ADAPTIVE WATER DEMAND FORECASTING FOR NEAR REAL-TIME MANAGEMENT OF SMART WATER DISTRIBUTION SYSTEMS

Michele Romano and Zoran Kapelan

Abstract

This paper presents a novel methodology to perform adaptive Water Demand Forecasting (WDF) for up to 24 hours ahead with the aim to support near real-time operational management of smart Water Distribution Systems (WDSs). The novel WDF methodology is exclusively based on the analysis of water demand time series (i.e., demand signals) and makes use of Evolutionary Artificial Neural Networks (EANNs). It is implemented in a fully automated, data-driven and self-learning Demand Forecasting System (DFS) that is readily transferable to practice. The main characteristics of the DFS are: (a) continuous adaptability to ever changing water demand patterns and (b) generic and seamless applicability to different demand signals. The DFS enables applying two alternative WDF approaches. In the first approach, multiple EANN models are used in parallel to separately forecast demands for different hours of the day. In the second approach, a single EANN model with a fixed forecast horizon (i.e., one hour) is used in a recursive fashion to forecast demands. Both approaches have been tested and verified on a real-life WDS in the United Kingdom (UK). The results obtained illustrate that, regardless of the WDF approach used, the novel methodology allows accurate forecasts to be generated thereby demonstrating the potential to yield substantial improvements to the state-of-the-art in near real-time WDS management. The results obtained also demonstrate that the multiple-EANN-models approach slightly outperforms the single-EANN-model approach in terms of WDF accuracy. The single-EANN-model approach, however, still enables achieving good WDF performance and may be a preferred option in engineering practice as it is easier to setup/implement.

Keywords: Water Demand Forecasting, Artificial Neural Network, Evolutionary Algorithm.

Michele Romano, Centre for Water Systems, College of Engineering, Mathematics and Physical Sciences, University of Exeter, Harrison Building, North Park Road, Exeter, Devon, EX4 4QF, United Kingdom, Email: mr277@exeter.ac.uk **(corresponding author)**

Zoran Kapelan, Centre for Water Systems, College of Engineering, Mathematics and Physical Sciences, University of Exeter, Harrison Building, North Park Road, Exeter, Devon, EX4 4QF, United Kingdom, Email: Z.Kapelan@exeter.ac.uk

1. INTRODUCTION

Water Demand Forecasting (WDF) is an important issue for water companies worldwide. It provides the basis for making operational, tactical and strategic decisions (Billings and Jones, 2008; Gardiner and Herrington, 1990) and can help to improve the performance of a Water Distribution System (WDS) by anticipating the corresponding system operation. However, forecasting water demand is a challenging task. Indeed, demand patterns from a WDS show daily, weekly and seasonal variations and are also influenced by socioeconomic and meteorological factors such as population characteristics or number of industrial establishment and air temperature or precipitations. The difficulties encountered because of the great variability of these factors have engendered a plethora of studies in an attempt to produce reliable demand forecasts (Donkor et al., 2012).

The variety of methods that have been proposed for modelling and forecasting water demand patterns can be broadly classified into linear and nonlinear (Zhang, 2001). Examples of linear methods are univariate time series analysis - such as exponential smoothing and autoregressive integrated moving average models, and linear regression models (e.g., Hughes, 1980; Anderson et al., 1980; Maidment et al., 1985; Zhou et al., 2000; Alhumoud, 2008). Examples of nonlinear methods are nonlinear regression models, bilinear models, threshold autoregressive models, Artificial Neural Network (ANN) or ANN-based models, fuzzy logic, extended Kalman filter and genetic programming, and model trees (e.g., Jain et al., 2001; Kim et al., 2001; Jain and Ormsbee, 2002; Bougadis et al., 2005; Altunkaynak et al., 2005; Cutore et al., 2008; Nasseri et al., 2011; Bennett et al., 2013a). The linear methods have been widely used because they are easy to develop and implement, in addition to being simple to understand and interpret. However, water demand data have varying degrees of nonlinearity, which may not be adequately handled by the linear methods. In this scenario, the nonlinear methods (especially those that make use of a data-driven approach such as ANNs) can help improving the WDF performance. For example, Jain et al. (2001), Jain and Ormsbee (2002) and Bougadis et al. (2005) observed that ANN models outperform regression and univariate time series analysis. Similarly, Adamowski (2008) developed and compared relative performance of: (i) 39 multiple linear regression models, (ii) 9 autoregressive integrated moving average models and (iii) 39 ANN models; his study concluded that the latter perform the best.

With regard to the forecast horizon (i.e., how far into the future demand is to be predicted) and to the periodicity (i.e., the time span between consecutive forecasts) of the forecast (Donkor et al., 2012), three main types of WDF can be discerned. These are: (i) the long-term WDF, (ii) the medium-term WDF and (ii) the short-term WDF (Gardiner and Herrington, 1990). The long-term WDF (i.e., annual forecasts for ten years or more) is useful for making strategic decisions on issues such as WDSs capacity expansion and design of new WDSs. The medium-term WDF (i.e., monthly to annual forecasts for one to less than ten years) is useful for making tactical decisions on issues such as extensions of existing WDSs and investment planning. The short-term WDF (i.e., hourly to monthly forecasts for up to one year) is useful for making operational decisions on issues such as WDS management and optimisation. Taking this into consideration, it is clear that water companies would benefit from methods that enable performing WDF for all the forecast horizons in the short-term to long-term range (i.e., from 1 hour to 20-30 years). However, WDF generally aims at the evaluation of water management and savings policies (see – e.g., White and Fane, 2002) and research in the WDF field has mainly focused on satisfying the mandate of water utilities to maintain a reliable supply of potable water to the customers and to ensure that this level of reliability is maintained in future years. Therefore, methods aimed at supporting near real-time WDS management tasks such as on-line pump scheduling and dynamic hydraulic modelling (i.e., forecast horizons from 1 to 24 hours) have received comparatively less attention (Donkor et al., 2012; Bakker et al., 2013).

The novel WDF methodology developed and presented in this paper aims to support the near real-time management of smart WDS. The term "smart" is used here to indicate a WDS where several data technologies (such as a data-driven WDF system) help to operate the WDS (SWAN, 2014). This methodology is exclusively based on the analysis of observed demand data and makes use of the Evolutionary Artificial Neural Networks (EANNs) to predict water demand for up to 24 hours in the future. The water demand data (i.e. time series of historical demands) are, in the general case, estimated by mass balance analysis - i.e., from inflow/outflow signals and storage volume changes in the studied WDS subsystem. EANNs are biologically-inspired computational models that use Evolutionary Algorithms (EAs) (see – e.g., Holland, 1975; Schwefel, 1995; Koza, 1992) in conjunction with ANNs. In this framework, EAs are often used for designing ANN models. Common approaches involve performing tasks such as connection weight optimisation (e.g., Keesing and Stork, 1991), architecture optimisation (e.g., Harp et al., 1991), parameter optimisation (e.g., Castillo et al., 2000) and input data selection (e.g., Reeves and Taylor, 1998).

A comprehensive (but not so recent) review of the different interactions/combinations between EAs and ANNs that have been proposed is given in Yao (1999). With particular regard to the water resources planning and management field, although EAs have been extensively applied to solve a wide range of problems (see – e.g., Nicklow et al., 2010), EANN applications are scarce. Noteworthy are, therefore, the works presented in Giustolisi and Simeone (2006) and Romano et al. (2013). Giustolisi and Simeone (2006) made use of EANNs for groundwater level prediction. Romano et al. (2013) made use of EANNs for forecasting pressure/flow values 15 minutes in the future and hence enable detection of pipe bursts and other events in WDSs. The key advantages of the EANNs identified in many of the aforementioned studies include: (i) their remarkable adaptability to dynamic environments (i.e., EANNs can adapt to an environment as well as to changes in the environment) and (ii) the fact that they dramatically reduce the effort required from a human expert to design an ANN model for a given problem whilst enabling replicating or outperforming the quality of the results achievable through human expert intervention.

The main advantage of the novel demand forecasting methodology presented here over aforementioned approaches is its self-learning ability which helps to adapt to the ever changing operating conditions in the WDS. This is of fundamental importance because, as stressed by Bakker et al. (2013), existing WDF approaches need to be improved with respect to adaptive functionality. The second advantage concerns the robustness of the WDF models building process. Indeed, some of the existing methods (e.g., Jain et al., 2001; Jain and Ormsbee, 2002; Bougadis et al., 2005) involve an arbitrary selection and use of various explanatory variables and/or different lags of the demand variable. This, in turn, provides a less rational basis for the inclusion of such variables in the WDF models (Donkor et al., 2012). The third advantage concerns the practicality of the methodology operationalisation - i.e., its application in an on-line environment. Unlike some of the methods proposed in the literature (e.g., Goodchild, 2003; Coomes et al., 2010), the methodology presented here does not make use of many or ad hoc explanatory variables, which pose the greatest challenge to practice in terms of collecting and keeping track of the data (Donkor et al., 2012; Bakker et al., 2013). Having said this, it is important to stress that the parsimoniousness of the EANN models resulting from the application of this methodology does not negatively affect the forecast quality.

This paper is organised as follows. After this introduction, the Demand Forecasting System (DFS), which implements the novel WDF methodology presented in this paper, is described in the methodology section. Specifically, an overview of the DFS is given first. This is then followed by four sub-sections presenting the theoretical background and methodological details of the various data analyses performed by the DFS. Once this is done, the case study section presents the results of the DFS tests on demand time series (i.e., signals) from three District Metered Areas (DMAs) and a Water Supply Zone (WSZ) in the United Kingdom (UK). The DMAs are WDS subsystems isolated from the rest of the WDS by closing appropriate (boundary) valves. Flows in and out of the DMA are normally fully metered. The WSZs are larger WDS subsystems containing a number of DMAs which are supplied either by a single water source or a group of water sources blended within service reservoirs. Finally, the main conclusions are drawn and acknowledgements given. Several abbreviations are used in this paper. A list of these abbreviations can be found in Table 1.

