Notes on EM

Jiding Zhang

Reference:

- EM: `http://cs229.stanford.edu/notes2020spring/cs229-notes8.pdf`
- CCP-EM: Arcidiacono and Miller [2011]

1 Standard EM

- objective: maximize the likelihood $\log p(x;\theta)$
- with latent variables z, we maximize:

$$l(\theta) = \log \sum_z p(x,z;\theta)$$

- not easy to maximize the **log of sums**
- instead, easy to maximize the **sum of logs**
- From Jensen's inequality:

$$\log p(x;\theta) = \log \sum_z Q(z) \frac{p(x,z;\theta)}{Q(z)} \geq \underbrace{\sum_z Q(z) \log \frac{p(x,z;\theta)}{Q(z)}}_{ELBO(x;Q,\theta)}$$

- equality takes place when $Q(z) = p(z|x;\theta)$
- also note,

$$\sum_z Q(z) \log \frac{p(x,z;\theta)}{Q(z)} = \sum_z Q(z) \log p(x,z;\theta) - \underbrace{\sum_z Q(z) \log Q(z)}_{\text{not a function of } \theta \text{ if } Q(z) \text{ fixed}}$$

- E-step: compute $Q(z)$: the posterior distribution of the latent variable
- M-step: fix $Q(z)$ and view that as a *weight* (i.e., not a function of θ). Maximize w.r.t. θ the sum of logs:

$$\sum_z Q(z) \log p(x,z;\theta) = \mathbb{E}_{z|x;\theta^{\text{old}}} \underbrace{\log p(x,z;\theta)}_{\text{complete data log likelihood}}$$

2 CCP-EM

- individual n, period t, with choice $d_{nt} = j \in \{1, \ldots, J\}$, observable state x_{nt}, unobservable state s_{nt}
- likelihood of observing data: action and the observable state

$$\mathcal{L}_t(d_{nt}, x_{n,t+1}|x_{nt}, s_{nt}; \theta, \pi, p) = \prod_j \left[l_{jt}(x_{nt}, s_{nt}, \theta, \pi, p) f_{jt}(x_{n,t+1}|x_{nt}, s_{nt}, \theta) \right]^{d_{jnt}}$$

- E-step:
  - update $q_{nst}^{(m)}$, the conditional probability of latent variable in state s in t
    - This is exactly $p(z|x;\theta)$, but the dynamic nature means that we should compute the posterior distribution of s in each period
    - posterior:

$$p(z|x;\theta) = \frac{p(z,x;\theta)}{\sum_z p(x|z;\theta)p(z;\theta)}$$

    - denominator:

$$\begin{aligned} L_n &= L(d_n, x_n|x_{n1}; \theta, \pi, p) \\ &= \underbrace{\sum_{s_1} \cdots \sum_{s_T}}_{\text{integrate over } z,\ S^T} \underbrace{\pi(s_1|x_{n1}) \left(\prod_{t=2}^{T} \pi(s_t|s_{t-1}) \right)}_{\text{density of } z} \underbrace{\left(\prod_{t=1}^{T} \underbrace{\mathcal{L}_t(d_{nt}, x_{n,t+1}|x_{nt}, s_t; \theta, \pi, p)}_{\mathcal{L}_{nt}(s_t)} \right)}_{p(x|z;\theta)} \end{aligned}$$

    - numerator:

$$\begin{aligned} &L_n(s_{nt} = s) \\ &= \sum_{s_{t'}: t' \neq t} \pi(s_1|x_{n1}) \left(\prod_{t'=2, t' \neq t, t+1}^{T} \pi(s_{t'}|s_{t'-1}) \right) \left(\prod_{t'=1, t' \neq t, t+1}^{T} \mathcal{L}_{nt'}(s_{t'}) \right) \pi(s|s_{t-1}) \mathcal{L}_{nt}(s) \pi(s_{t+1}|s) \mathcal{L}_{n,t+1}(s_{t+1}) \end{aligned}$$

    - If transitions of latent variables are i.i.d., i.e., $\pi(s) = \pi(s|s'), \forall s'$, and independent of observable states and actions, then all other period actions and state transitions are irrelevant and hence canceled out in the ratio. $q_{nst}^{(m)}$ could be simplified:

$$q_{nst}^{(m)} = \frac{\pi(s_t = s)\mathcal{L}_{nt}(s_t = s)}{\sum_{s'} \pi(s_t = s')\mathcal{L}_{nt}(s_t = s')}$$

  - update $\pi^{(m)}$, transition probabilities on latent variables
  - update $p^{(m)}(x, s)$, the CCP
    - note the CCP is on both the observable and latent states
    - update based on logit choice probability given value functions

- M-step: maximize

$$\sum_n \underbrace{\sum_t \sum_s}_{T\times S} \sum_j q_{nst}^{(m+1)} \log \underbrace{\mathcal{L}_t(d_{nt}, x_{n,t+1}|x_{nt}, s_{nt} = s; \theta, \pi^{(m+1)}, p^{(m+1)})}_{\text{likelihood of observing both decision and observable state given latent value}}$$

  - note, technically $\mathcal{L}_t$ is a conditional probability rather than joint, but the probability of having s is not a function of the structural parameters θ, so it does not affect the maximization

References

Peter Arcidiacono and Robert A Miller. Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica*, 79(6):1823–1867, 2011.