12th IEA Heat Pump Conference 2017

Developing synergies for automated optimal control of residential heat pumps

Hussain Kazmi [a, b*], Farah Cheaib [a], Johan Driesen [b]

[a]Data science group, Greenville Campus, Houthalen-Helchtern, 3530, Belgium
[b]ELECTA, ESAT, KU Leuven, 3000, Belgium

Abstract

Energy efficiency in the built environment has been identified as one of the key enabling technologies to meet global climate change targets. In this paper, we present promising results from a black box method to automatically characterize various aspects of heat pump operation in residential settings. Experimental data is gathered from heat pumps used to provide spatial heating and domestic hot water in recently refurbished net-zero energy houses. This is done by data-driven determination of the heat pump's performance and the impact of building occupants. These interactions, typically in the form of hot water consumption profiles and preferences for temperature set points, are learnt from sensor data. This allows the formulation of an explicit Markov Decision Process (MDP), which can be solved with the objective to maximize energy efficiency of local heat pump operation. In doing so, we show substantial gains over default policies (grounded in thermodynamics) but which don't consider occupant behaviour. Three key short-term benefits are envisaged from this research: first, leveraging such synergies allows the energy efficiency of heat pump operation to be improved by, on average, more than 10%. Second, automation unlocks the potential to circumvent the costly, non-generalizable model building step in model predictive control. Finally, it allows direct, unbiased benchmarking of theoretical performance of different types of heat pumps against real world performance.

Selection and/or peer-review under responsibility of the organizers of the 12th IEA Heat Pump Conference 2017.

Keywords: Smart buildings, heat pumps, reinforcement learning, optimal control, energy efficiency

1. Introduction

Policy instruments and increasing public awareness have created amenable conditions for energy efficiency measures and renewable energy sources [1, 2]. These measures and renewable sources are also becoming an increasingly economically viable option [3]. For the aging residential built environment, responsible for about 27% of the European energy demand [4], this has meant an accelerating proliferation of net-zero energy buildings (NZEB). These are grid-connected buildings that, typically over the course of a year, consume as much energy as they produce [5]. Frequently, NZEBs make heavy use of energy efficiency measures such as improved façade insulation techniques and higher efficiency heating mechanisms (e.g. heat pumps). This has the desirable side effect of driving these buildings towards complete electrification, severing the direct fossil fuel connection in its entirety (this argument disregards grid electricity generated using fossil fuels).

* Corresponding author. Tel.: +32 47 66 22614.
E-mail address: Hussain.kazmi@enervalis.com.

Research has shown however that these high efficiency buildings may not be the solution for efficiency standards as once hoped for. Theoretically designers of very high efficiency buildings in particular have often overestimated their efficiency in real life situations [6, 7]. This is primarily because manufacturers, in their quest for ever-increasing efficiency, have focused almost exclusively on thermodynamics and treated the human interaction component as a secondary concern. There have historically been good reasons for doing so. Thermodynamics is predictable and well-understood; the climate conditions prevalent in a geographical location are also stable with daily and seasonal patterns. On the other hand, human demand demonstrates substantial stochasticity [8]. This complicates optimizing for such behavior in the absence of additional information. With the rise of internet of things, this is now becoming possible.
Even if occupant behavior were to be known in advance, there is an additional problem in using traditional optimal control methods such as model predictive control (MPC) however. MPC techniques require the existence of known models for the heating systems involved [9, 10]. This means that models would be required for both the heat pump and the storage medium it is providing energy to (e.g. a hot water vessel or a building). Such model learning (or calibration) is an expensive, human-intensive step. Also, since there is no active learning in MPC, any errors in the model building process persist for the lifetime of the operational phase. Furthermore, MPC can't scale to the millions of different device configurations worldwide.
We present what we believe to be the first truly online framework capable of learning heat pump behavior and the systems it interacts with, demonstrated in real world settings. The framework, after having learnt such a model, is capable of solving for multiple objectives such as energy efficiency and price based optimization. Given only a standard set of sensors, the system is able to learn the dynamics of the system in real-world conditions and generalize to states never seen before. It is also able to improve the energy efficiency of the system, and can potentially offer these residual energy gains as flexibility to the electricity grid in case of need for demand response. By making optimization occupant-driven, the system reaches the highest possible energy efficiency and improves thermodynamic gains. These research findings can also be used to inform the dimensioning of such systems for future buildings and to identify possible improvement in the heat pump design itself.

