Short-term Prediction and Filtering of Solar Power using State-space Gaussian Processes

Sean Nassimiha*, Peter Dudfield, Jack Kelly**, Marc Peter Deisenroth* and So Takao* (presenting)**
* University College London, ** Open Climate Fix

Tackling Climate Change with Machine Learning: workshop at NeurIPS 2022

Why is solar power nowcasting important?

Nowcasting - Prediction in the very near future. Typically 2~6 hours ahead.

- Adoption of renewable energy sources have increased (which is good!)

However,

- Solar power is intermittent
- Back-up “spinning” reserves emit large amounts of CO_2 [1]

→ Crucial to increase accuracy of predictions with **uncertainties**

UK Photovoltaics (PV) dataset

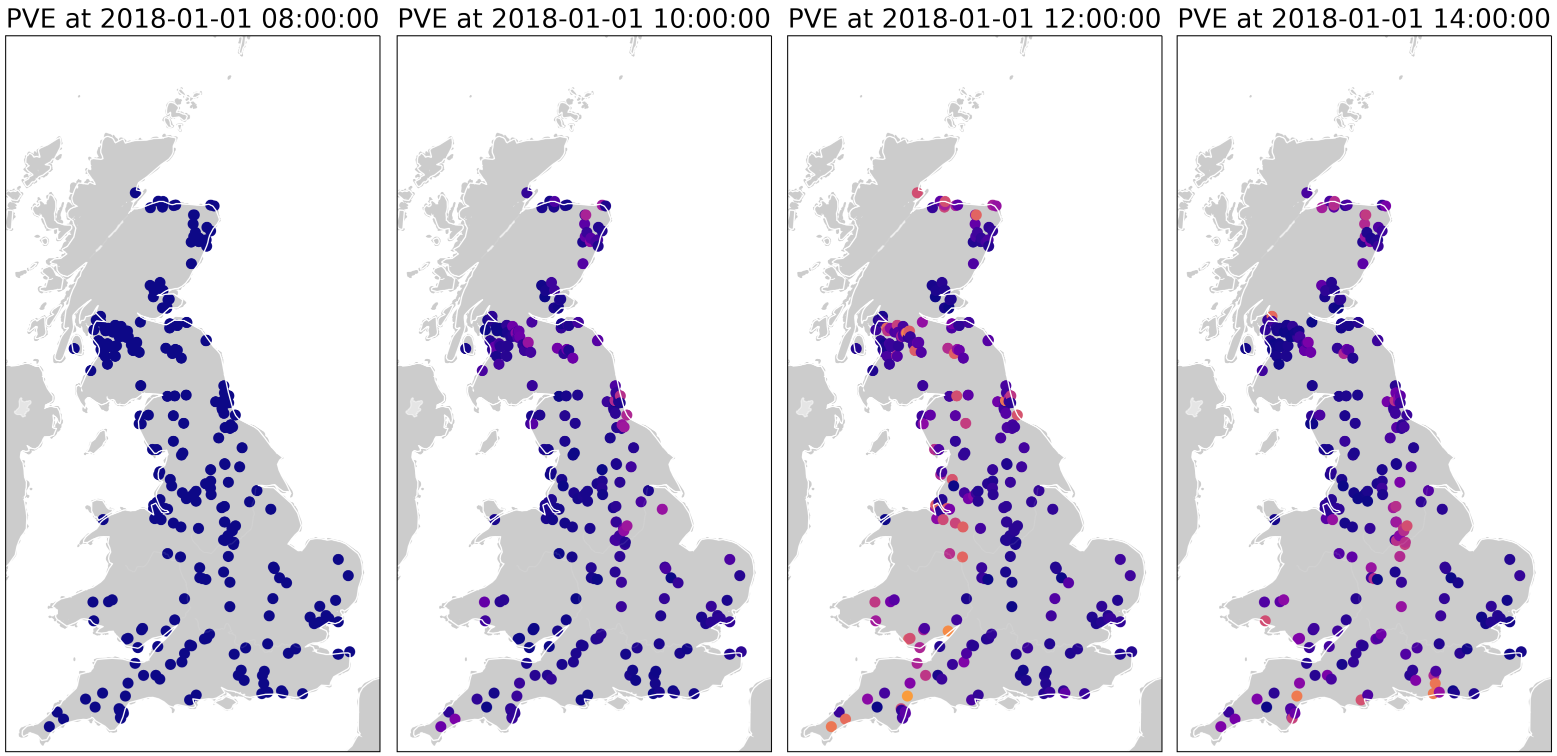

- Access to solar PV readings for 1311 stations across the UK*
- Select data to be between 08:00-16:00

