Active Learning for Matching Problems

Laurent Charlin, Richard Zemel and Craig Boutilier

UNIVERSITY OF TORONTO

Problem: Recommend items to users under matching constraints

- Learning user preferences eases user elicitation burden
- Active learning provides a further improvement

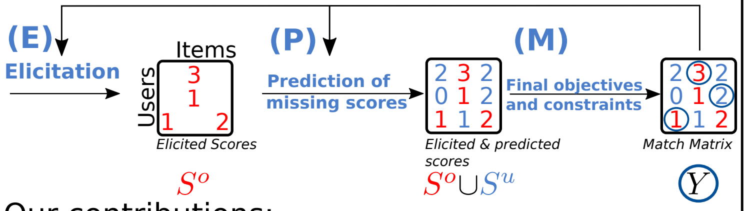

Our contributions:

- Matching-aware active learning methods
- Probabilistic matching procedure

(P) Score Prediction

Any score prediction model can be used

-Assume model keeps a distr. over $S^u : \Pr(S^u | S^0, X, \theta)$

(M) Matching Problem

$$\begin{bmatrix} 2 & 3 & 2 \\ 0 & 1 & 2 \\ 1 & 1 & 2 \end{bmatrix}$$

$$\text{max.} \quad J(Y) = \sum_r \sum_p s_{rp} y_{rp}$$

$$\begin{bmatrix} 2 & 3 & 2 \\ 0 & 1 & 2 \\ 1 & 1 & 2 \end{bmatrix}$$

$$\text{s.t.} \quad y_{rp} \in \{0,1\}, \quad \forall r, p$$

$$\sum_r y_{rp} = R_{targ}, \quad \forall p$$

$$\sum_p y_{rp} \geq P_{\min}, \ \sum_p y_{rp} \leq P_{\max}, \ \forall r.$$

Taking predicted score uncertainty into account:

$$\Pr(Y | S^o, X, \theta) = \int Y(S^u \cup S^o) \Pr(S^u | S^o, X, \theta) \, dS^u$$

- Intractable but we use sampling:

$$\begin{bmatrix} \textbf{S} & & & & & \\ 3 & 2 & 2 & 3 & 2 & 1 \\ 3 & 1 & 2 & 1 & 3 & 2 \\ 4 & 1 & 4 & 3 & 2 & 4 \end{bmatrix}$$

$$\begin{bmatrix} \text{Smaller Variance} & & & & & \\ 0.13 & 0.08 & 0.01 & 0.01 & 0.09 & 0.00 \\ 0.01 & 0.05 & 0.02 & 0.02 & 0.02 & 0.12 \\ 0.10 & 0.09 & 0.05 & 0.05 & 0.01 & 0.10 \end{bmatrix} \rightarrow \begin{bmatrix} \textbf{Y} & & & & & \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} \text{Larger Variance} & & & & & \\ 1.3 & 0.8 & 0.5 & 0.1 & 0.9 & 0.0 \\ 0.1 & 0.5 & 0.3 & 0.2 & 0.2 & 1.2 \\ 1.0 & 0.9 & 0.4 & 0.5 & 0.1 & 1.0 \end{bmatrix} \rightarrow \begin{bmatrix} \overline{\textbf{Y}} & & & & & \\ 0.20 & 0.78 & 0.04 & 0.94 & 0.02 & 0.02 \\ 0.48 & 0.20 & 0.10 & 0.00 & 0.98 & 0.24 \\ 0.32 & 0.02 & 0.86 & 0.06 & 0.0 & 0.74 \end{bmatrix}$$

$$\begin{bmatrix} \textbf{LP Solution} & & & & & \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}$$

Constraints:

- 2 items per user
- 1 user per item

(E) Elicitation: Active Learning Methods

Query the scores that will be most helpful for matching:

- Active learning methods should take the matching objective (J) into account
- Computing EVOI is expensive
- We propose methods that operate in matching space

