

“©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Predicting Household Water Consumption Events: Towards a Personalised Recommender System to Encourage Water-conscious Behaviour

Md Shamsur Rahim
Centre for Artificial Intelligence,
School of Software, FEIT
University of Technology Sydney
Sydney, NSW, Australia
mdshamsur.rahim@student.uts.edu.au

Khoi Anh Nguyen
School of Engineering and
Built Environment
Griffith University
Brisbane, QLD, Australia
k.nguyen@griffith.edu.au

Rodney Anthony Stewart
School of Engineering and Built
Environment
Griffith University
Brisbane, QLD, Australia
r.stewart@griffith.edu.au

Damien Giurco
Institute for Sustainable Futures
University of Technology Sydney
Sydney, NSW, Australia
damien.giurco@uts.edu.au

Michael Blumenstein
Centre for Artificial Intelligence,
School of Software, FEIT
University of Technology Sydney
Sydney, NSW, Australia
michael.blumenstein@uts.edu.au

Abstract— Recommender systems assist customers to make decisions; however, the modest adoption of digital technology in the water industry means no such system exists for household water users. Such a system for the water industry would suggest to consumers the most effective ways to conserve water based on their historical data from smart water meters. The advantage for water utilities in metropolitan areas is in managing demand, such as low pressure during peak hours or water shortages during drought. For customers, effective recommendations could save them money. This paper presents a novel vision of a recommender system prototype and discusses the benefits both for the consumers and the water utility companies. The success of this type of system would depend on the ability to anticipate the time of the next major water use so as to make useful, timely recommendations. Hence, the prototype is based on a long short-term memory (LSTM) neural network that predicts significant water consumption events (i.e., showers, baths, irrigation, etc.) for 83 households. The preliminary results show that LSTM is a useful method of prediction with an average root mean square error (RMSE) of 0.403. The analysis also provides indications of the scope of further research required for developing a commercially successful recommender system.

Keywords—personalised recommender system, urban water management, water conservation, heavy water consumption, prediction, LSTM.

I. INTRODUCTION

Both prolonged water shortages from drought and low pressure during hours of peak demand can be challenging for metropolitan water utilities [1]. This problem highlights the limitations of the existing water distribution infrastructure and the potential for updated urban water management based on smart technologies to ensure a safe, reliable, and sustainable water supply to consumers [2]. In response to these challenges, policymakers and infrastructure planners are interested in water

demand management (WDM). WDM aims to develop and implement strategies to ensure a more efficient water supply and to promote water conservation measures. There are five categories of WDM: (1) engineering (i.e., upgrading to more water-efficient appliances); (2) economics (e.g., imposing block water tariffs); (3) enforcement (i.e., imposing water restrictions); (4) encouragement (i.e., rewards or rebate schemes for water conservation); and (5) education (i.e., encouraging water conservation in consumers, such as taking shorter showers) [3]. However, identifying and implementing effective WDM strategies requires reliable and preferably real-time information [4]. Smart water meters (SWMs) accompanied by machine learning and data analytics technologies can play a vital role in identifying and implementing successful WDM strategies. For this reason, the importance of SWMs has increased dramatically over the past few years in Australia as well as in other countries [5]. Many studies have been undertaken to utilise the data generated from these SWMs.

Previous studies focusing on SWMs can be broadly divided into three categories: water demand forecasting, water end-use categorisation, and water-use feedback. Depending on the forecast horizon, existing forecasting models can be classified into three classes. On the one hand, long-term forecasting (e.g., the next 20 years) is beneficial for infrastructure and capital planning and medium-term forecasting (e.g., years or months), which is helpful for deciding water rates. On the other hand, very short-term forecasting (e.g., hourly intervals) is crucial for water utilities to plan how to manage demand in advance. Water end-use categorisation provides consumers with a better understanding of their consumption and also a leak detection facility. Lastly, water-use feedback compares water consumptions and provides general level tips that are not highly personalised. Although these three categories contribute in different dimensions, still they do not promote highly

individualised water conservation behaviour and demand management during peak hours.

