Sensitivity of Bayesian Inference to Data Perturbations

Lorenzo Masoero*, William T. Stephenson*, Tamara Broderick ▲MIT, CSAIL, *equal contribution

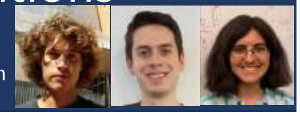

Problem: how can we quantify robustness of Bayesian inference to small changes in data?

Our contribution:

1. A new measure to quantify the impact of worst-case changes of data
2. Application I: how much can rounded data affect inferences?
3. Application II: how does likelihood misspecification affect inferences?

Setup

Bayesian formulation:

- Dataset $X = \{x_1, \dots, x_N\}$, parameter $\theta \in \Theta$
- Likelihood $p(X|\theta)$, prior $p(\theta)$
- Interested in posterior expectations:

$$E(X; h) := E_{p(\theta|X)}[h(\theta)]$$

We ask: How much can $E(X; h)$ change if we move each datapoint distance δ to form a new dataset X':

$$S_\delta(X; h) := \max_{X'} |E[X; h] - E[X'; h]| \ (*)$$

Approximation

Problem: computing $(*)$ is 😔:

- Typically non-convex in X'
- Single function evaluation = re-running inference
- Needs to be computed anew for different $h(\cdot)$

Solution: use a linear approximation to work out [2]:

$$S_\delta(X; h) \approx \delta \sum_{n=1}^{N} \left\| \mathrm{Cov}\left[h(\theta), \frac{\partial \log p(x_n \mid \theta)}{\partial x_n} \right] \right\|_2$$

Notice: covariance under original posterior $p(\theta \mid X)$

→ Easily computable given samples from MCMC for *any* h

Use cases

Case I: rounded data

- Masses of N rats measured at time T time steps are **rounded** to nearest integer gram
  → True data lies within $\delta = 0.5$ of observed x_n
- We are interested in learning the growth rate $\beta_{n,2}$ of the rats for $n = 1, \dots, N$

$$Y_{n,j} \sim N(\beta_{n,1} + \beta_{n,2} t_j, \sigma^2), j = 1, \dots, T$$
$$\beta_n \sim N(\mu_\beta, \mathbf{\Omega})$$

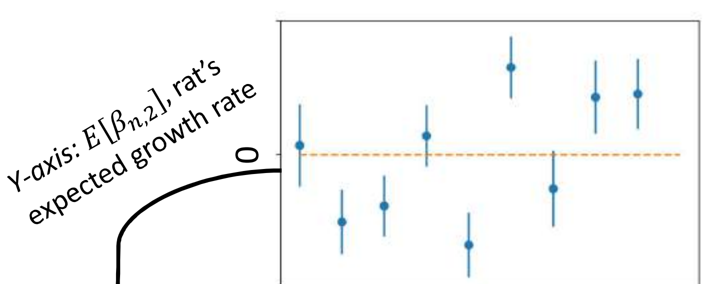

For each rat n, we show the original posterior mean, $E[\beta_{n,2}]$ (blue dot). For each rat, we compute its variation under worst-case perturbation (vertical bars). **Under our perturbation, some rats change from shrinking to growing!**

Case II: model misspecification

Intuition: dataset *X* resembles likelihood $p(X \mid \theta)$, and perturbed data *X'* resembles some other likelihood *q*

> ➢ **Conjecture**: data sensitivity is sensitivity to likelihood misspecification

Experiment: Generate data from a Gaussian mixture model with K=2 components. Our perturbation discovers brittleness of GMMs: under our worst case perturbation, the posterior concentrates on larger K as dataset size increases!

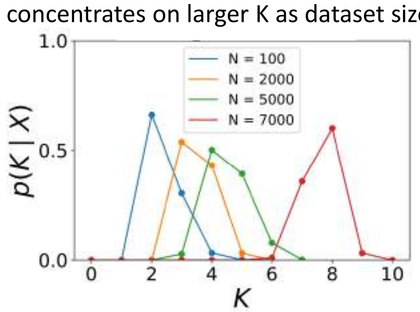

→**GMMs are extremely sensitive to misspecification!** We also show that a proposed robustification of GMM's [1] is robust (see our paper!)

References

[1] Miller, Jeffrey W and Dunson, David B, Robust Bayesian inference via coarsening, *Journal of the American Statistical Association*, 2018

[2] Giordano, R, Broderick, T, and Jordan, MI. Covariances, robustness, and variational Bayes, *Journal of Machine Learning Research*, 2018