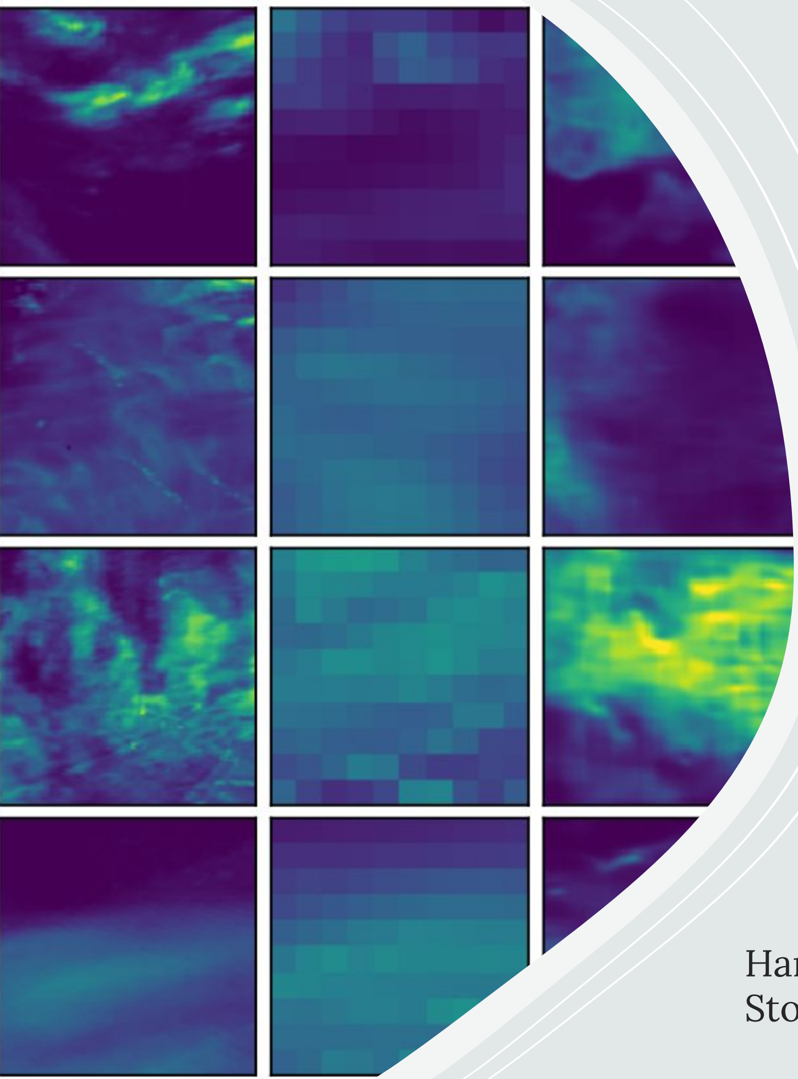

Stochastic Downscaling of Precipitation Forecasts with GANs

Andrew McRae

University of Oxford

on behalf of

Lucy Harris (ex-*Oxford*), Andrew McRae (*Oxford*), Mat Chantry (ECMWF)
Peter Dueben (ECMWF), Tim Palmer (*Oxford*)

Harris et al. (2022) “A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts.” JAMES

What is downscaling?

"Downscaling is any procedure to infer *high-resolution* information from *low-resolution* variables."

Motivation

We are trying to *make better weather forecasts* using machine learning.

- We do this by *post-processing* the output of existing physics-based forecast models.
- We do **not** *replace* existing models entirely.
- Could global model + ML be better than mediocre limited-area models?

Why precipitation (rainfall)?

- Exhibits much more small-scale variation than other fields (pressure, temperature, ...)
- Real-world impact, e.g., natural disasters due to flooding

