

## ggplot tutorial

Lets make up some data (don't worry about this step)

```
xval <- seq.int(1,40)
yval <- xval + rnorm(40,0,5)
```


```
data <- data.frame(xval,yval)
```

```
head(data) #only shows first 6 rows
```

```
##   xval      yval
## 1    1  3.9868733
## 2    2  1.7314658
## 3    3  1.1363760
## 4    4 -0.2467743
## 5    5  3.2290180
## 6    6  9.3313731
```


Here's a scatter plot:

```
ggplot(data=data, aes(x=xval, y=yval)) +  
  geom_point()
```




Here's a line plot:

```
ggplot(data=data, aes(x=xval, y=yval)) +  
  geom_line()
```




We could make both lines and points by writing:

```
ggplot(data=data, aes(x=xval, y=yval)) +  
  geom_point() +  
  geom_line()
```



Here's a bar plot:


```
ggplot(data=data, aes(x=xval, y=yval)) +
  geom_col()
```



### Adding a regression line

You can also put regression lines on your plots by running the `stat_smooth()` function. This function takes a “method” argument for the kind of regression. Below, I used `method="lm"` to specify that it's a linear model.

```
ggplot(data=data, aes(x=xval, y=yval)) +
  geom_point() +
  stat_smooth(method="lm")
```



### Visualizing groups

Now suppose the data came in *groups* and we want to visualize each group separately. For example, let's suppose we have multiple treatment groups with a single outcome measure. For simplicity, we'll just assign each point in our dataframe randomly to one of two groups.

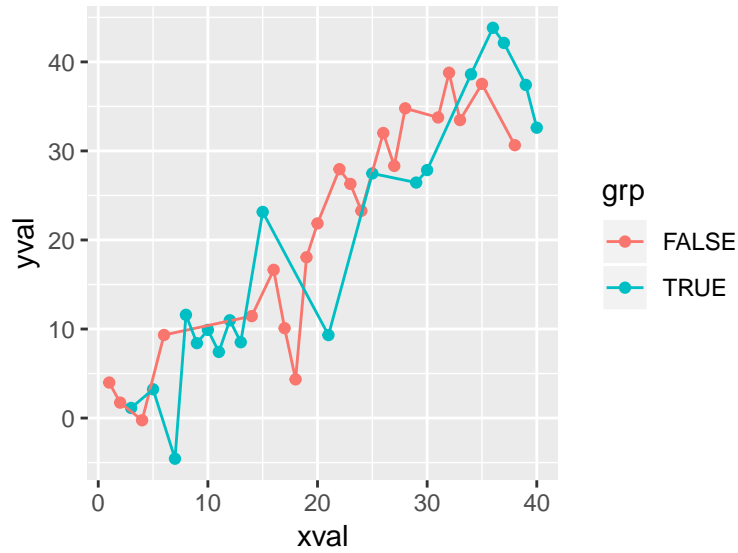
```
grp <- rnorm(nrow(data),0,1) > 0
data$grp <- grp
```

```
head(data)
```

```
##   xval      yval  grp
## 1    1  3.9868733 FALSE
## 2    2  1.7314658 FALSE
## 3    3  1.1363760  TRUE
## 4    4 -0.2467743 FALSE
## 5    5  3.2290180  TRUE
## 6    6  9.3313731 FALSE
```


There are a couple basic ways to visualize this. The first is to put all the data on the same plot, but separate the groups visually, such as by making them a different color.

```
ggplot(data=data, aes(x=xval,y=yval,color=grp)) +
  geom_point() +
  geom_line()
```



Another common way to visualize each group separately is to make subplots or “facets” for each group. To do this, we use the function `facet_wrap`.

```
ggplot(data=data, aes(x=xval,y=yval)) +
  geom_point() +
  geom_line() +
  facet_wrap(~grp)
```



### A couple other tips

- To change the name of the x- and y-axis titles, use `xlab("x-axis name")` and `ylab("y-axis name")`.
- To change the color of groups use `scale_color_manual(values=c("red","green",...))`.
- To change stylistic aspects of the plot that are not related to the data, such as the text size, the color of the background, etc. . . , you can modify the “theme”. If you wanted to increase the size of the x-axis text, you could write `theme(axis.text.x = element_text(size=__))`.
- Check out the ggplot cheatsheet, which has a list of functions and when to use them: <https://rstudio.com/wp-content/uploads/2015/03/ggplot2-cheatsheet.pdf>


