
Preference Learning for Real-World Multi-Objective Decision Making


Zhiyuan “Jerry” Lin, Stanford

Adam Obeng, Facebook

Eytan Bakshy, Facebook

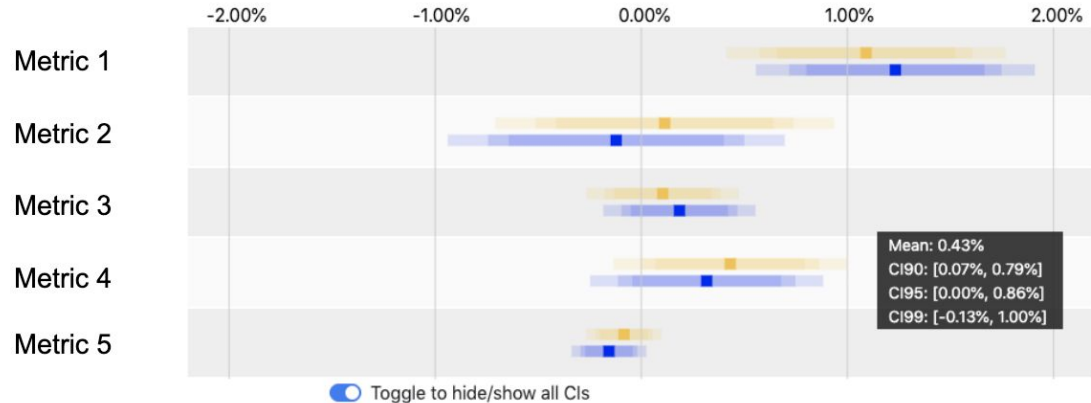
Multi-objective decision making

In Internet experimentation



User study

Preference elicitation interface



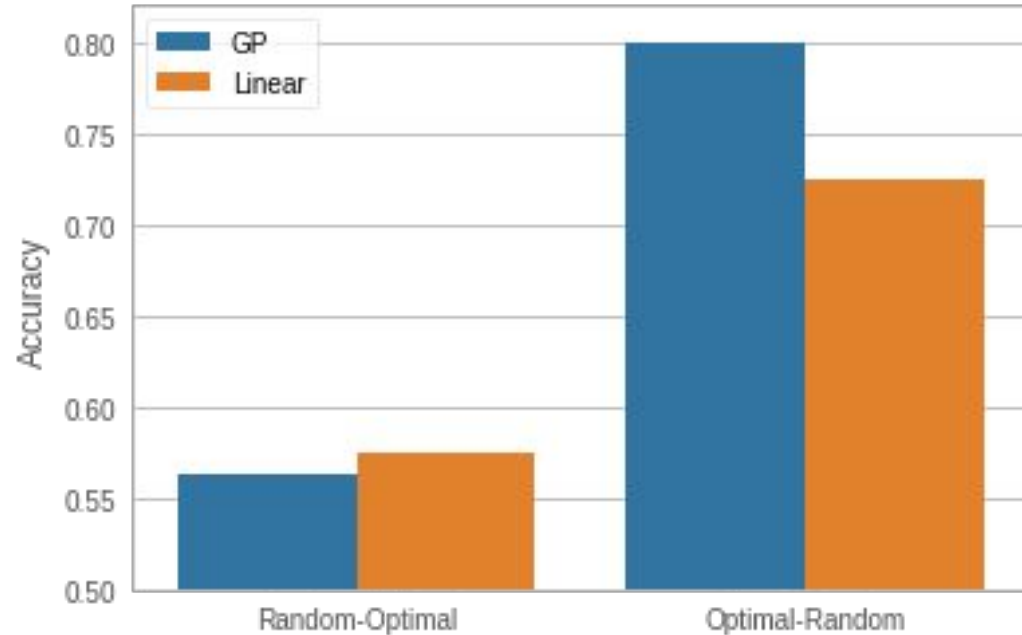
Which is your preferred outcome?

Top Outcome

Bottom Outcome

User study

Random vs. optimal outcomes



Summary

- Preference learning for real decision making utility functions in Internet experiments
- Pure exploration (i.e., random) strategy may not be the most efficient strategy for learning a good utility function in practice
- More complex models (e.g., Gaussian Processes) may better model the tradeoffs than simple linear models