The previous studies in this field provide a solid foundation to take urban water management to the next level by introducing the concept of a personalised recommender system (RS) [6] based on advanced machine learning and data analytics technologies. RSs are widely popular in the e-commerce, entertainment, services, and social domains due to their ability to help people make decisions in everyday life. Further, in a recent study, RSs were found to be effective in changing behaviours [7]. However, such systems are not commonly available for the water industry.

One goal of a “knowledge-based” water RS would be to promote more effective water conservation behaviour among consumers by providing a list of custom-tailored suggestions based on the historical water consumption data of the individual household. For example, in an effort to promote water conservation during dry seasons, Sydney Water suggested that consumers limit their shower time by four minutes¹. However, if a consumer’s showerhead is not water efficient, then keeping the shower time to four minutes does not help as much. Yet there is no way to suggest that a consumer might consider replacing their showerhead with one that is more water efficient. Similar types of promotions across many different water-saving measures are common for water utilities.

Another goal is to reduce or shift high water consumption events (i.e., shower, bathtub, clothes washer) during peak hours to prevent water shortages or low water pressure. Peak hour is a specific time when the overall demand for water is comparatively high than other times. For instance, existing data shows that a small fraction of consumers tend to perform irrigation and use bathtub during peak hours, which consumes a vast amount of water. Recommending an alternative time and providing incentives based on the extent to which the consumers have adopted the recommendations would be beneficial to managing water supply sustainably and to WDM.

This paper presents the concept behind a prototype for a novel RS that will hopefully play a significant role in promoting water conservation behaviour and formulating effective WDM strategies. However, among the different primary challenges for developing this type of system, identifying the patterns of massive water consumption events is a major one.

Predicting times of high water consumption events is the core of success in a personalised RS that encourages water-conserving behaviour and WDM. It is also one of the most challenging tasks. If the system fails to identify consumption patterns correctly, it will not be able to provide meaningful, personalised recommendations. This will result in customer dissatisfaction, and the system will lose its charm. Besides introducing the concept of a personalised RS, the goal of this paper is to identify patterns in water use and to predict significant water consumption events. For instance, if we could predict that the next clothes washing event is likely to take place at around 7 a.m. tomorrow based on historical water consumption data, which is in peak time, then a recommendation

could be made to shift the event to an off-peak hour. In this way, it would be possible to reduce water demand during peak hours, resulting in efficient pumping activities and increased operational efficiency. However, variations in water consumption patterns occur from time to time, making these predictions more challenging. To overcome this challenge, this paper presents the preliminary experimental results of applying long short-term memory (LSTM), a variant of a recurrent neural network (RNN), to identify patterns in heavy water consumption events and make predictions based on those patterns for future use. The initial results show that LSTM is a good candidate for solving this kind of problem.

The main contributions of this paper are as follows:

- a vision for a personalised recommender system along with a prototype that guides consumers towards more water-conserving behaviours. The benefits of the RS are also discussed from the perspective of both consumers and water utility operators.
- The preliminary findings from an LSTM application to predict high water consumption events.

II. RELATED WORK AND CONTEXT FOR THIS PAPER

Many studies have been conducted to utilise the data generated from SWMs. As the ultimate goal of this research is to promote water conservation behaviours among residential consumers, literature related to commercial and industrial SWMs are beyond the scope of this study.

Many approaches have been proposed to forecast short and long-term water demand based on the periodic patterns in water consumption data [1, 8-10]. Although these models deliver good accuracy for forecasting water demand, the results only benefit the water utilities because they neither help to reduce water consumption during peak hours nor contribute to water conservation behaviour by providing feedback to consumers.

In 2010, Stewart et al. [11] identified categorising the end uses of water as a critical challenge accompanied by a call for urgent attention to new approaches for classifying water end use as the existing methods of analysis were time-consuming. Since then, several classification systems have been proposed [1, 12-14]. Intelligent water management systems like Autoflow© can disaggregate high-resolution data from SWMs into different water end-use categories (i.e., showers, laundry, etc.) with accuracy ranging from 86-94.2% [15]. Although these types of systems are beneficial for detecting leaks or providing a better understanding of water consumption, they do not promptly suggest what consumers might do conserve water or how they might change their behaviour.

