

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/340355946>

Safe energy savings through context-aware hot water demand prediction

Article in *Engineering Applications of Artificial Intelligence* · April 2020

DOI: 10.1016/j.engappai.2020.103481

CITATION

1

READS

25

4 authors, including:



Josh Siegel

Michigan State University

60 PUBLICATIONS 473 CITATIONS

SEE PROFILE



Yongbin Sun

mit

24 PUBLICATIONS 1,187 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Cloud vehicle mirroring [View project](#)



Secure and efficient architectures for connectivity [View project](#)

Safe Energy Savings Through Context-Aware Hot Water Demand Prediction

Joshua E. Siegel*

Department of Computer Science and Engineering
Michigan State University
East Lansing, MI 48824
Corresponding author: jsiegel@msu.edu

Aniruddha Das*

Yongbin Sun
Shane Pratt
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139

* These authors contributed equally.

Abstract—Tank-style water heaters provide a critical utility but often waste energy. The trivial solution of idle shutdown encourages the formation of bacteria harmful to humans. We develop a system capable of learning and anticipating demand proactively while remaining sensitive to health concerns by combining a predictive autoregressive network capable of modeling hot water flow demand with a Cognitive Supervisor designed to minimize Legionella formation with minimal energy expenditure. We developed low-cost computing hardware to capture sensor data and run a predictive model and train it on real-world water flow data captured from a single home over eight months. This system has the potential to save energy in home and commercial applications without compromising health, and may be used to augment new and incumbent water heater installations in low to middle-high income countries.

I. MOTIVATION AND OPPORTUNITY

Hot water is a high-priority utility[1] critical to safety and comfort, though overheated or underutilized hot water wastes energy. Water heating consumes 15% of electricity and 25% of natural gas in the U.S.,[2] and in excess of 22% of energy in Canadian households[3], [4]. Low prices drive adoption of inefficient tank-style systems (97% of the U.S. market[5]), with a New Zealand study finding 34% losses for electric and 27% for gas systems. Plumbing design also impacts standby and distribution losses, which may reach 1,200 kWh/home/yr [3], driving economic and environmental costs.

Proactive shutdown during low-use periods reduces energy expenditure and cost, but accelerates the formation of malignant Legionella bacteria. Legionella can harm residents of a single home, with risk amplified within industrial water distribution systems and communal living environments including hospitals or nursing facilities[6].

Excess energy use, insufficient hot water, and tainted plumbing pose significant problems in low to upper-middle-income economies. Storage-type water heaters could be improved with a model-predictive, health-conscious controller capable of anticipating hot water demand or identifying bacterial growth conditions and modulating the supply to match predicted demand or to create a bactericidal environment. Accurate matching could eliminate the need for costly mixing valves, while designing the algorithm to operate on low-cost con-

trollers could unlock access to emerging markets with less-abundant energy and heightened cost sensitivity in addition to augmenting the controllers on the substantial installed base of tank-style heaters in upper middle-income countries.

In this article, we develop a self-learning water heater control algorithm capable of anticipating flow demand and efficiently and proactively heating water while remaining sensitive to Legionella formation. We create a home-specific water consumption model to accurately predict hot water flow, helping ensure water is optimally heated to maximize comfort with a minimum of excess energy expenditure. We further envision a context-aware watchdog sensitive to bacterial growth conditions and pair this system with the demand prediction model to form a “cognitive supervisor,” using Legionella growth patterns and long-timescale predictions to identify where the minimum input energy elevates the water temperature beyond the critical heating point required for curbing or reversing bacterial formation.

We begin with an exploration of prior art in Section II, detail our proposed solution in Section III and describe our data collection system and potential heater controller in Section IV. The control algorithm is proposed in Section VI and results are shown in Section VII.

II. PRIOR ART

Energy may be conserved by reducing water temperature during low use[7] or through electronic demand projection[8]. Another solution is manual heat modulation, though individuals often inaccurately estimate their consumption and continue to over-heat water (cold showers cause more immediate suffering than slow-accruing energy bills).

On-demand tankless heaters provide a hardware solution for energy reduction without compromising comfort, [9] though these systems are costly and poorly-suited to low-income economies. While low-cost, localized heating has been envisioned to provide on-demand hot water for showers in these regions, power requirements limit adoption[10]. In contrast, there is a large addressable base of inexpensive storage water heaters around the world.

Studies have found water consumption to be varied across geographies but repeatable at a local (home) scale, suggesting that personalized energy management could yield savings[11]. Schedule automation saves energy[12], though sudden demand spikes may go unmet due to water's high heat capacity. Theoretical savings suggest a 14.7% energy reduction using a timed system[12]. Heaters may use real-world data for optimization[13] or predictive load scheduling[14], [15]. Utility-level control has also been proposed to reduce peak energy demand,[16] along with centralized heater demand management and direct-load control programs [17], [1]. Predictive modeling has been evaluated[11], with some models using date/time aware filter-based models to forecast demand[17]. These approaches may increase the risk of Legionella and are not recommended for those with compromised health[18].

Heat is a factor in the growth of Legionella bacteria, with temperatures near 25°C accelerating growth [19] and temperatures $> 60^{\circ}\text{C}$ decreasing Legionella formation[6]. Setting a water temperature floor of 60°C imposes a high minimum energy cost, while intratank temperature variability means the outlet temperature may be higher[20], increasing scalding risk. Conventional scheduling struggles to balance energy savings with comfort, adapts poorly to unexpected use, and neglects bacterial development considerations.

Recently, Booyesen, et al. proposed techniques to modulate water heater energy input based on temperature matching, energy matching, and energy matching subject to bacterial growth constraints. Energy Matching with Legionella (EML) prevention matches tank outlet energy with demanded hot water energy by varying water temperature and flow rate. In this approach, water is heated to 60°C for 11 minutes at least once per day, just before the largest predicted outflow event[21]. Results for Energy Matching (EM) show a median energy reduction of 17.8%, whereas the EML with sterilization heating yielded 13.1% savings, both with no increase in perceptible "cold events," or insufficient hot water noticeable to consumers. The EML approach attains minimum energy expenditure in cases where large outflow events are directly correlated to the most significant water energy use, but this may not be optimal when large outflows of hot water are mixed to reduce the temperature (e.g. a tepid bath). A system capable of predicting true, at-the-tap energy demand may further reduce energy consumption by precisely timing the sterilization event, and allow improved efficiency in timing the Legionella sterilization event.

A context-aware, energy demand-based scheduler for existing tank-style water heaters has the potential to save energy without compromising safety or comfort. Such a system would anticipate demand for hot water, and predict far enough into the future to identify whether the demand would intrinsically cause the system to exceed the sterilization temperature for Legionella, and if not, where the smallest additional of extrinsic input energy will cause the system to exceed that limit, assuring the safety of the water stored in the tank. Section III describes our approach to creating such a system.

III. PROPOSED SOLUTION

Individuals are good at identifying water flow events but poorly estimate events' volume and duration [22]. Using real-world sensor data, it may be possible to predict future demand for both water flow rate and desired temperature.

A data-driven model could enable a hot water control system capable of anticipating outflows to efficiently and proactively modulate stored water energy to meet demand. Such a system could replace incumbent data-blind controllers that over- or under-heat water, and with the use of a proportional integral derivative (PID) control, further address inefficiencies inherent in these imprecise "bang-bang" hysteresis controllers[1]. By integrating a predictive algorithm and improved control hardware into existing heating systems, a more advanced controller could imbue incumbent infrastructure with machine intelligence at minimal cost. This approach would improve energy efficiency by matching hot water supply to demand, but would do little to ameliorate the safety concerns resulting from under-heated stagnant water.

Contextual rules could help to address the issue of bacterial growth to assure the safety of heated water. Rules known to the controller could identify conditions where bacteria formation is likely and trigger actions proactively to mitigate and reverse formation of bacteria. For example, the system might identify that water has recently exceeded the sterilization temperature for Legionella bacteria, and that it is safe to cool down to ambient temperatures as scheduled. In another scenario, the system might identify that the system is projected to idle with cold, stagnant water long enough for Legionella to grow, and as a result might explore opportunities for bacterial growth reformation (release of a treatment agent, or the addition of extrinsic heat to cease further formation).

Combining the elements of data-driven demand prediction, proportional heat control, and context-aware safety, we propose a learning, demand-responsive, Internet-connected energy control system for low-cost storage water heaters. Learned models anticipate demand to proactively heat water, as suggested in [23]. Atop this model, a "Cognitive Supervisor"[24], [25], [26] understands the the water heater's purpose (hot water delivery on demand) and its constraints (human susceptibility to and growth factors for Legionella bacteria) in context. This Supervisor is part of a Cognitive Protection System capable of monitoring system states to ensure adequate performance, and is uses similar models to a "Cognitive Firewall" capable of testing commands received over the Internet for benignness prior to execution[24], [25], [26].

The proposed heater control system uses real world data to learn demand models, and projects demand forward in time in order to make schedule modifications minimizing energy consumption while meeting demand requirements and obeying safety rules. This approach utilizes embedded intelligence to improve system efficiency and safety, building upon the efficiency and safety benefits realized by prior pervasive infrastructure computing implementations[13], [27].

