Supplementary material for AISTATS submission: Factorial HMM with Collapsed Gibbs Sampling for optimizing long-term HIV Therapy

Amit Gruber[1,†], Chen Yanover[1], Tal El-Hay[1], Anders Sönnerborg[2]
Vanni Borghi[3], Francesca Incardona[4], Yaara Goldschmidt[1]
[1]IBM Research Haifa Research Lab
[2]Karolinska Institute Karolinska University Hospital
[3]Modena University Hospital
[4]EuResist Network GEIE InformaPro S.r.l.

A Appendix: Collapsed Gibbs Sampling for Factorial HMM

Let $q_{t,m}, t = 1 \dots T, m = 1 \dots M$ be the hidden state of the m'th chain of a factorial HMM with M chains at time t and let O_t be the observation at that time. At each step, a single state $q_{t,m}$ is sampled given all the other states. Let $C^m_{q,q'}$ be the count of states q and q' at times t and $t+1$ respectively in chain m excluding the variable to be sampled, and let $C_{q1,\dots,q_M,o}$ the count of the combination of states $q_1, \dots q_m$ with outcome o (a potentially exponentially large table). Let α and β are Symmetric Dirichlet priors on the chain specific transition matrix and the emission matrix respectively. The sampling equations for $1 < t < T$ are:

$$\Pr(q_{t,m}|q_{-(t,m)}, o_t, \alpha, \beta) \propto \left[\frac{C^m_{q_{t,m},q_{t-1,m}} + \alpha}{\sum_q C^m_{q,q+t-1,m} + Q\alpha}\right] \quad (1)$$
$$\left[\frac{C^m_{q_{t+1,m},q_{t,m}} + \alpha}{\sum_q C^m_{q+t-1,m,q} + Q\alpha}\right]\left[\frac{C_{q_{t,\cdot,},O_t} + \beta}{\sum_o C_{qt,\cdot,o} + O\beta}\right]$$

where Q is the number of states and O is the number of possible outcomes. For $t = 1$ the first term is replaced with 1 and for $t = T$ the second term is replaced with 1.

Computationally, this algorithm is practical only when the dependency between O_t and the states $q_{t,\cdot}$ has some compact form, for example when O_t depends on only a few of these states at a time like in FResist or in [Dong et al., 2012], or when this dependence has some functional form (like in FResist).

The parameters $\Pr(q_{t+1}|q_t)$ and $\Pr(o_t|q^1_t \dots q^M_t)$ can be estimated by:

$$\Pr(q_{t+1}|q_t, m) = \frac{C^m_{q_{t+1},q_t} + \alpha}{\sum_q C^m_{q',q_t} + Q\alpha} \quad (2)$$

$$\Pr(O|q_1 \dots q_M) = \frac{C_{q_{t,\cdot,},O} + \beta}{\sum_o C_{qt,\cdot,o} + O\beta} \quad (3)$$

[†]Current affiliation: Facebook Inc.

References

[Dong et al., 2012] Dong, W., Heller, K., and Pentland, A. (2012). Graph-coupled hmms for modeling the spread of infection. In *Uncertainty in Artificial Intelligence*.