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

Session 3: Discusson of prior choices & primer on model checking

May 6, 2024

Schedule for today's session

Time	Activity
5 min	Course organisation
45 min	Discussion of workflow diaries
10 min	Break
30 min	Primer for next workflow steps

Course organisation

Halfway point

- we are approaching the halfway point of the course
- to see how you are progressing, we will take a look at your workflow diaries so far
- submit your in-progress diaries on MyCourses by next session (Mon 13.05.)

Workflow Diary Discussion

Discussion of workflow diaries

When it is your turn:

1. present your workflow diary via projector or screen
2. briefly summarise your data and research question
3. teaching staff will guide discussion of your implementation of the workflow steps

Your goals were:

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

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

Primer: Model Checking

Your next steps are:

Parameter estimate and model fit checks

- check model fitting diagnostics
- print and/or plot the marginal (and biviariate) posteriors for important parameters
- comment on the parameter estimates; do they make sense? are they surprising?

Posterior predictive checks

- create appropriate plots to check how model predictions compare to observations
- comment on any potential deficiencies in your model

Influence and sensitivity checks

- use cross validation to identify influential data points (justify if not applicable)
- perform a prior/likelihood sensitivity analysis and comment on results

Running Example

- randomised placebo-controlled trial
- effect of drug (sulindac) on intestinal polyps over 12 months
- 22 participants (11 treatment, 11 placebo)
- initial model: `n_polyps ~ treatment * month, family = "gaussian"`

Parameter estimate checks

- printing the summary output of a model fit is an important first check
- evaluating if the posteriors seem reasonable checks how well the model is capturing your implicit prior information

Parameter estimate checks

Initial model:

variable	mean	5%	95%	rhat	ess_bulk
b_Intercept	42.04	23.9	60.4	<1.01	4110
b_month	-0.55	-3.5	2.5	<1.01	3705
b_treatmentsulindac	-5.88	-20.9	9.1	<1.01	4293
b_month:treatmentsulindac	-2.08	-5.7	1.7	<1.01	3649

Posterior predictive checks

- visualisation based
- type depends on model properties, particularly outcome variable type (continuous, categorical, ordinal)
- `bayesplot` is a good starting point, but don't just use the default density plot, look deeper