Motivation
Sub-grid rainfall variation in different meteorological scenarios

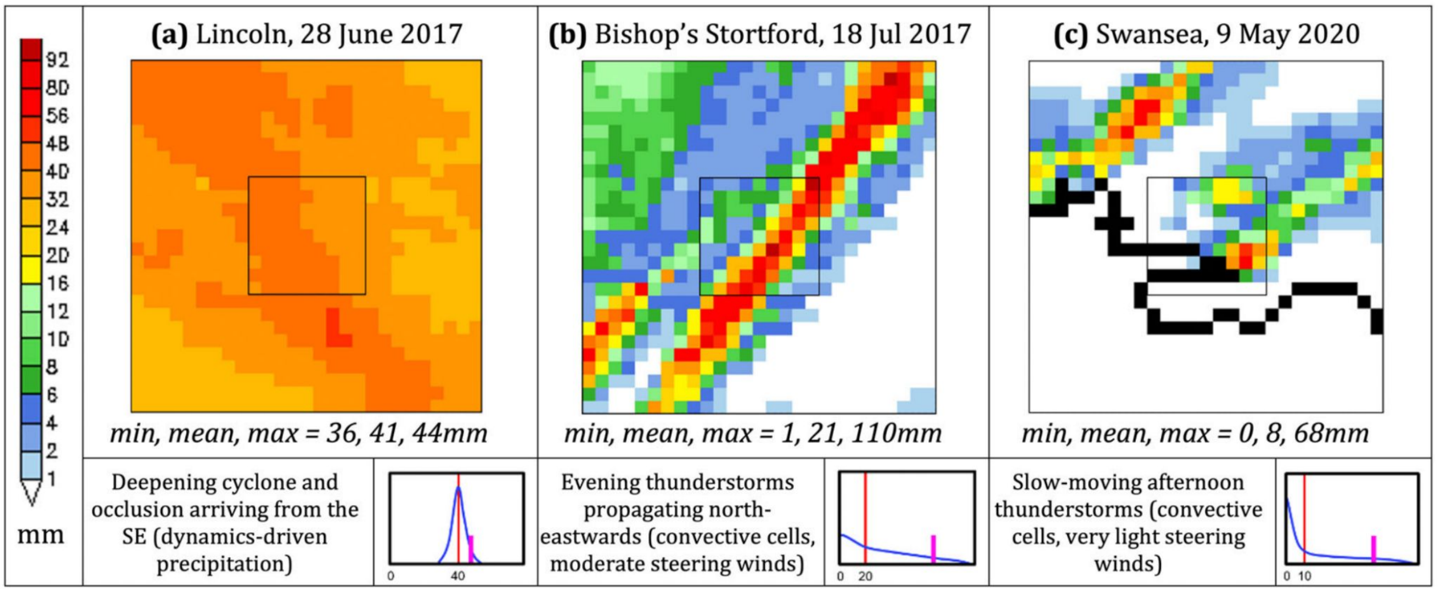

From Hewson & Pillosu (2021) "A low-cost post-processing technique improves weather forecasts around the world." *Nature Communications Earth & Environment*

Downscaling with neural networks

Can we use neural networks to produce **spatially-realistic**, **statistically-reliable** post-processed forecasts?

Why adversarial training?

- Exact locations of rainfall events cannot be predicted (at relevant spatial resolutions and forecast lead times)

- Traditional loss functions (e.g., MSE) are really bad metrics to optimise for, in the presence of substantial uncertainty at the grid-scale ("double-penalty effect")

- Adversarial training: automatically-trained discriminator network *acts as custom loss-function*, tailored to this problem

