Towards Using Neural Networks for Geoscientific Discovery

Presenter: Ben Toms

My Dissertation Collaborators


Prof. Elizabeth Barnes


Prof. Imme Ebert-Uphoff


Dr. Karthik Kashinath

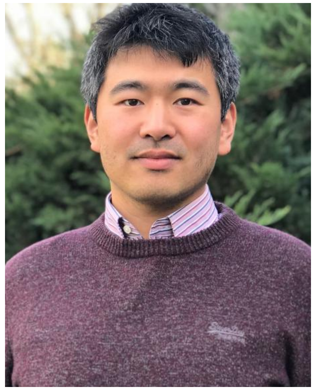
Prof. Da Yang


Dr. Prabhat


Prof. James Hurrell

What does a basic neural network look like?

Knowns

Input Sample

x_1

x_2

x_3

⋮

x_n

Input Label

y_1

y_2

Summary: connected layers of nonlinear regression

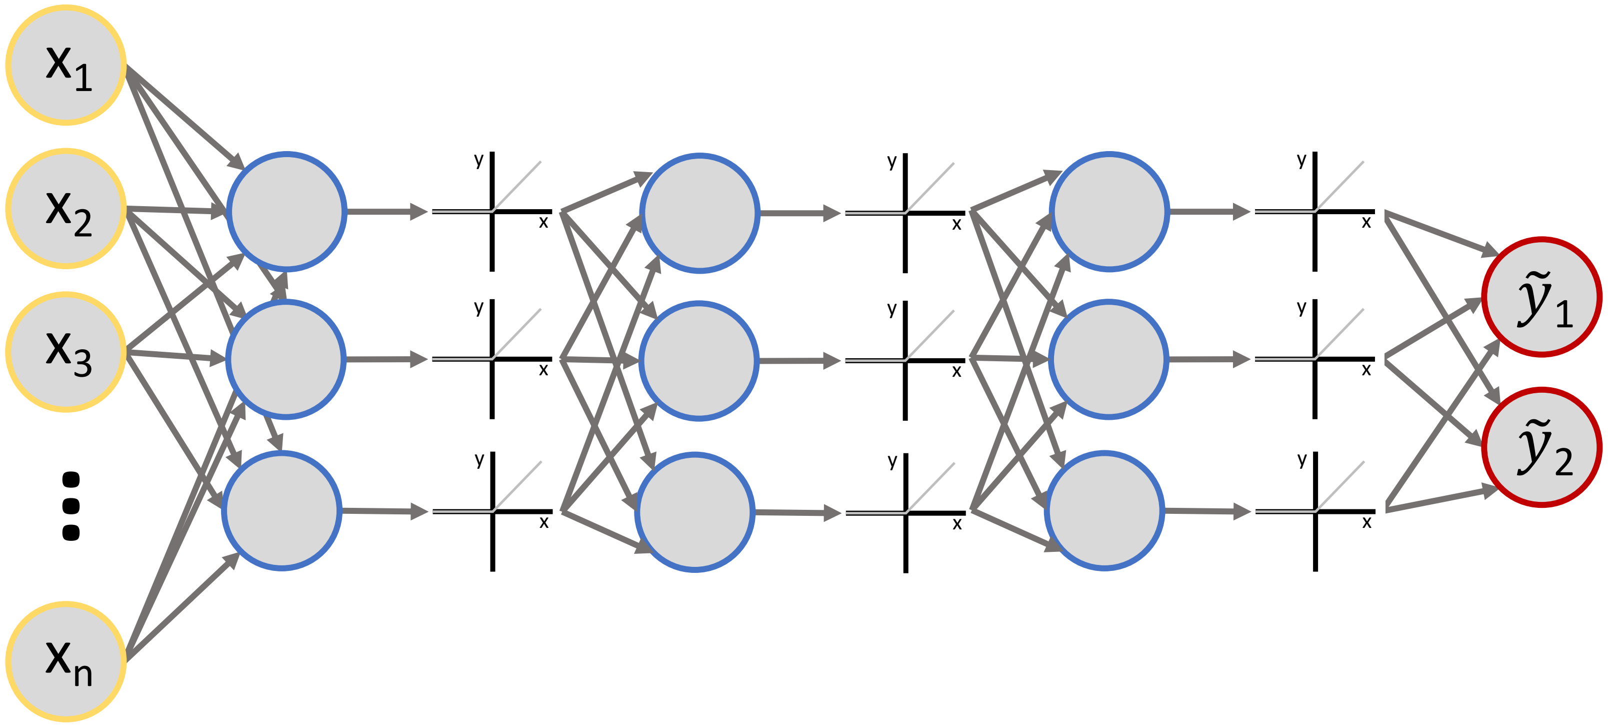

MIND BENDER

Physicist: The Entire Universe Might Be a Neural Network

"The idea is definitely crazy, but if it is crazy enough to be true? That remains to be seen."

VICTOR TANGERMANN | SEPTEMBER 9TH 2020

What if we can better our understanding of nonlinearities within the earth system by interpreting what neural networks learn?

We know this: oceanic patterns can be used to predict other weather and climate patterns years in advance

Sea-surface temperature anomalies

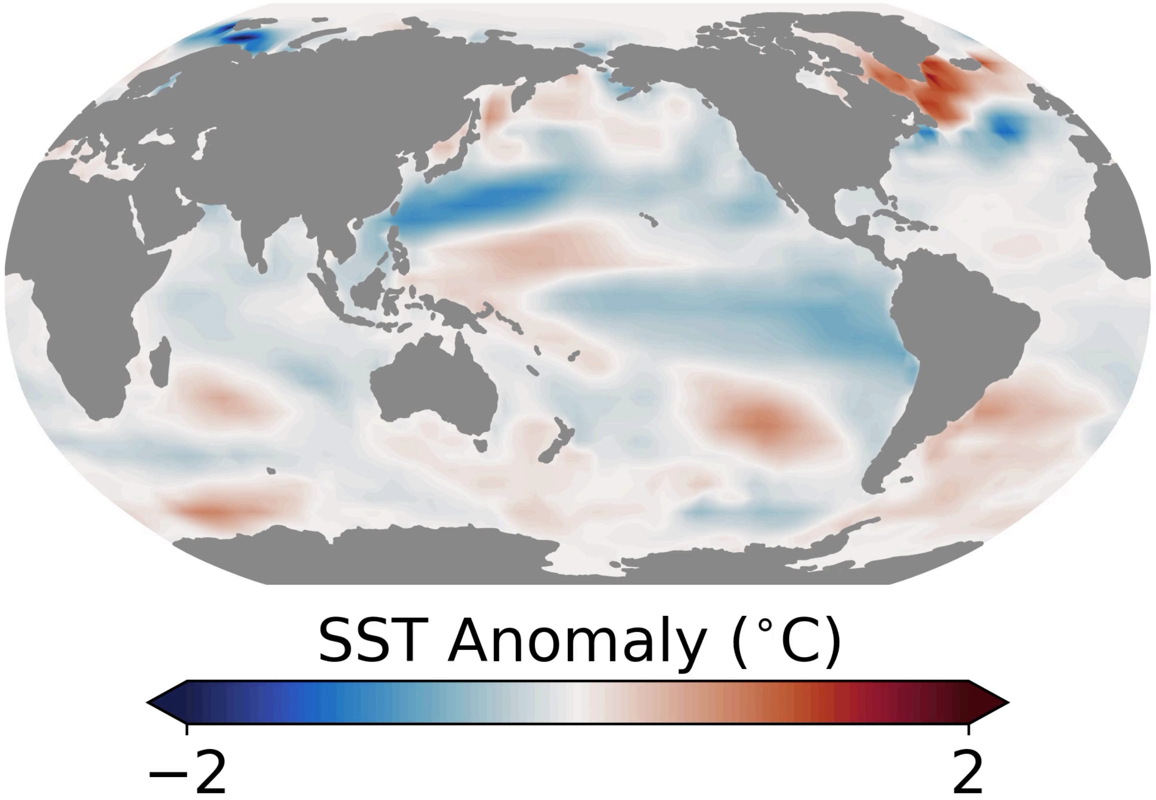

Continental surface temperature anomalies

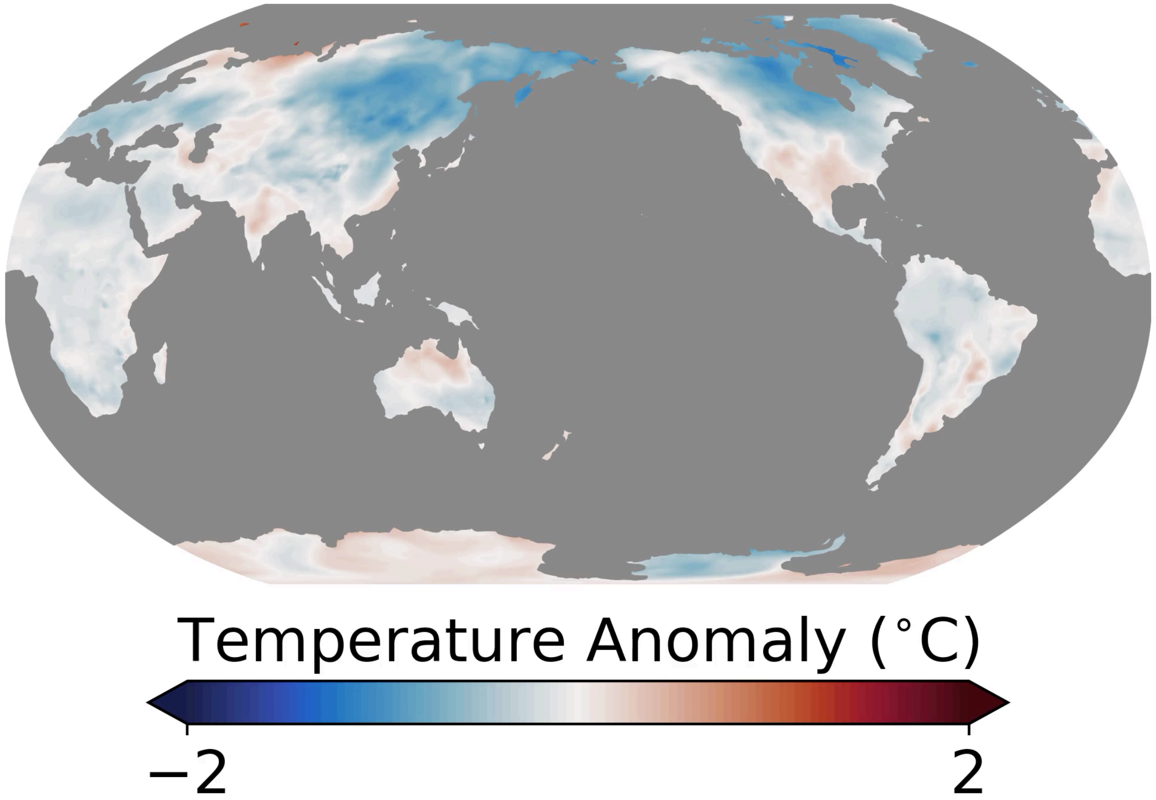

The end-goal:

Identify oceanic patterns that lend multi-year predictability, then assess their nonlinearity

Sea-surface temperature anomalies

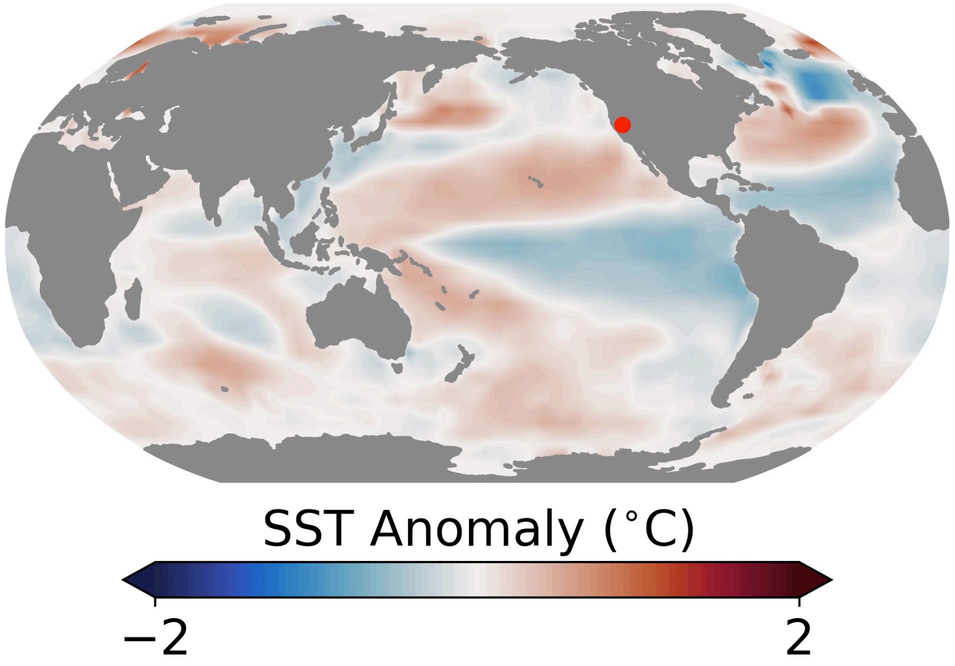

Regions of the anomalies that lend predictability

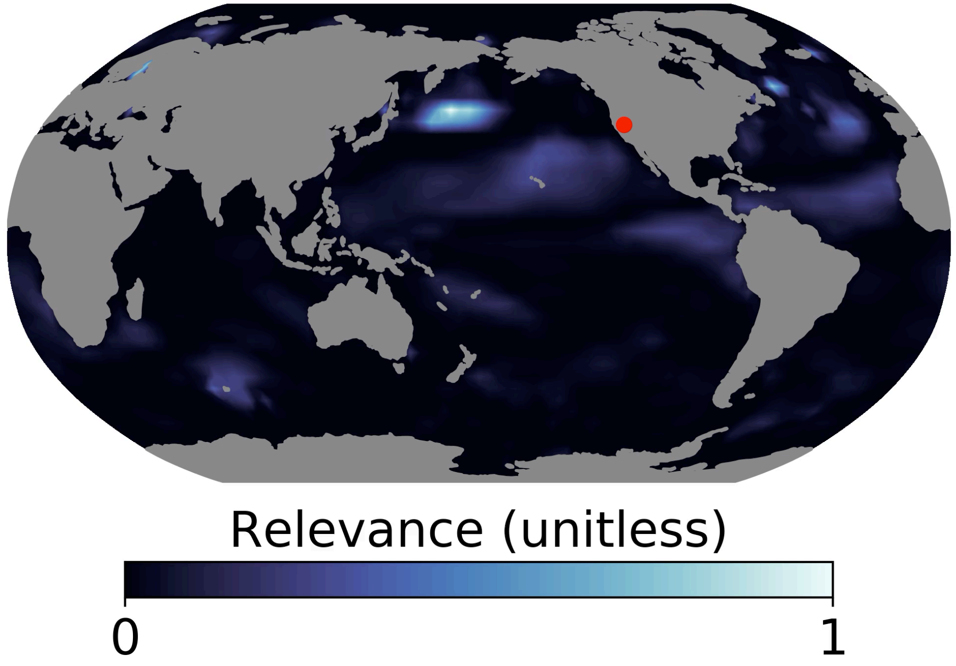

Problem #1

There isn’t a framework for clearly understanding how and why neural networks make their decisions for geoscientific applications.

