Bayesian regression models

Bruno Nicenboim / Shravan Vasishth

2020-03-17

A first linear model: Does attentional load affect pupil size?

Log-normal model: Does trial affect reaction times?

Logistic regression: Does set size affect free recall?

A first linear model: Does attentional load affect pupil size?

Data:
One participant's pupil size of the control experiment of Wahn et al. (2016) averaged by trial

Task:
A participant covertly tracked between zero and five objects among several randomly moving objects on a computer screen; multiple object tracking–MOT– (Pylyshyn and Storm 1988) task

Research question:
How does the number of moving objects being tracked (attentional load) affect pupil size?

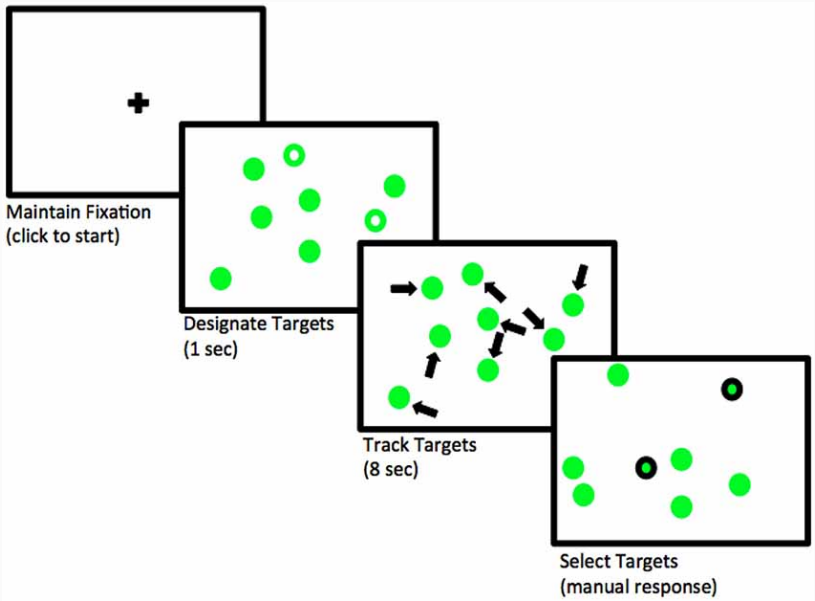

Figure 1: Flow of events in a trial where two objects needs to be tracked. Adapted from Blumberg, Peterson, and Parasuraman (2015); licensed under CC BY 4.0.

Assumptions:

1. There is some average pupil size represented by α.
2. The increase of attentional load has a linear relationship with pupil size, determined by β.
3. There is some noise in this process, that is, variability around the true pupil size i.e., a scale, σ.
4. The noise is normally distributed.

Formal model

Likelihood for each observation n:

$$p\_size_n \sim Normal(\alpha + c\_load_n \cdot \beta, \sigma) \tag{1}$$

where n indicates the observation number with $n = 1 \dots N$

How do we decide on priors?

Priors

- pupil sizes range between 2 and 5 millimeters,
- but the Eyelink-II eyetracker measures the pupils in arbitrary units (Hayes and Petrov 2016)
- we either need estimates from a previous analysis or look at some measures of pupil sizes

Pilot data:
Some measurements of the same participant with no attentional load for the first 100ms, each 10 ms, in `pupil_pilot.csv`:

```
df_pupil_pilot <- read_csv("./data/pupil_pilot.csv")
df_pupil_pilot$p_size %>% summary()
```

```
##    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##     852     856     862     861     866     868
```

Prior for α

$$\alpha \sim Normal(1000, 500) \tag{2}$$

Meaning:
We expect that the average pupil size for the average load in the experiment would be in a 95% central interval limited by approximately $1000 \pm 2 \cdot 500 = [20, 2000]$ units:

```
c(qnorm(.025, 1000, 500), qnorm(.975, 1000, 500))
```

```
## [1]   20 1980
```

Prior for σ

$$\sigma \sim Normal_{+}(0, 1000) \tag{3}$$

Meaning:
We expect that the standard deviation of the pupil sizes should be in the following 95% interval.

```
c(
  qtnorm(.025, 0, 1000, a = 0),
  qtnorm(.975, 70, 1000, a = 0)
)
```

```
## [1]   31 2290
```

Prior for β

$$\beta \sim Normal(0, 100) \tag{4}$$

Meaning:
We don't really know if the attentional load will increase or even decrease the pupil size, but we are only saying that one unit of load will potentially change the pupil size consistently with the following 95% interval:

```
c(qnorm(.025, 0, 100), qnorm(.975, 0, 100))
```

```
## [1] -196  196
```

Fitting the model

```
df_pupil_data <- read_csv("data/pupil.csv")
df_pupil_data <- df_pupil_data %>%
  mutate(c_load = load - mean(load))
df_pupil_data
```

```
## # A tibble: 41 x 4
##    trial  load p_size c_load
##    <dbl> <dbl>  <dbl>  <dbl>
## 1      1     2  1021. -0.439
## 2      2     1   951. -1.44
## 3      3     5  1064.  2.56
## 4      4     4   913.  1.56
## 5      5     0   603. -2.44
## # ... with 36 more rows
```

Specifying the model in brms

```
fit_pupil <- brm(p_size ~ 1 + c_load,
  data = df_pupil_data,
  family = gaussian(),
  prior = c(
    prior(normal(1000, 500), class = Intercept),
    prior(normal(0, 1000), class = sigma),
    prior(normal(0, 100), class = b, coef = c_load)
  )
)
```

```
plot(fit_pupil)
```

