Lecture 8: Graphical models for Text

4F13: Machine Learning

Joaquin Quiñonero-Candela and Carl Edward Rasmussen

Department of Engineering
University of Cambridge

http://mlg.eng.cam.ac.uk/teaching/4f13/

A really simple document model

Consider a collection of D documents with dictionary of M unique words.

- N_d: number of (non-unique) words in document d.
- w_{id}: i-th word in document d ($w_{id} \in \{1 : M\}$).
- $\boldsymbol{\beta} = [\beta_1, \ldots, \beta_M]^\top$: parameters of a Multinomial distribution over the dictionary of M unique words.

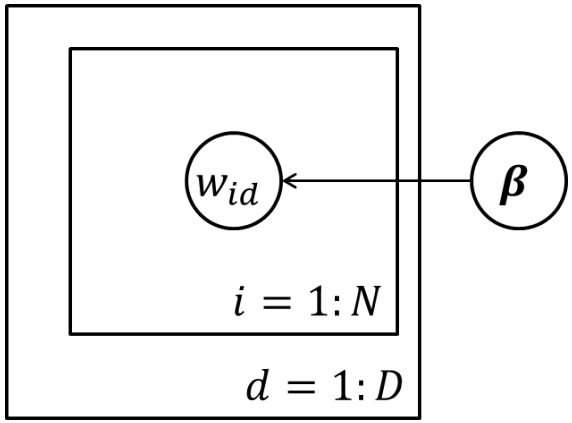

We can fit $\boldsymbol{\beta}$ by maximising the likelihood:

$$\hat{\boldsymbol{\beta}} = \operatorname{argmax} \prod_{d=1}^{D} \text{Mult}(c_{1d}, \ldots, c_{Md} | \boldsymbol{\beta}, N_d)$$

$$= \operatorname{argmax} \text{Mult}(c_1, \ldots, c_M | \boldsymbol{\beta}, N)$$

$$\boxed{\hat{\beta}_j = \frac{c_j}{\sum_{l=1}^{M} c_l}}$$

- $N = \sum_{d=1}^{D} N_d$: total number of (non-unique) words in the collection.
- c_{jd}: count of occurrences of unique word j in document d.
- $c_j = \sum_{d=1}^{D} c_{jd}$: count of total occurrences of unique word j in the collection.

Limitations of the really simple document model

- Document d is the result of sampling $\mathrm{N_d}$ words from the Multinomial $\boldsymbol{\beta}$.
- $\boldsymbol{\beta}$ estimated by maximum likelihood reflects the aggregation of all documents.
- All documents are therefore modelled by the global word frequency distribution.
- The generative model wastes mass, because it cannot specialize.
- All unique words do not necessarily *co-occur* in a given document.
- It possible that documents might be about different *topics*.

A mixture of Multinomials model

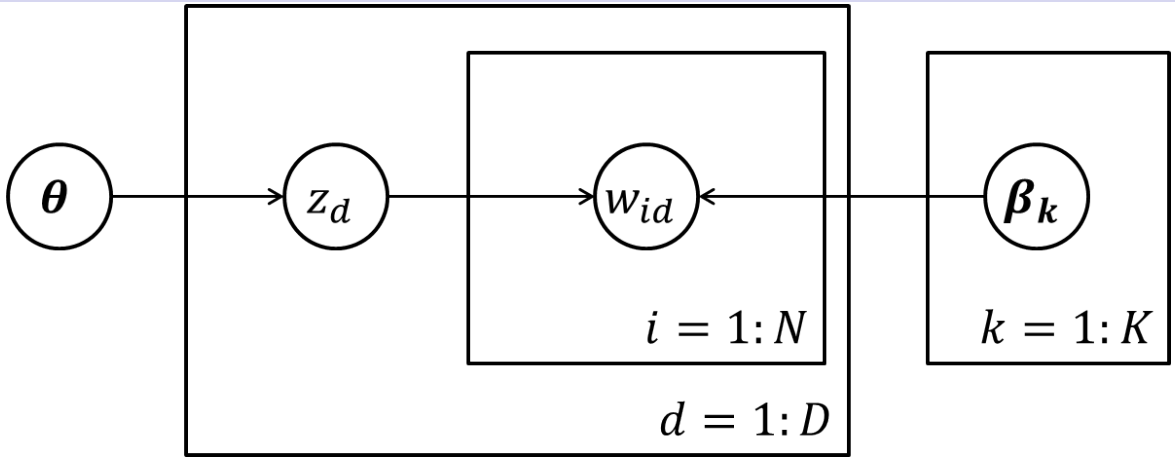

We want to allow for a mixture of K Multinomials parametrised by $\beta_1, \ldots, \beta_K$. Each of those Multinomials corresponds to a *document category*.

