

CS-E407520 - Special Course in Machine Learning and Data Science: Bayesian Workflows

Session 2: Discussion of Workflow Diaries & Primer on Prior Choices

April 29, 2024



Schedule for today's session

Time	Activity
45 min	Discussion of workflow diaries
10 min	Break
20 min	Discussion of workflow diaries (ctd.)
15 min	Primer for next workflow steps

Discussion of workflow diaries

Peer Discussion (10-15min)

Form groups of 2-3 students

If you are the presenter - take 3mins to give a quick overview of what you worked on

If you are the audience - ask one question to the person who just shared their work

Repeat the process until each of you has had the chance to present some of their work

Seminar Discussion (30min)

When it is your turn, please

1. present your workflow diary by sharing your screen
2. summarise briefly what you decided to work on in the previous week and show some results
3. report issues that you encountered
4. ask questions

Reminder: Seminar Discussions

Note

It is important to engage with the workflow steps, be prepared to present and actively participate in the seminar sessions to pass the course!

Presenting the current status of your project is a great opportunity for you to share your thoughts and get direct advice or help when you feel stuck - make the most out of it by actively preparing and engaging during the session.

It is ok to not be able to present max. 2 times from Session 2 to Session 6.

Seminar Discussion (30min)

When it is your turn, please

1. present your workflow diary by sharing your screen
2. summarise briefly what you decided to work on in the previous week and show some results
3. report issues that you encountered
4. ask questions

Let's take a break! (10 min)

Some suggestions for recharging during breaks

- Move your body
- Open a window or go outside
- Drink some water
- Try to avoid checking e-mails, messengers, or social media

Seminar Discussion (20min)

When it is your turn, please

1. present your workflow diary by sharing your screen
2. summarise briefly what you decided to work on in the previous week and show some results
3. report issues that you encountered
4. ask questions

This week in context

Choosing an initial model

This week: Prior Choices

Model checking

Extending Models & Model Selection

Interpreting and Presenting Model Results

Primer: Prior Choices

Your next steps are:

Choose a prior

- specify generative priors (that can be sampled from) for each parameter in your model
- justify these choices

Evaluate prior

- use prior predictive visual checks
- interpret the plots and document any issue with priors

Adjust prior

- based on observed issues, change your prior and repeat

The impact of the prior

- Priors are part of the model and can have various impacts
- The generic different level of informativeness
 1. flat prior
 2. Super vague proper prior, $\text{normal}(0, 10^6)$
 3. Weakly informative, $\text{normal}(0, 10)$
 4. Weakly informative, $\text{normal}(0, 1)$
 5. Specific informative prior $\text{normal}(0.4, 0.2)$
- But how informative any of those priors are, depends the data, scale, transformations, parameterisation, observation model type, question being asked...

The good of priors

- Improper priors are not bad per se: common starting point
 - However, once data and initial modeling question are in, good priors tend to...
1. Be generative (finite integral, allows prior predictive)
 2. Have appropriately fat tails (more forgiving when ignorant)
 3. Softly exclude nonsensical regions in parameter/outcome space
 4. Give sensible prior predictives (e.g., interpretable summary stats)
 5. Encode dependence structures (e.g. slopes and intercepts)
 6. Encode expert knowledge or lean on previous research/best practices

And the bad...

- Try to avoid:
 1. Too thin tails (unless carefully chosen)
 2. Invariance, maximum entropy, Jeffrey's priors
 3. Generic priors from PPL
 4. Many independent components

Converting knowledge into Priors

- Expressing prior knowledge mathematically is hard (even for seasoned modellers)
- Knowledge usually exists on either parameters, causal effects or predictions
- Packages such as PreliZ make it easy to encode information
- Chapter 15 gives a good workflow for converting knowledge to the prior

Prior predictive checks

- For examples of prior predictive checks in the outcome space, see for example chapter 3, 15
- in `brms`, use `sample_prior = "only"` to get draws from the prior
- `bayesplot` has several functions for prior predictive checking

Prior predictive checks

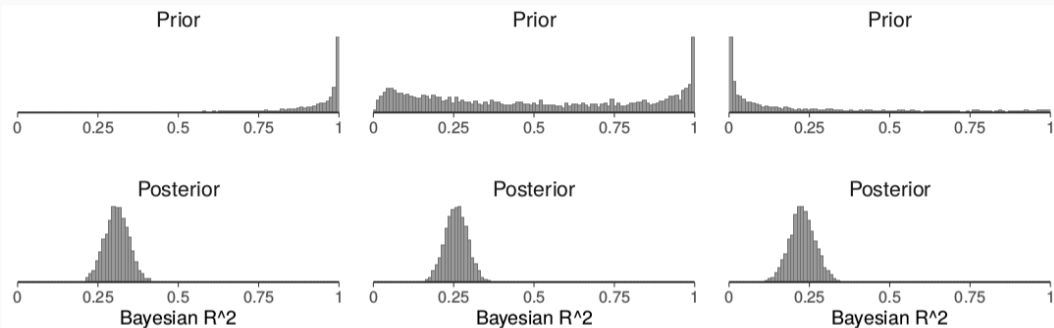


Figure 12.11 *Prior and posterior distribution of Bayesian R^2 for the regression predicting student grades from many predictors, using three different priors for the coefficients: (a) default weak prior, (b) normal prior scaled with the number of predictors, and (c) regularized horseshoe prior.*

From Gelman, Hill and Vehtari (2021)

- Sometimes it is easier to show predicted quantities via summary statistics such as the R^2 for regression models

Resources for the Week

- BRMS demo for generating from prior


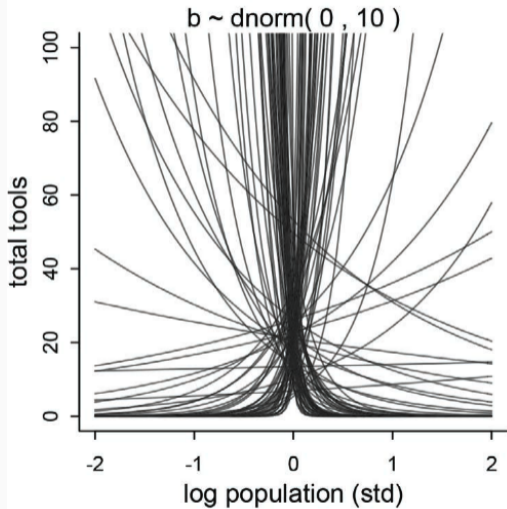
Choosing priors

- Stan prior recommendations
- Stan distribution visualizer
- Distribution explorer
- PreliZ

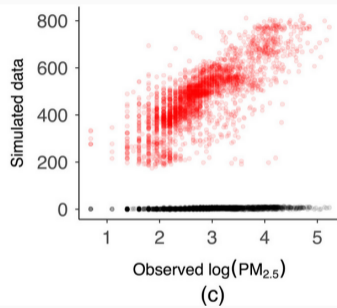
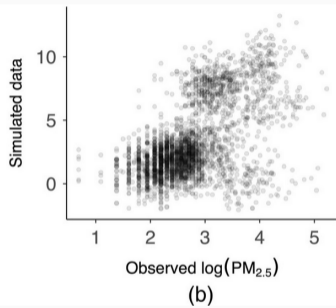

Evaluating priors

- bayesplot
- Aki's case study

Examples



Examples



From Gabry et al. (2019)