
Zero-Shot Building Control

Scott R. Jeen¹ Jonathan M. Cullen¹
University of Cambridge

Abstract

Heating and cooling systems in buildings account for 31% of global energy use, much of which are regulated by Rule Based Controllers (RBCs) that neither maximise energy efficiency nor minimise emissions by interacting optimally with the grid. Control via Reinforcement Learning (RL) has been shown to significantly improve building energy efficiency, but existing solutions require pre-training in simulators that are prohibitively expensive to obtain for every building in the world. In response, we show it is possible to perform safe, zero-shot control of buildings by combining ideas from system identification and model-based RL. We call this combination PEARL (Probabilistic Emission-Abating Reinforcement Learning) and show it reduces emissions without pre-training, needing only a three hour commissioning period. In experiments across three varied building energy simulations, we show PEARL outperforms an existing RBC once, and popular RL baselines in all cases, reducing building emissions by as much as 31% whilst maintaining thermal comfort.

1. Introduction

Heating and cooling systems in buildings account for 31% of global energy use, primarily in managing occupant thermal comfort and hygiene (Cullen & Allwood, 2010). Such systems are usually regulated by rule-based controllers (RBCs) that take system temperature as input, use a temperature setpoint as an objective, and actuate equipment to minimise the error between objective and state. Whilst usefully simple, RBCs do not maximise energy efficiency, nor do they interact with the grid to draw power when carbon intensity is low to minimise emissions. New techniques that demonstrate this capability across generalised settings would prove valuable climate change mitigation tools.

An option for obtaining optimal control policies across var-

¹Department of Engineering, University of Cambridge, Cambridge, UK. Correspondence to: Scott Jeen <srj38@cam.ac.uk>.

ied, complex settings is Reinforcement Learning (RL) (Sutton & Barto, 2018). Where existing advanced control techniques, like Model Predictive Control (MPC), require the specification of a dynamical model that can be expensive to obtain (Camacho & Alba, 2013), RL’s strength is in obtaining policies *tabula rasa*, and updating system understanding *online* as the environment evolves. Recent successes in complex physics tasks (Lillicrap et al., 2015), gaming (Silver et al., 2017), and robotics (Gu et al., 2017) have highlighted these characteristics. With this generality, one can imagine RL agents being placed in *any* energy-intensive setting, and feasibly learning to control them more efficiently, minimising emissions.

Recent applications of RL to building control have indeed shown marked energy efficiency improvements over conventional controllers. Chen et al. (2018) designed a Q-learning agent that controlled HVAC and window actuation in a residential building with 23% less energy than the existing RBC. Similarly, Zhang et al. (2019b) trained an A3C agent in a building energy simulation of an office, deployed it in the real building, and found heating demand was reduced by 16.7%. Despite these impressive results, most contributions require training an agent in a simulation of the real environment, often using millions of timesteps of data. Creating such simulations is time-consuming, and sometimes impossible if the correct building data cannot be obtained in advance, limiting deployment to the fraction of global buildings for which we have this data *a priori*. An alternative, as proposed by Lazic et al. (2018), is to train agents during a short commissioning period of 3-12 hours in the real environment. Doing so is suitably general for scaled deployment, but their agent only improved cooling costs by 9% in experiments conducted for one day, poorer performance than the agents trained in simulation. Eliciting the performance of model-free agents trained in simulation whilst training on limited data from real environments remains an open problem.

In this paper, our primary contribution is showing deep RL systems for building control can be deployed safely without pre-training. We achieve this by amalgamating existing works into a new approach called PEARL (Probabilistic Emission-Abating Reinforcement Learning) that performs system identification during a 3 hour commissioning pe-

riod, models the system dynamics using an ensemble of probabilistic networks, and selects actions by pairing probabilistic planning with Model Predictive Control (MPC). In experiments across three varied building simulations, we show our approach can reduce annual emissions by 31.46% compared with an RBC whilst maintaining thermal comfort, and outperforms all RL baselines used for comparison. We carry out further experiments exploring the sensitivity of results to commissioning period length, deep network configuration and optimiser choice, finding the results to be largely invariant to commissioning period length (after the first few weeks of operation) and network setup, but sensitive to optimiser choice. By sidestepping pre-training in simulation, systems like PEARL could be deployed at far greater scale than previously envisioned, widening their applicability as climate change mitigation tools.

2. Related Work

Previous research on RL for building control has focused mainly on model-free algorithms (Yu et al., 2021). In this setting, the agent uses data collected from the environment to learn a policy π that maps states s_t from a state-space \mathcal{S} to a probability distribution over the action-space \mathcal{A} i.e. $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$. To obtain an optimal policy, the agent must necessarily visit many states $s \in \mathcal{S}$ and trial many actions $a \in \mathcal{A}$. Doing so is data inefficient, with Deep-Q Learning (Mnih et al., 2013), Deep Deterministic Policy Gradient (DDPG) (Lillicrap et al., 2015), and Proximal Policy Optimisation (PPO) (Schulman et al., 2017) each taking of the order 10^7 interactions in complex environments to obtain optimal policies. Such sample inefficiency has been corroborated in the building control literature. Wei et al. (2017) train a Deep-Q agent to control the HVAC equipment of a 5-zone building with 35% reduction in energy cost, but require 8 years of simulated data for training. Similarly, Valladares et al. (2019) require 10 years of simulated data to train a Double-Q agent to control the HVAC in a university classroom with 5% energy savings. Such data-intensive training can only be achieved in bespoke building simulators that are time-consuming to create, or impossible to specify, in most cases.

In contrast, model-based RL algorithms have demonstrated better sample efficiency. Here, the agent learns a mapping from current state-action pair (s_t, a_t) to next state s_{t+1} , called the system dynamics or transition function i.e. $f_\theta : (s_t, a_t) \rightarrow s_{t+1}$. The agent then uses the model to predict the expected reward of a range of candidate action sequences $a_{t:t+H}$ to time horizon H .

A popular choice of function approximator is the Gaussian Process (GP) which provides uncertainty quantification in its predictions, and works well with small datasets (Williams & Rasmussen, 2006). PILCO learned the cartpole swing-

up task after 7 trials by using GPs to model the transition function (Deisenroth & Rasmussen, 2011), before Kamthe & Deisenroth (2018) demonstrated improved performance by combining GPs with model predictive control (MPC) to solve the same task after 9 seconds of interaction. Jain et al. (2018) used a similar approach to curtail hotel energy-use during a simulated demand response event having trained their model on only a few days of data. Although GPs are the favoured choice in the low-data regime, exact inference using a data set of size n has complexity $\mathcal{O}(n^3)$ which becomes intractable with more than a few thousand samples (Hensman et al., 2013). This limits their applicability for modelling building transition functions with arbitrarily large feature-spaces or training sets.

An alternative approach is to model the transition function using Deep Neural Networks (DNNs). Nagabandi et al. (2018) combined DNNs with MPC to solve the *Swimmer* task on the MuJoCu benchmark using 20x fewer datapoints than a model-free approach. In the building control setting, Zhang et al. (2019a) built a similar agent that reduced energy consumption in a data centre by 21.8% with 10x fewer training steps than a model-free algorithm. Jain et al. (2020) then tested the algorithm in-situ, finding an 8% energy reduction. Despite encouraging data efficiency, the performance of such agents is hampered by overfitting, or *model bias*, in the low-data regime, causing poor generalisation to unobserved transitions. Ding et al. (2020) attempt to mitigate model bias by deploying an ensemble of DNNs to model the transition function of a large multi-zone building showing they can achieve 8.2% energy savings in 10.5x fewer timesteps than model-free approaches. Although the ensemble allows for the quantification of epistemic uncertainty, aleatoric uncertainty remains unquantified, potentially limiting performance.

To the best of our knowledge, Lazic et al. (2018) is the only work that attempts to learn a zero-shot building control policy by interacting with the real environment. Their agent fits a linear model of a datacentre’s thermal dynamics using data obtained in a three hour commissioning period, and selects actions by optimising planned trajectories using MPC. During commissioning, the agent explores the state-action space by performing a uniform random walk in each control variable, bounded to a safe operating range informed by historical data. Their choice of model expedites learning, but limits agent performance as the building’s non-linear dynamics are erroneously linearised. In this study we aim to preserve the data efficiency of this approach whilst improving performance with expressive deep networks.

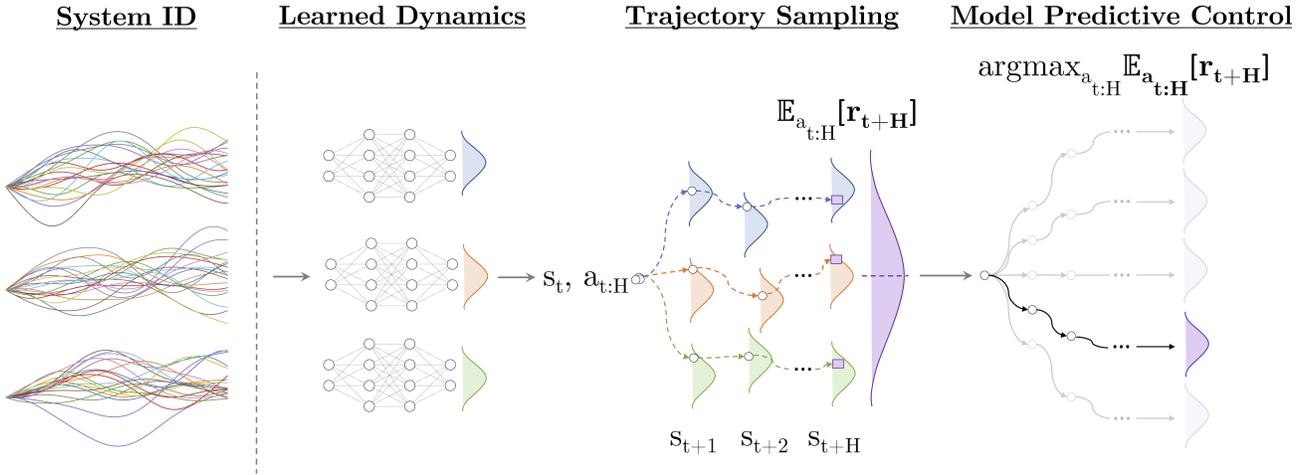


Figure 1. Our end-to-end deep reinforcement learning approach: PEARL. **System ID**: Temporary policy $\pi_{systemID}$ used to collect diverse dataset during commissioning period. **Learned Dynamics**: modelled as an ensemble of probabilistic deep neural networks. **Trajectory Sampling**: probabilistic planning used to predict future rewards (purple squares) that form a distribution at time horizon H . **MPC**: optimises for the trajectory with highest expected reward.

