Predicting Heating Demand and Sizing a Stratified Thermal Storage Tank using Deep Learning Algorithms

Aowabin Rahman[a], Amanda D. Smith[a]

[a]Site-Specific Energy Systems Laboratory, Department of Mechanical Engineering, University of Utah, Salt Lake City, Utah, USA, 84112

Abstract

This paper evaluates the performance of deep recurrent neural networks in predicting heating demand for a commercial building over a medium-to-long term time horizon ($\geq$ 1 week), and proposes a modeling framework to demonstrate how these longer-term predictions can be used to aid design of a stratified thermal storage tank. The building sector contributes significantly to primary energy consumption in the U.S, and as such, there is a need to predict heating demand in buildings over longer time horizons, and to develop methods that can facilitate installation, planning and management of distributed generation and thermal storage to meet these heating demands. Key objectives of this paper are: (a) Investigate how a deep recurrent neural network model performs in predicting heating demand in campus buildings at University of Utah over multiple weeks, and (b) Develop an optimization framework that which can provide definitive guidelines on sizing a stratified thermal storage tank without requiring high performance computing resources. The results showed that the predictions by the deep RNN are comparatively more accurate than those by a 3-layer MLP, and that these deep RNN predictions can adequately serve as proxy for future demand while considering sizing in the design of a complementary stratified thermal storage tank.

Keywords: Building Energy Modeling, Machine Learning, Recurrent Neural Networks, Deep Learning, Heating Load Prediction, Thermal Energy Storage

---

NOMENCLATURE

CHP	Combined Heating and Power
GP	Gaussian Processes
LSTM	Long short term memory
ML	Machine Learning
MLP	Multi-layered perceptron
NN	Neural network
PGU	Power Generation Unit
PI	Probability of Improvement
RMS	Root Mean Squared

Email address: amanda.d.smith@utah.edu (Amanda D. Smith)

RNN   Recurrent Neural Network

SMBO Sequential Model-Based Optimization

TS   Thermal Storage

α Learning rate in gradient descent algorithm

σ Sigmoid function serving as a gating function

σ_g Standard deviation associated with GP prediction

$\circ$ Element-wise vector multiplier.

γ Fraction of heat recovered in heat recovery unit that is routed to TS.

γ^{opt} Optimal value of γ at a given timestep t

γ^{GP} value of γ at a given timestep t predicted by GP

η_{PGU} Electric efficiency of PGU.

η_{rec} Thermal efficiency of the heat recovery unit.

ζ Factor accounting for energy losses outside of heat recovery unit.

ρ Density of water

a Acquisition function for Bayesian Optimization

A_c Cross-sectional area of each node (m^2)

c_p Specific Heat Capacity of water (J/KgK).

$\mathbf{c_t}$ Transient 'memory' value in LSTM function

d_i Inner diameter of heat exchanger (mm).

d_o Outer diameter of heat exchanger (mm).

e Mean squared error in predicting electricity consumption

E_{gen} Electricity provided to the building

f_t Targets in training data for GP

f_e Predictions made by GP

$\mathbf{g}$ Input activation function in LSTM

$\mathbf{h_t}$ Output of LSTM function at given timestep t

h_j^m Value of hidden node in a neural network in node j, layer m.

h_i Inner heat transfer coefficient (W/m^{2}K).

h_o Outer heat transfer coefficient (W/m^{2}K).

H Height of storage tank (m).

$\mathbf{i}$ Input gate in LSTM

k Thermal conductivity of water (W/mK).

k_{mat} Thermal conductivity of heat exchanger material (W/mK).

M_I Minimum size of training data corresponding to case I $Q_{rec} \geq Q_d$ for GP to predict $\mathbf{x^{opt}}$

M_{II} Minimum size of training data corresponding to case II $Q_{rec} < Q_d$ for GP to predict $\mathbf{x^{opt}}$

N Total number of nodes

$\mathbf{o}$ Output gate in LSTM

Q_d Heating demand in building

Q_{rec}	Heat recovered in heat recovery unit
Q_{st}	Heat stored in thermal storage tank
$Q_{T,bldg}$	Total heat delivered to the building
$Q_{TS,out}$	Heat delivered by thermal storage to building
s	Parameter to describe discrepancy between electricity consumption in test data and that in the corresponding training data.
T	Stored water temperature (K).
T_m	Mean stored water temperature (K), computed using temperatures specific to each node.
T_h	Hot water temperature (K).
T_c	Cold water temperature (K).
$T_{c,in}$	Inlet temperature of water in the cold heat exchanger (K).
T_i	Temperature of stored water in node i (K).
$T_{h,i}$	Temperature of hot water in node i (K).
$T_{h,in}$	Inlet temperature of water in the hot heat exchanger (K).
$T_{h,in}^{opt}$	Optimal inlet temperature of water in the hot heat exchanger (K) at a given timestep t
$T_{h,in}^{GP}$	Optimal inlet temperature of water in the hot heat exchanger (K) at a given timestep t
$T_{c,i}$	Temperature of cold water in node i (K).
UA	Overall heat transfer coefficient of heat exchanger (W/m^{2}K).
$\dot{V}_h$	Volume flow-rate inside the hot heat exchanger.
$\dot{V}_c$	Volume flow-rate inside the cold heat exchanger.
$\dot{V}_c^{\,opt}$	Optimal Volume flow-rate inside the cold heat exchanger.
$\dot{V}_c^{\,GP}$	Volume flow-rate inside the cold heat exchanger, as predicted by GP
$\dot{V}_c^{\,max}$	Maximum possible value of $\dot{V}_c$
s	Date-related variables used as inputs to the deep RNN model.
w_{ji}^m	Weight connecting j in layer m to node i in layer $m-1$
w	Weather variables used as inputs to the deep RNN model.
X	Input features to the deep RNN model.
x$_t$	Input to LSTM activation corresponding to a previous layer and current timestep t.
X$_t$	Training input for a given ML algorithm.
X$_e$	Test input for a given ML algorithm.
x^{opt}	Set of optimal operational variables at a given timestep.
x$_I^{opt}$	Set of optimal operational variables at a given timestep, corresponding to case I: $Q_{rec} \geq Q_d$
x$_{II}^{opt}$	Set of optimal operational variables at a given timestep, corresponding to case II: $Q_{rec} < Q_d$
x^{GP}	Set of approximations to optimal operational variables at a given timestep as predicted by GP
x$_I^{GP}$	Set of approximations to optimal operational variables at a given timestep as predicted by GP, corresponding to case I
x$_{II}^{GP}$	Set of approximations to optimal operational variables at a given timestep as predicted by GP, corresponding to case II

$\mathbf{X}$ Feature vector used as inputs to the deep RNN model.

y_p Predicted value of electricity consumption

y_a Actual value of electricity consumption.

$\mathbf{z_{I,train}}$ Training input for GP in optimization scheme corresponding to case I: $Q_{rec} \geq Q_d$

$\mathbf{z_{I,test}}$ Training input for GP in optimization scheme corresponding to case I: $Q_{rec} \geq Q_d$

$\mathbf{z_{II,train}}$ Training input for GP in optimization scheme corresponding to case II: $Q_{rec} < Q_d$

$\mathbf{z_{II,test}}$ Test input for GP in optimization scheme corresponding to case II: $Q_{rec} < Q_d$

1. Introduction

The building sector is responsible for a significant fraction of the primary energy consumption and greenhouse emissions in the U.S [1] - a good portion of which is contributed by space and water heating, as well as gas equipment usage [2]. With increasing application of distributed generation and storage systems in order to meet these demands, there is a need for forecasts of heating demands across different time horizons [3]. Such time horizons for forecasts could be (i) short-term (< 1 week), which is useful for real-time control and optimization of building energy components, short-term maintenance and immediate scheduling and management of generation capacity and storage [3–6] or (ii) medium-to-long term, which concerns planning, installation and management of distributed generation and storage systems [7], and decision-making related to demand response strategies [3].

This paper concerns the use of longer-term predictions in aiding design of a stratified thermal storage. Conventionally, deterministic energy simulation packages such as eQuest and EnergyPlus are used to estimate the heating and cooling loads in a building over a longer time horizon [8]. These physics-based models compute these loads by considering transient mass and energy balance between different connecting zones in a building. However, these energy simulation packages often require detailed knowledge of building construction and operational schedules - which are often not available in practice. Thus, these energy simulation packages often do not accurately predict future demands [9]. As these energy simulation packages require inputs which are often uncertain or difficult to obtain, they are often used as comparative tools, often prior to the building construction.

As such, machine learning (ML) models that predict future loads based on past observed data are often employed by energy researchers and modelers [10]. In prior literature, ML algorithms such as simple and multivariate linear regression [10, 11], non-linear regression [10–12], multi-layered perceptron neural networks [10, 13–15], auto-regressive techniques [13, 16], Gaussian Processes [17] and hybrid models combining ML models with deterministic thermal networks [18] have been used to predict electric, heating and cooling loads in buildings. While these methods, in general, have been successfully employed in short-term forecasts, comparatively limited work has been done in applying ML models for medium-to-long term forecasts. Making longer-term predictions is a more difficult objective, with comparatively higher associated errors [3, 11, 19].

Rahman and Smith compared the relative accuracies of several different machine learning algorithms in predicting fuel consumption (as simulated using EnergyPlus) in a small office, a supermarket and a restaurant, over a time period of up to one year at one-hour resolution [20]. The comparative study showed that with access to explicit knowledge of building schedules, a 3-layer multi-layered perceptron (MLP) neural network, in general, performs better in making point estimates than other multivariate linear regression, autoregressive models, ridge regression and Gaussian Processes. One possible reason as to why a 3-layer MLP performs comparatively better than other algorithms in making point estimates could be that MLP can model the complex, non-linear behaviors in energy consumption through adaptive feature extraction, without having an excessively large number of parameters that can contribute to overfitting . An extensive review of the application and relative performance of different machine learning algorithms is available in other literature [10, 20].

