

Ole Winther

Technical University of Denmark / University of Copenhagen

Title: Hierarchical Bayesian modelling for collaborative filtering

Abstract: The Netflix prize has generated a lot of interest for collaborative filtering in machine learning. Collaborative filtering is a good example of the large scale learning problems of web era. Perhaps surprisingly, many machine learning methods scale well to giga-byte data sets. In this talk I will present joint work with Ulrich Paquet and Blaise Thomson, Cambridge on a Bayesian hierarchical approach to the Netflix problem. The model combines an ordinal likelihood for the discrete ranks with movie-viewer low-rank matrix factorization. Inference in this model, which has between 25 to 100 millions parameters (depending upon the rank of the matrix factorization), is carried out with Gibbs sampling. This confirms previous work (Salakhutdinov and Mnih, ICML-2008) that Bayesian averaging is viable for large models and data sets. I will discuss how parallelization can be used to further scale up the problem. Our best model achieves a 6.32% improvement in terms of RMSE over the Netflix's own CineMatch system. This is to our knowledge the best single model result, i.e. with no use of blending of different algorithms.