Table 1. List of abbreviations.

ANN	Artificial Neural Network
DFS	Demand Forecasting System
DMA	District Metered Area
DoW	Day of the Week
e	Ensemble
EA	Evolutionary Algorithm
EANN	Evolutionary Artificial Neural Network
FP&IS	Fixed Parameters & Input Structure
MAPE	Mean Absolute Percentage Error
MSE	Mean Square Error
NAN	Not A Number
NS_{Index}	Nash-Sutcliffe index
r	Recursive
ToD	Time of the Day
U	With Updating
UK	United Kingdom
WDF	Water Demand Forecasting
WDR	Weight Decay Regularisation
WDS	Water Distribution System
WSZ	Water Supply Zone
WoutU	Without Updating
YWS	Yorkshire Water Services

2. METHODOLOGY

2.1. DFS overview

The DFS presented here enables performing short-term (i.e., up to 24 hours in the future) demand forecasting by using two alternative approaches. In the first approach, a total of g EANN models (e.g., g=24 - if hourly demand values are considered) are built (i.e., trained and tested) and used. Each EANN model in this ensemble has a different forecast horizon (e.g., 1 hour ahead, 2 hours ahead, etc.), thus it predicts the water demand at a particular Time of the Day (ToD) (e.g., at 1 a.m., at 2 a.m., etc.). The WDF for the next 24 hours is performed by running the resulting g EANN models in parallel. In the second approach, only one EANN model is built and used. This EANN model has a fixed forecast horizon (e.g., 1 hour ahead) and WDF for the next 24 hours is performed by using it in a recursive fashion. Note that, hereafter, the acronym *e*EANNs, where "*e*" stands for "*ensemble*", will be used to refer to the first approach. Similarly, the acronym *r*EANN, where "*r*" stands for "*recursive*", will be used to refer to the second approach.

Figure 1 provides a diagrammatic representation of the DFS. This figure shows that the DFS consists of four main components: (i) the data pre-processing module, (ii) the ANN optimisation module, (iii) the ANN building module and (iv) the WDF module. For the specific demand signal being analysed, the data pre-processing module prepares the raw data in order to facilitate/improve the EANN model(s) building process and hence achieve more accurate WDF. For the specific demand signal being analysed and for each particular forecast horizon being considered, the ANN optimisation module automatically selects the optimal ANN input structure (e.g., number of past demand values to be used and additional explanatory variables to be used) and ANN parameters (e.g., number of hidden neurons and number of training cycles), all with the aim to obtain the best possible WDF performance. Finally, for the specific demand signal being analysed and for each particular forecast horizon being considered, the ANN building and WDF modules are used to develop the actual EANN model (by using the optimised input structure and parameters set) and to perform forecasting, respectively.

Figure 1 also shows that the DFS has three main modes of operation: (i) the "Set-up" mode, (ii) the "Update" mode and (iii) the "Forecast" mode. These modes of operation define when the relevant data analyses in each DFS module are performed. The "Set-up" mode is used for tuning the data-driven DFS when it is initialised (i.e., used for the first time to analyse a specific demand signal). Later on, it is used periodically (e.g., every three months) when the DFS is re-initialised (to account for seasonal variations, growing demand over time, etc.; or following occasional operational/other changes in the WDS - e.g., increased demand due to a new network expansion). The "Update" mode is used regularly (e.g., every week) when the DFS is updated (to constantly capture the WDS's most recent operating conditions) thereby providing a continuously adaptive DFS. Finally, the "Forecast" mode is the normal operating mode used at every forecasting time (e.g., every hour - if the observed demands update the historical time series records of demand data that are stored into a "demand signals" database at hourly or sub-hourly intervals; or every 24 hours - if the observed demands update those historical time series records at daily intervals) to perform WDF.

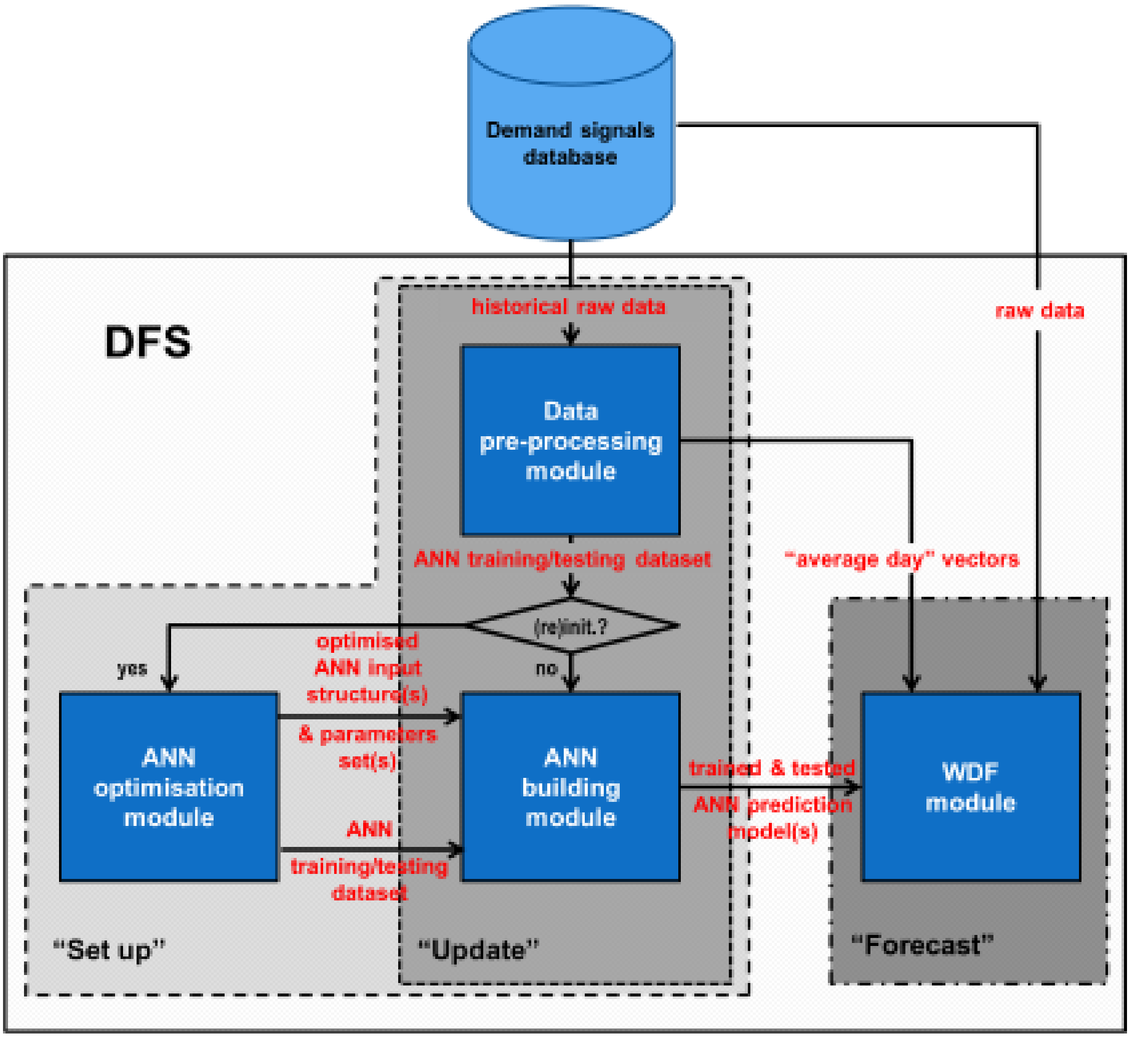

Figure 1. Diagrammatic representation of the Demand Forecasting System.

The DFS modules' methodological details are presented in the following four sub-sections.

2.2. Data pre-processing module

Demand data from a WDS are often imperfect - i.e., with frequent erroneous timestamps, large parts of missing data, etc. For this reason, the effective cleaning and pre-processing of raw data is important to achieve accurate WDF. In view of this, the main objective of this module is to, for each demand signal being analysed, assemble a valuable set of demand data to be used for building the EANN model(s). The secondary objective of this module is to compute, for each demand signal being analysed and for each Day of the Week (DoW) (i.e., Monday, Tuesday, etc.), an "average day" vector whose values will be used as surrogate demand predictions when the DFS cannot return an output (e.g., due to lack of incoming data – faulty sensor). As it can be seen from Figure 1, the assembled dataset is then passed onto the ANN optimisation module (if the DFS is being initialised/re-initialised) or directly to the ANN building module (if the DFS is being updated only). On the other hand, the computed "average day" vectors are passed on to the WDF module.

