Sensitivity of MCMC analyses to data removal

Tin Nguyen - BAYSM 2023

Bayesian models are used to encode structure (hierarchy, expert information etc) in social sciences.

We want to check major decisions made with these methods are well founded.

A concern not specific to Bayes: social scientists want to generalize beyond observed data.

- Do the findings of an RCT conducted over certain Mexican villages generalize beyond these places?

It is concerning if conclusions substantively changes after removing a small percentage of data.

- In Mexico RCT, removing 1 out of 16,500 observations changes the treatment effect sign.

Goal: Find a small subset such that dropping it changes the conclusion.

Problem: Directly leaving out every possible small subset and refitting is computationally impractical.

Idea: [Broderick et al 2020] uses an approximation that applies to optimization-based analyses.

Problem: MCMC bring new challenges: not an optimization, and is randomized.

Our work:

- Approximate what happens after dropping the worst small data subset -> We use ideas from [Broderick et al 2020] and local robustness literature [Gustafson 1996].
- Quantify the sampling variability of this approximation -> We use block bootstrap.
- We verify our methods theoretically and empirically

Experiments

- [Angelucci et al. 2015] gathered data on microcredit and profit in Mexico.
  - Using a Bayesian model, the estimated effect (MCMC mean) of microcredit on profit is -$4.5.
- Can we change the sign of MCMC mean after removing a small data subset?
  - The full framework can handle other decision criteria.

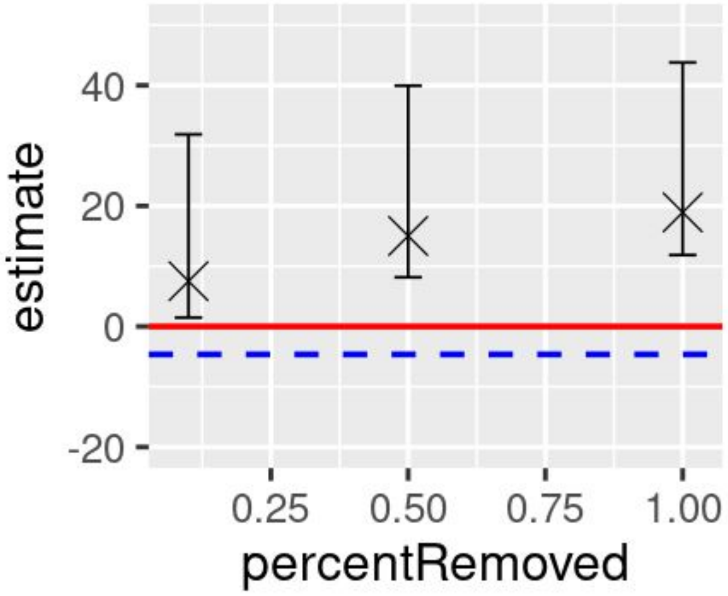

Blue: MCMC mean   Red: zero threshold

Interval: a confidence interval of what the MCMC mean would be after removing the worst small subset

-> The intervals **predict** that the MCMC mean will change sign

X: the actual MCMC mean after removing an implicated subset

-> The MCMC mean indeed changes sign, matching predictions!