Constrained Online Convex Optimization with Polyak Feasibility Steps

Spencer Hutchinson **Mahnoosh Alizadeh**

University of California, Santa Barbara

CONTRIBUTIONS

We study online convex optimization with functional constraints. Our new approach, that we call *Polyak feasibility steps*, enjoys $O(\sqrt{T})$ regret and anytime constraint satisfaction while only using first-order constraint feedback.

Reference	Regret	Constraint Guarantee	Known Strictly-feasible Point?
[1,2,3,4]	$O(\sqrt{T})$	$\sum_{t=1}^{T} g(x_t) \leq 0$	no
this work	$O(\sqrt{T})$	$g(x_t) \leq 0,\ \forall t \in [T]$	yes
this work	$O(\sqrt{T})$	$g(x_t) \leq 0,\ \forall t \geq C \log T$	no

PROBLEM SETUP

Interaction Model

At each round $t \in [T]$:

1. Choose action $x_t \in \mathbb{R}^d$.
2. Observe cost function $f_t: \mathbb{R}^d \to \mathbb{R}$.

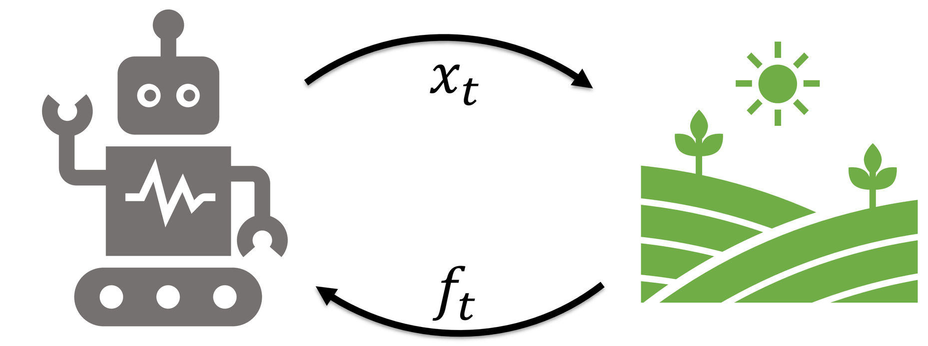

Learning Goals

Minimize regret with respect to feasible set:

$$R_T = \sum_{t=1}^{T} f_t(x_t) - \min_{x: g(x) \leq 0} \sum_{t=1}^{T} f_t(x)$$

Keep constraint violations $g(x_1), g(x_2), \ldots, g(x_T)$ small.

Assumptions

Cost gradients ∇f_t are bounded.
Action set $\mathcal{X}$ is bounded.
Constraint gradients ∇g bounded.
Exists $x \in \mathbb{R}^d$ such that $g(x) < 0$.

ALGORITHM

In each round $t \in [T]$:

1. Cost gradient descent:
$$y_t = x_t - \eta \nabla f_t(x_t)$$
2. Polyak feasibility step:
$$x_{t+1} = y_t - \frac{[g(x_t) + \nabla g(x_t)^\top (y_t - x_t) + \rho]_+}{\|\nabla g(x_t)\|^2} \nabla g(x_t)$$

Motivation

1. Polyak step-size:
$$y_t - \frac{g(y_t)}{\|\nabla g(y_t)\|^2} \nabla g(y_t)$$
2. First-order approximation of constraint:
$$g(y_t) \approx g(x_t) + \nabla g(x_t)^\top (y_t - x_t)$$
3. Tightening of constraint:
$$g(x_t) \leftarrow g(x_t) + \rho$$

