

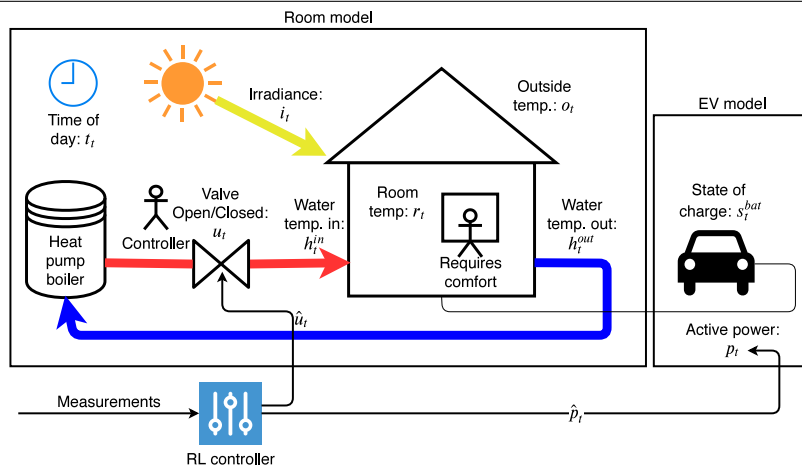
# Data-driven control of room temperature and bidirectional EV charging using deep reinforcement learning: Simulations and experiments

B. Svetožarević<sup>a,\*</sup>, C. Baumann<sup>a,b</sup>, S. Muntwiler<sup>b</sup>, L. Di Natale<sup>a</sup>, M.N. Zeilinger<sup>b</sup>, P. Heer<sup>a</sup>

<sup>a</sup> Urban Energy Systems Laboratory, Empa, Switzerland

<sup>b</sup> Institute for Dynamic Systems and Control, ETH Zurich, Switzerland

## GRAPHICAL ABSTRACT



## ARTICLE INFO

### Keywords:

Data-driven building control  
Deep reinforcement learning  
Room temperature control  
Thermal comfort  
EV charging  
Recurrent neural networks

## ABSTRACT

The control of modern buildings is a complex multi-loop problem due to the integration of renewable energy generation, storage devices, and electric vehicles (EVs). Additionally, it is a complex multi-criteria problem due to the need to optimize overall energy use while satisfying users' comfort. Both conventional rule-based (RB) controllers, which are difficult to apply in multi-loop settings, and advanced model-based controllers, which require an accurate building model, cannot fulfil the requirements of the building automation industry to solve this problem optimally at low development and commissioning costs. This work presents a fully data-driven pipeline to obtain an optimal control policy from historical building and weather data, thus avoiding the need for complex physics-based modelling. We demonstrate the potential of this method by jointly controlling a room temperature and an EV to minimize the cost of electricity while retaining the comfort of the occupants. We model the room temperature with a recurrent neural network and use it as a simulation environment to learn a deep reinforcement learning (DRL) control policy. It achieves on average 17% energy savings and 19% better comfort satisfaction than a standard RB room temperature controller. When a bidirectional EV is connected additionally and a two-tariff electricity pricing is applied, it successfully leverages the battery and decreases the overall cost of electricity. Finally, we deployed it on a real building, where it achieved up to 30% energy savings while maintaining similar comfort levels compared to a conventional RB room temperature controller.

\* Corresponding author.

E-mail address: [bratislav.svetozarevic@empa.ch](mailto:bratislav.svetozarevic@empa.ch) (B. Svetožarević).

<https://doi.org/10.1016/j.apenergy.2021.118127>

Received 11 June 2021; Received in revised form 8 September 2021; Accepted 22 October 2021

Available online 10 November 2021

0306-2619/© 2021 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Buildings account for one-third of global primary energy consumption and one-quarter of greenhouse gas (GHG) emissions. Consequently, they have been identified as a critical element to enable climate change mitigation [1]. When looking at the energy use during a building's lifecycle, about 80% of it stems from building operation [2]. However, over the last two decades, buildings have become much more complex to operate optimally due to the integration of renewable energy generation, transformation, and storage devices [3]. Also, due to the electrification of the mobility sector, electric vehicle (EV) chargers are installed in buildings, thus further increasing the control complexity [4]. Today, multiple energy-flows are possible within a single building, which gives rise to the need for system-wide optimal energy management. At the same time, the users' needs for comfort, such as indoor thermal and visual comfort, and having enough energy in the EV battery for the next trip, shall be satisfied.

This multi-loop, multi-criteria control problem has raised specific needs within the building automation (BA) industry related to delivering optimal performance at low development and commissioning costs. In the following text, we first provide an overview of the BA industry requirements for an optimal controller for modern buildings. Then, we list the limitations of the current widespread rule-based (RB) controllers and the advanced, state-of-the-art model-based controllers, of which Model Predictive Control (MPC) is the most famous representative. We describe why both control methods fail to satisfy the current BA industry requirements. Following that, we motivate the potential of deep reinforcement learning (DRL) algorithms for the BA industry. We briefly review the current work on DRL applied to room temperature and EV charging control, and we close the Introduction with an overview of this work and a summary of main contributions.

### 1.1. Current BA industry requirements

*BA requirement I — Multi-loop control policy:* Compared to the situation before the 2000s, renewable energy generation, transformation, and storage devices have now been vastly integrated into new or retrofitted buildings allowing for more energy-efficient and cleaner operations, in terms of GHG and, in particular, CO<sub>2</sub> emissions [5,6]. A typical set of these devices could include photovoltaic (PV) panels, battery storage, a heat pump, and thermal storage. Hence, the number of possible energy flows and the number of decision variables have increased. For example, electricity could be obtained either from the grid, a stationary battery, or PV panels. Similarly, when and which electricity source to use to charge an EV depends on several factors, such as weather prediction and price of the electricity. Therefore, control of a modern building is a multi-input-multi-output (MIMO), i.e., multi-loop, energy management problem, which requires finding a control policy for several *controlled variables* simultaneously while considering several *external factors* that influence it.

*BA requirement II — Building-EV coupling:* The building-mobility sector coupling allows for more efficient control solutions than when these two sectors are addressed separately [7]. For example, when the electricity price is low, the building management system (BMS) could decide to heat the room, charge the EV, or store it in a stationary battery for later use. On the other side, this coupling also brings challenges. The charging of EVs causes additional energy consumption for a building, increasing its total – and possibly peak – energy consumption. Furthermore, most EV chargers start charging with full power as soon as an EV is connected. Therefore, if multiple EVs are charged at the same time in a neighbourhood, the aggregated demand can be very high, potentially causing energy dispatching and grid stability issues.

A particularly interesting symbiosis between a building and an EV arises when the latter is bidirectional, i.e., the EV battery can be charged and discharged. In that case, the stored energy could be used as a source of electricity for a building [8]. In this case, the EV battery

expands the capacity of the stationary battery, if one is installed. The difference with the stationary battery lies in its availability — the battery of a bidirectional EV is only available when the EV is connected to the building. Therefore, a BMS can use the battery of a bidirectional EV for energy management when the EV is connected. However, a BMS shall also ensure to charge the EV battery to a satisfactory level before the next trip.

*BA requirement III — Occupants comfort:* In developed countries, people spend on average 80%–90% of their time indoors. Therefore, the influence of building systems on occupants' well-being is deemed critical [9]. Consequently, occupants put more and more stringent requirements for comfort to facility managers, which is passed indirectly to the BA industry. Therefore, the value of a building controller is not only measured in terms of saved energy but also how comfortable the indoor environment is to the occupants.

*BA requirement IV — Transferability:* Buildings differ from each other in terms of *construction properties* (e.g. floor layout, geometry, materials used, age), *installed building services* (e.g. heating, ventilation, and air-conditioning (HVAC) systems), *outside conditions* (climatic region, orientation), and *occupancy profiles*. Therefore, an ideal building controller shall be able to provide optimal performance not only for the building it is designed for, but also for other similar buildings. If the engineering effort to apply such a controller to a similar building is small or negligible in terms of expert knowledge and time required, then the controller is considered transferable.

*BA requirement V - Adaptability and continuous commissioning:* The dynamics of a building can change significantly during its lifetime for several reasons, such as a retrofit, a change in the occupancy profile, or ageing. An ideal building control shall detect a change in the building operation performance, e.g. if a building starts to consume more energy than it used to, and readjust its parameters, i.e. adapt to the new situation. This capability of a controller is also called continuous commissioning [10,11].

Overall, the control of a modern building is a complex MIMO control problem with the objective to provide the desired thermal comfort to the occupants and simultaneously ensure the EV is charged to a satisfactory level for the next trip, all while minimizing the overall energy consumption to reduce the costs.

### 1.2. Limitations of RB controllers

Traditionally, more than 90% of industrial BA controllers are RB, such as bang–bang or proportional-integral-derivative (PID) controllers. They have fixed predefined rules, simple architectures with straightforward implementation, and several parameters with clear guidance on how to tune them. Even though RB controllers (RBCs) are widely adopted in the BA industry, there are several limitations on their use for achieving optimal control of modern buildings.

*RBCs limitation I - Difficulty to achieve system-wide optimal performance:* RBCs are suitable for single output control loops, whether single-input-single-output (SISO) or multi-input-single-output (MISO). Therefore, applying these single-output controllers to solve a multi-output (MIMO) control problem is a challenging and often infeasible task in practice as MIMO systems typically have dependencies between their sub-systems that cannot be neglected [10,12,13]. A MIMO system cannot be typically addressed as a collection of individual SISO/MISO systems [14]. For example, the temperature of the thermal storage determines the available heating capacity over the next couple of hours, while the output heating power of a heat pump connected to the thermal storage determines at what pace the temperature of the storage could be increased. A similar analogy could be drawn for the EV battery and its charging and discharging power. Indeed, optimal control of MIMO systems requires applying advanced MIMO control techniques [14].

*RBCs limitation II — Absence of optimality guarantees:* Even for single-output problems, manual tuning cannot provide optimality guarantees:

control experts could tune the RBC, in particular PID, to provide close-to-optimal regulation performance by looking at the overshoot, rise time, stability margins, and disturbance rejection, but there is no mathematical optimization involved in the tuning of the parameters. Therefore, most of the RB controlled loops in buildings perform sub-optimally [10,12,13].

*RBCs limitation III — Difficulty to include prediction rules:* RBCs do not typically involve any prediction rule. A prediction rule could be defined, for example, for pre-scheduled dynamic comfort bounds, which change between narrow, e.g. [22 °C, 24 °C], and wider, e.g. [20 °C, 26 °C] constraints. Such dynamic bounds are typical for office buildings, where wider comfort bounds are allowed outside of office hours to save energy. However, as an RBC would react to the change of the bounds only at the time of their change, such control will violate comfort. On the other hand, a predictive controller would pre-heat the room for some time before the narrow bounds need to be reached. Defining and tuning the prediction rule in an RBC would require experimenting with a building and determining the time dominant constant of a particular room so that the pre-heating interval could be defined precisely. However, this interval depends on the day of the year, the state of the room, i.e., accumulated heat in walls, and weather prediction. Hence, determining it precisely for all combinations of these parameters over the year is a challenging task [12].

*RBCs limitation IV — Absence of self-adaptation:* RBCs need to be re-tuned after a change in building dynamics to regain the previous performance, which requires expert knowledge and incurs costs [10]. (See BA requirement V)

Overall, RBCs fail to satisfy all the needs of the BA industry for an efficient way of obtaining an optimal controller for a modern building — they can only provide sub-optimal performance and require expert knowledge during commissioning and maintenance.

### 1.3. Limitations of MPC controllers

Advanced controllers, on the other hand, in their classical and non-adaptive form, can overcome the first three limitations of RBCs. The most well-known representative of this type of controllers is MPC, which can calculate optimal MIMO control signals for several steps ahead while respecting the state, input, and/or output constraints. However, the performance of an MPC controller strongly depends on the quality of the underlying building model used to develop this controller. A building model of a poor quality, which does not represent well the true building dynamics, e.g. a simple grey-box model with some generic parameters, will lead to unacceptable control performance [15]. On the other hand, obtaining a high-quality building model is a complex and time-consuming task requiring expert knowledge. Therefore, the costs of developing and implementing an MPC controller are justifiable only for well-defined systems, where the same controller could be used on many instances of the same system. However, as buildings differ substantially from each other, the costs of developing and deploying MPC controllers outweighs the cost benefits, and, therefore, classical versions of MPC controllers have not yet been widely adopted in the BA industry [15–17].

