A Cost-Effective Framework for Preference Elicitation and Aggregation

Zhibing Zhao[1], Haoming Li[1], Junming Wang[1], Jeffrey Kephart[2], Nicholas Mattei[2], Hui Su[2], Lirong Xia[1]

1. Rensselaer Polytechnic Institute (RPI)
2. International Business Machine (IBM)

Accepted to UAI-2018

Introduction

- Key group makes decisions
- Budget/time limit
- Asking key group directly for preferences is very costly
- How do we elicit key group's preferences cost-effectively?

Our Contributions

- Cost-effectives Framework
  - Cost
  - Flexible, compatible with any:
    - Ranking model
    - Types of questions asked: pairwise, full ranking, top-k, etc
    - Information criteria
  - Algorithm computing winner distributions using randomized voting rules

Preliminaries: Plackett-Luce Model, One-Step Bayesian Experimental Design

- Plackett-Luce Model (with features):

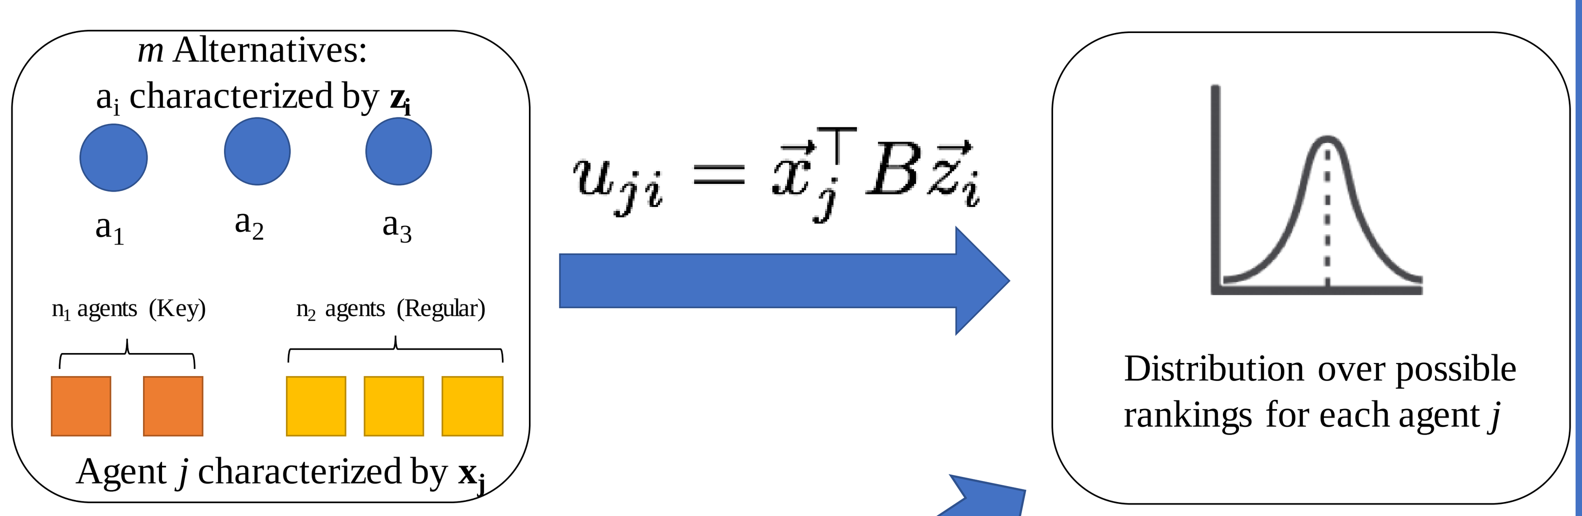

$$\Pr(R_j = a_1 \succ \ldots \succ a_m) = \prod_{i=1}^{m-1} \frac{\exp(u_{ji})}{\sum_{p=i}^{m} \exp(u_{jp})}$$

