscar Rodriguez de Rivera Ortega
Lecturer in Statistics, Faculty of Environment, Science and Economy
o.m.rodriguez-de-rivera-ortega@exeter.ac.uk

BioCLIP-world:

Sept 2025

Progress and challenges towards a multimodal embedding space for the tree of life

Ross Gardiner

(Postgrad, EI CDT, Exeter)

Supervisors: Benno Simmons, Sareh Rowlands.

Contact: rg654@exeter.ac.uk



Contents

- 01 What are Foundation Models?
- 02 Self-supervised learning primer
- 03 Foundation model examples
- 04 BioCLIP-world
- 05 Datasets and evaluation
- 06 Modelling

What is a Foundation Model?



I'm a computer vision engineer

Foundation models are typically **large scale architectures** which have been trained on **vast quantities of data** using **weak labels** and/or **self-supervision**.

E.g. GPT



I'm a user of AI systems

Foundation models are something I might use to **help me do work**, **train further** at **low cost** or **use to access AI** where I previously wouldn't have.

E.g. ChatGPT



I'm interested in automated species monitoring

A huge **burden** here is **labelling datasets**. Thus ecologists are increasingly **training their own bespoke models**. Foundation models can make this **faster and easier**.

Foundation models do not have to be computationally expensive (but some have been historically).

CLIP Objective + SelfSupervision

MONTH YEAR

Supervised learning

Discrete labels

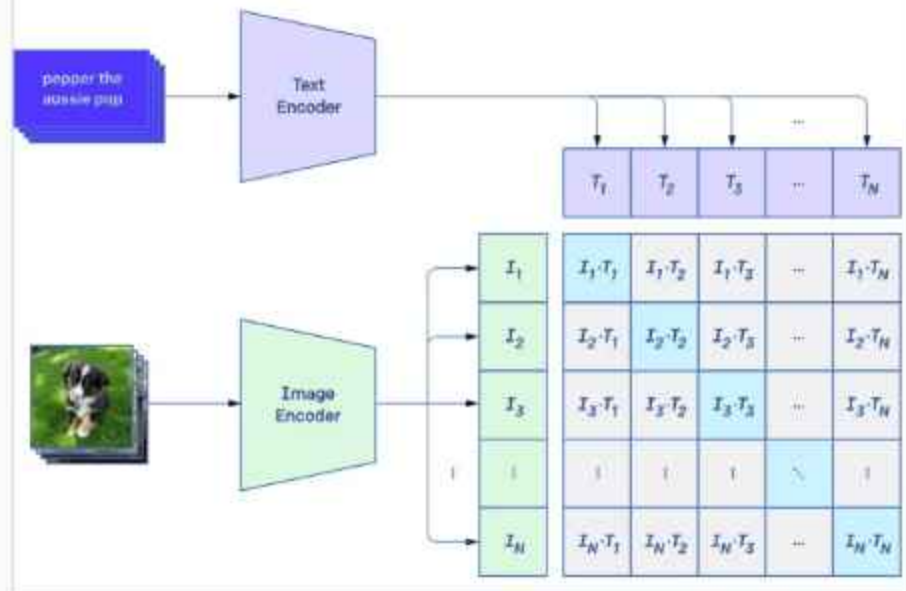


Self-supervision

(Weak labels)



1. Contrastive pre-training



Contrastive Language-Image pretraining is a way to align general text "this is a dog" with an image. This is many-to-many relationship. Which means that we don't need discrete labels.

CLIP allows the generalist relationship between image and text to be learned, which unlocks training data.

Foundation Model Examples

Model	Training Size	# Params	Year	Domain
GPT-OSS	~13 trillion text tokens	120B (oss-120b)	2025	General
Open-CLIP	~2 billion Image-text pairs	428M (L-variant)	2021	General
DINOv3	1.7 billion images	304M (Large variant)	2025	General
BioCLIP	10M Image-text pairs	150M	2023	Nature images and species names Natural world images
BioTrove-CLIP	40M Image-text pairs	150M	2024	Nature images and species names Natural world images
BioCLIP2	200M Image-text pairs	428M	2025	Nature images and species names Natural world images
BioCLIP-world	~60M Image-text-metadata triplets	?	202X	Nature images, species names and spatio-temporal priors Natural world

“Big tech”, e.g.
Facebook, OpenAI

Other research groups

ME!

Priors

Modelling prior knowledge has historically been important for building supervised AI models that **fit** well.

Importantly, this was a popular paradigm **before** large-scale pre-training took off.

Some priors give big **contextual hints** and can be interfaced with AI models with far **less computational demand** than images and text alone.

This is the philosophy behind BioCLIP-world:

*Finding a route towards a biodiversity image foundation model which both **fits the data** and is more **computationally efficient** to train.*

Prior knowledge...



I'm in the desert, so I'm less likely to see a polar bear.



It's the winter, so this sycamore may not have leaves; the image will look different.



It's night time, so I may see certain species and/or lighting conditions.

"Thinking like a naturalist"

Datasets, testing and Biodiversity

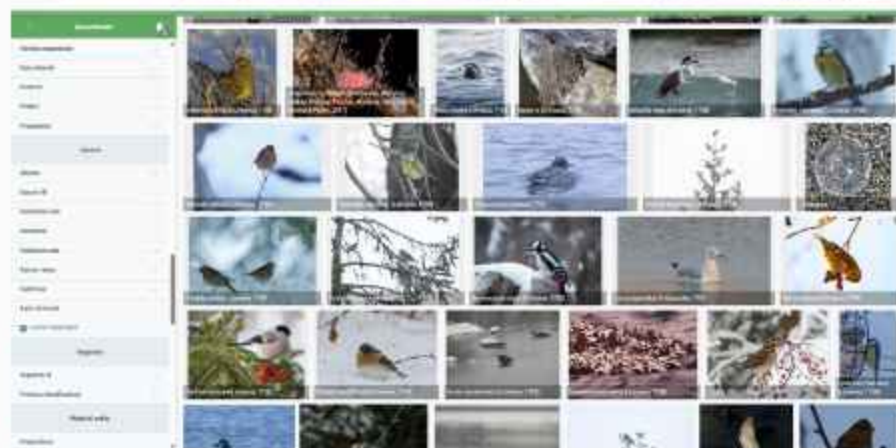
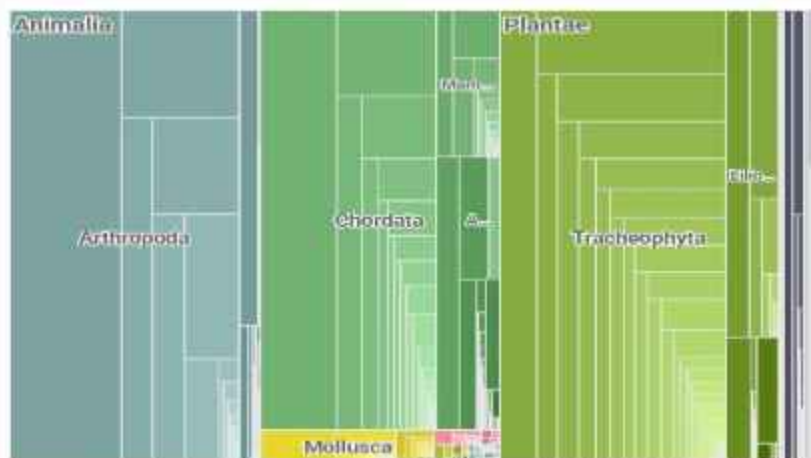
MONTH YEAR

A large dataset of images and species names was assembled using the **GBIF** - which hosts mostly citizen **science observations** from around the world.

- **Lat-lon locations**
- **Observation time-of-day**
- **Observation time-of-year**

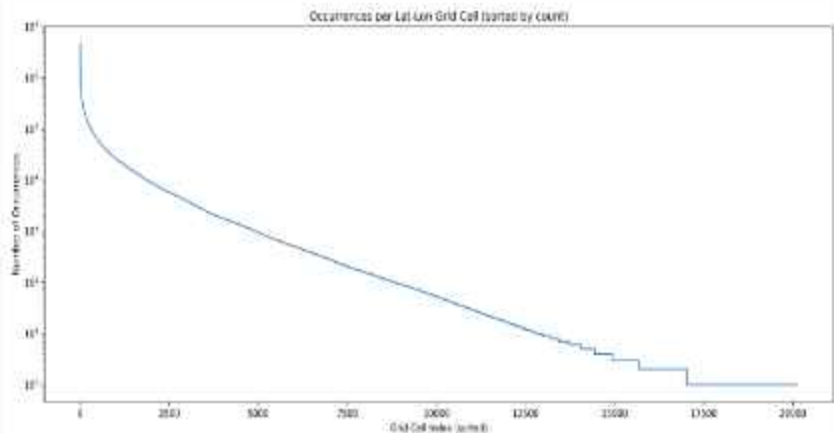
Built a processing pipeline to extract and manage observations.

Dealing with **long-tailedness** is challenging.

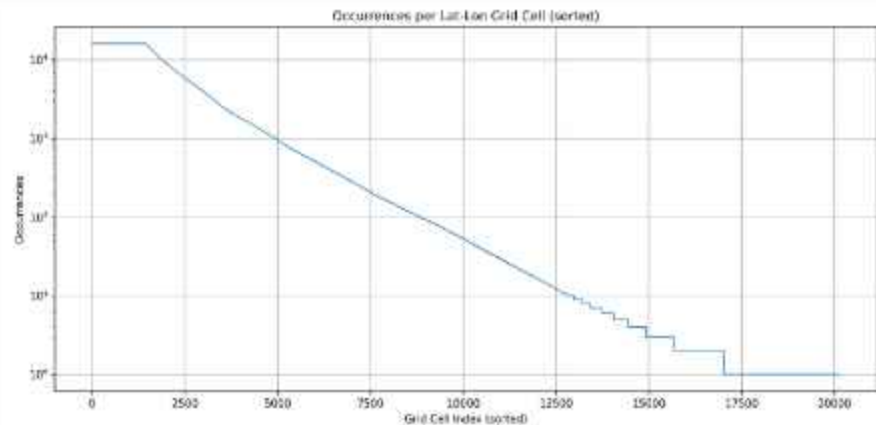
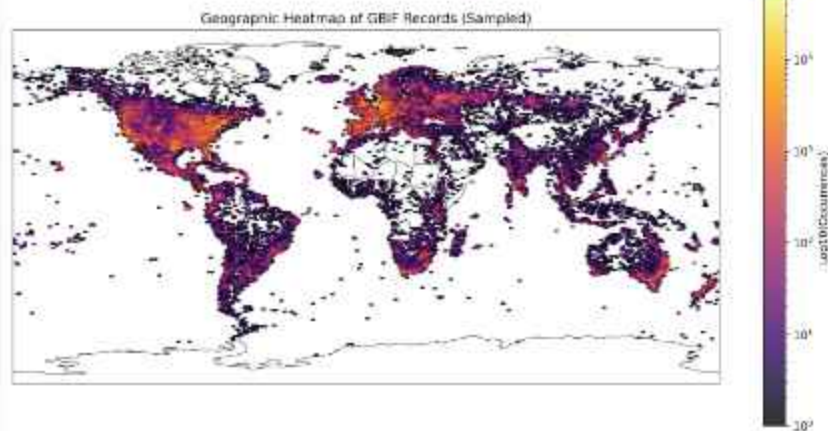


Dataset	Nr Training Pairs	Nr Species	Metadata
BioCLIP	10M	454K	no
BioTrove	40M	33K	no
BioCLIP2	200M	952K	no
Ours	60M	131K	yes

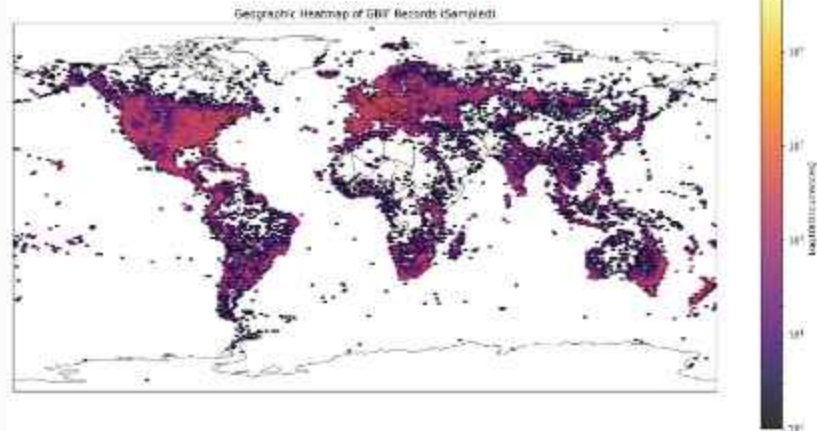
Long Tails



~160M Samples

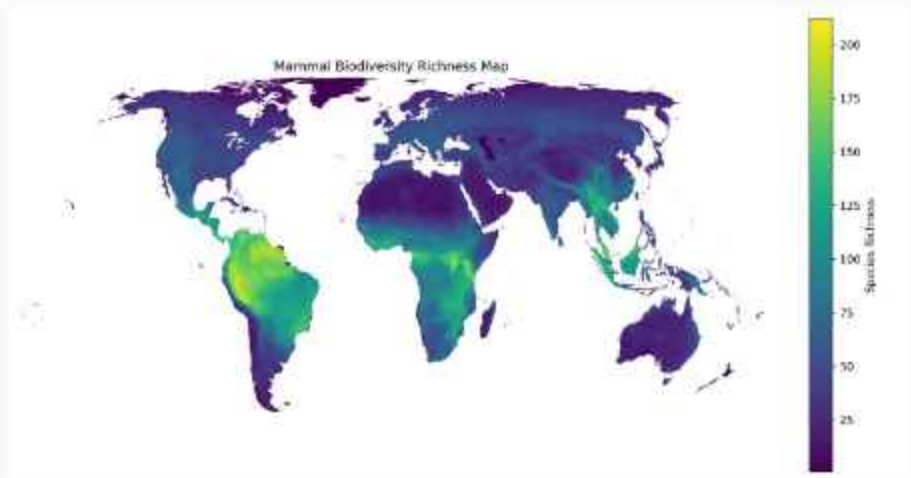
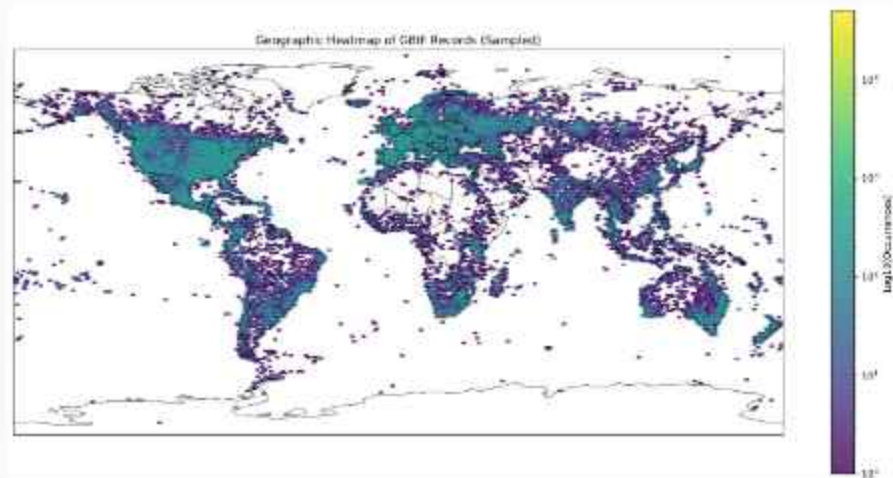


~60M Samples



Datasets, testing and biodiversity

MONTH YEAR



Biodiversity-rich regions do not necessarily align with **where** we have data...

Test sets developed to be **biodiversity-weighted** in the first effort to evaluate the impact of this.

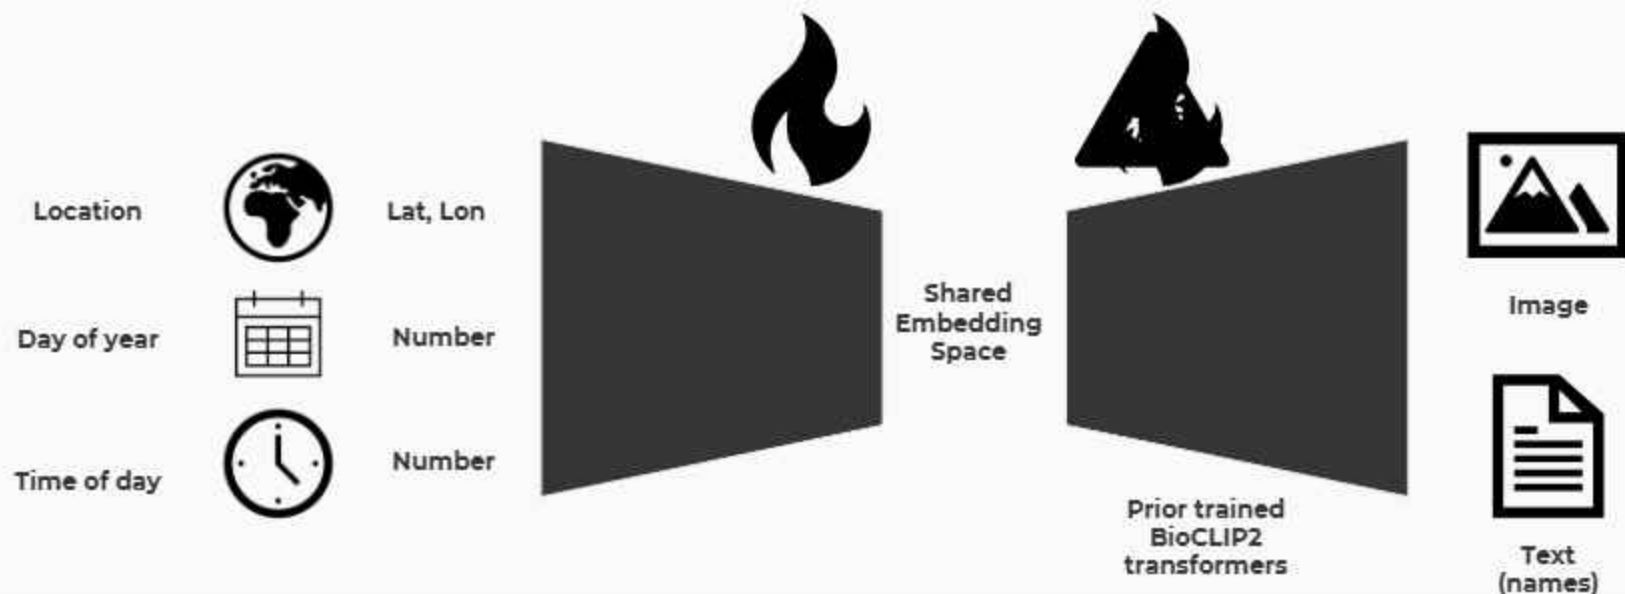
Three sets of **1902 mammal** species, **8022 birds** and **2034 amphibians**. Although these are in **early stages**.

Also created **balanced species** test sets across all species.



Encoder training

MONTH YEAR



Training works in two phases:

- 1. Image+text encoder remains frozen.** An encoder is found to map metadata to the current embedding.
 - a. Currently exploring **random fourier features** for this, but I'm **open to suggestions!**
- 2. Image and text encoders are allowed to move weights.** Can learn from the location encoding, but **compute is greatly increased.**
 - a. Some initial success in using **Low-Rank adaptation** to significantly reduce training burden.

Low-Rank Adaptation (LoRA)

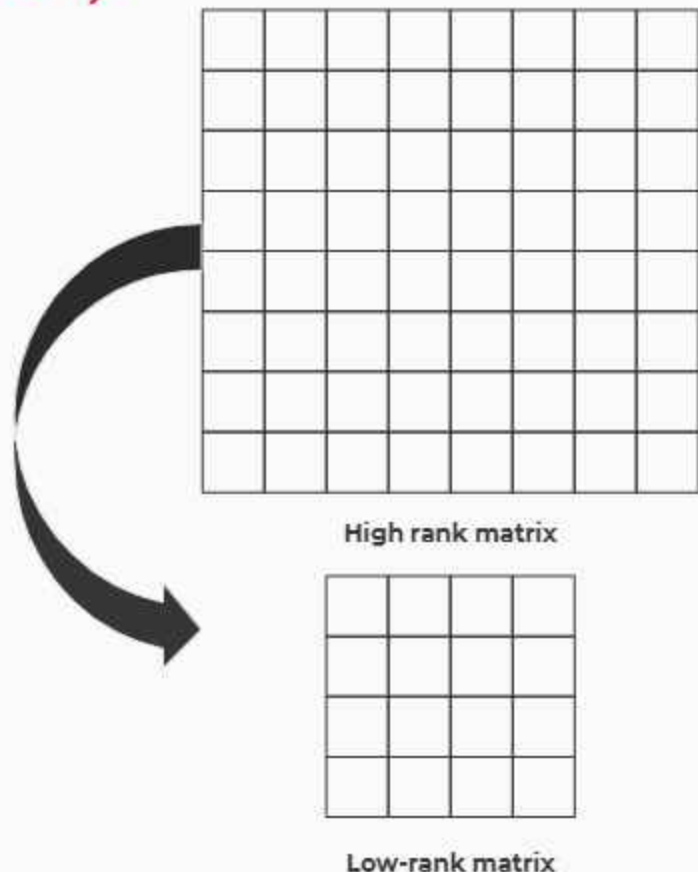
To **efficiently** train large models, we can use the technique, **LoRA**.

The **key idea** is that for a given matrix of weights inside the model, a lower-rank matrix (i.e. smaller) of weights can be used to **approximate** it, and thus **trained in its place**.

This reduces computational demand by as much as a **250x reduction in training params** (original paper). Greatly reducing the number of weights to update in backpropagation.

LoRA has mainly been developed/used for **language models**.

Recently applied to CLIP-style vision-language. Allows BioCLIP-world training stage 2, to be made efficient.





Final Remarks

- Foundation models for biodiversity monitoring are here/coming.
- I aim to advance these **and** push development towards a **data and computationally efficient route** through **multimodality**.
- I also aim to spread conversations around model **evaluation from a conservation perspective**.

P.S. Come and check out my poster in the main hall:

"Bridging Domain Gaps for Fine-Grained Moth Classification Through Expert-Informed Adaptation and Foundation Model Priors"



Uncovering Systematic Errors in Camera Trap Species Classification Models

Viktor Domazetoski

11.09.2025

QUEX INSTITUTE



University
of Exeter



THE UNIVERSITY
OF QUEENSLAND
AUSTRALIA

AI for Camera Traps

Animal Detection



Species Classification



Individual Reidentification



Illumination | Blur



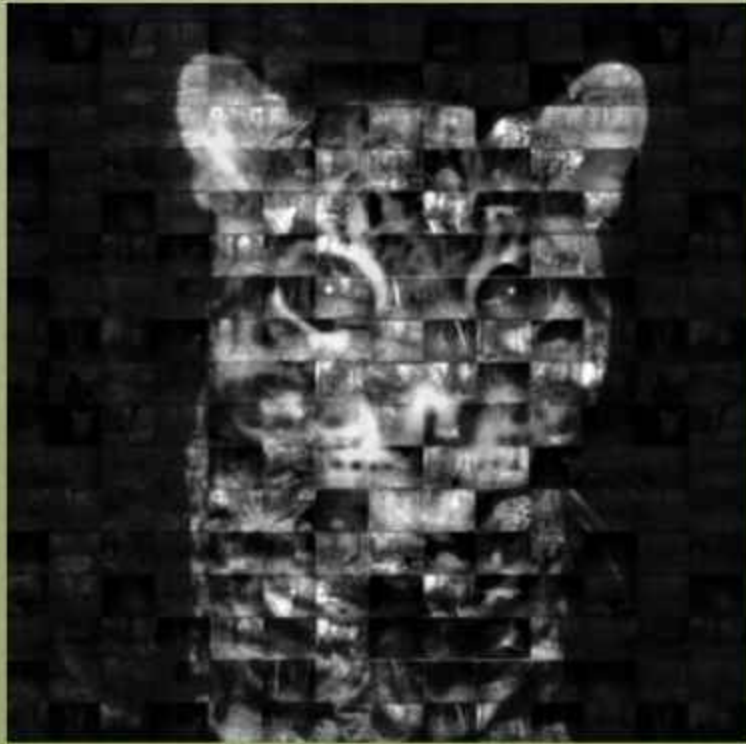
Perspective



AI for Camera Traps



Wildlife Insights



Camera Traps

Observations

Model

Prediction

Evaluation (Metrics)

Trap #1



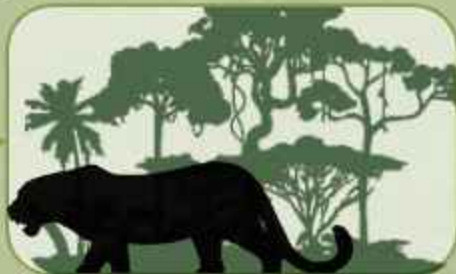
Trap #2



Trap #3



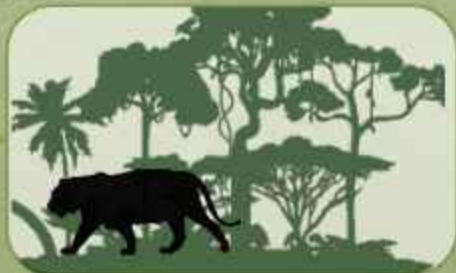
Obs. #1



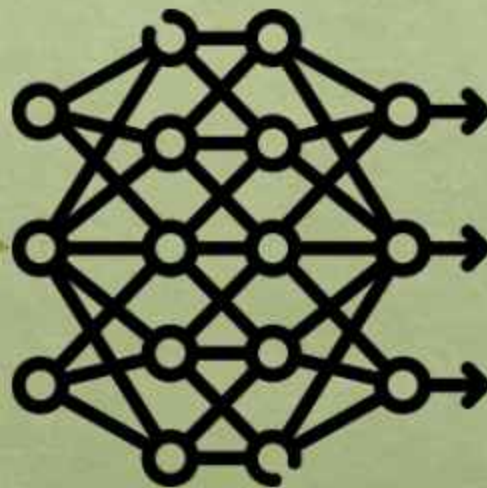
Obs. #2



Obs. #3



Obs. #4



Panthera tigris

Score: 0.33

Panthera pardus

Score: 0.78 Error: 0.22

Blank

Score: 0.07

Precision: 0.89

Recall: 0.85

Camera Traps

Observations

Model

Prediction

Evaluation (Metrics)

Trap #1



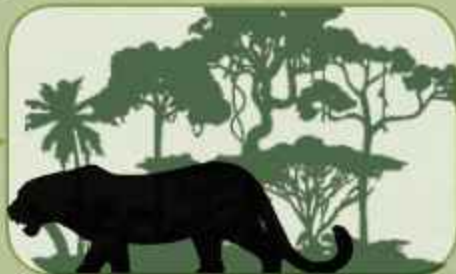
Trap #2



Trap #3



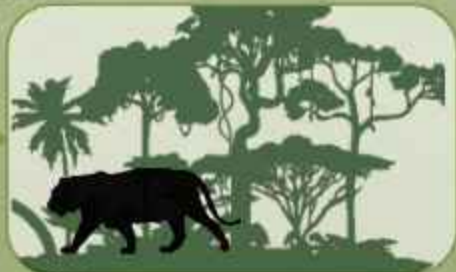
Obs. #1



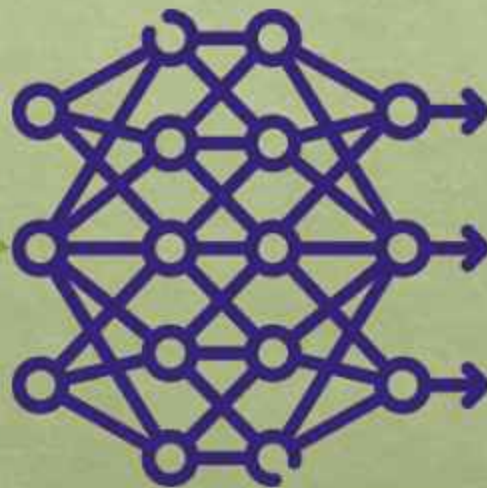
Obs. #2



Obs. #3



Obs. #4



Panthera tigris

Score: 0.33

Panthera pardus

Score: 0.78 Error: 0.22

Blank

Score: 0.07

Precision: 0.89

Recall: 0.85

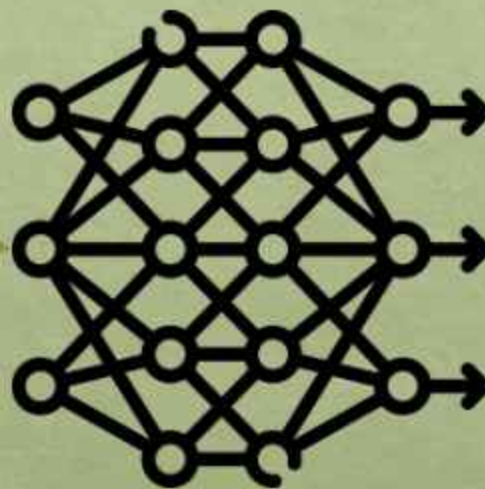
Camera Traps

Observations Species

Model

Prediction

Evaluation (Metrics)



Panthera tigris
Score: 0.33

Panthera pardus
Score: 0.78 Error: 0.22

Blank
Score: 0.07

Precision: 0.89

Recall: 0.85

Data

◆ Snapshot Safari 2024 Expansion (subset)

- ◆ 8 projects across South Africa
- ◆ 227 sites
- ◆ 60 species + empty
- ◆ 400,000 + observations



Model

- SpeciesNet

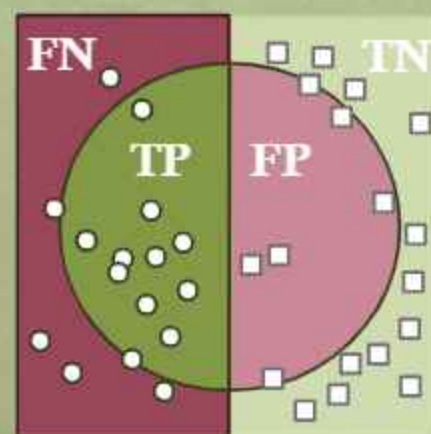
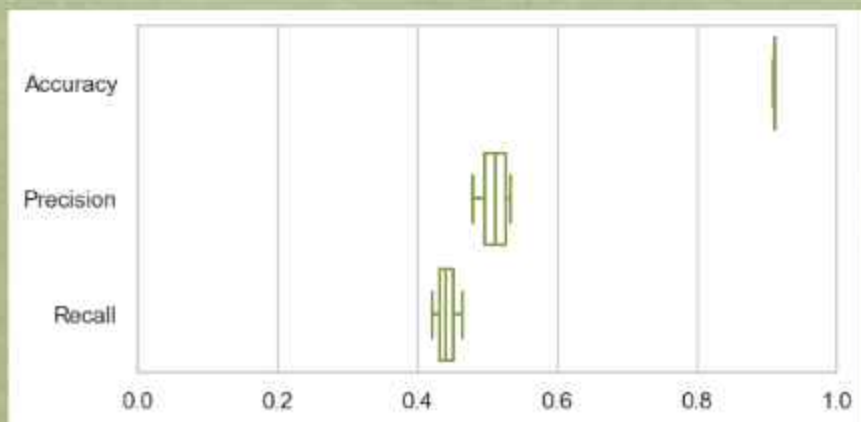
- Gadot, T. (2024). IET Computer Vision, 18(8).



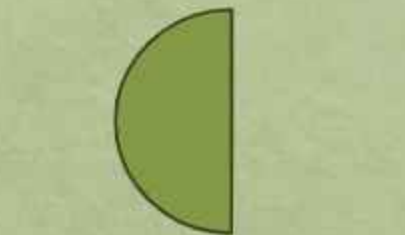
Species	# of images	Recall	Unknown Rate
African Buffalo	5047	82.94%	1.39%
Douglas's squirrel	5310	59.87%	12.81%
Nutria	5540	96.32%	0.14%
Eurasian Lynx	31	51.61%	3.23%

Performance Analysis

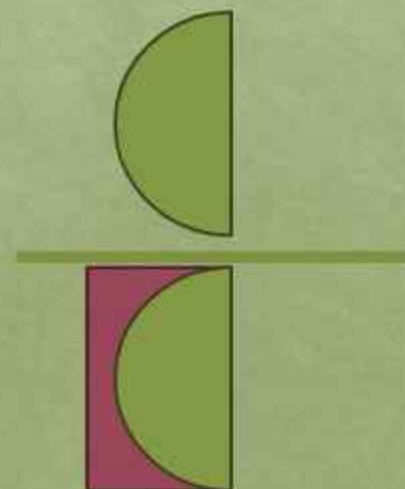
Overall Model Performance



$$\text{Precision} = \frac{TP}{TP + FP} =$$

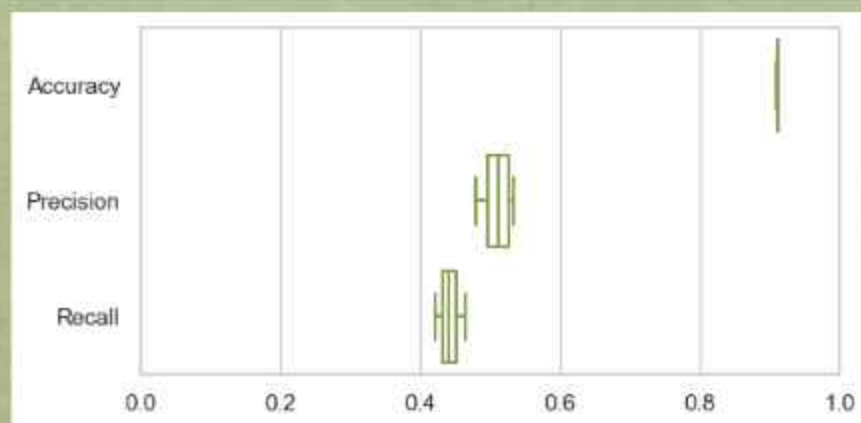


$$\text{Recall} = \frac{TP}{TP + FN} =$$

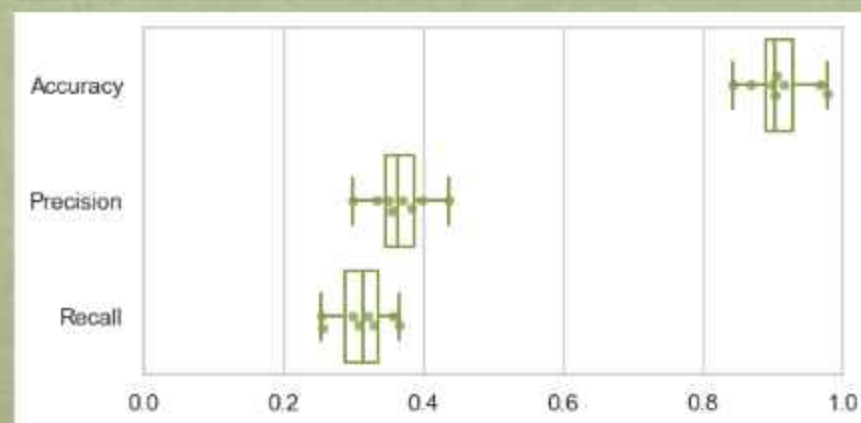


Performance Analysis

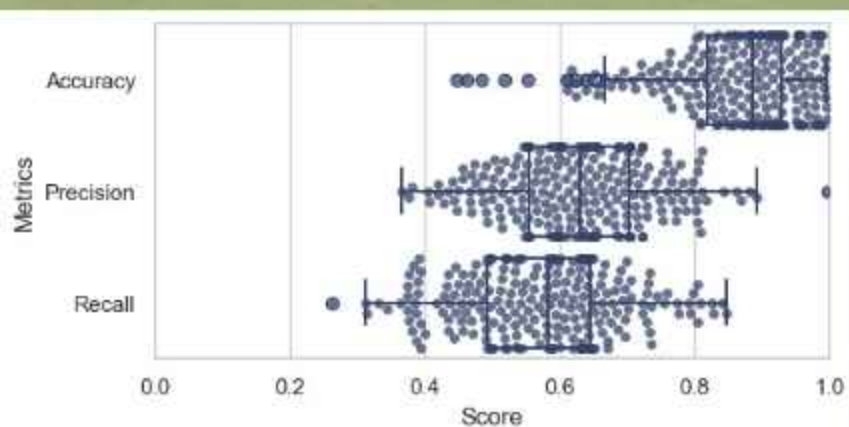
Overall Model Performance



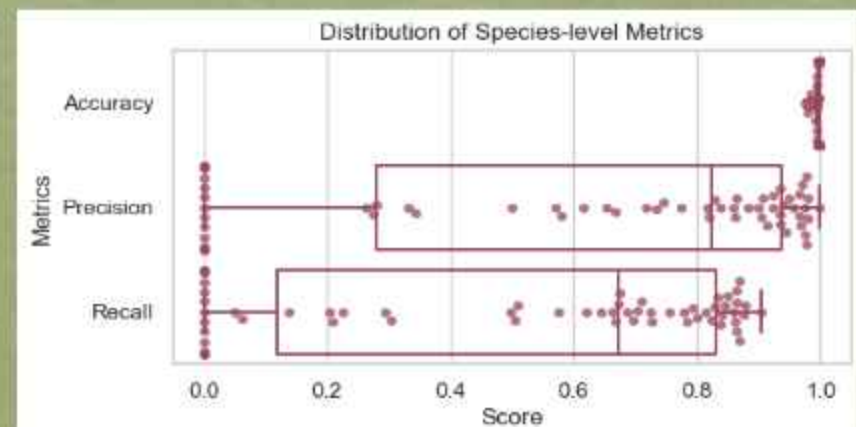
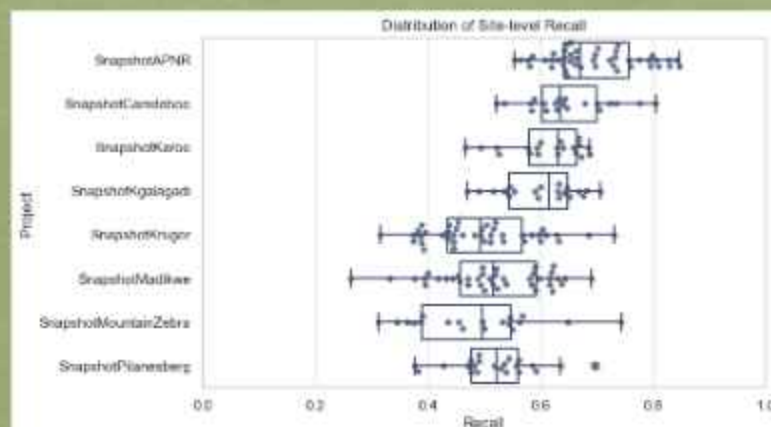
Project-level Performance



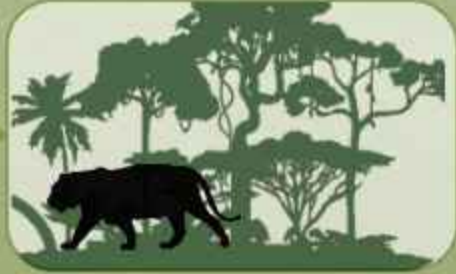
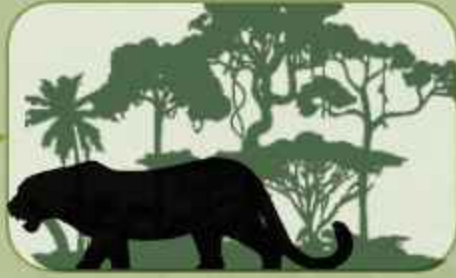
Site-level Performance



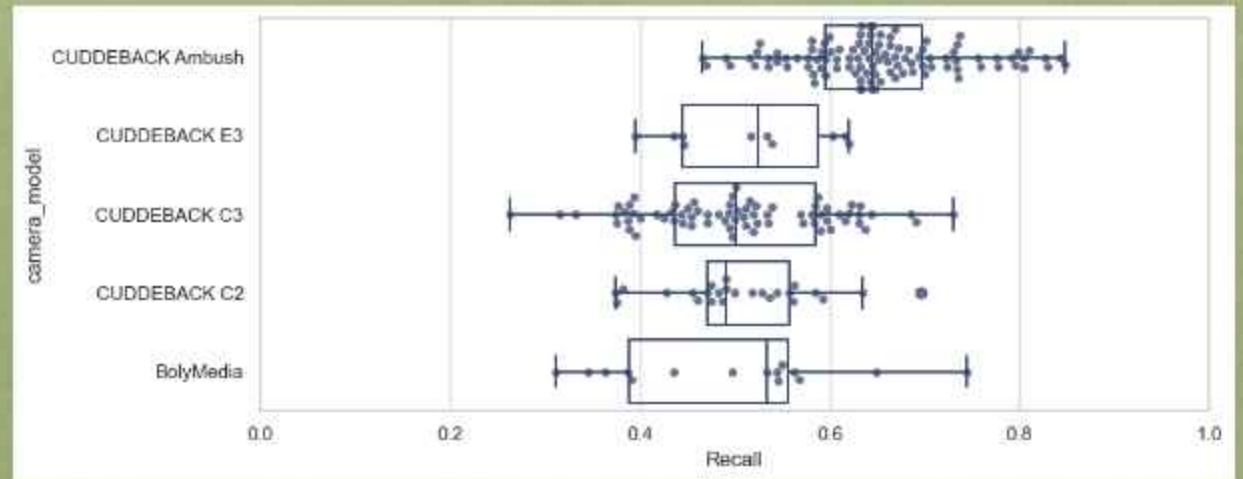
Species-level Performance



Site Covariates



Camera Model

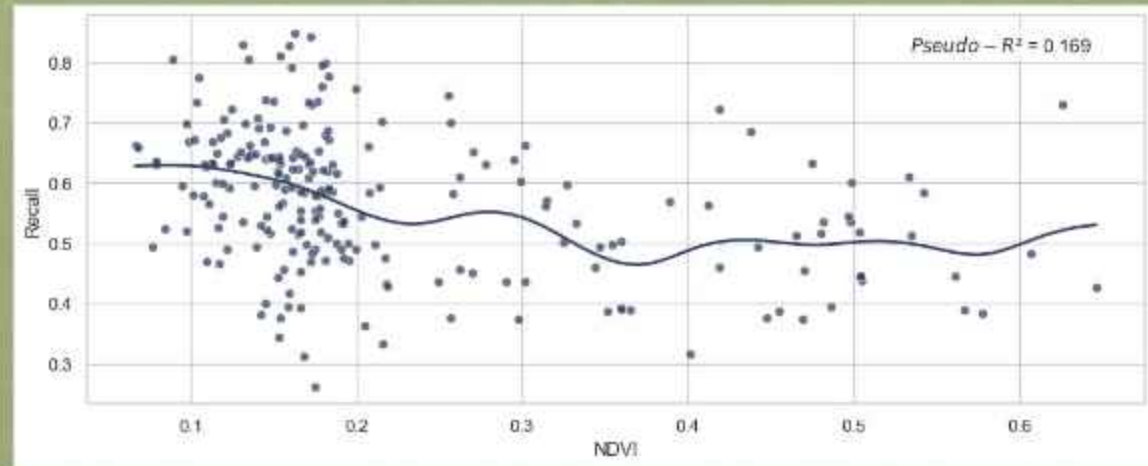
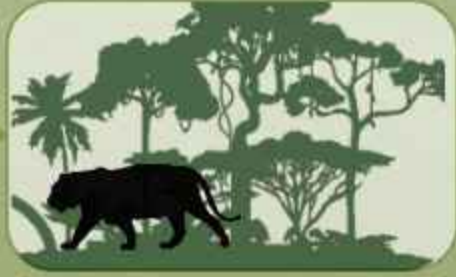
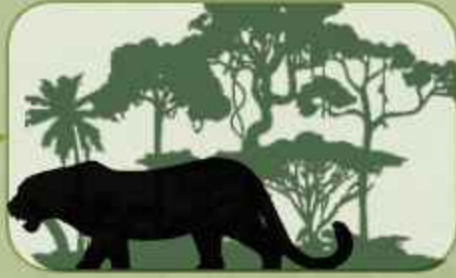


Site Covariates

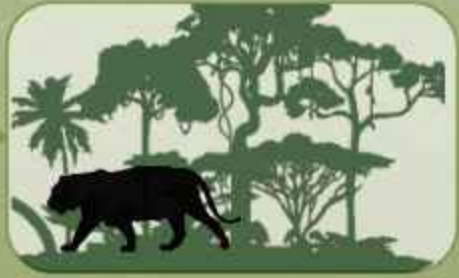
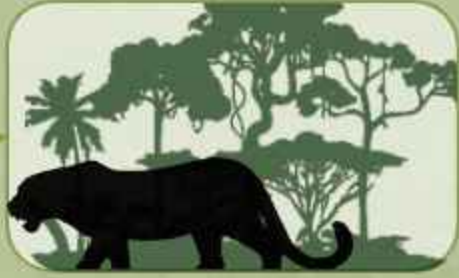
Camera Model



NDVI



Site Covariates



Camera Model



NDVI



Precipitation



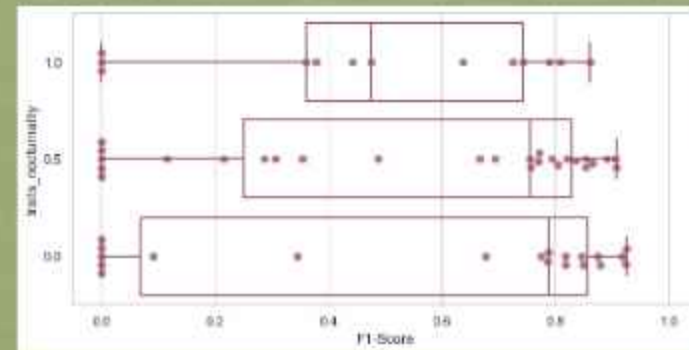
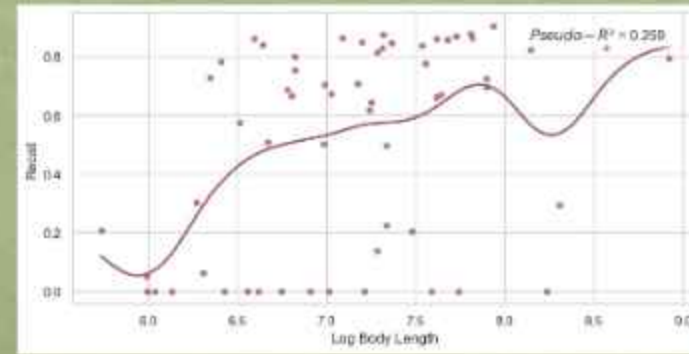
Geography



Species Covariates

Taxonomy

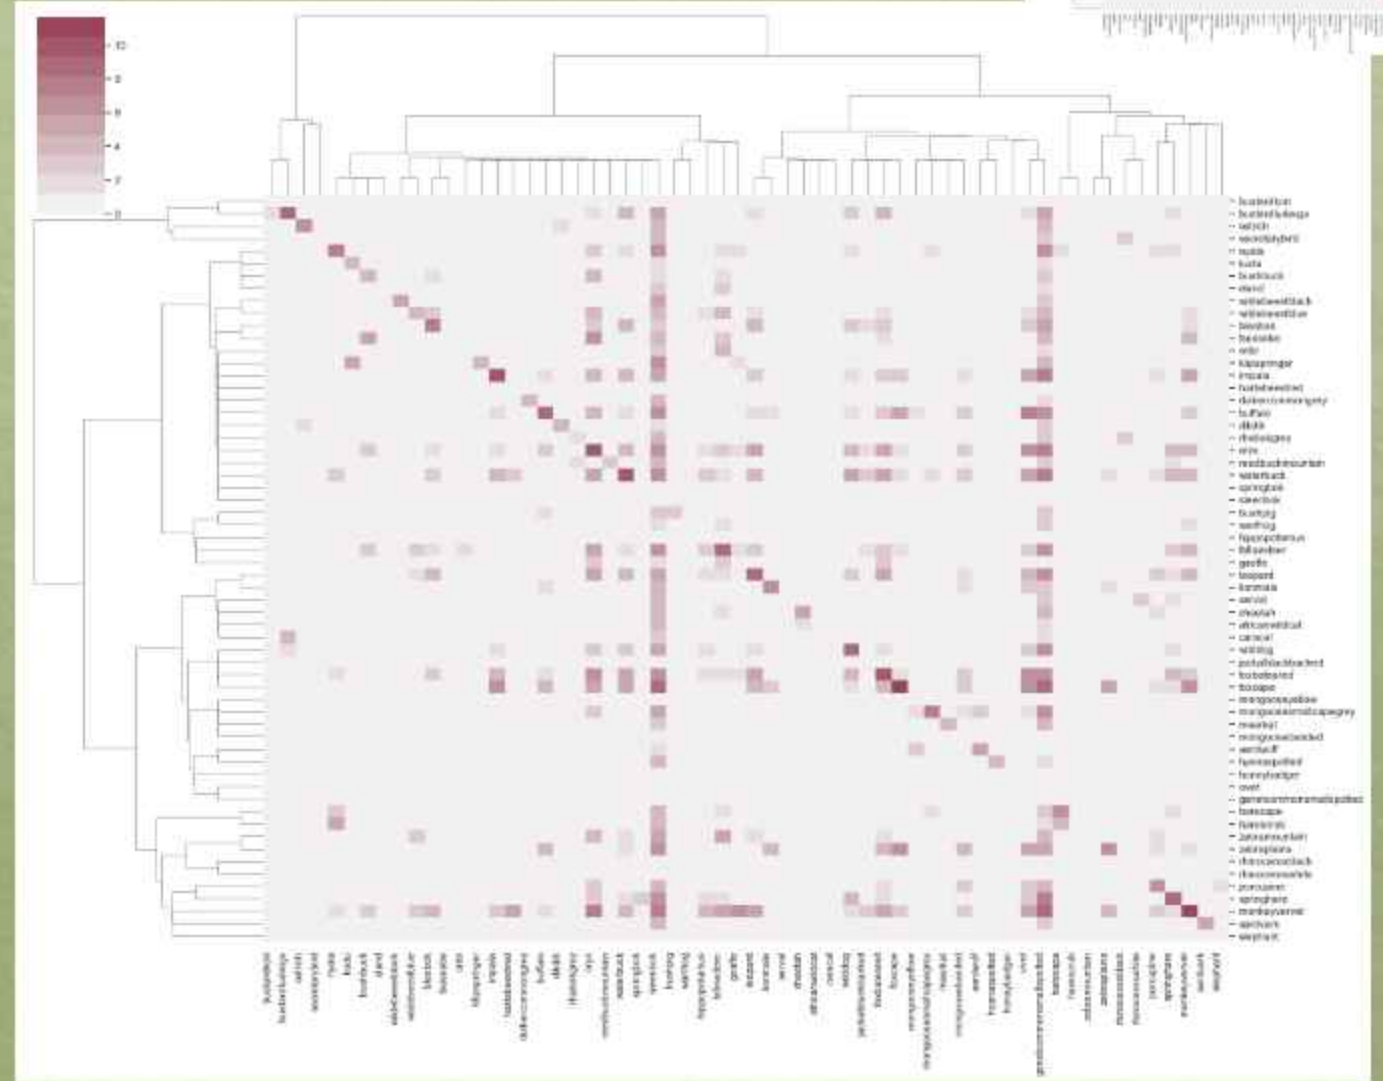
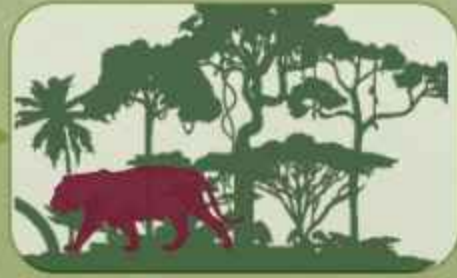
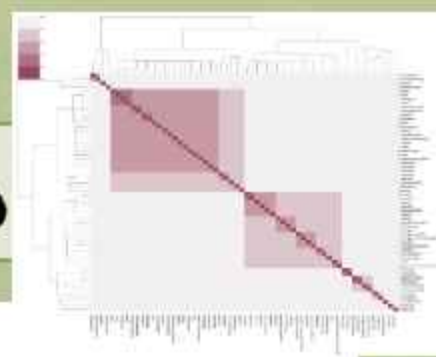
Traits



Species Covariates

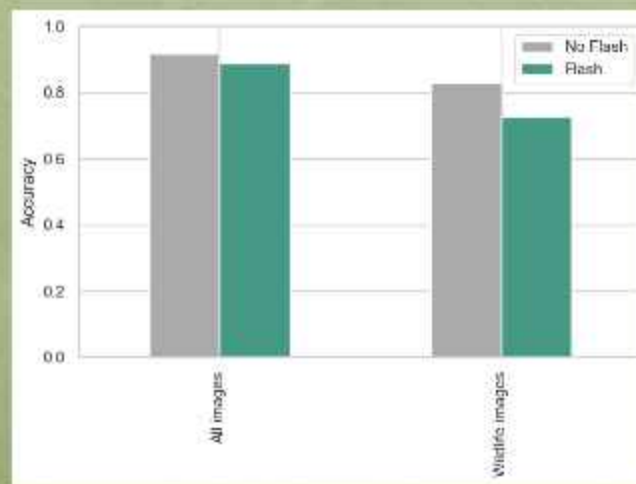
Taxonomy

Traits



Observation Covariates

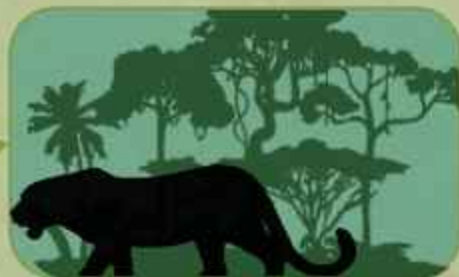
Time of day



Bounding Box Area



Exposure Time | ISO speed

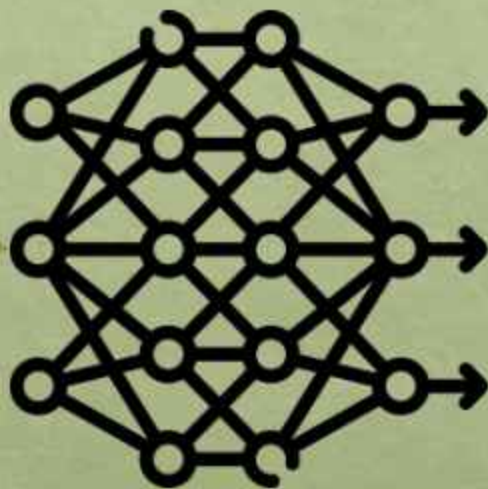
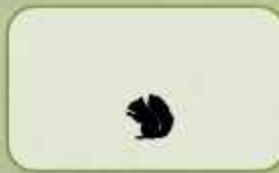


NDVI

Camera Model

Taxonomy

Traits



Panthera tigris

Score: 0.33

Panthera pardus

Score: 0.78 Error: 0.22

Blank

Score: 0.07



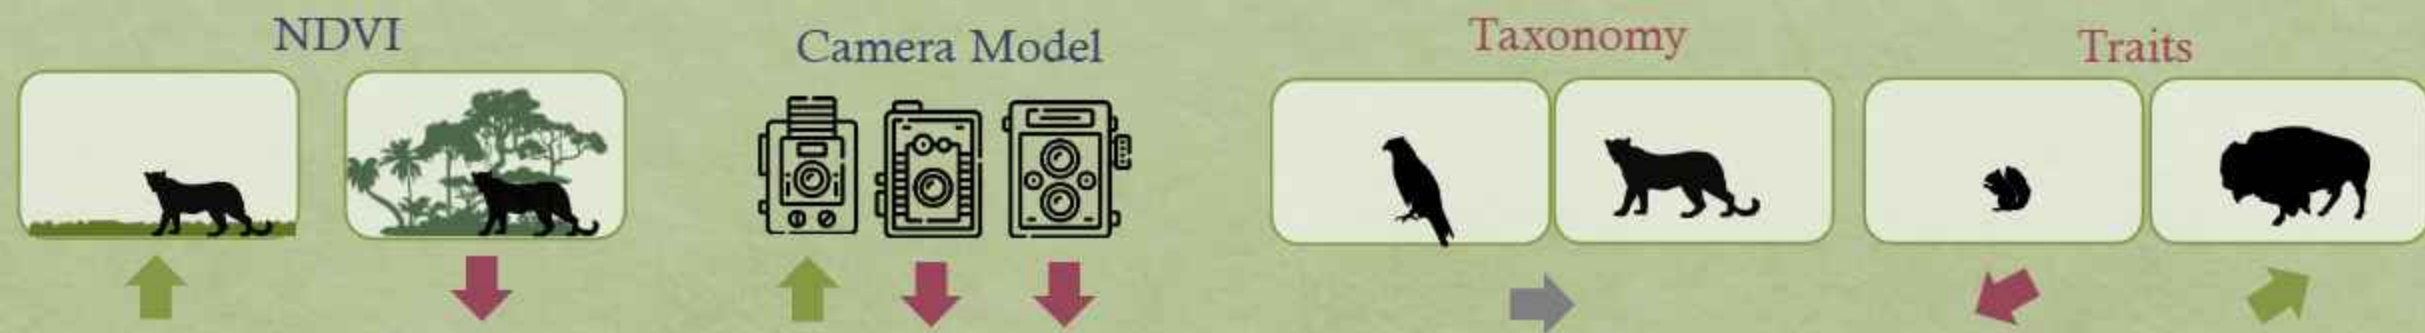
Time of day



Bounding Box Area



Exposure Time | ISO speed



$$\text{logit}(P(\text{Correct}_i)) = \beta_0 + \beta_1 \cdot \text{CameraModel}_i + \beta_2 \cdot \text{NDVI}_i + \beta_3 \cdot \text{Precipitation}_i + \beta_4 \cdot \text{Elevation}_i + \beta_5 \cdot \text{Nocturnality}_i + \beta_6 \cdot \text{ExposureTime}_i + u_{\text{site}[i]} + u_{\text{family}[i]}$$



Use cases

Site-level model



- ◆ Predict expected model performance at new site locations
- ◆ Use site level prediction probability in camera trap modeling processes (e.g Random Encounter Modelling)

Species-level model



- ◆ See relation between data coverage and performance.
- ◆ AI as an extinction filter

Observation-level model



- ◆ Bias mitigation methods
 - ◆ Pre-processing
 - ◆ In-processing
 - ◆ Post-processing
- ◆ Encode ecological knowledge in the model

Thank you

Get In Touch!

vd278@exeter.ac.uk



AI-Based Visual Monitoring of Small, Fast-Moving Nocturnal Fauna:

A Case Study in Bat Roost Detection

William Rees

PhD Researcher

Natural England Licenced Bat Worker



Supervisors:

Dr Benno Simmons (Lead)

Dr Orly Razgour

Dr Iain Soutar

Dr Alona Armstrong (Lancaster University)



Environmental
Intelligence CDT



University
of Exeter



UK Research
and Innovation

Why Monitor Bats

The UK is home to 18 bat species.

Globally bats represent 20% of all mammal species.

Bats provide vital ecosystem services, including:

- Pest suppression
- Pollination
- Bioindicators of broader ecosystem health

Locating roosts helps:

- Prevent human-bat conflict
- Support conservation and population recovery
- Monitor broader ecosystem health



Dusk Bat Surveys

BCT 4th Edition Bat Surveys for Professional Ecologists:

- “Surveys should usually be carried out using Night Vision Aids (NVAs)... as bats often emerge after it is too dark for surveyors to observe them.”



The Analysis Problem...



Typical recording length: ~2.5 hours per survey.

Manual review is time-intensive and unsustainable at scale (large-area or long-term monitoring).

The Invisible Gorilla Experiment (Simons & Chabris, 1999)

- ~50% of participants failed to notice a man dressed as a gorilla when asked to count basketball passes.

Marcot, B.G. et al. (2019) 'Efficacy of automated detection of motion in wildlife monitoring videos

- Motion detection identified 38% more events of target species (woodpeckers and potential predators), than were found by manual student reviews.

Existing Solutions

Audio triggers (BatRack, 2021):

- Bats often silent on exit → many missed events.

Motion detection (Motion Meerkat '15/'18; ThruTracker '21):

- Limited annotation integration
- Hard-to-tune motion thresholds → missed bats
- Many false positives to review

AI animal detectors (MegaDetector '20; SpeciesNet '24; BioCLIP-2 '25):

- Trained on PIR camera-trap stills; fast nocturnal species underrepresented
- Reference images skewed to bats in hand/roost, not in flight
- Small-object + motion blur → low recall/precision for bats

Research Objectives

1. Develop an annotation tool for reviewing bat emergence and re-entry events in long-form videos.
2. Develop an AI Classifier to filter out bycatch events (moths, fluttering vegetation etc) and reduce analysis workloads while achieving high recall of rare bat events.



Dataset Overview



Data

Scope: Dusk emergence surveys (Aug 2023–Aug 2024) from Professional Ecologists

Coverage: 21 sites, 62 scenes

Species: Wide species mix (small → large bats; slow → fast flyers)

Ground Truth Annotations: 130 emergences, 16 re-entries, 156 bat passes/internal flights

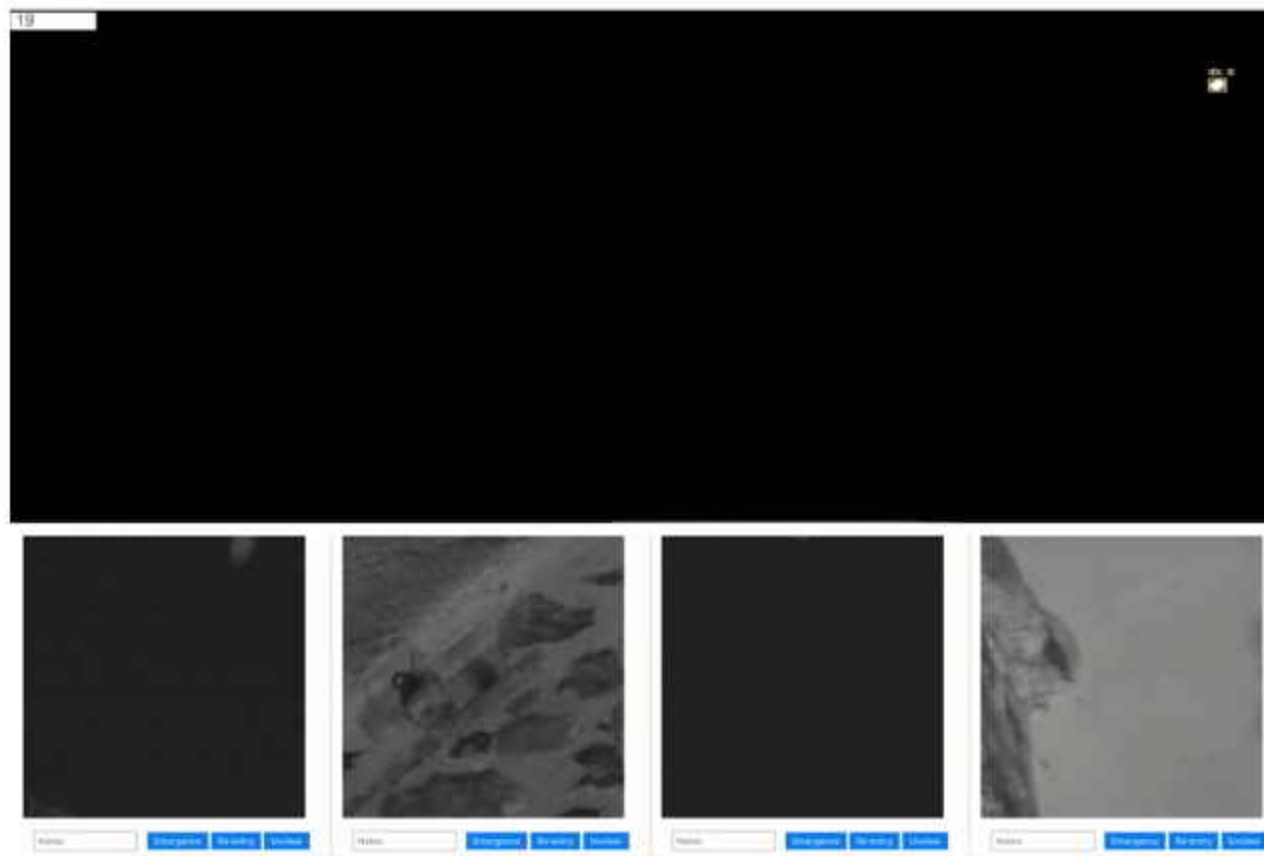
Annotation

80/20 split grouped by location to avoid leakage.

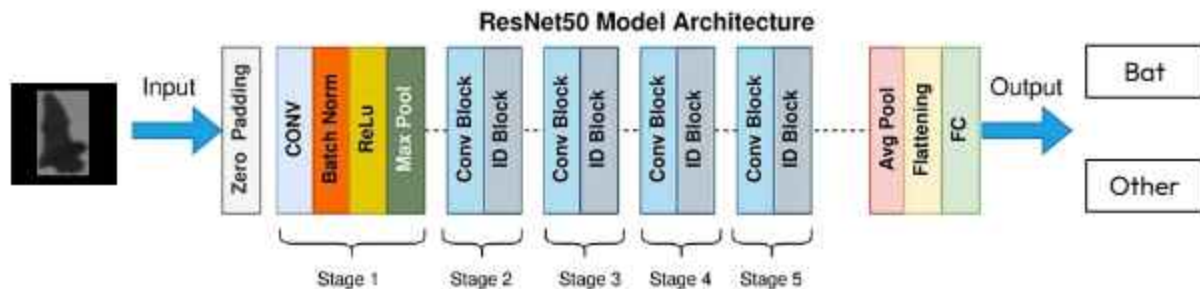
Training data: 1-min buffers around GT “bat” events, annotated Winter 2024/25

Test set: 41.5 h of full-length videos fully annotated.

Tracking and Annotation Interface



Small-Image Classifier



ResNet50 CNN adapted for monochrome 64×64 pixel inputs.

Early Layers: Identify basic features – edges, textures, small patterns.

- E.G: Detecting the ears on a long-eared bat.

Deeper Layers: Recognise larger shapes & overall structure of a bat.

- E.G: Distinguishing a bat from a moth based on the broad silhouette shape.

Inputs resized and zero-padded. Ensures the broad shape of objects is preserved.

Image Classifier Performance

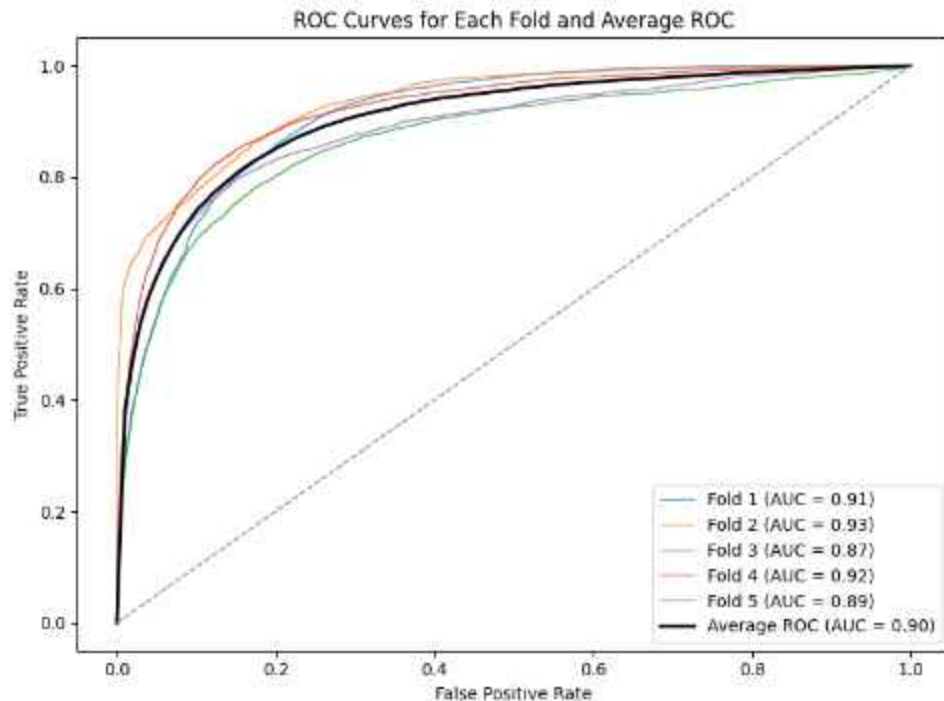
Model parameters tuned using a 5-fold (location-grouped) split of training data.

ROC Curve (Receiver Operating Characteristic)

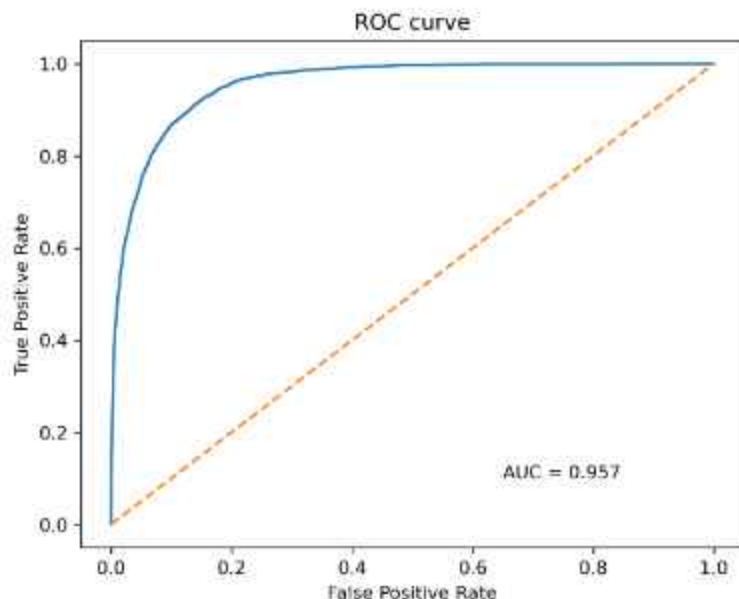
- Measures the ability of models to distinguish between bats and other objects across different decision thresholds.

ResNet50 outperformed alternative architectures.

Average AUC = 0.90
Exceptional Discrimination

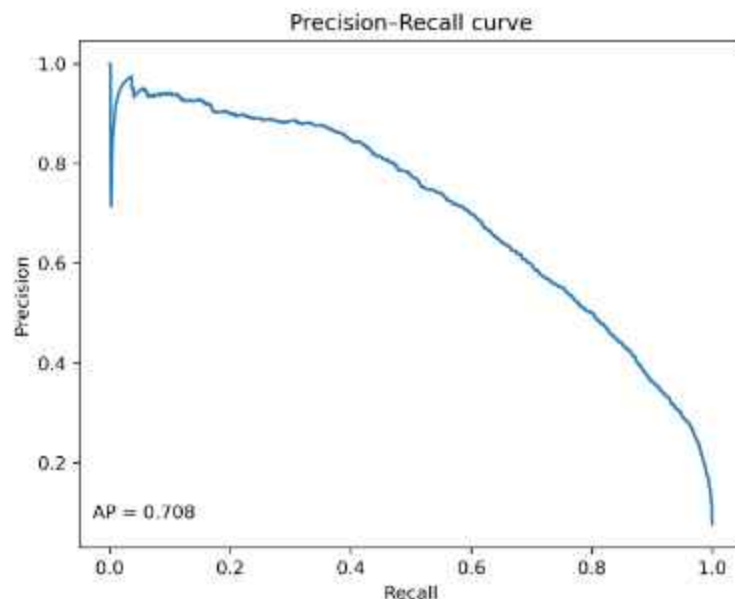


Event Level Predictions



Average AUC = 0.957

Event Discrimination > Frame Discrimination



Average Precision = 0.708

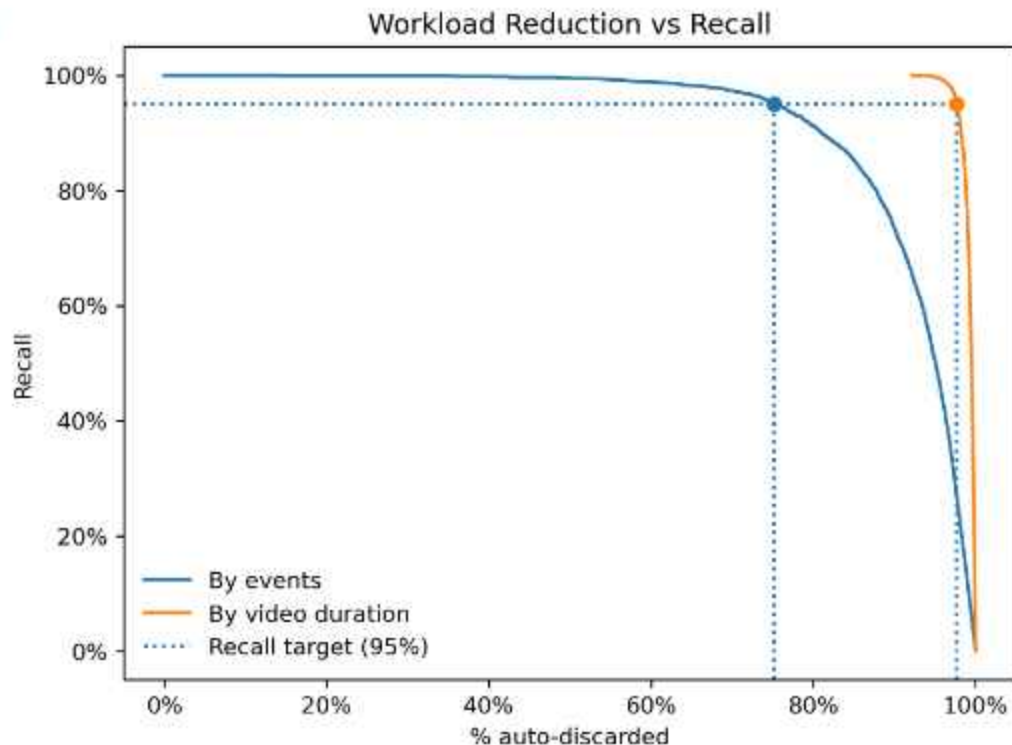
Positive Prevalence : 0.076

~ 9.3x improvement over Random event Classifier.

Workload Savings

Conservative Scenario (confidence threshold of 0.3)

- >100% recall of key events (bat emergencies).
- 95% recall of all positive events.
- 75.2% reduction in motion events requiring review.
- 97% reduction in video needing human review.

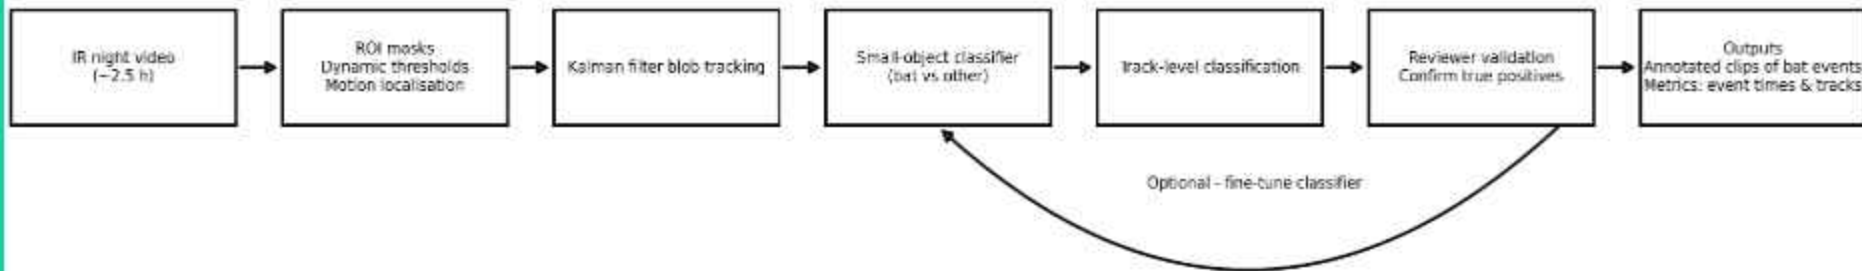


Thanks for Listening!

wr291@exeter.ac.uk



University
of Exeter





Geological mapping in the age of artificial intelligence

Pathways to progress in the 21st century

Charlie Kirkwood, Research Fellow
Institute for Data Science and AI
University of Exeter

For CEI launch 11/09/2025





A big thank you to everyone who came to or otherwise supported the AI for Geological Modelling and Mapping conference (AI-GMM) in Exeter in May 2024

Lots of brilliant posters and talks. You can watch the talks on youtube here:

www.youtube.com/@AI-GMM



University
of Exeter

Institute for Data
Science and Artificial
Intelligence



Engineering and
Physical Sciences
Research Council

A Geological Society of London special publication on this topic is in progress

We're still considering expressions of interest (although space is limited)

If you've got work you'd like to publish *where geological modelling and mapping intersect with statistics, machine learning, and artificial intelligence* please get in touch.

We want diversity in people, methods, & perspectives

You don't have to be pro-AI

Under consideration

Geological Society Special Publication

Geological Modelling and Mapping in the Age of AI

Edited by Charlie Kirkwood, Kristine Asch, Mark Lindsay, Florian Wellmann, Michael Hillier, Guillaume Caumon

Rapid developments in AI and data science are unlocking new opportunities for how we go about modelling and mapping the Earth. This timely Special Publication will assemble a collection of papers where geological modelling and mapping intersect with statistics, machine learning, and artificial intelligence, with the aim of informing how we can maximise the benefits of these technologies for the future of geological modelling and mapping.

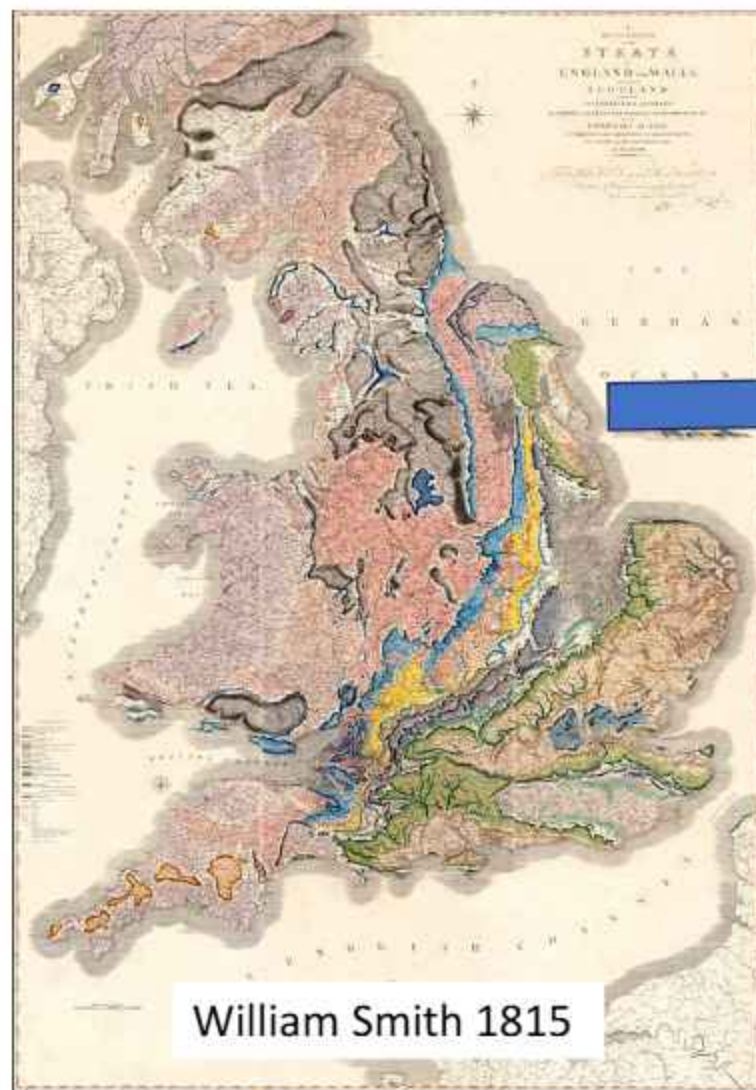
Papers need not be written in support of AI techniques; well-reasoned perspectives are encouraged from all angles - this Special Publication is open to diverse views from across the geological community.

Special topics for this SP, but not limited to, are:

- ✓ Applications of statistics, machine learning, and AI for geological modelling and mapping, with preference given to applications that can facilitate our transition to net zero emissions (the age of AI is also the age of climate crisis);
- ✓ Verification / validation of geological models and maps, including benchmarking between different techniques for modelling and mapping, assessing their closeness to reality;
- ✓ Perspectives on the role of AI and associated technologies in the future of geological modelling and mapping, and the implications for future geoscientific skills and culture;
- ✓ Educational pieces and low-code tutorials to help widen accessibility of these topics.

A short summary of this area,
from my perspective

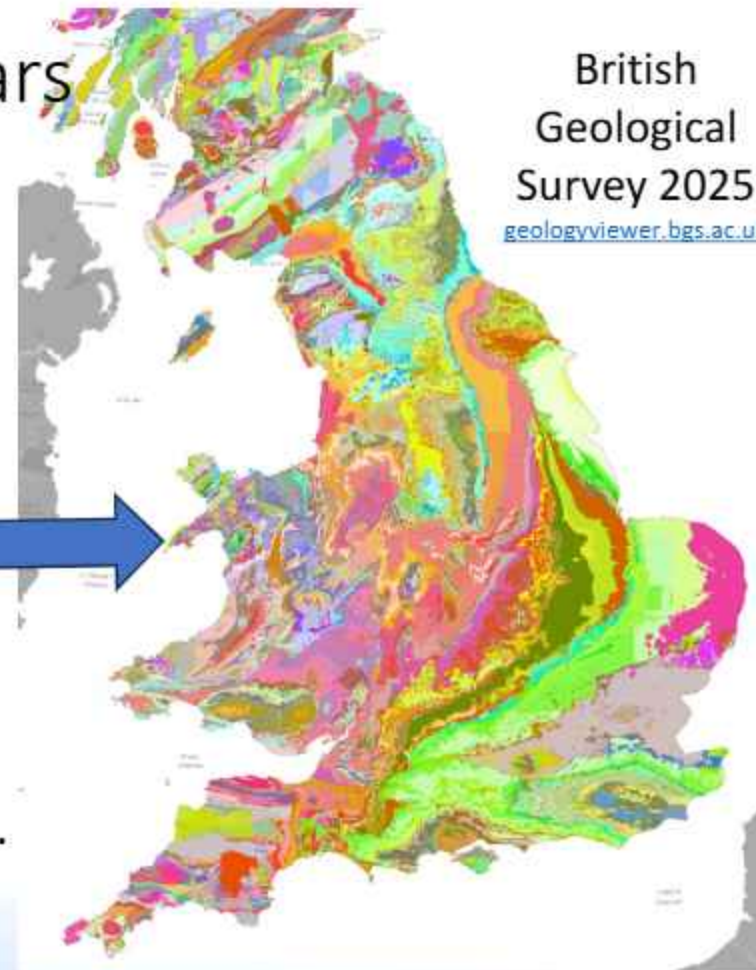
Geological mapping for the last 200 (and 10) years



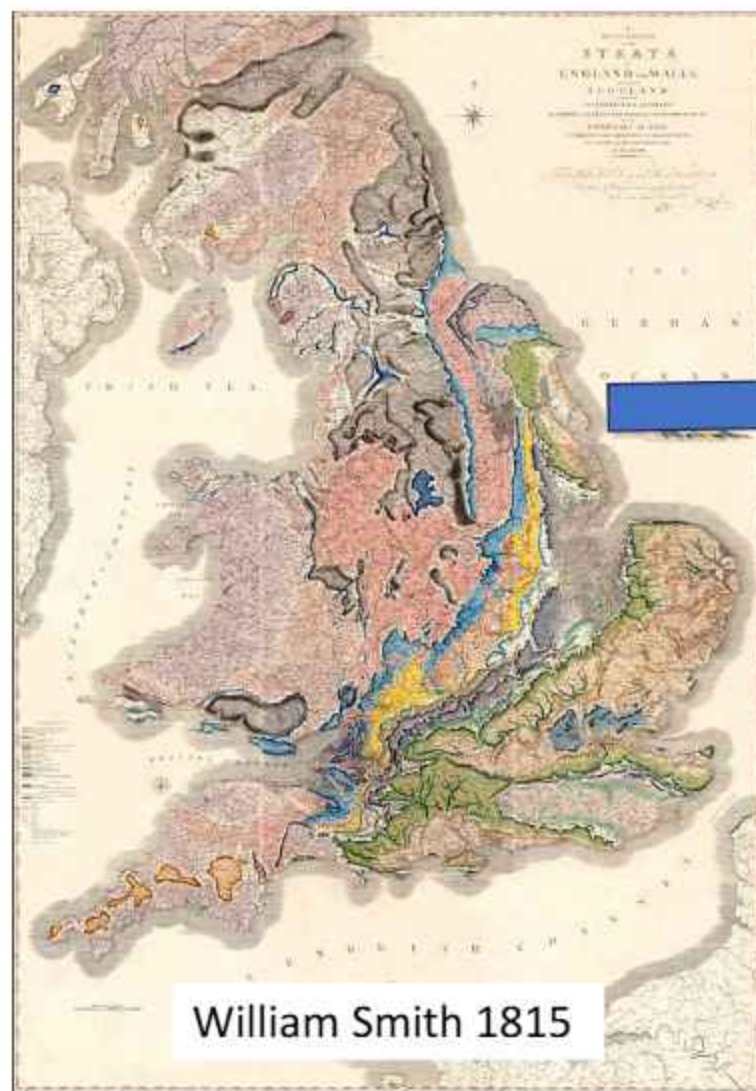
- For two centuries we've (mostly) been constructing geological maps by drawing lines to delineate boundaries between different units, or classes, of rock.

More lines

- What do these lines represent?
- Each line represents a change in geological conditions or properties. (e.g. in age, composition, texture).



Geological mapping for the last 200 (and 10) years



- For two centuries we've (mostly) been constructing geological maps by drawing lines to delineate boundaries between different units, or classes, of rock.

More lines

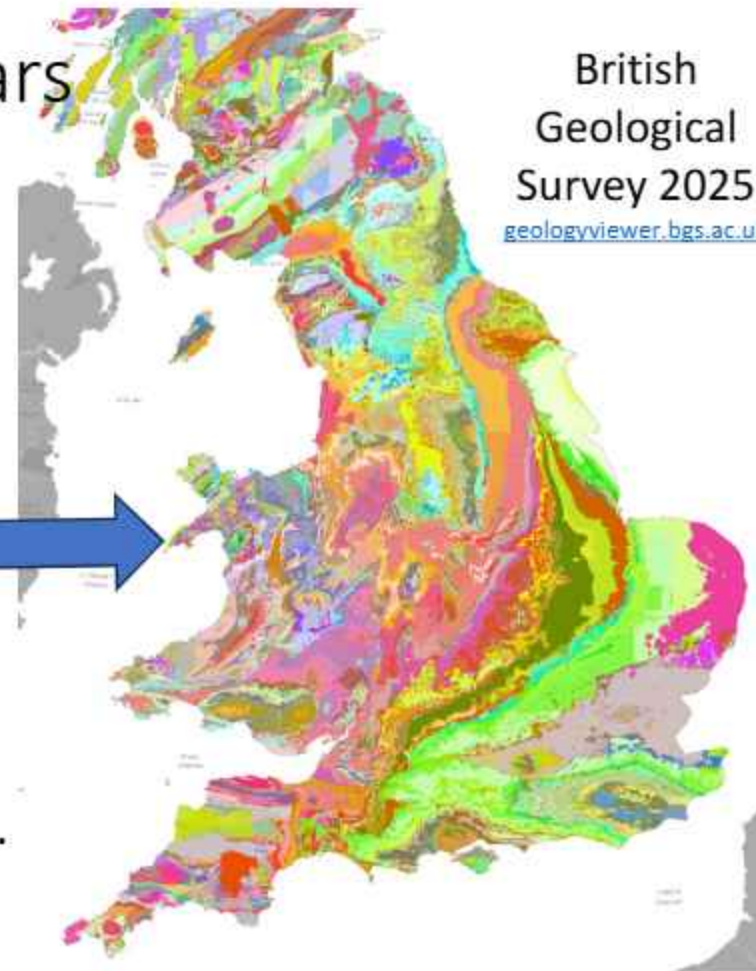
- What do these lines represent?
- Each line represents a change in geological conditions or properties. (e.g. in age, composition, texture).

But:

- Do changes in geology occur only where lines are drawn?
- Do all lines represent the same magnitude of change?
- For all properties at once?

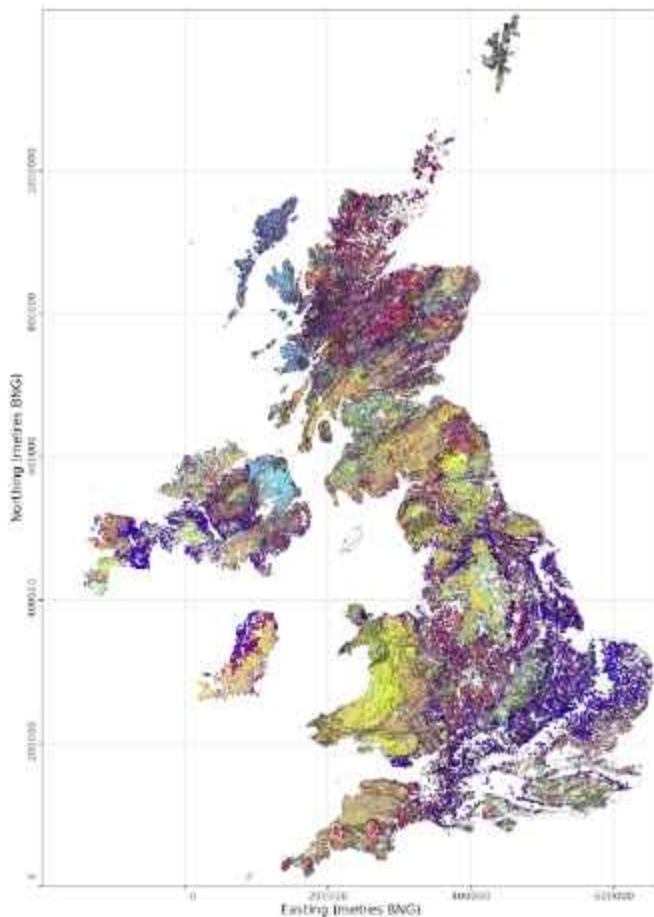
If not, which properties are changing, and by how much?

And from what, to what? **How informative can such maps be?**



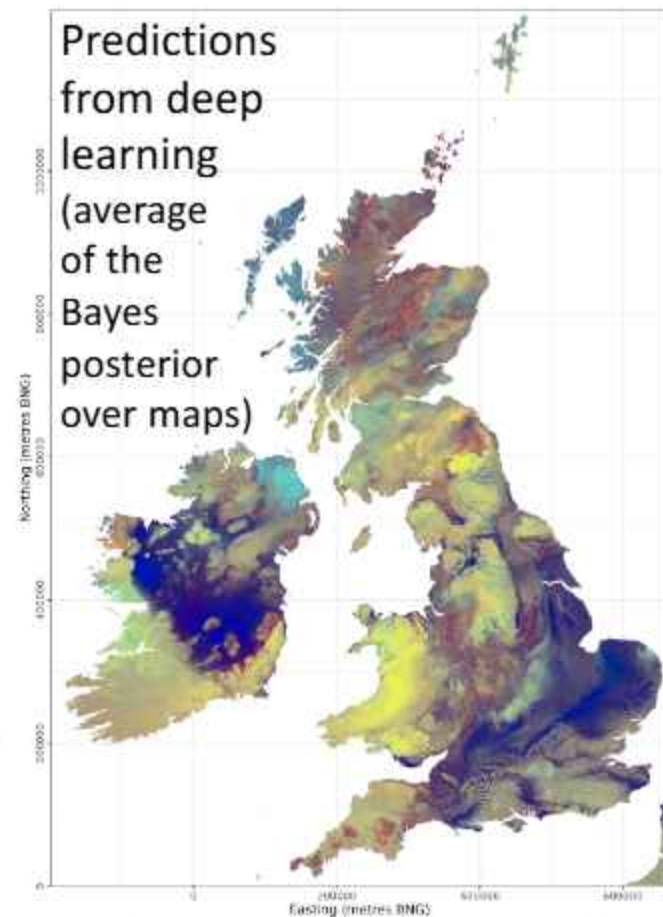
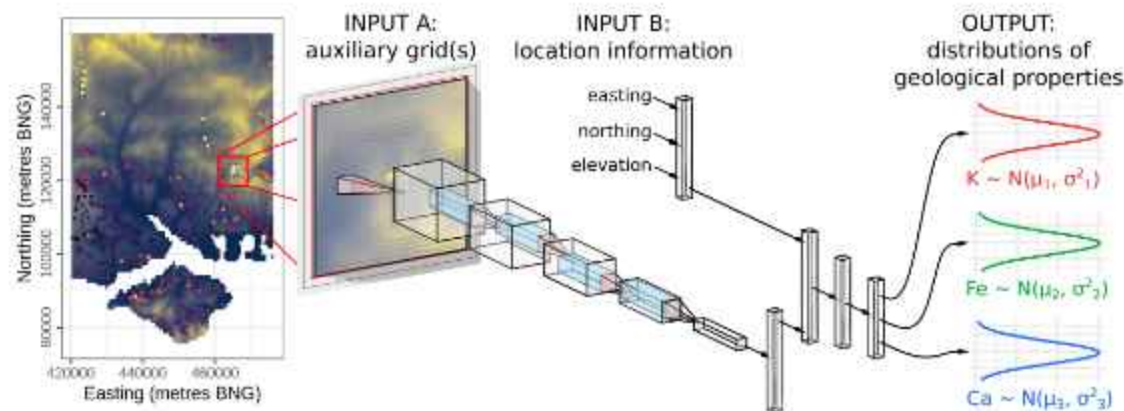
British
Geological
Survey 2025
geologyviewer.bgs.ac.uk

Alternative approach: map the geological properties themselves



Input observations of geological properties, in this case geochemical composition

- This is where AI comes in
- The diagram below is the neural network architecture we developed for 'Bayesian deep learning for spatial interpolation in the presence of auxiliary information' in Mathematical Geosciences, 2022
- It makes predictions by combining global location information with local contextual information (learned by computer vision)



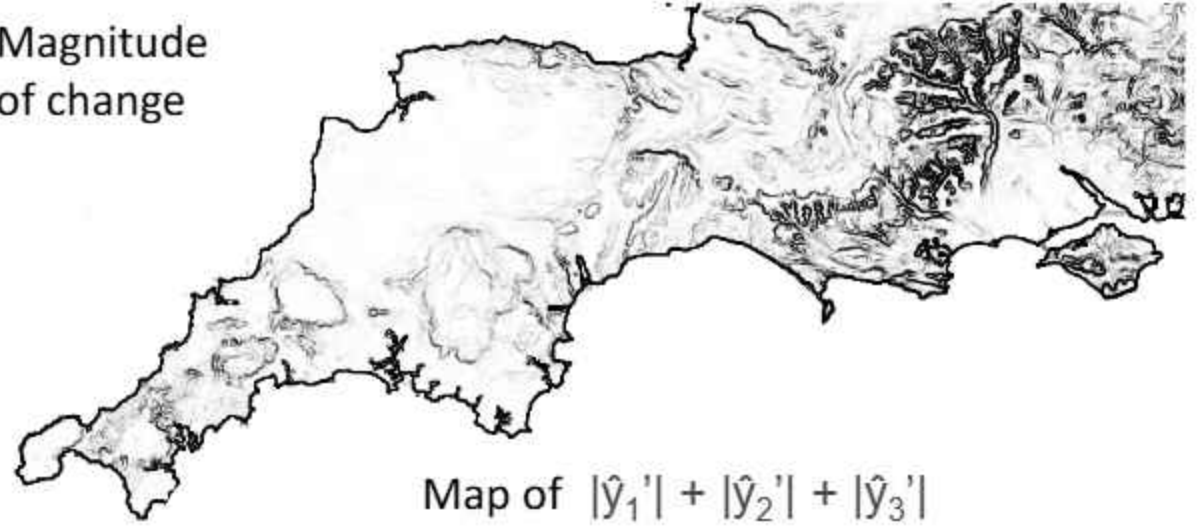
Output predictions of geological properties. Here red = potassium, green = iron, blue = calcium

Alternative approach: map the geological properties themselves, and the boundaries become self-evident

Geological
properties



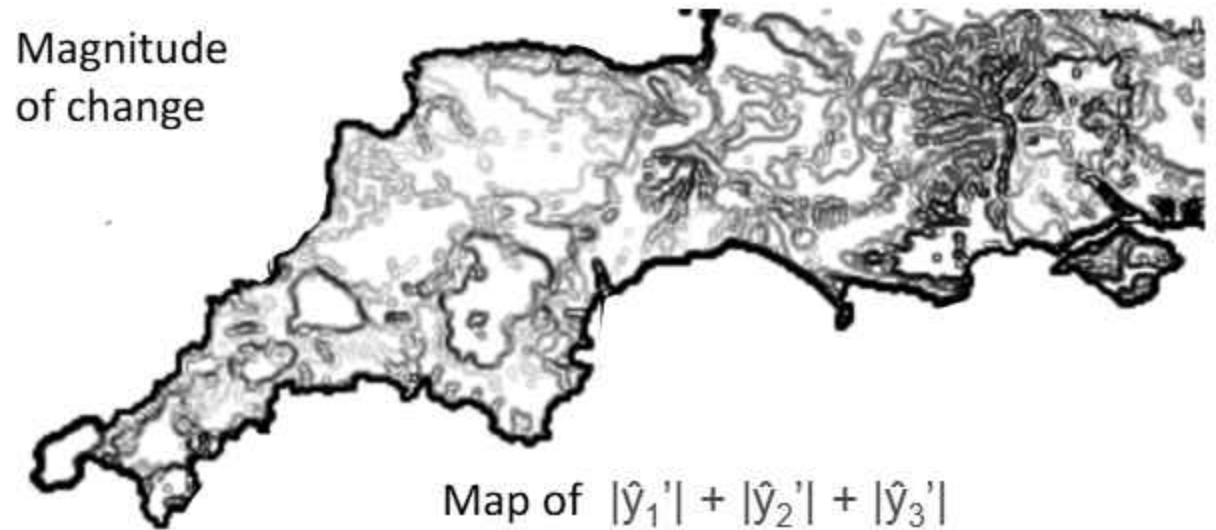
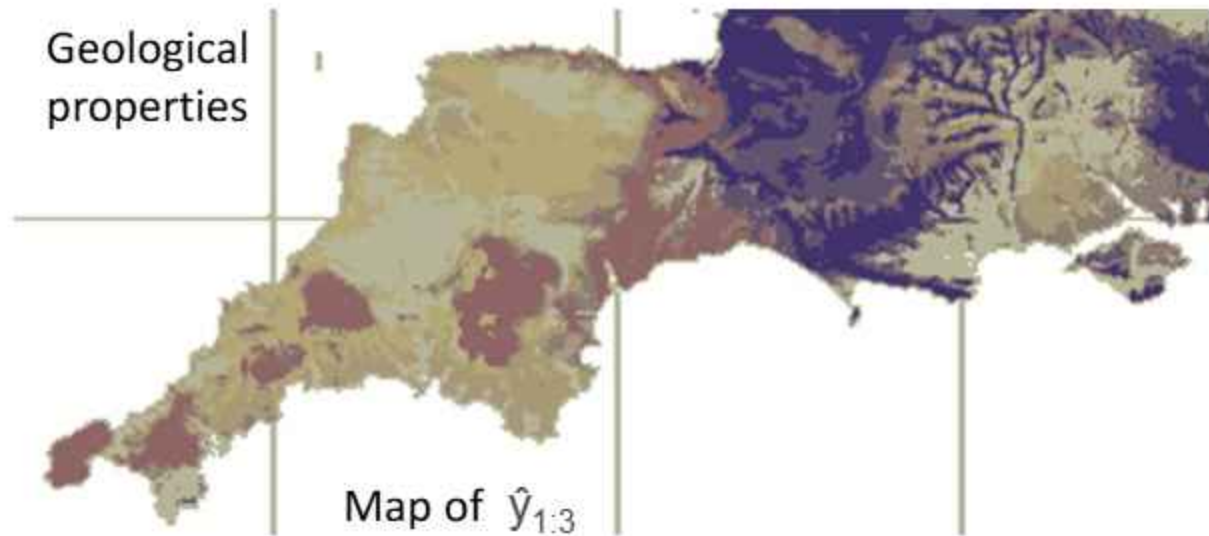
Magnitude
of change



Interestingly, mapping the magnitude of change in geological properties does not often naturally produce 'closed polygons'.

This makes logical sense when you consider that not all geological changes are of equal magnitude, and some are quite gradational.

To achieve closed polygons like our traditional geological maps requires **quantising** the geological properties (thresholding their values)

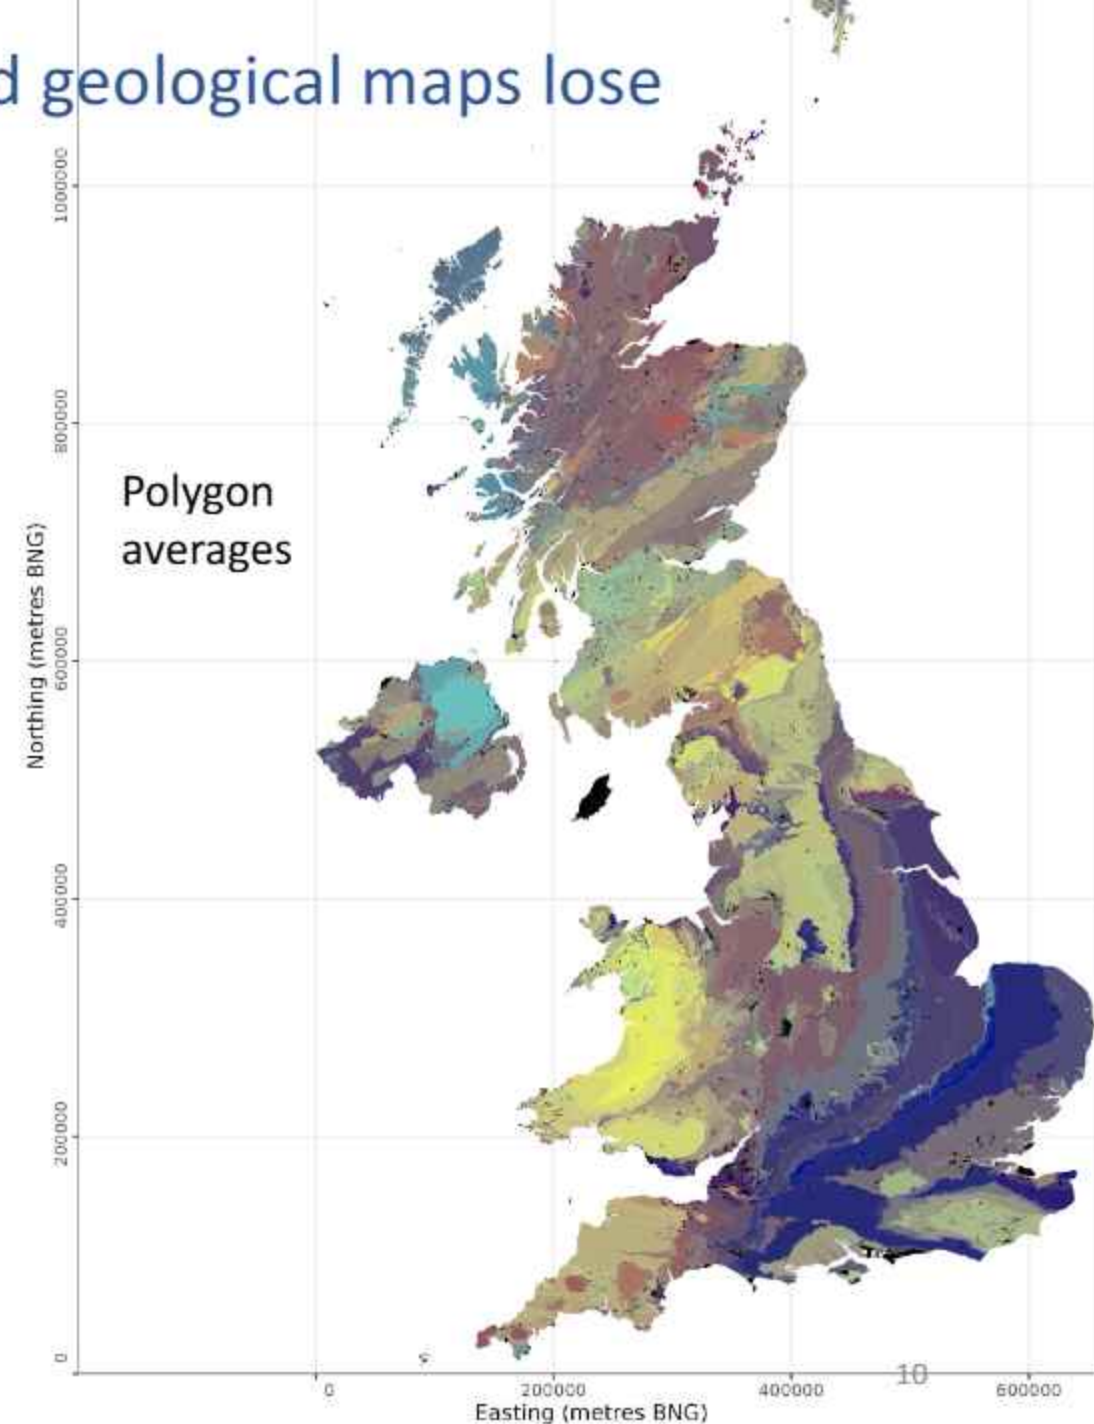
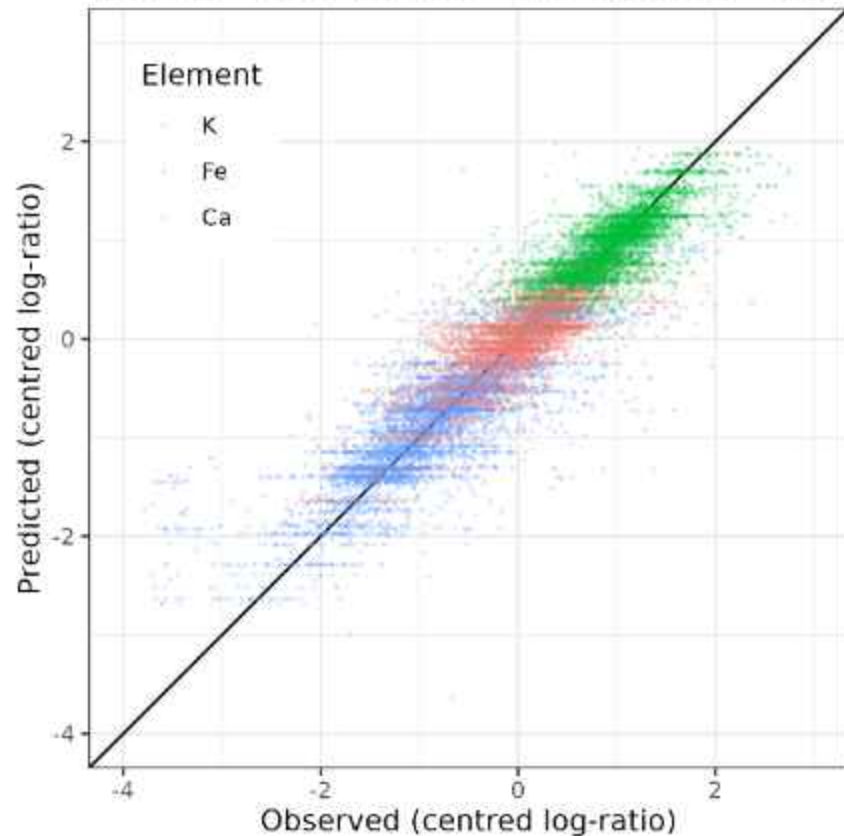


But quantising the map to achieve closed polygons loses information to quantisation error.

How much information do traditional classified geological maps lose to quantisation error?

Predicting chemical composition of K, Fe, Ca using polygon averages results in R^2 values of 0.64, 0.57, 0.66

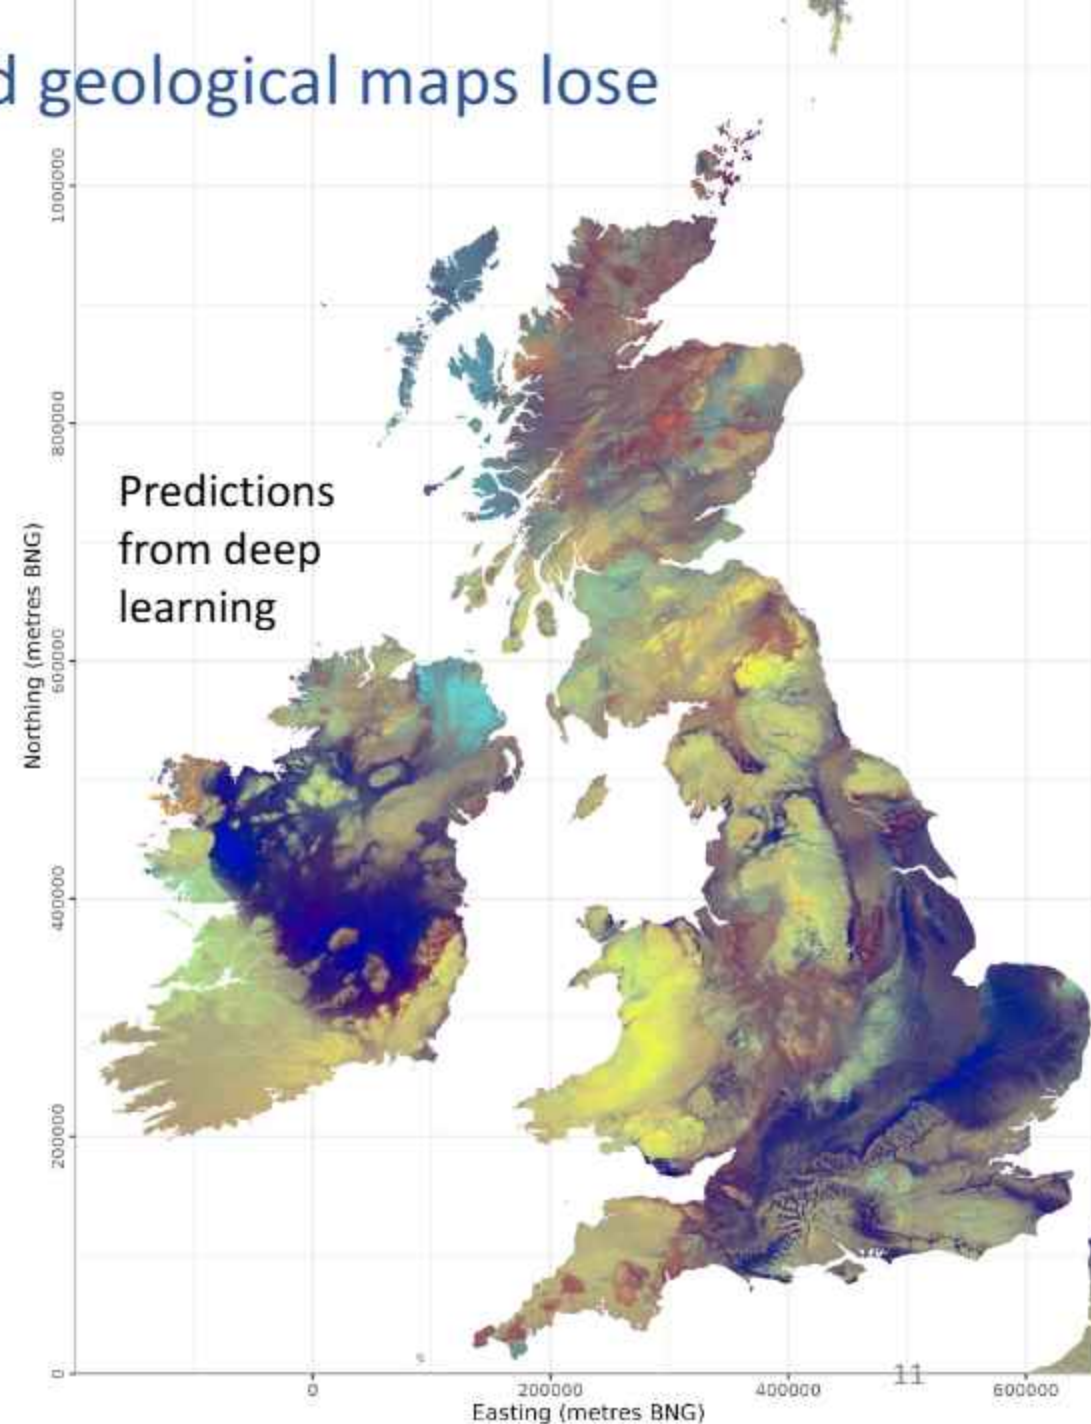
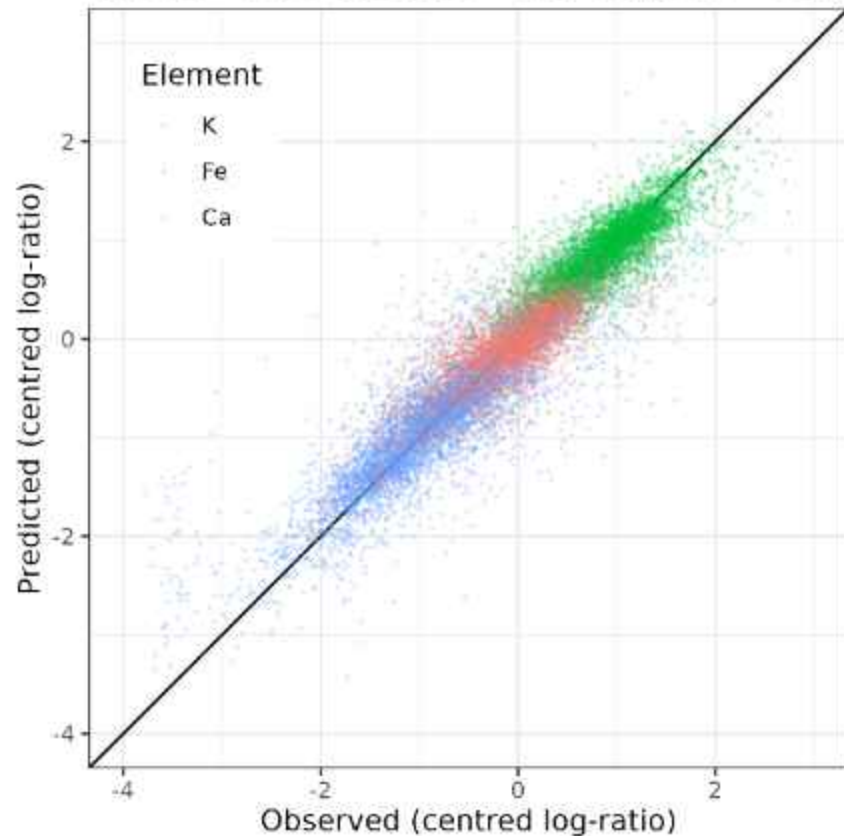
R^2 for K = 0.64 R^2 for Fe = 0.57 R^2 for Ca = 0.66
RMSE K = 0.32 RMSE Fe = 0.3 RMSE Ca = 0.47



How much information do traditional classified geological maps lose to quantisation error?

Predicting chemical composition of K, Fe, Ca using Bayesian deep learning results in R^2 values of 0.72, 0.66, 0.75

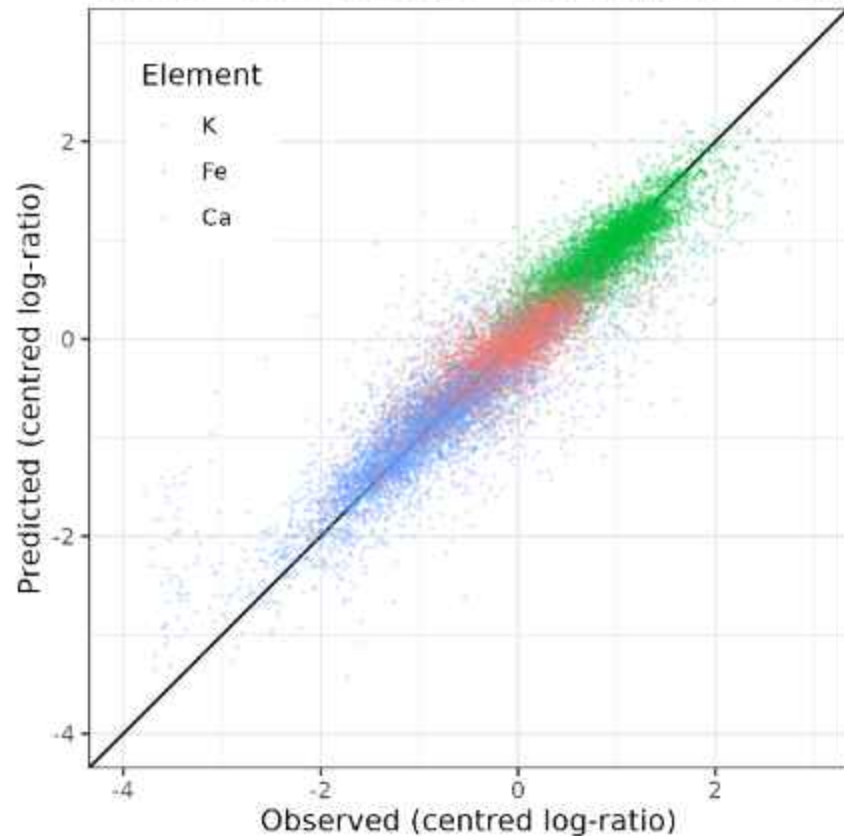
R^2 for K = 0.72 R^2 for Fe = 0.66 R^2 for Ca = 0.75
RMSE K = 0.29 RMSE Fe = 0.27 RMSE Ca = 0.41



How much information do traditional classified geological maps lose to quantisation error?

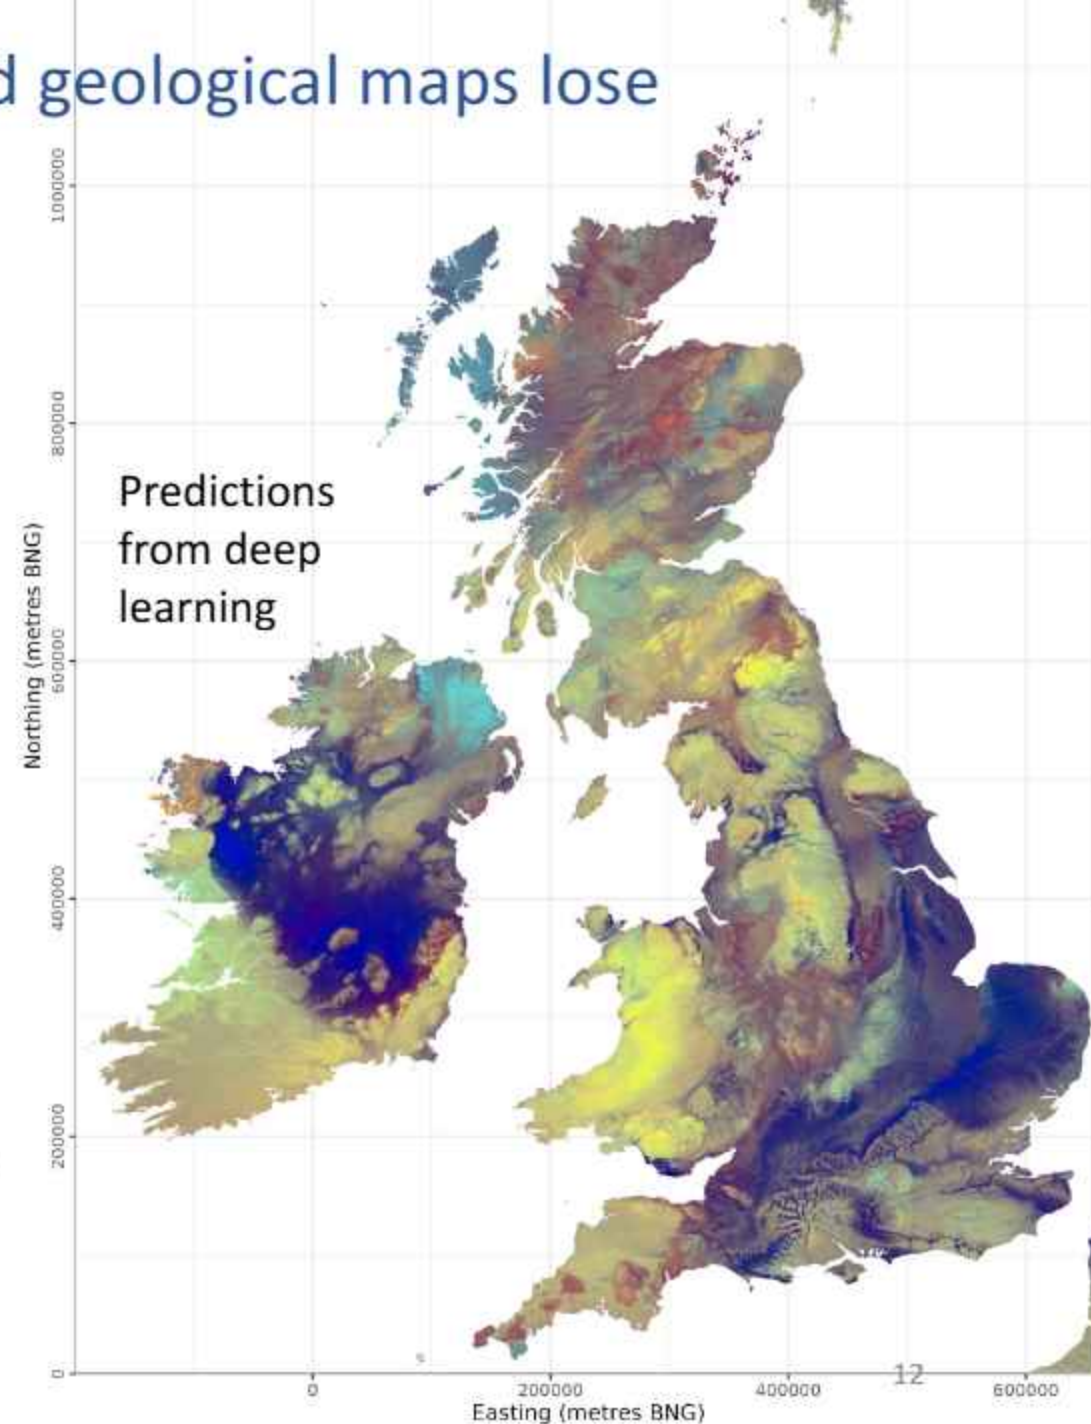
Predicting chemical composition of K, Fe, Ca using Bayesian deep learning results in R^2 values of 0.72, 0.66, 0.75

R^2 for K = 0.72 R^2 for Fe = 0.66 R^2 for Ca = 0.75
RMSE K = 0.29 RMSE Fe = 0.27 RMSE Ca = 0.41



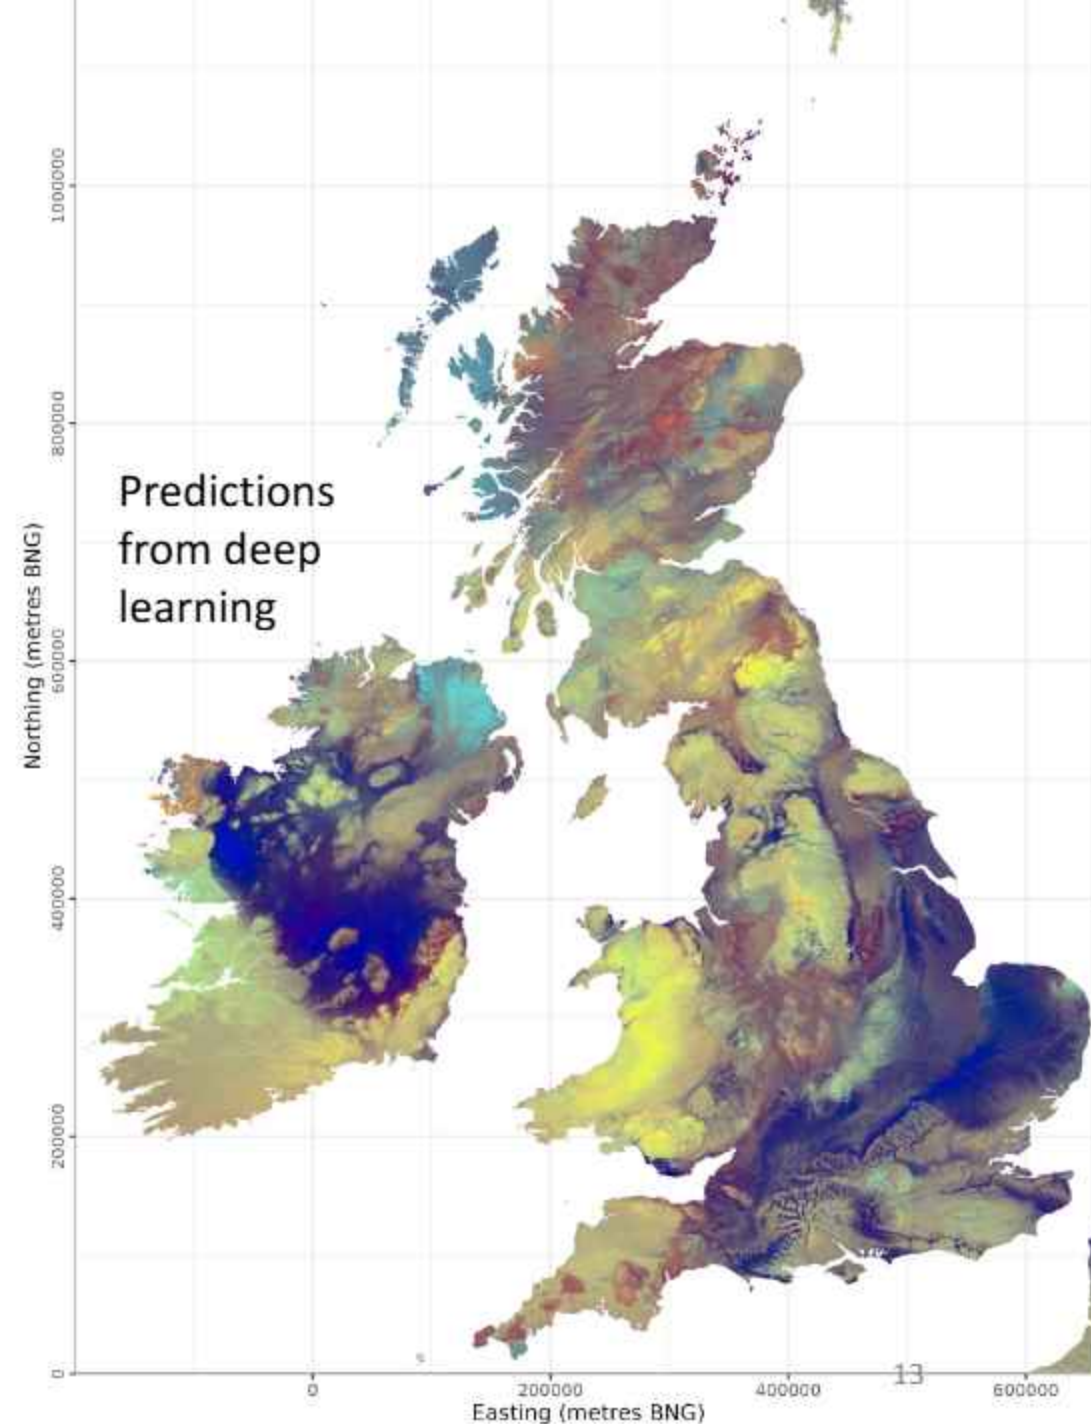
That's +8%, +9%, +9%
variance explained
or, 1/5th of the way
closer to perfect
information

And these are major
rock-forming
elements that should
be well represented
by the traditional map



Summary

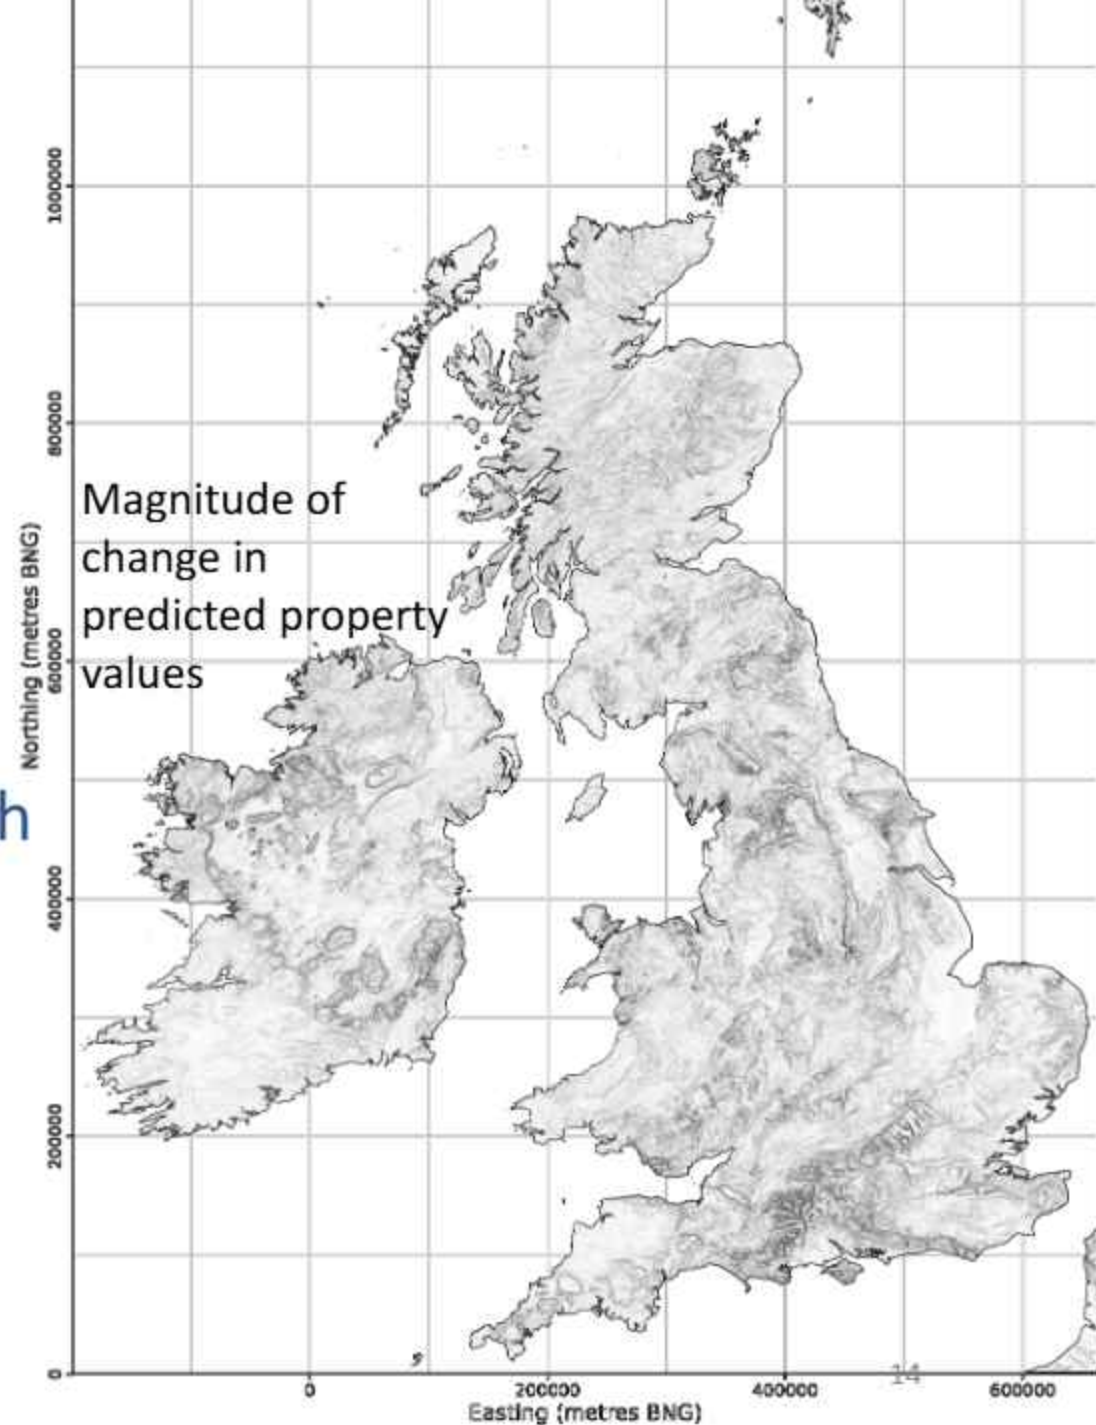
- A revolution in geological mapping, made feasible by developments in AI.
- Comparably significant to the Numerical Weather Prediction revolution in weather forecasting (which also started out using hand-drawn maps).
 - > probably **more** significant than the transition from Numerical Weather Prediction to AI weather prediction.
- Better geological maps – with high fidelity and quantified uncertainties – serve humanity by enabling us to interact with the Earth more precisely, minimising mistakes and waste. Quantified uncertainties allow us to better judge risks.
- Applies to energy, minerals, construction, water, agriculture, hazard mitigation, carbon storage, radioactive waste disposal, and more...
- The University of Exeter, and the Centre for Environmental Intelligence, brings together all the necessary skills and culture to steer geological mapping into the age of AI.



Thanks for listening

If you find it interesting, feel free to get in touch

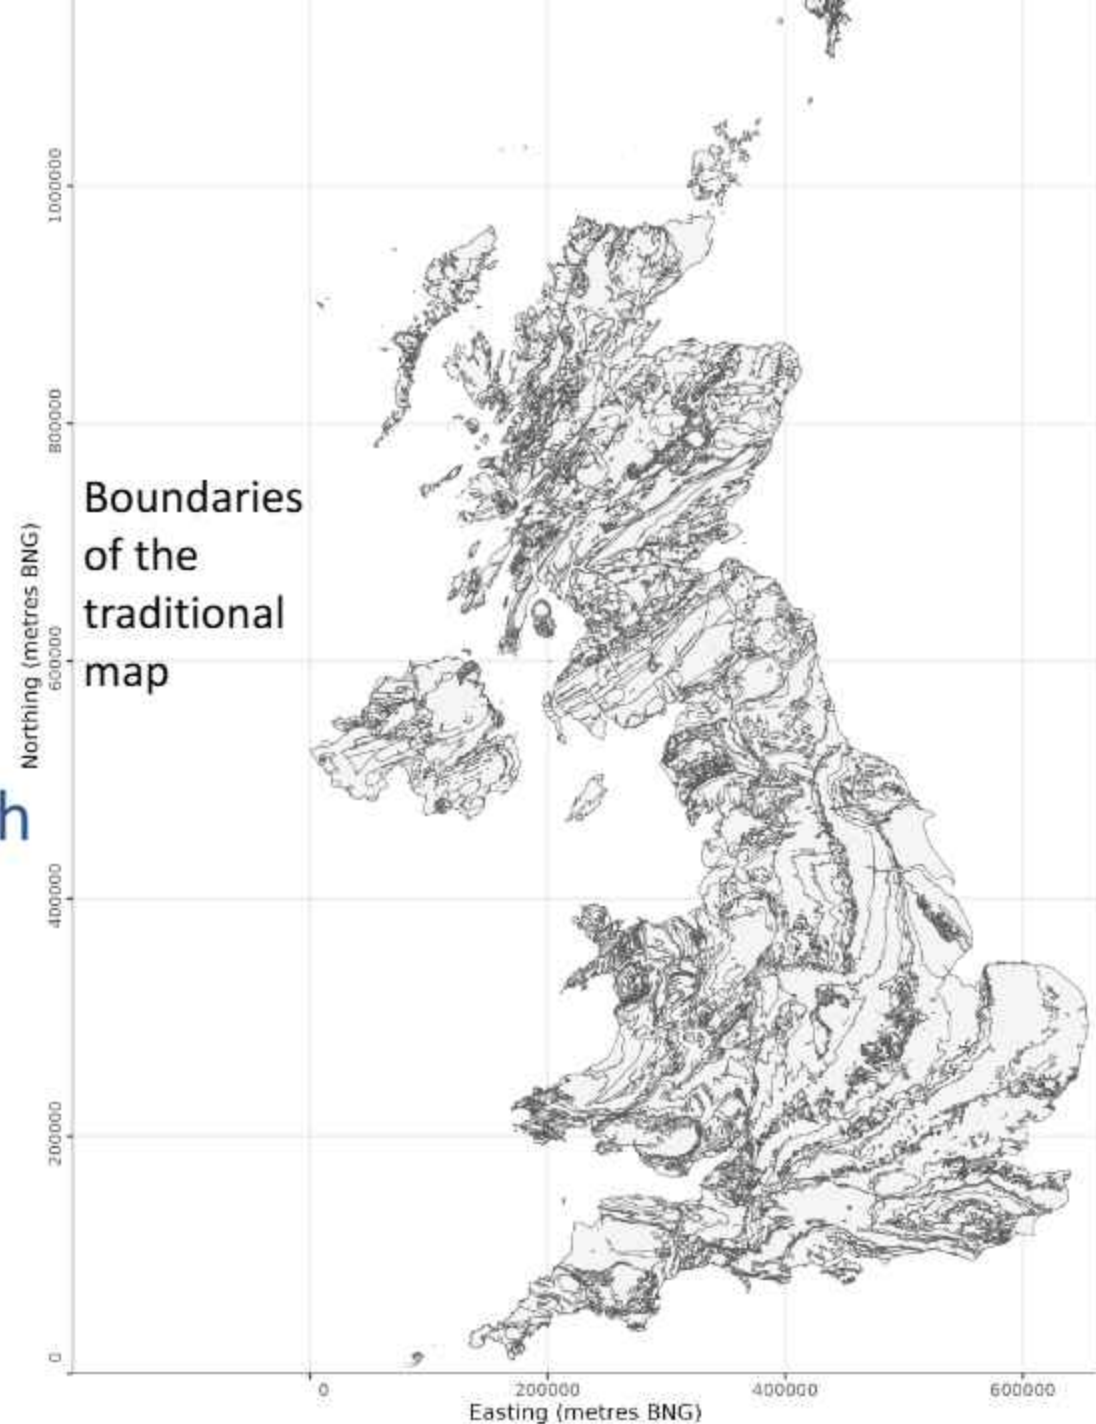
c.kirkwood@Exeter.ac.uk



Thanks for listening

If you find it interesting, feel free to get in touch

c.kirkwood@Exeter.ac.uk



Data-driven subgrid-scale parameterization in atmosphere and ocean models: a pattern-based approach

Frank Kwasniok

Department of Mathematics and Statistics
University of Exeter, UK

Environmental Intelligence Conference 2025
Exeter, UK
11 September 2025

Introduction

Modeling approaches in weather and climate science:

- Physics-based or forward or direct modelling
- Data-driven or inverse modelling
- Hybrid modelling

Issues:

- Deterministic versus stochastic dynamics
- Equations of motion versus machine learning-style approaches
- Sparsity
- Continuous versus discrete dynamics
- Dimension reduction, optimal coordinates

Stochastic subgrid-scale parametrization

Climate state vector:

$$\mathbf{u} = (\mathbf{x}, \mathbf{y})$$

True tendency of resolved variables:

$$\dot{\mathbf{x}} = \mathbf{R}(\mathbf{x}) + \mathbf{S}(\mathbf{x}, \mathbf{y})$$

Parametrisation of unresolved tendency:

$$\mathbf{S}(\mathbf{x}, \mathbf{y}) \sim \mathbf{f}(\mathbf{x}) + \boldsymbol{\eta}(\mathbf{x})$$

Canonical choice:

$$\mathbf{f}(\mathbf{x}) = \langle \mathbf{S}(\mathbf{x}, \mathbf{y}) | \mathbf{x} \rangle$$

Discrete version:

$$\mathbf{x}^{(t+\delta t)} = \mathcal{R}(\mathbf{x}^{(t)}) + \delta \mathbf{x}^{(t+\delta t)}$$

Non-Markovian stochastic reduced model

$$\mathbf{x}^{(t+\delta t)} = \underbrace{\mathcal{R}(\mathbf{x}^{(t)}) + \langle \mathbf{S} \rangle + \sum_{j=1}^J \mathbf{p}_j \mathbf{q}_j^T \mathbf{Z}^{(t)}}_{\text{deterministic}} + \underbrace{\sum_{j=1}^J s_j \mathbf{p}_j + \epsilon}_{\text{stochastic}}$$

$$\mathbf{Z}^{(t)} = (f_1, \dots, f_L)^T, \quad f_l = f_l(\mathbf{x}^{(t-r\delta t)}, \dots, \mathbf{x}^{(t)})$$

$$s_j^{(t)} = \sum_{m=1}^M \alpha_m s_j^{(t-m\delta t)} + \xi_j^{(t)}$$

Nonlinear principal prediction patterns

Minimise error function:

$$I = \left\langle \left\| \mathbf{x}_{n+1} - \mathcal{R}(\mathbf{x}_n) - \langle \mathbf{S} \rangle - \sum_{j=1}^J \mathbf{p}_j \mathbf{q}_j^T \mathbf{z} \right\|^2 \right\rangle$$

Symmetric eigenvalue problem:

$$\mathbf{G} \mathbf{C}^+ \mathbf{G}^T \mathbf{p}_j = \lambda_j \mathbf{p}_j,$$

$$\mathbf{q}_j = \mathbf{C}^+ \mathbf{G}^T \mathbf{p}_j$$

$$\mathbf{C} = \langle \mathbf{Z} \mathbf{Z}^T \rangle,$$

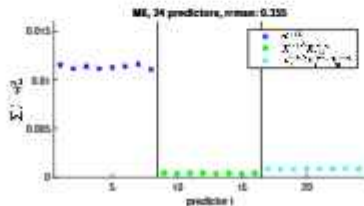
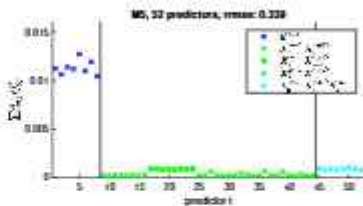
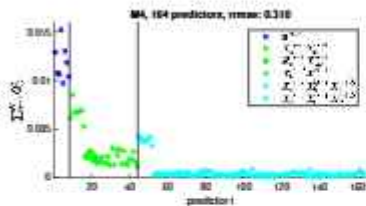
$$\mathbf{G} = \langle [\mathbf{x}_{n+1} - \mathcal{R}(\mathbf{x}_n) - \langle \mathbf{S} \rangle] \mathbf{Z}^T \rangle$$

$$\mathbf{p}_i^T \mathbf{p}_j = \delta_{ij},$$

$$\mathbf{q}_i^T \mathbf{C} \mathbf{q}_j = \lambda_i \delta_{ij}$$

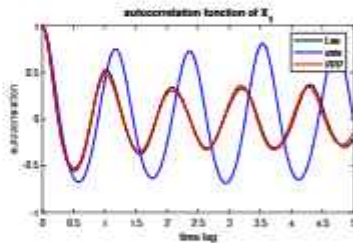
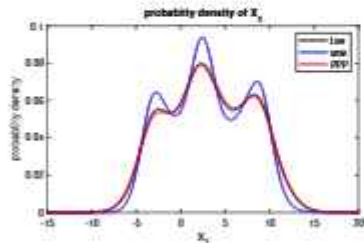
$$I = \left\langle \|\mathbf{x}_{n+1} - \mathcal{R}(\mathbf{x}_n) - \langle \mathbf{S} \rangle\|^2 \right\rangle - \sum_{j=1}^J \lambda_j$$

Model pruning / Predictor selection



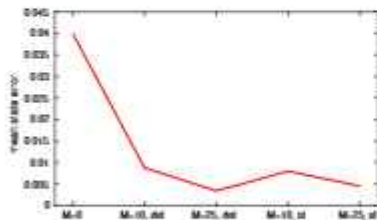
Long-term integration with closure

Probability density and autocorrelation function

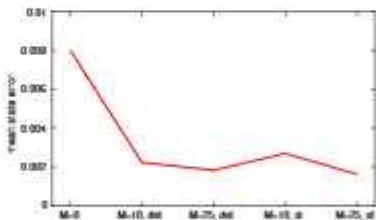


Long-term simulation – Mean state

T15 – T30

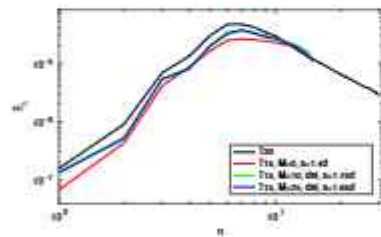


T21 – T30

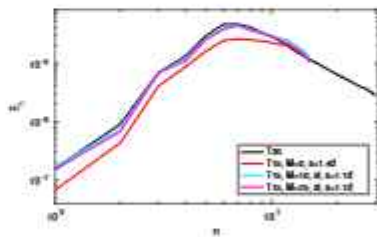


Long-term simulation – Kinetic energy spectrum

deterministic



stochastic



Conclusions and outlook

- Data-driven deterministic and stochastic subgrid-scale parameterization
- Key points:
 - Spatio-temporal approach including dimension reduction
 - Efficient and interpretable subgrid modelling
 - Identification of parameterization structure
 - Predictor selection / Sparsity via data-driven pruning procedure
 - Low computational cost for both training and implementation in reduced model runs
 - Localisation by combining pattern technique with a model-based clustering algorithm
 - Option of adding physical constraints

Measuring the difference between weather and climate models and the real world

Hugo Lambert

University of Exeter

- Weather and climate models produce widely differing predictions because they contain **parametrisations**.
- Measure the differences between parametrized models, process models and observations for a given **process** using **Continuous Structural Parametrization [2]**.
- Here, I analyse the land surface heat budget in GCMs and **FLUXNET** observations.

Comparing models and observations

- Requires rewriting model inputs, X , and outputs, Y , as functions of the same variables.

Comparing models and observations

- Requires rewriting model inputs, X , and outputs, Y , as functions of the same variables.
- Standard data reduction then identifies the most important X and Y using e.g. PCA. But that may not identify the inputs you want...

Comparing models and observations

- Requires rewriting model inputs, X , and outputs, Y , as functions of the same variables.
- Standard data reduction then identifies the most important X and Y using e.g. PCA. But that may not identify the inputs you want...
- Instead,
 - Write $Y = X\beta$, where β are regression coefficients.
 - Do PCA of $X\beta$.
 - Examine $Q = P\gamma$, where Q and P are the weights of the eigenvectors of the covariance matrices of Y and X and γ are the rotated β .

Comparing models and observations

- Requires rewriting model inputs, X , and outputs, Y , as functions of the same variables.
- Standard data reduction then identifies the most important X and Y using e.g. PCA. But that may not identify the inputs you want...
- Instead,
 - Write $Y = X\beta$, where β are regression coefficients.
 - Do PCA of $X\beta$.
 - Examine $Q = P\gamma$, where Q and P are the weights of the eigenvectors of the covariance matrices of Y and X and γ are the rotated β .

Note – this is not how you actually do the calculations [2].

Specific (low dimensional) example

Monthly mean 30° N – S **FLUXNET** observations [3], and model output from AMIP, a CESM2 land surface PPE [4] and a HadCM3 land surface PPE [1].

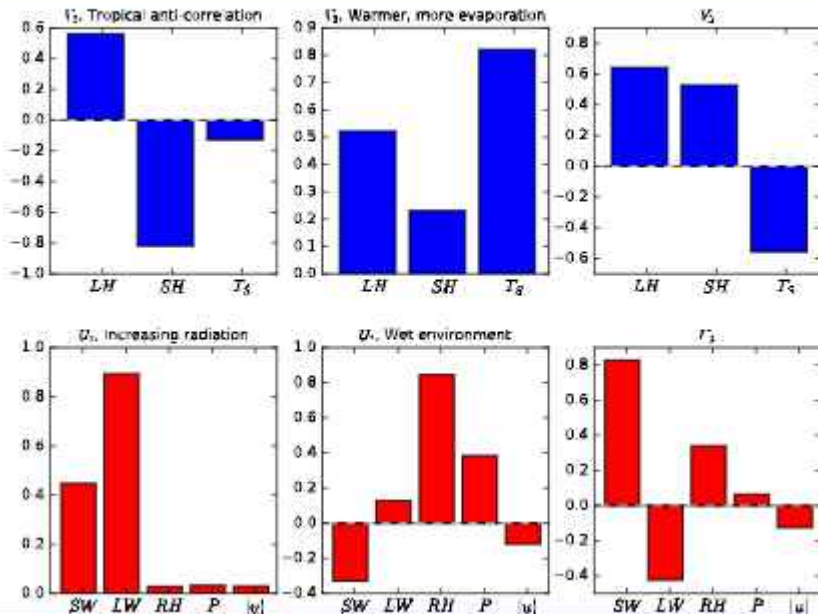
Define Outputs, Y , as upward LH and SH, and surface temperature.

and Inputs, X , as downward shortwave and longwave radiative fluxes, near surface relative humidity, precipitation and windspeed.

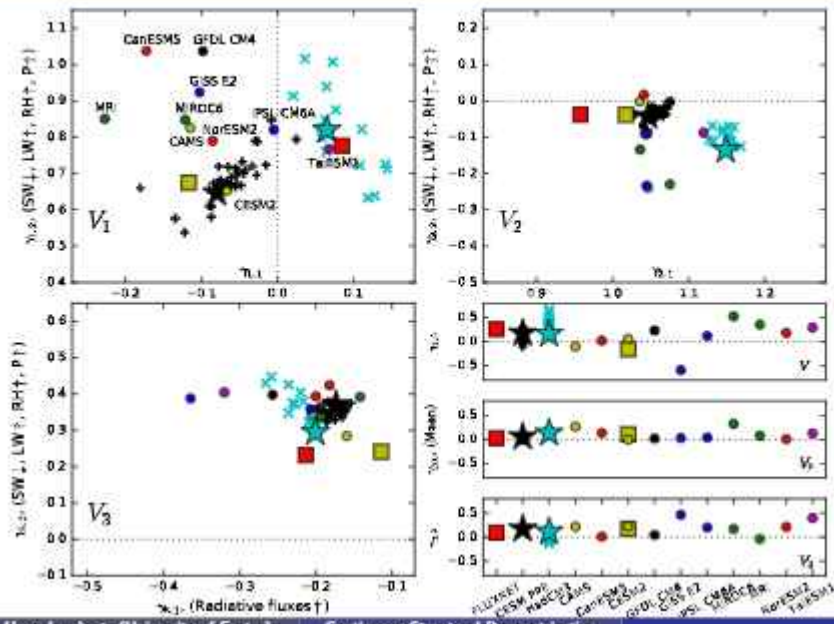
Find observation-model differences and **ask** if they matter.

Thanks to Anna Ukkola, Gab Abramowitz, Ben Booth, Ben Buchovecky, Monisha Natchiar, Clare Zarakas and Abigail Swann.

Input and output vectors



Linear analysis



Gaussian process emulators

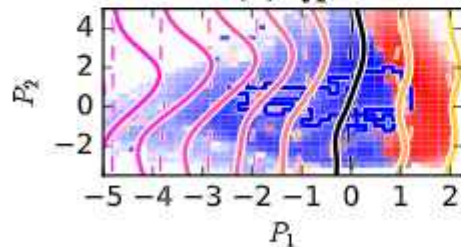
(a) CESM2 Q_1



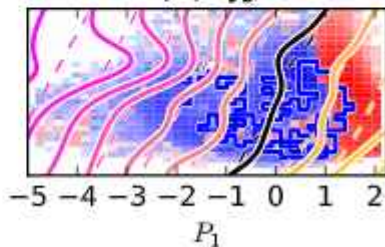
(b) GISS E2 Q_1



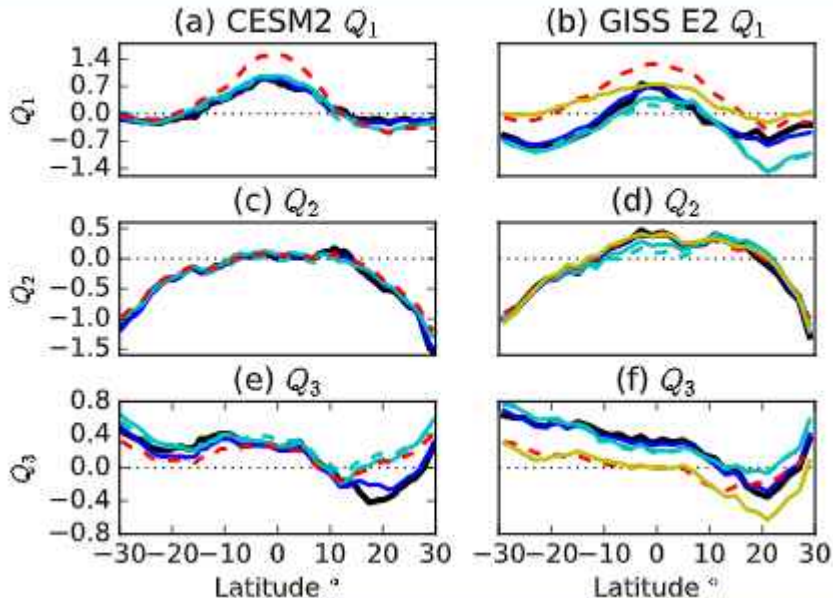
(c) Q_2



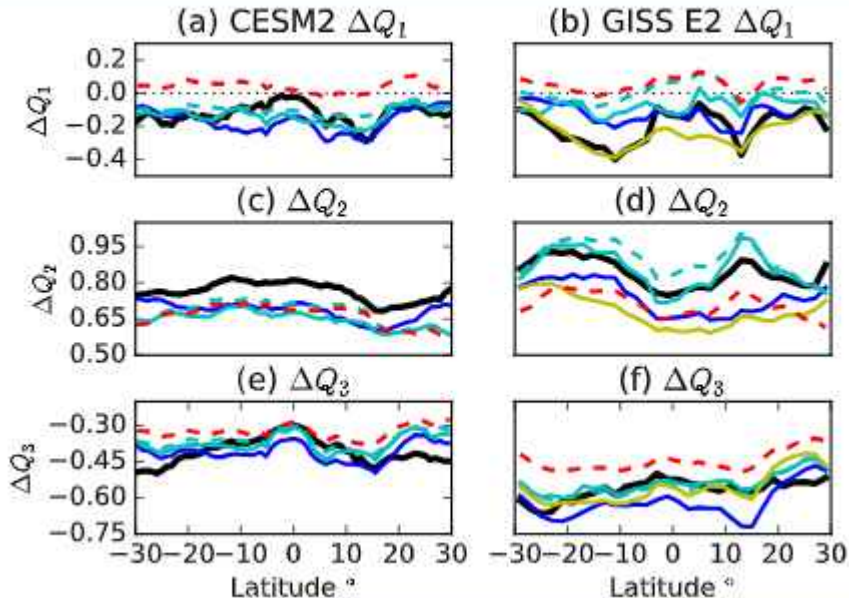
(d) Q_2



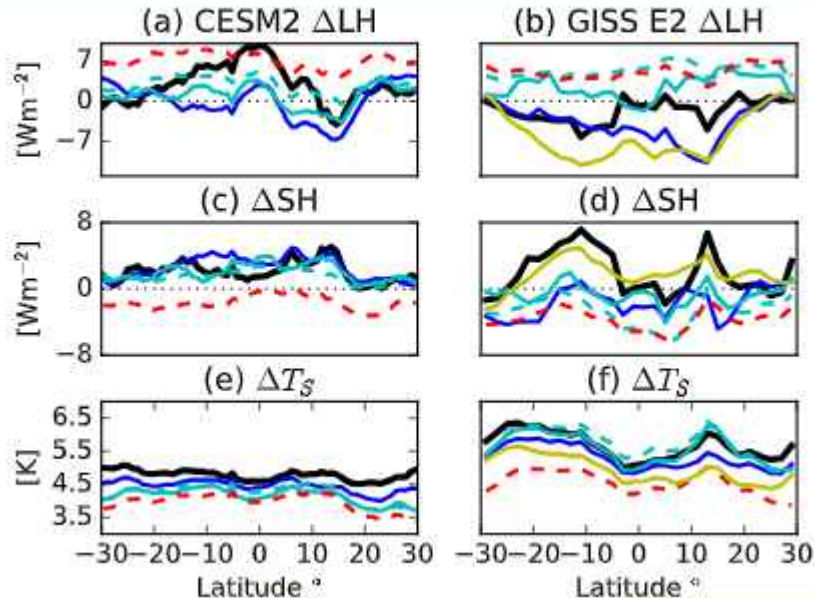
Verification of the present day



Predictions of changes in a +4 K climate



And for real variables...



We analysed **FLUXNET** observations and **GCMs** using an augmented **Continuous Structural Parametrization (CSP)** framework:

- The majority of models produce **too little latent heat** and **too much sensible heat** in response to increasing downward radiative fluxes. There are process differences in the land surface models. It's **not** just differences in model inputs.
- Lack of observations mean it is difficult to know what is happening in some regions, such as north of 10°N . I predict that CESM2 and GISS E2 will still not agree with each other after tuning to **FLUXNET**!
- CSP captures some features of **AMIP +4 K "climate change"**. A more **FLUXNET**-like model might produce cooler, wetter results.



B. B. B. Booth, C. D. Jones, M. Collins, I. J. Totterdell, P. M. Cox, S. Sitch, C. Huntingford, R. A. Betts, G. R. Harris, and J. Lloyd.
High sensitivity of future global warming to land carbon cycle processes.

Environ. Res. Lett., 7, 2012.
024002.



F. H. Lambert, P. G. Challenor, N. T. Lewis, D. J. McNeall, N. Owen, I. A. Boutle, H. M. Christensen, R. J. Keane, N. J. Mayne, A. Stirling, and M. J. Webb.

Continuous structural parameterization: A proposed method for representing different model parameterizations within one structure demonstrated for atmospheric convection.

Journal of Advances in Modeling Earth Systems,
12(8):e2020MS002085, 2020.



A. M. Ukkola, G. Abramowitz, and M. G. De Kauwe.
A flux tower dataset tailored for land model evaluation.

Earth Syst. Sci. Data, 14:449–461, 2022.



C. Zarakas, D. Kennedy, K. Dagon, D. Lawrence, A. Liu, G. Bonan, C. Koven, D. Lombardozzi, and A. Swann.

Land processes can substantially impact the mean climate state.

Geophys. Res. Lett., 2024.

Submitted.



University
of Exeter

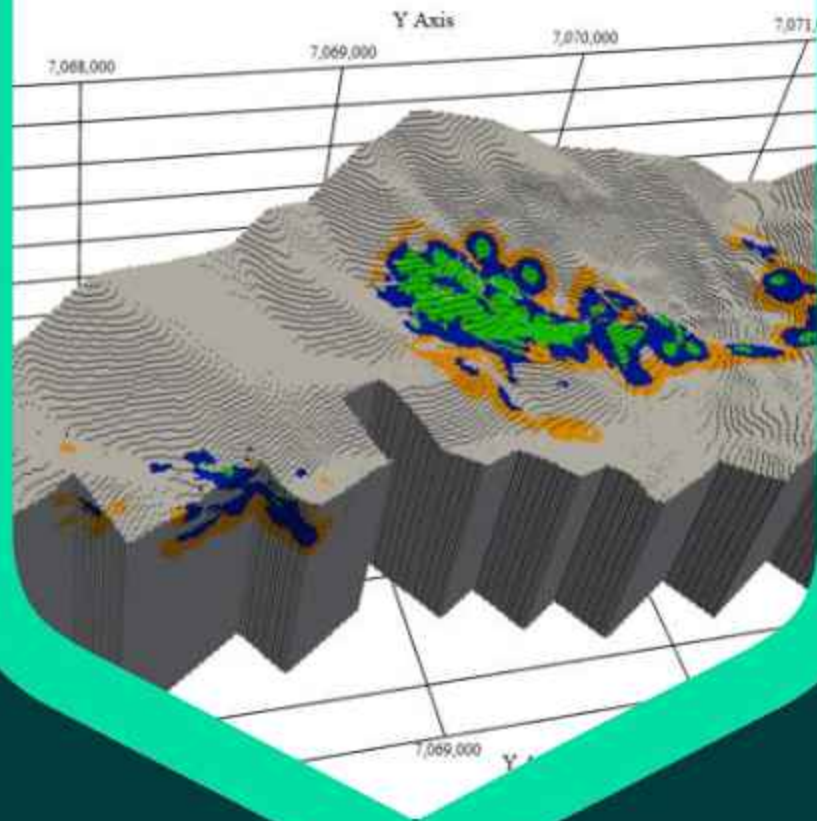
Camborne
School of Mines

Geostatistical uncertainty quantification in Earth Sciences... in 7 min!

Prof. Julian M. Ortiz, Ph.D., P.Eng.

Mark Cutifani / Anglo American Chair in Mining
Innovation | Camborne School of Mines (UK)

Co-Founder | CEO | Annapurna Suite (Canada)



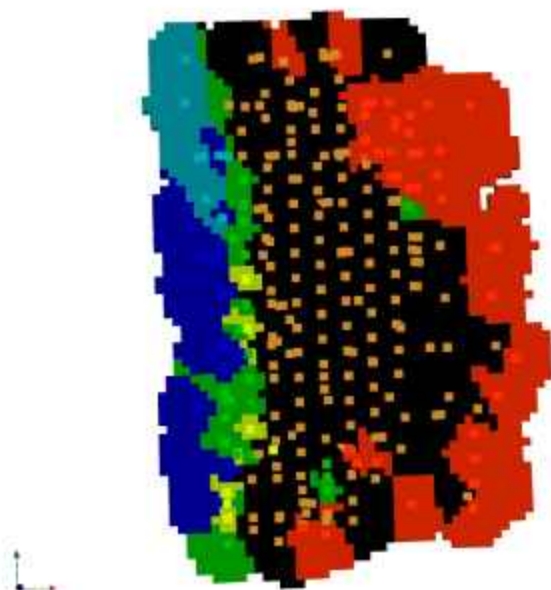
Motivation



University
of Exeter

Camborne
School of Mines

- In Earth Sciences we often need to **map** attributes in space
- We would like to know the **confidence** in our predictions
- **Continuous** variables often behave differently in different **categories**
- Decisions are made based on these maps and should consider **uncertainty**, as it translates to **risk**
- To complicate things:
 - We work in the **natural environment** → geological constraints
 - **Multiple attributes** may be involved in the decision → need to account for their relationships



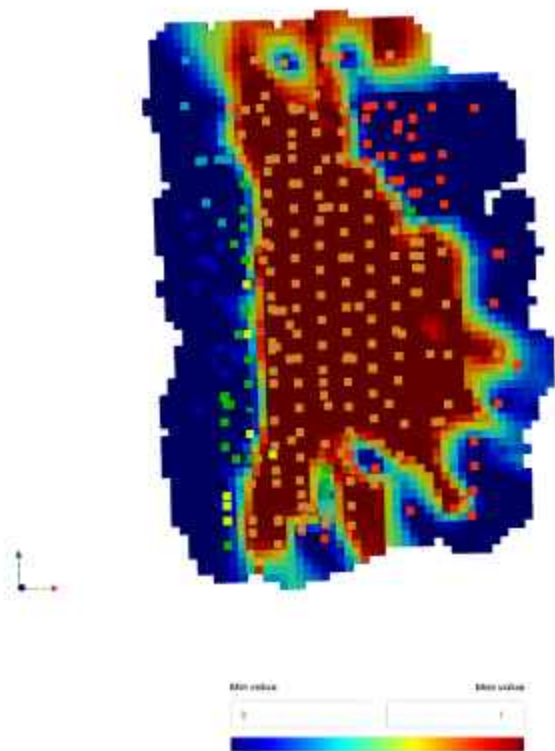
Motivation



University
of Exeter

Camborne
School of Mines

- In Earth Sciences we often need to **map** attributes in space
- We would like to know the **confidence** in our predictions
- **Continuous** variables often behave differently in different **categories**
- Decisions are made based on these maps and should consider **uncertainty**, as it translates to **risk**
- To complicate things:
 - We work in the **natural environment** → geological constraints
 - **Multiple attributes** may be involved in the decision → need to account for their relationships



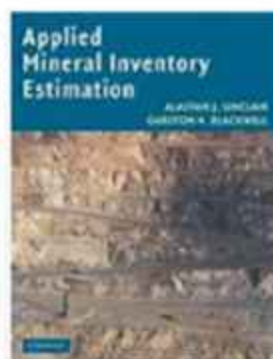
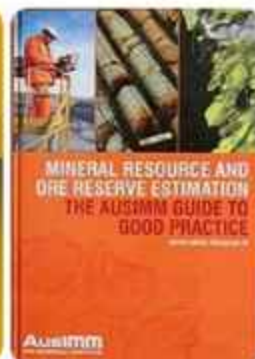
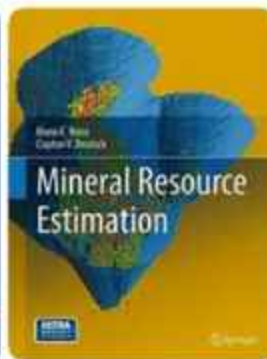
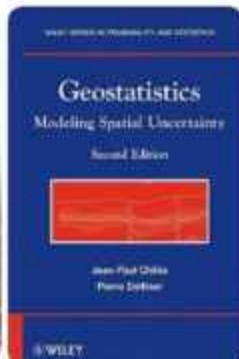
What is geostatistics?



University
of Exeter

Camborne
School of Mines

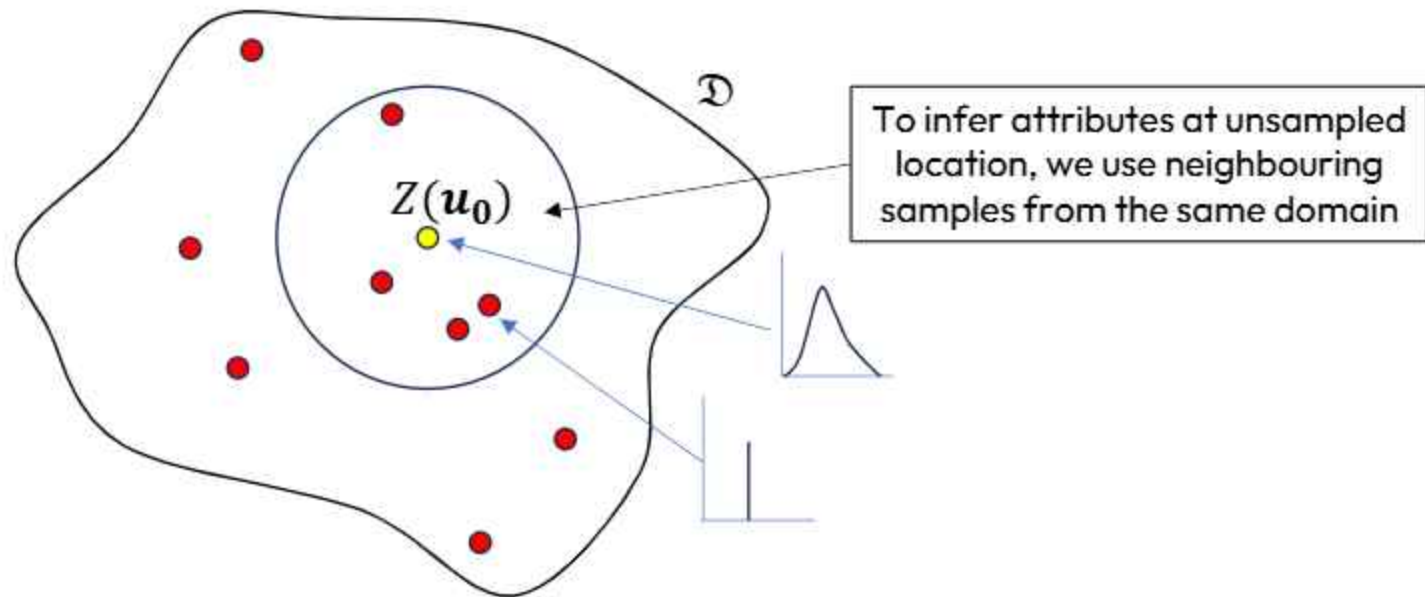
- The application of *probabilistic methods* to regionalized variables (Chiles and Delfiner, 2012)
- The branch of statistics concerned with the description and modelling of the *spatial distribution* of variables (Frenzel et al., 2023)
- Deals with *prediction* (estimation) and *uncertainty quantification* (through stochastic simulation)



How does it work?



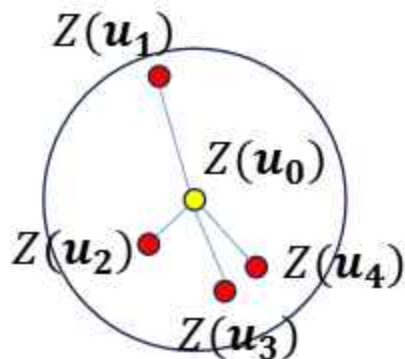
- The principles are simple: a variable in space is modelled as a **random function** $Z = \{Z(\mathbf{u}) \in \mathcal{D}\}$, that is, a collection of **random variables** $Z(\mathbf{u})$ that are spatially correlated.



How does it work?



- The principles are simple: a variable in space is modelled as a **random function** $Z = \{Z(u) \in \mathcal{D}\}$, that is, a collection of **random variables** $Z(u)$ that are spatially correlated.

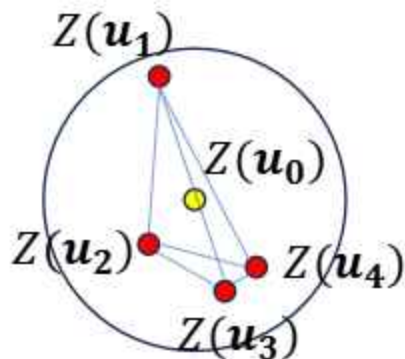


- Each sample is correlated with the unsampled location

How does it work?



- The principles are simple: a variable in space is modelled as a **random function** $Z = \{Z(u) \in \mathcal{D}\}$, that is, a collection of **random variables** $Z(u)$ that are spatially correlated.

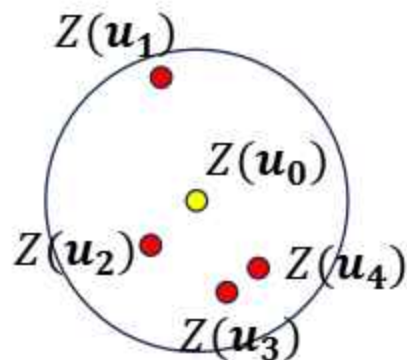


- Samples are correlated with each other

How does it work?



- The principles are simple: a variable in space is modelled as a **random function** $Z = \{Z(\mathbf{u}) \in \mathcal{D}\}$, that is, a collection of **random variables** $Z(\mathbf{u})$ that are spatially correlated.



We define a linear unbiased optimal estimate \rightarrow **Kriging**

- We need to determine the **weights**
- Linear estimate

$$Z^*(\mathbf{u}_0) = \lambda_0 + \sum_{i=1}^n \lambda_i Z(\mathbf{u}_i)$$

- Unbiased estimate

$$E\{Z^*(\mathbf{u}_0)\} = E\{Z(\mathbf{u}_0)\}$$

- Optimal estimate

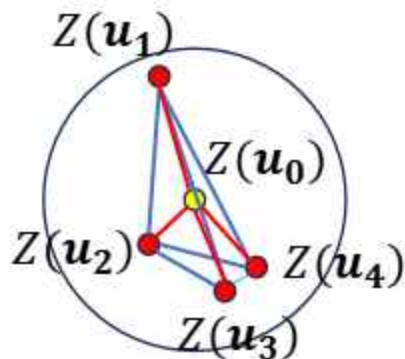
$$\min_{\{\lambda_i, i=0, \dots, n\}} \text{Var}\{Z^*(\mathbf{u}_0) - Z(\mathbf{u}_0)\}$$

- We assume behaviour of random variables is **statistically homogeneous** (domain \rightarrow stationarity)

How does it work?



- The principles are simple: a variable in space is modelled as a **random function** $Z = \{Z(\mathbf{u}) \in \mathbb{D}\}$, that is, a collection of **random variables** $Z(\mathbf{u})$ that are spatially correlated.



- System of equations

$$\underbrace{\begin{bmatrix} C_{11} & C_{12} & \cdots & C_{1n} \\ C_{21} & C_{22} & \cdots & C_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ C_{n1} & C_{n2} & \cdots & C_{nn} \end{bmatrix}}_{\text{covariances between the samples}} \underbrace{\begin{pmatrix} \lambda_1^{SK} \\ \lambda_2^{SK} \\ \vdots \\ \lambda_n^{SK} \end{pmatrix}}_{\text{weights}} = \underbrace{\begin{pmatrix} C_{10} \\ C_{20} \\ \vdots \\ C_{n0} \end{pmatrix}}_{\text{covariances between samples and estimation location}}$$

- Linear estimate

$$Z^*(\mathbf{u}_0) = \lambda_0 + \sum_{i=1}^n \lambda_i Z(\mathbf{u}_i)$$

- Estimation variance

$$\sigma_{SK}^2(\mathbf{u}_0) = \sigma_0^2 - \sum_{i=1}^n \lambda_i^{SK} C_{i0}$$

How does it work?



