1.5em 0pt

Calibrated Approximate Bayesian Inference

Hanwen Xing [1] Geoff K. Nicholls [1] Jeong Eun Lee [2]

[1]Department of Statistics, University of Oxford, UK

[2]Department of Statistics, University of Auckland, New Zealand

June 5, 2019

Motivation

- Approximation schemes are key to Bayesian inference.
- Does a α level approximate credible set have the right coverage?
- Let C_y and $\tilde{C}_y$ be the exact and approximate α level credible set,

$$\alpha = E_\pi(\mathbb{1}_{\phi \in C_y}) = \int_\Omega \mathbb{1}_{\phi \in C_y} \pi(\phi|y) d\phi.$$

$$\alpha = E_{\tilde{\pi}}(\mathbb{1}_{\theta \in \tilde{C}_y}) = \int_\Omega \mathbb{1}_{\theta \in \tilde{C}_y} \tilde{\pi}(\theta|y) d\theta.$$

Motivation

- We also define

$$b(y) = \Pr(\phi \in \tilde{C}_Y | Y = y) = \int_{\Omega} \mathbb{1}_{\phi \in \tilde{C}_y} \pi(\phi|y) d\phi$$

  be the *operational* coverage $\tilde{C}_Y$ achieves.
- We want to estimate $b(y_{obs})$, the true Bayesian coverage of the approximate credible set, as it measures the reliability of approximation at the observed data.

Method

- **Regression approach** Let $\{\phi_i, y_i\}_{i=1}^{M}$ be samples from the generative model $\pi(\phi)p(y|\phi)$, let $\tilde{C}_{y_i}$ be an approximate credible set for y_i, and $c_i = \mathbb{1}_{\phi_i \in \tilde{C}_{y_i}}$. Conditional on y_i,

$$c_i \sim Bernoulli(b(y_i)), \quad b(y_i) = \Pr(\phi_i \in \tilde{C}_{Y_i} | Y_i = y_i)$$

- **Weighted-sample approach** Estimate $b(y)$ by first approximately sampling from the exact posterior using Annealed Importance Sampling algorithm. This leverages our ability to draw samples from the approximate posterior $\tilde{\pi}(\phi|y_{obs})$ as a good starting point for the AIS iteration.

Example

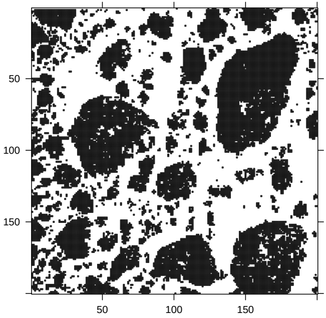

Figure: Icefloe image

We approximate the free boundary condition likelihood by a toroidal boundary condition likelihood.

Example

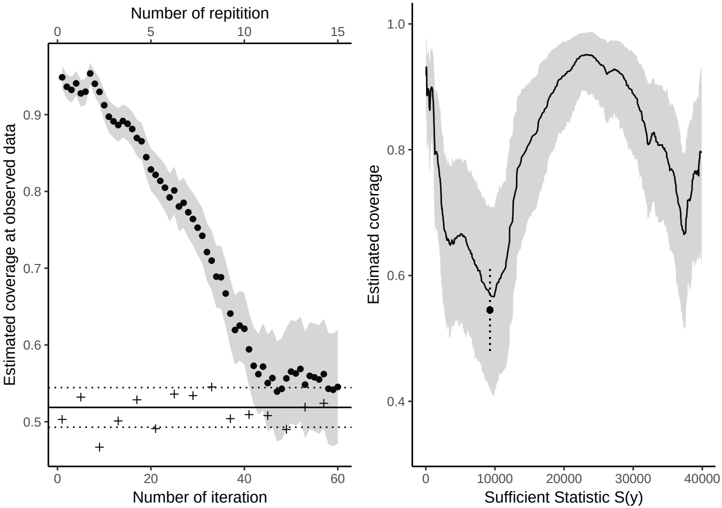

Figure: Left: AIS Right:Regression