Prior Predictive Checks: Grammaticality Judgment Example

Bayesian Mixed Effects Models with brms for Linguists

Job Schepens

2025-12-17

Table of contents

1 Prior Predictive Checks for Binary Data **1**
- 1.1 Setup . . . 2
- 1.2 Fitting Prior Only Model . . . 2
- 1.3 Prior Predictive Checks . . . 2
  - 1.3.1 Visual Checks . . . 2
  - 1.3.2 Prior Predictive Accuracy Distribution . . . 3
- 1.4 Prior Distributions . . . 3
  - 1.4.1 Intercept Prior . . . 3
  - 1.4.2 Effect Size Prior . . . 4
- 1.5 Random Effects Distributions . . . 5
  - 1.5.1 Subject Random Intercepts . . . 5
  - 1.5.2 Subject Random Slopes . . . 6
  - 1.5.3 Random Effects Comparison . . . 7
- 1.6 Interpretation . . . 7
  - 1.6.1 Good Signs (Prior is Reasonable) . . . 7
  - 1.6.2 Problems to Watch For . . . 7
- 1.7 Summary . . . 8

1 Prior Predictive Checks for Binary Data

This document demonstrates how to validate priors for a Bayesian binary model (grammaticality judgments) before fitting to data.

1.1 Setup

1.2 Fitting Prior Only Model

1.3 Prior Predictive Checks

1.3.1 Visual Checks

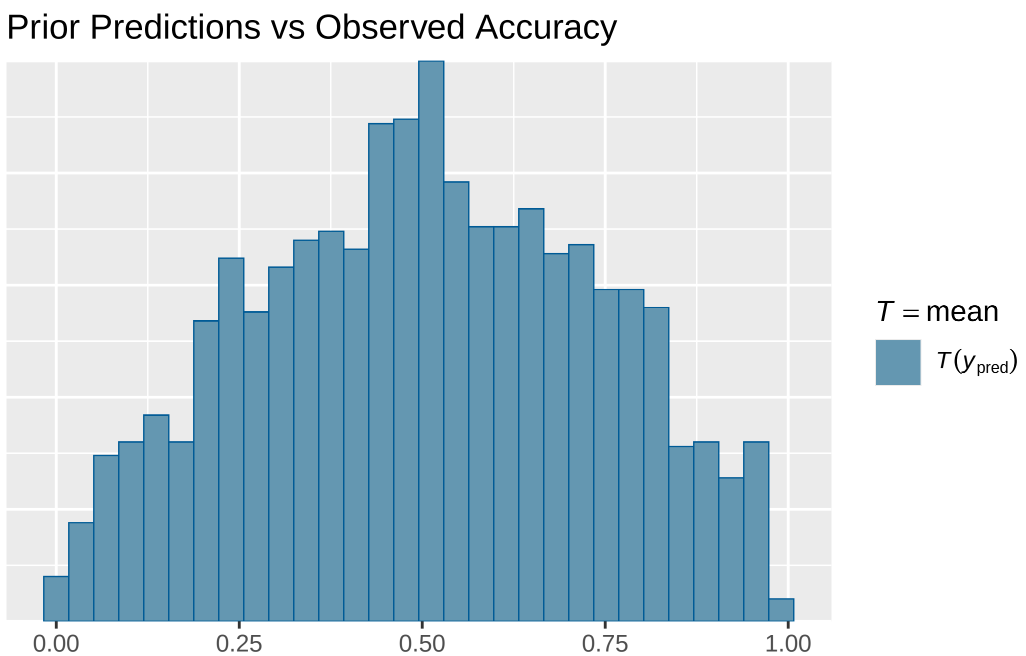

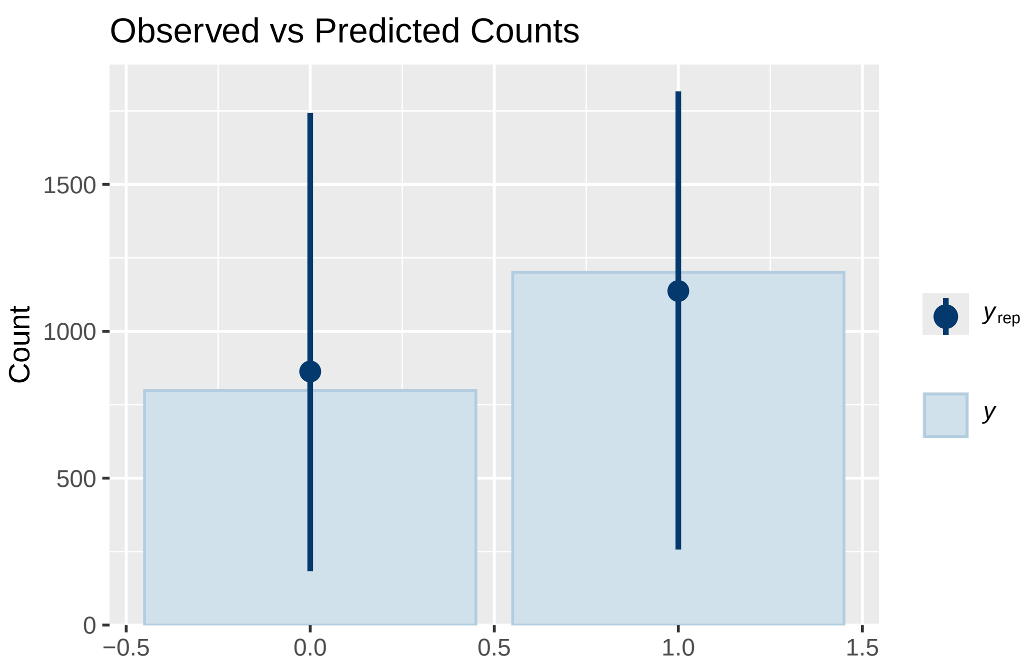

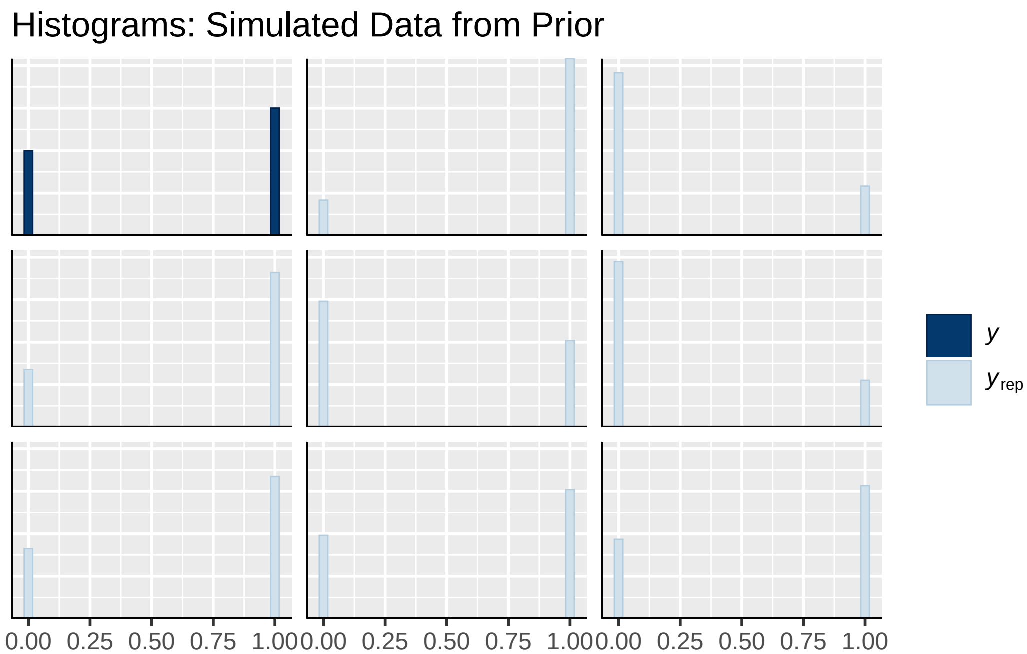

1.3.2 Prior Predictive Accuracy Distribution

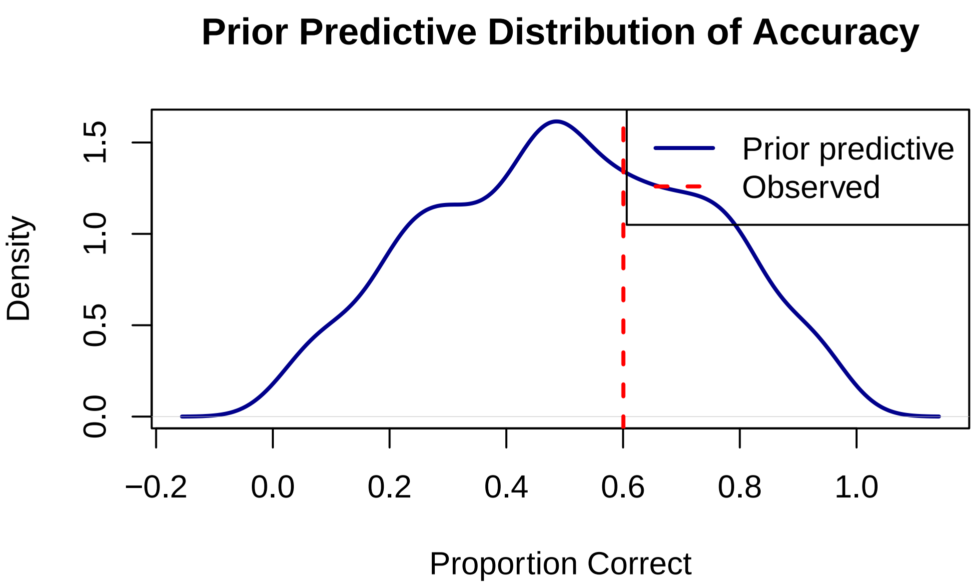

1.4 Prior Distributions

