Day-ahead Scheduling of Thermal Storage Systems Using Bayesian Neural Networks

Alexandre Capone * **Conrad Helminger** * **Sandra Hirche** *

* *Chair of Information-oriented Control (ITR), Department of Electrical and Computer Engineering, Technical University of Munich, Germany (e-mail: {alexandre.capone, conrad.helminger, hirche}@tum.de).*

Abstract: The increased need for energy efficiency in buildings requires sophisticated scheduling strategies. A considerable challenge when developing such strategies is to address the stochasticity of demand appropriately. In this paper, we propose a day-ahead scheduling technique for energy storage systems with heat pumps and backup resistance heaters under uncertain heat demand, which aims to minimize electricity costs as well as power grid congestion. We employ a Bayesian neural network to model the stochastic consumer demand, which takes historical measurements as training inputs, and is able to model complex stochastic patterns. The model is then employed to generate sample demands, which are used to approximate the expected costs. The minimization of the resulting cost function corresponds to a stochastic optimal control problem with quadratic costs and mixed integer constraints. In a numerical simulation of a single-family building, the proposed approach is shown to perform better than a standard neural network based scheduling scheme.

Keywords: Load forecasting, optimal control, stochastic modelling, machine learning, district heating, Non-Gaussian processes

1. INTRODUCTION

Improving the performance of building heat, ventilation and air-conditioning (HVAC) systems is essential to reduce overall energy consumption in buildings, as these correspond to approximately 40% of total building energy demand (Pérez-Lombard et al., 2008). Moreover, building HVAC systems have shown considerable potential to increase grid flexibility and reduce peak system demand (Kramer et al., 2016; Hao et al., 2014; Zheng et al., 2015; Aduda et al., 2016; Zhang et al., 2017). In order to harness the full potential of HVAC systems, day ahead scheduling strategies are of crucial importance. Such strategies, which in essence correspond to the solution of an optimal control problem, are widely employed, most notably in model predictive control methods (Lautenschlager and Lichtenberg, 2016; Afram and Janabi-Sharifi, 2014b; Arnold and Andersson, 2011) and energy market pricing approaches (Wittmann et al., 2011; Zamani et al., 2016). However, while the physical properties of HVAC systems are well understood (Afram and Janabi-Sharifi, 2014a), forecasting loads is considerably more difficult due to the markedly uncertain nature of energy demand. Common demand forecasting tools include first-principles models, statistical models that correlate energy consumption and system variables, and neural networks (NNs) (Idowu et al., 2016; Zhao and Magoulès, 2012; Ahmad et al., 2014). However, these methods neglect complex stochastic behaviors, e.g. heteroscedasticity due to varying uncertainty over time or multimodality due to building occupancy patterns, which are particularly relevant when considering single buildings instead of aggregate loads.

Recently, machine learning techniques have been increasingly employed to tackle control tasks, particularly in settings where uncertainty needs to be quantified and taken into account in the control design. Examples include Gaussian processes (Capone and Hirche, 2019; Beckers et al., 2019) and Bayesian neural networks (BNNs) (Depeweg et al., 2017). These approaches provide high model flexibility and have good generalization properties. So far, Gaussian processes have been applied in different building control settings (Nghiem and Jones, 2017; Jain et al., 2018; Yan et al., 2017). However, Bayesian neural networks have been employed comparatively less often (Chonan, 1996), despite being able to model more complex distributions (Depeweg et al., 2016).

In this paper, we present a day-ahead scheduling scheme for thermal storage systems, which takes complex stochastic patterns in consumer behavior into account. We employ a Bayesian neural network to compute stochastic demand forecasts, which are then employed to approximate the expected cost over a given time horizon. The resulting problem is a mixed-integer quadratic optimization problem, for which established solution tools exist.

The remainder of this paper is structured as follows. Section 2 describes the general problem setting. In Section 3, we introduce Bayesian neural networks, and discuss how they are employed to forecast demand. In Section 4, we then describe the stochastic optimal control formulation, which we use to compute the day-ahead scheduling strategy. A numerical illustration is presented in Section 5, after which we conclude with some final remarks, in Section 6.

2. PROBLEM SETTING

We consider a heat pump system with thermal storage and a low-efficiency backup resistance heater, which is described by the energy conservation equation

$$l_{p,t} + l_{s,t} + l_{b,t} = l_{c,t}, \tag{1}$$

where t denotes the time step, $l_{c,t} \in \mathbb{R}_+$ is the consumer load, $l_{p,t} \in \mathbb{R}_+$ is the heat pump power output, $l_s \in \mathbb{R}$ is the charge and discharge rate of the thermal storage, and $l_{b,t} \in \mathbb{R}_+$ is the power output of the electric backup heater. Here $\mathbb{R}_+$ denotes the positive real numbers. The heat demand $l_{c,t}$ is a time-dependent stochastic variable and is sampled according to an unknown probability distribution $p(l_{c,t})$. The heat pump power output $l_{p,t}$, the rate $l_{s,t}$, and the backup heater power output $l_{b,t}$ are controllable system inputs, and are subject to the power constraints

