

response to these new challenges, the UK government is aiming to help network operators and suppliers prepare for a low carbon economy through initiatives such as the £500m low carbon network fund (LCNF) and the £1.2bn smart network fund (SNF).

0(n)-0.387167isns na05449Td [(n)329.4662(o)-0.31555213.55383.22(n)03-0.21[(t)-0.139462fa05449Td [(n)328.5

4(t)-0.165609(r)-0.165609dunch g4(342.6812h(o)-36Td [(c)-0.264.39s4(342.8352be)-0.3075[(n)-0.387167(n)-0.38

where $F : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is some metric. In this paper we focus on the absolute p -norm,

$$E_p = \|f - x\|_p = \left(\sum_{i=1}^n |f_i - x_i|^p \right)^{1/p}, \quad (2.2)$$


for some $p \geq 1$ see [8, p. 52]. For example, this type of error includes the Mean Absolute Error (MAE) and the root mean square (RMS) error, which are simply constant multiples of the 1-norm and 2-norm errors respectively.

2.2 The Adjusted Error

For the purpose of developing forecasts to plan the discharging and charging patterns of storage devices on the LV network, it is more important to predict peaks at adrloon

5552(t-0.9836567y)55.32792adr

	Forecast			
Error	F1	F2	F3	F4
Absolute Error	0.82	0.20	0.99	1.00
Adjusted Error (w = 1)	0.82	0.20	0.79	1.00
Adjusted Error (w = 2)	0.82	0.20	0.48	1.00
Adjusted Error (w = 3)	0.82	0.20	0.20	1.00







Figure 4: Panels (a)–(c) correspond to households (a)–(c) respectively. Each panel depicts the daily averages of the 4-norm (Black) and adjusted 4-norm (Gray) errors for the three forecasts.

peaks. We use $w = 3$



forecasts as Poor.

2. Points in the upper-left quadrant represent forecasts whose mean flat forecast error is smaller than the mean 4-norm error forecast but larger than the mean adjusted 4-norm error. Since the small temporal re-alignment has reduced the error compared to the 4-norm error we refer to these forecasts as Good after adjustment.
3. Points in the top right quadrant represent forecasts whose mean 4-norm and mean adjusted 4-norm errors are both smaller than the mean flat forecast errors. We refer to these as Good forecasts.

The plot shows that the AA forecasts (filled circles) are in general superior to the LW forecasts (unfilled circles). The majority of the AA forecast are either good (360) or good after adjustment (208). Only 32 of the AA forecast are poor whereas 225 of the LW forecasts are poor. For the LW method, only 105 are good forecasts and just less than half (270) are good after adjustments. Of the 600 households, the LW forecast only out-performs the AA forecasts for 30 households in the 4-norm but for 46 households in the adjusted 4-norm. In Figure 5 we also include the data for the LW and AA forecasts of households (a), (b) and (c). In terms of our accuracy classification both the LW and AA are good forecasts for household (a) whereas the AA forecast is good after adjustment for households (b) and (c) while the LW forecast is poor for households (b) and (c). The large proportion of forecasts that are good after adjustment are particularly important. If only the 4-norm is used as an accuracy measure then these forecast methods could potentially be mistakenly rejected, despite their improved performance with respect to the adjusted norm.

he,Id s en

In this section we analyse the adjusted error in more detail. The choice of the adjustment window, w , is largely subjective and application specific. In section 3 we chose $w = 3$ based on the assumption that a reasonable forecast of household electrical energy usage should only misplace a peak by a maximum of an hour and a half. Other criteria, such as requiring the forecast to out-perform the flat forecast, can also be used to inform on a suitable adjustment limit. For a given application it may be necessary to consider the error as a function of w , as described in this section, in order to make a more informed decision on the size of the adjustment window. For smart control algorithms utilised in storage devices it is preferable to forecast a peak earlier rather than later, this guarantees that the battery is fully charged and thus able to more efficiently gadju90.4028326(a)5327.769(b)-0.3871(□)-0.0537(e)-0.239462(r)-327.71(t)--0.050728(lie)-0.23




Figure 6: The mean adjusted errors for (a) the AA forecast and (b) the LW forecast for the usage of three different households (a) (solid line), (b) (dotted) and (c) (dashed) as a function of w . The black marker on each line shows where the forecast errors equal the errors of the flat forecast.

being matched to the actuals. We focus on the AA forecast for our analysis, similar results hold for the LW forecast. As we increase w from 0 to 2 there are large decreases in the adjusted error of the forecast for household (a) due to the closeness (within 3 half hours) of the peaks in the forecast and actual usage. Moderate decreases in the forecast errors are also observed for household (c), although even with $w = 20$ (shifts of ± 10 half hours) the errors are relatively large compared with the errors in the forecasts for households (a) and (b). Household (b) has a slow rate of reduction as w increases. As shown in Section 3, the general behaviour of household (b) can be forecasted accurately and so the slow reduction is likely to be due to the matching of the small daily irregularities.

The adjusted error decreases with increasing w but this is likely to simultaneously increase the mean displacement of the forecast positions. Smaller displacement

manage the local networks it is vital that network operators understand how demand is changing and what practical solutions are available. Household smart meters are becoming an integral part of many government's low carbon agenda and many countries aim to have a meter in every home within the next decade. Smart meters provide a valuable opportunity for detailed data analytics and in particular for forecasts at the individual and low voltage substation level. Accurate household level forecasts can also be utilised for planning the smart control of storage devices to reduce peak demands. However, before useful household level forecasts can be developed an appropriate verification measure must be established to assess the accuracy of such forecasts.

In this paper we suggest such a measure for assessing the success of forecasts of volatile and noisy data. A standard treatment of the accuracy is to consider the p-norm of the error, but due to the "double penalty" effect such measures have been shown to be inadequate, especially when attempting to forecast peaks and troughs in the data. Any successful forecast method requires a degree of flexibility in the spatial/temporal positioning of the peaks. Our proposed solution, the adjusted p-norm error, allows for limited permutations of the forfo40Td [-0.166728(o)-356.71(f)0.352478(o)

Incentive³. We also wish to thank Scottish and Southern Energy (SSE) for providing the EDRP data for use in this project. JAW also acknowledges the EPSRC for support of MOLTEN (EP/I016058/1).

A The Averaged Adjustment (AA) Forecast

In this appendix we briefly describe the Averaged Adjustment (AA) forecast as implemented in this report. For clarity, we show how we forecast for one particular day, the other days of the week are forecasted in an analogous way. We assume that we have N daily usage profiles of half hourly resolution of the d^{th} day of the week ($d = 1, \dots, 7$) which we notate $G^{(k)} = (g_1^{(k)}, g_2^{(k)}, \dots, g_{11}^{(k)})^T$ for $k = 1, 2, \dots, N$, where $G^{(1)}$ is the previous week usage of the d^{th} day and $G^{(2)}$ is the usage over the d^{th} day from 2 weeks before etc. We create a base profile $F^{(1)} = (f_1^{(1)}, f_2^{(1)}, \dots, f_{11}^{(1)})^T$

[20] Alexander Schrijver. Combinatorial Optimization: Polyhedra and Efficiency. Springer, 2002.