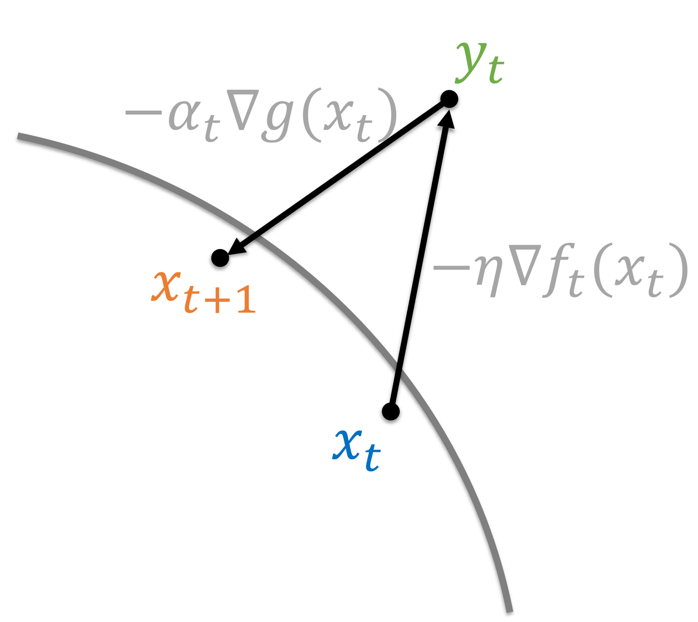

ANALYSIS

Regret Analysis

1. Represent as projection on to separating hyperplane:
$$x_{t+1} = \Pi_{\mathcal{H}_t}(y_t)$$
$$\mathcal{H}_t = \{x: g(x_t) + \nabla g(x_t)^\top (x - x_t) + \rho \leq 0\}$$
2. Projection shrinks distance to tightened feasible set:
$$\|\Pi_{\mathcal{H}_t}(y_t) - x\| \leq \|y_t - x\|, \qquad \forall x \in \mathcal{X}_\rho = \{x: g(x) + \rho \leq 0\}$$
3. Apply standard OGD analysis:
$$R_T = O(\sqrt{T}) + C\rho$$

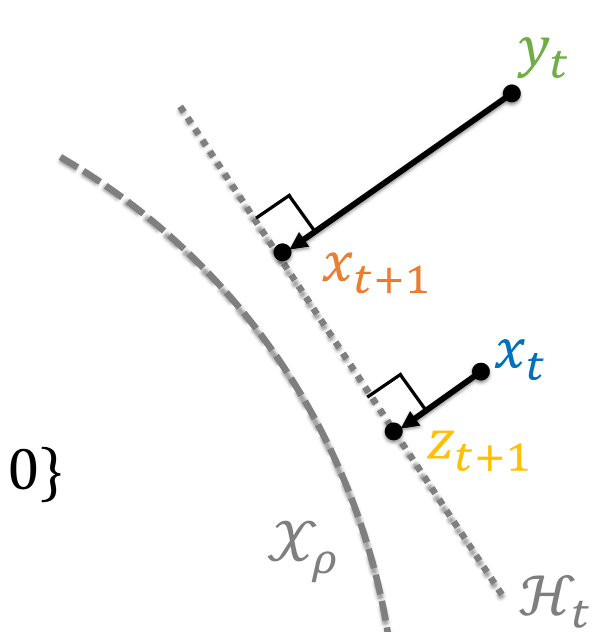

Feasibility Analysis

1. Polyak step-size *rapidly* shrinks distance to feasible set:
$$z_{t+1} = x_t - \frac{g(x_t) + \rho}{\|\nabla g(x_t)\|^2} \nabla g(x_t), \qquad \mathrm{dist}(z_{t+1}, \mathcal{X}_\rho) \leq \gamma \, \mathrm{dist}(x_t, \mathcal{X}_\rho)$$
2. Algorithm approximates Polyak step-size:
$$\|z_{t+1} - x_{t+1}\| \leq \eta G$$
3. Put everything together:
$$\mathrm{dist}(x_{t+1}, \mathcal{X}_\rho) \leq \gamma^t \mathrm{dist}(x_1, \mathcal{X}_\rho) + \frac{\eta G}{1-\gamma}, \qquad g(x_{t+1}) \leq C_1 \gamma^t \mathrm{dist}(x_1, \mathcal{X}_\rho) + C_2 \eta - \rho$$

