Emulating maximum gust speed

Team “Let it Rip”

Introduction

- We adopted two methods; emulating maximum wind gust speed at 10m with two ML models, and predicting return periods of extreme gust speeds using Extreme Value Theory (peak-over-threshold), for a specific 12km gridpoint in the UK.
    - We excluded spatial data given the time-constraints
- We found suitability of each method depends on the purpose - whether one is more interested in average gust trends or more damaging gusts.
- The MET office defines gusts as “the maximum three second average wind speed in any period, although it is considered up to 20 seconds in practical cases”. The database provides daily mean wind speed and gust speed across the UK over the period of data collection. In this analysis, we considered gust speeds above 21m/s as extreme gusts.

Background literature

- Informed which variables we included in our emulator
  “The higher the average wind speed, the gustier it is likely to be. This relationship depends on the local terrain, though: towns and cities generate more gusts through friction and solar heating. For the same average wind speed, urban areas are therefore likely to see much stronger gusts than open landscapes.
  Gusts are caused by turbulence from wind shear, friction (from the surface that the air is passing over) or instability in the air (from say surface heating from the sun).”
- Informed which models were historically deemed appropriate
  https://journals.ametsoc.org/view/journals/mwre/140/3/mwr-d-11-00075.1.xml:
    - Non-homogeneous Gaussian regression (NGR) statistical postprocessing method with appropriately truncated Gaussian predictive distributions.

  http://www.research.lancs.ac.uk/portal/en/publications/statistical-models-for-extreme-weather-events(27067bfa-c708-4936-a1e5-cca4d101b311).html
    - Poisson Process Reparameterization for Bayesian Inference for extreme weather events

Our code is available at: https://github.com/oscarkey/wind

Throughout, we’ve highlighted the limitations of our results in italics.

Emulation results

We first attempted to emulate the daily maximum wind gust speed at 10 m from a subset of the core variables from the Met Office model output.

We performed several preliminary experiments to investigate different approaches. These results are available in the file single_point.ipnb in the GitHub repository.

Initial data examination

We started by looking at the correlation between the input variables and max gust speed. In the plot below, we have computed the Pearson correlation coefficient between the named variable and the max gust speed over time, on each 12km grid point:

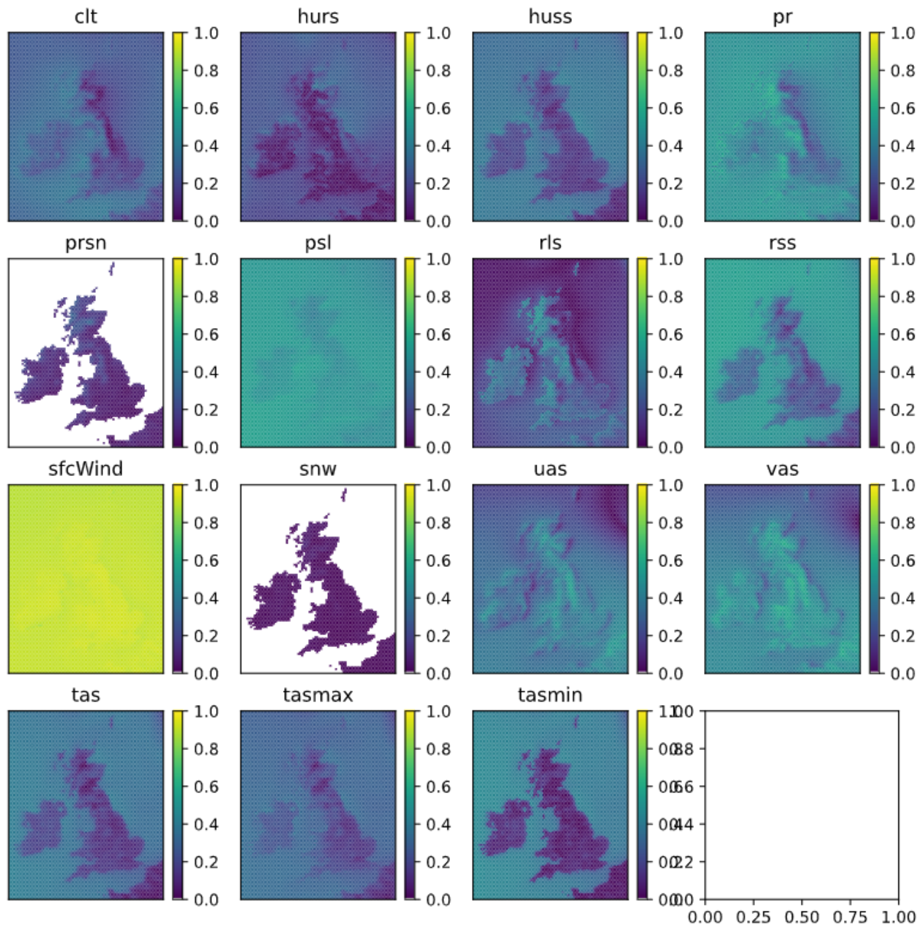

As we might expect, the sustained wind speed (sfcWind) and directional wind speeds (uas, vas) show a strong correlation. Psl (surface pressure) also showed a fairly strong correlation. This influenced the variables we included in the models below.