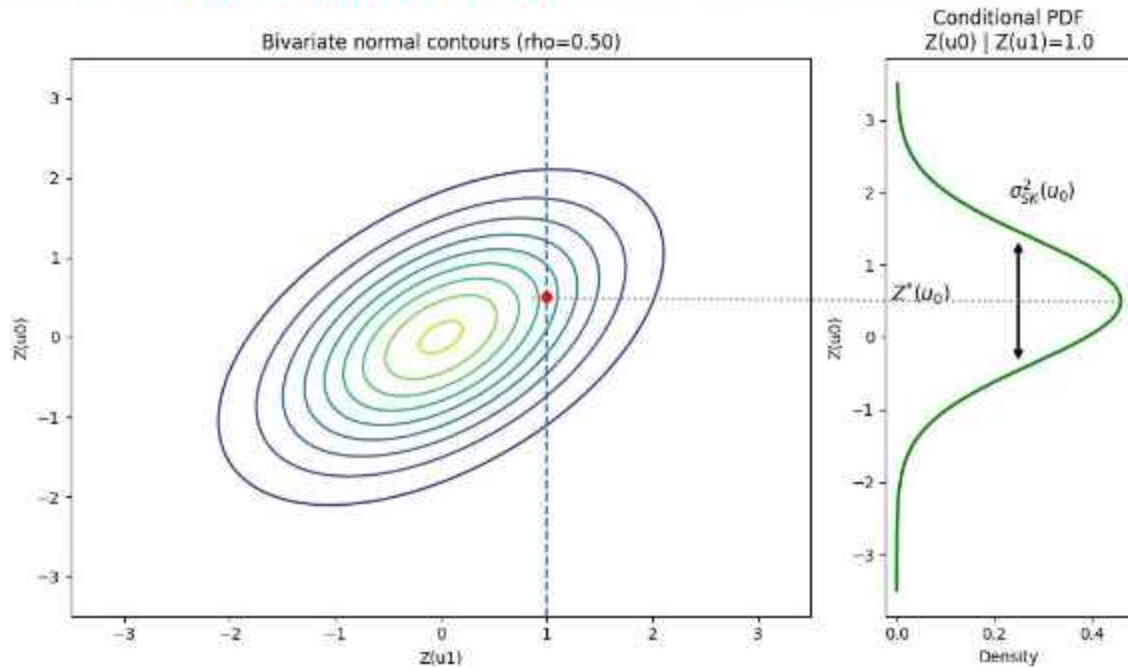
University
of Exeter

Camborne
School of Mines

- If variables are multiGaussian, the estimate and estimation variance are the conditional mean and conditional variance, thus we can infer the **conditional distribution** (uncertainty distribution).

$$Z^*(\mathbf{u}_0) = \lambda_0 + \sum_{i=1}^n \lambda_i Z(\mathbf{u}_i)$$

$$\sigma_{SK}^2(\mathbf{u}_0) = \sigma_0^2 - \sum_{i=1}^n \lambda_i^{SK} C_{i0}$$



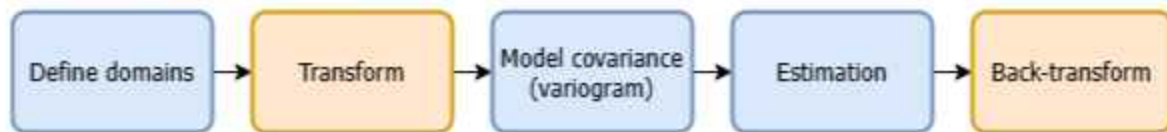
Typical workflow



University
of Exeter

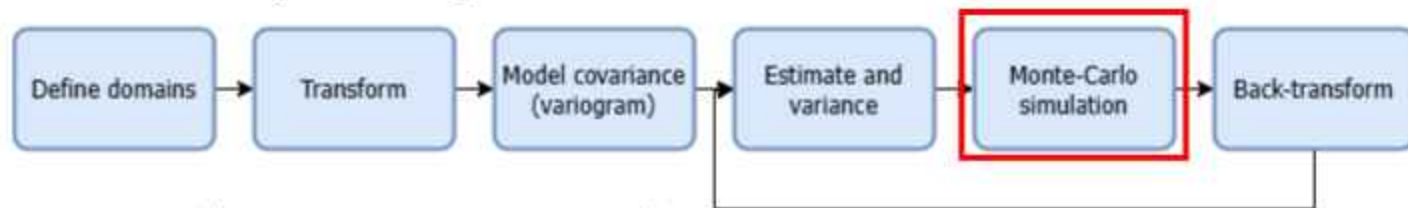
Camborne
School of Mines

- Continuous variables (estimation)



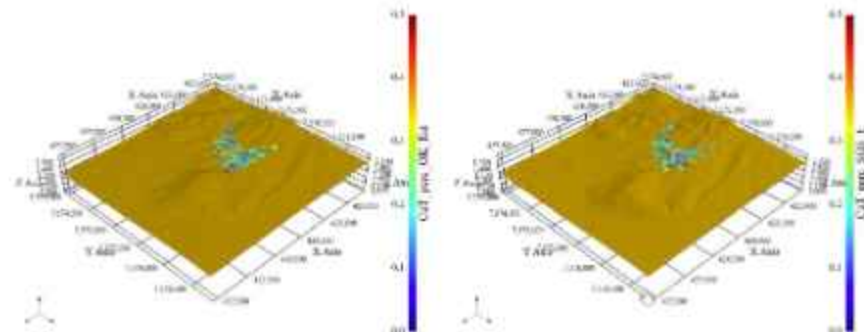
Kriging

- Continuous variables (simulation)



Sequential
simulation

Simulated value used to condition new
simulated nodes



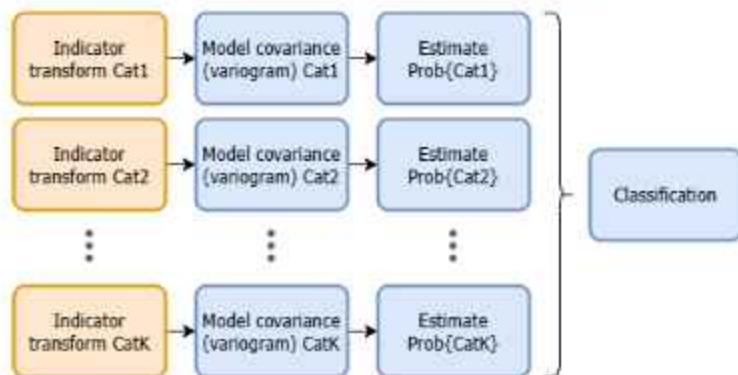
Typical workflow



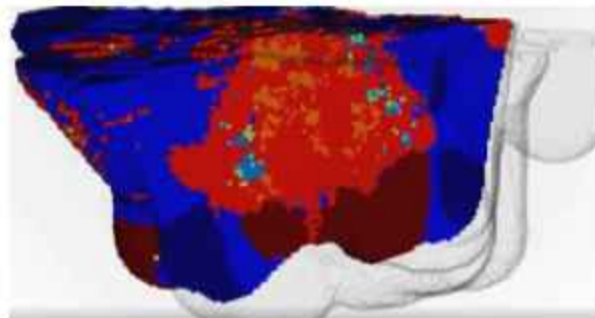
University
of Exeter

Camborne
School of Mines

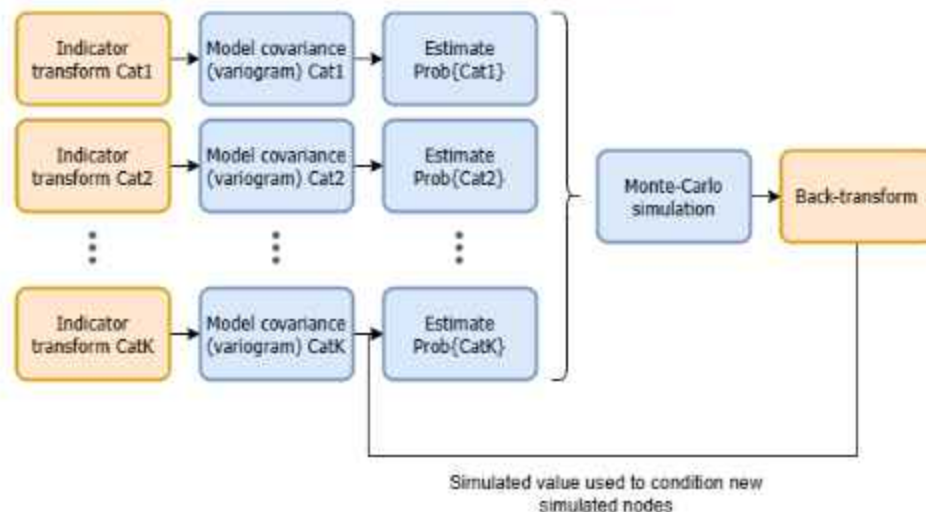
• Categorical variables (prediction)



Indicator
Kriging



• Categorical variables (simulation)



Indicator
Simulation

Implementation



University
of Exeter

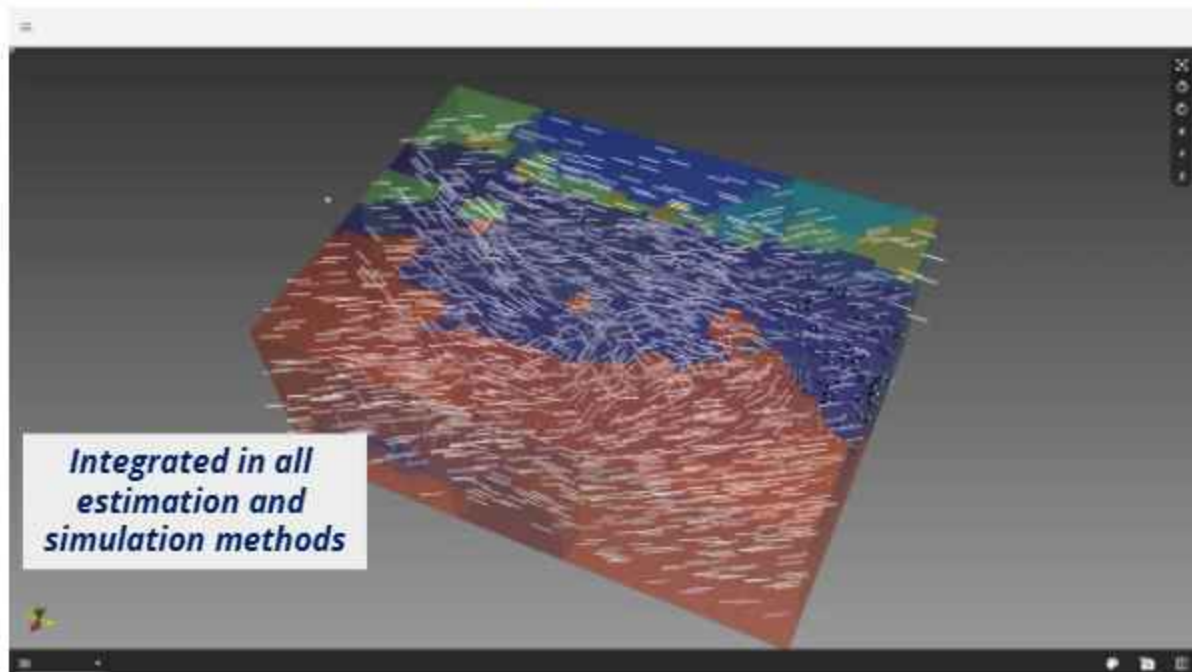
Camborne
School of Mines

Local Anisotropy Field (LVA) computation



*Advanced
methods*

Interpolation of local
orientations with Direction
Cosine Matrix / Quaternions
modelling



*Integrated in all
estimation and
simulation methods*

Implementation



University
of Exeter

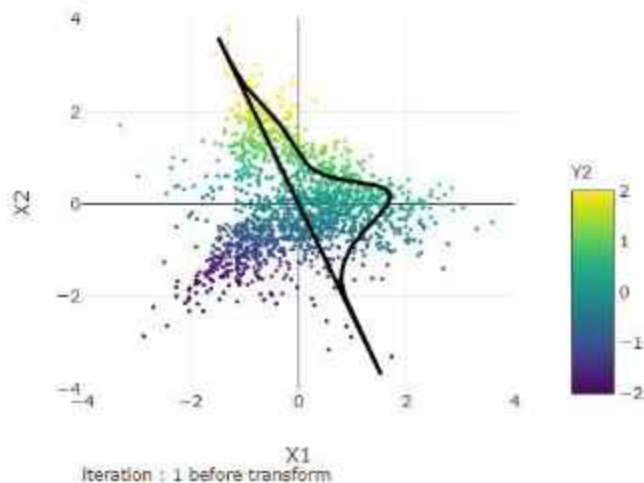
Camborne
School of Mines

Multivariate (PPMT) Simulation



Advanced
methods

- Normal score transform
- Iterative Sphering
- **Multi-Gaussian distribution**
- Independent component simulation
- Back-transform



<https://geostatisticslessons.com/lessons/ppmt>

Implementation



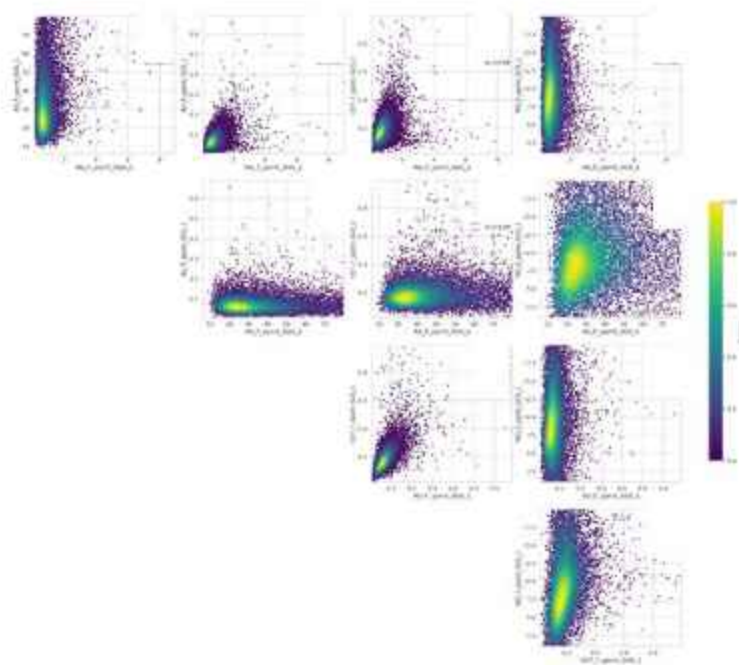
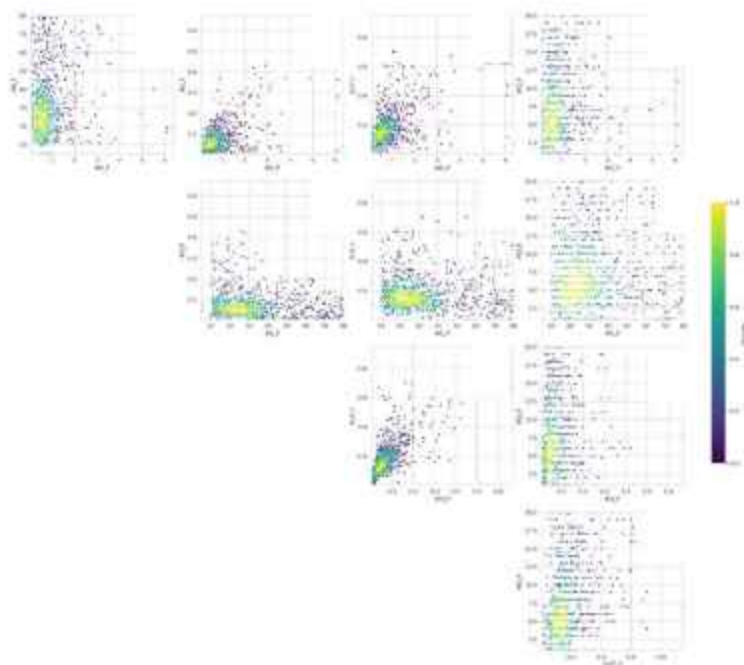
University
of Exeter

Camborne
School of Mines

Multivariate (PPMT) Simulation



*Advanced
methods*



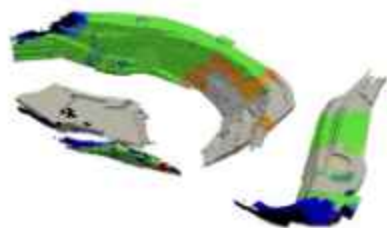
Example



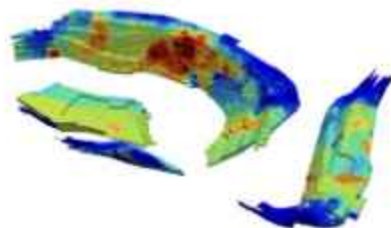
University
of Exeter

Camborne
School of Mines

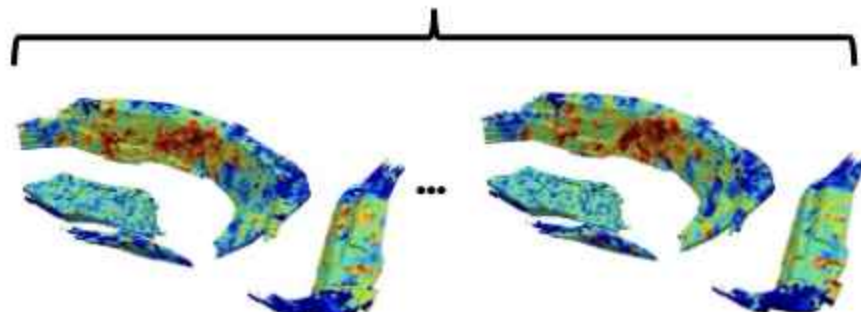
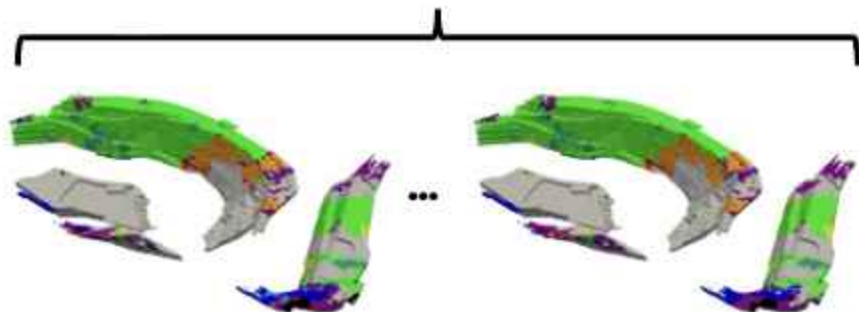
- Integrated categorical + continuous (multivariate) workflow to assess risk in decision (transfer function)



Domains



Attributes



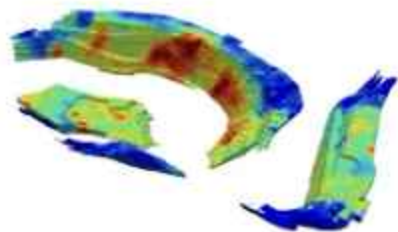
Example



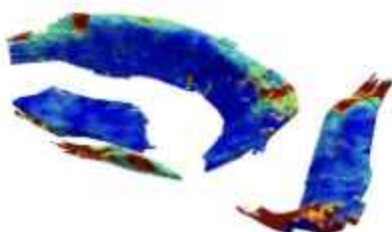
University
of Exeter

Camborne
School of Mines

- Integrated categorical + continuous (multivariate) workflow to assess risk in decision (transfer function)



Expected



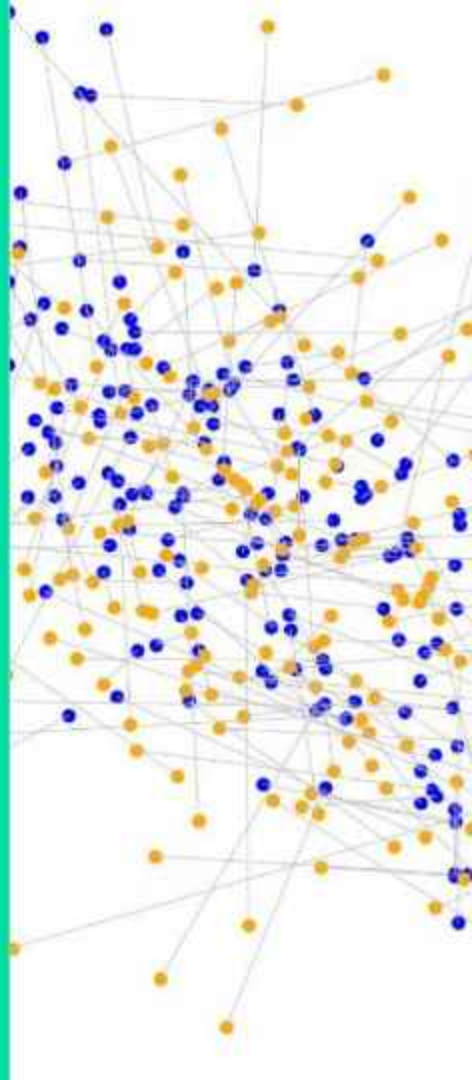
CoV



Risk

We can identify areas of high risk and inform decision making

Discussion



Discussion



University
of Exeter

Camborne
School of Mines

- **Geostatistics** provides methods for the estimation (prediction) and simulation (uncertainty quantification) of categorical and continuous variables
- Advanced techniques allow adapting these methods to **local anisotropies** and integrating **multivariate relationships** (statistical and spatial) into the modeling workflows
- Workflows combine the characterisation of **categorical** variables (soil types, rock types, soil use, lithologies, etc.) with the spatial distribution of **continuous** attributes within these categories (geochemical concentrations, hardness, humidity, specific gravity, etc.)
- Resulting models characterise the **joint uncertainty** and can be used for risk assessment and mitigation



University
of Exeter

Statistical Learning Applications in Earth Sciences

Presenter: Khan Muhammad

Environmental Intelligence @ Exeter Conference

University of Exeter

11-9-2025

Contents:

Background

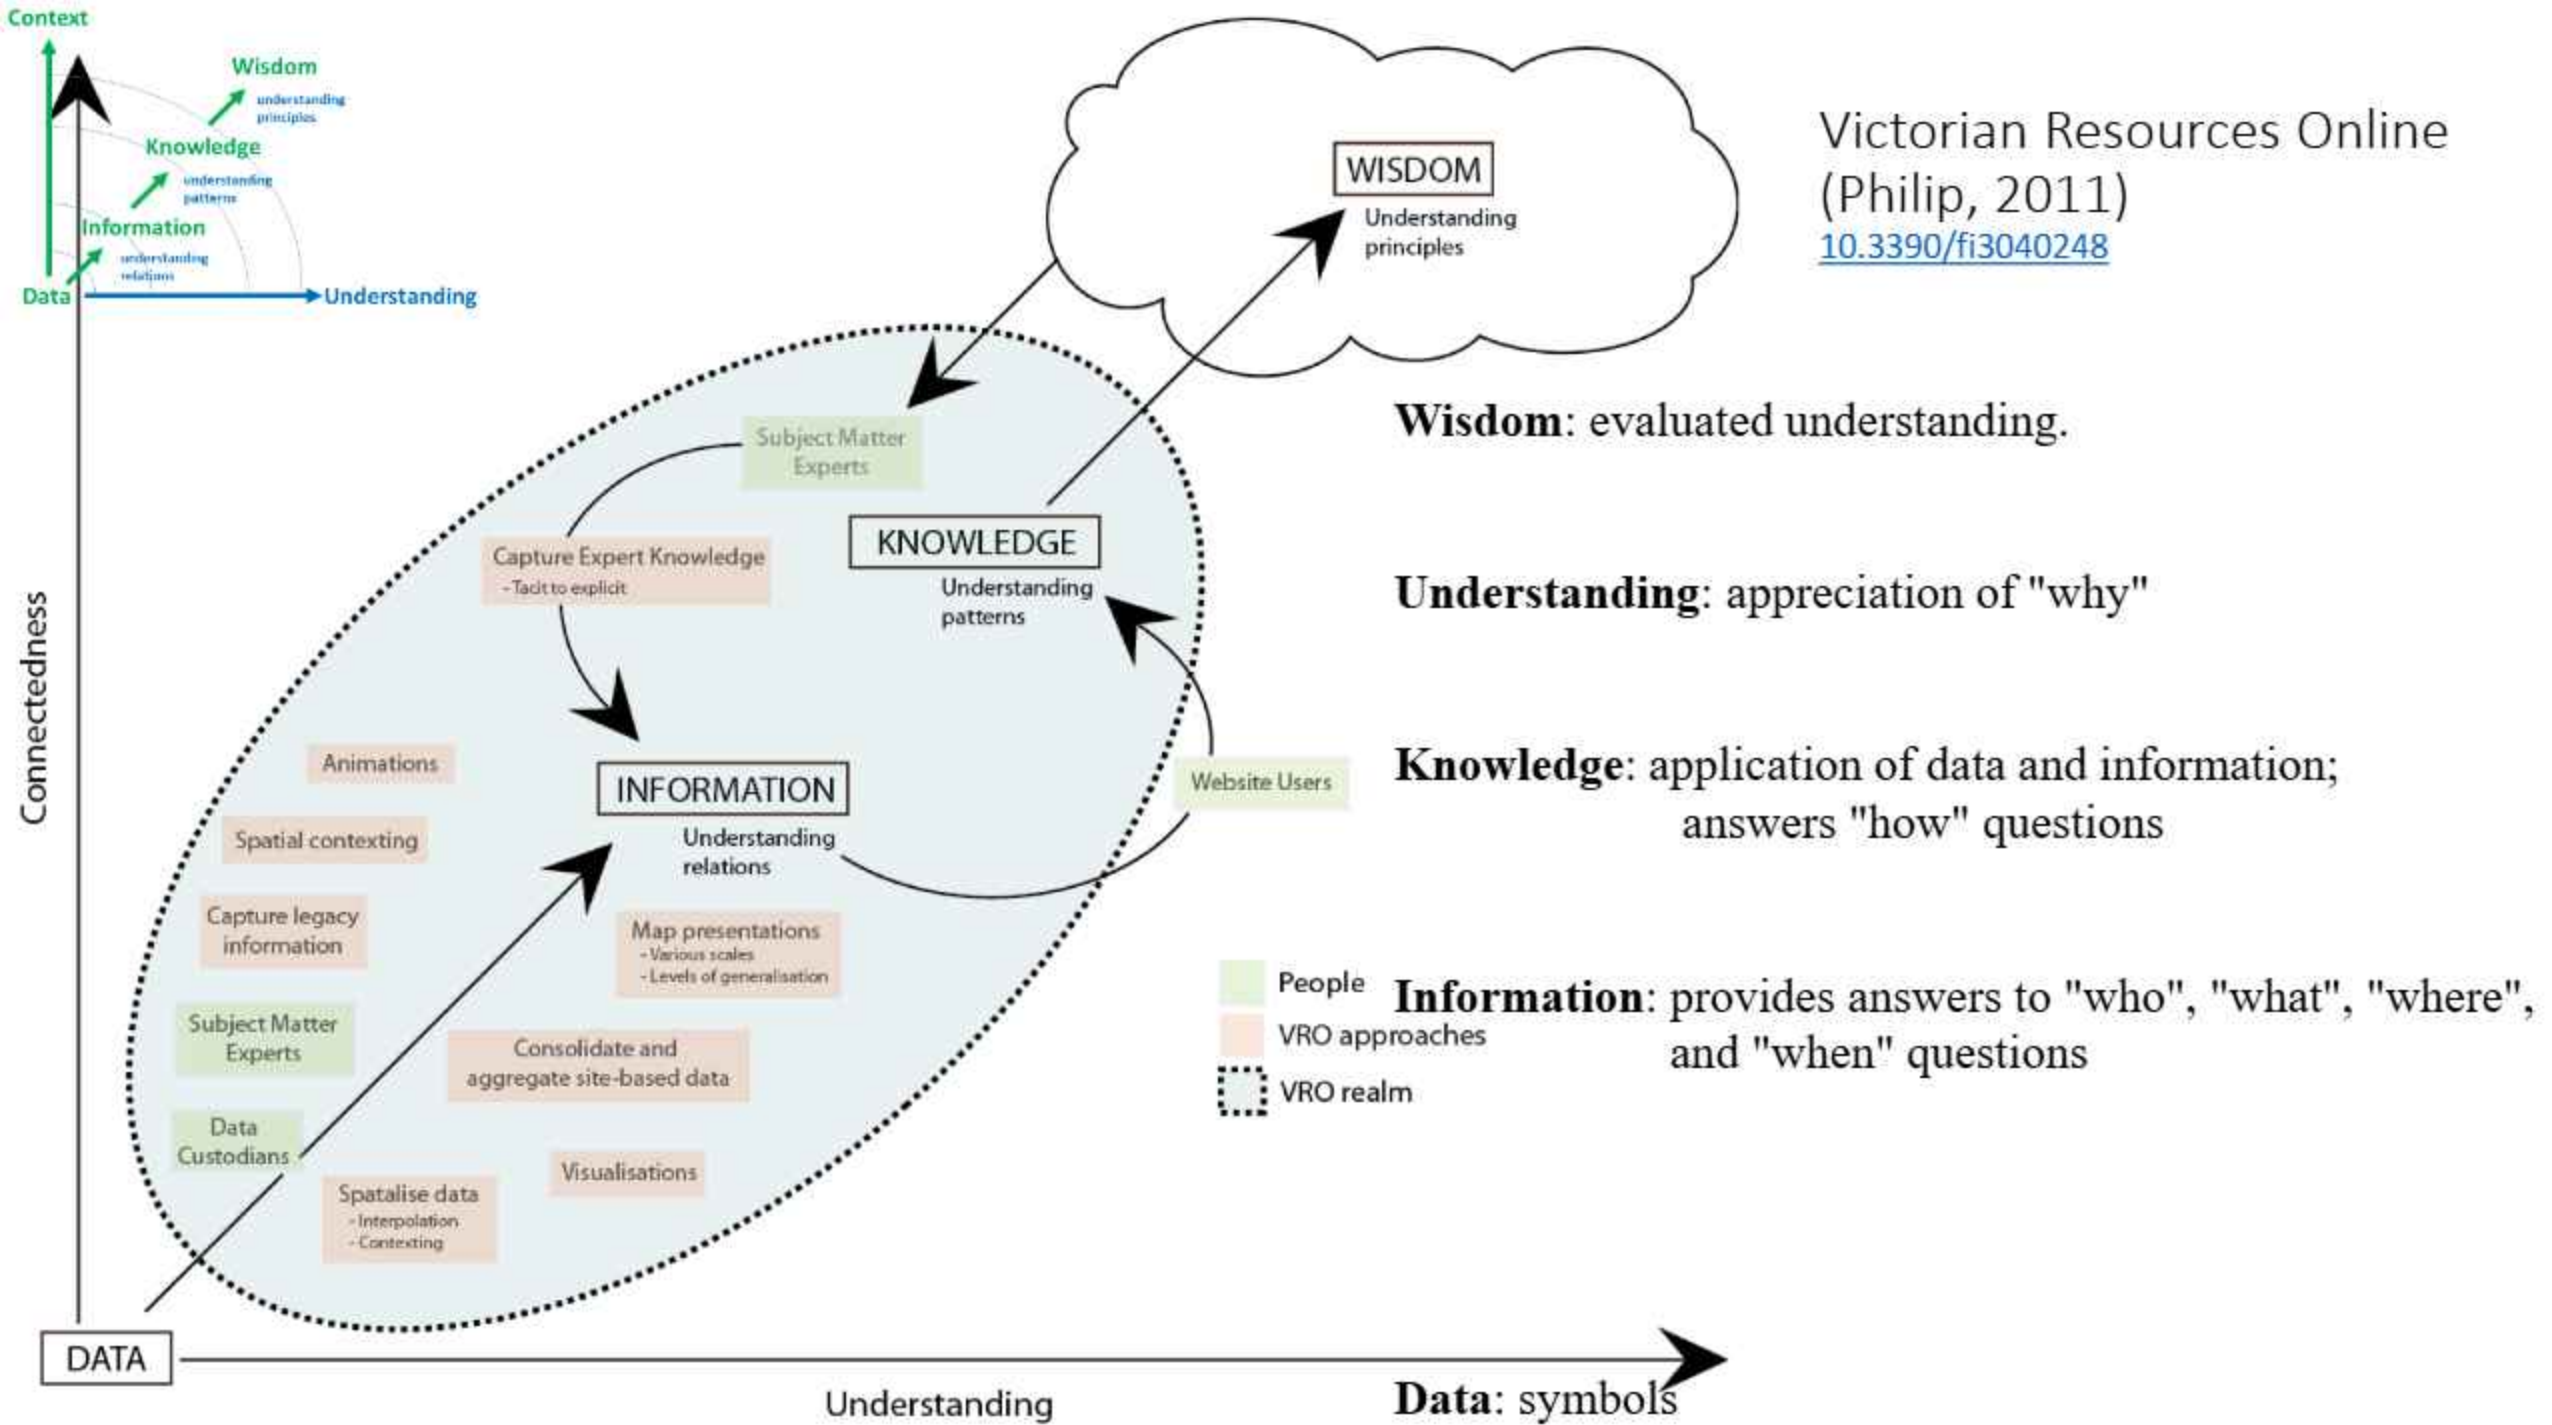
Geometallurgical Domaining

Production Scheduling

Conclusions



University
of Exeter



Develop AI based Mineral Resource Estimation Models

- 2D Remote Sensing Data

Pattern Recognition
Supervised/Unsupervised classification
Descriptive Analytics

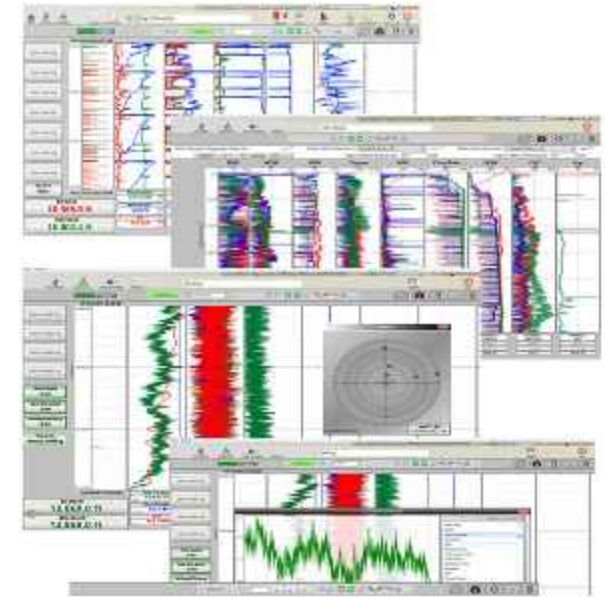
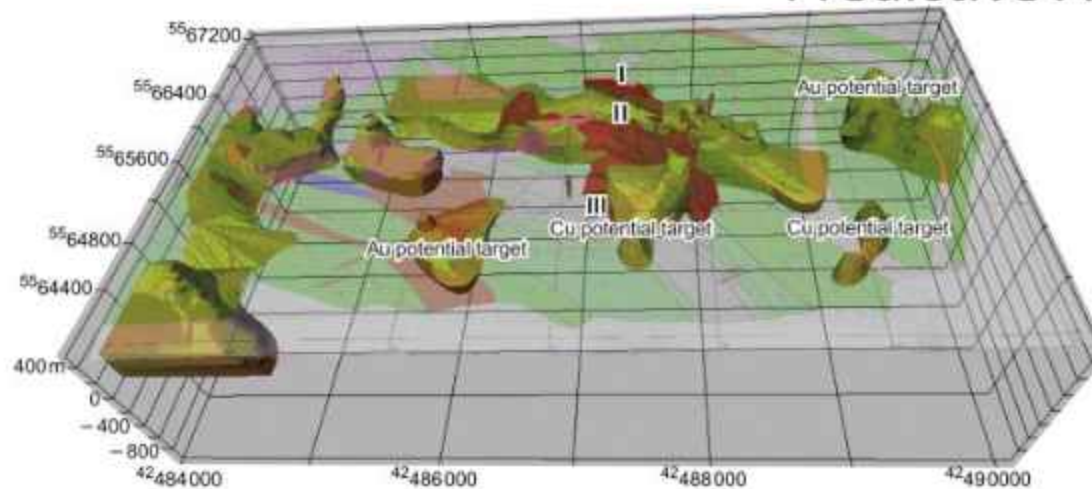
+

- Subsurface data

Spatial
Interpolation
techniques

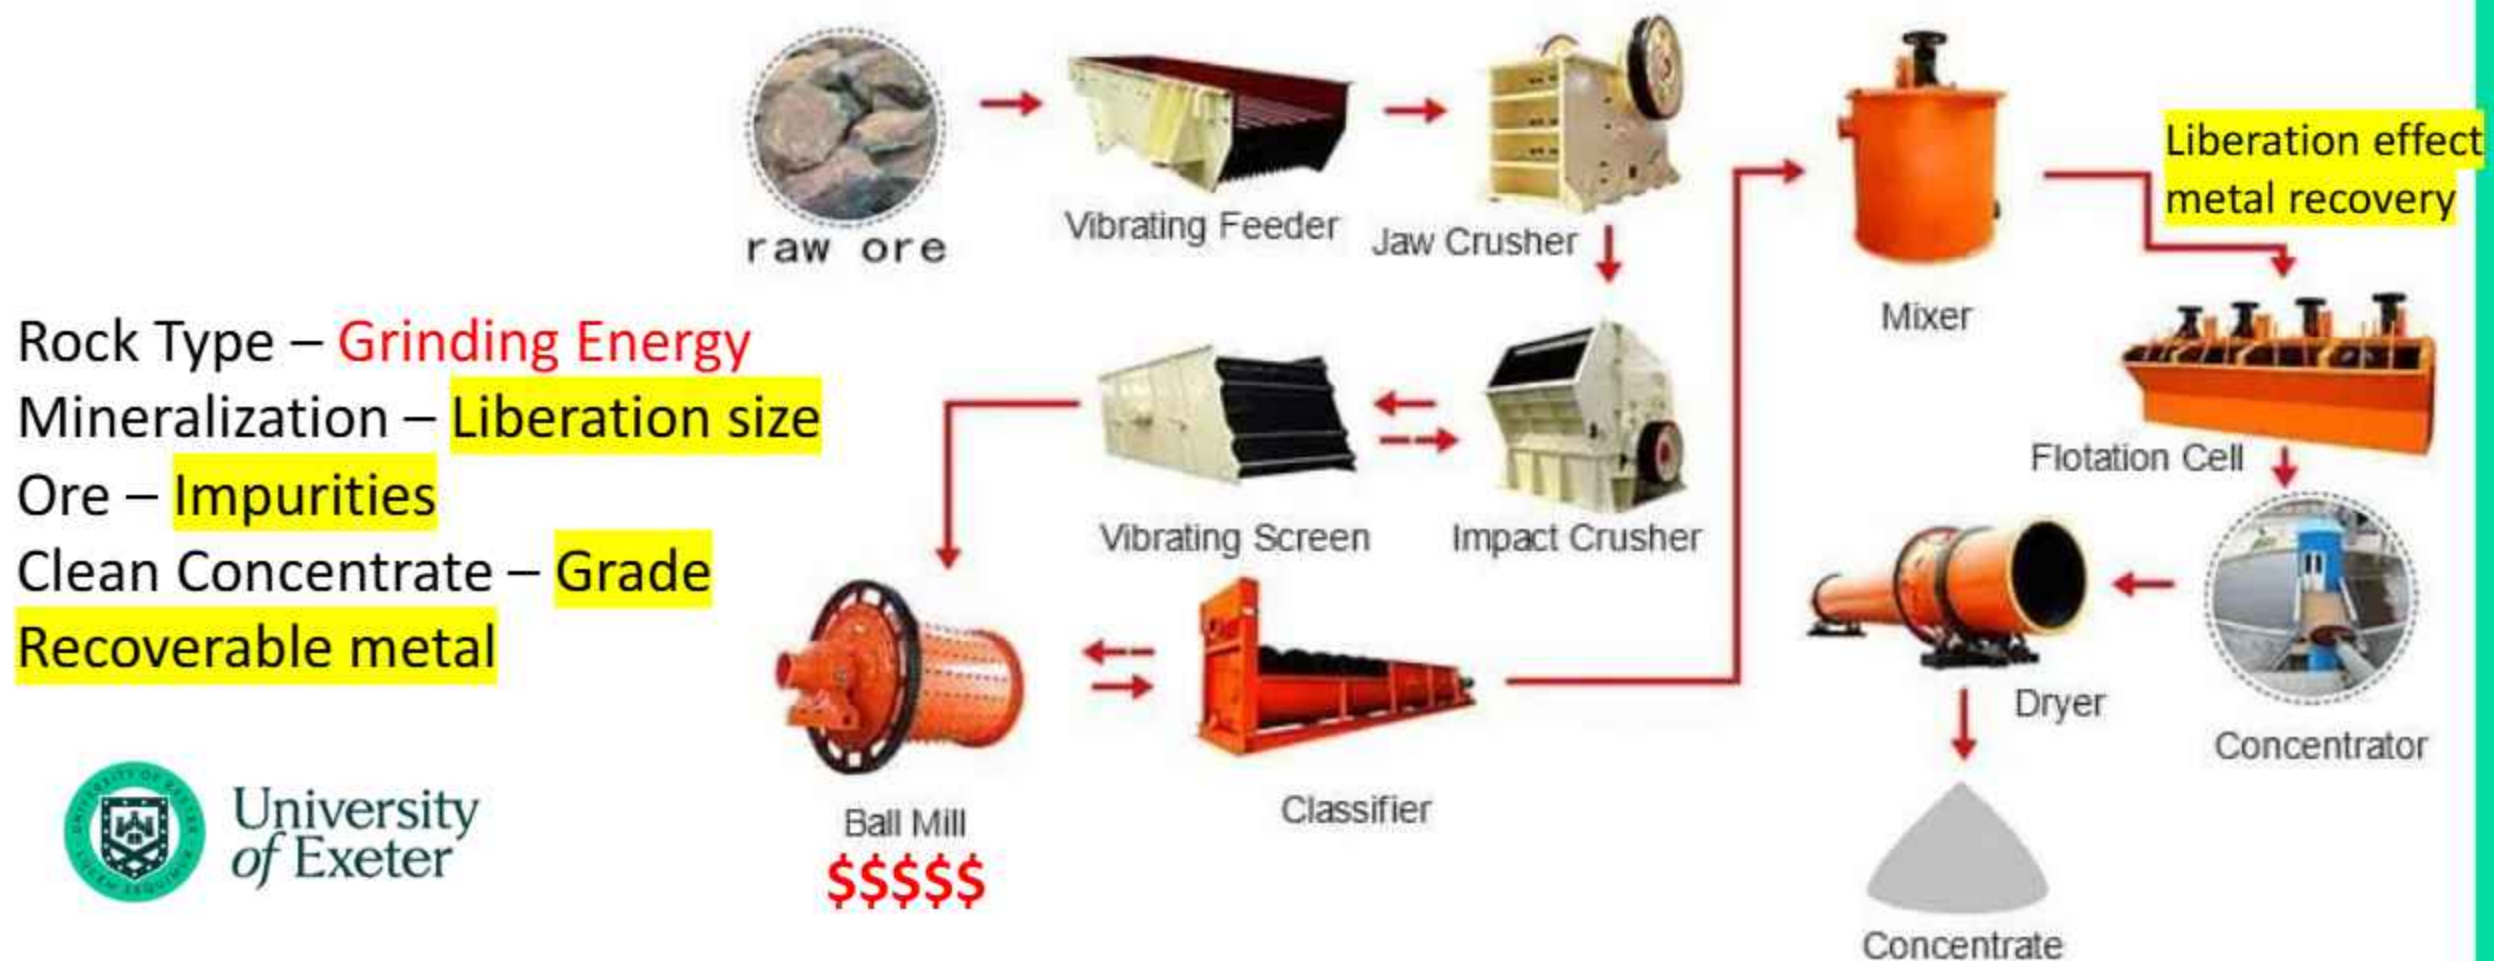
||

Predictive Analytics



Geometallurgical Domaining

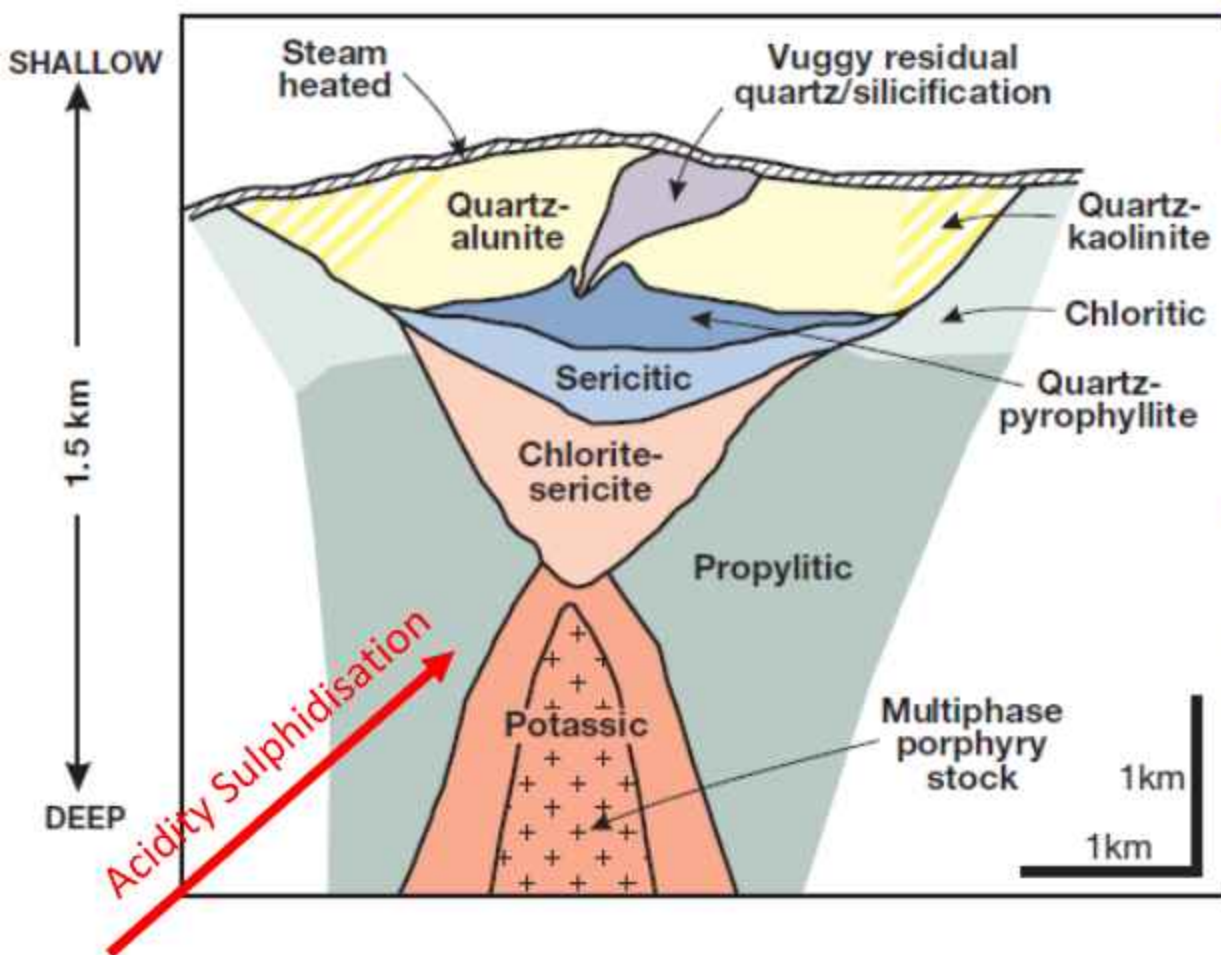
- Geometallurgical domains separate different processing streams for oxide, sulphide, or rock, based on strength and grindability.



University
of Exeter

Variables Selection

As+Sb+Hg



Argillic: Stable Ti/Al , Volatile enrichment + base leaching;
feldspar and mafic breakdown
 $(Al+Ti)/(Na+Ca+Mg)$, $(As+Sb+Hg)/S$

Qtz/Chl-sericite: Rb & Sr substitutes K & Ca; Volatiles
enrichment; B+Sn+W+Bi Mobility
 Rb/Sr , $(As+Sb+Hg)/(Ca+Mg)$,
 $(B+Sn+W+Bi)/(Fe+Ti+Al)$

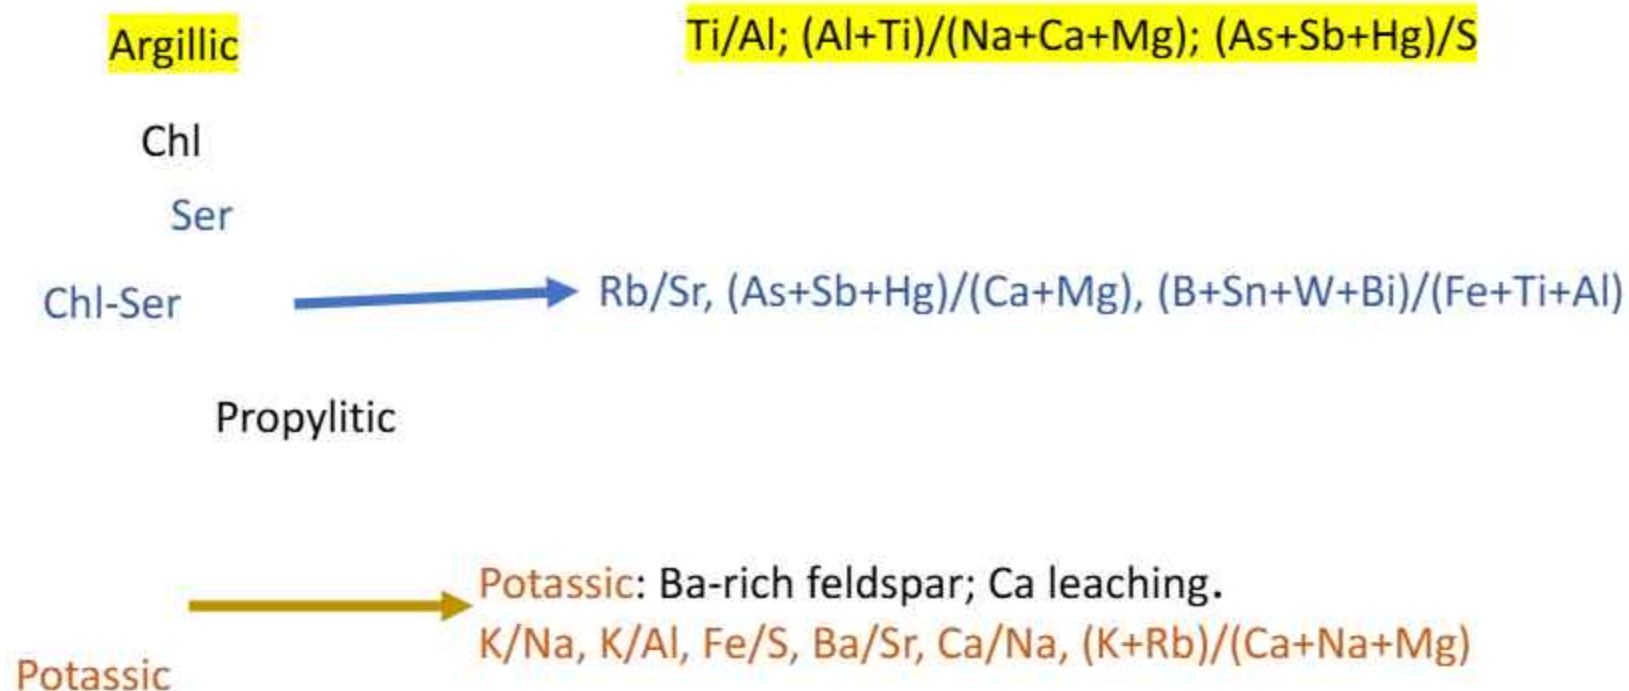


University
of Exeter

R. H. Sillitoe, "Porphyry Copper Systems*." [Online].

Available: <http://pubs.geoscienceworld.org/segweb/economicgeology/article-pdf/105/1/3/3464429/3.pdf>

Variables Selection 19 + 15 ratios = 34



K, Na, Ca, Mg, Al, Ti, S, Ba, Sr, Rb, B, P, As, Sb, Hg, Sn, W, Bi, Te



University
of Exeter

Aims and Objectives

Aim: Present a statistically valid and spatially contagious Unsupervised Geometallurgical domaining

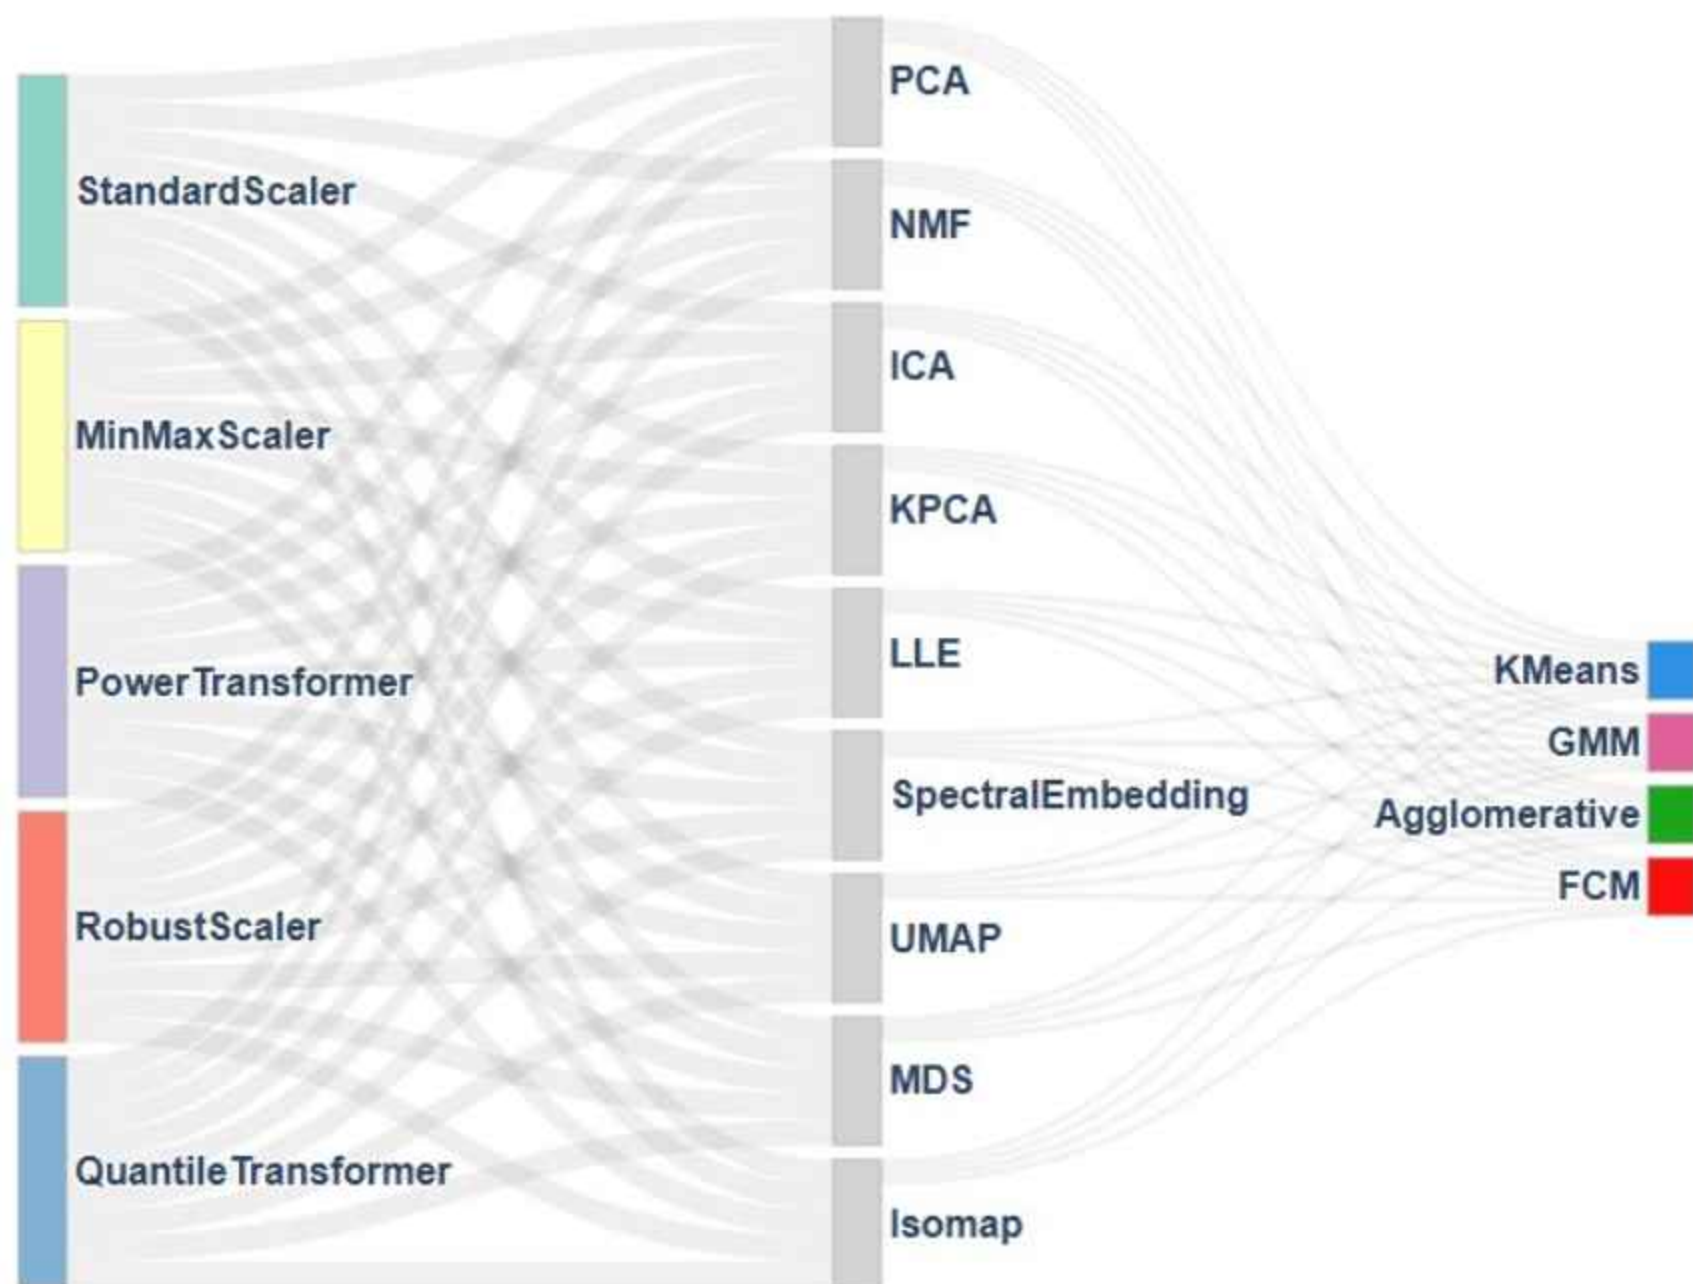
Objectives:

- Apply Feature Extraction and Classification
- Extract Core Samples with Least Neighbourhood Entropy
- Clean Core Domain zones: Pruning
- Train the ML model on core and apply it to non-core samples

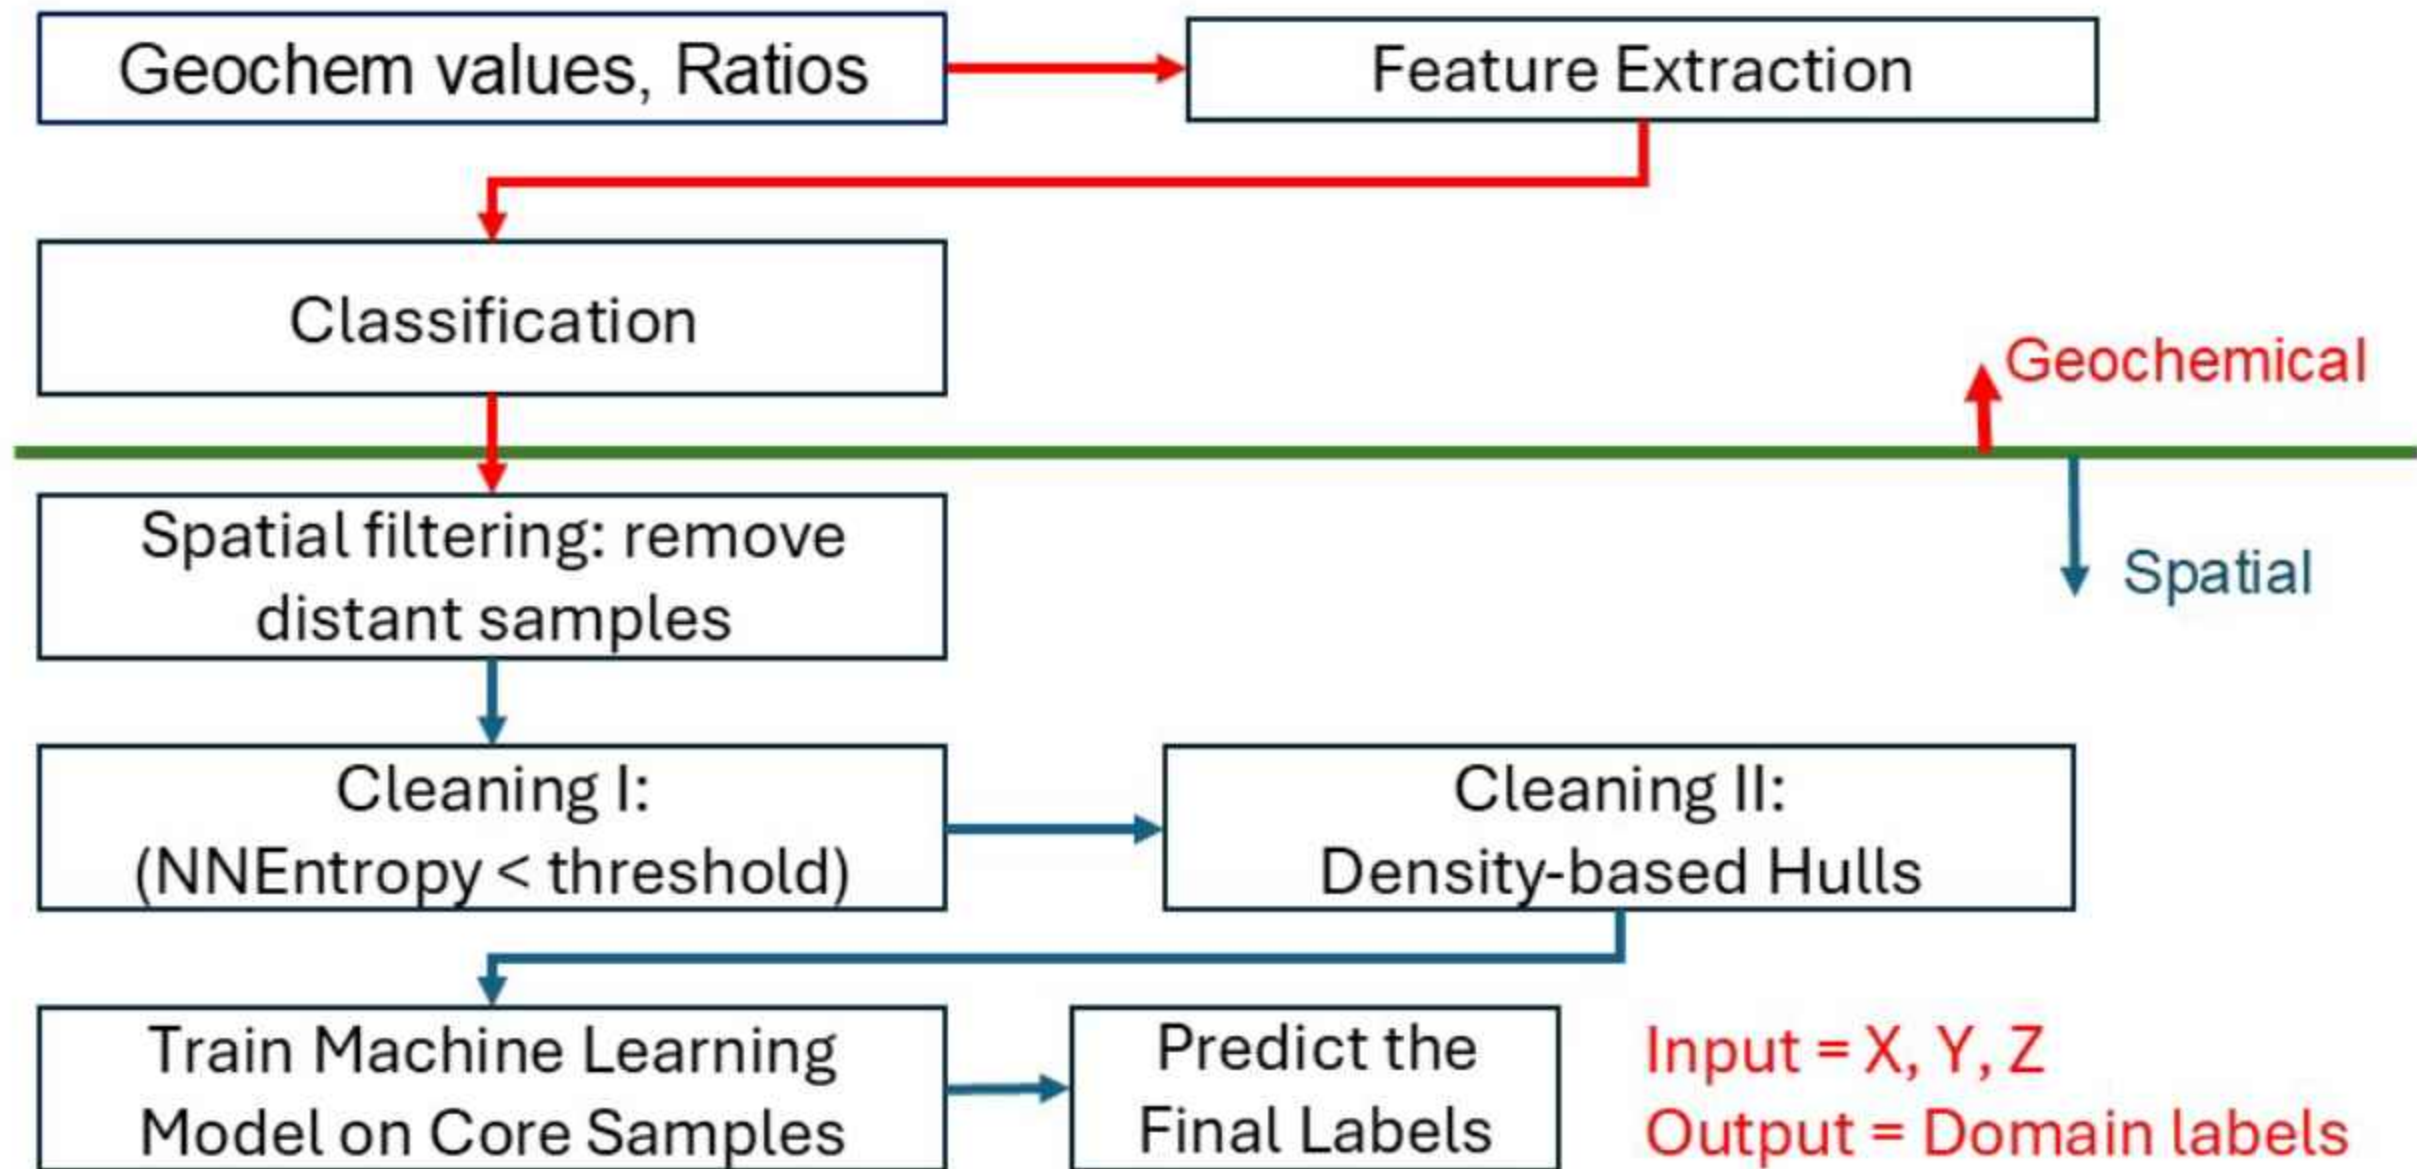


University
of Exeter

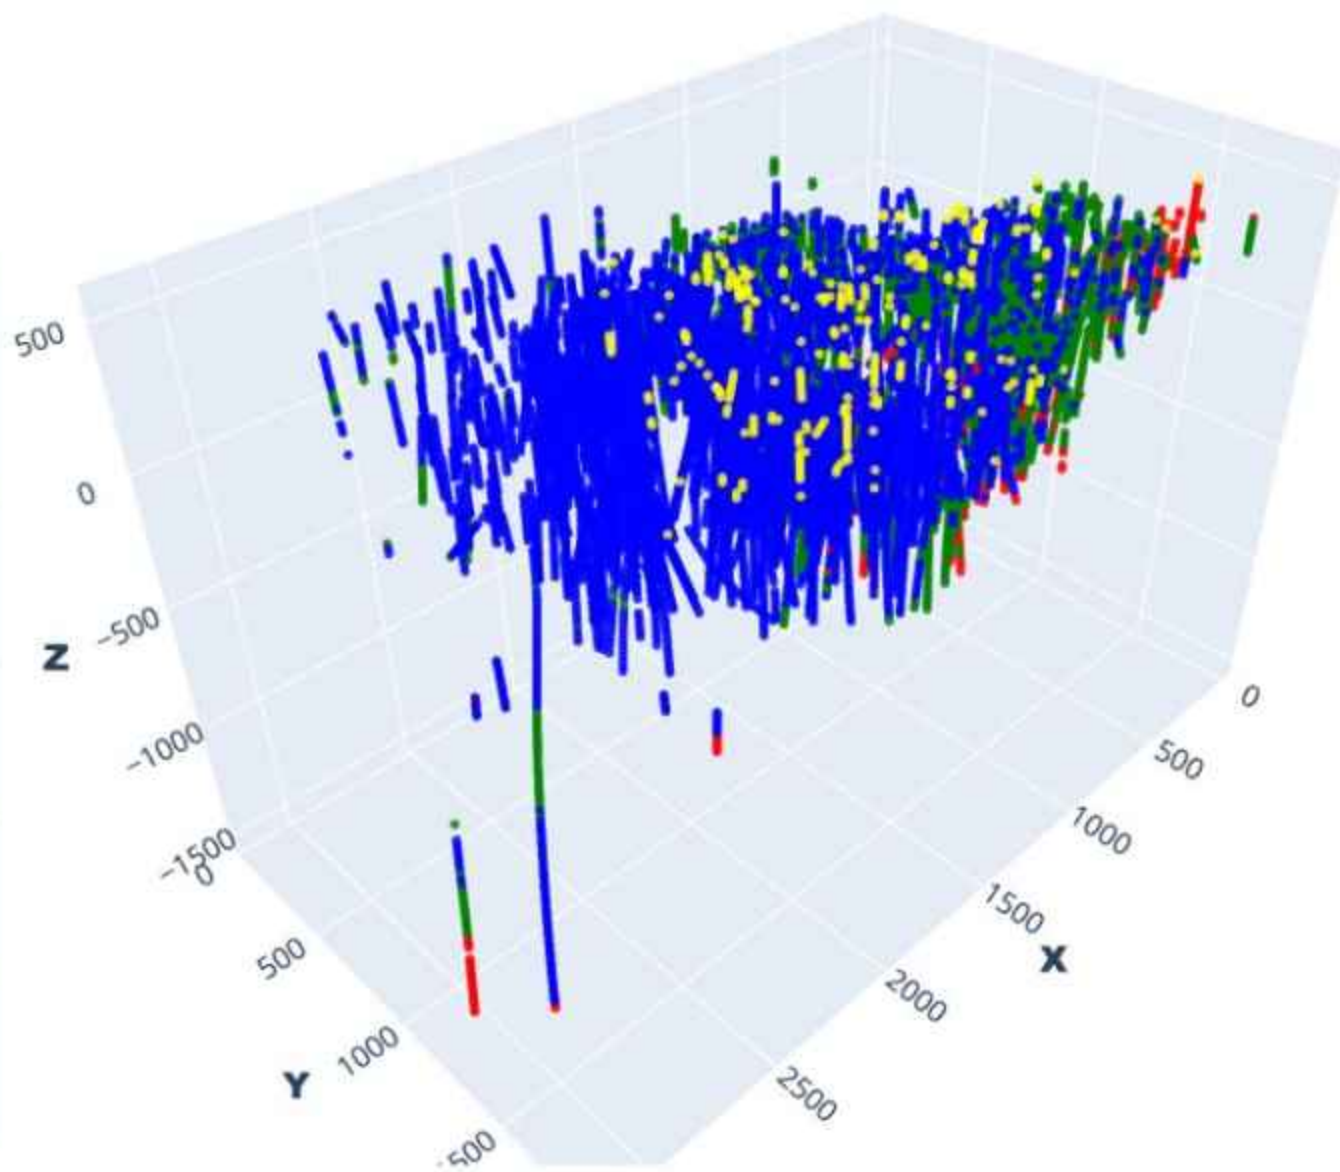
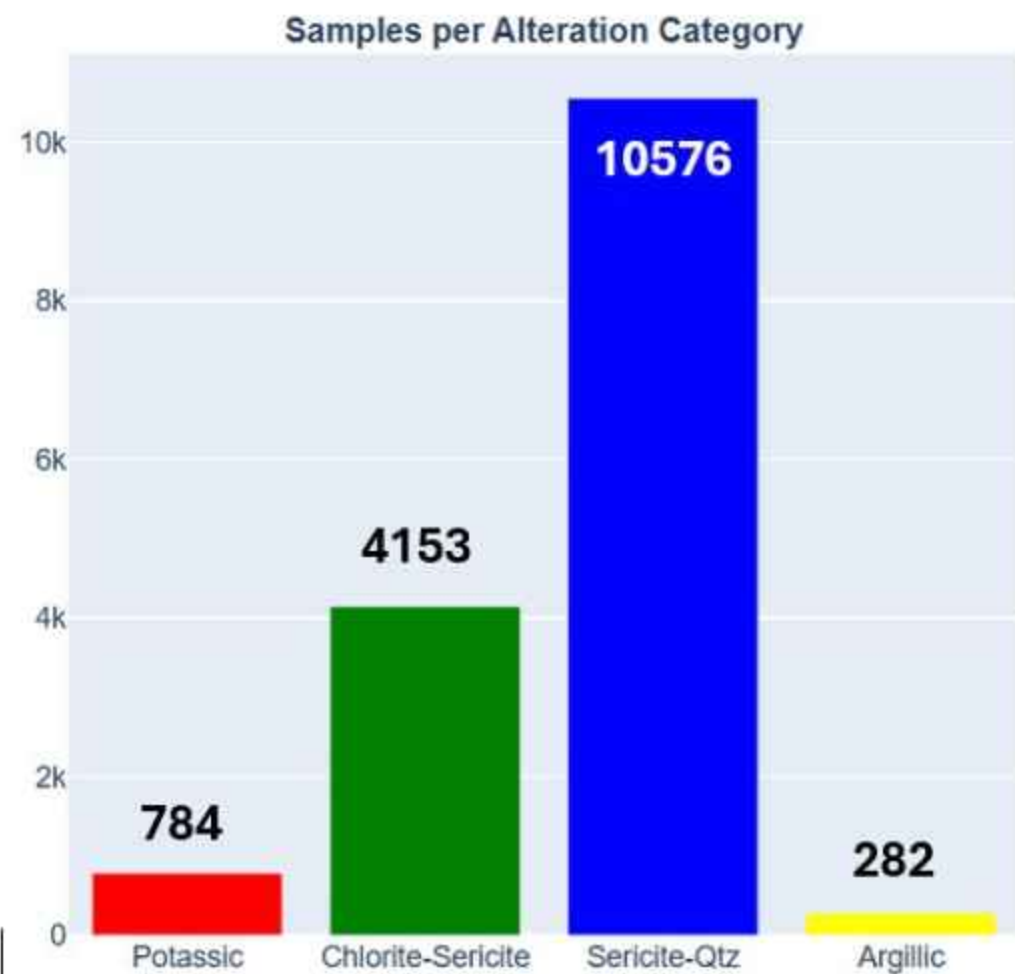
Preliminary Geochemical based Design Space of 180 Runs



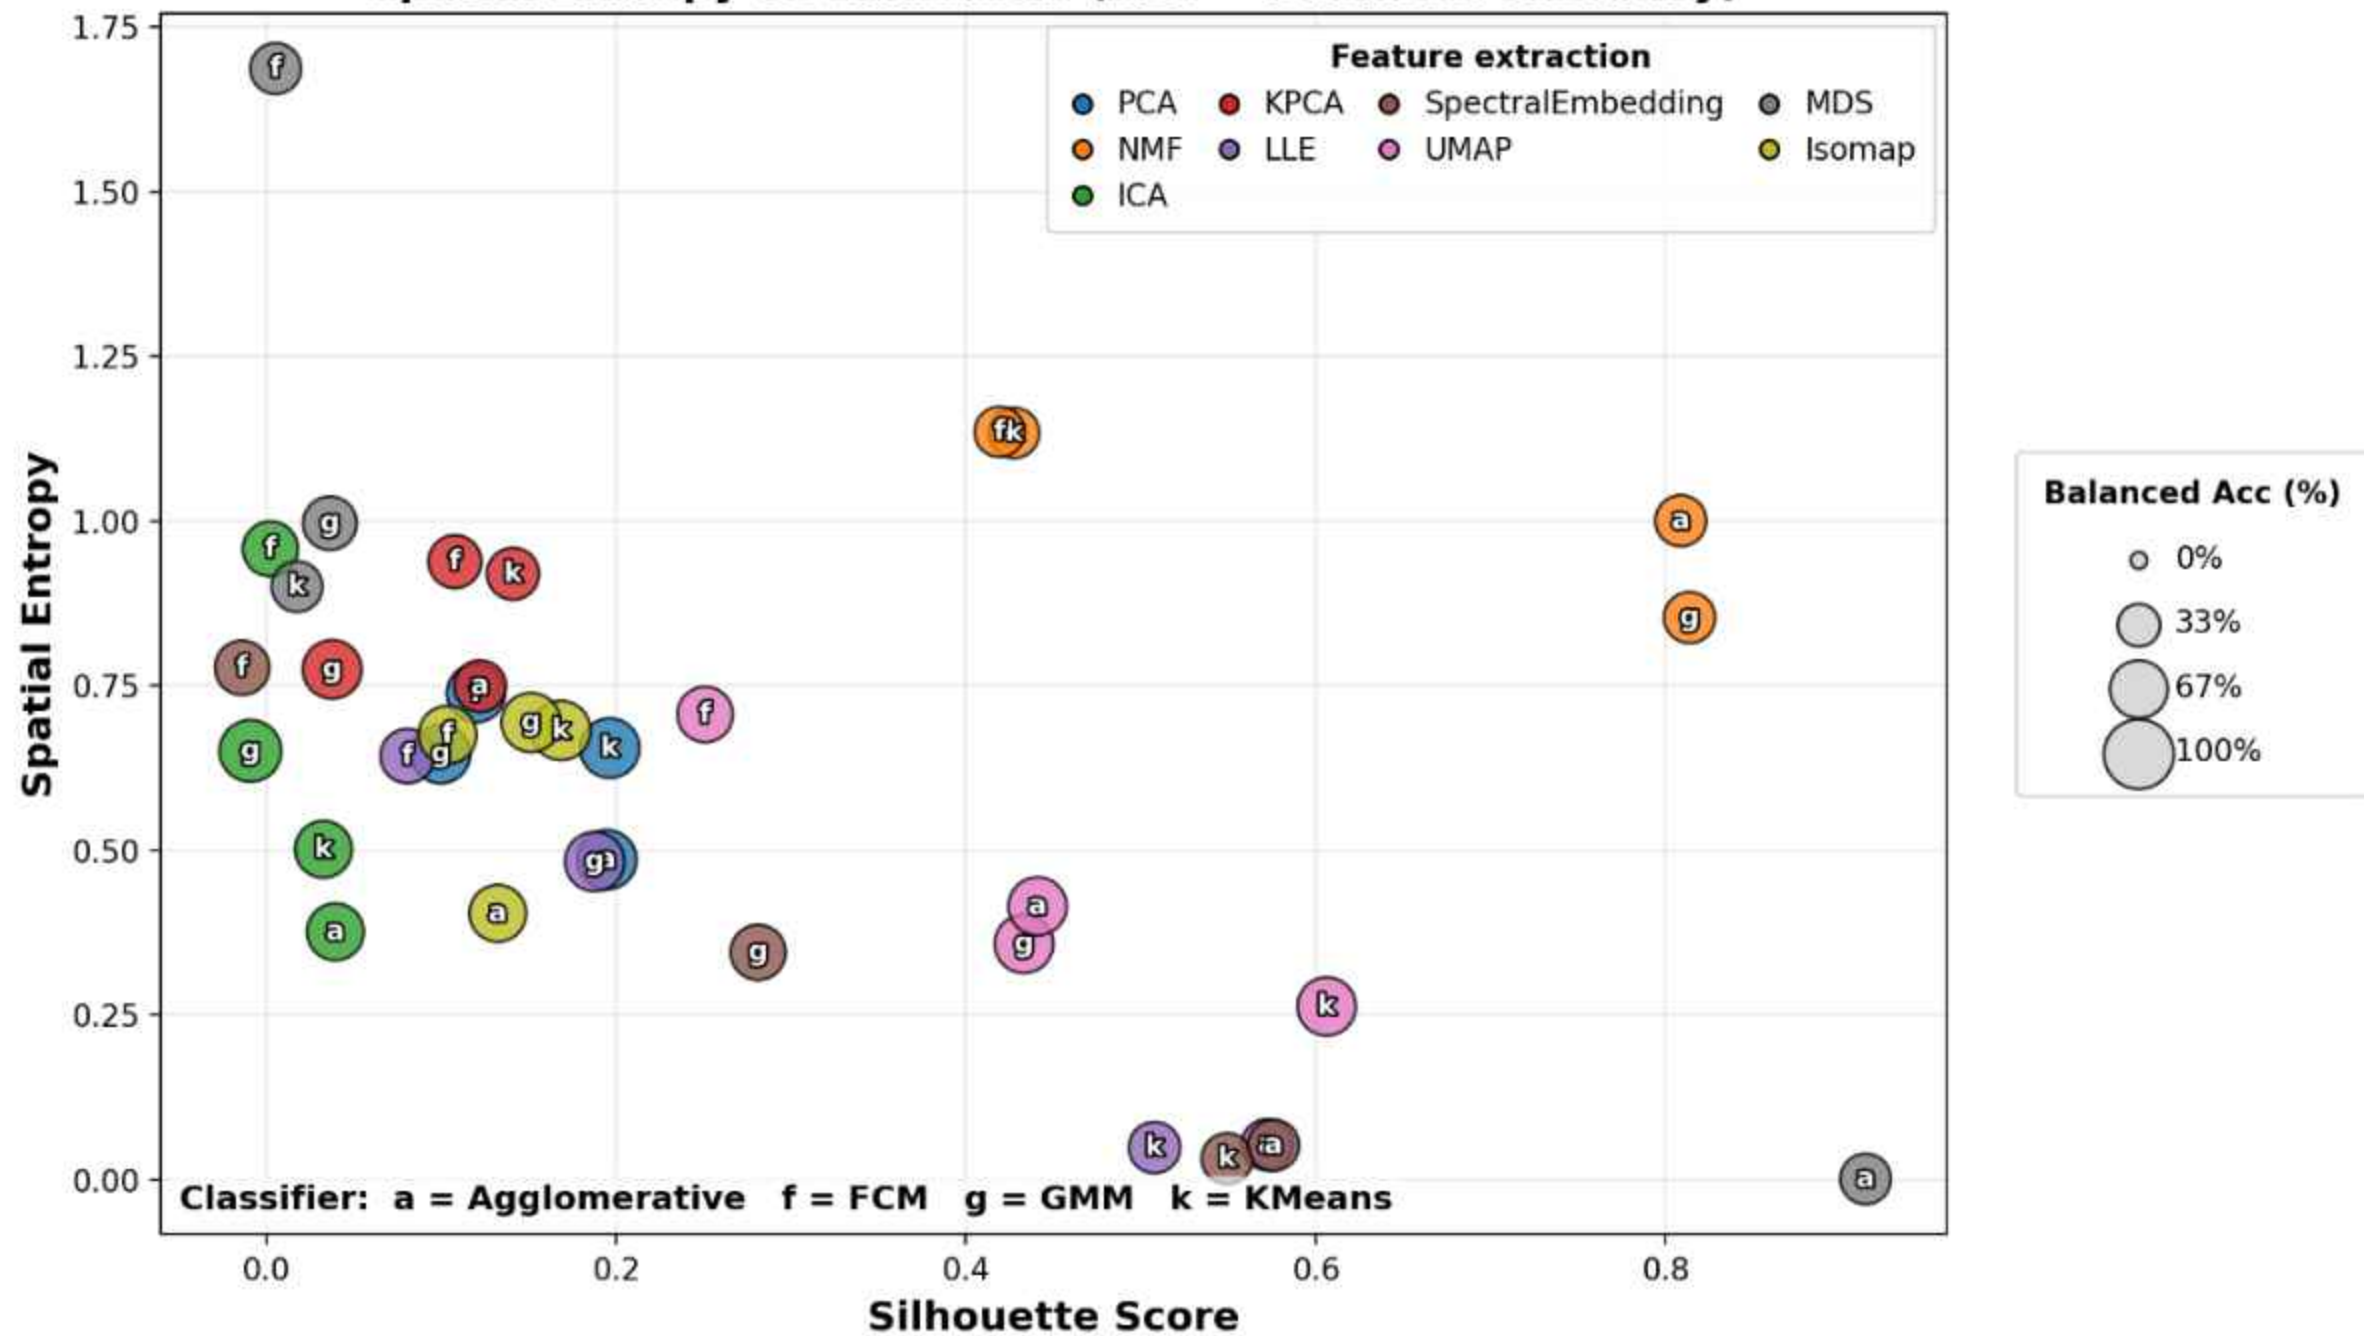
Unsupervised Approach for Spatially Contagious Domaining

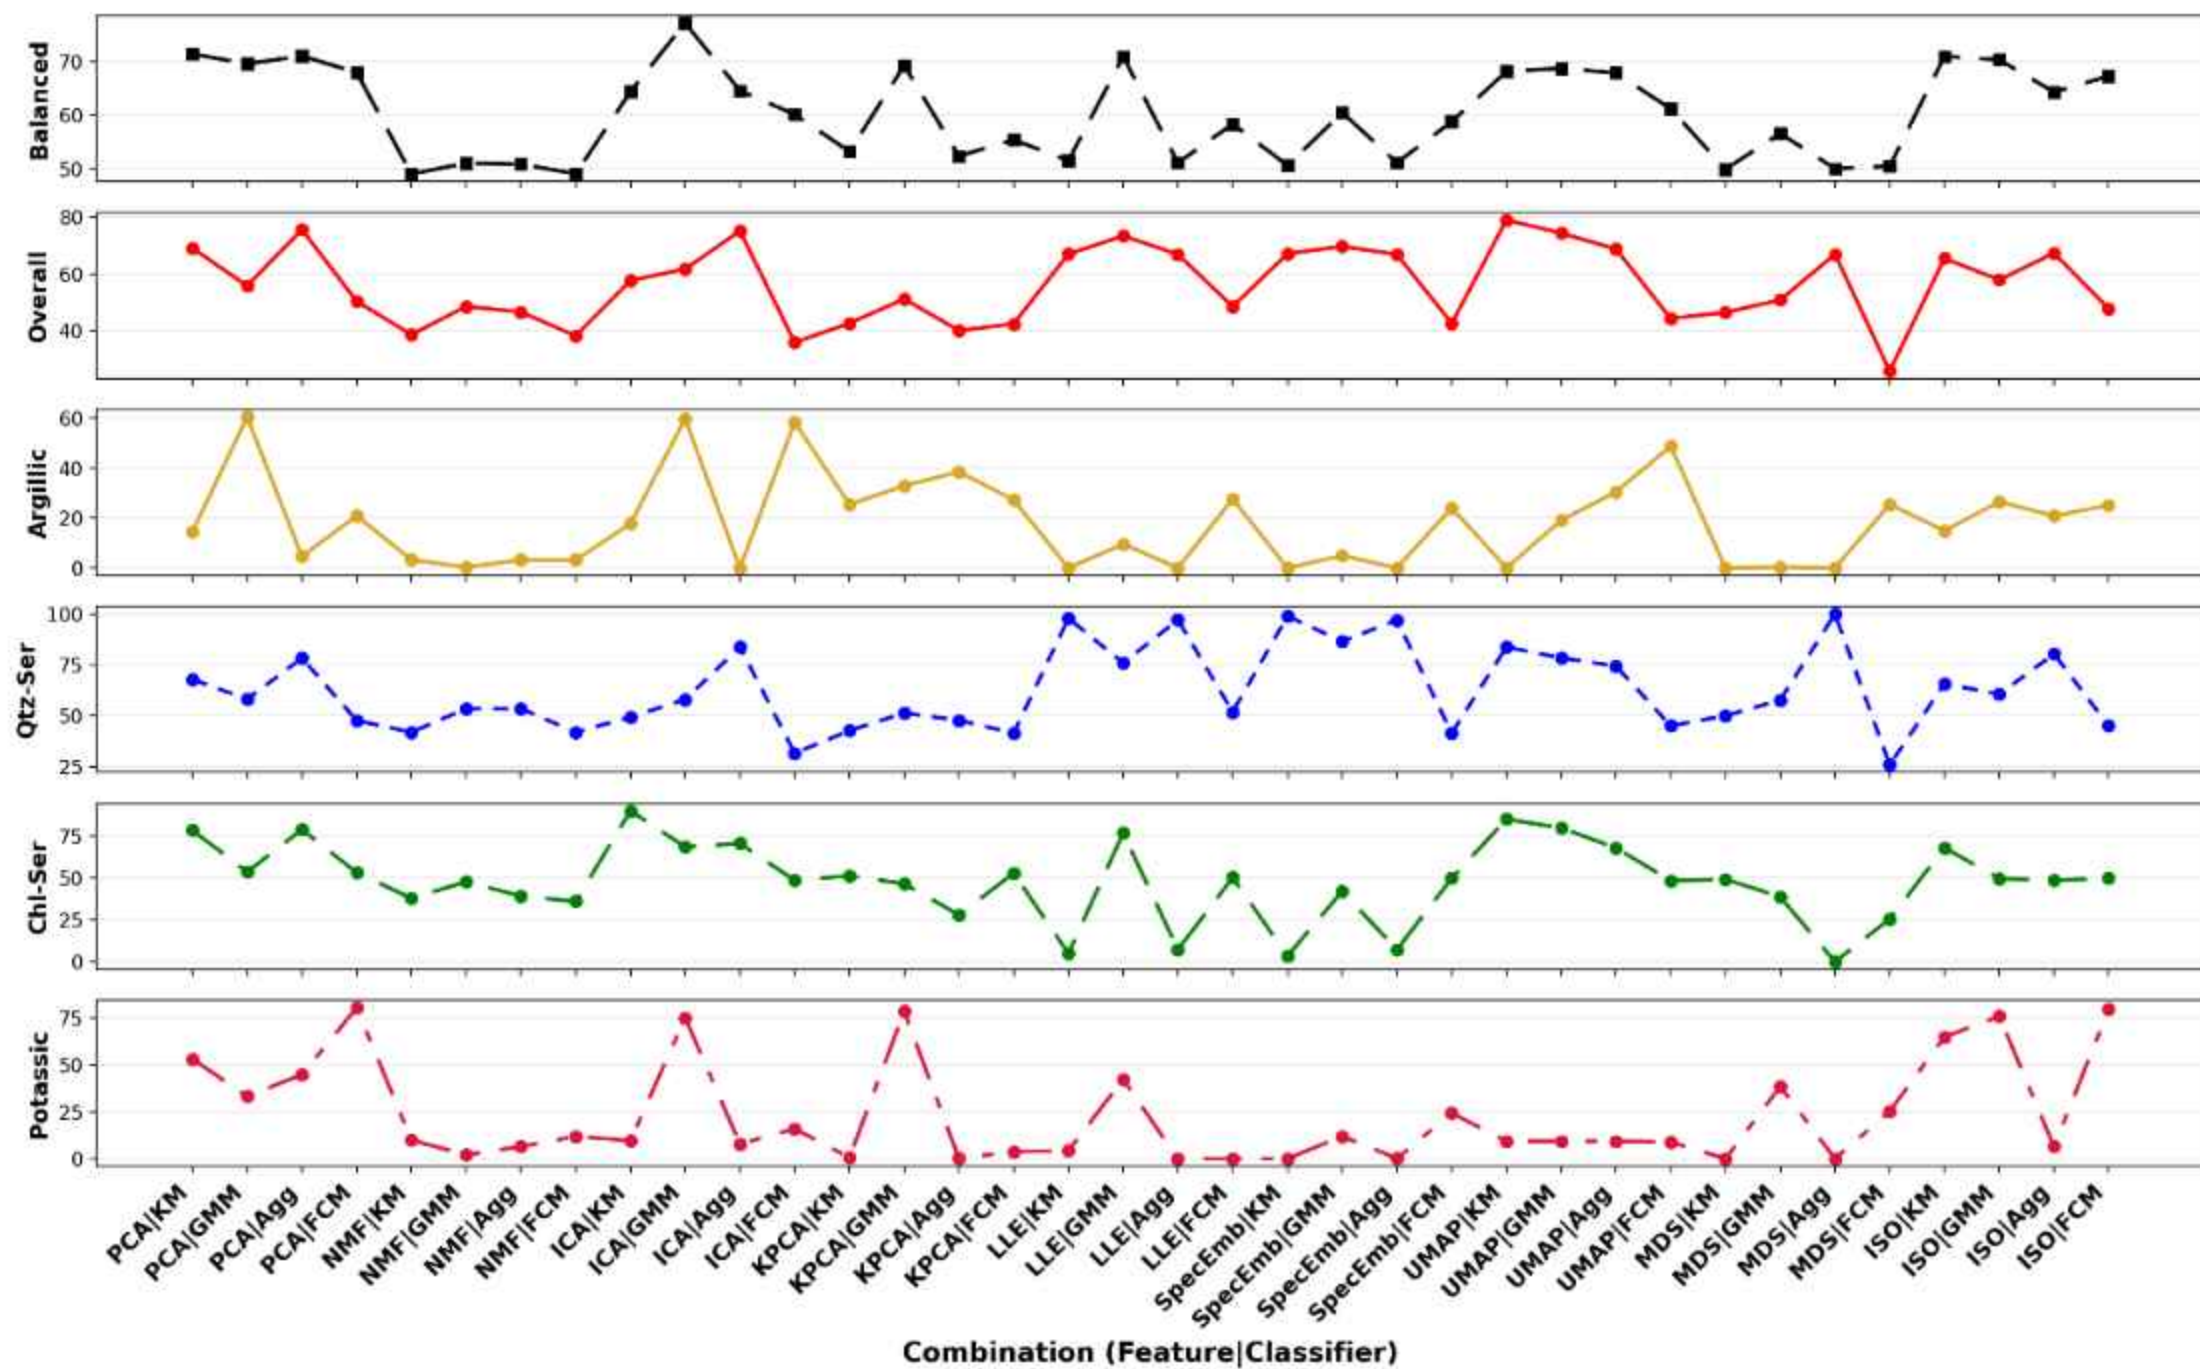


Alteration Category Sample Counts



Spatial Entropy vs Silhouette (size = Balanced Accuracy)





Cleaning based on Spatial Entropy

3-D Region Editor (Hull or Ellipsoid)

Domain: Domain 1 — Sericite-Qtz

Geometry mode:
☐ Hull ☒ Ellipsoid

Translate (m): X / Y / Z

Rotate (deg): Rx / Ry / Rz

Hull scale: Sx / Sy / Sz

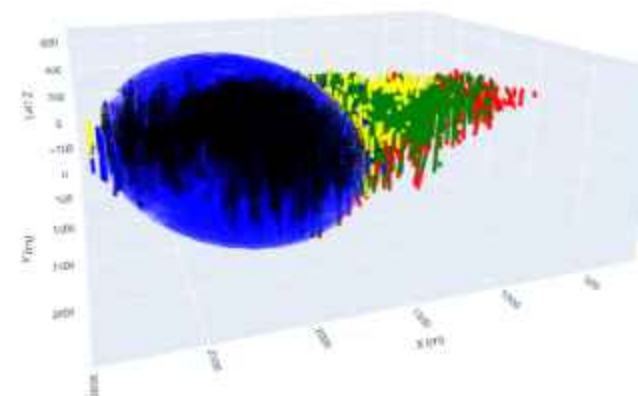
Ellipsoid axis scales: Ex / Ey / Ez

Accept for this domain

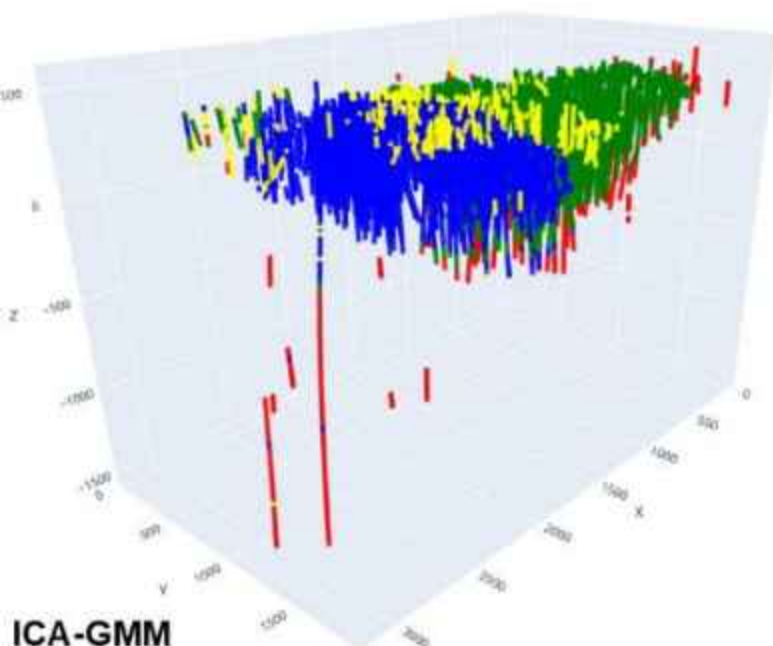
Finish & Continue

Domain 1 — Sericite-Qtz | Inside ellipsoid: 9964 | mean entropy: 0.9712

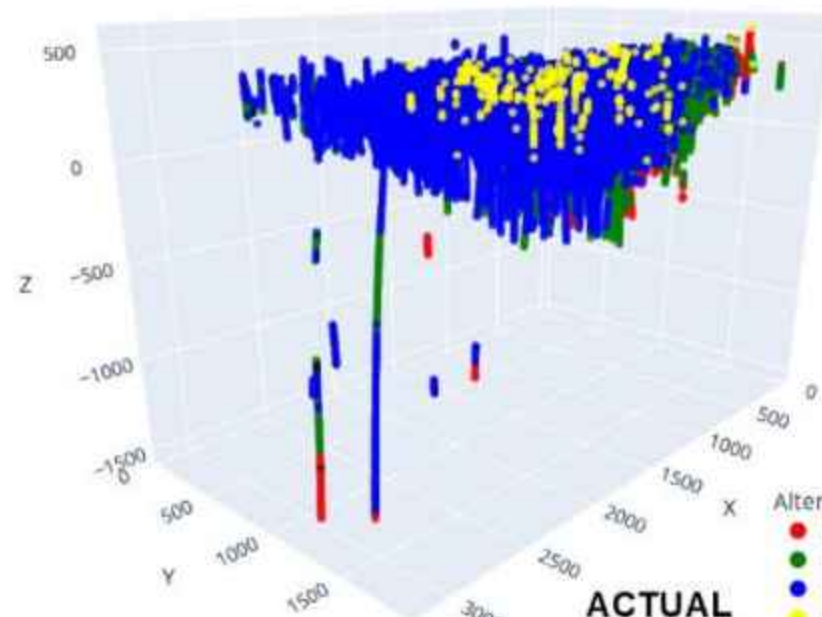
Edit region for Domain 1 — Sericite-Qtz



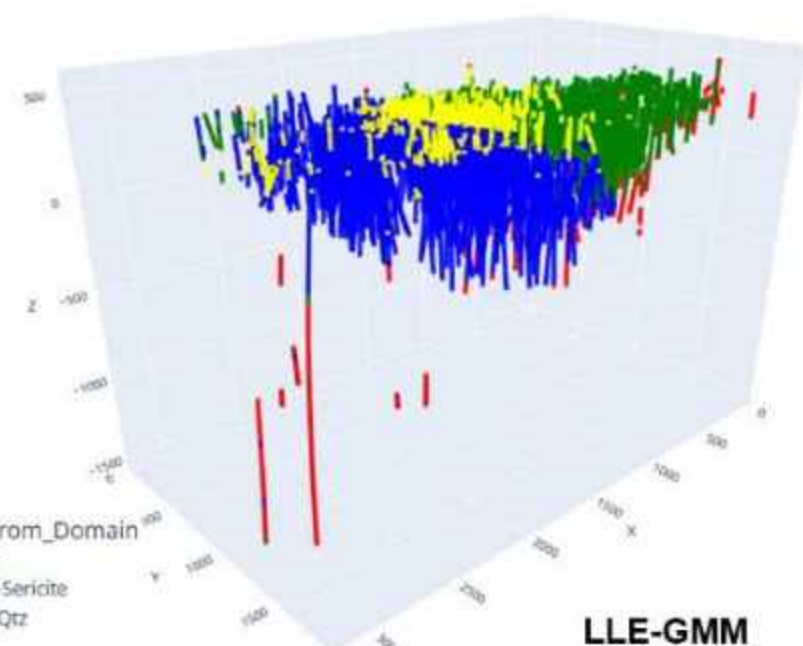
Final Results Comparison



ICA-GMM



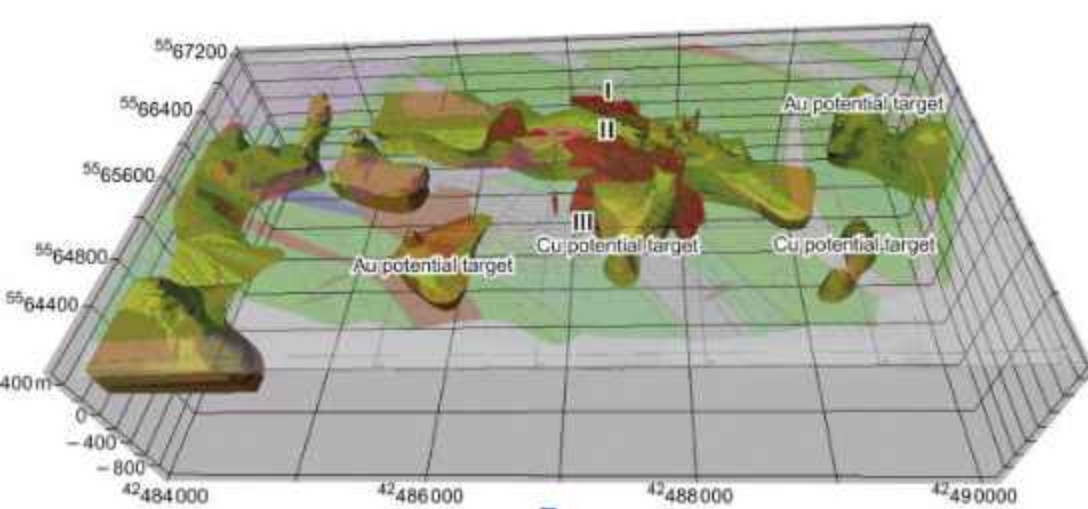
ACTUAL



LLE-GMM

Alteration_from_Domain

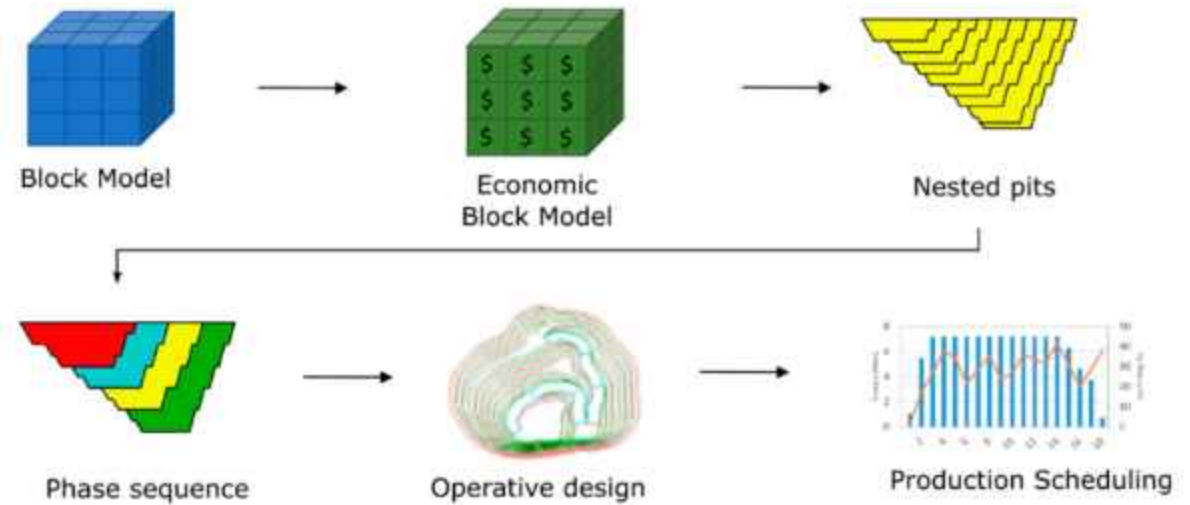
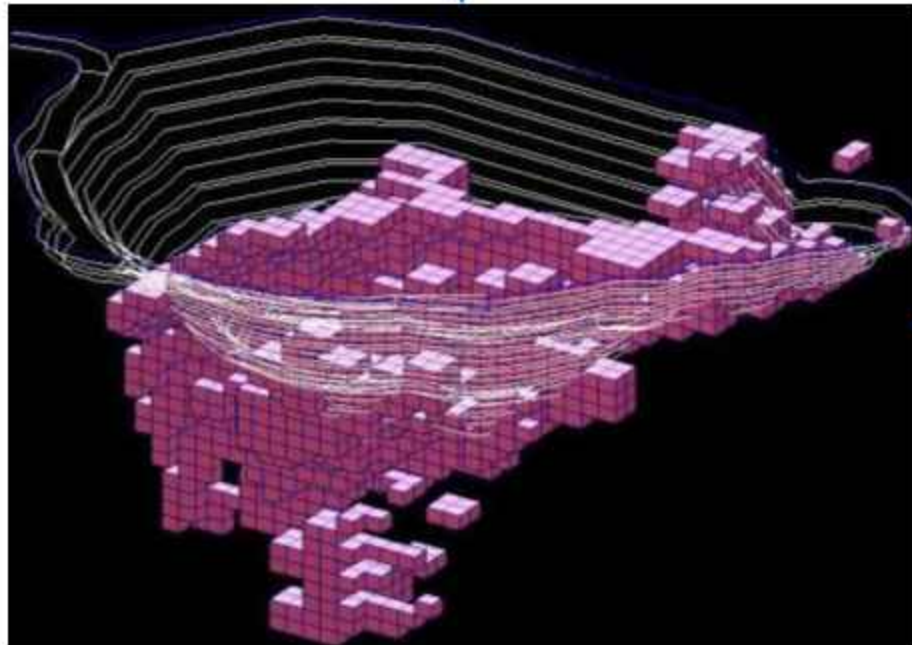
- Potassic
- Chlorite-Sericite
- Sericite-Qtz
- Argillic



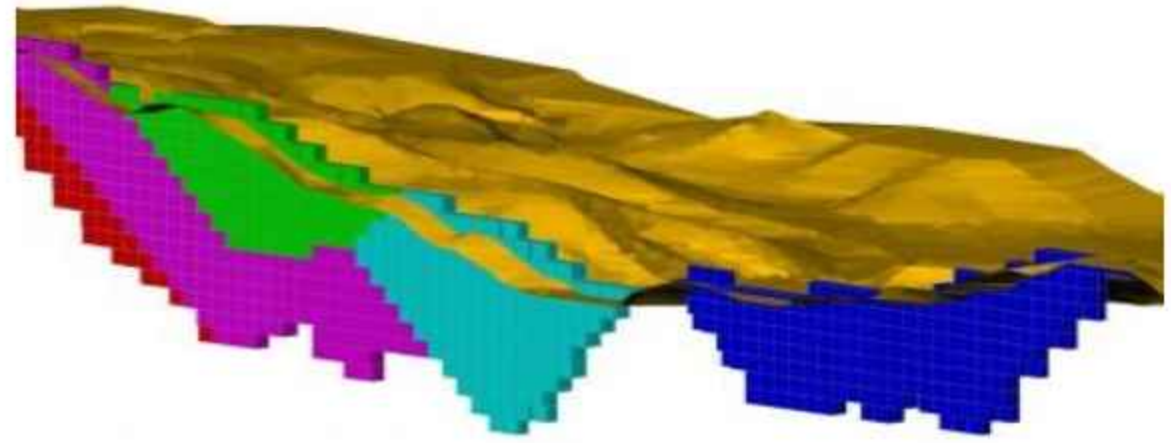
Production Scheduling of Mining operations

Optimization problem

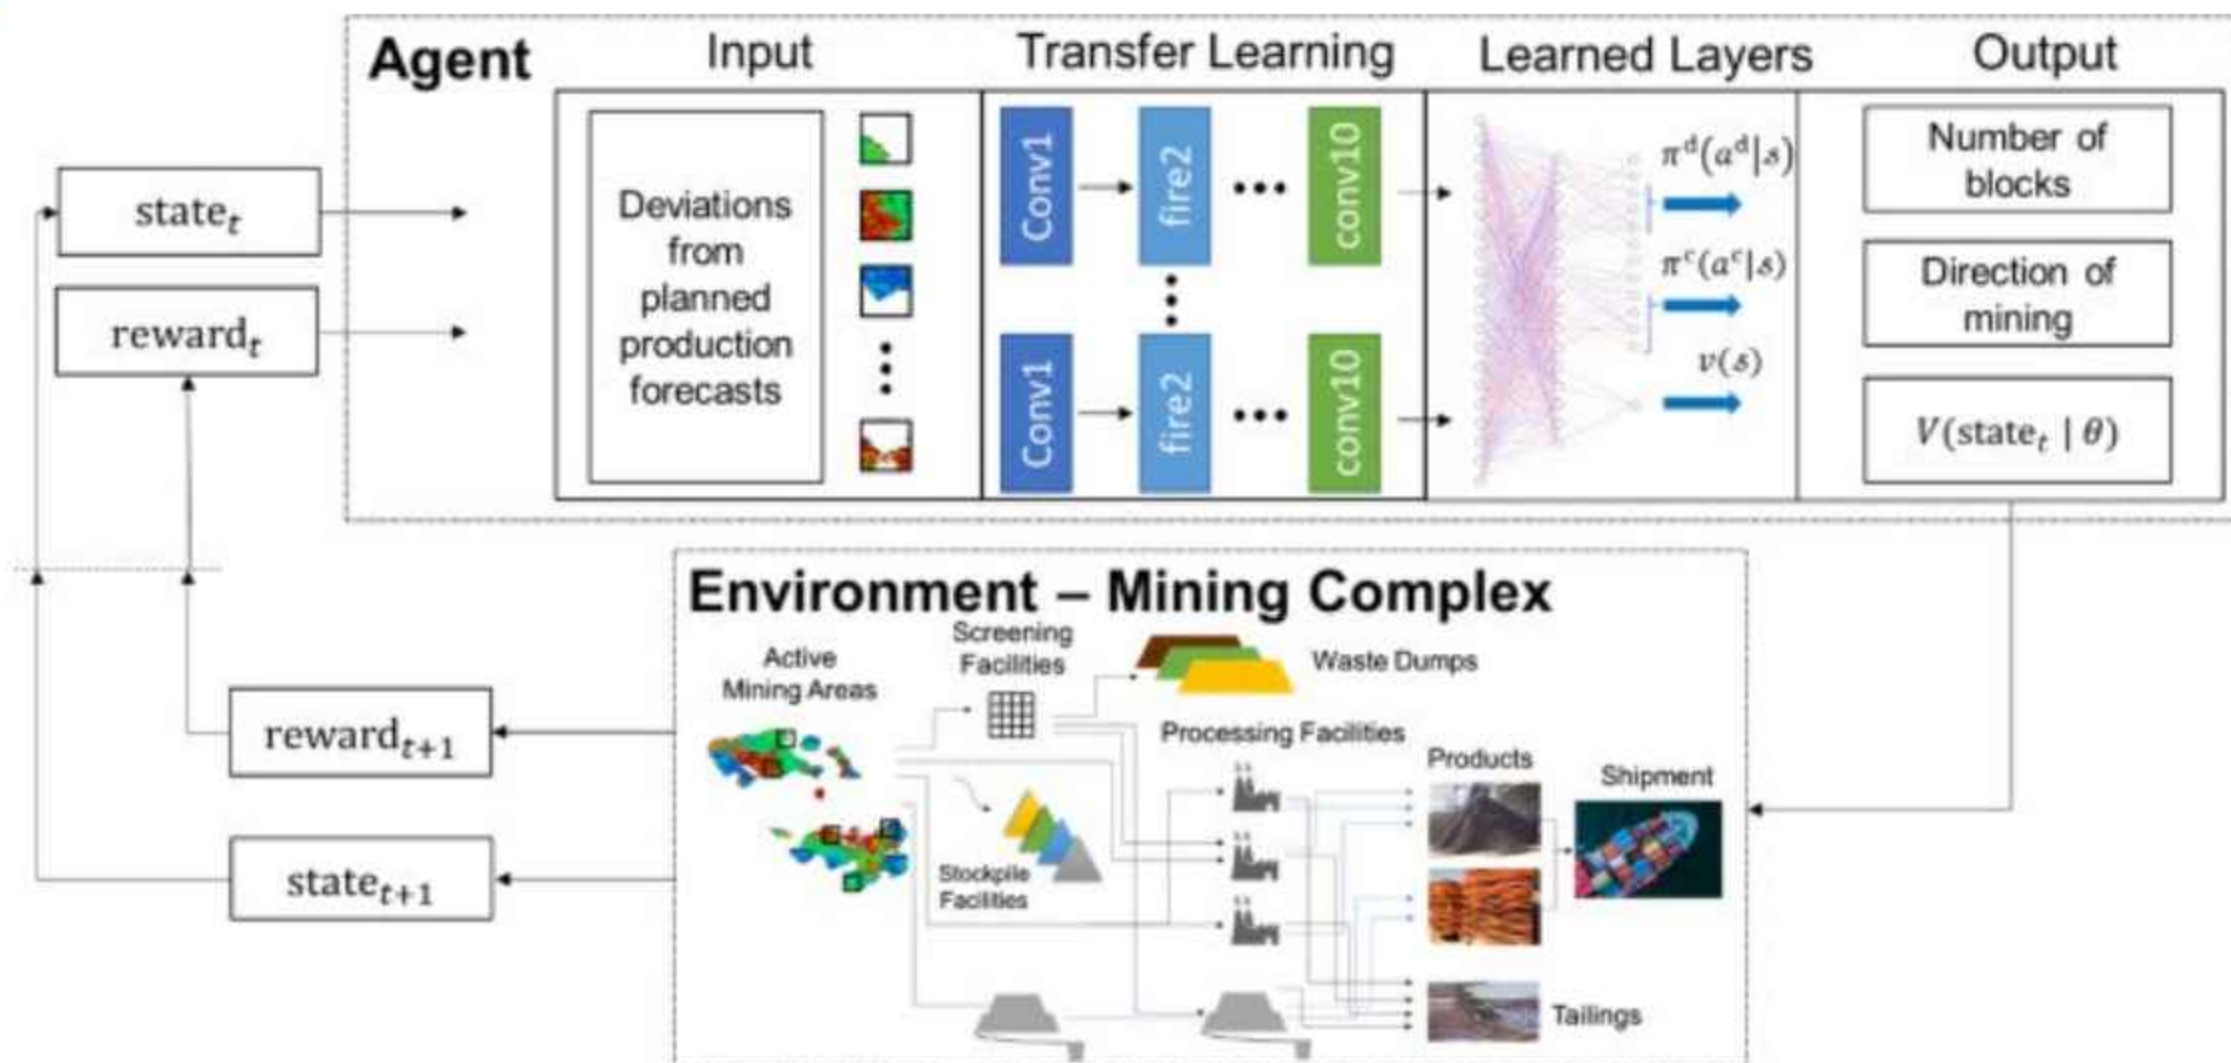
Prescriptive analytics



Mine Planning and AI
Prescriptive analytics
Extraction sequence of Blocks
Current work



Reinforced Learning of Mining Complex – current work



Conclusions

- Statistical Learning is key to sustainable resource Management
- Clean samples will lead to reliable statistical models for sustainable resource
- Feature Extraction and Classification works
- Realtime Reinforced Learning for Optimization of Mine Complex is the future challenge
- Noise free sensors at plant can enhance better learning strategies

THANK YOU



University
of Exeter



Improving Machine Learning Predictions of High CO₂ Climates

Dr Stephen Thomson

Dr Daniel Williams

Department of Mathematics and Statistics, University of Exeter

Predicting high CO₂ climates using machine learning



Predicting high CO₂ climates using machine learning

- Weather prediction systems using ML are becoming more common and more powerful

Predicting high CO₂ climates using machine learning

- Weather prediction systems using ML are becoming more common and more powerful
- Typically trained on ERA5 reanalysis data

Predicting high CO₂ climates using machine learning

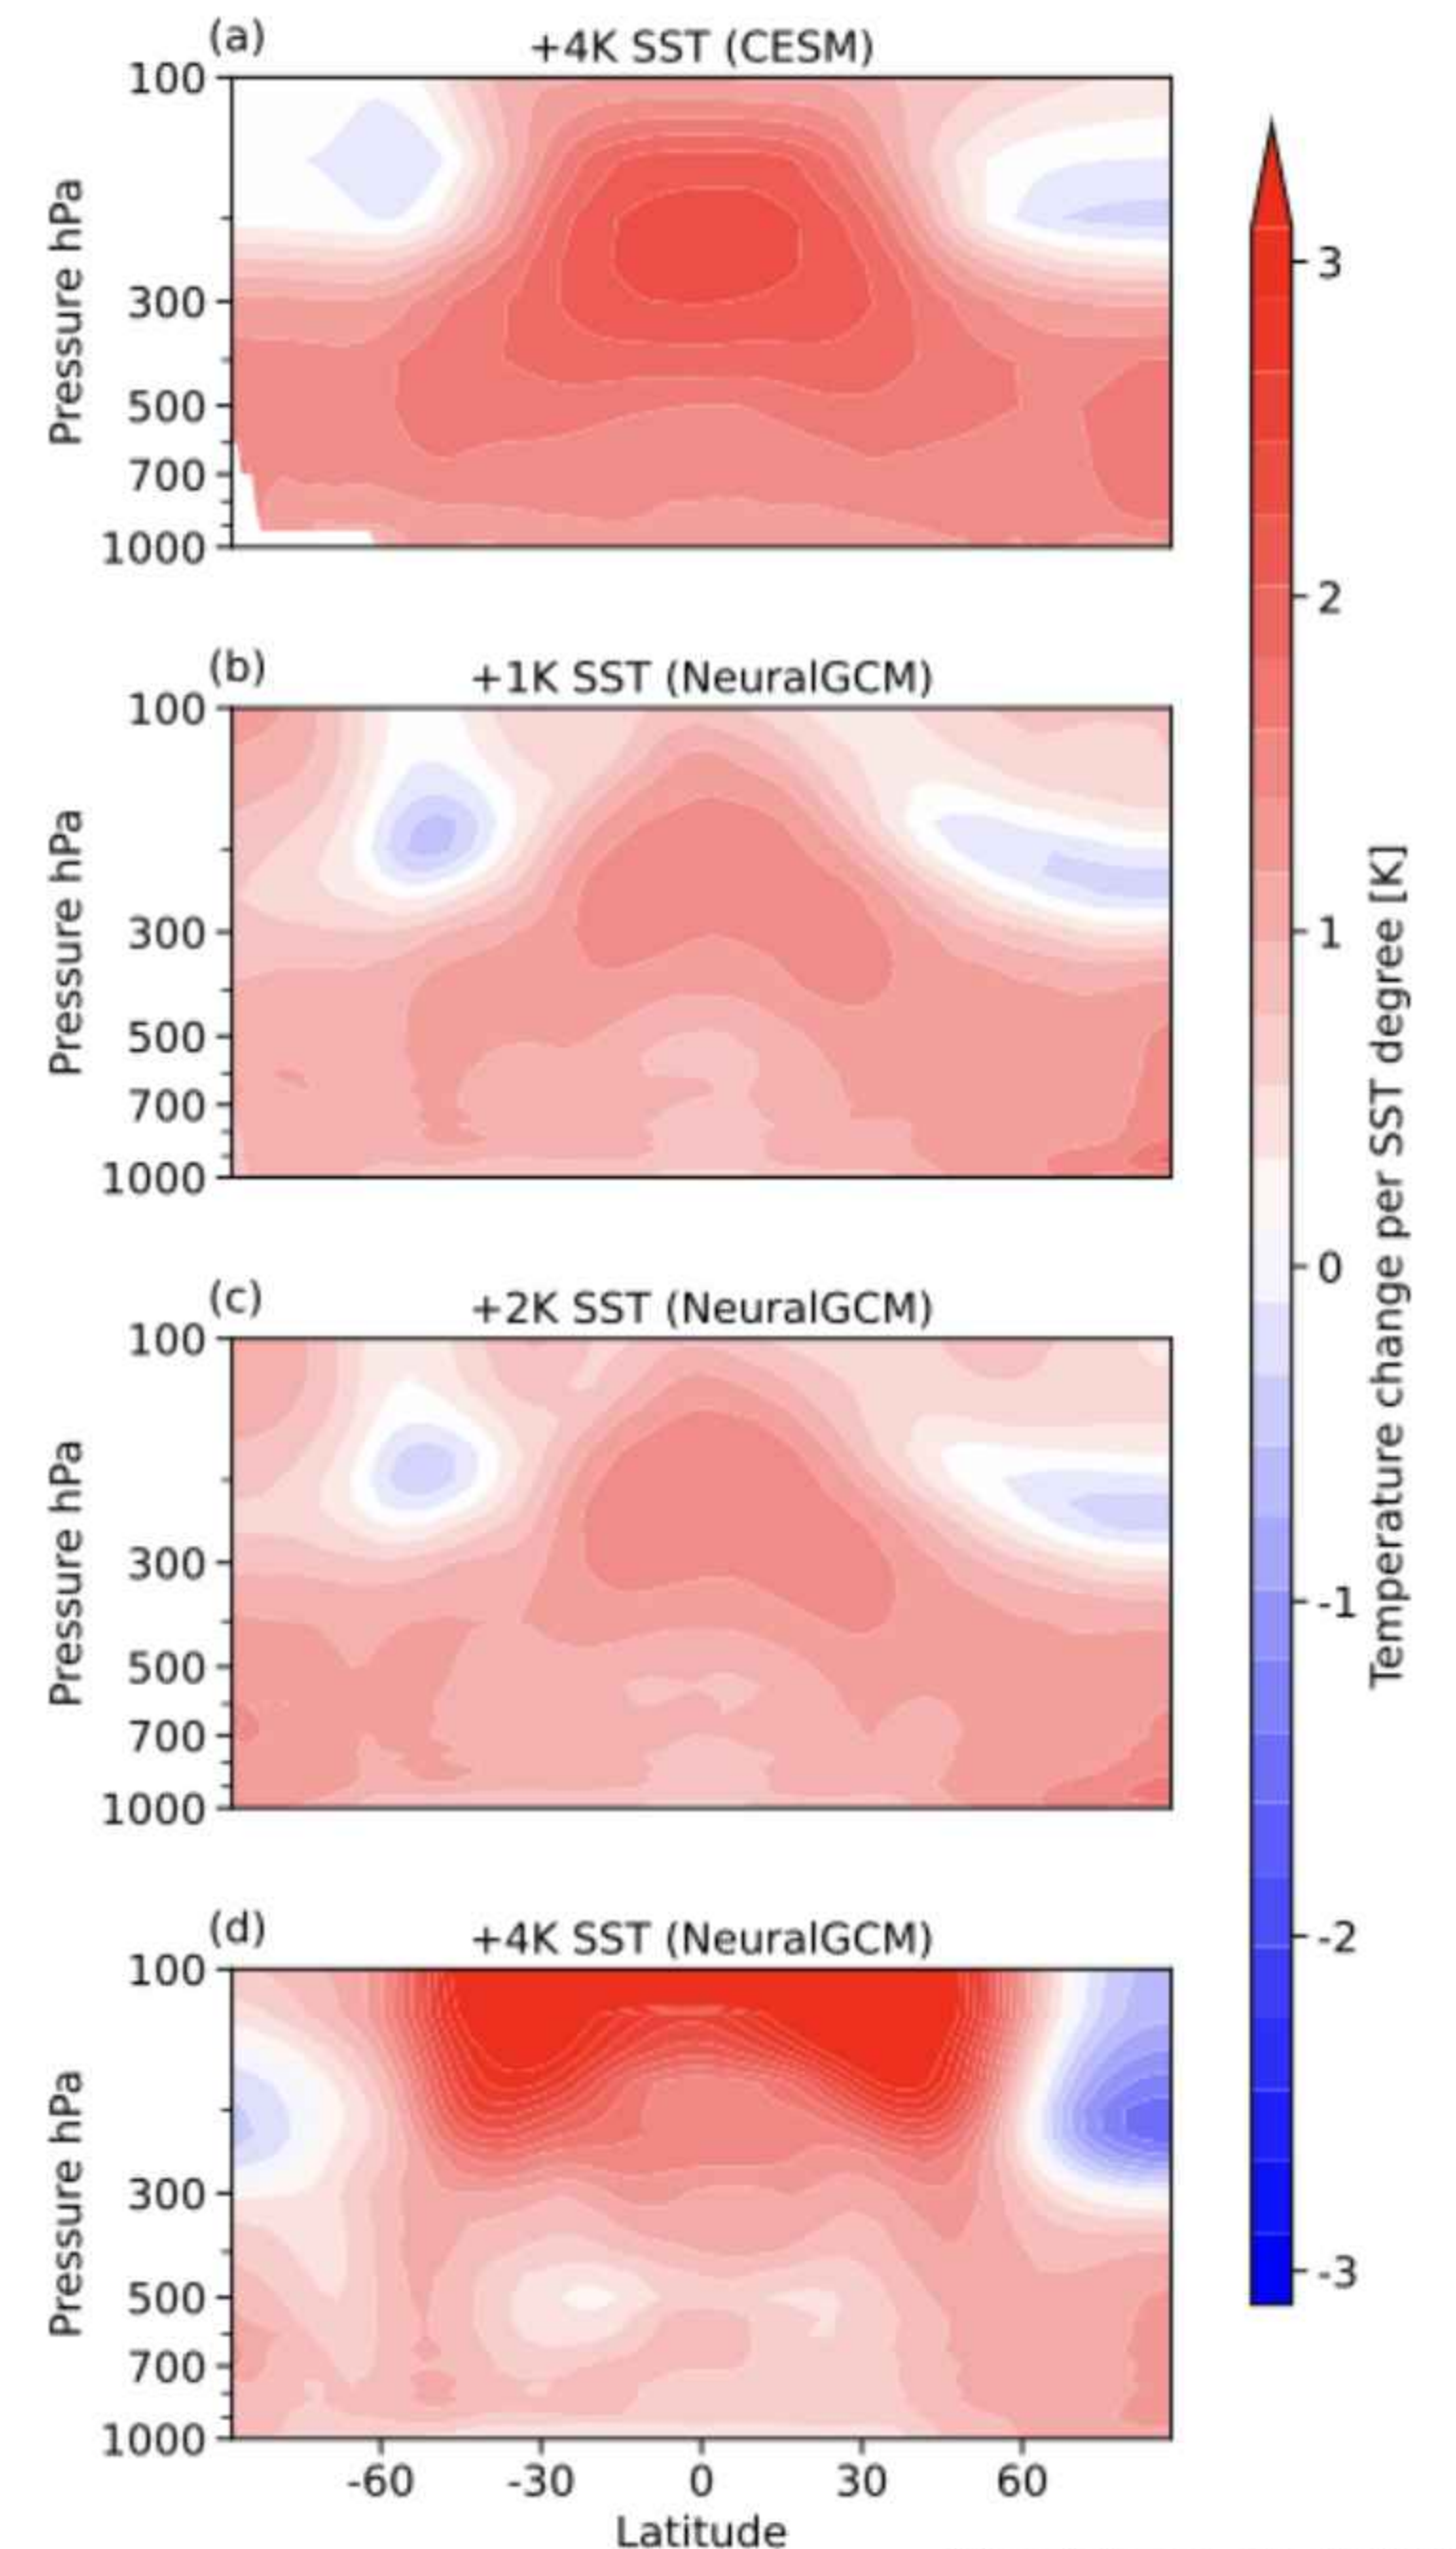
- Weather prediction systems using ML are becoming more common and more powerful
- Typically trained on ERA5 reanalysis data
- **Climate change is an out-of-sample test**

Predicting high CO₂ climates using machine learning

- Weather prediction systems using ML are becoming more common and more powerful
- Typically trained on ERA5 reanalysis data
- **Climate change is an out-of-sample test**
- **Neural-GCM** (Google) is a 'hybrid' ML model - **dynamics** with learned physics, and uses prescribed SSTs

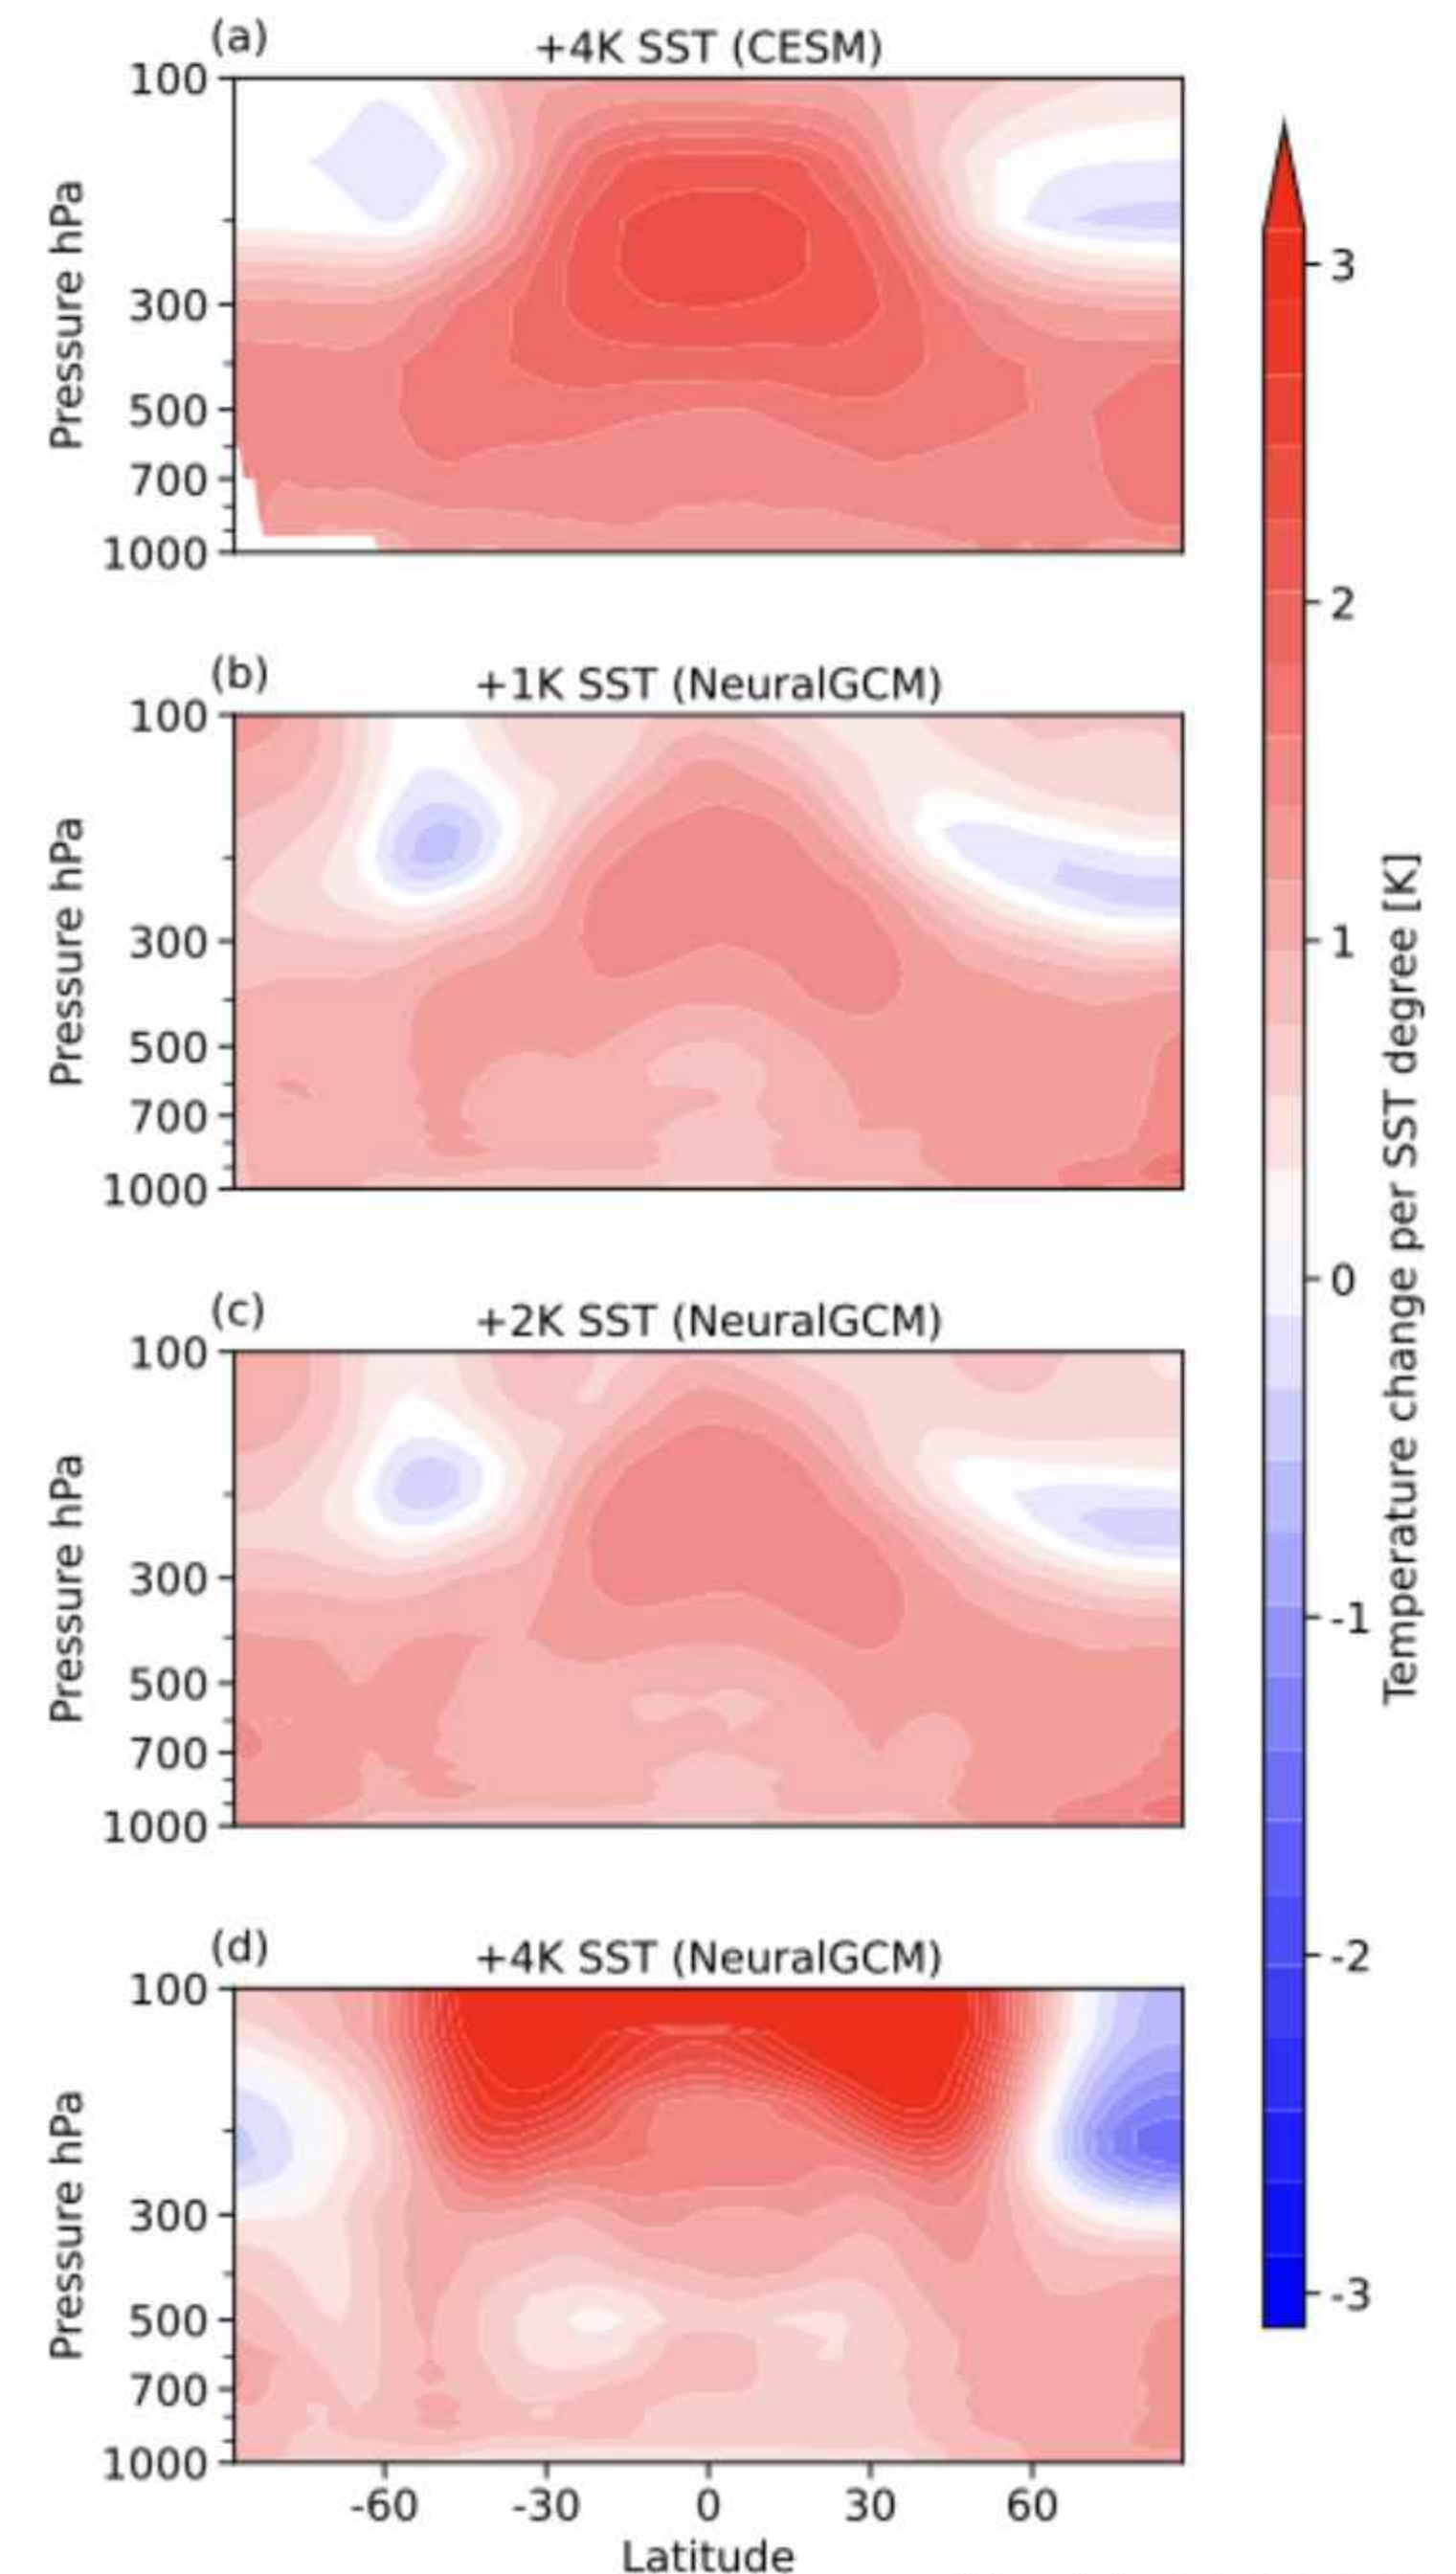
Predicting high CO₂ climates using machine learning

- Weather prediction systems using ML are becoming more common and more powerful
- Typically trained on ERA5 reanalysis data
- **Climate change is an out-of-sample test**
- **Neural-GCM** (Google) is a ‘hybrid’ ML model - **dynamics** with learned physics, and uses prescribed SSTs
- Neural-GCM’s response to an AMIP-4K warming experiment is
 - Not proportional to SST increase
 - Very different to a physical model



Predicting high CO₂ climates using machine learning

- Weather prediction systems using ML are becoming more common and more powerful
- Typically trained on ERA5 reanalysis data
- **Climate change is an out-of-sample test**
- **Neural-GCM** (Google) is a ‘hybrid’ ML model - **dynamics** with learned physics, and uses prescribed SSTs
- Neural-GCM’s response to an AMIP-4K warming experiment is
 - Not proportional to SST increase
 - Very different to a physical model
- **Is it possible to incorporate a representation of radiative transfer into ML models?**



Our approach

Our approach

- We have explored approaches for including radiative transfer into a simple ML framework

Our approach

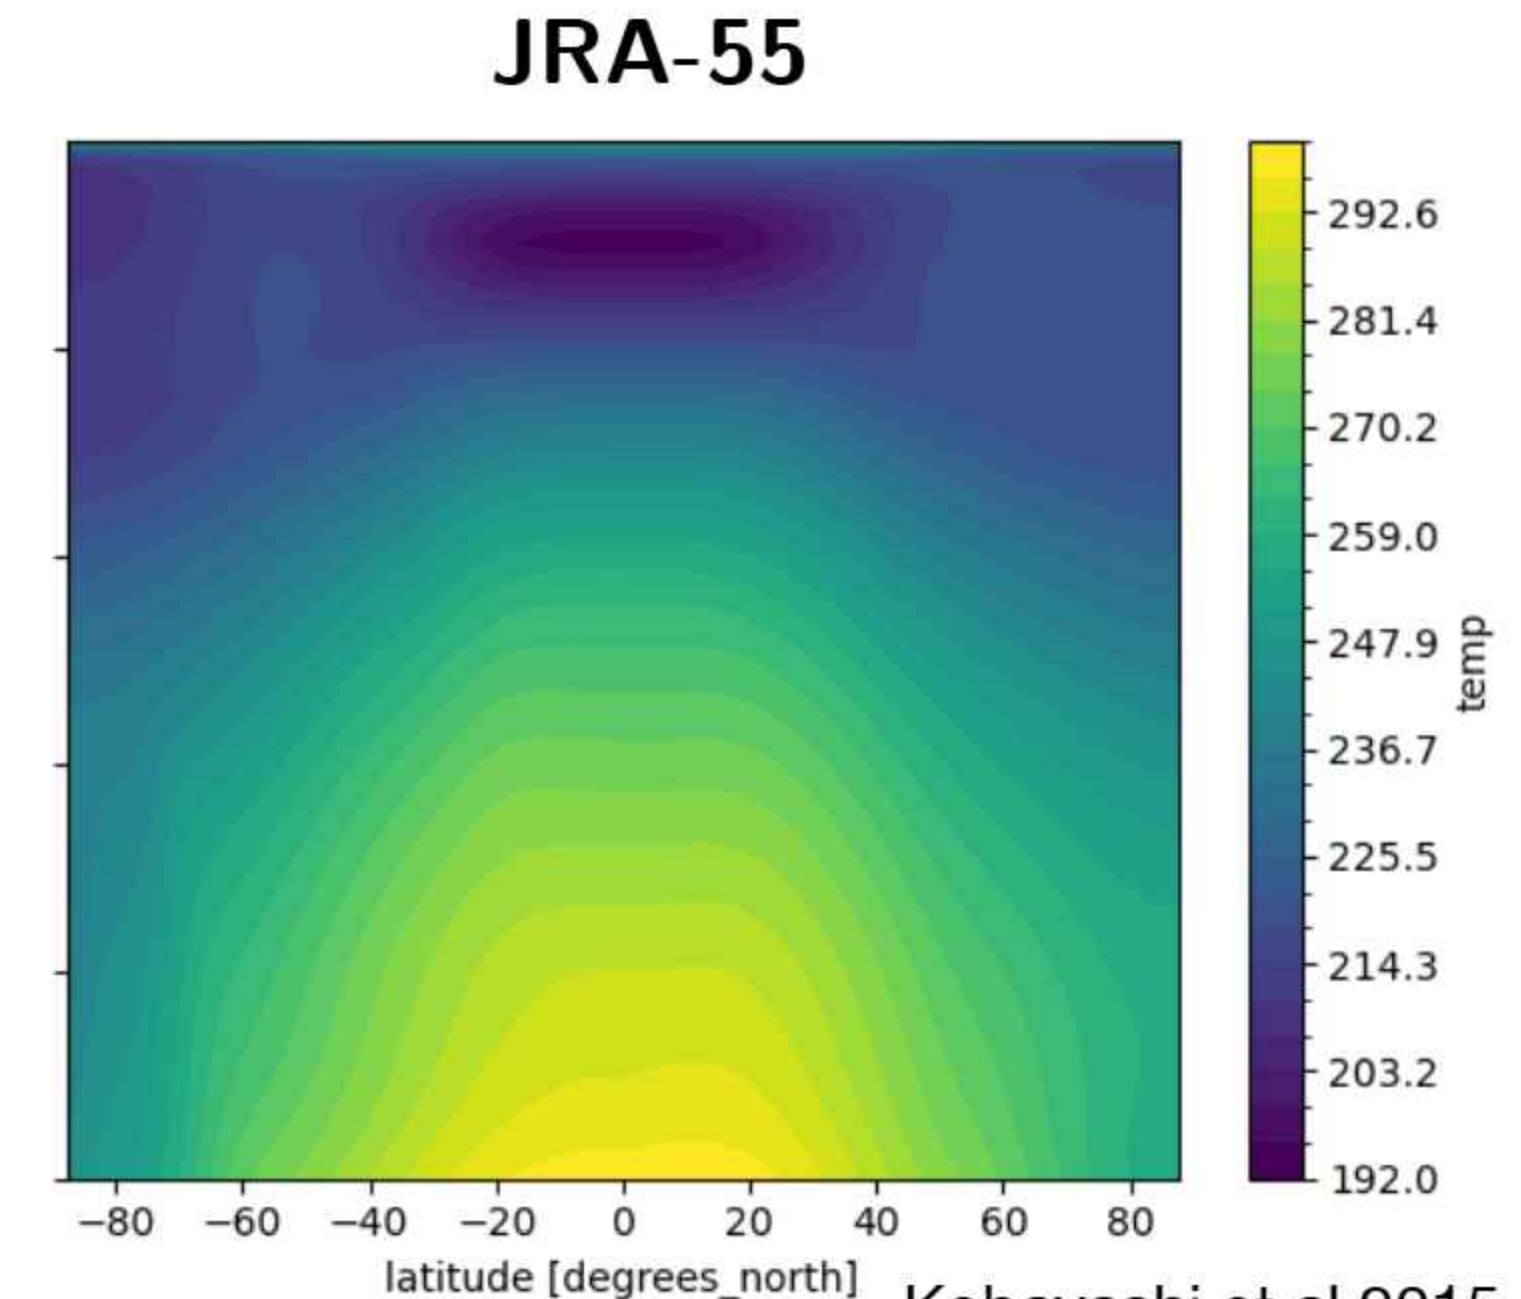
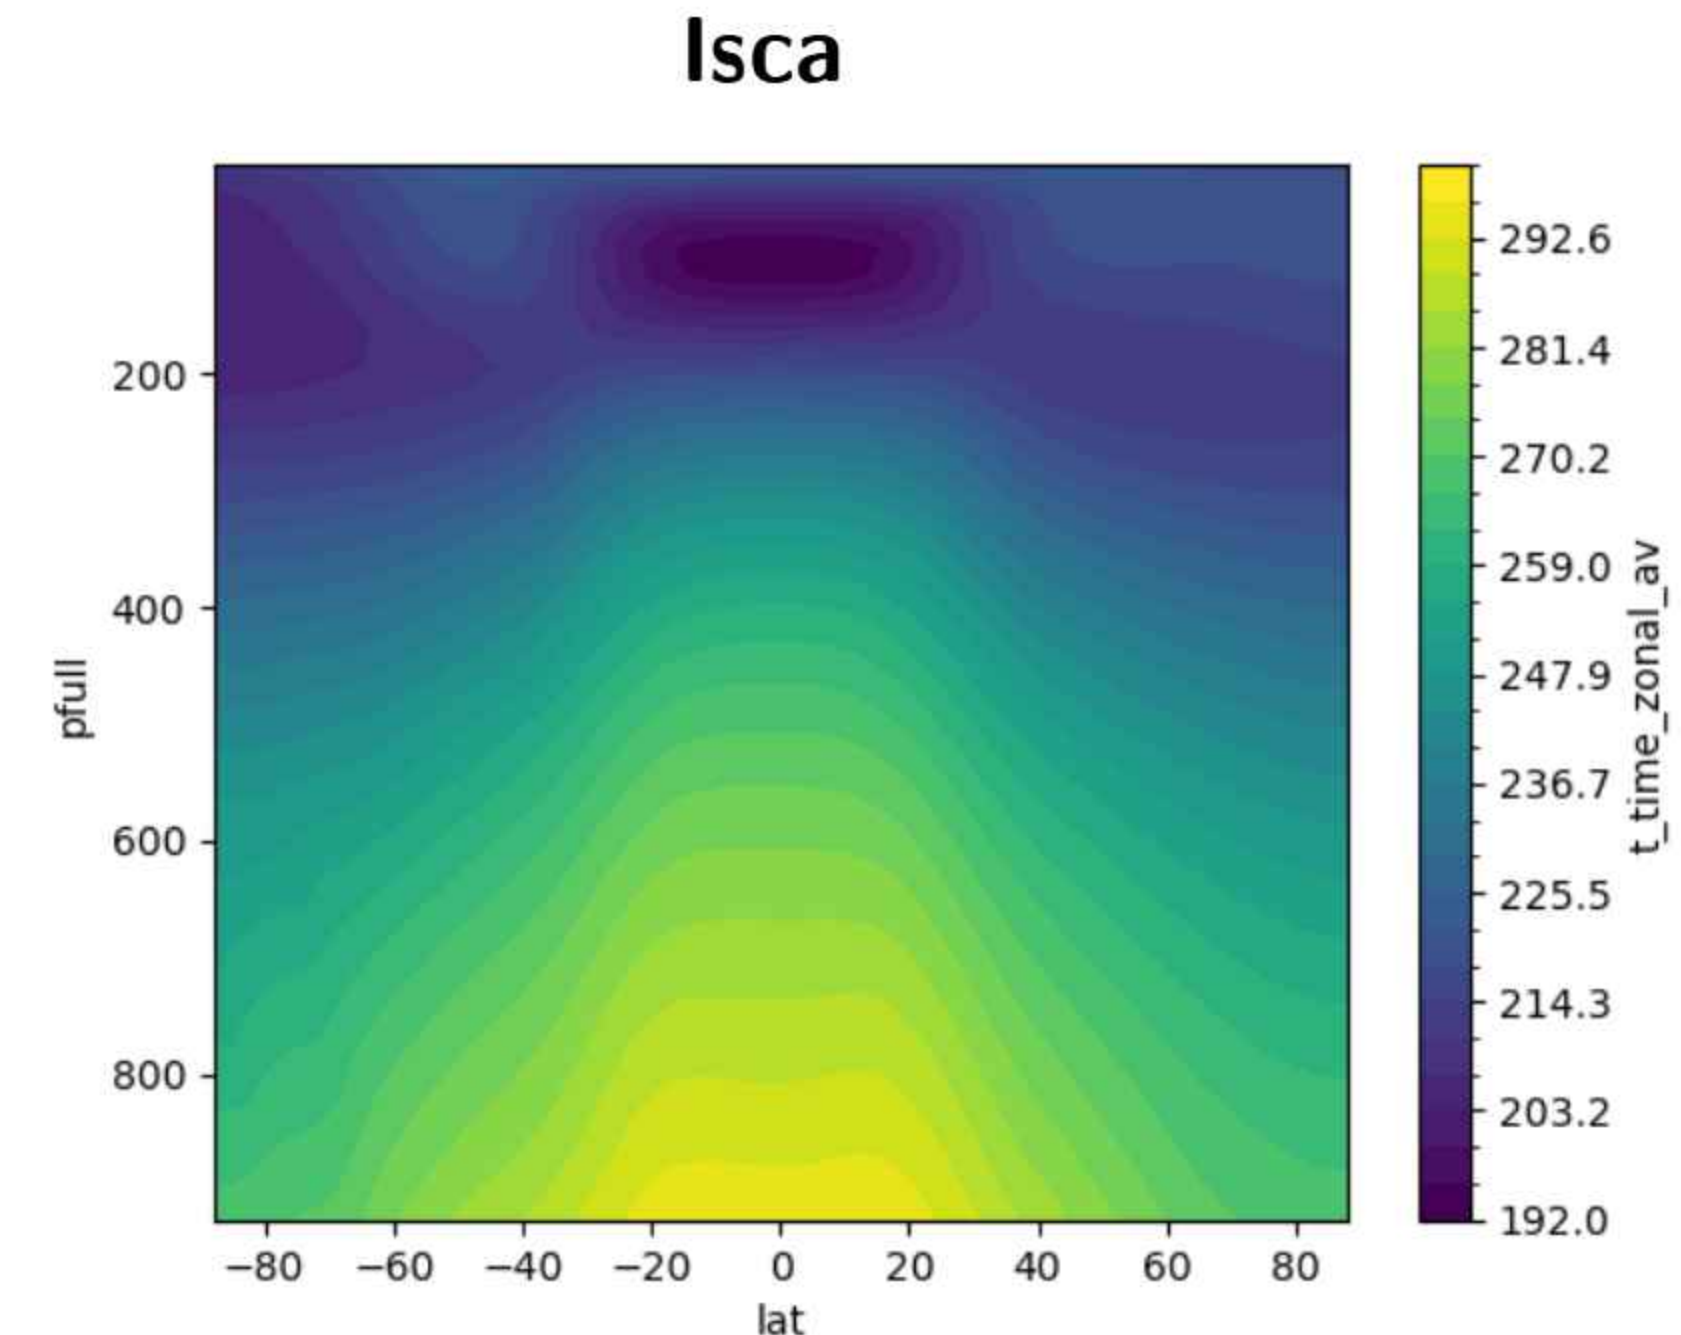
- We have explored approaches for including radiative transfer into a simple ML framework
- We are not using an MLWP system, but have built our own simple ML framework to illustrate the problems and explore ways forward

Our approach

- We have explored approaches for including radiative transfer into a simple ML framework
- We are not using an MLWP system, but have built our own simple ML framework to illustrate the problems and explore ways forward
- **Training data:** we have run 140 different GCM experiments using the Isca modelling framework
 - Earth-like configuration with continents, a mixed-layer ocean, and Socrates radiative transfer without clouds
 - 5 different CO₂ values (278 - 1600ppmv), 7 different rotation rates (2x Earth - 1/16 Earth) and 4 different obliquities (0 - 1.5 x Earth)

Our approach

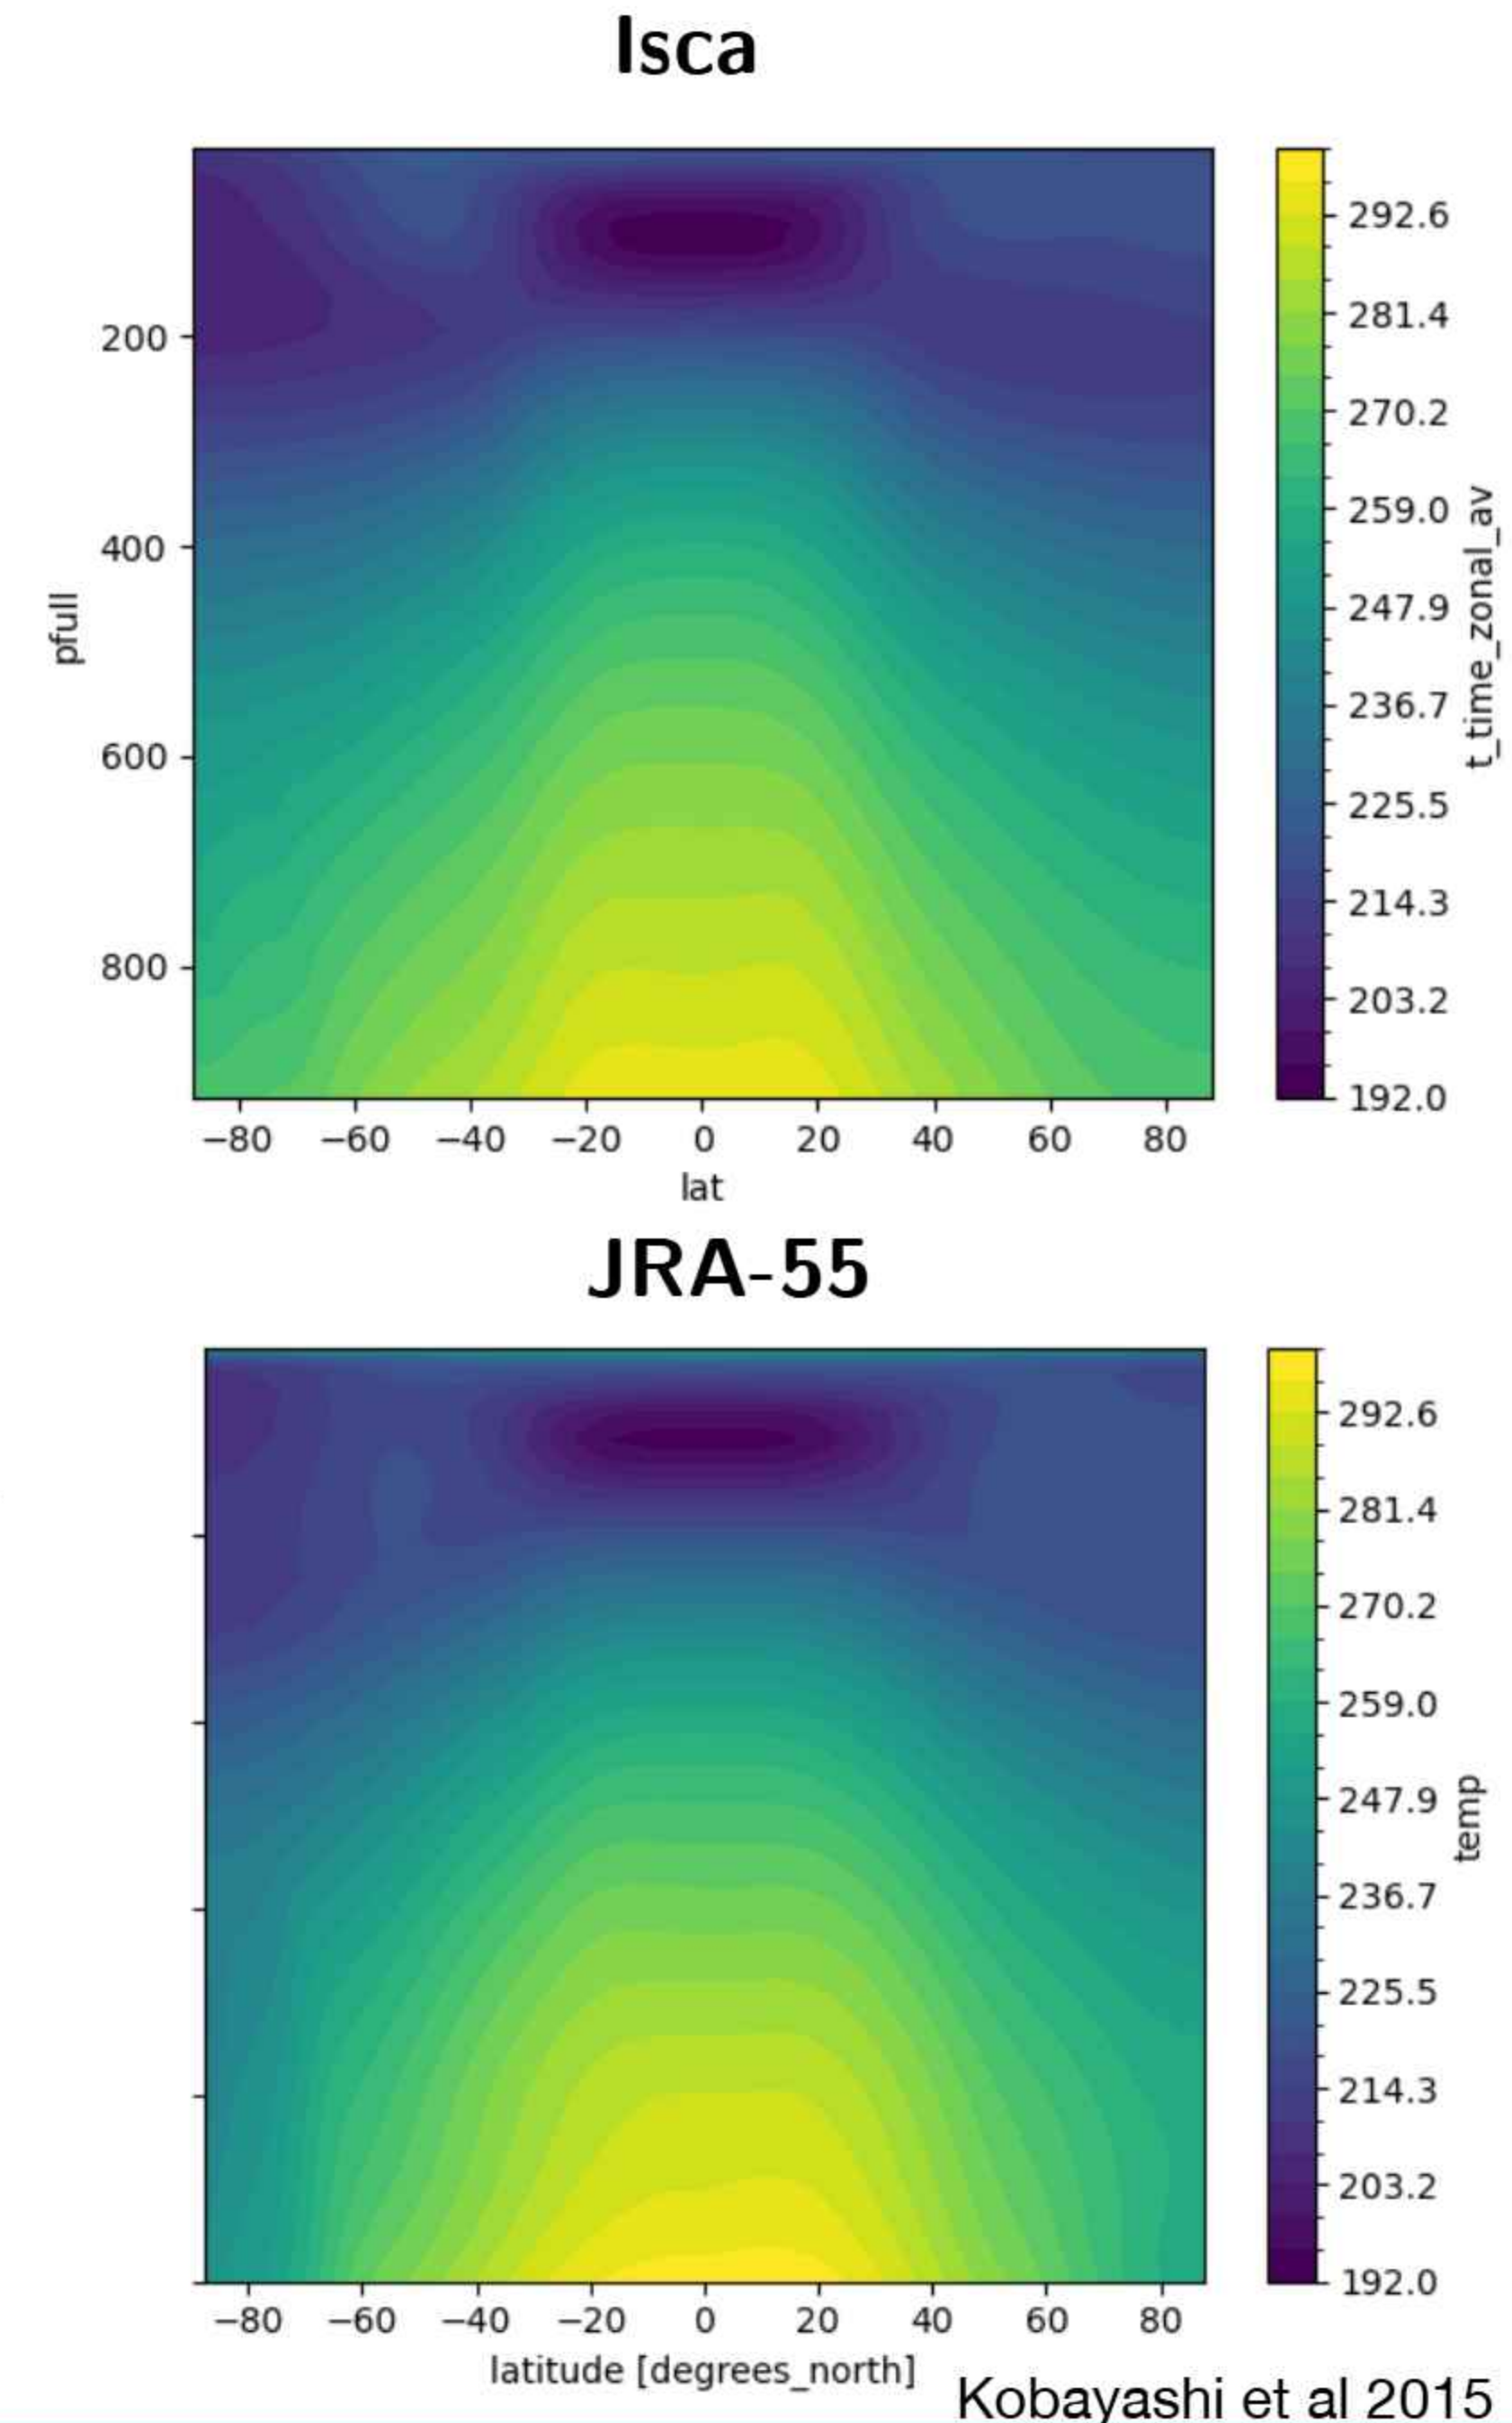
- We have explored approaches for including radiative transfer into a simple ML framework
- We are not using an MLWP system, but have built our own simple ML framework to illustrate the problems and explore ways forward
- **Training data:** we have run 140 different GCM experiments using the Isca modelling framework
 - Earth-like configuration with continents, a mixed-layer ocean, and Socrates radiative transfer without clouds
 - 5 different CO₂ values (278 - 1600ppmv), 7 different rotation rates (2x Earth - 1/16 Earth) and 4 different obliquities (0 - 1.5 x Earth)



Kobayashi et al 2015

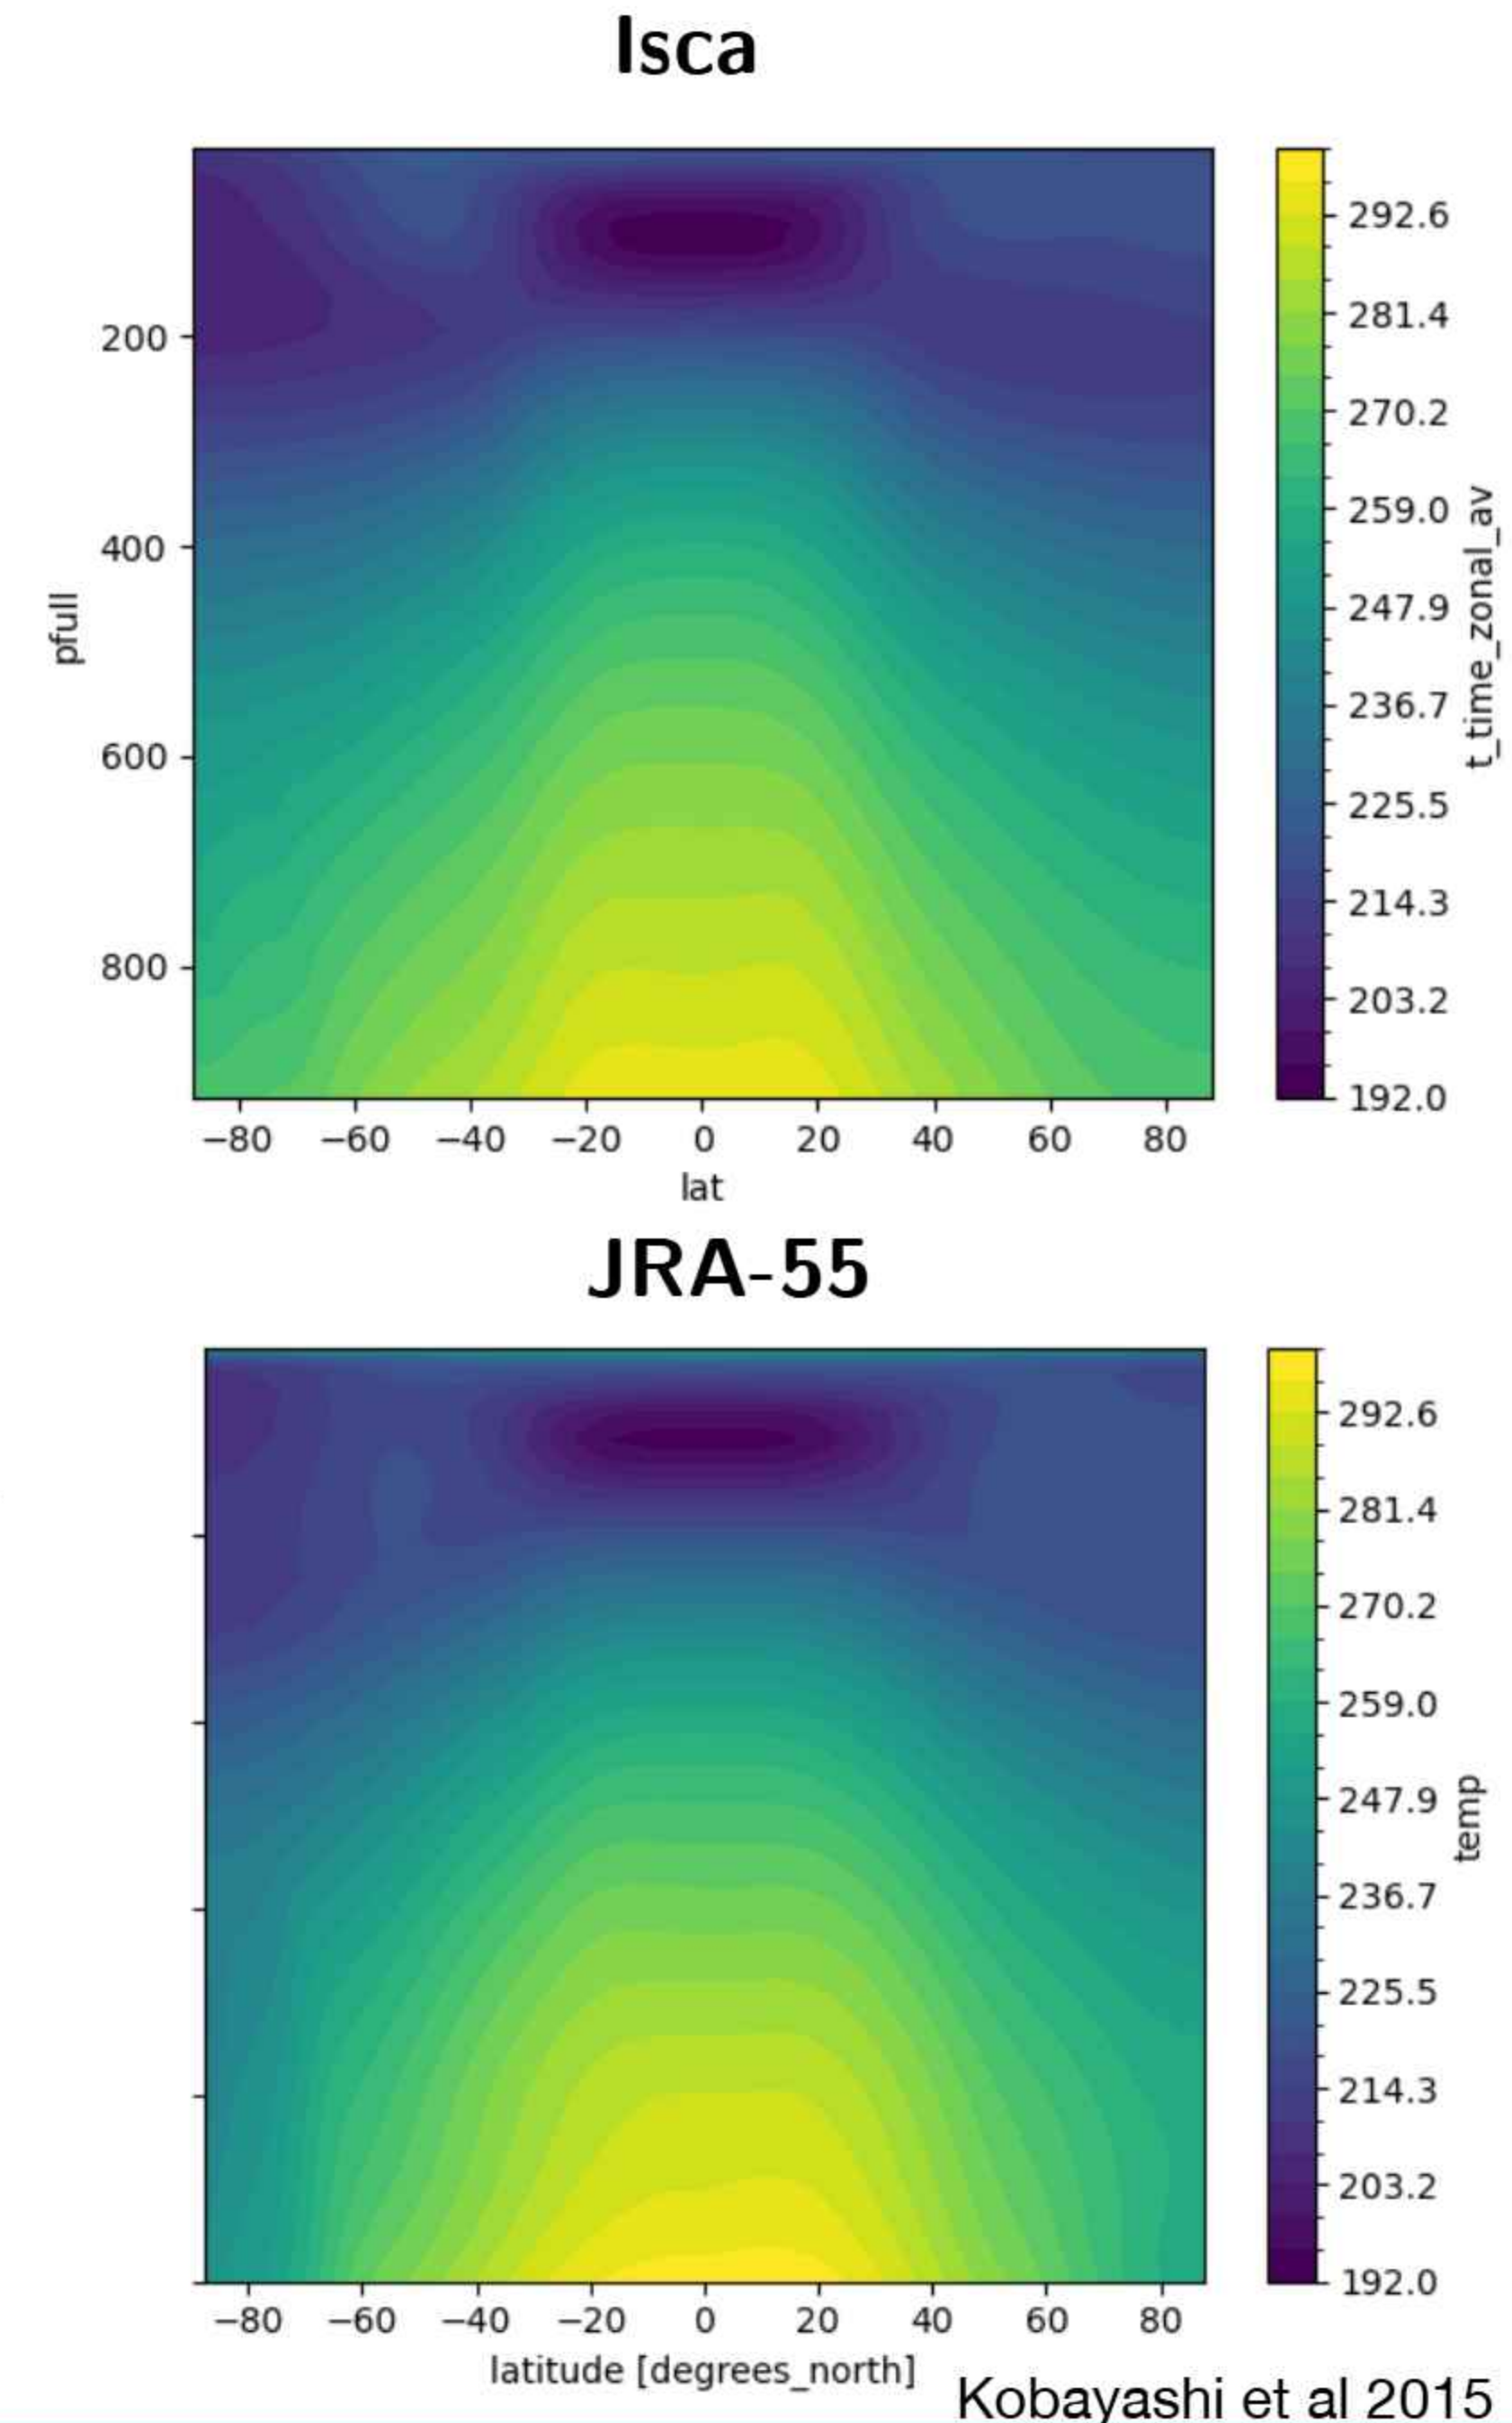
Our approach

- We have explored approaches for including radiative transfer into a simple ML framework
- We are not using an MLWP system, but have built our own simple ML framework to illustrate the problems and explore ways forward
- **Training data:** we have run 140 different GCM experiments using the Isca modelling framework
 - Earth-like configuration with continents, a mixed-layer ocean, and Socrates radiative transfer without clouds
 - 5 different CO₂ values (278 - 1600ppmv), 7 different rotation rates (2x Earth - 1/16 Earth) and 4 different obliquities (0 - 1.5 x Earth)
- Lots of analogous ML problems in exoplanetary research too!



Our approach

- We have explored approaches for including radiative transfer into a simple ML framework
- We are not using an MLWP system, but have built our own simple ML framework to illustrate the problems and explore ways forward
- **Training data:** we have run 140 different GCM experiments using the Isca modelling framework
 - Earth-like configuration with continents, a mixed-layer ocean, and Socrates radiative transfer without clouds
 - 5 different CO₂ values (278 - 1600ppmv), 7 different rotation rates (2x Earth - 1/16 Earth) and 4 different obliquities (0 - 1.5 x Earth)
- Lots of analogous ML problems in exoplanetary research too!
- **Our aim is to take in 3 input parameters - CO₂, rotation rate and obliquity and use ML to predict the zonal-mean and annual-mean temperature structure and the zonal-mean zonal wind**



Network structure

Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image

Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image

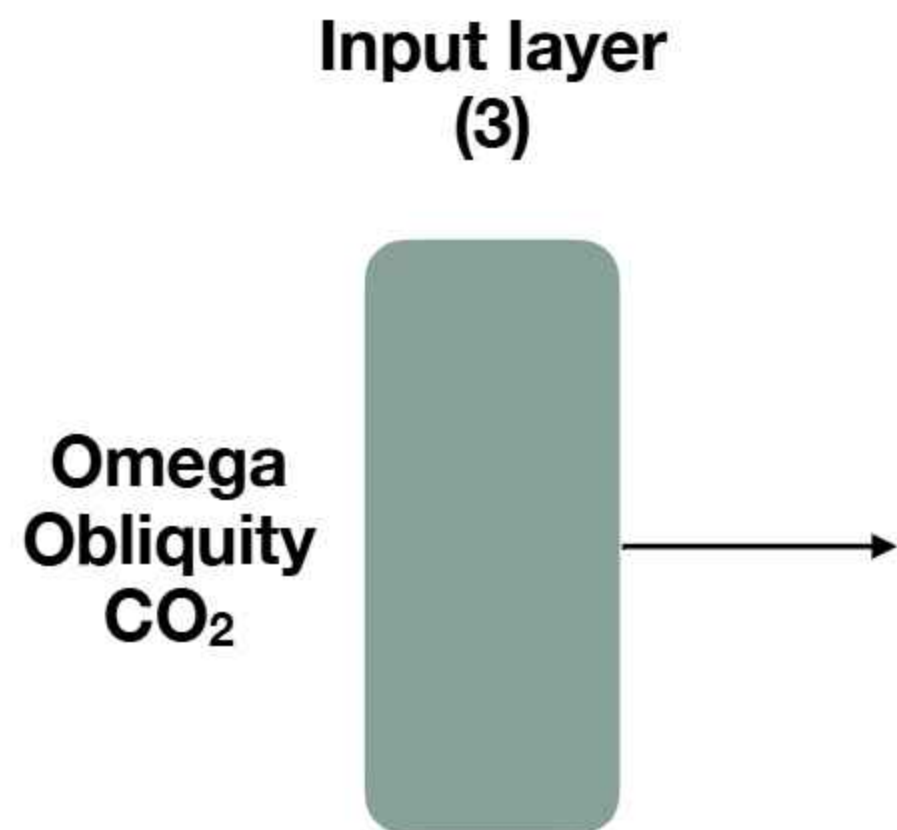
Omega
Obliquity
CO₂



University
of Exeter

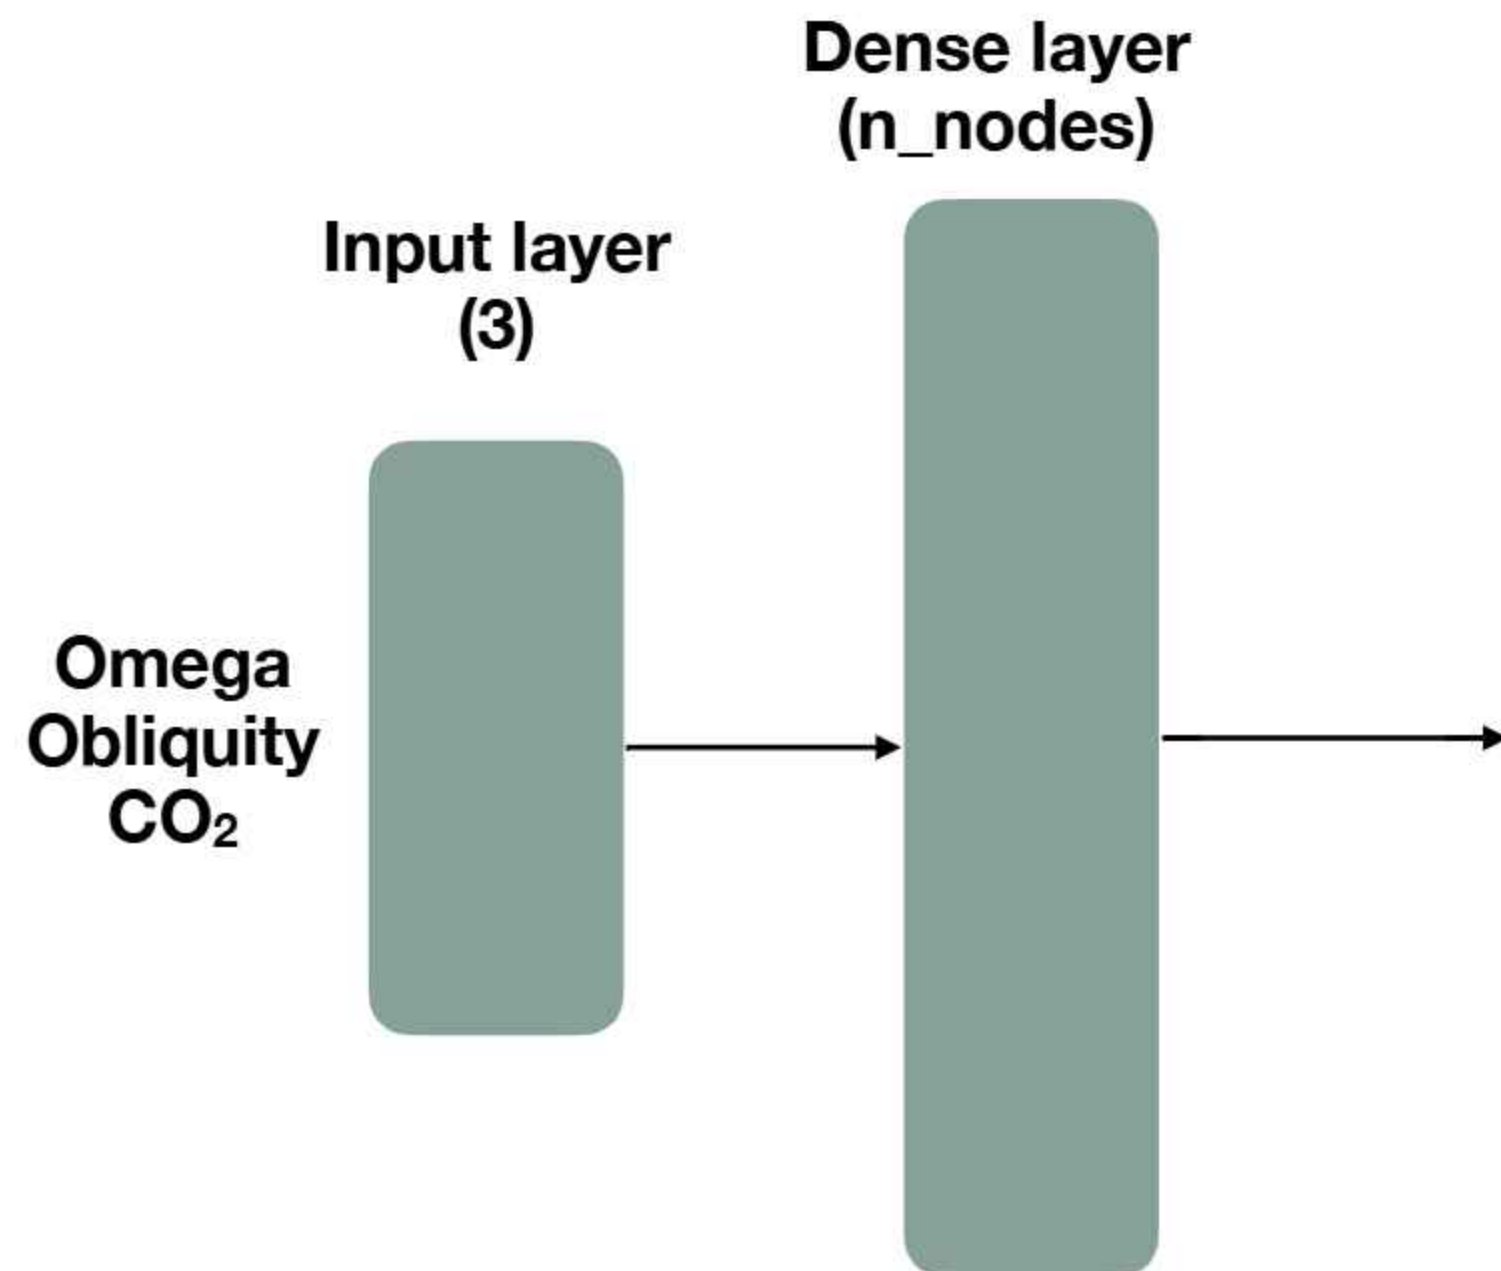
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



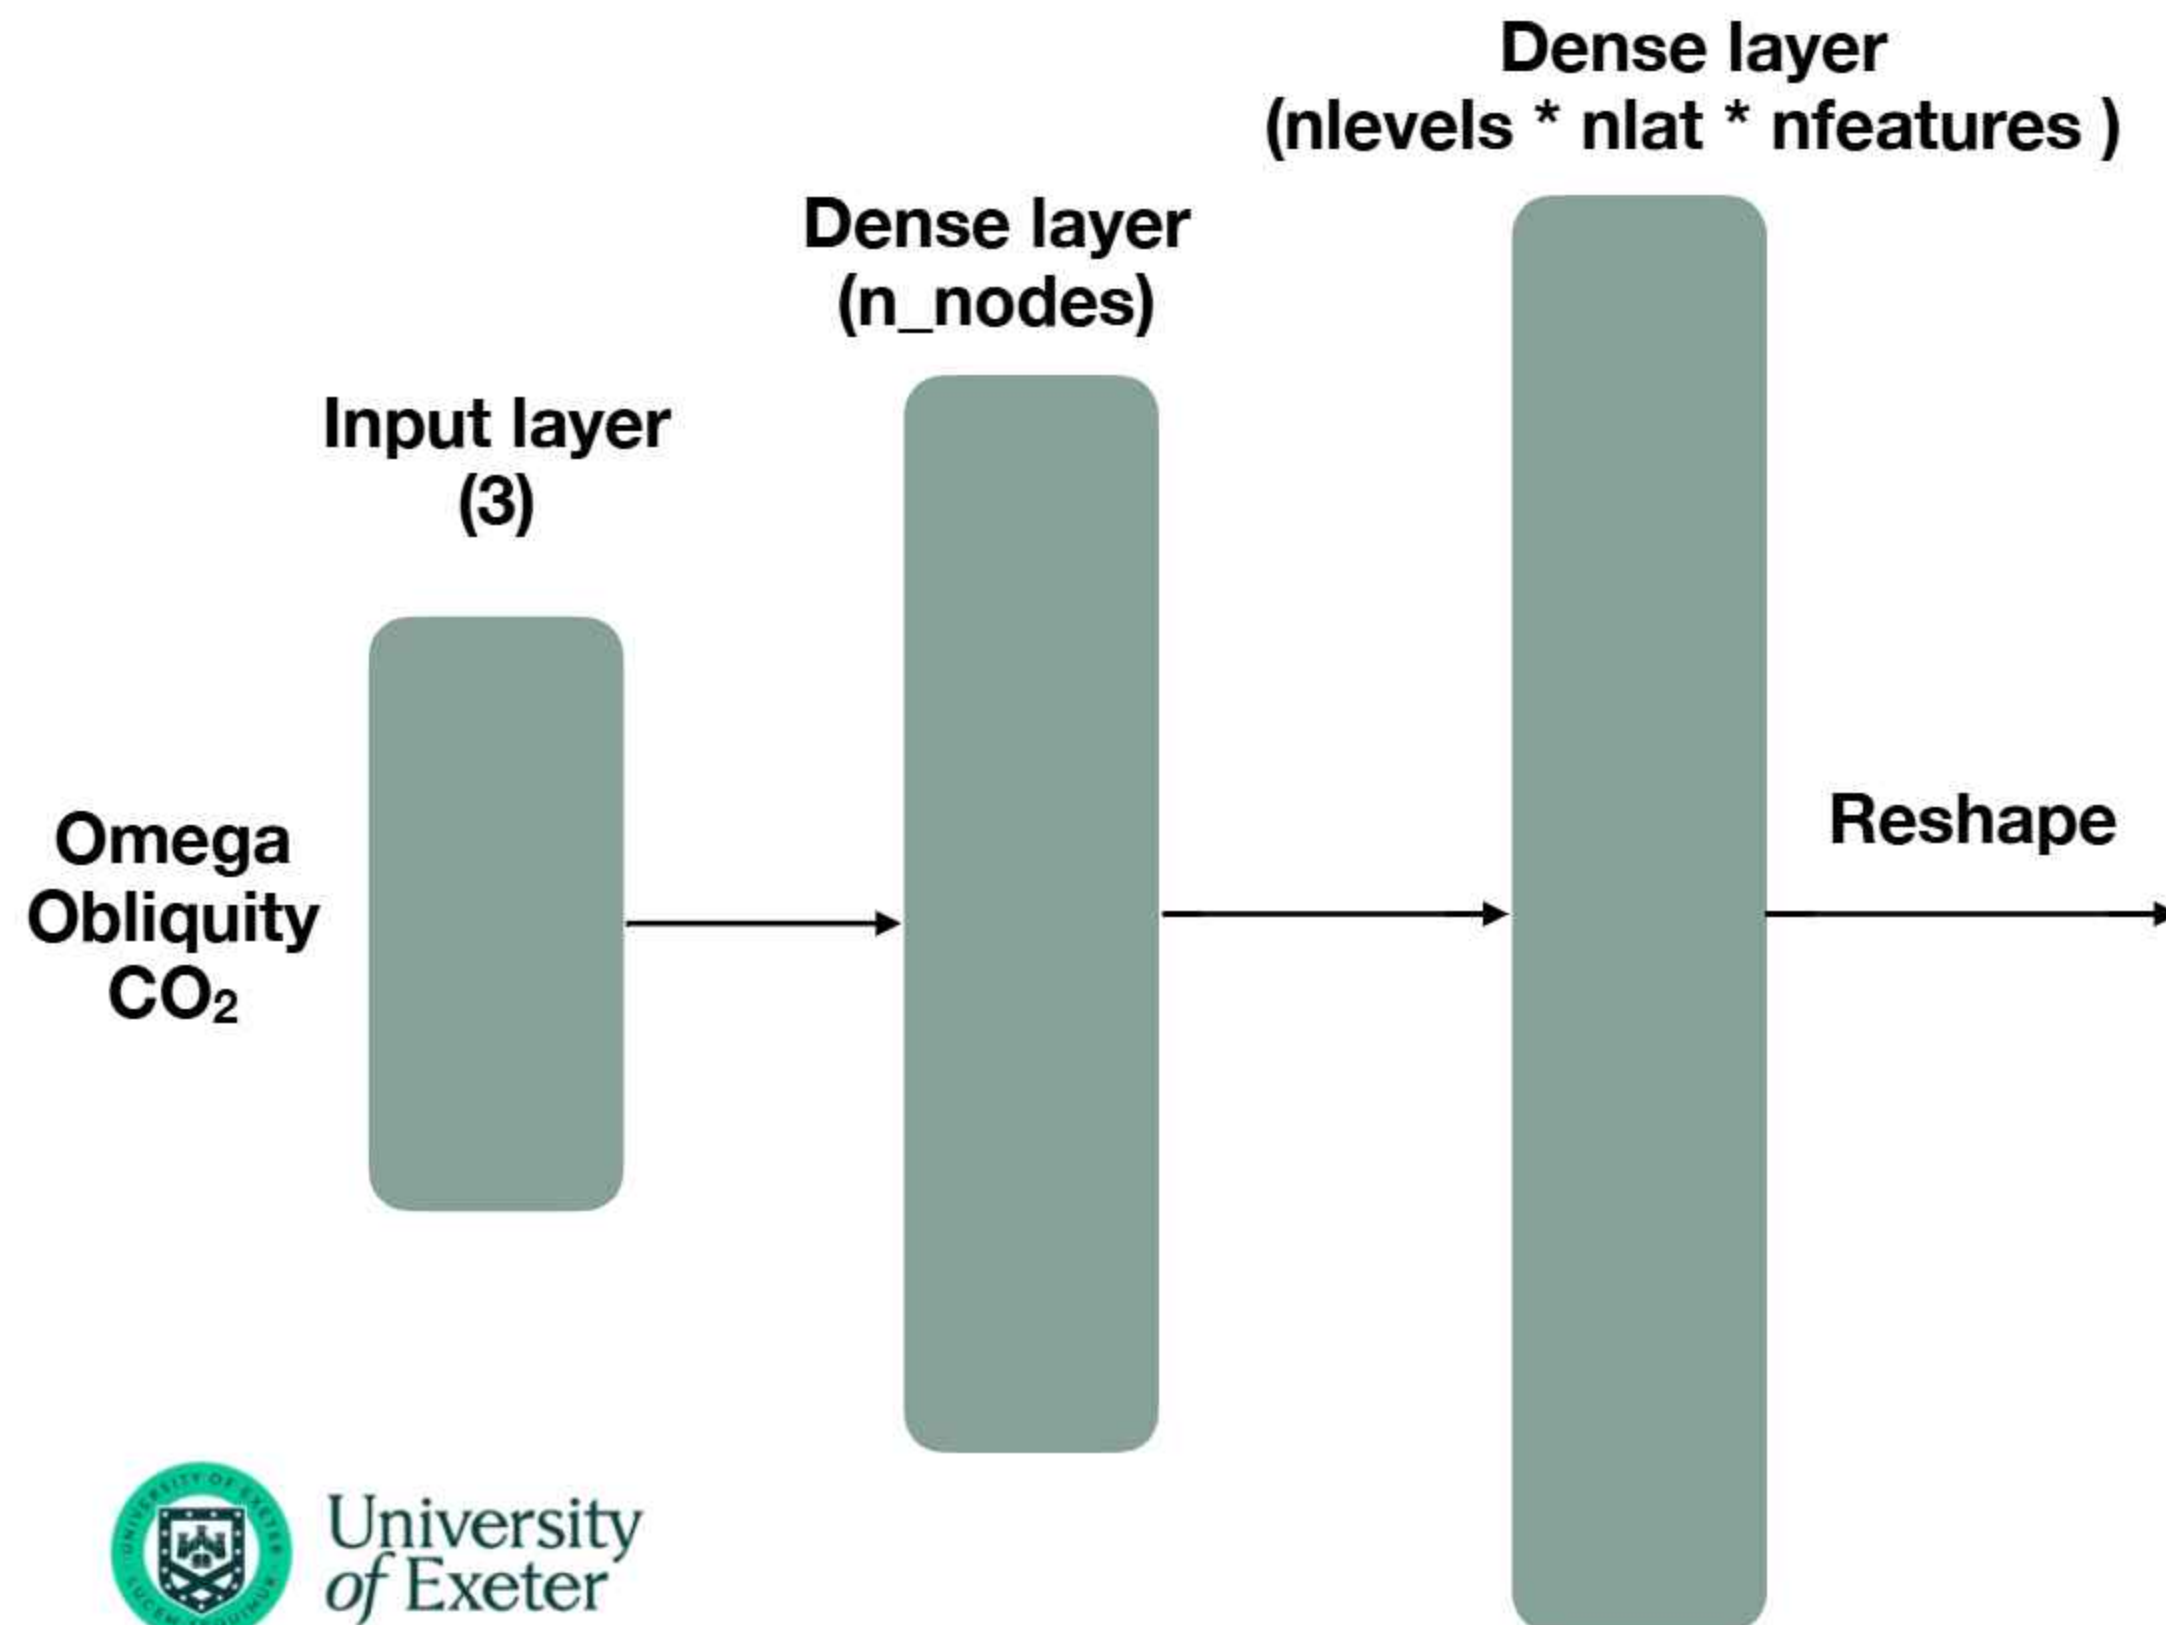
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



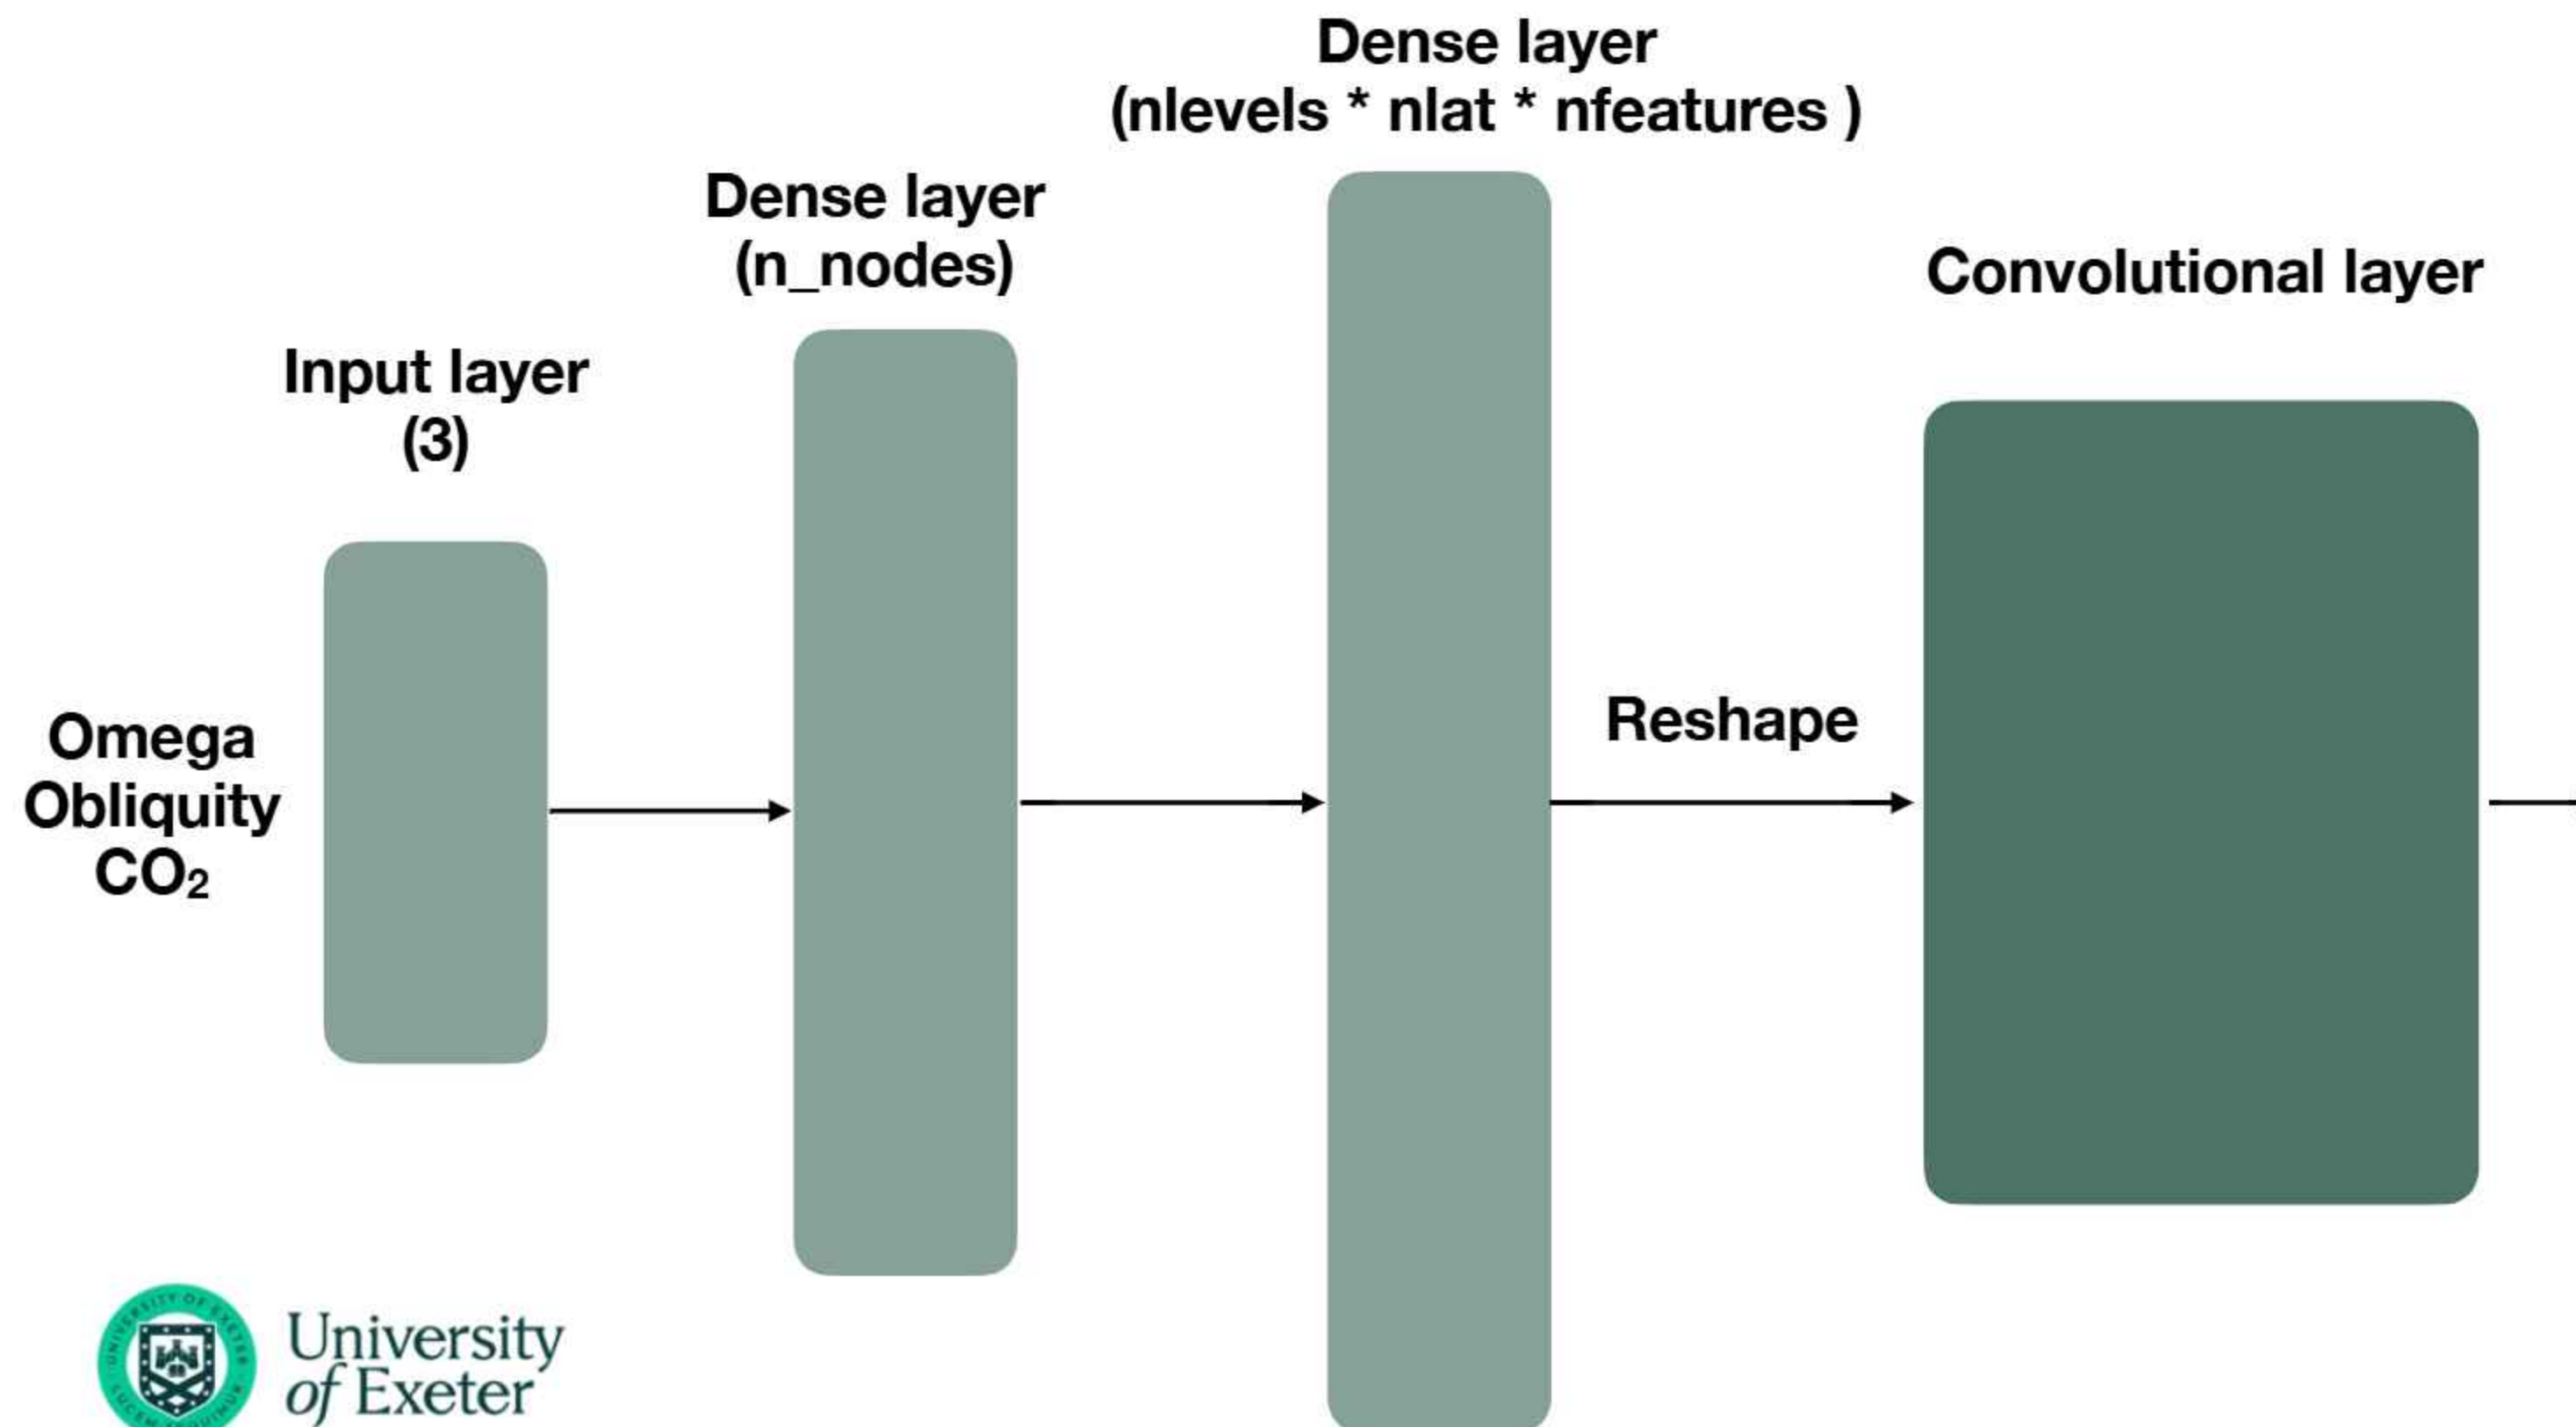
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