*https://huggingface.co/datasets/openclimatefix/uk_pv

UK Photovoltaics (PV) dataset

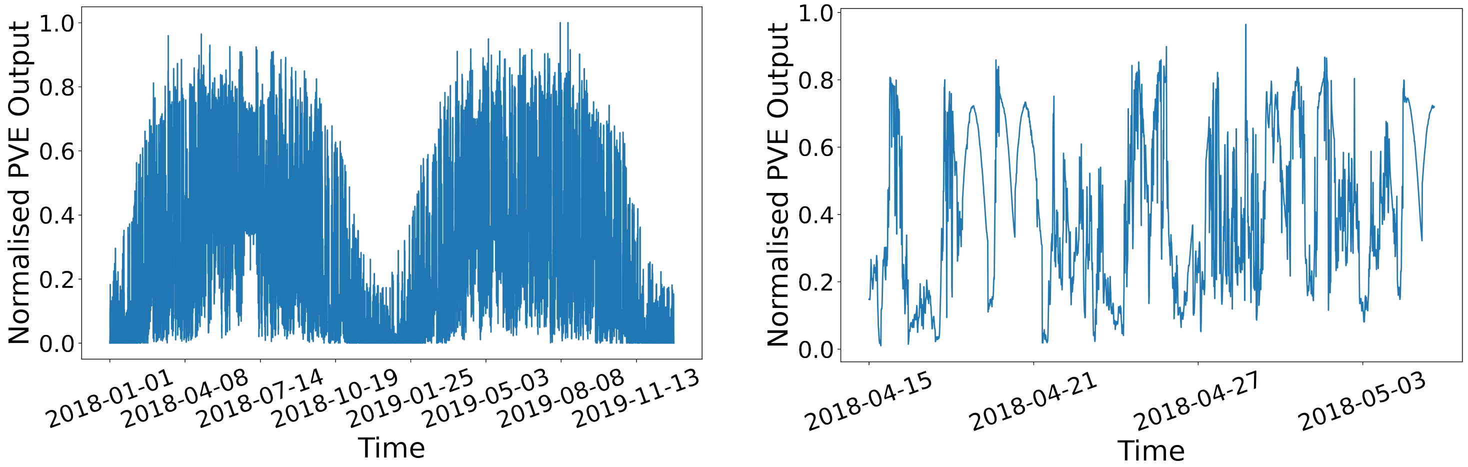

PV timeseries at a single station (annual and daily scales)

