Environmental Modelling & Software 66 (2015) 87–97

Contents lists available at ScienceDirect

Environmental Modelling & Software

journal homepage: www.elsevier.com/locate/envsoft

A probabilistic methodology for quantifying, diagnosing and reducing model structural and predictive errors in short term water demand forecasting

Christopher J. Hutton [a,*], Zoran Kapelan [b]

[a] Water and Environmental Management Research Centre, Department of Civil Engineering, Queen's School of Engineering, University of Bristol, Queen's Building, University Walk, Bristol BS8 1TR, UK
[b] Centre for Water Systems, College of Engineering, Mathematics and Physical Sciences, University of Exeter, UK

ARTICLE INFO

Article history:
Received 5 July 2014
Received in revised form
16 December 2014
Accepted 21 December 2014
Available online

Keywords:
Water demand
Forecast
Model calibration
Uncertainty
Bayesian
Real time

ABSTRACT

Accurate forecasts of water demand are required for real-time control of water supply systems under normal and abnormal conditions. A methodology is presented for quantifying, diagnosing and reducing model structural and predictive errors for the development of short term water demand forecasting models. The methodology (re-)emphasises the importance of posterior predictive checks of modelling assumptions in model development, and to account for inherent demand uncertainty, quantifies model performance probabilistically through evaluation of the sharpness and reliability of model predictive distributions. The methodology, when applied to forecast demand for three District Meter Areas in the UK, revealed the inappropriateness of simplistic Gaussian residual assumptions in demand forecasting. An iteratively revised, parsimonious model using a formal Bayesian likelihood function that accounts for kurtosis and heteroscedasticity in the residuals led to sharper yet reliable predictive distributions that better quantifies the time varying nature of demand uncertainty across the day in water supply systems.

1. Introduction

Understanding natural variability in urban water demand, the fundamental aleatory uncertainty affecting water supply systems (Hutton et al., 2014a), helps water utilities to satisfy consumer demand, whilst at the same time allowing them to try and minimise the costs associated with supplying sufficient water. Over decadal scales estimates of future water demand support strategic planning, allowing utilities to understand potential water shortages relating to climatic changes, and make capital investments in the water distribution and treatment infrastructure to meet future demand (Qi and Chang, 2011; Almutaz et al., 2013). At shorter timescales predicted water demand up to several days ahead forms a key input to near real-time control systems, and can contribute towards the reduction of energy consumption and cost associated with supplying water in distribution networks (Martinez et al., 2007; Bakker et al., 2013a). Furthermore, short term predictions of urban water consumption are important for burst detection, helping utilities to distinguish between actual demand and non-revenue water (Mounce et al., 2010).

Existing short-term Water Demand Forecasting (WDF) research (e.g. < 48 h), has mainly focussed on two aspects of the forecasting problem: identification of the best inputs to predict future demand – both endogenous and exogenous variables – and on identifying the best model structures to map these input variables to predict future demand (Adamowski, 2008; Herrera et al., 2010). Relatively few approaches, however, have attempted to quantify the uncertainty in demand forecasts over shorter timescales (Cutore et al., 2008), despite the fact that water demand is highly uncertain due to: (a) a range of difficult to constrain socio-demographic and economic factors known to affect water consumption (Arbués et al., 2003), which themselves vary both spatially and temporally; (b) the fact that residential demand is often not fully metered (e.g. <40% properties in the UK). Even when properties are metered, often they are not read frequently enough to quantify short term demand fluctuations. Demand uncertainty needs to be quantified adequately as it will propagate adversely to affect the accuracy of subsequently derive models, forecasts and control decisions (Hutton et al., 2014a, 2014b). The relative performance of demand

* Corresponding author.
E-mail addresses: c.j.hutton@bristol.ac.uk (C.J. Hutton), z.kapelan@exeter.ac.uk (Z. Kapelan).

http://dx.doi.org/10.1016/j.envsoft.2014.12.021
1364-8152/

forecasting models has typically been evaluated and compared with reference to global metrics of model performance such as Root Mean Square Error (RMSE; e.g. Herrera et al., 2010) that summarise model performance in an average sense over the whole dataset. Such metrics reveal little information about how a model performs poorly, where the key errors in model performance lie, and therefore provide little guidance upon how models may be improved. Furthermore, the statistical assumptions upon which demand forecasting model calibration that employs metrics such as RMSE is typically based (e.g. independent, and identically distributed (iid) Gaussian errors) are seldom reported, and therefore evaluated, in the demand forecasting literature. This is despite the fact that it is on the validity of these statistical foundations that the legitimacy of any model comparison is based.

Building upon the work of Hutton et al. (2014a), who presented a framework for considering the cascade of uncertainty from model calibration, through forecasting, to real time control in Water Supply Systems, this paper presents a probabilistic methodology for the development and calibration of short water demand forecasting models. The methodology is designed to develop more reliable short term WDF models, and provide quantitative information on model predictive uncertainty to the decision maker. A Bayesian approach is applied for model parameter calibration, and subsequent posterior predictive uncertainty quantified probabilistically. The framework emphasises the iterative application of residual error analysis during calibration, and evaluation of the reliability and sharpness of the predictive distributions in order to diagnose errors within the model structure and errors in the residual error assumptions made during calibration. Section 2 reviews short term water demand forecasting; Section 3 presents the overall methodology, followed by a case study implementing the methodology to forecast demand for 3 District Meter Areas in the UK (Section 4); Sections 5 and 6 then discuss and conclude the paper, respectively.

