Data analysis of agroforestry experiments

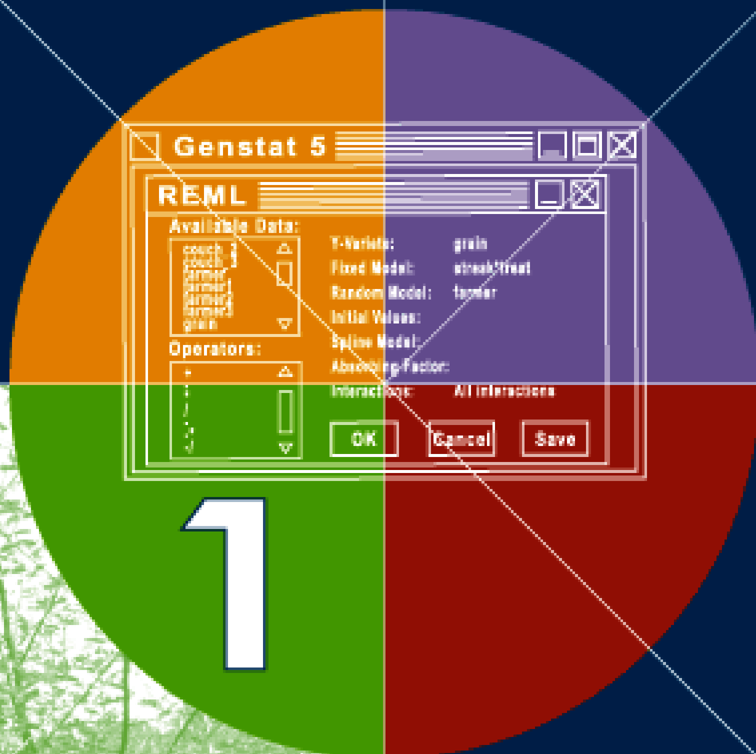

Workshop overview & session summaries

World Agroforestry Centre

TRANSFORMING LIVES AND LANDSCAPES

Data analysis of agroforestry experiments

1

by

Richard Coe

Roger Stern

Eleanore Allan

edited by

Jan Beniest

Janet Awimbo

Workshop overview & session summaries

World Agroforestry Centre
TRANSFORMING LIVES AND LANDSCAPES

The World Agroforestry Centre (ICRAF) is the international leader in Agroforestry – the science and practice of integrating 'working trees' on smallholder farms and in rural landscapes. Agroforestry is an effective and innovative means to reduce poverty, create food security, and improve the environment. The Centre and its many partners provide improved, high quality tree seeds and seedlings, and the knowledge needed to use them effectively. We combine excellence in scientific research and development to address poverty, hunger and environmental needs through collaborative programs and partnerships that transform lives and landscapes, both locally and globally.

ISBN 92 9059 145 5

ICRAF
The World Agroforestry Centre
United Nations Avenue
PO Box 30677
Nairobi, Kenya
Tel: + 254 2 524 000
Fax: + 254 2 524 001
Contact via the USA
Tel: + 1 650 833 6645
Fax: + 1 650 833 6646
E-mail: icraf@cgiar.org
Internet: www.worldagroforestrycentre.org

Design: Mariska Koornneef
Printed by: Kul Graphics Ltd, Nairobi, Kenya

Table of contents

The course structure and strategy

Introduction

These course notes are on the analysis of data from experiments. They result from a series of statistics training courses organized by ICRAF/World Agroforestry Centre. These courses were originally on the design and analysis of agroforestry experiments, but they have been used more widely than this.

The first component was on the design of experiments. This analysis course assumes familiarity with the main concepts from the design course. A brief review is given in Session 1.

The second component was on data entry and management. This is a key area because poor data management often limits the processing of data. In this analysis course the examples provided have been 'managed' so that the concepts related to the analyses could be illustrated easily. We anticipate that an initial phase in the course preparation will be to organize datasets from participants similarly. Hence, the data management component though normally undertaken prior to this component, is not a necessary prerequisite.

This course is divided into two parts. The first part is entitled **The Everyday Toolkit** and covers the concepts that we believe scientists should be able to understand fully and the corresponding analyses that they should be able to undertake unaided following the training.

The second part is called **Handling Complexities**. This examines how experimental data can be processed where there are complications. These complications are divided into three broad types. The first is due to complexities in the design which may either be due to a complex treatment structure, to difficulties in the layout of the trial, or to the way data were measured. For example, a measurement of farmers' responses may be on a 5-point scale ranging from very good to very poor. The analysis of this type of 'categorical data' is described here.

We do not consider on-farm trials as a special category and hence examples of them will be used throughout the course. However, their analysis is often complicated because of their combination of a complex layout (many farmers, with few plots per farm) and the nature of measurement. The complexities arise from the lack of control of factors that would be within the treatment structure in an on-station trial, and the fact that this lack of control occurs both within and between farms. The handling of these complexities is discussed in the course.

The second type of complexity is that which is due to the particular field of application. Particular features of agroforestry trials include 'repeated measures', both in time and space, and difficulties that arise from the need to measure multiple components (e.g. concerning both trees and crops) within each plot. Courses for other audiences need to replace this section, as each subject area has a set of problems and methods specific to it. As an example, we provide a parallel session that considers some of the complications that are commonly encountered in livestock experiments.

Finally we consider complexities that arise because of the nature of the data. Coping with zeros in the data and missing values are among the topics considered here.

Our main aim in this second part of the course is for scientists to be aware of the methods that now exist to handle complex data. These are methods where scientists, at least initially, might want to work jointly with statisticians.

Audience

This course is intended primarily for scientists undertaking agricultural research.

In targeting the course for scientists we are assuming some prerequisites. We assume scientists have some practical experience in the design and analysis of trials. They may have felt diffident in their write-up of an experiment for which they have been responsible, but they have an awareness of the process of conducting and processing an experiment.

We assume basic computing skills. Scientists who do not have regular access to a computer and who are not comfortable with the use of a word processor and a spreadsheet should not take this course. Most scientists will already have some experience in data processing using a statistical package. Those without this experience should make themselves aware of the capabilities of this type of software before the course.

We assume some basic statistical knowledge, usually from statistics courses taken while participants were students. We do not assume that they liked the course, or that they understood all the content. We hope that at least something was understood and a little is still remembered! The design course is also a good preparation for this training and we hope that participants who have followed it still remember (and use!) some of the concepts that were covered in that course.

We mentioned above that this course is divided into two parts, 'The Everyday Toolkit,' and 'Handling Complexities'. There is no reason why both parts have to run sequentially. If the material is too demanding, then the training could be split into two parts, with the first course concentrating on the Everyday Toolkit only. We consider this point further in the section on the duration of the course, where we suggest different ways that the training could be given.

We assume that participants are in posts where they are able to use the materials that are covered here regularly during the following year. This workshop should be seen as the start of the learning process, rather than as an end in itself. Even after an MSc, most statisticians find that they learn more in the subsequent years by practicing what they have learned. Similarly here, participants who follow the course by avoiding data analysis for the next year are unlikely to profit much from the training.

As with the design course, we anticipate that the materials and the approach to training might also be of interest to the supporting statisticians. They may be in the agriculture faculty of a university or in a research institute. They will typically find that the topics are familiar, but the practical way they are introduced and the way they are covered and justified may be new.

Resource persons

At least one of the resource persons should be a statistician and preferably there would be at least two statisticians involved for the full duration of the training.

The approach adopted in this course is sufficiently different to conventional textbooks or other courses on statistics that we suggest that a statistician should first attend this, or a similar course, as a participant, before being involved as a resource person.

Ideally the resource persons should not all be statisticians; there should also be one or two agriculture professionals. Sometimes someone in an institute has become a de-facto 'semi-statistician', in that they are willing to help colleagues on their statistical problems occasionally. Having initially attended this type of course, or the course on design, they can support the resource team, which consolidates their statistical skills, while providing excellent support to participants.

Resource people who are skilled at scientific data management and in data processing can help participants with their data preparation and software use. They would not normally be resource persons for the statistics sessions.

Datasets and data management

In the design course, one requirement was that participants provide a protocol for an experiment with which they are involved. The same is required here, but additionally participants should bring the data. Ideally they should bring the data as computer files and also a photocopy of the data on the original data collection sheets. The protocols at least should be sent ahead of the training so sessions can make maximum use of the data from participants.

The materials presented here include datasets from ICRAF, ICRISAT and ILRI experiments, and part of the preparation for the course involves substituting the examples from participants wherever possible.

The penultimate session in this course describes part of the course that is devoted to the analysis of the participants' own data. Usually, in addition to the sessions described here, there is an introductory session for participants to check that their data are organized, so they can be analysed during the course.