More recently, deep learning has been employed - both for short and long-term energy forecasting [3, 21, 22]. Deep learning is a particularly attractive option for longer-term forecasts as they are able to model expressive functions through multiple layers of abstractions [23]. Mocanu et al. used Conditional Restricted Boltzmann machine (CRBM) and Factored conditional Restricted Boltzmann machine (FCRBM) for longer-term forecasts of electricity consumption in a residential building at one-hour resolution [3]. Rahman et al. used a deep recurrent neural network

model to predict electricity consumption in commercial and residential buildings over medium-to-long term time horizon, without having access to building operational schedules [22]. The proposed model uses a sequence-to-sequence approach, similar to those used in machine translation and speech recognition applications [24–26] to generate surrogates for these transient operational schedules, while accounting for both short and long-term temporal dependencies in heating demand profile. While the deep RNN model has been employed to predict electricity consumption in commercial buildings with improved accuracy (compared to a multi-layered perceptron neural network), their performance in predicting heating demand in buildings is yet to be investigated. Thus, in this paper, the performance of the proposed deep RNN model in predicting heating demand over medium-to-long term time horizon investigated. Details of the deep RNN model are presented in section 3.

As mentioned, one of the goals of these longer term forecasts is to facilitate planning and management of distributed generation and storage systems. As a demonstration of a scheme where longer-term predictions made by an ML algorithm are used to aid design of an energy systems component, the predictions made by the deep RNN model developed in [22] were used for design optimization of a building-scale stratified thermal storage tank. A sensible stratified thermal storage tanks using water as a storage medium are a reasonable option for integration with a CHP unit [27–30], as they are inexpensive and have a high storage density in the low-temperature range (300 - 375 K) [31]. Presence of thermal stratification within a storage tank can significantly increase the exergy stored in the tank compared to a fully-mixed tank [32], and can potentially contribute to as much as 37% improvement in meeting heating demand in buildings [33].

Previous studies have shown that sizing of a thermal storage is critical for efficient and cost-effective integration of thermal storage with the CHP unit [29, 34–37]. Current TS sizing methods consider a demand profile over a design day at one-hour resolution, along with the specifications of other components in the CHP [29, 36]. However, these methods do not always consider the transient characteristics of the storage over a longer time period, i.e. time period over which the TS would be in operation. To simulate the transient behavior of TS (which, in turn, affect the CHP performance), physics-based models along with forecasts along with forecasts over the time period of interest need to be incorporated as part of the design process of the thermal storage.

In case of stratified thermal storage tanks using a sensible storage medium, the physics-based models are often one-dimensional, transient heat equations describing the temperature profile inside the tank [27, 38–41]. However, the characteristic timescale of these heat transfer processes are often in the order of 10 seconds to one minute [27, 40], whereas the time period of interest for which forecasts are made may span over more than one month. Thus, one of the key challenges in sizing a TS using forecasts over a longer time horizon is formulating an optimization scheme that can incorporate a TS model at a high temporal resolution ($\sim$ 10 s) with forecasts over a longer time horizon at one-hour resolution in a computationally efficient manner. In this paper, a scheme is proposed, which, for a given tank dimension, uses Bayesian optimization and Gaussian processes to determine the optimal operational variables of a stratified TS at a given time-step, while reducing the overall computational cost. As Bayesian optimization is computationally efficient, it was used to optimize the operational variables in a CHP with TS at each timestep (i.e. each hour) over a two-month time horizon. To reduce the computational cost even further, Gaussian process regression was applied to find approximations to the optimal operational variables, provided the Bayesian optimization method has been applied to a certain number of prior timesteps. The scheme is applied for multiple selections of tank sizes to evaluate the overall performance of the CHP with TS for each tank size. Details of the optimization framework are provided in section 4.

Therefore, based on the previous work, there is a research gap pertaining to evaluation of deep recurrent neural networks that are able to account for longer-term temporal dependencies in predicting heating demand in buildings - particularly when access to explicit schedule variables are not available. There is also a critical need to develop novel optimization frameworks that can utilize these longer-term predictions to aid optimal design of distributed generation and storage systems.

As such, the key objectives of this paper can be summarized as follows:

- Investigate the performance in deep RNN in predicting heating demand in multiple campus buildings at University of Utah over a period of seven months, compared to that of a 3-layered multi-layered perceptron model over the same time period.

- Develop an modeling scheme that uses RNN forecasts and a 1-D transient heat transfer model of stratified thermal storage tank to optimize the operational variables of a thermal-storage at each timestep (i.e. each hour).

Figure 1. Schematic Diagram of a multi-layered perceptron neural network - the colored boxes are used to distinguish between the different layers in the MLP. Layer 1 is the input layer where each node corresponds to input variables that contribute to energy consumption. The layer M is the output layer - in this context, this is the prediction of energy consumption. The relationship between a successive layers, e.g. a generic layer m and the previous layer $m-1$ is presented in equation 1.

- Evaluate the performance of the CHP with TS for different selections of tank sizes, based on RNN predictions and results from the optimization scheme.

2. Theoretical Background

2.1. Neural Networks

Neural networks are a class of machine learning algorithms that model non-linear relationship between an input and an output vector through multiple layers of interconnected processing elements. A schematic of a multi-layered perceptron (MLP) neural network is presented in figure 1. The output from a given node j in layer l can be computed as:

$$h_j^m = \phi\left(\sum_i^I w_{ji}^m h_j^{m-1}\right) = \phi(s_j) \tag{1}$$

Te error in the final layer can be calculated as follows:

$$e = (y_p - y_a)^2 \tag{2}$$

During training, the weights are updated using stochastic gradient descent as follows:

$$w_{ji}^m \leftarrow w_{ji}^m - \alpha \frac{\partial e}{\partial w_{ji}^m} \tag{3}$$

Here α is the learning rate. The partial derivatives $\frac{\partial e}{\partial w_{ji}^m}$ is determined through back propagation [23].

2.2. Recurrent Neural Networks and Long Short-term Memory

Recurrent neural networks (RNN's) are an extension to the conventional multi-layered perceptron neural networks that model temporal dependencies in target data through feedback connections. An activation function in an RNN that accounts for the temporal dependency with respect to the previous timestep $t-1$ can be expressed as follows:

$$h_{j,t}^m = \mathbf{w_1}\phi(h_{j,t-1}^m) + \mathbf{w_2}\phi(h_{i,t}^{m-1}) \tag{4}$$

Figure 2. Schematic diagram of a long short-term memory (LSTM) activation function

In the equation above, $h_{j,t}^m$ denotes the output of the activation function of node j in layer m at a given timestep t. However, as shown mathematically by Hochreiter and Schmidhuber [42], vanilla RNN's do not accomodate long term temporal dependencies due to the problem of vanishing gradients. As a solution to modeling these longer term temporal dependencies, the authors proposed the long-short term memory (LSTM) unit to be used in recurrent neural networks.

A schematic of an LSTM unit is presented in figure 2. As presented below, based on an input $(h_{t-1}, x_t) = (h_{j,t-1}^l, h_{i,t}^{m-1})$, the LSTM unit can adaptively scale the input, a transient memory state vector, and the output. Thus, the input to the LSTM unit $(h_{t-1}, x_t,)$, it is scaled by the input gate $\mathbf{i}$ as follows:

$$\mathbf{g} = \phi(\mathbf{w_{g1}h_{t-1}} + \mathbf{w_{g2}x_t} + \mathbf{b_g}) \tag{5}$$

$$\mathbf{i} = \sigma(\mathbf{w_{i1}h_{t-1}} + \mathbf{w_{i2}x_t} + \mathbf{b_i}) \tag{6}$$

In equation 6, σ is the sigmoid activation function applied to each element inside the parentheses. Subsequently, the transient memory vector $\mathbf{c_t}$ can be computed as follows using the forget fate $\mathbf{f}$:

$$\mathbf{f} = \sigma(\mathbf{w_{f1}h_{t-1}} + \mathbf{w_{f2}x_t} + \mathbf{b_f}) \tag{7}$$

$$\mathbf{c_t} = \mathbf{i} \circ \mathbf{g} + \mathbf{c_{t-1}} \circ \mathbf{f} \tag{8}$$

Here, $\circ$ is an element-wise multiplier. Similarly, the output gate $\mathbf{o}$ can be expressed as:

$$\mathbf{o} = \sigma(\mathbf{w_{o1}h_{t-1}} + \mathbf{w_{o2}x_t} + \mathbf{b_o}) \tag{9}$$

Finally, the output from the LSTM unit $\mathbf{h_t}$ can be expressed as:

$$\mathbf{h_t} = \mathbf{o} \circ \psi(\mathbf{c_t}) \tag{10}$$

The weights $\mathbf{w} = [\mathbf{w_{g1}}, \mathbf{w_{g2}}, \mathbf{w_{i1}}, \mathbf{w_{i2}}, \mathbf{w_{f1}}, \mathbf{w_{f2}}, \mathbf{w_{o1}}, \mathbf{w_{o2}}]$ and the bias vectors $\mathbf{b} = [\mathbf{b_g}, \mathbf{b_i}, \mathbf{b_f}, \mathbf{b_o}]$ are learned through back-propagation [23]. In this paper, the deep recurrent neural network uses the LSTM units to model temporal dependencies in heating demand profile.

2.3. Gaussian Process Regression

A Gaussian Process (GP) is simply a generalization of Gaussian Process distribution. GP defines a distribution over functions, and the inference takes place directly in the function space For a given test point $\mathbf{X_e}$, GP not only predicts a point estimate, but also a variance associated with the prediction. Given a training set $[\mathbf{X_t}, \mathbf{f_t}]$, the prediction $\mathbf{f_e}$ is made by considering the following joint distribution:

$$\begin{bmatrix} f_t \\ f_e \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(\mathbf{X_t}, \mathbf{X_t}) & K(\mathbf{X_t}, \mathbf{X_e}) \\ K(\mathbf{X_e}, \mathbf{X_t}) & K(\mathbf{X_e}, \mathbf{X_e}) \end{bmatrix}\right) \tag{11}$$

Thus, the covariance matrix is modeled using a Kernel $K(X_1, X_2)$, which, in this analysis is a squared exponential function indicating the proximity between X_1 and X_2 [43]. Thus, the point estimate associated with a prediction f_e can be expressed as :

$$\bar{f}_e = K(\mathbf{X_e}, \mathbf{X_t})[K(\mathbf{X_t}, \mathbf{X_t})]f_t \tag{12}$$

The variance associated with the prediction can be expressed as follows:

$$\sigma^2(f_e) = K(\mathbf{X_e}, \mathbf{X_e}) - K(\mathbf{X_e}, \mathbf{X_t})[K(\mathbf{X_t}, \mathbf{X_t})]^{-1}K(\mathbf{X_e}, \mathbf{X_t}) \tag{13}$$

In this paper the Bayesian Optimization package "bayesopt" was used, which implicitly calls the Gaussian Process algorithm. Furthermore, we also used Gaussian Process in the proposed optimization scheme to improve computational time. Further information about Gaussian Processes can be obtained in [43].