2. Methodology

In this paper, we present a model-based reinforcement learning framework to automatically characterize and optimize the operation of heat pumps used for providing spatial heating as well as hot water. The heat pumps under consideration are air source heat pumps connected to recently refurbished, net-zero energy residential buildings, each equipped with a 200 litre storage vessel. Additionally, each heat pump is equipped with a 2kW booster heater. However, being data-driven, the proposed framework is independent of the specific heat pump technology or the type of building and storage vessel it is providing energy to.
Model-based reinforcement learning generally consists of two iterative steps: learning and planning. Learning is the process of building a model for the system dynamics and its interactions with the environment and the reinforcement learner (also termed the agent). It usually consists of two steps: first, features are extracted from the time series data and, second a regression model is trained on the extracted features. The choice of regression model is problem-dependent. A generative model learnt in this way can then be used to simulate future system behaviour given arbitrary initial conditions and control actions. Once such a model has been learnt, optimization (planning) can take place. This optimization takes the form of generating many future scenarios (roll-outs) and then choosing the sequence of control actions that maximizes some objective function.
In model-based reinforcement learning, the quality of the model used for planning defines the quality of optimal control. Since the regression model is data driven and not grounded in thermodynamics, two sources of error can cause inaccurate predictions, (1) stochasticity inherent to the time series (from measurement noise etc.) and (2) sampling artefacts (from limited state space exploration). The effect of the latter decreases as the model collects more experiences through its interaction with the system. Alternatively, explicit exploration strategies can drive the system to unknown states to improve the quality of the model. The former – system stochasticity – depends on relevant influencing variables not being observed etc. and can only be improved up to a constant, after which the model quality doesn't improve. In the following, we describe these ideas in greater detail.

2.1. Learning phase

The learning problem encompasses more than just function approximation. Representations for the heat pump, the storage vessel and the building itself have to be learnt. Just as importantly, future occupant behaviour has to be predicted. Additionally, how these states are estimated during the operational phase also has to be considered (i.e. what data is available and how reliable it is). The representations learnt are somewhat dependent on the

objective function to be optimized. For instance, if the objective is to optimize for energy efficiency, we don't concern ourselves with the power profile of the heat pump; on the other hand, if the objective were to maximize the share of solar energy in heat pump operation, such a power profile would be useful.
To formulate the problem as a reinforcement learning task, we first formalize the notion of a Markov Decision Process (MDP) [11]. An MDP can be completely specified by the tuple containing *x*, an estimate for the state; *u*, the control action executed by the agent; *T*, the transition function defining the system and environment dynamics and *R*, the reward stream that an agent can expect to receive. We define each of these terms in greater detail next for both the case of the hot water vessel and the building's thermal mass.

1. State, *x*

Storage vessel. The storage vessel's state can theoretically be quantified with its energy content at any given time. Since the energy content in the vessel is not directly observable, we use the temperature of the water as a proxy for this state. An additional complexity arises here because in most real world scenarios domestic hot water vessels are equipped with only a single sensor. This single sensor fails to provide adequate information about important stratification effects in the vessel. Thus, researchers have frequently resorted to offline model learning due to lacking information on the state of the system and the nonlinear vessel dynamics (e.g. thermodynamic and mixing losses, stratification effects etc.). At any given time, the storage vessel state is given by eq. 1, where *i* is the amount of time since the last reheat cycle, *E* is the energy provided to the vessel at that time and *W* is the water consumption pattern. For this state estimate, a hot water flow meter and a temperature sensor mounted anywhere inside the vessel (preferably above midway point) is required.

$$x_v = [W_v^{t\,:\,t-i},\ T_{amb}^t,\ T_v^t,\ E_v^{t-i}] \tag{1}$$

Building thermal mass. Using (thermostat) temperature in a building to quantify its current state is – at least cosmetically – very similar to the vessel case. As in the case of a water vessel, there are spatial temperature differences in a building's temperature as well. A controller solely relying on the thermostat temperature to execute its control actions while being oblivious of the geometry of a building will fail to consider temperature differences inside a room because of possible solar gains and infiltration or transmission losses. There are however more subtle, complicating differences while controlling a building's thermal mass. Occupants seldom concern themselves with what the exact energy content of their storage vessel is, as long as the water at the outflow exceeds a certain threshold. However, they are very much engaged with the temperature in their living area. This means that state variations allowed in building's thermal mass have to be constrained to much tighter bounds. Furthermore, while hot water draws are sporadic, thereby allowing another dimension of flexibility, room thermostat temperatures have to constantly lie within the same reasonable temperature bounds. The state for the building thermal mass can be defined as in eq. 2, where *i* is the number of previous room temperature recordings considered for the prediction.

$$x_b^t = \left[T_{amb}^t,\ T_{setpoint}^t,\ T_{solar}^t,\ T_{room}^{t-1\,:\,t-i}\right] \tag{2}$$

Occupant behaviour and ambient temperature. The occupant's state is completely unknown in this problem, and it is only through the (usually infrequent) interactions with thermostats or hot water draws that their effect on the system can be estimated. The ambient temperature, on the other hand, is very well recorded for most inhabited locations and predictions can be obtained from a host of web-services. The ambient temperature is important in the state formulation because it affects both the heat losses and the coefficient of performance of the heat pump.

