Institution: University of Liverpool

Unit of Assessment: 11 – Computer Sciences & Informatics

Title of case study: National Gas Demand Forecasting

1. Summary of the impact (indicative maximum 100 words)

This impact case is based on economic impact through improved forecasting technology. It shows how research in pattern recognition by Professor Henry Wu at the School of Electrical Engineering and Computer Science led to significantly improved accuracy of daily national gas demand forecasting by National Grid plc. The underpinning research on predicting non-linear time series began around 2002 and the resulting new prediction methodology is applied on a daily basis by National Grid plc since December 2011. The main beneficiaries from the improved accuracy (by 0.5 to 1 million cubic meters per day) are UK gas shippers, who by conservative estimates save approximately £3.5M per year. Savings made by gas shippers benefit the whole economy since they reduce the energy bills of end users.

2. Underpinning research (indicative maximum 500 words)

Prediction and forecasting of time-varying quantities is of crucial importance for the operational efficiency of systems in a diverse range of areas, including the financial sector, medicine, manufacturing, transport, weather forecasting, and large-scale energy distribution networks.

Time-series prediction methods use historical data to build models that utilise recent observations of relevant system variables to make forecasts of their future values. Standard prediction methods are global and static; they build a mathematical model that assumes that the time-varying behaviour of the system variables remains the same over the operational life of the system. In other words, they assume that the system has one “mode” of behaviour. However, this assumption is often not realistic because in many real-world applications the relationships between system variables change over time. Although such global predictors make use of all the available time-series data to build a model, they do not differentiate between different system modes, and the resulting conflation results in less effective prediction.

To address these problems, Wu and Lau (Wu’s PhD student) proposed a novel local predictor in their Pattern Recognition paper [1] in 2008. Unlike global predictors, their local predictor produces forecasts by fitting different models for different windows of the historical data. To achieve this, the available historical time-series data is converted, based on the theory of chaotic dynamics, to data in a new (embedding dimension) space where the different modes of dynamics are easier to capture and analyse. In contrast to global predictors, the proposed method opens small local windows, aggregating more data which share similar characteristics in this new space, and fits individual models that best fit the individual prediction characteristics. The localised predictions are more accurate and flexible than the global ones. Furthermore, data complexity is an issue that very often causes forecasting difficulties in real-world problems. The work in [1] employs a powerful data regression methodology based on kernel methods, capable of tackling highly complex (nonlinear) datasets. The proposed method was thoroughly compared with other state-of-the-art methods using benchmark datasets and was found to outperform them.

The above work was further extended by Wu in [2] and [3]. For example, the work in [2] improved the behaviour of local predictors by introducing weighted data prediction. In addition to employing multiple local models, it proposed a method that prioritises the local windows, by using weights to balance how the windows contribute to the prediction.

3. References to the research (indicative maximum of six references)

Publications in the academic literature:

[1] Lau, K.W., Wu, Q.H., 'Local prediction of non-linear time series using support vector regression', Pattern Recognition, Vol.41, No.5, pp.1556- 1564, 2008. http://dx.doi.org/10.1016/j.patcog.2007.08.013

[2] Elattar, E.E., Goulermas, J.Y., Wu, Q.H., 'Electric load forecasting based on locally weighted support vector regression', IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and reviews, Vol. 40, No. 4, July 2010, pp.438-447. http://dx.doi.org/10.1109/TSMCC.2010.2040176

[3] Elattar, E.E., Goulermas, J.Y., Wu, Q.H., 'Generalized locally weighted GMDH for short term load forecasting', IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and reviews, Vol. 42, No. 3, May 2012, pp.345-356. http://dx.doi.org/10.1109/TSMCC.2011.2109378

4. Details of the impact (indicative maximum 750 words)

Gas demand forecasting has been a challenging problem faced by National Grid for many years. National Grid's published daily gas demand forecasts from the National Transmission System (NTS) are used by the market to aid trading and operational decisions in balancing supply and demand. The NTS takes gas from producers and delivers it to consumers. Gas shippers nominate quantities of gas entering and exiting the NTS. National Grid is responsible for the physical transportation of gas through the NTS, and for ensuring the physical balance of the total system (which is important both for efficiency and safety). Each gas shipper is financially responsible for the costs incurred for the management of an imbalance in its supply and demand or a difference between its gas nominations and actual flows.

