

A review of reinforcement learning for autonomous building energy management[☆]

Karl Mason^{*}, Santiago Grijalva

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA

ARTICLE INFO

Article history:

Received 18 December 2018

Revised 16 May 2019

Accepted 30 July 2019

Available online 2 August 2019

Keywords:

Reinforcement learning

Building energy management

Smart homes

Smart grid

Deep learning

Machine learning

ABSTRACT

The area of building energy management has received a significant amount of interest in recent years. This area is concerned with combining advancements in sensor technologies, communications and advanced control algorithms to optimize energy utilization. Reinforcement learning is one of the most prominent machine learning algorithms used for control problems and has had many successful applications in the area of building energy management. This research gives a comprehensive review of the literature relating to the application of reinforcement learning to developing autonomous building energy management systems. Energy savings of greater than 20% are reported in the literature for more complex building energy management problems when implementing reinforcement learning. The main direction for future research and challenges in reinforcement learning are also outlined.

© 2019 Elsevier Ltd. All rights reserved.

1. Introduction

Meeting the global energy needs in an efficient and sustainable manner is one of the most pressing issues facing society today. The energy consumption from buildings accounts for a large portion of the total global energy consumption. It is therefore imperative to find innovative ways to help reduce and optimize the energy consumed by buildings. Some of the main sources of energy consumption in buildings include: Heating Ventilation and Air Conditioning (HVAC), water heating, and lighting. Reducing the energy consumed by buildings has many benefits. The consumer is benefited with lower energy bills. Electricity providers benefit from reduced peak loads. There is also a wider benefit to society as a whole from reducing building energy consumption. Much of the energy that powers the grid comes from greenhouse gas emitting sources, e.g. coal. Therefore, a reduction in building energy consumption would also correspond to a reduction in the emission of greenhouse gasses.

The prevalence of digital systems in every aspect of modern society is largely due to the steady reduction in the cost and size of micro-processors, along with their increased computing power. The widespread availability of computing has manifested itself in the area of building energy consumption in the form of advanced sensors and actuators, and building energy management systems. These systems consist of monitoring and controlling the activity of each device within the building with the aim of reducing the demand placed onto the grid and the cost to the consumer. The difficulty of this management

[☆] This paper is for CAEE special section SI-aies.Reviews processed and recommended for publication to the Editor-in-Chief by Guest Editor Dr. Mustapha Hatti.

^{*} Corresponding author.

E-mail addresses: kmason35@gatech.edu (K. Mason), sgrijalva@ece.gatech.edu (S. Grijalva).

process is increased with the addition of photovoltaic (PV) panels, batteries, electric vehicles, and smart appliances. Building energy management systems are comprised of many components. Sensors are needed to monitor the building, e.g. temperature, humidity and load sensors. Communication is needed between devices in order to facilitate the monitoring and scheduling of their operation. This communication between devices is an example of the “Internet of Things” (IoT). Smart meters are needed to record energy consumption and to communicate with energy providers. Many devices are equipped with rudimentary control mechanisms such as rule-based decision systems. For example, the thermostat that controls a HVAC system turns heating and cooling on/off when the temperature of the room crosses a user-defined threshold. This paper will provide an overview of how more sophisticated machine learning control systems, in particular Reinforcement Learning (RL), have improved the operation of building energy management systems.

Reinforcement Learning (RL) is a form of machine learning that consists of an agent interacting in an environment, learning what actions to take depending on the state of the environment. The agent learns by trial and error and is rewarded for taking desirable actions. The environment is often modelled as a Markov Decision Process (MDP). RL algorithms date back to the 1960s and 1970s. These algorithms have been applied to a wide range of problems over the years, everything from traffic light control to watershed management. The combination of RL algorithms with deep neural networks has significantly increased the effectiveness of RL methods and has enabled them to be applied to tasks involving computer visions, e.g. self-driving cars. Recent notable achievements of RL have resulted in a significant amount of research into RL and its applications.

Many researchers have since begun applying reinforcement learning to some of the challenging areas of building energy management. RL has been applied to tasks such as HVAC control, water heater control, electric vehicle charging, lighting control and appliance scheduling. An advantage of applying RL to address these problems is that the algorithm learns by itself what the best control policy is. When implementing more traditional rule-based approaches, the designer must handcraft the thresholds that the system will adhere to. This naïve approach is not necessarily able to minimize energy consumption to the extent that RL can. There are several factors that increase the complexity of applying RL to these problems, such as identifying what state information is needed, conflicting objectives and simulator design. These will be discussed in detail throughout the paper. The main contributions of this research are:

1. To give a comprehensive review of the literature relating to the application of RL to building energy management.
2. To quantify the impact of RL on building energy management systems in terms of energy savings.
3. To establish the limitations of RL for building energy management and outline potential areas of future research.

The outline of the paper is as follows. [Sections 2](#) will give outline the research area of building energy management. An overview of reinforcement learning will be provided in [Section 3](#). [Section 4](#) will give a comprehensive account of the applications of reinforcement learning within the area of building energy management. [Section 5](#) will discuss some of the limitations and potential areas for future work concerning the application of reinforcement learning to building energy management. Finally, [Section 6](#) will outline what can be concluded as a result of this research.

2. Autonomous building energy management

The development of autonomous building energy management systems has received a lot of interest in recent years for a number of reasons, one of the primary reasons being the drive to increase the energy efficiency in buildings through operational methods. As stated in the introduction, energy consumption in buildings accounts for a significant portion of the total worldwide energy consumption. There are a number of factors that influence how energy efficient a building is, e.g. insulation, construction materials and overall design. The aim of developing autonomous building energy management systems is to further reduce both overall energy consumption and the cost of powering buildings by utilizing advancements in technologies such as:

- Autonomous control systems
- Internet of Things (IoT)
- Batteries
- Smart grid

2.1. HVAC

Heating Ventilation and Air Conditioning (HVAC) is one of the most energy intensive consumers of energy in buildings. HVAC is now responsible for the majority of building energy consumption in the US. This is due to expectation of thermal comfort. It is also well known that there is a strong correlation between outside temperatures and HVAC energy consumption. This energy consumption is expected to increase due to the increased number of extreme weather events observed globally. The most basic HVAC control system would be a threshold-based approach in which a thermostat is used to regulate the temperature. If the temperature crosses some predefined threshold, the HVAC system activates to heat/cool the environment as required. There have been a wide number of strategies proposed to control the operation of HVAC systems. Temperature is not the only variable that must be considered when implementing building energy management systems. It is also important to consider other factors such as air velocity and humidity in relation to thermal comfort and building energy management. [Section 4](#) will outline previous applications of RL to HVAC control.