Problem #2

If the framework is developed, it needs to be tested on multiple applications to ensure its reliability.

Problem #3

We can then start applying the framework to furthering our understanding of earth-system predictability.

A method for interpreting neural networks…

Step 1) Input a sample into a trained network…

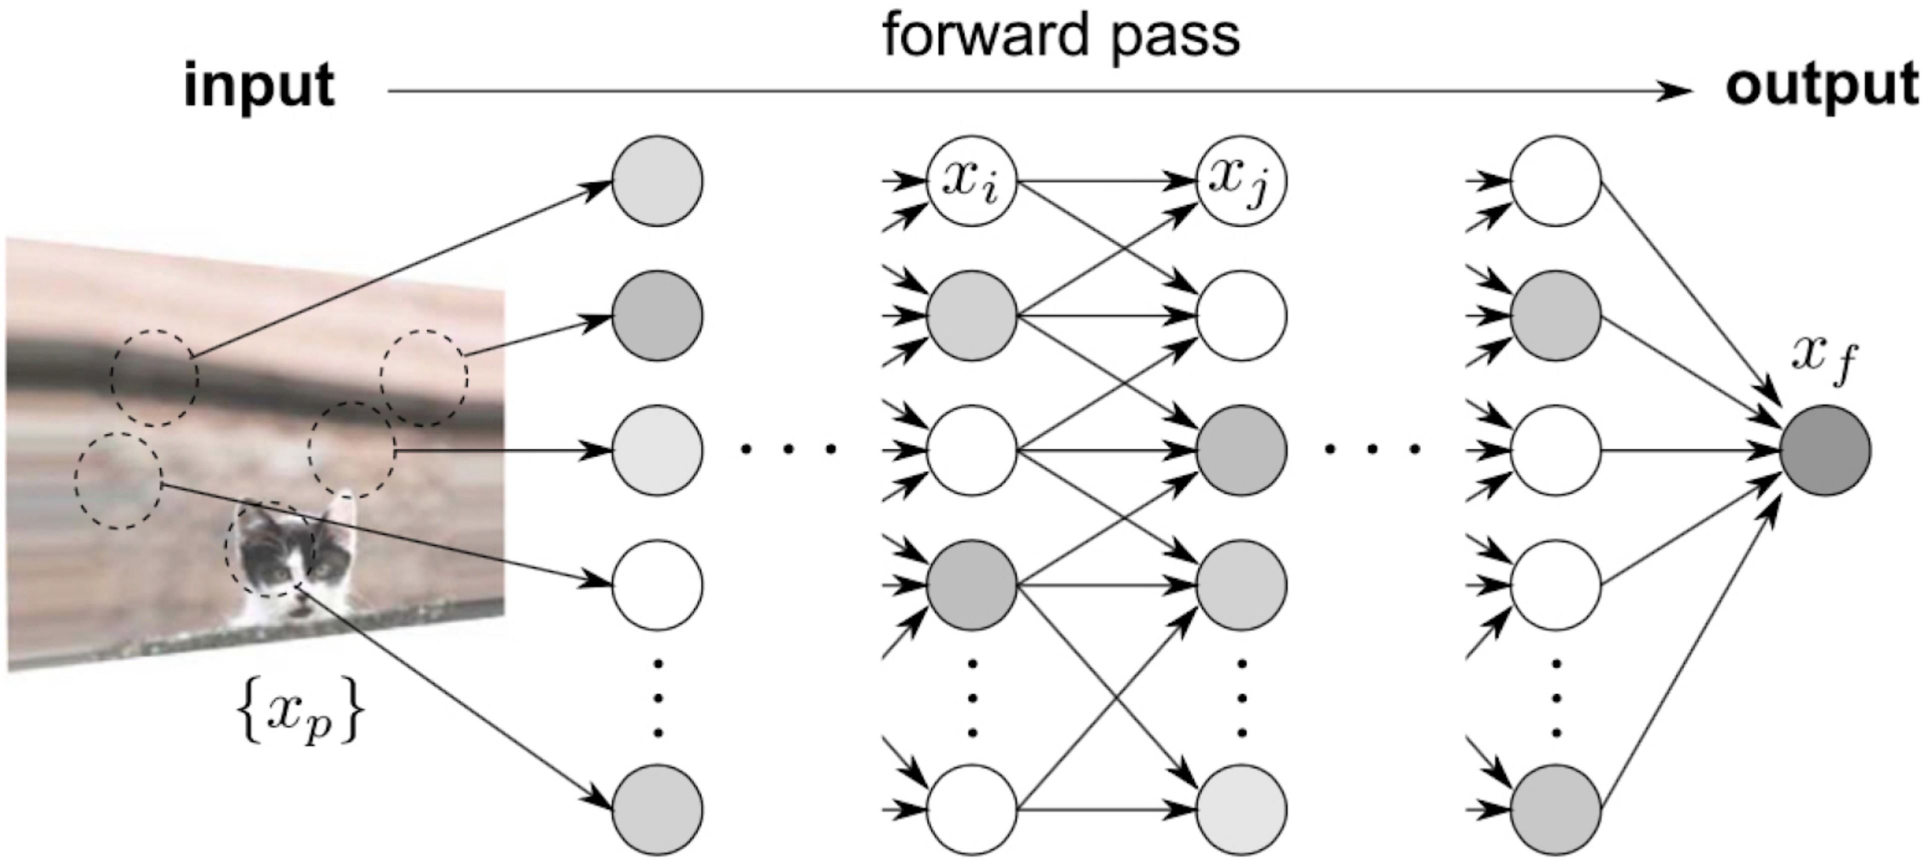

From Montavon et al. (2017); Layerwise Relevance Propagation

A method for interpreting neural networks…

Step 2) Trace backwards to find which inputs were most relevant for the output…

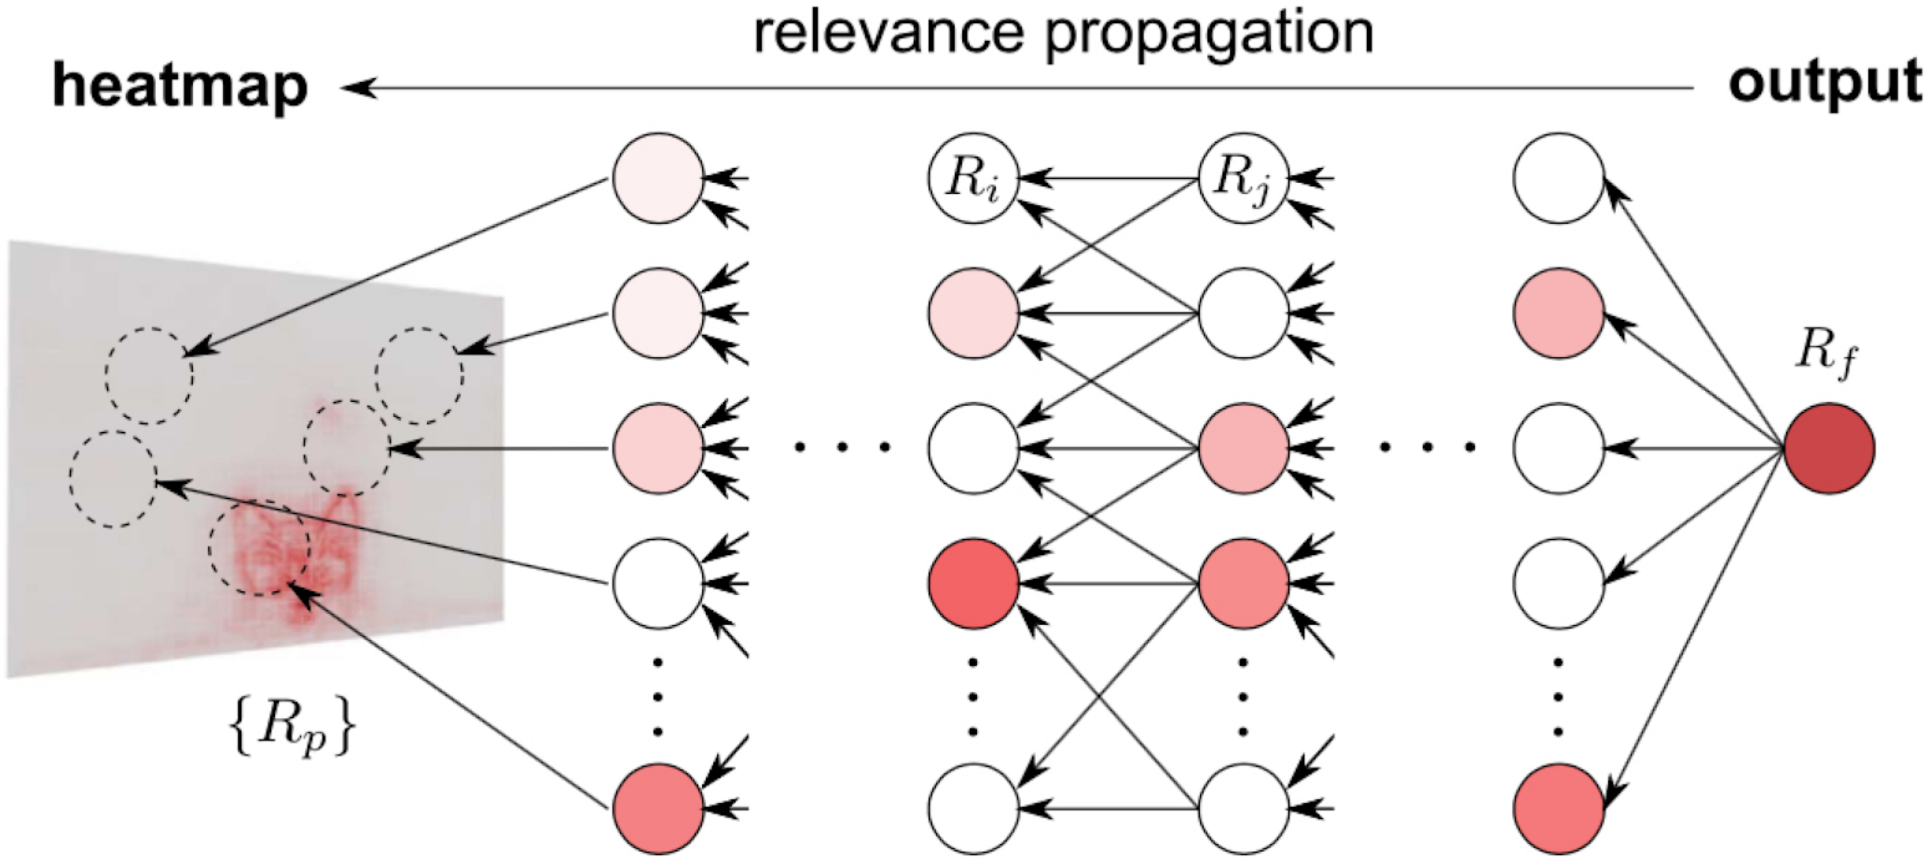

From Montavon et al. (2017); Layerwise Relevance Propagation

Testing the interpretations using a simple application