According to Liu et al. [16], the detailed feedback allowed by SWMs has positive impacts on more sustainable urban water management. Further, this study also reported that online-based feedback reports could be more effective than paper-based reports if the low engagement rate with online feedback could be overcome. However, the literature shows that attempts to incorporate online reports into feedback systems also lack timely personalised information on water-conserving behaviours [16].

¹ <http://www.sydneywater.com.au/SW/your-home/using-water-wisely/water-wise-tips/index.htm>

The limitations of the existing studies clearly show the research gap in an advanced personalised water-use feedback system. To promote water conservation behaviour among consumers, state of the art machine learning techniques must be combined with data analytics techniques within a personalised RS if they are to make a significant contribution to managing water more sustainably.

III. VISION OF THE RECOMMENDER SYSTEM

Water demand forecasting and categorising water events can be considered as the beginning of integrating artificial intelligence and machine learning into the water sector.

recommendations will be given to consumers: personalised and non-personalised.

Personalised recommendations will be made to the consumers based on water consumption patterns from the consumer’s historical data and their preferred *items*. Here, the term *item* refers to each of the activities that a consumer wishes to receive recommendations about. After analysing historical water events, water consumption patterns will be extracted. Then, based on these patterns, initial recommendations will be made to the consumers. Further, the system will collect feedback in two forms: *explicit* and *implicit*. By rating each of the recommendations, consumers will provide explicit feedback,

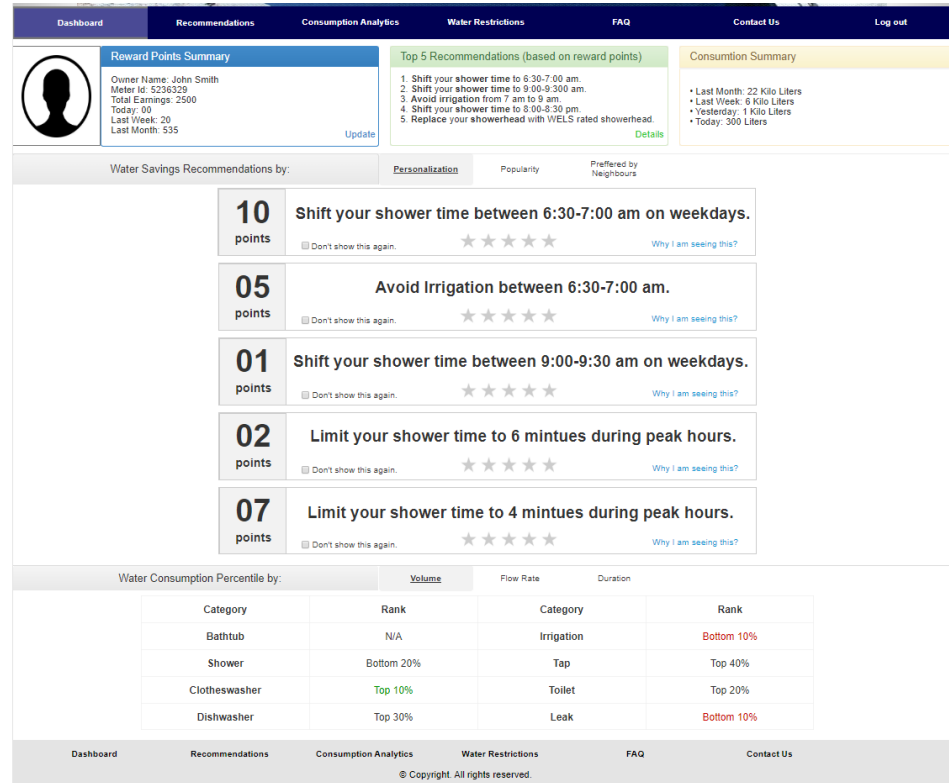


Fig. 1. Prototype of the Dashboard for the proposed web-based knowledge management system.

However, autonomous end-use analysis systems for water with higher accuracy, such as Autoflow©, raise further questions like “What’s next?” and “How can these analysed water events data be used for further improvement?”. To answer these questions, we present a vision for a recommender system that analyses water consumption data and recommends water-saving behaviour. However, the proposed system is not merely a recommender system. Rather, once fully-fledged, it will be a complete web-based knowledge management system. Figure 1 illustrates the dashboard of the system prototype.