Fig. 1. Appliance energy consumption.

2.2. Lighting and appliances

There are of course many other energy intensive devices within the buildings besides HVAC. Much of the energy consumed in domestic dwellings comes from: lighting, televisions, water heaters, washing machines, dryers, fridge freezers, etc. In order to reduce energy consumption of appliances, the design of these devices has been significantly improved in recent years. For example, some dryers are designed to have moisture sensors to prevent over drying. A significant amount of effort has focused on changing consumer behaviour, e.g. turning lights off when the room is not in use or turning off television sets when not in use. It is known that consumer behaviour plays a significant role in energy consumption. Another crucial issue with regards to the energy consumption of appliances is the variation in energy demand throughout the day. This affects the load demand placed on the electrical grid. This will be discussed further in [Section 2.5](#). This provides an opportunity for autonomous building energy management systems to help reduce these peak demands by managing energy consumption, e.g. dimming lights at peak times. Autonomous control systems are now being developed to effectively schedule appliances to minimize energy cost. [Fig. 1](#) shows the power consumed by many of the commonly found domestic appliances.

2.3. Electric vehicles and batteries

The automotive industry is also heavily influencing the behaviour of residential energy consumption due to the current transition from fossil fuel powered vehicles to electric vehicles. The need to draw power from the grid in order to charge Plug in Hybrid Electric Vehicles (PHEVs) significantly increases the electrical load placed on the grid. [Fig. 2](#) illustrate the proportion of power consumed by a typical electric vehicle when compared to the average US power consumption.

Vehicle to grid (V2G) technology is a promising solution to address this problem. Vehicle to grid technology works by selling electricity back to the grid when the car is not being used. A variation V2G is to consider the electric vehicle as a deferrable load. When the grid needs more power, charging of electric vehicle can stop for a few seconds or minutes. In this process there is no actual injection of power from the vehicle to the grid, but a pause in charging that can help the grid. V2G also has the potential to alleviate some of the load on the grid during peak demand. This further motivates the need for effective building energy management systems. For example, a potentially effective strategy would be to charge the vehicle overnight when electricity is cheapest and to sell energy back to the grid at 7 pm when demand is high. Electric vehicles could also contribute to address the problem of incorporating variable renewable energy into the grid. The variability of renewable sources is a challenge to incorporating technologies such as Photovoltaic (PV) into the grid. This will be discussed

Average Yearly US Power Consumption vs Estimated Consumption of EV (kWh)




Fig. 2. Average US power consumption and typical electric vehicle power consumption.

further in [Section 2.5](#). Battery technology in general is expected to play a major role in addressing some of the issues raised above. The Tesla Powerwall is a prime example of this. Energy management systems have already been developed that utilize battery technology to help manage the electrical demand.

2.4. Smart meters and the internet of things

Smart meters are a key development that will enable the large-scale deployment of autonomous building energy management systems. They enable communication with electricity providers that aids their operation. Smart meters also enable the control of appliances within the home. Smart meters are already in widespread use. Italy and Sweden were two of the first to countries to have completed their deployment of smart meters nation-wide, with many other countries making significant progress in their deployment.

The term “Internet of Things” or IoT, refers to a network of devices, appliances, sensors and electronics that can connect with one another. It is this communication between devices that enables building energy management systems to operate.

2.5. Electrical grid and renewables

One of the main challenges when developing autonomous building control systems is integrating it with the electrical grid. These systems need to schedule their operation in order to minimize the cost of electricity to the consumer. This is further complicated with many individuals installing PV solar panels, as previously discussed. This gives rise to the recent paradigm of “prosumers” (producers and consumers). Within the prosumer paradigm, individuals both consume energy from the grid and produce energy which is sold to the grid. This creates issues of instability within the grid which must be addressed. This is due to feeders with high levels of PV penetration producing an excess of electricity that the grid struggles to cope with.

Another factor which increases the complexity of implementing autonomous building control systems is the price of energy. Electricity prices vary based on demand and are therefore considered by many control systems. Weather also has a significant effect on energy consumption. As stated previously, energy consumption via HVAC increases significantly with both exceedingly high and low outside weather temperatures. In addition to this, sunny weather results in a higher output from PV systems, this further exacerbates the problems mentioned relating to grid instabilities.

The next section will outline Reinforcement learning, a popular machine learning method that has been applied to many problems relating to smart grid, building energy management, and smart homes.

3. Reinforcement learning

Reinforcement Learning (RL) is a subgroup of machine learning research that involves an agent learning by itself what actions to take in an environment so that it maximizes some reward [1]. This typically involves a significant amount of trial and error from the agent as it learns what actions result in the highest reward. [Fig. 3](#) illustrates this interaction with its environment. [Algorithm 1](#) presents a generic pseudocode, outlining the broad steps taken in a typical RL algorithm. The agent interacts with its environment in discrete time steps. An agent typically has a policy (π) which determines what




Fig. 3. Agent environment interaction.