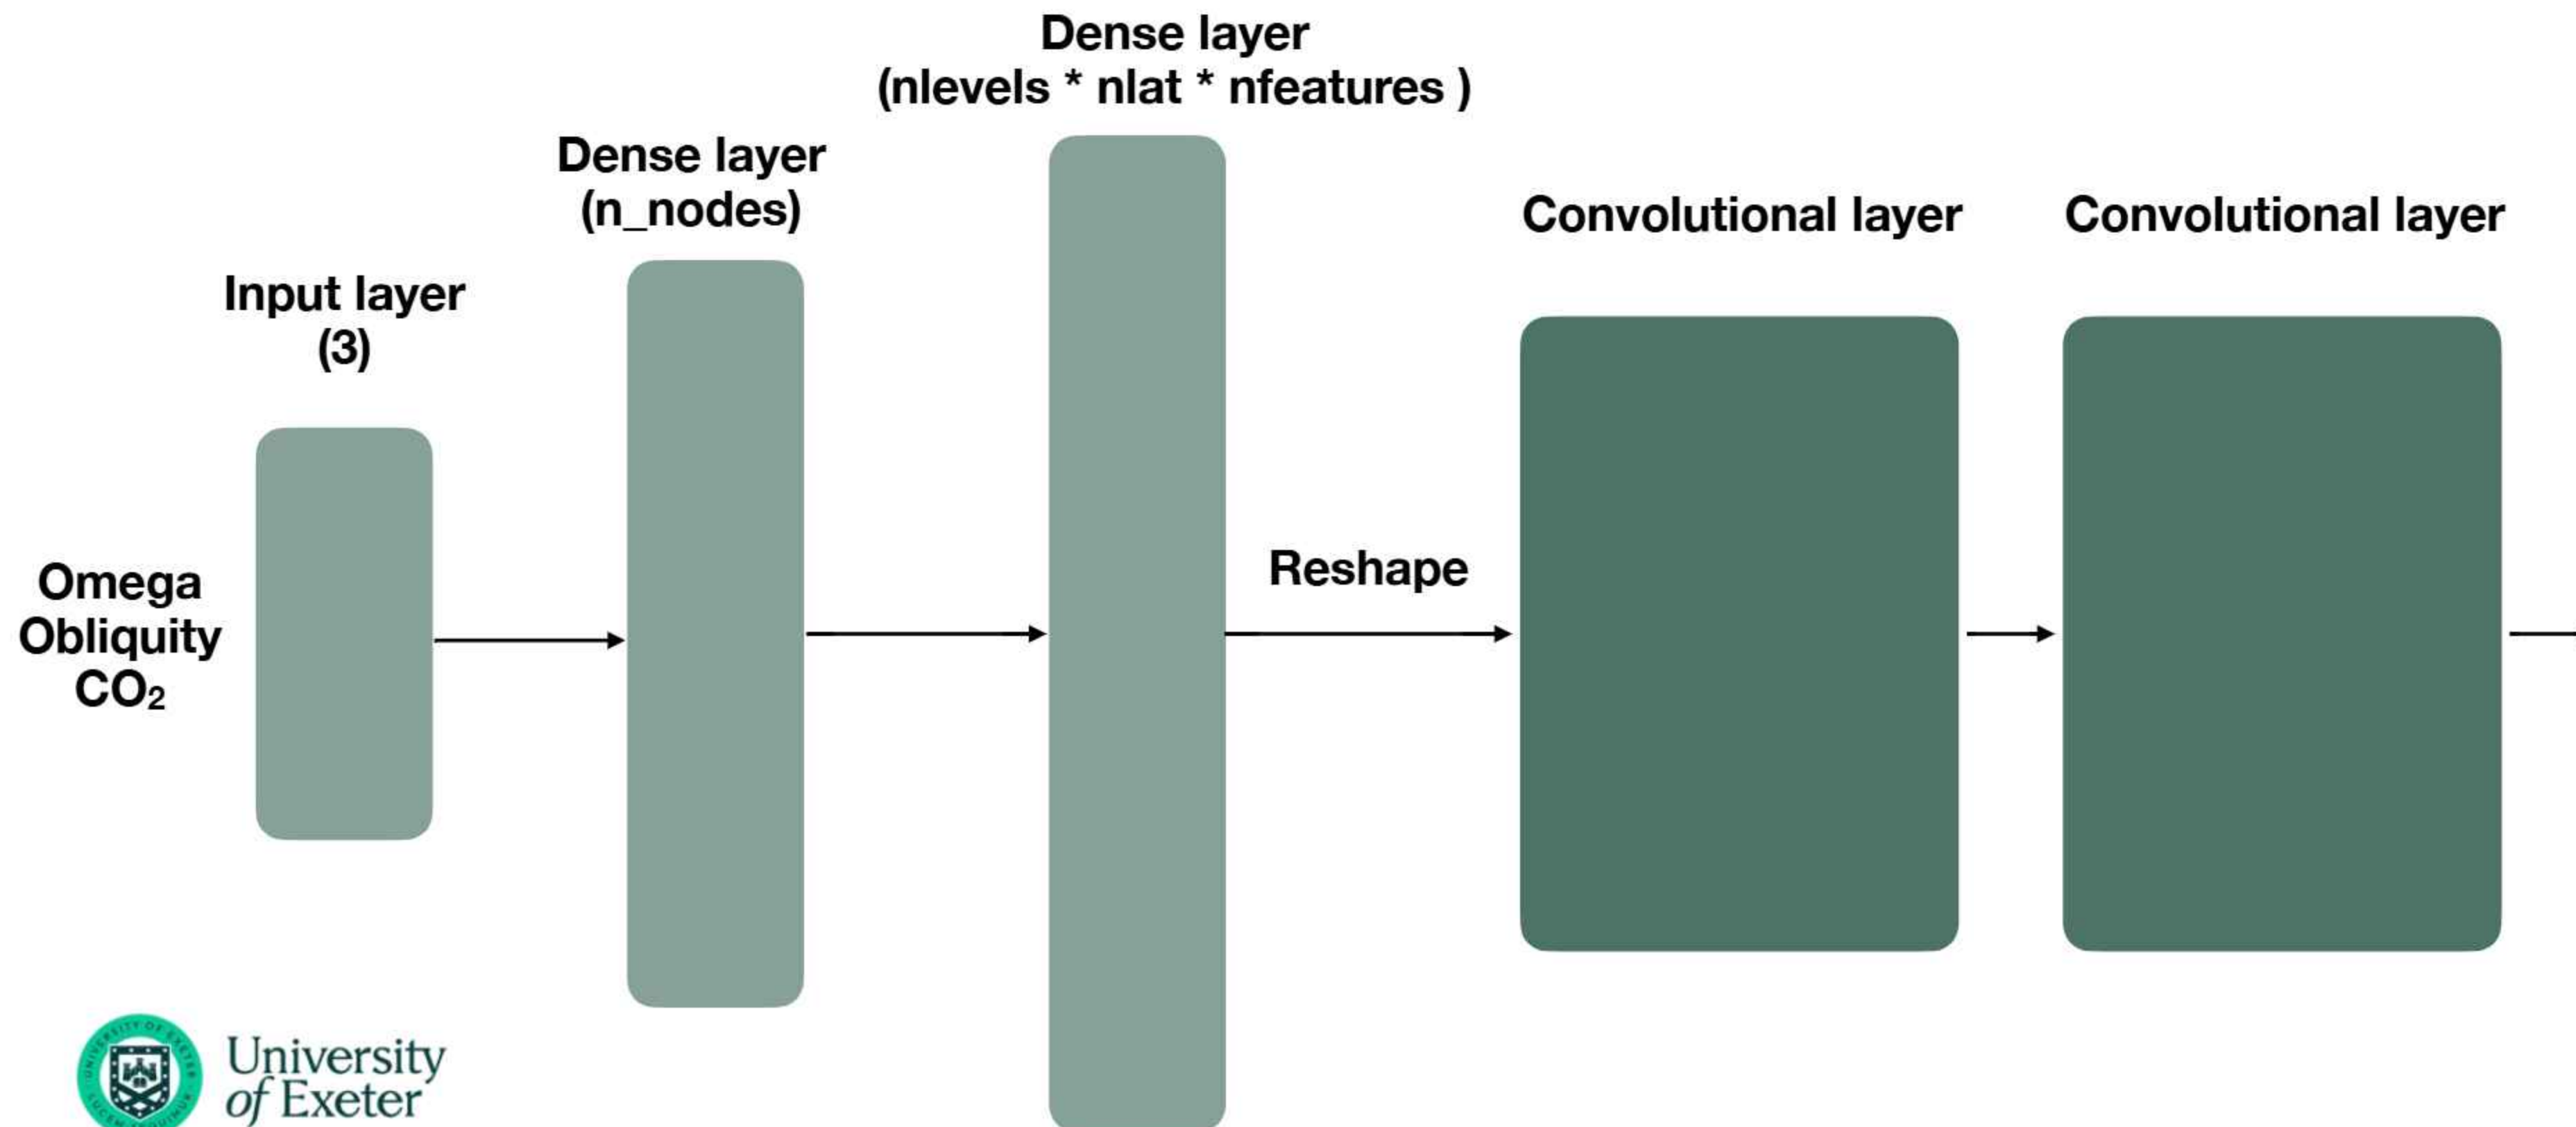
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



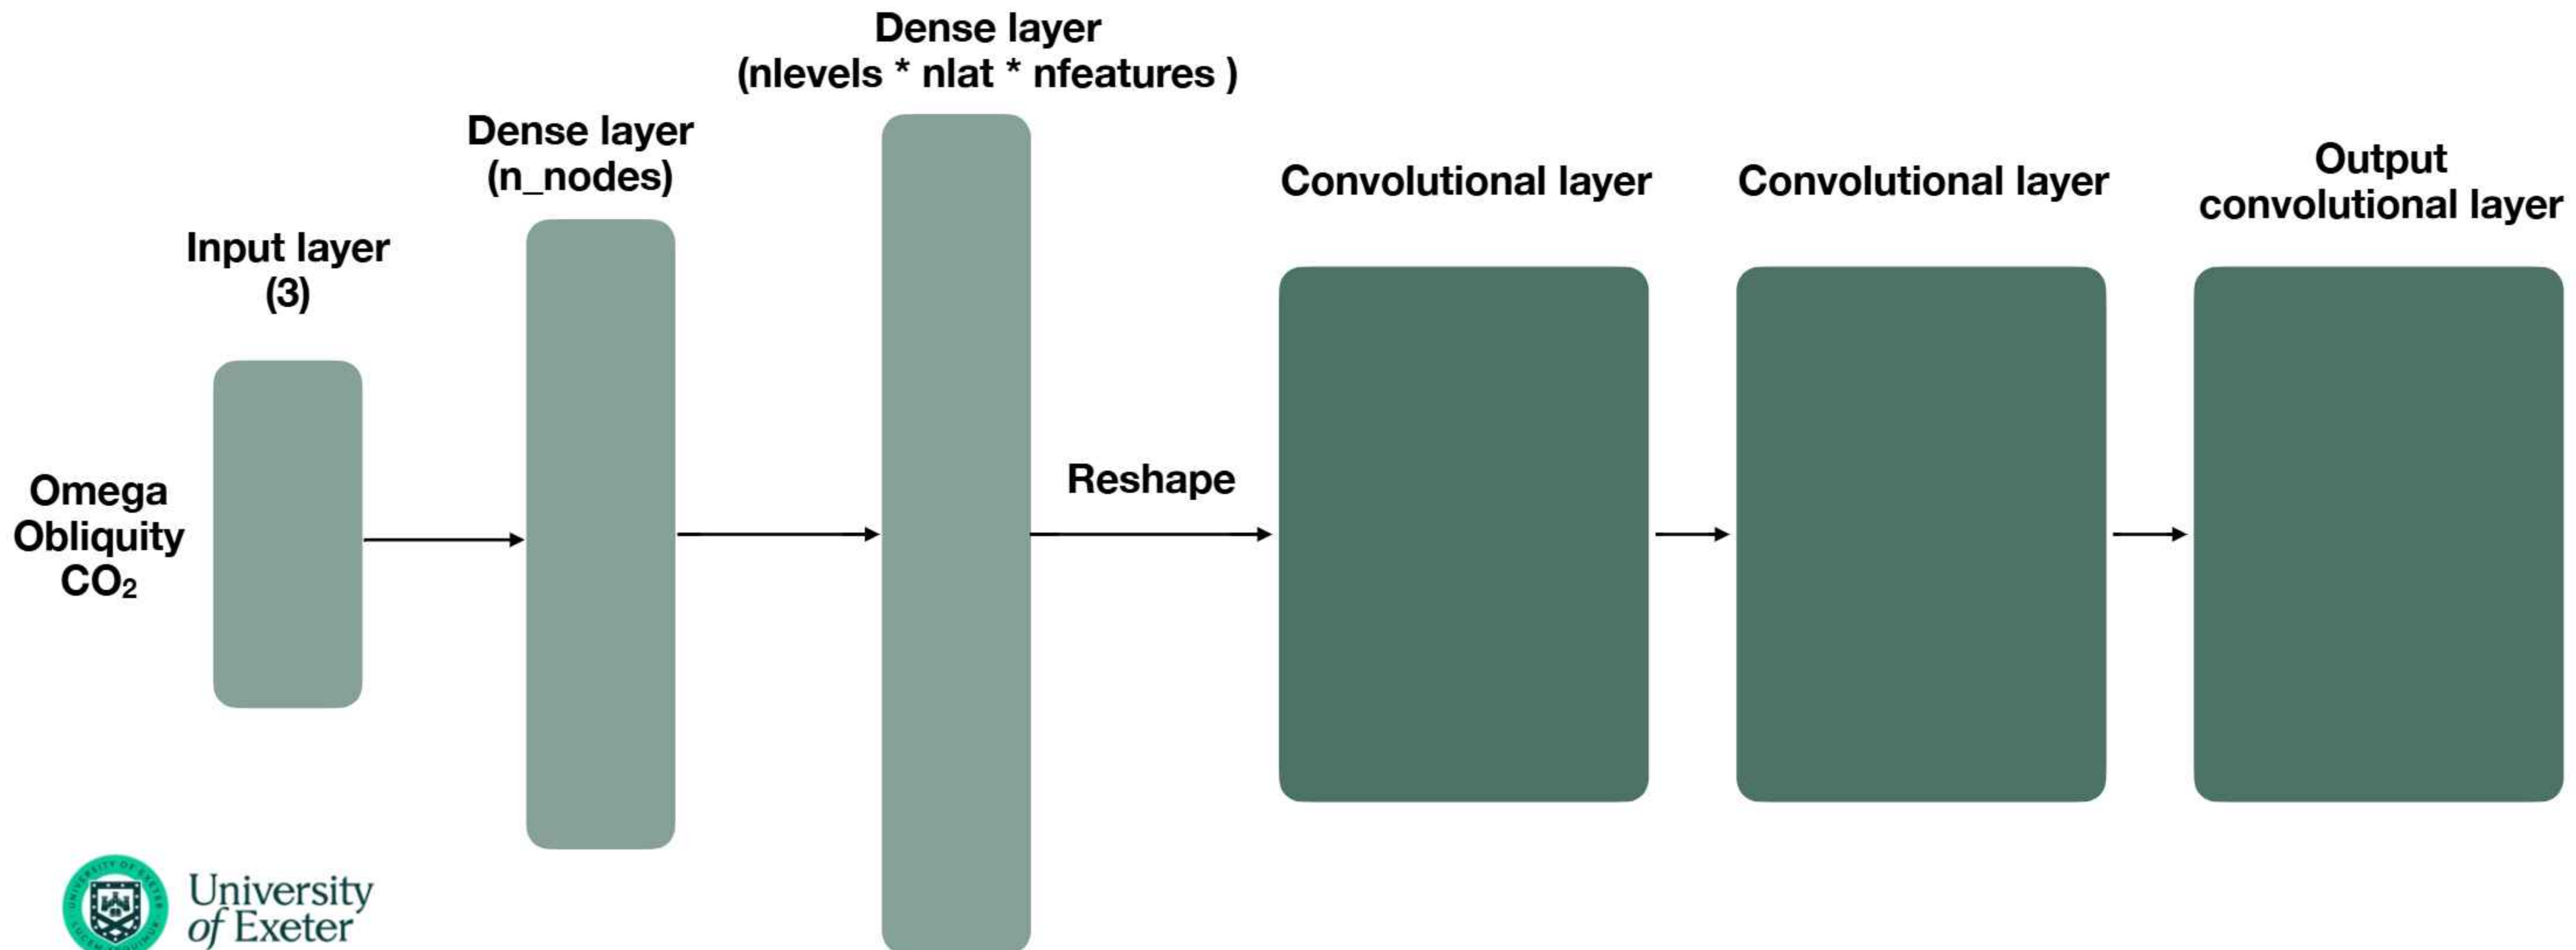
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



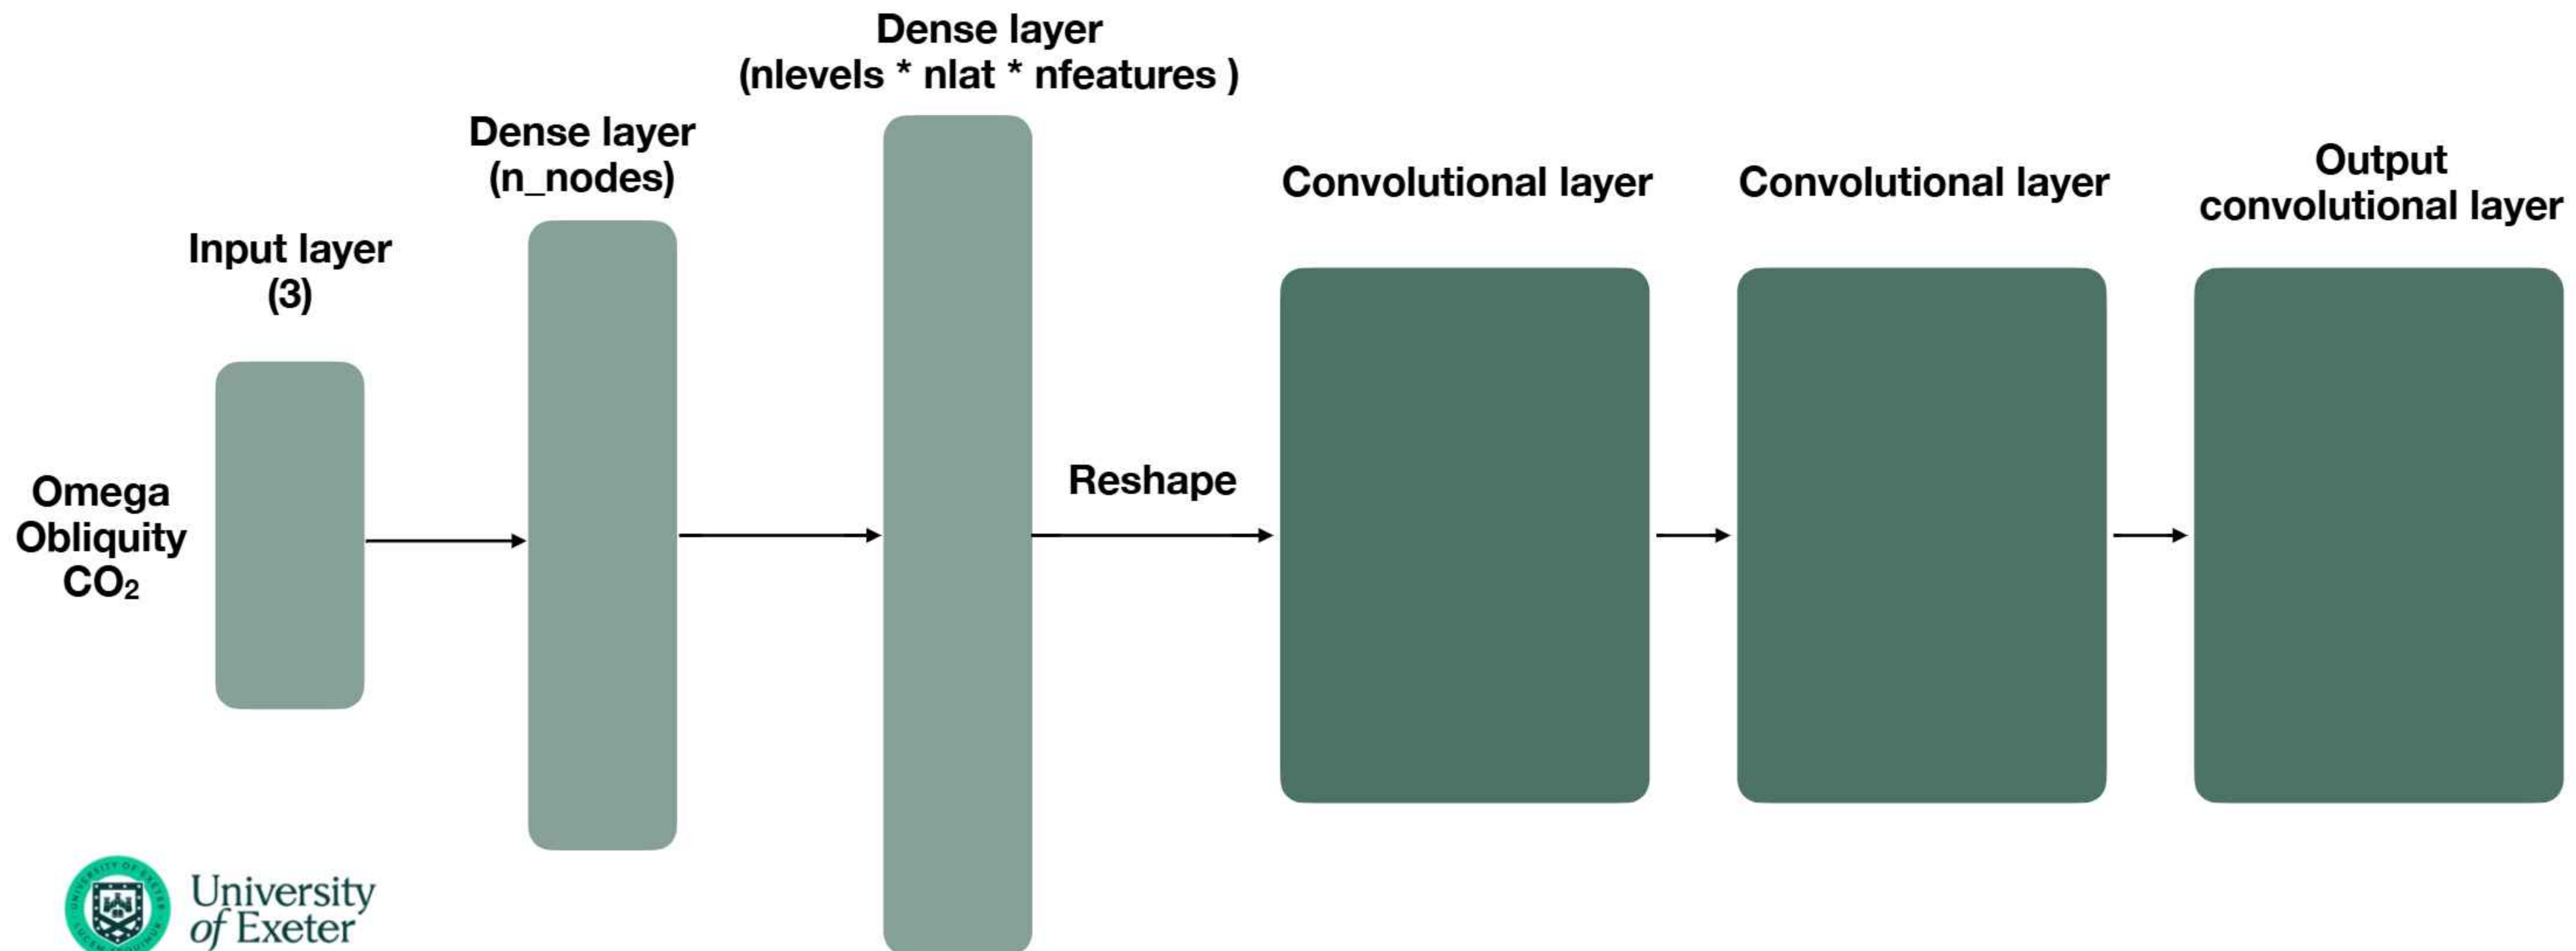
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image



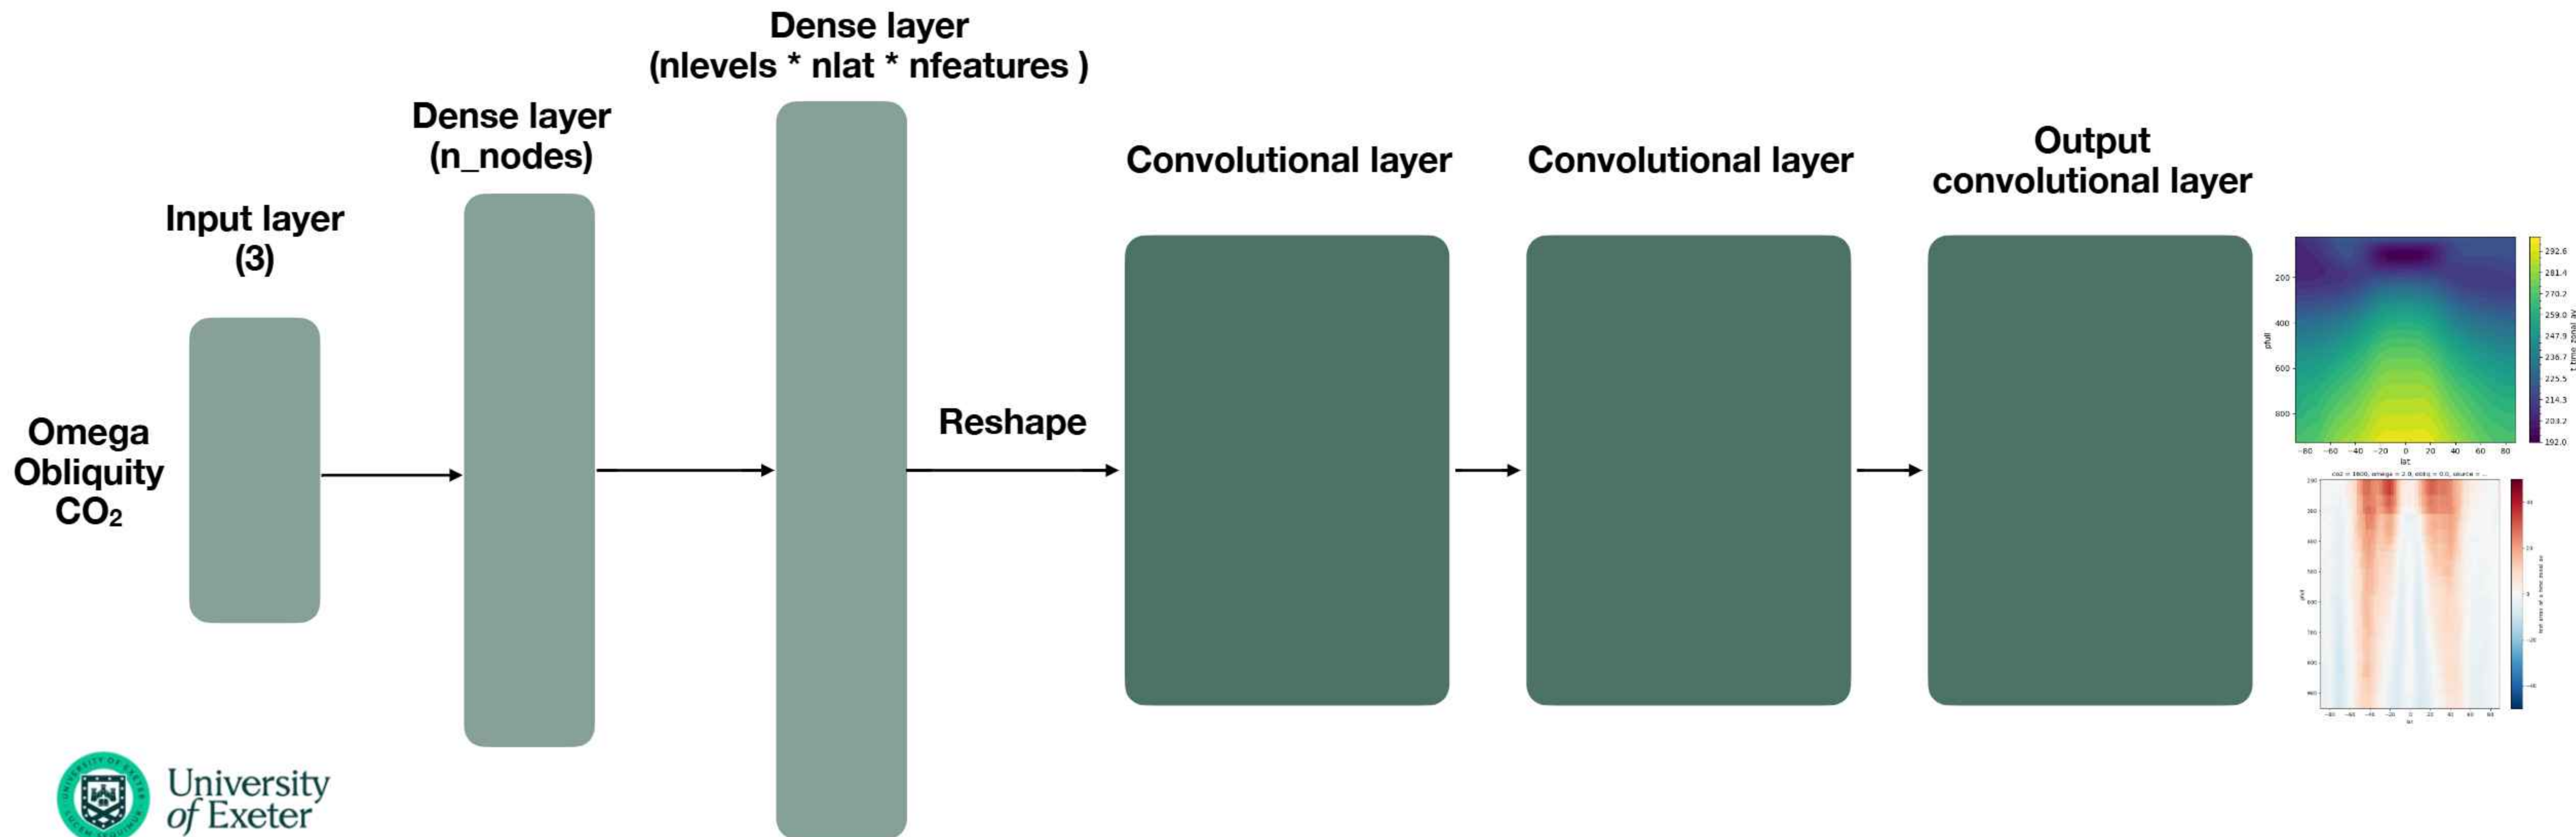
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image
- Zonal-mean temperature and zonal-mean zonal wind are predicted as **channels** in the final image



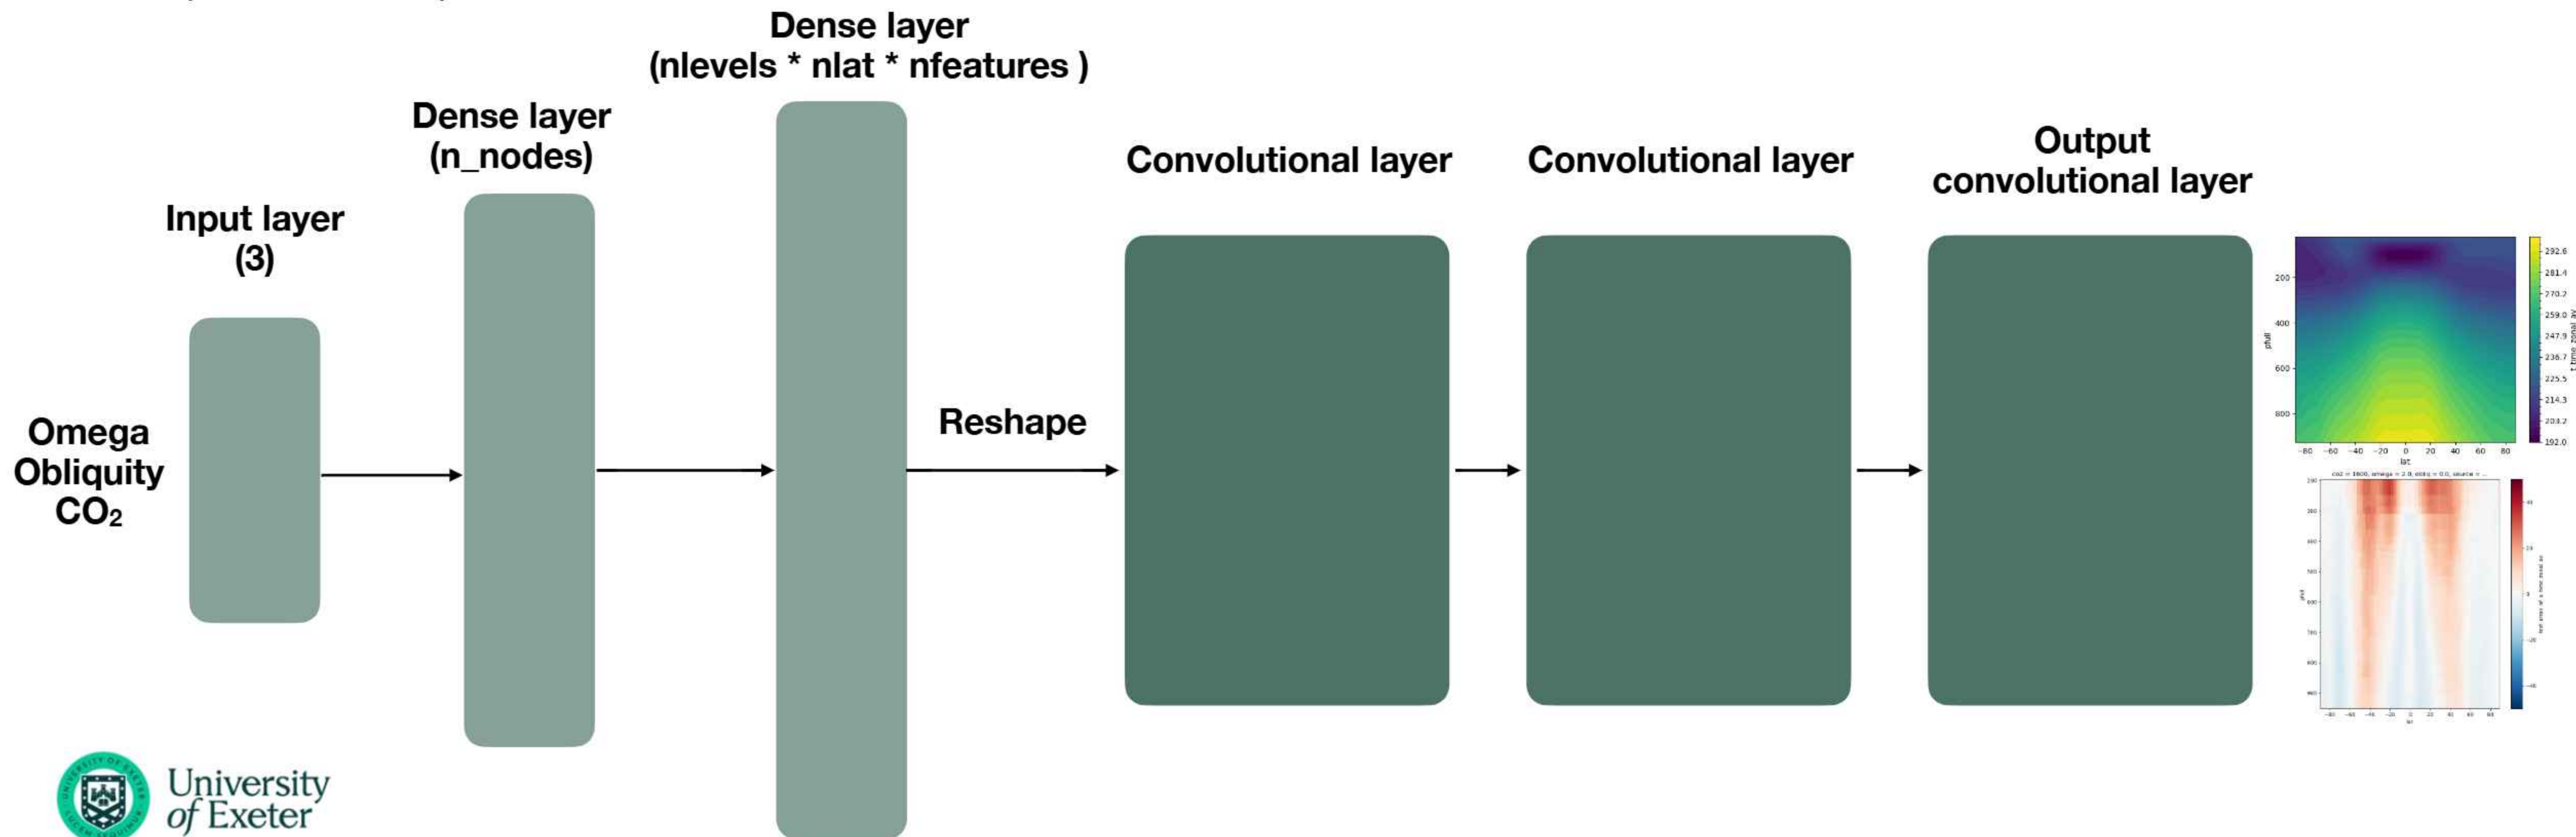
Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image
- Zonal-mean temperature and zonal-mean zonal wind are predicted as **channels** in the final image



Network structure

- For the ML we've used a **Convolutional Neural Network (CNN)**, with dense layers building up from the inputs to a single layer with enough dimensions to then be reshaped into a 2D image
- Zonal-mean temperature and zonal-mean zonal wind are predicted as **channels** in the final image
- We have explored the impact of how the data is normalised, how the input parameters are normalised etc (not shown)



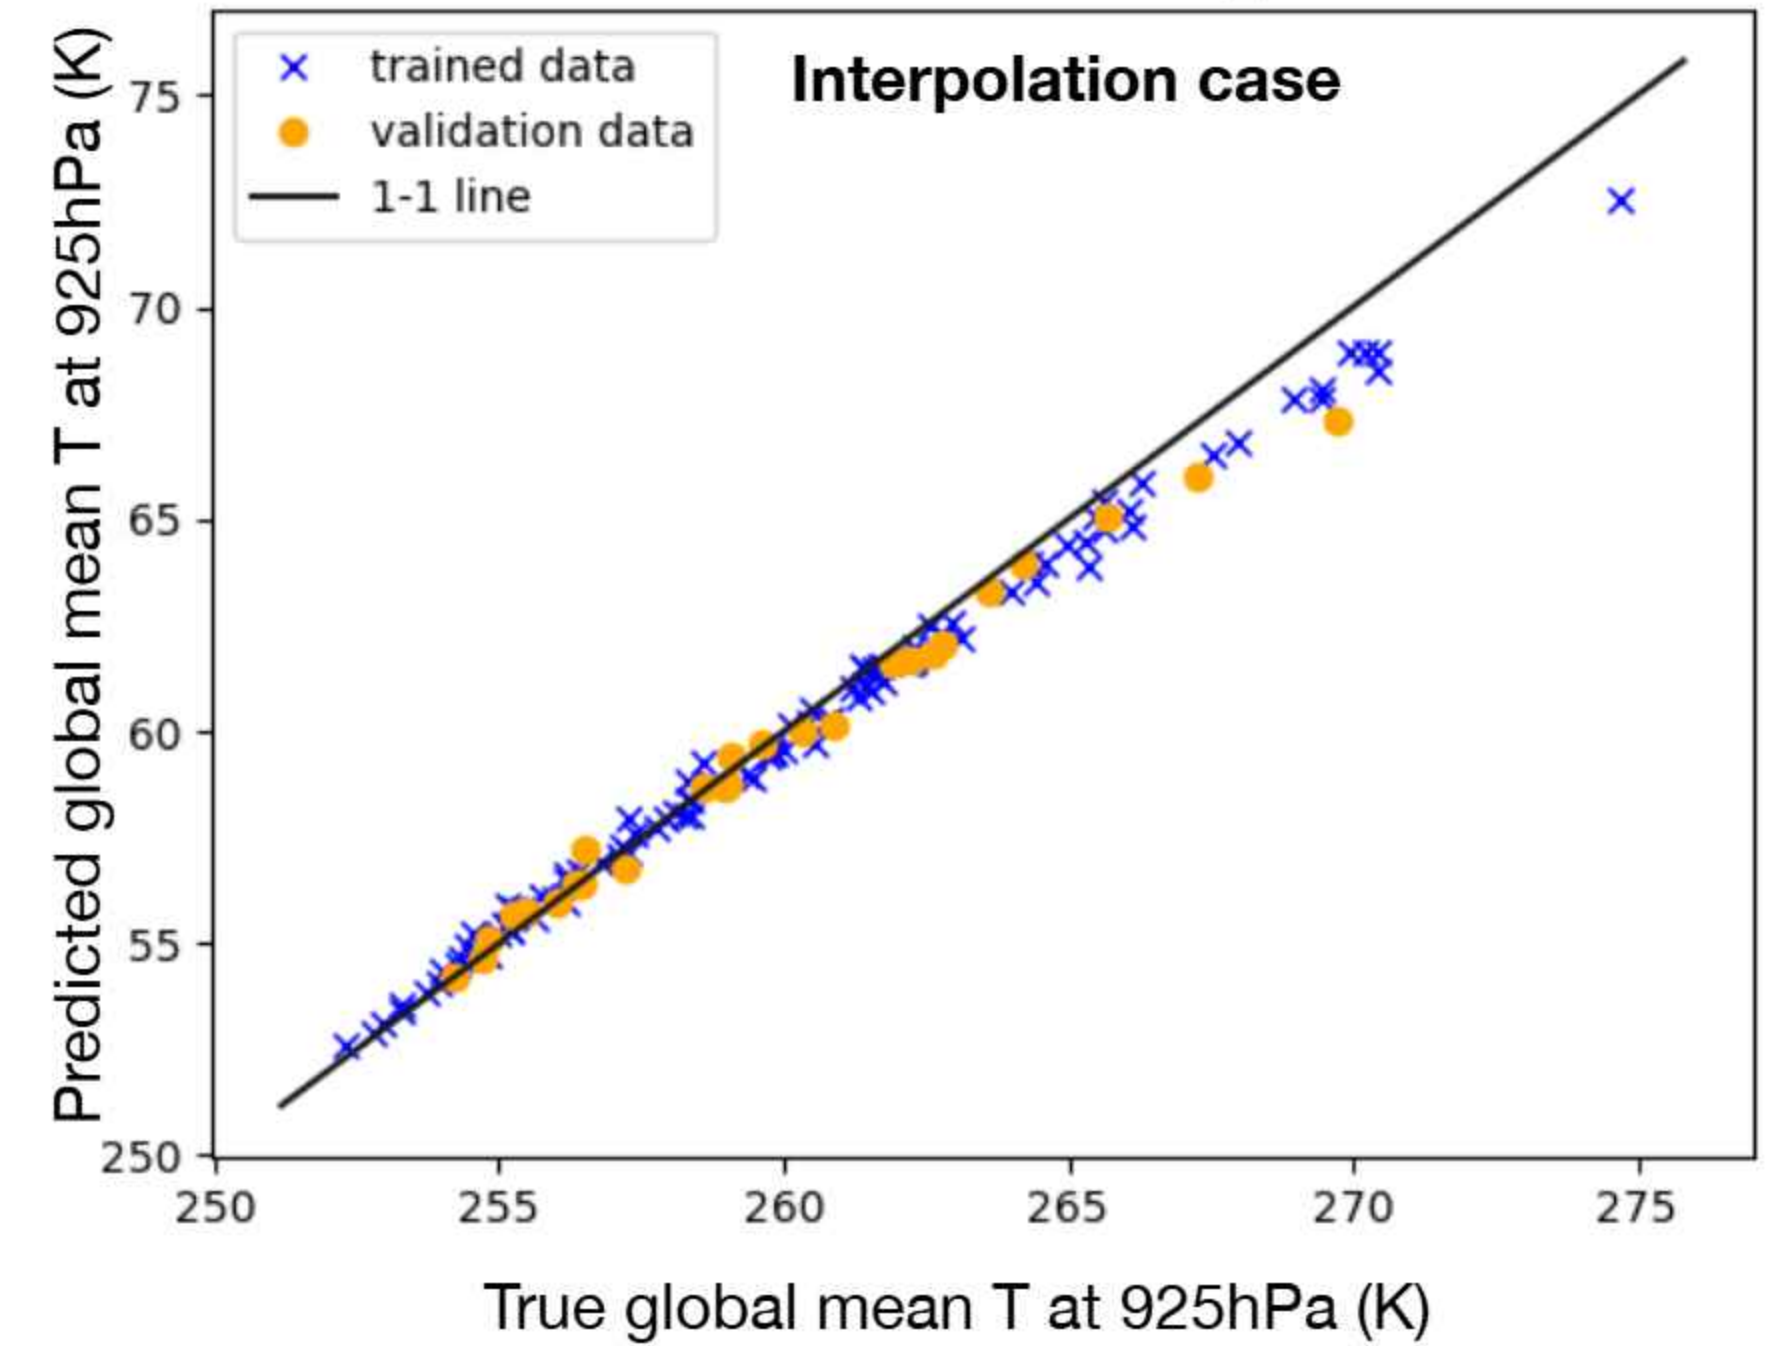
Network performance for interpolation vs extrapolation

Network performance for interpolation vs extrapolation

- Model performs well at predicting the zonal-wind and temperature when **interpolating**, i.e. when it sees a random subset of the full parameter space in the training phase

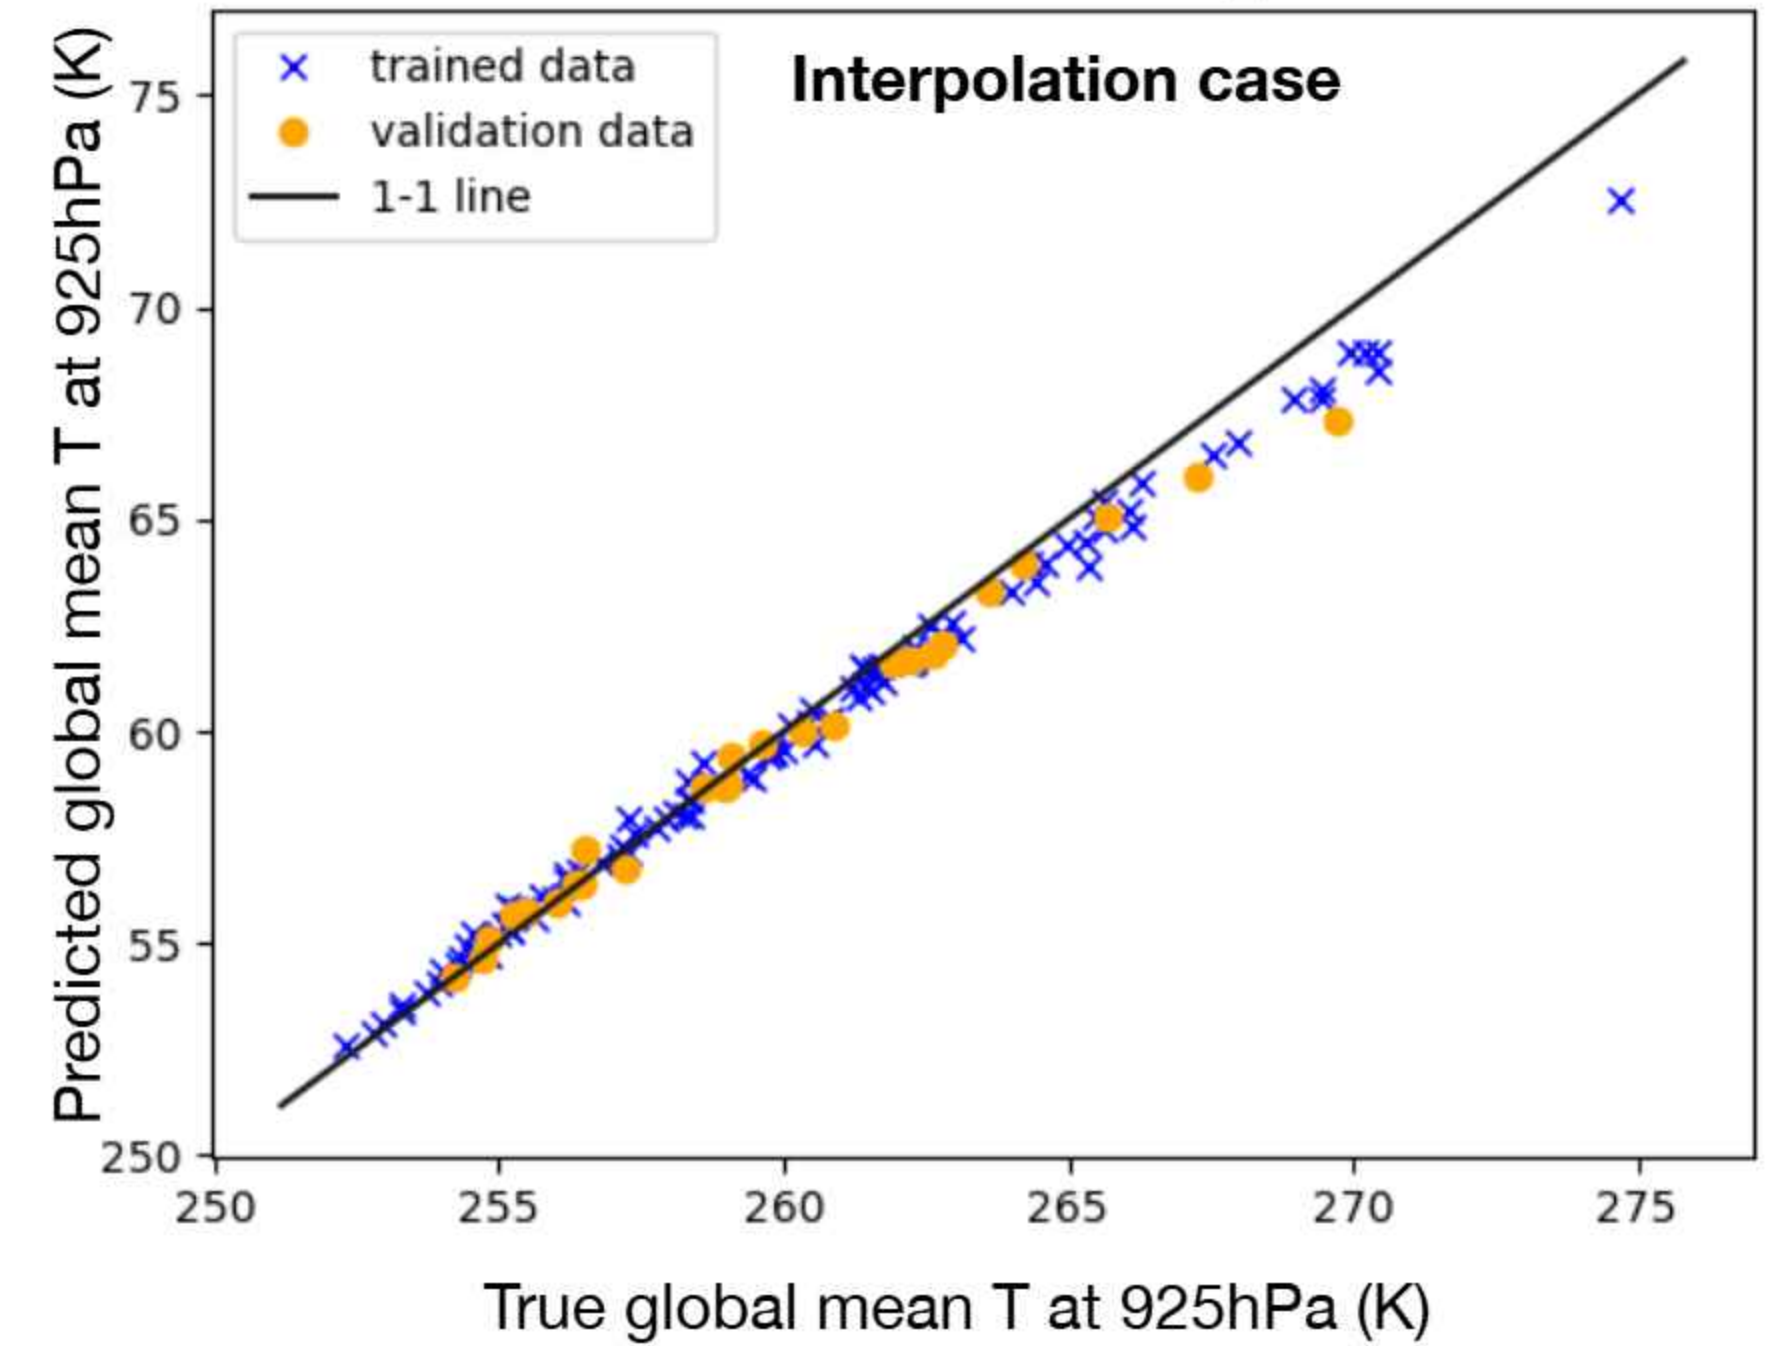
Network performance for interpolation vs extrapolation

- Model performs well at predicting the zonal-wind and temperature when **interpolating**, i.e. when it sees a random subset of the full parameter space in the training phase



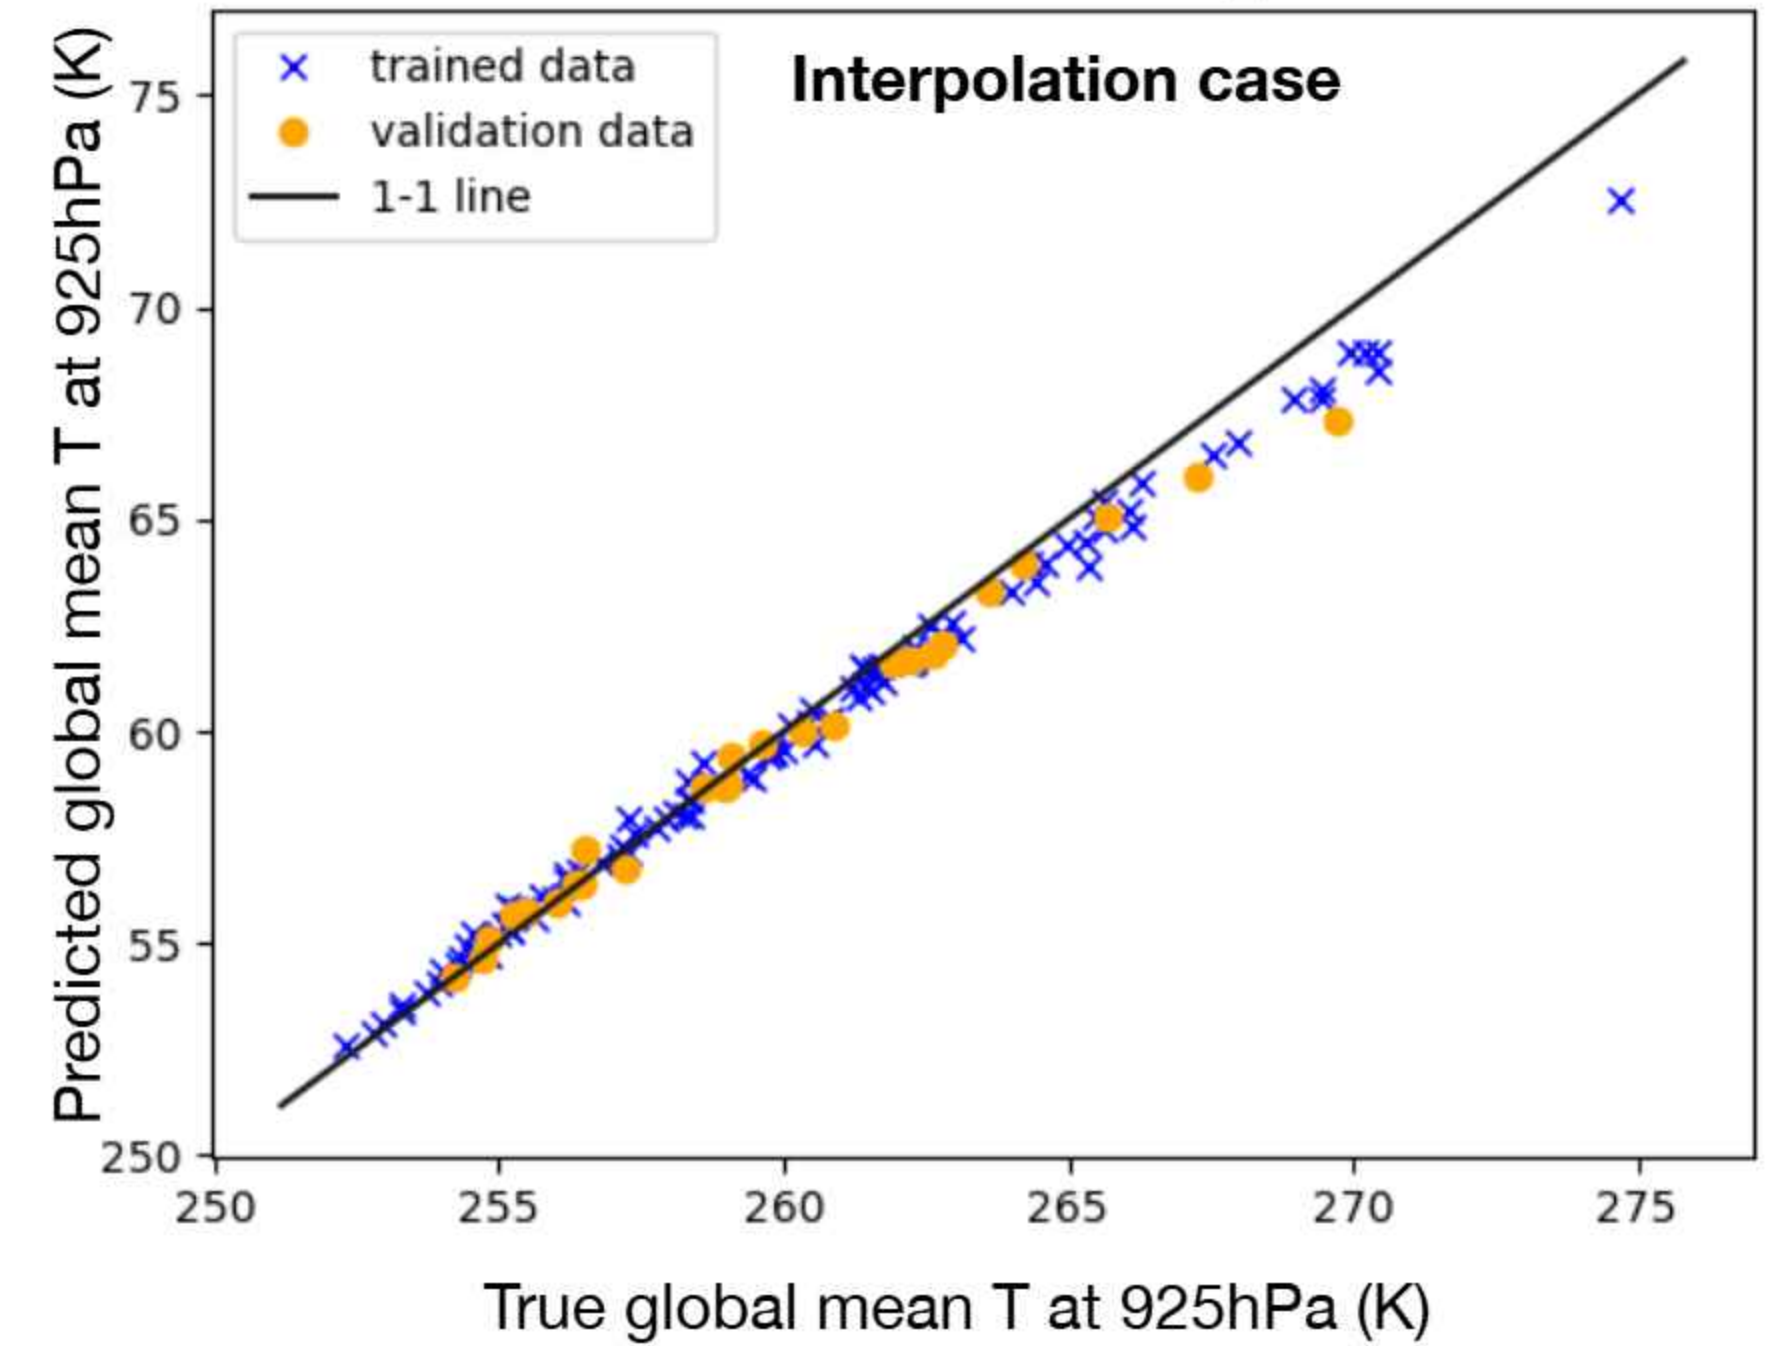
Network performance for interpolation vs extrapolation

- Model performs well at predicting the zonal-wind and temperature when **interpolating**, i.e. when it sees a random subset of the full parameter space in the training phase
- But the model performs significantly worse when it's asked to **extrapolate** to cases that it hasn't seen in training



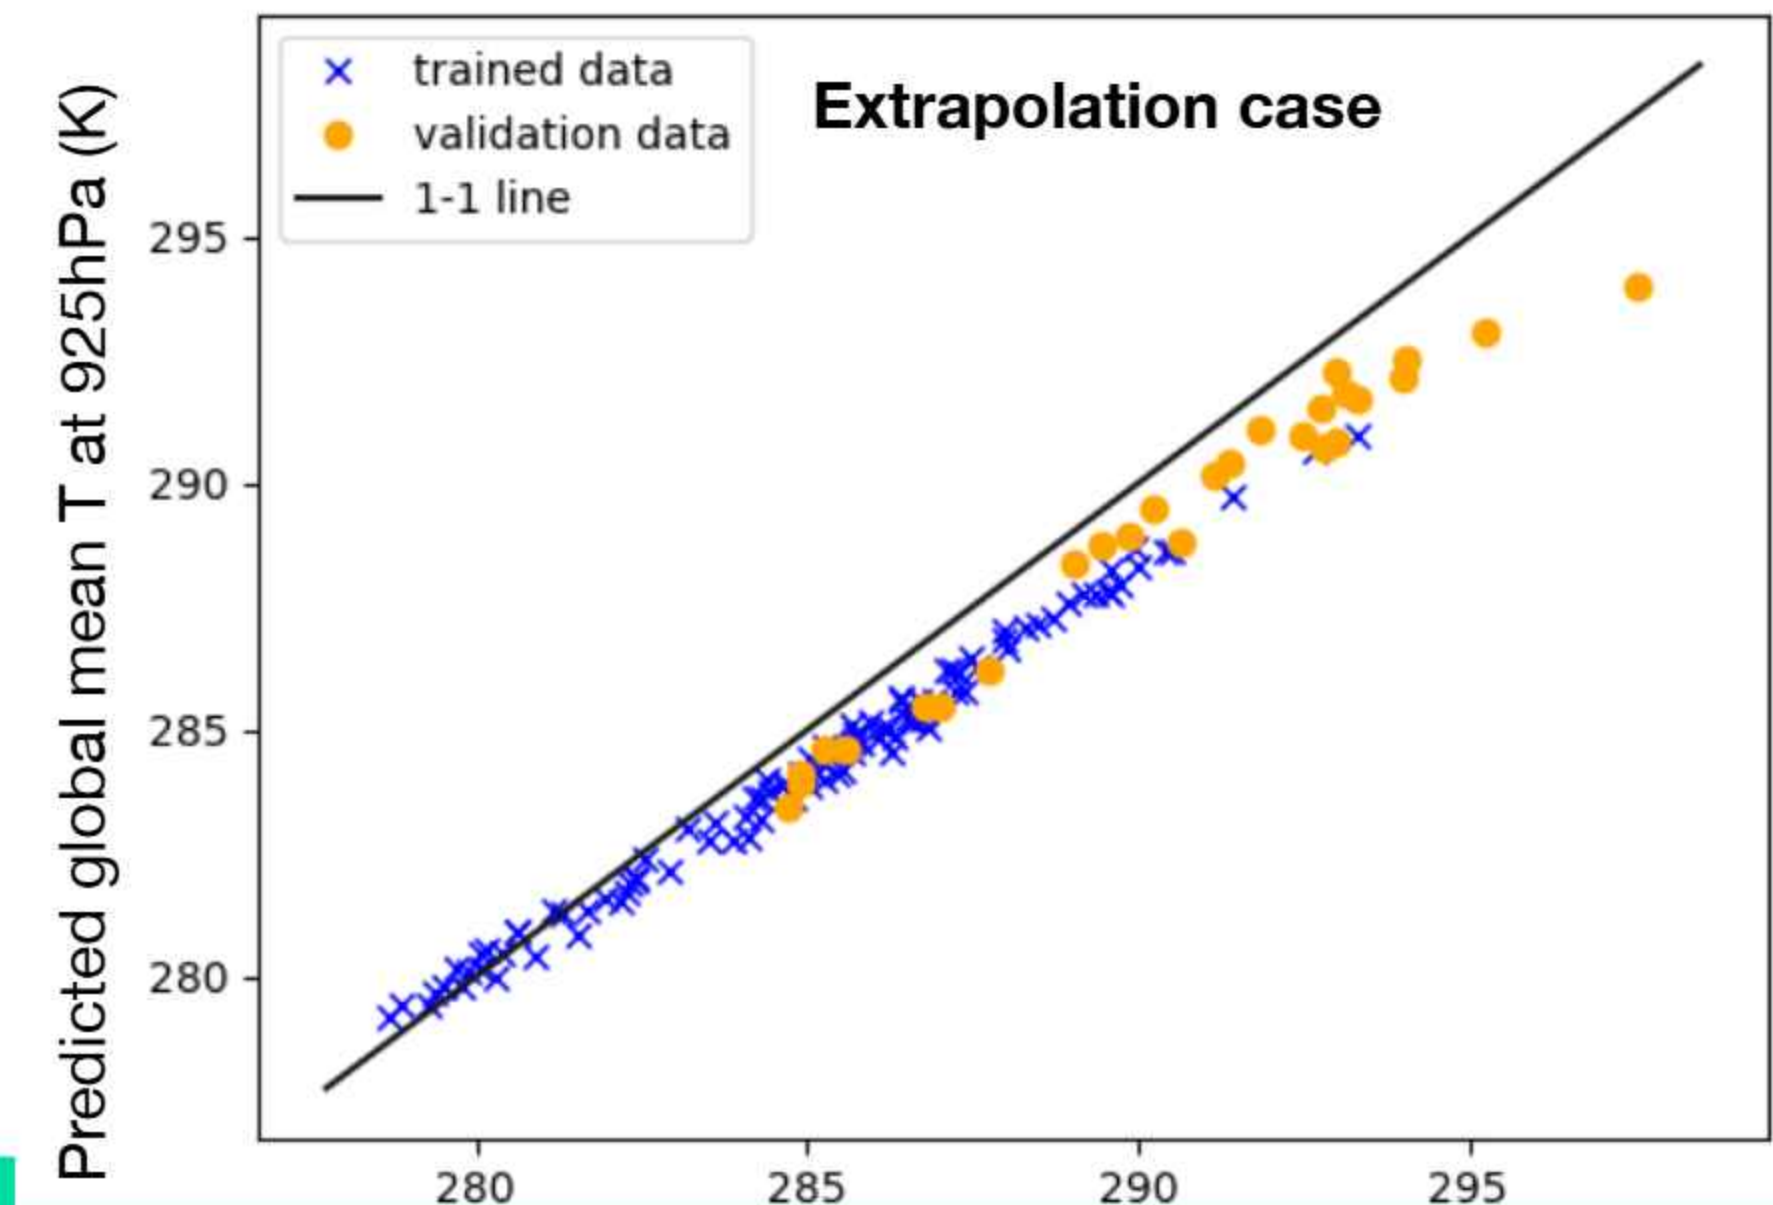
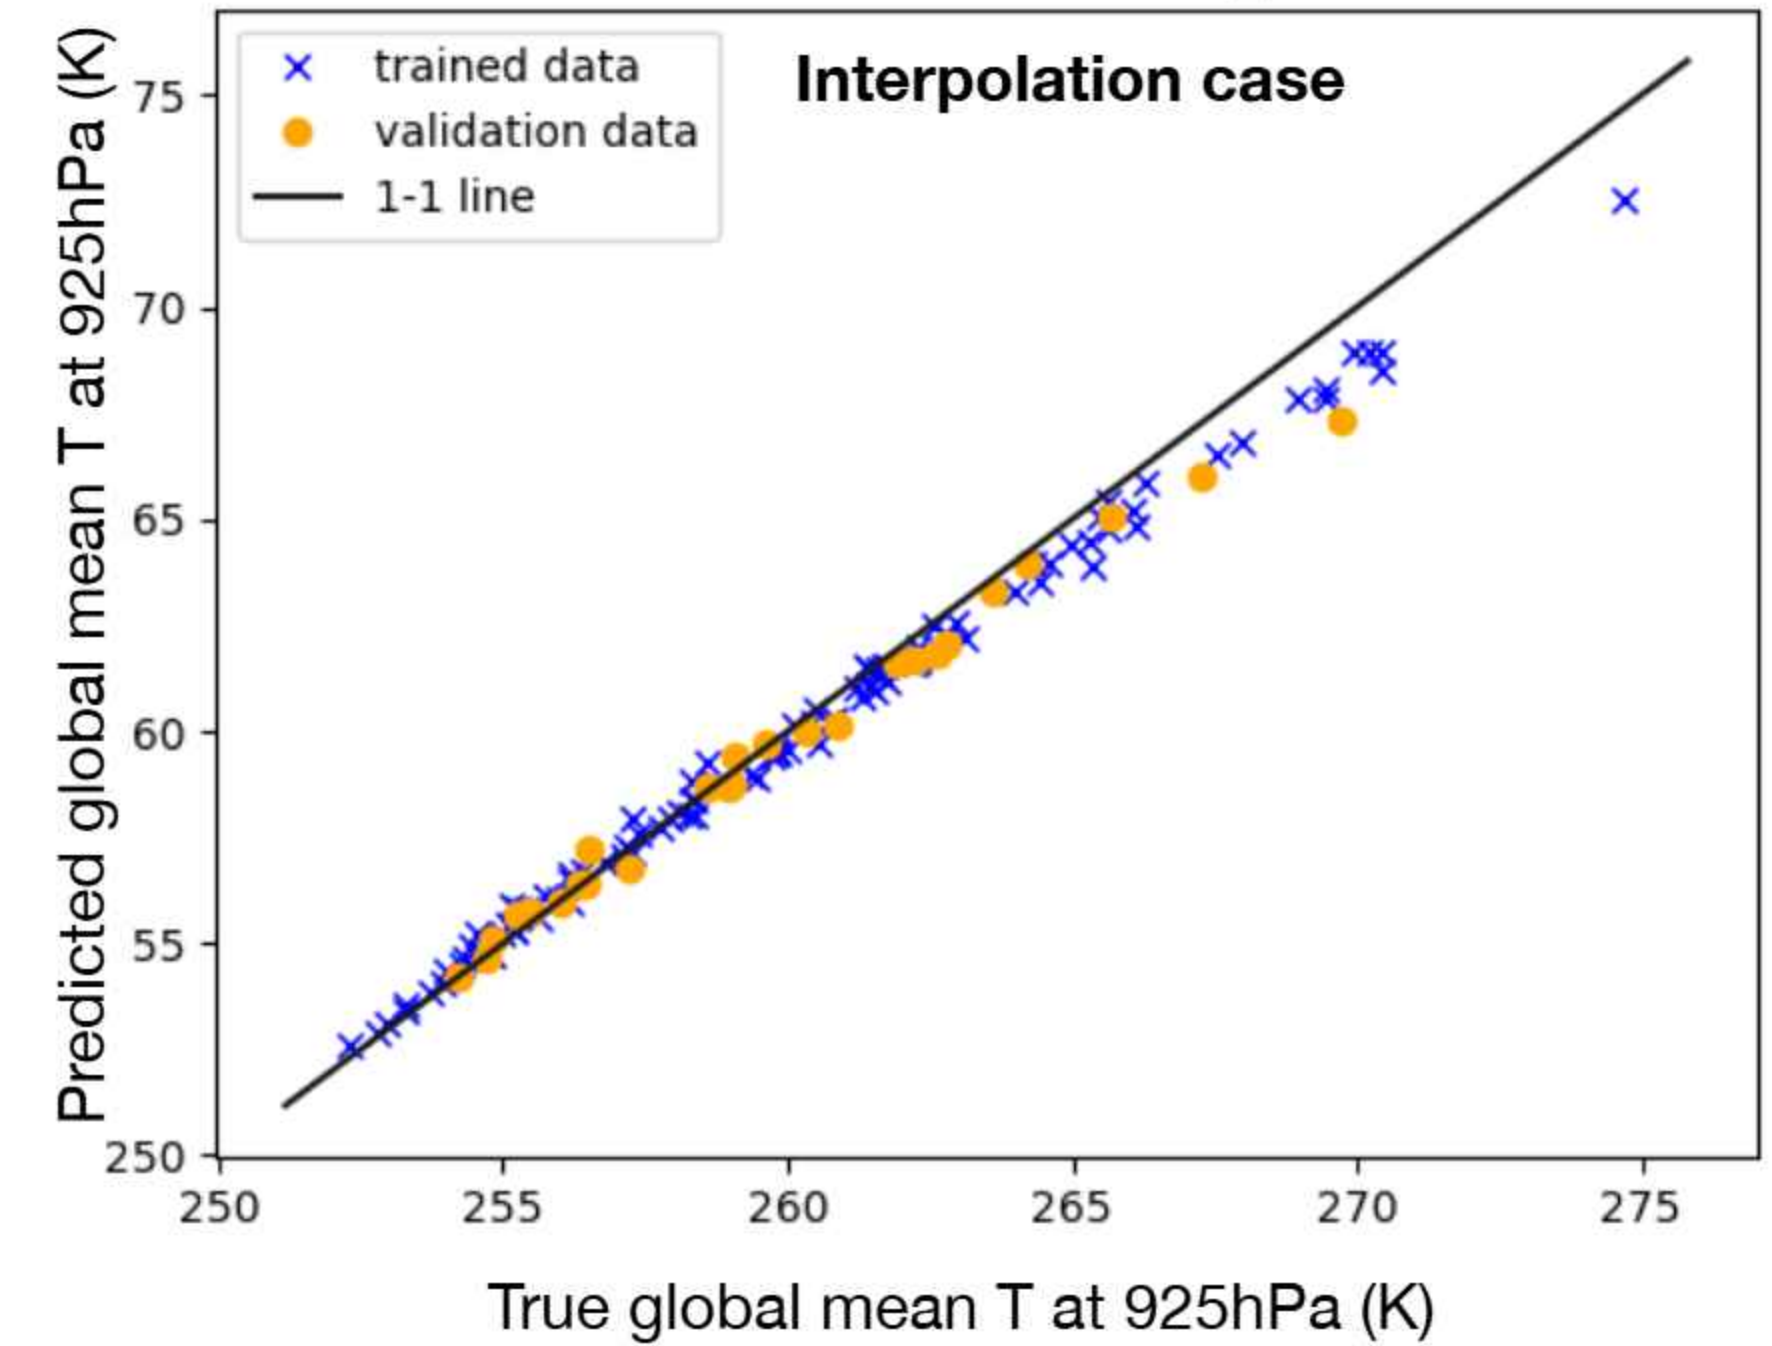
Network performance for interpolation vs extrapolation

- Model performs well at predicting the zonal-wind and temperature when **interpolating**, i.e. when it sees a random subset of the full parameter space in the training phase
- But the model performs significantly worse when it's asked to **extrapolate** to cases that it hasn't seen in training
- For the extrapolation case, we have trained the network on the 278, 335, 400 and 800ppmv CO₂ values, and ask it to predict the 1600ppmv CO₂ cases in validation



Network performance for interpolation vs extrapolation

- Model performs well at predicting the zonal-wind and temperature when **interpolating**, i.e. when it sees a random subset of the full parameter space in the training phase
- But the model performs significantly worse when it's asked to **extrapolate** to cases that it hasn't seen in training
- For the extrapolation case, we have trained the network on the 278, 335, 400 and 800ppmv CO₂ values, and ask it to predict the 1600ppmv CO₂ cases in validation



Incorporating radiative constraints...

Incorporating radiative constraints...

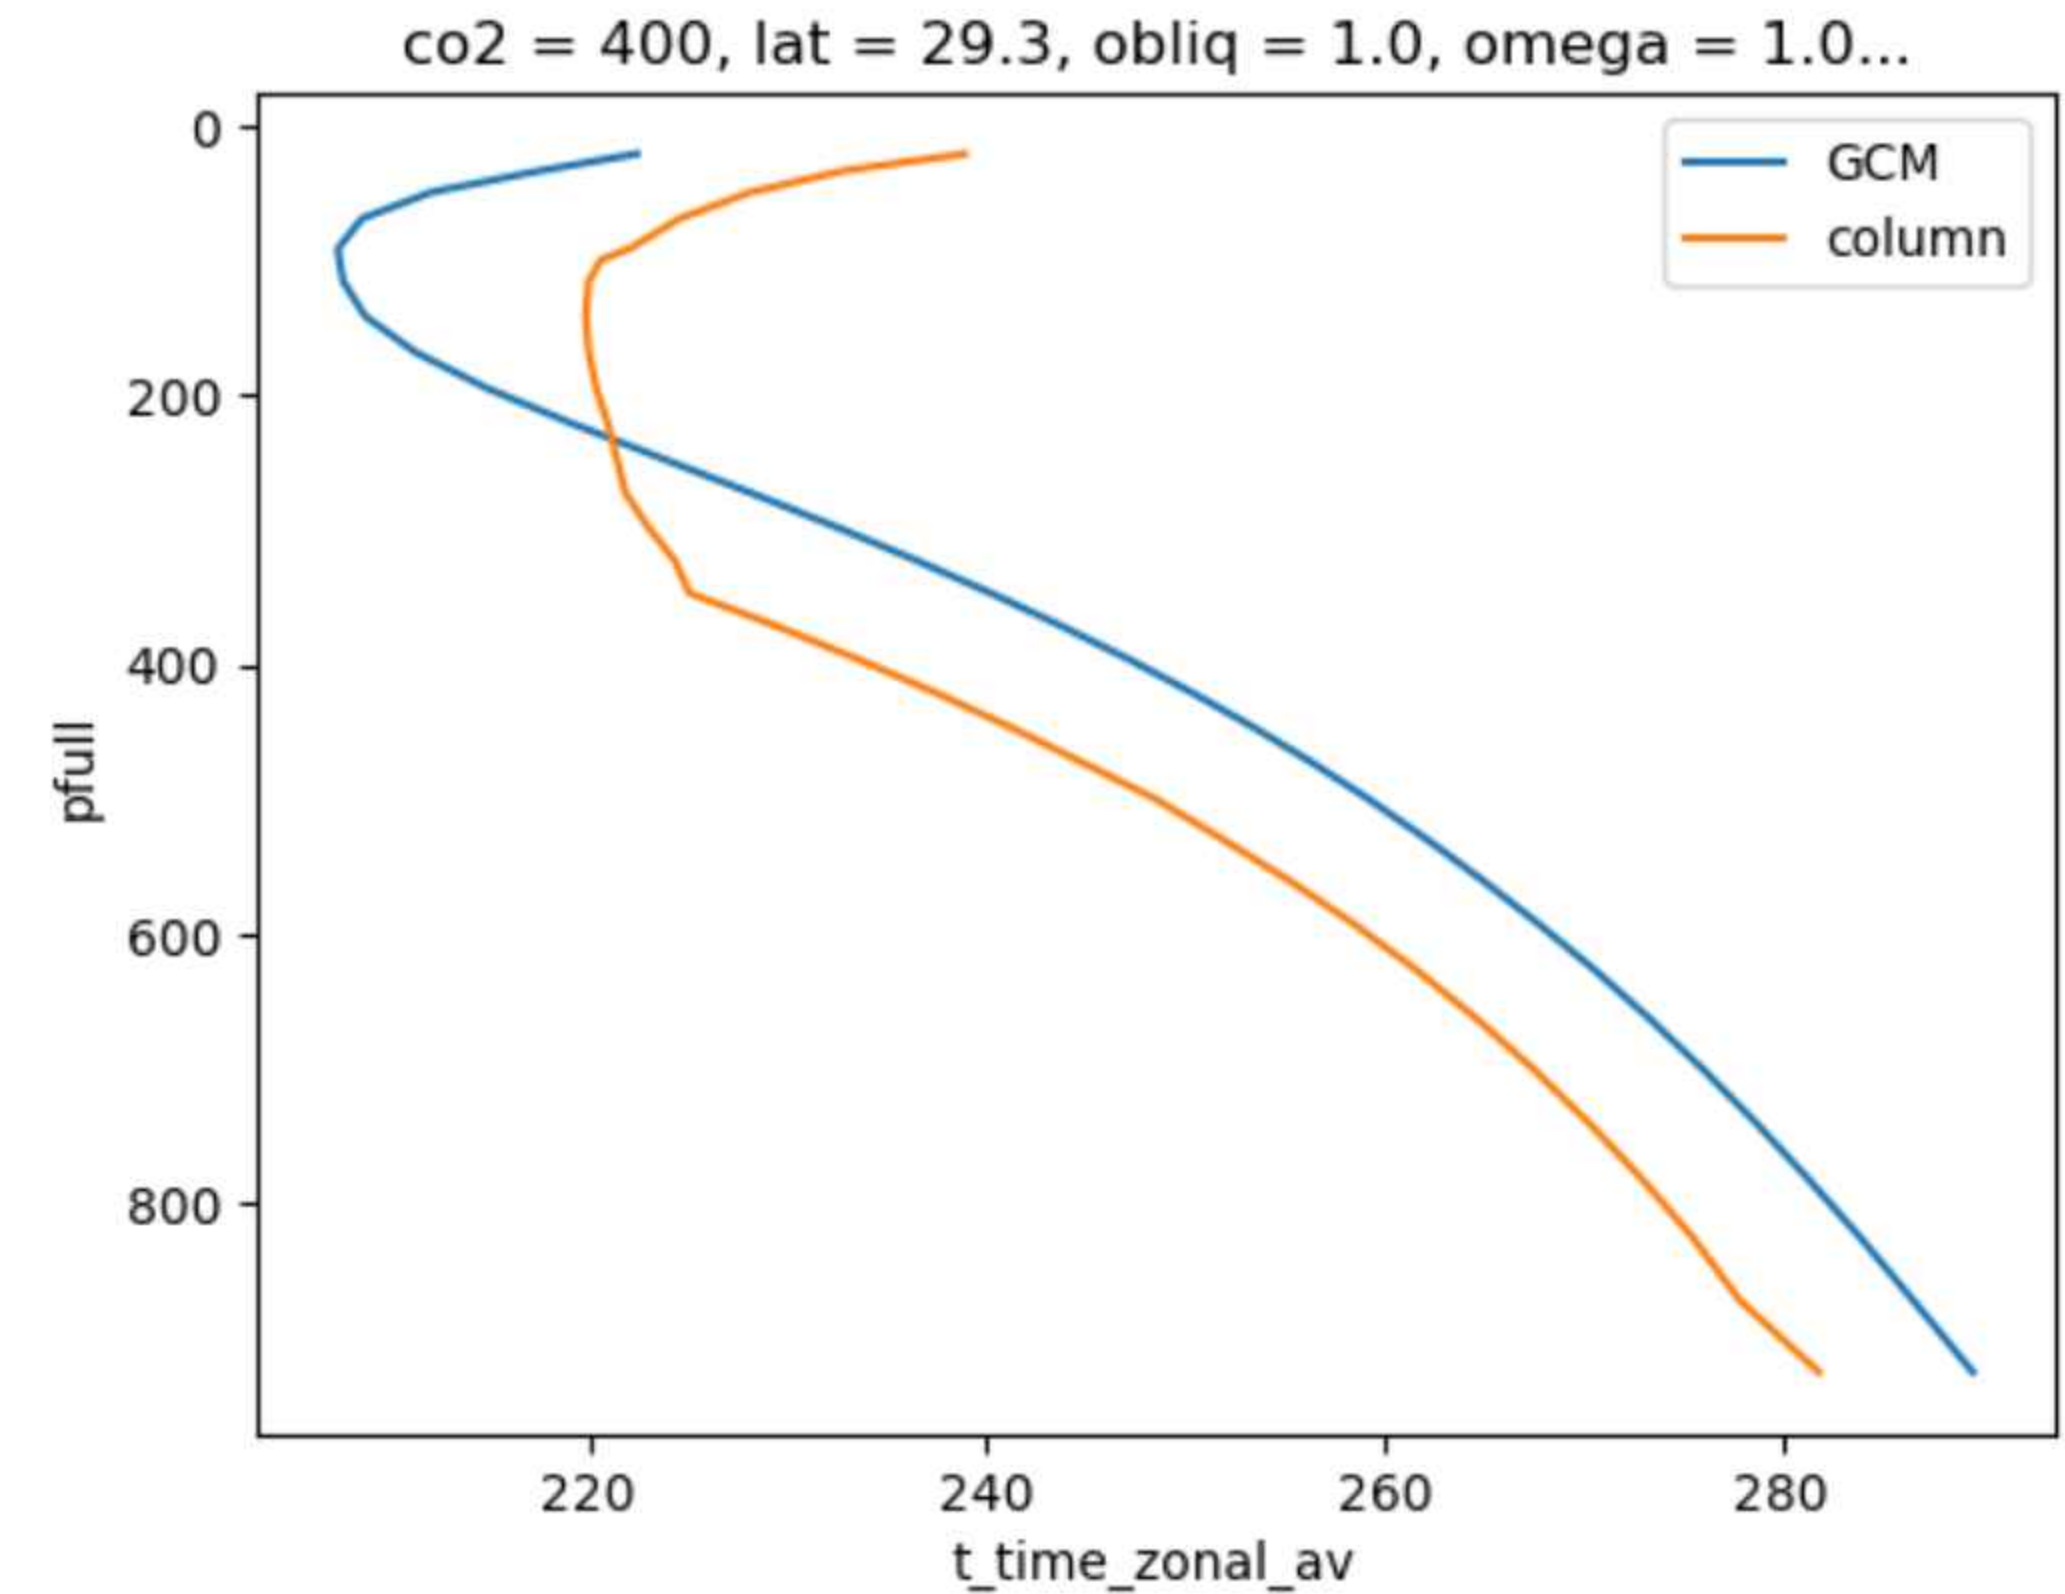
- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca

Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium

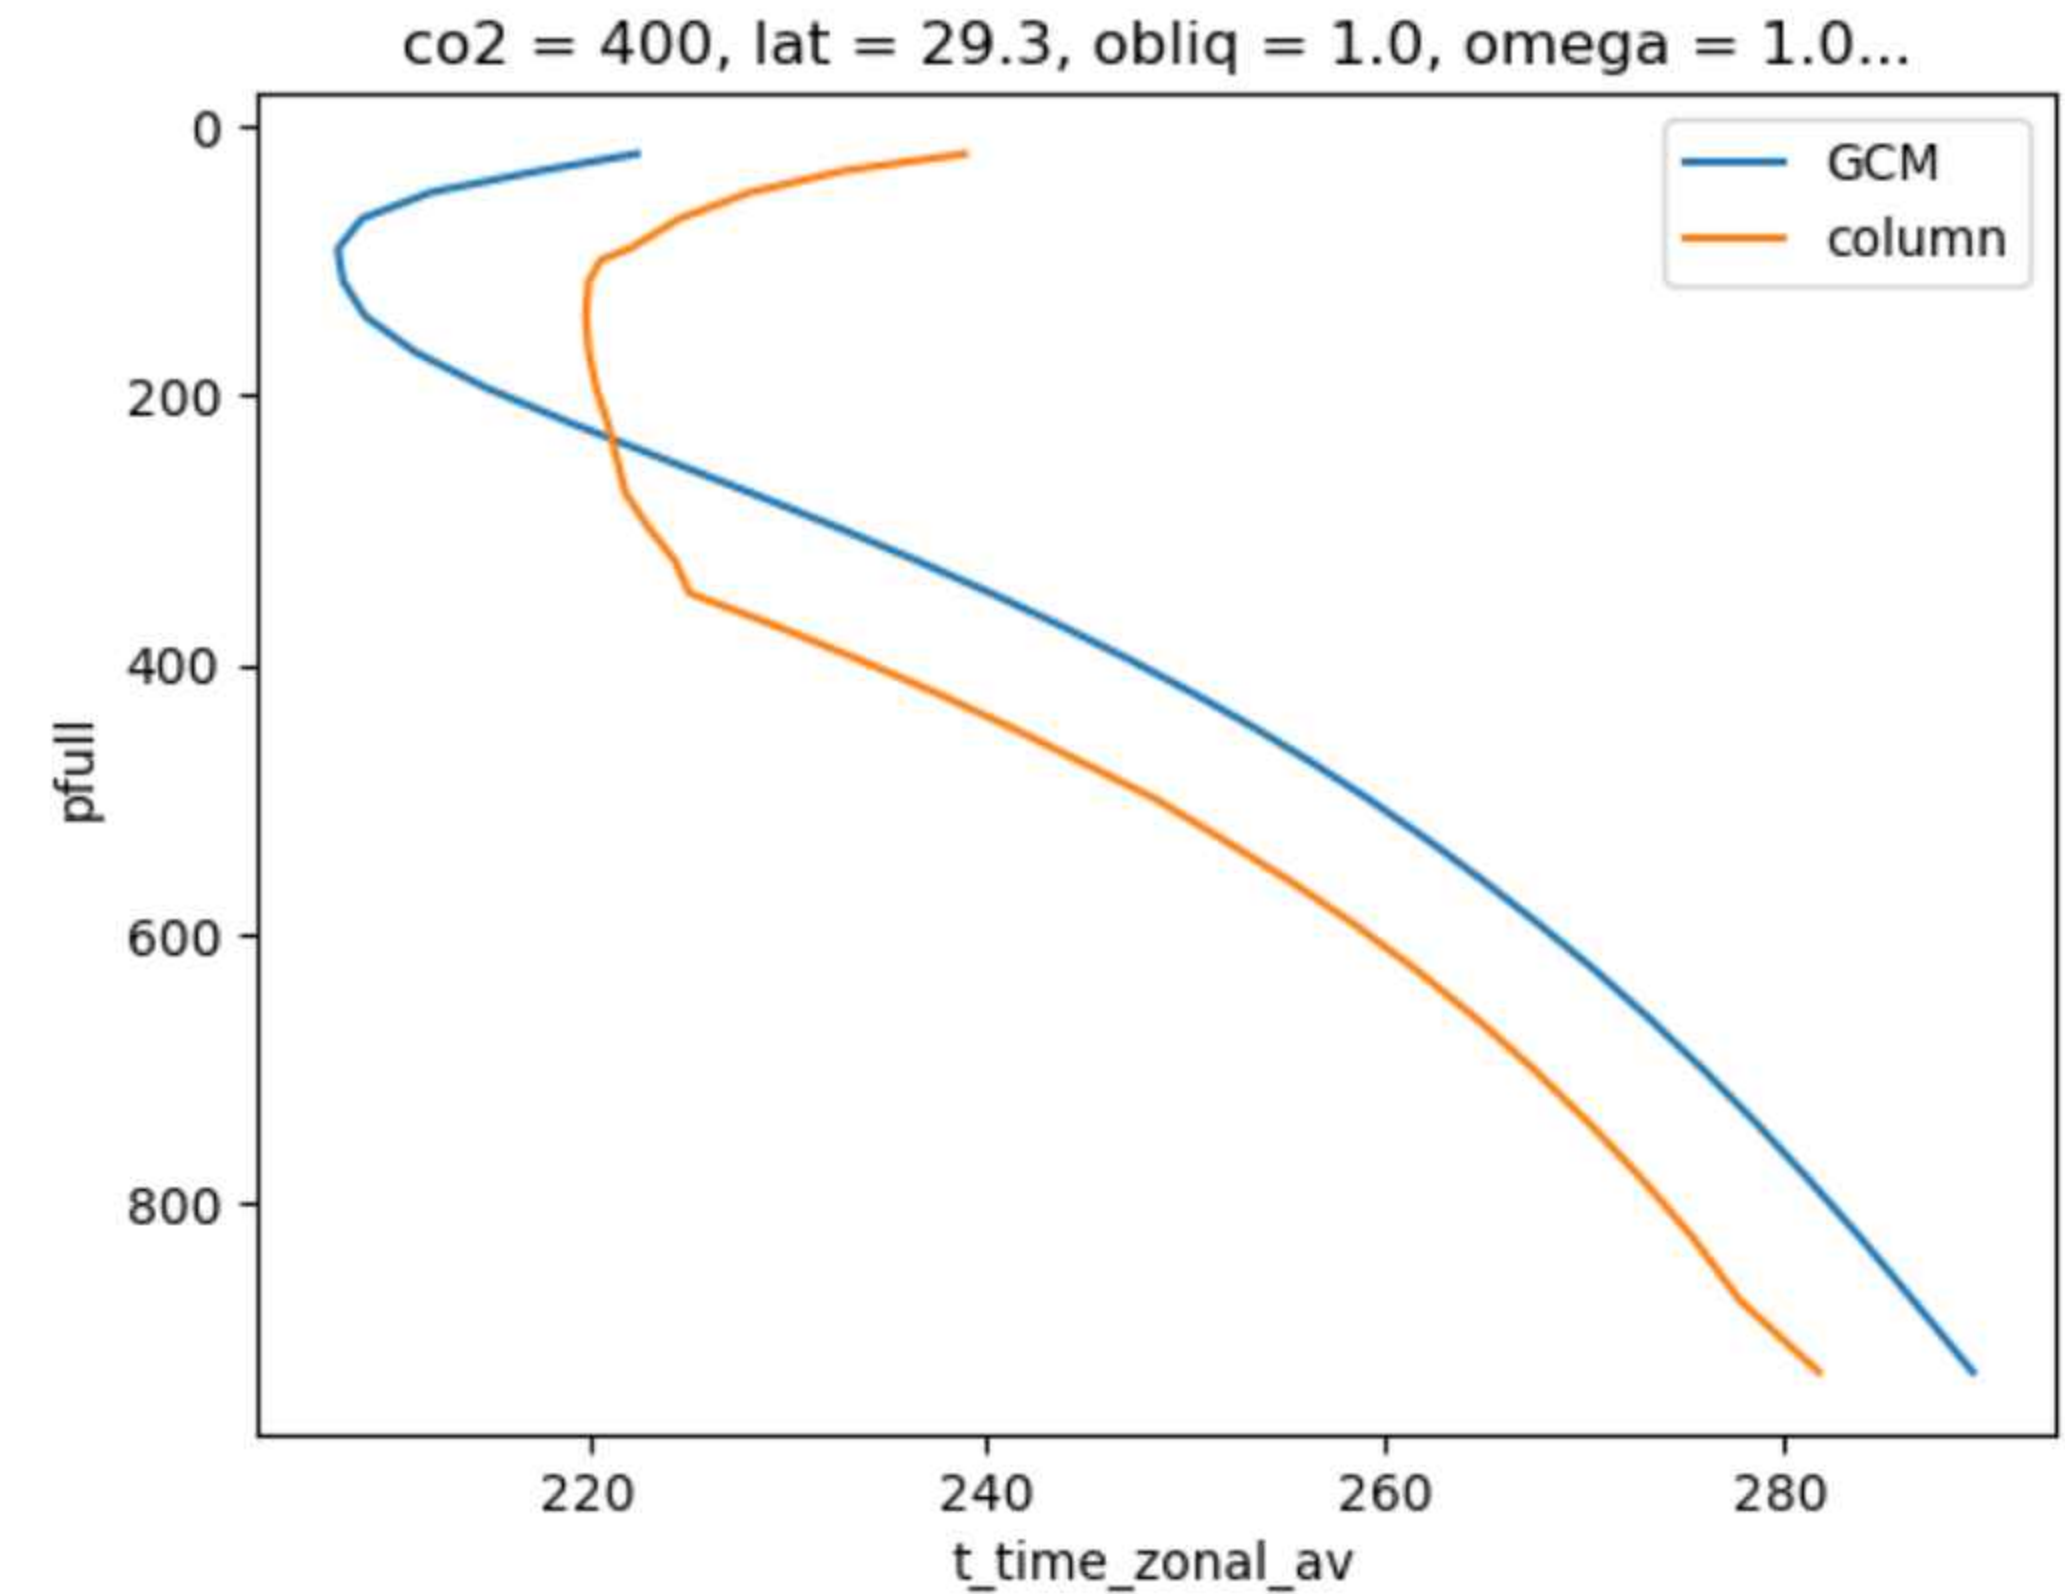
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium



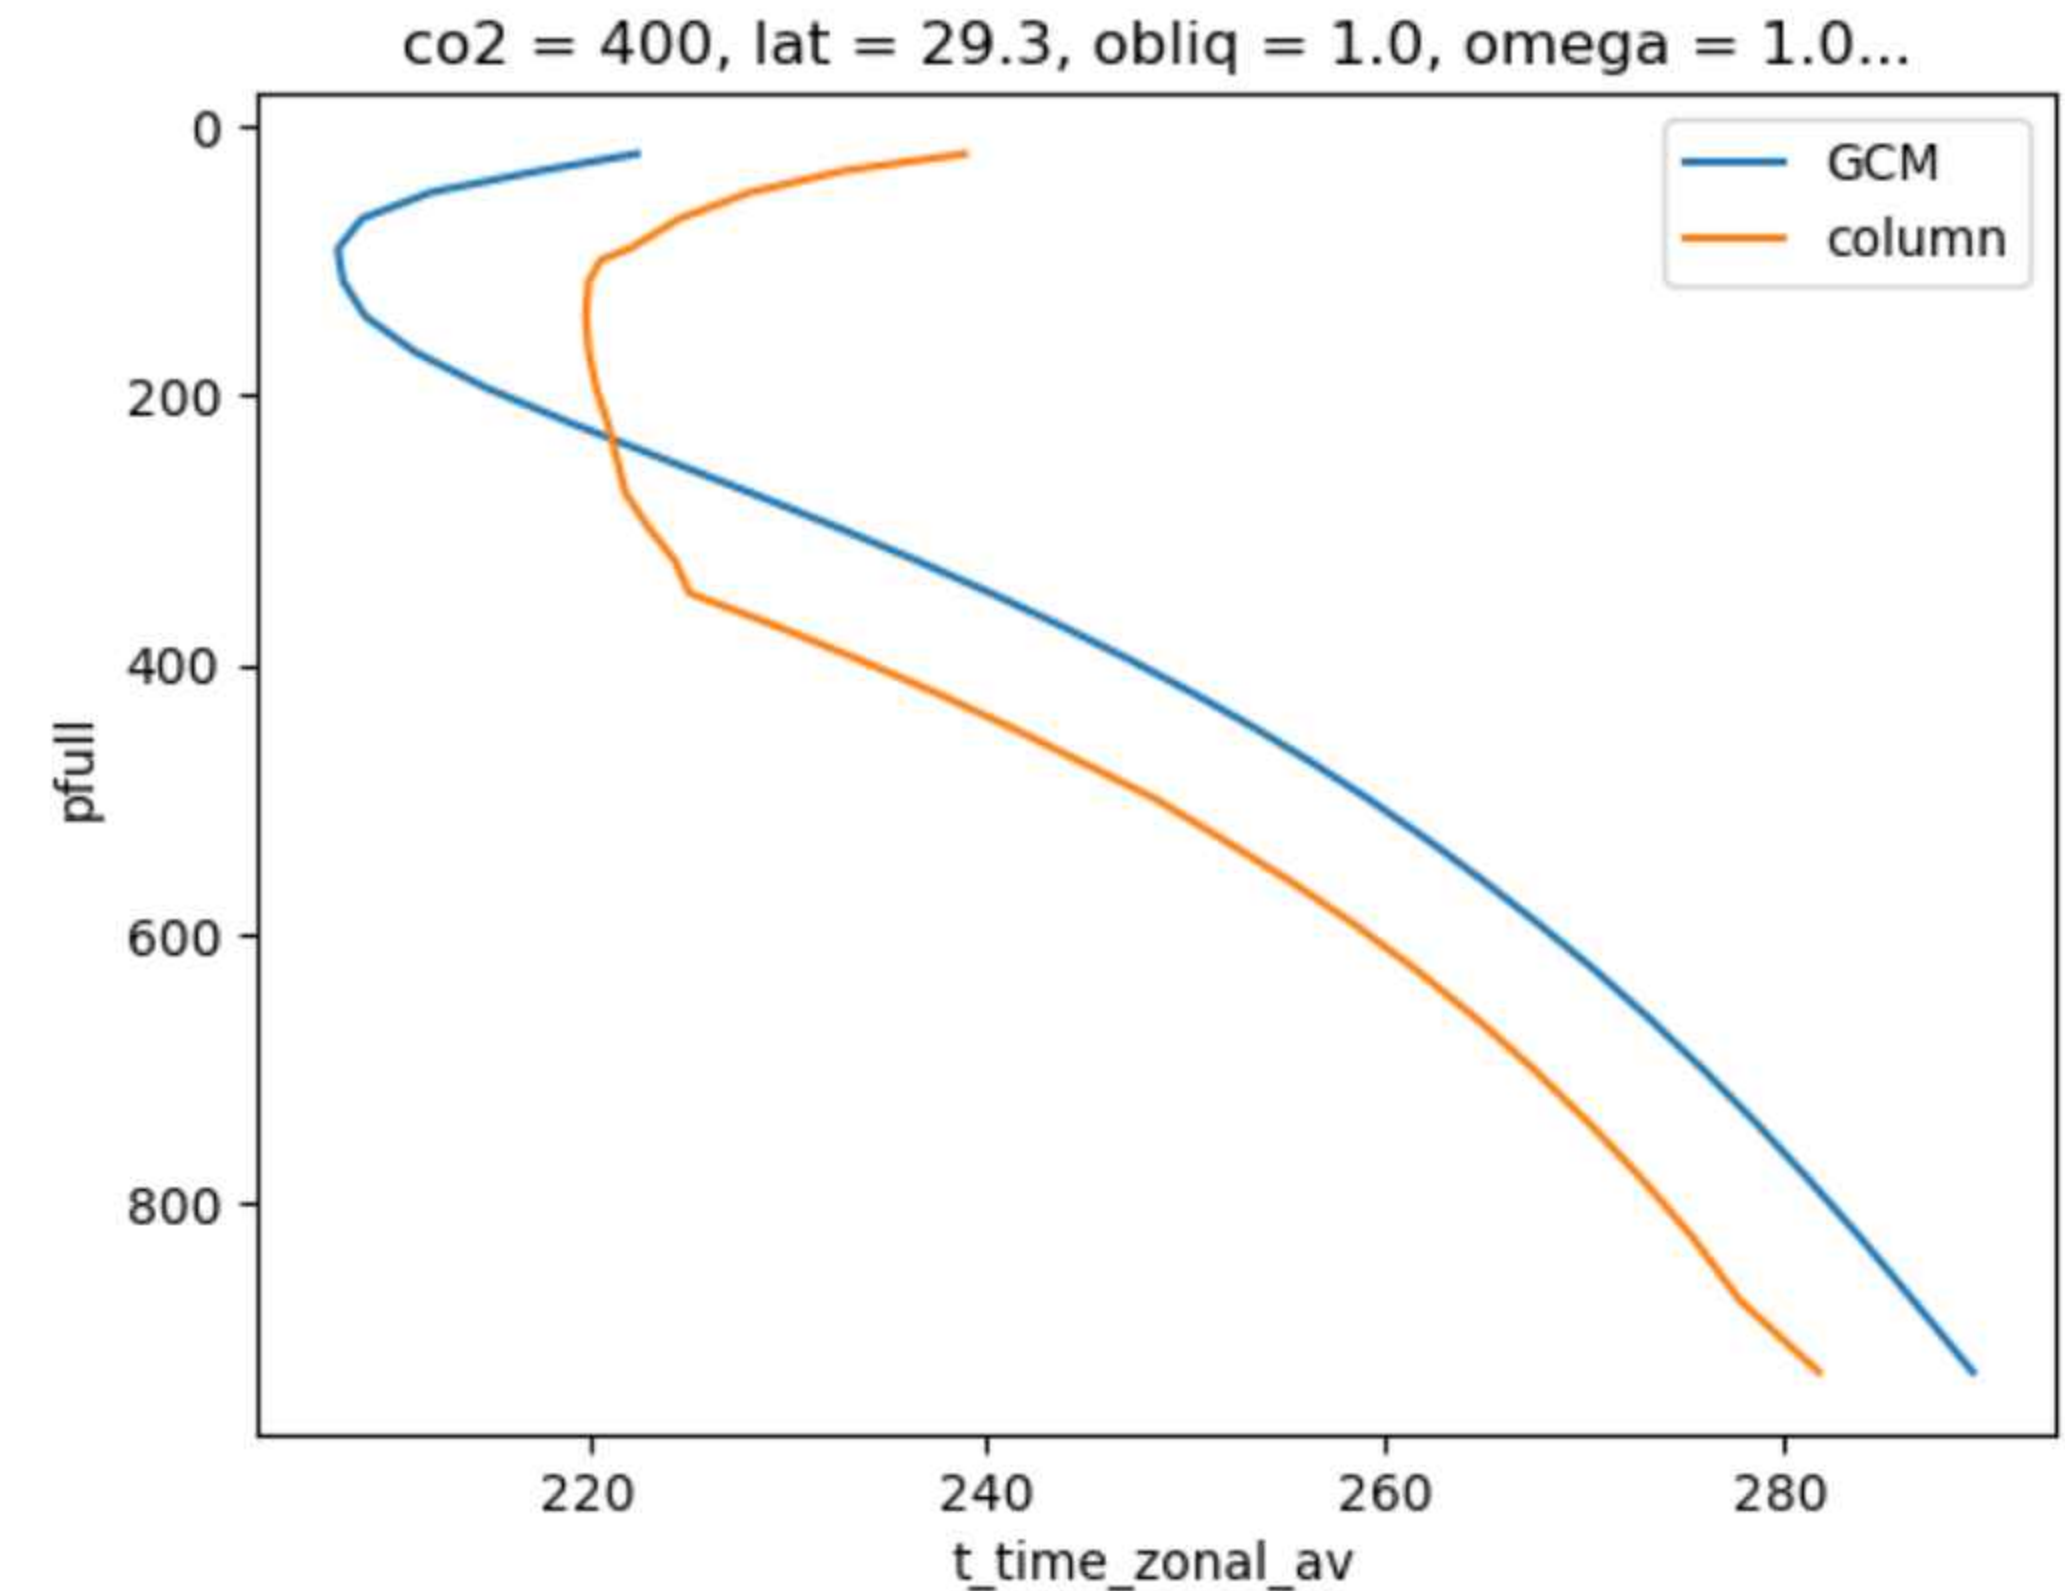
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model



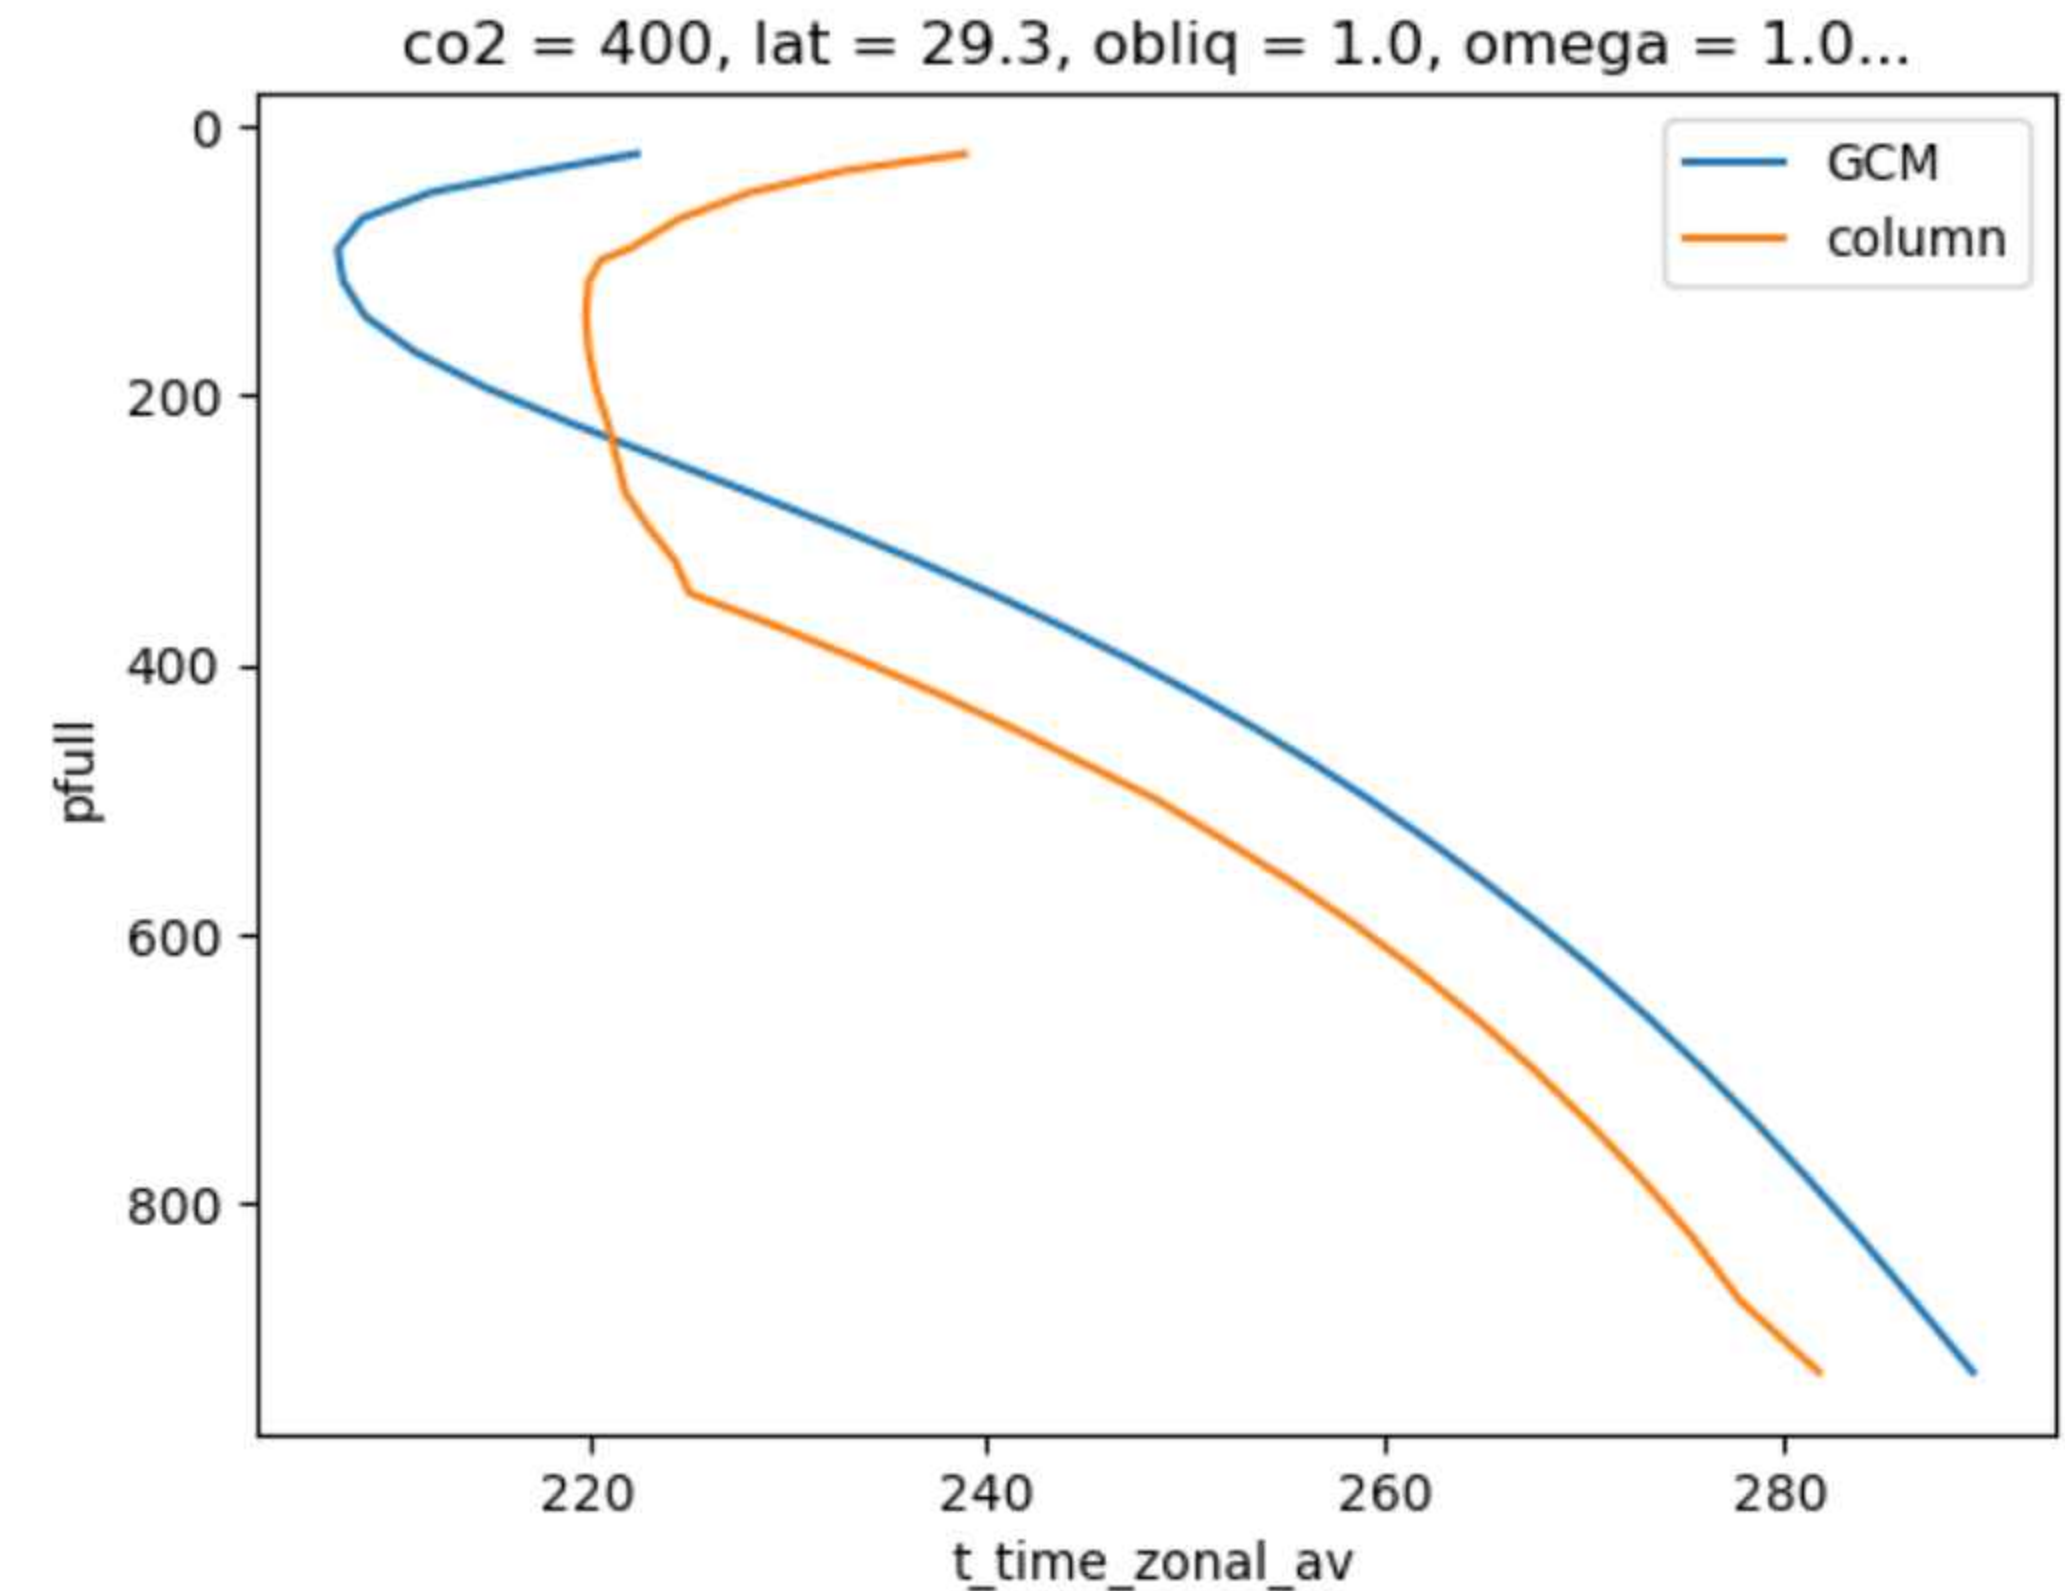
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model
- How can we incorporate the column model results?



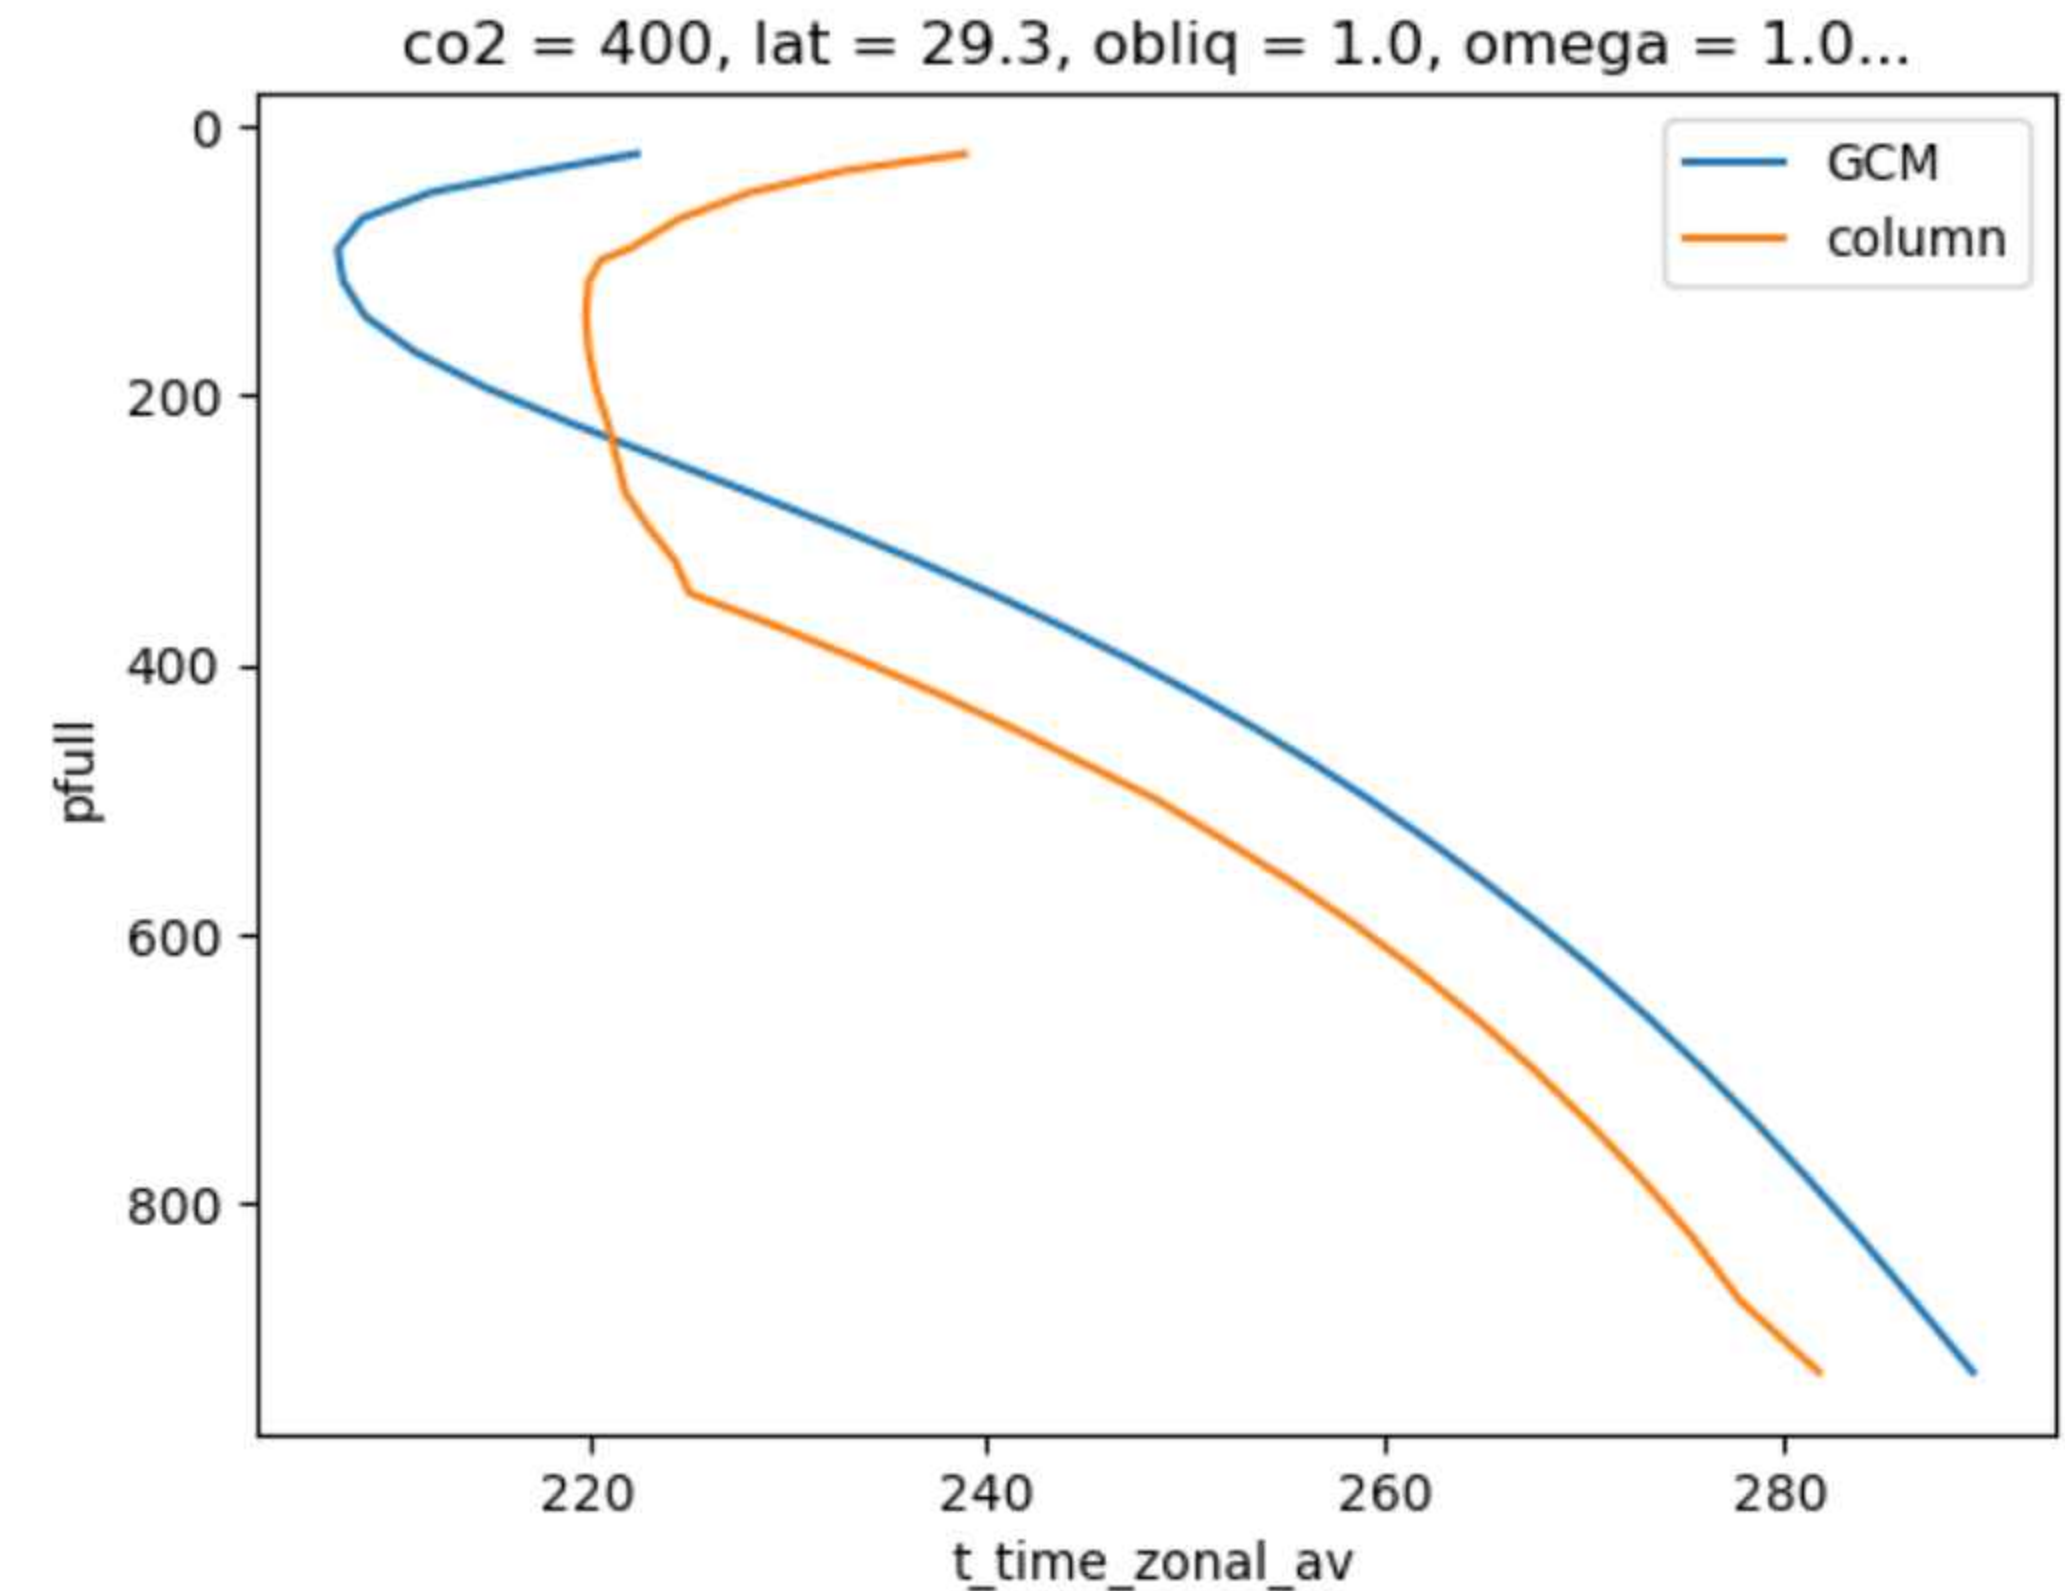
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model
- How can we incorporate the column model results?
 - Predict departure of temperature from column model - predict how far away the state is from radiative-convective equilibrium - **still extrapolation problem**



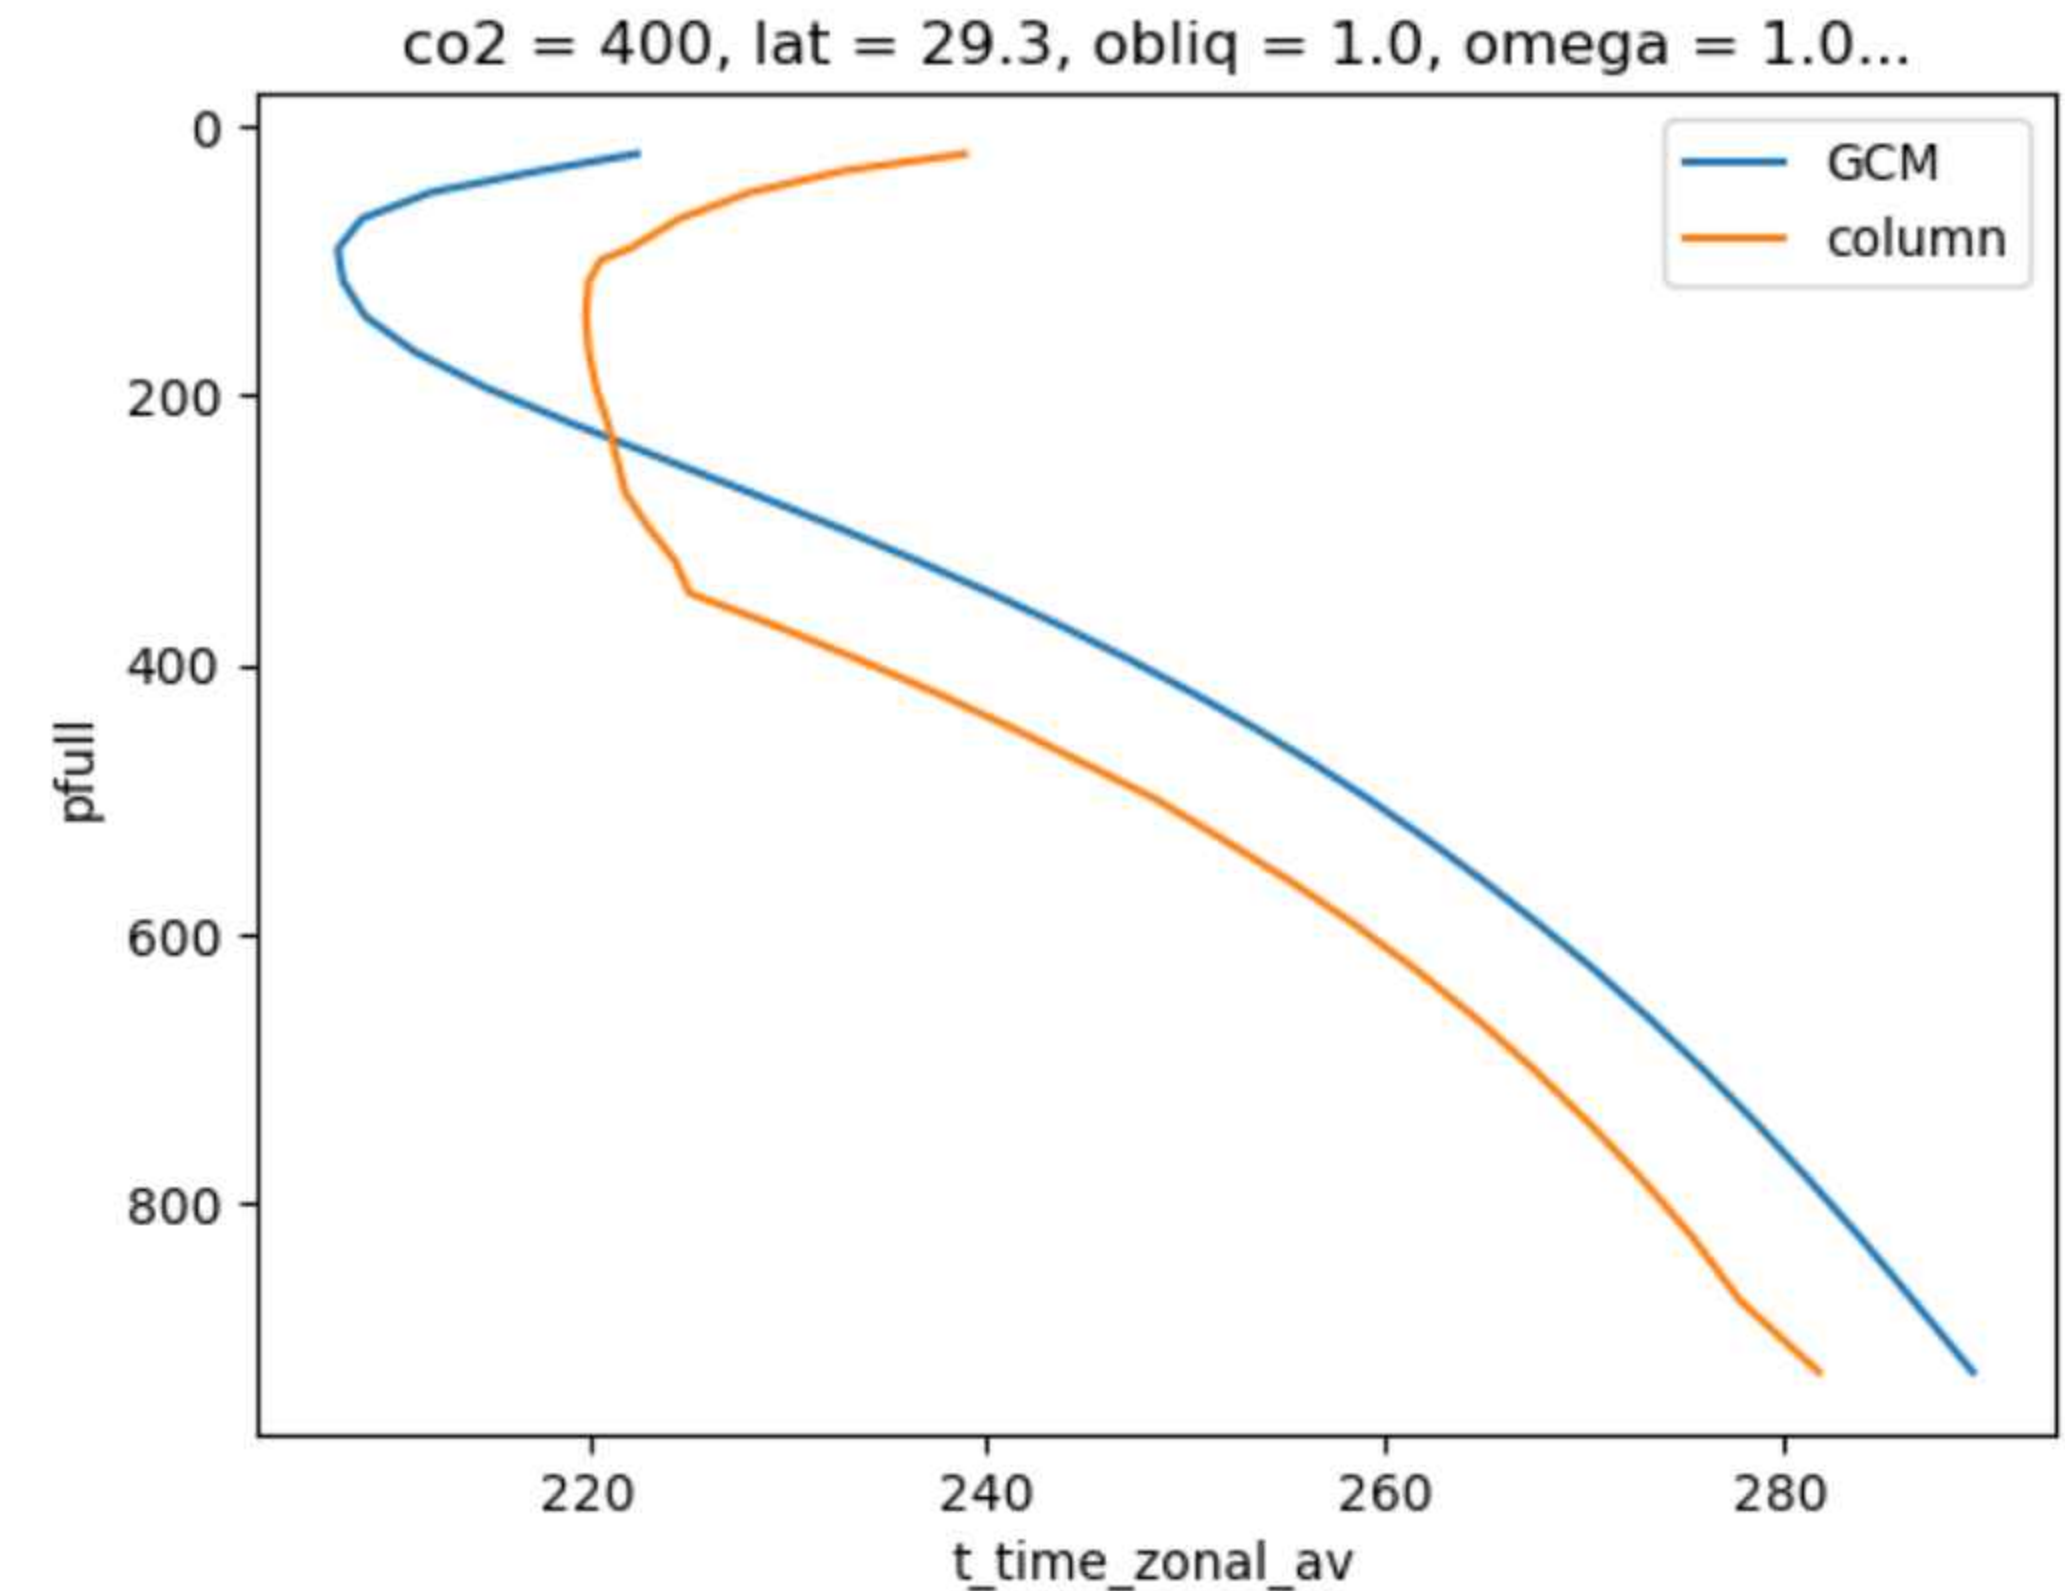
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model
- How can we incorporate the column model results?
 - Predict departure of temperature from column model - predict how far away the state is from radiative-convective equilibrium - **still extrapolation problem**
 - Incorporate the column model data into the network's **training...**



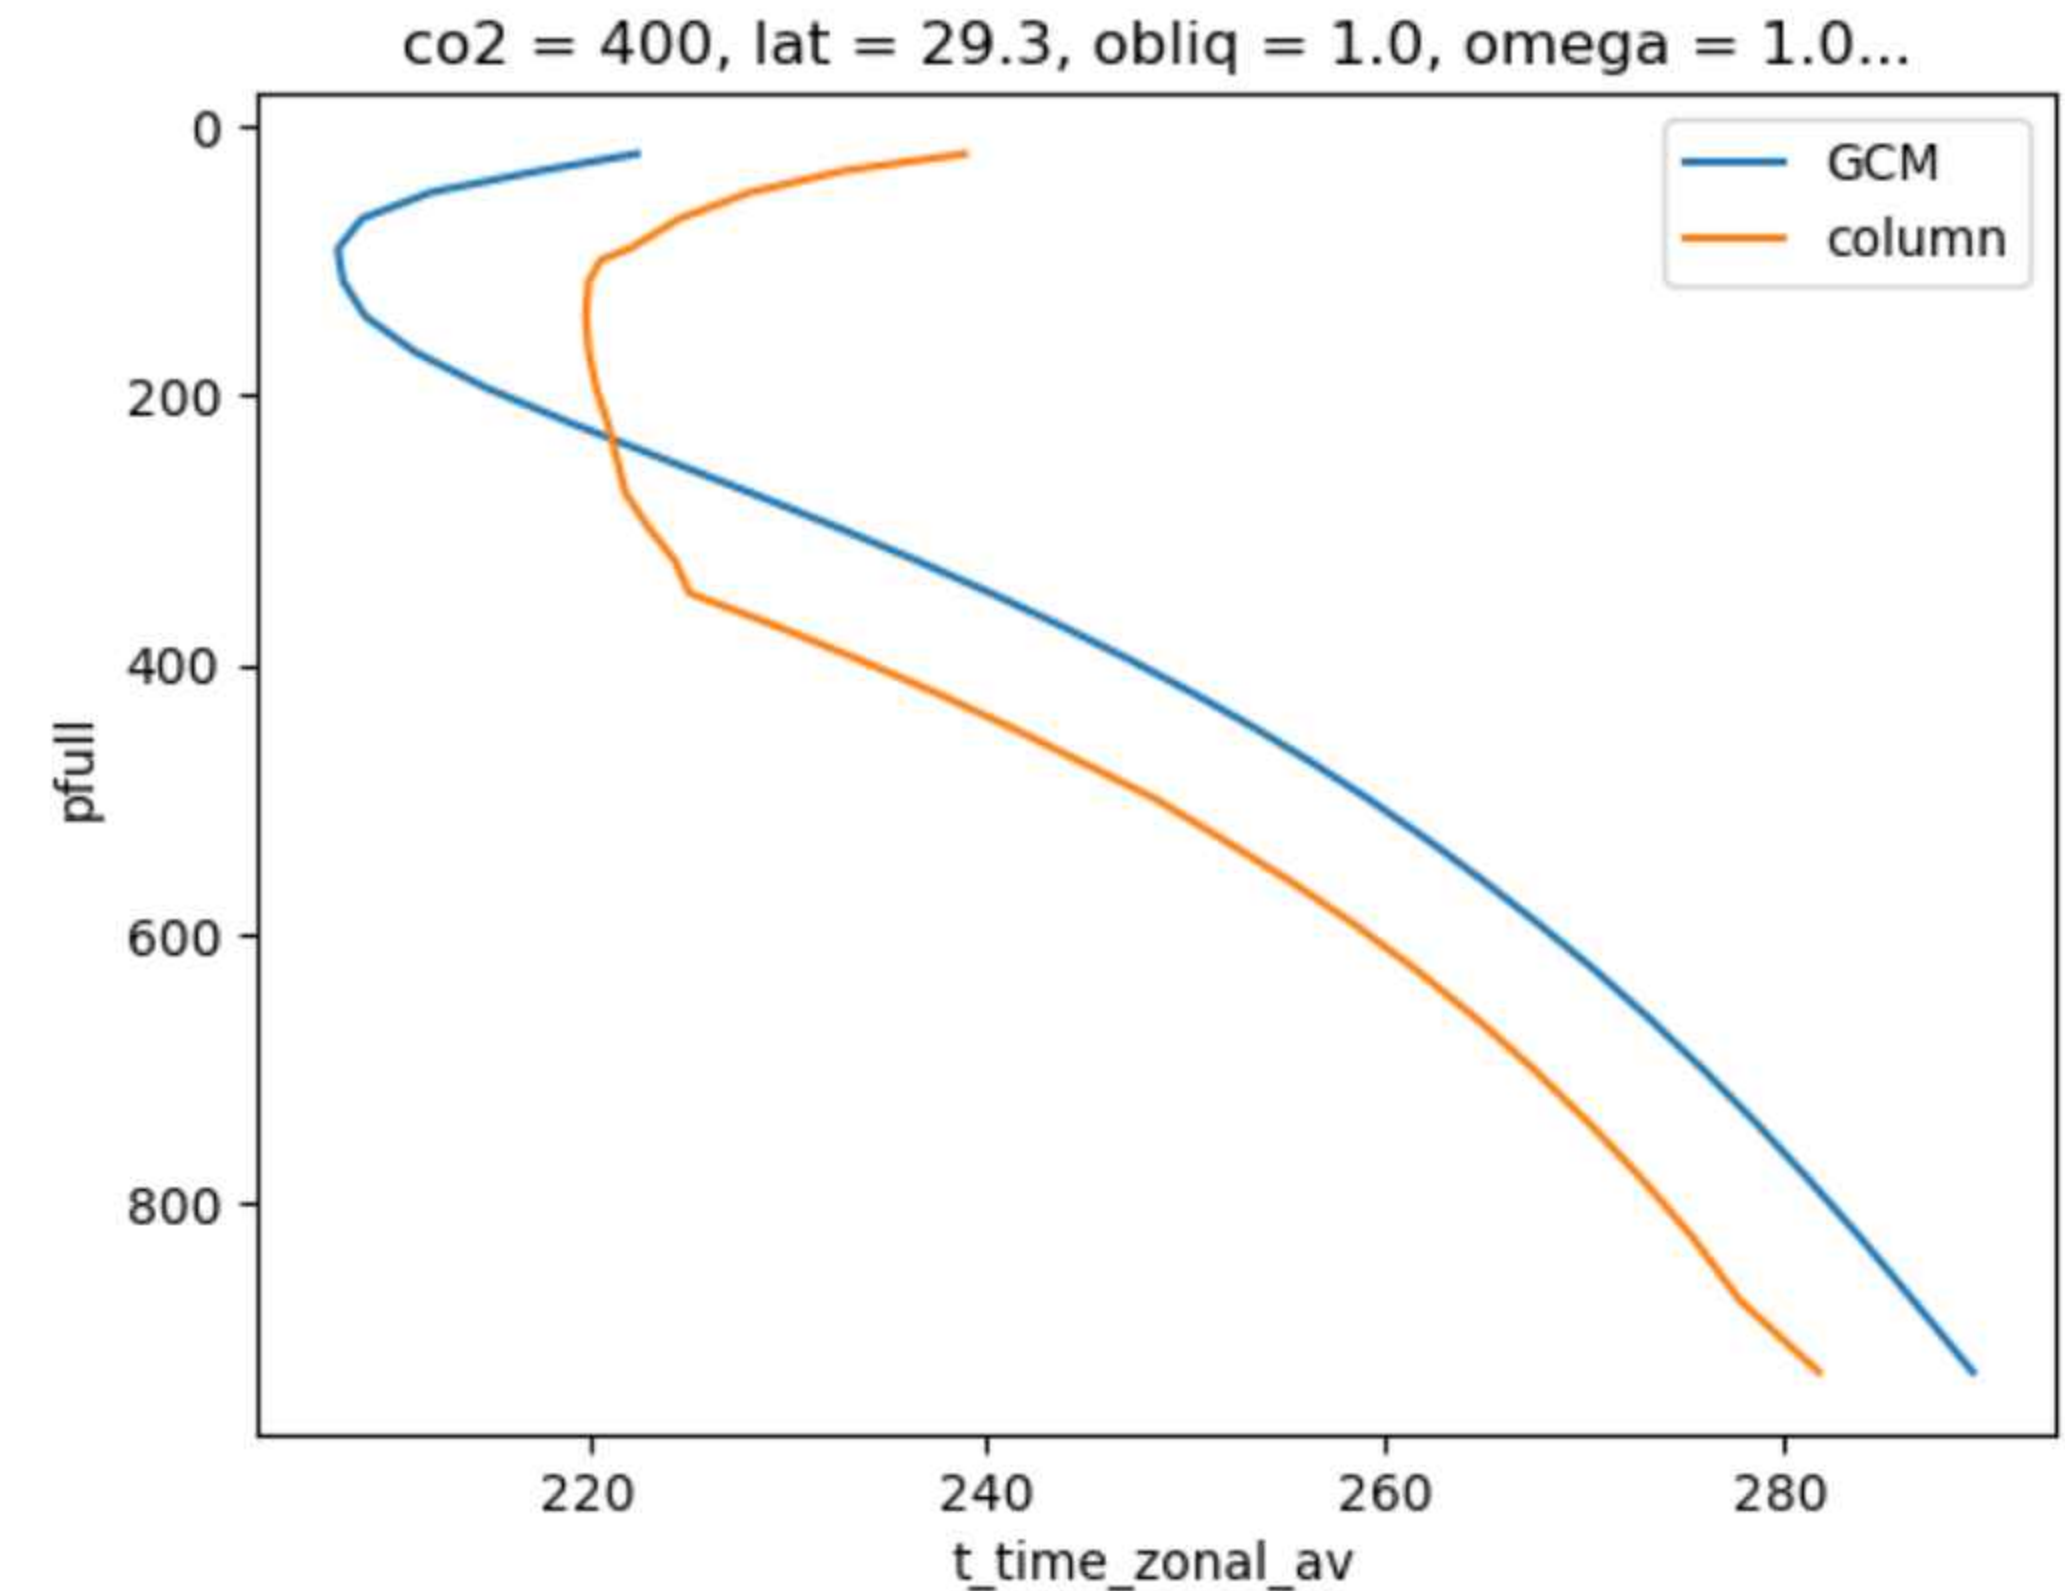
Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model
- How can we incorporate the column model results?
 - Predict departure of temperature from column model - predict how far away the state is from radiative-convective equilibrium - **still extrapolation problem**
 - Incorporate the column model data into the network's **training...**
- **We want ML to learn about things that are difficult for us to predict cheaply (i.e. dynamics etc)**

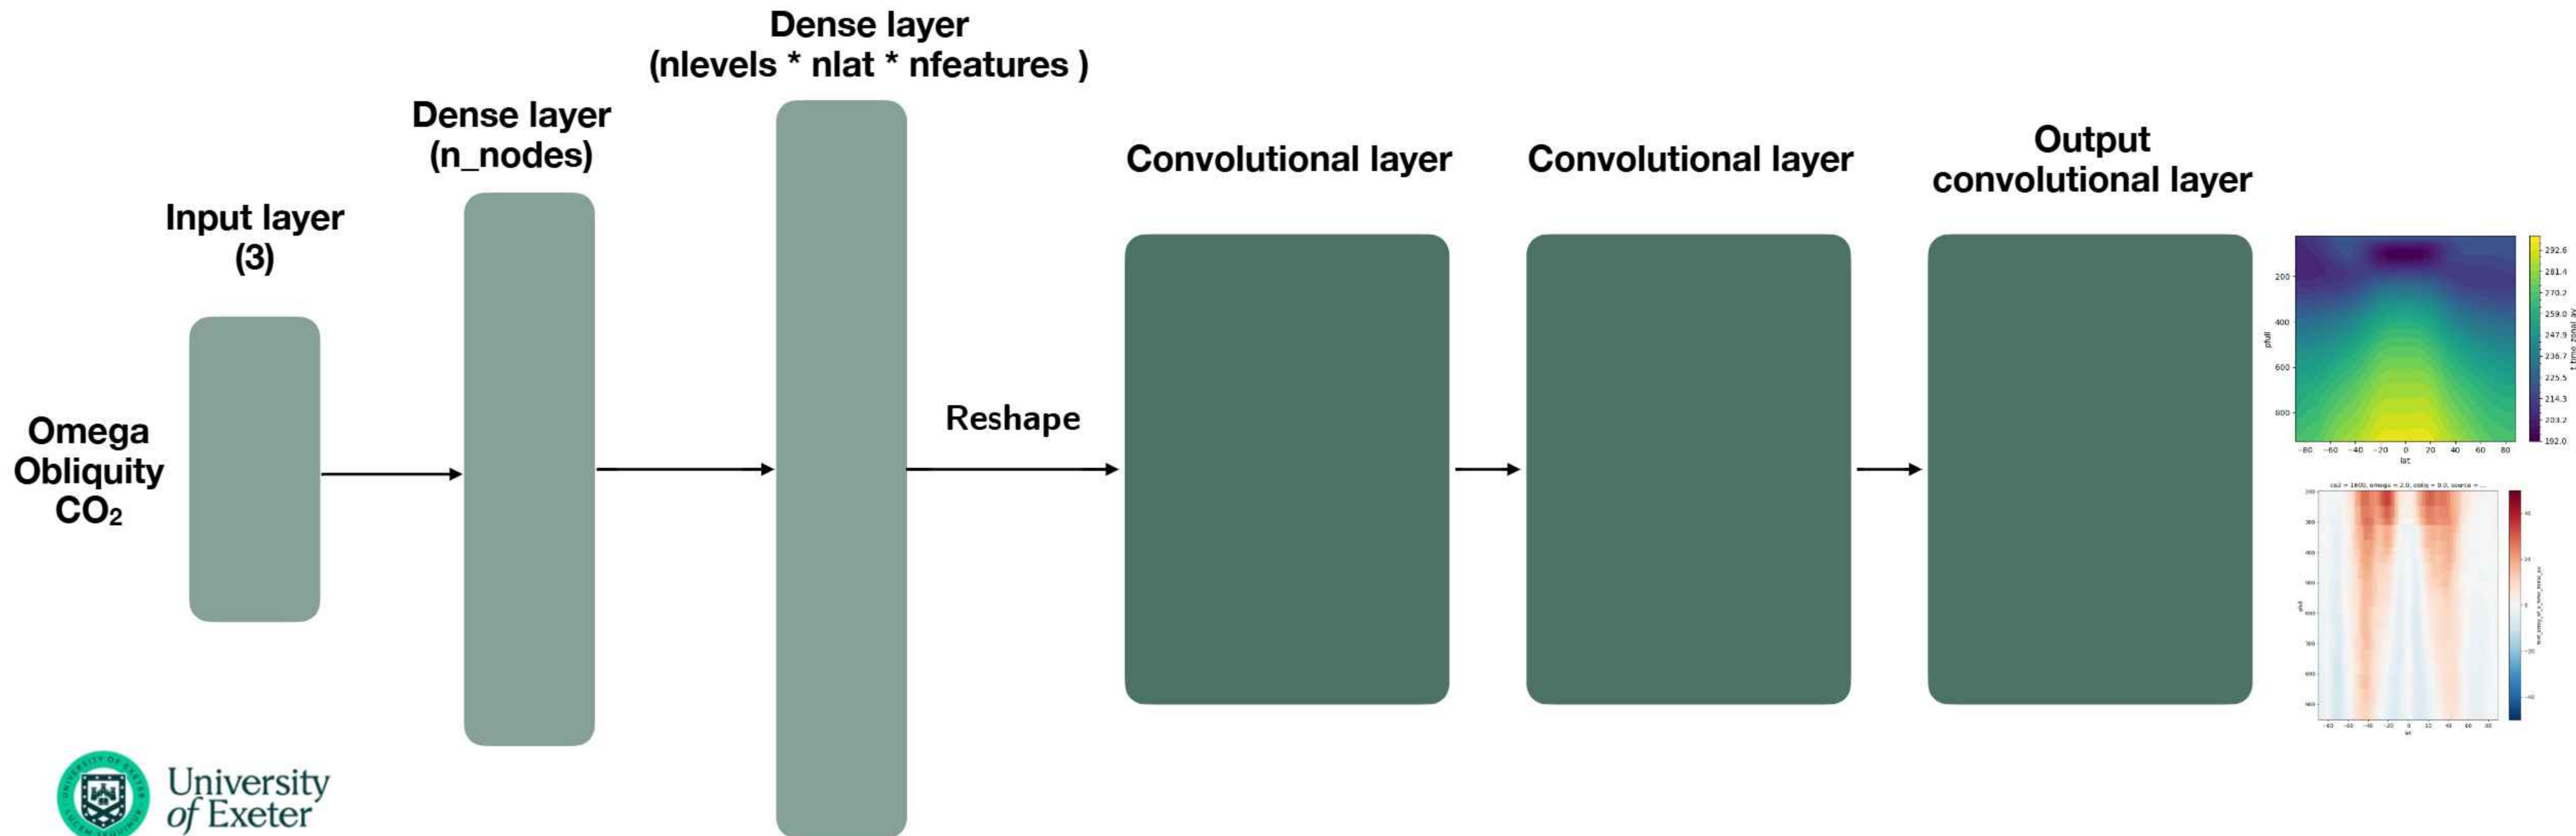


Incorporating radiative constraints...

- In order to incorporate radiative constraints, we have conducted a further 5 experiments using the column-model version of Isca
 - Single column without any fluid dynamics - radiative convective equilibrium
 - Column model is $\mathcal{O}(100)$ times cheaper than the full model
- How can we incorporate the column model results?
 - Predict departure of temperature from column model - predict how far away the state is from radiative-convective equilibrium - **still extrapolation problem**
 - Incorporate the column model data into the network's **training...**
- **We want ML to learn about things that are difficult for us to predict cheaply (i.e. dynamics etc)**
- **Radiative-convective equilibrium is very cheap compared with physical models, so less of a need to have the ML learn the details of this**

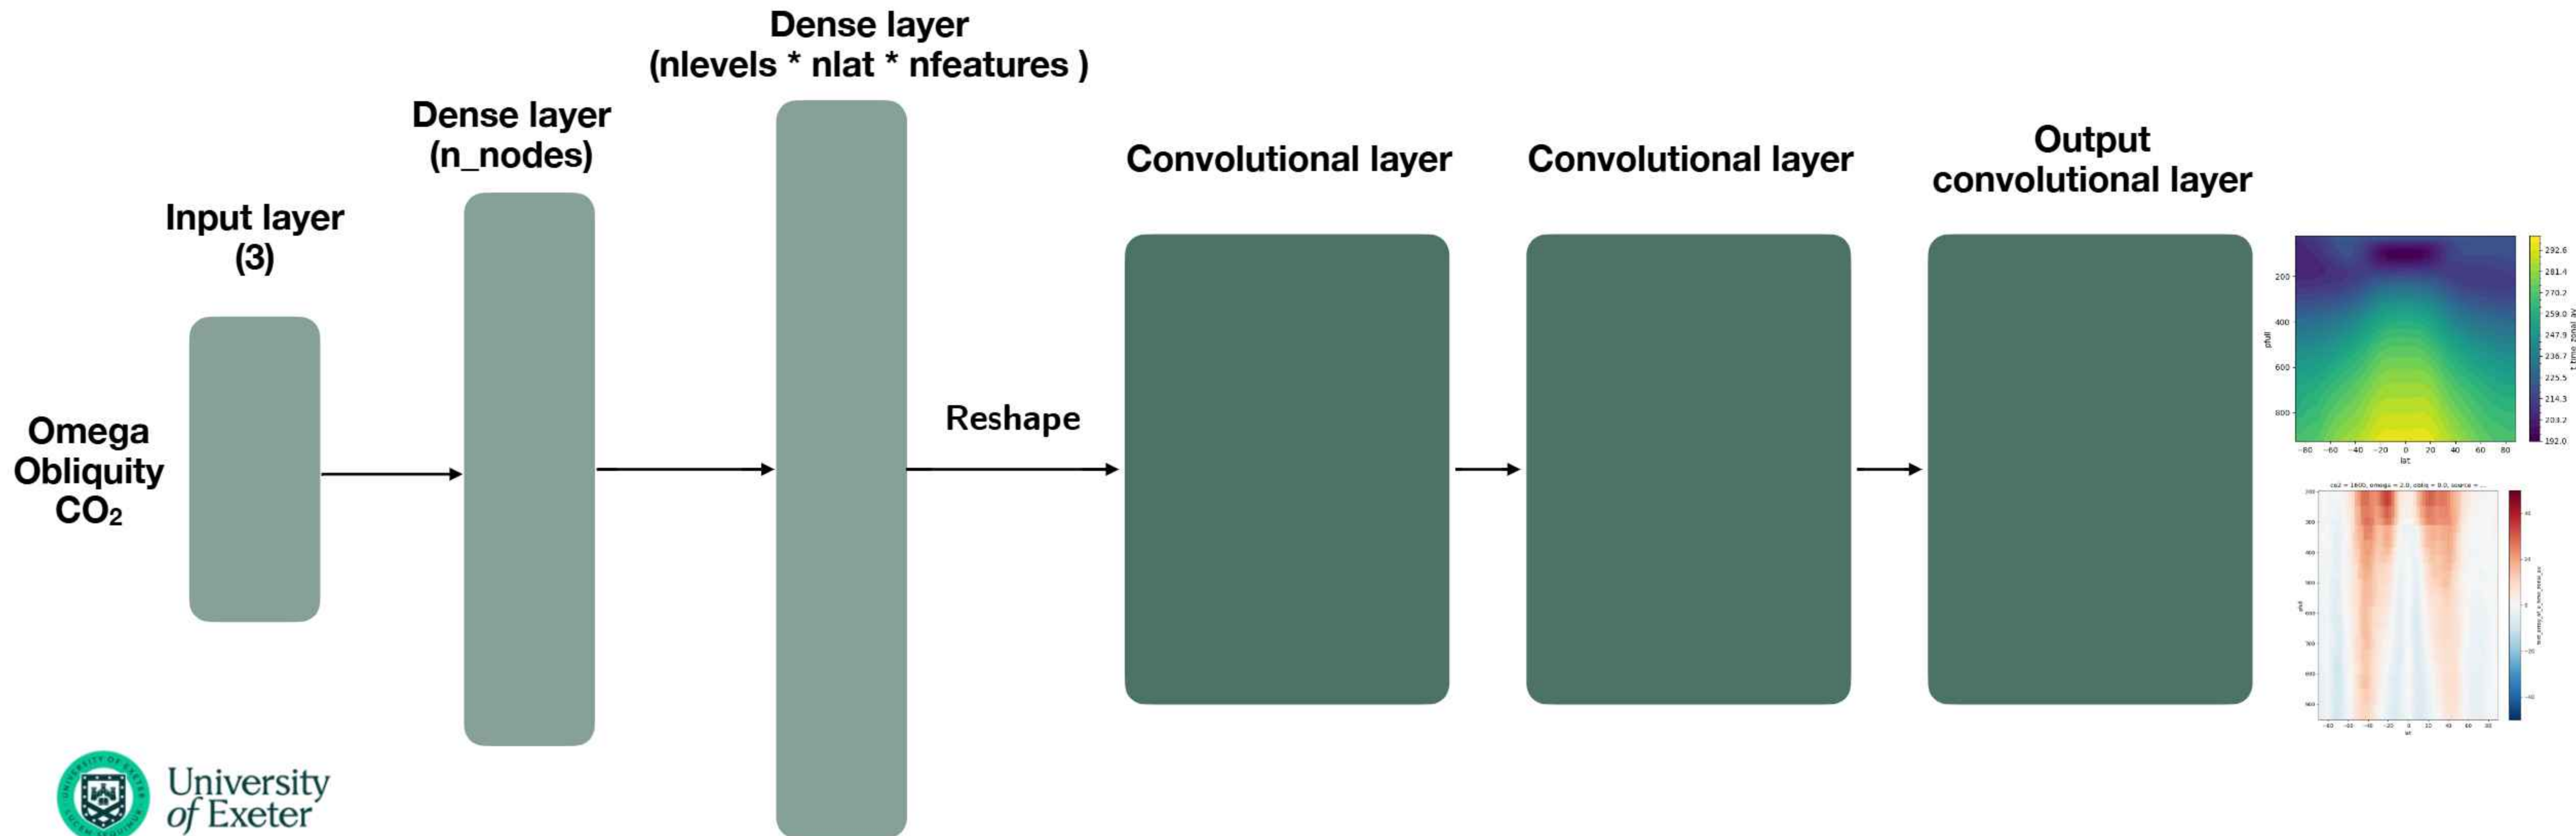


Using the column model as a constraint



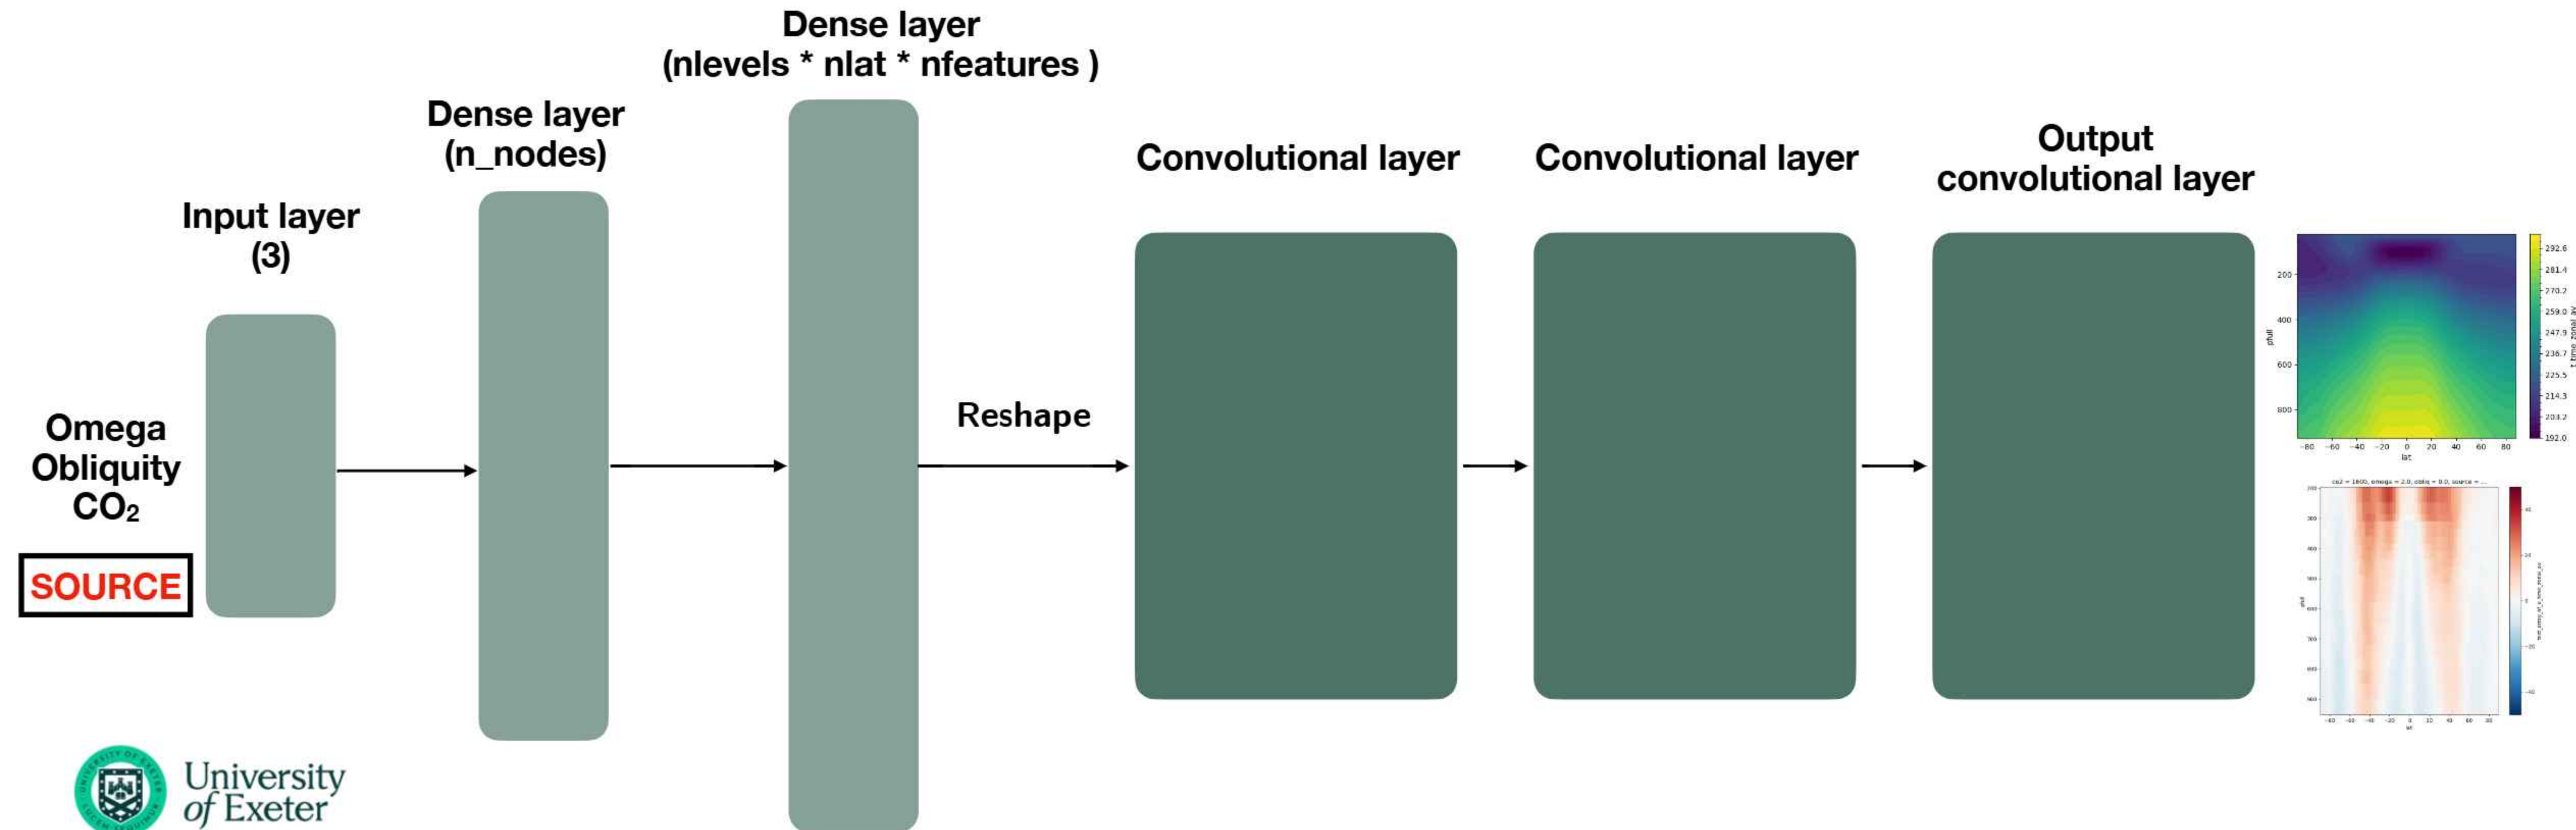
Using the column model as a constraint

- Add the 'source' as an input (0 for GCM or 1 for column), and ask the model to predict from either source



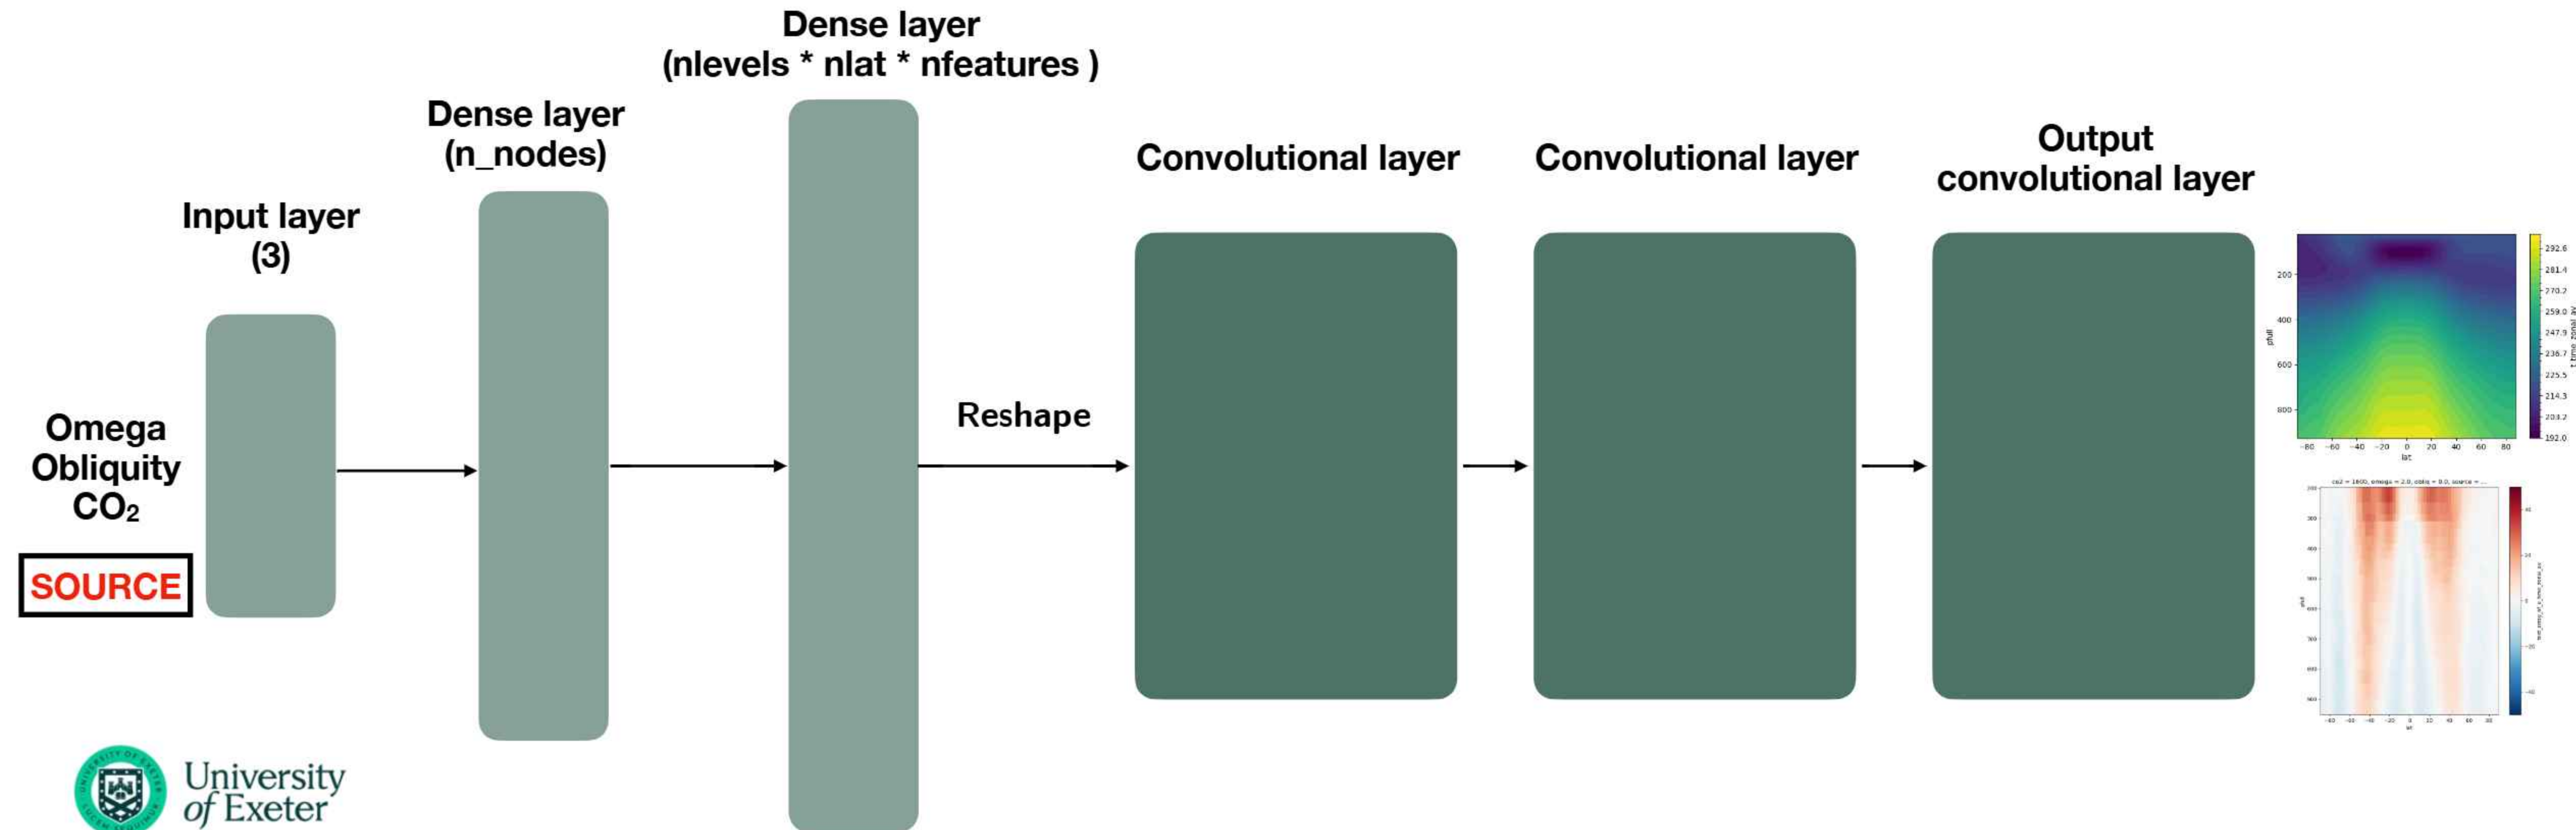
Using the column model as a constraint

- Add the 'source' as an input (0 for GCM or 1 for column), and ask the model to predict from either source



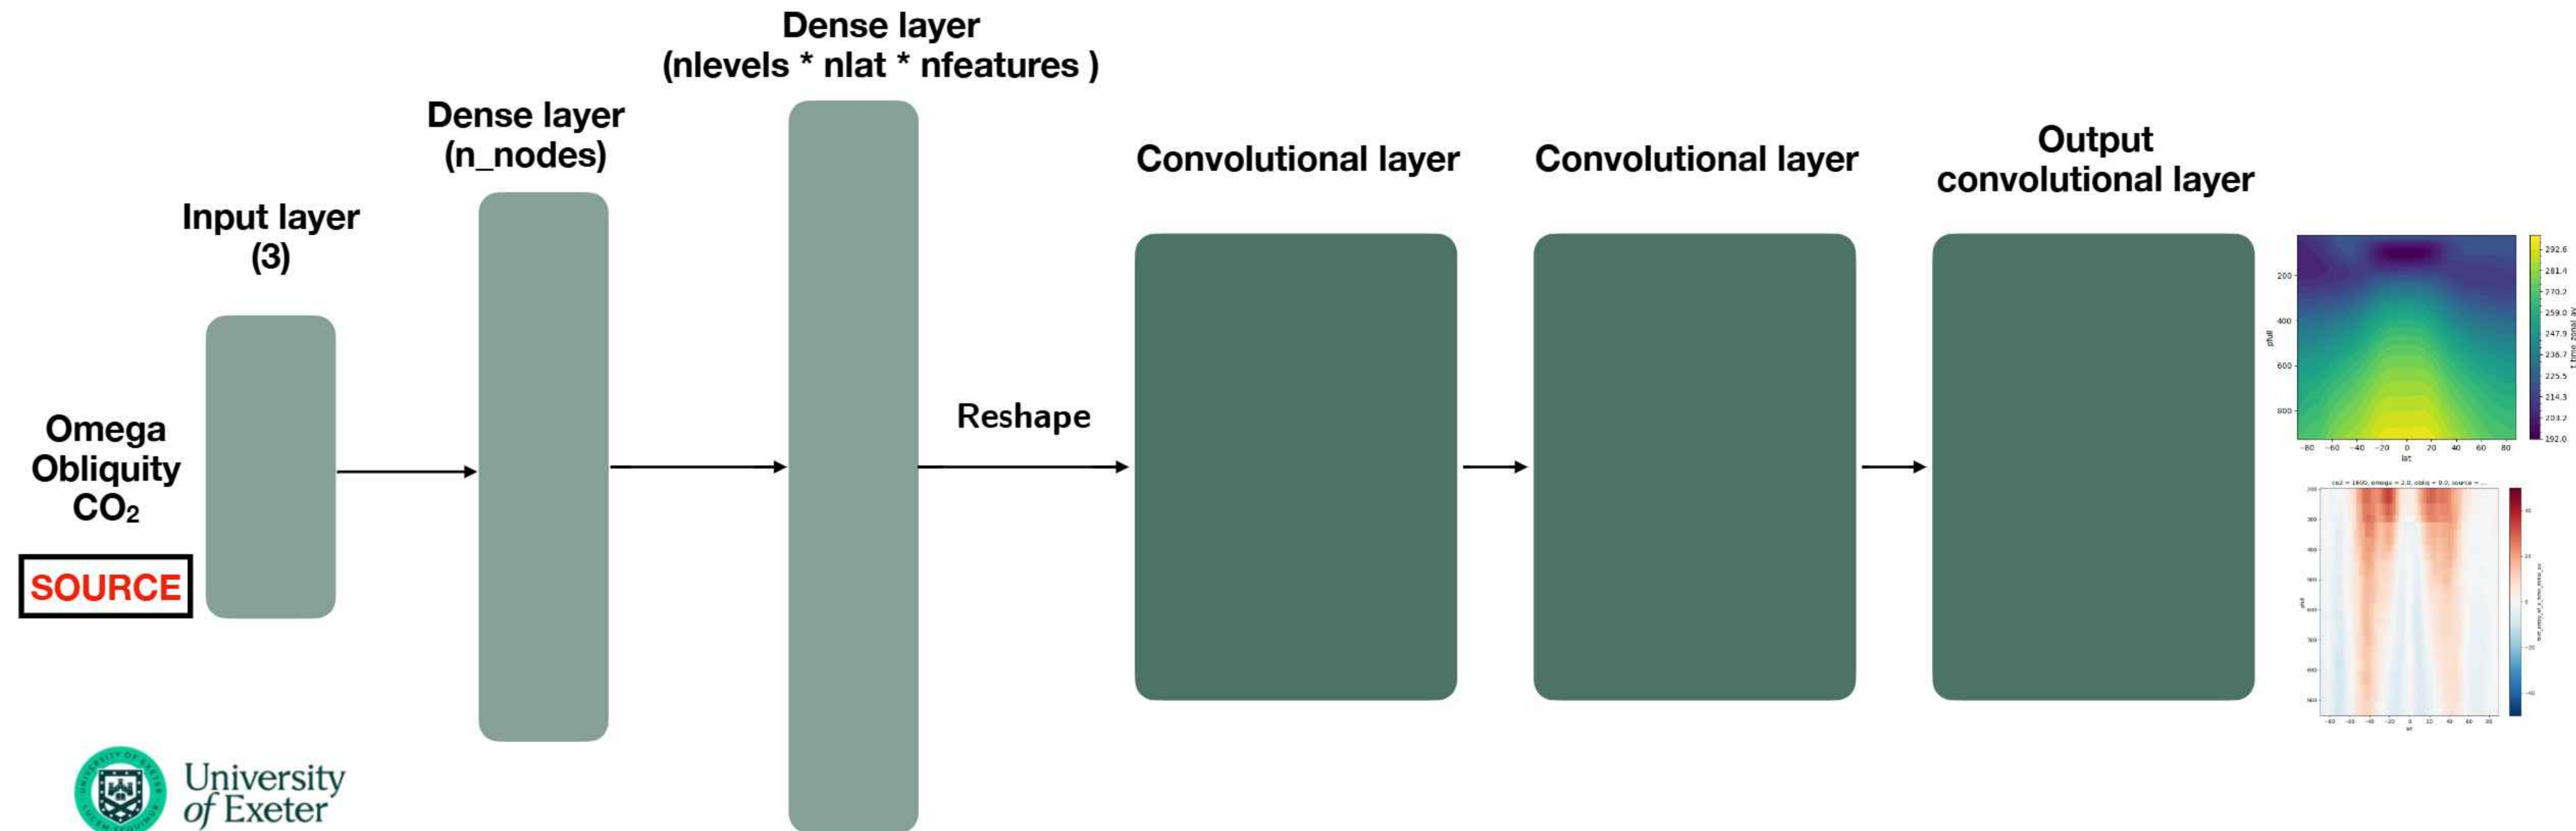
Using the column model as a constraint

- Add the 'source' as an input (0 for GCM or 1 for column), and ask the model to predict from either source
- In training we show the model all except the 1600ppmv case from the GCM, and ALL of the CO₂ values from the column model



Using the column model as a constraint

- Add the 'source' as an input (0 for GCM or 1 for column), and ask the model to predict from either source
- In training we show the model all except the 1600ppmv case from the GCM, and ALL of the CO₂ values from the column model
- Now when we ask it to predict the GCM at high CO₂, it has seen high CO₂ in training from the column but not the GCM



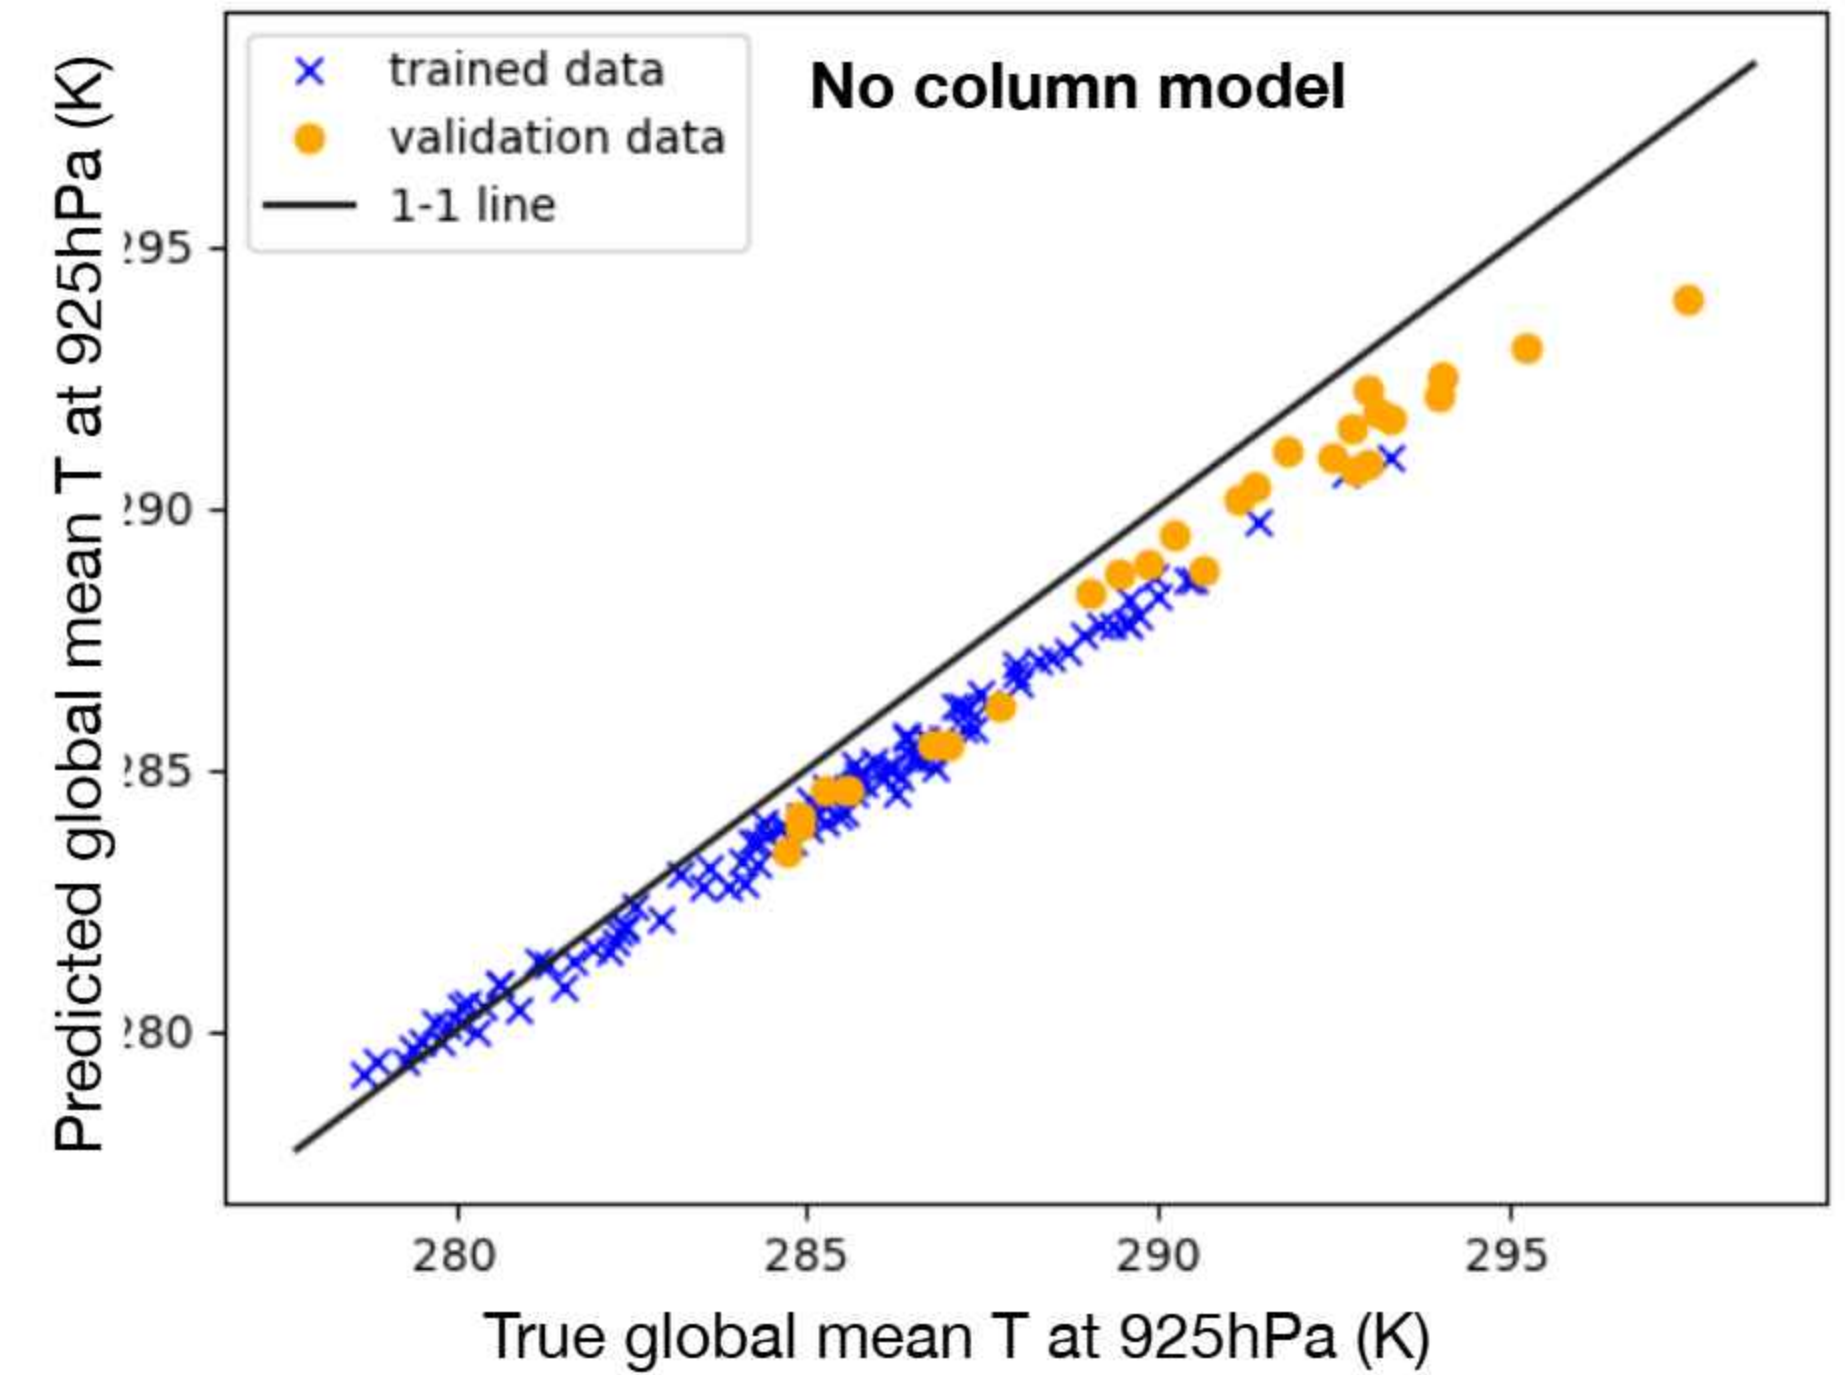
Improved high-CO₂ performance with column model training