$$0 \leq l_{p,t} \leq L_p \tag{2}$$
$$-L_s \leq l_{s,t} \leq L_s \tag{3}$$
$$0 \leq l_{b,t} \leq L_b, \tag{4}$$

where $L_p \in \mathbb{R}_+$, $L_s \in \mathbb{R}_+$, and $L_b \in \mathbb{R}_+$ are the maximum output power of the heat pump, thermal storage and backup heater, respectively. The thermal storage dynamics are given by the difference eqution

$$\kappa_{t+1} = \kappa_t + l_{c,t}\Delta t, \tag{5}$$

where κ_t is the level of charge of the thermal storage at time t. Moreover, the thermal storage is subject to the capacity constraint

$$0 \leq \kappa_t \leq K, \quad \forall\, t \in \{0, \cdots, N\}, \tag{6}$$

where $K \in \mathbb{R}_+$ is the thermal storage capacity. Note that even though this model is simple, the techniques presented in this paper are applicable to more detailed representations.

The role of the heat pump and backup heater is twofold. They are employed to charge the thermal storage when necessary, while simultaneously guaranteeing that demand is met by covering demand that is not covered by the thermal storage. As such, they are not controlled directly, but are rather only employed when the thermal storage output is not enough to satisfy demand. Moreover, we assume that the heat pump has higher efficiency than the backup resistance heater. Hence, the backup heater is only employed to satisfy peak demand when the combination of heat pump and storage output are not sufficient. This is guaranteed by introducing the nonlinear constraint

$$(l_{p,t} - L_p) l_{b,t} = 0, \tag{7}$$

which enables the backup heater output $l_{b,t}$ to be zero only when the heat pump output $l_{p,t}$ is equal to its maximum output L_p. We assume that the thermal storage charge and discharge rate $l_{s,t}$ is not subject to energy losses, and is computed as

$$l_{s,t} = \begin{cases} l_{c,t} - L_p - L_b & \text{if } \rho_t \leq l_{c,t} - L_p - L_b \\ l_{c,t} & \text{if } \rho_t \geq l_{c,t} \\ \rho_t & \text{otherwise} \end{cases} \tag{8}$$

where ρ_t is given by the bounded feedback control law

$$\rho_t = \begin{cases} -L_s & \text{if } \frac{\tilde{\kappa}_t - \kappa_t}{\Delta t} \leq -L_s \\ L_s & \text{if } \frac{\tilde{\kappa}_t - \kappa_t}{\Delta t} \geq L_s \\ \frac{\tilde{\kappa}_t - \kappa_t}{\Delta t} & \text{otherwise} \end{cases} \tag{9}$$

Equation (9) aims to steer the thermal storage charge level κ_t towards a desired level $\tilde{\kappa}_t$ at every time step t, whereas (8) ensures that the power constraints (2)-(4) are satisfied. Note that $\tilde{\kappa}_t$ is a reference input, and as such is not subject to constraints. Moreover, the following holds:

Proposition 1. Let $\tilde{\kappa}_t$, $t \in \{1, \cdots, N\}$ be the desired charge level, and assume $0 \leq \tilde{\kappa}_t \leq K$ holds for all t. Moreover, let κ_t be given by (5), (8) and (9), and assume κ_0 satisfies $0 \leq \kappa_0 \leq K$. Then constraint (6) is satisfied by all κ_t, $t \in \{1, \cdots, N\}$.

Proof. Assume the contrary is true. Then there exists a κ_t, such that either $0 > \kappa_t$ or $\kappa_t > K$ holds. Consider the case where $0 > \kappa_t$ holds. Without loss of generality, we assume $0 \leq \kappa_{t-1} \leq K$. Then, due to (6), this implies $(0 - \kappa_t)/\Delta_t > l_{s,t} = (\tilde{\kappa}_t - \kappa_t)/\Delta_t$, which in turn is equivalent to $0 > \tilde{\kappa}_t$, which is a contradiction, i.e., $0 \leq \kappa_t$ holds. A similar argument leads to $\tilde{\kappa}_t > K$, which implies the desired statement.

Hence, by choosing the desired charge levels such that $0 \leq \tilde{\kappa}_t \leq K$ holds, (6) is satisfied automatically, i.e., the capacity constraints do not need to be enforced directly.

Let $\tilde{\boldsymbol{\kappa}}_s := (\tilde{\kappa}_1, \cdots, \tilde{\kappa}_N)$ denote the subsumption of the desired storage charge levels. Furthermore, let $\nu_c, \nu_b \in \mathbb{R}_+$ respectively denote the heat pump coefficient of performance and the thermal efficiency of the backup heater. We aim to devise a charging strategy that minimizes electricity costs, while simultaneously reducing electricity consumption peaks. The latter requirement is of particular significance to avoid congestion and guarantee power system stability. This is achieved by minimizing a cost function of the form

$$C(\tilde{\boldsymbol{\kappa}}_s) = \mathbb{E}\left[\sum_{t=0}^{N} c_t \left(\frac{l_{p,t}}{\nu_c} + \frac{l_{b,t}}{\nu_b}\right) + \eta \left(\frac{l_{p,t}}{\nu_c} + \frac{l_{b,t}}{\nu_b}\right)^2\right], \tag{10}$$

where c_t is the electricity price at time t, which we assume to know. The scalar $\eta > 0$ is a design parameter that penalizes electricity consumption peaks.