Algorithm 1: Generic reinforcement learning algorithm.

```

Initialize Agent
while Episode e < Emax do
  Initialize Environment
  for Time t = 1 to T do
    Observe state s_t
    Select action a ∈ A
    Observe reward r and new state s'
    Update policy π
    Transition to new state s_{t+1}
  end
end
end
  
```

actions it will take. The goal is then to find the optimum policy $\pi^* \in \Pi$, where Π is the set of possible policies. The value of a policy π in a given state s is calculated using the value function in Eq. (1).

$$V^\pi(s) = E\left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s, \pi\right] \quad (1)$$

Where E is the expected future return and γ is the discount factor. In order for the agent to assess the value of its current state, it must also consider the expected future rewards it will obtain if it follows its current policy. Eq. (1) quantifies this mathematically. The discount factor γ determines how much weighting the agent gives to future rewards, as discussed in Section 3.1. The value function of state s for the optimum policy π^* is stated in Eq. (2), calculated using the Bellman equation.

$$V^{\pi^*}(s) = \max E[r_{t+1} + \gamma V^{\pi^*}(s_{t+1}) \mid s, \pi^*] \quad (2)$$

3.1. Problem types

The most common way to model an RL problem is as a Markov Decision Process (MDP). A MDP is a discrete time framework for modelling decision making. A MDP is defined as a tuple $\langle S, A, T, R, \gamma \rangle$, where S is the set of states the environment can be, A is the set of possible actions, $T = Pr_a(s, s') = Pr(s_{t+1} = s' \mid s_t = s, a_t = a)$ is the probability that taking action a in state s will lead to state s' at time $t+1$, and finally $R = R_a(s, s')$ is the reward received after transitioning from state s to s' . RL agents account for future rewards using a discount factor $\gamma \in [0, 1]$, a higher value makes the agent more forward thinking.

An extension to the standard MDP is the Partially Observable Markov Decision Process (POMDP). In POMDPs it is assumed that the underlying problem is a MDP, however the agent does not have complete knowledge of the state of the environment. A POMDP is defined as the tuple $\langle S, A, T, R, \Omega, O, \gamma \rangle$. Here Ω and O are the set of observations and the set of conditional observation probabilities respectively. When the environment transitions to state s' , the agent receives the observation $o \in \Omega$ with probability $O(o \mid s', a)$.

For problems that require multiple agents, RL is often combined with Multi-Agent Systems (MAS). This is referred to as Multi-Agent Reinforcement Learning (MARL). One of the key challenges of MARL is to ensure that agents coordinate their actions to achieve a global optimal result. This can be difficult to achieve in practice as each agent may only have partial information of the environment, making it a decentralized problem. These sorts of problems can be modeled as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP). Dec-POMDPs are defined by the tuple $\langle I, S, \mathbf{A}, T, R, \mathbf{\Omega}, O, h \rangle$. In this case, I is the set of agents, $\mathbf{A} = \times A_i$ is the joint action space, $\mathbf{\Omega} = \times \Omega_i$ is the joint observation space and finally h is the horizon.

Another type of RL problem domain are problems that consist of multiple tasks, i.e. Multi Task Learning (MTL). In this type of problem, each task T_j has a tuple $\langle D, T_j, R_j, O_j \rangle$ where D is the underlying domain, T_j is the task specific transition, R_j is the task specific reward and O_j is the task specific observation. This is closely related to Transfer Learning (TL), in which the goal is to transfer knowledge learned when solving one problem to solve a different problem. The distinction being that in MTL, the tasks are variants of the same underlying problem.

Many RL problems have multiple objectives. These problems are referred to as Multi-Objective RL (MORL) problems. In MORL, the MDP is extended to be a Multi-Objective MDP (MOMDP) where the agent receives a vector reward signal $\mathbf{R}(s, a) = (R_1(s, a), R_2(s, a), \dots, R_m(s, a))$ where m is the number of objectives. The goal here is to learn multiple policies simultaneously for each objective.

The task of learning, itself can be viewed as a learning problem. This is referred to as Meta learning. Meta learning has been applied to RL algorithms. For example the RL² algorithm consists of a fast and slow RL algorithm. The fast RL learner learns how to complete the task while the slow RL learner learns the optimum parameters of the fast learner.

Multi-Agent RL problems can also be either cooperative or competitive. In cooperative MARL, each agents' objective aligns with one another either in the form of a shared goal or goals that do not conflict with one another. In competitive environments, agents have conflicting objectives whereby a gain for one agent results in a loss for another.

3.2. Model-free vs model-based

RL algorithms can be sub categorized into two groups, model-based and model-free algorithms. In model-based algorithms, the agent learns a model of the environment based on its observing how the state of the environment changes when certain actions are taken. These observations are used to estimate the environments state transition function $T(s'|s, a)$ and reward R . Once the algorithm learns a model of the environment, it can be combined with a planning algorithm to decide what actions to take. Examples of model-based algorithms include: Explicit Explore and Exploit (E³), Prioritized sweeping, Dyna and Queue-Dyna.

Model-free approaches do not need to develop a model of the environment. Instead model-free approaches learn a policy, via trial and error, with the aim of approximating the optimum policy. Many of the routinely used RL algorithms are model-free, e.g. Q Learning and SARSA. Model-free approaches are more popular in the literature as they are generally less computationally expensive. Model-based approaches first require building an accurate model of the problem, which can be difficult. Once an accurate model of the environment is obtained, the algorithm must find the optimum policy by planning ahead.

3.3. Discrete and continuous search spaces

Many RL algorithms, such as Q learning, represent the values of each state action pair in a look up table (a Q table for Q learning). This approach works well for problems where the state action space is small. However, a fundamental limitation with the tabular approach for representing Q values is its scalability. It's clear that when the number of states and actions increase, the size of the table quickly becomes exceedingly large. The state action space becomes infinite for problems with continuous variables. The widely adopted solution to this is to replace the Q table with a function approximator, most commonly a neural network. When implementing a neural network to represent the Q function, the network reads in the state of the environment as input. The output of the network is the Q value for each action. This function approximation allows RL algorithms to deal with exceedingly large state spaces such as those with images. The image in Fig. 4 illustrates a deep neural network that can be used as a function approximator for RL algorithms.

Function approximation works well for problems with continuous state spaces. This approach still presents a problem for domains with continuous action spaces. One approach would be to discretize the action space, however this will limit the performance of the RL algorithm. A number of RL algorithms have been proposed to address problems with continuous action variables, including: Deep Deterministic Policy Gradient (DDPG) and Sequential Monte Carlo (SMC).

3.4. Q Learning

One of the most prominent RL algorithms is Q Learning [1]. The Q Learning algorithm has the properties of being an off-policy and model-free reinforcement learning method. Off-policy refers to an agent that learns the value of its policy independent of its actions [1]. When applied to discrete state action spaces, the agents policy is determined by its Q table, a matrix where $Q = S \times A$. Section 3.7 outlines algorithms in which reinforcement learning is combined with neural networks for problems with continuous state spaces. These Q values are used to select what action to take in a given state. Once the




Fig. 4. Neural network function approximator.

agent takes an action (a) for the current environmental state (s), the agent receives a reward (r). The agent then updates its Q table based on Eq. (3).

