
Bayesian nonparametric models for prediction in networks*

Sinead Williamson

Abstract

We are often interested in studying the interactions between entities for example, friendships or communications in social networks, interactions in biological networks, traffic patterns in road and communication networks. We are often interested in representing these networks in terms of some latent structure for example, explaining the pattern of emails sent between a collection of individuals in terms of the social groups of these individuals.

Most existing Bayesian approaches to network modeling focus on elucidating this underlying structure. While the resulting models are useful for characterization, they tend to perform poorly in terms of prediction of new interactions.

Prediction is an important modeling goal: For example we are often interested in who a person in a social network will interact with next. This may involve out-of-sample prediction: my next email may be to someone who is not represented in the training set. In this work, we focus on predictive models for networks, where both the number of nodes and the dimensionality of the latent structure are unbounded. We demonstrate that we can capture latent structure while maintaining predictive power, and discuss possible extensions.

*Machine Learning External Seminar, Gatsby Unit, May 20, 2015.