CREST - GENES
Cours doctoraux 2020 - 2021

Adaptivity in Online Learning

Tim van Erven
Leiden University

SCHEDULE	Monday (17/02/20)	Wednesday (19/02/20)	Thursday (20/02/20)
ENSAI - Salle 012	14h-17h	14h-17h	14h-16h

Abstract: Online learning methods for Online Convex Optimization (OCO) are normally designed to be robust to adversarial data. This means they come with a guaranteed upper bound on their regret that holds uniformly over all data sequences. But the set of all data sequences also includes data sequences that are completely irrelevant in practice, and we see that the matching lower bounds are usually witnessed by such irrelevant data sequences. Specifically, the lower bounds are usually proven for completely random data, from which there is nothing interesting to learn.
This has motivated the development of a series of adaptive methods, which come with data-dependent regret guarantees. These data-dependent guarantees are never worse than the previous uniform guarantees, but if the data happen to be relatively "easy", they can be orders of magnitude better. Another important type of adaptivity is to remove the need for a user to set hyperparameters like an upper bound on the norm of the optimal parameters or the Lipschitz constant of the losses. In this course I will introduce the mathematical theory of OCO from first principles and build up to state-of-the-art results in both types of adaptivity.

Suggested Literature:
* M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In Proc. of the 20th Annual International Conf. on Machine Learning (ICML), pages 928–936, 2003.
* S. de Rooij, T. van Erven, P. D. Grünwald and W. M. Koolen. Follow the Leader If You Can, Hedge If You Must. Journal of Machine Learning Research, vol. 15, pages 1281-1316, 2014.
* W. M. Koolen and T. van Erven. Second-order quantile methods for experts and combinatorial games. In Proc. of the 28th Annual Conf. on Learning Theory (COLT), pages 1155–1175, 2015.
* D. van der Hoeven, T. van Erven and W. Kotłowski. The Many Faces of Exponential Weights in Online Learning. Proc. of the 31st Conf. on Learning Theory (COLT), pages 2067-2092, 2018.
* T. van Erven and W. M. Koolen. Metagrad: Multiple learning rates in online learning. In Advances in Neural Information Processing Systems 29 (NIPS), pages 3666–3674, 2016.
* A. Cutkosky, F. Orabona. Black-Box Reductions for Parameter-free Online Learning in Banach Spaces. In Proc. of the 31st Conf. On Learning Theory (COLT), pages 1493-1529, 2018.

Further Literature:
* E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. Machine Learning, 69(2-3):169–192, 2007.
* W. M. Koolen, P. Grünwald and T. van Erven. Combining Adversarial Guarantees and Stochastic Fast Rates in Online Learning. Advances in Neural Information Processing Systems 29 (NeurIPS), pages 4457-4465, 2016.
* A. Cutkosky. Combining Online Learning Guarantees. ArXiv:1902.09003 preprint, 2019.