

**Planning and Analyzing
Qualitative On-Farm Research
with Microsoft Excel 2000:**

Single Year, Single Location

T.S. Murrell
Potash & Phosphate Institute

K.J. Moore
Iowa State University

Instructor Packet

PREFACE

This packet provides background material needed to teach an introductory training session on various aspects of on-farm research. The material is developed to take the student through all phases of conducting and analyzing a randomized complete block design experiment with a single factor and qualitative levels. Rate studies are not covered. The intention is to use a simple experimental approach, guiding the student through experimental design, data collection, data analysis, and interpretation. Target audiences are Certified Crop Advisers interested in conducting on-farm research in their trade areas.

If possible, students should come to this session with ideas of topics they would like to research. In the introductory and design sections, students are asked to share their ideas on the types of research they feel would be effective in their regions. In the data analysis section, an example data set is provided, to be worked through as a class. In longer workshops, students can create or bring their own data to analyze after the class example has been worked. After being given time to conduct the analyses, students can present their results and share their concerns and observations. Allowing the students to work through their own data set reinforces the concepts presented and builds confidence that the student can take the knowledge gained in the workshop and apply it to his/her situation.

EQUIPMENT LIST

This session is intended to be an interactive workshop, conducted with computers. Ideally, there should be as many students as there are computers. Microsoft Excel 2000 is the only software required. A list of major items needed by the instructor follows:

1. Flip chart or overhead projector with pens/markers
2. LCD projector
3. Computer for instructor
4. Computers for students (preferably one student per computer), with Microsoft Excel 2000 installed

CONTENTS

Introduction to on-farm research	5
Designing an experiment	7
Treatment selection	7
Asking a good question	7
Formulating your hypothesis	10
Creating a list of factors and levels for your experiment	10
Creating a list of treatments	11
Creating statistical hypotheses for your experiment	12
Error control in advance of analyses	13
Types of error	13
Selecting acceptable alpha levels	14
Creating a list of measurements	14
Controlling other factors	15
Experimental design	16
Experimental units	16
Design	19
Introduction to example data	28
Hypothesis	28
Factors and levels	28
Treatments	29
Statistical hypothesis	29
Error control	29
Experimental design	29
Yield data	29
Analyzing an experiment with Microsoft Excel 2000	31
Installing the Analysis ToolPak	31
Entering data into the spreadsheet	33
Creating a scatter plot	35
Numerical descriptive statistics for each treatment	41
Mean	41
Standard deviation	42
Creating a bar chart of treatment means	44
Performing an Analysis of Variance (ANOVA)	55

Understanding How an ANOVA Works	55
Running the Analysis of Variance Procedure	57
Interpreting Results from the Analysis of Variance	60
Least Significant Difference	62
Understanding the Least Significant Difference	62
Calculating the Least Significant Difference	63
Controlling Error in Using the Least Significant Difference	64
Selected References	65
Appendix	66

Introduction to On-Farm Research

Welcome to the workshop “Analyzing On-Farm Research with Microsoft Excel.” We are pleased that you were able to come today.

Interactive Activities

1. Have each instructor introduce themselves to the workshop participants.
2. Look around the room and be sure that students are well distributed around the computers. If not, ask some students if they would like to move to fill in vacant positions.

Deciding to begin a local research program is probably one of the most significant steps you can take to develop yourselves professionally. It is not for everyone. It takes commitment, patience, people-skills, a long-term vision, and a strong desire to know more about crop production in your area than anyone else. All of you who have shown up today have already made or are interested in making such a commitment. Let me just take a moment to encourage you to follow your interests. What are the benefits of creating a local research program?

1. **You become the local agronomic expert** by knowing more about what management practices work or don't work in your area.
2. **You differentiate yourself** from the dealership down the road. So many dealerships all try the same things in an effort to differentiate themselves – improved billing, customer appreciation dinners, advertising, etc.. Only a select few try to corner the market on agronomic information. You can almost count on your competition not doing it.
3. **You develop yourself professionally.** By conducting research to seek out useful, scientifically-based management practices, you gain the respect of your peers.
4. **You strengthen relationships with university extension.** Working with university extension scientists during your research, something that is strongly encouraged, creates

mutual respect and friendships and improves the flow of information back and forth between university and local research.

5. **You develop customer loyalty.** Your farmers will appreciate the opportunity to help you plan and conduct your research. They will also recognize your sincere efforts to help them discover what works best for them.

So as you can see, there are many benefits to local research programs that extend far beyond just the results of a particular study.

The types of experiments we are focusing on today are field-scale experiments with long strips. They are the most common type of experiment used by crop advisers and farmers. We consider studies with only qualitative treatments in this workshop. Qualitative treatments are those for which a numerical description has no real meaning, such as seed color, hybrids, or fertilizer type. We will not consider rate studies in this session. Rate studies require different analytical procedures than we describe here. This session will show you how to think about the treatments you want to investigate, how to lay them out in the field, and how to analyze the data statistically. We will be using Microsoft Excel 2000 to perform our data analysis. This software was selected because it is widely available, affordable, and requires no advanced statistical or computer programming skills to perform analyses.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

1. Ask students if they have been involved in putting out field trials. If so, what was their role and what types of questions were they examining?
2. Ask the students what they expect to get out of the workshop. List these on an overhead projector, flip chart, or computer screen.
3. Ask for specific questions that students want answered during the workshop. List these on an overhead projector, flip chart, or computer screen and refer to them during the workshop.

Designing an Experiment

TREATMENT SELECTION

Asking a Good Question

The first step in creating an experiment is to decide what it is you want to investigate. Imagine someone asking “Just exactly what are you trying to find out?” While it may be easy to come up with a quick list of questions, not all questions are worth the effort to investigate. Remember that conducting an experiment properly will take investments in time and money and will require quite a bit of coordination and quality control.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Have the class develop a list of resources necessary for conducting a field trial. Post this list for the class to see. The list should resemble **Table 1**.

Table 1. Partial list of resources needed to conduct a field experiment.

Field equipment	Specialized data collection equipment
Herbicide	Laboratory analyses
Insecticide	Computer software
Nutrients	Computer hardware
Seed	Grain hauling/handling
Labor	
Fuel	
Land rent	
Weigh wagon or yield monitor	

With all of these resources at stake, you want to be sure they are invested wisely. Here are some steps you can take:

1. **Become knowledgeable in the area in which you plan to conduct research.** Ask yourself what has been done in this area already? Without doing some review of what has been done before, you could end up investigating a question that has a well-known answer. You may also find that the question you want to investigate is a lot more involved than you first thought. If this happens, refine your question to become more specific. Sometimes it takes several different research projects to piece together an answer to a broader question.

Researching a topic well requires reviewing scientific principles and sound scientific studies. University and government sources are considered providers of unbiased information. Few people in production agriculture have time to visit a library on a university campus. Fortunately, more and more information is becoming available online. Table A1 in the appendix provides a partial list of websites that provide access to unbiased technical information.

2. **Keep your question simple.** It is easy to ask very complex questions, especially now that we are more aware of variability in the field and recognize increasingly that many factors can interact to affect yield and quality. Analyzing complex designs can be done, but it requires advanced statistical analysis software and good knowledge of statistical principles. For most applications a simple, straightforward approach is best.
3. **Select factors you think most limit yields or quality under current production practices in your area.** This gets back to the first suggestion – becoming knowledgeable in your research area. There are many factors that can influence yield and quality, but some have greater impacts than others. Knowledge about the factors you are considering is necessary to estimate their possible significance. Remember that you will be working under field conditions, where Mother Nature is running rampant. If the factors you are examining are not the most limiting for yield or quality, their effects may not be measurable. This is why it is important in any experiment to ensure, as much as possible, that other factors not of interest to your question are in sufficient quantity not to limit yield or quality
4. **Consider the possible impacts of your research.** The research you want to conduct may be scientifically valid, appropriate for your area, and have a good chance of producing meaningful results, but that's only part of the solution. For a new management practice to become widely adopted, it must also fit well into the logistic and economic aspects of the farmer's operation. Challenges in either of these areas will keep new management practices, even those with proven benefits, from being adopted.

Interactive Activities

1. Ask students if they are aware of examples of research that have led to important changes in management practices.

Examples:

- a. New hybrid development
 - b. New pesticide formulations
 - c. Fertilizer rate studies leading to recommendations
 - d. No-tillage practices
2. Ask students if they are aware of research projects that produced improved management practices that were poorly adopted.

Examples:

- a. Starter fertilizer is known to be beneficial in areas where it is not widely adopted because farmers do not want to worry about the logistics of nursing a fertilizer tank while planting
 - b. Applying nitrogen in the spring, prior to planting corn, has been shown to make most efficient use of the nutrient, but small numbers of days suitable for field work and wetter conditions in the spring lead many farmers to apply their nitrogen the previous fall, even though there are greater risks of nitrogen loss.
3. Ask students what can be done to ensure there will be interest in the research and increase the chances that results will be implemented. This question is important enough to warrant special discussion, below.

There are some things you can do to increase the probability that new management practices arising from your research will be adopted by farmers in the area. Below are a list of suggestions.

- a. **Form a farmer research advisory group.** As you visit with farmers in your area, inform them of your intentions to begin some research projects. Share with them why you want to begin the program and ask if they would like to be part of a group that helps determine what type of research gets done. Keep your group small and manageable, composed of the most enthusiastic and progressive of your farmer customers. Look for the farmers that ask you the agronomic questions that don't have easy answers.
- b. **Consider having an area for small plot research.** Such areas are good for many reasons:
 - i. Small plots, because they are more easily controlled, provide you with additional opportunities to investigate cause and effect relationships

- ii. They allow you to test some theories without risking yield or quality reductions on large areas
 - iii. Small plots, with alleys, create a good place to host field days. Farmers can walk through the plots and see first hand the responses you are discovering.
- c. **Hold meetings to discuss your research results.** Do this the first year, even if results showed little differences. The purpose of such a meeting goes beyond presenting results. It creates awareness of your efforts and shows that you are attempting to find better management practices for your customers.
- d. **Work with university Extension scientists.** They can share their experience in putting out research trials and can provide you some background and guidance on the topic of interest. The relationships formed can be invaluable for both parties.

Formulating Your Hypothesis

A hypotheses should be a simple statement that captures what you are trying to discover.

Examples:

1. Does site-specific nutrient management lead to higher yields than current whole-field management?
2. Are there yield differences among three new hybrids?
3. Does taking nitrogen credits from manure lower my yields?

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Ask students to provide examples of hypotheses of interest to them. Post these for the class to see and refer to them in the following section.

Creating a List of Factors and Levels for Your Experiment

Factors are the variables you want to include in you study. In this workshop, we will focus on studies examining only one factor. You may have more than one factor in an experiment, but Microsoft Excel is not readily capable of analyzing these studies in the way we will be designing them in this workshop. Levels are the various quantities or aspects of a given factor. Factors and levels lead to the treatments you apply.

Examples:

1. Does site-specific nutrient management lead to higher yields than current whole-field management?

Factor	Levels
Management approach	Whole field nutrient management Site-specific nutrient management

2. Are there yield differences among three new hybrids?

Factor	Levels
Hybrid	Hybrid 1 Hybrid 2 Hybrid 3

3. Does taking nitrogen credits from manure lower my yields?

Factor	Levels
Management approach	Manure applied and N credit taken Manure applied and no N credit taken

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Develop, with the class, a list of factors and levels for the hypotheses the class formed previously.

Creating a List of Treatments

For our single factor examples, the levels define the treatments. In multi-factor experiments, treatments will be specific combinations of levels of various factors.

Examples:

1. Does site-specific nutrient management lead to higher yields than current whole-field management?

Treatments:

- a. Whole field nutrient management, based on recommendations for the entire field
- b. Site-specific nutrient management, based on refined recommendations within a field

2. Are there yield differences among three new hybrids?
Treatments:
 - a. Hybrid 1
 - b. Hybrid 2
 - c. Hybrid 3

3. Does taking nitrogen credits from manure lower my yields?
 - a. Manure applied and its nitrogen content considered for inorganic nitrogen rate determination
 - b. Manure applied and its nitrogen content not considered for inorganic nitrogen rate determination

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Develop, with the class, a list of treatments using the factors and levels created by the class in the previous section.

CREATING STATISTICAL HYPOTHESES FOR YOUR EXPERIMENT

Preparing statistical hypotheses requires that you make some statement about statistical parameters in your study. You will most likely be concerned with the means, or averages, of your treatments. There are two types of statistical hypotheses: the null hypothesis and the alternate hypothesis. The *null hypothesis* usually states that no differences exist among treatment means.

Examples of null hypotheses:

1. The mean yields from site-specific management and whole-field management are the same.
2. The new hybrids of interest all have the same mean yield.
3. Mean yields are the same whether or not manure nitrogen is credited.

The *alternate hypothesis* is a statement that contradicts the null hypothesis. This hypothesis is accepted if the null hypothesis is rejected.

Examples of alternate hypotheses:

1. The mean yields from site-specific management and whole field management are not the same.
2. At least one new hybrid has a mean yield different from the others.
3. Mean yields associated with taking manure nitrogen credits are not the same as mean yields associated with ignoring nitrogen credits.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Develop, with the class, statistical hypotheses for the treatments the class created in the previous section.

ERROR CONTROL IN ADVANCE OF ANALYSES

Types of Error

When you conduct your experiment, you want to be sure that you draw the right conclusions. There are two types of errors you could make: type I and type II errors.

1. A *type I* error occurs when you find a statistically significant difference among treatment means when in fact, there really was none. In statistical terms, a type I error occurs if you reject the null hypothesis when it should have been accepted. Statisticians denote the probability of making a type I error with the Greek letter alpha, α . The probability is expressed in decimal form. For instance, wanting only a 5% chance of making a type I error is denoted as $\alpha=0.05$. People report this in a variety of ways, some examples are $\alpha=0.05$, $\alpha_{0.05}$, $\text{alpha}=0.05$, and $\text{alpha}_{0.05}$. Since creating a new management practice based on research results usually involves some risk, most experimenters want to be reasonably sure that the significant differences they discover are real. They don't want to start telling people to use their new management strategy, only to find it doesn't really work. To avoid making this error, most experimenters choose a fairly low alpha probability level. Alpha levels of 0.10, 0.05, and 0.01 (10%, 5%, and 1% chances, respectively), are commonly used. Controlling a type I error is usually of primary concern in an experiment.
2. A *type II* error occurs when you find no significant difference among treatment means, when in fact, one existed. In statistical terms, a type II error occurs if you accept the null hypothesis when it should have been rejected. This is usually a less serious error, since it means accepting the status quo rather than the new, riskier management practice.

Type I and type II errors are related. When you lower the probability of making a type I error, you increase the probability of making a type II error. The only way to avoid any chance of error is not to make a decision at all!

It is important to decide what levels of risk you are willing to accept from the outset. Stick to these during your analysis and interpretation. It can be tempting to compromise your original intentions if the data show some promise.

Selecting acceptable alpha levels

Selecting the probability levels in advance of statistical analyses is an important step. There are two primary objectives in the statistical analyses we will perform: 1) determining if there is a difference among treatments, and 2) determining which treatments are different.

An analysis of variance (ANOVA) is performed to determine if there were significant differences among treatment means. It tests whether the null or alternate hypothesis should be accepted. The alpha level selected for this analysis therefore controls the probability of stating that there was a significant difference, when in fact there was none.

If a significant difference is found, the next step is to determine which of the treatments were different. This will be accomplished by a procedure termed the protected least significant difference (protected LSD). Alpha levels selected for this procedure will control the probability of stating that one treatment mean was different from another, when in fact there was no real difference.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Discuss with the class to what extent they want to control error in the experiments they are considering. Have the class select alpha levels for their ANOVA and LSD analyses.

CREATING A LIST OF MEASUREMENTS

As you design your study, be thinking about what measurements should be taken during the experiment to help you draw the right conclusions. Certainly, yield will always be measured and recorded, but final yield is the result of many interacting factors. How can you be more confident that the treatments you applied were truly causing the yield responses observed? To help explain the yield responses seen, you may need to take some other measurements as well. For instance, in fertility studies, it is customary to take plant tissue samples and have them analyzed for nutrient content. Examining nutrient levels in the plants helps verify the observed impacts of the fertility treatments. If plants in one fertility treatment had nutrient levels below sufficiency levels and yield reductions were also measured, the experimenter has more confidence that the yield reductions observed were truly the result of the treatment. Deciding which additional measurements to take requires expertise in the field of study. If you aren't sure, consult with one of your university extension scientists or other trusted information sources.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Create a list of measurements that should be taken for the experiments the class has formulated.

In every experiment, there are a set of data that should be recorded. These are listed in **Table 3**.

Table 3. Data that should be recorded in every experiment.

