

Winning the CityLearn Challenge: Adaptive Optimization with Evolutionary Search under Trajectory-Based Guidance

Vanshaj Khattar and Ming Jin

Department of Electrical and Computer Engineering, Virginia Tech
vanshajk@vt.edu, jinming@vt.edu

Abstract

Modern power systems will have to face difficult challenges in the years to come: frequent blackouts in urban areas caused by high power demand peaks, grid instability exacerbated by intermittent renewable generation, and global climate change amplified by rising carbon emissions. While current practices are growingly inadequate, the path to widespread adoption of artificial intelligence (AI) methods is hindered by missing aspects of trustworthiness. The CityLearn Challenge is an exemplary opportunity for researchers from multiple disciplines to investigate the potential of AI to tackle these pressing issues in the energy domain, collectively modeled as a reinforcement learning (RL) task. Multiple real-world challenges faced by contemporary RL techniques are embodied in the problem formulation. In this paper, we present a novel method using the solution function of optimization as policies to compute actions for sequential decision-making, while notably adapting the parameters of the optimization model from online observations. Algorithmically, this is achieved by an evolutionary algorithm under a novel trajectory-based guidance scheme. Formally, the global convergence property is established. Our agent ranked first in the latest 2021 CityLearn Challenge, being able to achieve superior performance in almost all metrics while maintaining some key aspects of interpretability.

Introduction

Rapid urbanization in the past decades has led to a substantial increase in energy use that puts stress on the grid assets, while the integration of additional renewable generation and energy storage at the distribution level introduces both opportunities and new challenges (Rolnick et al. 2022). The cornerstone of addressing emerging issues is the deployment of proper control and coordination strategies, which have a potential impact on enhancing energy flexibility and resilience in the face of a surge in climate-induced demand (as already seen in places like California, where rolling blackouts are increasingly frequent during the summer) (Vazquez-Canteli et al. 2020). Current industry practice is heavily based on optimization models, such as energy dispatch and unit commitment, where parameters (e.g., technological and physical

constraints) are fixed throughout the lifecycle; however, such an approach is increasingly confronted by environmental uncertainty, renewable generation stochasticity, and the ever-increasing complexity of the distribution grid (Abedi, Gaudard, and Romerio 2019). On the other hand, there has been a surge in machine learning research, notably RL, because it allows the agent to act without the need to access the true model—a feature of particular interest for large-scale, complex systems, where it is not cost-effective to develop models of such high fidelity. Despite recent progress, real-world RL is still in its infancy (Dulac-Arnold et al. 2021).

Against this backdrop, the CityLearn Challenge aims to spur RL solutions for the control of modern energy systems by providing a set of benchmarks for urban energy management, load shaping, and demand response in a range of climate zones (Vazquez-Canteli et al. 2020). The agent is tasked with exploring and exploiting the best control strategy for energy storage distributed in a community of buildings. Performance is evaluated against standard metrics such as ramping costs, peak demands, and carbon emissions. The CityLearn encapsulates 4 of the 9 real-world RL challenges identified by (Dulac-Arnold et al. 2021), including 1) the ability to learn on live systems from limited samples—there is no training period; 2) dealing with system constraints that should never or rarely be violated—there are strict balancing equations for electricity, heating, and cooling energy; 3) the ability to provide quick action—there is a strict time limit for completing the 4-year evaluation on Google’s Colab; and 4) providing system operators with explainable policies—a necessity to facilitate real-world adoption and deployment.

In this paper, we describe our winning solution for the 2021 CityLearn Challenge based on the idea broadly categorized as *adaptive optimization*. Indeed, optimization (especially convex optimization) has become the de facto standard in industrial systems with profound theoretical foundations, and various formulations for control and planning applications (Boyd, Boyd, and Vandenberghe 2004). Such approaches can easily encode domain-specific constraints (in the form of nonlinear functions and variational inequalities) and can gracefully handle problems with millions of decision variables (Facchinei and Pang 2007). Although well established, optimization models, once built, typically do not adapt to real-world conditions, rendering current approaches rather “rigid.”

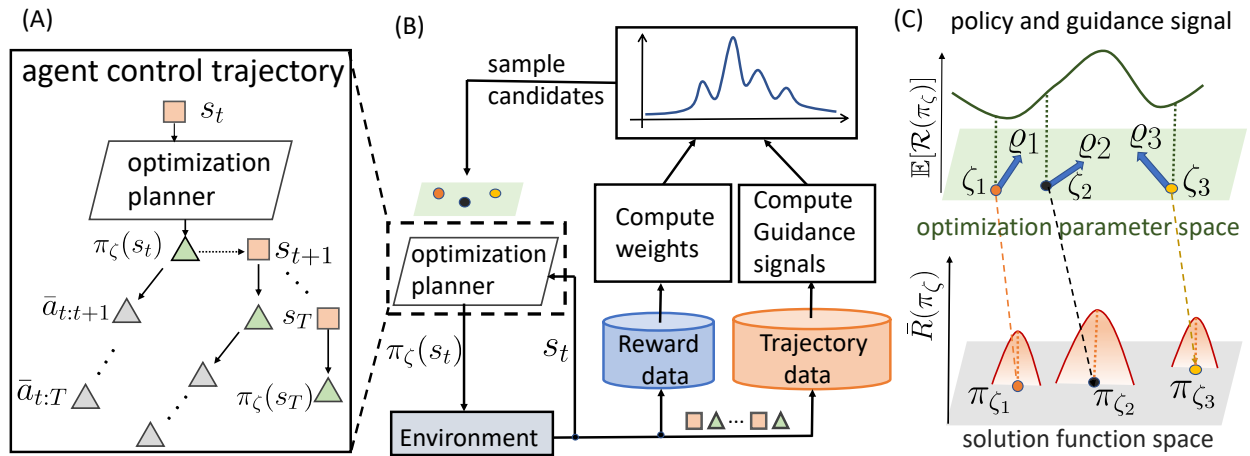


Figure 1: Overview. (A) Agent trajectory of observed states (square) and actual (green triangle)/planned (grey triangle) actions. At each step, the agent solves optimization (2) to plan ahead but only executes the most immediate action; $\bar{a}_{t:t'}$ represents the action planned at time t for time $t' \geq t$. (B) Illustration of the system, including interaction with the environment and adaptation of optimization parameters. Each iteration of k consists of evaluating N_k (three in this case) policies sampled from the candidate distribution. The observed rewards $\{R_j^k\}_{j \in [N_k]}$ and trajectories $\{s_t^{j,k}, a_t^{j,k}\}_{j \in [N_k]}$ are stored in some buffers, which are then used to compute weights $\{\nu_j^k\}_{j \in [N_k]}$ and guidance signals $\{\varrho_j^k\}_{j \in [N_k]}$ to update the distribution as in (5). (C) The two-level structure indicates the correspondence between an optimization parameter ζ and the solution function π_ζ (used as policy). The upper-level objective is the expected episodic rewards (1) and the lower-level objective is the objective function of (2), the extremum of which is the policy action (see (3)). The guidance signal ϱ is computed based on the trajectory of each candidate and applied on top of the original parameter during the update.