3. Deep RNN model

3.1. Model Inputs

The deep RNN model was applied to predict hourly heating demands (in KWh) in multiple campus building in University of Utah, Salt Lake City. The input variables to the RNN model were a concatenation of weather variables ($\mathbf{w}$) and date-related variables ($\mathbf{s}$). The weather variables were dry-bulb-temperature and humidity; whereas the date-related variables were (i) hour of day (from one to 24), (ii) day of week (iii) day in a month and month number. The day of the week was considered a concatenation of seven binary variables, each variable corresponding to a binary flag for each day of the week.

3.2. Model Formulation

The predictive model should ideally be adaptive (i.e. learn from data without human intervention) and able to model short and long-term temporal dependencies and non-linear patterns in energy consumption profiles, as well as model these transient profiles without having explicit knowledge of building schedules and building construction. In this analysis, the deep RNN model proposed by Rahman et al. [22] is used to predict heating demands at hourly interval.

Figure 3 presents a schematic diagram of the proposed deep RNN model, which has been previously been used to forecast electricity consumption in buildings. The proposed deep RNN model has, in general, performed better than a multi-layered perceptron (MLP) neural network in predicting hourly electricity consumption in commercial buildings; however, their performance in predicting heating demands is yet to be evaluated. The proposed model uses an encoder-decoder architecture-type that takes advantage of the sequential nature of electricity consumption, such as periodicity over a period of 24 hours. As such, layer 2 in the proposed model is analogous to an "encoder" layer that converts a given input sequence to a fixed vector representation, whereas layer 3 is analogous to a decoder layer that converts the vector representation to a target sequence. The Long short-term memory (LSTM) units in these two layers ensure that both long and short term temporal dependencies in heating demand profile are accurately modeled.

Intuitively, this means that the outputs from the layers 2 and 3 are transient variables that act-as surrogates to the schedule variables. The outputs in layer 3 are concatenated with the original inputs in layer 4 and introduced to a shared MLP layer in layer 5. The shared MLP layer accepts a vector at each timestep and effectively, it weights these

Figure 3. Schematic diagram of the proposed deep RNN model. Layers 2 and 3 contain LSTM units to model temporal dependencies and are analogous to the encoder-decoder architure. Outputs from layer 3 are concatenated with the original input in layer 4 before introducing to a shared MLP layer in layer 5. The final output is heating demand in 24-hour sequences.

transient variables according to the extent to which they affect the heating demand profile. Finally, the outputs in layer 6 are the heating demands in sequences of 24-hours.

The layers in the deep RNN model are described as follows:

Layer 1: Layer 1 contains the input features X at a temporal resolution of one-hour in 24-hour sequences, As such, The input data contains N samples where each sample is $\mathbf{X} \in \mathbb{R}^{24 \times d}$.

Layer 2: As mentioned, layer 2 converts the input sequence to a fixed vector representation:

$$\mathbf{c} = \mathbf{h_{t,e}} = LSTM(\mathbf{h_{t-1,e}}, \mathbf{X}) \tag{14}$$

Layer 3: Layer 3, i.e. the layer analogous to the decoder layer takes the fixed vector representation as input and generates a vector at each timestep that act as surrogates for schedule variables:

$$\mathbf{h_{t,d}} = LSTM(\mathbf{c}, \mathbf{h_{t-1,d}}) \tag{15}$$

Layer 4: Layer 4 concatenates these surrogates with the original input, which ensures that the dependencies with the original inputs are retained:

$$\mathbf{X'} = [\mathbf{X}; \mathbf{h_{t,d}}] \tag{16}$$

Layers 5 and 6: The concatenated input $\mathbf{X'}$ is introduced to a shared MLP layer (layer 5). The output in layer 6 is the heating demand obtained in 24-hour sequences.

$$y_{p,t=1,2...24} = MLP(\mathbf{X'_{t=1,2...24}}) \tag{17}$$

Table 1. Parameters for CHP system considered in this analysis

Parameter	Value
η_{PGU}	0.3
η_{rec}	0.95
ζ	0.8

Thus, the proposed deep RNN model attempts to model the transient effects - both long and short-term, which may have been caused by schedule variables such as occupancy and equipment schedules, without having explicit knowledge of these schedules. Unlike deterministic packages such as EnergyPlus, the proposed deep RNN model can therefore be used to estimate building heating demands in the immediate future without having access to actual building schedules, which usually difficult to obtain and predict in practice.

3.3. Model Optimization and Implementation

To prevent over-fitting, the deep RNN model is regularized using weight decay regularization and early-stopping. The hyper-parameters to be optimized in this model are (i) length of output vector from layer 2 (ii) Length of output vector from layer 3 - or the number of surrogates for transient variables (iii) Length of output from layer 5 and (iv) The selection of activation functions in layers 2, 3 and 5. These hyper-parameters were optimized using the hyperopt package [44, 45]. Details of the regularization and optimization methods are provided in [22].

The model was developed using the Keras API in Python on a Theano backend [46]. The accuracy of the heating demand predictions were evaluated using the root mean squared error, relative to the root mean squared average of the heating demand in the training set, as presented below:

$$e_1 = \frac{\sqrt{\sum_{i=1}^{T}(y_{a,e} - y_p)^2}}{\sqrt{\sum_{i=1}^{T}y_t^2}} \tag{18}$$

Here, $y_{a,e}$ is the actual heating demand in the test set, y_p represents the corresponding prediction made by the deep RNN mode and y_t denotes the heating demand in the training set.

4. Methodology

4.1. Overall Modeling Approach

The aim of the proposed modeling framework is to incorporate longer-term predictions of heating demand in a a building to optimize the design parameters of a sensible stratified thermal storage tank. Figure 4 shows the schematic of a building-scale combined heating and power (CHP) plant with thermal storage, as presented by Smith et al [34]. The heat recovered by the heat recovery unit can be expressed as [34]:

$$Q_{rec} = (\frac{1}{\eta_{PGU}} - 1)\zeta\eta_{rec}E_{gen} \tag{19}$$

Here, Q_{rec} refers to the heat recovered by the heat recovery unit, E_{gen} refers to the amount of electricity provided to the building by the CHP prime mover, η_{PGU} is the thermal efficiency (E_{gen} per unit of fuel) of the PGU, η_{rec} is the thermal efficiency of the heat recovery unit, and ζ is a factor to account for any energy losses outside of the heat recovery unit. In this analysis, $E_{gen} = 180$ kWh was considered. The values of the parameters η_{PGU}, η_{rec} and ζ are taken from those presented by Smith et al. [34] and are presented in Table 1.

The total heat supplied to the building, $Q_{T,bldg}$ can be expressed as:

$$Q_{T,bldg} = (1 - \gamma)Q_{rec} + Q_{TS,out} \tag{20}$$

Here $Q_{TS,out}$ is the heat delivered by the thermal storage tank and γ is the fraction of heat routed to the thermal storage from the heat recovery unit. It was assumed that if $Q_d > Q_{T,bldg}$, the additional heating demand is made up by an on-site boiler.
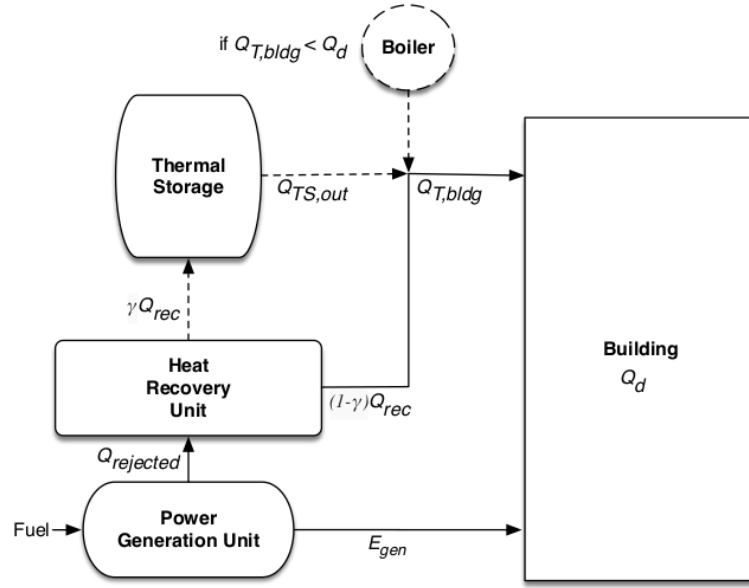
Figure 4. Schematic showing heat flow in a building CHP unit

Thus, in order to meet the heating demand Q_d at a given time t, the objective function for the optimization can be framed as follows:

$$O(\gamma(t), T_{h,in}(t), \dot{V}_c(t)) = -\frac{|Q_{T,bldg}(t) - Q_d(t)|}{Q_d(t)} + \mu\frac{Q_{st(t)}}{Q_d(t)} - \kappa\frac{\dot{V}_c(t)}{\dot{V}_{c,max}} \tag{21}$$

Here, Q_d is heat demand from the building, as obtained from the deep RNN model (as presented in section 3) and Q_{st} is the heat stored in the thermal storage tank within the given one-hour interval. The terms Q_{st} and $Q_{TS,out}$ are calculated using temperature profiles obtained using the thermal storage model (as described in section 4.2. Here The variables to be optimized at each timestep are: (i) The fraction of heat routed to the thermal storage tank, γ, (ii) The inlet temperature to the hot heat exchanger in the thermal storage tank, $T_{h,in}$ and (iii) The volume flow rate in the cold heat exchanger in the thermal storage tank, $\dot{V}_c$. It should be noted that these variables are optimized at one-hour intervals, and that while the deep RNN predictions are known beforehand, the thermal storage model is coupled and needs to be called at each one hour interval to compute Q_{st} and $Q_{TS,out}$.

