




Design and analysis of evaluation trials of genetic resources collections






IPGRI Technical Bulletins are published by the International Plant Genetic Resources Institute with the intention of putting forward definitive recommendations for techniques in genetic resources. They are specifically aimed at National Programme and genebank personnel.



Design and analysis of evaluation trials of genetic resources collections

A guide for genebank managers

Produced by the Statistical Services Centre (SSC),
University of Reading for the International Plant
Genetic Resources Institute (IPGRI)



The **International Plant Genetic Resources Institute** (IPGRI) is an autonomous international scientific organization, supported by the Consultative Group on International Agricultural Research (CGIAR).

IPGRI's **mandate** is to advance the conservation and use of genetic diversity for the well-being of present and future generations. IPGRI's headquarters is based in Rome, Italy, with offices in another 19 countries worldwide. It operates through three programmes: (1) the Plant Genetic Resources Programme, (2) the CGIAR Genetic Resources Support Programme, and (3) the International Network for the Improvement of Banana and Plantain (INIBAP).

The **international status** of IPGRI is conferred under an Establishment Agreement which, by January 2000, had been signed and ratified by the Governments of Algeria, Australia, Belgium, Benin, Bolivia, Brazil, Burkina Faso, Cameroon, Chile, China, Congo, Costa Rica, Côte d'Ivoire, Cyprus, Czech Republic, Denmark, Ecuador, Egypt, Greece, Guinea, Hungary, India, Indonesia, Iran, Israel, Italy, Jordan, Kenya, Malaysia, Mauritania, Morocco, Norway, Pakistan, Panama, Peru, Poland, Portugal, Romania, Russia, Senegal, Slovakia, Sudan, Switzerland, Syria, Tunisia, Turkey, Uganda and Ukraine.

The geographical designations employed and the presentation of material in this publication do not imply the expression of any opinion whatsoever on the part of IPGRI or the CGIAR concerning the legal status of any country, territory, city or area or its authorities, or concerning the delimitation of its frontiers or boundaries. Similarly, the views expressed are those of the authors and do not necessarily reflect the views of these participating organizations.

Citation: IPGRI. 2001. The design and analysis of evaluation trials of genetic resources collections. A guide for genebank managers. IPGRI Technical Bulletin No. 4. International Plant Genetic Resources Institute, Rome, Italy.

Cover: Testing chickpea accessions for drought resistance at ICRISAT, Patancheru, India. ICRISAT.

ISBN 92-9043-505-4

IPGRI
Via dei Tre Denari, 472/a
00057 Maccarese
Rome, Italy

© International Plant Genetic Resources Institute, 2001

Introduction to the Series

The Technical Bulletin series is targeted at scientists and technicians managing genetic resources collections. Each title will aim to provide guidance on choices while implementing conservation techniques and procedures and in the experimentation required to adapt these to local operating conditions and target species. Techniques are discussed and, where relevant, options presented and suggestions made for experiments. The Technical Bulletins are authored by scientists working in the genetic resources area. IPGRI welcomes suggestions of topics for future volumes. In addition, IPGRI would encourage, and is prepared to support, the exchange of research findings obtained at the various genebanks and laboratories.

Contents

Acknowledgements	5
1 Introduction	6
2 Setting objectives	8
3 Treatments	12
3.1 Terminology—factors in experiments	12
3.2 How many accessions per trial?	12
3.3 Control treatments	13
3.4 Practical considerations	13
4 Sites	15
5 Plots	16
6 Plot layout	16
6.1 General concepts	16
6.2 Blocking	17
6.3 Lattice and alpha designs: evaluating many accessions in small blocks	20
6.4 Augmented designs	23
7 Measurements	26
7.1 Levels of measurement	26
7.2 Measurements at the plant level	27
7.3 Measurements at the plot level	28
7.4 Measurements at the trial level	29
8 Data management	30
9 Analysis	34
9.1 A strategy for analysis	34
9.2 Exploratory analyses	35
9.3 Standard methods of analysis	38
9.4 Further methods of analysis	43
10 Conclusions	46
Bibliography	47
Appendix 1 Selecting the appropriate software	49
Appendix 2 A training strategy	52

Acknowledgements

We are grateful to IPGRI for the opportunity to prepare this material, and in particular to Luigi Guarino, who prepared the detailed specifications, arranged for reviewing and edited the report. The information in the guide has been prepared mainly by Roger Stern (r.d.stern@reading.ac.uk) and Sandro Leidi (a.a.leidi@reading.ac.uk). We are also grateful to other members of the Statistical Services Centre and the Department of Applied Statistics, particularly Carlos Barahona and Fiona Underwood, who prepared draft materials. Thanks also to Robert Curnow and Richard Ellis, who helped with initial discussions on the work, and to Derek Pike, Simon Berry and Alberto Leon, who advised on the detailed structure and needs of the guide. The final version of the guide has benefited considerably from many constructive comments from reviewers and we thank them for the time and interest they have shown in the work.

1 Introduction

This guide is for genebank managers who are considering undertaking evaluation trials on the genetic material in their care. We cover the stages involved in an experimental programme, from the determination of the objectives of each trial to the methods used for the analyses. The coverage can only give general guidelines and managers will need to interpret and adapt them for their particular crops.

The topics covered in this guide are broader than is usually considered to be “statistics”. Traditional statistics often begins with formulae which assume that data have already been collected. It thus concentrates on data analysis. It is our view that the comparative failure of many experimental programmes has been the result of insufficient time being devoted to the planning phases of the research. In particular, the objectives of experiments are often too vaguely stated. We therefore begin this guide with a discussion of how research objectives can be formulated and show how this can assist in defining the measurements to be taken and the analyses to be conducted.

Thus, this guide discusses some of the statistical issues that should be borne in mind when conducting an evaluation trial. Managers also will need to consider practical aspects of the way their crops should be grown. Usually, a compromise between statistical and practical considerations can be found. If they are ever in conflict, however, then practical considerations take precedence over the statistical. In such cases, it is important to revisit the objectives of the trial, to ensure that they can be realized.

Many of the trials that will be undertaken by, or for, genebank managers will have two features in common that set them apart from others. The first is that most genetic resources collections are made up of accessions which are genetically variable. It may therefore be necessary to collect data at the plant level, rather than at the plot level, because knowledge of the average value of an evaluation descriptor for an accession as a whole is not always sufficient.

The second characteristic is that the objective of the experimental programme is usually primarily to highlight promising material in the collection to potential users. Unlike the early stages in a breeding programme, there is no need to “select” certain accessions and “reject” the remainder. Hence the trials are often simply required to report on the genetic materials, rather than to make a strict comparison of the accessions against a known standard. This difference in emphasis simplifies some aspects of the research

strategy that is described in this guide. Sections 2 to 6, on setting objectives to the concepts of blocking, remain the same, as do the ideas of good data management described in Section 8. What is different is that there is increased flexibility in the measurements to be taken and the analysis is simpler. Readers may be comforted to hear that there is little need for formal significance tests if the main aim of the trial is simply to report on the potential of the different accessions.

This document is intended to stimulate discussion with and among genebank managers on how they could be using their genetic materials to the fullest. We have therefore reviewed the stages involved in the design and analysis of a trial, laying emphasis on the topics that distinguish the type of trial that we feel is appropriate for genebank managers. Sections 2 to 7 deal with the planning of a trial. We have described the setting of objectives (Section 2) in rather more detail than the choice of treatments, selection of sites and the type of plots to be used (Sections 3 to 5).

Section 6 is on the layout of the trial. This concentrates primarily on the use of lattices and other incomplete block designs, because there are usually many accessions to be included. Augmented designs are covered in more detail because they offer the possibility of using only a single replicate of the accessions and are not described in many textbooks.

Section 7 is on the measurements to be taken. Here we emphasize particularly the measurements that can be made at the plant level to capture the information about the variability between plants of the same accession. The gathering of information at this level of detail presents some data management problems, which are reviewed briefly in Section 8.

Finally, Section 9 reviews briefly how the data can be analyzed. We concentrate here on the description of a general strategy for the analysis, because the actual processing is easily handled. The main problems in these days of fast computers and user-friendly software are not how to do the analysis, but which analysis is appropriate, given the objectives of the trial, and how the results should be interpreted and presented.

There are many computer packages for the analysis of experimental data, ranging from spreadsheets to very expensive specialist software. Our view on some of the software that is available is given in Appendix 1. Appendix 2 describes how the information that is introduced here can form the basis for short training courses.

2 Setting objectives

The first part of an experimental protocol gives the justification and background to the proposed research. Here we assume that this justification has provided the case for one or a series of trials, and the next part of the protocol gives the objectives of the proposed trial. These objectives must be clearly and precisely stated. So,

“Evaluating the potential of landraces in the collection”

is not a good objective, because it is much too vague. It might, however, be considered as a higher-level “goal”, in the sense that different trials and other information-gathering exercises might all contribute to it.

If you cannot specify objectives precisely then you should question whether your first piece of research should be an experiment. Alternatives are surveys and participative studies. Do you know the precise needs of your clients? For example, the U.S. National Plant Germplasm System (NPGS) gives access to a wide range of information on germplasm at <http://www.ars-grin.gov/npgs/>. Do your clients find that the presentation of the information there, on your crops, is adequate for them to choose accessions? If so, then you could usefully collect the same information and analyze your own trials in a similar way. If not, then what other information would your clients like to help them in their decision-making?

If not enough is known about the clients' needs, then perhaps a first step is an open discussion with clients. This is often called a participatory exercise and can be structured sufficiently formally so that it becomes a recognized component of the research process. If you know your clients' general needs, but require clarification about their priorities for particular crops, then perhaps a questionnaire could be prepared, so the research starts with a survey.

If you decide that your first exercise is a survey, or a participatory exercise, then the details of this guide are not yet for you. However, all studies benefit from a similar level of care in the planning. Hence, if you do not have experience in survey data collection, then seek guidance before you embark on this part of your work.