Over time, stochastic [18], robust [19], and adaptive [20] MPC controllers have been developed to address or circumvent the need for a high-quality building model. Even though some directions are promising, in particular those of adaptive MPC controllers with online system identification [20], they have only been applied to single-zone temperature control problems and validated in simulation. Validation on real buildings and solving of more complex building problems is needed for these methods to be accepted by the BA industry.

### 1.4. State-of-the-art data-driven RB and MPC controllers

In recent years, due to the increased availability of stored sensor and actuator data in buildings, researchers have started exploiting the information contained in this past data to develop improved building controllers. Two trends can be observed: first, using data to improve classical control strategies, such as RB and MPC, and second, applying pure data-driven methods from the machine learning (ML) domain and adapting them to building control.

The first direction, data-driven autotuning of RBCs, even though interesting from the industry perspective due to potential direct applicability, has not yet been widely addressed in the literature — only some recent preliminary results exist [21,22].

Considerably more work has been published in the domain of learning-based MPC (LB-MPC) recently [23,24]. The most widely spread approach is to model the building dynamics with a neural network (NN) and use it as a model in the MPC framework. However, as NNs are non-linear models, the main challenge is to use them in a linear or convex fashion so that efficient solvers could be applied. One option is to design a NN that can be used within MPC by constraining the output of the model to be convex with respect to the control inputs [25]. Besides NNs, Jain et al. [16] uses Gaussian processes to learn a model, which is then used within MPC. Another approach uses random forests for modelling [26]. Recently, it has also been validated experimentally, and preliminary results are promising [27]. Even though initial results on data-driven MPC are promising, what is missing is the discussion on the scalability and transferability of these approaches across different buildings (see BA requirement IV).

### 1.5. Potential of DRL for building control

In terms of pure data-driven ML methods, reinforcement learning (RL), and in particular DRL, have emerged in recent years as approaches that can fulfil all the requirements for modern building control. Even though RL was established in the 1960s [28], complex problems remained out of reach until recently, when people started using NNs within the RL framework. Together with the increased availability of large data sets and extensive computational power on demand, this led to the popularization of DRL methods and the demonstration of successful solutions for complex real-world problems [29,30]. Mnih et al. showed that DRL algorithms could achieve human-level or even super-human level intelligence in playing Atari games [31] and the game of Go [32]. Since then, other problems have been solved at human or super-human levels in image recognition [29], natural language processing [33], and medicine [34]. Motivated by these achievements in DL and DRL, building and control engineers started applying these methods to building control [35–37]. There are several reasons why DRL is a promising framework to fulfil all the BA industry requirements for control of modern buildings.

*DRL potential I:* DRL algorithms operating on a continuous state space, such as deep deterministic policy gradient (DDPG) [38], can learn a continuous control policy to maximize a given reward function through interactions with the environment. As DDPG requires many interactions with the building, this is not feasible in practice, and people have to rely on models to learn the optimal policy. There are no particular requirements on the underlying model, such as convexity condition, as needed in MPC. As a building model, one could use any kernel-type model. NNs are particularly popular as they can capture the non-linear dynamics of the building [39,40]. After fitting the model to the past data, it is used as the simulation environment in the RL framework.

*DRL potential II:* There are no restrictions on how the reward function could be defined. Not only a single criterion but also multi-criteria reward functions and trading off requirements could be used. Hence, complex MIMO control policies could be obtained at once (see BA requirements I, II, and III).

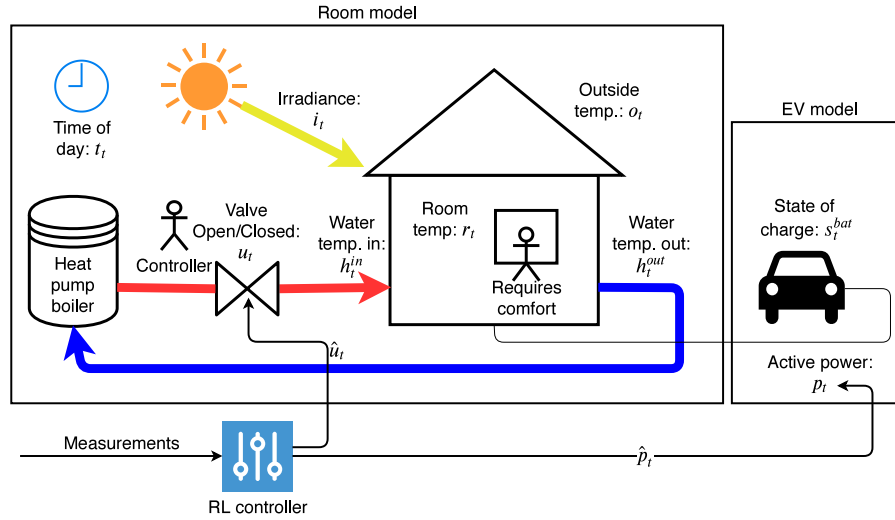


Fig. 1. Overview of the room model, bidirectional EV model, and joint deep reinforcement learning controller.

**DRL potential III:** Once the method is working for a certain room or building, it could also be applied to other rooms or buildings. The main part of the algorithm could be reused directly, thus demonstrating the transferability of the method (see *BA requirement IV*). This problem is known as transfer learning, and it has been already considerably addressed in general reinforcement learning [41]. However, only limited prior work was published recently on the transferability of DRL algorithms for building control [42].

**DRL potential IV:** Finally, if updated with the newly received measurement data, the DRL algorithm could be updated online to adapt to the new building dynamics, thus fulfilling the *BA requirement V* [43].

#### 1.6. State-of-the-art DRL-based room temperature and EV charging control

Most previous works on RL and DRL consider either controlling the building energy system, e.g. [27,44–49], or EV charging, e.g. [50–55]. There are a few works that control both the charging of an EV and a building energy system, e.g. [56–59]. In [56] for example, a building equipped with PV, an EV and an energy storage system is considered as a smart grid system, but no temperature control is addressed. The authors of [57] minimize the costs of electricity through improved operations of an HVAC system, an EV, a washing machine and a dryer. In [58,59], one-day ahead planning is used for building control, including an EV supporting bidirectional charging.

#### 1.7. Novelty and contribution of this work

In this work, we describe a fully black-box, data-driven, DRL-based method for the joint control of a room temperature and bidirectional EV charging (see Fig. 1). The main contributions of this work are the following.

First of all, the proposed data-driven pipeline requires only historical data to learn an optimal building control policy, and thus avoids the need for complex physics-based modelling required to develop advanced, model-based controllers (see Section 1.3) or expensive fine-tuning of the conventional, rule-based controller (see Section 1.2). As a black-box simulation environment that does not require any physics-based prior knowledge to train the policy, we use a Recurrent Neural Network (RNNs) model of the room thermal dynamics and a linear model of the EV battery. We applied Deep Deterministic Policy Gradient, which is a DRL algorithm in the continuous domain, to learn the control policy. Hence, this pipeline is a cost-effective way to obtain an optimal MIMO building control policy by only using available historical data of a building.

Secondly, we use the historical data from a real building, the DFAB HOUSE at Empa Duebendorf in Switzerland to obtain a close-to-reality simulation environment. We analyse the simulation results of the DRL policy in a heating season in terms of energy savings and occupant comfort and showed that it delivers better performance than a standard industrial RB controller. Furthermore, we considered an extended problem when bidirectional EV is connected to the building and the electricity price has two tariffs. We analysed the simulation results of the simultaneous control of room temperature and bidirectional EV (dis-)charging in terms of costs savings while minimizing the comfort violations for the desired comfort bounds and providing enough energy to the EV battery for the next trip. The obtained DRL-based control policy showed better performance compared with two standard industrial RB controllers — one for temperature regulation and another for EV (dis-)charging.

Thirdly, we validated experimentally the learnt DRL policy for room temperature control during the heating season. The DRL policy was directly transferred from simulation onto the real building, the DFAB HOUSE, and it was successfully regulating the temperature from the initial time of deployment, achieving up to 30% energy savings and better comfort satisfaction compared to a conventional, rule-based controller.

Fourthly, we discuss throughout the paper the potential of this approach to satisfy all the *BA Requirements (I–V)*.

#### 1.8. Structure

The paper is structured as follows: In Section 2, the case study used to showcase the proposed data-driven building control methodology and the data collection process are described. In Section 3, we present the methods used to model the room temperature and the SoC of the bidirectional EV. Further, we describe the definition of the RL environment and the reward functions for two different problems: (i) room temperature control and (ii) joint control of the room temperature and bidirectional EV charging. The simulation and experimental results are illustrated in Section 4. Finally, Section 5 provides an overview and concluding remarks of this work, as well as directions for future research.

## 2. Case study and data collection

The DFAB HOUSE, a three-storey residential building of the Empa demonstrator NEST in Duebendorf in Switzerland [61] (Fig. 2). NEST [62] (Fig. 2c) is a vertically integrated neighbourhood and a living lab.





Fig. 2. DFAB HOUSE. (a) The 2nd floor layout. The room temperature controller was developed for the room 471, on the bottom right. The locations of thermostats are marked with a square containing a “T”. (b) External view. (c) NEST demonstration building at Empa, Duebendorf, Switzerland.  
Source: Credits: R. Keller [60].

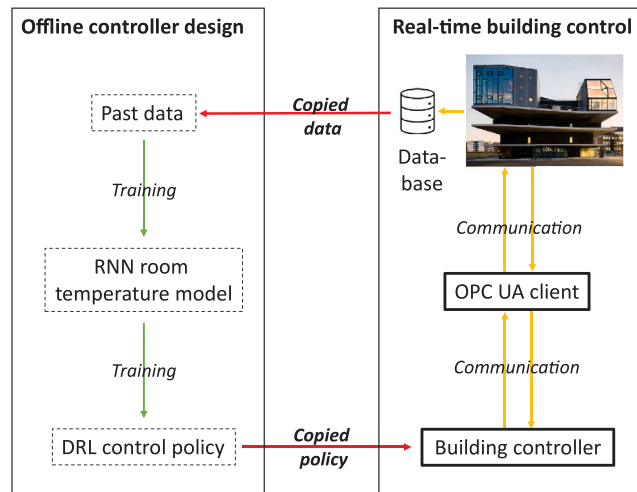


Fig. 3. Flow diagram of the methodology.

The DFAB HOUSE is operational since March 2019 and the corresponding sensor and actuator data is collected at 1 min resolution. We chose one bedroom (room 471, Fig. 2a) to apply our control algorithm. In this room, we can control the opening and closing of the valve that regulates the water flow into the floor heating system. As a bidirectional EV was not available at the time of this work, we emulated it based on the

past charging/discharging data of the stationary battery at NEST. For information on data preparation see Appendix D.

### 3. Methodology

An overview of the methodology is presented in the flow diagram of Fig. 3. Previously collected building and weather data, stored in a database, is firstly used to train the RNN room temperature model. This model serves as environment to learn the DRL policy offline, which is then copied to the building controller for real-time control. The communication with the building is done through an OPC UA client. In the following sections, we describe the control problem and present the data-driven pipeline used to get control policies from past data.

#### 3.1. The control problem and model overview

The overview of the system to be controlled is illustrated in Fig. 1. It consists of two parts: the room temperature model and the EV battery charging/discharging model. These two models are mainly independent, as they serve two different needs of the building occupants, namely to provide indoor comfort and enough battery capacity for the next trip, respectively. They are, however, linked through the overall building electricity demand. If the EV is being charged, the used energy indeed represents additional building energy demand. If the electric energy for heating/cooling is sourced from the bidirectional EV battery instead from the grid, then the overall building demand is reduced.

We can therefore formulate the control problem as: given the energy stored in the bidirectional EV battery, what would be the optimal room temperature control (heating or cooling) and optimal EV (dis-)charging strategy such that the overall costs for energy is minimized while satisfying the indoor comfort bounds and the minimum SoC of the EV at the moment of leaving. We assume that the EV leaves at 7:00 with a minimum of 60% SoC and returns at 17:00 with 30% SoC. The energy price is assumed to take a standard two-stage tariff profile, with a high price between 8:00 and 20:00 and a low price outside of this interval.