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r + \gamma \max_{a_t} Q(s_{t+1}, a_t) - Q(s_t, a_t)] \quad (3)$$

Where $\alpha \in [0, 1]$ is the learning rate and $\gamma \in [0, 1]$ is the discount factor. It is also common to implement Q Learning with an ϵ -greedy selection policy where a random action is selected with probability $\epsilon \in [0, 1]$. The purpose of this is to encourage exploration. Another option is to use a softmax action selection policy, outlined in Section 3.6.

3.5. SARSA

Another popular RL algorithm is SARSA (State Action Reward State Action) [2]. As is the case with Q Learning, SARSA is a model-free learning algorithm. Unlike Q Learning, SARSA is an on-policy learning approach meaning that it estimates the value of its policy at a given time while using that particular policy. SARSA updates its Q values using Eq. (4).

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha[r + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)] \quad (4)$$

Where each parameter has the same meaning as in Eq. (3).

3.6. Actor critic

Actor Critic (AC) is a RL algorithm that consists of two agents: an actor and a critic. The actor makes decisions based on its observations of the environment and current policy. The purpose of the critic is to observe both the state of the environment and the reward obtained from the environment based on the actor's decision. The critic then gives feedback to the actor [3]. Actor critic is an on-policy and model-free RL algorithm.

In actor critic learning, the actor takes an action a . This results in a Temporal Difference (TD) error, calculated using Eq. (5).

$$\delta_t = r_{t+1} + \gamma V(s_{t+1}) - V(s_t) \quad (5)$$

Where V is the current estimate of the state value. The critic then updates its estimate of the value function using Eq. (6).

$$V(s_t) \leftarrow V(s_t) + \alpha \delta_t \quad (6)$$

Where α is the learning rate. The actor also uses the TD error to update its preference (p) to select that action again via Eq. (7)

$$p(s_t, a_t) \leftarrow p(s_t, a_t) + \beta \delta_t \quad (7)$$

Where $\beta \in [0, 1]$. The actor then uses this preference in combination with a softmax selection method to choose its action. Its policy is outlined in Eq. (8).

$$\pi(s_t, a_t) = \frac{\exp(\frac{Q(s, a)}{\tau})}{\sum_{b=1}^B \exp(\frac{Q(s, a)}{\tau})} \quad (8)$$




Fig. 5. Actor critic environment interaction.

Where τ is the temperature and is used to calibrate the Softmax selection policy. Lower τ values increase the likelihood of selecting the best estimated action while higher values make each action selection equally likely. As stated in Section 3.4, ϵ -greedy is another popular action selection policy. While Fig. 3 illustrates how an agent generally interacts with its environment, Fig. 5 specifically illustrates how AC agents interact with their environment.

3.7. Recent methods

The algorithms outlined in the previous sections are still routinely used in the literature. There have been some significant advances made in RL that have led to new and powerful techniques. The most significant of these being the combination of RL with deep neural networks as previously mentioned in Section 3.3. One of the more popular recent methods is the Deep Q Network (DQN) algorithm [4]. DQN utilizes both an experience replay of past state action pairs to evaluate policies and a periodic update for the target network. The parameters (θ) of the policy function (Q) are updated by calculating the error between the policy network (Q) and the target network (Q) parameterised by θ^- , as shown in Eq. (9).

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s') \sim U(D)} (r_i + \gamma \max_{a'} \hat{Q}(s', a'; \theta_i^-) - Q(s, a; \theta_i))^2 \quad (9)$$

Where $\mathbb{E}_{(s,a,r,s') \sim U(D)}$ is an episode (s, a, r, s') sampled uniformly from the dataset D of previous experiences. The DQN algorithm was applied to a range of Atari games and found to exceed human level performance.

Another recent RL algorithm is the Asynchronous Advantage Actor Critic (A3C) algorithm [5]. As mentioned in Section 3.6, actor critic methods consists of a policy $\pi(s)$ (actor) and a value function $V(s)$ (critic). This algorithm utilizes multiple learners that interact with their environment at the same time. The “Advantage” portion of the A3C algorithm consists of estimating how much better an action was than expected, calculated as $A = R - V(s)$. The “Asynchronous” portion of A3C comes from the fact that each learner is running individually in parallel. This dramatically speeds up learning. Each learner is initialized to the global policy. They then interact with each individual environment and calculate the value and loss. This is then used to calculate the gradients for the policy (Eq. (10)) and value function (Eq. (11)).

$$d\theta \leftarrow d\theta + \nabla_{\theta'} \log \pi(a_i | s_i; \theta') (R - V(s_i; \theta'_v)) \quad (10)$$

$$d\theta_v \leftarrow d\theta_v + \partial (R - V(s_i; \theta'_v))^2 / \partial \theta'_v \quad (11)$$

Where θ and θ_v are the parameters of the global policy and value function, and θ' and θ'_v are the parameters of the thread specific individual policy and value function. These are used to asynchronously update the global θ and θ_v .

DQN and A3C are two state of the art deep RL methods. The aim of this section was to give an overview of some recent developments in RL and is by no means a comprehensive review of RL as this is outside of the scope of this paper. The next section will give a comprehensive account of the applications of RL to building energy management systems.

4. Autonomous building energy management via reinforcement learning

An overview of both the problem of building energy management and RL was provided in Sections 2 and 3. This section will now provide a comprehensive review of the literature that combines these two areas of research.

4.1. HVAC

Reinforcement learning has been successfully applied to many areas of building energy management, including HVAC control. Typically, the environment states for the RL algorithm include factors such as: time of day, outdoor temperature, indoor temperature, weather forecast and occupancy. Typical RL actions include: temperature set points, air flow control, heating control and cooling control. The RL algorithm also needs some sort of reward to operate, all previous studies in the literature calculate rewards based on energy cost, thermal comfort or a combination of both.