Matching Max-score query

$$YM: \quad \arg\max_{(rp) \in S^u} y_{rp} \hat{s}_{rp}$$

$$\overline{Y}M: \quad \arg\max_{(rp) \in S^u} \bar{y}_{rp} \hat{s}_{rp}$$

$$\begin{bmatrix} 3 & 3 & 2 \\ 0 & 2 & 3 \\ 3 & 1 & 2 \end{bmatrix}$$

Matching Max-entropy query

$$\overline{Y}E: \quad \arg \max_{(rp) \in S^u} \left[- \sum_{\bar{y}_{rp} \in \{0,1\}} \bar{y}_{rp} \log \Pr(\bar{y}_{rp}) \right]$$

Experimental Setup

- Simulate elicitation procedure
  - Every user starts with a few obs. scores
  - At each round each user is queried (batch)
- Bayesian PMF is used to predict scores

- Data
  - Jokes (Jester):
    - 300 users, 10 jokes
    - Match 10 to 30 users per joke
  - Dating (LibimSeTi.cz):
    - 250 users, 250 items (users)
    - Match 15 to 25 users per item
  - Conference (NIPS'10):
    - 1250 papers, 48 revs
    - Match 20 to 30 papers per rev.

Suitabilities

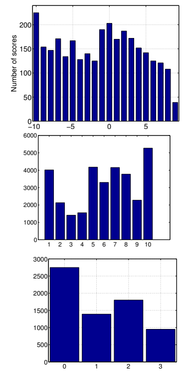

Experimental Results

Matching Performance

- Run the LP (M): $Y(S^o \cup \hat{S}^u)$
- Value of a match: $\sum_{rp} y_{rp} s_{rp}$

Difference in matching obj.

Num. rand. queries

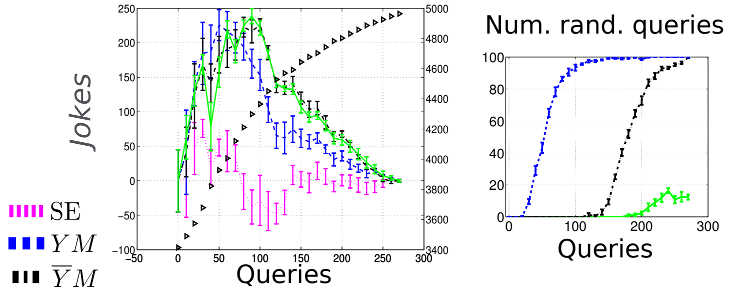

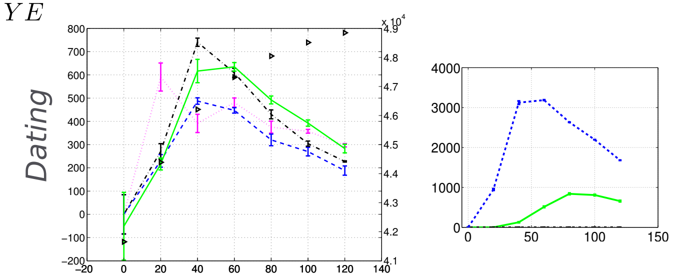

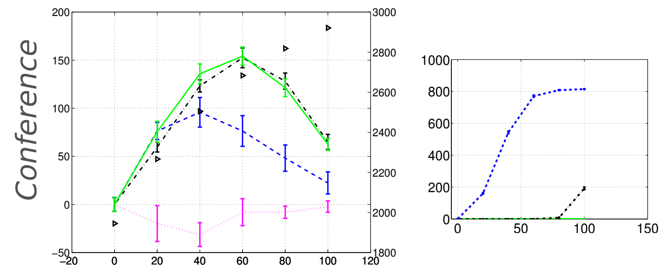

- Matching-aware active learning is a win
- Further results: our approach is robust to matching constraints, batch sizes.

Conclusion: Effective for determining high-quality matches with significantly less elicitation

Current work: EVOI extensions, different types of queries (side-information, higher level)