Emulator configuration and evaluation

In each case the input to the emulator was a subset of the following core outputs of the Met Office model, at daily time resolution:

- pr (precipitation)
- psl (pressure at sea level)
- sfcWind (mean sustained wind speed at 10m)
- uas (mean eastward wind speed)
- vas (mean northward wind speed)

We trained the model on 1080 contiguous days of data, starting at day 500 in the dataset.

We then evaluated the model on a different 1080 days of data, starting at various points in the future.
We considered each geographic grid point individually, so this approach will not take account of spatial correlations, which might help to improve performance.

Models considered:

Model	**Inputs**	**Notes**
Bayesian linear	Only sfcWind	Very simple model
Gaussian process	All five as above	Configuration selected by grid search

Evaluation metrics:
- MSE (Mean Squared Error) between predicted and true maximum gust speed for each day
  Gives general measure of goodness-of-fit
- IEE (Identified Extreme Events), the number of days on which the max gust speed >21m/s, and the upper credible interval of the model was also above 21m/s
  This tries to focus on the task of predicting dangerous events

Experiments at a single location
To start, we chose a location in the middle of the UK (y=40, x=55 in the dataset), as shown below, and considered only data from this point.

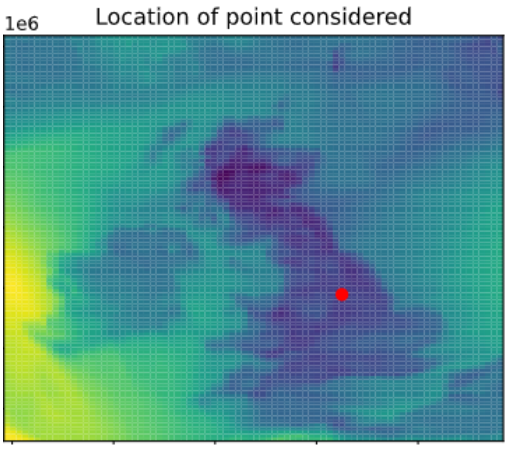

Below we show the “true” maximum gust speed from the Met Office model (blue) over time, plotted against the predicted max gust speed by the emulator (orange).
We evaluated on days 1580 - 2660 (indexed from all the days from 1980-2080), but only plotted 100 days below for clarity.
The orange shaded areas show the 95% credible intervals for the predictions.

Bayesian linear model

MSE: 4.03
IEE: 90%

GP model

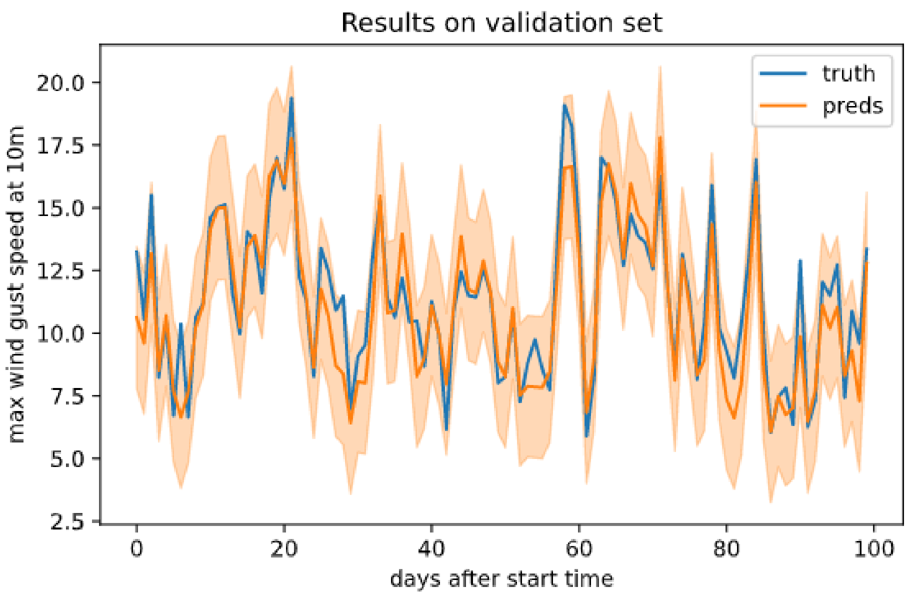

MSE: 2.60
IEE: 88%

Comments:

The Bayesian linear model, based only on mean sustained wind speed, does very well, probably because this variable is strongly weighted in the gust speed calculation in the Met Office model.