Preliminary experiments may be useful, even if the objectives cannot yet be specified for a full programme of evaluation trials. These preliminary trials have different types of objectives. They could relate to the training of staff who would be involved in the full experimental programme later. Other preliminary objectives relate to practical ways of taking measurements, e.g. how should

the measurements be taken, how much work is involved, how should samples be taken if some measurements are time consuming. Without such preliminaries, it is easy for scientists to overwhelm field staff, who then have to collect detailed information that will never be used.

Some experimental programmes decide, in hindsight, that the first year will be considered as a “pilot run” or learning experience. It is much better if pilot experiments are planned as such. Managers who are embarking on experimental research for the first time should not consider that planning an initial series of pilot experiments is a failure of some kind. Pilot studies are accepted as normal in other areas, such as survey work. Needless to say, the objectives still have to be stated clearly, even for a pilot study.

We now suppose that the preliminaries are over and you are ready for a full programme of evaluation trials. The objectives of each individual trial must be defined in a way that gives an idea of the size of the experiment, the measurements to be taken (e.g. from a published descriptor list) and where it is to be conducted. A sufficiently precisely stated objective would be as follows:

“Determining the resistance of all XXX races of wheat to disease YYY in semi-arid environments where the minimum temperatures do not fall below ZZZ degrees C.”

An example from the NPGS Web site illustrates one way in which results can be presented. This is given here to emphasize that when a trial is proposed the researchers should already have an idea of the type of presentation of results that they are aiming for. This example considers the resistance of 13 cowpea accessions to aphids. The results are given in Table 2.1 as a frequency table for the degree of resistance. Then there is a table showing the extent of resistance for each of the 13 accessions. Finally, we show the detailed information about one of the resistant accessions.

There is one general point about the statement of the objectives that typifies the trials undertaken by (or for) genebanks and has important implications for their design and analysis. The trials are normally to “evaluate” or to “determine” something and not specifically to “compare” or “find the best”. Thus we may be interested in listing all the accessions that have reasonable resistance to mildew, rather than choosing the best. Most books on the design of experiments assume that the design should be good at comparing accessions, and that the analysis should give tests on whether accessions are “significantly different”. These aspects of design and analysis are of only minor importance here. The difference in

Table 2.1. An example of the reporting results**Codes for APHID of *Vigna oblongifolia***

Code	Definition	No. of accessions
3	RESISTANT	2
5	INTERMEDIATE	5
7	SUSCEPTIBLE	6

APHID	Accession	Plant name
3	PI-284106	C.P.I. 21518
3	PI-299895	661
5	PI-322304	IRI 2138
5	PI-322343	IRI 2052
5	PI-354915	
5	PI-365092	B/53/332
5	PI-365093	DALRYMPLE
7	PI-181585	
7	PI-276474	C.P.I. 17855
7	PI-292872	No. C36-305
7	PI-300176	412
7	PI-305072	4
7	PI-352988	TVu 2836

PI 299895

Vigna oblongifolia var. *oblongifolia* FABACEAE

Collector identifier: 661

Maintenance site: Southern Regional PI Station (S9). NPGS received: 31-Aug-1964.

Inventory volume: 172. Form received: Seed. Accession backed up at second site.

Record entered: 09-Aug-1994.

Accession names and identifiers

661

Type: Collector

Availability

Material is available for distribution. The normal amount distributed is 50 seeds.

Narrative

Seeds

Source history

Type: Collected. Date: Feb-1964. From: South Africa.

Locality: Stutterheim Agricultural Research Station, Stutterheim Cape Province

Cooperators:

1.Oakes, A., USDA, Germplasm Resources Laboratory.

emphasis does not lead to a major change in the types of design that are proposed, but it does simplify the way the data are analyzed and presented.

Finally, it is useful to remember that managers do not have to undertake all their evaluation work themselves. They can commission some of the research from others. However, this does not absolve the managers from needing to understand the concepts in this guide. It is even more important that the objectives and all other details of an experiment are carefully specified if others will undertake the work. Otherwise the work will either not be done well, or will be done for the objectives of the group undertaking the work, rather than those of the genebank.

3 Treatments

3.1 Terminology—factors in experiments

First we need to clarify the meaning of the following basic terms, which are used throughout this guide and which are sometimes confused:

- treatments
- factors
- levels

We consider three illustrative examples:

1. A trial that evaluates 24 accessions.
2. A trial that evaluates 8 accessions under each of 3 different fertility regimes.
3. A trial that evaluates 4 accessions at 3 levels of spacing, for 2 planting dates.

These three trials all have 24 treatments. In the first, there is just a single factor, accession, which has 24 levels. Thus, in this simple case, whether we think of the different accessions as the *treatments*, or the *levels of a treatment factor* makes no difference.

In the second experiment there are 2 factors, namely accession, with 8 levels, and fertility, with 3 levels. Each treatment consists of the combination of a particular accession and a particular fertility level. Thus, there are 24 different combinations, or treatments. This is sometimes known as an 8 by 3 factorial treatment structure.

Similarly, the third trial has 3 factors and the 24 treatments are arranged in a 4 by 3 by 2 factorial treatment structure.

In this guide we concentrate on the first type of trial and assume that the goal is simply to evaluate different accessions. If managers wish to conduct trials where there is more than one treatment factor, they (6.2) that are not covered here.

3.2 How many accessions per trial?

The statement of the objectives should include an indication of the number of treatments—here the accessions—to be included in the trial. There is no prescribed limit for this. As in the early stages of a breeding programme, some trials may have many hundreds of accessions, each in small, perhaps single-line, plots.

Within a given site, there is sometimes a choice between putting all accessions in a single large trial or having a number of smaller ones. Here the guideline is to put them in the same trial **only if they all need to be compared with each other**; otherwise they can be distributed among a set of smaller experiments.

For example, if there are accessions that are known to be in different maturity groups and recommendations are required for each of a range of season lengths, then there is no requirement to compare short with long season varieties. They can therefore be included in separate trials. In contrast, disease resistance could be present in accessions of any season length, and it will then be more appropriate to evaluate all the accessions together. Even in this latter case, if clients expect recommendations at each of different season lengths, then a set of smaller trials should be used. If, however, they are likely to return with demands such as “*How good is the resistance of the best long-season accession relative to the accessions recommended from the short-season trial?*”, then it would have been better to have had all the accessions together in the same trial.

3.3 Control treatments

In addition to the accessions being evaluated there will often be one or more standard lines that are considered as “controls” or “checks”. Their presence and the way they are incorporated in the trial are determined by the objectives. For example, a trial on resistance to a given disease might include three controls: one resistant, one tolerant and one susceptible. If, however, there is interest in highlighting accessions that are highly resistant, then the only control might be a well-known resistant variety.

The controls are sometimes replicated more often than the other accessions. This is considered further in Section 6 on blocking. One example is so-called “augmented designs”. In these there is often only a single repetition of the tested accessions, with multiple repeats of one or more control lines.

Some trials may need controls simply as part of the “environmental”, or “site”, information. In such cases, the controls might not be in plots for formal comparisons with the accessions in the trials. They might be planted in guard rows or separate plots that are of a different size from the other accessions.

3.4 Practical considerations

One major practical concern in a trial with many accessions used to be whether it was possible to analyze the data in the first place. This is no longer a problem as modern statistical packages impose no limits on the number of treatments in an experiment. However, large trials may be more difficult to manage and there is sometimes a danger that the large volumes of data that are collected may be overwhelming to field staff, resulting in data of lower quality than would be the case with smaller trials.

Large experiments need large areas of land. They also often have a more complicated blocking structure (see Section 6). If this large area is quite heterogeneous, but a part of it is more homogeneous, then better (i.e. more precise) results may be obtained from a trial that uses only the homogeneous land.

4 Sites

When the objective of the trial is simply an assessment of the potential of different accessions, the trial will probably be carried out in an “ideal” environment, for example one that is managed so that there is no water stress or competition from weeds. For some objectives, however, for example relating to response to stress, the choice of sites is crucial and trials may be repeated over a range of sites and years. There is a voluminous literature on “genotype by environment” interaction, or the phenomenon whereby different genotypes react differently as the levels of a treatment factor change, leading to different rankings of the genotypes at the different factor levels.

This highlights the importance, for these objectives, of conducting trials in a range of different “environments”. When multisite trials are conducted, the data management aspects (considered in Section 8) become even more important and there are many alternative methods of analysis of the combined data. These issues are beyond the scope of this report, but they may be sufficiently important to be included in a training programme.

In contrast, for some disease studies, the trials may be best laid out in large pots in a greenhouse, and the results may be relatively independent of the actual site being used. The key point here is to ensure the presence of high and evenly spread disease pressure.

It is vital that information on the site be recorded and made available with the other results. This aspect is discussed in more detail in Section 7, which deals with taking measurements.

5 Plots

Plots will often be small, partly because sowing material is likely to be scanty. They are often a single row. In pot experiments, there may be just a few plants per pot. Often there are no guard rows, because it is reasonable to assume there is no interplot competition. If this assumption is not tenable, then the random allocation of the test materials can be restricted by placing together those accessions likely to have the same phenotypic characteristics (David *et al.* 1996).

6 Plot layout

6.1 General concepts

In this section we consider the layout of a trial within a particular site. To explain the main concepts we take a simple example with just six accessions: A, B, C, D, E and F. There are two main decisions to make:

1. the number of plots to be sown with each accession, i.e. the number of “replications”
2. how these replications of each accession will be placed in the field, i.e. the “blocking” to be used.

Replications, and how they are distributed within experimental layout, are important because they can be used to control




Fig. 6.1a. A Randomized Complete Block Design before randomization.




Fig. 6.1b. The same design after being randomized.

experimental error, due both to variation inherent in the material being tested and to variation in the site where the material is being tested.