Inputs

Global maps of sea-surface temperature

Outputs

The sign of surface temperature in the red box

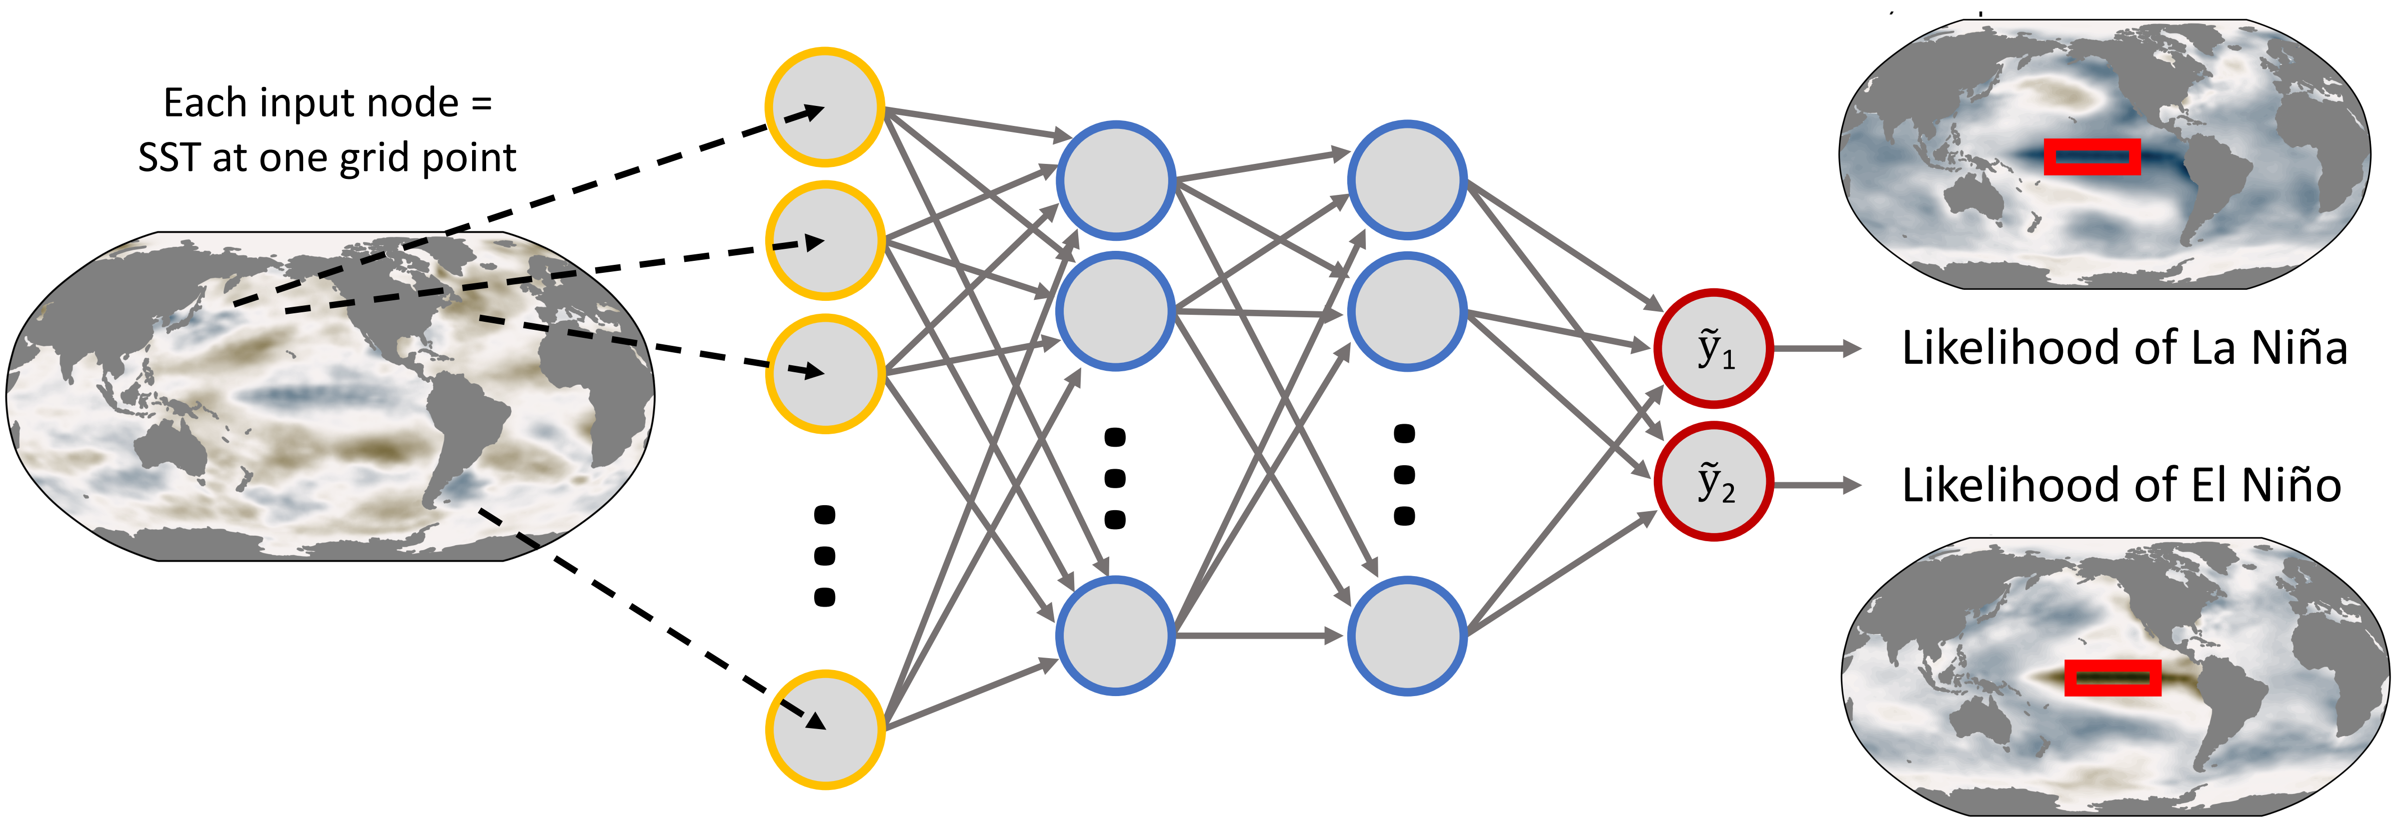

The interpretability method works well in this case

Where does the neural network focus its attention for each sample?

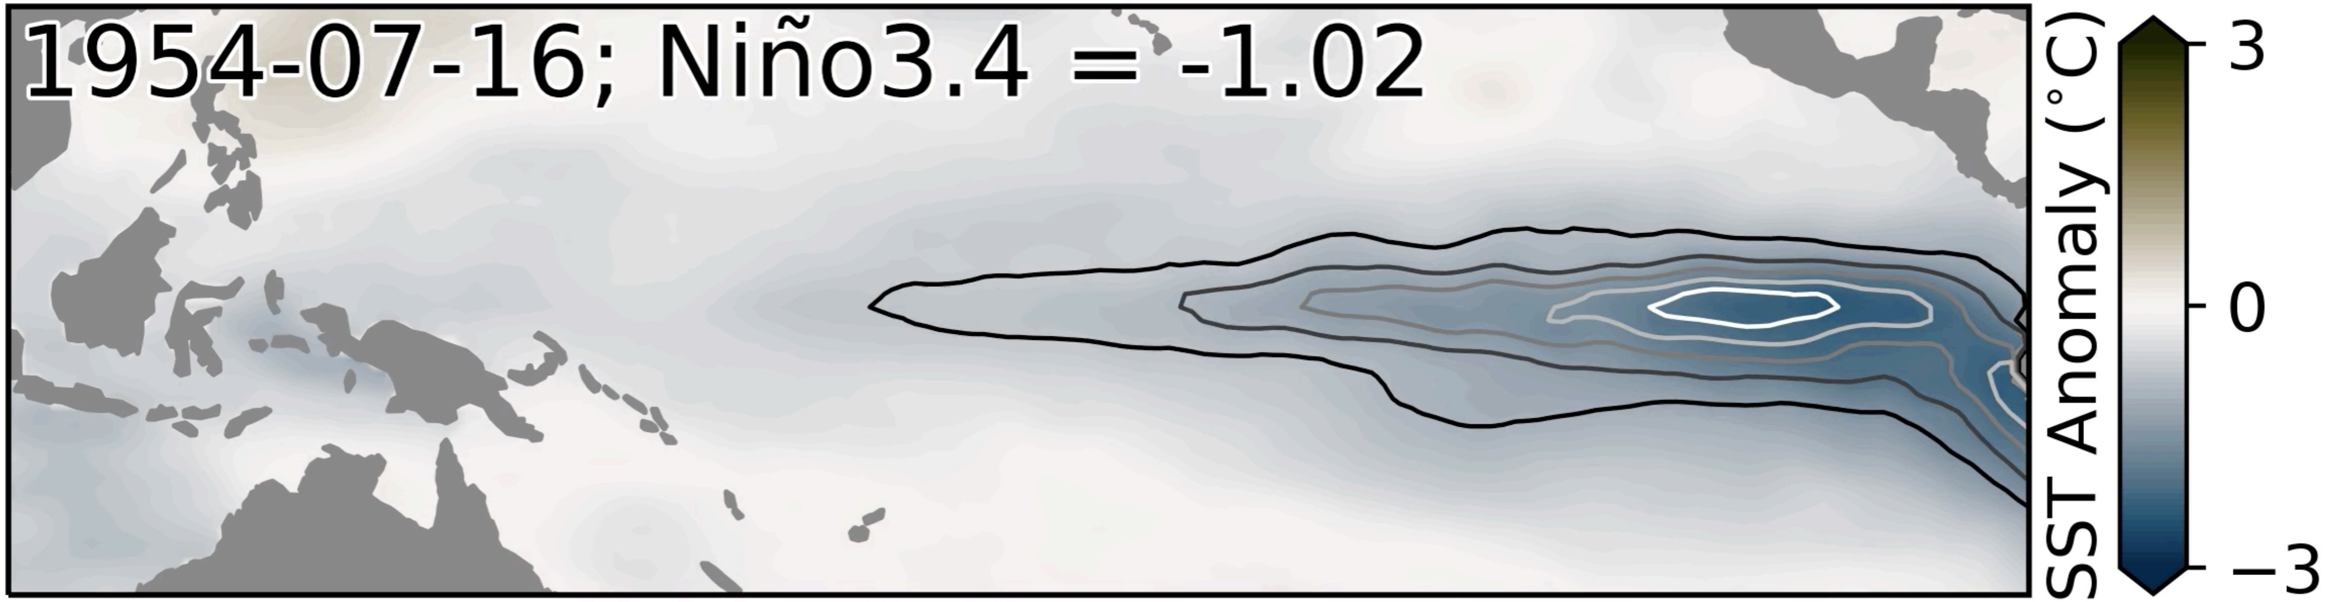

Problem #1

There isn't a framework for clearly understanding how and why neural networks make their decisions for geoscientific applications.

Problem #1

There isn’t a framework for clearly understanding how and why neural networks make their decisions for geoscientific applications.

Problem #2

If the framework is developed, it needs to be tested on multiple applications to ensure its reliability.