- Remove outliers and scale by total capacity

  $\Rightarrow$ Data between 0 and 1

- Timeseries exhibits strong annual and daily seasonality

- $\mathcal{O}(10^5)$ data points

Gaussian Processes

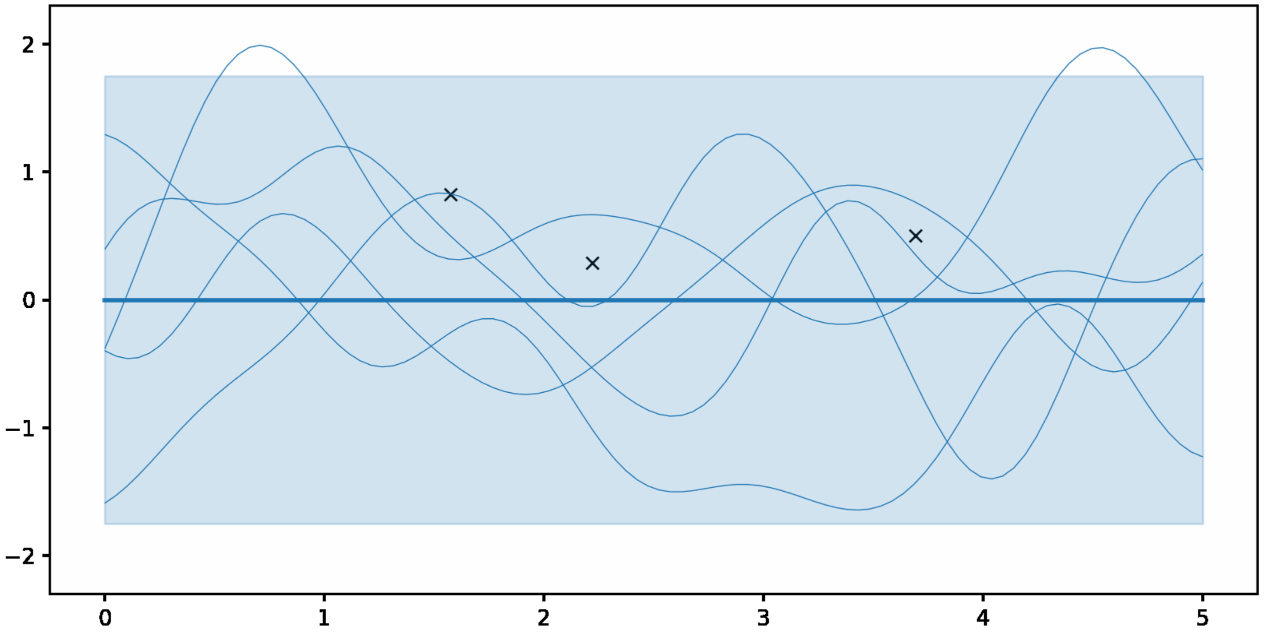

- GPs provide flexible priors for modelling data
- Can get predictions + uncertainty estimates

Gaussian Processes

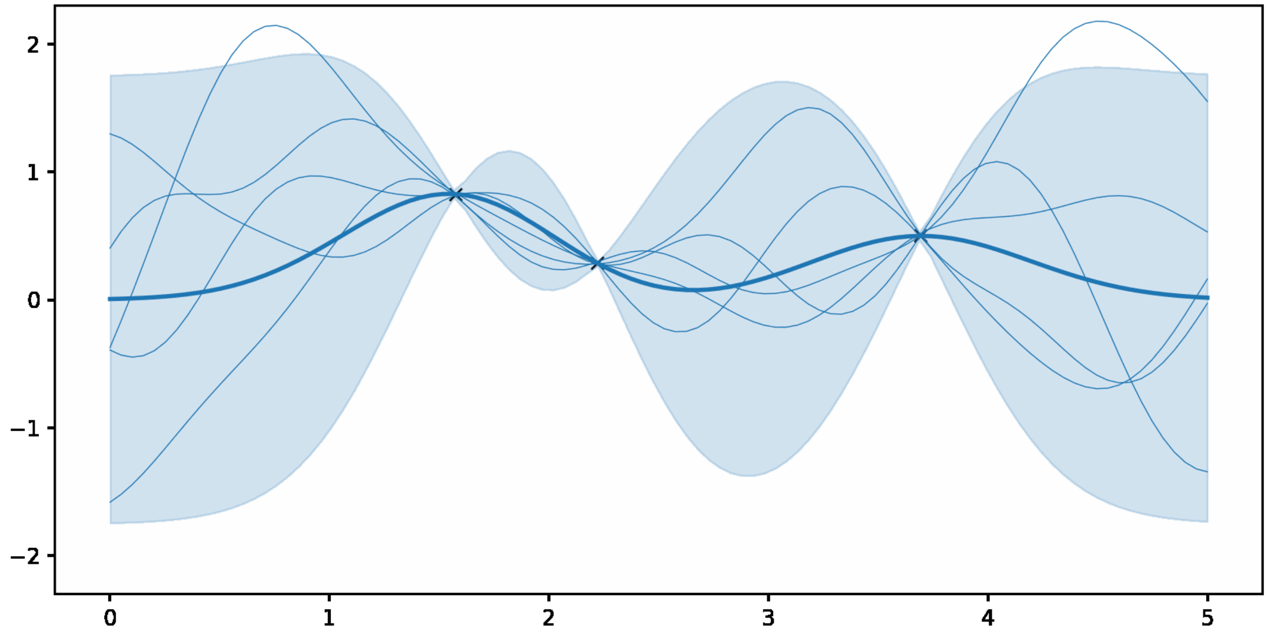

- GPs provide flexible priors for modelling data
- Can get predictions + uncertainty estimates
- However, inference cost is *cubic* in data size
- Use *state-space method* to reduce inference cost [2]