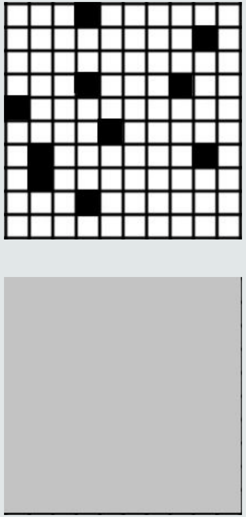

Truth

MSE-optimised forecast

Conditional GANs

Random noise

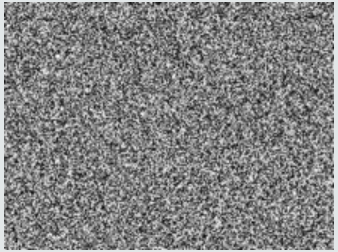

Input condition

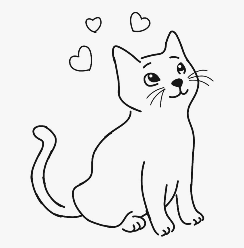

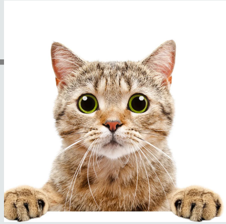

Ensemble of
possible solutions

Conditional GANs

Random noise

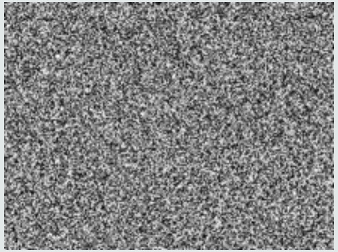

Input condition

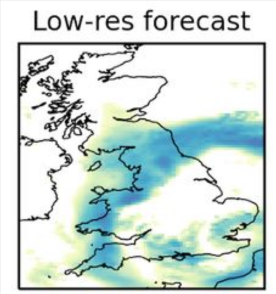

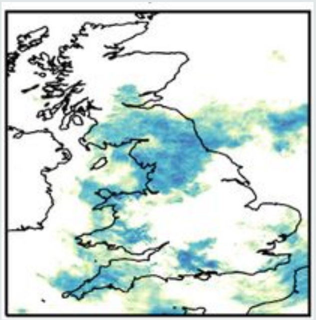

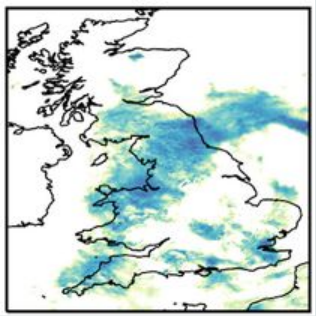

Ensemble of
possible solutions

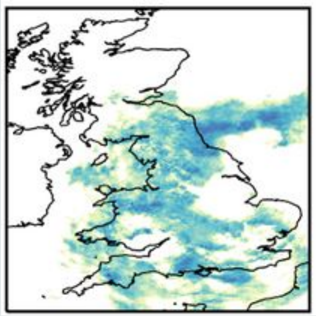

Prior work: Leinonen et al. (2020)

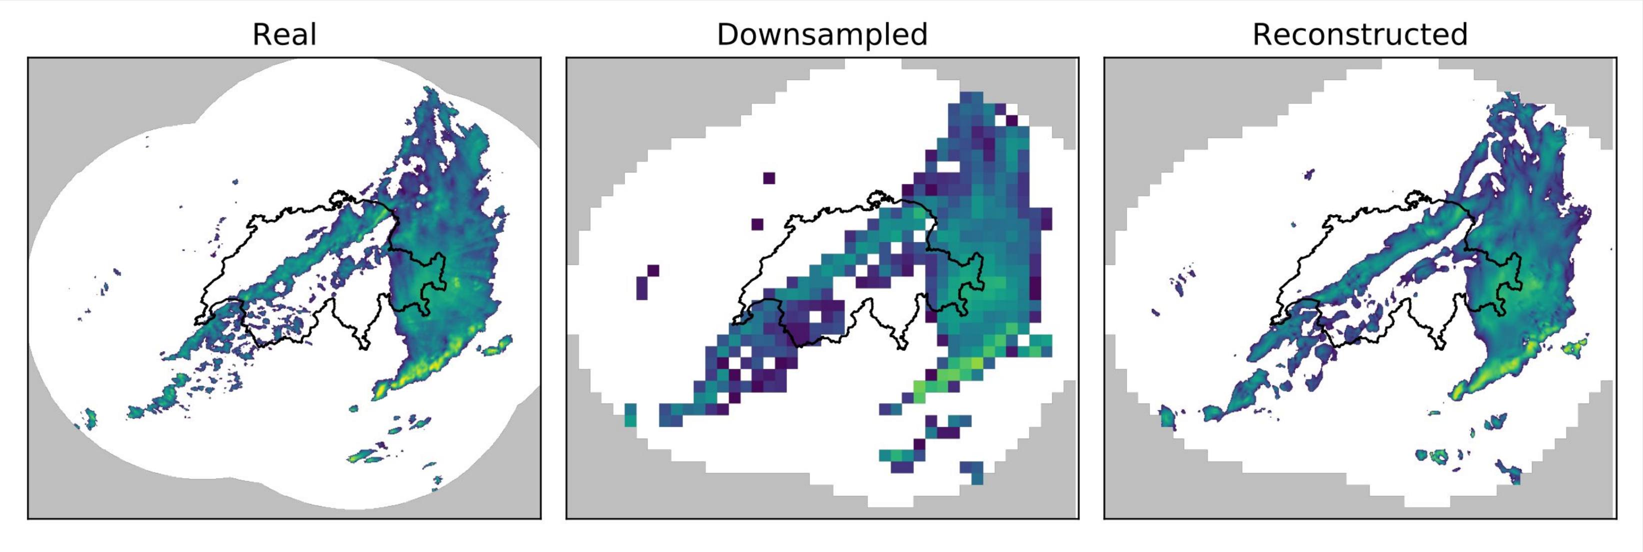

"Stochastic Super-Resolution for Downscaling Time-Evolving Atmospheric Fields with a GAN." *IEEE Transactions on Geoscience and Remote Sensing*

Prior work: Leinonen et al. (2020)

We build on this work <u>technologically</u>, but big scientific differences:

Leinonen: coarsened observations -> observations (forecast error <u>absent</u>)
Us: forecast data -> observations (forecast error <u>present</u>)

We found methods that give good results on the former problem but fail on the latter problem…
(VAEs without adversarial training)

See also: Price & Rasp (2022)

Our work

Network inputs:

- **IFS hi-res** forecasts (6-18hr lead-time)
  - Hourly data
  - Regridded to 0.1° (approx 10km)
  - Fields: total and convective precip, surface pressure, TISR, CAPE, TCLW, TCWV, u700, v700
- **High-resolution geographic data**
  - Orography and land-sea mask
  - Approx 1.25km res, regridded to 0.01°

Target:

