Online Optimization with Predictions and Non-convex Losses

Yiheng Lin^{*†}
IIS, Tsinghua University
Beijing, China

Gautam Goel[†]
California Institute of Technology
Pasadena, California, USA

Adam Wierman[†]
California Institute of Technology
Pasadena, California, USA

ABSTRACT

We study online optimization in a setting where an online learner seeks to optimize a per-round hitting cost, which may be non-convex, while incurring a movement cost when changing actions between rounds. We ask: *under what general conditions is it possible for an online learner to leverage predictions of future cost functions in order to achieve near-optimal costs?* Prior work has provided near-optimal online algorithms for specific combinations of assumptions about hitting and switching costs, but no general results are known. In this work, we give two general sufficient conditions that specify a relationship between the hitting and movement costs which guarantee that a new algorithm, Synchronized Fixed Horizon Control (SFHC), achieves a $1 + O(1/w)$ competitive ratio, where w is the number of predictions available to the learner. Our conditions do not require the cost functions to be convex, and we also derive competitive ratio results for non-convex hitting and movement costs. Our results provide the first constant, dimension-free competitive ratio for online non-convex optimization with movement costs. We also give an example of a natural problem, Convex Body Chasing (CBC), where the sufficient conditions are not satisfied and prove that no online algorithm can have a competitive ratio that converges to 1.

CCS CONCEPTS

• **Theory of computation** → **Online learning algorithms**; *Non-convex optimization*; • **Computing methodologies** → *Machine learning algorithms*;

KEYWORDS

Online non-convex optimization, online convex optimization (OCO), non-convex optimization, competitive analysis

ACM Reference Format:

Yiheng Lin, Gautam Goel, and Adam Wierman. 2020. Online Optimization with Predictions and Non-convex Losses. In *ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS '20 Abstracts)*, June 8–12, 2020, Boston, MA, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3393691.3394208>

^{*}This work was done when Yiheng Lin was visiting California Institute of Technology.

[†]This work was supported by NSF grants AitF-1637598 and CNS-1518941, with additional support for Gautam Goel provided by an Amazon AWS AI Fellowship.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS '20 Abstracts, June 8–12, 2020, Boston, MA, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7985-4/20/06.

<https://doi.org/10.1145/3393691.3394208>

1 INTRODUCTION

Online optimization is a classical area in online learning with a long and impactful history. In this paper, we study a variation of online optimization where the learner incurs a movement (switching) cost associated with the change in actions between consecutive rounds. Specifically, we study online optimization in a setting where an online learner interacts with the environment in a sequence of rounds $1 \dots T$. In each round, a cost function $f_t : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is revealed and the learner chooses a point $x_t \in \mathbb{R}^d$ in response. After picking its point, the learner pays a *hitting cost* $f_t(x_t)$ as well as a *movement (switching) cost* $c(x_t, x_{t-1})$, which penalizes the learner for changing its actions between rounds.

In many applications, the classical formulation of an online learner is too restrictive. It is not true that the learner has no information about future costs, instead the learner has the ability to derive (noisy) forecasts of future cost functions. As a result, there has been a great deal of work focused on designing algorithms for online learners that have access to predictions of future costs [3, 6, 7]. This line of work, initiated by [7], seeks to design algorithms which have competitive ratios that converge to 1 as the number of predictions available to the algorithm, w , grows. More specifically, in this line of work, at time t an online learner with prediction window w observes the cost functions $f_t \dots f_{t+w-1}$ before choosing the point x_t . Note that the case of $w = 1$ captures the standard SOCO setting. Given these predictions, the learner seeks to have a competitive ratio of the form $1 + g(w)$, where $g(w) \rightarrow 0$ as $w \rightarrow \infty$. Thus, as the number of predictions grows the cost of the learner converges to the offline optimal cost.

There has been considerable progress in the design of competitive algorithms for SOCO, both with and without access to predictions. However, at this point all existing results require specific assumptions on both the hitting costs and movement costs. In this paper, instead of studying a specific class of costs, we ask: *under what general conditions is it possible for an online learner to achieve near-optimal costs both with and without predictions?* In particular, is it possible to obtain constant-competitive algorithms without assumptions like strong-convexity and local polyhedrality; potentially even in the case of *non-convex costs*?

The case of non-convex costs is particularly tantalizing given the importance of non-convex losses for machine learning and the prominence of non-convex costs in applications such as power systems and networking. Techniques from non-convex optimization have been applied to a wide variety of problems in machine learning, including matrix factorization, phase retrieval, and sparse recovery. The Optimal Power Flow (OPF) problem at the core of the operation of power systems is also non-convex [8, 9]; thus requiring online non-convex optimization for real-time control. Non-convex optimization in online settings has also been studied in a variety

of other contexts, such as portfolio optimization [1, 5] and support vector machines [4, 10], among many others.

2 MAIN RESULTS

In this paper we introduce two general, sufficient conditions under which is possible to achieve a constant competitive ratio without predictions *and* to leverage predictions to achieve near-optimal cost, i.e., a $1 + O(1/w)$ competitive ratio. Importantly, these conditions do not require convexity of the hitting or movement costs.

