Received 4 April 2022; revised 19 May 2022; accepted 21 May 2022. Date of publication 27 May 2022; date of current version 7 June 2022. The review of this article was arranged by Associate Editor Xiaodong Liang.

Digital Object Identifier 10.1109/OJIA.2022.3178235

A Recommender System for Predictive Control of Heating Systems in Economic Demand Response Programs

DAVID TOQUICA [1], KODJO AGBOSSOU [1] (Senior Member, IEEE), ROLAND MALHAMÉ [2] (Fellow, IEEE), NILSON HENAO [1], SOUSSO KELOUWANI [3] (Senior Member, IEEE), AND MICHAËL FOURNIER [4]

[1]Department of Electrical and Computer Engineering, University of Quebec, Trois-Rivières, QC G8Z 4M3, Canada
[2]Department of Electrical Engineering, Polytechnique Montreal, Montreal, QC H3T 1J4, Canada
[3]Department of Mechanical Engineering, University of Quebec, Trois-Rivières, QC G8N 4M3, Canada
[4]Laboratoire des Technologies de l'érgie, Institut de Recherche Hydro-Québec, Shawinigan, QC G9N 7N5, Canada

CORRESPONDING AUTHOR: David Toquica (e-mail: david.camilo.toquica.cardenas@uqtr.ca)

ABSTRACT Flexibility from demand-side resources is increasingly required in modern power systems to maintain the dynamic balance between demand and supply. This flexibility comes from elastic users managing controllable loads. In this context, controlling Electric Space Heaters (ESHs) is of particular interest because it can leverage building inner thermal storage capacity to shift consumption while maintaining comfort conditions. Some economic Demand Response (DR) programs have considered exploiting EHSs flexibility potentials in recent years. However, these programs still struggle to engage customers due to the complexity of processing price signals for inexpert users. Therefore, it is necessary to develop automated tools for helping users to operate their loads. Accordingly, this paper presents a recommender system based on Gaussian processes to discover users' valuations of thermal comfort and perform the predictive control of their ESHs. The proposed method enables customers to participate in DR programs and impose their preferences through straightforward queries instead of directly changing control parameters. Validation results demonstrate that users maximize their utility by supplying noiseless and consistent data to the recommender system. Additionally, the suggested approach achieves a higher acceptance rate than other methods from the literature, such as persistency and support vector machines.

INDEX TERMS Heating systems, predictive control, preference learning, recommender system, utility maximization.

NOMENCLATURE

Acronyms

ESH	Electric space heater.
DR	Demand response.
GP	Gaussian process.
HMI	Human-machine interface.
MAP	Maximum a posteriori.
UX	User experience.
WTP	Willingness to pay.

Variables

c_t	Electricity cost at time t.
F	Scale factor of the thermal comfort valuation.
F_H	Set of n historical values of F.
$\mathbb{I}_{n \times n}$	Identity matrix of $n \times n$ dimension.
K	Covariance function.
P_t	Energy consumption of the ESH at time t.
P_{max}	Capacity of the ESH.
$P_{app,t}$	Consumption of household appliances at time t.
$P_{irr,t}$	Total solar irradiation at time t.
s	Consumption strategy.
T	Market period of t time slots.
u_t	User's utility at time t.
x_M	Set of average values of π and θ_{ext} for a market period.
X_H	Set of n historical average values of π and θ_{ext}.

δ_t	Form factor of the thermal comfort valuation at time t.
π_t	Electricity price at time t.
σ_H	Standard deviation of the components of X_H.
σ_M	Standard deviation of the estimation of F.
σ_n	Standard deviation of the noise in F_H.
θ_t	House indoor temperature at time t.
θ_{ref}	User's preferred temperature.
$\theta_{ext,t}$	Outdoor temperature at time t.
φ_t	Thermal comfort valuation at time t.

I. INTRODUCTION

The world is experiencing an energy and technology transition driven by environmental, social, and economic concerns. One of the goals of this transition is to decarbonize electric power generation by integrating renewable resources. However, it becomes challenging to keep the dynamic balance of the grid counting with these resources due to their intermittent and stochastic nature. Therefore, it is crucial to increase the flexibility of power systems to absorb imbalances and maintain reliable, resilient, and secure energy supply [1]. Flexibility can come from different sources like Demand Response (DR) programs, fast generation ramping, grid reconfigurations, and energy storage systems [2].

The interest in flexibility from DR programs is encouraged by the adoption of new smart technologies, especially in the residential sector [3]. This sector is also a huge energy consumer in many countries, and modifying its demand patterns can significantly improve the system operation. Accordingly, different DR programs have emerged to manage residential sector loads such as electric vehicles, Electric Water Heaters, Air-Conditioners, and Electric Space Heaters (ESHs) [4], [5]. Some of these programs have also incorporated energy storage systems to shift the demand without compromising the comfort [6]. In this regard, ESHs are advantageous since they can use building inner thermal energy storage to provide flexibility.

According to the International Energy Agency, ESHs can become a major source of flexibility in many regions [7]. Such opportunity has stimulated regulators to develop codes and standards for electrifying space heaters and accelerating the adoption of smart technologies. However, it is still necessary to engage customers and show them the economic advantages of providing flexibility with their ESHs. Some barriers to having more participants in DR programs are the lack of customers' knowledge, the technology cost, and the response fatigue [8]. In order to address these problems, it is suitable to automate the customers' decision process rather than take direct control of their loads. However, the automation is not simple since each customer has a different valuation of comfort and experiences various conditions.

A. RELATED WORKS

It is desirable to focus the DR programs on customers' preferences to avoid imposing conditions and to make the control actions more welcomed. In this context, the convenience of recommender systems for engaging residential customers into DR programs has been explored in the literature [18]. Indeed, recommendations can facilitate users' interactions with automated systems and help discover better consumption strategies. Previous findings on the application of recommender systems for residential energy management are listed in Fig. 1. Those applications highlight the challenge of collecting data from inconsistent user inputs even when they act rationally with stable and monotone preferences. This situation is due to noisy drivers like impatience, comfort bias, privacy concerns, and misperceptions. Accordingly, recommender systems have to build training datasets aiming to improve both the accuracy and serendipity of the suggestions [19].

For the specific case of ESHs, recommender systems have been used to aggregate energy consumption and infer thermal comfort perception [20]. A common technique for these recommenders is collaborative filtering that collects data from several users before suggesting control actions. These filters can identify average preferences and cluster customers, so they are suitable for office buildings or groups of residences [21]. However, these techniques are not aware of the specific context of each customer. Beyond that, they do not examine the effect of individual preferences on elastic consumption, which is crucial for customizing DR programs.

Aside from recommender systems, other strategies to include users' preferences into ESH control require querying customers directly about specific parameters. For active learning strategies, the control algorithm queries the customer when it faces unknown conditions and uses Bayesian updates to discover absolute preferences [22]. Then, customers must know their willingness to pay (WTP), parametrize their demand curves, and define temperature comfort limits explicitly for answering the queries and participating in DR programs. These scenarios are not practical to promote DR programs among inexpert customers. Moreover, direct queries are not handy for automating ESHs control since they do not tend to reduce annoyance [23].

There exist also approaches that infer occupancy instead of preferences to consider human-in-the-loop in ESHs control [24]. These intrusive methods require several sensors to detect occupancy and model customers' behavior. Avoiding user queries makes it hard for the modeling process since finding all variables that affect user behavior is not easy, and general assumptions result in low accuracy predictions. [25]. Furthermore, control of heating systems based on occupancy measurements becomes reactive and unsuitable for planning DR programs [26].