3. PEARL: Probabilistic Emission-Abating Reinforcement Learning

We consider the standard reinforcement learning setup in which an agent takes actions in an environment at discrete timesteps to maximise the cumulative sum of future rewards. We model the environment as an infinite-horizon Markov Decision Process (MDP), with state-space $\mathcal{S} \in \mathbb{R}^{d_s}$, action-space $\mathcal{A} \in \mathbb{R}^{d_a}$, reward function $r(s_t, a_t)$, and transition dynamics $p(s_{t+1}|s_t, a_t)$. Despite this formulation, the state s_t is not fully observed, and is instead partially observed via an observation space $\mathcal{O} \in \mathbb{R}^{d_o}$. To allow us to apply standard RL techniques for MDPs, the agent is passed a recent trajectory of observations as a state representation i.e. $s_t = (o_t, o_{t-1}, \dots, o_{t-h})$ where h is the history length (Kaelbling et al., 1998). The agent selects actions using its policy $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$, which it manipulates to maximise return $G_t = \sum_H \gamma^H r(s_{t+H}, a_{t+H})$, where $\gamma \in [0, 1]$ is a discount parameter, and H is a finite time horizon in the general case. In our case we do not discount and set $\gamma = 1$.

3.1. Prediction

We follow the schema of model-based reinforcement learning where our task is to fit a function \tilde{f}_θ that approximates the true forward dynamics of the system $f(s_{t+1}, (s_t, a_t))$ given a dataset of experience collected from the environment $\mathcal{D} = [s_{n+1}, (s_n, a_n)]_{n=1}^N$. We employ probabilistic DNNs to learn this mapping, which provide data-efficient approximations of complex system dynamics and allow agents

to incorporate prediction uncertainty into action selection (Gal et al., 2016; Higuera et al., 2018). Where traditional, *deterministic* DNNs output point predictions given an input, here our probabilistic DNNs output distributions over the output nodes parameterised by a multivariate Gaussian distribution with mean μ and diagonal covariance matrix Σ ; i.e. $\tilde{f}_\theta(s_t, a_t) = \mathcal{N}(\mu_\theta(s_t, a_t), \Sigma_\theta(s_t, a_t))$. The agent maximises the likelihood of a target variable being drawn from the predicted distribution i.e. it performs Maximum Likelihood Estimation (MLE):

$$Loss(\theta) = \sum_{n=1}^N -\log(P(s_n; \mu_\theta, \Sigma_\theta)) \quad (1)$$

By outputting a distribution over the next state our network can quantify *prediction* uncertainty. Ensembling multiple probabilistic DNNs, training each on different subsets of the data, and averaging over their predictions can quantify *model* uncertainty (Lakshminarayanan et al., 2017). Here, we employ K -many models, averaging the predictions to ensure both types of uncertainty are captured:

$$\tilde{f}_\theta(s_t, a_t) = \frac{1}{K} \sum_{k=1}^K \tilde{f}_{\theta_k}(s_t, a_t) \quad (2)$$

3.2. Control

Between interactions with the environment, our agent plans and optimises a sequence of actions using Model Predictive Control (MPC) (Camacho & Alba, 2013). In

the absence of a tractable, white-box dynamics model, MPC normally optimises via the random shooting method where a matrix of candidate action sequences $\mathbf{a}_{t:t+H} \doteq \{\mathbf{a}_t, \dots, \mathbf{a}_{t+H}\}$ is populated, passed through the dynamics model \tilde{f}_θ to time horizon H , and optimised by selecting the first action of the sequence producing highest expected reward: $\text{argmax}_{\mathbf{a}_{t:t+H}} \sum_{i=t}^{t+H} \mathbb{E}_{\tilde{f}_\theta} [r(s_i, a_i)]$. However, random shooting in its traditional setup cannot account for uncertainty in the predictions of a probabilistic dynamics model.

Instead, we combine the cross-entropy method of optimisation (Botev et al., 2013) with Trajectory Sampling (Chua et al., 2018). Like the random shooting method, a matrix of action sequences is populated by sampling actions from an initially arbitrary Gaussian distribution. Then, we duplicate P -many state-action pairs at the first planning timestep $(s_t^p, a_t) \forall p \in P$ called *particles*. Each particle is assigned one bootstrap from the dynamics function ensemble and iteratively passed through it to create next-state distributions which can be sampled: $s_{t+1}^p \sim \mathcal{N}(\mu_{t+1}^p, \Sigma_{t+1}^p; \theta)$. After all initial candidate sequences have been trialled, the parameters of the action-selection distribution are updated to reflect the mean and variance of sequences that produced the top 10% of expected rewards (the so-called *elite* sequences), from which a new action sequence matrix is sampled. After n optimisation iterations (in our case $n = 5$ was sufficient), the final elite sequence means represent the optimal action sequence, and the first action of this sequence is performed in the real environment.

3.3. System Identification

Following Lazic et al. (2018), we grant the agent a window, or *commissioning period*, to explore the state-action space and fit a dynamics function. During such time, the agent performs a uniform random walk in each control variable bounded to a safe operating range informed by historical data if available, or initialised conservatively and expanded if no data is available. Stepwise changes in action selection are bounded to mitigate potential hardware damage that may be caused by large swings in chosen actions.

Formally system identification is performed as follows. Let $a^i[t]$ be the value of action i at timestep t , with $i \in \mathcal{A}$. Let $[a_{min}^i, a_{max}^i]$ be the safe operating range for action i , and let Δ^i be the maximum allowable change in action a^i between timesteps. The system identification strategy, $\pi_{systemID}$, is therefore:

$$a^i[t+1] = \max(a_{min}^i, \min(a_{max}^i, a^i[t] + v^i)), \\ v^i \sim \text{Uniform}(-\Delta^i, \Delta^i) \quad (3)$$

Model parameters are updated iteratively when the number of timesteps observed in the environment equals the batch size of the model. After the commissioning period,

Algorithm 1 PEARL

```

Initialize memory  $\mathcal{D}$  and dynamics model  $\tilde{f}(s_t, a_t)_\theta$ 
# Commissioning Period
for t=0 to Commissioning Steps do
     $a_{t+1} \sim \pi_{systemID}(s_t, a_t)$ 
    store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
end for
update  $\tilde{f}(s_t, a_t)_\theta$  given  $\mathcal{D}$ 
# Control
for t=Commissioning Steps to T do
    for  $i = 1$  to No. Planning Sequences do
        sample action sequence:  $a_{t:t+H}^i \sim \pi_{CEM}$ 
        plan:  $s_{t+H}^p = \tilde{f}(s_t^p, a_{t:t+H}^i)_\theta \forall p \in P$ 
        evaluate exp. reward:  $\sum_{t=t}^{t+H} \frac{1}{P} \sum_{p=1}^P r(s_t^p, a_t)$ 
        update  $\pi_{CEM}$  using elite sequences
    end for
    execute  $a_t^*$  from sequence with highest exp. reward
    store transition  $(s_t, a_t, r_t, s_{t+1})$  in  $\mathcal{D}$ 
    if end of day then
        update  $\tilde{f}(s_t, a_t)_\theta$  given  $\mathcal{D}$ 
    end if
end for

```

parameters are updated at the end of each day.

Lazic et al. (2018) showed that active exploration produces a more informative dataset for dynamics learning, allowing the agent to sample regions of state-action space that may not be observed during normal operation. Doing so allowed their agent to reduce energy-use by 10% over one trained using a $\sim 50x$ larger dataset of non-exploratory, steady-state behaviour from the existing building controller. Because of their results, we chose to implement their method directly, leaving the task of further optimising the system identification policy for future work.

Figure 1 provides an illustration of the agent’s setup: system identification, prediction (learned dynamics and trajectory sampling) and control. We call the resulting algorithm Probabilistic Emission-Abating Reinforcement Learning (PEARL) - Algorithm 1

3.4. Reward Function

The reward $r[t]$ is a linear combination of an *emissions term* and a *temperature term*. Our goal is to motivate the agent to minimise emissions whilst satisfying thermal comfort in the building. Let $r_E[t]$ be the emissions-term reward at timestep t , $E[t]$ be the total energy consumption in the environment at time t , and $C[t]$ be the grid carbon intensity at time t , then the emissions-term reward is

$$r_E[t] = -\phi(E[t]C[t]) \quad (4)$$

where ϕ is a tunable parameter that sets the relative emphasis of emission-minimisation over thermal comfort. The reward is negative because our goal is to minimise emissions, or maximise the negative of emissions produced. Let $r_T^i[t]$ be the temperature-term reward at timestep t for thermal zone i , T_{obs}^i be the observed temperature in thermal zone i , and T_{low} and T_{high} be the lower and upper temperature bounds on thermal comfort respectively. The temperature reward is then given by

$$r_T^i[t] = \begin{cases} 0 & : T_{low} \leq T_{obs}^i \leq T_{high} \\ -\theta \min[(T_{low} - T_{obs}^i[t])^2, (T_{high} - T_{obs}^i)^2] & : \text{otherwise} \end{cases} \quad (5)$$

where θ is a second tunable parameter that sets the relative emphasis of thermal comfort over emission-minimisation. The second term in Equation 5 can be thought of as a *penalty term* that punishes the agent in proportion to deviations from the thermal comfort zone. This allows the agent to temporarily prioritise emission-minimization over thermal comfort if the savings are deemed sufficient. The total temperature reward $r_T[t]$ is obtained by summing the rewards across thermal zones i.e. $\sum_{i=1}^N r_T^i[t]$.