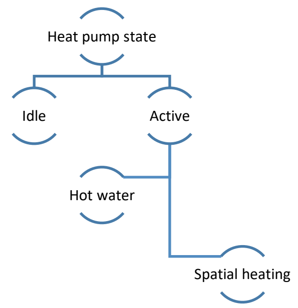

Heat pump. The heat pump can, at any given time, be in one of three possible states. It can either be idle or activated; however if it is in active mode then it can further be providing warmth to the building itself or heating up the hot water storage vessel:

2. Action, *u*

The reinforcement learner (agent) controls the heat pump to decide when to switch between its allowed states (idle to active etc.). Should the need arise to provide both draws at the same time (e.g. in winter), the agent should theoretically make the optimal choice between hot water and spatial heating automatically. However, there are default overrides in the heat pump system under consideration which always prioritize hot water over spatial heating. In case only one or the other is required, the problem can be simplified to a binary decision problem.

3. Transition function, T

The transition function defines the state transitions, given an initial state, $x,$ the agent's control action, u, the influence of observed environmental influences, ε_o (e.g. climate or occupant behaviour) and unobserved environmental influences, ε_n (stochastic dynamics and noise inherent to the system).

$$x' = f_T(x, u) + \varepsilon_o + \varepsilon_n \quad (3)$$

Storage vessel. The transition function depends on the current vessel state (i.e. the energy content embodied in the vessel presently), which in turn is a function of occupant behaviour (e.g. the consumption profile), and when and to what extent the storage vessel was reheated last. The transition function for the vessel state returns a temperature distribution over the entirety of the vessel volume, encapsulating both learnt stratification, as well as thermodynamic and mixing losses.

Building thermal mass. The subsequent states for the building's thermal mass show strongly correlated behaviour with past observations and climate fluctuations that define thermodynamic losses. At the same time however, occupant behaviour and preferences are hidden because of a lack of presence / motion detection sensors. Likewise, the opening and closing of windows is not observable making the state transitions quite stochastic given the observable variables.

Heat pump. The heat pump's state is defined by the control action it is taking. However, unlike an electric resistance heater, a heat pump usually has temporally correlated behaviour when reheating a building or the storage vessel. In this paper, we assume that time steps between subsequent control decisions are large enough to mask this dependence.

4. Reward function, R

The temporal reward stream, like the transition function, is derived from a function that depends on the state of the system and the control action applied (eq. 4).

$$r = f_R(x, u) + \varepsilon_o + \varepsilon_n \quad (4)$$

More concretely, given a system state and depending on the choice of objective function, one or more of the following functions are learnt: (1) the probability, p¸ of lost occupant comfort given the current system state and choice of an action, (2) the energy required, E, (in kWh) to reheat the storage vessel or building to a certain state, given an initial state, and (3) the power profile, P, the heat pump would follow to execute such an operation. Additionally, extrinsic rewards might be provided; one possibility for this is to provide an 'exploration bonus', e, to the agent in case a certain control action can improve either the reward or transition function model. A second possibility is to incorporate additional considerations deriving from market prices [12], flexibility signals [13] or time-of-use tariffs [14] etc.

In all these, occupant comfort is the core constraint for all objectives, violation of which leads to a large negative reward. Efficiency and economy are usually of secondary concern to building occupants as evidenced by the low elasticity of electricity demand [15, 16].

2.2. Planning phase

Once an appropriate representation of the MDP has been learnt, optimization can be performed. One possibility is to apply techniques from stochastic model-predictive control, assuming the learnt representation to be reasonably accurate and the only uncertainty arising from variations in occupant demand etc. However, this goes against the philosophy of a learning system that improves its representation over time. It could also lead to wildly optimistic or pessimistic control actions, depending on the arbitrary starting conditions used to learn the MDP. By interleaving learning with planning, it is possible to achieve all the core benefits of MPC without losing the flexibility of a reinforcement learning system. The simplest way to do this is through heuristics for the planning phase. Derivative-free optimization improves on these heuristics by expanding the search neighborhood and usually returning higher quality solutions.