EXPERIMENTS

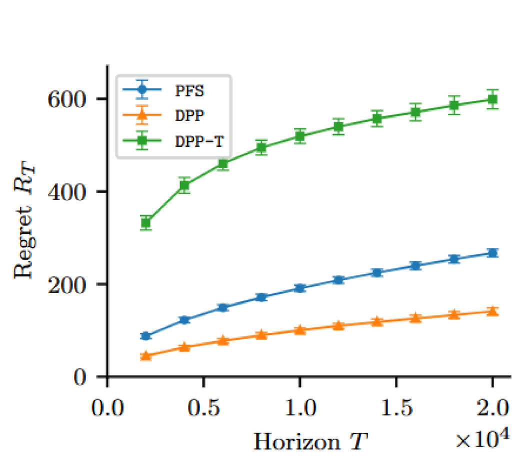

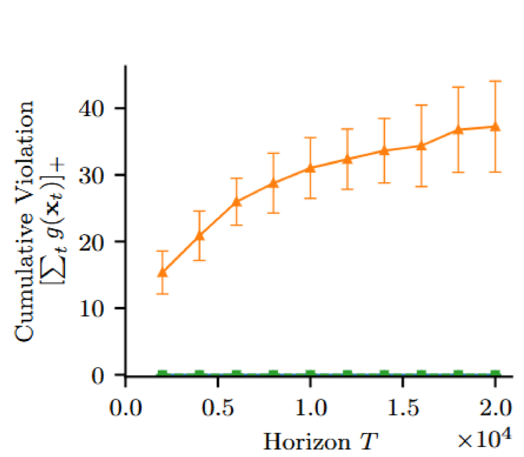

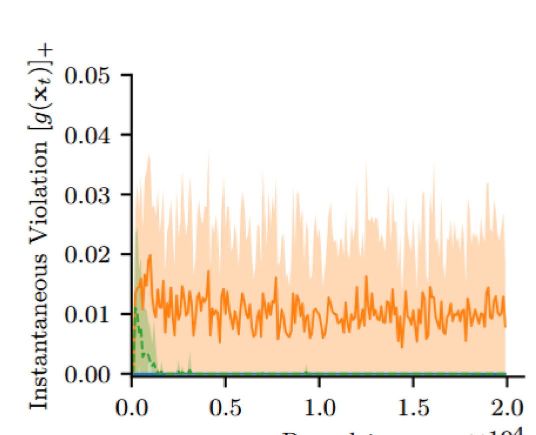

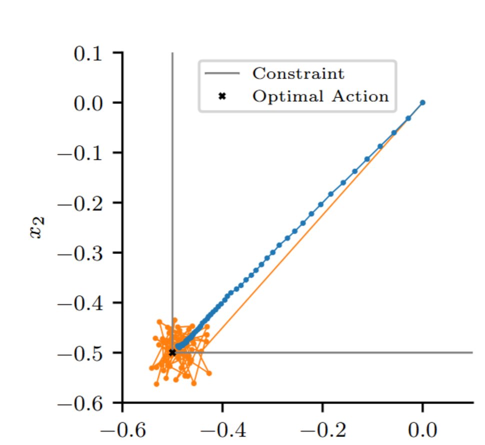

REFERENCES

[1] Mahdavi et al. Trading regret for efficiency: online convex optimization with long term constraints. JMLR 2012.
[2] Jenatton et al. Adaptive algorithms for online convex optimization with long-term constraints. ICML 2016.
[3] Yuan and Lamperski. Online convex optimization for cumulative constraints. NeurIPS 2018.
[4] Yu et al. Online convex optimization with stochastic constraints. NeurIPS 2017.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation under grant #2330154.