# **Enforcing Policy Feasibility Constraints through Differentiable Projection for Energy Optimization**

Bingqing Chen\* Carnegie Mellon University Pittsburgh, PA, USA bingqinc@andrew.cmu.edu Priya L. Donti\* Carnegie Mellon University Pittsburgh, PA, USA pdonti@cs.cmu.edu Kyri Baker University of Colorado, Boulder Boulder, CO, USA kyri.baker@colorado.edu

J. Zico Kolter Carnegie Mellon University Pittsburgh, PA, USA zkolter@cs.cmu.edu Mario Bergés Carnegie Mellon University Pittsburgh, PA, USA marioberges@cmu.edu

#### **ABSTRACT**

While reinforcement learning (RL) is gaining popularity in energy systems control, its real-world applications are limited due to the fact that the actions from learned policies may not satisfy functional requirements or be feasible for the underlying physical system. In this work, we propose PROjected Feasibility (PROF), a method to enforce convex operational constraints within neural policies. Specifically, we incorporate a differentiable projection layer within a neural network-based policy to enforce that all learned actions are feasible. We then update the policy end-to-end by propagating gradients through this differentiable projection layer, making the policy cognizant of the operational constraints. We demonstrate our method on two applications: energy-efficient building operation and inverter control. In the building operation setting, we show that PROF maintains thermal comfort requirements while improving energy efficiency by 4% over state-of-the-art methods. In the inverter control setting, PROF perfectly satisfies voltage constraints on the IEEE 37-bus feeder system, as it learns to curtail as little renewable energy as possible within its safety set.

#### **CCS CONCEPTS**

• Hardware  $\rightarrow$  Smart grid; Energy generation and storage; • Computing methodologies  $\rightarrow$  Reinforcement learning.

#### **KEYWORDS**

safe reinforcement learning, implicit layers, differentiable optimization, inverter control, smart building

#### ACM Reference Format:

Bingqing Chen, Priya L. Donti, Kyri Baker, J. Zico Kolter, and Mario Bergés. 2021. Enforcing Policy Feasibility Constraints through Differentiable Projection for Energy Optimization. In *The Twelfth ACM International Conference on Future Energy Systems (e-Energy '21), June 28-July 2, 2021, Virtual Event,* 

\*These authors contributed equally.



This work is licensed under a Creative Commons Attribution International 4.0 License.

e-Energy '21, June 28-July 2, 2021, Virtual Event, Italy © 2021 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-8333-2/21/06. https://doi.org/10.1145/3447555.3464874

Italy. ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3447555.

#### 1 INTRODUCTION

There has been increasing interest in using learning-based methods such as reinforcement learning (RL) for applications in energy systems control. However, a fundamental challenge with many of these methods is that they do not respect the physical constraints or functional requirements associated with the systems in which they operate. Therefore, there have been many calls for embedding safety guarantees into learning-based methods in the context of energy systems applications [20, 29, 65].

One common proposal to address this challenge is to provide machine learning methods with "soft penalties" to encourage them to learn feasible solutions. For instance, the authors of [14, 64] incentivize their RL-based building HVAC controller to satisfy thermal comfort constraints by adding a constraint violation penalty to the reward function. While such approaches often involve tuning some weight on the penalty term, recent work has proposed more theoretically-grounded approaches to choosing these weights; for instance, in the setting of approximating AC optimal power flow, the authors of [13, 25] interpret the weight on their constraint violation penalty as a dual variable, and learn it via primal-dual updates. Gupta et al. [32] adopt a similar approach in an inverter control problem. However, a challenge with these types of "soft penalty" methods in general is that while they incentivize feasibility, they do not strictly enforce it, which is potentially untenable in safety-critical applications.

Given this limitation, a second class of approaches has aimed to strictly enforce operational constraints. For instance, in some cases, the outputs of a machine learning algorithm can be clipped post-hoc in order to make them feasible. However, a challenge is that such post-hoc corrections are not taken into account during the learning process, potentially negatively impacting overall performance. More recent approaches based in deep learning have therefore aimed to enforce simple classes of constraints in a way that *can* be taken into account during learning; for instance, Zamzam and Baker [62] train a neural network to approximate AC optimal power flow (OPF), and enforce box constraints on certain variables via sigmoid activations in the last layer of the neural network. In general, however, existing approaches have only been able

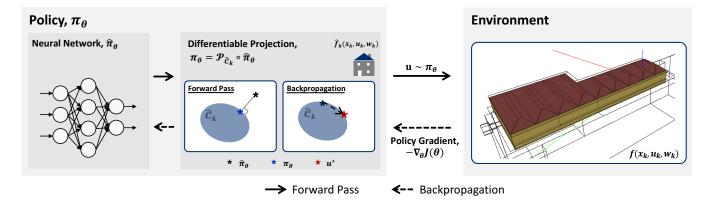


Figure 1: The PROF framework. Our policy consists of a neural network followed by a differentiable projection onto a convexified set of operational constraints,  $\hat{C}_k$  (which is constructed via an approximate model,  $\hat{f}_k$ , of the environment). The differentiable projection layer enforces the constraints in the forward pass, and induces policy gradients that make the neural network cognizant of the constraints in its learning.

to accommodate *simple* sets of constraints, prompting a need for methods that can incorporate broader classes of constraints.

In this work, we propose a method to enforce general convex constraints into RL-based controllers in a way that can be taken into account during the learning process. In particular, we construct a neural network-based policy that culminates in a projection onto a set of constraints characterized by the underlying system. While the "true" constraints associated with the system may be somewhat complex, we observe that simple, approximate physical models are often available for many systems of interest, allowing us to specify convex approximations to the relevant constraints. The projections onto these (approximate) sets can thus be characterized as convex optimization problems, allowing us to leverage recent developments in differentiable convex optimization [2, 5] to train our neural network and projection end-to-end using standard RL methods. The result is a powerful neural policy that can flexibly optimize performance on the true underlying dynamics, while still satisfying the specified constraints.

We demonstrate our PROjected Feasibility approach, PROF, on two settings of interest. Specifically, we explore a building operation setting in which the goal is to reduce energy consumption during the heating season, while ensuring the satisfaction of thermal comfort constraints. We additionally explore an inverter control setting where the goal is to mitigate curtailment, while satisfying inverter operational constraints and nodal voltage bounds. In both settings, we find that our controller achieves good performance with respect to the control objective, while ensuring that relevant operational constraints are satisfied.

To summarize, our key contributions are as follows:

• A framework for incorporating convex constraints. We propose a projection-based method to flexibly enforce convex constraints within neural policies (as summarized in Figure 1). By examining the gradient fields of the differentiable projection layer, we recommend the incorporation of an auxiliary loss for more robust results. We also show

- in an ablation study (Section 5.3) that propagating gradients through the differentiable projection layer is indeed conducive to policy learning.
- Demonstration on building control. In the building control setting, we show that PROF further improves energy efficiency by 10% and 4%, respectively, compared to the best-performing RL agents in [64] and [14]. By using a locally-linear assumption to approximate the building thermodynamics and thereby formulating the constraints as a polytope [15, 66], we largely maintain the temperature within the deadband, except when the control is saturated.
- Demonstration on inverter control. In the inverter control setting, PROF satisfies the voltage constraints 100% of the time over more than half a million time steps (1 week at one second per time step), with a randomly initialized neural network, compared to 22% over-voltage violations incurred by a Volt/Var control strategy. With respect to the objective of minimizing renewable generation curtailment, PROF performs as well as possible within its conservative safety set after learning safely for a day.

### 2 RELATED WORK

Our approach relies on recent developments in implicit neural network layers, and is thematically similar to several recent works in safe RL. We briefly discuss these topics, and refer interested readers to [20, 22, 29, 49, 65] for comprehensive reviews of relevant work in power and energy systems application domains.

Implicit layers. A neural network can be viewed as a composition of functions, or *layers*, with parameters that can be adjusted to improve performance on some task. While many of the layers commonly used within neural networks (e.g., convolutions or sigmoid functions) represent *explicit* functions that provide a direct mapping between inputs and outputs, there has recently been a great deal of interest in expanding the set of commonly-used layers to include those representing *implicit* functions [39]. This has included the creation of layers capturing optimization problems [2, 5, 19, 30, 56, 59],

physical equations [17, 18, 31], sequence modeling processes [6], and games [41]. In this work, we leverage advances in differentiable optimization in particular, namely by incorporating a differentiable convex optimization layer into our neural policy in order to project proposed control actions onto the feasible set of constraints.