In ICRAF courses, a key component of this introductory session was for participants, working in pairs, to exchange the sets of data and protocols that they have brought. Here the recipient is a participant who knows nothing of the trial other than the data and protocols given. Each participant then seeks explanation of any elements that are not clear. Queries are specifically related to the data, rather than whether the objectives are clearly specified.

Examples of queries are as follows:

- What are the meanings of particular columns of data?
- What are the units of measurement?
- Are the data in Excel?
- Can they be easily transferred to GenStat, or to another statistics package?
- Is there sufficient information of the experimental design to complete an analysis?

Support staff are on hand to help in the organization of the data.

This session was often held on the first afternoon of the workshop and resulted in the participant's data being available for analysis, from day 2. This helped both the individual participants and the resource persons.

Further details of this session are given in the notes on the data management component. The examples of datasets that we have provided are intended to demonstrate 'good-practice'. For consistency they are all provided as Excel files, though some were originally in Microsoft Access.

Duration

The design course was typically of 2 weeks duration and we assume here that this may be the case for the course covering both data management and analysis components. It is anticipated that the data management component would be for about 3 days. In this case the basic and further sections of statistics would take about 4 days each.

This would be highly intensive and other scenarios are possible. Where participants are relatively inexperienced the course might run for 3 weeks; the first week could be devoted to data management, the second to the basic analyses, and the third to the further methods.

If yet a further week were possible, making 4 weeks in total, we suggest that the additional week be devoted to the analyses and write-up of the scientists' own datasets. Here the first two weeks might be as above (in the three-week course), with the final 2 weeks shared between the further methods and the scientists' data.

If scientists are relatively inexperienced, but a course of only 2 weeks is possible, we suggest that it is preferable to cover one part of the material well rather than to attempt to describe all the topics. One scenario is to devote the first week to the data management plus an introduction to the statistical software. The second week would be devoted to the basic statistical analysis. Within this second week there might be a single session (half-day) reviewing the contents of the further statistical methods part.

Where training is within a country (so travel costs are low) then we suggest that dividing a four-week training course into two sections of 2 weeks each would be of even more use. In this case, we suggest that the first training might be on the data management component and the second could be on analysis.

Even within the 4 weeks scenario, it is important for participants to be realistic in their objectives from this training. We aim for participants to be comfortable with basic statistical methods. They should be able to conduct their own analyses when their problems are simple. They should know when a problem is simple and when it is complicated. For complex problems they should have some ideas about the methods of analysis and be comfortable in interactions with a statistician where one is involved. This is an ambitious list, but it is not the same as trying to turn participants into statisticians.

If training is for longer than 4 weeks then it will usually be broader than the course envisaged here. It is likely to include design, data management and analysis and also to consider the analysis of survey and monitoring data, as well as data from experiments. In this case, these notes could be used within the corresponding sessions.

Notes from individual sessions have also been used within statistics training courses for MSc and PhD students, who are starting their career in agricultural research.

If 4 weeks is the longest time that might be planned for this training, then what is the minimum? There is no minimum. The analysis course includes a one-hour review of the design course. This was a two-week course, so perhaps a one-hour seminar on analysis is a practical minimum for a formal presentation. One session at the end of the 'Everyday Toolkit' section is a 'Where are we now?' review, that is about this length. We also have a booklet on analysis that serves as a brief review of the main concepts in the course as a whole.

A one-hour seminar could be for a participant's own institute, or for a particular person. Within one institute a scientist gave his Director General a one-hour review of the importance of the data management and statistical issues for the effectiveness of the research of the organization. It is useful to keep your director informed, but not realistic to expect a director to spend one week on a training course.

Teaching style

Lecturing is kept to a minimum in this course.

It is planned that a typical half-day session would start with a lecture/demonstration of between 20 and 45 minutes. The main part of the session would then be devoted to practical and discussion work. Participants would usually divide into groups for this work, with each group having a defined set of tasks. Finally there would be a review period with a representative from each group presenting their findings.

Normally two participants are nominated each day to prepare a brief summary of the work undertaken and particularly the key points that were covered in the day. This is used to review the key points that have been introduced and also to build a report of the work actually covered in the workshop.

In the basic part of the course participants will normally conduct the full analyses of the data themselves. Most of the practical work is generally specified and participants have to translate the tasks into instructions for the statistics package. Reference documents on the statistical software are available to help with this translation.

In the further analyses, the style of the practical work is different, because the limited time does not allow for participants to become embroiled in computational difficulties. Where this is likely, an alternative is for participants to dictate the analyses, which either one of them, or a resource person then undertakes.

The notes in this manual are extensive, but this does not mean that all sessions must be included in the training. On the contrary, the existence of the notes permits resource people to omit coverage of certain topics, while still providing the necessary reference materials for the participants. We discussed different general scenarios in the section on the duration of the training and here we consider the potential problem that some sessions may take longer than foreseen and hence other sessions may have to be omitted.

The obvious sessions to omit are those in the 'Handling Complexities' part of the course. One possibility is to include all in the basic part and then move straight to Session 16 on the 'own data analysis' and Session 17 on the tasks following the workshop. If there is time for one further session then Session 15, which is on complications in the data, can also be included in the teaching.

In one course, where participants were not very experienced with computers and were finding many ideas to be new, they were split into groups, depending on how fast they wished to proceed. The slowest group only managed up to Session 4, plus Session 9, so they were at least able to present some results from their research.

Unlike the design course, this training needs computers and there must be sufficient time and computers for participants to have considerable hands-on practice. One computer between 2 participants is a minimum and often better than one each, because it encourages more discussion during practical work.

In versions of the training where participants have considerable time to process their own data it becomes useful to have one machine each, though this is not always practical.

There should also be a projection device that allows the whole group to view a computer screen, so analyses can be discussed jointly.

Participants can use an overhead projector for presentations. They can be prepared using presentation software, such as PowerPoint, or they can use the statistics package, or other software directly. Ideally the computers would be networked so the presentations that use a computer can use the same computer on which the materials were prepared.

Some of the work is carried out in small groups. Groups can work in different parts of the same room. Ideally there would also be additional small rooms for these discussions where needed. Again an ideal would be for one computer to be available per group, within each of these rooms.

Software

A powerful statistics package is required. The requirement is also for a package that minimizes the time that must be devoted to mastering its use. The aim is to teach the concepts and the practice of data analysis, not the skills of using a particular package.

We currently use GenStat as the most effective package for the analysis of experimental data. The current version (5^{th} release for Windows at the time the documentation was written) is also much easier than previous versions for data exploration and descriptive statistics, but it still remains more difficult than Excel for some of these tasks.

The course notes have been written to be as generic as we felt possible without making them more difficult for the reader. We expect resource persons to adapt the notes in general and adaptation for a different statistics package should not be difficult. Ideally the chosen package should include facilities to handle multilevel models (REML in GenStat and PROC Mixed in SAS). However this is a minor part of the course and therefore other packages could be considered.

In addition to a statistics package, Microsoft Excel is used here for data management, plus Microsoft Access and the ICRAF produced Logbook software for some courses. A word processor (possibly Word) and software for presentations (possibly PowerPoint) are also needed.

Course Content

Summary

Here the content of the course is summarized. This summary should be explained to participants at the start of the course as part of the introduction to motivate their involvement. This motivation is important for participants who have attended training courses on data analysis before and are still unsure of their capabilities.

If this training is put into perspective with the other components we find that:

a) With the design course, many participants felt that initially it missed their main priority, which was analysis. By the end of the design workshop they realied the importance of the design phase and had mastered the main concepts.

b) With the data management, some scientists felt initially that this was an underlying problem, but not as pressing as analysis. Others felt it was not a great issue and was something mainly for the technicians. The training period was again sufficient to impart the key concepts.

So now we start the training on analysis, which most scientists feel is their real need.

The purpose of this course is to emphasize concepts. It minimizes the formulae, they are needed more by the statistical software than the user. You should also minimize the time devoted to the statistical computing. This point may seem curious, because extensive use will be made of statistical software for the analysis. However, the software has become so much simpler to use, that the emphasis can remain on the ideas of the analysis and not on how to use the statistical package.

In the design workshop we saw that the objectives of the study led to the treatments to be applied, to the layout of the trial and to the measurements that needed to be taken. These 4 components (objectives, treatments, layout and measurements) are also the key to this workshop. Our aim here is to analyse the measurements, taking account of the layout and treatments, so we satisfy the objectives.

We will see that the requirements of an effective analysis are clear and often simple, if we concentrate on the analysis that is required to satisfy the objectives. Too many courses in the past have concentrated on helping participants to understand the analysis that is supplied by the statistical package. This has to be understood, but should not be the major concern of the course. We should assess what is needed for the analysis and demand it from the software, not look at what is provided and hope that it is needed.