4. Experimental Setup

4.1. Environments

We evaluate the performance of our proposed approach using *Energym*, an open-source building simulation library for benchmarking smart-grid control algorithms (Scharnhorst et al., 2021). *Energym* provides a *Python* interface for ground-truth building simulations designed in *EnergyPlus* (Crawley et al., 2001), and offers buildings with varied equipment, geographies, and structural properties. We perform experiments in the following three buildings:

Mixed-Use facility in Athens, Greece. 566.38m² surface area; 13 thermal zones; $\mathcal{A} \in \mathbb{R}^{12}$ and $\mathcal{S} \in \mathbb{R}^{37}$. Temperature setpoints and air handling unit (AHU) flowrates are controllable.

Office block in Athens, Greece. 643.73m² surface area; 25 thermal zones; $\mathcal{A} \in \mathbb{R}^{14}$ and $\mathcal{S} \in \mathbb{R}^{56}$. Only temperature setpoints are controllable.

Seminar Centre in Billund, Denmark. 1278.94m² surface area; 27 thermal zones; $\mathcal{A} \in \mathbb{R}^{18}$ and $\mathcal{S} \in \mathbb{R}^{59}$. Only temperature setpoints are controllable.

In all cases, environment states are represented by a combination of temperature, humidity and pressure sensors (among others). Full state and action spaces for each building are reported in Appendix A for brevity. Experiments were run in each environment for one year starting on

01/01/2017, advancing in k -minute timestep increments, with k being environment dependent. Weather and grid carbon intensity match the true data in each geography for this period.

4.2. Baselines

We compare the performance of our agent against several strong RL baselines, and an RBC:

Soft Actor Critic (SAC; (Haarnoja et al., 2018)), a state-of-the-art model-free RL algorithm, known for lower variance performance than other popular model-free algorithms like PPO and DDPG.

Proximal Policy Optimisation (PPO; (Schulman et al., 2017)), a seminal model-free algorithm that is popular in production and used in previous works as a baseline by Ding et al. (2020) and Zhang et al. (2019a).

MPC with Deterministic Neural Networks (MPC-DNN; Nagabandi et al. (2018)), a simple, high-performing model-based architecture. Varying implementations have been used by previous authors, notably Ding et al. (2020) and Zhang et al. (2019a). We use Nagabandi et al. (2018)’s implementation which optimised MPC using random shooting.

Pre-trained SAC, an SAC agent with hyperparameters optimised for each environment and trained for 10 years prior to test time. We use this as our *oracle*, attempting to represent optimal performance by allowing access to the building simulator *a priori* for pre-training.

RBC, a generic controller that modulates temperature setpoints following the heuristics outlined in Appendix B.1.

We ensure both model-based agents plan with the same number of candidate actions over the same time horizon H to ensure performance differences are a consequence only of their differing design. For each agent we adopt the implementations from the original papers, except for the number of network layers and their dimensions which are set to 5 and 200 respectively to attempt to capture the full complexity of the system dynamics. Full hyperparameter specifications are provided in Appendix B.2–PEARL’s trajectory sampling and MPC hyperparameters follow their original implementations by Chua et al. (2018).

5. Results and Discussion

5.1. Performance

Table 1 reports cumulative emissions, temperature infraction percentages, planning speeds, and mean rewards for our

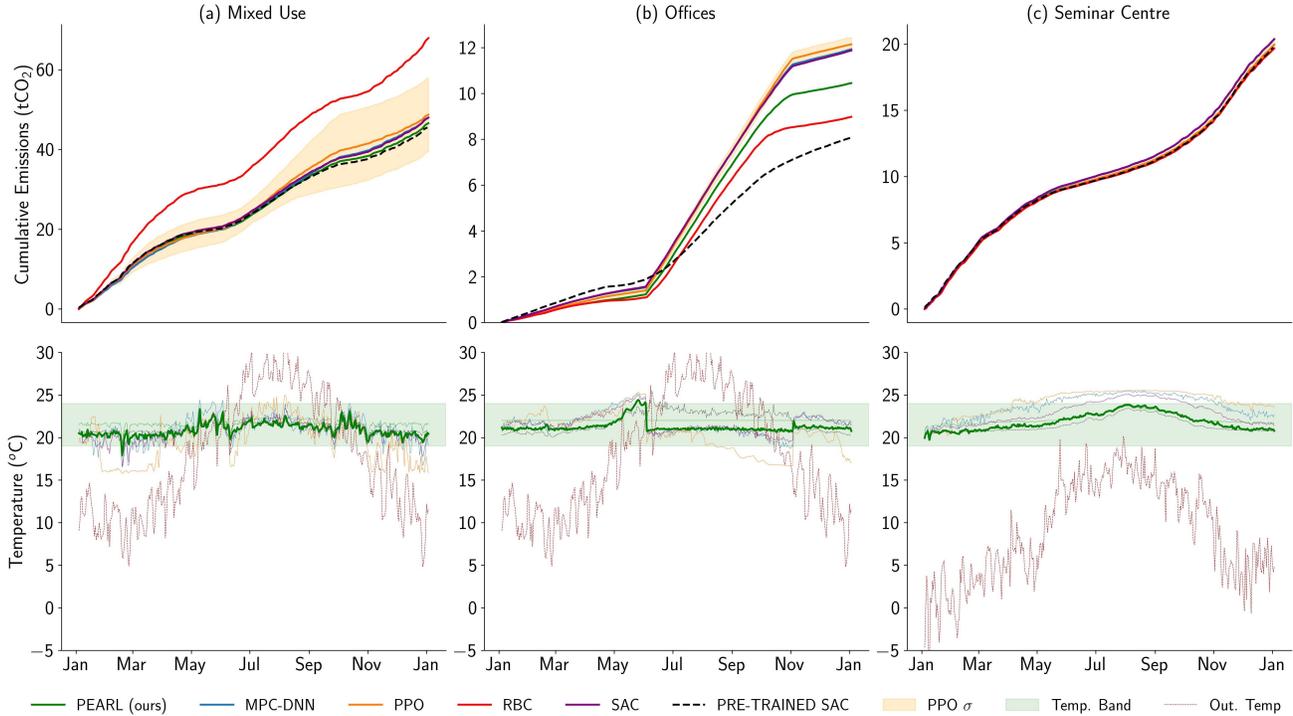


Figure 2. Top: Cumulative emissions produced by all agents across the (a) Mixed Use, (b) Offices, and (c) Seminar Centre environments. Curves represent the mean of 3 runs of the experiment, shaded areas are one standard deviation (too small to see in all cases except PPO). Bottom: Mean daily building temperature produced by all agents, the green shaded area illustrates the target temperature range [19, 24].

six controllers across the Mixed-Use, Office and Seminar Centre *Energym* environments. Results are reported as the mean \pm standard deviation for 3 runs of each experiment. The RBC performs identically across all experiments because both its policy and the environment are deterministic; varying RL agent performance is a consequence of policy and initialisation stochasticity. Comparisons with the Oracle are left to the end of the section.

Emissions. We find PEARL produces minimum emissions in the Mixed-Use environment, with cumulative emissions 31.46% lower than the RBC. In the Office block, the RBC exhibits lowest cumulative emissions, and PEARL outperforms all RL baselines. In the Seminar Centre, the RBC minimises emissions, and PPO marginally outperforms PEARL, but does so at the cost of erroneous temperature control, a theme discussed in the following section. Low external temperatures in the Seminar Centre make experiments there less informative as the optimal policy is to heat most of the year, providing little room for improved control. Stepwise emission totals for each environment are illustrated in the top half of Figure 2.

Temperature. We find that PEARL produces minimum daily mean temperature infractions in the Office and Sem-

inar centre environments, and is slightly outperformed by SAC in the Mixed-Use environment. The RBC is comparably performant across all environments, as would be expected. The RL baselines miss the thermal bounds regularly, with PPO and MPC-DNN exhibiting temperature infraction rates as high as 51.23% and 49.13% respectively. In some cases the strong emissions performance of these baselines is a direct consequence of shutting off HVAC equipment and forcing uncomfortable internal temperatures. This behaviour happens too frequently for any of the RL baselines to be considered safe. Mean daily building temperatures for each agent across the environments are plotted in the lower half of Figure 2.

Latency. We report action selection latency as a metric for comparing compute requirements for each agent, although absolute values here are hardware dependent. PPO is the most compute efficient, selecting actions in two thirds of the time required by MPC-DNN, and 41 times faster than PEARL on average; an expected result given PPO does not explicitly plan between actions taken in the environment. We note the latency of all agents is far smaller than the sampling period of each environment, meaning all implementations would prove adequate for real-world deployment. To minimise simulation time we limit both model-based

Table 1. Performance of all agents outlined in Section 4 across our three *Energym* environments. We define the temperature infraction metric as the percentage of days where mean building temperature falls outside the target range [19, 24], and latency as the mean compute time each agent requires to select an action given its policy. Results are averaged across 3 runs and presented as mean \pm standard deviation, except for pre-trained SAC which has converged on a policy prior to deployment with multiple runs showing the same performance.