Safe reinforcement learning. While (deep) RL methods in general lack safety or stability guarantees, there has been recent interest in learning RL-based controllers that attempt to maintain some notion of safety during training and/or inference - e.g., to satisfy physical constraints or avoid particularly negative outcomes [28]. These include methods that aim to determine "safe" regions of the state space by making smoothness assumptions on the underlying dynamics [3, 10, 57, 58], methods that combine concepts from RL and control theory [12, 21, 33, 42, 44, 46, 63], approaches based on formal verification logics [27, 34, 35], and methods that aim to bound some (discounted) cost function characterizing violations of safety constraints [1, 4, 54, 61]. While the particular notion of "safety" considered varies between settings, relevantly to the present work, several of these prior works employ some form of differentiable projection within the loop of deep RL. For instance, within the context of constrained Markov decision processes (C-MDPs), Yang et al. [61] project neural network-based policies onto a linearly-constrained set of policies with bounded cumulative discounted cost. In the context of asymptotic stability, Donti et al. [21] project the actions output by their controller onto a convex set of actions satisfying stability specifications obtained via robust control. In the setting of robotic motion planning, Pham et al. [45] project actions onto a linear set of robotic operational constraints, and apply separate updates to the neural network based on both pre-projection and post-projection actions. Similarly to this prior work, our approach employs differentiable projections within a neural network policy to enforce operational constraints over some planning horizon.

# 3 PRELIMINARIES

We now present background on technical concepts used by PROF, namely reinforcement learning and differentiable projection layers.

#### 3.1 Reinforcement Learning

The goal of RL is to learn an optimal control policy through direct interaction with the environment. The problem is usually formulated as a Markov decision process (MDP). At each time step k, the agent selects an action  $u_k$  given the current state  $x_k$ , using its policy  $\pi_\theta$  (Equation 1). In many modern RL techniques, the policy is commonly represented by a neural network parameterized by  $\theta$ . When the agent takes the action u, the state transitions to x' based on the system dynamics f (Equation 2), and the agent receives a reward  $r_k$  (or equivalently, incurs a cost  $c_k = -r_k$ ).

$$u \sim \pi_{\theta}(u_k|x_k),\tag{1}$$

$$x' \sim f(x_k, u_k). \tag{2}$$

RL algorithms optimize for a policy that maximizes the expected cumulative reward, or equivalently, minimizes the expected cumulative cost, where  $\gamma$  is a temporal discount factor:

$$\theta^{\star} = \arg\max_{\theta} \mathbb{E}_{\pi_{\theta}} \left[ \sum_{l=0}^{\infty} \gamma^{l} r_{k+l} \right] = \arg\min_{\theta} \mathbb{E}_{\pi_{\theta}} \left[ \sum_{l=0}^{\infty} \gamma^{l} c_{k+l} \right]. \quad (3)$$

To simplify notation, we will denote the expected cumulative cost as  $J(\theta)$ , i.e.,

$$J(\theta) = \mathbb{E}_{\pi_{\theta}} \left[ \sum_{l=0}^{\infty} \gamma^{l} c_{k+l} \right]. \tag{4}$$

There are three general approaches to RL, namely value-based methods, policy gradient methods, and actor-critic methods. Value-based methods, e.g., Q-learning and its variants, update the value function of state-action pairs using the Bellman equation and take the action that maximizes the value of an action selection policy (the Q function) through exploration. Policy gradient methods, e.g., Proximal Policy Optimization (PPO) [51], directly search for an optimal policy  $\pi_{\theta}^{\star}$  using estimates of policy gradients. Denoting the policy gradient as  $g := \nabla_{\theta} J(\theta)$ , the core idea of policy gradient algorithms is that they update  $\theta$  based on an estimate,  $\hat{g}$ , of the gradient, i.e.,

$$\theta \leftarrow \theta - \alpha \hat{q} \tag{5}$$

for some learning rate  $\alpha$ . Different algorithms vary in how they obtain  $\hat{g}$ . For instance, the learning objective for PPO, which we use in our building control experiment (Section 5), is given by the following equation, where  $\hat{A}_t$  is the generalized advantage estimate that can be estimated via any of the estimators in [50]:

$$\begin{split} J_{\text{PPO}}(\theta) &= \hat{\mathbb{E}}_{k} \left[ \min(w_{k}(\theta) \hat{A}_{k}, \text{clip}(w_{k}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{k}) \right], \\ w_{k}(\theta) &= \frac{\pi_{\theta}(u_{k}|x_{k})}{\pi_{\theta_{old}}(u_{k}|x_{k})}, \end{split} \tag{6}$$

and the estimate  $\hat{g}$  is constructed based on this learning objective. Actor-critic methods, e.g., Advantage Actor-Critic (A2C), are hybrids of the value-based and policy gradient approaches, using a policy network to select actions (the actor) and a value network to evaluate the action (the critic).

#### 3.2 Differentiable Projection Layers

As previously described, a neural network is a composition of parameterized functions (*layers*) whose parameters are adjusted during training via *backpropagation* (a class of gradient-based methods). Any function can be incorporated into a neural network as a layer provided that it satisfies two main conditions. The first condition is that it must have a *forward procedure* to map from inputs to outputs (i.e., do *inference*). The second is that it must have a *backwards procedure* to compute gradients of the outputs with respect to the inputs and function parameters, in order to enable backpropagation.

With that in mind, consider the  $L_2$ -norm projection  $\mathcal{P}_C : \mathbb{R}^n \to C$  that maps from some point in  $\hat{u} \in \mathbb{R}^n$  to its closest point in some constraint set  $C \subseteq \mathbb{R}^n$  as follows:

$$\mathcal{P}_C(\hat{u}) = \underset{u \in C}{\operatorname{argmin}} \frac{1}{2} \|u - \hat{u}\|_2^2.$$
 (7)

In cases where C is convex, Equation 7 is a convex optimization problem. The forward procedure of this operation can then be implemented by simply solving the optimization problem, e.g., using standard convex optimization solvers. Perhaps less evidently, it is also possible to construct a backwards procedure for this problem by using the implicit function theorem [40], as described in previous work (e.g., [2, 5]).

As an example, consider the case where C characterizes linear constraints, i.e.,  $C \equiv \{u : Au = b, Gu \le h\}$  for some  $A \in \mathbb{R}^{n_{eq} \times n}, b \in \mathbb{R}^{n_{eq} \times n}$ 

 $\mathbb{R}^{n_{\mathrm{eq}}}$ ,  $G \in \mathbb{R}^{n_{\mathrm{ineq}} \times n}$ , and  $h \in \mathbb{R}^{n_{\mathrm{ineq}}}$ . It is then possible to efficiently compute gradients through Equation 7 by implicitly differentiating through its KKT conditions, i.e., conditions that are necessary and sufficient to describe its optimal solutions. In particular, as described in [5], the KKT conditions for stationarity, primal feasibility, and complementary slackness for this case are given by

$$u^{\star} - \hat{u} + A^{T}v^{\star} + G^{T}\lambda^{\star} = 0$$

$$Au^{\star} - b = 0$$

$$\operatorname{diag}(\lambda^{\star})(Gu^{\star} - h) = 0,$$
(8)

where  $u^*$ ,  $\lambda^*$ , and  $v^*$  are the optimal primal and dual solutions. By the implicit function theorem, we can then take derivatives through these conditions at the optimum in order to obtain relevant gradients. Specifically, the total differentials of these KKT conditions are given by

$$du - d\hat{u} + dA^T v^* + A^T dv + dG^T \lambda^* + G^T d\lambda = 0$$

$$dAu^* + Adu - db = 0 \qquad (9)$$

$$diag(Gu^* - h)d\lambda + diag(\lambda^*)(dGu^* + Gdz - dh) = 0.$$

As described in [5], these equations can then be rearranged to solve for the Jacobians of any of the solution variables  $u^*$ ,  $\lambda^*$ ,  $\nu^*$  with respect to any of the problem parameters  $\hat{u}$ , A, b, G, h (or, in practice, to solve directly for these Jacobians' left matrix-vector product with some backward pass vector, in order to reduce space complexity).

While the above example is for the case of a linearly-constrained projection operation, these kinds of gradients can be computed for convex projection problems in general. For instance, Donti et al. [21] compute gradients through a projection onto a second order cone by differentiating through the fixed point equations of their solver, and Agrawal et al. [2] provide a method and library for differentiable disciplined convex programs. A key benefit of using these kinds of projection layers for constraint enforcement is that they allow gradients through the enforcement procedure to flow back to the neural network, thereby informing the parameter updates of this network during training.