Experiment	Cultural practices
Name of investigators	Tillage
Objectives	Soil conditions at planting
Plot plan (experimental layout)	Residue cover (%)
	Planting date, depth, rate, row spacing
	Variety/hybrid
Location	Type of planter
Township, range, section, and quarter	Fertilizer sources, rates, dates, and placement
Names and contact information of cooperators	Manure sources, rates, dates, and placement
	Manure nutrient composition
Site characteristics	Herbicide source, rate, date, and placement
Soil series present in experimental area	Insecticide source, rate, date, and placement
Soil physical properties (compacted areas, etc.)	
Soil chemical properties (complete soil tests)	
Topography (map is preferable)	Experimental units
Location of drainage tiles	Size (length, width, number of rows)
Cropping history	Emergence, plant height
Chemical and fertilizer use history	Planted population
Manure history	Harvest population
	Size of harvested area
Climatic conditions	Harvest date
Rainfall/irrigation daily amounts	Grain moisture
Maximum and minimum daily temperatures	Harvest weight

CONTROLLING OTHER FACTORS

The final step before laying out your experiment is to list factors that may interfere with the treatments you have chosen for your study. Decide which factors you can control and which ones you can't. For the factors you can control, come up with strategies to ensure that they are not at levels that may limit yield or quality responses. Coming up with the right factors requires knowledge in the topic area you are planning to investigate.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Develop, with the class, a list of factors that could interfere with the hypotheses developed by the class.

EXPERIMENTAL DESIGN

Experimental Units

Experimental units are the areas in the field that will contain your treatments. There are two important considerations: width and length.

Width

All experimental units should be a uniform width and placed side by side to the extent possible. Before the rest of the experiment can be designed, you must first determine how wide an experimental unit must be to accommodate a treatment. For instance, comparing site-specific to whole-field nutrient management requires fertilizer spreading equipment. For pneumatic applicators, part of the consideration will be the boom width of the spreader and the width of the harvester. You may find that there is no way to get everything to match up evenly. In this case, consider putting in border rows. These are rows that can be planted to solve logistical problems, but their data are not included in analyses. Including border rows is also advisable if applications of treatments might result in overlaps or gaps between passes. Let's consider a couple of examples where boom width defines the width of the experimental units.

Example: Applicator boom width 60 ft., 8 row harvester, 30 in. rows. For an 8-row harvester and 30 in. rows, the width of the harvest swath is 8×30 in., or 240 in. Dividing this by twelve converts inches to feet, or $240 / 12 = 20$ ft. Therefore, 1 pass of the harvester is 20 ft. wide. The boom is 60 ft. wide. Therefore, the boom width is evenly divided into 3 passes of the harvester.

The most obvious way to harvest this study is to align the harvester with the edge of the experimental unit and begin collecting data. In this way, one experimental unit is defined as being 60 ft. wide and is harvested every three passes. If a yield monitor is used, it would be advisable to record data from each pass separately. This will allow more flexibility in statistical analyses.

Although this is an obvious way to harvest the experimental units, there are some potential problems with doing this. In the field, application equipment cannot do a perfect job. There may be gaps or overlaps between passes of the applicator. This means that the first and third passes of the harvester may not be collecting yield data truly representative of the intended treatment. One option is to exclude the data from the first and third passes of the harvester in each treatment, leaving only the data from the second pass, located in the middle of the boom width (**Figure 1**). In this way, the first and third passes become border rows. However, this would allow only one pass of the harvester to be used in statistical analyses, limiting the amount of usable data.

An alternative scheme is to begin with the harvester straddling the edge of the first treatment. This places 4 rows of the harvester outside and 4 rows inside the fertilizer application area (**Figure 2**). Harvesting with this initial offset will place two harvest passes away from the edges of the application area, increasing the amount of usable yield data.

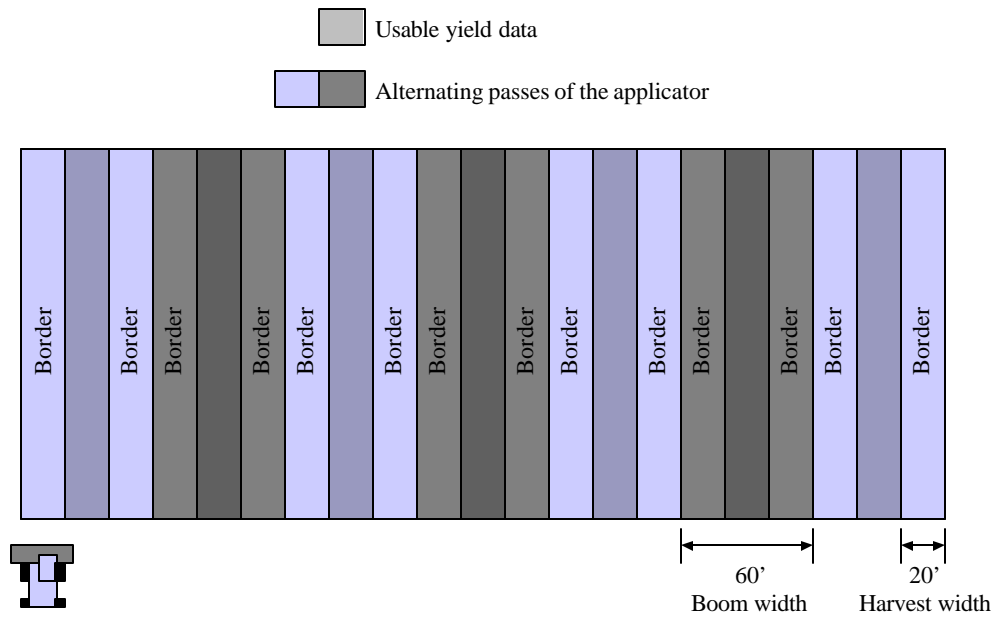


Figure 1. Harvest pattern of 60 ft. wide experimental units using the first and third harvest passes within each applicator pass as borders.

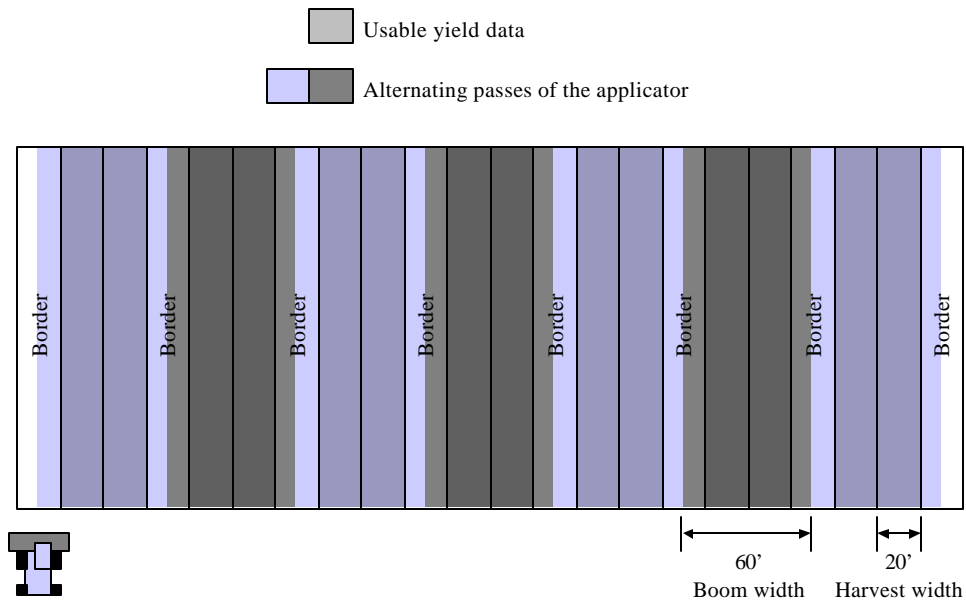


Figure 2. Harvest pattern of 60 ft. wide experimental units using an initial 10 ft. offset. Every third harvest pass is a border.

Example: Applicator boom width 70 ft., 8 row harvester, 30 in. rows. Often, the width of application equipment does not match up evenly with the width of harvest equipment. This requires more advanced planning to ensure that a maximum amount of usable yield data can be collected. A common mismatch is a 70 ft. application boom and a 20 ft. (8-rows with 30 in. spacing) harvest width. In this case, the width of the application boom equals the width of 3.5 passes of the harvester. A possible arrangement is shown in **Figure 3**. In this scenario, the harvester is aligned with the edge of the first spreader pass. Border rows are then placed at the edges of the applicator passes. This arrangement still allows 2 harvest passes per experimental unit, maximizing the amount of usable yield data from the experiment.

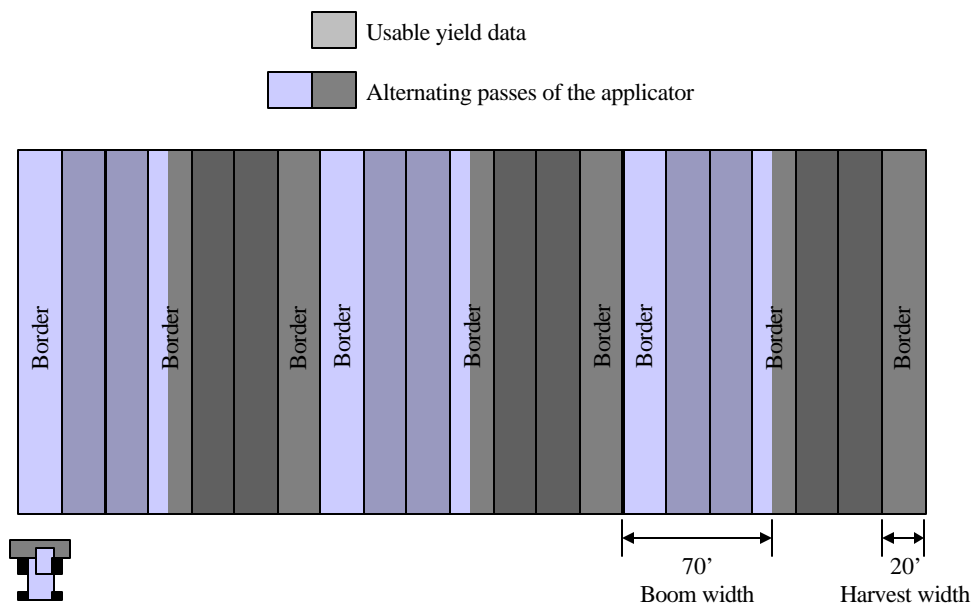


Figure 3. Harvest pattern of 70 ft. wide experimental units. After the initial harvest pass, the number of border rows alternates between 1 and 2.

Length

Practically, the length of experimental units are likely to be close to the length of the field in which the experiment is being conducted. Indications from previous studies indicate that long, narrow strips reduce experimental error (Wuest et al., 1994; Shapiro et al., 1989). Research conducted with wheat and barley under dryland conditions in the Pacific Northwest has shown that experimental units 750 ft. or longer have much less experimental error than plots 250 or 500 ft. in length (Wuest et al., 1994).

Interactive Activities

Equipment: Calculators and flip chart, overhead projector, or computer screen

Provide the class time to calculate appropriate sizes of experimental units for the treatments the students want to investigate. Allow each student to share with the class how he/she performed the calculations. As students do these calculations on their own, circulate around the room to provide assistance.

Design

Experimental design must be chosen carefully. As an experimenter, you want to maximize the power of your study. You want to have a good chance of detecting differences among treatments if such differences truly exist. The variability in the conditions of your experiment will greatly affect your ability to detect significant differences. Generally, the greater the variability, the less your chances of detecting treatment differences. Much of this variability will come from the site on which you decide to conduct research. Therefore, site selection is important. Try to choose a site that is representative of the area across which you want to generalize. Be sure you are familiar with the history and characteristics of the field you have chosen. What has the historical management been? Was there once a farmstead on the property? Was manure applied previously? Was the field previously divided into smaller fields? If the field has yield monitor data, how uniform have the yields been across the field? These questions and more can help you find a site that has the conditions you want.

Once you have chosen an appropriate site for your experiment, the next step is to try and control the variability within it. In this design phase, you are deciding where your treatments should be placed in the field. You want to do this in a manner that will minimize the impact of the field variability upon your ability to draw correct conclusions. There are many possible designs to choose from; however, in this workshop, we focus on only one: the randomized complete block design (RCBD). There are three primary components to this design: replication, randomization, and blocking.

Replication

Replication is repeating the number of times you put the treatment in the field. If you were to put each treatment in the field only once, you might end up putting each treatment in distinctly different areas of the field, biasing your results.

For instance, suppose you decide to examine two treatments. Treatment 1 is on the west half of the field, and treatment 2 is on the east half of the field. The west half of the field had large areas receiving manure about five years ago. The east half did not. Let's suppose that the residual effects of the manure could increase the yields of either treatment. However, because only treatment 1 is in the manured area, its average yield is greater than that of treatment 2. By just comparing the yield of the two treatments, you might conclude that treatment 1 was better than treatment 2, but was it really? How much of the higher yield of treatment 1 was from the effects

of the manure and how much was from the management associated with the treatment? Without replication, you have no way of knowing, and you could easily draw the wrong conclusions about the superiority of treatment 1.

Replication helps you measure the variability in your experiment. Variability is the spread in your data. For instance, if you put out 2 treatments only once and get corn yields of 150 bu/acre for treatment 1 and 160 bu/acre for treatment 2, you might conclude that treatment 2 was the better treatment. But would you get that same 10 bu/acre difference at the same yield levels if you were to repeat the experiment? Putting out the treatments again and again in the field helps you develop a picture of the range in yields produced by each treatment. A wider range in yields means greater variability. Quantifying how this variability affects your ability to draw conclusions is the job of the statistical analysis. Without some measure of variability, there is no way to analyze your data statistically. The more replications you have, the better able your statistical analysis will be able to assess variability, increasing the power of your experiment to detect treatment differences when they truly exist. Most field experimenters will use 4 replications at minimum. This is a commonly accepted number among applied scientists. However, you should include as many replications as you can, considering costs and risks associated with the treatments.

Randomization

Randomization is the process of mixing up the order of your treatments. It helps you draw conclusions that are representative of the area you have chosen. If you do not randomize, you can bias your results.

As an example, suppose you decide to examine two treatments. You know you want to replicate them and decide simply to alternate them one after another across the field. Such an arrangement may also correspond to your tile spacing, with one of the treatments always ending up over the tile line. The improved drainage for every occurrence of the treatment across the field would probably increase the treatment's average yield. However, much of this higher yield resulted simply from the treatment's position in the field, not from the treatment itself. Randomization minimizes the chances that something like this will happen in your experiment.

To randomize the placement of your treatments, simply write the number of each treatment on the back of business cards (or some other evenly-sized piece of paper), mix them up, and then draw one card at a time, at random. Each time a card is drawn, put it aside after recording the treatment. The order in which the treatments were selected is the order they will be placed in the field, for instance going from east to west.

Blocking

When you randomize treatments, they can end up anywhere in the field. By chance, more replicates of one treatment may end up in one area of the field than in another. If the field is fairly uniform, this may not be a problem. However, if you know that there is variability in the field, it is best to restrict the randomization of your treatments. You can do this by creating smaller areas, termed *blocks*, within which the treatments are randomized (**Figure 4**). The

purpose of blocking is to create smaller, more uniform areas where observed differences will largely be due to the treatments themselves. This will limit the influence of the treatment's position in the field. To create more uniform areas, blocks will most likely need to be smaller in size. This reinforces the need to restrict experiments to just a few treatments.

In the RCBD, blocks are areas that accommodate one complete set of randomized treatments, hence the term randomized **complete** block design. To meet the commonly accepted number of 4 replications for each treatment, 4 blocks are required (Figure 4).

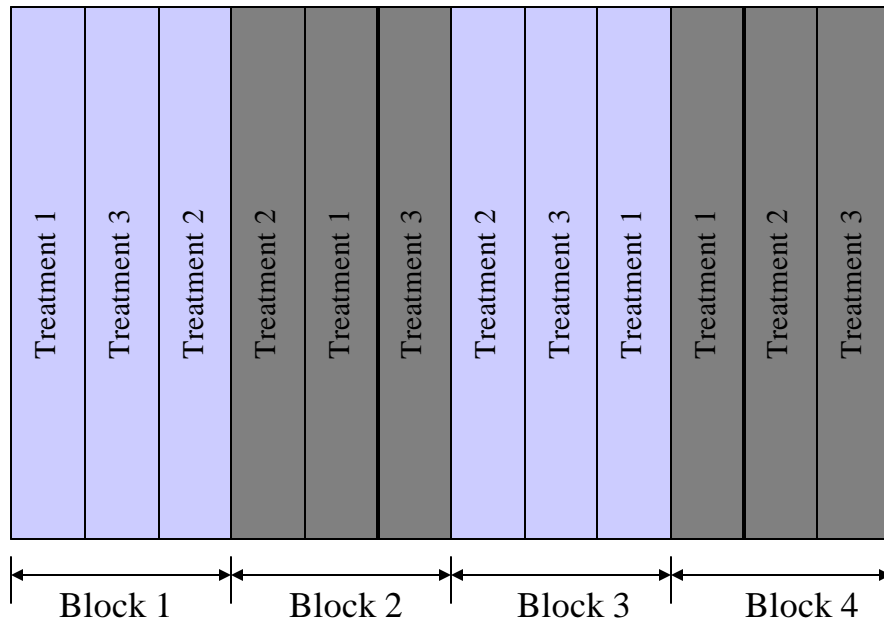


Figure 4. Randomization of 3 treatments within blocks, replicated 4 times.