B. PAPER CONTRIBUTIONS

This paper presents a recommender system that suggests adequate thermal comfort valuations for practical participation in economic DR programs. The recommender can be seen as a regressor in a supervised machine learning framework that takes previous users' inputs to predict their preferences

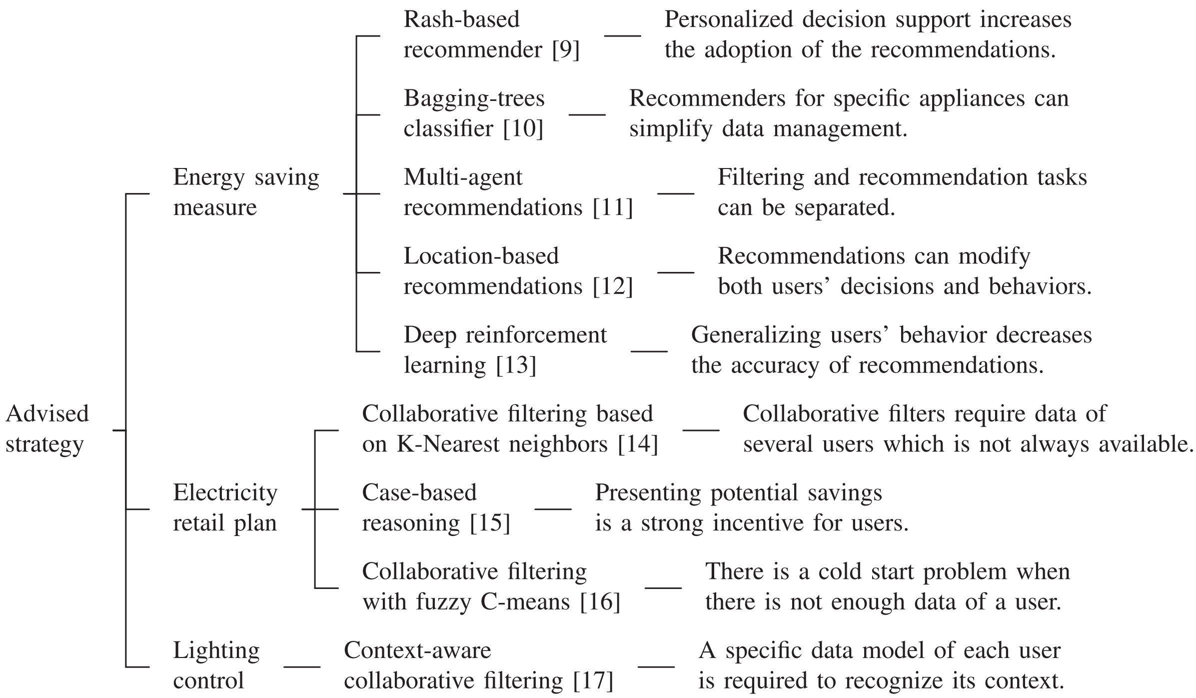

FIGURE 1. Findings of previous applications of recommender systems in home energy management systems.

TABLE 1 Characteristics of the Human-in-The-Loop Control Schemes for Residential ESHs

Approach	Allow customers to impose decisions	Suitable for inexpert customers	Non-intrusive	Context-aware	Helps reducing queries (annoyance)
Comfort assumption [28]		✓	✓		✓
Direct queries [29]	✓		✓	✓	
Occupancy modelling [26]		✓		✓	✓
Collaborative filtering [20]		✓	✓		✓
Active learning [22]	✓	✓	✓	✓	
Proposed approach	✓	✓	✓	✓	✓

in future conditions [27]. In the proposed method, users can accept or adjust the suggestions before automated control mechanisms take action. The human intervention is reduced in the long run as the recommender system learns users' preferences. Furthermore, the recommender presents user-friendly information, making it easy for the users to impose their preferences when they disagree with the suggestions. In this way, the recommender learns the specific context of a customer without querying for complex data or using intrusive methods. The features of the proposed approach are summarized in Table 1.

The proposed recommender system is integrated into a predictive control technique for the ESHs to participate in economic DR programs. This technique is based on a thermal dynamics model to verify the feasibility of consumption strategies. Previous studies have proved the feasibility of employing discrete-time linear models to represent thermal load dynamics, even in multi-zone buildings [30]. Considering the matters discussed concerning the proposed recommender system, the contributions of this paper can be summarized as follows:

- The design of a recommender system that helps customers complete the information of a predictive control scheme for ESHs. The system does not impose constraints but adapts the control mechanism to the users' context without intrusive methods. The recommendations are made over the weights of a payoff function, which reflects thermal comfort valuations and users' price-elasticity. The developed method builds on Gaussian process assumptions to update user preferences when they provide new information.
- A straightforward querying method based on the generalized optimal-choice axiom to retrieve preference information from inexpert users. The proposed method deploys user-friendly information for customers to compare similar options when making decisions. Then, they can directly modify the parameters of an ESH control mechanism and participate in DR programs. At the same time, they improve the recommender system training with relevant feedback. In contrast with the active-learning methods, users are not committed to providing information for each new DR condition.

C. PAPER ORGANIZATION

The rest of the paper is organized as follows: Section II presents the model-based predictive control formulation and the effects of user preferences on DR programs. Section III describes the recommender system, developed for the specific case of ESHs, and the characteristics of the queries to the users. Section IV discusses the case study and the validation method, followed by the concluding remarks in Section V.

II. MODEL-BASED PREDICTIVE CONTROL OF HEATING SYSTEMS

Economic DR programs send price signals to customers for relevant time windows [31]. Next, customers trust the information they receive and formulate their consumption strategy with certainty. This procedure allows for considering DR in planning and dispatching problems. Here, the DR configuration assumes no communication between residential customers, so their plans are not coordinated. For each customer, the consumption strategy s^* attempts to maximize the individual utility as the difference between customer benefit (thermal comfort) and electricity cost, as presented in (1). The decision variables correspond to the consumption at each time slot.

$$s^* = \underset{P_1, P_2, \ldots, P_T}{\text{argmax}} \sum_{t=1}^{T} (u_t - c_t) \tag{1}$$

$$\text{s.t.}: \quad 0 \leq P_t \leq P_{max} \quad \forall t \tag{2}$$

$$\theta_0 = \theta_T \tag{3}$$

where for t time slots, u_t is the benefit perceived from P_t energy consumption, and c_t is the energy cost. T is the market period and P_{max} is the capacity of the ESH system. The final temperature, θ_T must be equal to the initial condition, θ_0 to ensure the problem has an optimal substructure in the long run. The cost function, c_t with the known price signal, π_t is presented in (4) as a linear function since it is assumed there is no economy of scale for residential customers. The decision variable, P_t is continuous considering that the control signal can be modulated appropriately for thermostatically controlled heaters or heat pumps.

$$c_t = \pi_t P_t \tag{4}$$

It should be noted that the utility function, u_t is concave because there are no monotone preferences in temperature [32]. This function is not directly related to consumption but the actual temperature of the indoor air mass, θ_t. In fact, there is an optimal temperature, θ_{ref} that maximizes users' comfort so that the utility is proportional to the deviation from that reference, as presented in (5). The factor, φ_t encompasses the user's arc price-elasticity to weight the user's utility against the cost.

$$u_t = -\varphi_t (\theta_{ref} - \theta_t)^2 \tag{5}$$

In order to relate u_t to the consumption, it is necessary to include the thermal dynamic model of the system. For a single-zone structure, the dynamic response is similar to an RC circuit as described in the standard ISO 52016 [33]. In discrete-time, the thermal model is reduced to the state-space equation, (6), where $\theta_{ext,t}$ is the external temperature, $P_{irr,t}$ is the total solar irradiation, and $P_{app,t}$ is the energy consumption of other appliances in the house. The coefficients $\alpha_1, \alpha_2, \beta_1, \beta_2, \beta_3$ can be found with a least-squares method and updated adaptively when new measurements are available. For multi-zone buildings it is necessary to include the heat transfer coefficients between zones [34].

