1014-68-443 **Cynthia Rudin*** (`rudin@nyu.edu`), 4 Washington Place, Room 809, New York, NY 10003-6603.
Ranking with a P-Norm Push.

We are interested in supervised ranking with the following twist: our goal is to design algorithms that perform especially well near the top of the ranked list, and are only required to perform sufficiently well on the rest of the list. For instance, assume that you would like a recommendation for a good movie. Then, based on your own past movie preference data and the movie preference data of others, we wish to construct a ranked list of movies for you, making sure that the movies you are most likely to enjoy appear near the top.

Towards this goal, we provide a general form of convex objective that gives high-scoring examples more importance. This "push" near the top of the list can be chosen to be arbitrarily large or small. We choose ℓ_p-norms to provide a specific type of push; as p becomes large, the algorithm concentrates harder near the top of the list.

We will present our results on this problem: a generalization bound based on the p-norm objective, followed by a boosting-style algorithm and experimental results on UCI data. Then, a series of (not entirely difficult!) challenges will be presented, relating to the supervised ranking problem. This talk will be accessible to a wide mathematical audience. (Received September 15, 2005)