Real-time Multiattribute Bayesian Preference Elicitation with Pairwise Comparison Queries

Shengbo Guo & Scott Sanner

Dec. 12, 2009

NICTA is proudly supported by:

Australian Government
Department of Communications, Information Technology and the Arts
Australian Research Council

NICTA Members

NICTA Partners

Problem: PC Recommendation Example

ID	Manufacturer	CPU	Price
1	Apple	PowerPC G3 300	799
2	Dell	Intel Celeron 900	1119
3	Fujitsu	Intel Pentium 600	989
…	…	…	…
N	Toshiba	Intel Pentium 1000	1150

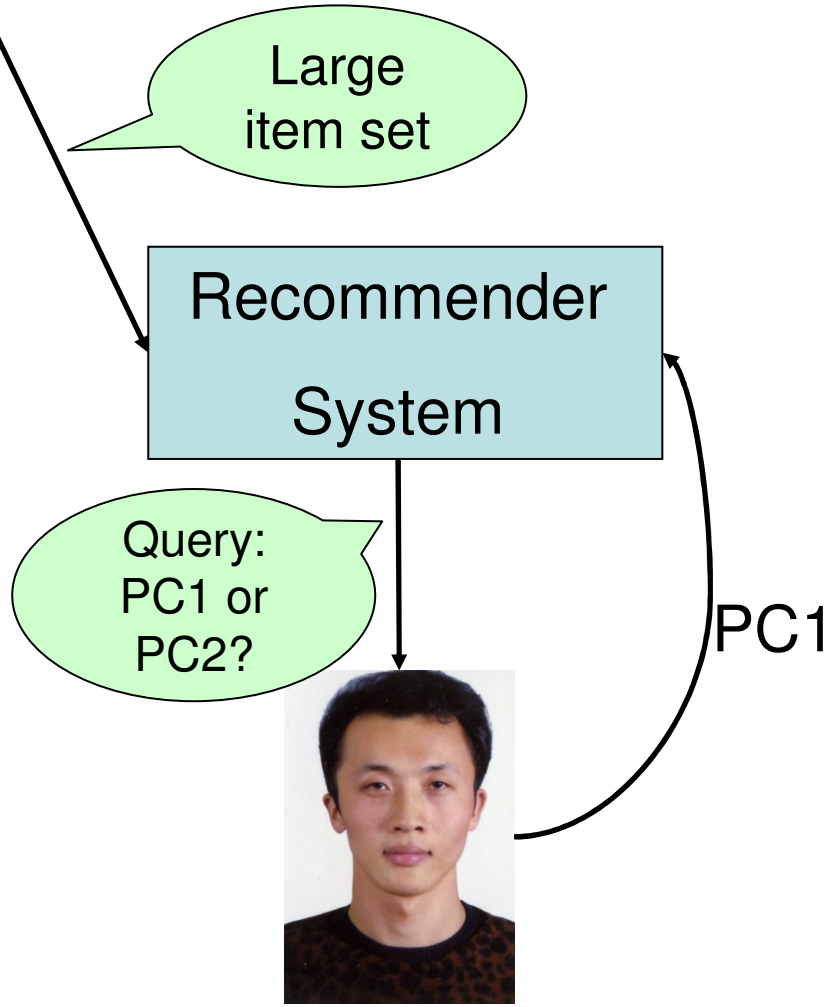

After 10~20 queries, make recommendation!

Utility Representation & Bayesian Update

- Additive utility representation factors across D attributes

$$u^*(\mathbf{x}) = \sum_{d=1}^{D} \mathbf{w}^*_{d,\#(\mathbf{x},d)}, \quad u(\mathbf{x}|\mathbf{w}) = \sum_{d=1}^{D} \mathbf{w}_{d,\#(\mathbf{x},d)}$$

- Bayesian update

$$P(\mathbf{w}|R^{n+1}) \propto P(q_{ij}|\mathbf{w}, R^n)P(\mathbf{w}|R^n)$$
$$\propto P(q_{ij}|\mathbf{w})P(\mathbf{w}|R^n)$$

