Journal of Physics: Conference Series

PAPER • OPEN ACCESS

Prediction of Domestic Hot Water Temperature in a District Heating Network

Prediction of Domestic Hot Water Temperature in a District Heating Network

Frédéric Montet[1], Lorenz Rychener[1], Alessandro Pongelli[2], Jean Hennebert[1], Jean-Philippe Bacher[2]

iCoSys Institute[1], *HEIA-FR*, Fribourg, Switzerland
ENERGY Institute[2], *HEIA-FR*, Fribourg, Switzerland

E-mail: `frederic.montet@hefr.ch, lorenz.rychener@hefr.ch, alessandro.pongelli@hefr.ch, jean.hennebert@hefr.ch, jean-philippe.bacher@hefr.ch`

Abstract.
With the fourth generation of district heating networks in sight, opportunities are rising for better services and optimized planning of energy production. Indeed, the more intensive data collection is expected to allow for load prediction, customer profiling, etc. In this context, our work aims at a better understanding of customer profiles from the captured data.

Given the variety of households, such profiles are difficult to capture. This study explores the possibility to predict domestic hot water (DHW) usage. Such prediction is made challenging due to the presence of two components in the signal, the first one bound to the physical properties of the DHW distribution system, the second one bound to the human patterns related to DHW consumption.

Our contributions include (1) the analysis of recurrent neural network architectures based on GRU, (2) the inclusion of state-based labels inferred in an unsupervised way to simulate domain knowledge, (3) the comparison of different features.

Results show that the physical contribution in the signal can be forecasted successfully across households. On the contrary, the stochastic "human" component is harder to predict and would need further research, either by improving the modelling or by including alternate signals.

1. Introduction

Since the innovation of Frédéric Sauter and its electric water heater called the Cumulus in 1915 [1], domestic hot water (DHW) has been heated up with various methods: on-demand, with solar panels, heat pump, and more [2, 3]. What mainly happened inside households is now getting deported with heat provided by district heating networks (DHN).

Today, the 4th generation of DHN is being developed. Such networks will include the understanding of their connected customers to produce the right heat load [4, 5]. Understanding the water usage and the heat loads is a complex task [6, 7, 8]. The network is composed of dwellings with many different heating systems having their own behavior. Furthermore, the usage of DHW depends on human decisions; a long-known source of stochasticity inherently hard to predict. To take an example, when a Swiss resident is taking a shower, its duration can vary from two to fifteen minutes. Such a DHW draw represents around 44 liters and 2,6 kWh on average[9]. Furthermore, the usage of DHW has a daily probability distribution varying greatly between households [10]. For a DHN operator, such draws can be absorbed during the cold season with the constant buffer required for heating purposes. In mild and warm seasons, the aforementioned buffer is much smaller. Planning the right heat load becomes a much more challenging task to meet unpredicted demand.

Therefore, this paper focuses on the understanding of DHW usage in households and its predictability. DHW temperature in the hot water tank is used as a proxy to represent both the heating system and its usage. Given the human nature to have daily rhythms [11], we hypothesize that the combination of human and household takes different states that a machine learning model should model successfully. If reliable, such a prediction could pave the way towards temperature reduction on the overall network during mild and warm seasons. The heat load planning would take into account the individual need of each household, thus saving a potentially significant amount of energy.

The paper presents machine learning models to predict DHW temperature within eight common Swiss households. The method provides insights on the data selection and the modelling which includes different feature sets. Finally, the results are discussed and put the behavior of the households into perspective.

2. Method

2.1. Selection of Eight Households

The data is sourced from a DHN connecting approximately 170 buildings of various natures (households, commercial, etc.) located in western Switzerland. With a combination of physical and stochastic properties, the numerous measurements provide an interesting playground for applied research on time series.

From this dataset, 8 households within 5 different semi-detached houses are selected for their representation in the data. In 1, letters A to E denote each house and indices $i \in \mathbb{N}$ distinguishes its households—i.e. all households A_i belong to the same house A.

The data was collected between September 2018 to October 2019. The 2 minutes sampling has been preserved to capture most DHW draws with accuracy. Each dwelling has its own heating system where DHW temperature is measured at the lower part of the water tank. Additionally, the dataset includes the outside temperature at each time step.

	A_1	A_2	B_1	B_2	C_1	C_2	D_1	E_1
Power $[kW]$	10	10	10	10	10	10	20	20
Energy $[kWh/year]$	10'000	10'000	10'000	10'000	15'000	15'000	40'000	40'000
Circulation pump	-	-	A	D	-	-	A	A

Table 1. Description of households with their power and energy subscriptions. Circulation pumps are either not present (-), activated (A) or deactivated (D)

2.2. Unsupervised Labelling of the DHW Temperature Time Series

To reinforce the distinction between the different DHW patterns, time series segments are clustered into four groups: heating, cooling, cooling with circulation pump and water usage. Post-clustering, each time step in the time series is associated with a group label.