Improved high-CO₂ performance with column model training

- When we include the column model experiments into our training data, **we substantially improve the extrapolation to high CO₂**

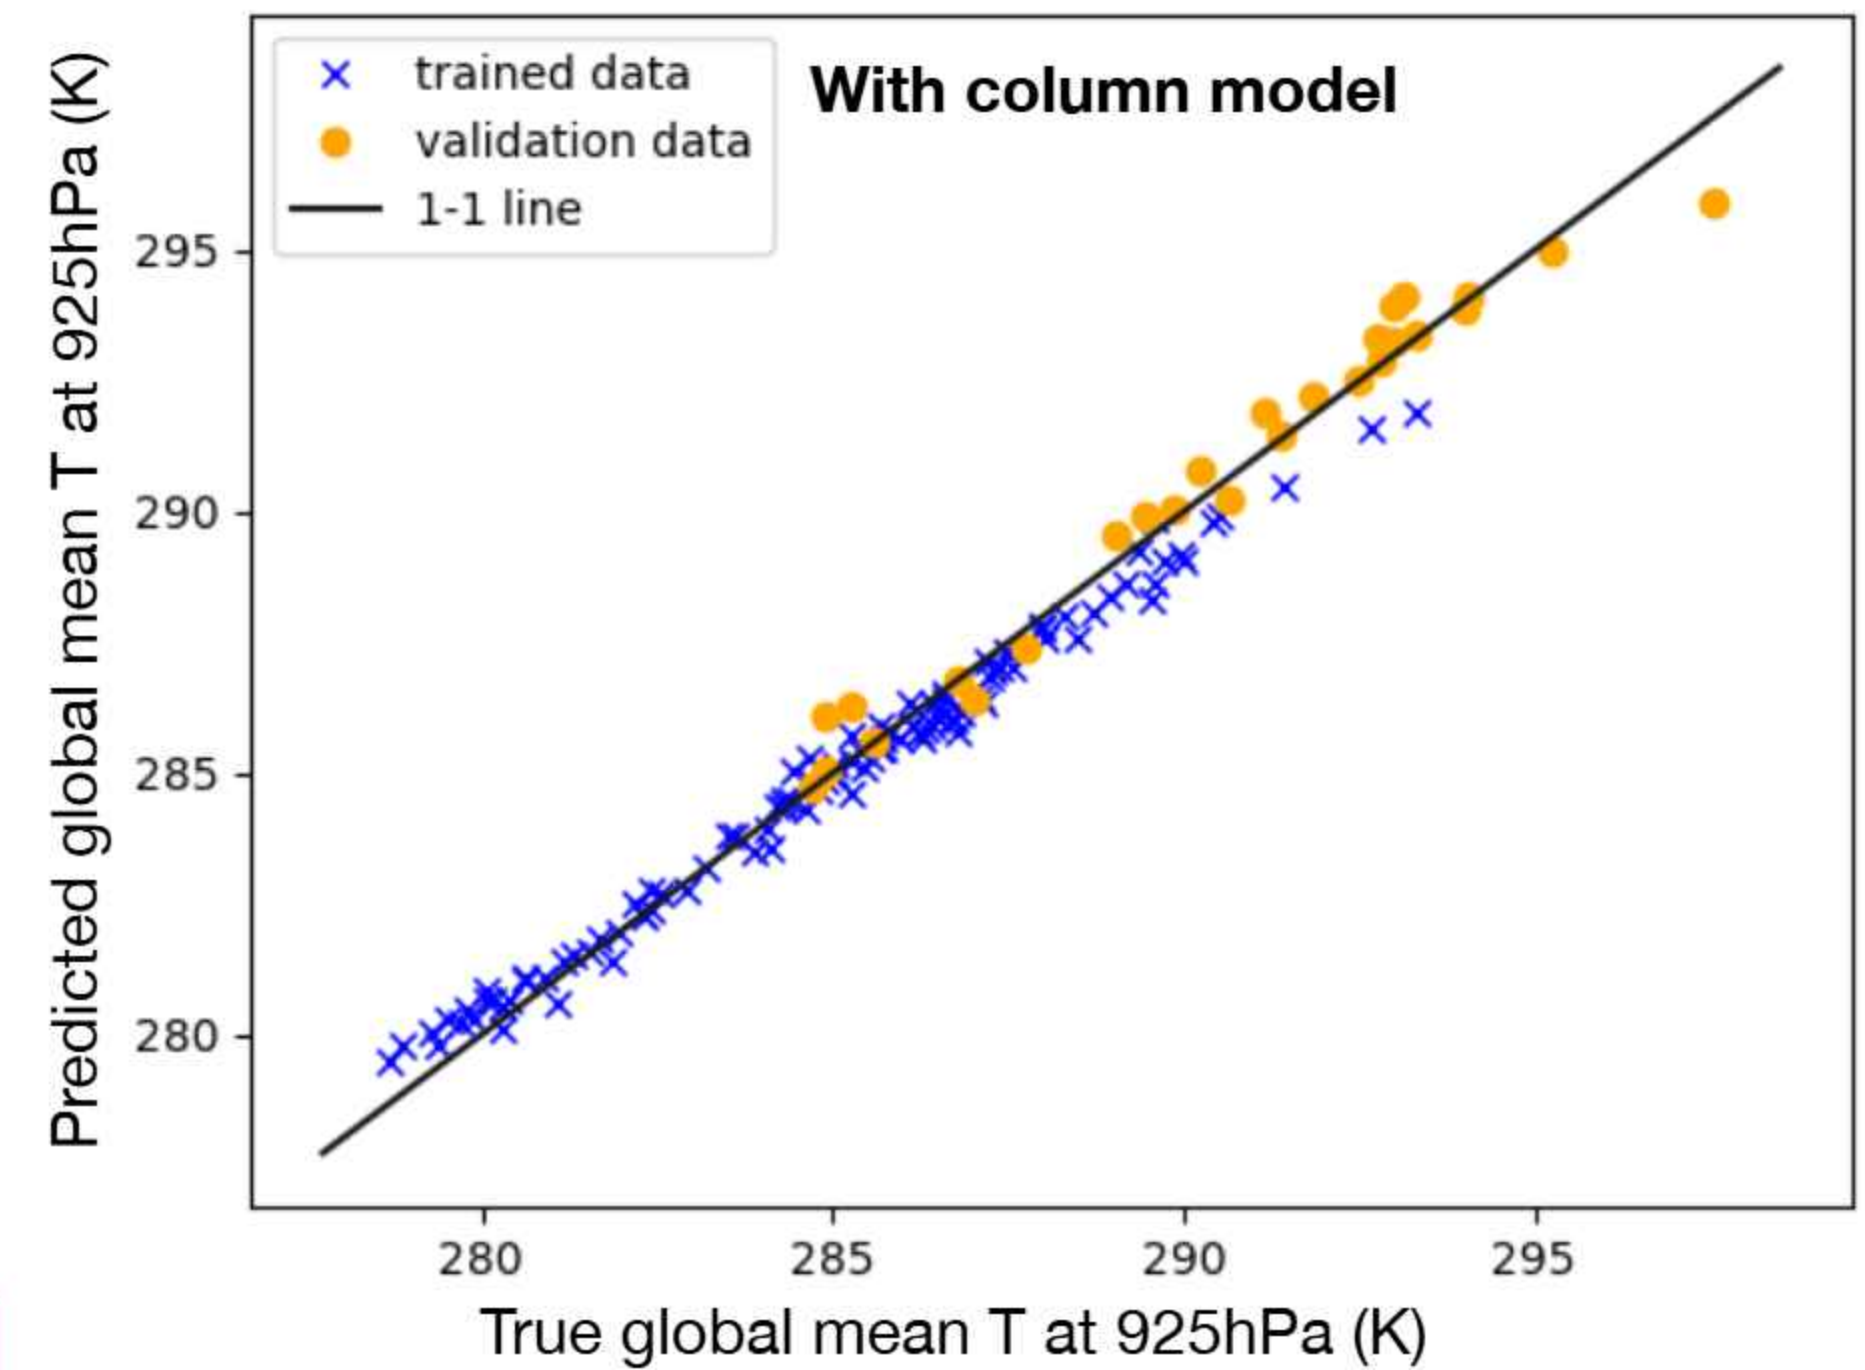
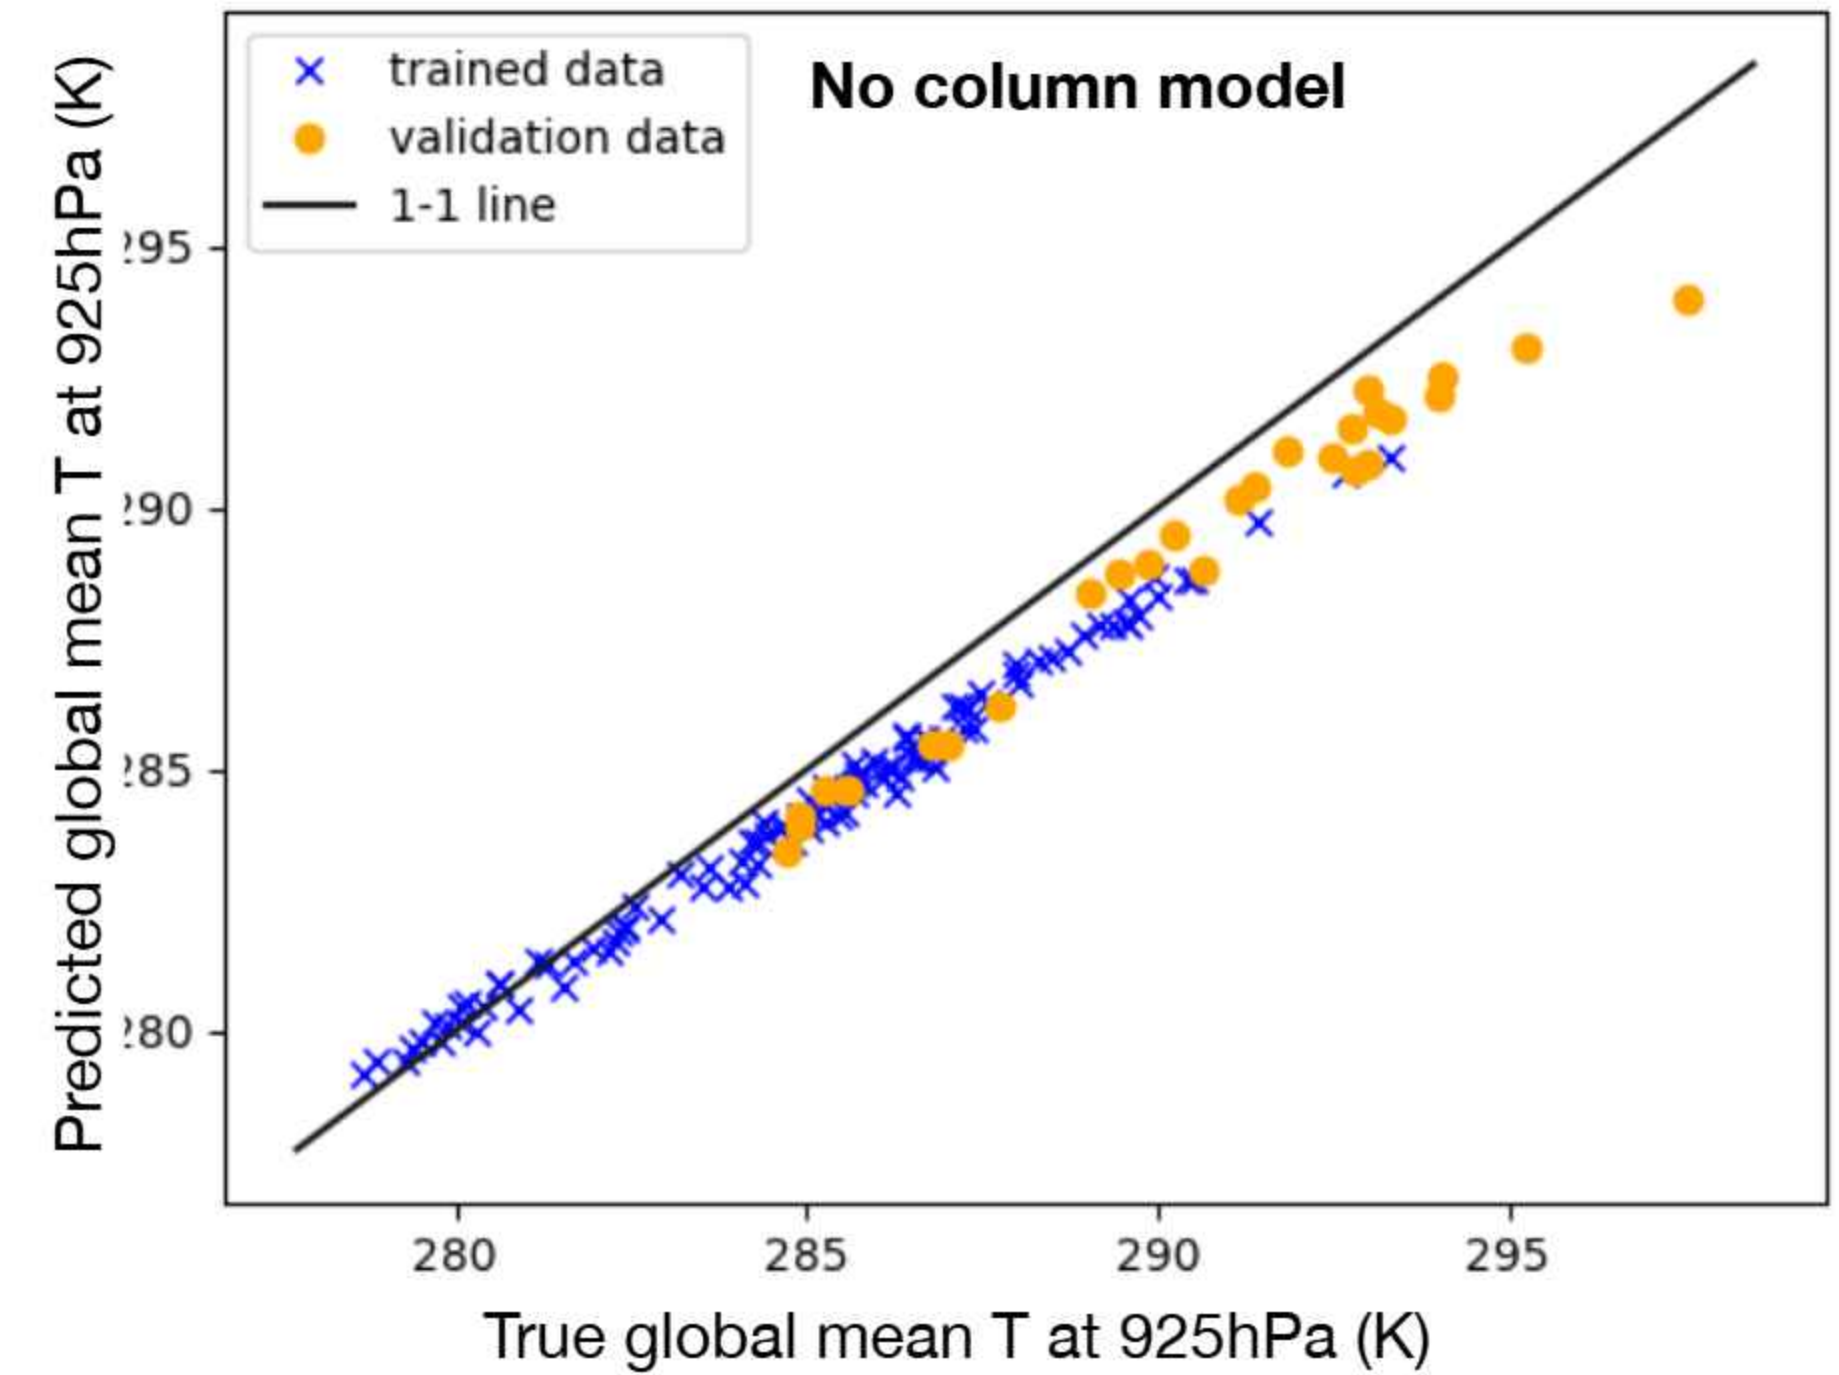
Improved high-CO₂ performance with column model training

- When we include the column model experiments into our training data, **we substantially improve the extrapolation to high CO₂**



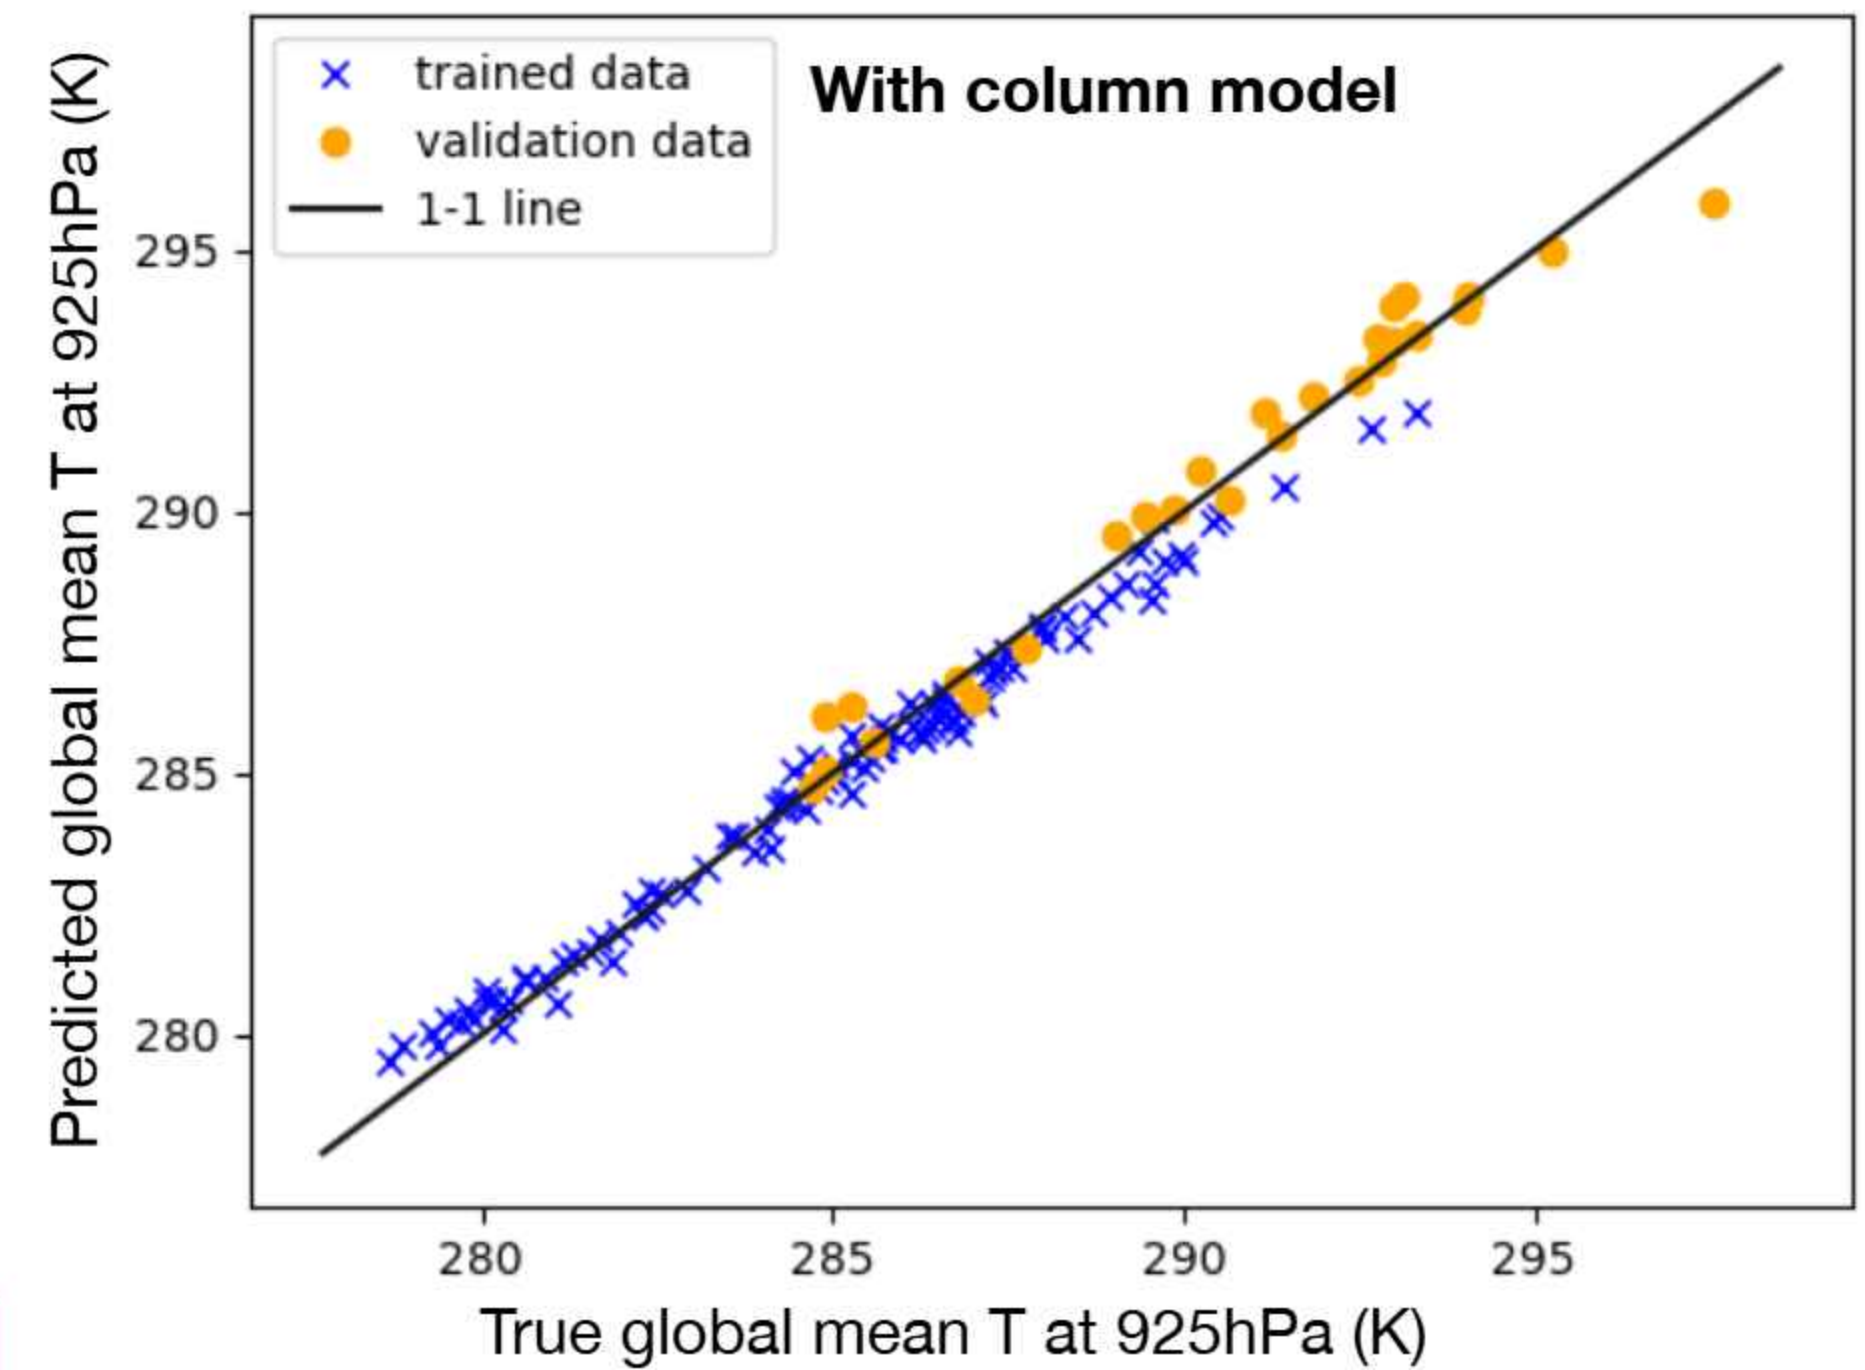
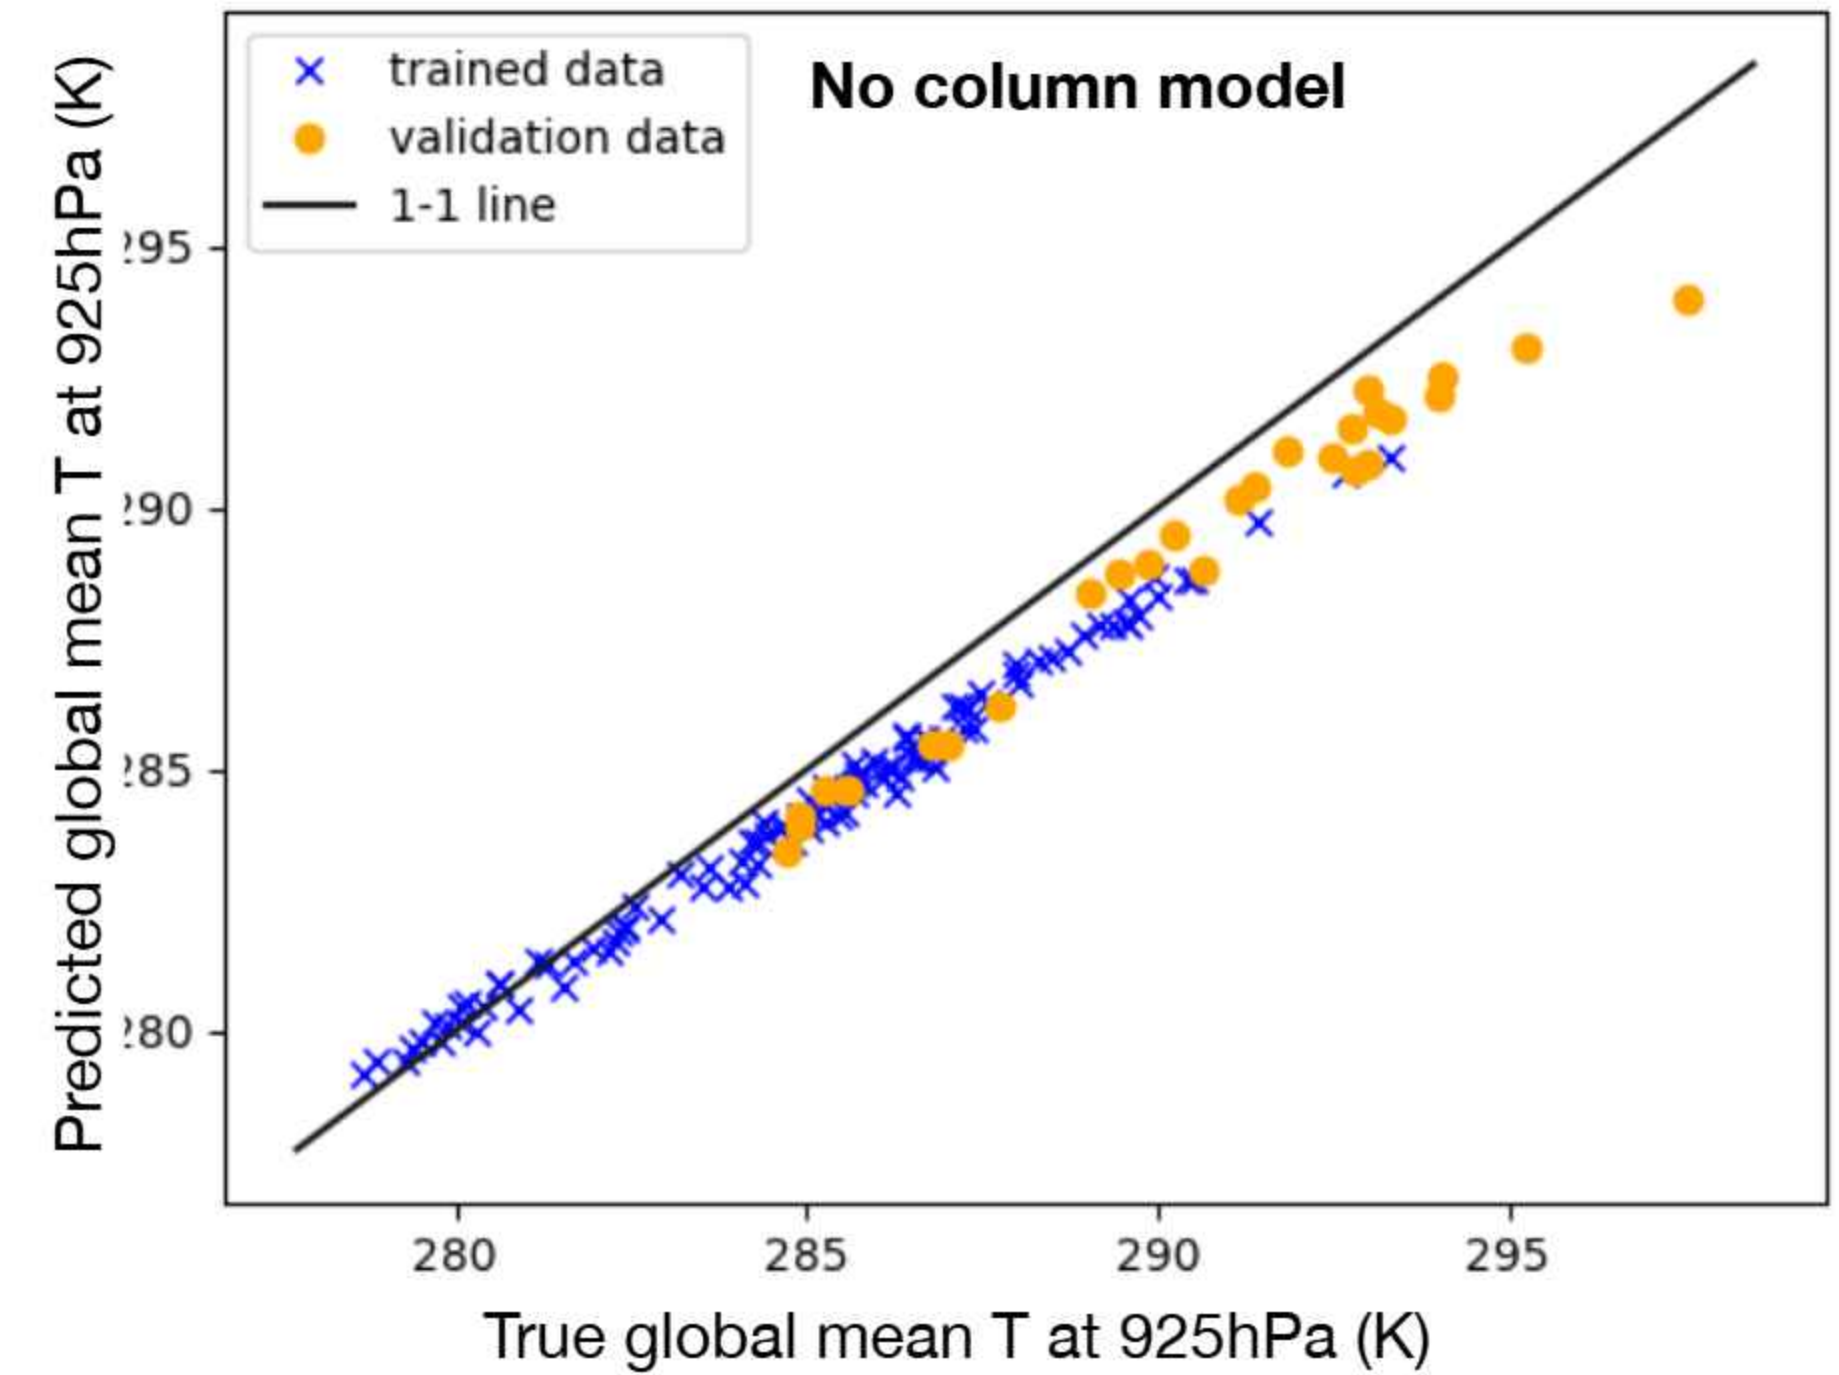
Improved high-CO₂ performance with column model training

- When we include the column model experiments into our training data, **we substantially improve the extrapolation to high CO₂**



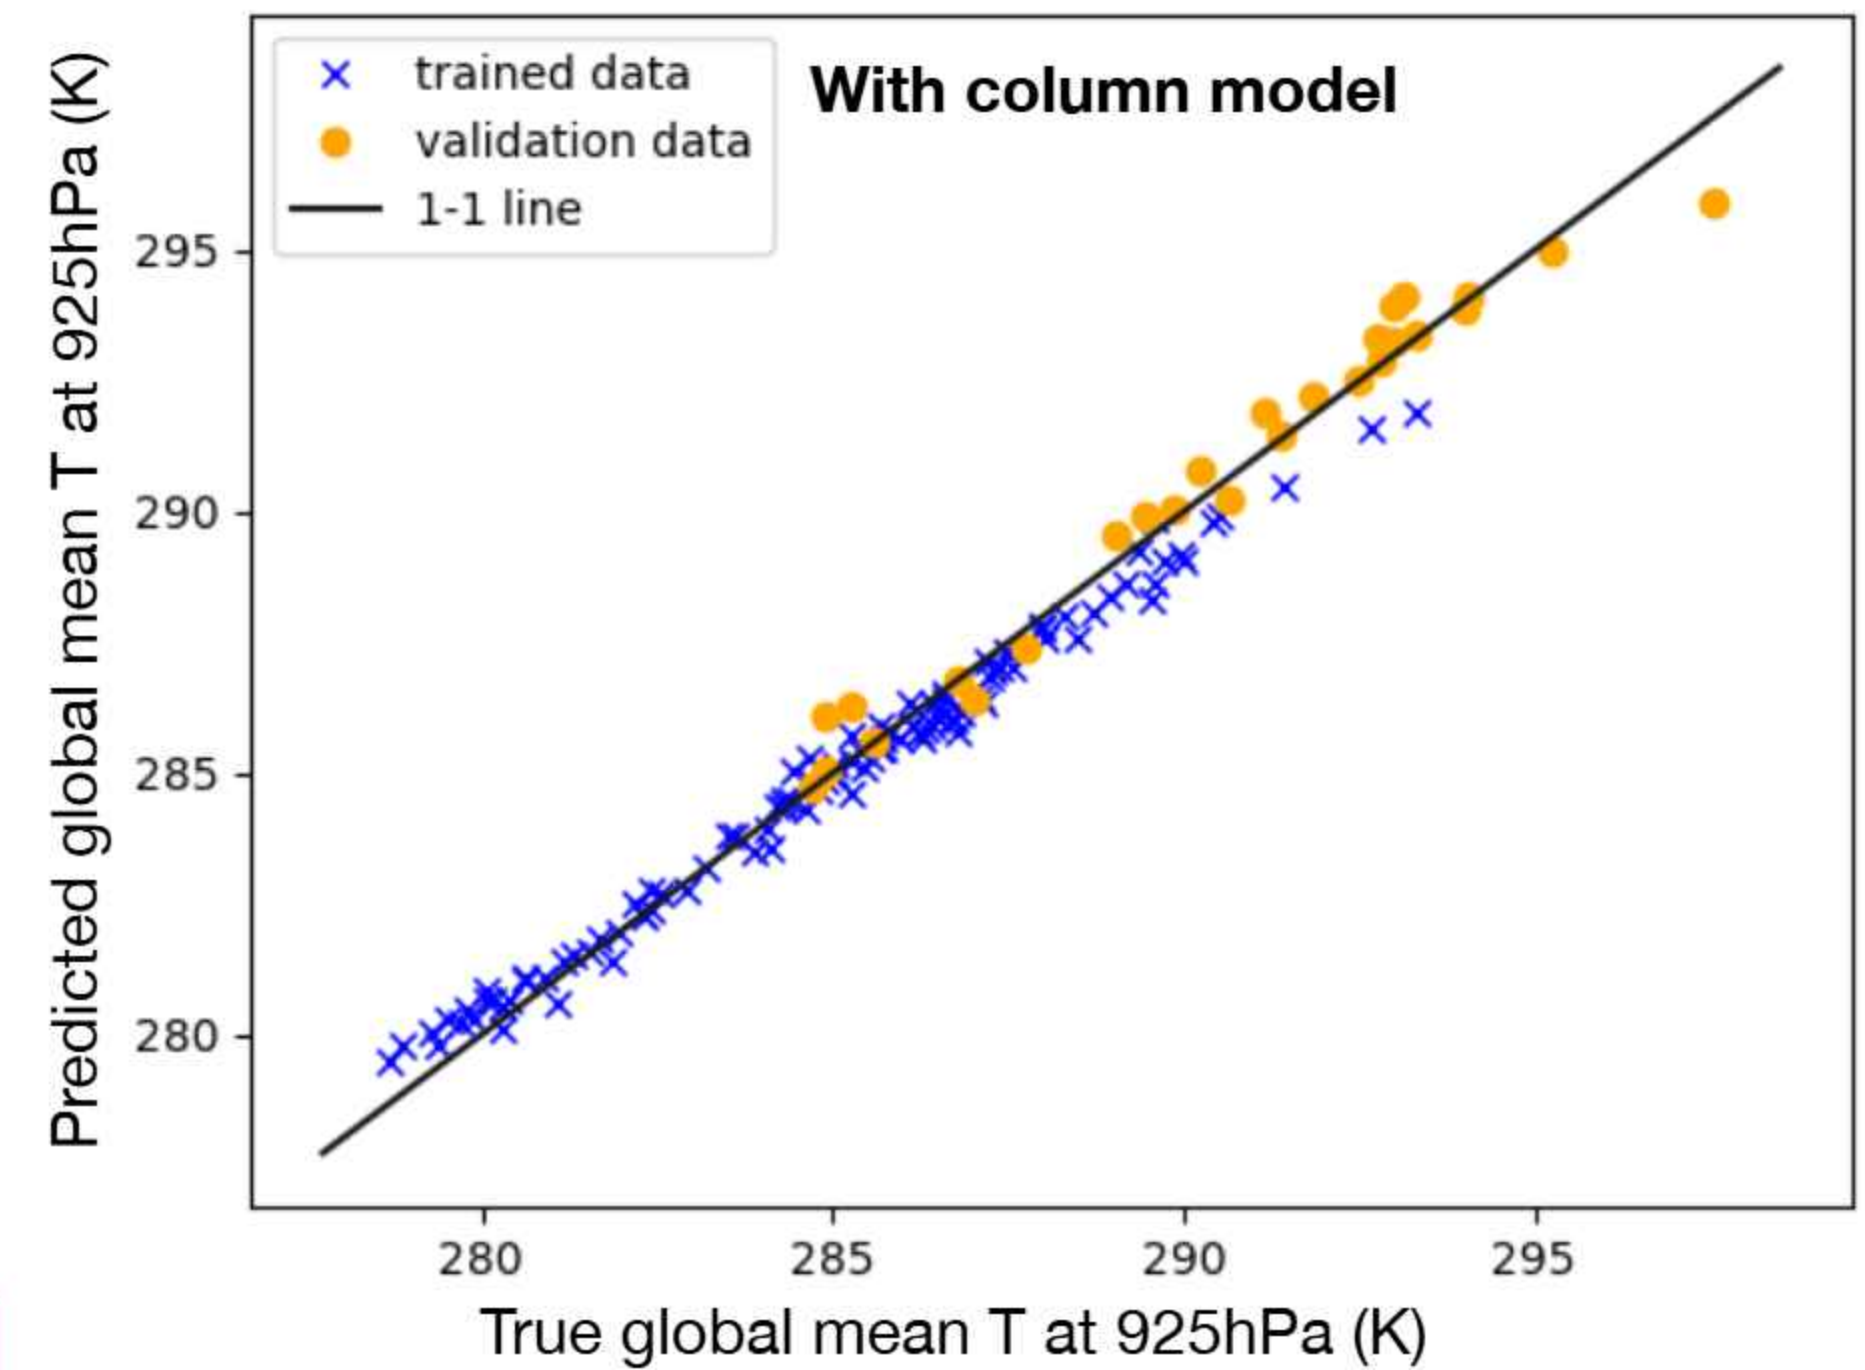
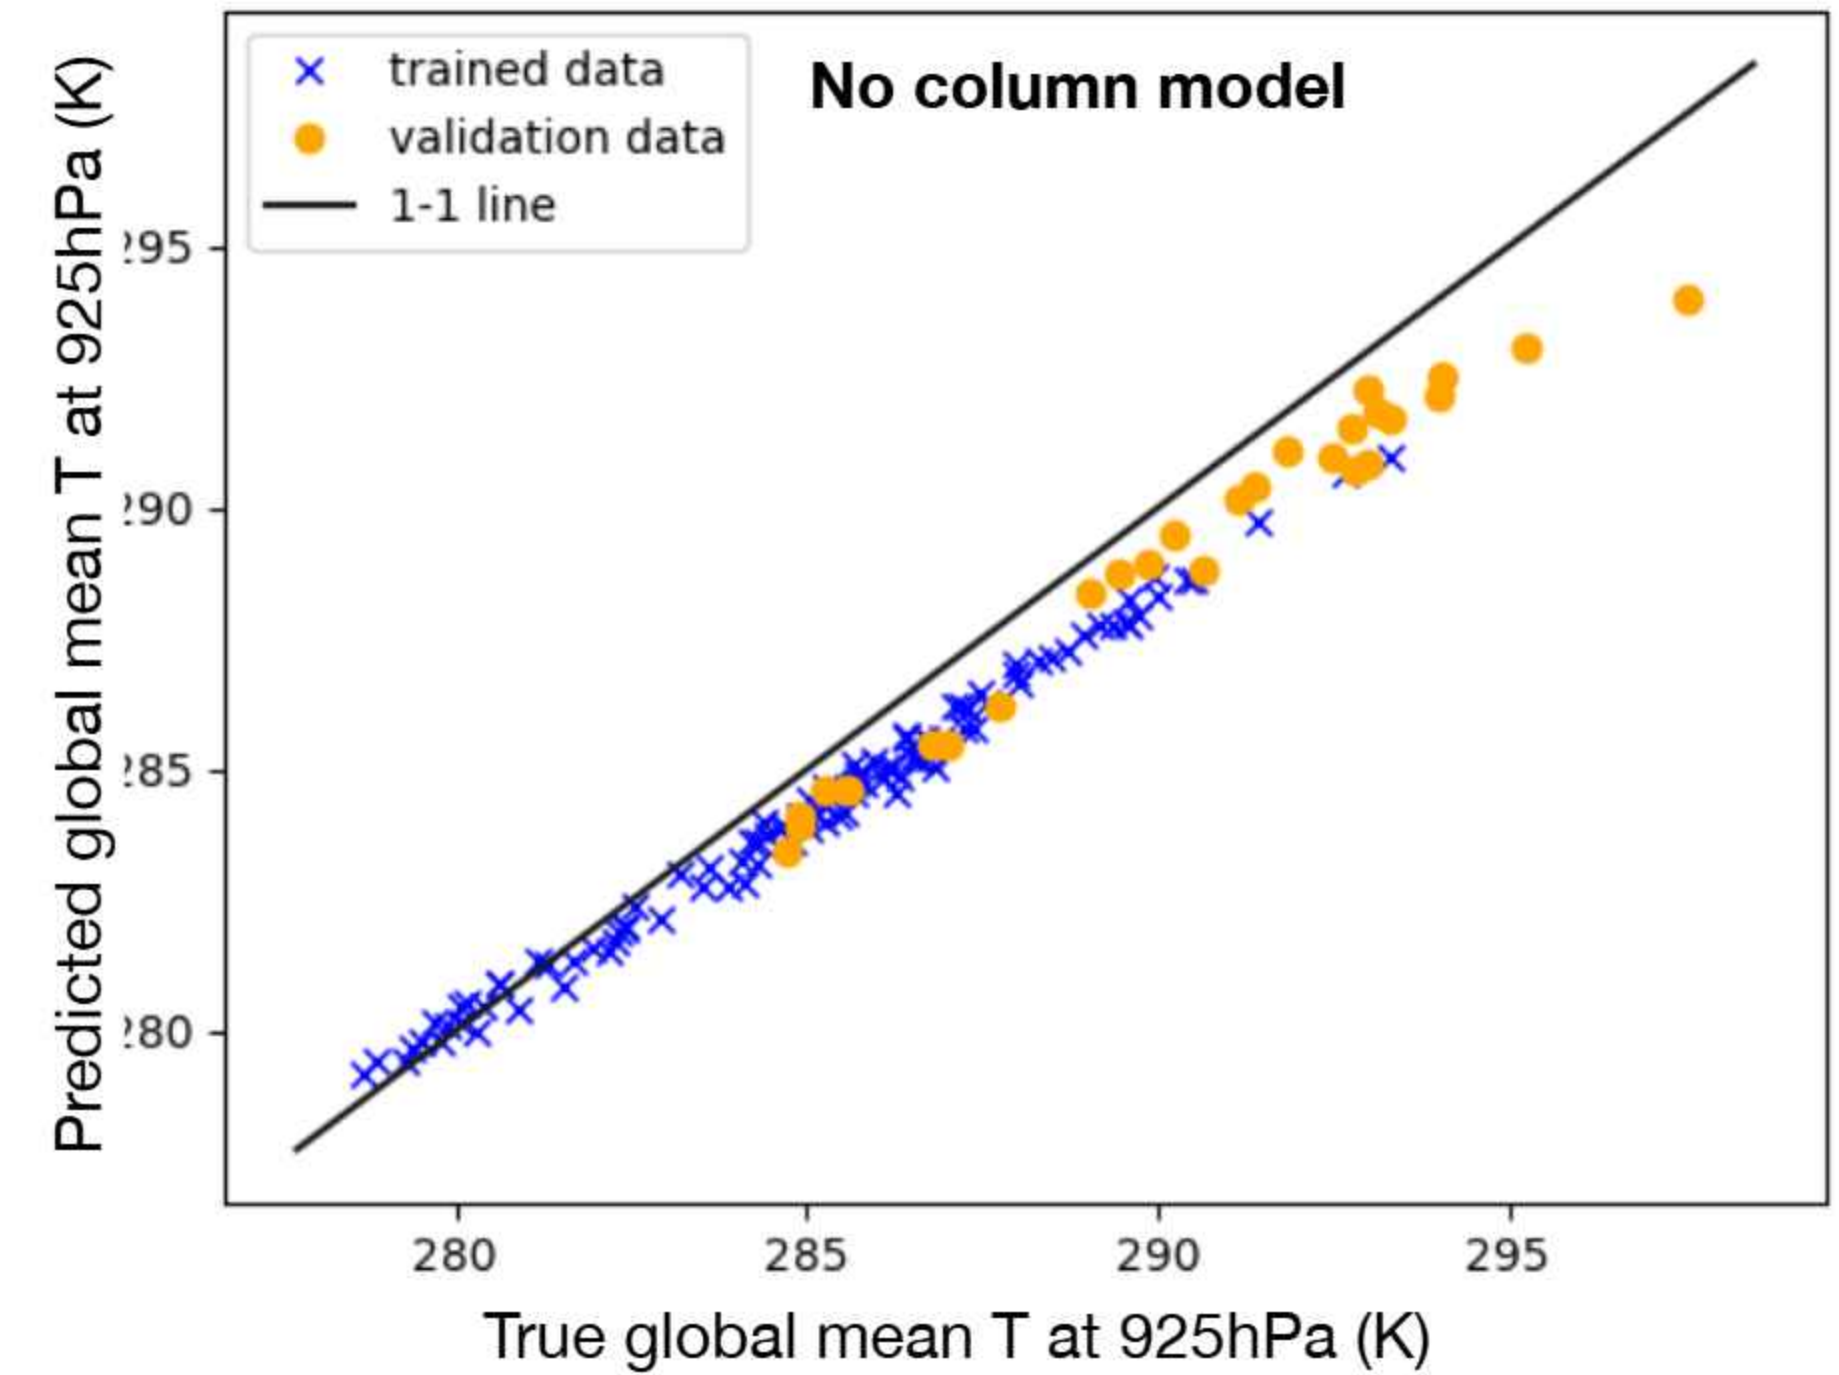
Improved high-CO₂ performance with column model training

- When we include the column model experiments into our training data, **we substantially improve the extrapolation to high CO₂**
- New network has seen high CO₂ in training, but only with the column model.



Improved high-CO₂ performance with column model training

- When we include the column model experiments into our training data, **we substantially improve the extrapolation to high CO₂**
- New network has seen high CO₂ in training, but only with the column model.
- Model has learned how the GCM predictions differ from the column at low CO₂ - now it can apply this at high CO₂



Incorporating dynamical constraints

Incorporating dynamical constraints

- Incorporating the column model has improved the extrapolation performance of the network for temperature, but makes very little difference to the wind predictions.

Incorporating dynamical constraints

- Incorporating the column model has improved the extrapolation performance of the network for temperature, but makes very little difference to the wind predictions.
- How can we incorporate a physical constraint that helps to constrain the winds using the newly-constrained temperatures?

Incorporating dynamical constraints

- Incorporating the column model has improved the extrapolation performance of the network for temperature, but makes very little difference to the wind predictions.
- How can we incorporate a physical constraint that helps to constrain the winds using the newly-constrained temperatures?

$$\underbrace{2\Omega \sin \theta u}_{\text{Coriolis}} + \underbrace{\frac{u^2 \tan \theta}{a}}_{\text{Metric term}} = - \underbrace{\frac{1}{\rho a} \frac{\partial p}{\partial \theta}}_{\text{Pressure gradient}}$$

Incorporating dynamical constraints

- Incorporating the column model has improved the extrapolation performance of the network for temperature, but makes very little difference to the wind predictions.
- How can we incorporate a physical constraint that helps to constrain the winds using the newly-constrained temperatures?

$$\underbrace{2\Omega \sin \theta u}_{\text{Coriolis}} + \underbrace{\frac{u^2 \tan \theta}{a}}_{\text{Metric term}} = - \underbrace{\frac{1}{\rho a} \frac{\partial p}{\partial \theta}}_{\text{Pressure gradient}}$$

- Take zonal-mean and vertical derivative WRT pressure (with some assumptions) gives the thermal wind relation for gradient wind balance:

Incorporating dynamical constraints

- Incorporating the column model has improved the extrapolation performance of the network for temperature, but makes very little difference to the wind predictions.
- How can we incorporate a physical constraint that helps to constrain the winds using the newly-constrained temperatures?

$$\underbrace{2\Omega \sin \theta u}_{\text{Coriolis}} + \underbrace{\frac{u^2 \tan \theta}{a}}_{\text{Metric term}} = - \underbrace{\frac{1}{\rho a} \frac{\partial p}{\partial \theta}}_{\text{Pressure gradient}}$$

- Take zonal-mean and vertical derivative WRT pressure (with some assumptions) gives the thermal wind relation for gradient wind balance:

$$2\Omega \sin \theta \frac{\partial \bar{u}}{\partial p} + \frac{2 \tan \theta}{a} \bar{u} \frac{\partial \bar{u}}{\partial p} \approx \frac{R}{ap} \frac{\partial \bar{T}}{\partial \theta}$$

Incorporating dynamical constraints

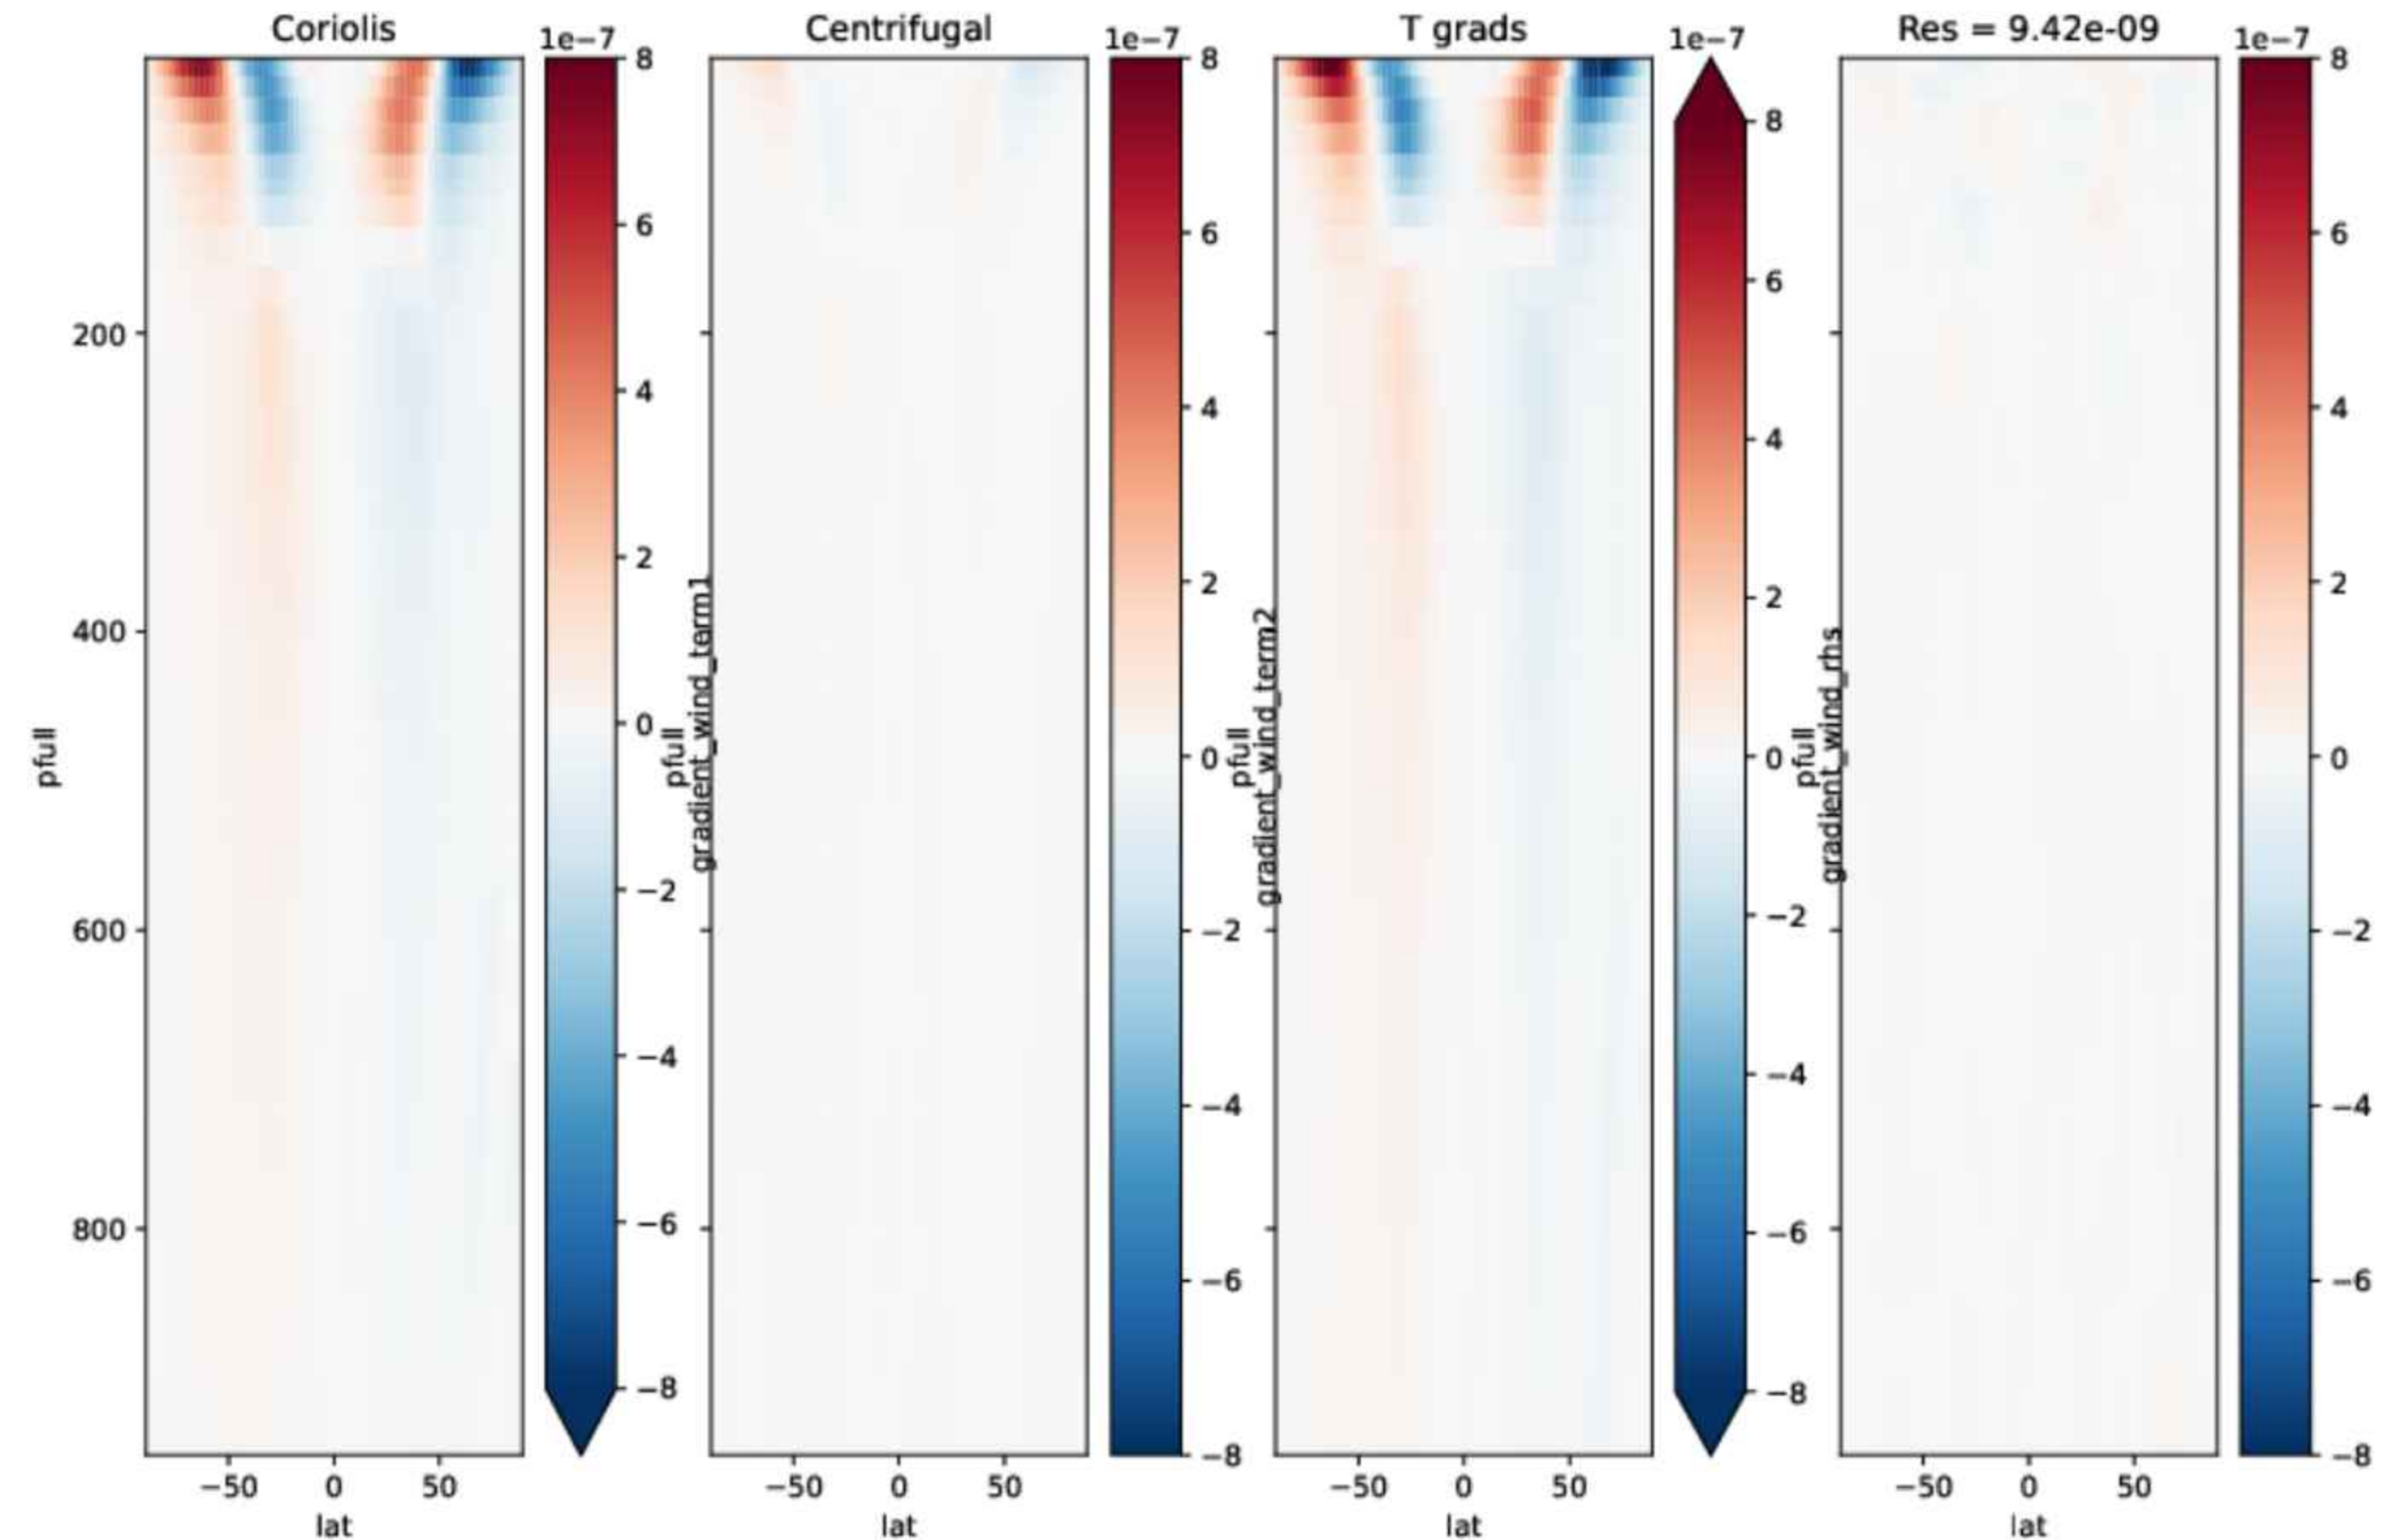
Incorporating dynamical constraints

- This constraint is well respected in the Isca data

Incorporating dynamical constraints

- This constraint is well respected in the Isca data

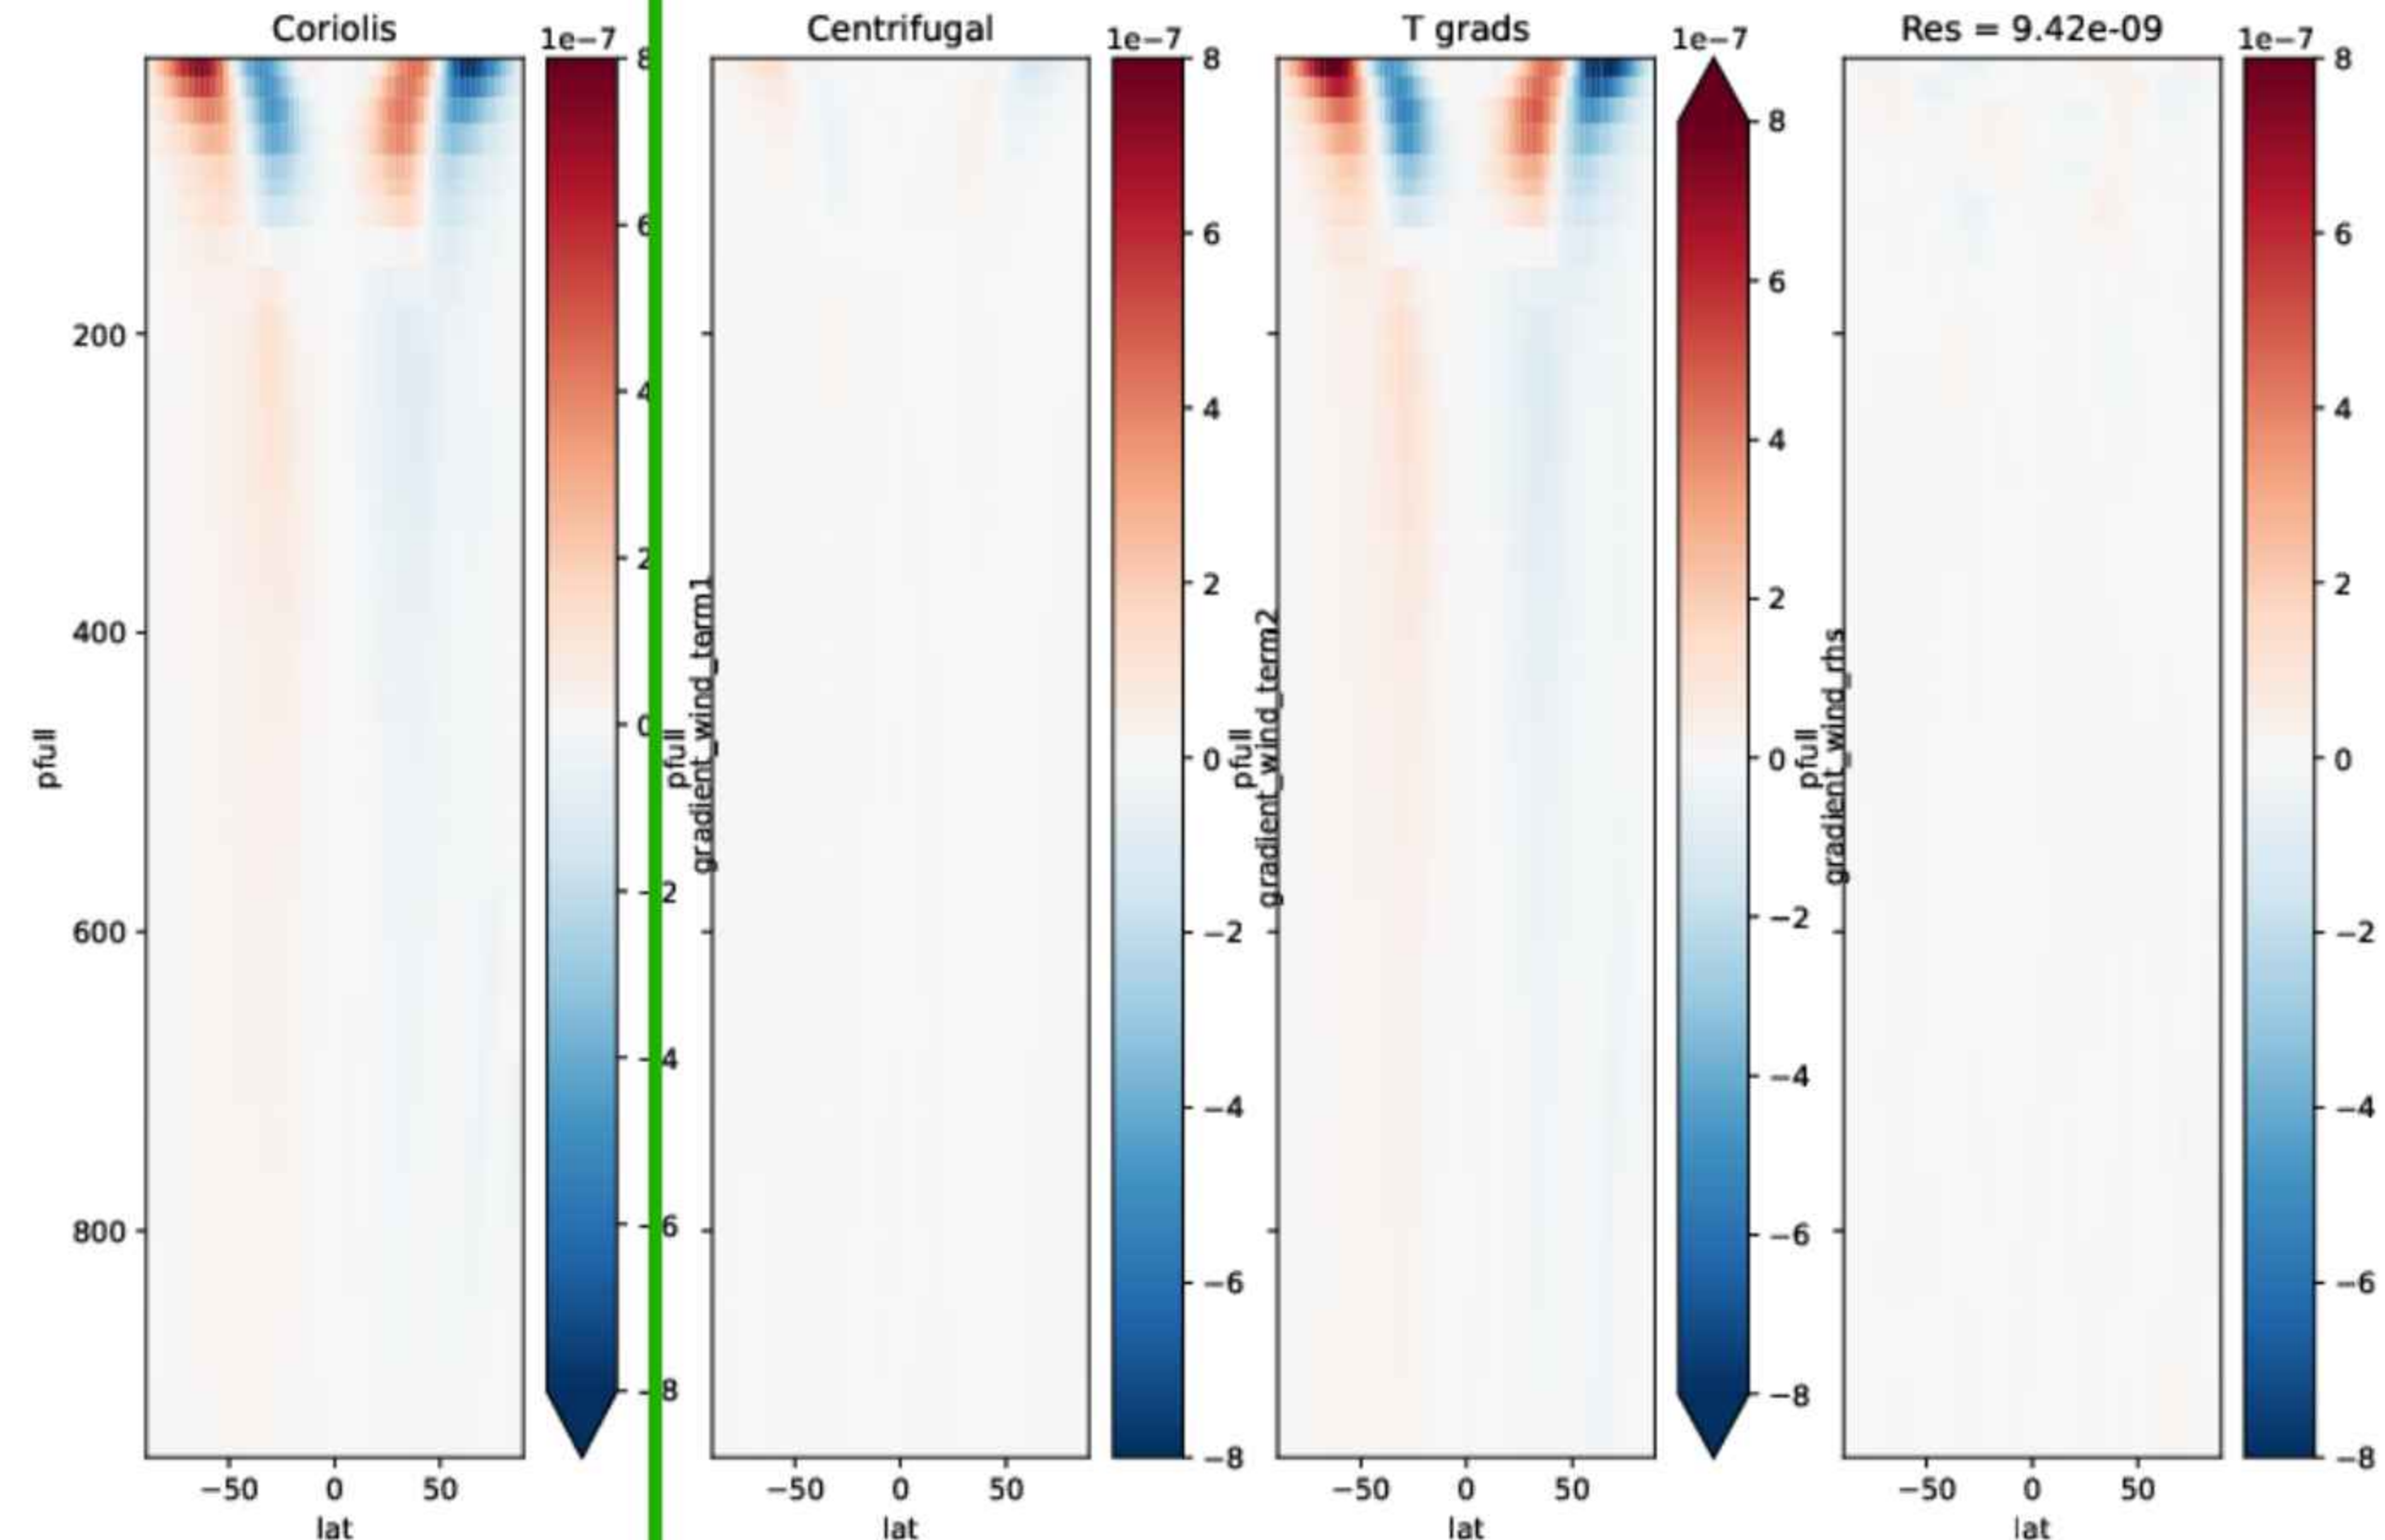
$$2\Omega \sin \theta \frac{\partial \bar{u}}{\partial p} + \frac{2 \tan \theta}{a} \bar{u} \frac{\partial \bar{u}}{\partial p} \approx \frac{R}{ap} \frac{\partial \bar{T}}{\partial \theta}$$



Incorporating dynamical constraints

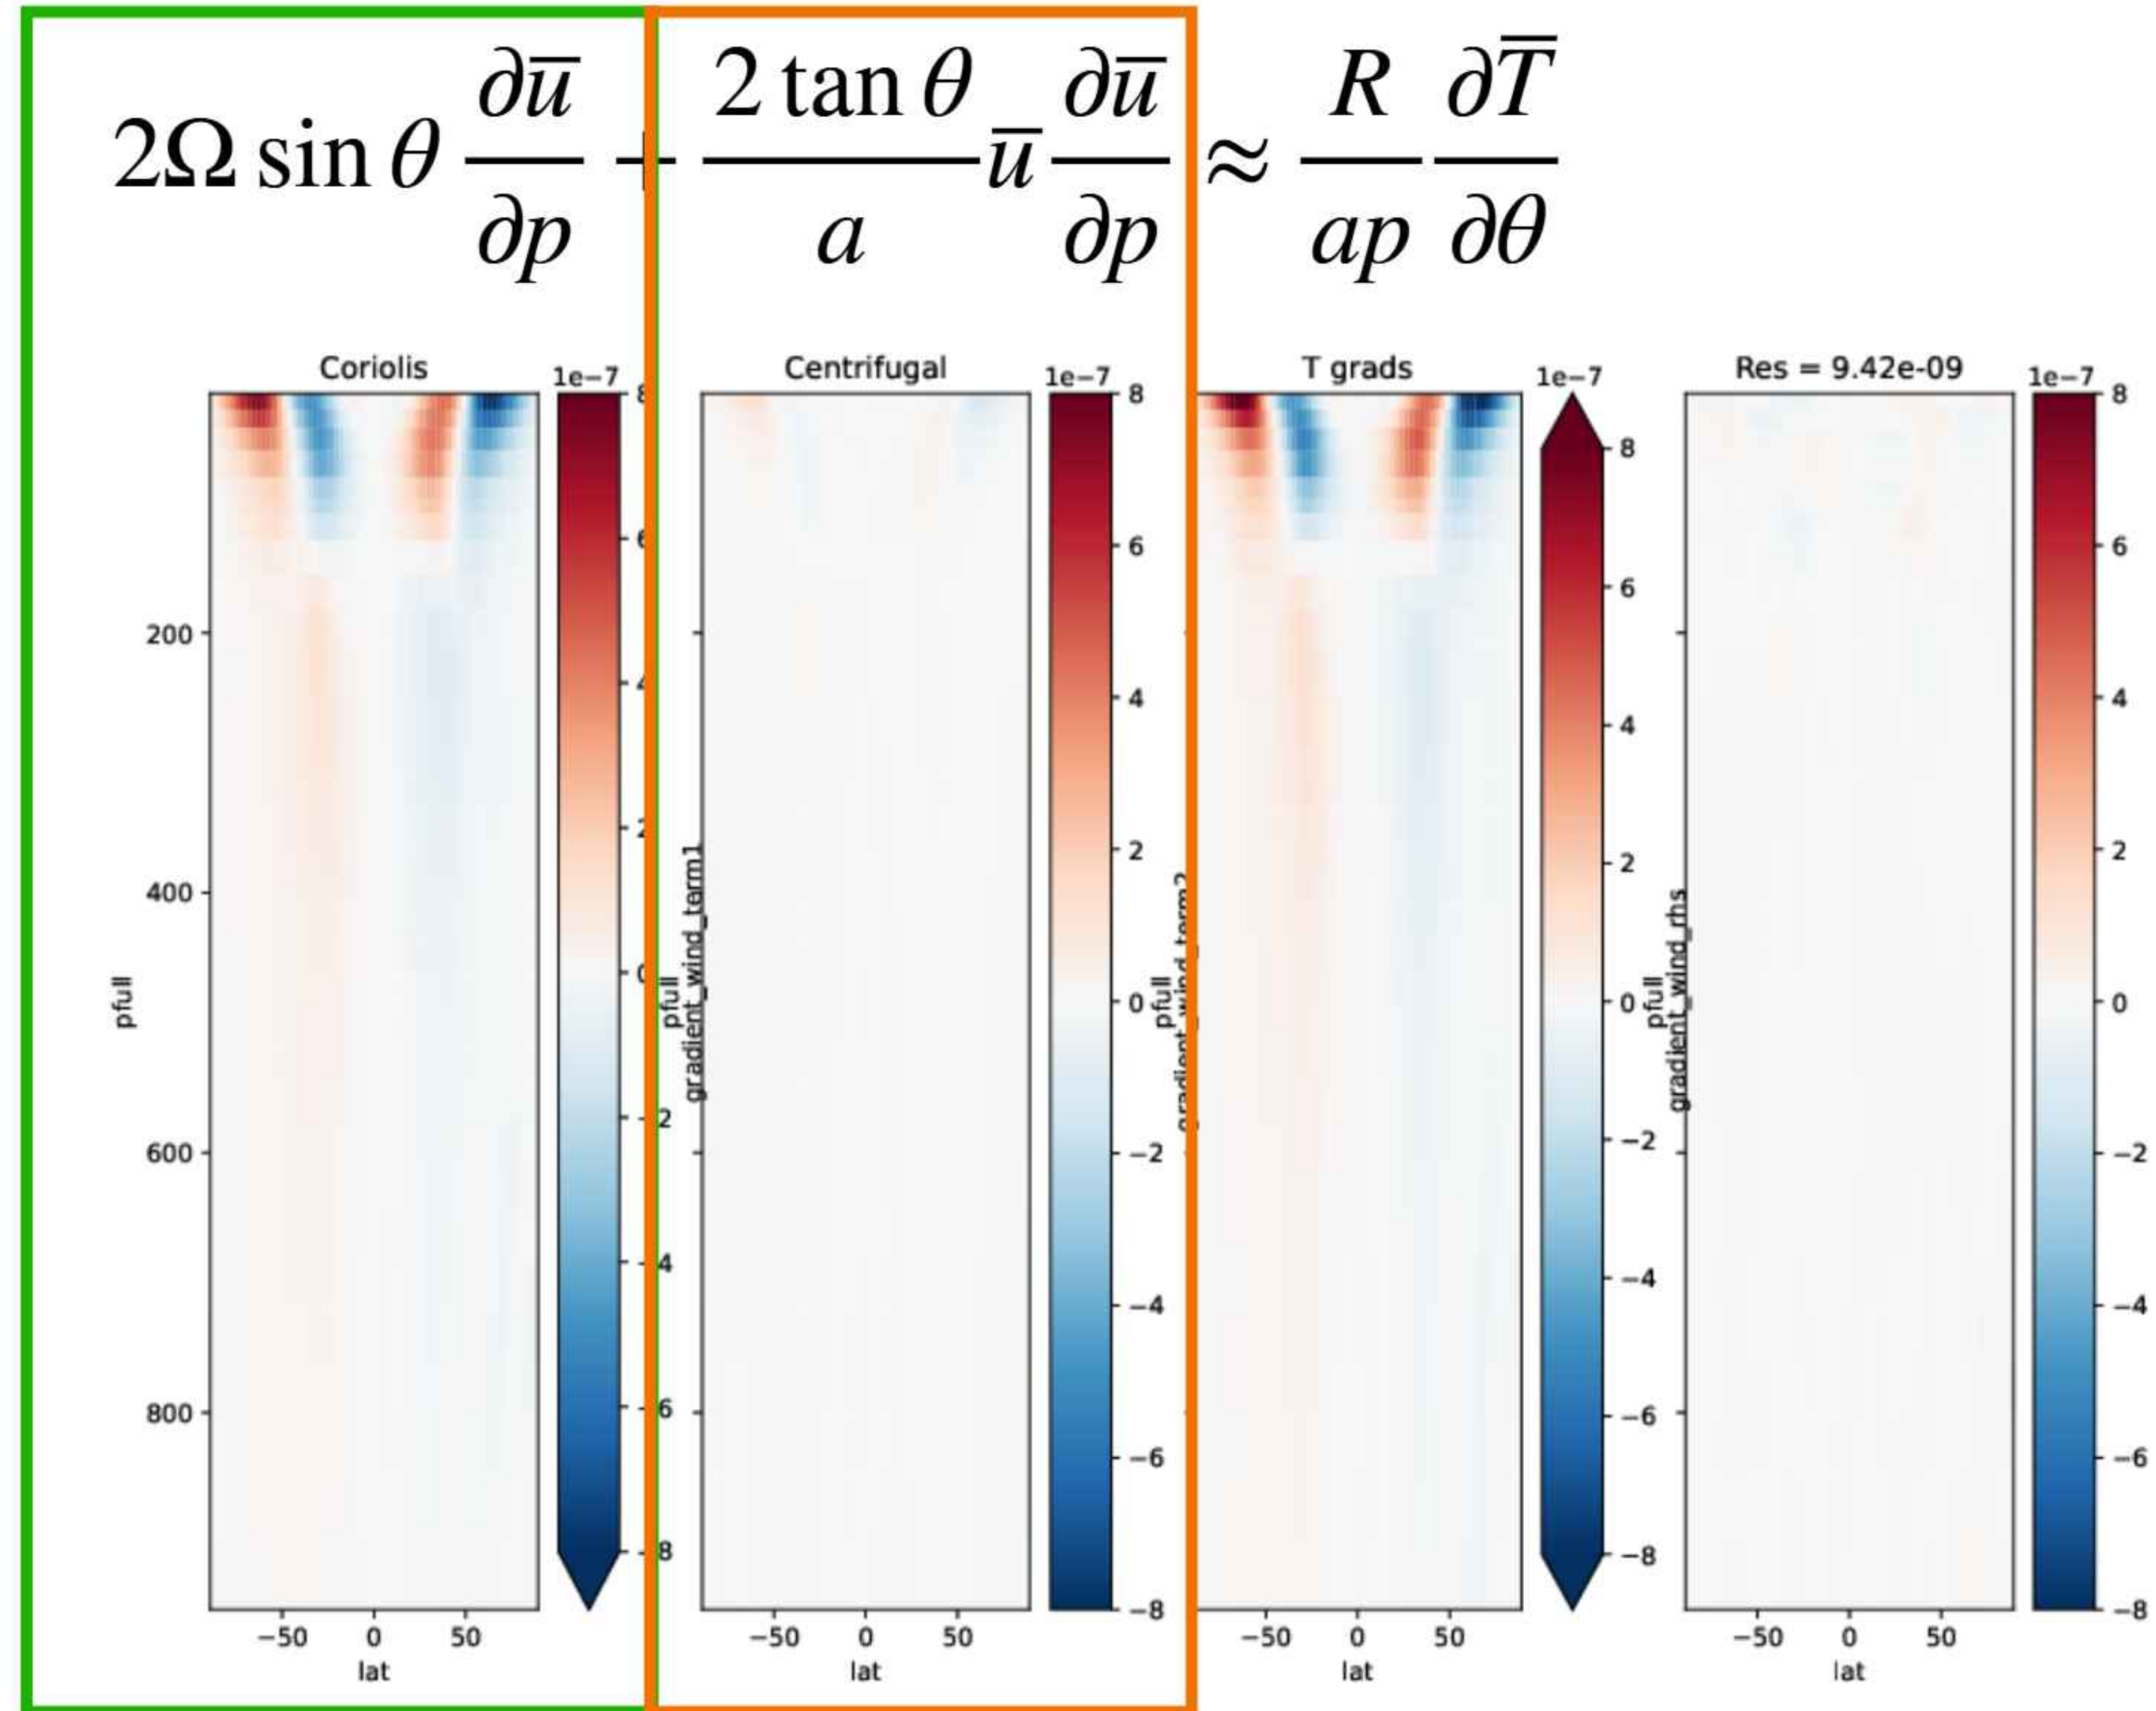
- This constraint is well respected in the Isca data

$$2\Omega \sin \theta \frac{\partial \bar{u}}{\partial p} + \frac{2 \tan \theta}{a} \bar{u} \frac{\partial \bar{u}}{\partial p} \approx \frac{R}{ap} \frac{\partial \bar{T}}{\partial \theta}$$



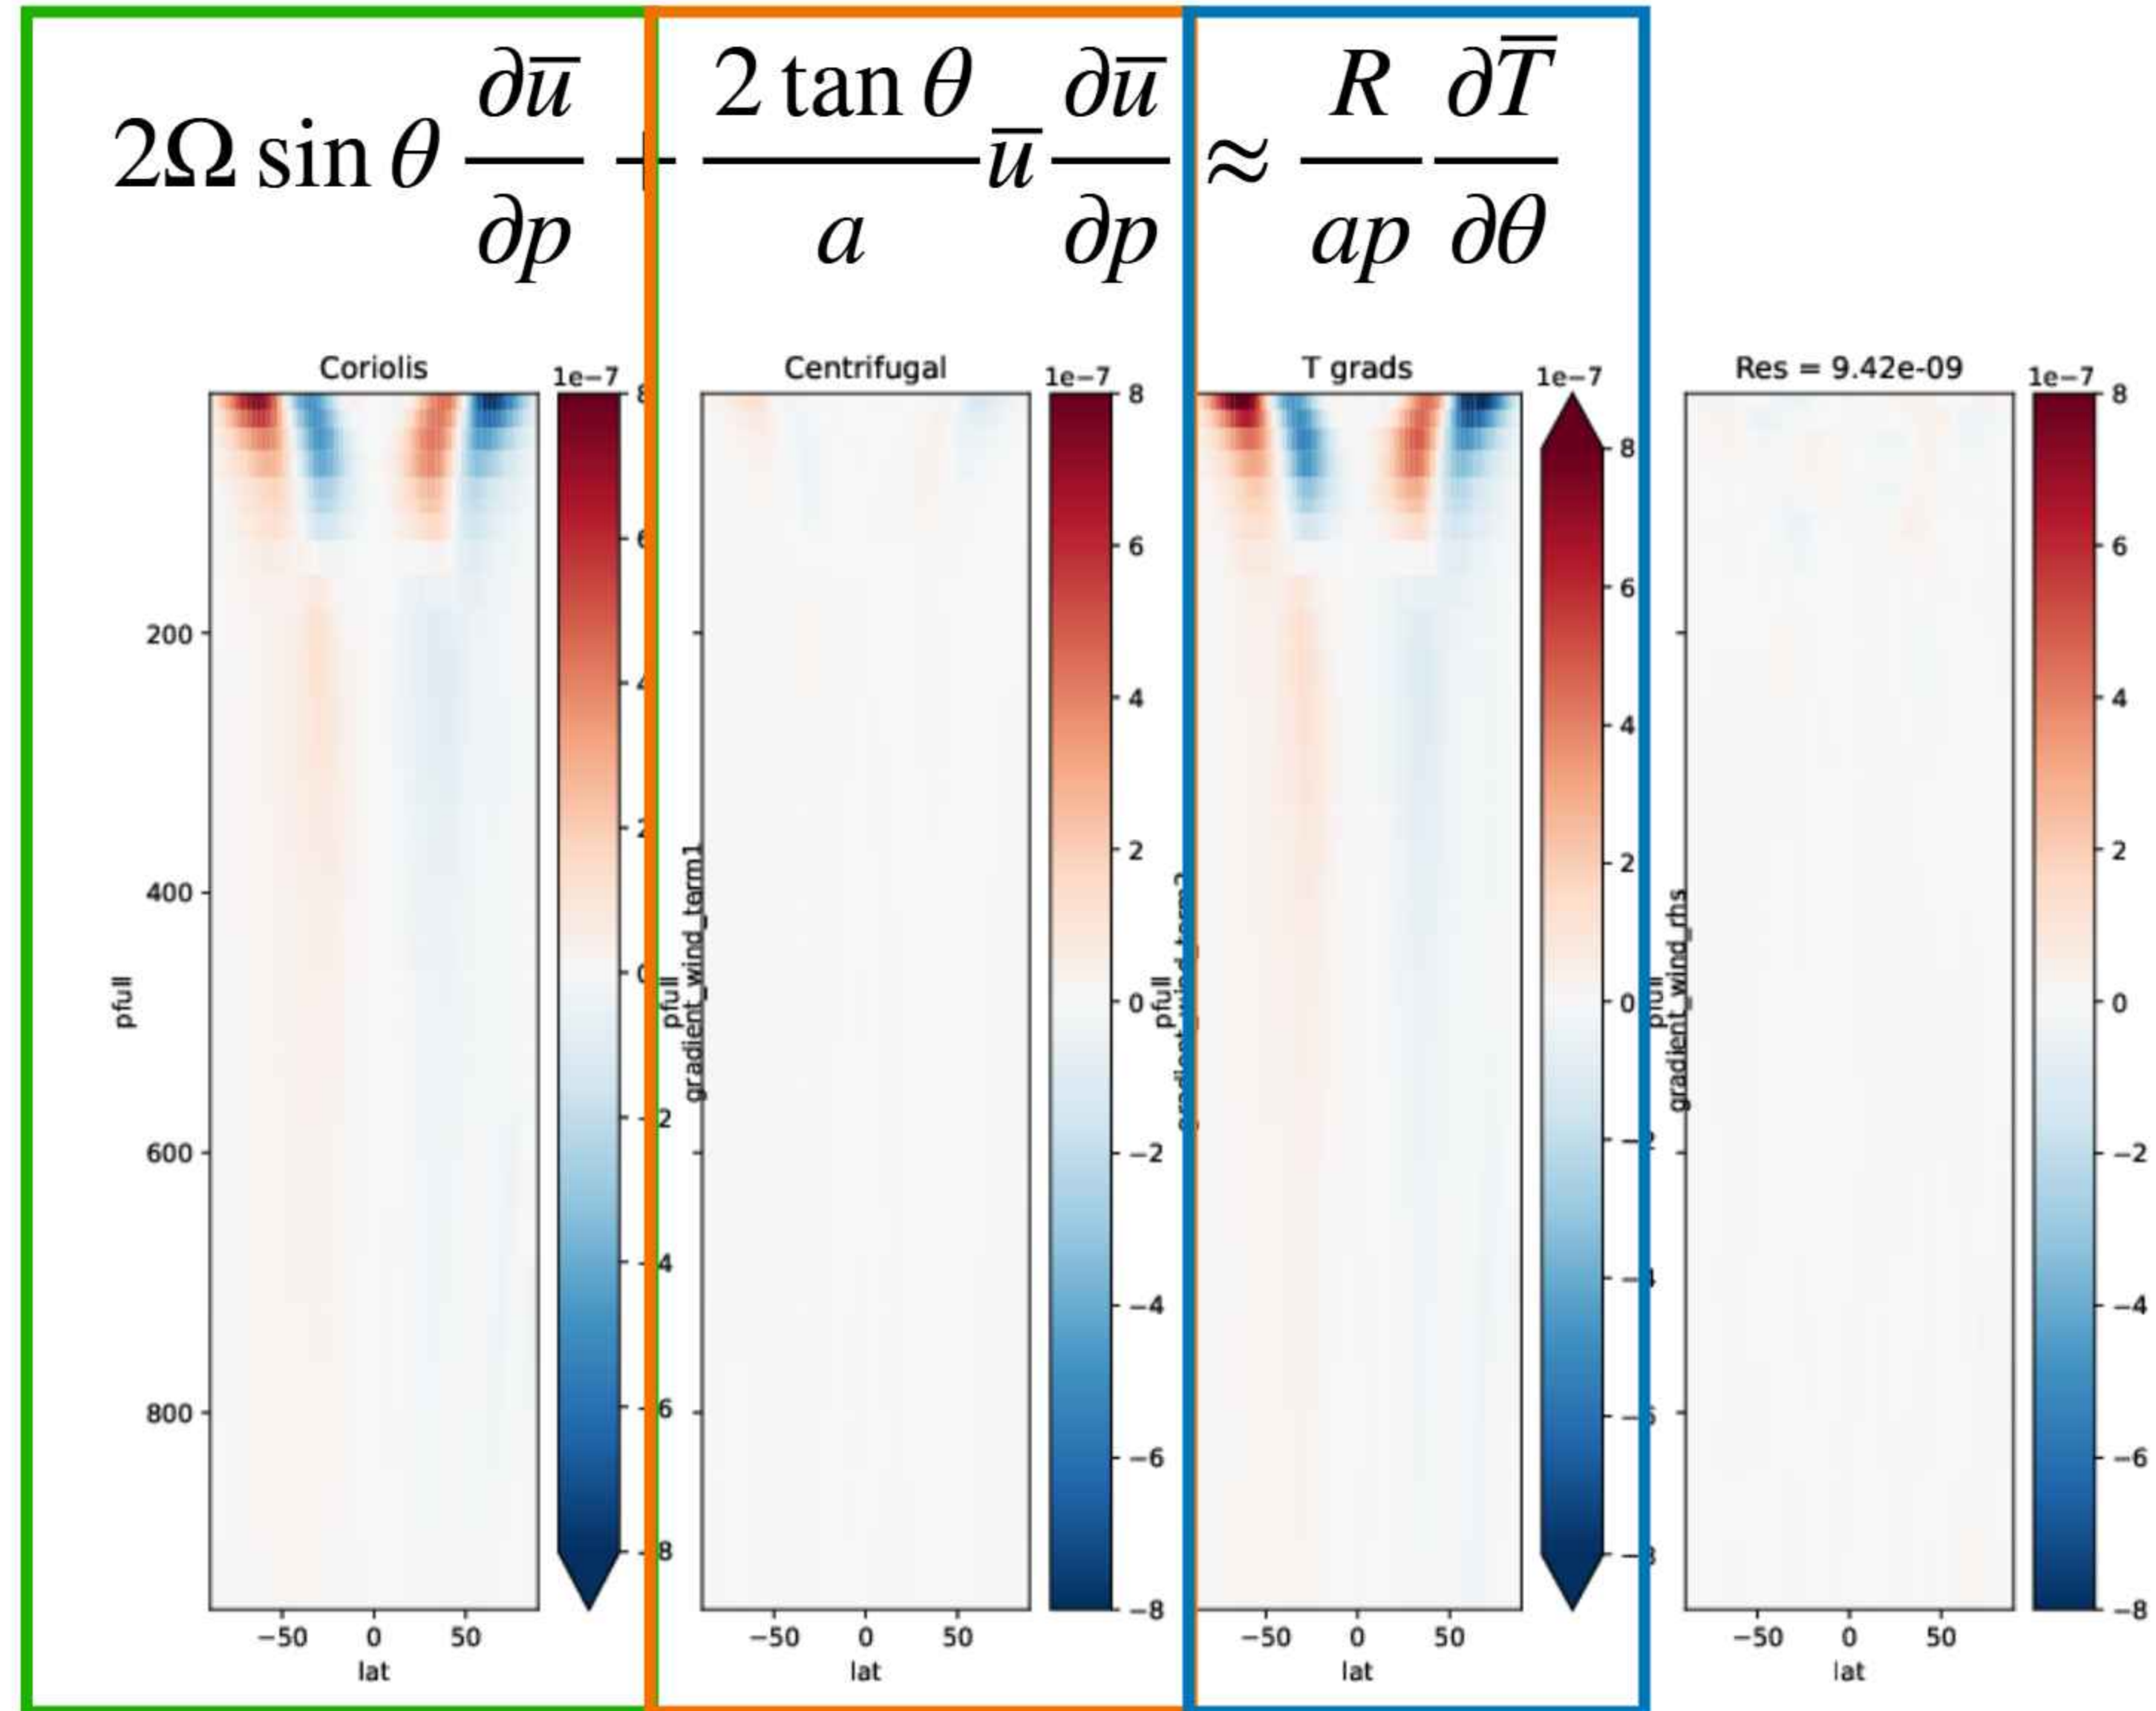
Incorporating dynamical constraints

- This constraint is well respected in the Isca data



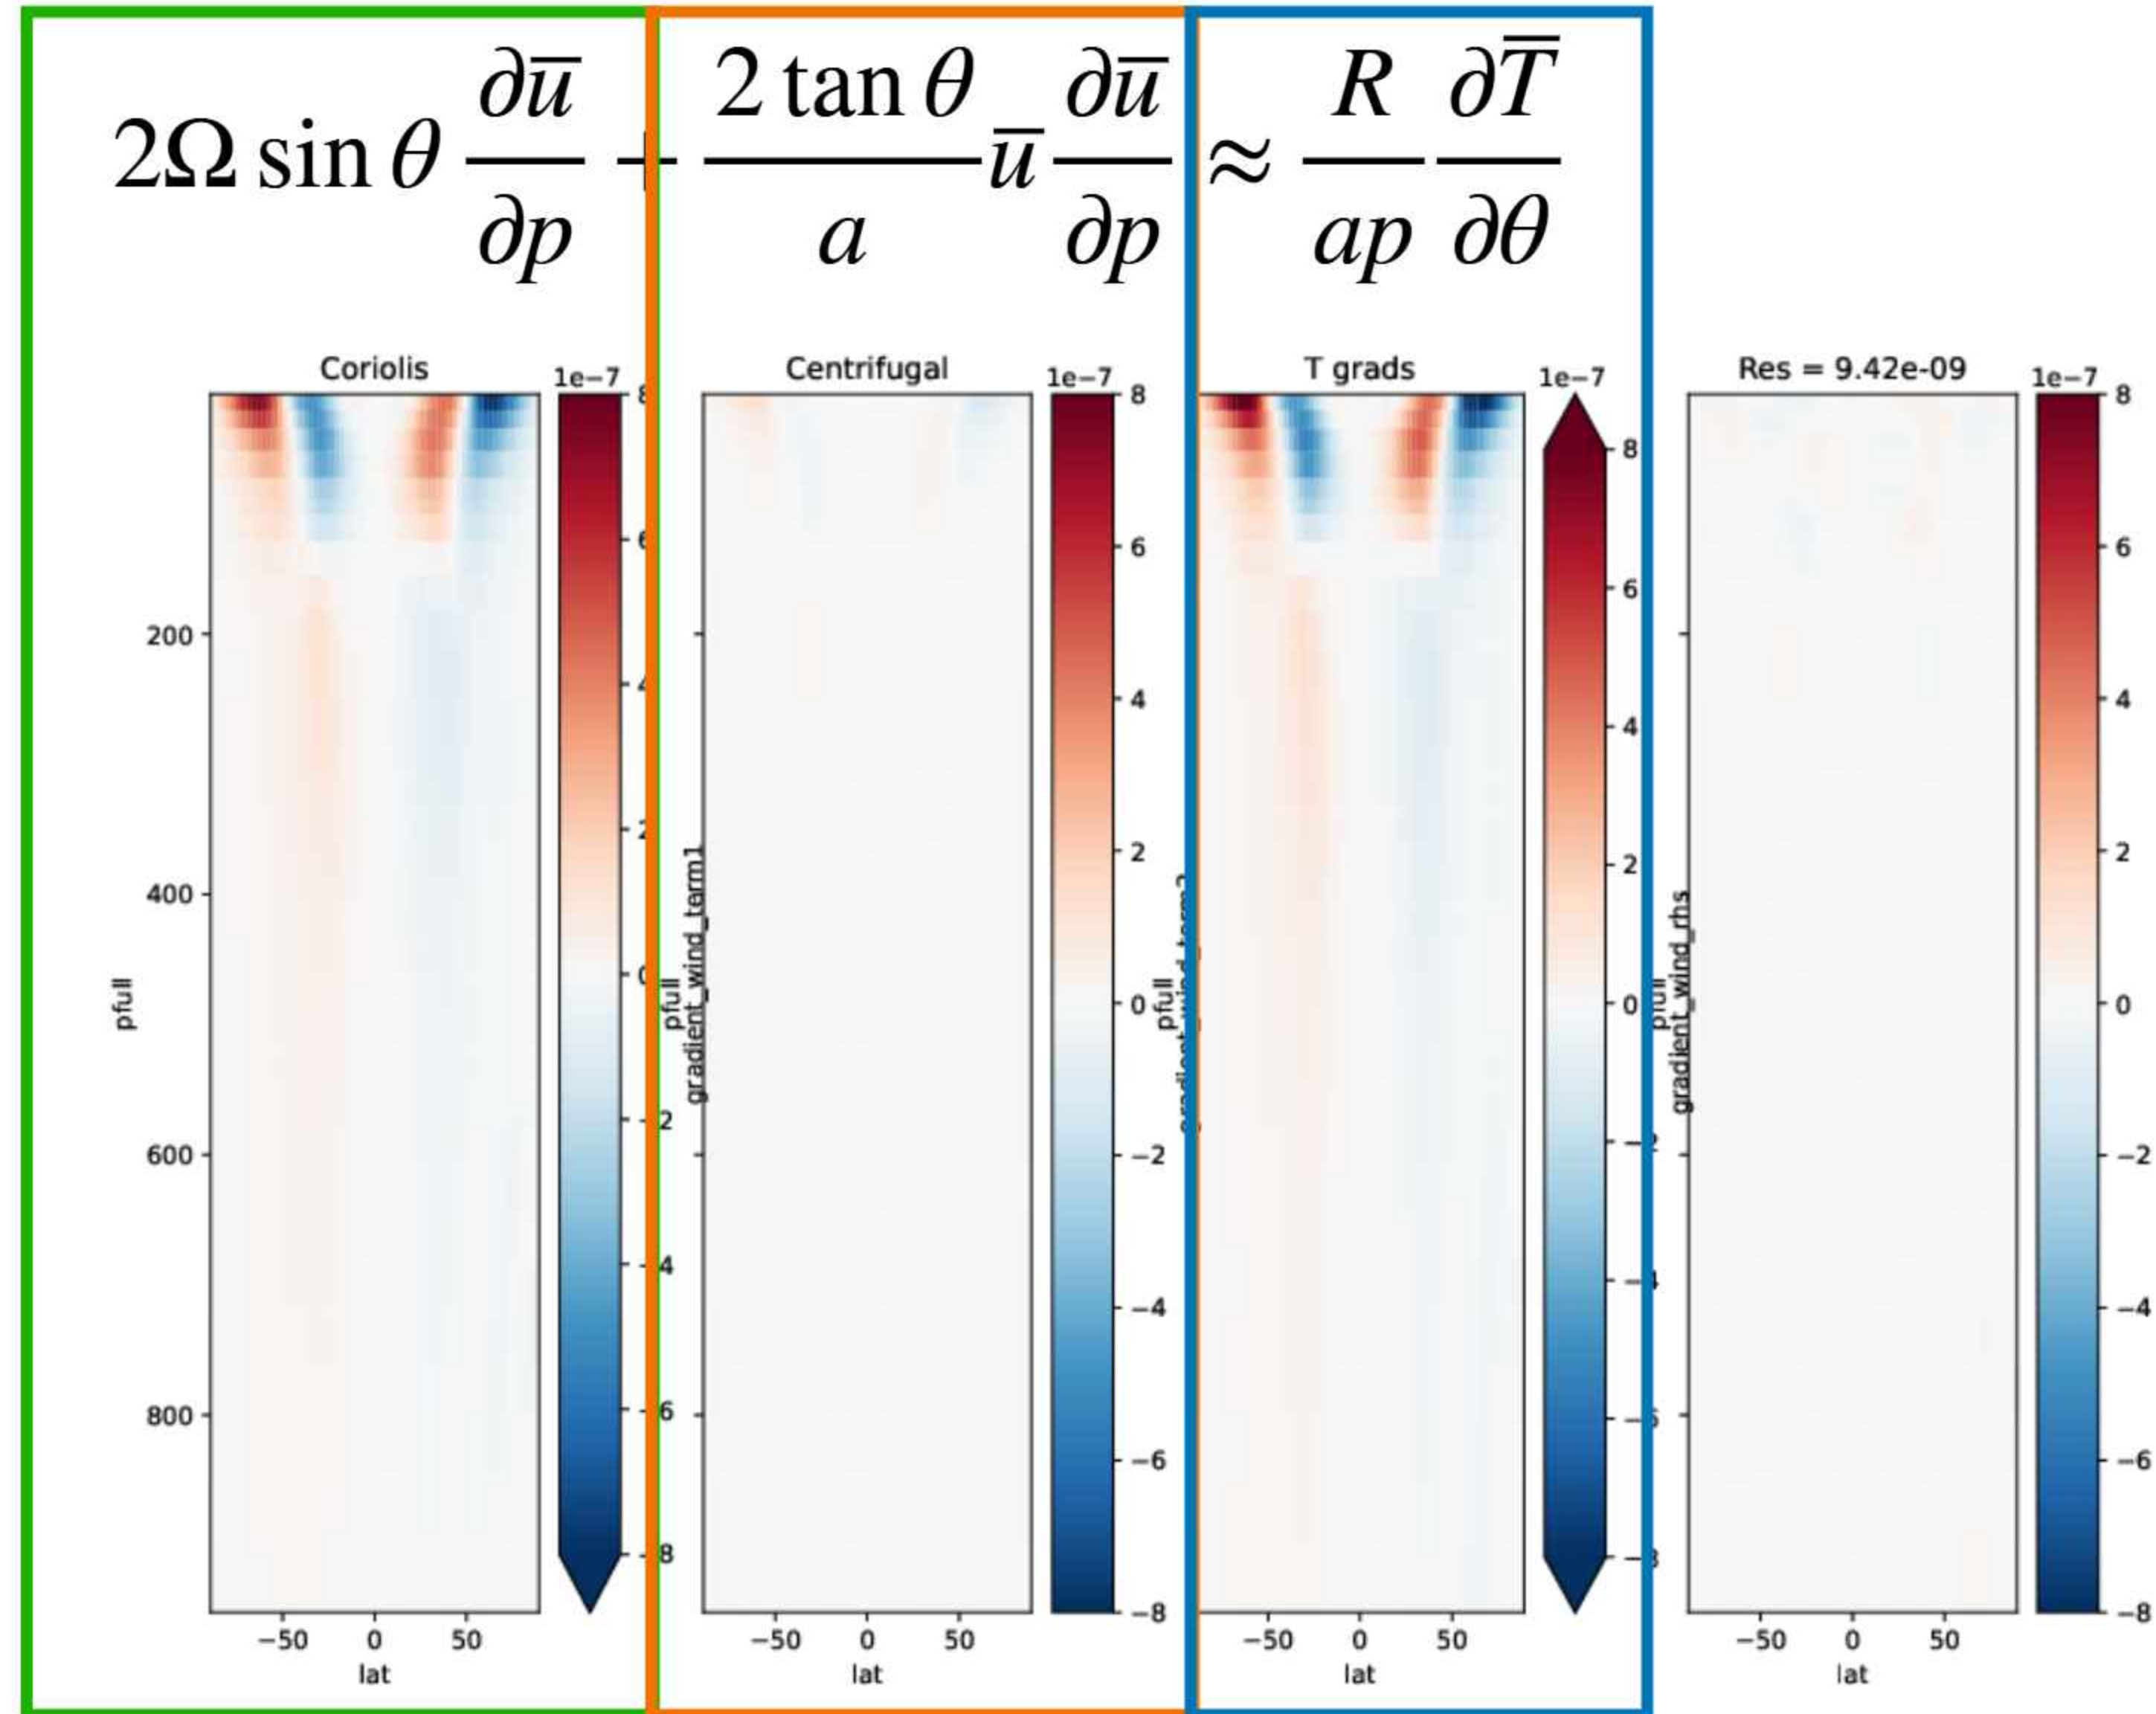
Incorporating dynamical constraints

- This constraint is well respected in the Isca data



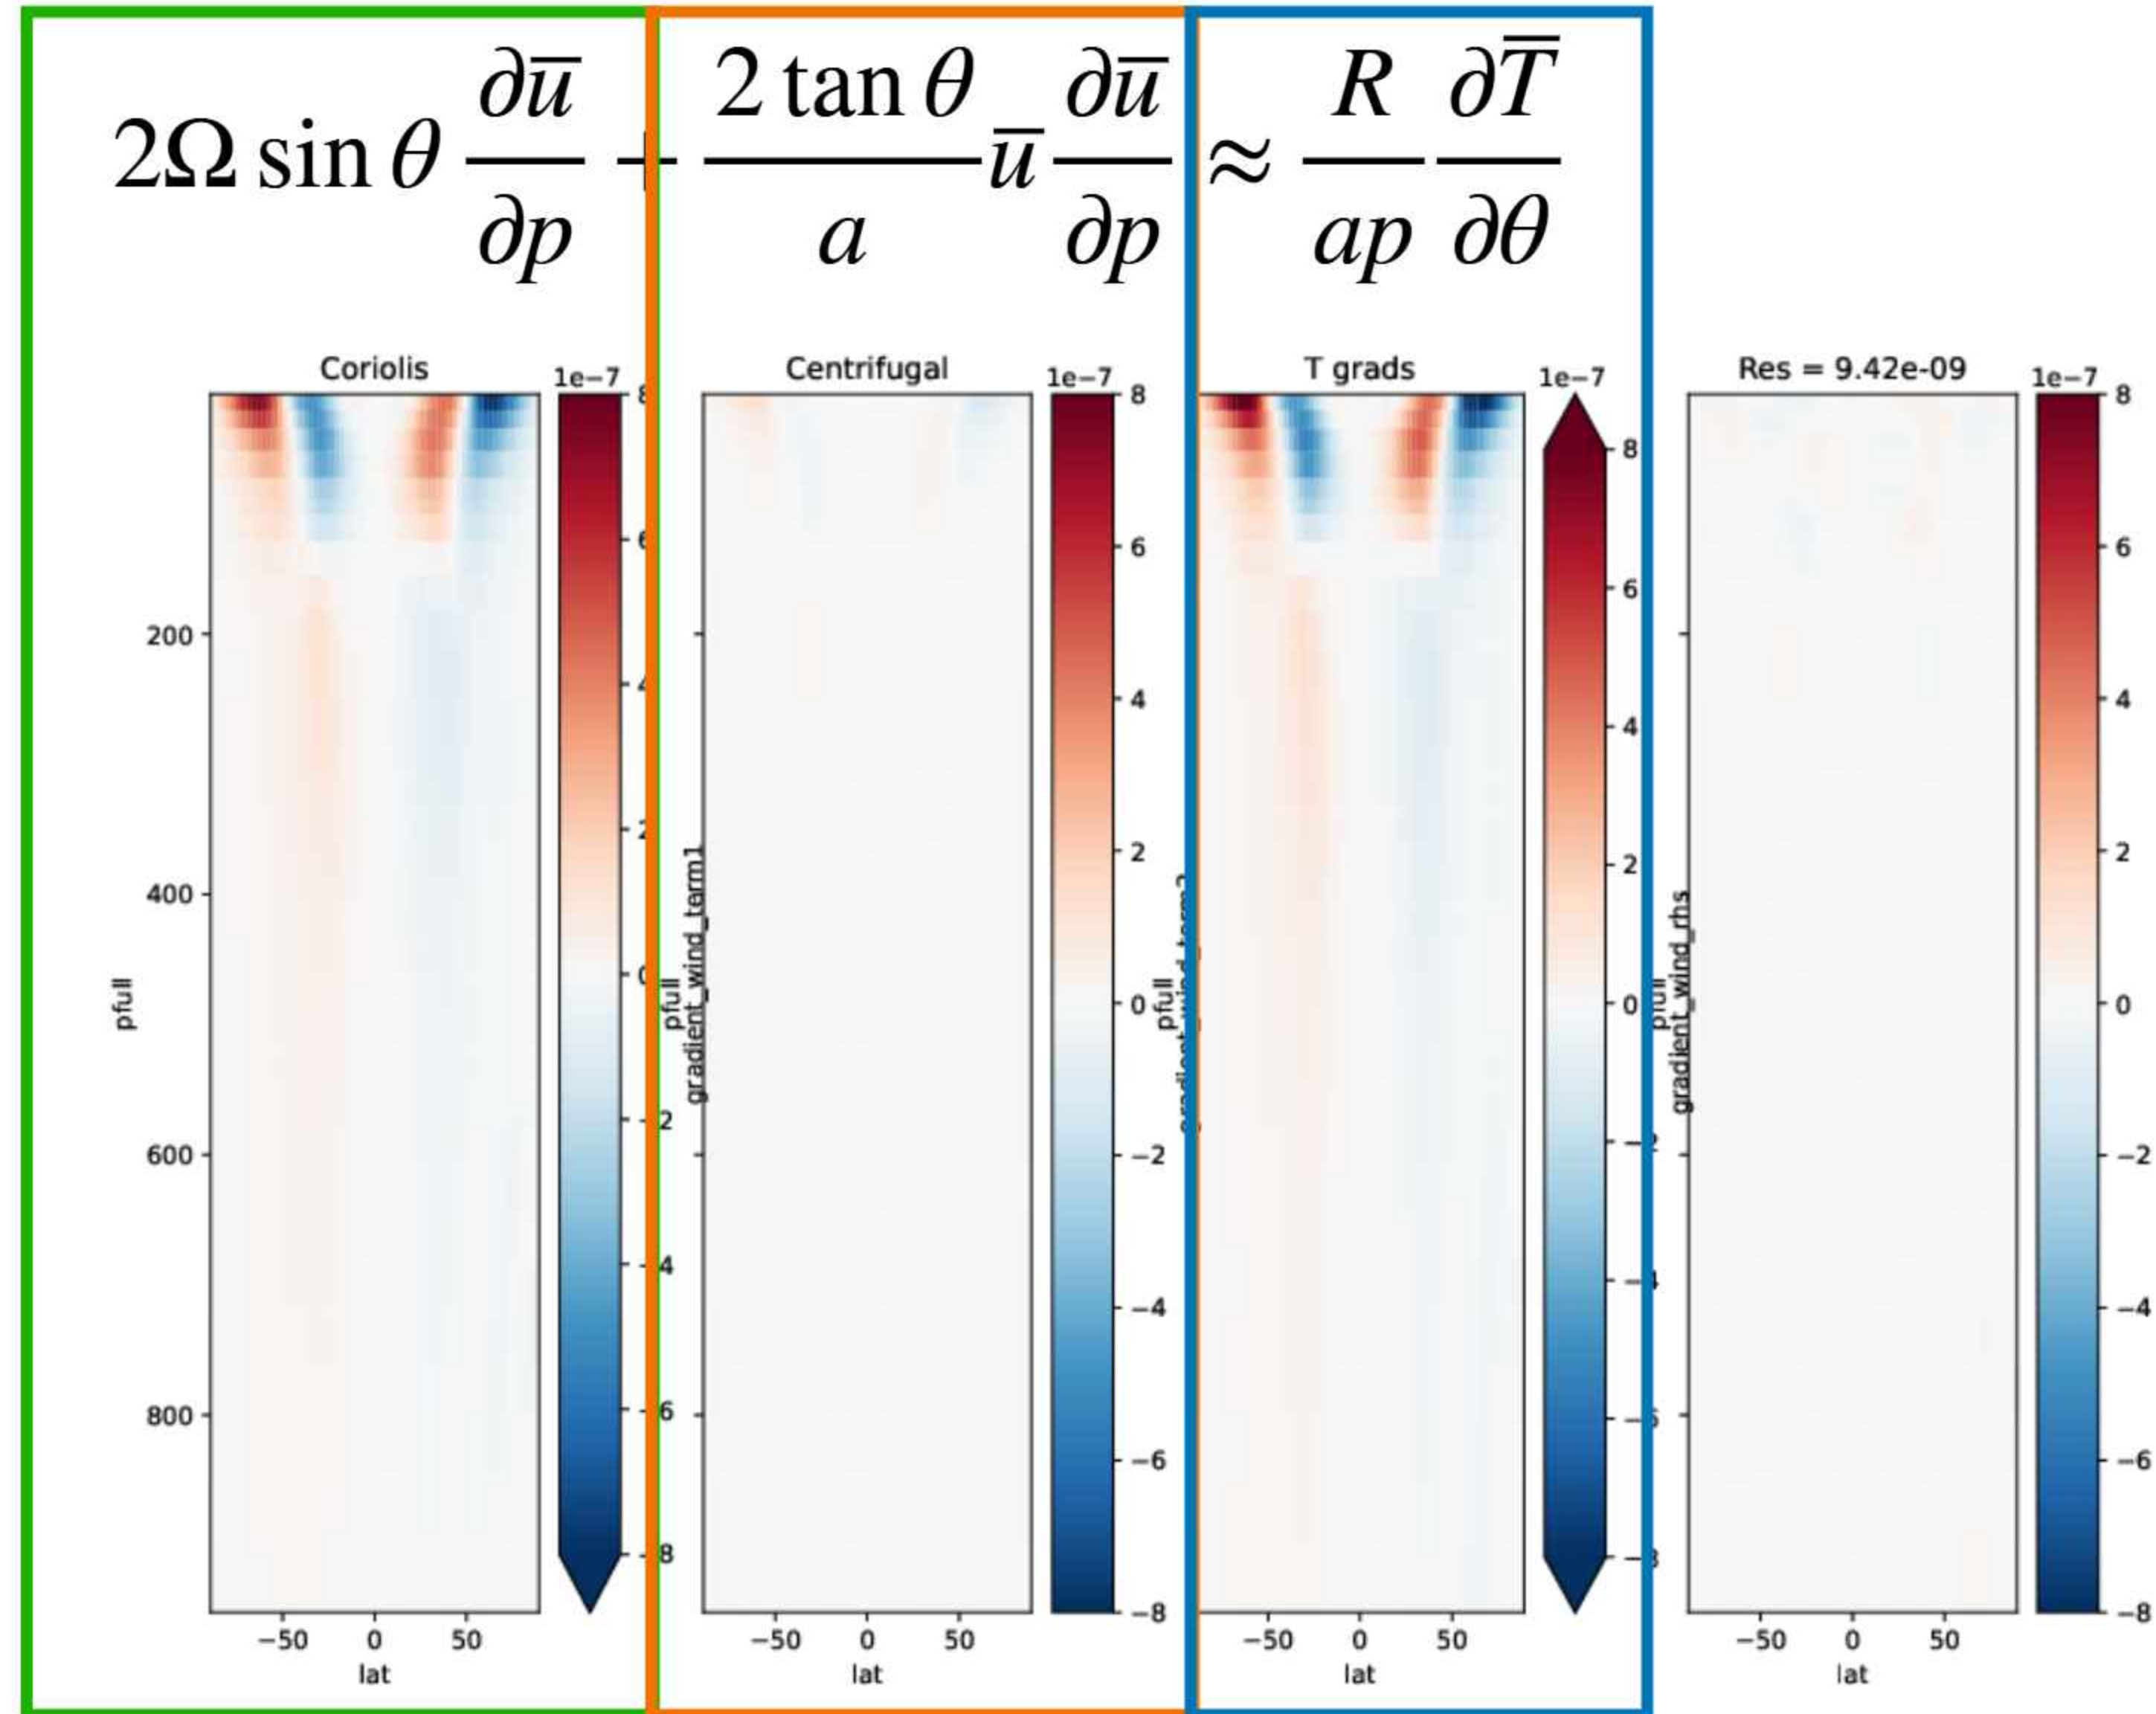
Incorporating dynamical constraints

- This constraint is well respected in the Isca data
- We have attempted to incorporate this constraint into the loss function of our network (physics-informed NN)



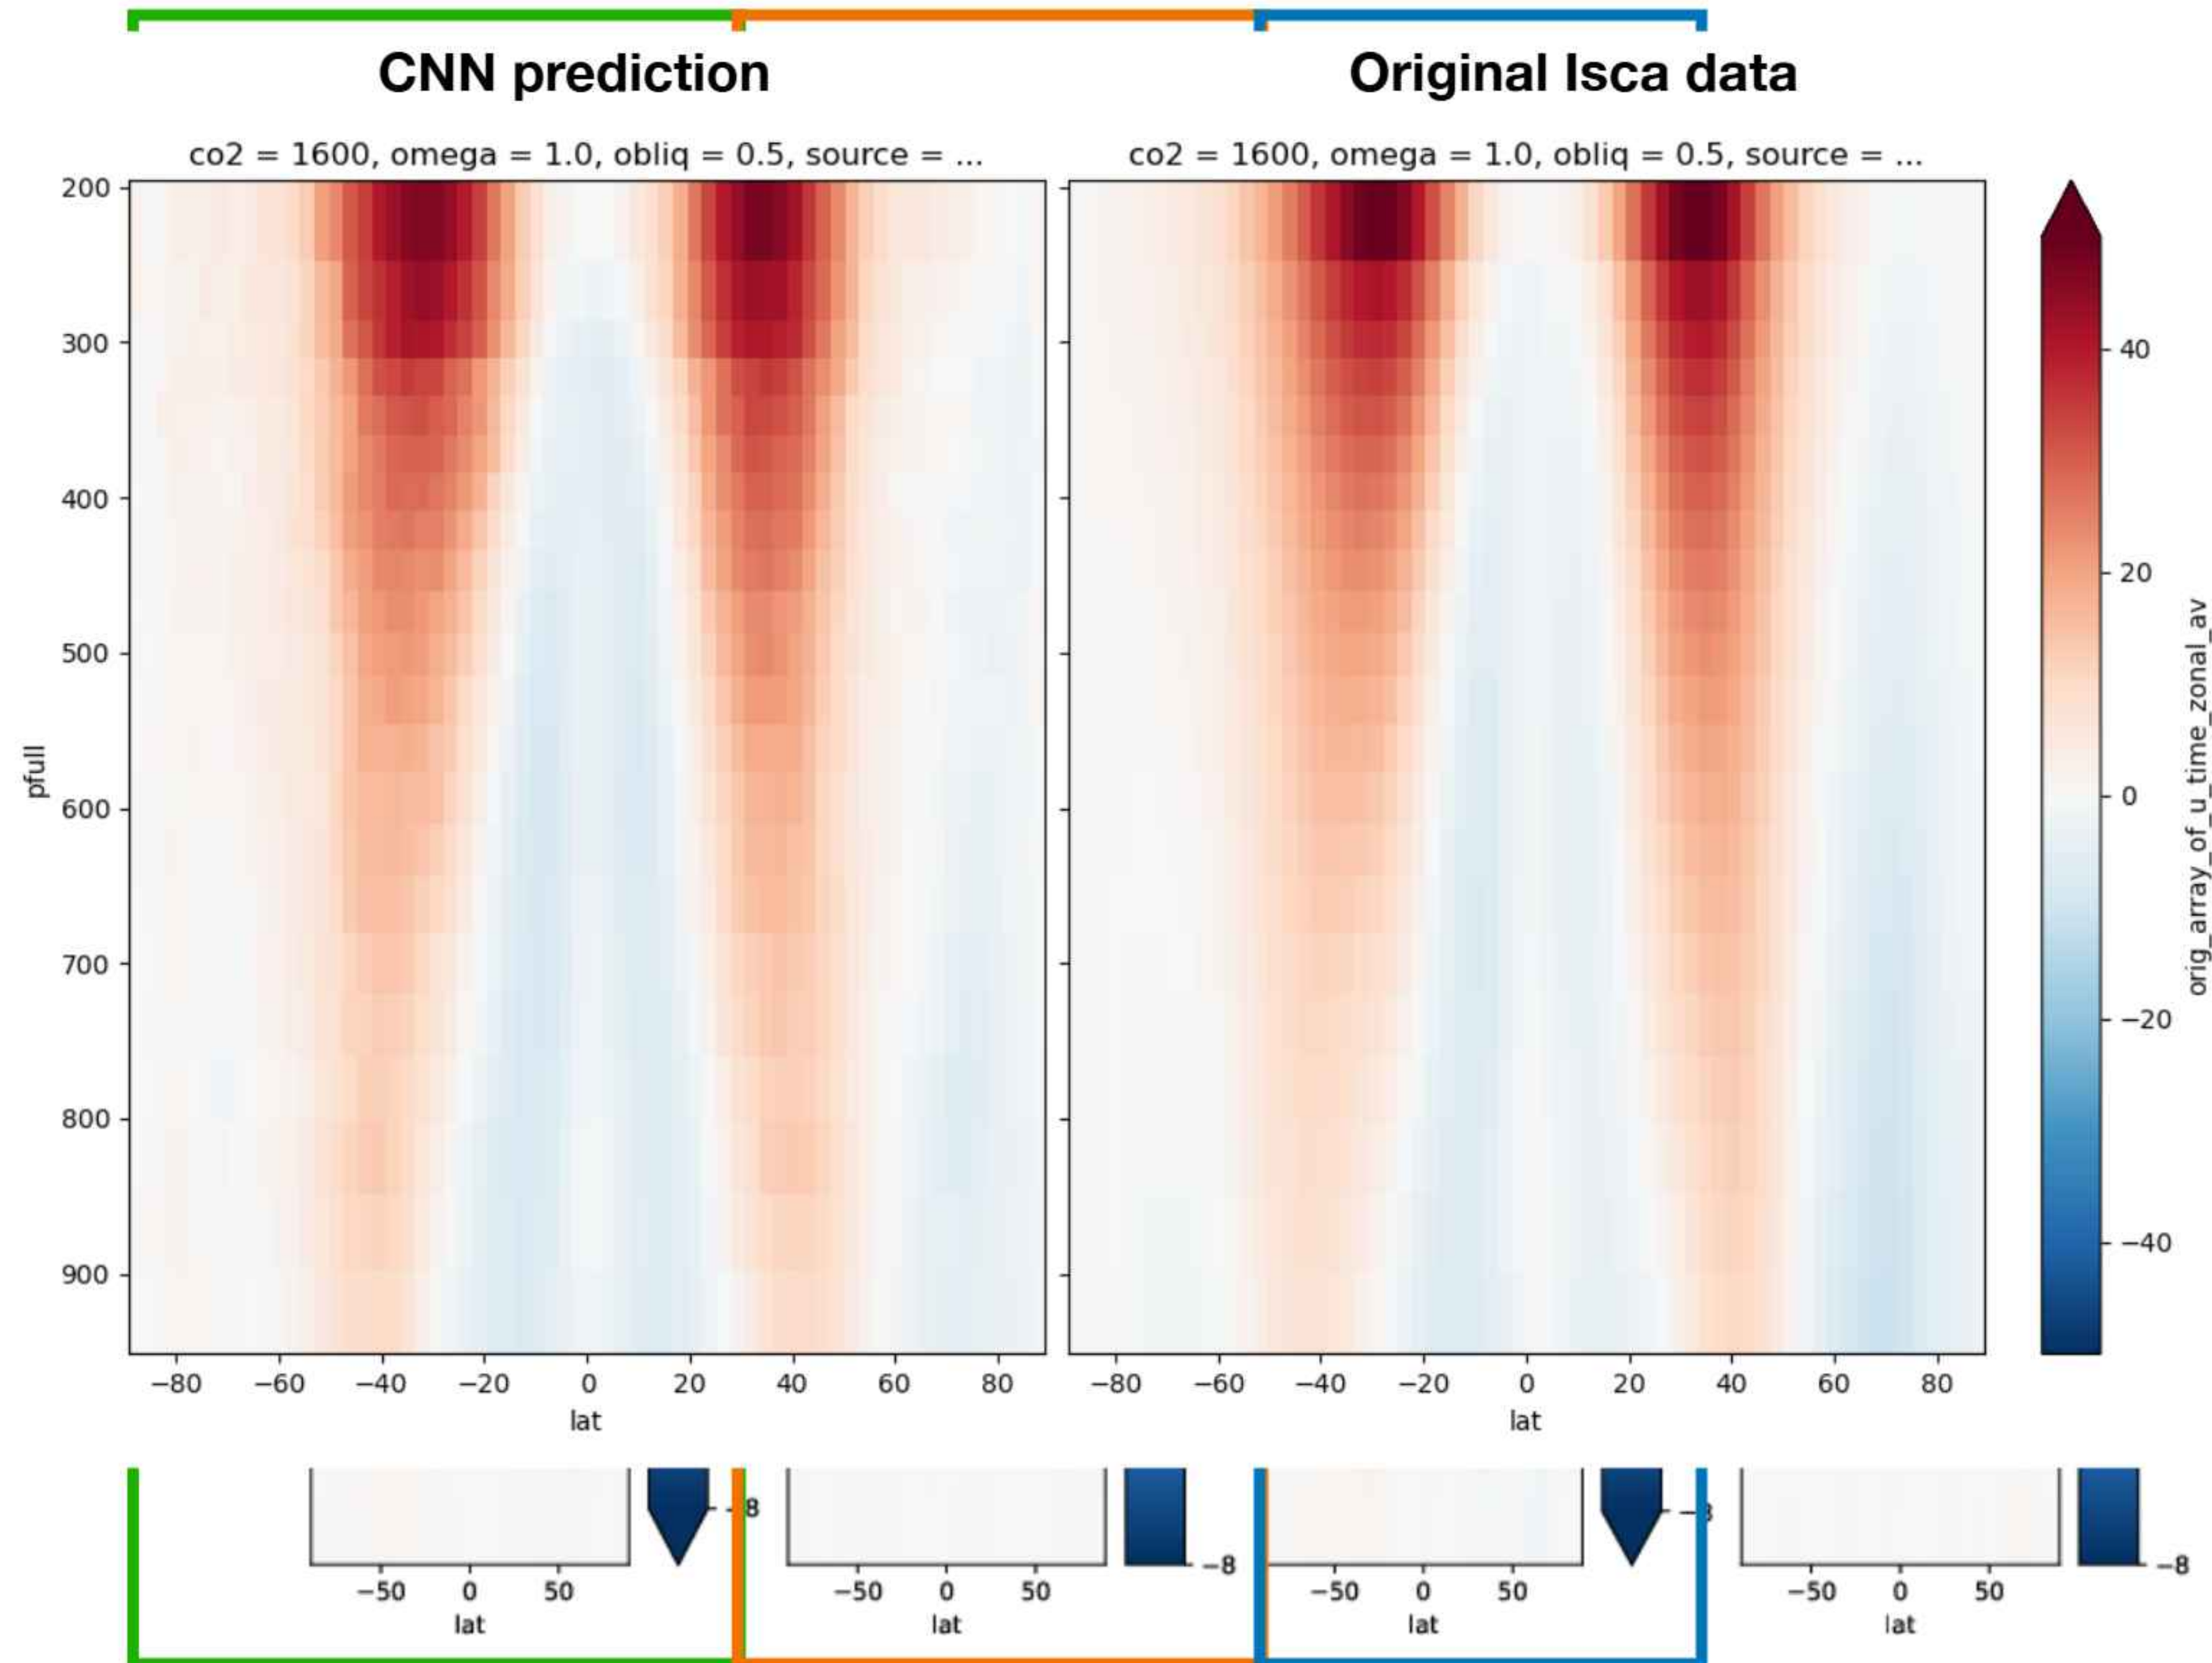
Incorporating dynamical constraints

- This constraint is well respected in the Isca data
- We have attempted to incorporate this constraint into the loss function of our network (physics-informed NN)
- Our network's predictions aren't smooth enough for the constraint to help (gradients in latitude and pressure are dominated by small-scale noise)



Incorporating dynamical constraints

- This constraint is well respected in the Isca data
- We have attempted to incorporate this constraint into the loss function of our network (physics-informed NN)
- Our network's predictions aren't smooth enough for the constraint to help (gradients in latitude and pressure are dominated by small-scale noise)



Removing small-scale noise

Removing small-scale noise

- To remove small-scale noise, we tried:

Removing small-scale noise

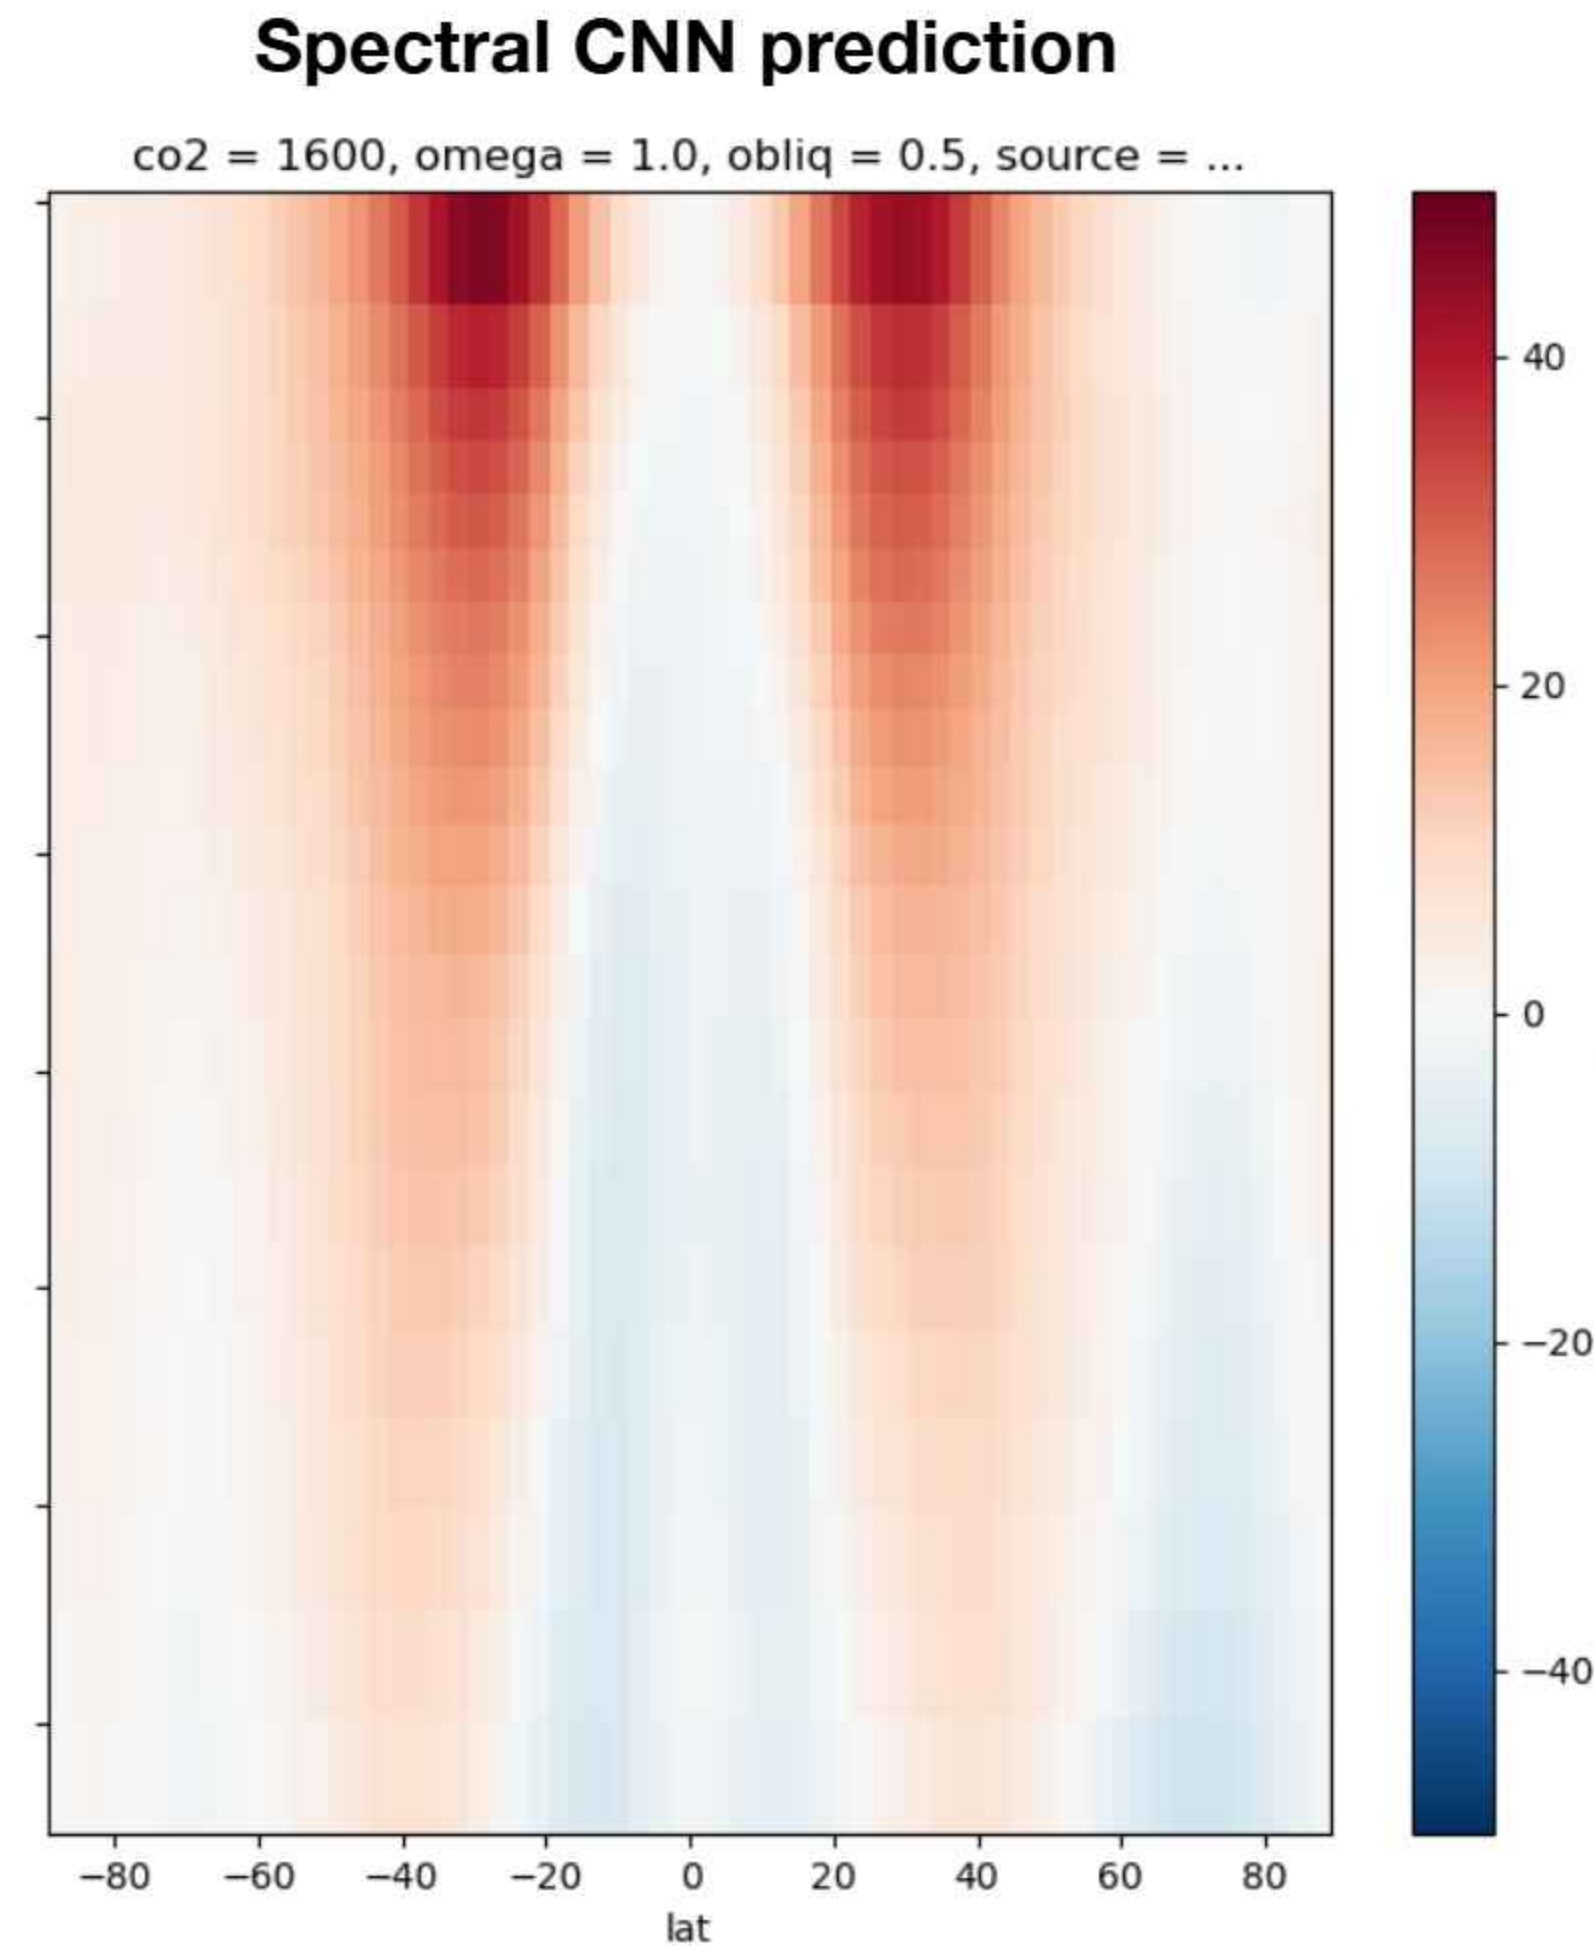
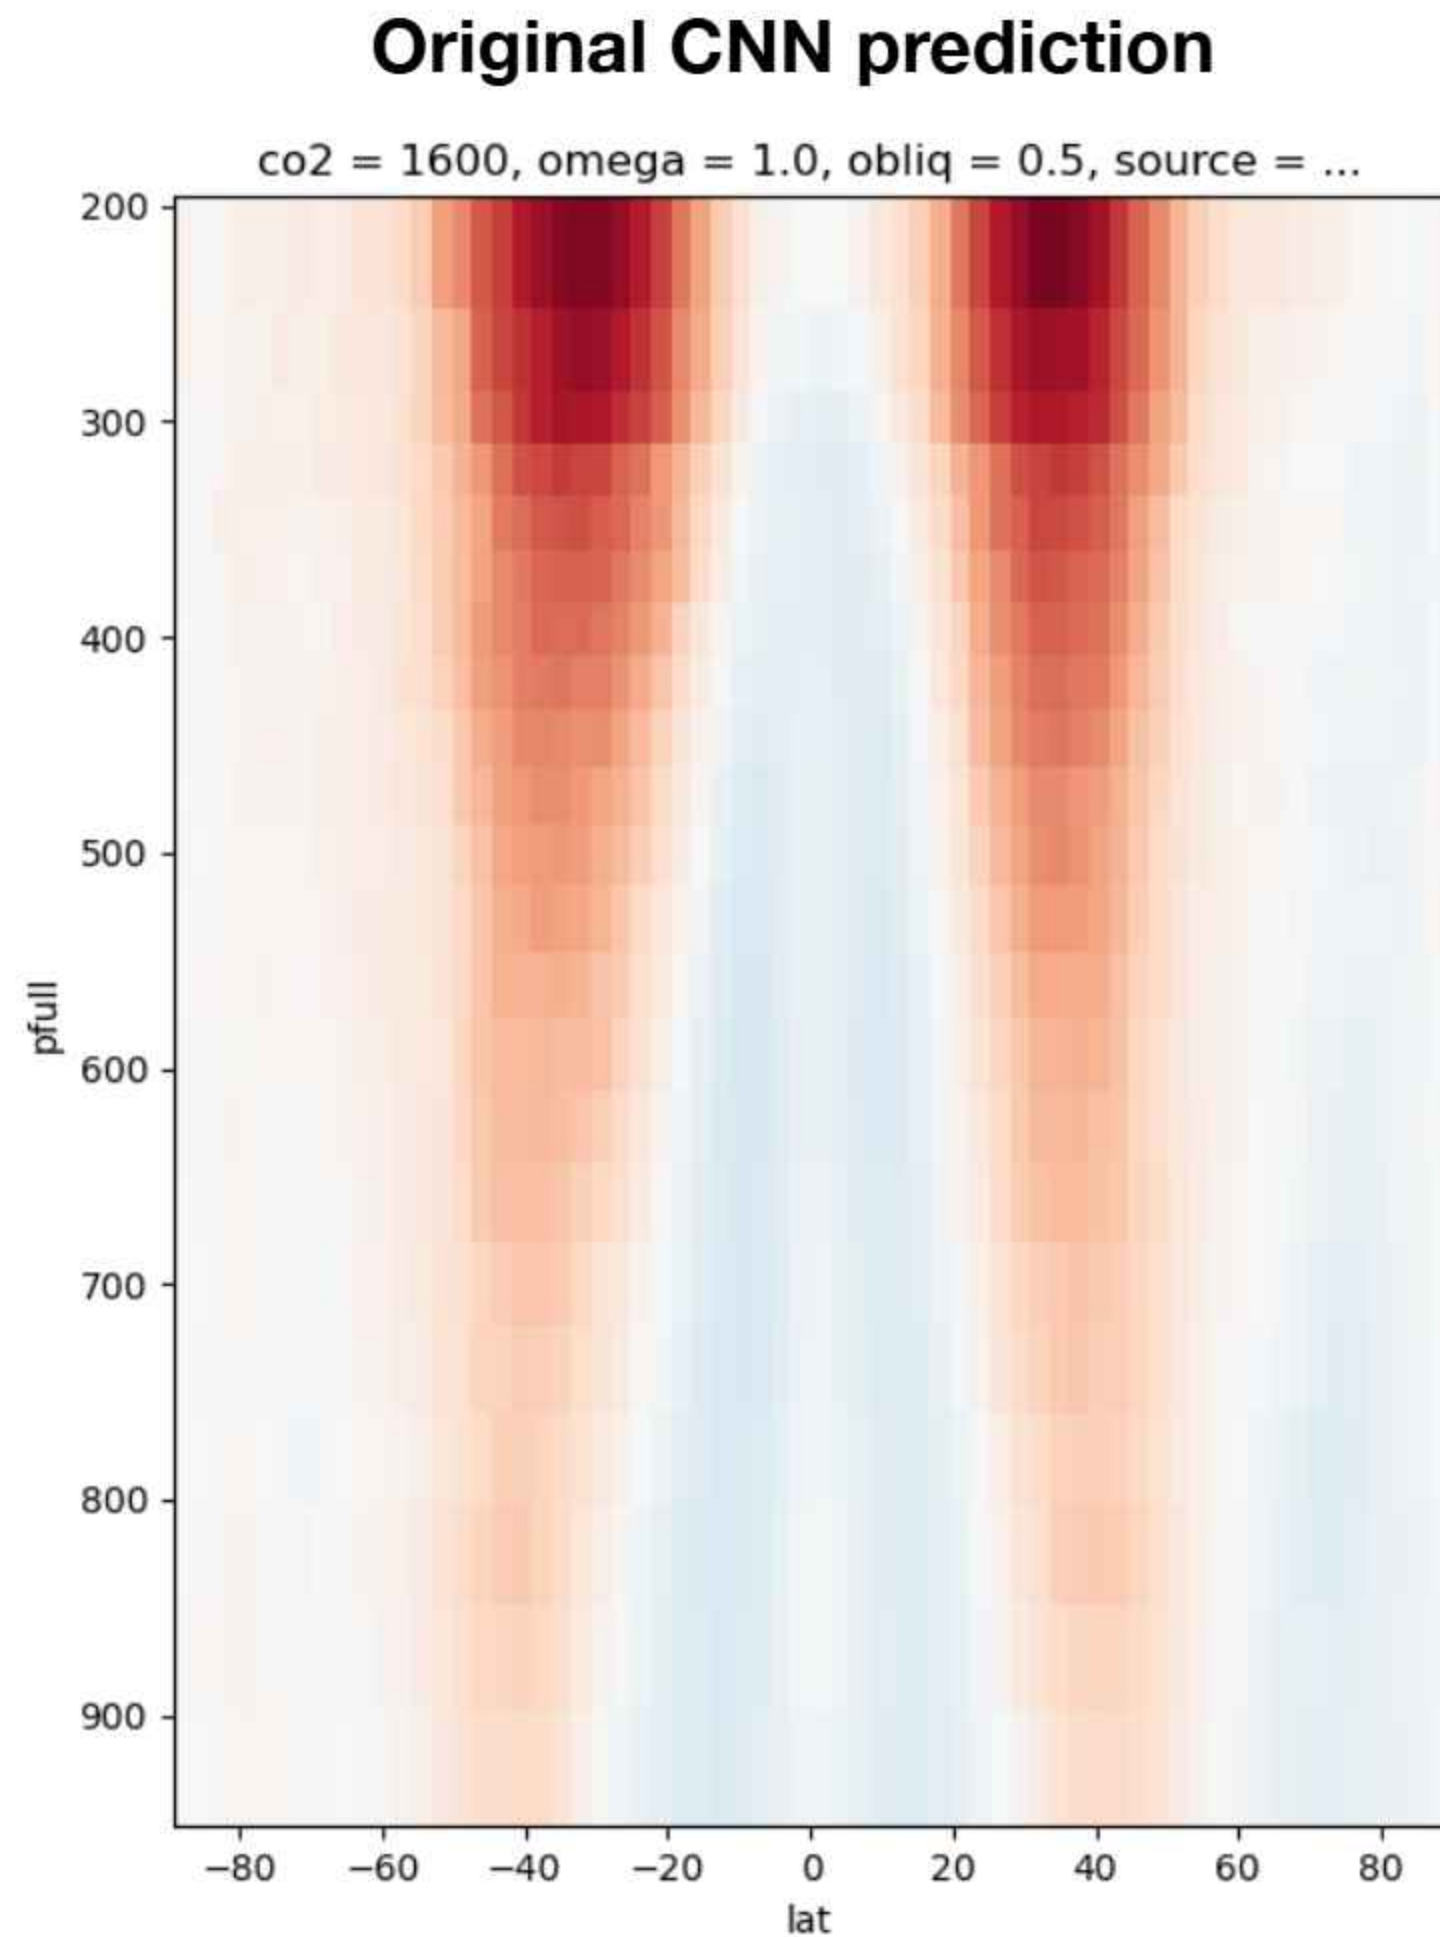
- To remove small-scale noise, we tried:
 - adding a term penalising pixel-to-pixel variation in the output

Removing small-scale noise

- To remove small-scale noise, we tried:
 - adding a term penalising pixel-to-pixel variation in the output
 - **Making predictions of the fields in spectral space, focussing on large-scale structure**

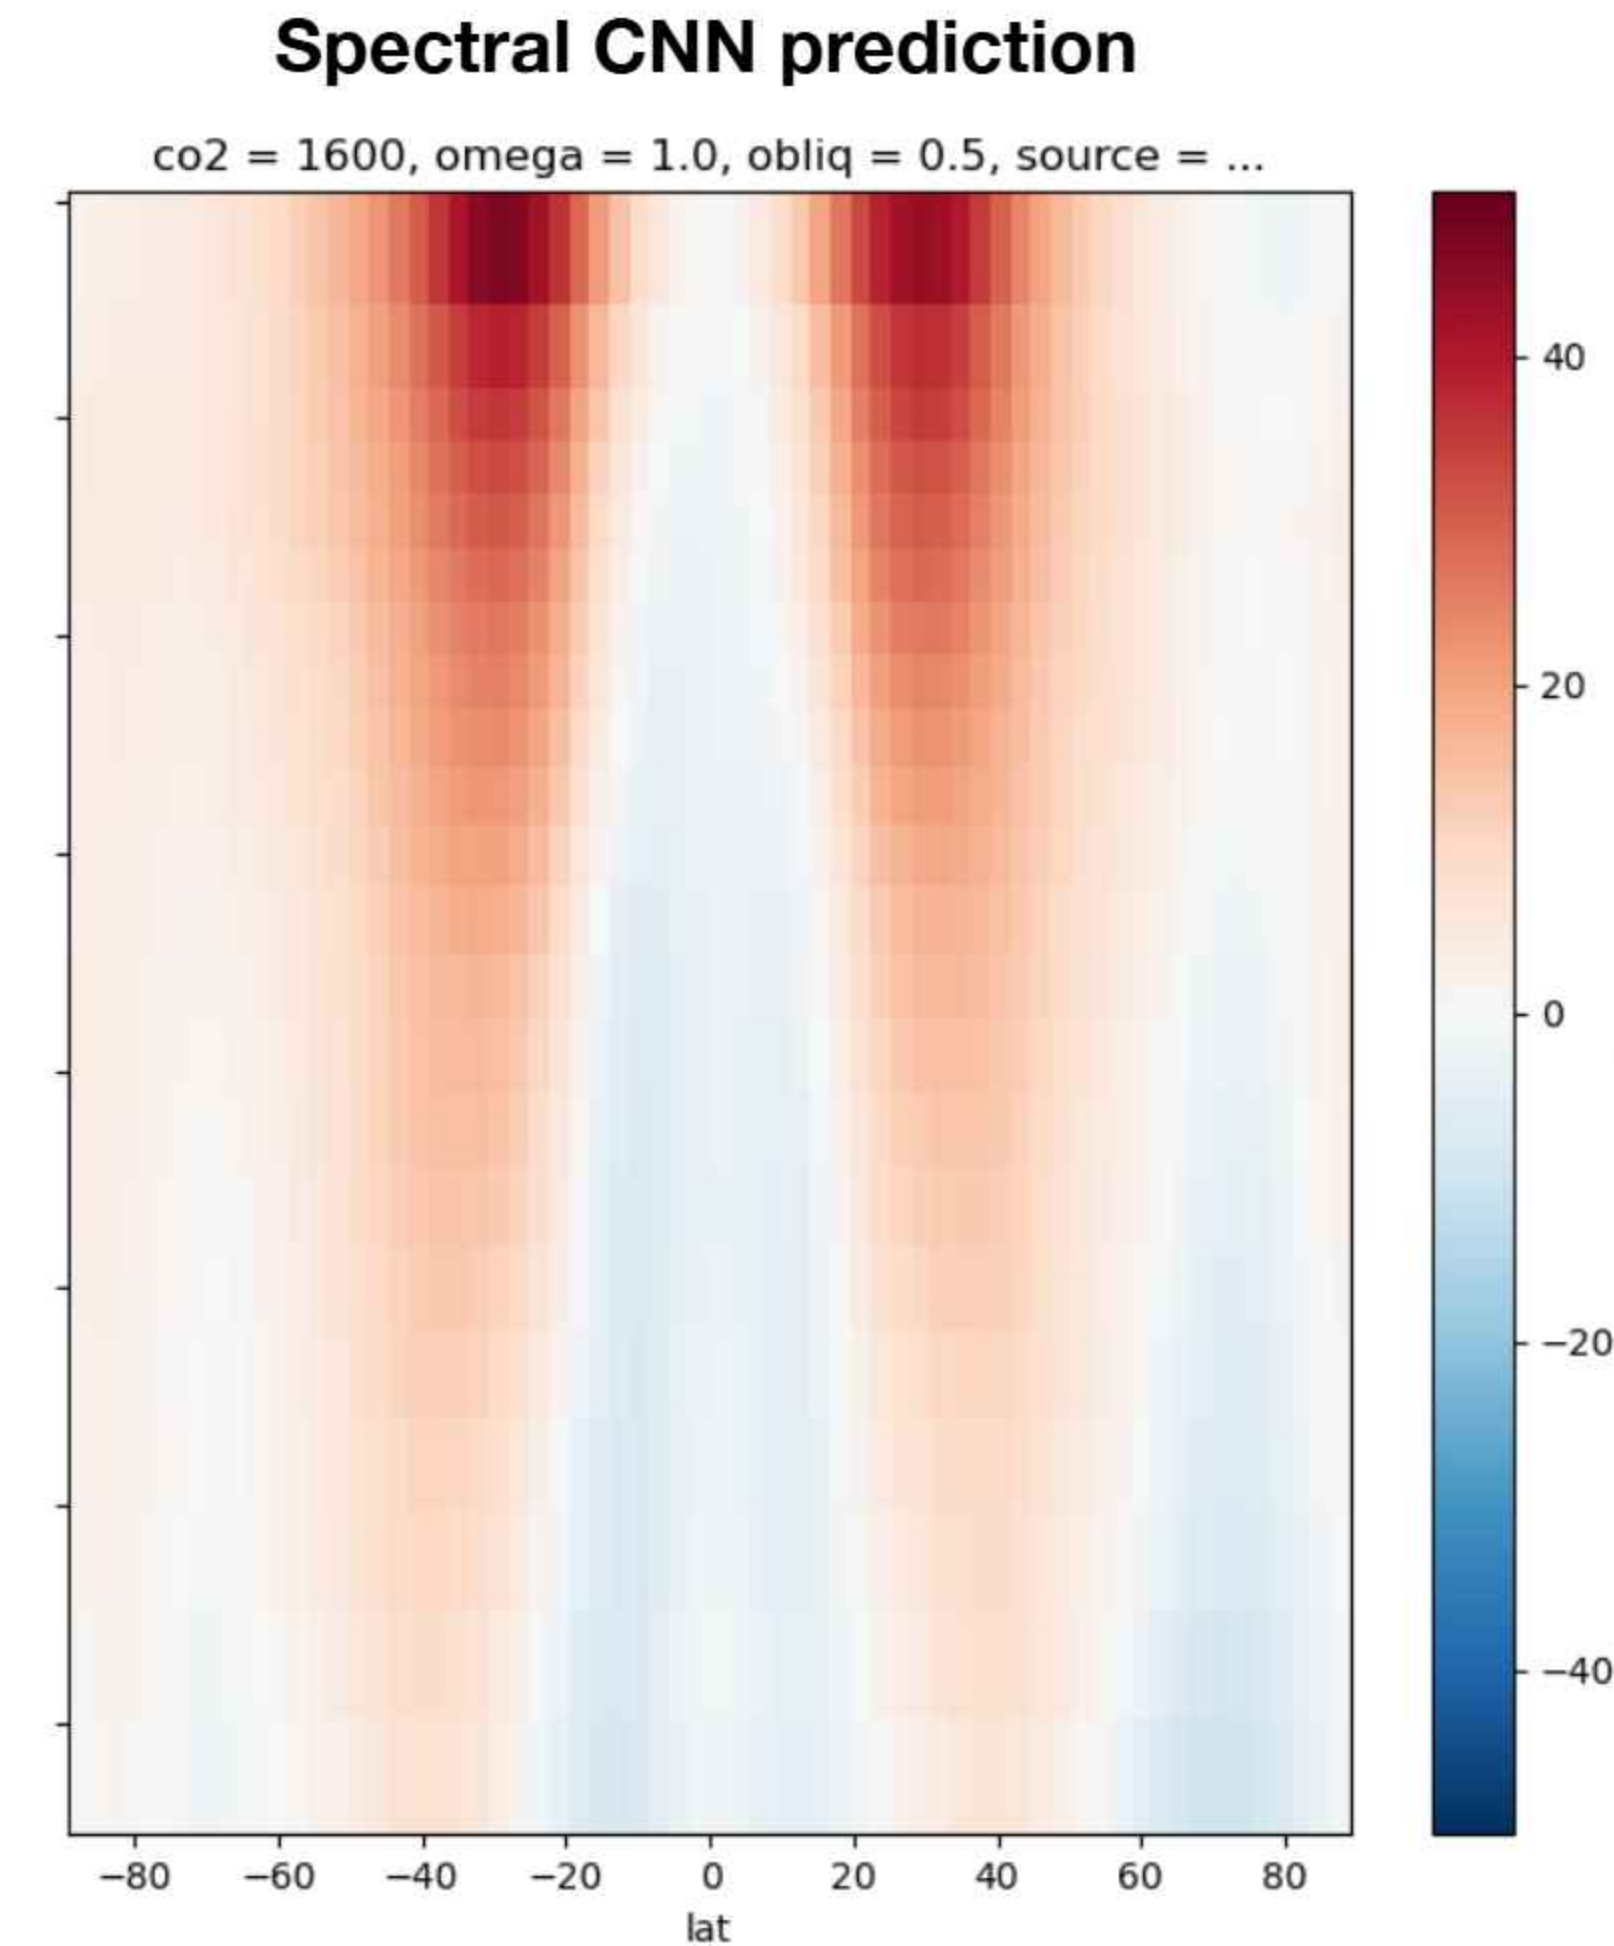
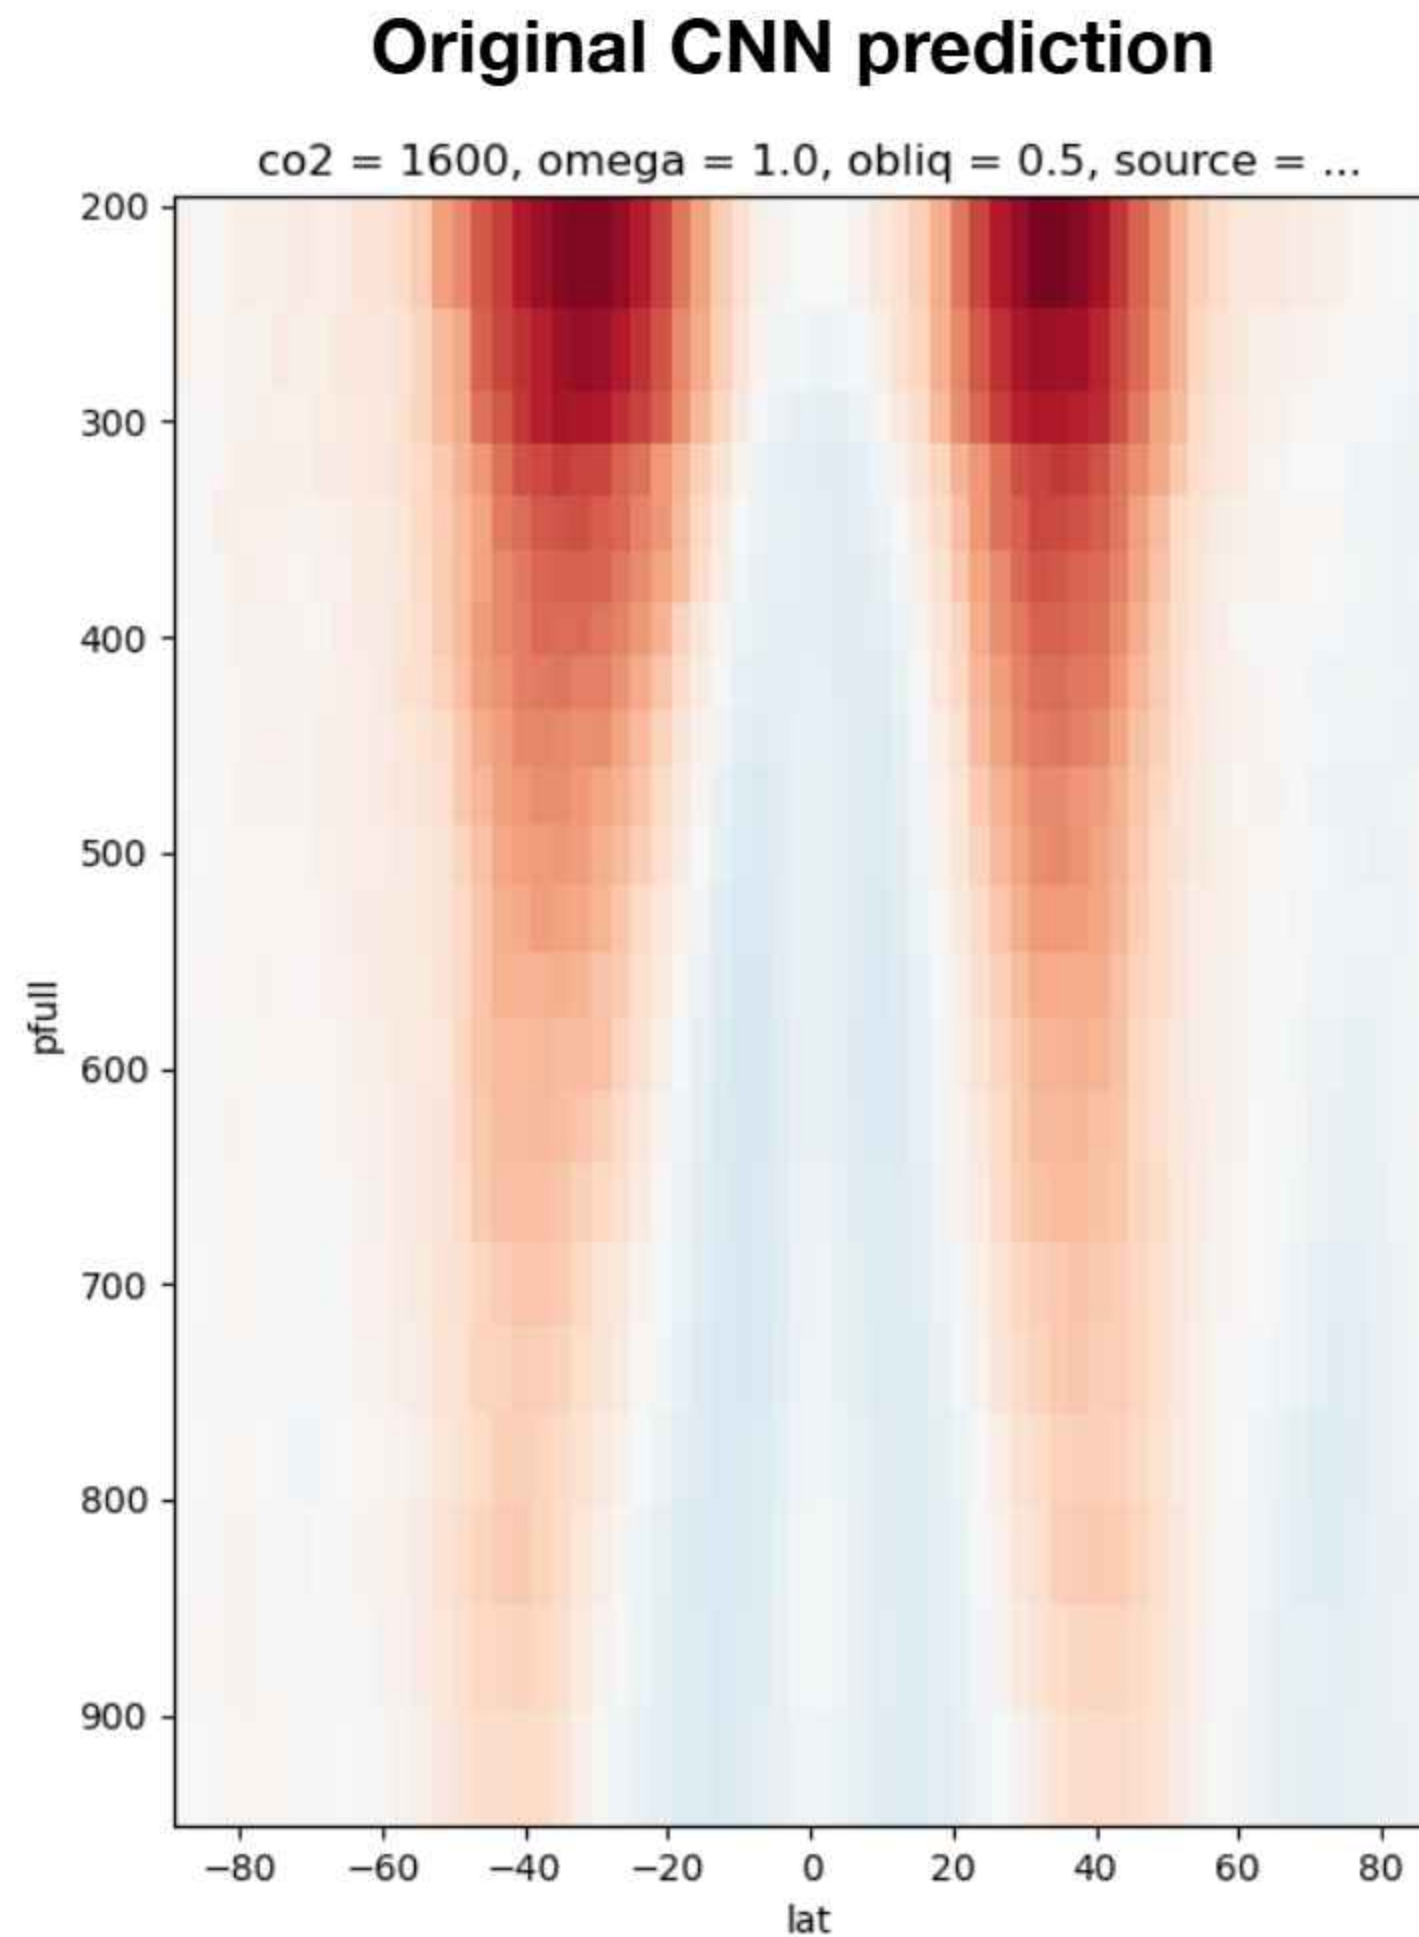
Removing small-scale noise

- To remove small-scale noise, we tried:
 - adding a term penalising pixel-to-pixel variation in the output
 - **Making predictions of the fields in spectral space, focussing on large-scale structure**



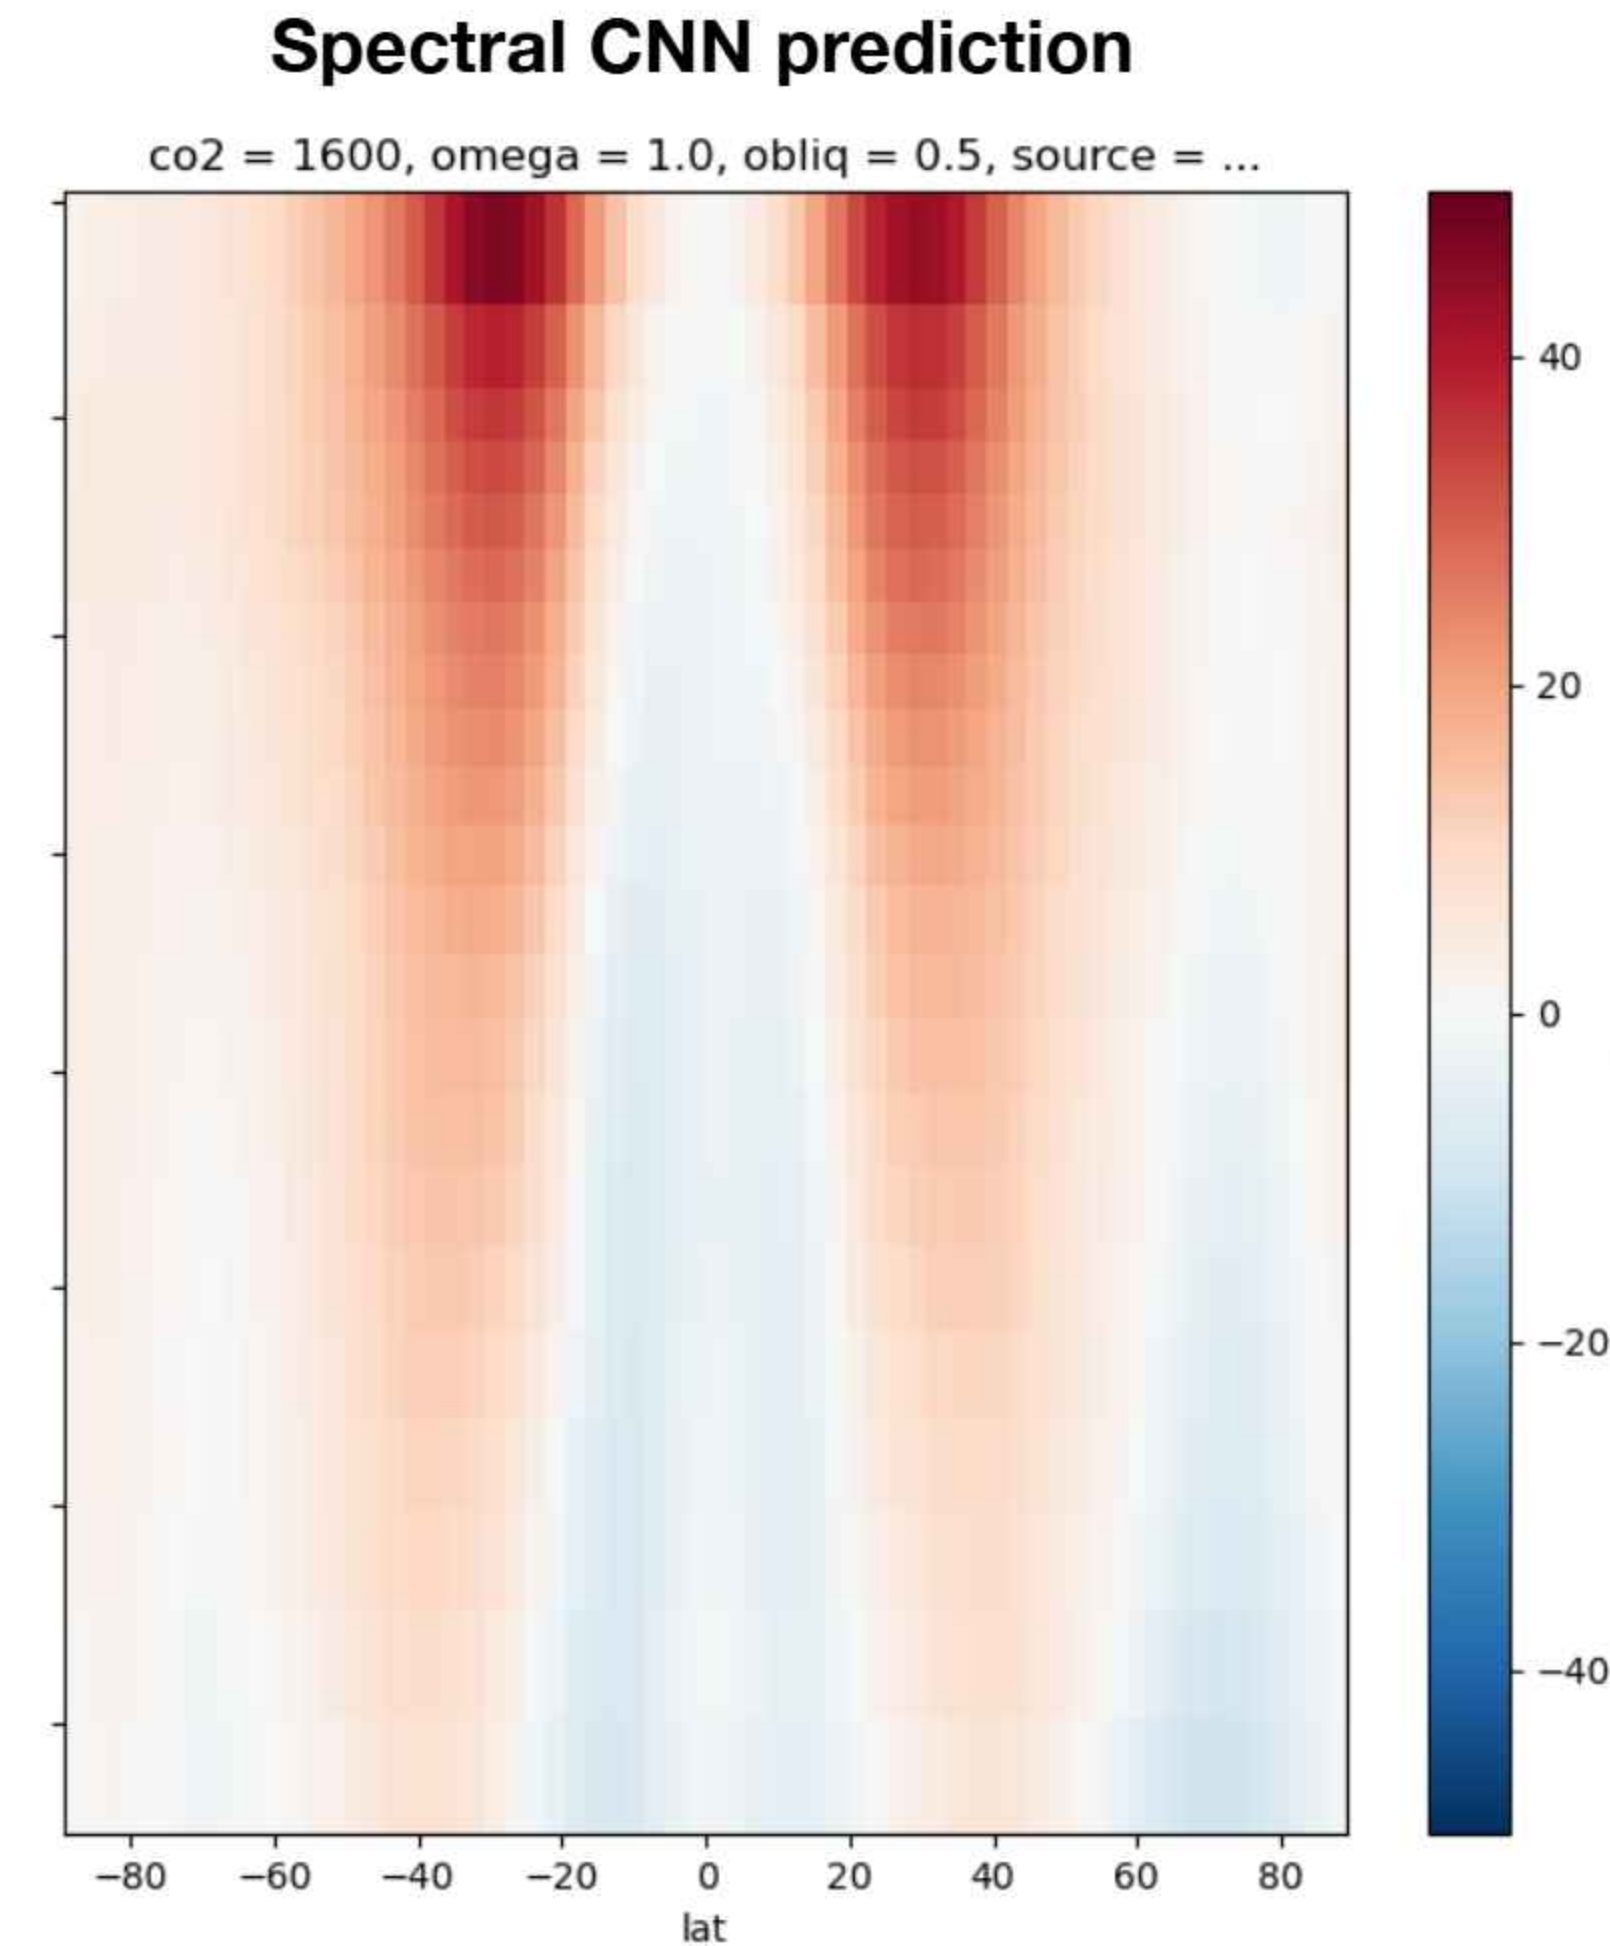
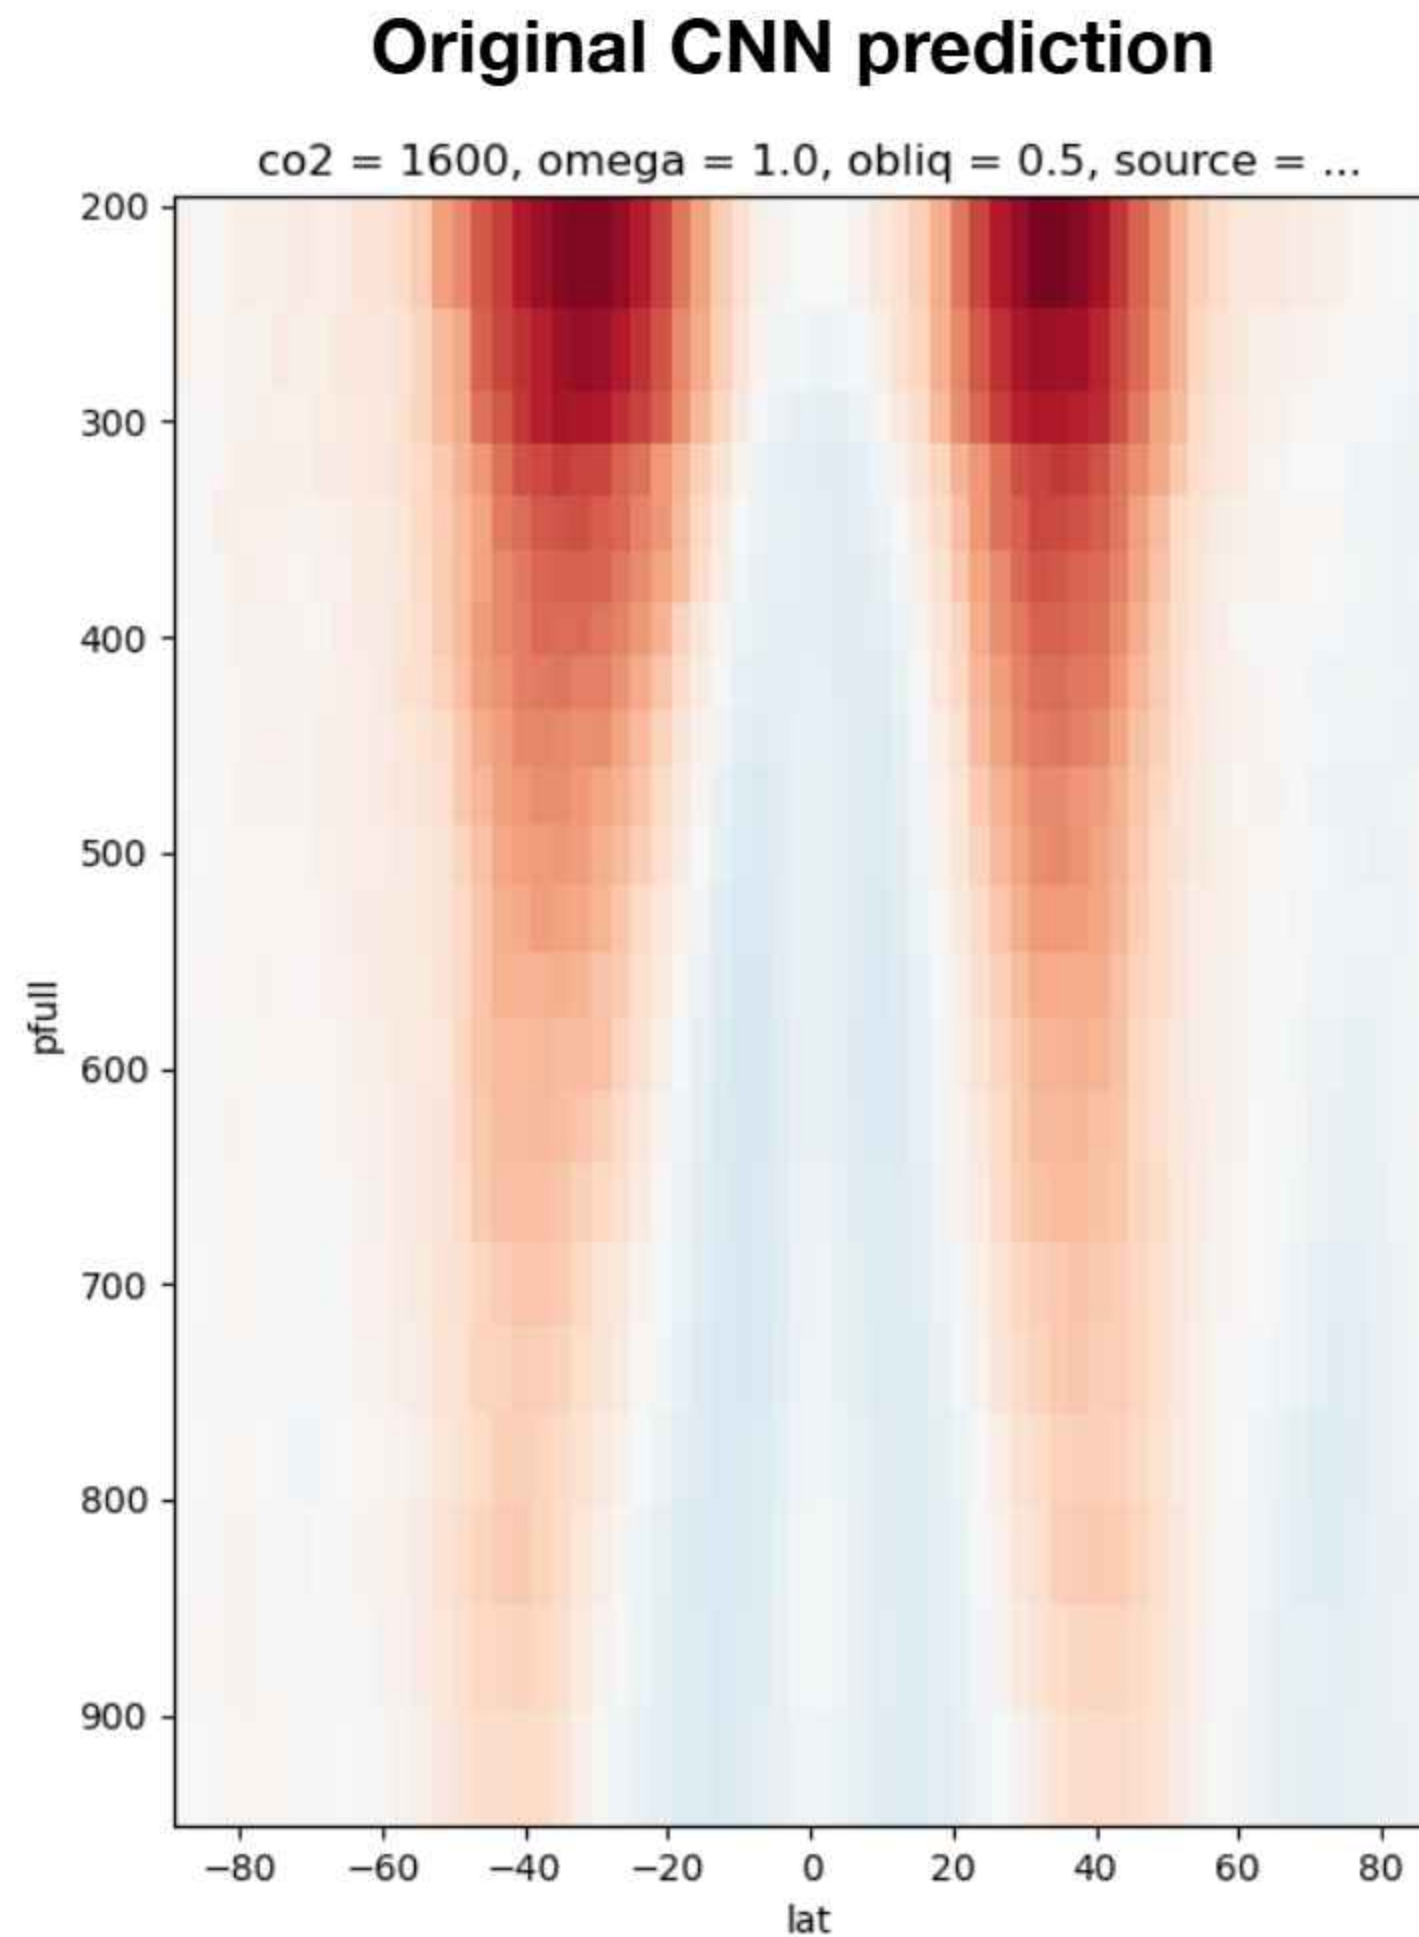
Removing small-scale noise

- To remove small-scale noise, we tried:
 - adding a term penalising pixel-to-pixel variation in the output
 - **Making predictions of the fields in spectral space, focussing on large-scale structure**
- **Spectral predictions** have substantially improved small scale noise...