Wi-Fi connectivity allows homeowners to view realtime

consumption data over the Internet, as 9% of energy conservation stems from energy awareness, [22], [13], [4] and idle energy use may go unnoticed if not brought to the consumer’s attention. Networking further allows the same controller replacement to be used for remote Internet control, for example, for a utility to manage resource consumption to reduce peak grid load, or for a homeowner to increase water temperature when freezing conditions are expected. The aforementioned Cognitive Firewall may be used to not only monitor locally-issued commands, but also to address resultant security concerns, simulating commands in context to assure their safety prior to execution on physical hardware.

With Internet connectivity, automated water heating may be controlled by models learned from home-specific data, stored online and combined with weather data, student athletic schedules, or coupled with other external sources useful for improving predictor performance.

Our demand-responsive hot water heater builds upon established technologies and consumer desire to create an efficient and safe solution to making water heating demand-predictive. Consumers will appreciate the cost savings and environmental benefit of energy efficient devices, and the concept of adaptive, demand-based home heating has been widely adopted (e.g. the Nest thermostat) with other demand-based appliances in testing[28], [13].

What differentiates this concept from the earlier-described utility- and demand-side proactive heating systems is direct connectivity, adaptive home preference models, and context awareness which improves system energy savings without compromising safety. Further differentiating this from the proposed concept of EML[21] is that the sterilization event may be timed precisely to minimize energy consumption, rather than in advance of the largest outflow event, which may not match the peak energy demand - particularly after mixing.

Section IV follows, describing the hardware solution used to capture model training data, and that may be used in the future to replace existing water heater controllers.

IV. EXPERIMENTAL SETUP

Hot water demand varies with mitigating factors including temporal, climate, regional, and cultural differences[22], but demand is largely predictable within a single home. We therefore developed a proof-of-concept system capturing hot water flow data from, and modeling behavior within, a single home. In this article, we consider flow as a surrogate for hot water energy in part to simplify the experimental design, and in part because energy modeling and matching is well described in [21]. The primary contribution of this article is the combination of the demand predictive model with the context-aware safety elements and Internet connectivity rather than the creation of a highly-precise energy model.

To capture flow data, we built an Internet of Things flow-metering system using a low-cost Raspberry Pi 3B microcomputer which serves as a data logger and web interface. Our sensing system employs inexpensive Hall-effect flow meters to capture consumption data useful for creating demand models.

Though the Pi 3B was used primarily to capture training data for demand modeling, the microcomputer was chosen also for its ability to serve as a replacement hot water controller able to locally operate the predictive model. The Pi 3B offers analog and digital outputs capable of triggering solid-state relays or variable solenoid valves, which can directly modulate electric and gas water heaters as envisioned in Section III. The Pi 3B also outputs SPI, I2C, PWM, and analog signals, which easily interface with the power electronics or control systems found in most other existing heater controllers.

Though the Pi 3B is a low-cost, low-power computer, it features multiple cores and can learn new or adapt existing neural network models locally as a background process. While the system is designed to meet stringent cost and power requirements, it also has sufficient processing to run pre-trained Deep Learning models onboard. The Pi 3B has already been shown to be an effective endpoint device for running deep learning models for embedded intelligence[27].

Local operation allows lower latency for the controller, improving energy efficiency relative to delayed commands received from a remote Cloud or Edge solution. Alternatively, the Pi can share data with a remote server in order to more rapidly learn individual and aggregate home’s models, with the server returning pretrained binaries to each end device.

The data collection and control system is shown in 1- 3. While we incorporated multiple hall-effect flow sensors to collect data for future water consumption studies, the models developed in Section VI consider only the data from the hot water tank’s outflow rate sensor. Multi-outlet instrumentation would be costly[22], so using fewer sensors is better representative of systems suitable for low-income countries. In the U.S., consumers already own single-point connected flow sensors (Moen Flo) that may be leveraged as input to this system, indicating a consumer willingness to install such devices in their homes and potential secondary applications for flow data.

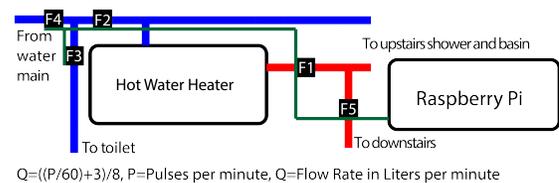


Fig. 1: The system uses mechanical water flow meters sending pulses to a Raspberry Pi 3B for counting. We only use data from F1, the hot water tank outlet, as it supplies the whole house.

Flow data are captured to the Pi as the impeller turns and are recorded to a file once per minute. While low-flowrate events might be missed with infrequent sampling, larger events driving demand for hot water such as bathing, laundry, and cooking, which compromise the largest outflow in most homes, show up clearly in the data. Bathing leads by volume (40% of total water usage) and cooking leads by number of discrete events.[22] The low flow rate reduces controller energy con-

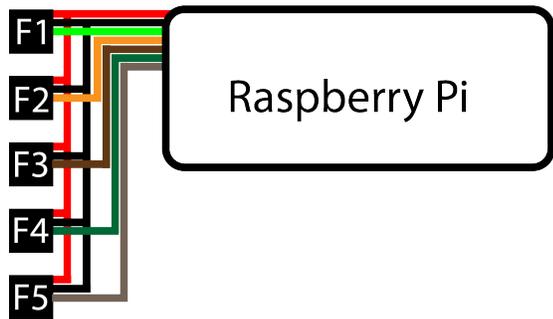


Fig. 2: Each sensor is wired directly to the Raspberry Pi for pulse counting.



Fig. 3: The control box appears on the left-hand side; a representative flow sensor (F1, hot water outlet) is shown on the right hand side.

sumption, data storage, and network bandwidth requirements relative to faster sampling.

As noted at the beginning of this section, this experimental setup considers flow as a surrogate for water temperature, which itself is integrated over time and used as a surrogate for energy. While this is an abstraction, it stands to reason that water must be heated only in advance of an outflow event. This is consistent with the collection methodology in [22]. In future iterations, the addition of data from cold water flow sensors and/or temperature sensors may develop a more accurate, heater-specific relationship between flow, temperature, and energy at all points in the system, similar to the temperature, flow, and system models proposed in [1], [29], [3]. With an improved heater model, it may be possible to build predictive models capable of meeting hot water demands even more efficiently.

The data collection process and sample plots are shown in Section V.

V. SAMPLE DATA

Data were collected from the experimental setup once per minute from January 18, 2018 to August 15, 2018. Once per minute was selected as being an appropriate window size to allow for the capture of small outflow events ($\geq 2m$) while remaining resilience to timing jitter during data capture and keeping storage, computation, and networking requirements reasonable for the Raspberry Pi 3. Due to intermittent device

inaccessibility due to power or network interruptions, there were sporadic data outages as might be present in a real-world system.

To validate the data collection system’s performance, we plotted hot water consumption by day (Figure 4) and by hour (Figure 5).

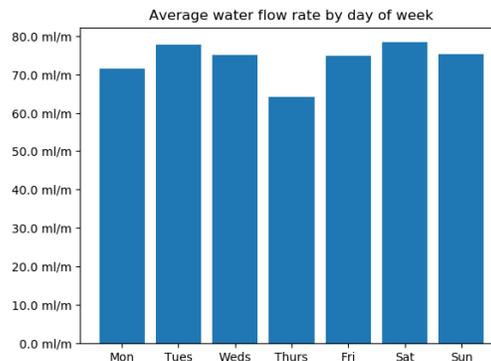


Fig. 4: Usage varies by weekday based on the family’s activities. Water increases on Tuesday, Wednesday, and Saturday due to exercise, whereas water consumption decreases on Thursday, when the family eats dinner out.

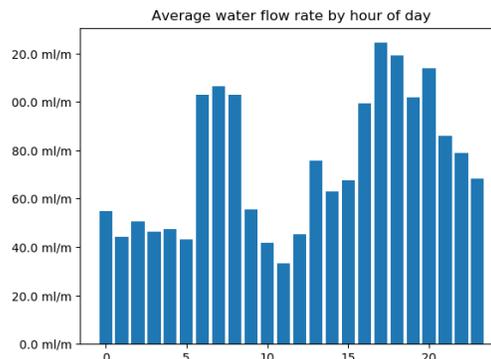


Fig. 5: Usage varies by time of day based on the family’s activities. Overnight hours show reduced water use, with heightened use in the morning and around mealtimes indicating bathing and cooking events.

In the data, we notice weekly and hourly trends that agree with the family’s behavior. For example, water consumption increases on Tuesday, Wednesday, and Saturday due to the family’s exercise-related showering, while consumption drops on Thursday as the family eats dinner away from home.

While additional data would be desirable to develop a commercial product, with seven months of information, there were sufficient data to begin developing a house-specific heating model considering weekly and seasonal variation. Further, the predictive element of this model has been proven effective for water heaters, thermostats, and other utility controls. These

data are therefore sufficient to develop a simple predictive model in order to test the novel contribution of a context-aware supervisory system to limit bacterial growth.

The prediction model and safety-centric Cognitive Supervisor are described in Section VI.

VI. CONTROL ALGORITHM

The control algorithm has two elements: a predictive model to anticipate future outflow events (as a surrogate for energy demand) and a Cognitive Supervisor considering Legionella risk and adapting the tank’s commanded temperature to reduce bacteria formation with the minimum increase in energy consumption. It is the combination of these two elements (anticipatory demand modeling and context-aware, energy-minimizing safety systems) that makes our proposed solution unique. The following subsections describe the design and development of each of these elements.