By Session 3 we will see that the types of presentation of the data that are needed to satisfy the objectives are usually quite simple and also easy to provide. In Session 4, which is devoted to exploratory analysis, we show that while simple summaries are important, it is also necessary to look for features of the data that might alter the way we tackle analysis or even modify objectives.

We will find that the objectives rarely imply the need for an analysis of variance table. So what is the role of an analysis of variance? We answer this question in Sessions 5, 6 and 7. Perhaps it is not surprising that many scientists do not make much use of the ANOVA table, since its role (in relation to the objectives of the trial) is unclear. Considering the measurements as:

data = pattern + residual

is a useful idea. The ANOVA table is a useful way of summarizing all the components of the 'pattern'. Part of this 'pattern' concerns the treatment effects and it is these treatments that relate to the objectives of the trial.

In Session 6 we introduce the idea of a 'standard error' and its role in relation to the objectives of a trial. Session 7 shows how the idea of writing

data = pattern + residual as a 'statistical model'

is helpful. It enables users to see that many different problems can be analysed in a similar way, because they have a similar 'model'.

The final main concept of Multiple Levels is introduced in Session 8. If 'everything' is at just one level, then the analysis is usually simple. To discuss what we mean by 'everything', we need to consider our three components of treatments, layout and measurements again. So, the analysis of an on-farm experiment is usually not so simple, because we take measurements at two levels (at least). Thus, in a typical on-farm trial we record some data at the farmer level - we ask them questions - and we take other measurements at the plot level.

The session on the presentation of the results concludes the first part of the workshop. We expect participants to try the ideas on their own data throughout the workshop and will be looking for suitable examples to include in the formal training. Session 17 is a direct challenge to the participants, to discuss whether the workshop will change anything once they leave.

The second part of the workshop, that is titled Handling Complexities, assumes the concepts mentioned above. It considers common complexities about which participants should be aware in order to do an effective analysis. It is not expected that participants will necessarily be able to handle all the complexities unaided. But they should be able to detect their existence. They should also know whether the particular complexity is one that they can handle unaided, or one that requires support.

We have divided the complexities into those that were planned, those that are common because of the particular field of application and those that were 'surprises'. In describing the complexities we again consider them in relation to the three components of our trial. Thus there may be complications in the way the treatments are applied, in the layout, or in the measurements.

When participants assess their own objectives from the workshop, it should be made clear that the sessions, particularly in the second section, could only introduce the subjects concerned. The aim of this workshop is to provide both skills and confidence in analysing data.

Sessions

Part 1: The Everyday Toolkit

1. Review of experimental design ideas

   A quick review of the main concepts and terminology of experimental design.
2. Objectives of analysis

   Identifying objectives for analysis, so that it proceeds in a focused and constructive way.
3. Software familiarization

   The remainder of the course relies on use of a statistical package, which is introduced at this point.
4. Descriptive analysis and data exploration

   Looking critically at data to reveal the important patterns, before any formal inference starts.
5. Analysis of variance

   Analysis of variance is introduced as a tool in exploratory statistics of data with complex structures, typical of experiments.
6. Statistical inference

   Making formal statistical inferences and explaining the role of ANOVA.
7. Models

   Introducing the idea of a statistical model, and the link between models and analysis of variance.
8. Multi-level data

   Data are harder to analyse when they are at multiple levels. The concepts are introduced with split-plot designs and within-plot samples.
9. Presenting the results

   Putting the information from analysis in a format suitable for various audiences.

Part 2: Handling Complexities

10. Review

    We review the first part of the course to set the scene for the second part.

11. Complexities introduced by the design

    Handling common analysis complexities that arise from the design of the experiment.

12. Dealing with categorical data

    We use an on-farm trial in which most of the measurements were categorical to illustrate modern approaches to the analysis of this type of data.

13. Getting more out of on-farm trials and other multilevel problems

    Handling complexities common in data from on-farm trials. We concentrate here primarily on the complexities that arise because of the availability of information at multiple levels.

14. Complexities due to the area of application

    There are currently two parallel sessions and we expect more to be provided in the future.

    Agroforestry experiments typically involve measurements of multiple components over long time periods. Methods for analysing these are covered.

    Mixed models and the analysis of crossover designs are described. They are common in livestock trials.

15. Complexities due the nature of the data

    These are problems that could not be anticipated - missing values, many zeros, non-constant variance and so on.

16. Analysis of participants' data

    This is a session in which participants work on their data. This may well be spread out through the course.

17. Action plan

    Planning the action participants will take on return to their institute.

Resource materials

For each main session we have provided a summary, a lecture note and a practical exercise. These are provided as printable (pdf) files that can be read on-line and can be printed when needed. These materials are also provided as Microsoft Word files for resource staff to adapt as needed.

The protocols for 18 trials are provided, and this includes all those that are used for illustration on the course. This document links to the Excel files that provide the datasets used on the course. The next table summarizes the trials.

Title	Type	Used	Features
Relay planting of *Sesbania sesban* and maize.	Station	Le.: 2	Factorial treatment structure. Unequal replication. Strong trend along field not accounted for by blocks.
Effect of *Tithonia diversifolia* and *Lantana camara* mulches on crop yields in farmers fields.	Type 2	Le.: 1, 6, 8, 13, 15 Ex.: 1, 2, 3, 4, 6, 7, 8	3 treatments, 1 complete replicate per farm.
Screening of suitable species for three-year fallow.	Station	Le:. 7, 11,5 Ex.: 1, 2, 4, 5, 7	Simple design and layout. Repeated in 2 years. Several responses and covariates.
Upperstorey/ understorey tree management trial.	Station	Le.: 4, 5 Ex.: 1, 2, 4, 5, 11	Complex treatment structure. Tree and crop components measured.
Leucaena trichandra seed production trial.	Station	Le.: 8, 11, 15, Ex.: 8	Large trial. Multi-level design and data. Spatial analysis possible.
Fruit trees survival.	Type 2	Ex.: 11,12	Treatment factors of different types. Repeated measures. Binary responses.
On-farm cropping with *sesbania* and *gliricidia*.	Type 2	Ex.: 13	Simple treatments. Not all farms have the same treatments. Multi-level covariates.
Roots and Competition (RAC).	Station	Le.: 14 Ex.: 14	Large trial with simple layout but multiple observations per plot.
Prototype hedgerow intercropping systems.	Station	Ex.: 14	Simple trial with responses measured in 12 successive seasons and a strong rainfall x treatment interaction.
Fertilizer, *Tithonia* and *Lantana* mulch as sources of phosphorus for maize.	Station	Le.: 11	Complex treatment structure. Unequal replication, treatments not orthogonal to blocks.
Calliandra feeding trial.	Station		Simple crossover trial.
Effects of organic and inorganic sources of nutrients on striga, weeds and maize.	Station		Mix of qualitative and qualitative treatments. Split plot design.
The influence of improved fallows on soil phosphorus fractions.	Type 1	Le.: 8, 11, Ex.: 1, 2, 4, 8	Factorial treatment structure. Split-plot design. Covariates measured.
Improved fallows and rock phosphate: farmers' experiences.	Survey, Type 3	Le.: 4, 12, Ex.: 12	Part of a large study with irregular design and qualitative responses.
On-farm trial with improved fallow and inorganic fertilizer.	Type 2/ 3		Factorial treatment structure. Irregular arrangement of treatments to farms.

There is also a series of about 20 'good-practice' guidelines. These are found on-line and printed booklets that are on a range of topics, some of which provide supporting materials for the course, or can be used as pre-course reading. Those on design and data management, relate to topics that are assumed in the present course. The guidelines that are particularly relevant to the analysis course materials are as follows:

Confidence and Significance: Key Concepts of Inferential Statistics

These are topics we would like to assume, but find they are often poorly understood. This guideline is an excellent candidate for pre-course reading.

Modern Approaches to the Analysis of Experimental Data

This provides a résumé of the topics covered in Sessions 1 to 8 of the course.

Informative Presentation of Tables, Graphs and Statistics

This provides an extended version of the material in Session 9, but more general than for experimental data.

Modern methods of Analysis

This introduces some of the key concepts from the second part of the course. It is more general than the course with suggested methods applying equally to the analysis of survey data.

Mixed Models and Multi-Level Data-Structures in Agriculture

This extends the materials introduced in Sessions 13 and 14.

Strategy

In this final section we describe aspects of the workshop strategy highlighted when drafts of this training material were used previously. We hope that the lessons learned will help future resource staff in planning their workshops.