Removing small-scale noise

- To remove small-scale noise, we tried:
 - adding a term penalising pixel-to-pixel variation in the output
 - **Making predictions of the fields in spectral space, focussing on large-scale structure**
- **Spectral predictions** have substantially improved small scale noise...
- Next step is to incorporate the **dynamical constraint** into this spectral network



Conclusions and future work

Conclusions and future work

- We have constructed a CNN to predict the zonal-mean temperature and zonal-mean zonal wind from GCM experiments over a parameter space of CO₂, obliquities and rotation rates

Conclusions and future work

- We have constructed a CNN to predict the zonal-mean temperature and zonal-mean zonal wind from GCM experiments over a parameter space of CO₂, obliquities and rotation rates
- We have demonstrated that incorporating constraints from radiative-convective equilibrium have substantially improved the CNN's ability to extrapolate to high CO₂ climates

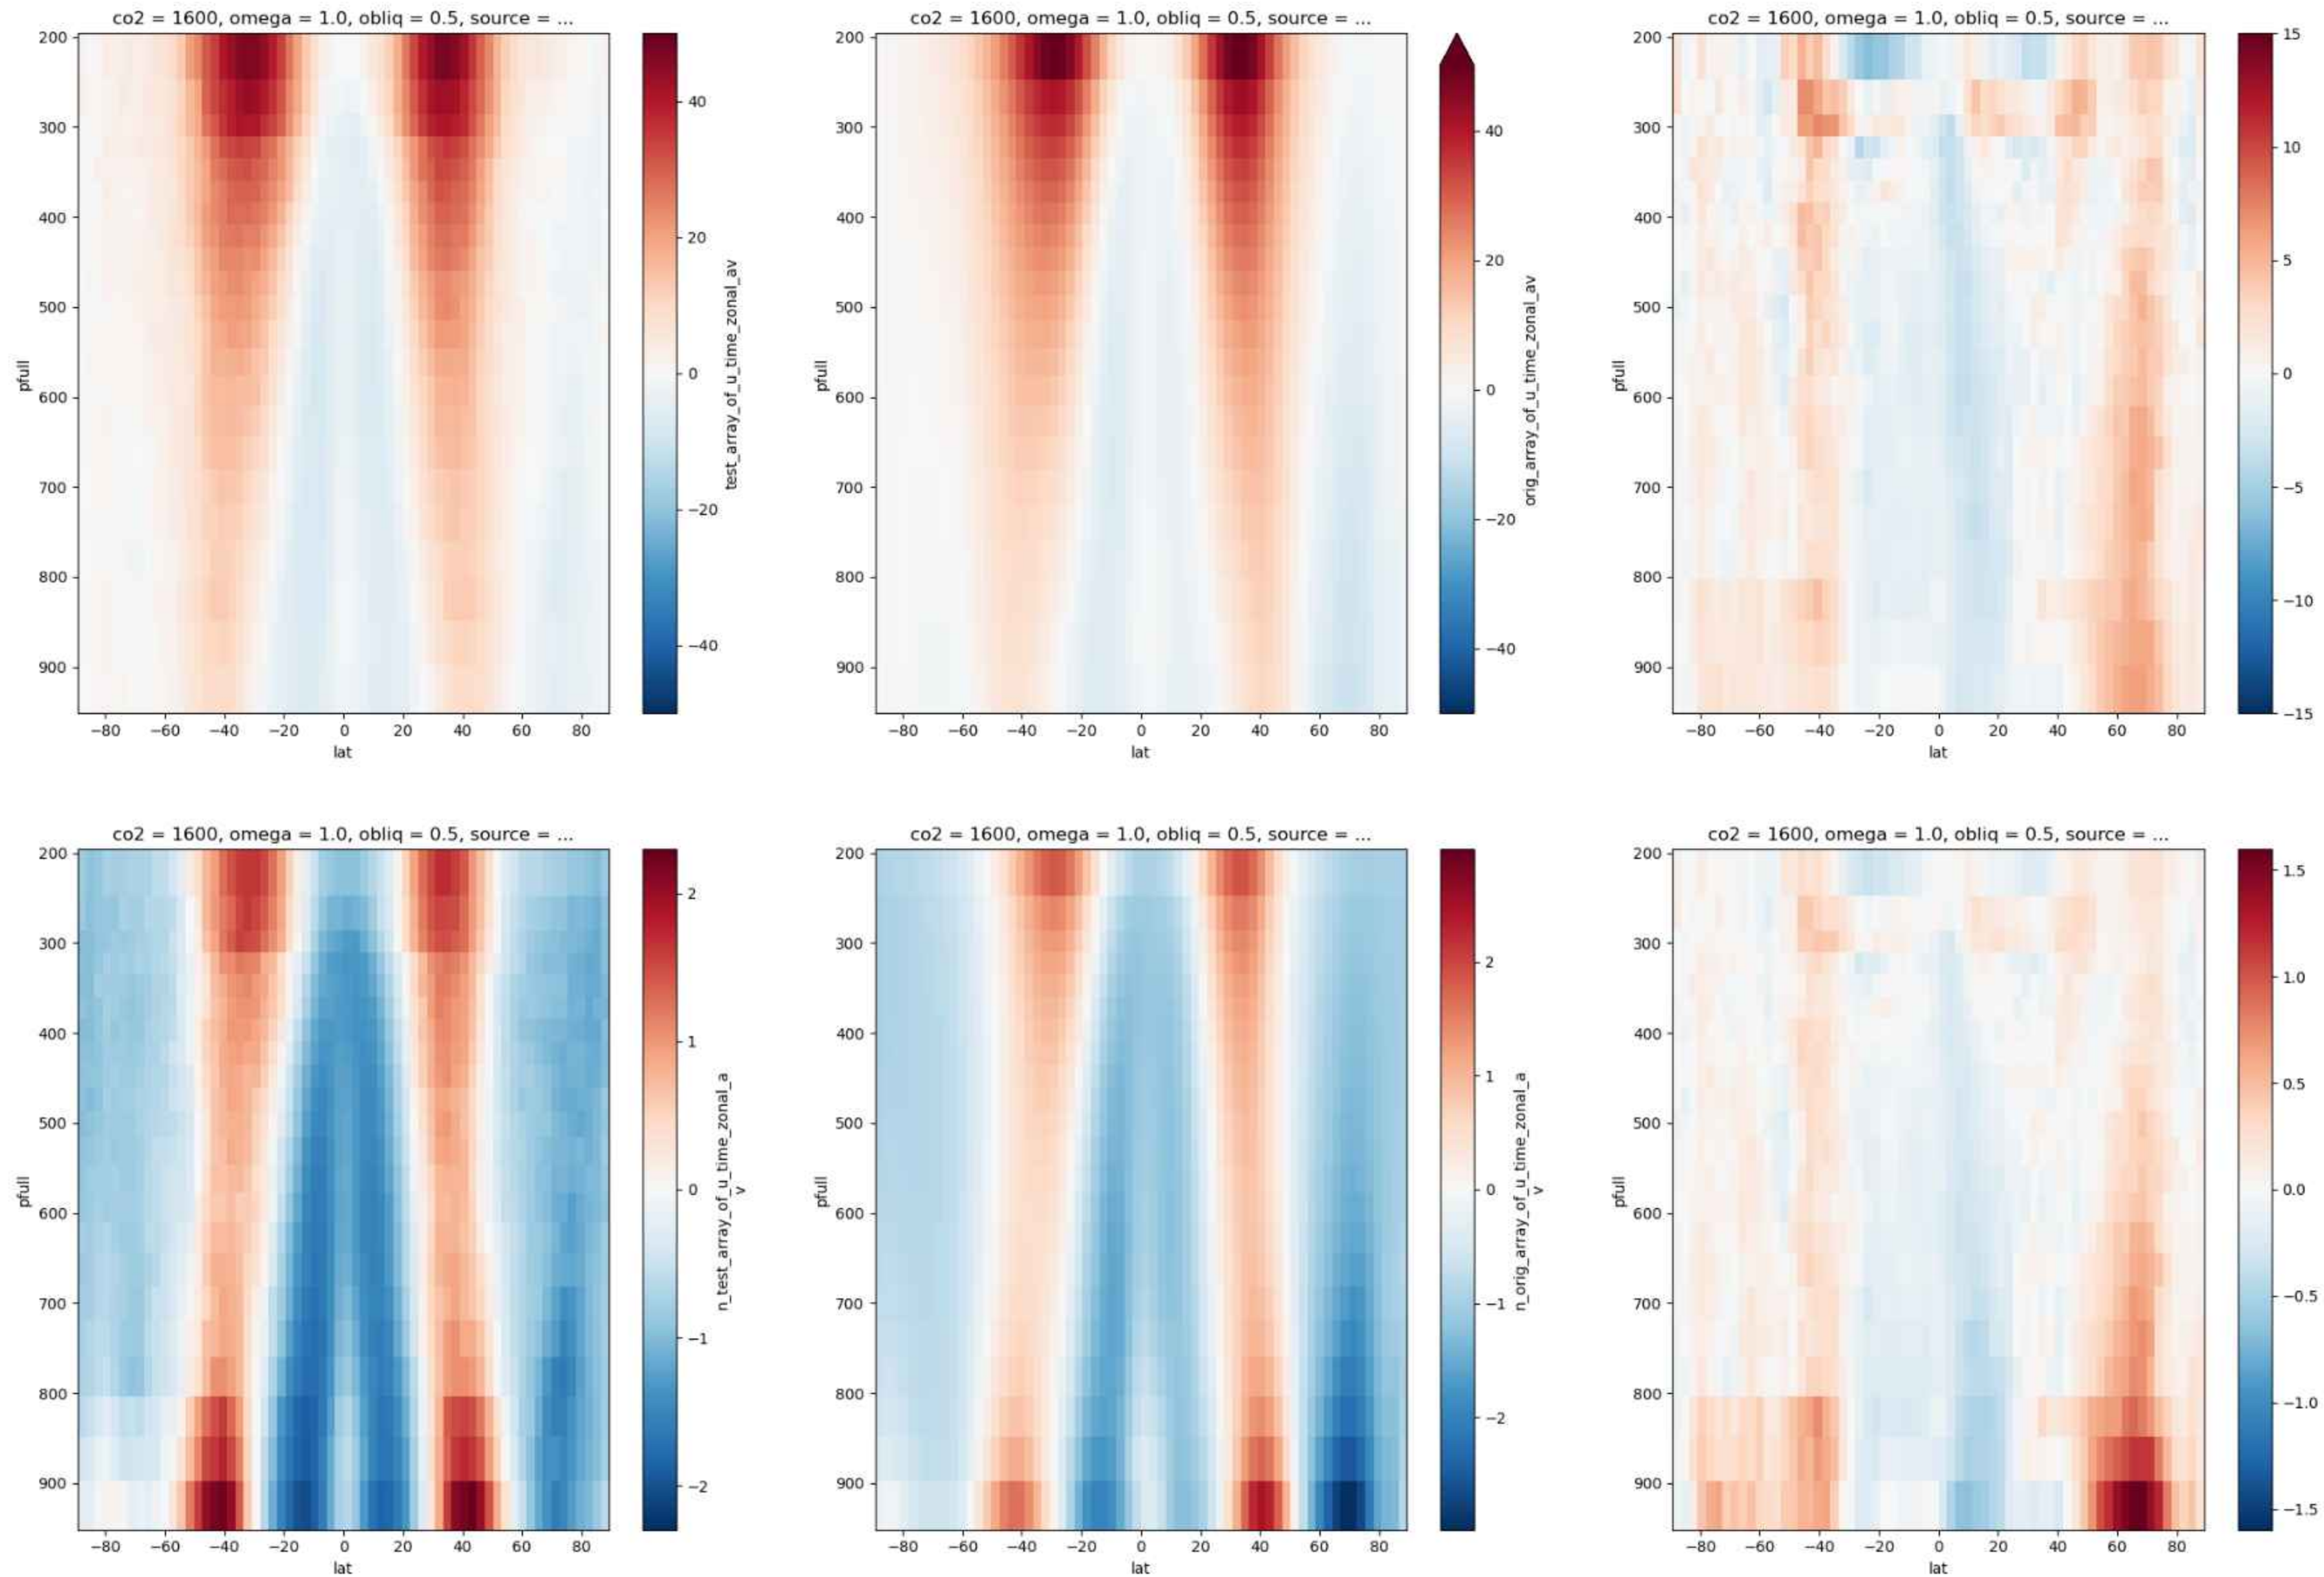
Conclusions and future work

- We have constructed a CNN to predict the zonal-mean temperature and zonal-mean zonal wind from GCM experiments over a parameter space of CO₂, obliquities and rotation rates
- We have demonstrated that incorporating constraints from radiative-convective equilibrium have substantially improved the CNN's ability to extrapolate to high CO₂ climates
- Incorporating results from cheap physical models (e.g. radiative convective equilibrium) may be a way to improve ML predictions of future climates

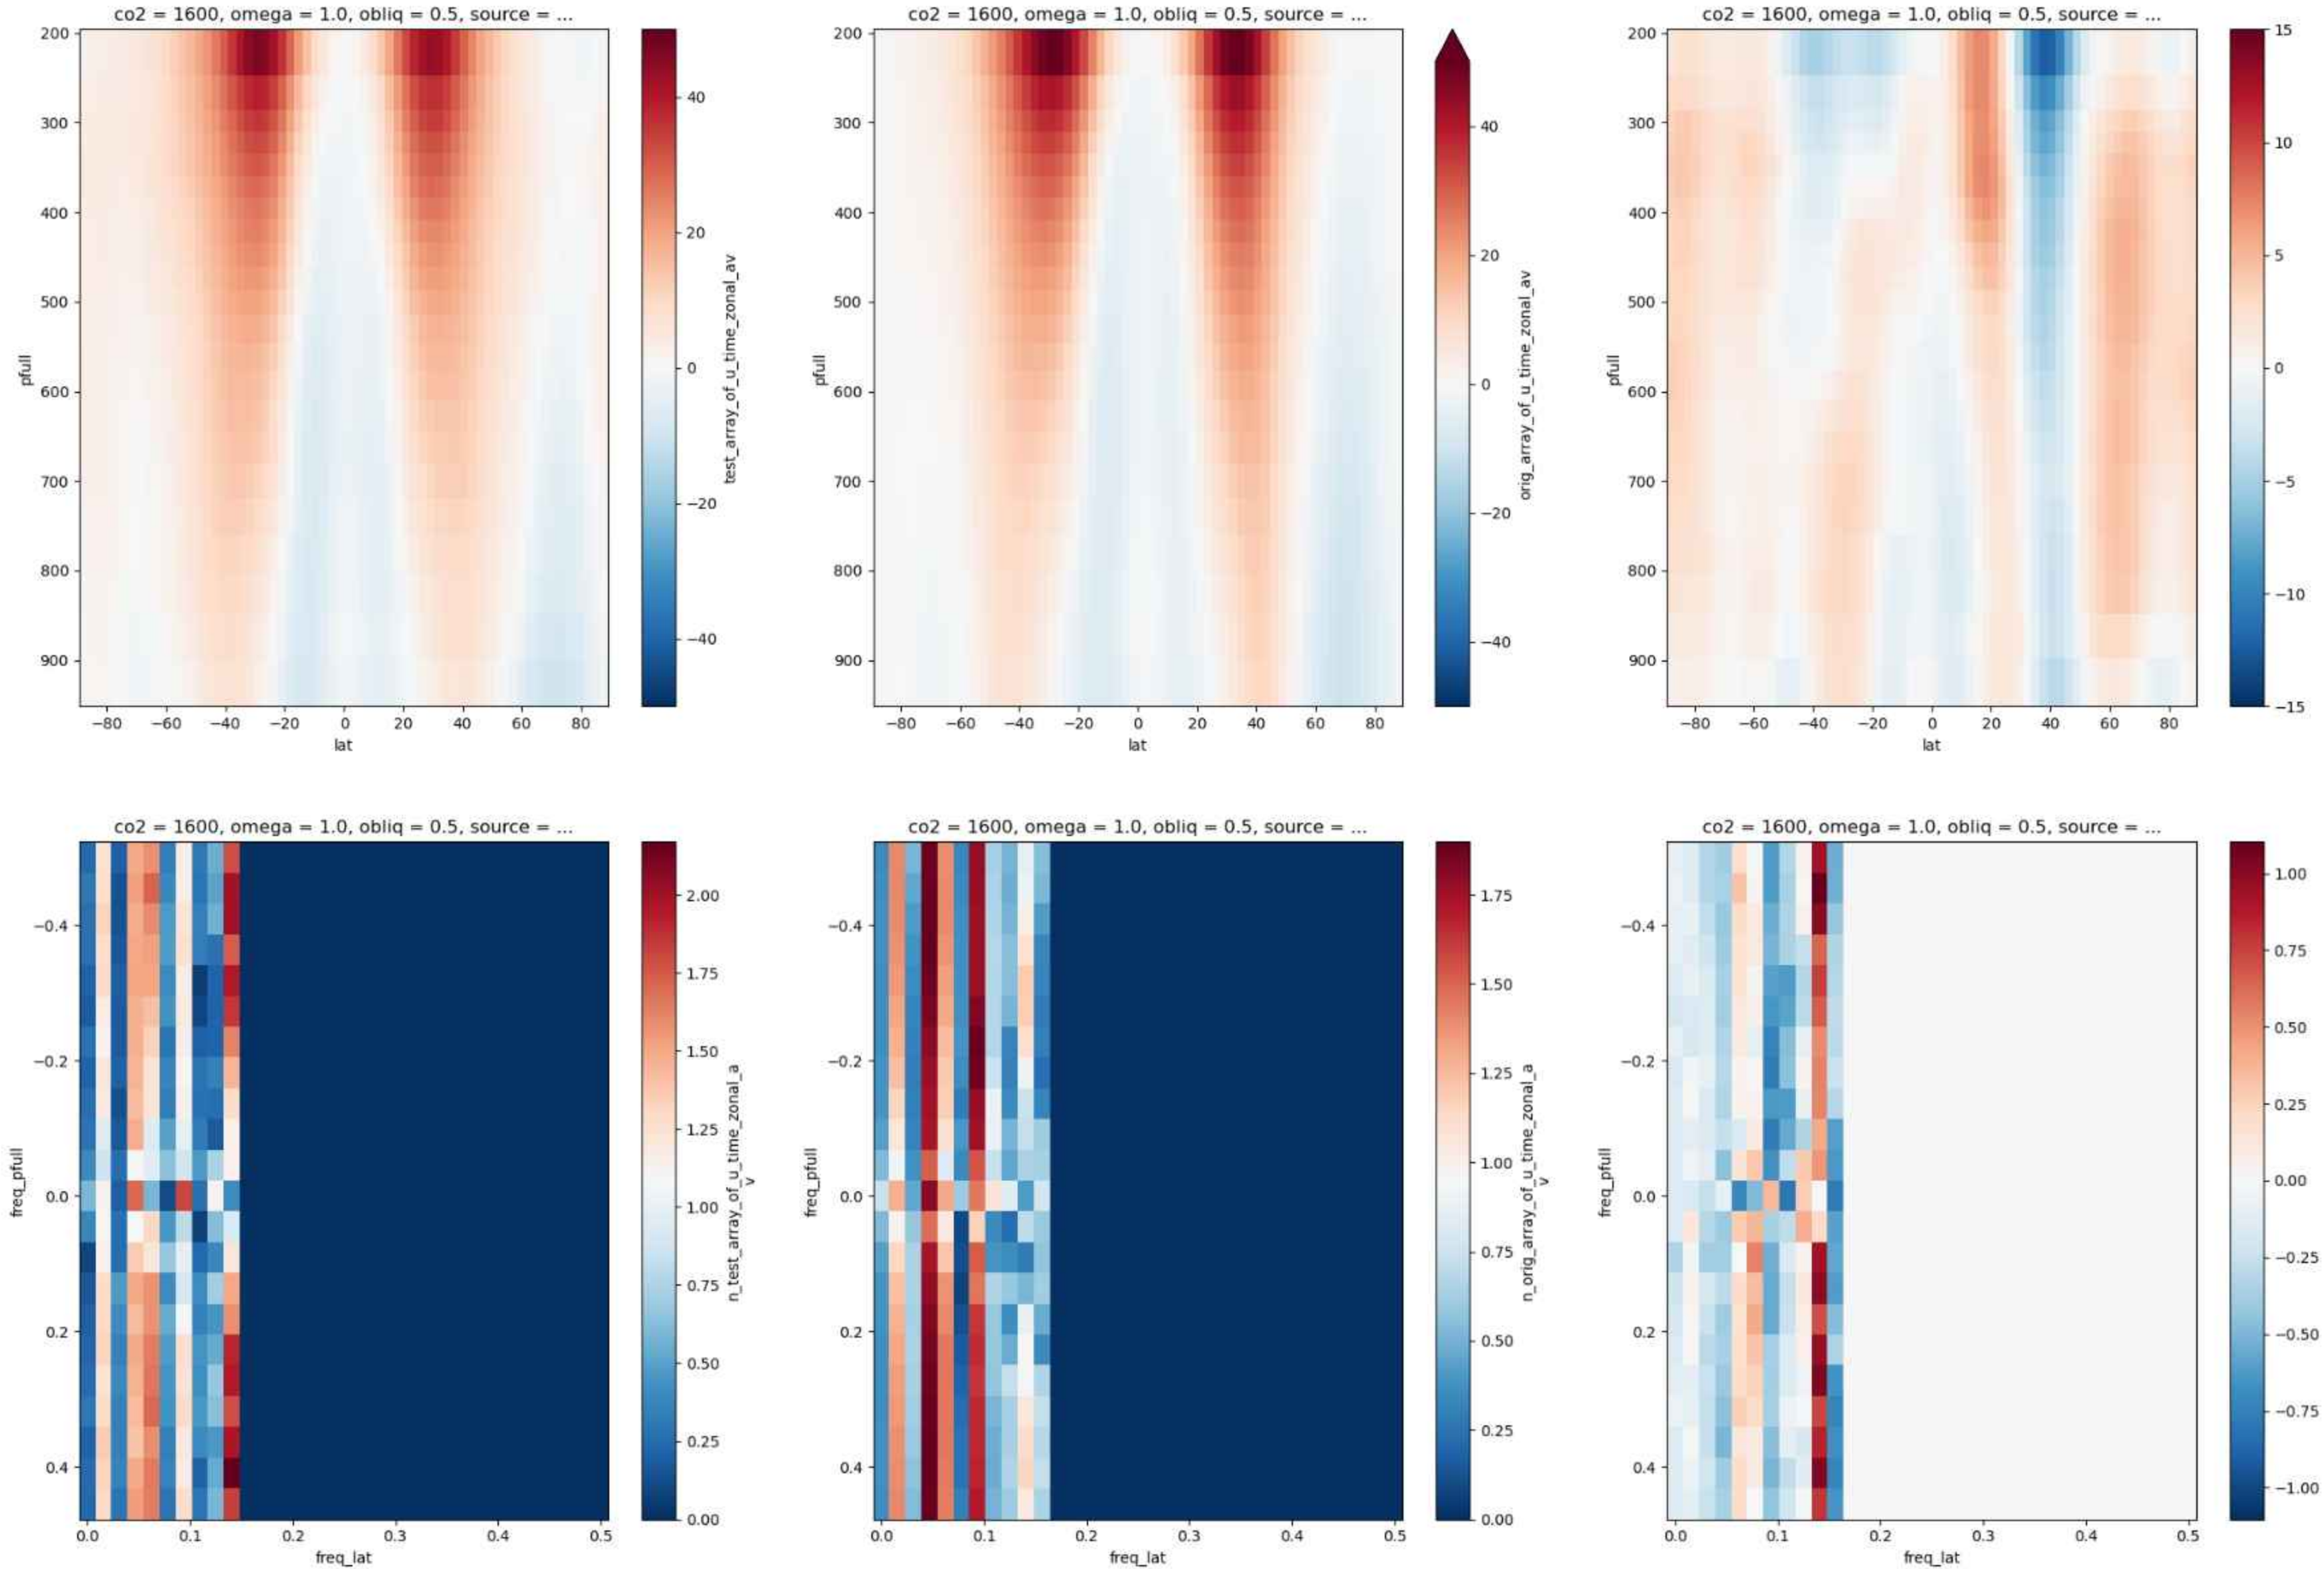
Conclusions and future work

- We have constructed a CNN to predict the zonal-mean temperature and zonal-mean zonal wind from GCM experiments over a parameter space of CO₂, obliquities and rotation rates
- We have demonstrated that incorporating constraints from radiative-convective equilibrium have substantially improved the CNN's ability to extrapolate to high CO₂ climates
- Incorporating results from cheap physical models (e.g. radiative convective equilibrium) may be a way to improve ML predictions of future climates
- **Future work:**
 - Making predictions in spectral space and applying gradient wind balance constraint
 - Incorporating additional variables, e.g. moisture, precipitation etc
 - Explore exoplanet applications, including variations in composition etc

Normalisation in real space...



Normalisation in spectral space...

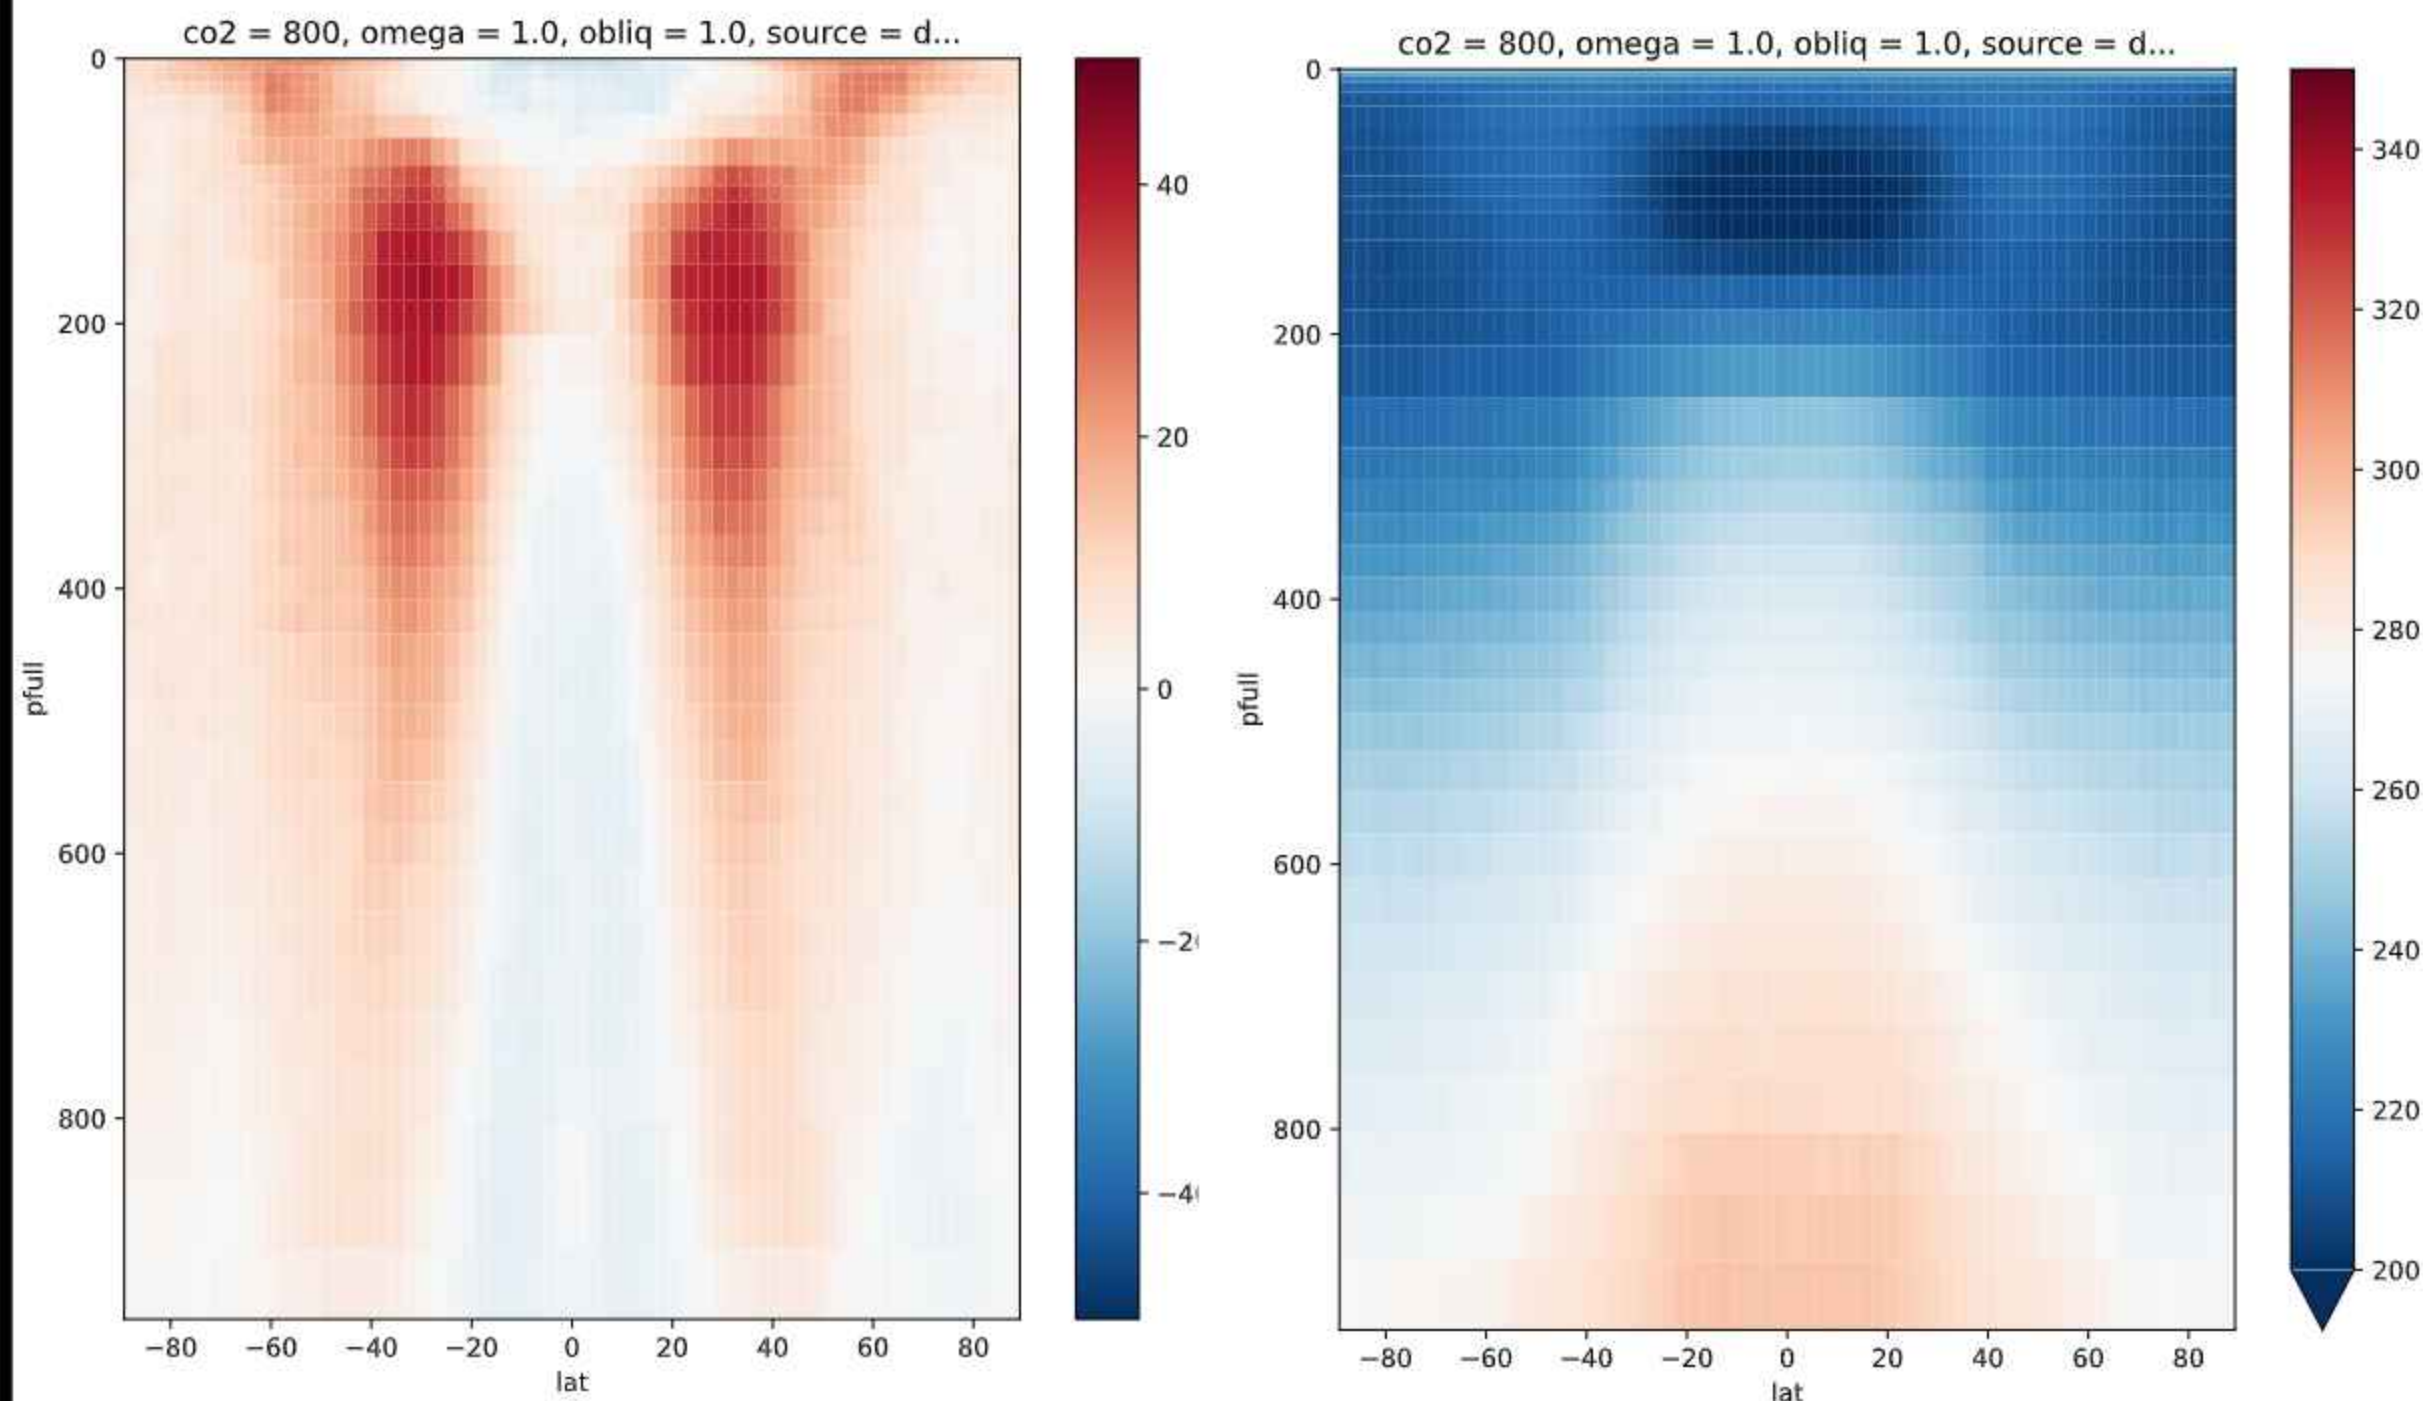


Exploring the data

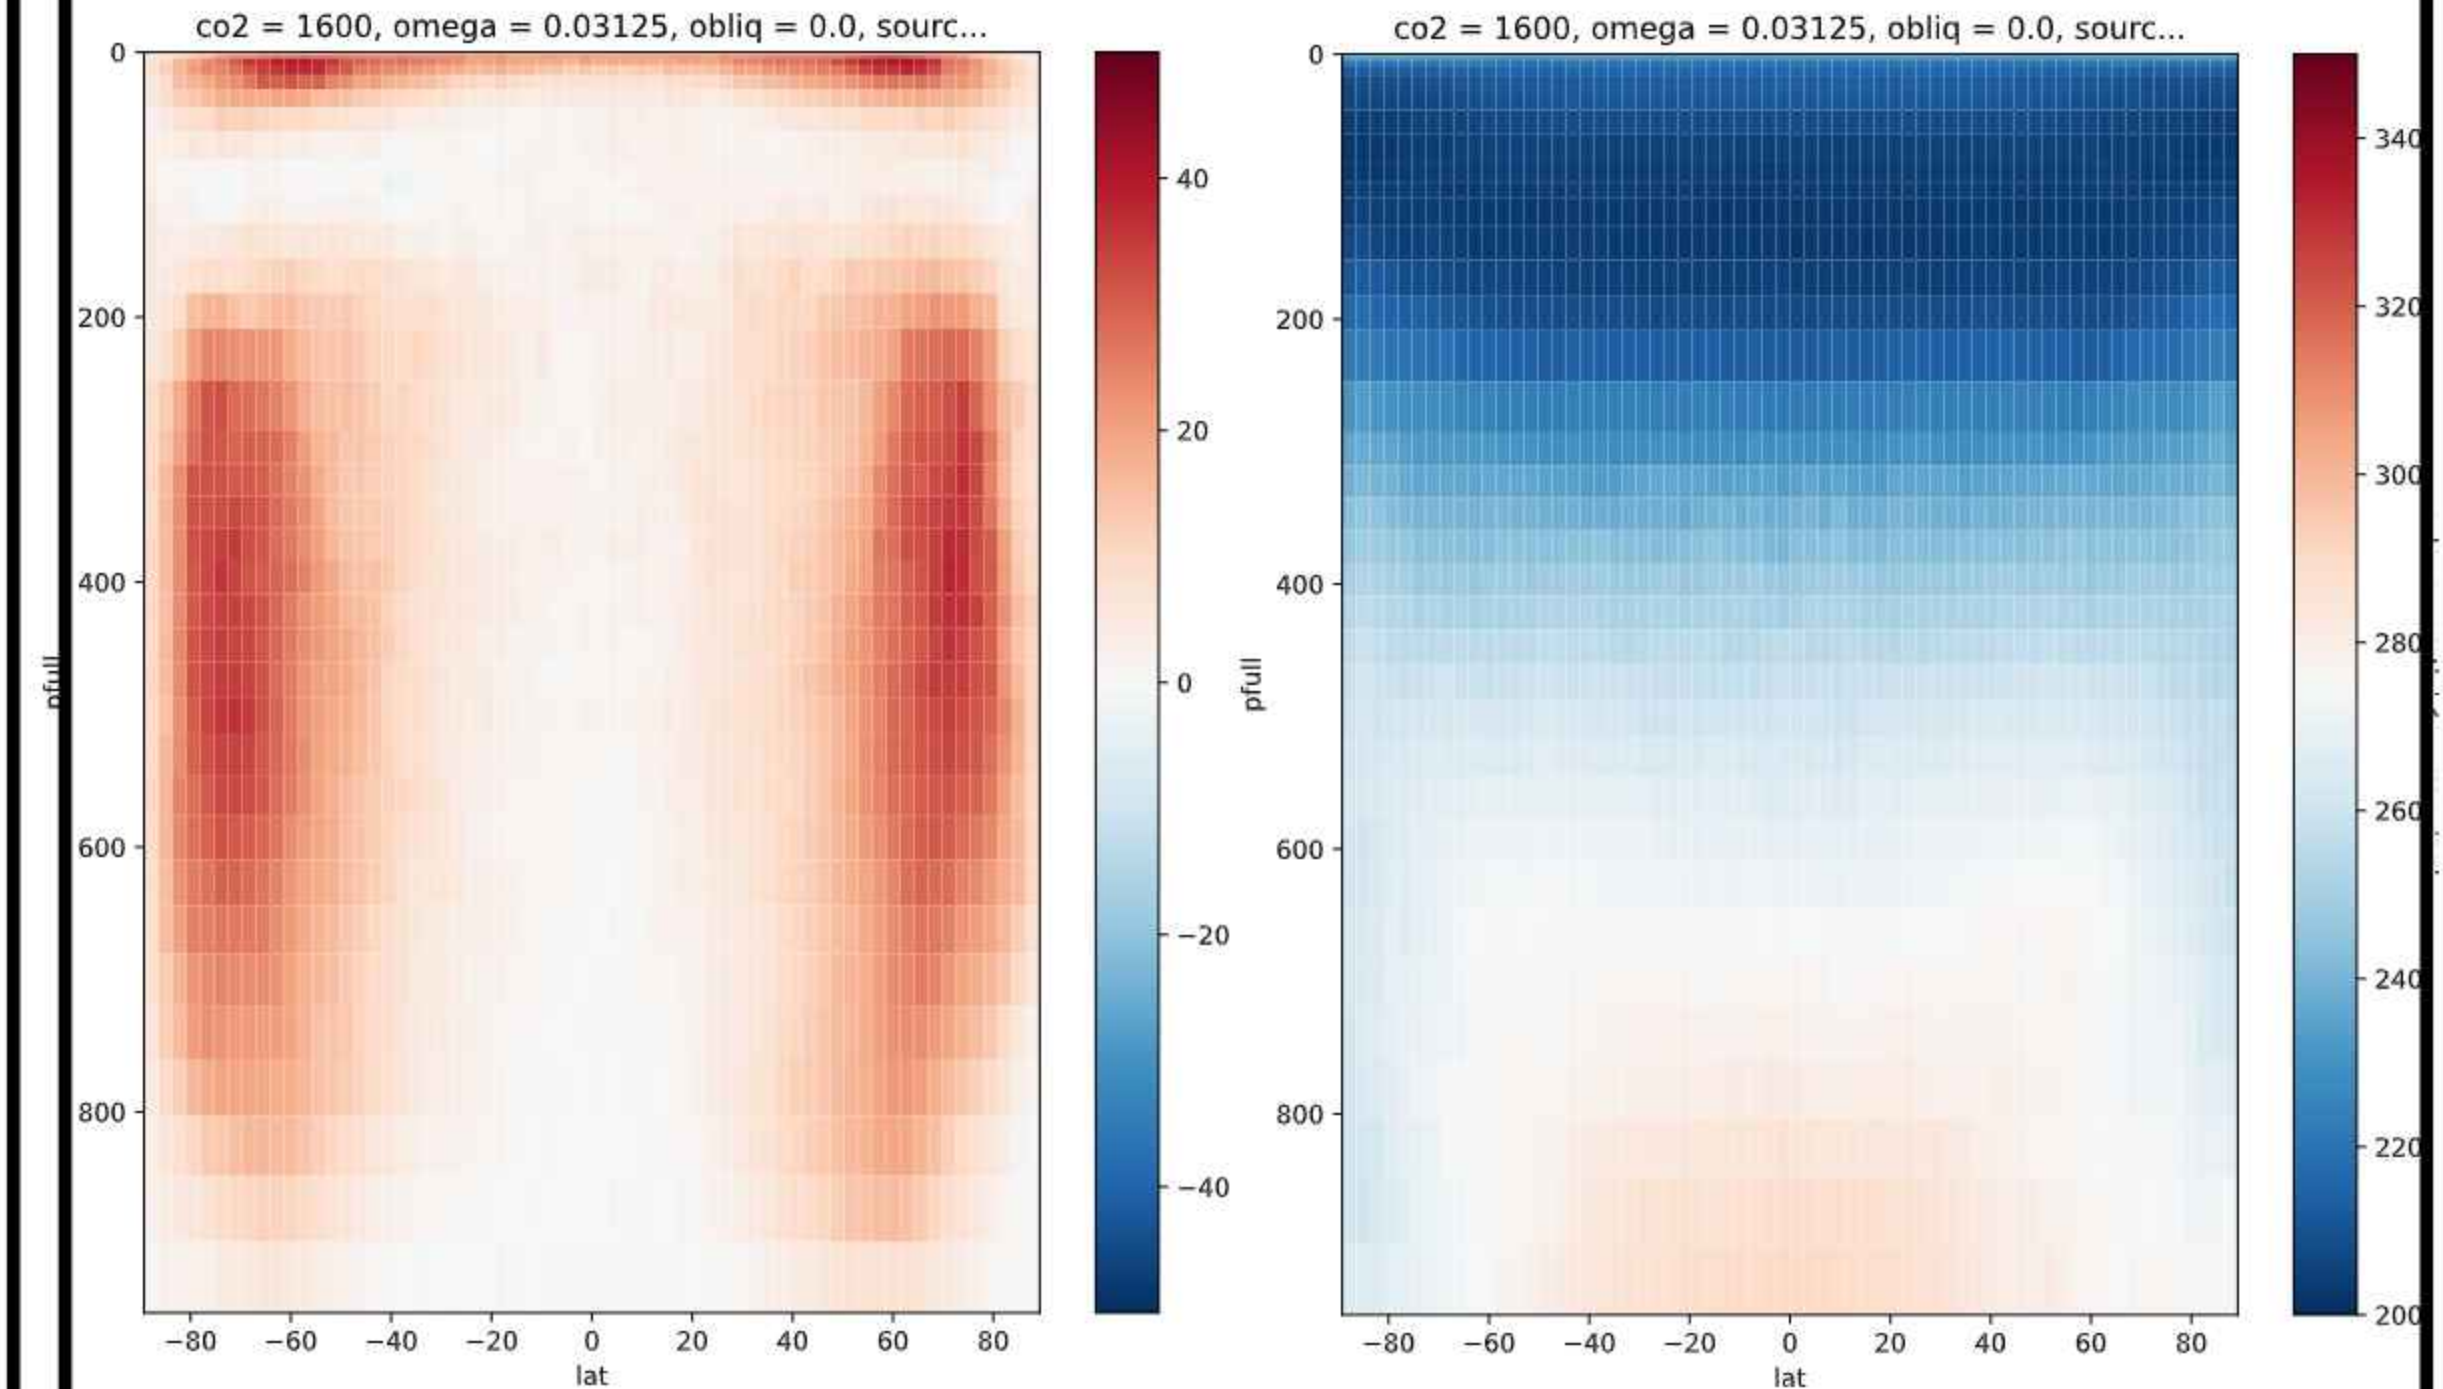
$$\Omega_0 = \frac{2\pi}{86400} s^{-1}$$



CO₂ = 800ppmv, obliquity = 1, omega = 1



CO₂ = 1600ppmv, obliquity = 0, omega = 0.03125

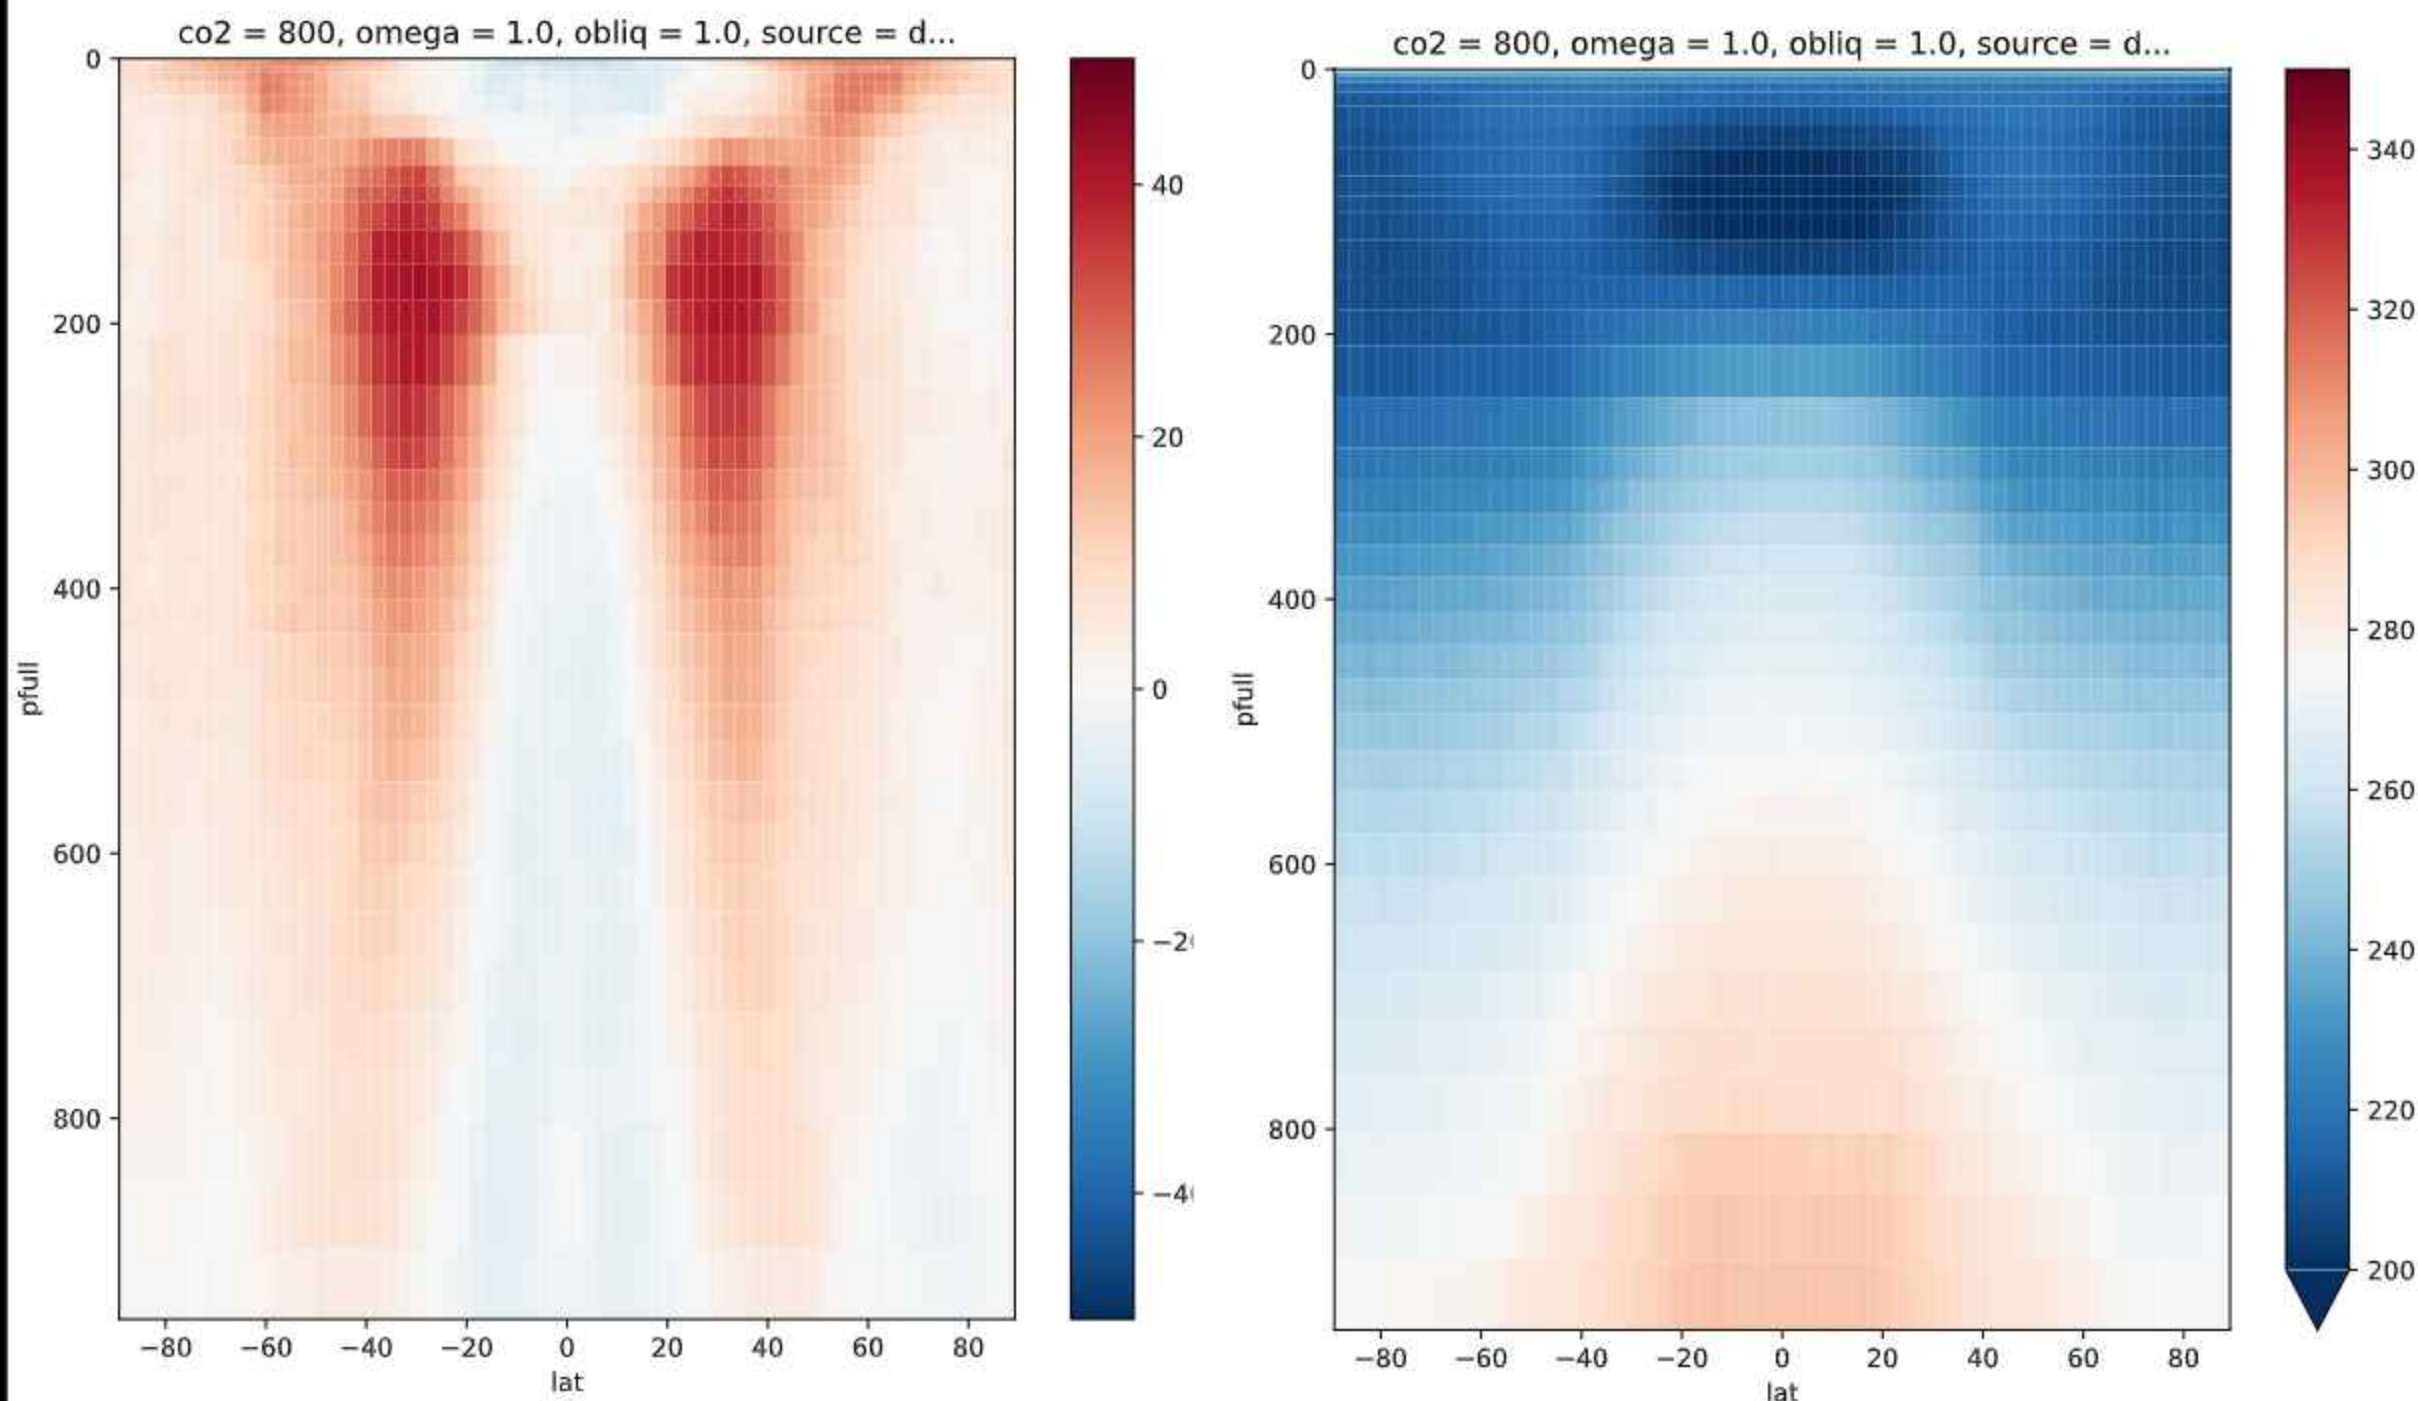


Exploring the data

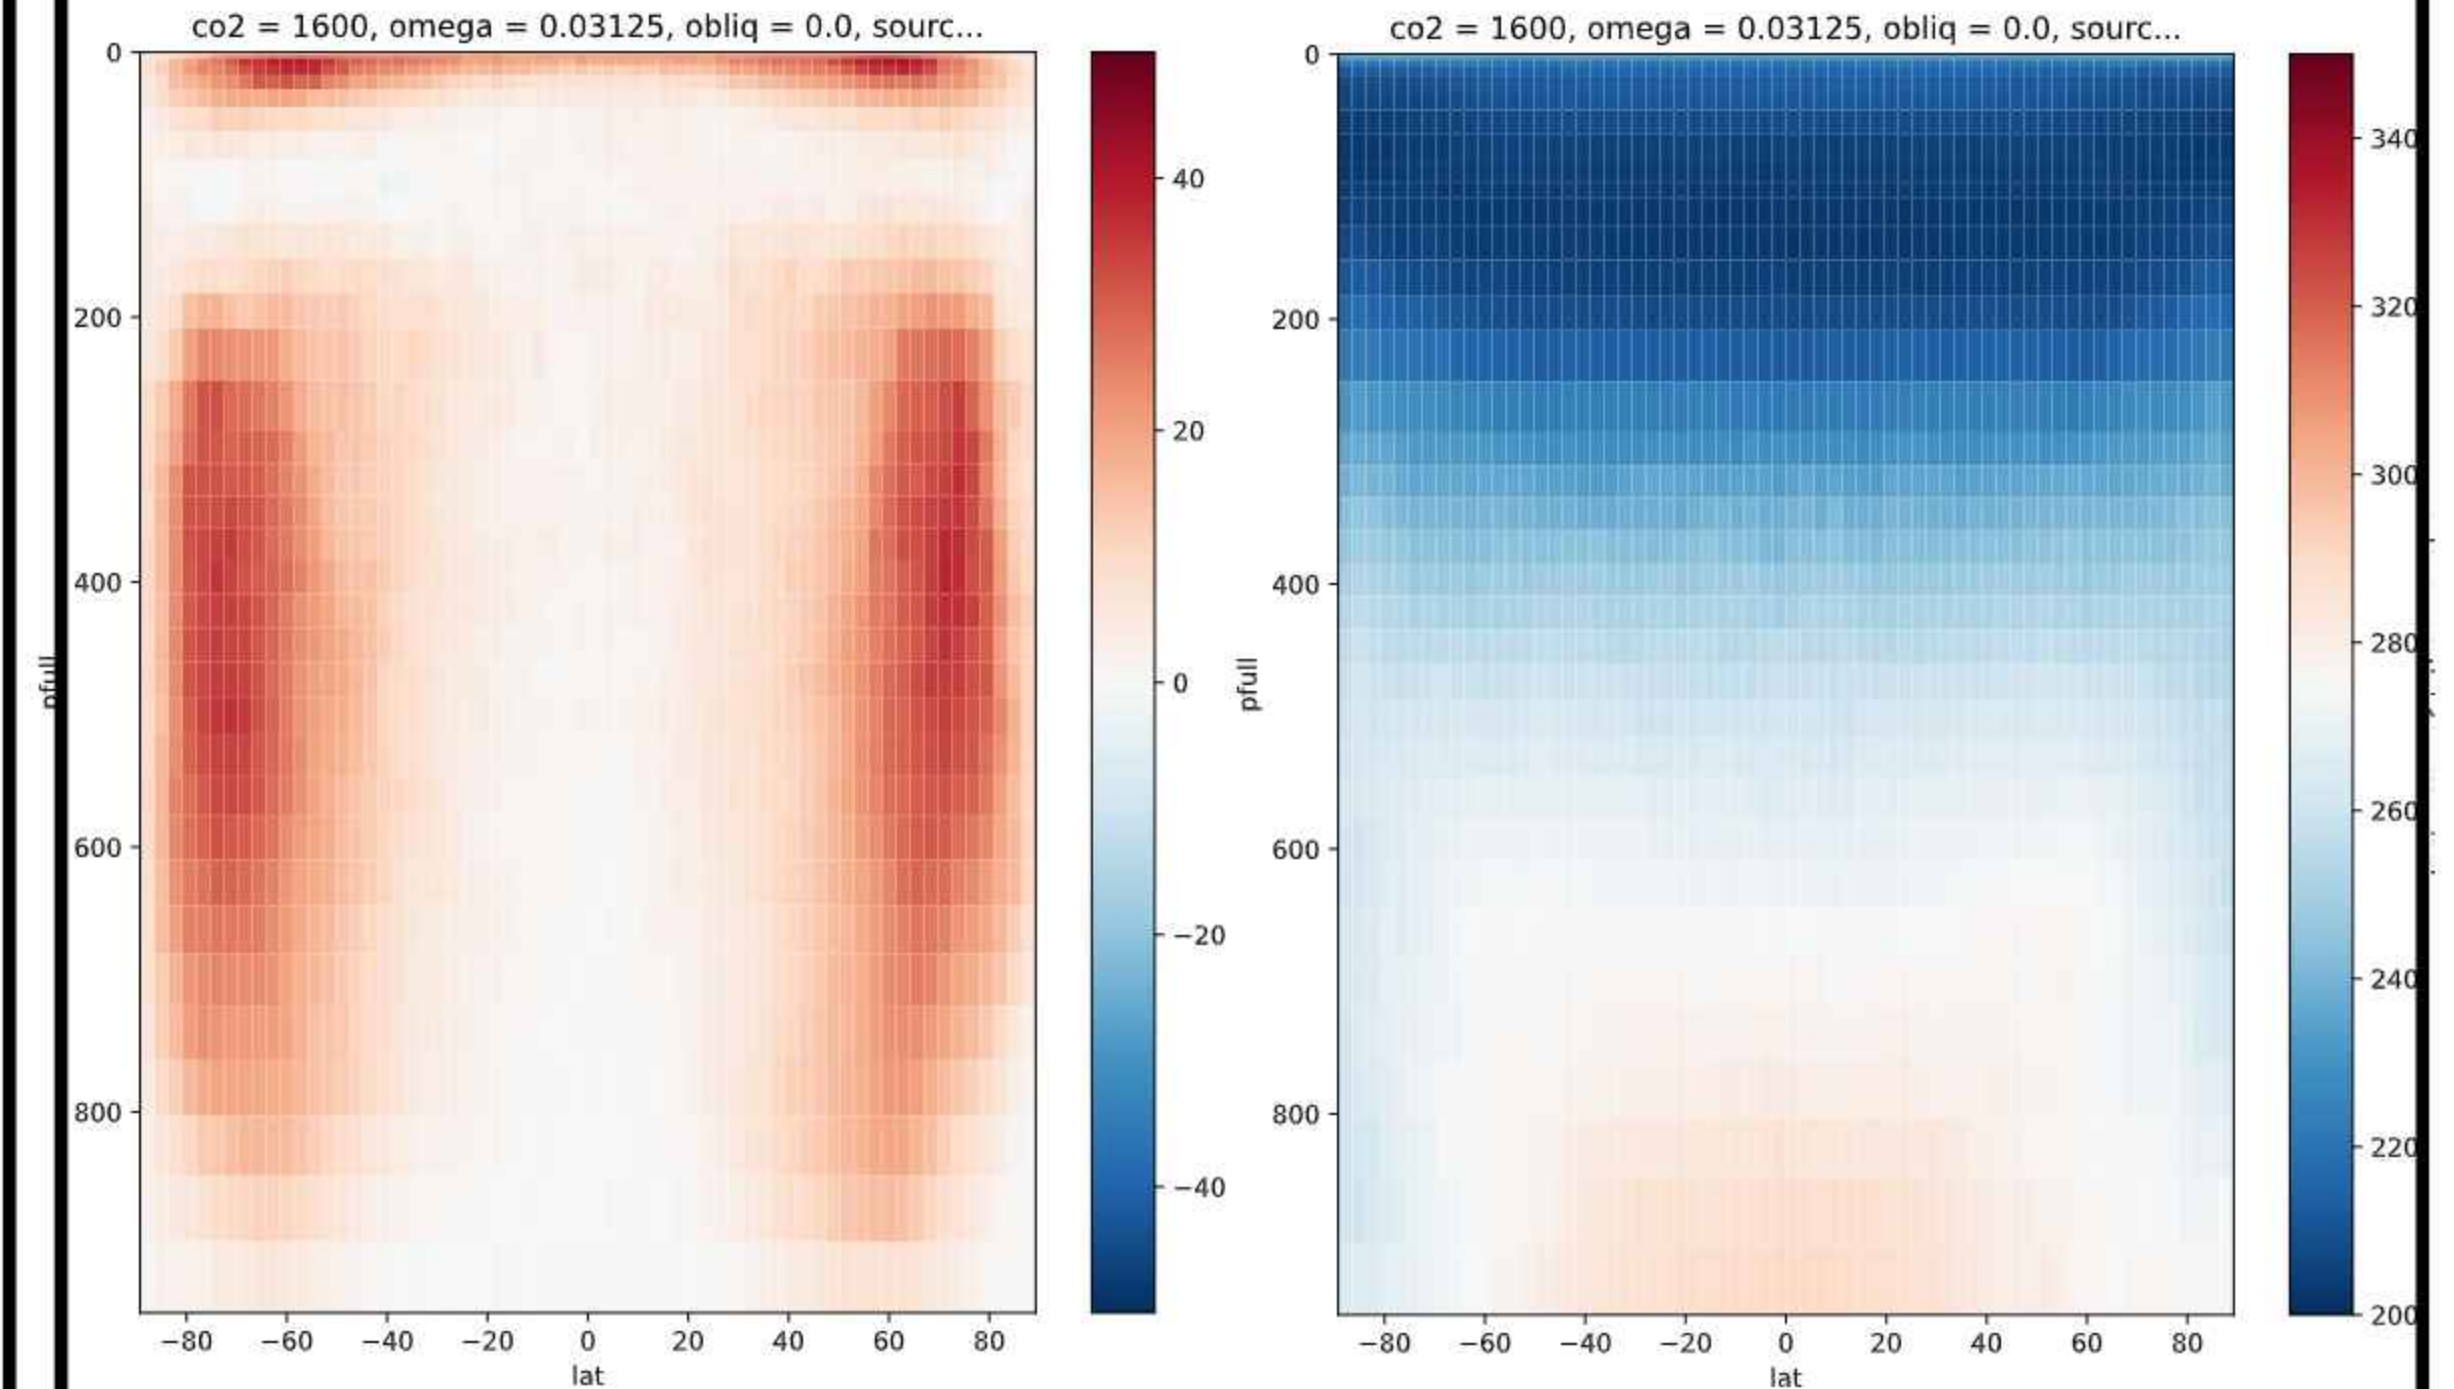
- 7 rotation rates (2, 1, 0.5, 0.25, 0.125, 0.0625, 0.03125) in units of Earth's $\Omega_0 = \frac{2\pi}{86400} s^{-1}$



CO₂ = 800ppmv, obliquity = 1, omega = 1



CO₂ = 1600ppmv, obliquity = 0, omega = 0.03125

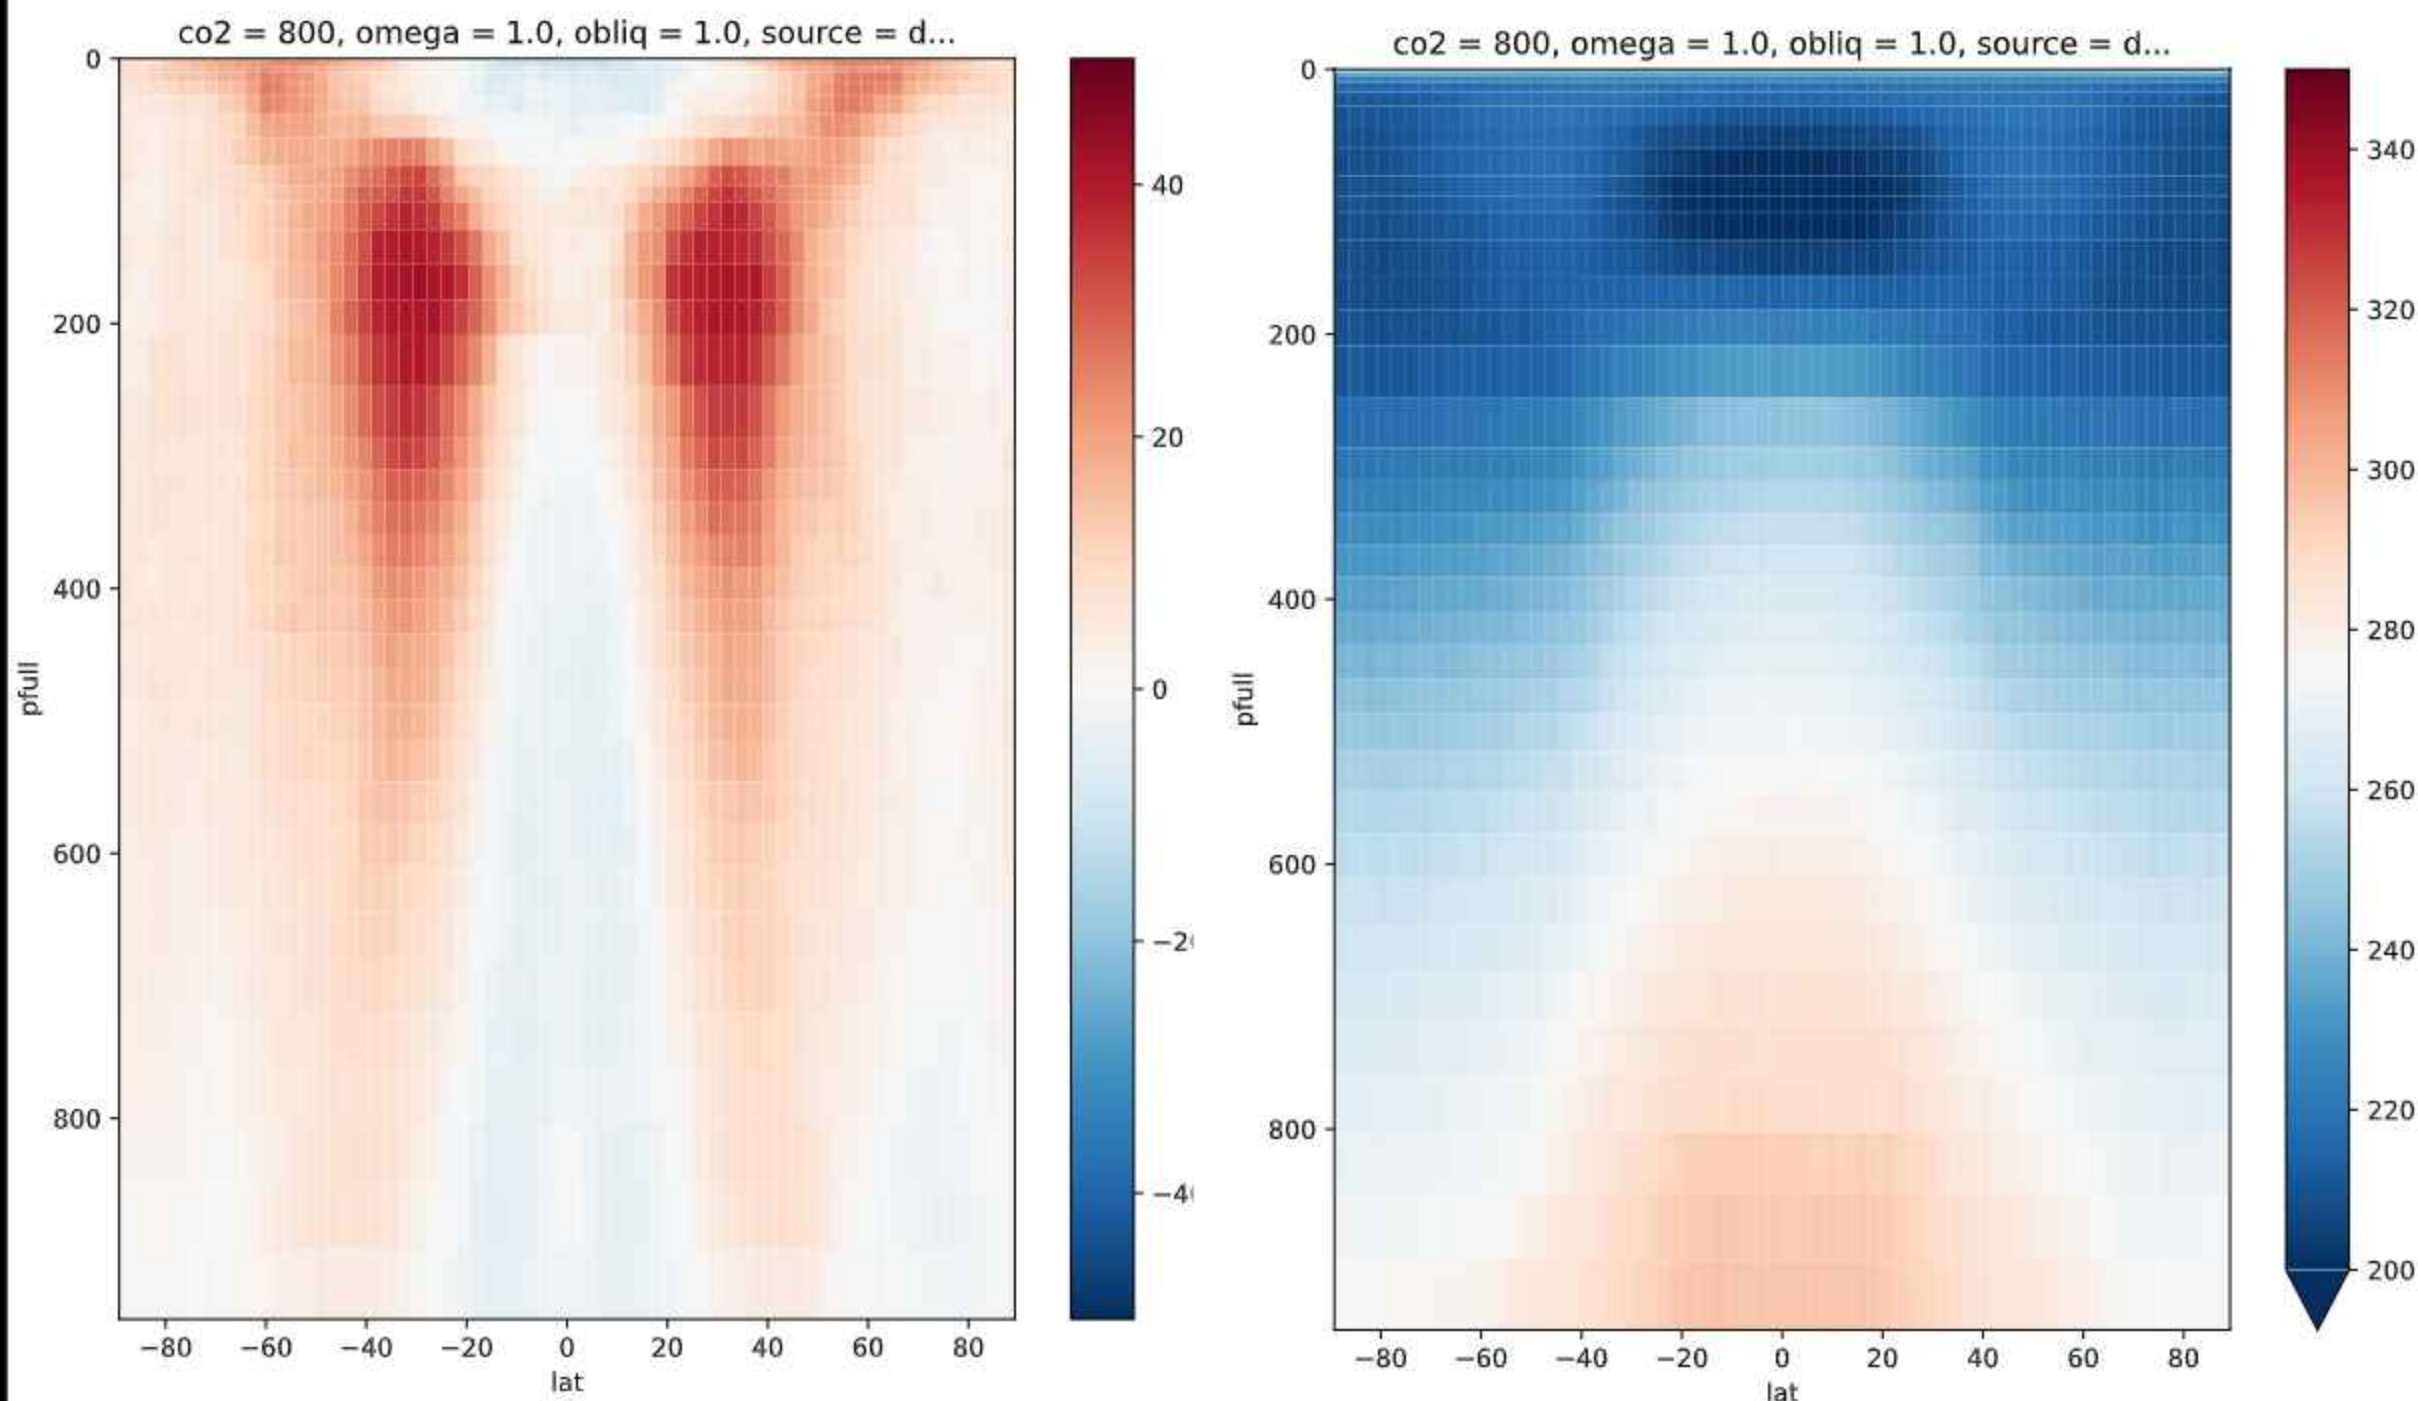


Exploring the data

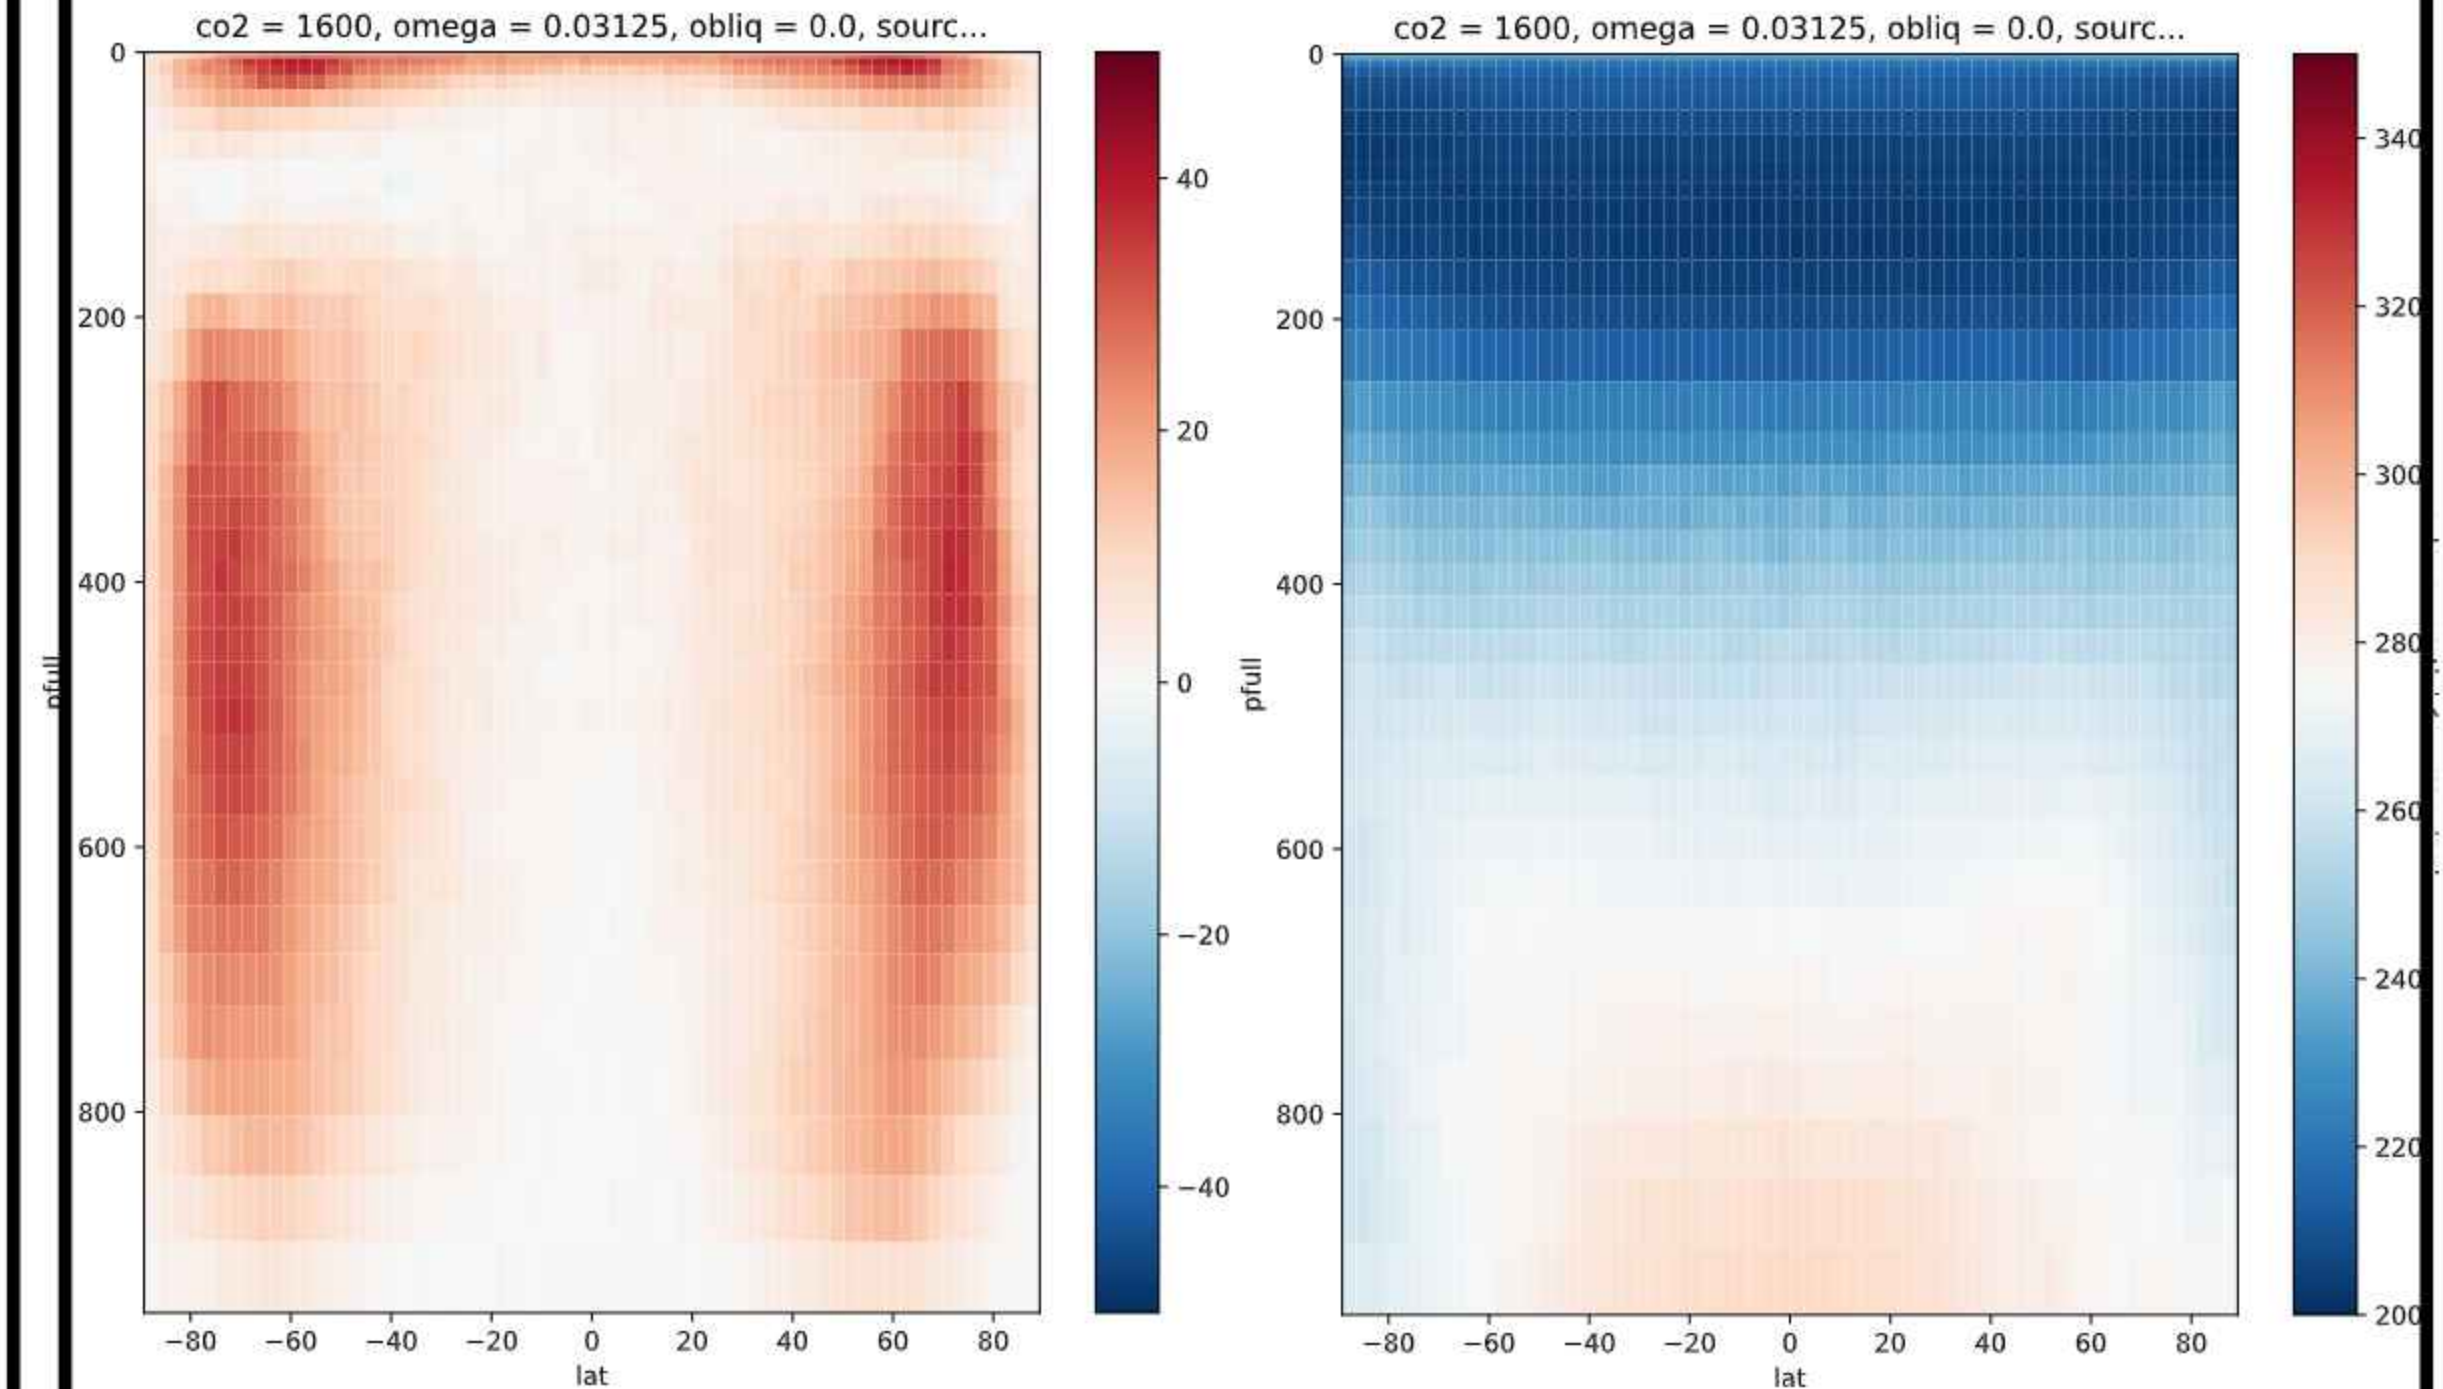
- 7 rotation rates (2, 1, 0.5, 0.25, 0.125, 0.0625, 0.03125) in units of Earth's $\Omega_0 = \frac{2\pi}{86400} s^{-1}$
- 4 obliquities (0, 0.5, 1, 1.5) in units of Earth's obliquity (23°)



CO₂ = 800ppmv, obliquity = 1, omega = 1



CO₂ = 1600ppmv, obliquity = 0, omega = 0.03125

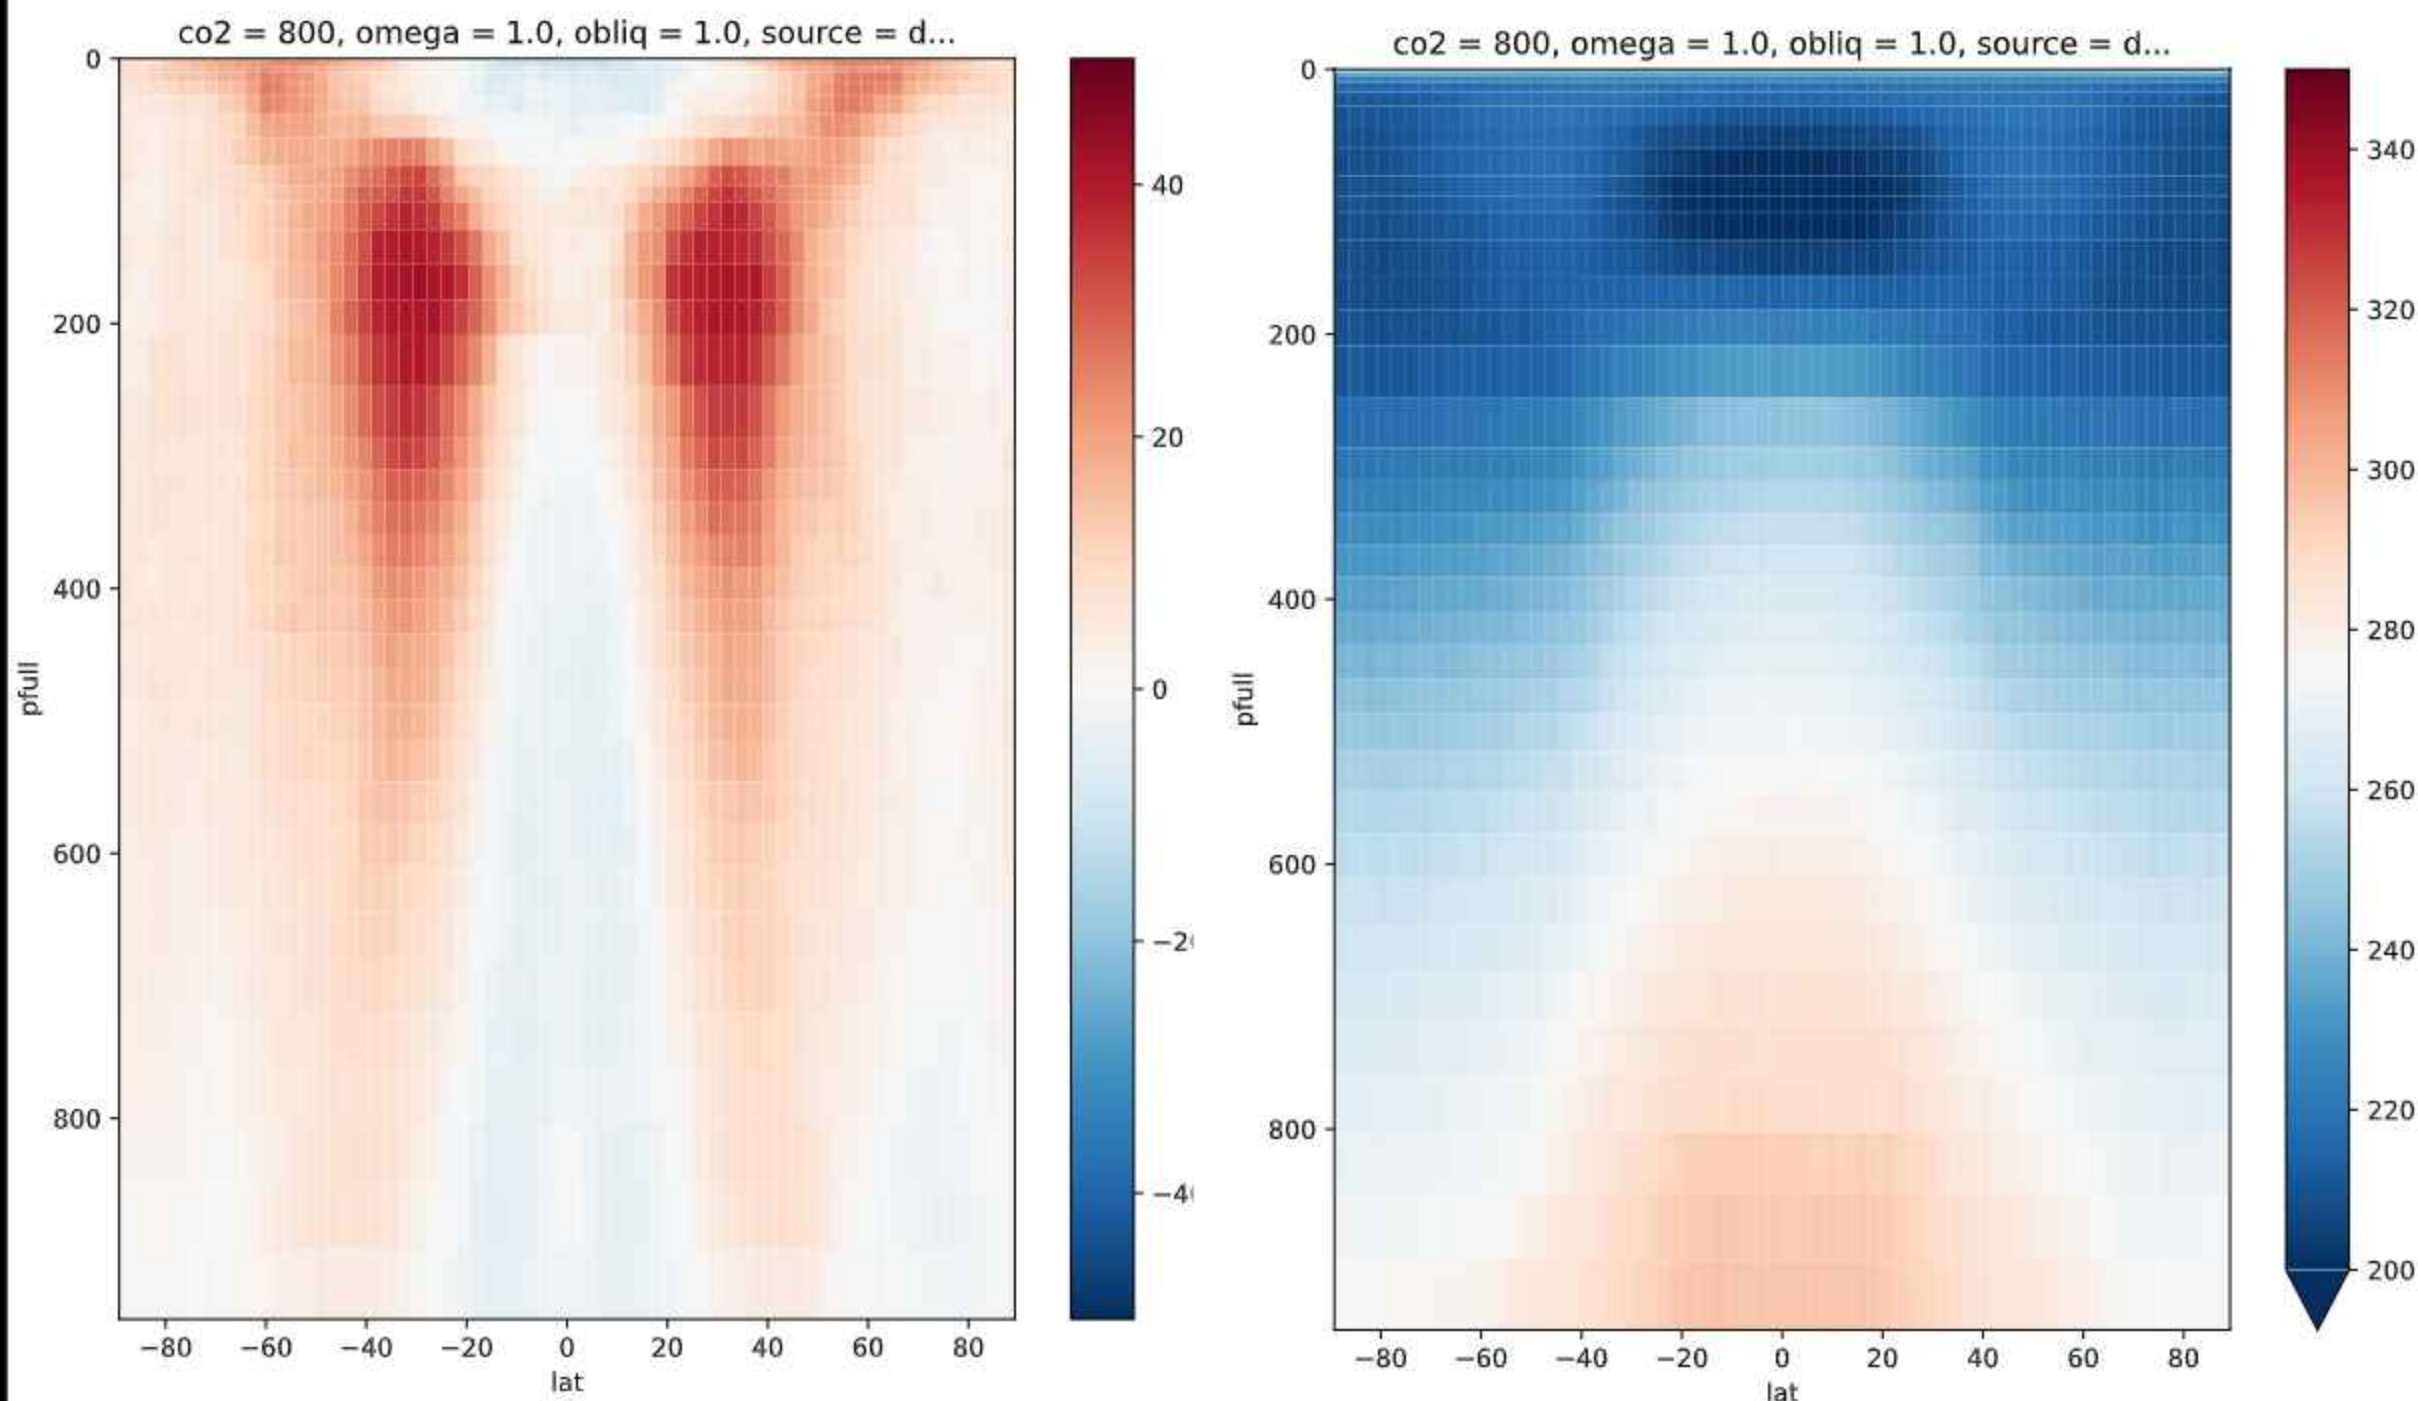


Exploring the data

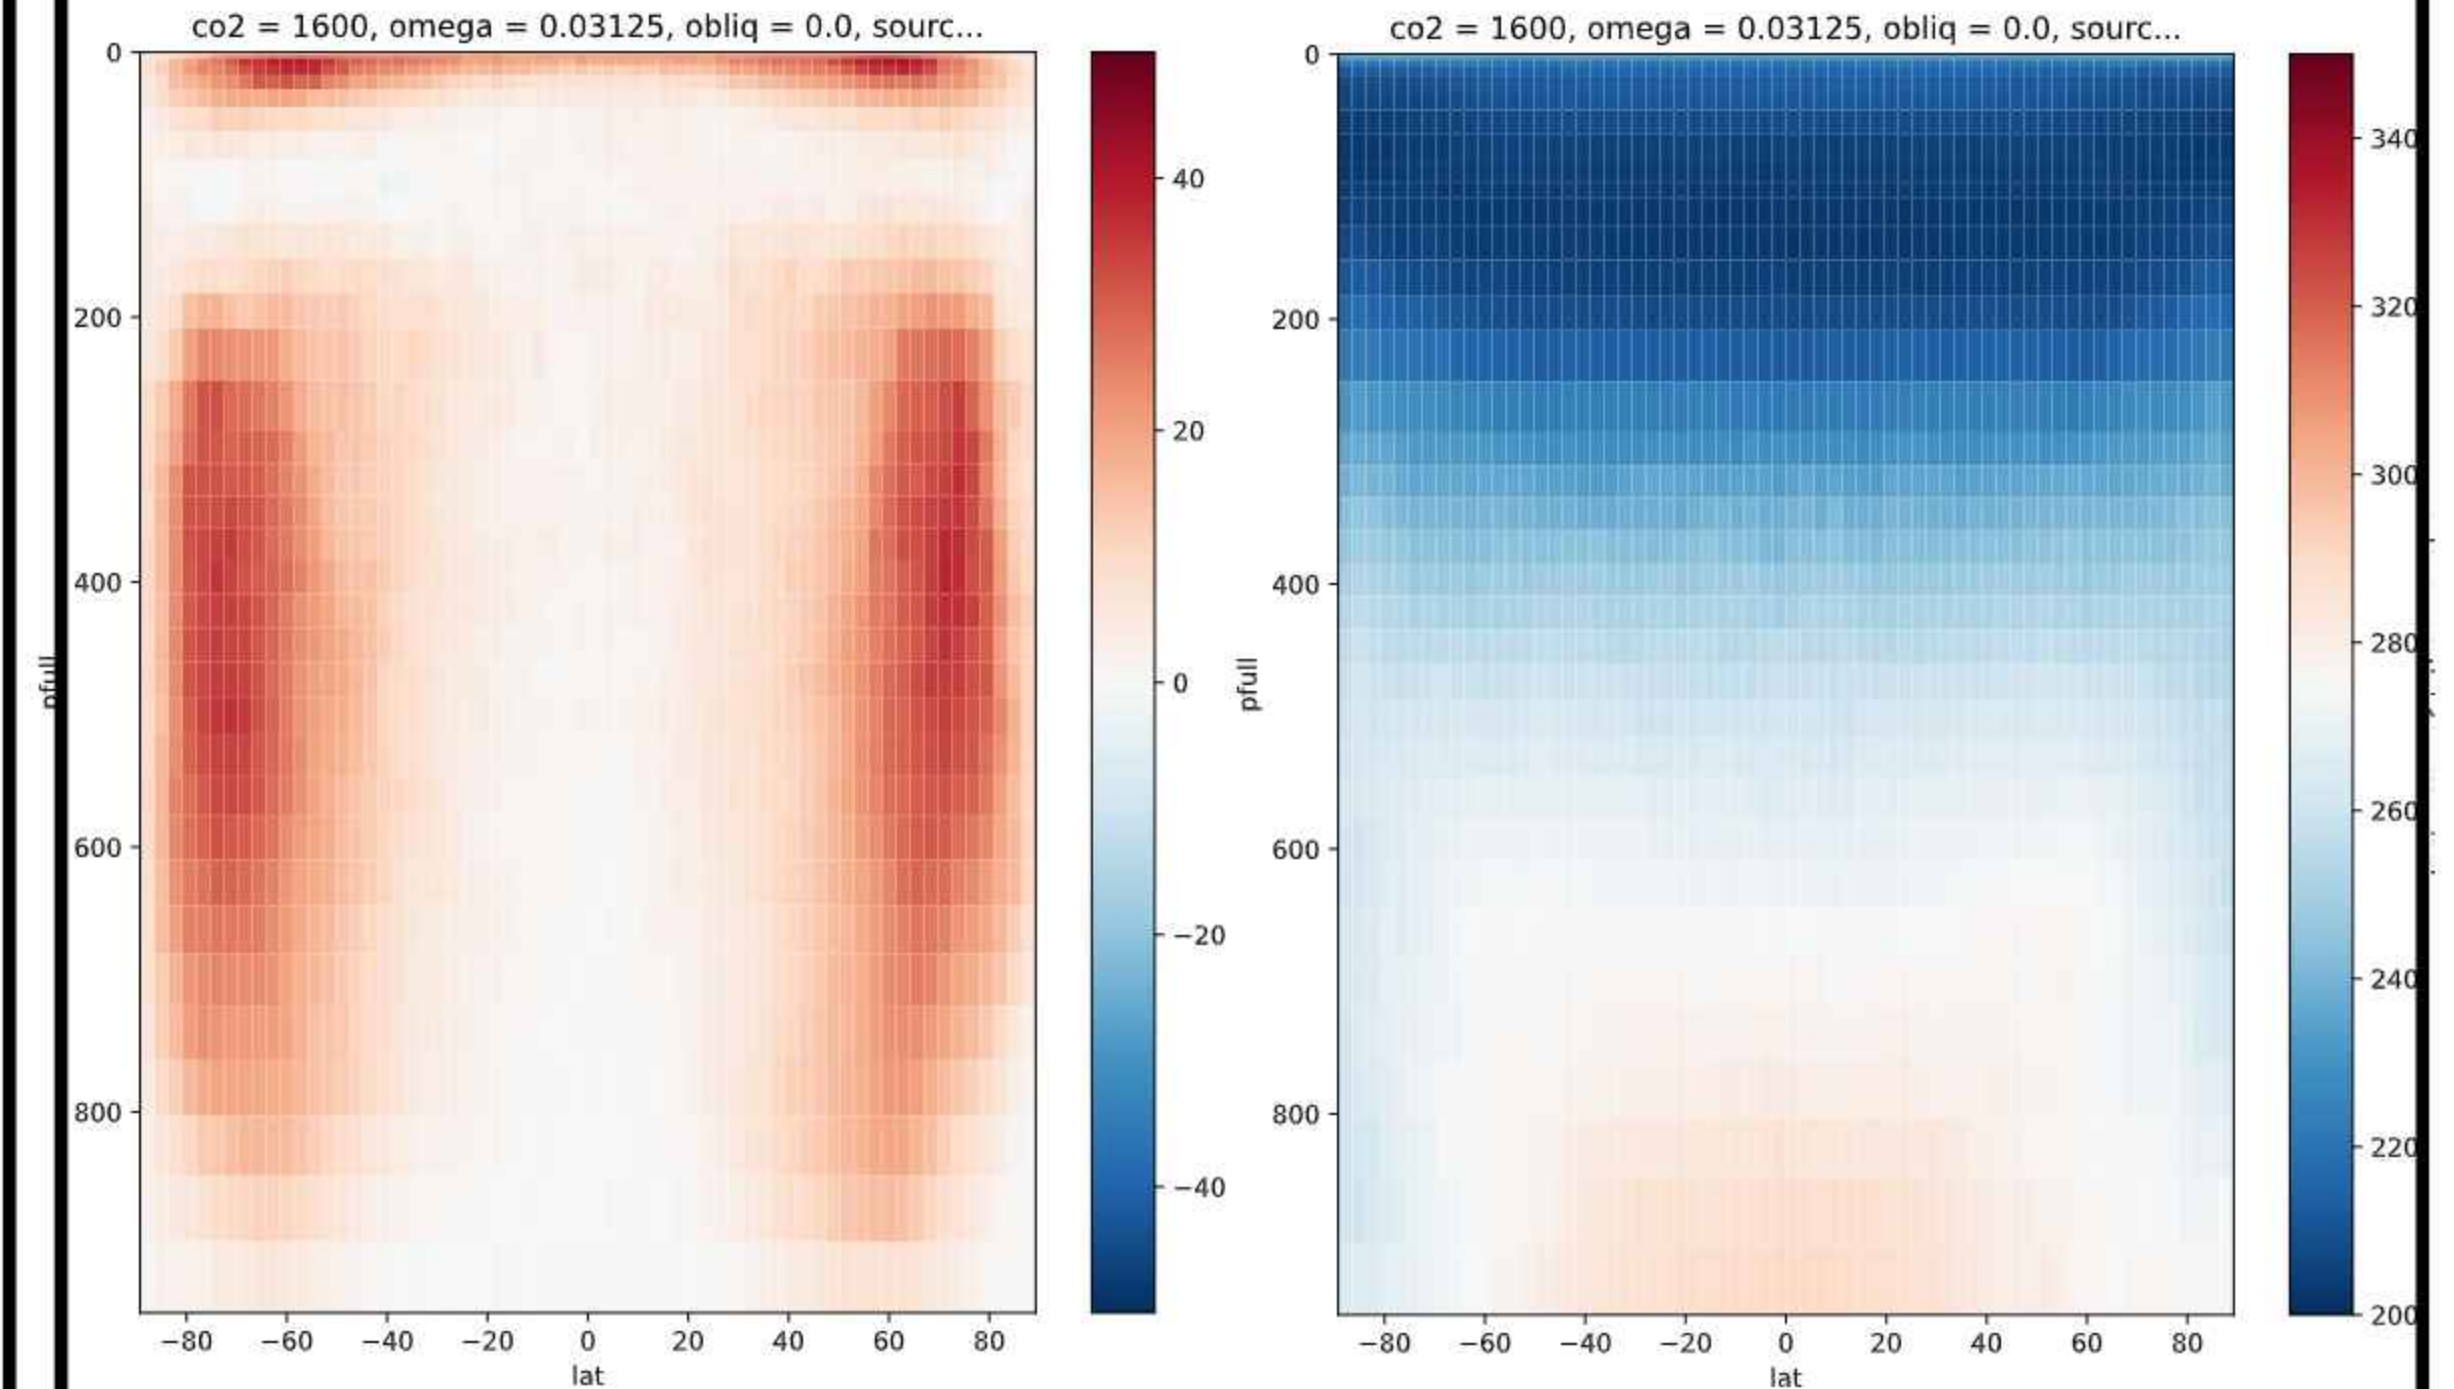
- 7 rotation rates (2, 1, 0.5, 0.25, 0.125, 0.0625, 0.03125) in units of Earth's $\Omega_0 = \frac{2\pi}{86400} s^{-1}$
- 4 obliquities (0, 0.5, 1, 1.5) in units of Earth's obliquity (23°)
- 5 CO₂ concentrations (278, 335, 400, 800, 1600ppmv)



CO₂ = 800ppmv, obliquity = 1, omega = 1



CO₂ = 1600ppmv, obliquity = 0, omega = 0.03125



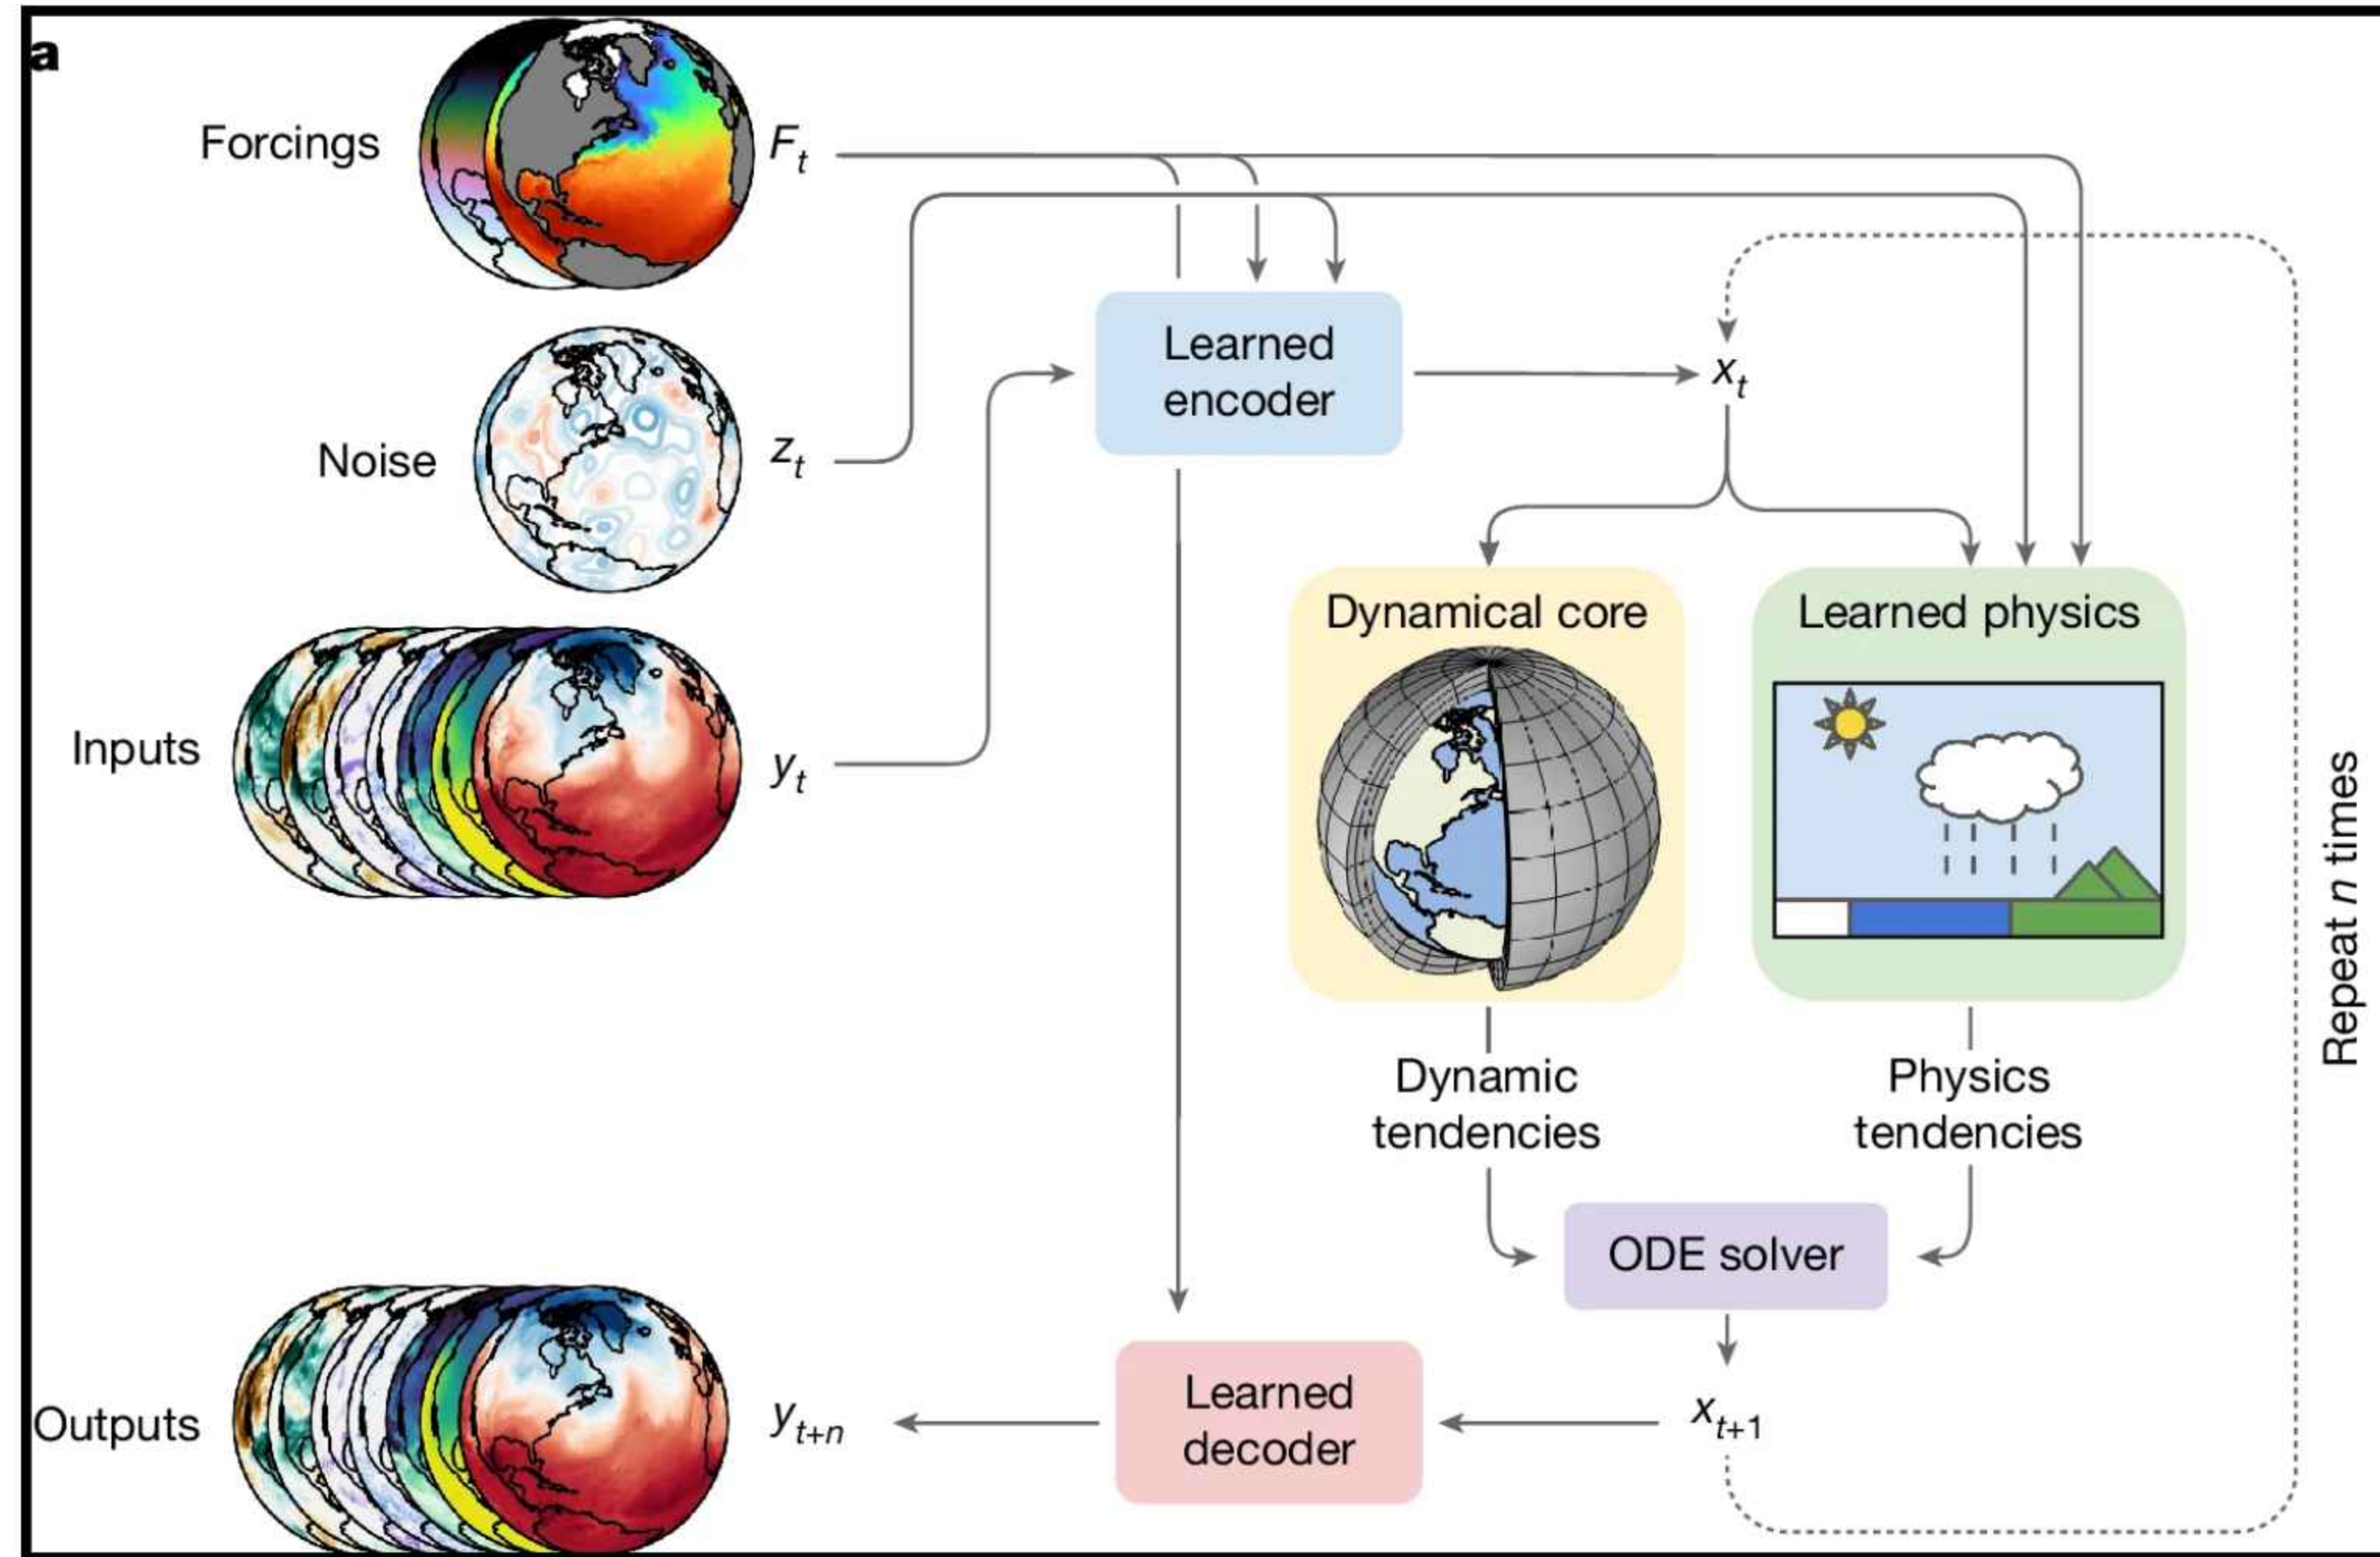
How can we bring in physical constraints?

How can we bring in physical constraints?

- Neural-GCM includes a dynamical core that can be differentiated in a ML sense

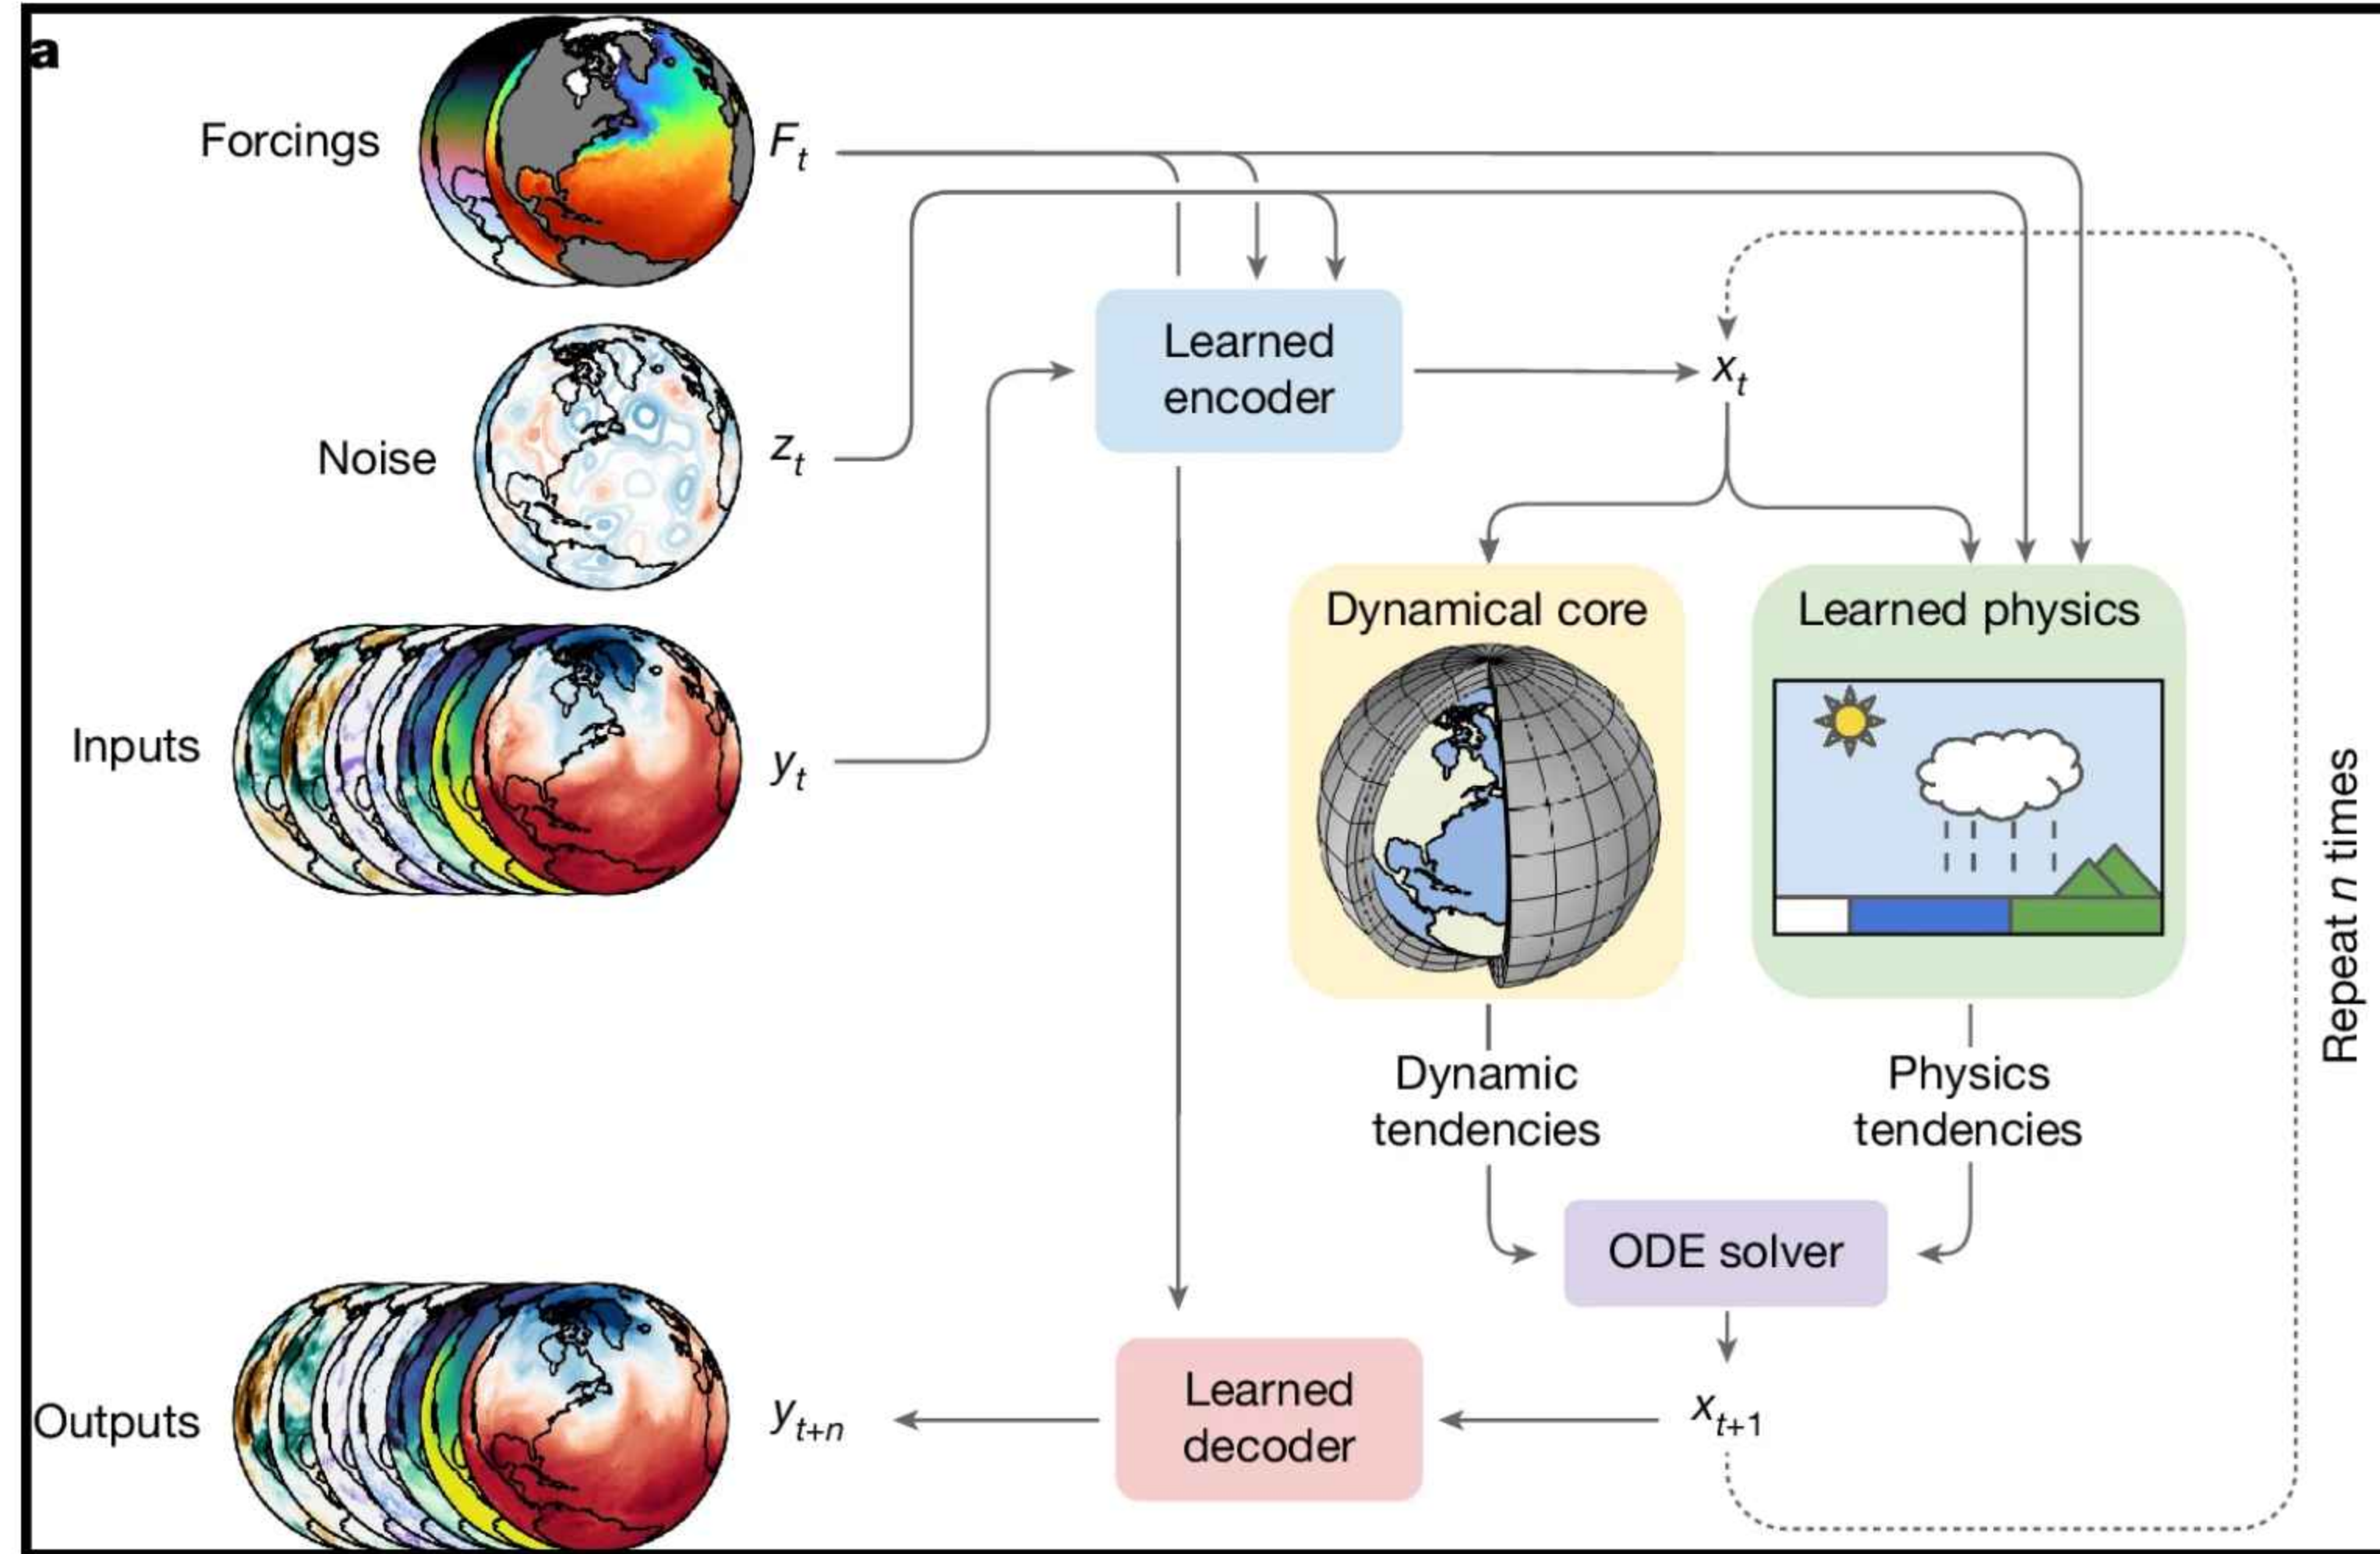
How can we bring in physical constraints?

- Neural-GCM includes a dynamical core that can be differentiated in a ML sense



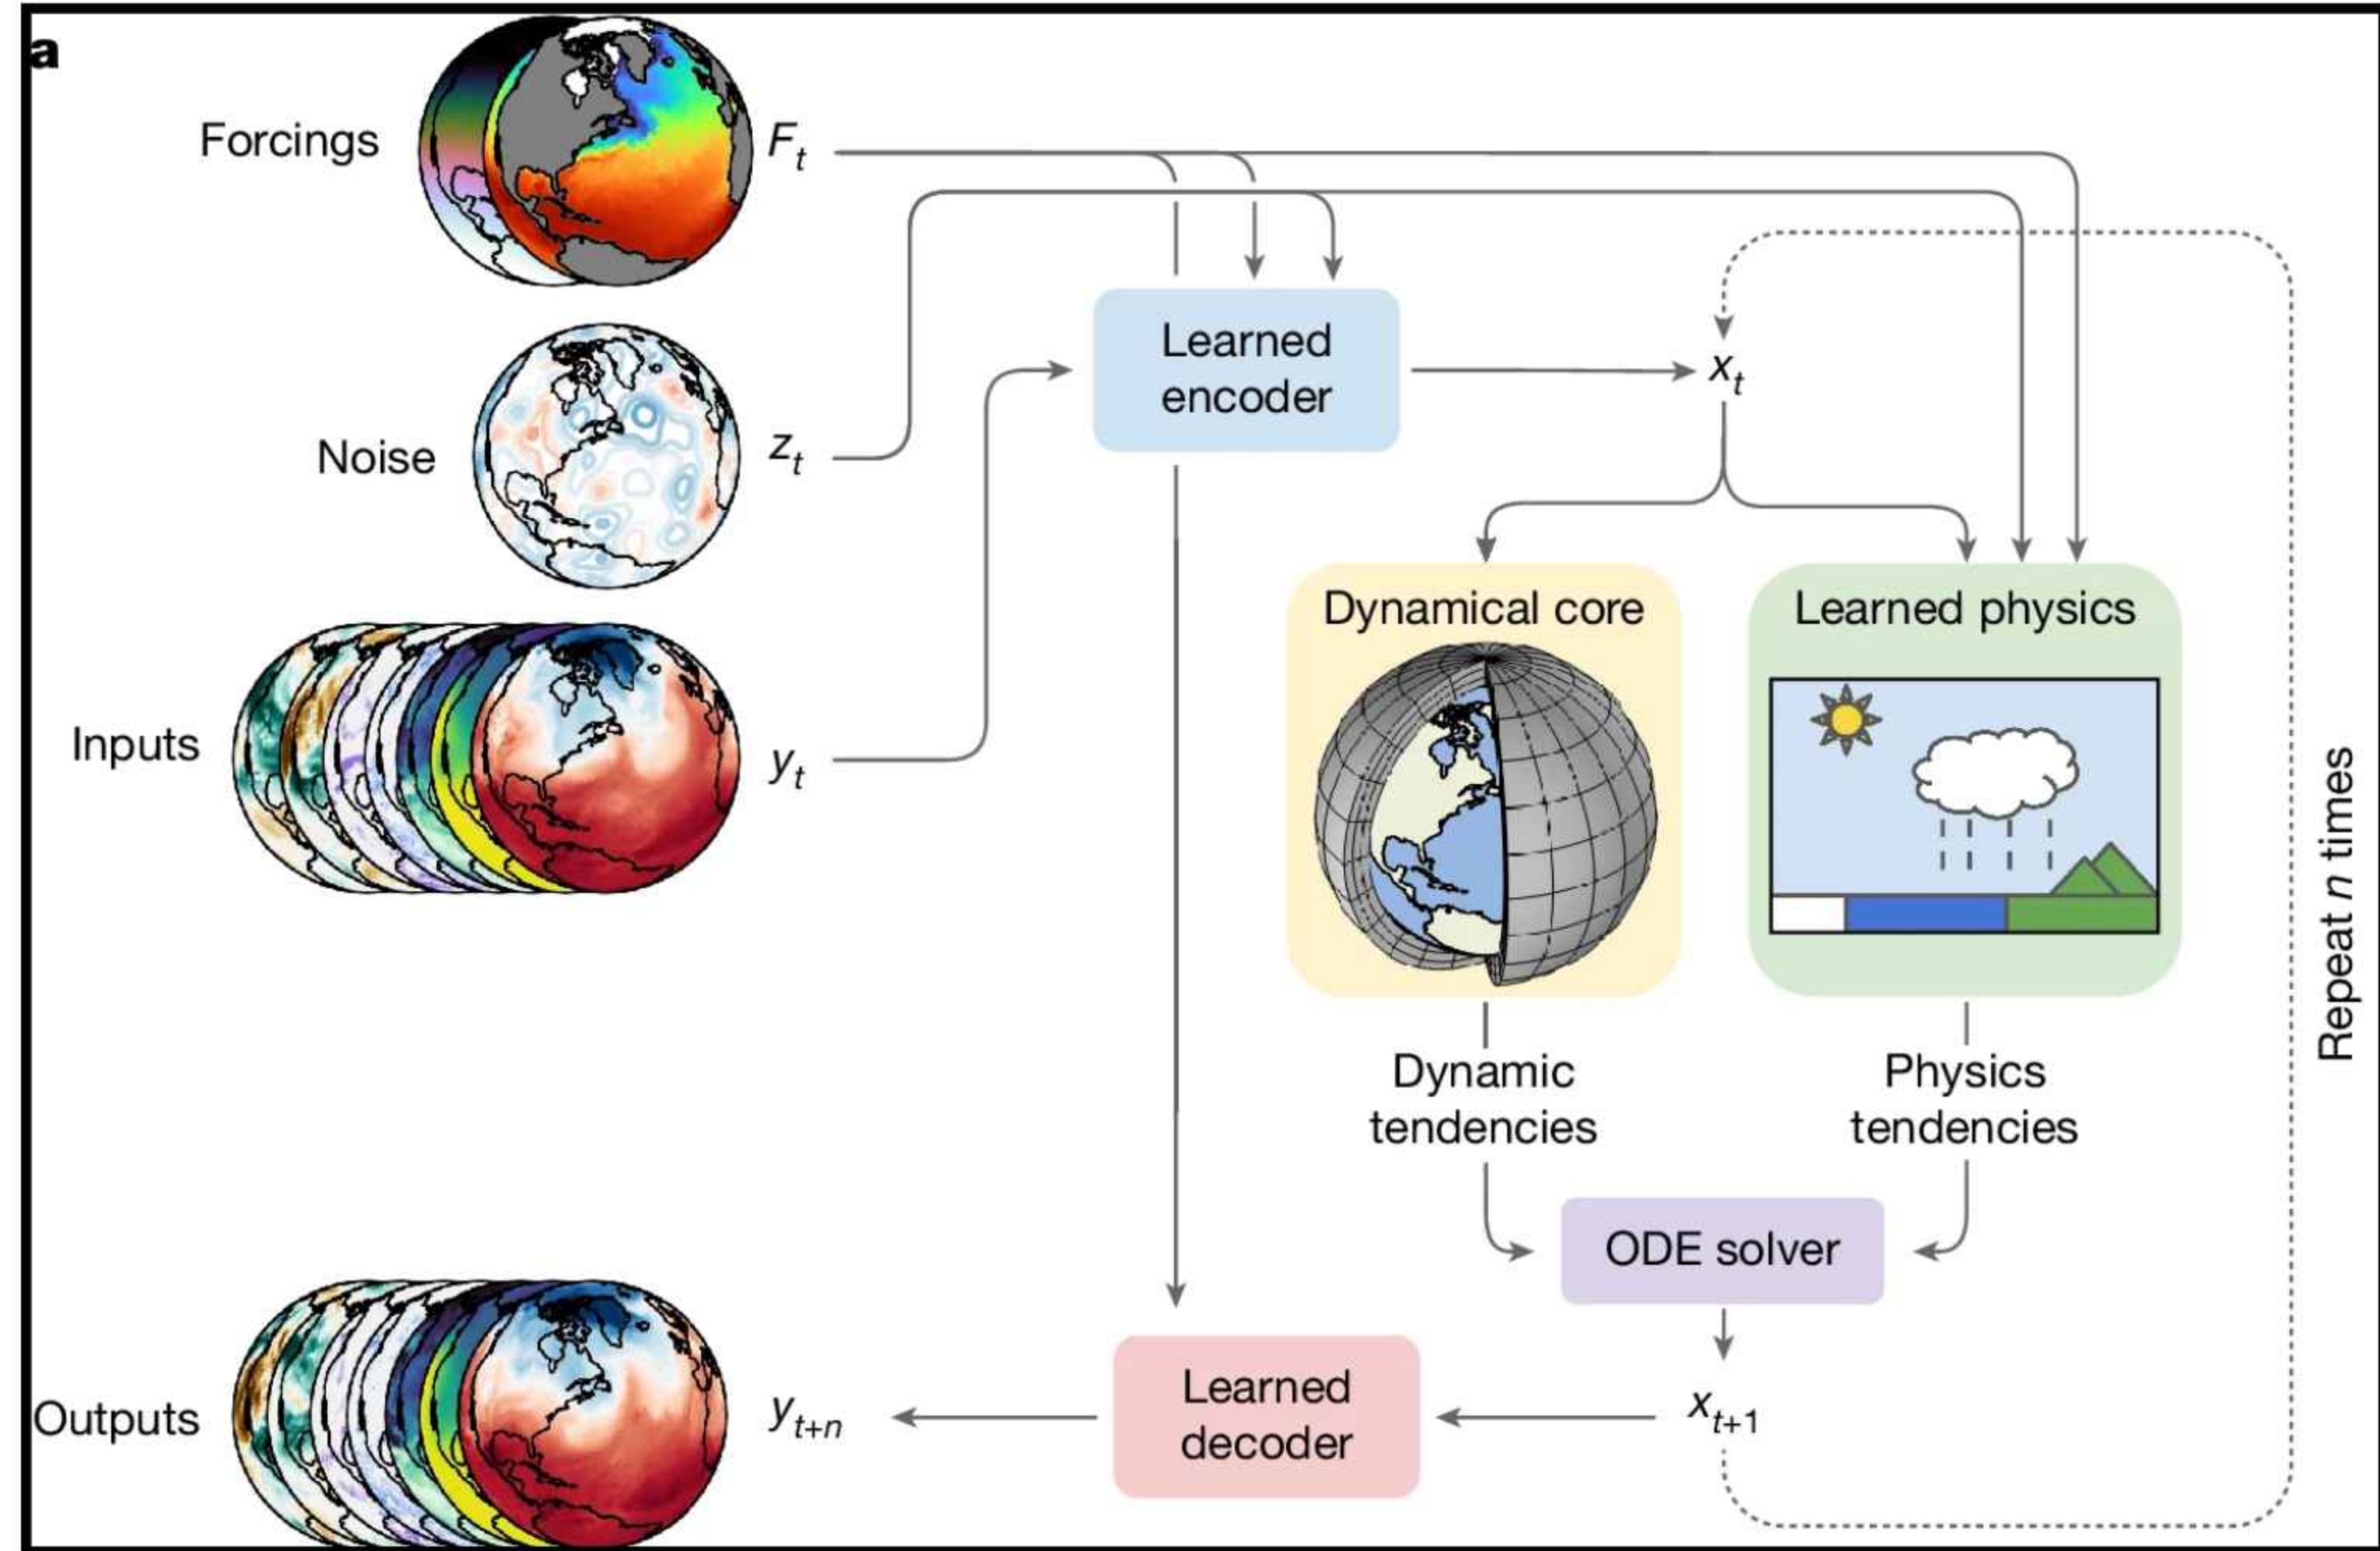
How can we bring in physical constraints?

- Neural-GCM includes a dynamical core that can be differentiated in a ML sense
- Allows physics to be learned online with changes in the dynamical core accounted for



How can we bring in physical constraints?

- Neural-GCM includes a dynamical core that can be differentiated in a ML sense
- Allows physics to be learned online with changes in the dynamical core accounted for
- **Could we do something similar for radiative transfer?**
 - Radiation schemes are very long and complex, and difficult to express analytically
 - A substantial effort to rewrite in differentiable form



AI for Sustainable Systems

Dr Fanlin Meng

f.meng2@exeter.ac.uk

Department of Management (Operations and Analytics)

University of Exeter

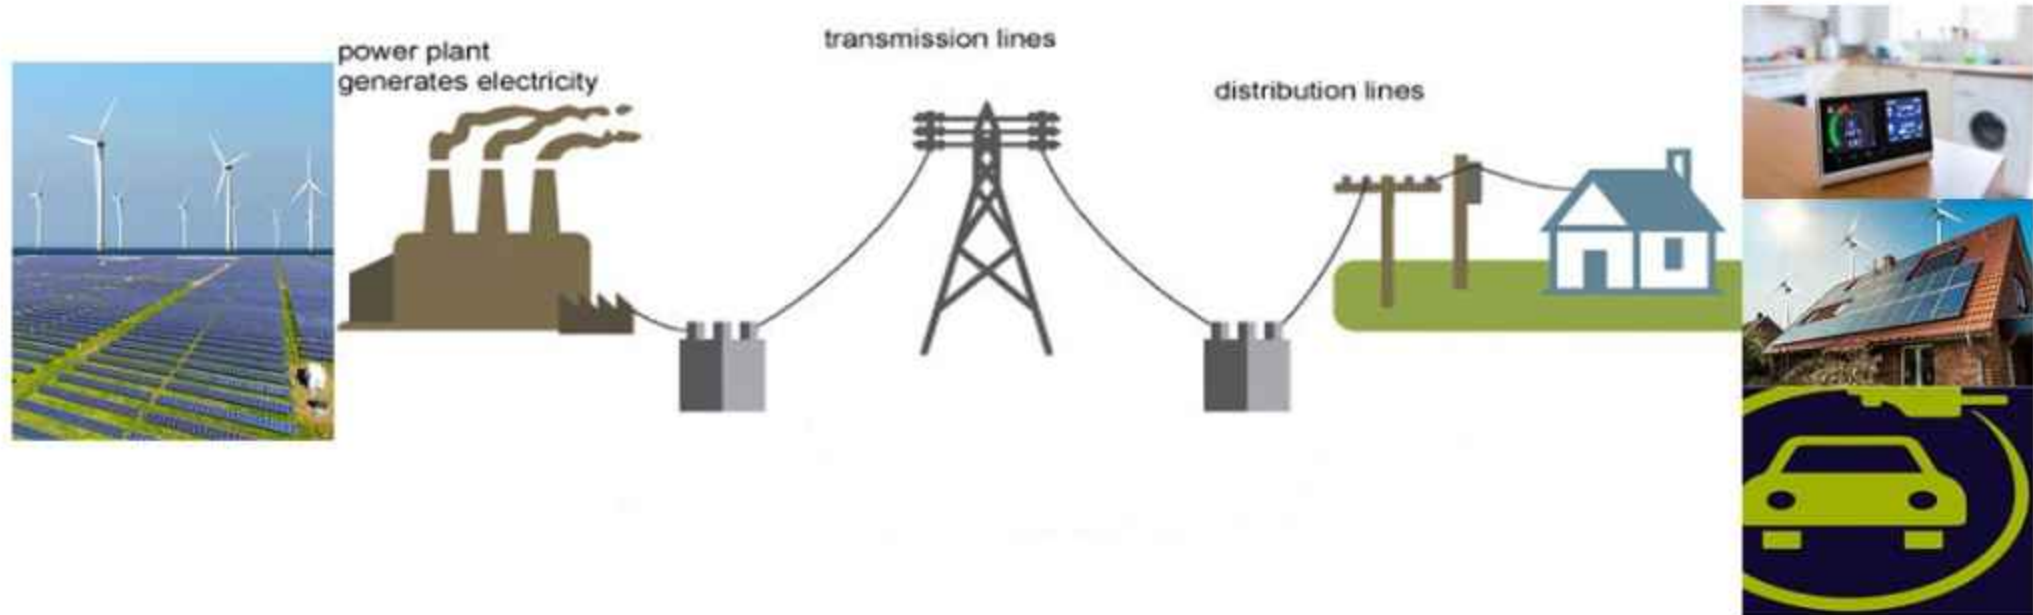
September 11, 2025

AI for Sustainable Systems

- Energy
- Sustainable buildings
- Transportation

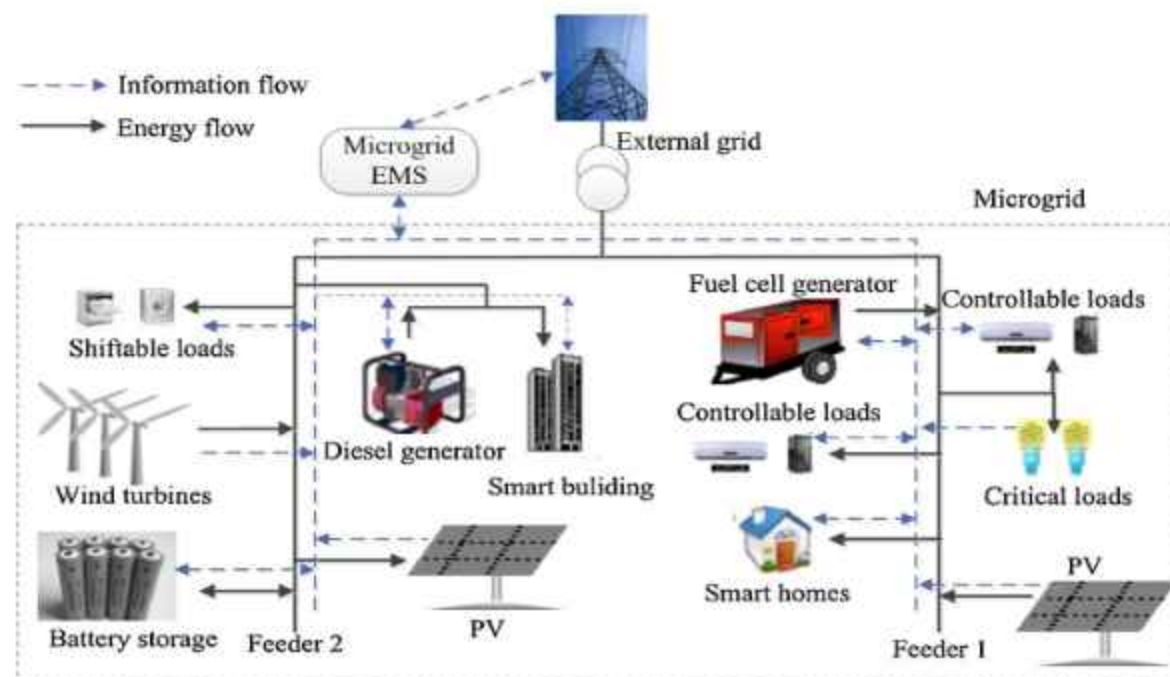
Energy System: An overview

- Minimize operating cost
- Balance supply and demand
- Reduce carbon emissions (net zero)



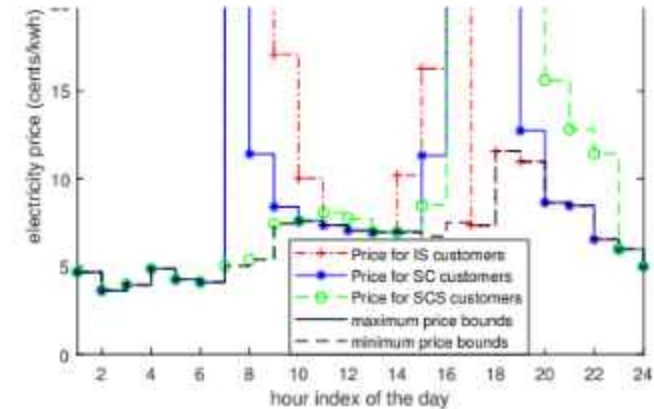
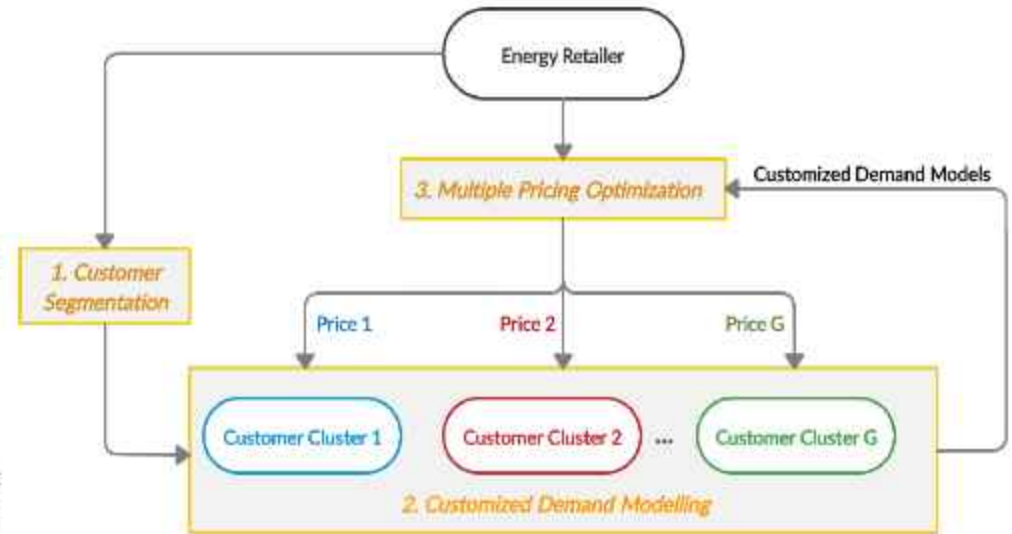
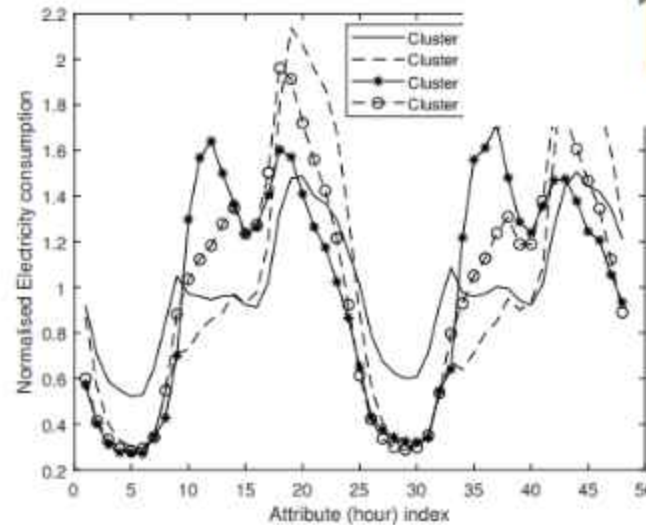
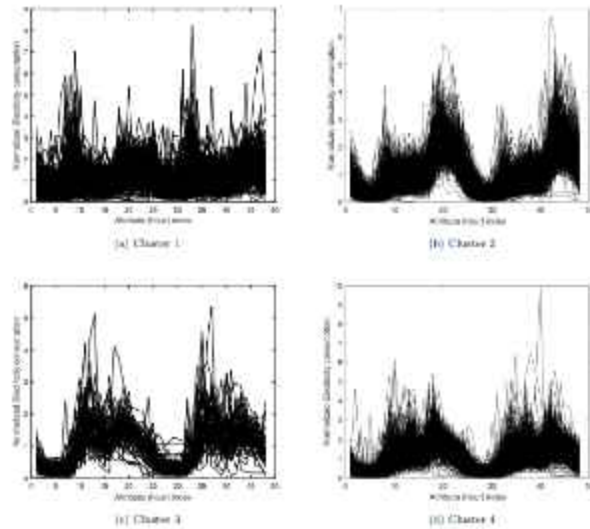
Key Challenges (Forecasting/ Smart Meter Data Analytics)

- Supply:
 - **Renewable generation forecasting**
 - ...
- Demand:
 - **Forecast (e.g., peak demand, EV charging)**
 - **Smart Meter Data Analytics**
 - **Privacy-preserving**
 - ...



Integrated Machine Learning and Optimisation for Smart Energy Pricing

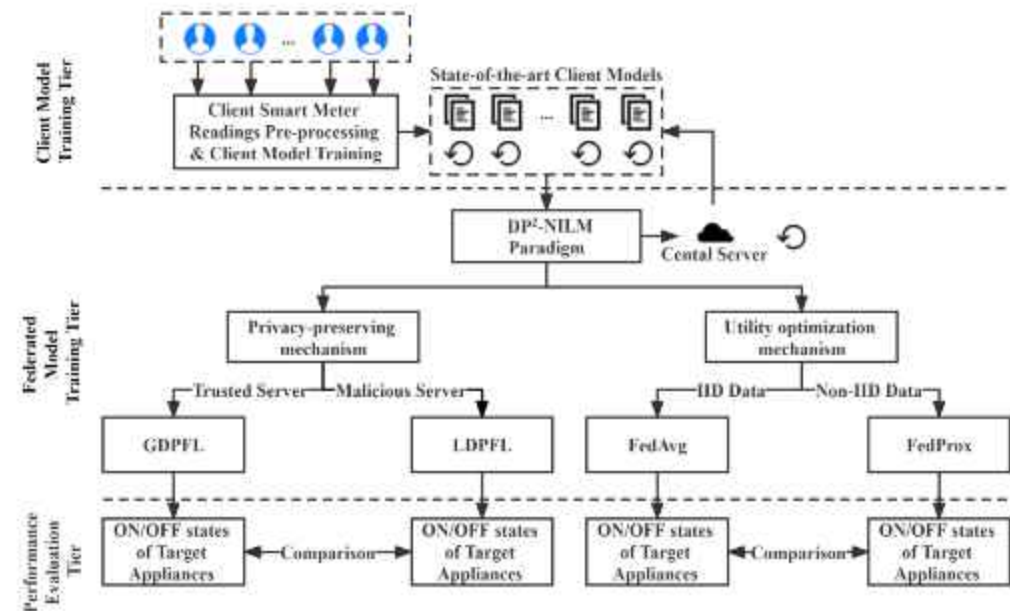
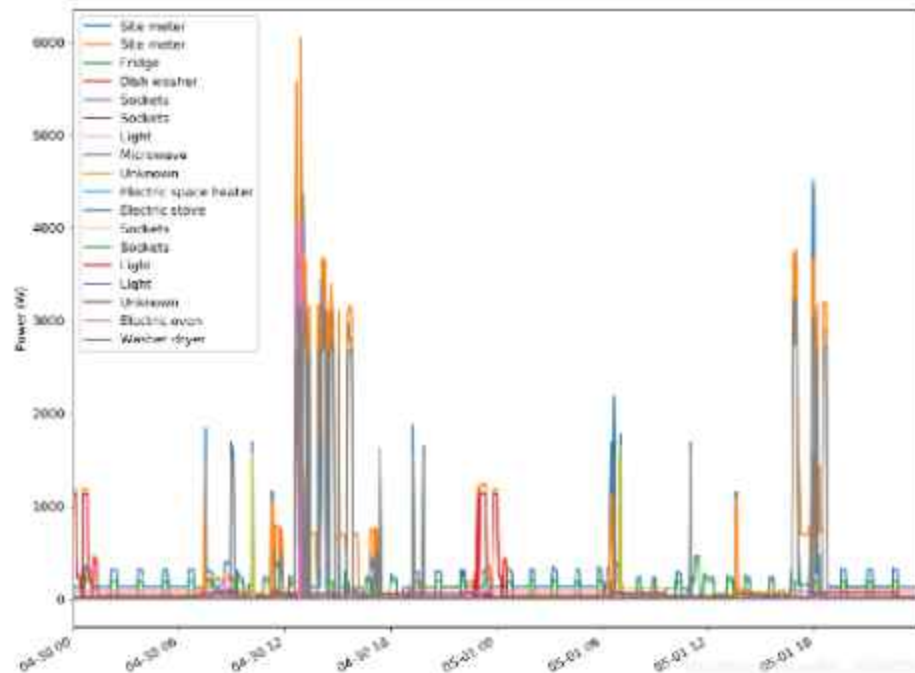
- Smart meter data
- Clustering based customer segmentation
- Customized energy pricing



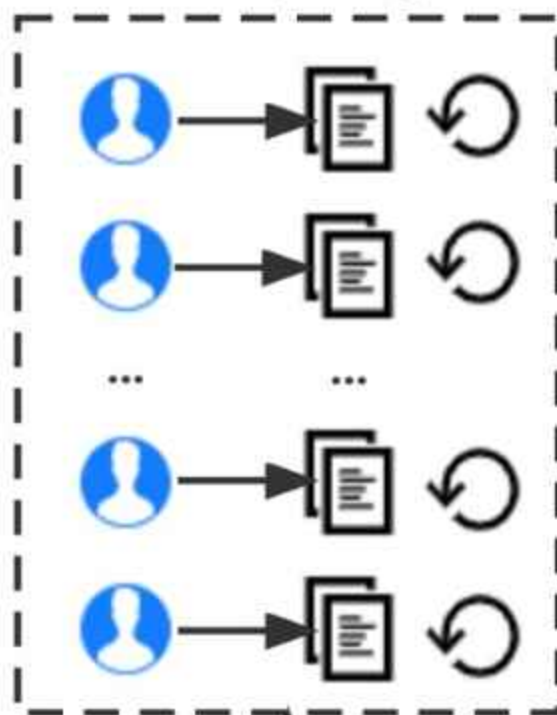
- Meng, F., Ma, Q., Liu, Z., & Zeng, X. J. (2023). Multiple dynamic pricing for demand response with adaptive clustering-based customer segmentation in smart grids. Applied Energy, 333, 120626.

Distributed and privacy-preserving framework for non-intrusive load monitoring

- **Non-intrusive load monitoring (NILM):** load disaggregation from house-level consumption (smart meters) to appliance-level consumption
- **Distributed and privacy-preserving ML framework:** federated learning + differential privacy



**1. Client Model Training
and Updating**



Central Server

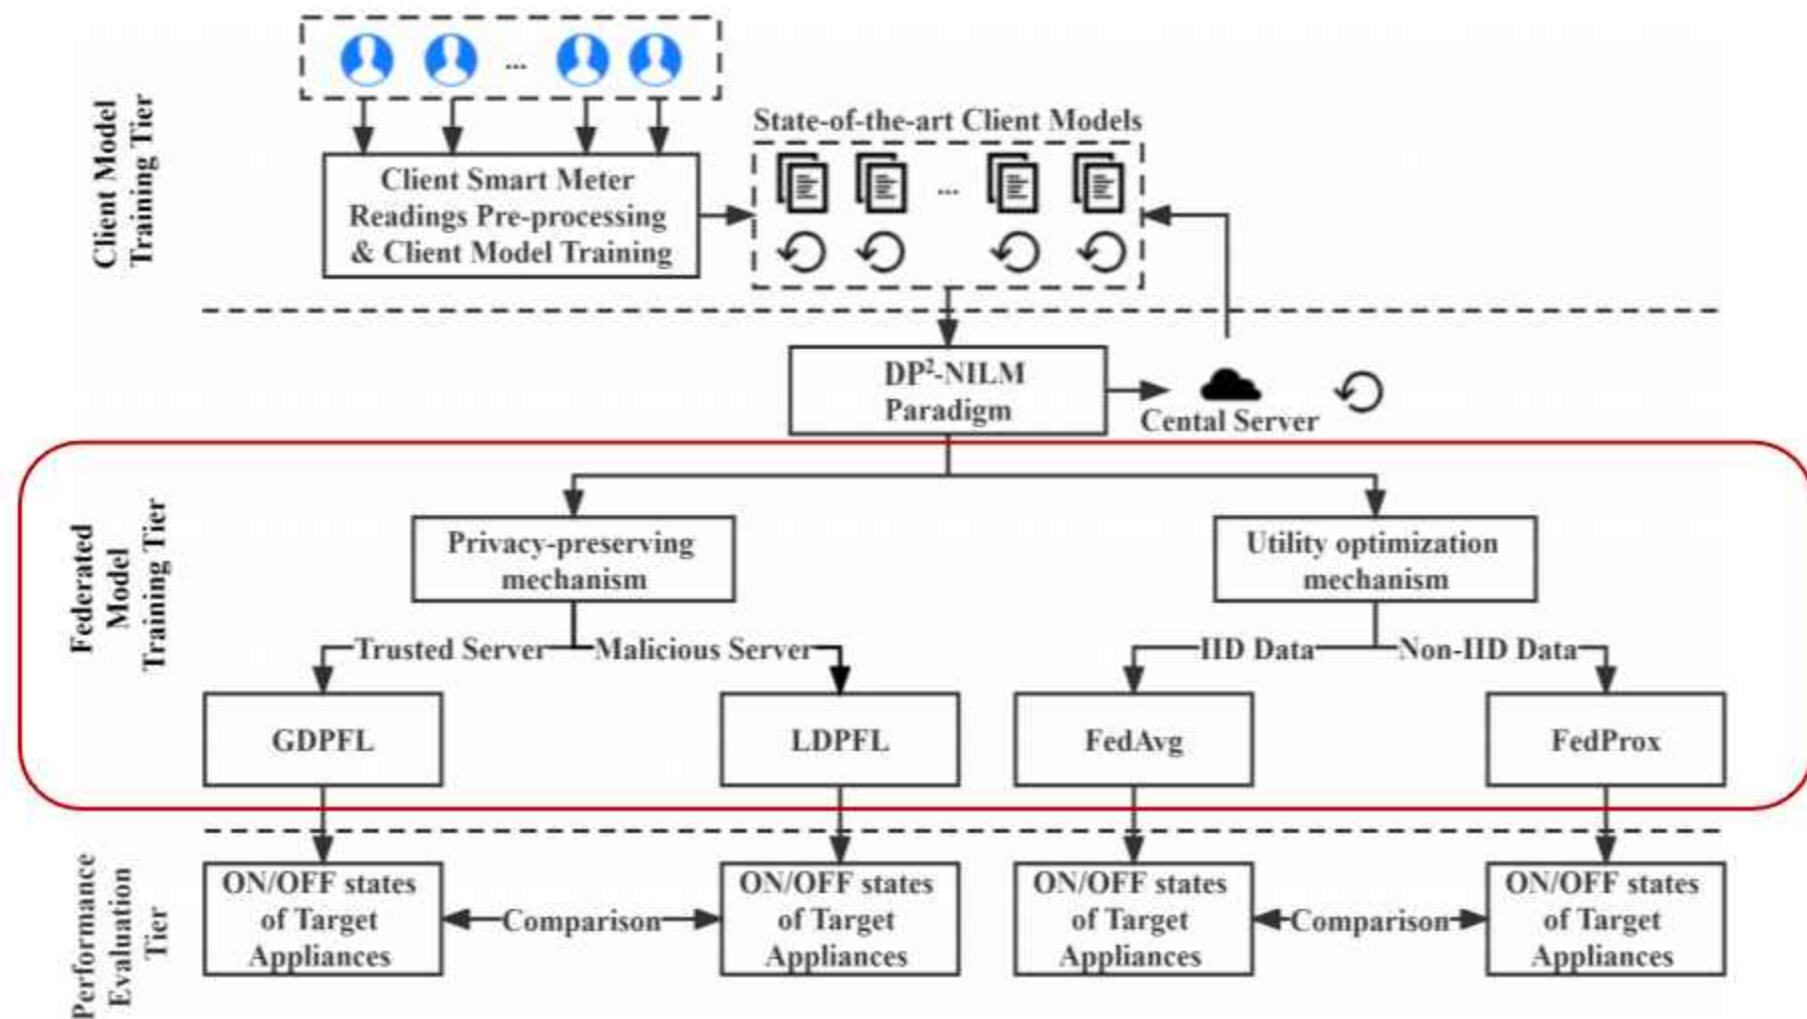
**2. Global Model
Updating**



3. Global Model Broadcasting

A federated deep learning framework

DP²-NILM



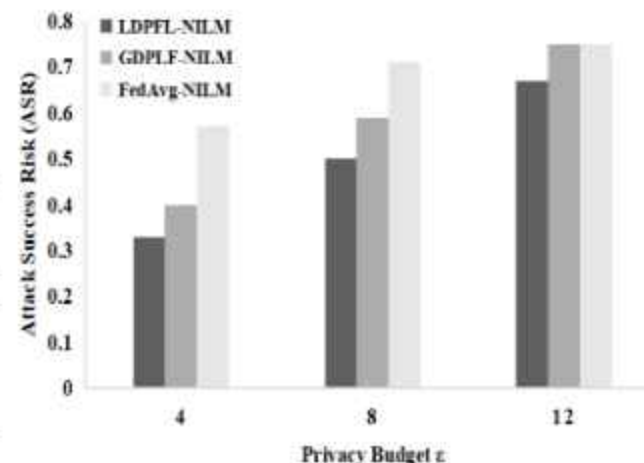
Proposed DP²-NILM framework

Utility Optimization mechanism

	Fridge				Dishwasher				Washing Machine			
	Accuracy	F ₁	Precision	Recall	Accuracy	F ₁	Precision	Recall	Accuracy	F ₁	Precision	Recall
FedAvg-NILM	0.65	0.63	0.50	0.85	0.97	0.75	0.92	0.64	0.98	0.71	0.83	0.62
FedProx-NILM	0.85	0.81	0.82	0.81	0.98	0.80	0.78	0.82	0.97	0.54	0.83	0.40
Evaluation	(↑ 20%)	(↑ 18%)	(↑ 32%)	(↓ 4%)	(↑ 1%)	(↑ 5%)	(↓ 19%)	(↑ 18%)	(↓ 1%)	(↓ 17%)	(-)	(↓ 22%)

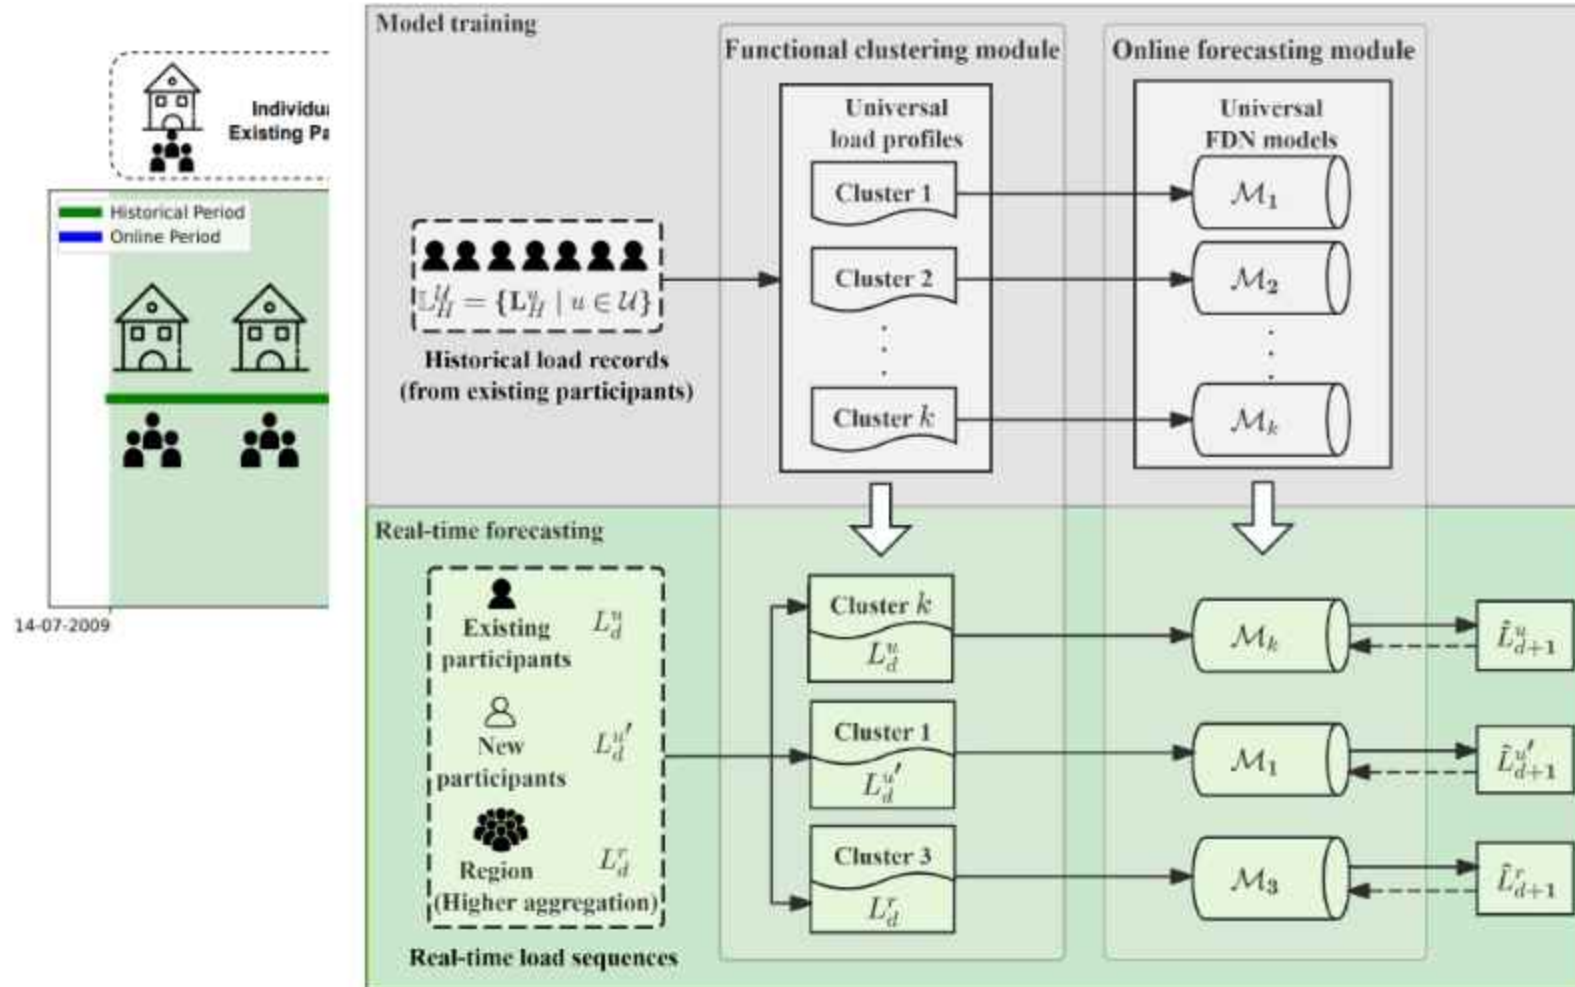
Privacy-preserving mechanism

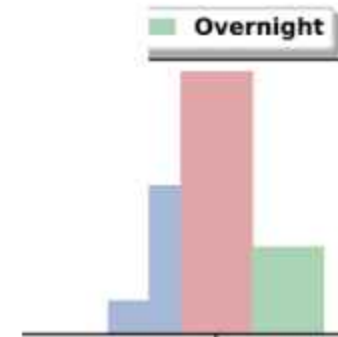
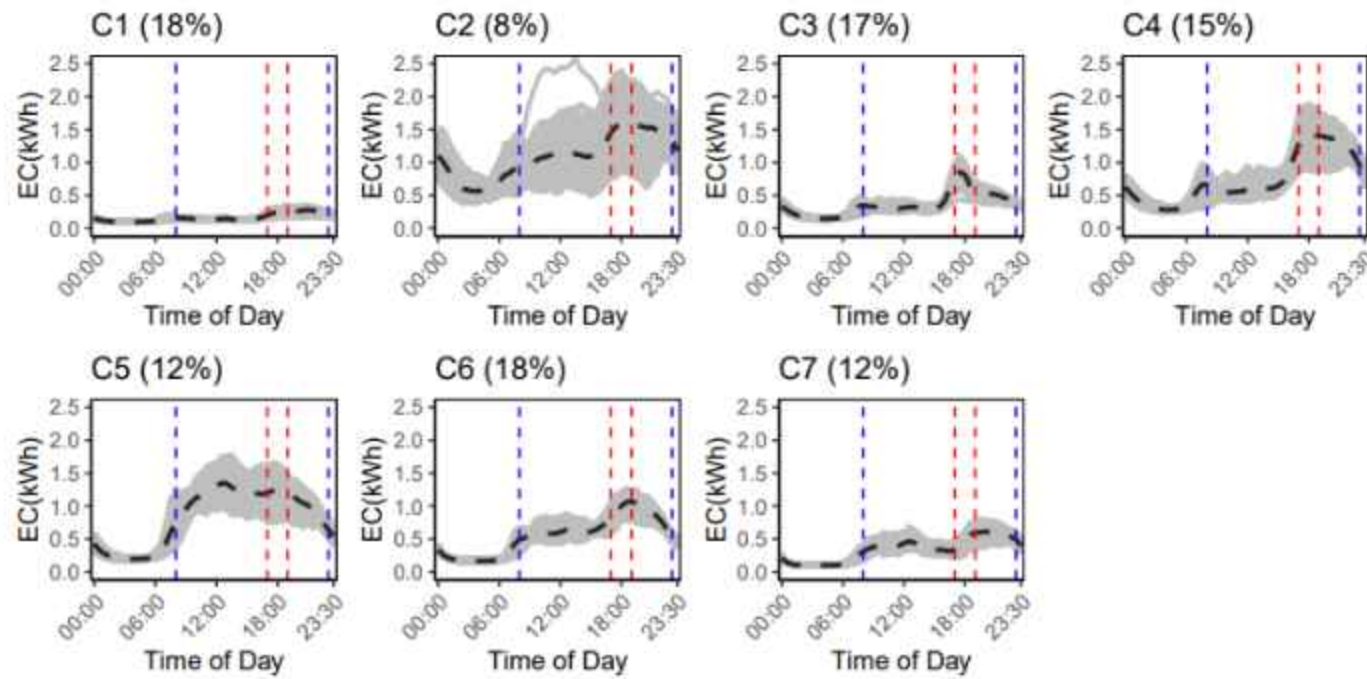
	Privacy Budget	Fridge				Dishwasher				Washing Machine				Privacy Guarantee	Trusted Server
		Accuracy	F ₁	Precision	Recall	Accuracy	F ₁	Precision	Recall	Accuracy	F ₁	Precision	Recall	Strong	No
FedAvg-NILM	\	0.65	0.63	0.50	0.85	0.97	0.75	0.92	0.64	0.98	0.71	0.83	0.62		
GDPFL-NILM	4	0.54	0.53	0.37	0.95	0.90	0.14	0.16	0.72	0.95	0.68	0.40	0.92		
	8	0.63	0.61	0.49	0.84	0.97	0.69	0.93	0.56	0.98	0.68	0.79	0.60		
	12	0.66	0.82	0.81	0.83	0.99	0.85	0.86	0.85	0.98	0.74	0.80	0.63		
LDPFL-NILM	4	0.58	0.40	0.40	0.38	0.93	0.11	0.21	0.39	0.94	0.10	0.11	0.34	Strong	No
	8	0.58	0.42	0.41	0.44	0.94	0.20	0.30	0.40	0.96	0.20	0.40	0.47		
	12	0.65	0.42	0.36	0.50	0.94	0.13	0.26	0.48	0.96	0.43	0.40	0.50		



Online Energy Forecasting

Online demand forecasting based on functional data analysis





The seven universal load profiles obtained from the functional clustering module. Peak period: 17:00–18:30; Daytime period: 8:00–16:30 and 19:00–23:00; Overnight period: 23:00–07:30. The black dashed line represents the cluster centroid of each profile.

Demand shift reflected in intra-day volatility score (IVS)

Sustainable Building

- thermal comfort: health and productivity
- thermal profiles vary across buildings and within the buildings
- targeted thermal and ventilation control



(a)



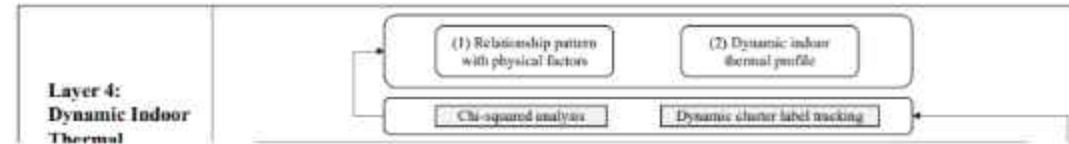
(b)



(c)

The seven-floor conceptual office modeled in DesignBuilder, (a) front elevation; (b) north-east elevation; (c) south elevation.

AI for sustainable building thermal management



Category summary for the dynamic cluster trajectories over the year. DT: dynamic trajectory.

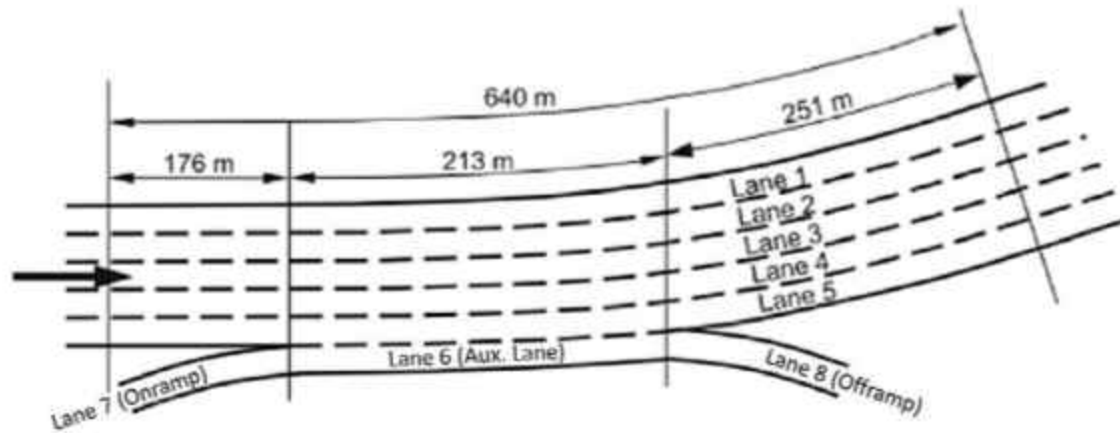
Spring Cluster	Summer Cluster	Autumn Cluster	Winter Cluster	Floor Level	Floor Area	Mean Temp.	DT No.
1	1	2	2	1 (low)	28.14 (med.)	19.58	1
		1	1	2 (low)	31.72 (large)	22.53	2
	2	1	2	4.50 (med.)	34.22 (large)	22.79	3
		2	2	7 (high)	30.03 (med.)	23.42	4
2	2	1	1	4 (med.)	32.11 (large)	24.56	5
		1	2	6.50 (high)	30.24 (med.)	24.14	6
		2	1	3.88 (low-med.)	24.87 (small)	24.68	7
		2	2	6.08 (high)	23.85 (small)	24.69	8



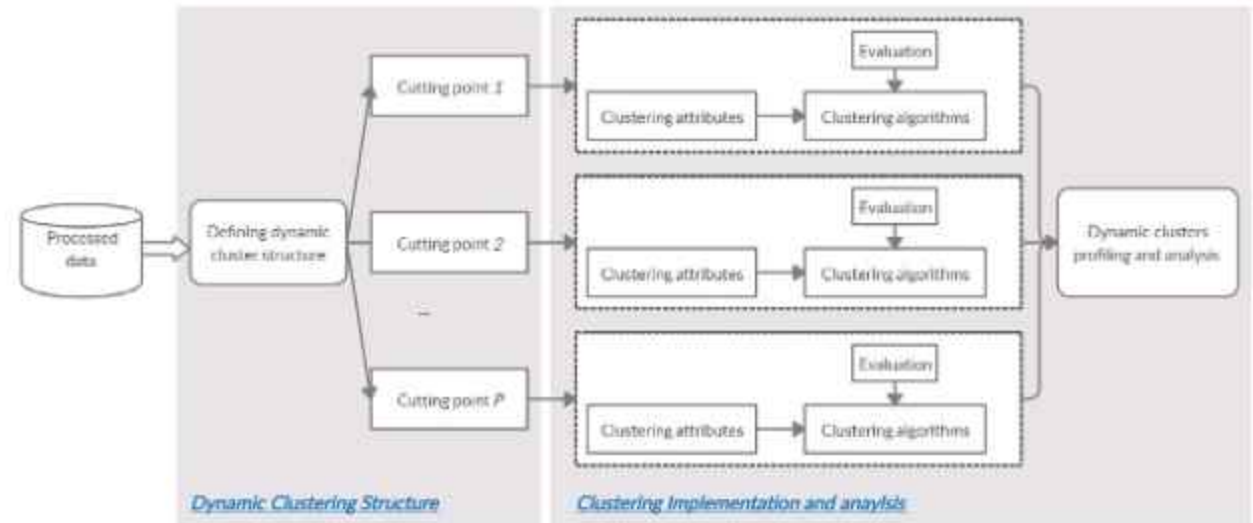
The system framework of clustering-based indoor temperature profiling.

Intelligent Transportation System (Driving styles)

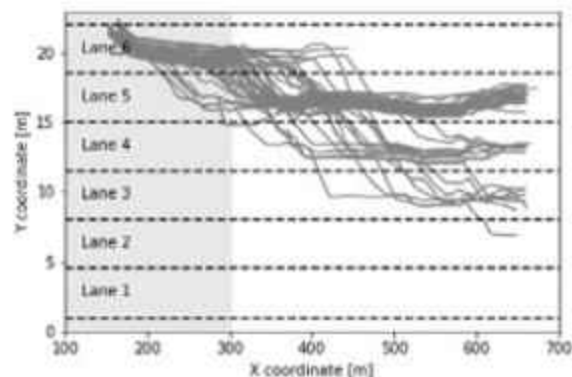
- Driving styles
- Efficiency (travel, energy, etc.)
- Safety
- Autonomous driving



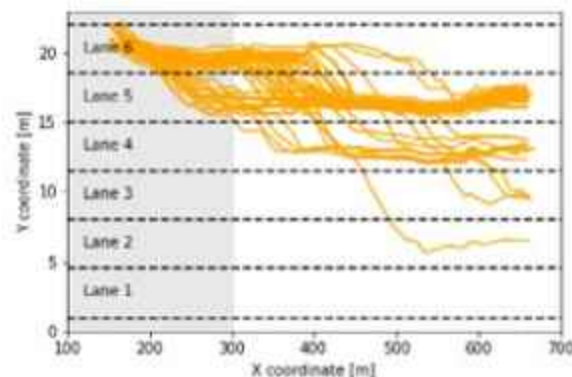
US 101 highway diagram.



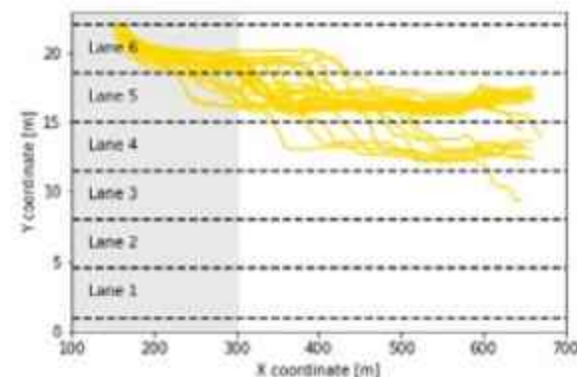
Dynamic clustering for driving styles analysis



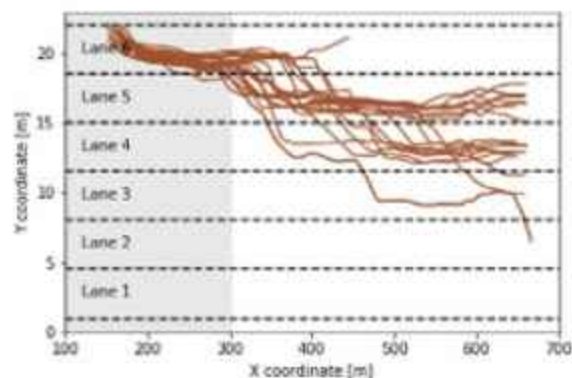
(a) Cluster $C_{1,300}$



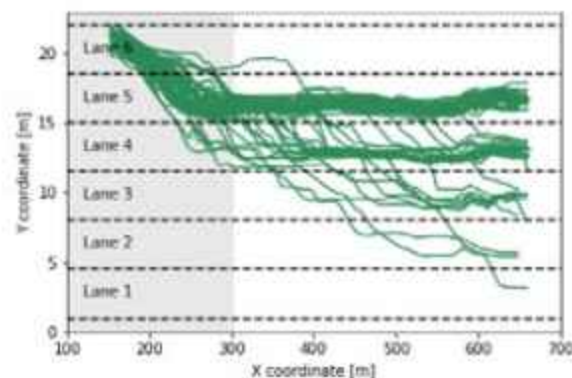
(b) Cluster $C_{2,300}$



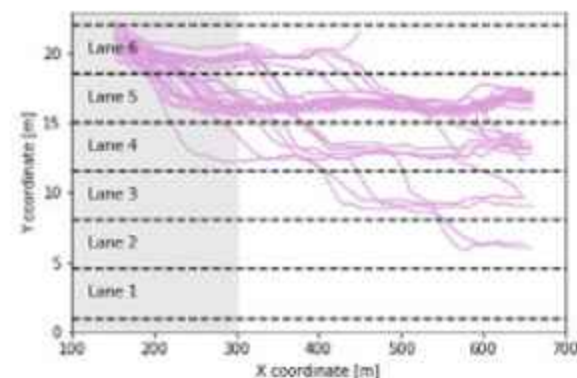
(c) Cluster $C_{3,300}$



(d) Cluster $C_{4,300}$

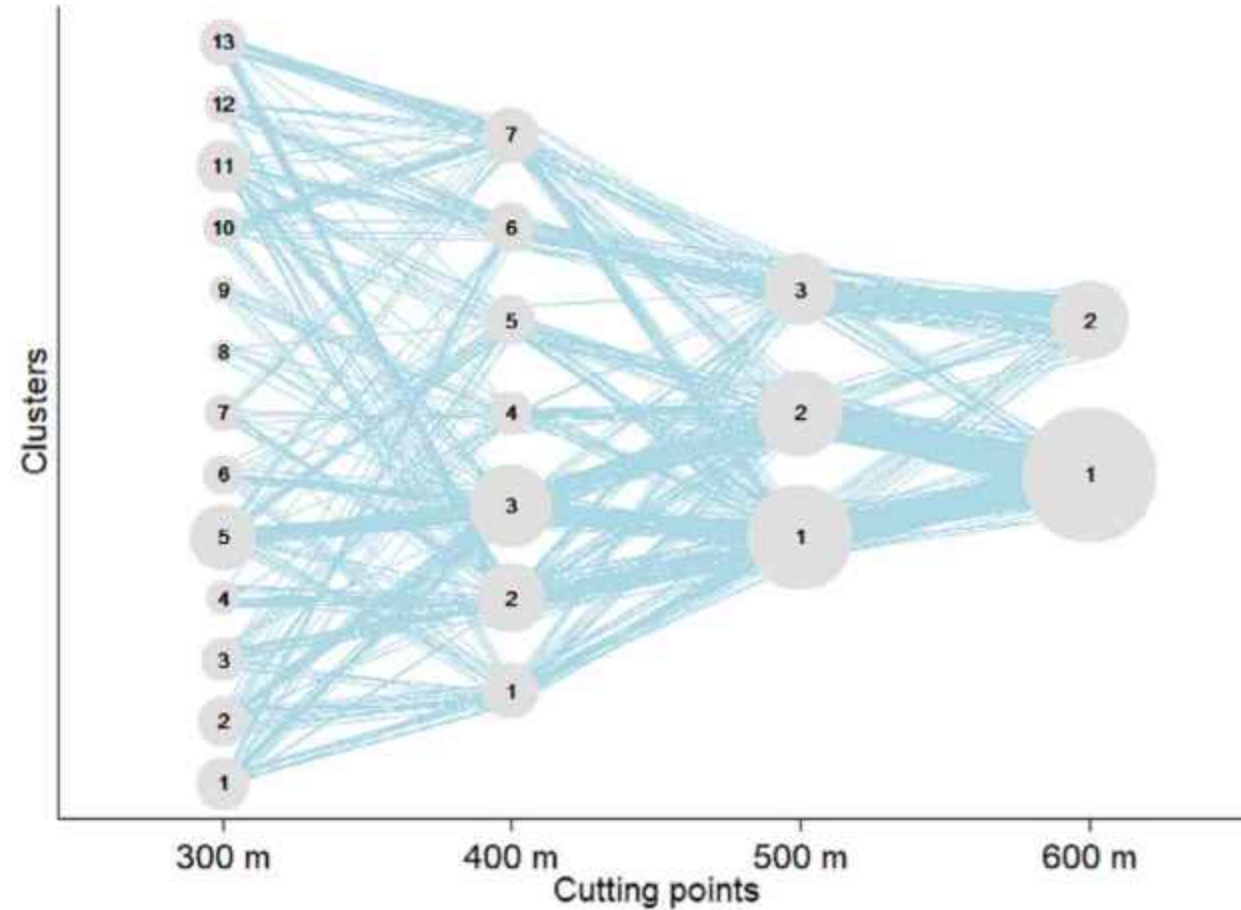


(e) Cluster $C_{5,300}$



(f) Cluster $C_{6,300}$

Dynamic clustering for driving styles analysis



Clusters transitions and driving styles switches.