The labelling sequence has three stages. First, temperature change points are localized by finding local minima and maxima within a 20 minutes window. Second, the time series is split at all minima and maxima, thus dividing cooling and heating periods. Finally, all segments are clustered using k-Means clustering with Dynamic Time Warping (DTW) as a distance measure. The number of clusters k varies between $k = \{3, 4\}$ depending on the presence of a circulation pump.

To evaluate the quality of the clustering procedure, the Silhouette scores [12] are computed on 100 random samples per cluster and averaged over five repetitions.

2.3. Modelling of the Eight Households with Different Feature Sets

Each house is modelled individually with Gated Recurrent Units (GRUs) [13], a simpler variant of LSTM [14]. Both are forms of Recurrent Neural Networks (RNN) able to capture temporal dependencies, which makes them widely applied for time series prediction.

Since the goal of this study is the exploration of DHW usage across households and its potential predictability, the use of a reference model has been favoured over the research of optimal architecture and modelling. The model has 2 layers of GRUs with 13 neurons and is trained with a 10% dropout for regularization.

The training set is one year long, from the beginning of September 2018 to the end of August 2019. The validation set is the full month of September 2019. The number of time steps given as input varies to fit the dynamics of each house heating system. The window length is found by using the autocorrelation or the similarity between observation and their lag time. The prediction output is given by a business need to predict the next hour; i.e. 30 time steps. Within such a time frame, a DHN operator could react early enough to meet the expected demand.

To observe the effect of each feature on the model output, different features sets F_i have been composed with: DHW temperature T, holidays H, weekdays W, clustering labels L described in section 2.2 and outside temperature O. First, a univariate set (for reference) with $F_1 = \{T\}$. Then, three multivariate sets with $F_2 = \{T, H, W, L\}$, $F_3 = \{T, H, W, O\}$ and $F_4 = \{T, H, W, O, L\}$. Features H, W, L are one-hot encoded. Those 4 features sets have been used to model each household, thus creating 32 models in total.

2.4. Evaluation of the Models with Hourly Historical Forecasting for a Month

To simulate a real scenario, hourly historical forecasts are computed for the duration of the validation set (September 2019). Then, the mean absolute percentage error (MAPE) is computed to allow for comparison between households and models.

3. Results

Table 2 shows the Silhouette scores of the heating event clustering [1]. As all scores are close to zero, it indicates a low membership of the samples to their cluster. Furthermore, this shows an overlap between clusters. The presence of a circulation pump in the system doesn't influence the scores.

	A_1	A_2	B_1	B_2	C_1	C_2	D_1	E_1
Silhouette score	0.10	0.07	0.08	0.09	0.10	0.13	0.14	0.06

Table 2. Silhouette scores of the events described in 2.2 for all households.

Most of the 32 models did converge acceptably within 12 to 20 epochs. Figure 1 shows the convergence of one of them. On some models, the validation loss did show an overfitted behavior which did not influence the study intent.

Table 3 presents the MAPE score for each model for all studied households during its evaluation on the validation. The low scores indicate good predictability of the data. The addition of multiple features to the models doesn't significantly affect the predictions, nor does

[1] The Silhouette score shows the degree of membership to a cluster. It ranges from -1 to 1; 0 indicating no membership.

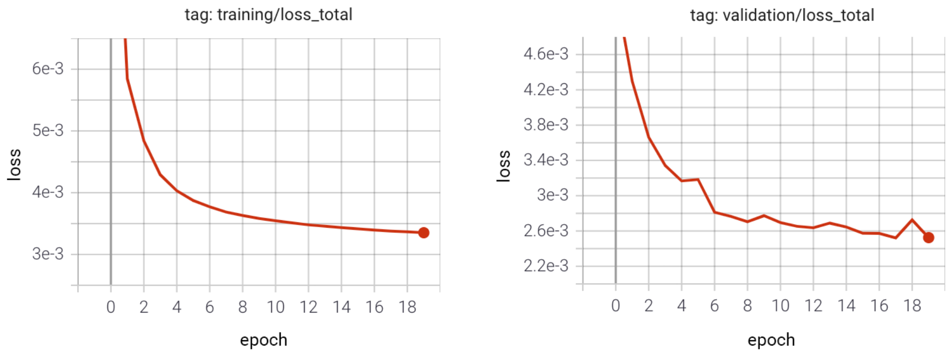

Figure 1. Training and validation loss of the univariate model $F_1 = \{T\}$ on B_2.

the presence of circulation pumps. Given the low Silhouette scores, which implies overlapping clusters, the addition of the cluster labels to the models doesn't bring much information to improve the model output. The higher MAPE when the feature set F_i includes the labels L confirm the latter statement.