Following Kurz's approach, we extract prior samples using `as_draws_df()` from the posterior package.

1.4.1 Intercept Prior

```
Intercept prior (log-odds scale):

       2.5%         50%       97.5%
-2.98728595  0.04194952  3.19327008
```

```
Intercept prior (probability scale):

      2.5%        50%      97.5%
0.04800357 0.51048584 0.96058023

Prior expects 5 - 96 % baseline accuracy (median 51 %)
```

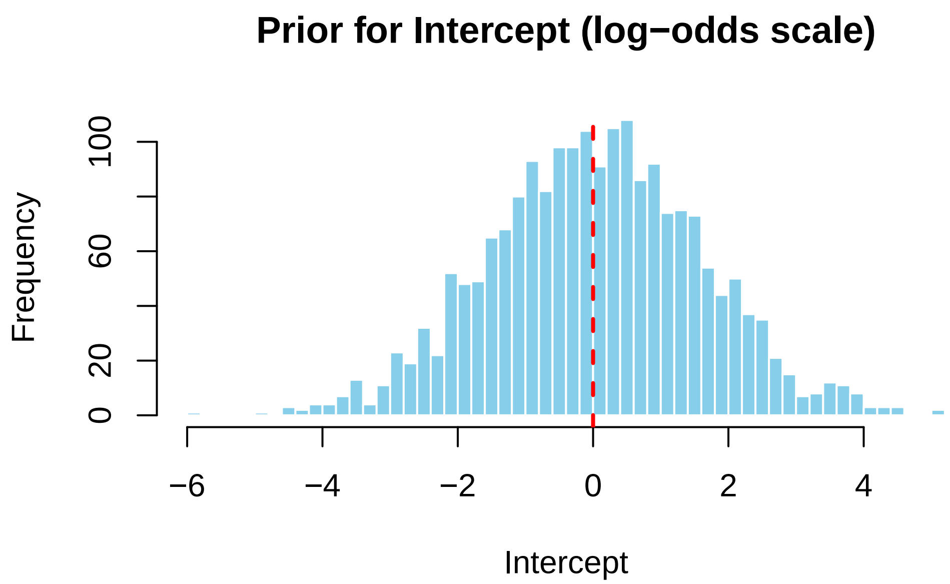

1.4.2 Effect Size Prior

```
Condition effect prior (log-odds scale):

       2.5%         50%       97.5%
-1.93961957 -0.01510226  1.96317214


Condition effect (odds multiplier):

     2.5%       50%     97.5%
0.1437586 0.9850112 7.1218829
```