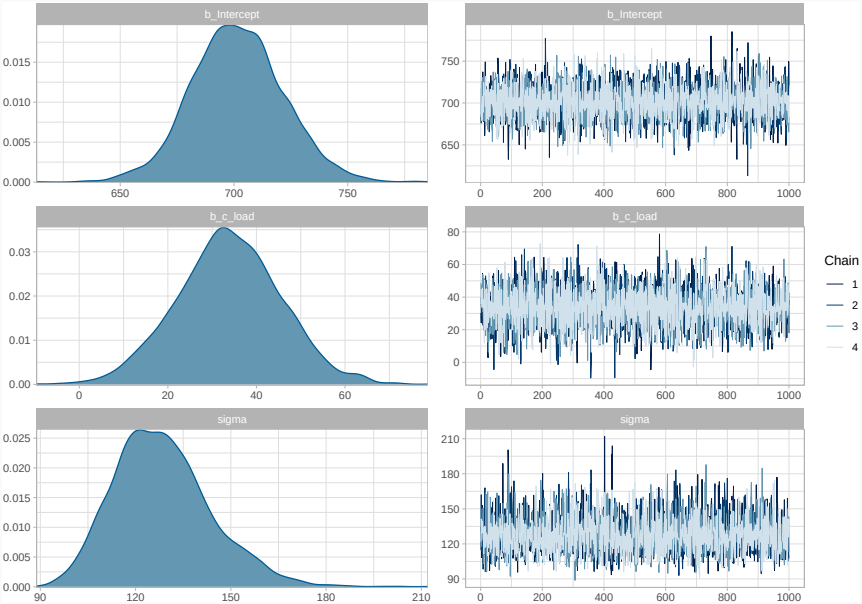

```
fit_pupil
```

```
##  Family: gaussian
##   Links: mu = identity; sigma = identity
## Formula: p_size ~ 1 + c_load
##    Data: df_pupil_data (Number of observations: 41)
## Samples: 4 chains, each with iter = 2000; warmup = 1000; thin = 1;
##          total post-warmup samples = 4000
##
## Population-Level Effects:
##           Estimate Est.Error l-95% CI u-95% CI Rhat Bulk_ESS Tail_ESS
## Intercept   701.53     20.10   662.27   742.58 1.00     3702     2751
## c_load       33.80     11.73    10.84    56.84 1.00     4126     2779
##
## Family Specific Parameters:
##       Estimate Est.Error l-95% CI u-95% CI Rhat Bulk_ESS Tail_ESS
## sigma   128.45     15.29   102.54   161.65 1.00     3066     2814
##
## Samples were drawn using sampling(NUTS). For each parameter, Bulk_ESS
## and Tail_ESS are effective sample size measures, and Rhat is the potential
## scale reduction factor on split chains (at convergence, Rhat = 1).
```

How to communicate the results?

Research question:
"What is the effect of attentional load on the participant's pupil size?"

We'll need to examine what happens with β (`c_load`):

How to communicate the results?

- The most likely values of β will be around the mean of the posterior, 33.8, and we can be 95% certain that the true value of β *given the model and the data* lies between 10.84 and 56.84.
- We see that as the attentional load increases, the pupil size of the participant becomes larger.

How likely it is that the pupil size increased rather than decreased?

```
mean(posterior_samples(fit_pupil)$b_c_load > 0)
```

```
## [1] 1
```

Take into account that this probability ignores the possibility of the participant not being affected at all by the manipulation, this is because $P(\beta = 0) = 0$.

Descriptive adequacy

```r
# we start from an array of 1000 samples by 41 observations
df_pupil_pred <- posterior_predict(fit_pupil, nsamples = 1000) %>%
  # we convert it to a list of length 1000, with 41 observations in each element:
  array_branch(margin = 1) %>%
  # We iterate over the elements (the predicted distributions)
  # and we convert them into a long data frame similar to the data,
  # but with an extra column `iter` indicating from which iteration
  # the sample is coming from.
  map_dfr(function(yrep_iter) {
    df_pupil_data %>%
      mutate(p_size = yrep_iter)
  }, .id = "iter") %>%
  mutate(iter = as.numeric(iter))
```

```
df_pupil_pred %>% filter(iter < 100) %>%
    ggplot(aes(p_size, group=iter)) +
  geom_line(alpha = .05, stat="density", color = "blue") +
    geom_density(data=df_pupil_data, aes(p_size), inherit.aes = FALSE, size =1)+
    geom_point(data=df_pupil_data, aes(x=p_size, y = -0.001), alpha =.5, inherit.aes = FALSE) +
    coord_cartesian(ylim=c(-0.002, .01))+ facet_grid(load ~ .)
```

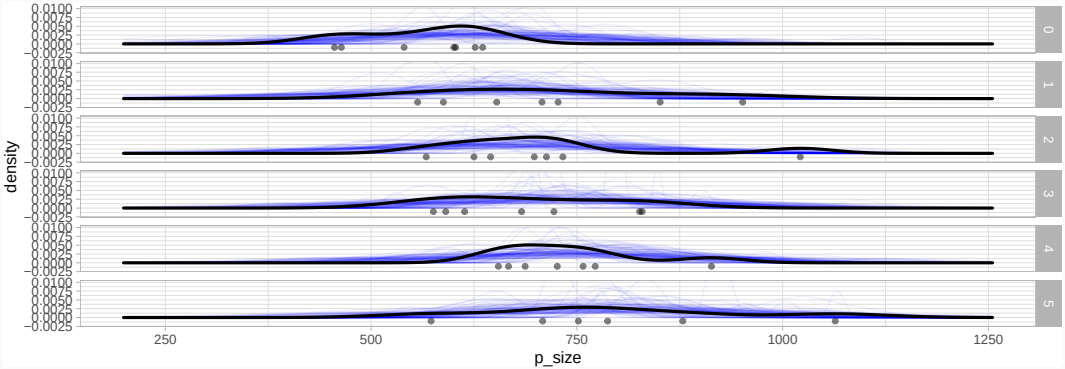

Figure 2: The plot shows 100 predicted distributions in blue density plots, the distribution of pupil size data in black density plots, and the observed pupil sizes in black dots for the five levels of attentional load.

Distribution of statistics

```
# predicted means:
df_pupil_pred_summary <- df_pupil_pred %>%
  group_by(iter, load) %>%
  summarize(av_p_size = mean(p_size))
# observed means:
(df_pupil_summary <- df_pupil_data %>%
  group_by(load) %>%
  summarize(av_p_size = mean(p_size)))
```

```
## # A tibble: 6 x 2
##    load av_p_size
##   <dbl>     <dbl>
## 1     0      561.
## 2     1      719.
## 3     2      715.
## 4     3      691.
## 5     4      740.
## # ... with 1 more row
```

```
ggplot(df_pupil_pred_summary, aes(av_p_size)) +
  geom_histogram(alpha = .5) +
  geom_vline(aes(xintercept = av_p_size), data = df_pupil_summary) +
  facet_grid(load ~ .)
```

