

GRDC Precision Agriculture Manual

August
06

Designing your own on-farm experiments

GRDC
Grains
Research &
Development
Corporation


Australian Government
Grains Research and
Development Corporation

Standards for Electromagnetic Induction mapping in the grains industry

Compiled by: Garry O'Leary, Victorian Department of Primary Industries (formerly CSIRO Land and Water) Edited by: Jo Peters, Farm Works

Executive summary

A workshop to establish a protocol for electromagnetic mapping applications in the grains industry was held in Mildura on 3-4 November 2004. The workshop developed a framework and four working groups completed the detail. This protocol consists of ten sections, intended to provide guidance for the standard measurement, recording and interpretation of soil conductivity data; particularly for third parties that might receive the data for subsequent analyses and development of recommendations. The ten sections are:

- EM survey objectives
- EM survey design
- Instrument set-up
- Best management practise
- Fundamental data set
- Data management and processing
- Presentation & reporting
- Soil testing
- Recommended statement of ethics
- Opportunities for future working groups
- Further Reading References Appendices

Designing your own on-farm experiments: how PA can help

Authors: Rob Bramley, Simon Cook, Matthew Adams, and Robert Corner, formerly CSIRO Land and Water

About this guide

When we first started investigating precision agriculture (PA), our CSIRO team quickly discovered that much of what we thought we knew about the way crop production works was difficult to apply at a sub-paddock scale. Like many farmers using a yield monitor for the first time, we found that there could often be as large as a four or five-fold variation in crop yield over fairly short distances within the same paddock.

While our existing knowledge provided valuable understanding of crop production at a regional scale, it was found wanting in explaining what we could now see using the new technology. We quickly realised that not only did we need to do some more field experiments, these experiments needed to be done in a very different way to what we were used to.

The solution was in the very technology that had alerted us to the problem - the yield monitor and variable rate fertilizer box.

After conducting experiments over the past few years we have concluded, that in addition to simply providing farmers with an appreciation of the variability in their production systems, one of the most powerful uses of PA for farmers is helping to conduct their own on-farm experiments.

This is a summary guide for farmers and their advisers on PA-based field experiments - their design, and the important issues to be considered in analysing the results.

Edited by Lloyd O'Connell

ISBN 1 875477 37 3 • GRDC Project CSO179

© This publication is copyright. Apart from any use as permitted under the Copyright Act 1968, no part may be reproduced without permission from the Grains Research and Development Corporation (GRDC).

Any recommendations contained in the publication do not necessarily represent GRDC policy. No person should act on the basis of the contents of this publication, whether as to matters of fact or opinion or other content, without first obtaining specific, independent professional advice which confirms the information contained in this publication.

Contents

1. questions farmers are asking
2. getting started
 - The essential elements of an on-farm experiment
3. simple experimental designs: VRT not required
 - One way strips
 - Two way strips
 - The donut
 - Using VRT with simple designs
 - Be aware of design shortcomings
4. designs for the adventurous: VRT and a yield monitor essential
 - The chequerboard
 - Waves
 - Step and wave
 - Setting up a VRT-controlled experiment
5. analysing the results: how to work out what happened
 - The five golden rules
 - Data analysis - The basic process
 - The nuts and bolts of data analysis
 - Financial analysis
 - Archiving data
6. what next?
 - Evaluating the results - What does the map mean?
 - Forming a plan of action
7. where to find support
 - Further reading
 - Useful web sites
8. glossary of terms
9. acknowledgments

1. Questions farmers are asking

Over the past few years many growers have become interested in precision agriculture (PA) and the various technologies on which it depends - global positioning systems (GPS), yield monitors, variable rate seeders or fertilizer applicators and so on. And many are starting to use these new technologies in their day to day farming activities. But the question farmers are asking is: "What can PA do for me?".

Doug Maitland and family, farmers from Western Australia's eastern wheat belt, have been asking this and other questions as they consider how PA can help their cropping operation. Doug's experiences and questions about PA typifies what other farmers are starting to ask of the new technology.

For some time, Doug has known that different parts of his farm have different productivities. Curious to know more about this variation, and thinking that PA might be able to help him see it more clearly, Doug bought a yield monitor and GPS in 1994 and has been mapping his paddocks ever since.

He has found that not only does the productivity of different paddocks vary, but he also now knows that yield varies widely within them, and as Figure 1 shows, some parts of his paddocks even made a loss rather than a profit.

Doug's maps showed that some parts of his paddocks respond very differently to fertilizer application. So even though he has access to the best available advice, the standard recommendations are not correct for all parts of his paddocks - even though they seem alright on average. This has left him facing some difficult questions:

- Should he increase or decrease the fertiliser rates over the paddock?
- Should he apply more in one part of the paddock than another?
- If so, how many different rates should he apply?
- Should he vary the rates of all his fertilizers or just N?
- What about the seed rate?
- Would he even be better off taking some of the poorest areas out of production?




Figure 1: Yield of barley, expressed in terms of gross margin for a 42 ha paddock in the wheatbelt of WA. Green areas represent those returning more than \$100/ha, yellow areas are those returning up to \$100/ha, whilst the red areas made a loss.

Doug realised that to get the answers to these questions, he needed to experiment a bit. He needed to understand the variation within paddocks in addition to differences between them. And to achieve that greater understanding Doug used PA to do various experiments which were aimed at enhancing his level of control over grain yield and quality.

Grain growers can gain a lot from a general understanding of spatial variability and the principles of how to manage it, but to put this into practice they need to be able to apply that knowledge to their farm and cropping system. This means testing the principles on their own paddocks. An objective of this guide is to explain how to carry out those on-farm experiments.


Doug needed to vary the inputs of his grain production (experimental treatments) and observe, measure and assess the outputs, or results of the experiment.

The key thing about PA is that it recognises that the relationship between inputs and outputs can vary over a particular area - the yield in one part of the paddock can differ from that in another, even when the same inputs are used. This could be a problem if we wanted to do an experiment using two adjacent paddocks.

For example, we might want to compare the effect of applying nitrogen (N) in a single or split application by applying one dressing in one paddock, and split applications in the other. If the two paddocks are not the same, then at the end of the experiment, we could not be sure that any differences in the measured performance of the paddocks was due to the different N management.

But the use of precision agriculture technologies allow us to get around this problem. Variable rate technology (VRT) allows us to vary the inputs within single paddocks, and yield monitors fitted with GPS enable us to measure the effects of doing so on yield or crop quality.

The first PA-based experiment that Doug ended up conducting - which has become known as the chequerboard - is described later in this guide. Since then Doug has conducted many experiments over his whole farm. He has found that the results of these experiments are telling him more and more about his farm and its productivity. As a result of this new knowledge, he now believes that he has better control over his production system and a much better chance of being profitable.



Impressive technology - but what can PA really do for farmers?

2. Getting started

The basic idea of experimentation is simply to measure the effect of a deliberate variation in an input. This idea is very familiar to farmers. But what may be new in the following sections, is the attention to detail necessary to conduct a good experiment. The more precise the laying out of the experiment and the more detailed the measurement of results (and of in-season monitoring), the more information you will get from the experiment and the more able you are to act on the results.

The essential elements of an on-farm experiment

With every experiment there are some basic but essential elements that must be considered before you start.

Asking a question:

What question do you want your experiment to address? It could be:

- Does superphosphate have any effect in paddock x?
- Does the effect of potash vary over paddock y?
- Does the effect of urea depend on the level of super?
- What would happen if I increase the seed rate?
- Would germination improve with deeper cultivation?

You should also consider whether you are interested in experimenting with more than one input - maybe urea and super, or variety and seed rate?

In general, our advice is to stick to just one or two variables in a single experiment.

A plan of action:

Having decided on the question that the experiment is going to address, you need to choose an appropriate experimental design or layout. Possibilities include strips, blocks or more complex patterns such as chequerboards or waves. Your choice may depend to some extent on whether you have access to VRT. It will also be dependent on how time-consuming and complex you think setting up and conducting the experiment is going to be.

Experimental design

The first thing to consider is the number of levels of each experimental variable you want to have in the experiment. For example if your experiment was designed to determine the effects of varying the rate of super, how many levels (rates) of super would you need to use?

A useful way of deciding how many input levels to use is to begin with the average level (which is probably your current 'best bet' option, or the rate suggested by a local adviser), and then choose the tolerance or variation either side that is appropriate to your situation. In other words, you have to decide how far either side of this average rate you are prepared to vary. This decision will depend on two things:

- Your expectation of how far from the average rate you need to go in order to see a difference in yield or quality; and,
- How willing you are to accept a possible loss of yield from the treatments which have levels lower than your average rate.

We recommend that, in general, and for a single variable, you have a maximum of four levels. So you might set two of these above your best bet rate and two below. Or for a three level experiment, you might have one above, one below and the third set at the rate you would have used if you were not experimenting.

Control treatments

A control treatment is really a 'do nothing' treatment. In some experiments, your objective will be to find out whether a particular set of treatments gives a better result than your current practice. In this case, your current practice is the control.

But many crop scientists believe that to identify the optimal rate of an input to use, and/or the full response curve, the control treatment should be a zero or 'apply nothing' treatment.

Our recommendation is that if you have no idea what the optimal level might be, then your experiment needs to include a zero treatment as a control. But if you think your current rate is about right, and only needs a bit of 'tweaking' to account for the variability of your individual paddocks, then use the current rate as the control treatment.

Number of levels and treatments

Your decision about the number of levels to use in an experiment and the sort of control you are going to include, in turn determines the number of treatments the experiment will have. For example, an experiment with two variables at four levels has eight treatments.