	AGENT	EMISSIONS (TCO ₂)	TEMP. INFRACTIONS	LATENCY (S/ACTION)	MEAN REWARD
MIXED-USE	RBC	68.09 \pm 0.00	2.47% \pm 0%	-	-
	PPO	48.80 \pm 10.35	48.49% \pm 6.08%	0.021 \pm 0.01	$-1.47 \times 10^7 \pm 2.84 \times 10^6$
	SAC	48.80 \pm 0.14	0% \pm 0%	0.028 \pm 0.02	$-6.02 \times 10^5 \pm 3.85 \times 10^4$
	MPC-DNN	48.03 \pm 0.46	13.42% \pm 0.59%	0.030 \pm 0.01	$-1.12 \times 10^6 \pm 3.47 \times 10^4$
	PEARL	46.67 \pm 0.09	0.55% \pm 0.08%	0.870 \pm 0.15	$-5.76 \times 10^5 \pm 2.12 \times 10^3$
	PRE-TRAINED SAC	45.49	0%	0.027	-4.77×10^5
OFFICE	RBC	8.98 \pm 0.00	1.64% \pm 0%	-	-
	PPO	12.14 \pm 0.31	31.51% \pm 7.19%	0.018 \pm 0.006	$-2.26 \times 10^6 \pm 1.07 \times 10^6$
	SAC	11.87 \pm 0.01	6.58% \pm 1.12%	0.025 \pm 0.01	$-2.75 \times 10^5 \pm 2.03 \times 10^4$
	MPC-DNN	11.93 \pm 0.022	9.86% \pm 0.84%	0.029 \pm 0.001	$-5.50 \times 10^5 \pm 2.89 \times 10^4$
	PEARL	10.45 \pm 0.07	1.52% \pm 0%	0.845 \pm 0.14	$-5.51 \times 10^4 \pm 1.82 \times 10^3$
	PRE-TRAINED SAC	8.08	2.47%	0.023	-2.63×10^4
SEM. CENTRE	RBC	19.74 \pm 0.00	0% \pm 0%	-	-
	PPO	20.01 \pm 0.27	51.23% \pm 14.99%	0.028 \pm 0.01	$-4.15 \times 10^6 \pm 9.67 \times 10^5$
	SAC	20.37 \pm 0.08	29.32% \pm 1.87%	0.033 \pm 0.02	$-1.95 \times 10^6 \pm 7.11 \times 10^4$
	MPC-DNN	20.44 \pm 0.02	49.13% \pm 0.56%	0.035 \pm 0.001	$-2.45 \times 10^6 \pm 3.26 \times 10^4$
	PEARL	20.02 \pm 0.10	0% \pm 0%	0.911 \pm 0.17	$-1.18 \times 10^6 \pm 1.50 \times 10^3$
	PRE-TRAINED SAC	19.75	0%	0.031	-1.14×10^6

agents to 25 planned action sequences per timestep (Appendix B.2). Were they deployed in situ, we would utilise the time between environment interactions fully to plan with far greater numbers of action sequences—a similar approach was taken by Silver et al. (2016). We suspect broadening the search space used for planning would further improve performance.

Reward. Mean annual reward captures an agent’s ability to minimise emissions *and* maintain thermal comfort; the primary measure of agent performance. PEARL exhibits maximum mean reward across all environments—24.35x, 4.66x and 2.57x higher on average than PPO, MPC-DNN and SAC respectively—suggesting it strikes this balance better than the other controllers. The reward curves for each agent are reported in Appendix B.3 for brevity.

Comparison with oracle. The pre-trained oracle outperforms the baselines and PEARL in all cases as expected. However, its performance is surprisingly close to PEARL’s, showing only 2.5% and 1.3% lower emissions in the Sem-

inar Centre and Mixed-Use environments, and exhibiting similar thermal performance. From these results one could conclude that PEARL has produced a near-optimal policy, but one cannot be certain the oracle has reached optimality given the unusual, non-convergent reward curves this problem setting creates (Appendix B.3). Indeed, the shape of the reward curve associated with an optimal policy is unclear, unlike the canonical episodic RL tasks (cartpole, mountain car etc.) where optimal solutions can be quantified by a fixed episodic return. This problem has been hitherto unexplored by the community and requires future work. Here, we can conclude that PEARL performs similarly to a pre-trained agent but cannot make concrete assertions about performance with respect to optimality.

Why is Mixed-Use performance an outlier? An important observation from Table 1 is that PEARL only outperforms the RBC in the Mixed-Use facility. Why would this be the case? The triviality of the Seminar Centre results has been discussed above, but one may expect PEARL to perform as well in the Office environment as it does in the

Mixed-Use environment. Unlike the Office environment, in the Mixed-Use facility the agent has access to thermostat setpoints *and* continuous AHU flowrate control. This greatly increases action-space complexity and moves the control problem away from a setting where simple heuristics can be readily applied. This would suggest that RL building controllers should only be deployed when the action space is sufficiently complex, likely owing to some access to continuous control parameters. Many previous works corroborate the usefulness of RL in complex state-action spaces (Garcia & Fernández, 2012; Hasselt, 2012; Bertsekas, 2019).

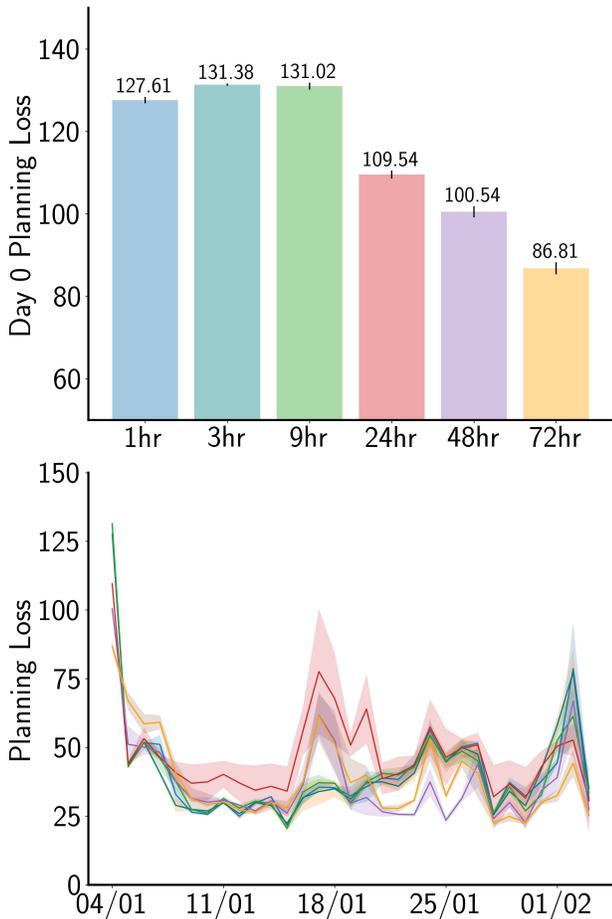


Figure 3. Planning loss on a holdout set of 100 randomly sampled state-action trajectories from another controller, given varying system identification duration, for three runs of the experiment. **Top:** mean planning loss at the end of the commissioning period, black lines represent one standard deviation. **Bottom:** mean planning loss throughout first month of operation, shaded areas represent one standard deviation. Experiments were performed in the Mixed-Use environment.

5.2. System Identification

We test the sensitivity of PEARL’s performance to system identification duration. We vary the commissioning period at six intervals between 1 and 72 hours, train the dynamics models on the collected dataset, and test their predictive accuracy on a holdout set of 100 randomly sampled state-action trajectories produced by another controller in the same environment. We measured the mean squared error between true and predicted state trajectories and sum across samples i.e. $\mathcal{L}(\hat{s}, s) = \frac{1}{N} \sum_{n=1}^N (\hat{s}_{t:t+H}^n - s_{t:t+H}^n)^2$, with $N = 100$. We report the results at the end of the commissioning period and throughout the first month of operation in Figure 3, we assume the effect of the commissioning period on model accuracy wanes beyond the first month when data collected through normal interaction far exceeds that from system identification. Note the model continues to be updated at the end of each day using data collected from the environment during that day.

We find that model accuracy at the end of the commissioning period generally correlates with length of system identification, as would be expected. However, as the agent continues beyond the commissioning period and collects more data we find that the accuracy of the models converge and by 21/01 there is little difference in cross-model accuracy. These results suggest the length of system identification period are a matter of practitioner preference. How important is immediate good performance to the practitioner and how much time are they willing to commit to commissioning? If one is time-poor and willing to sacrifice initial performance, they can utilise a short commissioning period in the knowledge that the model will be as accurate as any other after a few days.

5.3. Agent Decomposition

What components of PEARL allow it to perform safe control? We note two differences between PEARL and our model-based RL baseline MPC-DNN: 1) The use of probabilistic networks, rather than deterministic networks, and 2) The use of CEM optimisation, rather than random shooting (RS). Given control via MPC-DNN is unsafe and inefficient, we assume the improved performance of PEARL can be attributed to these design choices. To understand their relative importance, we vary the design of PEARL to either include or exclude the components and compare performance—Figure 4.