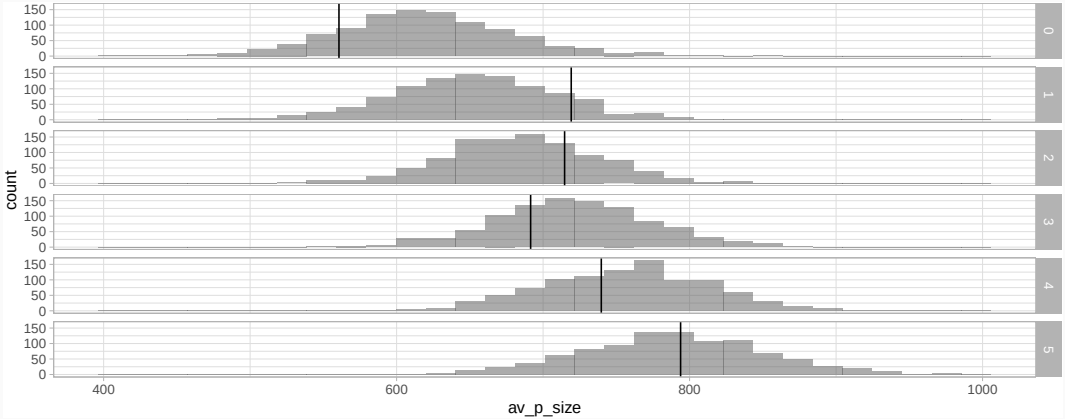

Figure 3: Distribution of posterior predicted means in gray and observed pupil size means in black lines by load.

- the observed means for no load and for a load of two are falling in the tails of the distributions.
- the data might be indicating that the relevant difference is between (i) no load, (ii) a load between two and three, and then (iii) a load of four, and (iv) of five.
- but beware of overinterpreting noise.

Value of posterior predictive distributions

- If we look hard enough, we'll find failures of descriptive adequacy.[1]
- Posterior predictive accuracy can be used to generate new hypotheses and to compare different models.

[1] all models are wrong

Exercises

4.6.1.1 Our priors for this experiment were quite arbitrary. How do the prior predictive distributions look like? Do they make sense?

4.6.1.2 Is our posterior distribution sensitive to the priors that we selected? Perform a sensitivity analysis to find out whether the posterior is affected by our choice of prior for the σ.

4.6.1.3 Our dataset includes also a column that indicates the trial number. Could it be that trial has also an effect on the pupil size? As in `lm`, we indicate another main effect with a + sign. How would you communicate the new results?

Log-normal model: Does trial affect reaction times?

We revisit the small experiment, where a participant repeatedly pressed the space bar as fast as possible, without paying attention to the stimuli.

New research question:
Does the participant tend to speedup (practice effect) or slowdown (fatigue effect)?

Formal model

Likelihood:

$$rt_n \sim LogNormal(\alpha + c\_trial_n \cdot \beta, \sigma) \tag{5}$$

Priors

$$\begin{aligned} \alpha &\sim Normal(6, 1.5) \\ \sigma &\sim Normal_+(0, 1) \\ \beta &\sim ... \end{aligned} \tag{6}$$

Prior for β

$$\beta \sim Normal(0, 1) \tag{7}$$

We edit our `normal_predictive_distribution_fast` from section and make it log-normal and dependent on trial:

```
lognormal_model_pred <- function(alpha_samples,
                                 beta_samples,
                                 sigma_samples,
                                 N_obs) {
    # pmap extends map2 (and map) for a list of lists:
    pmap_dfr(list(alpha_samples, beta_samples, sigma_samples),
             function(alpha, beta, sigma) {
                 tibble(
                     trialn = seq_len(N_obs),
                     # we center trial:
                     c_trial = trialn - mean(trialn),
                     # we change the likelihood:
                     # Notice rlnorm and the use of alpha and beta
                     rt_pred = rlnorm(N_obs, alpha + c_trial * beta, sigma))
             }, .id = "iter") %>%
    # .id is always a string and needs to be converted to a number
        mutate(iter = as.numeric(iter))}
```

This is our first attempt for a prior predictive distribution:

```
N_obs <- 361
N <- 800
alpha_samples <- rnorm(N, 6, 1.5)
sigma_samples <- rtnorm(N, 0, 1, a = 0)
beta_samples <- rnorm(N, 0, 1)
prior_pred <- lognormal_model_pred(
  alpha_samples = alpha_samples,
  beta_samples = beta_samples,
  sigma_samples = sigma_samples,
  N_obs = N_obs
)
```

```r
(median_effect <-
  prior_pred %>%
  group_by(iter) %>%
  mutate(diff = rt_pred - lag(rt_pred)) %>%
  summarize(
    median_rt = median(diff, na.rm = TRUE)
  ))
```

```
## # A tibble: 800 x 2
##    iter median_rt
##   <dbl>     <dbl>
## 1     1  1.40e- 5
## 2     2  2.12e-15
## 3     3 -6.36e- 1
## 4     4 -5.69e+ 0
## 5     5 -1.81e-16
## # ... with 795 more rows
```

```r
median_effect %>%
  ggplot(aes(median_rt)) +
  geom_histogram()
```

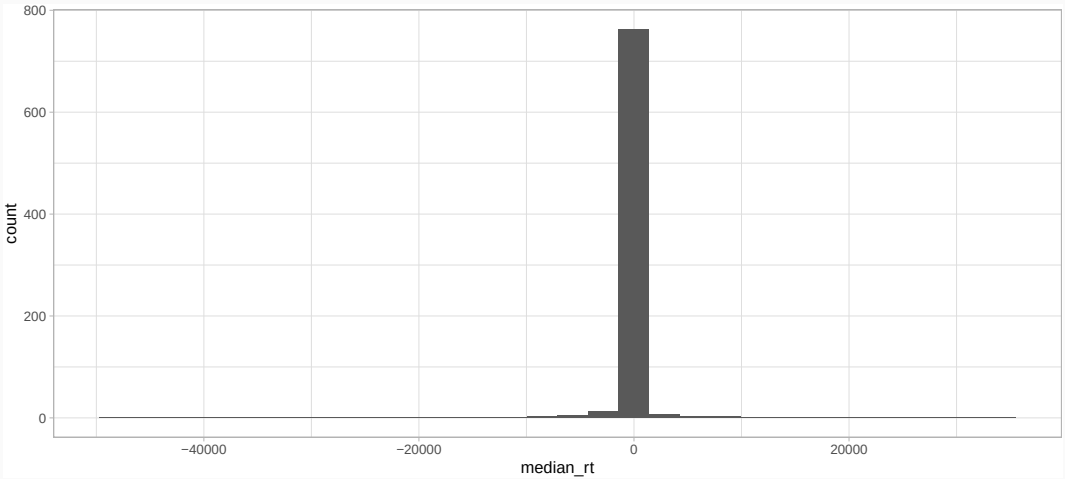

Figure 4: Prior predictive distribution of the median effect of the log-normal model with $\beta \sim Normal(0, 1)$.