Because the blocks themselves strive to be relatively uniform areas, each one should be treated uniformly. For instance, when plots are harvested, it may happen that the entire experimental area cannot be harvested in the time left on a particular day. It is important not to quit harvesting in the middle of a block. Instead, try to finish harvesting the block before stopping. If other blocks are not harvested until the next available time, that is acceptable. By blocking in the first place, you assumed that each block would have a different set of conditions. This also brings up another reason for choosing this type of design. It allows you to break up the experimental area into smaller, more manageable areas. Having such flexibility is often needed in field experiments.

Now that we have a complete understanding of the RCBD, it may be useful to point out that randomizing treatments within blocks can be done easily in Microsoft Excel. The RAND() worksheet function returns a random number when entered into a formula. The first step in this approach is to list all of your blocks in column A of an empty worksheet. Make sure that each

block is listed the same number of times as there are treatments. Next, enter the treatment names in column B for each block. As an example, consider a RCBD experiment with 3 treatments and 4 blocks.

Student Tasks

Enter the block and treatment designations into a blank spreadsheet

1. In cell A1, type in “Block”
2. In cell B1, type in “Treatment”
3. In cell C1, type in “Random”
4. Type the values for blocks and treatments as they are shown in **Figure 5**.


	A	B	C
1	Block	Treatment	Random
2	1	1	
3	1	2	
4	1	3	
5	2	1	
6	2	2	
7	2	3	
8	3	1	
9	3	2	
10	3	3	
11	4	1	
12	4	2	
13	4	3	

Figure 5. List of treatments in each block to be randomized.


Next, enter the formula “=RAND()” into the first cell in Column B and then copy it down so that every treatment listed in Column A has a corresponding random number listed in Column B.

Student Tasks

Calculate random numbers using the RAND() function

1. In cell C2, enter “=RAND()” and press **ENTER (Figure 6)**.
2. Move the cursor back to cell C2 and click the left mouse button to select it.
3. On the standard toolbar, click the  icon to copy the cell contents. The border around the cell will become animated.
4. Move the cursor to cell C3 and click and hold down the left mouse button. While holding down on the left mouse button, drag the cursor to cell C13. Once the

cursor is on cell C13, release the left mouse button. The selected cells will now be shaded (**Figure 7**). The process of clicking and holding down the left mouse button, dragging the cursor, and then releasing the left mouse button is referred to as “click and drag”. Now that we have this terminology at our disposal, we can simplify the directions just given to “click and drag from C3 to C13”.

5. On the standard toolbar, click the  icon to paste the contents of cell C3 into all of the highlighted cells. Random numbers will now fill the selected cells in column C (**Figure 8**). The numbers you generate will be different from those in Figure 8, since they are random.

	A	B	C
1	Block	Treatment	Random
2	1	1	=RAND()

Figure 6. Entering the RAND function into a cell

	A	B	C
1	Block	Treatment	Random
2	1	1	0.8889582
3	1	2	
4	1	3	
5	2	1	
6	2	2	
7	2	3	
8	3	1	
9	3	2	
10	3	3	
11	4	1	
12	4	2	
13	4	3	

Figure 7. Block selected to receive the contents of the copied cell.

	A	B	C
1	Block	Treatment	Random
2	1	3	0.3211658
3	1	1	0.6308023
4	1	2	0.7345911
5	2	2	0.301041
6	2	3	0.0228174
7	2	1	0.2044406
8	3	1	0.0486705
9	3	2	0.7372811
10	3	3	0.3341336
11	4	2	0.3319543
12	4	1	0.2225965
13	4	3	0.9337146

Figure 8. Random numbers generated in each cell.

The final step is to sort the treatments within the blocks, based upon the order of the random numbers generated in column C.

Student Tasks

Sort the treatments within the blocks, based upon random numbers

1. Click and drag from A2 to C13 to select this block of cells
2. From the main menu click on **Data**, then click on **Sort...**. This will bring up a dialog box labeled **Sort** (Figure 9).
3. Under the **My list has** label at the bottom of the box, click on the radio button to the left of the label **Header row** (Figure 10).
4. Under the label **Sort by**, click on the downward-pointing triangle on the right side of the white box. This will bring up the list of column names you typed in earlier. Click on **Block** (Figure 10).
5. Click on the radio button to the left of the label **Ascending**.
6. Under the label **Then by**, click on the downward-pointing triangle on the right side of the white box. Select **Random** (Figure 10).
7. Click on the radio button to the left of the label **Ascending**.
8. All of the selections in the Sort dialog box should look like those in **Figure 10** when completed. Verify that this is the case and click **OK**.
9. The output from this procedure will resemble that in **Figure 11**. Treatments have been randomized within each block.

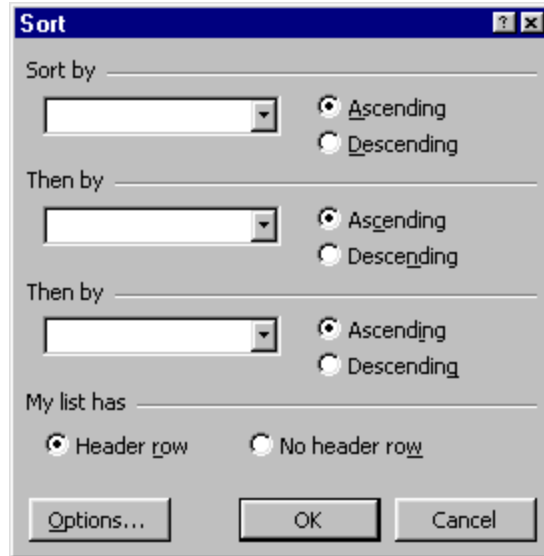


Figure 9. The Sort dialog box.

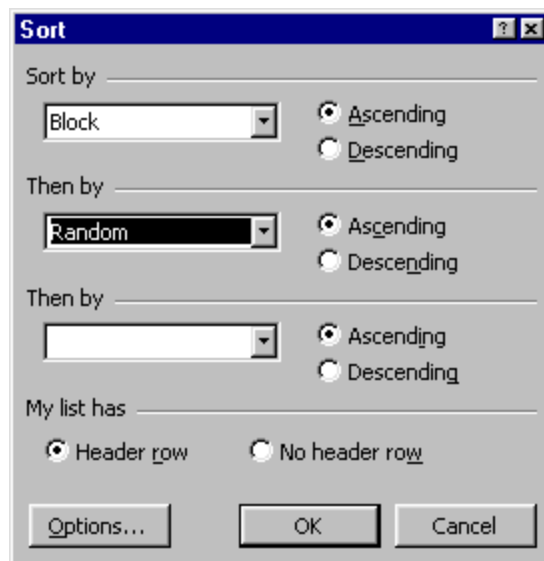


Figure 10. The Sort dialog box after setting parameters.

	A	B	C
1	Block	Treatment	Random
2	1	3	0.1303466
3	1	1	0.0195417
4	1	2	0.9848089
5	2	2	0.2014839
6	2	3	0.5236728
7	2	1	0.5422373
8	3	1	0.9221394
9	3	2	0.5073527
10	3	3	0.5774662
11	4	2	0.7327529
12	4	1	0.9362103
13	4	3	0.1712376

Figure 11. Treatments randomized within each block.

To assign treatments to experimental units, simply start at the top of the list and work down, completely assigning treatments to one block before moving to the next block.

Arrangement of experimental units and blocks

Arranging field-long experimental units and blocks within a field is done differently than for small plot experiments. In traditional small plot studies, blocks are arranged perpendicular to variation sources in the field. For instance, if there are two soil types in the field, one block is usually placed within each soil mapping unit, and a third block placed in the transition zone. Field-long strips, however, will likely cross many sources of variation. In such cases, each block should be placed in the field so that each experimental unit within the block encounters the same sources of variation across the field. For instance, in a field with rolling topography, arrange the experimental units lengthwise down the same slope (**Figure 12**).

There are both strengths and weaknesses to using field-long strips as experimental units. The strength of this approach comes from the large area and several sources of variation covered by the experiment. This allows conclusions to be applicable to the entire field or at least to large areas within the field. The weakness of these studies is reduced control over variation. Compared to small plot studies, there is more variability among plots (or strips) within blocks. This creates more noise in the analysis of variance for testing the significance of treatments, reducing the precision of the analysis.

Interactive Activities

Equipment: Flip chart, overhead projector, or computer screen

Have students randomize, replicate, and block their treatments for the experimental areas they wish to investigate.




Figure 12. Alignment of experimental units perpendicular to variability in slope.

Introduction to the Example Data Set

HYPOTHESES

There are two hypotheses for this experiment. The first is whether or not site-specific management is better than conventional whole-field management. The second is whether or not an integrated approach to site specific management is superior to site specific management concerned primarily with fertilizer management based on grid-sampling. In this context, integrated site-specific management uses a multi-disciplinary approach to manage all types of information, beyond just soil fertility. Weeds, insects, and other factors are also managed spatially and temporally.

FACTORS AND LEVELS

This is a single factor study with 3 levels. The factors and levels, created from the hypotheses, are listed in **Table 4**.

Table 4. Factors and levels for the example data set.

Factor	Levels
Management approach	Whole field management
	Site-specific nutrient management based on grid sampling
	Integrated site-specific management, incorporating temporal and spatial management of all inputs, not just fertilizer

TREATMENTS

Because this is a single factor experiment, the levels of the factor determine the treatments. Therefore, there are three treatments, which we will abbreviate whole field, site-specific, and integrated.

STATISTICAL HYPOTHESES

The statistical hypotheses for the treatment means are as follows:

Null hypothesis: the means of all the treatments are equal;

Alternate hypothesis: one or more of the treatment means are not equal.

ERROR CONTROL

Currently, site-specific management is an expensive venture. Therefore, there is a relatively high risk to its adoption as a new management practice. For these reasons, the experimenters wanted to be very sure that they minimized a type I error: they did not want to suggest that site-specific management practices were superior to conventional, whole field management practices if in fact they were not. To minimize the probability of making such an error, they set their alpha levels for both the ANOVA and the protected LSD at 5% ($\alpha=0.05$).

EXPERIMENTAL DESIGN

This experiment was conducted at the Southwest Experiment Station near Lamberton, Minnesota. The design used was a RCBD with 4 blocks (replications or reps), and three treatments. **Figure 13** shows the arrangement of experimental units, treatments, and blocks. This study examines corn and soybean responses at separate but adjacent locations. In the northeast part of the field is the corn experiment, and in the southwest half of the field is the soybean experiment. Each year, corn and soybeans are alternated at each location. This keeps the study areas in a corn and soybean rotation, commonly used by farmers in the area, and allows corn and soybean data to be collected each year. Between the two experimental areas is a grass drainage way. The northeast area of the field slopes from east to west and the southwest area of the field slopes from west to east. Therefore, experimental units are aligned in an east/west orientation, perpendicular to boundary changes in soil mapping units, drainage, etc.

YIELD DATA

The data we will use from the experiment comes from the northeastern part of the field planted to corn. These data are shown in **Table 5**. Notice the arrangement of data. Each treatment is a separate row in the table, and each block is a separate column. We will arrange these data similarly in a Microsoft Excel spreadsheet to analyze them.

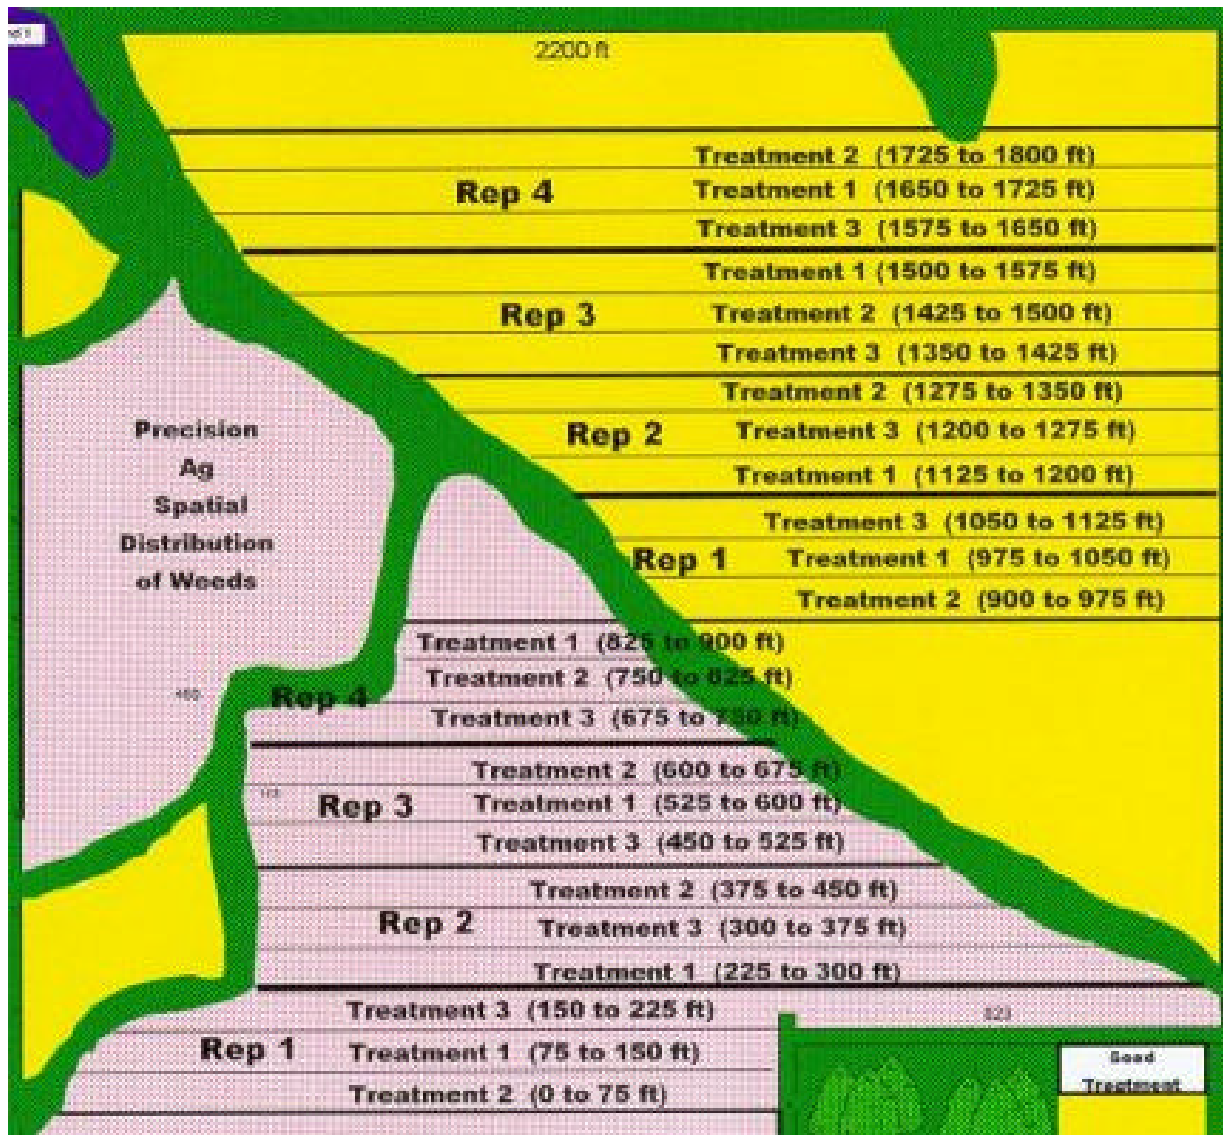


Figure 13. Experimental design yielding the example data.

Table 5. Experimental corn grain yield data from the example study.

Treatment	Block			
	1	2	3	4
	(bu/acre)			
Whole field	137.8	158.5	165.9	146.0
Site specific	147.1	153.3	170.5	161.4
Integrated	143.4	147.9	157.9	166.5

Analyzing an Experiment with Microsoft Excel 2000

INSTALLING THE ANALYSIS TOOLPAK

Before using Microsoft Excel to analyze field experiments, it may be necessary to install the Analysis ToolPak. Without this software installed, Microsoft Excel will not have the features available to perform the analyses. Installation is very simple and requires just a few steps.

Student Tasks

1. On the menu bar, click on **Tools**. This will bring down a drop list of options. Click on the **Add-Ins...** item, as shown in **Figure 14**.
2. After clicking on the Add-Ins... item, a dialog box will appear (**Figure 15**). This box has a list of items that can be or have already been installed. Installed items will be denoted by a checked box to the left of the item description. Uninstalled items will have a blank box.
3. In the list of add-ins, click the white box to the left of the **Analysis ToolPak** option and then click the **OK** button, as shown in **Figure 16**. This will install the Analysis ToolPak into Excel.