The system comprises two components: a recommender system and reward/rebate system. The following subsections discuss these two components in more detail.

A. The Recommender System

The primary goal of this system is to bring change to the water consumption behaviour of consumers by providing guidance. To achieve this objective, two main types of

which will be used to improve the relevancy of the recommendations in future. For instance, if a consumer submits a lower rating for recommendations related to shifting shower time, then subsequent recommendations will be less likely to contain this type of suggestion. Implicit feedback will be collected from the water consumption data. For instance, if the consumer’s water consumption patterns show that an average shower is around 6 minutes, recommending a new shower time of 4-5 minutes is likely to be more enticing than, say, 2 minutes. Implicit recommendations will particularly benefit from integrating demographics with a content-based technique and result in a better hybrid recommender system [17].

Some consumers may not be prone to providing explicit feedback. This makes it difficult to understand their preferences since only implicit feedback is available. To overcome this challenge, the system will employ another recommender

technique where the top recommendations made to consumers with similar consumption patterns are used as a compliment. This technique is known as collaborative filtering [18]. In this context, a hybrid RS that combines collaborative filtering and demographic techniques will perform better than relying too heavily on explicit feedback. Figure 2 illustrates the prototype of the personalised recommendations. To complement these recommendations, the vision is to also include the most popular

B. Reward/rebate Scheme

If a consumer changes their behaviour based on the recommendations made by the system, they should be rewarded. We argue that motivating consumers to be more water efficient over the long-term will be easier to achieve if coupled with an incentive scheme. Based on this idea, the system includes a reward/rebate scheme that offers virtual points for complying

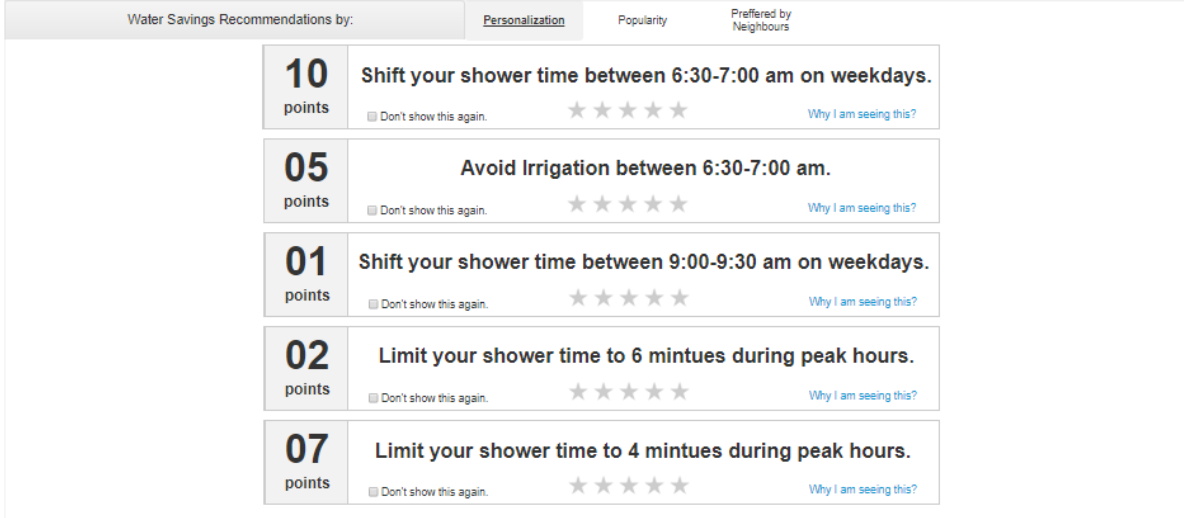


Fig. 2. Prototype of the personalized recommendations. recommendations among similar consumers.

Non-personalised recommendations are based on summary statistics. The motivation behind this type of recommendation is to provide a clear understanding of a consumer’s standing compared to other consumers based on three water consumption metrics for each category of water event: volume, flow rate, and duration. This idea is similar to what we see in electricity bill, but more comprehensive. Decile statistics can help to accomplish this goal, where nine values are used to arrange the data in ten equal parts. The following formula is used to calculate these nine values:

$$D_i = \frac{i(n+1)}{10} \quad (1)$$

where $i=1,2,3,\dots,9$ and n is the number of observations.

