#ICML2023

Machine Learning with **Social Purpose**

Shakir Mohamed

OUR PLAN

ML & EARTH SYSTEMS

Generative models and research in the Earth Systems create all the pathways to impact we want.

SOCIOTECHNICAL AI

Research in technology and society be done together: so deepen a sociotechnical AI research portfolio.

GLOBAL AI

A more global field and industry can shift machine learning to be more general.

01 | WORLDS IN THE MAKING

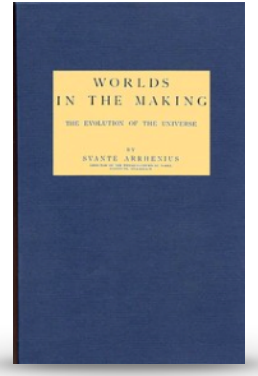

Title from Svante Arhenius's 1908 book where the problem of global heating was first investigated. Cover: Project Gutenberg

Earth System Models

The Earth System is a representation of all the processes on our planet.

Typically we can represent the earth system using the physical equations that represent the state of the world.

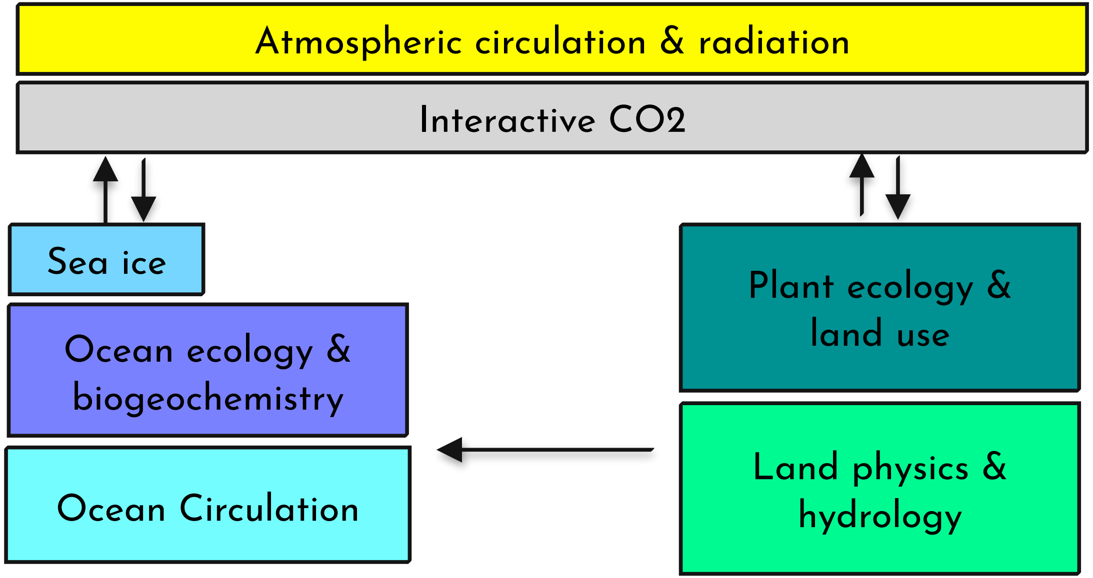

Shakir Mohamed

Forecasting Timescales

NOWCASTING

High-resolution (100m-1km) predictions of variables up to 2 hours ahead.

MEDIUM-RANGE

Global forecasts (10-80km) of atmospheric state up to 10 days head.

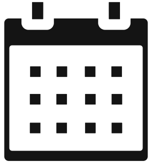

SUBSEASONAL TO SEASONAL

1-3 months

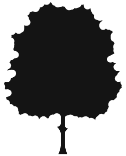

DECADAL

10 years

CLIMATE

100 years

Nowcasting

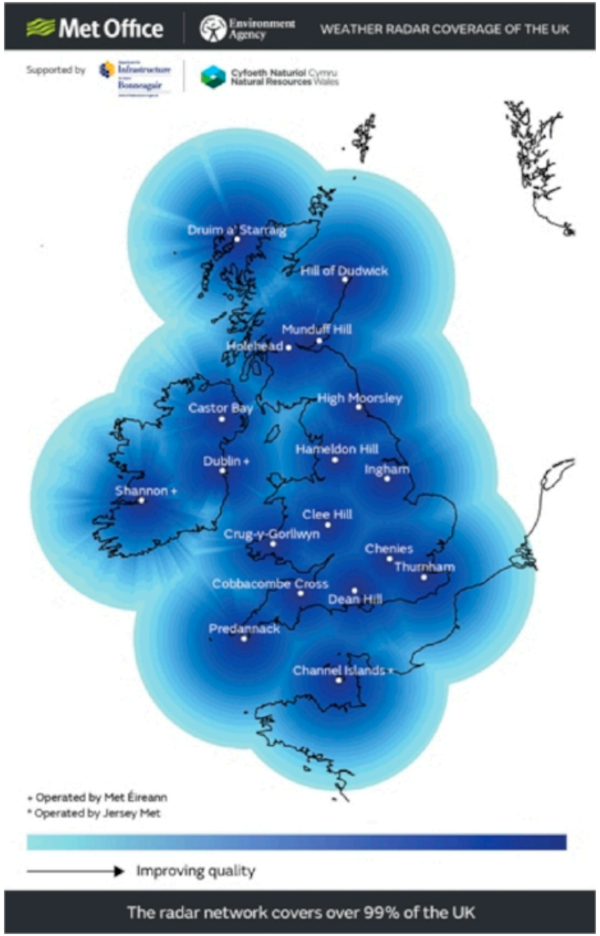

Shakir Mohamed

UK Radar Data

- Met Office RadarNet4 Data
- Every 5 minutes, 288/day
- 1536 x 1280 pixels at 1km x 1km grids
- Data from 2016-2020

Focus on heavy rain that is of specific interest to operational meteorologists: more rare, but where protection of life and property is high priority.