The first term in equation 21 corresponds to the discrepancy between the building heating demand and the heat supplied, whereas the second term ensures priority of thermal storage when $Q_{rec} > Q_d$. In equation 21, μ is a parameter that indicates whether amount of heat stored in the thermal storage is prioritized in the optimization process. As such, $\mu = 1.3$ was picked when $Q_{T,bldg} > Q_d$, and $\mu = 0$ otherwise. $\mu > 1$ was selected when $Q_{T,bldg} > Q_d$ as the magnitude of heat stored Q_{st} is comparatively lower than the other energy terms in the objective function. κ is a parameter to bias the optimization solution against high volume flow-rates, as high flow-rates would require high amount of pump power. In this analysis, $\kappa = 0.05$ was selected. Here, $\dot{V}_{c,max} = 20$ gpm, the maximum allowable flow-rate at a given timestep. The flow-rate in the hot heat exchanger can be expressed as:

$$\dot{V}_h(t) = \frac{\gamma Q_{rec}(t)}{\rho_h(t)c_{p,h}(t)(T_{h,in}(t) - T_{rec,in})} \tag{22}$$

Here, $T_{rec,in}$ is the inlet water temperature to the heat recovery unit, and was considered to be 300 K. The optimization is performed to determine the operational variables for 59×24 timesteps, corresponding to heating demand at one-hour resolution between January 01, 2017 12:00 AM and February 28, 2017 11:59 PM. Table 2 indicates the

Table 2. Search space for operational variables

Operational variable	Allowable Range
Flow-rate in cold heat exchanger, V_c	[0, 20] gal/min
Temperature of hot water entering tank, $T_{h,in}$	[310, 390] K
Fraction of heat routed to thermal storage, γ	[0, 0.5]

Table 3. Options for thermal storage tank selection

Option	Nominal Volume, L	Tank Diameter	Tank Height	Manufacturer
I	926	0.990	2.015	SolarBayer
II	1422	1.00	2.145	SolarBayer
III	2204	1.25	2.095	SolarBayer
IV	2371	1.25	2.245	SolarBayer
V	2920	1.25	2.695	SolarBayer
VI	4960	1.60	2.795	SolarBayer
VII	7192	2.438	1.829	Hydroflex Systems Inc.
VIII	8327	2.438	2.134	Hydroflex Systems Inc.

minimum and maximum values of operational values allowed for each timestep, i.e. search space for the optimization algorithm. The process is repeated for multiple discrete options of tank diameter and height, as presented in table 3.

4.2. Thermal Storage Model

As mentioned in Introduction, a stratified, sensible thermal storage tank is considered for integration with the CHP unit. The subsections below describe the specifications and mathematical modeling of the aforementioned stratified thermal storage tank.

4.2.1. Thermal Storage Tank Description

The thermal storage tank considered for this analysis is one with two heat exchangers: a hot and a cold heat exchanger, and no mass flow in and out of the tank [40, 41]. Figure 5 shows a schematic diagram of the thermal storage tank considered. The proposed type of storage tank requires little maintenance and is likely to avoid leakage problems. Water was considered as both the storage and heat transfer fluid. The heat exchangers were considered to run through the entire length of the cylindrical tank, with the flow inside the heat exchangers occurring in counter-flow arrangement, as shown in the figure. The specifications of the heat exchangers are taken from [47] and are provided in table 4.

As mentioned in section 4.1 [3], the optimization process to determine the operational parameters (i.e. γ, $\dot{V}_c$ and $T_{h,in}$) at each timestep is repeated for each selection of (D, H), where D is the diameter of the storage tank and H is the height. The selections of D and H were taken from [48] and [49]. The costs associated with the SolarBayer tanks, i.e. tanks between options I and VI, range approximately between \$1,250 (1050 euros) to \$5,550 (4500 euros), corresponding to the range of nominal volume between [927, 4960] L; whereas the costs associated with Hydroflex Systems storage tanks, i.e. options VII and VIII are approximately \$4200 and \$4600 respectively.

4.2.2. Mathematical Model of Thermal Storage Tank

The mathematical model of a stratified thermal storage proposed by Rahman et al. [41] is considered for this analysis. The proposed model is a 1-D, transient heat transfer model that describes the temperature profile inside the tank and the two heat exchangers. The proposed model was suggested as a screening tool for sizing of stratified thermal storage and is based on the following assumptions:

- The tank was assumed to be adiabatic, with no stored water leaving or entering the tank.

- The water inside the tank and the heat exchangers is incompressible.

Figure 5. Schematic used to develop mathematical model of stratified storage tank

Table 4. Specifications of heat exchanger considered

Material Conductivity (k_{mat})	30 (W/m K)
Coil pitch (p)	36.2 mm
Coil Diameter (D_{coil})	0.49 m
Inner tube diameter (d_i)	21.6 mm
Outer tube diameter (d_o)	26.9 mm

- The water inside the tank and the heat exchangers are at a sufficiently high pressure (> 0.5 MPa) such that no phase change occurs.

- The temperature variation in the radial direction was neglected. .

For a given node i, the temperature of stored water T_i can be expressed using the following heat equation:

$$m_i c_p \frac{dT_i}{dt} = UA_h(T_{h,i} - T_i) + \frac{kA_c(T_{i-1} - T_i)}{\Delta x} + \frac{kA_c(T_{i+1} - T_i)}{\Delta x} + UA_c(T_{c,i} - T_i) \tag{23}$$

Here m_i is the mass of water in node i, $T_{h,i}$ is the temperature of the heat exchanger in node i, UA_h is the heat transfer coefficient between the stored water and the hot heat exchanger in node i, UA_c is the heat transfer coefficient between the stored water and the cold heat exchanger in node i, and A_c is the cross sectional area of the tank.

The temperatures of the hot heat exchanger and cold heat exchanger in node i can be computed similarly as:

$$m_{h,i} c_{p,h} \frac{dT_{h,i}}{dt} = \dot{m}_h c_{p,h}(T_{h,i-1} - T_{h,i}) - UA_h(T_{h,i} - T_i) \tag{24}$$

$$m_{c,i} c_{p,c} \frac{dT_{c,i}}{dt} = -\dot{m}_c c_{p,c}(T_{c,i} - T_{c,i+1}) + UA_c(T_i - T_{c,i}) \tag{25}$$

Here $\dot{m}_h$ and $\dot{m}_c$ are the mass flow-rates in the hot heat exchanger and cold heat exchanger respectively. The boundary values are modeled as follows:

$$m_{h,1} c_{p,h} \frac{dT_{h,1}}{dt} = \dot{m}_h c_{p,h}(T_{h,in} - T_{h,1}) - UA_h(T_{h,1} - T_1) \tag{26}$$

$$m_{c,N} c_{p,c} \frac{dT_{c,N}}{dt} = -\dot{m}_c c_{p,c}(T_{c,N} - T_{c,in}) + UA_c(T_N - T_{c,N}) \tag{27}$$

The overall heat transfer coefficient UA was computed by considering a thermal circuit between the stored water and the given heat exchanger in node i:

$$UA = \frac{1}{\frac{1}{h_i A_i} + \frac{ln\frac{d_o}{d_i}}{2\pi k_{mat}\Delta x} + \frac{1}{h_o A_o}} \tag{28}$$

Here h_i is the heat transfer coeffecient corresponding to forced convection inside the heat exchangers, and was computed using correlation proposed by Pethukov for turbulent flow [50] and by Xin and Ebadian for laminar flow [51]. The heat transfer coefficient corresponding to the free convection at the outer surfaces of the heat exchanger was determined using Ali's correlation [52].

The heat balance equations 23 - 27 were solved using an implicit scheme in MATLAB with a timestep $\Delta t = 10s$ for a 10-node tank. At a given time t and node i, the solutions to the heat balance equations give the transient temperature profiles $T_i(t)$, $T_{h,i}(t)$ and $T_{c,i}(t)$. The heat delivered by the thermal storage tank over a 1-hour time interval can be expressed as:

$$Q_{TS,out} = \int_{t=0}^{t=1h} \rho(t)\dot{V}_c c_{p,h}(t)[T_{c,out}(t) - T_{c,in}]dt \tag{29}$$

Here the outlet temperature in the cold heat exchanger $T_{c,out}(t) = T_{c,1}(t)$ computed from the heat equations. We considered the inlet temperature to the cold heat exchanger, $T_{c,in}$ to be constant at 300 K. The heat stored in a given node i over the given 1-hour interval can be expressed as:

$$Q_{st,i} = \int_{t=0}^{t=1h} m_i c_{p,i}(t)[T(t) - T(t-1)]dt \tag{30}$$

Thus the total heat stored in the tank over the 1-hour interval is $Q_{st} = \sum_i^N Q_{st,i}$. At t= 0, we considered the stored water temperature to be at 330 K.

Further details on the implementation of the thermal storage tank are available in other literature [41].

4.3. Optimization Method
The subsections that follow describe the details of the optimization framweowork used in this analysis.

4.3.1. Selection of Optimization Technique
The overall goal of the optimization scheme is to optimize the operational parameters, i.e. $\gamma(t)$, $\dot{V}_c(t)$ and $T_{h.in}(t)$, at each timestep for a given diameter and height of tank, before repeating the process for multiple discrete combinations of (D, H). The bayesopt package in Python, which employs a Bayesian Optimization scheme [53] is selected to solve the optimization problem presented in section 4.1.

Bayesian optimization is a type of non-linear optimization scheme which maximizes or minimizes a given expensive function by constructing a probabilistic surrogate model for the aforementioned expensive function using prior function evaluations [44, 45, 53, 54]. This makes Bayesian optimization a computationally efficient algorithm for solving non-convex problems which often do not hae closed-form expressions [54], as it uses fewer evaluations of the computationally expensive function without relying on local gradients or Hessians [53]. Bayesian optimization is a particularly attractive option for the optimization problem presented here (which is non-convex, computationally expensive and does not have a closed-form expression), as we need to perform the optimzation at every timestep (i.e. every hour) for 59×24 hours ahead - and ideally would like to avoid too many evaluations of temperature profiles using the computationally expensive thermal storage model for every instance where the optimization algorithm is applied.

