

A Practical Guide to On-Farm Research

R.L. (Bob) Nielsen
Agronomy Department
Purdue Univ., West Lafayette, IN
ph. 765-494-4802, Email: rnielsen@purdue.edu

The notion of conducting on-farm research is not new. Many farmers have routinely conducted their own on-farm trials for years. In some states, farmers and land-grant universities have established formal partnerships to organize and conduct on-farm research trials. Today's GPS-enabled farming technologies can greatly simplify the logistics of conducting on-farm trials.

Why Conduct Field Research?

The purpose of conducting field crop research is to come up with fact-based answers to farming's challenging questions for which no previous answers exist. Effects of experimental treatments or variables on crop yield or other important outcomes are evaluated under controlled conditions and then those results are used to predict their future performance across the broader extent of agricultural production. On-farm research (OFR) not only seeks to identify answers to important questions but may also serve to validate previously discovered answers or convince growers that an alternative crop management practice is profitable for their own situations.

There are several potential barriers, or at least "speed bumps", for those who participate in OFR. One consideration is the time involvement required to successfully conduct such trials. It is easy to commit to a project in the winter and often more difficult to carry through with the task the following growing season. The logistics of conducting OFR trials can be burdensome or even prohibitive to growers. The good news today is that OFR is more feasible than ever before due to the availability of GPS-enabled crop production technologies including light-bar navigation, auto-steering, variable-rate controllers, and yield monitors.

Size and Scope of Field Research

Traditional small plot research is conducted on small uniform experimental areas that minimize the background "noise" that often plagues field research. Theoretically, small plot research enhances the researcher's ability to detect true and repeatable differences among the experimental treatments. Small plot research enables researchers to evaluate many treatments in a small area of land and thus minimizes the land resources required for field plot research. The small plot sizes often require specialized or small-scale research plot equipment.

On-farm research targets “real world” fields that, by virtue of their larger size, are typically more variable than smaller fields used for small-plot research. The greater within-field variability introduces a lot of background “noise” that can mask true differences in the measured responses between treatments. On-farm research allows for the use of commercial-scale field equipment and yield monitoring, but because individual plot size with OFR is larger (often equipment width by length of field), the number of treatments that can be evaluated per acre of land is fewer than with small-plot research.

It is worth noting that on-farm demonstrations are not synonymous with on-farm research. The purpose of demonstrations is not to identify or validate answers to research questions, but rather to simply gain experience with new technology or cultural practices. Sometimes demonstrations are designed purely to expose growers and others to new technology or cultural practices. Because on-farm demonstrations are not the same as research, yield responses or other data need not be measured or analyzed.

Background “Noise”

As stated at the outset, the purpose of conducting field crop research is to come up with fact-based answers to farming’s challenging questions for which no previous answers existed. It is important to recognize, though, that research is not simply about documenting history (e.g., my hybrid beat your hybrid last year in my variety trial by 10 bu/ac). Rather, it is about predicting future responses (e.g., based on my research, it is statistically probable that my hybrid will beat your hybrid every time they are planted in the same field).

Historical performance does not necessarily predict future performance because field research is always plagued by the confounding effects of background “noise” that tend to camouflage the effects of the treatments being evaluated. In other words, yield differences between treatments in a field experiment may simply be due to the background “noise” of the experiment. This background “noise” is also referred to as “experimental error”.

Background “noise” consists of variability among plots in your trial that is due to other, uncontrolled, often unknown yield influencing factors. Examples of background “noise” include human error in conducting the trial (i.e., that “loose nut behind the wheel”), variable soil characteristics within a field (soil texture, drainage, compaction, elevation), within-field variability for stresses (pest damage, herbicide injury, weather, etc.), and finally weather variability year to year, especially as it interacts with other yield influencing factors.

One of the challenges with OFR is to sort out the true effects caused by the treatments from those effects caused by “background noise”. You can never be 100% certain that measured yield differences in a trial are solely due to the treatments being evaluated. Fortunately, that’s why statistical analysis was invented!