In order to obtain a charging strategy that minimizes (10) efficiently, we require an accurate load forecast. To this end, based on Depeweg et al. (2016), we approximate the heat demand as

$$l_{c,t} \approx f(t, z_t, \boldsymbol{W}), \tag{11}$$

where the function $f : \mathbb{R} \times \mathbb{R} \times \mathbb{R}^{d_w}$ is a NN with weights $\boldsymbol{W} \in \mathcal{W} \subseteq \mathbb{R}^{d_w}$, and $z_t \sim \mathcal{N}(0, 1)$ is an unobservable stochastic variable. In order to compute the weights $\boldsymbol{W}$, we assume to have N historical measurements consisting of time labels $\boldsymbol{t} = (t^{(1)}, \cdots, t^{(N)})$ and corresponding load measurements $\boldsymbol{Y} = \left(y^{(1)}, \cdots, y^{(N)}\right)$, where $y^{(i)} = l_{c,t}^{(i)} + \varepsilon^{(i)}$ are noisy demand measurements with $\varepsilon^{(i)} \sim \mathcal{N}(0, \sigma)$ and $\sigma > 0$.

3. BAYESIAN NEURAL NETWORK MODEL

We now introduce Bayesian neural networks (BNNs) and discuss how they are employed to model loads. We begin by reviewing classical feedforward artificial neural networks (NNs), after which BNNs are discussed.

3.1 Standard Artificial Neural Networks

Given a time step $t \in \mathbb{R}$, an artificial feedforward NN computes an output $l \in \mathbb{R}$, whose corresponding input-

output relationship is given by

$$l = o_L \tag{12}$$

where L denotes the number of layers in the NN. The variable o_L denotes the output of the L-th layer, which is computed recursively as

$$o_{i,j} = \phi_{i,j}\left(\sum_{k=1}^{U_{i-1}} w_{k,j}^{i-1}\, o_{i-1,k}\right), \tag{13}$$

$$i = 1, \cdots, L, \quad j = 1, \cdots, U_i \tag{14}$$

$$o_0 := t, \tag{15}$$

where L denotes the number of layers in the network, $o_{i,j}$ is the i-th component of the output of the j-th NN layer, U_i denotes the number of neurons in the i-th layer, and $w_{i,j}^l$ is a weight corresponding to the $l-1$-th layer.
Typically, artificial NNs are trained by adjusting the weights using gradient descent, such that the output fits the training data or minimizes an ojective function tailored to the problem (Werbos, 1974).

3.2 Bayesian Neural Networks

A BNN extends the notion of NNs to accommodate stochastic variables. This is achieved by sampling the weights $w_{i,j}^l$ from a probability distribution, i.e., $w_{i,j}^l \sim p(w_{i,j}^l)$. We treat the weights

$$\boldsymbol{W} = \left\{w_{i,j}^l\right\}_{l\in\{1,\cdots,L\}, i\in\{1,\cdots,U_l\}, j\in\{1,\cdots,U_{l-1}+1\}}, \tag{16}$$

as independent and identically distributed stochastic variables, and assume that every weight has a normal prior distribution with mean zero and variance λ, i.e.,

$$p(\boldsymbol{W}) = \prod_{l=1}^{L}\prod_{i=1}^{U_l}\prod_{j=1}^{U_{l-1}+1} \mathcal{N}(0,\lambda). \tag{17}$$

After having collected N training samples $\boldsymbol{t}$ and $\boldsymbol{Y}$, the posterior distribution of the weights $\boldsymbol{W}$ and stochastic features $\boldsymbol{z} = \left(z^{(1)}, \cdots, z^{(N)}\right)$ is computed using Bayes' rule as

$$p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y}) = \frac{p(\boldsymbol{Y} \,|\, \boldsymbol{W}, \boldsymbol{z}, \boldsymbol{t})\, p(\boldsymbol{W})\, p(\boldsymbol{z})}{p(\boldsymbol{Y} \,|\, \boldsymbol{t})}. \tag{18}$$

Hence, the likelihood of obtaining the output data $\boldsymbol{Y}$ is given by

$$p(\boldsymbol{Y} \,|\, \boldsymbol{W}, \boldsymbol{z}, \boldsymbol{t}) = \prod_{n=1}^{N} p_{\mathcal{N}}\big(f(t^{(n)}, z^{(n)}; \boldsymbol{W}), \boldsymbol{\sigma}\big). \tag{19}$$

$$p(y^* \,|\, t^*, \boldsymbol{t}, \boldsymbol{Y}) = \iiint_{\boldsymbol{W}\ \boldsymbol{z}\ z^*} p_{\mathcal{N}}\big(y^* \,|\, \tilde{\boldsymbol{f}}(t^*, z^*; \boldsymbol{W}), \boldsymbol{\sigma}\big)\, p_{\mathcal{N}}(z^* \,|\, 0, \gamma) \; dz^*\, p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y})\, d\boldsymbol{z}\, d\boldsymbol{W}, \tag{20}$$

where

$$p_{\mathcal{N}}(a|m, v) := \frac{1}{\sqrt{2\pi v^2}} e^{-\frac{(a-m)^2}{2s^2}} \tag{21}$$

denotes the normal probability distribution with mean m and variance v evaluated at a.