Replication

It is a good idea to have some replication in the experiment. We recommend that you repeat all your treatments at least twice, and preferably more than that. This will enable you to be more certain about the effects of an individual treatment.

Space available

The number of treatments and levels to use, and whether or not to include replication, will be influenced by the amount of space you have available. In general, it is not a good idea to spread a single experiment over more than one paddock.

A field plan:

When you have designed your experiment, sketch it out on a piece of paper and mark out where in the paddock each treatment will be positioned. If you have VRT, it is then easy to turn your sketch into an application map using one of the commercially available software packages.

If you do not have VRT, then the edges of each treatment can be marked in the field or on fencelines for later survey.

Yield measurement:

The best way to measure results is with a yield monitor but if you don't have access to one, a weigh-trailer can be used. For simple experimental designs a weigh trailer will give perfectly adequate results, providing harvesting is done carefully and the header is not allowed to wander across treatments.

But if you have access to a yield monitor and a mapping capability, the resolution of the experiment will be greatly enhanced.

Analysing the results:

A word of warning: Do not expect a clear result from the yield map. It is nearly always necessary to analyse the yield variation using statistical procedures to provide a clear answer. Most spreadsheet computer packages, such as Microsoft Excel or Lotus 123, have at least some simple statistical routines that can be used to analyse the results of experiments with more simple designs. The more complex experimental designs require Geographical Information Systems (GIS) software and packages that use more advanced spatial statistical procedures.

All this may sound daunting, but your adviser may have access to these or know how to get it. So don't let concerns about analysis of results dissuade you from experimenting.

Record keeping:

Analysis, and especially the interpretation of results, is much easier if you can be certain about what you have done in the paddock during the year. Good record keeping is essential, especially when it comes to assessing the experimental results from an economic viewpoint. So in addition to keeping a written record of the treatments used in the experiment, you should also be recording details of your other normal farming operations, such as what you do in terms of tillage, weed and pest control and the rates of fertilizers and seed used.

Some optional extras:

In-season monitoring

Careful observation through the season can help explain the effects measured at harvest. For example, an outbreak of weeds or a period of waterlogging could affect the results, especially if it doesn't occur over the whole of the experimental area. If you have a differential GPS (accurate to a metre or so), use it to record where these things occur. If not, draw a mud-map.

Remote sensing

If you are in a program to acquire mid-season satellite or airborne imagery of your crop, it may well help to explain variation in your experimental results.

Soil mapping and testing

The results of experiments conducted in our PA research program suggest that variations in the soil often explain much of the variation in crop performance. So a soil map may help you explain the results of the experiment, especially if accompanied by soil property information.



VRT makes experimentation easy. You can drive across your paddock normally without worrying about the experimental area boundaries

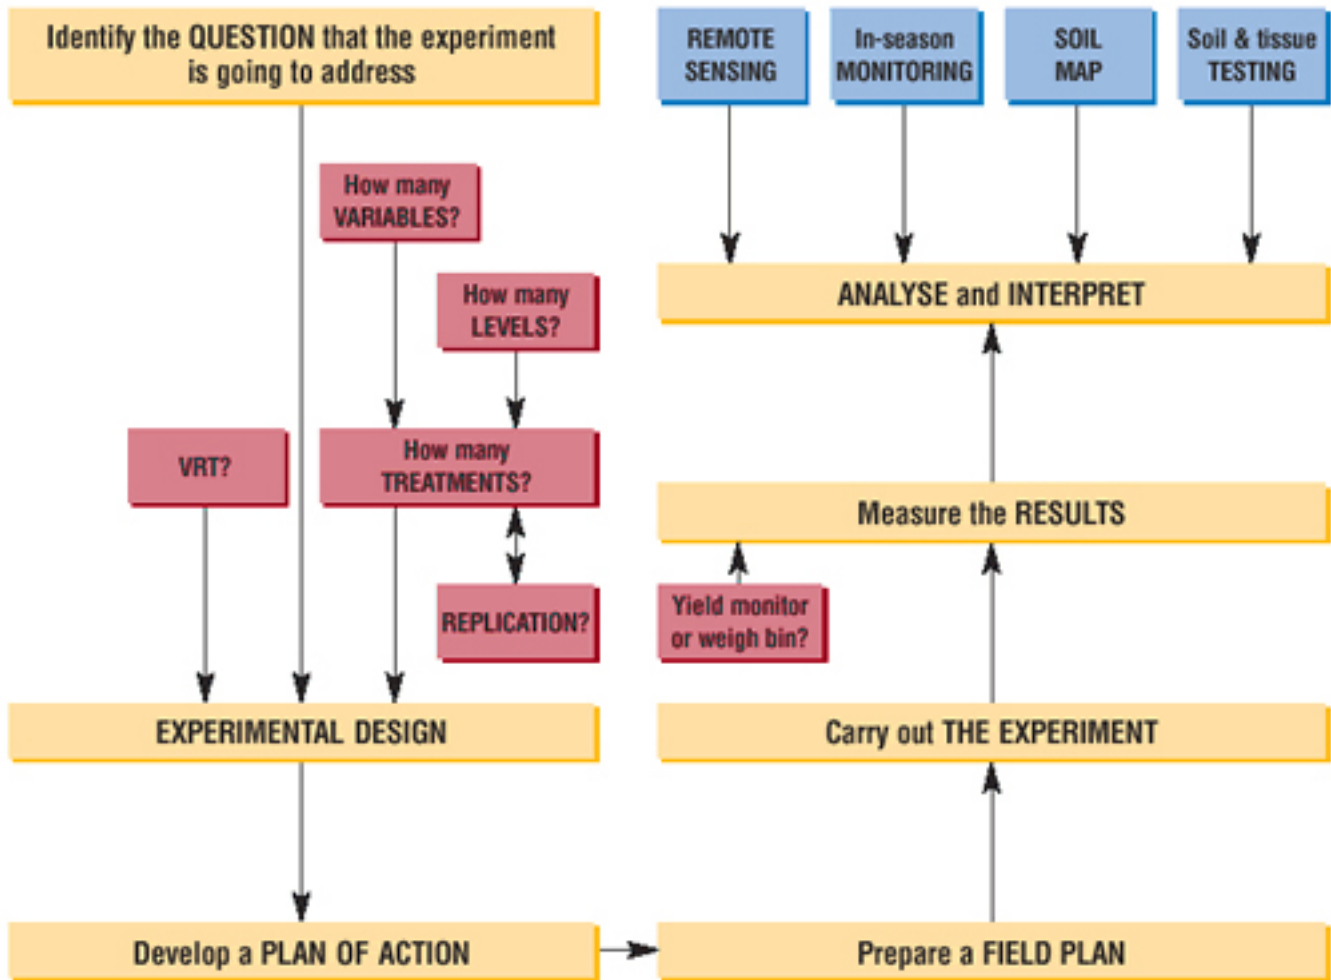


Figure 2: The stages of conducting on-farm experiments. Essential actions are in yellow; optional extras are in blue. Many of the essential actions require consideration of some important questions (shown in red).

3. Simple experimental designs – VRT not required

The following sections provide a step-by-step guide to installing and analysing your own experiments. These sections do not offer detailed instructions on the statistical analysis of experimental results. If you need more help on this topic, contact your local consultant/adviser or read the references listed at the back of this guide. Many of the necessary analytical tools are part of computer spreadsheet packages such as Microsoft Excel or Lotus 123.

This section is concerned with designs which do not require VRT, although the benefits of using VRT are discussed later. Whether or not you are using VRT, it is a good idea - but not essential - to use differentially corrected GPS if you can, as this makes identification of the various parts of the experiment much easier at harvest time.

One way strips

For farmers who do not own VRT, experimental designs must be simple enough to lay out with conventional airseeders, sprayers and spreaders. Simplicity is also a great advantage if the results are to be assessed using a weigh-bin rather than a yield monitor.

Most designs are best laid out using 'up-and-back' wheel-lines rather than 'round and round' for which the donut is suitable. Unless otherwise indicated, the width of the area covered by each treatment should be not less than three times the width of the seeder, spreader, boomspray or header, whichever is the widest. This is an important point as it insures the experimental results against the possible confounding effects of variable machinery performance in the up and back directions.

Select the area covered by each treatment as some multiple of the width of the seeder or boomspray. The essential point is that you know where the boundaries between treatments occur.

Strip designs are generally best harvested using a header equipped with a yield monitor, although a conventional header and weigh bin is satisfactory for most.

One way strips

Single strips: One variable, two levels; or, two variables, one level

Single strip experiments are the simplest option for a farmer interested in experimenting. They are appropriate for addressing a simple question such as: "Do I get better yields when I use seed rate A or seed rate B?". Because this experiment involves one variable (seed) and two levels (A and B) with no replication, it has just two treatments.

But this design (Figure 3a) can also be used for a question such as: "Do I get better wheat yields when I apply 100 kg of nitrogen per hectare as ammonium sulphate or as urea?". In this case, the experiment has two variables (ammonium sulphate and urea) but only one level (100 kg N/ha).

To lay out this experiment split the experimental area (usually a whole paddock) into two, and when moving from the first to the second treatment, change the setting on the air-seeder, spreader or boomspray. The location of each treatment can be easily marked on a fence line for later survey. (Don't forget to do this survey, even if it is no more complicated than measuring along a fenceline with a tape measure - you will need this information for analysis of results.)

This experiment is analysed by comparing the yields in one strip with the yields in the other. It is possible to gain additional information from this design if, at harvest time, you can divide each strip into a number of sections along the line of the strip and weigh the yield obtained in each section separately. You can then see if some parts of the paddock respond differently to others. This is especially useful if you are aware of changes in soil type within the paddock. In this case the experiment should be laid out so the strips run across the change in soil type.