2. Short term water demand forecasting and model development

Short term WDF modelling research has generally focussed on identifying the best model inputs, and on identifying the best models to combine these inputs and map them to predict future water demand. Approaches have applied either endogenous variables – e.g. past values of water demand (Alvisi et al., 2007; Cutore et al., 2008; Romano and Kapelan, 2014) – and/or exogenous variables such as temperature and precipitation (Zhou et al., 2002; Herrera et al., 2010; Adamowski, 2008). Unless past weather variables are used, temperature and precipitation variables may need to be forecasted as inputs to the demand forecasting model, which will contain additional uncertainty. Furthermore, as pointed out by Bakker et al. (2013b) it may be difficult to include weather variables reliably in a practical setting due to reliance on external systems.

A number of different data driven modelling approaches have been applied for short term WDF including multi-linear regression (MLR), Autoregressive (Integrated) Moving Average models (AR(I)MA; Adamowski, 2008; Zhou et al., 2002), and non-linear methods including multiple non-linear regression (MNLR; Adamowski et al., 2012), Artificial Neural Networks (ANNs; Romano and Kapelan, 2014) and variants thereof including dynamic ANNs (Ghiassi et al., 2008), Wavelet transform (WA-) ANNs (Adamowski et al., 2012), and Support Vector Machines (SVM; Herrera et al., 2010). A final class of models that may be considered more heuristic in approach have structures built upon observations made from exploratory data analysis. Such models share similarities with ARMA approaches, and generally include a component representing the average behaviour of the system, such as an average of past water demands (Herrera et al., 2010), and a persistence component representing local deviations in time, which may be represented through regression on recent prediction errors (Alvisi et al., 2007). Bakker et al. (2013b) applied a heuristic approach in which normalised water demands are used as input variables, and combined with multipliers for the specific day of the week and time of day to derive the forecast.

A number of papers have conducted comparative analysis between different data driven models. Adamowski et al. (2008) found that ANNs outperform linear regression and ARIMA models for peak daily water demand forecasting. SVM models have also be found to outperform 5 other model structures for 1 h ahead demand forecasts (Herrera et al., 2010), whilst WA-ANNs have been found to outperform MLR, MNLR, AIRMA and ANN models for daily water demand forecasting (Adamowksi et al., 2012). However, in this latter approach wavelet transformed data could also be applied as input to other model types. Whilst it is difficult to compare different WDF methodologies in different contexts, Mean Absolute Percentage Errors (MAPE) reported in the literature generally vary from 3 and 10% for lead times up to 24 h (Bakker et al., 2013; Romano and Kapelan, 2014; Alvisi et al., 2007), where the lowest errors reported by Bakker et al. (2013b) were found in the larger supply zones where deviant behaviour from the norm is more likely to be masked by average behaviour.

The relative performance of different WDF models has been judged mainly with reference to global metrics of model performance like MAPE, including Root Mean Square Error (RMSE), and Mean Absolute Error (MAE; Ghiassi et al., 2008; Herrera et al., 2010). Such metrics, however, provide limited scope for comparative analysis as they collapse all residual error information into a single value, and can therefore only tell us how good models are in an average sense. Gupta et al. (2008) argue that such metrics are therefore weak in a diagnostic sense, as they reveal little information about how and where within a simulation a model performs poorly. Such metrics therefore provide limited information to determine between competing models, and to guide subsequent model improvement.

To overcome the problems of model evaluation solely with global metrics, further investigation of the residual errors is required. Such exploration is important for two reasons. First, it is important to test the assumptions of Gaussianity, heteroscedasticity and independence of residual errors that are (implicitly) assumed during the model fitting exercise (Engeland et al., 2005). This is particularly important as it is on these assumptions that the validity of the model fit, and in turn the validity of any subsequent model comparison, is based. Second, context specific residual error analysis helps to identify how and where a model performs poorly. Residual analysis, however, is not routinely applied (or at least not fully reported) in the literature during WDF model development.

A further need to analyse in more detail the residual errors is that urban water demand is highly uncertain due to limited spatial and temporal metering coverage, and also because of a range of factors that influence water consumption (and leakage), which themselves vary spatially and temporally (Arbués et al., 2003). Water demand uncertainty is also the key aleatory uncertainty that propagates into, and influences that accuracy of Water Distribution System model predictions (Hutton et al., 2014a). However, despite this uncertainty, and despite the wider application of uncertainty quantification methods in Urban Water Systems' modelling (Kapelan et al., 2007; Alvisi and Franchini, 2010; Hutton et al., 2014a, 2014b; Breinholt et al., 2012; Deletic et al., 2012) and hydraulic/hydrological modelling more generally (Liu and Gupta, 2007; Beven, 2008), few approaches have moved beyond deterministic WDF modelling. In a forecast setting, where models are to