1. Heuristics

Heuristics for local optimization build off multiple notions of heat pump operation which are not usually optimised. The default behaviour usually follows one such rule-based mechanism which prioritises occupant comfort, but in a way that is oblivious to the environmental dynamics or occupant demand. Reheating the vessel or building every time the thermostat temperature falls below a certain threshold might ensure the occupant seldom suffers discomfort, but is also quite inefficient. The primary idea behind further optimizing this using the learnt MDP is to only reheat when necessary. This minimizes thermodynamic and ambient losses, and also forces the heat pump to operate with lower temperature water at the inflow, thereby further increasing the heat pump COP. By incentivizing reheat cycles when the ambient temperature (and consequently the heat pump COP) is higher, efficiency can be further increased.

2. Derivative-free optimization

Multiple heuristics for optimization using occupant demand and climate conditions can be derived, however these might lead to varying performance levels. As an example, COP is higher when temperature is usually higher so reheating at this time might reduce energy expenditure. However, reheating at this time could be unnecessary (based on historic occupant demand) and would lead to higher thermodynamic losses.

This is just one example of the complexity of the reward landscape. The different reward streams and the nonlinear dynamics of the storage vessel and building mean that the problem is non-convex. In light of this, we use population-based metaheuristics to perform derivative-free optimization for planning using the learnt representations. If the objective is to maximize energy efficiency, then the optimizer initializes a large population of initial solutions (i.e. policies or sequences of control actions). This initialization can be random, drawn from a historic prior or based off heuristics. Afterwards, local neighbourhood search is performed. Different members of the population are responsible for diversification, while the additional search step corresponds to intensification. These two aspects can be handled by any meta-heuristic algorithm or combination of algorithms, such as genetic algorithms [17] and swarm intelligence [18] etc. This leads to a solution, which performs better at achieving the optimal control objectives than pure heuristics-based controllers while avoiding making linearizing assumptions that would be required for a more traditional optimization approach.

2.3. Bringing it all together

The intertwined learning and planning workflows take the following form for optimal control of both the storage vessel and the building's thermal mass:

1. Initialize MDP, $M \leftarrow (X, U, T, R)$, $\pi \leftarrow \pi_{default}$
2. Repeat forever:
   (a) Update experiences: $x_t, u_t, r_{t+1}, x_{t+1}, ...$
   (b) Update the transition function, $T : X \times U \times X \rightarrow \mathbb{R}^2$
   (c) Update the reward function, $R : X \times U \times X \rightarrow \mathbb{R}^2$
   (d) Predict O, the occupant behaviour
   (e) Predict T_a, the ambient temperature
   (f) Simulate possible future scenarios, given x_t
   (g) Update $\pi \leftarrow argmax_u(\mathbb{E}[r + e]_t)$
   (h) Execute π

The transition and reward function can be learnt with an appropriate choice of a function approximation algorithm. Neural networks offer one such possibility. A limitation of standard neural networks for regression is that they don't return the mean and variance for the prediction. For this purpose, we can learn an ensemble of neural networks, which allows us to not just estimate the expected transition or reward but also the uncertainty around this prediction. This measure of uncertainty is important because it can help avoid making over-optimistic projections and guide future search (via exploratory incentives). In doing so, the proposed algorithm not only interleaves learning with planning, but it also incentivizes active exploration which improves the quality of optimal control over time. It is this combination that makes the proposed algorithm flexible enough to perform optimal control of *any* heat pump device.

3. Results

The results section is divided into three parts: (1) learning the hot water systems; (2) learning the building thermal mass, and (3) an application to energy efficiency.

3.1. Hot water

In this section, we investigate the efficacy of the proposed system to learn the hot water vessel's behaviour and its interactions with the occupant and the heat pump.

1. Ambient losses and stratification effects

Ambient loss from the vessel is primarily a function of the vessel state, ambient conditions and physical properties of the vessel. In this work, we estimate the ambient losses for a given vessel state based on measured data. Nonlinear dynamics of water mean that hot water rises to the top of the vessel because of differences in density, but there is also stratification inhibiting mixing between different layers. Learning this information is critical because, in the current regime of incomplete sensing, it might otherwise lead to over-estimates of hot water remaining in the vessel, thereby leading to lost occupant comfort. Fig. 1 illustrates the model's representation for a vessel state given water flow of 100 litres after a reheat cycle to 50℃ for three different cases.

