Computer-Aided Methods for Evaluating Online Optimization Algorithms

Erwan Meunier
ICTEAM, UCLouvain
erwan.meunier@uclouvain.be

Julien M. Hendrickx
ICTEAM, UCLouvain
julien.hendrickx@uclouvain.be

This abstract presents a line of research on the analysis of online optimization algorithms using the Performance Estimation Problem framework.

1 Online Optimization

Online optimization (also known as Online Learning) is a prominent area of research where the objective is to minimize functions f_t (usually convex and L-Lipschitz) received sequentially over a finite-time horizon $t \in [T]$, where $f_t = f, \forall t \in [T]$ is constant, corresponding to the classical offline setting. Facing different objective functions at each time-step can be useful to capture some dynamic problems where the optimal solution of the problem (e.g. the physic of the problem) is not fixed once and for all. This kind of problem arise in many applications such as in Control [4] or in Machine Learning [2] to mention a few.

A paradigmatic method to solve online optimization is the so-called *Online Gradient Descent* (OGD) [3] which updates the current minimizer estimate x_t by doing an *online* gradient step (i.e. the subgradient $\partial f_t(x_t)$ is indexed by the time) controlled by a step-size strategy η_t, and projecting back onto the convex bounded feasible set Ω:

$$x_{t+1} \leftarrow P_\Omega(x_t - \eta_t \partial f_t(x_t)).$$

Performance metrics are manifold for online optimization algorithms, the most widely used is the Static Regret defined as:

$$Regret = \sum_{t=1}^{T} f_t(x_t) - \min_{u\in\Omega} \sum_{t=1}^{T} f_t(u)$$

which compares the total cost incurred at each time-step in the Online scenario with the total cost incurred if the sequence of functions was known in advance (offline setting).

The classical aim is to reach a sublinear regret and for a proper choice of step-sizes, one can show that the regret for OGD is indeed $O(DG\sqrt{T})$ [3]. However, the degree of conservatism and tightness of online algorithms have not been completely analyzed in many scenarios, which is important to evaluate and improve their performance.

2 Performance Estimation Problem

To obtain an accurate performance analysis of online optimization methods, we exploit the *Performance Estimation Problem (PEP)* framework [1], which allows computing exact worst-case bound of first-order optimization algorithms.

In principle, a tight performance bound on an algorithm could be obtained by running it on every single instance, function, and initial condition allowed by the setting considered and selecting the worst performance obtained. The PEP formulates this abstract idea as a real optimization problem that maximizes the error measure of the algorithm result, over all possible functions and initial conditions allowed.

In most-cases PEP optimization problems can be cast into easy to solve Semidefinite Programs (SDP).

3 Contributions and Directions of research

First we show that state-of-the-art upper-bounds on the Static Regret are not tight. Then, we give a Burer-Monteiro non-convex PEP formulation returning the lowest dimensional exact worst-case instance for OGD and we demonstrate experimentally that it is often in 2 dimensions.

Finally, under some mild assumptions and given a fixed step-size strategy, we propose an analytical formula to build a worst-case instance with f_t being piecewise-linear and giving the associated tight worst-case bound for OGD. This formula paves the way for building exact worst-case instances for broader classes of algorithms, such as *Online Mirror Descent* [3].

Looking ahead, we outline several research avenues, including Decentralized Online Optimization and the development of methods through step-size strategy optimization. In the first case, we demonstrate, based on numerical experiments, that existing upper bounds on regret are overly conservative by several orders of magnitude.

Acknowledgment E. Meunier is supported by the French Community of Belgium through a FRIA fellowship (F.R.S.-FNRS).

References

[1] A. B. Taylor, J. M. Hendrickx, and F. Glineur, “Smooth strongly convex interpolation and exact worst-case performance of first-order methods,” *Mathematical Programming*, vol. 161, no. 1, pp. 307–345, Jan. 2017.

[2] S. C. Hoi, D. Sahoo, J. Lu, and P. Zhao, “Online learning: A comprehensive survey,” *Neurocomputing*, vol. 459, pp. 249–289, 2021.

[3] E. Hazan, *Introduction to Online Convex Optimization* (Adaptative Computation and Machine Learning series), 2nd ed. The MIT Press, Sep. 6, 2022, 248 pp.

[4] E. Hazan and K. Singh, *Introduction to online non-stochastic control*, 2024. arXiv: 2211.09619.