A 2006 study applied a combined RL – Model Predictive Control (MPC) approach to control HVAC operation at the Energy Resource Station Laboratory building in Ankeny, Iowa [6]. It was found that the combined approach performed better than either Q learning or MPC individually. An actor critic – neural network learning approach was applied to adjust the signal of a local control for HVAC control in 2008 by Du and Fei [3]. This study reports significant improvements from a combined PID actor critic learning approach than a stand-alone PID controller. A 2013 study by Urieli and Stone apply Tree Search to control a HVAC heat pump [7]. The Tree Search method resulted in a 7–14% saving when compared to the standard rule-based control. In 2015 Barrett and Linder applied Q learning to the problem of HVAC control combined with Bayesian Learning for occupancy prediction [1]. The results outline a 10% energy savings improvement over a programmable control method. This work was then extended to apply parallel Q learning to HVAC control. Another 2015 study by Ruelens et al. implemented a combined Auto-Encoder (AE) Q learning to control a heat pump for a HVAC system [8]. The authors reported energy savings in the range of 4–11%, comparable to the previous study.

Yang et al. implemented a Batch Q learning with a neural network to control a PV powered heating system [9]. The results report a 10% improvement on a standard rule-based system. In 2017 Wei et al. utilized deep neural network and applied Deep RL (DQN) to the problem of HVAC control and report energy saving improvements over conventional Q learning in the range of 20–70% [4]. Another 2017 study by Wang et al. applied a Monte Carlo Actor Critic with Long Short Term Memory (LSTM) neural network to the task of HVAC control [10]. Their results indicate a 15% thermal comfort improvement and a 2.5% energy efficiency improvement when compared to other methods. More recently, Marantos et al. proposed using a Neural Fitted Q-Iteration (NFQ) approach for HVAC control in 2018 [11]. This approach also used a neural network to approximate the Q function. The results of this study report significant savings in terms of both energy consumption and thermal comfort when compared to rule-based controllers. In 2018, Zhang et al. applied Asynchronous Advantage Actor Critic (A3C) (a deep RL approach) to control the HVAC system for a simulation of the Intelligent Workplace building in Pittsburgh, USA [5]. The authors report a 15% energy saving improvement upon the base case. Chen et al. applied Q learning to control both the HVAC and window systems [12]. The authors report energy savings of 13 and 23% and lowered discomfort ratings by 62 and 80% in the two buildings tested. In a 2018 study, Patyn et al. compared the performance difference of Q learning on three different neural network architectures (convolutional neural networks (CNN), LSTM networks and a feed forward multi-layer neural network) for HVAC control [13]. Their results indicate that the LSTM and multi-layer network perform best however the LSTM training time was twice as long.

This section has outlined numerous successful applications of RL to HVAC control with a variety of RL algorithms. The most common RL algorithm implemented was Q learning. The vast majority of studies report energy savings in the range of $\approx 10\%$ when compared to rule-based approaches. There is also a recent trend to apply deep RL to HVAC control, which has outperformed traditional tabular RL. This trend is unsurprising given the recent success and attention that deep learning has received.

4.2. Water heater

As discussed in Section 2, water heaters consume a significant amount of energy. This section will explore some of the applications of RL to control water heaters with the aim of reducing energy costs. Some of the state variables include: time of day, current water temperature and forecast usage. The action that the RL agent makes is generally to turn the heater on or off. The reward given to the agent is the electricity consumption.

A 2014 study by Al-Jabery et al. applied Q learning and a fuzzy Q learning to control a water heater [14]. Their results indicate that their proposed fuzzy Q learning algorithm gives smoother convergence than the standard Q learning. Al-Jabery et al. then explored applying an Actor Critic – Q learning approach to control a water heater in 2017, where they included the grid load in the state space for the RL agent [15]. This study reports energy cost savings between 6 and 26% using RL. In a more recent 2018 study, Ruelens et al. applied Q iteration to a physical water heater in a laboratory setting and found that Q iteration was able to reduce the energy consumption by 15% over 40 days when compared to a thermostat controller [16]. De Somer et al. applied Q learning to control the heating cycle of a domestic water heater to maximize the consumption from a local PV source. The authors then evaluate the Q learning agent on 6 residential houses and report a 20% increase in the amount of local PV consumption [17]. Kazmi et al. did a similar study in the Netherlands where deep RL was applied to 32 homes [18]. The authors report a 20% reduction in the energy consumption for water heating and no loss in user comfort.

Although there are significantly fewer studies in the literature applying RL to water heaters than HVAC. Based on the papers reviewed, there appears to be a significant reduction in energy consumption when applying RL to water heaters. Many of the studies report reductions of $\approx 20\%$ when compared to the baseline. A similar trend can be observed that recent research is more focused on the application of deep RL as it has demonstrated its effectiveness.

4.3. Home management systems

The studies outlined in Sections 4.1 and 4.2 consider applying RL to HVAC and water heaters respectively in isolation. This section explores applications of RL within the home to manage multiple appliances, lighting, PV, battery, etc. Applying RL in this manner is needed as the task of building energy management is a complex problem with multiple factors that must be managed to reduce the overall energy consumption. Similar to the previously mentioned sections, the states for the RL algorithm when exploring building energy management in a more holistic manner, generally consist of the time of day, temperature information and the current usage state of the various appliances. Many of these studies also include other state information such as electricity prices, grid load and information in relation to PV panels such as solar irradiance. The actions available to the RL agent in the studies outlined in this section are similar to those outlined in the previous sections, i.e. turning a device or appliance on or off. In the cases of studies which incorporate batteries, the actions are to charge/discharge the battery or to do nothing.