AI for Sustainable Systems

Dr Fanlin Meng

f.meng2@exeter.ac.uk

Department of Management (Operations and Analytics)

University of Exeter

September 11, 2025

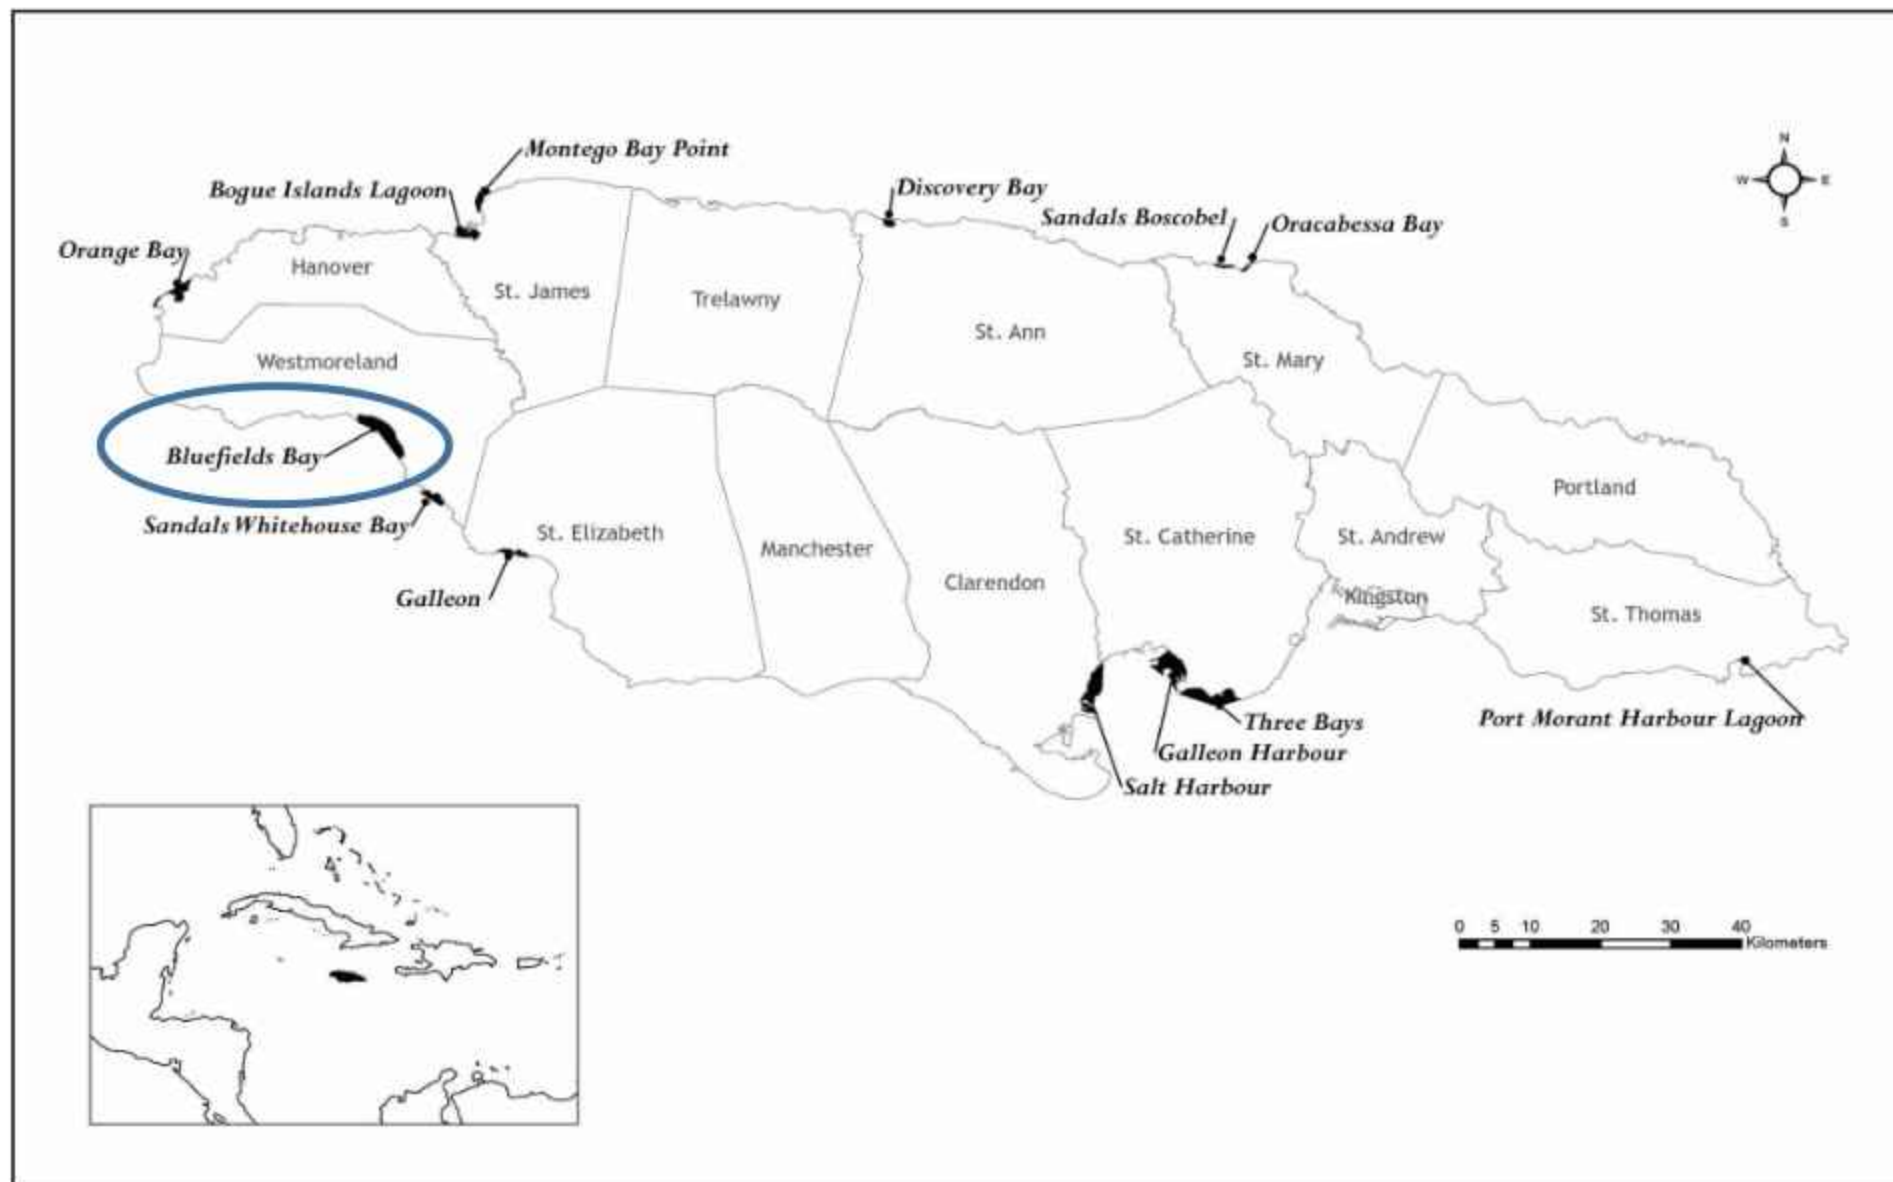
Social-Ecological networks and risk: how is perceived and ecological risk associated with network structure?

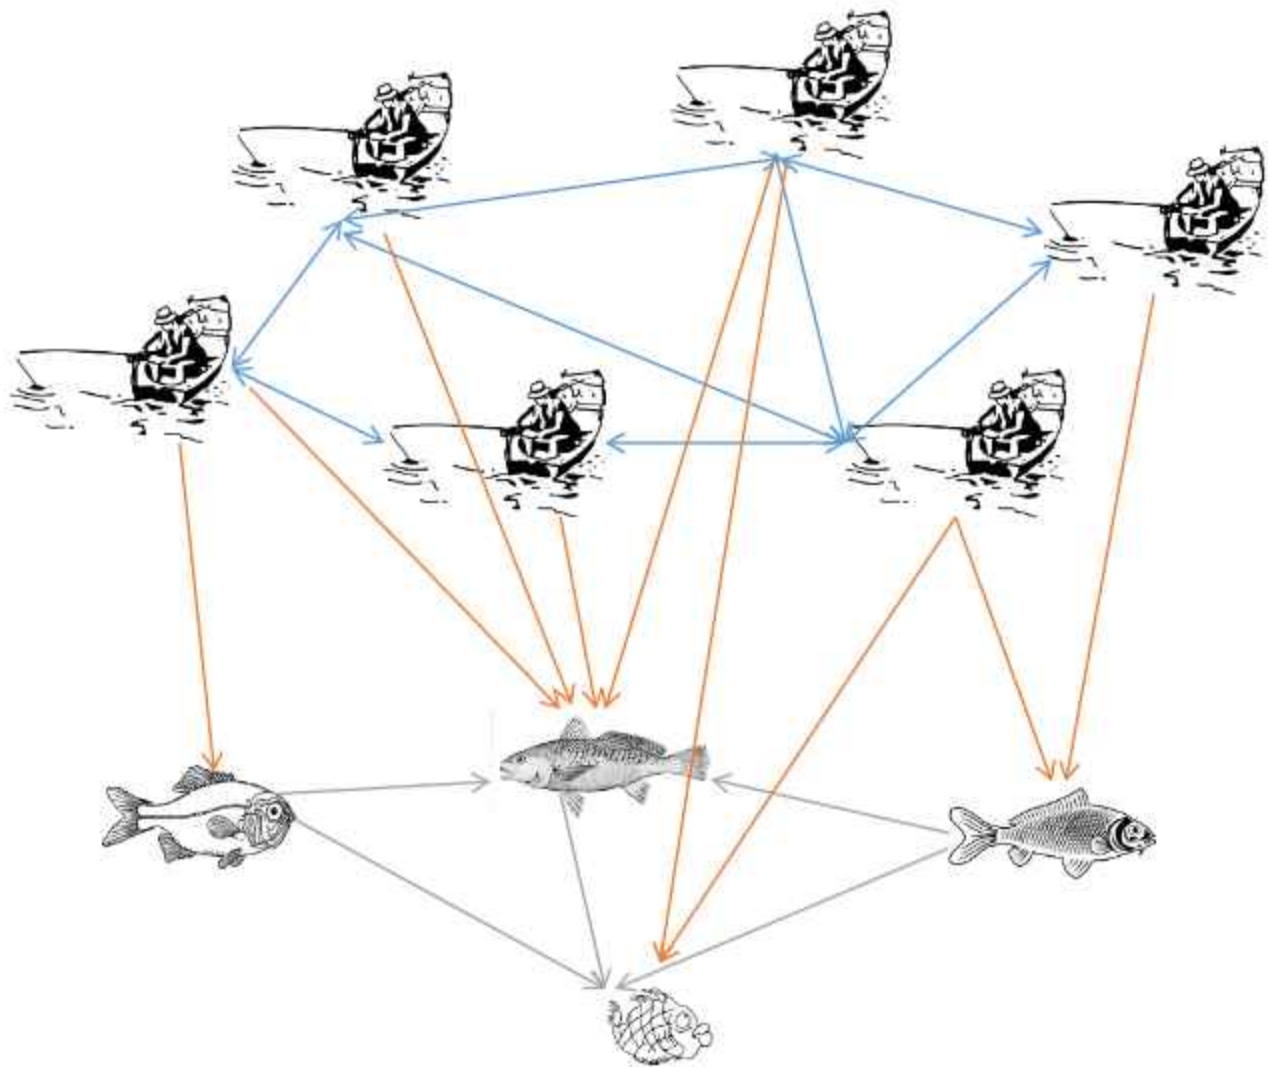
Steven Alexander, Michele Barnes, Ramiro Berardo, Laura Dee, Marie-Josée Fortin, Angela Guerrero, Kate Helmstedt, Lorien Jasny, Aislyn Keyes, Francois Massol, Spencer Wood, Laura Roldan Gomez







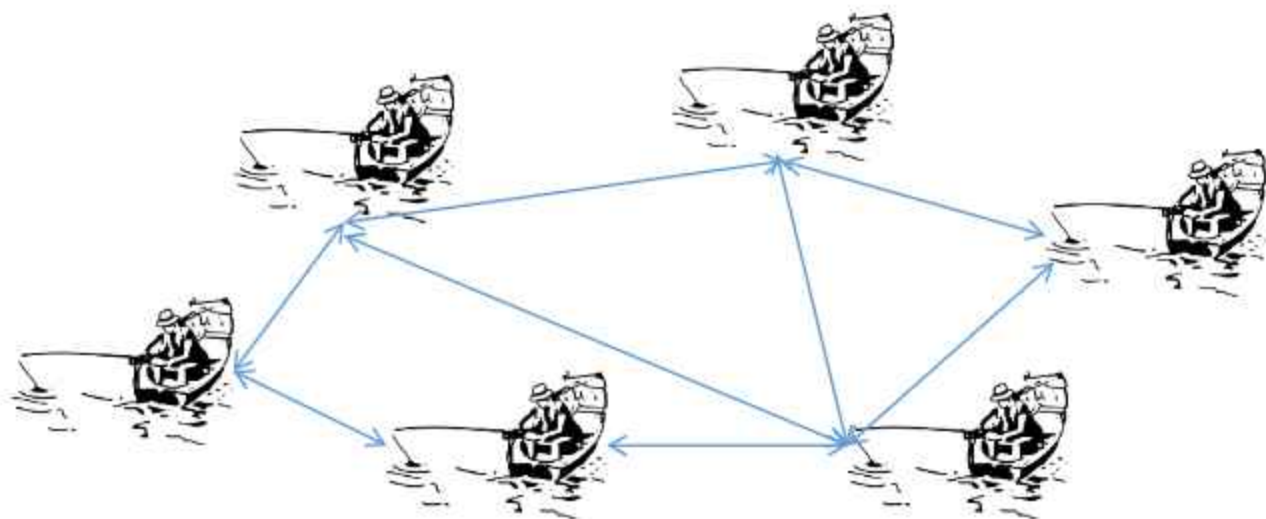




Social Network

Information Sharing Ties

- Do you exchange information with anyone, which is useful for your fishing & time at sea? (e.g., told you about practices, good fishing spots, equipment, timing and seasons, etc.)
- If yes, who have you provided information to in the last year?
- If yes, who have you received useful information from in the last year?



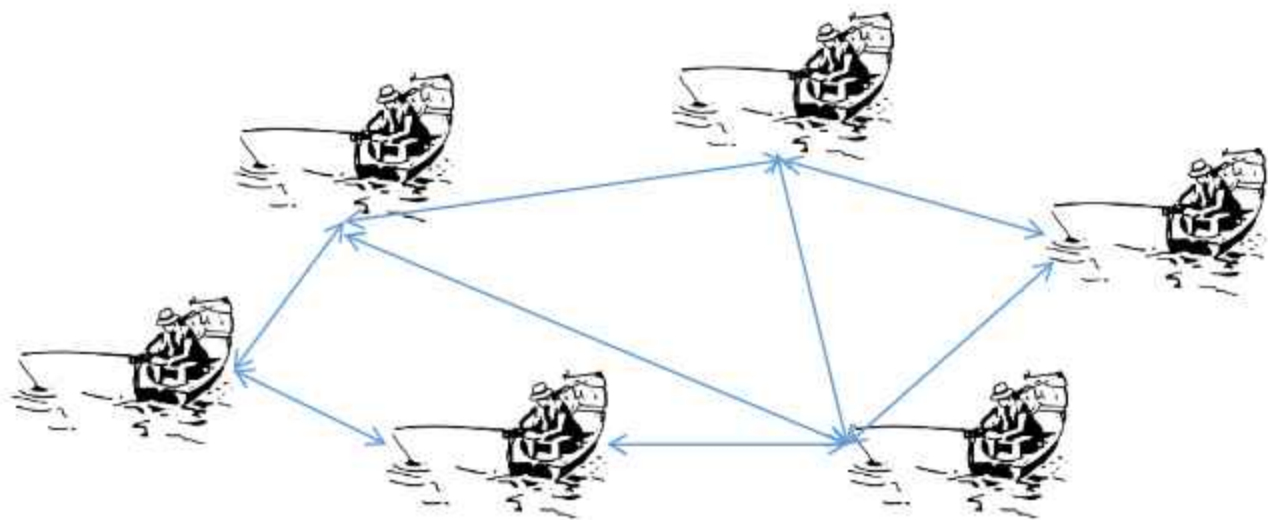
Fisherman Attributes

- Are you a member of a fisherfolk co-operative or organization?
- Perception of social risk – Since the establishment of the fish sanctuary, to what extent has illegal fishing been an issue? (5pt scale, non-issue to serious)
- Perception of ecological risk – In your opinion, how has your catch changed within the last five years? (Smaller Fish)

Social Network

Information Sharing Ties

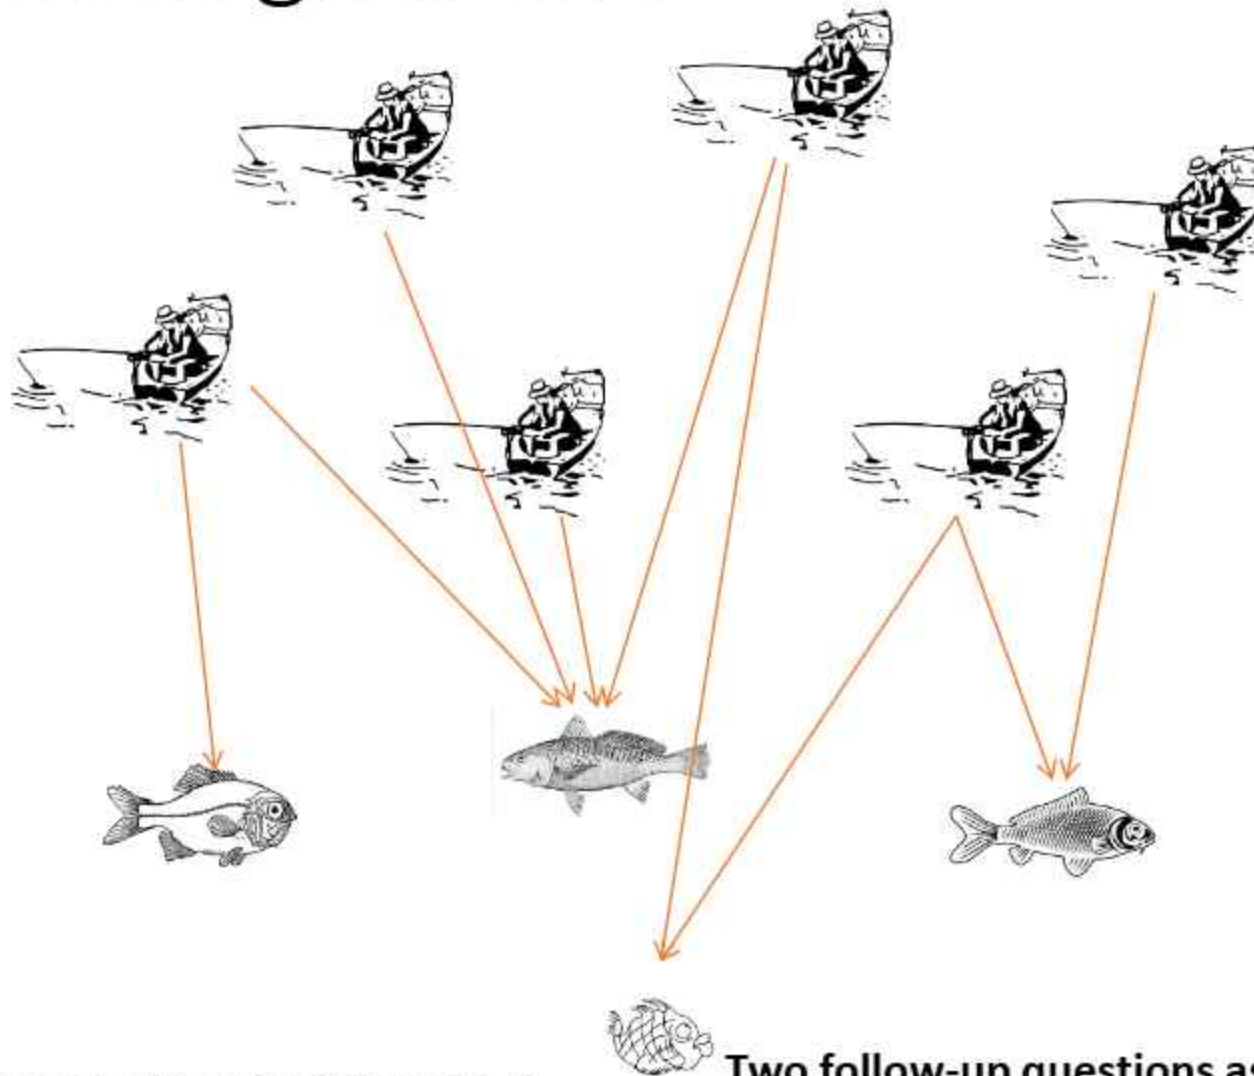
- Do you exchange information with anyone, which is useful for your fishing & time at sea? (e.g., told you about practices, good fishing spots, equipment, timing and seasons, etc.)
- If yes, who have you provided information to in the last year?
- If yes, who have you received useful information from in the last year?



Fisherman Attributes

- Are you a member of a fisherfolk co-operative or organization?
- Perception of social risk – Since the establishment of the fish sanctuary, to what extent has illegal fishing been an issue? (5pt scale, non-issue to serious)
- Perception of ecological risk – In your opinion, how has your catch changed within the last five years? (Smaller Fish)

Social-Ecological Ties



Participants were asked to identify up to 5 fish species they target.

Two follow-up questions asked participants specifically whether they target lobster and/or conch.

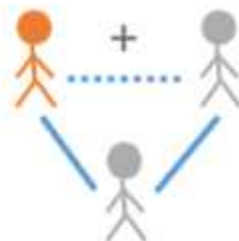
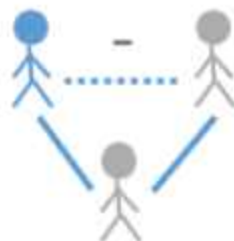
Risk Hypothesis

- Berardo and Scholz (2010)
- 'Risk' means the social risk of individuals/organizations defecting on their commitments
- Low-risk situations favour the creation of bridging structures
- High-risk favour reciprocity and bonding structures

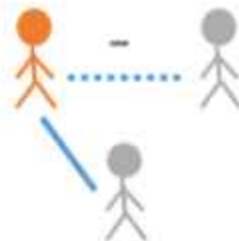
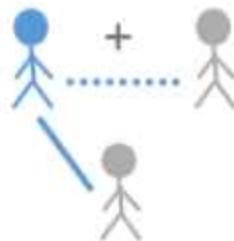
Risk Hypothesis -Social motifs

Individual responses to social risk

Social bonding
(triangle)



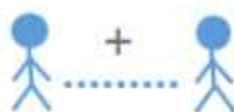
Social bridging
(open two-path)



Activity



Homophily



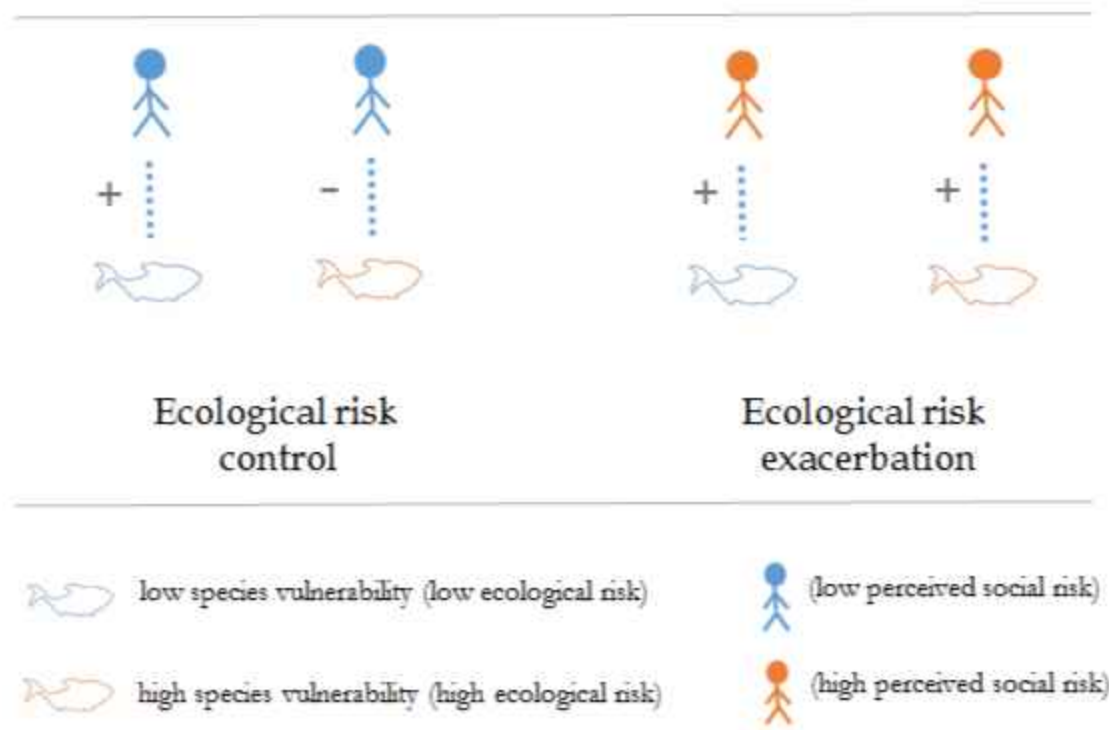
(low perceived social risk)



(high perceived social risk)

Risk Hypothesis – Socio-Ecological motifs

Individual responses to ecological risk



Risk Hypothesis – Socio-Ecological motifs

Individual responses to social-ecological risk

Risk control



Risk “spreading”



low species vulnerability (low ecological risk)



high species vulnerability (high ecological risk)



(low perceived social risk)



(high perceived social risk)

ERG Models!

ERG Models!

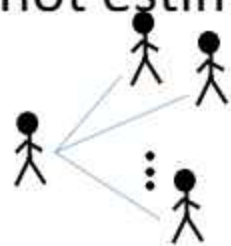
- Control terms:
 - Number of isolates (constrained, not estimated)

ERG Models!

- Control terms:
 - Number of isolates (constrained, not estimated)
 - Edges

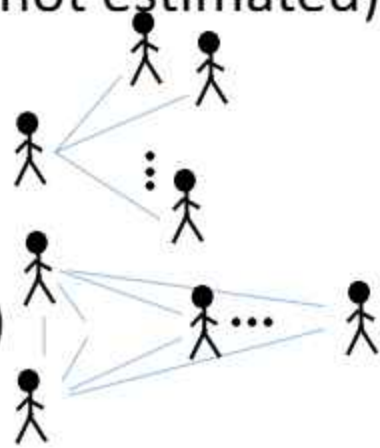
ERG Models!

- Control terms:
 - Number of isolates (constrained, not estimated)
 - Edges
 - Weighted popularity (gwdegree)



ERG Models!

- Control terms:
 - Number of isolates (constrained, not estimated)
 - Edges
 - Weighted popularity (gwdegree)
 - Weighted social triangles (gwesp)



ERG Models!

- Control terms:

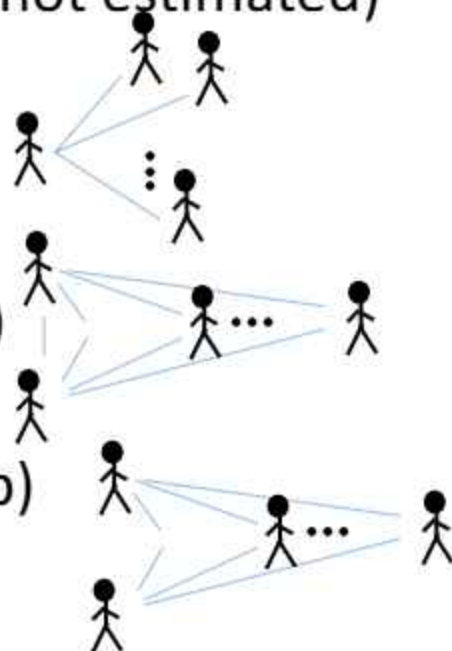
- Number of isolates (constrained, not estimated)

- Edges

- Weighted popularity (gwdegree)

- Weighted social triangles (gwesp)

- Weighted open two-paths (gwnsp)



ERG Models!

- Control terms:

- Number of isolates (constrained, not estimated)

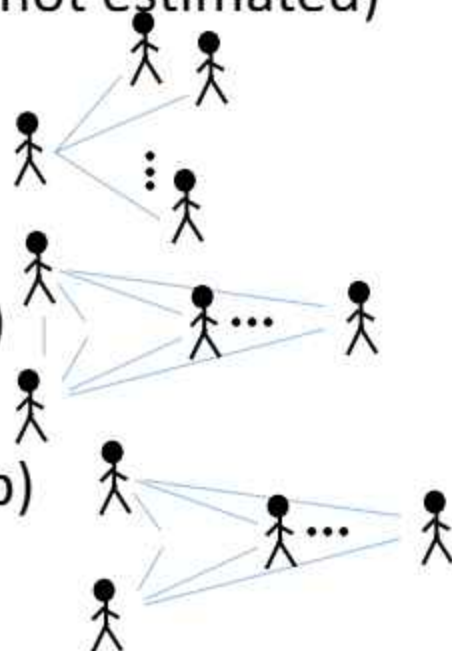
- Edges

- Weighted popularity (gwdegree)

- Weighted social triangles (gwesp)

- Weighted open two-paths (gwnsp)



- Member of a fishery org



ERG Models!

- Risk: (modelled separately for Social and Ecological risk)

ERG Models!


- Risk: (modelled separately for Social and Ecological risk)
 - Number of ties  — 


ERG Models!

- Risk: (modelled separately for Social and Ecological risk)
 - Number of ties 
 - Homophily 

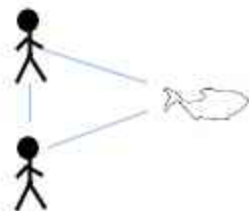
ERG Models!

- Risk: (modelled separately for Social and Ecological risk)

- Number of ties 


- Homophily 


- S-E Triangles



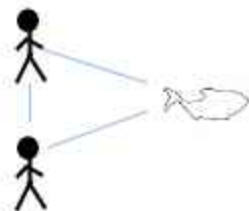
ERG Models!

- Risk: (modelled separately for Social and Ecological risk)

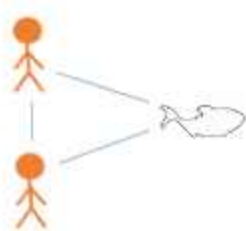
- Number of ties 

- Homophily 

- S-E Triangles





- S-E Triangles (risk homophily)

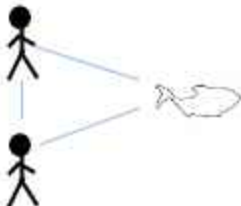


ERG Models!

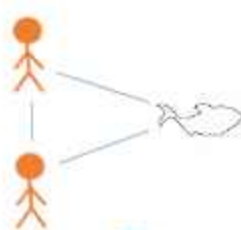
- Risk: (modelled separately for Social and Ecological risk)

- Number of ties 

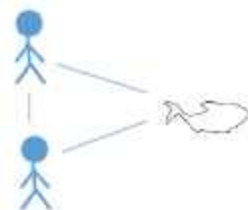
- Homophily 

- S-E Triangles 

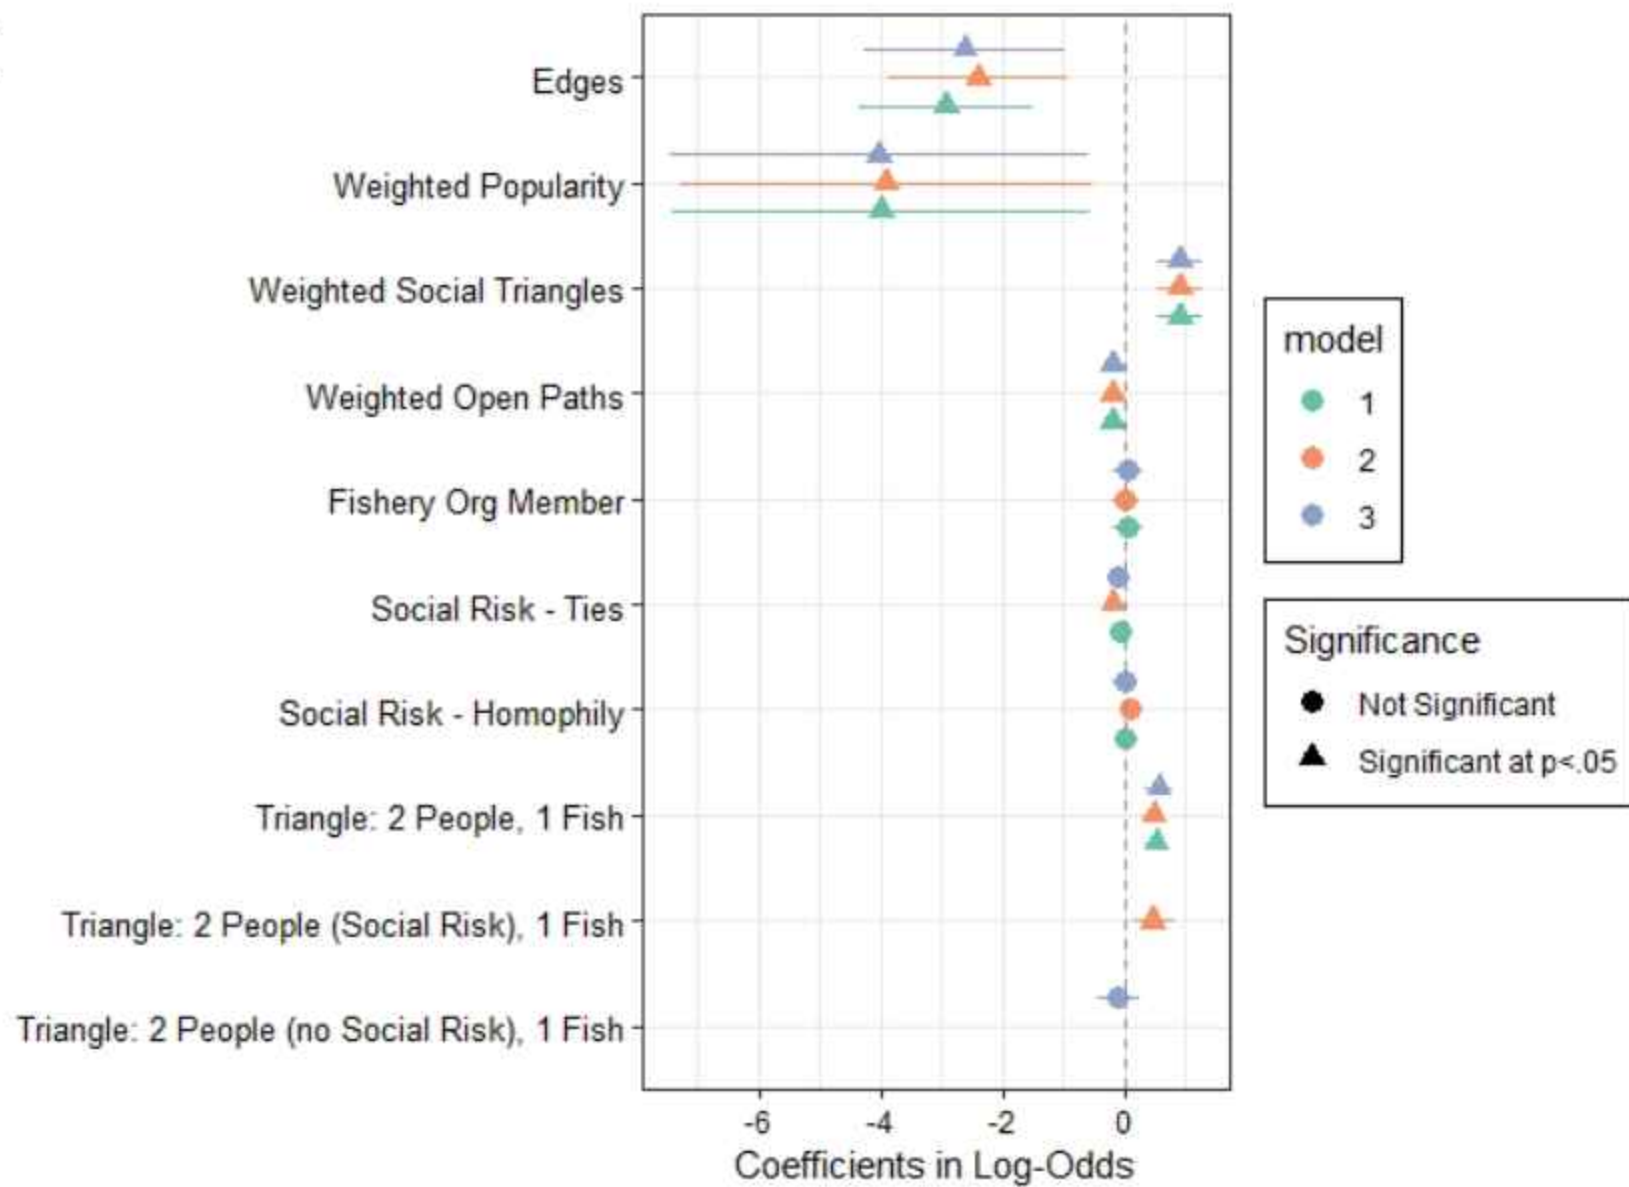
- S-E Triangles (risk homophily)



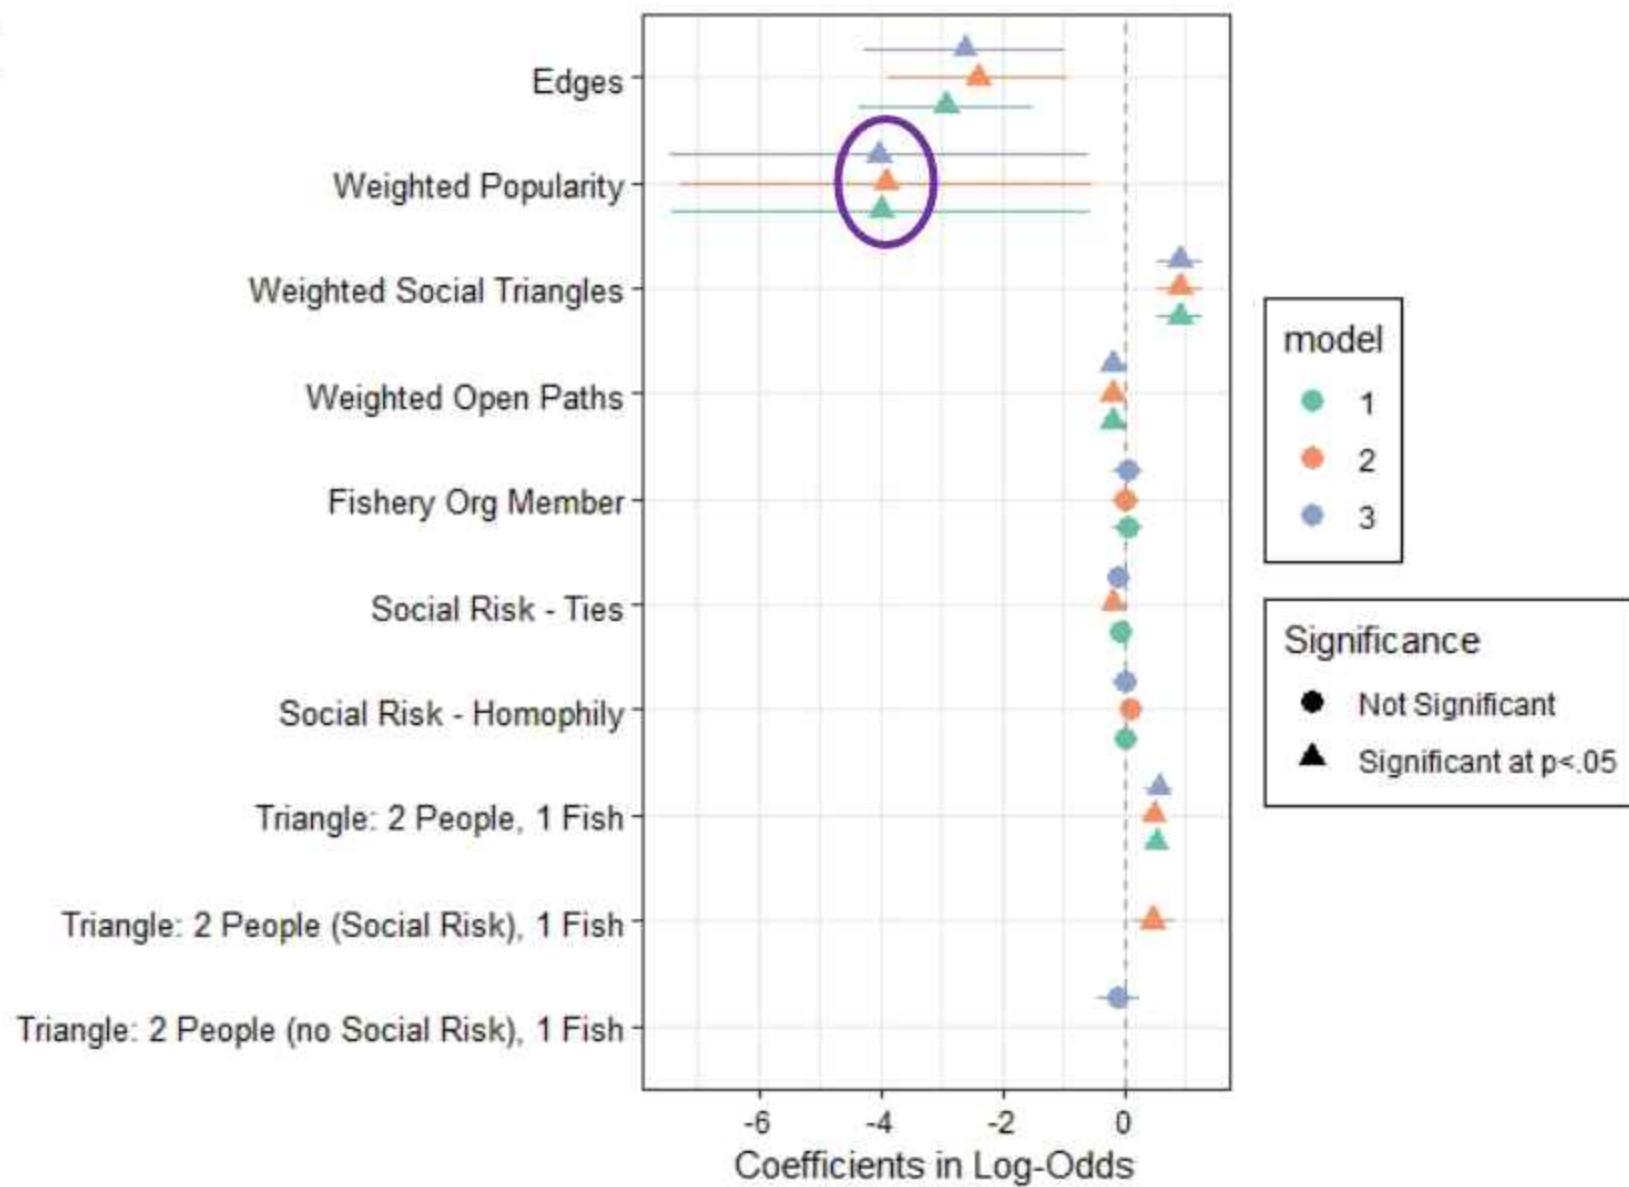
- S-E Triangles (no risk homophily)



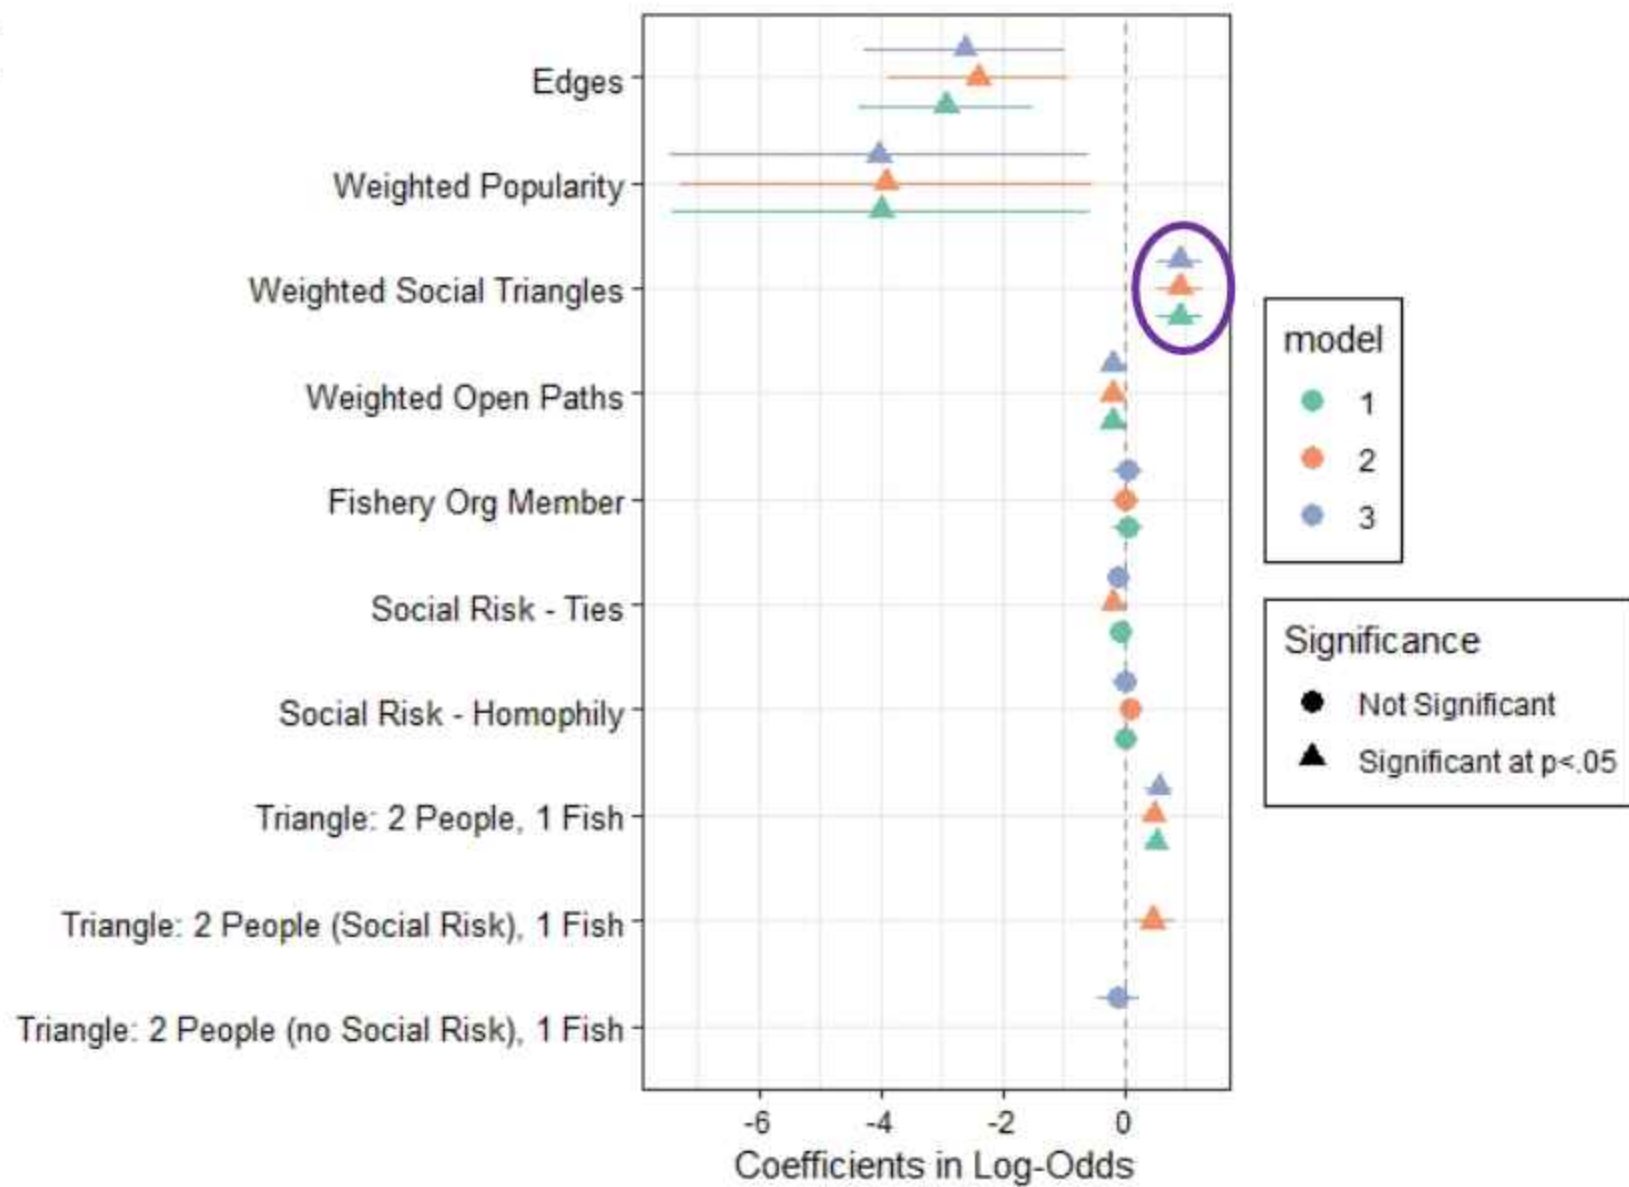
Social Risk



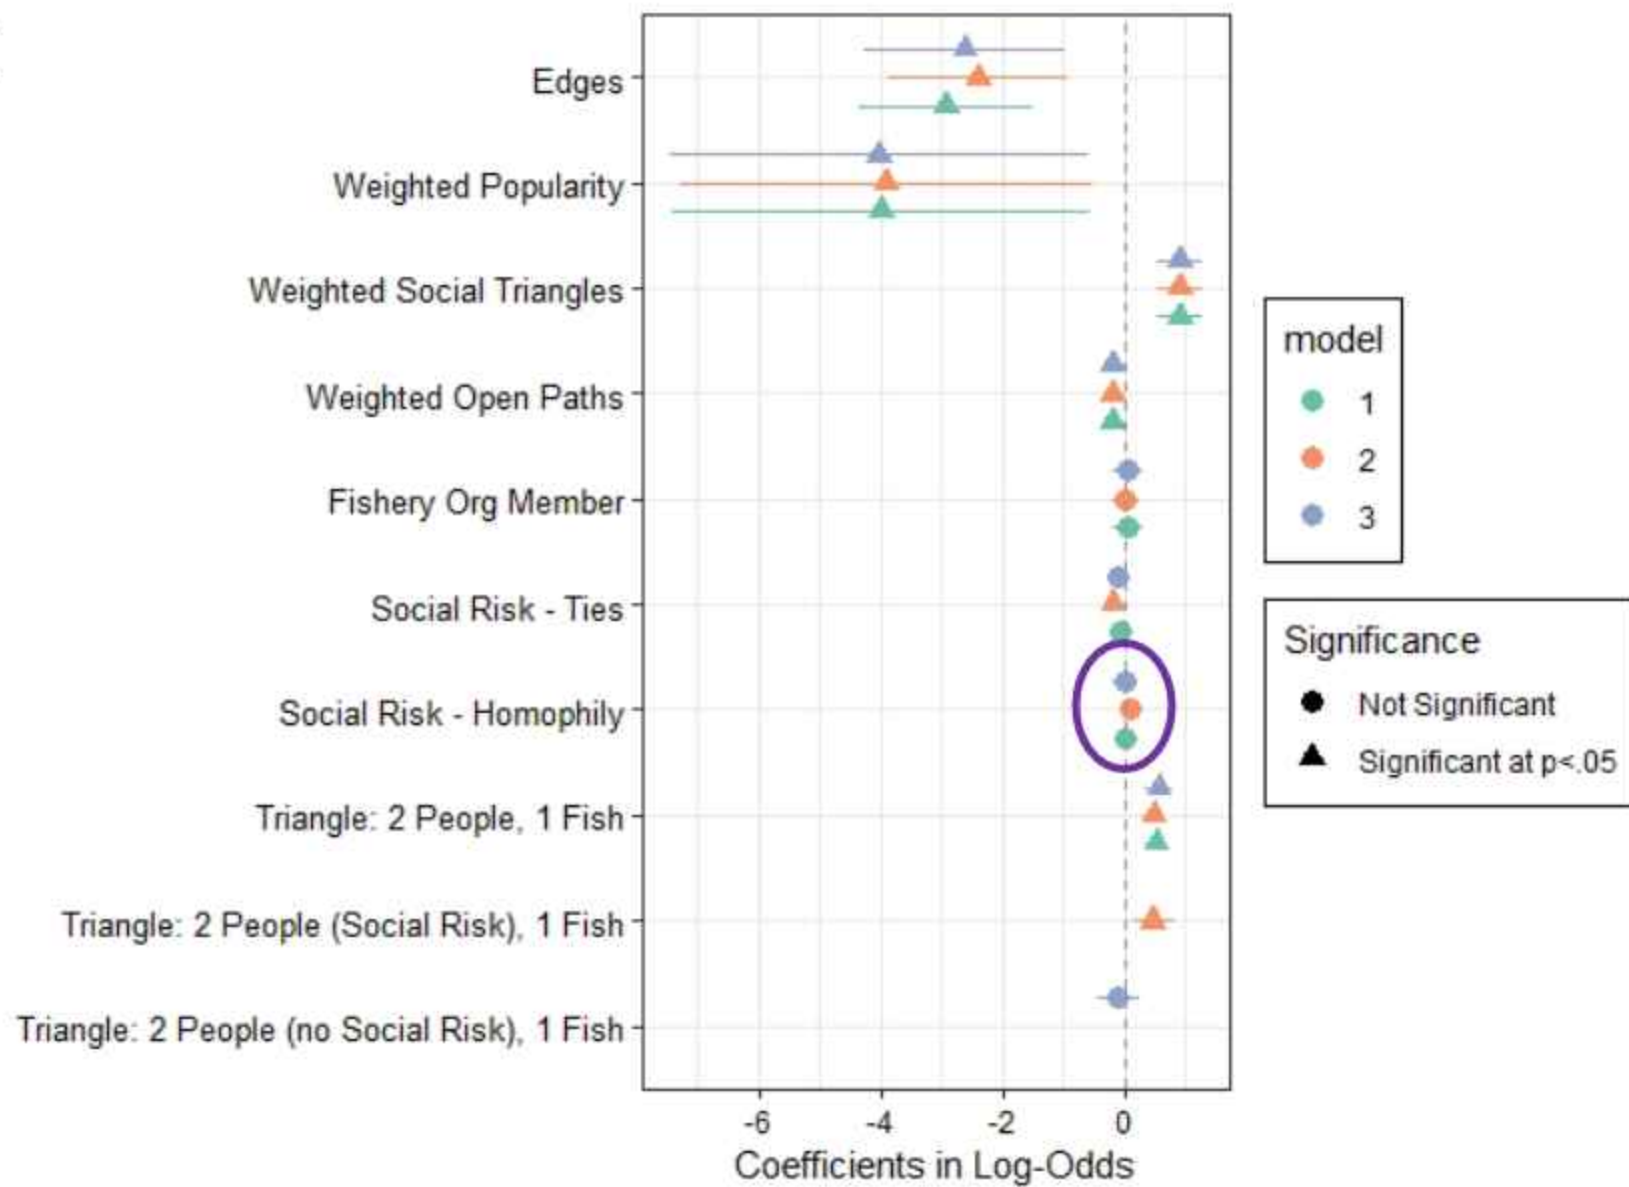
Social Risk



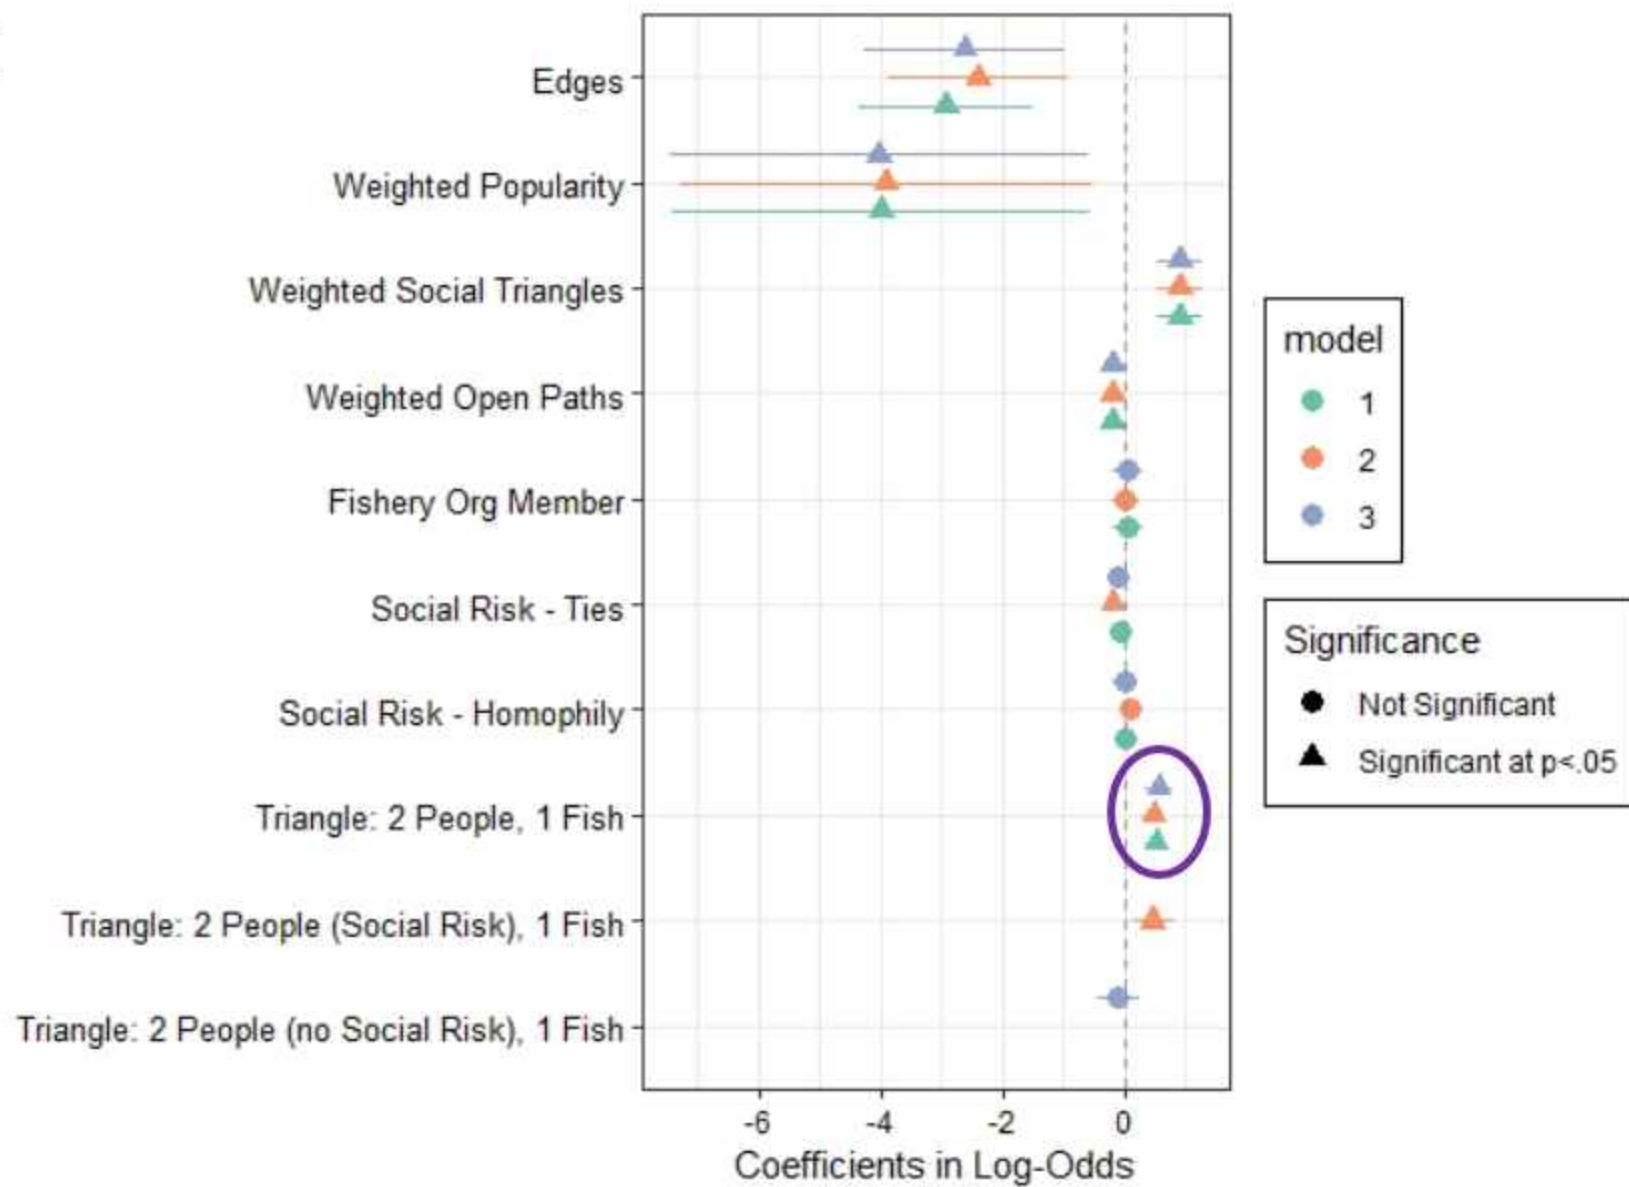
Social Risk



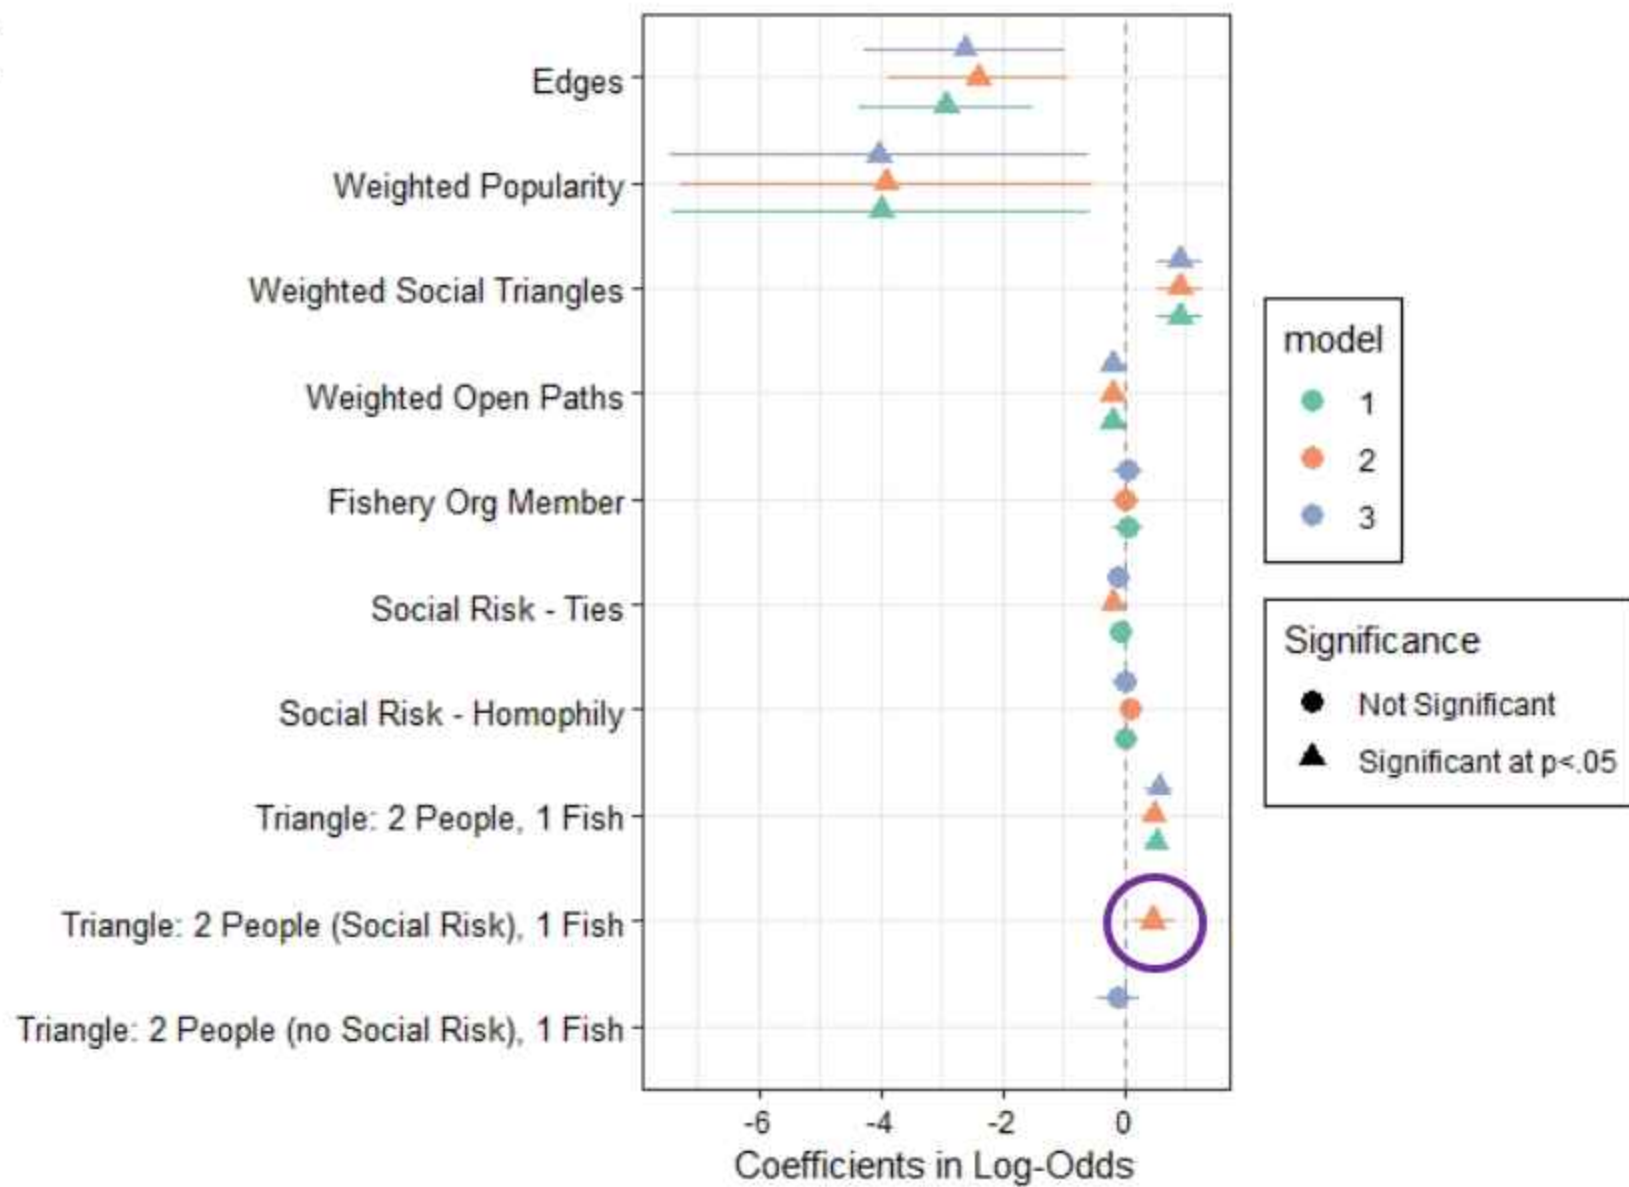
Social Risk



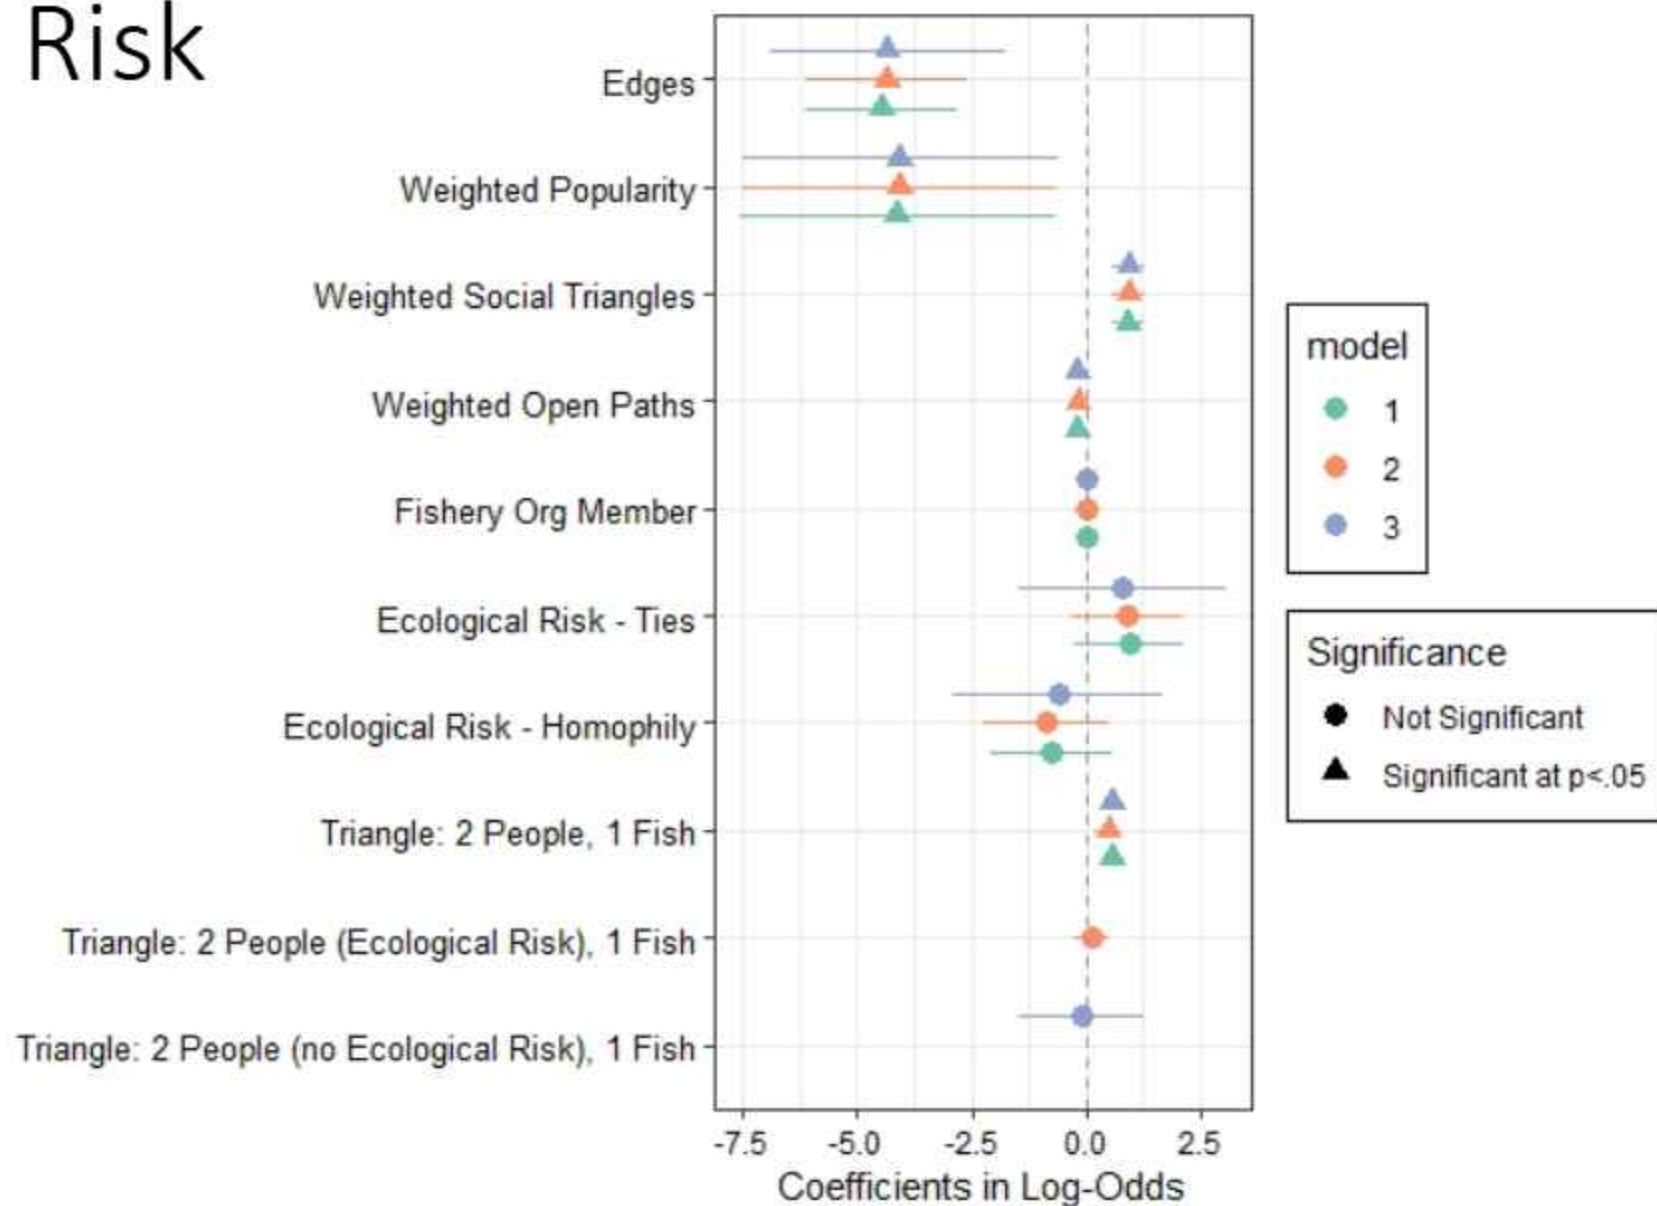
Social Risk



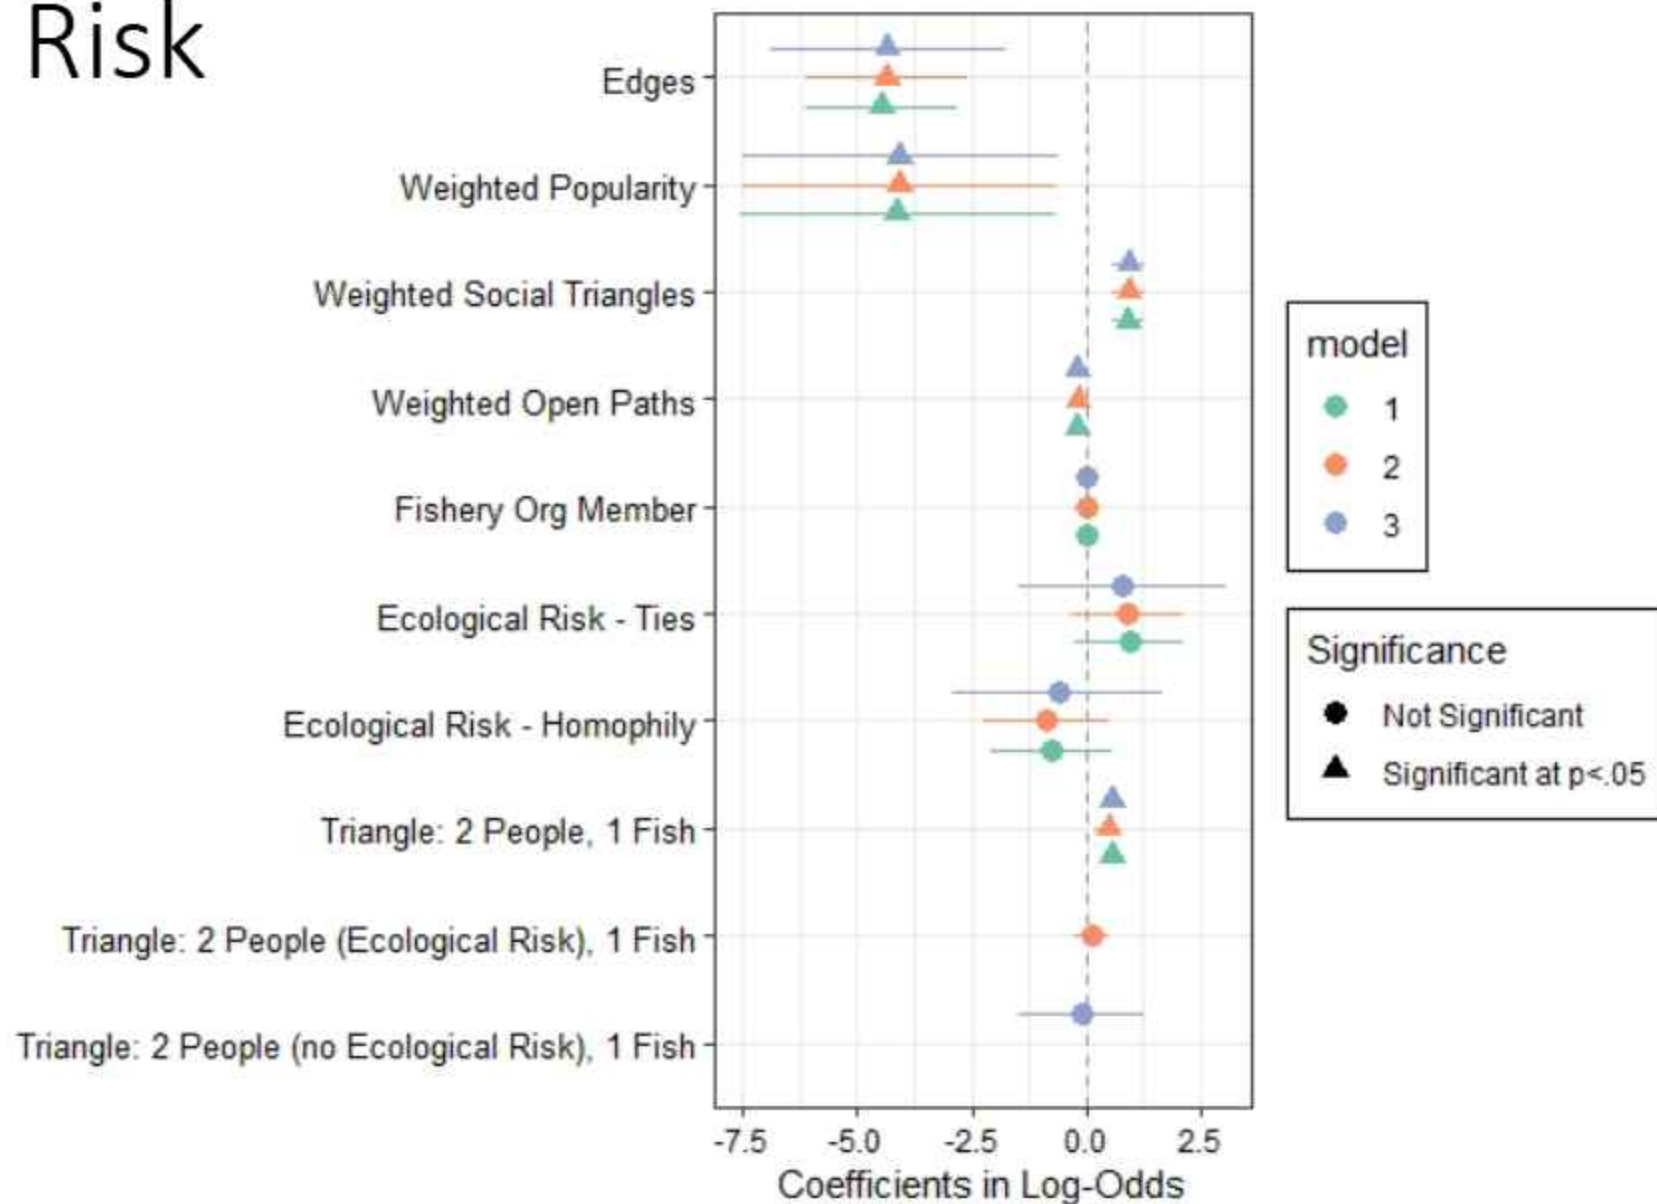
Social Risk



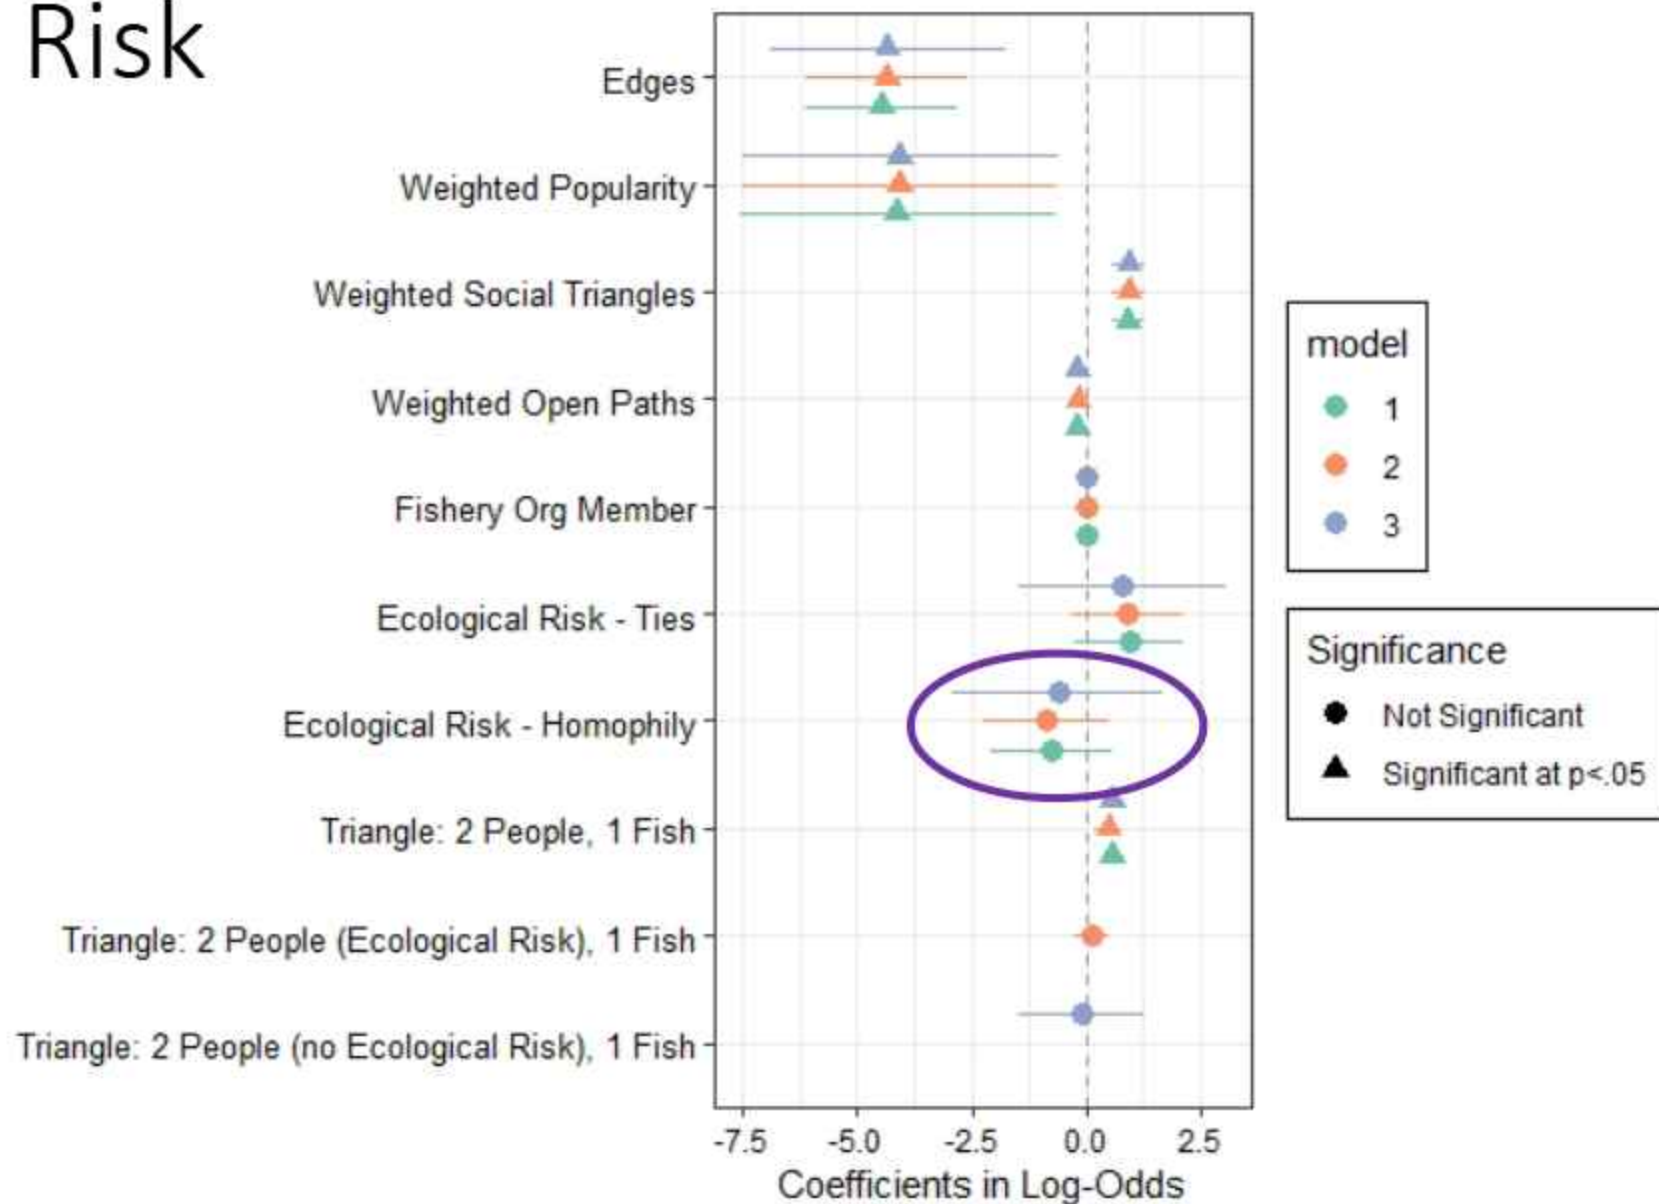
Ecological Risk



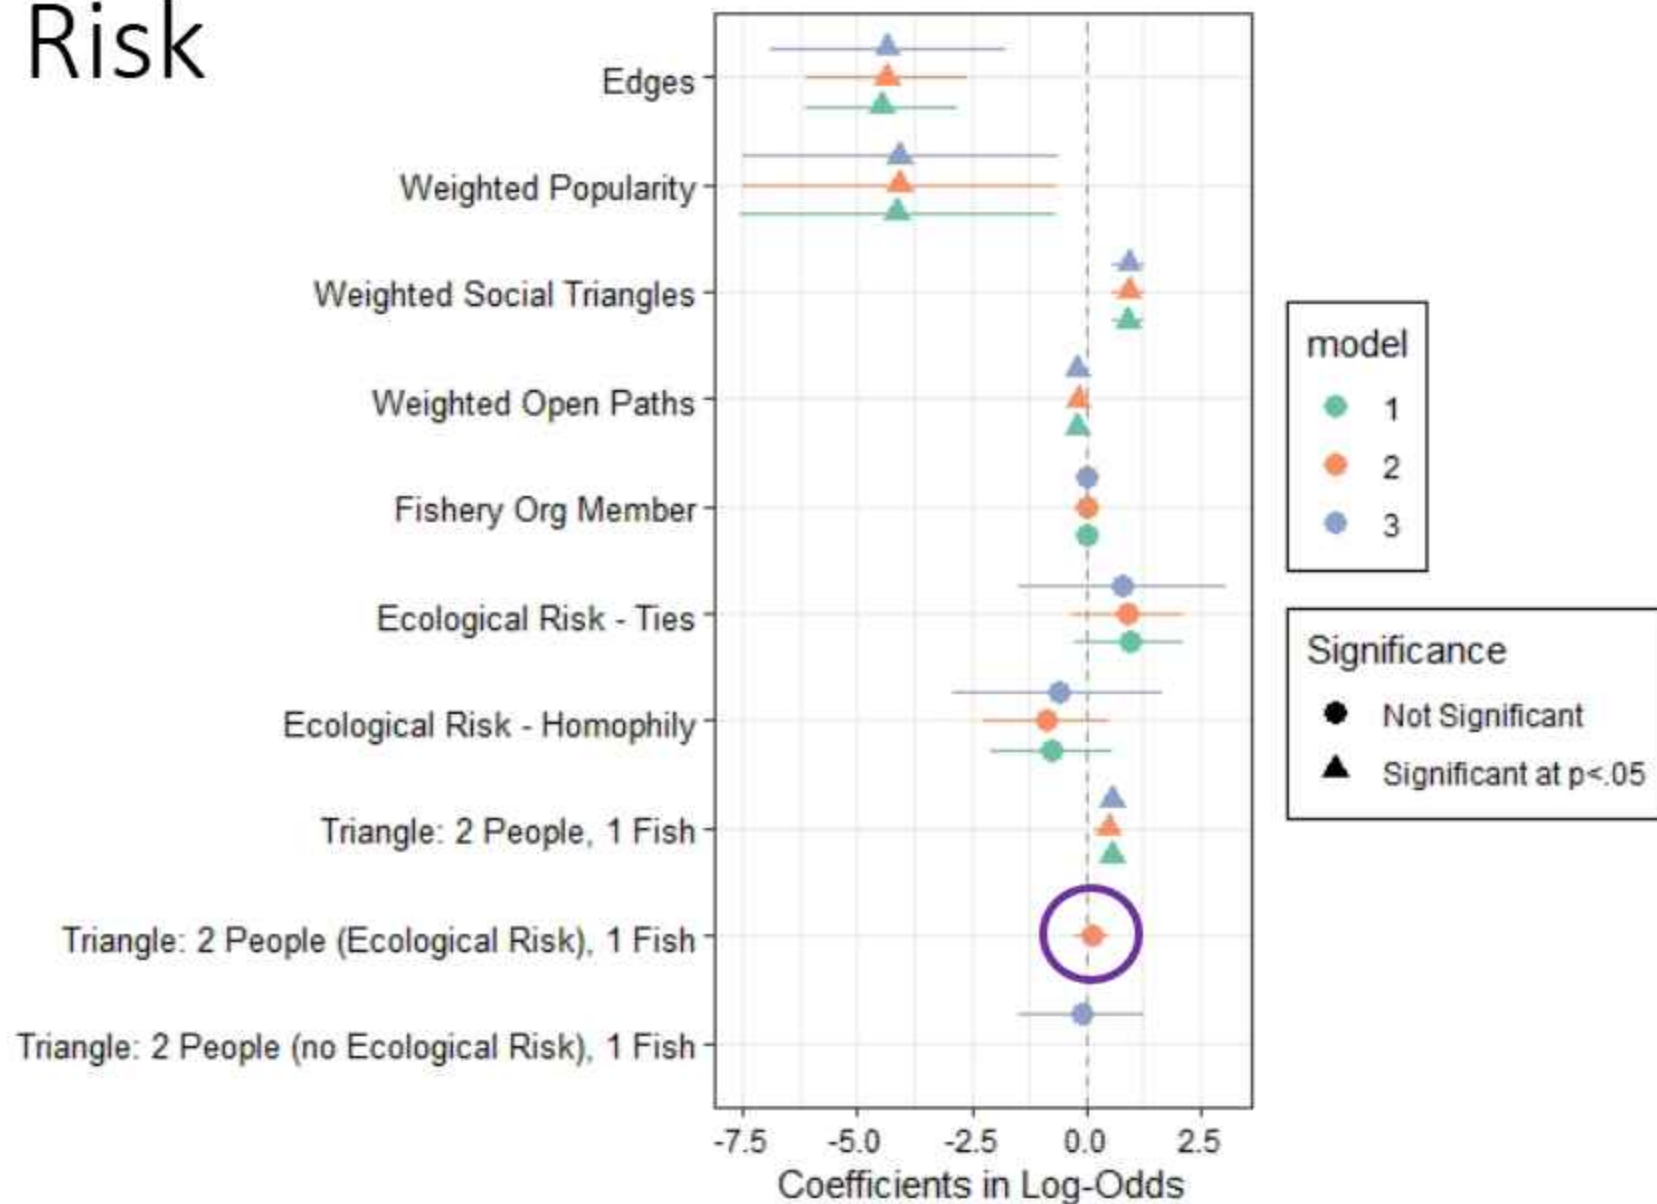
Ecological Risk



Ecological Risk



Ecological Risk



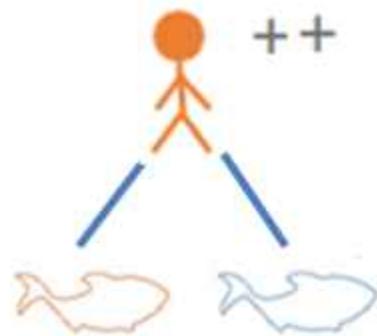
Risk Hypothesis?

- Limited support for perceptions of social risk (yes triangles, no homophily)
- No support for perceptions of ecological risk

Next steps

- Change/refine motifs
- Add in more attribute data controls
- Add Kenya and Papua New Guinea cases
- Full models of socio-ecological network

Social-Ecological Models



Social-Ecological networks and risk: how is perceived and ecological risk associated with network structure?

Steven Alexander, Michele Barnes, Ramiro Berardo, Laura Dee, Marie-Josée Fortin, Angela Guerrero, Kate Helmstedt, Lorien Jasny, Aislyn Keyes, Francois Massol, Spencer Wood, Laura Roldan Gomez



Board Gender Diversity and Carbon Emissions: A Comparative Study Using Panel Regression and Machine Learning Approaches

**Mohammad Hassan Shakil^{1&2}, Arne Johan Pollestad¹, Khine Kyaw¹
& Ziaul Haque Munim³**

¹NTNU Business School, Norwegian University of Science and Technology,
Trondheim, Norway.

²Smith School of Enterprise and the Environment, University of Oxford, United
Kingdom.

³Faculty of Technology, Natural and Maritime Sciences, University of South-Eastern
Norway, Horten, Norway.



NTNU

Norwegian University of
Science and Technology



Background and motivation

- Board characteristics are widely recognized as key determinants of firms' financial and non-financial performance
- Board diversity enhances firms' decision-making and risk-management capabilities, which helps in improving firms' financial and non-financial outcomes (Jizi & Nehme, 2017; Reguera-Alvarado & Bravo-Urquiza, 2020).
- Previous studies explore the impact of board attributes on firm performance (Bennouri et al., 2018), ESG disclosure (Nguyen & Nguyen, 2023), carbon emissions (Barroso et al., 2024; Khatri, 2024) and applied traditional regression models
- Few studies investigate the board diversity-carbon emission nexus, there is little evidence on how machine learning models can explain the complicated diversity-carbon emissions relationship.

Research questions

(1) Does board gender diversity influence carbon emissions performance?

*(2) Does **ESG controversies** moderate the relationship between board gender diversity and carbon emissions performance?*

Theory

(1) Stakeholders Theory (Freeman, 1984)

(2) Legitimacy Theory (Suchman, 1995)

Research Methodology

Sample selection

- STOXX 600 firms in Europe from 2016 to 2022

Data sources

- Carbon emissions score of LSEG Workspace formerly known as Refinitiv (Target variable/Dependent variable)
 - Board gender diversity is as measured as the percentage of female members on the board
 - Board specific control variables include board member compensation, board size, and board member tenure, CEO duality
 - Firm-specific control variables include Tobin's Q, market risk, leverage, liquidity, cost of debt, log of total assets.
-

Research Methodology

Target variable

Carbon emissions score	The emissions score of LSEG indicates the percentile rank score of a firms' commitment and effectiveness towards reducing carbon emission in the production and operational processes relative to its industry. The higher the score, the more efforts are made to reduce carbon emissions.
-------------------------------	---

Independent variables

Board gender diversity	Percentage of female members on firm board.
-------------------------------	---

Moderating variable

ESG controversies	Environmental, social and governance controversies score of LSEG between zero to hundred. Higher score means low controversies, lower score means high controversies.
--------------------------	---

Board specific control variables

Board member compensation	Total compensation of board members.
Board size	Total number of board members at fiscal year-end
Board tenure	Average number of years each board member has been on the board
CEO duality	CEO serves as the chairman or the chairman has been the CEO (coded 1 if yes, or 0 otherwise).

Firm specific control variables

Tobin's Q	"The book value of total assets minus the book value of equity and balance sheet deferred taxes plus the market value of equity, all divided by the book value of total assets" (Aouadi & Marsat, 2018 p. 1031).
Market risk	CAPM beta is used as proxy for market risk.
Leverage	Total Debt/Total equity
Liquidity	Quick ratio
Cost of debt	Weighted average cost of debt
Total assets	Natural logarithm of total assets

Research Methodology

-
- Fixed effects and random effects panel regression
 - Stochastic XGBoost (XGB) algorithm, an advanced ensemble-based machine-learning technique to explore whether board gender diversity helps predict firms' carbon emissions performance.
 - PDP to examine the average marginal predictability of board gender diversity on carbon emissions score.
-



Results (Panel regression model)



Table 4: Regression results of fixed and random effects

Dependent/Target variable	Carbon emissions performance	Carbon emissions performance	Carbon emission performance	Carbon emissions performance
Independent variable				
Board gender diversity	.18*** (.06)	.24*** (.05)	.24** (.11)	.27*** (.1)
Control variables				
Board compensation	.66 (.87)	.47 (.78)	.67 (.88)	.47 (.79)
Board size	.31 (.38)	.43 (.27)	.32 (.38)	.43 (.27)
Board tenure	-.68 (1.63)	-.35 (.65)	-.68 (1.63)	-.35 (.65)
CEO duality	4.23* (2.41)	4.32** (1.81)	4.29* (2.41)	4.35** (1.8)
Tobin's Q	.83 (.72)	.87* (.5)	.82 (.72)	.86* (.5)
Market risk	3.27** (1.47)	3.72*** (1.27)	3.28** (1.47)	3.73*** (1.27)
Leverage	-2.16** (.89)	-2.01** (.78)	-2.18** (.89)	-2.02** (.79)
Liquidity	-1.71* (.99)	-1.34 (.93)	-1.71* (.99)	-1.34 (.93)
Cost of debt	-13.99 (28.25)	-9.17 (26.74)	-13.91 (28.47)	-9.98 (26.88)
Total assets	28.19*** (4.52)	18.65*** (1.71)	28.17*** (4.51)	18.51*** (1.78)
ESG controversies			.02 (.04)	.01 (.04)
Moderating variable				
Board gender diversity*ESG controversies			-.001 (.001)	-.0003 (.001)
Constant	-230.3*** (47.79)	-136.55*** (18.03)	-232.24*** (48.03)	-135.73*** (18.98)
Model statistics				
Observations	1782	1782	1782	1782
R ²	0.1385	0.2631	0.1388	0.2631
Sargan-Hansen statistic		17.521		20.045
Chi ²		11		13
Hausman test (p-value)		0.0934		0.0941
Fixed effect	Yes	No	Yes	No
Random effect	No	Yes	No	Yes

Note: Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Results
(Machine
Learning)



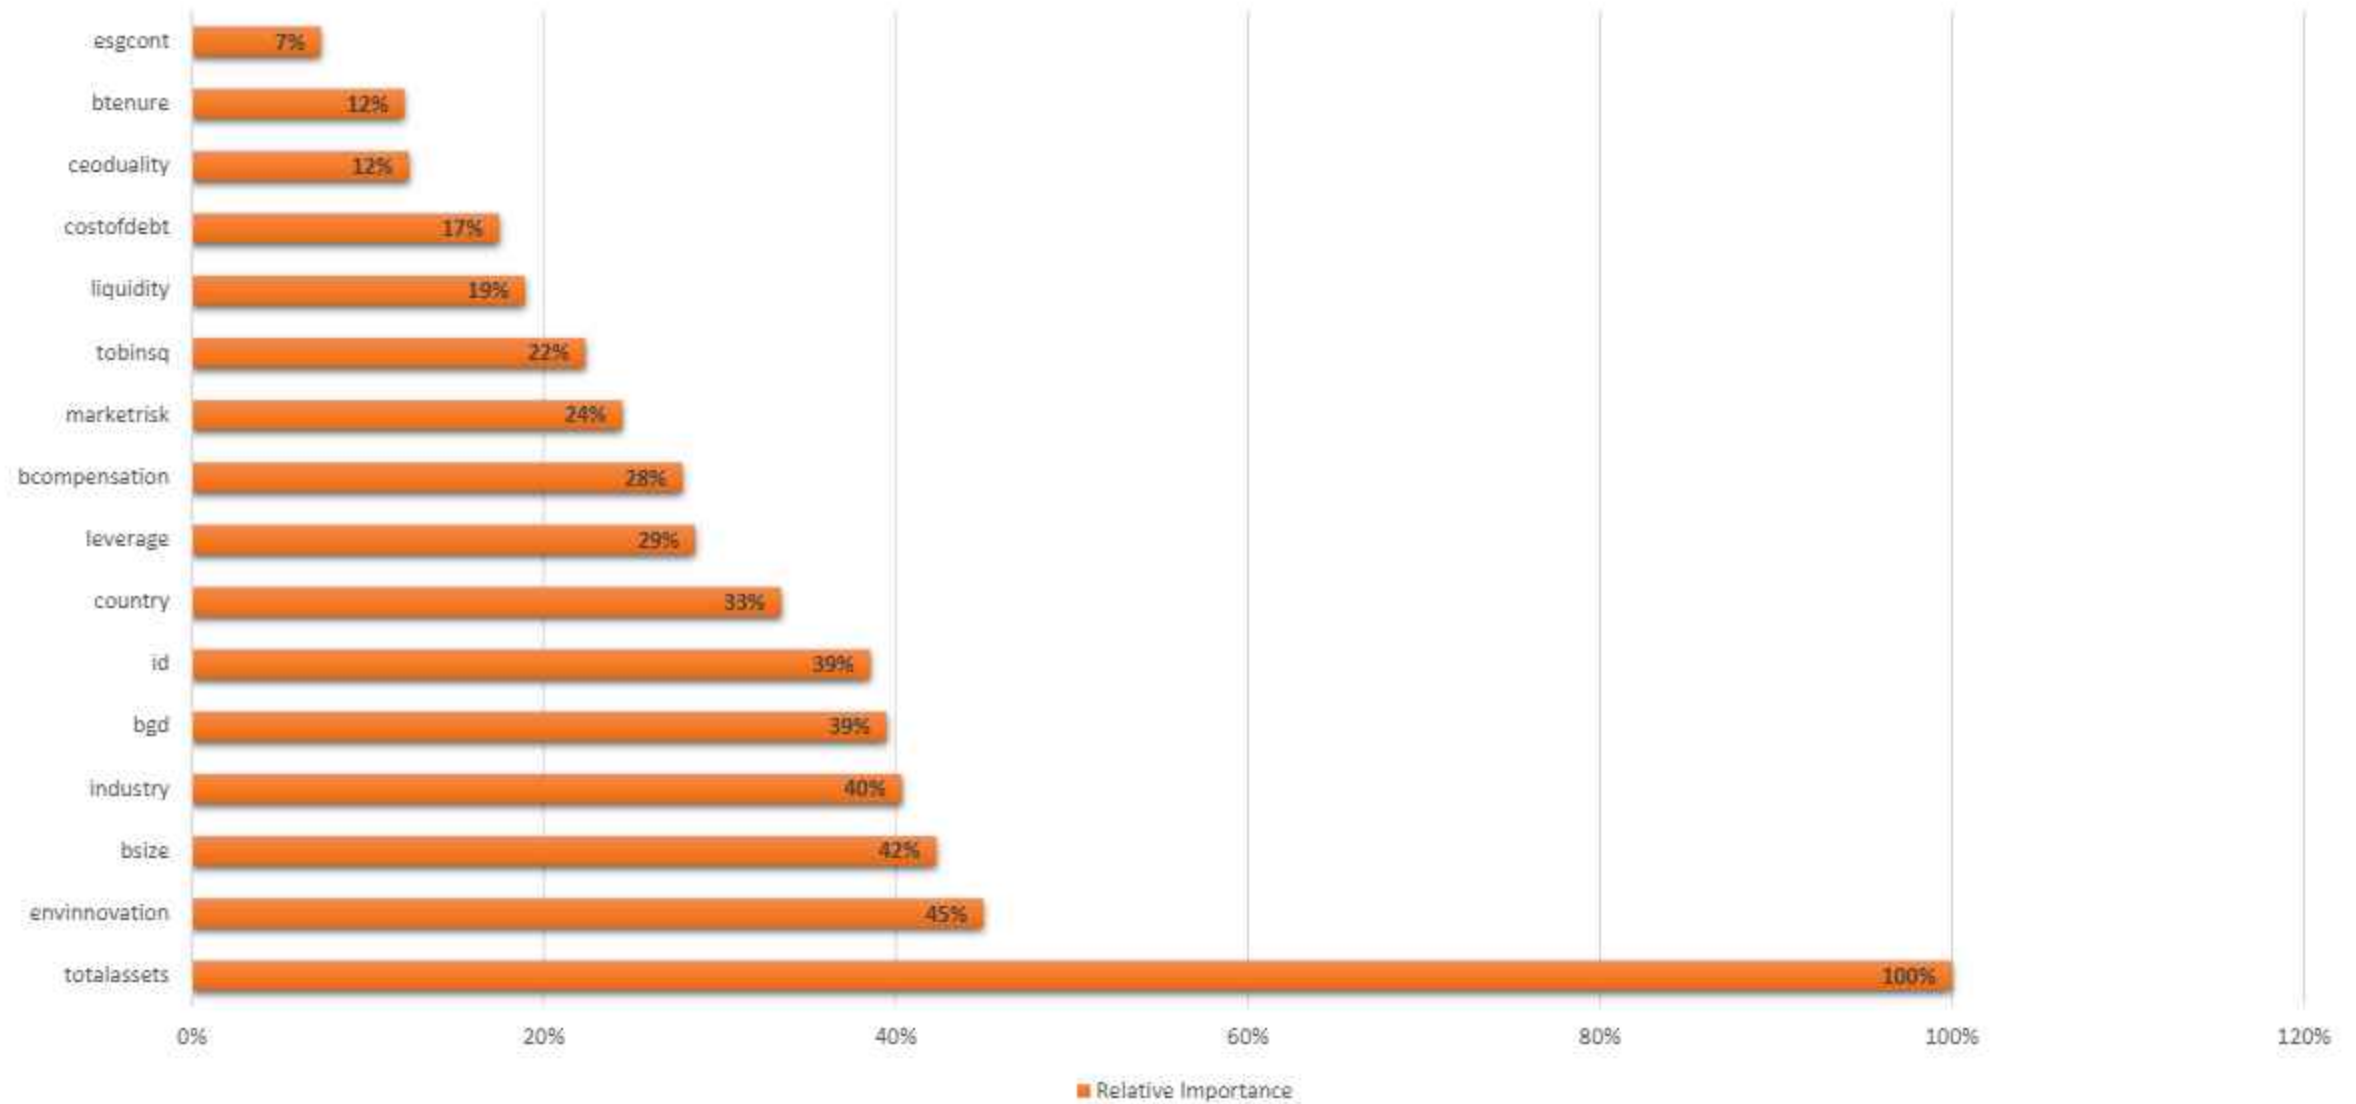
Table 7: Carbon emissions predictions for 2022 (Holdout)

Model	Training obs.	MAE	RMSLE	R^2
XGBoost	2,546	9.5726	0.2233	0.5751

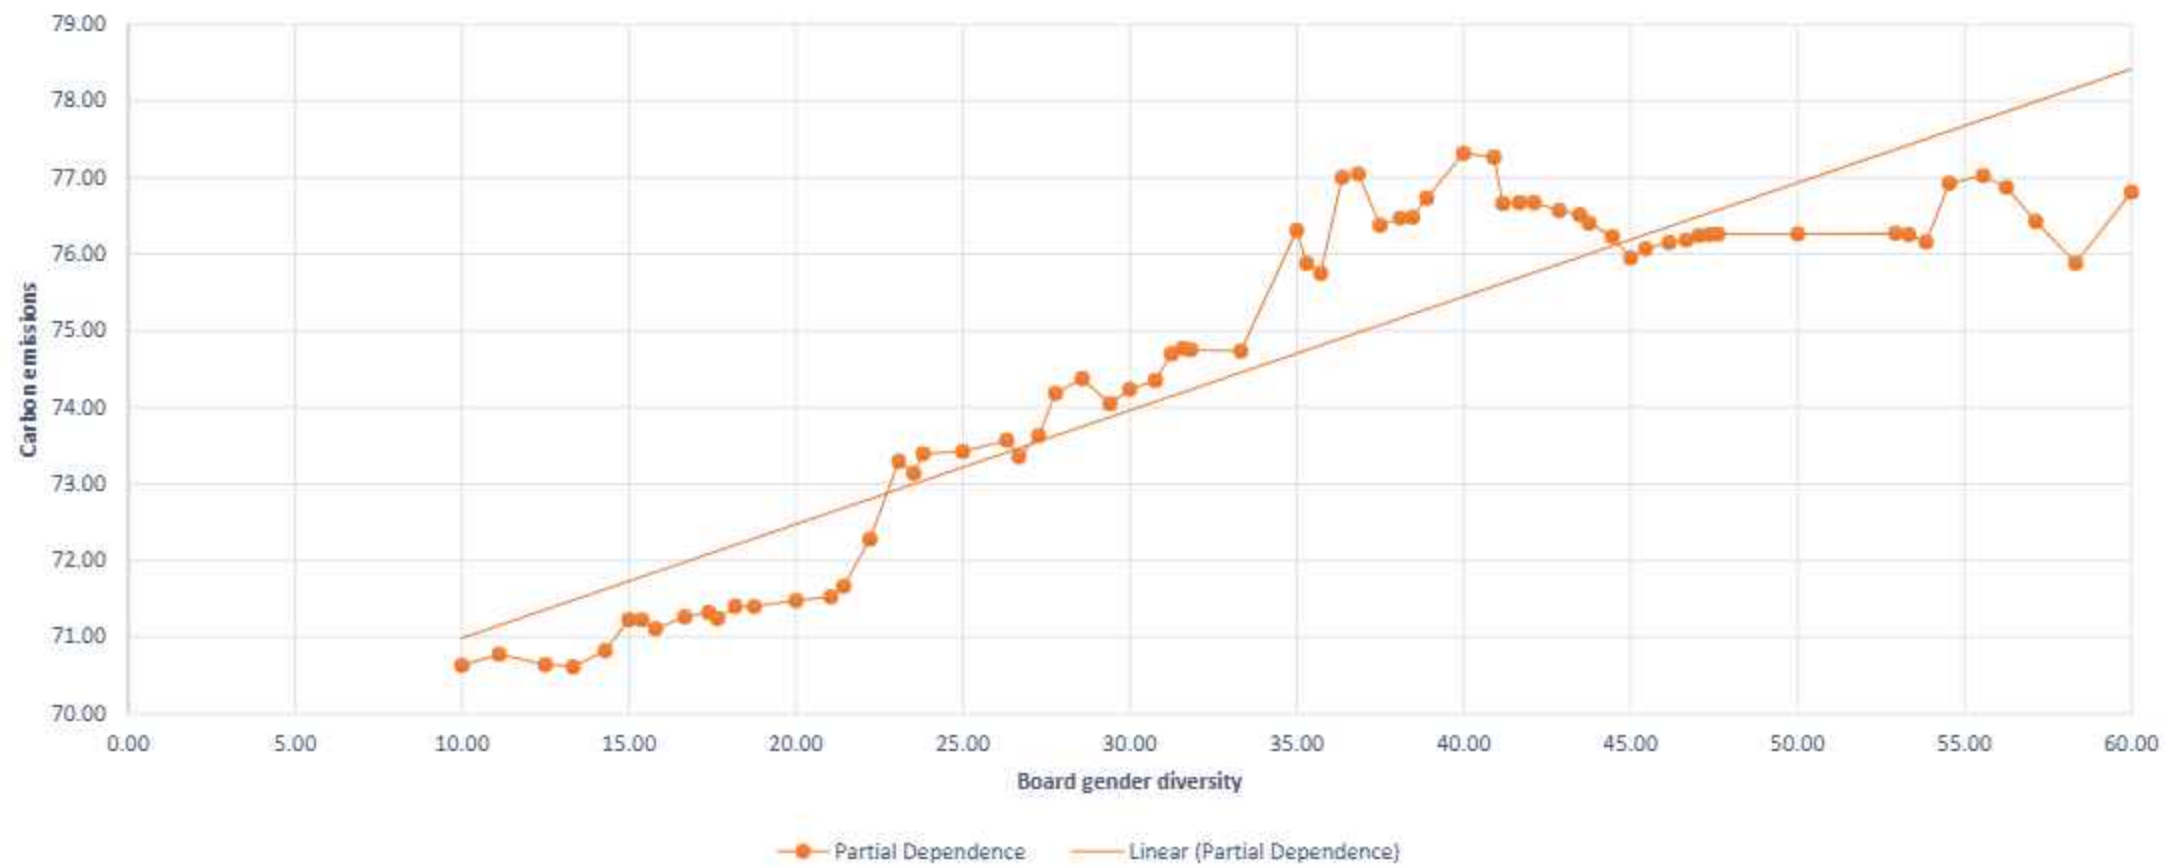
Table 8: Carbon emissions predictions using 5-fold Cross-validation

Model	Training obs.	MAE	RMSLE	R^2
XGBoost	2,090	9.0350	<i>0.2573</i>	<i>0.7184</i>

Feature Impact



Partial Dependence Plot



Further
analysis



Table 9: Carbon
emissions predictions
(Cross validation and
Holdout)

METHOD (RANDOM FOREST)	HOLDOUT (2022)	CROSS VALIDATION (2016-2022)
MAE	12.1686	12.4988
RMSLE	0.2599	0.3251
R Squared	0.3914	0.5102
Training observations	2,546	2,090

Thank You!



Forecasting climate driven healthcare surges using machine learning

Steven Squires

Postdoctoral Research Fellow

Department of Computer Science

University of Exeter

s.squires@exeter.ac.uk



Health System Resilience to Climate-related Surges (HERCS)



Strengthening the Kenya health system's capacity to anticipate, respond and adapt to prevent climate-driven health crises



Core Research Team

1. Peter Hailey, Co-PI
2. **Dr. Matt Fortnam, Co-PI**
3. Nancy Balfour-Director
4. **Dr. Tinkle Chugh, WP3 lead**
5. Dr. Anne Khisa, PDRF
6. Jackson Wachira, PDRF
7. **Dr. Steven Squires, PDRF**

Additional Exeter people:

Dr. Theo Economou
Dr. Katy Sheen

HERCS aim



HERCS will pioneer participatory methods with health workers and communities, **combined with artificial intelligence (AI)**, to enhance **health system climate resilience**. We aim to:

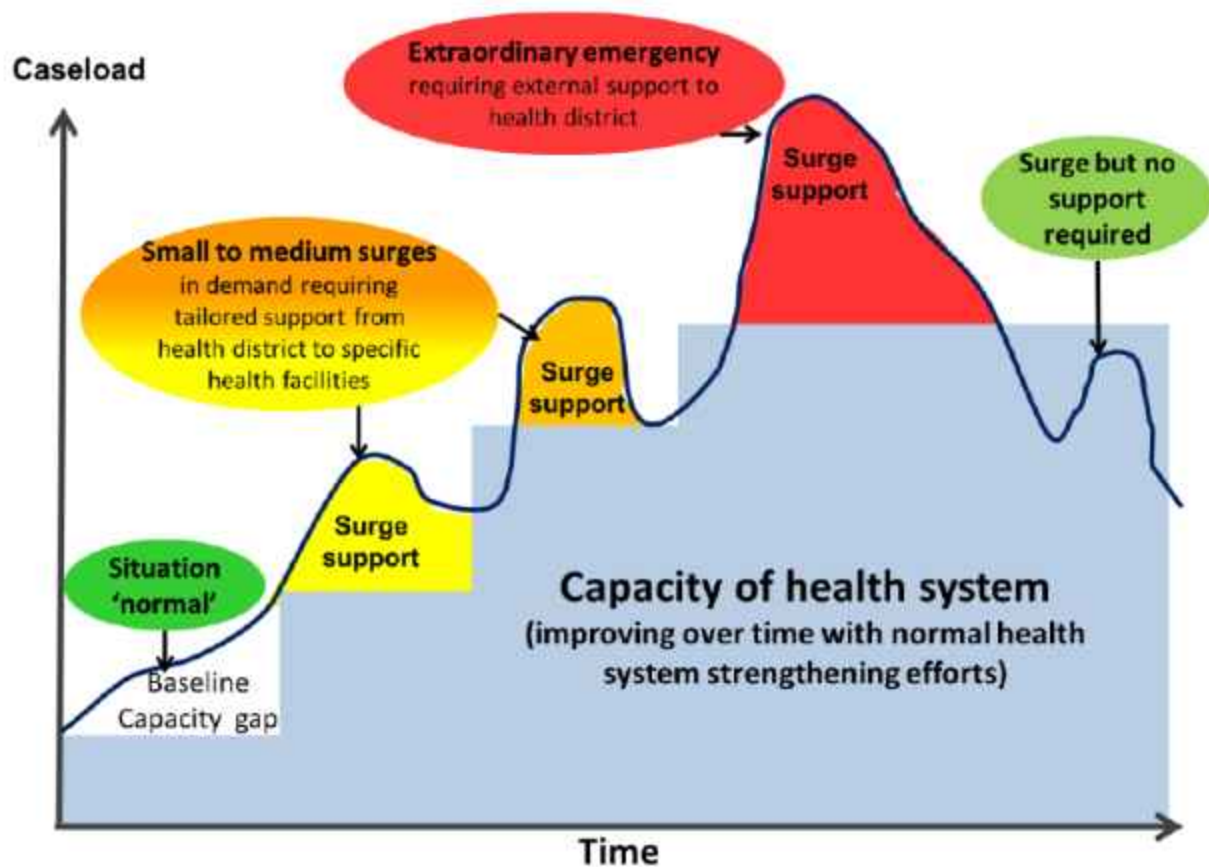
1. Strengthen capacities to **anticipate** and respond to climate-driven surges in cases and demand in communities and health facilities
2. **Develop capabilities to forecast the probability of climate-driven surges in demand at health facilities**
3. Strengthen the long-term resilience of health facilities and communities to the impacts of intensifying weather extremes and climate change on health

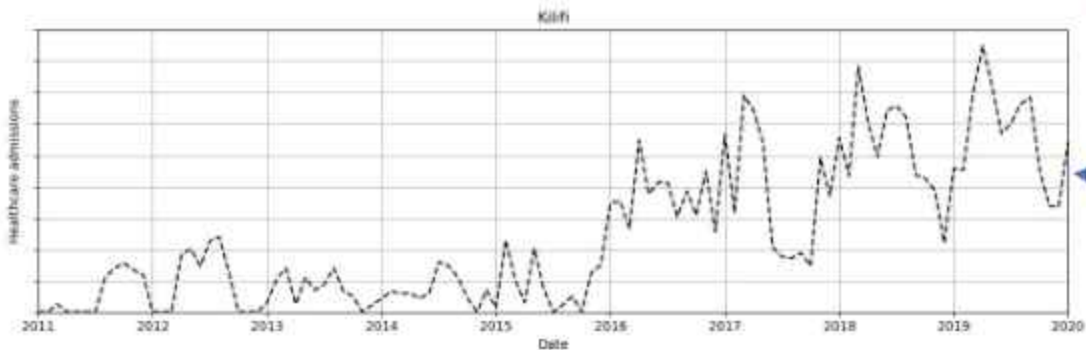
Impact



1. **Enhanced** real-time and early warning information on climate-driven health surges
2. Health workers and communities **empowered** to use health information, assess climate risks, take anticipatory and early action, and build their long-term climate resilience
3. Health resilience assessment and forecasting tools that are **scalable** to fragile, climate-vulnerable counties in Kenya

Problem



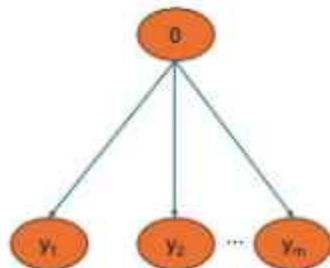
University
of Exeter

Planned methods

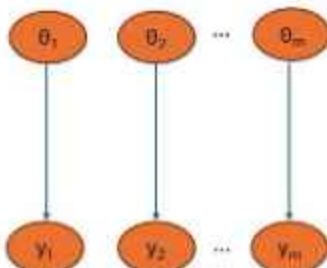
Inputs: Climate variables. Socio-economic variables. Other related variables.

Methodology: Bayesian hierarchical models. Other approaches to explore.

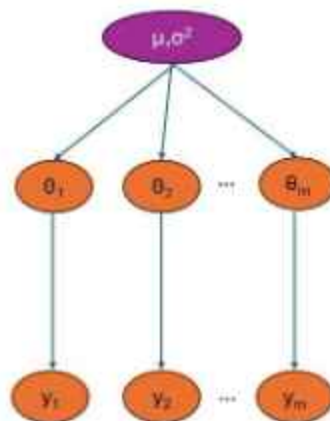
Pooled analysis



Independent analysis



Hierarchical analysis



Summary

Want to:

- Accurately predict healthcare surges with uncertainty estimates
- Understand causes of healthcare surges especially relating to climate
- Use predictions and improved understanding to improve healthcare resilience

Using:

- Climate related data
- Socio-economic data
- Any other data related to healthcare

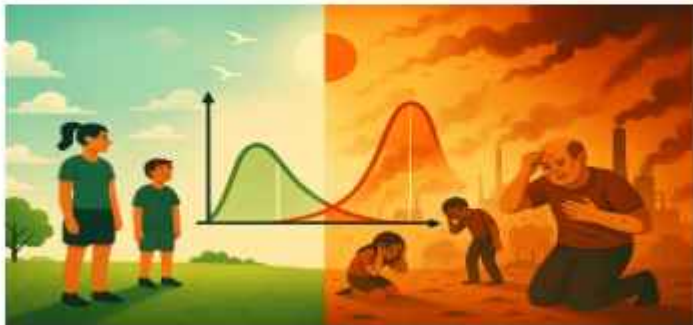
We would love to discuss/collaborate with anyone interested in any aspect:

s.squires@exeter.ac.uk

Statistical models for temporally distributed effects from multiple environmental stressors on health

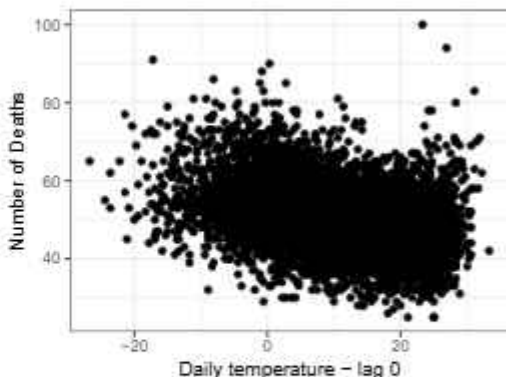
Theo Economou (t.economou@exeter.ac.uk)

D. Parliari, F. Kekkou, A. Tzyrkalli, C. Giannaros, C. Sarran



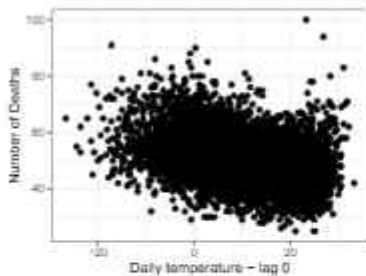
Effects from environmental variables on health outcomes

- **Goal:** quantify association between environmental variables with health metrics.
- E.g. Daily mortality (1987–2000) in Chicago and daily mean temperature[†].
 - (Plot excludes deaths from 1995 heatwave)
- The effect of temperature is not immediate, but rather **distributed over a few days**.

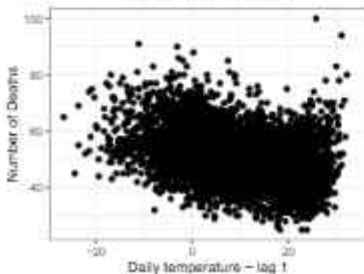
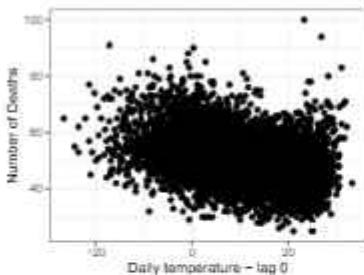


[†] Aurelio Tobias and Lina Madaniyazi (2024). "todatasets". In: *Mendeley Data*. DOI: [10.17632/xjgrahut5h.4](https://doi.org/10.17632/xjgrahut5h.4)

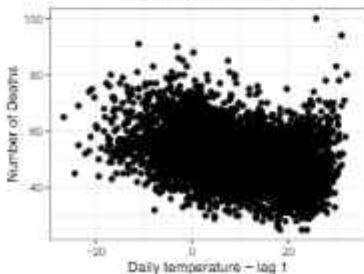
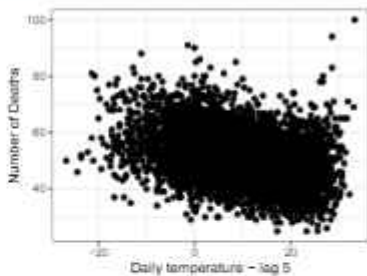
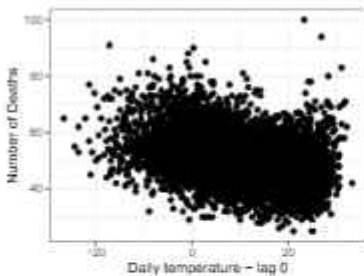
Temporally distributed effects



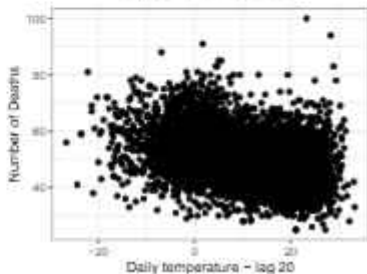
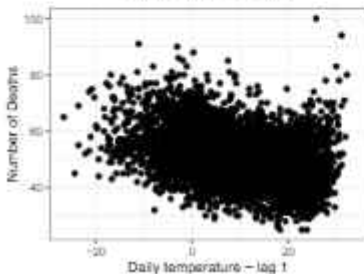
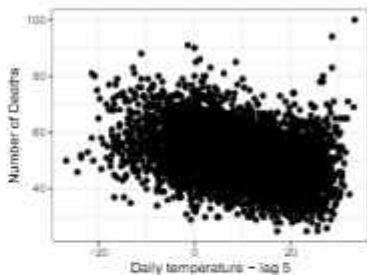
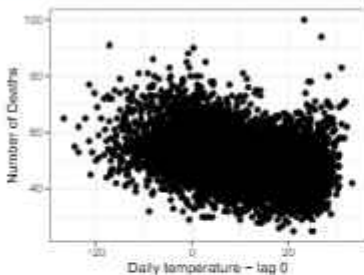
Temporally distributed effects



Temporally distributed effects



Temporally distributed effects



Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

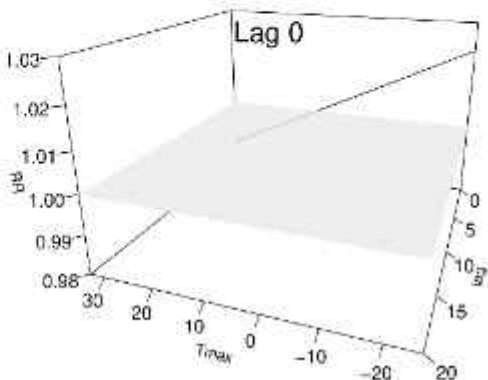
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

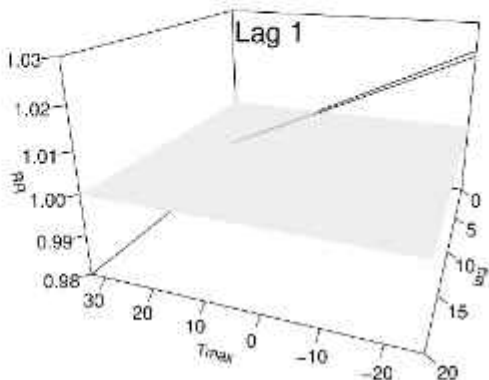
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

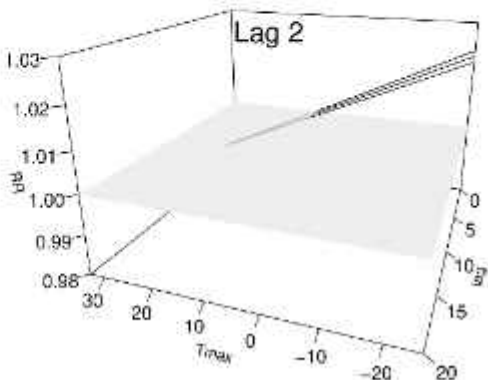
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

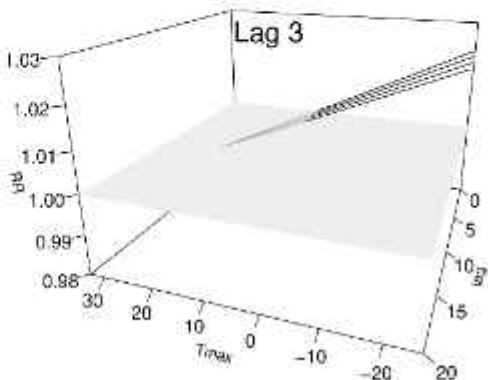
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

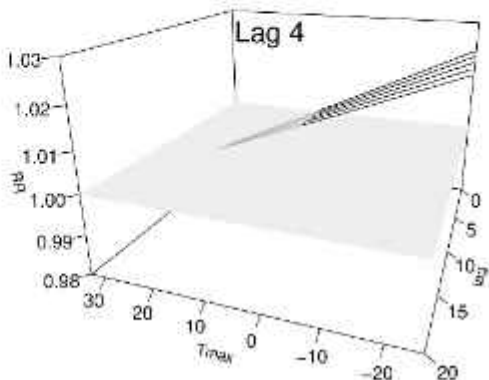
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Distributed Lag Models

- **Regression:** Linear lagged effects modelled **additively**[†]. For lags $\ell = 0, 1, \dots, L$:

$$\log(\text{mean \# deaths on day } t) = \alpha + \beta_0 x_t + \beta_1 x_{t-1} + \dots + \beta_L x_{t-L}$$

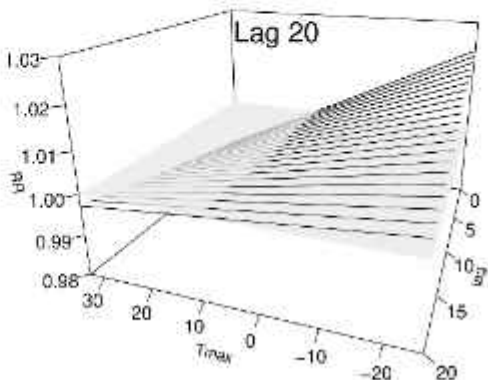
- Similarity in nearby lags:

$$\beta_\ell = \text{constrained}$$

- **Relative risk**

$$RR(x, \ell) = \exp\{\beta_\ell x\}$$

- $RR(x, \ell) > 1$ means **higher mortality than average**



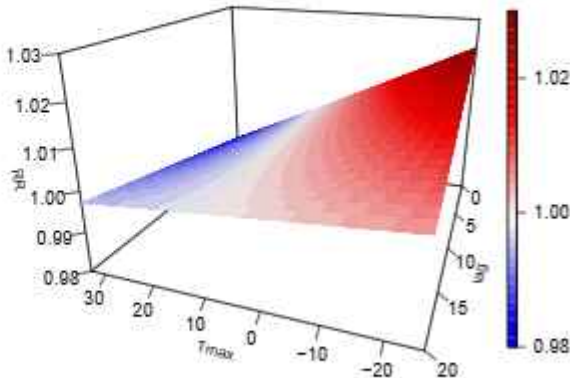
[†] Shirley Almon (1965). "The Distributed Lag Between Capital Appropriations and Expenditures". In: *Econometrica* 33, p. 178

Interpretation of the estimates in terms of risk

- The estimates indicate that the **risk is highest at low temperatures at short lags** (coldwaves?).
- Also, **lower than average mortality risk is evident at high temperatures over short lags** (no heatwave effect).
- However, the resulting **surface** is constrained by assumption of linearity in the effect of x :

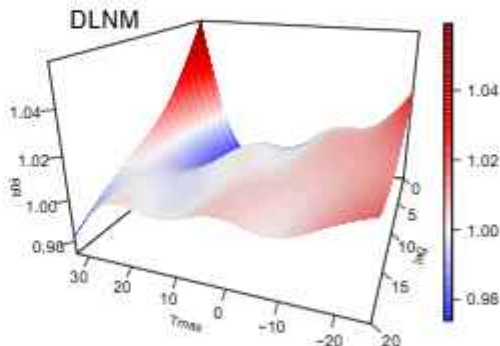
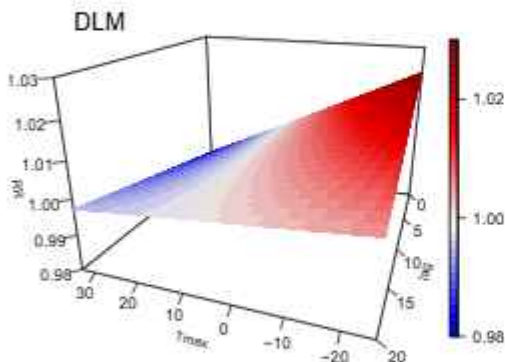
$$RR(x, \ell) = \exp\{\beta_{\ell}x\}$$

3D "surface".



Non-linear lagged effects

- **Distributed Lag Non-linear Models** or **DLNMs**[†] relax the linearity assumption:

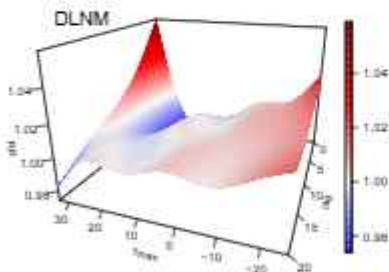


- The DLM completely misses the **spiked heatwave effect at high temperatures and short lags!**

[†] A Gasparini et al. (2010). "Distributed lag non-linear models". In: *Statistics in Medicine* 29 (21), pp. 2224–2234. URL: <https://doi.org/10.1002/sim.3540>

Generalized Additive Models (GAMs)

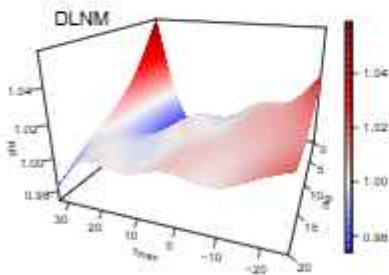
- Easy to fit DLNMs as GAMs in R package `mgcv`[†]:
 - Penalisation ensures optimal flexibility.



[†]Simon N. Wood (2017). *Generalized Additive Models*. Chapman and Hall/CRC. ISBN: 9781315370279. URL: [10.1201/9781315370279](https://doi.org/10.1201/9781315370279)

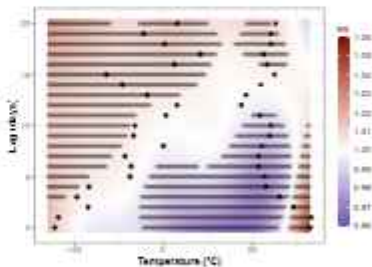
Generalized Additive Models (GAMs)

- Easy to fit DLNMs as GAMs in R package `mgcv`[†]:
 - Penalisation ensures optimal flexibility.



[†]Simon N. Wood (2017). *Generalized Additive Models*. Chapman and Hall/CRC. ISBN: 9781315370279. DOI: [10.1201/9781315370279](https://doi.org/10.1201/9781315370279)

- Bayesian inference enables thorough **uncertainty quantification**[‡].
- grey points imply **RR is (statistically) significantly not 1**.



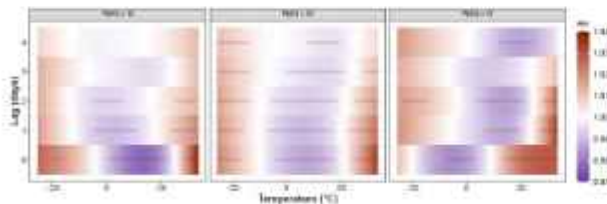
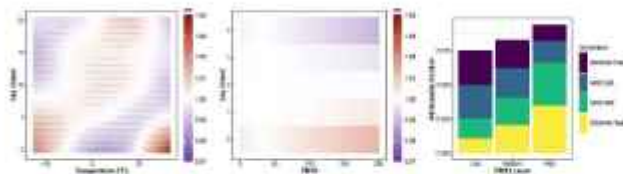
[‡]Andrew Gelman et al. (Nov. 2013). *Bayesian Data Analysis*. Chapman and Hall/CRC. ISBN: 9780429113079. DOI: [10.1201/b16018](https://doi.org/10.1201/b16018)

Synergistic effects

- Often, the temperature acts in **synergy** with other stressors such as PM10 (air quality).

Synergistic effects

- Often, the temperature acts in **synergy** with other stressors such as PM10 (air quality).



Temperature-PM10 synergy for Chicago.

- Straightforward to fit synergistic models in `mgcv`.
- Different stressors can have **different lag periods**.
- Hierarchical formulation:**

`temp + PM10 + temp:PM10`

Categorical synergies

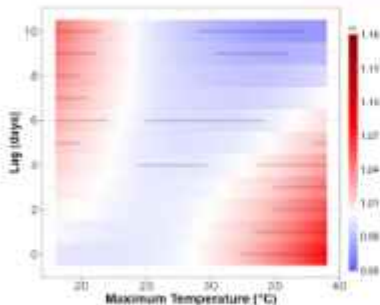
- Synergy with categorical variables[†] achieved via

overall lagged effect + deviation

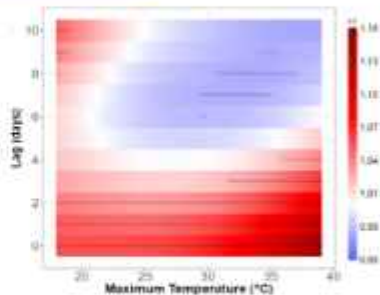
Categorical synergies

- Synergy with categorical variables[†] achieved via

overall lagged effect + deviation



Warm period – “normal days”

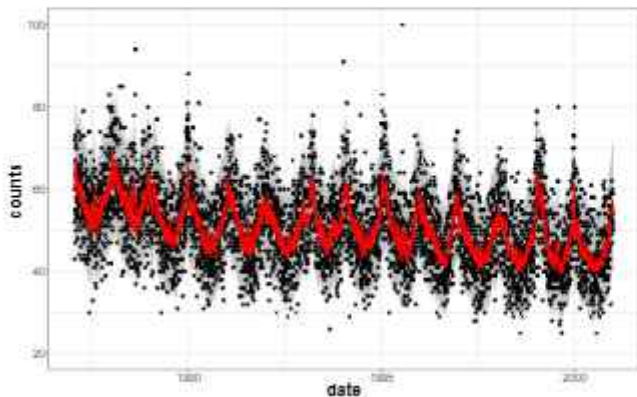


Warm period – “heatwave”

[†] Fragkeskos Kekkou et al. (2025). “Temperature extremes and human health in Cyprus: Investigating the impact of heat and cold waves”. In: *Environment International* 199, p. 109451. doi: [10.1016/j.envint.2025.109451](https://doi.org/10.1016/j.envint.2025.109451)

Flexibility and model validation

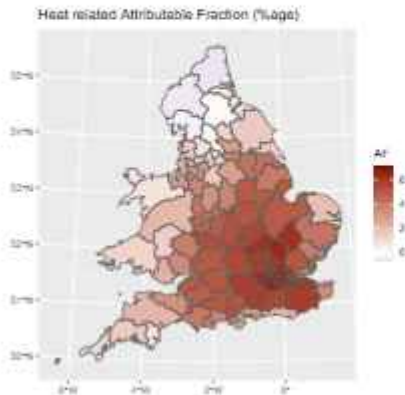
- Ample flexibility via **smooth functions interpreted as random effects**.



- **Temporal structures** such as different seasonal cycle per year.

Flexibility and model validation

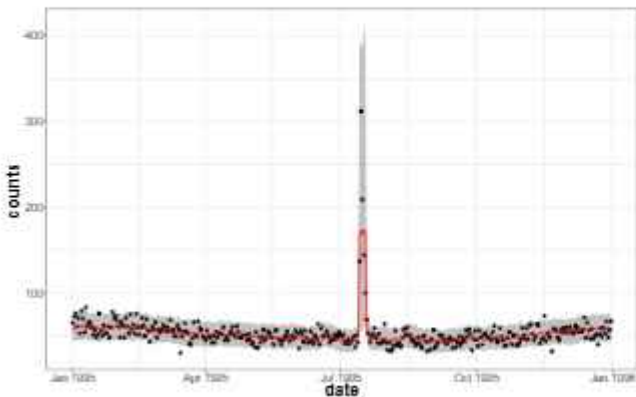
- Ample flexibility via **smooth functions interpreted as random effects.**



- **Temporal structures** such as different seasonal cycle per year.
- **Spatial structures:** e.g., different risk profile per region.

Flexibility and model validation

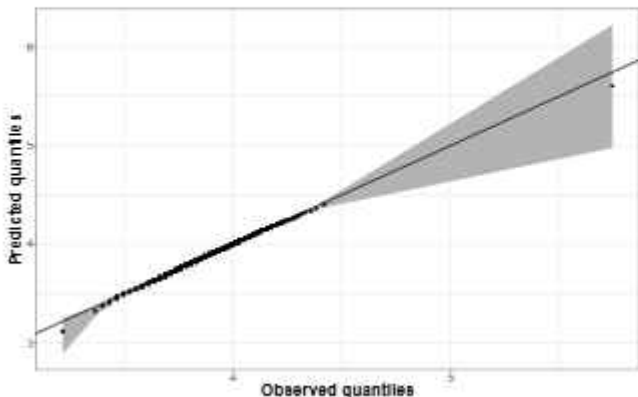
- Ample flexibility via **smooth functions interpreted as random effects**.



- **Temporal structures** such as different seasonal cycle per year.
- **Spatial structures**: e.g., different risk profile per region.
- **Mixture distribution** to capture extremes.

Flexibility and model validation

- Ample flexibility via **smooth functions interpreted as random effects**.



- **Temporal structures** such as different seasonal cycle per year.
- **Spatial structures**: e.g., different risk profile per region.
- **Mixture distribution** to capture extremes.
- **Predictive model checking** to ensure model is a good fit to the data.

- Fitting DLNMs in `mgcv` provides a very flexible approach[†] for quantifying the lagged effects of covariates on health data:
 - synergistic effects
 - objective penalisation of the non-linear functions
 - quantification of uncertainty via Bayesian (Monte Carlo simulation)
 - addition of confounding effects (e.g. seasonal cycle and population characteristics).
 - and temporally/spatially varying risk profiles.

[†] Theo Economou et al. (2025). "Flexible distributed lag models for count data using `mgcv`", in: *The American Statistician* 0, pp. 1–18, [doi: 10.1080/00031305.2025.2505514](https://doi.org/10.1080/00031305.2025.2505514)

- Fitting DLNMs in `mgcv` provides a very flexible approach[†] for quantifying the lagged effects of covariates on health data:
 - synergistic effects
 - objective penalisation of the non-linear functions
 - quantification of uncertainty via Bayesian (Monte Carlo simulation)
 - addition of confounding effects (e.g. seasonal cycle and population characteristics).
 - and temporally/spatially varying risk profiles.
- Thank you!

[†] Theo Economou et al. (2025). "Flexible distributed lag models for count data using `mgcv`", in: *The American Statistician* 0, pp. 1–18, [doi: 10.1080/00031305.2025.2505514](https://doi.org/10.1080/00031305.2025.2505514)

Accurate Enough? Opening up sensing of the Cryosphere

Ellie Fox

PhD Student in Environmental Intelligence – Geography Department,
University of Exeter

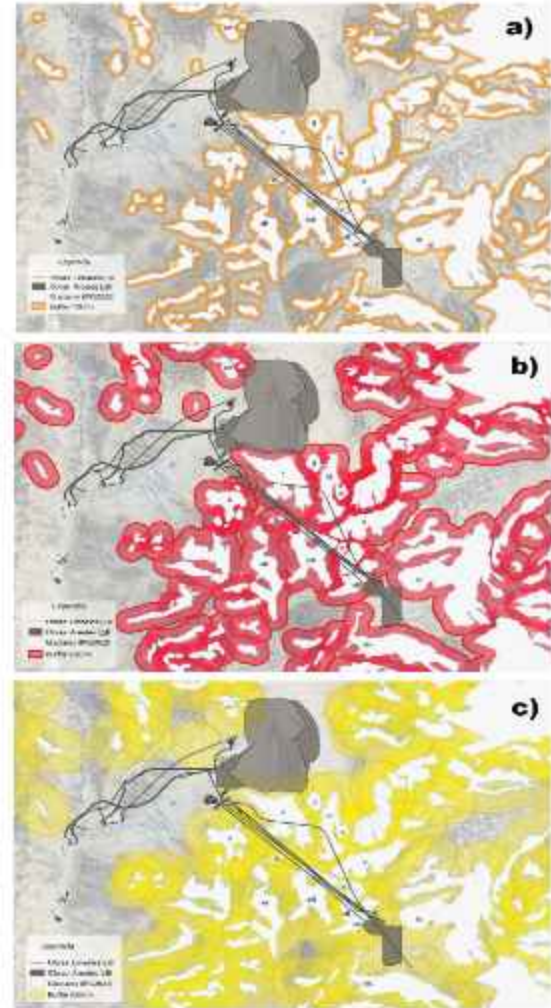
Environmental Intelligence for complex socioecological challenges

- Threats to glaciers in Chile from climate change and mining.
- ‘Small’ data ethnographic research to understand the power dynamics of this issue.
- Informs approaches to how big data can be used to support environmental justice.



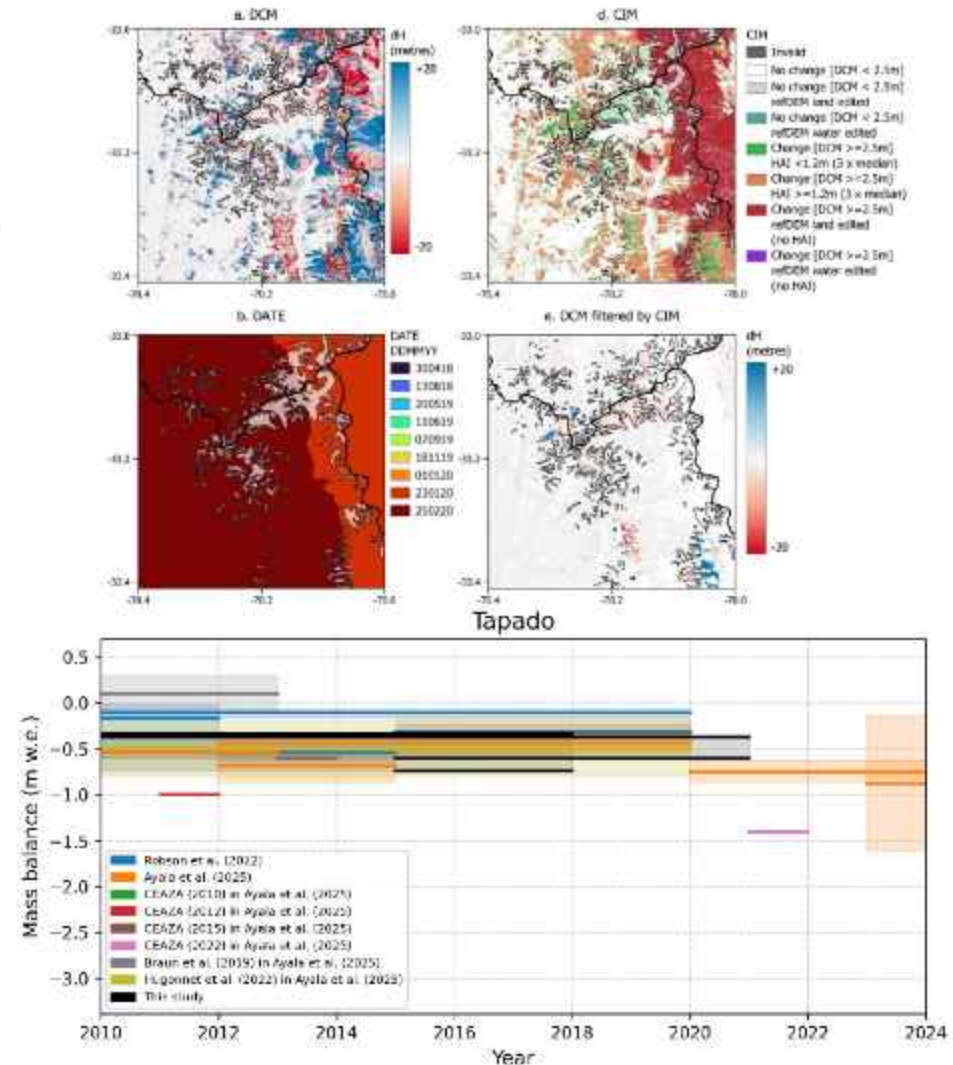
The Challenges of Accuracy

- Presents challenges for using AI and big data in environmental analysis:
 - More accurate less meaningful? (Braun, 2021; 2024); Virtue of vagueness (Cullum, 2017).
- Presents challenges for using AI and big data in policy and legal settings:
 - Cohen (2001); Moon (2012); Fox and Schwartz-Marin (in Review).



E.g.1 TanDEM-X Change Maps

- Global Elevation change dataset (2010/15 – 2016/22)
- Freely available online, easy to use portal, shows clear glacier change
- Issues – date availability, uncertainty
- Analysis shows data is in line with and extends published estimates.



E.g.2 Citizen sensing



E.g.3 More than human sensing practices

- Emotions, animals, plants and weather out of place (Gabrys, 2019; Howe, 2019)
- Global weirding (Turnbull et al., 2022)



Ellie: How did a puma arriving in the city lead to the Corporation being founded?

Juan: Well, in 2019, a puma appeared in Sector El Arrayan, which is where we live, and the puma climbed a tree, and everyone was really alarmed, they were really worried, because there was a puma in people's homes, and with kids you know, can you imagine? Also because here the wells had dried up, and there are areas where they've had to deepen the well to 100 metre because there was no water. So, with the puma coming, glacier melt, wells drying up, amongst other things, the residents got talking, we started to meet up, and we said: right, something must be happening here, and it has to do with the mining company. This is how it was born.

Summary

1. Integrating social science from the beginning of the research process:
 - Enables the power relations in how data and models are used to be understood
 - Leads to novel understandings of problems and solutions
2. Embracing uncertainty, vagueness, and going beyond the exclusively quantitative approaches, expands our ideas of:
 - What data is relevant to understand environmental change
 - What actions are possible to address environmental injustices

AI in Practice: Access to compute and data

Adrian Hines

Director of JASMIN

Science and Technology Facilities Council

The UK Compute Landscape

A variety of options for access to compute resource within the UK

Local group resources

Institutional systems

Large-scale discipline-based systems

National systems

Commercial Cloud

- How do you decide what compute facilities to use for your work?
- Does this support the scale of compute and storage that you need?
- Are there better suited options?



Department for
Science, Innovation
& Technology



UK Research
and Innovation

**International
Partnerships**

EuroHPC

Belfast

NI-HPC

Daresbury

Hartree Centre

**Oxford &
Harwell**

JASMIN
GridPP
Iris

Bristol

Isambard-AI
(AIRR)
Isambard-3

Exeter

Met Office

Edinburgh
1st National
Supercomputing
Centre (NSC)

★ **New National
Supercomputer**
Archer2
DiRAC
Cirrus

Durham

DiRAC
Bede

Leicester

DiRAC

Cambridge

Dawn (AIRR)
DiRAC
CSD3

London

Materials and
Molecular
Modelling Hub

Aldermaston

Atomic Weapons
Establishment
(AWE)



Science and
Innovation
Future
Establishment

Source: UK Compute Roadmap, DSIT, July 2025

JASMIN

Supports data analysis for environmental science community

- Large scale, data-intensive science

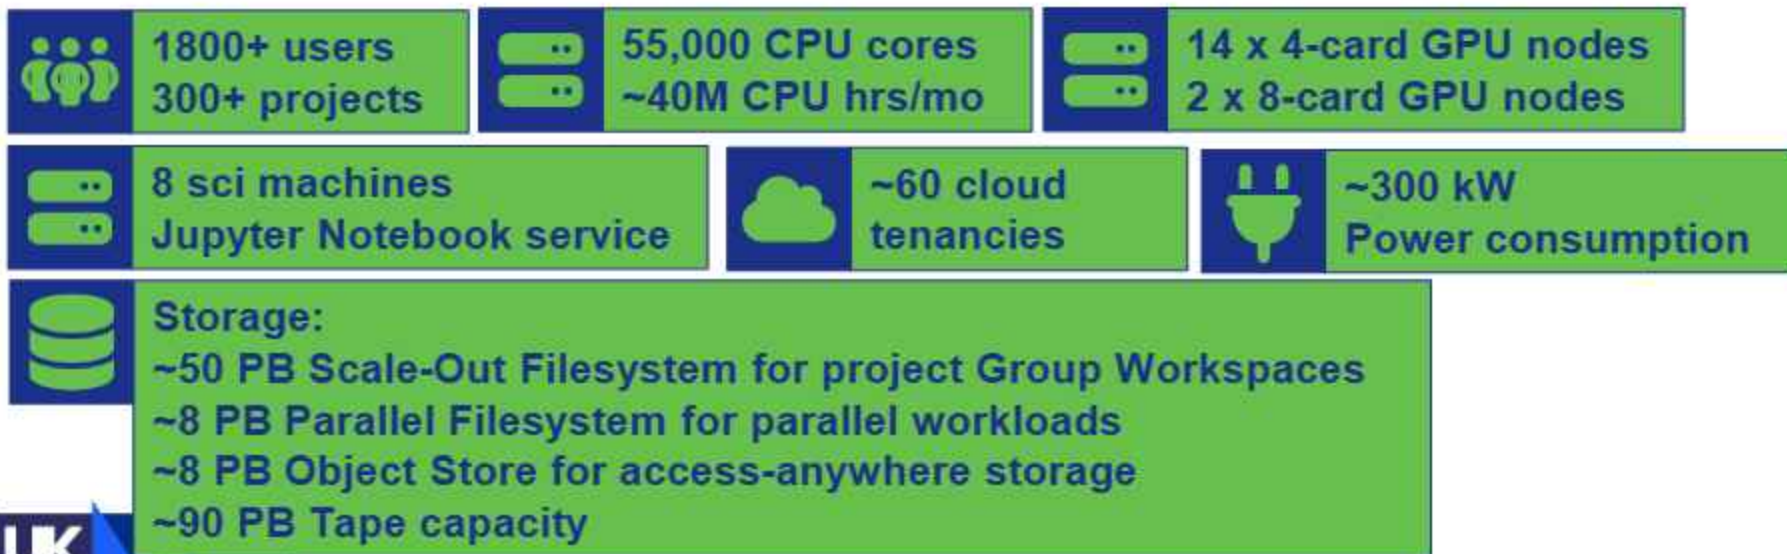
Designed for performance

- Centred more around storage & data analysis than traditional supercomputer
- Philosophy of bringing the compute to the data

Flexible compute capabilities

- Interactive and batch compute
- CPU and GPU
- JASMIN Cloud

Operated by STFC on behalf of NERC



Environmental sustainability

- Use of computing in research is energy-intensive
- Pressing need to reduce environmental impact

Breaking down the environmental impacts of computing

Environmental impact of computing

Powering the computer

Long-term
data storage

Life-cycle footprint
of the hardware

Global GHG emissions of data centres

~100 Mt of CO₂e / year

Equivalent to American commercial aviation

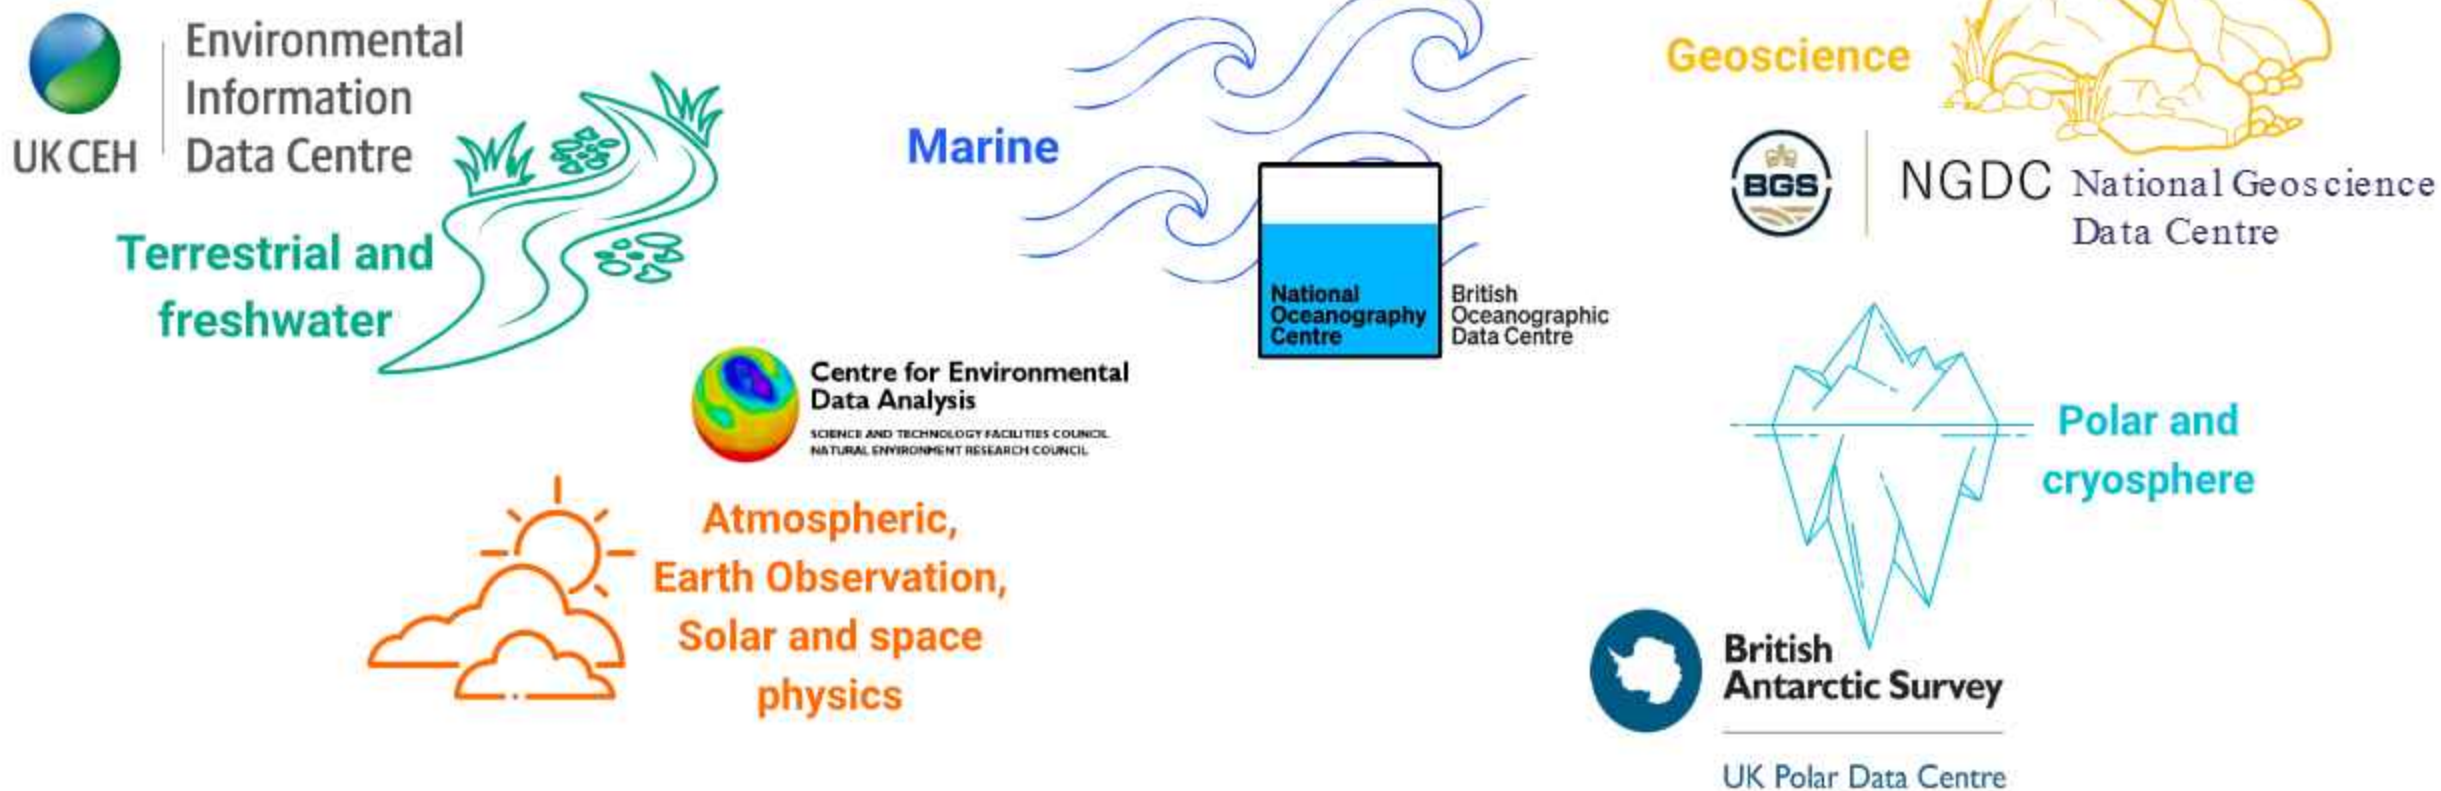


Source: Loic Lannelongue, University of Cambridge

- How do you ensure that you are using compute efficiently?
- Is your code optimised for the hardware that you are using?
- Is the hardware the most efficient option?
- Does your experimental design and data management plan consider environmental impact?

The NERC Environmental Data Service

A focal point for environmental scientific data and information – via a network of distributed data centres, with domain specific expertise.



- Do you have access to the right data for your work?
- Are there data that the EDS may hold that could be of use?

Interdisciplinarity vs Insularity in AI “Ethics”

Catie Johnston

cj505@exeter.ac.uk

PhD Candidate, CDT Environmental Intelligence

Environmental Intelligence @ Exeter Conference

AI Ethics in Practice

The “Ethics” Conversation

Conversation often stays within technical boundaries. Important, but framed as technical problems with technical solutions.

Technosolutionism

Belief that social and political problem can be solved via sufficiently advanced technologies.

E.g. Algorithmic Bias

E.g., Transparency

E.g., Trustworthiness

These are important but substantive and technical in nature, centring questions of AI safety, efficiency, and accuracy. Solutions to these problems are also technical in nature, and **the arising feedback loop results in AI development that is very insular.**

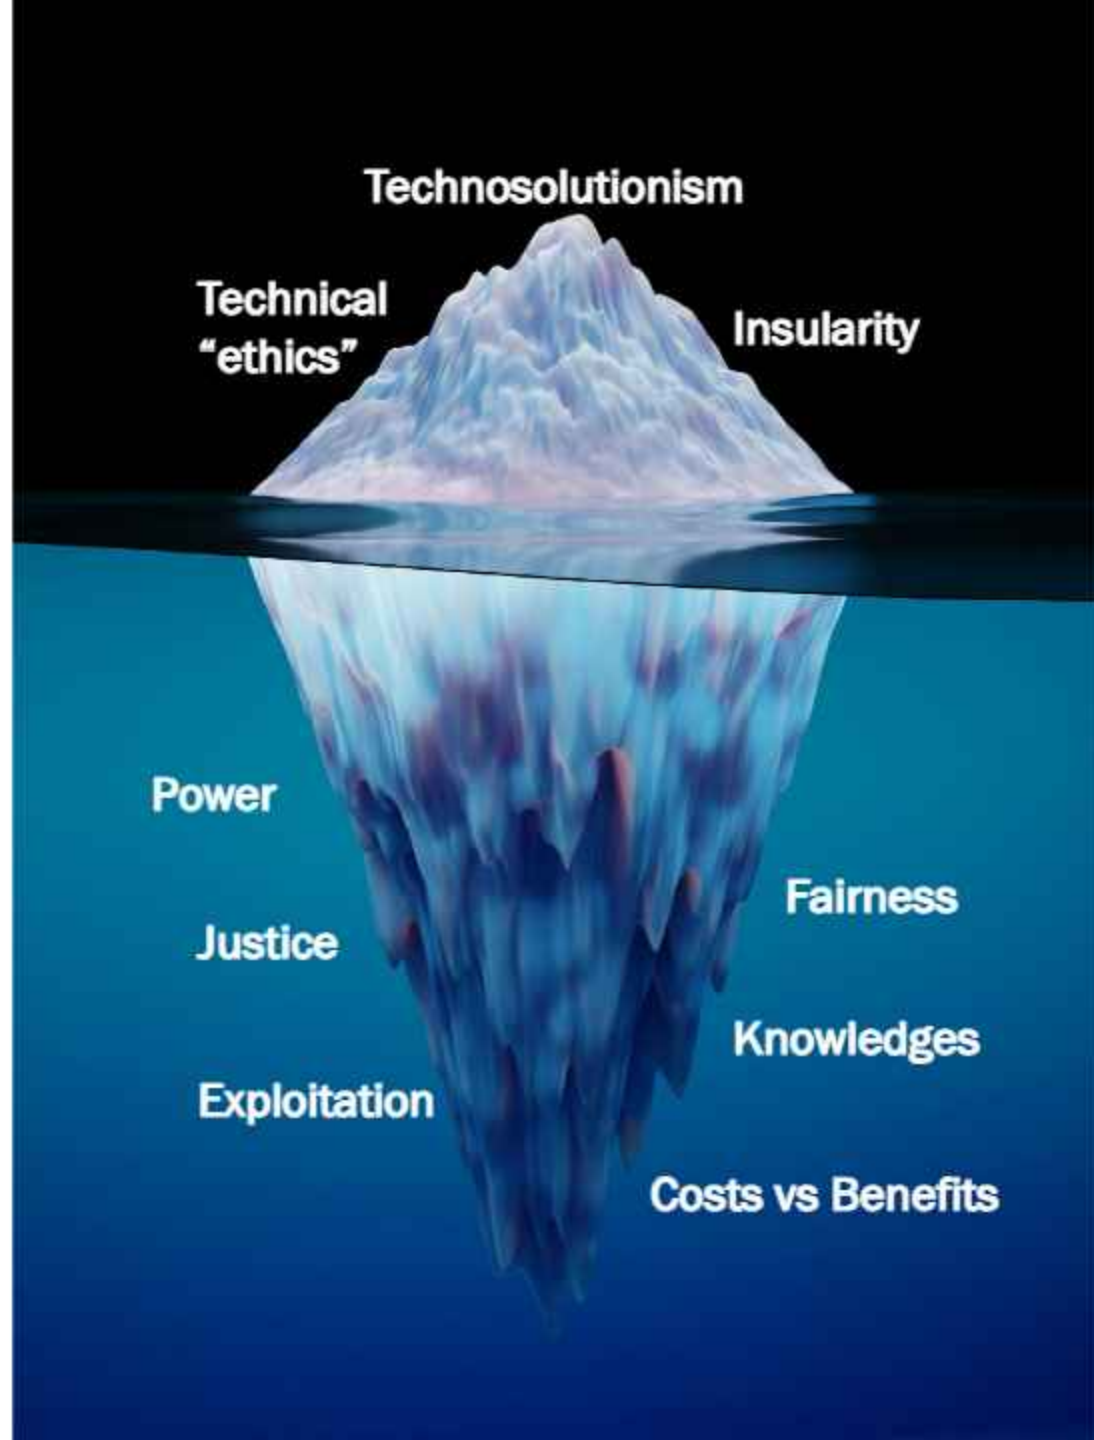
Misses the Big Picture

Technical innovation alone does not ensure responsible, socially just, or environmentally sound outcomes.

Need to shift focus from narrow, technical questions of capability and accuracy to broader concerns of power, legitimacy, and epistemic justice.

Delicate community considerations are often overlooked in favour of more technical or logistical challenges.

Who decides what fairness means? Whose knowledges are being listened to? Who decides when AI is an appropriate solution? Who benefits from AI systems? Who bears the cost?



From Insularity to Interdisciplinarity

Tickbox Approach to Ethics

- In an insular system, consideration of “ethics” risks becoming a tickbox exercise
- Solutions to “technical ethics” neglect broader harms (e.g., labour exploitation, environmental costs, cost-benefit imbalances, and power dynamics)
- Dominated by technical framing.
- Self-referential, internal governance, “ethics-washing”

Need for Criticality

- Sharing knowledge across disciplines
- Ethical inquiry requires plurality of knowledges
- Interdisciplinary practices (e.g., involving computer science, environmental sciences, social and political sciences) are a practical and actionable first step.
- Redistributing authority (instead of just adding experts) is critical for avoiding tokenism.

Beyond Interdisciplinarity

Insularity

- Narrow
- Ethics-washing
- Technosolutionist, self-referential



Interdisciplinarity

- Broader, critical, plural
- Can be tokenistic
- Dominated by technical frames



Inclusion

- Justice-oriented
- Redistributive
- Responsive to real harms

Catie Johnston

cj505@exeter.ac.uk

မုန့်

$$C \leftarrow \bigcirc \check{R} \equiv \leftarrow P \mid \check{C} \leftarrow \equiv \alpha \check{A} \leftarrow \equiv \check{C} \alpha \check{R} \check{C} \leftarrow \check{C} \check{A} \equiv \leftarrow P \check{C} \mid \check{C} \leftarrow \check{C} \hat{A} \hat{U} \hat{V} \hat{W}$$
$$\begin{array}{c} ?^{-} \leftarrow A \equiv \alpha \check{R} \parallel \\ \wedge \check{C} \check{S} A \check{I} A a \end{array}$$

WUWYUJAJHq



« ČA |ČAW|Č =L8 |RŁ R-Ě|
←-|RŁ←- oApts- =|= RŁA|Č |RŁČ

- Established over 230 years ago
- Trusted source of geospatial data
- Supporting national infrastructure, public services and innovation
- Operating globally



The original
reference
for Britain's
geography



The most
advanced
map of GB
—built for
everyone



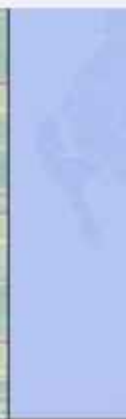
« ČA=♠ A | |=♂İ\$ ČA♠ ČA Pts=||A/4=Pts | €ö Č←||RČA\$ R.R= A
 |€= |€||Č←=||R←A1'A |R= R A | UUA Pts R R←A=1/4 ||R←A
 1'Č=||€||A= S=||€S AŵUUAUA Pts Č||A Č|=S= IJ

A close-up of a person's face, wearing glasses and drinking from a clear plastic water bottle. A semi-transparent notification box with a dashed line pointing to the bottle contains the text "12 km to go".

Maps in everyday life: how does OS data power navigation apps



$\check{C}A_2\check{C} \leftarrow \frac{1}{4} \check{O}PtS^- \quad | \quad \check{C}A_2\check{C} \leftarrow \frac{1}{4} \check{O}PtS^-$
 $\check{C}A_2\check{C} \leftarrow \frac{1}{4} \check{O}PtS^- \quad | \quad \check{C}A_2\check{C} \leftarrow \frac{1}{4} \check{O}PtS^-$



AI ←

Strategic focus and associated actions necessary to realise the potential of AI across the business



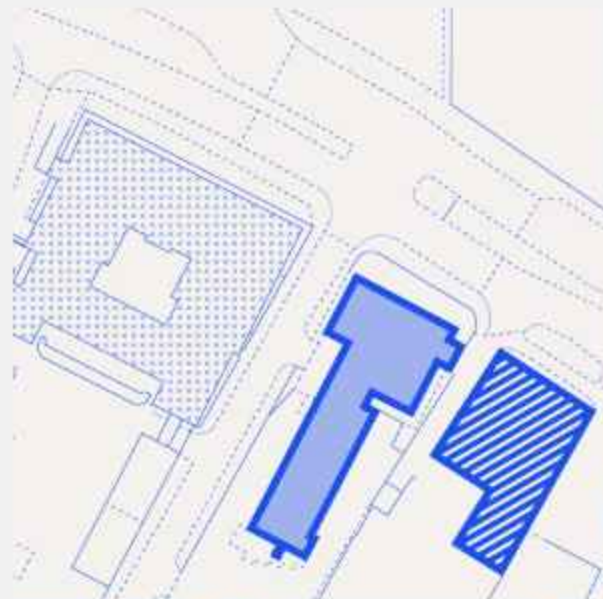
OFFICIAL

AI Plan

Shaping the Future of AI at OS

$\alpha\tilde{\Delta}\sigma^- \leftarrow$

- $\varphi \check{C} \parallel \equiv \leftarrow \parallel \check{R} \alpha \check{C} A \leftarrow \S A \leftarrow 1/4 \blacksquare \parallel \check{R} \check{C} A a$
- $\varphi \check{C} \alpha \parallel \check{R} L A i \models \S \blacksquare 1/4 \parallel \check{R} \check{R} I A \leftarrow \S A \hat{S} 1/4 \check{C} \leftarrow 1/4 I A V \check{R} \parallel$
- $V = \parallel \check{C} \parallel \check{R} L A \leftarrow \equiv \text{♂} \parallel \check{R} \leftarrow A \leftarrow \S A \check{C} 1/4 \leftarrow \check{R} 4 \alpha$
 $C \text{♩} / \check{C} \alpha \check{C} \leftarrow 1/4 \check{C}$



$$w\check{c} \equiv v \equiv \blacksquare \leftarrow \S^- \equiv \frac{J\check{c}}{K} \leftarrow A \equiv \S \check{c} \parallel$$

[illegible]

9⁻ ← A.C. ← C. = α. || ČA. ČA. = A. Ts. G. R. α. ČA. || Č. || A.
 = || A. || α. || A. 1/4 = || || A. || = || R. α. Č. = G. R. ← ||

$$\mathfrak{c} \equiv 1/\left\| \frac{\mathbb{A}}{A} \right\| = A \mathfrak{o}^- \leftarrow \mathbb{L} \check{\mathbb{C}} \left\| \check{\mathbb{A}} \check{\mathbb{A}} \right\| \left\| \mathbb{S}^- \right\|$$





䷧ ← ䷦

$\text{? } \check{C}P\check{S} \equiv \frac{1}{4} \vdash \frac{\check{H}}{\check{H}} \check{C} \check{A} \frac{1}{4} \check{C} \vdash \check{H} \check{H} \check{A} \leftarrow \check{S} \check{A} \leftarrow \check{d} \check{H} \check{H} \check{A} \vdash \check{A}$
 $\text{L}\check{C} \equiv \vdash \frac{\check{H}}{\check{H}} \check{R} \alpha \check{S} \vdash \frac{\check{H}}{\check{H}}$

$$a \leftarrow \frac{1}{4} | \check{C} - \| \check{C} A \| = S_{\check{C}} \frac{1}{4} | \check{C} A \| = L \check{O} A \frac{1}{4} | P_{\check{C}} - \| \check{C} \leftarrow A$$

$$\begin{array}{l} \text{c} \check{\text{C}} \text{S} \quad | \quad \check{\text{C}} \text{A} \text{S} \check{\text{C}} \text{A} \text{P} \text{S} \check{\text{C}} \text{A} \quad | \quad \text{P} \text{S} \text{S} \quad | \quad \text{A} \quad \text{A} \check{\text{C}} \quad | \quad \text{A} \quad \text{A} \\ \text{P} \text{S} \quad | \quad \text{S} \quad | \quad \check{\text{R}} \check{\text{C}} \quad \leftarrow \quad \text{S} \check{\text{C}} \quad | \quad \check{\text{R}} \quad \leftarrow \quad \text{A} \text{P} \text{S} \quad \alpha \check{\text{R}} \text{L} \end{array}$$

Climate Heat, Narratives and Environmental Intelligence



University
of Exeter

Global Systems
Institute



Environmental
Intelligence CDT

Manju Bura,

PhD Candidate, Environmental Intelligence and
Science and Technology Studies,

El Centre for Doctoral Training

11th September 2025

Ethnography, STS and Computational Text Analysis



Translation



Nepal struggles with blistering heat wave
 Extreme heat is taking its toll on the country's health and economy, particularly for the poor and underprivileged who lack cooling fans and air conditioning.



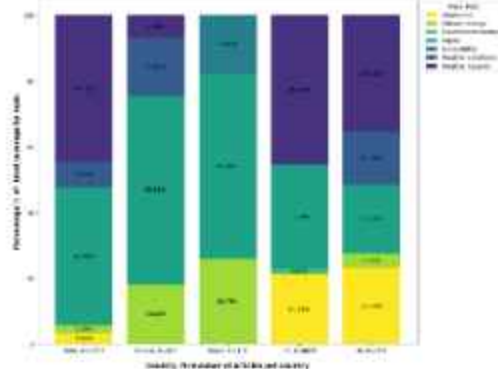
Lessons from Nepal: Coping with hotter and longer heatwaves



Topics in Nepali news about heatwaves



Figure 1. More of heat topics in online news coverage of heatwaves by country.



Beyond bias – Unalignment?

In AI and NLP, we often hear about bias.

If we can reduce bias, we can make our outputs and models fairer.



UNREPRESENTATIVE
DATA



SKEWED
DATA



INCOMPLETE
DATA

Not just bias, some worlds and practices don't translate neatly.



Translation into dominant imaginaries erases other ways of knowing.



University
of Exeter

'Fixing bias' then is about recognising this politics of translation.

Why does this matter for EI?

- Who gets to have their worlds, experiences, and ways of knowing recognised in EI?
 - AI and NLP (Natural Language Processing) tools are very good at pattern recognition, but they also flatten
 - Ethnography shows us that narratives are enacted differently in different worlds.
 - Both AI/NLP and ethnography do the work of translation.
- So how do we make different worlds visible (or commensurable) without erasing differences?
 - Bringing ethnography and NLP becomes interesting. We need methods that hold tension instead of collapsing it.
 - Combining methods with reflexivity, making space for interpretive plurality and acknowledging the politics of translation.
 - Our conceptualisation of the dramatological index to bring ethnography and NLP into dialogue to trace how translation happens in text.



- Temperatures are rising → more people are living with extreme heat.
- Unaligned narratives matter → many ways of living with heat don't fit neatly into heatwave stories.
- Translation is political → when these are reframed into dominant terms, something is always lost.
- Ethnography + AI together → can surface unaligned worlds, trace translations, and expand EI in just and reflexive ways.

Role of sociology and STS in EI goes beyond critiquing AI, and towards expanding EI practices for justice in climate knowledge.



University
of Exeter



A Probabilistic Vision-Language Framework for Remote Sensing Land Cover Classification

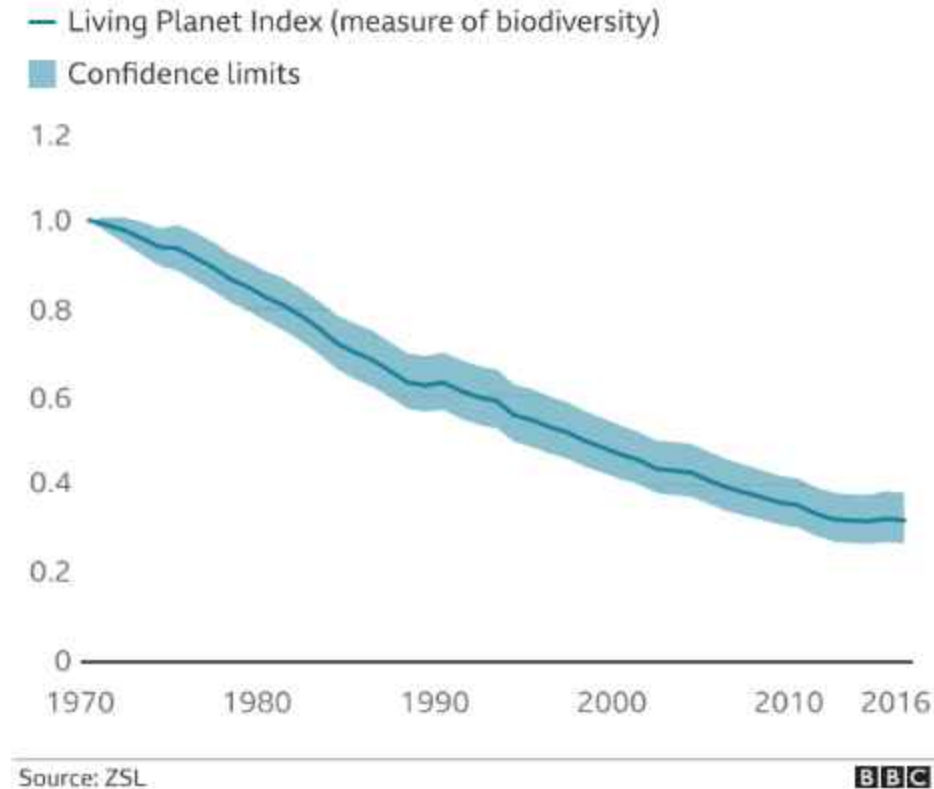
Brianna Pickstone

Supervisory Team: Prof Karen Anderson, Dr Sareh Rowlands, Prof Richard Delahay

Environmental Intelligence Conference 2025
Session 4: Machine Learning For Earth Observation

INTRODUCTION

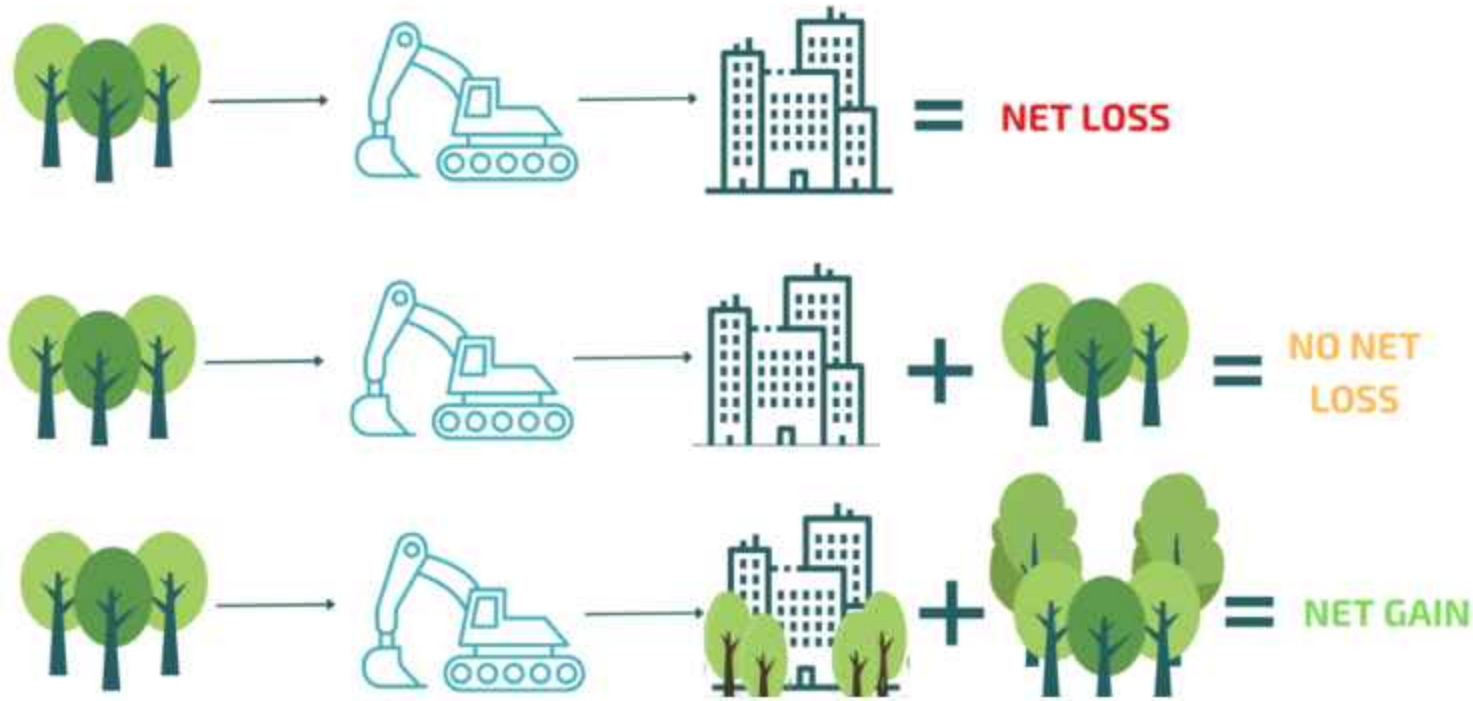
How wildlife has declined, 1970-2016



source: State of Nature Report - ZSL (2023)



BIODIVERSITY NET GAIN



Biodiversity Net Gain (BNG)

source: MGISS (2024)



MANDATORY IN
ENGLAND –
FEBRUARY 2024



SUSTAINABILITY-
FOCUSED
DEVELOPMENT



NET GAIN
OF 10%

MEETING BNG REQUIREMENTS



Baseline Biodiversity
Assessment



Calculation of Biodiversity
Units



habitat type, area, condition,
distinctiveness, and strategic
significance

MEETING BNG REQUIREMENTS

Expertise of ecologists to:

- assess baseline habitats
- plan ecological enhancements
- monitor biodiversity outcomes post-development



95% of the 337 individual respondents have “no or very limited capacity to ensure, most, if not all applications are assessed by an ecologist”

Snell & Oxford (2021)



A challenge that has been identified as a significant resource gap by Planning Advisory Service (2023); Rampling et al. (2023); Snell & Oxford (2021)

THE MISSING PIECE?



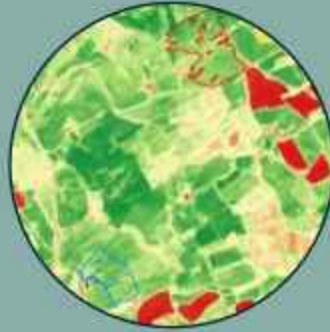
Remote Sensing



Machine Learning



Habitat Mapping



Condition Assessments



Habitat Monitoring

BNG NEEDS...



Consistent
framework



Streamlined



Transparent

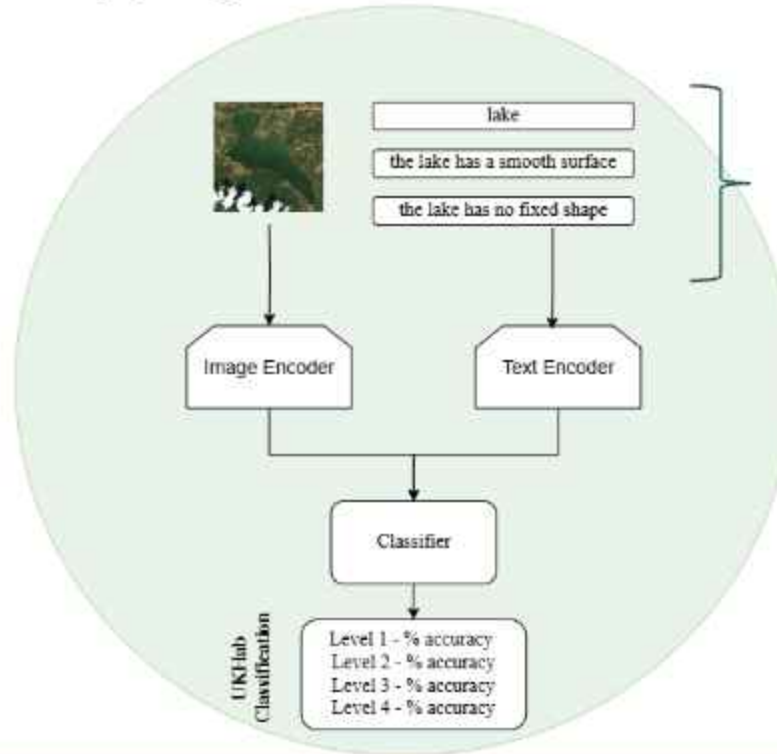


Reproducible

OVERARCHING AIM

- Assess the performance of Vision Language Models (VLM) and open-source remote sensing data for habitat mapping in the UK

Can probabilistic priors of existing land cover products support effective reclassification to UKHab categories?



What level of detail within the text data is needed?

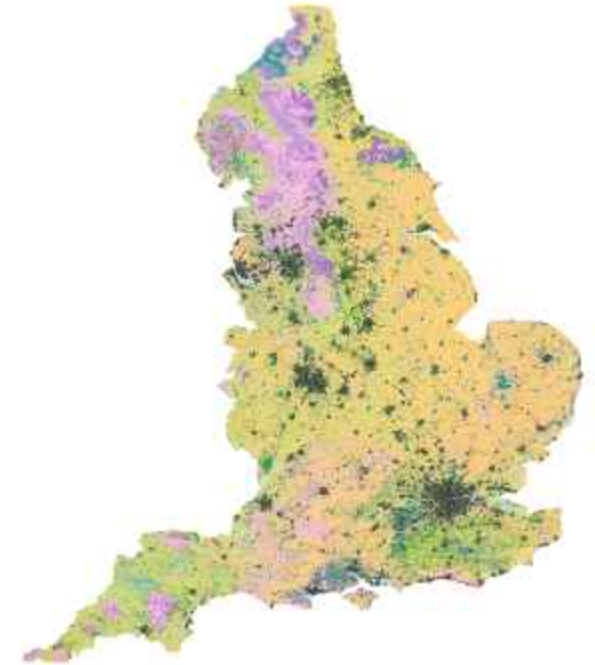
LAND COVER MAPS

UKHab – 5 Levels of Classification

- Level 1 - Terrestrial
- Level 2 - Grasslands
- Level 3 – Acid Grassland
- Level 4 – Upland acid grassland
- Level 5 – Montane acid grasslands



UKCEH Land Cover Map
(UK BAP)



Living England Land Cover Map
(UK BAP)

Why Vision Language Model?