3.3 Weight Training

In general, optimizing (20) with respect to the weights is intractable. Hence, based on Depeweg et al. (2016), we approximate the exact posterior distribution $p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y})$ as

$$q(\boldsymbol{W}, \boldsymbol{z}) = \prod_{l=0}^{L}\prod_{i=1}^{U_l}\prod_{j=1}^{U_{l+1}-1} p_{\mathcal{N}}\left(w_{i,j}^l \,|\, m_{w_{i,j}^l}, v_{w_{i,j}^l}\right) \times \prod_{n=1}^{N} p_{\mathcal{N}}\left(z^{(n)} \,|\, m_{z^{(n)}}, v_{z^{(n)}}\right). \tag{22}$$

In order to compute the means m and the variances v of the Gaussian factors of $q(\cdot)$, we minimize the α-divergence between $p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y})$ and $q(\boldsymbol{W}, \boldsymbol{z})$ (Minka et al., 2005; Depeweg et al., 2016)

$$\mathrm{D}_\alpha\left[p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y}) \| q(\boldsymbol{W}, \boldsymbol{z})\right] = \frac{1}{\alpha(\alpha-1)}\left(1 - \int p(\boldsymbol{W}, \boldsymbol{z} \,|\, \boldsymbol{t}, \boldsymbol{Y})^\alpha q(\boldsymbol{W}, \boldsymbol{z})^{(1-\alpha)}\right) d\boldsymbol{W}\, d\boldsymbol{z}, \tag{23}$$

where $\alpha \in \mathbb{R}_+$. The α-divergence measures the discrepancy between the true probability $p(\boldsymbol{W}, \boldsymbol{z}\,|, \boldsymbol{t}, \boldsymbol{Y})$ and the approximation $q(\boldsymbol{W}, \boldsymbol{z})$, and corresponds to a generalization of Kullback-Leibler divergence (Van Erven and Harremos, 2014). Since minimizing (23) is generally intractable, we employ black-box α-divergence minimization (Hernández-Lobato et al., 2016), which aims to approximate the minimum of (23) by minimizing

$$E_\alpha(q) = -\log Z_q - \frac{1}{\alpha} \times \sum_{n=1}^{N} \log \mathbb{E}_q\left[\left(\frac{p_{\mathcal{N}}\big(y^{(n)} \,|\, f(t^{(n)}, z^{(n)}; \boldsymbol{W}), \boldsymbol{\sigma}\big)}{g(\boldsymbol{W}) h(z^{(n)})}\right)^\alpha\right], \tag{24}$$

where

$$\log Z_q = \sum_{l=0}^{L}\sum_{i=1}^{U_l}\sum_{j=1}^{U_{l+1}-1}\left(\frac{1}{2}\log\left(2\pi v_{w_{i,j}^l}\right) + \frac{m_{w_{i,j}^l}^2}{v_{w_{i,j}^l}}\right) + \sum_{n=1}^{N}\left(\frac{1}{2}\log\left(2\pi v_{z^{(n)}}\right) + \frac{m_{z^{(n)}}^2}{v_{z^{(n)}}}\right), \tag{25}$$

$$g(\boldsymbol{W}) = \exp\left(\sum_{l=0}^{L}\sum_{i=1}^{U_l}\sum_{j=1}^{U_{l+1}-1}\frac{1}{N} \times \left(\frac{\lambda v_{w_{i,j}^l}}{\lambda - v_{w_{i,j}^l}}(w_{i,j}^l)^2 + \frac{m_{w_{i,j}^l}}{v_{w_{i,j}^l}} w_{i,j}^l\right)\right), \tag{26}$$

$$h(z^{(n)}) = \exp\left(\frac{\gamma v_{z^{(n)}}}{\gamma - v_{z^{(n)}}}(z^{(n)})^2 + \frac{m_{z^{(n)}}}{v_{z^{(n)}}} z^{(n)}\right), \tag{27}$$

and $\mathbb{E}_q[\cdot]$ indicates that the expected value is computed according to the surrogate posterior distribution (22). Furthermore, the expectation term in (24) is approximated by Monte Carlo integration

$$\mathbb{E}_q\left[\left(\frac{\mathcal{N}\big(y^{(n)} \,|\, f(t^{(n)}, z^{(n)}; \boldsymbol{W}), \boldsymbol{\sigma}\big)}{g(\boldsymbol{W}) h(z_n)}\right)^\alpha\right] \approx \frac{1}{K}\sum_{k=1}^{K}\left(\frac{\mathcal{N}\big(y^{(n)} \,|\, f(t^{(n)}, z^{(n,k)}; \boldsymbol{W}^{(k)}), \boldsymbol{\sigma}\big)}{g(\boldsymbol{W}^{(n,k)}) h(z^{(n,k)})}\right)^\alpha. \tag{28}$$

This approach scales to large datasets, as (24) can be minimized by employing gradient-descent Hernández-Lobato et al. (2016).