Another prior for β

$$\beta \sim Normal(0, .01) \tag{8}$$

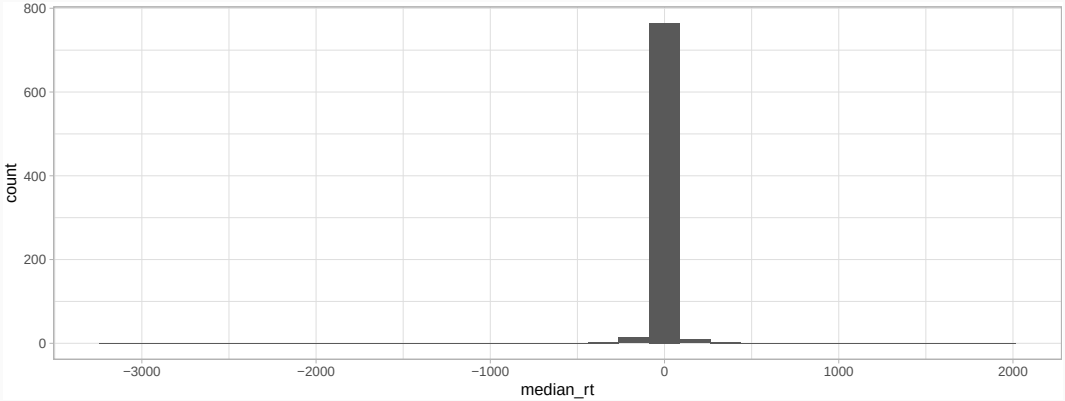

Figure 5: Prior predictive distribution of the median effect of the log-normal model with $\beta \sim Normal(0, .01)$.

Prior selection

Prior selection might look daunting and a lot of work. However...

- priors can be informed by the estimates from previous experiments;
- this work is usually done only the first time we encounter an experimental paradigm;
- we will generally use very similar (or identical priors) for analyses dealing with the same type of task;
- when in doubt, do a sensitivity analysis.

Fitting the model

```
df_noreading_data <- read_csv("./data/button_press.csv")
df_noreading_data <- df_noreading_data %>%
  mutate(c_trial = trialn - mean(trialn))
fit_press_trial <- brm(rt ~ 1 + c_trial,
  data = df_noreading_data,
  family = lognormal(),
  prior = c(
    prior(normal(6, 1.5), class = Intercept),
    prior(normal(0, 1), class = sigma),
    prior(normal(0, .01), class = b, coef = c_trial)
  )
)
```

```r
posterior_summary(fit_press_trial)[, c("Estimate", "Q2.5", "Q97.5")]
```

```
##                Estimate       Q2.5       Q97.5
## b_Intercept     5.11844     5.1058     5.13064
## b_c_trial       0.00052     0.0004     0.00065
## sigma           0.12330     0.1147     0.13295
## lp__        -1603.65601 -1606.7664 -1602.27805
```

```
plot(fit_press_trial)
```

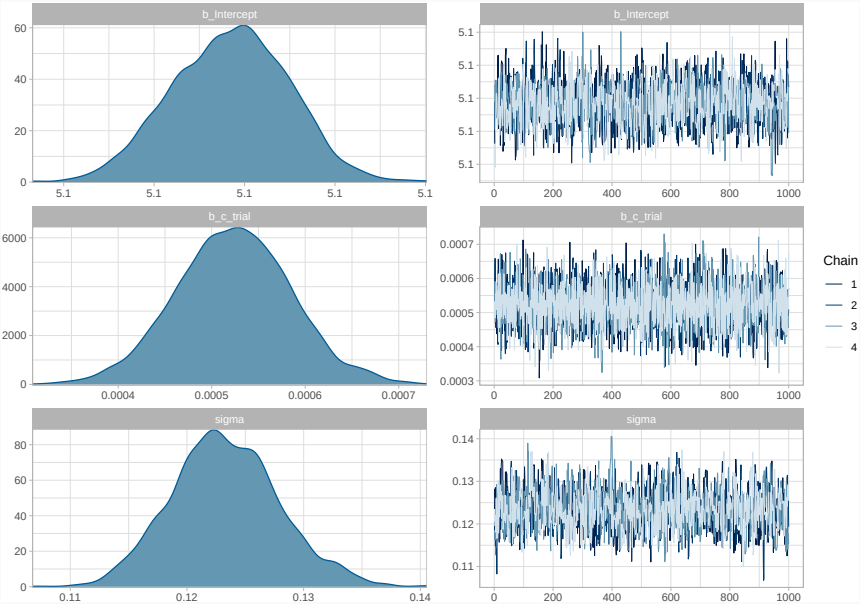

How to communicate the results?

We focus on the effect of trial:

- $\widehat{\beta} = 0.00052$, 95% CrI = $[0.0004, 0.00065]$.
- But in most cases, the effect is easier to interpret in milliseconds.

We calculate an estimate if we consider the difference between reaction times in a trial at the middle of the experiment (when the centered trial number is zero) and the previous one (when the centered trial number is minus one).

```
alpha_samples <- posterior_samples(fit_press_trial)$b_Intercept
beta_samples <- posterior_samples(fit_press_trial)$b_c_trial
effect_middle_ms <- exp(alpha_samples) -
  exp(alpha_samples - 1 * beta_samples)
## ms effect in the middle of the expt (mean trial vs. mean trial - 1 )
c(mean = mean(effect_middle_ms), quantile(effect_middle_ms, c(.025, .975)))
```

```
##  mean  2.5%   98%
## 0.087 0.067 0.109
```

Alternatively we consider the difference between the second trial and the first one:

```
first_trial <- min(df_noreading_data$c_trial)
second_trial <- min(df_noreading_data$c_trial) + 1
effect_beginning_ms <- exp(alpha_samples + second_trial * beta_samples) -
  exp(alpha_samples + first_trial * beta_samples)
## ms effect from first to second trial:
c(mean = mean(effect_beginning_ms), quantile(effect_beginning_ms, c(.025, .975)))
```

```
##  mean  2.5%   98%
## 0.080 0.062 0.097
```

There is a slowdown in both cases.