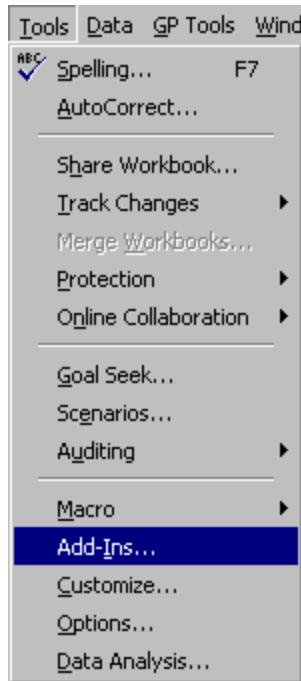


Figure 14. Drop list of options under Tools.

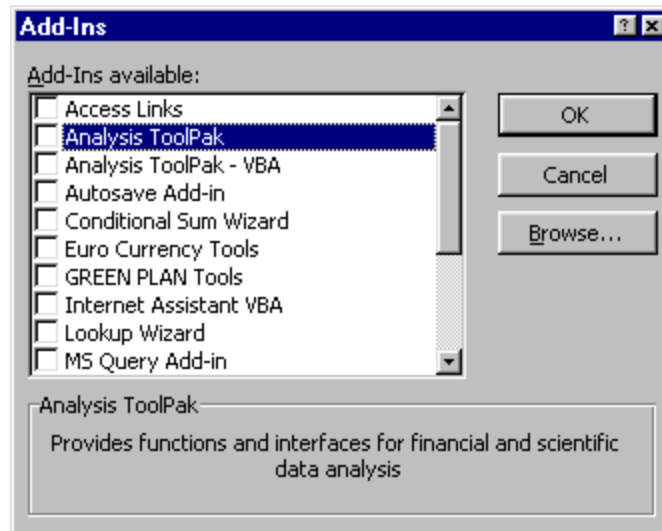


Figure 15. List of available add-ins.

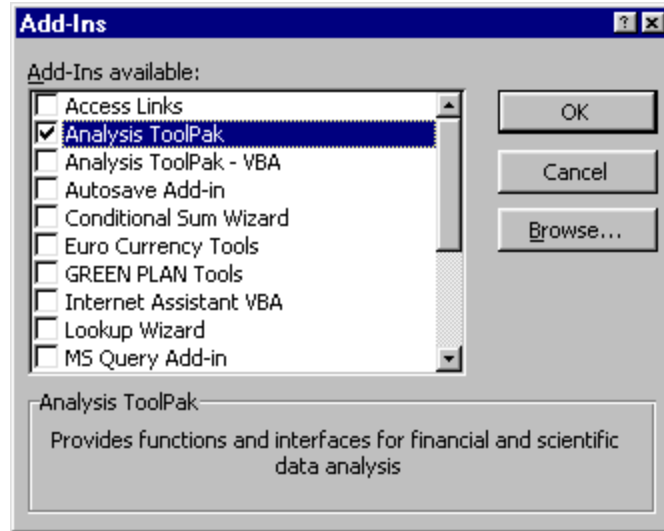


Figure 16. **Add-Ins** dialog box with the **Analysis ToolPak** option selected for installation

ENTERING DATA INTO THE SPREADSHEET

When you first open Microsoft Excel, a new workbook is created. This workbook contains one or more worksheets, each labeled Sheet1, Sheet2, etc. (**Figure 17**). You can add or subtract new sheets from the workbook. Each worksheet has a label, which may be found on a tab attached to the bottom of each worksheet. Throughout this workshop, the class will all use the same labels for each spreadsheet, to facilitate navigating through the workbook together.

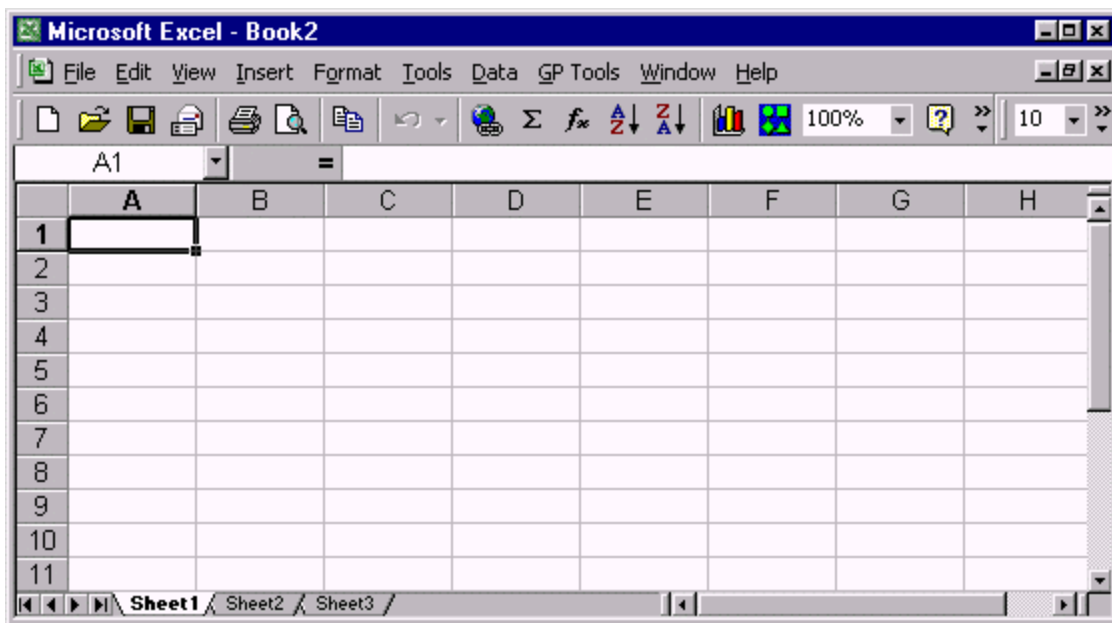


Figure 17. New workbook showing 3 worksheets, labeled Sheet1, Sheet2, and Sheet3.

Student Tasks

Change the name “Sheet1” to “data”.

1. You can easily change a label by double clicking the left mouse button on it. Double-clicking will highlight the sheet label, as shown in **Figure 18**.
2. Once the label is highlighted, you can type in a new name for the label. Type in the word “data”, since this will be the worksheet containing the example data (**Figure 19**).

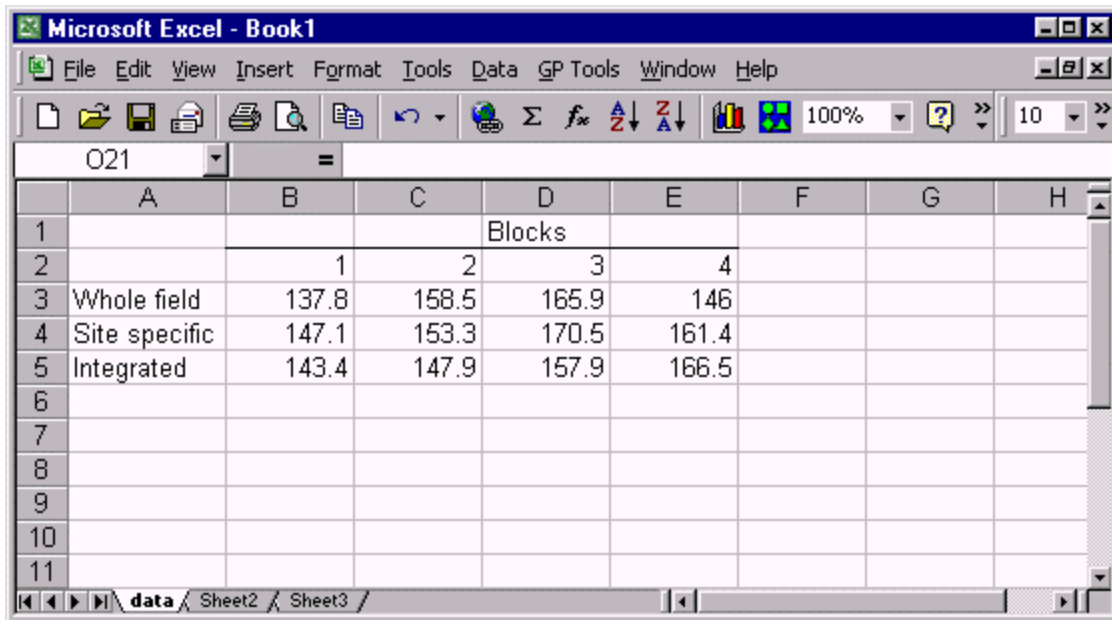


Figure 18. Highlighted sheet label after double clicking the left mouse button.



Figure 19. New label “data” typed into the spreadsheet tab.

The newly created “data” spreadsheet will hold the data from the experiment. To perform all of the analyses most easily, it is suggested that data be arranged with each treatment as a separate row and replications (or blocks) as separate columns. **Figure 20** shows the arrangement of data in our example. Treatment description is in column A. Yield values for Blocks 1, 2, 3, and 4 are in columns B, C, D, and E, respectively.



	A	B	C	D	E	F	G	H
1				Blocks				
2		1	2	3	4			
3	Whole field	137.8	158.5	165.9	146			
4	Site specific	147.1	153.3	170.5	161.4			
5	Integrated	143.4	147.9	157.9	166.5			
6								
7								
8								
9								
10								
11								

Figure 20. Arrangement of example data in the spreadsheet labeled “data”.

Student Tasks

Enter the data as shown in Figure 20. After entering data in a cell, use the arrow keys to navigate to the next cell.

Tip

If the column width is too small to show the entire entry in a cell, position the mouse in the gray area at the top of the spreadsheet at the dividing line between two columns. The cursor will change to a double headed arrow \leftrightarrow . Hold down on the left mouse button while moving the cursor to the right (**Figure 21**). A box will appear showing the column width. Move the cursor to the right until the desired width is obtained. This action changes the width of the column to the left of the cursor.

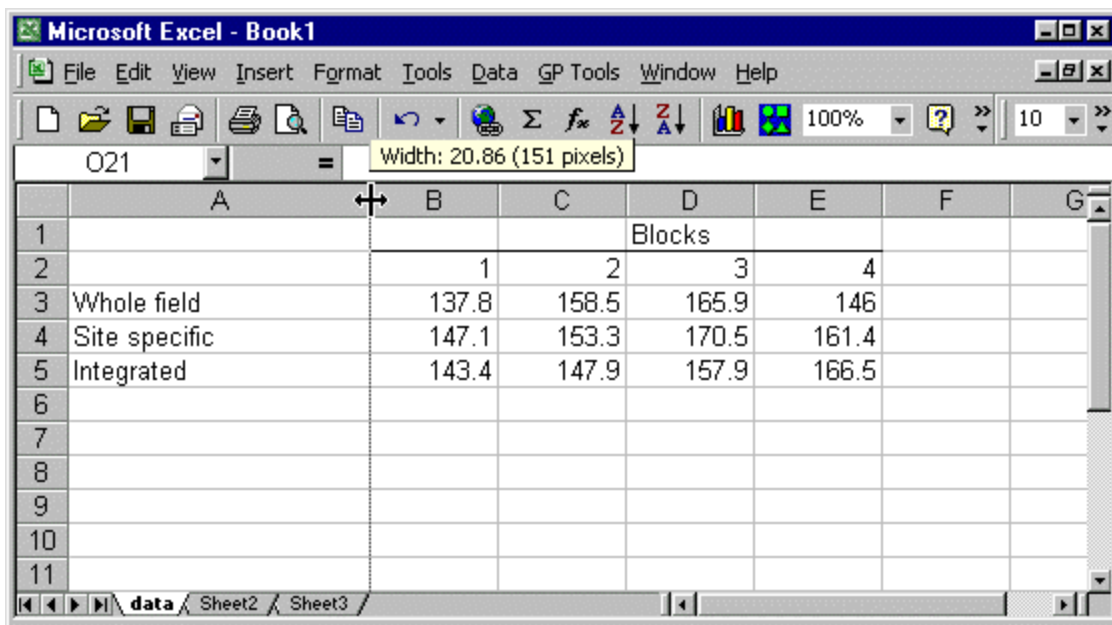



Figure 21. Widening a column to accommodate longer cell entries.


CREATING A SCATTER PLOT

Before conducting a statistical analysis, it is always a good idea to take a good look at your data. There is an old adage in statistics: “garbage in, garbage out”, meaning that for an analysis to be meaningful, quality data must go into it. One way to look over your data is to put all of the individual observations on a graph. This type of graph is called a scatter graph.

Student Tasks

Create a scatter graph of the example data.

1. First select the data you want to plot. In our example, this will be the block of data from cell A3 (top left corner of block) to E5 (bottom right corner of block). Click and drag from A3 to A5. The selected cells will now be shaded (**Figure 22**).
2. Click on the icon for the Plot Wizard on the Standard Toolbar at the top of the page . This will bring up the Chart Wizard window entitled Chart Wizard Step 1 of 4 – Chart Type (**Figure 23**). This box has two tabs at the top labeled Standard Types and Custom Types. Click the **Standard Types** tab if it is not already selected. On the left side of this box, there is a list of options under the title **Chart type**. Select **XY (scatter)** as your chart type (**Figure 24**). This will create a selection of plot types to choose from on the right side, labeled **Chart sub-type**.

Click on the topmost option , which looks like a graph with points only, no lines. After you have made this selection, click **Next**.

3. After clicking Next, a new window appears entitled Chart Wizard Step 2 of 4 – Chart Source Data (**Figure 25**). This box has two tabs labeled Data Range and Series. Click on the Data Range tab if it is not already selected. To the left of the label **Data range:**, verify that the selected cells equal “=data!\$A\$3:\$E\$5”. To the right of the **Series in:** label are two choices, Rows or Columns. Select **Columns**, as shown in Figure 24 then hit the **Next** button.
4. After clicking Next, a new window appears, entitled Chart Wizard – Step 3 of 4 – Chart Options (**Figure 26**). This box has 5 tabs labeled Titles, Axes, Gridlines, Legend, and Data Labels. Click on the Titles tab, if it is not already selected.
 - a. Under the label **Chart Title** type **Scatter plot of raw data**.
 - b. Under the **Value (X) axis:** label type **Treatment**
 - c. Under the **Value (Y) axis:** label type **Yield (bu/acre)**
 - d. Click the **Next** button.
5. After clicking Next, a new window appears, entitled Chart Wizard – Step 4 of 4 – Chart Location (**Figure 27**). Click the blank circle to the left of the label **As new sheet**. This will highlight the field to the right of the label. Type in **Scatter graph**. This will create a new worksheet named Scatter graph, which will contain the newly created scatter plot. Click the **Finish** button.

	A	B	C	D	E
1				Blocks	
2		1	2	3	4
3	Whole field	137.8	158.5	165.9	146
4	Site specific	147.1	153.3	170.5	161.4
5	Integrated	143.4	147.9	157.9	166.5

Figure 22. Cells selected for creating a scatter graph

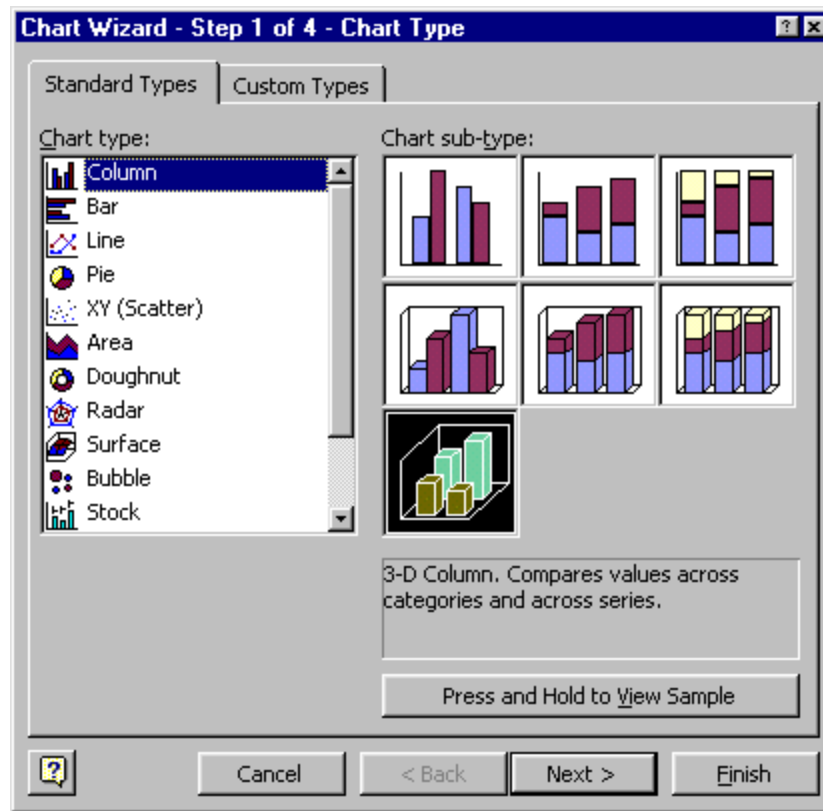


Figure 23. Chart Wizard window: **Step 1 of 4 – Chart Type.**

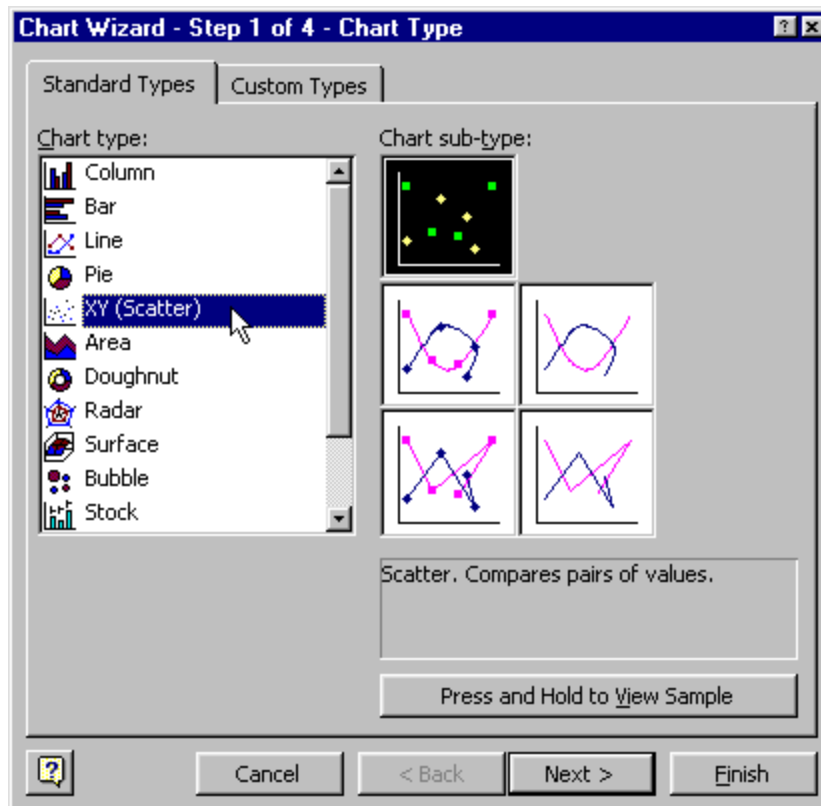


Figure 24. Selecting the **XY (scatter)** chart type.