A. Prediction Model

The prediction model’s purpose is to anticipate water outflow events. The proposed algorithm is structured in the form of a regression model, using data from a single home’s hot water tank outflow history as input to estimate future outflows. While the described model in this section considers the rate and volume of outflow events rather than temperature or energy considerations at the heater, flow is a reasonable surrogate metric to prove model feasibility and to test the incorporation of the “Cognitive Supervisor.” Future variations of this model may incorporate relationships between energy input and thermal properties, with thermodynamic models learned from water heating systems and relating heater energy consumption with flow rate and temperature potentially enabling more precise control and system-wide energy optimization.

1) *Problem Formulation:* To predict the water flow rate, we designed an autoregressive Deep Learning framework that ingests time series water flow data as input and returns a prediction of the expected flow for the upcoming 24 hours.

From previous studies which identify trends in habitual water use, we anticipate that future flow values will be highly correlated to recently-preceding flow values and/or flow values from similar times on previous days. Therefore, a critical first step is to transform the data into a format better-suited to capturing this potential correlation than conventional, less-structured time series. By changing the data representation, we may subsequently develop a more accurate model capable of capturing patterns and predicting future flow values dependent on time of day, day of week, and seasonal effects.

To convert the data to a more robust representation, we take the time series data points and construct an $m \times n$ matrix where m is the number of days considered and n is the number of flow samples per day. Only the most recent data fitting into this space is used and data extending further back in time than $(m \cdot n - 1)$ samples is excluded.

Given the nature of the data and the probable correlation of samples to others in their local neighborhood, a Convolutional Neural Network (CNN) model is an ideal choice for

the predictive model. CNN’s are well suited to learn filters capturing latent correlation among data points related by time (and, in the updated representation, space) than might be extracted from a single-point time-series. Our model draws inspiration from PixelCNN models [30] due to their ability to capture the correlation among neighboring datapoints and the autoregressive nature of our problem.

2) *Training process:* As a first step, the data were split into training, testing, and validation sets to avoid cross-contamination and overfitting. The most recent 20% of data were kept as outsample data for testing. The data were split sequentially, rather than randomly, to capture the inherent time-dependence of the data. Of the remaining 80%, 10% of the data were used for validation, while the remainder were used for training.

The data were magnitude-normalized by dividing the training, testing and validation set by the highest value in the training set (51.5864). Normalization leads to more stable training dynamics and allows backpropagation to arrive at an optimum more easily, resulting in faster training of the model. The mean of the normalized training data were 0.001253, with this low value for flow indicating that water remains stagnant the majority of the time.

Due to data unavailability and sparsity resulting from sensor, network, or computing outages, there were discontinuities within the data i.e. regions in the data set where two consecutive samples were collected at times separated by hours, days, or even weeks. In the sample house data, we found three such discontinuous sections. In analysis, these sections of data are treated as disparate segments in an effort to not confuse the model with data that may be less-well-correlated than might be expected by the model (a likely consequence of considering all data as contiguous regardless of the existence of gaps in the time series).

Unlike heating or cooling which are often active even when homeowners are away, water is often stagnant within a home due to the occupants being away or otherwise not engaging with plumbing. In the case of our sample home, most of the data points in the hot water flow time series were 0. As a result, feeding data sequentially into most predictive models would result in batches comprising solely of zero values as input. Predicting future values from long periods of no-flow, typical models would resort to the null solution of expecting 0.0 or very small values for the entire predicted time series. In the case of our sample home, there were so many sequential 0 values that early attempts at modeling the data would output only minuscule values for predicted flow independent of the input series. While these highly-invariant models may work for thermostats, they are less useful for anticipating water demand, which is by nature “spikier” than other utilities.

To mitigate this challenge, when generating a batch of training data, we select a segment with probability proportional to the length of that segment. Then, a predefined constant determines the percentage of “positive samples” (samples with non-zero water demand at the timestamp to be predicted) from that segment that will be included in the batch sent

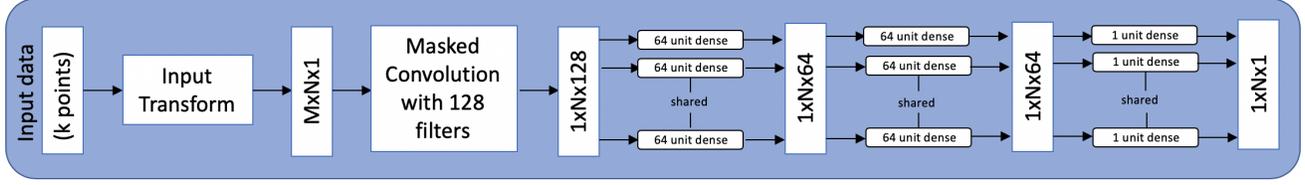


Fig. 6: This figure shows the structure of our developed autoregressive network.

to the model as input. We randomly sample this percentage of non-zero samples from the selected segment and fill the remaining samples required for the batch labels/outputs with zero-valued samples. The corresponding inputs that the model should approximate the function for are created by taking the preceding $\min(p, m \cdot n - 1)$ (where p is the number of values available before the sampled point) values of each of the sampled points and performing the aforementioned input transformation to create the corresponding $m \times n$ matrices for each sample. These transformed $m \times n$ matrices are passed to the model as batch inputs (Figure 6). This constructed batch with the sampled outputs and corresponding created inputs is then fed to the model for training, with the model attempting to minimize the mean-squared error (MSE) between the predicted time series and the ground-truth time series.

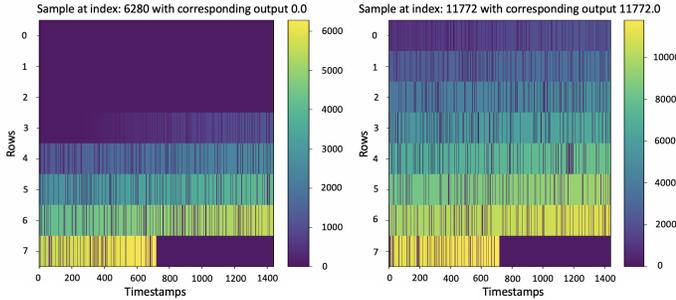


Fig. 7: Sample matrix outputs for a monotonically increasing segment of length 20,000 with interspersed zeros. **Left:** Note in the case of the sample at index 6,280, the data do not fill the matrix and has zero values for the sections it couldn't fill (and beyond the timestamp at which prediction begins). **Right:** In the case of the sample at index 11,772, we see the matrix is filled everywhere until the timestamp beyond which the model makes predictions, due to an abundance of non-zero datapoints relative to the model's desired input.

For a concrete example, assume the training data contain three segments of length 50, 1,000 and 20,000. When we generate a training batch of size 32 from these raw data, we are probabilistically most likely to pick the segment of length 20,000. Assuming the segment of length 20,000 is chosen and the predefined constant is set to 15%, we first sample five points within the segment that have non-zero values and then sample the remaining 27 points from the zero-valued elements in the segment to create a vector 32 elements long. The

combination of these sampled values comprise our training input paired with the segment's known output label. Assuming that the model uses $m = 7$, and one of the points sampled was at index 50 of the segment, we take the previous 50 points and input transform them into the $m \times n$ matrix format. Similarly, if one of the points was sampled at index 15,000 but our model can only look back in time by $7 \cdot 1,440 - 1 = 10,079$ timesteps, we take the most recent 10,079 points and convert these into the $m \times n$ matrix format. These matrices are fed to the model together with their corresponding output labels.

Figure 7 illustrates this subselection process by showing sample matrices of dimensions $m \times n = 8 \times 1,440$ generated from synthetic data.

During training, in the case where all preceding demand values are zero, we construct the training batch labels by randomly sampling from these zero values and feed the samples into the model as outlined earlier. In the sample house data, $\sim 12\%$ of the water demand was non-zero valued and therefore, given at least one day of historic demand data, we never encountered a situation where we only fed zero-valued batches to the model resulting in the prediction of a trivial solution.

Note: The constant determining the percentage of non-zero ground truth values to be included in a batch was empirically determined to perform best as 15% non-zero values, when tested using our sample data set. This value struck a balance between predicting the trivial solution and accurately following the trends in the real data, and was the value used to train the models in all subsequent experimentation.

While the model is training, the optimal heater behavior is to remain in an "always heating" stage to ensure safety and comfort. Once the model is learned, we retain the most recent flow values for at least the last week even if this is in excess of the model's desired input window. Doing so allows the model to be resistant to possible data unavailability, allowing the retained values to help back-fill values in the event that data is missing. While this is all that is required for $m \leq 7$, we will need to store up to $(m - 7) \cdot n - 1$ further values as history to utilize the models complete lookback capabilities. Therefore, we will have to maintain a buffer of size b such that b takes on values as denoted by Equation 1.

$$b = \begin{cases} 7 \cdot n & m \leq 7 \\ m \cdot n - 1 & m > 7 \end{cases} \quad (1)$$

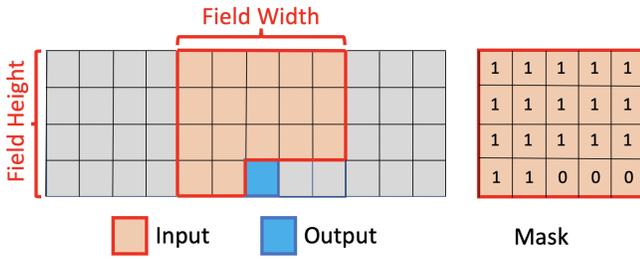


Fig. 8: **Left:** Application of the masked convolution on the input $m \times n$ matrix. **Right:** The mask used to segment valid values of the input patch or convolution kernels.

