Eliciting Preferences By Comparing Candidates

Haoming Li, Lirong Xia & Zhibing Zhao
Rensselaer Polytechnic Institute

Introduction

The Problem:

- **Given:** Agents and candidates;
- **Goal:** Learn agents' preferences and make recommendation;
- **Method:** Ask the agents to compare candidates. Answers can be non-deterministic.

Mon	: "Burger King > McDonald's"
Tue	: "Wendy's > Burger King"
Wed	: "Wendy's > McDonald's"
Thu	: "McDonald's > Wendy's"
Fri	: "McDonald's > Burger King"
…	…

Mathematically:

- $z_{ni} = z(x_i, s_n)$ is a set of features relating to candidate i for agent n that depends on attributes of the candidate, x_i, and characteristics of the agent, s_n;
- $U_{ni} = \beta z_{ni} + \varepsilon_{ni}$ is the utility agent n obtains from candidate i, where
- β is a corresponding vector of coefficients of the features.

Assuming z is known, our goal is to estimate β by asking the agents to compare candidates.

Methods

Algorithm:

Randomly ask for an initial set of data D^1
For $t = 1 \dots T$ do:

1. Compute $\hat{\beta}^t = MAP(D^t)$ using MCMC (Metropolis algorithm);
2. Ask for pairwise comparison h^t that maximizes the expected information gain $\mathrm{E}[G(Pr(\cdot\,|D^t + \{\pi^t\}))]$, where π^t is the answer to h^t;
3. Update $D^{t+1} = D^t + \{\pi^t\}$.

"Quality" and Information Criteria:

- $G(Pr(\cdot\,|D^t + \pi^t))$ evaluates the "quality" of the posterior of parameter, the information gain from answer π^t.
- Standard criteria from Bayesian experimental design literature, such as D-Optimality (Shannon information gain) and E-Optimality, favor "steep" posterior distribution.

D-Opt	E-Opt	Proposed
Minimizes determinant of covariance matrix of the Gaussian	Minimizes maximum eigenvalue of covariance matrix of the Gaussian	Maximizes minimum $\frac{\lvert mean(U_{na} - U_{nb})\rvert}{std(U_{na} - U_{nb})}$ for any pair of candidates a and b

Inspired by Student's t-test, we favor posterior distribution from which the min certainty of any comparison is high.

Results

Below, we compare the root mean squared error between the true β and estimated β.

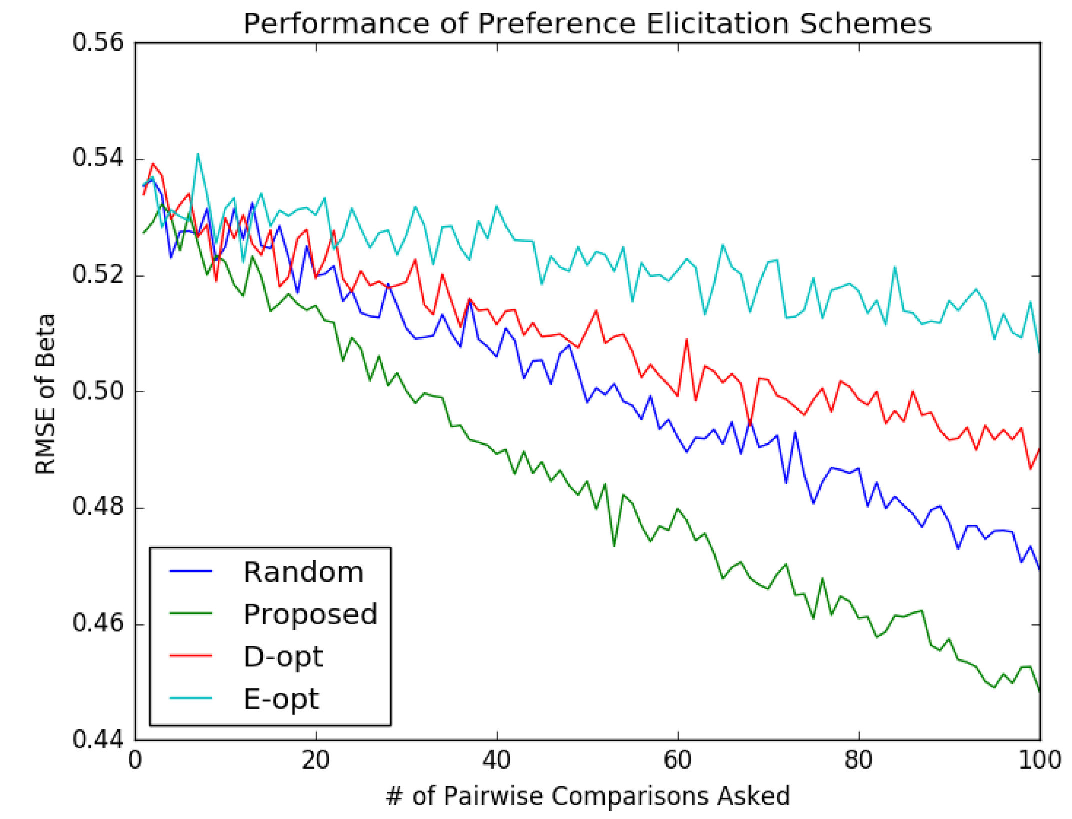

Below, we compare the number of overlaps between the top-3 candidates computed from the true β and that from estimated β.