This will provide consumers with insights into their water consumption behaviours as well as showing areas for improvement. Figure 3 shows the part of the prototype that

with the recommendations. The number of points assigned to each recommendation can be assigned by the water utility and varied in different seasons to reflect the most valuable suggestions based on current needs. To tally the points earned, a custom algorithm operating at the customer account level will automatically calculate the recommendations followed from the consumption data. Depending on the policy set by the operator, consumers then claim a water rebate based on their reward points. Water utilities might allocate a fraction of the money saved by using the system to their rebate program. In this way, both the consumers and water utilities will be benefitted from the system. Further benefits of the system are discussed in the following section.

C. System Benefits for Water Consumers and Utilities

1) *Benefits to consumers:* The end consumers of water are the biggest beneficiaries of this type of recommender system, save both water and money. For instance, a recommendation by

Water Consumption Standings by:		Volume	Flow Rate	Duration
Category	Rank	Category	Rank	
Bathtub	N/A	Irrigation	Bottom 10%	
Shower	Bottom 20%	Tap	Top 40%	
Clotheswasher	Top 10%	Toilet	Top 20%	
Dishwasher	Top 30%	Leak	Bottom 10%	

Fig. 3. Prototype of the Non-personalized recommendations.

displays the summary statistics for each water event category grouped by volume, flow rate, and duration.

the system to limit shower times to 4 minutes during the dry season is likely to be more effective than an awareness program. Evidence to this effect comes from Moglia et al. [19],

where smart metering feedback to the consumer reduced water consumption by approximately 20%. The reward points/rebate program will also encourage consumers to conserve water, reducing their overall water bill coupled with further discounts for wise water use.

Further, a system such as this should be able to increase, or at least maintain, consumer satisfaction levels by providing customised recommendations. Moreover, consumers can directly communicate with their water utility company through implicit and explicit feedback on the recommendations being made. For the water utility, this feedback could be used to improve the relevance of future recommendations.

From a consumer's viewpoint, this recommender system has the potential to benefit them both economically and through a higher quality of service.

2) *Benefits to the water utilities*: More and more, water utilities are facing new challenges: water shortages, low water pressure during peak hours, and rising supply costs. This may mean water utilities must continue pumping activities even during peak hours for electricity, or operate extra pumps to meet demand in peak times. According to Nguyen [1], advising consumers to avoid or limit heavy water consumption events, such as irrigation and showers, could result in a significant reduction in water and energy consumption during peak hours in urban areas. However, it would also mean a reduction in the amount of greenhouse gases released into the environment. With that in mind, this system could play a significant role in conveying desired usage guidelines to consumers smartly and efficiently. For instance, the system can recommend the optimal time to undertake a heavy water consumption event. By doing this, it may be possible to reduce urban water consumption significantly during peak hours. This alone would improve quality of service, minimise operating cost, increase profit, and serve as a better management and planning tool.

Perhaps most importantly, the system would provide a better understanding of consumers and their opinions on particular types of water-conserving behaviour – what works, what does not. Based on this knowledge, the system could play an important role in devising and conveying water restrictions to consumers at little cost compared to other awareness campaigns.

D. Challenges for developing the RS

Developing an RS for the water industry is unique and highly complex in comparison to other domains. For instance, an e-commerce RS can recommend similar items based on a previous purchase history or by product association. Further, the provider's primary, and often singular, goal is to increase sales and turn more profit. While the goal for a water utility is just the opposite – to conserve water. Additionally, consumers have very different patterns of water use, so it may not be possible to inform recommendations with the use of association rules between different water events, such as recommending shower time based on common laundry times. Thus, developing an RS for the water industry needs to be considered from a new viewpoint. These among a few other challenges are summarised below.

- 1) Identifying water consumption patterns to anticipate the time of the next water usage event of interest.
- 2) Extracting implicit feedback from water consumption data.
- 3) Clustering consumers with similar consumption patterns across different water events (e.g., showers, bathtub, irrigation, laundry).
- 4) Ensuring the system is scalable and reliable.
- 5) Integrating recommender techniques based on demographic data.
- 6) Limiting the number of recommendations to the most important items.
- 7) Determining the best recommender technique.