Statistical analysis procedures allow researchers to mathematically identify and isolate background “noise” so that the true treatment effects are more clearly detectable. Statistical analysis helps you conclude whether the observed differences are real and then assigns a probability level that your conclusion is correct.

Planning an OFR Trial

Well-designed trials follow a systematic approach. A meaningful question or hypothesis is developed. The research project is planned and conducted to objectively (without bias) test the question. Data are carefully measured and recorded. Results are statistically interpreted to answer the research question.

When developing a meaningful question or hypothesis, keep it simple, simple, simple! Trials require time, energy and money. Complex trials involve more of each. Ask yourself: Are you a researcher or do you work for a living? The simplest research questions involve a simple yes/no answer (e.g., herbicide 'A' versus herbicide 'B', or fungicide treated corn versus non-treated corn.)

When selecting treatments to evaluate, include a control or check treatment. A logical choice for a control may be your standard or normal practice. The purpose of a check treatment is to answer the question: Does an alternative management practice yield the same as or better than your usual one? Recognize, though, that including a control or check treatment is NOT for the purpose of adjusting yields of other treatments (e.g., as is sometimes done in a non-replicated variety trial). Adjusting plot yields based on the performance of a check treatment across the field assumes that all treatments (including the check) respond similarly to changing field conditions. That assumption is often false.

When the objective of an OFR trial is to identify an optimum rate of an input such as seeding rate or fertilizer rate, include a fairly wide range of treatment levels and then develop regression-based response curve¹ to fit the yield response to seeding rate. For example, to identify the optimum seeding rate for corn, don't simply compare two seeding rates. Rather, establish plots that represent a range of seeding rates (e.g., seeding rates of 18, 24, 30, 36, and 42 thousand seeds per acre).




Fig. 1. Example of a quadratic regression response curve that describes the response of relative yield to harvest plant population.

¹ Microsoft™ Excel™ can be used not only to generate the scatter graph that depicts the responses but can also calculate the regression response equation.

A poorly planned (statistically and logistically) OFR trial has a high risk of failure. Three important components of planning an OFR trial are 1) Request help, 2) Replicate the treatments, and 3) Randomize the sequence of treatments.

Important research decisions include field selection, treatment replication, treatment randomization, plot layout & size, and treatment choice. If research is not your vocation, then don't hesitate to **request help** from those who conduct research for a living. Such folks include university researchers & Extension specialists, industry researchers & agronomists, and independent crop consultants.

Repeating or **replicating treatments** in a trial enables the researcher to mathematically separate the true treatment effects from those due to "background noise". At the very least, treatment replicates help you evaluate whether treatment effects are consistent.

If spatial variability can be identified ahead of time (e.g., soil types), then try to position the replicates such that each "rep" of plots is reasonably uniform within itself. The goal being to best minimize the "noise" level among plots within a single rep.

Randomizing the sequence of treatments within a "rep" decreases the odds that spatial variability (foreseeable or not) will influence the treatment effects. For example, plots in low ground might be stressed more (i.e., soggy soils) than plots in high ground (better drainage) in a wet year, or vice versa in a dry year. Treatments that fall in those areas of the field may be influenced without your knowledge.

An example of a simple on-farm research trial is depicted in the following Fig. 2. The trial consists of three experimental treatments (e.g., varieties) that are replicated three times in the trial. The sequence of the three treatments is random within each of the three replicates (e.g., flip of a coin).

The width of each treatment plot (strip) would typically be equal to that which best matches the field equipment to be used to manage the trial. For example, if the planter was 30 ft wide and the combine header was 15 ft wide, then a suitable plot width would be 30 ft. The length of each treatment plot would typically be equal to the length of the field unless the grower wants to only use a shorter length in the field.