One way strips

Repeated strips: One variable, two levels

This design (Figure 3b) is basically a replicated version of that shown in Figure 3a and answers the same sorts of questions. But it is the preferred option for this type of simple experiment. This design is only suitable for 'up-and-back' wheel-lines. You should aim for at least three or four replicates.

The results can be analysed using simple statistical routines such as analysis of variance (ANOVA). Alternatively, simple comparisons of average yield for each treatment can be made over the whole paddock. The replicated design also allows comparison of treatment effects across the paddock which provides information on the variability of the results over the paddock.

This is done by comparing adjacent strips across the paddock, or adjacent sections of strips down the paddock.



Figure 3(a): Experimental design for a simple two treatment, single strip experiment. Figure 3(b): Experimental design for a replicated strip experiment with two treatments.

Highly replicated strips

This design (Figure 4a) is very similar to the repeated strips design shown in Figure 3b. But its highly replicated nature enables an understanding of treatment effects over the whole paddock through the production of a yield difference map (Figure 4b).




Figure 4(a): Design for a split-planter or adjacent strip experiment. The design can be separated into a separate plan for each treatment. This is important for producing a yield difference map (Figure 4b).

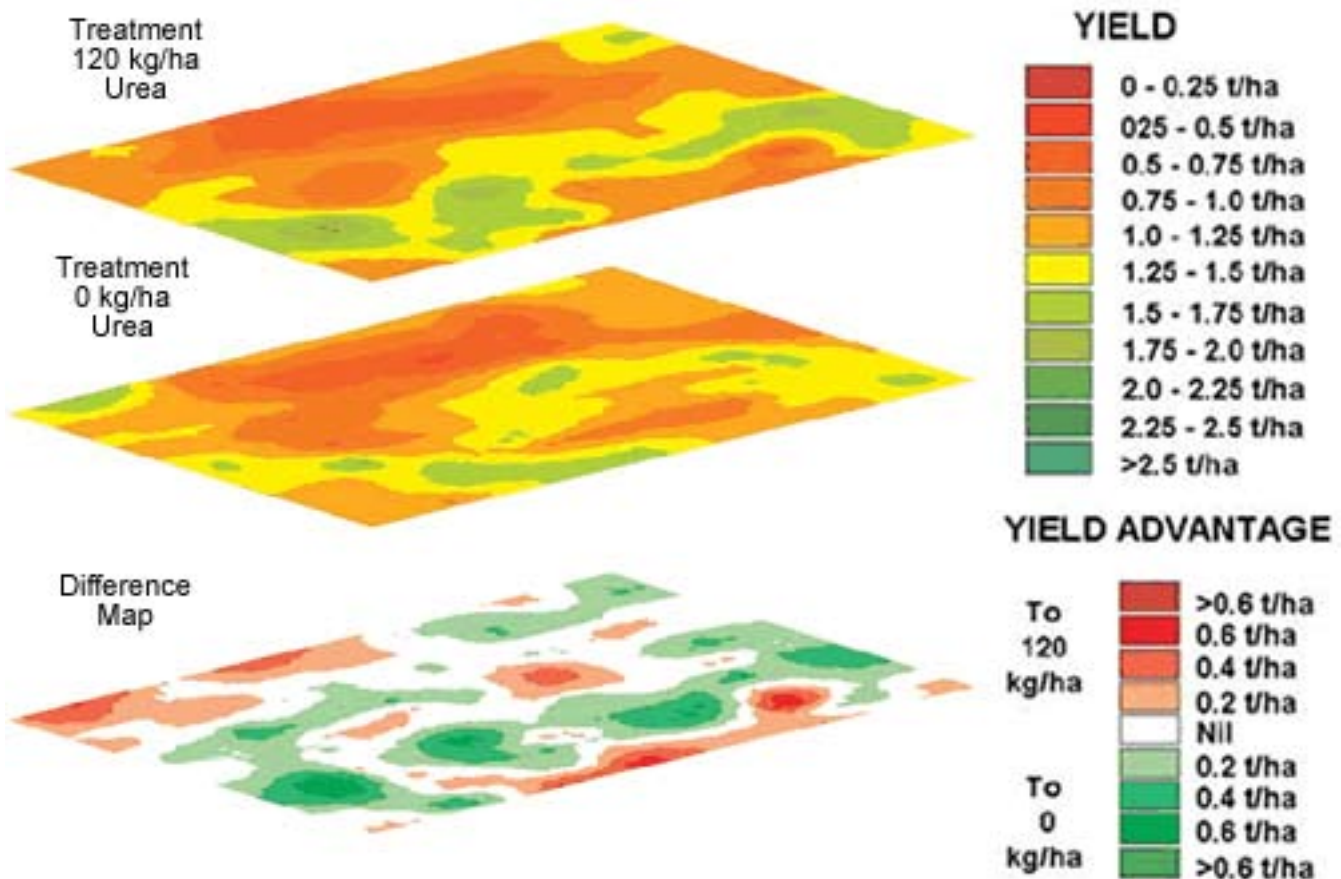


Figure 4(b): The analysis of a chequerboard experiment using yield difference maps. A GIS is used to produce two yield maps for the paddock using statistical interpolation routines - one for the treatment being analysed (in this case, 120 kg urea/ha) and one for the control treatment. These are then subtracted one from the other to produce a third map - the yield difference map. This shows the relative performance of the treatment across the paddock.

The important point is that the application map shown on the left hand side of Figure 4a can be split into separate application maps for each treatment. This is easy to do in GIS. Experiments with this design must be harvested with a header equipped with a yield monitor and GPS; it is further improved if some form of machinery guidance or VRT is used in laying the experiment out.

This technique (also called the split planter design) was developed by Pioneer Hi-Bred International in the US to assist growers to select the best hybrids for their farms.

In the case of either a variety comparison or, for example, comparing the benefits of applying either mined gypsum or phosphogypsum, remember that the experiment is a two variable/single level experiment and so great care is needed to ensure the level is the same for each variable.

For a variety experiment, you must make certain that the seeding rates are such that the same stand of each variety is achieved in each treatment. And in the gypsum example, you need to ensure that the same amount of product (the level of calcium), is being applied to each treatment area.

This sort of experiment is laid out as though you were using the design shown in Figure 3b, but with as many strips as possible given the width of the paddock, and with each strip being three times the width of the header to be used at harvest.

If you have GPS and/or VRT, you could apply all of treatment 1 followed by all of treatment 2. If you were doing a fertilizer rate experiment you could adjust the rate every third pass, but this could result in a lot of additional work.

An important consideration in this design is that when you get to the stage of analysing the results, you will need to produce separate yield maps for each treatment over the whole paddock. To do this, GIS software is essential because the analysis depends on spatial interpolation routines or geostatistics (see below). In general, these maps are very useful, but if the distance between pairs of the same treatment gets too large, they start to falter.

This is the reason why this design works best if only two treatments are used.

Spatial interpolation is a statistical method which essentially amounts to filling in the gaps. In fact it is a means of estimating what is in the gaps based on a knowledge of what surrounds them.

So using Figure 4a as an example, spatial interpolation is used to estimate the yield that would have been achieved from treatment 2 in the strips that actually contained treatment 1, and vice versa. Separate yield maps are produced for each treatment for the whole paddock. These are then subtracted one from the other to produce a yield difference map (Figure 4b).

You can then see how the relative performance of one treatment compared to the other throughout the paddock. It then becomes possible to split the paddock into sections in which one variety or treatment outperforms the other(s). This information can then be used in subsequent years in designing an application map for use with VRT.

Repeated strips, one variable, multiple levels

This design is essentially the same as that shown in Figure 3b, except more than two levels are used. It is used to answer questions such as: “What is the best rate of potash to use on lupins?”. Figure 5 shows an example in which each treatment is replicated twice and within each replicate, the order of the treatments has been mixed-up. This mixing ensures that the experiment is designed so that any experimental variation (the thing you want to measure) is independent of any other source of variation.




Figure 5: Repeated one-way strip design for a single variable experiment with four levels.

This experimental design is especially suitable for fertilizer trials because it will provide more information about the nature of the response to the fertilizer, and with careful choice of levels, should identify the optimum level of fertilizer required. More than four levels is probably unnecessary unless you wish to define the whole response curve. One of these levels should ideally be a control.

The experiment can be analysed in several ways: contrasting the results from the various treatments; ANOVA; fitting simple regression equations that describe the effect of the variable on yield as a function of level; and, where a yield monitor and GIS is available, using a ‘moving window’ regression technique.

Two way strips

If you want to investigate the effects of two variables, but do not have access to VRT, two way strips (Figure 6) should enable you to adequately answer the question of interest. This could be something like: “If I increase my seeding rate, do I also need to apply more N?”. This sort of experiment should not be conducted without replication.

If this design is used over a whole paddock in an experiment where the number of replicates is maximised, it is possible to assess variable response over the paddock. This is especially the case if the experiment is harvested with a header equipped with a yield monitor and GPS. This design is not recommended if harvesting is to be done with a weigh bin.




Figure 6: Replicated two way strip design for an experiment with two variables, each with four levels (denoted 1, 2, 3, 4 for one variable and A, B, C, D for the other). Note that the experiment has 16 (4 x 4) treatments, replicated twice.