We find the performance of PEARL to be more sensitive to the choice of optimiser than the choice of network. The average performance of agents with CEM optimiser is 24.2% higher than those using RS and only 14.6% lower than pre-trained agent performance. Surprisingly, the performance of PEARL appears invariant to the choice of network, with agents composed of deterministic and probabilistic networks

performing similarly when optimiser choice is controlled for. This is in contradiction to many works suggesting probabilistic modelling of dynamics function improves performance over deterministic models, particularly in complex, partially-observed state-action spaces like those exhibited in this study (Deisenroth & Rasmussen, 2011; Kamthe & Deisenroth, 2018; Chua et al., 2018; Levine, 2018). Given the wealth of literature corroborating their usefulness, and their low cost of inclusion, we continue to endorse probabilistic dynamics functions despite this result, as they may improve performance in settings beyond those used in this study. Future work could look to specify settings in which better performance can be expected from probabilistic models.

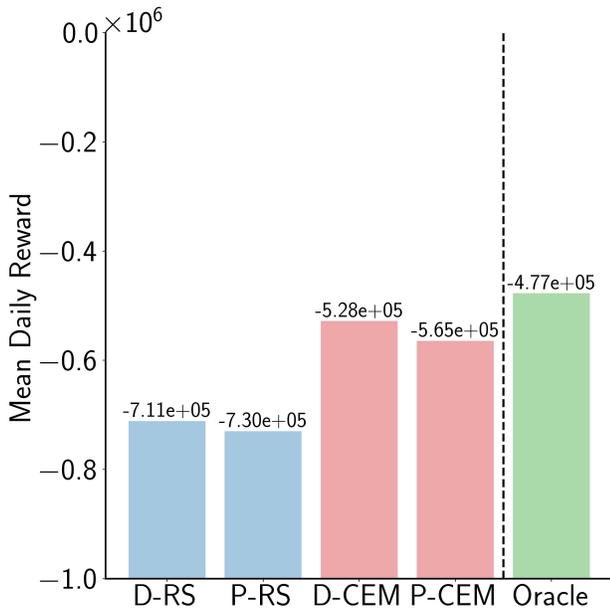


Figure 4. Mean daily reward for four instantiations of PEARL varying the choice of network between Deterministic (D) and Probabilistic (P), and the choice of optimiser between Random Shooting (RS) and Cross-Entropy Maximisation (CEM). The oracle (pre-trained SAC) is plotted as a baseline. Experiments were performed in the Mixed-Use environment for one year.

6. Limitations and Future Work

In this section we contextualise our findings by highlighting limitations in the experimental setup, and propose avenues for future work to improve upon these results.

6.1. Limitations

Sensor Coverage. The *Energym* simulations used in this work contain large state-spaces with comprehensive sensor coverage (Appendix A). Such sensor deployment in the real

world is unrealistic outside of advanced facilities and it is unclear how agent performance would be affected by fewer sensors.

State Sampling. Our MDP, as formulated in (Section 3) assumes access to a temporally consistent stream of state samples. However, in practice it is common for sensors to fail, or go periods without logging data, violating this property of the MDP. As it stands, it is unclear how PEARL would react to periods where the state-space is unobserved, and methods for dealing with such times safely are as yet unexplored.

Action-space Complexity. In Section 5.1 we hypothesise that PEARL’s improved performance over the RBC in the Mixed-use facility is a consequence of higher action-space complexity. In practice, it may be difficult to find the appropriate level of complexity; for example, we are unlikely to be able to access continuous control of AHU flowrate in every building, and, in pursuit of more complex action-spaces, we are likely to hit a maxima after which agent performance suffers rather than benefits.

Environments. We were only able to perform experiments in three environments, one of which provided results that were less informative than the other two. As has been the case in previous studies, research in the smart building space is hamstrung by limited access to building energy simulations, despite attempts to unify the space (Findeis et al., 2022). The results reported in this work would be improved by conducting experiments in a wider variety of buildings, were they easily accessible or producible.

6.2. Future Work

State Non-stationarity. In this work, and adjacent areas like Continual Learning (Kirkpatrick et al., 2017), the agent must adapt well to a non-stationary environment. In our case, such non-stationarity is usually a consequence of changing seasons or occupant preferences. This work does not deal with the problem of non-stationarity rigorously, instead assuming our simple solution of updating network parameters at the end of each day captures the changing environment. Indeed, previous works choose to update parameters on different timescales (days vs. months) and using different techniques (daily data vs. sampled historical data). Exploring the most appropriate methods for dealing with non-stationarity in the building control space could improve agent performance during out-of-distribution events.

Inconsistent State Sampling. As discussed in Section 6.1, temporally inconsistent state sampling is a significant problem in practice, but has been ignored thus far in the literature. Work that looks to handle periods of sensor downtime safely and efficiently would prove useful contributions.

Scaled Experiments. Much work in the theoretical RL lit-

erature relies on performing experiments in many varied environments – up to 15 in the *DeepMind Control Suite* (Tassa et al., 2018) and 50 in *Meta-World v2* (Yu et al., 2019). To make progress in the RL for building control space, we need a unified problem setting with many environments. Doing so would allow researchers to make algorithmic innovations more reliably, and would convince practitioners of the safety and efficiency of the techniques.

7. Conclusion

In this work we consider the task of learning policies *tabula rasa* that minimise emissions in buildings whilst ensuring thermal comfort, a considerably harder task than pre-training models in simulation before deployment. By combining existing works in a new approach called PEARL (Probabilistic Emission-Abating Reinforcement Learning) we have shown it is possible to safely and efficiently control buildings *online* without pre-training in simulation. We showed our approach can reduce emissions by 31.46% compared with an existing RBC after a commissioning period of only three hours, but posit that this performance is correlated strongly with action-space complexity which cannot be assumed in practice. When compared with existing RL baselines, our algorithm performs favourably, showing reduced emissions in all cases bar one, whilst maintaining thermal comfort more effectively. We suggest that the positive performance of PEARL is linked closely with the choice of optimiser and, to our surprise, that modelling the uncertain system dynamics probabilistically has little effect. We close with limitations related to simulation sensor coverage and sampling rates, and the number of building environments we could access to test agent generalisability. This work adds to the growing literature on RL for energy-system control, highlighting its usefulness as a climate change mitigation tool.

Acknowledgements. We thank Arduin Findeis, Nantas Nardelli and Srinivasan Keshav for helpful feedback throughout. This work was supported by an EPSRC DTP Studentship (reference EP/T517847/1) and Emerson Electric.

References

Bertsekas, D. *Reinforcement learning and optimal control*. Athena Scientific, 2019.

Botev, Z. I., Kroese, D. P., Rubinstein, R. Y., and L’Ecuyer, P. The cross-entropy method for optimization. In *Handbook of statistics*, volume 31, pp. 35–59. Elsevier, 2013.

Camacho, E. F. and Alba, C. B. *Model predictive control*. Springer Science & Business Media, 2013.

Chen, Y., Norford, L. K., Samuelson, H. W., and Malkawi,

A. Optimal control of hvac and window systems for natural ventilation through reinforcement learning. *Energy and Buildings*, 169:195 – 205, 2018.

Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *arXiv preprint arXiv:1805.12114*, 2018.

Crawley, D. B., Lawrie, L. K., Winkelmann, F. C., Buhl, W. F., Huang, Y. J., Pedersen, C. O., Strand, R. K., Liesen, R. J., Fisher, D. E., Witte, M. J., et al. Energyplus: creating a new-generation building energy simulation program. *Energy and buildings*, 33(4):319–331, 2001.

Cullen, J. and Allwood, J. The efficient use of energy: Tracing the global flow of energy from fuel to service. *Energy Policy*, 38(1):75–81, 2010.

Deisenroth, M. and Rasmussen, C. E. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472, 2011.

Ding, X., Du, W., and Cerpa, A. E. Mb2c: Model-based deep reinforcement learning for multi-zone building control. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 50–59, 2020.

Findeis, A., Kazhamiaka, F., Jeon, S., and Keshav, S. Beobench: A toolkit for unified access to building simulations for reinforcement learning. In *e-Energy ’22*, pp. 374–382, 2022. URL <https://doi.org/10.1145/3538637.3538866>.

Gal, Y., McAllister, R., and Rasmussen, C. E. Improving pilco with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, volume 4, pp. 25, 2016.

Garcia, J. and Fernández, F. Safe exploration of state and action spaces in reinforcement learning. *Journal of Artificial Intelligence Research*, 45:515–564, 2012.

Gu, S., Holly, E., Lillicrap, T., and Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 3389–3396. IEEE, 2017.

Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, 2018. URL <https://arxiv.org/abs/1801.01290>.

Hasselt, H. V. Reinforcement learning in continuous state and action spaces. In *Reinforcement learning*, pp. 207–251. Springer, 2012.