Comparing nonlinear and linear approaches

Nonlinear

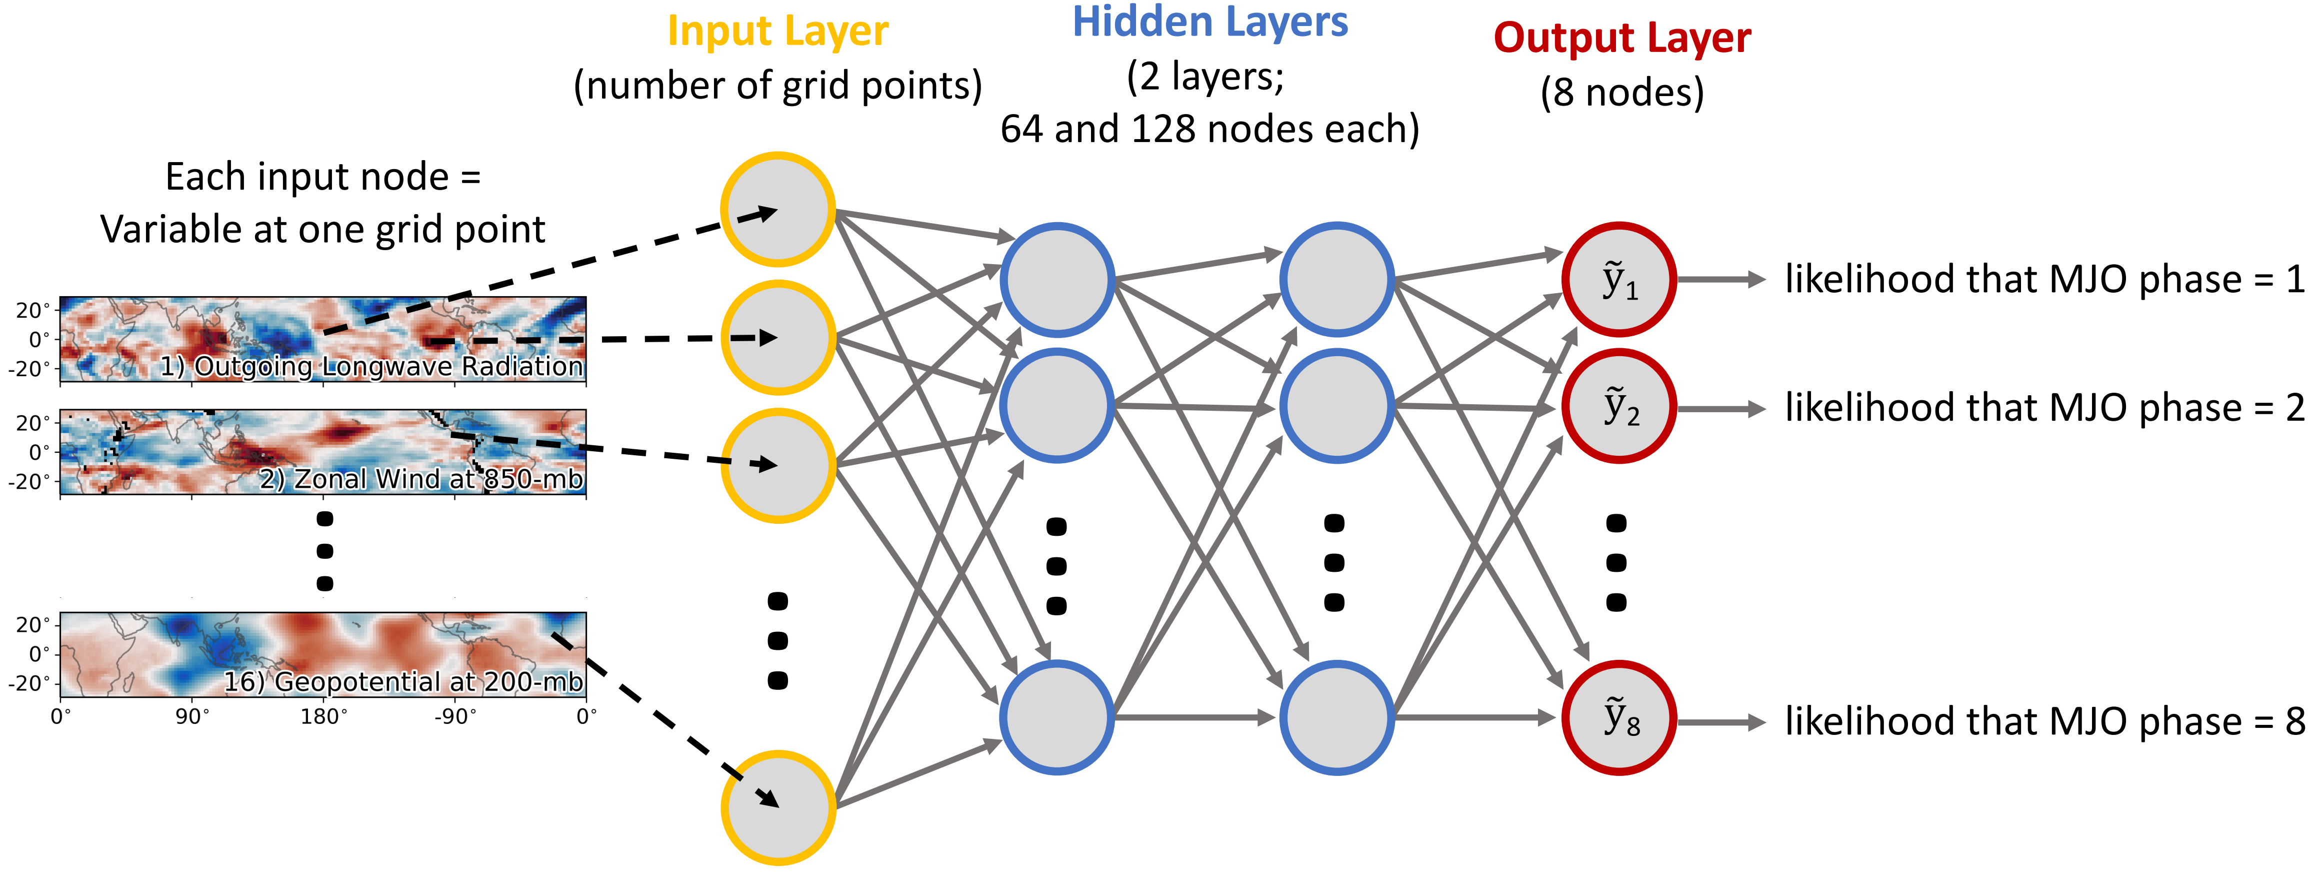

Comparing nonlinear and linear approaches

Linear

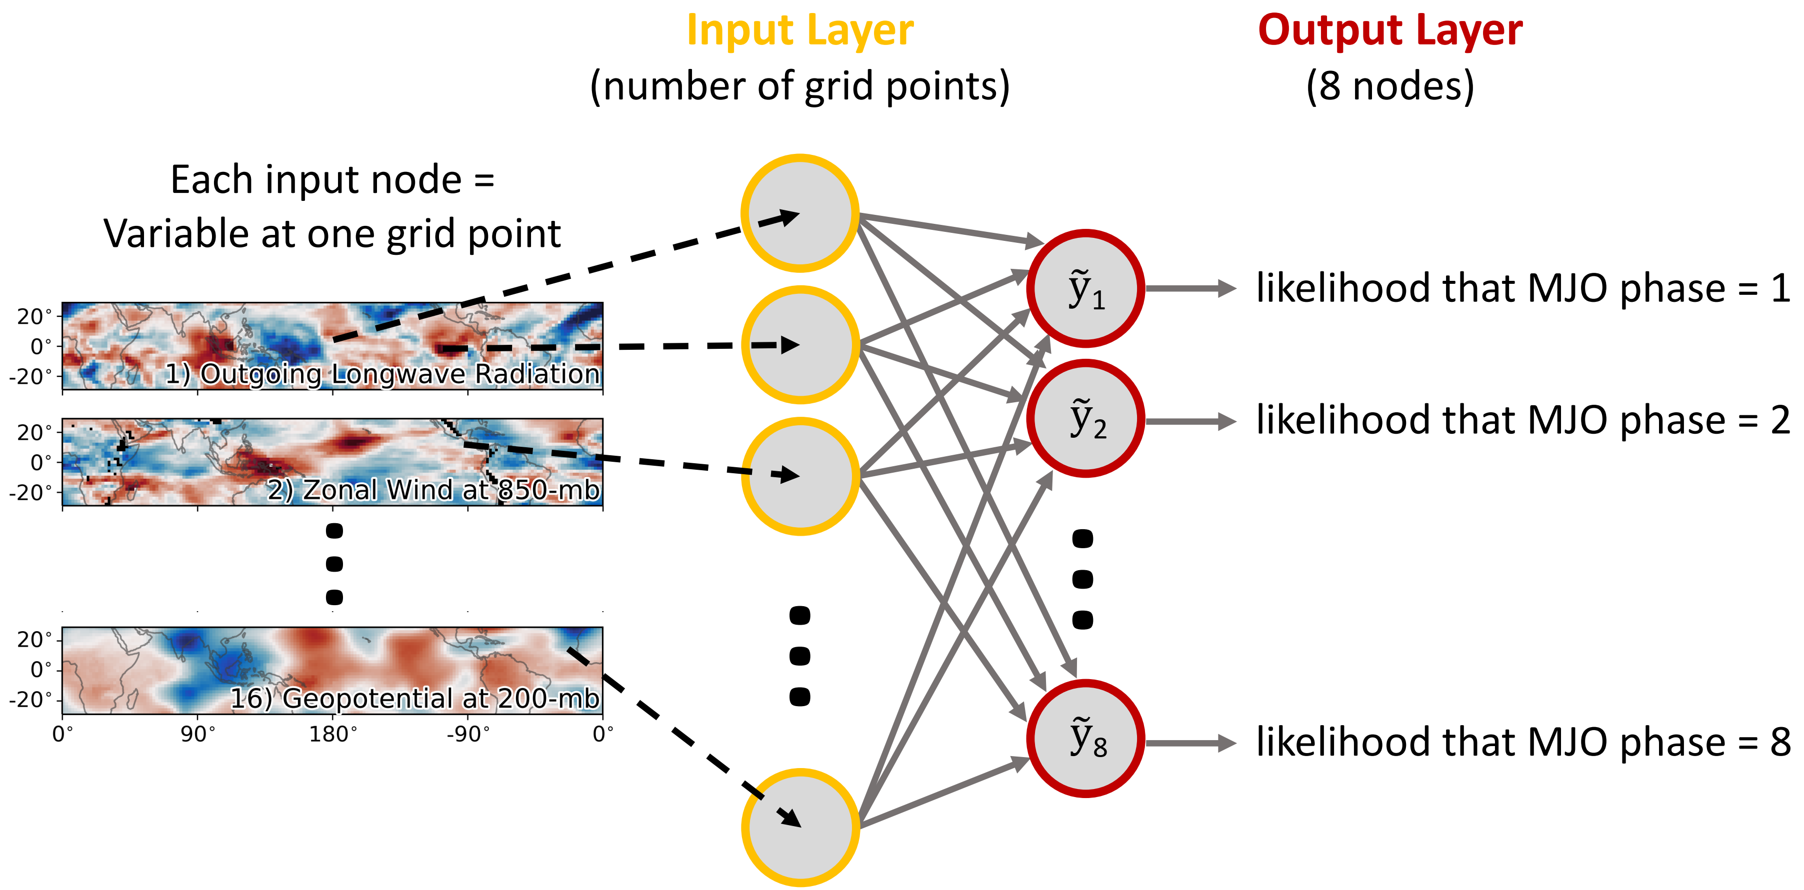

Comparing nonlinear and linear approaches

Nonlinear

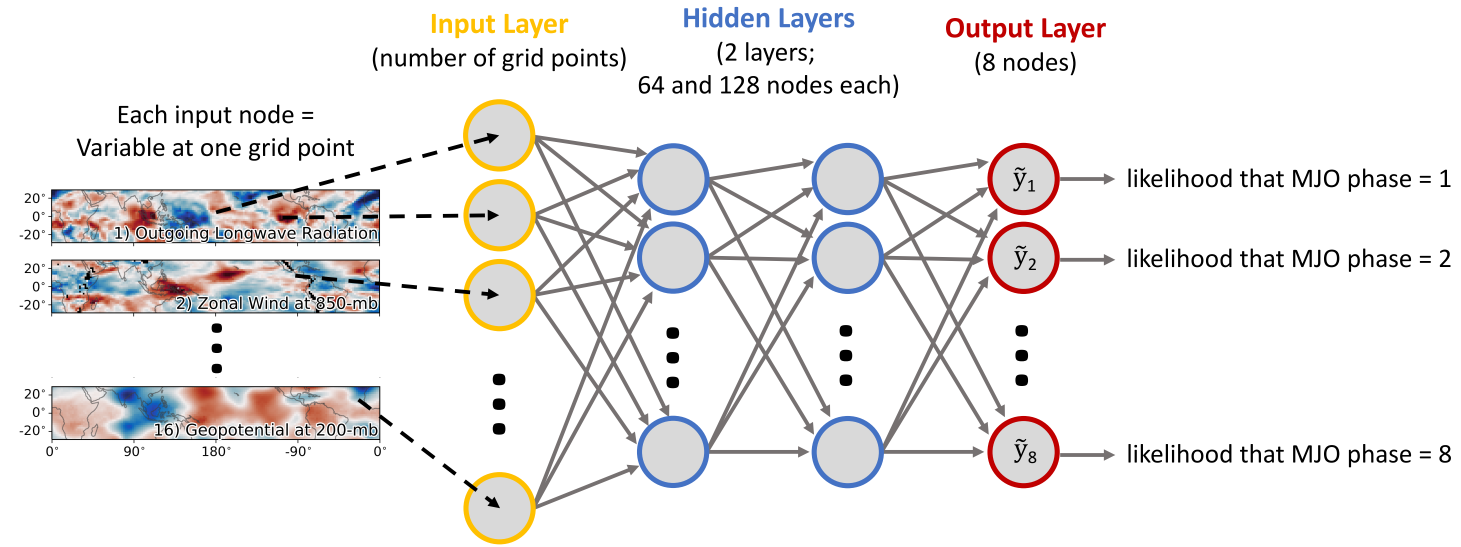

Linear

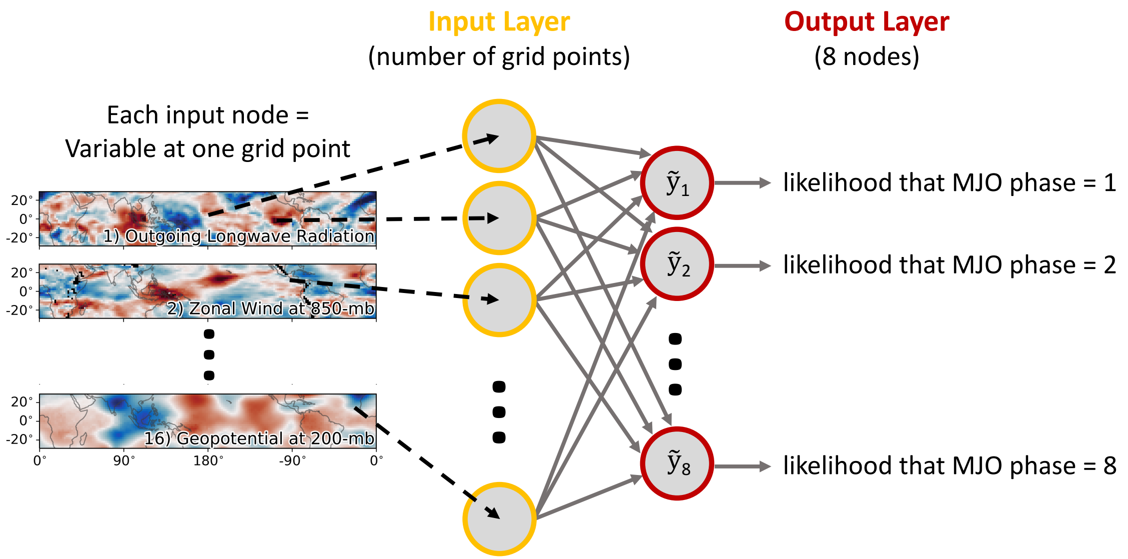

Nonlinear vs linear accuracy

The neural network approach is more accurate than the linear approach.

