Lecture 8, Wednesday, 4/13/2022

⊙ Sinkhorn algorithm

Let $a \in \mathcal{P}_m$, $b \in \mathcal{P}_n$, $C \in \mathbb{R}^{m \times n}$, $C \geq 0$.

$\mathcal{A}(a,b) = \{P \in \mathbb{R}^{m \times n} : P \geq 0,$ row sum of $P = a$, col. sum of $P = b\}$.

K-discrete OT: $\min_{P \in \mathcal{A}(a,b)} E[P]$, $E[P] = \langle P, C \rangle$.

Entropy regularized OT: $\min_{P \in \mathcal{A}(a,b)} E_\varepsilon[P]$.

$E_\varepsilon[P] = \langle P, C \rangle + \varepsilon \langle P(\log P - 1), \mathbf{1}_{m \times n} \rangle = \sum_{i,j} P_{ij} C_{ij} + \varepsilon \sum_{i,j} P_{ij}(\log P_{ij} - 1)$.

$\exists!\, P_\varepsilon = \arg\min_{\mathcal{A}(a,b)} E_\varepsilon$. $P_\varepsilon \to P_0 \in \mathcal{P}_c(a,b)$ with max. entropy.

Method of Lagrange multiplier

$$\mathcal{L}(P, \lambda) = \mathcal{L}(P, (f, g)) = E_\varepsilon[P] - \sum_{i=1}^{m} f_i \left(\sum_{j=1}^{n} P_{ij} - a_i\right) - \sum_{j=1}^{n} g_j \left(\sum_{i=1}^{m} P_{ij} - b_j\right).$$

$\frac{\partial \mathcal{L}}{\partial f_k} = 0 \Rightarrow \sum_j P_{kj} = a_k \ \forall k$. $\frac{\partial \mathcal{L}}{\partial g_\ell} = 0 \Rightarrow \sum_i P_{i\ell} = b_\ell \ \forall \ell$.

$\frac{\partial \mathcal{L}}{\partial P_{k\ell}} = 0 \quad C_{k\ell} - \varepsilon \log P_{k\ell} - f_k - g_\ell = 0 \quad \forall k, \ell$.

$$P_{k\ell} = e^{-(C_{k\ell} - f_k - g_\ell)/\varepsilon} = e^{-C_{k\ell}/\varepsilon} e^{-f_k/\varepsilon} e^{-g_\ell/\varepsilon}$$

Set $K_{k\ell} = K_{\varepsilon,k\ell} = e^{-C_{k\ell}/\varepsilon}$,
$u_k = u_{\varepsilon,k} = e^{-f_k/\varepsilon}$,
$v_k = v_{\varepsilon,\ell} = e^{-g_\ell/\varepsilon}$.

$$\left.\begin{aligned} P_{k\ell} &= K_{k\ell} u_k v_\ell \geq 0 \\ P &= \mathrm{diag}(u)\, K\, \mathrm{diag}(v) \end{aligned}\right\} \quad \mathrm{diag}(u) = \begin{bmatrix} u_1 & & 0 \\ & \ddots & \\ 0 & & u_m \end{bmatrix}$$

The problem to solve now becomes the following:
Given $K \in \mathbb{R}^{m \times n}$, $K > 0$. Find diagonal matrices $\mathrm{diag}(u)$ and $\mathrm{diag}(v)$ with $u \in \mathbb{R}^m$, $u > 0$ and $v \in \mathbb{R}^n$, $v > 0$, such that $P = \mathrm{diag}(u)\, K\, \mathrm{diag}(v) \in \mathcal{A}(a,b)$, i.e., row sum of $P = a$, col. sum of $P = b$.
(Note: We assume $a > 0$ and $b > 0$.)

$$\sum_{l=1}^{n} P_{kl} = a_k \Rightarrow \sum_{l=1}^{n} K_{kl} u_k v_l = u_k (Kv)_k = a_k$$

$$\Longrightarrow u \odot (Kv) = a, \quad \text{similarly } v \odot (K^T u) = b.$$

Remark strictly speaking, we should consider the inequality constraints $P_{ij} \geq 0 \ \forall i,j$ in the Lagrangian function: adding $\sum_{i,j} \mu_{ij} P_{ij}$. And use the KKT conditions. But, the complementary slackness: $\mu_{ij} P_{ij} = 0$ imply that $\mu_{ij} = 0$ since $P_{ij} > 0$ by the conditions $\frac{\partial \mathcal{L}}{\partial P_{ij}} = 0 \ \forall i,j$.

Theorem Assume $a > 0$ and $b > 0$. Then the unique minimizer $P_\varepsilon = \underset{\mathcal{A}(a,b)}{\operatorname{argmin}} E_\varepsilon[\cdot]$ is given by $P_\varepsilon = \operatorname{diag}(u) K \operatorname{diag}(v)$ for some $u \in \mathbb{R}_+^m$ and $v \in \mathbb{R}_+^n$ such that $P_\varepsilon \in \mathcal{A}(a,b)$. Moreover, (u,v) is unique up to a scaling factor.