Posterior predictive checks

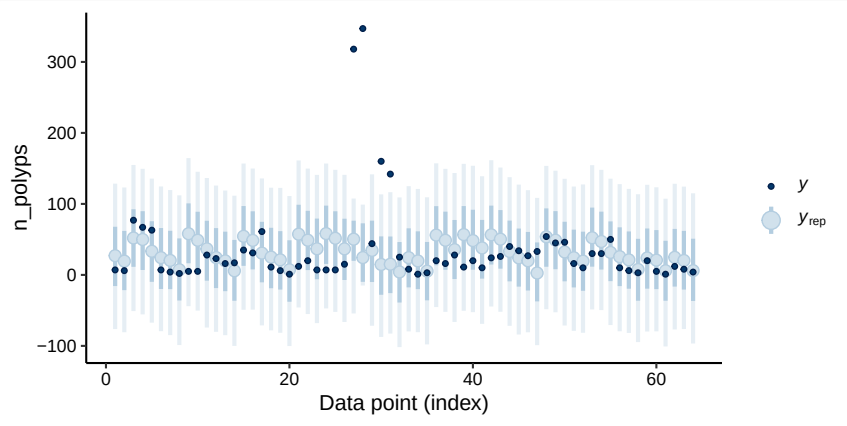

LOO intervals: Initial model

Posterior predictive checks

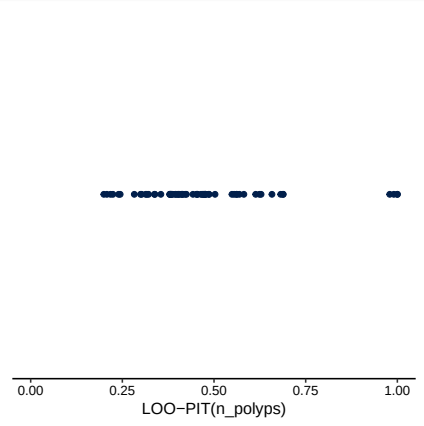

LOO PIT: Initial model

Posterior predictive checks

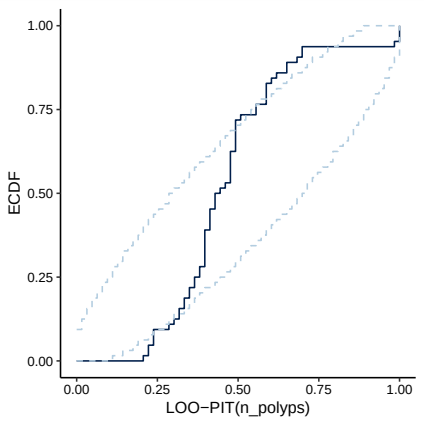

LOO PIT ECDF: Initial model

Posterior predictive checks

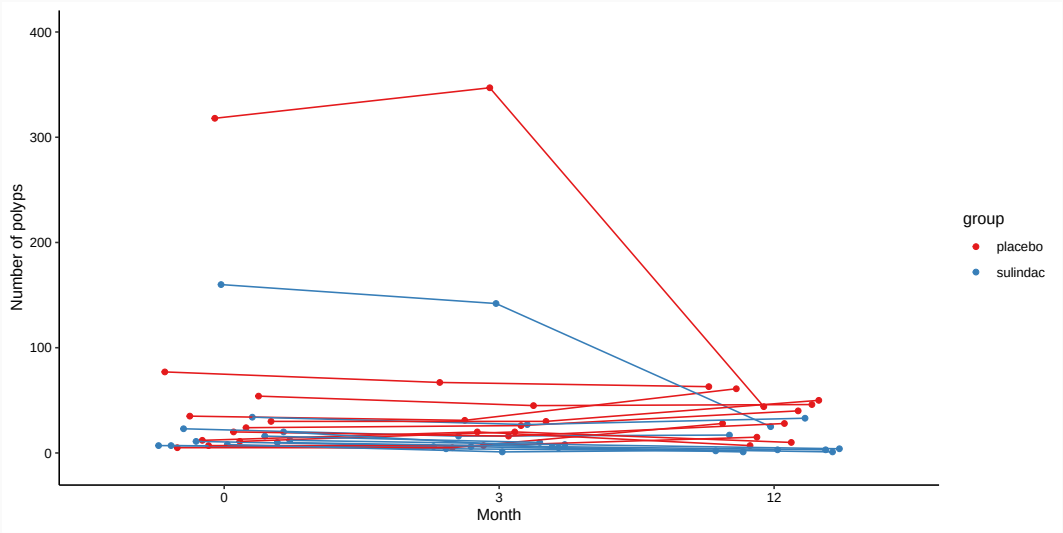

Observations

