Using Neural Networks for Data Cleaning in Weather Datasets

Jack R. P. Hanslope & Laurence Aitchison
University of Bristol
jack.hanslope@bristol.ac.uk

Abstract

In climate science, we often want to compare across different datasets. Difficulties can arise in doing this due to inevitable mismatches that arise between observational and reanalysis data, or even between different reanalyses. This misalignment can raise problems for any work that seeks to make inferences about one dataset from another. We considered tropical cyclone location as an example task with one dataset providing atmospheric conditions (ERA5) and another providing storm tracks (IBTrACS). We found that while the examples often aligned well, there were a considerable proportion (around 25%) which were not well aligned. We trained a neural network to map from the wind field to the storm location; in this setting misalignment in the datasets appears as "label noise" (i.e. the labelled storm location does not correspond to the underlying wind field). We found that this neural network trained only on the often noisy labels from IBTrACS had a denoising effect, and performed better than the IBTrACS labels themselves, as measured by human preferences. Remarkably, this even held true for training points, on which we might have expected the network to overfit to the IBTrACS predictions.

Introduction

- Often we want to compare across datasets
- This can can be difficult if the datasets are misaligned
- Which may be the case if one or more is a *reanalysis*
- We studied tropical cyclone tracking and observed a misalignment
- Atmospheric conditions provided by one dataset (ERA5) and storm locations by another (IBTrACS)
- We trained a U-Net to determine location from atmospheric conditions
- This network almost always matched human perception of storm location

Methods

- We use only wind data; we could be more accurate by using more variables but goal is not to maximise accuracy
- Restrict ourselves to only one storm basin
- Network takes as input a wind field
- Network outputs storm location
- We grid the output space and perform classification; this enables the network to output probabilities

Results

- Qualitative results obtained by plotting a wind field (see Figure 1)
- IBTrACS location is denoted by the orange cross
- U-Net confidence is indicated by the colourmap

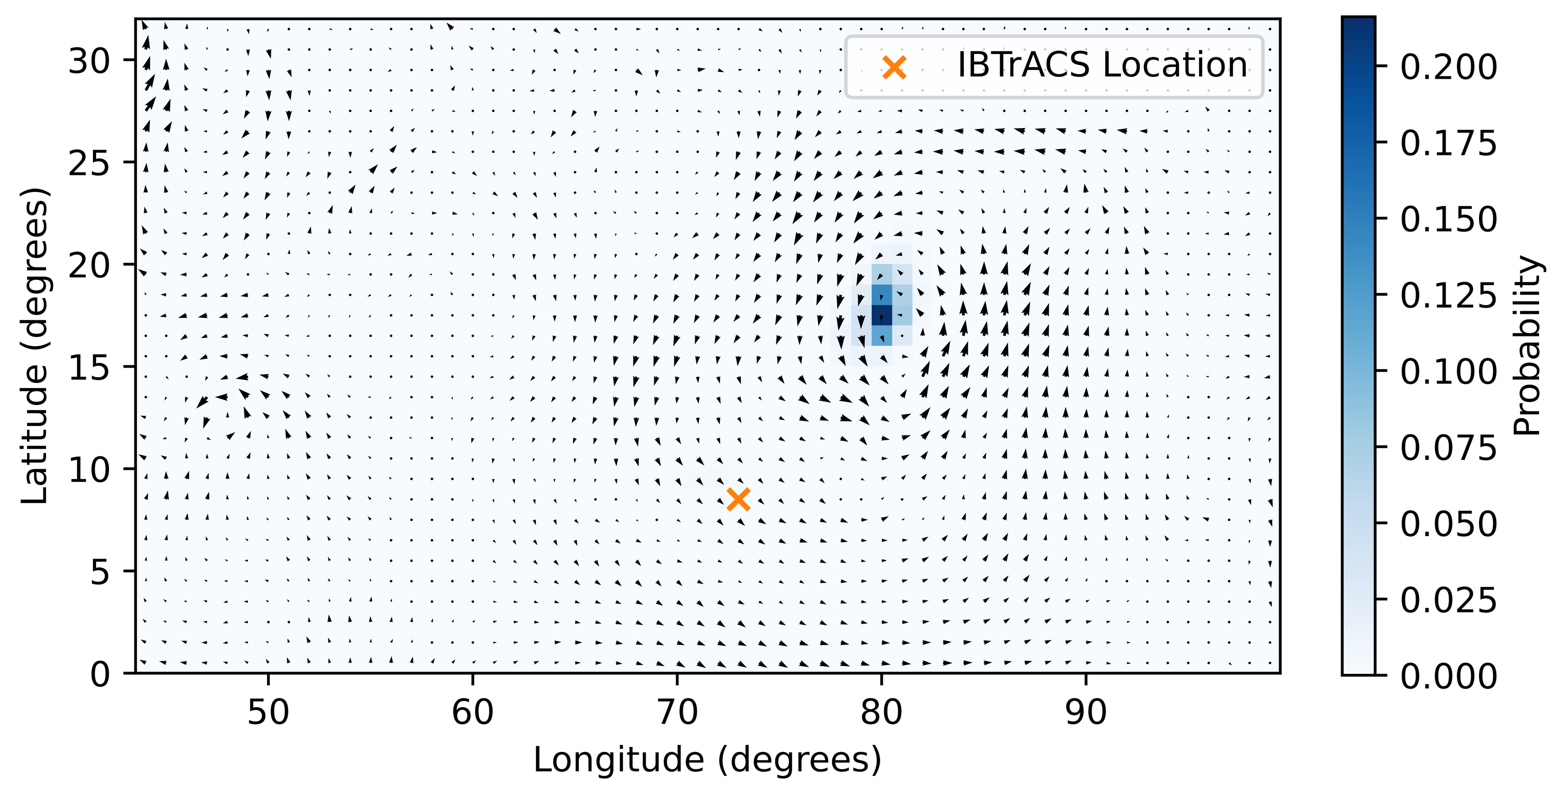

Figure 1: Wind Field for 1980-10-19. IBTrACS location of storm centre is marked by orange cross and probability predicted by the U-Net is shown by the colourmap. This example is hand picked.