State-space reformulation (Hartikainen et. al. [2])

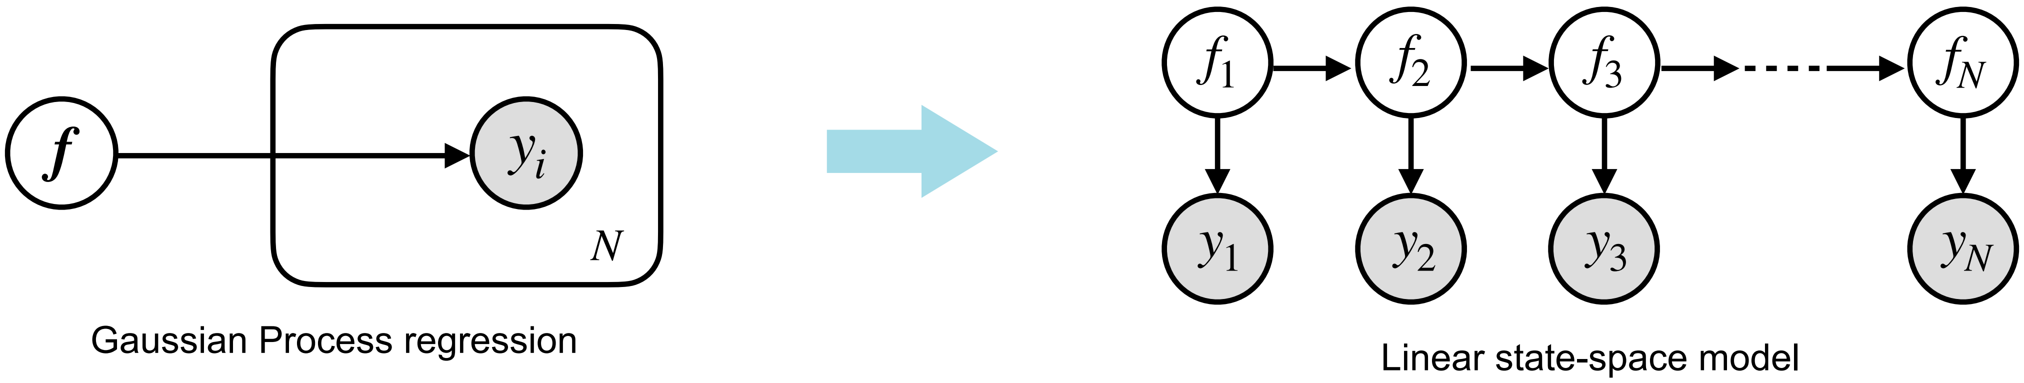

- Use **Kalman filtering** to condition on data as they arrive
- **Linear** cost in data size

Model

- We used a quasi-periodic Matérn-3/2 GP to model the latent states. The kernel is:

$$k(t, t') = k_{\text{matern3/2}}(t, t') + k_{\text{matern3/2}}(t, t')\, k_{\text{periodic}}(t, t').$$

- Use *Beta likelihood*

$$p(y \mid f) = \prod_{i=1}^{N} \mathscr{B}\big(y_i \mid \alpha(f_i), \beta(f_i)\big),$$

to model data in $[0, 1]$.

