Authors: Patrick Mc Sharry and Georgios Anastasiades

Smith School of Enterprise & the Environment | University of Oxford
Hayes House, 75 George St, Oxford, OX1 2BQ
Tel: +44 1865 614 943 | Fax: +44 1865 614 960
http://www.smithschool.ox.ac.uk/crf

Highlight results

- **Production of an quantile forecast model conditionally to predictibility information**

Background:

Due to the large fluctuations in wind speed caused by the unpredictable and dynamic nature of the earth's atmosphere, there is great variability in wind power production. This inherent variability of wind speed is the main cause of the uncertainty observed in wind power generation. Recently, scientists have been directly or indirectly attempting to model this uncertainty and produce improved forecasts of wind power production. Moreover, recent research has focused on producing probabilistic or density forecasts, because the point forecast methods are not able to fully quantify the uncertainty related to the prediction. Only a fully probabilistic framework will give us the opportunity to model the uncertainty related to the prediction, and avoid the intrinsic uncertainty involved in a point forecasting calibrated model.

Quantile forecasting models:

Our approach aims to produce very short-term wind power quantile forecasts (it has been shown that statistical time series models may outperform sophisticated meteorological forecasts for short lead times within six hours), using only wind power data. The proposed methodology uses a linear quantile regression model with endogenous explanatory variables extracted from univariate wind power time series. More specifically, four novel variability indices (Figure 1) are produced which serve to capture the underlying volatile nature of the wind power series. These indices, together with some lagged versions of the wind power series, are used as explanatory variables in four quantile regression models. The proposed models have the following two advantages:

- Simplicity: Our approach uses only a single wind power time series to produce probabilistic forecasts, without including any exogenous variables (such as wind speed, temperature, atmospheric pressure etc.) Moreover, the relatively few model parameters are easy to optimize using a one-fold cross validation technique.

- Flexibility: The proposed models can be independently optimized to forecast either specific quantiles (for example the tails of the distribution which capture the ‘rare’ events) or the whole predictive density of the wind power generation. Moreover, the methodology focuses on improving the predictions of either the first step ahead or the average over all forecast horizons. Focusing on different quantiles and forecast horizons will result in producing a different set of variability indices.

Evaluation study:

In order to train and evaluate the probabilistic forecast performance of the proposed models we used wind power series from three wind farms in Denmark. We chose one high, one low, and one average variability wind farm, in order to better understand the ability of each model to produce probabilistic forecasts under different circumstances. The models were optimized using two years of wind power data (based on a historical dataset of 2007-2008 with a temporal resolution of 15 minutes). Their quantile and density forecast performances were evaluated using the next two years (2009-2010). Our approach produced 19 different quantile forecasts (for the 5%, 10%,…,95% quantiles) for 24 forecast horizons (from 15 minutes up to six hours ahead), for each wind farm. The models were trained in order to improve the skill (overall quality) of the quantile forecasts using appropriate quantile evaluation scores (quantile loss function). We optimize the models’ parameters to improve (i) the first lead time and (ii) the average over all forecast horizons quantile forecast performance, for each individual quantile. The out-of-sample quantile and density forecast performances were evaluated using both density and quantile evaluation skill scores, for both the first lead time and the average over all forecast horizons.