To address this fundamental limitation, we exploit that the solution of optimization lies on a manifold implicitly defined by a general equation (Dontchev and Rockafellar 2009). The crux of our idea is to shape this manifold by adapting the parameters of the optimization model while extracting insights from trajectory data to design guidance signals (see Fig. 1 for a detailed illustration). Key differences between our method and well-established optimization techniques (e.g., stochastic optimization (Powell 2020), bi-level optimization (Dempe and Zemkoho 2020)) include:

- 1) we only allow access to the environment through interactive samples (reward, states, etc.) but not the true dynamics or reward function;
- 2) we make full use of the control trajectory of a Markov decision process (MDP) to obtain a guidance signal.

While 1) is similar to the RL setup, 2) can be viewed as an *augmentation of zeroth-order search methods with insights extracted from control data*, which, to the best of our knowledge, is the *first of its kind*. In general, zeroth-order algorithms, such as simultaneous perturbation (Spall 2005; Mania, Guy, and Recht 2018), evolutionary algorithms (Salimans et al. 2017; Zhou, Yu, and Qian 2019), and Bayesian optimization (Snoek, Larochelle, and Adams 2012; Frazier 2018) are natural candidates for RL and easy to implement, but can potentially suffer scalability problems (Ghadimi and Lan 2013). Nevertheless, the parameters of an optimization model (i.e., variables to be learned) usually have clear interpretations. Thus, we design a mechanism to leverage domain knowledge for the design of appropriate guidance during the search for these parameters; such guidance can be specified

as a general function of trajectory data (including observed states and actions of an MDP). The method works well in an online environment without an extensive training period, which is particularly advantageous in a real-world setting where an offline environment for model training is usually not available. According to independent evaluations, the proposed method achieved the highest performance in the recent 2021 CityLearn Challenge. To demonstrate effectiveness against existing techniques, we further validate the method by comparing it with a range of baselines. Key contributions are as follows:

- A framework of adaptive optimization for online control, winning the **1st place in the 2021 CityLearn Challenge**
- A novel evolutionary search (ES) algorithm with a guidance function based on state-action-trajectory data
- Theoretical analysis of asymptotic convergence to global optima with noisy function evaluations
- Empirical comparison against a range of baselines

Related Work

Optimal control and stochastic optimal control are well-known approaches to sequential decision-making problems (Bertsekas 2019). Convex optimization is another avenue (Agrawal et al. 2020). Most existing works assume a known dynamic model of the system. Various large-scale stochastic programs have been proposed in the literature to deal with future uncertainty (Prékopa 2013; Powell 2020). The major drawback is the computational burden of rapidly expanding scenario trees in multi-stage stochastic programming.

Our method relieves computation by using plug-in estimators, a.k.a., deterministic approximation of future uncertainty within convex optimization.

Recently, RL has gained popularity in various domains (Chen et al. 2022; Haydari and Yilmaz 2020; Nian, Liu, and Huang 2020). To contextualize the present approach, we make a few remarks about the relation to model-based RL (MBRL). In particular *implicit* MBRL, where the entire procedure (e.g., model learning and planning) is optimized for optimal policy computation (Moerland, Broekens, and Jonker 2020). However, unlike existing works (e.g., (Tamar et al. 2016; Karkus, Hsu, and Lee 2017; Racanière et al. 2017; Guez et al. 2018; Schrittwieser et al. 2020) that build a model based on (recurrent/convolutional) neural networks (NNs) with primary restrictions to discrete state and action space, our method learns how to plan by solving optimization and adapting its parameters; hence, it is amenable to a wide range of applications with continuous states and actions. The present work is closely related to (Ghadimi, Perkins, and Powell 2020; Agrawal et al. 2020), which also use convex optimization as a policy class to handle uncertainty. In particular, convex optimization control policies are learned in (Agrawal et al. 2020) with implicit differentiation (Agrawal et al. 2019). We extend their method to the RL setting and propose a novel ES algorithm for learning.

Differentiated from existing ES-based strategies (Szita and Lörincz 2006; Salimans et al. 2017; Khadka and Tumer 2018; Gangwani and Peng 2018; Conti et al. 2018) or zeroth-order optimization (Snoek, Larochelle, and Adams 2012; Liu et al. 2020), our ES is augmented with a guidance function that depends on the data collected as the policy interacts with the environment; hence, it can be viewed as a type of *MDP-augmented ES*. The guidance mechanism is also flexible enough to allow the effective incorporation of domain knowledge, as demonstrated in CityLearn. We further provide theoretical justification for this rather complex scheme.

Preliminaries

Problem Setup

Consider an MDP $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r)$, where \mathcal{S} is the (possibly infinite) state space, \mathcal{A} is the set of actions, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{M}(\mathcal{S})$ is the transition probability kernel with $\mathcal{M}(\mathcal{S})$ denoting the set of all probability measures over \mathcal{S} , and $r(s, a)$ gives the corresponding immediate reward (can be time-dependent). The goal in episodic RL is to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes cumulative rewards over a finite time horizon:

$$\max_{\pi \in \Pi} \mathbb{E}[R(\pi)], \quad (1)$$

where $R(\pi) := \sum_{t=0}^T r_t(s_t, \pi(s_t))$ is the episodic reward, with $s_t \in \mathcal{S} \subseteq \mathbb{R}^{n_s}$ denoting the state at time t and T as the horizon. The expectation is taken over initial state distributions and transition dynamics (under deterministic policy π). We require access to a random sample $R(\pi)$ of the episodic reward as well as trajectory data $\{(s_t, a_t)\}_{t \in [T]}$ to be used to later compute the guidance signal. Here, we use the shorthand $[T] = \{1, \dots, T\}$.