#### 4 APPROACH

We now describe PROF, which incorporates differentiable projections onto convex(ified) sets of operational constraints within a neural policy.

#### 4.1 Problem Formulation

Consider a discrete-time dynamical system

$$x_{k+1} = f(x_k, u_k, w_k),$$
 (10)

where  $x_k \in \mathbb{R}^s$  is the state at time  $k, u_k \in \mathbb{R}^a$  is the control input,  $w_k \in \mathbb{R}^d$  is an uncontrollable disturbance (which we assume to be observable), and  $f : \mathbb{R}^s \times \mathbb{R}^a \times \mathbb{R}^d \to \mathbb{R}^s$  denotes the system dynamics. Letting  $X_k$  and  $\mathcal{U}_k$  denote the allowable state and action space, respectively, we can define the set of all feasible actions over the planning horizon T as  $C_k$ , where

$$C_{k} = \left\{ u_{k:k+T-1} \middle| \begin{array}{l} x_{i+1} = f(x_{i}, u_{i}, w_{i}), \\ x_{i} \in X_{i}, u_{i} \in \mathcal{U}_{i} \end{array} \middle| \forall i \in \{k, ..., k+T-1\} \right\}.$$
(11)

Our goal is then to learn a policy that optimizes the control objective, J, while enforcing the operational constraints. To simplify notation, we denote  $\mathbf{u} = u_{k:k+T-1}$ . In the case of a deterministic policy, i.e.,  $\mathbf{u} = \pi_{\theta}$ , the learning problem is simply

$$\min_{\theta} J(\theta) \quad \text{s.t. } \pi_{\theta} \in C_k.$$
(12)

In the case of a stochastic policy, e.g.  $\mathbf{u} \sim \mathcal{N}(\boldsymbol{\mu}, \operatorname{diag}(\sigma^2))$ ,  $[\boldsymbol{\mu}, \boldsymbol{\sigma}] = \pi_{\theta}(x_k)$ , we can write the problem as

$$\min_{\theta} J(\theta) \quad \text{s.t. } \mathbf{u}, \boldsymbol{\mu} \in C_k.$$
(13)

In this case, it is necessary to sample actions around  $\mu$  in order to estimate policy gradients. At the same time the actions sampled from  $\pi_{\theta}$  might fall outside of  $C_k$ . Thus, we enforce that both  $\mu$  and the sample action  $\mathbf{u}$  satisfy the constraints.

#### 4.2 Approximate Convex Constraints

In practice, there are two key challenges inherent in solving Equations 12–13 as written. The first is that the disturbances  $w_i$  are not known ahead of time, meaning that the optimization problem must be solved under uncertainty. One approach to addressing this, from the field of robust control [67], involves constructing an uncertainty set over the disturbance, and then optimizing for worst-case or expected cost under this uncertainty set. Here, we simply assume a predictive model of the disturbances is available. (By re-planning frequently, we observe that the prediction errors have limited empirical impact on performance in the two applications we study.) We will use the notation  $\hat{w}_k$  to denote our forecast of the disturbance if k is a future time step, and the true value of the disturbance if k is the present or a prior time step.

The second challenge pertains to the form of the set  $C_k$ , which may be poorly structured or otherwise difficult to optimize over. In particular, our framework relies on obtaining convex approximations to the constraints in order to enable differentiable projections (see Section 3.2). Fortunately, for many energy systems applications, some approximate model  $\hat{f_k}$  is often available based on domain knowledge that allows  $C_k$  to be approximated as a convex set, despite the complex nature of the true dynamical system.

Thus, letting  $\hat{f}_i$  denote our approximations of the dynamics and  $\hat{w}_i$  denote the (forecast or known) disturbance at each i = k, ..., k + T - 1, we define our approximate convex constraint set as

$$\hat{C}_{k} = \left\{ u_{k:k+T-1} \middle| \begin{array}{c} x_{i+1} = \hat{f}_{i}(x_{i}, u_{i}, \hat{w}_{i}), \\ x_{i} \in \mathcal{X}_{i}, u_{i} \in \mathcal{U}_{i} \end{array} \right. \forall i \in \{k, ..., k+T-1\} \right\}.$$
(14)

We note that f and w are approximated *solely* for the purposes of constructing approximate constraint sets, and are not used otherwise during training and inference (i.e., our neural policy interacts with the *true* dynamics and disturbances during training and inference).

#### 4.3 Policy Optimization

Let  $\hat{\pi}_{\theta}$  be any (e.g., fully-connected or recurrent) neural network parameterized by  $\theta$ . Our policy entails passing the output from the neural network to the differentiable projection layer  $\mathcal{P}_{\hat{C}_k}$  characterized by the approximate constraints, which enforces that the

#### Algorithm 1 PROF

```
1: procedure MAIN(env, J) // input: environment, control objective
          init neural network \hat{\pi}_{\theta}, replay memory \mathcal{M}
 2:
          specify RL algorithm \mathcal{A}, batch size M, update interval K
 3:
          specify planning horizon T
 4:
          // online execution
 5:
          for k = 1, ... do
 6:
               observe state x_k
 7:
               predict future disturbances \hat{w}_{k:k+T-1}
               construct constraint set \hat{C}_k, policy \pi_{\theta} = \mathcal{P}_{\hat{C}_k} \circ \hat{\pi}_{\theta}
 9:
               compute u_k = Inference(\pi_\theta, x_k, T)
10:
11:
               execute action ENV.STEP(u_k)
               save memory.Append(x_k, u_k, \hat{w}_{k:k+T-1})
12:
               // update policy every K time steps
13:
               if mod(k, K) = 0 then
14:
                     \hat{\pi}_{\theta} = \text{TRAIN}(\hat{\pi}_{\theta}, J, \mathcal{M}, \mathcal{A})
15:
               end if
16:
          end for
17:
    end procedure
18:
19:
    procedure Inference(\pi_{\theta}, x_k, T)
20:
          // input: neural policy, current state, planning horizon
21:
          select action u_{k:k+T-1} \sim \pi_{\theta}
22:
            // only return the current action; replan at each time step
23:
          return u_k
    end procedure
24:
25:
     procedure TRAIN(\hat{\pi}_{\theta}, J, \mathcal{M}, \mathcal{A})
26:
          // input: neural policy, objective, replay memory, RL algorithm
27
          init \mathcal{L}(\theta) = 0
28:
          for i = 1, ..., M do
29:
               sample x, u, w \sim \mathcal{M}
30:
               construct constraint set \hat{C}_k, policy \pi_{\theta} = \mathcal{P}_{\hat{C}_k} \circ \hat{\pi}_{\theta}
31:
               compute training loss
32:
                     \mathcal{L}(\theta) += J(\theta) + \lambda \|\pi_{\theta}(x) - \hat{\pi}_{\theta}(x)\|_{2}^{2}
          end for
33:
          train \hat{\pi}_{\theta} via \mathcal{A} to minimize \mathcal{L}
34:
          return \hat{\pi}_{\theta}
35:
36: end procedure
```

resultant action is feasible with respect to these constraints. The overall (differentiable) neural policy is then given by

$$\pi_{\theta}(x_k) = \mathcal{P}_{\hat{C}_k} \circ \hat{\pi}_{\theta}(x_k).^{1} \tag{15}$$

The key benefit of embedding a differentiable projection into our policy is that it enforces constraints in a way that is visible to the neural network during learning. In this work, we implement the differentiable projection using the cvxpylayers library [2].

We construct the following loss function, which is a weighted sum of the control objective J and an auxiliary loss term to be explained shortly in this section.  $\lambda > 0$  is a hyperparameter.

$$\mathcal{L}(\theta, x_k) = J(\theta) + \lambda \|\pi_{\theta}(x_k) - \hat{\pi}_{\theta}(x_k)\|_2^2. \tag{16}$$

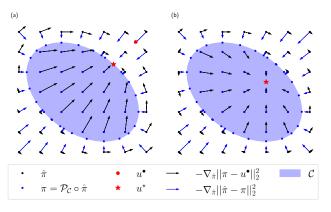


Figure 2: Illustrative example of gradients from the differentiable projection layer.  $u^{\bullet}$  and  $u^{\star}$  denote unique optimal actions minimizing some convex control objective J in the unconstrained and constrained settings, respectively;  $\nabla_{\hat{\pi}} \| \pi - u^{\bullet} \|_2^2$  is thus a proxy for  $\nabla_{\hat{\pi}} J$ . (a)  $u^{\bullet} \notin C$ . The gradients  $\nabla_{\hat{\pi}} J$  point towards  $u^{\star}$  as desired, such that  $\pi = \mathcal{P}_{\hat{C}} \circ \hat{\pi}$  will reach this optimal point. (b)  $u^{\bullet} = u^{\star}$  on the interior of C. The gradients  $\nabla_{\hat{\pi}} J$  do not cause  $\hat{\pi}$  (or its projection) to update towards the interior. Adding a weighted auxiliary loss term, e.g.,  $\|\pi - \hat{\pi}\|$ , can help direct updates towards the interior.