IV. PROPOSED APPROACH AND EXPERIMENTAL SETTINGS

To make useful recommendations, the system needs to be able to anticipate when heavy water consumption events are likely to occur with as much accuracy as possible. This section presents the proposed approach for predicting these events using LSTM. The following subsections discuss the approach in further detail.

A. Dataset collection and pre-processing

The dataset used in this preliminary study originated from Autoflow© [14], which is intelligent software that autonomously classifies residential water end use with a high accuracy rate. After analysing the raw dataset collected from the smart meters of 83 households for seven months from February 2010 to August 2010, Autoflow© provided the classifications of water end-use events: taps, dishwashers, leaks, evaporating coolers, clothes washers, showers, toilets, irrigation, and bathtubs.

The next step was to prepare the dataset for further analysis. As the focus of this study is to predict heavy water consumption events, random events with less water consumption, such as leaks, evaporating coolers, taps, and toilets were removed. Cyclic events, such as clothes washing and dishwashers were aggregated since water consumption is regular enough to be considered as a single event with an identified cycle. The last step was dimension reduction. The software provides different features include start date, start time, end date, end time, category, duration, volume, max flow, mode, and cyclic event. For our purposes, only the start date, start time and category attribute were necessary. Hence, the rest of the attributes were ignored and the start date and time attributes were aggregated to further reduce dimensionality.

B. Dataset Preparation for LSTM

The pre-processed dataset was then prepared for training the LSTM network. This involved framing the dataset as a supervised learning problem and encoding the input variables to predict the event at the current hour (t) given the events at the prior time steps. Three columns of lag observations ($t-1$, $t-2$, $t-3$) and one column of forecast observations ($t+1$) were created using Pandas shift function, which creates copies of columns by either pushing the values forward or backward. The rows that contained *NaN* values were dropped after this step. As a last step, the category feature was label encoded using an integer

encoder, and the dataset was transformed into a supervised learning problem [20].

C. LSTM Configuration

LSTMs are special kind of RNN, introduced by Hochreiter & Schmidhuber in 1997 [21]. As LSTMs are capable of learning long-term dependencies and are able to solve a variety of problems, they have become widely popular in different domains including time series forecasting. A common architecture for an LSTM model consists of a memory component, also called a cell, and three “regulators”, known as gates, to control the flow of information to the LSTM model.

The input to the input gate is first squashed between -1 and 1 using a *tanh* activation function expressed as

$$g = \tanh(b^g + x_t U^g + h_{t-1} V^g) \quad (2)$$

where b^g is the input bias, and U^g and V^g are the weights for the input and previous cell’s output, respectively.

Next, the squashed input is multiplied element-wise by the input gate’s output, which consists of a hidden layer of sigmoid-activated nodes, as a value of between 0 to 1. When multiplied element-wise, the gate decides which inputs to take or discard. In short, the input gate acts as a filter for the inputs. The input gate is formulated as

$$i = \sigma(b^i + x_t U^i + h_{t-1} V^i) \quad (3)$$

The next stage of the LSTM cell is the forget gate, which determines what information should be removed from the cell. This task is accomplished by a sigmoid layer called the “forget layer” using the following equation:

$$f = \sigma(b^f + x_t U^f + h_{t-1} V^f) \quad (4)$$

The output from the forget gate is multiplied element-wise with s_{t-1} to provide an internal recurrence loop with which to learn the relationship between the inputs separated by time. The product of the multiplication determines which previous steps should be remembered and which should be forgotten.

Using these three gates, LSTMs are designed in such a way that they can remember information for a long time without any difficulty. Due to this ability, LSTMs are a good choice for solving prediction problems associated with heavy water consumption events.

As a preliminary study to investigate whether and how LSTMs could be used to predict significant water consumption events, we deployed a simple LSTM network for each of the 83 households. The LSTMs were implemented using Keras [22], a high-level neural network API, that runs on top of TensorFlow [23].