In 1999 the workshop was run in three different countries in Africa. The first was in Cotonou, Benin in September 1999. This was for staff from the national programme in Benin, INRAB. The second workshop took place at ICRAF headquarters in Nairobi and was a regional course to researchers of agroforestry. The third workshop was an in-country training in Conakry, Guinea for researchers from the national programme, IRAG. In Guinea the workshop was given in two phases, firstly to a key group from all stations, and then with members of the key group acting as resource persons, to scientists in individual stations.

In Benin a draft version of the notes was used. Comments on the notes were sent to ICRAF, so that changes could be made for the second workshop. The revisions following this second workshop resulted in the first version of the course notes being produced in early 2000.

This process of revisions is described because we expect further changes in the future. Indeed we would welcome any suggestions for changes, revisions of sessions and additional sessions. One major addition to the notes following the ICRAF course was a plan to include extra materials on a 'Type 3 on-farm experiment'. This is almost a survey and is a useful indication that the same strategy of analysis will also be of some use in the processing of survey type data.

One feature of the Benin workshop was that 1 or 2 participants were nominated each day to produce a report of the day's proceedings. This was presented and discussed as the first item on the succeeding day and was then written up, by the participant, to be included in the workshop proceedings. It was quickly established that these reports should summarize briefly what was covered but then concentrate on the discussion points that were raised, rather than the topics that were already in the notes. The full workshop only covered Sessions 1 to 9 formally and the discussions often raised concerns that are considered in the later sessions. Topics included the role of multiple-comparison procedures, as well as the importance of the coefficient of variation, coping with zeros in the data, and ways of incorporating climatic data into an analysis.

By week 2 these reports became a key element of the workshop. They were of a high standard and typically generated a good debate of whether the essential points of the previous day's proceedings had been understood. They gave considerable confidence to the participants and showed that some of them had sufficient understanding of the key concepts that they could now contribute to training courses themselves. This helped the discussions on future plans (Session 17) , which were also in the second week.

In the earlier design course, participants often divided into small discussion groups. They prepared short presentations that generated a discussion at the end of most sessions. This approach was equally effective here for the sessions that did not involve computer use. However it was clearly a distraction to the participants in the sessions that involved computers and we abandoned the attempt. This was partly because the computers encourage participants to work singly or in pairs, rather than in small groups. It was also because the shape of the tables, in a large rectangle, plus the lack of easy discussion areas close to the machines, did not lend itself to this approach.

Comments on some of the specific sessions are as follows:

In the Benin course the first two days were spent on Sessions 1 and 2, namely on a review of design and of the specification of the objectives for the studies. An important result from these sessions was a realization by the participants that a review of their procedures for evaluating protocols was necessary.

GenStat was re-introduced on the third day, as most participants were already familiar with an earlier version. The changes in the new version were popular to the extent that, by the end of the course, the participants suggested the setting up of a GenStat support group.

Sessions 4 and 5 were two sessions that benefited from the use of local examples instead of those used in the notes. We used a single example in each session, which was copied to each computer.

The ideas of simple inference were (as usual) found to be difficult by the participants. 'Do participants understand the difference between a standard deviation and a standard error' is a key question? Can they explain what is meant by a confidence interval for the mean? These ideas were introduced in Session 6. It was strongly suggested that these topics be prepared as a booklet for pre-course or pre-session reading. (This is now available.)

In week 2 participants spent a long time (1.5 days) on Session 17. They made detailed plans on the in-station training that was needed over the next year. These plans covered design and data management as well as data analysis. It was agreed that this training within the stations was of more immediate priority than further training of this key group.

Acknowledgements

The European Union, World Bank and ICT have funded the courses that have motivated production of this material. Some of the development work at Reading was funded by DFID grants. Some development at ICRAF was funded by grants from the Netherlands Ministry of Foreign Affairs through their Direct Support to Educational Institutions Programme (DSO).

Many individuals have contributed to the materials in assisting with courses and reviewing documents, these include Jane Poole, Marie Rarieya, Jan Beniest, Ian Dale, Wim Buysse, Gaston Kokode and Sekou Beavogui.

Session 1

Review of experimental design

Session objectives

- To review the main concepts of experimental design - objectives, treatments, experimental layout and measurements - in the context of experimental data analysis.
- To remind workshop participants of the key terminology used in experimental design.

Summary

The design of an experiment starts with a statement of **objectives**. These must be clear, complete, relevant to the problem being addressed and capable of being met by an experiment. The objectives determine the three other components of the experimental design:

1. The **treatments**. The treatments are the conditions being deliberately compared in the experiment. The objectives may call for a simple collection of treatments, a factorial set of treatments, a range of quantitative levels of some factor, or combinations of these. The objectives will also determine suitable control treatments.
2. The **layout** of the experiment. This includes the choice of site or sites, of plots within sites and of measurement locations (samples) within plots. It also includes the arrangement of plots in blocks and the allocation of treatments to plots.
3. The **measurements** to be taken. The measurements include those taken directly to meet the objectives of the experiment and those taken to help understand the variation or pattern in other measurements.

Both on-farm and on-station trials can be described in terms of objectives, treatments, layout and measurements.

Activities

Introductory lecture

A short lecture to review the main features of design and remind participants of the key terminology used in experimental design. Point out that this is a one-hour review of material that has previously been covered in a two-week course.

Practical Exercises

Participants meet in small working groups and look at some experimental protocols; one provided by the workshop resource persons, and one, or several of their own. The purpose of this exercise is to see that all participants understand the basic concepts and key terminology of experimental design and that they can apply this to any particular trial.

Start with a simple protocol provided by the course. The following four are suitable:

1. 'Relay planting of *Sesbania sesban* and maize'
2. 'Effect of *Tithonia diversifolia* and *Lantana camara* mulches on crop yields in farmers fields'
3. 'Screening of suitable species for three-year fallow'
5. 'Upperstorey/understorey tree management trial'

The group should follow the instructions given on the practical sheet. The group reports back in a plenary session with a *brief* summary of the 4 elements of the design, and their comments on the completeness and clarity of them. Emphasize that it is not necessary to repeat the whole of the background to the trial. The group discussion needs to be no more than 10 minutes, plus a 3-minute report for each group. Plenary discussion will then be used to clarify any design points. This part should only take 30 minutes because the protocols provided by the course must be clear and simple.

Next the groups look at a participant's protocol again, following the practical sheet. The group reports back in a plenary to describe the problems they had understanding the experiment and *not* to present experimental details. This part will take longer than the first, since experiments tend to be poorly described for outsiders and thus may yield quite some discussion when it comes to understanding the design and its objectives. This part easily takes 1 hour to complete.

Supporting documents

- R. Coe, L. Nelson and R. Stern. 1996. *Design of agroforestry experiments.* Guidelines for training workshop resource persons. ICRAF/World Agroforestry Centre.
- R. Coe, R. D. Stern, E. Allan. 2002. *Session 1 - Review of experimental design*. Lecture notes. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Coe, R. D. Stern, E. Allan. 2002. *Experiments portfolio.* (Part 4). Data analysis of agroforestry experiments. ICRAF/ World Agroforestry Centre.
- R. Coe, R. D. Stern, E. Allan. 2002. *Session 1 – Review of experimental design*. Exercise guidelines (Part 3). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.

Session 2

Objectives and steps in data analysis

Session objectives

- To list and describe the sequence of steps involved in the analysis of an experiment.
- To determine objectives for analysis of data from an experiment based on the objectives of the experiment.
- To describe the tentative numbers, tables and graphs that will be needed to meet the analysis objectives.
- To describe the logic behind the training workshop structure and content, in the context of analysis objectives.

Summary

The general steps in the analysis of an experiment are:

1. Determining the objectives of the analysis.

2. Preparing the data by making the data files suitable for the planned analysis.

3. Descriptive analysis:
   i. Calculation of summary tables and graphs as earlier defined when setting analysis objectives.
   ii. Exploratory analysis to identify any unexpected patterns or results.

4. Confirmatory or formal analysis:
   i. Ensuring the planned tables, graphs, etc. are appropriate.
   ii. Adding measures of precision.
   iii. Improving the estimates of critical quantities.

5. Interpretation of analysis results.

6. Reporting of experimental results.

Defining the objectives of analysis will involve:

- Identifying the exact comparisons to be made or relationships to be estimated.
- Determining the exact data that are needed to make them (e.g. do we need comparisons of yield for each season or totaled over all seasons?).
- Designing the tables and graphs that will be used to present the results.

The objectives of analysis are determined by the objectives of the trial. However the analysis objectives are distinct from trial objectives in that:

- The objectives of the trial may have been stated in a rather vague way.
- The objectives set out in the original protocol may have other, unstated, objectives added, if these can usefully be met with the data available.
- It may not be possible to meet all the original objectives of the trial, either because the trial design does not allow it or because something unexpected has happened to prevent it.