3.4 Load Forecasting with BNNs

After computing the weights $\boldsymbol{W}$, as described in Section 3.3, we are able to forecast consumer loads $\delta_{c,t}^{(i)}$ by sampling $\zeta^{(i)} \sim \mathcal{N}(0,1)$ and computing

$$\delta_{c,t}^{(i)} = f(t, \zeta^{(i)}, \boldsymbol{W}), \quad i \in \mathbb{N}. \tag{29}$$

Here we use $\delta_{c,t}^{(i)}$ and $\zeta^{(i)}$ to denote samples that are drawn after training the NN weights $\boldsymbol{W}$, and should not be confused with the training samples $l_{c,t}^{(i)}$, $z^{(i)}$.

4. STOCHASTIC OPTIMAL CONTROL

We now provide a stochastic optimal control formulation that approximates the minimizer of (10). By computing M sample consumer load forecasts as described in Section 3.3, we are able to approximate the expected cost (10) using Monte Carlo integration. By additionally taking the system constraints (1)-(9) into account, we obtain the stochastic optimization problem

$$\begin{aligned}
\min_{\boldsymbol{\kappa}_s} \ & \sum_{i=1}^{M}\sum_{t=0}^{N} c_t \left(\frac{l_{p,t}^{(i)}}{\nu_c} + \frac{l_{b,t}^{(i)}}{\nu_b} \right) + \eta \left(\frac{l_{p,t}^{(i)}}{\nu_c} + \frac{l_{b,t}^{(i)}}{\nu_b} \right)^2, \\
\text{s.t.} \quad & \forall\, t \in \{1, \cdots, N\}, \quad \forall\, i \in \{1, \cdots, M\}, \\
& \delta_{c,t}^{(i)} + l_{p,t}^{(i)} + l_{s,t}^{(i)} + l_{b,t}^{(i)} = 0, \\
& (l_{p,t}^{(i)} - L_p) l_{b,t}^{(i)} = 0, \\
& \kappa_{t+1}^{(i)} = \kappa_t^{(i)} + l_{s,t}^{(i)} \Delta t, \\
& l_{s,t}^{(i)} = \begin{cases} \delta_{c,t}^{(i)} - L_p - L_b & \text{if } \rho_t^{(i)} \le \delta_{c,t}^{(i)} - L_p - L_b \\ \delta_{c,t}^{(i)} & \text{if } \rho_t^{(i)} \ge \delta_{c,t}^{(i)} \\ \rho_t^{(i)} & \text{otherwise} \end{cases} \\
& \rho_t^{(i)} = \begin{cases} -L_s & \text{if } \frac{\tilde{\kappa}_t - \kappa_t^{(i)}}{\Delta t} \le -L_s \\ L_s & \text{if } \frac{\tilde{\kappa}_t - \kappa_t^{(i)}}{\Delta t} \ge L_s \\ \frac{\tilde{\kappa}_t - \kappa_t^{(i)}}{\Delta t} & \text{otherwise} \end{cases} \\
& 0 \le \ l_{p,t}^{(i)} \le L_p \\
& 0 \le \ l_{b,t}^{(i)} \le L_b, \\
& 0 \le \tilde{\kappa}_t^{(i)} \le K, \quad \forall\, \tau \in \{0, \cdots, N\} \\
& \delta_{c,t}^{(i)} = f(t, \zeta^{(i)}, \boldsymbol{W}),
\end{aligned} \tag{30}$$

where $\zeta^{(i)} \sim \mathcal{N}(0,1)$, $i \in \{1, \cdots, M\}$ are sampled before the optimization and are treated as fixed parameters. This corresponds to a mixed-integer quadratic program, for which established solution methods exist Misener and Floudas (2013).

5. NUMERICAL EXPERIMENTS

We employ the techniques detailed in Sections 3 to predict the heat demand of a single family home on a weekend day and on a weekday, and design a corresponding thermal storage charging strategy.

The maximum power output of the heat pump, thermal storage, and backup heater are $L_p = 20\,\text{kW}$, $L_s = 15\,\text{kW}$, and $L_b = 50\,\text{kW}$, respectively. The thermal storage has a capacity of $K = 10\,\text{kWh}$. The heat pump has a coefficient