#### 3.2. Data-driven modelling of the room temperature and bidirectional EV charging

In this section, the model of the room temperature and the weather are described. Then, the two models are combined to obtain the final room temperature model, and we provide details on RNN architecture, model training, and hyperparameter tuning. Finally, we describe the bidirectional EV charging/discharging model.

##### 3.2.1. Room temperature model

We consider the temperature control of a single room (a single zone, i.e. room 471) at the DFAB HOUSE. The room temperature  $r_t$  depends on the outside temperature  $o_t$ , solar irradiance  $i_t$ , in-/out-flowing water temperature of the pipes  $h_t^{in}$ ,  $h_t^{out}$ , and the valve position  $u_t$  (see Fig. 1). Here, the index  $t$  denotes the time of the measurement. Since we will be using the room temperature model as a simulation environment for the RL agent, we need a model that predicts all uncontrollable (independent) variables. These are all of the above variables apart from the state of the valve  $u_t$ .

One way to solve the modelling task would be to fit the data with a multivariate time-series prediction model in an end-to-end fashion. This would allow predicting the evolution of all the variables based on their past values. Since the data collection at the DFAB HOUSE only started in March 2019, there was less than a year of operation and available historical data at the time of this work. To make the most out of this limited amount of data, we took a few more considerations into account that led us to partition the room model into different sub-models. They are discussed in the subsequent sections.




Fig. 4. Inlet and outlet water temperatures of the cooling panel over two weeks in summer.

**Remark 1.** The control framework described here could also be applied to different types of heating and cooling systems, where heating and cooling is provided by two different devices, e.g. an electric heater and an AC unit.

To get a smooth time variable, we transform the linear time of the day with a sine and a cosine function,  $t^s = \sin(t)$ ,  $t^c = \cos(t)$ , with  $t$  linearly going from 0 to  $2\pi$  during each day. Note that one could also define the time in a linear fashion, numbering the time intervals during each day. However, this induces jumps at midnight from the last to the first interval. Introducing the smooth sine and cosine time variables allows us to transfer the idea that these intervals are close to each other to the model.

### 3.2.2. Weather model

While there is a correlation between, e.g. the room temperature and outside temperature, the influence of the room temperature on the weather is non-existent. Therefore, to avoid that output of the weather model depends on the room state variables, we train an individual model of the weather. Such a model could be useful if no weather prediction data is available on site, but only past observed weather data could be taken as inputs. This model predicts the weather variables (outside temperature and irradiance) based on the past values of those variables and the time of day. Note that the weather model takes the  $n$  previous values of the input series into account to produce the output.

The temperatures of the water entering and leaving the cooling/heating system over a few weeks in summer are shown in Fig. 4. It can be seen that the water temperature coming from the heat pump is kept almost constant, but not always at the same level, which depends on the average outside temperature over a day. Since we are only interested in predictions with a rather short the horizon of one day at most, we decided to use a constant predictor for the water temperature variables. While this is a valid assumption for the inflow temperature, the outflow temperature is much more dynamic. However, we retained this assumption for the sake of the simplicity of this model.

### 3.2.3. Final room temperature model

Putting everything together, we can now build the *full* model of the room by combining the previously defined sub-models: the weather model, the water temperature model, and the room temperature prediction model. As mentioned previously, this model takes into account the  $n$  previous values of the input series  $s_{t-n+1:t}^{room}$  and the same number of control inputs  $u_{t-n+2:t+1}$  to compute the output. In this paper we use RNNs for this purpose, as explained in the next section.

### 3.2.4. RNN model

RNNs are commonly used in time series predictions to capture its time dependencies and tendencies [63]. Fig. 5 illustrates how a single step prediction is made and this scheme is naturally expanded to multi-step predictions. In that setting, part of the input is unknown and relies on the previous outputs of the model. It is then merged together with the known input part and fed to the RNN to build the next




Fig. 5. Simplified model structure for the case of predictions based on three previous inputs ( $n = 3$ ). On the left, the model uses the provided inputs to make a prediction  $\hat{O}_{k+1}$ . On the right, it extracts the true output  $O_{k+1}$  from the data, which can then be compared to the prediction to compute the loss and train the network.

prediction. Repeating this procedure allows one to get predictions for any number of steps for weather and room temperature models. Note that in practice, we train the actual recurrent model to only predict the difference in the prediction state, not the absolute state.

To optimize the loss, we use the ADAM [64] optimizer with a base learning rate  $\eta$  to minimize the mean-square-error (MSE) between the predictions and the ground truth. The training of the model lasted for  $n_{ep}$  episodes (see Table 2). We also monitor the losses on the training and on the validation set to get an idea about the amount of overfitting. The data used to fit the model is shuffled to avoid seasonal dependencies between the data in consecutive batches.

The hyperparameters that are used to tune the recurrent models are listed in Table 2. There are a few more parameters that we choose heuristically, for example, a number of neurons in each recurrent layer. To compare the performance of the models trained with different hyperparameters, we use the following objective. We predict 6 h (i.e. 24 timesteps of 15 min) into the future and take MSE between this prediction and the ground truth. For this process the validation data is used. The main idea is to find a model that generalizes well over multiple consecutive predictions and over unseen data. For the actual optimization, a Tree Parzen Estimator [65] is used, which is implemented in the Python library hyperopt [66].

### 3.2.5. Bidirectional EV charging/discharging model

We use a stationary battery available at NEST in order to emulate the battery of a bidirectional EV. This battery has a maximum capacity of 96 kWh at a SoC of 100%. However, we will restrict it to lie within the interval [20.0%, 80.0%] for safety reasons (the details on the safety are discussed later in Section 3.3.2). Furthermore, we limit the charge and discharge rate to  $\pm 100$  kW. Both stated maximum capacity and maximum (dis-)charging rate are also found in the following EV models: Tesla Models S and X [67], BMW i3 [68], and Mercedes-Benz EQC [69].

The change in SoC is modelled to be proportional to the active power applied, but the proportionality factor can be different for charging and discharging. We also allow for a constant discharging rate when the battery is not used, i.e. if the applied active power is zero, the battery slowly decreases its SoC due to losses. Let  $s_t^{bat} \in [20.0\%, 80.0\%]$  be the SoC at time  $t$ , let  $p_t \in [-100 \text{ kW}, 100 \text{ kW}]$  be the average active power from time  $t-1$  to time  $t$ . Finally let  $\Delta s_t^{bat} := s_t^{bat} - s_{t-1}^{bat}$  be the change in SoC at time  $t$  compared to time  $t-1$ . Therefore, the change in SoC, or charging/discharging of the EV battery, can be defined as:

$$\Delta s_t^{bat}(p_t) \approx \Delta s_t^{bat}(p_t) := \alpha_0 + \alpha_1 p_t + \alpha_2 \max\{0, p_t\} \quad (1)$$

where  $\alpha_i, i = 0, 1, 2$  are the variable coefficients that can be fitted to the data using least squares. Finally, we can define the battery model as:

$$\hat{s}_{t+1}^{bat} := s_t^{bat} + \Delta s_{t+1}^{bat}(p_{t+1}) \quad (2)$$

It models how the SoC evolves when an active power of  $p_{t+1}$  is applied. We consider the model to be charging if the active power is positive and discharging otherwise.

### 3.3. RL environment

In RL, an agent is learning a control policy through interaction with an environment. Letting  $s_t$  and  $a_t$  be the state and the action at time  $t$ , respectively, the environment maps them into the new state and the reward  $R_t: (s_t, a_t) \mapsto (s_{t+1}, R_t)$ . In this work, we trained our agents in an episodic framework, with a fixed episode length of  $l_{ep} := 48$ . With one timestep corresponding to 15 min, this corresponds to an episode length of 12 h.

In the following sections, we define the environment of our particular problem using the previously described room temperature and EV (dis-)charging models.

#### 3.3.1. Room temperature environment

To describe the evolution of the room temperature according to the actions taken by the agent, we use the room temperature model defined in Section 3.2.1. To initialize the environment in each episode, we sample an initial condition from the historical data in the database and we then use the model to simulate the behaviour of the room under the agent's policy for the length of the episode. This episodic framework allows us to control the errors of the model, since we know how well it performs. Further, to incorporate stochasticity, a disturbance term  $d(t)$  is added to the output of the deterministic model. The latter is modelled by an auto-regressive (AR) process that was fitted based on the residuals of the NN model. This ensures that the disturbance is realistic, i.e. as seen in the past data.

The reward of the agent controlling the room temperature is defined as follows:

$$r_t^{room}(s_t^{room}, a_t^{room}) := \underbrace{-a_t^{room} \cdot |h_t^{in} - h_t^{out}|}_{\text{Energy usage}} - \underbrace{\alpha \cdot c_t^{pen}(r_t)}_{\text{Comfort violation}} \quad (3)$$

$$= -e_t^{room} - \alpha \cdot c_t$$

where we defined  $e_t^{room} := a_t^{room} \cdot |h_t^{in} - h_t^{out}|$  and  $c_t^{pen}$  denotes the penalty function for room temperatures that are outside the comfort bounds. The parameter  $\alpha > 0$  determines the weight of the temperature bound violation compared to the energy usage. The penalty function  $c_t^{pen}$  is defined as follows:

$$c_t^{pen}(r_t) := \begin{cases} 0 & r_{min} \leq r_t \leq r_{max} \\ r_{max} - r_t & r_t > r_{max} \\ r_t - r_{min} & r_t < r_{min} \end{cases} \quad (4)$$

Note that this function is always positive and increases linearly with  $r_t \rightarrow \pm\infty$  as soon as the temperature gets out of the defined comfort bound  $[r_{min}, r_{max}]$ .

#### 3.3.2. EV battery environment

To build the RL environment for the EV battery, the battery model described in Section 3.2.5 is used. The SoC of the battery at a given time  $t$ ,  $s_t^{bat}$ , is used as state and the active power  $a_t^{bat} := p_{t+1}$  as action of the environment. Note that the subscripts do not match since we defined  $p_{t+1}$  as the active power applied from  $t$  to  $t+1$ , but this is also at time  $t$ .

Besides restricting the active power, we also want to restrict the SoC of the battery to lie within a certain range. Since the battery model learnt from the data is piece-wise linear and strictly increasing, it can be inverted and used to build a fallback controller. We implemented two functions in the fallback controller. First, the fallback controller prevents the SoC from falling out of the previously defined safety range,  $[20.0\%, 80.0\%]$ . The actions are not directly used but clipped using the safety guaranteeing function  $f^{safe}$  that will clip the chosen actions to the required range for the constraints to be fulfilled:

$$e_t^{bat} := -f^{safe}(a_t^{bat}) \quad (5)$$

More details on how this function is defined can be found in Appendix C.3. Furthermore, the fallback controller achieves a specified




Fig. 6. EV battery environment demonstration.

SoC at the desired future time  $t_{des}$  by restricting the battery to be charged at high power when the SoC is too low when approaching  $t_{des}$ . This makes it easy to build an environment for RL: we can choose the reward as the negative active power applied per timestep and we do not need additional penalties contained within the reward that penalize SoCs outside of the given bounds or not reaching the SoC goal at time  $t_{des}$ . In this way, we omit choosing a heuristic factor for balancing the energy used and the SoC constraint violation (see Appendix C for details on SoC constrain violation).

Fig. 6 shows how the resulting environment behaves under two different heuristic agents that apply a constant action. One is discharging and the other is charging at a constant rate. Note that in this case, we chose  $t_{des}$  as the end of the episode, i.e.  $t_{des} = l_{ep} := 48$ . One can see that the agent that constantly wants to discharge arrives at the minimum SoC after a few steps and needs to charge the battery at full capacity when approaching the end of the episode. The safety controller built into the environment prevents the SoC from falling below the minimum and charges the battery before the end of the episode, even if the agents continue to discharge.

#### 3.3.3. Joint room temperature and EV battery environment

As both subsystems evolve independently, we simply use the room and battery environments described above to map their current state and action to the next state and reward and then concatenate them. Since the reward was one-dimensional in both cases, we combined the two in a weighted sum as follows:

$$r_t^{joint}(s_t^{joint}, a_t^{joint}) := \underbrace{-p(t) \cdot (\alpha_{bat} \cdot e_t^{bat} + e_t^{room})}_{\text{energy usage}} - \underbrace{\alpha \cdot c_t^{pen}(r_t)}_{\text{comfort violation}} \quad (6)$$

where  $p(t)$  denotes a suitable energy price function that may vary over the course of a day, but is the same for different days. Note that compared to the room temperature environment, in this case, we are no longer interested in energy minimization but in price minimization. Note, also, that coefficient  $\alpha$  here is introduced to balance out the consumption of the battery and the room, which have different scales.