3) *Predictive Operation:* A convolutional operation was selected as the basis of the predictive function, because convolution is well-suited to data where all points are likely to be a linear combination of the points that precede it, either immediately before or at similar times on previous days. Using convolutional filters allow us to track small windows of time relative to the entire time history, thereby making it easier to capture the correlation between semi-local datapoints and future values. To ensure that the convolution operation does not look at the “future” values, the convolution operation performed was masked (Figure 8) to avoid contamination of the training process with future data.

When obtaining predictions for the desired time, we take an input time series and convert it into the $m \times n$ matrix format. This matrix is then passed through the network to predict the next value. We obtain this predicted value by reading the value returned in the timestamp of interest in the final $1 \times N \times 1$ tensor (as noted in Figure 6). We then iterate through the data, using this predicted value in conjunction with the initial input time series (minus the single oldest data point, to maintain the $m \times n$ of the input data) to make the next prediction in the sequence. This process is continued until we have predicted to the desired temporal distance into the future.

The performance of the model depends strongly on the dimensions of the $m \times n$ matrix and relatedly, the dimensions of the convolutional kernel. The height of the kernel is always equal to m so that we can look at datapoints from all the preceding days. However, the kernel width can vary for the same value of m . The kernel generally acts as a receptive field and the width influences how many samples before and after the timestamp of interest are we looking at (e.g. a kernel width of 61 will look at 30 timestamps before and after the timestamp of interest). This dependence is explored in depth when considering the predictive model’s results in Section VII

B. Cognitive Supervisor

In conjunction with the predictive model, a Cognitive Supervisor considers the context of the water heater with regards to Legionella formation, human safety, and energy demands.

The Cognitive Supervisor is a model-based simulator anticipating demand for hot water over the next 24 hours at every time step and computing the risk of cultivating Legionella.

The Supervisor, described in [24], [25], [26], uses context information (in this case, rules about acceptable levels of Legionella formation and a priori knowledge of growth conditions) to identify and mitigate the risk of bacterial growth. The Supervisor does this by commanding the water tank to heat standing water, even if there is no anticipated demand in excess of the sterilization temperature within the next day. It further integrates with the predictive model’s anticipated future demand and coupled energy input to identify the period in the coming day during which the smallest delta in input energy would cause the water to exceed the critical bactericidal temperature for the duration necessary to assure safety.

The proposed anti-bacterial Supervisor uses a “watchdog” timer with a sliding 24-hour window, setting a safety flag to “true” if the temperature exceeds a predefined bactericidal limit for a known duration, and resetting that flag to “false” after water sits for 24 hours with temperatures below the target. If heat in excess of the Legionella-lethal limit is not expected in the coming day, the watchdog selects an optimal point to increase the temperature exceeding the sterilization limit based upon predicted demand and the heater’s energy model (e.g. choosing the tank temperature minimizing the delta in energy between the anticipated demand and the lowest-energy “safe” state). The model may improve as new information is learned about Legionella growth, for example to incorporate ambient temperature, water supply quality and chemical treatments, pipe materials, or system flow rates known local to the system or captured from remote network resources. Eventually, variational techniques may be used to permute the expected heating schedule in search of the minimum energy difference between that required to meet household needs and that meeting Legionella safety requirements, rather than timing sterilization in advance of the largest outflow event as in[21]. This approach is ideally suited to use cases including typical home or industrial use, or more complex cases such as those with intermittent demand (e.g. vacation homes or offices that close on the weekend).

VII. RESULTS

In this section, we test the envisioned model for various hyperparameters and for the best model, visualize the predicted results and provide quantitative analysis of the model’s performance in predicting hot water outflows.

Subsection VII-A considers the input region’s hyperparameters (width and height), with Subsections VII-A1-VII-C using the identified optimal hyperparameters for patch dimensions to compute their results. Subsections VII-A-VII-C compute anticipated water flow for a period of 24 hours in advance of the present, which is the length of the window required for the Cognitive Supervisor to identify the need for and optimal location of sterilization heating, while Subsection VII-C considers the model’s ability to predict farther into the future, which could be more useful for long-term demand projections, e.g. those used for utilities to schedule power generation or fuel ordering.

A. Ablation Study on Input Region Dimension

The developed predictive model takes as input the dimensions of a local region patch to be examined. The size of this patch determines model performance and is a hyperparameter that must be tuned based on each home’s sample data and the desired applications’ characteristics.

In this subsection, we therefore present multiple design options for the local region patch. Specifically, we consider the the width of input region, which defines the range of correlated time steps across different days, and the height of input region. The input region height determines the length of previous days our model takes into consideration when making future predictions.

We consider different values for both width and height, and then visually compare the results before continuing on to quantitative evaluation.

1) Study on the range of time steps:

To identify the optimal kernel width, the sample data were considered and future flow was predicted forward in time using the training data as input and the results were compared against testing data as ground truth. Figure 9 shows results comparing the ground truth (blue) and predicted (orange) flow rate for various receptive fields, where the kernel width (number of days, m) is permuted. Examining Figure 9, we see that as the kernel width increases, the predicted data appears to better capture the correlation between the datapoints and makes predictions more in-line with the ground truth values. Intuitively, it makes sense that additional time history would improve the model’s predictive performance.

The models with smaller receptive fields generally output a running mean but fail to capture the extent of the variability of the data whereas those models with larger receptive fields output results better tracking daily, hourly, and minute-by-minute variability.

This is likely a result of the nature of the data. As water demand need not occur at the same time everyday, having a larger kernel width allows the model to see more of the data from the past, and make better predictions of the future. The models with smaller kernel widths are more likely to miss the water demand if there is large variability in the timing of the demand whereas those with larger kernel widths will still be able to see and use the past demand.

However, the downside to constantly increasing the kernel width is the large subsequent increase in the number of trainable parameters (model with the kernel width of 15 had $\approx 28,000$ parameters compared to $\approx 136,500$ of the model with the width of 121) which causes the model to become slower to train, more complex to run on the constrained computing environment of the Raspberry Pi 3, and increases the risk of overfitting.

Before examining the performance results quantitatively, note that due to the high heat capacity of water, it is more critical to track water demand on a longer (hourly) time-scale rather than minute-by-minute. This is because water in a tank takes a very long time to heat and cool off. In essence, the water tank stores heat energy sufficiently long that we do not

Receptive Field Width	Test #1	Test #2	Test #3	Test #4	Avg.
15	0.000848	0.000721	0.000771	0.004479	0.001705
31	0.000274	0.000975	0.000623	0.005376	0.001812
61	0.002013	0.001832	0.000987	0.004808	0.002409
91	0.000703	0.002951	0.000902	0.004688	0.002311
121	0.000409	0.001010	0.000955	0.004397	0.001692

TABLE I: This table shows the numerical scores obtained by the models with varying kernel width on each of the tests. The widest field (most days of prior history) performs the best on average, but the smallest receptive field does perform well.

need to consider events consuming hot water for a period on the order of minutes. Instead, we must consider longer-duration events and the related volume (a Riemann sum or integral of flow over time). The length of time considered is a function of tank size, since a bigger tank takes longer to heat up but also longer to draw down. As a result, even though we minimized MSE in training, we will use a different metric to evaluate model performance at runtime.

This metric is computed by looking at the points within a window and computing the MSE inside that window after best aligning the predictions to the ground truth values. This acts as a form of similarity score between the shapes and magnitudes of the two sets of values under the window.

The reported score is the value that results from summing over all the similarity scores between the predictions and the ground truth at each possible window position as the window slides across the data. We consider each datapoint within a window only once (i.e. the sliding window strides at intervals equal to its length), and divide this sum by the number of days in the future being predicting. This acts as a better measure of performance for our purposes than simply using the MSE across the entire length of the predictions and the ground truth. Similar to MSE, smaller values indicate improved performance.

Using this metric as shown in Table I, we see that the model with kernel width of 121 had the best performance overall. Notably, the model attained good results with much smaller smaller receptive fields (15) and comparatively poor performance with kernel width of 91, despite results that qualitatively appear reasonable.

The model performs well using the similarity metric with a smaller receptive field due to the contents of the training home’s input data. As mentioned in Section VI-A2, most of the flow data are 0, and from Figure 9, we see that the outputs from most learned models tends to be conservative. These models report small values resulting in small MSE for most future prediction regions. This behavior illustrates precisely why an improved similarity metric was necessary.