Proof The characterization of P_ε follows from the calculations above using the method of Lagrange multipliers. See a theorem below for the uniqueness. QED.

Now, we know $K = e^{-C/\varepsilon} = K_\varepsilon$. We need to find $u = u_\varepsilon \in \mathbb{R}_+^m$ and $v = v_\varepsilon \in \mathbb{R}_+^n$ such that $P = \operatorname{diag}(u) K \operatorname{diag}(v) \in \mathcal{A}(a,b)$. This means $u = \frac{a}{Kv}$ and $v = \frac{b}{K^T u}$ (i.e., $u_i = a_i / (Kv)_i$, and $v_j = b_j / (K^T u)_j$).

Sinkhorn algorithm Initialize $v^{(0)}$ ($= \mathbb{1}_n$, e.g.)

$$u^{(k)} = \frac{a}{Kv^{(k-1)}}, \quad v^{(k)} = \frac{b}{K^T u^{(k)}} \quad (k = 1, 2, \cdots).$$

row sum of $u^{(k)} \odot Kv^{(k-1)} = a$, col. sum of $v^{(k)} \odot K^T u^{(k)} = b$.

Or: Initialize $u^{(0)} > 0$ and set

$$v^{(k)} = \frac{b}{K^T u^{(k-1)}} \text{ and } u^{(k)} = \frac{a}{Kv^{(k)}} \ (k = 1, 2, \cdots).$$

Some remarks

① The alternate iteration is equivalent to row-column alternate normalization (process) as initially studied in Richard Sinkhorn (1964) (only for $m = n$, $a_i = b_i = 1/n$).

Notation. $G \in \mathbb{R}^{m \times n}$:

$r_i(G)$ = row-i sum of G, $c_j(G)$ = col.-j sum of G, $\forall i, j$.

Definition Let $a \in \mathbb{P}_m$, $b \in \mathbb{P}_n$ and $G \in \mathbb{R}^{m \times n}$ with $a > 0$, $b > 0$, and $G > 0$.

(1) Let $\lambda_i = \frac{1}{a_i} r_i(G)$ $(1 \le i \le m)$ and define $G^a \in \mathbb{R}^{m \times n}$ by

$$G^a_{ij} = G_{ij} / \lambda_i, \quad 1 \le i \le m, \ 1 \le j \le n.$$

Call G^a the row-normalization of G w.r.t. a.

(2) Let $\mu_j = \frac{1}{b_j} c_j(G)$ $(1 \le j \le n)$ and define $G^b \in \mathbb{R}^{m \times n}$ by $G^b_{ij} = G_{ij} / \mu_j$, $1 \le i \le m$, $1 \le j \le n$.

Call G^b the column-normalization of G w.r.t. b.

Note: $\lambda_i = (\text{row-}i \text{ sum of } G)/a_i$.

$$\begin{bmatrix} & G^a & \\ \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \end{bmatrix} = \begin{bmatrix} & G & \\ \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \\ \cdots & \cdots & \cdots \end{bmatrix} \begin{matrix} 1/\lambda_1 \\ \vdots \\ 1/\lambda_i \\ 1/\lambda_m \end{matrix}$$

raw-i sum of $G^a = \sum_j G^a_{ij} = \sum_j G_{ij}/\lambda_i$

$= a_i \sum_i G_{ij} / (\text{raw-}i \text{ sum of } G) = a_i, \ \forall i.$

So, Sinkhorn's alg. is just the alternate raw-column normalization process. In fact, this is the original Sinkhorn process/iteration.

⊙ Sinkhorn (1964) mentioned that Lloyd Welch used the normalization algorithm in an unpublished report of Inst. of Defense Analysis, no time was given. Sinkhorn-Knopp (1967) extended the result. So, the method is sometimes called the Sinkhorn-Knopp algorithm.

⊙ Sinkhorn's alg. (with $v^{(0)}$ given) produces $A^{(k)} = \text{diag}(u^{(k)}) K \, \text{diag}(v^{(k-1)})$ and $B^{(k)} = \text{diag}(u^{(k)}) K \, \text{diag}(v^{(k)})$ $(k = 1, 2, \cdots)$ that converge to $P \in \mathcal{A}(a, b)$; see below. But for any finite k, $A^{(k)}$ and $B^{(k)}$ may not be in $\mathcal{A}(a, b)$. One idea is to "project" some $A^{(k)}$ or $B^{(k)}$ to $\mathcal{A}(a, b)$. Altschuler et al. (2017)

Suggested the following alg. to compute some $P \in \mathcal{A}(a,b)$ that is close to a positive matrix F (like $A^{(k)}$ or $B^{(k)}$): $\alpha \wedge \beta = \min(\alpha, \beta)$

- $x_i = \dfrac{a_i}{r_i(F)} \wedge 1 \quad (i=1,\cdots,m)$
- $F' = \mathrm{diag}(x) F$
- $y_j = \dfrac{b_j}{c_j(F')} \wedge 1 \quad (j=1,\cdots,n)$
- $F'' = F' \mathrm{diag}(y)$
- $P = F'' + [a - r(F')][b - c(F'')]^T / \|a - r(F'')\|_1$.