- **NIMROD**
  - 1km C-band radar-based rainfall, adjusted with gauge measurements.
  - Accumulated to hourly data
  - Regridded to 0.01° (approx 1km)

TRAINING: 2016, 2017, 2018
TESTING: 2019
EVALUATION: 2020

Domain: 49.5 - 59 latitude, -7.5 - 2 longitude (British Isles)

Our work

Both generator and discriminator are residual convolutional neural networks, approx. 8 residual blocks deep. (Full architectural details in paper)

Can obtain meaningful results in hours, but final paper runs took ~3 days to train (A-100), ~1-2 days for model checkpoint selection

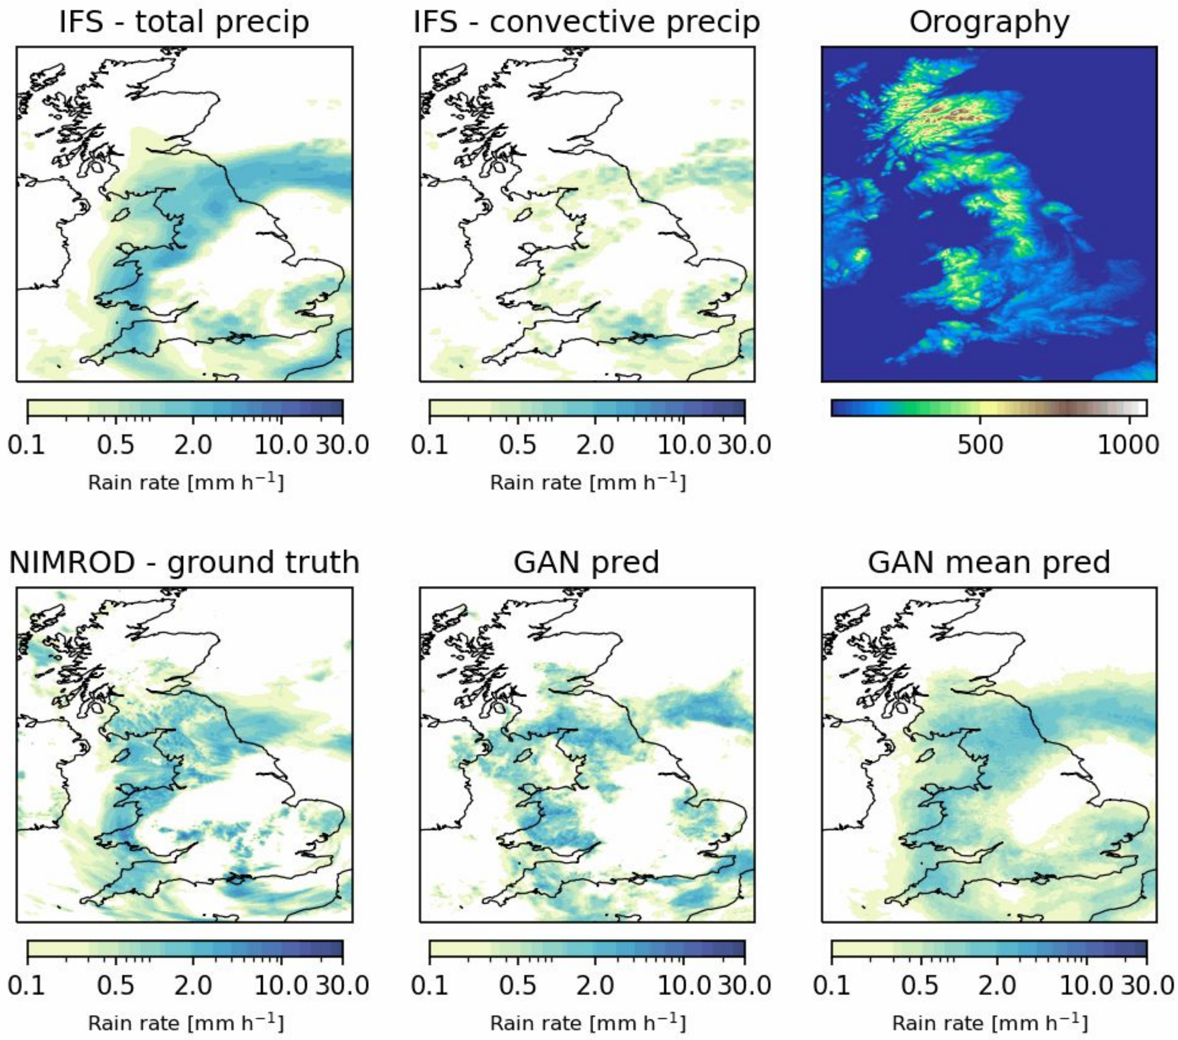

Note: multiple GAN samples, indicating **uncertainty**

NOT time series

Examples

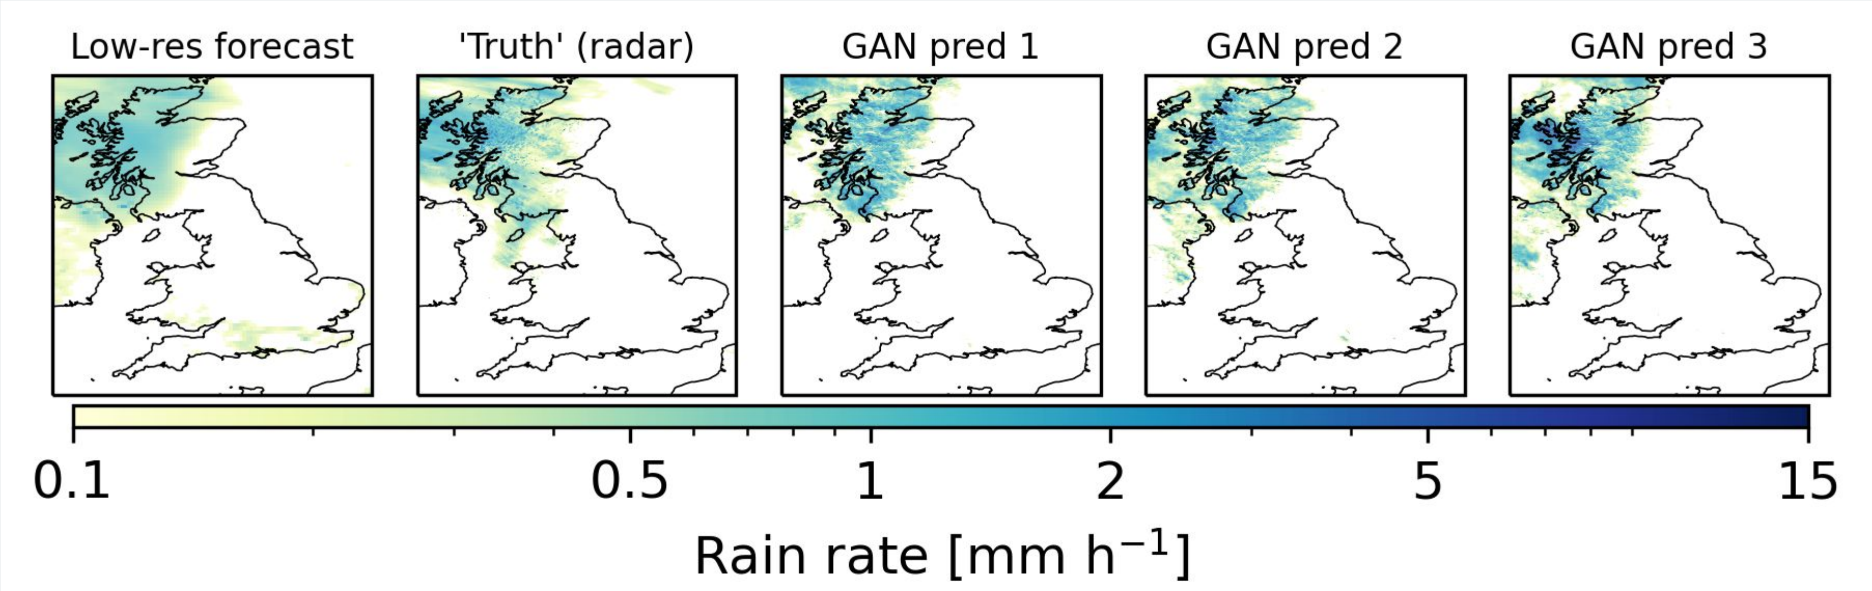

Low-res forecast
'Truth' (radar)
GAN pred 1
GAN pred 2
GAN pred 3
0.1
0.5
1
2
5
15
Rain rate [mm h−1]