Across the studied households from Table 1, scores vary greatly. Those differences are visible independently of the household's context; being in the same semi-detached house (e.g. B_1, B_2), between semi-detached houses with different energy subscription (e.g. A_i, D_1) or across (e.g. A_1, E_1). It should also be noted that scores are all staying within acceptably low values.

Feature set	A_1	A_2	B_1	B_2	C_1	C_2	D_1	E_1	Mean
$F_1 = \{T\}$	1.10	1.82	1.05	2.61	3.43	1.31	0.56	2.22	1.76
$F_2 = \{T, H, W, L\}$	1.56	2.18	1.49	2.97	4.46	1.76	1.00	2.27	2.21
$F_3 = \{T, H, W, O\}$	1.10	1.95	0.99	2.58	2.80	1.37	0.87	2.06	1.72
$F_4 = \{T, H, W, O, L\}$	1.56	2.47	1.69	2.96	3.89	1.63	0.97	2.19	2.17

Table 3. MAPE Score of the validation set by household and feature set; lower is better. The features given to each model are a combination of the DHW temperature (T), holidays (H), weekdays (W), clustering labels (L) and outside temperature (O).

4. Discussion

Modelling the DHW temperature as a proxy to predict human behavior is difficult. When confronted with a signal composed of both a physical system and a human component behaving in spikes, GRU models are not replicating the latter. This behavior is visible in Figure 2; the predictions are not following the spikes (DHW usages) from the measurements.

Increasing the model complexity with external features is of limited help. The addition of holidays, weekdays, clustering labels, and outside temperature to the DHW temperature did not significantly improve the models' performances. With additional features, the training procedure got more expensive for similar, or worse, performances.

The addition of more features could have worsened the performance for three main reasons. First, the DHW temperature T might be independent of the other variables. Thus, their addition doesn't add information. Second, the extra features H, W, O, supposed to reinforce the human behavioral signal, are not as relevant as expected. Third, the feature L simulating expert knowledge might be erroneous, thus increasing the overall entropy. When associated, those features might allow the creation of spikes in the prediction but worsening the overall score.

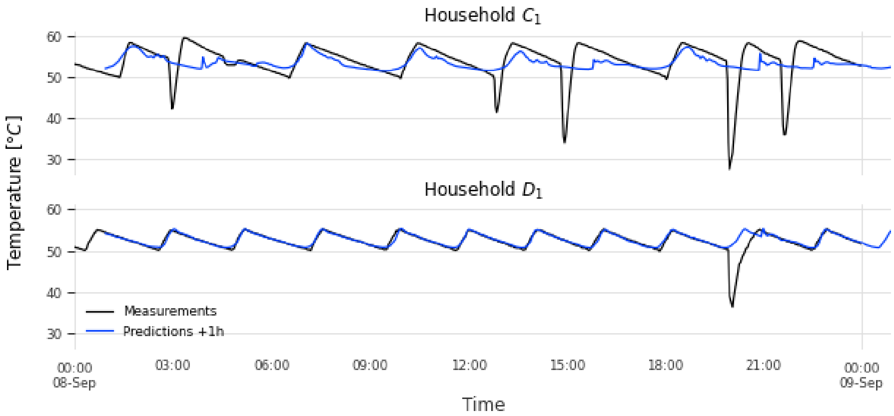

Figure 2. Comparison of two hourly predictions for a full day. The upper plot shows a difficult prediction scenario and the lower plot, an easier one.

In this context, keeping a simple prediction system seems to be a reasonable approach for the non-stochastic part. Across all households, the univariate models did successfully learned the physical properties of the heating system; with or without a circulation pump. Any other added features did only marginally affected the MAPE scores.

Modelling human behavior, within the DHW usage, has shown to be a non-trivial task. The comparison between the upper and lower plot from Figure 2, together with their MAPE scores in Table 3, highlights the differences between households. Even though the models recover after a DHW usage, they are unable to model the spikes representing DHW. With the simplistic model used in this study of the combination of human stochasticity and the physical properties of the system, only the regular cycles of water cooling and heating can be predicted. Therefore the human behavior in DHW is still considered a deviation from normality and to be treated as anomalous within this context.

