

From practice to theory in data science at Amazon

by

Dr. Charles Elkan

Amazon Fellow and Head, Machine Learning, Seattle/Silicon Valley

Amazon

Location: South Hall 202

Date: Monday, March 14

Time: 2-4 pm

Abstract

In this talk I will discuss examples of how Amazon serves customers and improves efficiency using learning algorithms applied to large-scale datasets, and open research questions that are inspired by these applications. These questions are unsolved issues that arise repeatedly and should be the topic of more research in the academic community. One research question is how to learn multilabel classifiers in time that is sublinear in the number of features, of examples, and of labels. Another is how to distinguish between users that are influenceable, versus merely likely to respond. A third question is how to measure and maximize the long-term benefit of movie recommendations and other recommendations. Note: All information in the talk will be already publicly available, and any opinions expressed will be strictly personal.

Bio

Charles Elkan is on leave from being a professor of computer science at the University of California, San Diego, working as Amazon Fellow and leader of machine learning for Amazon in Seattle and Silicon Valley. In the past, he has been a visiting associate professor at Harvard and a researcher at MIT. His published research has been mainly in machine learning, data science, and computational biology; the MEME algorithm that he developed with Ph.D. students has been used in over 3000 published research projects in biology and computer science. He is fortunate to have had inspiring undergraduate and graduate students who are in leadership positions now such as vice president at Google and professor at the University of Washington.