The above is achieved by performing the following steps: (i) retrieving, for the specific demand signal being analysed, the latest *m* days (e.g., 90 days) of past historical raw data from the "demand signals" database, (ii) checking and correcting erroneous timestamps, (iii) creating a uniformly spaced time series, (iv) replacing blank entries with missing value indicators (NAN – 'Not A Number'), (v) assigning ToD (i.e., a value between 1 and *g*, where 1 corresponds to midnight and *g* is the number of demand/NAN values in one day) and DoW (i.e., a value between 1 and 7) indices to each demand value, (vi) rearranging the resulting *m*-day time series into *m* vectors (i.e., one vector for each day with *g* demand/NAN values), (vii) using a heuristics-based procedure to discard vectors containing large parts of missing data, (viii) using the linear interpolation to fill in the missing values in each of the remaining vectors, (ix) grouping these vectors according to their relevant DoW and performing in sequence three statistical tests aimed at gradually filtering out vectors containing outliers and values that are not consistent with the expected demand variations assuming WDS normal operations (e.g., in the absence of pipe bursts or other unusual demands), (x) assembling the set of demand data for EANN model(s) training/testing by using all the remaining vectors and (xi) computing the "average day" vectors (each containing a total of *g* ToD averages) by using the remaining vectors in each DoW-group. Details of the above data analyses can be found in Romano et al. (2012) and Romano (2012).

It is important to stress that here, however, the data pre-processing module is implemented in such a way that allows the DFS user to decide how many of the three statistical tests mentioned in step (ix) to perform. As a consequence, not only is the DFS user able to adjust (i.e., relax or make more stringent) the parameters employed by these statistical process control -based (Shewhart, 1931) tests (e.g., number of standard deviations from the mean that define the confidence limits for outliers detection) but also to perform only one or two statistical tests or not to perform these tests at all. The main reason for this is to enable the user to have the maximum flexibility in setting/finding a desired/most suitable working definition of "anomalous data" and hence the DFS striking a balance between accuracy of the demand predictions made with spurious data (e.g., outliers caused by data communication problems), data recorded during abnormal WDS operations (e.g., in the presence of a pipe burst) and data recorded during normal WDS operations. Indeed, it has to be noted that by filtering out vectors containing outliers and values that are not consistent with the expected demand variations assuming WDS normal operations and then making use of those pre-processed data series to train the EANN models, the DFS framework is only reliant on the (limited) extrapolation capabilities of the EANN model(s) when having to make predictions using anomalous demand data.

2.3. ANN optimisation module

The objective of this module is to, for each demand signal analysed and for each forecast horizon considered, automatically select the ANN input structure and set of parameters that, when used for developing the relevant short-term ANN prediction model, enables it to yield the best WDF performance. The main reason for doing this is that different demand signals (e.g., from different DMA types - rural, residential, etc.) will have to be analysed in the studied WDS. As these signals may show extremely varying patterns, the use of a pre-defined ANN input structure and parameters set may lead to developing prediction models that exhibit sub-optimal forecasting performance. Furthermore, in order to accurately predict demand at increasing forecast horizons (e.g., from 1 hour to 24 hours), increasing the complexity of the ANN model (e.g., more hidden neurons) and/or increasing the complexity of its input structure (e.g., more past demand values in input to the ANN model) may be required. Taking all this into consideration, the potential benefits resulting from the use of the approach proposed here are two-fold: the quality of the ANN models' predictions improves and the DFS becomes tailored to the specific demand signal/forecast horizon to which it is applied.

The above is achieved here by using an EA-based optimisation strategy. The automatically selected ANN input structure(s) and parameters set(s) are then passed onto the ANN building module (see Figure 1) where they will be repeatedly (i.e., at every DFS updating) used for training and testing the (demand signal and forecast horizon specific) ANN prediction model(s), until being replaced by newly selected ones when the DFS is re-initialised.

2.3.1. ANN model building issues

Several issues have to be considered in order to build ANN models that exhibit good WDF performance for different demand signals/forecast horizons. These issues include the choice of: (i) the ANN structure, (ii) the transfer function, (iii) the training algorithm, (iv) the ANN parameters and (v) the ANN input structure. A number of preliminary sensitivity analysis type tests (the detailed results of which are not shown here due to space restrictions) were performed in order to investigate these issues for the problem at hand. A brief overview of the tests performed and of the main findings from these tests is given below (see Romano, 2012).

With regard to the first three issues, the aforementioned tests investigated their influence on the WDF performance and training speed. The investigated ANN structures included: (i) the Feed Forward ANNs (Bishop, 1995) with one and two hidden layers, (ii) the Jordan ANN (Jordan, 1986) and (iii) the Elman ANN (Elman, 1990). The investigated transfer functions (for the Feed Forward ANN models only) included the logistic and the hyperbolic tangent transfer functions for the neurons in the hidden layer(s) and the logistic, hyperbolic tangent and linear transfer functions for the neuron in the output layer. The investigated training algorithms (for the Feed Forward ANN models only) included the Back Propagation method, the Conjugate Gradient and the Levenberg-Marquardt algorithm (Rumelhart et al., 1986; Masters, 1995). The Feed Forward ANNs with a hyperbolic tangent transfer function for the neurons in the single hidden layer and a linear transfer function for the neuron in the output layer, trained using the Back Propagation method were identified as the most suitable candidates (i.e., faster training and better predictive accuracy). Furthermore, an approach whereby the DFS makes use of the same ANN structure, transfer function and training algorithm for every ANN model (i.e., for any demand signal and forecast horizon) was found not to affect the WDF performance significantly. In the light of the results obtained, the aforementioned ANN structure, transfer function, training algorithm and approach were selected for use in the DFS. Bearing this in mind, it is important to stress that this choice was also supported by the theoretical consideration that one hidden layer Feed Forward Back Propagation ANNs are capable of arbitrary non-linear function approximation (see – e.g., Hornick et al., 1989).

With regard to the selection of the ANN parameters, the aforementioned tests investigated the influence of the number of hidden neurons on the forecasting performance. The tests revealed that training an ANN model using too few hidden neurons leads to poor performance but also, using an arbitrary large number of hidden neurons leads to overfitting the data (i.e., such ANN model closely approximates the training dataset but it lacks the power to generalise - i.e., it fails on the unseen testing dataset). In view of this finding, the approach selected for use in the DFS involves the use of the early stopping (Weigend, 1994) and the Weight Decay Regularisation - WDR - (Bishop, 1995) techniques. These techniques have been successfully used (e.g., Moody, 1992) to allow striking a balance between ANN learning and generalisation. Early stopping involves controlling the number of training cycles while WDR involves applying a penalisation coefficient α (i.e., coefficient of WDR) to the weights of the ANN model. In this scenario, for each ANN model the right number of hidden neurons, the right number of training cycles and the right value of the coefficient of WDR have to be accurately chosen in order to achieve the best WDF performance.

With regard to the ANN input structure, the aforementioned tests investigated the influence on the WDF performance of input structures including combinations of the following pieces of information: (i) a certain number of past demand values (i.e., LagSize), (ii) the ToD index associated with the forecast horizon converted into a field type form (i.e., ones and zeros) and (iii) the DoW index associated with the forecast horizon also converted into a field type form. This is shown in Figure 2. The results obtained revealed that not only the LagSize but also the use (or omission) of the other considered explanatory variables strongly influence the ANN models' WDF performance. Furthermore, they showed that no general rule for the selection of the right input structure can be applied. Thus, similarly to what found for the ANN parameters, for each ANN model the right

LagSize and the right combination of the other considered explanatory variables have to be accurately chosen in order to achieve the best WDF performance.

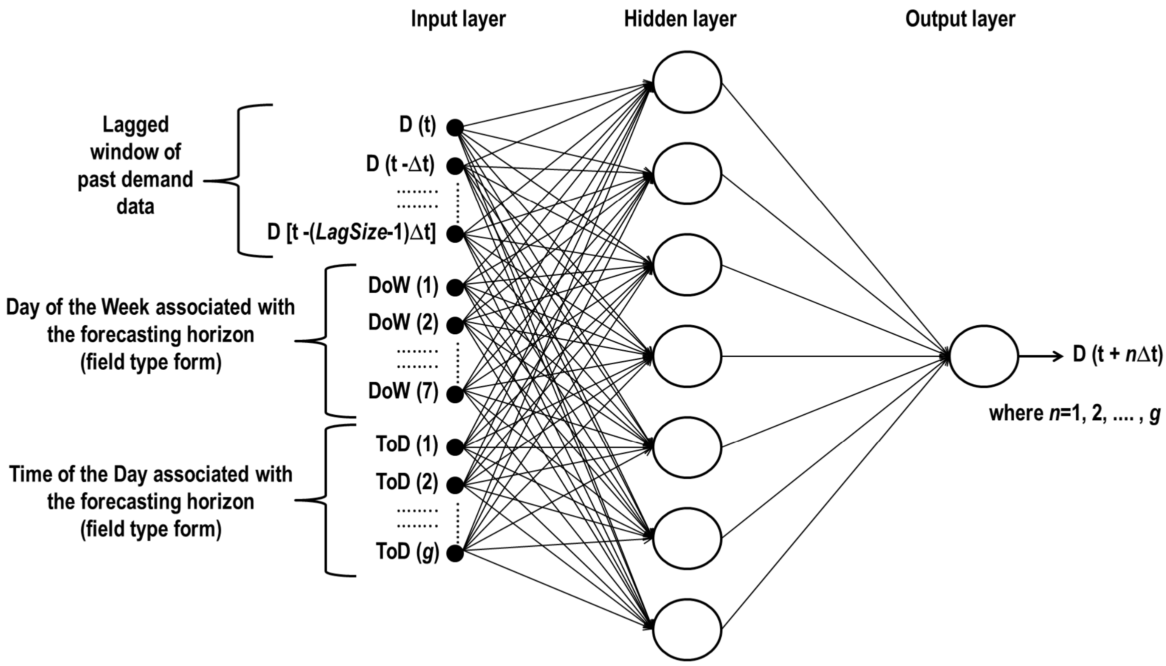

Figure 2. Artificial Neural Network for short-term Water Demand Forecasting showing a generic example of the investigated input structures.

