Distilling an Analog Ensemble into a Deep Neural Network

Salvatore Candido[1], Aakanksha Singh[1], and Luca Delle Monache[2]

[1]Loon
[2]Scripps Institution of Oceanography

November 22, 2022

Abstract

We discuss an efficient implementation of the analog ensemble algorithm, the distilled analog ensemble, which is achieved by distilling the post-processing transformation generated by the analog ensemble into a deep neural network. While the analog ensemble has been shown to be able to improve deterministic forecasts and create calibrated probabilistic predictions in many contexts, a common issue with operationalizing a large, global analog ensemble-based system is the amount of data (a corpus of historical forecasts) and post-processing latency required to process that data in the time-critical path of producing a forecast. Deep neural networks are high capacity function approximators, and we demonstrate that we are able to train a network that memorizes the post-processing behavior of the analog ensemble on a particular corpus of forecasts. This technique breaks the scale factor between the size of the historical forecast corpus (larger is better for forecast skill improvements) and the calculation required to post-process the current forecast in real-time operations. We show that the distilled analog ensemble is able to improve European Centre for Medium-Range Weather Forecasts (ECMWF) high-resolution deterministic forecasts of winds in the lower stratosphere using as ground-truth either the ECMWF analysis or observations from Loon high altitude balloons. In this case, rather than requiring terabytes of historical forecast data to apply the conventional analog ensemble, we can perform the post-processing that improves forecast quality on the fly doing computationally efficient forward passes through a pre-trained network that has a data size of only 100's of kilobytes.

DISTILLING AN ANALOG ENSEMBLE INTO A DEEP NEURAL NETWORK

SALVATORE CANDIDO[1], AAKANKSHA SINGH[1], AND LUCA DELLE MONACHE[1,2]

[1] Loon, Mountain View, California, United States
[2] Scripps Institution of Oceanography, San Diego, California, United States

OBJECTIVES

Billions of people are still without internet access

Loon is a network of stratospheric balloons delivering connectivity to under-served communities

Accurate wind predictions in the lower stratosphere are crucial to steer balloons over desired areas

LOON STRATOSPHERIC WIND OBSERVATIONS ARE AVAILABLE ON REQUEST TO LOON AT LOON-STRATOSPHERIC-DATA@GOOGLE.COM

METHODS

Implement an analog ensemble (AnEn) to generate accurate stratospheric wind predictions and reliable uncertainty quantification.

The analog ensemble is **generated from a historical data set** including the European Centre for Medium-Range Weather Forecasts deterministic high resolution (ECMWF HRES) prediction and analysis.

Ground-truth data sets: **ECMWF HRES analysis** (to generate AnEn and verification) and balloon observations (for verification).

A **cloud-based distributed strategy** is explored to generate 0-10 days predictions over a three-dimensional global domain in near real time.

A **deep neural network** is used to distill the analog ensemble with the goal of scaling on tens of years of historical forecasts without loss of skill.