- One-Step Bayesian Experimental Design:
  - Design h: an agent (regular) and a question
  - Cost: $w(h)$
  - $\pi(B)$: A distribution over B
  - G: A measure of information quality of $\pi(B)$
    - D-optimality, E-optimality, MPC
  - $h^* = \max \dfrac{E[G(\pi_{\text{posterior}}(B))] - G(\pi_{\text{prior}}(B))}{w(h)}$

Algorithm

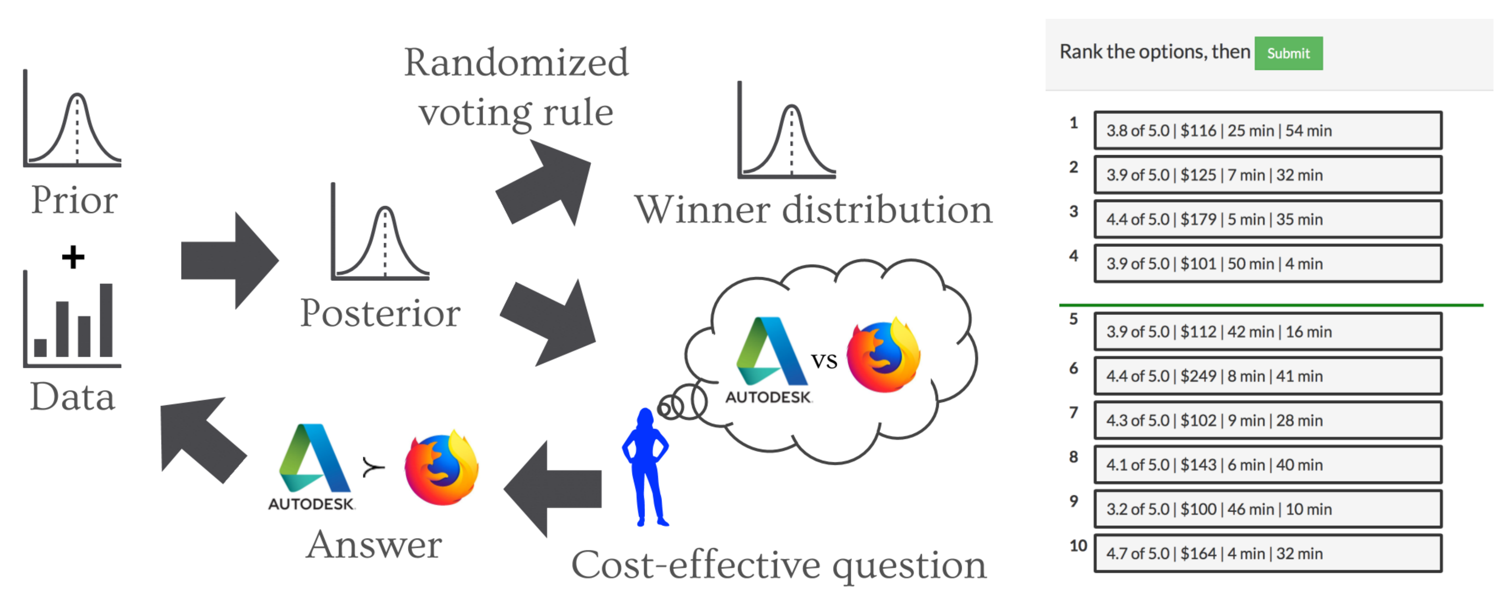

- **Input**: Budget W, randomized voting rule r, cost function w(h), information criterion $G(\pi(B))$, the set of designs H where for any $h \in H$, $w(h) \le W$.
- **Output**: A predicted preference if $n_1 = 1$, or a distribution of winning alternatives for group decision when $n_1 \ge 2$.
- **Initialization**: Randomly initialize dataset D^t.
- **Repeat**:
  - Compute/approximate $\pi(B^t | D^t)$
  - Compute $h^{*t} \in H$
  - Implement h^{*t}. Let R^t denote the answer. Then $D^{t+1} \leftarrow D^t \cup \{R^t\}$ $H \leftarrow H - h^{*t}$, $W \leftarrow W - w^t$
  - Remove all h's from H where $w(h') > W$
- **Until** $H = \emptyset$
- Compute the predicted preference when $n_1 = 1$ or a distribution of winning alternatives according to voting rule r