Canonical Approaches and Solution Functions

The proposed method is based on canonical stochastic programming methods (Powell 2020). For example, in multi-stage stochastic programming (Pflug and Pichler 2014), the action at state s_t is computed as

$$\arg \max_{a_t \in \mathcal{A}} \left(\tilde{r}_t(s_t, a_t) + \max_{\{a_{t'}\}_{t'=t+1}^T} \tilde{\mathbb{E}} \left[\sum_{t'=t+1}^T \tilde{r}_{t'}(s_{t'}, a_{t'}) \middle| s_t, a_t \right] \right),$$

where $\tilde{r}_t : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the surrogate reward function and $\tilde{\mathbb{E}}[\cdot]$ is the surrogate expectation operator (e.g., model-based scenario trees) designed to approximate the true environment. The above formulation can also be seen as finding the solution to a Bellman equation in dynamic programming. The main limitations, nevertheless, are the high computational cost of evaluating the expectation operator and the potential model mismatch due to approximations.

A simpler yet more practically appealing method, widely adopted in today's industries, is to use deterministic approximations of the future and capture the dependence of future states on prior decisions through constraints as part of the lookahead model (Powell 2020): $\pi_\zeta(s_t)$ is given by

$$\begin{aligned} \arg \max_{\bar{a}_t \in \mathcal{A}} \quad & \max_{\{\bar{s}_{t'}, \bar{a}_{t'}\}_{t'=t+1}^T} \bar{R}(s_t, \bar{a}_t, \{\bar{s}_{t'}, \bar{a}_{t'}\}_{t'=t+1}^T; \zeta) \\ \text{s. t.} \quad & g_i(s_t, \bar{a}_t, \{\bar{s}_{t'}, \bar{a}_{t'}\}_{t'=t+1}^T; \zeta) \leq 0 \quad ; \quad \forall i \in \mathcal{I} \\ & h_i(s_t, \bar{a}_t, \{\bar{s}_{t'}, \bar{a}_{t'}\}_{t'=t+1}^T; \zeta) = 0 \quad ; \quad \forall i \in \mathcal{E}, \end{aligned} \quad (2)$$

with the objective $\bar{R}(s_t, \bar{a}_t, \{\bar{s}_{t'}, \bar{a}_{t'}\}_{t'=t+1}^T; \zeta)$ defined as

$$\bar{r}_t(s_t, \bar{a}_t; \zeta) + \sum_{t'=t+1}^T \bar{r}_{t'}(\bar{s}_{t'}, \bar{a}_{t'}; \zeta),$$

where $\{\bar{s}_{t'}\}_{t'=t+1}^T$ and $\{\bar{a}_{t'}\}_{t'=t+1}^T$ are optimization variables corresponding to the planned states and actions, \bar{r}_t is surrogate reward, and the feasible set is defined by $\{g_i\}_{i \in \mathcal{I}}$ and $\{h_i\}_{i \in \mathcal{E}}$. We denote the parameters of the objective function and the constraints collectively by $\zeta \in \mathcal{Z} \subset \mathbb{R}^d$. The dependencies of future states on current and *planned* states/actions are encoded as constraints in (2) as part of the lookahead model. Many examples can be found in, e.g., (Borrelli, Bemporad, and Morari 2017). We remark that some optimization parameters may be provided by predictors based on the current state s_t , the parameter of which is also collected by ζ in this case.

The policy $\pi_\zeta(s_t)$, a.k.a., the solution function (Dontchev and Rockafellar 2009), provides the action at the current state s_t as the optimal solution to (2). Since this function is generally set-valued (Dontchev and Rockafellar 2009), we make the following assumption.

Assumption 1. For each $\zeta \in \mathcal{Z}$ and $s_t \in \mathcal{S}$: a) the objective function in (2) is continuous, strictly convex, g_i is continuous and convex for each $i \in \mathcal{I}$, and h_i is affine for each $i \in \mathcal{E}$; b) the feasible set of (2) is closed, absolutely bounded, and has a nonempty interior.

The above assumption can be satisfied by imposing proper conditions on the design of the surrogate model, i.e., objective and constraints in (2). Note that in our approach, we make

no convexity assumption about the true dynamics or rewards of the environment, which can be seen as a blackbox. The convexity condition is only stipulated for the surrogate model for computational efficiency. Our goal is simply to learn the parameters of the optimization model in order to have a good decision-making capability.¹ An immediate consequence of the above assumption is that the solution to (2) is unique; furthermore, it implies continuity with respect to parameters.

Lemma 1. *The solution function $\pi_\zeta(s_t)$ defined in (2) is continuous with respect to parameter ζ for each $s_t \in \mathcal{S}$.*

The proof is a direct application of the Berge maximum theorem (Berge 1997). To conclude this section, let us make some comments on the construction of the surrogate model in (2). By analogy to reward design (Prakash et al. 2020), the objective function should be chosen to promote desirable behaviors. The set of constraints introduces inductive bias on the transition dynamics of the environment. It is beneficial, though oftentimes unlikely and non-essential, that the surrogate model matches the functional forms of true reward or dynamics, an idea shared in model-based RL (Moerland, Broekens, and Jonker 2020). It is, nevertheless, desirable to ensure the computational efficiency of (2) to provide actions quickly—hence the choice of convex programs.

Policy Adaptation With ES Under Guidance

The potential mismatch between the surrogate model and the real environment and errors due to predictions may adversely affect the decision quality of (2). Thereby, we aim to adapt the parameters of the surrogate model to shape the solution function. The task of finding the optimal parameter within the set of solution functions $\Pi = \{\pi_\zeta : \zeta \in \mathcal{Z}\}$ can be compactly written as:

$$\Upsilon := \arg \max_{\zeta \in \mathcal{Z}} \mathbb{E} \left[\sum_{t=0}^T r_t(s_t, \pi_\zeta(s_t)) \right], \quad (3)$$

where Υ is the set of global optima for policy parameters. Note that ζ is not part of the true reward (which remains unknown to the agent) but only the parameters of the surrogate model that implicitly defines the policy in (2). Since $\pi_\zeta(s_t)$ is given by an optimization (2), (3) can also be viewed as a bi-level problem (c.f., (Dempe and Zemkoho 2020)): the outer level aims to learn parameters to maximize rewards, while the inner level defines policy action as a solution to (2). The key challenge in solving (3) as a bi-level problem, however, is that the outer level objective is not analytically revealed (thus prohibiting any direct differentiation approach) and can be nonconvex with respect to ζ .

Guided Evolutionary Search

This section discusses the proposed ES algorithm (Algorithm 1), inspired by the method of generations (Zhitljavsky

¹Perhaps surprisingly, despite the fact that (2) is convex, the policy (given by the solution function) can be highly nonconvex with high representational capacity. In particular, the contemporary work by (Jin et al. 2023) shows a “universal approximation” property of the solution functions of linear programs (LPs).