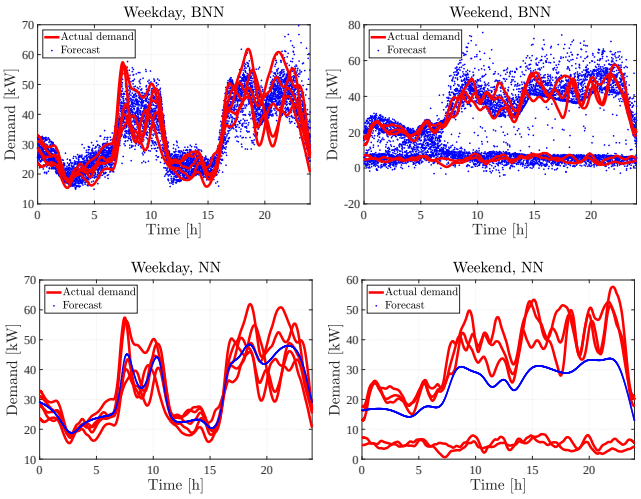

Fig. 1. Forecasts superimposed over actual consumer demand curves.

of performance of $\nu_c = 4$, whereas the backup resistance heater has an efficiency of $\nu_b = 0.95$. The consumer demand is generated by a mathematical model of the form

$$l_{s,t} = \beta_t + \sum_{k=1}^{N_a} \epsilon_{k,t} \boldsymbol{a}_{k,t}^{\mathrm{T}} (\boldsymbol{R}_{k,t} + \boldsymbol{I})(\boldsymbol{\psi}_{k,t} + \boldsymbol{r}_{k,t}), \tag{31}$$

where $\beta_{k,t} \in \mathbb{R}_+$, $\boldsymbol{a}_{k,t} \in \mathbb{R}_+^{d_p}$ and $\boldsymbol{\psi}_{k,t} \in \mathbb{R}_+^{d_p}$ are fixed parameters, $\epsilon_{k,t}$ are binary random variables, and the entries of $\boldsymbol{R}_{k,t}$ and $\boldsymbol{r}_{k,t}$ are sampled from a uniform distribution between 0 and 1. The parameters are chosen such that the demand (31) for weekdays generally feature peak consumptions during the morning and the evening hours. Moreover, times of peak correspond to a higher demand variance, whereas less variance is present at times of low demand. The binary random variables $\epsilon_{k,t}$ are chosen such that weekend days are bimodal, which reflects the possibility of consumers not being at home. Hence, the demand probability distribution exhibits heteroscedasticity and multimodality.

We perform load forecasts using a Bayesian NN with 25×25 hidden neurons, for which we employ 100 demand curves as training data. In order to emphasize the benefits of the proposed approach, we compare both the demand forecast as well as the scheduling results to those obtained with a standard NN with an identical architecture. The BNN and NN were trained by applying Adam (Kingma and Ba, 2014) for $8 \cdot 10^3$ and $20 \cdot 10^3$ epochs, respectively. The forecast results obtained with a BNN and a standard NN are respectively shown in the top and bottom plots in Figure 1. The left-hand side plots correspond to demand on a weekday, whereas the plots on the right-hand side are forecasts on a weekend day. In order to illustrate the accuracy of the predictions, we superimpose over the forecasts six demand curves that were not used for training. Both on weekdays and on weekend days, the BNN is able to capture complex stochastic phenomena, such as bimodality and heteroscedasticity. However, on the weekdays, the BNN data fit is tighter and generates less outliers than on the weekend day, which is due to the less complex nature of the stochastic variables. Moreover, on the weekend day, the two modes are relatively close to each other between 4 o'clock and 5 o'clock, which

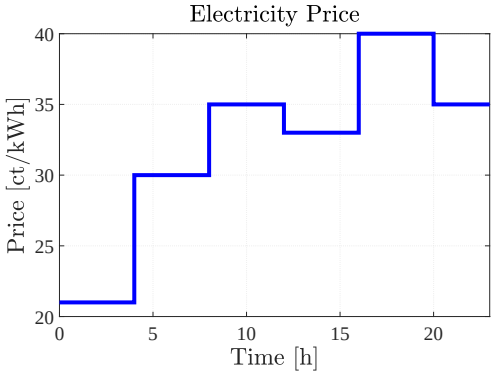

Fig. 2. Electricity price during simulation.

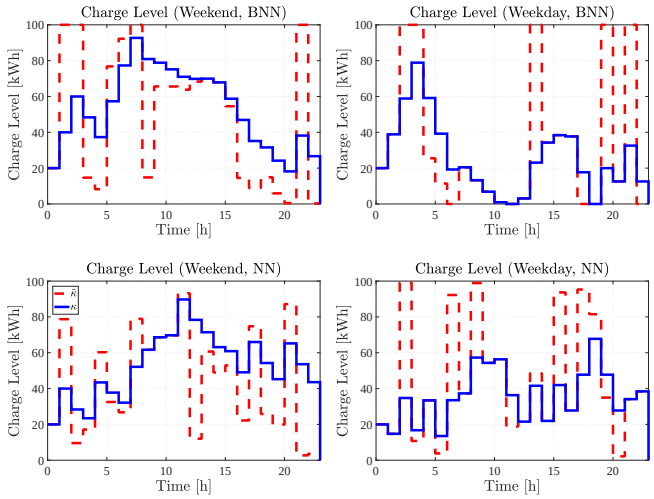

Fig. 3. Desired charge level and mean charge level of thermal storage obtained using BNN and standard NN.

causes the resulting samples to overlap slightly. However, the data fit follows the qualitative patterns of the test data. By contrast, the standard NN only approximates the mean of the data. Moreover, it assumes that the mean corresponds to the correct prediction with full confidence. This shortcoming is particularly accentuated on weekend days, where the predicted trajectory fails to cover any of the modes of the data.