### 3.4. DRL algorithm

In this work, we used the Deep Deterministic Policy Gradient (DDPG) algorithm [38]. It is model-free, off-policy, and uses an actor-critic setting. Unlike standard Q-learning, it naturally handles continuous state and action spaces, which was one of the main reasons this




Fig. 7. Piece-wise linear EV battery charging/discharging model.

algorithm was chosen. This choice was also motivated by previous work using this algorithm, for example in [70–73]. There exists an implementation of DDPG based on the python deep learning library Keras [74] in another library called Keras-RL [75].

Four neural networks are used within the DDPG algorithm: an actor (taking actions) and a critic network (evaluating these actions) and corresponding target networks for each of them. Note that the actor and its target network have the same architecture but different weights. The same applies to the critic and its target network. In our case, a fully connected neural network with two layers of 100 units and the Rectified Linear Unit (ReLU) activation function was used for both the actor and the critic. To perturb the actions chosen by the actor network with exploration noise, an Ornstein–Uhlenbeck process (see e.g. [76]) was used. As for the RNN training in the modelling section, we used the ADAM optimizer [64] to update the parameters of the neural networks. The discount factor  $\gamma$  was fixed to 0.99. Note that a few more hyperparameters, like the learning rate for the ADAM optimizer and the number of training episodes, were adjusted manually. This could be avoided using automatic hyperparameter tuning, as it was done in the case of the neural network models in .

## 4. Results

In this section, the results of different elements of the proposed data-driven DRL-based control learning pipeline are presented. First, the evaluations of the room temperature and bidirectional EV (dis-)charging models are shown and analysed. Then, the simulation results of applying the DRL algorithm to the room temperature control are illustrated, followed by the results on the joint control of the room temperature and EV (dis-)charging operations. Finally, the experimental results demonstrating the DRL agent applied to the real building are presented.

### 4.1. Simulation results

#### 4.1.1. Evaluation of the EV battery model

The piece-wise linear EV battery charging/discharging model, together with the real data collected at NEST used for fitting, can be seen in Fig. 7.

The 6 h ahead SoC prediction using the EV battery model described in Section 3.2.5 is shown in Fig. 8a. Note that the ground truth is shown for comparison and was not used to fit the model. We also performed a more detailed analysis of the prediction performance of the battery model by analysing the mean absolute error (MAE) and maximum absolute error for a different number of prediction steps, up to 12h prediction interval (Fig. 8b). The prediction captures the dynamics very well, with an MAE of the SoC of less than 0.75% after 6 h. On average, after 12h, the prediction will be less than 1% away from the true SoC.

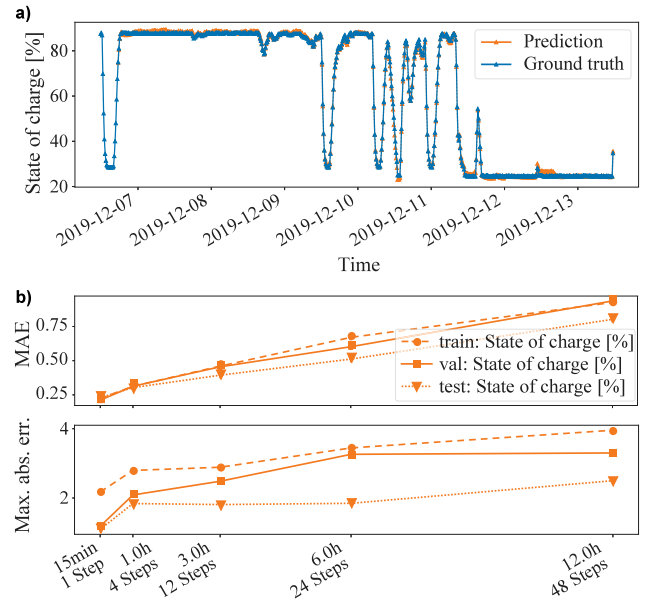


Fig. 8. Prediction performance of the EV battery model. (a) A sample week from the training set. (b) Quantitative evaluation.

#### 4.1.2. Evaluation of the weather model

We compare two methods for the weather model: a linear model and a recurrent neural network model. As a linear model, we chose a 5-fold cross-validated multi-task Lasso estimator from SKLearn [77]. For the RNN, we used the same configuration as the other RNNs in this study (see Section 3.2.4). Both models used the same inputs to make the predictions, i.e. data from the previous 19 steps. Further, we used clipping at 0 for the irradiance in both cases for a fair comparison. Note that this makes the model previously described as linear actually only piece-wise linear.

Fig. 9 shows how the weather model performs when evaluated on the test set for one specific initial condition. It can be observed that the piece-wise linear model makes smoother predictions and diverges faster than the RNN model. The quality of predictions drops with the longer horizon and, overall, the RNNs provide better predictions, even though the linear model is comparable on short horizons.

Note that, by investing more thoughts into the piece-wise linear model, e.g. by manual feature engineering, one might obtain a linear model that may be able to outperform the RNN. On the other hand, as the dataset grows with time, it is easy to increase the size of the RNN to make it more powerful, which is not the case for the linear model, which is another reason the RNN was favoured.

#### 4.1.3. Evaluation of room temperature model

The performance of the room temperature model is shown in Fig. 10. A quantitative evaluation of the model is shown in Fig. 10b, where the temperature prediction is done over a whole week. The MAE and maximum absolute errors are 0.5 °C and 2.3 °C after 12 h, respectively. As this RNN model showed a satisfactory fit, we selected it as an environment to train the DRL agent.

Note that the quality of the room model influences the final control performance. One known issue is that black-box models, i.e. non-physics based models, do not extrapolate well for unseen data. In our case, the room temperature model could be outputting physically inconsistent behaviours in the worst case. For example, on a winter day with low solar irradiance and the heating turned off, a black-box model might predict an increase of the room temperature. Such inconsistent physical outputs of the room temperature model can influence the control policy search negatively, as the DRL agent could learn that it could heat the room by closing the heating valves. Therefore, the






Fig. 9. Weather models: (a) Prediction over several days — example. (b) Quantitative comparison.

more physically-consistent behaviour a room model expresses for the test data, the better control performance of the DRL agent is expected. However, a detailed analysis of the physical inconsistency of the room temperature model for some input data is outside of the scope of this work.

#### 4.1.4. Evaluation of the DRL agent for room temperature control

We evaluated the DRL agent for both heating and cooling seasons, either by taking two different agents, one for each season, or by letting a unique agent learn the global control policy. We tested both approaches and obtained better results for the separate agents. The reasons for better results in the case of heating only or cooling only agent is that it makes the problem less complex. In that way, the deep learning (DDPG) agent is able to find a better control policy.

It actually turned out that for heating cases only, the optimization of the DDPG agent was much harder than in the case of searching for a global control policy and required some manual hyperparameter tuning to perform well. Therefore, we decided to switch to a reference tracking mode by setting  $r_{max} = r_{min} = 22.5$  °C. This makes it easier for the agent to know what actions are beneficial for temperature control since the temperature bound violation is only exactly zero for  $r_t = r_{max} = 22.5$  °C. As soon as  $r_t$  differs, the comfort violation will increase and the agent is penalized. We trained the RL agent for 20'000 steps and the evaluation is shown in Fig. 13, where the agent is compared to the following three controllers: one always opening the valves, one always closing them, and a rule-based bang–bang controller without hysteresis, which is a standard industrial controller. One can observe that the DDPG agent achieves on average 17% energy savings and 19% better comfort satisfaction compared to the rule-based controller.

A simulated case example is shown in Fig. 11. The DDPG agent can accurately control the room temperature by starting to open the valve before the RB controller, i.e. before the temperature reaches the setpoint, and opening them only a little to avoid overshooting. One can observe that the DDPG agent obtained the least comfort violations while using less energy than the rule-based agent. The quantitative analysis of this example shows 36% energy saving and 13% better comfort (see Fig. 12).

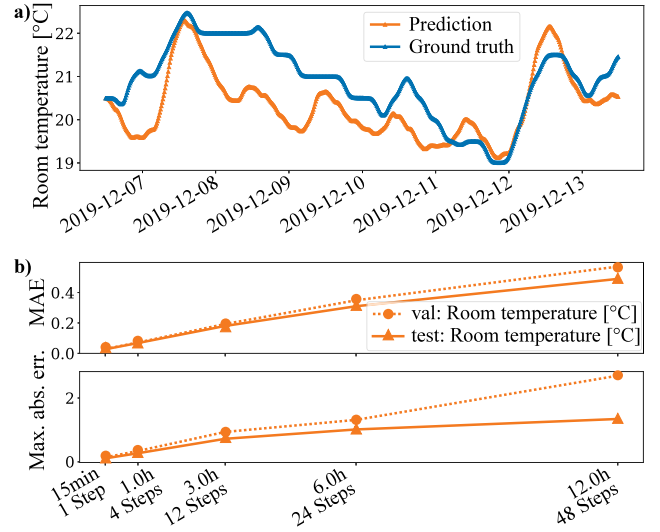


Fig. 10. Room temperature prediction models: Comparison of the piece-wise linear model and the RNN as trained for the weather model. (a) Qualitative comparison — example. (b) Quantitative evaluation.

#### 4.1.5. Evaluation of the joint room heating and EV charging control

As in the previous case of room temperature control, we again use three controllers as a comparison for the evaluation:

- **Valves Open, Charge:** This agent always leaves the valves open, as the *Valves Open* agent in the previous setting, but additionally always charges the battery at full power instantaneously upon arrival of the EV until it is full.
- **Valves Closed, Discharge:** This agent does the opposite of the previous one, i.e. it never opens the valves and constantly tries to discharge the battery at full power.
- **Rule-Based:** This agent does the same as the previous Rule-Based agent for the heating and constantly charges the battery at full power.

The performance of a MIMO DDPG agent trained on the joint environment is shown in Fig. 16. For the room temperature control, we used the same parameters as in Section 4.1.4 and we considered only heating cases. While again being able to reduce the comfort violations and the heating energy usage compared to the RB agent, the DDPG agent also achieved lower costs. As expected, the agent that never turns the heating on and discharges the battery uses the least energy, which also resulted in the lowest costs. Additionally, comfort violations are less pronounced than in the case of constant heating (constant valve kept on), but still worse than in both RB and DDPG controlled cases.

A simulated example is shown in Fig. 14. The DDPG agent manages to regulate the comfort better by using the energy stored in the EV battery. Compared to the RB controller for heating, which heats at the maximum power while the temperature is lower than the reference temperature of 22.5 °C, the DDPG controller actively regulates the valves so that better tracking is achieved. In terms of EV battery management, the energy from the EV battery is immediately used at the beginning of the interval until the minimum level of 20% of SoC is reached, which makes sense due to the lower electricity tariff at this time. Then, before the start of the next trip, the fallback controller charges the EV battery to the required SoC. The DDPG control output is presented in full red line, while the constrained DDPG is shown in the dashed light red line. The quantitative analysis of this DDPG agent is shown in Fig. 15, where it achieves 63% energy savings, 71% better comfort, and 53% costs savings compared to two RB controllers, for a certain weighting factor between the energy cost savings and comfort satisfaction. Note that this result is specific to the weighting factor used

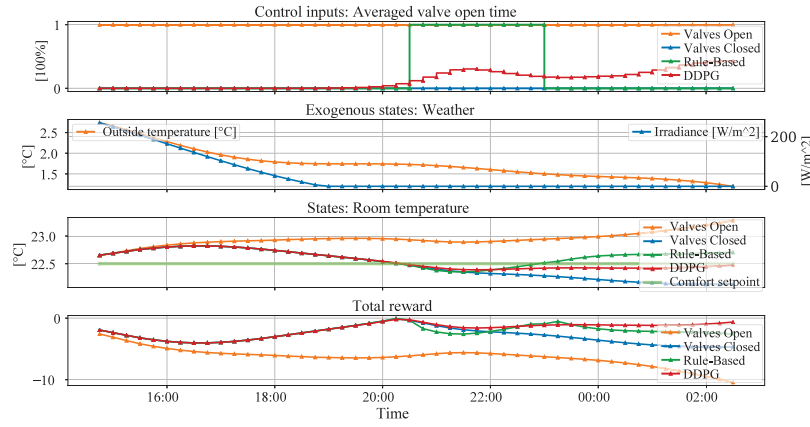


Fig. 11. Heating case demonstration. The inflow water temperature was 30.4 °C.