We then train our policy (Equation 15) to minimize this cost using standard approaches in deep reinforcement learning. The full algorithm is presented in Algorithm 1.

4.3.1 Visualization of gradient fields. To provide more intuition on the differentiable projection layer and our cost function, we visualize the gradient fields in a hypothetical example with a deterministic policy and a planning horizon of T=1. Specifically, for the purposes of illustration, let  $u^{\bullet}$  and  $u^{\star}$  denote unique optimal actions minimizing some convex control cost J in the unconstrained and constrained settings, respectively:

$$u^{\bullet} \sim \pi_{\theta^{\bullet}}; \quad \theta^{\bullet} = \arg\min_{\theta} J(\theta)$$
  
 $u^{\star} \sim \pi_{\theta^{\star}}; \quad \theta^{\star} = \arg\min_{\theta} J(\theta) \quad \text{s.t. } u \in C_k.$ 

In Figure 2, we then plot the gradient fields in two cases: (a)  $u^{\bullet} \notin C_k$ , and (b)  $u^{\bullet} \in C_k$ . Note that  $u^{\bullet}$  and  $u^{\star}$  are assumed to be known here for illustrative purposes only, and are not known during training.

In particular, we plot the gradients (black arrows) of  $\|u^{\bullet} - \mathcal{P}_{C_k} \circ \hat{\pi}\|_2^2$  with respect to the output of the neural network  $\hat{\pi}$ . These indicate the direction in which the neural network would be incentivized to update in order to minimize the system cost. If no differentiable projection were embedded within the policy, all the gradients would point towards  $u^{\bullet}$  without regard for the constraints. Instead, in the case of  $u^{\bullet} \notin C_k$  (Figure 2a), the gradients through the differentiable projection layer point towards  $u^{\star}$  instead of  $u^{\bullet}$ . More specifically, if  $\hat{\pi}_{\theta}(x_k) \in C_k$ , then the projection layer is simply the identity, and the gradients point directly towards  $u^{\star}$ ; otherwise, the gradients point along the boundary of  $C_k$  in the direction of  $u^{\star}$ .

<sup>&</sup>lt;sup>1</sup>We use the notation  $f \circ g(x) \coloneqq f(g(x))$  to denote function composition.

This case is of particular interest, as in many practical applications some operational constraint will be binding. As a concrete example, the ultimate energy-saving strategy for building operations is to keep all mechanical systems off (i.e.,  $u^{\bullet}=0$ ), which obviously violates occupants' comfort requirements and is outside the set of allowable actions (i.e.,  $u^{\bullet} \notin C_k$ ). Thus, the problem is to find a policy that uses the mechanical system as little as possible without violating comfort requirements. Given the common case where the control objective is convex, this then lies on the boundary of the constraint set (i.e.,  $u^{\star} = \mathcal{P}_{C_k} \circ u^{\bullet}$ ).

We also depict the case where the solution of the unconstrained problem already satisfies the constraints, i.e.,  $u^{\bullet} = u^{\star} \in C_k$  (Figure 2b). If this is generally the case for a particular application, we note that a constraint enforcement approach (ours or otherwise) is likely not needed, and indeed utilizing gradients through the projection layer may actually degrade performance. Specifically, if  $\hat{\pi}_{\theta}(x_k) \notin C_k$ , the gradients do not point towards the interior of the constraint set, meaning that  $\pi_{\theta}(x_k) = \mathcal{P}_{C_k} \circ \hat{\pi}_{\theta}(x_k)$  will lie on the boundary of the constraints despite the optimal solution being in the interior. This can be amended by augmenting the loss function with a (weighted) auxiliary term such as  $\|\pi_{\theta}(x_k) - \hat{\pi}_{\theta}(x_k)\|_2^2$  whose gradients (blue arrows) point towards the interior.

It may not be known a priori whether or not  $u^{\bullet}$  is in the constraint set in general or at any given time, except when domain experts are fully clear on the structure of the solutions for specific applications. In particular,  $C_k$  is time-varying, making it difficult to know for sure whether or not the constraints will indeed be binding at any given time. For robustness, we therefore recommend incorporating the auxiliary loss  $\|\pi_{\theta}(x_k) - \hat{\pi}_{\theta}(x_k)\|_2^2$  within the RL training cost, unless it is known from domain knowledge that the constraints will certainly be active. As such, we formulate the training cost function as previously given in Equation 16.

# 5 EXPERIMENT 1: ENERGY-EFFICIENT BUILDING OPERATION

There is significant potential to save energy through more efficient building operation. Buildings account for about 40% of the total energy consumption in the United States, and it is estimated that up to 30% of that energy usage may be reduced through advanced sensing and control strategies [24]. However, this potential is largely untapped, as the heterogeneous nature of building environments limits the ability of control strategies developed for one building to scale to others [14]. RL can address this challenge by adapting to individual buildings by directly interacting with the environment.

The most important constraint in building operation is to maintain a satisfactory level of comfort for occupants, while minimizing energy consumption. It is common in the RL-based building control literature to penalize thermal comfort violations [14, 64], which incentivizes but does not guarantee the satisfaction of these comfort requirements. In comparison, our proposed neural policy can largely maintain temperature within the specified comfortable range, except when the control is saturated.

We evaluate our policy in the same simulation testbed as [14, 64], following the same experimental setup as [14]. Specifically, we first pre-train the neural policy by imitating a proportional-controller

(P-controller). We then evaluate and further train our agent in the simulation environment, using a different sequence of weather data.

## 5.1 Problem Description

Simulation testbed. We utilize an EnergyPlus (E+) model of a 600m<sup>2</sup> multi-functional space (Figure 3a), based on the Intelligent Workplace (IW) on Carnegie Mellon University (CMU) campus, located in Pittsburgh, PA, USA. The system of interest is the water-based radiant heating system, of which a schematic is provided in Figure 3b. In this experiment, we control the *supply water temperature* so as to maintain the state variable, i.e., the *zone temperature*, within a comfortable range during the heating season. In the existing control, the supply water (SW) is maintained at a constant flow rate, and its temperature is managed by a P-controller. For more information on the simulation testbed, refer to [64].

Approximate system model. We approximate the environment as a linear system as follows:

$$x_{k+1} \approx \hat{f}(x_k, u_k, w_k) = Ax_k + B_u u_k + B_d w_k,$$
 (17)

where  $x_k$  represents the zone temperature and  $u_k$  represents the supply water temperature.  $w_k$  includes distributions from weather and occupancy. While building thermodynamics are fundamentally nonlinear, the locally-linear assumption works well for many control inputs [47]. We identify the approximate model parameters A,  $B_u$ , and  $B_d$  with prediction error minimization [47] on the same data used to pre-train the RL agent (see Section 5.2). The root mean squared error (RMSE) of this model on a unseen test set is  $0.14^{\circ}$ C.

Objective. Since our goal is to minimize energy consumption, we define the control cost at each time step as the agent's control action, i.e. *supply water temperature*, which is linearly proportional to the heating demand, i.e.,  $c_k = u_k$ .

In contrast to the objectives in [14, 64], which are defined as weighted sum of energy cost and some penalty on thermal comfort violations, we consider the thermal comfort requirement as hard constraints, in the form of Equation 13.

*Constraints.* To maintain a satisfactory comfort level, we require the zone temperature to be within a deadband  $X = \{x \mid 21.9^{\circ}C \le x \le 25.5^{\circ}C\}$  when the building is occupied, based on the building code requirement of 10% Predicted Percentage of Dissatisfied (PPD) [23]. We allow for a wider temperature range during unoccupied hours. For the action, the allowable range of supply water temperature for the physical system is  $\mathcal{U} = \{u \mid 20^{\circ}C \le u \le 65^{\circ}C\}$ .