- Quantitative analysis is difficult as the true ground truth (storm location) is unobserved
- We showed an author a wind field with two storm locations and asked for preference for one, the other or neither (see Figure 2)
- One location was from model output and other was from IBTrACS
- Not indicated which was which
- Performed one-tailed binomial tests to determine if, when there was a preference, one source was preferred over the other
- Results (see Table 1) strongly suggest U-Net is more performant

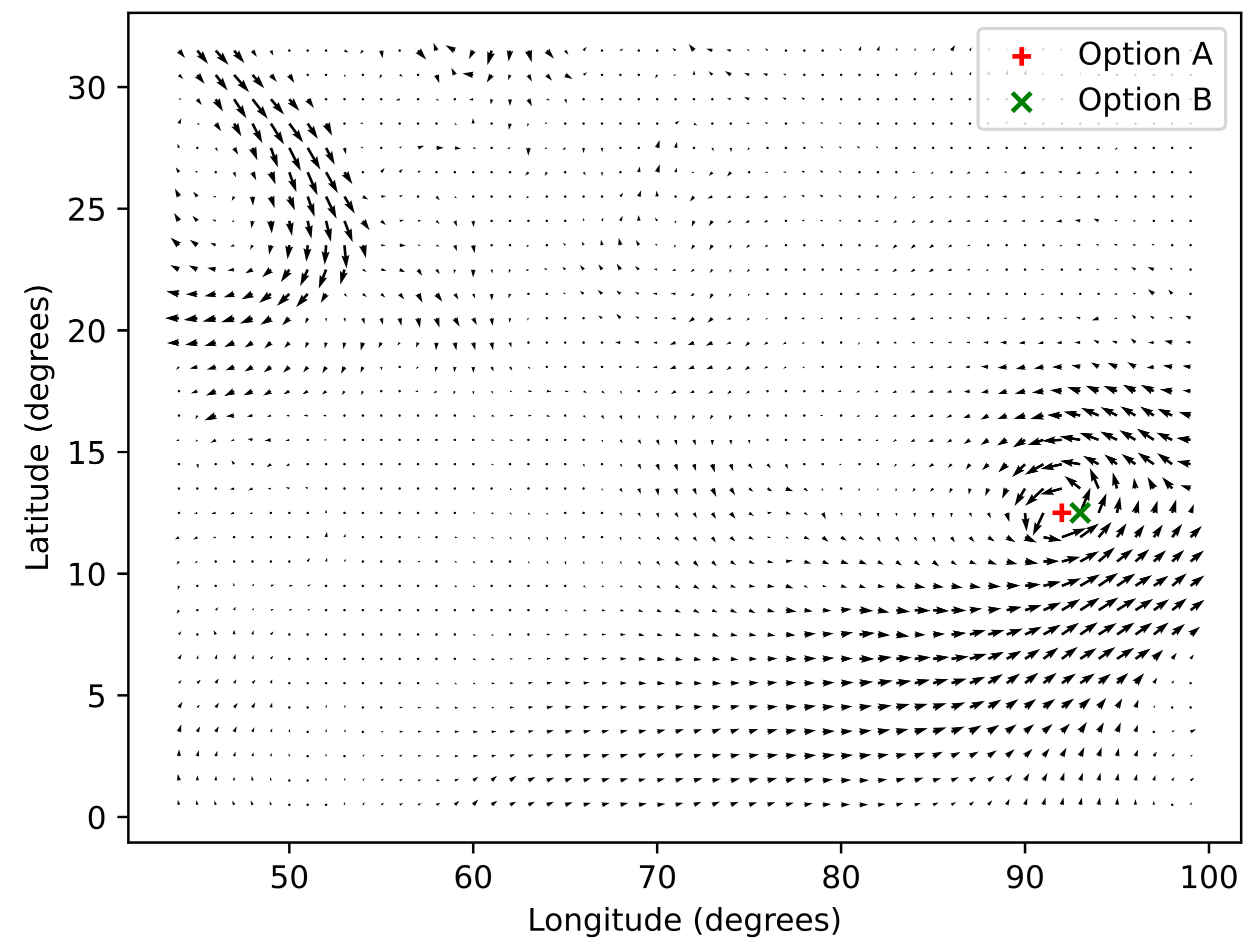

Figure 2: Example of a wind field shown to the author for comparison.

Table 1: Preferences of U-Net location, IBTrACS location and no preference for test set and train set. 200 images were shown for each set. The p-value is calculated as a one tailed binomial test using only the instances where one or other location was preferred.

Preference	Test Set	Train Set
U-Net	46	49
IBTrACS	13	15
Neither	141	136
Total	200	200
p-value	9.6×10^{-6}	1.2×10^{-5}