The Bayesian optimization method assumes that the computationally expensive function is drawn from a Gaussian Process prior [53]. As such, based on these priors and observations (i.e. function evaluations), the optimization uses an acquisition function $a(\mathbf{x})$ to probe for the subsequent point by solving the proxy optimization $\mathbf{x_{next}} = \text{argmax}_x a(\mathbf{x})$. In this analysis, the probability of improvement (PI) was used as the acquisition function [53].

Details on theoretical aspects of Bayesian optimization are available in other literature [53, 54]

4.3.2. Reducing Computational Cost using Gaussian Processes

While the Bayesian optimization technique has a reduced number of function evaluations every timestep, the optimization algorithm itself needs to be called for a large number of timesteps (i.e. the forecast period of 59×24 hours for which the thermal storage is being sized). To reduce the computational time even further, we propose to use Gaussian Process regression to estimate the solutions to the Bayesian optimization process at a given timestep without explicitly performing the optimization, provided that optimizer has been called a certain threshold number of instances under similar conditions prior to that timestep.

Thus, it was assumed that the optimal operational parameters $\mathbf{x^{opt}}(t) = [\gamma^{opt}(t), V_c^{opt}(t), T_{h,in}(t)]$ at a given timestep is a function of heating demand $Q_d(t)$ and the mean temperature of stored water in the tank, $T_m(t)$ at the given timestep. The GP was applied separately for case I: $Q_{T,bldg} \geq Q_d$ (i.e. when the heat supplied by the CHP unit exceeds the building demands and case II: $Q_{T,bldg} < Q_d$.

Thus for case I: $Q_{T,bldg} \geq Q_d$, the solutions from Gaussian Process regression can be expressed as:

$$\mathbf{x_I^{GP}} = f(\mathbf{z_I}) \sim \mathcal{N}[m(\mathbf{z_I}), K(\mathbf{z_{I,train}}, \mathbf{z_{I,test}})] \tag{31}$$

Here $\mathbf{x_I^{GP}}$ are the solutions from Gaussian process regression, $\mathbf{z_I} = [Q_{d,I}^*, T_{m,I}^*]$ where $Q_{d,I}^* = \frac{Q_{d,I}}{Q_{d,I}^{max}}$ is the non-dimensional heating demand and $T_{m,I}^* = \frac{T_m(t) - T_m(t=0)}{T_{h,in}^{max} - T_m(t=0)}$ is the non-dimensional mean stored water temperature when $Q_{rec} \geq Q_d$.

The predictions from the Gaussian Process Regression (which are considered as potential solutions for the optimization process), $\mathbf{x_I^{GP}}(\mathbf{t}) = [\gamma^{GP}(t), V_c^{GP}(t), T_{h,in}^{GP}(t)]$, along with the associated standard deviations $[\sigma_g(\gamma_g(t)), \sigma_g(V_c(t)), \sigma_g(T_{h,in}(t))]$, are obtained using the scikit package in Python. As mentioned, the GP regression is only applied when the length of the training data for case I, M_I, exceeds a minimum threshold, and the test point for which the GP regression is going to be applied, $\mathbf{z_{I,test}}$ has been "observed" in the training set. This means for GP to applied, $Q_{d,I,test}^*(t) \leq max(\mathbf{Q_{d,I,train}^*})$ and $T_{m,I,test}^*(t) \leq max(\mathbf{T_{m,I,train}^*})$.

The solutions to the GP regression process are only considered acceptable approximations to the solutions from Bayesian optimization if the standard deviations associated with the GP predictions are within given thresholds. For this analysis, the GP predictions were considered as acceptable approximations, i.e. $\mathbf{x_I^{opt}}(\mathbf{t}) = \mathbf{x_I^{GP}}(\mathbf{t})$, if $\sigma_g(\gamma) < 0.02$, $\sigma_g(V_c) < 0.5$ gpm and $\sigma_g(T_{h,in}) < 3$ K. If the conditions are not satisfied, or if the confidence on the GP predictions are not within the acceptable thresholds, the Bayesian optimization process is called upon, and the data corresponding to the given timestep is stored for subsequent GP predictions. Furthermore, to ensure that the size of the training set is growing with increasing number of timesteps, the Bayesian optimization method is called upon for every fifth timestep.

For case II, i.e. when $Q_{rec} \geq Q_d$, an identical process is followed to find the V_c^{opt}. For case II, $\gamma^{opt} = 0$ and $T_{h,in}^{opt} = T_{h,in}^{min}$ was set, such that when the heating demand is higher than than the maximum capacity of the heat recovery unit, all of the heat from the heat recovery unit is routed to meet the demands of the building. For this analysis, $M_I = M_{II} = 80$ was considered.

To summarize, the GP regression is only used to approximate the solutions of the Bayesian Optimization process only when (i) The size of the training data for each individual case exceeds some minimum threshold, (ii) A point similar to the one for which the prediction is being made has been observed in the training set and (iii) The confidence associated with these predictions are within acceptable thresholds. Algorithm 1 in section 8 presents the proposed optimization scheme in its entirety.

Thus, the optimization framework provides a means of linking heat transfer processes which occur at very small timescales with design optimization processes that may require demand forecasts over a significantly longer timescale. Results obtained using this optimization framework are presented in section 5.

5. Results and Discussion

5.1. Forecasting heating demand in campus buildings

Table 5 shows the performances of the proposed deep RNN model and the 3-layer MLP model in forecasting heating demand in multiple campus buildings at University of Utah, Salt Lake City. The table shows that in general,

Table 5. Relative errors of deep RNN and multi-layered perceptron neural network in predcting heating demand in buildings between January 01, 2017 to July 31, 2017. The corresponding absolute errors in kWh are indicated in parentheses.

Building	Meter Type	s-value	max($\mathbf{y_t}$), kWh	e_1 Deep RNN Model	e_1 MLP Model
Merrill Engineering Building	HTW Meter, Energy Total	0.143	510	0.271 (62.6 kWh)	0.278 (64.2 kWh)
Social Behavior Science Hall	HTW Meter, Zone 1 Energy Totalizer	0.279	1430	0.355 (191 kWh)	0.552 (298 kWh)
Marriott Center for Dance	HTW Meter, Zone 1 Energy Totalizer	0.263	0.161	0.347 (0.00900 kWh)	0.439 (0.0113 kWh)
Emery Building	HTW Meter, Zone 1 Energy Totalizer	0.136	102	0.300 (11.0 kWh)	0.353 (12.9 kWh)

the proposed RNN model perofrms better than the MLP and that the gain in accuracy in using deep RNN model generalizes fairly well for different types of load profiles. The table also shows that the proposed deep RNN algorithm is agnostic to the type of meter reading used as targets, and in general, is able to model heating demand profiles of different magnitudes comparatively better than a 3-layer MLP.

The relative gain in accuracy in using the deep RNN, however, could be different depending upon the the heating demand profile. Figures 6 - 9 show the predictions of the deep RNN model and the MLP model, along with the ground truth for the forecast period of January 01, 2017 12:00 AM to July 31, 2017 11:59 PM. As the deep RNN model contains LSTM units that can model longer-term temporal dependencies in the transient demand profile, it may be likely that the improvement in accuracy is more significant in demand profiles where these longer-term temporal dependencies exist [22]. To quantitatively evaluate the effects of these longer term dependencies, Rahman et al. proposed the following evaluation metric as an indicator of discrepancy between the heating demands in the training and the test [22]:

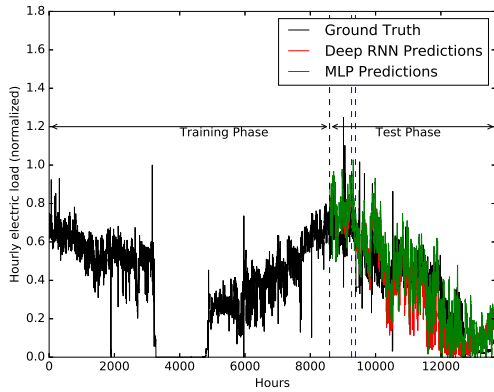$$s = \frac{\sqrt{\sum_{i=1}^{T}(y_{a,e} - y_t')^2}}{y_{a,e}^{max} - y_{a,e}^{min}} \tag{32}$$

Figure 6. Predictions of heating demand in Merill Engineering Building by deep RNN model ($e_1 = 0.271$) and MLP model ($e_1 = 0.278$ between January 01, 2017 and July 31, 2017. The maximum hourly load (in training) is 510.0 kW-h and s-value, $s = 0.43$.
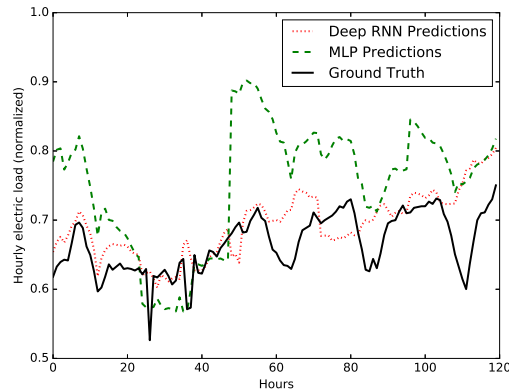
Figure 7. Predictions of heating demand in Merill Engineering Building by deep RNN model and MLP model between January 21, 2017 and January 26, 2017.