The donut

Many farmers prefer to sow crops in a 'round-and-round' pattern, rather than 'up and back'. The donut is a simple experimental design (Figure 7) which accommodates this preference. It is aimed at the same sorts of questions as the simple strip designs and is analysed in the same way. As with the strip designs, you should keep the donut simple and ensure that each treatment is at least as wide as three times the width of the seeder, spreader, boomspray or header, whichever is the widest. So go for three or more runs at one rate, remembering to record how many it was, and then change the rate.

If you feel a bit more adventurous, you could introduce an element of replication into the design by changing the rate back to the original level, so that the treatment resembles three or four (Figure 7b) concentric rings. Alternatively, you could have more than one level (that is, where more than two treatments are being compared, such as in Figure 7c).




Figure 7: Simple donut designs for experiments with one variable and two levels. Design 'a' has no replication whilst the treatments are replicated twice in design 'b'. Note that in practice, the round and round sowing pattern results in a spiral rather than the concentric square rings which, for convenience, we have shown here. As with the strips, the important point is that you know where your treatment boundaries lie. Design 'c' is more complex with one variable but at four levels.

As with the strip designs, if the paddock can be divided up into sections, donut experiments can yield more information. But we do not favour donut designs unless they have at least some replication of treatments and some separation of the paddock into sections because they are biased towards the centre of the paddock. This is illustrated in Figure 7c which shows that the treatment in the centre of the paddock only occurs in the centre.

In some cases, this may not be a problem. For example, for an experiment comparing two seeding rates in a paddock of uniform soil type, the simple two level comparison (Figure 7a) is quite good because the two treatments always occur adjacent to one another. In this case, the two treatments can be compared around the paddock.

But the donut design is not recommended for fertilizer rate trials, or for paddocks where there is known soil variability, because you can't always be confident that each treatment will cover at least some of each soil type.

Similarly, a donut design with more than one level is not recommended for fertilizer rate trials without replication because the comparison between the treatments is likely to be subject to the effects of underlying variation - such as a different clay content or availability of soil water at the edge compared to the centre.

The assumption underlying the donut design is a bit like saying that the greens on a golf course are the same as the fairways! So ask yourself whether or not you would be better to go for strips. If you do stick with the donut, incorporate as much replication (for example, Figure 7b) into the design as possible.

If you do decide on a donut experiment, it is important to know exactly where the changes from one treatment to another occur. If you have VRT, or scouting GPS, this is easy. If not, it can be estimated roughly in terms of the distance from the paddock boundary, D , where:

$D = \text{Number of runs} \times \text{width of spreader bar}$

And if you do have access to VRT, alternative designs would, in most cases, be better because the donut design has many layout and analysis difficulties.

Using VRT with simple designs

Variable rate technology (VRT) was originally developed to enable farmers to apply prescriptions within paddocks. If you have VRT, it can also make experimentation a lot easier. There is nothing to stop you from using VRT to lay out experiments with any of the strip or donut designs discussed above. There are a number of reasons why you might like to do this:

- Laying out the experimental design is done automatically, so you do not have to spend valuable time during seeding setting out designs and marking fencelines;
- Automatic location and control give you the flexibility to vary rates in any way that you choose. This can be useful for the inclusion of small zero control treatments within the trial;
- Automatic logging keeps a detailed record of exactly what went where; and,
- Automatic location and control allows you to try some more intricate designs.

Basically VRT makes the laying out of experiments easy - the technology can read a pre-designed map and use it to control the application of experimental variables. You can drive normally, and no longer have to worry about marking out treatment boundaries on fencelines. With differential GPS, the machine always knows exactly where it is (within a metre or two), and receives continuous instruction about what rate of seed, fertilizer or herbicide to put there. So you do not have to keep stopping to change rates.



VRT capability not only allows prescription application of inputs, it also simplifies on-farm experimentation.

Most VRT machines also have the capacity to carry machinery guidance systems which prevent overlaps and misses, and are also able to record what rate was actually applied. Again, this has several benefits.

First, the amount of experimental error (or noise) in the analysis of your experiment is reduced because the rate that you applied is known more accurately. Having confidence about what was actually applied assists in the analysis and interpretation of the results.

Second, knowledge of what was actually applied, and the ability to compare this with what you thought you were applying, enables you to check on potential problems with the spreader, such as blockages in the delivery system.

Overall, the beauty of VRT is that it enables you to keep much better control over your experiment and get a much better understanding of the reasons for different results in different parts of the paddock. In other words, it allows you to split the paddock up into sections.

Assuming that the experiment provides you with a useful result, the subsequent division of paddocks into sections requiring different management promotes much better control over the farming system.

Be aware of design shortcomings

The designs discussed in this section all provide useful ways of answering simple questions about agronomic performance. But be aware of their short-comings.

With the exception of the adjacent strip (split planter) design all designs covered in this section favour contrasts in one direction or another.

For example, the donut is biased towards the middle of the paddock. Bias is introduced into the strip designs by the orientation of the strips and the possibility that all treatments are not equally affected by an underlying source of yield variation such as a change in soil type.

A method to overcome these possible biases is to use 'embedded' designs, such as the chequerboard or wave (see next section). But these methods require VRT, high quality yield maps and more sophisticated statistical analysis in GIS.

A machine (like a computer!) will only do what it has been told to do. So, even if you are using a simple design with VRT, you still need to draw up an experimental plan before you start.

4. Designs for the more adventurous – VRT and a yield monitor essential

One of the benefits of VRT is that it enables experimentation with more intricate designs than strips or donuts. In fact with VRT, just about any design is possible. But it is wise to keep the designs as broad as possible so the experimental areas do not become so small that the risk of measurement error (noise) is increased. VRT also makes the use of zero treatment controls more palatable as they can be confined to smaller areas.

Simplicity is all very well, but there can be significant practical benefits of using the fancier designs. Aside from the advantages of using VRT in experimental layout, VRT-dependent designs also have some important advantages over the simpler designs:

- All can be highly replicated to the extent that all treatments occur in all parts of the paddock. As a consequence, both whole-paddock and site-specific results are clearer;
- With VRT it becomes possible to take account of underlying variation, such as changes in slope. It therefore becomes possible to repeat experiments in different management zones within the same paddock; and,
- Consequently it is possible to design VRT-based experiments which are essentially free of the problems of bias.

In this section we discuss the three designs that we have used so far in our research.

The chequerboard

Many people - including researchers - often have a sense of foreboding creep over them when they first see a chequerboard design (Figure 8); “How are we going to analyse the results of that, never mind lay it out in the first place?” is a common reaction. In fact, the chequerboard is very simple, and it is the design that Doug Maitland chose for his experiments.

The chequerboard design is very similar to the strip design shown in Figure 5, and is used to address the same sort of question, such as: “Is the locally recommended rate of N for durum wheat really the best rate for me to use in my paddock? - I know I need to be exact about this because of the narrow target protein range for durum”.




Figure 8: A chequerboard design for an experiment with one variable, four levels and 225 replicates.

The chequerboard has one variable and four levels - as was the case with the design in Figure 5. But there are two important differences between the chequerboard and the replicated strip design.

- Depending on the size of the paddock and your machinery, the chequerboard design maximises the replication in the experiment; and,
- Every treatment in the chequerboard is surrounded by at least one treatment with the other three levels of the experimental variable. Within the experiment, every possible area measuring two treatments wide by two treatments long contains treatments with the variable at all possible levels (note that chequerboards with more than one variable are also possible, but are also more complicated to design and analyse).

The continuous repetition in the chequerboard design means that when the experiment is analysed in GIS, it is possible to estimate the yield response to the experimental variable at every point in the paddock. So in addition to producing a yield map depicting the experimental results (which in an ideal case would closely resemble the experimental design), it is possible to produce a map of crop responsiveness. This is important when the question of risk in farming is considered.

One of the reasons why Doug likes the chequerboard is because it is very low risk. While Doug likes experimenting, he also has to earn a living. So to minimise the risk of losing money through experimenting, Doug chose the levels to use in his chequerboard by deciding on the rate of N he thought would be appropriate if he was not conducting an experiment - his local adviser had a fair bit to do with helping Doug select this rate.

Knowing the area of the paddock, he could then calculate how much he would have to spend on fertilizing the paddock with urea. For the experiment, Doug chose four levels that were either side of his recommended rate; two that were quite close to it and the other two a bit further away.



The success of precision farming depends on a farmer's willingness to investigate all the factors affecting crops. This includes the use of new technology and existing monitoring packages.

As a result, the cost of fertilizing the chequerboard was about the same as the cost of fertilizing the paddock had there been no experiment. So the experiment was roughly cost-neutral.

In practice, the selection of these levels is simply a matter of your attitude to risk - the more spread out they are, the riskier the experiment may be. On the other hand, the more spread out they are, the more information about yield responsiveness you are likely to get. But the key point is that overall, the chequerboard is a low risk experiment.

In addition, the crop responsiveness map that can be generated with the chequerboard provides useful information when making farm management decisions based on your attitude to risk.

Setting up a chequerboard experiment

As with the simple strip experiments the first thing is to decide on the experimental area and measure its size. Next, you need to consider how big each of the 'checks' should be. They should be no smaller than the equivalent of three widths of the header. Whether the checks are squares or rectangles does not really matter, but their width must be a multiple of the header width otherwise there will be a lot of unnecessary smoothing of results and consequent loss of data when the experiment is harvested.

We have found squares to be convenient and the GIS analysis is more straightforward with squares. This means that if your header is 10 metres wide, you need to make the checks no smaller than 30 metres by 30 metres.

Having decided on the size of the checks and the levels to be used in each treatment, you need to construct an application map so the VRT machine knows what to put where. Most VRT machines come with software that lets you do this. Once you have your application map, you are ready to start the experiment.