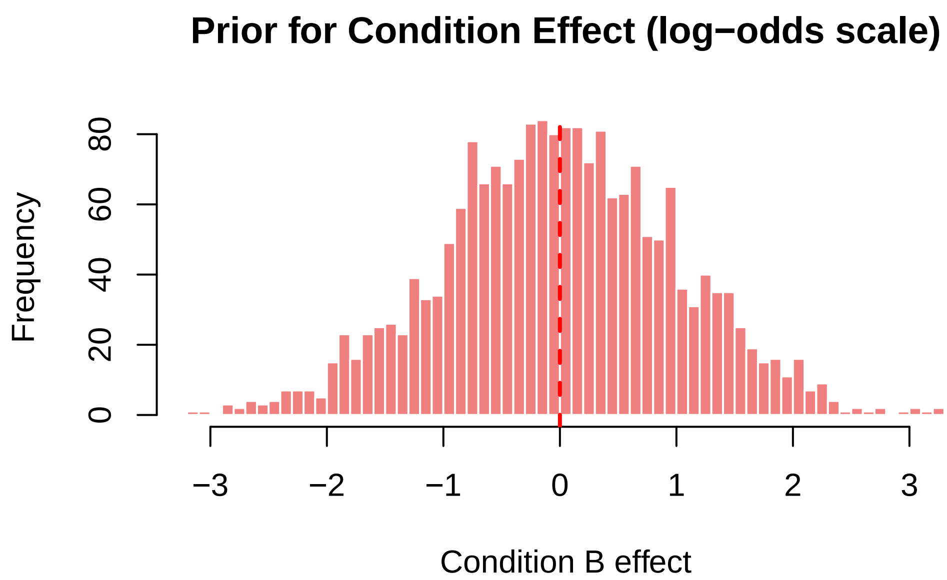

1.5 Random Effects Distributions

For prior predictive checks, we examine the **implied distribution** of subject-specific parameters by extracting the hyperprior SDs and simulating random effects.

1.5.1 Subject Random Intercepts

```
Subject random intercept SD prior:

      2.5%        50%      97.5% 
0.02900055 0.68343780 3.86970012 


Implied subject random intercepts (log-odds scale):

        2.5%          50%        97.5% 
-1.326931618  0.006294215  1.392768861 


Implied subject-specific accuracy (probability scale):

     2.5%       50%     97.5% 
0.2167034 0.5120586 0.8076354 

Prior implies subject accuracy ranges from 22 % to 81 %
```

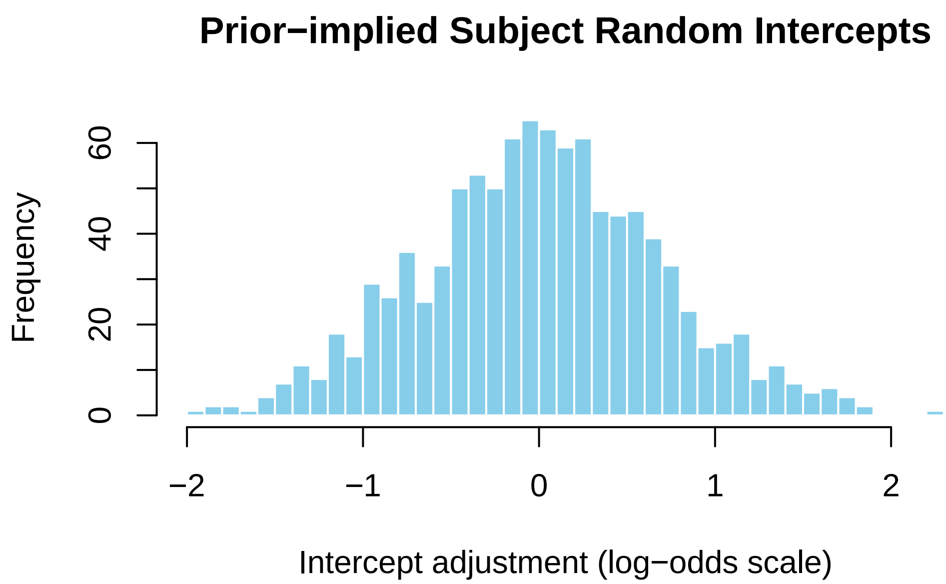

1.5.2 Subject Random Slopes

```
Subject random slope SD prior:

      2.5%        50%      97.5%
0.02717447 0.68593886 3.63193775


Implied subject random slopes (log-odds scale):

       2.5%         50%       97.5%
-1.36609138  0.03762538  1.30733868


Implied subject-specific effects (odds multipliers):

     2.5%       50%     97.5%
0.2551021 1.0383422 3.6963235

Weak effect subjects (2.5%):  0.26 × odds multiplier

Average effect subjects (50%):  1.04 × odds multiplier

Strong effect subjects (97.5%):  3.7 × odds multiplier
```

