

load libraries and set bayesplot theme

make sure you have these packages installed

```
library("rstan")
library("ggplot2")
library("bayesplot")

theme_set(bayesplot::theme_default())
```

Load and examine data

```
# Loads vector of counts 'y'
source("count-data.R")


N <- length(y)
print(N)
```

```
## [1] 500
```

```
print(y)
```


```
## [1] 0 3 5 0 4 7 4 2 3 6 7 0 0 3 7 5 5 0 4 0 4 4 6 3 7 5 3 0 0 2 0 1 0 1 5
## [36] 4 4 2 3 6 4 5 0 7 7 4 4 4 0 6 1 5 6 5 6 7 3 6 2 3 0 2 0 6 6 0 3 4 4 5
## [71] 5 0 5 7 5 5 6 4 2 3 4 6 4 6 6 4 0 6 5 5 7 0 1 6 7 0 5 0 0 5 6 5 1 0 7
## [106] 1 2 6 5 4 0 4 0 4 4 6 3 0 0 3 3 4 2 5 3 4 3 2 5 2 4 4 0 2 7 5 7 5 5 7
## [141] 7 0 4 6 0 4 6 7 4 0 4 1 5 0 3 5 7 6 0 5 5 6 7 6 7 3 4 3 7 7 2 5 4 5 5
## [176] 0 6 2 4 5 4 0 0 5 5 7 7 0 3 0 3 3 6 1 4 2 0 4 7 5 5 0 3 7 0 6 6 4 1 6
## [211] 7 6 0 3 6 4 7 0 5 5 4 0 0 2 4 6 0 5 0 2 7 2 7 5 4 6 2 4 0 4 0 0 3 5 4
## [246] 3 5 5 7 7 0 6 4 5 1 5 3 5 5 5 0 2 7 6 2 3 2 5 4 7 6 7 3 3 4 4 6 4 6 7
## [281] 1 5 6 3 3 6 3 4 0 7 0 3 6 5 0 0 0 5 4 4 0 4 7 5 5 3 3 0 0 5 4 0 7 6 0
## [316] 6 2 0 6 1 0 4 0 4 3 0 4 5 5 7 6 6 5 4 7 0 6 4 7 7 5 0 1 4 7 6 4 5 4 7
## [351] 2 5 2 6 3 2 7 4 3 4 6 6 6 6 7 1 0 0 7 7 4 2 4 5 5 7 4 1 7 6 5 6 5 4 0
## [386] 0 7 0 0 5 6 6 3 6 0 0 0 4 4 3 0 7 5 4 2 7 0 4 0 0 2 4 5 0 4 2 5 2 0 6
## [421] 6 3 6 0 2 5 0 0 0 6 0 0 6 5 4 6 4 5 5 4 0 3 4 3 3 5 3 4 5 7 0 0 1 4 6
## [456] 3 5 7 6 6 5 0 5 4 0 0 2 6 0 6 0 4 5 6 3 4 2 3 4 0 5 0 0 0 0 3 4 7 6 7
## [491] 7 3 4 4 7 4 5 2 5 6
```

```
qplot(y)
```



Compare our data to draws from Poisson with same mean

```
x <- rpois(N, lambda = mean(y))  
qplot(x)
```




```


plotdata <- data.frame(
  value = c(y, x),
  variable = rep(c("Our data", "Poisson data"), each = N)
)

# Frequency polygons
ggplot(plotdata, aes(x = value, color = variable)) +
  geom_freqpoly(binwidth = 0.5) +
  scale_x_continuous(name = "", breaks = 0:max(x,y)) +
  scale_color_manual(name = "", values = c("gray30", "purple"))

```



```
# Side-by-side bar plots  
ggplot(plotdata, aes(x = value, fill = variable)) +  
  geom_bar(position = "dodge") +  
  scale_x_continuous(name = "", breaks = 0:max(x,y)) +  
  scale_fill_manual(name = "", values = c("gray30", "purple"))
```