The chequerboard design can be analysed in a number of ways. The average response for the whole paddock can be determined from the average yield obtained in the replicate treatments. But a much more useful analysis is to use moving window regression which is one way of assessing the responsiveness around the paddock. Another is to use spatial smoothing techniques in GIS to construct a separate yield map for the paddock for each treatment, and then produce maps depicting the difference between the yield obtained for the whole paddock for each treatment.

Waves

One of the problems with any sort of experiment involving treatments in fairly small areas, is that you can run into difficulty with smearing or leakage between treatments - both during the layout and yield measurement stages. In other words, it is easy to accidentally apply one treatment to an area intended for another, or to get the yields from adjacent treatments mixed up.

This is why we advise that treatments should occupy areas no less than three times the width of your equipment. If there is some smearing, you then have a fall-back position by just considering the yield data for the middle swath.

GIS allows you to analyse the results from the middle swath as though they represent the whole treatment area. This is quite acceptable even in simple designs which you may well be analysing without GIS.

But waves remove this problem altogether because there are no specific levels applied as single treatments. VRT is used to apply the experimental variable in a wave pattern across the paddock.

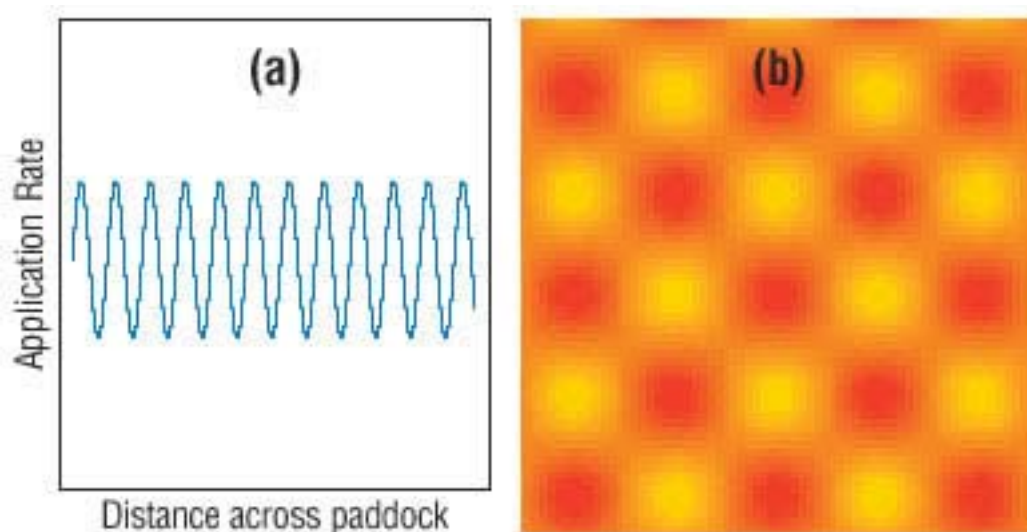


Figure 9(a): Variation in the rate of application of an experimental variable with distance in a wave design. In this example, the wave is a sine-wave. **(b):** Application map for an experiment using the wave pattern shown in 9a.

Figure 9a indicates how the machine calculates what to apply, and Figure 9b shows what the application map looks like when the wave is applied to a paddock. The wave design is similar to an egg carton. The eggs sit in the valleys (the red areas in Figure 9b), and are separated from each other at regular intervals by the peaks (the yellow areas in Figure 9b).

For the mathematically inclined:

A convenient waveform to use for wave experiments is the sine-wave which can be calculated as follows:

$$\text{Rate} = R + Z$$

where

$$Z = A \times (\sin(6.28 \times X)/W) \times (\sin(6.28 \times Y)/W)$$

R is the rate you would use if you were not experimenting.

In these equations, A denotes amplitude, W denotes wavelength and X and Y are the coordinates (easting and northing) of each point for which an application rate (Rate) is calculated.

It is important when making these calculations that X, Y and W all have the same units. We have found it convenient to set X and Y by determining the coordinates of the lower left and upper right corners of the paddock using the GPS.

Waves are otherwise very similar to the chequerboard and are suitable for determining how yield responds to a variable such as fertilizer or seed rate across a paddock. The decisions that you have to make in implementing this design relate to the wavelength and amplitude .

- The wavelength - or distance between peaks - is set by the characteristics of your header and the problem of data convolution (that is, the error in measuring grain yield at a given point in the paddock caused by factors such as the time it takes harvested grain to reach the yield monitor and variations in grain yield along the cutter bar).

Our recommendation is to have your wavelength set no less than 100 m. This means that there is a minimum of 50 m between points receiving the highest rate and those receiving the lowest.

- The amplitude is the range of application rates over which you are prepared to experiment.

Similar to selecting levels in the chequerboard, you should identify the rate that you would be advised to use if you were not experimenting and decide - based on your attitude to risk - how far from this rate you are prepared to go.

The amplitude is equal to half the difference between the highest and lowest rate. If you wish, the lowest and highest levels chosen can be replaced later by zeros and a very high rate.

You also have to decide how dense you would like your application points to be. We have found that when using a wavelength (W) of 100 m and a VRT applicator 13 m wide, calculating an application rate every six metres gives good results. This reduces the potential problem that can be caused with large rate changes occurring on-the-go.

The golden rule when setting up a wave or chequerboard experiment is **DRIVE SLOWLY!**

When you have calculated your application rates, they need to be stored in a file for input into the VRT machine. Some machines read the data from this file, whilst others need a map.

Because the application map for two variables can be overlain in GIS, wave experiments can also accommodate more than one variable. And because the wave has a symmetrical design in all directions - providing the wavelength is the same for all variables - the experiment remains balanced.

Waves are not suitable for experiments designed to assess the most appropriate variety to grow in a paddock (the chequerboard could be used for this, depending on the capability of the seeder to accommodate more than one seed type). The highly replicated strips design is the most suitable for variety comparisons.

As with the chequerboard, analysis of wave experiments can be done using moving window regression of yield on re-sampled points or spectral analysis.

Step and wave

It might be the case that you want to do an experiment in a paddock that you already know has different areas of productivity. You may already be working with different prescriptions for each area but are still uncertain how uniform these areas are. One way of finding out is to use the step and wave design.

For example, if you have three distinct productivity areas within a paddock (or several areas but with each one characterised by one of three levels of productivity), the step and wave design illustrated in Figure 10 could be useful. It can be thought of as a series of wave experiments - one for each area.

The relative size of the steps and waves depends on how certain you are about each. If you are certain about the nature of the contrast between the different areas within the paddock, then the height of the step can be assigned accordingly. But if you are uncertain about the degree of contrasts, it is better to reduce the step height and increase the wave amplitude.

Selection of the wavelength and amplitude within each step, and the average rate that you would apply if you were not experimenting is done in the same way as the single wave. The calculation of application rates is also the same with a different value of R for each zone within the paddock. The analysis of the experiment is also done the same way as in a single wave design.

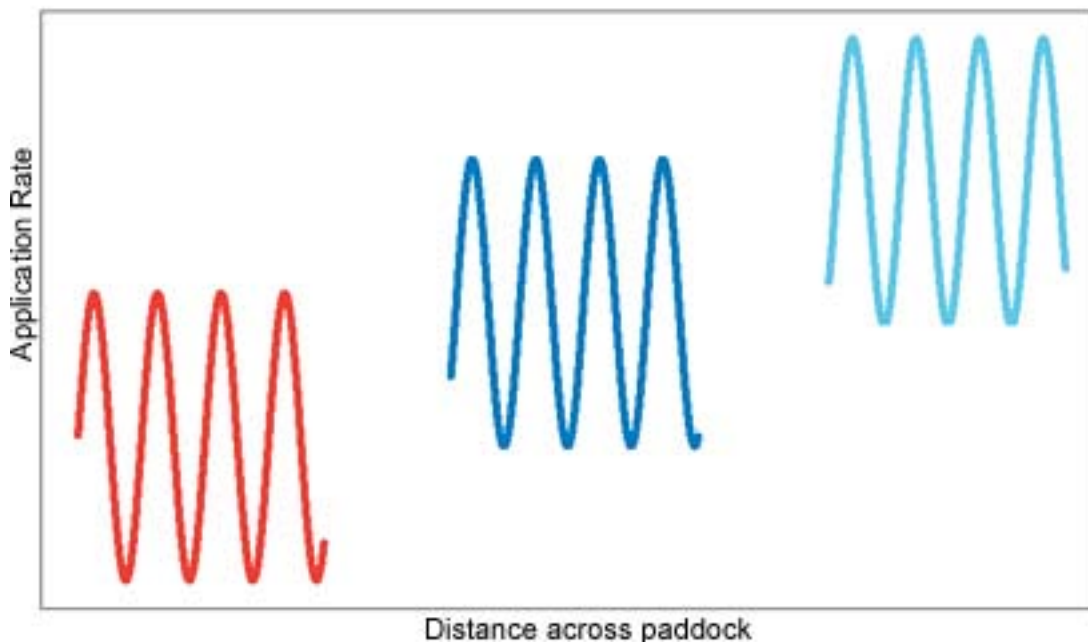


Figure 10. Variation in the rate of application of an experimental variable with distance in a step and wave design. Again, the wave form used here is a sine-wave.

For the mathematically inclined:

A convenient waveform to use for wave experiments is the sine-wave which can be calculated as follows:

Rate = R + Z where $Z = A \times (\sin(6.28 \times X)/W) \times (\sin(6.28 \times Y)/W)$ R is the rate you would use if you were not experimenting.

In these equations, A denotes amplitude, W denotes wavelength and X and Y are the coordinates (easting and northing) of each point for which an application rate (Rate) is calculated.