Figure 6.1 gives a simple example, where a total of 18 plots is used, and each accession has been replicated three times. The layout is in 3 blocks, with each block containing 6 plots. Thus each block contains one replicate. This is a very common design called the “randomized complete block design”, or RCBD for short. It is simple in its layout and easy to analyze. The plan is shown in Fig. 6.1a before randomization, and Fig. 6.1b after being randomized.

In Section 6.2 we consider the importance of blocking. A block is often thought to be synonymous with a replicate, because of the popularity of the randomized complete block design. It is important to understand the difference between blocking and replication, because many experiments envisaged for genebank accessions will not be in RCBD designs. This is because when there are many accessions to be evaluated in each trial, the RCBD design is relatively ineffective as a device for the control of error variation. This concept is discussed further in Section 6.2, where the subject of incomplete blocks is introduced.

A popular design for testing accessions is called a “lattice”. Its use is described in Section 6.3. When lattices are used for genetic resources evaluation, they will often comprise just two replicates, to maximize the number of accessions that can be evaluated on a given area of land. In the same section we discuss alpha designs, which are an extension of lattices to blocks with a different number of plots.

In Section 6.4 we consider the “augmented design”, a type of design that allows land to be used even more efficiently. Augmented designs have just a single replicate of the test accessions. They are therefore of particular value when there is a shortage of seed for the accessions, or of land. Augmented designs also include one or more check varieties and these are replicated more than once in the experiment. We believe that augmented designs may be of particular use for germplasm evaluation. They are not common, perhaps because they have not been described in detail in the standard literature. Hence a special reference section has been included.

6.2 Blocking

The purpose of blocking is to group plots within a part of the field that is as homogeneous as possible. This enables evaluation of accessions with greater precision than if the position of the plots were not restricted in this way.

Using the simple example above, with 6 varieties, we suppose that a small experimental field can accommodate 18 plots and that the soil has an inherent fertility gradient that changes smoothly from left to right. The trial could therefore be laid out with 3 blocks of 6 plots each, along the fertility gradient, as illustrated in Fig. 6.1.

In practice, it is often necessary to form small blocks with fewer plots per block than the number of accessions. This may be due to heterogeneous field conditions, or because there are many accessions to evaluate. If the soil fertility in the field were very patchy, then a possible approach to try to preserve soil homogeneity of plots within blocks would be to halve the size of blocks from 6 to 3 plots, as shown in Fig. 6.2.




Fig. 6.2. An incomplete block design.

The blocks are now “incomplete” as each contains only 3 of the 6 accessions. Blocks and replicates are no longer equivalent, since there are still 3 replicates per accession, but 6 blocks.

In some trials it is useful to compare the performance of the test accessions with that of control varieties. This could be accommodated within Fig. 6.2, if one of the labels refers to the control. An alternative is to put one or more control varieties into each block. An example is shown in Fig. 6.3.




Fig. 6.3. An incomplete block design with an added control.

In the plans depicted in Figs. 6.2 and 6.3, we have relaxed the condition (from Fig. 6.1) that a block has the same number of plots as there are accessions. We now take this relaxation one step further. As the main aim of blocking is to design trials where there is little within-block heterogeneity, it is sometimes useful if blocks are of different sizes. This refinement is sometimes built into the planning stage, because homogeneous areas of land are not always of the same size. Sometimes it occurs during an experiment, because of a failure with some of the accessions. Figures 6.4 and 6.5 give two examples.




Fig. 6.4. An incomplete block design with different block sizes.

Figure 6.4 typifies the situation where accession B failed completely. Perhaps it was destroyed by disease. This is not exactly a missing value and the results on accession B should still be reported. However, as far as the formal analysis is concerned, it is the same as if the accession were not part of the same trial. Hence the analysis proceeds with the blocks of different size. Thus, blocks I, III and VI have 3 plots and II, IV and V have 4 plots each. In the situation typified in Fig. 6.5, part of the field has been lost, perhaps due to waterlogging.




Fig. 6.2. An incomplete block design with a problem in part of the field.

The set of plans typified by Figs. 6.2 to 6.5 emphasizes the flexibility that is possible in designing trials that are appropriate for particular situations. Thus:

- checks may be replicated differently from the test accessions
- test accessions may have different levels of replication
- blocks may be of different sizes.

There is no difficulty with the analysis of data from any of these trials. In the pre-computer age, the analysis of data from incomplete block designs was difficult, but this should no longer deter anyone from their use.

In this section we have discussed the principles of incomplete block designs, because genebank managers will often have too many accessions in a single trial for the use of the randomized complete block design to be recommended. In the next section we look at a special case of an incomplete block design, called a lattice, which remains popular with breeders and should be useful for the types of trial that a genebank manager might be interested in carrying out.

6.3 Lattice and alpha designs: evaluating many accessions in small blocks

Lattices are special cases of incomplete block designs. Here we just consider square lattices, where the number of accessions is a perfect square, for example 9, 16, 25, 144 or 900. In a square lattice, the block size is fixed as the square root of the number of accessions. So, with 900 accessions the blocks would be of 30 plots each. Thus lattices provide simple designs for situations where there are many accessions and blocks are reasonably small compared with the size of the trial. Their analysis is slightly simpler than the general incomplete block designs, though this is now of little concern. Lattices are limited in the range of situations in which they can be used, compared with the general incomplete block designs described in Section 6.2. We consider how to address these limitations at the end of this section.

To illustrate using a lattice we take a very simple case. Suppose that 9 accessions are to be evaluated. We continue with the situation, depicted in Section 6.1, where we have 18 plots and this permits us to use 2 replicates in a 3x3-lattice arrangement. We think it is likely that lattices with 2 replicates will be a common design. Fig. 6.6 shows a possible design, prior to randomization.

Some readers may wonder what is special about a "lattice" compared with the general incomplete block designs discussed in Section 6.2. For a brief explanation, note that accession A is in the same block as accessions B and C in the second replicate and with

Block	Replicate 1			Replicate 2		
	I	II	III	IV	V	VI
	A	B	C	A	D	G
	D	E	F	B	E	H
	G	H	I	C	F	I

Fig. 6.6. A 3x3 lattice.

two different accessions, namely D and G, in the first. For comparison, in the randomized block design of Fig. 6.1, each accession was in the same block as **all the other accessions** in each replicate. In Fig. 6.6, if there were 4 replicates, rather than just 2, then accession A could be in the same block as **each of the other 8 accessions** in **just one** of these replicates. So could all the other accessions. This is therefore “fair” to all the accessions and would give a “balanced” design that is quite easy to analyze. The plan above, with just two replicates, is called a “partially balanced” design and is also not difficult to analyze.

In practice an experimenter may design a simple lattice, such as is shown in Fig. 6.6, and then find that some of the complications depicted in Section 6.2 occur. Fortunately, this is no longer a serious problem. Current methods of analysis often do not take advantage of the relative simplicity of a lattice design, and hence are equally able to analyze situations where there are some complications.

It is also possible to adapt standard lattice designs where checks are to be included and it is thought to be appropriate to replicate the checks more than the test accessions, as was described in Section 6.2. Figure 6.7 gives an example, again before randomization, where a check has been added to each block. In this example, there are therefore 6 replicates of the check variety and two of each of the test accessions.

Block	Replicate 1			Replicate 2		
	I	II	III	IV	V	VI
	A	B	C	A	D	G
	D	E	F	B	E	H
	G	H	I	C	F	I
	Z	Z	Z	Z	Z	Z

Fig. 6.7. A 3x3 lattice with an added control (before randomization).

It is clear that very large numbers of accessions can be evaluated by using lattices. However, lattices lack the flexibility required in many practical situations. This can lead to a different number of accessions from that which the researcher would wish to use and enforces a rigid block size that may not be appropriate to local field conditions.

The development of a more general class of designs called “alpha designs” has removed these restrictions. Computer-aided design makes it possible to produce alpha designs that are flexible enough to accommodate a large number of accessions with fewer replicates than the number of blocks, and also blocks of different sizes, i.e. containing different numbers of plots.

As an example, Fig. 6.8 shows a design with 36 plots and 18 accessions. We suppose that field heterogeneity at the proposed experimental site is such that we would like 6 blocks with 4 accessions, and 4 blocks containing 3 accessions. For good measure, we have also chosen to add a check variety to each block.

		Replicate 1				
Block		I	II	III	IV	V
	A	B	C	D	E	
	F	G	H	I	J	
	K	L	M	N	O	
	P	Q	R	Z	Z	
	Z	Z	Z			
		Replicate 2				
Block		VI	VII	VIII	IX	X
	A	D	E	B	C	
	G	J	F	H	I	
	O	M	N	K	L	
	R	P	Q	Z	Z	
	Z	Z	Z			

Fig. 6.8. Example of 18 accessions evaluated in an alpha design with blocks of different sizes.

So far, these designs still assume that the test accessions are replicated within the trial. Where there is little seed for some of the accessions, then the concepts of incomplete blocks allow for unequal replication of the accessions, as was shown in Fig. 6.5, and this can accommodate the fact that there may be just a single replicate of some of these accessions.

In Section 6.4, we extend this idea, and consider designs where there is just a single replicate of most, or all, of the test accessions. Such designs permit assessment of a very large number of accessions on a relatively small area.

6.4 Augmented designs

Augmented designs are appropriate for evaluation stages when hundreds or even thousands of accessions are being studied in the same experiment, using a limited amount of sowing material, perhaps enough for one replicate only. They cope with environmental heterogeneity by placing replicates of controls systematically in the experiment. As control plots may be said to keep a check on environmental variation, they have been called “checks”, and in this section we use “check” for “control”.

It is rare for patterns of heterogeneity in soil fertility and disease pressure of experimental sites to be known in advance. Thus if many accessions are grown in unreplicated plots, some external means of local adjustment is required to assess and possibly to adjust plot means for any environmental variability across the trial site. The usual method is to arrange replicated check plots in a systematic pattern. So replicated checks of established varieties measure the variation in a trait across the trial and the value of the trait for the unreplicated accessions can be assessed against its value in adjacent checks.

