Shiny

Colin Rundel

2019-11-25


Shiny
www.rstudio.com

Shiny

- is a framework for writing web apps in R
- easiest way to create interactive tools using R
- can be run locally in RStudio, hosted on RStudio's shinyapps.io or your own server, or even embedded in RMarkdown documents.

Shiny Resources

All of these can be found at RStudio's Shiny homepage:

- Shiny Gallery - https://shiny.rstudio.com/gallery/
- Shiny Articles - https://shiny.rstudio.com/articles/
- Function Reference - https://shiny.rstudio.com/reference/shiny/1.4.0/
- Mastering Shiny [WIP] - https://mastering-shiny.org/

A "Fair" Coin and the Beta-Binomial model

Let's imagine a situation where I have a coin that I claim is fair, you are less certain - you can ask me to flip the coin n times and we will record the number of "heads" as x. After this experiment what should you believe about the fairness of my coin?

We can model this using the conjugate beta-binomial model,

$$\begin{aligned} \pi(p) &\sim \text{Beta}(\alpha, \beta) \\ x \,|\, p &\sim \text{Binom}(p, n) \end{aligned}$$

it then follows that

$$p \,|\, x \sim \text{Beta}(\alpha + x, \beta + N - x)$$

Live Demo

Approximate Bayesian Computation (ABC)

Offered without any theory or proof, consider the following algorithm:

1. Sample m values from each of the given prior distribution(s), $\boldsymbol{\theta}^{(m)}$.
2. For each of the m draws, simulate the data from the assumed data generative process, $\boldsymbol{x}^{(m)}$.
3. If the simulated data is sufficient close to the observed data, keep $\boldsymbol{\theta}^{(m)}$.
4. The distribution of the kept $\boldsymbol{\theta}^{(m)}$s will approximate the posterior distribution $\boldsymbol{\theta}|\boldsymbol{x}$.

Beta-Binomial ABC

Our particular use case provides a straight forward application of this method,

1. Draw from the prior:

```
p = rbeta(m, alpha, beta)
```

1. Simulate the data:

```
x_sim = rbinom(m, n, p)
```

1. Keep the matches:

```
abc_post = p[x_sim == x]
```

Why do we care?