1. This problem asks us to do a one-way classification analysis to compare the population mean sugar beet yields for 6 different types of nitrogen sources. This analysis will be carried out as follows:

- We will first perform an overall F test to see if there any differences in the population means.
- If we conclude at least one of the population means is different, we will then perform a follow-up Tukey analysis to see which population means are different from the others. We can also make a recommendation as to which nitrogen source(s) maximize(s) the population mean yield.

Let's begin by looking at the data from the experiment; I used side-by-side boxplots to do this:

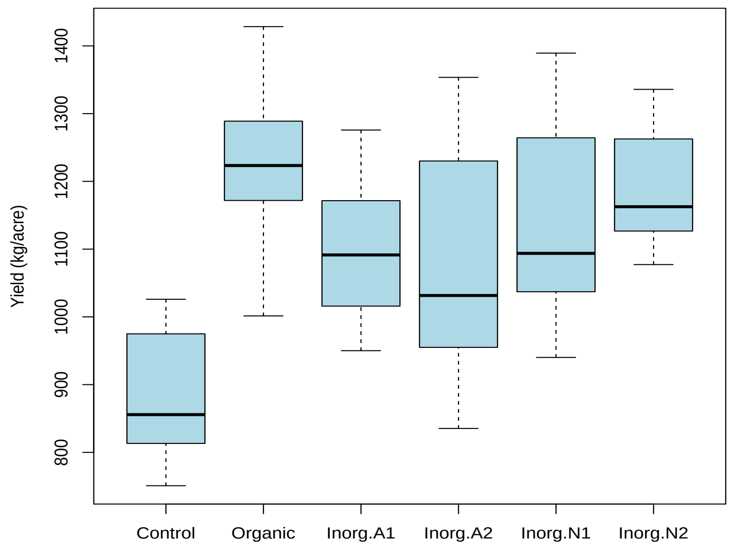

Here are some first impressions:

- I would not expect to see a plot like this if the 6 population mean sugar beet yields were really equal. I suspect the overall F test will reject H_0 decisively.
- Remember that one of the assumptions in a one-way classification analysis is that the population variances are equal. The variation I see in the boxplots is somewhat similar among the groups. Of course, we only have 10 observations per treatment group, which is not that much information.

I used R to produce the ANOVA table for a one-way classification analysis:

```
> fit = lm(Yield ~ Source)
> anova(fit)
Analysis of Variance Table

          Df Sum Sq Mean Sq F value  Pr(>F)
Source     5 738684  147737    9.19 2.3e-06
Residuals 54 868151   16077
```

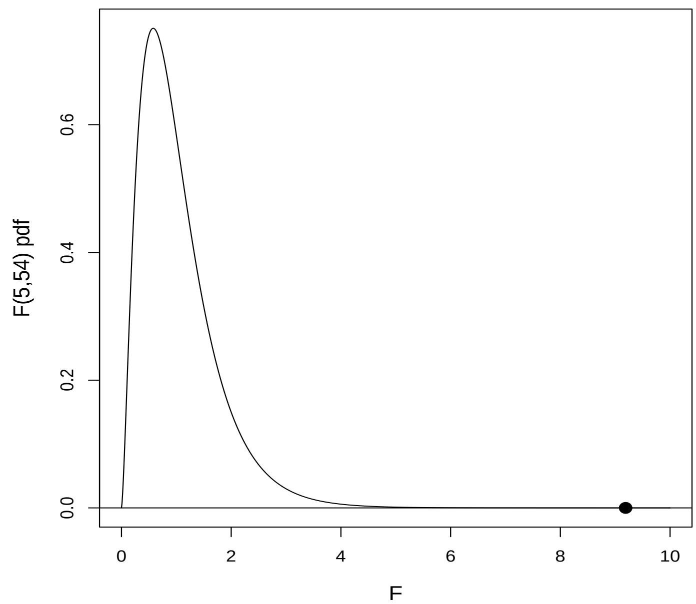

Overall *F* test: The F statistic in the ANOVA table ($F \approx 9.19$) is used to test

$$H_0 : \mu_1 = \mu_2 = \mu_3 = \mu_4 = \mu_5 = \mu_6$$
$$\text{versus}$$
$$H_1 : \text{the population means are not all equal.}$$

The null hypothesis says the sugar beet population mean yields are the same for all 6 nitrogen sources. The alternative hypothesis says that at least one population mean is different than the others.

If H_0 really was true, then we would expect to see F around 1. That is not what we see here. This value of F is more consistent with H_1. In the figure above, we plot $F = 9.19$ on the $F(5, 54)$ pdf, which describes the sampling distribution of F when H_0 is true. Not surprisingly, it is far out in the right tail with

$$\text{p-value} = 2.3 \times 10^{-6} = 0.0000023.$$

We have very strong evidence against H_0 and conclude that at least one population mean sugar beet yield is different than the others.

Follow-up analyses: The next step is to perform a follow-up analysis where we examine the

$$\binom{6}{2} = 15$$

pairwise comparisons. That is, we will write confidence intervals for the pairwise population mean differences

$$\Delta_{ii'} = \mu_i - \mu_{i'},$$

for $1 \leq i < i' \leq 6$, and determine which intervals include/exclude "0." Pairwise intervals which exclude "0" refer to population mean pairs which are judged to be different. We use Tukey's method to do this. I assume a familywise confidence level of 95%, meaning that we are 95% confident that all 15 pairwise intervals contain their true population mean difference.