Reporting results

We can

- present the posterior mean and the 95% credible interval;
- assess if the observed estimates are consistent with the prediction from our theory;
- assess the practical relevance of the effect for the research question; (only after 100 button presses we see a slowdown of 9 ms on average ($0.09 \cdot 100$), with a 95% credible interval ranging from 6.7 to 10.86);
- establish the presence or absence of an effect (Bayes factor)

Exercises

4.6.2.1 Estimate the slowdown in milliseconds between the last two times the subject pressed the space bar in the experiment.

4.6.2.2 How would you change your model (keeping the log-normal likelihood) so that it includes centered log-transformed trial numbers or square-root-transformed trial numbers (instead of centered trial numbers)? Does the effect in milliseconds change?

Logistic regression: Does set size affect free recall?

We'll look at the capacity limit of working memory to illustrate one special case of GLMs, logistic regression.

Subset of the data of Oberauer (2019):

Data One participants recall success (1 success, 0 failure)

Task: word lists of varying lengths (2, 4, 6, and 8 elements** were presented, and the participant was asked to recall a word given its position on the list

Research question: How does the number of items to be held in working memory affects recall accuracy?

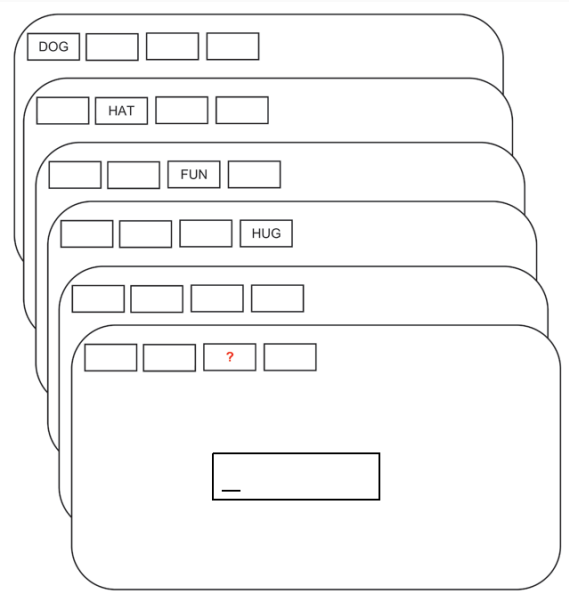

Figure 6: Flow of events in a trial with memory set size 4 and free recall. Adapted from Oberauer (2019); licensed under CC BY 4.0.

```r
df_recall_data <- read_csv("./data/PairsRSS1_all.csv") %>%
  # We ignore the type of incorrect responses (the focus of the paper)
  mutate(correct = if_else(response_category == 1, 1, 0)) %>%
  # and we only use the data from the free recall task:
  # (when there was no list of possible responses)
  filter(response_size_list + response_size_new_words == 0) %>%
  # We select one subject
  filter(subject == 10) %>%
  mutate(c_set_size = set_size - mean(set_size)) %>%
  select(subject, set_size, c_set_size, correct, trial)
```

```r
# Set sizes in the dataset:
df_recall_data$set_size %>%
  unique()
```

```
## [1] 4 8 2 6
```

```r
# Trials by set size
df_recall_data %>%
  group_by(set_size) %>%
  count()
```

```
## # A tibble: 4 x 2
## # Groups:   set_size [4]
##   set_size     n
##      <dbl> <int>
## 1        2    23
## 2        4    23
## 3        6    23
## 4        8    23
```

```
df_recall_data
```

```
## # A tibble: 92 x 5
##   subject set_size c_set_size correct trial
##     <dbl>    <dbl>      <dbl>   <dbl> <dbl>
## 1      10        4         -1       1     1
## 2      10        8          3       0     4
## 3      10        2         -3       1     9
## 4      10        6          1       1    23
## 5      10        4         -1       1     5
## # ... with 87 more rows
```

The likelihood for the logistic regression model

Recall that the Bernoulli likelihood generates a 0 or 1 response with a particular probability θ (here N = 10 trials with 50% chances of getting a one):

```
# We use as.numeric to get zeros and ones rather than FALSE and TRUE
rbernoulli(n = 10, p = 0.5) %>% as.numeric()
```

```
##  [1] 1 0 1 0 1 1 0 0 0 0
```

Formal model

The likelihood for each observation n:

$$correct_n \sim Bernoulli(\theta_n) \tag{9}$$

- θ_n is bounded to be between 0 and 1

How do we fit a regression model?

The generalized linear modeling framework

- A **link function** $g(\cdot)$ connects the linear model (real numbers ranging from $(-\infty, +\infty)$) to the quantity to be estimated (here, the probabilities θ_n in $[0, 1]$).
- A (common) link function in this case is the **logit link**:

$$\eta_n = g(\theta_n) = \log\left(\frac{\theta_n}{1-\theta_n}\right) \tag{10}$$

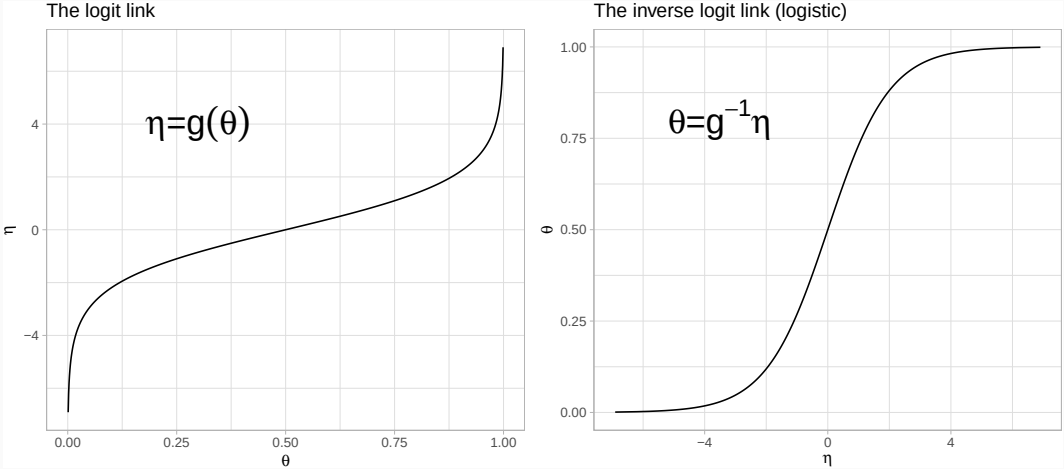

Figure 7: The logit and inverse logit (logistic) function.

