Sequential Optimal Inference for Experiments with Bayesian Particle Filters

Remi Daviet

Wharton Marketing Department, University of Pennsylvania

Introduction

- Behavioral experiments are bounded by time and resources considerations
- Researchers need to optimize the amount of relevant information with each question

Questions:

- What is "relevant information" ?
- How to optimize the question ?
- Can it be done adaptively ?

Importance

- Topic emerged in the 70s
  - see Chaloner and Verdinelli [1995] for a review of the Bayesian approach
- A whole field is dedicated to it (Experimental Design)
- Problem well defined, solution is not
- Increase in model complexity has lead to a need to create methods for adaptive designs:
  - DOSE: Imai and Camerer [2019]
  - DEEP: Toubia et al. [2013]
  - ADO: Cavagnaro et al. [2010]

Current Adaptive Methods

Adaptive method	DOSE	DEEP	ADO	SOI (this paper)
Estimation in continuous space		✓		✓
Model Selection			✓	✓
Exact optimization.	✓		✓	✓
General inference method				✓

Table: Comparison of various adaptive methods available in the literature.

- Our method (SOI) is general and has several advantages:
  - Compatible with complex models
  - Multiple objectives (estimation, prediction, model selection, ...)
  - Fast computation allowing for real time estimation

Optimal design ?

You are a researcher, we can define a utility for the observations in an experiment (e.g. relevance information) : $u(answer|question)$

e.g.: chose between the following lotteries:

- 50% chance of getting 20USD
- 20% chance of getting 10USD

Is this question useful ? How to define useful ?

We can use the Kullback–Leibler divergence between prior beliefs and posterior beliefs

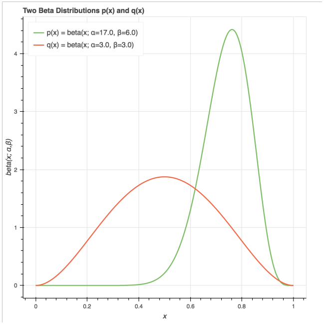

Bayesian Information

We can use the Kullback–Leibler divergence between prior beliefs and posterior beliefs

- Inference: Between the prior and the posterior on the parameters

$$p(\theta) \longrightarrow p(\theta | obs, question)$$

- Prediction: Between the prior and the posterior on the answer y^* to a particular question

$$p(y^*) \longrightarrow p(y^* | obs, question)$$

- Model selection: Between the prior and the posterior on models probabilities

$$p(model) \longrightarrow p(model | obs, question)$$

Expected utility

Since we do not know the answer when designing the question, we use expected utility

$$EU(question) = \sum_{answers} u(answer|question)p(answer|question)$$

Or in continuous answer space:

$$EU(question) = \int_{answers} u(answer|question)p(answer|question)d_{answer}$$

Issue:

- Generally requires a complicated integral over the parameter space Θ

Issue

Problem : Generally requires a complicated integral over the often high dimensional parameter space Θ. Example for parameter estimation :

$$\max_{\eta} EU(\eta) = \max_{\eta} \iint \log\left(\frac{p(y|\theta,\eta)}{p(y|\eta)}\right) p(y|\theta,\eta)p(\theta)d\theta dy.$$

η: question (design), θ: model's parameter, y: answer

How to solve this computational problem in between questions ?

Solution

Introducing Sequential Monte Carlo (SMC):
Provides at any time a set of P draws $\theta^{(p)}$ called particles from the prior/posterior distributions.

Benefits:

- Can be used to approximate the integral in the optimization problem

$$\max_{\eta} \frac{1}{P} \sum_{p=1}^{P} \sum_{y \in \mathcal{Y}} \log \left(\frac{p(y|\theta^{(p)}, \eta)}{p(y|\eta)} \right) p(y|\eta, \theta^{(p)})$$

- Handles multimodality well
- Computations are parallelizable

Implementation

The Sequential Optimal Inference (SOI) method:

- Draw P particles from prior
- Repeat:
  - Find optimal next question using particles
  - Observe answer
  - Update particles to reflect posterior (SMC update)

Implementation

Current applications:

- Purchase prediction (Prediction): Daviet (Original paper with theory)
- Choice with context effects (Parameter inference): Bergmann, Daviet, Fehr
- Neural normalization (Model selection): Daviet, Webb
- Social preferences (Model selection): Imai, Bose, Daviet, Nave, Camerer

Note: nobody in New York yet :(

Results

Application:

Uli gave me 30 questions (after harsh negotiations) to identify the indifference set of a given subject (2 options: red/green).

He then proceeded to ask preferences (ranking) between the 2 "indifference" options and a 3rd option (blue). We can thus "see" the indifference curve.

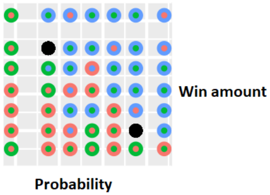

Results

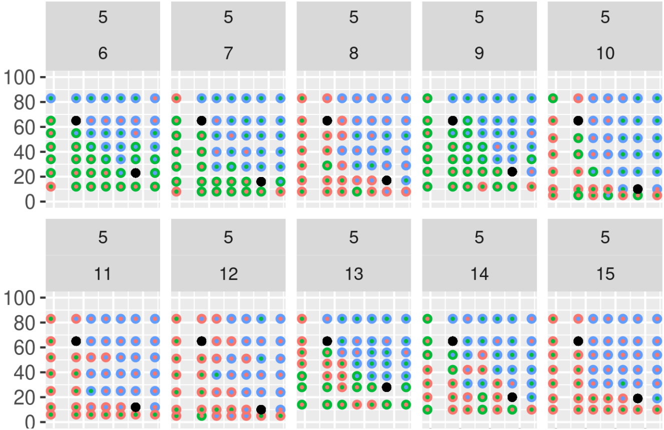

5
6
5
7
5
8
5
9
5
10
5
11
5
12
5
13
5
14
5
15
100
80
60
40
20
0
100
80
60
40
20
0

Convergence speed: SOI (red) vs. D-Optimal (green) vs. random (blue)

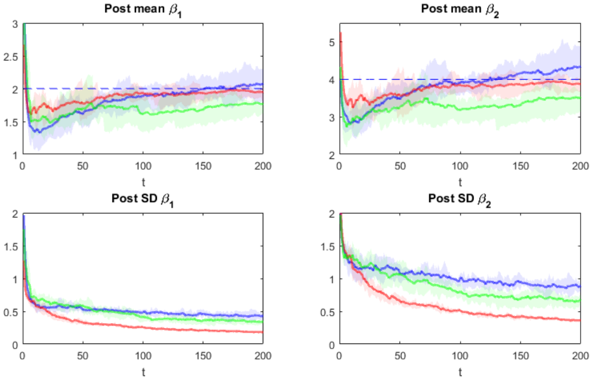

Challenges

- How to facilitate adoption ?
  - Currently Matlab and Python algorithm are provided.
- Maximizing over multiple questions in advance ?
  - Some approximate approaches are proposed (see paper).
- Possible strategic manipulation ?
  - Some different incentive scheme can be used (see paper).

Thank you & references I

References:

Daniel R Cavagnaro, Jay I Myung, Mark A Pitt, and Janne V Kujala. Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science. *Neural computation*, 22(4):887–905, 2010.

Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995.

Taisuke Imai and Colin F Camerer. Estimating time preferences from budget set choices using optimal adaptive design. *Working paper*, 2019.

Olivier Toubia, Eric Johnson, Theodoros Evgeniou, and Philippe Delquié. Dynamic experiments for estimating preferences: An adaptive method of eliciting time and risk parameters. *Management Science*, 59(3):613–640, 2013.