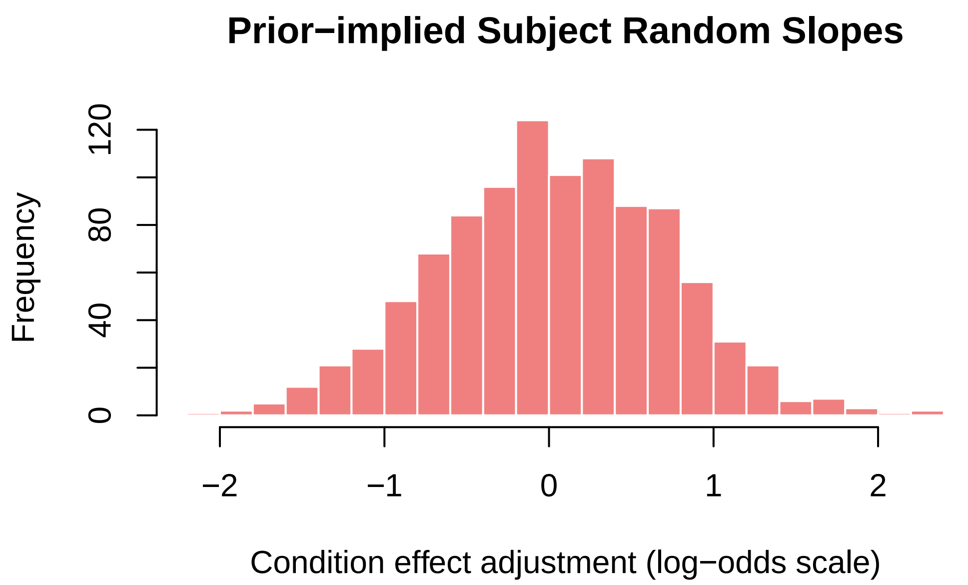

1.5.3 Random Effects Comparison

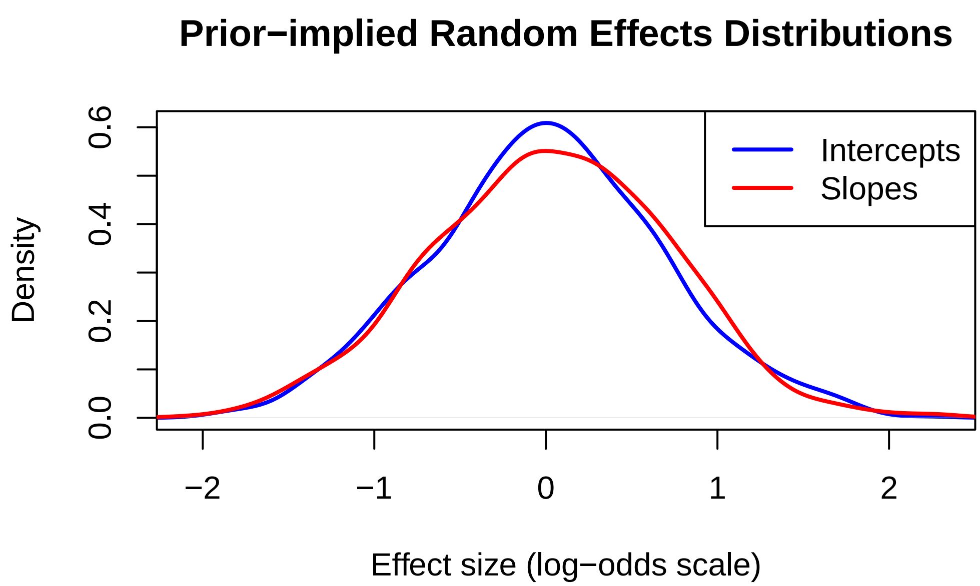

1.6 Interpretation

1.6.1 Good Signs (Prior is Reasonable)

- Prior generates ~50% baseline accuracy (Intercept 0)
- Condition effect varies but is typically moderate
- Between-subject accuracy ranges 30-70% or similar
- No impossible values (all between 0 and 1)

1.6.2 Problems to Watch For

- Mean accuracy » 90%: intercept prior too high
- All subjects 50% ± 1%: intercept SD too small

- Very weak prior predictions: slopes prior too small

1.7 Summary

Before fitting your model to actual data, always validate that your priors: 1. Generate reasonable predictions for your domain 2. Allow sufficient flexibility for the data to inform the posterior 3. Respect constraints (e.g., 0-1 for probabilities) 4. Account for variation across subjects/items

Adjust priors as needed and rerun these checks.