

Foundations of Graphical Models: Homework 2

Due: F 2018-10-26

The total page limit is three pages, though you may use extra pages for figures, and your code can be any length. Please use the \LaTeX template on the website. You should submit both the writeup and code to Courseworks.

Problem

Implement either a mixed-membership model or a matrix factorization model. You can use any inference method you'd like. For example, you can implement variational inference for latent Dirichlet allocation or MAP estimation for matrix factorization.

Apply your code to data and discuss what you learned. We encourage you to apply your code to data sets that aren't widely used in machine learning. We have also included data sets on the course website.

You can plot and discuss whatever you'd like. Among the plots, we would like to see a figure depicting the convergence of your inference procedure (for example, if you are performing variational inference you can plot the ELBO as a function of iteration). We would also like to see a plot that checks the model on held-out data.

Online Data Sets

MovieLens One of the online data sets contains movie ratings submitted by users on MovieLens, a movie recommendation website. The data set contains 100,000 ratings applied to 9,000 movies by 600 users. There's a `README.txt` file in the folder that contains more information about the data set (you can alternatively read an [online version here](#)).

AP The other provided data set contains the text of 2,246 articles from the Associated Press. After unzipping, you'll find three files in the folder `ap`: `ap.dat`, `ap.txt`, and `vocab.txt`.

`ap.txt` is an XML file that contains the full text of every article.

You'll probably want to work with `ap.dat`, which contains the counts of each word for each article. Every line is a different article in a bag-of-words format. The first number in each line is the total number of words in that article. Following this number, the rest of the line contains word counts in the format `word_index:count`. For example, the line "5 0:1 5:2 140:1 2031:1" indicates a 5-word article that has 1 occurrence of word 0, 2 occurrences of word 5, 1 occurrences of word 140, and 1 occurrence of word 2031.

Finally, the file `vocab.txt` contains a list of each word, zero-indexed to match the indices in `ap.dat`.