In the equation above, $y_{a,e}$ is the actual heating demand in the test set, $y_{a,e}^{max}$ is the maximum value of the heating demand $y_{a,e}$ in the test set, $y_{a,e}^{min}$ is the minimum value in $y_{a,e}$, and y_t' is the corresponding heating demand in the training data over the same forecast period (in this case, this would be between Janaury 01, 2016 to July 31, 2016). Thus, we assume that the discrepancy between heating demands in the training and test is due to longer-term temporal dependencies. Figure 6 and 8 illustrate that the deep RNN performs markedly better than MLP's for higher values of metric s. This supports the claim that the deep RNN's can adequately account for temporal dependencies in heating demand profile, and is consistent with earlier findings in [22], obtained for electricity consumption profiles.
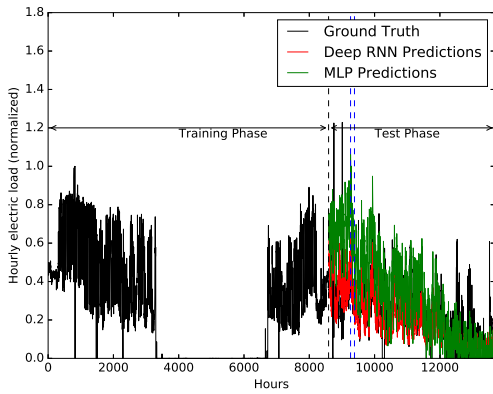
Figure 8. Predictions of heating demand in Social Behavior Science Hall by deep RNN model (e_1 = 0.355) and MLP model (e_1 = 0.552 between January 01, 2017 and July 31, 2017. The maximum hourly load (in training) is 1430 kW-h and s-value, $s = 0.279$.
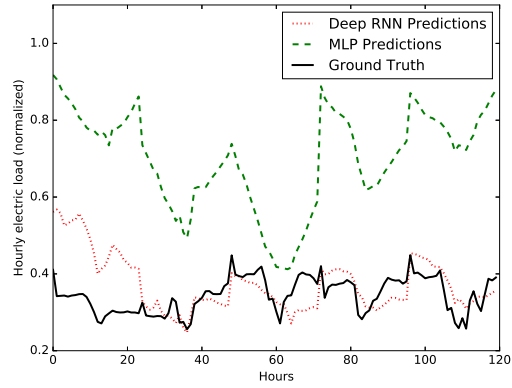

Figure 9. Predictions of heating demand in Social Behavior Science Hall by deep RNN model and MLP model between January 21, 2017 and January 26, 2017.

As observed in table 5 and 6, the Merill Engineering Building (MEB) profile is a case where relative errors by the proposed RNN and the MLP model have the closest proximity. Figure 10 shows the breakdown of errors e_1 corresponding to each month in the forecast period. The plot shows that, for the MEB demand profile, the proposed error e_1 is comparatively lower for the deep RNN model for January and February, whereas the MLP model does better for the summer months. Thus, the predictions made by deep RNN were used for sizing the building-scale thermal storage (as described in the following section), for which the forecast period between January 01, 2017 12:00 AM and February 28, 2017 11:59 PM was considered.
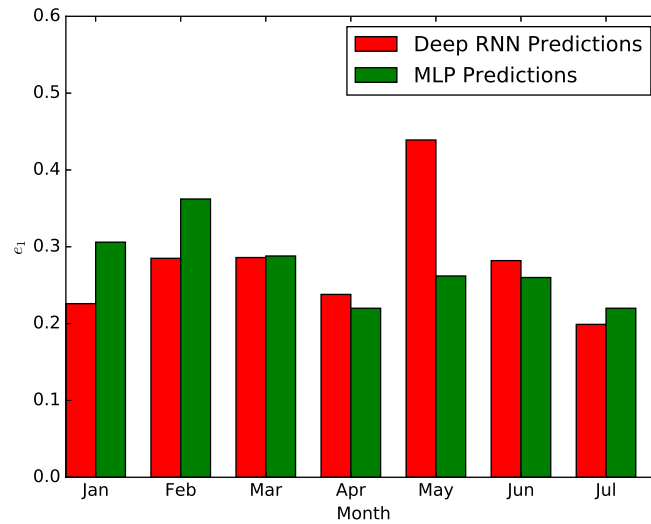

Figure 10. Errors e_1 corresponding to predictions by deep RNN and MLP with respect to different months within the forecast period.

5.2. Optimization of thermal storage using deep RNN predictions

The optimization scheme proposed in 4.3 is now applied for different tank size ons mentioned in table 3. The performance of the CHP with thermal storage (CHP-TS) was evaluated using the following metric, as proposed by Smith et al. [34] :

$$R_h = \frac{Q_{TS,bldg}}{Q_d} \tag{33}$$

The proposed metric R_h indicates the fraction of heating demand that is met by CHP-TS at a given time. For a given tank size, the R_h was computed at each timestep where $Q_d > Q_{rec}$, the mean of which was considered as the representative R_h for the given tank. Figure 11 shows how the metric R_h varies for different tank sizes, indicated here by the nominal volume of the tank. The plot shows that when R_h was calculated using heating demands predicted by the deep RNN model, R_h varies linearly ($R^2 = 0.9150$) with respect to the tank volume. Thus, this plot allows the modeler to estimate the performance of CHP-TS for a selection of tank size, given space and cost limitations.
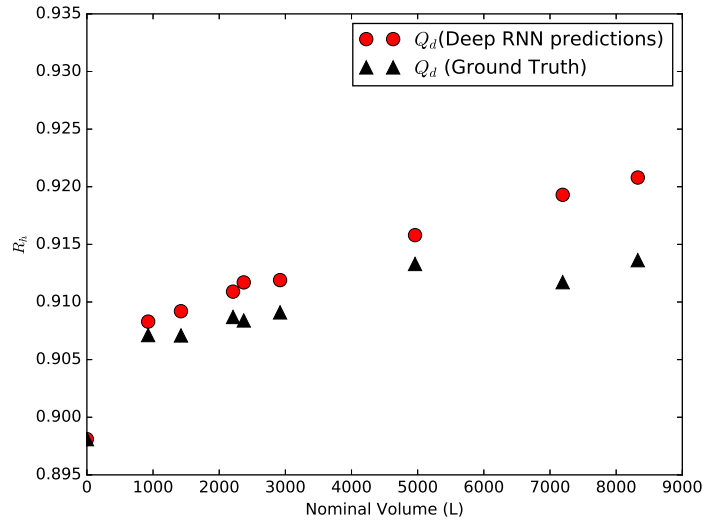

Figure 11. Comparison of R_h vs. Nominal Tank Volume when (i) Q_d was obtained using deep RNN predictions and (ii) Actual data for Q_d was used

Let us the case where give space and cost constraints, the energy modeler selects a nominal tank volume of 4960 L (i.e. on VI). Figures 12 - 15 present compare the heating demand and the heat supplied to the building for the given tank size - considering both cases where the heating demand from deep RNN predictions and the actual demand were used. In both cases, the heat supplied to the building follows the demand with reasonable accuracy - the root mean squared average of the discrepancies between $Q_{TS,bldg}$ and Q_d were found to be 11.6% and 11.2% respectively for the aforementioned two cases (i.e. figures 12 and 14). This indicates that the optimization scheme proposed performs fairly well, given that the optimization scheme had to be applied for 59×24 timesteps and a Gaussian Process approximation of the mal operational variables were considered within the scheme.

At a nominal volume of 4960 L, the discrepancy between the R_h values corresponding to the two cases was a margin of 0.0024 (i.e. 0.24%). This can be considered reasonable, particularly because: (i) the discrepancy is partly contributed by the demand at certain timesteps being significantly higher than the capacity of CHP with TS, and (ii) in practice it is extremely unlikely that the future demand would be known at one-hour resolution, and unlike conventional energy simulation packages, the deep RNN can provide these predictions that serve as proxy for future demand without having access to information about building operational construction or schedules. The maximum margin of improvement in R_h with respect to a CHP plant with no thermal storage (corresponding to nominal volume = 0 in figure 11) is 2.27%, corresponding to a nominal tank volume of 8327 L.

Thus, the physical significance of the results in figure 11 is that in general, the performance of the CHP-TS can be simulated reasonably well - particularly for comparatively smaller tank sizes, when the deep RNN predictions were used as inputs to the optimization framework. This makes the optimization framework useful, as it allows the performance analysis of a CHP-TS without the knowledge of future heating demand at a high (i.e. 1-hour) resolution, which, in practice, is likely to be unknown. However, the plot shows that the discrepancies between the R_h values obtained using the deep RNN predictions and those using the actual demand are comparatively higher for higher tank sizes, as observed from the las two points. One possible explanation for this could be that as the larger tanks have a greater capacity to fulfill the building heating demands, the prediction errors associated with the deep RNN manifest in larger discrepancies between the two R_h values.

It should be emphasized that the goal of this presented paper is not to optimize the tank size directly based on energy/cost considerations, but to optimize the operational parameters when the thermal storage is integrated with the CHP unit, so as to simulate, for a given thermal tank configuration, how the CHP-TS will perform given deep RNN estimations of future demand. As the thermal storage tank ons are usually discrete and the number of options are usually small for a given thermal storage application, it may be feasible to simulate the performance of CHP-TS using the optimization framework for each on. Subsequently, based on these performance estimations (as exhibited in figure 11), the building owner/engineer can make a decision on tank selection, knowing cost and space constraints. It should be noted that to directly optimize the tank size based on cost considerations is complex (as life cycle, maintenance and other factors need to be considered) and is beyond the scope of this analysis.
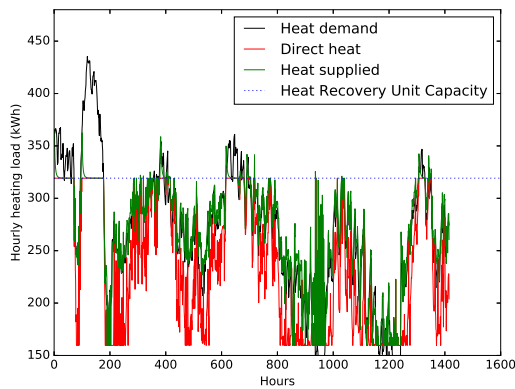

Figure 12. Comparison of heating demand (as predicted by deep RNN), heat delivered directly by the heat recovery and total heat supplied to the building between January 01, 2017 and February 28, 2017. Nominal tank volume = 4960 L
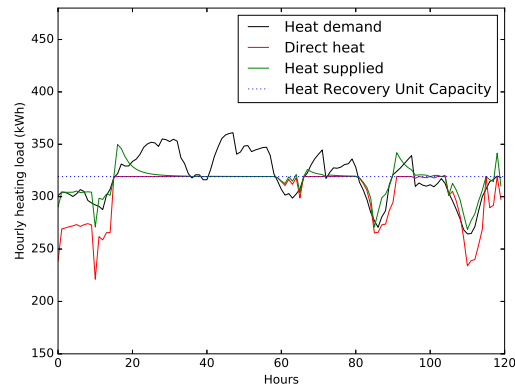
Figure 13. Comparison of heating demand (as predicted by deep RNN), heat delivered directly by the heat recovery and total heat supplied to the building over a 5-day period between January 21, 2017 and January 26, 2017. Nominal tank volume = 4960 L

5.3. Performance of Proposed Optimization Scheme

To analyze the relative benefits of using the proposed optimization scheme with GP regression, a comparison of the accuracy and computational time obtained using the proposed optimization scheme with those obtained when the proposed scheme was not applied were studied (i.e. lines 5 to 27 in algorithm 1 were skipped). Figures 16 and 17 compare the number of calls to the Bayesian optimization method (i.e. line 28 in algorithm 1) values of R_h respectively for these two different cases across multiple time horizons, corresponding to a tank volume of 4960 L.