The analysis objectives will evolve as the analysis proceeds.

Activities

Introductory lecture

A short lecture in two parts to:

a) introduce the steps in experimental analysis, and
b) link the setting of analysis objectives to experiment objectives.

Practical exercises

Participants meet in small working groups and look at some experimental protocols, one provided by the workshop resource persons and one, or several of their own. The purpose of these exercises is to make sure that all participants can describe analysis objectives based on an experimental protocol, and apply this to any particular trial. Groups and protocols should remain the same as for Session 1 so that participants are already familiar with these trial designs.

The first exercise uses a protocol provided by the course and should take about 1 hour to complete.

The second exercise is for participants to look at their own experimental protocols that will be used in the data analysis sessions and determine the analysis objectives for their own data.

Exactly how this second exercise will be completed depends on what participants have brought with them to the training workshop. If everyone has a suitable protocol and data set, participants can work in pairs, with each pair looking at the two protocols. A selection of participants are requested to present their findings in a plenary session through a brief presentation. Participants are encouraged to start a file for the analysis of their own data, which by the end of the course should contain a complete record of what they did with the data, starting with the analysis objectives. This second exercise should take approximately 1.5 hours to complete.

Concluding remarks

In a final plenary session, workshop organizers explain how the remainder of the workshop is structured in terms of the steps in data analysis. The 'Everyday Toolkit' part of the training workshop (Sessions 1 – 9) closely follows the steps outlined; proceeding through descriptive statistics and formal statistics to the final reporting stage. The 'Handling Complexities' part (Sessions 11 – 16) merely extends the range of situations that can be handled by both the descriptive and formal analysis methods.

Supporting documents

- R. Coe, R. D. Stern, E. Allan. 2002. *Objectives and steps in data analysis*. Lecture notes (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Coe, R. D. Stern, E Allan. 2002. *Objectives and steps in data analysis*. Exercise guidelines (Part 3). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Coe, R. D. Stern, E. Allan. 2002. *Experiments portfolio*. (Part 4) Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.

Session 3

Software familiarization

Session objectives

- To become familiar with the main statistics software package (GenStat) that will be used during the training workshop.
- To list the advantages of this software package in the analysis of agroforestry experiments.
- To review the role and use of supporting software needed to facilitate data analysis using this main statistics software package.
- To transfer data and results between these software packages.

Summary

The first part of this session is devoted to practical work, introducing 'GenStat for Windows' to the participants. This is the only introductory session on the software. The other sessions will deal with the practical use of GenStat for the different types of analyses.

The lecture describes the key points that make GenStat the obvious package to be used for this training workshop. This includes GenStat's powerful and flexible facilities for the analysis of simple experimental designs, facilities that are used in Sessions 4 to 8. The substance of GenStat is consistent with the structures discussed in the workshop. We have already mentioned the fact that an experiment is characterized by the treatments structure, the plot layout and the measurements that are taken. This translates into the treatments, blocks and data for analysis, within GenStat's ANOVA system.

Many datasets contain complications and this is the theme of the sessions in the second part of the workshop. It is therefore essential to use a powerful statistics package and GenStat's regression and REML facilities to demonstrate this power.

The workshop emphasizes the need for software that permits analyses that are dictated by the objectives of the research, without great effort by the user. GenStat is appropriate,

although its facilities for data manipulation and plotting are not as intuitive as those of Excel.

The third exercise of the session is for participants to look at their own data. This is to be imported into GenStat if possible and then used to consolidate the simple use of GenStat described earlier.

Activities

Introductory Lecture

A brief five-minute lecture is conducted, introducing GenStat to participants, and giving an outline of the session. The Introduction Section of the lecture note for Session 3 could be used for this.

Practical exercises

Exercise 1—Introduction to GenStat

The purpose of this exercise is for participants to follow the introductory tutorial on GenStat as described in the GWIM software manual.

Review/demonstration

Answer any questions/problems participants may have encountered during the previous GenStat tutorial.

Resource person to give projected computer demonstration on importing data into GenStat and using the **'restrict'** command. Within the demonstration of data importing it would be useful to point out to participants that the easiest way to import from Excel to GenStat is to mark the data range initially in Excel. You highlight the required data, go to **Name Box** in the top left corner of the sheet and write in a name for the range (e.g. 'data'). Then in GenStat you just select the worksheet that contains this range (i.e. it will appear as a sheet named 'data').

Use the 'Screening of suitable species for three-year fallow' trial for the demonstration, restricting data to look at only one season (e.g. 1991). Demonstrate how data can be restricted using different menu commands, and by rows or columns.

Practical exercises

Exercise 2

In this exercise participants are able to have a go at importing data into GenStat and using the 'restrict' command, as demonstrated previously. This exercise may be omitted if time is short, as participants will be using these procedures throughout the rest of the course.

Lecture/Discussion

A lecture and/or discussion on the structure and facilities of GenStat (using the lecture notes provided) will be conducted. This could include a projected computer demonstration of some of GenStat's features. The use of GenStat for data analysis and its main features that justify its use for analysis of data in agroforestry experiments are discussed.

Practical exercises

Exercise 3—Datasets for analysis

Participants look at their individual datasets and discuss how these need to be prepared for further analysis in the context of the use of GenStat for this purpose. Their data should be appropriate for being imported into GenStat by the end of this session.

Concluding remarks

The final part of this session provides an opportunity for workshop organizers and participants to discuss the use of the various software packages in the context of data analysis.

One final additional remark to make is that in previous workshops participants have requested more time on the specific use of GenStat. This should not usually be necessary, but could be accommodated by a second GenStat practical, given later in the course, which would cover the remaining sections of Part 1 of GWIM. However, the preferable option is for participants to use some of their spare time during the workshop to try all, or part of the ANOVA section of GWIM. The GWIM manual contains 'fast track' options, which could be tried by participants who are more proficient at computations .

Supporting documents

- R. D. Stern. E. Allan, R. Coe. 2002. *Software familiarization.* Lecture note (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Statistical Services Centre, University of Reading, U.K. 1999. *GenStat for Windows Introductory Manual.*
- R. D. Stern, E. Allan, R. Coe. 2002. *Software familiarization.* Exercise guidelines (Part 3). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Mead, R. N. Curnow, A. M. Hasted. Chapman and Hall. 1993. *Statistical methods in agriculture and experimental biology (2nd Edition).*

Session 4 Descriptive analysis and data exploration

Session objectives

- Demonstrate that tables and graphical summaries can reveal patterns and trends related to the objectives of the analysis and study.
- Show how data exploration can identify unexpected patterns, and/or observations.
- Encourage close inspection of the data and interpretation of exploratory methods, prior to formal statistical analysis.

Summary

Data analysis is more than just formal methods such as analysis of variance. An important part, which is all too often ignored, is initial investigation of the data. This preliminary investigation allows the researcher to look closely at the information in his data. The necessary tools include tables and graphs of descriptive statistics and the use of exploratory methods such as **boxplots** and **scatterplots.**

With descriptive statistics the researcher can quickly explore the patterns of interest, such as the treatment effects. If the tables and graphs needed to address the objectives have already been identified it is easy to 'fill in the numbers' in the presentation, and observe the pattern of response.

Careful examination of the data also gives an insight into the variability of the treatment pattern, and highlights any **outliers** that need further clarification. Suspected patterns such as block effects or fertility effects can be explored, and sometimes, unexpected patterns will emerge which the researcher may want to explore further.

Activities

Introductory Lecture

A short lecture, showing the role of descriptive summaries and exploratory methods as useful tools for close preliminary investigation of data shall be given. The idea of data = pattern + residual is introduced and the methods used to examine both 'pattern' and 'residual' are discussed. The data from the 'Upperstorey/understorey tree management' trial are used to illustrate the different points.

Practical Exercises

Exercise 1 - Computer practical where participants work in pairs on one of the four example datasets used in previous sessions with a view to familiarizing themselves with the data structure, identifying patterns which are associated with the objectives of the experiment and other possible patterns which might be expected. Consideration should also be given to variability, and detection of 'strange' values.

This is followed by short presentations by four pairs of participants, one pair for each dataset. These presentations should cover:

(a) The questions related to the objectives of the study - completed tables or graphs, methods used and findings.

(b) Other interesting questions that one wanted to explore and what they showed. Also any unexpected findings.

(c) Some assessment of variability in the data.

Round off with a discussion to elicit what methods are useful for different types of investigations and data. There should also be some discussion as to which other questions should be addressed.

Exercise 2 – Descriptive and exploratory analysis of participants' own data. This is a continuation of the second practical of Session 2 where participants identified the objectives of analysis in relation to the objectives of the study.