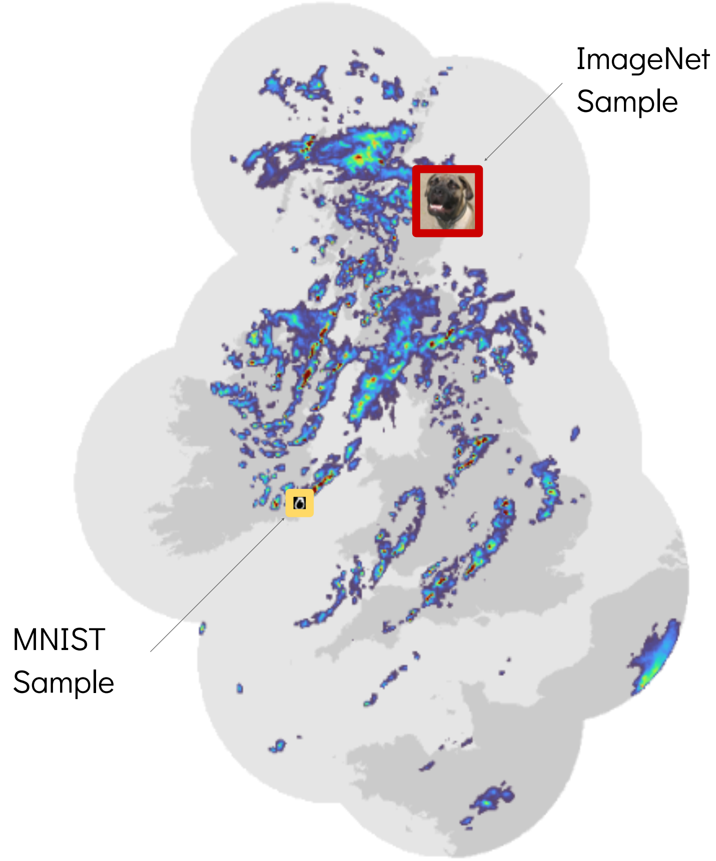

Deep Generative Models of Rain

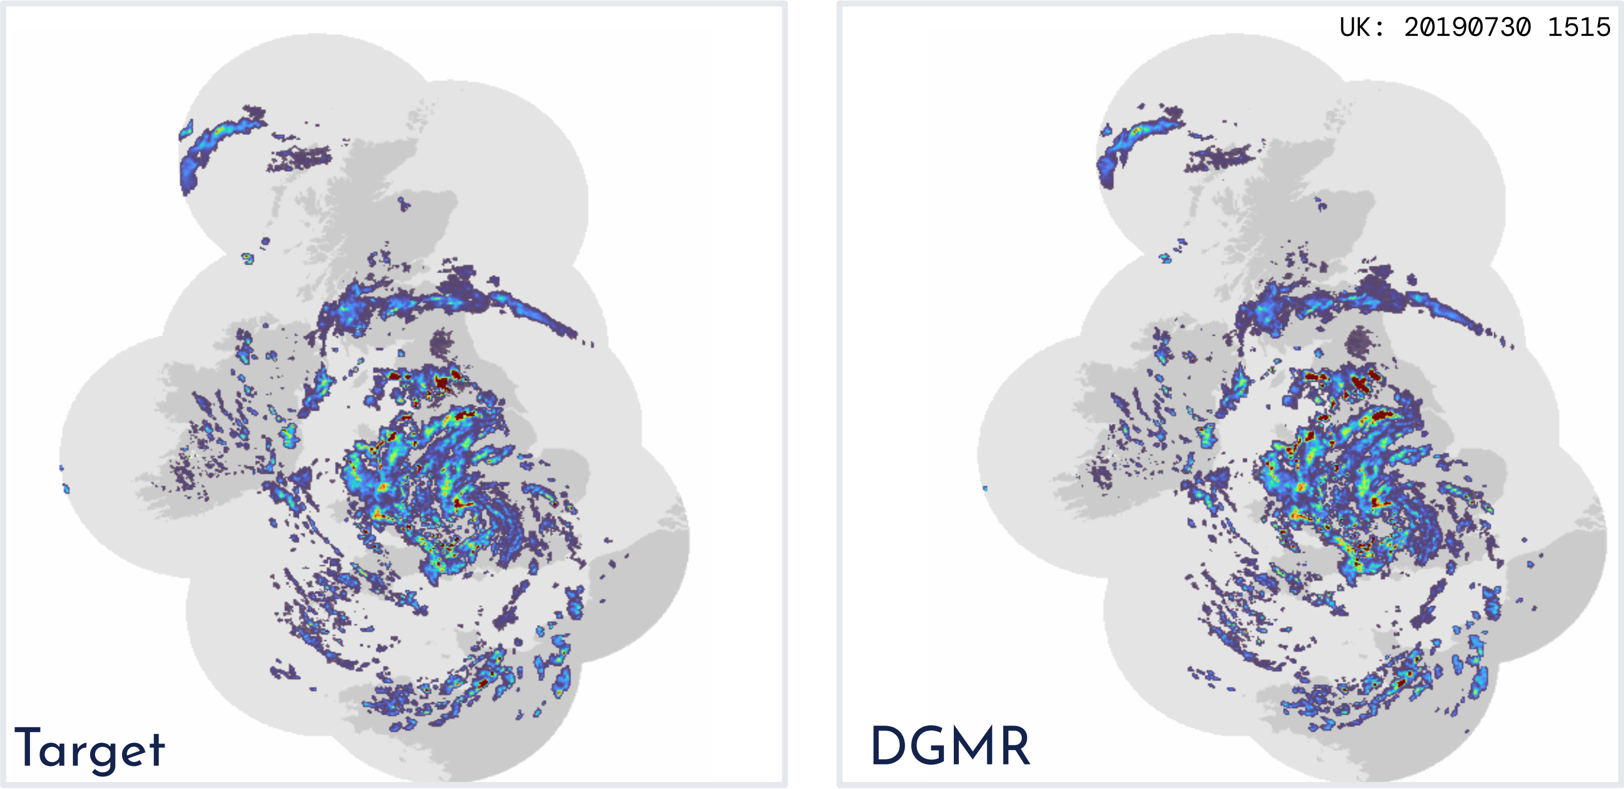

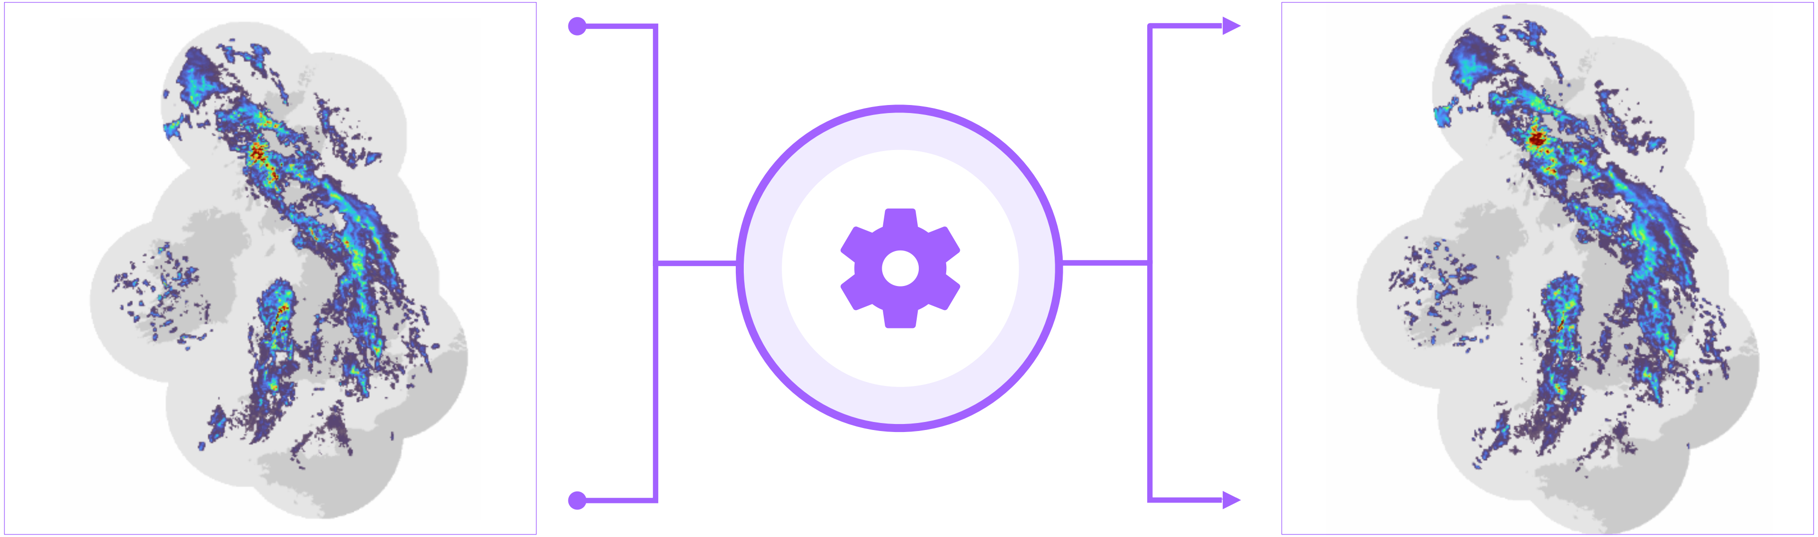

Context
Past 20mins

Deep Generative
Model of Rain

Nowcast
Next 90mins

Paper: Skillful precipitation nowcasting using deep generative models of radar.

24/06/2019 at 16:15

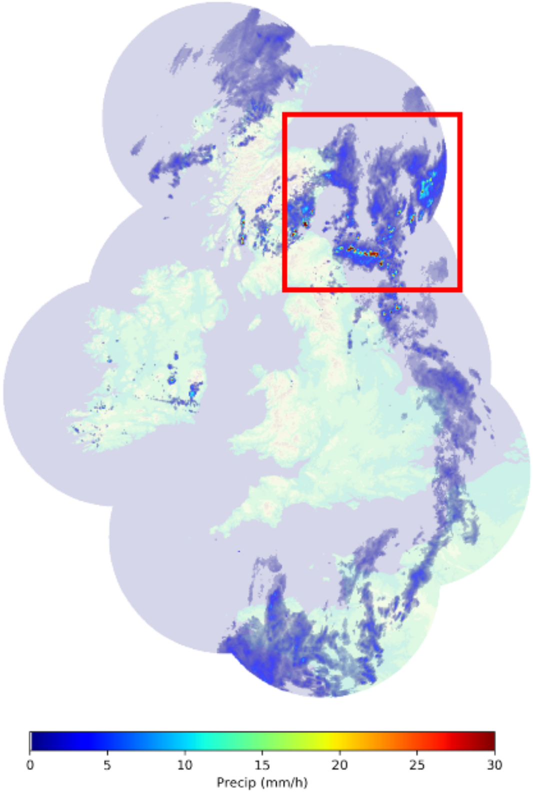

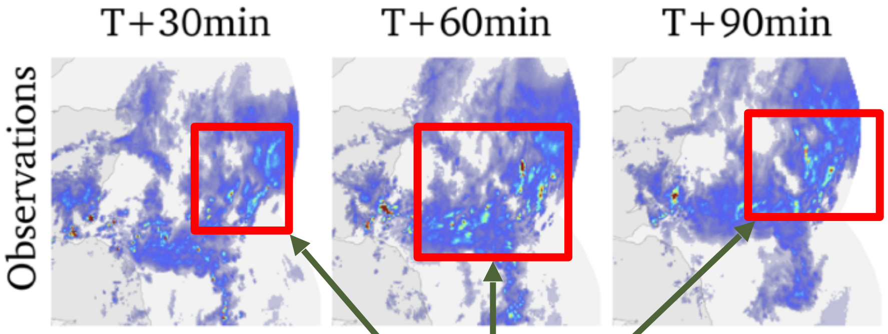

Important (and difficult) to predict convective cells

Difficult case chosen by the Chief Forecaster who is independent of the project team.