To fit the LSTM, we split the dataset into a training and a test set. Given that training separate LSTM models for each household would consume too much time, we speed up the process by fitting 60 per cent of the data to the model, and using the remaining 40 per cent for evaluation. Next, the training and test sets were split into input and output variables and the inputs were reshaped into the three-dimensional format required by

LSTMs. For the purposes of this preliminary investigation, we defined a simple LSTM network with 50 neurons in the first hidden layer and one neuron in the output layer to predict events in one-time step ahead. We fit the model over 50 training epochs with a batch size of 60. Mean absolute error (MAE) served as the loss function, and efficient Adam [24] (a version of stochastic gradient descent) was used as the optimisation function. Once the models were fitted, prediction and evaluation were conducted over the entire test sets.

V. EXPERIMENTAL RESULTS AND DISCUSSION

To verify the models’ fitness, we calculated the loss scores for each training and validation epoch and recorded them in a history trace. Plotting these values provide an understanding of the fitness of the model. If a model performs better on the training dataset and poor on the test dataset, then it is considered as underfit model. This means the training loss is lower than the validation loss. In the case of overfit model, the model is good for the training set and keeps improving; however, for the validation set, after a point the performance begins to degrade. In case of the good fit model, the train and validation loss decrease and almost stabilise at some point.

Figure 4 plots the history trace for one household. The model stabilises after the 10th epoch, which indicates it is a good fit for the data.

Next, we measured the model’s performance using root mean square error (RMSE). Note that LSTMs are stochastic, so each run may result in a different RMSE score. Therefore, we trained and evaluated the models five times for each household and calculated the RMSE each time. The average and




Fig. 4. Training and validation loss of the LSTM model for a sample household.

lowest RMSE of the five runs for each household is reported in Figure 5.

The overall average was 0.403 with the lowest RMSE at 0.007, which verifies the suitability of LSTM for predicting heavy water consumption events. Further, the difference between the average RMSE and best RMSE is very low, which shows the consistency of the model. More interestingly, there was great variation in the RMSE for each household. This is due to the different levels of consistent and inconsistent water consumption behaviours among the households.




Fig. 5. Comparison of Average and Best RMSE for each household.

From this preliminary study, we conclude that LSTM is a good candidate for predicting household water consumption events. However, this only a preliminary study; there is a great deal of further scope for research and improvement, which is discussed in the following section.

VI. CONCLUSIONS AND FUTURE WORK

This paper presented a novel concept for a personalised recommender system to promote water conservation behaviour among consumers for sustainable urban water management. The potential benefits both for the consumers and water utilities from this system are great, but its success largely depends on the ability of the system to accurately capture and then forecast water consumption behaviour by consumers – especially significant water consumption events like bathtub, showers, irrigation, and so on. As a step towards realising a personalised recommender system, we built a simple LSTM network to predict high water consumption events for 83 households. Our preliminary experiments demonstrate that LSTM is a good candidate for solving this type of problem. However, there is much work still to be done.

For example, currently the system requires a new LSTM network to be trained for each household. Given the vast deployment of SWMs, this approach is not scalable to the real world, and this must be addressed. The next step is to explore clustering techniques to group households based on water consumption patterns to overcome this limitation. Additionally, we plan to look to training separate models and multilabel classification to predict each type of event as further improvements. Once the foundations of the model are more secure, we will progress comprehensive development of the recommender system.

ACKNOWLEDGEMENTS

Project support is acknowledged from the Australian Research Council (ARC) Linkage Project (LP160100215) that also includes funding from the following industry partners: Yarra Valley Water, City West Water, Southeast Water, and Aquiba.

REFERENCES

[1] O. S. Khoi Anh Nguyen, Rodney Anthony Stewart, Hong Zhang, "Water demand forecasting with AUTOFLOW© using State-Space approach," presented at the 8th International Congress on Environmental Modelling and Software, 2016.

[2] K. A. Nguyen, R. A. Stewart, H. Zhang, O. Sahin, and N. Siriwardene, "Re-engineering traditional urban water management practices with smart metering and informatics," *Environmental Modelling & Software*, vol. 101, pp. 256-267, 2018.

[3] A. Liu, D. Giurco, and P. Mukheibir, "Urban water conservation through customised water and end-use information," *Journal of Cleaner Production*, vol. 112, pp. 3164-3175, 2016.