Here $r_i(F)$ = row-i sum of F,
$c_j(F')$ = col-j sum of F',
$r(F')$ = the vector of $r_i(F')$
$c(F'')$ = the vector of $c_j(F'')$

This is not the closed distance projection. But the convergence (together with Sinkhorn) is proved.

Theorem Let $a \in \mathbb{P}_m, b \in \mathbb{P}_n$, $a>0$, $b>0$. Let $K \in \mathbb{R}^{m\times n}$ with $K \geq 0$. Let $u, \bar{u} \in \mathbb{R}^m$, $v, \bar{v} \in \mathbb{R}^n$, $u>0$, $\bar{u}>0$, $v>0$, $\bar{v}>0$. Suppose $P = \mathrm{diag}(u) K \mathrm{diag}(v) \in \mathcal{A}(a,b)$ and $\bar{P} = \mathrm{diag}(\bar{u}) K \mathrm{diag}(\bar{v}) \in \mathcal{A}(a,b)$. Then $\bar{P} = P$. Moreover, $\exists\, p>0$ such that $u = p\bar{u}$ and $v = \frac{1}{p}\bar{v}$.

[This means that the pair $u>0, v>0$ such that $(u_i K_{ij} v_j) \in \mathcal{A}(a,b)$ is unique up to a multiplicative factor.]

Proof $P = \mathrm{diag}(u) K \mathrm{diag}(v) \Leftrightarrow P_{ij} = u_i K_{ij} v_j \ (>0) \ \forall i,j$
Similarly, $\bar{P}_{ij} = \bar{u}_i K_{ij} \bar{v}_j \ \forall i,j$. $K_{ij} = (1/\bar{u}_i)\bar{P}_{ij}(1/\bar{v}_j)$. So,

$P_{ij} = (u_i/\bar{u}_i)\,\bar{P}_{ij}\,(v_j/\bar{v}_j)$ Let $\alpha_i = u_i/\bar{u}_i$ and $\beta_j = v_j/\bar{v}_j$.

Then $P_{ij} = \alpha_i \bar{P}_{ij} \beta_j \;\; \forall i,j$. It suffices to show that $\exists\, p>0$ s.t. $\alpha_i = p$ and $\beta_j = \frac{1}{p}\;\; \forall i,j$.

Let $\alpha_{i_m} = \min_i \alpha_i$ and $\beta_{j_M} = \max_j \beta_j$. Since $\bar{P} \in \mathcal{A}(a,b)$ and $P \in \mathcal{A}(a,b)$, we have

$$\alpha_{i_m}\beta_{j_M} = \frac{\alpha_{i_m}\beta_{j_M}}{b_{j_M}} \sum_i \bar{P}_{i j_M} \overset{(A)}{\le} \frac{1}{b_{j_M}} \sum_i \alpha_i \bar{P}_{i j_M} \beta_{j_M}$$

$$= \frac{1}{b_{j_M}} \sum_i P_{i j_M} = 1.$$

$$\alpha_{i_m}\beta_{j_M} = \frac{\alpha_{i_m}\beta_{j_M}}{a_{i_m}} \sum_j \bar{P}_{i_m j} \ge \frac{1}{a_{i_m}} \sum_j \alpha_{i_m} \bar{P}_{i_m j} \beta_j$$

$$= \frac{1}{a_{i_m}} \sum_j P_{i_m j} = 1.$$

Hence, $\alpha_{i_m}\beta_{j_M} = 1$, and all above are equal. (A) is an equality. So,

$$\alpha_{i_m} \sum_i \bar{P}_{i j_M} = \sum_i \alpha_i \bar{P}_{i j_M}, \quad \sum_i \underbrace{(\alpha_i - \alpha_{i_m})}_{\ge 0} \underbrace{\bar{P}_{i j_M}}_{>0} = 0 \Rightarrow \alpha_i = \alpha_{i_m}\; \forall i$$

Similarly, $\beta_j = \beta_{i_M} = \frac{1}{\alpha_{i_m}} \;\; \forall j$.

Set $p = \alpha_i \;\forall i$. Then $u = p\bar{u}$ and $v = \frac{1}{p}\bar{v}$. clearly, then $P = \bar{P}$

QED