While it may appear from this description that we have only simple box constraints on both the state and action, we highlight the fact that actions are coupled over time through the building thermodynamics [66]. More concretely, a future state depends on all past actions. Thus, a box constraint on  $x_{k+l+1}$  is in fact a constraint on  $u_{k:k+l}$ . In this case, assuming  $\hat{f}$  to be a linear system,  $\hat{C}_k$  is then a set of linear inequalities, which can be geometrically interpreted as a polytope. We refer interested readers to [15, 66] for more details on this formulation. In fact, it was experimentally demonstrated in [15] that projecting actions onto the polytope constructed with an approximate linear model was sufficient to maintain temperature

<sup>&</sup>lt;sup>2</sup>A polytope can be characterized as a set  $S = \{x \in \mathbb{R}^n | Ax \le b\}$ .

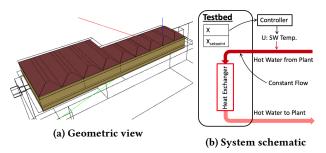


Figure 3: Building simulation testbed (reproduced from [14]).

within the deadband in a real-world residential household (though [15] did not then differentiate through this projection).

Control time step. The EnergyPlus model has a 5-minute simulation time step. Following [14, 64], we use a 15-min control time step (i.e., each action is repeated 3 times) and a planning horizon of T = 12 (i.e., a 3 hour look-ahead).

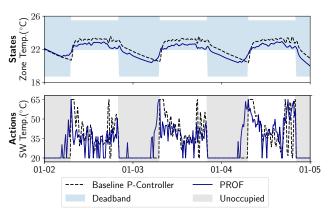
# 5.2 Implementation Details

Offline pre-training. We pre-train a long short-term memory (LSTM) recurrent policy (without a subsequent projection) by imitating a P-controller operating under the Typical Meteorological Year 3 (TMY3) [60] weather sequence, from Jan. 1 to Mar. 31. We min-max normalize all of the state, action, and disturbance, and use a learning rate of  $10^{-3}$ . Specifically, we use the pre-trained weights after training on the expert demonstrations for 20 epochs following the same procedures as [16]. We refer readers to [16] for more details on the neural network architecture, training procedures, loss, and performance evaluation.

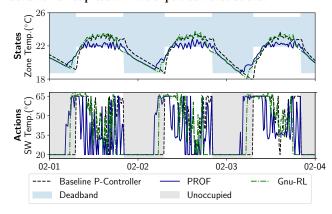
Online policy learning. We optimize the policy with PPO [51] over the weather sequence in 2017 from Jan. 1 to Mar. 31. We use  $\lambda=10$  (see Equation 16), a learning rate of  $5\times 10^{-4}$ , and RMSprop [55] as the optimizer.<sup>3</sup> We update the policy every four days, by iterating over those samples for 8 epochs with a batch size of 32. For hyperparameters, we use a temporal discount rate of  $\gamma=0.9$ ,  $\epsilon=0.2$  (see Equation 6), and a Gaussian policy (see Equation 13) with  $\sigma$  linearly decreased from 0.1 to 0.01.

# 5.3 Results

After pre-training on expert demonstrations from the baseline P-controller, our agent directly operated the simulation testbed based on actual weather sequences in Pittsburgh from Jan. 1 to Mar. 31 in 2017. Figure 4a shows the behavior of our agent at the onset of deployment over a 3-day period. The baseline P-controller reactively turns on heating when the environment switches from unoccupied to occupied, which results in thermal comfort violations in the mornings. In comparison, PROF preheats the environment such that the environment is already at a comfortable temperature when occupants arrive in the morning. Notably, the differentiable projection layer manages to enforce this preheating behavior despite this behavior not being present in the expert demonstrations.



(a) The differentiable projection layer enforces preheating behavior to ensure deadband constraints are never violated, even though this behavior is not present in the expert demonstrations.



(b) The agent has found a more energy-efficient control strategy by maintaining temperature at the lower end of the deadband.

Figure 4: Behavior of our proposed agent (a) at the onset of deployment, with pre-trained weights based on expert demonstrations and (b) after a month of interacting with and training on the environment.

Figure 4b shows the behavior of our agent in comparison with Gnu-RL [14], having interacted with and trained on the environment for a month. Gnu-RL is updated via PPO, similarly to the current work, and incorporates domain knowledge on system dynamics. In comparison to Gnu-RL [14], which ends up trying to maintain temperature at the setpoint, PROF learns an energy-saving behavior by maintaining the temperature at the lower end of the deadband. This explains the further energy savings compared with Gnu-RL [14]. However, we also notice that the temperature requirement may be violated on cold mornings. This happens when the control action is saturated, i.e., full heating over the 3-hour planning horizon is not sufficient to bring temperature back to the comfortable range. (In principle, even these constraint violations could be mitigated by increasing the length of the planning horizon.)

Table 1 summarizes the performance of our agent with comparison to the RL agents in [14, 64]. Our proposed agent (averaged

 $<sup>^3{\</sup>rm The~code}$  is available at https://github.com/INFERLab/PROF.

Table 1: Performance comparison. Our method saves energy while incurring minimal comfort violations.

	Heating PPD		
	Demand	Mean	SD
	(kW)	(%)	(%)
Existing P-Controller [64]	43709	9.45	5.59
Agent #6 [64]	37131	11.71	3.76
Baseline P-Controller [14]	35792	9.71	6.87
Gnu-RL [14]	34687	9.56	6.39
LSTM & Clip + No Update	37938	8.55	3.39
LSTM & Clip	$36068 \pm 2187$	$9.18 \pm 0.67$	3.49
PROF (ours)	$33271 \pm 1862$	$9.68 \pm 0.48$	3.66

over 5 random seeds) saves 10% and 4% energy compared to the best-performing agents in [64] and [14], respectively.

We also compare our method to two ablations: (1) LSTM & Clip + No Update, which uses the same pre-trained weights and the projection layer to enforce feasible actions, but does not update the policy, and (2) LSTM & Clip, which uses the same pre-trained weights and the projection layer to enforce feasible actions during inference, but does not propagate gradients through the differentiable projection layer in the policy updates. We find that LSTM & Clip slightly improves upon LSTM & Clip + No Update, but is less performant compared to PROF. This affirms our hypothesis that the gradients through the differentiable projection layer are cognizant of the constraints and are thus conducive to policy learning.

# **6 EXPERIMENT 2: INVERTER CONTROL**

Distributed energy resources (DERs), e.g., solar photovoltaic (PV) panels and energy storage, are becoming increasingly prevalent in an effort to curb carbon dioxide emissions and combat climate change. However, DERs interfacing with the power grid via power electronics, such as inverters, also introduce unintended challenges for grid operators. For instance, over-voltages have become a common occurrence in areas with high renewable penetration [53], and power electronics-interfaced generation has low-inertia and requires active control at much faster timescales compared to traditional synchronous machines [43].

To alleviate these issues, IEEE standard 1547.8-2018 [9] recommends a Volt/Var control strategy in which the reactive power contribution of an inverter is based on local voltage measurements. As will be clear in our empirical evaluation, this network-agnostic heuristic based on local information alleviates, but does not avoid, over-voltage issues. Given that the optimal solution needs to be obtained at the system-level and that the problem needs to be solved at very short timescales, a common paradigm is to address the problem in a quasi-static fashion [38] adopted in works such as [7, 32, 38], where one chooses a policy over the next time period, e.g., 15 minutes-1 hour, and uses the policy without update for fast inference. In this work, we adopt the same paradigm and consider real-time control on a 1-second timescale of both active (P) and reactive (Q) power setpoints at each inverter.

We envision that a neural policy can learn from its prior experiences, in contrast to the traditional *fit-and-forget* approach [20],

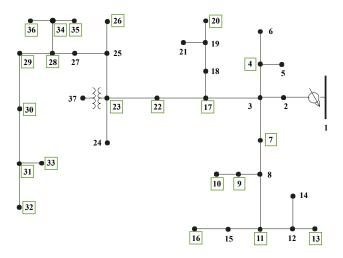


Figure 5: IEEE 37-bus feeder system, where the solar PV systems are indicated by green rectangles.

and is capable of making decisions faster compared to solving optimization problems. Our primary contribution compared to existing work is the ability to enforce physical constraints within the neural network. In fact, we successfully enforce voltage constraints 100% of the time with a randomly initialized neural network, over more than half a million time-steps (i.e., 1 week with a one-second time step). The assumed control and communication scheme is consistent with the new definitions for smart inverter capabilities under IEEE standard 1547.1-2020 [37].