24/06/2019 at 16:15

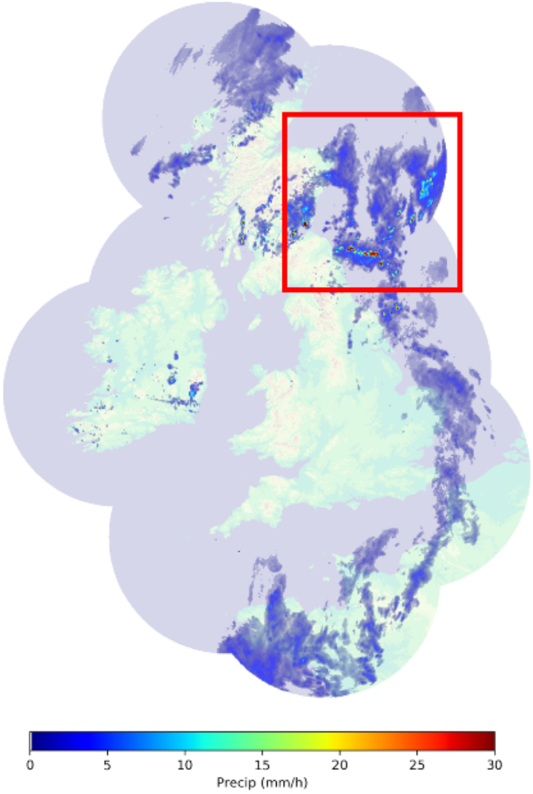

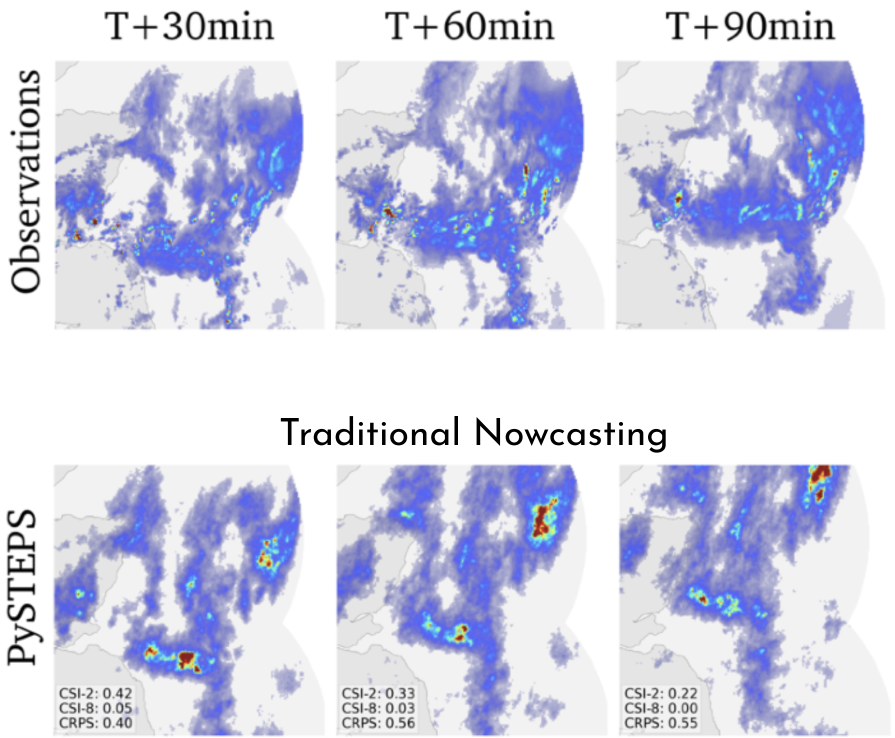

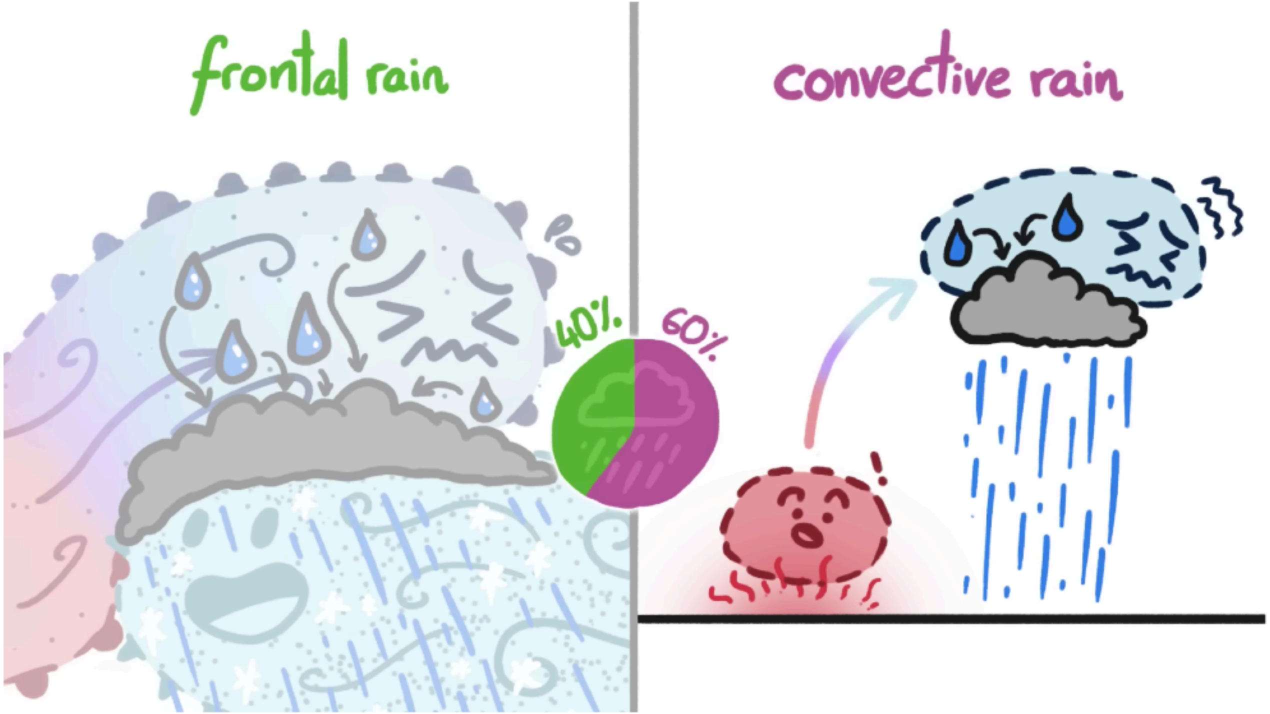

Why Weather Forecasts Suck

MinuteEarth
2.79M subscribers
Join
Subscribe
12K
Share
Download

https://youtu.be/snCo0Z0dt-k

Shakir Mohamed

24/06/2019 at 16:15

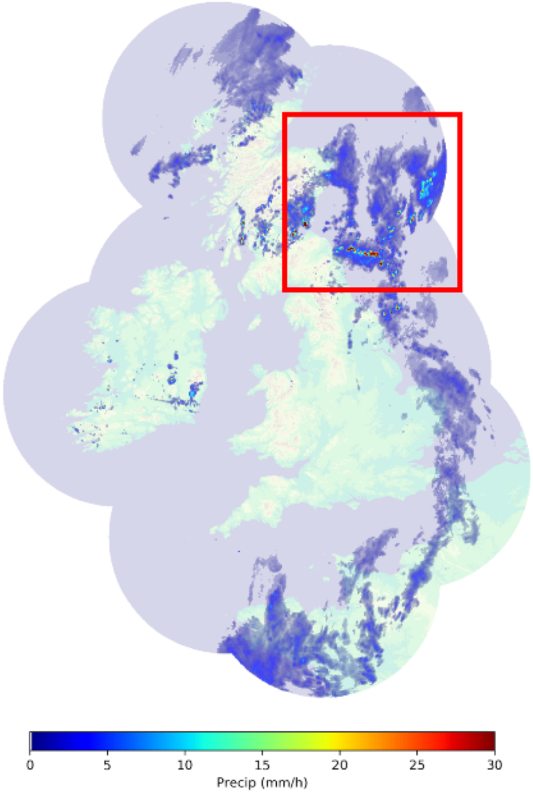

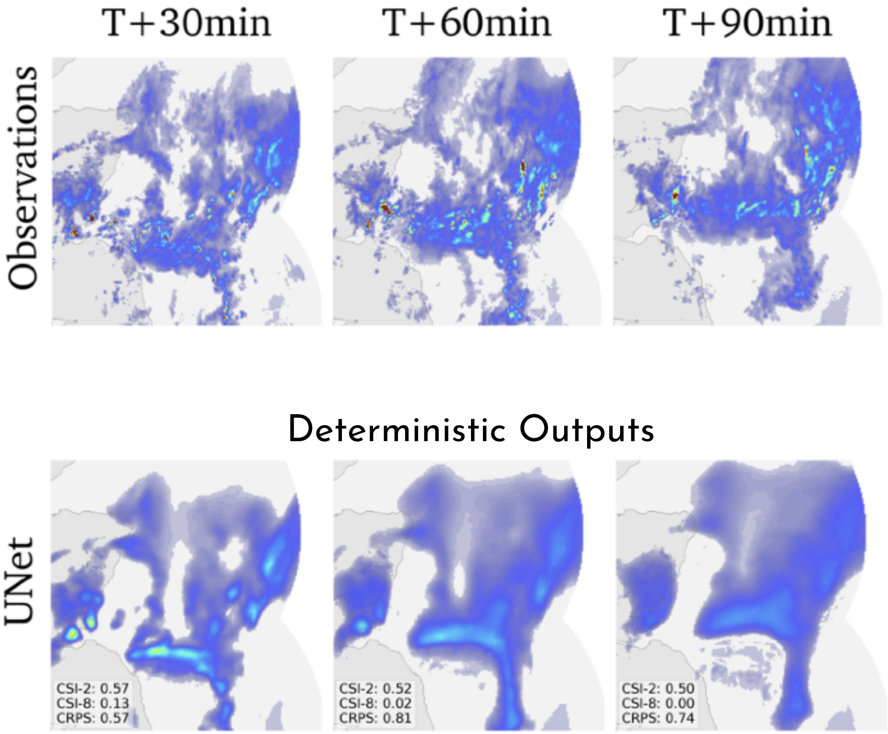

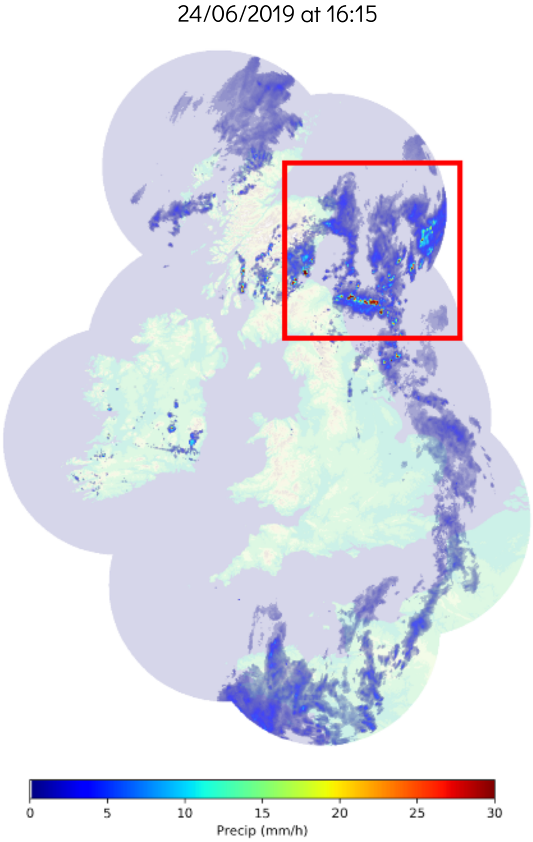

24/06/2019 at 16:15
0
5
10
15
20
25
30
Precip (mm/h)

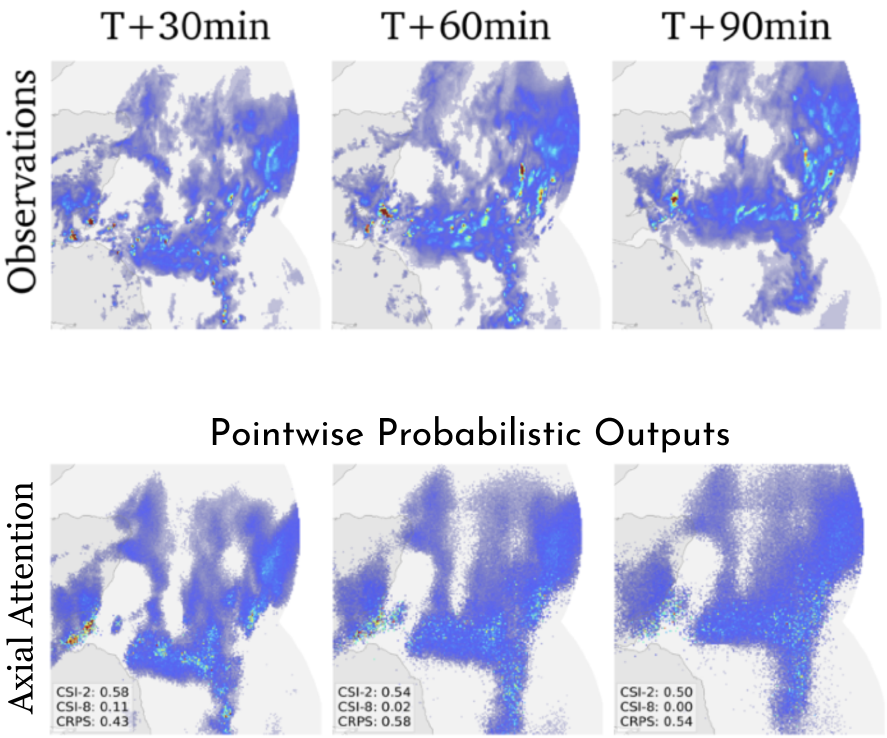

T+30min
T+60min
T+90min
Observations
Pointwise Probabilistic Outputs
Axial Attention
CSI-2: 0.58
CSI-8: 0.11
CRPS: 0.43
CSI-2: 0.54
CSI-8: 0.02
CRPS: 0.58
CSI-2: 0.50
CSI-8: 0.00
CRPS: 0.54