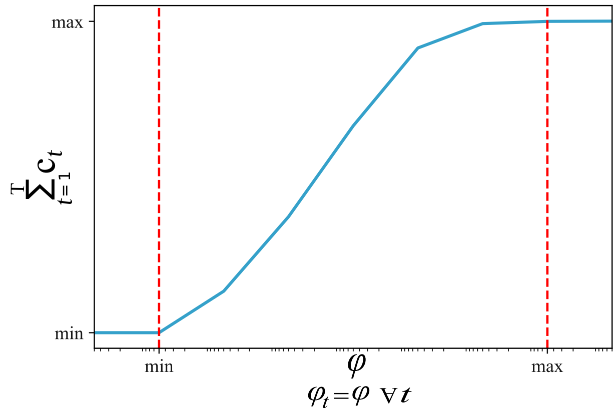

FIGURE 2. Agreed cost as a function of the thermal comfort valuation.

$$\theta_t = \alpha_1 \theta_{ext,t} + \alpha_2 \theta_{t-1} + \beta_1 P_{irr,t} + \beta_2 P_{app,t} + \beta_3 P_t \tag{6}$$

The constraint in (3) can be relaxed to preserve the tractability of the problem, considering that the optimization problem has the form of a linear quadratic regulator with this thermal model. Next, to perform predictive control, it is necessary to acquire the forecasts for $\theta_{ext,t}$ and P_{irr} from external information services. Likewise, it is required to either make a forecast for $P_{app,t}$, or gather such information from other smart controllers in the house if available. All the forecasted variables used to develop the consumption strategy represent a source of error. However, assuming a well-designed DR mechanism, it is reasonable to formulate a risk-neutral consumption strategy relying on the expected values of the forecasted variables. [35].

Finally, the factor φ_t of the utility function must be set according to the customer preferences. In this case, the choice of this factor is automated by the recommender system taking into account the historical user choices. However, in the automation, it is crucial to consider that the control problem formulation imposes limits on this factor where its increase or decrease does not change the final cost, as shown in Fig. 2. The cost is maximum when the indoor temperature is equal to the reference θ_{ref}, and it cannot be zero due to the constraint in (3). The curve in Fig. 2 can be stretched or contracted depending on the parameters of the thermal model.

III. RECOMMENDER SYSTEM

Given the configuration of the economic DR mechanisms, residential customers cannot collect data from their peers. This

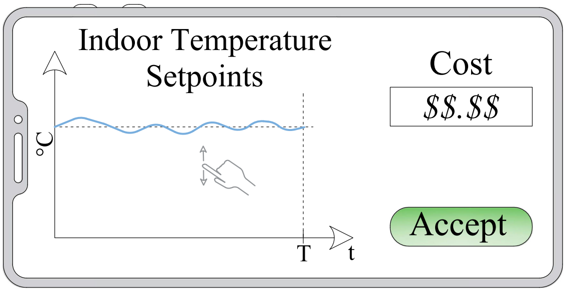

FIGURE 3. Basic components of the user interface.

also preserve their data privacy. Thus, they cannot implement collaborative filtering techniques to obtain values for φ_t [36]. Consequently, the option is to implement a content-based recommender considering only the previous choices of the same customer. The thermal comfort valuation φ_t is divided into a scale factor F and a form factor δ_t, as presented in (7), to capture context-aware preferences and intraday changes. The exponential transformation allows covering a wide range of values in compact numbers.

$$\varphi_t = e^{-F\delta_t} \tag{7}$$

Previous studies on ranking problems have shown that users' preferences fit Mallow's model because a customer tends to choose the same values under the same querying conditions [37]. Thus, that model is frequently used for distributions of ranked preferences with different distance metrics of the permutations. Following the same principle of Mallow's model (law of large numbers), the continuous parameter, F is modeled by a normal distribution, and the recommender system becomes a Gaussian Process (GP). The feasibility of modeling thermal preferences through a GP has been explored formerly in [38]. For δ_t, which changes during the day, it is convenient to set a profile in an interval between (0,1] to capture household occupancy information. This profile is also obtained from previous customer choices within similar days.

This content-based approach can face difficulties in the early stages when there is not enough data for training. Some solutions to this cold-start problem have been proposed based on Bayesian optimization and inverse reinforcement learning [39], [40]. In both cases, it is necessary to use a prior distribution of a function or its parameters and update it when new observations are available.

A. DATA DISAGGREGATION

The user interface should display simple information since the customer only needs to adjust thermostat set-points for the next market period while inspecting the electricity cost. The basic components of the user interface are presented in Fig 3. From the user's inputs, it is possible to obtain the φ_t profile for the utility function using the Karush-Kuhn-Tucker conditions. Next, it is necessary to disaggregate the values of F and the profiles of δ_t. To do this for a particular day, F is considered as the maximum value of customer choices in the product $(F\delta_t)$ since $\delta_t \in (0, 1]$.

For the next market period, the recommended value F is obtained from a GP regression. Then, δ_t for $t \in [1, \ldots, T]$ is taken by weighting previous values of similar days. Various measures of similarity have been explored in [41] considering several features. Since δ_t captures occupancy information, it is suitable to consider similarity according to the day of the week and to give more weight to the adjacent days. For instance, the occupancy of the next Friday can be closer to the one of the last Friday, or the occupancy of a Saturday can be similar to previous weekends.

B. GAUSSIAN PROCESS

From the machine learning perspective, the recommender system is a regressor that maps user preferences under given conditions. The general formulation of the GP regressor is presented in [42]. The parameter F is influenced by features like the electricity price and the external temperature. Therefore, its estimation, $\hat{F}$ can be defined by (8) below. Since this parameter is considered constant during the market period, i.e. one day, it is suitable to use the daily average values $\bar{\theta}_{ext}$ and $\bar{\pi}$ as feature variables in the regression.

$$\hat{F} = f(\bar{\pi}, \bar{\theta}_{ext}) = f(x) \tag{8}$$

In order to simplify the notation of the GP, the set of n days historical features is outlined by $X_H = [x_1, \ldots, x_n]^T \in \mathbb{R}^{n\times 2}$ and the next market-period features are $x_M \in \mathbb{R}^{1\times 2}$. Likewise, $F_H = [F_1, \ldots, F_n] \in \mathbb{R}^n$ stands for corresponding historical observations of F obtained from previous user's choices. The GP is presented in (9),

$$\begin{bmatrix} F_H \\ F \end{bmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n} & K(X_H, x_M) \\ K(x_M, X_H) & K(x_M, x_M) \end{bmatrix}\right) \tag{9}$$

where K is the covariance function that maps the feature vectors, x, into Gram matrices, and σ_n represents the noise of the observations. Since preferences are expected to be stationary, there is no need for a tracking system and the prior mean is set to zero. This canonical notation is favourable for the GP because the conditional distributions are well-known in the literature [42]. In this case, the interest is in the distribution of F given F_H, X_H, and x_M, so the first and second-order moments are presented in (10) and (11). In the GP, the expected value, $\mathbb{E}[F|x_M, X_H, F_H]$ is also the maximum a posteriori (MAP) estimator of F, which is both an adequate solution for regression problems and a good recommendation in some situations as presented later in this document.

