Can a neural network learn atmospheric dynamics for extreme rainfall?

Case study using Explainable AI in Western Norway

Robin Guillaume-Castel, Stefan Sobolowski, Camille Li, Joshua Dorrington

robin.guillaume-castel@uib.no

BJERKNES CENTRE for Climate Research

LEAD AI

Co-funded by the European Union

This is a neural network:

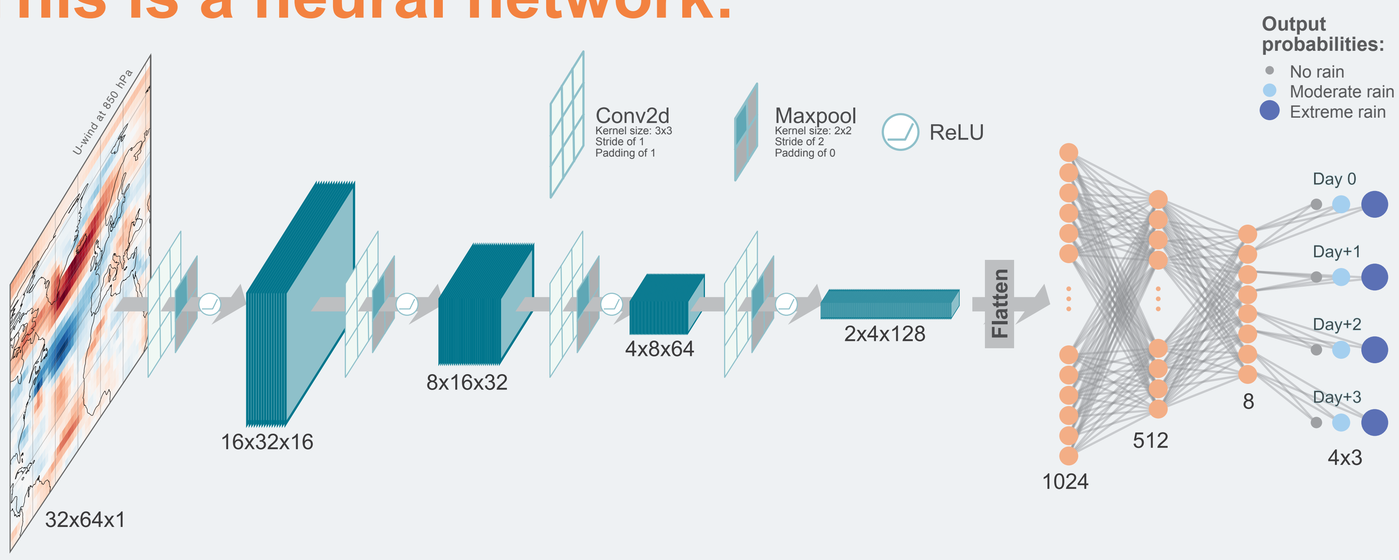

It is trained to predict the occurrence of **daily extreme rainfall** (90th percentile) in **Western Norway** over **multiple days** from maps of **atmospheric variables** (u850 in this case).

Data used

- ERA5 daily data 1959-2021
- Northern hemisphere regridded to 32x128
- Subset over the Atlantic and Europe
- 23 011 samples
- 70% train, 10% validation, 20% test

Performance

There is room for improvement in the model performance but it still does some good predictions.

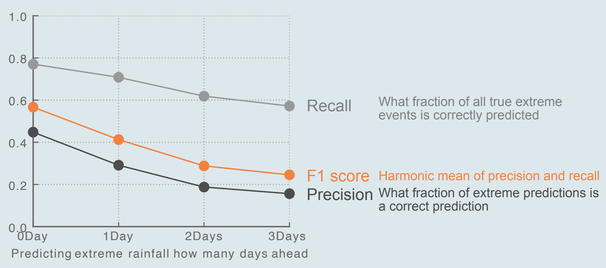

In test data: **257** existing extreme events, predicted four times by the neural network.