Algorithm 1: Evolutionary search under guidance

Input: Hyperparameters $\{N_k\}$, uniform distribution μ , initial point z_0

- 1: Initialize $P_1(d\zeta) \sim \exp(\|\zeta - z_0\|)\mu(d\zeta)$
- 2: **for** $k = 1, 2, \dots$ **do**
- 3: Sample N_k candidates from the distribution p_k :
 $\zeta_1^k, \zeta_2^k, \dots, \zeta_{N_k}^k \stackrel{iid}{\sim} p_k$
- 4: **for** $j = 1, \dots, N_k$ **do**
- 5: Deploy policy $\pi_{\zeta_j^k}$ for one episode and observe an episodic reward $R_j^k \leftarrow R(\pi_{\zeta_j^k})$
- 6: Compute the guidance signal ϱ_j^k by (4)
- 7: **end for**
- 8: Update the distribution p_{k+1} for the next iteration according to (5).
- 9: **end for**

2012). At each iteration k , the algorithm randomly samples a set of N_k parameter candidates, $\zeta_1^k, \dots, \zeta_{N_k}^k \stackrel{iid}{\sim} p_k$. For each candidate $j \in [N_k]$, we evaluate the corresponding policy in the environment and observe an episodic reward $R_j^k \sim \mathcal{R}(\pi_{\zeta_j^k})$, where $\mathcal{R}(\pi_{\zeta_j^k})$ denotes the distribution of episodic reward for policy $\pi_{\zeta_j^k}$, as well as all the state and action pairs $\{s_t^{j,k}, a_t^{j,k}\}_{t \in [T]}$ in the past episode. Based on trajectory data, we compute a guidance signal

$$\varrho_j^k = \phi(\{s_t^{j,k}, a_t^{j,k}\}_{t \in [T]}), \quad (4)$$

where $\phi : (\mathcal{S} \times \mathcal{A})^T \rightarrow \mathcal{Z}'$ is a function that may have complicated dependence on past states and actions with \mathcal{Z}' as the range. The design of such a guidance function is often based on domain knowledge (to be discussed later in CityLearn). Then, we update the distribution for the next iteration as

$$p_{k+1}(d\zeta) = \sum_{j=1}^{N_k} \nu_j^k Q_k(\zeta_j^k, \varrho_j^k, d\zeta), \quad (5)$$

where

$$\nu_j^k = \frac{\exp(R_j^k)}{\sum_{j=1}^{N_k} \exp(R_j^k)} \quad (6)$$

are the weights obtained by taking the softmax over candidate rewards. The probability measure $Q_k(\zeta_j^k, \varrho_j^k, d\zeta)$ is the transition probability given candidate ζ_j^k and guidance signal ϱ_j^k . Hence, $p_{k+1}(d\zeta)$ is a mixture of distributions weighted by observed rewards in the current iteration k , which can be sampled by the standard superposition method: at first the index j is sampled from the discrete distribution $\{\nu_j^k\}$, followed by sampling from $Q_k(\zeta_j^k, \varrho_j^k, d\zeta)$. For example, in our algorithm for CityLearn,

$$Q_k(z, \varrho, d\zeta) \sim \exp(\|\zeta - z - \alpha_k \varrho\|/\iota_k)\mu(d\zeta), \quad (7)$$

where $\|\cdot\|$ is the Euclidean norm, $\mu(d\zeta)$ is a uniform measure over \mathcal{Z} , and $\iota_k > 0$ and $\alpha_k \geq 0$ are such that their sum over time is bounded. Other candidates are possible and can still ensure convergence to global optimal, as long as certain

conditions are met; intuitively, we require that the span of Q_k decreases over time but not so rapidly that it fails to reach a global optimum. Note that to simplify the presentation, in the above algorithm, we assume that each candidate policy is evaluated only on one episode; extending this to the case of multiple episodes is straightforward (e.g., we would instead take the average of the evaluations among the episodes in the computation of weights (6)).

Theoretical Analysis

We now analyze the convergence property of the sequence generated by Algorithm 1. The following notations are used: $f(\zeta) = \mathbb{E}[R(\pi_\zeta)]$ is the expected episodic reward of policy π_ζ , $\Upsilon = \arg \max_{\zeta \in \mathcal{Z}} f(\zeta)$ is the set of global maximizers (may not be unique), $f^* = \max_{\zeta \in \mathcal{Z}} f(\zeta)$ is the global maximum, and $\lambda(d\zeta)$ is some measure over Υ ; $\mathbb{B}(\zeta, \epsilon) = \{\zeta' \in \mathcal{Z} : \|\zeta' - \zeta\| \leq \epsilon\}$ is a ball centered at ζ with radius ϵ , $\mathbb{B}^*(\epsilon) = \{\zeta \in \mathcal{Z} : \min_{\zeta' \in \Upsilon} \|\zeta' - \zeta\| \leq \epsilon\}$ is a set of points that are ϵ away from the optimal solution set Υ ; $\delta_\zeta(dz)$ is the probability measure concentrated at the point ζ . We use \Rightarrow to denote the weak convergence of measures. We can consider $\|\cdot\|$ as any norm (e.g., Euclidean norm).

The measures $p_{k+1}(d\zeta)$, $k \in \mathbb{N}$ defined in (5) are distributions of random points ζ_j^{k+1} , for any $j \in [N_{k+1}]$, conditional on the results of preceding evaluations of $\{R_j^k\}_{j \in [N_k]}$ and realizations of $\{\zeta_j^k, \varrho_j^k, \xi_j^k\}_{j \in [N_k]}$. Let $P_k(d\zeta_1, \dots, d\zeta_{N_k})$ represent their unconditional joint distributions at iteration k , and

$$\tilde{P}_k(d\zeta) = \int_{\mathcal{Z}^{N_{k-1}}} P_k(d\zeta, dz_2, \dots, dz_{N_k})$$

is the unconditional marginal distribution (note that we introduce z for ζ as the need arises in integration).

The formalism of the guidance signal requires some basics of random process and measure theory (details can be found in the appendix). Essentially, the guidance signal $\varrho_j^k \in \mathcal{Z}'$ is a random variable (adapted to the σ -algebra generated by the trajectory within an episode) with probability measure $M_k(\zeta_j^k, d\varrho)$. Note that $M_k(\zeta_j^k, d\varrho)$ is dependent on ζ_j^k because the stochastic process that generates the trajectory depends on policy $\pi_{\zeta_j^k}$, but $M_k(\zeta_j^k, d\varrho)$ is conditionally independent of all other candidates $\{\zeta_{j'}^k\}_{j' \neq j}$. For analysis, we make the following assumptions.

