

Machine Learning for Predicting Legislative Votes: The Ideal Point Topic Model

Sean Gerrish

2/22/2011

Quantitative political scientists often study roll-call data, historical records of legislators' votes on a set of issues. Until recently, models of roll-call data could not predict votes on entirely new bills – those on which no votes have been cast. In this talk, I will describe a model for predicting Senate votes using the text of bills. Themes I will cover include ideal point models, for representing senators and bills in a latent political space; and topic models, for finding textual themes in collections of documents. I will also describe variational inference, which we use to fit an approximate posterior to the observations.

No prior knowledge of these areas is necessary: just come with a healthy interest in probability and applying machine learning to real-world datasets.