2.3.2. EA-based optimisation strategy

This part of the methodology focuses on the selection of the optimal ANN input structure and parameters in the DFS, which is essentially a combinatorial problem. The use of a manual trial and error procedure is not be feasible bearing in mind that the DFS has to deal with different forecast horizons and, potentially, many hundreds of different demand signals in large real-life WDS. Similarly, the use of a full enumeration procedure would be far too computational expensive. Therefore, an EA-based optimisation strategy was selected. The main reason is that EAs do well in large search spaces by working only with a sample population and have the power to discover good solutions rapidly for difficult high-dimensional problems (De Jong, 2007). Specifically, similarly to Romano et al. (2013), an Evolutionary Strategy algorithm (Schwefel, 1995) is used here.

The parameters of the Evolutionary Strategy algorithm are: (i) the number of parents per generation – μ, (ii) the number of offspring per generation – λ, (iii) the total number of fitness function evaluations – $N_{f.f.e.}$ (i.e., termination condition), (iv) the probability of a parameter being perturbed – $P_{mut.}$, (v) the standard deviation of normal (i.e., Gaussian) perturbation – σ (i.e., mutation strength) and (vi) the selection operator – "+" or "," (see Beyer and Schwefel, 2002). These parameters were chosen as shown in Table 2 after limited sensitivity type analysis. Note that, although the detailed results of this analysis are not shown here, it is worth reporting that the main finding confirmed the observation from research on meta-EAs that most EAs are fairly insensitive to exact parameter settings (see – e.g., De Jong, 2007). Indeed, for the range of parameters tested (i.e., μ equal to 5, 15 or 25, λ equal to 50, 100 or 200, $N_{f.f.e.}$ equal to 255, 515 or 1025, $P_{mut.}$ equal to 0.4, 0.6 or 0.8, σ equal to 0.2, 0.4 or 0.6, and selection operator equal to "+" or ","), the Evolutionary Strategy algorithm allowed finding ANN input structure and parameters sets that led to the development of EANN prediction models with good forecasting performances (as indicated by $NS_{Index} > 0.9$ for the various testing datasets – see below) in a computationally efficient manner (i.e., a single Evolutionary Strategy algorithm run completing in less than 15 minutes on a standard dual core personal computer with 3.48 Gb of RAM).

Table 2. Values of the Evolutionary Strategy algorithm's parameters.

Parameter	**Value**
Number of parents per generation – μ	15
Number of offspring per generation – λ	100
Number of fitness function evaluations – $N_{f.f.e.}$	515
Probability of a parameter being perturbed – $P_{mut.}$	0.6
Mutation strength – σ	0.4
Selection operator	+

Considering the ANN parameters and the variables that define the ANN input structure shown in Table 3 as the decision variables for the problem at hand, the aim of the Evolutionary Strategy used here can be stated as follows: to automatically find the set of decision variables that minimises the ANN model prediction error on the testing dataset (i.e., a randomly chosen sub-set – e.g., 30% – of the assembled EANN training/testing dataset – see the data pre-processing module sub-section). Note that, for each decision variable, Table 3 shows the range of values used in optimisation. These ranges were selected after carrying out a number of preliminary tests (not shown here due to space restrictions) aimed at defining the size of the search space that is likely to enable finding optimal solutions.

Table 3. Decision variables and associated ranges of variability.

Decision variable	**Range of values used in optimisation**
Number of hidden neurons	10 - 100
Number of training cycles	50 - 500
Value of the coefficient of Weight Decay Regularisation – α	10^{-5} - 10^{3}
LagSize (i.e., number of past demand values in input to the ANN prediction model)	2 - 168
Time of the Day	use/do-not-use
Day of the Week	use/do-not-use

For each generation (i.e., cycle of the Evolutionary Strategy algorithm), the ANN model prediction error on the testing dataset is computed by using the Nash-Sutcliffe index - NS_{Index} - (Nash and Sutcliffe, 1970). This index is a normalised statistic that determines the relative magnitude of the residual variance compared to the measured data variance. The index values range between $-\infty$ and 1, with 1 being the optimal value. A zero value indicates that model predictions are as accurate as the mean of the observed data, whereas a negative value occurs when the observed mean is a better predictor than the model. The NS_{Index} is commonly used in the literature and recommended by many (e.g., ASCE, 1993).

2.4. ANN building module

The objective of this module is to build a short-term ANN prediction model for the specific demand signal analysed and for each particular forecast horizon considered. This objective is achieved by training and testing the ANN prediction model(s) using the EANN training/testing dataset assembled in the data pre-processing module and the optimised ANN input structure(s) and parameters set(s) selected in the ANN optimisation module. The resulting EANN model(s) is(are) then passed onto the WDF module (see Figure 1) where it(they) will be used to predict future demands.

Each EANN model developed here takes as an input a number (i.e., *LagSize*) of past demand values. Furthermore, depending on the input structure selected in the ANN optimisation module, it may have the following additional inputs: (1) the ToD index associated with the forecast horizon and (2) the DoW index associated with the forecast horizon. The output of each EANN prediction model is the predicted demand at the particular forecast horizon being considered (see Figure 2).

The EANN model(s) training dataset consists of a sub-set (i.e., *Train%* - e.g., 80%) of the EANN training/testing dataset assembled in the data pre-processing module. The remaining data (i.e., *Test%*) form the testing dataset which is used to evaluate the EANN model(s) performance. The goodness-of-fit measure used to compare the predicted demand values with their observed counterpart in the testing dataset is the NS_{Index} (see the previous sub-section).

Note that, as it can be observed from Figure 1, when this module runs in the “Update” mode, the EANN model(s) building process continues to make use of the ANN input structure(s) and parameters set(s) automatically selected at DFS (re)initialisation. The rationale is that, in the absence of operational/other WDS changes, it is expected that a demand signal will be affected only by relatively minor changes in the interval between two DFS re-initialisations. Thus, continuing to make use of the optimised ANN input structure(s) and parameters set(s) for updating the prediction model(s) is likely not to affect its(their) WDF performance significantly. Bearing this in mind, it is possible to state that, in principle, the added computational burden of using the EA-based optimisation strategy at short regular time intervals (e.g., weekly) is not justified. Using the EA-based optimisation strategy periodically (e.g., every three months - when the DFS is re-initialised), on the other hand, enables the DFS to take into account factors such as the seasonal demand variations and growing demand over time. Also, it would hardly pose a computational problem even if hundreds of demand signals have to be analysed and the computing power is scarce. This is because the DFS re-initialisations can be scheduled to run at different times for different demand signals (i.e., in a sequential fashion) during, for example, a three month period.

2.5. WDF module

The objective of this module is to predict water demand for the next 24 hours every time the DFS runs in the “Forecast” mode (i.e., at every forecasting time). This is achieved by using the EANN model(s) trained and tested in the ANN building module and one of the two WDF approaches described before - i.e., the *e*EANNs and *r*EANN approaches. For each EANN model, a number (e.g., as equal to the selected optimal LagSize in the case of the *e*EANNs approach) of latest raw demand values is retrieved from the “demand signals” database. Once this is done, this data is subjected to a pre-processing procedure involving the following steps: (i) checking and correcting erroneous timestamps, (ii) creating a uniformly spaced time series and (iii) replacing blank entries with missing value indicators. Finally, demand prediction for the particular forecast horizon being considered is performed.

3. CASE STUDY ANALYSES

3.1. Case study description

The results of two data analyses carried out on a single real-life case study are reported here. The main objective of these analyses was to test, evaluate and illustrate the capabilities of the DFS that implements the novel WDF methodology presented in this paper. More specifically:

- The first data analysis aimed at testing and evaluating the capabilities of the *e*EANNs and *r*EANN approaches when a one day time interval between consecutive forecasts was considered (i.e., 24 hour forecast periodicity).

- The second analysis aimed at testing and evaluating the capabilities of the *e*EANNs and *r*EANN approaches when a 1 hour time interval between consecutive forecasts was considered (i.e., 1 hour forecast periodicity).

Four different scenarios were investigated for all of the above. More specifically, the DFS was run: (1) making use of the ANN optimisation module and performing the weekly DFS updating (Scenario 1), (2) making use of the ANN optimisation module and without performing the weekly DFS updating (Scenario 2), (3) without using the ANN optimisation module and performing the weekly DFS updating (Scenario 3) and (4) without using the ANN optimisation module and without performing the weekly DFS updating (Scenario 4). These scenarios enabled: (i) assessing the self-learning capabilities of the developed DFS in terms of ability to tailor itself to the analysed demand signal/considered forecast horizon and also in terms of ability to continuously adapt its parameters as conditions in the WDS change and (ii) evaluating if all this, in turn, results in more accurate WDF.

The signals analysed here represent water demands in three Yorkshire Water Services (YWS) DMAs and a single YWS WSZ covering significant parts of two towns in the Yorkshire county. The three DMAs being studied are deemed representative of many UK DMAs. They have different characteristics and varying sizes. As an ensemble, they contain light industrial, urban and rural regions. Their individual total mains length varies between 16.2 and 25 km and the number of domestic properties varies between 1,129 and 3,493. The overall number of commercial properties varies between 103 and 340 and the number of commercial properties with an annual demand in excess of 400 m^3 located in each of these DMAs varies between 31 and 193. Furthermore, one of these DMAs (i.e., DMA1) contains three major metered consumers (i.e., with an average annual demand in excess of 10,000 m^3). The average daily inflows into these DMAs (for the analysis period considered – see later) were as follows: (DMA1) 25.6 l/s, (DMA2) 24.6 l/s and (DMA3) 6.4 l/s. The WSZ being studied is also deemed representative of other UK WSZs. It has a population of over 70,000 and contains urban, industrial and rural regions. This WSZ is exclusively served by a service reservoir with an average daily outflow (for the relevant analysis period considered – see later) as equal to 5,553 l/s. Note that the studied DMAs are all leaf DMAs (i.e., without water exports), they all have only one inlet and there is no water storage in any of them. In the light of these characteristics, the flow measured at the inlet of each of these DMAs was assumed as equal to the actual DMA demand (= consumption + leakage). Similarly, the flow measured at the service reservoir outlet was assumed equal to the actual WSZ demand.