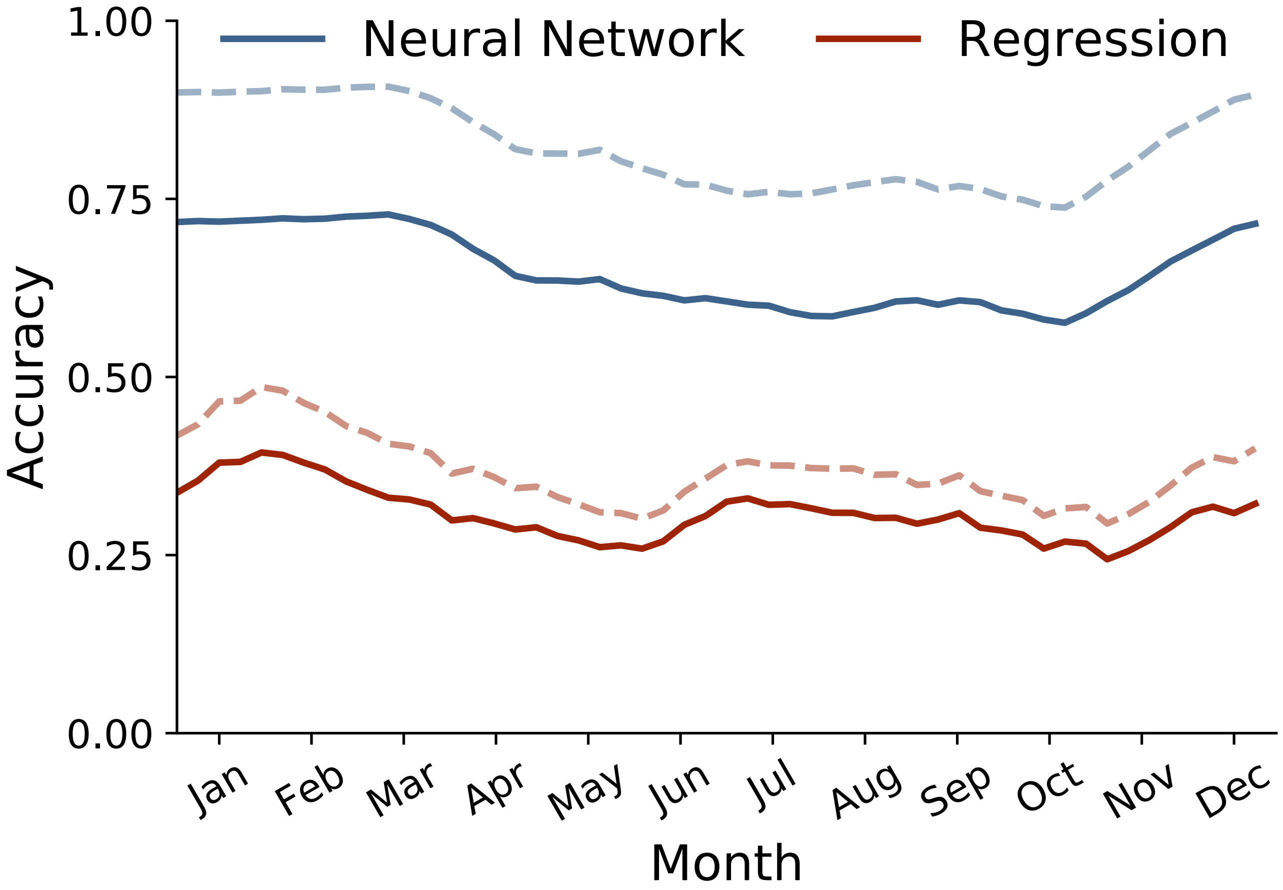

Case-by-case examples of MJO nonlinearity

Composite Phase 7 Cloud Pattern

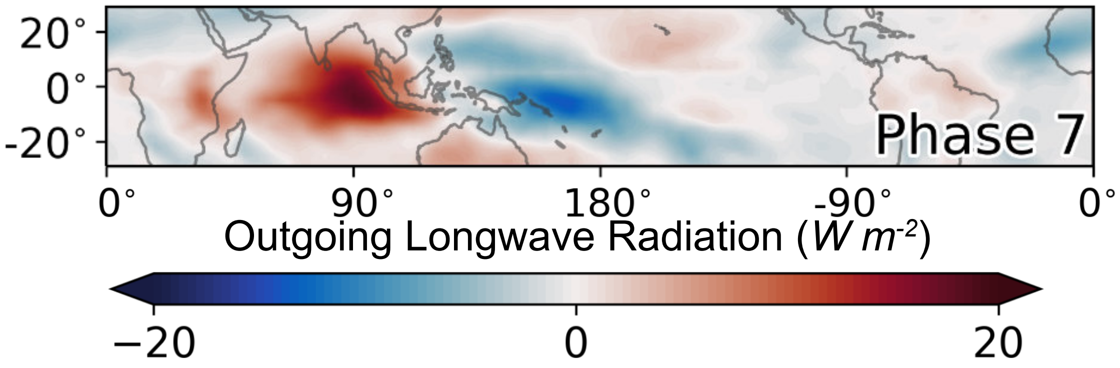

Example Phase 7 - January 16, 1989

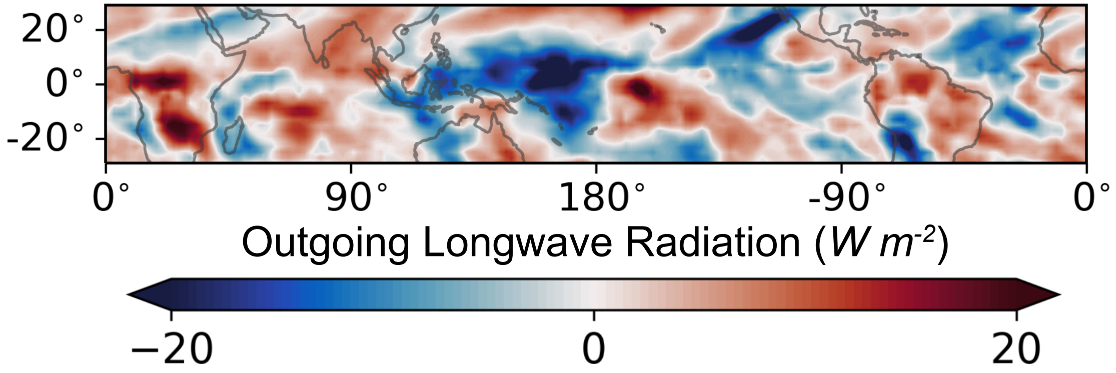

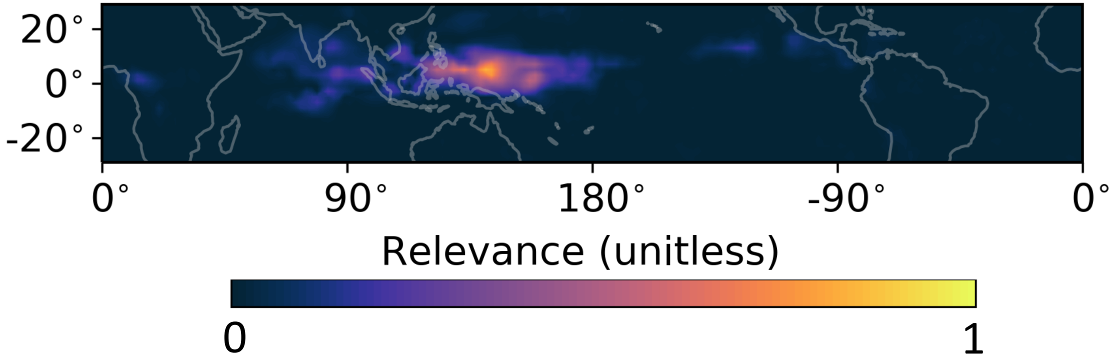

Case-by-case examples of MJO nonlinearity

Composite Phase 7 Cloud Pattern

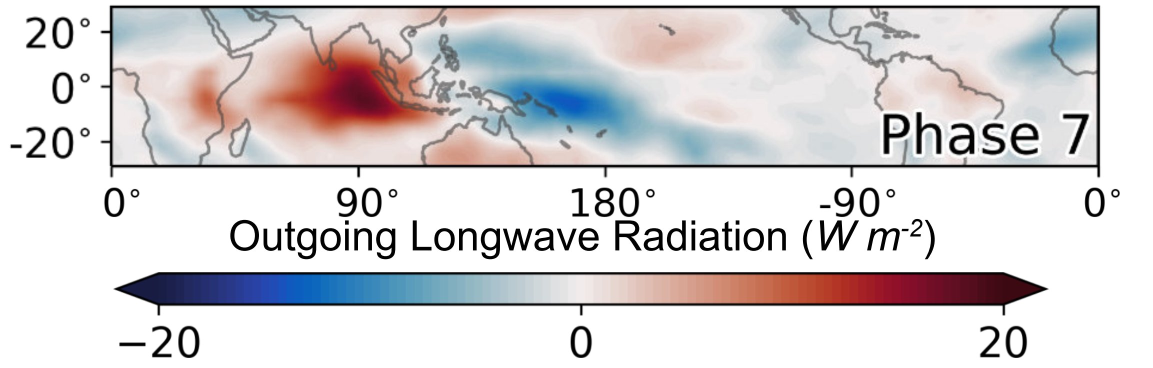

Example Phase 7 – November 27, 1998

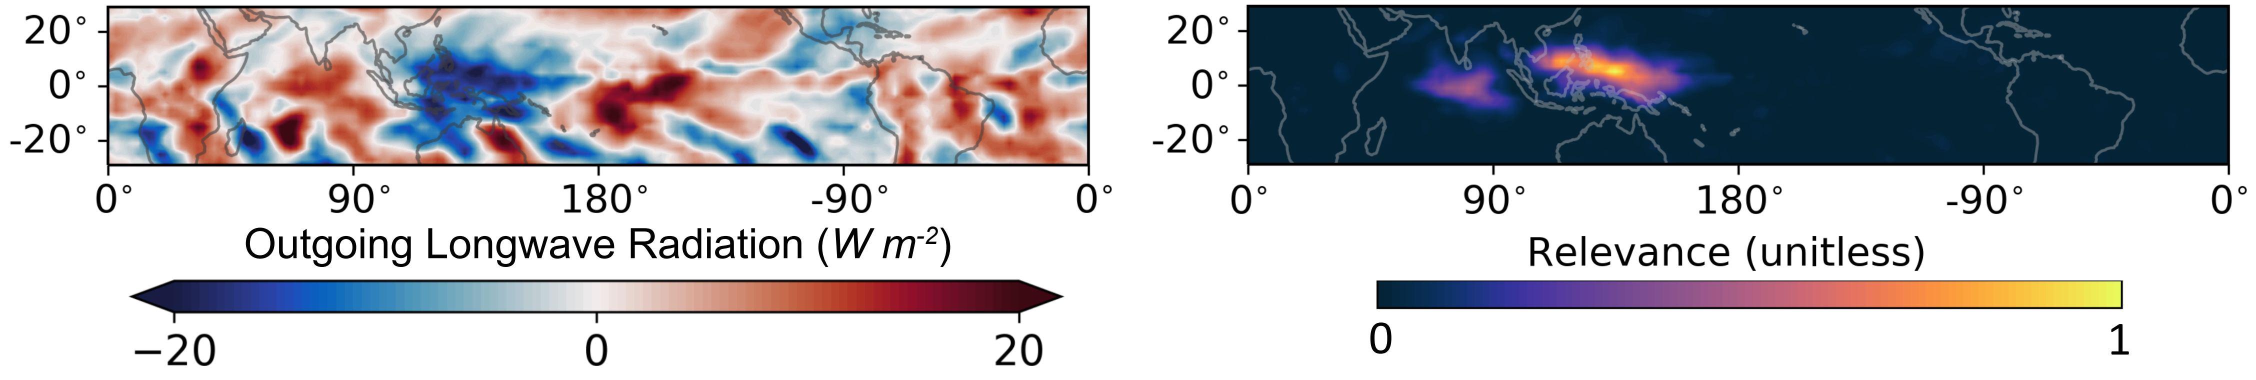

Case-by-case examples of MJO nonlinearity

Composite Phase 7 Cloud Pattern

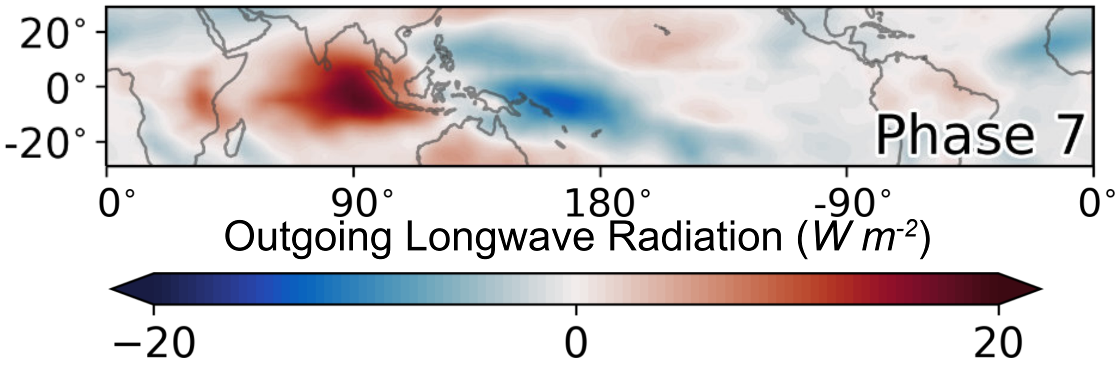

Example Phase 7 – February 2, 1985

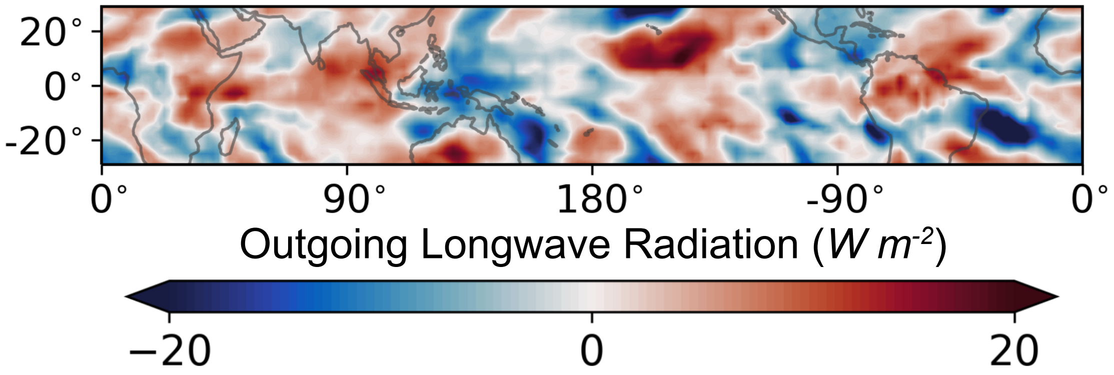

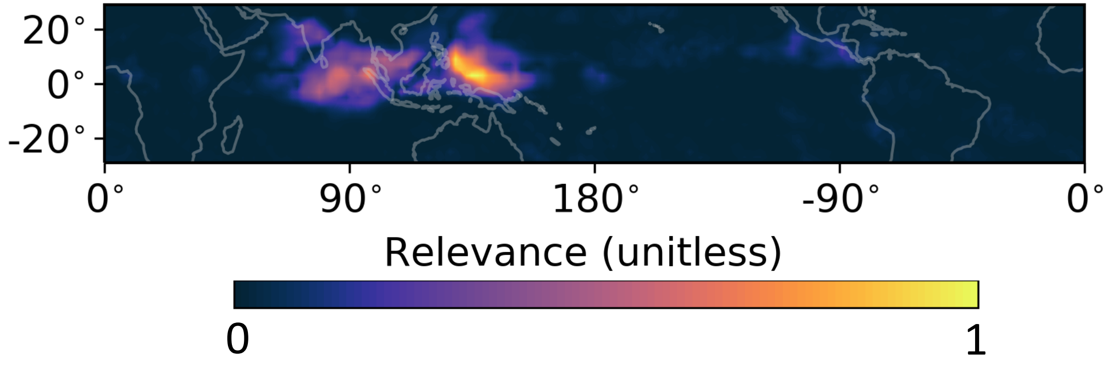

Problem #1

There isn't a framework for clearly understanding how and why neural networks make their decisions for geoscientific applications.

Problem #2