Formal model

The likelihood for each observation n:

$$\eta_n = \log\left(\frac{\theta_n}{1-\theta_n}\right) = \alpha + \beta \cdot c\_set\_size \tag{11}$$

$$\theta_n = g^{-1}(\eta_n) = \log\left(\frac{\exp(\eta_n)}{1+\exp(\eta_n)}\right) \tag{12}$$

$$correct_n \sim Bernoulli(\theta_n) \tag{13}$$

Priors for logistic regression

- α represents the *log-odds* of correctly recalling one word in a random position for the average set size of five (because we centered the predictor and since $5 = \frac{2+4+6+8}{4}$). (It's telling us how difficult the task is. Let's assume (a 50/50 chance) with a great deal of uncertainty:

We use `qlogis(p)` for the inverse logit or logistic function:

```
qlogis(.5)
```

```
## [1] 0
```

Prior for α

$$\alpha \sim Normal(0, 4) \tag{14}$$

```
samples_logodds <- tibble(alpha = rnorm(100000, 0, 4))
samples_prob <- tibble(p = plogis(rnorm(100000, 0, 4)))
ggplot(samples_logodds, aes(alpha)) + geom_density()
ggplot(samples_prob, aes(p)) + geom_density()
```

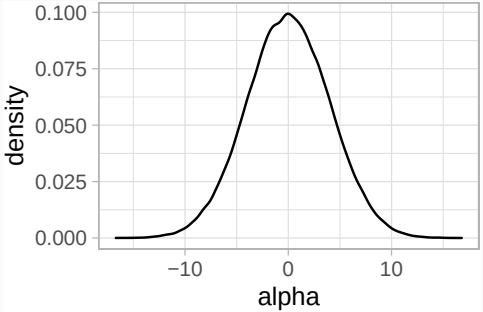

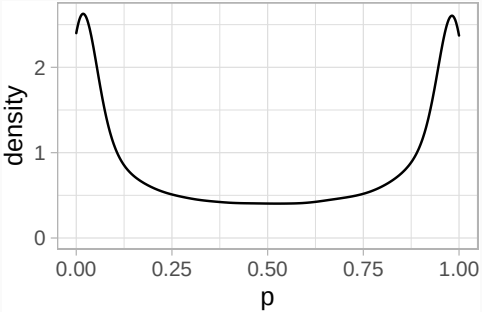

Prior for α

We try with:

$$\alpha \sim Normal(0, 1.5) \tag{15}$$

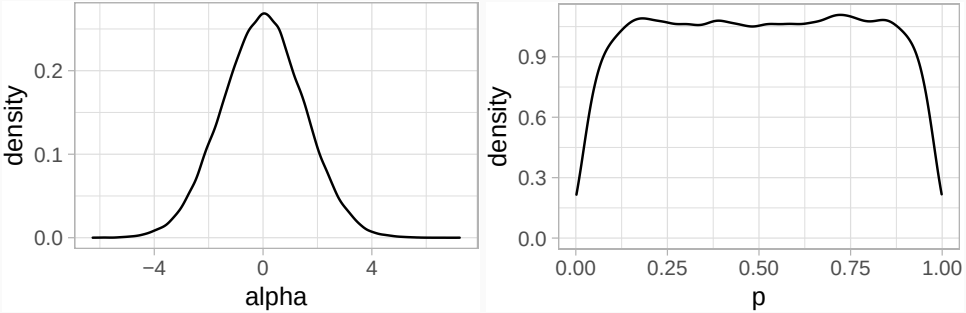

Figure 9: Prior for $\alpha \sim Normal(0, 1.5)$ in log-odds and in probability space.

Prior for β

- β represents the effect in log-odds of increasing the set size.

(a) $\beta \sim Normal(0, 1)$
(b) $\beta \sim Normal(0, .5)$
(c) $\beta \sim Normal(0, .1)$
(d) $\beta \sim Normal(0, .01)$
(e) $\beta \sim Normal(0, .001)$

Edited version of the earlier `normal_predictive_distribution_fast`:

```
logistic_model_pred <- function(alpha_samples,
                                beta_samples,
                                set_size,
                                N_obs) {
    map2_dfr(alpha_samples, beta_samples,
             function(alpha, beta) {
                 tibble(set_size = set_size,
                                      # we center size:
                        c_set_size = set_size - mean(set_size),
                                      # change the likelihood:
                                      # Notice the use of a link function for alpha and beta
                        theta = plogis(alpha + c_set_size * beta),
                        correct_pred = rbernoulli(N_obs,  p = theta))
             }, .id = "iter") %>%
    # .id is always a string and needs to be converted to a number
        mutate(iter = as.numeric(iter))
}
```

Let's assume 800 observations with 200 observation of each set size:

```
N_obs <- 800
set_size <- rep(c(2, 4, 6, 8), 200)
```

We iterate over the four possible standard deviations of β:

```
alpha_samples <- rnorm(1000, 0, 1.5)
sds_beta <- c(1, 0.5, 0.1,0.01, 0.001)
prior_pred <- map_dfr(sds_beta, function(sd) {
    beta_samples <- rnorm(1000, 0, sd)
    logistic_model_pred(alpha_samples = alpha_samples,
                        beta_samples = beta_samples,
                        set_size = set_size,
                        N_obs = N_obs) %>%
        mutate(prior_beta_sd = sd)
})
```

And we calculate the accuracy for each one of the priors we want to examine, for each iteration, and for each set size.

```
(mean_accuracy <- prior_pred %>%
   group_by(prior_beta_sd, iter, set_size) %>%
   summarize(accuracy = mean(correct_pred)) %>%
   mutate(prior = paste0("Normal(0, ",prior_beta_sd,")")))
```

```
## # A tibble: 20,000 x 5
## # Groups:   prior_beta_sd, iter [5,000]
##   prior_beta_sd  iter set_size accuracy prior
##           <dbl> <dbl>    <dbl>    <dbl> <chr>
## 1         0.001     1        2    0.255 Normal(0, 0.001)
## 2         0.001     1        4    0.27  Normal(0, 0.001)
## 3         0.001     1        6    0.24  Normal(0, 0.001)
## 4         0.001     1        8    0.255 Normal(0, 0.001)
## 5         0.001     2        2    0.435 Normal(0, 0.001)
## # ... with 2e+04 more rows
```

```
mean_accuracy %>%
  ggplot(aes(accuracy)) +
  geom_histogram() +
  facet_grid(set_size ~ prior)
```