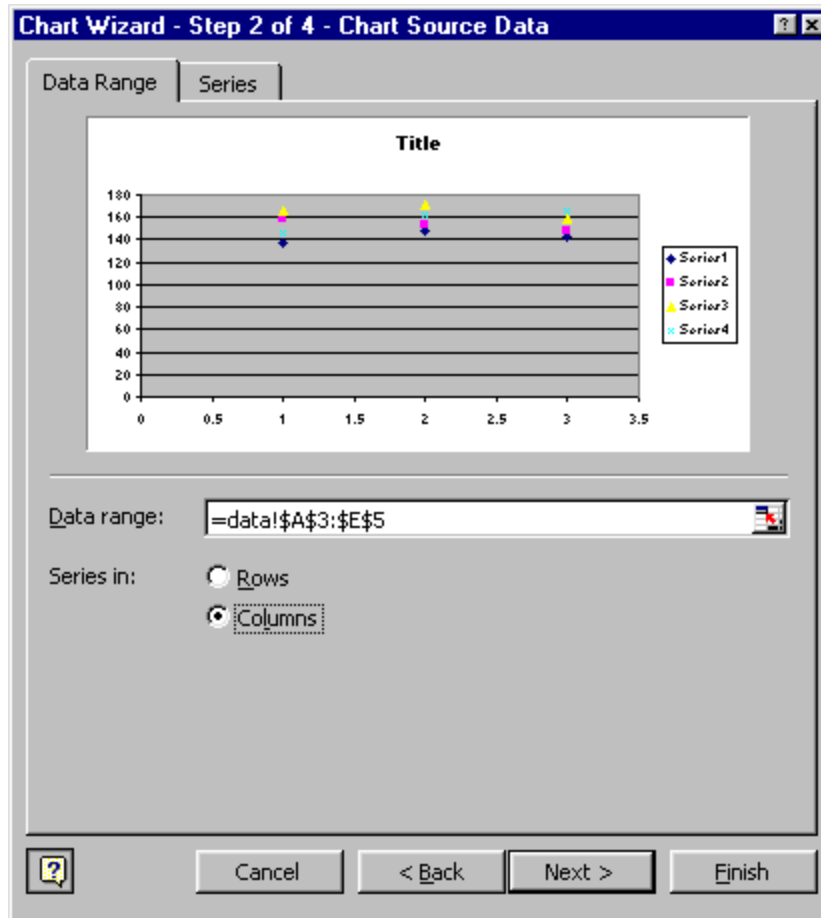


Figure 25. Chart Wizard window: **Step 2 of 4 – Chart Source Data.**

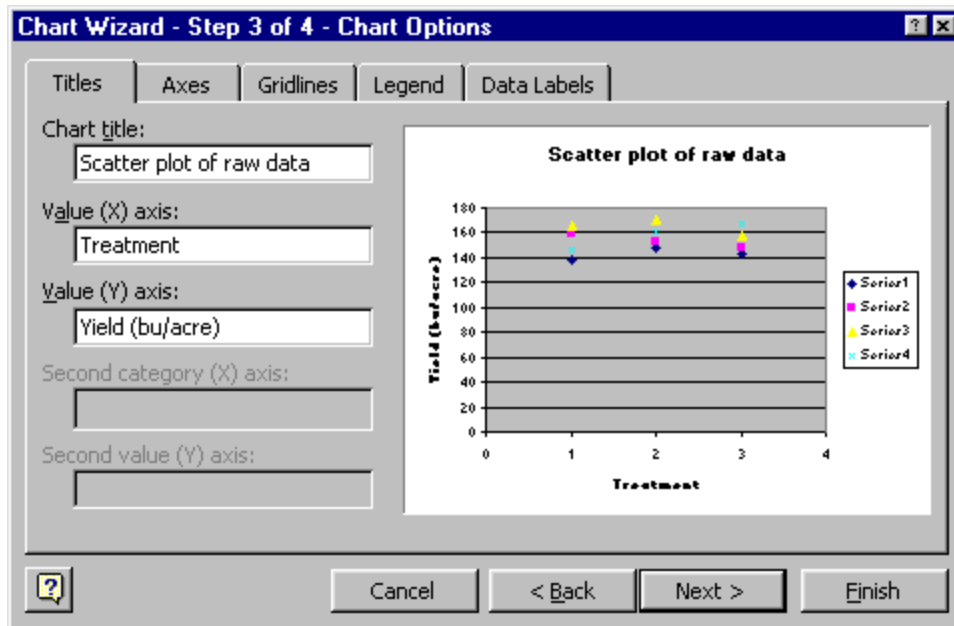


Figure 26. Chart Wizard window: **Step 3 of 4 – Chart Options**.

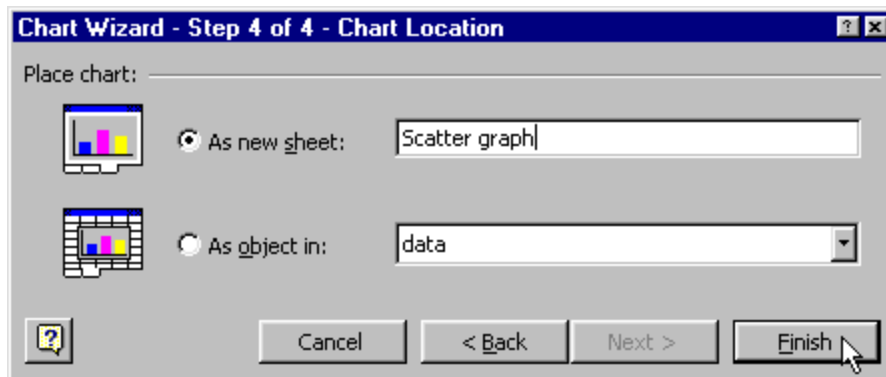


Figure 27. Chart Wizard window: **Step 4 of 4 – Chart Location**.

The scatter plot created is in **Figure 28**. This graph has two axes: the horizontal or x-axis and the vertical or y-axis. The x-axis delineates the treatments. The y-axis shows yields from all replicates for a particular treatment. The software creating the scatter plot does not recognize the treatment names in column A in the spreadsheet. Instead, it uses the numbers 1, 2, and 3 as the labels for the whole-field, site-specific, and integrated treatments, respectively. Each of the 4 series in the graph shows the yields for each treatment in a particular block. For instance, series 1 shows the yields for the whole field, site-specific, and integrated treatments for block 1.

Spend some time looking at the data in the scatter plot. Examining this graph will give you some insight into the variability in your experiment. For instance, look at how the treatments performed in each block. In blocks 1 and 3, (series 1 and 3, respectively), the highest yields came

from the site-specific treatment. In block 2 (series 2), whole-field management produced the best yield while in block 4 (series 4), the integrated treatment was best.

Looking at the scatter plot also provides some general indications of how each treatment performed. Visually, it appears that the site-specific treatment may have performed slightly better than the other treatments. However, there is quite a bit of variability in the yields produced by each treatment. In the next section, we will provide some descriptive statistics for each treatment.




Figure 28. Scatter plot of the example data.

NUMERICAL DESCRIPTIVE STATISTICS FOR EACH TREATMENT

There are many numerical descriptive statistics we can use to explore our data. However, in this workshop, we will concern ourselves with only two: 1) mean and 2) standard deviation.

Mean

The mean is the sum of all values divided by the number of values. For instance, the mean of the whole-field treatment is found by adding up all of the yields associated with this treatment from all 4 blocks:

$$\text{Whole field treatment mean yield} = \frac{137.8 + 158.5 + 165.9 + 146}{4} = 152.05 \text{ bu/acre.}$$

The mean provides us with some measure of the central tendency for the treatment. It's the way we would answer the question, "What was the yield of your whole-field management treatment?" To which we would respond, "It averaged about 152 bu/acre." This simple statement says a lot about how large the yields tended to be for that treatment.


Standard Deviation


The scatter plot in Figure 28 shows that while the mean of the whole-field treatment was 152.05 bu/acre, all of the measured yields were either higher and lower than that. In other words, there was variability in the yields for that treatment. The next question, then, is "How much variability was there?" A way to quantify the variability is to measure how far each yield was from the mean yield for the treatment. Yields closer to the mean are less variable than yields farther from the mean. What we need is one number that will summarize how large this variability is for all of the yields in the treatment.

Standard deviation is just such a summary number. Its calculation considers how far all of the values are from the mean. The larger the standard deviation, the greater the variability for a given mean. It's also a useful statistic because its units are the same as the units of the mean, in this case bu/acre. The standard deviation is often reported behind the mean after plus and minus signs. For instance, the standard deviation for the whole-field treatment is 10.9 bu/acre. Reporting the mean of the treatment with the standard deviation is denoted 152.05 ± 10.9 bu/acre. This is a more meaningful way of reporting means, because it also provides a measure of the variability.

Student Tasks

Calculate the mean and standard deviations of each treatment.

1. Click on the tab for the data worksheet.
2. In cell G2, type "Mean" and in cell H2 type "Std Dev." (**Figure 29**).
3. In cell G3 type "=`AVERAGE(B3:E3)`" (**Figure 29**).
4. In cell H3 type "=`STDEV(B3:E3)`" (**Figure 30**).
5. Click and drag from cell G3 to cell H3 to select these cells (**Figure 31**).
6. On the standard toolbar, click on the  icon to copy the contents of these cells. This will create a moving border around the selected cells (**Figure 32**).
7. Click and drag from cell G4 to cell H5 to select these cells (**Figure 33**).

8. On the standard toolbar, click on the  icon to paste the formulas into the remaining cells.
9. The results of this procedure should look like **Figure 34**.

	C	D	E	F	G	H
1		Blocks				
2	2	3	4		Mean	Std Dev
3	158.5	165.9	146		=AVERAGE(B3:E3)	
4	153.3	170.5	161.4			
5	147.9	157.9	166.5			

Figure 29. Using the AVERAGE function in Excel to calculate a mean.

Mean	Std Dev
152.05	=STDEV(B3:E3)

Figure 30. Using the STDEV function in Excel to calculate a standard deviation.

	C	D	E	F	G	H
1		Blocks				
2	2	3	4		Mean	Std Dev
3	158.5	165.9	146		152.05	10.87532
4	153.3	170.5	161.4			
5	147.9	157.9	166.5			

Figure 31. Selecting the cells to be copied.

Mean	Std Dev
152.05	12.55773

Figure 32. Moving cell borders created after copying cell contents.

Mean	Std Dev
152.05	12.55773

Figure 33. Selecting the cells into which to paste the contents of copied cells.

Mean	Std Dev
152.05	12.55773
158.075	10.14376
153.925	10.3442


Figure 34. Calculated means and standard deviations for all three treatments.


CREATING A BAR CHART OF TREATMENT MEANS

An important task in analyzing and summarizing data from an experiment is to create a graph of the treatment means. This provides an effective way of quickly grasping overall differences in magnitude among effects of the treatments. To do this, we will be using a bar chart. This type of chart is appropriate for the types of treatments to which we have limited ourselves in this workshop. We have also seen the importance of providing some measure of the variability in the experiment, using standard deviation. Including the standard deviation of each treatment in the graph will be accomplished by creating error bars. Error bars are lines drawn from the top of each bar and extend upward by one standard deviation and downward by one standard deviation. Longer error bars are associated with greater variability. Use of error bars is thus very powerful, providing a quick visual assessment of variability.

Student Tasks






Create a bar chart of the treatment means and standard deviations.


1. Click and drag from cell G3 to cell G5 to highlight the 3 cells containing the treatment means (**Figure 36**).
2. Click on the  icon for the Chart Wizard on the Standard Toolbar at the top of the page. This will bring up the Chart Wizard dialog box entitled Chart Wizard Step 1 of 4 – Chart Type (**Figure 37**). This box has two tabs at the top labeled Standard Types and Custom Types. Click the **Standard Types** tab if it is not already selected. On the left side of this box, there is a list of options under the title **Chart type**. Select **Column** as your chart type (Figure 37). This will create a selection of plot types to choose from on the right side, labeled **Chart sub-type**.

Click on the  option in the top left corner, which looks like a graph with two pairs of two bars. After you have made this selection, click **Next**.

3. After clicking Next, a new window appears entitled Chart Wizard Step 2 of 4 – Chart Source Data (**Figure 38**). This box has two tabs labeled Data Range and Series. Click on the Data Range tab if it is not already selected. To the left of the label **Data range:**, verify that the selected cells equal “=data!\$G\$3:\$G\$5”. To the right of the **Series in:** label are two choices, Rows or Columns. Select **Columns**,

as shown in Figure 38.

4. Click on the Series tab. To the right of the label **Category (X) axis labels:** at the bottom of the dialog box, click on the  icon on the right side of the blank box (**Figure 39**). This will shrink the size of the dialog box.
5. Click and drag from cell A3 to cell A5. This will enter “=data!\$A\$3:\$A\$5” in the blank box (**Figure 40**). Click on the  icon on the right side of the blank box. This will return the dialog box to its original size (**Figure 41**). The cells you selected are now entered as Category (X) axis labels. Click **Next**.
6. After clicking Next, a new window appears, entitled Chart Wizard – Step 3 of 4 – Chart Options (**Figure 42**). This box has 6 tabs labeled Titles, Axes, Gridlines, Legend, Data Labels, and Data Table. Click on the Titles tab, if it is not already selected.
 - a. Under the label **Chart Title** type **Treatment means**
 - b. Under the **Value (X) axis:** label type **Treatment**
 - c. Under the **Value (Y) axis:** label type **Mean yield (bu/acre)**
 - d. The data entered should look like that in **Figure 43**. If so, click the **Next** button.
7. After clicking Next, a new window appears, entitled Chart Wizard – Step 4 of 4 – Chart Location (**Figure 44**). Click the blank circle to the left of the label **As new sheet**. This will highlight the field to the right of the label. Type in **Bar chart**. This will create a new worksheet named Bar chart, which will contain the newly created bar graph of the treatment means. Click the **Finish** button.
8. Click on the Bar chart tab to view the newly-created spreadsheet (**Figure 45**).
9. Move the cursor to one of the bars and click the RIGHT mouse button. This will bring up a menu. Click on **Format Data Series** (**Figure 46**). This brings up a dialog box with 6 tabs labeled Patterns, Axis, Y Error Bars, Data Labels, Series Order, and Options (**Figure 47**).
10. Click on the Y Error Bars tab. At the bottom of the dialog box, click the radio button the left of the label Custom: (**Figure 48**). Move the cursor to the box to the right of the + label and click. Next, click on the  icon on the right side of the box. This will shrink the size of the dialog box. Click on the “data” spreadsheet tab to go back to the spreadsheet where standard deviation was calculated for each treatment. Click and drag from cell H3 to H5 (**Figure 49**). This will enter “=data!\$H\$3:\$H\$5” in the box (Figure 49). Click on the  icon, which will return the dialog box to its original size. The cells selected are now shown in the box to the right of the + label (**Figure 50**).
11. Move the cursor to the box to the right of the - label and click. Next, click on the  icon on the right side of the box. This will again shrink the size of the dialog box. Click on the “data” spreadsheet tab to go back to the spreadsheet where

standard deviation was calculated for each treatment. Click and drag from cell H3 to H5, just as before. This will enter “=data!\$H\$3:\$H\$5” in the box. Click on the  icon, which will return the dialog box to its original size. The cells selected are now shown in the box to the right of the - label (**Figure 51**). Click **OK**.