01

image classification methods often struggle with the complexity and variability of real-world ecosystems

02

introduces the semantic understanding of objects and their relationships, moving beyond simple pattern recognition

03

this has been noted to allow for better generalisation across different habitat mapping scenarios (e.g. seasonal, and location variations) – offering a more scalable approach

Li et al. (2024)

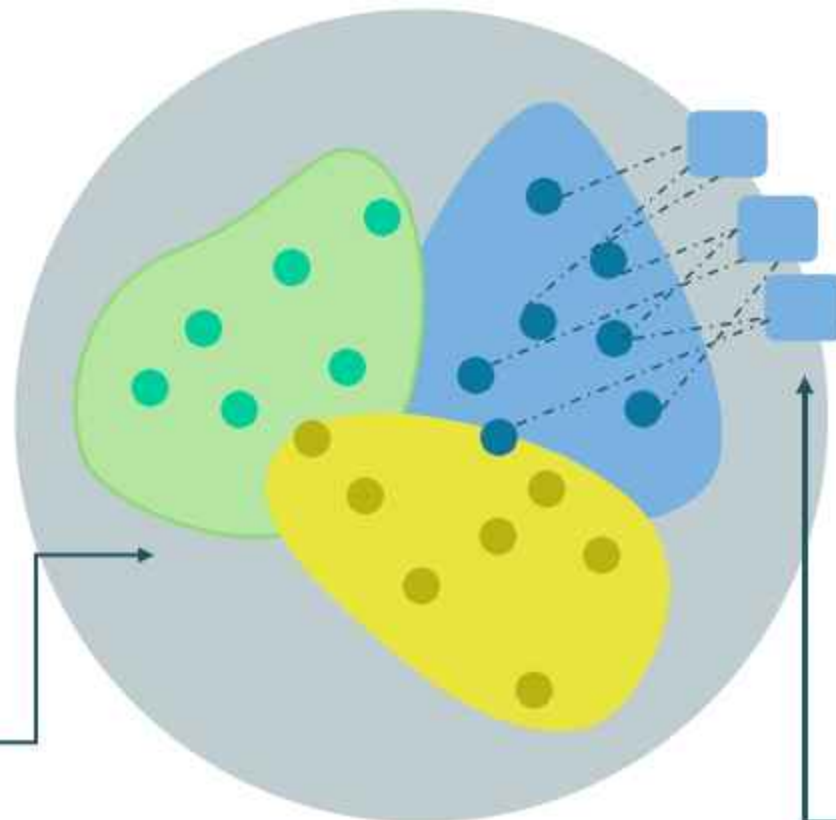
VLM

Input Image
13 x 13 x 14



Sentinel-2 &
UK LiDAR

Vision Encoder
Deep residual 3D CNN
network



Visual-Linguistic Alignment
Supervised Contrastive Learning
(Shared Semantic Space)

Rivers & Lakes

This is a multispectral image of rivers &
lakes

The rivers & lakes contains inland
surface waters (freshwater ecosystems)

Text Encoder
language-model transformer
using Byte Pair Encoding

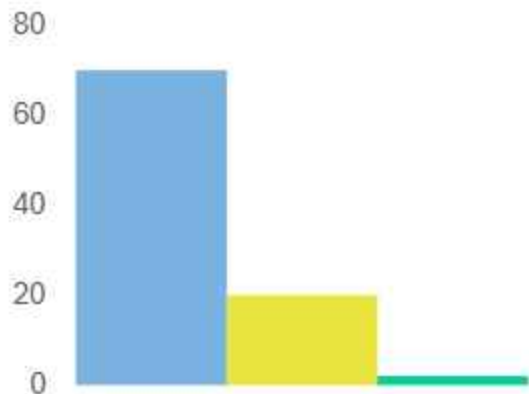
Adapted LDGnet Model (Zhang et al. 2023)

Introducing Probabilistic Priors

1.

Output class probabilities from the VLM:

$$L_j = P(\text{data} \mid UKHAB_j)$$



Raw logits for each class

2.

Living England
Classifications

$$P(UKHAB \mid UKBAP) = \frac{P(UKBAP_j \mid UKHAB) P(UKHAB) \times a_{proj}}{P(UKBAP)}$$

3. Adjust the Loss Function with Bayesian Priors

- Adjust the model's raw predictions (logits) using the Living England's prior:
 - $Adjusted\ Logits = Raw\ Logits + \log(P_{prior})$
 - Loss = CrossEntropy(Adjusted Logits, Ground Truth)

Preliminary Results

Habitat Class	Text Level 1		Text Level 4	
	F1 Score	F1 Score (Probabilistic Prior)	F1 Score	F1 Score (Probabilistic Prior)
Urban	0.81	0.95	0.87	0.97
Grassland	0.94	0.98	0.95	0.99
Rivers & Lakes	0.27	0.81	0.60	0.89
Woodland & Forest	0.76	0.94	0.83	0.96
Cropland	0.92	0.99	0.95	0.99
Sparsely Vegetated Land	0.46	0.91	0.71	0.93
Heathland & Shrub	0.06	0.85	0.49	0.88
Wetland	0.00	0.88	0.24	0.84
Overall Accuracy (weighted)	0.88	0.97	0.92	0.98

Conclusions & Future Work

- Priors from Living England can be used within a VLM loss function to align UKBAP with UKHab, demonstrating that ecological knowledge improves classification performance
- Priors especially helped **rare / under-represented classes** (Wetland, Rivers & Lakes, Sparsely Vegetated).
- Although text granularity improves classification performance, it has less impact once priors are introduced
- Further work:
 - Test **generalisation**: apply model + priors to new regions or unseen habitats.

REFERENCES

- DEFRA. (2023, November 29). *Biodiversity net gain*. GOV.UK. <https://www.gov.uk/government/collections/biodiversity-net-gain>
- Li, X., Wen, C., Hu, Y., Yuan, Z., & Zhu, X. X. (2024). Vision-Language Models in Remote Sensing: Current progress and future trends. *IEEE Geoscience and Remote Sensing Magazine*, 2–36. <https://doi.org/10.1109/MGRS.2024.3383473>
- State of Nature. (2023). *State of Nature*. https://stateofnature.org.uk/wp-content/uploads/2023/09/TP25999-State-of-Nature-main-report_2023_FULL-DOC-v12.pdf
- Snell, L., & Oxford, M. (2021). *Survey of LPAs Ability to Deliver Biodiversity Net Gain in England. Do LPAs currently have the necessary expertise and capacity?*
- UK Habitat Classification Working Group (2018). *UK Habitat Classification – Habitat Definitions V1.0 at <https://ecountability.co.uk/ukhabworkinggroup-ukhab>*
- Zhang, Y., Zhang, M., Li, W., Wang, S., & Tao, R. (2023). Language-Aware Domain Generalization Network for Cross-Scene Hyperspectral Image Classification. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–12. <https://doi.org/10.1109/TGRS.2022.3233885>

Dynamic assignment of photosynthetic parameters using machine learning for modelling primary production from satellite observations

Dr Mayra Rodriguez, Dr David Moffat, Dr Shubha Sathyendranath, Dr Gemma Kulk

Environmental Intelligence Conference

September 2025

Image from Pančić & Kjørboe, 2018



Phytoplankton take up ~50 Gton of carbon each year, playing a crucial role in the global carbon cycling

POLICY RELEVANCE

Improving estimates of phytoplankton-driven carbon fluxes is essential for:

- Understanding the ocean's role in **climate regulation**
- Reducing uncertainty in **carbon budget assessments**
- Informing **climate policy** and Earth system models

Addressing parameter uncertainty in earth observation models for a better estimation of phytoplankton primary productivity

Maximum Photosynthetic
Rate (P_m^B)

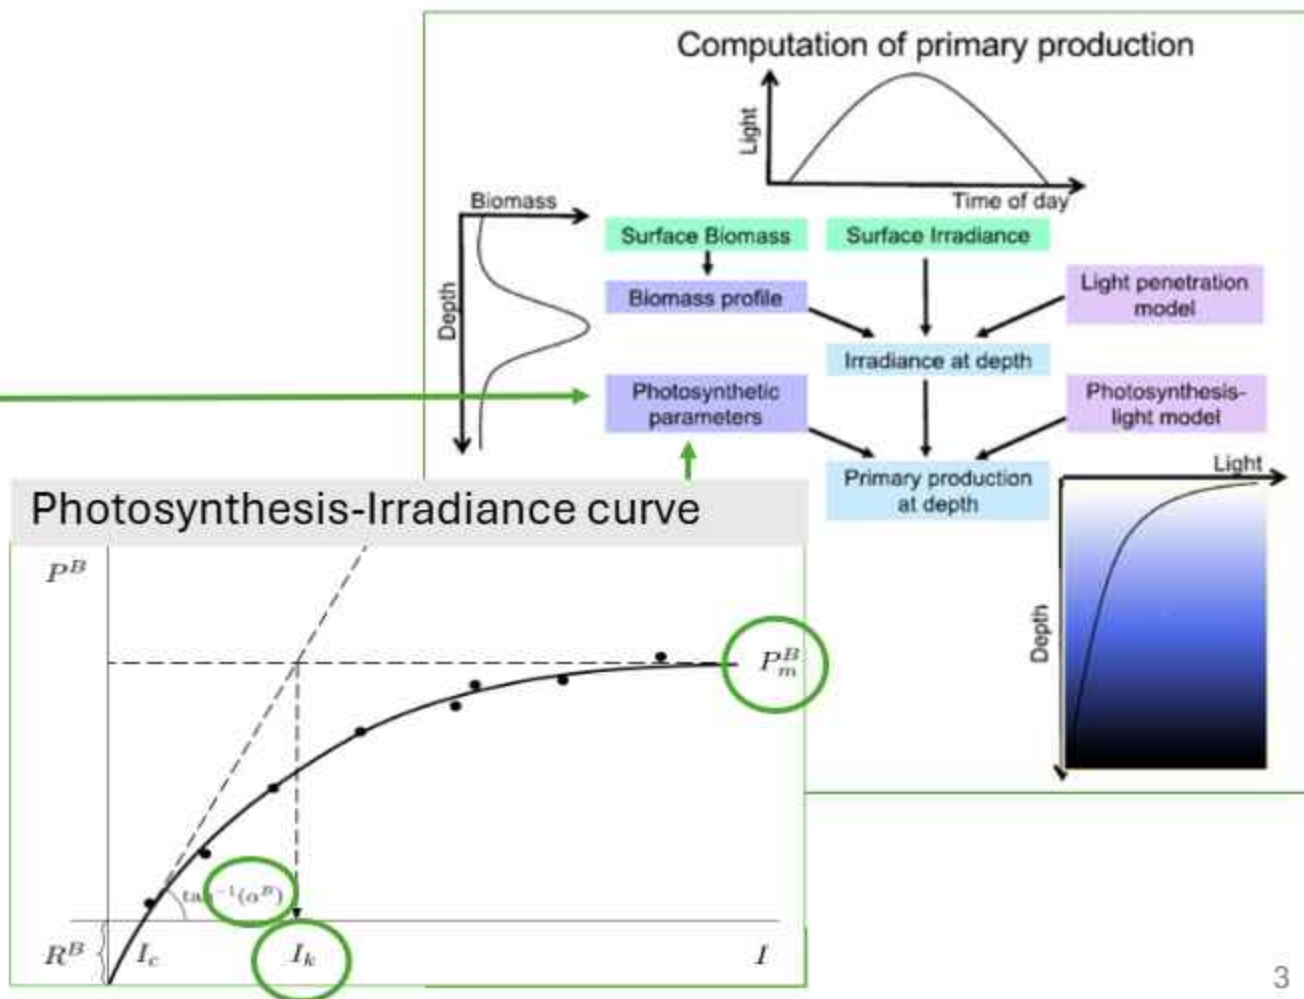
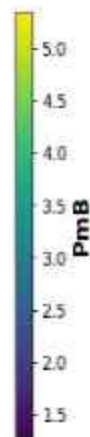
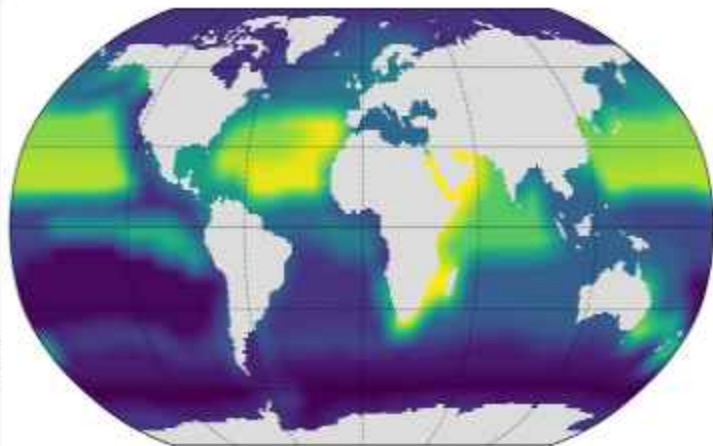
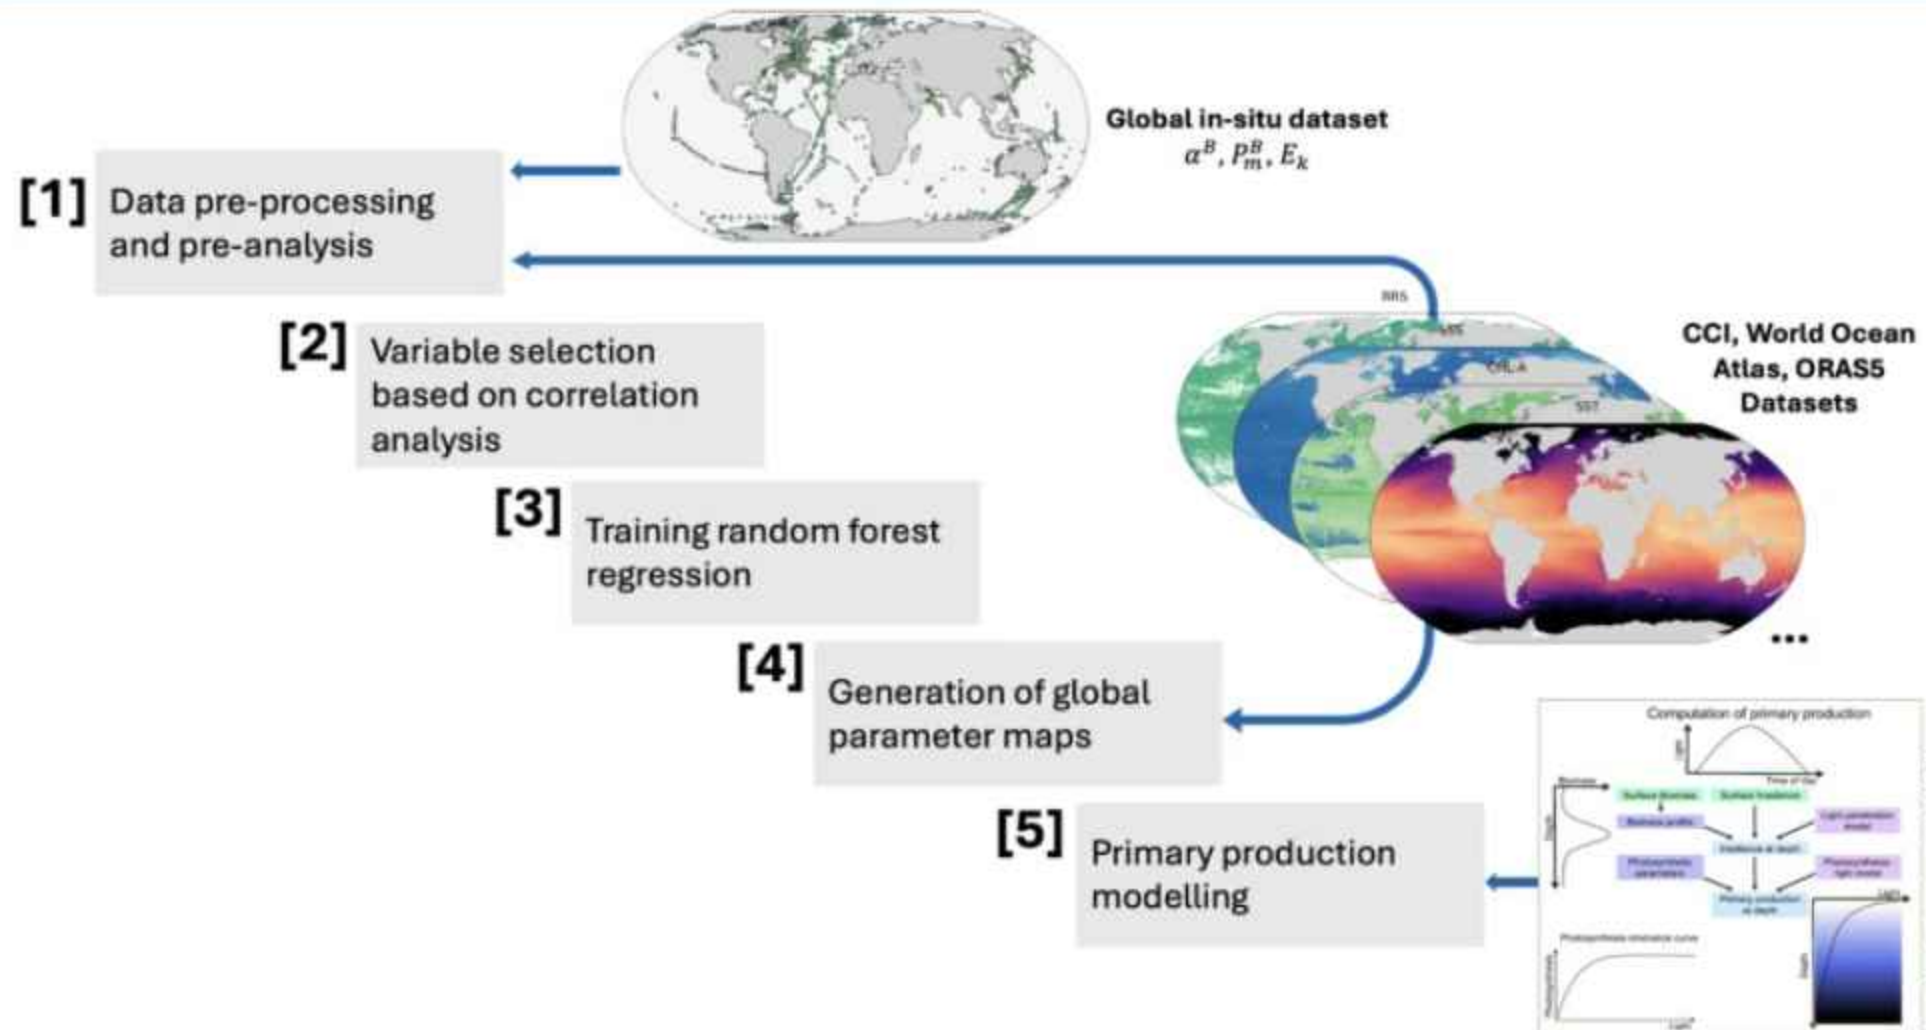
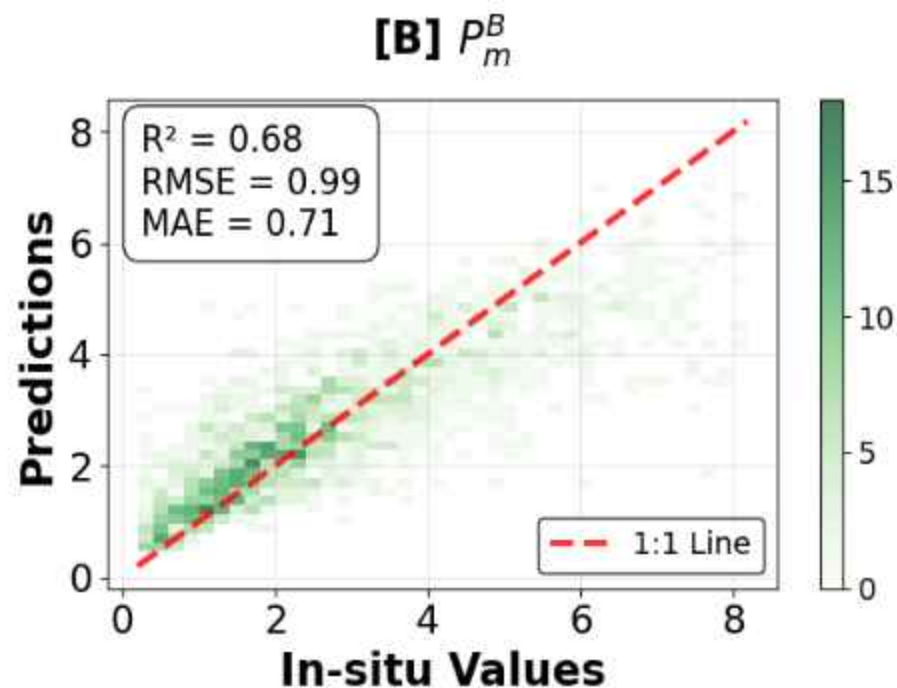
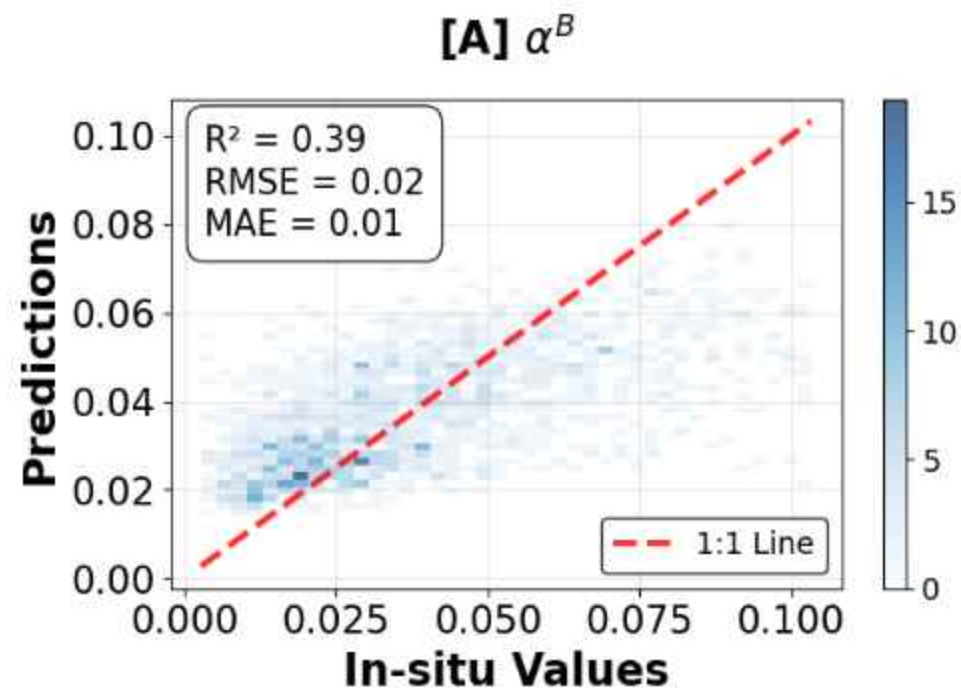


Image from Sathyendranath & Platt, 2020

Dynamic assignment of photosynthetic parameters using random forest regression models



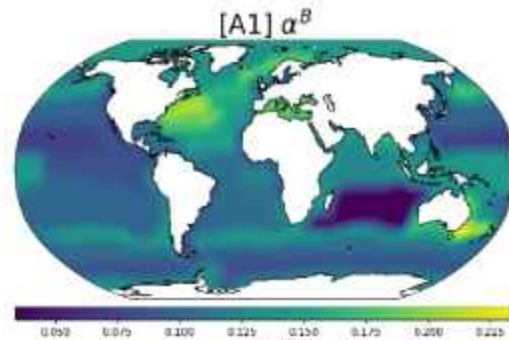
Random forest models show good predictive performance



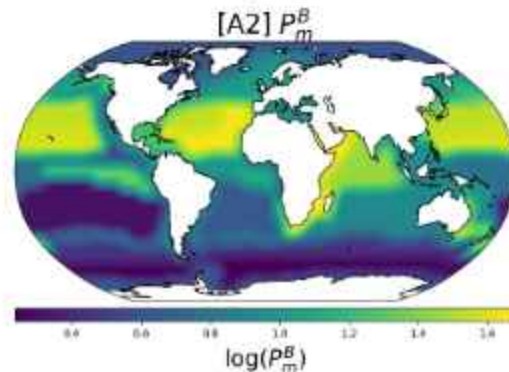
 **Loss Function**
Mean Absolute Percentage Error (MAPE)

Random forest-based global parameter maps provide enhanced spatial and temporal resolution

Current maps used



$\log(\alpha^B)$



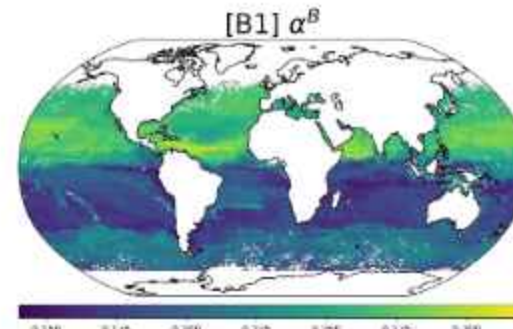
$\log(P_m^B)$

Winter

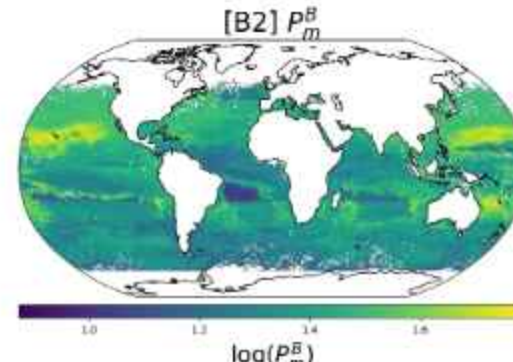
Per province, per season product



Random forest regression maps



$\log(\alpha^B)$



$\log(P_m^B)$

January 2010

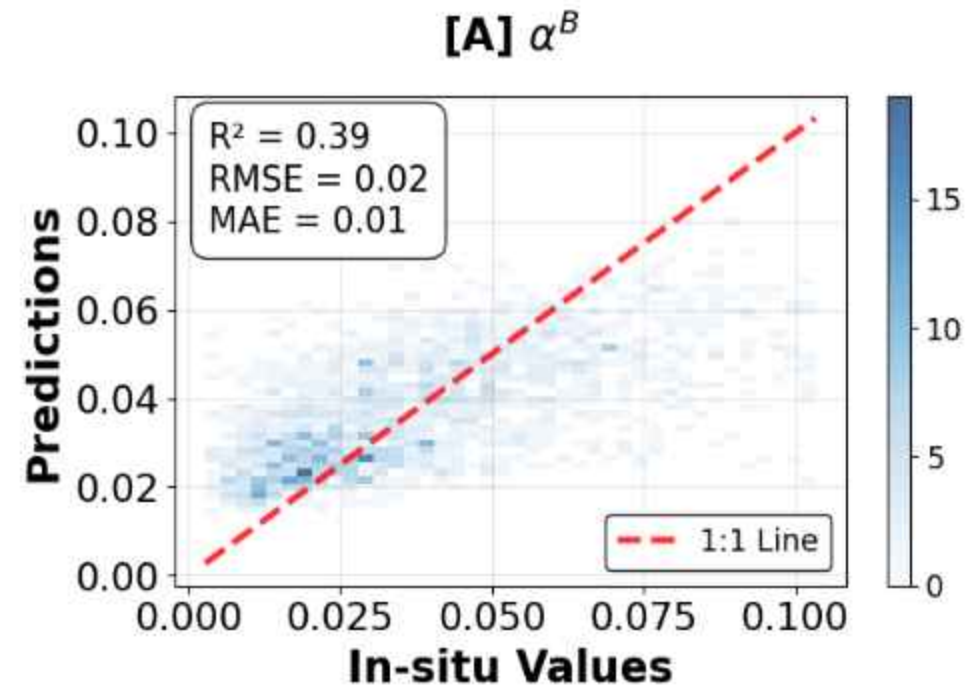
Per pixel, monthly product

Understanding model limitations and prediction errors

Observation density, environmental variability, and model limitations all contribute to prediction error.

[3] Random Forest Bias

Random Forest regression biases predictions toward the mean, reducing accuracy at distribution extremes.



Conclusions

- Machine learning enables creation of global maps of key photosynthetic parameters, enhancing both spatial and temporal resolution.
- Further model refinement and expanded in situ datasets are needed to fully capture local-scale variability and improve model reliability.

Check out our latest publications!
Dr Mayra Rodriguez
Earth Observation Scientist
mro@pml.ac.uk



Thank you!

PML

Plymouth Marine
Laboratory

cbiomes

Simons Collaboration on Computational
Biogeochemical Modeling of Marine Ecosystems



SCOPE

Geospatial foundation models for marine applications

Rémy Vandaele

Centre for Environmental Intelligence, University of Exeter

11 September 2025

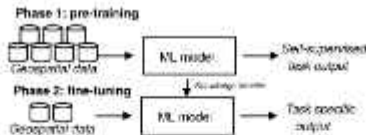
What is a Geospatial Foundation Model?



At use: similar to typical geospatial ML models

- ▶ **Input:** geospatial data (satellite imagery, weather data, etc.)
- ▶ **Output:** Task-specific (land cover classification, weather forecasting, etc.)

What is a Geospatial Foundation Model?



At use: similar to typical geospatial ML models

- ▶ **Input:** geospatial data (satellite imagery, weather data, etc.)
- ▶ **Output:** Task-specific (land cover classification, weather forecasting, etc.)

Key difference: how the model is trained

- ▶ **Phase 1: Pretraining (Self-Supervised Learning)**
 - ▶ Train the model on a **large** dataset with a similar input domain.
 - ▶ The task is often (masked) reconstruction as no labels are required
 - ▶ Learn generalizable features from large scale geospatial data
- ▶ **Phase 2: Fine-tuning**
 - ▶ Adapt the pretrained model to a specific downstream task
 - ▶ Learn task specific features on the smaller task-specific dataset
- ▶ Input domain is the same at pre-training and fine-tuning

IBM/NASA GFM for terrestrial data: Prithvi-EO-2.0

- ▶ Transformer-based geospatial foundation model
- ▶ Pretrained on 1TB of multispectral satellite imagery (HLS product).
 - ▶ Masked auto-encoder task (image reconstruction)
 - ▶ Pre-processed product of Sentinel-2/Landsat 7 images
- ▶ Fine-tuned on various remote sensing applications
 - ▶ Cloud gap imputation, flood mapping, wildfire scar segmentation, crop segmentation

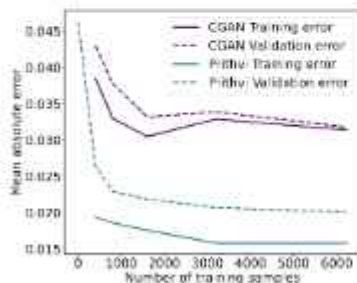


Fig. 7: Comparison of Prithvi on cloud gap filling using a range of training subsamples.

	IoU (fire scar)	F1 (fire scar)
U-Net (DeepLabV3) [81]	71.01	83.05
VIT-base [19]	69.04	81.69
Prithvi (not pretrained)	72.26	83.89
Prithvi (pretrained)	73.62	84.81

Table 3: Prithvi model performance for the segmentation of wildfire scars

- ▶ GFM's outperform non GFM approaches
- ▶ Especially on small datasets

GFM Ocean project

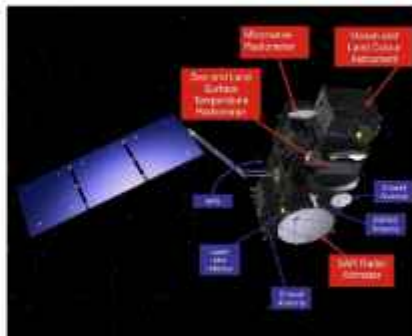


PML

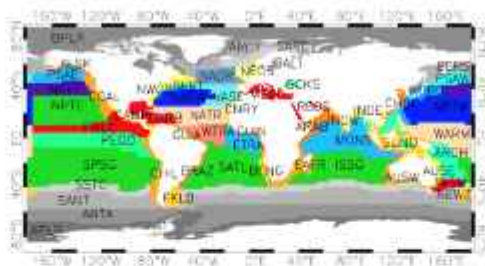
Plymouth Marine
Laboratory

Develop a GFM for marine applications

- ▶ Sentinel-3 data
- ▶ Prithvi model architecture
- ▶ Build a new pre-training dataset and pre-train the network
- ▶ Evaluate its performance on two downstream tasks
 - ▶ Algae bloom detection
 - ▶ Estimation of phytoplankton primary production



GFM pre-training



Masked auto-encoder task

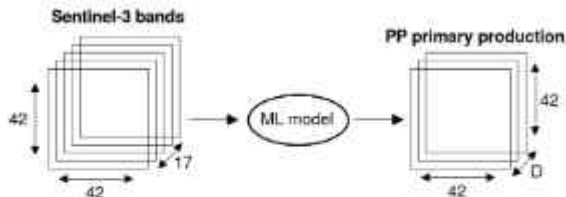
Data considered

- ▶ 16 OLCI water reflectance bands
- ▶ (1 sea-surface temperature band)
- ▶ (Additional encodings)
 - ▶ Coordinates, month, hour
- ▶ 42×42 crop size (300m resolution)
- ▶ < 20% of clouds

Sampling

- ▶ 2017-2021 timeframe
- ▶ Same number of images per Longhurst region (6000)
 - ▶ Similarity of environmental conditions
- ▶ Same number of images per month (≈ 100)
- ▶ $\approx 500,000$ images in total

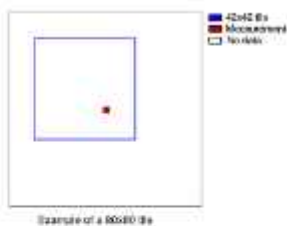
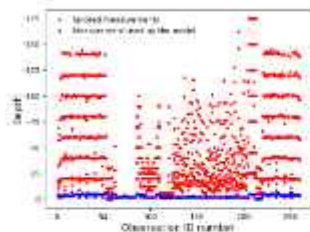
Estimation of phytoplankton primary production



Use the GFM to estimate phytoplankton primary production from Sentinel-3 images

- ▶ Same images/bands than the pre-training
- ▶ Phytoplankton primary production is typically measured in grams of carbon per square meter per day
- ▶ Can be measured at D different sea depths
- ▶ Pixel-wise regression problem

PP dataset collection



1. Collection of phytoplankton primary production measurements from 6 studies
 - ▶ 279 observations
 - ▶ 1 observation = same location and date
 - ▶ Primary production was measured at different depths and different numbers of depths
2. Collection of images of size 80×80 centered around the location of each observation
 - ▶ OLCI | SST | Encodings
 - ▶ Median of all data obtained within ± 3 days
3. The mask consists in a 3×3 pixel square of the **highest** measurement, the rest remains unlabeled

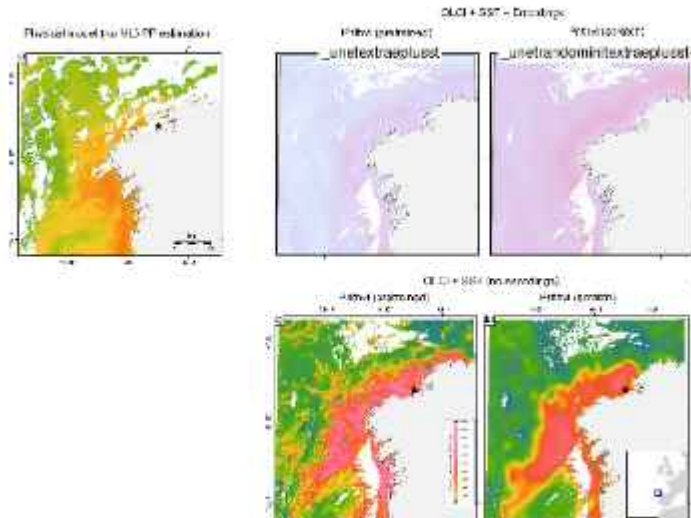
Results on the PP dataset

Methodology

- ▶ 10 42×42 tiles are selected per image
 - ▶ Tile must contain at least 90% of data (cloud, missing value,...)
 - ▶ 120 observations remain
- ▶ Comparison with
 - ▶ Model with (pretrained) and without pre-training (scratch)
 - ▶ Random Forest pixel regression
 - ▶ Null model (outputs the average of the training measurements)
- ▶ Validation of typical hyperparameters using 5 fold cross validation
- ▶ RMSE error criterion

Model	Data	RMSE
Null model		0.5804
Random Forests	OLCI + SST	0.4104
Prithvi (scratch)	OLCI	0.4332
Prithvi (pretrained)	OLCI	0.382
Prithvi (scratch)	OLCI + SST	0.4214
Prithvi (pretrained)	OLCI + SST	0.3924
Prithvi (scratch)	OLCI + SST + Encodings	0.3239
Prithvi (pretrained)	OLCI + SST + Encodings	0.3631

Test on a new area



- ▶ Encodings bring overfitting
- ▶ Scratch approach does not predict the whole 42×42 tile correctly

Conclusions and future steps

Current work and observations:

- ▶ Building of a large Sentinel-3 dataset
- ▶ Pre-training of a geospatial foundation model
- ▶ Evaluation on downstream tasks
 - ▶ Estimation of phytoplankton primary production
 - ▶ Algae bloom detection
- ▶ GFM shows better results
 - ▶ Better validation loss
 - ▶ Generalizes better on the large scale test image

Next steps:

- ▶ Additional experiments on the PP estimation problem
 - ▶ Robustness to noise and missing values
 - ▶ Evaluation on another large scale test image
- ▶ Experiment with new pre-training tasks
- ▶ Publish the paper and models



University
of Exeter

Remote Sensing Single Image Super-Resolution Benchmarking with Transfer Learning Algorithms

Dr Saptarshi Das

*Centre for Environmental Mathematics,
Faculty of Environment, Science and Economy, University of Exeter,
Penryn Campus, Cornwall TR10 9FE, United Kingdom*

Background of Single Image Super-Resolution (SSIR)

- **Single image super-resolution (SSIR)** is a method to increase image quality with limited amount of training dataset.
- Aim is to create a robust SSIR model using the **Deep Convolutional Neural Networks** and **transfer learning**
- Image quality to increase by a factor of **2x** upscale and **4x** upscale size from the low-resolution input image.
- The model has to be generalized across **various landcover** (water, buildings, beach, etc), **time** (day, night), and **weather** (clear, cloud, snow).
- Design **multiple error metrics** (SSIM, PSNR, NRMSE) with visualization to evaluate the super-resolved images quality.



University
of Exeter

Experiment Design & Methods

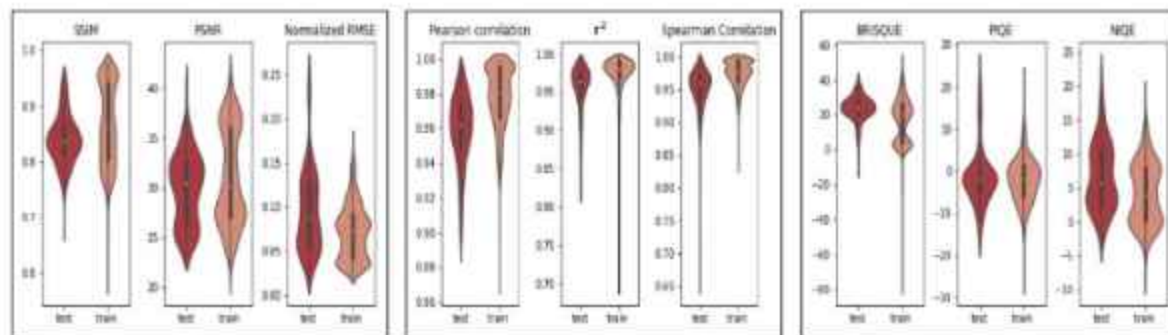
- **Dataset:** size (960x720),
 - **10 different scenes** (Agriculture, Airport, Beach, Buildings, Forest, Land, Parking, Playground, Road, Water),
 - **2 different times** (day, and night), and
 - **3 weather conditions** (clear, cloud, snow)
- **Data Pre-processing:** 1934 images in total, 1720 for training, 214 for testing.
 - Input to the model will be the rescaled and blurred version of the target images
- **Working Principle:** Three different CNN architectures were tested.
 - Difference in Convolution Layer, Batch Norm and PReLU.

Image Comparison Metrics

- Structural Similarities:
$$\text{SSIM} = \frac{((2\mu_x\mu_y + c_1)(2\sigma_{xy} + c_2))}{((\mu_x^2 + \mu_y^2 + c_1)(\sigma_x^2 + \sigma_y^2 + c_1))}.$$
- Peak Signal to noise ratio:
$$\text{PSNR} = 10 \log_{10} \left(\frac{R^2}{MSE} \right).$$
- Blind/Referenceless Image Spatial Quality Evaluator (BRISQUE)
- Natural Image Quality Evaluator (NIQE)
- Perception-based Image Quality Evaluator (PIQE)

- Pearson/Spearman Correlation Coefficient, R^2
- Normalized root mean squared error (NRMSE)

SSIR 2x and 4x Model Performance



- Coefficient of determination (R^2), Pearson/Spearman correlation coefficients are expected **to be around zero**
- SSIM, PSNR fluctuates
- NRMSE, BRISQUE, PQIE, NQE are expected to be distributed **around zero**.

Fig. 1. SSIR 2x model error distribution violin plots

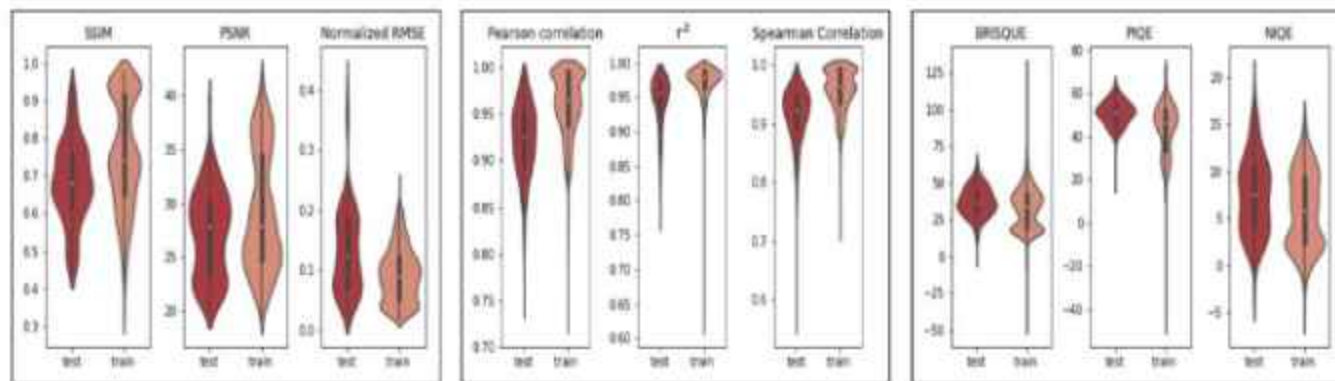


Fig. 2. SSIR 4x model error distribution violin plots

Difference Between Ground Truth and Super-resolved



- For Beach (2x) and Airport (4x) both images, the errors are low.
- The 4x upscaling has slightly more error.

Fig. 3. a) Ground truth, b) Predicted image, c) Difference in pixel between ground truth and predicted 2x upscaling.



Fig. 4. a) Ground truth, b) Predicted image, c) Difference in pixel between ground truth and predicted 4x upscaling.



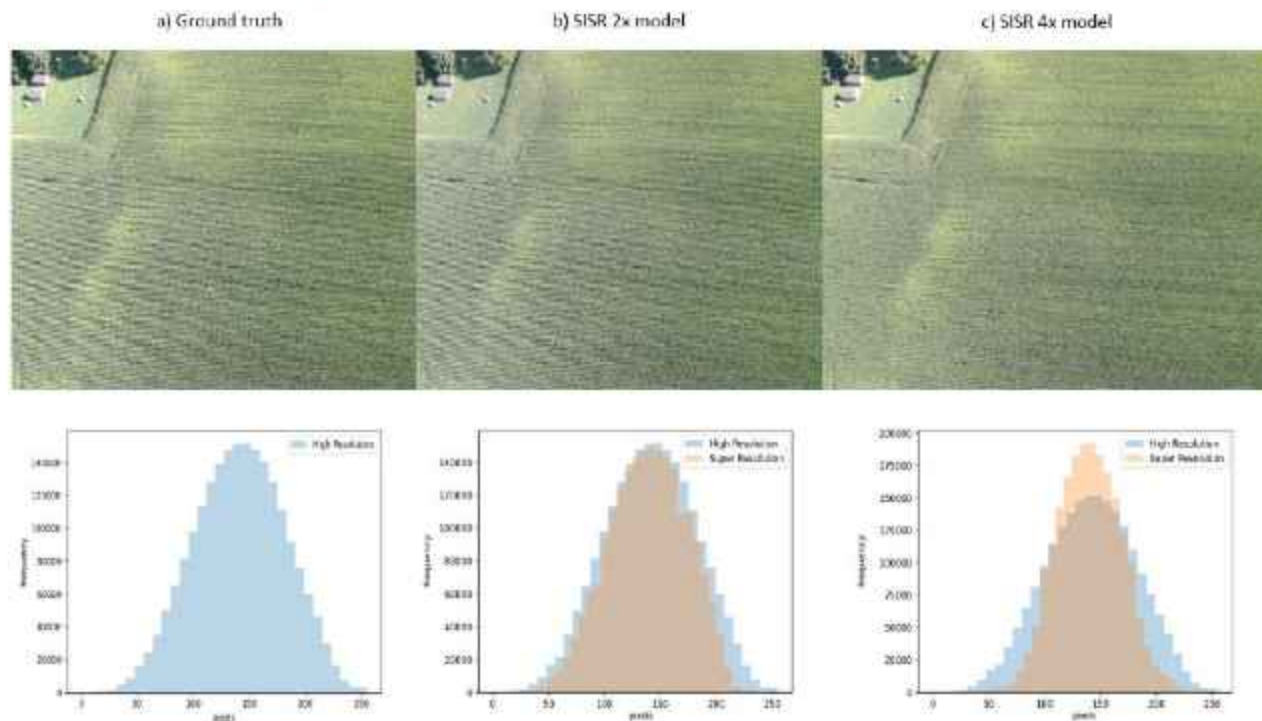
PSNR/SSIM 2x and 4x Comparisons on 10 Types of Scenes

Image scene	Average of SSIM		Average of PSNR
Agriculture	0.8108	↓	27.7904
Airport	0.8879	↗	31.4834
Beach	0.9196	↑	33.0430
Buildings	0.8552	↘	29.5510
Forest	0.8396	→	30.9009
Land	0.8802	↑	33.6734
Parking	0.8702	→	30.2096
Playground	0.8698	↗	31.8662
Road	0.8901	↑	33.0660
Water	0.8862	↗	32.1204

Image scene	Average of SSIM		Average of PSNR
Agriculture	0.5844	●	24.9742
Airport	0.7608	●	28.6035
Beach	0.8781	●	33.4427
Buildings	0.7466	●	27.5010
Forest	0.6963	●	28.7400
Land	0.7921	●	32.2157
Parking	0.7643	●	28.2997
Playground	0.8000	●	30.9382
Road	0.8232	●	31.7126
Water	0.8141	●	30.4455

Fig. 5. SISR model average SSIM and PSNR a) 2x upscaling b) 4x upscaling

Pixel Distribution Comparison between SSIR and the Truth



University
of Exeter

Fig. 6. Row 1 has ground truth, 2x upscaled, and 4x upscaled image. Row 2 shows the pixel histogram plot for the same ground truth, 2x upscaled, and 4x upscaled image.

Inferential Statistics

- Weather Hypothesis

Null: The SSIM error mean for weather conditions clear, cloudy, and snowy does not differ from the group mean.

TABLE I
WEATHER HYPOTHESIS STATISTICS

<i>Model</i>	<i>Name</i>	<i>Df</i>	<i>SumSq</i>	<i>MeanSq</i>	<i>Fvalue</i>	<i>pr(> F)</i>
2x upscale	Weather	2	0.026	0.0129	2.5680	0.0769
	Residuals	1931	9.6790	0.0050		
4x upscale	Weather	2	0.35	0.1748	8.525	0.0002
	Residuals	1931	39.60	0.0205		

- Time Hypothesis:

Null: The SSIM error means for images taken during the day and night do not differ from the group mean.

TABLE II
TIME HYPOTHESIS STATISTICS

<i>Model</i>	<i>Name</i>	<i>Df</i>	<i>Sum - Sq</i>	<i>Mean - Sq</i>	<i>F - value</i>	<i>pr(> F)</i>
2x upscaling	Time	1	0.015	0.0146	2.916	0.0879
	Residuals	1932	9.690	0.0050		
4x upscaling	Time	1	0.05	0.504	2.441	0.118
	Residuals	1932	39.90	0.0206		



Visual Comparison of Results between CNNs



Fig. 7. 2x upscaling: Ground truth, SISR, EDSR, LapSRN, and FSRCNN



Fig. 8. 4x upscaling: Ground truth, SISR, EDSR, LapSRN, and FSRCNN

Comparison of Error Metrics

TABLE III
AVERAGE ERROR METRICS WITH UPSCALING FACTORS 2X AND 4X.

<i>Upscale</i>	<i>Models</i>	<i>SSIM</i>	<i>PSNR</i>	<i>NRMSE</i>	<i>Pearson's – correlation</i>	<i>R²</i>
2x	SISR model	0.84098	29.7897	0.0952	0.9590	0.9539
	EDSR	0.8335	30.0897	0.0916	0.9591	0.9593
	LapSRN	0.8287	29.9039	0.0934	0.9574	0.9583
	FSRCNN	0.8377	30.1546	0.0905	0.9596	0.9591
4x	SISR model	0.6845	27.0929	0.1348	0.9214	0.9414
	EDSR	0.6507	26.8456	0.1369	0.9131	0.9444
	LapSRN	0.6273	26.5504	0.1413	0.9072	0.9423
	FSRCNN	0.6242	26.4633	0.1428	0.9056	0.9421



University
of Exeter

Enhanced Deep Residual Network (EDSR), Fast Super-Resolution Convolutional Neural Network (FSRCNN), and Laplacian Super-Resolution Network (LapSRN)

Limitations and Future Scope of Work

- Accuracy changes across different type of scenes. More examples in training set may help increasing the accuracy (e.g. agriculture, forest).
- Generative Adversarial Network (GAN) architecture may be useful instead of CNNs.
- Beside remote sensing other application areas are:
 - Surveillance and forensics,
 - CCTV camera,
 - SDTV to HDTV.



Thank you for your attention.

Questions?

Reference:

Thiruppandiaraj, E. and Das, S., 2023, December. Remote Sensing Single Image Super-Resolution Benchmarking with Transfer Learning Algorithms. In *2023 7th International Conference on Electronics, Materials Engineering & Nano-Technology (IEMENTech)* (pp. 1-6). IEEE, doi:

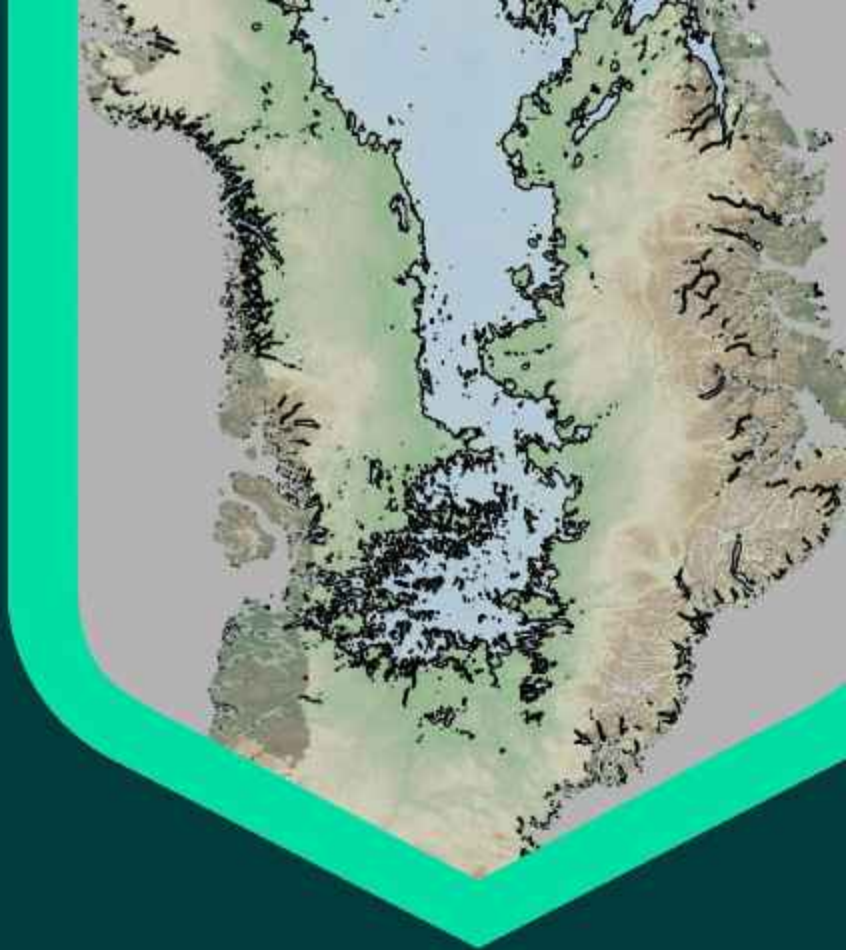
<https://doi.org/10.1109/IEMENTech60402.2023.10423556>.



University
of Exeter

A Quantile Regression forest estimate of Greenland's subglacial topography

Steven Palmer & Charlie
Kirkwood



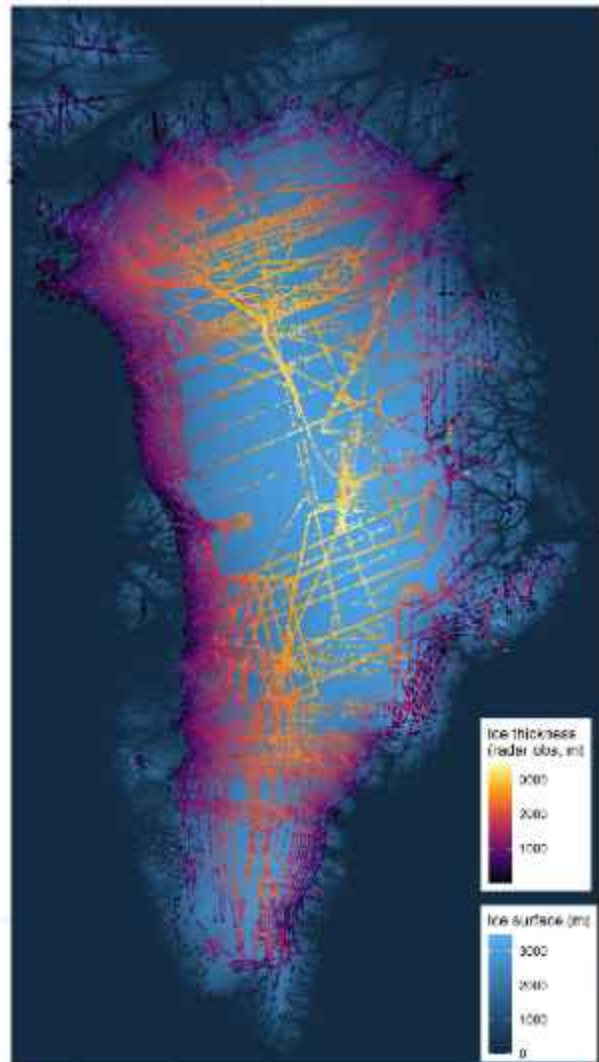
Why is ice sheet bed topography important?

Greenland's ice represents 7m potential contribution to global mean sea level – (~10 x more ice in Antarctica – but more challenging target!)

Models used to make projections of ice sheet behaviour require accurate and continuous information about sub-ice bed elevations

~23 million observations of ice depth collected along flight lines by airborne radar surveys (NASA Operation Ice Bridge, CReSIS, NERC-funded projects etc), but data collection highly anisotropic

Previous attempts to derive sub-ice topography contain artefacts and do not robustly estimate uncertainty (Morlighem et al., 2017)

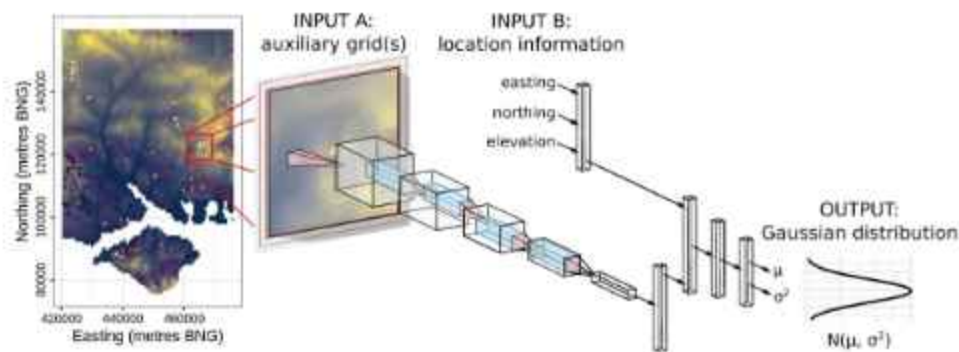


Improved interpolation of RES observations and their uncertainties

Goal: Predict bed elevation in unobserved regions and estimate uncertainty.

How: Train a **ML model** on existing RES and auxiliary information e.g. surface velocity, surface elevation to **predict bed topography**, while also **quantifying the prediction interval**.

Result: Improved knowledge of bed elevations between flightlines and improved uncertainty estimates.



Kirkwood et al. (2022)

Approach

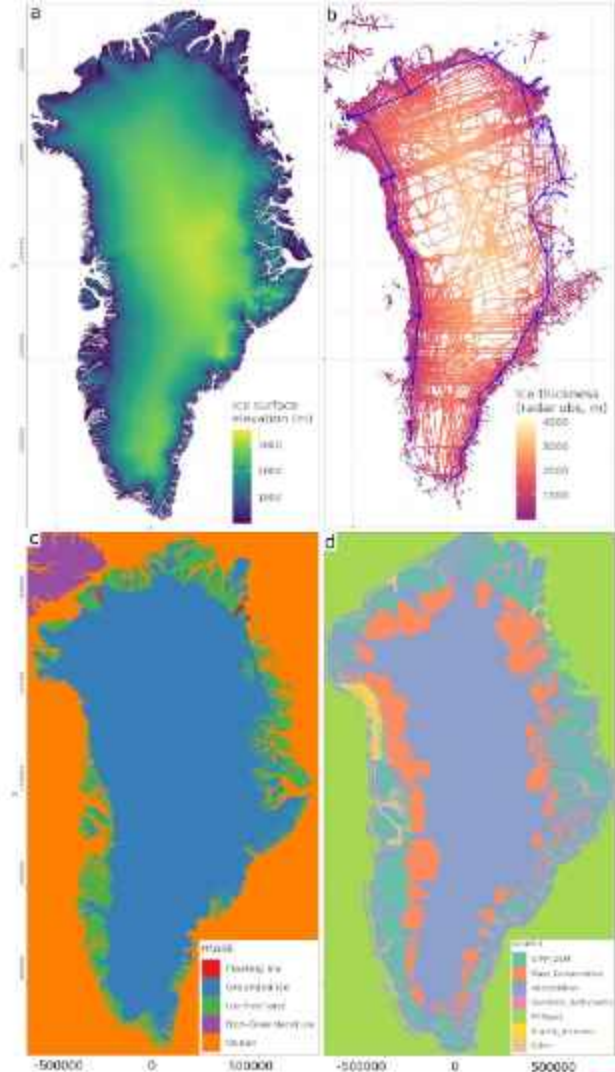
We used the same RES observations used by BedMachine v3 (Morlighem et al, 2017).

Mostly NASA OIB data processed by CReSIS (Leuschen et al., 2010)

BMv3 elevations derived via several approaches in contrast to our new method

Primarily **kriging interpolation** for interior ice and **mass conservation** flow modelling near the ice margin

PROMICE radar observations (Sorenson et al., 2018) used as 'held out' observations to compare the performance of our QRF with BMv3



QRF input features (covariates)

Our QRF models bed elevation as a function of *position in space* and:

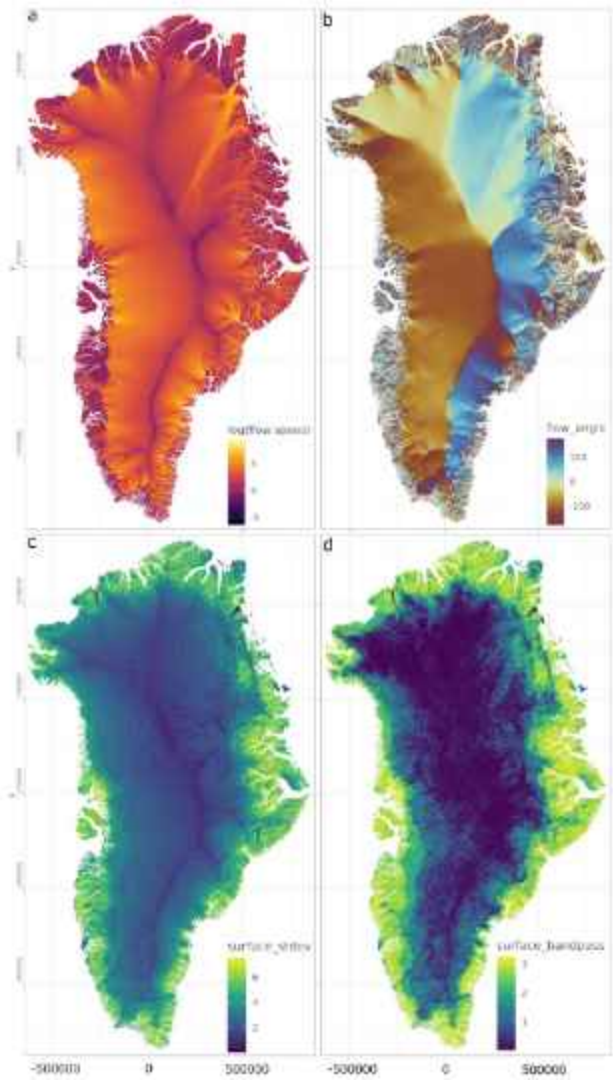
- ice surface elevation,
- ice surface flow speed,
- ice surface flow angle ('circular' representation uses sin and cosine)
- ice surface roughness
- ice surface bandpass (intensity of 'mid frequency' roughness)

Position in space is provided to the QRF as 16 features. These are Easting and Northing, but also rotations every 22.5 degrees, so that 16 compass angles are provided to the model. This makes the interpolation smoother, as spatial decision thresholds are not forced to only be aligned North-South or East-West.

Each tree is provided with all these features and grows by placing decision thresholds optimally on whichever feature will minimise variance in the daughter nodes at each split.



University
of Exeter



QRF predictions

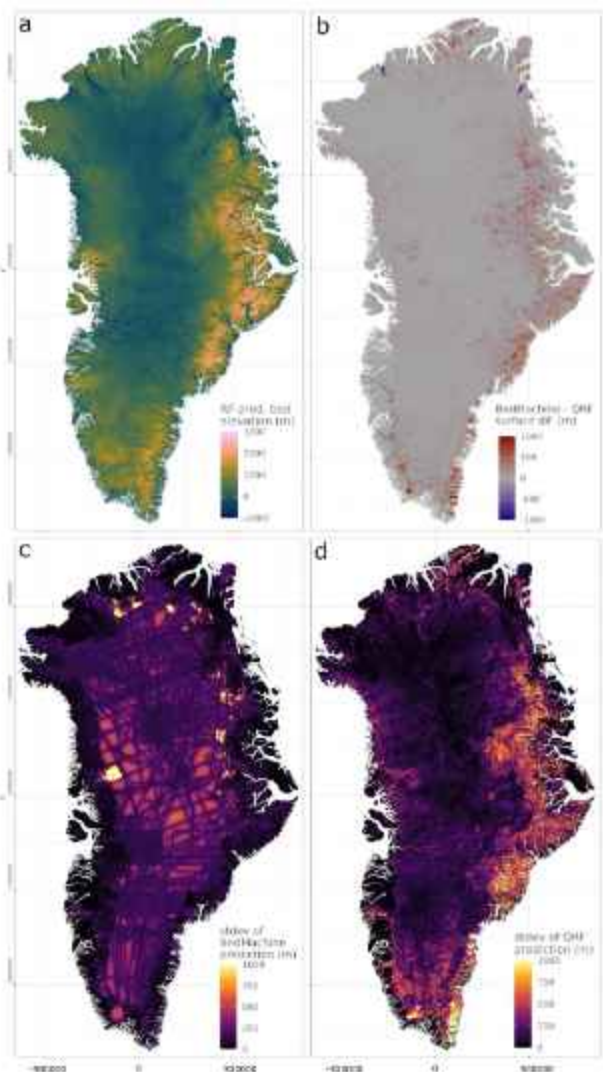
We have derived a new, highly computationally efficient, bed elevation dataset for Greenland (a)

Comparisons with BMv3 (b) show that there is close agreement for much of the ice sheet

Our estimated ice volume is $3.01 \pm 0.004 \times 10^6 \text{ km}^3$

BMv3 estimates $2.99 \pm 0.021 \times 10^6 \text{ km}^3$

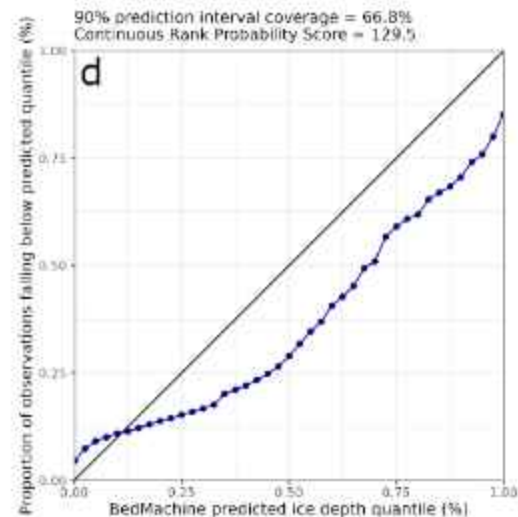
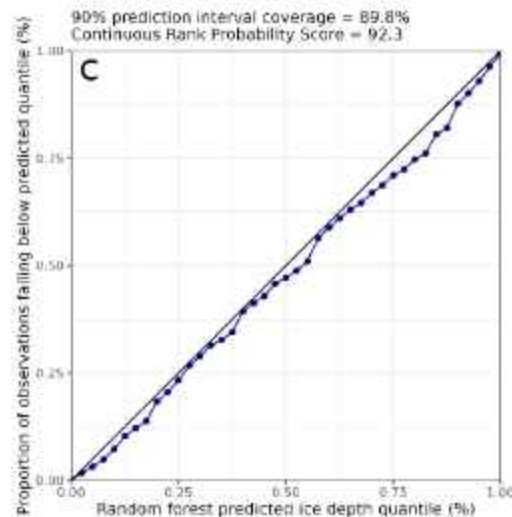
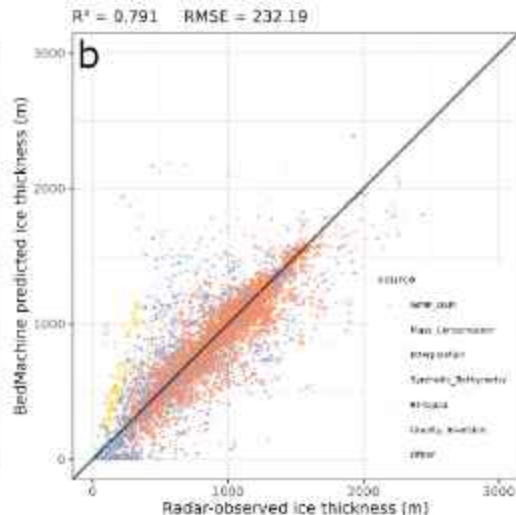
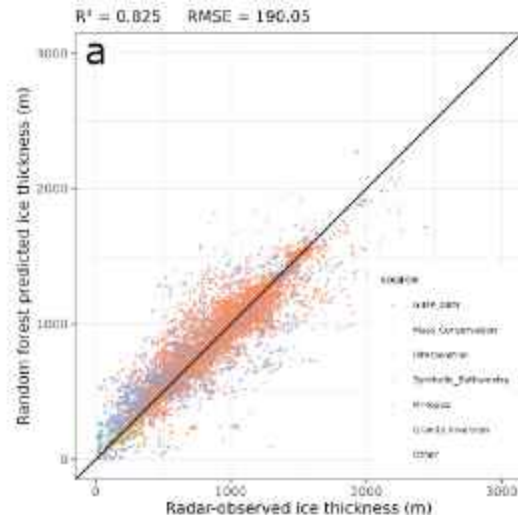
Estimates of uncertainty show very different patterns (c & d)



Comparisons with held-out data

The QRF performs ~18% better than BMv3 when compared with the held out PROMICE data

However, comparison is complicated by the fact that BMv3 is derived through a combination of approaches



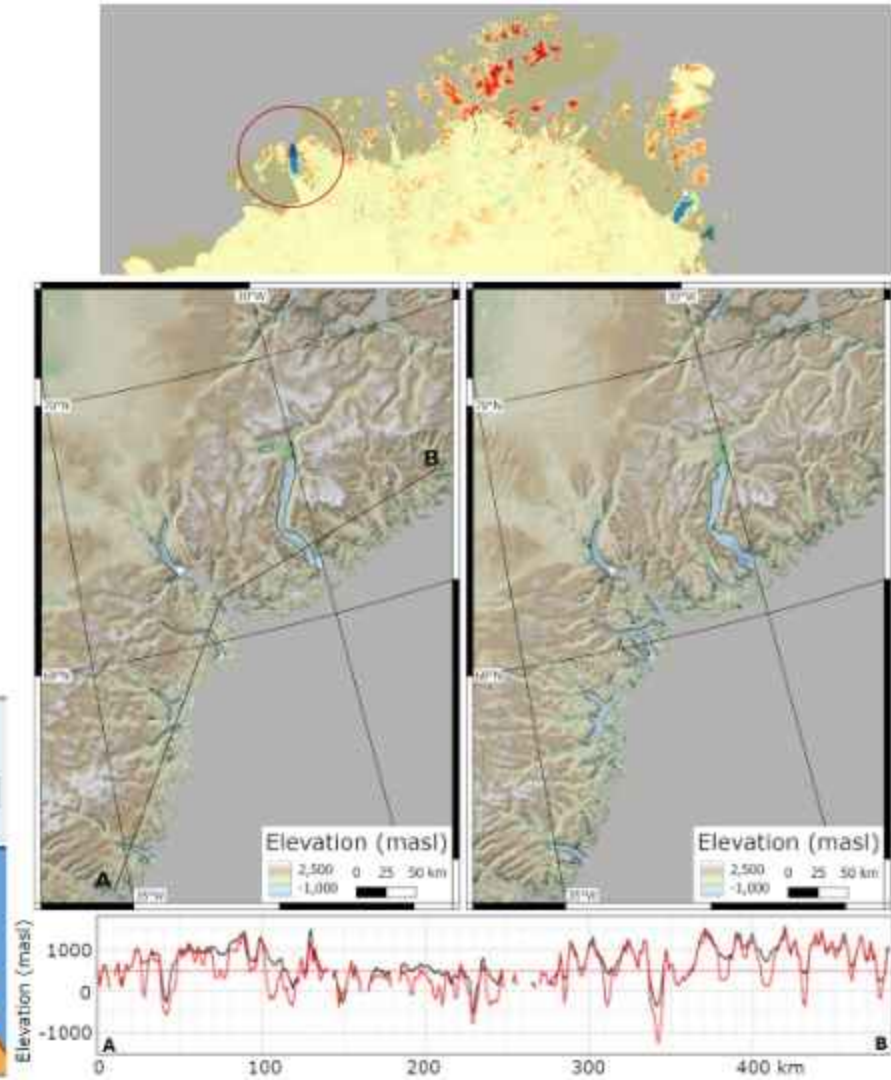
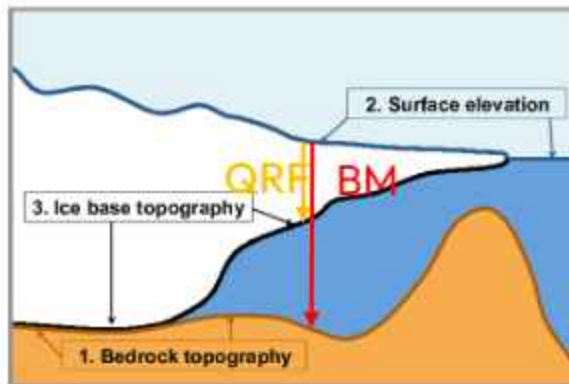
Floating ice

- The QRF *overestimates* bed elevations where ice is floating (e.g. Petermann Glacier and 79N glacier)
- The QRF has no way of knowing that the ice is not in contact with the bed
- We want to explore how this could be improved using the ice mask and gravity (and/or surface slope) as an input feature

In other areas, the QRF predicts deeper ice!



University
of Exeter



References

Breiman, L. (2001). Random forests. *Machine learning*, 45, 5-32.

Kirkwood, C. et al. (2022). Bayesian deep learning for spatial interpolation in the presence of auxiliary information. *Mathematical Geosciences*, 54(3), 507-531.

Leuschen, C. et al. (2010 updated 2016) Icebridge records of ice thickness, version 1. Boulder, CO: *NASA National Snow and Ice Data Center Distributed Active Archive Center*

Morlighem, M. et al. (2017). BedMachine v3: Complete bed topography and ocean bathymetry mapping of Greenland from multibeam echo sounding combined with mass conservation. *Geophysical research letters*, 44(21), 11-051.

Palmer, S. J. et al. (2025a) A quantile regression forest estimate of Greenland's subglacial topography, *Journal of Glaciology* pp. 1 – 25 DOI: <https://doi.org/10.1017/jog.2025.10071>

Palmer, S.J. & Kirkwood (2025b) Harnessing Machine Learning to Improve Ice Sheet Bed Mapping, *Philosophical Transactions of the Royal Society A*, in review.

Sørensen, et al (2018). Circum-Greenland, ice-thickness measurements collected during PROMICE airborne surveys in 2007, 2011 and 2015. *GEUS Bulletin*, 41, 79-82.



University
of Exeter

Abstract geometric lines in the top left corner, consisting of several thin, light brown lines that intersect to form various polygons and shapes, creating a complex, layered effect.

WHEN FICTION REFLECTS FICTION: CONTRARIAN CLAIMS OF CLIMATE CHANGE IN ENTERTAINMENT MEDIA

Francisco Gonzalez, Chico Camargo, Ranadheer
Malla, Travis Coan

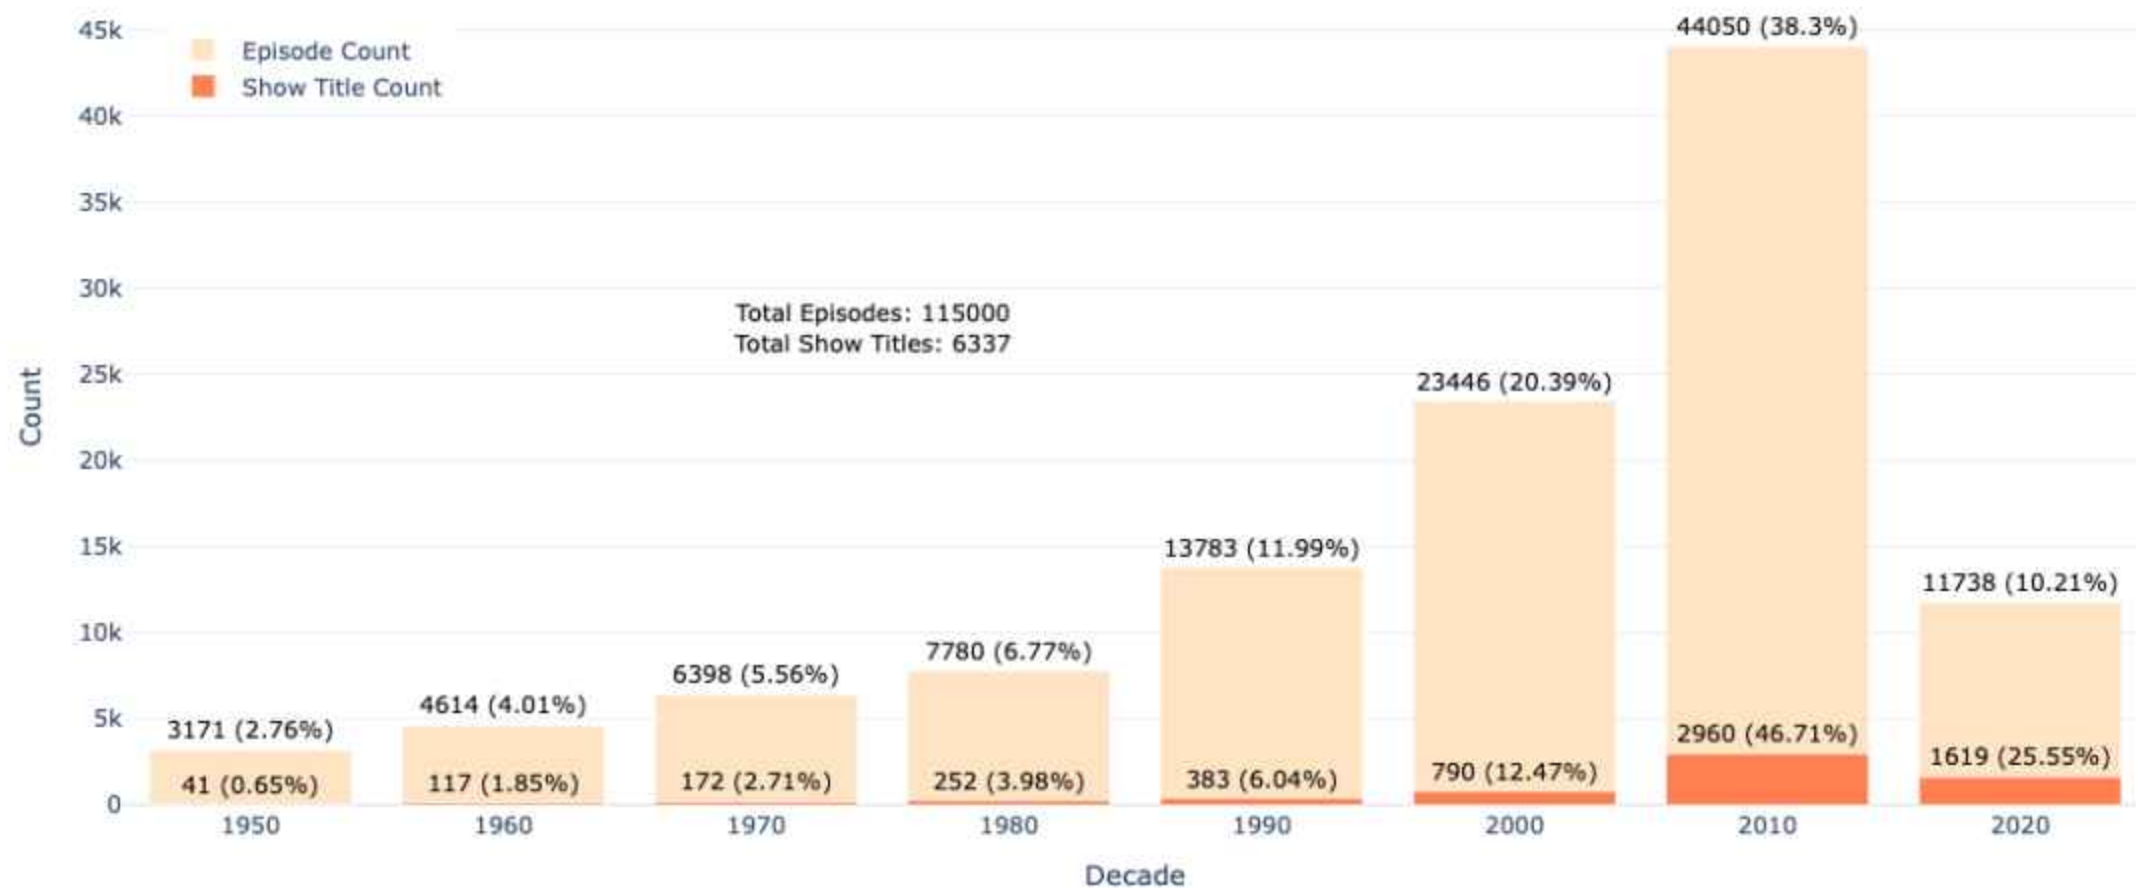
University of Exeter

OBJECTIVE

Identify comments about climate change within scripts of films and television shows, examining whether or not these comments contain contrarian claims.



Episode and Show Title Count by Decade



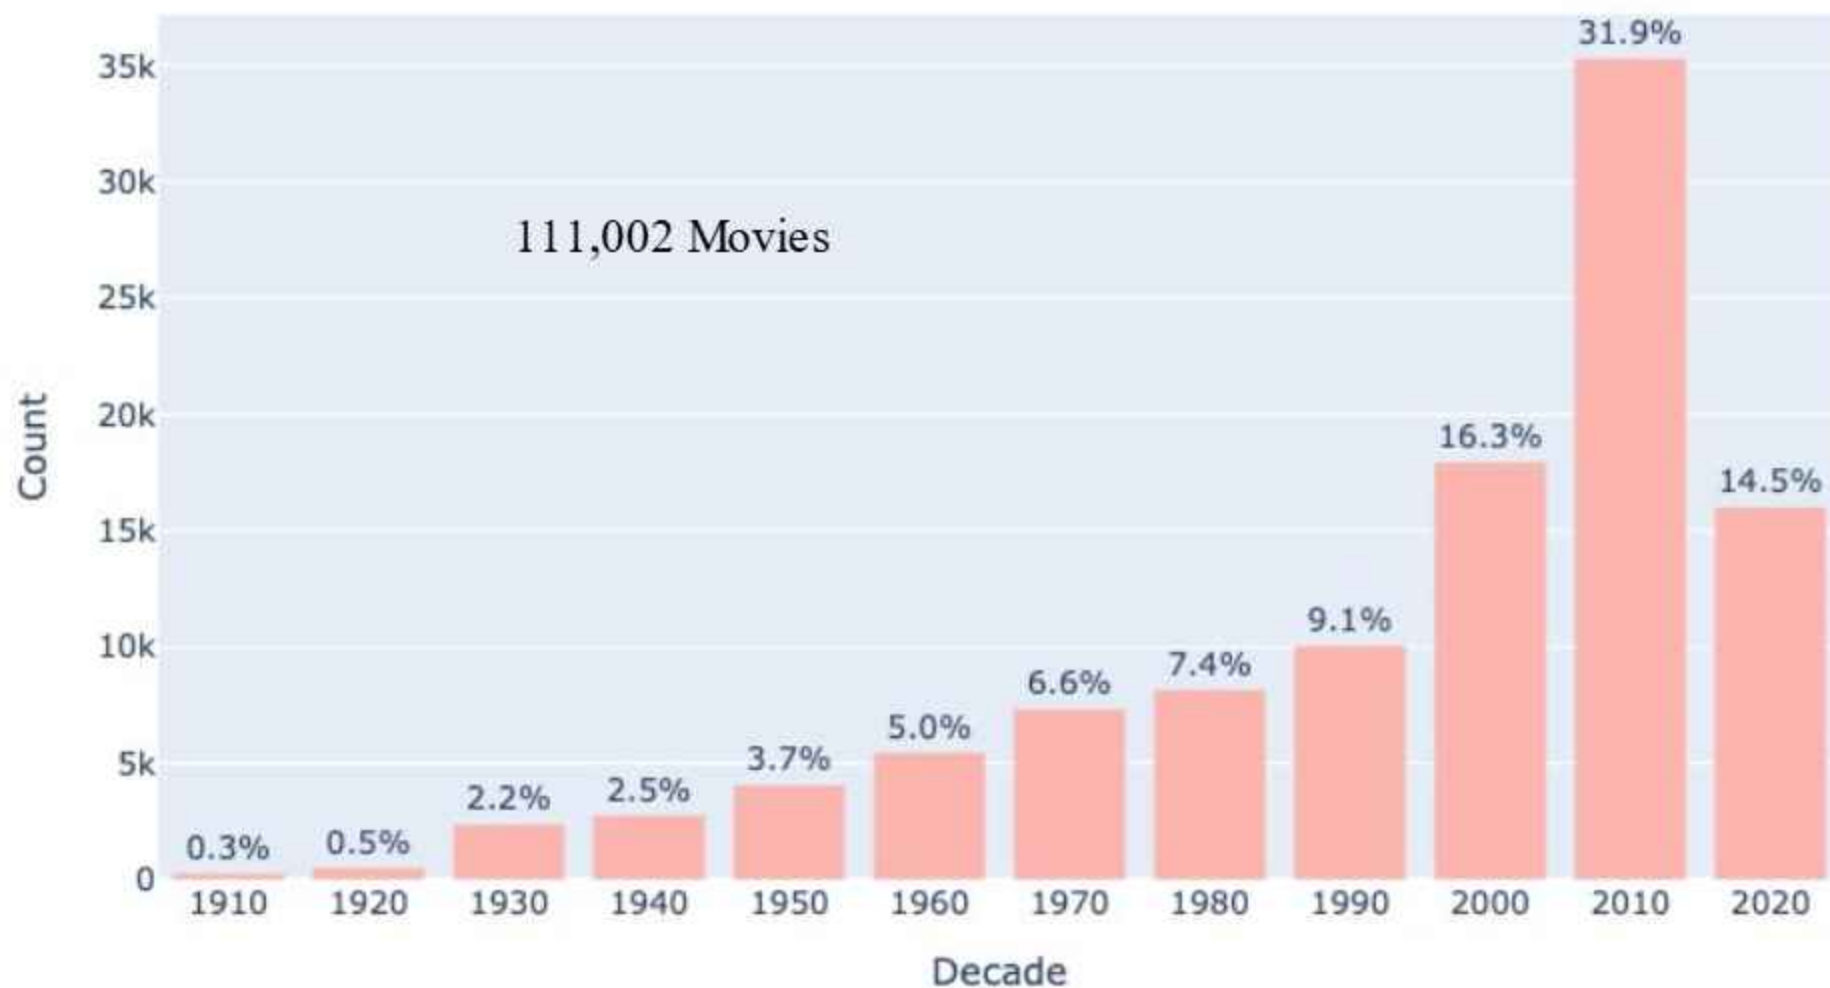
TV SHOW DATA

20XX

Pitch Deck

4

Movie Count by Decade



MOVIE DATA

20XX

Pitch Deck

5

KEYWORD FILTERING

Scripts with mentions

TV Episodes:	1044
Movies:	1397
<i>Total:</i>	<i>2441</i>

Number of Mentions

Keyword	Movies	TV Shows	Total
Climate Change	1225	835	2060
Global Warming	1184	740	1924
Climate Crisis	45	26	71
Total	2454	1601	4055

Average of mentions per script: 1.66

Table 3: Individual count of climate-related keyword mentions across movies and TV shows, excluding climate documentaries.

CARDS RESULTS

CARDS identified a total of **133** claims across **95** scripts (3.3%TV shows - 3.9%Movies)



VALIDATION

124 of the predicted contrarian claims were true contrarian claims (93%)
9 are not (7%)

F1 = **0.88** (overall)



MODEL'S RESPONSE

2) Looking for climate-related claims:

- **The only climate-related claim is the boy saying "Global warming is fake"**
- **However, this is presented as part of a fictional story**
- The context is about a boy's influence over a girl, **not making a serious climate claim**
- The statement is not being endorsed or argued for, just presented as part of a character's dialogue

3) Checking against the codebook:

- The text doesn't make any actual claims about climate change
- It's not arguing for or against any climate positions
- The climate statement is merely a character's line in a story
- None of the other categories in the codebook apply to this narrative

4) Conclusion:

While the text contains a character who makes a climate-related statement, it doesn't actually make any substantive claims about climate change itself. The narrative is focused on character dynamics rather than climate arguments. Therefore, this falls under the category of "No relevant claim detected" (0_0_0).

TOP CONTRARIAN CLAIMS

Label	Description	Freq.
1_0_0	Global warming is not happening.	34
6_2_0	Climate change (science or policy) is a conspiracy (deception).	28
1_3_0	We are experiencing cold weather, therefore climate change is not happening.	20
6_1_0	Climate proponents are alarmist/wrong/political/biased/hypocritical (people or groups).	17
3_0_0	The impacts of climate change will not be bad and might even be beneficial.	12
6_1_3	Politicians/government/UN are alarmist/wrong/political/biased.	12
4_1_5	Climate regulation limits individual liberty, freedom, and undermines capitalism.	11

Table 6: Top verified contrarian categories and their frequencies across the scripts. This table only shows the classes with double-digit frequency.

CONCLUSIONS SO FAR



MENTIONS ARE VERY
LOW



RATIO OF MENTIONS
AND CONTRARIAN
CLAIMS



SCRIPTS WITH
CONTRARIAN CLAIMS ALSO
HAD PRO-CLIMATE CLAIMS



University
of Exeter



Met Office

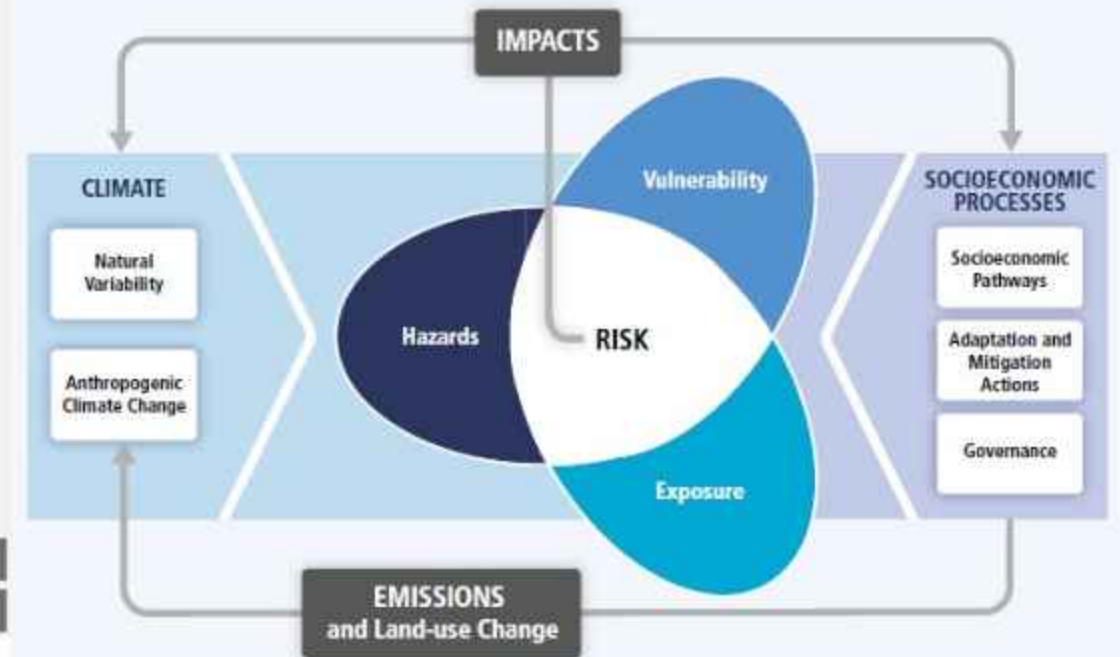
Upscaling Climate Services Using Knowledge Graph and LLM

Dr Hailun Xie
Arjun Biswas
Dr Hang Dong
Prof Hywel Williams
Dr Anrijs Abele
Dr Fai Fung

1. Climate Services

Climate services are the provision and use of **climate data, information and knowledge** to assist decision-making.

Climate-Related Risks, Opportunities, and Financial Impact



Challenge of climate services: **a huge demand in the society, not enough experts**

2. HACID Knowledge Graph

HACID - human and artificial collective intelligence for decision making

HACID KG – Ontology → Instances

Table 3. List of data sources for the climate services

Data	Source	Numerosity	How	Format
CMOR Tables	GitHub	<ul style="list-style-type: none">1,273 MIP variables (physical variables specialised by masking/aggregation methods)2,068 CMOR variables (further specialised by time granularity)90 units of measure	Git	JSON
CVs	GitHub	<ul style="list-style-type: none">49 institutions134 models	Git	JSON
CMIP5 Datasets	ESGF (e.g., CEDA node)	<ul style="list-style-type: none">152 simulations6,365 datasets	API	JSON
CORDEX Domains	GitHub	<ul style="list-style-type: none">14 domains	Git	CSV
CORDEX Datasets	ESGF (e.g., CEDA node)	<ul style="list-style-type: none">1,330 dynamical downscaling runs162,191 datasets	API	JSON
Climdex	GitHub , climdex.org	<ul style="list-style-type: none">63 indices9 categories	Git, Web	Markdown, HTML

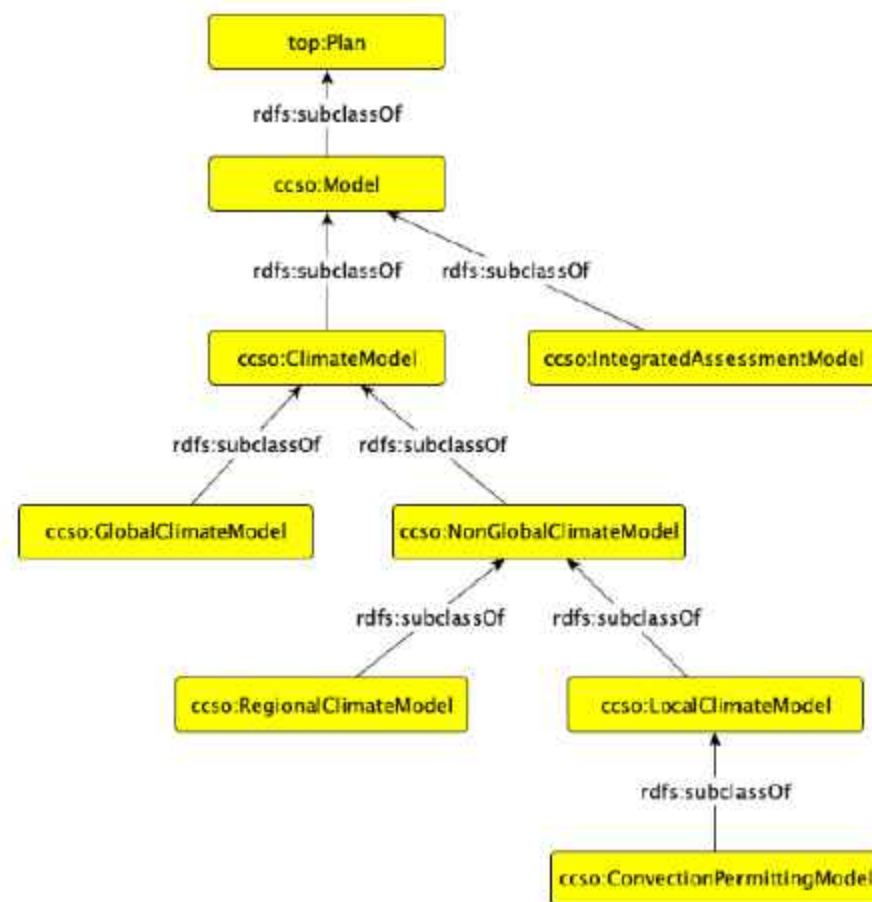


Figure 5. Taxonomy of Climate Models.

3. Aim

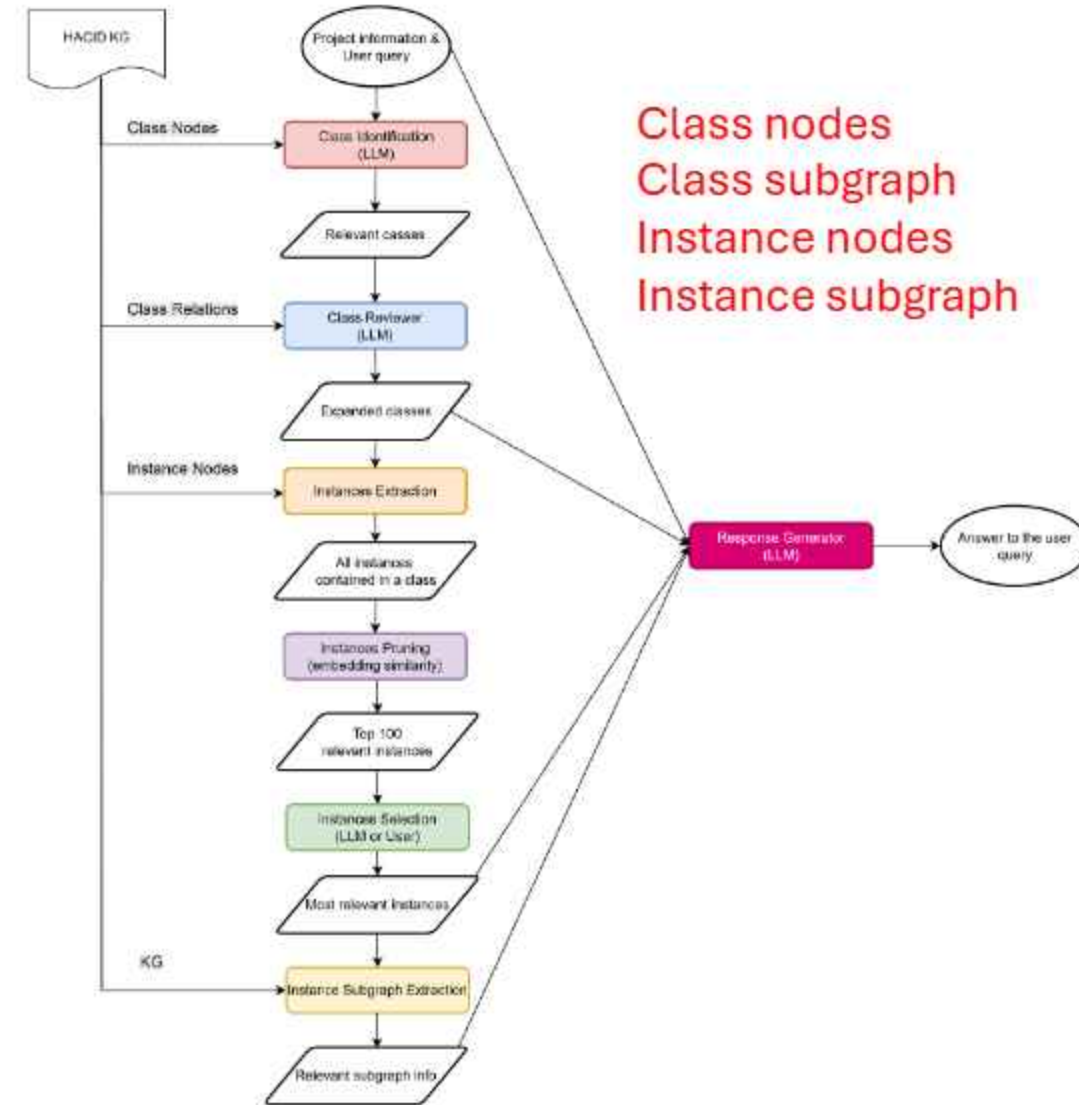
Upscaling Climate Services using AI

Key research question:

Is it possible to generate recipes/answers for specific climate services cases automatically using LLM and HACID KG ? And how ?

4. Methodology - GraphRAG

- HACID KG preprocessing (class & instances)
- LLM for class identification
- LLM for class expansion
- LLM for instance selection
- LLM for response generation
- User Interaction



5. Results - GraphRAG

Query 1:

Picking up everything related to temperature and heat

"Overheating assessment for University of Exeter as a result of climate change. What are the climate variables and metrics relevant to the case?"

Answer:

The climate variables relevant to the overheating assessment for the University of Exeter include northward atmosphere heat transport, tendency of specific humidity due to stratiform precipitation, northward ocean heat transport, tendency of air temperature due to stratiform precipitation, relative humidity, specific humidity, ocean heat y transport, sea surface temperature, air temperature, sea water temperature anomaly, convective rainfall rate, precipitation flux, tropopause air temperature, sea water conservative temperature, rainfall rate, atmosphere energy content, land ice temperature, and tendency of air temperature due to model physics. These variables play a significant role in assessing the impact of climate change on overheating phenomena.

Context from knowledge graph:

The project related to this case is an overheating assessment for the University of Exeter as a result of climate change. The involved classes are Variable and Parameter, which refer to meaningful abstractions and measurable quantities in context. The selected instances within the class include various climate variables such as heat transport, specific humidity, air and sea temperatures, and precipitation-related measures. These are associated with specific units of measurement, indicative of their roles in climate models and assessments.

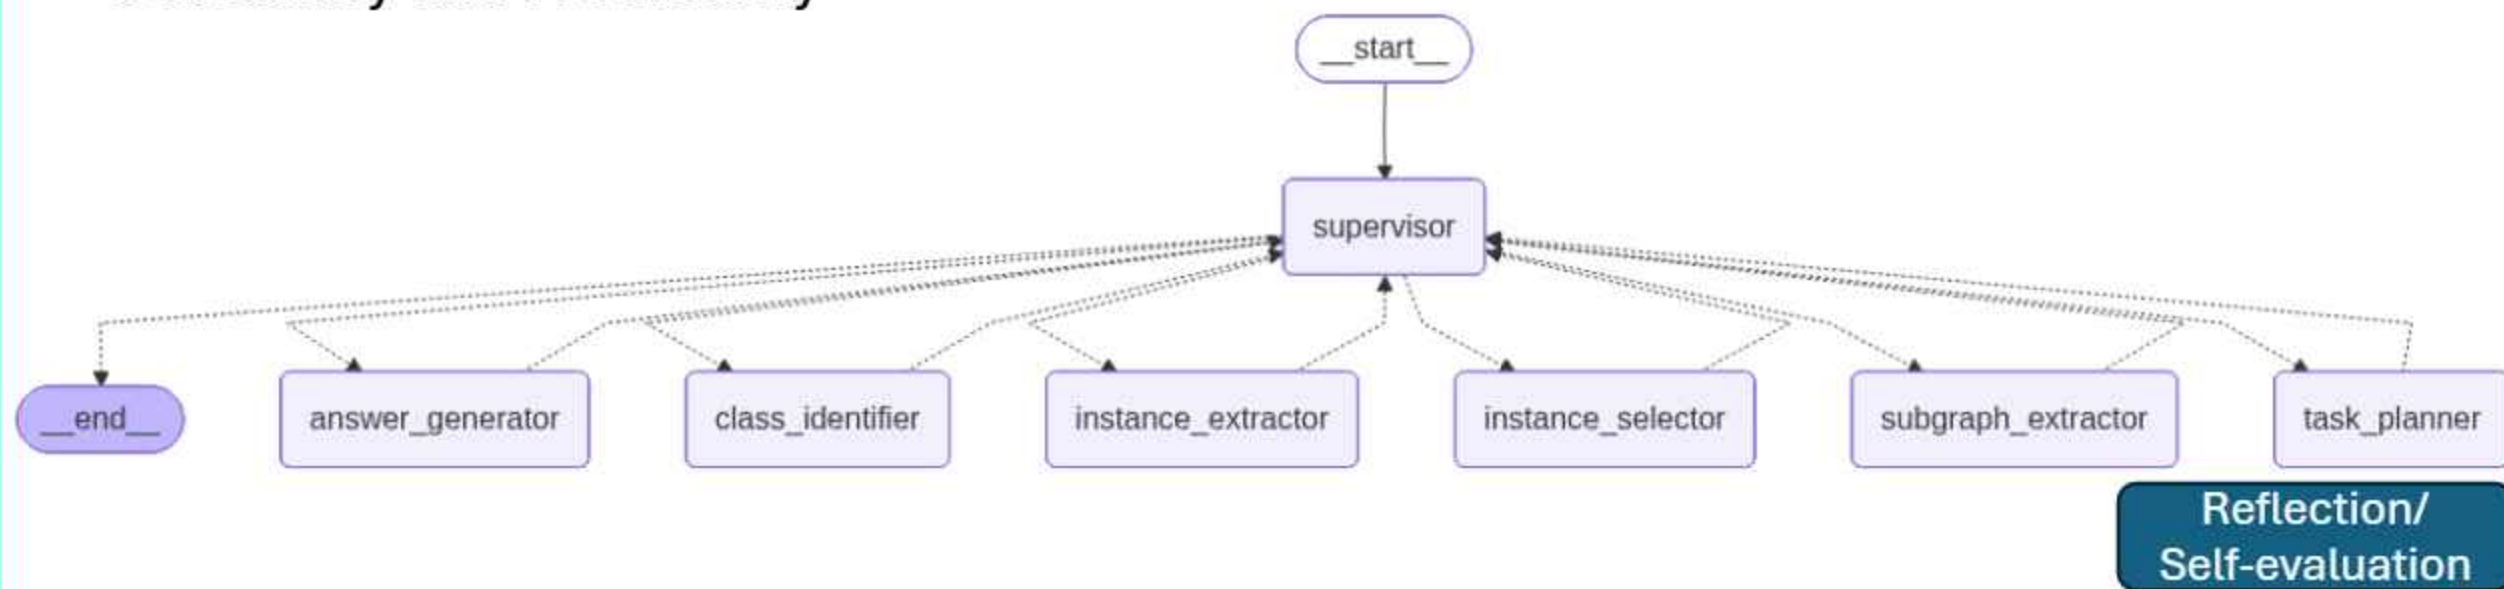
Further reasoning:

The context provided indicates a focus on various aspects of the climate system, including heat transport both in the atmosphere and ocean, different forms of temperature measurements, humidity metrics, and rainfall data. Understanding these variables' interactions and magnitudes can help in quantifying and predicting overheating risks associated with climate change impacts at the University of Exeter. The inclusion of both atmosphere and ocean dynamics highlights the complex interplay in climate change phenomena that must be considered for comprehensive assessments.

6. Methodology – Agentic GraphRAG

- Capture dependencies among questions in a climate service case
- Planning, Action, Reflection, Refinement
- Scalability and Modularity

The proposed solution uses supervisor based architecture to coordinate several agents/tools to solve the problem.



7. Result – Agentic GraphRAG

Air temperature is the bulk temperature of the air, not the surface (skin) temperature. The quantity with standard name heat_index_of_air_temperature is the perceived air temperature when relative humidity is taken into consideration (which makes it feel hotter than the actual air temperature). Heat index is only defined when the ambient air temperature is at or

Query 1:

temperature. The air temperature excess is the air temperature minus the air temperature threshold, where only positive values are included in the integral. Its integral with respect to time is often called after its units of "degree-days". The

"Overheating assessment for University of Exeter as a result of climate change. What are the climate variables and metrics relevant to the case?"

Specific and relevant answers

Answer:

The climate variables and metrics relevant to the overheating assessment for the University of Exeter as a result of climate change include air temperature, heat index of air temperature, relative humidity, specific humidity, air temperature threshold, air temperature anomaly, integral with respect to time of air temperature excess, and duration of sunshine.

Context from knowledge graph:

The selected instances within the class of variables related to the project include air temperature, heat index of air temperature, relative humidity, surface specific humidity, specific humidity, air temperature threshold, air temperature anomaly, integral with respect to time of air temperature excess, and duration of sunshine. These variables are associated with units of measure like Kelvin (K), dimensionless (1), and seconds (s).

Further reasoning:

The variables identified are critical for assessing thermal comfort and the risks of overheating due to climate change. Air temperature and heat index are directly related to perceived and actual temperatures respectively. Specific humidity and relative humidity provide insights into moisture levels, which impact human comfort and heat perception. The threshold, anomaly, and integral of air temperature excess help in understanding historical deviations and potential future exceedances of temperature expectations, which are crucial for planning and mitigation strategies. Duration of sunshine also plays a role in thermal conditions.

8. Result – GPT4o

Query 1:

"Overheating assessment for University of Exeter as a result of climate change. What are the climate variables and metrics relevant to the case?"

1. Air Temperature

- Daily maximum, minimum, and average air temperatures
- Record high temperatures and trends over time

2. Temperature Extremes

- Frequency, duration, and intensity of heatwaves
- Number of days exceeding specific temperature thresholds

3. Cooling Degree Days (CDD)

- Measures the demand for cooling based on temperature data

5. Relative Humidity

- Impacts perceived temperature and comfort levels

6. Solar Radiation and Insolation

- Amount of solar energy received, influencing building heat gain

7. Precipitation and Humidity Levels

- Affect outdoor and indoor climate conditions indirectly

8. Climate Change Projections

- Future temperature increases based on climate models and scenarios (e.g., RCP pathways)

These variables help in assessing how climate change may increase the likelihood and severity of overheating within university buildings and outdoor spaces, guiding adaptive measures.

Broad, but very general

8. Future works

- Quality and completeness of knowledge graph
- Context engineering – information pruning
- Knowledge graph reasoning
- Precise – we need one definitive answer, not 100 plausible answers. (domain knowledge)

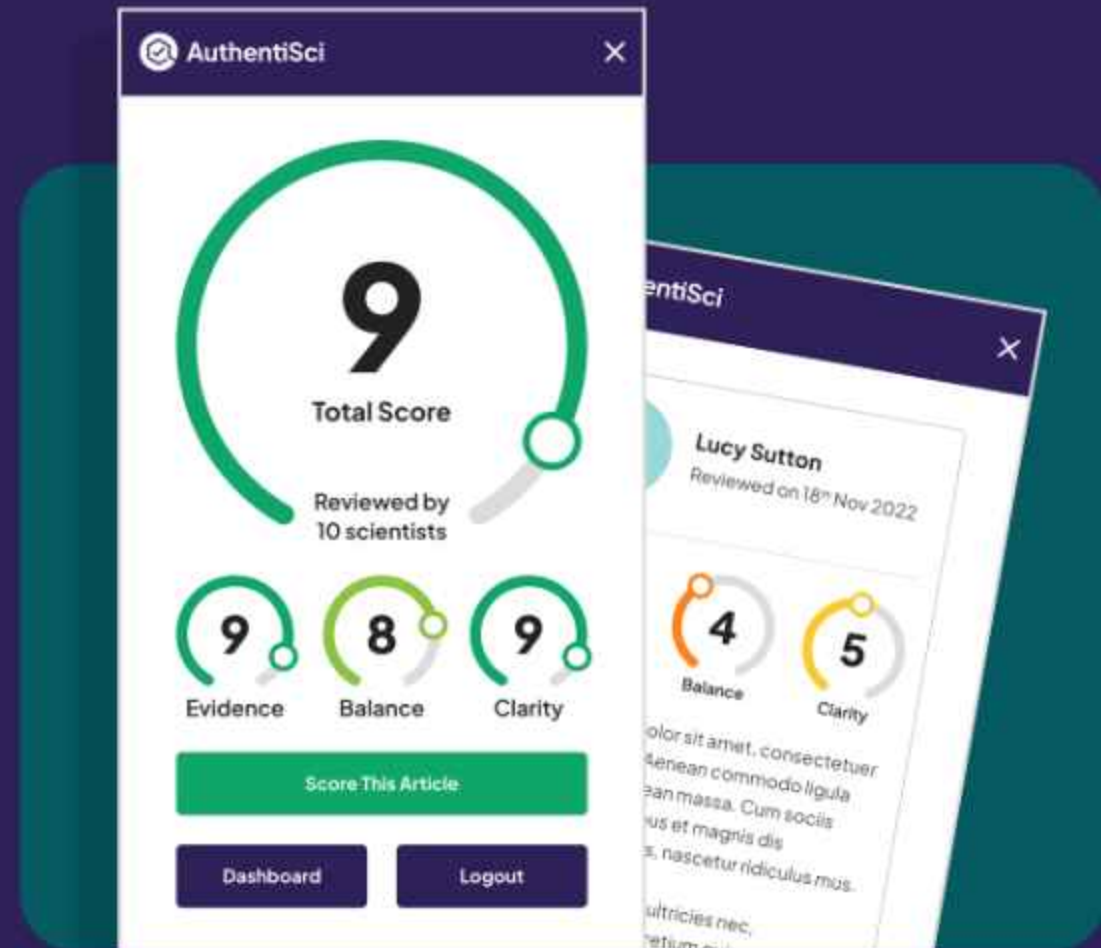
Bridging the Communities of Science and Society

AuthentiSci harnesses global scientific expertise to promote education and counter misinformation

MAX PLANCK INSTITUTE
FOR GRAVITATIONAL PHYSICS
(ALBERT EINSTEIN INSTITUTE)



LINDAU
NOBEL LAUREATE
MEETINGS



The Misinformation Problem

“Misinformation and disinformation” ranked as most severe global risk

World Economic Forum Global Risks Perception Survey 2023-2025

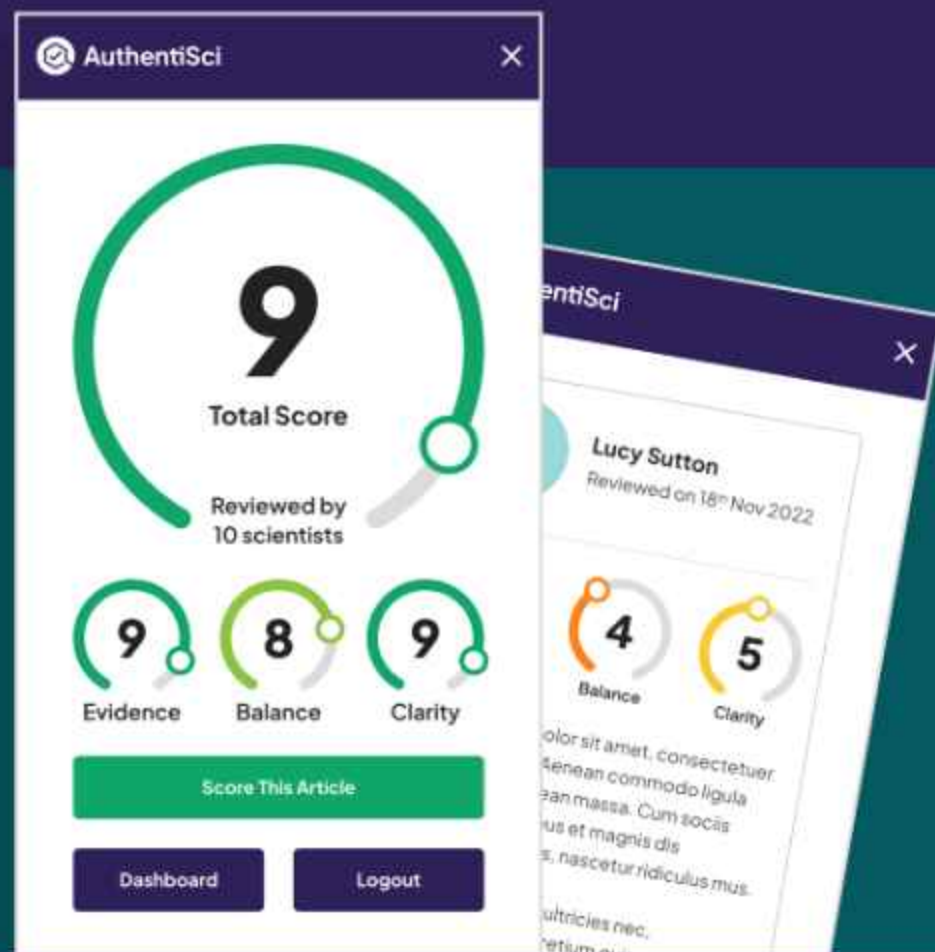
Scientists trusted most to “do what is right” and to “tell me the truth”

Edelman Trust Barometer 2024-2025

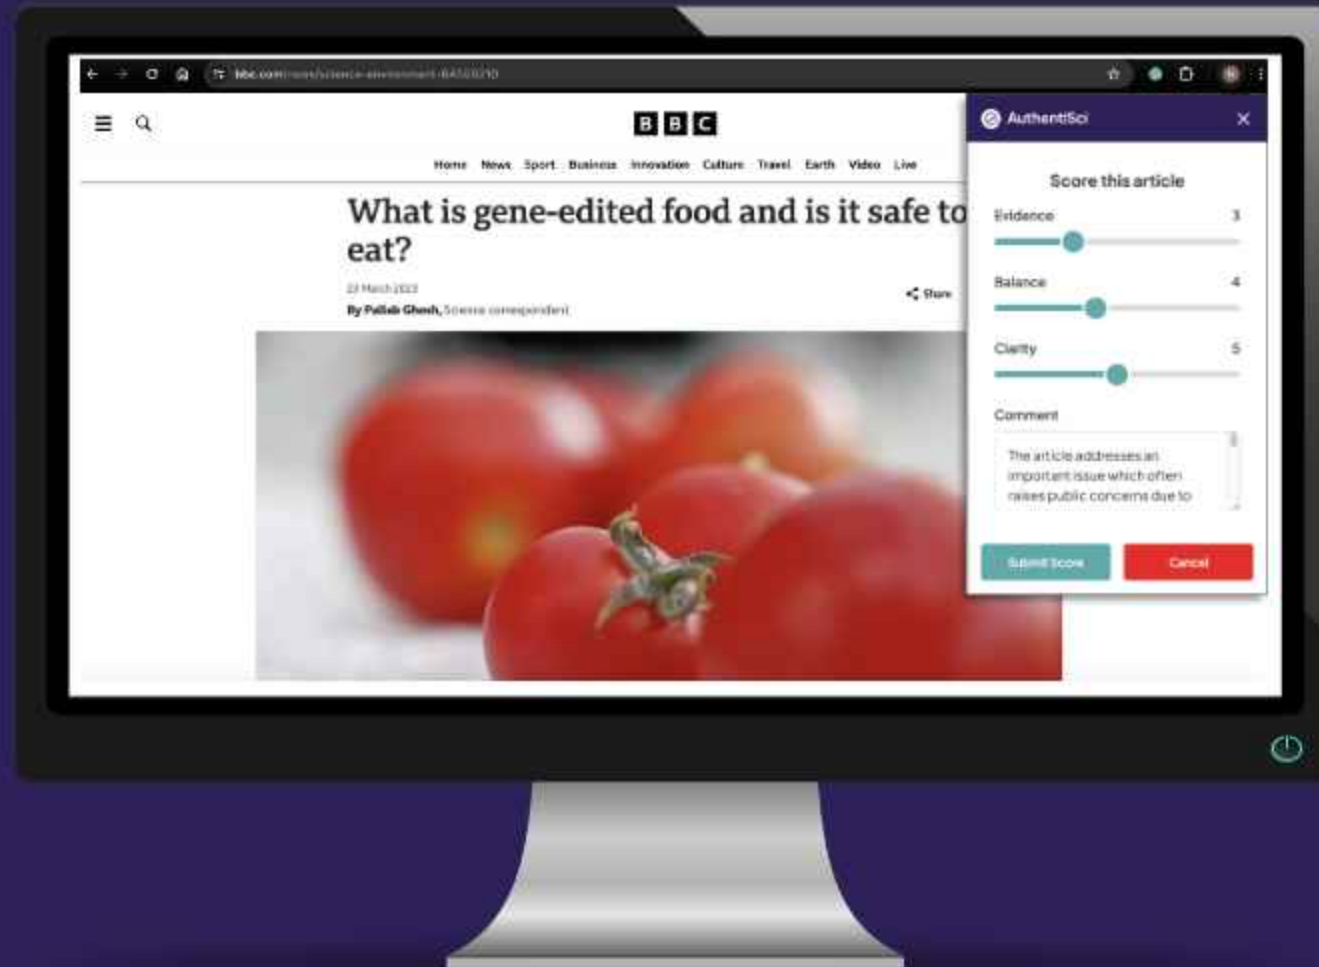
- Misinformation is devastating to climate action, public health, and societal progress;
- Voice of global scientists unheard;
- Fact-checking services are slow and narrow-focused.



Misleading media shapes opinion before experts can respond



Verified Experts, Real-Time Reviews



- Verified scientists (via ORCID) review media for Evidence, Balance, Clarity;
- Reviews are moderated, educational, and sharable;
- Free and accessible to all.


What is gene-edited food and is it safe to eat?

23 March 2023

By Pallab Ghosh, Science correspondent

🔗 Share



 **AuthentiSci** ✕

Score this article

Evidence	<div><div></div><div></div></div>	3
Balance	<div><div></div><div></div></div>	4
Clarity	<div><div></div><div></div></div>	5

Comment

The article addresses an important issue which often raises public concerns due to

Submit Score

Cancel





[Home](#) [News](#) [Sport](#) [Business](#) [Innovation](#) [Culture](#) [Travel](#) [Earth](#) [Video](#) [Live](#)

What is gene-edited food and is it safe to eat?

23 March 2023

By Pallab Ghosh, Science correspondent

↻ Share



AuthentiSci



Nicholas Clifton

Reviewed on
05/05/2024



Evidence



Balance



Clarity

The article addresses an important issue which often raises public concerns due to the unfamiliarity and historical associations of genetic alterations, particularly when they relate to something so critical to our health - i.e. food. The article does not sufficiently explain the advantages of gene editing. The brief explanations of

Score Breakdown

Reviews





[Home](#) [News](#) [Sport](#) [Business](#) [Innovation](#) [Culture](#) [Travel](#) [Earth](#) [Video](#) [Live](#)

What is gene-edited food and is it safe to eat?

23 March 2023

By Pallab Ghosh, Science correspondent

🔗 Share



AuthentiSci ✕



Request Score again

Reviewer login

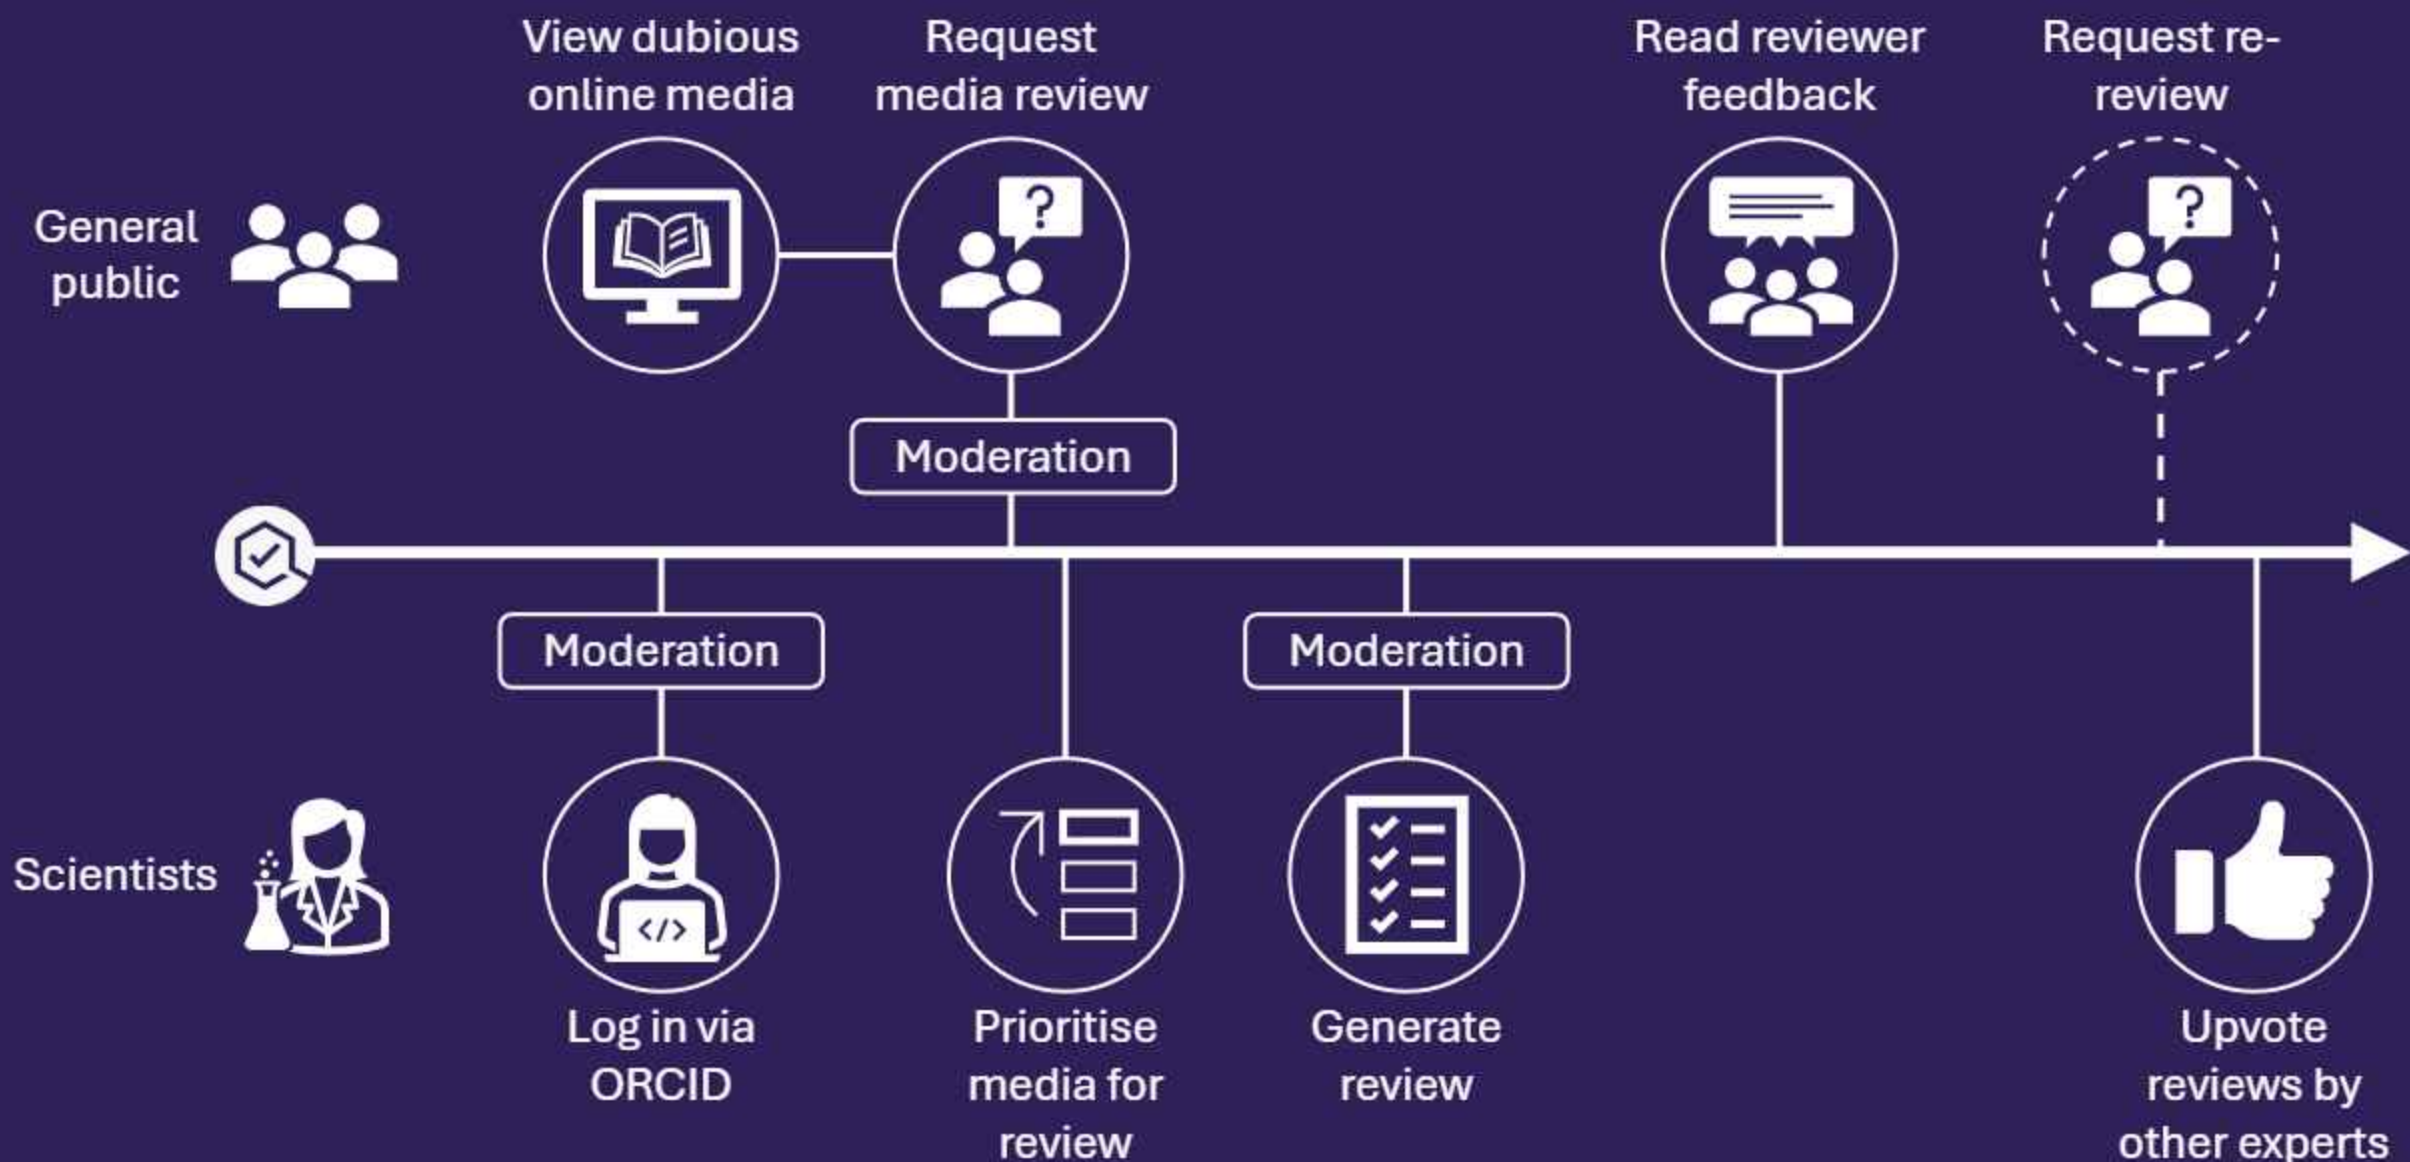
Find out more

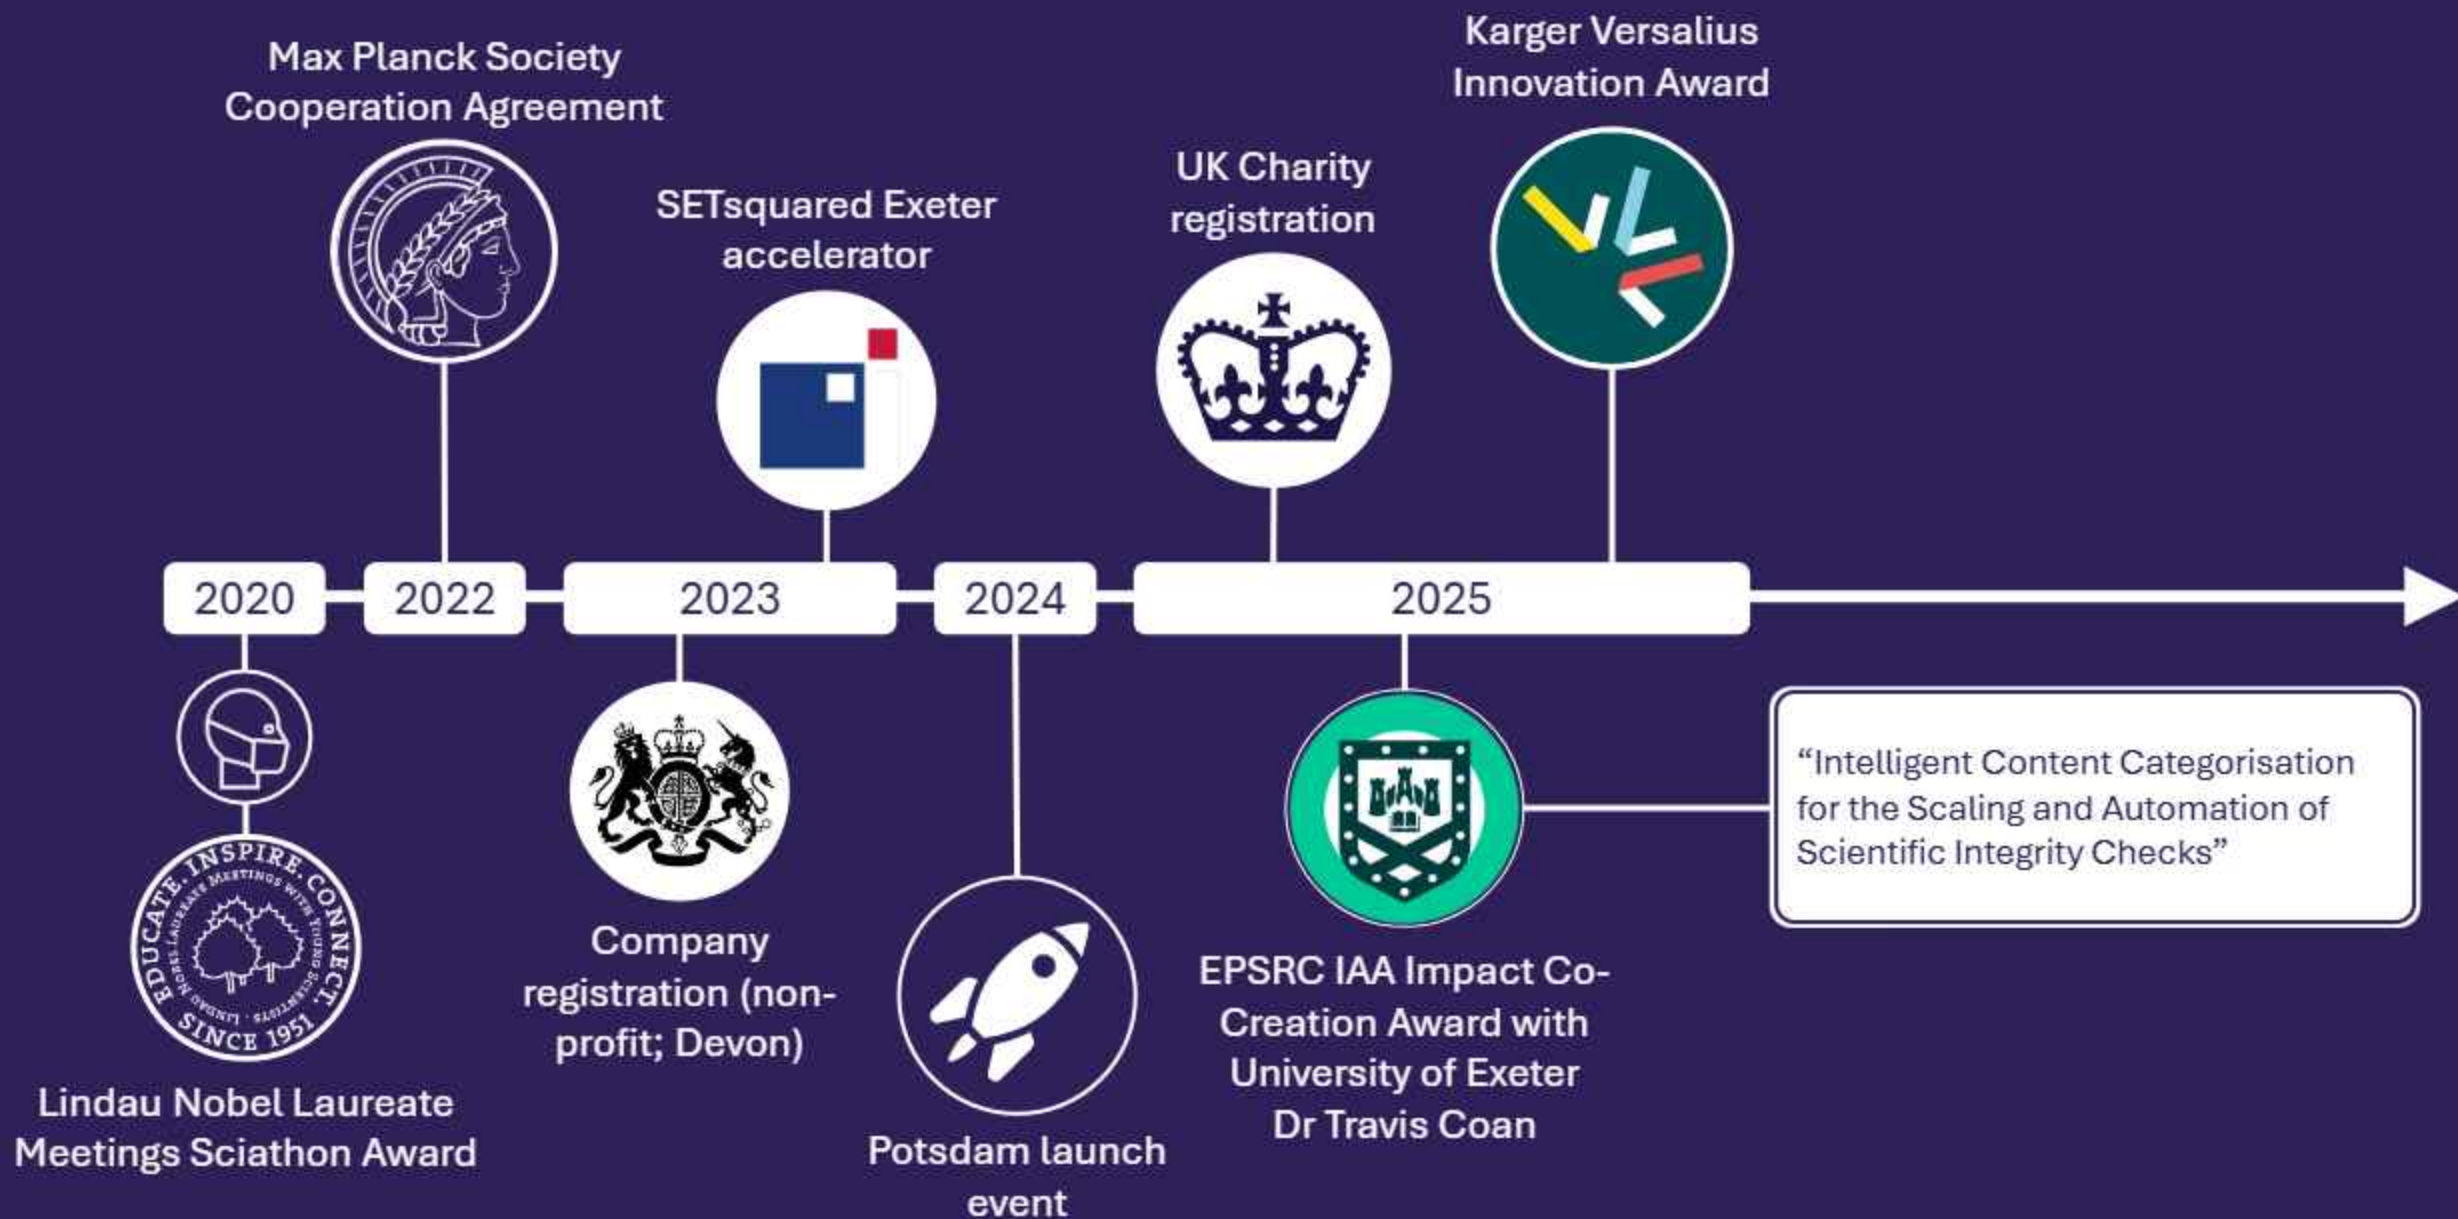
Score Breakdown

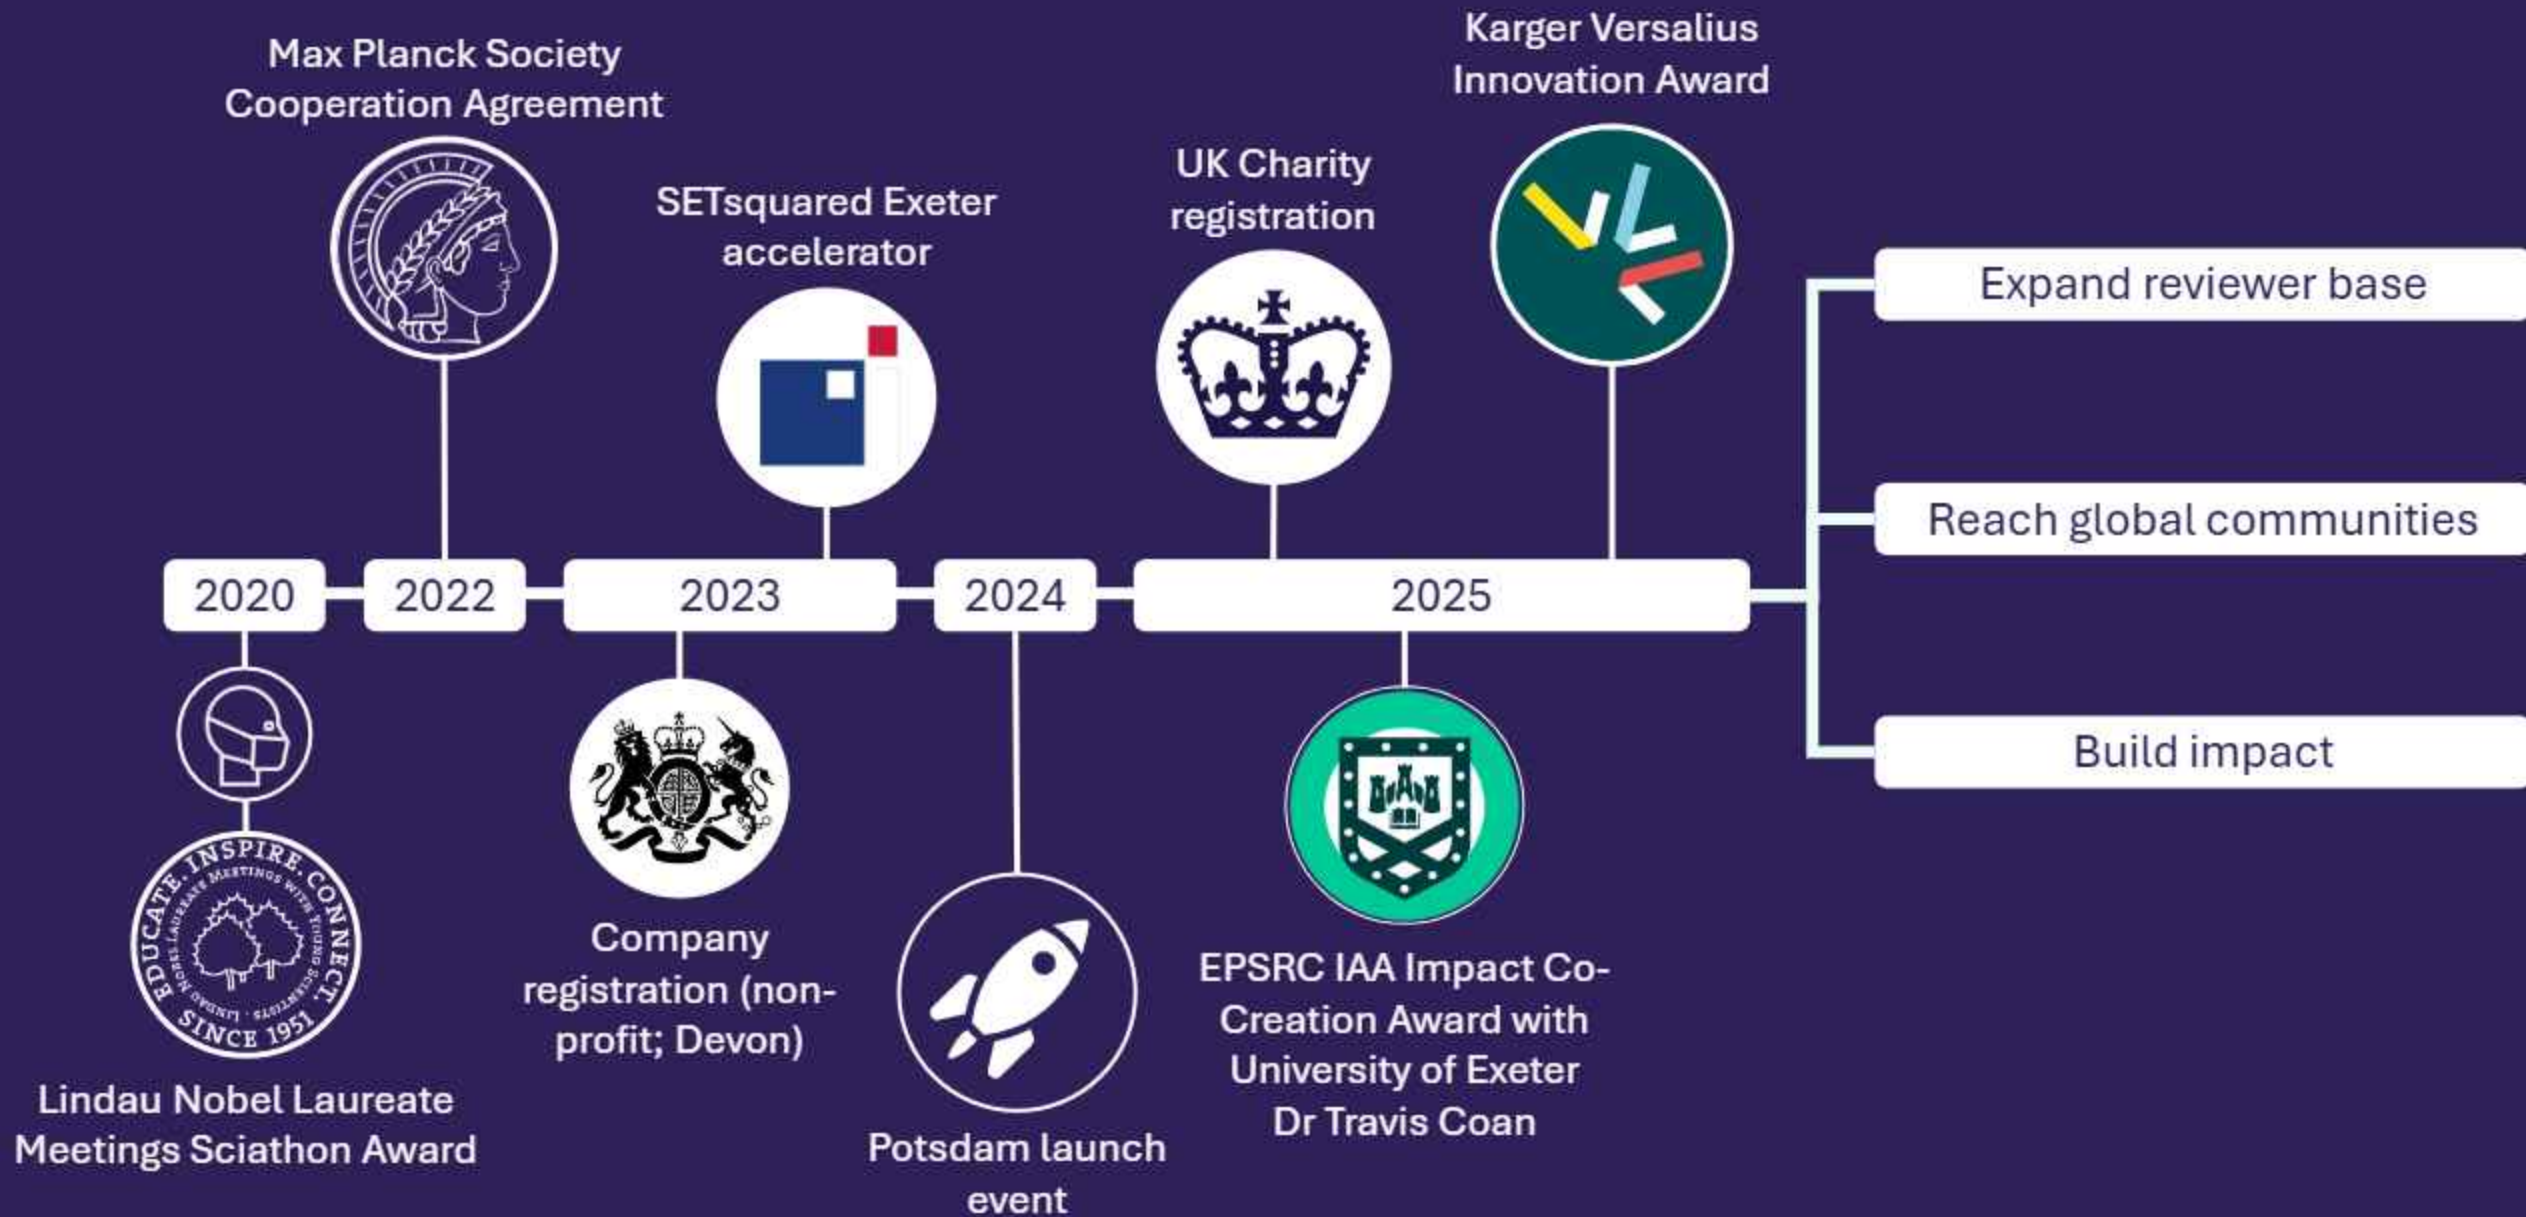
Reviews



Functionality overview







Why Contribute?

- Responsibility
- Public reviewer profile
- Local charity
- University impact case
- Ease of use

71%

of you say that lack of time prevents
you from contributing more to public
engagement

88%

of you encounter miscommunication
of your field in the media at least once
a month

76%

of you say that experiencing scientific
inaccuracies in the media is a strong
motivator to contribute more to public
understanding of science

Register for an Account

A screenshot of the ORCID ID sign-in page. It features the ORCID ID logo at the top, followed by the text "Sign in to ORCID". Below this are two input fields: "Email or 16-digit ORCID ID" and "Password". A "SIGN IN" button is positioned below the password field. Links for "Forgot your password or ORCID ID?" and "Don't have an ORCID ID yet? Register now" are provided. At the bottom, there is a link to "Access through your institution" with a small icon of a building.

- ORCID identification
- Academic Affiliation
- At least one peer-reviewed publication (public)

A screenshot of a research article snippet from a database. The title is "Cortical neuronal hyperexcitability and synaptic changes in SGCE mutation-positive myoclonus dystonia". It includes a date "2023-04-19", a source "Journal article", and a DOI "10.1093/brain/awac365". The contributors listed are Alessandra Sperandeo, Claudia Tamburini, Zoe Noakes, Daniel Cabezas de la Fuente, Francesca Keefe, Olena Petter, William Plumbly, Nicholas E Clifton, Meng Li, and Kathryn J Peall. The source is identified as "Crossref". A "Show more detail" link is present. A teal arrow points from the "Register now" link in the ORCID form to this snippet.



Scan to log in with ORCID
(<https://authentisci.com/auth/login>)



University
of Exeter

LLM-assisted classification of claims in news coverage of net zero

Tristan Cann, Sylvia Hayes,
Ranadheer Malla, Travis Coan

Using semantic similarity to measure the echo of strategic communications

Research | Open access | Published: 12 March 2023
Volume 14, article number 20, (2023) | [Cite this article](#)



Visual portrayals of fun in the sun in European news outlets misrepresent heatwave risks

Saffron O'Neill , Sylvia Hayes, Nadine Strauß, Marie-Noëlle Doutreix, Katharine Steentjes, Joshua Ettinger, Ned Westwood, James Painter

First published: 18 October 2022 | <https://doi.org/10.1111/geoj.12487> | Citations: 4

Computer Science



Social Science

{'score': 0.8562073111534119, 'label': 'other hob'}



Computer-assisted classification of contrarian claims about climate change

Travis G. Coan, Constantine Boussalis, John Cook  & Mirjam O. Nanko

Scientific Reports 11, Article number: 22320 (2021) | [Cite this article](#)



What is a claim?

We communicate our understanding of the world around us through language.

Claims are statements or assertions for a particular position:

Extreme weather events are becoming more frequent and intense due to climate change.

Winning the global race to net zero is essential to securing economic growth, energy security, skilled jobs and inward investment.

WWF warns that delaying investment in a greener economy by 10 years would double the amount of money needed.

Local news claims about net zero



We consider a case study of net zero coverage in UK local news.

We gathered perspectives from specialists in this space, incorporating their experience into an initial codebook of key themes in local news.

The Public Interest News Foundation give a list of local news sources – we collected 80,000 articles about net zero, before sampling 1000 to see how the stakeholder claims align with those identified by an LLM.

Local News Map

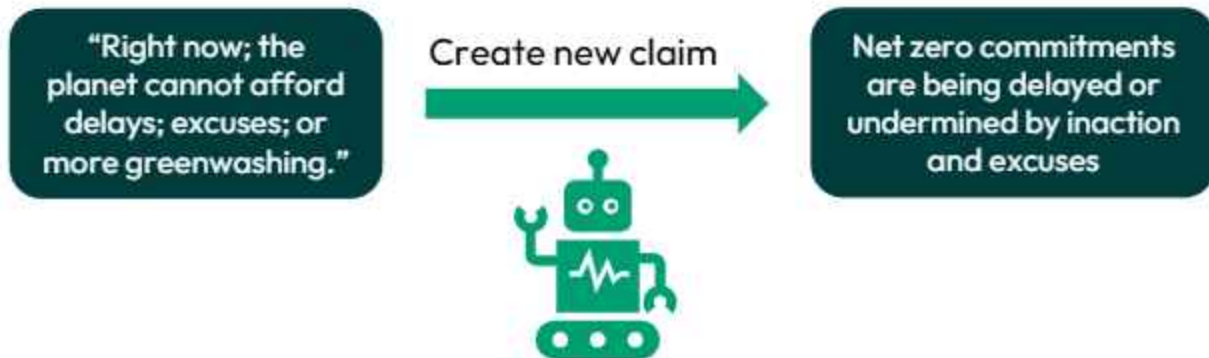
A project to map out all of the UK's professional providers of local news.



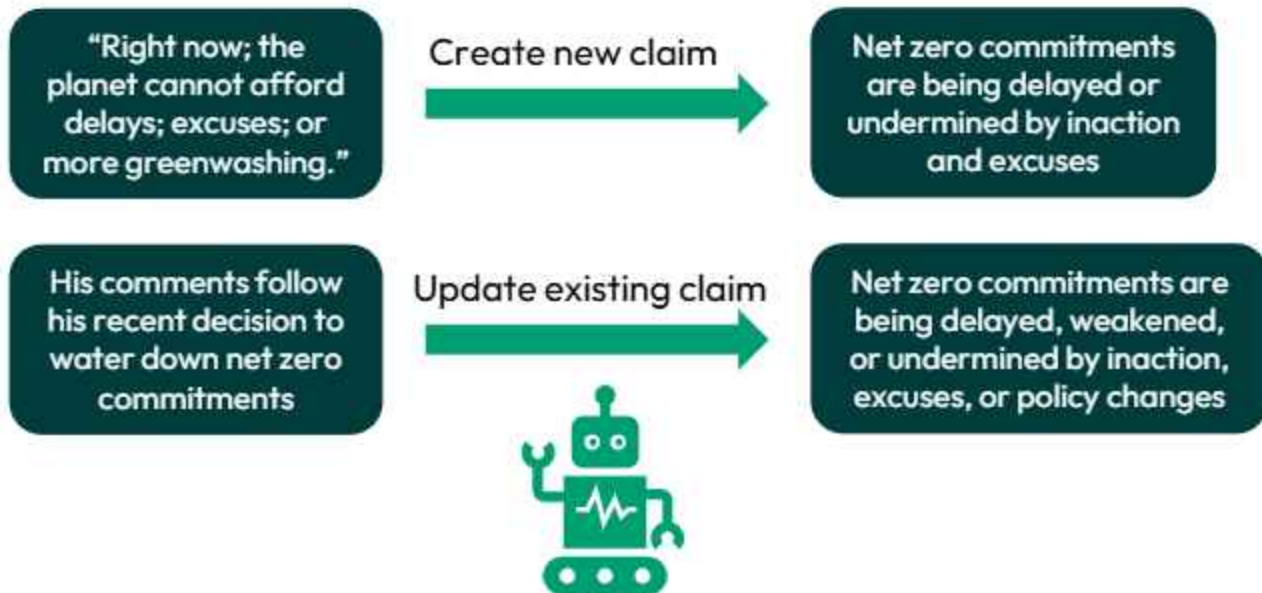
The GenAI codebook

"Right now; the planet cannot afford delays; excuses; or more greenwashing."

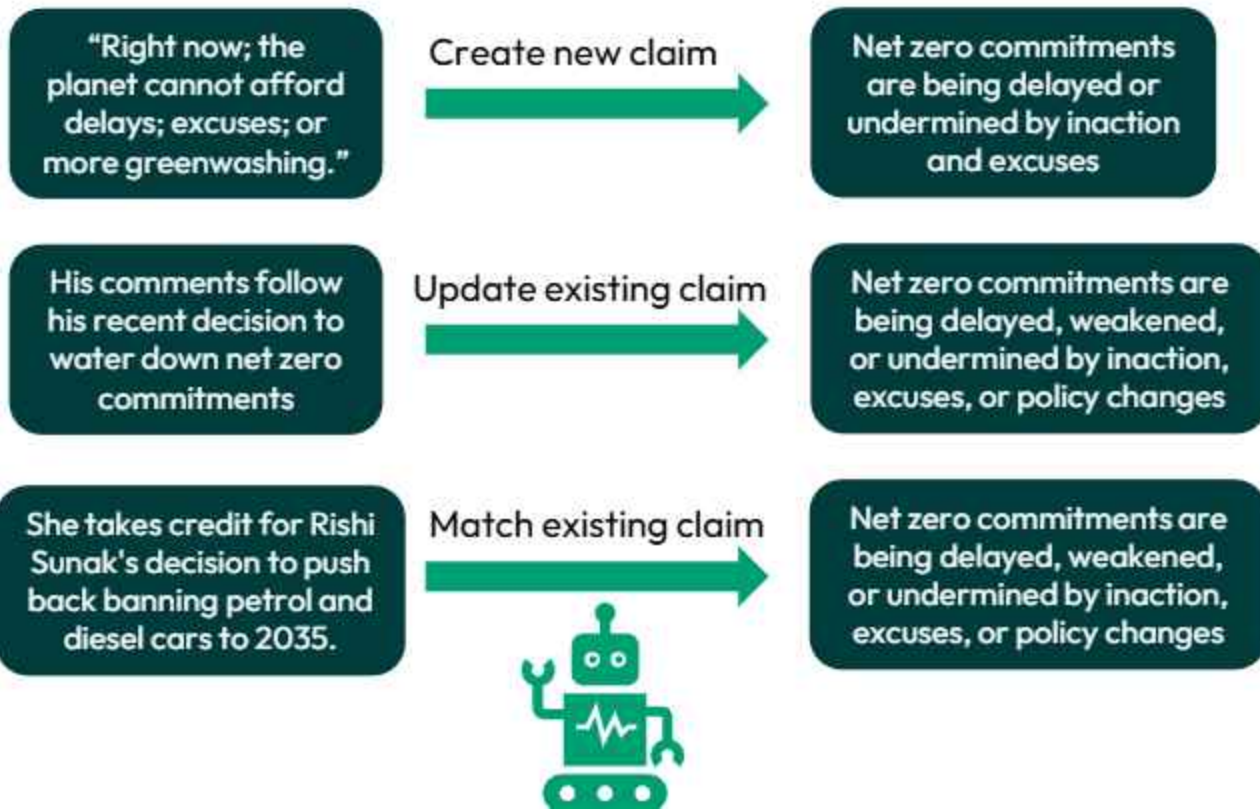
The GenAI codebook



The GenAI codebook



The GenAI codebook



Local news claims about net zero

Automatically recovered claims vary in quality:

- “External crises strengthen the business case for net zero investment”
 - “Protests that disrupt fossil fuel operations without providing solutions undermine efforts to achieve net zero”
 - “Net zero can be achieved through technological innovation”
-
- “Government faces challenges in revising net zero strategies”
 - “Local government sets net zero target or strategy with political support”
 - “Specific projects are implemented to advance net zero goals”

The human in the loop

We use expert judgment to refine the codebook and combine specific LLM-defined claims into more general “super claims”.

Examples:

- Net zero implementation is being hindered by contradictions and failures in politics such as greenwashing
- The costs and benefits of net zero are not equally distributed
- The UK still need to continue using and developing fossil fuels for energy security

The final codebook of 18 claims contains key positions for and against the drive towards net zero.

Summary



- LLMs can accelerate the process of codebook development and application.
- Keeping humans in the loop is **vital** for extracting usable insights.
- Next steps: validating a fine-tuned classifier for this codebook and expanding our analysis to the full 80k articles.

t.j.b.cann@exeter.ac.uk

c3ds@exeter.ac.uk

Codebook prompt – illustrative example



TASK DETAILS:

1) FIND ALL RELEVANT CONTENT:

- Scan the entire article for any mention or claim related to net zero.
- Quote the statement (verbatim) to preserve the original wording.
- A "claim" is an assertion that takes a position on net zero. Claims typically express opinions, predictions, interpretations, or causal relationships rather than purely objective facts.

2) CATEGORIZE EACH CLAIM:

- For each claim, decide whether it matches or partially matches one or more existing codebook categories.
- If it clearly falls under an existing category, choose `action_type: "match_existing_category"`.
- If it mostly matches but introduces a meaningful difference that warrants broadening the category, choose `action_type: "update_existing_category"` and explain how you updated it.
- WHEN NOT TO UPDATE (Examples):
 - * KEEP SEPARATE: If you have a category "Net zero will bankrupt businesses" and encounter a claim about "Net zero will bankrupt me", do NOT update to "Net zero will lead to bankruptcy" as these represent fundamentally different targets of impacts requiring separate categories.

3) NEUTRALITY AND OBJECTIVITY:

- Do not judge or verify the truthfulness of any claim. Simply present it as is and categorize it.
- Because many articles come from local news sources, you may see a lot of opposition to developments at the local level.

Codebook prompt – illustrative example



TASK DETAILS:

1) FIND ALL RELEVANT CONTENT:

- Scan the entire article for any mention or claim related to net zero.
- Quote the statement (verbatim) to preserve the original wording.
- A "claim" is an assertion that takes a position on net zero. Claims typically express opinions, predictions, interpretations, or causal relationships rather than purely objective facts.

2) CATEGORIZE EACH CLAIM:

- For each claim, decide whether it matches or partially matches one or more existing codebook categories.
- If it clearly falls under an existing category, choose `action_type: "match_existing_category"`.
- If it mostly matches but introduces a meaningful difference that warrants broadening the category, choose `action_type: "update_existing_category"` and explain how you updated it.
- WHEN NOT TO UPDATE (Examples):
 - * KEEP SEPARATE: If you have a category "Net zero will bankrupt businesses" and encounter a claim about "Net zero will bankrupt me", do NOT update to "Net zero will lead to bankruptcy" as these represent fundamentally different targets of impacts requiring separate categories.

3) NEUTRALITY AND OBJECTIVITY:

- Do not judge or verify the truthfulness of any claim. Simply present it as is and categorize it.
- Because many articles come from local news sources, you may see a lot of opposition to developments at the local level.

Codebook prompt – illustrative example



TASK DETAILS:

1) FIND ALL RELEVANT CONTENT:

- Scan the entire article for any mention or claim related to net zero.
- Quote the statement (verbatim) to preserve the original wording.
- A "claim" is an assertion that takes a position on net zero. Claims typically express opinions, predictions, interpretations, or causal relationships rather than purely objective facts.

2) CATEGORIZE EACH CLAIM:

- For each claim, **decide whether it matches or partially matches one or more existing codebook categories.**
- If it clearly falls under an existing category, choose `"action_type: "match_existing_category"`.
- If it mostly matches but introduces a meaningful difference that warrants broadening the category, choose `"action_type: "update_existing_category"` and explain how you updated it.
- WHEN NOT TO UPDATE (Examples):
 - * KEEP SEPARATE: If you have a category "Net zero will bankrupt businesses" and encounter a claim about "Net zero will bankrupt me", do NOT update to "Net zero will lead to bankruptcy" as these represent fundamentally different targets of impacts requiring separate categories.

3) NEUTRALITY AND OBJECTIVITY:

- Do not judge or verify the truthfulness of any claim. Simply present it as is and categorize it.
- Because many articles come from local news sources, you may see a lot of opposition to developments at the local level.

Codebook prompt – illustrative example



TASK DETAILS:

1) FIND ALL RELEVANT CONTENT:

- Scan the entire article for any mention or claim related to net zero.
- Quote the statement (verbatim) to preserve the original wording.
- A "claim" is an assertion that takes a position on net zero. Claims typically express opinions, predictions, interpretations, or causal relationships rather than purely objective facts.

2) CATEGORIZE EACH CLAIM:

- For each claim, decide whether it matches or partially matches one or more existing codebook categories.
- If it clearly falls under an existing category, choose `action_type: "match_existing_category"`.
- If it mostly matches but introduces a meaningful difference that warrants broadening the category, choose `action_type: "update_existing_category"` and explain how you updated it.
- WHEN NOT TO UPDATE (Examples):
 - * KEEP SEPARATE: If you have a category "Net zero will bankrupt businesses" and encounter a claim about "Net zero will bankrupt me", do NOT update to "Net zero will lead to bankruptcy" as these represent fundamentally different targets of impacts requiring separate categories.

3) NEUTRALITY AND OBJECTIVITY:

- Do not judge or verify the truthfulness of any claim. Simply present it as is and categorize it.
- Because many articles come from local news sources, you may see a lot of opposition to developments at the local level.

Data-Augmented LLMs for Soil Science

Kuangdai Leng

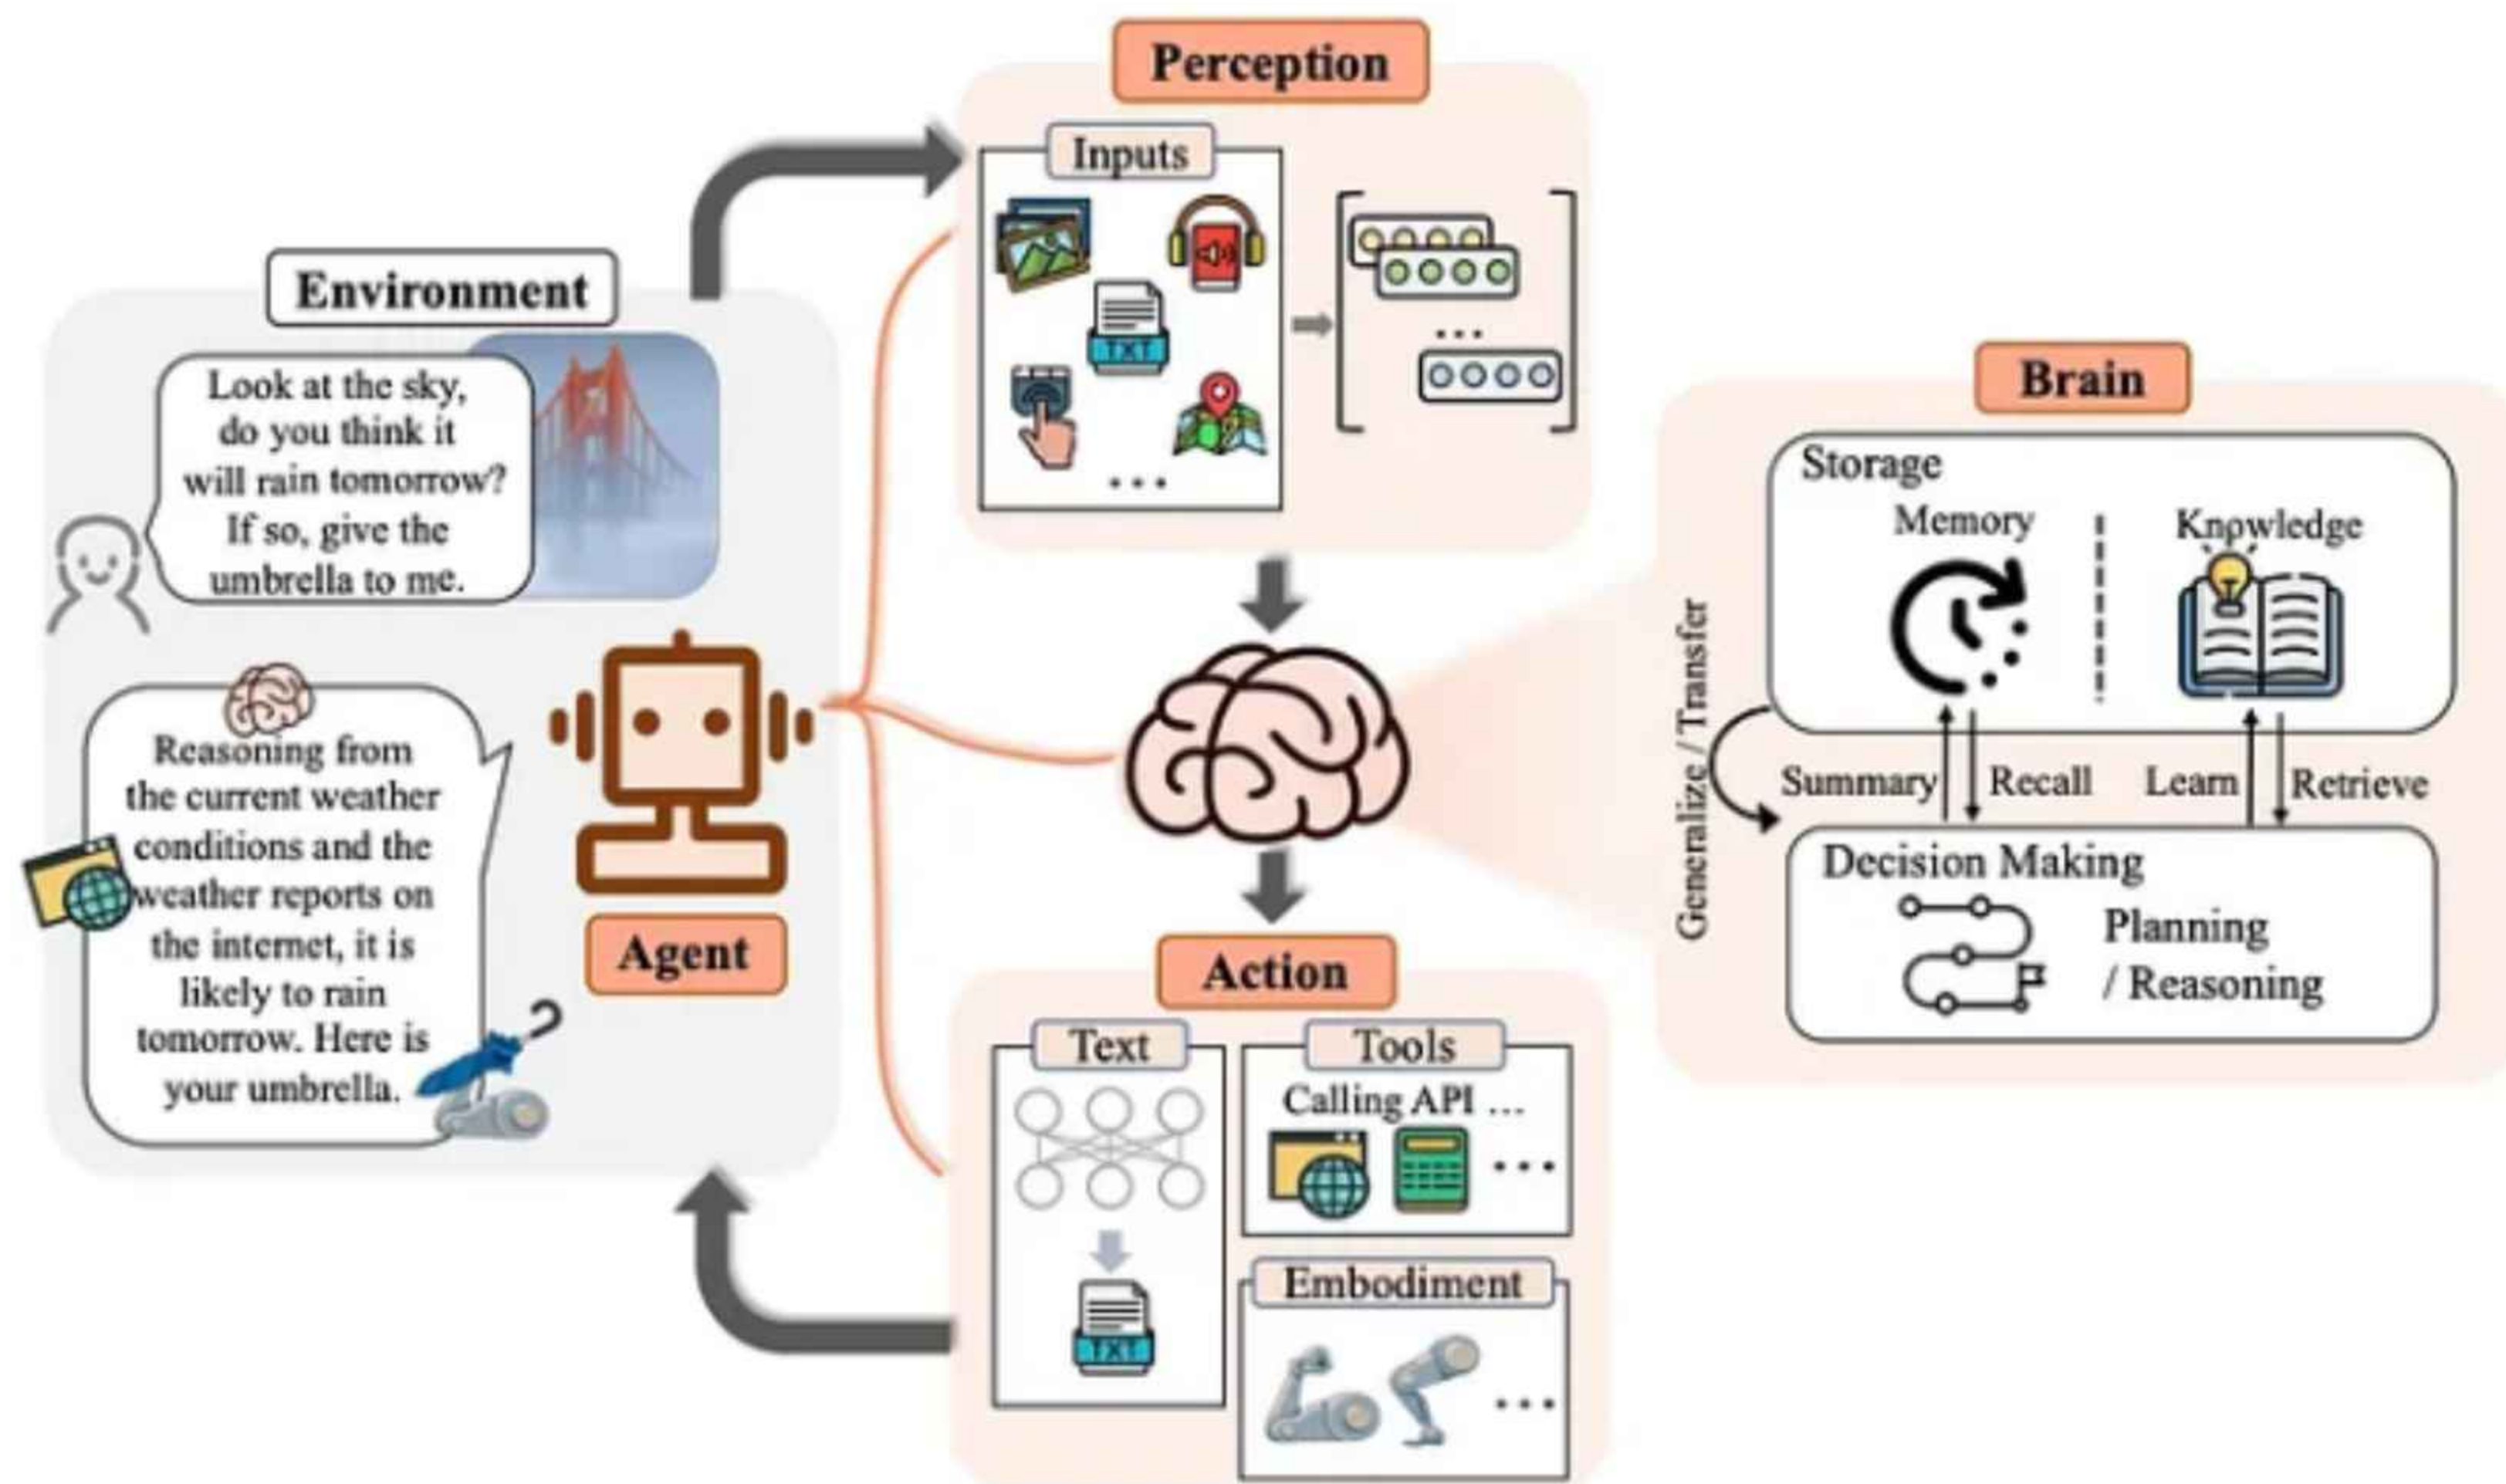
Earth Rover Program

El@Exeter, 11 Sep 25

EARTH
ROVER
PROGRAM

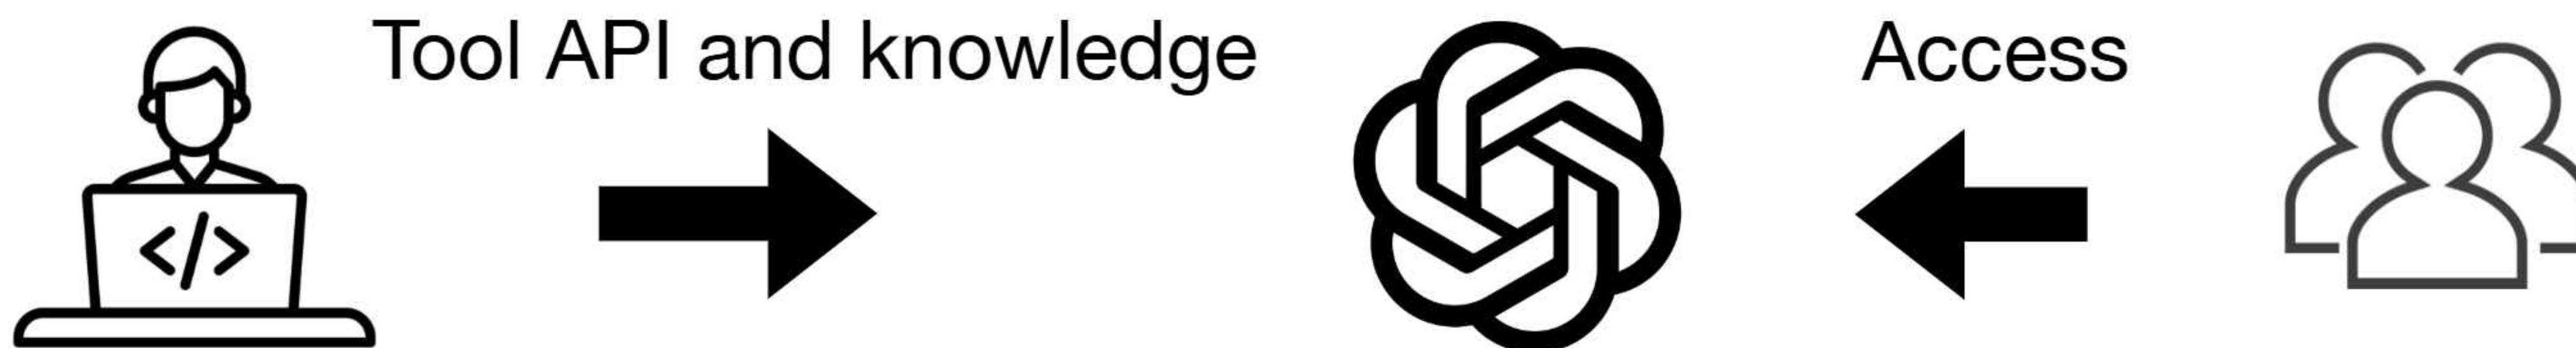
(Multi-)Agent Orchestration

- Empower next-gen applications with
 - LLMs of emergent capabilities
 - Databases
 - Knowledge bases
 - User memory
 - Tools — Maps, news, symbolic math, ...

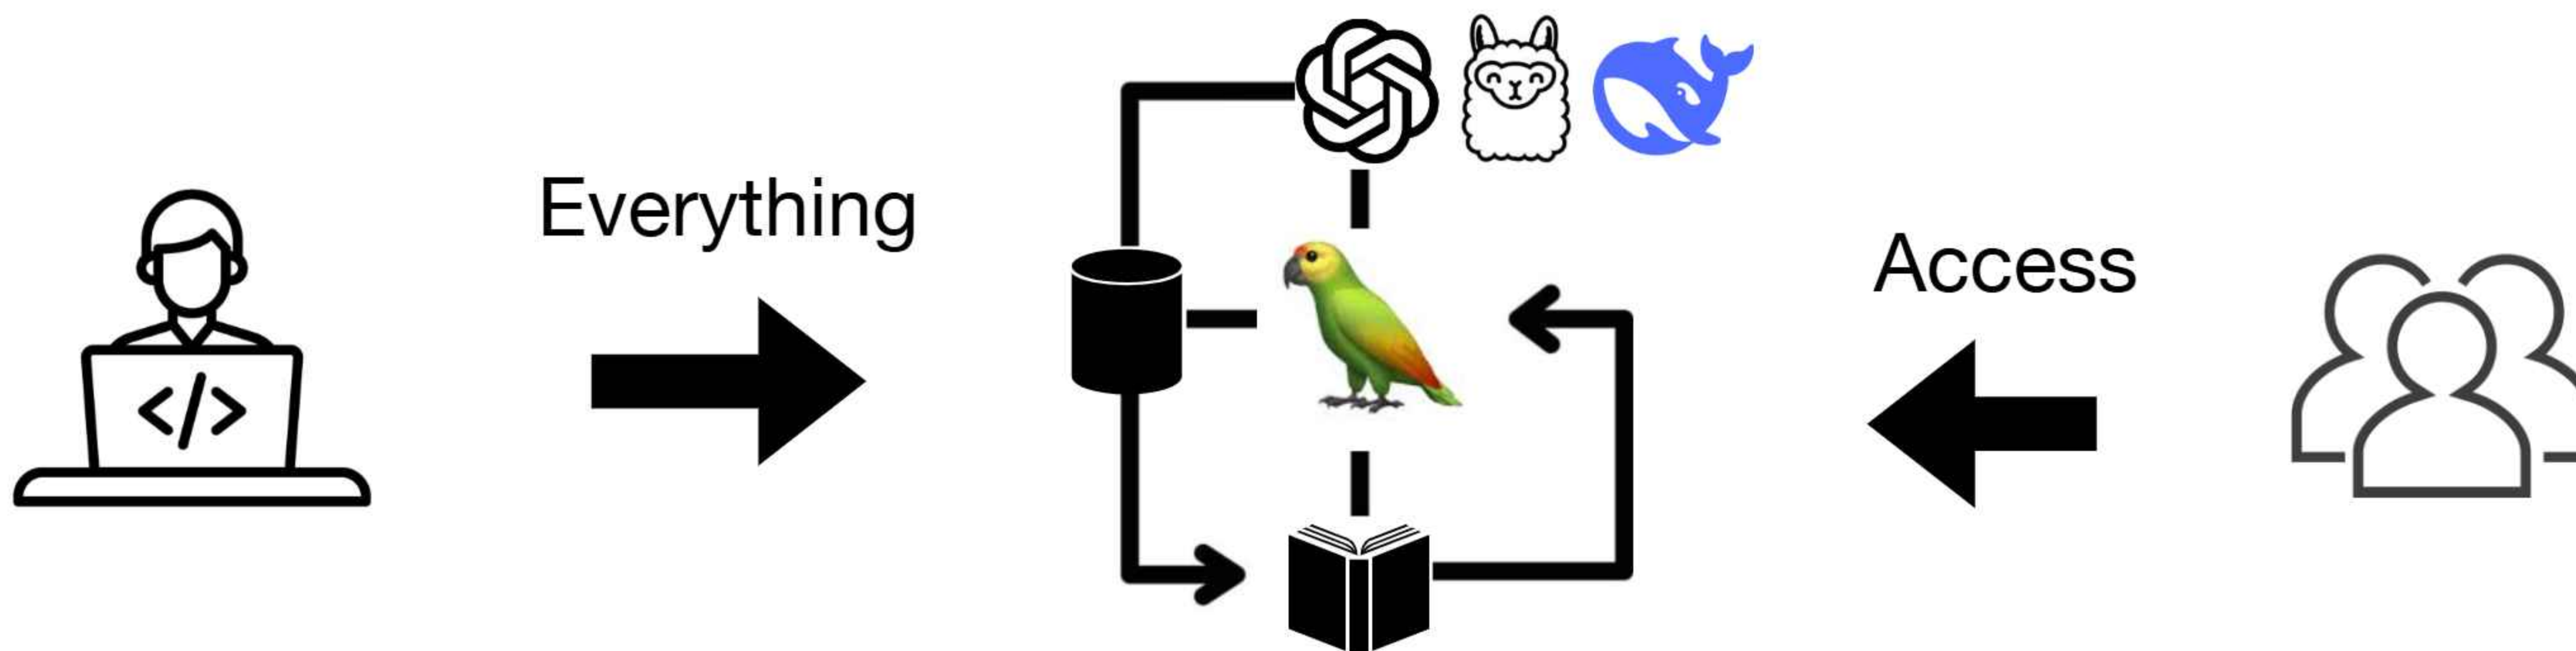


Two Practical Modes

- Tool-augmented GPT (OpenAI, Claude, DeepSeek)



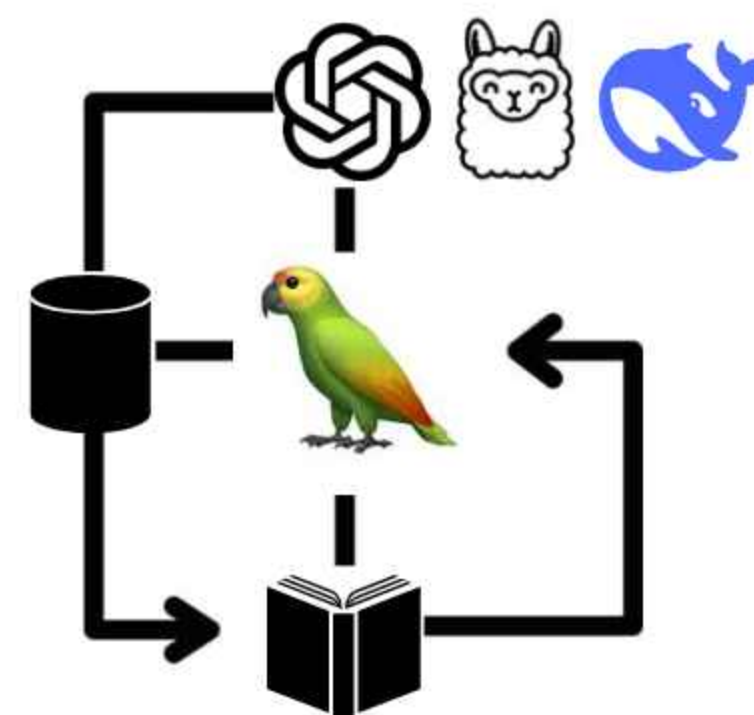
- LLMs as components of a framework



Two Practical Modes

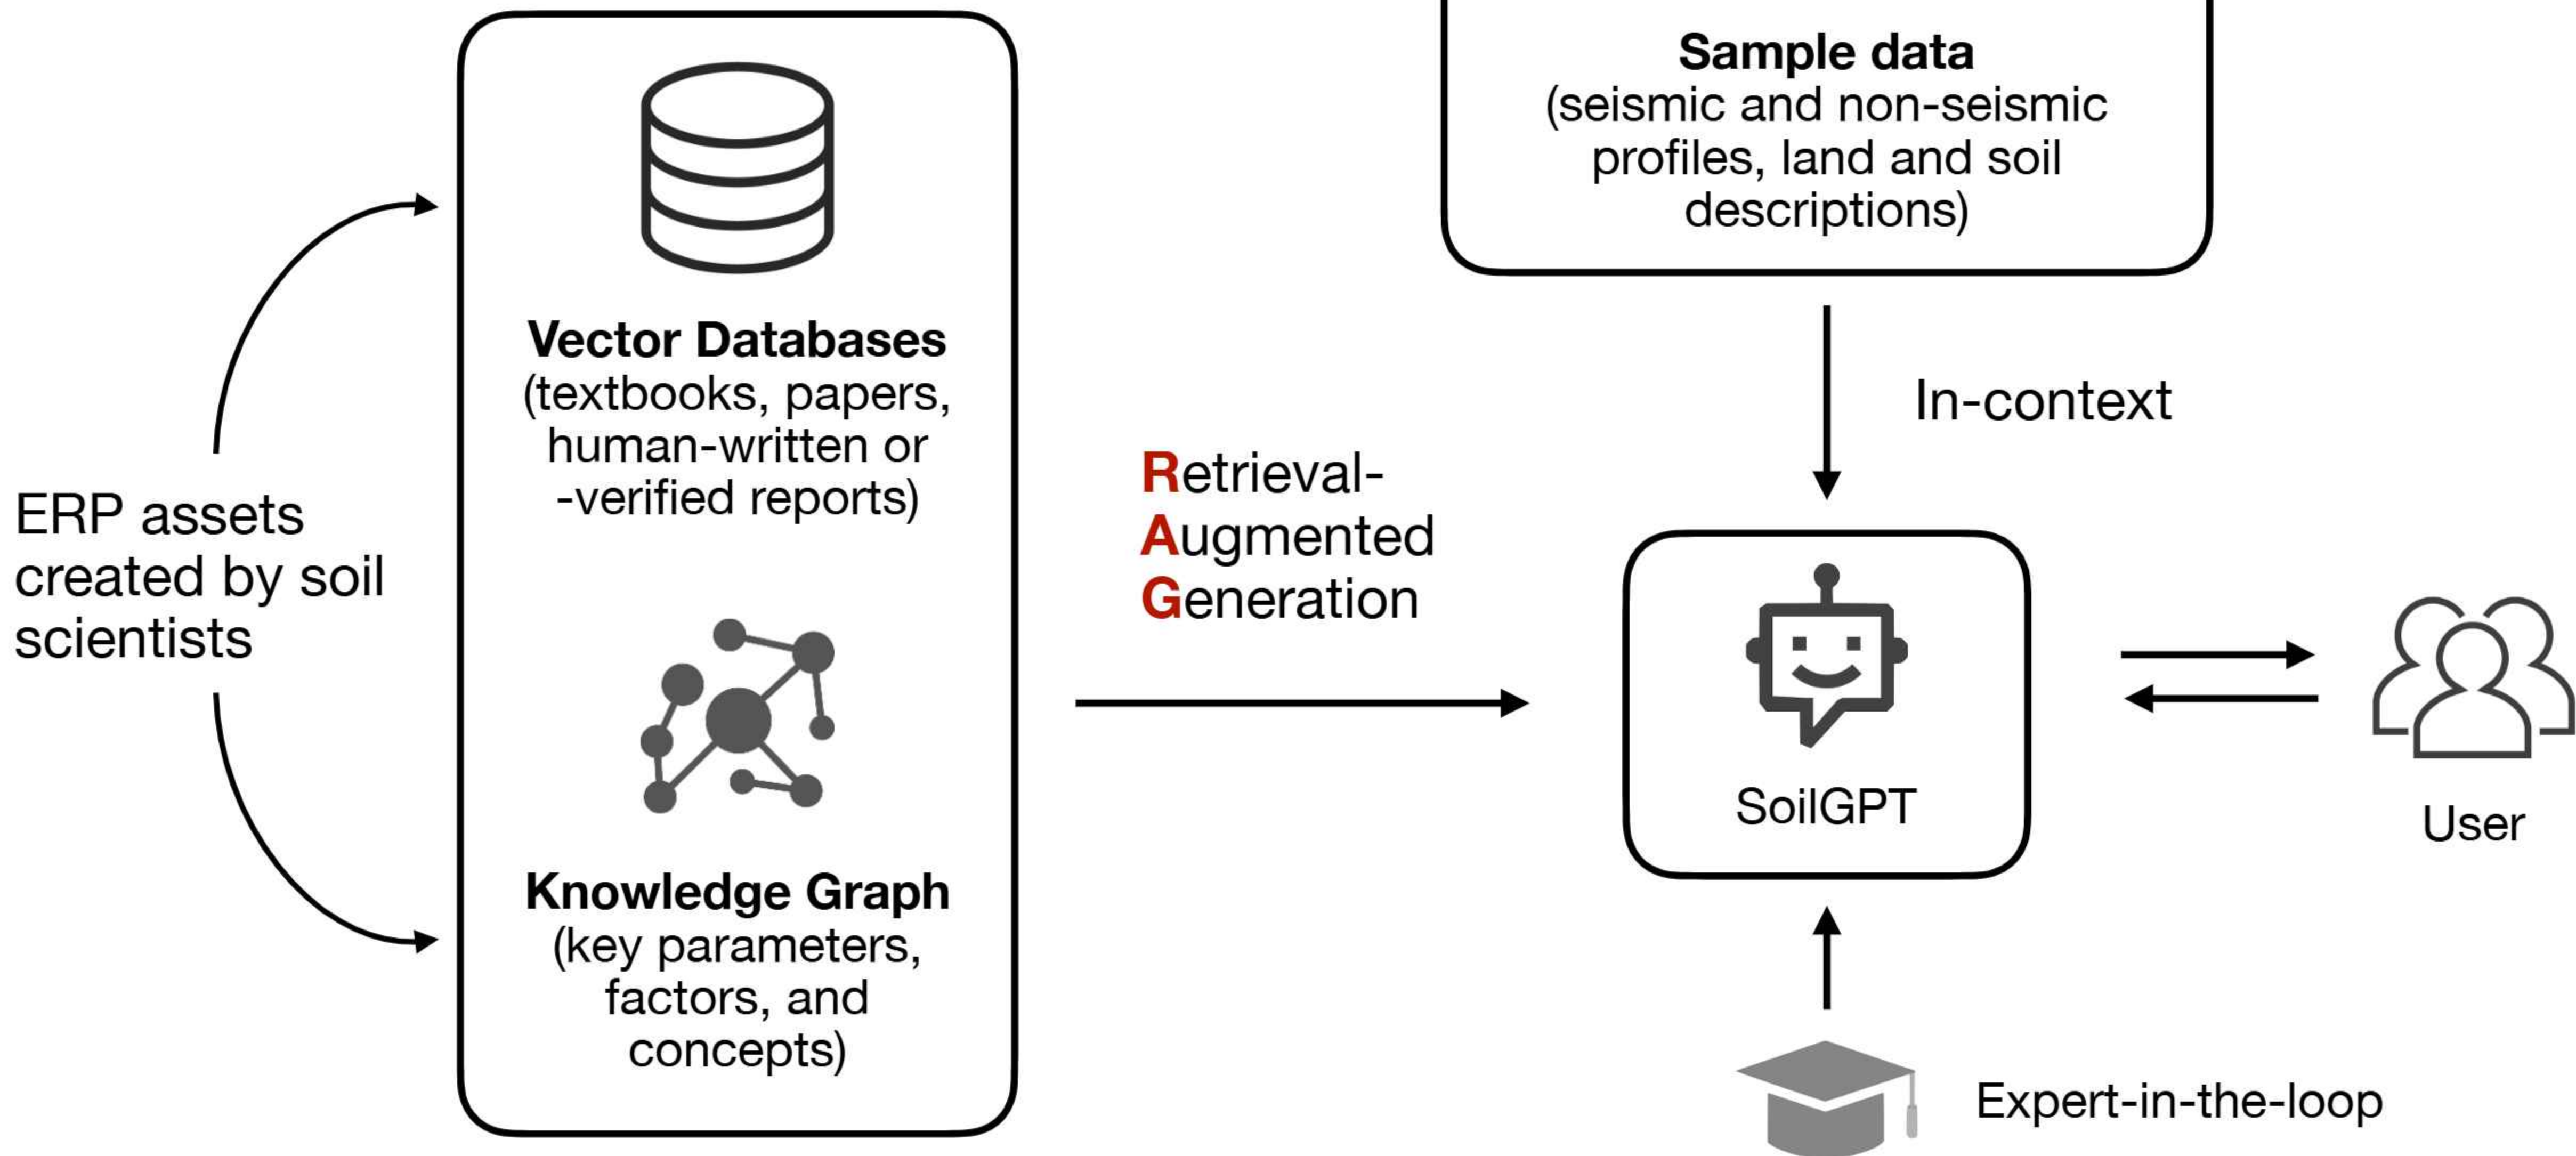


- LLM-specific
- Lighter development
- Simpler tasks
- GUI by LLM provider
- **Users pay for tokens**



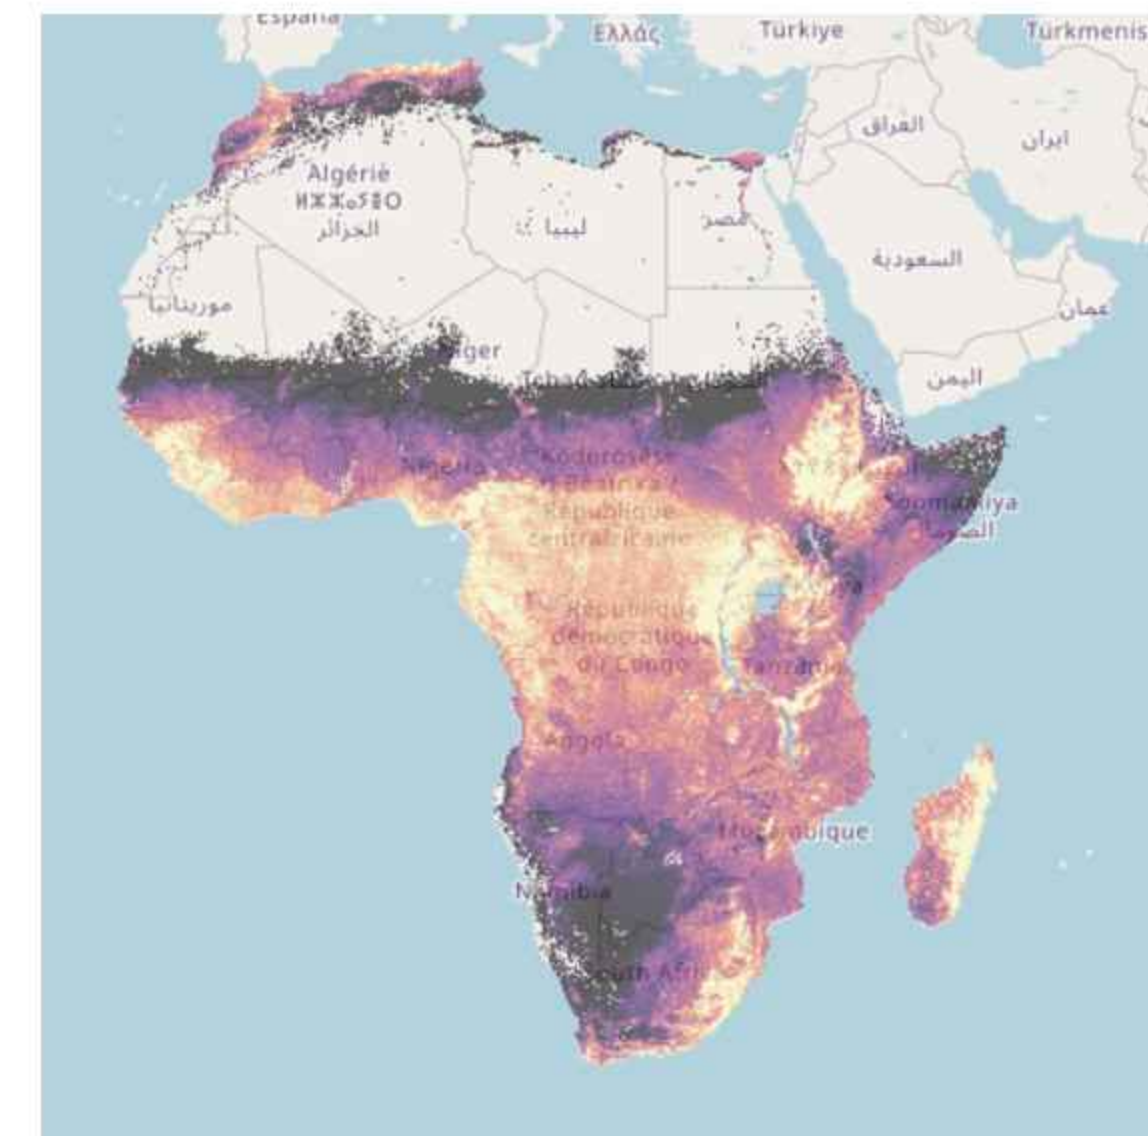
- LLM-independent
- Heavier development
- Complex tasks
- GUI by developer
- **Developers pay for tokens**

SoilGPT



SoilGPT-Africa

- We provide ChatGPT with
 - APIs for geographic reasoning (Nominatim + GADM)
 - APIs for retrieving pointwise soil data
 - APIs for visualising soil data on a map
 - Knowledge about the dataset
- ISDA Soil dataset (as a demo)
 - Data quantity: 45000+ Samples
 - 20+ columns (physical, chemical, fertility)



www.isda-africa.com

SoilGPT-Africa



TRY IT



SoilGPT-Africa

By K Leng 人

ChatGPT with access to soil data and visualisations across Africa
using custom tools.