The original historical water demand time series were made up of flow readings averaged (by the flow sensors themselves) over a 15 minute sampling interval. In the case of the three DMAs, the utilised time series data referred to the 181 day period (i.e., approximately 6 months) between the 3rd of May 2011 and the 30th of October 2011. In the case of the WSZ, the utilised time series data referred to the 181 day period between the 30th of November 2010 and the 29th of May 2011. For the purposes of the data analyses performed here, however, the original historical water demand time series were re-sampled (by averaging) at 1 hour time intervals (thus mimicking the situation whereby only 24 demand values per day are available – i.e., g=24). Having said this, it is important to stress that, although historical time series were used, the demand data were fed to the DFS in a simulated on-line fashion (i.e., as the DFS would have been used in real-life – see the following sub-sections for further details). However, given that only approximately six months of demand data were considered here, the periodic (e.g., every 3 months) DFS re-initialisation was not performed.

3.2. Results and discussion

3.2.1. Daily forecasting analysis

This sub-section summarises the analysis done and the results obtained when forecasting demands with a 24 hour forecast periodicity. This corresponds to a real-life situation where the observed demand data are communicated to the control room once a day. Note that, here, it was assumed that this daily data transfer occurred soon after observing the last demand value of the day at 23:00 p.m.. Hence, the 1 to 24 hours ahead demand predictions were made for the next day, with the 1 hour ahead prediction always corresponding to 00:00 a.m. and the 24 hours ahead prediction always corresponding to 23:00 p.m..

When the ANN optimisation module was used, the DFS was firstly initialised using the first 90 days (i.e., m=90) of data in each relevant water demand dataset (i.e., from the 3rd of May 2011 to the 31st of July 2011 - in the case of the three DMAs, and from the 30th of November 2010 to the 27th of February 2011 - in the case of the WSZ). Bearing in mind that (after the DFS initialisation) the demand data were fed to the DFS to simulate on-line operation, the DFS was at that point used once a day for forecasting the 1 to 24 hours ahead future demand values for the entire methodology's validation period (i.e., 91 day period from the 1st of August 2011 to the 30th of October 2011 in the case of the three DMAs, and from the 28th of February 2011 to the 29th of May 2011 in the case of the WSZ) by subjecting it to weekly updates (i.e., Scenario 1). Following all this, the "Forecast" process was repeated without subjecting the DFS to the weekly updates (i.e., Scenario 2). As an example of the ANN parameters and variables that define the ANN input structure automatically selected during the DFS initialisation, Table 4 shows the selected DMA1 values for the two WDF approaches and different forecasting horizons analysed.

Table 4. Example of the automatically selected Artificial Neural Network parameters/input structures.

Water Demand Forecasting approach	Forecast horizon [h]	Time of the Day variable	Day of the Week variable	LagSize [#]	Hidden neurons [#]	Training cycles [#]	Coefficient of Weight Decay Regularisation [#]
***r*EANN**	1	used	not-used	36	100	450	0.1
***e*EANNs**	1	used	not-used	30	50	400	0.1
	2	used	not-used	36	40	500	0.1
	3	used	used	30	70	450	1
	4	used	not-used	30	100	400	0.1
	5	used	not-used	36	50	150	0.00001
	6	used	not-used	42	50	250	0.1
	7	used	not-used	36	60	150	0.0001
	8	used	not-used	48	40	250	0.1
	9	used	used	48	80	250	0.1
	10	used	not-used	48	70	300	0.1
	11	used	not-used	42	30	300	0.1
	12	used	not-used	42	100	300	1
	13	used	not-used	48	80	250	0.1
	14	used	not-used	36	30	500	1
	15	used	not-used	48	40	500	0.1
	16	used	not-used	42	90	400	0.1
	17	used	not-used	42	10	250	0.1
	18	used	not-used	30	20	350	0.1
	19	used	not-used	30	60	200	0.00001
	20	used	not-used	36	90	400	0.1
	21	used	not-used	30	80	450	0.1
	22	used	not-used	30	50	200	0.0001
	23	used	not-used	30	80	250	0.001
	24	used	not-used	18	40	350	0.0001

When the ANN optimisation module was not used, the DFS was also firstly initialised using the first 90 days of data in each relevant water demand dataset. Here, however, Fixed Parameters and a fixed Input Structure (i.e., FP&IS) were used for all the ANN prediction models. These fixed ANN parameters and input structure were chosen as follows. The number of hidden neurons was set equal to 60, the number of training cycles was set equal to 400 and the coefficient of WDR was set equal to 10. The ANN input structure included 24 past demand values (i.e., LagSize=24) and the ToD and DoW indices associated with the forecast horizon. Note that the selection of these particular ANN parameters and input structure was found, after a series of preliminary tests, to ensure that, for all the analysed demand signals and for all the forecast horizons, the resulting ANN prediction models were able to perform reasonably well (i.e., closely approximate the training datasets whilst allowing good generalisation performance). This was evaluated using the NS_{Index}. The preliminary tests performed here involved training and testing 100 ANN models for each demand signal and for each forecast horizon by varying (at discrete intervals) the decision variables shown in Table 3 within their associated ranges of variability. The detailed results of these tests are not shown here due to space restrictions. Once initialised, the DFS was used daily for performing 1 to 24 hours ahead prediction of future demand values for the entire validation period by subjecting it to weekly updates (i.e., Scenario 3). The "Forecast" process was then repeated without subjecting the DFS to the weekly updates (i.e., Scenario 4).

Table 5 summarises the forecasting performances obtained on the validation datasets for the four scenarios, two WDF approaches (which, as the ANN optimisation module was not always used, are generically called 1 ANN approach and 24 ANNs approach where appropriate in this and in the following sub-section) and all demand signals analysed. The WDF performances are expressed in terms of the NS_{Index} and in terms of the Mean Square Error (MSE) and the Mean Absolute Percentage Error (MAPE) indices (see – e.g., Donkor et al., 2012; Bennett et al., 2013b). The MSE and MAPE goodness-of-fit measures are used in addition to the NS_{Index} because: (1) MSE penalises WDF models that exhibit large deviations, hence it is useful for complementing the NS_{Index} in identifying models that fit the data well and (2) MAPE is independent of the actual demand signals' magnitude, hence it is useful for comparing the performance of WDF models for different demand signals (e.g., from different DMAs or WSZs) and different WDSs. Note that lower values of both MSE and MAPE are better.

The results shown in Table 5 lead to the following observations. Firstly, regardless of the WDF approach used, the best forecasting performances were obtained for Scenario 1. This provides evidence that the advanced self-learning methodological framework (i.e., involving EA-based optimisation and continuous adaptation) presented in this paper is sound and enables developing good forecast-quality models without requiring a high degree of human intervention. Indeed, although the careful human expert selection of the ANN parameters set and input structure (i.e., Scenarios 3 and 4) resulted in prediction models that performed reasonably well (i.e., average NS_{Index} values in the 0.92-0.94 range, and average MAPE values in the 7.80%-9.48% range), these models exhibited constantly worse performance (i.e., average NS_{Index} values in the 0.95-0.97 range, and average MAPE values in the 5.28%-6.50% range) than their optimised counterparts (i.e., Scenarios 1 and 2). Furthermore, although less remarkably, systematic WDF performance improvements were observed by comparing the relevant results obtained for the "with updating" and "without updating" scenarios.

Secondly, the results shown in Table 5 seem to suggest that more accurate forecasts can be achieved by using multiple ANN models. Indeed, in all the different scenarios investigated, the 24 ANNs approach outperformed the 1 ANN approach. Despite this, with specific regard to the EA-based optimisation case, it is important to stress that the *r*EANN approach still enabled achieving relatively good WDF performances. Also, compared to the *e*EANNs approach, this approach requires less computational effort (as only one ANN prediction model has to be optimised and then trained and tested). All this suggests that the *r*EANN approach could nonetheless be satisfactorily and conveniently used by water companies.

Table 5. Results of the 24 hour forecast periodicity analysis (refer to Table 1 for the meaning of the abbreviations used).