Approximation of Posterior Distribution

- Posterior Distribution is hard to compute
- Approximated by $N(B_{\text{CML}}, \Sigma)$:[Pauli et al., 2011]
  - Composite marginal log-likelihood: $\text{CLL}(B) = \sum_{\lambda=1}^{q} \ln \Pr(\mathcal{E}_\lambda | B)$
  - Set of events: $\{\mathcal{E}_1, \ldots, \mathcal{E}_q\}$
  - $B_{\text{CML}} = \arg\max_{B \in \Theta} \text{CLL}(B)$
    $\Sigma^{-1} = -\nabla^2 \text{CLL}(B)$

Randomized Voting Rule

	a_1	a_2	a_3
plurality	*2*	*1*	*0*
Borda	*5*	*3*	*1*

An example of plurality and Borda scores. Under plurality, a_1 wins with a probability of 2/3. a_2 wins with a probability of 1/3.

- Computing winner(s) with non-deterministic preferences using deterministic voting rule is hard
- Theorem 1: For any $1 \le i \le m$, $\Pr_r(a_i) \propto \sum_j^{n_1} EX_{ji}$
- Winner distributions under Probabilistic plurality and Probabilistic Borda are proven to be easy to compute

Experiment Results

Part One: Measuring Cost Function

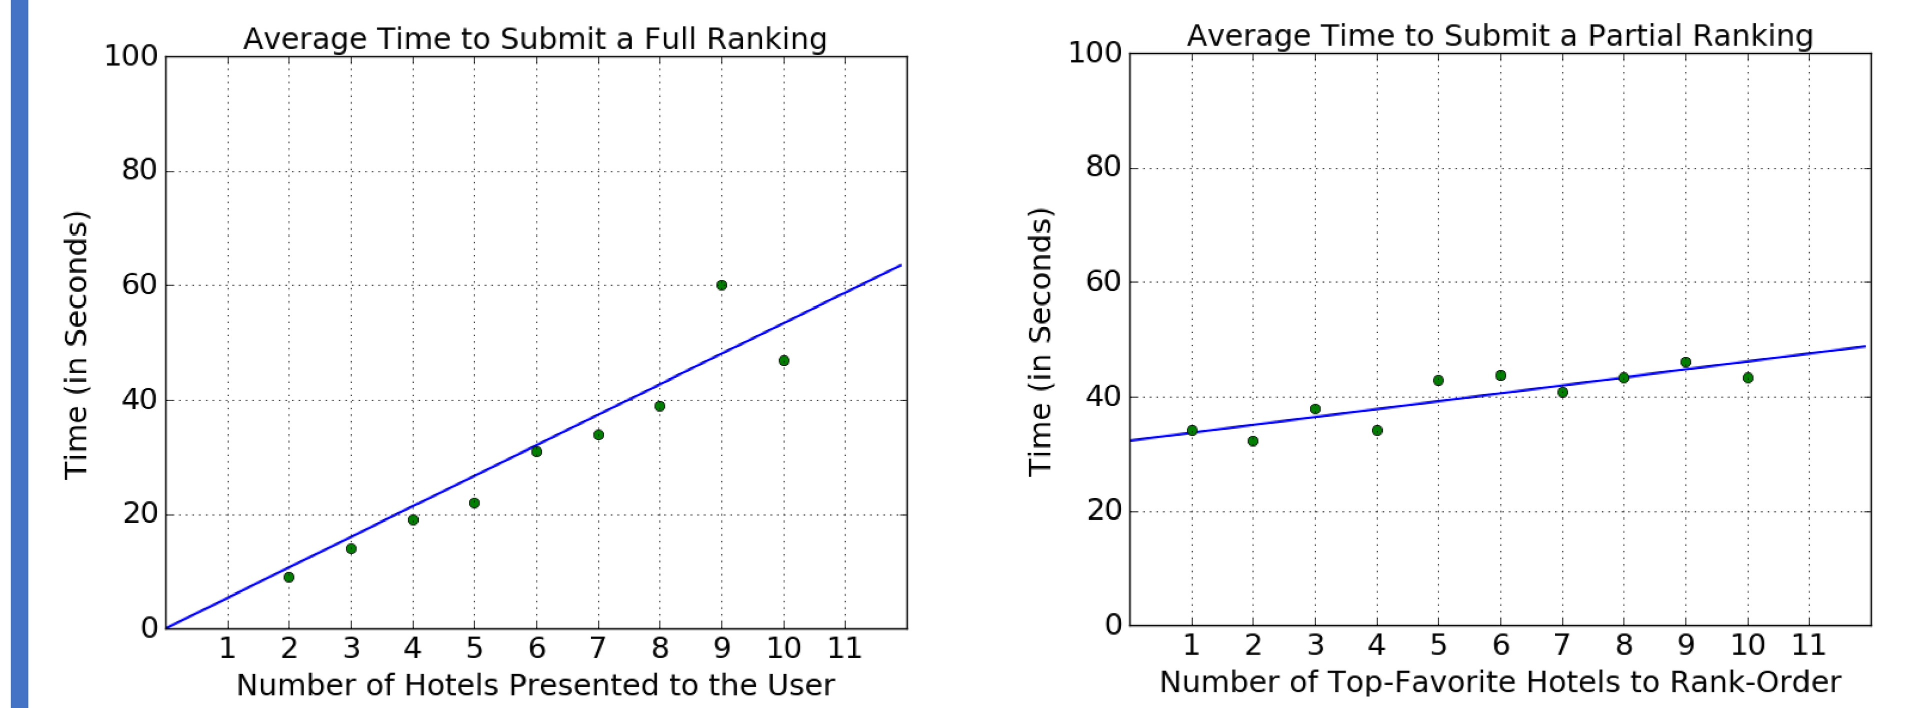

The left subfigure shows the average time a user spent to submit a full ranking over 2,…, 10 alternatives; the right subfigure shows the average time a user spent to give her ranked top 1,…,10 alternatives when 10 alternatives were proposed.

