Search or Article-id

(Help | Advanced search)

All papers | Go!

Mathematics > Optimization and Control

Stochastic convex optimization with bandit feedback

Alekh Agarwal, Dean P. Foster, Daniel Hsu, Sham M. Kakade, Alexander Rakhlin

(Submitted on 8 Jul 2011 (v1), last revised 8 Oct 2011 (this version, v2))

This paper addresses the problem of minimizing a convex, Lipschitz function f over a convex, compact set $\xset$ under a stochastic bandit feedback model. In this model, the algorithm is allowed to observe noisy realizations of the function value $f(x)$ at any query point $x \in \xset$. The quantity of interest is the regret of the algorithm, which is the sum of the function values at algorithm's query points minus the optimal function value. We demonstrate a generalization of the ellipsoid algorithm that incurs $\otil(\poly(d)\sqrt{T})$ regret. Since any algorithm has regret at least $\Omega(\sqrt{T})$ on this problem, our algorithm is optimal in terms of the scaling with T.

Subjects:	**Optimization and Control (math.OC)**; Learning (cs.LG); Systems and Control (cs.SY)
Cite as:	**arXiv:1107.1744** [math.OC]
	(or **arXiv:1107.1744v2** [math.OC] for this version)

Submission history

From: Alekh Agarwal [view email]
[v1] Fri, 8 Jul 2011 22:18:05 GMT (525kb,D)
[v2] Sat, 8 Oct 2011 06:06:43 GMT (526kb,D)

Which authors of this paper are endorsers?

Link back to: arXiv, form interface, contact.

Download:

- PDF
- Other formats

Current browse context:
math.OC
< prev | **next >**
new | recent | 1107

Change to browse by:

cs
     cs.LG
     cs.SY
math

References & Citations

- NASA ADS

Bookmark (what is this?)

Science WISE