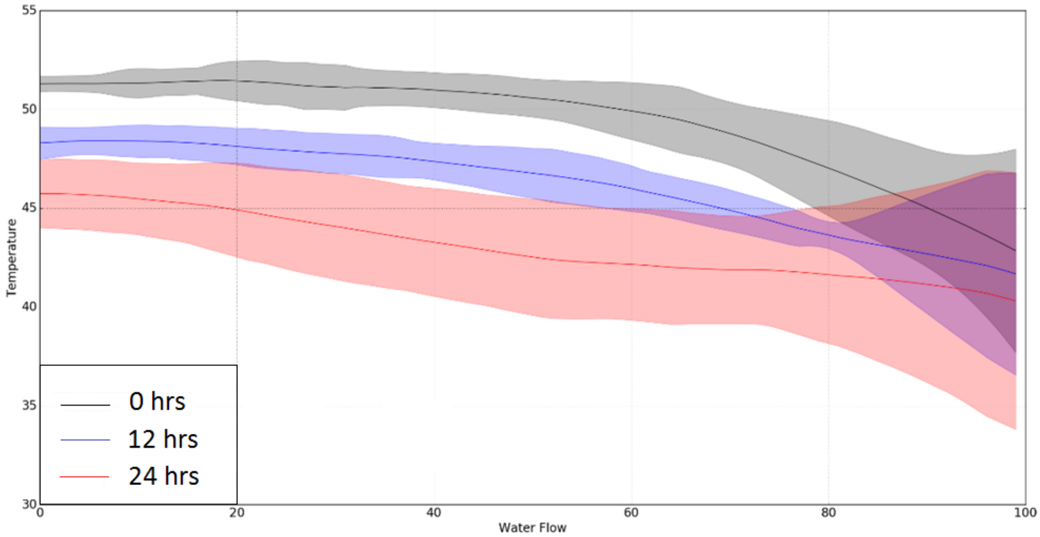

Fig. 1. Stratification and thermodynamic losses, as learnt by the vessel model

Variation in temperature shows the temperature drop with increasing water consumption from the storage vessel. The reduction in the starting temperature for the two curves (red and blue) identify the thermodynamic losses corresponding to an idle period of 12 and 24 hours. The shaded regions correspond to the uncertainty in the model's prediction: the higher the uncertainty, the less certain the agent is in its state estimation or prediction. It is interesting to note that uncertainty is lowest for cases of low flow (i.e. uncertainty increases with consumption) and for less delay since the last reheat cycle (i.e. uncertainty increases as the time between consumption and reheat cycle increases). This is intuitive, since the agent learns its representation from occupant behaviour and it stands to reason that occupants will consume hot water before 24 hours have elapsed since the last reheat cycle. Nevertheless, the generalization potential is outlined and, given sufficient training data and a robust exploration strategy, the agent will be able to learn the correct representation. The reduction in uncertainty over time is given in Fig. 2.

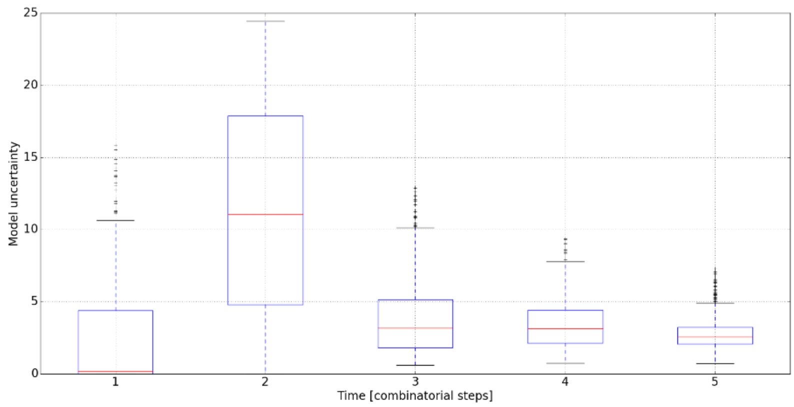

Fig. 2. Temperature vessel model uncertainty reduction over time (the initial high confidence corresponds to neural network initialization parameters)

2. Uncertainty during operation

During physical operation of the storage vessel, a mid-point sensor is made available to both train the state transition model and also to serve as validation (for unseen data). Fig. 3 plots the observations against the predictions, with the shaded bars representing the uncertainty bounds around the prediction.