If the framework is developed, it needs to be tested on multiple applications to ensure its reliability.

Problem #1

There isn't a framework for clearly understanding how and why neural networks make their decisions for geoscientific applications.

Problem #2

If the framework is developed, it needs to be tested on multiple applications to ensure its reliability.

Problem #3

We can then start applying the framework to furthering our understanding of earth-system predictability.

The simple neural network design for multi-year forecasts:

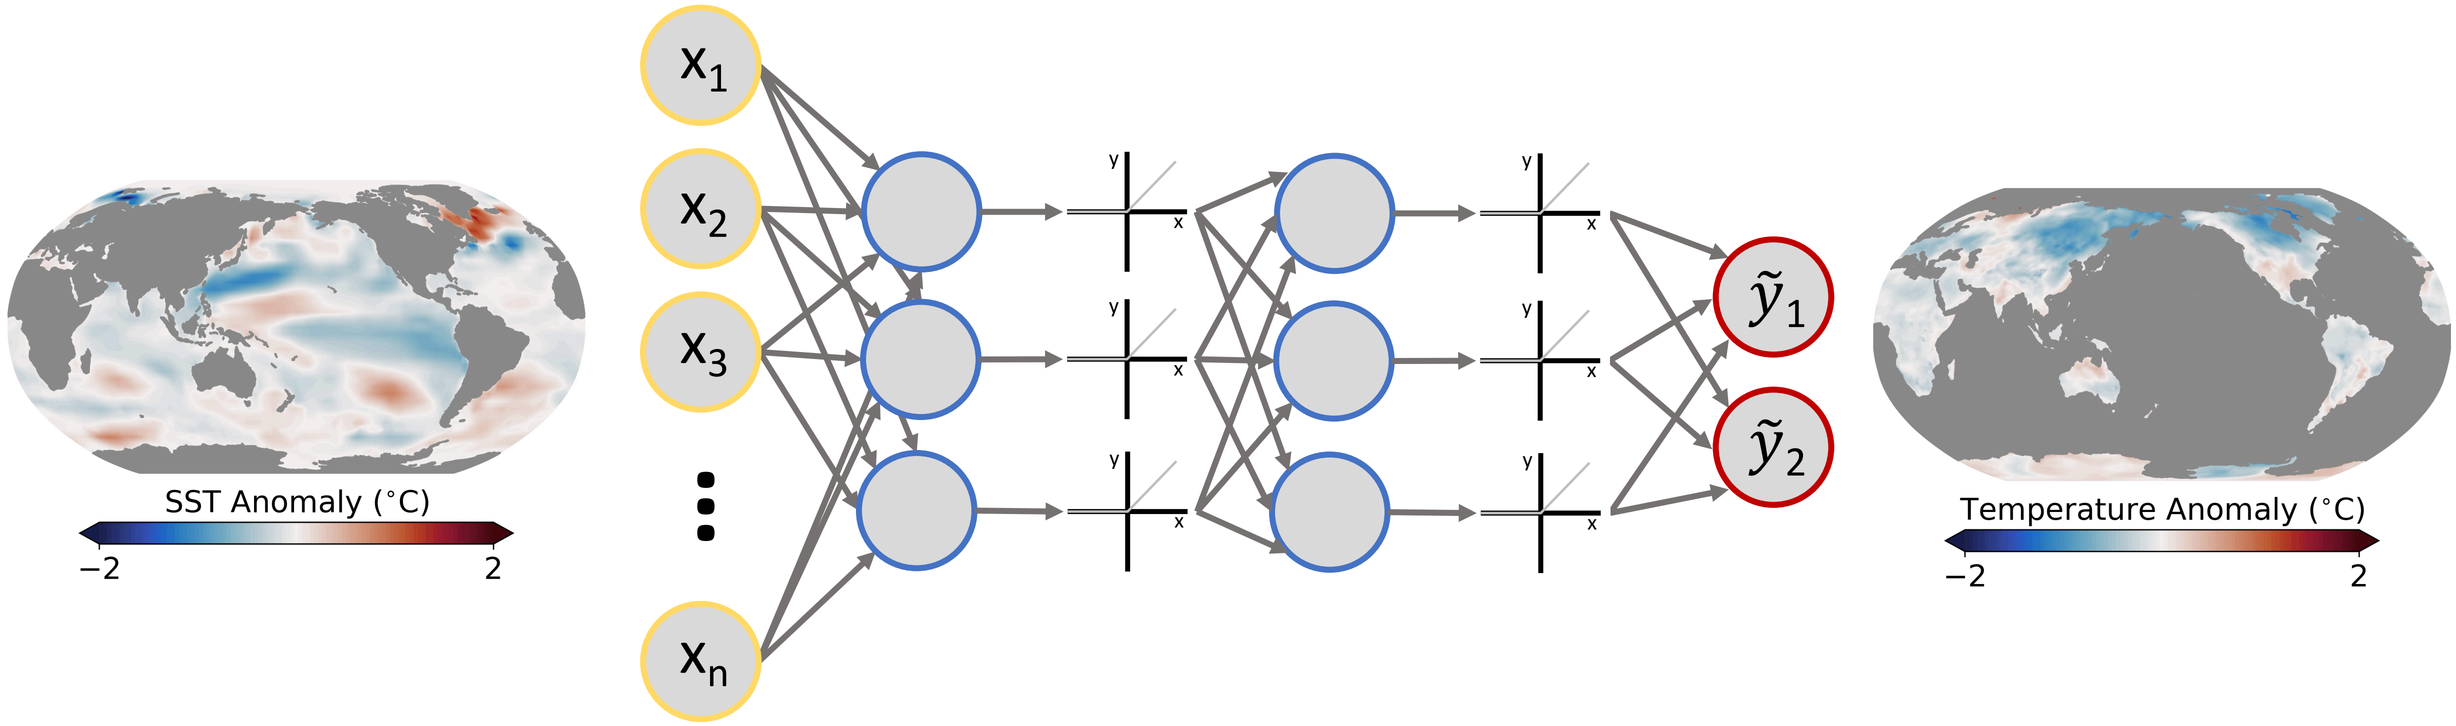

Looking at where the neural network looks to make its predictions….

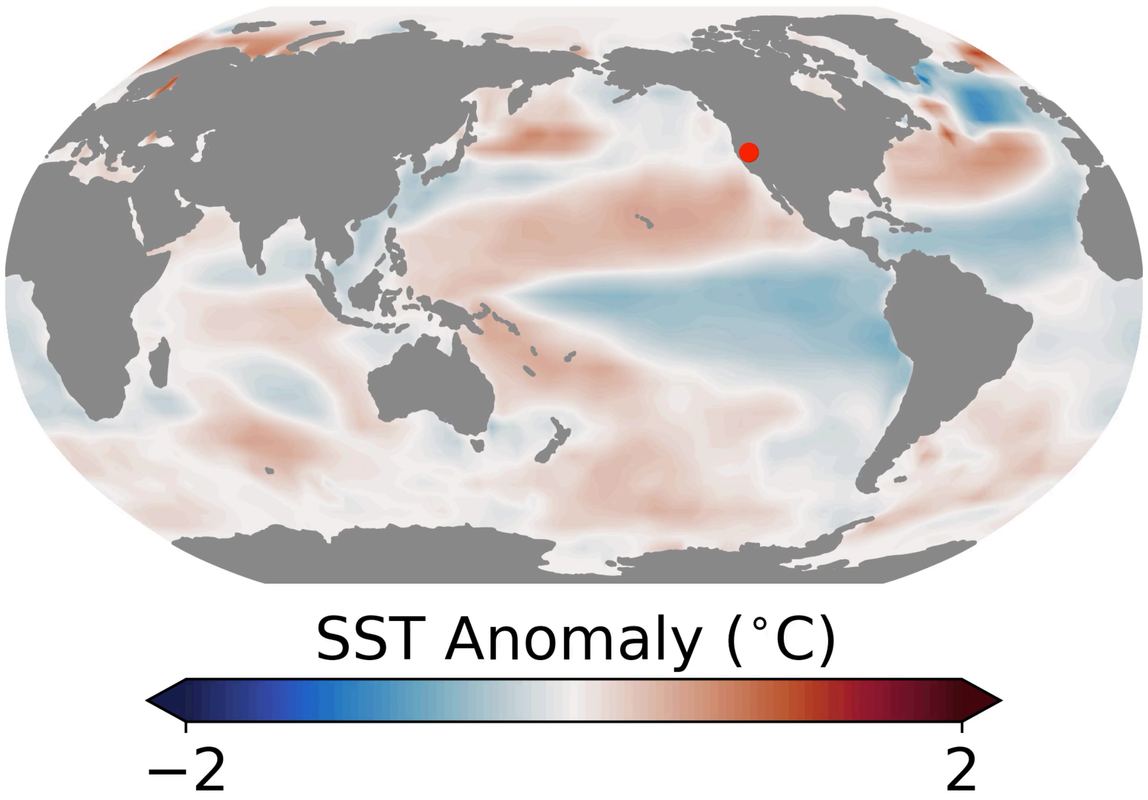

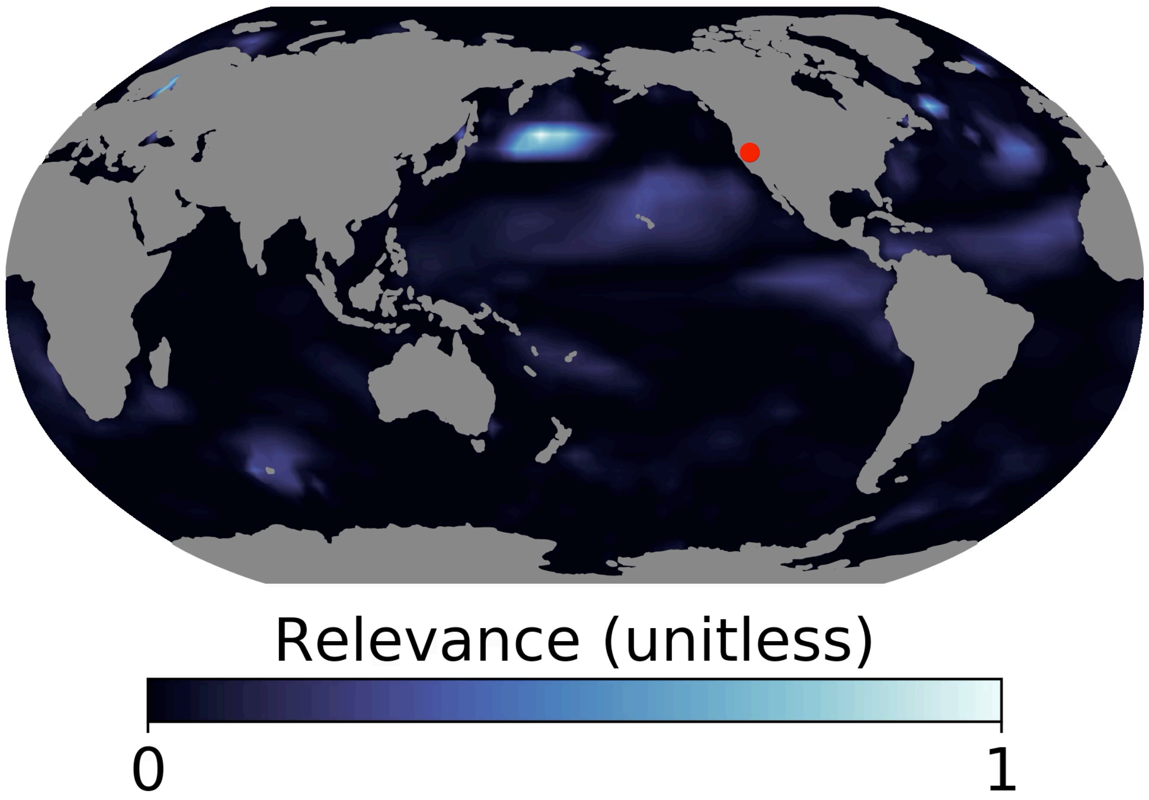

Composite interpretation maps for different locations

The interpretations can be used to understand which oceanic patterns lead to predictability at any location.