Here is the Tukey analysis from R:

```
> TukeyHSD(aov(fit),conf.level=0.95)
  Tukey multiple comparisons of means
    95% family-wise confidence level

                    diff    lwr upr p adj
Inorg.A1-Control  221.6   54.1 389 0.003
Inorg.A2-Control  202.0   34.5 370 0.010
Inorg.N1-Control  256.4   88.9 424 0.000
Inorg.N2-Control  304.3  136.7 472 0.000
Organic-Control   349.4  181.9 517 0.000

Inorg.A2-Inorg.A1 -19.6 -187.1 148 0.999
Inorg.N1-Inorg.A1  34.8 -132.7 202 0.990
Inorg.N2-Inorg.A1  82.7  -84.9 250 0.692
Organic-Inorg.A1  127.8  -39.7 295 0.231
Inorg.N1-Inorg.A2  54.4 -113.2 222 0.929
Inorg.N2-Inorg.A2 102.2  -65.3 270 0.473
Organic-Inorg.A2  147.4  -20.1 315 0.115
Inorg.N2-Inorg.N1  47.9 -119.7 215 0.958
Organic-Inorg.N1   93.0  -74.5 261 0.576
Organic-Inorg.N2   45.2 -122.4 213 0.967
```

Interpretation: The analysis shows that all 5 of the nitrogen sources have population means which are larger than the control group (plots that received no nitrogen). The following are equivalent findings:

- all pairwise confidence intervals with the control contain only positive values; e.g., (54.1, 389), etc.
- the adjusted p-values (`p adj`) for comparisons with the control group are all very small.

Note that all intervals involving two nitrogen sources other than the control group include "0," which does not provide statistical evidence (at the familywise 95% confidence level) that the respective population mean yields are different.

We conclude that all population means for the nitrogen groups are larger than the population mean for the control group. However, there are no statistical differences among the 5 nitrogen population means themselves. Therefore, if we wanted to advise the experimenter on which nitrogen source(s) to use to maximize population mean yield, our recommendation should be to use any of the nitrogen sources other than the control group.

Assumptions: There are four statistical assumptions we are making with this analysis:

1. We have random samples of plots from each population (i.e., from all plots that would be given a particular nitrogen source).
2. The samples are independent. This is reasonable because the nitrogen sources were randomly assigned to the plots.
3. The yields of each nitrogen source are normally distributed.
4. The population variances for the 6 nitrogen source yields are equal; i.e.,
$$\sigma_1^2 = \sigma_2^2 = \cdots = \sigma_6^2 = \sigma^2.$$

We could look at qq plots for each of the six samples (6 qq plots in total) to check the normality assumption for each population. Personally, I don't think this will give us all that much information becuase the sample sizes (10 per group) are so small. The overall F is robust to mild normality departures anyway. The equal population variance assumption is critical. We could do a formal test for this too (Bartlett's or Levene's), but again the sample sizes are so small that this may provide limited information. An important part of any statistical analysis like this is clearly stating which assumptions you are making.

Summary: We performed a one-way classification analysis to determine if there are any differences in the population mean sugar beet yields for five nitrogen sources and a control group. The overall F test provided strong evidence the population means were not all equal. A follow-up analysis using Tukey pairwise confidence intervals revealed, at the 95% confidence level, that there were significant differences between each nitrogen group and the control group. However, there were no significant differences in the population mean yields for the 5 nitrogen sources themselves. If the goal is to maximize population mean yield, on the basis of the data in this experiment, we would advise the experimenter to use any of the nitrogen sources other than the control.

R CODE:

```
# Enter the data
Control = c(814.8,813.2,974.9,862.0,750.8,769.0,1026.0,849.4,946.3,997.9)
Organic = c(1235.3,1185.9,1117.0,1171.8,1284.7,1211.5,1288.9,1001.4,1428.4,1373.6)
Inorg.A1 = c(1157.5,1236.1,1074.3,1171.5,1031.3,1015.9,950.1,1108.5,1275.8,999.4)
Inorg.A2 = c(955.0,1039.4,1318.6,926.9,1230.1,835.3,1013.8,1128.3,1023.7,1353.5)
Inorg.N1 = c(1070.0,1153.1,940.1,998.5,1264.3,1351.1,1117.5,1389.3,1037.1,1047.3)
Inorg.N2 = c(1077.2,1137.7,1187.4,1335.8,1262.6,1126.7,1081.6,1134.6,1272.0,1231.3)

# Create side-by-side boxplots
boxplot(Control,Organic,Inorg.A1,Inorg.A2,Inorg.N1,Inorg.N2,
    xlab="",names=c("Control","Organic","Inorg.A1","Inorg.A2","Inorg.N1","Inorg.N2"),
    ylab="Yield (kg/acre)",col="lightblue")

# Concatenate all the data
Yield = c(Control,Organic,Inorg.A1,Inorg.A2,Inorg.N1,Inorg.N2)
# Create a treatment indicator variable
Source = c(
    rep("Control",length(Control)),
    rep("Organic",length(Organic)),
    rep("Inorg.A1",length(Inorg.A1)),
    rep("Inorg.A2",length(Inorg.A2)),
    rep("Inorg.N1",length(Inorg.N1)),
    rep("Inorg.N2",length(Inorg.N2))
    )
# Inform R that Source is a factor variable
Source = factor(Source)
# Analysis of variance
fit = lm(Yield ~ Source)
anova(fit)

# Create F pdf with overall F statistic marked
f = seq(0,10,0.01)
plot(f,df(f,5,54),type="l",lty=1,xlab="F",ylab="F(5,54) pdf",cex.lab=1.25)
abline(h=0)
points(x=9.19,y=0,pch=19,cex=1.5)

# Follow-up Tukey analysis
TukeyHSD(aov(fit),conf.level=0.95)
```