

Latent Dirichlet Allocation

David Mareček

📅 October 26, 2022



EUROPEAN UNION
European Structural and Investment Fund
Operational Programme Research,
Development and Education


Charles University
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

Many of the slides in this presentation were taken from the presentations of Carl Edward Rasmussen (University of Cambridge)

Mixture of Categoricals Model




()

| ()

With the Expectation-Maximization algorithm we have essentially estimated θ and β_k by maximum likelihood.

Bayesian Mixture of Categoricals Model




An alternative, Bayesian treatment infers these parameters starting from priors, e.g.:

- (α) is a symmetric Dirichlet over category probabilities,
- (β_k) are symmetric Dirichlets over vocabulary probabilities.

What is different?

- We no longer want to compute a point estimate of θ and w_{nd} .
- We are now interested in computing posterior distributions.

Limitations of the mixture of categoricals model



A generative view of the mixture of categoricals model:

1. Draw a distribution over topics from a (θ) .
2. For each topic z_d , draw a distribution over words from a (β_k) .
3. For each document d , draw a topic z_d from a (θ) .
4. For each document d , draw words w_{nd} from a (β_k) .

Limitations:

- All words in each document are drawn from one specific topic distribution.
- This works if each document is exclusively about one topic, but if some documents span more than one topic, then “blurred” topics must be learnt.

Latent Dirichlet Allocation

Topics

gene	0.04
dna	0.02
genetic	0.01
...	

life	0.02
evolve	0.01
organism	0.01
...	

brain	0.04
neuron	0.02
nerve	0.01
...	

data	0.02
number	0.02
computer	0.01
...	

Documents

Seeking Life's Bare (Genetic) Necessities

COLD SPRING HARBOR, NEW YORK—How many **genes** does an **organism** need to **survive**? Last week at the genome meeting here,* two genome researchers with radically different approaches presented complementary views of the basic genes needed for **life**. One research team, using **computer** analyses to compare known **genomes**, concluded that today's **organisms** can be sustained with just 250 genes, and that the earliest life forms required a mere 128 **genes**. The other researcher mapped genes in a simple parasite and estimated that for this organism, 800 genes are plenty to do the job—but that anything short of 100 wouldn't be enough. Although the numbers don't match precisely, those **predictions**


"are not all that far apart," especially in comparison to the 75,000 **genes** in the human genome, notes Siv Andersson, a professor at Uppsala University in Sweden, who arrived at the 800 number. But coming up with a consensus answer may be more than just a **genetic numbers game**, particularly as more and more **genomes** are completely mapped and sequenced. "It may be a way of organizing any newly **sequenced genome**," explains Arcady Mushegian, a **computational** molecular biologist at the National Center for Biotechnology Information (NCBI) in Bethesda, Maryland. Comparing an

Stripping down. Computer analysis yields an estimate of the minimum modern and ancient genomes.


* Genome Mapping and Sequencing, Cold Spring Harbor, New York, May 8 to 12.

SCIENCE • VOL. 272 • 24 MAY 1996

Topic proportions and assignments




Latent Dirichlet Allocation: what we observe



In reality, we only observe the documents.

The other structure are hidden variables.

Latent Dirichlet Allocation: what we observe

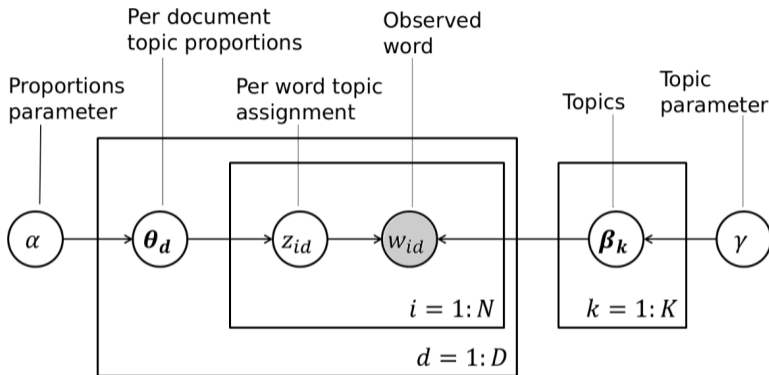


Our goal is to infer the hidden variables.

This means computing their distribution conditioned on the documents

$$\left(\theta_1, \theta_2, \dots, \theta_K \mid \text{document} \right)$$


Latent Dirichlet Allocation: graphical model



Nodes are random variables; edges indicate dependence.

Shaded nodes indicate observed variables.

Mixture of Categoricals vs. LDA



A generative view of LDA:

1. For each document draw a distribution over topics from a (α) .
2. For each topic draw a distribution over words from a (γ) .
3. Draw a topic for the n -th word in document d from a (θ_d) .
4. Draw word w_{nd} from a (β_k) .

Differences with the mixture of categoricals model:

- In LDA, every word in a document can be drawn from a different topic,
- and every document has its own distribution over topics θ_d .

Gibbs sampling algorithm

- Initialize θ randomly for all words in all documents
- Choose random word and sample a new category based on all other words in all other documents.
- The distribution across categories is the predictive distribution of the posterior Dirichlet distribution (integration across all possible θ and ϕ).
- Perform these small changes in many iterations across the data. The algorithm will converge to a good solutions.