Fig. 12. Heating case rewards — example. Quantitative analysis of the results in Fig. 11.




Fig. 13. Heating agent evaluation — average over the whole heating season.

in the reward function. On average, when tested over 10'000 historical intervals, the MIMO DDPG controller achieved 12% better comfort satisfaction, 11% energy savings, 63% less EV charging at home, and 42% energy costs savings compared to two standard RB controllers, for the same weighting factor.

#### 4.2. Experimental results

The DRL control agent, which was obtained in Section 4.1.4 for the heating season and tested in simulation, was applied on the real building, the DFAB HOUSE, in room 471, for two weeks in February

2020. The performance of the DRL controller was compared with the performance of the room temperature bang–bang RB controller implemented in the same room over a subsequent week. The time-series results are shown in Fig. 17. Both controllers are aiming at the setpoint 22.5 °C. Due to the chosen weighting factor emphasizing energy savings, the DDPG controller is using less energy, at the cost of comfort, keeping the temperature slightly under the setpoint (−0.3 °C on average). On the other hand, the RBC is staying closer to the setpoint (−0.1 °C on average), but it is using more energy.

As the ambient conditions were naturally different for both controllers, we compared them using the Heating Degree Days (HDD) as a normalization variable. As per definition, the HDD of a given day represents how far from 18 °C the daily average temperature is [78]. In other words, higher heating degree days mean lower average outside temperature, for which we naturally expect more energy to be needed. The outside temperature was indeed approximately 4 °C lower during the DDPG experiment, which forced the controller to use more energy and made it hard to compare both experiments without a normalization procedure.

The daily energy used by both the DDPG and the RBC during five experimental days each are plotted against the corresponding HDD in Fig. 18. We can see that the DDPG controller outperforms the RBC: at HDD levels of around 7 and 12.5, the energy savings are 28% and 26%, respectively. On the other hand, we can also observe that while both controllers used between 6 and 8 kWh during three days, the average outside temperature was much lower (about 4 °C colder) during the DDPG experiment. In other words, the DDPG algorithm was able to use the same energy budget and maintain similar comfort levels to the RB approach but in harsher conditions.

Additionally, the points in Fig. 18 exhibit linear-like behaviours. To leverage that fact, we fitted a linear regression to both controllers to capture their global behaviour. This allowed us to clearly picture the gap between the RB algorithm and our proposed method, which on average saves around 25%–30% energy.

## 5. Conclusion and discussion

In this paper, we introduced a fully data-driven DRL-based method to obtain optimal control policies for MIMO building control problems. We demonstrated the method on the joint control of room temperature and bidirectional EV (dis-)charging to minimize the energy consumption and maximize occupants thermal comfort while ensuring enough energy stored in the EV upon leaving for the next trip. We demonstrated the method on a real building case study — the DFAB HOUSE at Empa Dübendorf in Switzerland with available past operational data of less than a year.

We show in simulation that the trained DRL agents are capable of saving on average 17% energy over the whole heating season while

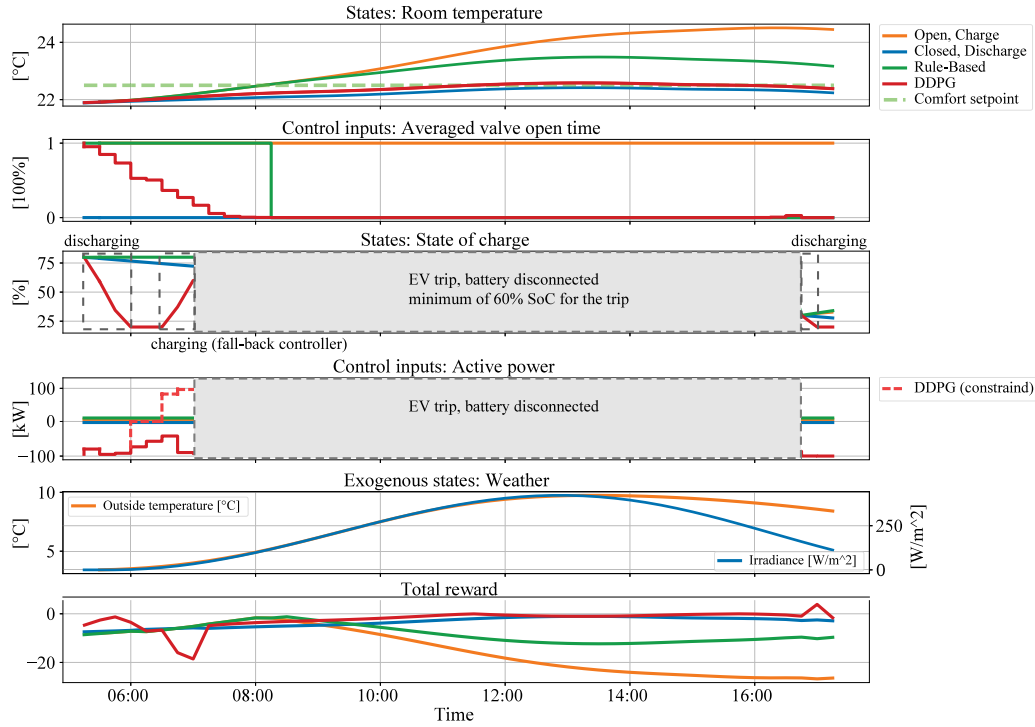


Fig. 14. Joint room heating control and EV charging. Evaluation of control agents in simulation — example. Inflow water temperature: 29.4 °C.

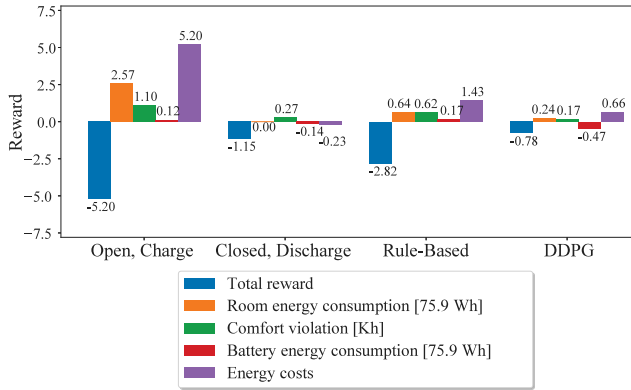


Fig. 15. Joint EV charging and room heating control agent evaluation — Statistics of the example plotted in Fig. 14.




Fig. 16. Joint EV charging and room heating control agent evaluation over a total of 10'000 steps.

providing 19% better comfort satisfaction compared to a classical rule-based controller. When an EV is additionally connected to the building and two tariff electricity pricing is considered, the DRL agents can successfully leverage its battery and decrease the overall cost of electricity. The obtained DRL control agent achieved 12% better comfort satisfaction, 11% energy savings, and 42% energy costs savings compared to two standard RB controllers, one controlling the room temperature and another controlling the bidirectional EV (dis-)charging. This result is specific to the weighting factor used in the DRL algorithm to balance the energy cost savings and comfort satisfaction. Finally, we demonstrate a successful transfer of the learnt DRL policy from simulation onto the actual building achieving up to 30% energy savings while maintaining similar comfort compared to a conventional RB room temperature controller over one week during the heating season. Note that one should keep in mind that, while the methodology is generic and might be applied to any building and under any external conditions, those particular energy savings were obtained on the DFAB house over a week long experiment, and further analysis would be required to understand the effect of such control policies in general.

The data-driven DRL-based control method proposed in this work is a viable approach to satisfy all the BA industry requirements for control of modern buildings, as defined in Section 1.1. We demonstrated that this method could match the first three BA industry requirements. In terms of the fourth requirement on transferability (and usability) for similar control problems in other buildings, we can argue in favour of the developed method; One can reuse the same NN and DRL architectures to obtain control policies in other conditions.

Indeed, we applied the same methodology to another room at the DFAB HOUSE, and we obtained comparable results. Thus, we believe that this method has a strong potential to work for any building or room, and could thus be a stepping stone towards obtaining transferable model-free data-driven room temperature control policies. As such, we also believe it to be valuable for the BA industry as it minimizes the engineering efforts to obtain a custom-tailored controller for a different room and building.

However, a few points still need to be addressed before this method can achieve widespread transferability to any building or room.

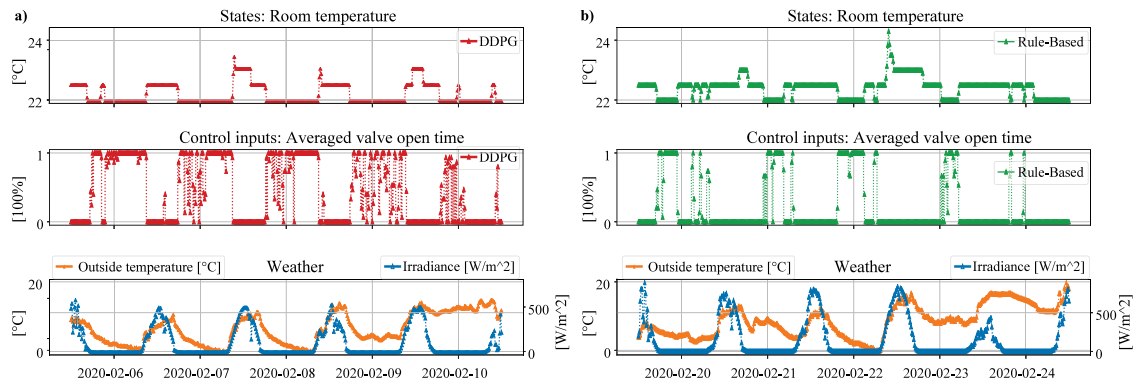


Fig. 17. Comparison of DRL (DDPG) control agent (a) and classical, RB controller (b) at the DFAB HOUSE room 471 at Empa, Duebendorf in Switzerland.




Fig. 18. Experimental comparison of the DDPG and rule-based controller performances in term of required energy over the heating degree days. The DDPG provides in average 27% energy saving and better comfort at the same time. The numbers in the plot correspond to days in February 2020.

The availability and quality of the building model is the first point to be addressed in future work. As demonstrated in this paper, the building model could be built as an RNN model, which could be directly applied to another room with the same setting, i.e. the same HVAC equipment and the same number of sensors and actuators. However, rooms generally differ in terms of HVAC equipment and the number of sensors and actuators. Thus, to model a different room, a certain engineering effort needs to be invested into linking the new inputs and outputs to the RNN model and fitting it. This process could be simplified and even automated if a linked, i.e. semantic, database of a building exists.

Secondly, the availability of past building operational data is a requirement to apply our black-box pipeline. While this may not be an issue for existing building with operational data stored in databases, it could be an issue for new or retrofitted buildings. A potential solution to this could be to apply transfer learning to the modelling part and learn the dynamics of a new building with the available fewer data. Similarly, transfer learning could be applied to “jump-start” the learning of the control policy for another building, given already existing proven policies in other buildings. This is also directly related to the last *BA industry requirement* on (self-) adaptability and continuous commissioning of building controllers. Indeed, transferring a controller to another or a retrofitted building is in essence very similar and a very interesting direction for future work.

Thirdly, the questions of stability and robustness of the proposed methodology could be addressed in future work, which is an open problem and very hard to solve (see, e.g. [79,80]). Since data-driven DRL-based control policies are based on NNs, there may exist adversarial examples for which the room temperature violates the predefined comfort bounds. In our case however, we are not dealing with safety-critical systems, such as in autonomous driving; the water-based radiant heating/cooling systems are stable by design — if constantly kept on,

the room temperature would increase/decrease and saturate at the water temperature circulating in the system. The maximum/minimum room temperatures that could be achieved would lay outside standard comfort bounds, but they are typically not dangerous for human health or catastrophic for the heating/cooling system. Therefore, in our case, we rather talk about a soft constraint satisfaction problem, and the controller does not always have to be robust. Moreover, comfort bound violations also occur with other types of controllers, such as RB or MPC, due to unlearned dynamics or strong disturbances, such as from solar gains through windows. In other words, no room temperature controller would guarantee comfort bounds satisfaction at all times, in particular, if the bounds are very narrow, for example, of  $\pm 1$  °C. One interesting method that could be explored to create more robust policies is policy distillation [81], where different control policies are learnt for different input data (e.g. one controller per season/month in the case of buildings) and then a common policy that shall work in any conditions is extracted from them.

