

1067-A0-35

Robert M. Bell*, AT&T Labs, Florham Park, NJ. *Lessons from the Netflix Prize.*

In October 2006, the DVD rental company Netflix released more than 100 million user ratings of movies for a competition to predict new ratings based on prior ratings. The size of the data (over 17,000 movies and 480,000 users) and the nature of human-movie interactions produced many modeling challenges. One allure to data analysts around the world was a \$1,000,000 prize for a team achieving a ten percent reduction in root mean squared prediction error relative to Netflix's existing algorithm. Besides producing a photo finish worthy of a movie, the 33-month competition spurred numerous advances in the science of recommender systems and machine learning, more generally. After describing some of the techniques used by the leaders, I will offer lessons and raise some questions about building massive prediction models; the role of statistics, computer science, and mathematics in such endeavors; and prizes as a way to advance science. This is joint work with Chris Volinsky and Yehuda Koren, current and former colleagues at AT&T Labs-Research. (Received June 10, 2010)