If there is a chance of treatment effects "bleeding over" into adjacent rows (e.g., adjacent plots with different fertilizer rates), then one should plan a plot width that will allow you to harvest the centers of the plots for yield data and leave border rows for gleaning. For example, the combination of a 12-row planter and a 6-row combine header would allow you to plant 12-row plots, but harvest the center 6 rows for yield data and leave the alternating 6 border rows between harvested plot areas for gleaning.

		Border rows
Rep 1	Plot 1	Treatment 2
	Plot 2	Treatment 1
	Plot 3	Treatment 3
Rep 2	Plot 4	Treatment 2
	Plot 5	Treatment 3
	Plot 6	Treatment 1
Rep 3	Plot 7	Treatment 3
	Plot 8	Treatment 2
	Plot 9	Treatment 1
		Border rows

Fig. 2. Example of a replicated field trial with three experimental treatments randomly assigned within each replicate.

Some GIS software programs allow you to design the physical layout of the treatment strips before you head out to the field. This can be very helpful in determining how many treatment strips will fit in a field, determining where the treatment replicates should be located to minimize variability within the replicates, determining the randomization of the treatment strips, and in some cases facilitating the application of the treatments themselves.

After harvest, these layers of GIS information can be merged with the yield data to facilitate the summary of the results. I regularly use ArcView™ v3.x with the available EFRA and XTools extensions, though that is a costly program unless you have access to a site license like we have at Purdue. Other potential software programs that may be of use include AgLeader™ SMS Advanced, FarmWorks™ and MapShots™, though I have little experience using these programs in this fashion.

Admittedly, the practical logistics of implementing both the replication and randomization of treatment plots in an on-farm trial can be a proverbial pain in the rump for growers. Lightbar navigation or autosteer technologies, coupled with GIS-planned plot maps, can greatly reduce the logistical headaches of implementing randomly located treatments.

For example, if the plot layout depicted in the preceding Fig. 2 was for a trial evaluating rates of preplant nitrogen fertilizer, knowing that the 150 lb preplant N rate treatment was assigned to plots 2, 6, and 9 would allow you to navigate specifically to those plots to apply that rate of preplant N (lightbar or autosteer navigation) before you need to change the application settings for the next rate of preplant N in a field trial. Not only would this facilitate the application of the preplant N rates, but would assure you that your subsequent corn planting will line up with the preplant N rate strips.

The power of locations

One of the major limitations with on-farm research is that growers often “go it alone” in conducting their own on-farm trials rather than collaborating with other farmers to conduct similar trials in the same year and then pooling data. We know that weather greatly influences yield itself, but weather can also greatly influence yield responses to experimental treatments. Results from only one location in a single year may mislead your interpretation of the effects of the treatments you are evaluating simply because of the weather patterns that year. Most of the time, field trials need to be repeated over locations or over years simply to experience a greater range of weather patterns. Sometimes, a 1-year, 1-location experiment becomes simply an experience for the farmer with little real impact on future crop production decisions.

Multiple locations of the same trial in the same year or over several years improve our ability to identify consistent and repeatable treatment effects in field research. Encourage several or more farmers to participate in evaluating the same set of experimental treatments and pooling the data. Over the past three years, for example, we’ve had more than 40 farmers from all over Indiana collaborate with us on our nitrogen rate trials. Such a large number of participating farmers has greatly aided our efforts in identifying optimum N rates for Indiana corn growers.

Take plenty of notes

During the season, take notes on any possible “noise” that may influence the outcome of the trial, especially if the “noise” is not distributed equally over the field. This includes field operations (what and when), weather events (especially rainfall), and pest problems (disease, weeds, insects). Walk the field throughout the season. Note the uniformity of stand establishment throughout plot area. Document the occurrence of any oddities that could unduly influence the yield of individual plots (e.g., deer beds, your neighbor kid’s ATV tracks, raccoon parties, ponded areas, stalk rot patches).