It is important when making these calculations that X, Y and W all have the same units. We have found it convenient to set X and Y by determining the coordinates of the lower left and upper right corners of the paddock using the GPS.

A summary of basic design rules

How many replicates?

In general, the more replicates you have, the more detail you will get about the yield response to the experimental variable. As an absolute minimum, you should aim to get at least three in any given direction. The trade-off is between getting a clear picture of the response curve for the whole paddock (to fertilizer N for example), and getting a detailed picture of variation in N response within the paddock. Analysis is made easier if the replicates are ordered in a regular pattern: 0, 1, 2, 3, 0, 1, 2, 3 ...

What is the minimum plot size?

The thing that limits minimum plot dimension is the size of your seeder, boomspray or header. Some spreaders or boomsprays can split rates, but there is no point in having plots only half a cutter bar wide as it is impossible to separate out the treatments in the yield data after harvest. Furthermore, if you intend to harvest across-track, you will need plenty of space to allow for smearing.

Plot orientation

Since you want to separate out the effects of the treatment from any underlying trends, it makes sense to try and avoid confusing the two during experimental layout. A good example of this is laying out the experiment across slope, so that each part of the slope gets the same treatment. Orientation becomes less critical if you have plenty of replicates, as you will have plenty in all directions anyway.

Overlapping treatments

As you get more confident with the experimental process, you may want to speed things up by running two experiments in the same paddock, or at least, consider two or more variables such as seed and nitrogen. This is perfectly reasonable, providing you follow the same rules as above and prevent confounding the two together. This can be done by setting out the first treatment one way and the second across it.

Setting up a VRT-controlled experiment

There are seven basic steps for setting up a VRT controlled experiment for a paddock:

- Select the treatments of interest (seed, N, P, spray etc);
- Set the budget for the paddock;
- From experience, soil test data or other information to hand, and bearing in mind your own attitude to risk, decide on the average, maximum and minimum rates that you are prepared to vary;
- Using an outline map or photo image of the paddock as guidance, draw up a diagram of your preferred experimental plan. Do not worry too much about detail at this stage, just get the design about right taking into consideration factors such as a reasonable size of plots, sensible number of replicates, informative range of treatments, etc;
- Draw up the experimental design in GIS, and use this to calculate the total bill for seed and fertilizer over the paddock (do not forget to include the controls). If necessary, tweak the values until they seem reasonable;
- When you are happy with the design, and know how much it will cost to install, save the map in the format required to run the VRT; and,
- Allow for a review closer to seeding. If you insist, it is not too late to change settings as the season breaks. Just recalculate the totals for the map.

5. Analysing the results - how to work out what happened

When you have harvested the experiment, you are obviously going to want to know what happened:

- Did the experiment work?
- Do the results tell you something important?
- If so, is this something you did not know before?
- How should you change your management in the future in order to accommodate this new knowledge?
- Do you need to do another experiment, or maybe repeat this one?

Before we think about the way in which you will answer these and other questions, we need to be very clear about the answer to the first one: The experiment did work! Experimental results will always reflect what happened during the period of the experiment. It is true that you may have been unlucky and had a flood, drought, or even a fire during the life of the experiment, but the effects of these will be reflected in the results.

There are five basic rules to guide you through the process of analysing on-farm experiments.

The five golden rules

1. Don't expect too much from your first experiments

Doug started experimenting to help answer some fairly basic questions about his paddock: Did urea produce a yield response?; did yield respond equally over the whole paddock?; and, what about protein?

He is proceeding in fairly simple steps over a number of years. Even so, the answers are sometimes unclear - the areas involved are so large that the unexpected nearly always happens.

Uncertainty about experimental outcomes does not need to be a problem if, like Doug, you take a patient approach. Each experiment adds a bit more to the picture, and as each one is fairly simple, the chance of mistakes is reduced.

Remember that a result that shows no difference between the treatments is still a useful result, even if it might be a bit disappointing initially. In the case of a fertilizer trial for example, a lack of treatment differences tells you that the site was not responsive. This could have been due to drought or a major infestation of weeds or pests, or maybe the season went well - you just didn't need any fertilizer!

2. Don't expect the answer to jump out at you from the raw yield map

Probably the first thing you will see in the yield map is.....well, nothing obvious!

It is rare to see an obvious effect of treatment in the raw (unprocessed) yield map - even the control plots may be difficult to pick out. Don't panic! Treatment effects can be quite subtle, and easily masked by major variations within the paddock. This is also the reason why you should not expect the treatment effects to be obvious when inspecting the experimental paddock during the season.

3. Separate out treatment effects from noise

The first thing that Doug noticed when he inspected the yield map from the first experiment was the large impact of underlying trends within the paddock. This is illustrated in Figure 11. Over some parts, shown on this map as light blue or yellow, yield barely reached half a tonne per hectare, whereas just 200 m away, the same crop was producing over four times as much.




Figure 11. The total variation in experimental results is due to the effects of the treatments, spatial effects and noise. Even after taking out noise (which accounts for much of the fuzziness in this map), the spatial effects can be very large, potentially masking the treatment effects. For example, the trend in the eastern half of this paddock ranges from less than 0.5 t/ha (light blue) to over 2 t/ha (dark blue).

While Doug is interested in these broad trends, the question he wanted his experiment to address was: “What was the effect of urea?”. To properly answer this, the treatment effects must be separated from the underlying spatial variation. We did this by linking yield variation to one of three factors:

- Treatment effects: The variation caused by changes in urea application rates;
- Underlying (spatial) variation: The broader variations due mainly to changes in soil type, degree of weed infestation or other inherent site factors; and;
- Noise: The variation we cannot account for because of things like measurement error or harvesting problems.

These consider the variation to be characterised as:

Total variation = treatment effects + spatial effects + noise

The result is insight into both the treatment effect (where did urea have an impact?) and the underlying spatial effect (how does the yield potential vary within the paddock?).

The urea response map showed Doug that while a few areas responded strongly, much of the paddock responded poorly, and that some areas actually responded negatively - apparently because of haying-off. This pattern will not always be true. There was a poor finish to this season, but it gave Doug detailed new insight into how the urea worked with the crop.



After cropping information has been collected farmers can use a range of computer software programs to create application maps. These are used to direct inputs in the following season's crops.

4. Treatment effects can vary from year to year

Perhaps the largest source of uncontrolled variation in yield maps is caused by the weather. Our experience is that change is normal - only relatively small areas are consistent performers.

5. Treatments can interact

Crop growth can be limited by any one of a number of factors and may be affected by several at the same time. Changing one limitation (for example, the rate of potash) may modify the influence of another (for example, the rate of N). For example, Doug found that in some parts of a paddock he could improve the response to superphosphate by increasing the seed rate, because the denser plant populations were able to squeeze out weeds on the better fertilized soils.

A word of caution about headlands

Headlands are a potential source of problems in generating yield maps because the harvester may pass over them many times. The easiest way to minimise headland aberrations in yield maps is to turn off the yield monitor when you harvest headlands.

Data analysis – The basic process

The basic process of data analysis is illustrated in Figure 12.

- Data preparation is aimed at getting the best yield map possible by removing holes, spikes and other known errors in the original yield data. This is then followed up with various mathematical procedures which create a continuous map surface from yield values measured at specific points - that is, interpolation. After a while, using these methods will become standard procedure, so record what you are doing as you go.
- Estimating the whole-paddock effect enables you to answer the basic question as to whether the treatments had any effect over the paddock as a whole? This can be estimated by comparing the yield within and between treatment plots, using conventional analysis of variance (ANOVA).
- Estimating the site-specific effects is probably what is of greatest interest to a grower using PA. So this stage aims to look more closely within the paddock to produce a map showing the site-specific effects.

- Estimating the likely causes of the variation within the paddock is the final stage in the process. This is done by comparing the map of site-specific effects with other information, such as the soil map. Is there an association between the two? If so, how strong is it?