[4] O. Sahin, R. A. Stewart, and M. G. Porter, "Water security through scarcity pricing and reverse osmosis: a system dynamics approach," *Journal of Cleaner Production*, vol. 88, pp. 160-171, 2015.

[5] C. D. Beal and J. Flynn, "Toward the digital water age: Survey and case studies of Australian water utility smart-metering programs," *Utilities Policy*, vol. 32, pp. 29-37, 2015.

[6] P. Resnick and H. R. Varian, "Recommender systems," *Communications of the ACM*, vol. 40, no. 3, pp. 56-58, 1997.

[7] O. Yürüten, "Recommender Systems for Healthy Behavior Change," 2017.

[8] S. L. Zubaidi, J. Dooley, R. M. Alkhattar, M. Abdellatif, H. Al-Bugharbee, and S. Ortega-Martorell, "A Novel approach for predicting monthly water demand by combining singular spectrum analysis with neural networks," *Journal of Hydrology*, vol. 561, pp. 136-145, 2018.

[9] J. E. B. Pesantez, Emily Zechman; Kaza, Nikhil "Modeling and Forecasting Short-Term Water Demand Reported by Smart Meters," presented at the 1st International WDSA / CCWI 2018 Joint Conference, Kingston, Ontario, Canada, 2018.

[10] V. B. Platsko, Peter van, "Identification, Prediction, and Explanation of Outdoor Residential Water Consumption Using Smart Meter Data2018," presented at the 1st International WDSA / CCWI 2018 Joint Conference, Kingston, Ontario, Canada, 2018.

[11] R. A. Stewart, R. Willis, D. Giurco, K. Panuwatwanich, and G. Capati, "Web-based knowledge management system: linking smart metering to the future of urban water planning," *Australian Planner*, vol. 47, no. 2, pp. 66-74, 2010.

[12] S. A. McKenna, F. Fusco, and B. J. Eck, "Water Demand Pattern Classification from Smart Meter Data," *Procedia Engineering*, vol. 70, pp. 1121-1130, 2014.

[13] M. L. Abadi *et al.*, "Predictive Classification of Water Consumption Time Series Using Non-homogeneous Markov Models," presented at the 2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA), 2017. Available: <https://ieeexplore.ieee.org/ielx7/8255765/8259747/08259791.pdf?tp=&number=8259791&isnumber=8259747>

[14] K. A. Nguyen, R. A. Stewart, H. Zhang, and C. Jones, "Intelligent autonomous system for residential water end use classification: Autoflow," *Applied Soft Computing*, vol. 31, pp. 118-131, 2015.

[15] A. Yang, H. Zhang, R. Stewart, and K. Nguyen, "Enhancing Residential Water End Use Pattern Recognition Accuracy Using Self-Organizing Maps and K-Means Clustering Techniques: Autoflow v3.1," *Water*, vol. 10, no. 9, 2018.

[16] A. Liu, D. Giurco, and P. Mukheibir, "Advancing household water-use feedback to inform customer behaviour for sustainable urban water," *Water Science and Technology: Water Supply*, vol. 17, no. 1, pp. 198-205, 2017.

[17] F. Ricci, L. Rokach, and B. Shapira, "Recommender Systems: Introduction and Challenges," in *Recommender Systems Handbook*, 2015, pp. 1-34.

[18] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using collaborative filtering to weave an information tapestry," *Communications of the ACM*, vol. 35, no. 12, pp. 61-70, 1992.

[19] M. Moglia, S. Cook, and S. Tapsuwan, "Promoting Water Conservation: Where to from here?," *Water*, vol. 10, no. 11, 2018.

[20] J. Brownlee. (2017, 20/12/2019). *How to Convert a Time Series to a Supervised Learning Problem in Python*. Available: <https://machinelearningmastery.com/convert-time-series-supervised-learning-problem-python/>

[21] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735-1780, 1997.

[22] F. Chollet, "Keras: The python deep learning library," *Astrophysics Source Code Library*, 2018.

[23] M. Abadi *et al.*, "Tensorflow: a system for large-scale machine learning," in *OSDI*, 2016, vol. 16, pp. 265-283.

- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.