5.1 Optimal control

We consider a setting where the electricity price changes every 4 hours, as given in Figure 2.

We employ a BNN with 25×25 hidden neurons, and use $M = 50$ Monte Carlo samples to approximate the expected cost (10). In order to illustrate the advantages of the BNN-based forecast, we compare the optimal control strategy to the one obtained using a standard NN prediction with the same network architecture.

We perform 60 simulations to approximate the true expected cost (10) for a weekend day and a weekday. The resulting desired charge levels $\tilde{\kappa}$ and the mean charge levels κ_t are given in Figure 3. The corresponding mean heat pump and backup heater power outputs are provided in Figures 4 and 5, respectively.

The BNN-based algorithm produces mean electricity costs of 7 Euro on the weekday, and 4 Euro on a weekend day.

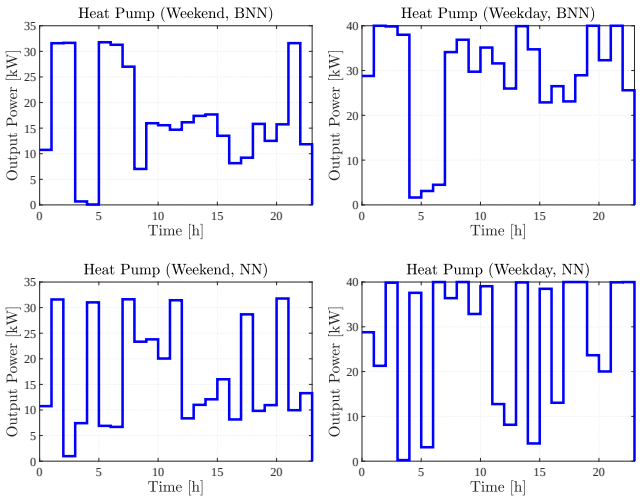

Fig. 4. Mean heat pump power output obtained using BNN and standard NN.

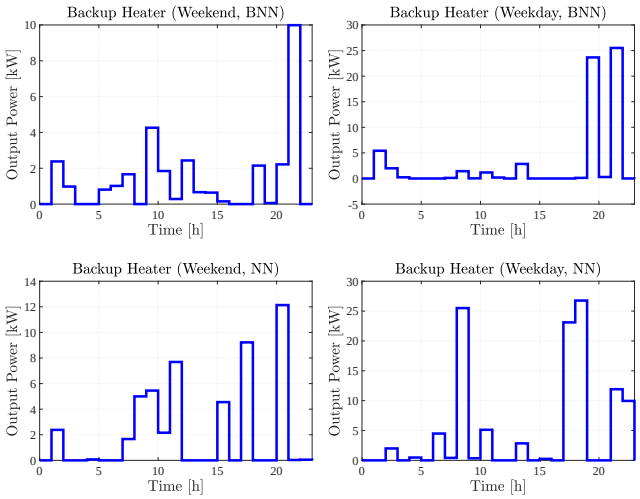

Fig. 5. Mean backup heater output obtained using BNN and standard NN.

The corresponding penalties for peak demand are $3 \cdot 10^5$ and $2 \cdot 10^5$, respectively. By contrast, when employing a standard NN, electricity costs of 10 Euro and 5 Euro are achieved respectively on weekdays and on weekends, with corresponding peak demand penalties of $7 \cdot 10^5$ and $3 \cdot 10^5$. Hence, the BNN-based algorithm performs better than the standard NN-based one, both with respect to electricity costs and peak shaving. A direct cause of this discrepancy lies in the usage of the backup heater, which is employed more frequently and at higher output power by the NN-based solution on both days. This happens because the BNN provides information about the demand uncertainty at every time step, which in turn makes the resutling scheduling strategy less susceptible to deviations from the mean. This contrasts with the NN-based strategy, which assumes that the mean is fully correct at every time step. Similarly, the BNN-based strategy aims to have a fully charged thermal storage before the final period of the simulation on both days, which is when the electricity price is highest. During the high price period, the storage is discharged, which helps alleviate the overall costs. This

differs from the NN-based strategy, which is less accentuated in this aspect.

6. CONCLUSION

We presented a Bayesian neural network-based day-ahed scheduling strategy for thermal storage systems with heat pump and backup resistance heater. The proposed technique leverages stochastic load forecasts to approximate the expected system cost function using Monte Carlo integration, and subsequently optimizes over the storage schedule. In numerical simulations, we have shown that the Bayesian neural network accurately approximates complex stochastic consumer behavior, which in turn leads to better performance in day-ahead scheduling compared to a standard neural network-based approach.

REFERENCES

Aduda, K., Labeodan, T., Zeiler, W., Boxem, G., and Zhao, Y. (2016). Demand side flexibility: Potentials and building performance implications. *Sustainable cities and society*, 22, 146–163.

