Bayesian Cognitive Models

Paul Bürkner

Idea of Bayesian Ideal Observers

General research questions:

- Do humans process information in an ideal (Bayesian) way?
- In what aspects do the responses deviate from optimality?
- What does that tell us about cognitive processes?

Why is Bayesian ideal?

- Bayesian models are fully probabilistic
- If correctly specified they take all uncertainty into account
- Decision are thus based on all available information
- If one fails to follow the rules of probability, a Dutch book can be made against you so that in the long run you always loose

Example: Sensor Fusion

Suppose we have a multisensory observations x_a and x_v generated from the same source y:

$$x_a \sim N(y, \sigma_a) \quad \text{and} \quad x_v \sim N(y, \sigma_v)$$

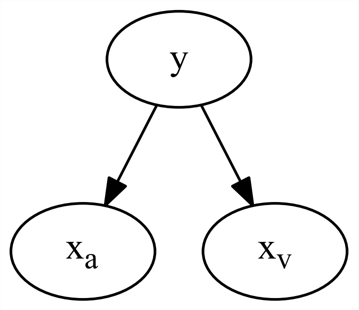

Bayesian Ideal Observers: Analysis Workflow

Generate the sensory input from the generative model

- As if we were the subject receiving senory input

Fit the sensory input using Bayesian cognitive model(s)

- As if we were the subject processing the sensory input

Extract the expected responses from the fitted model(s)

Compare the distributions of observed and expected responses

- A mismatch indicates a deviation from the expected behavior

Sensor Fusion: Analytic Posterior

Having observed x_a and x_v and knowing the corresponding standard deviations σ_a and σ_v, the posterior distribution of y is:

$$y \sim N(\mu_y, \sigma_y)$$

where

$$\mu_y = \frac{w_a x_a + w_v x_v}{w_a + w_v} \quad \text{and} \quad \sigma_y = \frac{1}{\sqrt{w_a + w_v}}$$

with weights $w_a = \sigma_a^{-2}$ and $w_v = \sigma_v^{-2}$

Under a quadratic loss function, μ_y is the optimal point estimate

The subject only reports his / her estimate $\hat{y}$ of y, which may or may not be an ideal combination μ_y of the two sources

We may also directly fit the Bayesian cognitive model in Stan without having to worry about analytic solutions

Requirements for Experiments

Subjects must have enough information to process the input in an ideal way

- Otherwise optimality is not achievable

Expected outcome distributions must be computable

- Analysing ideal observers using fully Bayesian inference is more flexible than trying to find the posterior analytically

Expected outcome distributions must vary across compared cognitive models

- Otherwise responses cannot provide any evidence

Ways to compare distributions

Comparing expected and observed distributions is of central interest

Potential ways to compare two distributions:

- Differences in means: $\mu_1 - \mu_2$
- Differences in standard deviations: $\sigma_1 - \sigma_2$
- Kullback-Leibler (KL) Divergence:

$$KL(p||q) = \int \log\left(\frac{p(x)}{q(x)}\right) p(x)\, d\, x$$

Example: Uncertainty in the Causal Structure

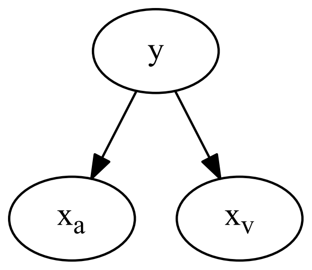

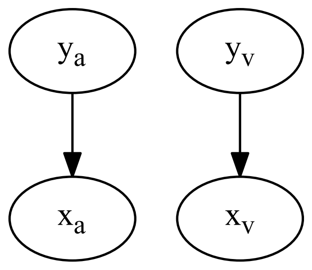

Each of the two causal models is true with a certain probability

Uncertainty in the Causal Structure: Statistical Model (1)

Let C indicate if a common ($C = 1$) or uncommen ($C = 0$) source is true:

$$\text{If } C = 1: \; x_a \sim N(y, \sigma_a) \quad \text{and} \quad x_v \sim N(y, \sigma_v)$$

$$\text{If } C = 0: \; x_a \sim N(y_a, \sigma_a) \quad \text{and} \quad x_v \sim N(y_v, \sigma_v)$$

The model parameters to be estimated are y, y_a, y_v, and C

This model cannot be fit in Stan as C is a discrete parameter

Uncertainty in the Causal Structure: Statistical Model (2)

We don't model C directly but instead $\pi_C = p(C = 1)$

The Bayesian cogntive model comes a *mixture* model:

$$x_a \sim \pi_C \, N(y, \sigma_a) + (1 - \pi_C) \, N(y_a, \sigma_a)$$

$$x_v \sim \pi_C \, N(y, \sigma_v) + (1 - \pi_C) \, N(y_v, \sigma_v)$$

The model parameters to be estimated are y, y_a, y_v, and π_C

This model can be fit in Stan as all parameters are continuous

Mixture models in Stan

Example: Normal mixture model with two components:

$$x \sim \pi\, N(\mu_1, \sigma_1) + (1 - \pi)\, N(\mu_2, \sigma_2)$$

For a single real value x, the Stan code is as follows:

```
model {
  real ps[2];
  ps[1] = log(pi) + normal_lpdf(x | mu1, sigma1);
  ps[2] = log(1 - pi) + normal_lpdf(x | mu2, sigma2);
  target += log_sum_exp(ps);
}
```

Example: Neural Correlates of Bayesian Belief

Paper by Hu et al. (2015): Predict stop signal probabilities

Create a generative model:

- $s_k = 1$ if trial k was a stop trial and $s_k = 0$ otherwise
- r_k is the expected probability of trial k being a stop trial
- This implies $s_k \sim \text{bernoulli}(r_k)$

In trial 1 we use a simple prior distribution:

$$r_1 \sim \text{beta2}(\mu, \phi)$$

where μ and ϕ represent prior mean and precision

Updating of Bayesian Belief

For the following trials, the prior will depend on the past:

$$r_k \sim \alpha\,\delta(r_{k-1}) + (1-\alpha)\,\text{beta2}(\mu, \phi)$$

The parameter α indicates the probability that the subject expects r_k to be the same as $k-1$

Given the mean $E(r_{k-1}|s)$ of the estimated posterior distribution of r_{k-1}, we can compute the mean $E(r_k)$ of the prior distribution of r_k as follows:

$$E(r_k) = \alpha E(r_{k-1}|s) + (1-\alpha)\mu$$

This can be used to predict behavioral and neuroimaging data

Updating of Bayesian Belief: Stan Code Snippet

```
// specify the likelihood
target += bernoulli_lpmf(s | r);
// specify the priors
target += beta2_lpdf(r[1] | mu, phi);
for (k in 2:N) {
  real ps[2];
  ps[1] = log(alpha) + beta2_lpdf(r[k] | r[k - 1], tau);
  ps[2] = log(1 - alpha) + beta2_lpdf(r[k] | mu, phi);
  target += log_sum_exp(ps);
}
```

References

- Vilares, I., & Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, 1224(1), 22-39.
- Hospedales, T., & Vijayakumar, S. (2009). Multisensory oddity detection as Bayesian inference. *PloS one*, 4(1), e4205.
- Hu, S., Ide, J. S., Zhang, S., & Chiang-shan, R. L. (2015). Anticipating conflict: neural correlates of a Bayesian belief and its motor consequence. *Neuroimage*, 119, 286-295