

# AdaBoost: The Expected but Untrue, and the Unexpected but True

Cynthia Rudin

New York University, 4 Washington Place, Room 809, New York, NY 10003

When the AdaBoost algorithm was introduced by Freund and Schapire in 1997, it became immediately popular for its empirical success in solving classification problems. Some fundamental convergence properties of AdaBoost remained unsolved until very recently, so I would like to describe the bizarre twists in the recent history of AdaBoost that have led to our current understanding. I will focus on two topics, a theory which was widely believed yet false, and an amazing property which no one could have guessed.

**The Expected but Untrue:** Since the development of AdaBoost eight years ago, there have been at least 3 large-scale empirical studies and several theoretical results to indicate that AdaBoost maximizes the “margin”. This fundamental question regarding AdaBoost’s convergence was finally solved last year, via the analysis of AdaBoost as a dynamical system. Not only does the nonlinear iterated map defined by AdaBoost yield rich dynamics such as stable cycles and chaos, it also provides the key to understanding the convergence of AdaBoost’s margins. The result is the opposite of what was thought to be true; AdaBoost does not necessarily maximize the margin.

**The Unexpected but True:** AdaBoost was designed for solving specifically the classification problem. It is quite surprising that AdaBoost is equally superior for a completely different application: the bipartite ranking problem. In fact, a very recent proof shows that AdaBoost asymptotically achieves the same AUC as RankBoost in the case where the constant hypothesis is included in the set of AdaBoost’s weak classifiers.

Portions of this talk are joint work with Corinna Cortes, Ingrid Daubechies, Mehryar Mohri, and Robert E. Schapire.