However, the model with the kernel width of 91 likely underperformed quantitatively due in part to some outputs being large. In Figure 9, the y-limits were kept constant to maintain uniformity and aid easier comparative analysis. Points that went beyond the limits were truncated visually, but these points were still used in calculation of the error

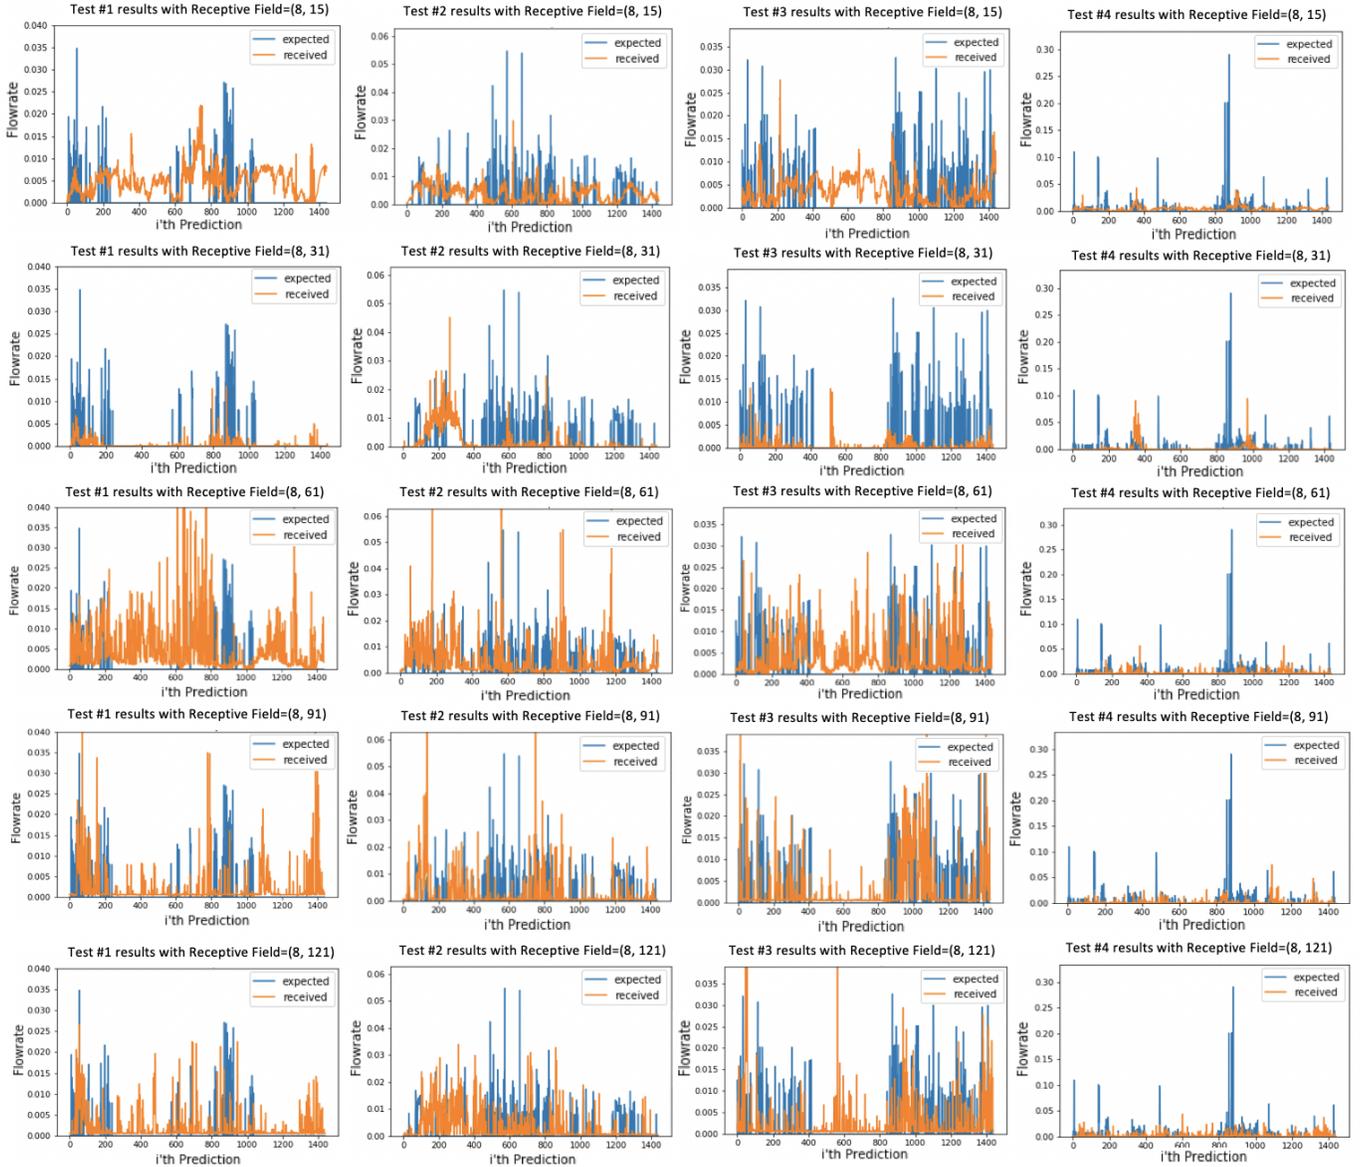


Fig. 9: Test results on varying kernel width. The maximum number of days (m) and therefore, the height of the convolutional kernel was kept constant but the width of the kernel was varied to change the time window under consideration. Each row corresponds to a different model configuration and each column is the performance of the corresponding model on a particular test. We see that as we increase the kernel width, the model better learns to capture the trends in the data and begins to make predictions that more accurately match the ground truth values.

metric which resulted in the inconsistencies (Test #2) in the quantitative results.

2) *Sensitivity to inclusion of previous days*: The results in this subsection consider the optimal input patch dimensions identified in Subsection VII-A1. We identified the optimal kernel width as being 121, and will keep the width constant as we study the effect of varying the number of days included in the reference time history.

Observing Figure 10, we see that for smaller kernel heights (*one* and *two*), the graphical results poorly capture the data trend as the model lacks significant contextual information.

However, once we start considering reference data \geq *one* week prior, the model begins to successfully learn the trends in the data and make predictions that better-mirror the ground truth values.

While looking at data more than one week back does improve predictor performance, it doesn't improve results significantly compared to the increasing the number of learnable parameters (Models with kernel height of 1, 2, 8, 22 and 29 have \approx 28,000, 43,500, 136,500, 353,345 and 461,700 parameters respectively). Increasing the learnable parameter