				DMA1	DMA2	DMA3	WSZ	Averages
Updating	1 ANN	EA	NS_{Index}	0.96	0.97	0.98	0.94	**0.96**
			MSE [(l/s)^2]	4.89	1.69	0.19	275317.25	
			MAPE [%]	6.39	4.09	5.64	6.23	**5.59**
		FP&IS	NS_{Index}	0.91	0.94	0.93	0.91	**0.92**
			MSE [(l/s)^2]	12.49	3.69	0.54	398415.71	
			MAPE [%]	9.83	6.40	9.24	8.84	**8.58**
	24 ANNs	EA	NS_{Index}	0.96	0.98	0.98	0.95	**0.97**
			MSE [(l/s)^2]	4.80	1.54	0.19	211383.83	
			MAPE [%]	6.30	3.80	5.35	5.68	**5.28**
		FP&IS	NS_{Index}	0.93	0.95	0.94	0.93	**0.94**
			MSE [(l/s)^2]	10.14	3.15	0.46	309517.58	
			MAPE [%]	9.07	5.69	8.68	7.75	**7.80**
Without Updating	1 ANN	EA	NS_{Index}	0.94	0.97	0.97	0.93	**0.95**
			MSE [(l/s)^2]	7.60	1.71	0.22	342227.43	
			MAPE [%]	8.47	4.19	5.99	7.36	**6.50**
		FP&IS	NS_{Index}	0.89	0.94	0.92	0.91	**0.92**
			MSE [(l/s)^2]	14.78	3.77	0.55	413164.15	
			MAPE [%]	11.93	6.48	10.12	9.37	**9.48**
	24 ANNs	EA	NS_{Index}	0.96	0.98	0.98	0.95	**0.97**
			MSE [(l/s)^2]	5.97	1.55	0.20	226066.10	
			MAPE [%]	7.73	3.82	5.51	6.26	**5.83**
		FP&IS	NS_{Index}	0.93	0.95	0.94	0.92	**0.94**
			MSE [(l/s)^2]	10.15	3.16	0.49	357520.00	
			MAPE [%]	10.54	5.70	9.72	8.63	**8.65**

As an example of the good quality forecasts achievable by using the optimised ANN prediction model(s) and both the alternative WDF approaches, Figure 3 shows a comparison of the WSZ demand values observed over a three day period during the first week of the methodology's validation period with their predicted counterparts. Note that the rationale for showing this comparison for the WSZ demand signal is that the DFS exhibited the worst WDF performances in terms of NS_{Index} for this particular signal (see Table 5).

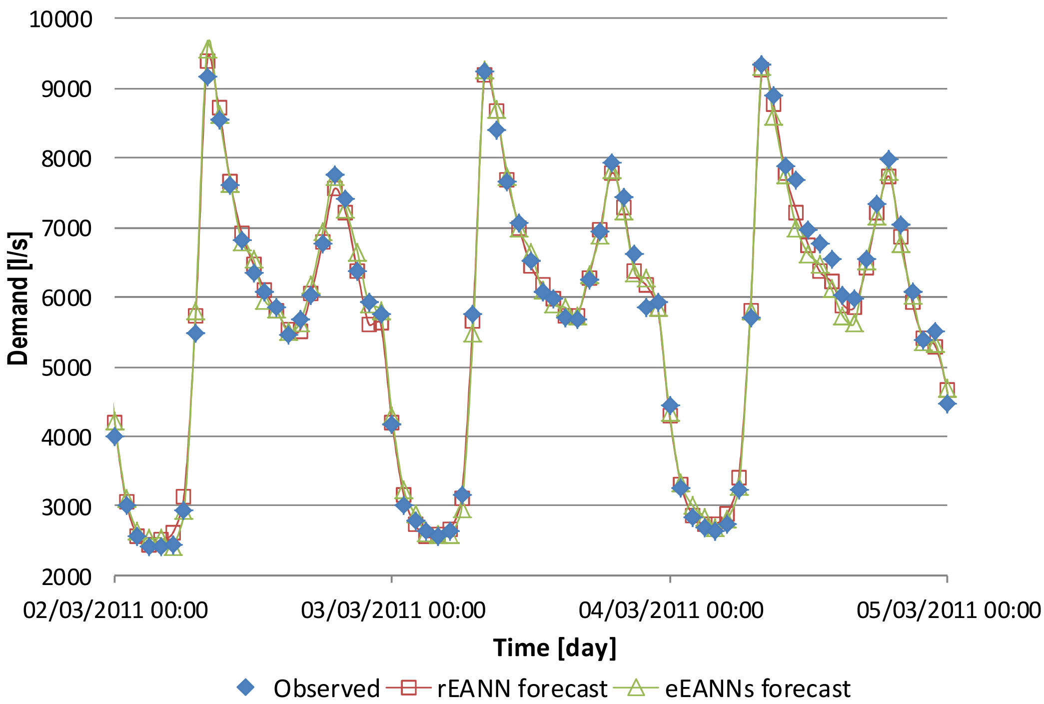

Figure 3. Three day comparison of the observed Water Supply Zone demand values with the corresponding demand values predicted by using recursive and ensemble Evolutionary Artificial Neural Network models.

3.2.2. Hourly forecasting analysis

This sub-section summarises the analysis done and the results obtained when forecasting demands with a 1 hour forecast periodicity. This corresponds to a real-life situation where the observed demand data are communicated to the control room every hour and enabled evaluating how the DFS performs when near real-time demand data are available. Here, the two WDF approaches that the DFS enables implementing and Scenarios 1 to 4 were investigated by running the DFS in the same way as described in the daily forecasting analysis sub-section with the only difference that, after the DFS initialisation, the DFS was run every hour producing the demand forecasts for the next 1 to 24 hours.

Figures 4 and 5 show the averaged (across all four signals) NS_{Index} and MAPE values respectively, computed by comparing the observed demand values with all (i.e., for the entire validation period) 1 hour ahead forecasted values, all 2 hours ahead forecasted values and so forth up to all 24 hours ahead forecasted values. These two figures clearly confirm the validity of the corresponding findings from the previous sub-section with regard to the use of the ANN optimisation module, the DFS weekly updating mechanism and the two WDF approaches. They also provide further and more reliable evidence that the use of the proposed WDF methodology results in accurate forecasts. Indeed, with regard to the reliability of the evidence gathered, it has to be noted that the results obtained by using a 24 hour forecast periodicity were dependent on the particular ToD the DFS was run. On the contrary, the results obtained here do not depend on any arbitrary assumption and enable drawing a more complete picture of the DFS performances for each scenario investigated, WDF approach implemented and particular forecast horizon considered.

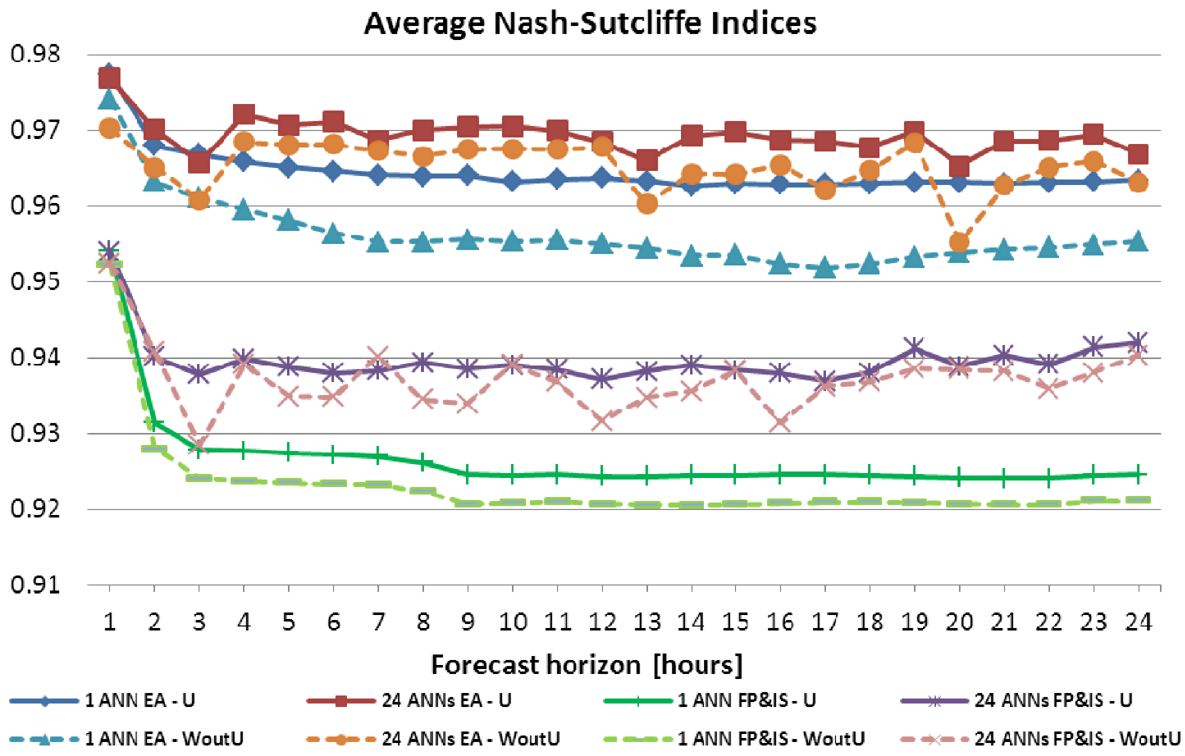

Figure 4. Validation datasets average Nash-Sutcliffe indices for each forecast horizon considered, scenario investigated and Water Demand Forecasting approach implemented (refer to Table 1 for the meaning of the abbreviations used in the legend).

Figures 4 and 5 also show how the best performances were obtained for the first hour ahead. This behaviour was expected and it is understandable that the farther in the future a prediction has to be made the less accurate the prediction will be. Notwithstanding, especially when the ANN optimisation module was used and the DFS was updated weekly (i.e., Scenario 1), the performances appeared to almost plateau after the second/third hour ahead. This suggests that the developed DFS can be used confidently by the water companies for supporting the near real-time management of their WDSs. For example, when using the DFS for supporting near real-time pump scheduling, the water companies’ personnel can be confident that the level of uncertainty of the farther in the future forecasts (e.g., 24 hours ahead) will unlikely result in a service reservoir with insufficient water to satisfy the consumers demand.