Condition I: Order of Growth. *The hitting costs f_t and movement cost c satisfy $f_t(x) \geq \lambda(c(x, v_t) + c(v_t, x))$ for all x , where v_t is a global minimum of f_t .*

Condition II: Approximate Triangle Inequality. *The movement cost c satisfies $c(x, z) \leq \eta(c(x, y) + c(y, z))$ for all x, y, z .*

The first sufficient condition is an order of growth condition that ensures the hitting cost functions grow at least as quickly as the switching costs as one moves away from the minimizer. The second condition requires that the switching costs satisfy an approximate version of the triangle inequality. Nearly all assumptions made in previous papers on online optimization with movement costs are special cases of these conditions. While we do not prove that these conditions are necessary, we show that an important class of online optimization problems, namely Convex Body Chasing, violates the conditions, and that furthermore it is impossible for an online learner to leverage predictions to achieve near-optimal costs in this class.

To show that these two conditions are sufficient, we propose a novel algorithm, Synchronized Fixed Horizon Control (SFHC), and show that it is constant-competitive whenever the two conditions hold, including both when the cost functions are convex and non-convex. More specifically, we introduce two variants of SFHC, Deterministic SFHC and Randomized SFHC.

The design of SFHC is inspired by the design of Averaging Fixed Horizon Control (AFHC) [7], which has served as the basis for many algorithms in this space. Like AFHC, Deterministic SFHC works by averaging the choices of w different subroutines. However, the subroutines are very different than AFHC. At each time step τ , one of the subroutines of AFHC optimizes the cost over the window $[\tau, \tau + w - 1]$ given the starting state $x_{\tau-1}$ (see [7]). The SFHC subroutines perform a similar optimization, but with an additional constraint that the point selected at the end of the window is “synced” to the minimizer of the hitting cost at that timestep. These synchronization points ensure that, when the sufficient conditions hold, the algorithm does not drift too far from the actions of the offline optimal. Thus, rather than optimize cost, SFHC is designed to track the offline optimal (which also implicitly leads to achieving good cost). The key difference between Deterministic SFHC and Randomized SFHC is that Randomized SFHC chooses an action of a subroutine uniformly at random rather than averaging the choices of the subroutines. It is perhaps surprising that randomization helps in the case of non-convex costs given that [2] shows that randomization cannot help in the case of SOCO.

In the case when costs are convex, Deterministic SFHC provides a competitive ratio of $\max\left(1 + \frac{\eta + \eta^2}{2\lambda}, \eta^2\right)$ without access to predictions and a competitive ratio of $1 + O(1/w)$ in the case of predictions.

In the case when costs are non-convex, Deterministic SFHC maintains a competitive ratio of $\max\left(1 + \frac{\eta + \eta^2}{2\lambda}, \eta^2\right)$ without access to predictions but provides a competitive ratio of $C + O(1/w)$ in the case of predictions, where $C > 1$. Thus, it does not leverage predictions to ensure near-optimal cost. However, randomization can be used to improve the result in the case of predictions. Specifically, Randomized SFHC provides a competitive ratio of $1 + O(1/w)$ for general non-convex functions that satisfy our sufficient conditions, given an oblivious adversary. Further, the result extends (with slight modifications to the design of Randomized SFHC) to the case of a semi-adaptive adversary. These results represent *the first constant-competitive guarantees for online optimization with movement costs and non-convex losses*.

REFERENCES

- [1] David Ardia, Kris Boudt, Peter Carl, Katharine M Mullen, and Brian Peterson. 2010. Differential evolution (deoptim) for non-convex portfolio optimization. (2010).
- [2] Nikhil Bansal, Anupam Gupta, Ravishankar Krishnaswamy, Kirk Pruhs, Kevin Schewior, and Cliff Stein. 2015. A 2-competitive algorithm for online convex optimization with switching costs. In *Proceedings of the Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- [3] Nianguan Chen, Anish Agarwal, Adam Wierman, Siddharth Barman, and Lachlan LH Andrew. 2015. Online convex optimization using predictions. *ACM SIGMETRICS Performance Evaluation Review* 43, 1 (2015), 191–204.
- [4] Seyda Ertekin, Leon Bottou, and C Lee Giles. 2010. Nonconvex online support vector machines. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33, 2 (2010), 368–381.
- [5] Pavlo Krokmal, Jonas Palmquist, and Stanislav Uryasev. 2002. Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of risk* 4 (2002), 43–68.
- [6] Yingying Li, Guannan Qu, and Na Li. 2018. Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit. *arXiv preprint arXiv:1801.07780* (2018).
- [7] Minghong Lin, Zhenhua Liu, Adam Wierman, and Lachlan LH Andrew. 2012. Online algorithms for geographical load balancing. In *Proceedings of the International Green Computing Conference (IGCC)*. 1–10.
- [8] Steven H Low. 2014. Convex relaxation of optimal power flow—Part I: Formulations and equivalence. *IEEE Transactions on Control of Network Systems* 1, 1 (2014), 15–27.
- [9] Steven H Low. 2014. Convex relaxation of optimal power flow—Part II: Exactness. *IEEE Transactions on Control of Network Systems* 1, 2 (2014), 177–189.
- [10] Llew Mason, Peter L Bartlett, and Jonathan Baxter. 2000. Improved generalization through explicit optimization of margins. *Machine Learning* 38, 3 (2000), 243–255.