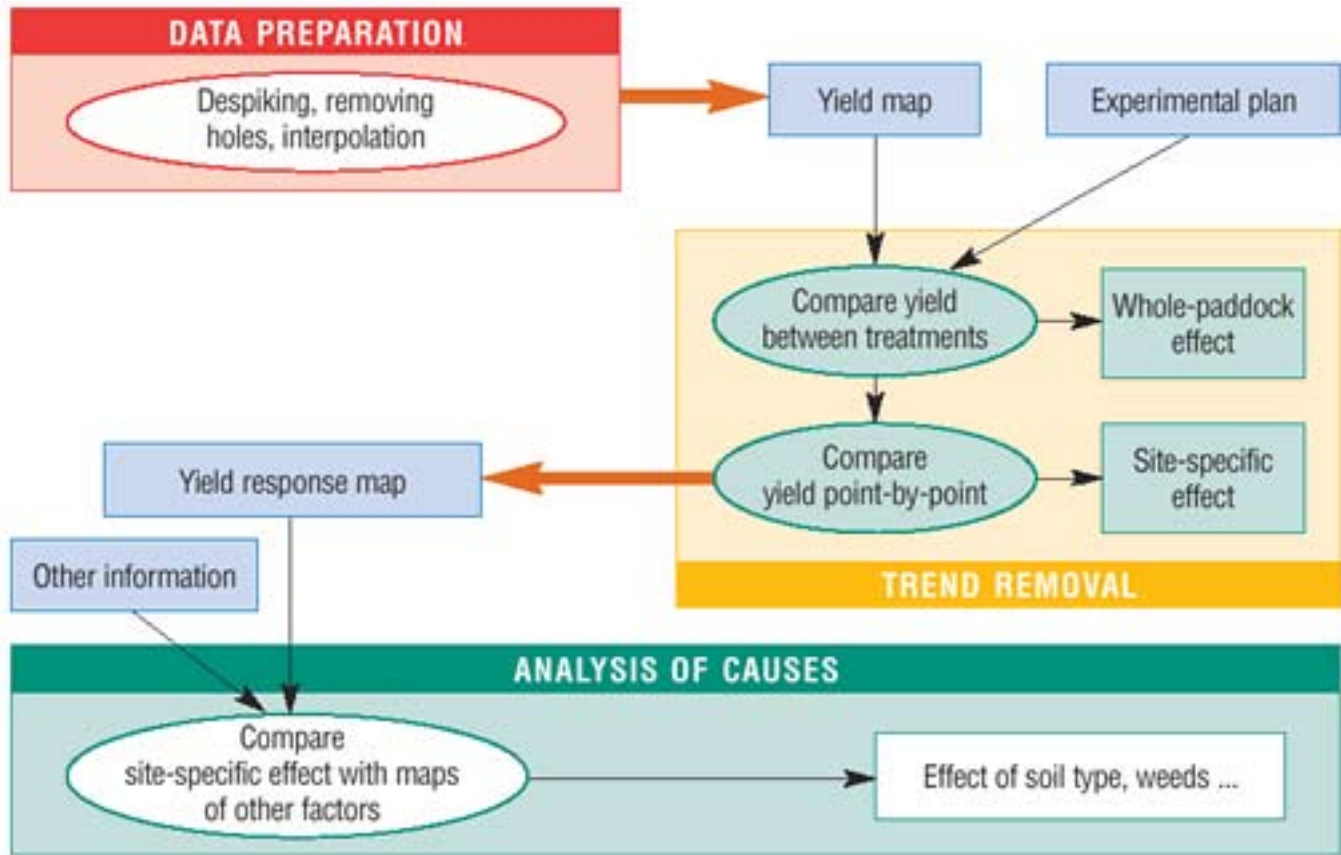


Figure 12: The process of analysis of experimental results.

You may be able to complete the basic analysis of yield data to determine the treatment effects yourself, or you may need help from your adviser or one of the specialist PA service providers. Analysing and comparing your results as soon as practical after harvest is a priority, as they will be a valuable input to the decisions for the next cropping season.

Financial analysis

Of course experimenting, and for that matter the use of precision agriculture in general, is all very well, but is it helping you to make more money than you would using a more traditional approach? Plotting the map in terms of gross margin (as in Figure 1), or return on investment (ROI), really shows how that variation affects your back pocket, and this is usually the thing that will persuade you to try and do something about it.

So a final stage of analysis could be basic financial analysis, to provide information about which parts of the paddock paid their way, and which did poorly. After all, it is not yield which is the real measure of success so much as return on investment.

Gross margin maps are produced by calculating, for each point in the paddock, the price you would get for your crop given the yield at that point, and deducting the costs incurred in producing it. These sorts of analyses are quite easy in GIS.

In addition to gross margin maps, you may find it useful to express your yield map and experimental results in terms of return on investment.

A simple way of estimating ROI is to multiply the yield map by a price ratio. (Note that since this ignores fixed costs, it is a simplification).

An alternative is a map which reviews the results of an experiment to select the level of input which appeared to give the best rate of return from the input. An example of this, from Doug Maitland's chequerboard experiment is shown in [Figure 13](#). It was produced by comparing yield for each level of three treatments, converting for prices and choosing the best one for each point over the paddock. This would have been the best option for this paddock. In fact it shows that applying urea according to this plan, as opposed to the conventional uniform application would have netted Doug an additional \$25 per hectare.

Archiving data

Archiving, like insurance, usually only appears worthwhile in hindsight. But archiving is easy, and will prove its worth in years to come. It is a good idea to ensure not only that you store raw data, but also data in a format you will be able to retrieve later, regardless of possible changes in your monitoring or mapping system. This may require some additional storage and processing capacity, but developments in software and hardware have made it relatively easy to add hard disk storage, tape drives, or CD writers.

Suggestions for establishing a file directory to enable easy storage and retrieval of your files are included in the PA Manual on this CD (see [section 4.2.2](#)).




Figure 13: Optimal rates of urea application. The rates are 0 (red), 40 (yellow), 80 (green) or 120 kg/ha (blue). These rates were determined from the results of Doug Maitland's chequerboard experiment and are optimal in that, assuming a year similar to the one in which the experiment was conducted, they would maximise Doug's return on investment in urea.

Here are some tips:

1. Store digital copies of yield maps you already have in a library using a separate directory for each farm, each harvest.
2. In addition save at least one copy of the original, unadulterated yield monitor data file (for example, a file copy of the SRAM card from an AgLeader® yield monitor). This is your data archive. This original yield monitor data should be saved even if problems were encountered in collecting the data because programmers are still developing ways to restore, clean up, and correct erroneous data points. Typically this involves reassigning yields to correct locations, correcting harvest area measurements, or removing other anomalies in the data. It is usually easier to correct these problems when using original data.
3. Keep a record of file names, file locations and data summaries for each paddock along with other information for the paddock. It may help you interpret the maps later.
4. Save the yield data for each paddock as delimited text files via the yield monitor software's export function

(include latitude, longitude, grain yield, grain moisture as a minimum). While less ideal than working from the unadulterated yield monitor data, it provides a further layer of backup in case the yield monitor data files become corrupt, or the data formats are no longer supported.

5. Consider storing a second backup copy of your data off-site as insurance against disaster.

6. What next?

Doug's map from his chequerboard experiment shows where wheat responded to urea. It shows that some areas responded well, and others poorly. Overall, it shows that the FERTILIZER response was disappointing - but the experiment ITSELF was not, because it gave Doug some useful indicators about how to improve his paddock management.

Evaluating the results

Doug knows that the map provides information that is uncertain. The map says what happened, but does not say what will happen. The paddock is still calling for Doug's management skill and experience to get the best out of it!

What has changed is that for the first time, Doug has hard numbers to work with instead of visual estimates. He knows that in some areas, low yield was associated with low soil pH. He knows that urea produced a negative yield response over some areas. By looking at past results, he will know that some parts of the paddock do badly three years out of four.

The results do not predict - they can only look back - but they can help Doug predict much better.

What caused the variation?

Experimentation is much more valuable when it can point a finger at likely causes. After all, if you know why the crop varied in a certain way last year, you can make a much better guess at how it will perform next year.



After information has been collected it is analysed using mapping software.

There may be plenty of other information to help you decide the cause - such as soil maps, soil test data, field observations and previous yield maps. The art is to work out what additional information you could get during the experiment to help you clinch the matter. This need not mean a major time commitment. Simply ask yourself, "What other information will really sort out one cause from another?". Some examples could be:

- Would a targeted soil test identify sub-surface soil acidity?
- Does the air-seeder put seed too deep or maybe too shallow over certain parts of the paddock?
- Is a recurring yield depression due to weeds? Or could it be disease? Can this be confirmed?
- Could poor response to N or P be related to a micro-nutrient deficiency? Does the crop show any symptoms in the affected area?
- Would a head count help determine haying-off as the problem?

What's likely to happen next year?

The whole purpose of experimentation is to help you make a better guess next season. The more certain you are about next year's events, the better you can manage the paddock. Of course, the biggest uncertainty is the weather. You cannot control it and it injects a huge uncertainty in your results. But that does not mean that the results are invalid - just be aware of their limitations. After an initial shock, certainty will increase with each year's experiment, but if you want absolute certainty, be prepared for a very long haul!

Forming a plan of action

After completing an experiment or two, Doug feels increasingly confident to make one of the following basic statements about the response of his land to investment:

- “I am certain this area responds well;
- I am certain this area responds poorly; or,
- I cannot be certain either way.”

The more Doug is inclined towards the first statement, the more certain he will be, given the normal lottery of season, disease and other mishap, that his land will repay money.

So he will probably be inclined to put effort into identifying what changes to his management are possible, and to spend more on it. Conversely, he will be less inclined to spend the same money on the poor areas. As for the “not sure either way”, these areas will reduce over time as he acquires more detailed observations.

Doug might respond to these options in a variety of ways.

If the paddock is dominated by a responsive area, he might decide to keep the paddock as one and increase inputs over the entire paddock. Conversely, he might decide to reduce inputs, and perhaps repeat an experiment over part of it to confirm his view.

If, as is normally the case, different parts of the paddock behave in very different ways, Doug is faced with the decision whether to treat these differently. With VRT, this is less of a problem as the machinery can operate according to a programable application map.

Changes can be made incrementally over different parts of the paddock, depending on how certain Doug feels about their likely response.

Without VRT the decision becomes more permanent. Should the paddock be split and unresponsive areas dropped? Should some areas be taken out of cultivation completely? These decisions become clearer with each year of yield mapping.

Implementing change

It can be hard to ignore the results of on-farm experiments. The knowledge that substantial parts of a paddock repaid your investment in fertilizer poorly can be persuasive evidence for change. But there are many reasons for poor crop performance.

Only you can decide how certain you feel about the need for change. In some cases, the experimental results may confirm what you already suspected. In others it may be prudent to keep a watching brief until you feel confident enough to act.

An advantage of VRT is changes can be introduced gently. Fencelines can remain as they are, but land managed as if they are not there. Paddocks can be left intact, and rates of fertilizer, seed or spray varied within them, or between them. Small paddocks can effectively be lumped together to be managed as one.

A second possibility is that change can be introduced gradually, depending on how certain you are that it is needed. And of course, it's also possible to put in some small experiments at the same time to help you monitor performance.