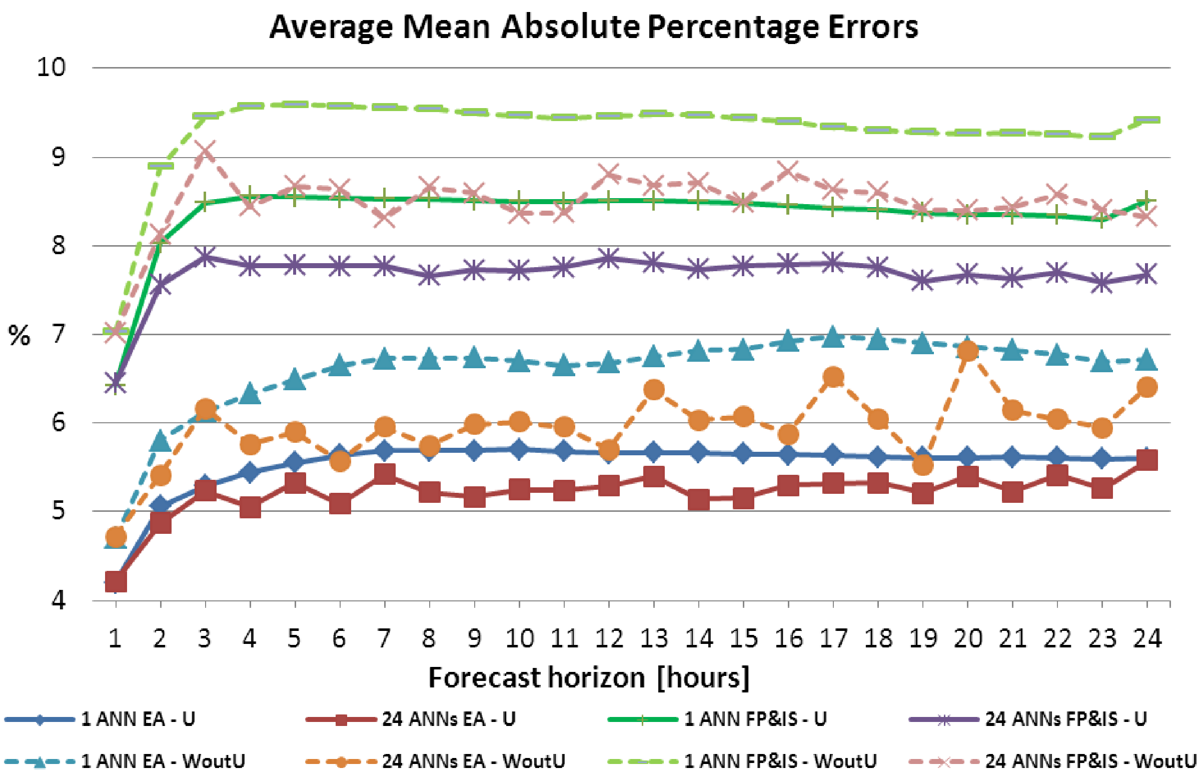

Figure 5. Validation datasets average Mean Absolute Percentage Errors for each forecast horizon considered, scenario investigated and Water Demand Forecasting approach implemented (refer to Table 1 for the meaning of the abbreviations used in the legend).

As mentioned in the previous paragraph, Figures 4 and 5 show the DFS performances in terms of NS_{Index} and MAPE values aggregated for the four analysed demand signals. In order to provide the aforementioned complete picture of the DFS performances, it is therefore important to stress here that the patterns of the NS_{Index} and MAPE values observed for each analysed signal were very similar to those showed in Figures 4 and 5. Likewise, the patterns of the MSE values observed for each analysed signal did not make an exception to this.

3.2.3. Further observations

The results obtained in the case study clearly show the good performances of the WDF methodology presented. Nevertheless, those results should be, ideally, compared with the results reported in other studies from the literature and obtained using different short-term WDF methodologies. Such an assessment is attempted here. However, it has to be born in mind that direct results comparison is not possible as the tests carried out in the studies that will be mentioned below involved the use of different approaches, datasets (of varying forecastability and length), explanatory variables and also diverse periodicities and forecast horizons. This said, as an indication only, note that the hourly WDF results reported in the studies by Zhou et al. (2002), Herrera et al. (2010) and Bakker et al. (2013), the daily WDF results reported in the study by Cutore et al. (2008), and the weekly WDF results reported in the studies by Jain et al. (2001) and Bougadis et al. (2005) all showed values of the computed NS_{Index} or Coefficient of Determination smaller than 0.90. In addition, the hourly WDF results reported in the studies by Alvisi et al. (2007) and Kim et al. (2012) showed MAPE values in the 4% to 15% range. Given these reported figures, taking into consideration that $NS_{Index} \leq$ Coefficient of Determination (Murphy, 1995), and stressing once again the limitations of the comparative analysis tried here, it can be stated that performance improvements in terms of WDF accuracy/reliability over previously developed methods appear to be achievable by using the DFS.

Bearing in mind the above, it has to be stressed that, in this paper, a comparison between the performance attained using the proposed self-learning and adaptive methodological framework and the performance attained using ANN models with fixed ANN parameters and input structure (with and without updating) is provided for both the daily and hourly forecasting analyses. The latter ANN models can be seen as the state-of-the-art in the short-term WDF research field (see – e.g., Jain et al., 2001; Bougadis et al., 2005) and thus provide a suitable benchmark for (more robust) performance comparison. In view of this, the contribution of the novel WDF methodology implemented in the DFS is further highlighted.

It is also worth noting here that the novel WDF methodology implemented in the DFS is of generic nature. It could therefore be applied for performing WDF at different (i.e., shorter or longer) forecast horizons and with varying periodicities. Allowing for this, the inclusion of additional determinants (if required), such as weather-related and/or socio-economic explanatory variable, in the relevant ANN prediction models could benefit from the automatic EA-based ANN input structure selection framework presented in this paper. Finally, yet again because of its generic nature, the presented methodology has the further potential to be applied for the prediction of other water-demand-related forecast variables that are of interest to water companies (e.g., daily peak demand, daily or monthly total system demand, annual per capita demand, etc.) and/or for the prediction of other signals coming from a WDS (e.g., various water quality parameters and pressures - although this may require, for example, the use of more advanced data pre-processing techniques – see Romano et al., 2010).

4. CONCLUSIONS

Reliable short-term demand forecasting is of paramount importance for making informed operational decisions and hence for supporting the near real-time management of smart WDS. A novel short-term (i.e., up to 24 hours in the future) WDF methodology has been developed and presented here. As the developed methodology is exclusively based on the analysis of water demand signals, it results in parsimonious WDF prediction models that do not require a large number of explanatory variables and, hence, lengthy/complex data collection processes. Furthermore, by making use of EANN models, it provides water companies with an effortless but theoretically sound and robust tool for building those models.

The above short-term WDF methodology is implemented in a fully automated, data-driven DFS that enables implementing two different WDF approaches whereby one or multiple EANN models are used. The DFS has advanced self-learning capabilities. That is, not only it can tailor itself to the particular demand signal being analysed and forecast horizon being considered but it is also able to dynamically recalibrate its parameters as conditions in the WDS change. All this, in turn, entails: (i) ability to work in an on-line environment and (ii) easy operationalisation and scalability.

The developed DFS was tested on a real-life UK case study involving the use of different water demand data streams recorded at different time periods over the year. Data analyses conducted considered both 24 and 1 hour forecast periodicity. For each of these analyses, the two WDF approaches were tested and four different scenarios mainly aimed at assessing the DFS self-learning capabilities were investigated.

The results obtained provide evidence that, regardless of the WDF approach used, the methodological framework presented in this paper is sound and enables developing good forecast-quality models requiring a minimal degree of human intervention. This, in turn, indicates that the developed DFS has the potential to yield substantial improvements to the state-of-the-art in near real-time WDS management. With regard to the WDF approach used, the results obtained show that the ensemble *e*EANNs approach outperformed the recursive *r*EANN approach. However, given that the *r*EANN approach requires less effort for its implementation and still enables achieving relatively good WDF performances, its use (if favoured by a water company) can be considered too.

The future work should involve further testing and validation of the proposed DFS on longer time periods and a larger number of demand signals reflecting the wider range of operational conditions in different WDSs. The testing and validation of the DFS in an actual on-line environment should also be performed. Furthermore, the EA-based ANN input structure selection framework presented here could also be improved by letting the Evolutionary Strategy algorithm select the optimal LagStructure (i.e., the combination of past demand values that best explain the future demand variations) rather than the LagSize. Indeed, the current method does not allow including, for example, t-24 and t-168 hour lags only and a LagSize as equal to 24 means that the ANN model input includes all of the last 24 past demand values. This methodological improvement would likely result in the generation of smaller/simpler ANN models that are faster to train and may also show better generalisation performances. In addition to all this, the DFS could also be further enhanced by considering information about holidays when assembling the training/testing datasets and/or as an extra input to the ANN prediction models. This methodological improvement would likely enable overcoming an important DFS limitation (i.e., inability to account for holidays). Finally, the application of the methodology presented here for forecasting water demand at different forecast horizons and with varying periodicities could be investigated, as well as its use for forecasting other water-demand-related variables and WDS signals.

5. ACKNOWLEDGEMENTS

The data used in the paper have been collected as part of the Neptune project funded by the UK Engineering and Physical Sciences Research Council (EP/E003192/1) and provided by Mr Ridwan Patel from Yorkshire Water Services which is gratefully acknowledged.

REFERENCES

Adamowski, J.F., 2008. Peak daily water demand forecast modeling using artificial neural networks. Journal of Water Resources Planning and Management 134 (2), 119–128.

Alhumoud, J., 2008. Freshwater consumption in Kuwait: analysis and forecasting. Journal of Water Supply Research and Technology-AQUA 57 (4), 279–288.

Altunkaynak, A., Özger, M., Çakmakci, M., 2005. Water consumption prediction of Istanbul city by using fuzzy logic approach. Water Resources Management 19 (5), 641–654.