24/06/2019 at 16:15

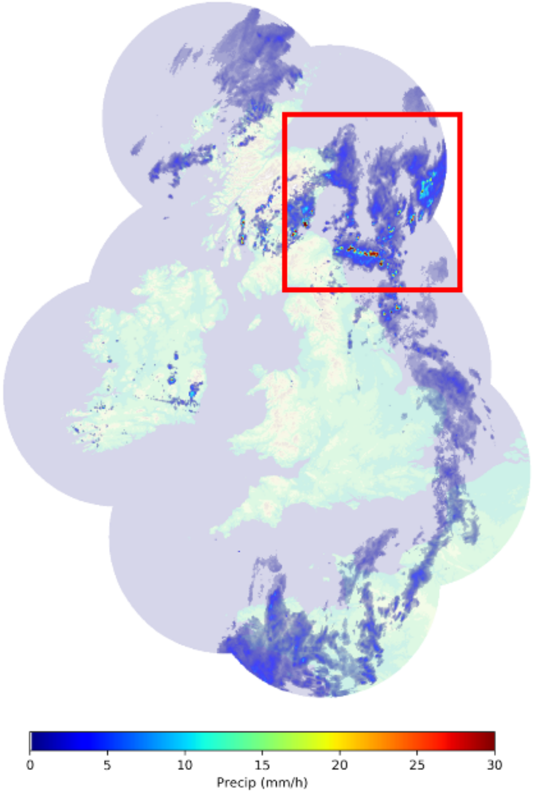

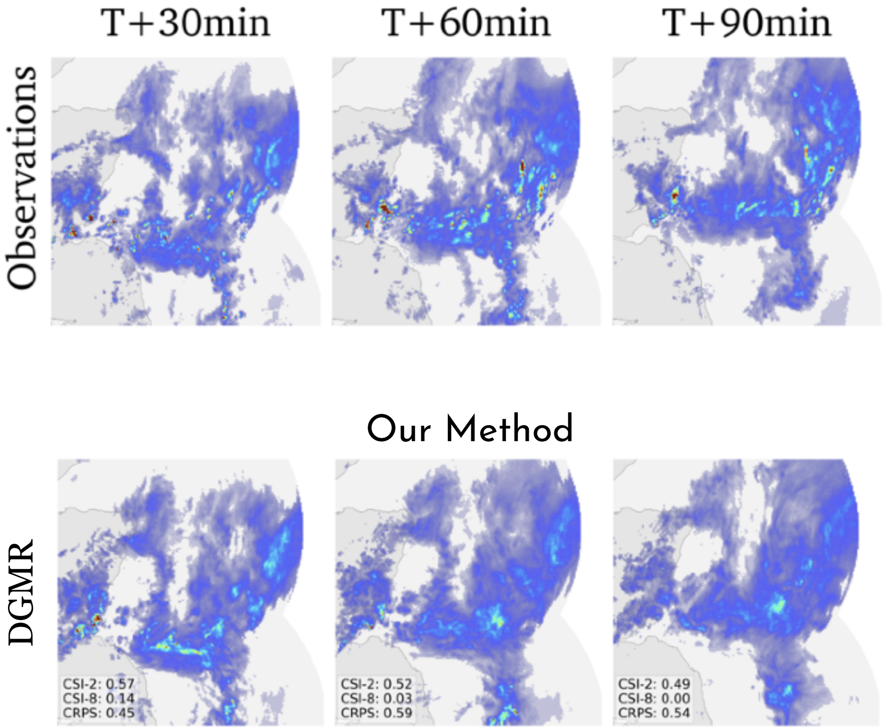

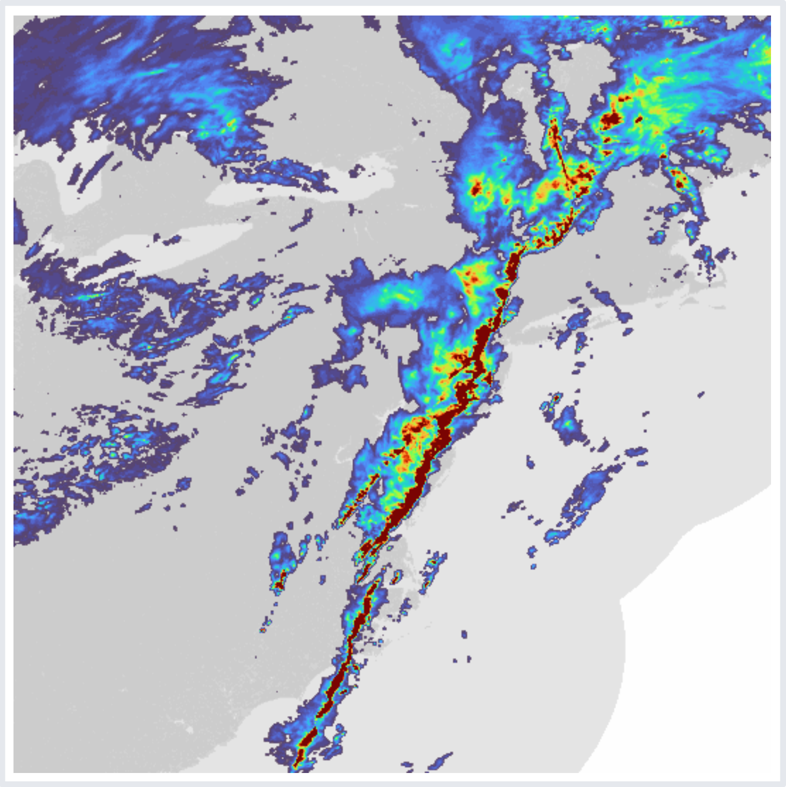

Target

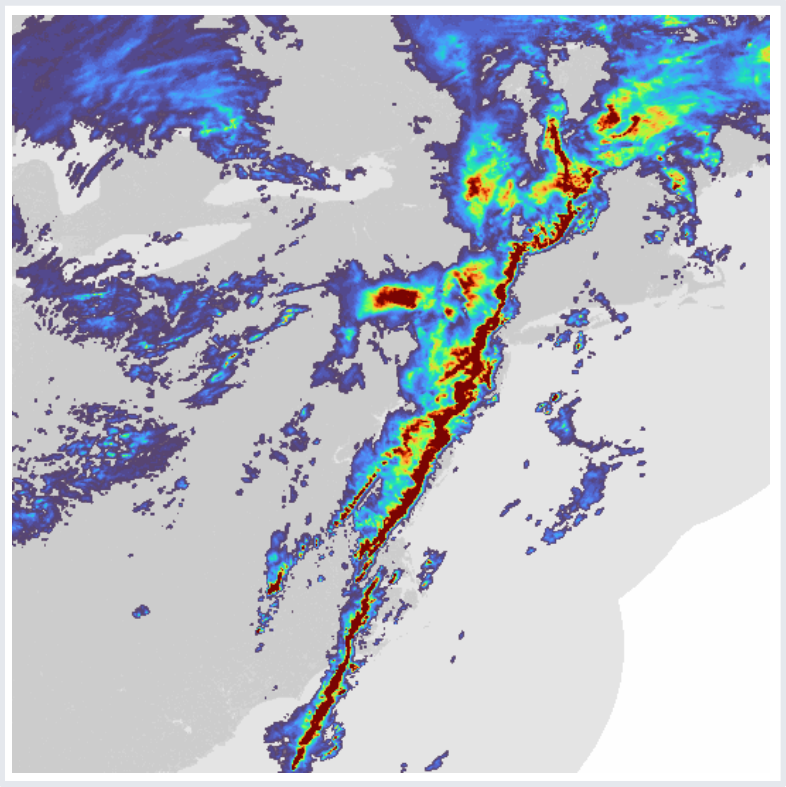

DGMR

US-East: 20190415 0930

Intercomparison

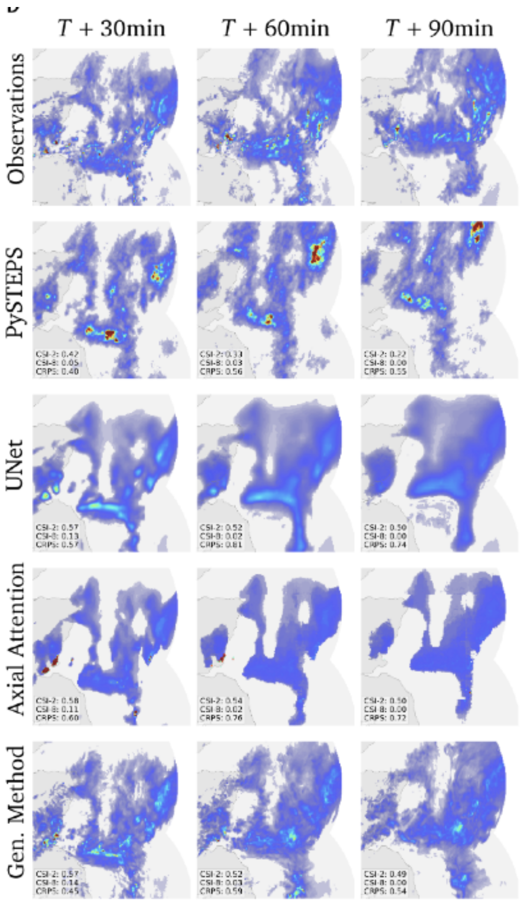

Postage stamp plots to assess sample spread and uncertainty

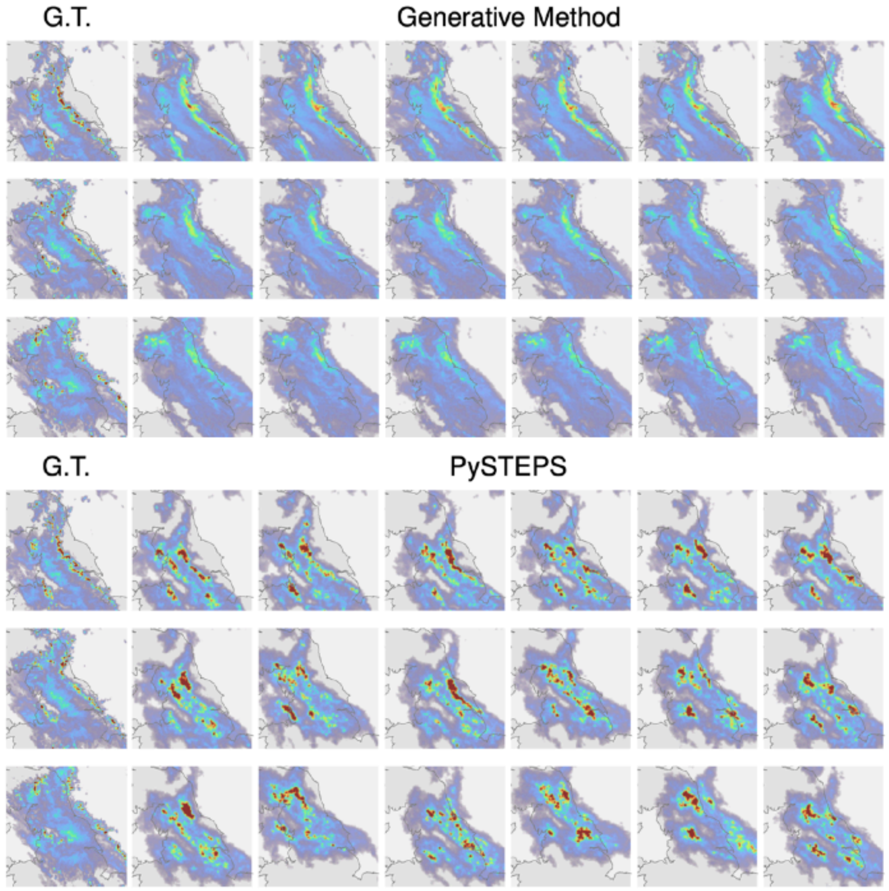

Critical Success Index (CSI)

CSI allows us to measure location accuracy of the forecast at various rain rates, and is a single summary of binary classification performance that rewards both precision and recall.

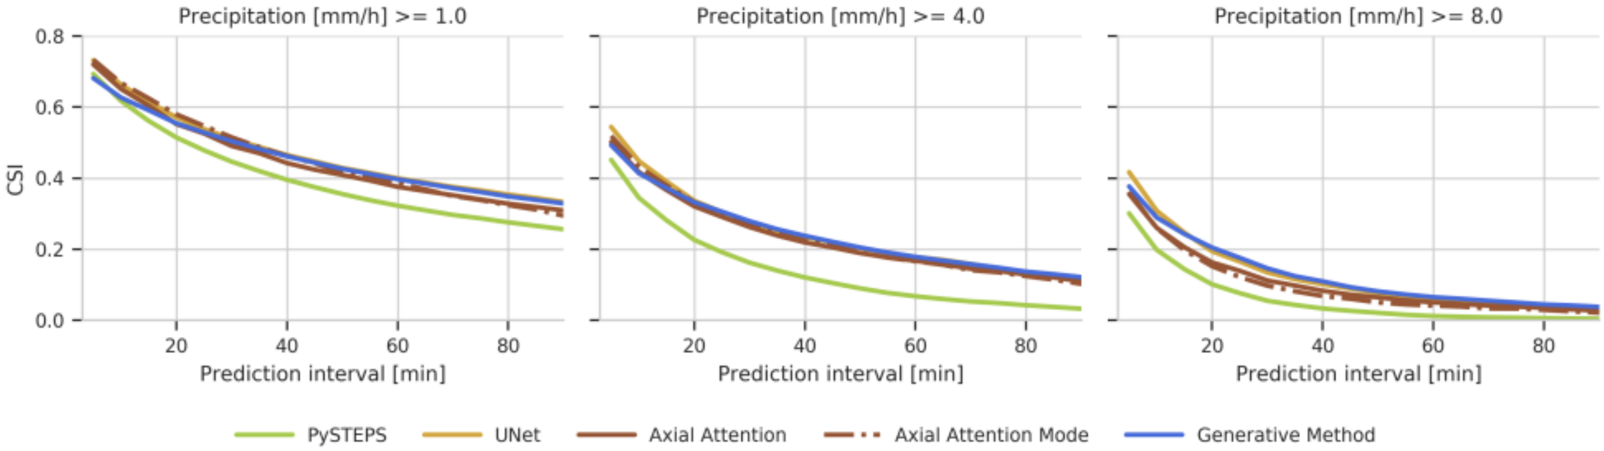

$$CSI = \frac{TP}{TP+FP+FN} = \frac{f_1}{2-f_1}$$

CSI doesn't account for all the ways a model can make predictions or can 'cheat' in making predictions (e.g., by blurring).

Power Spectral Density (PSD)

PSD measures how power is distributed across a range of spatial frequencies in each model's forecasts and to compare the spectrum with observed data.

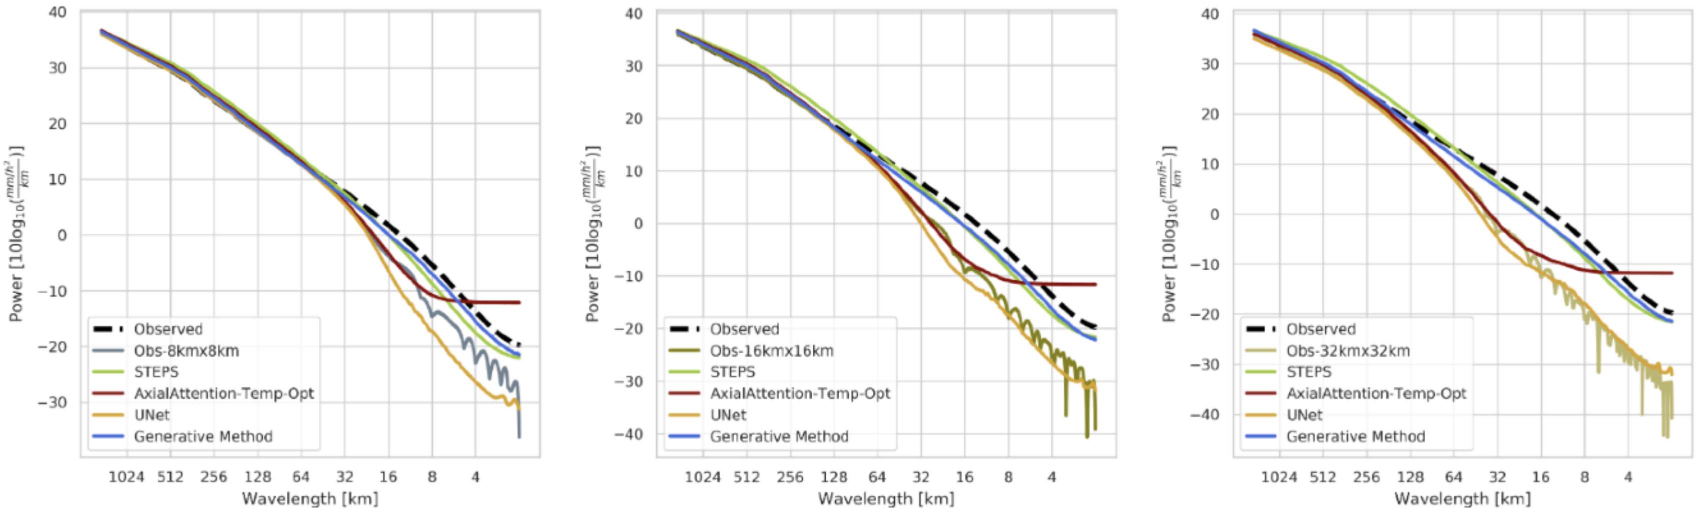

After 30mins, other models make predictions at a resolution on 8x8km, whereas the generative approach maintains predictions at the resolution of the data.

Continuous Ranked Probability Score (CRPS)

CRPS is a proper scoring rule for univariate distributions, which we use to score the per-grid-cell marginals of a model's predictive distribution against observations.

$$\mathbb{E}[F - Obs] - \frac{1}{2}\mathbb{E}[F - F']$$

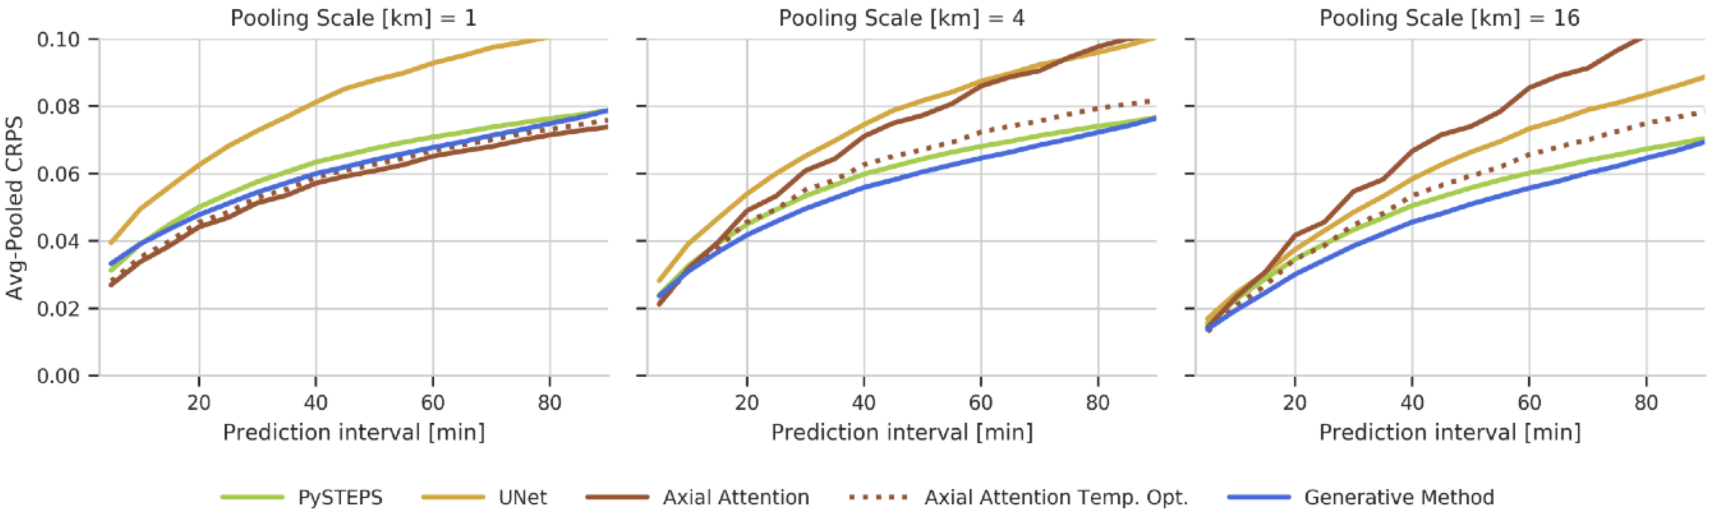

Also show pooled versions, which are scores on neighbourhood aggregations that show whether a prediction is consistent across spatial scales.