The GP model does have a better fit overall (lower MSE), but when it comes to predicting extreme wind events it has a tendency to underestimate the magnitudes of more extreme gusts.
We did not repeat the experiments over different random seeds, or try different train/test splits, to compute the significance of the results.

Both the true and predicted gust speeds at this location do not reach severe gust levels (21 m/s), so the IEE result may not be representative of the performance in more windy locations, or on occurrence of less frequent, more severe gust events (see Appendix A). Ideally, this model would be repeated for a variety of locations.

We further evaluated the GP model on time periods in the future, to confirm the accuracy:

Time period (days after 1980)	**MSE**	**IEE**
20000 - 21080	2.45	92%
25000 - 26080	2.63	93%
34000 - 35080	2.63	97%

Comments: Even though the model was trained in 1980, it appears to make good predictions in the future.

Again, unclear how extreme the gusts were at this location in the future, and we didn't check different seeds or dataset splits. See Appendix A for graphs indicating gust variation over the UK.

Estimating return periods for dangerous gusts (EVT)

- Using Peak-over-Threshold, considering 47mph (21m/s) as the threshold, as according to the MET this denotes gust speed for extreme gales and storms
    - https://www.axa.co.uk/home-insurance/tips-and-guides/what-to-do-before-during-after-high-winds/
    - https://www.abi.org.uk/products-and-issues/topics-and-issues/flooding/storms/
- Fitting a GEV distribution using the inbuilt python library (pyextremes) to the gusts beyond the threshold, and calculating return periods based on that.

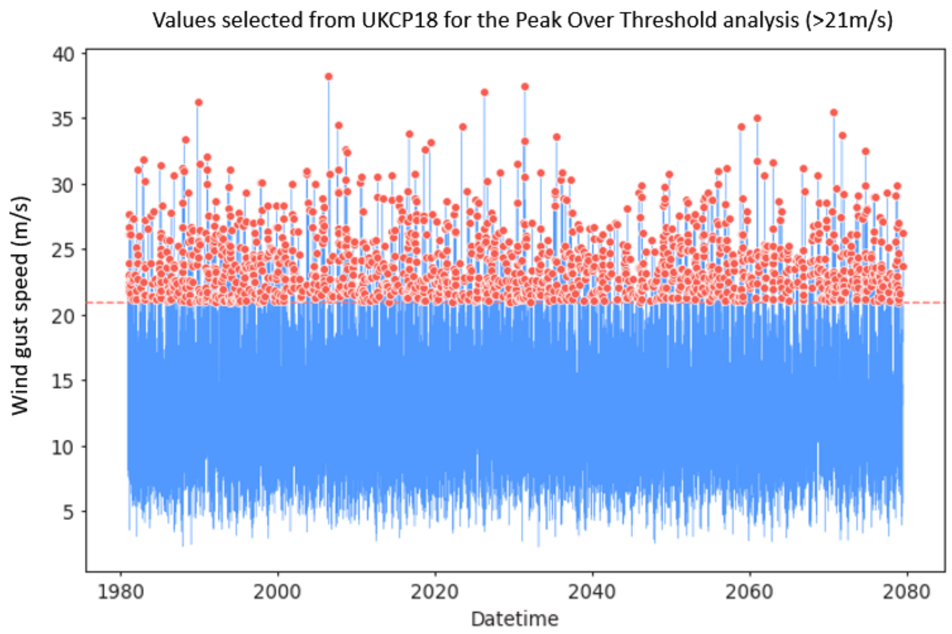

Key for graph above: Values considered ‘extreme’ for the Peak Over Threshold analysis over the UKCP18 dataset.

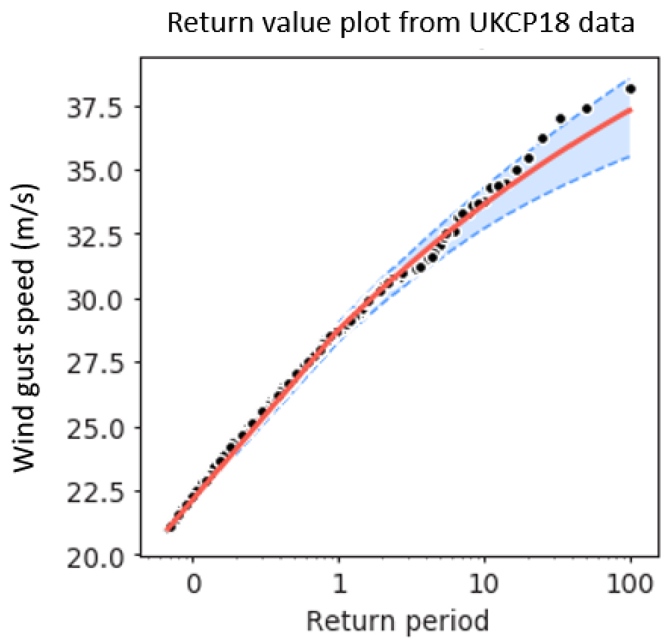

Key for graph above: GEV fit (red line) with a 95% conf. Interval band with points denoting true MET projected values