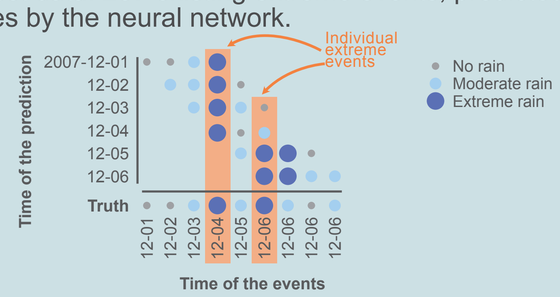

97 events were correctly predicted every time, and 3 out of 4 times for **149** events. The best predictions are in **winter** and **autumn.**

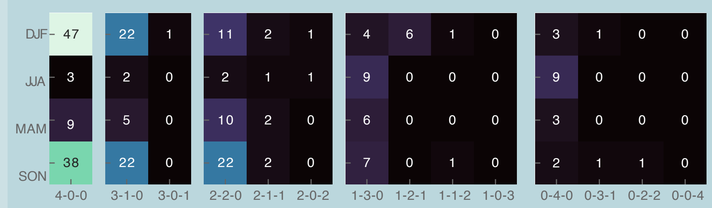

N of predictions for each extreme event: Extreme-Moderate-No rain.

Is the neural network making predictions based on atmospheric dynamics ?

Using Explainable AI - **Integrated Gradients**

Mean of DJF True positive cases

What the inputs are:

Normalized U850 input **features.**

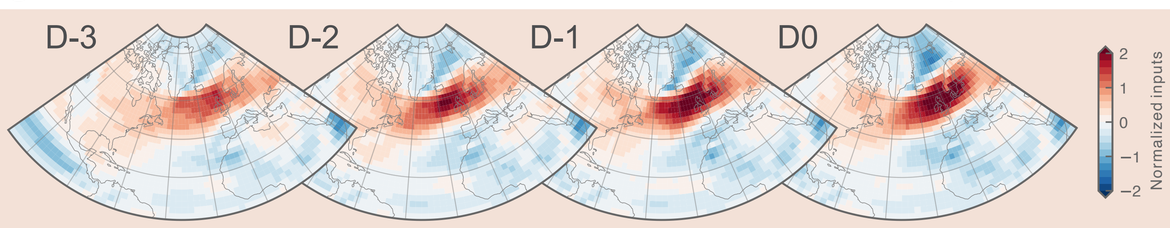

Where the model focuses on:

Sensitivity of the prediction to the different input features. This is independent of the feature amplitude.

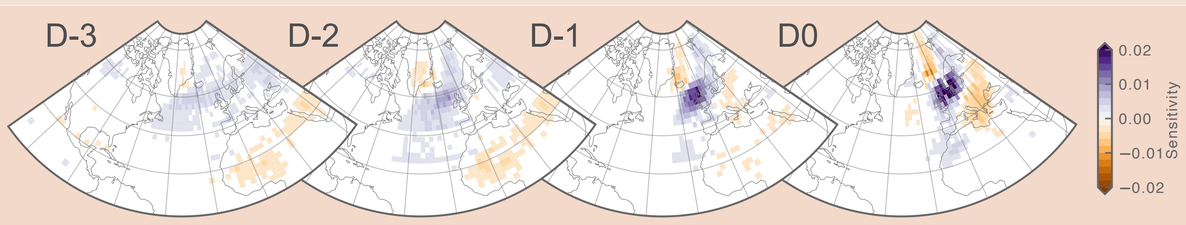

Where the model should focus on:

Dynamical precursors determined from statistical composite analysis. From Dorrington et al. (2024).

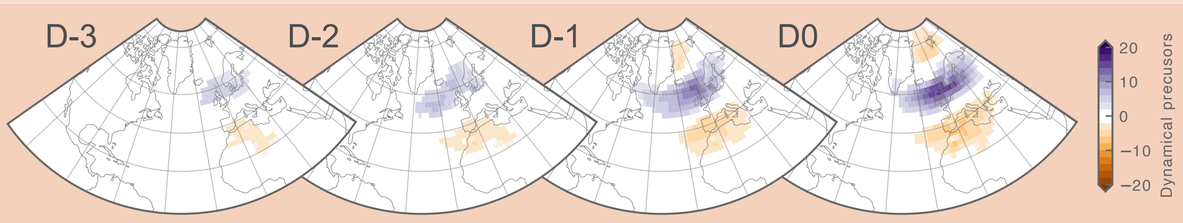

What is the prediction based on:

Attribution of each feature to how much it contributes to the prediction. This corresponds to the sensitivity times the input features. Values are relative to the maximum for better readability.