- Hensman, J., Fusi, N., and Lawrence, N. D. Gaussian processes for big data. *arXiv preprint arXiv:1309.6835*, 2013.
- Higuera, J. C. G., Meger, D., and Dudek, G. Synthesizing neural network controllers with probabilistic model based reinforcement learning. *CoRR*, abs/1803.02291, 2018. URL <http://arxiv.org/abs/1803.02291>.
- Jain, A., Nghiem, T., Morari, M., and Mangharam, R. Learning and control using gaussian processes. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCPS)*, pp. 140–149, 2018.
- Jain, A., Smarra, F., Reticcioli, E., D’Innocenzo, A., and Morari, M. Neuropt: Neural network based optimization for building energy management and climate control. In *Proceedings of the 2nd Conference on Learning for Dynamics and Control*, volume 120 of *Proceedings of Machine Learning Research*, pp. 445–454, The Cloud, 6 2020. PMLR.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1-2):99–134, 1998.
- Kamthe, S. and Deisenroth, M. Data-efficient reinforcement learning with probabilistic model predictive control. In *International conference on artificial intelligence and statistics*, pp. 1701–1710. PMLR, 2018.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, mar 2017.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30, 2017.
- Lazic, N., Lu, T., Boutilier, C., Ryu, M., Wong, E. J., Roy, B., and Imwalle, G. Data center cooling using model-predictive control. In *Proceedings of the Thirty-second Conference on Neural Information Processing Systems (NeurIPS-18)*, pp. 3818–3827, Montreal, QC, 2018.
- Levine, S. Reinforcement learning and control as probabilistic inference: Tutorial and review, 2018. URL <https://arxiv.org/abs/1805.00909>.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning, 2015.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Nagabandi, A., Kahn, G., Fearing, R. S., and Levine, S. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7559–7566. IEEE, 2018.
- Scharnhorst, P., Schubnel, B., Fernández Bandera, C., Salom, J., Taddeo, P., Boegli, M., Gorecki, T., Stauffer, Y., Peppas, A., and Politi, C. Energym: A building model library for controller benchmarking. *Applied Sciences*, 11 (8):3518, 2021.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- 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*, 550(7676):354–359, 2017.
- Sutton, R. and Barto, A. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., Casas, D. d. L., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., Lillicrap, T., and Riedmiller, M. Deepmind control suite, 2018.
- Valladares, W., Galindo, M., Gutiérrez, J., Wu, W.-C., Liao, K.-K., Liao, J.-C., Lu, K.-C., and Wang, C.-C. Energy optimization associated with thermal comfort and indoor air control via a deep reinforcement learning algorithm. *Building and Environment*, 155:105 – 117, 2019. ISSN 0360-1323.
- Wei, T., Wang, Y., and Zhu, Q. Deep reinforcement learning for building hvac control. In *Proceedings of the 54th Annual Design Automation Conference 2017, DAC ’17*, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349277.
- Williams, C. K. and Rasmussen, C. E. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.

- Yu, L., Qin, S., Zhang, M., Shen, C., Jiang, T., and Guan, X. A review of deep reinforcement learning for smart building energy management. *IEEE Internet of Things Journal*, 8(15):12046–12063, Aug 2021. ISSN 2372-2541.
- Yu, T., Quillen, D., He, Z., Julian, R., Narayan, A., Shively, H., Bellathur, A., Hausman, K., Finn, C., and Levine, S. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning, 2019.
- Zhang, C., Kuppannagari, S. R., Kannan, R., and Prasanna, V. K. Building hvac scheduling using reinforcement learning via neural network based model approximation. In *Proceedings of the 6th ACM international conference on systems for energy-efficient buildings, cities, and transportation*, pp. 287–296, 2019a.
- Zhang, Z., Chong, A., Pan, Y., Zhang, C., and Lam, K. P. Whole building energy model for hvac optimal control: A practical framework based on deep reinforcement learning. *Energy and Buildings*, 199:472–490, 2019b.

A. Energym Environments

Table 2: State-space \mathcal{S}_{mixed} of the *Mixed-use Energym* environment.

Variable	Type	Lower Bound	Upper Bound	Description
Bd_T_AHU1	scalar	10	30	AHU 1 temperature (°C).
Bd_T_AHU2	scalar	10	30	AHU 2 temperature (°C).
Ext_Irr	scalar	0	1000	Direct normal radiation (W/m ²).
Ext_P	scalar	80000.0	130000.0	Outdoor air pressure (Pa).
Ext_RH	scalar	0	100	Outdoor relative humidity (%RH).
Ext_T	scalar	-10	40	Outdoor temperature (°C).
Fa_Pw_All	scalar	0	50000.0	Total power consumption (W).
Z02_Fl_Fan	scalar	0	1	Zone 2 fan flow rate.
Z02_RH	scalar	0	100	Zone 2 relative humidity (%RH).
Z02_T	scalar	10	40	Zone 2 temperature (°C).
Z02_T_Thermostat_sp_out	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_Fl_Fan	scalar	0	1	Zone 3 fan flow rate.
Z03_RH	scalar	0	100	Zone 3 relative humidity (%RH).
Z03_T	scalar	10	40	Zone 3 temperature (°C).
Z03_T_Thermostat_sp_out	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_Fl_Fan	scalar	0	1	Zone 4 fan flow rate.
Z04_RH	scalar	0	100	Zone 4 relative humidity (%RH).
Z04_T	scalar	10	40	Zone 4 temperature (°C).
Z04_T_Thermostat_sp_out	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_RH	scalar	0	100	Zone 5 relative humidity (%RH).
Z05_T	scalar	10	40	Zone 5 temperature (°C).
Z05_T_Thermostat_sp_out	scalar	16	26	Zone 5 thermostat setpoint (°C).
Z08_Fl_Fan	scalar	0	1	Zone 8 fan flow rate.
Z08_RH	scalar	0	100	Zone 8 relative humidity (%RH).
Z08_T	scalar	10	40	Zone 8 temperature (°C).
Z08_T_Thermostat_sp_out	scalar	16	26	Zone 8 thermostat setpoint (°C).
Z09_Fl_Fan	scalar	0	1	Zone 9 fan flow rate.
Z09_RH	scalar	0	100	Zone 9 relative humidity (%RH).
Z09_T	scalar	10	40	Zone 9 temperature (°C).
Z09_T_Thermostat_sp_out	scalar	16	26	Zone 9 thermostat setpoint (°C).
Z10_Fl_Fan	scalar	0	1	Zone 10 fan flow rate.
Z10_RH	scalar	0	100	Zone 10 relative humidity (%RH).
Z10_T	scalar	10	40	Zone 10 temperature (°C).
Z10_T_Thermostat_sp_out	scalar	16	26	Zone 10 thermostat setpoint (°C).
Z11_Fl_Fan	scalar	0	1	Zone 11 fan flow rate.
Z11_RH	scalar	0	100	Zone 11 relative humidity (%RH).
Z11_T	scalar	10	40	Zone 11 temperature (°C).
Z11_T_Thermostat_sp_out	scalar	16	26	Zone 11 thermostat setpoint (°C).

Table 3: Action-space \mathcal{A}_{mixed} of the *Mixed-use Energym* environment.

Action	Type	Lower Bound	Upper Bound	Description
Z02_T_Thermostat_sp	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_T_Thermostat_sp	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_T_Thermostat_sp	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_T_Thermostat_sp	scalar	16	26	Zone 5 thermostat setpoint (°C).

continued on next page

Table 3 – continued from previous page

Action	Type	Lower Bound	Upper Bound	Description
Z08_T_Thermostat_sp	scalar	16	26	Zone 8 thermostat setpoint (°C).
Z09_T_Thermostat_sp	scalar	16	26	Zone 9 thermostat setpoint (°C).
Z10_T_Thermostat_sp	scalar	16	26	Zone 10 thermostat setpoint (°C).
Z11_T_Thermostat_sp	scalar	16	26	Zone 11 thermostat setpoint (°C).
Bd_T_AHU2_sp	scalar	10	30	AHU 2 temperature setpoint (°C).
Bd_FL_AHU2_sp	scalar	0	1	AHU 2 flow rate setpoint.
Bd_T_AHU1_sp	scalar	10	30	AHU 1 temperature setpoint (°C).
Bd_FL_AHU1_sp	scalar	0	1	AHU1 flow rate setpoint.

Table 4: State-space S_{office} of the *Offices Energym* environment.

Variable	Type	Lower Bound	Upper Bound	Description
Bd_Pw_All	scalar	0	5000	Building power consumption (W).
Ext_Irr	scalar	0	1000	Direct normal radiation (W/m2).
Ext_RH	scalar	0	100	Outdoor relative humidity (%RH).
Ext_T	scalar	-10	40	Outdoor temperature (°C).
Fa_Pw_All	scalar	0	10000.0	Total power consumption (W).
Fa_Pw_HVAC	scalar	0	10000.0	HVAC power consumption (W).
Fa_Pw_PV	scalar	0	2000.0	PV power production (W).
Z01_FL_Fan_sp_out	scalar	0	1	Zone 1 fan flow setpoint.
Z01_T	scalar	10	40	Zone 1 temperature (°C).
Z01_T_Thermostat_sp_out	scalar	16	26	Zone 1 thermostat setpoint (°C).
Z02_FL_Fan_sp_out	scalar	0	1	Zone 2 fan flow setpoint.
Z02_T	scalar	10	40	Zone 2 temperature (°C).
Z02_T_Thermostat_sp_out	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_FL_Fan_sp_out	scalar	0	1	Zone 3 fan flow setpoint.
Z03_FL_Fan1_sp_out	scalar	0	1	Zone 3 fan 1 flow setpoint.
Z03_T	scalar	10	40	Zone 3 temperature (°C).
Z03_T_Thermostat_sp_out	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_FL_Fan_sp_out	scalar	0	1	Zone 4 fan flow setpoint.
Z04_T	scalar	10	40	Zone 4 temperature (°C).
Z04_T_Thermostat_sp_out	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_FL_Fan_sp_out	scalar	0	1	Zone 5 fan flow setpoint.
Z05_T	scalar	10	40	Zone 5 temperature (°C).
Z05_T_Thermostat_sp_out	scalar	16	26	Zone 5 thermostat setpoint (°C).
Z06_FL_Fan_sp_out	scalar	0	1	Zone 6 fan flow setpoint.
Z06_T	scalar	10	40	Zone 6 temperature (°C).
Z06_T_Thermostat_sp_out	scalar	16	26	Zone 6 thermostat setpoint (°C).
Z07_FL_Fan_sp_out	scalar	0	1	Zone 7 fan flow setpoint.
Z07_T	scalar	10	40	Zone 7 temperature (°C).
Z07_T_Thermostat_sp_out	scalar	16	26	Zone 7 thermostat setpoint (°C).
Z15_FL_Fan_sp_out	scalar	0	1	Zone 15 fan flow setpoint.
Z15_T	scalar	10	40	Zone 15 temperature (°C).
Z15_T_Thermostat_sp_out	scalar	16	26	Zone 15 thermostat setpoint (°C).
Z16_FL_Fan_sp_out	scalar	0	1	Zone 16 fan flow setpoint.
Z16_T	scalar	10	40	Zone 16 temperature (°C).
Z16_T_Thermostat_sp_out	scalar	16	26	Zone 16 thermostat setpoint (°C).
Z17_FL_Fan_sp_out	scalar	0	1	Zone 17 fan flow setpoint.
Z17_FL_Fan1_sp_out	scalar	0	1	Zone 17 fan 1 flow setpoint.
Z17_T	scalar	10	40	Zone 17 temperature (°C).

