


Problem setting

How is y influenced by x_1 and x_2 ?



How good is my model?

How can I estimate the elpd?

What are the challenges?



Methods for approximating the elpd in LOGO CV



Case study: Item response dataset

Setting: Verbal aggression questionnaire

- 316 participants x 24 items
- Ordinal response: Each respondent judges whether they *would want to* or *would actually curse / scold / shout* in four frustrating situations.
- Balanced grouping structure: each participant has 24 observations

Two models: different complexities

Model 1: (btype || id) $\Rightarrow \dim(\phi_j) = 3$

Model 2: (btype + mode + situ || id) $\Rightarrow \dim(\phi_j) = 7$

Take away:

- Model 1: MC + PSIS is accurate and cheapest – no gradient evaluations required
- Model 2: MC + PSIS breaks down; LA + PSIS is the best choice

“A bus fails to stop for me”