Fit basic Poisson model

Even though we already suspect it won't be a good model for this data, it's still a good idea to start by fitting the simplest Poisson model. From there we can then identify in which ways the model is inadequate.

```
# create list
stan_list <- list()
stan_list$Y <- y
stan_list$N <- N

# stan code
stan_code <- c("
data {
  int N;           // Number of observations
  int<lower=0> Y[N]; // Count data (integer array)
}

parameters {
  real<lower=0> lambda; // Poisson rate parameter
}

model {
  // Prior
  lambda ~ exponential(0.1);

  // Likelihood
  Y ~ poisson(lambda);

```




```

## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/include/stan/math/prim/math
##   size_t fft_next_good_size(size_t N) {
##     ~
## In file included from file34f394dd571.cpp:8:
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##   typedef typename Array::index_range index_range;
##     ~
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##   typedef typename Array::index index;
##     ~
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##   typedef typename Array::index_range index_range;
##     ~
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##   typedef typename Array::index index;
##     ~
## 8 warnings generated.


```

```
plot(fit, plotfun="trace", pars='lambda')
```



Look at posterior distribution of lambda

```
color_scheme_set("brightblue") # check out bayesplot::color_scheme_set
lambda_draws <- as.matrix(fit, pars = "lambda")
mcmc_areas(lambda_draws, prob = 0.95) # color 95% credible interval
```



Compare posterior of lambda to the mean of the data

```
means <- c("Posterior mean" = mean(lambda_draws), "Data mean" = mean(y))
print(means, digits = 3)
```

```
## Posterior mean      Data mean
##           3.66           3.66
```

The model gets the mean right, but, as we'll see next, the model is quite bad at predicting the outcome.


Graphical posterior predictive checks

Extract `y_rep` draws from the fitted model object

```
y_rep <- as.matrix(fit, pars = "y_rep")
```



Compare histogram of y to histograms of several y_{reps}

```
ppc_hist(y, y_rep[1:8, ], binwidth = 1)
```



Compare density estimate of y to density estimates of a bunch of y_{reps}

```
ppc_dens_overlay(y, y_rep[1:50, ])
```



clear environment and load normal data


```
rm(list = ls())  
mydata <- read.csv("practice_data.csv")
```

```
X <- mydata$x  
N <- length(X)  
print(mean(X))
```


```
## [1] 48.91559
```

```
qplot(X)
```

```
## `stat_bin()` using `bins = 30`. Pick better value with `binwidth`.
```



```
Y <- rnorm(N, mean=mean(X), sd=sd(X))  
qplot(Y)
```




```

plotdata <- data.frame(
  value = c(X, Y),
  variable = rep(c("Our data", "Normal data"), each = N)
)

ggplot(plotdata, aes(x=value)) +
  geom_histogram(data=subset(plotdata, variable == 'Our data'),
    fill = "red", alpha = 0.25) +
  geom_histogram(data=subset(plotdata, variable == 'Normal data'),
    fill = "blue", alpha = 0.25)

```



```

# create list
stan_list <- list()
stan_list$X <- X
stan_list$N <- N

```

```

# stan code
stan_code <- c("
data {
  // write code here
}

parameters {
  // write code here
}

model {
  // write code here
}
")

```



```

## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/StanHeaders/inc
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## In file included from /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##     typedef typename Array::index_range index_range;
##     ^
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##     typedef typename Array::index index;
##     ^
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##     typedef typename Array::index_range index_range;
##     ^
## /Library/Frameworks/R.framework/Versions/3.4/Resources/library/BH/include/boost/multi_array/concept_
##     typedef typename Array::index index;
##     ^
## 8 warnings generated.

## Warning in .local(object, ...): some chains had errors; consider specifying
## chains = 1 to debug

## [[1]]
## Stan model '4cfe54917e686dc5089576fcaed71af8' does not contain samples.
##
## [[2]]
## Stan model '4cfe54917e686dc5089576fcaed71af8' does not contain samples.
##
## [[3]]
## Stan model '4cfe54917e686dc5089576fcaed71af8' does not contain samples.
##
## [[4]]
## Stan model '4cfe54917e686dc5089576fcaed71af8' does not contain samples.

```