Prior Predictive Checks: Reaction Time Example

Bayesian Mixed Effects Models with brms for Linguists

Job Schepens

2025-12-17

Table of contents

1 Prior Predictive Checks for Reaction Time Data

This document demonstrates how to validate priors for a Bayesian RT model before fitting to data.

1.1 Why Prior Predictive Checks Matter

Before fitting your model to actual data, validate that your priors generate sensible predictions. This is crucial because:

- A prior that's too restrictive can prevent the model from learning from data
- A prior that's too permissive might not regularize your estimates
- It's much easier to revise priors before fitting than after

1.2 Setup

1.3 Fitting Prior Only Model

```
## Prior Predictive Checks

### Visual Checks

::: {.cell}
::: {.cell-output-display}
```

```
![](02_prior_predictive_checks_rt_files/figure-pdf/pp-checks-1.pdf)
:::

::: {.cell-output-display}
![](02_prior_predictive_checks_rt_files/figure-pdf/pp-checks-2.pdf)
:::

::: {.cell-output-display}
![](02_prior_predictive_checks_rt_files/figure-pdf/pp-checks-3.pdf)
:::

::: {.cell-output-display}
![](02_prior_predictive_checks_rt_files/figure-pdf/pp-checks-4.pdf)
:::

::: {.cell-output-display}
![](02_prior_predictive_checks_rt_files/figure-pdf/pp-checks-5.pdf)
:::
:::

## Prior Distributions

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

::: {.cell}

:::

### Intercept Prior

::: {.cell}
::: {.cell-output .cell-output-stdout}
```

Intercept prior (log scale):

:::

::: {.cell-output .cell-output-stdout}

```
2.5%      50%    97.5%
```

2.875128 6.020023 8.923694

:::

::: {.cell-output .cell-output-stdout}

Intercept prior (RT scale in milliseconds):

:::

::: {.cell-output .cell-output-stdout}

```
  2.5%        50%      97.5%
```

17.72769 411.58799 7507.77385

:::

::: {.cell-output-display}

:::
:::

Effect Size Prior

::: {.cell}

::: {.cell-output .cell-output-stdout}

Condition effect prior (log scale):

:::

::: {.cell-output .cell-output-stdout}

```
   2.5%          50%       97.5%
-0.99889479 -0.01414723 0.94791795
```

:::

::: {.cell-output .cell-output-stdout}

Condition effect prior (RT scale in ms):

:::

::: {.cell-output .cell-output-stdout}

```
 2.5%       50%     97.5%
0.3682863 0.9859524 2.5803317
```

:::

::: {.cell-output-display}
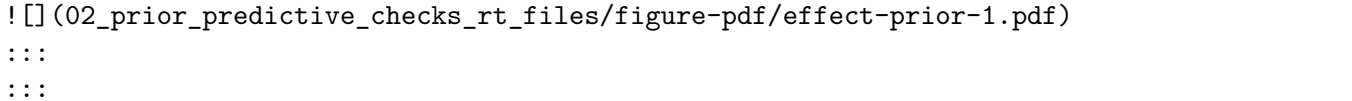
:::
:::

Residual Noise Prior

::: {.cell}
::: {.cell-output .cell-output-stdout}

Residual noise prior (sigma, log scale):

:::

::: {.cell-output .cell-output-stdout}

```
  2.5%        50%      97.5%
```

0.02573465 0.69876976 3.81711126

:::
:::

Random Effects Distributions

For prior predictive checks, we examine the **implied distribution** of subject-specific parameter

Subject Random Intercepts

::: {.cell}
::: {.cell-output .cell-output-stdout}

Subject random intercept SD prior:

:::

::: {.cell-output .cell-output-stdout}

```
  2.5%        50%      97.5%
```

0.02519548 0.67840017 3.61983168

:::

::: {.cell-output .cell-output-stdout}

Implied subject random intercepts (log scale):

:::

::: {.cell-output .cell-output-stdout}

```
    2.5%          50%        97.5%
```

-1.317150788 0.006247821 1.382502744

:::

::: {.cell-output .cell-output-stdout}

Implied subject-specific RTs (milliseconds):

:::

