We learned in class how to do hypothesis tests for comparing the means of two populations from single samples of size n_1 and n_2 of these populations. This involves several steps. First we make sure that the data are approximately normally distributed (or at least do not show strong skewness or contain extreme outliers). Then we compute the sample means and standard deviations for both samples, $\bar{x}_1, \bar{x}_2, s_1, s_2$. The *test statistic* in this case is t value of the difference in the sample means:

$$t = \frac{(\bar{x}_1 - \bar{x}_2) - (\mu_1 - \mu_2)}{\sqrt{\frac{s_1^2}{n_1} \frac{s_2^2}{n_2}}}$$

Our null hypothesis will be that there is NO DIFFERENCE in the populations means. In other words we hypothesize that $\mu_1 - \mu_2 = 0$. We then look up the propbability that we would observe the computed vlue of t if the null hpothesis were true. If this probability is very small, then we conclude that the null hypothesis is NOT true and our data supports the conlcusion there IS a difference in the population means.

A bit of work, but we know how to do it. The trickiest part is looking things up on the (incomplete) tables in the back of the book. Luckily, computers are great at storing tables and doing boring calculations. R even has a built-in `t.test` command to do exactly what we want!.

Imagine a data set from a hospital for post-surgery recovery times in days. Seven patients were randomly divided into a control group of 11 that received standard care, and a treatment group of 10 that received a new kind of care. The data is:

```
>  treatment <- c(15,10,13,7,9,8,21,9,14,8)
>  placebo  <-  c(15,14,12,8,14,7,16,10,15,12,12)
```

We could use R to help do our analysis. First check the data for strangeness, outliers, skewness. A side-by-side box should show that everything is quasi-normal.

two groups:

```
>  boxplot(treatment, placebo)
```

Next we need means and standard deviations and sample sizes.

```
>  x1 <- mean(treatment)
>  x2 <- mean(placebo)
>  s1 <- sd(treatment)
>  s2 <- sd(placebo)
```

```
> n1 <- length(treatment)
> n2 <- length(placebo)
```

We can easily compute the difference in the means and Confidence Interval for this difference.

```
>  SE <- sqrt( s1^2/n1 + s2^2/n2)
>  t <- (x1 - x2)/SE
>  t
```

```
[1] -0.5334211
```

Not a very large t score. Maybe there reallyis no difference in the means? Now all we need to do is look up this t value on our table to get the probability. How many degrees of freedom? Our (VERY) conservative estimate would be $df = 9$, since the smallest sample size is $n_1 = 10$. We can get R to do the looking up easily.

```
>  p <- pt(t,9)
>  p
```

```
[1] 0.303331
```

OK. What did we learn? Given our data, assuming that there is NO DIFFERENCE in the means, there is a 30% chance of seeing the observed difference. Pretty high chance. Therefore, we cannot reject the null hypothesis. Our data does not support the contention that there IS a difference in the means.

Whew, a lot of work - even with a computer. But R can do all this for us in one line. Try this

```
> t.test(treatment,placebo,alt="less") #Do a t.test on the data, alt hypoth: control < placebo
```

```
        Welch Two Sample t-test

data:  treatment and placebo
t = -0.5334, df = 15.587, p-value = 0.3006
alternative hypothesis: true difference in means is less than 0
95 percent confidence interval:
     -Inf 1.988347
sample estimates:
mean of x mean of y
 11.40000  12.27273
```

Wow! All that work reduced to one line. The result is only slightly different, but this is because R was less conservative in its estimate of the number of degrees of freedom. The answer, that $p \sim 0.3$ still stands. There is no evidence of significant differences in the means.

Question 0.1. Do ALL of this again using the data in Table 7.5 of the text book. In other words, work the problem out step by step, and then have R do it all for you in one line. What do you find? Report the results so that someone NOT taking M214 can understand the result.

Question 0.2. Let's look at some real data. There is a data set in R for a study of the effects of vitamin C on the growth of inscisors (front teeth) in guinea pigs. (Yes, people do study the strangest things.) We want to use our knowledge of statistics to understand this data set. In particular we want to assess, with hypothesis tests, whether the data supports the following contentions: (a) vitamin C helps guinea pig teeth grow and (b) giving vitamin C in Orange Juice is better for guinea pig tooth growth than giving vitamin as ascorbic acid (pure vitamin C). Type `help(ToothGrowth)` in R to get a full description of the data set.

Load the data:

```
>   data(ToothGrowth)
>   attach(ToothGrowth)
>   names(ToothGrowth)

[1] "len"  "supp" "dose"
```

We can start to answer the questions by looking at the data. Box-plots on the two factors: Dosage amount and Delivery Means will help us begin to answer the questions.

```
>   boxplot(len~dose)
>   boxplot(len~supp)
```

From the box-plots, two things are apparent. First, the data looks relatively normal, there are no signs of severe skewness or the presence of outliers in any. Second, it looks pretty clear that increases in the dosage of vitamin C given to the guinea pigs leads to increases in tooth length. Whether or not there is a statistical difference between the method of delivery (OJ versus straight vitamin C) is open to debate.

Lets be statistical and use R to do standard two sample t-tests of the differences in the means across both dosage and delivery method. First, lets split the data up.

```
>  big_dose = len[dose==2.0]
>  med_dose = len[dose==1.0]
>  low_dose = len[dose==0.5]
```

```
> vc = len[supp=='VC']
> oj = len[supp=='OJ']
```

To Do: Use a hypothesis test to investigate the differences in mean tooth length for big dose and med dose. Do this by hand first, using R to do the calculations for you. Once that is done, use R to do the whole thing for you in one line. Are your results and R's results consistent? Explain both to someone who is not taking Mth214.

To Do: Repeat the above, this time comparing the means for OJ and VC. Do it two ways: by hand - use R to do the calculations and then using R to do everything. Again, do you and R agree? Is there evidence for differences due to the delivery mechanism of Vitamin C? Explain.