If the actions of gas shippers are out of balance, then National Grid takes residual balancing actions: it buys or sells gas in the 'On-the-day Commodity Market' (OCM), which is part of the gas trading arrangements introduced in the UK in October 1999. When the NTS is short of gas National Grid tends to force prices up; when the NTS has an oversupply of gas, the price is forced down. These residual balancing trades in the OCM determine marginal system buy and sell prices. At the end of any day, gas shippers that are out of balance are automatically balanced through a 'cash-out' procedure in which the shipper is made to buy or sell the required quantity of gas at the marginal system buy or sell price for that day.

More accurate published gas demand forecasts by National Grid allow gas shippers to pursue actions that are more likely to be balanced. The more balanced the actions of gas shippers, the less balancing actions need to be taken by National Grid, and in turn there is a reduced associated cost for gas shippers. This is important for the economy as a whole because the costs incurred by gas shippers are ultimately passed on to end users, resulting in higher energy bills.

In the future, it will be even more difficult to make accurate gas demand predictions, as the decarbonization agenda promotes a greater dependence on renewable energy sources, which inherently suffer from extreme variability. The intermittent energy generated by these renewable sources directly translates to variability in the demand for traditional energy sources like gas.

In order to improve the accuracy of forecasting, National Grid operates within a framework which includes penalties for inaccurate prediction and incentives for reduction of forecasting errors. As evidence for the large scale and significance of the problem: in April 2013 Ofgem (the regulator for electricity and gas markets in the UK) proposed an incentive scheme with an annual reward of up to £10M for National Grid. The level of the reward will depend on the quality of National Grid's daily gas demand forecasts (see [B] in Section 5).

Earlier methods used for national gas demand forecasting by National Grid were based on global prediction methods, as described above. Forecasts were calculated using a globally fitted model to the historical data using standard regression techniques. The results were often unsatisfactory (in fact, National Grid had to pay significant penalties for inaccurate daily forecasts).

Professor Henry Wu has a long-standing collaborative research relationship with National Grid plc of over 15 years. In July 2011, the joint project “Daily Gas Demand Forecasting” between Wu and National Grid started. The aim of the project was to apply the methodology developed in [1] to daily gas demand forecasting. The project was funded by National Grid plc with cash contribution of £100,115 (see [C] in Section 5). On the National Grid side the project was led by Chris Aldridge (commercial analyst) with a project team of five engineers. National grid collected local and national gas demand data and Professor Wu jointly with postdoctoral researcher Dr Li and PhD student Zhu applied the local predictor techniques developed in [1] to the data provided by National Grid and developed the required software. Initial investigations showed that Wu's approach, which combines local prediction techniques with kernel methods to deal with non-linearities, is well-suited to the problem of forecasting gas demand and showed improvements over National Grid's previous forecasting method.

A live trial of a prototype of the new model ran from October 2011 until December 2011. The results showed an improvement in forecast accuracy (over the old method, which was still used for published forecasts) by 0.5 to 1 million cubic meters (mcm) per day. Daily gas demand varies through the year but is on average around 240mcm. This volume of gas (0.5 to 1 mcm), priced at typical current wholesale gas price of 70 pence per therm, is worth £130k to £260k per day.

The new forecasting method was implemented in December 2011 and since then is in daily operational use. The main direct beneficiaries are gas shippers, since they bear costs from inaccurate forecasts. In 2012 balancing actions were taken on around 200 days, with an average cost of £700k/day. The improvement in performance over previous methods is conservatively estimated at about 5 days less on which balancing actions would have been needed. This corresponds to an estimated saving of £3.5M per annum for gas shippers, and ultimately for end user's energy bills. These economic impacts have arisen directly from the application of the techniques developed in [1].

Based on the first project discussed above, and because of its huge success, a second phase started in October 2012 and will last for two years. It was funded by National Grid plc with cash contribution of £200,443. The project is again led by Professor Henry Wu and Chris Aldridge, where the Liverpool team is responsible for algorithms and software system development and the National Grid team is responsible for daily live trials of the developed software. Within this project, the models now used will be refined, and further economic benefits in the forecasting of gas demand within the transmission network are expected.

5. Sources to corroborate the impact (indicative maximum of 10 references)

[A] Description of the impact of Professor Henry Wu’s research on daily gas demand forecasting can be corroborated in a letter from Commercial Analyst, Gas Incentives and Strategy at National Grid. This contact can also be contacted to verify the impact.

[B] Information about Ofgem incentives for accurate daily gas demand forecast by National Grid Plc in Ofgem provides evidence to corroborate the impact of Professor Henry Wu’s research.

[C] Description of joint project “Daily Gas Demand Forecasting” between Henry Wu and National Grid plc on page 399 of Annual Report 2011/12 “Innovation Funding Incentive: Gas Transmission R& D Programme Detailed Reports” by NationalGrid.