## Your turn

Now it's your turn. We're going to look at the built-in R dataset showing UC Berkeley admissions of men and women from 1973, broken down by department. It comes from a lawsuit alleging a gender bias in admissions. Here you're going to figure out: were UC Berkeley admissions biased?

```
data(UCBAdmissions)
admit <- data.frame(UCBAdmissions)
```

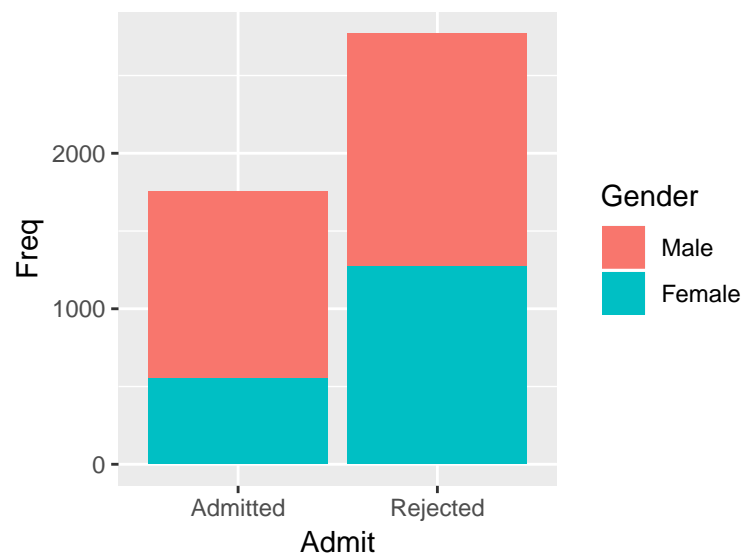
First, let's just look at how many men and women applied in total.

```
ggplot(data=admit, aes(x=Gender, y=Freq)) +
  geom_col()
```



Next, let's look at the proportion of admissions by gender.


```
ggplot(data=admit, aes(x=Admit, y=Freq, fill=Gender)) +
  geom_col()
```



Ugh, the default color scheme makes 'admission' red and 'rejected' green, which is counter-intuitive. So before we go any further let's fix that (hint: use `scale_fill_manual()`). You can find color names at

<http://www.stat.columbia.edu/~tzheng/files/Rcolor.pdf>.

```
ggplot(data=admit, aes(x=Gender, y=Freq, fill=Admit)) +  
  geom_col() +  
  scale_fill_manual(values=c("green3", "red3"))
```

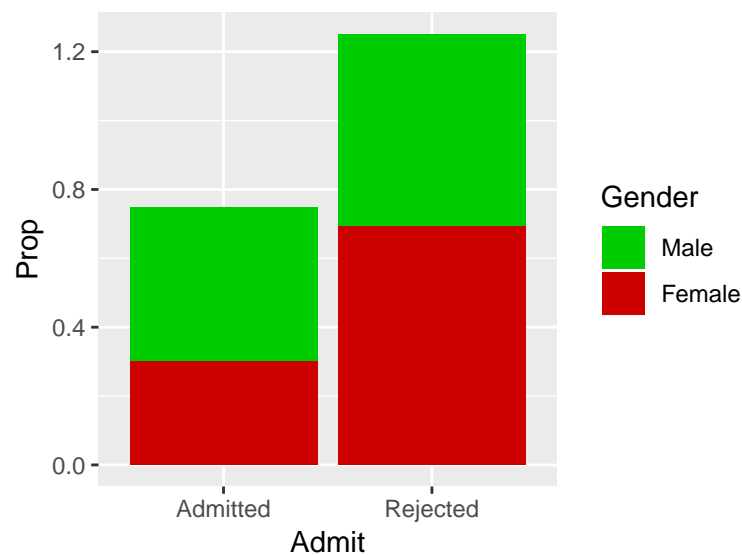


Looks pretty biased, but it's kind of hard to compare, because the total number of men and women who applied is different — so the bars are different heights. So let's turn these into proportions.

```
admit <- admit %>%  
  group_by(Gender) %>%  
  mutate(Prop = Freq/sum(Freq))  
View(admit)
```