Assumption 2. *The followings hold:*

- (a) $R_j^k = f(\zeta_j^k) + \xi_j^k$, where $\xi_j^k \stackrel{iid}{\sim} F_k(d\xi)$ for any $k \in \mathbb{N}$ are independent with distribution $F_k(d\xi)$ bounded on a finite interval $[-c_\xi, c_\xi]$ and $\mathbb{E} \exp(\xi_j^k) = 1$;
- (b) $|f(\zeta)| \leq c_f$ for all $\zeta \in \mathcal{Z}$ and \mathcal{Z} is compact;
- (c) there exists $\epsilon > 0$ such that f is continuous on $\mathbb{B}^*(\epsilon)$;
- (d) $Q_k(z, \varrho, d\zeta) = q_k(z, \varrho, \zeta) \mu(d\zeta)$, with $\sup_{z, \varrho, \zeta \in \mathcal{Z}} q_k(z, \varrho, \zeta) \leq L_k < \infty$ for all $k \in \mathbb{N}$, where μ is a probability measure on \mathcal{Z} such that $\mu(\mathbb{B}^*(\epsilon)) > 0$ for any $\epsilon > 0$; for any $z \in \mathcal{Z}$, the sequence of probability measures $Q_k(z, \varrho, d\zeta)$ weakly converges to $\delta_z(d\zeta)$;
- (e) $\{N_k\}$ is a sequence of natural numbers $N_k \in \mathbb{N}$ such that $N_k \rightarrow \infty$ for $k \rightarrow \infty$;

- (f) $\sup_{z, \varrho} M_k(z, d\varrho) < \infty$ for all $k \in \mathbb{N}$;
- (g) $\tilde{P}_1(\mathbb{B}(\zeta, \epsilon)) > 0$ for all $\epsilon > 0$ and $\zeta \in \mathcal{Z}$;
- (h) for any $\epsilon > 0$, there are $\delta > 0$ and a natural \bar{k} such that $\tilde{P}_k(\mathbb{B}^*(\epsilon)) \geq \delta$ for all $k \geq \bar{k}$.

Let us comment on the assumptions above. Condition (a) requires that the evaluation noise be independent and bounded; the expectation requirement can be satisfied for truncated log-normal distributions (Thompson 1950). The iid requirement can be relaxed to mixing processes at the cost of more complex analysis (Doukhan 2012); the boundedness condition, on the other hand, seems necessary to keep iterates in the vicinity of global maximum if they are already there. Condition (b) is non-restrictive for practical problems. Condition (c) is natural since π_ζ is continuous by Lemma 1, and can be satisfied if the true reward functions r_t are continuous for all $t \in [T]$. Assumptions (d), (e), (f), (g), and (h) formulate necessary requirements on the parameters of the algorithm. Intuitively, conditions (d), (e), and (f) stipulate that the search becomes more “focused” over time in order to concentrate on the global optima; however, conditions (g) and (h) indicates that the decrease of span cannot be too fast in order not to miss the global optima. Condition (f) on the guidance function ϕ can be met by proper smoothing if needed. Condition (e) can be relaxed to $N_k = N$ for some finite integer N for all $k \in \mathbb{N}$, but the convergence will only be towards the vicinity of Λ due to the finite sample effect (see Lemma 4 in the appendix, which states the rate to be on the order $N^{-1/2}$). Unlike (c), (d), (e), (f), and (g), condition (h) is not constructive; hence, we provide some verifiable conditions sufficient for (h) to hold in Corollary 1 (also see Corollary 2 in the appendix). Next, we analyze the update rule of Algorithm 1.

Lemma 2. *The probability distribution $P_{k+1}(d\zeta_1, \dots, d\zeta_{N_{k+1}})$ can be written in terms of the distribution $P_k(d\zeta_1, \dots, d\zeta_{N_k})$ as:*

$$\int_{\Omega^{N_k}} \chi_k(d\omega_{N_k}) \prod_{j=1}^{N_{k+1}} \left\{ \beta(\omega_{N_k}) \sum_{i=1}^{N_k} \Lambda(z_i, \varrho_i, \xi_i, d\zeta_j) \right\}, \quad (8)$$

where $\Omega = \mathcal{Z} \times \mathcal{Z}' \times [-c_\xi, c_\xi]$,

$$\omega_{N_k} = \{z_1, \dots, z_{N_k}, \varrho_1, \dots, \varrho_{N_k}, \xi_1, \dots, \xi_{N_k}\} \in \Omega^{N_k},$$

$$\chi_k(d\omega_{N_k}) = P_k(dz_1, \dots, dz_{N_k}) \prod_{j=1}^{N_k} F_k(d\xi_j) M_k(z_j, d\varrho_j),$$

$$\beta(\omega_{N_k}) = \frac{1}{\sum_{j=1}^{N_k} \exp(f(z_j) + \xi_j)}, \quad \text{and}$$

$$\Lambda(z, \varrho, \xi, d\zeta) = \exp(f(z) + \xi) Q_k(z, \varrho, d\zeta).$$

The proof is immediate by recognizing that the bracket term in (8) is the conditional distribution $p_{k+1}(d\zeta_j)$ defined in (5) and the integration is over the distribution from the preceding iteration. We take the product over N_{k+1} candidates since they are drawn iid from p_{k+1} .

Now, we provide the main result on the convergence of $\tilde{P}_k(d\zeta)$ to some distribution $\lambda(d\zeta)$ over the global optima.

Theorem 1. *Suppose that Assumption 2 holds true, and let $\{\tilde{P}_k\}$ be the sequence of unconditional marginal distributions determined by Algorithm 1. Then, the distribution sequence weakly converges to some measure λ over the optimal set, i.e., $\tilde{P}_k \Rightarrow \lambda$ as $k \rightarrow \infty$.*

The key stage of proof is to show that there exists a subsequence in $\{\tilde{P}_k\}$ that weakly converges to

$$\vartheta_m(d\zeta) = \frac{\exp(mf(\zeta))\kappa(d\zeta)}{\int \exp(mf(z))\kappa(dz)}$$

for some measure κ , where m is the index of the subsequence. The above distribution is effectively a softmax function over function values and converges to the extrema as $m \rightarrow \infty$.

All the conditions in Assumption 2 are natural with the exception of (h), which requires some further justification. In the following, we present a sufficient condition for (h) with a proper design of $Q_k(z, \varrho, d\zeta)$, which applies to the case of noisy function evaluations; see the appendix for another example in the case of noiseless function evaluations.