## CRedit authorship contribution statement

**B. Svetozarevic:** Conceptualization, Methodology, Visualization, Writing – original draft, Supervision. **C. Baumann:** Methodology, Software, Data curation, Visualization, Writing – review & editing. **S. Muntwiler:** Methodology, Writing – review & editing, Supervision. **L. Di Natale:** Methodology, Visualization, Writing - review & editing. **M.N. Zeilinger:** Review & editing, Supervision. **P. Heer:** Review & editing, Project administration, Funding acquisition.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgement

This research was funded by Empa Duebendorf in Switzerland, Swiss National Science Foundation under NCCR Automation (grant agreement 51NF40\_180545), and funds from the Bosch Research Foundation im Stifterverband.

## Appendix A. Variables

The main variables used in the paper are summarized in Table 1.

## Appendix B. Hyperparameters of the RNN

The hyperparameters of the RNN model of the room temperature dynamics are shown in Table 2.



**Table 1**

Overview of the variables used in the model.

| Variable [Unit]                                    | Symbol  | Interval        |
|--|---|-----------------|
| Outside temperature [°C]                           | $o_t$   | [−15.0, 40.0]   |
| Room temperature [°C]                              | $r_t$   | [10.0, 40.0]    |
| In- and Out-flowing heating water temperature [°C] | $h_t^{in}, h_t^{out}$                               | [10.0, 100.0]   |
| Irradiance [W m <sup>−2</sup> ]                    | $i_t$   | [0.0, 1300.0]   |
| Valve open/close state                             | $u_t$   | [0.0, 1.0]      |
| Battery SoC [%]                                    | $s_t^{bat}$   | [0.0, 100.0]    |
| Battery power [kW] [kW]                            | $p_t$   | [−100.0, 100.0] |
| Safe (active) power applied to the battery [kW]    | $e_t^{bat}$   | [−100.0, 100.0] |
| Sine and Cosine of time of day                     | $t_t^s = \sin(t), t_t^c = \cos(t)$                  | [−1.0, 1.0]     |
| Combined time of day                               | $t_t = (t_t^s, t_t^c)$                              |                 |
| Combined weather variables                         | $w_t = (o_t, i_t)$                                  |                 |
| Room state variable space                          | $s_t^{room} = (w_t, r_t, h_t^{in}, h_t^{out}, t_t)$ |                 |
| Full room variable space                           | $s_t^{full} = (s_t^{room}, u_{t+1})$                |                 |
| Joint state space of room and battery              | $s_t^{joint} = (s_t^{room}, s_t^{bat})$             |                 |

**Table 2**

Hyperparameters of the RNN.

| Symbol      | Meaning   | Domain         | Room temp. model | Weather model |
|-------------|---|----------------|------------------|---------------|
| $n_l$       | Number of recurrent layers                                  | $\mathbb{N}$   | 3                | 1             |
| $n_e$       | Number of neurons per layer                                 | $\mathbb{N}$   | 30               | 60            |
| $n_{ep}$    | Number of epochs in training                                | $\mathbb{N}$   | 10               | 80            |
| $\eta$      | Learning rate   | $\mathbb{R}^+$ | 1.544e−5         | 6.163e−5      |
| $\sigma_i$  | Standard deviation of the Gaussian noise added to the input | $\mathbb{R}^+$ | 3.633e−6         | 0.01202       |
| <i>Cell</i> | Type of recurrent cell used                                 |                | <i>LSTM</i>      | <i>GRU</i>    |

### Appendix C. SoC constraint satisfaction

Assume we have the piece-wise linear battery model as described in Section 3.2.5, omitting the *bat* superscript and using  $p = p_{t+1}$  for clarity:

$$\begin{aligned}\hat{s}_{t+1} &= s_t + \Delta s(p) \\ &= s_t + \alpha_0 + \alpha_1 p + \alpha_2 \max\{0, p\}\end{aligned}\quad (7)$$

We assume that the coefficients  $\alpha_i$  have the following properties to make sure the model can be inverted and is physically meaningful:

- $\alpha_1 > 0$ : The more the battery is discharged (i.e. the more negative the active power), the faster the SoC decreases.
- $-\alpha_1 < \alpha_2 < 0$ : The slope of  $\Delta s(p)$  is always positive, but smaller for charging ( $p > 0$ ), so that one cannot gain energy from charging and discharging again.
- $\alpha_0 \leq 0$ : The battery slowly loses energy when it is not used.

Using these properties, we explain below how to make sure that bounds on the SoC will be satisfied.

#### C.1. Minimum and maximum SoC constraints

We require the SoC of the battery to lie within predefined bounds  $[s_{min}, s_{max}]$  at any time. Assuming we start from  $s_t$ , it suffices to show that the next SoC,  $s_{t+1}$ , given the previous SoC, stays within the bounds, and then apply the argument recursively. For the maximum constraint, we have to make sure that:

$$\begin{aligned}s_{t+1} &= s_t + \alpha_0 + \alpha_1 p + \alpha_2 \max\{0, p\} \leq s_{max} \\ \Leftrightarrow \alpha_1 p + \alpha_2 \max\{0, p\} &\leq s_{max} - s_t - \alpha_0\end{aligned}\quad (8)$$

Let us define the following helper function:

$$h(p) := \begin{cases} \alpha_1 + \alpha_2 & p > 0 \\ \alpha_1 & \text{else} \end{cases}\quad (9)$$

Note that it is positive for all values of  $p$  because of the properties on the coefficients  $\alpha_i$ . Now we can rewrite the equation above as:

$$\begin{aligned}p \cdot h(p) &\leq s_{max} - s_t - \alpha_0 \\ \Leftrightarrow p &\leq \frac{s_{max} - s_t - \alpha_0}{h(p)}\end{aligned}\quad (10)$$

To get a bound for  $p$  from this equation that does not contain  $p$  itself, we need to make a case distinction:

- **Case 1:**  $s_{max} - s_t - \alpha_0 > 0$  This means that the SoC at the next step will be lower than the maximum SoC when  $p = 0$ , we can therefore discharge as much as we want, i.e. we do not need to handle the case  $p < 0$ , so we only look at  $p > 0$  and we have  $h(p) = \alpha_1 + \alpha_2$
- **Case 2:**  $s_{max} - s_t - \alpha_0 < 0$  This means that the SoC at the next step will be higher than maximum SoC when  $p = 0$ , we therefore need to discharge in any case, i.e.  $p < 0$ , which means  $h(p) = \alpha_1$

Putting the two cases together, we get the following bound on the active power  $p$ :

$$p \leq c^{bat}(s_{max}, s_t; \alpha_i) := \begin{cases} \frac{s_{max} - s_t - \alpha_0}{\alpha_1 + \alpha_2} & s_{max} - s_t - \alpha_0 > 0 \\ \frac{s_{max} - s_t - \alpha_0}{\alpha_1} & \text{else} \end{cases}\quad (11)$$

Note that in the edge case  $s_{max} - s_t - \alpha_0 = 0$  both cases return the same, i.e. the bound is continuous. Applying the same chain of reasoning to the  $s_{min}$  case, one can derive the following:

$$p \geq c^{bat}(s_{min}, s_t; \alpha_i)\quad (12)$$

Note that this case is using the exact same function.

#### C.2. Achieving the goal SoC

We want to ensure that the battery is charged for some minimum desired amount  $s_{des}$  at a given time  $t_{des}$ . Assuming we are now at time  $t$ , i.e. the SoC is  $s_t$ , and assuming we can charge for a maximum of  $p_{max}$ , then at the next timestep, the SoC has to be at least  $s_{des} - (t_{des} - t - 1) \cdot \Delta s(p_{max})$ , where  $(t_{des} - t - 1)$  denotes the number of remaining timesteps after this step. Now we can use the results derived in the previous section from Eq. (12) and get:

$$p \geq p_{min}^{des} := c^{bat}(s_{des} - (t_{des} - t - 1) \cdot \Delta s(p_{max}), s_t; \alpha_i)\quad (13)$$

Note that, if we start with an SoC that is already too low to achieve the goal SoC, the bounds will require an active power  $p > p_{max}$ , which is not possible, and  $p_{max}$  would be applied.

### C.3. Constraining battery controller

Now, we can finally combine all the previous considerations to define the controller that constrains the active power for the battery charging and discharging. We consider the following constraints:

- Direct constraints:  $p_{min} \leq p \leq p_{max}$
- SoC constraints:  $s_{min} \leq s_t \leq s_{max}$
- Charging constraint:  $s_t \geq s_{des}$  for  $t = t_{des}$

Note that we still use  $p := p_{t+1}$ . Using the formulas defined before, the last two constraints can be converted to constraints on  $p$  as shown in Eqs. (11), (12) and (13). Combining these constraints with the direct constraints on  $p$  and choosing always the tightest one yields:

$$\begin{aligned} p &\geq \tilde{p}_{min} := \max\{p_{min}, p_{min}^{des}, c^{bat}(s_{min}, s_t; \alpha_i)\} \\ p &\leq \tilde{p}_{max} := \min\{p_{max}, c^{bat}(s_{max}, s_t; \alpha_i)\} \end{aligned} \quad (14)$$

Finally, we can define our safety controller that assures that the chosen action, i.e. the active power, lies in the appropriate range.

$$f^{safe}(p) := clip_{[\tilde{p}_{min}, \tilde{p}_{max}]}(p) \quad (15)$$

where  $clip$  is the clipping function defined as follows:

$$clip_{[a,b]}(p) := \begin{cases} a & p \leq a \\ b & p \geq b \\ p & \text{else} \end{cases} \quad (16)$$

Note that the function  $f^{safe}(p)$  implicitly depends on a lot of parameters, i.e.  $p_{min}, p_{max}, s_{min}, s_{max}, s_{des}, s_t, t, t_{des}$  and the parameters of the model  $\alpha_i$  and not only on  $p$ .

## Appendix D. Data preparation

### D.1. DFAB data

The following variables are measured inside the DFAB unit and are processed as follows, before their usage in the data-driven learning process.

- **Room temperature** ( $r_t$ ): The room temperature contained a few data points at exactly 0 °C which were removed. Furthermore, sequences of constant temperature that lasted for at least one day were removed, too. In the next step, spikes in the temperature of a magnitude of at least 1.5 °C were extracted and deleted. Finally, we applied Gaussian smoothing with a standard deviation of 5.0.
- **Valves** ( $u_t$ ): The data measured for each individual valve only stated if the valve was open (1) or closed (0). Subsampling resulted in values in the interval [0.0, 1.0]. Since this series naturally contains long sequences of constant values, i.e. 0 or 1, we only removed constant sequences which lasted for at least 30 days.
- **Water temperatures** ( $h_t^{in}, h_t^{out}$ ): The water temperature of the heating water flowing into and out of the rooms was processed by removing all data points that did not lie in the range [10.0 °C, 50.0 °C] were removed and then smoothing with a Gaussian filter with a standard deviation of 5.0 was applied.

### D.2. Weather data

Outside temperature and solar irradiance are measured by the weather station at NEST. They were processed in the following way.

- **Outside temperature** ( $o_t$ ): First, we remove values that are constant for more than 30 min. In the next step, we fill values that are missing by linear interpolation between the last and the next known value, but only if the time interval of missing values was less than 45 minutes. Finally, we smooth the data with a Gaussian filter with a standard deviation of 2.0.