$$\mathbb{E}\left[F|x_M, X_H, F_H\right] = K(x_M, X_H)\left[K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n}\right]^{-1} F_H \tag{10}$$

$$\sigma_M^2 = K(x_M, x_M) - K(x_M, X_H)\left[K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n}\right]^{-1} \times K(X_H, x_M) \tag{11}$$

In practice, the inverse of $[K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n}]$ is computed with the Cholesky decomposition since the covariance matrices are Hermitian positive-definite. For this case, it is convenient to consider the radial basis function as the covariance function, which has two parameters σ_H and L, presented in (12). This function is helpful because user preferences are expected to be stationary, discarding seasonal effects, vacancies, and other long-run behavior changes [43]. Thus, the covariance function is stationary considering only the similarity in the features space.

$$k(x, x') = \sigma_H^2 e^{-\frac{1}{2}(x-x')^T L^{-1}(x-x')} \tag{12}$$

where $L \in \mathbb{R}^{2\times 2}$ is a diagonal matrix in which large entries imply a low correlation between the feature and the covariance. σ_H^2 corresponds to the expected value of the squared norm of X_H elements. Then, the required Gram matrices are composed as follows:

$$K(X_H, X_H) = \begin{bmatrix} k(x_1, x_1) & \dots & k(x_1, x_n) \\ \vdots & \ddots & \vdots \\ k(x_n, x_1) & \dots & k(x_n, x_n) \end{bmatrix} \tag{13}$$

$$K(x_M, X_H) = [k(x_M, x_1), \dots, k(x_M, x_n)] \tag{14}$$

$$K(x_M, x_M) = [k(x_M, x_M)] \tag{15}$$

the set of parameters, L can be tuned by maximizing the likelihood of historical data using quasi-Newton methods. The likelihood for the formulated GP is presented in (16),

$$\begin{aligned} \log(\mathbb{P}(F_H|X_H)) = & -\frac{1}{2} F_H^T \left[K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n}\right]^{-1} F_H \\ & -\frac{1}{2}\log\left(K(X_H, X_H) + \sigma_n^2 \mathbb{I}_{n\times n}\right) \\ & -\frac{n}{2}\log(2\pi) \end{aligned} \tag{16}$$

where n is the number of samples in F_H. Subsequently, it is possible to estimate the MAP estimator of F for the next market period. At this point, it is necessary to analyze when the MAP expresses an accurate recommendation. First, utility functions are considered ordinal and users are assumed to be rational, which means they have *strict preferences* for cost savings [44]. Besides, customers can have a range of acceptable temperatures for which the perceived utility is the same. This means, in brief, a consumption strategy s_1 is strictly preferred over s_2 ($s_1 \succ s_2$) if it has a lower cost and manages the temperature setpoints inside an acceptable margin. However, this does not imply that the user knows what the optimal strategy is. In fact, users can merely accept consumption strategies cheaper than their WTP without searching for the optimal. Therefore, recommending the MAP, obtained from historical data, will lead to an acceptable value of $\hat{F}$ but not to the best.

It is convenient to suggest the MAP when the GP becomes stationary, assuming stationary preferences of rational users. Conversely, it is favorable to suggest a higher value than the MAP (lower-cost strategy) while the GP has not reached a steady state. For example, the MAP plus one standard deviation. In this way, the recommender system improves the serendipity by showing unknown alternatives to the user. The statistical moments of the GP are compared before and after an update through a pre-defined threshold to check stationarity. The recommender system is summarized in Algorithm 1.

Algorithm 1: Recommender Algorithm.

input : Previous accepted parameters ($F\delta_t$) with the corresponding features ($\bar{\theta}_{ext}, \bar{\pi}$), future prices and weather conditions for the next market period
output: Parameters for the next market period ($\hat{F}\hat{\delta}_t$)
begin
  Divide the historical user's accepted parameters into F and $\delta_t \in (0, 1]$;
  Compute the covariance matrix with the available data using Eq. 12;
  Compute $\mathbb{E}[F|x_M, X_H, F_H]$ using Eq. 10;
  Check weak-sense stationarity regarding the previous statistical moments considering an adequate threshold;
  if *The process is in steady state* **then**
    $\hat{F} = \mathbb{E}[F|x_M, X_H, F_H]$;
  else
    $\hat{F} = \mathbb{E}[F|x_M, X_H, F_H] + \sigma_M$;
  Set $\hat{\delta}_t$ according to past similar days
end

The process of moving from the MAP can be seen as an exploration that is needed to find the optimal strategy from stochastic feedback [39]. When the process converges, the recommender system can stop the exploration and start the exploitation of the acquired knowledge about user's preferences, i.e. suggesting only the MAP. Accordingly, it is suitable to move with the standard deviation, which is big when the agent has few data points and gets smaller quickly when the user gives consistent responses.

A drawback of the formulated GP is its application to online learning because the covariance computation and the matrix inversion are impractical for large datasets. The GP regressor has $\mathcal{O}(n^2)$ memory cost and $\mathcal{O}(n^3)$ computational complexity, where n is the number of samples [45]. This issue can be dealt with a recursive estimation, similar to the Kalman filter, by considering the previous posterior as the prior and using a linear model, which relates states and observations [46]. In this way, the recommender system can avoid storing extensive historical data and keep only information from the last market period.

C. CUSTOMER QUERIES

Considering that querying can create annoyance to the user, it is not suitable to ask for approval of a suggested power consumption plan at each time. Actually, when the customer does not make corrections, it is assumed that it has accepted the suggestion, and the related data is stored to train the recommender system. The customer interacts through a Human-Machine Interface (HMI) that displays the suggested DR agreement for the next market period (temperature set-points and cost) and allows changing the set-points. Once the

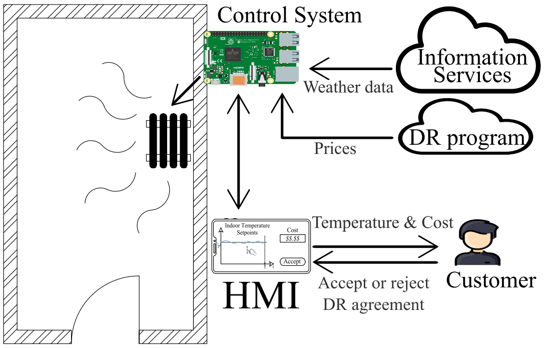

FIGURE 4. Information flows in the proposed recommender system.

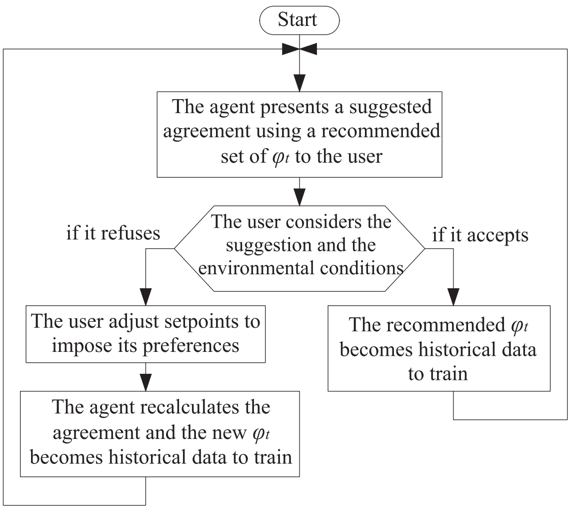

FIGURE 5. Flow chart of the user interactions.

user changes a parameter, it is necessary to recalculate the cost. This interaction is sketched in Fig. 4.

