

Many data analysis applications require the solution of optimization problems involving a sum of large number of functions. We consider the problem of minimizing a sum of n functions over a convex constraint set. Algorithms that carefully sub-sample to reduce n can improve the computational efficiency, while maintaining the original convergence properties. For second order methods, we first consider a general class of problems and give quantitative convergence results for variants of Newtons methods where the Hessian or the gradient is uniformly sub-sampled. We then show that, given certain assumptions, we can extend our analysis and apply non-uniform sampling which results in modified algorithms exhibiting more robustness and better dependence on problem specific quantities, such as the condition number.