Corollary 1. *Under Assumption 2 (except (h)), and let the transition probability $Q_k(z, \varrho, d\zeta)$ be*

$$Q_k(z, \varrho, d\zeta) = c_k(z, \varrho)\psi((\zeta - z - \alpha_k\varrho)/\iota_k)\mu(d\zeta), \quad (9)$$

where $c_k(z, \varrho) = (\int \psi((\zeta - z - \alpha_k\varrho)/\iota_k)\mu(d\zeta))^{-1}$ is the normalization term, ψ is a continuous symmetrical finite density, and

$$\iota_k > 0, \quad \sum_{k=1}^{\infty} \iota_k < \infty, \quad \alpha_k \geq 0, \quad \sum_{k=1}^{\infty} \alpha_k < \infty.$$

Then, there exists a sequence of natural numbers $\{N_k\}$ such that $\{\tilde{P}_k\}$ weakly converges to λ .

Our analysis accounts for the effect of trajectory-based guidance, which is a novel contribution to the ES literature. By examining the proof, we can relax the condition that $\sum_{k=1}^{\infty} \alpha_k < \infty$, i.e., continue applying the guidance without the need to diminish its impact in the long run, as long as the guidance signal “approximately” points to the global optima in the proximity (see the appendix for exact conditions). However, such guidance can be difficult to design or even verify in practice; thus, it is still advisable to relinquish human knowledge and let data drive the decision, eventually.

Results from the CityLearn Challenge

Challenge overview. The competition has an online setup with a simulation period of 1 or 4 years, where agents exploit the best policies to optimize the coordination strategy. The goal of each agent is to minimize environmental costs, such as ramping costs, peak demands, 1-load factor, and carbon emissions. The state space contains information such as daylight hours, outdoor temperature, storage device state of charge (SOC), net electricity consumption of the building, carbon intensity of the power grid, among a total of 30 continuous states. The agent is allowed to control the charging/discharging actions of storage devices for domestic hot water (DHW), chilled water, and electricity (i.e., 3 continuous actions per building). The environment is seen as a blackbox to the agent

	ZO-iRL(ours)	ICD-CA	IDLab
Total score	0.944	1.070	1.070
Coord. score	0.915	1.107	1.098

Table 1: Total and coordination scores of the top 3 teams in 2021 CityLearn Challenge.

as a standard RL setup, where the transition dynamics depend on the responses of various devices (e.g., air-to-water heat pumps, electric heaters) as well as the energy loads of buildings, which include space cooling, dehumidification, DHW demand, and solar generation.

Evaluation. The submission of each team is evaluated on a set of metrics, including: (1) ramping: $\sum |e_t - e_{t-1}|$, where e is the net electricity consumption at each time step; (2) 1-load factor: average net electricity load divided by maximum electricity load; (3) average daily peak demand; (4) maximum peak electricity demand; (5) total electricity consumed; (6) carbon emissions. The competition evaluates performance by computing the ratio of costs with respect to a rule-based controller (RBC)—lower ratios indicate better performances.² The average of the above metrics for the full simulated period is the *total score*, while the average of the metrics (1)-(4) is *coordination score*. The performance of the top 4 teams is listed in Table 1. Refer to (Vazquez-Canteli et al. 2020) for more details on the contest.

ZO-iRL: zeroth-order implicit RL.³ As our method is designed for single-agent episodic RL, we first reduce the original task that consists of a single period of 1 or 4 years into episodes of 24 hours. We use the per-step reward $-\max(0, e_t)$ ³ as recommended by (Vazquez-Canteli et al. 2020), where e_t is the net electricity consumption (or generation if $e_t < 0$). This reward favours consumption patterns that are smoothly averaged without demand peaks, aligned with multiple metrics used in the evaluation, such as the 1-load factor and peak electricity demand. Another reduction is from multi-agent RL to single-agent RL, where each building’s policy is updated independently, reducing the problem to decentralized control with additive rewards; such a reduction is computationally efficient for large-scale problems (De Nijs et al. 2021). We omit the notational dependence on candidate j and iteration k when presenting the method.

Optimization planner. We instantiate the optimization in (2) as follows. The planned states \bar{s}_t consist of state variables such as net electricity consumption and SOC of storage devices; the action $\bar{a}_t \in \mathcal{A}$ is the action of the MDP; the surrogate reward

$$\bar{r}_t(\bar{s}_t; \zeta) = -|e_t - e_{t-1}| - \theta_t e_t$$

is a combination of the negated ramping cost and the “virtual” electricity cost, where $\zeta = \{\theta_t \in [0, 5]\}_{t \in [24]}$ can

²Note that the RBC controller is ubiquitous in traditional building control systems and is a simple form of “take action a_h in hour h ,” where a_h is a constant independent of current states except for the hour of the day ($h \in [24]$).

³We name our method ZO-iRL because the policy is implicitly determined by solving an optimization problem and the learning algorithm is zeroth-order in an RL setting.

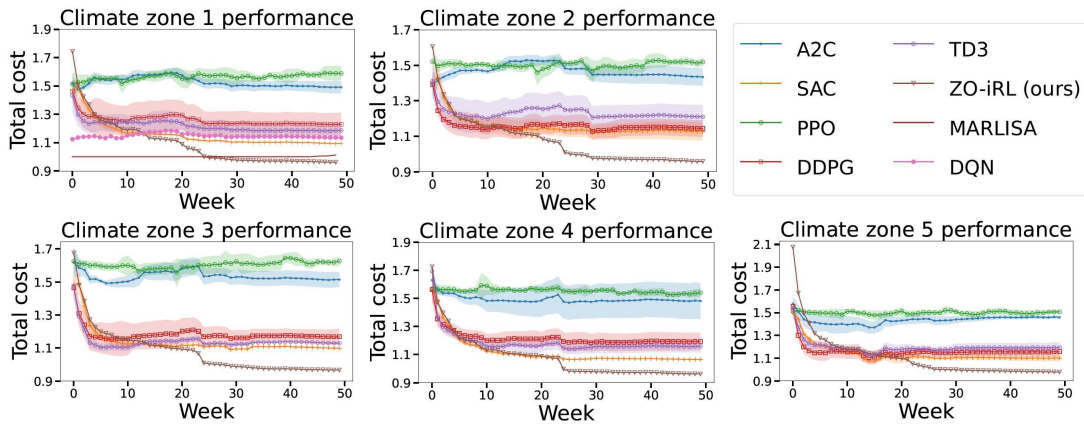


Figure 2: Learning curves of ZO-iRL and baselines. We perform 10 runs on each baseline to obtain performance plots with standard deviations for Climate Zones 1 to 5.