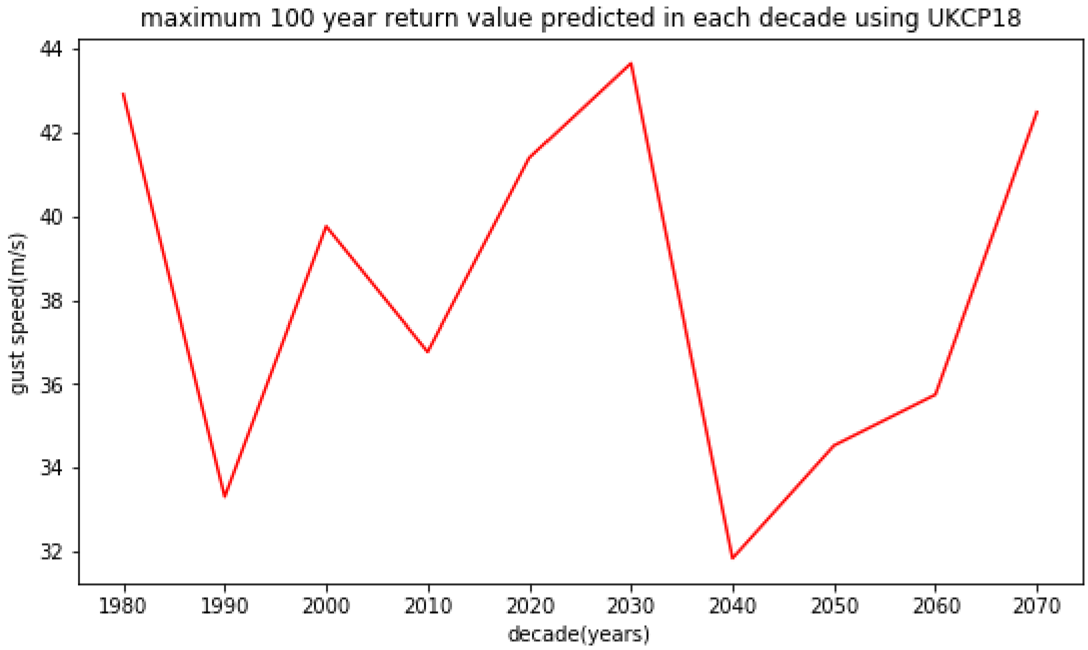

Key for graph above: maximum wind gust speed calculated using Peak Over Threshold on the data of each decade provided by UKCP18.

Comparison of approaches

- Extreme value theory (EVT) is able to handle a wider range of temporal data, which allows it to 'see' less frequent, more severe gust events.
- EVT appears better at identifying extreme values than the emulator, which seems to be able to track the long term trends for wind gusts, but does not track the magnitude as well as the EV density does.
- The selection of approaches depends on the purpose of use. If one is more interested in an insurance application or engineering purposes e.g. how extreme is a wind gust likely to be within the next 50 years, or within the lifetime of an engineered structure, then the EV approach is most likely a better suited approach.

Future Developments

- Experiments across multiple locations (including spatial data)
- Including additional variables such as long-wave flux radiation.
- Formal comparison (e.g. MSE for extreme gusts above 21ms for the emulator) and comparing that with the same metric for the EV fit.
    - Frequency/ Magnitude analysis: Determine how many extreme gust events the emulator identifies (looks like ~90%) and how well it emulates the magnitude of such events vs the EV density.

Appendix A

To identify how representative the training and validation data used are, we plotted the number of gust events of each magnitude (m/s), and number of gust events per x and y location as a histogram (for data 1980-2010).
The left-hand plot shows the total number of gust events of each magnitude. Note that since the data was taken directly from the dataset, this includes gust speed outputs that aren't necessarily large enough to be labelled as gusts.
The centre and right-hand plots show the number of severe gust events (> 21m/s) at each x coordinate, and each y coordinate on the UK 12km grid map. Coordinate points have been transformed to begin at 0. This shows, for example, points on the south coast (y=0) experienced twice the number of severe gust events than average during this time period. This supports our suggestion that any emulator model should be run for each separate location, or region (sharing similar geographical features).