Figure 16 shows that the GP scheme significantly reduces the fraction of timesteps for which the Bayesian optimization is actually performed, which in turn, reduces the number of timesteps for which the optimizer calls the computationally expensive thermal storage model. This reduction in calls to "bayesopt" package increases with increasing time horizon, and it directly corresponds to a reduction in overall computational time. For instance, figure 16 shows that at a time horizon of 59 days, i.e. when the number of timesteps over which the optimization method without the GP scheme would need to operate is 59×24, the reduction in the number of calls to the optimizer is by a factor of 0.332. At that time horizon, the computational time corresponding to the proposed scheme with the

Figure 14. Comparison of heating demand (as predicted by deep RNN), heat delivered directly by the heat recovery and total heat supplied to the building between January 01, 2017 and February 28, 2017. Nominal tank volume = 4960 L


Figure 15. Comparison of heating demand (as predicted by deep RNN), heat delivered directly by the heat recovery and total heat supplied to the building over a 5-day period between January 01, 2017 and January 06, 2017. Nominal tank volume = 4960 L

GP was found to be 4.43 hours compared to 14.41 hours without the proposed scheme - resulting in a reduction in computational time by a factor of 0.3072. Thus, the GP significantly reduces the calls to the Bayesian optimization function, which corresponds to a reduction to the overall computational cost.


Figure 16. Number of calls to the bayesopt package vs. time horizon over which mization is performed. Nominal tank volume = 4960 L, and $M_I = M_{II} = 80$
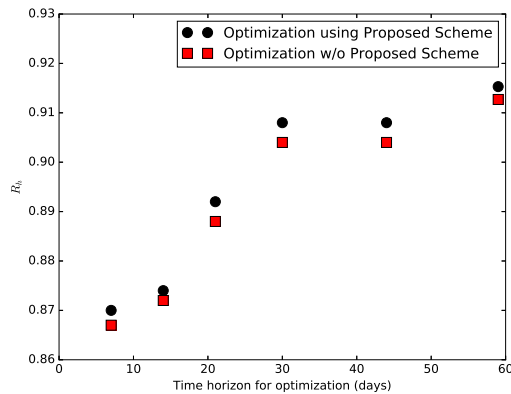

Figure 17. R_h vs. time horizon over which mization is performed. Nominal tank volume = 4960 L , and $M_I = M_{II} = 80$

Figure 17 shows how the values of R_h obtained using the proposed optimization scheme, compared with those obtained without the scheme, vary with different time horizons. At a time horizon of 59 days, the discrepancy between the R_h values obtained with and without the scheme is a margin of 0.0026 (0.26%). Thus, the physical interpretation of the results obtained in figures 16 and 17 is that the proposed mization framework using Gaussian Processes can use the deep RNN predictions to simulate the behavior of CHP-TS with significantly improved computational efficiency and with little loss in accuracy.

6. Further Discussions, Limitations, and Future Work

This paper evaluates the accuracies of deep recurrent neural network in prediction longer-term heating demand in multiple campus buildings at one-hour resolution, and proposes a computationally-efficient optimization framework

that can be used to aid the design of a stratified thermal storage tank. As the results in section 5 indicate, the deep RNN performs well in predicting heating demand in campus buildings, and the proposed optimization framework using Gaussian Processes can significantly improve the computational time. However, the deep RNN model and the proposed optimization scheme has the following limitations:

- The predictive deep RNN model assumes knowledge of future weather data while making predictions over medium-to-long term,, and does not account for uncertainties associated with inputs.

- The predictions of heating demand made by the deep RNN are specific to the past data of the building it was trained on. While the deep RNN model can account for longer-term temporal dependencies in the training data, it is likely to perform poorly if there are drastic changes in the test set pertaining to the building structure, equipment or occupancy that are not represented in the training set.

Based on these limitations, there is scope for future work pertaining to development of deep neural networks that can account for uncertainties in inputs. Another potential direction of future research could be development of machine learning models to act as surrogates of EnergyPlus modules that require uncertain and/or unknown inputs, so as to effectively integrate a machine learning framework with EnergyPlus. Furthermore, the proposed optimization scheme can potentially be applied, modified and adapted to aid design of other types of thermal storage, as well as other distributed generation components, such as solar collectors and heat exchangers in heat recovery units.

7. Conclusions

In this paper, we investigated the performance of a deep recurrent neural network model (previously proposed in [22]) in predicting heating demand in campus buildings at University of Utah over a longer time horizon, and proposed an optimization scheme to illustrate how longer-term predictions can be used to assist sizing of a thermal storage tank. The following conclusions can be drawn from this analysis:

- The deep RNN model, in general, performs better than a 3-layer MLP in predicting heating loads in campus buildings at University of Utah. The improvement in accuracy is more significant when the demand profile exhibited long-term temporal dependencies.

- The longer-term predictions can serve as proxy for the future demand, and can allow estimation of performance characteristics of a CHP unit with a stratified thermal storage tank for different tank sizes.

- The proposed optimization scheme provides a feasible means of incorporating longer-term predictions and a physics-based model at a high-temporal resolutions to assist sizing of a building-scale thermal storage tank.

8. Appendix

Algorithm 1 Proposed Optimization Scheme

1: **Define** tank diameter D and height H.
2: **Initialize** temperature of stored water and water inside the heat exchangers for all nodes, i.e. set $\mathbf{T} = \mathbf{T}(t = 0)$, $\mathbf{T_h} = \mathbf{T_h}(t = 0)$ and $\mathbf{T_c} = \mathbf{T_c}(t = 0)$
3: Define search space for optimization of operational parameters $\mathbf{x} = [\gamma, V_c, T_{h,in}]$
4: **for** $t = 0, 1, 2, ...t, ...T$ **do**
5:     **if** $Q_{rec} > Q_d$ **then**
6:         **if** $MOD(t, 5) \neq 0$ & $size(\mathbf{z_{II,train}}) > M_{II}$ & $Q_d^* < max(\mathbf{Q_{d,II}^*})$ & $T_m^* < max(\mathbf{T_{m,II}^*})$ **then**
7:             Use GP to predict V_c^{GP} and $\sigma_g(V_c^{GP})$ after training the GP model on $[\mathbf{z_{II,train}}, \mathbf{x_{II,train}}]$.
8:         **end if**
9:         **if** $\sigma_g(V_c^{GP}) > 0.5$ gal/min **then**
10:             Set $V_c^{opt} = V_c^{GP}$
11:             Set $T_{h,in}^{opt} = T_{h,in}^{min} = 310$ K
12:             Set $\gamma^{opt} = 0$
13:             Update $T_i(t)$, $T_{h,i}(t)$ and $T_{c,i}(t)$ using $\mathbf{x^{opt}} = [\gamma^{opt}, V_c^{opt}, T_{h,in}^{opt}]$
14:             Set $t = t + 1$, **go to line 4** (i.e. beginning of for loop) and *continue* loop.
15:         **end if**
16:     **else**
17:         **if** $MOD(t, 5) \neq 0$ & $size(\mathbf{z_{I,train}}) > M_I$ & $Q_d^* < max(\mathbf{Q_{d,I}^*})$ & $T_m^* < max(\mathbf{T_{m,I}^*})$ **then**
18:             Use GP to product γ^{GP}, $T_{h,in}^{GP}$, V_c^{GP} and $\sigma_g(\gamma^{GP})$, $\sigma_g(T_{h,in}^{GP})$, $\sigma_g(V_c^{GP})$ after training the GP model on $[\mathbf{z_{I,train}}, \mathbf{x_{I,train}}]$.
19:         **end if**
20:         **if** $\sigma_g(V_c^{GP}) > 0.5$ gal/min & $\sigma_g(\gamma^{GP}) > 0.02$ & $\sigma_g(T_{h,in}^{GP}) > 3$ K **then**
21:             Set $V_c^{opt} = V_c^{GP}$
22:             Set $T_{h,in}^{opt} = T_{h,in}^{GP}$
23:             Set $\gamma^{opt} = \gamma^{GP}$
24:             Update $T_i(t)$, $T_{h,i}(t)$ and $T_{c,i}(t)$ using $\mathbf{x^{opt}} = [\gamma^{opt}, V_c^{opt}, T_{h,in}^{opt}]$
25:             Set $t = t + 1$, **go to line 4** (i.e. beginning of for loop) and *continue* loop.
26:         **end if**
27:     **end if**
28:     Call bayesopt package to determine $\mathbf{x^{opt}(t)} = [\gamma^{opt}(t), V_c^{opt}(t), T_{h,in}(t)]$ that maximizes the objective function in equation 21.
29:     Update $T_i(t)$, $T_{h,i}(t)$ and $T_{c,i}(t)$ using $\mathbf{x^{opt}} = [\gamma^{opt}, V_c^{opt}, T_{h,in}^{opt}]$.
30:     **if** $Q_{rec} > Q_d$ **then**
31:         Concatenate $\mathbf{z_{II,train}} \leftarrow [\mathbf{z_{II,train}}, \mathbf{z_{II}}]$
32:         Concatenate $\mathbf{x_{II,train}} \leftarrow [\mathbf{x_{II,train}}, \dot{V}_c^{opt}]$
33:     **else**
34:         Concatenate $\mathbf{z_{I,train}} \leftarrow [\mathbf{z_{I,train}}, \mathbf{z_I}]$
35:         Concatenate $\mathbf{x_{I,train}} \leftarrow [\mathbf{x_{I,train}}, \gamma^{opt}, T_{h,in}^{opt}, \dot{V}_c^{opt}]$
36:     **end if**
37: **end for**