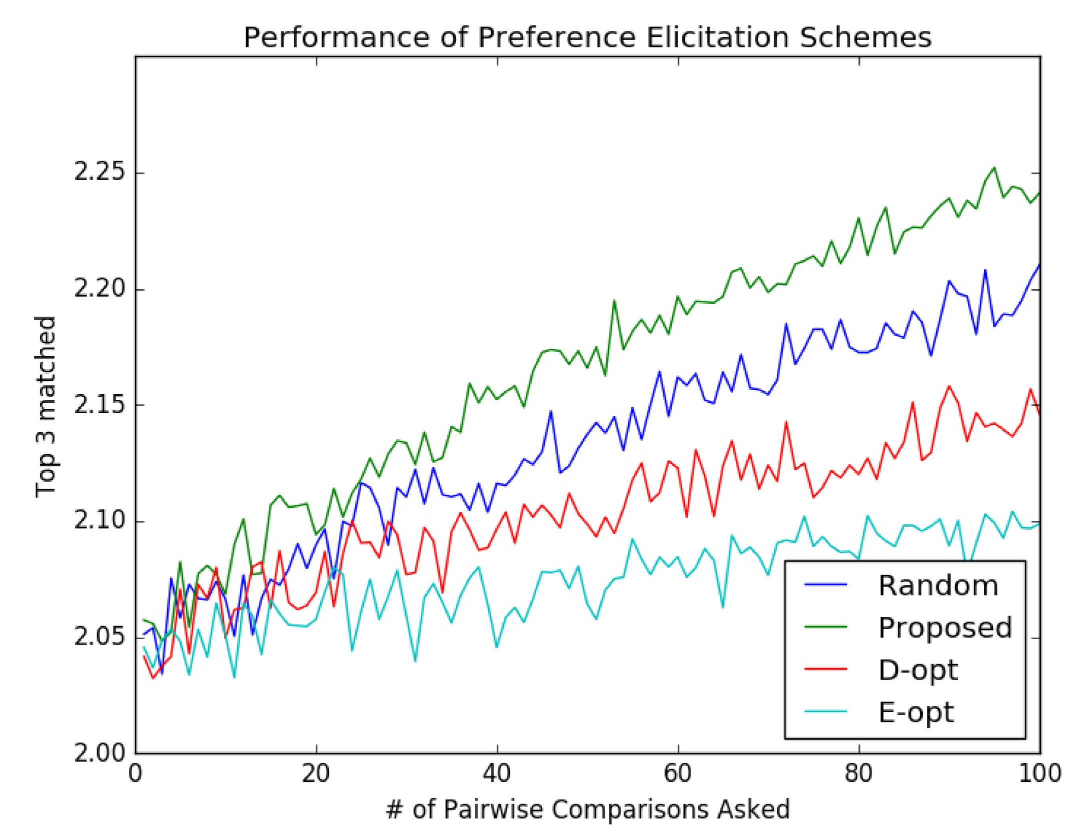

The synthetic data is generated with 10 candidates, 1 agent, 3 features z. We ran a total of 3316 experiments.

Conclusion

Our experiment of the elicitation criteria proposed by Azari Soufiani et al. shows that it improves the estimation relative to other elicitation criteria, when eliciting preferences by comparing candidates.

We see the methodology being adopted in other preference elicitation applications in the future, such as recommendation systems, product prediction, etc.

Future Works:

1. Implement the algorithm in the Cognitive and Immersive Situations Room at CISL@RPI.
2. Can we better elicit preferences by asking a larger set of questions, such as to choose from a set of 5 candidates?

Acknowledgements
This work is supported by Cognitive and Immersive Systems Lab - collaboration between IBM and RPI, a research center in Cognitive Horizon Network.

cognitive + immersive systems lab

References

Azari Soufiani, Hossein, David C. Parkes, and Lirong Xia. Preference Elicitation For General Random Utility Models. In *Uncertainty in Artificial Intelligence: Proceedings of the 29th Conference*, page 596-605, Corvallis, OR: AUAI Press.