Variable crop appearance is often your first clue that background “noise” is developing. Take the time to sketch a map of these extraneous “noise” factors (e.g., ponded areas in some plots but not others). Use GPS-enabled field mapping devices to draw the boundaries of “noisy” areas. Even a low-cost Garmin™ GPS unit (e.g., an eTrex™ model) will allow you to easily mark the geographic position of suspicious areas within a field. These notes may help you decide whether to abandon plots altogether or help you “clean” the yield monitor data after harvest.

If you have access to aerial imagery taken mid- to late season, it can help greatly in detecting spatially variable crop stress within an OFR trial. Color or infra-red photography works equally well. Geo-referenced images are nice, but are not necessary.

A simple digital camera, a small airplane, and immunity from air sickness is sometimes all that is needed!

Minimize the opportunity for harvest “noise”. Check the accuracy of the weigh wagon scales. Check the accuracy of grain moisture meters. Calibrate the yield monitor to the grain conditions of the field to be harvested. Yield monitor calibration also includes calibrating the machine’s grain moisture sensors and machine vibration settings. Triple-check the yield monitor settings! Record any changes you make to the settings in case you need to change them back to their original values. On some yield monitors (e.g., AgLeader PF3000), make sure the unit is set to record the yield data to the memory card.

Sometimes, anomalies occur in yield monitor data that have nothing to do with the treatment effects. If you have good notes or other supporting evidence, there is nothing wrong with “cleaning” these anomalies from the dataset using GIS software. A good example of justifiable data cleaning is the removal of yield data points from within marked boundaries of ponded areas that yielded unusually low simple due to water damage. Similarly, there’s nothing wrong with abandoning whole plots or an entire trial if you do not trust the conditions of the trial.

Analyzing the results

There are a couple of options for analyzing data measured from an OFR trial. One is to simply compare treatment means or averages and decide whether differences are “real” or not. You could flip a coin to decide whether the differences are “real” or not and be 50% certain your decision was correct. Alternatively you could statistically analyze the data and make your decision with a little more certainty.

Data analysis and interpretation can be challenging if the research project was not well designed and/or maintained. Excel™ is capable of simple statistical analyses, but is not intuitive to set up. AgStats02 (Washington State Univ., online at <http://pnwsteep.wsu.edu/agstatsweb>) is a simple to use program, but limits the analysis to one year and one location.

Another data analysis option is to collaborate with university specialists who have access to more comprehensive statistical analysis packages. Most on-farm trials are relatively simple and quick to run through statistical analyses. Find a specialist you can work with, agree ahead of time how to format the data, send/give the data to the specialist, sit back and wait for the results.

Statistical analysis of data allows you to calculate a value that is used to estimate whether the measured difference between two treatments is “real” or simply a result of background “noise”. This value is called the Least Significant Difference and is usually simply abbreviated LSD. The LSD value is calculated based on a chosen probability level designed to minimize your risk of making a wrong conclusion (e.g., when you decide that two hybrid yields are different, you may want to be 95% certain you are correct.)

If two treatment means (averages) differ by more than the LSD value, then you can conclude that the difference is truly due to the treatment effects AND will likely occur again in the future. If the treatment means differ by less than the LSD value, then the observed difference is likely due simply to random chance or background “noise”.

Example 1: In Fig. 3 that follows, the LSD value calculated from the statistical analysis of the data was 11 bu/ac. None of the pairs of treatment means differ by more than the 11 bu/ac LSD value (180 – 170, 180 – 176, 176 – 170), so the appropriate conclusion is that 1) the treatment effects on yield were similar, 2) the observed differences are likely due simply to random chance or background “noise”, and 3) the apparent trends in treatment yields (A > B > C) would likely NOT be repeated in subsequent trials comparing these same treatments.




Fig. 3. Example of a replicated trial that was statistically analyzed and a LSD value calculated for comparing treatment means.