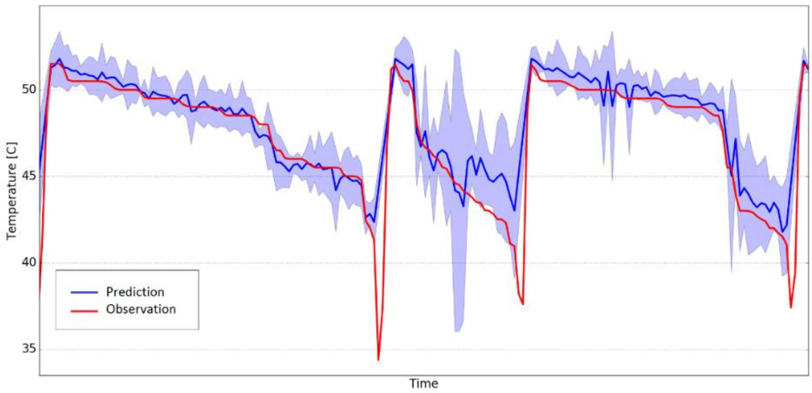

Fig. 3. Predicted vs. observed water temperature in the storage vessel

The error between observation and prediction is usually less than 0.5℃, which is close to the sensor's tolerance, and almost always within the confidence bounds returned by the model. The model has learnt to predict sudden temperature drops also, corresponding to water consumption by the user. The uncertainty for these is usually higher than in more conventional states; however this means that next time the agent encounters such a state, its estimation and prediction capabilities will have improved.

3. Energy consumption by the heat pump

While the current vessel state embodies the energy content present inside the vessel, the electrical energy that would be required by the heat pump to reheat the vessel to an arbitrary temperature must be determined as well. This value represents both the electrical load that has to be minimized over a time horizon while ensuring user comfort, as well as the flexibility potential (capacity) that can be offered to the electric grid at any point in time.

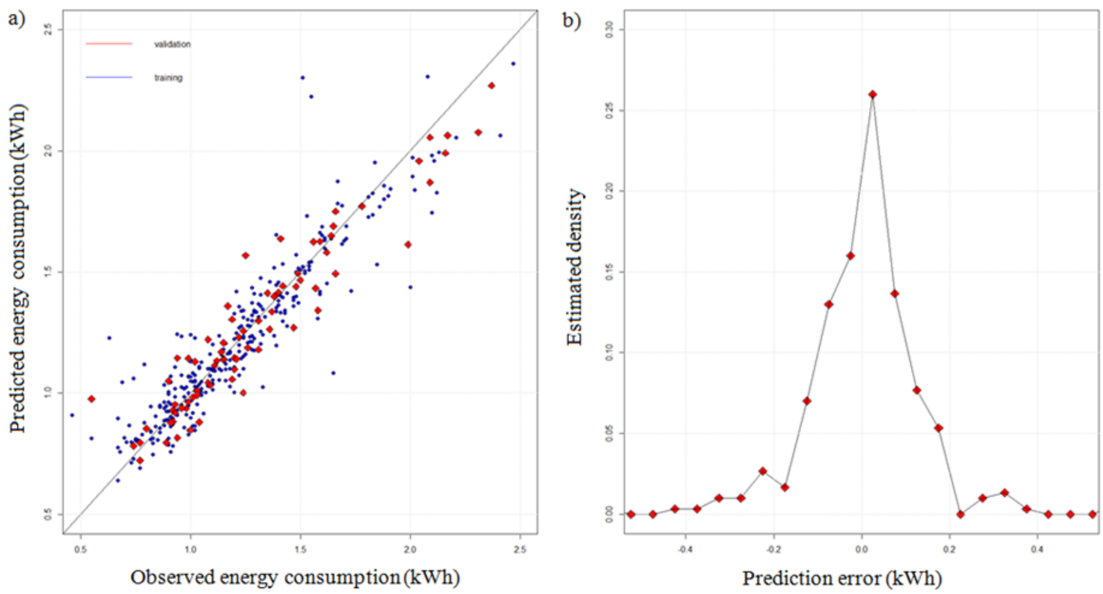

Fig. 4. Predicted vs. observed energy consumption of the heat pump for hot water production

Fig. 4a shows the results of this energy consumption learnt over time from real data, split into training and testing sets. The histogram of errors (Fig. 4b) shows that the prediction error is zero-mean and around 10% in mean absolute percentage error (MAPE) terms.

3.2. Spatial heating

In this section, we take a closer look at the proposed algorithm's performance on thermal mass of buildings.

1. Learning transmission losses, solar gains and the heat pump influence

To demonstrate that the model has learnt an approximately accurate representation of the thermal mass of the building, the model predictions are compared with actual sensor observations over a week. This includes both the transmission losses to the ambient, as well as solar gains leading to temperature variations in the building. Fig. 5 illustrates the results of the model as learnt from historic data. The shaded regions in the plot correspond to time periods in which the building was being heated up by the heat pump. The results are plotted out for two different houses, and illustrate one example where the model was able to accurately learn the building behaviour and another where it wasn't. In the absence of more information, multiple hypothesis can be presented. One likely explanation is that in the first house, the temperature follows a periodic profile which the learner has been able to capture and makes correct temperature predictions. In the second house, the overall temperature variation is lower but more chaotic which means a learner predicting cyclic patterns fails.