Posterior predictive checks

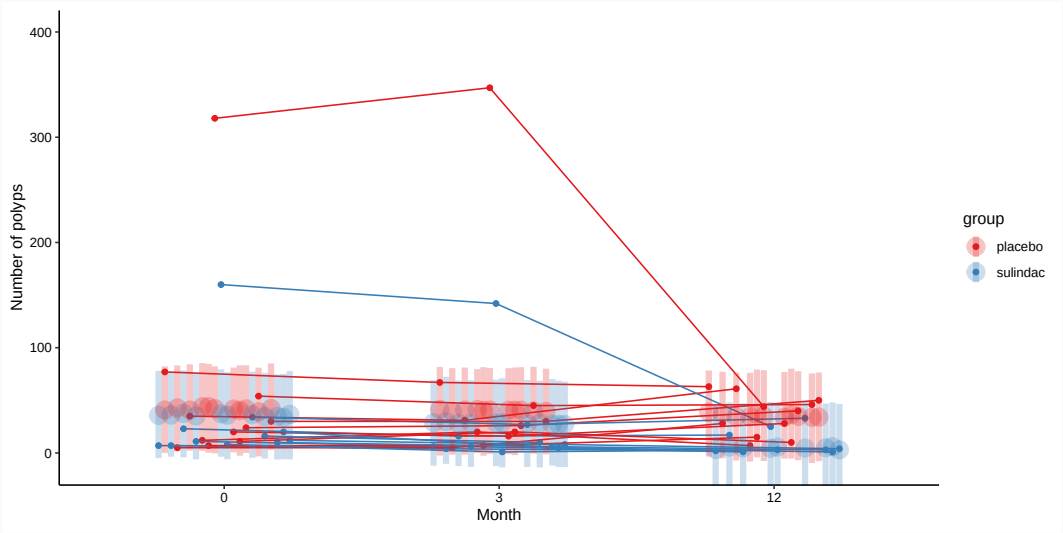

Predictions: Initial model

Posterior predictive checks

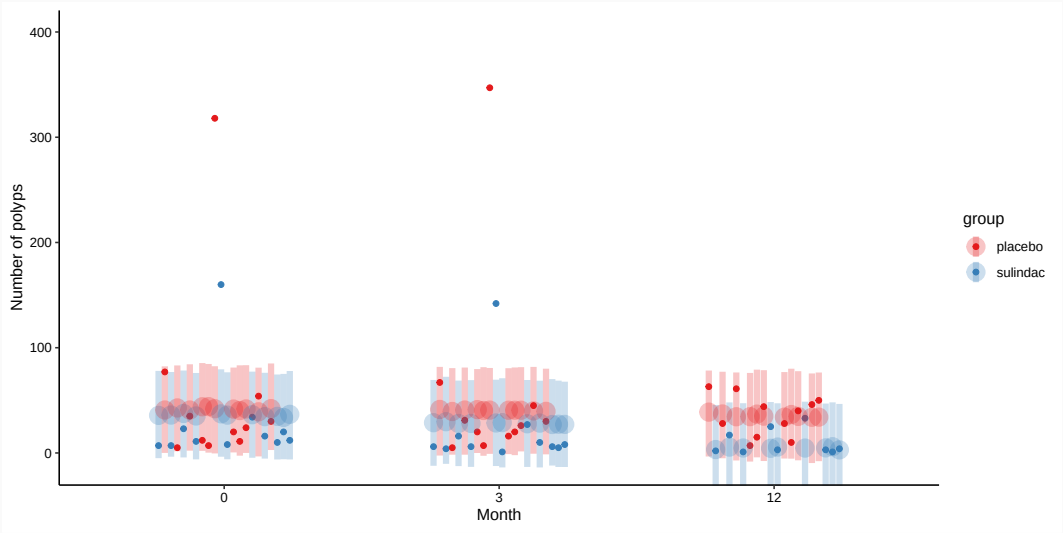

Predictions: Initial model

Posterior predictive checks

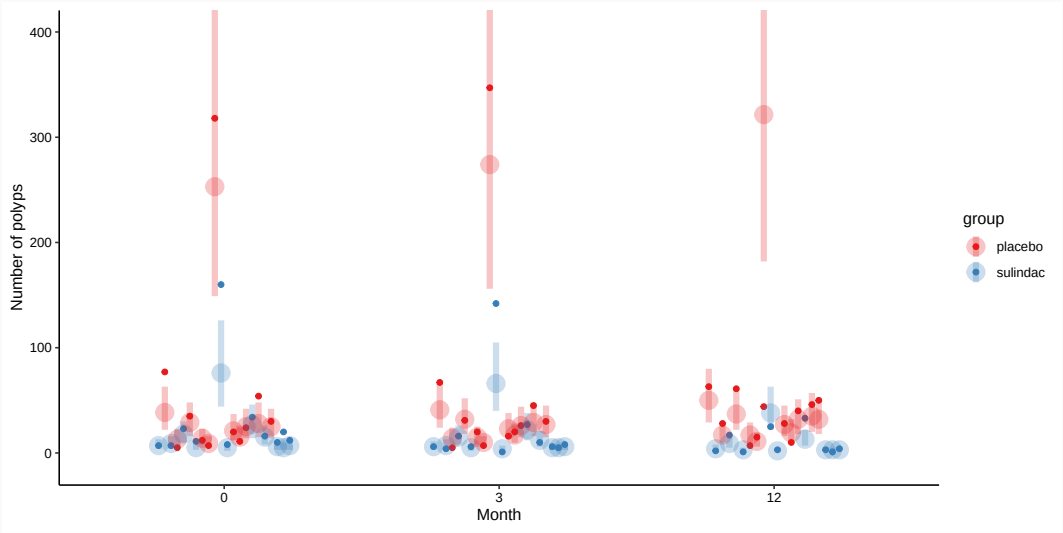

Predictions: Later model

Posterior predictive checks

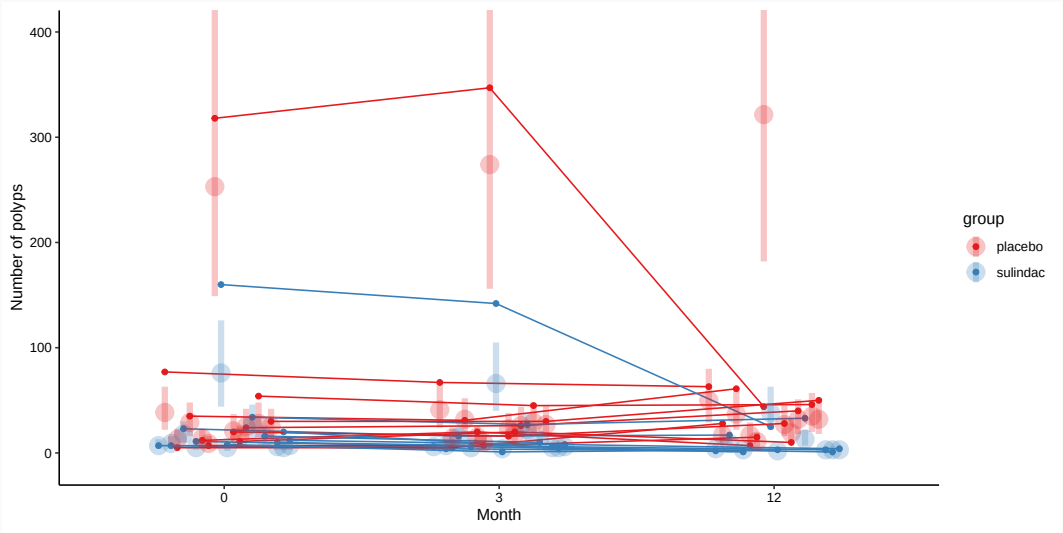

Predictions: Later model

Posterior predictive checks

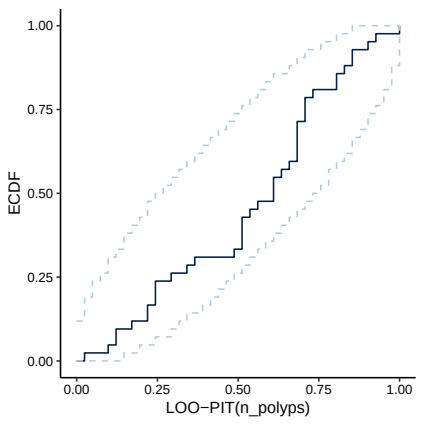

Calibration: Later model

Influential observation checking

- the posterior may shift drastically when an influential observation is left out
- leave-one-out cross-validation can indicate if any observations are highly influential
- the `loo` package implements efficient approximate LOO-CV (available in `brms`)

