

1172-49-138

Melanie Weber* (mw25@math.princeton.edu) and **Suvrit Sra**. *Frank-Wolfe Methods for efficient geodesically convex optimization.*

Many applications involve non-Euclidean data, where exploiting Riemannian geometry can deliver algorithms that are computationally superior to standard nonlinear programming approaches, resulting in an increasing interest in Riemannian methods in the machine learning community. In this talk, we focus on Riemannian Frank-Wolfe (RFW) methods, a class of projection-free algorithms for constraint geodesically convex optimization. The efficiency of the algorithm depends crucially on three aspects, (1) its iteration complexity, (2) the complexity of computing the Riemannian gradient, and (3) the complexity of the Riemannian “linear” oracle (RLO), a subroutine at the heart of the algorithm, which determines the update. For (1), we will see that the classical RFW algorithm converges at a sublinear rate but attains a linear rate in special cases. For (2), we introduce RFW variants that utilize batch estimates of the gradient, significantly improving over full gradient estimates. Lastly, for (3), we will show that, while in general a nonconvex semi-definite problem, the RLO can be computed in closed form for some matrix-valued tasks. Furthermore, we show that RFW retains its sublinear convergence rate, even if the RLO is only solved approximately. (Received August 23, 2021)