Joaquin Quiñonero Candela Microsoft Research, Cambridge

Title: Probabilistic Machine Learning in Computational Advertising

Abstract:

In the past few years online advertising has grown at least an order of magnitude faster than advertising on all other media. This talk focuses on advertising on search engines, where accurate predictions of the probability that a user clicks on an advertisement crucially benefit all three parties involved: the user, the advertiser, and the search engine. We present a fully probabilistic classification model that has the ability of learning from terabytes of web usage data. The model explicitly represents uncertainty allowing for fully probabilistic predictions. Observing 2 positives out of 10 instances or 200 positives out of 1000 instances leads in both cases to an average of 20%, but in the first case the uncertainty about the prediction should be larger. The only way of learning about an advertisement is to show it to the user, which often comes at an opportunity cost since this might not be the best ad to show. The talk describes the way in which the model presented allows for a natural way of exploring the available ad inventory.