Relative Economic Value

Shows the relative economic value in the decision to take, or not to take, a precautionary action in response to different rainfall thresholds.

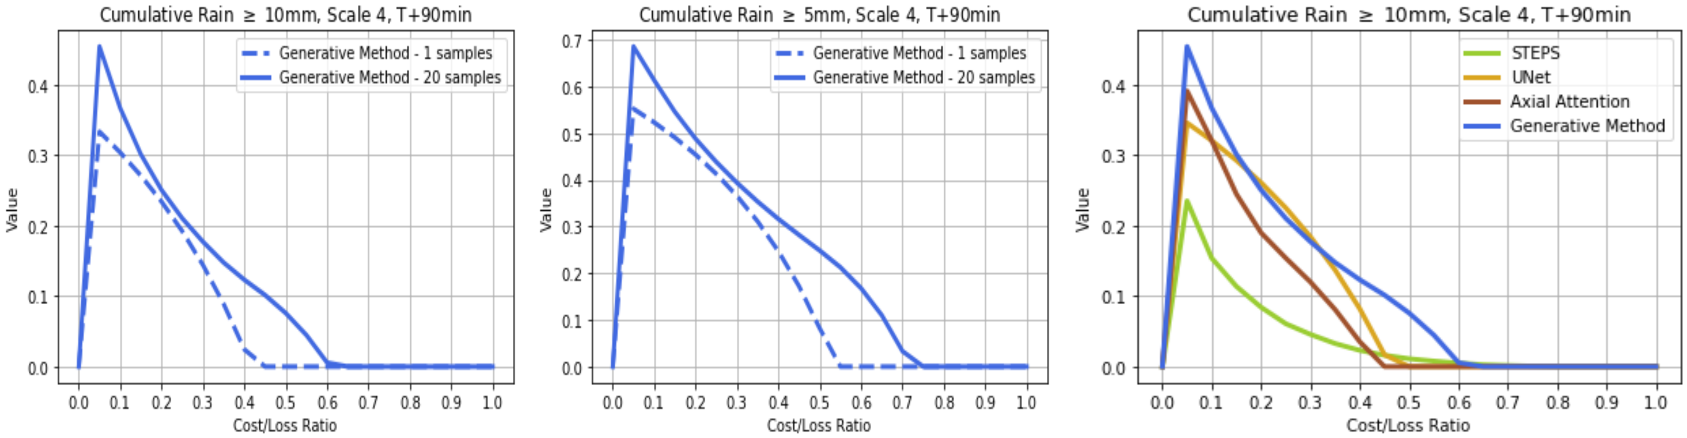

This evaluation uses a cost-loss decision model. If we take precautions we incur a fixed cost C; if we don't and a weather event occurs, we incur a loss L. We can compute the value which is the ratio of expenses for the forecast versus a perfect forecast. Value is only a function of C/L.

Expert Judgements

Worked with expert meteorologists who work in the 24/7 operational forecasting centre. Developed a two stage assessment to understand quality and value.

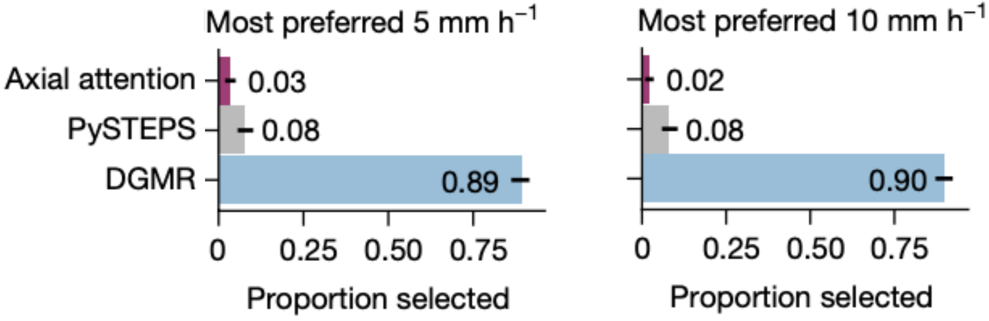

Participant Comments

- “I like things to look slightly realistic even if they’re not in the right place so that I can put some of my own physics knowledge into it.”
- “I would prefer the model to underdo intensities but get a much better spatial variation”
- “This looks much higher detail compared to what we’re used to at the moment. I’ve been really impressed with the shapes compared with reality. I think they’re probably better than what we’re currently using. The shapes in particular, some of them do look really high resolution”

Forecast Quality, Consistency & Value

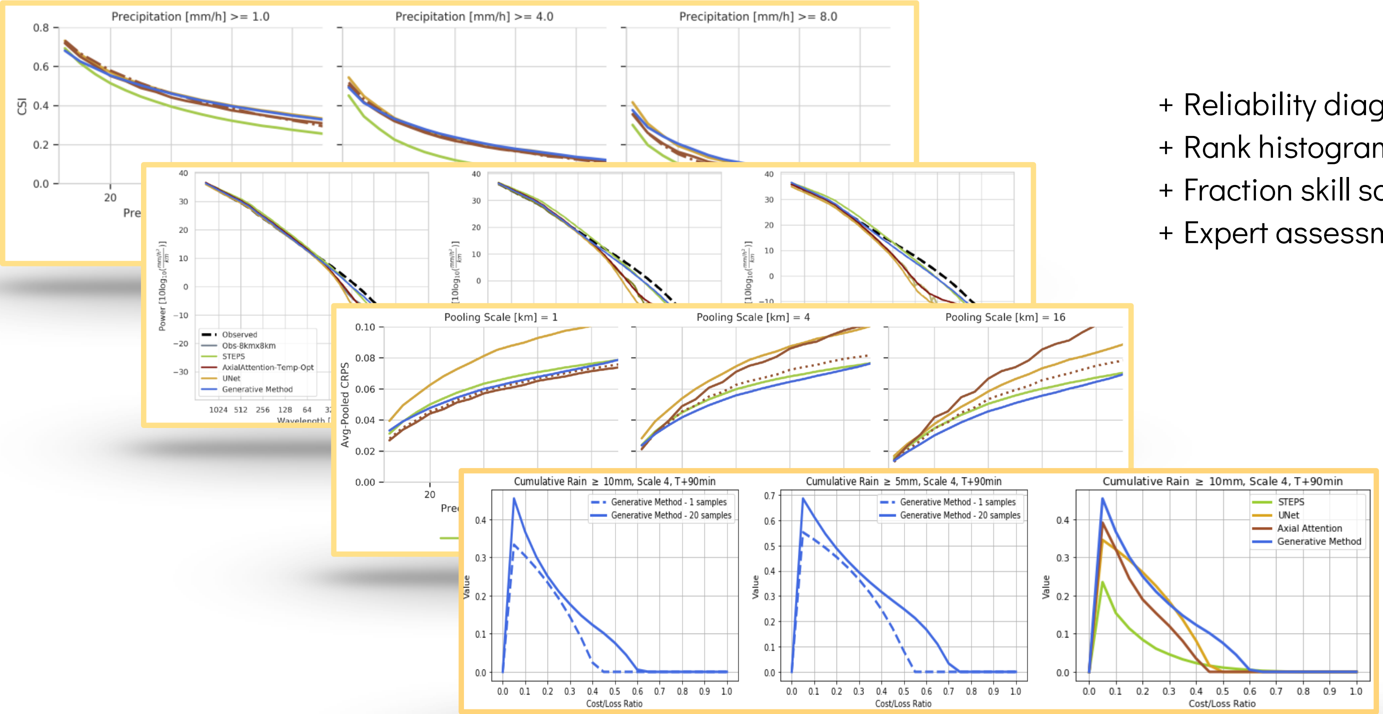

+ Reliability diagram
+ Rank histogram
+ Fraction skill score
+ Expert assessments

Paper: Skillful precipitation nowcasting using deep generative models of radar.

Global Forecasting

Shakir Mohamed

Medium-Range Weather Forecasting

Task: Predict atmospheric state at 6 hour intervals for the next 10 days, at high resolution.

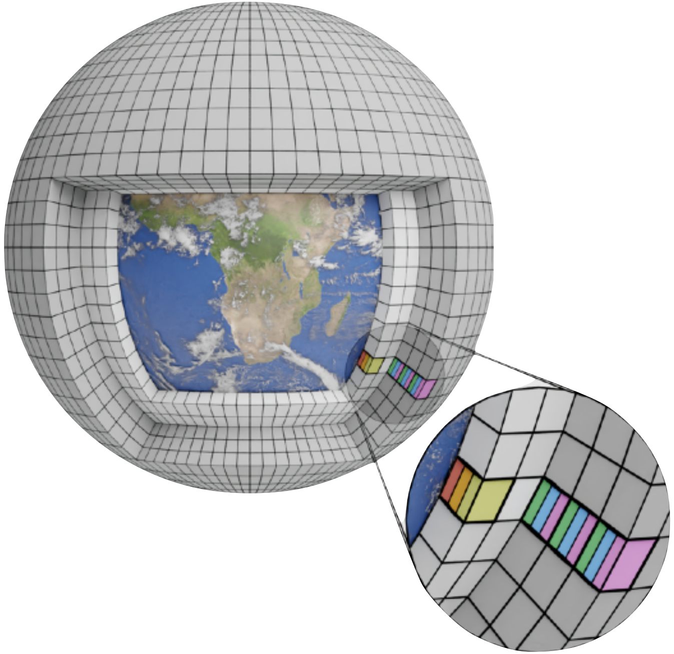

A location (or "pixel") in this grid is a column that contains:

- 5 **surface** variables (incl. 2m temperature, 10m winds, precipitation, sea-level pressure)
- 6 **atmospheric** variables each at 37 **vertical** pressure levels (incl. geopotential, temperature, wind, humidity)
- 227 variables per grid point and a total of 235,000 targets at any point.

ECMWF Forecast Data

ECMWF produce several data sets. Two we use:

- HRES: A **deterministic forecast called HRES**, and an ensemble forecast called ENS.
- ERA 5: A r**eanalysis dataset** called ERA5 from 1979-present at 0.25deg resolution.

Atmospheric **reanalysis** combines past weather observations from a range of sources with model information to provide a **complete and consistent record** of meteorological conditions.

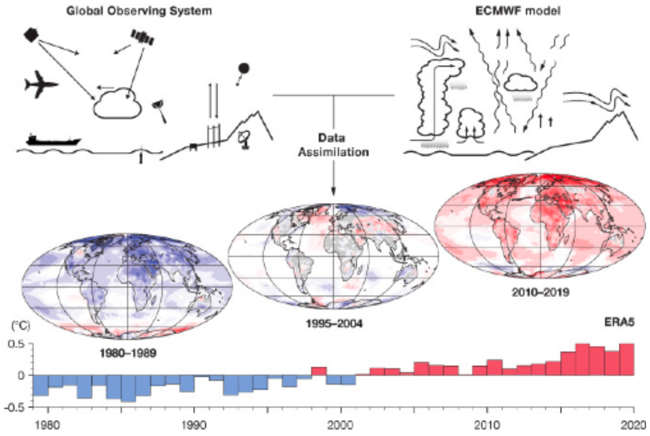

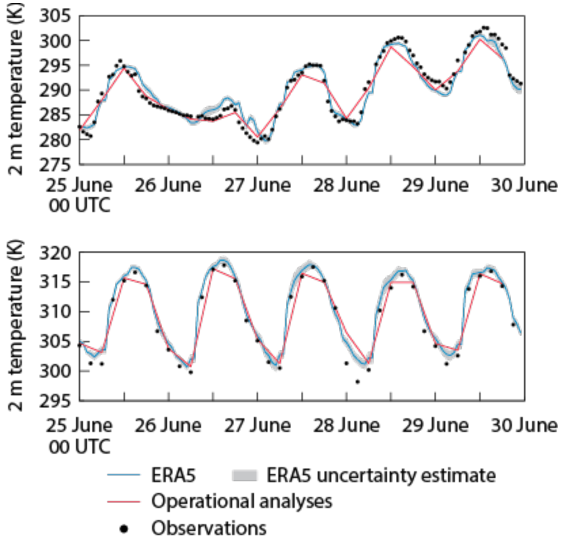

www.ecmwf.int/en/newsletter/147/news/era5-reanalysis-production

Performance Scorecard

Scorecards summarise differences in performance between models across multiple attributes.