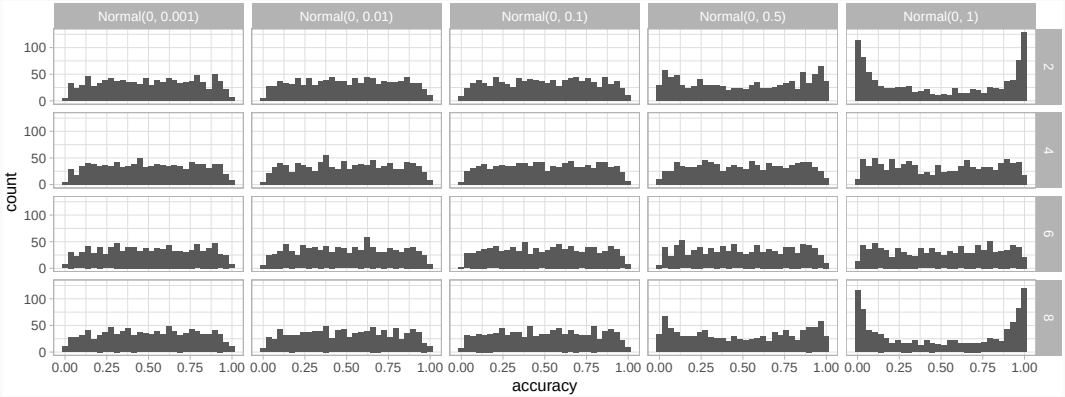

Prior predicted differences in accuracy

```
(diff_accuracy <- mean_accuracy %>%
    arrange(set_size) %>%
    group_by(iter, prior_beta_sd) %>%
    mutate(diffaccuracy = accuracy - lag(accuracy) ) %>%
    mutate(diffsize = paste(set_size,"-",  lag(set_size))) %>%
    filter(set_size >2))
```

```
## # A tibble: 15,000 x 7
## # Groups:   iter, prior_beta_sd [5,000]
##   prior_beta_sd  iter set_size accuracy prior            diffaccuracy
##           <dbl> <dbl>    <dbl>    <dbl> <chr>                   <dbl>
## 1         0.001     1        4    0.27  Normal(0, 0.001)       0.015
## 2         0.001     2        4    0.42  Normal(0, 0.001)      -0.015
## 3         0.001     3        4    0.32  Normal(0, 0.001)      -0.0400
## 4         0.001     4        4    0.825 Normal(0, 0.001)       0.0650
## 5         0.001     5        4    0.94  Normal(0, 0.001)      -0.01
##   diffsize
##   <chr>
## 1 4 - 2
## 2 4 - 2
```

```
diff_accuracy %>%
  ggplot(aes(diffaccuracy)) +
  geom_histogram() +
  facet_grid(diffsize ~ prior)
```

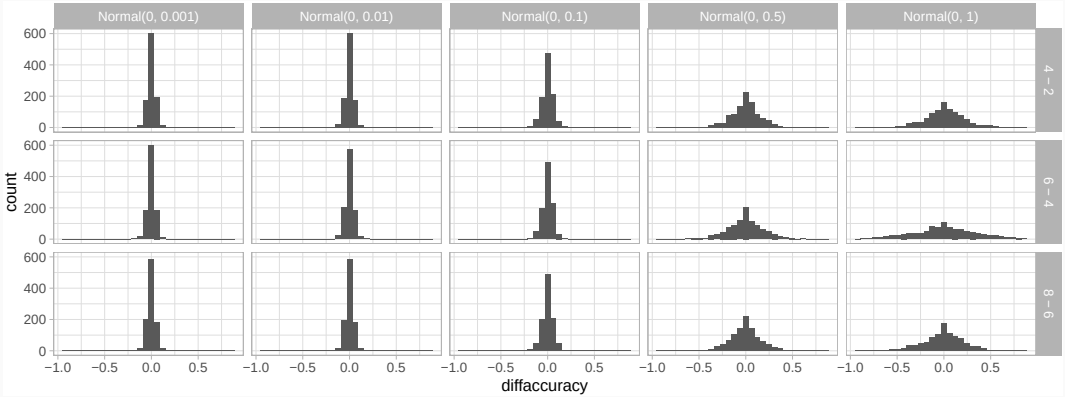

$$\alpha \sim Normal(0, 1.5)$$
$$\beta \sim Normal(0, 0.1) \tag{16}$$

The brms model

```
fit_recall <- brm(correct ~ 1 + c_set_size,
  data = df_recall_data,
  family = bernoulli(link = logit),
  prior = c(
    prior(normal(0, 1.5), class = Intercept),
    prior(normal(0, .1), class = b, coef = c_set_size)
  )
)
```

```r
posterior_summary(fit_recall, pars = c("b_Intercept", "b_c_set_size"))
```

```
##                Estimate Est.Error  Q2.5  Q97.5
## b_Intercept        1.92     0.298  1.36  2.524
## b_c_set_size      -0.18     0.081 -0.34 -0.028
```

```
plot(fit_recall)
```

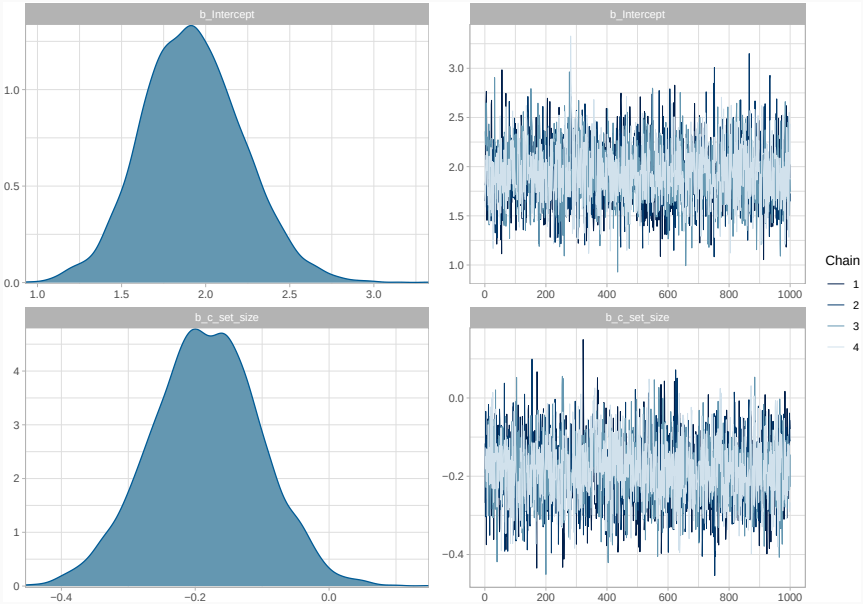

How to communicate the results?