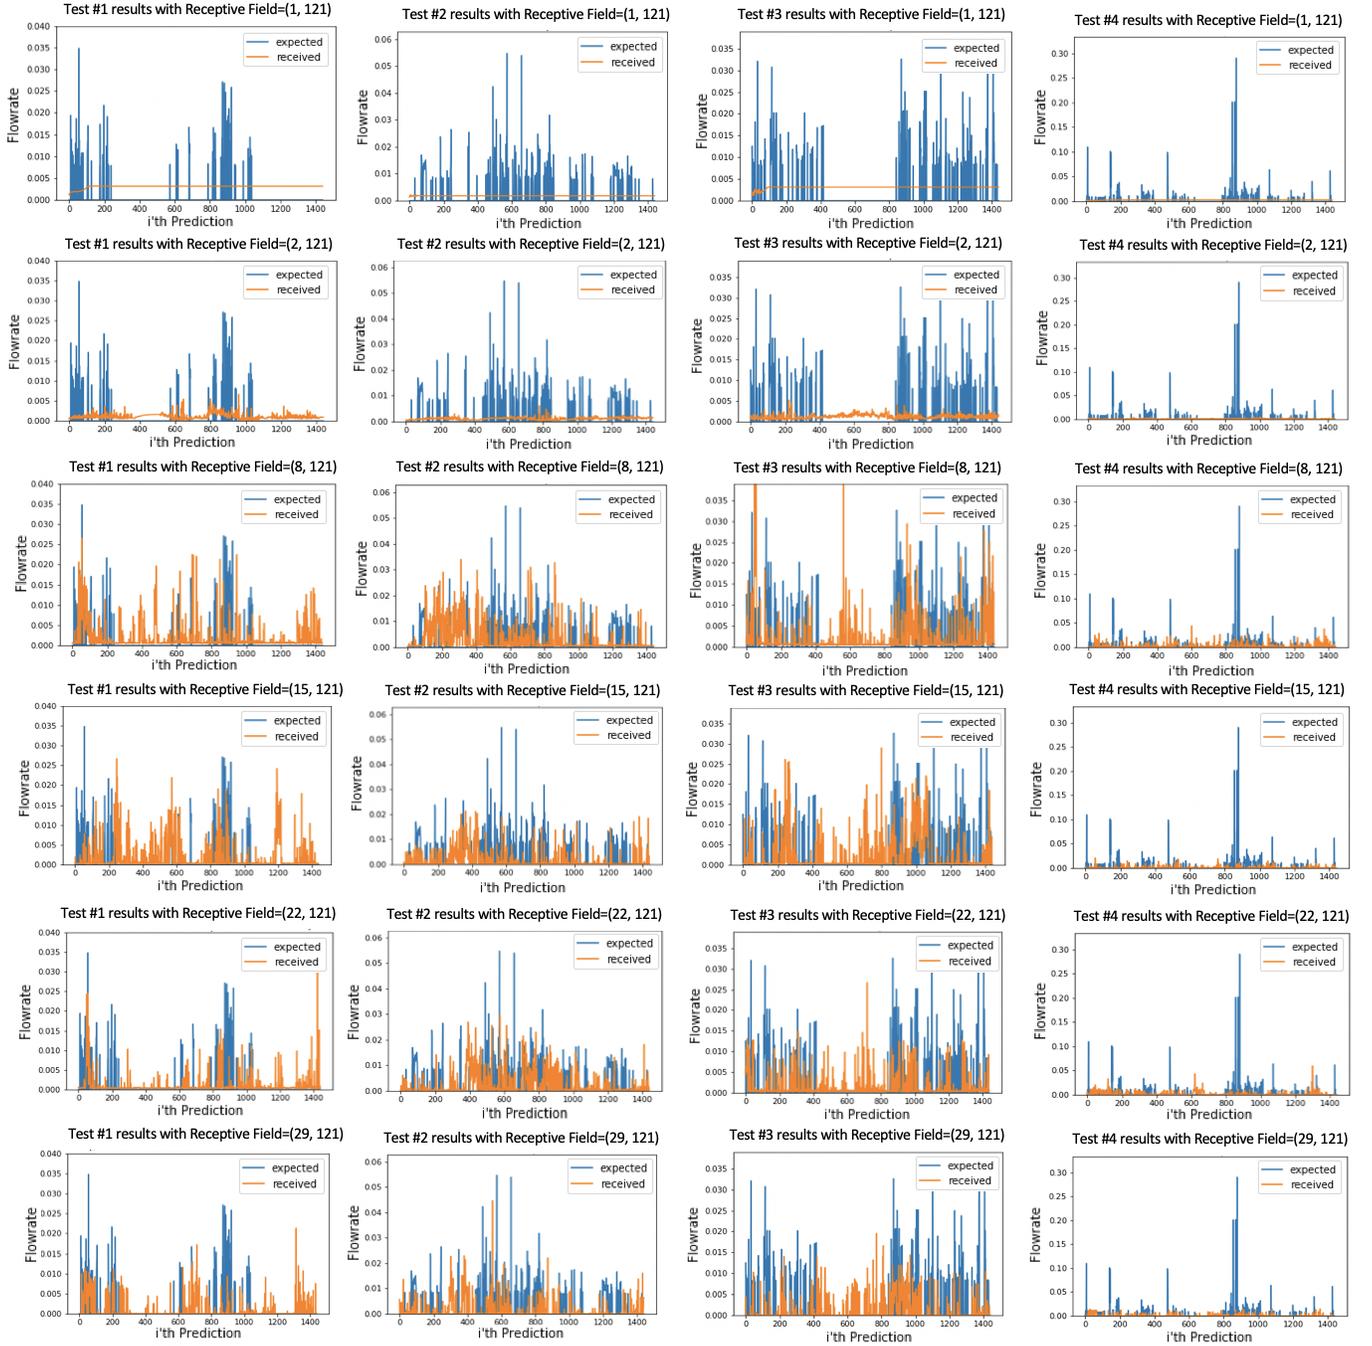


Fig. 10: Test results on varying kernel height. The maximum number of days (m) and therefore, the height of the convolutional kernel was varied while the width of the kernel was kept constant and equal to the best value from Section VII-A1 (kernel width of 121) to better study the effect of changing the number of previous days that are taken under consideration. Similar to Figure 9, each row corresponds to a different model configuration and each column is the performance of the corresponding model on a particular test. We see that as we look beyond a week, the models begin to generate more accurate predictions which is likely a result of them having access to the target days demand in earlier weeks.

count results in an increased tendency to overfit to the data and a related increase in time to train the model (the model with the kernel height of 29 took $\approx 2.5\times$ the amount of time to train compared with the model having a height of 8).

Looking at the results in Table II, we notice that as the kernel heights increased beyond one week, the models begin to outperform those models with the smaller receptive fields both qualitatively and quantitatively. The best results come from

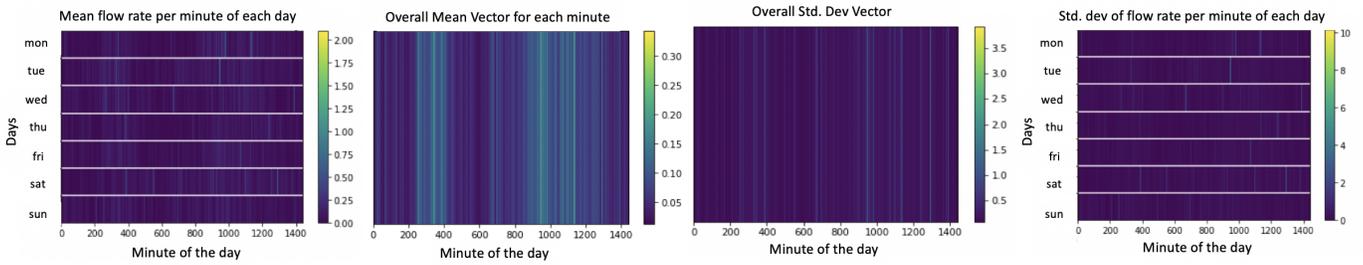


Fig. 11: This figure shows heatmap vectors representing the mean and standard deviation of unnormalized flow rates for every given day and minute. The color bars to the right of the images indicate the differing scales used in the each image. The values shown correspond with those indicated in Figure 5 and Figure 4. The algorithm predicts the next row of the constructed $m \times n$ matrix, and this representation provides a reasonable graphical understanding of what the next predicted row may resemble. These vectors may be used to construct another performance metric for the model by comparing the model’s predictions with the mean water demands for the corresponding days and minutes. Examining the similarity between the mean vector of the data and that of the model predictions’ mean vectors and reporting the percentage of model predictions that lie within a certain number of standard deviations gives us another qualitative and quantitative way of measuring the models performance.

Receptive Field Height	Test #1	Test #2	Test #3	Test #4	Avg.
1	0.000415	0.000687	0.000634	0.004721	0.001614
2	0.000278	0.000677	0.000612	0.004789	0.001589
8	0.000408	0.001009	0.000955	0.0043965	0.001692
15	0.000524	0.000782	0.000777	0.004571	0.001663
22	0.000317	0.000736	0.000597	0.004686	0.001584
29	0.000321	0.000740	0.000627	0.004606	0.001573

TABLE II: Table showing the numerical scores obtained by the models with varying kernel width on each of the tests. The tallest receptive field performs the best, but is the most computationally-intensive, which could be a challenge when training or running within a constrained computing environment.

the model that looked back the furthest (the model with the kernel height of 29, considering just over four weeks of prior data, which balances data volume, predictive performance, and computational complexity). Here again, the the smallest kernel heights (*one* and *two*) again perform impressively. This behavior occurs for the same reasons identified in Section VII-A1.

B. Evaluation on the Whole Day Prediction

In this subsection, we evaluate the quality of the data projected one day into the future, the length of prediction necessary to support the Cognitive Supervisor’s Legionella watchdog.

We begin by taking the water demand data we have and finding the mean water demand for each day and minute of the week. We then use these computed means to find the corresponding standard deviations of the water demand on each of the days and minutes. The resultant values are expressed as vectors in Figure 11

The resulting vectors help determine whether the model is performing well by considering how closely the model’s predictions match the output mean vectors (Figure 11). Additionally, it will also give us a sense of what the models predictions will look like for unseen data.

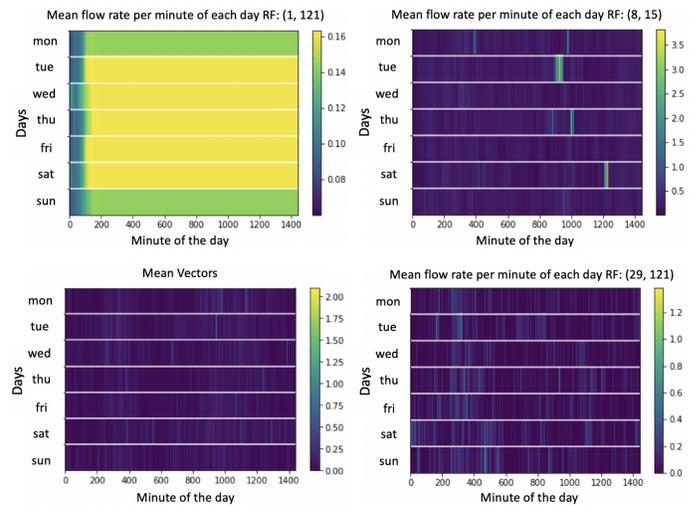


Fig. 12: Output mean vectors for model predictions. **Top-left:** Mean vectors for model with kernel dimensions of 1×121 . **Top-right:** Mean vectors for model with kernel dimensions of 8×15 . **Bottom-left:** Mean vectors for the data. Same as the one shown in Figure 11. Shown here for ease of comparison. **Bottom-right:** Mean vectors for model with kernel dimensions of 29×121 .

For each set of predicted outputs, we generate a corresponding set of mean vectors to help visualize how the model makes its predictions. Additionally, we compute whether each prediction that the model made was within two standard deviations of the mean demand on that day and time, and return the percentage of such inliers.

The model with the kernel height, width of 29 and 121 respectively had the highest percentage of inliers at 93.32%. In addition to its generated mean vectors we also show the generated mean vectors for models with kernel dimensions of 1×121 and 8×15 as examples of poorer predictions

(Figure 12).

Examining Figure 12, we see that the predictions generated by the model with an input kernel of dimension 29×121 most closely resemble the ground truth mean vectors both in range and configuration. We note that the predictions generated by the model with kernel dimensions of 1×121 poorly track the ground truth vectors' behavior, often maintaining the same conservatively small value throughout.

For the model with kernel dimensions of 8×15 , the generated mean prediction vectors appear more reasonable, however, its percentage of inliers is a low 74.62%. Understanding why requires further inspection of the image. Note that the predictions are consistently higher than those expected in the mean vectors, with the color bar indicating a non-zero minimum value and a maximum value nearly twice the mean vector. These results also align with the models corresponding results shown in Figure 9.