As a result, at the start of each market period, the system recommends a combination of temperature and cost to the customer to either accept (do nothing) or refuse (adjust setpoints). At the end, the agreed parameters become data to train the recommender system in both cases. This procedure is presented in Fig. 5. It is worth mentioning that the customer can also modify the reference temperature, θ_{ref}. However, such modification is not considered in the current market period but in the next one. Consequently, a single parameter is queried and the customer can compare similar options when making decisions.

The HMI does not present penalties for deviation from the DR agreement in the proposed approach because the system is performing a predictive control. The ultimate goal of the HMI is to present user-friendly information and allow users to set their preferences. Thus, it is convenient to show only the temperature set-points and the total cost, and let the user modify just one of those. For more expert users, the HMI can include penalties information, if it influences their decisions.

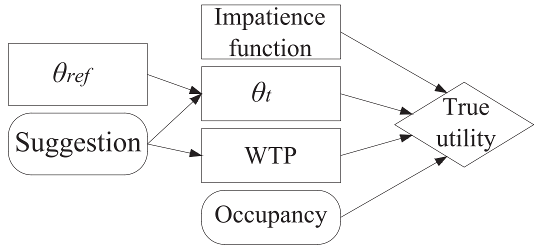

FIGURE 6. Decision network for the customer.

IV. EXPERIMENTAL SETUP AND VALIDATION

Since a customer becomes committed to the DR mechanism, the accuracy of suggestions is crucial in the proposed mechanism. Therefore, the metrics for the recommender system must consider both the decision support and the deviation of the suggestions. In this case, we use the user's acceptance rate for the decision support and the difference from the final agreement in kWh as the accuracy metric. We perform an offline validation simulating the customer decision process to validate the recommender system. Besides, we compare the proposed approach with other two recommender techniques frequently used in the literature, under the same conditions:

- A persistency method that recommends user's accepted values for the last market period [47]. Considering stationarity in the preferences, a customer can choose the same options for consecutive days when external features (π and θ_{ext}) have slight changes.
- A Support Vector Regressor (SVR) that has the same kernel (covariance function) as the GP [48]. For this case, the acceptable margin from the hyperplane is set to 0.1, the regularization parameter for deviations is considered 1, and the training stop criteria is a tolerance of 0.001.

The variables that influence customer behavior are depicted in the decision network in Fig. 6. The reference temperature, θ_{ref}, the temperature setpoints, θ_t, the WTP, and the impatience function are decision nodes for the user, while the suggestion and the occupancy profile are chance nodes. The impatience function is merely a step function. Thus, the user stops interacting with the HMI after a defined number of queries. Formally, the queries reduce customers' comfort, and they stop interacting when their utility starts decreasing at each interaction [38]. Highly involved users adjust setpoints until they find an appropriate agreement, giving less noisy data to the recommender system. It is worth mentioning that the offline validation presented here does not allow for measuring the user experience (UX) with the HMI.

The decision network is compiled for the simulated customer in Algorithm 2. The decision process starts with a predefined WTP and the range of acceptable temperatures. The reference temperature is fixed beforehand. The HMI specifies the step-size in which φ_t can increase or decrease.

As a result, more patient customers with lower WTPs make more changes in the HMI while very impatient ones with

Algorithm 2: User Decision Process.

input : Minimum acceptable temperatures, WTP, θ_{ref}, Maximum queries
output: $F\delta_t$ profile
begin
  Check the suggested agreement: temperature setpoints and cost ;
  while *the queries are tolerable **and** the agreement is not acceptable* **do**
    if *the cost is higher than the WTP* **then**
      increase all $F\delta_t \quad \forall t \in [1, ..., T]$
    if *the setpoints are lower than the acceptable temperatures* **then**
      decrease $F\delta_t$ for all unacceptable temperatures θ_t
  end
end

higher WTPs accept any suggested agreement at first. Due to the interest in the acceptance rate, a sensibility analysis is carried out with different values of WTP.

A. DATA DESCRIPTION

The weather data for θ_{ext} and P_{irr} corresponds to 90 winter days (20th December to 20th March, 2018) in the city of Trois-Rivieres, Quebec, Canada. The P_{app} data is related to a real house in the same location. This data has been used to find the parameters of the thermal model by ordinary least-squares regression. The resulting values are presented in (17) below. θ_{ext} and θ, P_{app} and P, and P_{irr} are expressed in Celsius degrees, kWh, and kW/m^2, respectively.

$$\begin{aligned}\theta_t \approx\ & 0.08316\theta_{ext,t} + 0.99168\theta_{t-1} + 0.00016P_{irr,t} \\ & + 0.07142P_{app,t} + 0.11064P_t\end{aligned} \tag{17}$$

The value of φ_t can vary from -10 to 10 at steps of 0.2. Finally, the market period, T is 24 hours and the signals are discretized at 5-minute intervals. The price signal corresponds to the spot market in Ontario, Canada during the same days in 2018 [49]. This signal is presented in Fig. 7 for the simulation period. The average price is 0.0532 CAD/kWh. The range of acceptable temperatures, used to simulate customer's behavior, is the same every day, as presented in Fig. 8. The threshold to check stationarity in the GP is 0.1 in the mean and 0.01 in the variance. The recommendation for the first day is zero for all recommenders, and thus, the prior for the GP is $\mathcal{N}(0, 1)$.

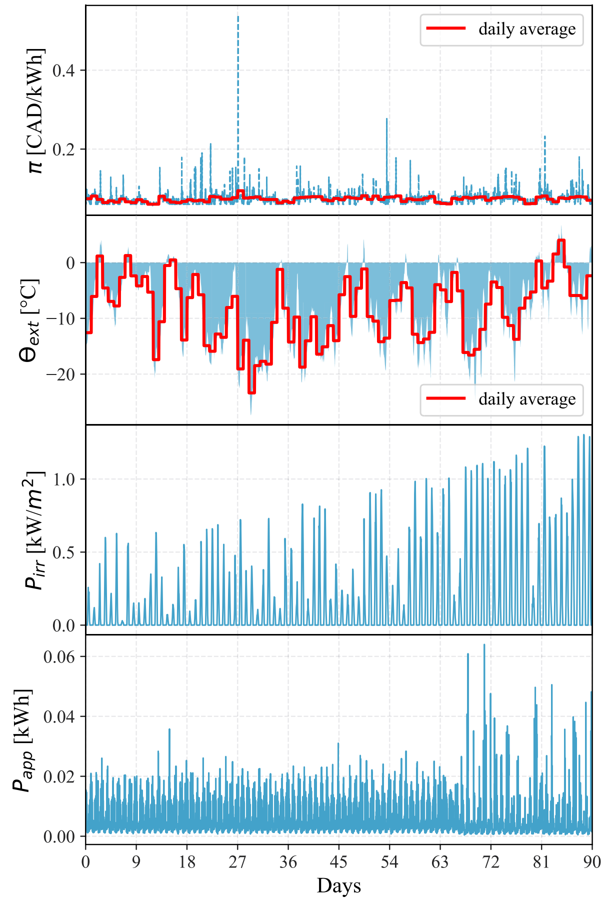

FIGURE 7. Signals used for the off-line validation.

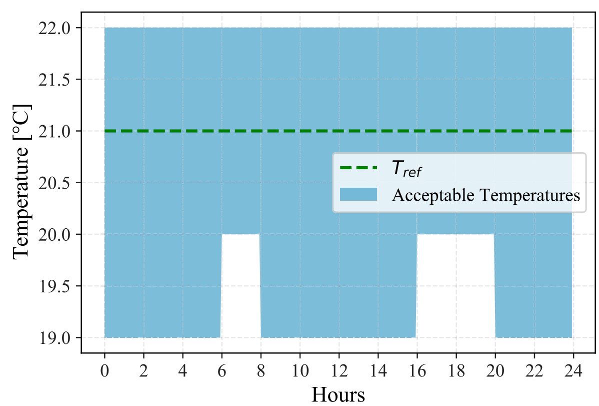

FIGURE 8. Range of acceptable indoor temperatures.