Part Two: Measuring Cost-effectiveness of Framework

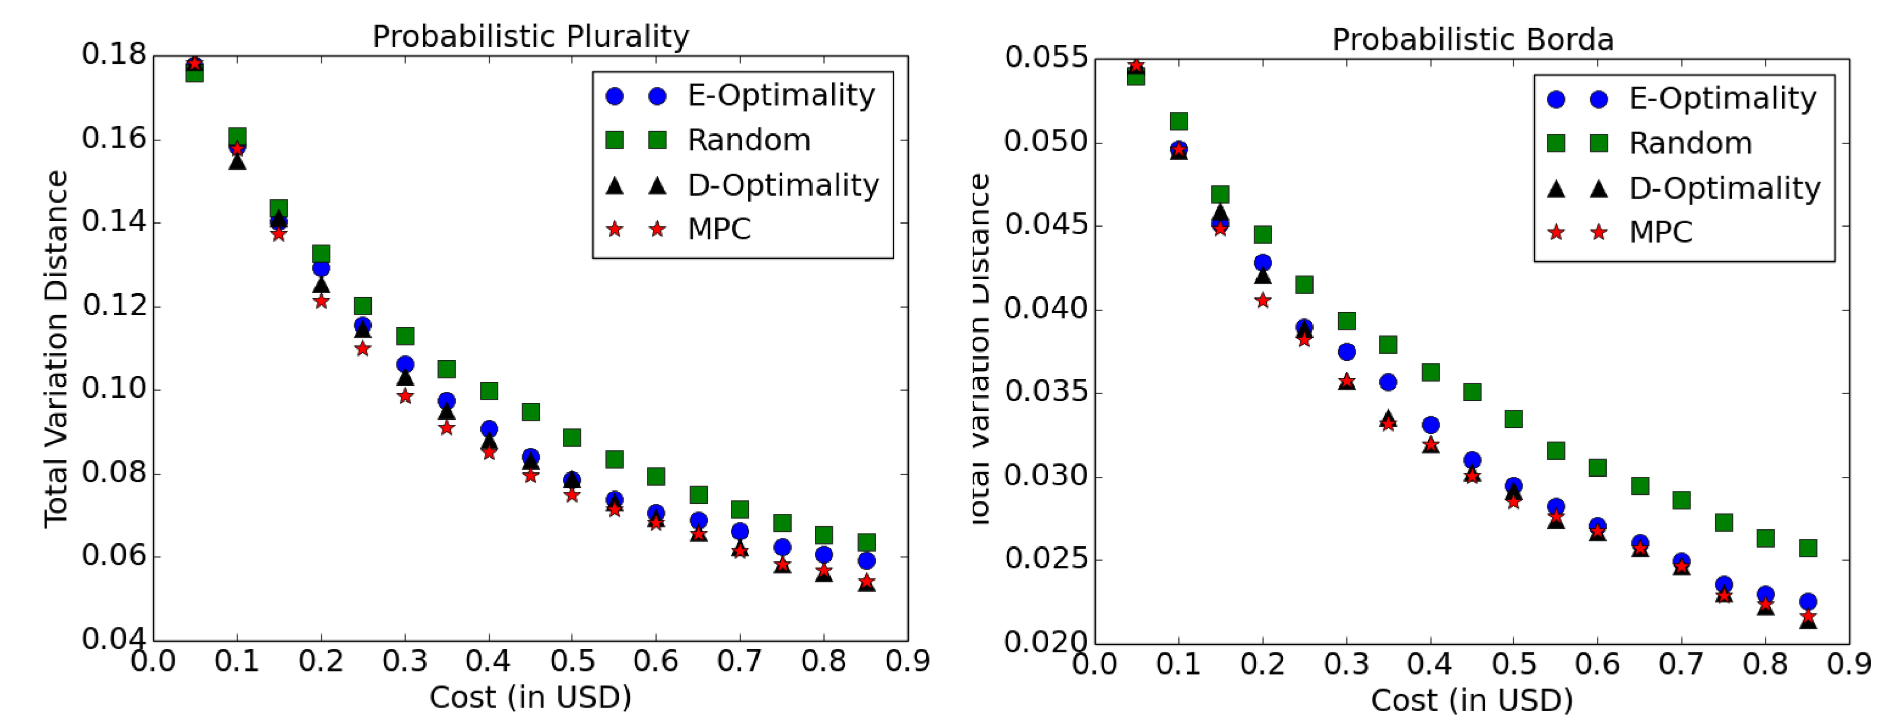

Total variation distance for probabilistic plurality (left) and probabilistic Borda (right).

