1135-49-3031 **Prem Talwai*** (pmt55@cornell.edu). *Stochastic Proximal Gradient Methods for Metric Learning.* Preliminary report.

In the past decade, nuclear norm regularization has been frequently applied in multi-task learning to obtain low-rank linear predictors. However, a common challenge involved in these applications is the formulation of a computationally efficient optimization procedure for solving the resulting non-smooth, semidefinite convex programming problem. This paper develops a novel stochastic proximal gradient algorithm that applies a symmetric, low-rank stochastic gradient at each gradient descent step, and thereby eliminates the necessity of an expensive subsequent semidefinite projection. The developed algorithm was applied to two regularized "triplet loss" objective functions to learn distance metrics for measuring facial similarity. The resulting learned metrics were evaluated on the popular Labeled Faces in the Wild (LFW) face recognition database, and demonstrated over 90% accuracy in correctly matching face images of the same identity. (Received September 26, 2017)