Concluding Remarks

Some wrap-up by resource person of techniques concerning 'problems encountered, where do we go from here?' is needed at the end of this session.

Supporting documents

- R. D. Stern. E. Allan, R. Coe. 2002. *Descriptive analysis and data exploration.* Lecture note (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Statistical Services Centre, University of Reading, U.K. 1999. *GenStat for Windows Introductory Manual*.
- R. D. Stern, E. Allan, R. Coe. 2002. *Descriptive analysis and data exploration*. Exercise guidelines (Part 3). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Coe, R. D. Stern, E. Allan. 2002. *Experiments portfolio*. (Part4). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.

Session 5 Analysis of variance as a descriptive tool

Session objectives

- To understand the role of the Analysis of Variance (ANOVA) table.
- To understand and use residuals.
- To investigate how to reduce unexplained variation in the data.

Summary

The tables and graphs required to meet the objectives of an analysis rarely include an analysis of variance table. This session explains the role of the ANOVA table in the analysis of experimental data. It is introduced here, as part of descriptive statistics, as a technique to look at the whole model (blocks and treatments). It is used to identify the components that are important. If we consider that

data = pattern + residual,

then the ANOVA table extracts all the pattern, and hence also provides an opportunity to look at the residuals, i.e. the part of the data that can not be explained by the model, or pattern, in the experiment. Furthermore the pattern may be broken down into parts so that we can understand the relative importance of different components of it.

Activities

Introductory Lecture

Lecture on the use of the ANOVA table, explained using a simple example. This lecture should be conducted as an interactive session.

Practical

The first part of the practical repeats material covered in the lecture. It should not be necessary for participants to report findings from this.

The second part asks participants to use descriptive ANOVA on their own datasets. Reporting conclusions from this is valuable. It is likely that 'inference' results will be reported, giving an entry point for discussing the distinctive roles of exploratory and formal analysis and the extent to which their analysis objectives are met by the former.

Supporting documents

- R. D. Stern, E. Allan, R. Coe. 2002. *The analysis of variance as a descriptive tool*. Lecture note (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. D. Stern, E. Allan, R. Coe. 2002. *The analysis of variance as a descriptive tool*. Exercise guidelines (Part 3). Data analysis of agorforestry experiments. ICRAF/World Agroforestry Centre.

Session 6 Ideas of simple inference

Session objectives

- To review the main concepts of statistical inference.
- To show what statistical inference adds to the descriptive summary.
- To show how the use of statistical inference can assist in the realistic interpretation of experimental results.
- To introduce the assumptions necessary for inferences to be valid.

Summary

There are a few key ideas that are needed to understand the basics of statistical inference. First is the idea that there are parameters, which describe important characteristics of populations. These have to be estimated from data. Estimates have properties, the most important being the precision of the estimate. The standard error is a measure of precision of an estimate. A confidence interval is a convenient way of interpreting an estimate and its standard error. Statistical tests are used to test hypotheses about parameters, and result in a significance level that indicates the amount of evidence against the hypothesis.

Standard output from running analysis of variance on data from an experiment allows us to make inferences by (a) providing information on precision, and (b) making it possible to test certain hypotheses about the 'pattern' part of the data. The hypothesis testing can be useful in helping to decide which components of the 'pattern' we wish to concentrate on. It usually has limited value in directly interpreting treatment effects and meeting analysis objectives.

Contrasts (comparisons of carefully selected treatment combinations) are introduced to help meet specific analysis objectives.

Activities

Lecture

The main concepts of statistical inference are reviewed in a short lecture of approximately 45 minutes. Only 5 to 10 minutes are given to the general ideas. The rest of the time is devoted to a demonstration and discussion of the concepts for an on-farm example. The lecture note describes this case study, so it can be used for reference afterwards.

Practical exercise

Exercise

The practical has the same objectives as the lecture, namely to consolidate the ideas of statistical inference and to apply these ideas to some real examples.

Presentation of results and discussion

The reports review the conclusions from each case study, with particular emphasis on the extra component that is added by the statistical inference, compared to the use of descriptive statistics.

Supporting documents

- R. D. Stern, E. Allan, R. Coe. 2002. *Ideas of simple inference*. Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Statistical Services Centre, The University of Reading, UK. 2001. *Confidence & Significance. Key Concepts of Inferential Statistics.* Booklet.
- Practical sheet

Session 7 An introduction to statistical modelling

Session objectives

- To introduce the concept of a statistical model for describing data.
- To practice fitting models and interpreting the resulting output.
- To demonstrate the relationship between a modelling approach to data analysis and simple ANOVA methods for balanced data structures.

Summary

The key points for this session are:

- Data can be explored and variation explained by building a model with design effects (representing the layout and treatments) and other potential sources of variation.

- The modelling approach is flexible and powerful, overcoming some limitations of simpler methods (ANOVA) and providing the basis of many more advanced methods used to handle various complexities.

- The steps in using a model for data analysis are: exploring data to find a suitable model type, fitting the model, checking the assumptions of the modelling, and interpreting the fitted model. The process is iterative.

Models with few effects are not difficult to fit with current software, but more complicated models require some careful model building. This session introduces a very broad and deep subject and all future sessions refer back to it.

Activities

There is a lot in this session and many new ideas are introduced. It will therefore be long (maybe a whole day) and should be broken down into a number of sections. A reasonable sequence would be:

Lecture/demo 1: covers the simple regression fitting and interpretation and 'Steps in modelling' sections.

Part 1 of the practical: this reinforces the material from the lecture.

Lecture/demo 2: Review the practical, then proceed to the sections on 'Models with factors ' and 'Analysing other designs'.

Part 2 of the practical: this repeats the lecture material. Participants may well have their own examples to use here.

Lecture/demo 3: covers the remaining material.

Part 3 of the practical: note that the methods needed to answer the first question have been covered. Those needed for the second part can be introduced through discussion once participants have worked out what they are trying to achieve. The ideas of 'comparison of regression lines' can be introduced in a general discussion at the end of the practical if participants are up to it!

It is worth getting participants to present results from each part of the practical.

Supporting documents

- E. Allan, R. Coe, R. D. Stern. 2002. *An introduction to statistical modelling*. Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Statistical Services Centre, University of Reading, U.K. 1999. *GenStat for Windows Introductory Manual*.
- Practical instructions.

Session 8 An introduction to multiple levels

Session objectives

- To introduce the concept of hierarchical structure in data.
- To introduce the idea of variability at the different levels and variance components.
- To demonstrate these ideas and show how they affect the analysis through the use of a split-plot experiment.
- To recommend ways of dealing with different complexities associated with multiple levels.

Summary

The layout of many of the experiments involves multiple levels or layers, for example sites, blocks, plots and trees. Treatments may be applied to one or more of these levels, and measurements made on one or more of the levels. Some of the problems of analysis of data from experiments arise because of the multiple layers in the design. However, there are situations in which the multiple layers need not lead to complex analysis. Two are considered here, others will be covered in Session 11.

In a split-plot experiment, treatments are applied at two levels but measurements are made at only one (the sub-plot). Provided the design is orthogonal the analysis is straightforward and the only complication is that the precision of treatment comparisons differs, depending on the particular comparison.

When treatments are applied at one level, but measurements are taken at a level lower than the treatment, then a split-plot type of analysis can also be performed. An alternative is to summarize the lower level data to the level at which treatments were applied and carry out an analysis at that level.

When the data can effectively be summarized to a single level the 'general linear model' approach to analysis is appropriate. Models that have more random (error) terms are needed for more complex situations.

Activities

Lectures

Lecture 1

The first lecture of approximately 30 minutes covers the analysis of the split-plot experiment (via interactive demonstration) and through it, the ideas of hierarchical structure and variance components. This uses MCH example pages 133-136. The lecture is followed by the first practical.

Lecture 2

The second lecture of 45 minutes deals with other scenarios involving measurements at different levels. and layout at different levels will all be discussed via interactive computer demonstrations.

Practical exercises

Exercise 1

Hands-on computer practical work so that participants can understand how to perform correct split-plot analyses using GenStat software of about 45 minutes.

Exercise 2

Further practical work of one hour, dealing with multiple observations per experimental unit and layout issues.

Supporting documents

- Statistical Services Centre, University of Reading, U.K. 1999. *GenStat for Windows Introductory Manual.*
- E. Allan, R. D. Stern, R. Coe. 2002. *An introduction to the ideas of multiple levels.* Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- R. Mead, R. N. Curnow, A. M. Hasted. Chapman and Hall. 1993. *Statistical methods in agriculture and experimental biology (2nd Edition).*
- Practical guidelines.