C. Capability for Long-Term Future Data Prediction

To enable long-term scheduling and demand prediction useful for utility-scale optimization, we are interested in determining how far ahead in time the proposed model can predict while remaining useful. In this subsection, we predict data for multiple days into the future and evaluate their performance.

To obtain these long-term predictions, we use the best model discovered in Section VII-A, i.e. the model with the receptive field height and width of 29 and 121, respectively. The data that we will be conditioning on to make predictions is the same as the data that was used for Test #3 in Figure 9 and Figure 10, but the prediction length now extends up to two weeks forward from the simulated “present” date. The results of this experimentation are shown in Figure 13

In Figure 13, we see that the model behaves as expected and diverges over time. The graph showing the progression of the error metric mentioned in Section VII-A1 over time. We note an upward trend, indicating that the model’s predictions become progressively less accurate as we predict further into the future.

This behavior is expected, and is common to many autoregressive models. Since we are making predictions for future values based on all the past data (in this case, up to $29 \cdot 1440$ datapoints), as we predict further into the future, we are increasingly basing our predicted values off of other predicted datapoints. As the quality of predicted data is inherently going to be less than that of the observed data, the slight differences tend to add up and the error accumulates over time.

We can see this reflected qualitatively in Figure 13, where the error metric is small for several days with a small, pseudo-linear increase at the outset. The sudden increase in the error metric on the 8th day is a result of multiple spikes in water demand that the model was unable to predict, likely as a result of the size of the input kernel which must make decisions based on ≥ 7 days of predicted data concatenated with the original data. As the time projected forward increases, the relative ratio of known-good data to predicted and uncertain data decreases, leading to poorer performance. By the 14th

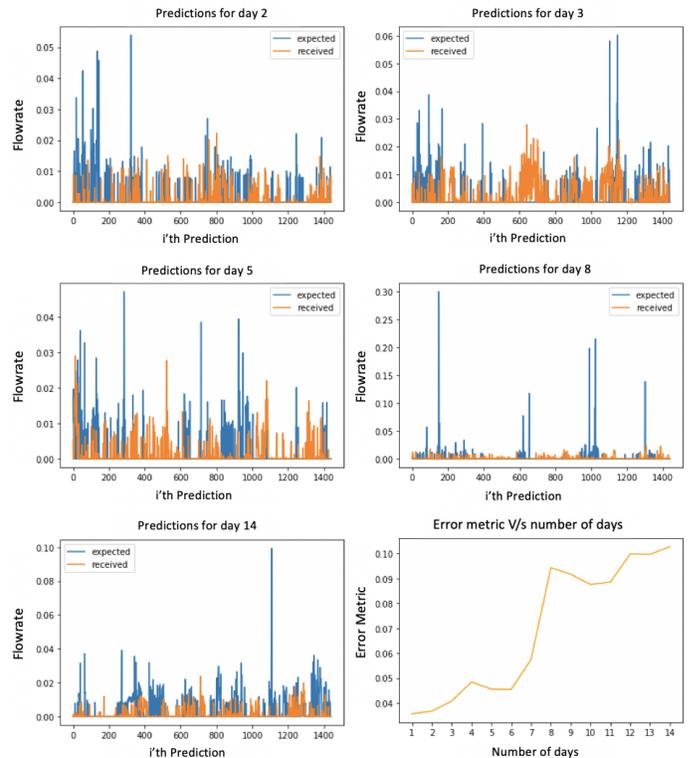


Fig. 13: Obtained results when the model is used to predict further than a single day. **Top-left:** 2 days prediction, **Top-right:** 3 days prediction. **Center-left:** 5 days prediction. **Center-right:** 8 days prediction. **Bottom-left:** 14 days prediction. **Bottom-right:** Graph showing the progression of the error metric as we predict further into the future.

day, the model misses significant periods of demand and/or underestimates the degree of demand at that time.

From this result, we find strong initial performance and a plateau allowing for reasonable flow (and by association heat energy) prediction for the first projected week after the current time. This significant duration suggests that it is feasible for a heater to proactively vary the upcoming temperature profile to stave off Legionella formation with a minimal change in input energy, thereby saving natural resources and money without exposing humans to risk unnecessarily. The performance up to one week in advance without significant divergence also indicates that it may be possible for utilities to predict future demand in order to better match supply with demand.

VIII. CONCLUSION AND FUTURE WORK

We successfully demonstrated a predictive model for anticipating water outflow events based on historic data within a single home and proposed a Cognitive Supervisor capable of using these predicted values to stave off or reverse the formation of malignant Legionella bacteria with a minimal increase in tank heating energy. We also recommended the use of a low-cost connected hot water tank heater controller based on a Raspberry Pi microcomputer to instrument and

modulate the power to incumbent water heating systems, thereby providing a cost-effective solution for increasing the energy efficiency of hot water delivery without compromising comfort or safety.

This work demonstrates the potential for predictive models to anticipate water heating demand in order to comfortably and safely reduce energy expenditure. The described control algorithm may be used to optimize energy consumption for low-cost systems without mixing valves, subject to health constraints imposed by Legionella bacteria. Extension from flow-based models to energy-driven models (considering tank heating dynamics and temperature/flow relationships) will lend this predictive model further utility in emerging and developed markets.

The impact of safety-conscious energy demand modeling is significant. While savings will vary based on the individual heater, the learned model, and use case, our solution’s savings should fall between the savings for the best-available technologies today, Booyesen’s EM and EML schemes (17.8% and 13.1%, respectively)[21].

In cases where water use is predictable and there are few idle periods, Booyesen’s EM scheme’s 17.8% median energy savings should be attainable. In cases where water utilization is repeatable and where the highest outflow volume corresponds to the highest energy event, but where stored water may stagnate for periods, our proposed model and Booyesen’s EML should perform similarly (13.1%).

For use cases where demand is intermittent, conventional predictive modeling tends to be less effective and long idle periods require additional tank sterilization energy. This is true, for example, in vacation homes or offices that shut down for > 24 hours on the weekend. It is in these scenarios, or situations in which the largest outflow event of the day is not necessary the most energetic event, that our proposed Cognitive approach will yield savings falling between the EM and EML models. This scenario may occur where there exist high-energy, limited-outflow events, e.g. due to hot and cold water mixing for bathing. Our ability to project demand well into the future and to monitor the expected energetics of the tank output will help us minimize the delta in energy required to exceed the Legionella sterilization temperature, balancing input energy and water safety.

We can compute rough, first-order economic, energetic, and environmental savings enabled by a proactive water heating model responsive to bacterial growth. To compute potential savings within the United States, we first estimate the number of buildings by type, including 5.8M commercial buildings[31], of which $\approx 4.6M$ [32] have one or more days where there is low- to no- water demand. We further estimate there to be 128M households [33] and 9M vacation homes [34]. Lastly, we approximate 8.143M households in poverty [35] that cannot use hot water or maintain low temperatures. From these data, we estimate the percentage of regularly-occupied homes and offices and conclude that $\approx 85\%$ of the buildings are regularly occupied. If the regularly-occupied buildings save 17.8% using our model (assuming similar per-

formance to the EM model), and the under-utilized buildings save 14.4% (assuming our Cognitive Supervisor approach can gain 10% efficiency over the EML model by more-optimally timing the sterilization event), we find a 17.3% potential savings relative to business as usual.

The average occupied household uses 64 gallons of hot water per day [36] and the average office utilizes 112 gallons of hot water per day (calculated using the average daily hot water consumption of a person in an office and the average number of people per office from[37], [31], [38]).

Using Energy.gov’s “Energy Cost Calculator for Electric and Gas Water Heaters” [39], we see that a representative home and office using heaters with Federal Energy Management Program (FEMP) recommended performance levels have annual hot water energy costs of 4,750kWh (428\$ at 0.09\$ per kWh) and 8,313kWh (748\$ at 0.09\$ per kWh) respectively. With energy savings of 17.3% we save 821.8kWh or 73.96\$ per representative house and 1438.15kWh or 129.43\$ per representative office annually. In addition to the energy expenditures, our carbon footprint also decreases by 1.56lb per kilowatt-hour saved[40].

In the U.S. alone, homeowners would save over \$10B annually in water heating costs, and save almost 80MMT of CO_2 from being released into the atmosphere. For commercial buildings, the savings are over \$750M.

At a global scale, the savings and impact of proactive water heating adds up. The use of predictive models will have the most impact at scale, where Cloud-aggregated data may be used to control aspects of energy management systems[41]. Connectivity amplifies these potential savings. For example, when water must be heated in excess of user tap demand to prevent bacterial growth, other connected devices (dish washers or washing machines) may be scheduled to run, taking advantage of energy that would otherwise be lost to cooling.

We know that mean performance is more critical to predict than high-frequency demand spikes and troughs, so future work will consider the optimal balance of tracking highly-variable data at different timescales (minute-by-minute or hour-by-hour) with computational complexity, as the high latent heat of water means that water takes a long time to heat and cool. Capturing data from multiple homes and over a longer period will allow us to explore repeatability of model performance across homes, and the impact of seasonal effects on model performance. With these data borne out, we may continue to consider the implications of utility-scale demand prediction and Internet-enabled water heater control using Cognitive Supervisor enabled safety models.