To improve stochastic modelling two directions are worth considering: a better model or a more elaborate method. In the former, further investigation in the time series clustering methods could provide good results in association with a state-of-the-art model like *Transformers for Time Series* [15]. In the context of a DHN with numerous connected households, training time could become computationally expensive. Furthermore, DHW prediction is referring to the prediction of spikes (i.e. water usage) where false predictions could induce unused heat loads in the DHN central. The latter, a more elaborate method, favours the prediction of by-signals (e.g. DHW spikes number per day) such as presented by Rathnayaka et al. [16]. For instance, the univariate model (modelling the physical system) could be used as a filter to obtain only the stochastic part of the compound signal. Then, a further modelling method could be envisioned.

5. Conclusion

In this paper, we highlighted the complexity of modelling DHW temperature in the context of a DHN connecting numerous households in a newly obtained dataset. The DHW temperature signal embeds two components: an underlying physical system and its haphazard human usage. Out of those two components, only the physical system part has been successfully modelled. Trials to augment the signal with additional features did only worsen the prediction scores, which led to the exclusion of the initial hypothesis. Given the good prediction results obtained to model the DHW physical system across households, the research suggests filtering the full signal with a univariate model prediction; leaving only the stochastic part of it for further research.

References

[1] Cumulus water heater specifications, procédés sauter ag, 1929

[2] Manokar A M and Karthick A 2021 *Environmental Science and Pollution Research* 1–14

[3] Cabeza L F 2021 2 - advances in thermal energy storage systems: methods and applications *Advances in Thermal Energy Storage Systems (Second Edition)* Woodhead Publishing Series in Energy ed Cabeza L F (Woodhead Publishing) pp 37–54 second edition ed ISBN 978-0-12-819885-8 URL `https://www.sciencedirect.com/science/article/pii/B9780128198858000024`

[4] Lund H, Werner S, Wiltshire R, Svendsen S, Thorsen J E, Hvelplund F and Mathiesen B V 2014 *Energy* **68** 1–11 ISSN 0360-5442 URL `https://www.sciencedirect.com/science/article/pii/S0360544214002369`

[5] Connolly D, Lund H, Mathiesen B V, Werner S, Möller B, Persson U, Boermans T, Trier D, Østergaard P A and Nielsen S 2014 *Energy policy* **65** 475–489

[6] Willis R M, Stewart R A, Giurco D P, Talebpour M R and Mousavinejad A 2013 *Journal of Cleaner Production* **60** 107–115 ISSN 0959-6526 special Volume: Water, Women, Waste, Wisdom and Wealth URL `https://www.sciencedirect.com/science/article/pii/S0959652611003064`

[7] Widén J, Lundh M, Vassileva I, Dahlquist E, Ellegård K and Wäckelgård E 2009 *Energy and Buildings* **41** 753–768 ISSN 0378-7788 URL `https://www.sciencedirect.com/science/article/pii/S0378778809000413`

[8] Mostafavi N, Gándara F and Hoque S 2018 *Energy and Buildings* **158** 1683–1693 ISSN 0378-7788 URL `https://www.sciencedirect.com/science/article/pii/S0378778817329973`

[9] Shahmohammadi S, Steinmann Z, King H, Hendrickx H and Huijbregts M A 2019 *Journal of Industrial Ecology* **23** 1186–1195 (*Preprint* `https://onlinelibrary.wiley.com/doi/pdf/10.1111/jiec.12858`) URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/jiec.12858`

[10] de Santiago J, Rodriguez-Villalón O and Sicre B 2017 *Energy and Buildings* **141** 341–348 ISSN 0378-7788 URL `https://www.sciencedirect.com/science/article/pii/S0378778817305911`

[11] Fuentes E, Arce L and Salom J 2018 *Renewable and Sustainable Energy Reviews* **81** 1530–1547 ISSN 1364-0321 URL `https://www.sciencedirect.com/science/article/pii/S1364032117308614`

[12] Rousseeuw P J 1987 *Journal of Computational and Applied Mathematics* **20** 53–65 ISSN 0377-0427 URL `https://www.sciencedirect.com/science/article/pii/0377042787901257`

[13] Cho K, van Merrienboer B, Gulcehre C, Bahdanau D, Bougares F, Schwenk H and Bengio Y 2014 Learning phrase representations using rnn encoder-decoder for statistical machine translation (*Preprint* `1406.1078`)

[14] Hochreiter S and Schmidhuber J 1997 *Neural computation* **9** 1735–1780

[15] Cohen M, Charbit M, Corff S L, Preda M and Nozière G 2020 End-to-end deep metamodeling to calibrate and optimize energy loads (*Preprint* `2006.12390`)

[16] Rathnayaka K, Malano H, Arora M, George B, Maheepala S and Nawarathna B 2017 *Resources, Conservation and Recycling* **118** 1–12 ISSN 0921-3449 URL `https://www.sciencedirect.com/science/article/pii/S0921344916303317`