B. RESULTS AND DISCUSSION

Given the price signal, performing a sensitivity analysis on WTP from 1 to 5 CAD per day is reasonable. The maximum number of queries is set at 20 to scan a wide range of φ_t. The acceptance rate results during the 90 days for all combinations are summarized in Table 2. he proposed GP recommender leads to a higher acceptance rate in most cases, especially when the user's WTP is low thanks to the exploration process that constantly looks for a lower cost. A user with a WTP of 5 CAD/day accepts all the suggestions within the range of comfortable temperatures since the cost never exceeds that value. This case is useful to validate that the recommendations are fitting and do not diverge from training data.

A method with a higher acceptance rate does not imply that it has the most accurate suggestions, as customers accept any consumption strategy that meets their temperature and cost conditions. Thus, it is also relevant to consider the recommenders' deviations from the final DR agreement of the customer. The energy deviations are presented in Fig 9

TABLE 2 Acceptance Rate of the Recommender Systems

WTP [CAD/day]	Maximum queries	GP	SVR	Persistency
1	1	6/90	0/90	0/90
	5	46/90	0/90	0/90
	10	51/90	0/90	0/90
	20	54/90	12/90	1/90
2	1	7/90	0/90	1/90
	5	46/90	0/90	2/90
	10	51/90	2/90	2/90
	20	54/90	15/90	5/90
3	1	47/90	43/90	40/90
	5	60/90	44/90	46/90
	10	64/90	46/90	47/90
	20	67/90	46/90	47/90
4	1	83/90	85/90	65/90
	5	84/90	85/90	85/90
	10	84/90	85/90	85/90
	20	88/90	86/90	86/90
5	1	90/90	90/90	90/90
	5	90/90	90/90	90/90
	10	90/90	90/90	90/90
	20	90/90	90/90	90/90

in absolute values. It is not appropriate to separate positive and negative deviations since we are not considering penalties on the DR configuration. The SVR method has a lower variation for impatient customers because it does not make exploration. When impatience ceases to be a relevant noisy driver, the serendipity given by the GP method is beneficial to find better consumption strategies. In all cases, the GP-based recommender has a lower median deviation.

Patient customers give more consistent feedback the recommenders offer according to the WTPs. When it is not feasible to have a consumption strategy cheaper than WTP, users reject any suggestion and start changing temperatures. In such situations, patient customers will lose time trying to find other alternatives. An experienced user with some idea of the feasibility of the consumption strategies may find it relevant to have boundaries information on the HMI. For instance, indicating the high prices period can reduce customer exploration at that time.

Regarding the deviation in temperature, we present the number of hours where the suggestion was outside the comfortable range in Fig. 10. For impatient customers, the GP-based method suggest more uncomfortable temperatures because it is trying to reduce the cost. When the limit of queries increases, the temperature deviation of all recommenders decreases, and the GP becomes similar to the SVR method. In that scenario. the persistency method shows more temperatures outside the comfort limits, even for high WTP.

The GP recommender seems to be a good option for most cases. However, it requires a larger hardware infrastructure than other methods like a persistency recommender that can be implemented in a simple data buffer. Recommending the last values of both F and δ_t can suit impatient customers since their noisy inputs are inadequate for training other algorithms. Conversely, this procedure results in significant deviations for

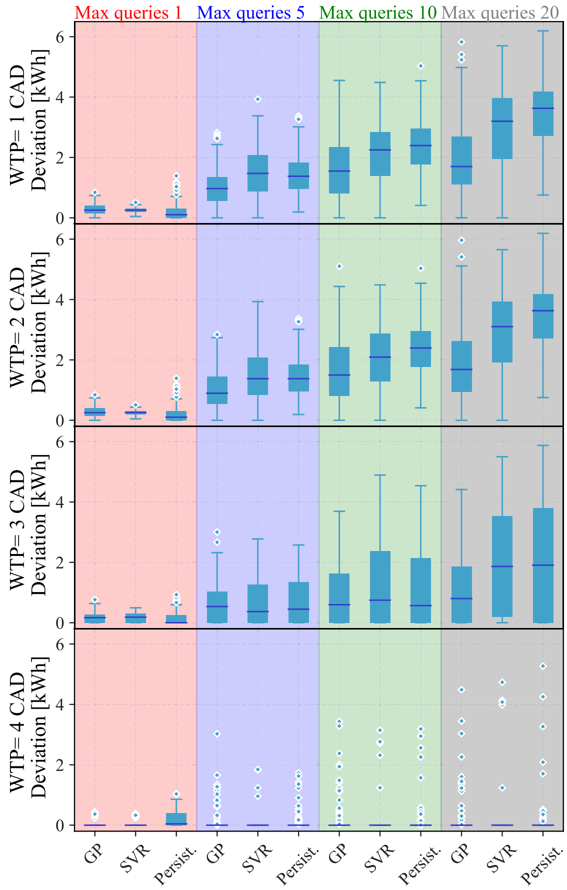

FIGURE 9. Daily consumption deviation between the suggested strategy and the DR agreement.

patient users. Hence, it is advisable to acquire detailed information from customers and install a recommender system depending on their specific context.

C. LIMITATIONS AND OPPORTUNITIES

The proposed approach allows incorporating a prior distribution to improve the performance of the recommendations. If there exists a characterization of residential demand price-elasticity in a particular location, it is possible to pre-configure the φ_t profile for that place. Besides, such information can be used to weight F against δ_t instead of using the (0,1] range. Collecting data from a group of users can help fit better the preferences for individuals, as it happens in platforms like Youtube [50].

One of the challenges for the recommenders is to deal with Big Data. When the agent obtains more information from the user, it is necessary to implement the recursive algorithms. Another alternative way to deal with this issue is to implement adaptive learning techniques and filter the training samples. However, this also requires patient customers who give less noisy data to train because their inconsistent preferences can be treated as outliers.

The offline validation setup used here is suitable to analyze the recommender system in controlled scenarios. In future

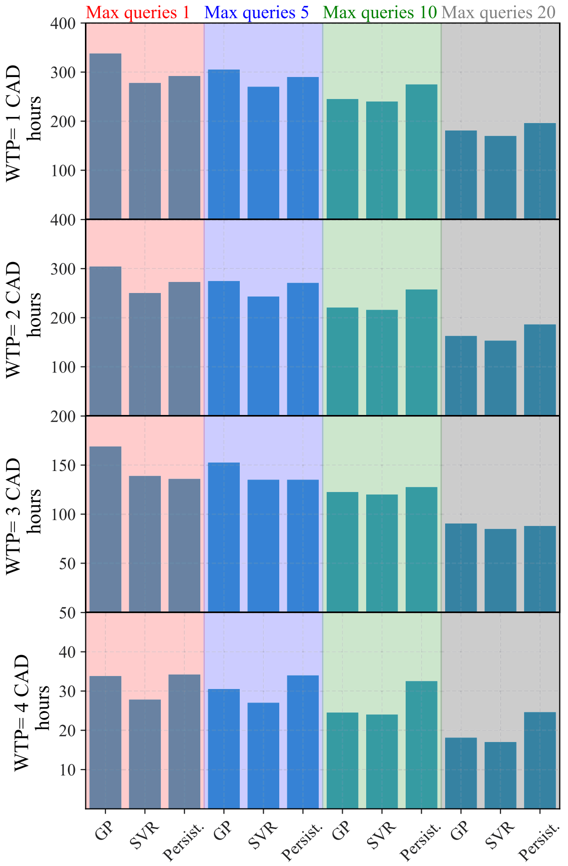

FIGURE 10. Total duration of suggested temperatures outside the comfort range in the 90-day simulation.

analysis, an online validation should be performed with clusters of actual customers and A/B testing methods to check the influence of external information and the HMI limitations. In that case, the thermal model must be adapted to each customer situation.

