Use PLSA for Text Mining

- PLSA would be able to generate
  - Topic coverage in each document: p(π_d = j)
  - Word distribution for each topic: p(w|θ_j)
  - Topic assignment at the word level for each document
  - The number of topics must be given in advance
- These probabilities can be used in many different ways
  - θ_j naturally serves as a word cluster
  - $\pi_{d,j}$ can be used for document clustering $j^* = \arg\max_j \pi_{d,j}$
  - Contextual text mining: Make these parameters conditioned on context, e.g.,
    - p(θ_j |time), from which we can compute/plot p(time| θ_j)
    - p(θ_j |location), from which we can compute/plot p(loc| θ_j)