In 2018, a study by Reymond et al. applied Fitted Q Iteration to learn to schedule a number of household appliances including a heat pump, water heater and dishwasher [19]. Their results indicate that independent learning performs 9.65% better than a centralized learning approach. Wei et al. implemented a dual iterative Q Learning algorithm for residential battery management [20]. The authors report a 32.16% savings in energy cost when compared to the baseline. A 2015 study by Guan et al. utilize temporal difference learning to control the energy storage of a battery in the presence of a PV panel [21]. It was found that temporal difference learning resulted in improvements of 59.8% reduction in energy costs. Wan et al. implemented actor critic learning with two deep neural networks as a residential energy management system [22]. In this study, the RL algorithm learned when to purchase electricity from the utility provider based on previous prices, house hold loads and current battery charge. The authors report an 11.38% reduction in energy cost by using actor critic learning when evaluated over 100 days. Remani et al. applied Q learning to schedule multiple devices within the home including lighting, clothes dryer, dish washer, etc. [23]. The authors implemented a price based demand response and included a PV panel in the system model and reported an approximate daily energy cost saving of 15%. Wen et al. proposed a demand response energy management systems for small buildings that enables automated device scheduling in response to electrical price fluctuations [24]. The authors implement Q learning with the aim of taking advantage of the estimated 65% of potential energy savings for small buildings by efficient device scheduling and report improvements upon the baseline. Bazhenkov and Goubko utilize inverse reinforcement learning (IRL) to predict consumer appliance usage and report a higher accuracy using IRL than other machine learning methods such as random forest [25]. A 2018 study by Mocanu et al. implement both Deep Q Learning (DQL) and Deep Policy Gradients (DPG) to optimize the energy management system for 10, 20 and 48 houses [26]. This study explored the use of electric vehicles, PV panels and building appliances. The authors report electricity cost savings of 27.4% for DPG and 14.1% for DQL.

A combined approach was proposed by Wang et al. that implements Q learning to manage a battery in a residential setting with installed PV [27]. Q learning was combined with a model-based control algorithm to minimize energy costs. This hybrid approach demonstrated significant cost savings in the range of 23–72%. Rayati et al. utilize Q learning for home energy management in a setting with PV installation and energy storage [28]. This study also considers home comfort and CO₂ emissions when learning the optimal control policy. The authors report a maximum energy saving of 40%, peak load reduction of 17% and CO₂ social cost reduction of 50%. Sheikhi et al. implement Q learning to control a smart energy hub that consists of a combined heat and power, auxiliary boiler, electricity storage and heating storage [29]. The authors report a saving of 30% and 50% for energy cost and peak loads respectively.

The results outlined in the literature relating to the application of RL to home energy management varies significantly more than in the previous two sections. The highest reported energy cost reduction was 72% when implementing RL with a model-based controller [27]. There is a greater capacity to generate savings when applying RL to home energy management systems with batteries and solar than simply applying RL to devices. This is reflected in the studies mentioned here. In terms of algorithms, Q learning is again the most routinely used method.

4.4. Smart homes and the electrical grid

While the studies in the previous sections have demonstrated the effectiveness of RL when applied within the home, many studies in the literature examine how RL can also have a positive impact when considering how these smart homes are incorporated into the grid. This section will outline the applications of RL that relate to smart homes and the grid. Many of these studies evaluate control systems for a community of homes, rather than single homes in isolation, and therefore utilize multi-agent RL in many cases. There are many studies that focus on the application of RL to the operation of a microgrid to address problems such as thermal generator scheduling and incorporating PV into a microgrid. These are outside of the scope of this paper and will therefore not be discussed. This section solely focuses on studies that explore incorporating smart homes into the electrical grid.

Jiang et al. implemented a hierarchical multi-agent Q learning approach for dynamic demand response and distributed energy resource management in a microgrid [30]. This microgrid also includes wind power and a battery. This study reports a 19% reduction in energy costs for the whole community. Kim et al. implemented a multi-agent Q learning approach for residential energy consumption based on electricity prices in a microgrid [31]. Significant savings were reported over the baseline myopic purchasing strategy. A multi-agent RL approach was implemented by Anvari-Moghaddam et al. to integrate

intelligent homes into the grid where batteries and renewables are also considered [32]. The authors' proposed structure consisted of: a central coordinator agent, a building management agent, a renewables agent, a battery agent and a services agent. The authors report an energy cost saving of 5%. A recent 2018 paper by Prasad and Dusparic implement a multi-agent deep RL approach to enable homes to share energy with one another in order to minimize cost for the community as a whole [33]. The authors implemented a multi-agent DQN approach where each agent controlled a house. The agent could consume/store energy, request energy from or grant energy to neighbours, deny energy to neighbours or purchase from the grid. Each home was had an installed battery. The results indicate savings of 97 and 156kWh during the winter and summer respectively for a 10 home community.

A 2018 study by Kofinas et al. implemented fuzzy Q learning to control the operation of a battery and desalination plant in conjunction with a PV panel for a microgrid that also includes residential electricity consumers [34]. Claessens et al. applied Q learning with a convolutional neural network for residential load control [35]. It was found that the proposed approach provided near optimum performance when compared to a theoretically optimal benchmark.

The studies outlined in this section focus on how RL can reduce energy costs for multiple smart homes and how RL can help integrate these homes into the grid. As was observed in the previous sections, the studies here also demonstrate the effectiveness of RL for these problems. Multi-agent RL is a commonly used approach to learn policies for multiple homes as it is well suited to the distributed nature of this control task. Based on these studies, electricity cost savings can range from 5% to 19%. In terms of utilizing RL to balance residential loads with the grid, it was demonstrated that RL performs with near optimum performance and significantly reduces electricity costs for the grid.

5. Discussion

A common theme across all of the various applications of RL within building energy management systems is that RL can provide significant savings in each problem it is applied to. As new and effective RL algorithms are developed, these new methods gradually make their way into the smart homes literature, e.g. deep RL methods. This section will outline some of the limitations of RL for building energy management and also some directions for future research.

5.1. Limitations

There are however limitations of applying RL to building energy management systems. The overwhelming majority of studies outlined in this paper discuss applications of RL to simulated versions of problems relating to building energy management. While this is perfectly acceptable and a natural way to apply RL to these problems, this approach relies heavily on accurate simulator design and data that is representative of real-world scenarios. Since RL is an online learning algorithm, it could be directly applied to a building energy management system for a physical building without ever learning within a simulated environment. The problem with this approach is that in order for the RL algorithm to learn an effective control policy that minimizes energy cost, it must learn by trial and error. This would result in an initial exploration period where the RL agent would evaluate different policies, many of which would have an unacceptably high energy cost associated with it. Accurate simulators are needed for this reason, so that the RL agent can learn in simulation which policies are best. These pre-trained agents can then further refine their policies when implemented in physical systems. This is a viable approach as was demonstrated by the application of RL to the HVAC operation at the Energy Resource Station Laboratory building in Ankeny, Iowa [6].