12. The bar chart now has error bars that show the standard deviation for each treatment mean (**Figure 52**).

G
Mean
152.05
158.075
153.925

Figure 36. Highlighted cells to be used to create a bar chart.

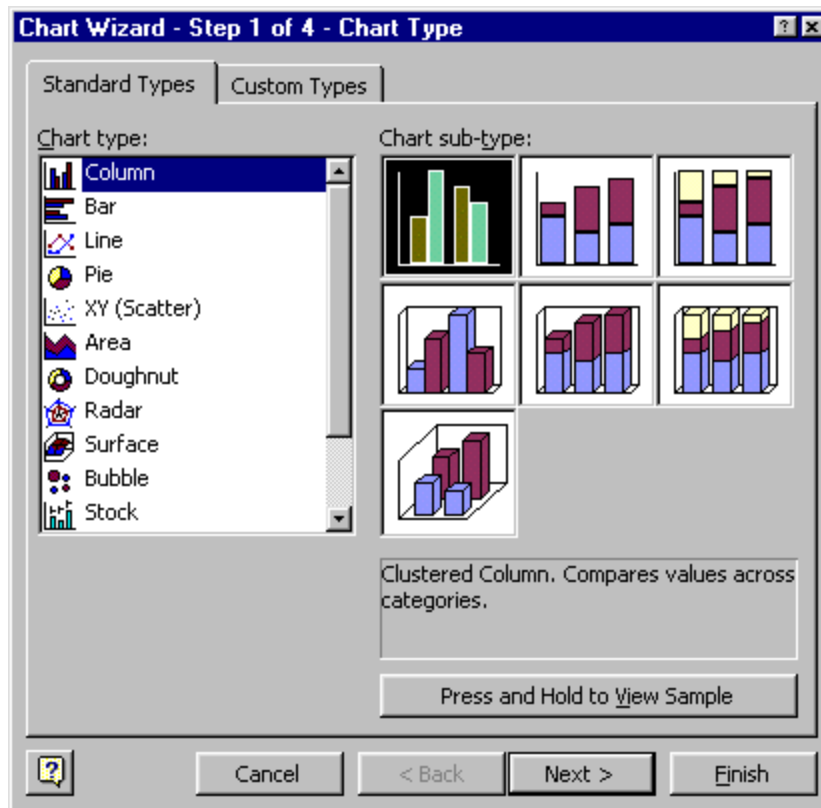


Figure 37. Chart Wizard – Step 1 of 4 – Chart Type.

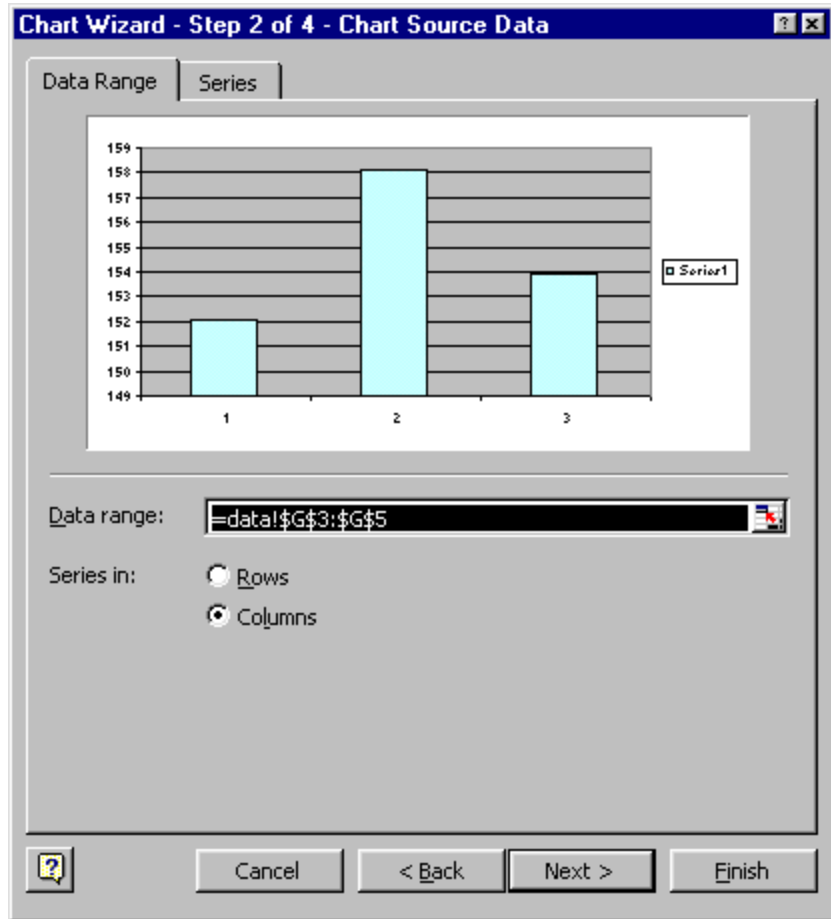


Figure 38. Chart Wizard – Step 2 of 4 – Chart Source Data.

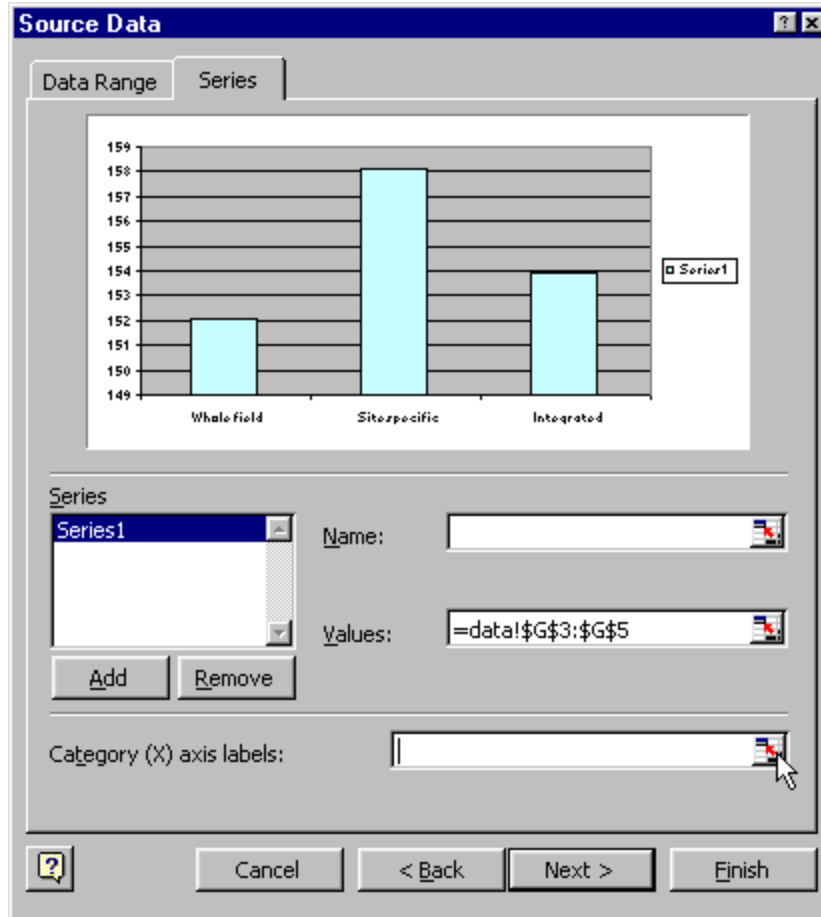


Figure 39. Contents of the Series tab.

	A	B	C	D	E	F
1				Blocks		
2		1	2	3	4	
3	Whole field	137.8	158.5	165.9	146	
4	Site specific	147.1	153.3	170.5	161.4	
5	Integrated	143.4	147.9	157.9	166.5	
6						
7	Chart Wizard - Step 2 of 4 - Chart Source Data - Category (X) ax...					
8	=data!\$A\$3:\$A\$5					
9						

Figure 40. Cells selected as the category (x) axis labels.

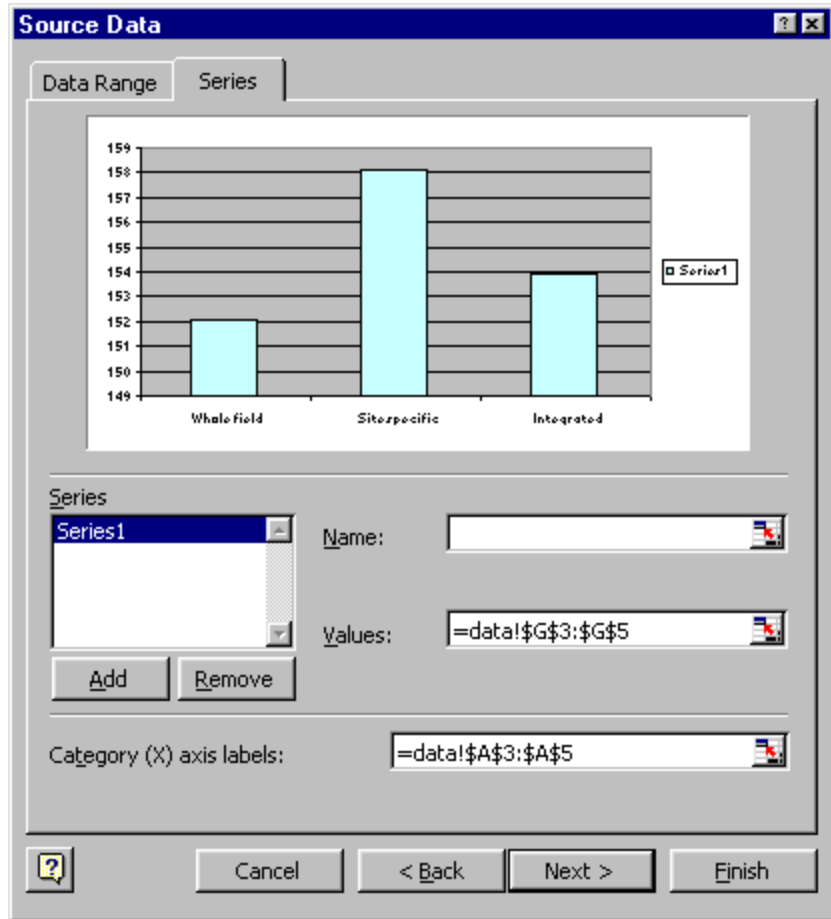


Figure 41. Series tab with data range entered for the category (X) axis labels.

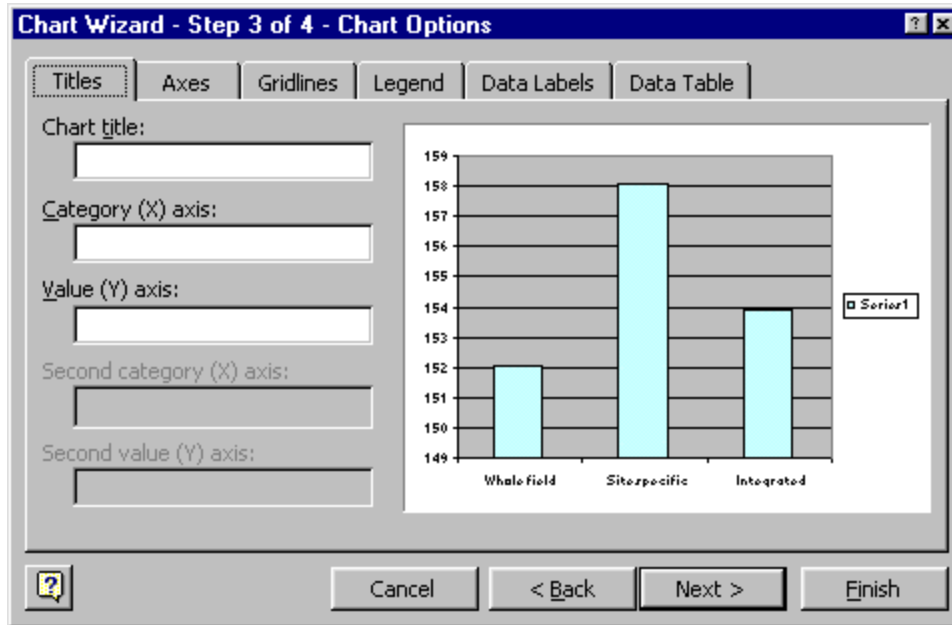


Figure 42. Chart Wizard – Step 3 of 4 – Chart Options.

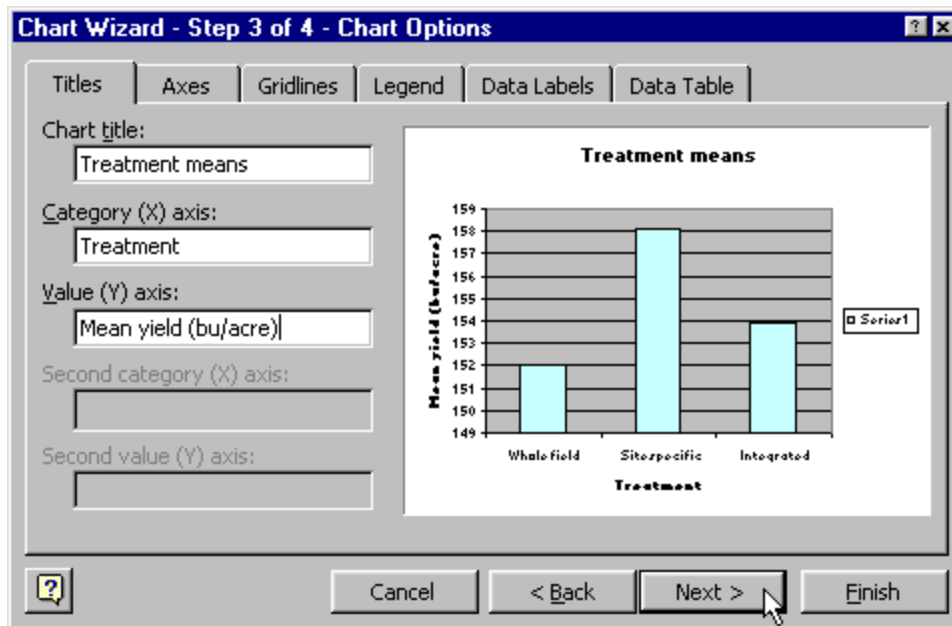


Figure 43. Data entered for Chart title, category (X) axis, and value (Y) axis.

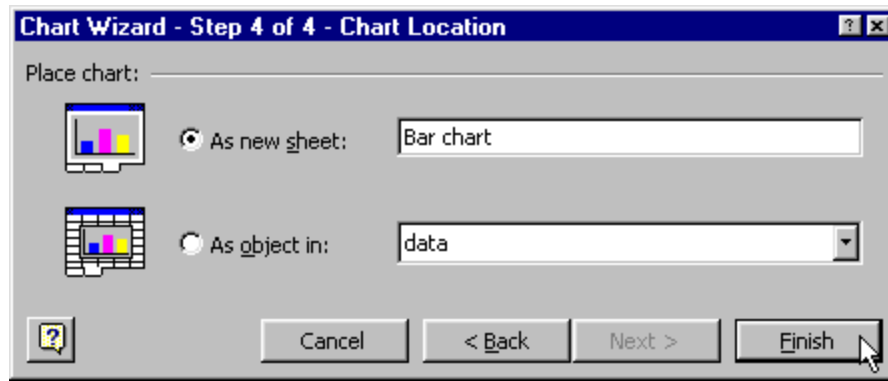


Figure 44. Chart Wizard – Step 4 of 4 – Chart Location.

