
A Universal Catalyst for First-Order Optimization*

Julien Mairal

Abstract

We introduce a generic scheme for accelerating first-order optimization methods in the sense of Nesterov. Our approach consists of minimizing a convex objective by approximately solving a sequence of well-chosen auxiliary problems, leading to faster convergence. This strategy applies to a large class of algorithms, including gradient descent, block coordinate descent, SAG, SAGA, SDCA, SVRG, Finito/MISO, and their proximal variants. For all of these approaches, we provide acceleration and explicit support for non-strongly convex objectives. In addition to theoretical speed-up, we also show that acceleration is useful in practice, especially for ill-conditioned problems where we measure significant improvements. [This is a joint work with Hongzhou Lin and Zaid Harchaoui.]

*Machine Learning External Seminar, Gatsby Unit, January 26, 2016.