---

9. Acknowledgments

References

[1] Building energy databook, last Accessed: 2017-11-23.
  URL https://openei.org/doe-opendata/dataset/buildings-energy-data-book
[2] Iec smart grid standardization roadmap, last Accessed: 2017-11-23.
  URL http://www.iec.ch/smartgrid/downloads/sg3$_r$oadmap.pdf
[3] E. Mocanu, P. H. Nguyen, M. Gibescu, W. L. Kling, Deep learning for estimating building energy consumption, Sustainable Energy, Grids and Networks 6 (2016) 91–99.
[4] L. Friedrich, A. Afshari, Short-term Forecasting of the Abu Dhabi Electricity Load Using Multiple Weather Variables, Energy Procedia 75 (2015) 3014–3026. doi:10.1016/j.egypro.2015.07.616.
  URL http://linkinghub.elsevier.com/retrieve/pii/S1876610215013843
[5] A. Dedinec, S. Filiposka, A. Dedinec, L. Kocarev, Deep belief network based electricity load forecasting: An analysis of macedonian case, Energy.
[6] E. A. Bakirtzis, C. K. Simoglou, P. N. Biskas, D. P. Labridis, A. G. Bakirtzis, Comparison of advanced power system operations models for large-scale renewable integration, Electric Power Systems Research 128 (2015) 90–99.
[7] D. Kolokotsa, The role of smart grids in the building sector, Energy and Buildings 116 (2016) 703–708. doi:10.1016/j.enbuild.2015.12.033.
  URL http://linkinghub.elsevier.com/retrieve/pii/S0378778815005447
[8] US Department of Energy, EnergyPlus Documentation (2013).
[9] H. S. Rallapalli, A comparison of energyplus and equest whole building energy simulation results for a medium sized office building, Ph.D. thesis, Arizona State University (2010).
[10] A. Foucquier, S. Robert, F. Suard, L. Stphan, A. Jay, State of the art in building modelling and energy performances prediction: A review, Renewable and Sustainable Energy Reviews 23 (2013) 272–288.
[11] N. Fumo, M. Rafe Biswas, Regression analysis for prediction of residential energy consumption, Renewable and Sustainable Energy Reviews 47 (2015) 332–343.
[12] W. Charytoniuk, M. S. Chen, P. Van Olinda, Nonparametric regression based short-term load forecasting, IEEE Transactions on Power Systems 13 (3) (1998) 725–730.
  URL http://ieeexplore.ieee.org/xpls/abs$_a$ll.jsp?arnumber = 708572
[13] H.-x. Zhao, F. Magouls, A review on the prediction of building energy consumption, Renewable and Sustainable Energy Reviews 16 (6) (2012) 3586–3592.
[14] D. C. Park, M. A. El-Sharkawi, R. J. Marks, L. E. Atlas, M. J. Damborg, others, Electric load forecasting using an artificial neural network, Power Systems, IEEE Transactions on 6 (2) (1991) 442–449.
[15] C. Robinson, B. Dilkina, J. Hubbs, W. Zhang, S. Guhathakurta, M. A. Brown, R. M. Pendyala, Machine learning approaches for estimating commercial building energy consumption, Applied Energy 208 (Supplement C) (2017) 889 – 904.
[16] P. A. Gonzlez, J. M. Zamarreo, Prediction of hourly energy consumption in buildings based on a feedback artificial neural network, Energy and Buildings 37 (6) (2005) 595–601.
[17] V. M. Z. Yeonsook Heo, Gaussian process modeling for measurement and verification of building energy savings, Energy and Buidings (53) (2012) 7–18.
[18] X. L, T. Lu, C. J. Kibert, M. Viljanen, Modeling and forecasting energy consumption for heterogeneous buildings using a physicalstatistical approach, Applied Energy 144 (2015) 261–275.
[19] K. Yun, R. Luck, P. J. Mago, H. Cho, Building hourly thermal load prediction using an indexed ARX model, Energy and Buildings 54 (2012) 225–233.
[20] A. Rahman, A. D. Smith, Predicting fuel consumption for commercial buildings with machine learning algorithms, Energy and Buildings 152 (2017) 341 – 358.
[21] C. Fan, F. Xiao, Y. Zhao, A short-term building cooling load prediction method using deep learning algorithms, Applied Energy 195 (Supplement C) (2017) 222 – 233.
[22] A. Rahman, V. Srikumar, A. D. Smith, Predicting electricity consumption for commercial and residential buildings using deep recurrent neural networks, Applied Energy 212 (2018) 372–385.
[23] Y. B. I. Goodfellow, A. Courville, Deep learning, book in preparation for MIT Press, http://www.deeplearningbook.org, Last Accessed = 2017-11-23 (2016).
[24] K. Cho, B. Van Merrinboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078.
  URL http://arxiv.org/abs/1406.1078
[25] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, in: Advances in neural information processing systems, 2014, pp. 3104–3112.
[26] Z. C. Lipton, A critical review of recurrent neural networks for sequence learning, CoRR abs/1506.00019, http://arxiv.org/abs/1506.00019, Last Accessed = 2017-11-23.

[27] G. Angrisani, M. Canelli, C. Roselli, M. Sasso, Calibration and validation of a thermal energy storage model: influence on simulation results, Applied Thermal Engineering 67 (1-2) (2014) 190–200. doi:10.1016/j.applthermaleng.2014.03.012.

[28] Thermal stratification within the water tank, Renewable and Sustainable Energy Reviews 13 (5) (2009) 1014–1026.

[29] I. Dincer, On thermal energy storage systems and applications in buildings, Energy and buildings 34 (4) (2002) 377–388.

[30] Optimal deployment of thermal energy storage under diverse economic and climate conditions, Applied Energy 119 (2014) 488–496.

[31] S. M. Hasnain, Review on Sustainable Thermal Energy Storage Technologies, Energy Conversion and Management 39 (2) (1998) 1127–1138.

[32] A. Celador, M. Oriozola, J. Sala, Implications of the modelling of stratified hot water storage tanks in the simulation of chp plants, Energy Conversion and Management 52 (2011) 3018–3026.

[33] C. A. Cruickshank, Evulation of a Stratified Multi-Tank Thermal Storage for Solar Heating Applications, Ph.D. thesis, Queen's Unviersity, Kingston, Ontario, Canada (2009).

[34] A. D. Smith, P. J. Mago, N. Fumo, Benefits of thermal energy storage option combined with chp system for different commercial building types, Sustainable Energy Technologies and Assessments 1 (2013) 3–12.

[35] E. S. Barbieri, F. Melino, M. Morini, Influence of the thermal energy storage on the profitability of micro-chp systems for residential building applications, Applied Energy 97 (2012) 714–722.

[36] H. Ren, W. Gao, Y. Ruan, Optimal sizing for residential chp system, Applied Thermal Engineering 28 (5) (2008) 514–523.

[37] V. Pandiyarajan, M. C. Pandian, E. Malan, R. Velraj, R. Seeniraj, Experimental investigation on heat recovery from diesel engine exhaust using finned shell and tube heat exchanger and thermal storage system, Applied Energy 88 (1) (2011) 77–87.

[38] Nakahara, Water thermal storage tank- part 2-mixing model and storage estimation for temperature-stratified tanks (1988).

[39] Kleinbach, Performance study of one-dimensional models for stratified thermal storage tanks (1993).

[40] A. Rahman, A. D. Smith, N. Fumo, Performance modeling and parametric study of a stratified water thermal storage tank, Applied Thermal Engineering 100 (2016) 668–679.

[41] A. Rahman, N. Fumo, A. D. Smith, Simplified modeling of thermal storage tank for distributed energy heat recovery applications, in: ASME 2015 9th International Conference on Energy Sustainability.

[42] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Computation 9 (1997) 1735–1780.

[43] C. E. Rasmussen, r. o. K. I. Williams, Gaussian processes for machine learning, MIT Press, Cambridge, Mass., 2006.
URL http://www.books24x7.com/marc.asp?bookid=12939

[44] J. Bergstra, D. Yamins, D. D. Cox, Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms, in: Proceedings of the 12th Python in Science Conference, 2013, pp. 13–20.

[45] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins, D. D. Cox, Hyperopt: a python library for model selection and hyperparameter optimization, Computational Science & Discovery 8 (1) (2015) 014008.

[46] F. Chollet, keras, https://github.com/fchollet/keras (2015).

[47] W. Logie, E. Frank, M. Rommel, Investigation of Immersed Coil Heat Exchangers in regard to Heat Transfer and Storage Stratification, in: EuroSun, 2010, pp. 1–10.

[48] SolarBayer GmbH, Stratification Buffer Tank SPS, uRL = www.solarbayer.com, Last accessed = 2018-06-2018.

[49] Hydroflex systems inc., water storage systems, last Accessed: 2018-01-11.
URL http://www.iec.ch/smartgrid/downloads/sg3$_r$oadmap.pdf

[50] F. Castiglia, P. Chiovaro, M. Ciofalo, M. D. Liberto, P. Maio, I. D. Piazza, M. Giardina, F. Mascari, G. Morana, G. Vella, Modelling flow and heat transfer in helically coiled pipes . Part 3 : Assessment of turbulence models , parametrical study and proposed correlations for fully turbulent flow in the case of zero pitch, Tech. rep., Agenzia Nazionale per le Nuove Tecnologie l'Energie e lo Sviluppo Economico Sostenibile (2010).

[51] L. Cattani, Numerical investigation of the convective heat transfer enhancement in coiled tubes, in: 2012 COMSOL Conference in Milan, 2012.

[52] M. Ali, Free convection heat transfer from the outer surface of vertically oriented helical coils in glycerol-water solution 40 (8) (2003) 615–620.

[53] J. Snoek, H. Larochelle, R. P. Adams, Practical bayesian optimization of machine learning algorithms, in: Advances in neural information processing systems, 2012, pp. 2951–2959.

[54] E. Brochu, V. M. Cora, N. de Freitas, A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, CoRR abs/1012.2599. arXiv:1012.2599.
URL http://arxiv.org/abs/1012.2599