Afram, A. and Janabi-Sharifi, F. (2014a). Review of modeling methods for hvac systems. *Applied Thermal Engineering*, 67(1-2), 507–519.

Afram, A. and Janabi-Sharifi, F. (2014b). Theory and applications of HVAC control systems–a review of model predictive control (MPC). *Building and Environment*, 72, 343–355.

Ahmad, A., Hassan, M., Abdullah, M., Rahman, H., Hussin, F., Abdullah, H., and Saidur, R. (2014). A review on applications of ANN and SVM for building electrical energy consumption forecasting. *Renewable and Sustainable Energy Reviews*, 33, 102–109.

Arnold, M. and Andersson, G. (2011). Model predictive control of energy storage including uncertain forecasts. In *17th Power Systems Computation Conference (PSCC 2011)*.

Beckers, T., Kulić, D., and Hirche, S. (2019). Stable Gaussian process based tracking control of euler–lagrange systems. *Automatica*, 103, 390 – 397.

Capone, A. and Hirche, S. (2019). Backstepping for partially unknown nonlinear systems using Gaussian processes. *IEEE Control Systems Letters*, 3, 416–421.

Chonan, Y. (1996). A bayesian nonlinear regression with multiple hyperparameters on the ashrae ii time-series data. *ASHRAE Transactions 1996*.

Depeweg, S., Hernández-Lobato, J.M., Doshi-Velez, F., and Udluft, S. (2016). Learning and policy search in stochastic dynamical systems with bayesian neural networks. *arXiv preprint arXiv:1605.07127*.

Depeweg, S., Hernández-Lobato, J.M., Doshi-Velez, F., and Udluft, S. (2017). Uncertainty decomposition in bayesian neural networks with latent variables. *arXiv preprint arXiv:1706.08495*.

Hao, H., Lin, Y., Kowli, A.S., Barooah, P., and Meyn, S. (2014). Ancillary service to the grid through control of fans in commercial building hvac systems. *IEEE Transactions on Smart Grid*, 5(4), 2066–2074.

Hernández-Lobato, J.M., Li, Y., Rowland, M., Hernández-Lobato, D., Bui, T., and Turner, R. (2016). Black-box α-divergence minimization.

Idowu, S., Saguna, S., Åhlund, C., and Schelén, O. (2016). Applied machine learning: Forecasting heat load in district heating system. *Energy and Buildings*, 133, 478–488.

Jain, A., Nghiem, T., Morari, M., and Mangharam, R. (2018). Learning and control using gaussian processes. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*, 140–149. IEEE.

Kingma, D.P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.

Kramer, M., Jambagi, A., and Cheng, V. (2016). A model predictive control approach for demand side management of residential power to heat technologies. In *IEEE International Energy Conference*, 1–6.

Lautenschlager, B. and Lichtenberg, G. (2016). Data-driven iterative learning for model predictive control of heating systems. *IFAC-PapersOnLine*, 49(13), 175–180.

Minka, T. et al. (2005). Divergence measures and message passing. Technical report, Microsoft Research.

Misener, R. and Floudas, C.A. (2013). Glomiqo: Global mixed-integer quadratic optimizer. *Journal of Global Optimization*, 57(1), 3–50.

Nghiem, T.X. and Jones, C.N. (2017). Data-driven demand response modeling and control of buildings with Gaussian processes. In *2017 American Control Conference (ACC)*, 2919–2924.

Pérez-Lombard, L., Ortiz, J., and Pout, C. (2008). A review on buildings energy consumption information. *Energy and buildings*, 40(3), 394–398.

Van Erven, T. and Harremos, P. (2014). Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7), 3797–3820.

Werbos, P.J. (1974). Beyond regression: New tools for predicting and analysis in the behavioral sciences.

Wittmann, M., Eck, M., Pitz-Paal, R., and Müller-Steinhagen, H. (2011). Methodology for optimized operation strategies of solar thermal power plants with integrated heat storage. *Solar Energy*, 85(4), 653 – 659. SolarPACES 2009.

Yan, B., Li, X., Shi, W., Zhang, X., and Malkawi, A. (2017). Forecasting building energy demand under uncertainty using gaussian process regression: Feature selection, baseline prediction, parametric analysis and a web-based tool. In *Proceedings of the 15th IBPSA Conference*.

Zamani, A.G., Zakariazadeh, A., and Jadid, S. (2016). Day-ahead resource scheduling of a renewable energy based virtual power plant. *Applied Energy*, 169, 324 – 340.

Zhang, S., Cheng, Y., Fang, Z., Huan, C., and Lin, Z. (2017). Optimization of room air temperature in stratum-ventilated rooms for both thermal comfort and energy saving. *Applied Energy*, 204, 420–431.

Zhao, H.x. and Magoulès, F. (2012). A review on the prediction of building energy consumption. *Renewable and Sustainable Energy Reviews*, 16(6), 3586–3592.

Zheng, M., Meinrenken, C.J., and Lackner, K.S. (2015). Smart households: Dispatch strategies and economic analysis of distributed energy storage for residential peak shaving. *Applied Energy*, 147, 246–257.