- $z_d \in \{1 : K\}$ assigns all words in document d to one of the K categories.
- $\theta_j = p(z_d = j)$ is the probability any document d is assigned to category j.
- $\theta = [\theta_1, \ldots, \theta_K]$ is also the parameter of a Multinomial over the K categories.

We have introduced a new set of *hidden* variables z_d.

- How do we fit those variables? What do we do with them?
- We are actually not interested in them: We are only interested in θ and β.

Jensen's Inequality

For any concave function, such as $\log(x)$

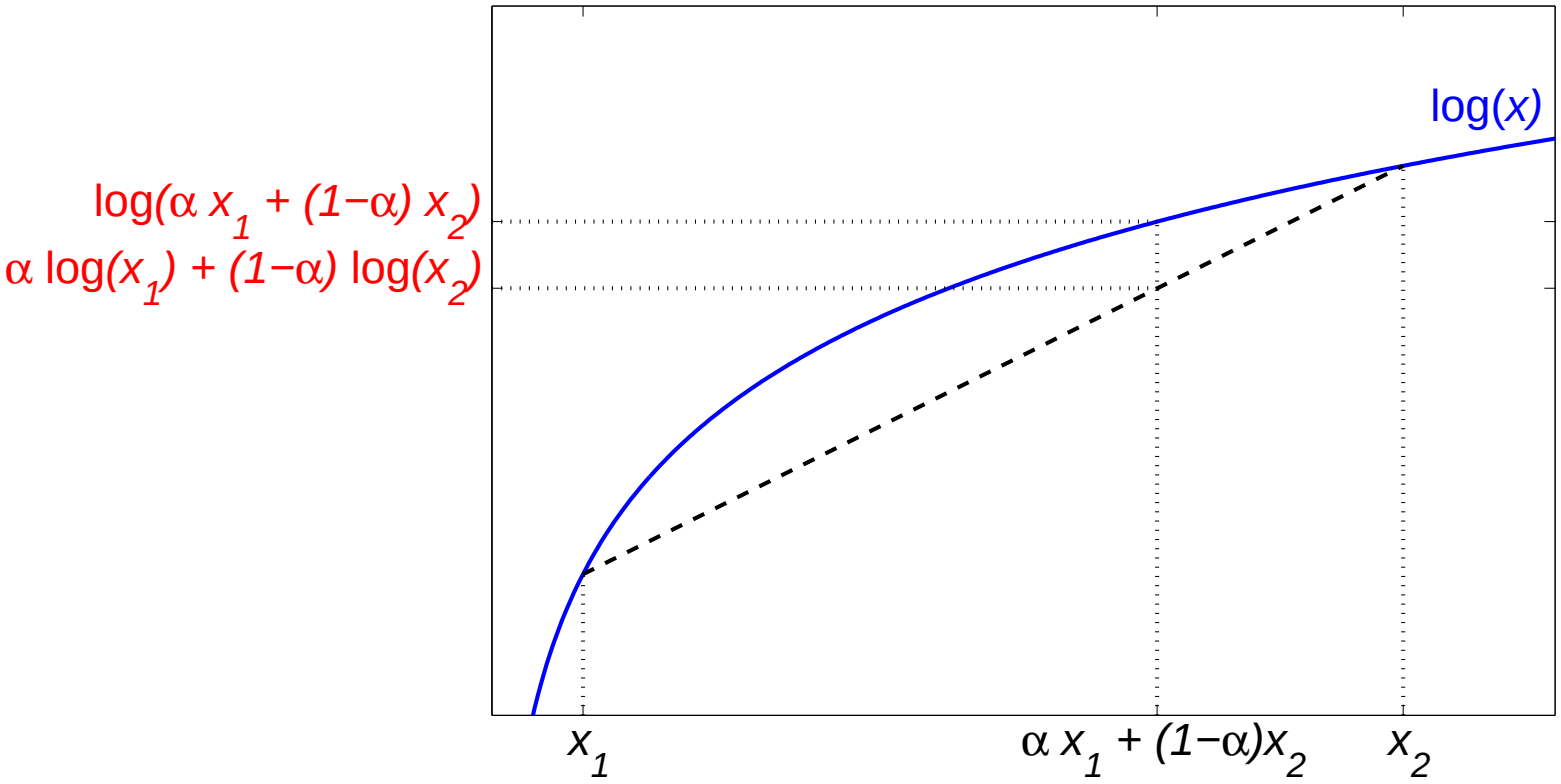

For $\alpha_i \geqslant 0$, $\sum \alpha_i = 1$ and any $\{x_i > 0\}$

$$\log\left(\sum_i \alpha_i x_i\right) \geqslant \sum_i \alpha_i \log(x_i)$$

Equality if and only if $\alpha_i = 1$ for some i (and therefore all others are 0).