Results

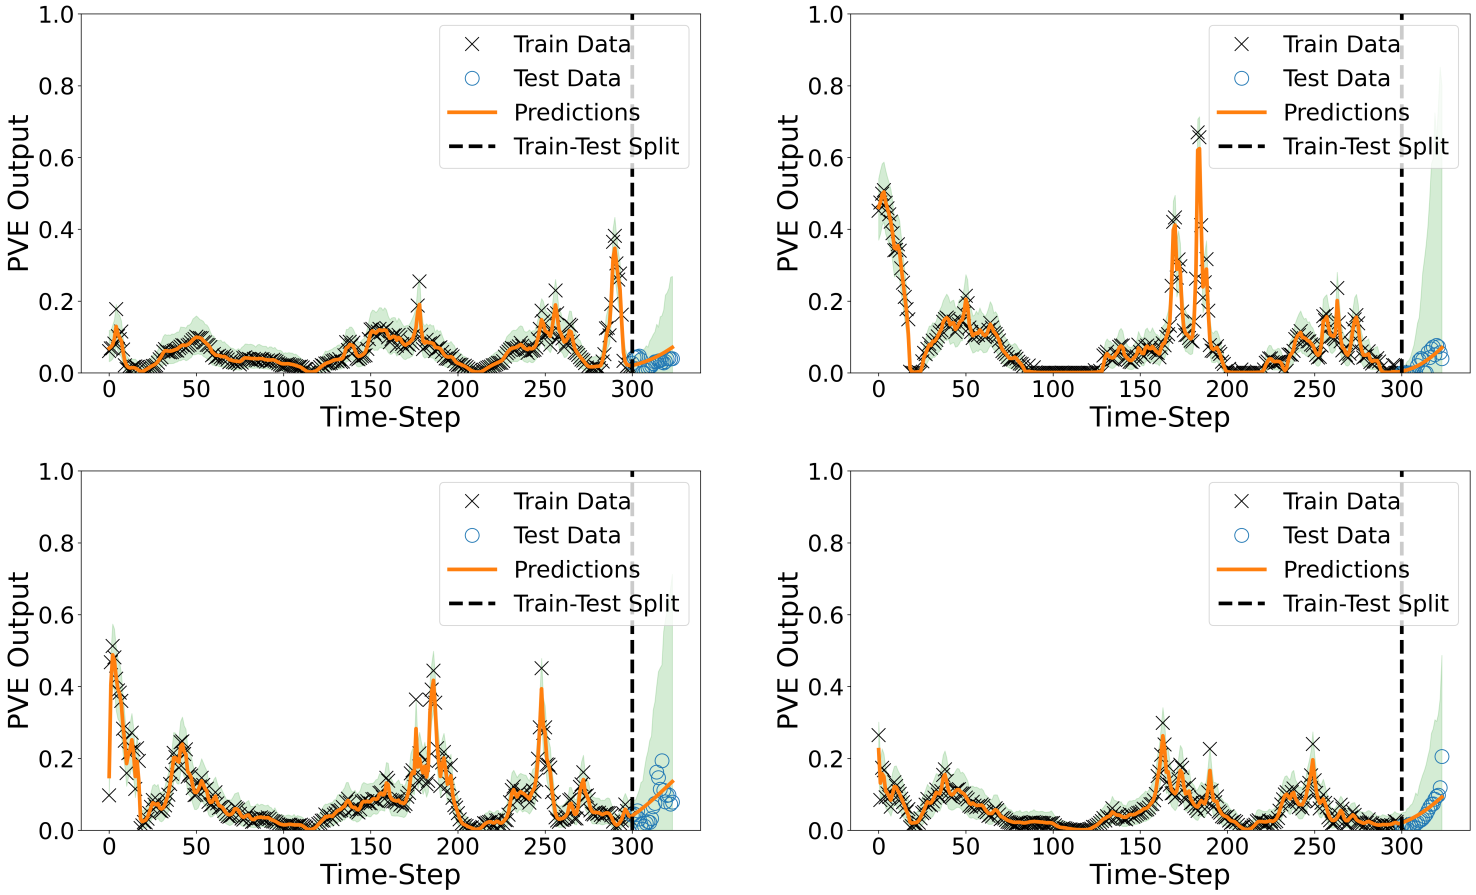

PV predictions on date 2018-01-29 at four random locations

- Surprisingly good predictions on a "calm" day
- **Note:** we've only used the past PV readings as regressors