Additional work will also consider controller generalizability and model transferability (using a common pre-trained model and adapting the model’s output layers for specific homes) as well as architectures for connectivity allowing models to be trained and shared remotely but used locally. Federated learning may be a reasonable technique to train networks using distributed compute hardware and sensor samples.

The techniques described in this manuscript have the potential to save energy not only in hot water heating, but for

other utilities and appliances. The combination of accurate predictive demand modeling and context-aware systems is a unique enabler of safe and efficient infrastructure with the potential to save energy and reduce disease globally.

REFERENCES

- [1] M. Pipattanasomporn, M. Kuzlu, and S. Rahman, "An algorithm for intelligent home energy management and demand response analysis," *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 2166–2173, 2012.
- [2] T. P. Wenzel, J. Koomey, G. Rosenquist, M. Sanchez, and J. Hanford, *Energy data sourcebook for the US residential sector*. University of California. Lawrence Berkeley Laboratory, 1997.
- [3] C. Aguilar, D. White, and D. L. Ryan, "Domestic water heating and water heater energy consumption in Canada," *Canadian Building Energy End-Use Data and Analysis Centre*, 2005.
- [4] T. Ueno, F. Sano, O. Saeki, and K. Tsuji, "Effectiveness of an energy-consumption information system on energy savings in residential houses based on monitored data," *Applied Energy*, vol. 83, no. 2, pp. 166–183, 2006.
- [5] Energy Star, "Energy star water heater market profile and program design guide," 2009.
- [6] P. Borella, M. T. Montagna, V. Romano-Spica, S. Stampi, G. Stancanelli, M. Triassi, R. Neglia, I. Marchesi, G. Fantuzzi, D. Tatò *et al.*, "Legionella infection risk from domestic hot water," *Emerging infectious diseases*, vol. 10, no. 3, p. 457, 2004.
- [7] K. W. Scott, "Water heater control system," Sep. 4 1979, uS Patent 4,166,944.
- [8] D. R. Stettin and F. W. Sterber, "Heater control device and method to save energy," Sep. 25 2001, uS Patent 6,293,471.
- [9] M. A. Hoeschele and D. A. Springer, "Field and laboratory testing of gas tankless water heater performance," *ASHRAE Transactions*, vol. 114, no. 2, pp. 453–461, 2008. [Online]. Available: <https://search.ebscohost.com/login.aspx?direct=true&db=a9h&AN=43729952&site=eds-live&scope=site>
- [10] R. T. Prado and O. M. Gonçalves, "Water heating through electric shower and energy demand," *Energy and Buildings*, vol. 29, no. 1, pp. 77–82, 1998.
- [11] M. Roux and M. J. Booyens, "Use of smart grid technology to compare regions and days of the week in household water heating," in *2017 International Conference on the Domestic Use of Energy (DUE)*, Apr. 2017, pp. 276–283.
- [12] A. Marthinus J. ; Engelbrecht, J. A. A. ; Molinaro, "Proof of concept : large-scale monitor and control of household water heating in near real-time," in *International Conference on Applied Energy*, 2013.
- [13] J. López-de Armentia, D. Casado-Mansilla, S. López-Pérez, and D. López-de Ipiña, "Reducing energy waste through eco-aware everyday things," *Mobile Information Systems*, vol. 10, no. 1, pp. 79–103, 2014.
- [14] P. Du and N. Lu, "Appliance commitment for household load scheduling," *IEEE Transactions on Smart Grid*, vol. 2, no. 2, pp. 411–419, Jun. 2011.
- [15] A. Sepulveda, L. Paull, W. G. Morsi, H. Li, C. P. Diduch, and L. Chang, "A novel demand side management program using water heaters and particle swarm optimization," in *2010 IEEE Electrical Power Energy Conference*, Aug. 2010, pp. 1–5.
- [16] M. W. Gustafson, J. S. Baylor, and G. Epstein, "Direct water heater load control-estimating program effectiveness using an engineering model," *IEEE Transactions on Power Systems*, vol. 8, no. 1, pp. 137–143, Feb. 1993.
- [17] R. Bischke and R. Sella, "Design and controlled use of water heater load management," *IEEE Transactions on Power Apparatus and systems*, no. 6, pp. 1290–1293, 1985.
- [18] Energy.gov, "Savings project: Lower water heating temperature," <https://www.energy.gov/energysaver/services/do-it-yourself-energy-savings-projects/savings-project-lower-water-heating>, accessed: 2018-08-31.
- [19] D. Van der Kooij, H. R. Veenendaal, and W. J. Scheffer, "Biofilm formation and multiplication of Legionella in a model warm water system with pipes of copper, stainless steel and cross-linked polyethylene," *Water research*, vol. 39, no. 13, pp. 2789–2798, 2005.
- [20] B. Lévesque, M. Lavoie, and J. Joly, "Residential water heater temperature: 49 or 60 degrees Celsius?" *Canadian Journal of Infectious Diseases and Medical Microbiology*, vol. 15, no. 1, pp. 11–12, 2004.
- [21] M. Booyens, J. Engelbrecht, M. Ritchie, M. Apperley, and A. Cloete, "How much energy can optimal control of domestic water heating save?" *Energy for Sustainable Development*, vol. 51, pp. 73 – 85, 2019. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0973082618306562>
- [22] W. Kempton, "Residential hot water: a behaviorally-driven system," *Energy*, vol. 13, no. 1, pp. 107–114, 1988.
- [23] M. C. Mozer, "The neural network house: An environment that adapts to its inhabitants," in *Proc. AAAI Spring Symp. Intelligent Environments*, vol. 58, 1998.
- [24] J. E. Siegel, S. Kumar, and S. E. Sarma, "The future internet of things: Secure, efficient, and model-based," *IEEE Internet of Things Journal*, vol. 5, no. 4, pp. 2386–2398, Aug. 2018.
- [25] J. Siegel and S. Sarma, "A cognitive protection system for the internet of things," *IEEE Security Privacy*, vol. 17, no. 3, pp. 40–48, May 2019.
- [26] J. Siegel, "Cognitive protection systems for the internet of things," *Homeland Defense and Security Information Analysis Center Journal*, vol. 5, no. 4, pp. 16–20, 2018.
- [27] J. E. Siegel, S. Pratt, Y. Sun, and S. E. Sarma, "Real-time Deep Neural Networks for internet-enabled arc-fault detection," *Engineering Applications of Artificial Intelligence*, vol. 74, pp. 35–42, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0952197618301246>
- [28] D. Ventura, D. Casado-Mansilla, J. López-de Armentia, P. Garaizar, D. López-de Ipiña, and V. Catania, "ARIIMA: a real iot implementation of a machine-learning architecture for reducing energy consumption," in *International Conference on Ubiquitous Computing and Ambient Intelligence*. Springer, 2014, pp. 444–451.
- [29] S. Shao, M. Pipattanasomporn, and S. Rahman, "Development of physical-based demand response-enabled residential load models," *IEEE Transactions on power systems*, vol. 28, no. 2, pp. 607–614, 2013.
- [30] A. Van den Oord, N. Kalchbrenner, L. Espeholt, O. Vinyals, A. Graves *et al.*, "Conditional image generation with pixelcnn decoders," in *Advances in neural information processing systems*, 2016, pp. 4790–4798.
- [31] Center for Sustainable Systems, University of Michigan., "Commercial buildings factsheet," <http://css.umich.edu/factsheets/commercial-buildings-factsheet>, accessed: 2019-07-15.
- [32] Eia.gov, "Commercial buildings energy consumption survey (cbecs)," <https://www.eia.gov/consumption/commercial/data/2012/bc/cfm/b6.php>, accessed: 2019-07-15.
- [33] Statista.com, "Number of households in the u.s. from 1960 to 2018 (in millions)," <https://www.statista.com/statistics/183635/number-of-households-in-the-us/>, accessed: 2019-07-15.
- [34] Hostfully.com, "The state of the u.s. vacation rental industry in 2018," <https://www.hostfully.com/blog/state-u-s-vacation-rental-industry-2018/>, accessed: 2019-07-15.
- [35] Center for Poverty Research, University of California, Davis., "What is 'deep poverty'?" <https://poverty.ucdavis.edu/faq/what-deep-poverty>, accessed: 2019-07-15.
- [36] Energy.gov, "New infographic and projects to keep your energy bills out of hot water," <https://www.energy.gov/articles/new-infographic-and-projects-keep-your-energy-bills-out-hot-water>, accessed: 2019-07-15.
- [37] Engineeringtoolbox.com, "Hot water consumption per occupant," https://www.engineeringtoolbox.com/hot-water-consumption-person-d_91.html, accessed: 2019-07-15.
- [38] Bureau of Labor Statistics, "Employment by major industry sector," <https://www.bls.gov/emp/tables/employment-by-major-industry-sector.htm>, accessed: 2019-07-15.
- [39] Energy.gov, "Energy cost calculator for electric and gas water heaters," <https://www.energy.gov/eere/femp/energy-cost-calculator-electric-and-gas-water-heaters-0>, accessed: 2019-07-15.
- [40] Epa.gov, "Greenhouse gases equivalencies calculator - calculations and references," <https://www.epa.gov/energy/greenhouse-gases-equivalencies-calculator-calculations-and-references>, accessed: 2019-07-15.
- [41] R. Evins and N. David, "Using simple predictive models to improve control of complex building systems," in *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments*. ACM, 2017.