Now plot the proportions:

```
ggplot(data=admit, aes(x=Admit, y=Prop, fill=Gender)) +  
  geom_col() +  
  scale_fill_manual(values=c("green3", "red3"))
```



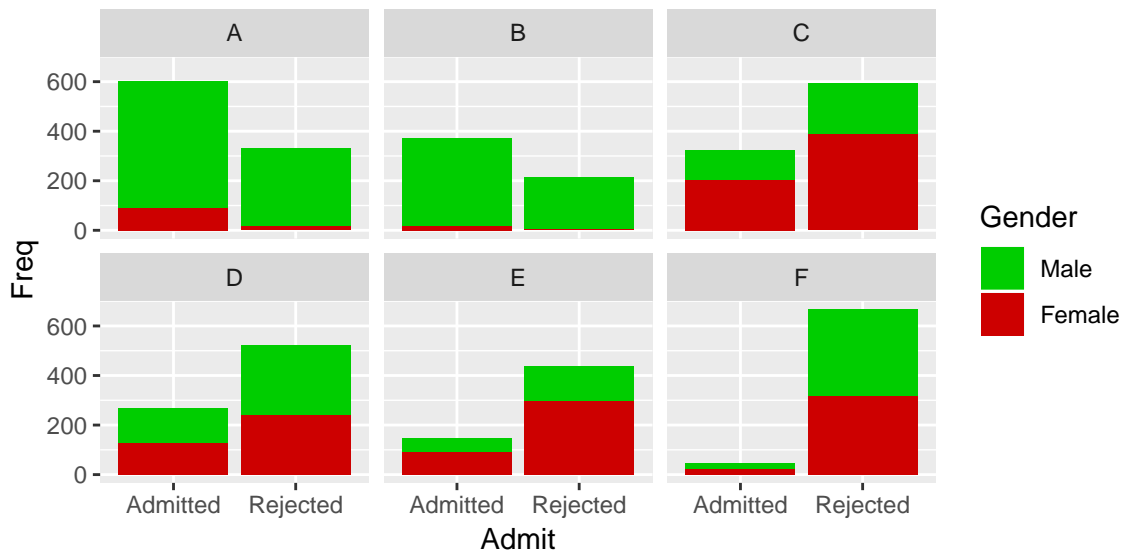
Should we conclude that the admissions are biased? The answer is no – at least, not yet. The reason is that

there's a missing piece of information here, which is the admission rates by department. Let's start off with a faceted plot of all the frequencies of admission by department (colored by gender).

Here we use the function `facet_wrap` to make subplots ("facets") for each department.

```
ggplot(data=admit, aes(x=Admit, y=Freq, fill=Gender)) +
  geom_col() +

  scale_fill_manual(values=c("green3", "red3")) +
  facet_wrap(~Dept)
```