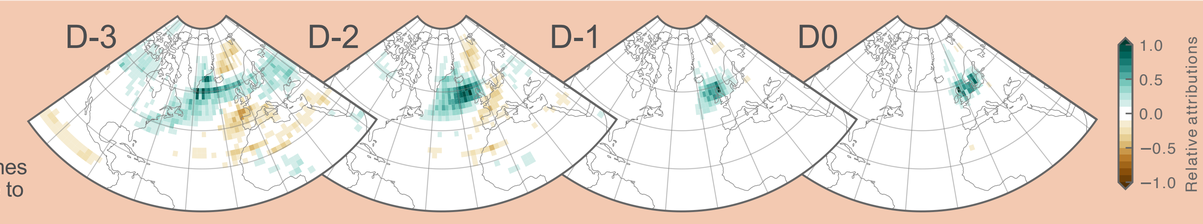

Yes (?)

The neural network seems to pick up reasonably physical areas

It can also pick up different physical cases:

Attributions for two individual extreme events:
Contours show geopotential height at 500hPa

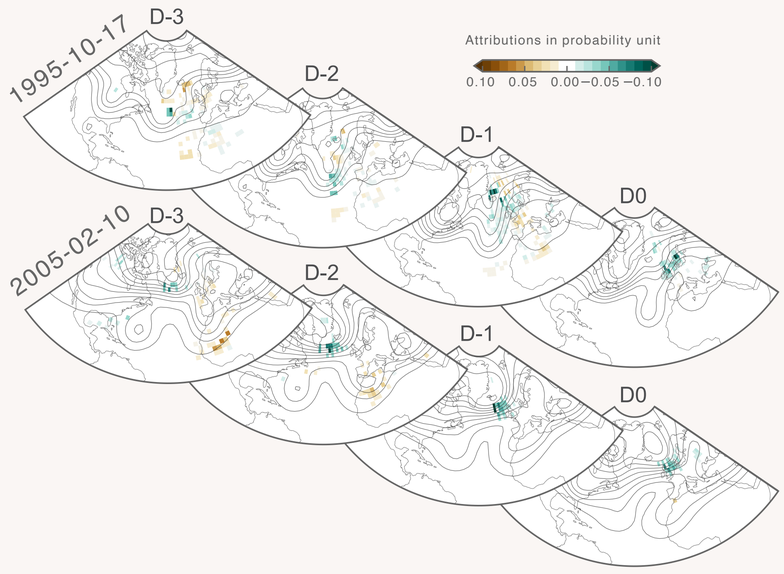

Yes, but:

The architecture of the model changes the shape of the attributions

The overall regions seem similar across architectures, but there are clear artifacts: without the max pool, we can see the kernel shapes; the attributions of the transformer shows the 4x4 patches.

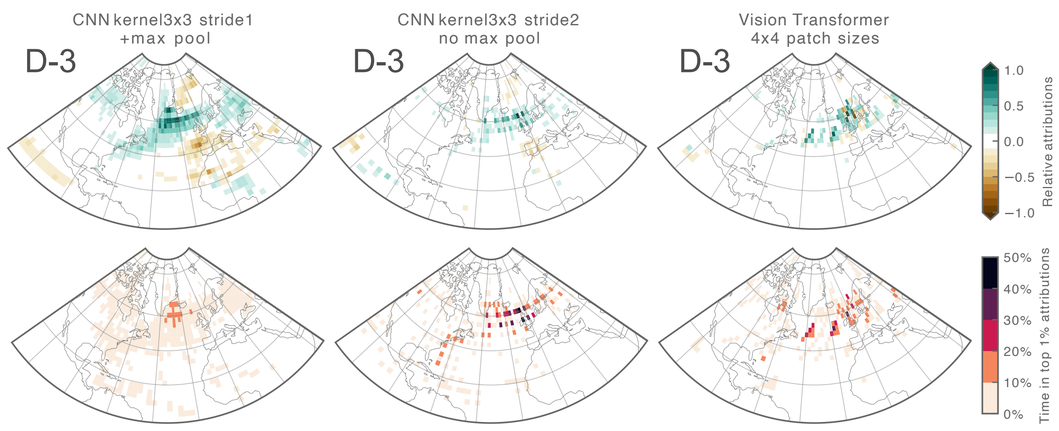

Summary

- We trained a neural network (NN) to **predict extreme rainfall** in Western Norway.
- Using explainable AI, we showed that the NN seems to **make decisions based on physically relevant regions.**
- However, our results are **very dependent on the architecture** of the NN, which means further interpretability requires caution.

Allégaten 55, NO-5007 Bergen | Tel: +47 55 58 98 03 | Fax: +47 55 58 43 30 | post@bjerknes.uib.no | bjerknessenteret.no

UNIVERSITY OF BERGEN
NORCE
NERSC Nansen Center, Bergen
HAVFORSKNINGSINSTITUTTET INSTITUTE OF MARINE RESEARCH