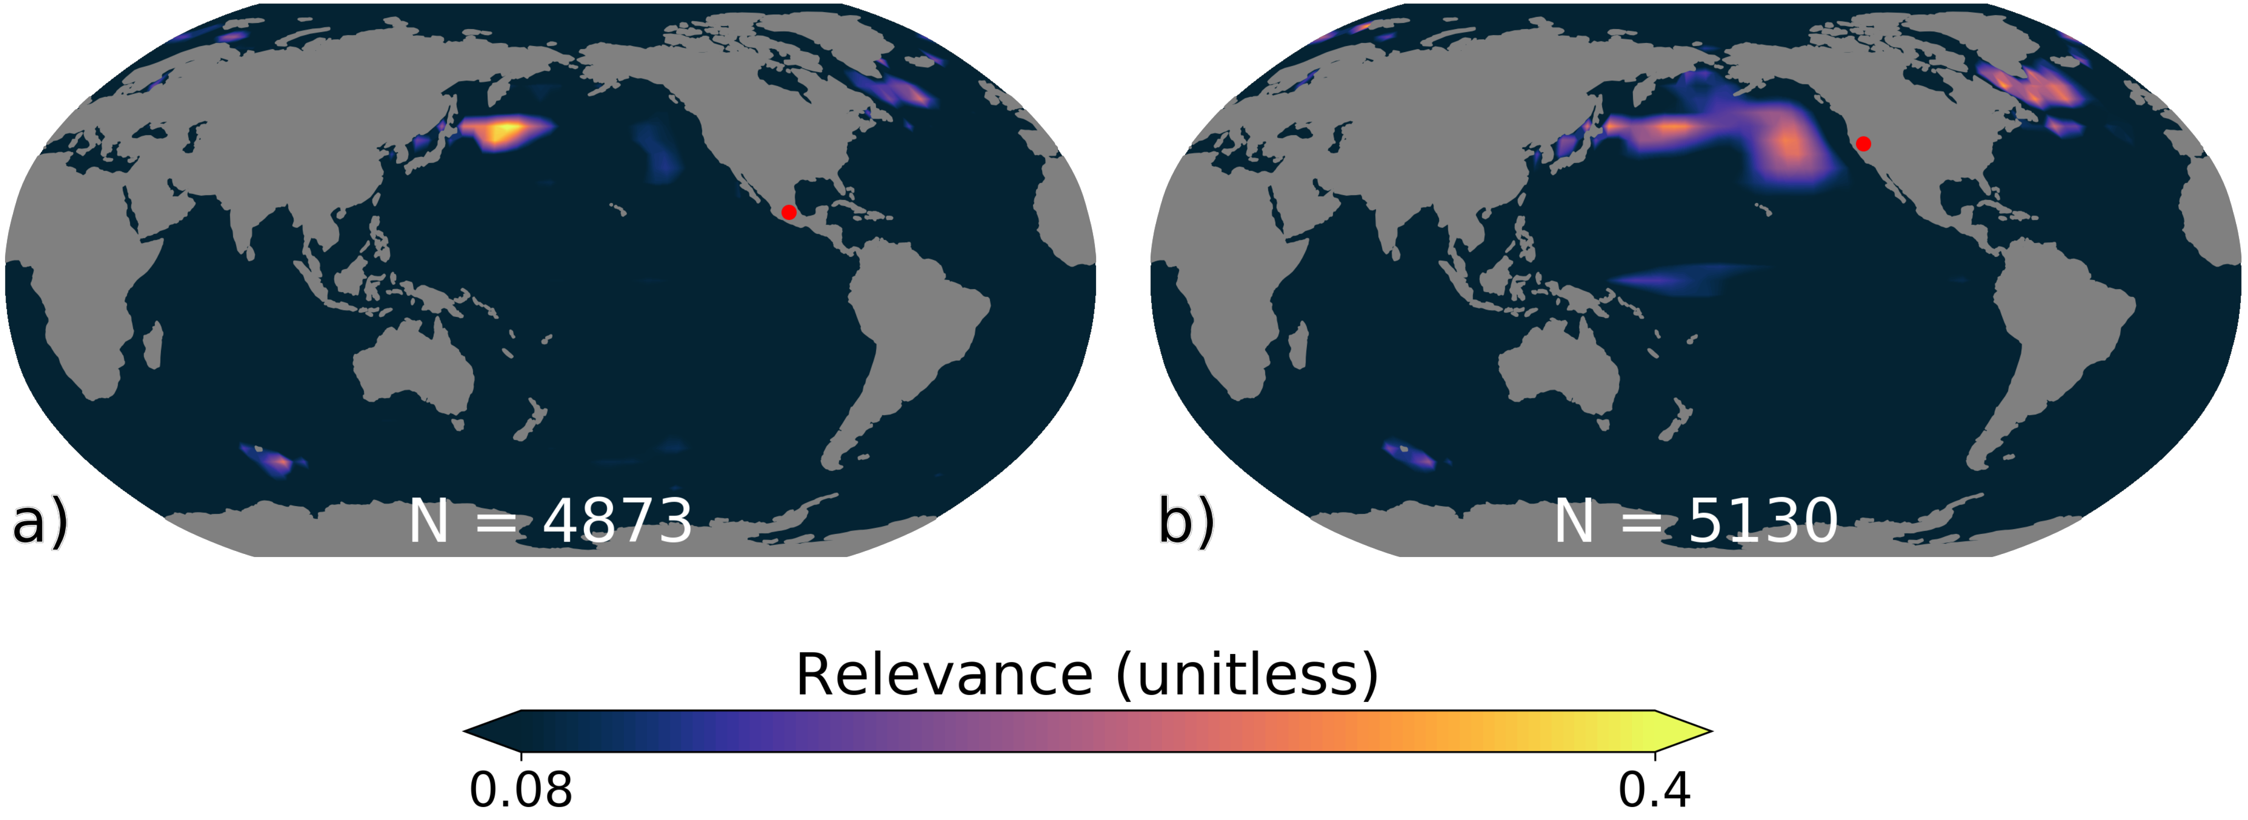

Interpretation map clusters for one location

Distinct regimes of predictability become apparent when the LRP heatmaps are clustered into their dominant patterns.

average pattern

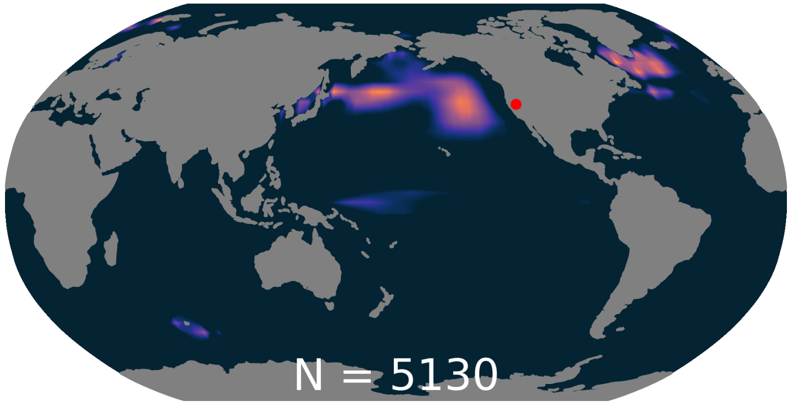

Interpretation map clusters for one location

Distinct regimes of predictability become apparent when the LRP heatmaps are clustered into their dominant patterns.

average pattern

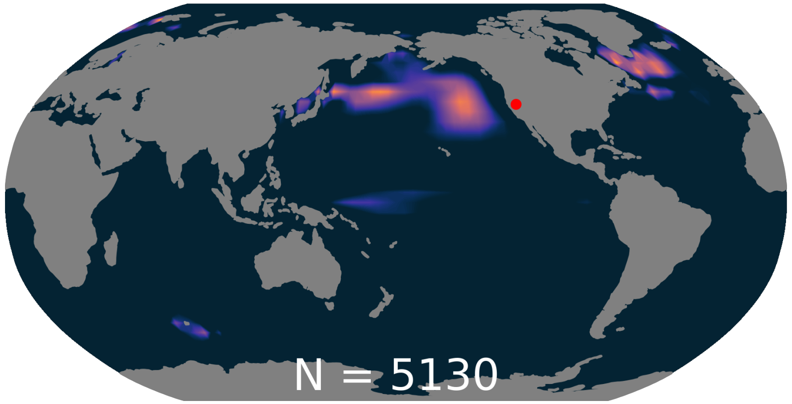

most common

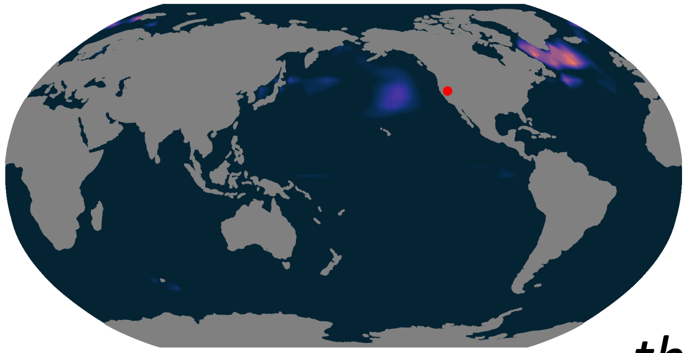

second most common

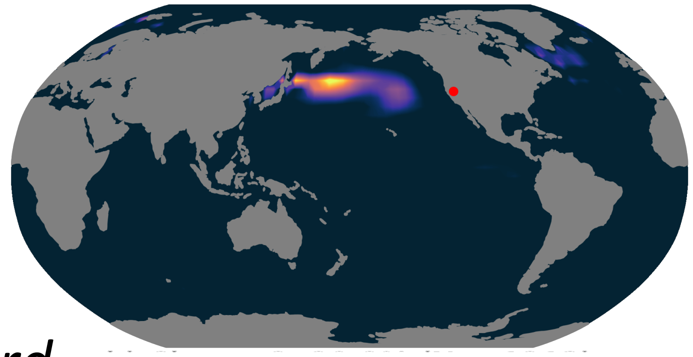

third most common

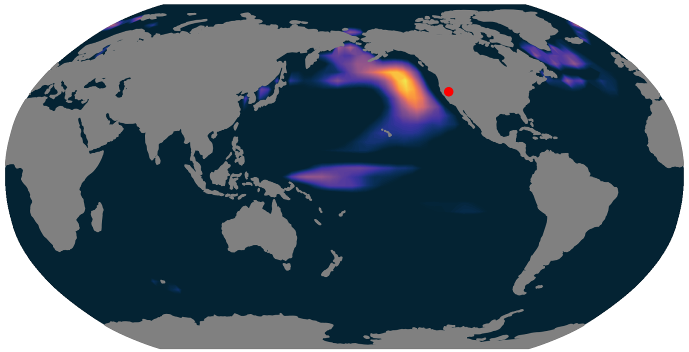

Another interesting application:

Identifying time-evolving patterns of climate change

(talk with CSGF fellow Jamin Rader for more information)