V. CONCLUSION

Space heater systems have the potential for becoming the principal source of flexibility from the residential sector in cold-weather countries like Canada. Demand response programs can exploit this flexibility to build a more sustainable grid. However, residential customers lack adequate expertise in controlling these devices according to market signals. Therefore, it is necessary to provide automated tools that engage users. These tools require knowledge about users' preferences to reach convenient agreements. Accordingly, this paper presents a recommender system that allows automated systems to suggest an appropriate valuation of thermal comfort and make an agreement in the context of an economic demand response program. The proposed method is based on a Gaussian Process that explores cheaper options to increase serendipity in the suggestions. This process is trained with the customer's previous choices. The information required from the customer is reduced to a single scale parameter in each transaction to provide simple queries and easily comparable options. In fact, in the presented approach, users need to supply more information only when they refuse the recommendation, not during all the training as in active learning methods. The offline validation of the recommender system shows that patient customers, who express more consistent preferences, give less noisy data and have a higher acceptance rate. From this perspective, this work contributes to analyzing recommender systems requirements according to customers' specific contexts. Future work focuses on integrating the proposed recommender system into complete home energy management systems to fully automate residential demand response.

ACKNOWLEDGMENT

The authors would like to thank the Laboratoire des technologies de l'énergie d'Hydro-Québec, the Natural Science and Engineering Research Council of Canada, and the Foundation of Université du Québec à Trois-Rivières.

REFERENCES

[1] International Renewable Energy Agency (IRENA), "Power system flexibility for the energy transition. Part I: Overview for policy makers," IRENA, Abu-Dhabi, UAE, Tech. Rep. 978-92-9260-089-1, 2018. [Online]. Available: https://www.irena.org/

[2] J. Cochran *et al.*, "Flexibility in 21st century power systems," National Renewable Energy Laboratory (NREL), Golden, CO, USA, Tech. Rep. NREL/TP-6A20-6172, 2014.

[3] L. Martirano *et al.*, "Demand side management in microgrids for load control in nearly zero energy buildings," *IEEE Trans. Ind. Appl.*, vol. 53, no. 3, pp. 1769–1779, May/Jun. 2017.

[4] Z. Liu, Q. Wu, K. Ma, M. Shahidehpour, Y. Xue, and S. Huang, "Two-stage optimal scheduling of electric vehicle charging based on transactive control," *IEEE Trans. Smart Grid*, vol. 10, no. 3, pp. 2948–2958, May 2019.

[5] J. Hu, G. Yang, K. Kok, Y. Xue, and H. W. Binder, "Transactive control: A framework for operating power systems characterized by high penetration of distributed energy resources," *J. Modern Power Syst. Clean Energy*, vol. 5, no. 3, pp. 451–464, May 2017. [Online]. Available: http://link.springer.com/10.1007/s40565-016-0228-1

[6] B. Gao, W. Zhang, Y. Tang, M. Hu, M. Zhu, and H. Zhan, "Game-theoretic energy management for residential users with dischargeable plug-in electric vehicles," *Energies*, vol. 7, no. 11, pp. 7499–7518, 2014.

[7] International Energy Agency (IEA), "Tracking buildings 2020," IEA, Paris, Tech. Rep., 2020. [Online]. Available: https://www.iea.org/reports/tracking-buildings-2020

[8] J.-H. Kim and A. Shcherbakova, "Common failures of demand response," *Energy*, vol. 36, no. 2, pp. 873–880, Feb. 2011. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0360544210007176

[9] A. Starke, M. Willemsen, and C. Snijders, "Effective user interface designs to increase energy-efficient behavior in a rasch-based energy recommender system," in *Proc. 11th ACM Conf. Recommender Syst.*, New York, NY, USA, 2017, pp. 65–73. [Online]. Available: https://dl.acm.org/doi/10.1145/3109859.3109902

[10] A. Alsalemi *et al.*, "Achieving domestic energy efficiency using micro-moments and intelligent recommendations," *IEEE Access*, vol. 8, pp. 15047–15055, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/8959214/

[11] D. Jiménez-Bravo, J. Pérez-Marcos, D. De la Iglesia, G. Villarrubia González, and J. De Paz, "Multi-agent recommendation system for electrical energy optimization and cost saving in smart homes," *Energies*, vol. 12, no. 7, Apr. 2019, Art. no. 1317. [Online]. Available: https://www.mdpi.com/1996-1073/12/7/1317

[12] P. Wei, S. Xia, and X. Jiang, "Energy saving recommendations and user location modeling in commercial buildings," in *Proc. 26th Conf. User Model., Adaptation Personalization*, New York, NY, USA, 2018, pp. 3–11. [Online]. Available: https://dl.acm.org/doi/10.1145/3209219.3209244

[13] P. Wei, S. Xia, R. Chen, J. Qian, C. Li, and X. Jiang, "A deep-reinforcement-learning-based recommender system for occupant-driven energy optimization in commercial buildings," *IEEE Internet Things J.*, vol. 7, no. 7, pp. 6402–6413, Jul. 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9001078/

[14] Y. Zhang, K. Meng, W. Kong, and Z. Y. Dong, "Collaborative filtering-based electricity plan recommender system," *IEEE Trans. Ind. Informat.*, vol. 15, no. 3, pp. 1393–1404, Mar. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8412100/

[15] T. Pinto, R. Faia, M. Navarro-Caceres, G. Santos, J. M. Corchado, and Z. Vale, "Multi-agent-Based CBR recommender system for intelligent energy management in buildings," *IEEE Syst. J.*, vol. 13, no. 1, pp. 1084–1095, Mar. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/8533391/

[16] F. Luo, G. Ranzi, X. Wang, and Z. Y. Dong, "Social information filtering-based electricity retail plan recommender system for smart grid end users," *IEEE Trans. Smart Grid*, vol. 10, no. 1, pp. 95–104, Jan. 2019. [Online]. Available: https://ieeexplore.ieee.org/document/7997914/

[17] P. Kar, A. Shareef, A. Kumar, K. T. Harn, B. Kalluri, and S. K. Panda, "ReViCEE: A recommendation based approach for personalized control, visual comfort & energy efficiency in buildings," *Building Environ.*, vol. 152, pp. 135–144, Apr. 2019. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0360132319300459

[18] Y. Himeur *et al.*, "A survey of recommender systems for energy efficiency in buildings: Principles, challenges and prospects," *Inf. Fusion*, vol. 72, pp. 1–21, Aug. 2021. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1566253521000178

[19] T. Silveira, M. Zhang, X. Lin, Y. Liu, and S. Ma, "How good your recommender system is? A survey on evaluations in recommendation," *Int. J. Mach. Learn. Cybern.*, vol. 10, no. 5, pp. 813–831, 2019. [Online]. Available: http://dx.doi.org/10.1007/s13042-017-0762-9

[20] D. Daum, F. Haldi, and N. Morel, "A personalized measure of thermal comfort for building controls," *Building Environ.*, vol. 46, no. 1, pp. 3–11, Jan. 2011. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0360132310001915

[21] A. Ghahramani, F. Jazizadeh, and B. Becerik-Gerber, "A knowledge based approach for selecting energy-aware and comfort-driven HVAC temperature set points," *Energy Buildings*, vol. 85, pp. 536–548, Dec. 2014. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0378778814007932

[22] M. Shann and S. Seuken, "An active learning approach to home heating in the smart grid," in *Proc. Int. Joint Conf. Artif. Intell.*, 2013, pp. 2892–2899.