#### 6.1 Problem Description

The problem we are considering here is to control active and reactive power setpoints at each inverter in order to maximize utilization (i.e., minimize curtailment) of renewable generation, while satisfying the maximum and minimum grid voltage requirements. Here, we first define the considered test case and input data, and describe the model of the network. We refer readers to [7] for more details on the problem set-up.

*IEEE 37-bus test case.* We evaluate our method on the IEEE 37-bus distribution test system [36], with 21 solar PV systems indicated by green rectangles in Figure 5. We utilize a balanced, single-phase equivalent of the system, and simulate the nonlinear AC power flows using PYPOWER [48]. For the simulation, the solar generation and loads are based on 1-second solar irradiance and load data collected from a neighborhood in Rancho Cordova, CA [8] over a period of one week (604800 samples).

Approximate system model. Denote the number of buses, excluding the slack bus (e.g., the distribution substation), as N, the net active and the reactive power as  $\mathbf{p} \in \mathbb{R}^N$  and  $\mathbf{q} \in \mathbb{R}^N$ , and the voltage at all buses as  $\mathbf{v} \in \mathbb{R}^N$ . We linearize the AC power flow equations around the flat voltage solution, i.e.  $\bar{\mathbf{v}} = \mathbb{1}$ , using the method in [11]. The reference active and reactive power corresponding to  $\bar{\mathbf{v}} = \mathbb{1}$  is denoted as  $\bar{\mathbf{p}}$  and  $\bar{\mathbf{q}}$ . The linearized grid model,

 $\hat{f}$ , is given by Equation 18, where  $\mathbf{R}, \ \mathbf{B} \in \mathbb{R}^{N \times N}$  represent system-dependent network parameters that can be either estimated from linearization (e.g., [11]) or data-driven methods:

$$\mathbf{v} \approx \hat{f}(\mathbf{p}, \mathbf{q}) = \bar{\mathbf{v}} + \mathbf{R}(\mathbf{p} - \bar{\mathbf{p}}) + \mathbf{B}(\mathbf{q} - \bar{\mathbf{q}})$$

$$= \bar{\mathbf{v}} + \underbrace{\begin{bmatrix} \mathbf{R}, \mathbf{B} \end{bmatrix}}_{\mathbf{H}} \underbrace{\begin{bmatrix} \mathbf{p} - \bar{\mathbf{p}}, \\ \mathbf{q} - \bar{\mathbf{q}} \end{bmatrix}}_{\mathbf{u}}.$$
(18)

A notable advantage of the method in [11] is that the resulting model has bounded error with respect to the true dynamics. By incorporating the error bound when constructing the safety set, the safety set is guaranteed to be a conservative under-approximation of the true safety set, and thus allow us to satisfy voltage constraints 100% of the time.

*Policy.* Our policy takes as input the voltage from the previous time-step, load, and generation at all the buses, and outputs active and reactive power setpoints at each inverter. (This is a deterministic policy; see Equation 12.) Note that while the grid model (Equation 18) contains all N buses, only those with inverters are controllable.

Our neural architecture is similar to the one used in [32], which consists of a utility-level network, and inverter-level networks for individual inverters. The utility-level network collects information from all nodes, and broadcasts an intermediate representation to all inverter-level networks. Using this information along with its local observations, each inverter makes its local control decisions, which are then projected onto the constraints (discussed below).

Objective. The objective is to minimize the curtailment of solar generation, or equivalently to maximize the utilization of the available solar power,  $p_{av}$ . Specifically, letting  $\mathcal I$  denote the set of buses with inverters, the objective is

$$J(\theta) = \min_{\mathbf{p}_{I}, \mathbf{q}_{I}} \sum_{i \in I} [p_{av,i} - p_{i}]_{+}, \quad \text{where } [\mathbf{p}_{I} \quad \mathbf{q}_{I}] = \pi_{\theta} \quad (19)$$

Constraints. For an individual inverter, i, with rated power  $s_i$  and an available power (from available solar generation)  $p_{av,i}$ , the feasible action space is

$$\mathcal{U}_i(k) = \{(p_i, q_i) : 0 \le p_i \le p_{av,i}(k), p_i^2 + q_i^2 \le s_i^2\}$$
  
$$\mathcal{U}(k) := \mathcal{U}_1(k) \times \cdots \times \mathcal{U}_{|\mathcal{I}|}(k).$$

At the same time, the voltage at each bus should remain between 0.95-1.05~p.u. The primary challenge of satisfying voltage constraints is that the voltage at each bus depends on actions of neighboring nodes, i.e.

$$\mathcal{X} = \{ v \mid 0.95 \times \mathbb{1} \le \mathbf{v} \approx \bar{\mathbf{v}} + \mathbf{H}\mathbf{u} \le 1.05 \times \mathbb{1} \},\$$

where the sparsity pattern of H is characterized by the admittance matrix. We jointly project actions from all inverters at each time step k onto the constraints  $\mathcal{U}(k) \cap \mathcal{X}$ .

# 6.2 Implementation Details

We evaluate PROF by executing it over the 1-week dataset (at 1 second) once. Similarly to other quasi-static approaches, we update the policy every 15-minutes. Similarly to [32], we optimize the neural policy with stochastic samples by directly differentiating

through the objective (Equation 19) and the linearized grid model (Equation 18). However our method differs in that Gupta et al. [32] characterized the constraints as a regularization term, and learned the policy via primal-dual updates. We incorporate the constraints directly via the differentiable projection layer and thus guarantee constraint satisfaction.

We use  $\lambda$ =10 (see Equation 16), a learning rate of  $10^{-3}$ , and RMSprop [55] as the optimizer. At every 15 minutes, we sample 16 batches of data with size of 64 from the replay memory. We keep a replay memory size of 86400, i.e., samples from the previous day. For the both the utility-level network and the inverter-level network, we use fully-connected layers with ReLU activations. The utility-level network has hidden layer sizes (256, 128, 64), and each inverter-level network has hidden layer sizes (16, 4) and outputs active and reactive power. On top of the neural network, we implement the differentiable projection layer, following the constraints described in Section 6.1.

We compare our methods to three baselines, (1) a Volt/Var strategy following IEEE 1547.8 [9], (2) the optimal solution with respect to the linearized grid model, and (3) the optimal solution with respect to the true AC power flow equations.

#### 6.3 Results

The performance of PROF in comparison to the three baselines is summarized in Figure 6. For clarity, we only show the maximum voltage over all buses; under-voltage is not a concern for this particular test case.

We see that the Volt/Var strategy violates voltage constraints 22.3% of time, mostly around noon when the solar generation is high and there is a surplus of energy. Since the Volt/Var baseline does not adjust active power, there is no curtailment.

In comparison, PROF satisfies the voltage constraints throughout the experiment, even with a randomly initialized neural policy. While PROF performs poorly on the first morning, it quickly improves its policy. In fact, the behavior of PROF is barely distinguishable from the optimal solution with respect to the linearized grid model, after learning safely for a day. This implies that PROF learned to control inverters as well as possible given its approximate model, which constructs a conservative under-approximation of the true safety set.

The optimal baseline with respect to the true AC power flow equations unsurprisingly achieves the best performance with respect to minimizing curtailment, as it can push the maximum voltage to the allowable limit in order to maximally reduce the amount of curtailed energy. However, inverter control is a task that requires near real-time inputs, and we find that running this baseline can be prohibitively slow. Specifically, we evaluate the computation time of different operations by averaging over 1000 randomly sampled problems from our dataset on a personal laptop. For PROF, on average, a forward pass in the neural network (excluding the projection layer) took 4.5 ms and the differentiable projection operation took 8.6 ms. The computation cost of the differentiable projection could be further reduced by using customized projection solvers such as the ones in [5, 21] that avoid the "canonicalization" costs introduced by general-purpose solvers such as the one we use [2]. In comparison, solving the optimization baseline with respect to the

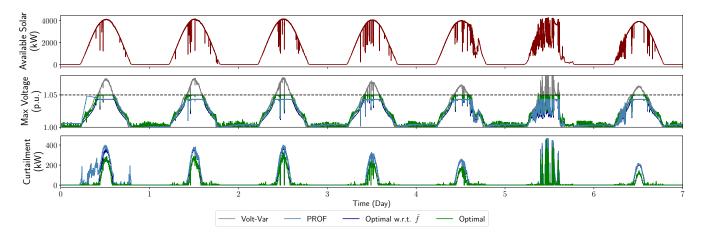


Figure 6: PROF satisfies voltage constraints throughout the experiment, and learns to minimize curtailment as well as possible within its conservative safety set,  $\hat{C}_k$ , after learning safely for a day.

true AC power flow equations took 1.02s on the same machine, which is even longer than the 1s control time-step.