Session 9

Writing up and presenting results

Session objectives

- To show how to present results of analyses in tables and graphs.
- To show how to describe the statistical methods used.
- To show how to include results of analyses in a text.
- To remind participants that the analysis, and write-up do not end with the presentation of statistical results.

Summary

It will usually be necessary to convey the results of your analysis to other people. The traditional way to do this is through writing of a scientific paper or report. Some of the challenges in this are:

- Keeping the statistics in perspective. The statistical analysis is a means to meet the objective, not the reason for the work. However some detail of statistics is needed to validate your conclusions.
- Designing and drawing effective tables and graphs. These are included to 'tell a story' and therefore must be constructed so that the story is clear.
- Reporting on methods in sufficient detail to make it clear what you have done without taking up too much space.
- Describing the results in the text so as to add to, not repeat, what is in tables and graphs.
- Designing the whole report to make it suitable for the intended audience.

Activities

Lectures

Lecture 1

The first lecture of one hour, is given by a science writer or editor, covers key elements and common mistakes made in scientific writing.

Lecture 2

The second lecture of 30 minutes deals with the important points of including statistical results in reports, papers and presentations.

Practical exercise

Challenge

Challenge to participants which asks them to:

i. take one objective from one of the studies they have been analyzing, and
ii. prepare a report and presentation which meets that objective, but is limited to a single graph or table and a maximum of 100 words of text.

The aim of this challenge is to force participants to think about the most effective way of describing the results while being concise. The reports have to be complete, presenting all the information necessary to meet the objective but can only include necessary information. Participants are told that it is not necessary to include any of the background or methods used in the study. How hard the challenge is depends on the study and objectives chosen.

At this stage in the course it would be good if everyone were to complete the challenge and present their results, but that may not be possible due to time constraints. It will also be necessary to decide whether the challenge is issued using the participants' own data or examples provided during the course. Using participants' own data is preferable.

The output can be presented in two ways: a single printed page containing the text and table/graph or a single transparency. Each participant makes a presentation using the single transparency [and limited to 100 words???] and hands in the single page. The work could be assessed and small prizes awarded. It would certainly be useful to have some resource people other than statisticians (a presentation specialist and one or two senior scientists) in the session to comment on the output. Participants should be allowed 3 – 4 hours to complete this challenge.

Supporting documents

- R. Coe, R. D. Stern. E. Allan. 2002. *Writing up and presenting results*. Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre
- Statistical Services Centre. University of Reading, UK. 2000. *Informative presentation of tables, graphs and statistics.*

Session 10 Where are we now?- Review of basic statistics

Session objectives

- To review the concepts presented in the previous sessions.
- To develop a framework for the methods of analysis and to introduce the remaining session topics.

Summary

The first part of the course has shown scientists that analysis of much of the data arising from many experiments can be handled effectively by them if:

- They have a clear idea of the objectives of the analysis.
- They have a well-prepared data set ready.
- They have access to and expertise in using a suitable statistics software.
- They are familiar with some key concepts and methods.

The second part of the course focuses on complexities – in layout, treatments or measurements – that may make it difficult or inappropriate to use the simple methods discussed so far. The aim of the course now changes, with the emphasis of helping scientists to recognize when they have a complex problem, whereby they understand the approaches, but not necessarily all the details that are appropriate to the analysis.

Activities

Introduction

A brief introduction of 5 – 10 minutes to explain the objectives of this review session in the context of what has been covered so far and what will be covered under the remaining sessions of the training workshop.

Practical exercise

Discussion groups

Divide into 7 groups each of 2 to 3 participants that will review the material of the following sessions of the training workshop:

- Session 2: Setting objectives
- Session 4: Descriptive statistics
- Session 5: Analysis of variance
- Session 6: Statistical inference
- Session 7: Modelling
- Session 8: Multiple levels
- Session 9: Presentation

Presentations

One person in each group makes a five-minute presentation of the key points covered in each session using not more than three up to a maximum of five slides. A presentations chairperson is instructed to keep exact time for this session. Presentations should be made as interesting as possible.

For each of these presentations, a chief discussant, selected from another group, comments on the presentation using a single slide (one to two minutes) and leads a five minute discussion on the presentation. Chief discussants must have seen the slides of the presentation before it is presented and focus their presentation and the discussion on whether they agree with the key points and indicate which important ones have been omitted.

Lecture/demonstration

At the end of the presentations and discussions, the outcomes of what has been covered until now are presented. Using some examples indicating how the methods introduced so far can be applied, this lecture further shows what remains to be covered in the remaining sessions of the training workshop if data are to be exploited fully. The time frame for presenting these outcomes should take about 30 minutes.

Supporting documents

- R. Coe, R. D. Stern and E. Allan. 2002. *Where are we now? Review of basic statistics.* Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Practical sheet

Session 11 Design and analysis complexity

Session objectives

- To understand how features of the design can lead to complications in the analysis.
- To understand some methods for spotting these complications.
- To learn how to handle some complications with a statistical model.

Summary

Design complexities - that is, complexities of analysis, which are a consequence of the design use, and could be predicted without looking at the data - are common in real experiments. They may be due to the layout, treatments or measurements.

It is necessary for the analyst to (a) understand that the feature of the design may lead to a problem in analysis, or at least require an analysis other than the simple ones used so far, (b) know how to spot if the problem exists, and (c) understand strategies for handling the problem.

In many examples we can find a model, which elaborates on the simple models used so far and which handles the complexity. In this session we introduce the examples of more advanced treatment structure models to handle complex sets of treatments, the 'generalized linear model' to handle data which is not measured on a continuous scale, and show how complex layout/ treatment combinations can be tackled with 'general linear models'.

Activities

Lecture

A lecture is given covering different complexities of measurements and layout. More complex models than the general linear model are introduced; in particular generalized linear modelling and REML models, both of which will be demonstrated by computer analysis.

There is a lot of material in the lecture note. Given 'as is' this would be more than a single lecture. If the material is interspersed with demonstrations and discussions of many of the points that are not covered in detail, then there is enough here for several sessions. Resource people will have to determine how much time to spend on the material according to the interests and skills of the participants.

Practicals

There are two practicals, each aimed at addressing particular objectives.

Practical 1 introduces some of the skills, which will be needed for complex analysis. All participants need to have some exposure to these.

This is a hands-on computer practical with participants working in pairs, taking approximately 1 hour.

Practical 2 gives participants the opportunity to attempt a more complex analysis with the assistance of a resource person.

There are two objectives here:

i. to understand the complexity of the problem and how it can be tackled,

ii to be able to interpret output from complex analyses.

This work should therefore be done in groups, with each group including a resource person, so that participants can learn from discussion. Participants should be allowed 1.25 hours to complete this task.

Supporting documents

- R. Coe, R. D. Stern, and E. Allan. 2002. *Design issues which add complexity to the analysis.* Lecture note. (Part2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.
- Statistical Services Centre, University of Reading, U.K. 1999. *GenStat for Windows Introductory Manual.*

Session 12 Dealing with categorical data

Session objectives

- Explain what categorical data are and what the problems for analysis are.
- Explain the special case of binary data.
- Explain possibilities and shortcomings of inference from contingency tables.
- Give an introduction to logistic regression.
- Give options for analysis of other categorical data.

Summary

The session looks at datasets in which the responses of interest fall into a small number of distinct categories (such as yes/no, dead/alive, or red/blue/green) rather than being measured on a continuous scale.

Summary and descriptive tables for such data were described in Session 4. In this session, models for describing variation in these data are developed and interpreted. The logistic regression model is appropriate for binary data. Other types of categorical response variables can sometimes be analysed by transforming them to one or more binary responses.

Activities

Lecture/demonstration/discussion. The discussion is carried out following the lecture note.

Practical exercise.

Follow up discussion as needed.

Supporting Documents

- E. Allan, R. Stern, R. Coe. 2002. *Dealing with categorical data.* Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre.

Session 13 Getting more out of on-farm trials and multilevel problems

Session objectives

- To understand that the multiple levels of variation common in data from on-farm trials can lead to analysis complexities.
- To be able to diagnose the existence of these problems.
- To find a suitable analysis strategy, including the use of models that represent the multiple levels of variation.

Summary

Data from simple experiments are analysed by comparing the amount of variation in the 'pattern' with the random variation, this being the plot-to-plot variance not due to treatments or other known factors. In on-farm trials and other complex designs there is more than one type of variation (e.g. within and between farms). The analysis is therefore likely to need to use an analysis that recognizes these.

The complexity requires the following steps.

1. Understanding that the problem might actually exist.
2. Spotting if it really is a problem. This might be by suitably chosen diagnostic tables, or perhaps from unexpected error messages that the analysis software produces.
3. Finding an analysis that either:
   - avoids the problem by 'moving the analysis to a single level', or
   - allows for the problem by using a model that includes several variation terms, fitted using REML.