continued on next page

Table 4 – continued from previous page

Variable	Type	Lower Bound	Upper Bound	Description
Z17_T_Thermostat_sp_out	scalar	16	26	Zone 17 thermostat setpoint (°C).
Z18_Fl_Fan_sp_out	scalar	0	1	Zone 18 fan flow setpoint.
Z18_T	scalar	10	40	Zone 18 temperature (°C).
Z18_T_Thermostat_sp_out	scalar	16	26	Zone 18 thermostat setpoint (°C).
Z19_Fl_Fan_sp_out	scalar	0	1	Zone 19 fan flow setpoint.
Z19_Fl_Fan1_sp_out	scalar	0	1	Zone 19 fan 1 flow setpoint.
Z19_T	scalar	10	40	Zone 19 temperature (°C).
Z19_T_Thermostat_sp_out	scalar	16	26	Zone 19 thermostat setpoint (°C).
Z20_Fl_Fan_sp_out	scalar	0	1	Zone 20 fan flow setpoint.
Z20_Fl_Fan1_sp_out	scalar	0	1	Zone 20 fan 1 flow setpoint.
Z20_T	scalar	10	40	Zone 20 temperature (°C).
Z20_T_Thermostat_sp_out	scalar	16	26	Zone 20 thermostat setpoint (°C).
Z25_Fl_Fan_sp_out	scalar	0	1	Zone 25 fan flow setpoint.
Z25_Fl_Fan1_sp_out	scalar	0	1	Zone 25 fan 1 flow setpoint.
Z25_T	scalar	10	40	Zone 25 temperature (°C).
Z25_T_Thermostat_sp_out	scalar	16	26	Zone 25 thermostat setpoint (°C).

Table 5: Action-space $\mathcal{A}_{offices}$ of the *Offices Energym* environment.

Action	Type	Lower Bound	Upper Bound	Description
Z01_T_Thermostat_sp	scalar	16	26	Zone 1 thermostat setpoint (°C).
Z02_T_Thermostat_sp	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_T_Thermostat_sp	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_T_Thermostat_sp	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_T_Thermostat_sp	scalar	16	26	Zone 5 thermostat setpoint (°C).
Z06_T_Thermostat_sp	scalar	16	26	Zone 6 thermostat setpoint (°C).
Z07_T_Thermostat_sp	scalar	16	26	Zone 7 thermostat setpoint (°C).
Z15_T_Thermostat_sp	scalar	16	26	Zone 15 thermostat setpoint (°C).
Z16_T_Thermostat_sp	scalar	16	26	Zone 16 thermostat setpoint (°C).
Z17_T_Thermostat_sp	scalar	16	26	Zone 17 thermostat setpoint (°C).
Z18_T_Thermostat_sp	scalar	16	26	Zone 18 thermostat setpoint (°C).
Z19_T_Thermostat_sp	scalar	16	26	Zone 19 thermostat setpoint (°C).
Z20_T_Thermostat_sp	scalar	16	26	Zone 20 thermostat setpoint (°C).
Z25_T_Thermostat_sp	scalar	16	26	Zone 25 thermostat setpoint (°C).

Table 6: State-space $\mathcal{S}_{seminar}$ of the *Seminar-Thermostat Energym* environment.

Variable	Type	Lower Bound	Upper Bound	Description
Ext_Irr	scalar	0	1000	Direct normal radiation (W/m2).
Ext_RH	scalar	0	100	Outdoor relative humidity (%RH).
Ext_T	scalar	-15	40	Outdoor temperature (°C).
Ext_P	scalar	80000	130000	Outdoor air pressure (Pa).
Z01_RH	scalar	0	100	Zone 1 relative humidity (%RH).
Z01_T	scalar	10	40	Zone 1 temperature (°C).
Z01_T_Thermostat_sp_out	scalar	16	26	Zone 1 thermostat setpoint (°C).
Z02_RH	scalar	0	100	Zone 2 relative humidity (%RH).
Z02_T	scalar	10	40	Zone 2 temperature (°C).

continued on next page

Table 6 – continued from previous page

Variable	Type	Lower Bound	Upper Bound	Description
Z02_T_Thermostat_sp_out	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_RH	scalar	0	100	Zone 3 relative humidity (%RH).
Z03_T	scalar	10	40	Zone 3 temperature (°C).
Z03_T_Thermostat_sp_out	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_RH	scalar	0	100	Zone 4 relative humidity (%RH).
Z04_T	scalar	10	40	Zone 4 temperature (°C).
Z04_T_Thermostat_sp_out	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_RH	scalar	0	100	Zone 5 relative humidity (%RH).
Z05_T	scalar	10	40	Zone 5 temperature (°C).
Z05_T_Thermostat_sp_out	scalar	16	26	Zone 5 thermostat setpoint (°C).
Z06_RH	scalar	0	100	Zone 6 relative humidity (%RH).
Z06_T	scalar	10	40	Zone 6 temperature (°C).
Z06_T_Thermostat_sp_out	scalar	16	26	Zone 6 thermostat setpoint (°C).
Z08_RH	scalar	0	100	Zone 8 relative humidity (%RH).
Z08_T	scalar	10	40	Zone 8 temperature (°C).
Z08_T_Thermostat_sp_out	scalar	16	26	Zone 8 thermostat setpoint (°C).
Z09_RH	scalar	0	100	Zone 9 relative humidity (%RH).
Z09_T	scalar	10	40	Zone 9 temperature (°C).
Z09_T_Thermostat_sp_out	scalar	16	26	Zone 9 thermostat setpoint (°C).
Z10_RH	scalar	0	100	Zone 10 relative humidity (%RH).
Z10_T	scalar	10	40	Zone 10 temperature (°C).
Z10_T_Thermostat_sp_out	scalar	16	26	Zone 10 thermostat setpoint (°C).
Z11_RH	scalar	0	100	Zone 11 relative humidity (%RH).
Z11_T	scalar	10	40	Zone 11 temperature (°C).
Z11_T_Thermostat_sp_out	scalar	16	26	Zone 11 thermostat setpoint (°C).
Z13_RH	scalar	0	100	Zone 12 relative humidity (%RH).
Z13_T	scalar	10	40	Zone 13 temperature (°C).
Z13_T_Thermostat_sp_out	scalar	16	26	Zone 13 thermostat setpoint (°C).
Z14_RH	scalar	0	100	Zone 14 relative humidity (%RH).
Z14_T	scalar	10	40	Zone 14 temperature (°C).
Z14_T_Thermostat_sp_out	scalar	16	26	Zone 14 thermostat setpoint (°C).
Z15_RH	scalar	0	100	Zone 15 relative humidity (%RH).
Z15_T	scalar	10	40	Zone 15 temperature (°C).
Z15_T_Thermostat_sp_out	scalar	16	26	Zone 15 thermostat setpoint (°C).
Z18_RH	scalar	0	100	Zone 18 relative humidity (%RH).
Z18_T	scalar	10	40	Zone 18 temperature (°C).
Z18_T_Thermostat_sp_out	scalar	16	26	Zone 18 thermostat setpoint (°C).
Z19_RH	scalar	0	100	Zone 19 relative humidity (%RH).
Z19_T	scalar	10	40	Zone 19 temperature (°C).
Z19_T_Thermostat_sp_out	scalar	16	26	Zone 19 thermostat setpoint (°C).
Z20_RH	scalar	0	100	Zone 20 relative humidity (%RH).
Z20_T	scalar	10	40	Zone 20 temperature (°C).
Z20_T_Thermostat_sp_out	scalar	16	26	Zone 20 thermostat setpoint (°C).
Z21_RH	scalar	0	100	Zone 21 relative humidity (%RH).
Z21_T	scalar	10	40	Zone 21 temperature (°C).
Z21_T_Thermostat_sp_out	scalar	16	26	Zone 21 thermostat setpoint (°C).
Z22_RH	scalar	0	100	Zone 22 relative humidity (%RH).
Z22_T	scalar	10	40	Zone 22 temperature (°C).
Z22_T_Thermostat_sp_out	scalar	16	26	Zone 22 thermostat setpoint (°C).
Grid_CO2	scalar	0	10	Grid CO2 emission intensity (g/kWh).

Table 7: Action-space $\mathcal{A}_{seminar}$ of the *Seminar-Thermostat Energym* environment.