#### 7 DISCUSSION AND CONCLUSIONS

In this work, we have presented a method, PROF, for integrating convex operational constraints into neural network policies for energy systems applications. In particular, we propose a policy that entails passing the output of a neural network to a differentiable projection layer, which enforces a convex approximation of the operational constraints. These convex constraint sets are obtained using approximate models of the system dynamics, which can be fit using system data and/or constructed using domain knowledge. We can then train the resultant neural policy via standard RL algorithms, using an augmented cost function designed to effect desirable policy gradients. The result is that our neural policy is cognizant of relevant operational constraints during learning, enhancing overall performance.

We find in both the building energy optimization and inverter control settings that PROF successfully enforces relevant constraints while improving performance on the control objective. In particular, in the building thermal control setting, we find that our approach achieves a 4% energy savings over the state of the art while largely maintaining the temperature within the deadband. In the inverter control setting, our method perfectly satisfies the voltage constraints over more than half a million time steps, while learning to minimize curtailment as much as possible within the safety set.

While these results demonstrate the promise of our method, a key limitation is in its computational cost. In particular, computing a projection during every forward pass of training and inference is decidedly more expensive than running a "standard" neural network. A fruitful area for future work – both in the context of our method, and in the context of research in differentiable optimization layers as a whole – may be to improve the speed of such differentiable projection layers. For instance, this might entail developing special-purpose differentiable solvers [5, 21] for optimization problems commonly encountered in energy systems applications, developing approximate solvers that do not rely on obtaining optimal solutions

in order to compute reasonable gradients, or employing cheaper projection schemes such as  $\alpha$ -projection [52] where possible.

Additionally, the success of our method (and many other constraint enforcement methods) depends fundamentally on the quality of the approximate model used to characterize the constraint sets. In particular, this determines the extent to which the resultant approximate constraint sets are a good representation of the true operational constraints. While we were able to employ reasonably high-quality approximation schemes in the context of this work, future work on safely updating the models or the constraint sets directly [26] may greatly improve the quality of the solutions.

More generally, while our work highlights one approach to enforcing physical constraints within learning-based methods, we believe this is only the start of a broader conversation on closely integrating domain knowledge and control constraints into learning-based methods. In particular, strictly enforcing physical constraints will be paramount to the real-world success of these methods in energy systems contexts, and we hope that our paper will serve to spark further inquiry into this important line of work.

#### 8 ACKNOWLEDGMENTS

This material is based, in part, on work supported by Carnegie Mellon University's College of Engineering Moonshot Award for Autonomous Technologies for Livability and Sustainability (ATLAS). This work was also supported by the U.S. Department of Energy Computational Science Graduate Fellowship (DE-FG02-97ER25308), the Center for Climate and Energy Decision Making through a cooperative agreement between the National Science Foundation and Carnegie Mellon University (SES-00949710), the Computational Sustainability Network, and the Bosch Center for AI. The work of K. Baker is supported by the National Science Foundation CAREER award 2041835.

#### **REFERENCES**

 Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained Policy Optimization. In Proceedings of the 34th International Conference on Machine Learning.

- [2] Akshay Agrawal, Brandon Amos, Shane Barratt, Stephen Boyd, Steven Diamond, and J Zico Kolter. 2019. Differentiable Convex Optimization Layers. In Advances in Neural Information Processing Systems. 9558–9570.
- [3] Anayo K. Akametalu, Shahab Kaynama, Jaime F. Fisac, Melanie Nicole Zeilinger, Jeremy H. Gillula, and Claire J. Tomlin. 2014. Reachability-based safe learning with Gaussian processes. In 53rd IEEE Conference on Decision and Control, CDC 2014.
- [4] Eitan Altman. 1999. Constrained Markov Decision Processes. Vol. 7. CRC Press.
- [5] Brandon Amos and J Zico Kolter. 2017. OptNet: Differentiable Optimization as a Layer in Neural Networks. In Proceedings of the 34th International Conference on Machine Learning. 136–145.
- [6] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. 2019. Deep equilibrium models. arXiv preprint arXiv:1909.01377 (2019).
- [7] Kyri Baker, Andrey Bernstein, Emiliano Dall'Anese, and Changhong Zhao. 2017. Network-cognizant voltage droop control for distribution grids. *IEEE Transactions on Power Systems* 33, 2 (2017), 2098–2108.
- [8] J. Bank and J. Hambrick. 2013. Development of a high resolution, real time, distribution-level metering system and associated visualization modeling, and data analysis functions. (2013). National Renewable Energy Laboratory, Tech. Rep. NREL/TP-5500-56610.
- T.S. Basso. 2014. In IEEE 1547 and 2030 Standards for Distributed Energy Resources Interconnection and Interoperability with the Electricity Grid. National Renewable Energy Laboratory.
- [10] Felix Berkenkamp, Matteo Turchetta, Angela P. Schoellig, and Andreas Krause. 2017. Safe Model-based Reinforcement Learning with Stability Guarantees. In Advances in Neural Information Processing Systems.
- [11] Saverio Bolognani and Florian Dörfler. 2015. Fast power system analysis via implicit linearization of the power flow manifold. In 2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 402–409.
- [12] Ya-Chien Chang, Nima Roohi, and Sicun Gao. 2019. Neural Lyapunov Control. In Advances in Neural Information Processing Systems. 3245–3254.
- [13] Minas Chatzos, Ferdinando Fioretto, Terrence WK Mak, and Pascal Van Hentenryck. 2020. High-Fidelity Machine Learning Approximations of Large-Scale Optimal Power Flow. arXiv preprint arXiv:2006.16356 (2020).
- [14] Bingqing Chen, Zicheng Cai, and Mario Bergés. 2019. Gnu-RL: A precocial reinforcement learning solution for building HVAC control using a Differentiable MPC policy. In Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation. 316–325.
- [15] Bingqing Chen, Jonathan Francis, Marco Pritoni, Soummya Kar, and Mario Berg'es. 2020. COHORT: Coordination of Heterogeneous Thermostatically Controlled Loads for Demand Flexibility. In Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation. 31-40.
- [16] Bingqing Chen, Ming Jin, Zhe Wang, Tianzhen Hong, and Mario Bergés. 2020. Towards Off-policy Evaluation as a Prerequisite for Real-world Reinforcement Learning in Building Control. In Proceedings of the 1st International Workshop on Reinforcement Learning for Energy Management in Buildings & Cities. 52–56.
- [17] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural ordinary differential equations. In Advances in neural information processing systems. 6571–6583.
- [18] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. 2018. End-to-end differentiable physics for learning and control. In Advances in Neural Information Processing Systems. 7178–7189.
- [19] Josip Djolonga and Andreas Krause. 2017. Differentiable Learning of Submodular Models. In Advances in Neural Information Processing Systems. 1013–1023.
- [20] Roel Dobbe, Patricia Hidalgo-Gonzalez, Stavros Karagiannopoulos, Rodrigo Henriquez-Auba, Gabriela Hug, Duncan S Callaway, and Claire J Tomlin. 2020. Learning to control in power systems: Design and analysis guidelines for concrete safety problems. Electric Power Systems Research 189 (2020), 106615.
- [21] Priya L Donti, Melrose Roderick, Mahyar Fazlyab, and J Zico Kolter. 2021. Enforcing robust control guarantees within neural network policies. In *International Conference on Learning Representations*.
- [22] Ján Drgoňa, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, Juraj Oravec, Michael Wetter, Draguna L Vrabie, et al. 2020. All you need to know about model predictive control for buildings. Annual Reviews in Control (2020).
- [23] PO Fanger. 1986. Thermal environment—Human requirements. Environmentalist 6, 4 (1986), 275–278.
- [24] Nicholas EP Fernandez, Srinivas Katipamula, Weimin Wang, YuLong Xie, Mingjie Zhao, and Charles D Corbin. 2017. Impacts of commercial building controls on energy savings and peak load reduction. Technical Report. Pacific Northwest National Lab.(PNNL), Richland, WA (United States).
- [25] Ferdinando Fioretto, Terrence WK Mak, and Pascal Van Hentenryck. 2020. Predicting AC optimal power flows: Combining deep learning and lagrangian dual methods. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34. 630–637.