An example of an augmented design is given in Fig. 6.9. This shows 45 plots, of which 15 have been allocated to the check varieties. Thus, although one-third of the experimental area is occupied by checks, as many as 30 test accessions can be evaluated in this trial, with the checks providing a means of adjustment for environmental variation. If each plot area is 5x1 m the net experimental area is just 225 m². The proportion of plots occupied by checks is normally about 15–20% in larger experiments, including perhaps 1000 accessions.

In the design shown in Fig. 6.9, the 45 plots are divided into 9 sets with 5 plots in each. Within each set of 5 plots, the central plot is used for 1 of the 3 check varieties and these 9 checks are laid out in what is called a “Latin square”, the distinguishing feature of which is that each entry appears in each row, and in each column. In Fig. 6.9 the middle plots of each block of 5 are in the Latin square arrangement. The analysis of the data from the check plots provides a system for adjustment and measurement of precision of the test accessions.

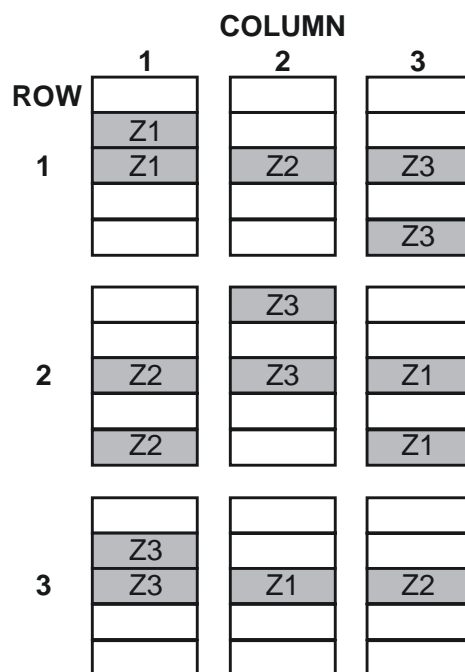


Fig. 6.9. A 3x3 square arrangements with a total of 15 checks (shaded).

Thus the checks act as baselines against which to compare accessions and they also allow a certain degree of extrapolation to the performance of accessions in other environments where the performance of the same check variety is known.

The main concept discussed in this section is the desirability of trials that include only a single replication of some or all of the accessions. The example above is one possible layout that can be used, but there are others. Instead of grouping the plots into blocks as is shown in Fig. 6.9, we could consider the use of adjacent check plots for the adjustment of each unreplicated entry. This removes almost all restrictions on the arrangement of plots in any precise shape over the area of the trial. That is, square or rectangular layouts as illustrated in Fig. 6.9 are no longer necessary. An attractive system combines the ideas of incomplete block designs, which would be used for the checks, with single replicates of the test accessions.

Using 2 check varieties (Z1 and Z2) to adjust the single replicates of test lines, as many as 1560 winter wheat lines were evaluated at Plant Breeding International in Cambridge, UK (Besag and Kempton 1986), in a rectangular field of just over 2 ha split into 1.5x4.5 m plots. In this trial, 16% of the area was allocated to the check plots. Figure 6.10 illustrates the layout, with lower case letters standing

for test accession and upper case letters for those test accessions whose raw record is to be adjusted using the 6 adjacent check plots.

a	B	Z1	c	d	e	f	g	Z2	h	I
j	K	Z2	L	M	N	O	P	Z1	q	R
s	T	Z1	u	v	x	y	w	Z2	aa	Bb

Fig. 6.10. Field layout of part of the experiment from Besag and Kempton (1986).

A single-replicate design such as illustrated in this section should be of considerable interest to genebank managers. However, there are two issues that require further discussion. First, all the examples we have found in the literature refer to yield-related variables. It is not clear whether the same method of adjustment can be used for the other types of traits that would normally be measured in the evaluation trials. Second, the adjustments assume that the accessions react in a similar manner as the checks to the heterogeneity of growing conditions, but this may well not be the case.

7 Measurements

7.1 Levels of measurement

In this section we distinguish among measurements that are taken at three different levels:

1. **Individual plants.** A key feature of many genetic resources evaluation trials is that they involve genetically heterogeneous material, i.e. there is variation among the plants in each accession. Hence some data, for example on plant height, will be needed at the “plant level”, i.e. collected on a number of individual plants in each plot, to summarize the information from accessions fully.
2. **Plots.** Each accession is planted on a small piece of land or sometimes in a pot. Some measurements are normally made at this “plot level”, e.g. by recording a single value for the whole plot, as with yield.
3. **The trial site as a whole.** To interpret the results of a trial it is important to provide details of the environment within which the trial was conducted. Thus, we also collect data at the site or “trial level”, for example site location, date of sowing and rainfall.

If there are a number of replicates for each accession, then the results may be **reported** as averages for the accession. We could consider this as an additional “accession level”, but that is not a level at which we take measurements. So, the accession level is part of the analysis, rather than a feature of the taking of measurements, and is not considered in this section.

We distinguish between the three levels at which measurements can be made in the subsections below. For some measurements, such as plant height or rainfall, it is obvious at what level they will be recorded. An important part of the planning concerns the way to record the information that could be from more than one level. For example, soil characteristics are measurements that are usually given for the trial as a whole, i.e. at the trial level. However, it is possible, though more expensive, to provide this information at the plot level too. Similarly, disease severity can be collected at either the plot level or from individual plants.

The decisions on what information to collect at the plant level are particularly important, and this is therefore one of the longer sections in this guide. Recording data on individual plants is very time consuming, and this may be time wasted if we then simply average the values for the whole plot. However, just taking

measurements at the plot level (which is all that is needed for most agronomic experiments) may not provide breeders with the detailed information they require on accessions that contain promising but diverse material.

7.2 Measurements at the plant level

Measurements may be taken at the plant level for three main reasons:

1. Information is needed on a plot basis, but there is no alternative but to take the measurements at a plant level. For example, with some plants, even if we are only interested in an overall (average) disease score for the plot, the only effective way of measuring is first to give the score of each plant, and then take the average. With some variables there may be a choice in the level of measurement. For example, to measure the average leaf area per plant in a given plot, one way would be to strip the leaves of each plant separately and record the area for each plant. The second would be to strip the leaves of all the plants and just give a total leaf area. Then this total, divided by the number of plants, gives the average. The second method is less work and is therefore to be preferred, unless the data are required on a plant basis.
2. The plant level information is required because the average or total value per plot does not provide sufficient information to evaluate the accession. If we are calculating disease score, then we may be more interested in the proportion of plants with a score of 3 or less (i.e. fairly healthy) than in the mean.
3. The plant level information is required because we wish to study the relationship between different measurements on the same plant. For example, we might find that 15% of the plants of an accession produce seed and 20% have dark green leaves. Perhaps the key question is how many plants produce seed and also have dark green leaves. This can only be found by a simultaneous examination of the results on each plant.

In general terms the argument against devoting too many resources to measurements at the plant level is that they are time-consuming and hence expensive to collect. Also there is no point in collecting data that will not be used. The accessions are sown at the plot level and hence data should be recorded and analyzed at the same level.

In many agronomic experiments there is little within-plot variability and little, if any, of the data are collected at the plant level.

However, in the trials envisaged here, we expect considerable within-plot variability and there is therefore a stronger case for the collection of these detailed observations.

Economies can sometimes be made by limiting the detail that is recorded for each plant. For example, if the objectives of an experiment relate to the detection of accessions containing plants that show resistance then it may not be necessary to use a detailed scale of measurement, say a 1 to 9 scale. A smaller 3-point scale, recording the disease level as just 1, 2 or greater than 2 may be sufficient and much quicker to record.

Set against this economy is the need to record on a standard scale and the fact that, if the requirements were to change, the chosen scale might not be informative. In the example above, the scale with 1, 2, 3, 4 and greater than 4 might later prove to have been more desirable, because there was high disease pressure and hence virtually no plants with a score less than 3.

7.3 Measurements at the plot level

The treatments, i.e. the different accessions, are normally sown at the plot level. Hence, as the objectives relate to the different accessions, this is the obvious level at which to take measurements. It is also the level at which most of the formal analyses will be undertaken.

One obvious measurement is the number of germinating plants per plot. This is a useful variable in its own right and is often also of use when analyzing other, subsequent measurements. Also measured at this level are all traits that are similar for the plants in an accession.

Sometimes, observations are made on individual plants, but the recording is only made at the plot level. For example, individual plants may be observed for evidence of disease, but only a single score, say between 1 and 9, is recorded for the plot as a whole. Alternatively, we could simply record the number of plants with a disease score of 3 or less. Note the difference between recording the disease score at the plot level and that of recording the disease score of each plant, or of a sample of plants, and then taking the mean score. The latter case is a recording at the plant level, and is discussed in Section 7.2.

We might also take measurements of just the “best” plant in the plot. This could be the height, or number of tillers of the largest plant in the plot. It could be the disease score of the healthiest plant. This might be useful in addition to the disease score of the plot, because that might indicate the disease pressure. We could also include the disease score of the most diseased plant. Information on the best and worst, taken together, gives an idea of variability of the plants in the accession.

Measurements on the best plant might be useful because we consider it to be an appropriate way of indicating the potential of the accession. If we are comparing accessions, then a comparison is not easy to make in a fair way if different accessions have very different germination rates. However, strict comparisons are not usually a major objective of these trials, which are more often to investigate “potential” in a breeding or other use programme. For this purpose, the information on the best plant per plot may be appropriate.

7.4 Measurements at the trial level

In this category we include all observations that are made to characterize the trial as a whole. For example, its location and other site particulars would be included here. Basic environmental data on the trial location should in fact be recorded routinely. The possibility of genotype by environment interaction highlights the fact that the relationship among different genotypes may depend on the particular environment in which the trial was conducted. Interpretation and use of the results therefore need some indication of the particular environment under which the trial was carried out. Also included here could be indications of the methodologies and practices used for collecting data, and the people involved.