	SAC	A2C	PPO	TD3	MARLISA	ZO-iRL
ramping cost	1.145 (0.015)	1.189 (0.002)	1.638 (0.005)	1.178 (0.026)	<u>1.022</u> (0.010)	0.781 (0.005)
1-load factor	1.158 (0.002)	1.146 (0.002)	1.168 (0.003)	1.142 (0.004)	<u>1.026</u> (0.006)	1.010 (0.008)
avg. daily peak	1.180 (0.002)	1.184 (0.008)	1.242 (0.001)	1.193 (0.006)	<u>1.015</u> (0.001)	0.996 (0.001)
peak demand	1.077 (0.008)	1.088 (0.007)	1.132 (0.006)	1.098 (0.013)	<u>1.000</u> (6e-5)	0.962 (0.005)
net electric. peak	0.995 (0.001)	0.994 (0.002)	1.003 (7e-5)	0.997 (7e-4)	1.000 (5e-4)	1.006 (2e-4)
carbon emissions	1.000 (0.001)	<u>1.000</u> (0.002)	1.009 (7e-5)	1.004 (7e-4)	1.001 (5e-4)	1.007 (2e-4)
total score	1.092 (0.003)	1.101 (0.003)	1.199 (0.001)	1.102 (0.006)	<u>1.011</u> (0.002)	0.962 (0.001)

Table 2: Comparison with baselines: SAC (Kathirgamanathan et al. 2020) and MARLISA (Vazquez-Canteli, Henze, and Nagy 2020) have been officially implemented for CityLearn; other baselines are implemented by (Raffin et al. 2021). The reported values are the average and standard deviation (in brackets) across 10 independent runs on Climate Zone 1 data.

be viewed as virtual electricity prices to be learned to encourage desirable consumption patterns (e.g., load flattening and smoothing). Intuitively, a higher value of θ_t discourages planned electricity consumption in the corresponding hour t .

The inequalities are grouped into technological constraints (e.g., maximum/minimum cooling power) and constraints on states and actions. The equalities are grouped into physics accounting for energy balances (i.e., consumption equal to supply) and technology (e.g., SOC update rules). Further details are provided in the appendix. Note that to set up the optimization (2), we also need to provide predictions of energy demands and solar generation. For simplicity, our predictors are based on a simple averaging scheme that takes the average in the corresponding hours of the last 2 weeks of data; thus, there are no specific parameters to learn.

Transition and guidance. We use (7) as the transition probability, with variance $\iota_k = 0.4/k^2$ that is initialized to 0.4 and decreases by k^2 in each iteration. The guidance signal ϱ is computed as follows. By the end of each episode, we examine the net electricity usage in the past 24 hours, e_t for $t \in [24]$ and find the top 2 hours with the most electricity usage, denoted by t_1 and t_2 . Then, the guidance signal ϱ_t is 0.02 if $t \in \{t_1, t_2\}$ and $-0.04/22$ otherwise. Note that we have centered the signal ($\sum_t \varrho_t = 0$) by assigning negative values for hours other than peaks. We choose $\alpha_k = 1$ for all k over the entire 4-year period, as there is no training phase

in the CityLearn Challenge and we prefer to adapt quickly during the test phase; this is not a violation of our theory, as we can choose to diminish α_k after a while to still satisfy the condition $\sum_{k=1}^{\infty} \alpha_k < \infty$.

Results. For baselines, we use the implementation of (Raffin et al. 2021) with the default ADAM optimizer, where the policy is an NN architecture with tanh activation and two layers of 256 units each. From Table 2, we see that ZO-iRL has achieved the lowest cost ratios (i.e., best scores) of all, which is consistent with the official result of the competition (Table 1). In particular, as shown in Fig. 2, ZO-iRL is able to find a good policy in the first few months, while baselines seem to struggle; we speculate that more samples would eventually improve the performance of baselines, and all methods may benefit from schemes to handle the potentially nonstationary environment due to seasonal patterns.

Conclusion and Future Directions

We proposed a novel adaptive optimization framework that adapts the optimization parameters using an evolutionary search algorithm that incorporates insights from trajectory data as guidance. The method outperforms several baselines and ranked first in the 2021 CityLearn Challenge. The proposed work could be extended to methods like Bayesian optimization or first-order methods such as actor-critic.

Acknowledgments

We would like to thank all the ZORL/ROLEVT team members at the CityLearn Challenge 2021: Qasim Wani, Zhiyao Zhang, and Mingyu Kim. We would also like to thank the organizers of the CityLearn Challenge: Dr. José R. Vázquez-Canteli, Dr. Zoltan Nagy, Dr. Gregor Henze, and Sourav Dey for organizing this competition and inspiring novel RL-based approaches to the contemporary challenges facing the power and energy sector. This project is funded by NSF, the Commonwealth Cyber Initiative (CCI), C3.ai Digital Transformation Institute, and the U.S. Department of Energy.