- [26] Jaime F Fisac, Neil F Lugovoy, Vicenç Rubies-Royo, Shromona Ghosh, and Claire J Tomlin. 2019. Bridging hamilton-jacobi safety analysis and reinforcement learning. In 2019 International Conference on Robotics and Automation (ICRA). IEEE, 8550–8556.
- [27] Nathan Fulton and André Platzer. 2019. Verifiably safe off-model reinforcement learning. In International Conference on Tools and Algorithms for the Construction and Analysis of Systems. Springer, 413–430.
- [28] Javier Garcia and Fernando Fernández. 2015. A comprehensive survey on safe reinforcement learning. Journal of Machine Learning Research 16, 1 (2015), 1437– 1480.
- [29] Mevludin Glavic. 2019. (Deep) Reinforcement learning for electric power system control and related problems: A short review and perspectives. *Annual Reviews* in Control 48 (2019), 22–35.
- [30] Stephen Gould, Richard Hartley, and Dylan Campbell. 2019. Deep declarative networks: A new hope. arXiv preprint arXiv:1909.04866 (2019).
- [31] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. 2019. Hamiltonian neural networks. In Advances in Neural Information Processing Systems. 15379–15389.
- [32] Sarthak Gupta, Vassilis Kekatos, and Ming Jin. 2020. Deep Learning for Reactive Power Control of Smart Inverters under Communication Constraints. In 2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm). IEEE, 1-6.
- [33] Minghao Han, Yuan Tian, Lixian Zhang, Jun Wang, and Wei Pan. 2019. H<sub>∞</sub> Model-free Reinforcement Learning with Robust Stability Guarantee. CoRR (2019).
- [34] Mohammadhosein Hasanbeig, Daniel Kroening, and Alessandro Abate. 2020. Deep reinforcement learning with temporal logics. In *International Conference on Formal Modeling and Analysis of Timed Systems*. Springer, 1–22.
- [35] Nathan Hunt, Nathan Fulton, Sara Magliacane, Nghia Hoang, Subhro Das, and Armando Solar-Lezama. 2020. Verifiably safe exploration for end-to-end reinforcement learning. arXiv preprint arXiv:2007.01223 (2020).
- [36] IEEE. [n.d.]. 37 node distribution test feeder. https://ewh.ieee.org/soc/pes/dsacom/testfeeders/. Online.
- [37] IEEE. 2020. IEEE Standard Conformance Test Procedures for Equipment Interconnecting Distributed Energy Resources with Electric Power Systems and Associated Interfaces. IEEE Std 1547.1-2020 (2020), 1–282. https://doi.org/10.1109/IEEESTD.2020.9097534
- [38] Mana Jalali, Vassilis Kekatos, Nikolaos Gatsis, and Deepjyoti Deka. 2019. Designing reactive power control rules for smart inverters using support vector machines. IEEE Transactions on Smart Grid 11, 2 (2019), 1759–1770.
- [39] Zico Kolter, David Duvenaud, and Matthew Johnson. 2020. Tutorial: Deep Implicit Layers - Neural ODEs, Deep Equilibirum Models, and Beyond. http://implicit-layers-tutorial.org/.
- [40] Steven G Krantz and Harold R Parks. 2012. The implicit function theorem: history, theory, and applications. Springer Science & Business Media.
- [41] Chun Kai Ling, Fei Fang, and J Zico Kolter. 2018. What game are we playing? end-to-end learning in normal and extensive form games. arXiv preprint arXiv:1805.02777 (2018).
- [42] Biao Luo, Huai-Ning Wu, and Tingwen Huang. 2014. Off-Policy Reinforcement Learning for H<sub>∞</sub> Control Design. IEEE Transactions on Cybernetics 45, 1 (2014), 65–76.
- [43] Federico Milano, Florian Dörfler, Gabriela Hug, David J Hill, and Gregor Verbič. 2018. Foundations and challenges of low-inertia systems. In 2018 Power Systems Computation Conference (PSCC). IEEE, 1–25.
- [44] Jun Morimoto and Kenji Doya. 2005. Robust Reinforcement Learning. Neural Computation 17, 2 (2005), 335–359.
- [45] Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. 2018. Optlayer-practical constrained optimization for deep reinforcement learning in the real world. In 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE. 6236–6243.
- [46] Lerrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust Adversarial Reinforcement Learning. In Proceedings of the 34th International Conference on Machine Learning. JMLR. org, 2817–2826.
- [47] Samuel Privara, Jiří Cigler, Zdeněk Váňa, Frauke Oldewurtel, Carina Sagerschnig, and Eva Žáčeková. 2013. Building modeling as a crucial part for building predictive control. Energy and Buildings 56 (2013), 8–22.
- [48] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas. 2011. MATPOWER: Steady-State Operations, Planning and Analysis Tools for Power Systems Research and Education. *IEEE Transactions on Power Systems* 26, 1 (2011), 12–19.
- [49] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. 2019. Tackling Climate Change with Machine Learning. arXiv preprint arXiv:1906.05433 (2019).
- [50] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. 2015. High-dimensional continuous control using generalized advantage estimation. arXiv preprint arXiv:1506.02438 (2015).
- [51] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. arXiv preprint arXiv:1707.06347

- (2017).
- [52] Sanket Shah, Sinha Arunesh, Varakantham Pradeep, Perrault Andrew, and Tambe Milind. 2020. Solving online threat screening games using constrained action space reinforcement learning. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34. 2226–2235.
- [53] N. Stringer, A. Bruce, I. MacGill, N. Haghdadi, P. Kilby, J. Mills, T. Veijalainen, M. Armitage, and N. Wilmot. 2020. Consumer-Led Transition: Australia's World-Leading Distributed Energy Resource Integration Efforts. *IEEE Power and Energy Magazine* 18, 6 (2020), 20–36. https://doi.org/10.1109/MPE.2020.3014720
- [54] Majid Alkaee Taleghan and Thomas G. Dietterich. 2018. Efficient Exploration for Constrained MDPs. In 2018 AAAI Spring Symposia.
- [55] Tijmen Tieleman and Geoffrey Hinton. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural networks for machine learning 4, 2 (2012), 26–31.
- [56] Sebastian Tschiatschek, Aytunc Sahin, and Andreas Krause. 2018. Differentiable Submodular Maximization. In Proceedings of the 27th International Joint Conference on Artificial Intelligence. 2731–2738.
- [57] Matteo Turchetta, Felix Berkenkamp, and Andreas Krause. 2016. Safe Exploration in Finite Markov Decision Processes with Gaussian Processes. In Advances in Neural Information Processing Systems.
- [58] Akifumi Wachi, Yanan Sui, Yisong Yue, and Masahiro Ono. 2018. Safe Exploration and Optimization of Constrained MDPs Using Gaussian Processes. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 32.
- [59] Po-Wei Wang, Priya Donti, Bryan Wilder, and Zico Kolter. 2019. SATNet: Bridging deep learning and logical reasoning using a differentiable satisfiability solver. In

- Proceedings of the 36th International Conference on Machine Learning, 6545-6554.
- [60] Stephen Wilcox and William Marion. 2008. Users manual for TMY3 data sets. National Renewable Energy Laboratory Golden, CO.
- [61] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. 2020. Projection-based constrained policy optimization. arXiv preprint arXiv:2010.03152 (2020).
- [62] Ahmed S Zamzam and Kyri Baker. 2020. Learning optimal solutions for extremely fast AC optimal power flow. In 2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm). IEEE. 1–6.
- [63] Kaiqing Zhang, Bin Hu, and Tamer Basar. 2020. Policy Optimization for H<sub>2</sub> Linear Control with H<sub>∞</sub> Robustness Guarantee: Implicit Regularization and Global Convergence. In Learning for Dynamics and Control. PMLR, 179–190.
- [64] Zhiang Zhang and Khee Poh Lam. 2018. Practical Implementation and Evaluation of Deep Reinforcement Learning Control for a Radiant Heating System. In Proceedings of the 5th Conference on Systems for Built Environments (Shenzen, China) (BuildSys '18). ACM, New York, NY, USA, 148–157.
- [65] Zidong Zhang, Dongxia Zhang, and Robert C Qiu. 2019. Deep reinforcement learning for power system applications: An overview. CSEE Journal of Power and Energy Systems 6, 1 (2019), 213–225.
- [66] Lin Zhao, Wei Zhang, He Hao, and Karanjit Kalsi. 2017. A geometric approach to aggregate flexibility modeling of thermostatically controlled loads. *IEEE Transactions on Power Systems* 32, 6 (2017), 4721–4731.
- [67] Kemin Zhou and John Comstock Doyle. 1998. Essentials of Robust Control. Vol. 104. Prentice hall Upper Saddle River, NJ.