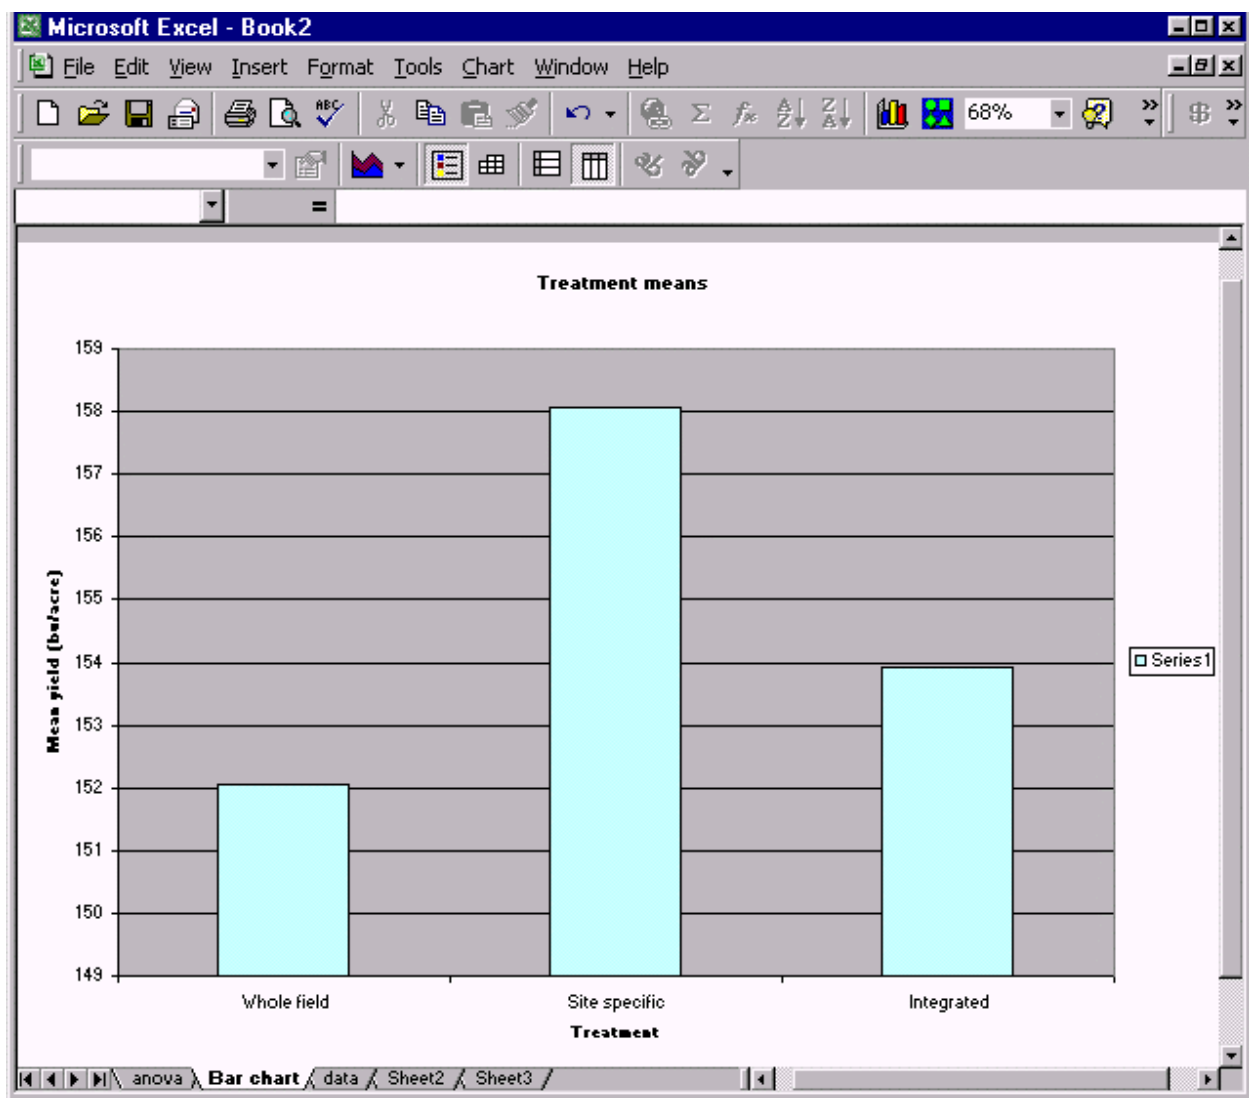


Figure 45. Newly created “Bar chart” spreadsheet.




Figure 46. Selecting **Format Data Series...** in the drop down menu appearing after right-clicking on a bar in the graph.

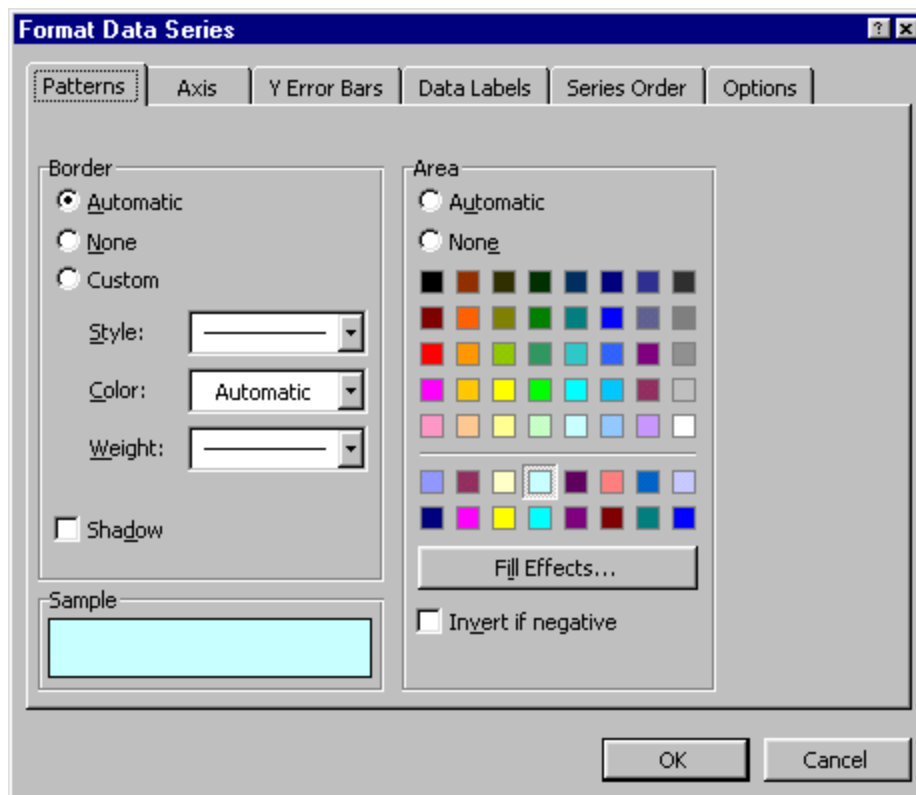


Figure 47. The Format Data Series dialog box.

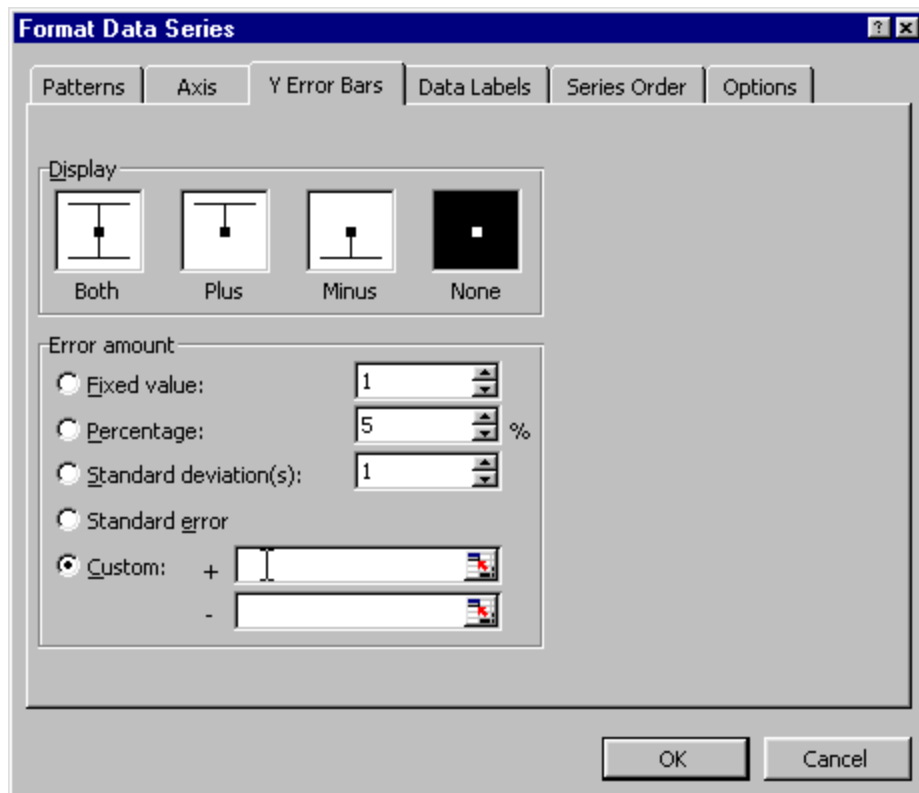


Figure 48. Selecting **Custom** for tailoring the error bars to the example data.

	A	B	C	D	E	F	G	H
1				Blocks				
2		1	2	3	4		Mean	Std Dev
3	Whole field	137.8	158.5	165.9	146		152.05	10.87532
4	Site specific	147.1	153.3	170.5	161.4		158.075	8.784752
5	Integrated	143.4	147.9	157.9	166.5		153.925	8.958341
6								
7								
8								
9								

Format Data Series - Custom +	
+	=data!\$H\$3:\$H\$5
-	

Figure 49. Selecting the data to be used for the positive part of the error bars.



Figure 50. Data range entered for the positive part of the error bars.

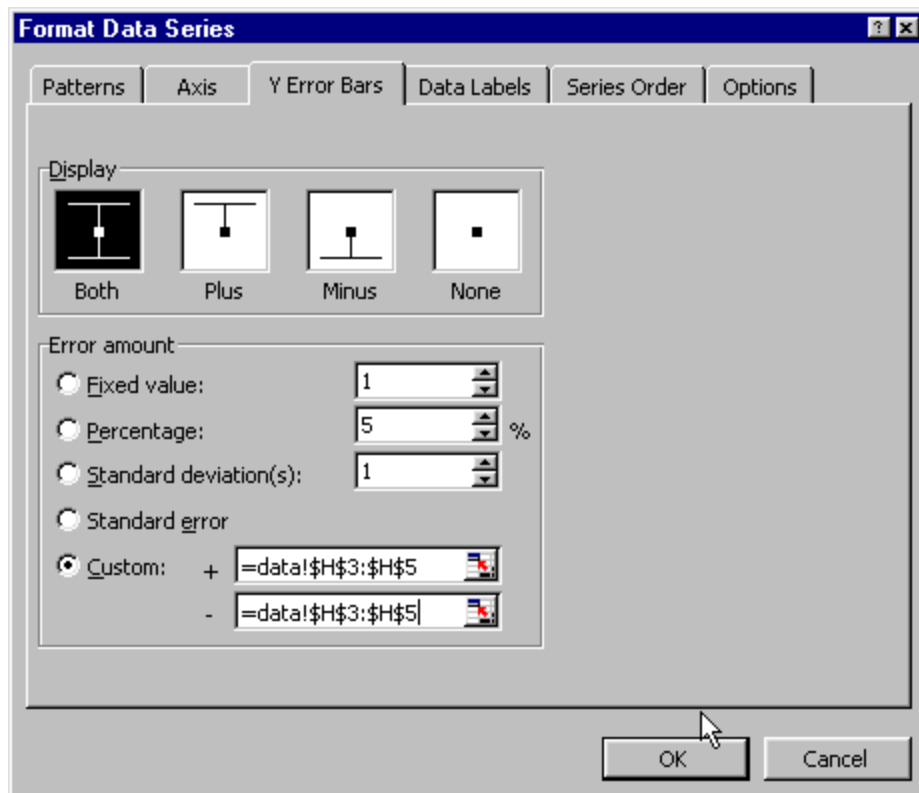


Figure 51. Data range entered for both the positive and negative parts of the error bars.




Figure 52. Revised bar chart showing error bars for each treatment.

PERFORMING AN ANALYSIS OF VARIANCE (ANOVA)

Understanding How an ANOVA Works

The scatter plot in Figure 28 and the bar chart shown in Figure 52 demonstrate together that there are 3 primary sources of variability in our experiment: 1) treatments, 2) blocks, and 3) the overall variability in the experiment, termed experimental error. A plot of the treatment means in Figure 52 shows that not all treatment means are equal; there is some variability among them. The mean yields of each block are also variable. They were calculated by averaging the yields of the three treatments in each one and were 142.8, 153.2, 164.8, and 158.0 for blocks 1, 2, 3, and 4, respectively. Figure 28 and its discussion in the previous section revealed that treatments produced different results, depending on the blocks they were in. The variability resulting from this interaction of treatments and blocks is termed experimental error. Therefore our RCBD design has three sources of variation: treatments, blocks, and the interaction of treatments and blocks, or experimental error.

It is important to remember that the magnitude of the variability we observe in our experiment is unique to our experiment. If we were to conduct the experiment again, we might find that the observed variability was different. If we conducted the experiment again and again, we might finally understand what the true variability of this experiment was. Obviously, we will never know the true variability. There simply aren't resources enough to keep conducting the same experiment over and over again. Therefore, the variability in our experiment (experimental variability) is only an estimate of what the true variability would be if we ran the experiment over and over, perhaps hundreds of times.

The variability in the experiment will determine whether or not we will be able to detect significant differences among our treatment means. So the logical question to ask is "How much variability is associated with this experiment?" We really want to know the true variability. Since we don't have resources to determine it, we use our experimental variability to estimate it. Each of the three sources of variability outlined above provides its own estimate of the unknown true variability. This means that we could estimate the true variability by looking only at the variability in the treatment means, or we could estimate variability from just the block means. Our last option is to estimate variability from the interaction of blocks and treatments (experimental error, shown graphically in our scatter plot, Figure 28). So which estimate is the best?

Our best estimate of the true variability for the experiment is the experimental error. The reason for this is that it considers the variability in each yield point. The other two estimates are based on means (either those of the treatments or the blocks), so they are more limited in their ability to estimate the true variance for the experiment. Therefore, the experimental error becomes the standard against which the variance estimates from treatment means or block means can be compared. For instance, if the variance in the treatment means provides a very different estimate of variance from that of the experimental error, then something significant has occurred. What is this significant event? Statistical theory provides the answer.

For the variability in treatment means to provide a good estimate of the true variance in the experiment, the treatment means must essentially all be the same, according to statistical theory. If the variability in treatment means does not provide a good estimate of the true variance, then at least one of the treatment means is different from the others. Statistical theory, then, provides the link between variability and statistical differences among treatment means.



What we need is a procedure that will compare the variance of the treatments to the variance of the experimental error to see if they are the same or different. Such a procedure is appropriately called an analysis of variance, or ANOVA. In the ANOVA procedure, comparing two variances is accomplished by the F-test. We will discuss more about the F-test after we have run the ANOVA procedure.

Our discussion has been limited to treatment means. Remember, however, that another estimate of the variance came from the variability in block means. The way that the ANOVA is constructed for our RCBD design, we do not have a valid F-test to decide whether to accept or reject a null hypothesis based on equality of block means. Although an F-test will be provided in the output of the ANOVA procedure, it is not a valid test for our conditions.

Running the Analysis of Variance Procedure

Student Tasks

Perform an analysis of variance for the example data set.

1. Before beginning the analysis of variance, be sure that your cursor is somewhere on the spreadsheet labeled “data”.
2. On the main menu at the top, select **Tools**. From the drop down menu, click **Data Analysis...** (**Figure 53**).
3. In the dialog box titled Data Analysis, click **Anova: Two-Factor Without Replication** and click **OK** (**Figure 54**).
4. In the dialog box titled Anova: Two-Factor Without Replication, click on the  icon on the right side of the white box located to the right of the label **Input Range:**. The dialog box will be reduced.
5. Click and drag from cell A2 to cell E5 to select the data to be included in the analysis. After doing this, the cell range will be shaded and will appear in the dialog box as \$A\$2:\$E\$5 (**Figure 55**). Click on the  icon in the right side of the box. This will return the dialog box to its original size.
6. Because we selected the treatment names in column A, click on the checkbox to the left of **Labels** to put a check in the box (**Figure 56**). Clicking this option will produce treatment names in the output from the analysis.
7. To the right of the **Alpha** label, type in the alpha level, in decimal form. For our example, enter **0.05** (**Figure 56**).
8. Under Output options, click on the radio button to the left of the label **New Worksheet Ply:** (**Figure 56**). Click **OK**.
9. Click on the spreadsheet tab labeled “anova”. This is the new spreadsheet created by Excel that contains the ANOVA table (**Figure 57**).
10. Resize column A so that it is wide enough to show all of the text in the cells.

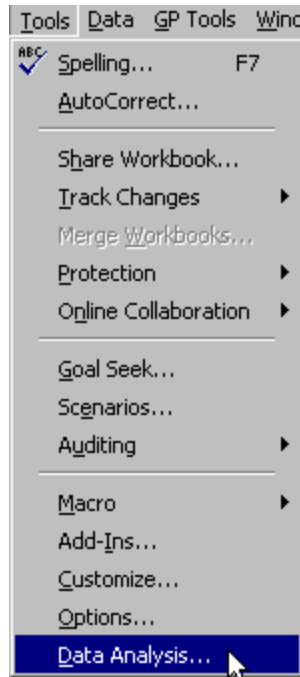


Figure 53. The **Tools** menu.

