Gradient descent with a general cost

Pierre-Cyril Aubin

INRIA SIERRA, France

pierrecyrilaubin@gmail.com

How to design algorithms that extend classical methods (and their convergence theory) beyond the Euclidean setting? How to integrate different families of algorithms into a unified framework? How to systematically approximate optimization problems? We show that c-transforms allow to build majorizing surrogates, which can be tackled through alternating minimization. The "five-point property" of Csiszar and Tusnady gives (sub)linear convergence rates for the values of the iterates. This property corresponds to a novel notion of c-semiconvexity, extending relative strong convexity, and intimately related to the MTW tensor . The mirror/Riemannian/natural gradient descent algorithms can all be cast in this formalism, leading to novel convergence conditions and rates.

Joint work with Flavien Léger (INRIA MOKAPLAN, France).