The Expectation Maximization (EM) algorithm

Given a set of observed (visible) variables V, a set of unobserved (hidden / latent / missing) variables H, and model parameters θ, optimize the log likelihood:

$$\mathcal{L}(\theta) = \log p(V|\theta) = \log \int p(H, V|\theta) dH, \tag{1}$$

where we have written the marginal for the visibles in terms of an integral over the joint distribution for hidden and visible variables.

Using *Jensen's inequality* for any distribution of hidden states $q(H)$ we have:

$$\mathcal{L} = \log \int q(H) \frac{p(H, V|\theta)}{q(H)} dH \geqslant \int q(H) \log \frac{p(H, V|\theta)}{q(H)} dH = \mathcal{F}(q, \theta), \tag{2}$$

defining the $\mathcal{F}(q, \theta)$ functional, which is a lower bound on the log likelihood.

In the EM algorithm, we alternately optimize $\mathcal{F}(q, \theta)$ wrt q and θ, and we can prove that this will never decrease $\mathcal{L}$.

The E and M steps of EM

The lower bound on the log likelihood:

$$\mathcal{F}(q,\theta) = \int q(H)\log\frac{p(H,V|\theta)}{q(H)}dH = \int q(H)\log p(H,V|\theta)dH + \mathcal{H}(q), \quad (3)$$

where $\mathcal{H}(q) = -\int q(H)\log q(H)dH$ is the entropy of q. We iteratively alternate:

E step: optimize $\mathcal{F}(q,\theta)$ wrt the distribution over hidden variables given the parameters:

$$q^{(k)}(H) := \underset{q(H)}{\operatorname{argmax}}\ \mathcal{F}\left(q(H),\theta^{(k-1)}\right). \quad (4)$$

M step: maximize $\mathcal{F}(q,\theta)$ wrt the parameters given the hidden distribution:

$$\theta^{(k)} := \underset{\theta}{\operatorname{argmax}}\ \mathcal{F}\left(q^{(k)}(H),\theta\right) = \underset{\theta}{\operatorname{argmax}}\ \int q^{(k)}(H)\log p(H,V|\theta)dH, \quad (5)$$

which is equivalent to optimizing the expected complete-data likelihood $p(H,V|\theta)$, since the entropy of $q(H)$ does not depend on θ.