**Table 3**

Python libraries, used with Python [83] version 3.6.6.

| Library          | Version |
|------------------|---------|
| Numpy [84]       | 1.18.1  |
| TensorFlow [82]  | 1.14.0  |
| Keras [74]       | 2.3.1   |
| Hyperopt [66]    | 0.2.3   |
| Pandas [85]      | 0.25.3  |
| SkLearn [77]     | 0.22.1  |
| Matplotlib [86]  | 3.1.2   |
| OpenAI gym [87]  | 0.15.4  |
| Keras-RL [75]    | 0.4.2   |
| SciPy [88]       | 1.4.1   |
| Statsmodels [89] | 0.10.2  |

- **Irradiance** ( $i_t$ ): Since the irradiance data series naturally contains values that are constant for a long time, e.g. zero at night, we only remove a series of data points if they are constant for at least 20h. Then again we fill missing data points by interpolation and smooth the data as was done with the temperature data.

### D.3. EV battery data

The data of the battery consists of the state of charge (SoC) and the active power used to charge or discharge the battery. The two time series were processed as follows.

- **State of charge** ( $s_t^{bat}$ ): Since the SoC cannot lie outside of the interval [0.0%, 100.0%], we remove all values that lie outside that range including the boundary values. Further, if the data is exactly constant for at least 24h, we assume something went wrong with the data collection and remove the data of that time interval.
- **Active power** ( $p_t$ ): In this case, we do not have strict boundaries for the values, so we only remove values where the series was constant for at least 6 h.

## Appendix E. Implementation

The work was implemented in Python version 3.6.6 and is not compatible to versions 3.5 and lower since f-strings were used. The main libraries that were used are listed in Table 3. Note that the most recent version of all libraries was used, except for TensorFlow [82] because of a dependency on another library, Keras-RL [75]. In most cases, the produced code is Pep-8. The actual code can be accessed at [https://github.com/BratislavS/MIMO\\_DRL\\_Building\\_control](https://github.com/BratislavS/MIMO_DRL_Building_control). There is also information available on how to run the code.

### E.1. Data whitening

As another data processing step, we whitened the data, i.e. it was scaled to have mean 0.0 and variance 1.0 before training the models. This is a standard procedure in machine learning and helps to avoid a bias in the feature importance while also allowing task-independent weight initialization in the neural network training. Since this was done manually, without the use of an existing library, this resulted in a few complications. For example, the reinforcement learning environment took the original actions as input and then had to scale them, feed them to the model and scale the output of the model back to the original domain to get the output for the agent.

## References

- [1] Change IC, et al. Mitigation of climate change. Contribution of working group III to the fifth assessment report of the intergovernmental panel on climate change, 1454, 2014, <https://www.ipcc.ch/report/ar5/wg3/>.
- [2] Ramesh T, Prakash R, Shukla K. Life cycle energy analysis of buildings: An overview. *Energy Build* 2010;42(10):1592–600, <https://doi.org/10.1016/j.enbuild.2010.05.007>.
- [3] Shaikh PH, Nor NBM, Nallagownden P, Elamvazuthi I, Ibrahim T. A review on optimized control systems for building energy and comfort management of smart sustainable buildings. *Renew Sustain Energy Rev* 2014;34:409–29, <https://doi.org/10.1016/j.rser.2014.03.027>.
- [4] Huang P, Lovati M, Zhang X, Bales C. A coordinated control to improve performance for a building cluster with energy storage, electric vehicles, and energy sharing considered. *Appl Energy* 2020;268:114983, <https://doi.org/10.1016/j.apenergy.2020.114983>.
- [5] Chel A, Kaushik G. Renewable energy technologies for sustainable development of energy efficient building. *Alex Eng J* 2018;57(2):655–69, <https://doi.org/10.1016/j.aej.2017.02.027>.
- [6] Chwieduk D. Towards sustainable-energy buildings. *Appl Energy* 2003;76(1–3):211–7, [https://doi.org/10.1016/S0306-2619\(03\)00059-X](https://doi.org/10.1016/S0306-2619(03)00059-X).
- [7] Zhou Y, Cao S, Hensen JL, Lund PD. Energy integration and interaction between buildings and vehicles: A state-of-the-art review. *Renew Sustain Energy Rev* 2019;114:109337, <https://doi.org/10.1016/j.rser.2019.109337>.
- [8] Liu C, Chau K, Wu D, Gao S. Opportunities and challenges of vehicle-to-home, vehicle-to-vehicle, and vehicle-to-grid technologies. *Proc IEEE* 2013;101(11):2409–27, <https://doi.org/10.1109/JPROC.2013.2271951>.
- [9] Park JY, Nagy Z. Comprehensive analysis of the relationship between thermal comfort and building control research-A data-driven literature review. *Renew Sustain Energy Rev* 2018;82:2664–79, <https://doi.org/10.1016/j.rser.2017.09.102>.
- [10] Salisbury TI. A survey of control technologies in the building automation industry. *IFAC Proc Vol* 2005;38(1):90–100.
- [11] Verhelst J, Van Ham G, Saelens D, Helsen L. Model selection for continuous commissioning of HVAC-systems in office buildings: A review. *Renew Sustain Energy Rev* 2017;76:673–86, <https://doi.org/10.1016/j.rser.2017.01.119>.
- [12] Stluka P, Parthasarathy G, Gabel S, Samad T. Architectures and algorithms for building automation—An industry view. In: *Intelligent building control systems*. Springer; 2018, p. 11–43, [https://doi.org/10.1007/978-3-319-68462-8\\_2](https://doi.org/10.1007/978-3-319-68462-8_2).
- [13] Samad T, Bauer M, Bortoff S, Di Cairano S, Fagiano L, Odgaard PF, Rhinehart RR, Sánchez-Peña R, Serbezov A, Ankersen F, et al. Industry engagement with control research: Perspective and messages. *Annu Rev Control* 2020;49:1–14, <https://doi.org/10.1016/j.arcontrol.2020.03.002>.
- [14] Skogestad S, Postlethwaite I. Multivariable feedback control: analysis and design, vol. 2. Citeseer; 2007, <https://dl.acm.org/doi/abs/10.5555/525126>.
- [15] Privara S, Cigler J, Váňa Z, Oldewurtel F, Sagerschnig C, Žáčková E. Building modeling as a crucial part for building predictive control. *Energy Build* 2013;56:8–22, <https://doi.org/10.1016/j.enbuild.2012.10.024>.
- [16] Jain A, Nghiem T, Morari M, Mangharam R. Learning and control using Gaussian processes. In: *2018 ACM/IEEE 9th international conference on cyber-physical systems (ICCCPS)*. IEEE; 2018, p. 140–9, <https://doi.org/10.1109/ICCCPS.2018.00022>.
- [17] Seralle G, Fiorentini M, Capozzoli A, Bernardini D, Bemporad A. Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities. *Energies* 2018;11(3):631, <https://doi.org/10.3390/en11030631>.
- [18] Oldewurtel F, Parisio A, Jones CN, Gyalistras D, Gwerder M, Stauch V, Lehmann B, Morari M. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy Build* 2012;45:15–27, <https://doi.org/10.1016/j.enbuild.2011.09.022>.
- [19] Xu X, Wang S, Huang G. Robust MPC for temperature control of air-conditioning systems concerning on constraints and multitype uncertainties. *Build Serv Eng Res Technol* 2010;31(1):39–55, <https://doi.org/10.1177/0143624409352420>.
- [20] Tanaskovic M, Sturzenegger D, Smith R, Morari M. Robust adaptive model predictive building climate control. *Ifac-Papersonline* 2017;50(1):1871–6, <https://doi.org/10.1016/j.ifacol.2017.08.257>.
- [21] Khosravi M, Schmid N, Eichler A, Heer P, Smith RS. Machine learning-based modeling and controller tuning of a heat pump. In: *J Phys: Conf Ser*. 1343, (1):IOP Publishing; 2019, 012065, <https://doi.org/10.1088/1742-6596/1343/1/012065>.
- [22] Khosravi M, Behrunani V, Myszkowski P, Smith RS, Rupenyan A, Lygeros J. Performance-driven cascade controller tuning with Bayesian optimization. *IEEE Trans Ind Electron* 2021, <https://doi.org/10.1109/TIE.2021.3050356>.
- [23] Aswani A, Gonzalez H, Sastry SS, Tomlin C. Provably safe and robust learning-based model predictive control. *Automatica* 2013;49(5):1216–26, <https://doi.org/10.1016/j.automatica.2013.02.003>.
- [24] Aswani A, Master N, Taneja J, Culler D, Tomlin C. Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control. *Proc IEEE* 2011;100(1):240–53, <https://doi.org/10.1109/JPROC.2011.2161242>.
- [25] Chen Y, Shi Y, Zhang B. Optimal control via neural networks: A convex approach. 2018, arXiv:1805.11835. <https://ui.adsabs.harvard.edu/abs/2018arXiv180511835C>.
- [26] Smarra F, Jain A, de Rubeis T, Ambrosini D, D'Innocenzo A, Mangharam R. Data-driven model predictive control using random forests for building energy optimization and climate control. *Appl Energy* 2018;226:1252–72, <https://doi.org/10.1016/j.apenergy.2018.02.126>.
- [27] Büning F, Huber B, Heer P, Aboudonia A, Lygeros J. Experimental demonstration of data predictive control for energy optimization and thermal comfort in buildings. *Energy Build* 2020;211:109792, <https://doi.org/10.1016/j.enbuild.2020.109792>.
- [28] Sutton RS, Barto AG. Reinforcement learning: An introduction. MIT Press; 2018, [https://doi.org/10.1016/S1364-6613\(99\)01331-5](https://doi.org/10.1016/S1364-6613(99)01331-5).
- [29] LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature* 2015;521(7553):436–44, <https://doi.org/10.1038/nature14539>.
- [30] Arulkumaran K, Deisenroth MP, Brundage M, Bharath AA. Deep reinforcement learning: A brief survey. *IEEE Signal Process Mag* 2017;34(6):26–38, <https://doi.org/10.1109/MSP.2017.2743240>.
- [31] Mnih V, Kavukcuoglu K, Silver D, Rusu AA, Veness J, Bellemare MG, Graves A, Riedmiller M, Fidjeland AK, Ostrovski G, et al. Human-level control through deep reinforcement learning. *Nature* 2015;518(7540):529–33, <http://dx.doi.org/10.1038/nature14236>.
- [32] Silver D, Schrittwieser J, Simonyan K, Antonoglou I, Huang A, Guez A, Hubert T, Baker L, Lai M, Bolton A, et al. Mastering the game of go without human knowledge. *Nature* 2017;550(7676):354–9, [doi:10.1038/nature24270](https://doi.org/10.1038/nature24270).
- [33] Young T, Hazarika D, Poria S, Cambria E. Recent trends in deep learning based natural language processing. *IEEE Comput Intell Mag* 2018;13(3):55–75, <https://doi.org/10.1109/MCI.2018.2840738>.
- [34] Esteva A, Robicquet A, Ramsundar B, Kuleshov V, DePristo M, Chou K, Cui C, Corrado G, Thrun S, Dean J. A guide to deep learning in healthcare. *Nature Med* 2019;25(1):24–9, <https://doi.org/10.1038/s41591-018-0316-z>.
- [35] Wang Z, Hong T. Reinforcement learning for building controls: The opportunities and challenges. *Appl Energy* 2020;269:115036, <https://doi.org/10.1016/j.apenergy.2020.115036>.
- [36] Vázquez-Canteli JR, Nagy Z. Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Appl Energy* 2019;235:1072–89, <https://doi.org/10.1016/j.apenergy.2018.11.002>.
- [37] Mason K, Grijalva S. A review of reinforcement learning for autonomous building energy management. *Comput Electr Eng* 2019;78:300–12, <https://doi.org/10.1016/j.compeleceng.2019.07.019>.
- [38] Lillicrap TP, Hunt JJ, Pritzel A, Heess N, Erez T, Tassa Y, Silver D, Wierstra D. Continuous control with deep reinforcement learning. 2015, arXiv preprint arXiv:1509.02971. <https://ui.adsabs.harvard.edu/abs/2015arXiv150902971L>.
- [39] Ruano AE, Crispim EM, Conceicao EZ, Lúcio MMJ. Prediction of building's temperature using neural networks models. *Energy Build* 2006;38(6):682–94, <https://doi.org/10.1016/j.enbuild.2005.09.007>.
- [40] Mustafaraj G, Lowry G, Chen J. Prediction of room temperature and relative humidity by autoregressive linear and nonlinear neural network models for an open office. *Energy Build* 2011;43(6):1452–60, <https://doi.org/10.1016/j.enbuild.2011.02.007>.
- [41] Taylor ME, Stone P. Transfer learning for reinforcement learning domains: A survey. *J Mach Learn Res* 2009;10(7). <https://www.jmlr.org/papers/volume10/taylor09a/taylor09a.pdf>.
- [42] Xu S, Wang Y, Wang Y, O'Neill Z, Zhu Q. One for many: Transfer learning for building HVAC control. In: *Proceedings of the 7th ACM international conference on systems for energy-efficient buildings, cities, and transportation*. 2020, p. 230–9, <https://doi.org/10.1145/3408308.3427617>.
- [43] Mocanu E, Mocanu DC, Nguyen PH, Liotta A, Webber ME, Gibescu M, Sloatweg JG. On-line building energy optimization using deep reinforcement learning. *IEEE Trans Smart Grid* 2018;10(4):3698–708, <https://doi.org/10.1109/TSG.2018.2834219>.
- [44] Wei T, Wang Y, Zhu Q. Deep reinforcement learning for building HVAC control. In: *2017 54th ACM/EDAC/IEEE design automation conference (DAC)*. 2017, p. 1–6, <https://doi.org/10.1145/3061639.3062224>.
- [45] 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, <https://doi.org/10.1016/j.enbuild.2018.03.051>. URL <http://www.sciencedirect.com/science/article/pii/S0378778818302184>.
- [46] Chen Y, Shi Y, Zhang B. Modeling and optimization of complex building energy systems with deep neural networks. 2017, ArXiv E-Prints arXiv:1711.02278. <https://ui.adsabs.harvard.edu/abs/2017arXiv171102278C>.
- [47] Afram A, Janabi-Sharifi F, Fung AS, Raahemifar K. Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy Build* 2017;141:96–113, <https://doi.org/10.1016/j.enbuild.2017.02.012>. URL <http://www.sciencedirect.com/science/article/pii/S0378778816310799>.
- [48] 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). <https://doi.org/10.3390/pr5030046>. URL <https://www.mdpi.com/2227-9717/5/3/46>.

