Regret bounded by gradual variation for online convex optimization

**Tianbao Yang · Mehrdad Mahdavi · Rong Jin ·
Shenghuo Zhu**

Abstract Recently, it has been shown that the regret of the Follow the Regularized Leader (FTRL) algorithm for online linear optimization can be bounded by the total variation of the cost vectors rather than the number of rounds. In this paper, we extend this result to general online convex optimization. In particular, this resolves an open problem that has been posed in a number of recent papers. We first analyze the limitations of the FTRL algorithm as proposed by Hazan and Kale (in Machine Learning 80(2–3), 165–188, 2010) when applied to online convex optimization, and extend the definition of variation to a gradual variation which is shown to be a lower bound of the total variation. We then present two novel algorithms that bound the regret by the gradual variation of cost functions. Unlike previous approaches that maintain a single sequence of solutions, the proposed algorithms maintain two sequences of solutions that make it possible to achieve a variation-based regret bound for online convex optimization.

To establish the main results, we discuss a lower bound for FTRL that maintains only one sequence of solutions, and a necessary condition on smoothness of the cost functions for obtaining a gradual variation bound. We extend the main results three-fold: (i) we present a general method to obtain a gradual variation bound measured by general norm; (ii) we extend algorithms to a class of online non-smooth optimization with gradual variation bound;

T. Yang (✉) · S. Zhu
NEC Laboratories America, Cupertino, USA
e-mail: tyang@nec-labs.com

S. Zhu
e-mail: zsh@nec-labs.com

M. Mahdavi · R. Jin
Department of Computer Science and Engineering, Michigan State University, East Lansing, USA

M. Mahdavi
e-mail: mahdavim@cse.msu.edu

R. Jin
e-mail: rongjin@cse.msu.edu

🖄 Springer