References

- Abedi, A.; Gaudard, L.; and Romerio, F. 2019. Review of major approaches to analyze vulnerability in power system. *Reliability Engineering & System Safety*, 183: 153–172.
- Agrawal, A.; Amos, B.; Barratt, S.; Boyd, S.; Diamond, S.; and Kolter, J. Z. 2019. Differentiable convex optimization layers. *Advances in neural information processing systems*, 32.
- Agrawal, A.; Barratt, S.; Boyd, S.; and Stellato, B. 2020. Learning convex optimization control policies. In *Learning for Dynamics and Control*, 361–373. PMLR.
- Berge, C. 1997. *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*. Courier Corporation.
- Bertsekas, D. 2019. *Reinforcement learning and optimal control*. Athena Scientific.
- Borrelli, F.; Bemporad, A.; and Morari, M. 2017. *Predictive control for linear and hybrid systems*. Cambridge University Press.
- Boyd, S.; Boyd, S. P.; and Vandenberghe, L. 2004. *Convex optimization*. Cambridge university press.
- Chen, X.; Qu, G.; Tang, Y.; Low, S.; and Li, N. 2022. Reinforcement learning for selective key applications in power systems: Recent advances and future challenges. *IEEE Transactions on Smart Grid*.
- Conti, E.; Madhavan, V.; Petroski Such, F.; Lehman, J.; Stanley, K.; and Clune, J. 2018. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. *Advances in neural information processing systems*, 31.
- De Nijs, F.; Walraven, E.; De Weerd, M.; and Spaan, M. 2021. Constrained multiagent Markov decision processes: A taxonomy of problems and algorithms. *Journal of Artificial Intelligence Research*, 70: 955–1001.
- Dempe, S.; and Zemkoho, A. 2020. *Bilevel optimization*. Springer.
- Dontchev, A. L.; and Rockafellar, R. T. 2009. *Implicit functions and solution mappings*, volume 543. Springer.
- Doukhan, P. 2012. *Mixing: properties and examples*, volume 85. Springer Science & Business Media.
- Dulac-Arnold, G.; Levine, N.; Mankowitz, D. J.; Li, J.; Paduraru, C.; Gowal, S.; and Hester, T. 2021. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Machine Learning*, 110(9): 2419–2468.
- Facchinei, F.; and Pang, J.-S. 2007. *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media.
- Frazier, P. I. 2018. Bayesian optimization. In *Recent Advances in Optimization and Modeling of Contemporary Problems*, 255–278. INFORMS.
- Gangwani, T.; and Peng, J. 2018. Policy Optimization by Genetic Distillation. In *International Conference on Learning Representations*.
- Ghadimi, S.; and Lan, G. 2013. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4): 2341–2368.
- Ghadimi, S.; Perkins, R. T.; and Powell, W. B. 2020. Reinforcement Learning via Parametric Cost Function Approximation for Multistage Stochastic Programming. *arXiv preprint arXiv:2001.00831*.
- Guez, A.; Weber, T.; Antonoglou, I.; Simonyan, K.; Vinyals, O.; Wierstra, D.; Munos, R.; and Silver, D. 2018. Learning to search with mctsnet. In *International conference on machine learning*, 1822–1831. PMLR.
- Haydari, A.; and Yilmaz, Y. 2020. Deep reinforcement learning for intelligent transportation systems: A survey. *IEEE Transactions on Intelligent Transportation Systems*.
- Jin, M.; Khattar, V.; Kaushik, H.; Sel, B.; and Jia, R. 2023. On Solution Functions of Optimization: Universal Approximation and Covering Number Bounds. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- Karkus, P.; Hsu, D.; and Lee, W. S. 2017. Qmdp-net: Deep learning for planning under partial observability. *Advances in neural information processing systems*, 30.
- Kathirgamanathan, A.; Twardowski, K.; Mangina, E.; and Finn, D. P. 2020. A Centralized Soft Actor Critic Deep Reinforcement Learning Approach to District Demand Side Management through CityLearn. In *Proceedings of the 1st International Workshop on Reinforcement Learning for Energy Management in Buildings & Cities*, 11–14.
- Khadka, S.; and Tumer, K. 2018. Evolution-guided policy gradient in reinforcement learning. *Advances in Neural Information Processing Systems*, 31.
- Liu, S.; Chen, P.-Y.; Kailkhura, B.; Zhang, G.; Hero III, A. O.; and Varshney, P. K. 2020. A primer on zeroth-order optimization in signal processing and machine learning: Principals, recent advances, and applications. *IEEE Signal Processing Magazine*, 37(5): 43–54.
- Mania, H.; Guy, A.; and Recht, B. 2018. Simple random search of static linear policies is competitive for reinforcement learning. *Advances in Neural Information Processing Systems*, 31.
- Moerland, T. M.; Broekens, J.; and Jonker, C. M. 2020. Model-based reinforcement learning: A survey. *arXiv preprint arXiv:2006.16712*.
- Nian, R.; Liu, J.; and Huang, B. 2020. A review on reinforcement learning: Introduction and applications in industrial process control. *Computers & Chemical Engineering*, 139: 106886.

- Pflug, G. C.; and Pichler, A. 2014. *Multistage stochastic optimization*, volume 1104. Springer.
- Powell, W. 2020. Reinforcement Learning and Stochastic Optimization: A Unified Framework for Sequential Decisions. *Princeton NJ*.
- Prakash, B.; Waytowich, N.; Ganesan, A.; Oates, T.; and Mohsenin, T. 2020. Guiding safe reinforcement learning policies using structured language constraints. *UMBC Student Collection*.
- Prékopa, A. 2013. *Stochastic programming*, volume 324. Springer Science & Business Media.
- Racanière, S.; Weber, T.; Reichert, D.; Buesing, L.; Guez, A.; Jimenez Rezende, D.; Puigdomènech Badia, A.; Vinyals, O.; Heess, N.; Li, Y.; et al. 2017. Imagination-augmented agents for deep reinforcement learning. *Advances in neural information processing systems*, 30.
- Raffin, A.; Hill, A.; Gleave, A.; Kanervisto, A.; Ernestus, M.; and Dormann, N. 2021. Stable-baselines3: Reliable reinforcement learning implementations. *The Journal of Machine Learning Research*, 22(1): 12348–12355.
- Rolnick, D.; Donti, P. L.; Kaack, L. H.; Kochanski, K.; Lacoste, A.; Sankaran, K.; Ross, A. S.; Milojevic-Dupont, N.; Jaques, N.; Waldman-Brown, A.; et al. 2022. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2): 1–96.
- Salimans, T.; Ho, J.; Chen, X.; Sidor, S.; and Sutskever, I. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*.
- Schrittwieser, J.; Antonoglou, I.; Hubert, T.; Simonyan, K.; Sifre, L.; Schmitt, S.; Guez, A.; Lockhart, E.; Hassabis, D.; Graepel, T.; et al. 2020. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839): 604–609.
- Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25.
- Spall, J. C. 2005. *Introduction to stochastic search and optimization: estimation, simulation, and control*, volume 65. John Wiley & Sons.
- Szita, I.; and Lörincz, A. 2006. Learning Tetris using the noisy cross-entropy method. *Neural computation*, 18(12): 2936–2941.
- Tamar, A.; Wu, Y.; Thomas, G.; Levine, S.; and Abbeel, P. 2016. Value iteration networks. *Advances in neural information processing systems*, 29.
- Thompson, H. 1950. Truncated normal distributions. *Nature*, 165(4194): 444–445.
- Vazquez-Canteli, J. R.; Dey, S.; Henze, G.; and Nagy, Z. 2020. CityLearn: Standardizing Research in Multi-Agent Reinforcement Learning for Demand Response and Urban Energy Management. *arXiv preprint arXiv:2012.10504*.
- Vazquez-Canteli, J. R.; Henze, G.; and Nagy, Z. 2020. MARLISA: Multi-Agent Reinforcement Learning with Iterative Sequential Action Selection for Load Shaping of Grid-Interactive Connected Buildings. In *Proceedings of the 7th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, 170–179.
- Zhigljavsky, A. A. 2012. *Theory of global random search*, volume 65. Springer Science & Business Media.
- Zhou, Z.-H.; Yu, Y.; and Qian, C. 2019. *Evolutionary learning: Advances in theories and algorithms*. Springer.