Influential observation checking

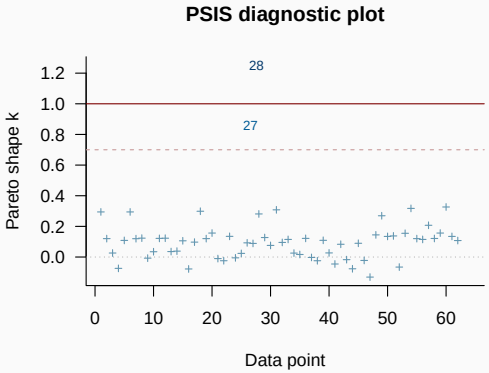

Pareto k diagnostics: Initial model

Prior and likelihood sensitivity checks

- small changes to the prior and/or likelihood can impact the posterior
- `priorsense` implements one type of sensitivity analysis

Prior and likelihood sensitivity checks

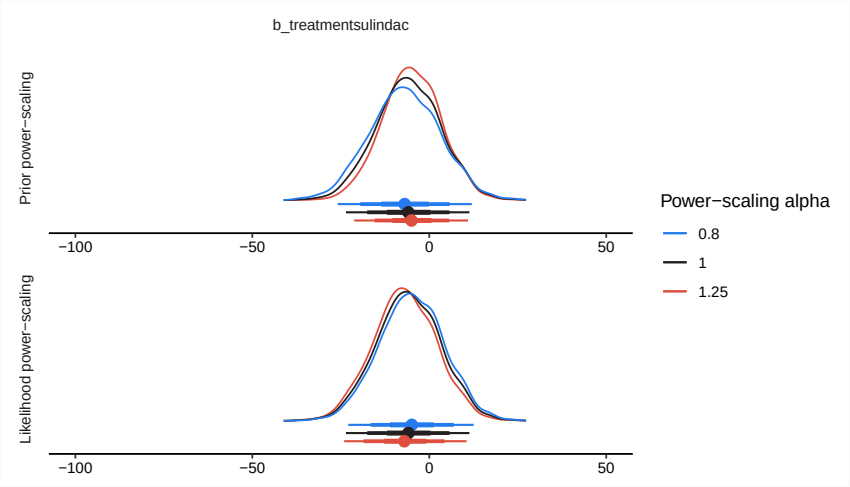

Power-scaling sensitivity: Initial model

Prior and likelihood sensitivity checks

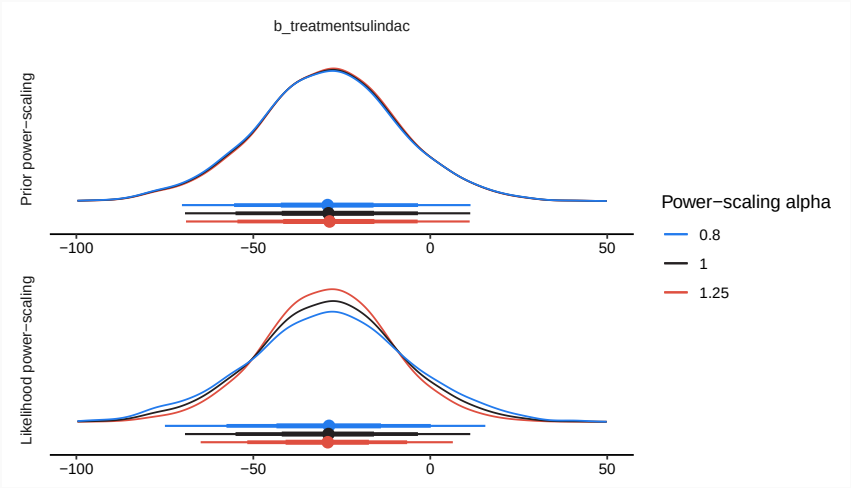

Power-scaling sensitivity: Adjusted priors

Resources

Case studies

- diabetes
- roaches
- birthdays
- nabiximols

Readings

- Workflow book Chapter 5, 9, 11

Relevant papers

- Säilynoja et al. (2022); Calibration checks
- Kallioinen et al. (2023); Power-scaling sensitivity checks
- Vehtari et al. (2016); Approximate leave-one-out CV