One experimental strategy is to repeat experiments in a variety of different environments. An alternative is to use a limited set of differing environments and then use a crop model to estimate the responses in other environments. An example of such a model is RoDMod (Watkinson *et al.* 1994). This is a rate of development model to characterize genotypic variation in flowering responses to photoperiod and temperature. These experimental and modelling approaches actually complement each other. Both need environmental information to exploit the experimental results fully.

Often, control or check accessions are part of the trial, and measurements on the controls are therefore made at the plot level. If, instead, checks are included just to characterize the environment (and not to compare formally with the other accessions), then they may be grown on plots of different size, or in guard rows. The results would then be recorded at the trial level.

8 Data management

In this section on data management we discuss two topics:

1. the checking and entry of the raw data into a computer
2. the subsequent organization of the records for analysis.

Experimental data are often poorly managed and a small booklet entitled *Data Management Guidelines for Experimental Projects* (SSC 2000) gives some information. If the data are well managed then a genebank manager should be able to supply **all** details of an experiment to a client, without any extra work. By full details we mean all the protocol information, plus all the raw data at the site (or experiment) level, the plot level and the plant level, i.e. at whatever level they were recorded.

In setting this requirement as an indication of good data management we are not assuming that all the raw data will necessarily be supplied on demand. It is the *capability* to supply which indicates that data management is adequate.

In Section 7 we showed that many evaluation trials could include the collection of data at the plant level. In Section 9 we show that most of the analysis will be at the plot level. There is then the question of how the data should be entered into the computer. In the past, in such situations, field officers have sometimes had the task of calculating the summaries at the plot level with a calculator. Then the data entry is in the appropriate form for analysis.

Though simple, this strategy is incorrect. If data are collected at the plant level, then they should be entered into the computer at this level. This is what is meant by the raw data. If data are recorded straight into a hand-held computer in the field, then the entry of the **raw** data into the computer is automatic. Otherwise, they should be typed directly from the field record form. Once entered, the computer can be used to calculate the means (or any other summary statistic) and to organize the data in the correct form for the analysis.

This stage is illustrated with an example from Gomez and Gomez (1984). This was a trial with 8 varieties of rice and 3 replicates. Data were recorded on the number of panicles per hill, on a sample of 12 hills in each of the 24 plots. Table 8.1 shows the data as given on page 547 of that book, while part of the data, as entered into Microsoft Excel[®], are shown in Table 8.2.

These data would normally be entered onto separate sheets of a spreadsheet workbook. At the plot level, only the three columns—called Plot, Rep and Variety—have been entered. In real examples there would often be more columns of data to be entered that were collected directly at the plot level.

Table 8.1. Example showing a textbook arrangement of data collected at plant or hill level

Variety	Number of panicles per hill		
	Rep. I	Rep. II	Rep. III
IR22	5, 8, 12, 14, 10, 10 6, 10, 8, 11, 11, 8	10, 13, 10, 13, 11, 11 12, 5, 10, 7, 14, 5	7, 6, 11, 10, 7, 8 8, 8, 10, 10, 6, 11
IR160-27-3	11, 11, 11, 12, 4, 12 8, 14, 8, 7, 9, 9	13, 4, 4, 7, 5, 7 11, 8, 7, 8, 10, 9	8, 7, 9, 10, 5, 5 9, 10, 4, 9, 12, 11
BPI-76-1	4, 5, 8, 5, 8, 4 5, 9, 6, 6, 7, 10	6, 8, 4, 5, 6, 10 8, 3, 7, 8, 7, 11	8, 7, 6, 5, 6, 7 6, 8, 6, 6, 5, 4
C4-63	8, 10, 9, 7, 9, 7 9, 13, 13, 5, 7, 5	9, 7, 9, 5, 8, 9 8, 10, 6, 5, 6, 5	8, 10, 7, 6, 7, 6 9, 8, 6, 4, 5, 7
RD-3	7, 12, 7, 11, 12, 7 7, 6, 5, 9, 8, 9	9, 7, 6, 8, 4, 8 8, 9, 8, 9, 6, 7	9, 3, 4, 6, 5, 3 9, 7, 9, 6, 6, 7
IR480-5-9	7, 7, 6, 11, 7, 8 8, 8, 9, 6, 4, 14	8, 10, 7, 6, 8, 8 10, 5, 7, 5, 8, 7	7, 6, 9, 7, 11, 8 12, 7, 8, 9, 8, 9
Jaya	8, 9, 12, 7, 7, 3 10, 10, 8, 7, 9, 8	8, 6, 7, 8, 9, 9 14, 8, 9, 11, 6, 7	10, 4, 8, 9, 4, 6 7, 4, 3, 4, 4, 6
IR20	5, 5, 10, 9, 7, 5 9, 10, 9, 6, 12, 8	8, 8, 8, 3, 13, 13 7, 12, 9, 9, 8, 11	5, 12, 10, 9, 7, 9 8, 7, 5, 8, 10, 7

At the hill level, the data are entered giving the hill number, the plot and the number of panicles. The remaining columns give the minimum, mean, maximum and standard deviation of the number of panicles per plot. They have not been entered, but rather have been calculated. This is done using either the tabulation facilities in the spreadsheet or equivalent facilities in a statistics package. These calculated values are now ready for analysis at the plot level.

If further summaries are required, then they can be derived as needed. For example, the analysis in Section 9 indicated the need to process the data on the number of hills with 10 or more panicles, within each plot. These were calculated from the hill-level data and are in the last column of Table 8.2.

We emphasize the necessity of a well-defined data entry and management strategy because it has so often been a weak point in agricultural research. As a simple example, consider the entry of the dataset given above. In the entry of the 288 values at the hill level a typing error originally gave a count of 1 for hill 5, rather than the correct value of 10. This was easy to spot and correct. Had it not been corrected at this stage, it would probably have been spotted later, when the minima were calculated in Table 8.2, or in the exploratory analysis, to be described in Section 9.

Suppose, however, that the mean count per plot had been done with a pocket calculator and the hill-level data had not been entered into the computer. The mean is now 8.7 rather than 9.4, and the mistake probably would not have been spotted.

The International Centre for Agroforestry (ICRAF) has been working on the development of a management system for research data, called Logbook. This will be tested on data from agroforestry trials in 2000. It is, however, potentially more general and could be investigated as a possible system to support the management of the raw data from all kinds of germplasm evaluation trials. Logbook does not introduce new software, but rather provides a system for users who are not database experts to exploit standard software such as MS Excel or Access[®] more effectively.

One reason for emphasizing the importance of the Logbook system is that it has the potential to manage disparate sets of data. Thus, data from different levels in the same trial, from a series of different trials, plus other research information from a survey or participative study, can all be managed together. This facilitates the use of combined information about any particular crop.

9 Analysis

9.1 A strategy for analysis

The first step in the analysis is to identify the tables and graphs that will present the results of the trial clearly. An example was given earlier, in Section 2, where the results were shown for a study of aphid resistance in cowpea. These results gave:

1. a frequency table of the resistance scores for the 13 accessions
2. the (mean) score for each accession
3. a link to the detailed history and information of accessibility of each accession.

We are not implying that you will present the final results in the way you expect initially, because the analysis of the data may indicate that a different presentation is needed or is more desirable. But it is useful to identify initial objectives for the analysis and presentation.

Table 9.1 presents a more complicated layout of a table that might result from a study of five accessions on disease resistance.

Table 9.1. Presentation of results from a hypothetical experiment to evaluate five varieties

Accession	Mean disease score	Standard deviation	Percentage of highly resistant plants (score of 1)	Percentage of resistant plants (score of 1–3)
D	2.4	0.5	0	100
A	2.7	1.2	10	85
C	4.2	1.0	0	10
E	4.6	2.5	2	24
B	8.5	0.3	0	0

In Table 9.1 the standard deviation of the disease scores is given as well as the mean, and is different for the five varieties. For example, if the aim was to choose accessions that included some highly resistant plants, then accessions A and E might be chosen, even though accession D has the best mean score.

A graph might be used to show resistance to two diseases. Here it is possible that more details might be provided for each accession, if requested by the user. This would include the name of the accession, which connects to yet further details as shown in the example from the NPGS Web site in Section 2.

The first formal step in the analysis is normally exploratory and is described in Section 9.2. Its aim is to visualize the data, perhaps to review the presentations you thought were appropriate. An exploratory analysis is also to see if there are any oddities in the data. These have to be dealt with before continuing with the analysis.

Then, if the trial has a standard design, the data are normally subjected to what is called an “Analysis of Variance” (ANOVA). This is described in Section 9.3 and the form of this part of the analysis will depend on the design of the trial, as described in Section 6. This step is not obligatory and, for some trials, simple tabulation of the data is all that is necessary to give results, such as those shown in the tables in Section 2 of this report.

One possible complication in the analysis is in trials where measurements were made at the plant level, while the ANOVA is undertaken at the plot level. The plant-level data are therefore normally summarized at the plot level prior to this analysis. How this can be done was described in Section 8.

A range of further analyses is possible and some ideas are mentioned in Section 9.4.

9.2 Exploratory analyses

The first stage in the analysis is usually exploratory. Two examples are shown below for datasets that are analyzed further in Section 9.3. The first example is the experiment with 8 varieties of rice and 3 replicates, from Gomez and Gomez (1984), that was used in Section 8. Data were recorded on the number of panicles per hill, on 12 hills in each of the 24 plots. Here the raw data on the 288 hills are presented as box-plots, with a separate box-plot for each experimental plot. The box-plot is a 5-number summary of each set of data (Fig. 9.1). For example, for the first plot, it shows the minimum is 5 panicles, the median is 10 and the maximum is 14. The other two points give the lower and upper “quartiles,” or the values of the variable such that 25% and 75% of the values of the variable fall below that value. The 25% point (first or lower quartile) is 8 and 75% point (third or upper quartile) is 11.