::: {.cell-output .cell-output-stdout}

```
 2.5%       50%     97.5%
```

110.2634 414.1676 1640.1214

:::

::: {.cell-output .cell-output-stdout}

Prior implies subject RTs range from 110 to 1640 ms

:::

::: {.cell-output-display}

:::
:::

Subject Random Slopes

::: {.cell}
::: {.cell-output .cell-output-stdout}
Subject random slope SD prior:

:::

::: {.cell-output .cell-output-stdout}
```
 2.5%       50%     97.5%
```
0.0249993 0.6998014 3.6529919

:::

::: {.cell-output .cell-output-stdout}
Implied subject random slopes (log scale):

:::

::: {.cell-output .cell-output-stdout}
```
   2.5%          50%       97.5%
```
-1.39369961 0.03838577 1.33375954

:::

::: {.cell-output .cell-output-stdout}
Interpretation: Condition effect varies by subject

:::

::: {.cell-output .cell-output-stdout}

Small effect subjects (2.5%): 0.248 × multiplier

:::

::: {.cell-output .cell-output-stdout}

Average effect subjects (50%): 1.039 × multiplier

:::

::: {.cell-output .cell-output-stdout}

Large effect subjects (97.5%): 3.795 × multiplier

:::

::: {.cell-output-display}

:::
:::

Residual Noise Distribution

::: {.cell}
::: {.cell-output-display}

:::

::: {.cell-output .cell-output-stdout}

Interpretation: Red line shows typical residual SD (0.3) for RT data Our prior puts 25.4 % probability below this value “‘

::: :::

1.3.1 Random Effects Comparison

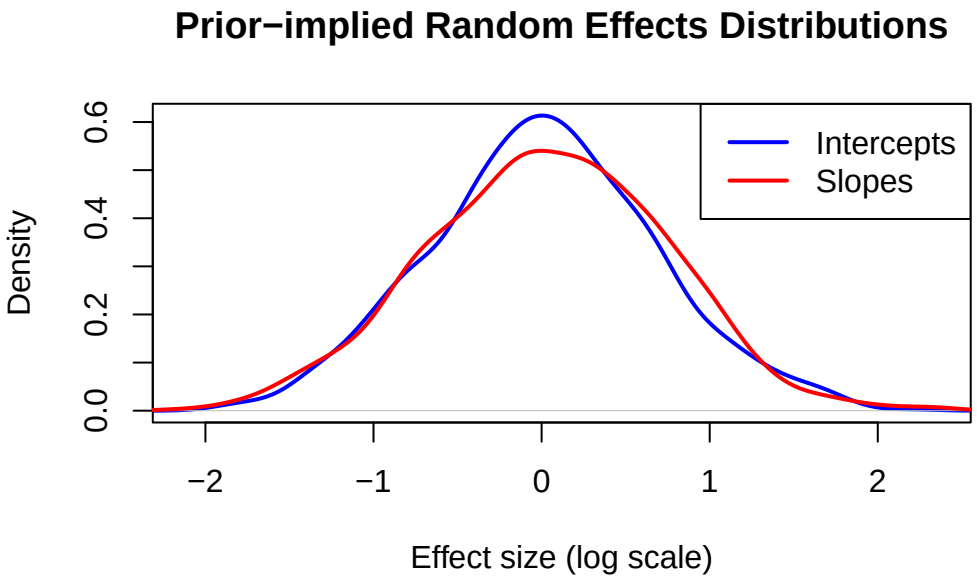

1.4 Interpretation

1.4.1 Good Signs (Prior is Reasonable)

- Prior generates log-RTs around 6 (400ms)
- 95% interval roughly 200-1100ms (plausible RT range)
- Condition effect typically $<$ 150ms difference
- Between-subject variation is moderate

1.4.2 Problems to Watch For

- Mean RT » 1000ms: intercept prior too high
- 95% interval 10ms-50s: priors too wide
- No variation between subjects: SD priors too small

1.5 Summary

Before fitting your model to actual data, always validate that your priors: 1. Generate reasonable predictions 2. Allow the data to inform the posterior 3. Respect domain knowledge constraints

Adjust priors as needed and rerun these checks.