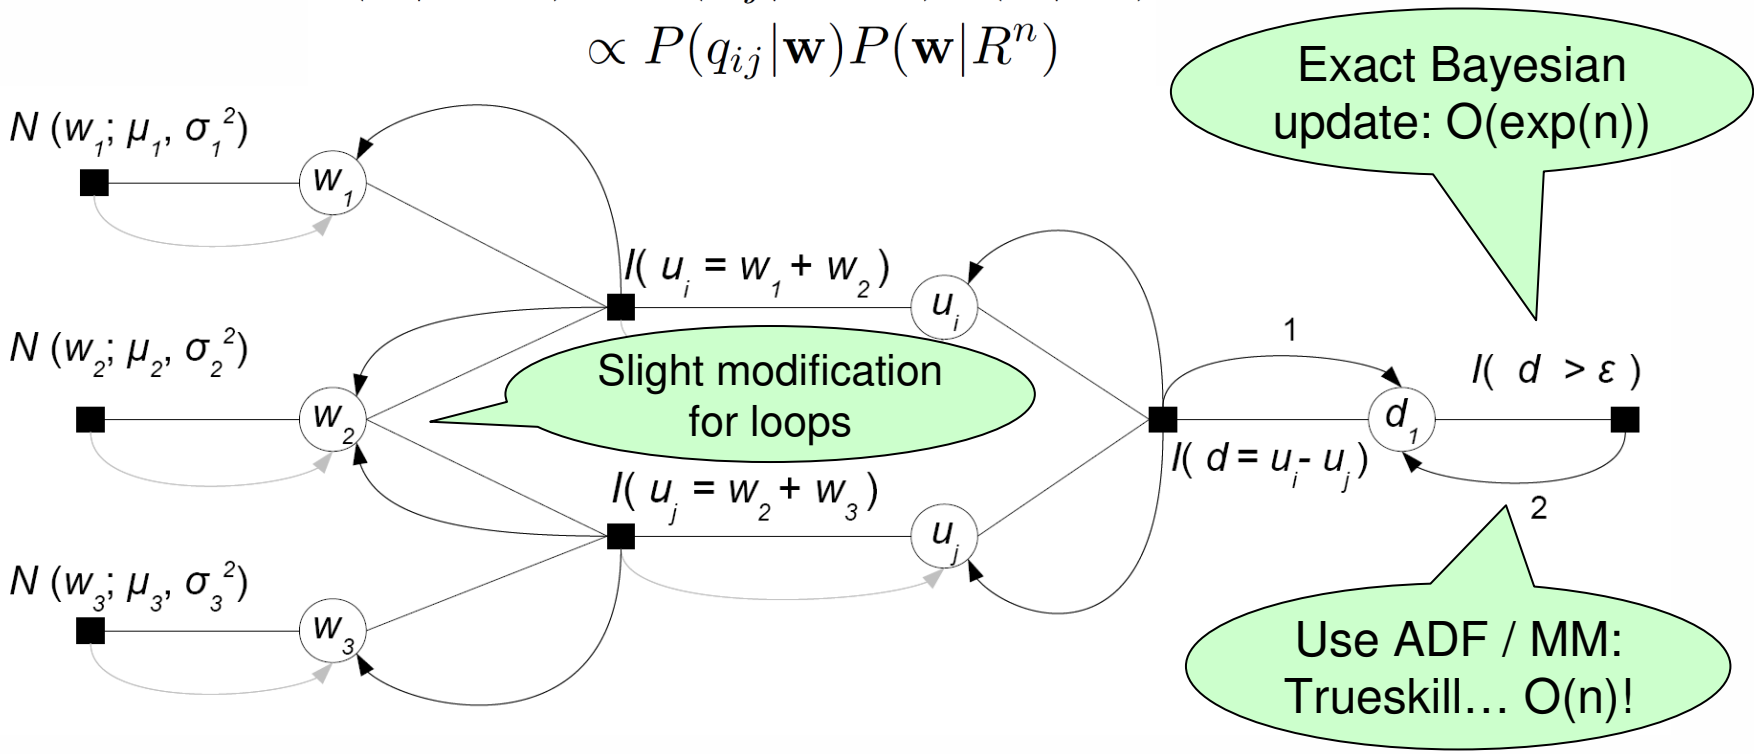

Query Strategies

- Expected Loss

$$\widehat{\mathrm{EL}}(k,R) = \mathrm{E}_{P(\mathbf{w}|R)}\left[\max(0, u(k|\mathbf{w}) - u(i_R^*|\mu^R))\right]$$
$$= (\mu_{i_R^*} - \mu_k)(1 - \Phi_{\mu_k,\sigma_k^2}(\mu_{i_R^*})) - \frac{\sigma_k}{\sqrt{2\pi}}\exp\left(-\frac{(\mu_{i_R^*} - \mu_k)^2}{2\sigma_k^2}\right)$$

- Maximal Expected Loss

$$\mathrm{MEL}(R) = \max_k \widehat{\mathrm{EL}}(k,R)$$

- Expected Value of Information (EVOI)

$$\mathrm{EVOI}(R,i,j) = -\mathrm{MEL}(R) + \mathrm{E}_{P(\mathbf{w}|R)}\sum_{q_{ij}}\left[P(q_{ij}|\mathbf{w})\mathrm{MEL}(R\cup\{q_{ij}\})\right]$$
$$= -\mathrm{MEL}(R) + \sum_{q_{ij}}\left[\mathrm{E}_{P(\mathbf{w}|R)}P(q_{ij}|\mathbf{w})\right]\mathrm{MEL}(R\cup\{q_{ij}\})$$

- Restricted EVOI: restrict item *i* to be the current best one… $O(n^2)$ → $O(n)$!

Experimental Results