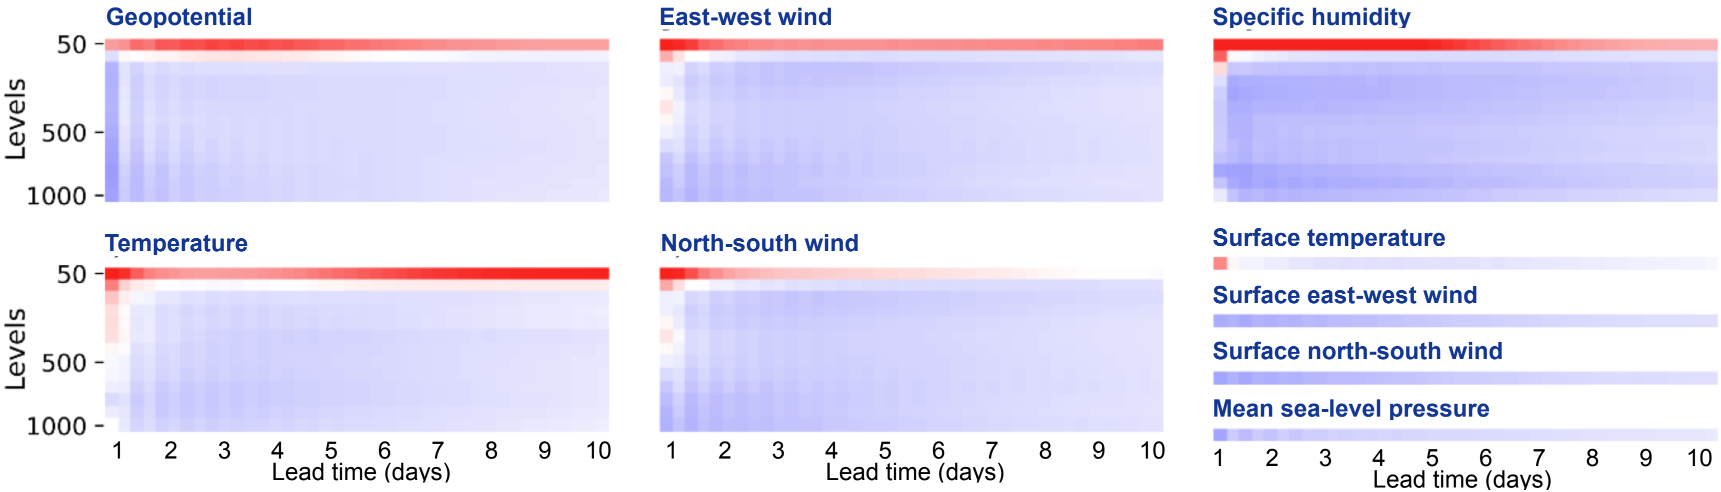

Blue: **GraphCast** better, Red: **HRES** better

GraphCast is better on 90% of 2760 targets.

Paper: GraphCast: Learning skillful medium-range global weather forecasting

Machine Learning Weather Predictions

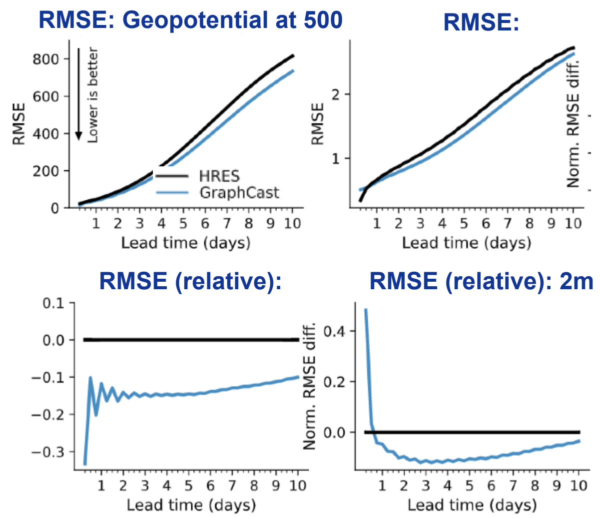

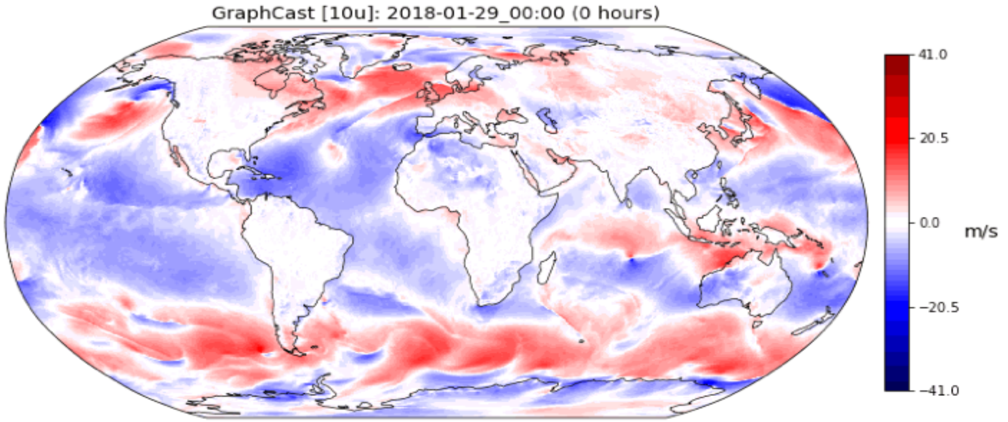

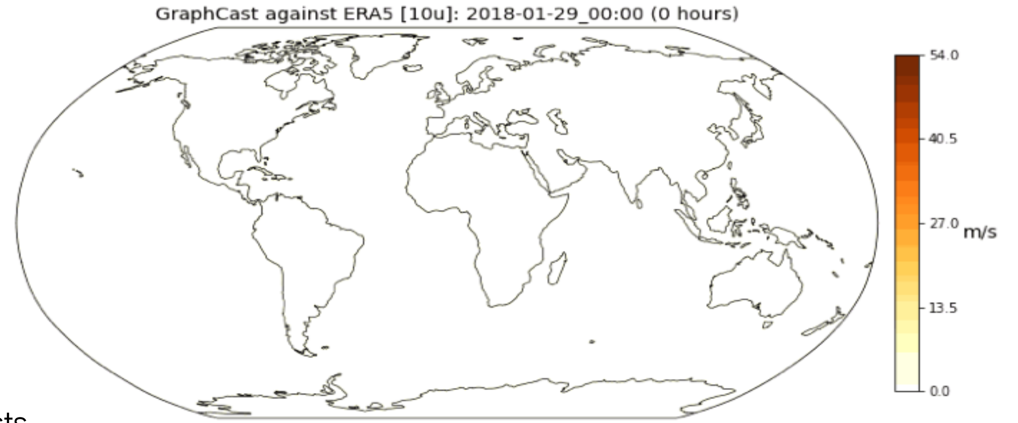

ML predictions we can outperform operational forecasts.

Paper: GraphCast: Learning skillful medium-range global weather forecasting

More to do

- **Better precipitation forecasts**, which is a key priority since most models are weak here.
- **Seasonal and decadal forecasting**, which have been so poor historically that there are few applications.
- **Downscaling** that allows higher-resolution processes to be represented from lower resolution simulation.
- Medium-range **forecasting in renewables** and net-zero.
- Causal methods for **detection and attribution** of significant events.
- Addressing **biases** in climate models, improving **uncertainty**, data sharing and distributed **computation**

02 | DREAMSCAPES OF MODERNITY

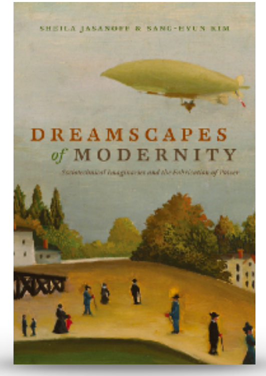

Title from Sheila Jasanoff and Sang-Hyun Kim's 2015 book that helps develop a sophisticated understanding of research and the politics of science and technology. Cover: Univ. Chicago Press

Equity in Forecasting

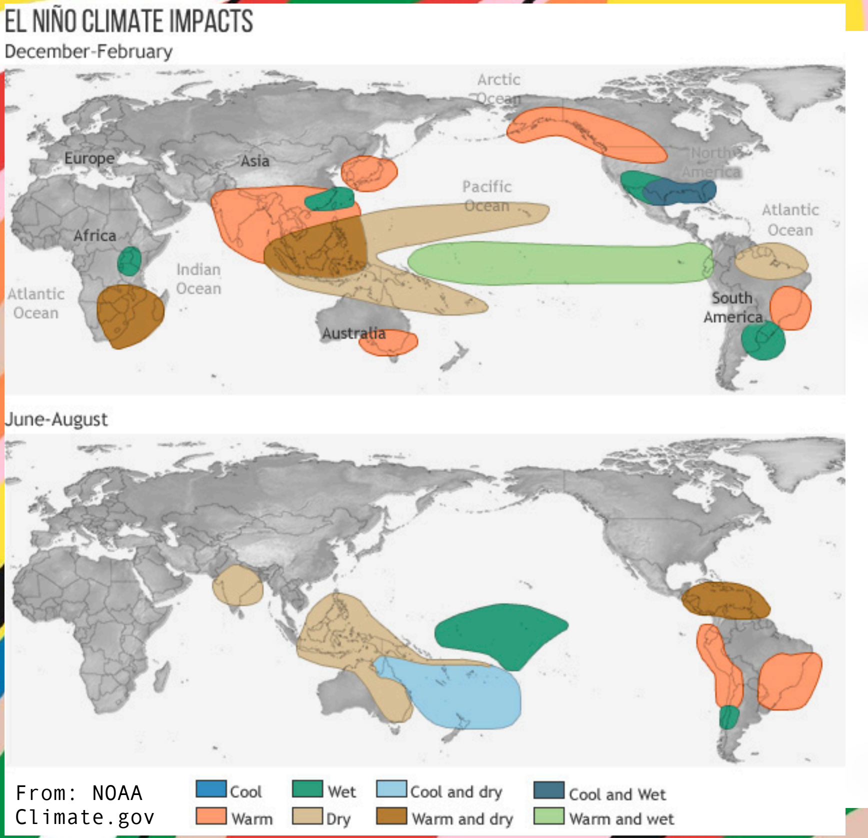

Integrating Climate Forecasts and Societal Decision Making: Challenges to an Emergent Boundary Organization

Shardul Agrawala, Kenneth Broad and David H. Guston

Science, Technology, & Human Values
Vol. 26, No. 4, Special Issue: Boundary Organizations in Environmental Policy and Science (Autumn, 2001), pp. 454-477 (24 pages)

The Use of Seasonal Climate Forecasting in Policymaking: Lessons from Northeast Brazil

Maria Carmen Lemos, Timothy J. Finan, Roger W. Fox, Donald R. Nelson & Joanna Tucker

Climatic Change **55**, 479–507 (2002) | Cite this article

Published: September 2002

Effective and Equitable Dissemination of Seasonal-to-Interannual Climate Forecasts: Policy Implications from the Peruvian Fishery during El Niño 1997–98

Kenneth Broad, Alexander S. P. Pfaff & Michael H. Glantz

Climatic Change **54**, 415–438 (2002) | Cite this article

Shakir Mohamed

Sociotechnical Systems

Our technical work is deeply intertwined with our social world, and rarely separable.