Action	Type	Lower Bound	Upper Bound	Description
Z01_T_Thermostat_sp	scalar	16	26	Zone 1 thermostat setpoint (°C).
Z02_T_Thermostat_sp	scalar	16	26	Zone 2 thermostat setpoint (°C).
Z03_T_Thermostat_sp	scalar	16	26	Zone 3 thermostat setpoint (°C).
Z04_T_Thermostat_sp	scalar	16	26	Zone 4 thermostat setpoint (°C).
Z05_T_Thermostat_sp	scalar	16	26	Zone 5 thermostat setpoint (°C).
Z06_T_Thermostat_sp	scalar	16	26	Zone 6 thermostat setpoint (°C).
Z08_T_Thermostat_sp	scalar	16	26	Zone 8 thermostat setpoint (°C).
Z09_T_Thermostat_sp	scalar	16	26	Zone 9 thermostat setpoint (°C).
Z10_T_Thermostat_sp	scalar	16	26	Zone 10 thermostat setpoint (°C).
Z11_T_Thermostat_sp	scalar	16	26	Zone 11 thermostat setpoint (°C).
Z13_T_Thermostat_sp	scalar	16	26	Zone 13 thermostat setpoint (°C).
Z14_T_Thermostat_sp	scalar	16	26	Zone 14 thermostat setpoint (°C).
Z15_T_Thermostat_sp	scalar	16	26	Zone 15 thermostat setpoint (°C).
Z18_T_Thermostat_sp	scalar	16	26	Zone 18 thermostat setpoint (°C).
Z19_T_Thermostat_sp	scalar	16	26	Zone 19 thermostat setpoint (°C).
Z20_T_Thermostat_sp	scalar	16	26	Zone 20 thermostat setpoint (°C).
Z21_T_Thermostat_sp	scalar	16	26	Zone 21 thermostat setpoint (°C).
Z22_T_Thermostat_sp	scalar	16	26	Zone 22 thermostat setpoint (°C).

B. Agents

B.1. Rule Based Controller (RBC)

The RBC’s objective is to maintain temperature at 22°C in the building. To achieve this it selects actions as follows:

$$a_t^i[t + 1] = \begin{cases} a_t^i[t] + 0.5^\circ C : & T_{obs}^i \leq 21.2^\circ C \\ a_t^i[t] - 0.5^\circ C : & T_{obs}^i \geq 22.8^\circ C \\ a_t^i[t] : & 21.2^\circ C \leq T_{obs}^i \leq 22.8^\circ C \end{cases} \quad (6)$$

B.2. Hyperparameters

Table 8: PEARL and MPC-DNN hyperparameters.

Hyperparameter	PEARL	MPC-DNN
Horizon (H)	20	20
Adam stepsize (α)	0.0003	0.0003
Num. epochs	25	25
Minibatch size	32	32
No. of candidate actions	25	25
No. of particles	10	n/a
Exploration parameter (Δ)	0.05	0.05
Network Fully Connected Layers	5	5
Network Layer Dimensions	200	200
Network Activation Function	Tanh	Tanh
Ensemble Models (K)	5	n/a
System ID Minutes	180	n/a
Reward Emissions Parameter (ϕ)	1	1
Reward Temperature Parameter (θ)	1000	1000

Table 9: PPO hyperparameters.

Hyperparameter	PPO
Batch Size	32
Adam stepsize (α)	0.0003
Epochs per update	50
Discount (γ)	0.99
GAE parameter (λ)	0.95
Number of actors	1
Clipping parameter	0.2
Actor Network Fully Connected Layers	5
Actor Network Layer Dimensions	200
Actor Network Activation Function	Tanh
Critic Network Fully Connected Layers	5
Critic Network Layer Dimensions	200
Critic Network Activation Function	Tanh
Reward Emissions Parameter (ϕ)	1
Reward Temperature Parameter (θ)	1000

Table 10: SAC hyperparameters.

Hyperparameter	SAC
Batch Size	256
Adam stepsize (α)	0.0003
Epochs per update	20
Discount (γ)	0.99
Target Smoothing Coefficient (τ)	0.005
Hidden Layers (all networks)	2
Layer Dimensions	256
Activation Function	ReLU
Reward Scale	2
Reward Emissions Parameter (ϕ)	1
Reward Temperature Parameter (θ)	1000

Table 11: SAC Oracle hyperparameters.

Hyperparameter	Mixed-use	Offices	Seminar Centre
Batch Size	256	256	1024
Adam stepsize (α)	0.0003	0.0015	0.0080
Epochs per update	1	1	1
Discount (γ)	0.9864	0.7539	0.8533
Target Smoothing Coefficient (τ)	0.07815	0.2532	0.2890
Hidden Layers (all networks)	2	2	2
Layer Dimensions	256	256	256
Activation Function	ReLU	ReLU	ReLU
Reward Scale	2	20	20
Reward Emissions Parameter (ϕ)	1	1	1
Reward Temperature Parameter (θ)	1000	1000	1000

B.3. Reward Curves

Reward curves for each agent, in each environment, for 5 runs are illustrated in Figure 5.

C. Hardware

All experiments were completed locally on a machine with the following specifications:

- OS: Windows 10
- RAM: 32GB 4X8GB 2666MHz DDR4
- CPU: Intel Core™ i7-8700 (6 Cores/12MB/12T/up to 4.6GHz/65W)
- GPU: 1X NVIDIA GeForce RTX 3050 Ti

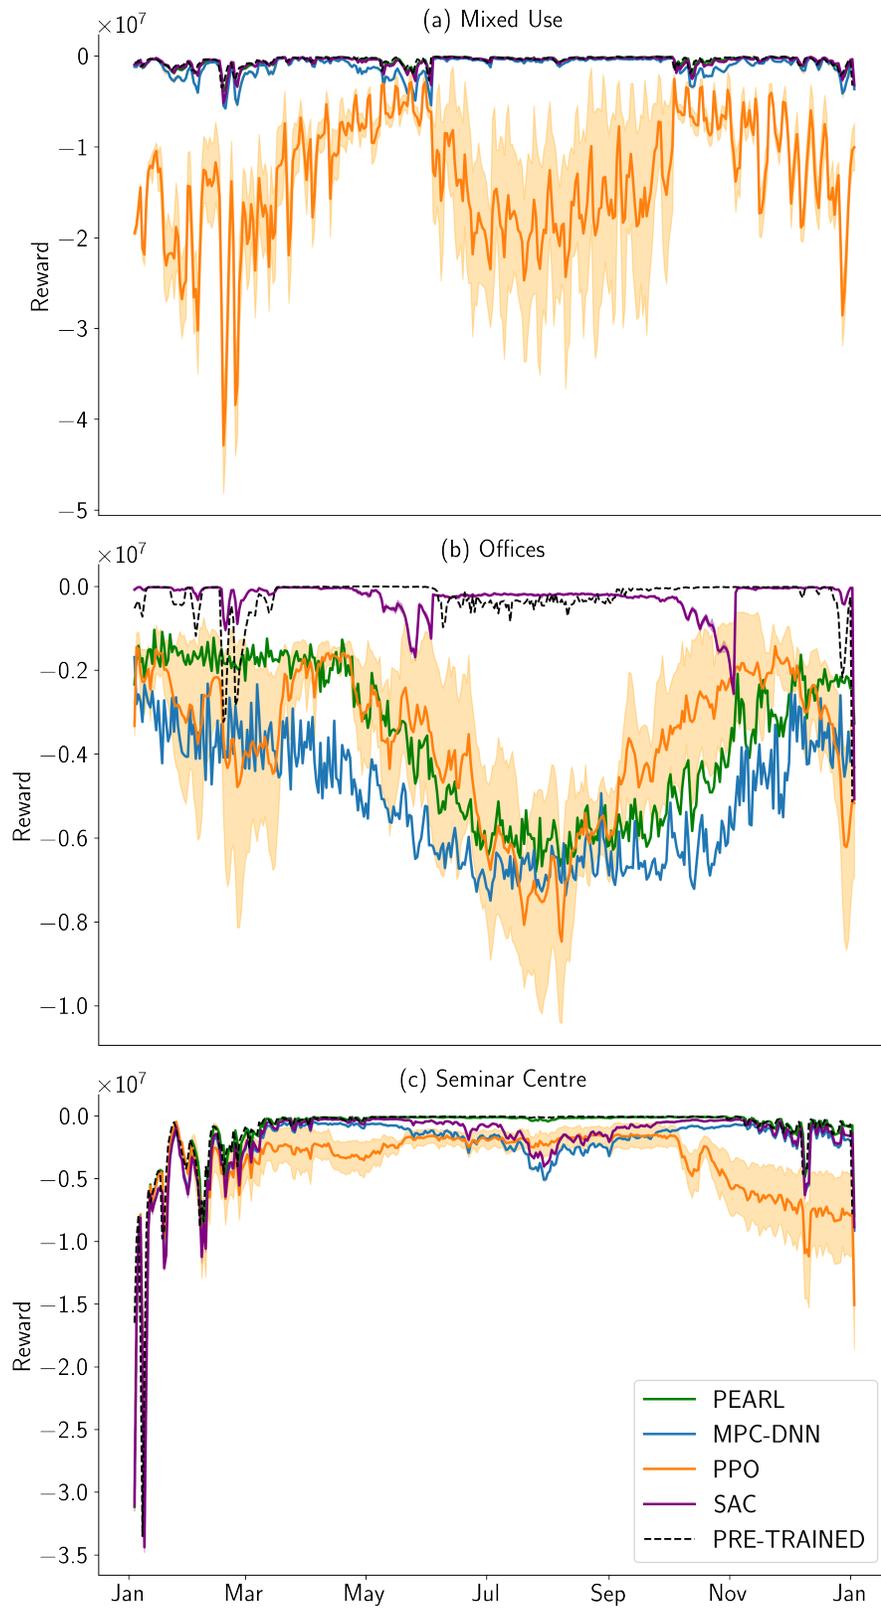


Figure 5. Reward curves for all RL agents across (a) Mixed Use, (b) Office, and (c) Seminar Centre environments.