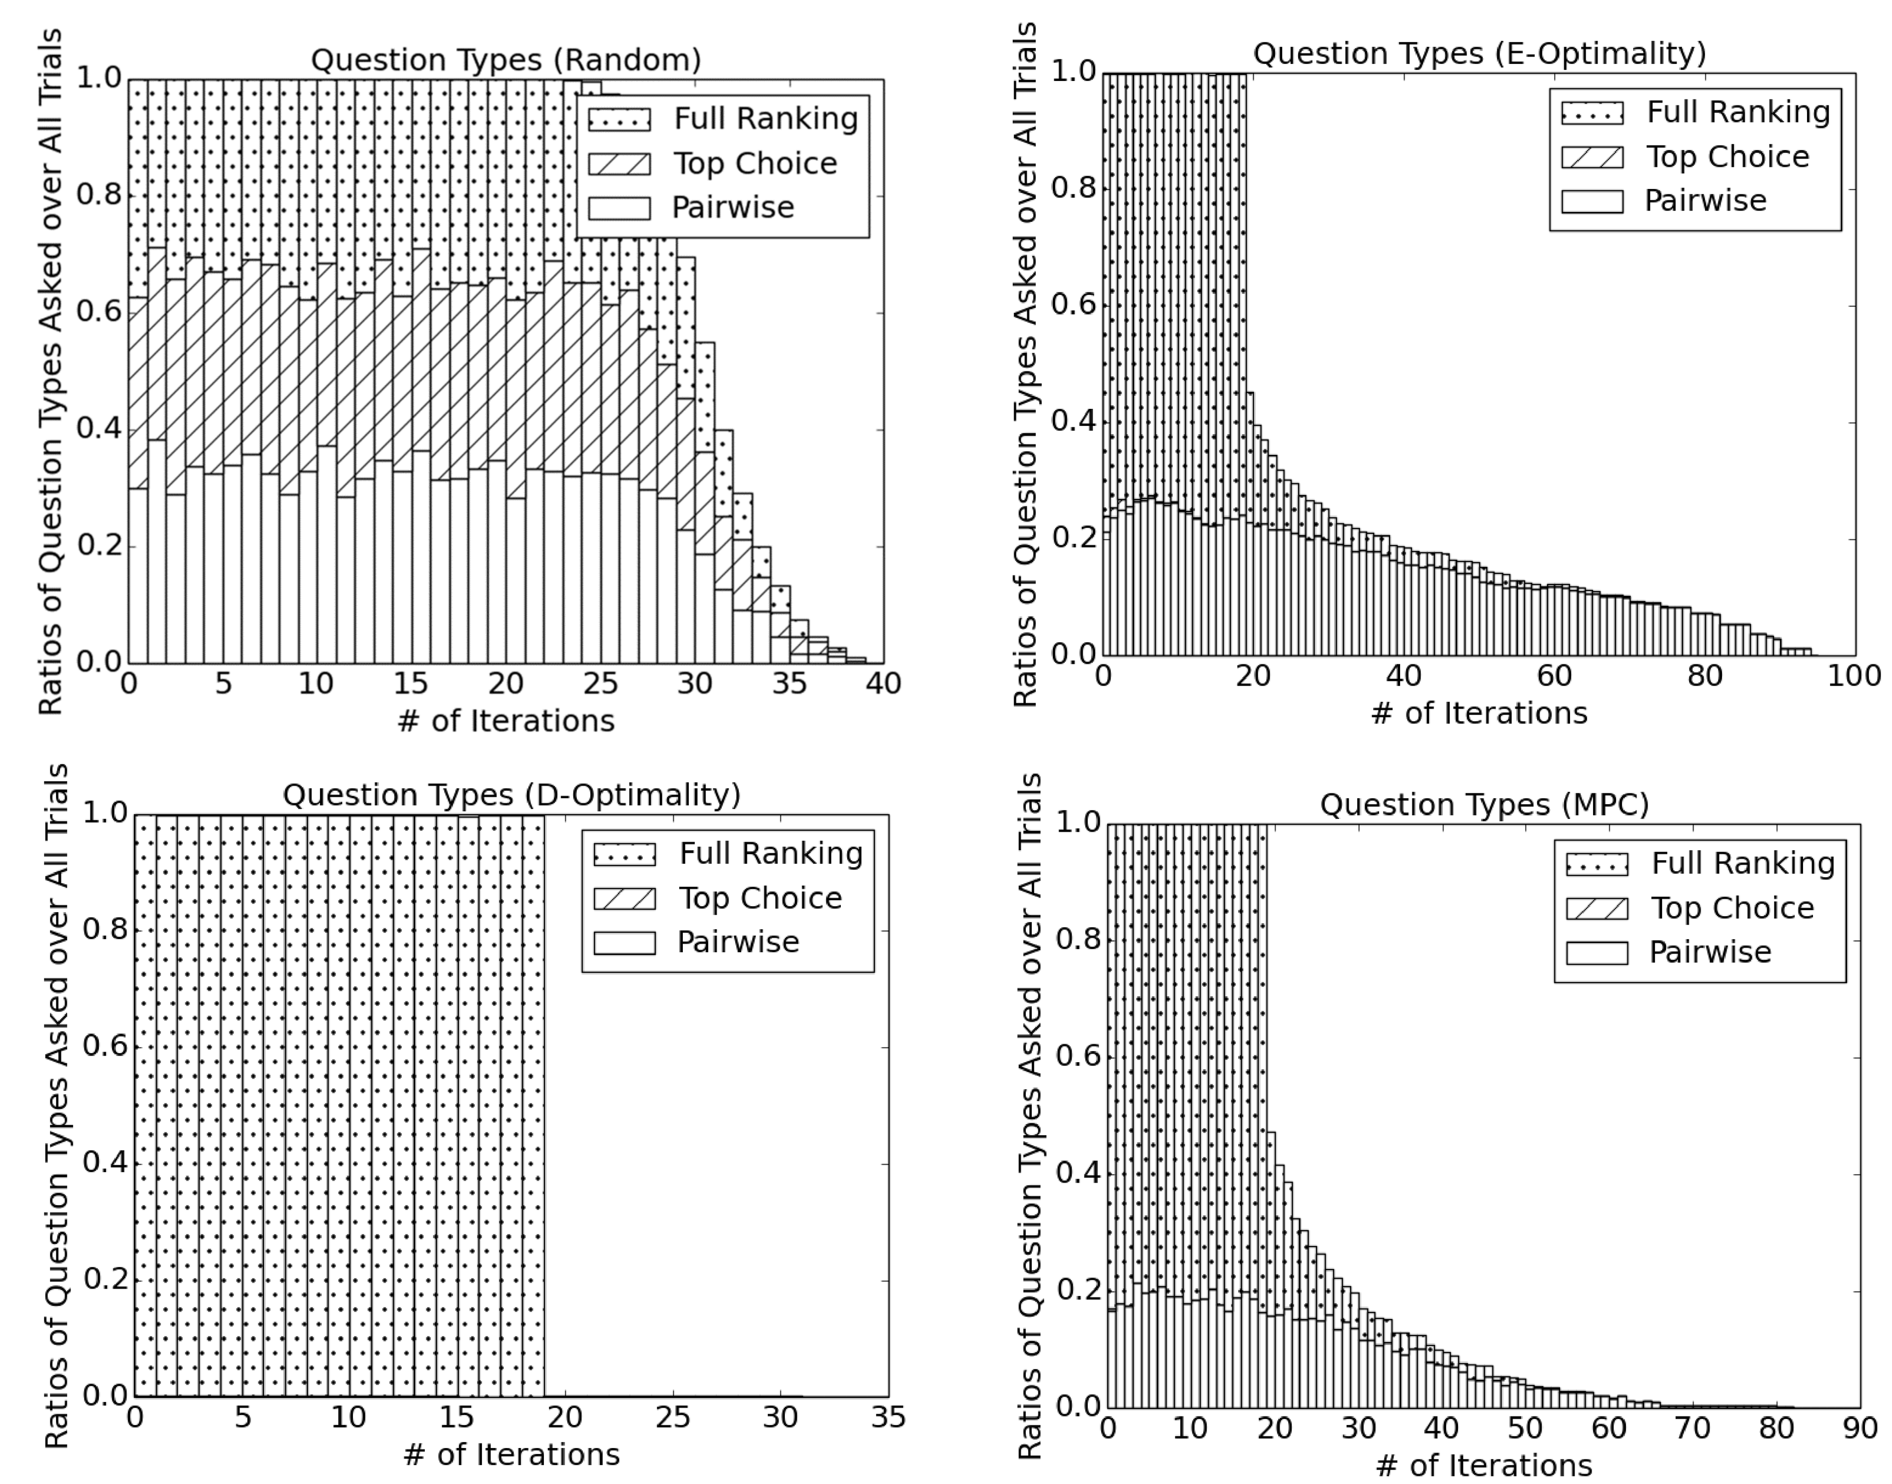

Types of questions chosen by the three information criteria and random.