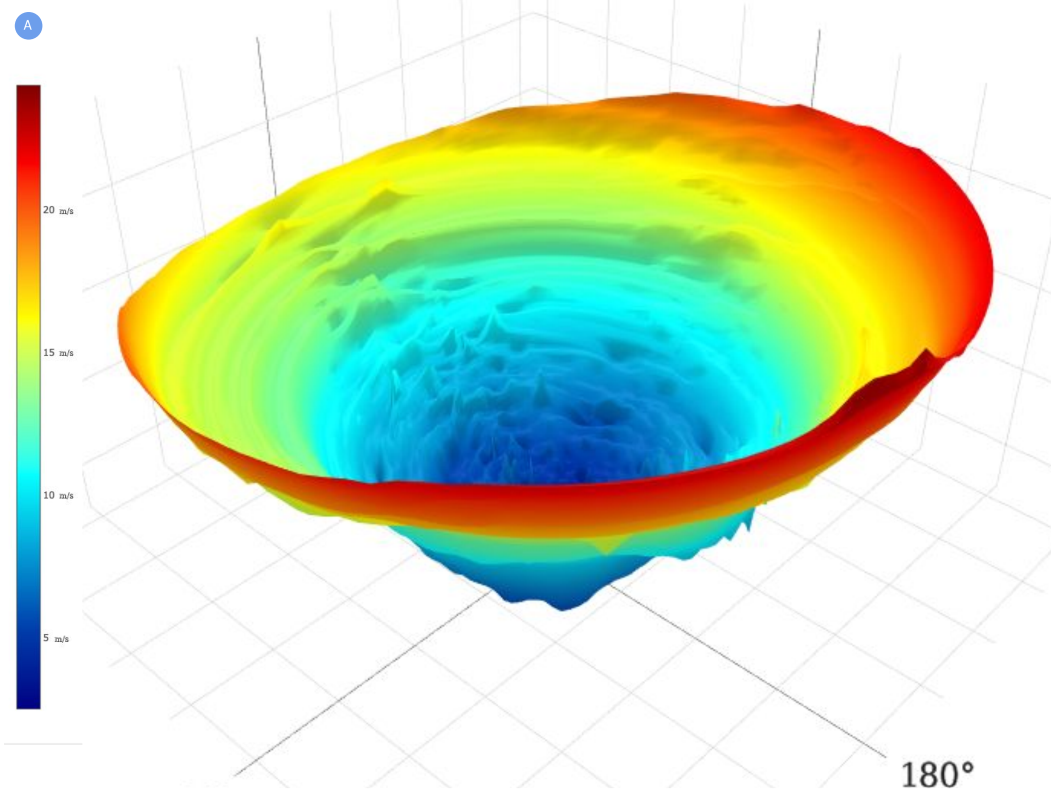

SPEED PREDICTION FOR A PARTICULAR LOCATION AND LEAD TIME SWEPT OVER SPEED (DISTANCE FROM ORIGIN) AND DIRECTION (ROTATION AROUND Z-AXIS)

TOP-DOWN VIEW OF PLOT A

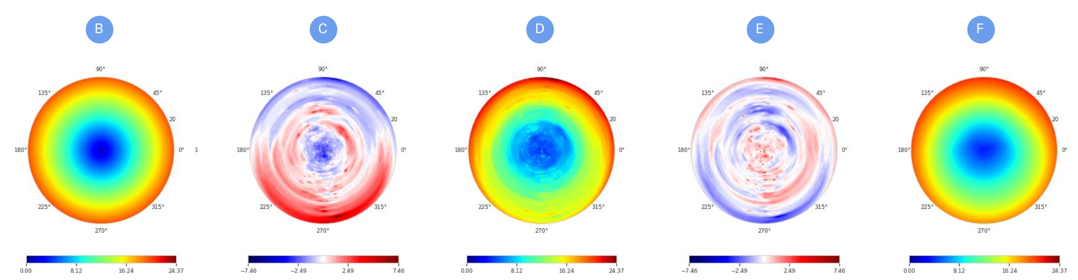

IDENTITY OPERATOR, I.E., NO AnEn POST-PROCESSING

THE DIFFERENCE BETWEEN PLOT B & D

AnEn POST-PROCESSING

THE DIFFERENCE BETWEEN PLOT D & F

POST-PROCESSING LEARNED BY THE DISTILLED AnEn

RESULTS | DETERMINISTIC

AnEn reduces HRES CRMSE for the prediction of wind speed and direction, particularly at longer lead times. Improvements are larger for wind direction. The distilled AnEn has very similar performance metrics as the conventional AnEn.