Results

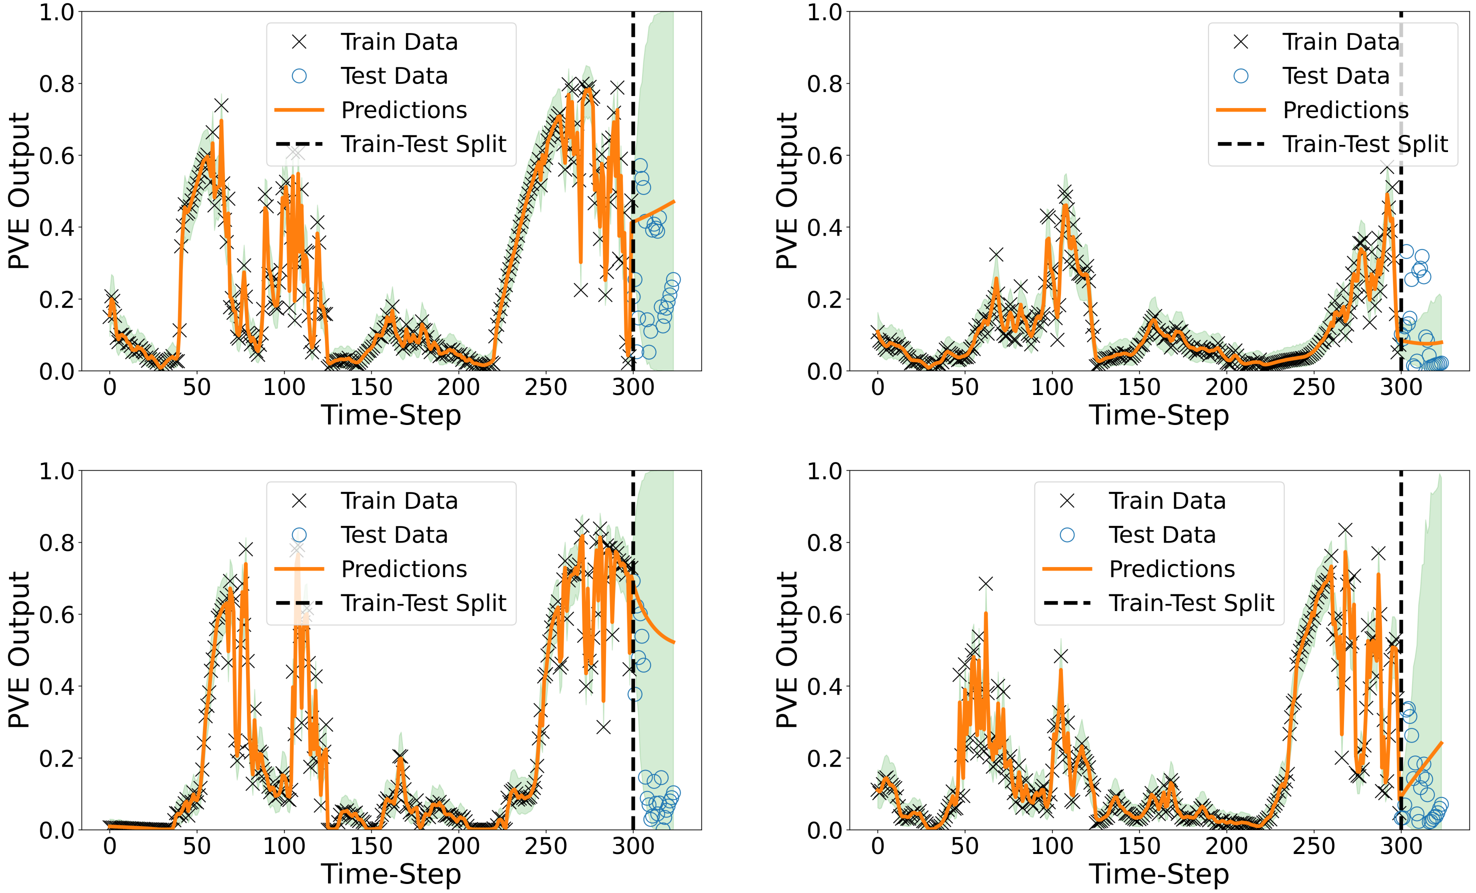

PV predictions on date 2018-02-01 at four random locations

- However, predictions are under-confident when timeseries is highly volatile
- We also considered spatio-temporal GPs to take into account spatial correlations, but achieved worse

Results

Model	MAE $\downarrow$ (mean $\pm$ std)	NLPD $\downarrow$ (median $\pm$ m.a.d.)
Persistence	0.119 ± 0.060	N/A
Yesterday	0.152 ± 0.091	N/A
Hourly smoothing	0.125 ± 0.061	N/A
Simple exponential smoothing	0.117 ± 0.058	-11.1 ± 11.1
Seasonal exponential smoothing	0.110 ± 0.049	-12.2 ± 10.4
Vector autoregression	0.129 ± 0.071	N/A
Our model	$\mathbf{0.109 \pm 0.050}$	$\mathbf{-12.9 \pm 13.8}$