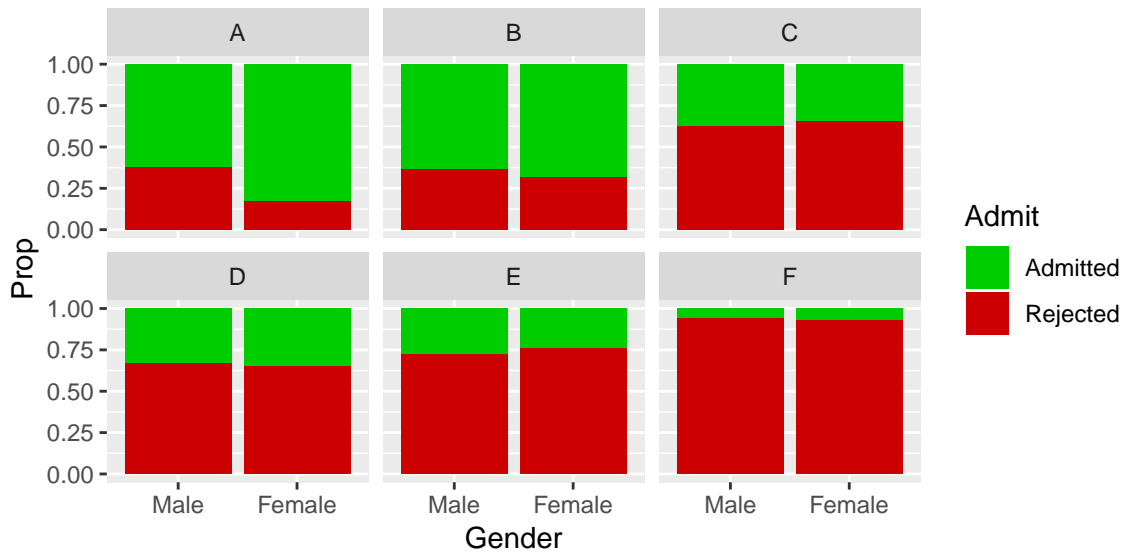
Again, because the total number of applicants of men and women varies by department, this is hard to interpret. So let's again turn each subplot into proportion instead of an overall frequency.

```
admit <- admit %>%
  group_by(Dept, Gender) %>%
  mutate(Prop=Freq/sum(Freq))
```

Now plot the proportions:

```
ggplot(data=admit, aes(x=Gender, y=Prop, fill=Admit)) +
  geom_bar(stat='identity') +
  scale_fill_manual(values=c("green3", "red3")) +
  facet_wrap(~Dept)
```





The data suddenly seem to tell an entirely different story. It seems that when you actually look at the proportion of men and women admitted to different departments, the rates of acceptance are essentially identical! This can be explained if women tend to apply to departments with lower acceptance rates. Let's demonstrate this a little more cleanly, by finding the proportion of women who applied to each department and plotting how selective those departments were overall.

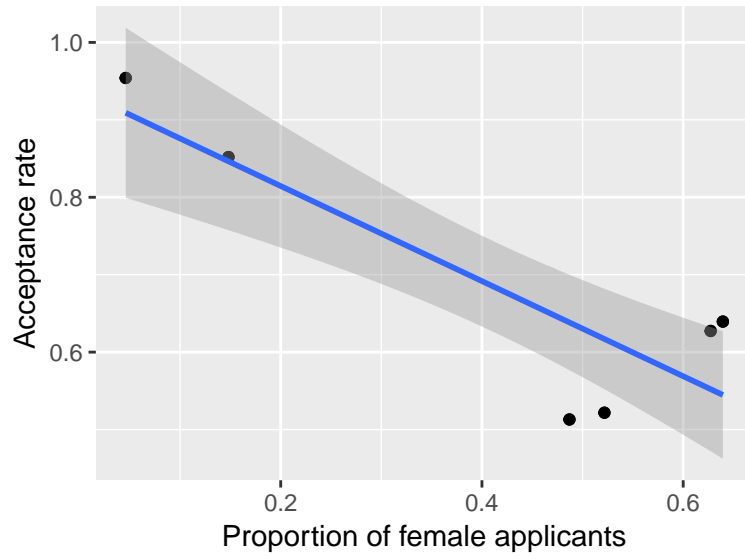
```
admit <- admit %>%
  group_by(Dept, Admit) %>%
  filter (Admit=="Admitted") %>%

  group_by(Dept, Gender) %>%
  mutate(Pct_Women = sum(Freq) * (Gender=="Female")) %>%

  group_by(Dept) %>%
  mutate(Pct_Women = max(Pct_Women)/sum(Freq)) %>%
  mutate(Prop=max(Freq)/sum(Freq))
```

So did women just apply to more selective departments?

```
ggplot(data=admit, aes(x=Pct_Women, y=Prop)) +
  geom_point() +
  xlab("Proportion of female applicants") +
  ylab("Acceptance rate") +
  stat_smooth(method="lm")
```



That was the legal conclusion. See: [https://en.wikipedia.org/wiki/Simpson%27s\\_\\_paradox](https://en.wikipedia.org/wiki/Simpson%27s__paradox)