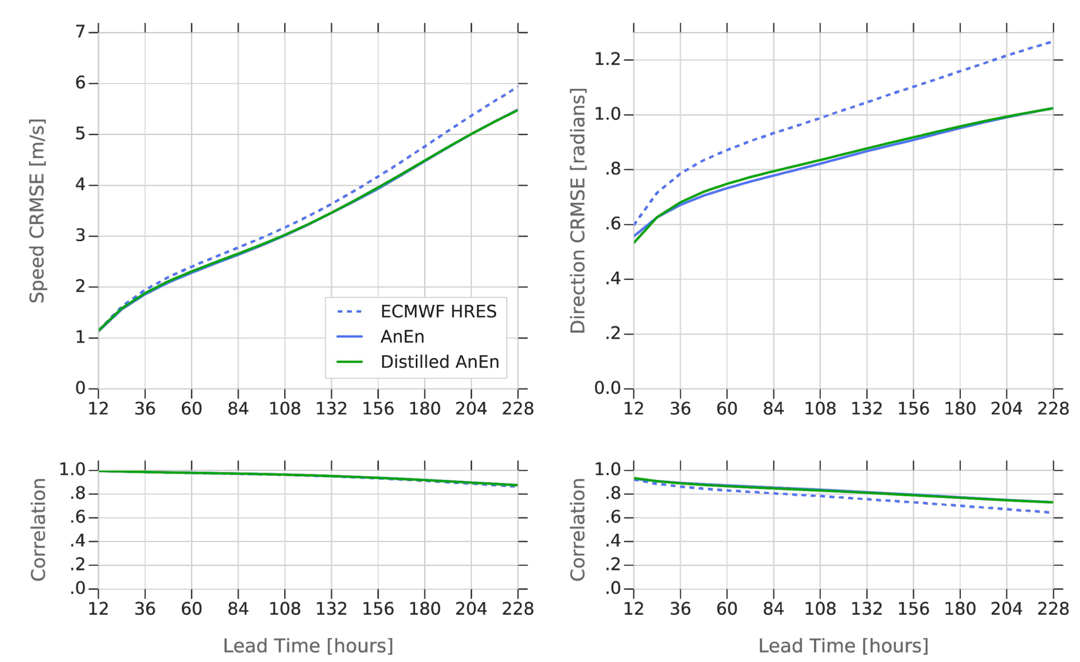

The above is true when either HRES analysis (above) or observations from Loon balloons (below) are used as ground-truth.

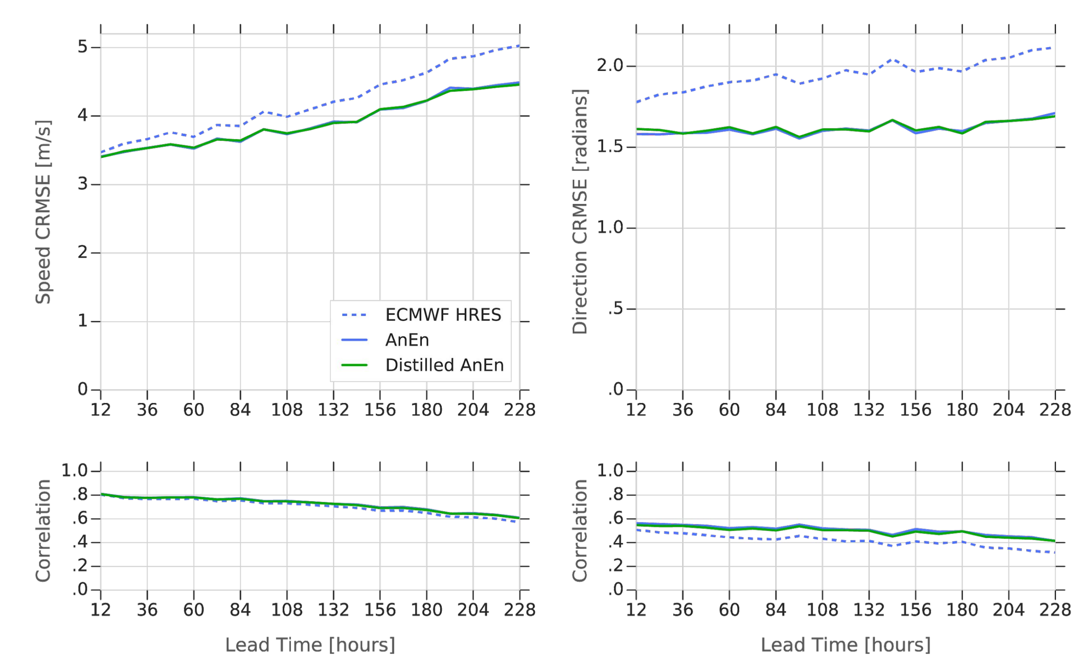

RESULTS | PROBABILISTIC

For probabilistic predictions, AnEn is compared to a persistence ensemble (PeEn). The PeEn consists of selecting the last 20 available ground-truth values to generate a 20-member ensemble.