Alvisi, S., Franchini, M., Marinelli, A., 2007. A short-term, pattern-based model for water-demand forecasting. Journal of Hydroinformatics 9 (1), 39–50.

Anderson, R., Miller, T., Washburn, M., 1980. Water savings from lawn watering restrictions during a drought year in Fort Collins, Colorado. Journal of the American Water Resources Association 16 (4), 642–645.

ASCE, 1993. Criteria for evaluation of watershed models. Journal of Irrigation and Drainage Engineering 119 (3), 429–442.

Bakker, M., Vreeburg, J.H.G., van Schagen, K.M., Rietveld, L.C., 2013. A fully adaptive forecasting model for short-term drinking water demand. Environmental Modelling and Software 48, 141–151.

Bennett, C., Stewart, R.A., Beal, C.D., 2013a. Ann-based residential water end-use demand forecasting model. Expert Systems with Applications 40 (4), 1014–1023.

Bennett, N.D., Croke, B.F.W., Guariso, G., Guillaume, J.H.A., Hamilton, S.H., Jakeman, A.J., Marsili-Libelli, S., Newham, L.T.H., Norton, J.P., Perrin, C., Pierce, S.A., Robson, B., Seppelt, R., Voinov, A.A., Fath, B.D., Andreassian, V., 2013b. Characterising performance of environmental models. Environmental Modelling and Software 40, 1–20.

Beyer, H.-G., Schwefel, H.-P., 2002. Evolution strategies: a comprehensive introduction. Journal of Natural Computing 1 (1), 3–52.

Billings, B., Jones, C., 2008. Forecasting urban water demand. 2nd edition. American Waterworks Association, USA.

Bishop, C.M., 1995. Neural networks for pattern recognition. Oxford University Press, New York, USA.

Bougadis, J., Adamowski, K., Diduch, R., 2005. Short-term municipal water demand forecasting. Hydrological Processes 19 (1), 137–148.

Castillo, P.A., Carpio, J., Merelo Guervós, J.J., Rivas, V., Romero, G., Prieto, A., 2000. Evolving multilayer perceptrons. Neural Processing Letters 12 (2), 115–127.

Coomes, P., Rockaway, T., Rivard, J., Kornstein, B., 2010. North America residential water usage trends since 1992. Water Research Foundation Report Project No. 4031, Denver, USA.

Cutore, P., Campisano, A., Kapelan, Z., Modica, C., Savic, D., 2008. Probabilistic prediction of urban water consumption using the SCEM-UA algorithm. Urban Water Journal 5 (2), 125–132.

De Jong, K.A., 2007. Evolutionary computation: a unified approach. Massachusetts Institute of Technology Press, Cambridge, USA.

Donkor, E., Mazzuchi, T., Soyer, R., Roberson, J., 2012. Urban water demand forecasting: a review of methods and models. Journal of Water Resources Planning and Management, [Pub. ahead of print doi: http://dx.doi.org/10.1061/(ASCE)WR.1943-5452.0000314].

Elman, J.L., 1990. Finding structure in time. Cognitive Science 14 (2), 179–211.

Gardiner, V., Herrington, P., 1990. Water demand forecasting. 1st edition. Taylor and Francis Ltd, London, UK.

Giustolisi, O., Simeone, V., 2006. Optimal design of artificial neural networks by a multi-objective strategy: groundwater level predictions. Hydrological Sciences Journal 51 (3), 502–523.

Goodchild, C., 2003. Modelling the impact of climate change on domestic water demand. Journal of the Chartered Institution of Water and Environmental Management 17 (1), 8–12.

Harp, S., Samad, T., Guha, A., 1991. Towards the genetic synthesis of neural networks. In: Proceedings of the 4th International Conference on Genetic Algorithms, San Diego, USA.

Herrera, M., Torgo, L., Izquierdo, J., Perez-Garcia, R., 2010. Predictive models for forecasting hourly urban water demand. Journal of Hydrology 387 (1-2), 141–150.

Holland, J.H., 1975. Adaptation in natural and artificial systems. University of Michigan Press, Ann Arbor, USA.

Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. Neural Networks 2 (5), 359–366.

Hughes, T., 1980. Peak period design standards for small western U.S. water supply. Journal of the American Water Resources Association 16 (4), 661–667.

Jain, A., Ormsbee, L., 2002. Short-term water demand forecast modeling techniques - conventional methods versus AI. Journal of the American Water Resources Association 94 (7), 64–72.

Jain, A., Varshney, A., Joshi, U., 2001. Short-term water demand forecast modeling at IIT Kanpur using artificial neural networks. Water Resources Management 15 (5), 299–321.

Jordan, M.I., 1986. Serial order: a parallel distributed processing approach. Technical Report No. 8604, Institute for Cognitive Science, University of California, USA.

Keesing, R., Stork, D.G., 1991. Evolution and learning in neural networks: the number and distribution of learning trials affect the rate of evolution. Advanced in Neural Information Processing Systems 3, 805–810.

Kim, J., Shin, G., Choi, D., 2012. Development of water demand forecasting simulator and its performance assessment. In: Proceedings of the 14th Water Distribution Systems Analysis Conference, Adelaide, Australia.

Kim, J.H., Hwang, S.H., Shin, H.S., 2001. A neuro-genetic approach for daily water demand forecasting. KSCE Journal of Civil Engineering 5 (3), 281–288.

Koza, J.R., 1992. Genetic programming. Massachusetts Institute of Technology Press, Cambridge, USA.

Maidment, D., Miaou, S., Crawford, M., 1985. Transfer function models of daily urban water use. Water Resources Research 21 (4), 425–432.

Masters, T., 1995. Advanced algorithms for neural networks: a C++ sourcebook. John Wiley and Sons, New York, USA.

Moody, J.E., 1992. The effective number of parameters: an analysis of generalization and regularization in nonlinear learning systems. In: Advances in neural information processing systems: vol. 4. Morgan Kaufmann, San Francisco, USA.

Murphy, A.H., 1995. The coefficients of correlation and determination as measures of performance in forecast verification. Weather and Forecasting 10, 681–688.

Nash, J.E., Sutcliffe, J.V., 1970. River flow forecasting through conceptual models, part I - a discussion of principles. Journal of Hydrology 10 (3), 282–290.

Nasseri, M., Moeini, A., Tabesh, M. 2011. Forecasting monthly urban water demand using extended Kalman filter and genetic programming. Expert Systems with Applications 38 (6), 7387–7395.

Nicklow, J., Reed, P., Savić, D.A., Dessalegne, T., Harrell, L., Chan-Hilton, A., Karamouz, M., Minsker, B., Ostfeld, A., Singh, A., Zechman, E., 2010. State of the art for genetic algorithms and beyond in water resources planning and management. Journal of Water Resources Planning and Management 136 (4), 412–432.

Reeves, C.R., Taylor, S.J., 1998. Selection of training data for neural networks by a genetic algorithm. In: Parallel problem solving from nature, Springer-Verlag Berlin and Heidelberg GmbH & Co. K, Berlin, Germany.

Romano, M., 2012. Near real-time detection and approximate location of pipe bursts and other events in water distribution systems. Thesis (PhD), University of Exeter, Exeter, UK.

Romano, M., Kapelan, Z., Savić, D.A., 2010. Pressure signal de-noising for improved real-time leak detection. In: Proceedings of the 9th International Conference on Hydroinformatics, Tianjin, China.

Romano, M., Kapelan, Z., Savić, D.A., 2012. Automated detection of pipe bursts and other events in water distribution systems. Journal of Water Resources Planning and Management. [Epub ahead of print doi: 10.1061/(ASCE)WR.1943-5452.0000339].

Romano, M., Kapelan, Z., Savić, D.A., 2013. Evolutionary algorithm and expectation maximisation strategies for improved detection of pipe bursts and other events in water distribution systems. Journal of Water Resources Planning and Management. [Epub ahead of print doi: 10.1061/(ASCE)WR.1943-5452.0000347].

Rumelhart, D.E., Hinton, G.E., Williams, R.J., 1986. Learning internal representations by error propagation. In: Parallel distributed processing: explorations in the microstructure of cognition, vol. 1: foundations. Massachusetts Institute of Technology Press, Cambridge, USA.

Schwefel, H.-P., 1995. Evolution and optimum seeking. John Wiley and Sons, New York, USA.

Shewhart, W., 1931. Economic control of quality of manufactured product. Van Nostrand Reinhold, New York, USA.

SWAN, 2014. Introduction to smart water networks. Available at http://www.swan-forum.com/swan-documents.html. Last accessed 04/01/2014.

Weigend, A., 1994. On overfitting and the effective number of hidden units. In: Proceedings of the 1993 Connectionist Models Summer School. Lawrence Erlbaum Associates Inc., Mahwah, USA.

White, S.B., Fane, S.A., 2002. Designing cost effective water demand management programs in Australia. Water Science & Technology 46 (6-7), 225–232.

Yao, X., 1999. Evolving artificial neural networks. Proceedings of the Institute of Electrical and Electronics Engineers 87 (9), 1423–1447.

Zhang, G., 2001. An investigation of neural networks for linear time-series forecasting. Computers and Operations Research 28 (12), 1183–1202.

Zhou, S., McMahon, T.A., Walton, A., Lewis, J., 2000. Forecasting daily urban water demand: a case study of Melbourne. Journal of Hydrology 236 (3-4), 153–164.

Zhou, S., McMahon, T.A., Walton, A., Lewis, J., 2002. Forecasting operational demand for an urban water supply zone. Journal of Hydrology 259 (1-4), 189–202.