EM algorithm for coin flipping problem

Cheng Li

We have K coins. The mixing proportions are $w_1, w_2, \ldots, w_K$. The probability of the k's coin getting heads is p_k. First we randomly pick a coin, then we flip this coin for D times. And we repeat this process for N rows. In this way, we can generate N data points $\{x_1, x_2, \ldots, x_N\}$, each of which is a D dimensional vector. $x_i \in \{Head, Tail\}^d$. Suppose the number of heads in x_i is H_i. We define the hidden variables z_{ik}, representing the component assignment for data point x_i using a vector of size K. If x_i is drawn from the kth component, $z_{ik} = 1$ while the remaining are all 0.

- E step: Compute $< z_{ik} >$ with current parameters $\theta = \{p_k, w_k\}$.

$$< z_{ik} >= prob(z_{ik} = 1 | x_n, \theta) = \frac{w_k \cdot pdf(x_i | p_k)}{\sum_{j=1}^{K} w_j \cdot pdf(x_i | p_j)}$$

- M step: update w_k and p_k

$$w_k = \frac{\sum_{i=1}^{N} < z_{ik} >}{N}$$

$$p_k = \frac{\sum_{i=1}^{N} < z_{ik} > H_i}{\sum_{i=1}^{N} < z_{ik} > D}$$

To match coind H/T distribution, you will need to use binomial distributions for $pdf(x | bias, D)$

References

[1] Dawen Liang, Technical Details about the Expectation Maximization (EM) Algorithm, 2012