The box-plots are useful in presenting the data and also to indicate possible outliers. In the graph, attention is drawn to three observations, namely hills 72, 78 and 175. The three values indicated are given in bold in Table 9.2. This presentation shows, for example, that the value of 14 panicles for one of the hills for the variety Jaya in Rep 2 is surprising, given that the other 11 values range from 6 to 11 panicles.




Fig. 9.1. Box-plots of the individual values of the number of panicles per hill from 24 plots (raw data in Table 8.1).

The second example is a simple 5x5 lattice, using an example of crop yields from a soybean experiment, taken from Cochran and Cox (1957: pp. 406, 412).

The graph in Fig. 9.2 shows the values from all 100 plots, with the x-axis giving the 25 varieties and a different symbol for each replicate. The means are also marked on the plot, connected by an interpolating line. Scrutiny of this graph indicates some variety

Table 9.2. Data on number of panicles per hill from three plots outlying observation (in bold)

Variety	Replicate	Number of panicles											
IR480	1	4	6	6	7	7	7	8	8	8	9	11	14
Jaya	1	3	7	7	7	8	8	8	9	9	10	10	12
Jaya	2	6	6	7	7	8	8	8	9	9	9	11	14

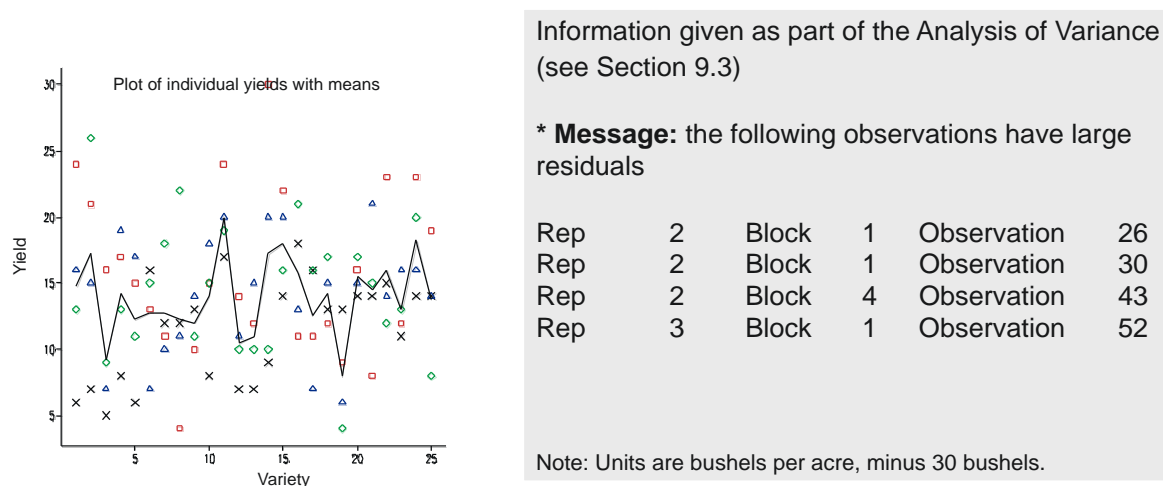


Fig. 9.2. Exploratory graph and diagnostics following an Analysis of Variance of the lattice data from Cochran and Cox (1957).

differences and also some observations that should be examined critically. For example, variety 11 has consistently high yields, in marked contrast to variety 19. Variety 14 has a reasonably high mean, but primarily because of a very high yield (30 units) in one of the plots.

The graph in Fig. 9.2 has been drawn with the x -axis in the order that the varieties were given. Would the display be more useful with the varieties in descending order of mean yield? Probably. And this kind of question is exactly why the people conducting an experiment should remain closely involved in the analysis. Experimental results cannot be exploited fully if they are just sent away for analysis. Where statisticians are involved, they should become part of the research team. They should analyze the data **with** the genebank managers and germplasm users and not simply **for** them.

This second example typifies the situation where there is highly structured data. Here the data are from a lattice, i.e. each repetition of 25 plots has been divided into 5 blocks, each with 5 plots. In such cases, exploratory graphs such as Fig. 9.2 remain of use, but the data exploration should also continue after taking account of the data “model”, which in this case includes the blocks. In Section 9.3 we consider the more formal analysis and part of one message from this analysis is included in Fig. 9.2. This indicates that some unusual observations could usefully be examined to see if there is a problem. This does not imply that these observations should be eliminated, only that they deserve close scrutiny.

We return to the subject of data exploration at the end of Section 9.4.

9.3 Standard methods of analysis

We start this subsection with some general points on the computer software that might be used for the analysis as a whole. Some different packages are compared in Appendix 2, but here we discuss desirable properties of whatever software is to be adopted for the analysis. In Section 9.2 we saw that the software should encourage a critical approach to the data analysis. This will partly be by giving tables of the raw data but mainly by graphs. It is useful to distinguish between “exploratory graphs” that help the research team to visualize and scrutinize their data, and “presentation graphs”. The latter are used following the analysis to present the results to others. So, for the data scrutiny, the software should have good facilities for exploratory graphics.

A formal analysis is often needed and a simple requirement of the software is that it can analyze all the designs used in the experimental research and present the results clearly for each design. We have indicated above (see Fig. 9.2) that the formal analysis should also encourage, or at least permit, the analysts to continue their critical approach to the data analysis.

Finally, real experimental data almost always include some complications. Sometimes these complications are predictable, for example it is more complicated to analyze a lattice design than a randomized block design. Often there are (additional) complications that arise when the experiment is being conducted. This may be at the planning stage, for example not having enough seed to have equal replication of each accession, or during the experiment, for example finding plots destroyed by animals.

A different type of complication is that typical data from the trials envisaged here often include information on counts (such as germination numbers) or categories (such as disease scores). Simple text-book analyses and some statistical software do not handle this type of complication easily. They are restricted to data analysis for continuous, quantitative variables — such as yield — that may be assumed to have a so-called “Normal” distribution, or the famous “bell curve”, where the exact shape of the distribution is defined by a function which has only two parameters, the mean and the standard deviation.

So we need software that presents the results from a formal analysis clearly and can easily handle typical complications. In this section we illustrate these two points. We consider first the formal analysis of the Lattice example, plotted in Fig. 9.2. Then we take a part of the analysis for the data given in Tables 8.1, 8.2 and Fig. 9.1 to illustrate one “complication” that may be common in the type of experiments that are envisaged here.

An analysis of the lattice data shown in Fig. 9.2 is given in Table 9.3. It is useful to assess what this formal analysis provides in addition to the simple tables that might be suggested directly from the objectives of the trial. In this trial, the objectives would probably lead to the presentation of the mean yields of the varieties in order. It is comforting to note that such means are indeed given by the standard analysis shown in Table 9.3. There are some additional results and we consider briefly in turn how these are used. In this discussion we are making general points rather than just analyzing this particular set of data. We see from Table 9.3 that the results consist of four parts. There is an ANOVA table, information about possible problem observations, the variety means and also the standard error of the difference:

1. There is much to discuss concerning the ANOVA table in a course on analysis. Here we note just one point concerning the Variety line that is in bold in the table. Here the F value of 3.38, and the final probability of effectively zero, indicate there are real differences between the yields of the different varieties. This provides what is effectively a “passport” to report differences between the varieties, knowing that there are real differences to report.
2. The message on residuals was described in the previous section. It reminds us to look critically at the data at all stages, and not merely at the beginning of the analysis.
3. The means can now be scrutinized and action taken, depending on the precise aims of the trial. These are not just the simple means of the four observations for each variety, but have been adjusted for the particular blocks of the experiment. Looking at these adjusted means we see that Variety 11 is in first place, as could be predicted from the exploratory analysis. Second is Variety 2. However, we note from the graph in Fig. 9.2 that the yields for this variety were very different in the different repetitions. We therefore look in more detail and see that one of the warning residuals concerned this variety.
4. The last element of this analysis is the standard error, which is used to put the differences between the means into perspective. We note, for example, that the difference in mean yield between Variety 11 and the next best variety is more than one standard error, but the mean yields of the next seven varieties in the list are quite close.

In practice we would possibly not proceed much further with the analysis of this one trait at this stage, but might now use these

Table 9.3. Standard of Analysis of Variance for the lattice data from Cochran and Cox (1957)

Source of variation	d.f.	S.S.	M.S.	F	Probability
Replicates	3	226.19	75.40	3.66	
Blocks within reps	16	474.00	29.62	2.18	
Varieties within block					
Variety	24	1103.24	45.97	3.38	<0.001
Residual	56	761.56	13.60		
Total	99	2564.99			

* **Message:** the following units have large residuals *

Rep 2	Block 1	Obs 26	6.64 s.e.	2.76
Rep 2	Block 1	Obs 30	-7.46 s.e.	2.76
Rep 2	Block 4	Obs 43	6.44 s.e.	2.76
Rep 3	Block 1	Obs 52	6.81 s.e.	2.76

*** Ordered Treatment means ***

Order	Variety	Mean
1	11	22.10
2	2	19.31
3	15	18.67
4	14	17.89
5	24	17.65
6	22	17.02
7	1	16.66
8	21	15.36
9	4	14.69
10	16	14.57
11	23	13.92
12	25	13.18
13	13	13.16
14	18	13.14
15	20	12.90
16	12	12.76
17	5	12.73
18	7	11.89
19	6	11.73
20	10	11.55
21	17	11.48
22	8	11.30
23	3	11.22
24	9	9.52
25	19	5.36

Average standard error of difference
between means = 2.87

results in conjunction with further information on other traits. This full information from all the traits might also help in our assessment of whether any of the suspect observations indicate a real problem with the data from a particular plot.

Our second task is to show how the analysis can proceed if there is a complication. Here we return to what might be a common problem. We may not be interested in just the mean response, but also in knowing whether an accession includes promising plants. The rice data from Table 8.2 are used to illustrate some methods of analysis in this case.

There is a second issue here, in that in the analysis we could use either the raw data from the 288 hills, or the summary values (see Table 8.2) that were calculated for the 24 plots. There we calculated various summary statistics, including the mean number of panicles per hill and the “spread” of the number of panicles per hill within each plot, as measured by the standard deviation.