5.2. Future directions

The first and most obvious trend that will be experienced relating to the application of RL to building energy management is the move towards deep RL algorithms. These deep RL methods are capable of learning more complex policies that are more sophisticated than those represented by shallow neural network or look up tables. As the volume of data collected by sensors will only increase, deep RL methods will become necessary to develop effective policies when interacting in environments with very large state action spaces. It is worth noting here that increasingly powerful computers will also enable researchers to train RL agents with increasingly complex policies.

Another potential route for future research in RL for building energy management is to apply many of the variants of traditional RL, some of which were outlined in Section 3.1. For example, based on our extensive literature review, there are no examples within the literature that apply multi task RL to building energy management. Multi-task learning could be potentially useful for learning different HVAC control policies depending on how many occupants are in the house. Transfer learning could be useful for transferring knowledge learnt when learning to schedule appliances to scheduling water heaters.

Some of the problems relating to building energy management can be framed as multi-objective RL problems. The most obvious example of this is the trade of between thermal comfort and energy cost. There are numerous examples in the literature that apply multi-objective optimization algorithms to address this problem, however none of these studies utilize any state of the art multi-objective RL algorithms, e.g. Pareto Q learning.

An additional avenue for future research would be the application of meta learning to RL problems in building energy management, e.g. the RL² algorithm mentioned in Section 3.1. There are no applications of meta reinforcement learning

to building energy management. Methods such as RL² could be useful for learning more effective control policies for the various control tasks within the studies outlined in this paper.

In terms of the tasks that RL is applied to, many studies in the literature are beginning to explore implementing multi-agent RL to control a community of homes rather than a single home in isolation. Developing control policies for groups of homes has the added benefit of allowing homes to pool their resources in a manner that results in greater savings for everyone. A natural route for this research to follow into the future would be to apply multi-agent RL to larger communities of homes to explore the scalability of multi-agent RL in this domain. Conducting game theory analysis in this area would also be a promising area of research. Decentralized multi-agent control policies for large communities of homes would rely heavily on cooperation to achieve cost savings for all, it would therefore be worthwhile to explore the issue of cooperation and defection in such systems.

Based on the extensive literature review conducted, all of the current research relating to applications of RL in building energy management evaluates systems in which the environment does not significantly change outside of its normal environmental conditions. This raises the question of how well do these learned control policies perform when faced with significant environmental changes, e.g. extreme weather events, battery degradation or failure, PV panel installation or failure, homes where the number of occupants increases or decreases, etc. It is not clear based on the current research, how the learned control policies can adapt to these significant environmental changes, i.e. are the learned policies brittle and would result in catastrophic failure with sub optimal performance in the event of these changes? If the RL agent is robust and capable of adapting to such environmental changes, how long would it take for the RL agent to relearn new policies in such scenarios? These questions would be worth investigating in depth in future studies.

Another potential future direction of research would be to conduct an experimental comparison of multiple RL algorithms along with other control algorithms. This would provide a more complete picture of the performance of RL in relation to other methods. There have been many other approaches proposed in the literature for control problems, e.g. Model Predictive Control, Evolutionary Neural Networks, Bayesian methods, Swarm Intelligence, Fuzzy Control, Radial Basis Functions and Genetic Programming. It would be insightful to compare the performance of these methods to RL. There are also a vast number of algorithms in the literature that focus on energy optimization problems rather than control, e.g. Genetic Algorithms and Particle Swarm Optimization. Future research could investigate combining such methods with RL for building energy management.

6. Conclusion

The key concepts behind developing building energy management and an introduction into reinforcement learning have been presented. The primary contribution of this research was to provide a comprehensive account of the applications of reinforcement learning to building energy management. As a result of this literature review, it is clear that:

1. Reinforcement learning algorithms significantly improve the energy efficiency of homes. Energy savings vary significantly depending on the specific application. Broadly speaking RL can typically provide savings of $\approx 10\%$ for HVAC applications, $\approx 20\%$ for water heaters and $> 20\%$ for more complete buildings energy management systems.
2. The vast majority of studies that implement RL for problems relating to building energy management are in simulation only. With accurate simulator design however, RL is a viable approach for building energy management outside of simulation.
3. Much of the recent research for RL and building energy management is focused on applications on deep RL algorithms due to their increased effectiveness over traditional approaches. This is likely to become more pronounced into the future.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

- [1] Barrett E, Linder S. Autonomous hvac control, a reinforcement learning approach. In: Joint European conference on machine learning and knowledge discovery in databases. Springer; 2015. p. 3–19.
- [2] Rummery GA, Niranjan M. On-line Q-learning using connectionist systems, 37. University of Cambridge, Department of Engineering Cambridge, England; 1994.
- [3] Du D, Fei M. A two-layer networked learning control system using actor–critic neural network. *ApplMathComput* 2008;205(1):26–36.
- [4] Wei T, Wang Y, Zhu Q. Deep reinforcement learning for building hvac control. In: Proceedings of the 54th annual design automation conference 2017. ACM; 2017. p. 22.
- [5] Zhang Z, Chong A, Pan Y, Zhang C, Lu S, Lam KP. A deep reinforcement learning approach to using whole building energy model for hvac optimal control. In: 2018 building performance analysis conference and SimBuild; 2018.
- [6] Liu S, Henze GP. Experimental analysis of simulated reinforcement learning control for active and passive building thermal storage inventory: part 2: results and analysis. *Energy Build* 2006;38(2):148–61.
- [7] Urieli D, Stone P. A learning agent for heat-pump thermostat control. In: Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems. International Foundation for Autonomous Agents and Multiagent Systems; 2013. p. 1093–100.
- [8] Ruelens F, Iacovella S, Claessens BJ, Belmans R. Learning agent for a heat-pump thermostat with a set-back strategy using model-free reinforcement learning. *Energies* 2015;8(8):8300–18.