Shakir Mohamed

Asks us to adjust the conceptual apertures we use in our work

Ask our technical and engineering work to account for a wider and more expansive set of considerations

Bringing focus and manageability to the seeming vastness of social considerations

A Sociotechnical Stack

Community & Cooperation

Governance & Accountability

Deployment & Commercialisation

Research & Innovation

A Sociotechnical Stack

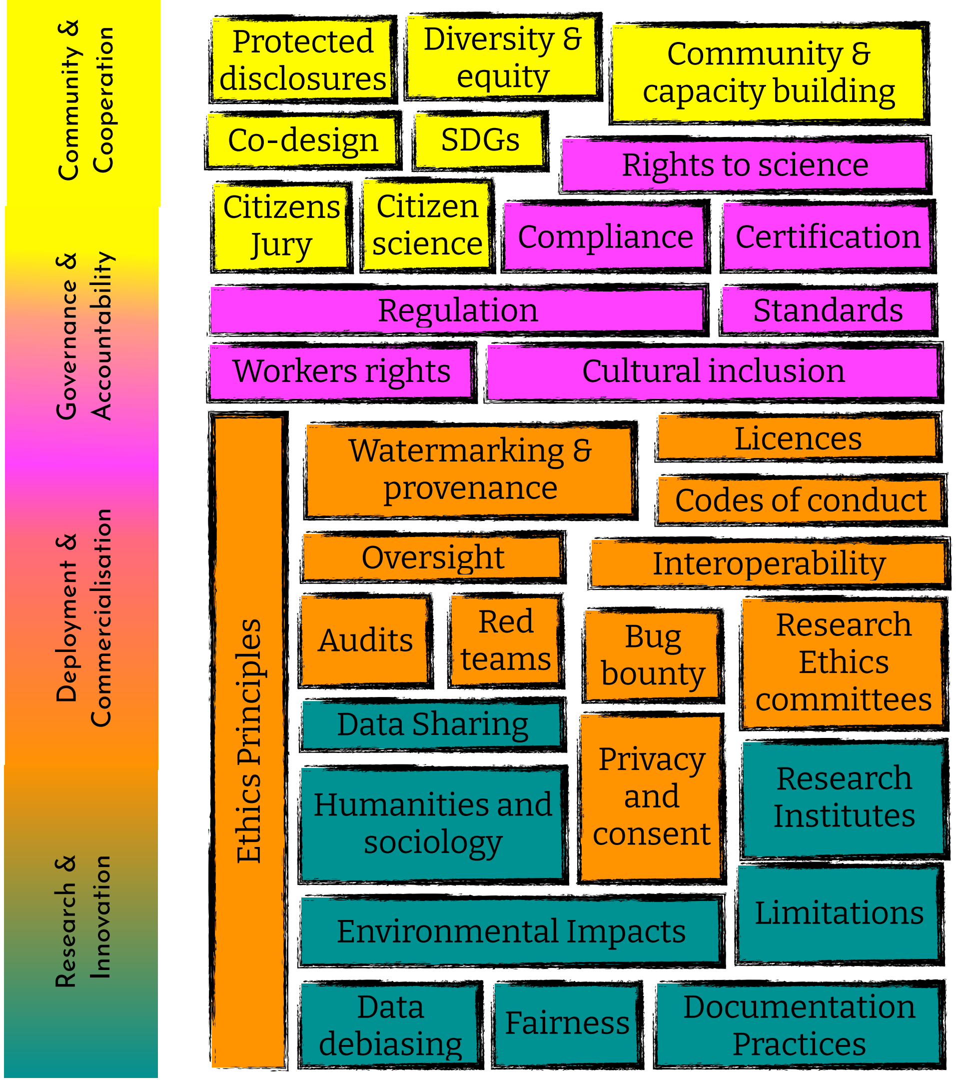

Shakir Mohamed

Participatory AI

PARTICIPATION

The ways that broader communities of people, especially of those most vulnerable, are involved in technology design

Participation means including people in the design of our methods, and being open to changing what we work and how we work, based on their input.

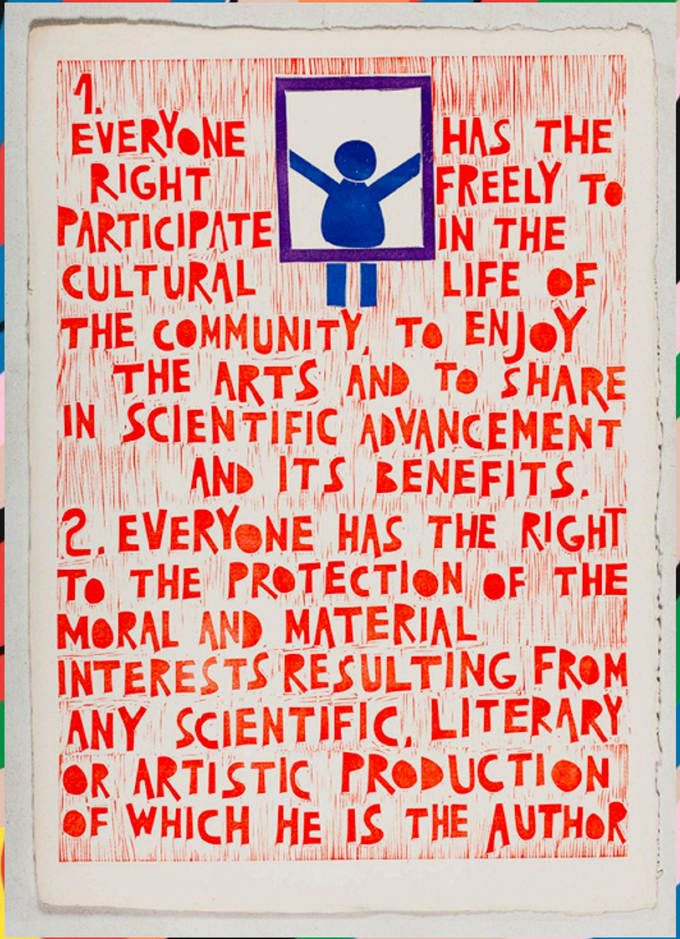

Paper: Power to the People? Opportunities and Challenges for Participatory AI.

More to do

- Establishing more **participatory and community-centred work** to showcase and enable their effective use
- More work on **human-AI interactions** and evaluations.
- More **historical and decolonial work** that continues to provide public memory for the need for a high degree of rigour.
- New instantiations of **social safety and care** that are tied to broad-scope impact assessments and release strategies.
- Evolving practices for **documentation**, **standards** and testing, verification and monitoring, **data rights** and **sovereignty**, and social and **democratic influence**.

03 | THE HUMAN CONDITION

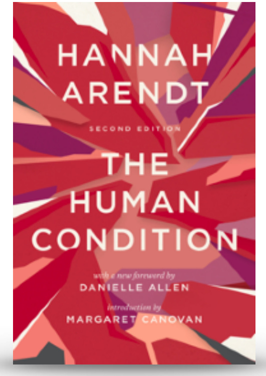

Title from Hannah Arendt's 1958 book that is a critical view of science; and concerned with the vital role of our action in the societies we live in. Cover: Univ. Chicago Press

Empirical Likelihood

Testing Problem

Check whether the mean of an unknown distribution P is equal to a known constant c

n data points $x_1, \ldots, x_n \in \mathbb{R}^d$

Mean condition $\mathbf{c} \in \mathbb{R}^d$

Empirical distribution $\hat{P}_n(x) = \sum_i n^{-1}\delta_{[x=x_i]}$

Weighted distribution $P_{\boldsymbol{\pi}}(x) = \sum_{i=1}^{n} \pi_i \mathbb{I}_{[x_i=x]}$

Likelihood $\prod_i \pi_i \qquad \sum_i \pi_i = 1, \pi_i \geq 0$

Objective $\min_{\{\boldsymbol{\pi} | \sum_i \pi_i=1, \pi_i \geq 0\}} D[\hat{P}_n || P_\pi] \quad \text{s.t.} \sum \pi_i x_i = c$

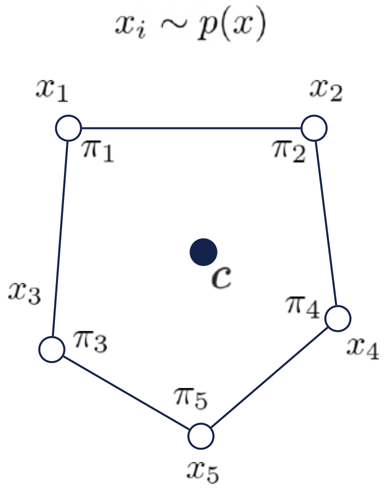

Generalised Empirical Likelihood (GEL)

General moment condition

$$\mathbb{E}_{X\sim P_\pi}[m(X;c)] = \mathbf{0}$$

Generalised Empirical Likelihood (GEL)

$$\min_{\{\boldsymbol{\pi} | \sum_i \pi_i = 1, \pi_i \geq 0\}} D[\hat{P}_n \| P_\pi] \quad \text{s.t.} \quad \mathbb{E}_{X\sim P_\pi}[m(X;c)] = \mathbf{0}$$

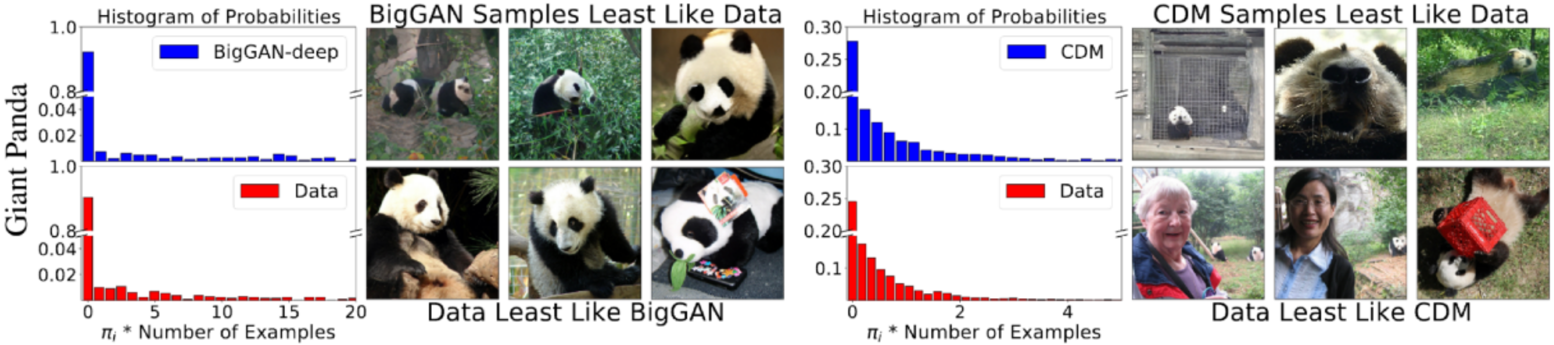

Use to identify examples not represented by the model

Paper:Understanding Deep Generative Models with Generalized Empirical Likelihoods

Fairness and Participation

Medical and other data represent people

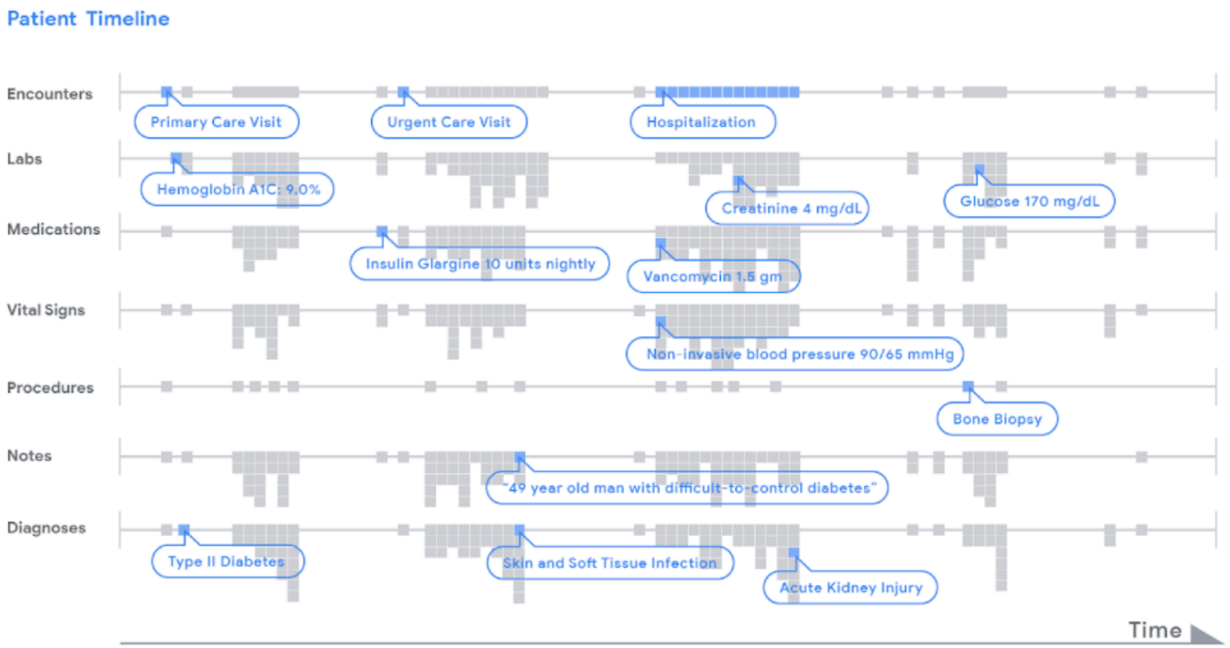

Knowing legal gender or self-identified race or age becomes the basis of fairness analyses.