- [49] Ruelens F, Claessens BJ, Vandael S, De Schutter B, Babuška R, Belmans R. Residential demand response of thermostatically controlled loads using batch reinforcement learning. *IEEE Trans Smart Grid* 2017;8(5):2149–59, <https://doi.org/10.1109/TSG.2016.2517211>.
- [50] Shi W, Wong VWS. Real-time vehicle-to-grid control algorithm under price uncertainty. In: 2011 IEEE international conference on smart grid communications (SmartGridComm). 2011, p. 261–6, <https://doi.org/10.1109/SmartGridComm.2011.6102330>.
- [51] Chiş A, Lundén J, Koivunen V. Optimization of plug-in electric vehicle charging with forecasted price. In: 2015 IEEE international conference on acoustics, speech and signal processing (ICASSP). 2015, p. 2086–9, <https://doi.org/10.1109/ICASSP.2015.7178338>.
- [52] Chiş A, Lundén J, Koivunen V. Scheduling of plug-in electric vehicle battery charging with price prediction. In: IEEE PES ISGT Europe 2013. 2013, p. 1–5, <https://doi.org/10.1109/ISGTEurope.2013.6695263>.
- [53] Sadeghianpourhamami N, Deleu J, Develder C. Definition and evaluation of model-free coordination of electrical vehicle charging with reinforcement learning. *IEEE Trans Smart Grid* 2019;1, <https://doi.org/10.1109/TSG.2019.2920320>.
- [54] Ko H, Pack S, Leung VCM. Mobility-aware vehicle-to-grid control algorithm in microgrids. *IEEE Trans Intell Transp Syst* 2018;19(7):2165–74, <https://doi.org/10.1109/ITS.2018.2816935>.
- [55] Vandael S, Claessens B, Ernst D, Holvoet T, Deconinck G. Reinforcement learning of heuristic EV fleet charging in a day-ahead electricity market. *IEEE Trans Smart Grid* 2015;6(4):1795–805, <https://doi.org/10.1109/TSG.2015.2393059>.
- [56] Kim S, Lim H. Reinforcement learning based energy management algorithm for smart energy buildings. *Energies* 2018;11(8). <http://dx.doi.org/10.3390/en11082010>, <https://www.mdpi.com/1996-1073/11/8/2010>.
- [57] Nguyen HT, Nguyen DT, Le LB. Energy management for households with solar assisted thermal load considering renewable energy and price uncertainty. *IEEE Trans Smart Grid* 2015;6(1):301–14, <https://doi.org/10.1109/TSG.2014.2350831>.
- [58] Kim B, Ren S, van der Schaar M, Lee J. Bidirectional energy trading and residential load scheduling with electric vehicles in the smart grid. *IEEE J Sel Areas Commun* 2013;31(7):1219–34, <https://doi.org/10.1109/JSAC.2013.130706>.
- [59] Nguyen DT, Le LB. Joint optimization of electric vehicle and home energy scheduling considering user comfort preference. *IEEE Trans Smart Grid* 2014;5(1):188–99, <http://dx.doi.org/10.1109/TSG.2013.2274521>.
- [60] Keller R. DFAB HOUSE. 2021, [https://www.swiss-architects.com/de/architecture-news/bau-der-woche/building-digitally-living-digitally?utm\\_source=newsletter&utm\\_medium=email&utm\\_campaign=2417](https://www.swiss-architects.com/de/architecture-news/bau-der-woche/building-digitally-living-digitally?utm_source=newsletter&utm_medium=email&utm_campaign=2417). [Accessed 20.05.2021].
- [61] DFAB HOUSE – digital fabrication and living, empa, duebendorf, Switzerland. 2021, <https://www.empa.ch/web/nest/digital-fabrication>, [Accessed: 20.05.2021].
- [62] Nest – exploring the future of buildings, swiss federal laboratories for materials science and technology - EMPA, duebendorf, Switzerland. 2021, <https://www.empa.ch/web/nest>, [Accessed: 20.05.2021].
- [63] Lipton ZC, Berkowitz J, Elkan C. A critical review of recurrent neural networks for sequence learning. 2015, arXiv preprint [arXiv:1506.00019](https://arxiv.org/abs/1506.00019).
- [64] Kingma DP, Ba J. Adam: A method for stochastic optimization. 2014, ArXiv E-Prints [arXiv:1412.6980](https://arxiv.org/abs/1412.6980). <https://ui.adsabs.harvard.edu/abs/2014arXiv1412.6980K>.
- [65] Bergstra J, Bardenet R, Bengio Y, Kégl B. Algorithms for hyper-parameter optimization. In: Proceedings of the 24th international conference on neural information processing systems. NIPS'11, Red Hook, NY, USA: Curran Associates Inc.; 2011, p. 2546–54, <https://hal.inria.fr/hal-00642998/>.
- [66] Bergstra J, Komer B, Eliasmith C, Yamins D, Cox DD. Hyperopt: a python library for model selection and hyperparameter optimization. *Comput Sci Discov* 2015;8(1):014008, <https://doi.org/10.1088/1749-4699/8/1/014008>. URL <http://stacks.iop.org/1749-4699/8/i=1/a=014008>.
- [67] Dow J. Tesla releases 'long range plus' model s/x with 390/351 mile range, new wheels. 2020, <https://electrek.co/2020/02/14/tesla-releases-long-range-plus-model-s-x-with-390-351-mile-range-new-wheels/>, [Accessed: 20.05.2021].
- [68] Brodie J. AutoExpress: BMW i3 updated with more range and new trim options. 2018, <https://www.autoexpress.co.uk/bmw/i3/104793/bmw-i3-updated-with-more-range-and-new-trim-options>, [Accessed: 20.05.2021].
- [69] Taylor M. Better late than never as EQC leads mercedes-benz EV assault. 2018, <https://www.forbes.com/sites/michaeltaylor/2018/09/04/better-late-than-never-as-eqc-leads-mercedes-benz-ev-assault/?sh=25d6245c5f5a>, [Accessed: 20.05.2021].
- [70] Qiu C, Hu Y, Chen Y, Zeng B. Deep deterministic policy gradient (DDPG)-based energy harvesting wireless communications. *IEEE Internet Things J* 2019;6(5):8577–88, <https://doi.org/10.1109/JIOT.2019.2921159>.
- [71] Vecerik M, Hester T, Scholz J, Wang F, Pietquin O, Piot B, Heess N, Rothörl T, Lampe T, Riedmiller M. Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. 2017, ArXiv E-Prints [arXiv:1707.08817](https://arxiv.org/abs/1707.08817). <https://ui.adsabs.harvard.edu/abs/2017arXiv170708817V>.
- [72] Xiong X, Wang J, Zhang F, Li K. Combining deep reinforcement learning and safety based control for autonomous driving. 2016, ArXiv E-Prints [arXiv:1612.00147](https://arxiv.org/abs/1612.00147). <https://ui.adsabs.harvard.edu/abs/2016arXiv161200147X>.
- [73] Yang Q, Zhu Y, Zhang J, Qiao S, Liu J. UAV air combat autonomous maneuver decision based on DDPG algorithm. In: 2019 IEEE 15th international conference on control and automation (ICCA). 2019, p. 37–42, <https://doi.org/10.1109/ICCA.2019.8899703>.
- [74] Chollet F, et al. Keras. 2015, <https://keras.io>.
- [75] Plappert M. Keras-rl. GitHub Repository 2016. <https://github.com/keras-rl/keras-rl>.
- [76] Finch S. Ornstein-uhlenbeck process. 2004, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.710.4200>.
- [77] Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E. Scikit-learn: Machine learning in Python. *J Mach Learn Res* 2011;12:2825–30, [https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf?source=post\\_page-----](https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf?source=post_page-----).
- [78] De Coninck R, Helsen L. Practical implementation and evaluation of model predictive control for an office building in Brussels. *Energy Build* 2016;111:290–8, <https://doi.org/10.1016/j.enbuild.2015.11.014>.
- [79] Dulac-Arnold G, Levine N, Mankowitz DJ, Li J, Paduraru C, Gowal S, Hester T. An empirical investigation of the challenges of real-world reinforcement learning. 2020, arXiv preprint [arXiv:2003.11881](https://arxiv.org/abs/2003.11881).
- [80] Zhao W, Qeralta J, Westerlund T. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In: 2020 IEEE symposium series on computational intelligence (SSCI). IEEE; 2020, p. 737–44.
- [81] Chebotar Y, Waters A. Distilling knowledge from ensembles of neural networks for speech recognition. In: Interspeech. 2016, p. 3439–43.
- [82] Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X. TensorFlow: Large-scale machine learning on heterogeneous systems. 2015, URL <http://tensorflow.org/>. <https://arxiv.org/abs/1603.04467>.
- [83] Rossum G. Python reference manual. Technical report, NLD: CWI (Centre for Mathematics and Computer Science); 1995, <https://dl.acm.org/doi/abs/10.5555/869369>.
- [84] Oliphant TE. A Guide To NumPy, vol. 1. Trelgol Publishing USA; 2006, <https://ecs.wgtn.ac.nz/foswiki/pub/Support/ManualPagesAndDocumentation/numpybook.pdf>.
- [85] McKinney W. Pandas: a foundational python library for data analysis and statistics. *Python High Perform Sci Comput* 2011;14. [https://www.dlr.de/sc/portaldat/15/resources/dokumente/pyhpc2011/submissions/pyhpc2011\\_submission\\_9.pdf](https://www.dlr.de/sc/portaldat/15/resources/dokumente/pyhpc2011/submissions/pyhpc2011_submission_9.pdf).
- [86] Hunter JD. Matplotlib: A 2D graphics environment. *Comput Sci Eng* 2007;9(3):90–5, <https://doi.org/10.1109/MCSE.2007.55>.
- [87] Brockman G, Cheung V, Pettersson L, Schneider J, Schulman J, Tang J, Zaremba W. OpenAI Gym. 2016, [arXiv:arXiv:1606.01540](https://arxiv.org/abs/1606.01540).
- [88] Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Jarrod Millman K, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey C, Polat I, Feng Y, Moore EW, Vand erPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P, Contributors S. SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods* 2020. <https://doi.org/10.1038/s41592-019-0686-2>.
- [89] Seabold S, Perktold J. Statsmodels: Econometric and statistical modeling with python. In: 9th Python in science conference. 2010, <https://pdfs.semanticscholar.org/3a27/6417e5350e29cb6bf04ea54785601d5a215.pdf>.