Examples

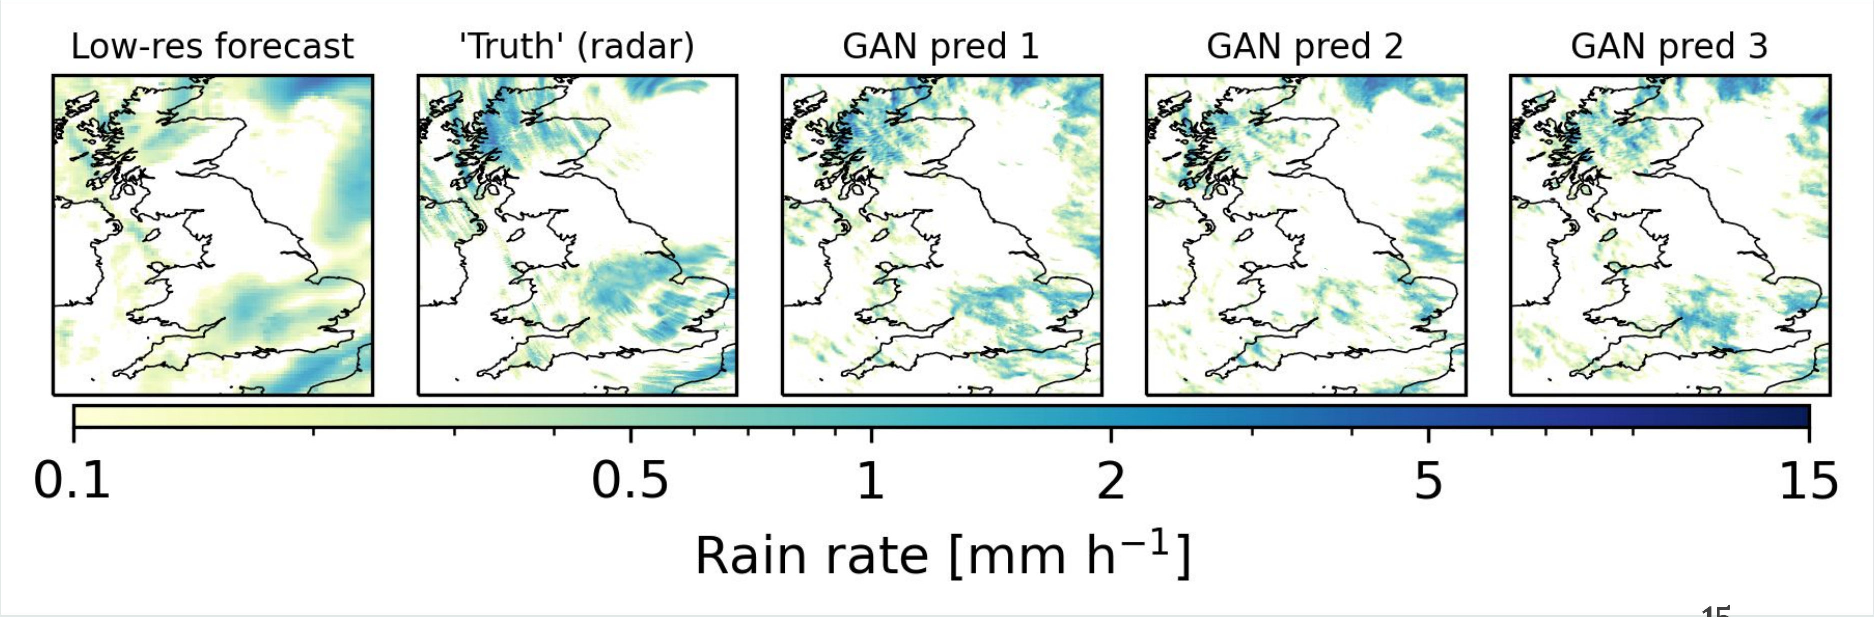

Low-res forecast
'Truth' (radar)
GAN pred 1
GAN pred 2
GAN pred 3
0.1
0.5
1
2
5
15
Rain rate [mm h^{-1}]

Comparison to ecPoint approach

- Gives a probabilistic *point* rainfall prediction within a forecast gridbox, via multiplicative “correction factor”
- 100+ “weather types” (decision tree based on local meteorological variables); separate PDF for each type

Strengths:

- Computationally extremely cheap
- Excellent calibration

Limitations:

- Only uses information from parent grid box
- Output has no information about spatial relationships

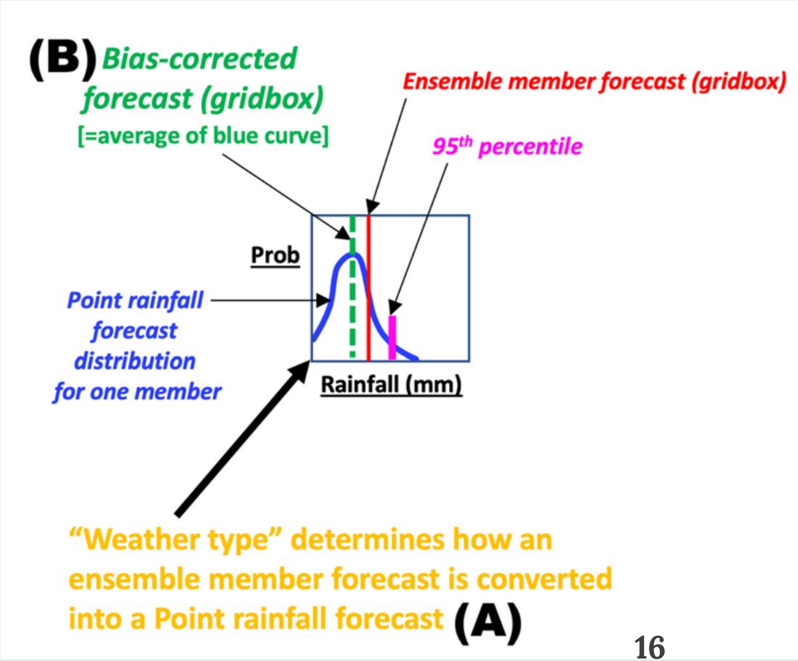

Quantitative metrics

(256 hourly examples from 2020; 100 ensemble members drawn)