For illustration, we suppose that our main interest is not in the mean number of panicles per hill, but in varieties with a high proportion of hills with 10 or more panicles. This is the same concept as choosing accessions with a high proportion of resistant plants (mentioned in Section 9.1), using the hypothetical data from Table 9.1. How should the analysis proceed in this case?

Table 9.4. Analysis of mean number of panicles per hill

	d.f.	S.S.	M.S.	F value	Probability
Reps	2	53.52	26.8	3.62	
Variety	7	191.06	27.3	3.69	0.018
Residual	14	103.59	7.4	1.47	
Hills	264	1331.33	5.04		
Total	287	1679.5			

Variety	Mean
IR22	9.3
IR160-27-3	8.6
IR20	8.4
IR480-5-9	7.9
C4-63	7.6
Jaya	7.5
RD-3	7.3
BPI-76-1	6.5
Standard Error of Difference (SED)	0.64
Least Significant Difference (LSD)	1.38

We still begin with an analysis of the mean number of panicles per hill. The results are shown in Table 9.4.

If our interest is in choosing varieties with more panicles on average, then we see that the varieties IR22, IR160 and IR20 are possible choices.

A key point for our specific objective is whether the variability of the observations is the same for all varieties. This is a general issue in data analysis. We often have situations where we assume that the mean may be different for different treatments, but the variability is the same. This point is important here because if the spread (i.e. the “shape”) of the data is the same for all accessions, then we can use the means (as shown in Table 9.4) to choose between accessions even when the mean is not of direct interest.

For this set of data, we are able to assess whether the variability is the same for all varieties, because we have calculated the standard deviation for each of the 24 plots (Table 8.2). These values can be analyzed as above, i.e. subjected to standard analysis of variance. The analysis is not given here because there is no evidence that the spread is different for the different varieties.

An alternative approach for this objective is to summarize the number of hills with 10 or more panicles directly. The number of hills with 10 or more panicles was calculated earlier and is shown in the last column of Table 8.2. These values can then be analyzed. The percentages are shown in Table 9.5 and lead again to the choice of the same three varieties, with some indication that variety IR22, for which 58% of the hills had 10 panicles or more, is superior to IR160 and IR20.

Table 9.5. Direct analysis of the percentage of hills with 10 or more panicles

Variety	Percentage	Standard error
IR22	58	0.081
IR160-27-3	36	0.079
IR20	28	0.074
Jaya	17	0.062
IR480-5-9	17	0.062
C4-63	14	0.057
RD-3	8	0.046
BPI-76-1	8	0.046

The key point is that there are often alternative methods of analysis. We have looked at two. Usually, they lead to the same conclusions. When this is not the case, it is interesting to assess what aspects of the data are being used differently by each method.

9.4 Further methods of analysis

In the subsections above we have outlined the way in which we assume that most sets of data will be processed initially. Here we mention some further methods of analysis. However, we caution against the assumption that advanced statistical methods will be of dramatic value in the processing of the type of data that is collected in these trials. We propose that the most important advances in the analysis and presentation of the results will come through discussions with breeders and others who are interested in making use of the genetic materials from the genebank.

We consider that the data available through the NPGS website provide the current “standard” and hence this type of presentation can be used to suggest possible improvements. These will inevitably be in two directions. The first is to provide more information about the environments in which the experiments were performed. This involves reporting the experiment-level information described in Section 7.1. The second is to provide more detail, i.e. more information about the plot-level and even the plant-level information. It is likely that some users will require summaries of the data, on which they can take direct action, while others might request access to the raw data, so they can conduct their own analyses. Providing access to the raw data has not been possible until recently, because of the volume of the data, but now it is just the difficulties of effective data management and questions of property rights that have to be resolved.

Our emphasis on the search for simple methods of presentation is not to deny that there have been important recent advances in statistical methods that will help in the analysis and presentation of the data from these trials. We first mention three areas that should help in the analysis of traits that are to be processed singly.

1. The first is that the data collected in this type of trial are often not Normally distributed. In the past, the only method of analysis for data that did not satisfy the assumption of a standard ANOVA was to transform the data and hope that the transformed values were now Normally distributed. There are now better methods of analysis for non-Normal data. These include, in particular, the facility to handle binomial data, such as was used in Section 9.3 to process the number of hills with 10 or more panicles. This extension of the Analysis
-

of Variance to the “Analysis of Deviance” can also be used for more general “categorical” data. Disease score on a 1 to 9 scale and binomial data (a special case where there are just two alternatives) are examples of such data.

2. The second area is that of spatial analysis of data from field trials, which refers to a class of techniques used to analyze and predict values of a variable distributed in space. This is particularly important for trials that involve repeated use of controls that are well distributed over the experimental area, with just one or two repetitions of the test accessions.
3. The third topic is that of multilevel modelling. Here the two important levels will be the plot level and the plant level. There might be considerable interest in accessions whose plants are highly variable, because this could indicate the presence of interesting genetic traits. The analysis of data over different levels is not new, and recent advances, driven partly by research into the analysis of survey data, concern particularly the modelling of variances as well as the means.

We have not considered here a higher level, when trials are repeated over different years and sites and a combined analysis of the accessions is required. This is a large and important subject going under the name of genotype by environment interaction. However, one feature of many modern methods of analysis is their complexity, combined with a lack of utilization of the detailed characteristics of each site. We would argue that reports of the type of experiment envisaged in this guide may wish to take a much simpler approach. That is to report each experiment separately, but to include all details of each environment in the report. Thus managers should insist that the experiment-level information, as described in Section 7, is always presented as part of the report.

A further topic that has not been covered is the methods of analysis that are appropriate for data that are “repeated measures”. These are trials where measurements are made on the same trait at intervals throughout the growing period.

Finally, on the methods of analysis we note that most trials will involve the collection of information on multiple traits. The methods of analysis described so far are for the presentation of the results on each trait separately. However, breeders may be interested in accessions, or even plants, that show simultaneous resistance to a number of diseases. This introduces the huge subject of “multivariate analysis”. We caution managers that they should not give enthusiastic analysts free reign in this area. One powerful weapon for the manager is the “so what” test, which

may be used to quiz the analysts when they present their results. It is used in the following way: “*So what in simple terms have you learned, given the objectives of this trial?*” or “*So what have you learned that is not clear from a simple analysis of this trial?*” or even just “*So what?!*” We propose that multivariate methods may have an important role to play in the analysis of the data from these trials. But they should be considered at the beginning of the analysis, as a means of data exploration (see Section 9.2), rather than as something complicated, to be done after the standard analyses have been performed. Some trials, or sets of trials, may be used to look for groups of accessions that behave similarly. This points to techniques in “cluster analysis” that look for natural groupings in data. And recently there have been advances in ways of visualizing (just looking at, but in interesting ways) multivariate data. These methods are highly interactive and might help breeders to pinpoint “odd” accessions that are different in a way that makes them interesting.

10 Conclusions

In this guide we have given information on statistical aspects of the design and analysis of trials that could enable genebank managers to publicize the details of accessions in their genebank. We have provided the information at the level that we believe can support managers who need to decide on their experimental strategy and then remain involved with the research as it proceeds.

More detailed information is needed by researchers or managers if they are actually to conduct the type of trials envisaged in this guide. We have used examples in this guide from some of the standard textbooks that could be consulted. An alternative would be to provide special training workshops and we consider a possible structure for a training strategy in Appendix 2. This guide might provide preliminary reading for such a workshop.

In Section 9, on analysis, we mentioned that the statistician's role should be as a member of the research team, working with managers and other researchers and not as a group apart. The same applies to the production of this type of guide itself. Statisticians have prepared this document, with no direct input from the genebank managers that it is supposed to help. We hope that it succeeds in stimulating discussion and look forward to the teamwork that should help future versions be more informative.

Bibliography

Main

- Cochran, W.G. and G.M. Cox. 1957. *Experimental Designs*. John Wiley and Sons, New York, USA.
- Gomez, K.A. and A.A. Gomez. 1984. *Statistical Procedures for Agricultural Research*. John Wiley and Sons, New York, USA.
- Kempton, R.A. and P.N. Fox, eds. 1997. *Statistical Methods for Plant Variety Evaluation*. Plant Breeding Series 3, Chapman & Hall, London, UK.
- Mead, R., R.N. Curnow and A.M. Hasted. 1993. *Statistical Methods in Agriculture and Experimental Biology*. Chapman and Hall, London, UK.
- Pearce, S.C., G.M. Clarke, G.V. Dyke and R.E. Kempson. 1988. *Manual of Crop Experimentation*. Charles Griffin, London, UK.
- Snedecor, G.W. and W.G. Cochran. 1989. *Statistical Methods*. Iowa State University Press, Ames, Iowa, USA.
- SSC. 2000. *Data Management Guidelines for Experimental Projects*. Statistical Services Centre (SSC), University of Reading (<http://www/rdg.ac.uk/ssc>).
- Watkinson, A.R., R.J. Lawn, R.H. Ellis, A. Qi and R.J. Summerfield. 1994. RoDMod A computer program for characterising genotypic variation in flowering responses to photoperiod and temperature. CSIRO Division of tropical crops and pastures. St Lucia, Queensland, Australia.