[23] X. Jin, K. Baker, D. Christensen, and S. Isley, "Foresee: A user-centric home energy management system for energy efficiency and demand response," *Appl. Energy*, vol. 205, pp. 1583–1595, Nov. 2017.

[24] L. Rueda *et al.*, "Online unsupervised occupancy anticipation system applied to residential heat load management," *IEEE Access*, vol. 9, pp. 109806–109821, 2021.

[25] A. A. Adamopoulou, A. M. Tryferidis, and D. K. Tzovaras, "A context-aware method for building occupancy prediction," *Energy Buildings*, vol. 110, pp. 229–244, Jan. 2016. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0378778815303212

[26] R. Tang, S. Wang, and K. Shan, "Optimal and near-optimal indoor temperature and humidity controls for direct load control and proactive building demand response towards smart grids," *Automat. Construction*, vol. 96, pp. 250–261, Dec. 2018. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0926580518303133

[27] R. C. Bagher, H. Hassanpour, and H. Mashayekhi, "User trends modeling for a content-based recommender system," *Expert Syst. Appl.*, vol. 87, pp. 209–219, Nov. 2017. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0957417417304384

[28] R. Adhikari, M. Pipattanasomporn, M. Kuzlu, and S. R. Bradley, "Simulation study of transactive control strategies for residential HVAC systems," in *Proc. IEEE PES Innov. Smart Grid Technol. Conf. Europe*, 2016, pp. 1–5.

[29] N. Mohammad and A. Rahman, "Transactive control of industrial heating-ventilation-air-conditioning units in cold-storage warehouses for demand response," *Sustain. Energy, Grids Netw.*, vol. 18, 2019, Art. no. 100201. [Online]. Available: https://doi.org/10.1016/j.segan.2019.100201

[30] D. T. Vedullapalli, R. Hadidi, and B. Schroeder, "Combined HVAC and battery scheduling for demand response in a building," *IEEE Trans. Ind. Appl.*, vol. 55, no. 6, pp. 7008–7014, Nov./Dec. 2019.

[31] F. Liberati, A. D. Giorgio, A. Giuseppi, A. Pietrabissa, E. Habib, and L. Martirano, "Joint model predictive control of electric and heating resources in a smart building," *IEEE Trans. Ind. Appl.*, vol. 55, no. 6, pp. 7015–7027, Nov./Dec. 2019.

[32] P. Samadi, H. Mohsenian-Rad, R. Schober, and V. W. Wong, "Advanced demand side management for the future smart grid using mechanism design," *IEEE Trans. Smart Grid*, vol. 3, no. 3, pp. 1170–1180, Sep. 2012.

[33] T. Zakula, M. Bagaric, N. Ferdelji, B. Milovanovic, S. Mudrinic, and K. Ritosa, "Comparison of dynamic simulations and the ISO 52016 standard for the assessment of building energy performance," *Appl. Energy*, vol. 254, 2019, Art. no. 113553. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306261919312279

[34] R. Chi, Y. Hui, R. Wang, B. Huang, and Z. Hou, "Discrete-time-Distributed adaptive ILC with nonrepetitive uncertainties and applications to building HVAC systems," *IEEE Trans. Syst., Man, Cybern. Syst.*, to be published, doi: 10.1109/TSMC.2021.3113090. [Online]. Available: https://ieeexplore.ieee.org/document/9552438/

[35] S. Kuhn and R. Schultz, "Risk neutral and risk averse power optimization in electricity networks with dispersed generation," *Math. Methods Operations Res.*, vol. 69, no. 2, pp. 353–367, May 2009. [Online]. Available: http://link.springer.com/10.1007/s00186-008-0264-3

[36] J. Vuelvas, F. Ruiz, and G. Gruosso, "Limiting gaming opportunities on incentive-based demand response programs," *Appl. Energy*, vol. 225, pp. 668–681, Sep. 2018.

[37] V. Vitelli, O. Sorensen, M. Crispino, A. Frigessi, and E. Arjas, "Probabilistic preference learning with the Mallows rank model," *J. Mach. Learn. Res.*, vol. 18, no. 158, pp. 1–49, 2018. [Online]. Available: http://jmlr.org/papers/v18/15-481.html

[38] A. Perrault and C. Boutilier, "Experiential preference elicitation for autonomous heating and cooling systems," in *Proc. Int. Joint Conf. Auton. Agents Multiagent Syst.*, 2019, pp. 431–439.

[39] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of Bayesian optimization," *Proc. IEEE*, vol. 104, no. 1, pp. 148–175, Jan. 2016.

[40] N. Bazenkov and M. Goubko, "Advanced planning of home appliances with consumer's preference learning," in *Proc. Russian Conf. Artif. Intell.*, 2018, pp. 249–259. [Online]. Available: http://link.springer.com/10.1007/978-3-030-00617-4_23

[41] C. Aggarwal, A. Hinneburg, and D. Keim, "On the surprising behavior of distance metrics in high dimensional space," *Amer. J. Dermatopathology*, vol. 30, no. 4, pp. 420–434, 2001. [Online]. Available: http://link.springer.com/10.1007/3-540-44503-X_27

[42] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. vol. 7. London, U.K.: MIT Press, 2006.

[43] D. Rafailidis and A. Nanopoulos, "Modeling users preference dynamics and side information in recommender systems," *IEEE Trans. Syst., Man, Cybern. Syst.*, vol. 46, no. 6, pp. 782–792, Jun. 2016. [Online]. Available: http://ieeexplore.ieee.org/document/7194815/

[44] P. Glimcher and E. Fehr, *Neuroeconomics: Decision Making and the Brain*, vol. 63, 2nd ed., London, U.K.: Elsevier, 2014.

[45] F. Perez-Cruz, S. Van Vaerenbergh, J. J. Murillo-Fuentes, M. Lazaro-Gredilla, and I. Santamaria, "Gaussian processes for nonlinear signal processing: An overview of recent advances," *IEEE Signal Process. Mag.*, vol. 30, no. 4, pp. 40–50, Jul. 2013.

[46] M. Schürch, D. Azzimonti, A. Benavoli, and M. Zaffalon, "Recursive estimation for sparse Gaussian process regression," *Statist.*, May 2019. [Online]. Available: http://arxiv.org/abs/1905.11711

[47] J. Beel, S. Langer, M. Genzmehr, and A. Nürnberger, "Persistence in recommender systems: Giving the same recommendations to the same users multiple times," in *Proc. Int. Conf. Theory Pract. Digit. Lib.*, 2013, pp. 386–390. [Online]. Available: http://link.springer.com/10.1007/978-3-642-40501-3_43

[48] S.-H. Min and I. Han, "Recommender systems using support vector machines," in *Proc. 5th Int. Conf. Web Eng.*, 2005, pp. 387–393. [Online]. Available: http://link.springer.com/10.1007/11531371_50

[49] The Independent Electricity System Operator (IESO), "Data directory: Realtime market price report," 2019, Accessed: Feb. 12 2019. [Online]. Available: http://www.ieso.ca/en/Power-Data/Data-Directory

[50] S. Baluja, "A simple and efficient method to handle sparse preference data using domination graphs: An application to YouTube," *Procedia Comput. Sci.*, vol. 80, pp. 2302–2311, 2016. [Online]. Available: http://dx.doi.org/10.1016/j.procs.2016.05.424