Example 2: In Fig. 4 that follows, the LSD value calculated from the statistical analysis of the data was 8 bu/ac. Based on that LSD value, you can confidently conclude that Trtmt A significantly out-yielded Trtmt B (185 – 176 = 9 bu/ac) and will likely do so again in future field trials, but was statistically similar to Trtmt C (185 – 182 = 3 bu/ac). Treatment C was also statistically similar to Treatment B (182 – 176 = 6 bu/ac).

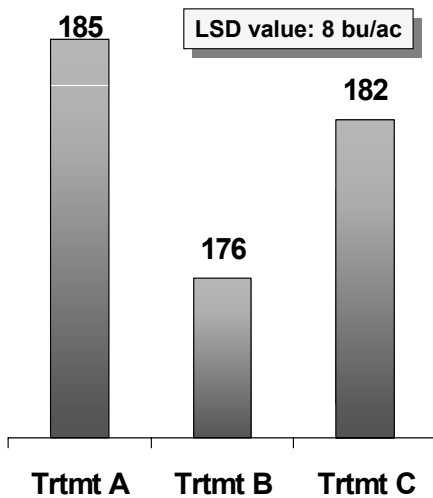


Fig. 4. Example of a replicated trial that was statistically analyzed and a LSD value calculated for comparing treatment means.

In the absence of statistical analyses, a replicated trial still allows you to assess the consistency of treatment effects. For example, if your hybrid consistently outyields my hybrid in every replicate of the trial, chances are that the average difference between the two is indeed significant and repeatable in the future.

For example, in Fig. 5 that follows, Treatment C consistently yielded more than Treatment A in all 4 replicates of the trial. Therefore, the two treatments are likely statistically significant and one would expect that Treatment C would outyield Treatment A in all future trials. However, Treatment C outyielded Treatment B in only two of the four replicates. Even though the average yield of Treatment C (32 bu/ac) is greater than the average yield of Treatment B (25 bu/ac), one cannot confidently conclude that Treatment C is truly superior to Treatment B.




Fig. 5. Example of a replicated trial, but without statistical analysis of the data.

Bottom Line

On-farm research can help answer questions important to growers, but requires sound planning and attention to detail. Background “noise” can play havoc with your ability to detect true treatment effects. Sound research design plus statistical analyses can help isolate background “noise” and improve your success in answering questions with on-farm research.

References

- AgStats02. A simple statistical analysis program for on-farm testing. Pacific Northwest Conservation Tillage Systems Information Source. [on-line] Available at <http://pnwsteep.wsu.edu/onfarmtesting> [URL accessed 2/4/09].
- Farm Works Software. <http://www.farmworks.com> [URL accessed 2/4/09].
- Gomez, K.A. and A.A. Gomez. 1984. Statistical Procedures for Agricultural Research, 2nd Edition. Interscience. Available at <http://www.amazon.com>.

Mullen, Robert, Ed Lentz, Greg Labarge, and Keith Diedrick. 2008. Statistics and Agricultural Research. Ohio State Univ. Fact Sheet. [on-line] Available at http://ohioline.osu.edu/anr-fact/pdf/Statistics_ag_research.pdf [URL accessed 2/4/09].

Indiana Spatial Data Portal. 2008. Indiana University. [on-line] Available at http://www.indiana.edu/~gisdata/isdp_dl/map/m10000.html [URL accessed 2/4/09].

Little, Thomas M. and F. Jackson Hills. 1978. Agricultural Experimentation : Design and Analysis. John Wiley & Sons. Available at <http://www.amazon.com>.

MapShots Software. <http://www.mapshots.com> [URL accessed 2/4/09].

Don't forget, this and other timely information about corn can be viewed at the Chat 'n Chew Café on the Web at <http://www.kingcorn.org/cafe>. For other information about corn, take a look at the Corn Growers' Guidebook on the Web at <http://www.kingcorn.org>.

© 2008, Purdue University

It is the policy of the Purdue University that all persons shall have equal opportunity and access to its programs and facilities without regard to race, color, sex, religion, national origin, age, or disability. Purdue University is an Affirmative Action employer. This material may be available in alternative formats.