If we want to talk about the effect estimated by the model in log-odds space, we summarize the posterior of β in the following way:

- $\widehat{\beta} = -0.18$, 95% CrI = $[-0.34, -0.03]$.

Effect in proportions

Average accuracy for the task:

```
alpha_samples <- posterior_samples(fit_recall)$b_Intercept
av_accuracy <- plogis(alpha_samples)
c(mean = mean(av_accuracy), quantile(av_accuracy, c(.025, .975)))
```

```
## mean 2.5%  98%
## 0.87 0.80 0.93
```

Effect in proportions

Effect of our manipulation

```
beta_samples <- posterior_samples(fit_recall)$b_c_set_size
effect_av_set_size <- plogis(alpha_samples) - plogis(alpha_samples - beta_samples)
c(mean = mean(effect_av_set_size), quantile(effect_av_set_size, c(.025, .975)))
```

```
##   mean   2.5%    98%
## -0.019 -0.037 -0.003
```

Notice the interpretation here:

if we increase the set size from the average set size minus one to the average set size (5), we get a reduction in the accuracy of recall of -0.02*, 95% CrI =* $[-0.04, 0]$.

Effect in proportions

Recall that the average set size, 5, was not presented to the subject!

Decrease in accuracy from a set size of 2 to 4:

```
set4 <- 4 - mean(df_recall_data$set_size)
set2 <- 2 - mean(df_recall_data$set_size)
effect_4m2 <- plogis(alpha_samples + set4 * beta_samples) -
  plogis(alpha_samples + set2 * beta_samples)
c(mean = mean(effect_4m2), quantile(effect_4m2, c(.025, .975)))
```

```
##    mean    2.5%     98%
## -0.0295 -0.0540 -0.0057
```

We see that increasing the set size does have a detrimental effect in recall, as we suspected.

Descriptive adequacy

We could also make predictions for other conditions not presented in the actual experiment, such as set sizes that weren't tested:

- We extend our dataset adding rows with set sizes of 3, 5, and 7: we add 23 trials of each new set size
- Notice is that **we need to center our predictor based on the original mean set size**

```r
df_recall_data_ext <- df_recall_data %>%
  bind_rows(tibble(
    set_size = rep(c(3, 5, 7), 23),
    c_set_size = set_size - mean(df_recall_data$set_size)
  ))
df_recall_pred_ext <- posterior_predict(fit_recall,
  newdata = df_recall_data_ext,
  nsamples = 1000
) %>%
  array_branch(margin = 1) %>%
  map_dfr(function(yrep_iter) {
    df_recall_data_ext %>%
      mutate(correct = yrep_iter)
  }, .id = "iter") %>%
  mutate(iter = as.numeric(iter))
```

```r
(df_recall_pred_ext_summary <- df_recall_pred_ext %>%
  group_by(iter, set_size) %>%
  summarize(accuracy = mean(correct)))
```

```
## # A tibble: 7,000 x 3
## # Groups:   iter [1,000]
##    iter set_size accuracy
##   <dbl>    <dbl>    <dbl>
## 1     1        2    0.826
## 2     1        3    0.913
## 3     1        4    0.957
## 4     1        5    1
## 5     1        6    0.826
## # ... with 6,995 more rows
```

```r
# observed means:
(df_recall_summary <- df_recall_data %>%
  group_by(set_size) %>%
  summarize(accuracy = mean(correct)))
```

```
## # A tibble: 4 x 2
##   set_size accuracy
##      <dbl>    <dbl>
## 1        2    1
## 2        4    0.957
## 3        6    0.913
## 4        8    0.609
```

```
ggplot(df_recall_pred_ext_summary, aes(accuracy)) +
  geom_histogram(alpha = .5) +
  geom_vline(aes(xintercept = accuracy), data = df_recall_summary) +
  facet_grid(set_size ~ .)
```

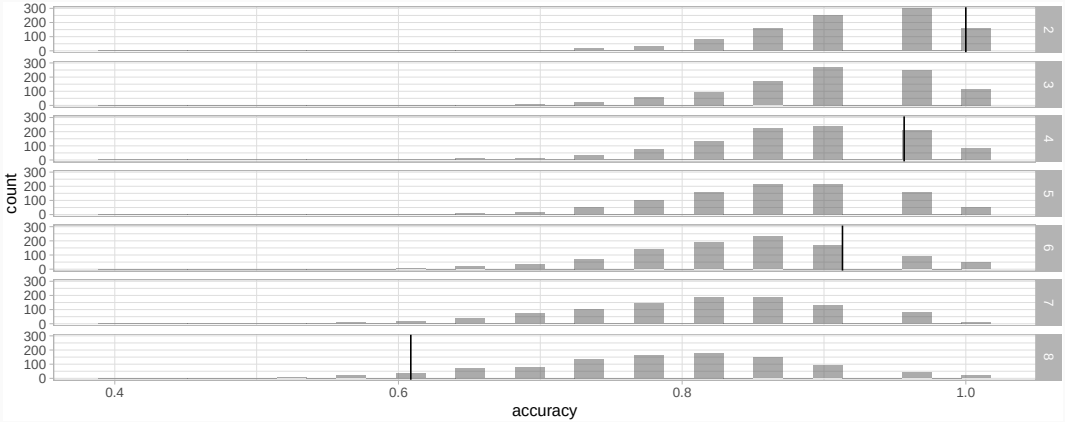

References

Blumberg, Eric J., Matthew S. Peterson, and Raja Parasuraman. 2015. “Enhancing Multiple Object Tracking Performance with Noninvasive Brain Stimulation: A Causal Role for the Anterior Intraparietal Sulcus.” *Frontiers in Systems Neuroscience* 9: 3. https://doi.org/10.3389/fnsys.2015.00003.

Hayes, Taylor R., and Alexander A. Petrov. 2016. “Mapping and Correcting the Influence of Gaze Position on Pupil Size Measurements.” *Behavior Research Methods* 48 (2): 510–27. https://doi.org/10.3758/s13428-015-0588-x.

Oberauer, Klaus. 2019. “Working Memory Capacity Limits Memory for Bindings.” *Journal of Cognition* 2 (1): 40. https://doi.org/10.5334/joc.86.

Pylyshyn, Zenon W., and Ron W. Storm. 1988. “Tracking Multiple