Model	CRPS (pixelwise), mm/hr	RMSE (ensemble mean), mm/hr	RMSE (individual members), mm/hr
GAN	0.0856	0.404	0.528
VAE-GAN	0.0852	0.405	0.499
ecPoint approach	0.0895	0.423	0.644
Det CNN (MSE loss)	0.1347 (MAE)	0.404	
Lanczos interpolation	0.1412 (MAE)	0.447	

no. samples where
pixel value is
smaller than truth

Rank histogram plot (100 ens members)

rank

$$r = \frac{N_s}{N_p}$$

total no.
predictions

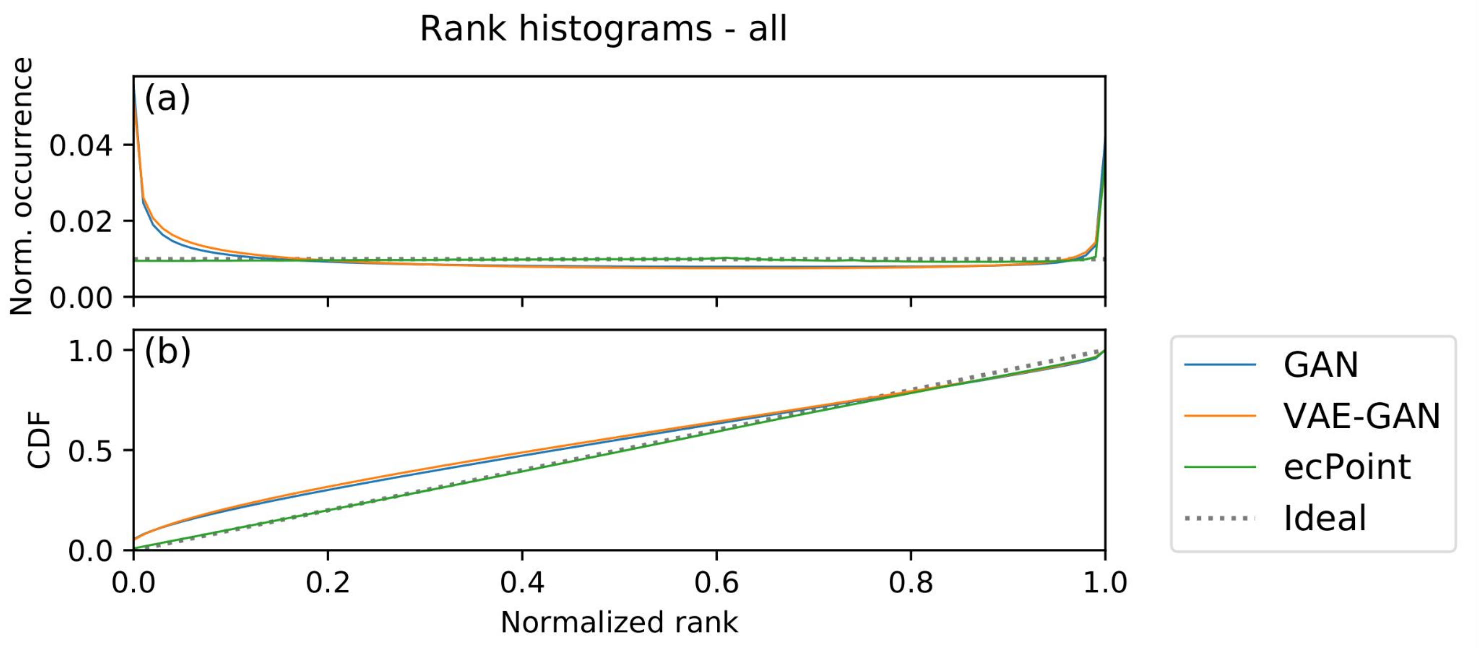

Rank histogram, extreme events only

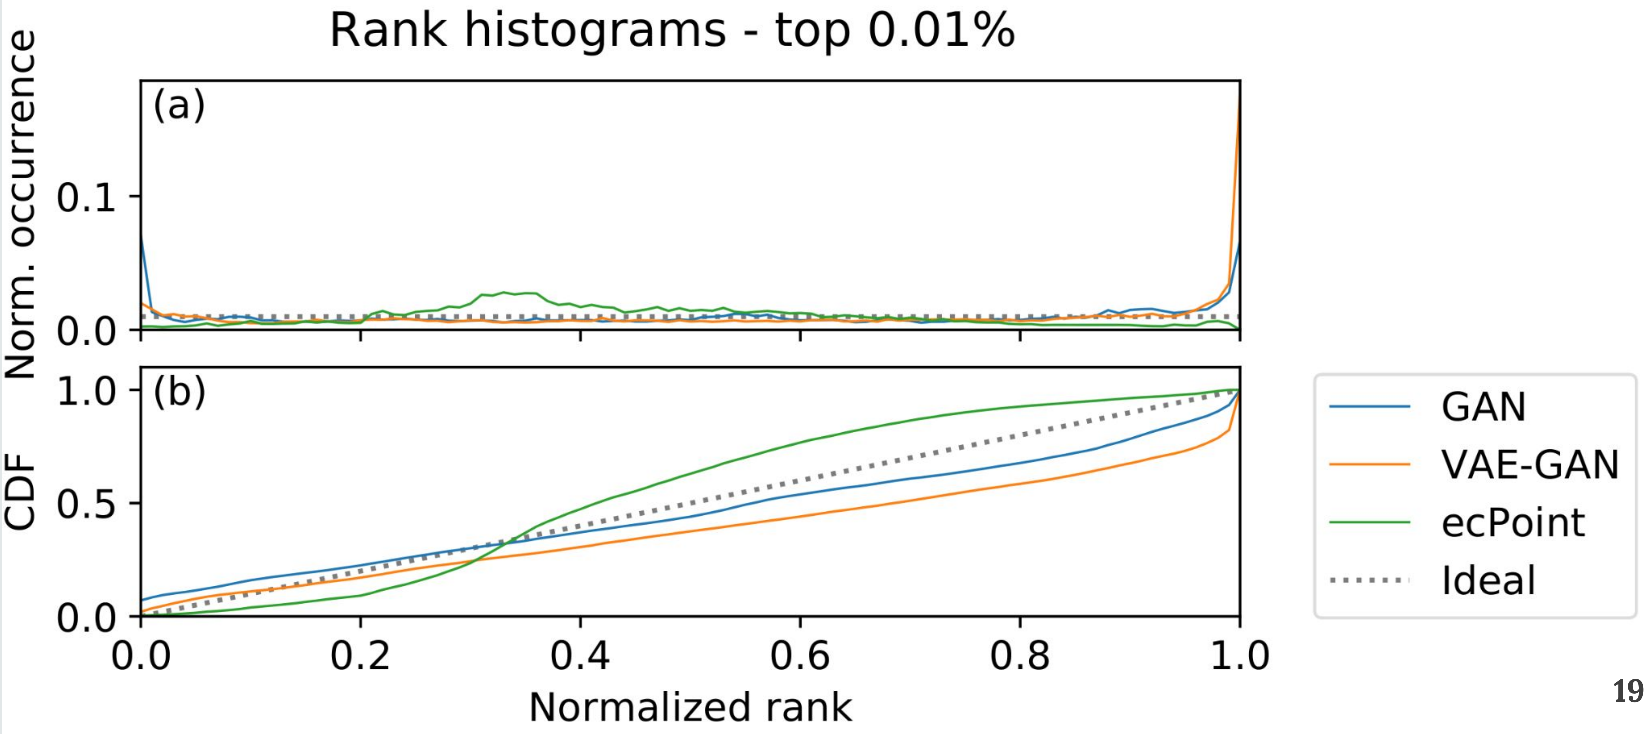

Lots more details in the paper...

- Spatial metrics
  - Spatially-pooled CRPS scores
  - Ensemble Fractions Skill Score
  - Power spectra
- ROC curves, precision–recall curves
- Precise GAN setup
  - Wasserstein GAN + gradient penalty
  - Network architecture
- Data subsampling for training
- Content loss term (ensemble-mean MSE, or CRPS)
- VAE-GAN
- ...

Conclusions

- GAN produces sharply varying but spatially coherent higher-resolution precipitation forecasts
- Similar point-wise accuracy to ecPoint approach (better CRPS, worse calibration)
- Once trained, low computational cost (0.2s/sample)

Related + future work:

- Testing on other geographic regions (USA, Kenya) and datasets (Fenwick Cooper + Bobby Antonio, Oxford)
- Applying to tropical cyclone data (Emily Vosper + Peter Watson, Bristol)
- Restore temporal network aspects; perhaps use IFS ensemble input instead of IFS high-res

Harris et al. (2022) "A Generative Deep Learning Approach to Stochastic Downscaling of Precipitation Forecasts." *JAMES* (google 'Harris McRae Chantry')

Code available on Github/Zenodo (link in paper)