EM as Coordinate Ascent in $\mathcal{F}$

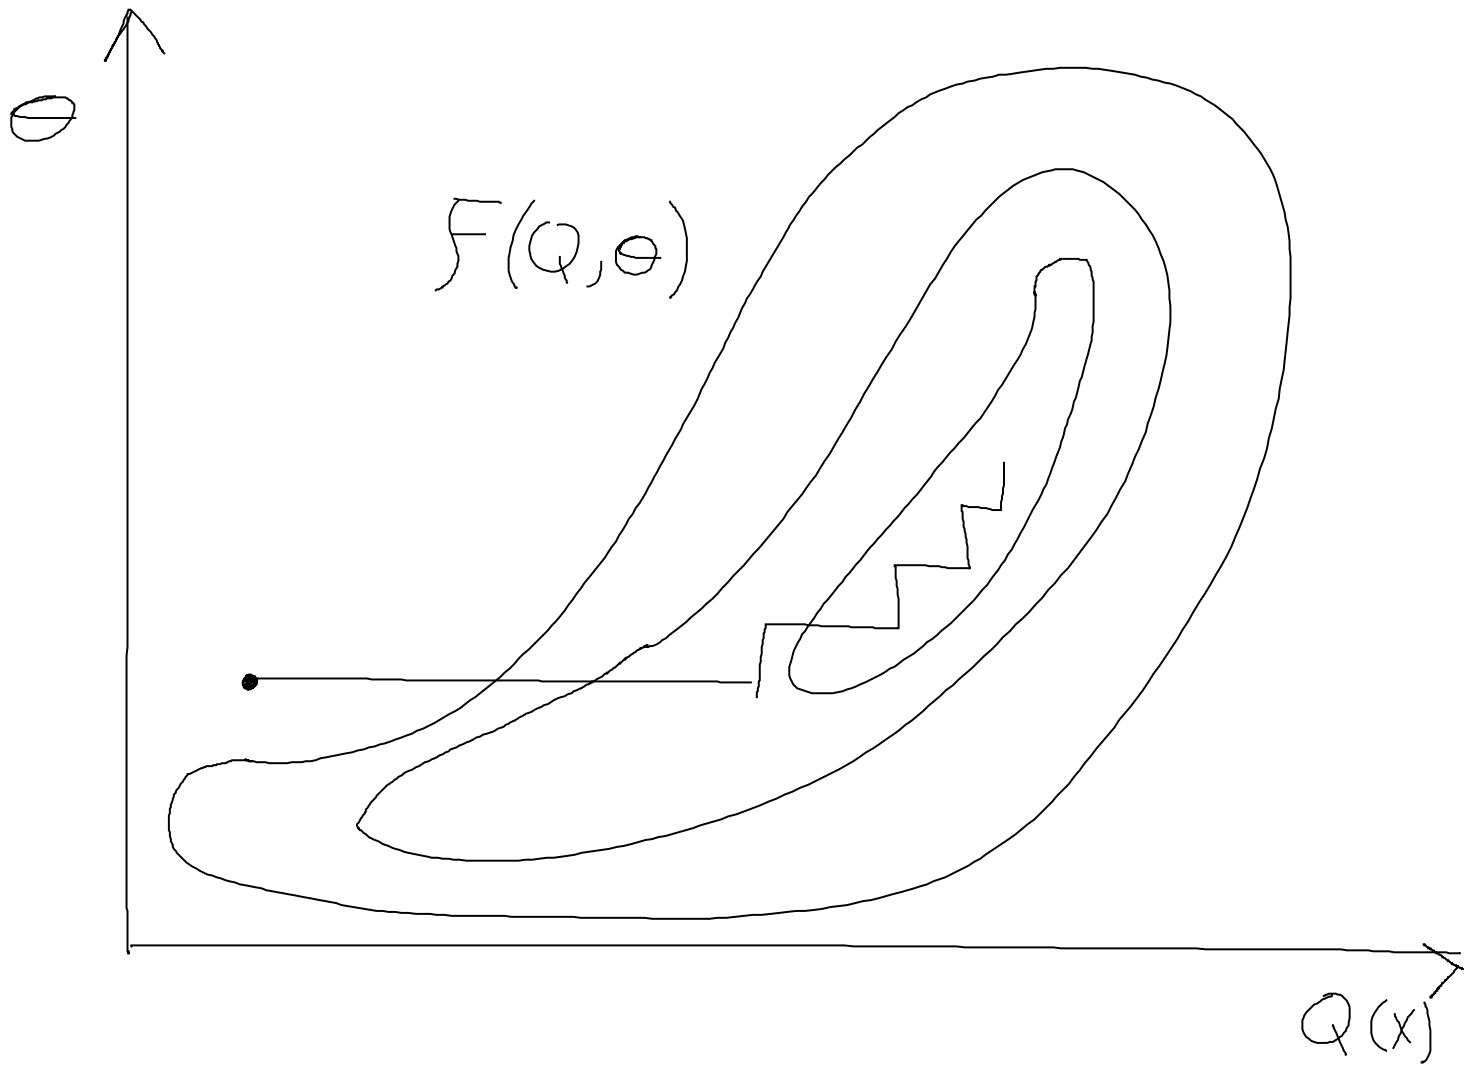

The EM algorithm never decreases the log likelihood

The difference between the cost functions:

$$\begin{aligned}\mathcal{L}(\theta) - \mathcal{F}(q,\theta) &= \log p(V|\theta) - \int q(H) \log \frac{p(H,V|\theta)}{q(H)} dH \\ &= \log p(V|\theta) - \int q(H) \log \frac{p(H|V,\theta)p(V|\theta)}{q(H)} dH \\ &= -\int q(H) \log \frac{p(H|V,\theta)}{q(H)} dH = \mathcal{KL}\big(q(H), p(H|V,\theta)\big),\end{aligned}$$

is called the Kullback-Liebler divergence; it is non-negative and only zero if and only if $q(H) = p(H|V,\theta)$ (thus this is the E step). Although we are working with the wrong cost function, the likelihood is still increased in every iteration:

$$\mathcal{L}\big(\theta^{(k-1)}\big) \underset{\text{E step}}{=} \mathcal{F}\big(q^{(k)}, \theta^{(k-1)}\big) \underset{\text{M step}}{\leqslant} \mathcal{F}\big(q^{(k)}, \theta^{(k)}\big) \underset{\text{Jensen}}{\leqslant} \mathcal{L}\big(\theta^{(k)}\big),$$

where the first equality holds because of the E step, and the first inequality comes from the M step and the final inequality from Jensen. Usually EM converges to a local optimum of $\mathcal{L}$ (although there are exceptions).