Augmented designs

- Besag, J.E. and R.A. Kempton. 1986. Statistical analysis of field experiments using neighbouring plots. *Biometrics* 42:231-251.
- Briggs, K.G. and L.H. Shebeski. 1967. Implications concerning the frequency of control plots in wheat breeding nurseries. *Can. J. Plant Sci.* 48:149-153.
- David, O., R.A. Kempton and I.M. Nevison. 1996. Designs for checking interplot competition in variety trials. *J. Agric. Sci. (Camb.)* 127:285-288.
- Kempton, R.A. and A.C. Gleeson. 1997. Unreplicated trials. Chapter 6 in *Statistical Methods for Plant Variety Evaluation* (R.A. Kempton and P.N. Fox, eds.). Plant Breeding Series 3, Chapman & Hall.
- Kempton, R.A. and M. Talbot. 1988. The development of new crop varieties. *J. Royal Stat. Soc. A* 151:327-341.
-

- Lin, C.S. and G. Poushinsky. 1983. A modified augmented design for an early stage of plant selection involving a large number of test lines without replication. *Biometrics* 39:553-561.
- Lin, C.S. and G. Poushinsky. 1985. A modified augmented design (type 2) for rectangular plots. *Can. J. Plant Sci.* 65:743-749.
- Lin, C.S., G. Poushinski and P.Y. Jui. 1983. Simulation study of three adjustment methods for the modified augmented design and comparison with the balanced lattice square design. *J. Agric. Sci. (Camb.)* 100:527-534.
- May, K.W., G.C. Kozub and G.B. Schaalje. 1989. Field evaluation of a modified augmented design (type 2) for screening barley lines. *Can. J. Plant Sci.* 69:9-15.
- Schaalje, G.B., D.R. Lynch and G.C. Kozub. 1987. Field evaluation of a modified augmented design for early stage selection involving a large number of test lines without replication. *Potato Res.* 30:35-45.
- Townley-Smith, T.F. and E.A. Hurd. 1973. Use of moving means in wheat yield trials. *Can. J. Plant Sci.* 53:447-450.
-

Appendix 1 Selecting the appropriate software

Our main conclusion from a brief survey of the software is that there is no ideal package for the design and analysis of germplasm evaluation trials. However, we are not necessarily searching for a single “winning” package. Organizations may have a strategy that includes a range of packages. One scenario would be to use Excel for data entry and possibly for some graphics. Then some combination of Genstat, Agrobases and SAS could be used for the randomization of the trials and for the analyses.

One recent development is the ease and similarity of use of different statistics packages. This has two important consequences. The first is that little time need be devoted to instruction in any particular package. The second is that more than one package can be used in a complementary way. It is therefore no longer essential that the same package be used on a training course that is needed subsequently. The Website <http://www.statistics.com/vendors/index.html> has information on, and links to, many statistical analysis software packages, including some of the ones discussed here.

Specific points are as follows:

1. Microsoft's Excel[®] is likely to be familiar to most users, and to be available on most computers. If used with discipline, it can be the package of choice for data entry and for some of the data management tasks. It may also be useful for some of the simple analyses, particularly where simple tabulation is all that is required. There are statistical functions within MS Excel, but some have problems. There are also packages that are advertised as add-ons to MS Excel. However, none that we know of offer the facilities that are needed for the comprehensive analysis of experimental data. Hence, we do not recommend MS Excel as a serious package for statistical analysis of germplasm evaluation trials.
 2. MSTAT (<http://www.msu.edu/user/bricker/mstat.htm>) is a popular package for the design and analysis of experimental designs. It can be given credit for introducing many scientists in developing countries to the use of the computer for data analysis. It does have problems, however. It now looks old-fashioned, unless and until there is a Windows version. It is limited in the designs that can be randomized and analyzed. A serious defect is that it does not permit (or even encourage) a critical attitude to data analysis. It is the only commonly used statistics package, that we know of, that does not allow users to access the residuals, to check on the validity of their analyses. Despite being cheap, we are
-

not able to recommend MSTAT for the applications described in this manual.

3. Agrobase (<http://www.agronomix.mb.ca>) has probably the best facilities for supporting the design and layout of the types of experiments discussed in this guide. It has some facilities for analysis, but they are weaker than those of packages discussed below. It is in Microsoft Windows™ and is possibly easier to use than the other packages discussed below. If Agrobase is adopted, there should also be access to more powerful software for the cases when its facilities for analysis are insufficient.
 4. Genstat (http://www.nag.co.uk/stats/tt_soft.asp, <http://www.mimas.ac.uk/stats/maps/genstat.html>) has excellent general facilities for the analysis of experimental data. It is also good in encouraging a critical attitude to data analysis. It has some facilities for design and randomization, but these are weaker than those in Agrobase and not as easy to use. The output is of high quality, not voluminous but highly informative. Genstat is produced by statisticians at Rothamsted Experimental Station, UK and is aimed at the type of trials discussed in this guide.
 5. S-PLUS (<http://www.splus.mathsoft.com>) is a modern package, with the best graphics facilities for the display of experimental and other data. The Windows implementation is comprehensive and it has therefore recently become accessible to a much wider audience than statisticians. It is an extremely powerful object-oriented package with comprehensive facilities for data analysis. For users with experience of this package it is easy to extend the facilities. However, numerical output is not displayed in a form that is as clear as Genstat for the analysis of experimental data. If funds are not limiting it is an exciting package to include, though primarily for organizations that have strong statistical support.
 6. SAS (<http://www.sas.com>) remains the giant among statistics packages. It provides comprehensive facilities for data analysis for the types of trial considered here. Its latest release, version 8, has a user interface which is easy to use for non-statisticians. As with S-PLUS, we suggest that organizations where there is statistical support, and finance is not limiting, should include SAS within their statistical armory.
 7. There are some specialist packages that could be evaluated as additions to the software considered above. We mention two here, but there are others. The first is called CycDesign (<http://www.ffp.csiro.au/software/>) and is produced by the
-

- Commonwealth Scientific and Industrial Research Organization, Australia (CSIRO). Running under MS Windows 95 or NT, it is aimed at designing and randomizing a wide range of alpha and related designs. However, it has no facilities for subsequent analysis. The second package is called ASREML (<ftp://ftp.res.bbsrc.ac.uk/pub/aar>; discussion group at asreml@chiswick.anprod.csiro.au). This is very fast and has the capability to analyze very large unbalanced designs and also some spatial designs that may be too large for, or are not supported by, SAS or Genstat. It is not, however, particularly friendly to the beginner, in either its input or output stages.
8. There are many other statistics packages, but most do not have sufficient facilities for the analysis of experimental data. SPSS (<http://www.spss.com>) is rightly recommended for the analysis of survey data. Other packages that we would claim do not rival those described above, in this context, include Systat, Statistica, JMP, Statgraphics and Minitab.
 9. Some packages, including Minitab (<http://www.minitab.com>), are rightly popular to support the teaching of statistics, and could be considered here in that context. However Minitab's bias in experimental design and analysis is toward industrial experiments, and hence it could not be the only package in use subsequently.

Thus, for data entry we suggest MS Excel, with consideration given to the Logbook software that links to MS Access[®], for data management. Alternatively, Agrobase and SAS provide good data management facilities themselves. For analysis, if just one package were obtained, then we would suggest Genstat as being the most appropriate all-rounder for experimental data. It is user extendable, and hence could also be made simpler for any designs or analyses that were to be recommended for routine use by managers. Consideration should also be given to Agrobase, particularly if more than one package is to be obtained. Our reticence concerning Agrobase stems from its lack of flexibility in analysis. It handles the standard analyses easily, but we are not clear that this will be sufficient for the presentation of the results of these experiments. It is also not a cheap package, given its limitations.

Neither Genstat nor Agrobase are expensive by SAS or SPSS standards, but they do imply the kind of cost that could possibly be best negotiated by IPGRI on behalf of a project or a group of institutes.

Appendix 2 A training strategy

We consider one possible scenario for training, which could be combined with the development of a strategy for experimental work by genebank managers. This would consist of an initial workshop, primarily devoted to the planning of the trials. Participants bring information about their genebank and proposed experimental protocols with them. The workshop then includes a detailed discussion of the proposed protocols.

The trials are then conducted and are followed by a second workshop that is primarily concerned with the description of methods of analysis of the data. This is combined with detailed analyses of the data from the current trials.

There would be two types of participants:

1. Most would be genebank managers, or the scientists who are being commissioned to undertake evaluation trials on their behalf. These would be highly participative workshops and hence managers would have to be personally closely involved in the experimental programme to be eligible to attend. They would be asked to supply proposed or past protocols with which they are personally involved with their application.
2. The second type of participant would be the statisticians who would provide support for the work. Much material in this type of workshop, particularly on design, will be new to many statisticians. If a sequence of workshops is anticipated, then some of the statisticians might become resource persons in later workshops. There should not be too many statisticians at the workshop, say a maximum of a fifth of the participants, and they would not be expected to provide protocols as a passport to participation.

The workshop on design could cover in more detail the material that is described in Sections 2 to 7. This would be done in an interactive manner with continued small-group discussions on each area in relation to the proposed experimental protocols. Topics that are not in this guide, but that might be important in a workshop, include sampling methods for plant-level information.

There would also be brief sections (say half a day) on data management, particularly on data entry and on the concepts underlying the analyses.

The second workshop would review the methods of data entry and management and then cover data analysis and the presentation of results. About half the time could be devoted to instruction, with the remainder being used for analyses of the data from the current

experiments. This workshop might also review briefly the concepts of design to assess future designs in the light of the results from the current experiments.

There are other possible types of training. A 1-week workshop might be appropriate for senior managers. These would be for senior staff who need to know the general concepts of experimental strategy (this includes both design and analysis), but who would not be closely involved in the details of individual experiments. This could cover the material included in this guide, but perhaps extend the topics on alternative strategies for information collection (Section 2).

At a different level, instruction on data entry and management is sometimes usefully conducted as a “within-institute” training course. This might consist of an initial visit to discuss viable strategies for the institute, followed by instruction of up to a week on the agreed procedures. This might include direct instruction to staff who will be responsible for the data management, plus the establishment of a regular, short, internal course for assistants undertaking the data entry. It may be appropriate to combine this type of “roving” workshop with the guidance on the data entry and management of the trials that is included in the training workshops described above.



FUTURE
HARVEST
<www.futureharvest.org>

IPGRI is
a Future Harvest Centre
supported by the
Consultative Group on
International Agricultural
Research (CGIAR)

ISBN 92-9043-505-4