6. Representing an HMM as an exponential family

Def. 16, Sec. 1 $\Rightarrow$ the joint p.d. for a Markov chain model with strictly positive prob's can be written as

$$p(s) = p(s_1, \ldots, s_n) = \frac{1}{Z} \prod_{i=2}^{n} g_i(s_{i-1}, s_i) = \frac{1}{Z} \exp\left[\sum_{i=2}^{n} u_i(s_{i-1}, s_i)\right]$$

Remark 1 The factors g_i, resp. the potentials u_i define the model uniquely. The inverse is not true.

Remark 2 The normalising factor Z is defined by

$$Z(u) = \sum_{s \in K^n} \exp\left[\sum_{i=2}^{n} u_i(s_{i-1}, s_i)\right]$$

and can be computed by an algorithm similar to the one described in Sec. 3

Denote:

- $\vec{\varphi}_i(s_i) \in \mathbb{R}^K$ a binary valued indicator vector, which denotes the state s_i („one out of K" coding), i.e.

$$\vec{\varphi}_i(s_i = k) = (0, \ldots, \underset{\substack{\uparrow \\ \text{pos. } k}}{1}, \ldots 0)$$

- U_i denotes the $K \times K$ matrix with values $u_i(s_{i-1}, s_i)$

Then, the joint p.d. can be written as

$$p(s) = \frac{1}{Z(u)} \exp \sum_{i=2}^{n} \langle \vec{\varphi}_{i-1}(s_{i-1}), U_i \cdot \vec{\varphi}_i(s_i) \rangle$$

$$= \frac{1}{Z(u)} \exp \sum_{i=2}^{n} \langle \varphi_i(s_{i-1}, s_i), U_i \rangle$$

where

$\varphi_i(s_{i-1}, s_i) = \vec{\varphi}_{i-1}(s_{i-1}) \otimes \vec{\varphi}_i(s_i)$ is a $K \times K$ binary valued indicator matrix and

$\langle \Phi, U \rangle = tr(\Phi^T U)$ denotes the Frobenius inner product.

Finally, denote $\Phi = (\varphi_2, \varphi_3, \ldots, \varphi_n)$ and $U = (U_1, U_2, \ldots, U_n)$.

The joint p.d. of a Markov chain model can be written as

$$p(s) = \frac{1}{Z(U)} \exp \langle \Phi(s), U \rangle$$

The joint p.d. of an HMM can be written as

$$p(s) = \frac{1}{Z(U)} \exp \langle \Phi(x, s), U \rangle$$

by using similar notations.

7. Supervised learning, ML-estimator

Given an i.i.d. sample of pairs of sequences

$$\mathcal{T} = \{(x^j, s^j) \mid x^j \in F^n,\ s^j \in K^n,\ j = 1, \dots, \ell\},$$

estimate the model parameters of the HMM by the maximum likelihood estimator

$$u^* \in \arg\max_u \prod_{(x,s)\in\mathcal{T}} p_u(x,s)$$

$$= \arg\max_u \frac{1}{|\mathcal{T}|} \sum_{(x,s)\in\mathcal{T}}' \log p_u(x,s)$$

i.e. find optimal $u_i^*(s_{i-1}, s_i)$, $\tilde{u}_i^*(x_i, s_i)$, or, equivalently, $p(s_{i-1}, s_i)$, $p(x_i, s_i)$

Intuitive answer u^* is given by

$$p_{u^*}(s_{i-1}, s_i) = \beta(s_{i-1}, s_i)$$
$$p_{u^*}(x_i, s_i) = \beta(x_i, s_i)$$

where β-s denote frequencies of corresponding events in $\mathcal{T}$.
Let us prove correctness. Log-likelihood of $\mathcal{T}$ is

$$L(u) = \frac{1}{|\mathcal{T}|} \sum_{(x,s)\in\mathcal{T}}' \left[\langle \varphi(x,s), u \rangle - \log Z(u) \right]$$

$$= \langle \psi, u \rangle - \log Z(u),$$

where

$$\psi = \mathbb{E}_{\mathcal{T}}(\varphi) = \frac{1}{|\mathcal{T}|} \sum_{(x,s)\in\mathcal{T}}' \varphi(x,s)$$

Lemma 1 The log-partition function $\log Z(u)$ of an HMM (with strictly positive p.d.) is convex in u.

Proof

$$\nabla \log Z(u) = \frac{1}{Z(u)} \sum_{x,s} \exp\langle \varphi(x,s), u\rangle \, \varphi(x,s) \overset{!}{=} \mathbb{E}_u(\varphi)$$

The components of $\mathbb{E}_u(\varphi)$ are the pairwise marginal prob's on the edges of the model.

$$\nabla^2 \log Z(u) = \mathbb{E}_u(\varphi \otimes \varphi) - \mathbb{E}_u(\varphi) \otimes \mathbb{E}_u(\varphi)$$

$$= \mathbb{E}_u\left[(\varphi - \mathbb{E}_u(\varphi)) \otimes (\varphi - \mathbb{E}_u(\varphi))\right]$$

The expectation of a positive semidefinite matrix is p.s.d. $\Rightarrow$ $\log Z(u)$ is convex. ∎

The log likelihood $L(u)$ is concave as a consequence and has global maxima only, which are given by

$$\nabla L(u^*) = \frac{1}{|\mathcal{T}|} \sum_{(x,s)\in\mathcal{T}} \varphi(x,s) - \mathbb{E}_{u^*}(\varphi)$$

$$= \mathbb{E}_{\mathcal{T}}(\varphi) - \mathbb{E}_{u^*}(\varphi) = 0$$

Recall that the components of $\mathbb{E}_u(\varphi)$ are the pairwise marginal prob's of the model $p_u(x,s)$. Hence, the optimiser u^* defines the model that has precisely the same pairwise marginal prob's as the empirical marginal frequencies of $\mathcal{T}$.

The concavity of $L(u)$ also ensures the consistency of the estimator

Theorem 1 (w/o proof)
The maximum likelihood estimator for HMMs is consistent, i.e.

$$\mathbb{P}_u \left(\| u^*(\Upsilon) - u \| > \varepsilon \right) \xrightarrow{|\Upsilon| \to \infty} 0$$

holds for every $\varepsilon > 0$.