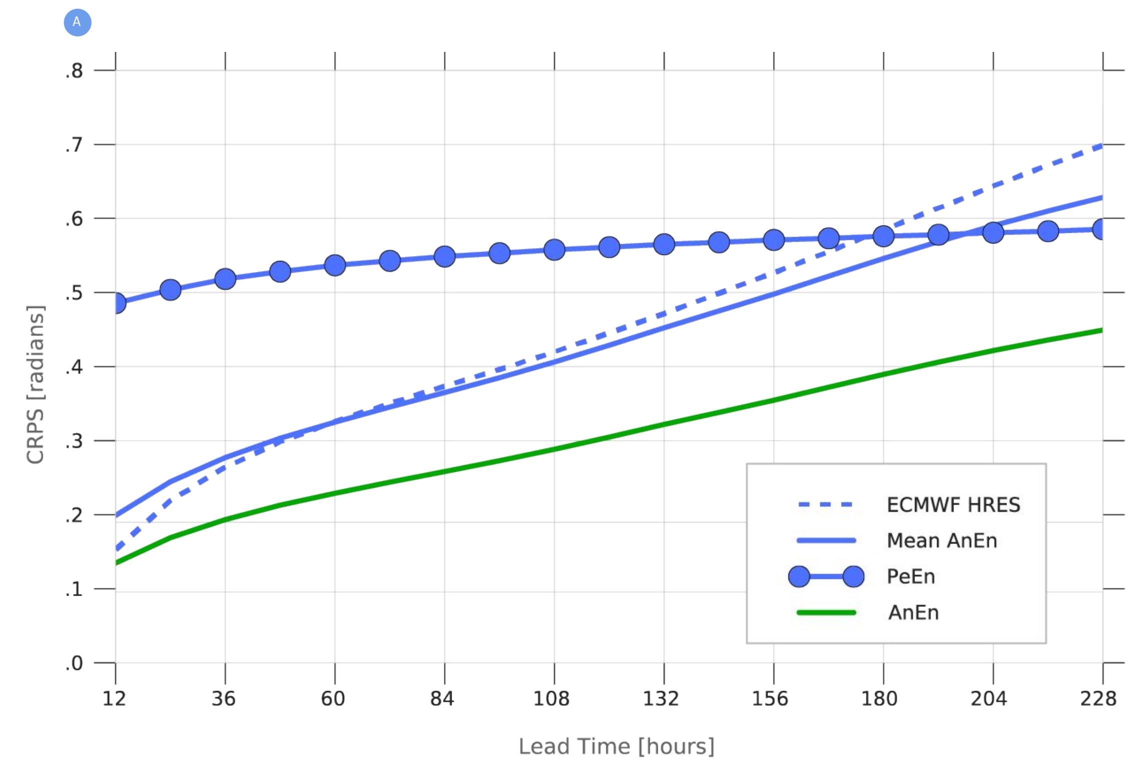

AnEn MATCHES THE GROUND-TRUTH CUMULATIVE DISTRIBUTION FUNCTION SIGNIFICANTLY BETTER THAN PeEn, AS EVIDENT FROM THE (A) CONTINUOUS RANK PROBABILITY SCORE (CRPS; THE LOWER THE BETTER)

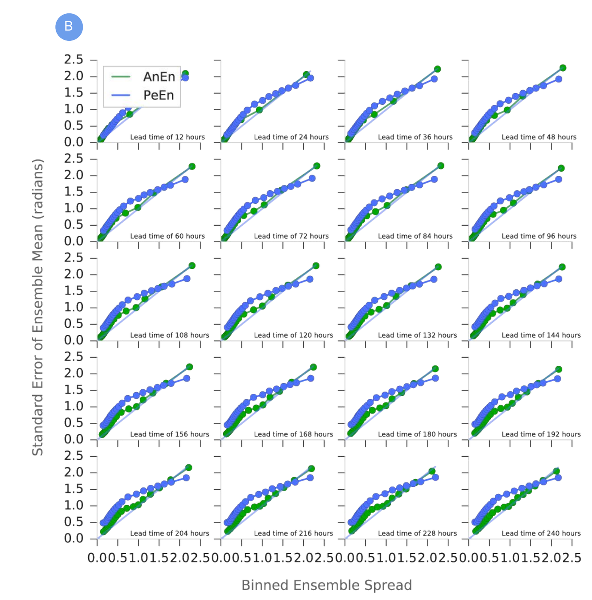

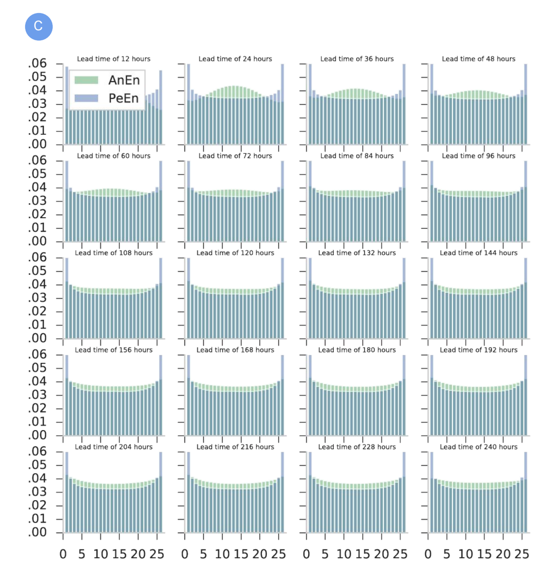

AnEn IS ABLE TO BETTER QUANTIFY THE PREDICTION UNCERTAINTY AS EVIDENT FROM THE (B) BINNED- SPREAD/SKILL PLOTS, AND IT PRODUCES A MORE STATISTICALLY CONSISTENT ENSEMBLE AS EVIDENT FROM THE (C) RANK HISTOGRAMS THAN PeEn

KEY POINTS

The AnEn **generates accurate predictions** of lower-stratosphere winds & reliably quantifies prediction uncertainty.