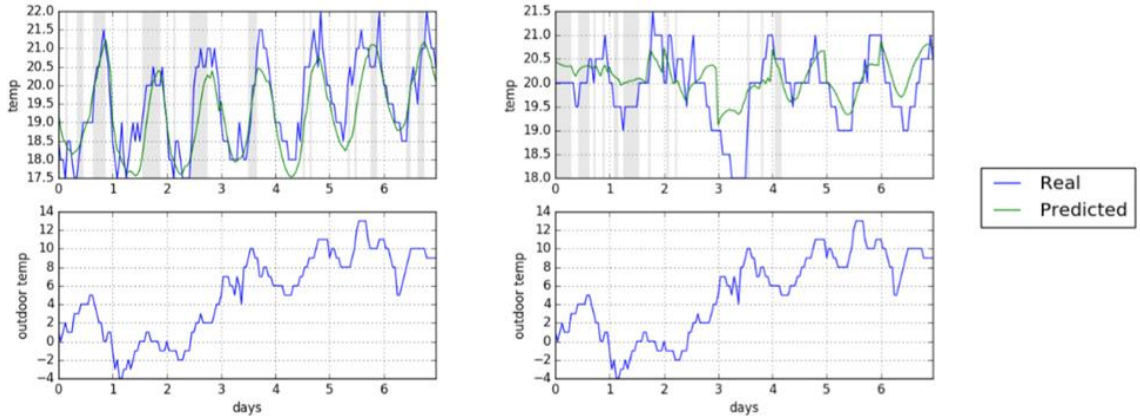

Fig. 5: Predicted vs. observed temperature in two different buildings; ambient temperature is plotted for reference

To investigate this matter further, we tried different parameterizations of learning methods, including linear regression, regularized polynomial regression, neural networks and a combination of multiple learning methods. We found that while a polynomial regression method worked quite well for most houses on average, however by combining multiple methods, we were able to improve the worst case performance at the cost of a slight reduction in average predictive performance. One reason why polynomial regression worked better than the neural network was because of limited amount of data used in training. The comparison is presented in Fig. 6.

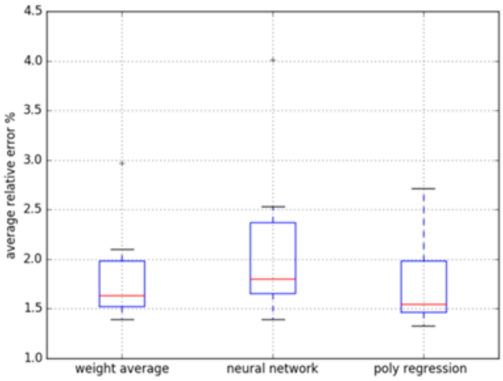

Fig. 6. Predictive performance comparison for multiple learning algorithms

2. Heat pump energy consumption

Similar to the issues encountered in learning a model for the building thermal state, the energy consumption prediction has to be made using incomplete sensor data. Fig. 7 shows the results of the model's performance, again for two different houses with different consumption profiles. The R^2 values for both houses are substantially different, indicating that while such black box models can accurately learn the behaviour for some houses, generalization might be trickier and further investigation into causal effects for this divergence is required.

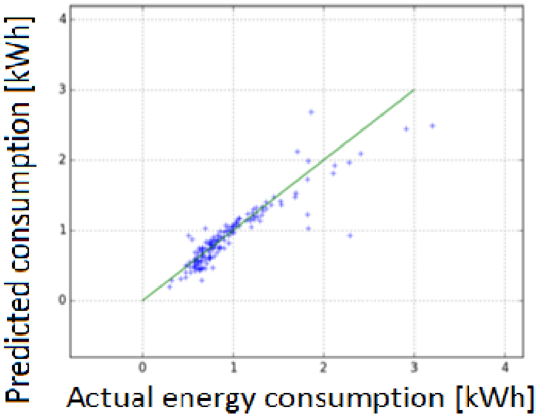

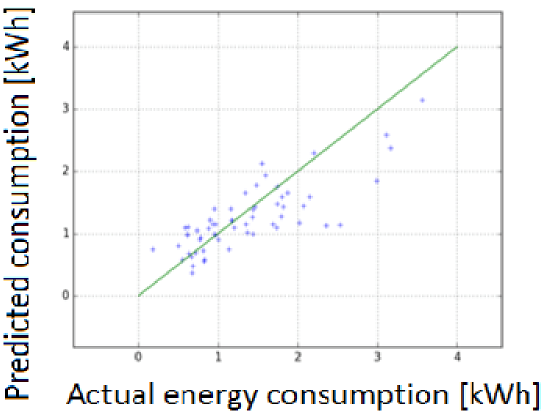

Fig. 7. Predicted vs. observed energy consumption of the heat pump for spatial heating