EM and Mixtures of Multinomials

The the Mixture model for text, the latent variables are

$$z_d \in \{1, \ldots, K\}, \quad \text{where} \quad d = 1, \ldots, D$$

which for each document encodes which mixture component generated it.

E-step: for each document d, set q to the posterior

$$q_d(z_d) \propto p(z_d = k|\theta) \prod_{i=1}^{N_d} p(w_i|\beta_{w_i k}) = \theta_k \text{Mult}(c_{1d}, \ldots, c_{Md}|\boldsymbol{\beta}_k, N_d) = r_{kd}$$

M-step: Maximize

$$\begin{aligned}
\sum_{k=1}^{K} q_d(z_d = k) \log p(\{w_{id}\}, z_d) &= \sum_k r_{kd} \log \prod_{d=1}^{D} \prod_{i=1}^{N_d} p(w_i|\beta_{w_i k}) p(z_d = k) \\
&= \sum_k r_{kd} \Big(\sum_{d=1}^{D} \log \prod_{j=1}^{M} \beta_{jk}^{c_{jd}} + \log \theta_k \Big) \\
&= \sum_{k,d} r_{kd} \Big(\sum_{j=1}^{M} c_{jd} \log \beta_{jk} + \log \theta_k \Big) = F(\theta, \beta)
\end{aligned}$$

EM: M step for mixture model

Need Lagrange multipliers to constrain the maximization and ensure proper distributions.

$$\theta_k = \operatorname{argmax} F(\theta, \beta) + \lambda(1 - \sum_{k=1}^{K} \theta_k)$$

$$= \frac{\sum_{d=1}^{D} r_{kd}}{\sum_{k'=1}^{K} \sum_{d=1}^{D} r_{k'd}}$$

$$\beta_{jk} = \operatorname{argmax} F(\theta, \beta) + \sum_{j=k}^{K} \lambda_k(1 - \sum_{j=1}^{M} \beta_{jk})$$

$$= \frac{\sum_{d=1}^{D} r_{kd} c_{jd}}{\sum_{j'=1}^{M} \sum_{d=1}^{D} r_{kd} c_{j'd}}$$

A Bayesian mixture of Multinomials model

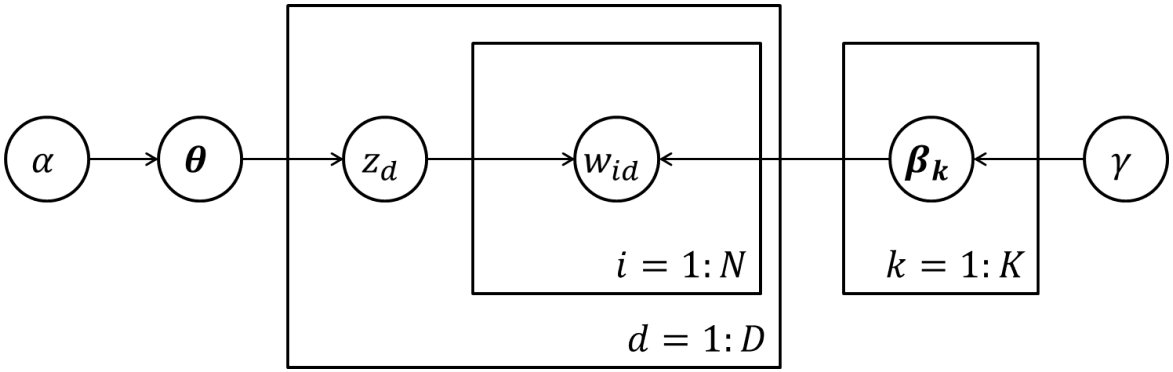

With the EM algorithm we have essentially estimated $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ by maximum likelihood. An alternative, Bayesian treatment is to introduce *hyperpriors*.

- $\theta \sim \text{Dir}(\alpha)$ is a symmetric Dirichlet over category probabilities.
- $\boldsymbol{\beta}_k \sim \text{Dir}(\gamma)$ is a symmetric Dirichlet over unique word probabilities.

What is different?

- We no longer want to compute a point estimate of θ or $\boldsymbol{\beta}$.
- We are now interested in computing the *posterior* distributions.