Distilling AnEn into a **deep neural network** allows scaling on tens of years of historical forecasts without loss of skill.

A **cloud-based distributed computing** implementation builds global three- dimensional predictions in tens of minutes.

The described capabilities allow for an operational **real-time implementation** of the distilled analog ensemble.

WIND SPEED FORECAST CORRECTION

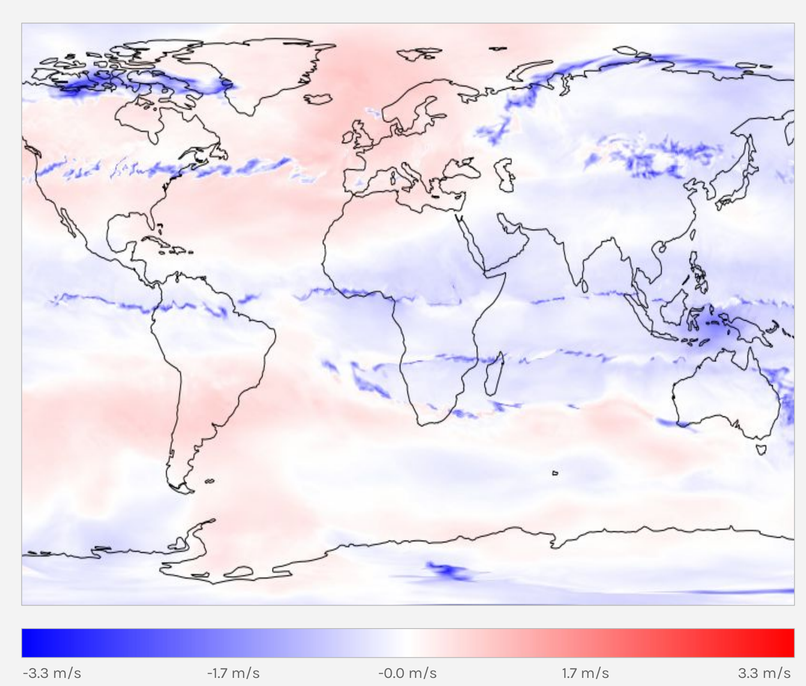

FORECAST CHANGES BY THE DISTILLED AnEn

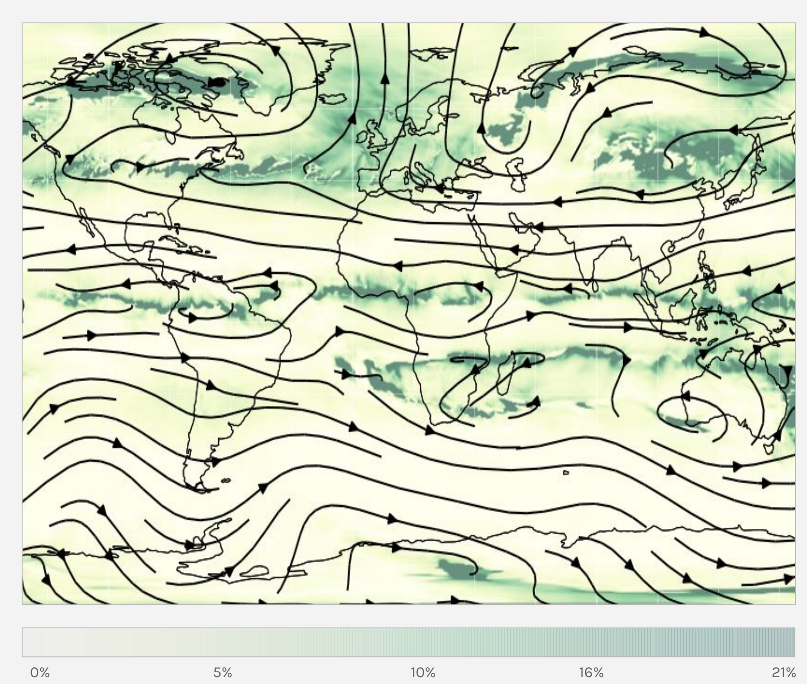

TRAINING PERIOD: JULY 2016 – SEPTEMBER 2017. VERIFICATION PERIOD: OCTOBER 2017 – SEPTEMBER 2018