Repeating experiments

In many cases, on-farm experimentation will raise more questions. If so, why not put in another experiment? With experience, the process gets easier, especially if you have VRT so that on-farm experimentation may become standard for your paddocks.

But repeating experiments on the same land can be tricky. Residual effects of earlier treatments can confuse the picture, especially with fertilizer experiments, and those involving P or lime in particular. There are several things you can do to reduce this problem:

- Avoid experiments over the same paddock for a year or two;
- Install designs which do not confound one another (for example east-west strips one year and north-south the next);
- Put in additional controls so you can identify the residual effects of the previous experiment; and,
- Run the experiment in parallel with a paddock which has not been experimented with in the past few years and see if you can detect any difference between them.

7. Where to find support

Further reading

Increasingly, the rural press and grains industry magazines are becoming a useful source of information on precision agriculture. The GRDC has published two Groundcover supplements that describe the work underway in its national PA initiative.

Recommended reading of a more technical nature includes:

Adams, M.L. and Cook, S.E. 1997. Methods of On-farm Experimentation Using Precision Agriculture Technology. (ASAE Paper #97-3020. St. Joseph, MI : ASAE).

Bramley, R.G.V., Cook, S.E. and McMahon, G.G. (eds.). 1997. 'Precision Agriculture: What Can it Offer the Sugar Industry?'. 101 pp. (CSIRO Land and Water, Townsville).

McBratney, A.B. 1985. The role of geostatistics in the design and analysis of field experiments with reference to the effect of soil physical properties in the field. p. 3-8. In D.R. Nielsen and J. Bouma (eds.) Soil Spatial Variability. (Pudoc, Wageningen).

Pierce, F.J. and Sadler, E.J. (eds.). 1997. The State of Site-Specific Management for Agriculture. 430 pp. (ASA/CSSA/SSSA, Madison, Wisconsin, USA).

Robert, P.C., Rust, R.H. and Larson, W.E. (eds.). 1996. 'Precision Agriculture: Proceedings of the 3rd International Conference on Precision Agriculture', 1222 pp. (ASA/CSSA/SSSA, Madison).

Robert, P.C., Rust, R.H. and Larson, W.E. (eds.). 1995. 'Site-Specific Management for Agricultural Systems: Proceedings of the 2nd International Conference on Precision Agriculture'. 993 pp. (ASA/CSSA/SSSA, Madison).

Stafford, J.V. (ed.). 1997. 'Precision Agriculture '97: Proceedings of the 1st European Conference on Precision Agriculture', 2 Volumes. (Bios Scientific Publishers, Oxford).

Useful web sites

In no particular order, we recommend the following web sites as further sources of useful information and other contacts.

<http://www.pioneer.com/usa/technology/precision.htm> is the precision agriculture web site of Pioneer Hi-Bred International - contains further information on the split-planter design discussed on page 8 and additional useful information.

<http://precision.agri.umn.edu/> is the Precision Agriculture home-page of the University of Minnesota, a centre of excellence in PA research.

<http://www.usyd.edu.au/su/agric/acpa/pag.htm> is the home-page of the Australian Centre for Precision Agriculture, University of Sydney.

<http://www.kondinin.com.au/> is the home page of the Kondinin Group and provides access to a range of information of use and interest to Australian farmers.

<http://www.grdc.com.au/> is the home page of the Australian Grains Research and Development Corporation.

<http://www.clw.csiro.au/> is the home page of CSIRO Land and Water.

8. Glossary of terms

Amplitude: In a wave experiment, this describes the range of levels of an experimental variable that you are prepared to experiment with. It is defined as half the difference between the maximum and minimum rate. A wave experiment examining response to urea with a zero control and maximum rate of 200 kg/ha has an amplitude of 100.

Analysis of variance (ANOVA): This is statistical procedure that distinguishes between different sources of variation, especially treatment and “noise”. It may be used to establish whether or not statistically significant differences exist between the treatments used in an experiment.

Average: The average (sometimes called the mean) is formally defined as the expected value of something being measured. It is calculated by totalling all the measured values and dividing by the number of values measured.

Bias: This is not a desirable feature of some experimental designs! It occurs when a particular area within the experiment is accidentally favoured relative to other areas, or when a particular treatment in an experiment is given greater prominence than the other treatments.

Chequerboard: A highly replicated experimental design, usually with one variable and four levels.

Controls: The “do nothing” or “zero” treatments that are crucial to the success of most experiments.

Convolution: A term commonly used to describe the smearing of grain yield data over space. It is caused by the mixing which occurs between the time (and position) at which the plant is cut, and the time when it is measured as clean grain.

Deconvolution: The process of “unsmearing” which attempts to separate the mixing effects due to grain processing in the header from “true” yield variation.

Differential correction: A signal received from a fixed and precisely surveyed point which enables the accuracy of GPS to be increased from about ± 50 metres to about \pm one metre.

Donut: Experimental design suitable for growers who prefer “round and round” trafficking to “up and back” trafficking.

Error: The uncontrollable source of variation that is inherent in all forms of measurement.

Geostatistics: A group of statistical procedures which aim to understand variation over space.

Geographical information systems (GIS): Software packages allowing organisation, analysis and illustration of maps, images and points.

Global positioning system (GPS): A positioning system based on an array of 24 satellites which orbit the earth. This is a key technology on which precision agriculture depends since it enables precise spatial location within paddocks.

Inverse distance weighting (IDW): An interpolation method in which unsampled data points are estimated from values sampled at nearby points according to a weighting system based on the distance between the sampled and unsampled points.

Interpolation: A term describing a number of mathematical procedures which are used to create continuous map surfaces from values measured at points.

Kriging: A sophisticated interpolation methodology which, in addition to estimating data values at unsampled points, also gives a measure of the precision of the estimate.

Levels: The rates or amounts of a variable that are used in an experiment.

Mean: See “average”.

Moving window regression: A technique for analysing large, highly replicated experiments occupying quite large areas (ie whole paddocks) in which a “window” is repeatedly placed over the experimental area, the size of the window being defined as some multiple of the smallest area in which all treatments occur.

Noise: See “error”.

Precision Agriculture (PA): A suite of technologies (Yield monitor, GPS, VRT and GIS) which promote improved management of agricultural production by accounting for variations in crop performance in space. Also sometimes called “precision farming”, “site-specific management” or “information-intensive farming”.

Regression: A statistical procedure aimed at identifying cause and effect relationships.

Remote sensing: An all-encompassing term for a range of technologies which allow indirect measurement of something, often from afar. Examples include satellite or airborne imagery, or electromagnetic (EM38) soil survey.

Replicates: Repeated treatments within an experiment.

Residual effects: The effects (normally) of fertilizer application which carry forward from one year to the next.

Response curve: A curve or equation to a curve describing the marginal increase in crop yield obtained as a result of marginal increases in the application of fertilizer, spray or soil amendment.

Sine wave: A wave form that is particularly suitable for generating wave experimental designs.

Soil testing: The process of sampling and (chemically) analysing soils with a view to making informed fertilizer management decisions.

Spectral analysis: A developing methodology for analysing yield variation. This looks for effects in relation to their frequency over space by “tuning” analysis to the wavelength of treatments.

Splines: An interpolation technique in which a (yield) surface is “fitted” to all the points at which measured (yield) values exist.

Strips: The basic building blocks of many simple experimental designs. These may be one-way or two-way.

Synergy: The practice of using VRT for the variable application of more than one input (usually fertilizer and/or seed) to a crop production system.

Treatments: The net result of combining variables and levels. An experiment in which five levels of a single variable is being applied has five treatments.

Variable rate technology (VRT): A new technology which enables agricultural inputs such as fertilizers to be varied continuously within or between paddocks. It comprises one or more of the following components in a spreader or sprayer: variable rate capacity; dGPS; digital map control.

Variables: The thing(s) being varied in the experiment (sometimes called “experimental variables”).

Wavelength: The minimum distance in a wave experiment between points receiving the same treatment.

Waves: Highly replicated experimental designs based on wave patterns, and which require the use of VRT.

Yield monitor: The cornerstone of PA - a sensor installed within a header which enables continuous measurement of yield during harvest. Must be connected to dGPS to generate maps.

Yield response: The marginal increase in yield obtained from application of (usually) fertilizer. See also “Response curve”.

9. Acknowledgments

The majority of the ideas and experimental methodologies discussed in this guide have been formulated in the course of conducting research in collaboration with a number of farmers, researchers and agribusiness people whose input has been greatly appreciated.

In particular, our thanks go to Rinex Technology for instigating the chequerboard trial at Wyalkatchem, to Precision Farming Australia who conducted the trial at Corrigin, and to Morris Industries who enabled the large scale wave trials at Wyalkatchem. Differential GPS has been variously supplied by OMNISTAR and RACAL, and soil analysis provided by Wesfarmers CSBP. The Australian Centre for Precision Agriculture, University of Sydney, provided grain convolution data. However, none of our work would have been possible without the unstinting patience, advice and support of Doug, Noela, Scott and Boyd Maitland at Wyalkatchem.

Production of this guide was funded by the Australian Grains Research and Development Corporation and CSIRO Land and Water. The comments of Ed Blanchard, Belinda Eastough, Doug Maitland, Andrew Wood, Mark Palmer, Mick Poole, Harm van Rees and Nigel Bodinnar on an earlier draft are gratefully acknowledged, as is the assistance of Ms Sonja Slatter (CSIRO Tropical Agriculture) with layout and formatting of earlier drafts. Thanks are also due to Lloyd O'Connell (Australian Grain) for editing.

For more information, please visit:
www.grdc.com.au