- [9] Yang L, Nagy Z, Goffin P, Schlueter A. Reinforcement learning for optimal control of low exergy buildings. *Appl Energy* 2015;156:577–86.
- [10] Wang Y, Velswamy K, Huang B. A long-short term memory recurrent neural network based reinforcement learning controller for office heating ventilation and air conditioning systems. *Processes* 2017;5(3):46.
- [11] Marantos C, Lamprakos CP, Tsoutsouras V, Siozios K, Soudris D. Towards plug & play smart thermostats inspired by reinforcement learning. In: *Proceedings of the workshop on INTElligent embedded systems architectures and applications*. ACM; 2018. p. 39–44.
- [12] Chen Y, Norford LK, Samuelson HW, Malkawi A. Optimal control of hvac and window systems for natural ventilation through reinforcement learning. *Energy Build* 2018;169:195–205.
- [13] Patyn C, Ruelens F, Deconinck G. Comparing neural architectures for demand response through model-free reinforcement learning for heat pump control. In: *2018 IEEE international energy conference (ENERGYCON)*. IEEE; 2018. p. 1–6.
- [14] Al-Jabery K, Wunsch DC, Xiong J, Shi Y. A novel grid load management technique using electric water heaters and q-learning. In: *Smart grid communications (SmartGridComm)*, 2014 IEEE international conference on. IEEE; 2014. p. 776–81.
- [15] Al-Jabery K, Xu Z, Yu W, Wunsch DC, Xiong J, Shi Y. Demand-side management of domestic electric water heaters using approximate dynamic programming. *IEEE Trans Comput-Aided Des Integr Circuits Syst* 2017;36(5):775–88.
- [16] Ruelens F, Claessens BJ, Quaiyum S, De Schutter B, Babuška R, Belmans R. Reinforcement learning applied to an electric water heater: from theory to practice. *IEEE Trans Smart Grid* 2018;9(4):3792–800.
- [17] De Somer O, Soares A, Vanthournout K, Spiessens F, Kuijpers T, Vossen K. Using reinforcement learning for demand response of domestic hot water buffers: A real-life demonstration. In: *Innovative smart grid technologies conference Europe (ISGT-Europe)*, 2017 IEEE PES. IEEE; 2017. p. 1–7.
- [18] Kazmi H, Mehmood F, Lodewyckx S, Driesen J. Gigawatt-hour scale savings on a budget of zero: deep reinforcement learning based optimal control of hot water systems. *Energy* 2018;144:159–68.
- [19] Raymond M, Patyn C, Rădulescu R, Deconinck G, Nowé A. Reinforcement learning for demand response of domestic household appliances. In: *Proceedings of the adaptive and learning agents workshop (at AAMAS 2018)*; 2018.
- [20] Wei Q, Liu D, Shi G. A novel dual iterative q-learning method for optimal battery management in smart residential environments. *IEEE Trans Ind Electron* 2015;62(4):2509–18.
- [21] Guan C, Wang Y, Lin X, Nazarian S, Pedram M. Reinforcement learning-based control of residential energy storage systems for electric bill minimization. In: *Consumer communications and networking conference (CCNC)*, 2015 12th annual IEEE. IEEE; 2015. p. 637–42.
- [22] Wan Z, Li H, He H. Residential energy management with deep reinforcement learning. In: *2018 international joint conference on neural networks (IJCNN)*. IEEE; 2018. p. 1–7.
- [23] T R, Jasmin EA, Ahamed TPI. Residential load scheduling with renewable generation in the smart grid: a reinforcement learning approach. *IEEE Syst J* 2018;1–12. doi:10.1109/JSYST.2018.2855689.
- [24] Wen Z, O'Neill D, Maei H. Optimal demand response using device-based reinforcement learning. *IEEE Trans Smart Grid* 2015;6(5):2312–24.
- [25] Bazenkov N, Goubko M. Advanced planning of home appliances with consumers preference learning. In: *Russian conference on artificial intelligence*. Springer; 2018. p. 249–59.
- [26] Mocanu E, Mocanu DC, Nguyen PH, Liotta A, Webber ME, Gibescu M, et al. On-line building energy optimization using deep reinforcement learning. *IEEE Trans Smart Grid* 2018.
- [27] Wang Y, Lin X, Pedram M. A near-optimal model-based control algorithm for households equipped with residential photovoltaic power generation and energy storage systems. *IEEE Trans Sustainable Energy* 2016;7(1):77–86.
- [28] Rayati M, Sheikhi A, Ranjbar AM. Optimising operational cost of a smart energy hub, the reinforcement learning approach. *Int J Parallel Emergent Distrib Syst* 2015;30(4):325–41.
- [29] Sheikhi A, Rayati M, Ranjbar A. Dynamic load management for a residential customer; reinforcement learning approach. *Sustainable Cities Soc* 2016;24:42–51.
- [30] Jiang B, Fei Y. Smart home in smart microgrid: a cost-effective energy ecosystem with intelligent hierarchical agents. *IEEE Trans Smart Grid* 2015;6(1):3–13.
- [31] Kim B-G, Zhang Y, Van Der Schaar M, Lee J-W. Dynamic pricing and energy consumption scheduling with reinforcement learning. *IEEE Trans Smart Grid* 2016;7(5):2187–98.
- [32] Anvari-Moghaddam A, Rahimi-Kian A, Mirian MS, Guerrero JM. A multi-agent based energy management solution for integrated buildings and micro-grid system. *Appl Energy* 2017;203:41–56.
- [33] Prasad A., Dusparic I. Multi-agent deep reinforcement learning for zero energy communities. *arXiv:181003679v1* 2018.
- [34] Kofinas P, Vouras G, Dounis AI. Energy management in solar microgrid via reinforcement learning using fuzzy reward. *Adv Build Energy Res* 2018;12(1):97–115.
- [35] Claessens BJ, Vranckx P, Ruelens F. Convolutional neural networks for automatic state-time feature extraction in reinforcement learning applied to residential load control. *IEEE Trans Smart Grid* 2018;9(4):3259–69.

Karl Mason is a Postdoctoral Research Fellow at The Georgia Institute of Technology. He completed both his Ph.D. and M.Sc. at The National University of Ireland, Galway in 2018 and 2015 respectively. His research interests span a range of topics including: machine learning, neural networks, reinforcement learning, evolutionary computing and their applications.

Santiago Grijalva is the Southern Company Distinguished Professor of Electrical and Computer Engineering and Director of the Advanced Computational Electricity Systems Laboratory (ACES) at The Georgia Institute of Technology. He was previously the Founding Director of the Power System Engineering Center at NREL. He received his PhD (2002) and MS (1999) from the University of Illinois at Urbana Champaign.