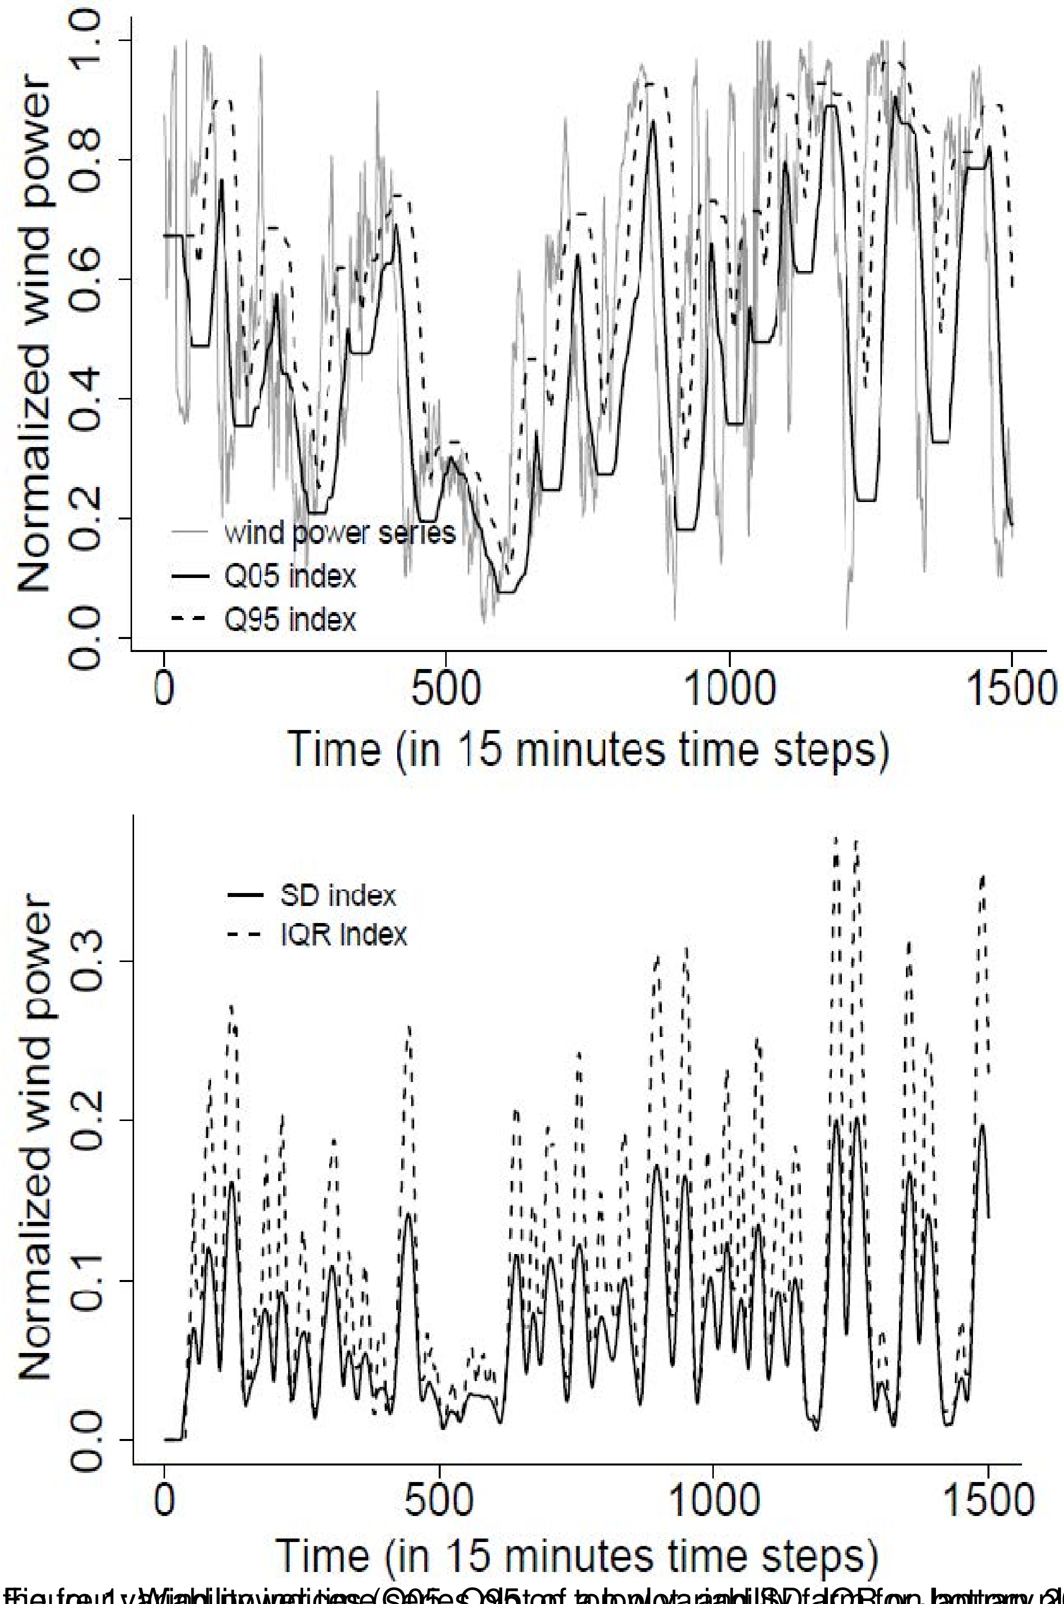

Figure 1: Wind power time series of a wind farm for January 2007, together with the four variability indices (Q05, Q95 on top plot, and SD, IQR on bottom plot)

Out-of-sample results:

The results show that the proposed quantile regression models can outperform existing simple benchmarks such as persistence and unconditional distribution (climatology). When focusing on

forecasting each quantile separately, the performance gain with respect to the above benchmarks is maintained above 53% (for all quantiles) for the models optimized just for the first lead time and above 12% (for 15 out of 19 quantiles) for the models optimized on the average performance over all forecast horizons. When focusing on the whole predictive density, the models optimized just for the first lead time outperform the above benchmarks by at least 58% and the models optimized on the average performance over all forecast horizons by (on average) 12%. A simpler quantile regression model, which consists of only some lagged versions of the wind power series, was also used as a benchmark. This benchmark highlights the importance of using the constructed variability indices.

Further work:

It is expected that substantial emphasis will be placed on the further development of more variability indices that can capture a different spectrum of predictive density. Such indices may be combined by using a non-linear quantile regression model which is expected to improve the skill of the probabilistic forecasts (but also increase the required computational effort).

Bibliography:

[1] Davy, R., Milton, J., Russel, C. & Coppin, P. (2010). Statistical downscaling of wind variability from meteorological fields. Boundary-Layer Meteorol, 135:161-175.

[2] Gneiting, T. & Raftery, A. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102:359–378.

[3] Koenker, R. & Basset, G. (1978). Regression quantiles. Econometrica, 46(10):33-50.

[4] Nielsen, H., Madsen, H. & Nielsen, T. (2006). Using quantile regression to extend an existing wind power forecasting system with probabilistic forecasts. Wind energy, 9:95-108.