3.3. Investigating energy efficiency

As explained in section 2, once a reliable model for the system dynamics has been learnt, optimal control can be performed. This optimal control takes the form of searching through the action-space, given a certain starting state. Fig. 8 illustrates results of simulations for providing hot water, where occupant demand was simulated using historic data. Heuristics already demonstrate a sizable reduction in energy consumption over the default threshold-based mechanism, however incorporating the complete reinforcement learning framework with derivative-free optimization leads to even greater savings. In our simulations, we were able to achieve close to 20% energy efficiency gains. However, these rely heavily on the consumption profiles and the temperature defaults. It is therefore important to note here that energy efficiency gains achievable through such a system decrease with increasing hot water consumption and also with more relaxed occupant comfort bounds by default (in which case, the proposed system can improve occupant comfort at similar energy consumption levels).

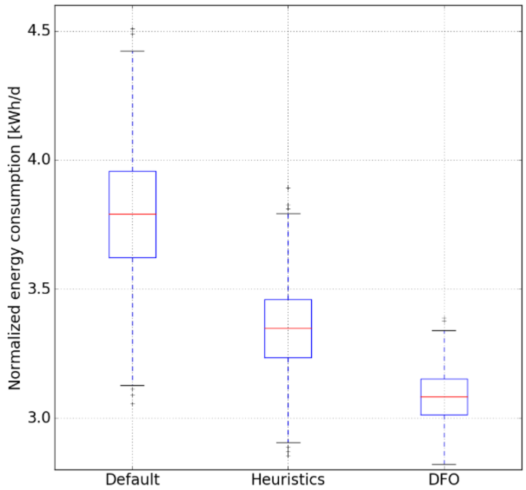

Figure 8: Simulated energy efficiency gains

4. Conclusions

In this paper we have investigated the feasibility of learning a heat pump's behaviour in completely online settings. With no prior information about the heat pump, the storage vessel or building type it was connected to, we were able to learn a usable representation using sensor data. The sensor requirements for the proposed system are minimal, with no additional sensors installed beyond what comes with the default heat pump configuration. Furthermore, being a reinforcement learning system, over time the uncertainty in the system showed a downward trend: as the agent observed more data, its estimates and predictions improved. We have used these models to present a practical application by showing that optimization gains are possible by implementing the full reinforcement learning framework. Such systems can help further reduce the operational costs of heat pump systems, while also making them more amenable to being more active participants of the electricity grid by responding to imbalances and congestions.

Acknowledgements

This work was done with support from VLAIO (formerly IWT) and KIC InnoEnergy.

References

[1] European Commission. Efficiency first: a new paradigm for the European energy system. 2016.
[2] European Commission. The European Union leading in renewables. 2016.
[3] IRENA. The power to change: solar and wind cost reduction potential to 2025. 2016.
[4] Santamouris, M. Energy consumption in the European built environment.
[5] Zero Energy Buildings: a critical look at the definition. P. Torcellini, S. Pless, M. Deru. Pacific Grove, California : s.n., 2006. ACEEE Summer Study.
[6] Majcen, Dasa. Predicting energy consumption and savings in the housing stock: A performance gap analysis in the Netherlands . 2016.
[7] A. Pegels, A. Figueroa. The human factor in energy efficiency. 2015.
[8] US Department of Energy. Modeling and simulation of human behavior in buildings. 2015.
[9] Boyd, S. Model Predictive Control.
[10] Model Predictive Control of Building Heating System. J. Siroky, S. Privara, L. Ferkl. 2010. CLIMA .
[11] R. Sutton, A. Barto. Reinforcement learning: an introduction. s.l. : MIT Press, 1998.
[12] Demand reduction in building energy systems based on economic model predictive control. J. Ma, J. Qin, T. Salsbury, P. Xu. 2012, Chemical engineering science, pp. 92 - 100.
[13] Smart grid task force. Regulatory recommendations for the deployment of flexibility. 2015.

[14] Hogan, W. Time of use rates and real time prices. 2014.
[15] NERA. An econometric assessment of elecitricity demand in the US. 2015.
[16] RMI. The economics of demand flexibility.
[17] E. Anderson, M. Ferris. Genetic Algorithms for combinatorial optimization: the assembly line balancing problem. 1993.
[18] Caro, G. Di. The Ant Colony Optimization (ACO) Metaheuristic: a swarm intelligence framework for complex optimization tasks.