Paper: Fairness for Unobserved Characteristics: Insight from Impacts on Queer Communities

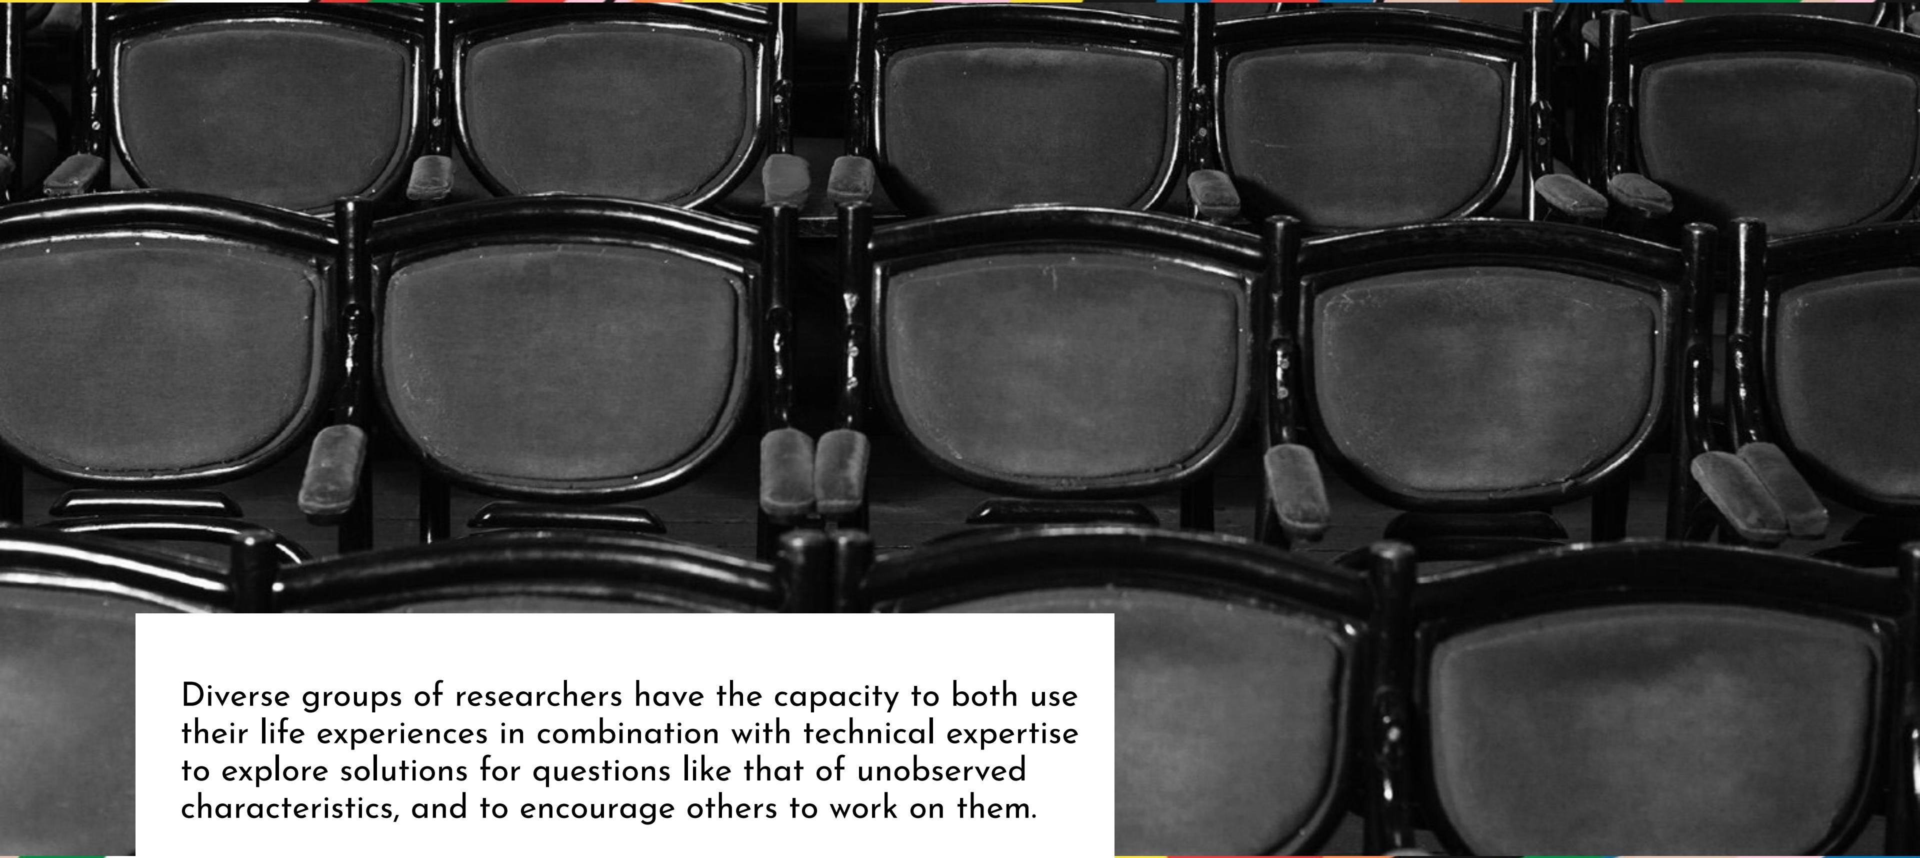

Diverse groups of researchers have the capacity to both use their life experiences in combination with technical expertise to explore solutions for questions like that of unobserved characteristics, and to encourage others to work on them.

Queer Fairness

Assessing fairness for unobserved characteristics

Several Areas of Vulnerability

Privacy; Censorship; Inclusive language; Fighting Online abuse; Health; Mental Health; Employment

REUSABLE METHODOLOGY

Find ways for different types of communities, those who are marginalised, and those who are most vulnerable, to become part of addressing the problem.

Paper: Fairness for Unobserved Characteristics: Insight from Impacts on Queer Communities

AI for Everyone?

INTERCULTURAL ETHICS

How technology can support society and culture, rather than becoming an instrument of cultural oppression and colonialism.

POLITICAL COMMUNITY

Power in strengthening varied forms of political community, who can create new forms on understanding and elevate intercultural dialogue.

Paper: Decolonial Theory as Sociotechnical Foresight in AI

Deep Learning Indaba

2017, Johannesburg, South Africa

Deep Learning Indaba

2018, Stellenbisch, South Africa

Deep Learning Indaba

2019, Nairobi, Kenya

Deep Learning Indaba

2022,Tunis, Tunisia

Deep Learning Indaba

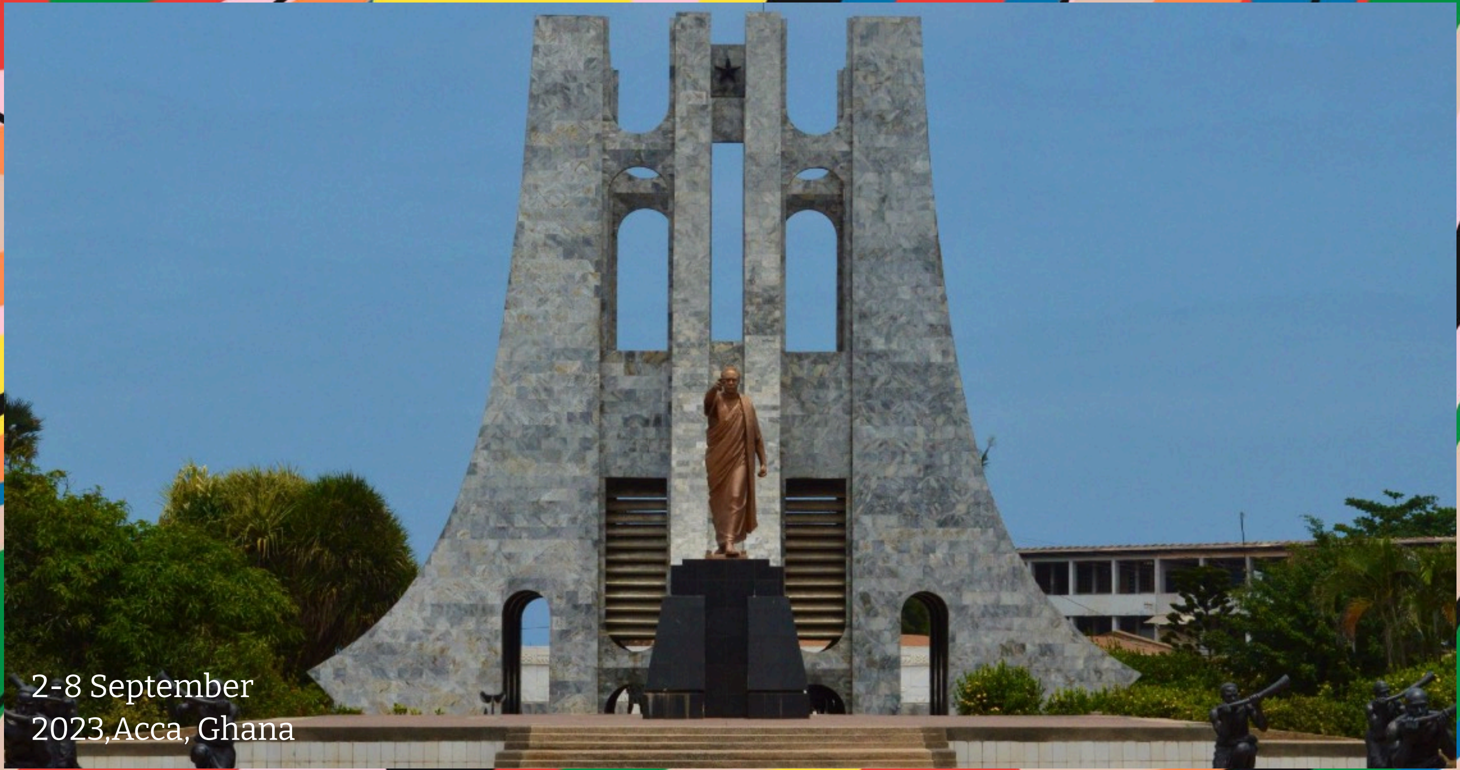

04 | WRAP UP

Title from Svante Arehenius's 1908 book where the problem of global heating was first investigated. Cover image: Project Gutenberg

YOUR SOCIAL PURPOSE

An expanded view of what is within our responsibilities that infuses our work with social purpose.

RECLAIM YOUR AGENCY

Develop a view of the sociotechnical stack, intervene where you best can, and support richer participation in AI.

GLOBAL AI

Support for grassroots transformation and wider ownership of AI is working. Keep doing more.

Some Resources

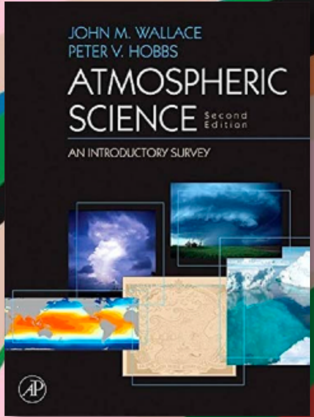

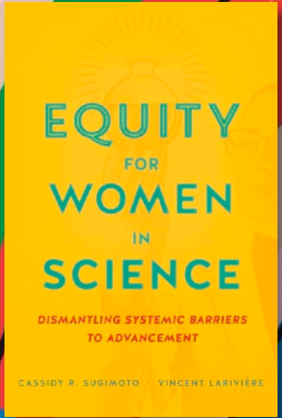

AI, Labor, and the Economy

Guidelines for AI and Shared Prosperity

Explore PAI's Guidelines for AI and Shared Prosperity: tools to design and deploy AI systems in service of workers' rights and well-being.

Article | Open Access | Published: 29 September 2021

Skilful precipitation nowcasting using deep generative models of radar

Suman Ravuri, Karel Lenc, Matthew Willson, Dmitry Kangin, Remi Lam, Piotr Mirowski, Megan Fitzsimons, Maria Athanassiadou, Sheleem Kashem, Sam Madge, Rachel Prudden, Amol Mandhane, Aidan Clark, Andrew Brock, Karen Simonyan, Raia Hadsell, Niall Robinson, Ellen Clancy, Alberto Arribas & Shakir Mohamed

Nature **597**, 672–677 (2021) | Cite this article

GraphCast: Learning skillful medium-range global weather forecasting

Remi Lam[*,1], Alvaro Sanchez-Gonzalez[*,1], Matthew Willson[*,1], Peter Wirnsberger[*,1], Meire Fortunato[*,1], Alexander Pritzel[*,1], Suman Ravuri[1], Timo Ewalds[1], Ferran Alet[1], Zach Eaton-Rosen[1], Weihua Hu[1], Alexander Merose[2], Stephan Hoyer[2], George Holland[1], Jacklynn Stott[1], Oriol Vinyals[1], Shakir Mohamed[1] and Peter Battaglia[1]

[*]equal contribution, [1]DeepMind, [2]Google

We introduce a machine-learning (ML)-based weather simulator—called "GraphCast"—which outperforms the most accurate deterministic operational medium-range weather forecasting system in the

Philosophy & Technology
https://doi.org/10.1007/s13347-020-00405-8

RESEARCH ARTICLE

Check for updates

Decolonial AI: Decolonial Theory as Sociotechnical Foresight in Artificial Intelligence

Shakir Mohamed[1] · Marie-Therese Png[2] · William Isaac[1]

Power to the People? Opportunities and Challenges for Participatory AI

Abeba Birhane
Mozilla Foundation & University College Dublin
Ireland
abeba.birhane@ucdconnect.ie

William Isaac
DeepMind
UK
williamis@deepmind.com

Vinodkumar Prabhakaran
Google
USA
vinodkpg@google.com

Mark Díaz
Google
USA
markdiaz@google.com

Madeleine Clare Elish
Google
USA
mcelish@google.com

Iason Gabriel
DeepMind
UK
iason@deepmind.com

Shakir Mohamed
DeepMind
UK
shakir@deepmind.com

Machine Learning with **Social Purpose**

Shakir Mohamed