Identifying time-evolving patterns of climate change

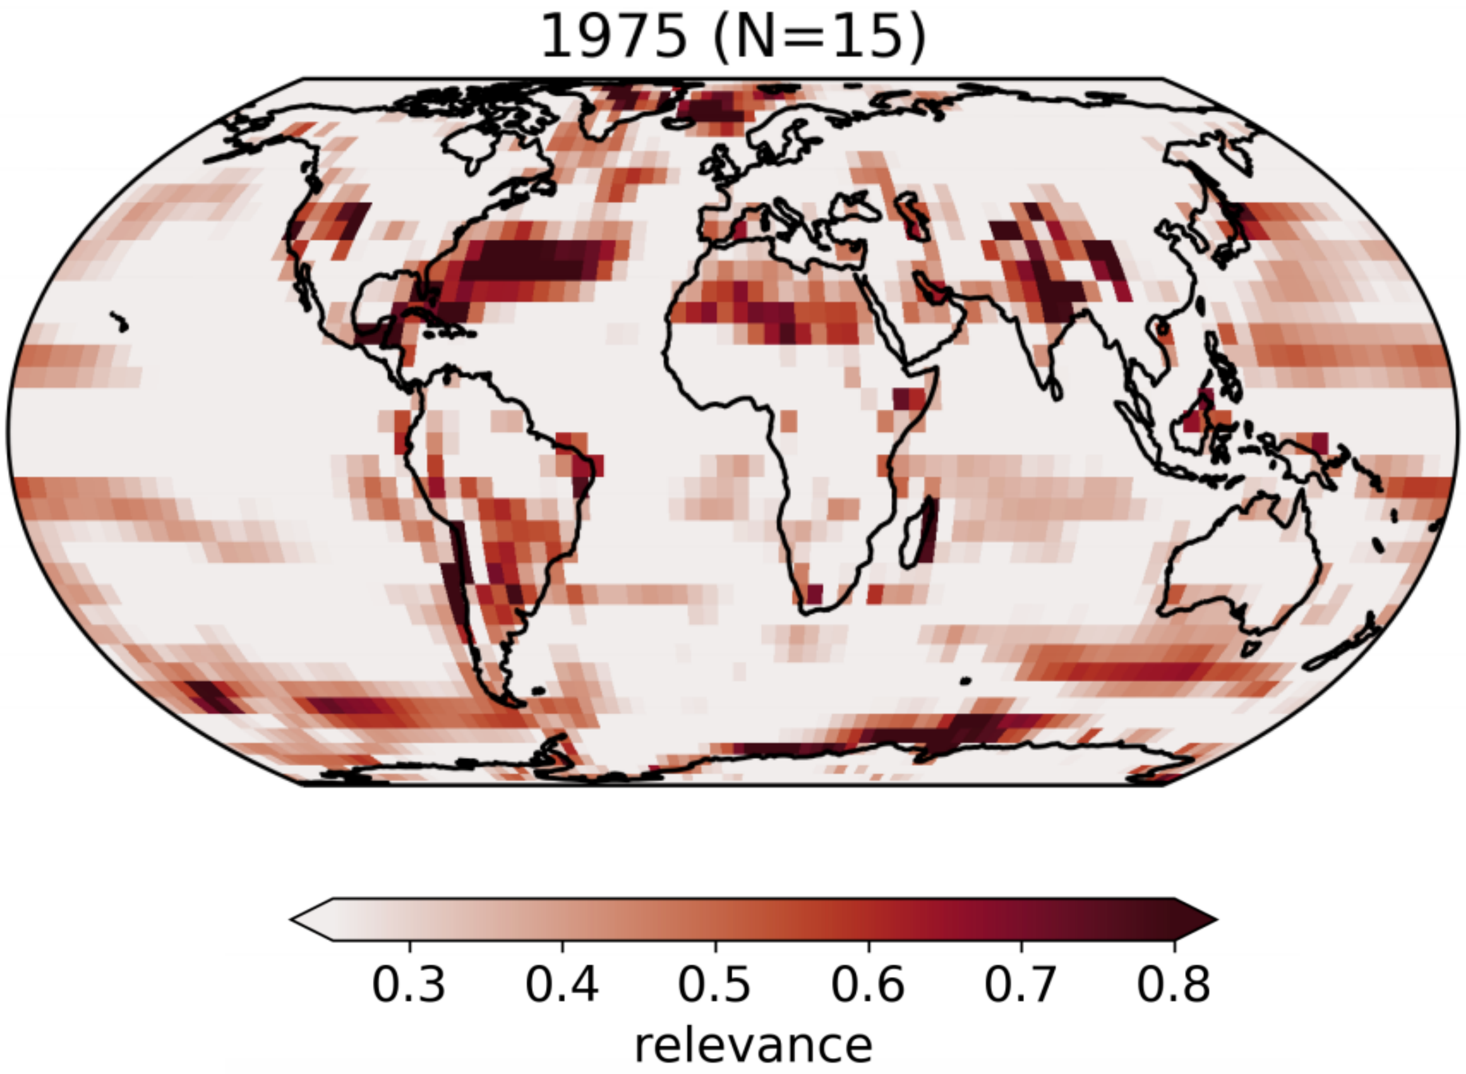

From Barnes et al. (2020); talk with CSGF fellow Jamin Rader for more information

Identifying time-evolving patterns of climate change

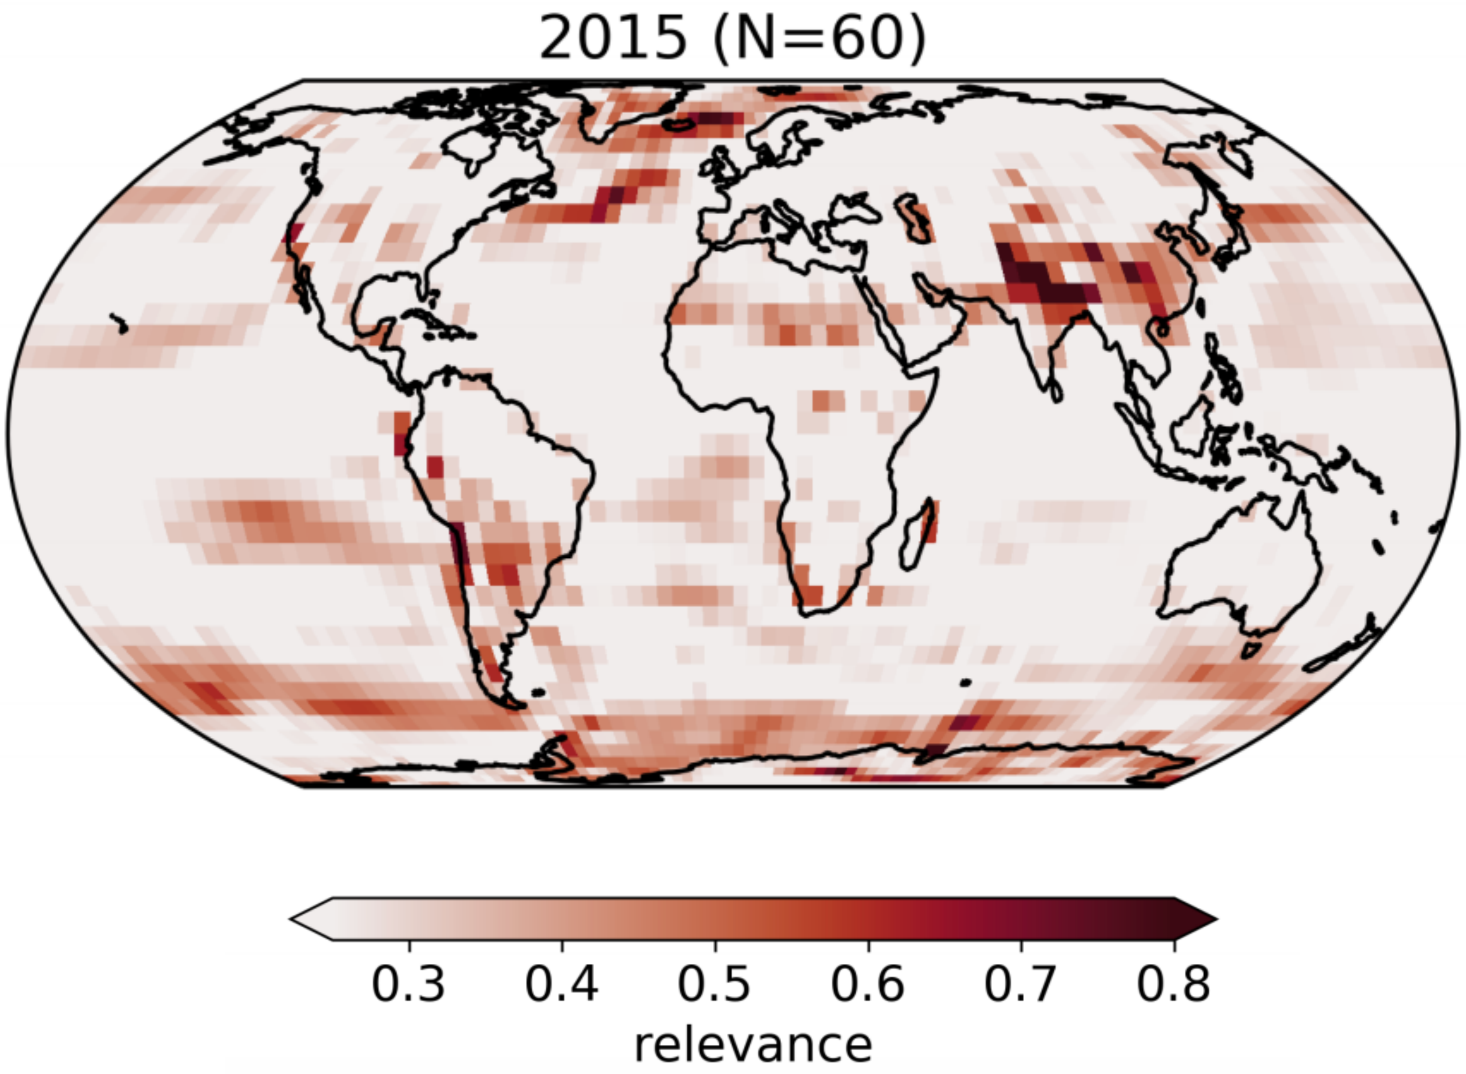

From Barnes et al. (2020); talk with CSGF fellow Jamin Rader for more information

Identifying time-evolving patterns of climate change

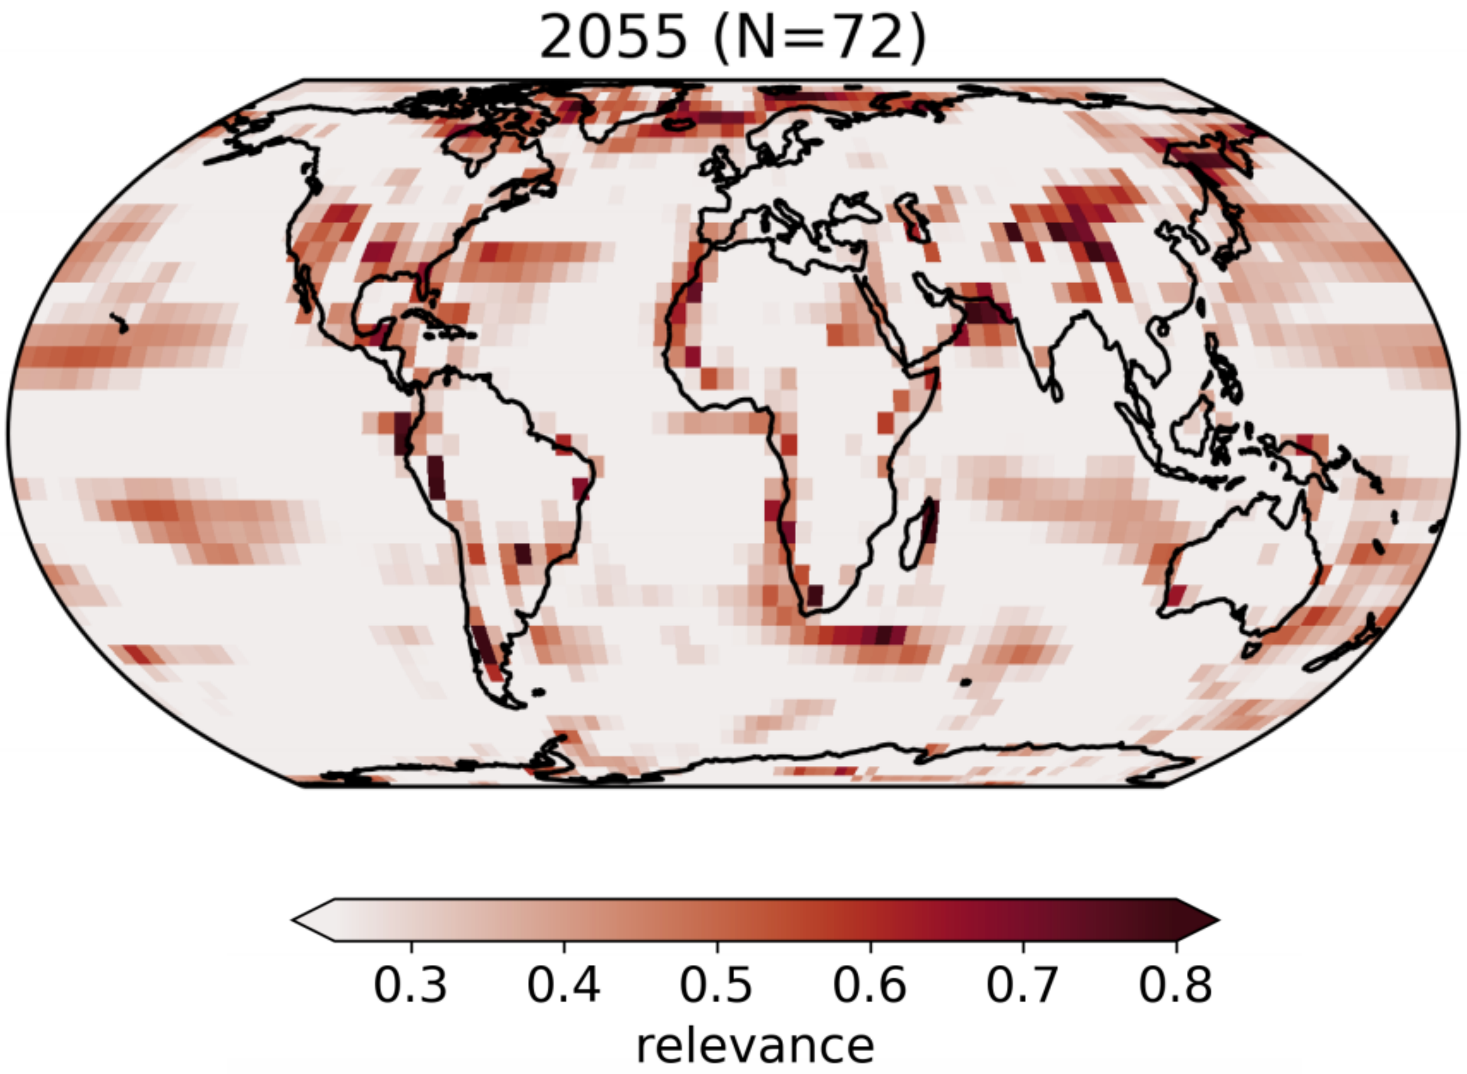

From Barnes et al. (2020); talk with CSGF fellow Jamin Rader for more information

3:37 PM
You
You
Zane Martin
Ben Toms
Zachary Labe
WeiTing Hsiao
Charlotte Connolly
Jamin Rader
Daniel Hueholt
Kirsten Mayer
Antonios Mamalakis
Emily Gordon
Group Photo