- Results are *marginally* better than several baseline models we tested on
- Metrics used: mean absolute error (MAE), negative log-predictive density (NLPD)

Summary

- Being able to predict solar power production a few hours in advance is incredibly useful for power grid planning, which may help to reduce CO2 emission and costs associated with the use of solar power
- We focussed on the use of Gaussian processes for prediction and quantifying uncertainties
- Using only the PV time series as inputs, we can get good predictions on “calm” days
- However, predictions are under-confident on “noisy” days
- At the current stage though, hard to justify its use over simpler baselines

Discussion

- **Cloud cover**, as well as other weather variables, play an important role in solar power output

  ⇒ Currently exploring how to incorporate this information

- We may also try to "learn" underlying latent dynamics using e.g. Latent Force Models [3] or Neural SDEs [4]

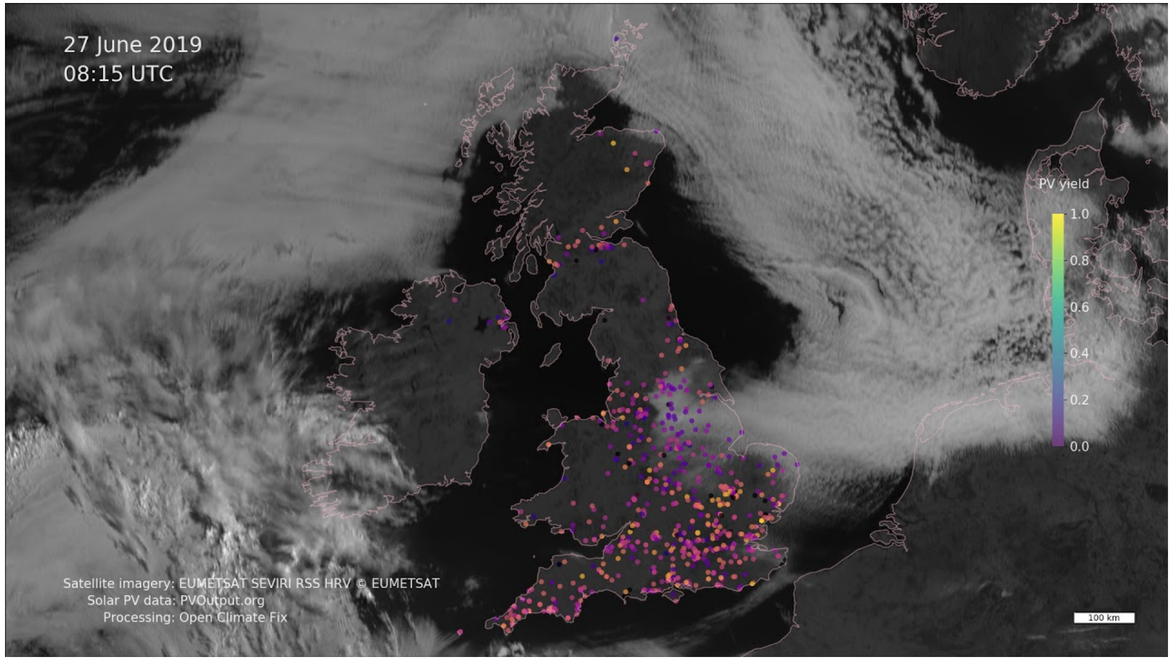

Link: https://www.youtube.com/watch?v=IOp-tj-IJpk

References

[1] Open Climate Fix. Nowcasting. https://www.openclimatefix.org/projects/nowcasting/

[2] Hartikainen, Jouni, and Simo Särkkä. "Kalman filtering and smoothing solutions to temporal Gaussian process regression models." *International workshop on machine learning for signal processing*. IEEE, 2010.

[3] Alvarez, Mauricio, David Luengo, and Neil D. Lawrence. "Latent force models." *Artificial Intelligence and Statistics*. PMLR, 2009.

[4] Li, Xuechen, et al. "Scalable gradients for stochastic differential equations." *International Conference on Artificial Intelligence and Statistics*. PMLR, 2020.