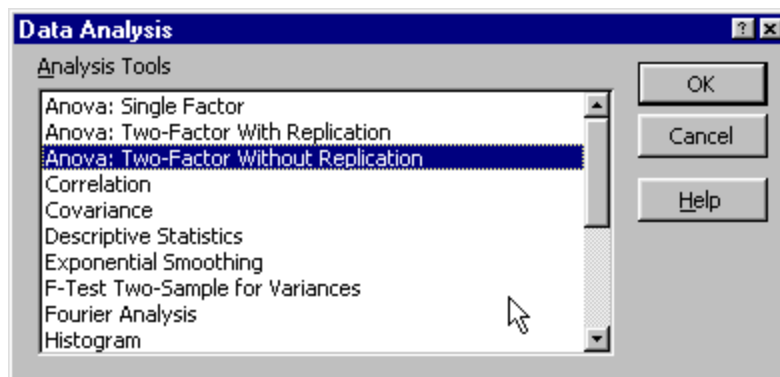


Figure 54. Selecting the **Anova: Two-Factor Without Replication** option.

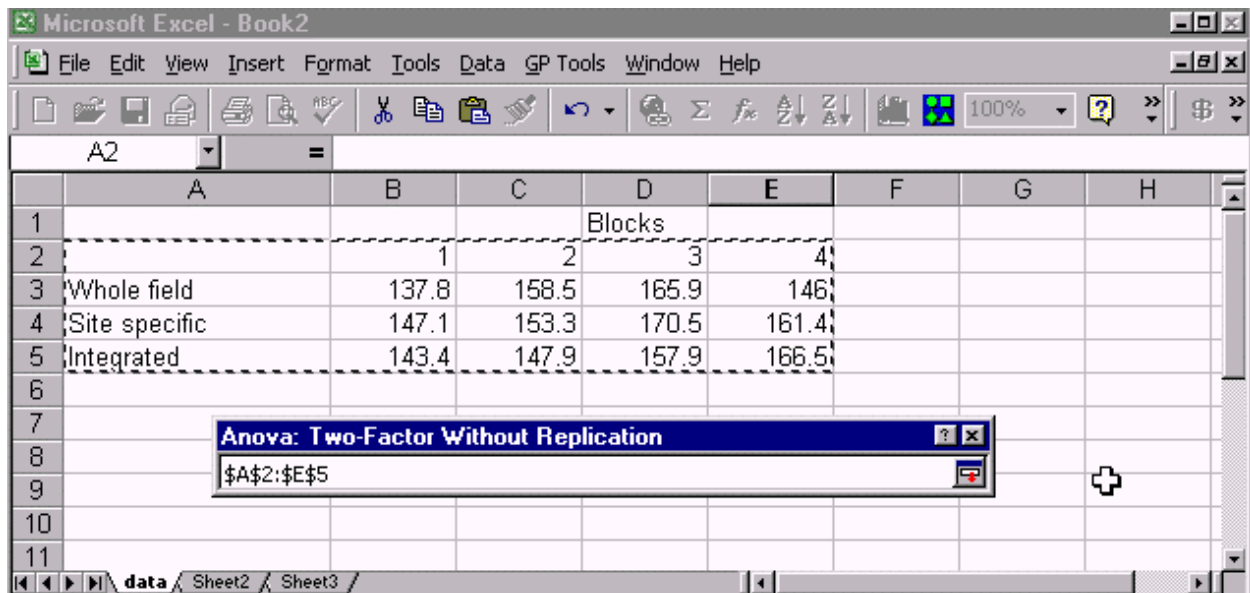


Figure 55. Selecting the cell block for analysis.

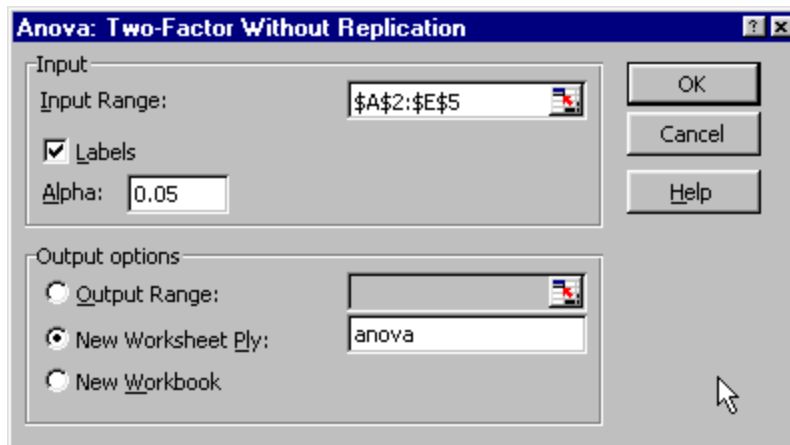


Figure 56. Selected options for **Anova: Two Factor Without Replication**.

	A	B	C	D	E	F	G
1	Anova: Two-Factor Without Replication						
2							
3	<i>SUMMARY</i>	<i>Count</i>	<i>Sum</i>	<i>Average</i>	<i>Variance</i>		
4	Whole field	4	608.2	152.05	157.6967		
5	Site specific	4	632.3	158.075	102.8958		
6	Integrated	4	615.7	153.925	107.0025		
7							
8		1	3	428.3	142.7667	21.92333	
9		2	3	459.7	153.2333	28.09333	
10		3	3	494.3	164.7667	40.65333	
11		4	3	473.9	157.9667	113.9033	
12							
13							
14	ANOVA						
15	<i>Source of Variation</i>	<i>SS</i>	<i>df</i>	<i>MS</i>	<i>F</i>	<i>P-value</i>	<i>F crit</i>
16	Rows	76.05167	2	38.02583	0.684955	0.539594	5.143249
17	Columns	769.69	3	256.5633	4.621444	0.052978	4.757055
18	Error	333.095	6	55.51583			
19							
20	Total	1178.837	11				
21							

Figure 57. Completed analysis of variance.

Interpreting Results from the Analysis of Variance

Figure 28 shows the output from the ANOVA procedure. It is important to have a general idea what various parts of the output mean. This helps you do some error checking to ensure that the analysis has been performed correctly. During the procedure, you may accidentally not select all of the data that you intended to analyze, or you may have selected the wrong data. There are many ways to make mistakes. Knowing how to make sense of the output is therefore critical to ensuring a proper analysis has been performed.

The output contains two tables: a summary table and an ANOVA table. The summary table is at the top of the spreadsheet we labeled “anova,” and the ANOVA table is underneath it.

Summary table

This table provides some descriptive statistics of the treatments and blocks.

Summary. The first column is labeled “Summary” and lists the names of our treatments (whole field, site specific, and integrated) and our blocks (1, 2, 3, and 4).

Count. The column labeled “Count” indicates how many observations went into the summary statistics on the right side of the table. The count for each treatment is the number of times it has been replicated. This equals the number of blocks containing each treatment, which is 4 in our experiment. The count for each block is the number of treatments contained in each block, which is 3.

Sum. The column labeled “Sum” contains calculations that add together all of the yields comprising each count for each treatment or block. For instance, the sum for the whole-field management is the yield obtained by adding together the yields associated with this treatment in each of the 4 blocks.

Average. This is the mean of each treatment or each block. Block means are calculated by averaging together the 3 treatments in each block. Treatment means are calculated the same way we discussed in the section “Numerical Descriptive Statistics for Each Treatment.” Check to be sure that the treatment means on the “anova” spreadsheet are the same as the means in the “data” spreadsheet. For example, the mean of the whole field treatment is 152.05 in the summary table in the “anova” spreadsheet. This is the same as the mean we calculated for this treatment in the “data” spreadsheet, cell G3 (Figure 31).

Variance. This is the variability associated with each treatment and block. Taking the square root of the variance calculates the standard deviation. Check to be sure that the square roots of the variances in the ANOVA table match the standard deviations we calculated in the section “Numerical Descriptive Statistics for Each Treatment.” For example, the variance of the whole field treatment in the ANOVA table is 157.6967. The square root of this variance is 12.55773, which equals the standard deviation we calculated for this treatment in cell H3 on the “data” spreadsheet (Figure 34).

ANOVA table

Source of Variation. In the first paragraph of the section “Understanding How an ANOVA Works,” three sources of variation were defined for our RCBD experiment: treatments, blocks, and experimental error. Remember how we arranged the data in our “data” spreadsheet prior to analysis. Treatments were the rows and blocks were the columns. Therefore, the sources of variation listed in the ANOVA table as Rows, Columns, and Error correspond to treatments, blocks, and experimental error, respectively.

SS. This is short for “sums of squares” and is part of the calculation for the variance.

df. This is an abbreviation for “degrees of freedom” and is concerned with restrictions on randomization during the design phase. For rows, this number should be one less than the number of treatments, or $3 - 1 = 2$. For columns, this number should be one less than the number of blocks, or $4 - 1 = 3$. The df for error is the multiplication of the df for rows and columns, or $2 \times 3 = 6$. Knowing how the degrees of freedom are calculated provides further verification for the correctness of the analysis.

MS. This is an abbreviation for “mean square” and is the variance associated with each source: treatments (rows), blocks (columns), and experimental error (error). The mean square for the experimental error is often referred to as the mean square error (MSE).

F. This is the result of an F-test comparing the variance of the treatment with the variance of the experimental error. For treatments (rows), it is the mean square for the treatment (38.02583) divided by the mean square error (55.51583), or $38.02583/55.51583 = 0.684955$. There is also an F-test given for blocks, but for a RCBD experiment, this test is not valid.

F crit. For the F-test to have meaning, its results must be interpreted. During the ANOVA procedure, we input an alpha level of 0.05. By entering this value, we decided we wanted to have only a 5% chance of stating that there was a significant difference among treatment means if in fact there was none. Providing the procedure with this information created a critical value for the F-test that had to be met or exceeded for us to find statistical significance at the alpha level we selected. The critical F level calculated for our treatments (rows) was 5.143249. We compare this to the F-test for our treatments (rows), 0.684955. Since 0.684955 is less than 5.143249, we conclude that there was no significant difference among treatments at an alpha level of 0.05. The critical level provided for blocks (columns) is not used since the F-test is not valid.

P-value. The P-value indicates the probability of making a type I error if the result of the F-test in the experiment is used as a critical value. It therefore provides the alpha level for rejecting the null hypothesis, calculated from the results of the experiment. In our example, if we chose to state that there was a significant difference among the treatment means, we would do so with a 53.9594% chance of drawing the wrong conclusion. As before, although a p-value is provided for the blocks, this test is not valid for our RCBD experiment.

LEAST SIGNIFICANT DIFFERENCE

Understanding the Least Significant Difference

When the F-test for treatments in the ANOVA turns out to be significant, you can conclude that one or more pairs of treatments are different. However, the ANOVA does not tell you specifically which ones are different and how they differ. Another statistical method called a mean comparison procedure is used to answer these questions.

The least significant difference (LSD) is probably the most often used mean comparison procedure for interpreting agronomic research. It is easy to calculate and because only one value is required, easy to apply. The LSD represents the smallest difference between two means that can be considered significantly different at a given probability level. If the difference between two treatment means exceeds the LSD then they are considered to be different. Otherwise, they are considered the same.

The LSD often is used for variety trials and other experiments where a large number of qualitative treatments are compared. Qualitative treatments are those for which a numerical description has no real meaning, such as seed color, hybrids, or fertilizer type. Other mean comparison procedures are more appropriate for quantitative treatment factors like fertilizer and herbicide application rates.

Calculating the Least Significant Difference

To calculate the LSD, there are two primary steps: 1) calculating a multiplicative coefficient, called a critical t value and 2) using the critical t value to calculate the LSD.

To calculate the critical t value, we will use a statistical function in excel. The function name is TINV. The syntax is: =TINV(alpha level, df of the mean square error). The alpha level we specify controls the probability of making a type I error when comparing two means. This is called the comparisonwise error rate. In our example, we will choose an alpha level of 0.05 for the same reasons we gave for the ANOVA analysis. The second input needed in the TINV function is the degrees of freedom of the mean square error. In our experiment, this was 6 (ANOVA table, df associated with the error source of variation).

Student Tasks

Calculate the critical t value

1. If the “anova” spreadsheet is not currently selected, click on the spreadsheet tab labeled “anova” to view it.
2. In cell A23, input the following: “Critical t value”.
3. In cell B23, input the following: “=TINV(0.05,6)” (**Figure 58**)



Figure 58. Entering the TINV function into a cell.

The critical t value we just calculated will now be used to calculate the LSD. In this calculation, we will need to find a square root. The function used in Excel to calculate a square root is named SQRT. Its syntax is: =SQRT(number). The number can be a constant or an equation. To calculate the LSD, we will use a formula with the following syntax:

$$=\text{Critical t value}*\text{SQRT}(2*\text{mean square error}/\text{number of blocks})$$

The critical t value we calculated above is in cell B23. The mean square error in the ANOVA table is in cell D18. The number of blocks in our experiment is 4. The “*” is the symbol Excel

uses for multiplication, while the “/” symbol is used for division. Therefore the equation we will want to enter into the excel spreadsheet to calculate the LSD for our example is

$$=B23*SQRT(2*D18/4)$$

Student Tasks

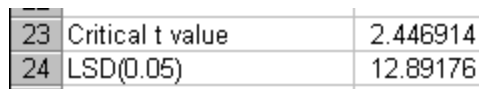
Calculate the LSD

1. If the “anova” spreadsheet is not currently selected, click on the spreadsheet tab labeled “anova” to view it.
2. In cell A24, input the following: “LSD(0.05)”.
3. In cell B24, input the following: “=B23*SQRT(2*D18/4)” (**Figure 59**)
4. The completed spreadsheet should look like **Figure60**.



23	Critical t value	2.446914
24	LSD(0.05)	=B23*SQRT(2*D18/4)

Figure 59. Entering the formula for the LSD into a cell.



23	Critical t value	2.446914
24	LSD(0.05)	12.89176

Figure 60. Completed calculations for the LSD.

To use the LSD we calculated (12.89176 bu/acre), examine the two means you want to compare. If their difference is greater than 12.89176 bu/acre, then they are significantly different.

Controlling Error in Using the Least Significant Difference

There is a potential problem with using the LSD that you should be aware of. The actual probability of making a Type I error increases with the more comparisons you make. So if you compare all possible pairs of treatments in an experiment your probability of making a Type I error is likely to be higher than you intended. Therefore, it is wise to limit your use of the LSD to those comparisons that are of particular interest. Another way to lessen the risk is to conduct an LSD test only if the ANOVA F-test for treatments is significant. This is sometimes called a protected F-test and is a good practice to follow with experiments where there is a good chance that the null hypothesis is true.

SELECTED REFERENCES

- Freund, R.J. and W.J. Wilson. 1997. Statistical methods. Rev. Ed. Academic Press, San Diego, CA.
- Shapiro, C.A., W.L. Kranz, and A.M. Parkhurst. 1989. Comparison of harvest techniques for corn field demonstrations. *Am. J. Agr. 4*:59-64.
- Steel, R.G.D. and J.H. Torrie. 1980. Principles and procedures of statistics: A biometric approach. McGraw-Hill, New York, NY.
- Wuest, S.B., B.C. Miller, J.R. Alldredge, S.O. Guy, R.S. Karow, R.J. Veseth, and D.J. Wysocki. 1994. Increasing plot length reduces experimental error of on-farm tests. *J. Prod. Agric. 7*:211-215.

APPENDIX

Table A1. Partial listing of websites providing electronic access to sound technical information.

Technical resources and websites	Free items (A = abstract) (F = full article) (R = research report) (T = titles)
Agricola (searches agricultural journals) www.nal.usda.gov/ag98/ag98.html	A, T
American Society of Agronomy www.soils.org	A, T
Crop Science Society www.crops.org	A, T
Iowa State University Extension www.exnet.iastate.edu/pubs/	F, T
Kluwer Academic Publishers www.wkap.nl/kaphtml.htm/HOME PAGE	A, T
Michigan State University Extension www.msue.msu.edu/	F, T
North Dakota State University Extension www.ag.ndsu.nodak.edu/farmranch.htm	F, T
Ohio State University Extension www.ag.ohio-state.edu/	F, T
Potash & Phosphate Institute www.ppi-ppic.org	F, R, T
Purdue University Extension persephone.agcom.purdue.edu/AgCom/Pubs/	F, T
Soil Science Society of America www.soils.org	A, T
South Dakota State University Extension www.abs.sdstate.edu/plantsci/ext/plantsci.htm	F, T

Springer-Verlag link.springer.de/forum.htm	A, T
United States Department of Agriculture – Agricultural Research Service www.ars.usda.gov	F, R, T
University of Illinois Extension www.extension.uiuc.edu/ag/ag-resources.html	F, T
University of Minnesota Extension www.extension.umn.edu/crops/	F, T
University of Nebraska Extension www.ianr.unl.edu/pubs/browse.htm	F, T
University of Wisconsin Extension www1.uwex.edu/topics/new/Agriculture_Publications.cfm	F, T