What data can
you access?

What is soil
texture? Explain
with an example.

Show me crop
condition
in Tanzania.

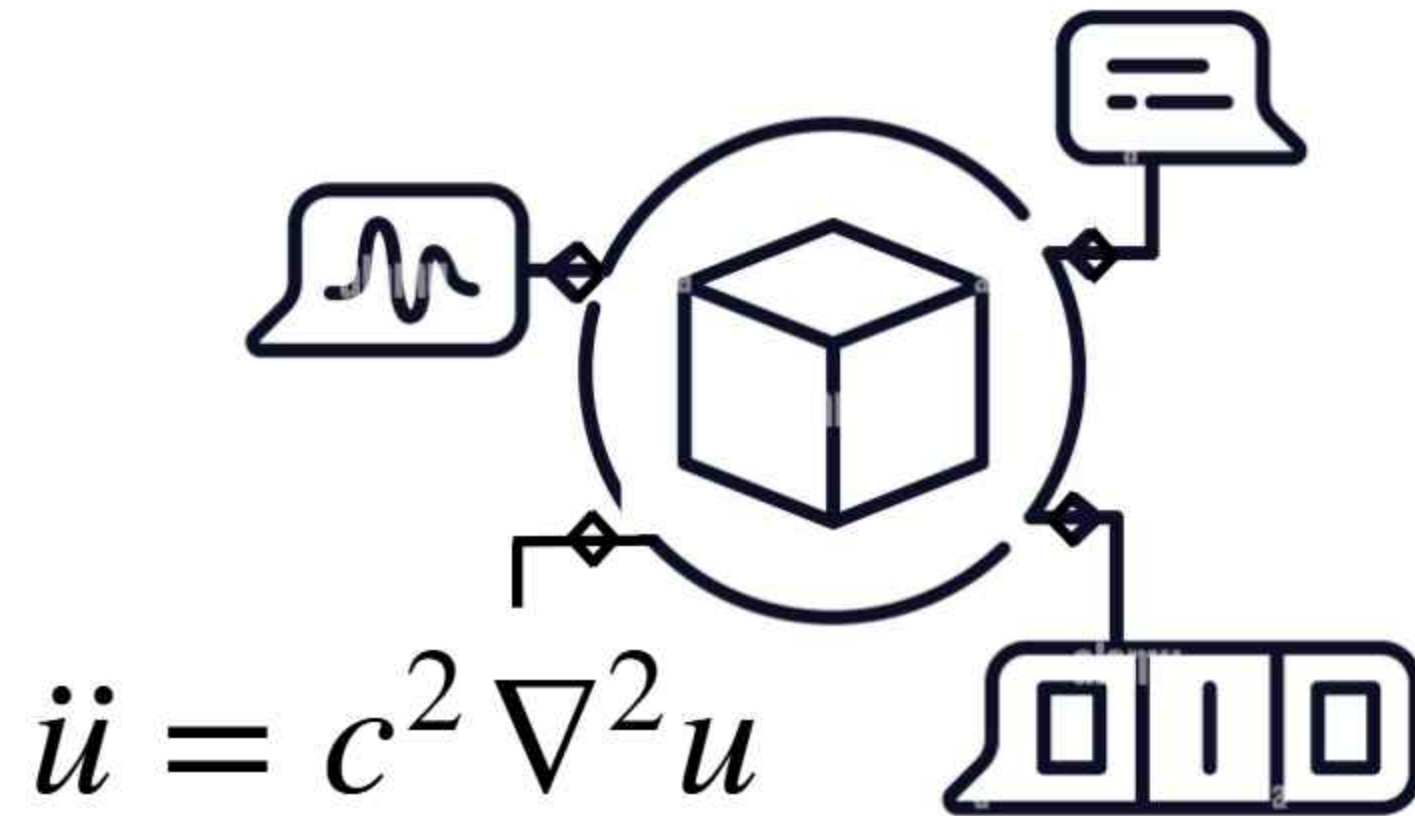
I want to grow
coffee in Sidama,
Bensa, Ethiopia....

AI Focus @ ERP

- Our AI Focus aims for impactful, science-driven solutions by bridging soil data, physics, and domain knowledge.

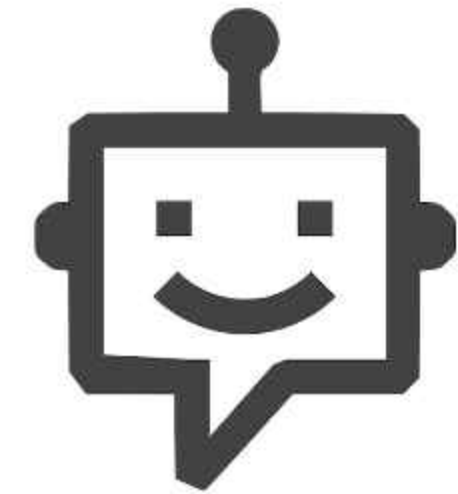


ML-Ready
Global Soil Data



Physics-informed,
Multi-modal ML

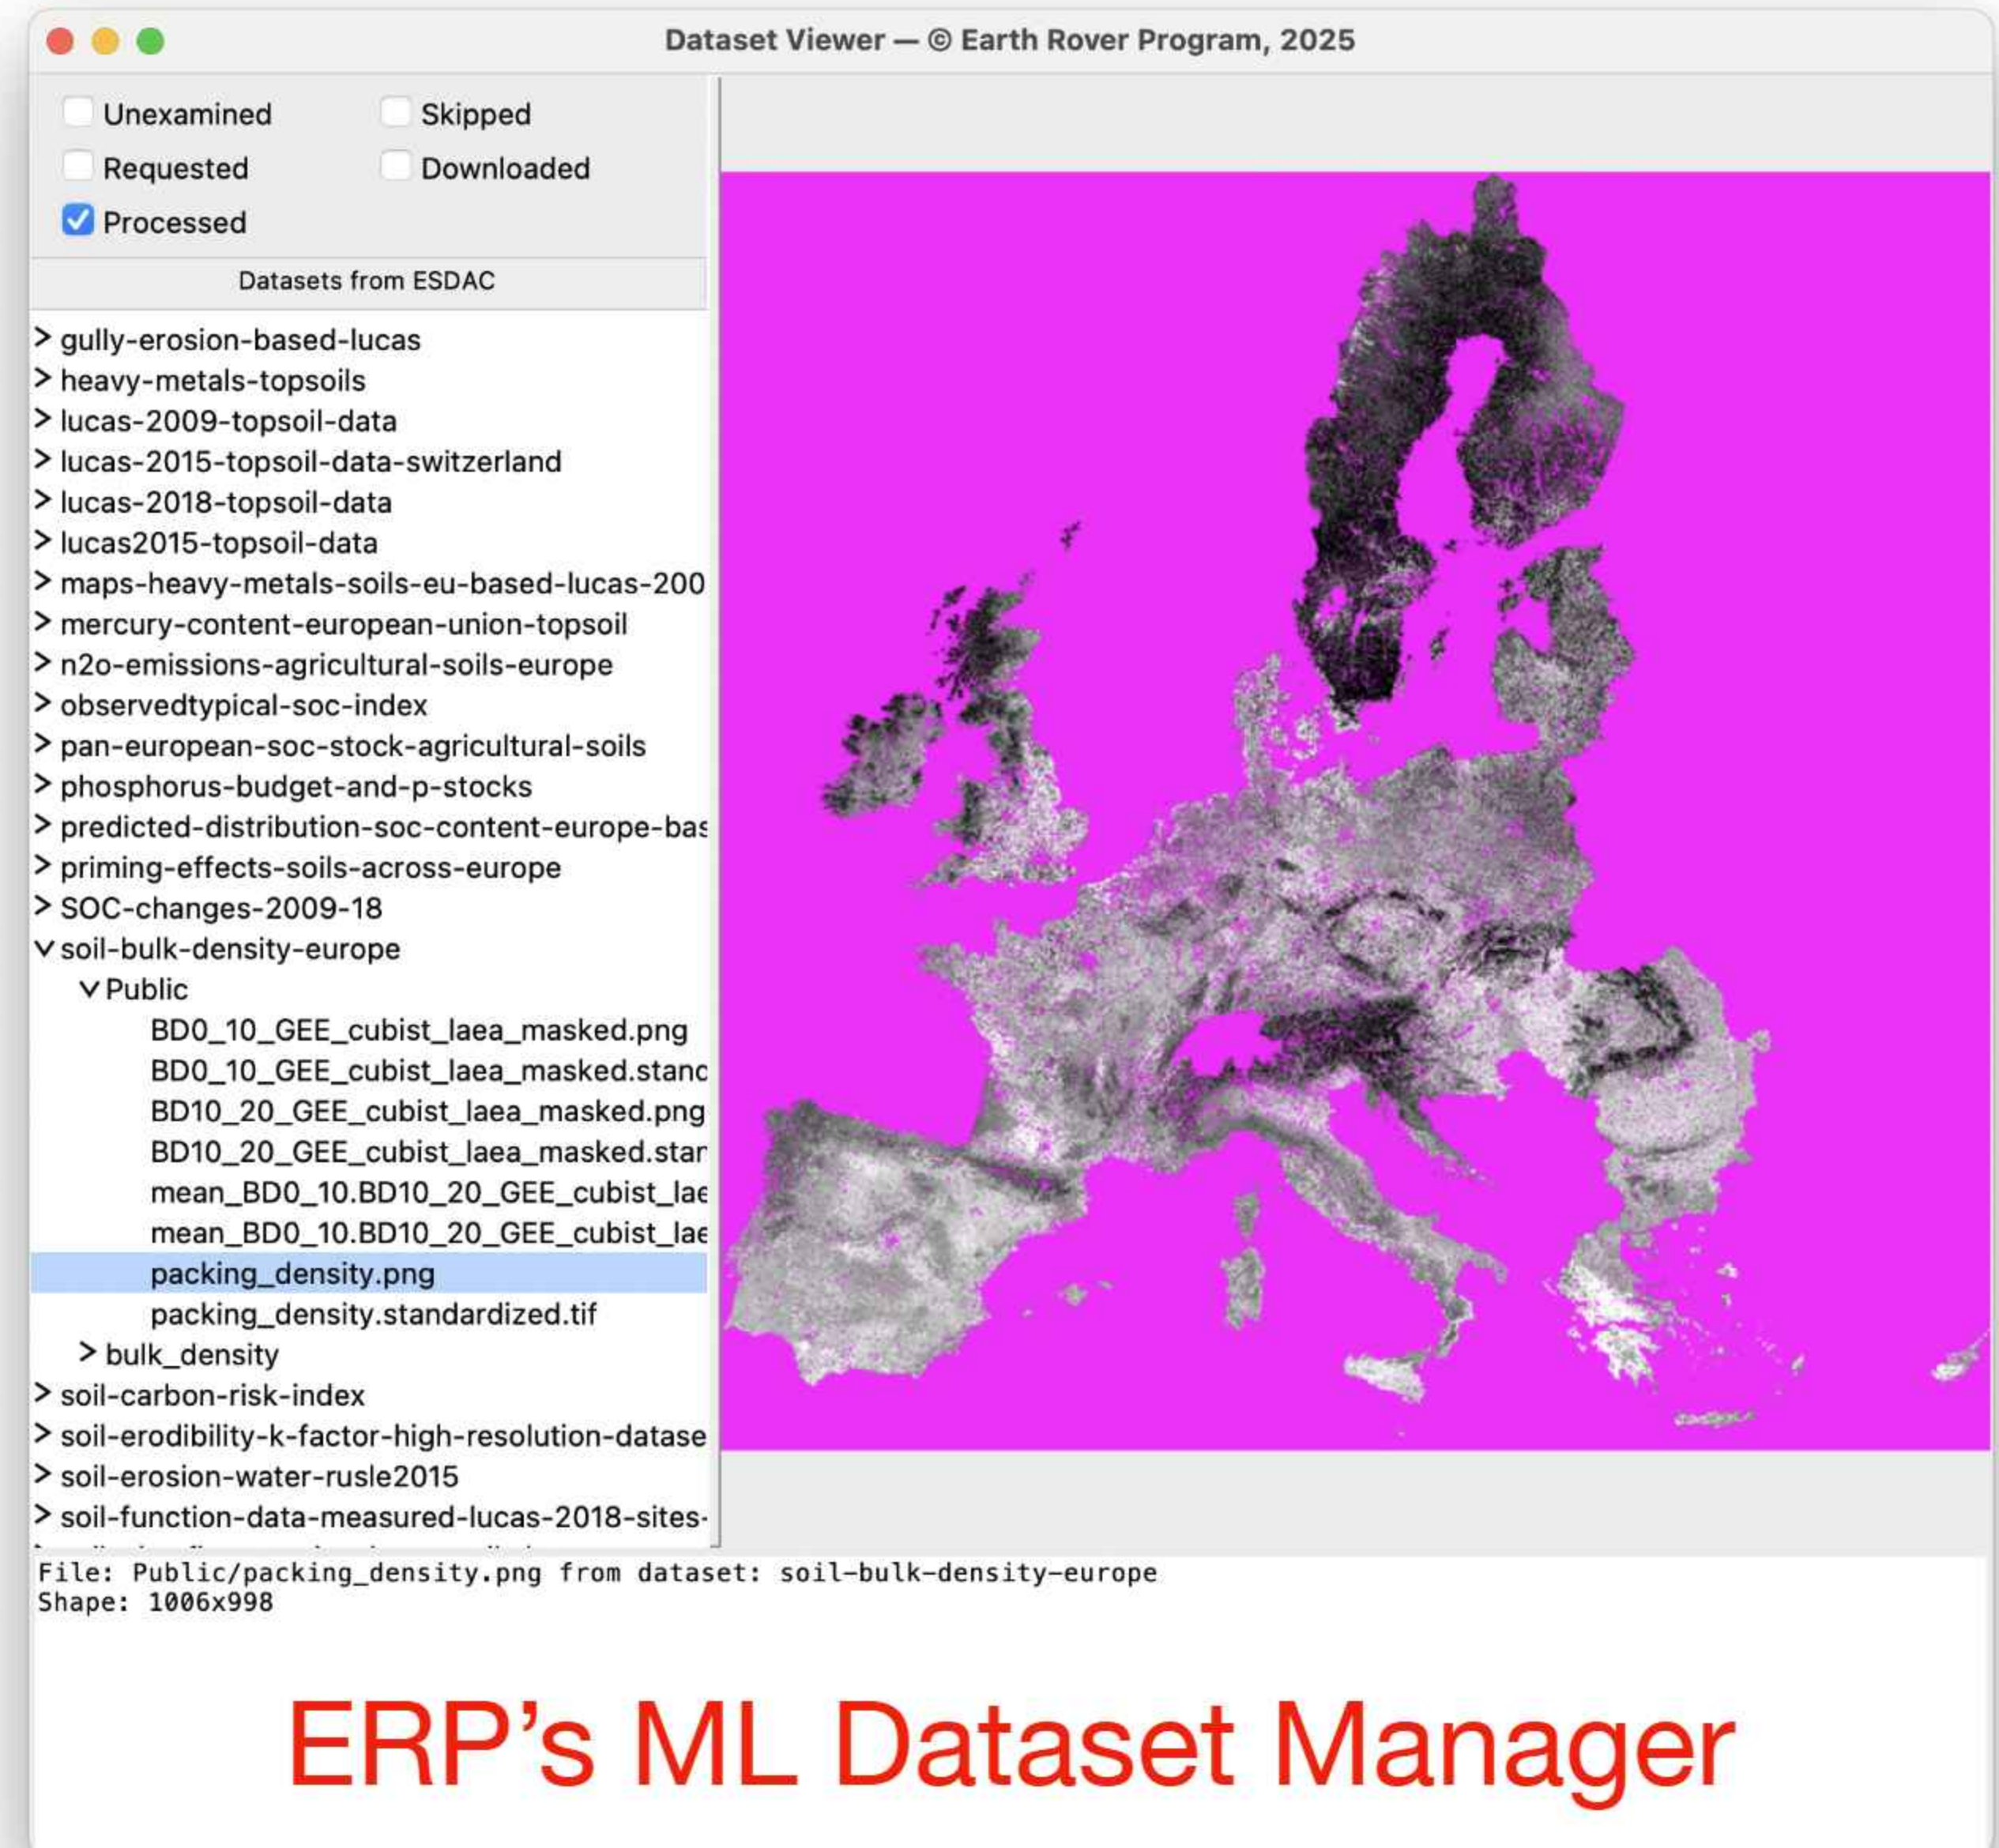
Data + Knowledge



SoilGPT

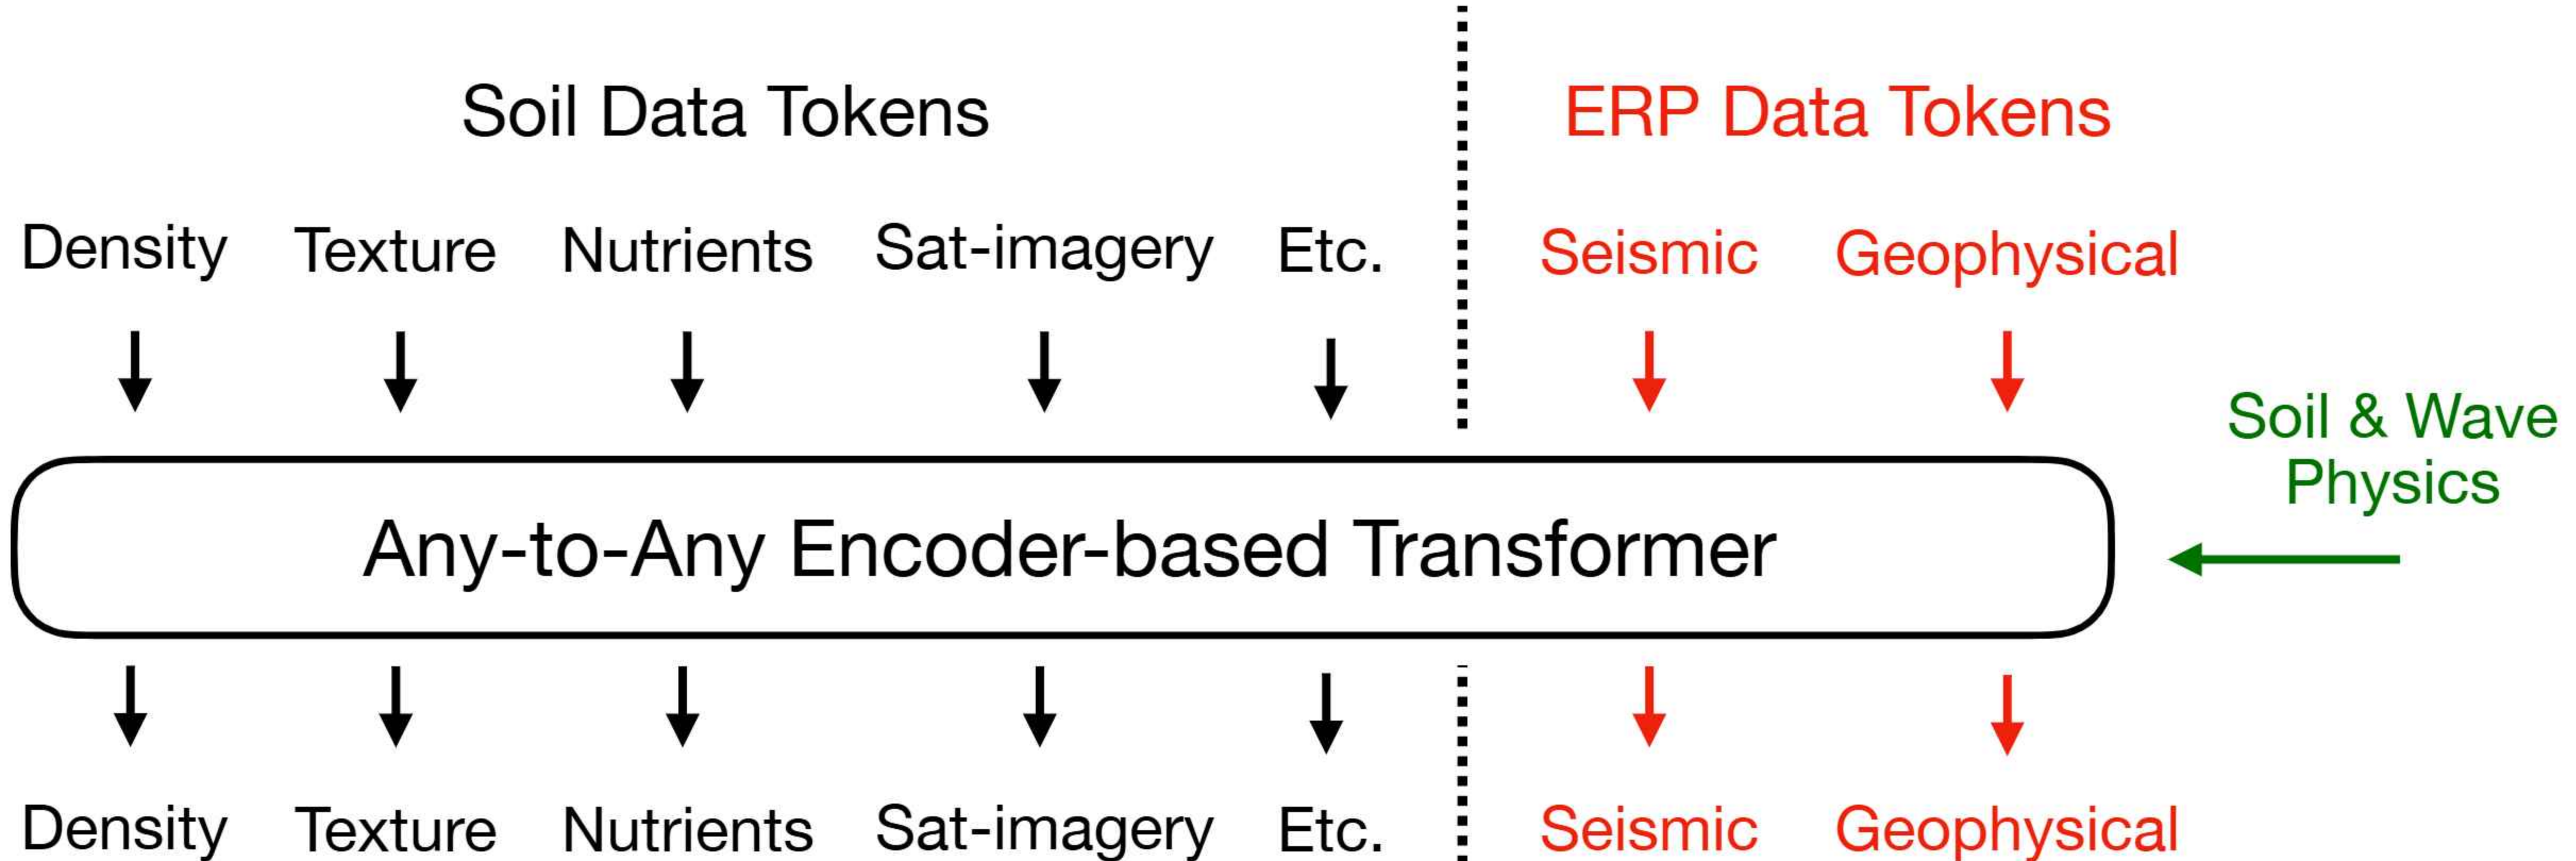
Soil Data

- We develop a **multi-agent system** to *automatically* process a large amount of available soil datasets
 - Global coverage: Europe, Africa, North America, South America, ...
 - Multidisciplinary: physical, chemical, biological properties; nutrients; functions; contamination & threats...
- **ML-Ready Corpus**
 - Sample-based structure with unified format and access
 - Explicit representation of uncertainty
 - Rich natural language annotations



Physics-informed, Multi-modal ML

- We unify **Soil Science** and **Seismology** through generalised **Masked Language Modelling (MLM)** on multi-modal tokens, integrating data-driven learning with soil & wave physics.



Thanks!



University of Exeter

Centre for Water Systems



Upgrading Sustainable Drainage Systems

Ben Hamilton

B.Hamilton3@exeter.ac.uk

SuDS upgrades – the context



<https://www.thames21.org.uk/sustainable-drainage-systems/>



University of Exeter

Centre for Water Systems

What is Sustainable Drainage?

- Attenuation for flood/CSO mitigation
 - Less “peaky” loads
 - Delivery over longer duration
- Water quality improvements
 - Within SuDS
 - At treatment works
- Biodiversity Net Gain
- Air quality
- Amenity



STORMWATER INTELLIGENCE GROUP

SuDS upgrades – the options

Source control

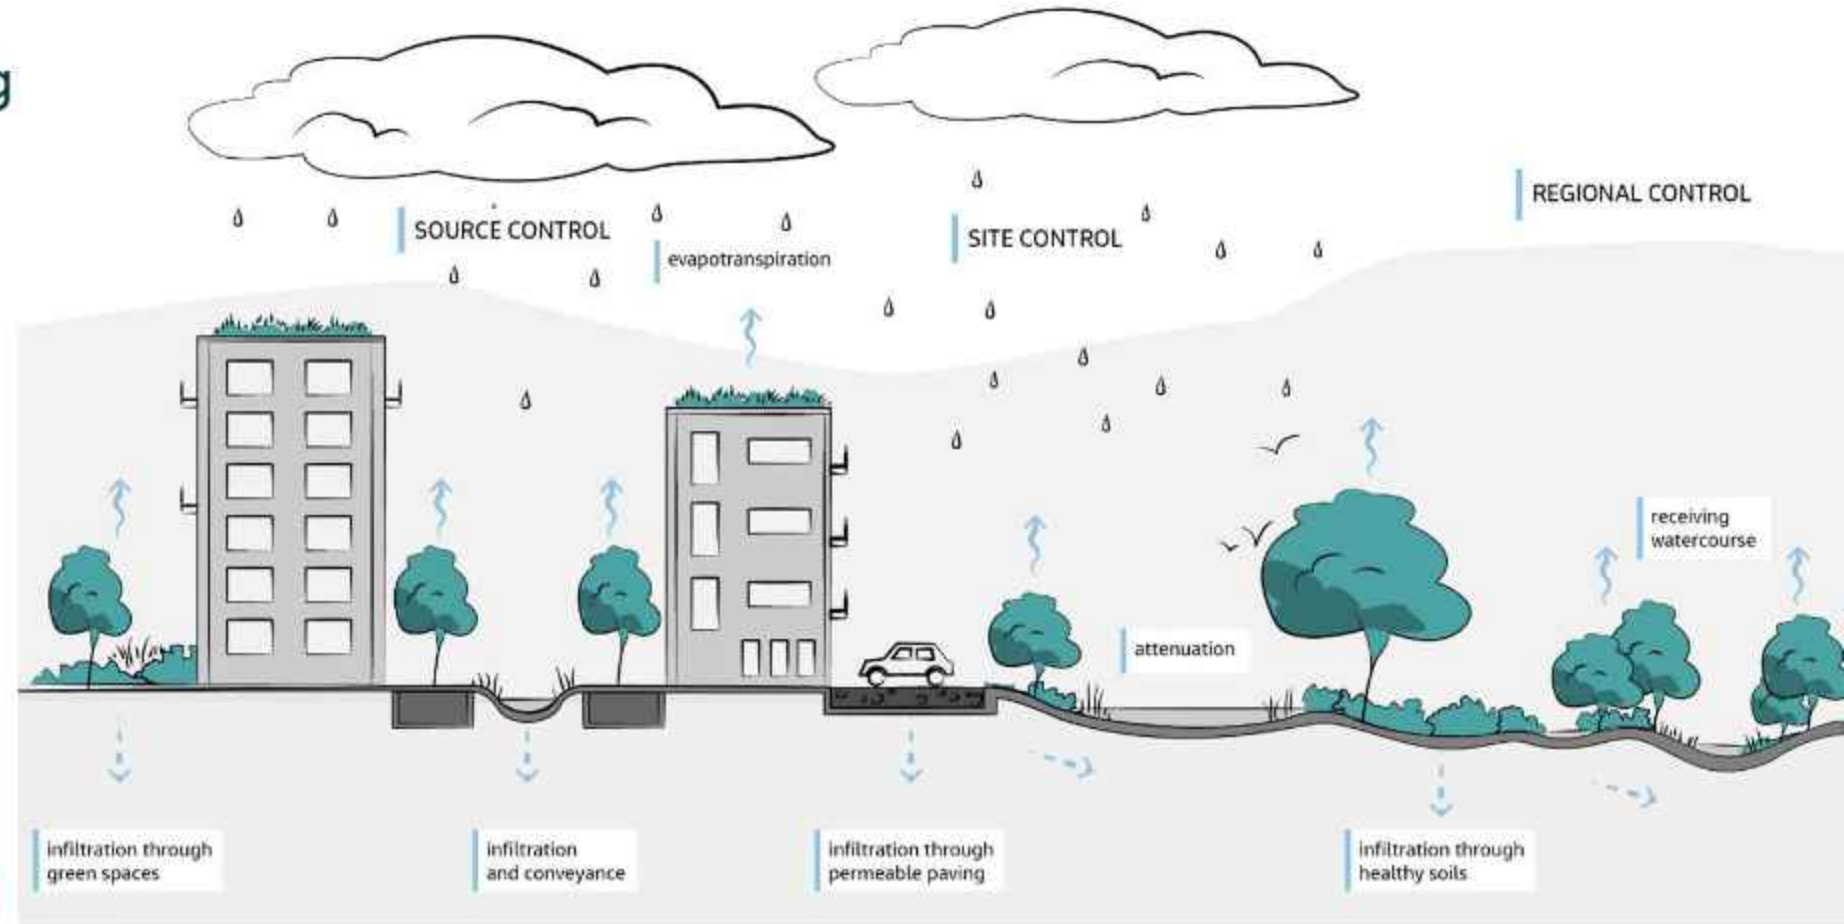
- Rainwater harvesting
- Blue/Green roofs
- Permeable paving
- Rain gardens

Site control

- Attenuation tanks
- Swales and ponds
- Soakaways
- Filter strips

Regional control

- Retention/Detention
- Constructed wetland



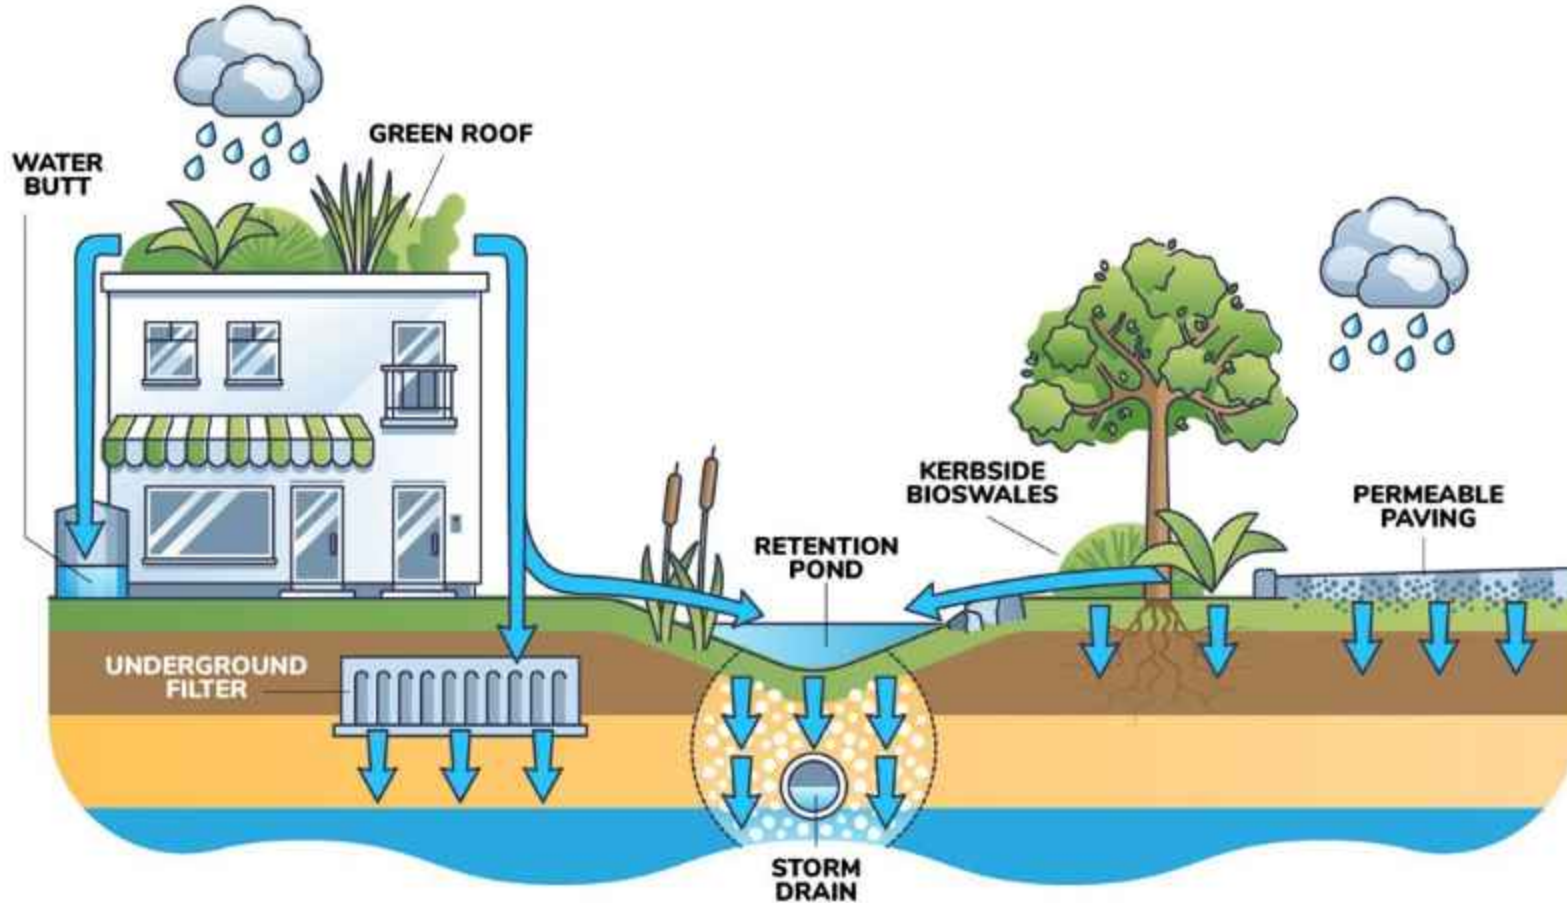
<https://naturalengland.blog.gov.uk/2023/11/08/green-infrastructure-creating-habitable-towns-and-cities-fit-for-the-future/>



University of Exeter

Centre for Water Systems

SuDS upgrades – the challenges



<https://mrms.co.uk/the-natural-approach-sustainable-drainage-systems-suds/>



University of Exeter

Centre for Water Systems

Systemic

- Climate change
- Urbanisation

Water systems

- Flood mitigation
- Sewer overflow
- Biodiversity
- Water quality/pollution

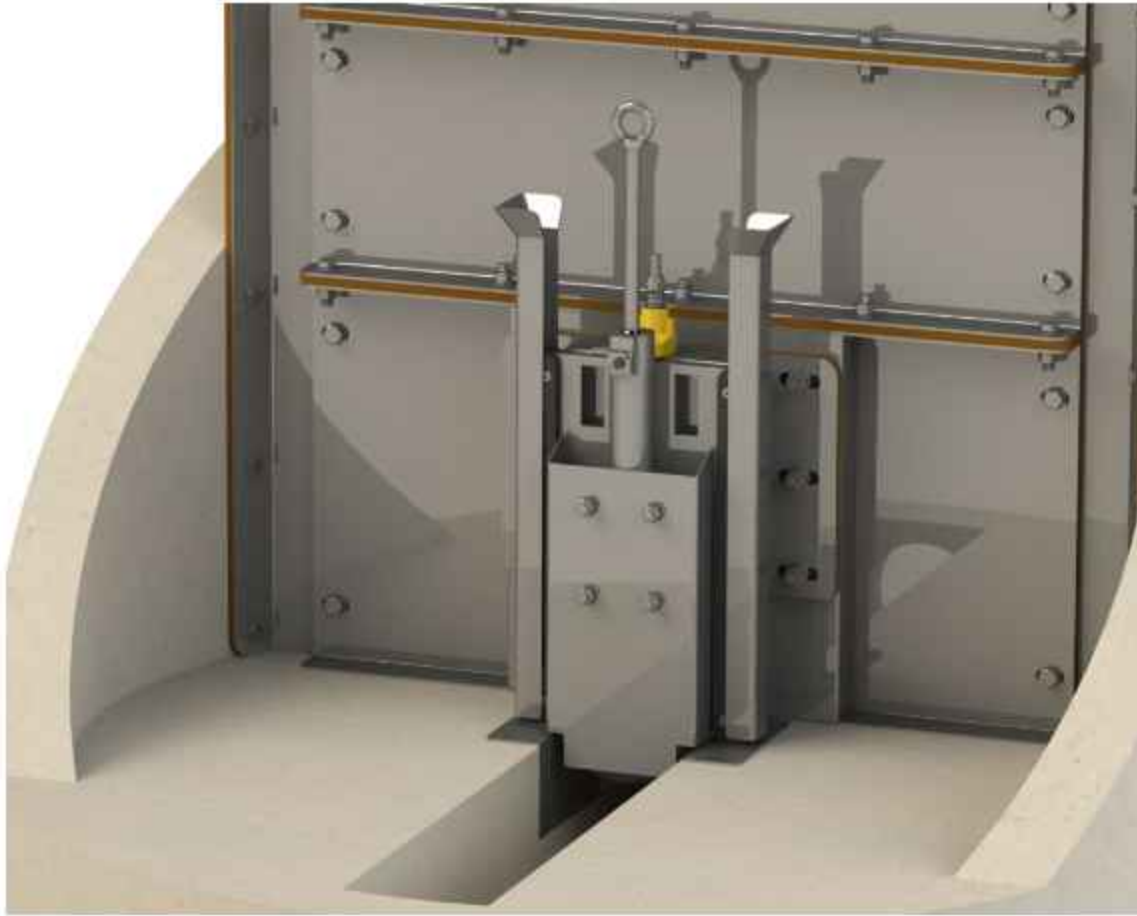
Infrastructure

- Accessibility
- Aging asset
 - Capacity
 - Maintenance
- Passive assets



STORMWATER INTELLIGENCE GROUP

SuDS upgrades – our retrofit, real time focus

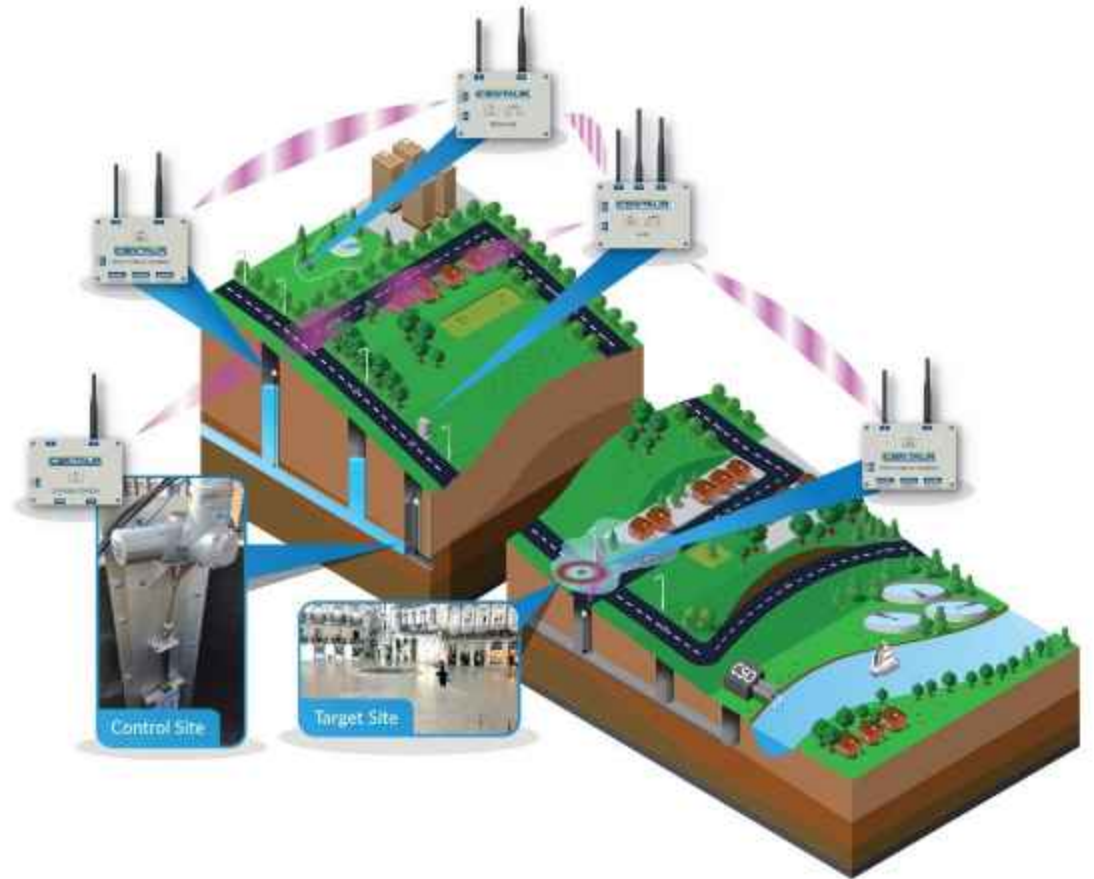


<https://www.radio-data-networks.com/wastewater-flow-regulators/>



University of Exeter

Centre for Water Systems



<https://floodlist.com/europe/artificial-intelligence-and-urban-flood-prevention>

SuDS upgrades – Campus sensing-RTC pathway



Synergistic sensing – CREWW

- Understanding the Lower Hooper Valley

Armory rain gauges

- Are all devices fit for purpose?

Identifying Sustainable Drainage assets

- Planning applications review
- ‘As built’ schematics sourced

Benchmarking performance

- Sewer Level Monitoring

Campus Real Time Control

- Radio Data Network Sewer Flow Regulator



University of Exeter

Centre for Water Systems



STORMWATER INTELLIGENCE GROUP

SuDS upgrades – RTC lab tests

Lab methodologies testing:

- Water tightness
 - Seal damage
 - Range in head
- Blockage response
 - Rags
 - Debris
 - Sediment
- Actuation

All are welcome to drop by!



Photo credit: Ben Hamilton



University of Exeter

Centre for Water Systems



STORMWATER INTELLIGENCE GROUP



University
of Exeter

Centre for
Water Systems



STORMWATER INTELLIGENCE GROUP



University
of Exeter

*Centre for Resilience in Environment,
Water and Waste*

Can we predict harmful algal blooms in lakes?

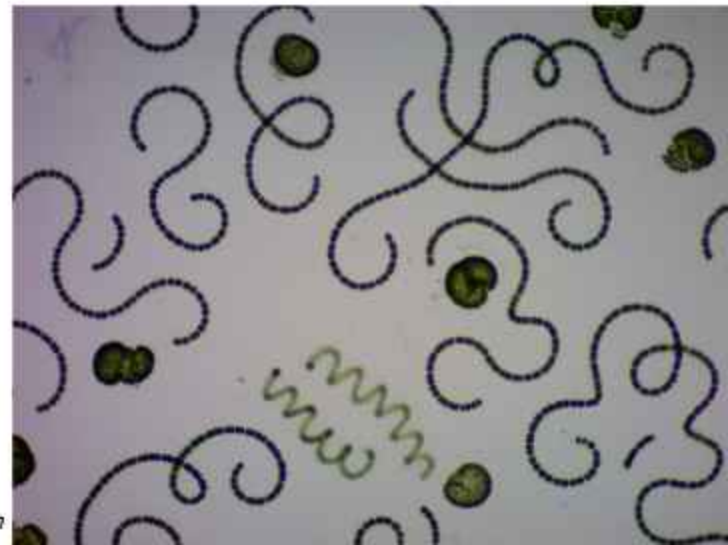
**Harnessing AI and heterogeneous datasets
for decision tools and digital twins**

Dr Diego Panici
Senior Lecturer in Hydrological and Water Quality Modelling

What is the issue?

Harmful algal blooms

- HABs form in lakes or reservoirs
- Although there are different types of algal blooms, most dangerous 'algae' are cyanobacteria also known as blue-green algae
- Cyano-HABs particularly thrive when:
 - High load of nutrients (N and P) concentration
 - Warm water temperature
 - High sunlight levels



What is the issue?



University
of Exeter

Then we have...

➤ Effects can be nasty:

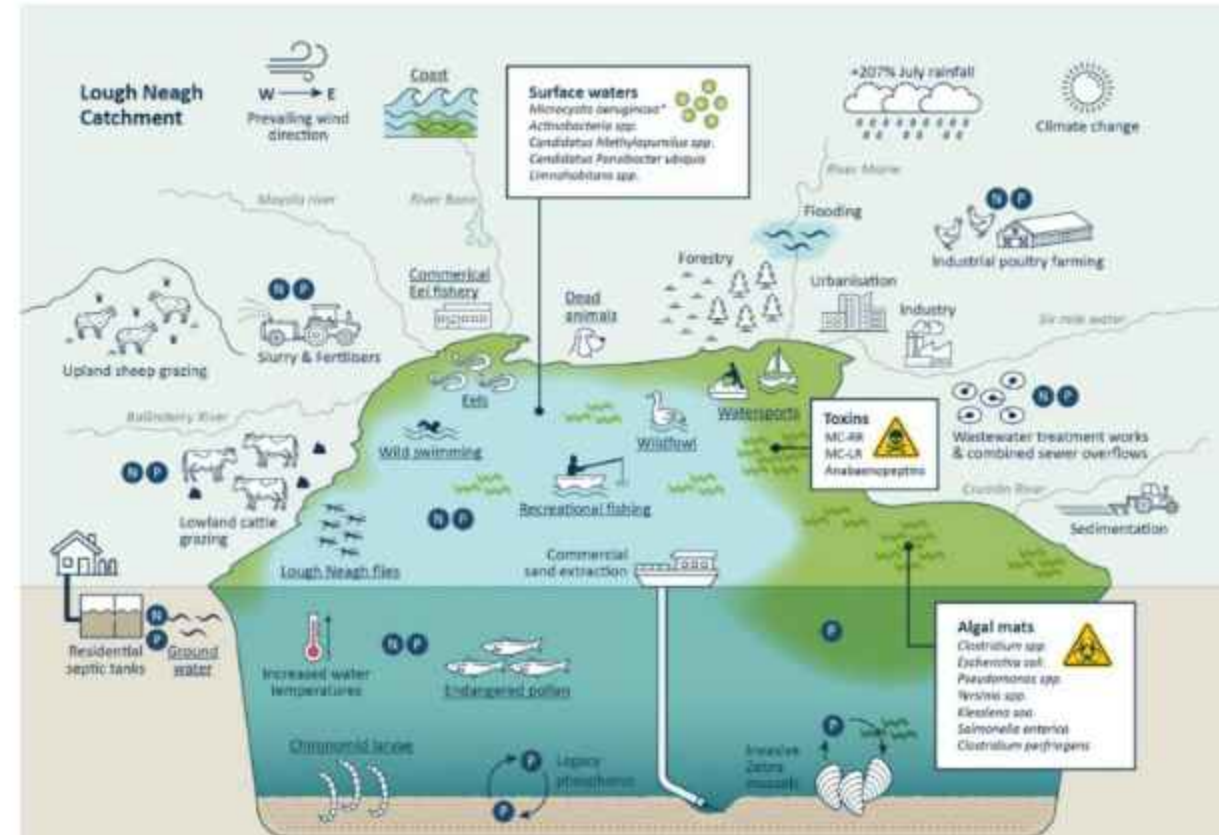
- Creates hypoxic (or even anoxic) conditions
- Can be lethal for wildlife and toxic for human consumption
- Huge disruption to water supply from cyanobacteria decay



What is the issue?

For example, Lough Neagh...

- Regular occurrence of HABs
- Ongoing (made the news)
- Severe disruption to 40% of Northern Ireland water resources



Source: Reid et al. (2024)

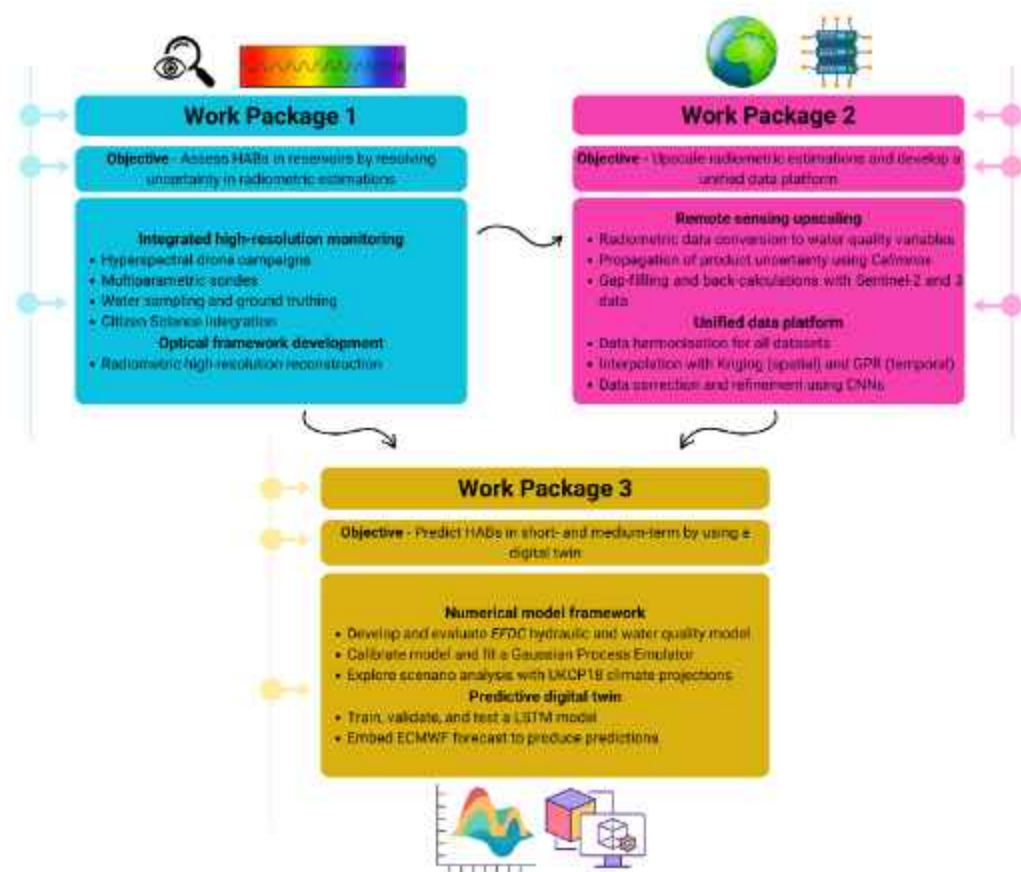
So what?

What if we could leverage multiple sources of data in space and time to predict a bloom?

Prediction and Early Detection of harmful ALgal blooms in lakes and reservoirs



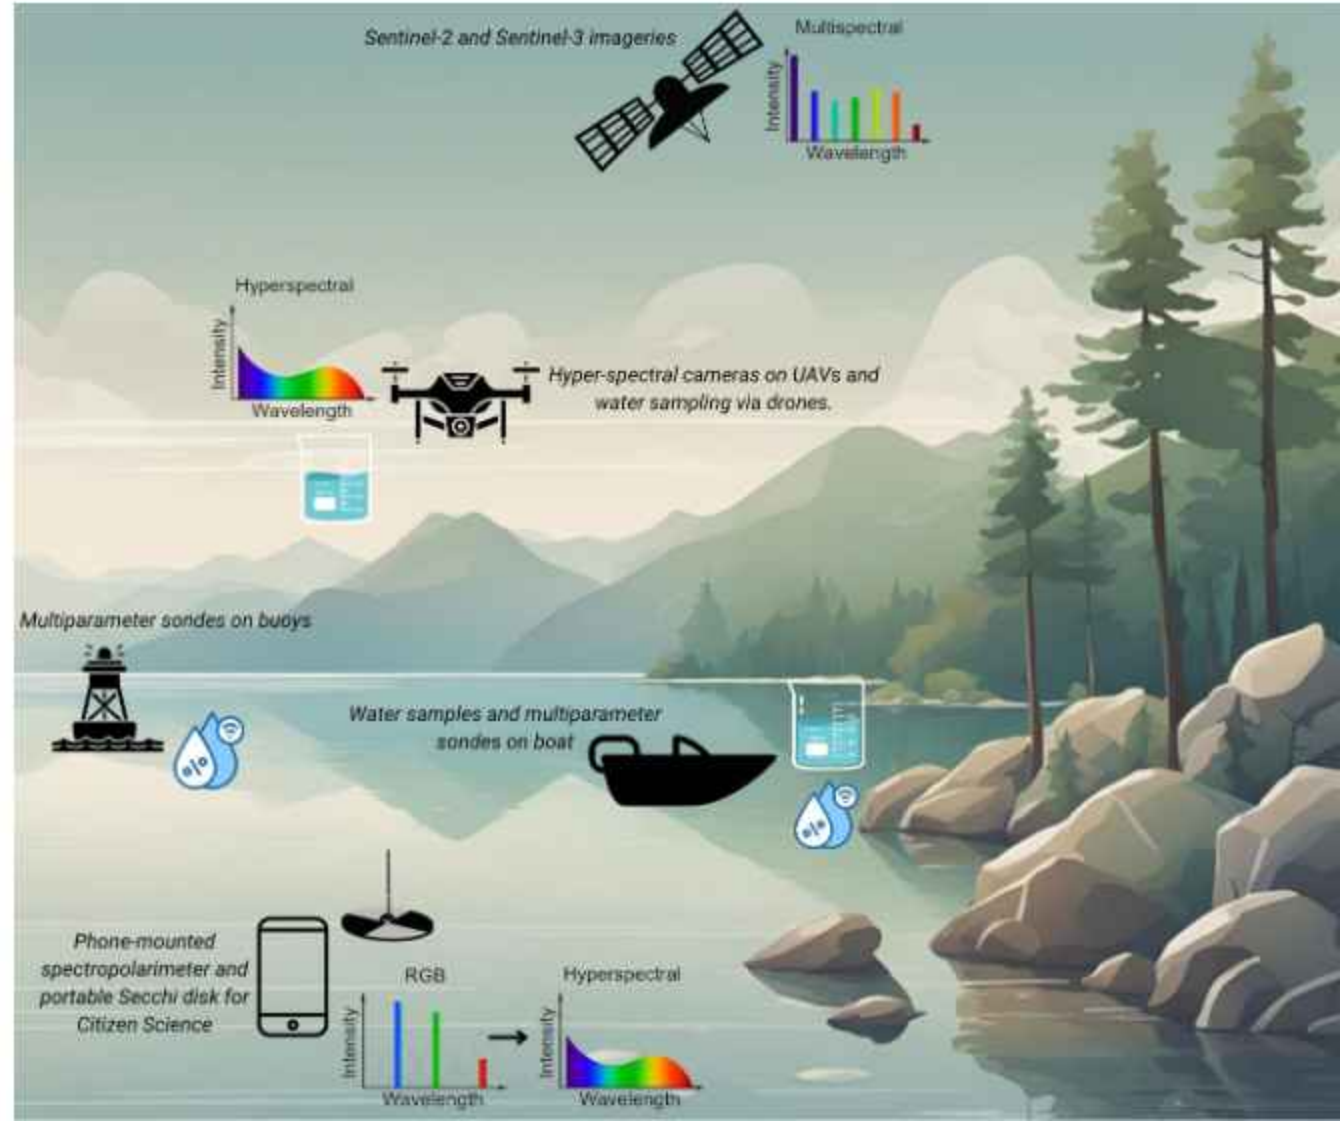
- WP1 – multi-scale nested monitoring
- WP2 – Upscaling and unified data framework
- WP3 – Numerical modelling and digital twin development



Monitoring

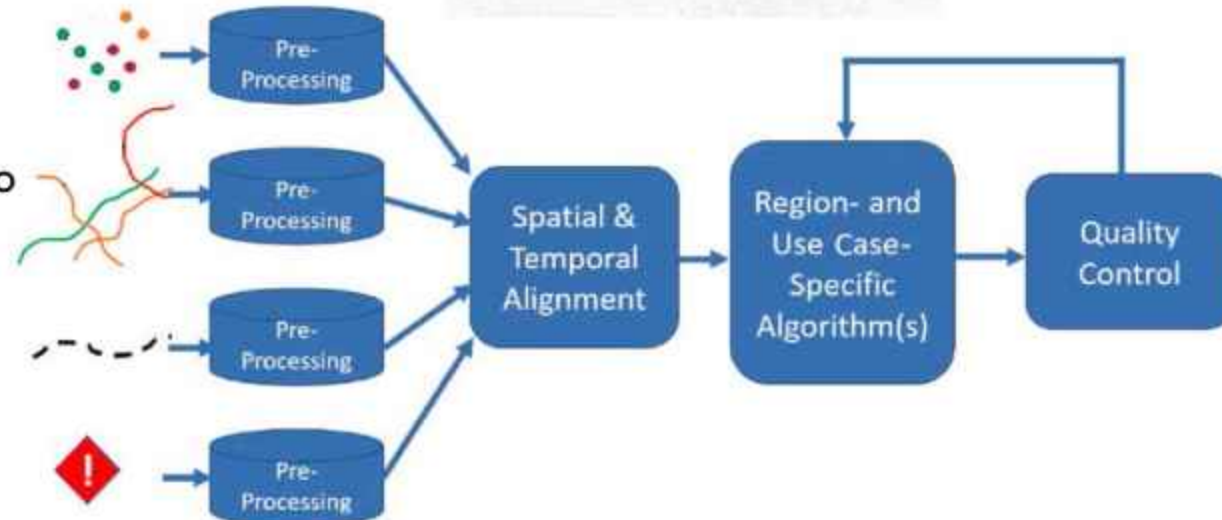
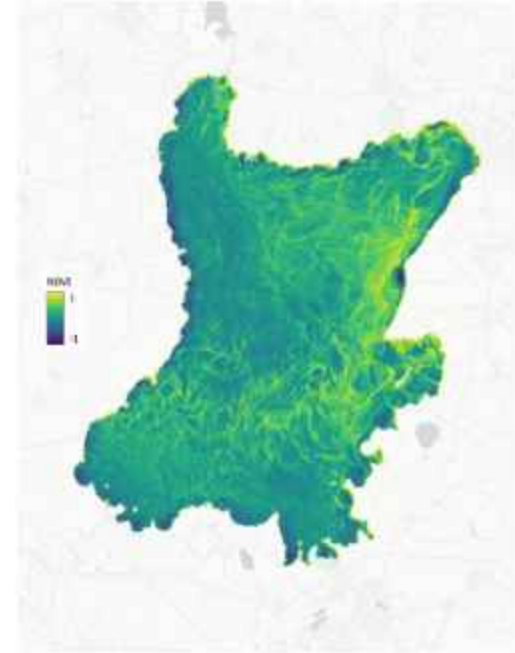
➤ Data sources

- Drone-based hyperspectral imaging (*spatial detail*)
- High-frequency multiparametric sondes (*temporal resolution*)
- Water sampling (*validation and reference*)
- Satellite/remote sensing (*scaling up*)
- Citizen science (*broad coverage, engagement*)



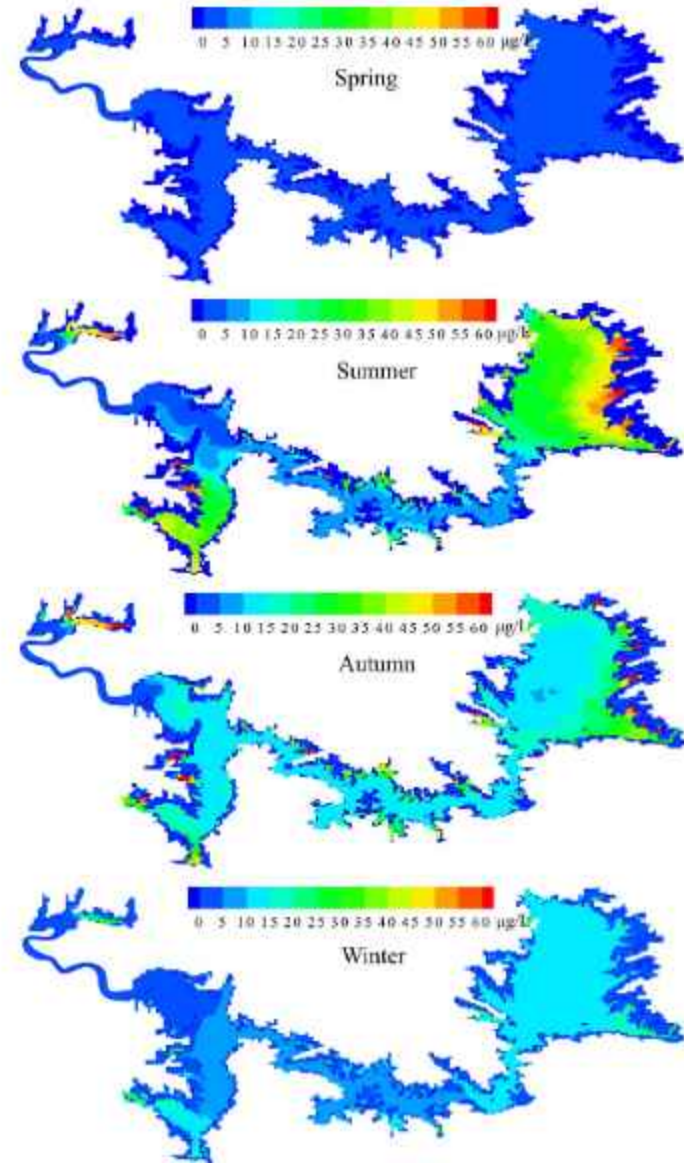
Upscaling

- Remote sensing uncertainty estimation and reduction
 - Calibration of remote sensing (e.g., haze or reflections near the shore) validated against observations in WP1
 - Target critical spectral wavelengths not fully captured by current multispectral satellites, to identify cyano-HABs
 - Propagate uncertainty to quantify reliability of estimations
- Data fusion
 - Integrating multiple heterogeneous datasets in space and time
 - Multi-scale linkage (punctual measurements ↔ lake-level estimations)
 - Gap-filling using interpolation and AI-based techniques into a unified spatial-temporal map



Detecting and predicting

- Numerical model for scenario analysis
 - Developing fully coupled hydrodynamic-water quality model
 - History-matching of recorded events to calibrate model parameters
 - Scenario analysis including climate change (e.g., UKCP18 climate projections) and water quality inputs
- Digital twin
 - Leverage the large amount of data (observed and simulated)
 - Including near-real-time early detection, climate projections, and what-if scenarios
 - Integration of medium-term weather forecast data (e.g., from ECMWF) for medium-term predictions



So what is the *Environmental Intelligence* vision?



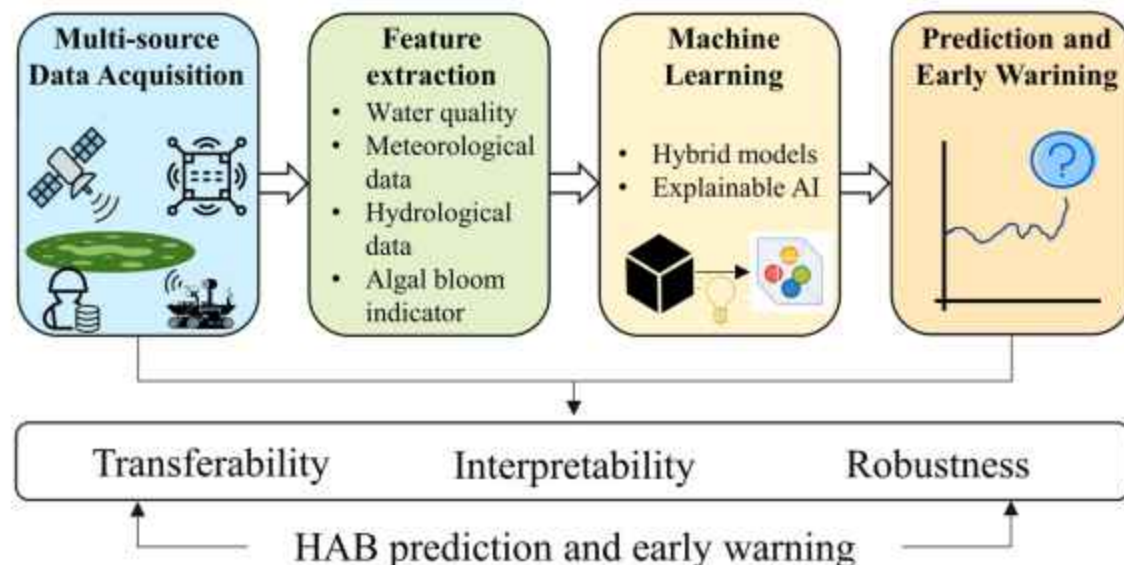
➤ Heterogeneous datasets

- Link multi-scale heterogeneous datasets (satellite ↔ drones ↔ in-situ ↔ citizen science) into a single workflow
- Spatial-temporal interpolation and gap-free maps: interpolation techniques or deep-learning approaches (e.g., CNN for spatial feature extraction and LSTM for temporal models)?

➤ Digital twin development

- Potentially 100s or 1000s of datasets to build upon
- We want to leverage this data for long-range temporal predictions
- Will need to intrinsically have spatial-temporal structure (and an uncertainty estimation framework)

Wang et al. (2025)



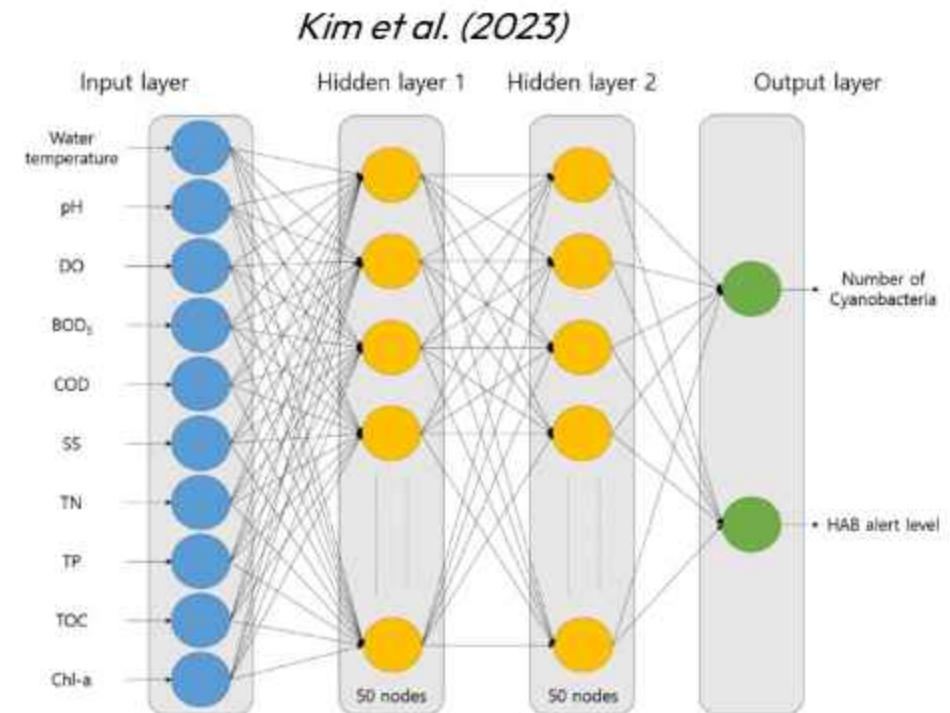
So what is the *Environmental Intelligence* vision?

➤ Key challenges

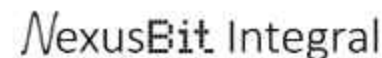
- Data imbalance (blooms occur only at a fraction of the total time observation)
- Hybrid approach of physics-based models and data science
- Uncertainty quantification and confidence

➤ Where do we go from here?

- What type of predictive model should we prioritise?
- What's the most effective way to propagate (and communicate!) uncertainty?
- How can we generalise this framework to out-of-project lakes/reservoirs?



Thank you!





University
of Exeter

Advancing Non-Target Water Chemistry Through Environmental Intelligence

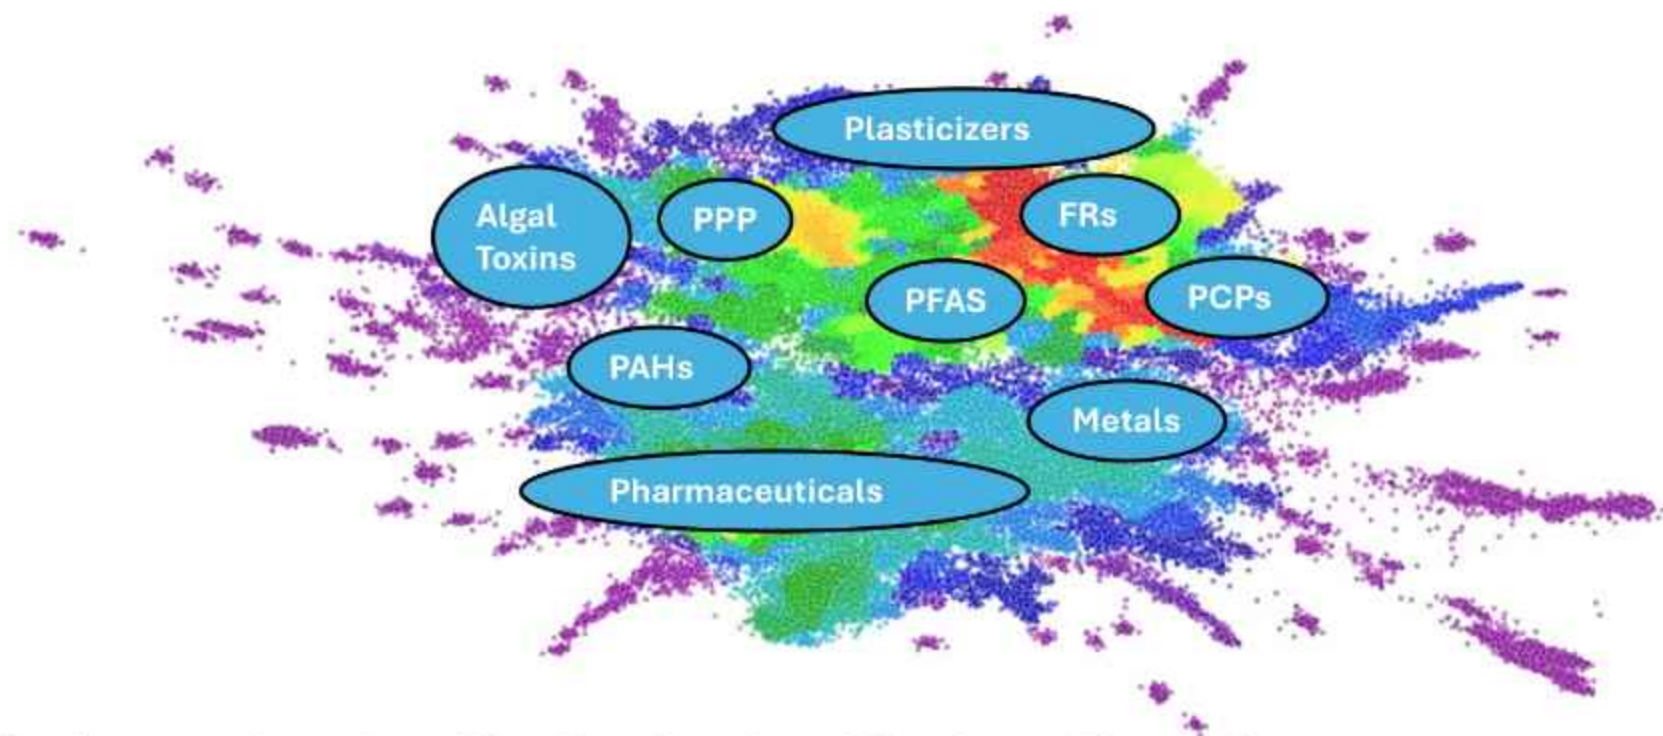
Garry Codling –EI Presentation 2025

Regulatory Landscape

Chemical Category	Examples
Heavy Metals	Mercury, cobalt, etc.
Organic BPA & Siloxanes	Bisphenol A, siloxanes
PFAS / PFCs	Perfluorinated compounds (broad PFAS class)
PBDEs	Flame retardants (polybrominated diphenyl ethers)
PPCPs & Pharmaceuticals	Metformin, antibiotics, analgesics, beta-blockers, DEET, cotinine
Emerging Micropollutants	Caffeine, artificial sweeteners, triclosan/carban, phthalates, fragrances
Illicit Drugs & Metabolites	Norfentanyl, xylazine, carfentanil

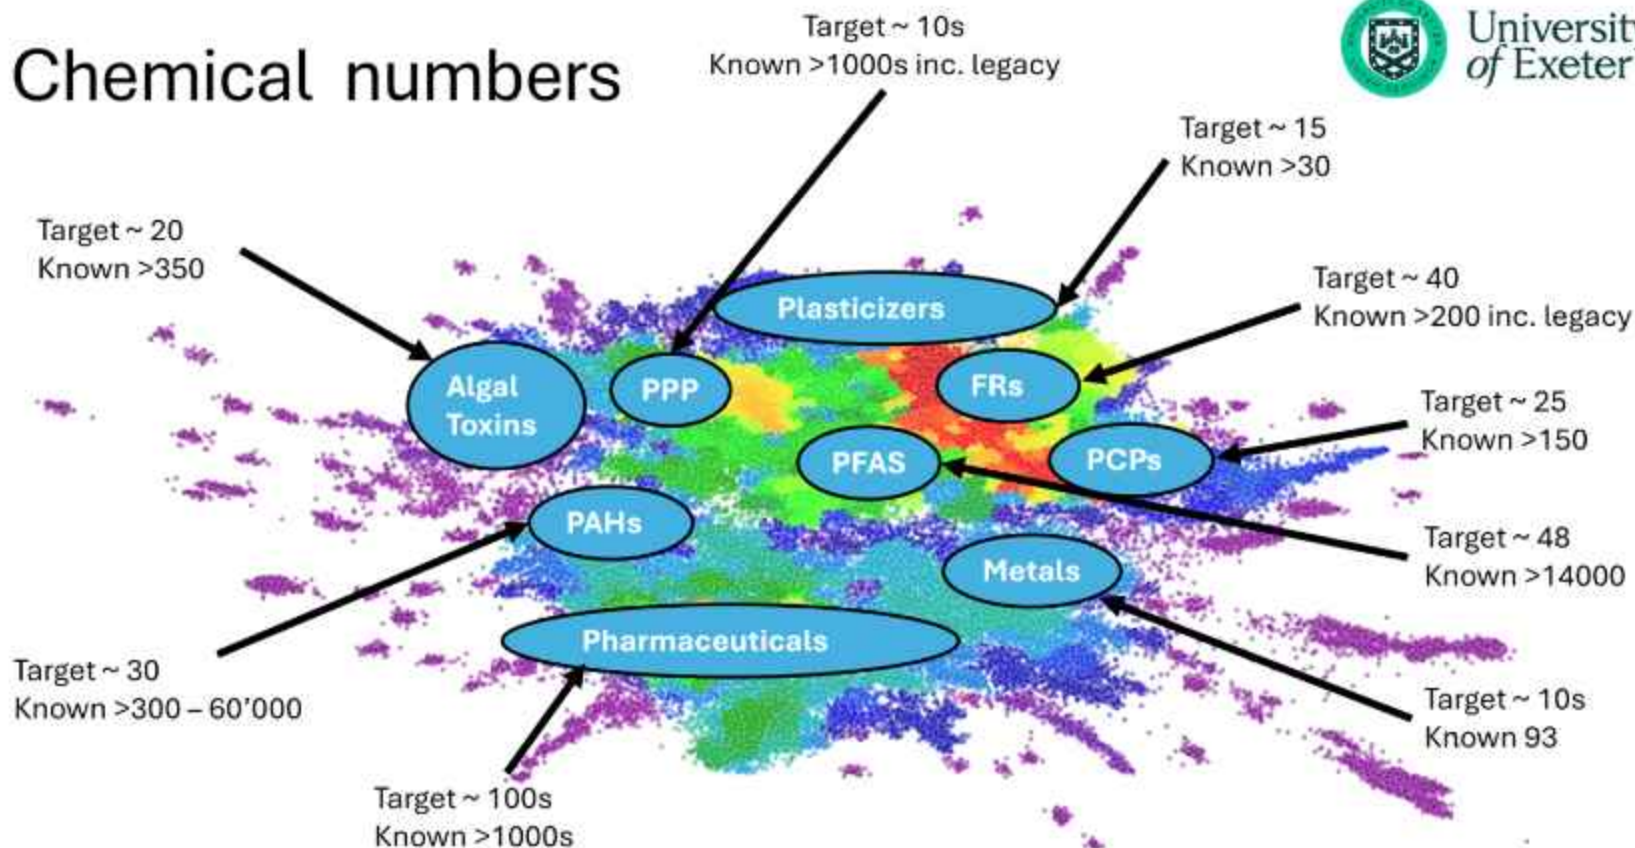
We can't monitor what we don't know

Chemical Space



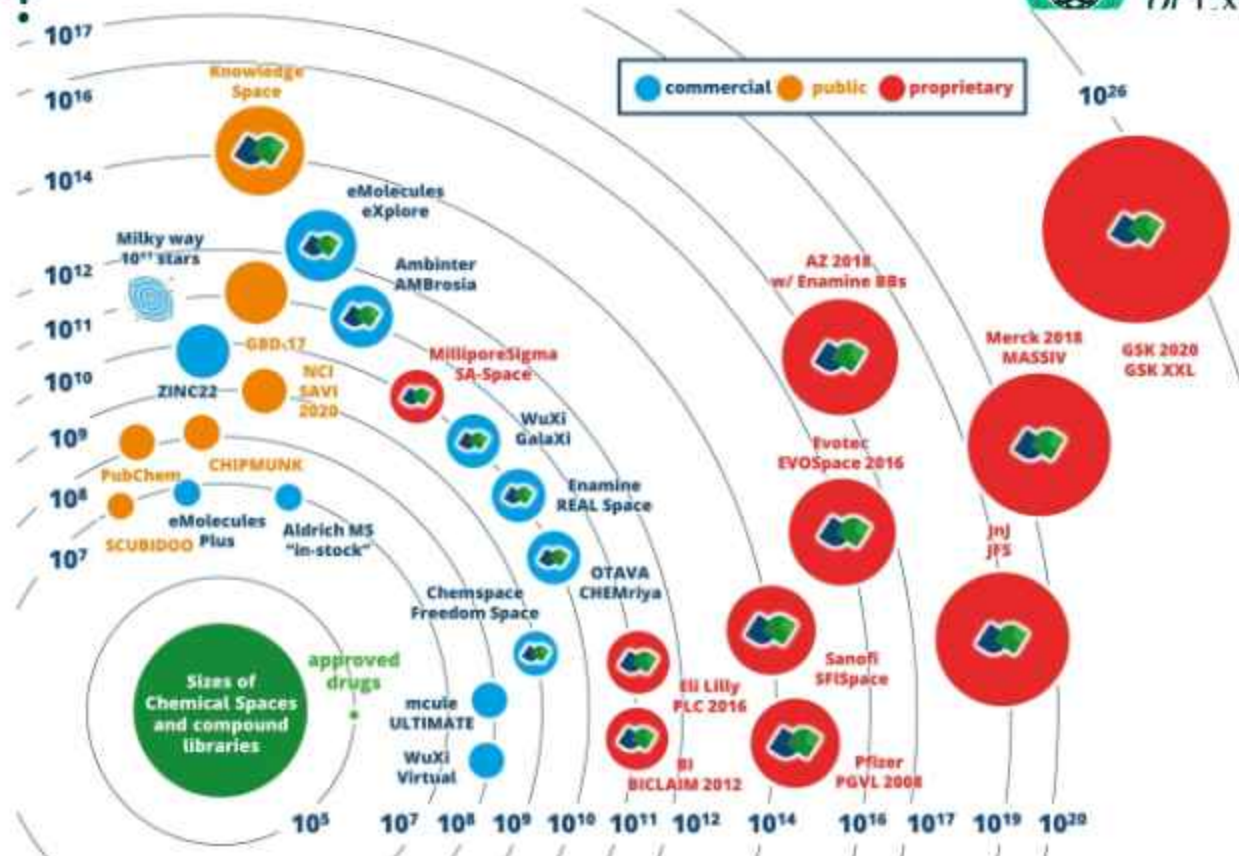
PPP = plant protection products; FRs = Flame Retardants; PCPs = Personal Care Products

Chemical numbers



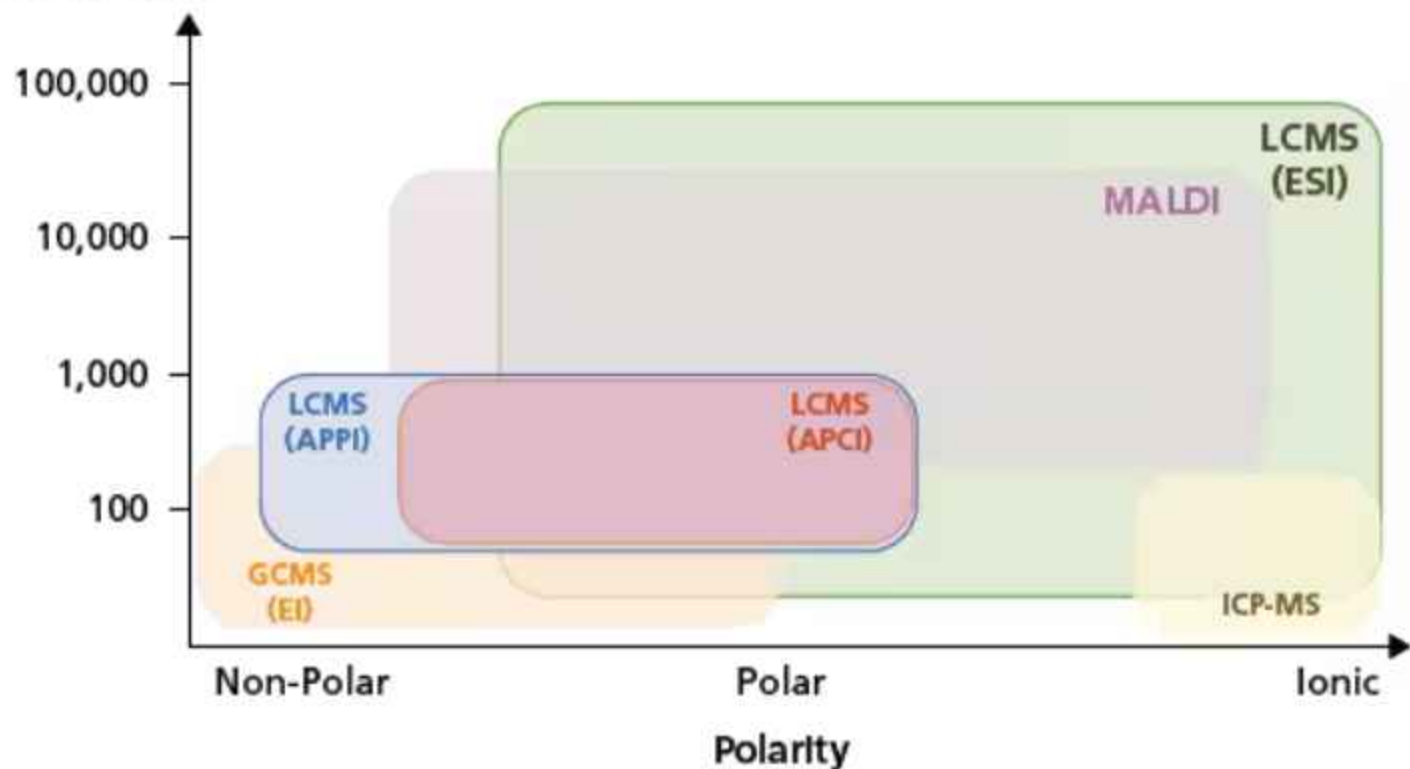
PPP = plant protection products; FRs = Flame Retardants; PCPs = Personal Care Products

How many compounds are we looking at?

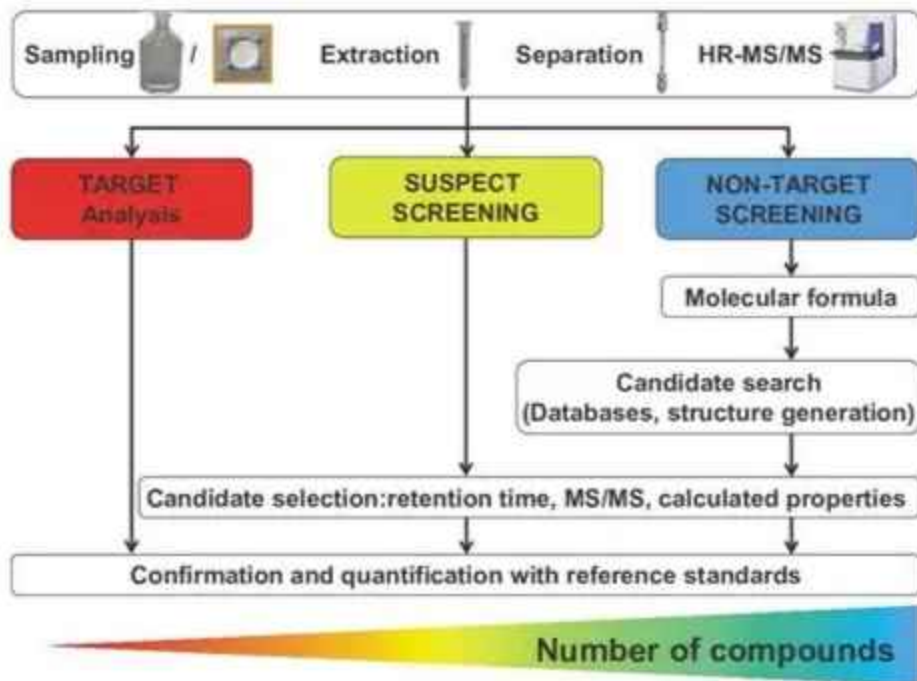


Chemical Space to be captured

Molecular Mass



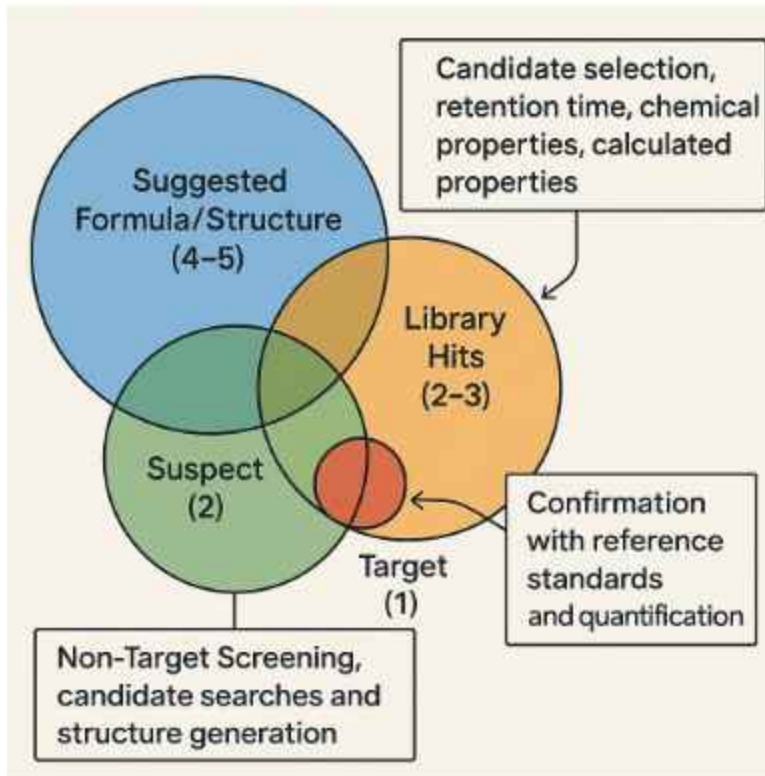
Workflow of Chemistry



Data Processing

Data from full-scan non target is highly complex providing:

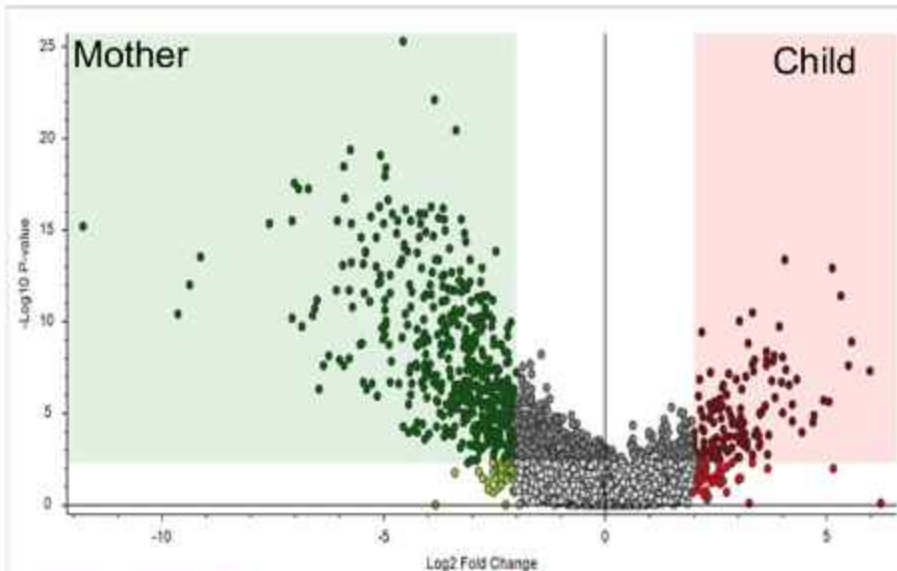
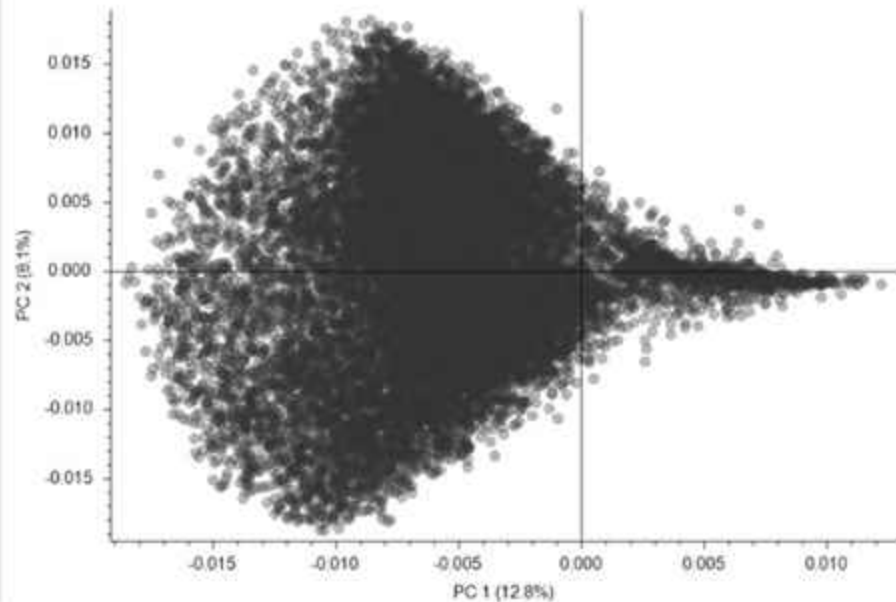
- Target data with high confidence
- Suspect data of intermediate confidence
- Formula and structures that direct towards chemical features such as:
 - Chlorinated (probably carcinogenic)
 - Brominated (probably bioaccumulates)
 - C-F bonds (PFAS compounds)



Confidence level of compound identification from 1 (high) to 5 (low).

One sample set

70'000 Features



3'000 Potential Fluorinated

Chemistry



University
of Exeter

soluble corrosive flammable
safety insoluble
biodegradable
persistent degradable volatile
inorganic metal inert
neutral harmful limit
bioaccumulative toxic non effluent
precaution risk standard threshold
hazardous carcinogenic
reactive disruptor residue regulation
heavy bioavailable mutagenic
monitoring emission herbicide assessment fungicide
contaminant pollutant pesticide
endocrine

How Environmental Intelligence Can Enhance Analytical Chemistry for Viable Outcomes



DEVELOP
NEW
WORKFLOWS



HANDLE
DATA-DENSE
SPECTRA



IDENTIFY ALL
CHEMICAL
FEATURES &
REMOVE
NOISE



VALIDATE
USING
MULTIPLE
METHODS



MERGE
SEARCHES
ACROSS
TOOLKITS



PRIORITIZE
CHEMICALS



PROVIDE
SUCCINCT
FOCUS FOR
REGULATORS



COMBINE
WITH
PREDICTIVE
TOXICITY
ALGORITHMS

Smart agents for water resource planning

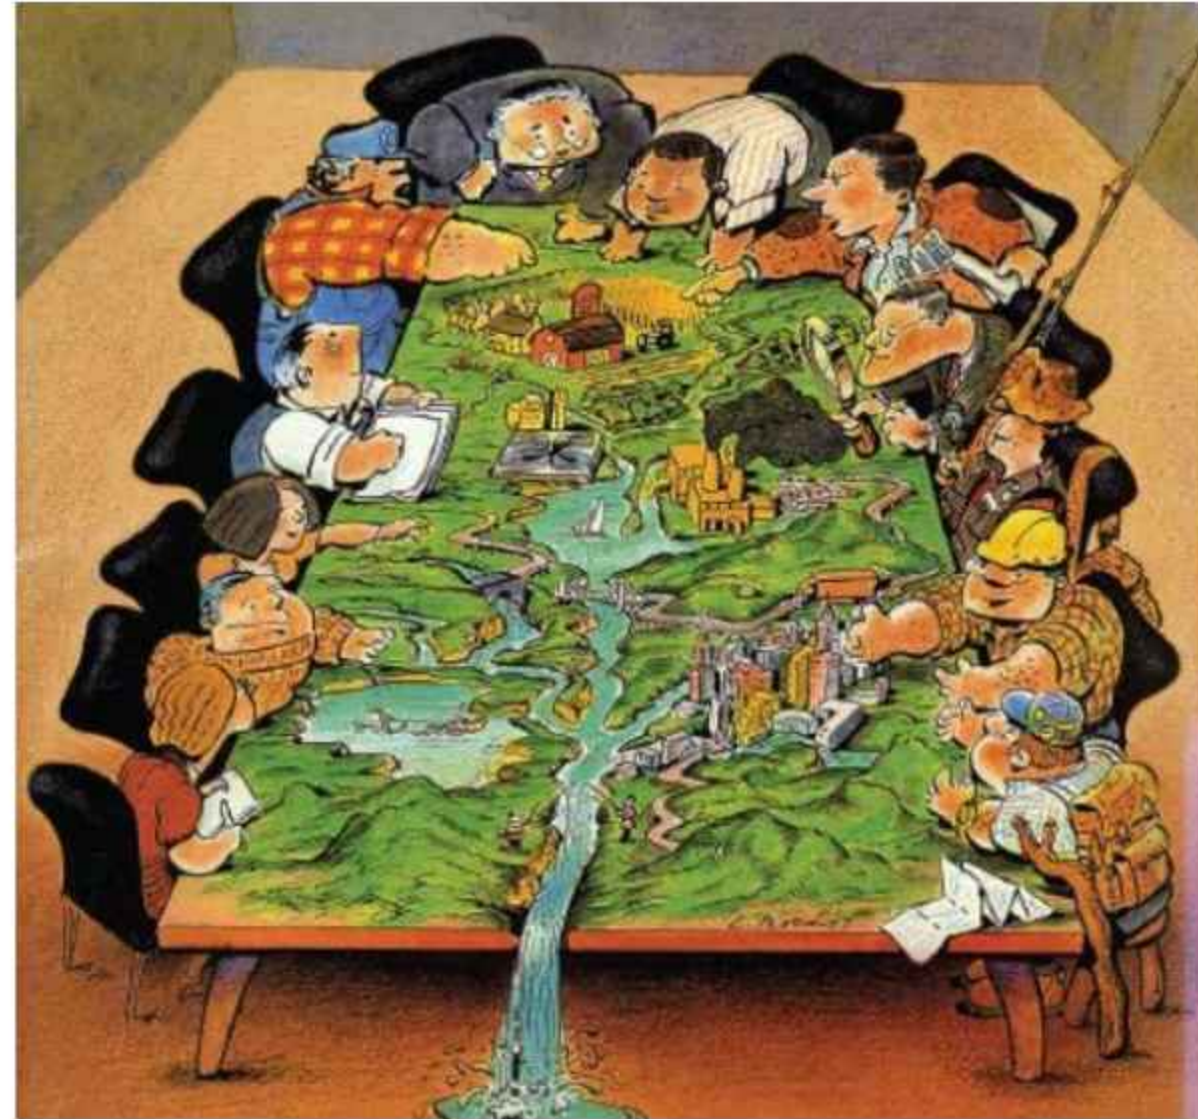
Environmental Intelligence Conference
September 2025

Mikiyas Etichia, Plešek Zdeněk, Hywel T. P. Williams



Why simulate water resource systems?

- ❑ Water resource systems involve multiple stakeholders and competing uses—such as drinking, irrigation, industry, and energy—where interventions often lead to trade-offs.
- ❑ Climate change, population growth, and demand-supply uncertainty further intensify the pressure on limited resources.
- ❑ Water resource simulation models help address these challenges by providing approximate representations to support efficient management of freshwater resources.

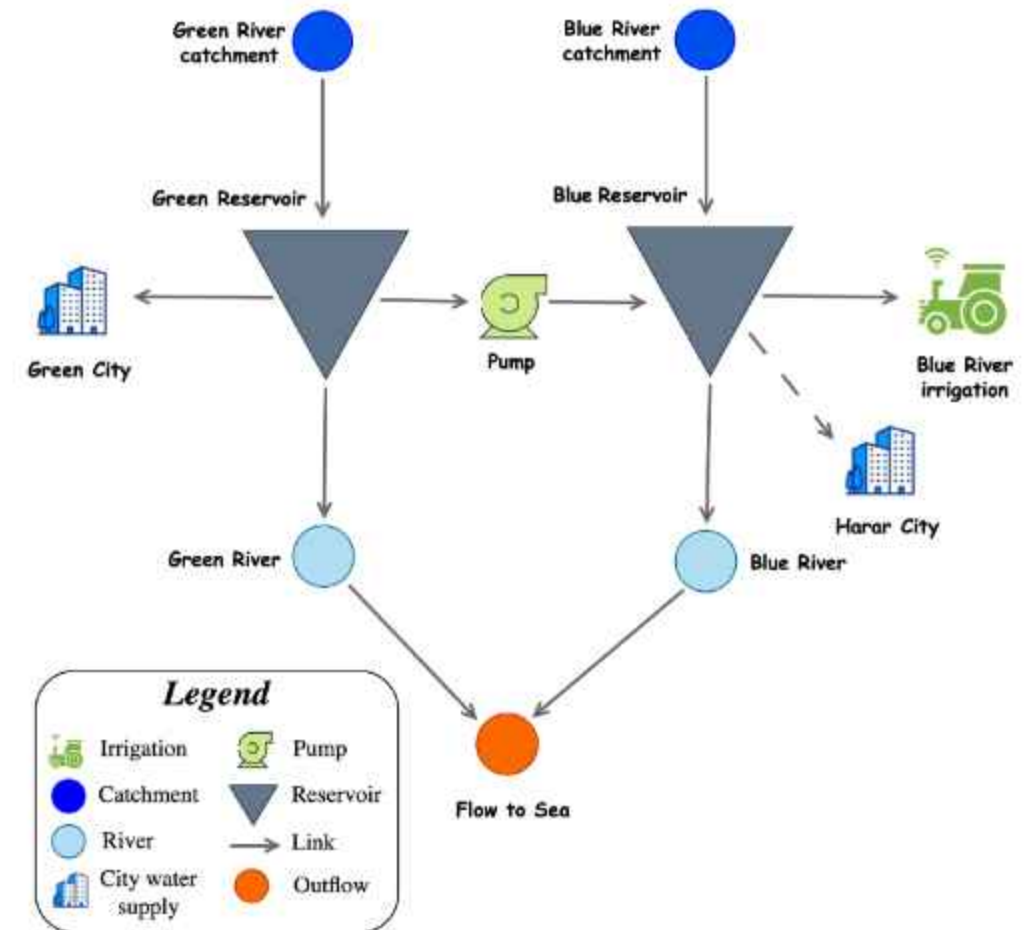


Loucks, D.P., van Beek, E., 2017

- ❑ Traditional water resource system models are expert-driven, and reliant on technical intermediaries and limited accessibility for non-experts.
- ❑ We use LLMs to create natural language interfaces to water simulations, reducing barriers to access and enabling participatory water resource planning.

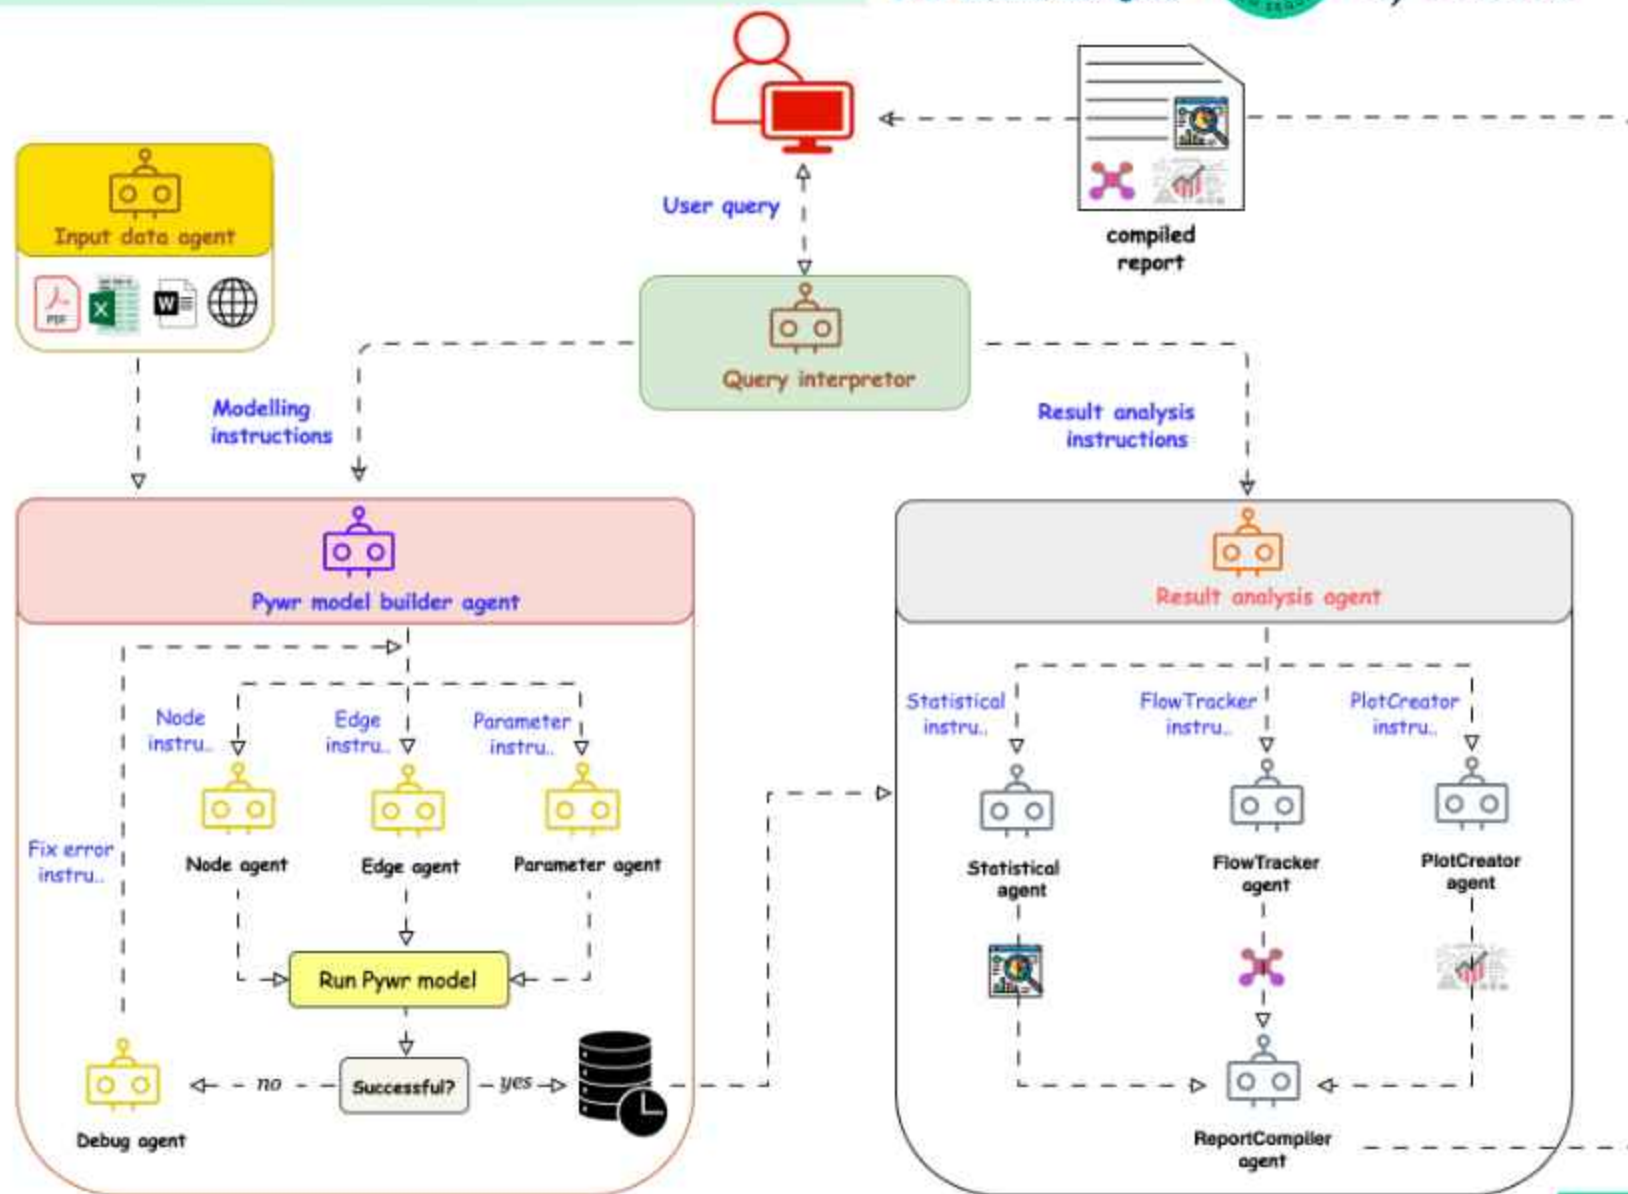


- ❑ Pywr is a dynamic modelling library for network-based resource allocation that enables advanced river basin simulation.
- ❑ Water resource system models are represented as a network of nodes and links.
- ❑ Water allocation is simulated at each time step (daily, weekly, or monthly), with decisions on storage, release, and use optimized using linear programming.
- ❑ Pywr models have three main components:
 - **Nodes:** Represent elements like rivers, dams, irrigation demands, gauges, etc.
 - **Edges:** Connect nodes.
 - **Parameters:** Supply data like river flows, dam capacities, demands, etc.



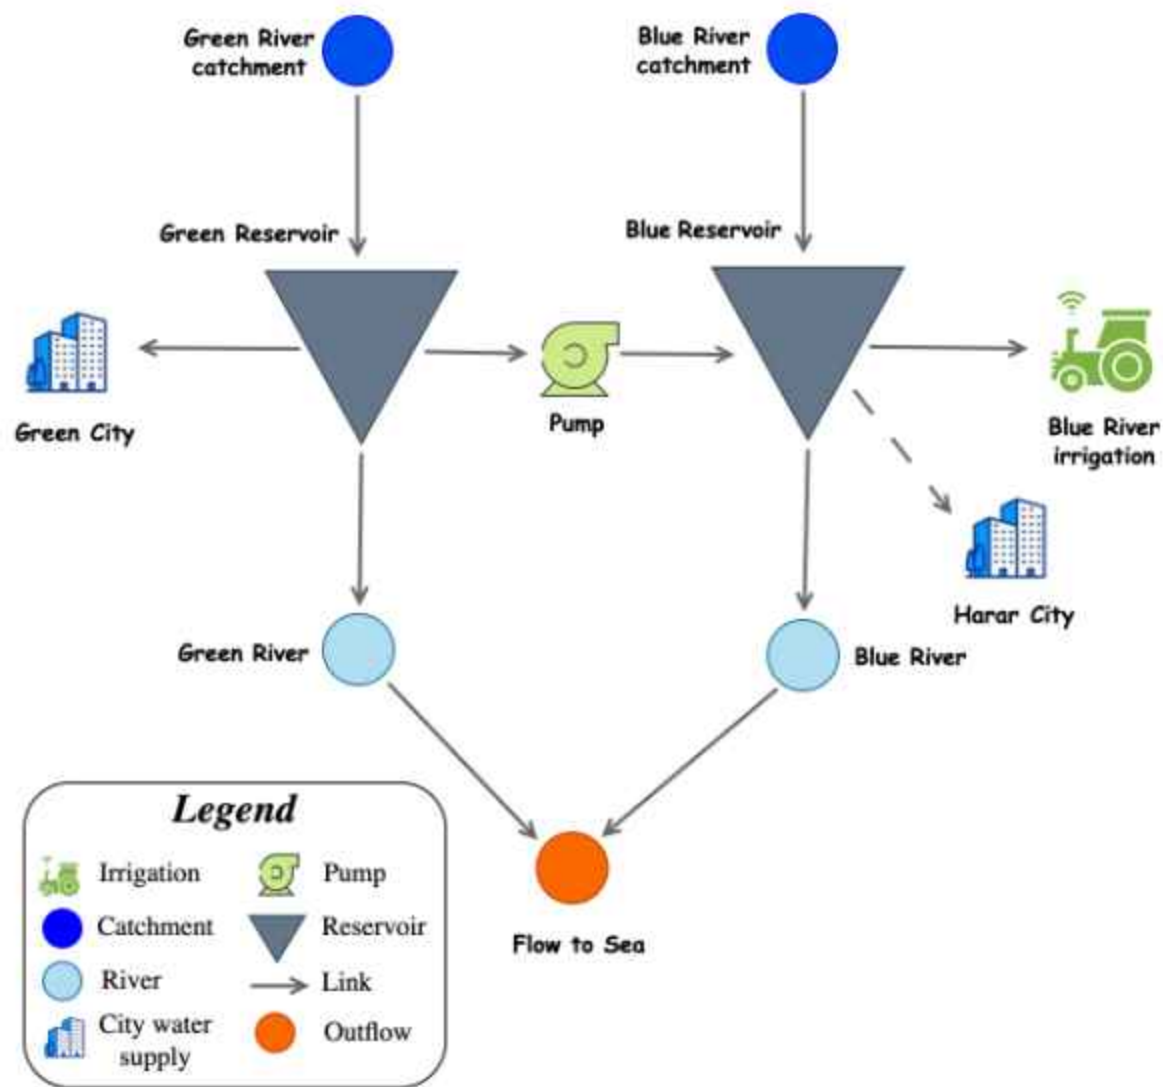
Schematic diagram of a hypothetical Twin River Basin

- ❑ Built with LangGraph for structured agent workflows.
- ❑ Maintains state and memory to enhance context and interaction.
- ❑ Adapts responses based on the user's background and expertise.
- ❑ Powered by GPT4-4o.



A workflow for water resource model building and analysis

- ❑ The synthetic Twin River Basin consists of distinct Green and Blue River catchments, each dammed, supplying water for Blue River irrigation and City water supplies, ultimately flowing to the sea.
- ❑ There is a plan to divert water for a new city called Harar from Blue Reservoir.
- ❑ Farmers are concerned that the planned diversion for Harar city may affect their water supply.



Schematic diagram of Twin River Basin

LLM-agent output



LLM-agent
output



User profile: A farmer in the Blue River Irrigation command area, concerned about the potential impact of new water demand on their supply.

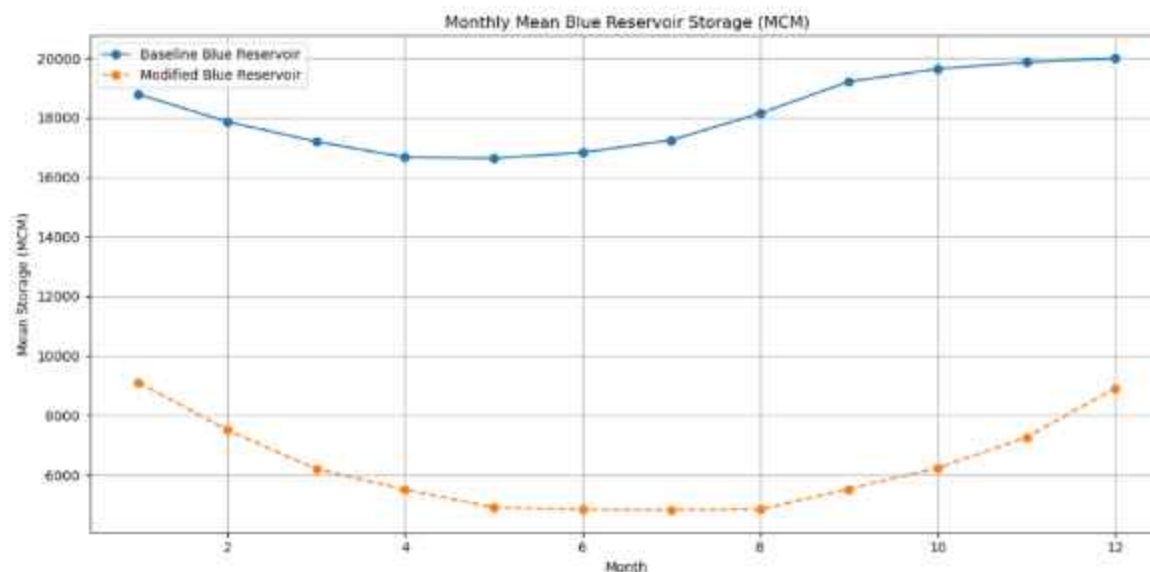
User query: What is the impact of adding Harar city's demand abstracted from the Blue Reservoir on Blue River Irrigation?

Report on the impact of adding Harar on water allocation in the Twin Rivers Basin

We've put Harar City's new water needs into our Twin Rivers Basin water system in the model, and it shows us some big changes in how water is divided up and how much stays in Blue reservoirs. While your water supply is not currently affected, the added demand increases the risk of future supply shortfalls.

Here's what our analysis shows:

- **Harar City's water supply:** The model shows that Harar City is consistently getting all the water they need from the Blue Reservoir. They're receiving a steady flow of about **20 million cubic meters of water every single day**. This means their new water needs are being met reliably, even when there isn't much rain coming into the reservoir.
- **Impact on Blue Reservoir's storage:** Now, this is the part that directly affects you. Taking water for Harar City is significantly reducing how much water is stored in Blue Reservoir, especially during the dry months.
 - During these critical dry periods, the average amount of water in the Blue Reservoir is showing a drop of **more than 70%** compared to before Harar's demand.
 - To be more precise, from **May through August**, when water is usually scarce, the storage in the Blue Reservoir is reduced by a significant amount, ranging from about **half (51.5%) to nearly three-quarters (73.3%) less** than what was typically there.



The graph shows two lines tracking Blue Reservoir water levels annually. The **Blue Line** (solid) represents the average monthly water held *before* Harar City's demand. The **Orange Line** (dashed) shows average monthly water held *now*, with Harar City's new demand.

What to notice on the graph:

- **The Big Drop:** The orange line is significantly below the blue line, especially from **May through August** (months 5-8), indicating much less water during drier months.
- **Year-Round Impact:** The orange line is consistently lower, meaning the reservoir generally holds less water throughout the year due to the new demand.



University of Exeter

Centre for Water Systems



Establishing a living lab for water at the University of Exeter

Associate Professor Peter Melville-Shreeve
(on behalf of Chris Sweetapple)

C.Sweetapple@exeter.ac.uk



Setting the scene: A world leader in water management

- Twice ranked global number one in the Times Higher Education (THE) Impact Rankings for Clean Water and Sanitation (2023 & 2024)
- Metrics based on:
 - ✓ Research on clean water and sanitation
 - ✓ Water consumption tracking
 - ✓ Water usage and care
 - ✓ Landscape planting to minimise water usage
 - ✓ Water re-use
 - ✓ Water in the community



Towards a 'living lab' for water



119 ultra low-flush 'Propelair' toilets installed in 7 buildings on campus in 2018

- Monitoring implemented
- Captured evidence of impact on:
 - Water use (*2,287m³ annual reduction*)
 - Sewer network condition (*no evidence of increase in blockages*)
 - Water bills (*£12,580 annual reduction, at 2018 prices*)
 - Sewer discharges (*6m³/day reduction*)
- Ultra low-flush toilets becoming 'business as usual' for new installations

Melville-Shreeve, P. et al. (2021) 'Campus study of the impact of ultra-low flush toilets on sewerage networks and water usage', Water, 13(4), p. 419. doi:10.3390/w13040419.



University of Exeter

Centre for Water Systems



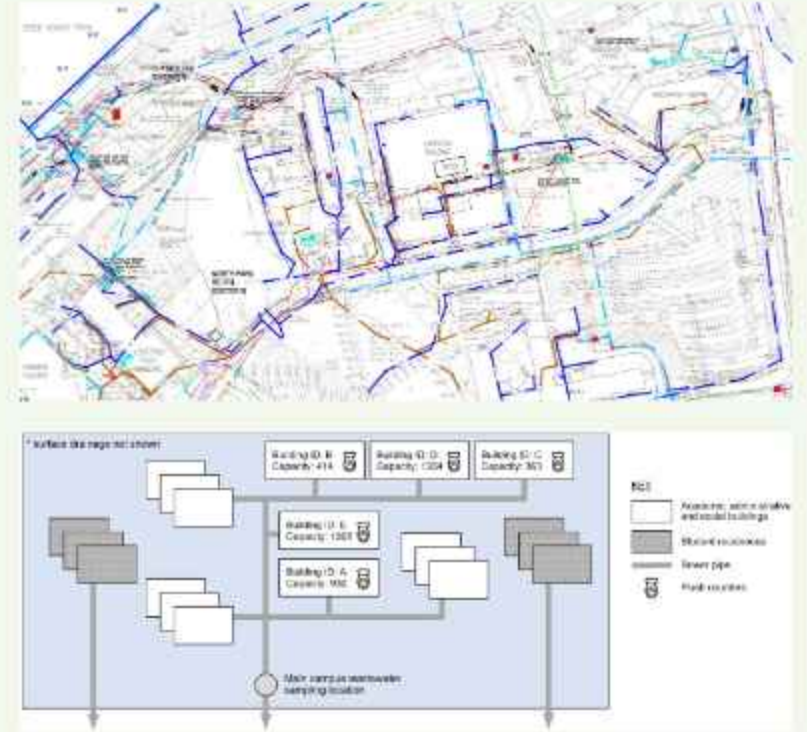
STORMWATER INTELLIGENCE GROUP

Towards a 'living lab' for water

Near-to-source wastewater surveillance pilot for SARS-CoV-2 detection

- Wastewater samples collected and analysed for SARS-CoV-2, ammoniacal nitrogen and orthophosphate.
- Wastewater flow rate measured at sampling locations
- In conjunction with flush count data, provided better understanding of impact of variable populations on trends provided by wastewater-based epidemiology.

Sweetapple, C. et al. (2022) 'Building knowledge of university campus population dynamics to enhance near-to-source sewage surveillance for SARS-COV-2 detection', Science of The Total Environment, 806, p. 150406. doi:10.1016/j.scitotenv.2021.150406.



University of Exeter

Centre for Water Systems

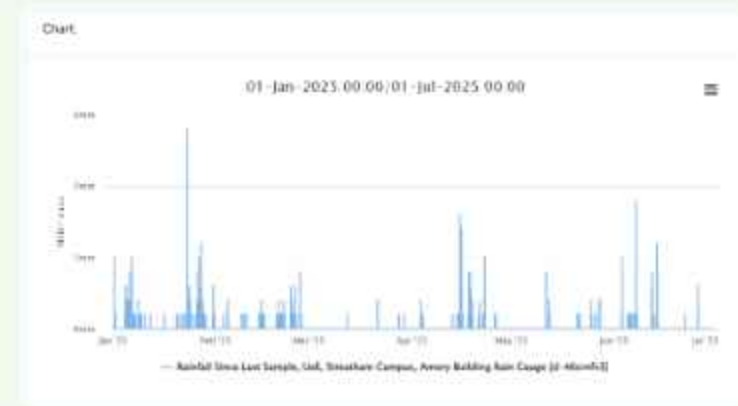
Towards a 'living lab' for water



Grey water use and water-use monitoring in the new net zero CREWW building



Rainfall monitoring



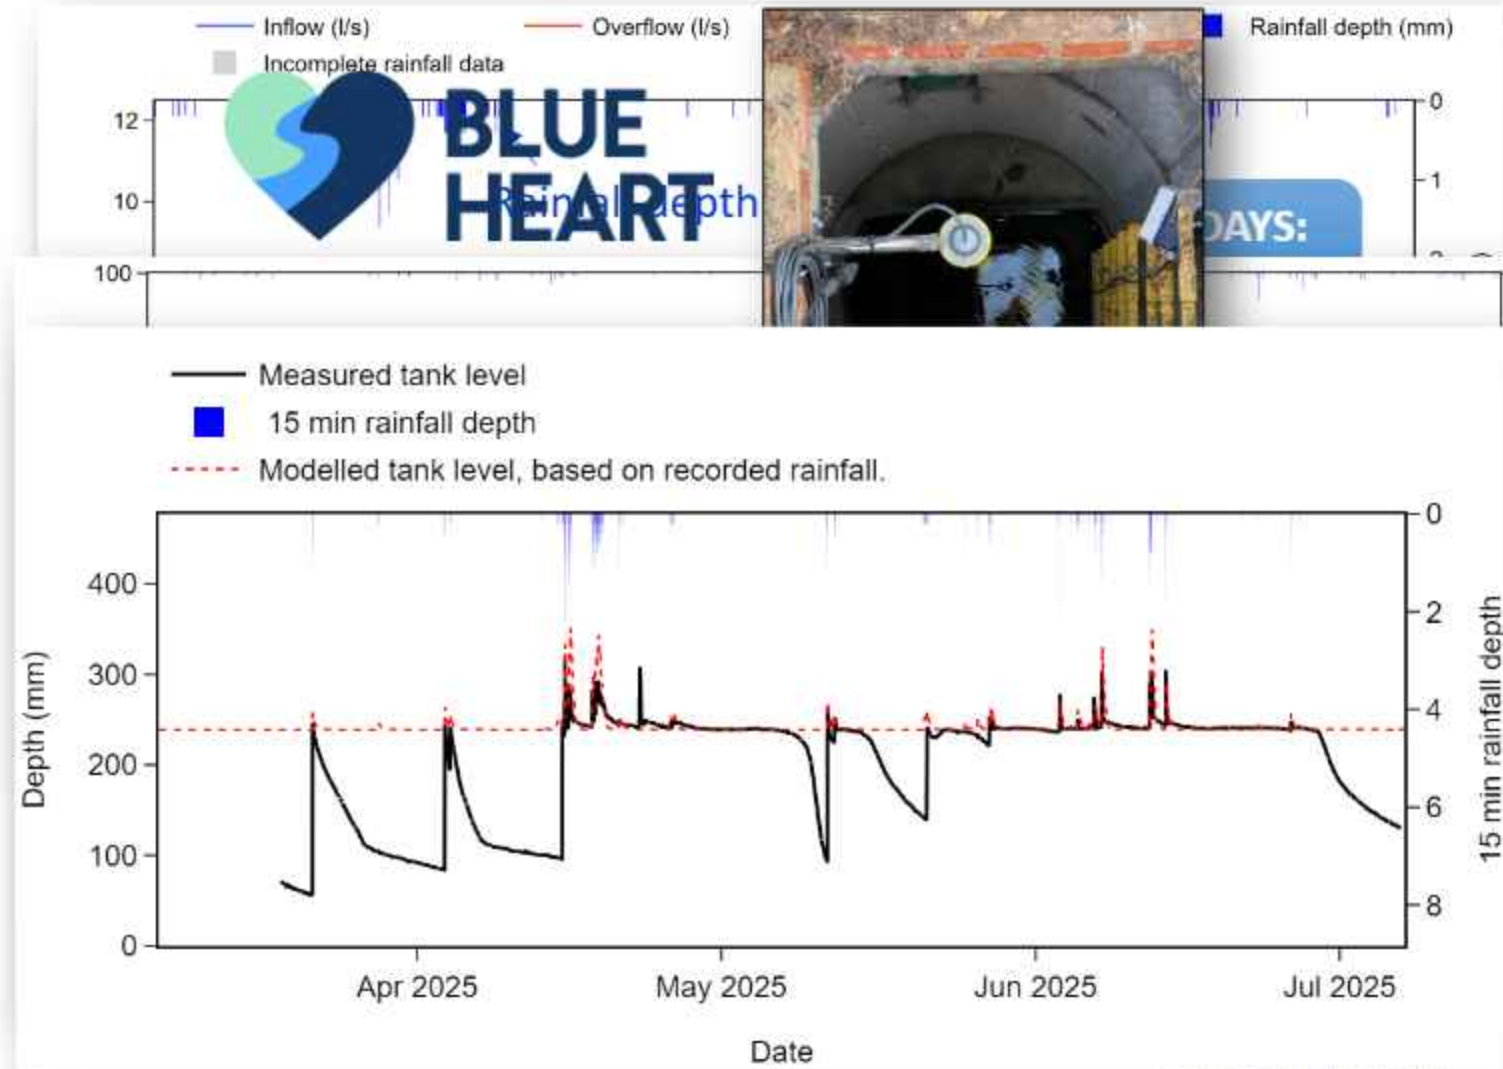
University of Exeter

Centre for Water Systems

Monitoring Stormwater Systems on Campus

Attenuation tank monitoring

- Case study site identified
- Theoretical performance modelled using campus rainfall data
 - Design criteria met, BUT not more than 38% full in 13 years
 - Potential for greater day-to-day benefit
 - Opportunities for 'smart' technologies, including real time control
- Level sensor installed to validate results



University of Exeter

Centre for Water Systems

Future plans?

- Trial new real time control technology to manage surface water runoff – including lab tests and full scale deployment
- Extended monitoring of attenuation assets, including levels and flows
 - Provide evidence of current (under) performance
 - Provide evidence of achievable performance with improved control.
 - Inform future funding bids



→ Development of a 'smart' campus
→ "Net Water Zero Proposal"



University of Exeter

Centre for Water Systems



STORMWATER INTELLIGENCE GROUP



University
of Exeter

Centre for
Water Systems



STORMWATER INTELLIGENCE GROUP



University
of Exeter



Dr. Jawad
Fayaz



Lecturer (Assistant Professor)
Dept of Computer Science,
University of Exeter (UoE), UK



Honorary Lecturer
Dept of Civil, Environmental &
Geomatic Engineering
University College London (UCL), UK

Holistic Machine Learning Frameworks for Proactive Water Distribution and Leakage Management

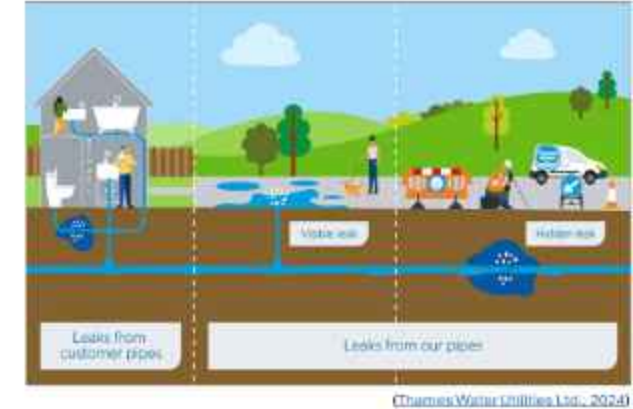
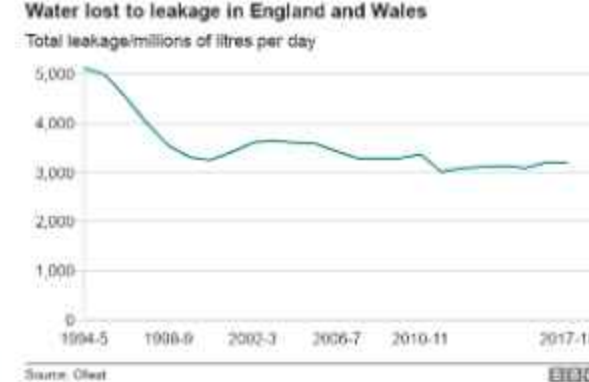
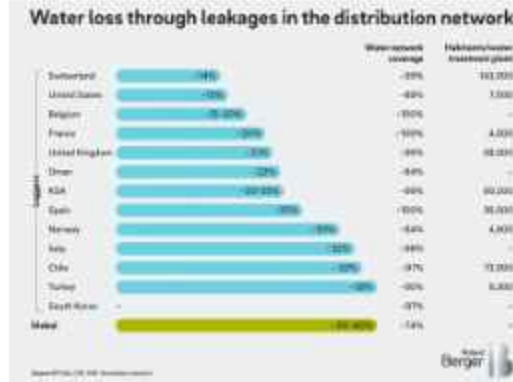
MANCHESTER
PRIZE



Department for
Science, Innovation
& Technology

Global Challenge

- Over 45 million cubic meters of water are lost globally every day
- Aging infrastructure and poor maintenance are leading causes of water loss
- Reducing leakage can help achieve sustainable water management and address water scarcity challenges

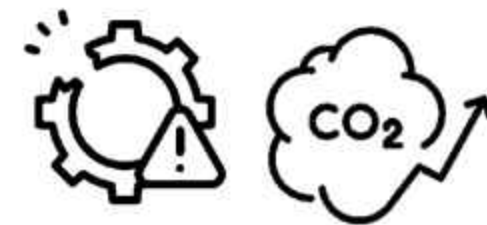


53L per person per day are lost in England and Wales (~20% of supply)

70% of leaks are not seen by customers and often are deep underground

£60M is estimated to be the annual cost of water leaks to UK water companies

30% of leaks are reported by the customers when they reach surface level (water pooling)



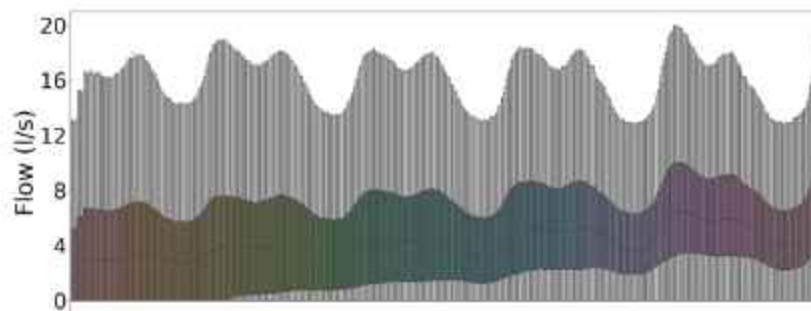
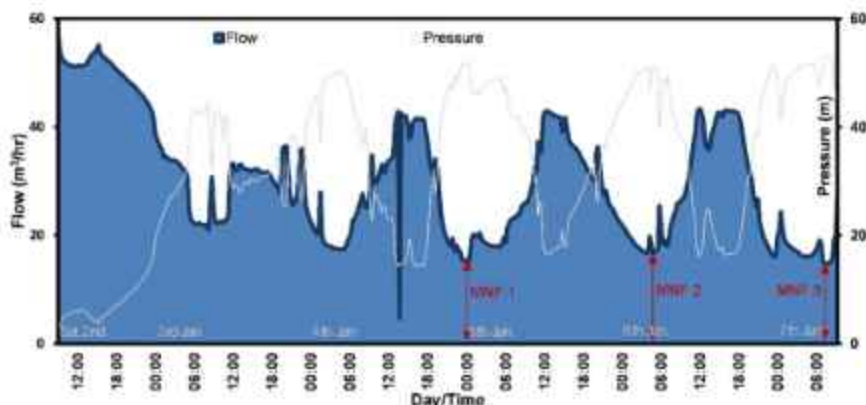
CIWEM Chartered Institution of Water and Environmental Management

Plan to:

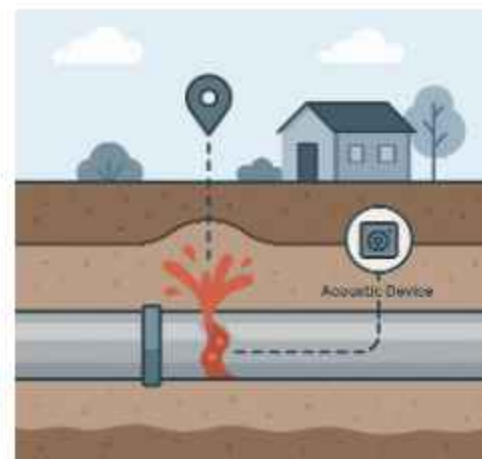
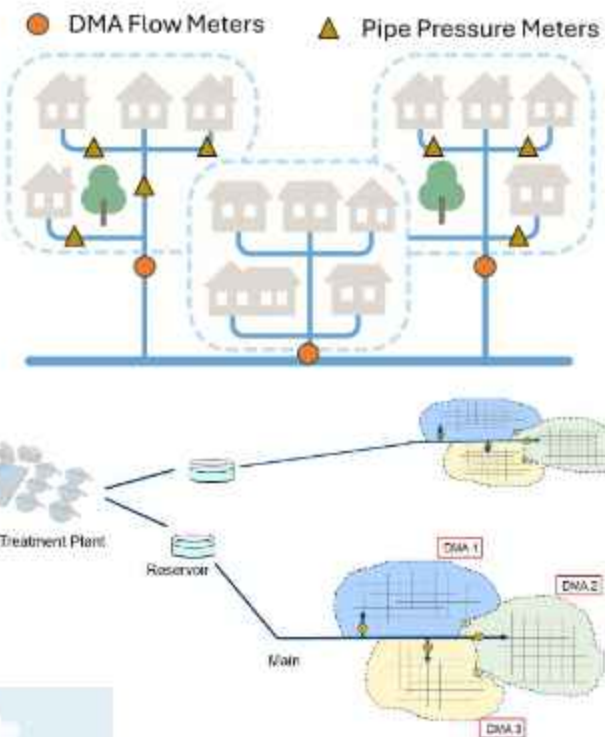
- 1) Roll out proactive leakage control strategies
- 2) Reduce 50% of leakages by 2041
- 3) £400 million invested for water savings

Current State of Practice (Only Reactive)

- **DMAs:** subdivisions of water networks (1000–2000 properties each; 22,000+ in UK)
- **Flow meter:** installed at each DMA inlet to measure inflow
- **Pressure sensors:** 5–6 per DMA to track pressure variations



- Current practice: **Minimum Night Flow (MNF)**
- MNF: lowest flow between 2 to 4 am, averaged over **7 to 14 days**
- Leakage indicator: MNF rises above set thresholds
- Localization relies on **acoustic devices** to pinpoint leaks
- Requires manual deployment throughout the DMA
- Labor-intensive and can miss small or intermittent leaks

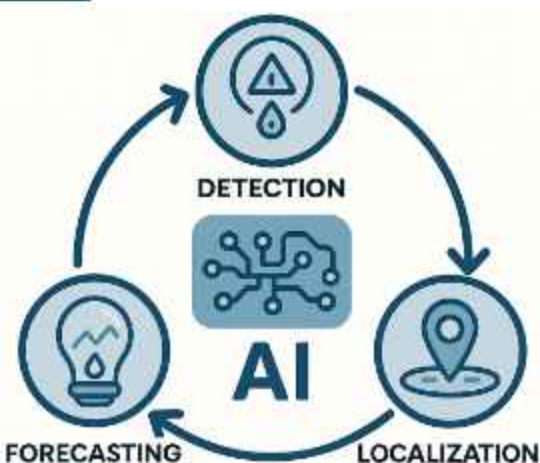


Towards Smart Water Management

- **Regulation:** Ofwat requires **15% leakage reduction by 2030**
- **Plan:** expand flow & pressure sensor networks to meet target
- **Challenge:** Fixing entire infrastructure = costly, slow, complex

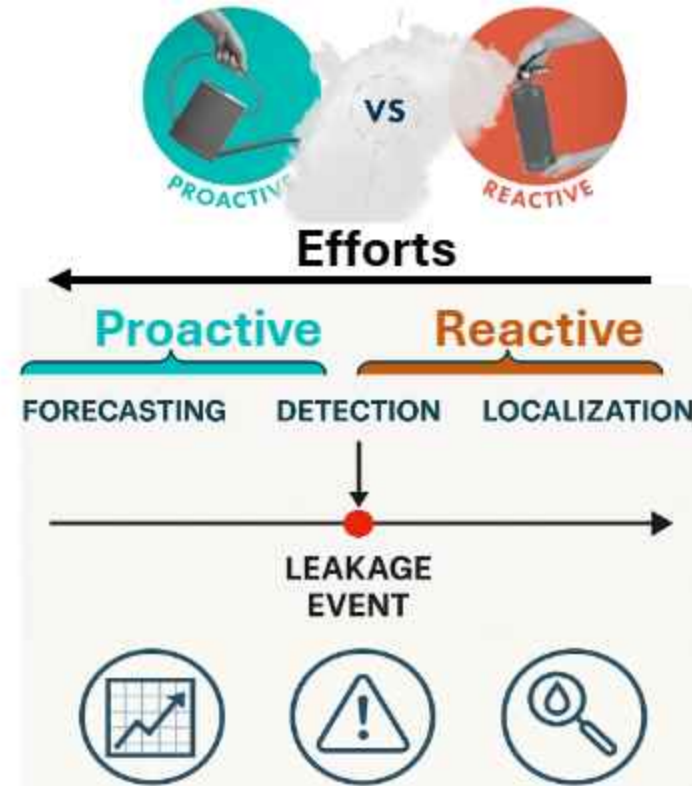
Step 1:

- ❑ Use existing networks + add low-cost sensors
- ❑ **AI integration:** smarter, faster leakage management without full overhaul
- ❑ **Forecasting:** predict leakage risk from historical flow, pressure, environment
- ❑ **Detection:** real-time anomaly monitoring to catch leaks early
- ❑ **Localization:** pinpoint leaks (pipe/node level) for targeted repairs



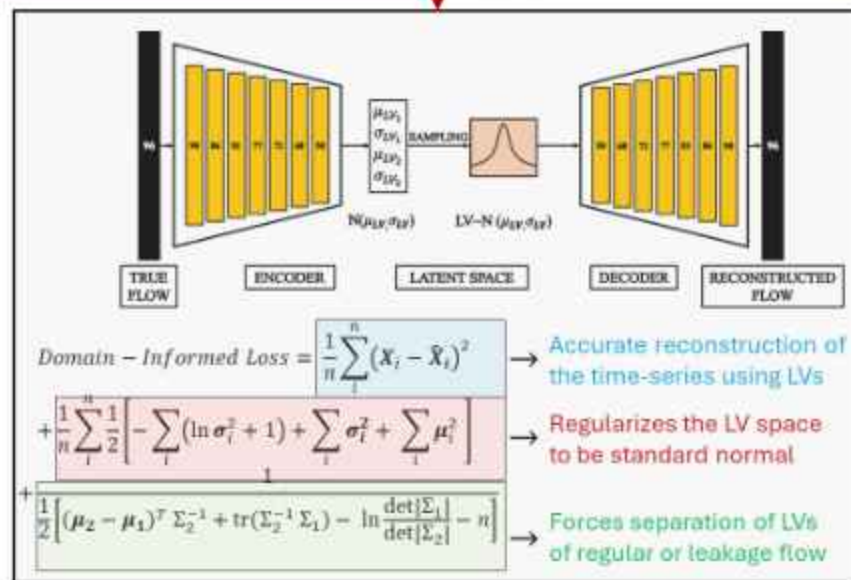
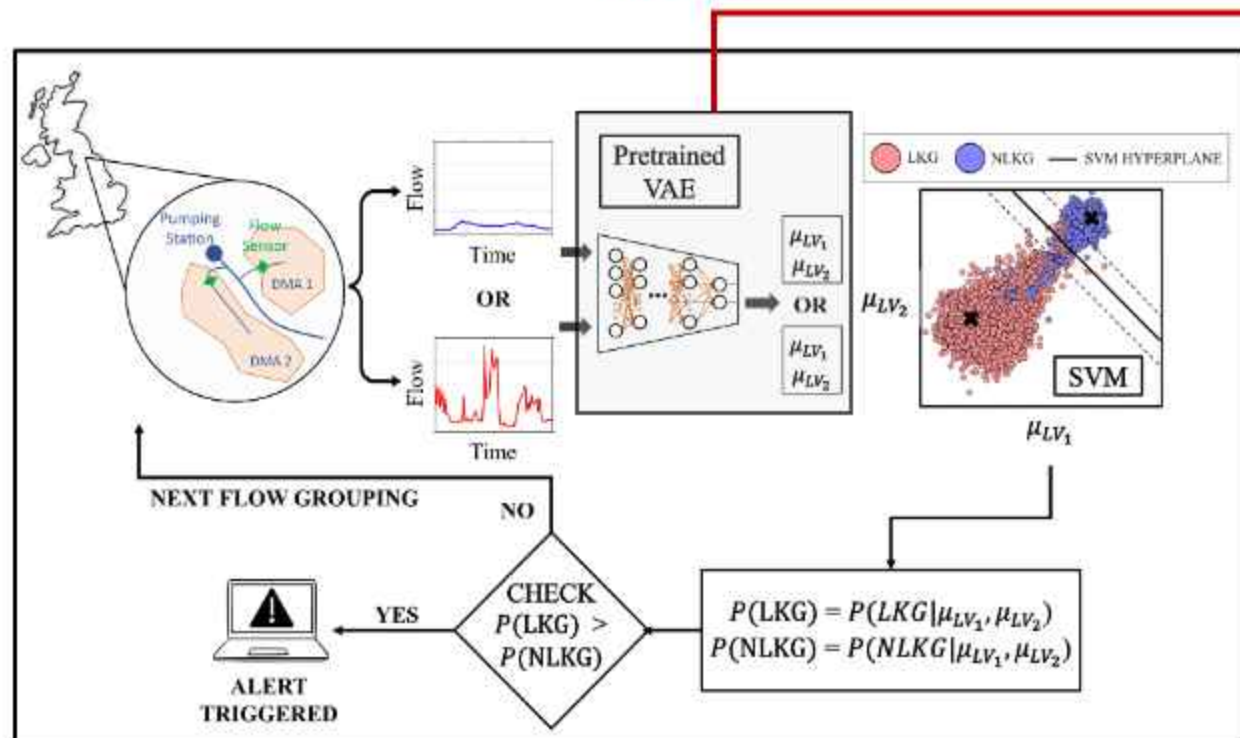
We propose real-time AI deep learning-based frameworks for:

- ✓ **Detection:** Domain-informed VAE + SVM → distinguish between usual flow and leakage
- ✓ **Localization:** GAT + Transformer → estimate pressures, pinpoint leak location for rapid repairs
- ✓ **Forecasting:** Hybrid class-based attention + Bayesian updating → predict future flow & leak risk
- ✓ **Decision-making:** Deep RL agent → prioritize maintenance & inspections to reduce failures

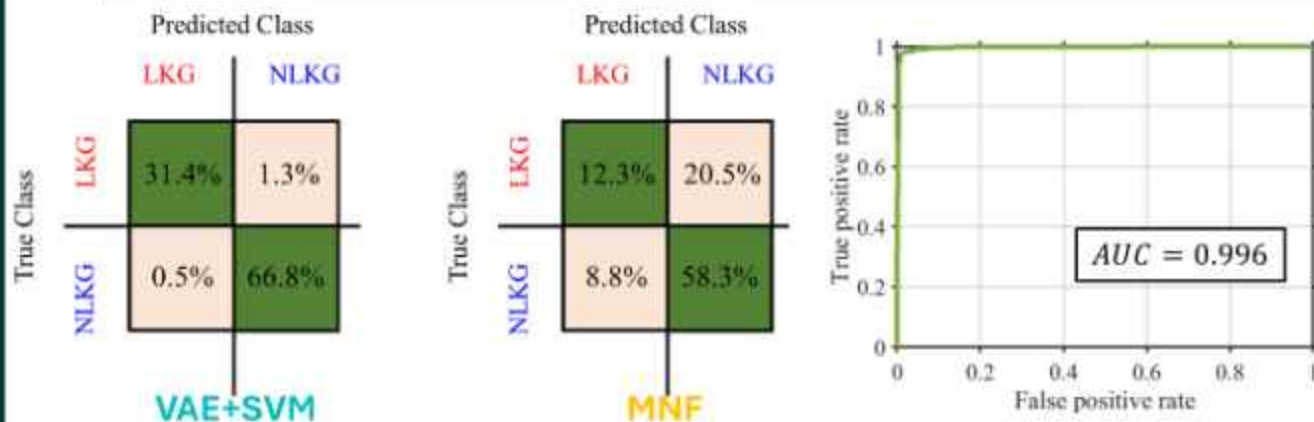


Reactive Leakage Detection

- Domain-informed variational autoencoder and support-vector-machines based
- Uses flow recordings from **previous 1 day** to detect probability of leakage



VAE compresses the signal into a regularized 2D latent space preserving most relevant structure and variance which is then separated easily by SVM hyperplane

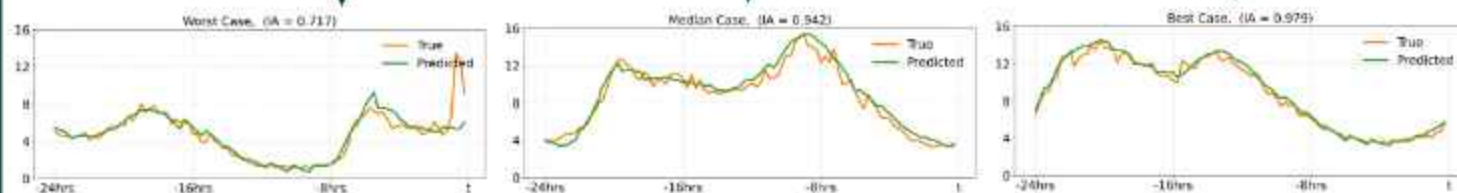
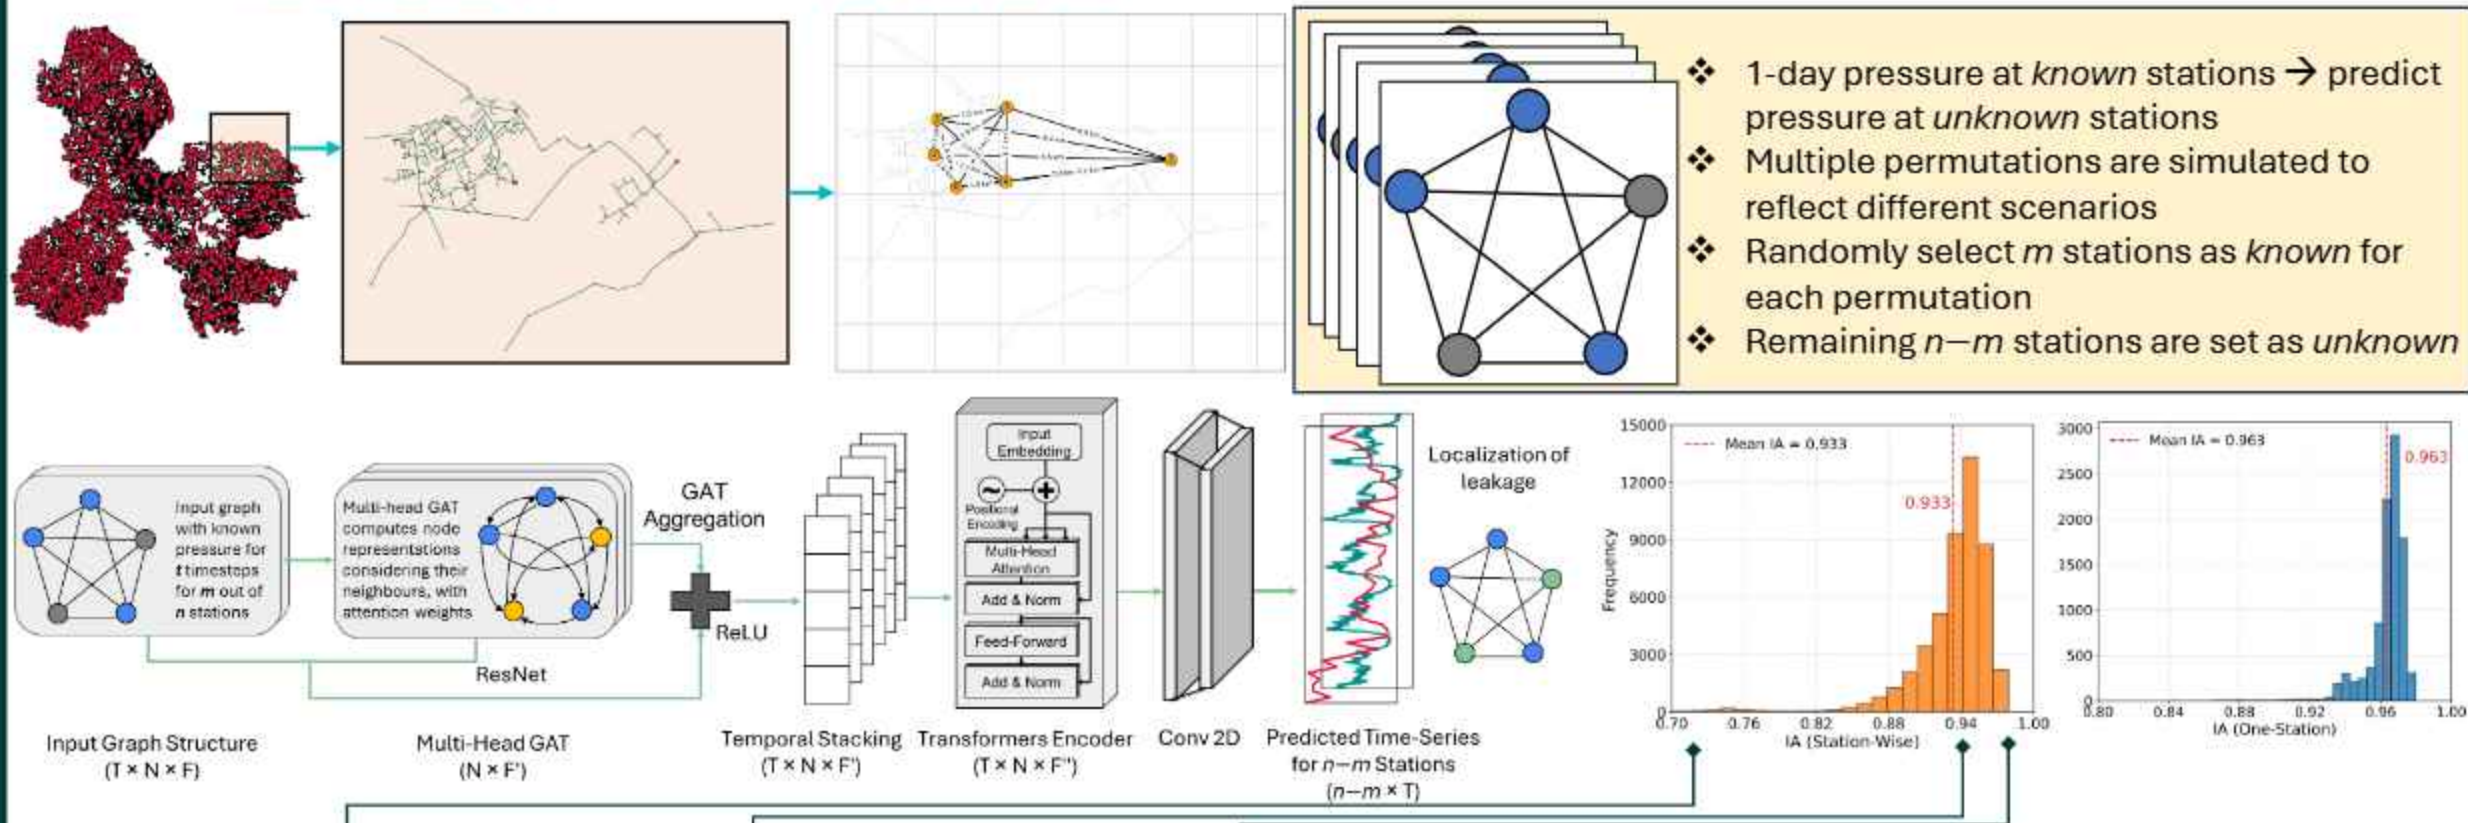
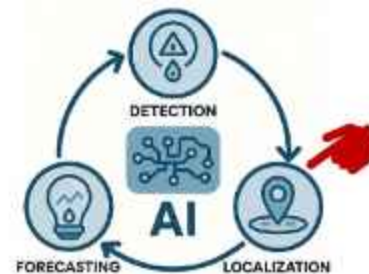


Strong detection rates:

- ❖ Accuracy: **> 98%** (70.6%)
- ❖ AUC **> 0.99** (0.73)
- ❖ True Positive Rate (LKG): **98.4%** (58.3%)
- ❖ True Negative Rate (NLKG): **98.1%** (74%)

Spatial Prediction (Localization)

- **Graph model:** nodes = sensor junctions, edges = pipes
- **Leak localization:** GNN-based spatial regression on pressure anomalies
- Uses **node features** (pressure) + **edge features** (pipe diameter, distances) + topology

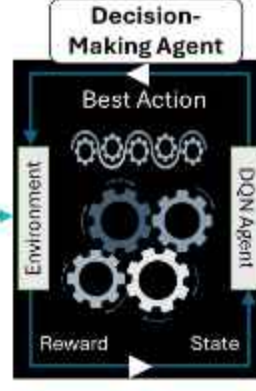
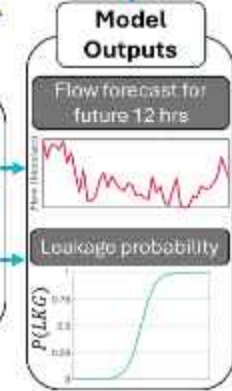
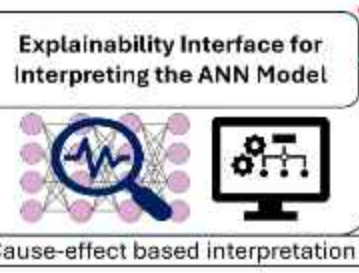
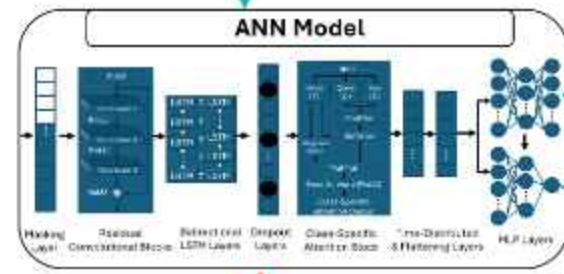
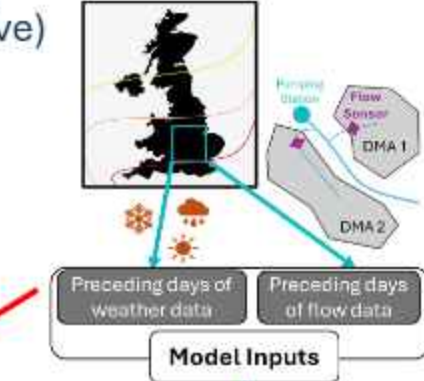
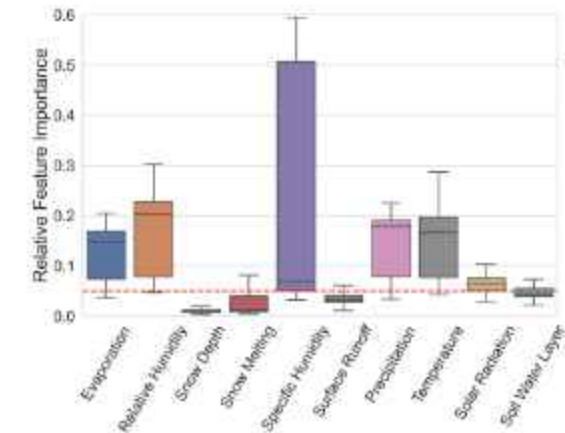


- ✓ Average index of agreement (IA) — **0.963** for one-station and **0.933** station-wise
- ✓ **Leak localization:** inferred by combining large residuals at sensor nodes with predicted low-pressure zones

Forecasting Framework

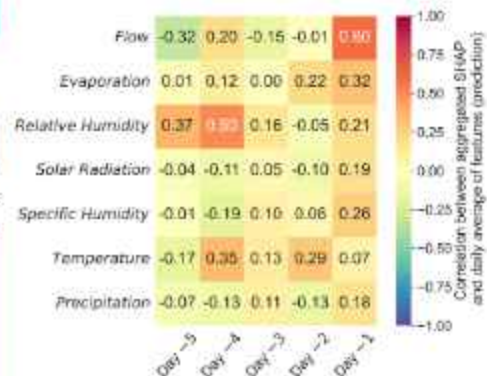
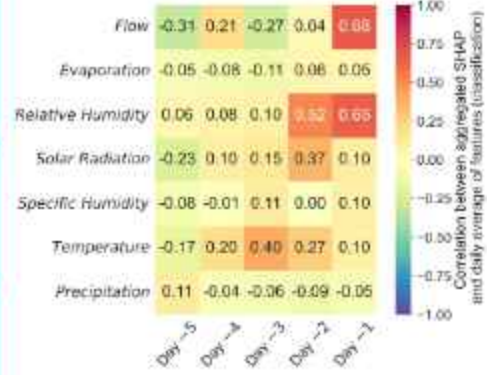
Analysed for **feature importance** to predict water flow and leakages

- ✓ Humidity (specific and relative)
- ✓ Precipitation
- ✓ Temperature
- ✓ Solar radiation
- ✓ Evaporation



Feature influence: quantified using Shapley values across 12-hour forecasting window (48 points)

$$\Phi_i(\text{day}) = \sum_{k=1}^K \sum_{t=1}^T \sum_{j=1}^N |\phi_{i,j,t,k}|$$



$$\text{Attention}^c(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d_k}} + \mathbf{W}_c\right) \mathbf{V}$$

$$L_{\text{adaptive}}(w_i, \varepsilon_i, \delta) = \frac{\sum_{i=1}^N w_i \cdot L_{\text{huber}}(\varepsilon_i, \delta)}{\sum_{i=1}^N w_i}$$

$$L_{\text{binary}} = -\frac{1}{N} \sum_{i=1}^N [v_i \cdot y_i \cdot \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

\mathbf{W}_c : class-specific weight matrix that increases attention scores for features related to LKG

w_i : LKG is assigned a higher weight, making its errors more impactful on model updates

v_i : weight for each instance, ensuring higher penalties for misclassifying leakage cases

SHAP correlations: show direction & strength of feature influence relative to daily averages

□ This helps to **counteract the masking bias**, which otherwise may make NLKG cases dominate learning

Bayesian Updating

➤ Temporal scale for 2-weeks future prediction

Prior → forecasts 12 hrs after 2 weeks in future

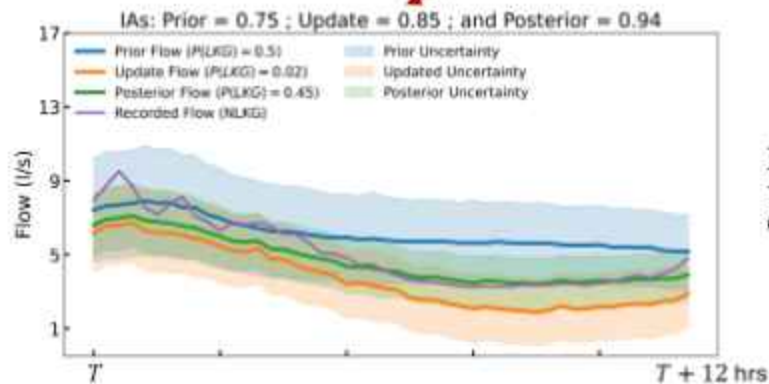
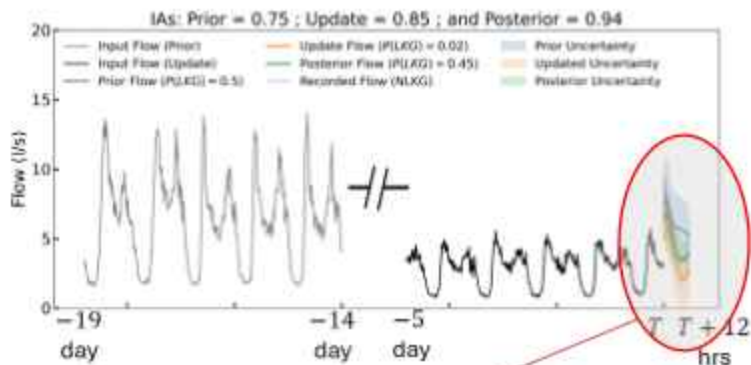
$$\{y_{t:t+12\text{hrs}}, P(LKG)\} = \text{ANN}(y_{t-19\text{ days}:t-14\text{ days}}, \mathbf{W}_{t-19\text{ days}:t-14\text{ days}}, \text{Pipe Features})$$

Update → forecasts 12 hrs in future

$$\{y_{t:t+12\text{hrs}}, P(LKG)\} = \text{ANN}(y_{t-5\text{ days}:t}, \mathbf{W}_{t-5\text{ days}:t}, \text{Pipe Features})$$

Posterior ~ MCMC(*Update* × *Prior*)

➤ Predictive Performance



Balanced Test Set

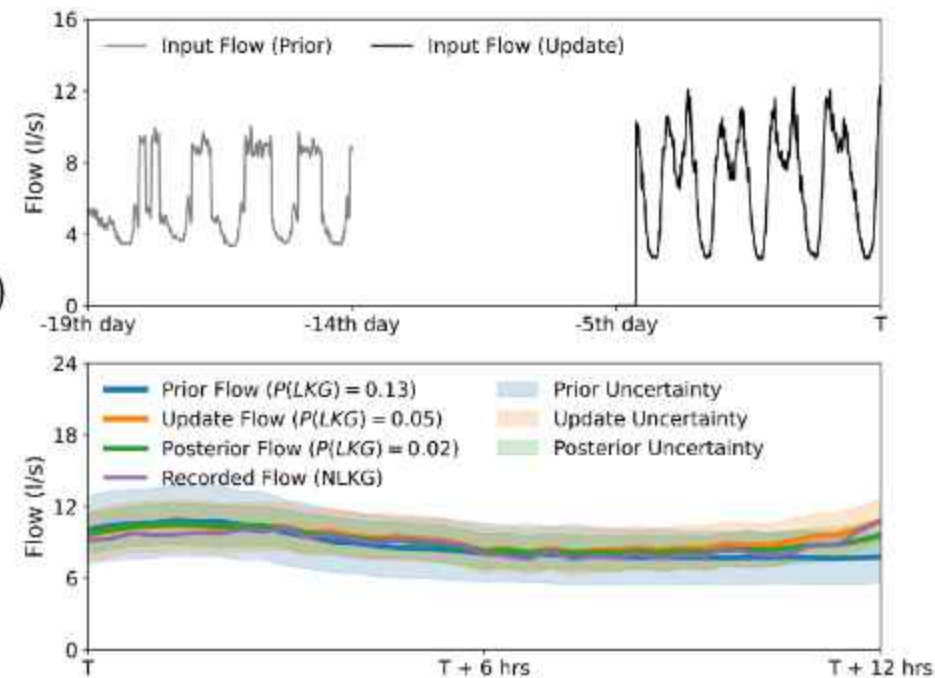
True Label	NLKG	2761 (98.89%)	31 (1.11%)
	LKG	327 (11.95%)	2409 (88.05%)
		NLKG	LKG
		Predicted Label	

50% (LKGs)
50% (NLKGs)

Imbalanced Test Set

True Label	NLKG	645 (99.33%)	5 (0.77%)
	LKG	3 (0.02%)	31 (91.38%)
		NLKG	LKG
		Predicted Label	

5% (LKGs)
95% (NLKGs)

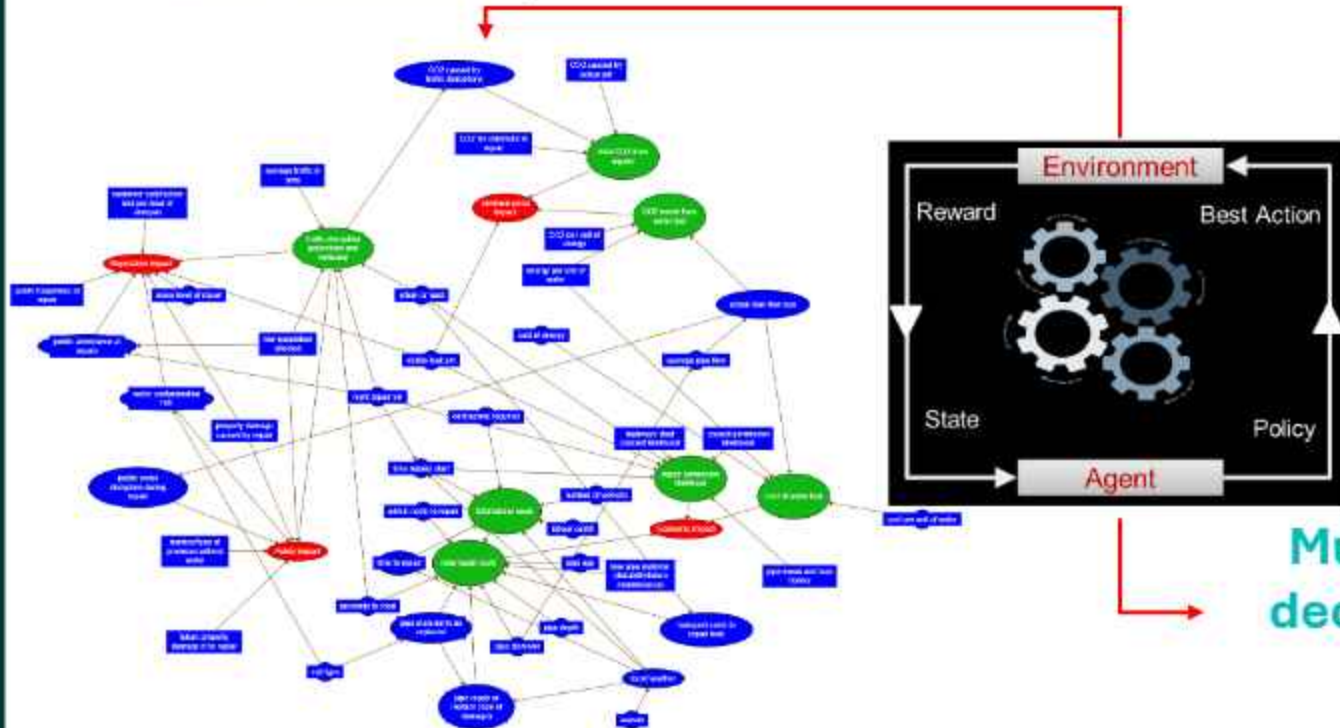


- ✓ We can predict 15-minute water flow demands for the next 12 hrs with **index of agreement IA > 80%**
- ✓ We can predict whether there is going to be a leakage in the next 12 hours with **accuracy > 90%**

- Fayaz J., McMillan L., Cardenas V., and Varga L. (under review). "Real-Time Weather-Adaptive Water Flow and Leakage Forecasting using an Interpretable and Unified Deep Neural Network". *Applied Soft Computing*.
- McMillan L., Fayaz J., and Varga L., (2023). "Flow Forecasting for Leakage Burst Prediction in Water Distribution Systems using Long Short-Term Memory Neural Networks and Kalman Filtering". *Sustainable Cities and Society*.
- Rapp M., and Fayaz J., (under review). "Uncertainty-Aware Decision Support for Demand Forecasting and Leakage Prediction in Water Distribution Systems Using Unified Temporal Fusion Transformers". *Reliability Engineering and System Safety*.

Decision-Making Agent

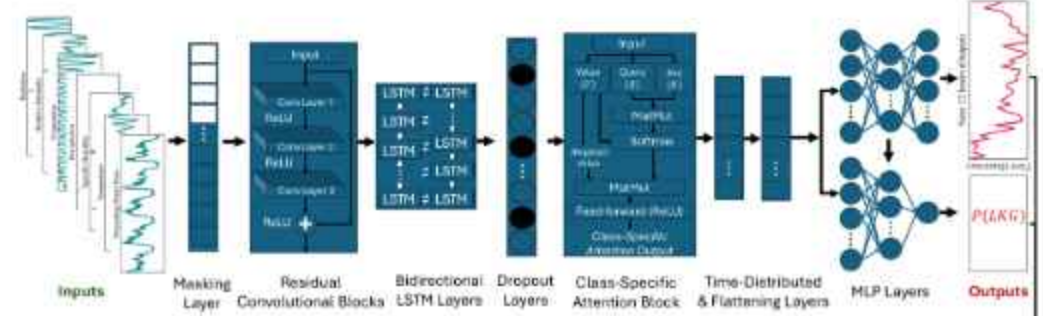
- ❑ Uses forecasted posterior water demands and leakage probabilities to decide whether and when to send repair team



Multi-objective
decision-making
agent



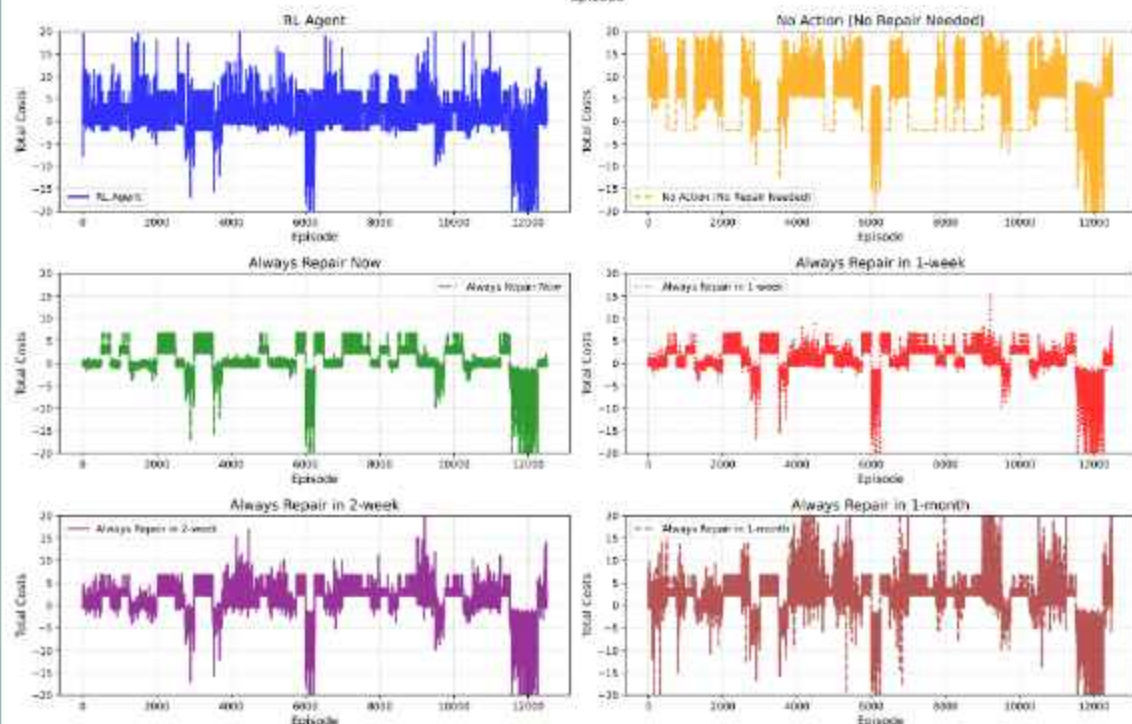
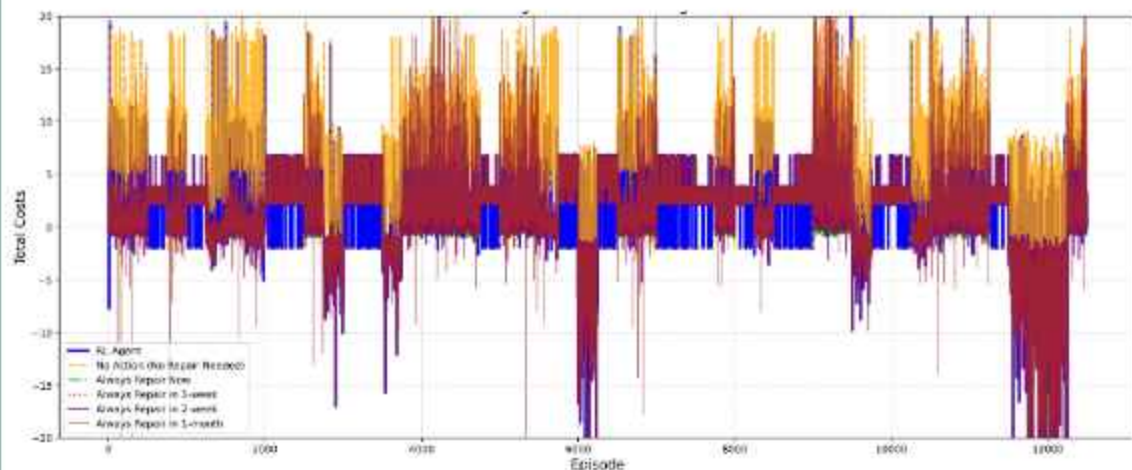
- ✓ Economic impact → £
- ✓ Reputation impact → CMEX score
- ✓ Public impact → Hours of exposure to nuisance
- ✓ Environmental impact → kg CO₂e



- 1) Send repair team now
- 2) Send repair team in 1 week
- 3) Send repair team in 2 weeks
- 4) Send repair team in 1 month
- 5) Don't send repair team

- Simulation environment considered variables that indirectly penalize or reward decisions
- Data obtained from > 500 reports and articles from water companies and councils
- Different types of continuous and discrete variables and probability distributions

Decision-Making Reinforcement Learning Agent



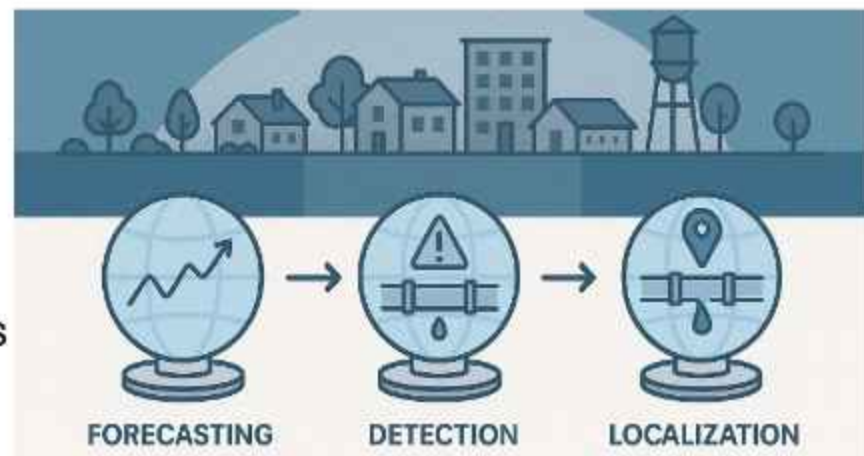
$$J(a_t, s_t) = I_{True\ Negative}(R_{TN}) + I_{other}(C_{econ} + C_{rep} + C_{pub} + C_{env} + P)$$

$$\theta^* = \arg \min_{\theta} \mathbb{E}_{\tau \sim \pi_{\theta}} \left[\sum_{t=0}^T \gamma^t J(a_t, s_t) \right]$$

- Uses the predicted flow values to decide **optimal action** out of 5 options
- The total cost contains **normalised summed-up costs** from the 4 considered impacts
- Lower cost means better action
- **Outperforms the other fixed strategies** throughout various episodes thereby helping decision-makers

Summary

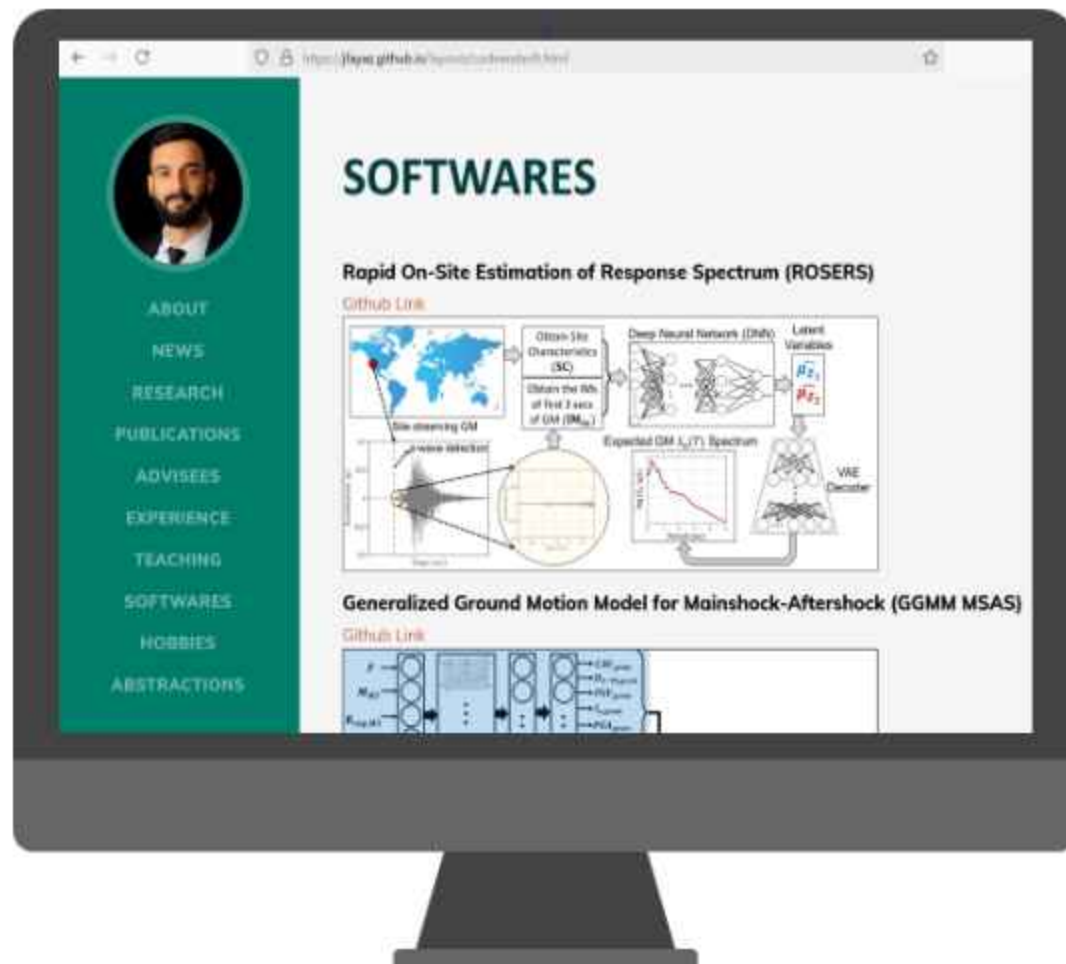
- ✓ **Three specialized AI frameworks** were developed:
 - ✓ **Forecasting** future leak risks using hybrid attention-based neural networks and Bayesian updating
 - ✓ Forecasted flow tied to **multi-objective decision-making** through deep reinforcement learning agent
 - ✓ **Detection** of ongoing leakages via domain-informed variational autoencoder and support vector machines
 - ✓ **Localization** of leakages through graph neural networks and transformer-based spatial pressure reconstructions
- ✓ High performance achieved on test sets:
 - ✓ **>90% accuracy (IA)** in forecasting 12 hrs ahead leakages
 - ✓ **>98% accuracy** with **AUC >0.99** for detecting previous leakages
 - ✓ **>0.93 accuracy (IA)** in spatial reconstruction of 24-hr pressure profiles
- ✓ Frameworks are designed to work with **limited existing sensor networks**, minimizing the need for costly upgrades.
- ✓ Future work aims to **integrate forecasting, detection, and localization into a unified, end-to-end decision-support framework** — enabling proactive, real-time water network management (from specialized AI models to holistic smart water management)
- ✓ The integrated system will support **automated leak alerts, dynamic repair scheduling, and adaptive sensor deployment**, pushing the transition towards **resilient and self-healing water distribution networks**





University
of Exeter

<https://jfayaz.github.io/>



THANKS FOR ATTENDING



University
of Exeter

Dr. Jawad Fayaz

j.fayaz@exeter.ac.uk

<https://jfayaz.github.io>