Activities

Lecture

A lecture is given that shows that multiple levels of variation may give rise to complexities in analysis and the strategies that can be used to overcome it. The simple strategy of 'move everything to one level' can be explained as one that is effective in many cases, but not universally applicable. The alternative that builds a model is similar in scope to the strategies used for design complexities in the previous session.

Practical exercise

Exercise

The exercise sets out some precise questions to be tackled. Participants should be encouraged to analyse the root of difficulties they have and try to understand what the analysis options are doing to overcome them.

Participants will need help with GenStat, but should clearly articulate what they are trying to do before being shown the commands to accomplish it.

Supporting documents

- *R. Coe, E. Allan, R. D. Stern. 2002.* Getting more out of on-farm trials and multilevel problems *Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre*

Session 14 Complications in agroforestry trials

Session objectives

- To recognize some common characteristics of agroforestry trials that may complicate analysis and find strategies to cope with them.
- To understand the nature of 'repeated measures' in an experiment and simple methods to analyse them.
- To discuss approaches for dealing with the multiple components from some agroforestry trial data.

Summary

Agroforestry trials have some common characteristics that may complicate the analysis compared with more 'traditional' agronomy experiments. Three are examined in this session.

1. They tend to be long term, giving rise to long series of observations from the same set of plots.
2. They often involve measurement and comparison of different parts of a single plot (e.g. under and away from a line of trees).
3. They generate multiple products.

The first two are aspects of the same statistical problem - measurements repeated in space or time. If this data structure is not recognized then invalid conclusions may result. There are several approaches to handling this problem. The simplest and most generally applicable requires calculating meaningful summaries of the repeated observations for each plot, then analysing these using standard methods. This is an example of the strategy of 'moving everything to a single level'.

Activities

Lecture

There will be an introductory lecture covering common complexities in agroforestry trials then focusing on analysis of repeated measures. Decide whether to dwell on the mechanics of doing the calculations (not covered in the lecture notes). Some of the data manipulations may be confusing and distract participants from the important statistical ideas.

Practical exercise

Exercise

A practical discussing an example of repeated measures in space or in time.

Lecture

A brief lecture on handling multiple components.

Supporting documents

- R. Coe . 2002. *Complexity in agroforestry trials*. Lecture Note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre

Session 15

Complications in data

Session objectives

- To understand some of the common problems that become apparent when starting to analyse data from experiments.
- To take remedial action as to continue with an effective analysis.

Summary

In the lecture we list the types of problem, such as, missing values that commonly occur. We then look at the possible strategies that scientists can employ for dealing with the problems. They range from ignoring the problem to making a separate research study of the problem itself.

The practical examines ways of solving problems in particular cases, and these solutions are then discussed.

Activities

Lecture

An introductory lecture of 45 minutes illustrated by a demonstration, using the contents of the lecture note and also using appropriate examples selected by workshop resource persons.

Practical exercise

Exercise

There could be 5 or 6 groups, investigating different complications that occur in their own datasets or in those used in the workshop. Each group prepares a presentation on the complications that they have considered, the methods used for the solution, the ease with which they can be applied, and their success in resolving the problem. This exercise should take about 1.5 hours.

It is important that presenters are clear on the objectives of their presentation. It will be tempting for them to discuss their data analysis instead. While some discussion is needed, the objective is to describe methods of resolving complications.

Presentations, followed by discussion for 1 hour, (about 10 minutes for each group) should provide instruction to the whole group on practical ways of dealing with complications.

Review

This would probably be by one of the resource persons for 15 minutes and would summarize the key points of the session and particularly of the discussion.

Supporting documents

- R. D. Stern, R. Coe, E. Allen. 2002. *Complications in the data*. Lecture note. (Part 2). Data analysis of agroforestry experiments. ICRAF/World Agroforestry Centre
- Practical sheet

Session 16 Data analysis

Session objectives

- To analyse and present the results for participants own data.
- To consolidate the concepts introduced in the workshop by analysing these data with support and guidance from other workshop participants and resource persons.

Summary

The session is concerned solely with participants analysing their own data and presenting the results.

Activities

This session is one that should not necessarily wait until the end of the workshop. Scientists will have described the data they have brought, in the introduction to the workshop. In some course programmes there will be time within the workshop devoted specifically to the analysis of these problems. Alternatively, or additionally, a day towards the end of the workshop may be devoted to this area.

Some participants may also have made arrangements to continue their own work in the resource centre, following the formal close of the workshop. Others may wish to continue the analysis in their own institute, but seek a formal review of their work with a well-defined time scale to encourage the completion of the analysis.

Despite the many possibilities, it is important that each workshop defines the programme for these activities carefully. Realistic objectives must be set for what can be accomplished by each participant during the time permitted. There will normally be presentations by participants at the end of this session. These must be kept brief, for example 5 minutes per presentation, with 5 minutes discussion. Otherwise the presentations can become as ineffective as many conference sessions! Written information that can be circulated prior to the presentations, should be encouraged.

Supporting documents

- Lecture note.

Session 17 On your return

Session objectives

- To develop realistic personal action plans for future data analysis work, based on the outcomes of the training workshop.
- To develop a strategy for research support in data management and statistics at either the individual, group, institutional or national levels.

Summary

This session consists of presentations and discussions on what participants can do differently following this workshop. Changes could be on a personal level, or be for the institute or the national programme. They can affect data that will now be analysed and presented, a changed computing or software or training strategy, etc.

Scientists need to have impact. They need to promote technologies from their research that help farmers. The reason for this workshop is that the lack of skills to analyse data fully, is often quoted as a barrier to the completion of research projects. So the ideal is for scientists to feel that they may now be able to have more impact than before. However, increasing the skills of a few individuals who attend a course is unlikely to lead to much change, hence the need for this session.

Problems can be raised particularly if this workshop has suggested routes to a solution. Difficulties, where discussion is needed to suggest solutions should be raised as early as possible, so there is time for the necessary discussion. It is however, all too easy to raise problems rather than solutions. Participants should therefore be encouraged to concentrate on what they **can** do, rather than on what they can **not** do.

Activities

We would like to split the discussion into two main themes. The first follows the main theme in this workshop, the improvement of scientists' capabilities for data management and statistical analyses, so that they can complete their research activities more effectively on their return.

Most scientists have data that could usefully be analysed and written up. We would like to know of scientists' plans over the next 6 months and one year. We encourage scientists to be reasonable in their ambitions and hope that we can contact them after these periods to see how much progress they have made and whether there is any support we can give.

The analysis of data and the preparation of reports and other publications are necessary, but not sufficient. The research only becomes useful when someone reads the report or paper and takes some action to put the findings into practice. We would like the subsequent use of the results to be considered at the same time. We accept that some research is not intended to be applied directly by farmers; it is perhaps to assist our understanding of part of the process. However we do think that concern by scientists on how their results might be used is helpful in the writing of them, as often these are written in different ways for the different types of reader.

The second theme is that of the development of a strategy for research support, in particular for the areas of data management and analysis that have been discussed in this workshop. But it could also relate to support on design, and possibly to other related topics that participants feel could improve their effectiveness to conduct and complete their research.

However, to be able to make concrete proposals we would like to encourage scientists to limit the range of topics that are included in this discussion of strategy. For example, we accept there are problems of computer access, of training for higher degrees, of hierarchical structures in research institutes and so on. We cannot solve all these problems!

We do feel, however, that there are possibilities for improvements in the 'system' for statistical support, both in aspects that you can do yourselves, on your return and in ways that you can have access to support from others. It is these aspects we would like to explore briefly here. They relate to the ways computers are used, the software that is employed, who does the different tasks and so on. The starting point is your current strategy.

The following table might be used:

	Now		**Future**	
	Software	Strategy	Software	Strategy
Data entry				
Data management				
Data analysis				
Presentation of graphs				
Report writing				

How do you get support from others? This could be from colleagues, from statisticians, etc. They could be in your institute locally or centrally. They could be employed by your programme or by another organization.

What is your current situation and what would you like that you think is feasible?

On this strategy area, we anticipate that there might be an initial 15 minute introduction and then a 45 minute discussion at a point that is perhaps halfway through the workshop. It should not be too early, because it would be helpful if participants were able to relate the question of support to the topics in the workshop. If there is a day for Session 16, 'Own analyses', it could be within this day.

This would set the scene for a full discussion of the topics in this session. This would be towards the end of the workshop, though not too late, giving enough time for action if clear proposals are identified. We anticipate that the project report would summarize the key points from this discussion.

ISBN 92 9059 145 5

1

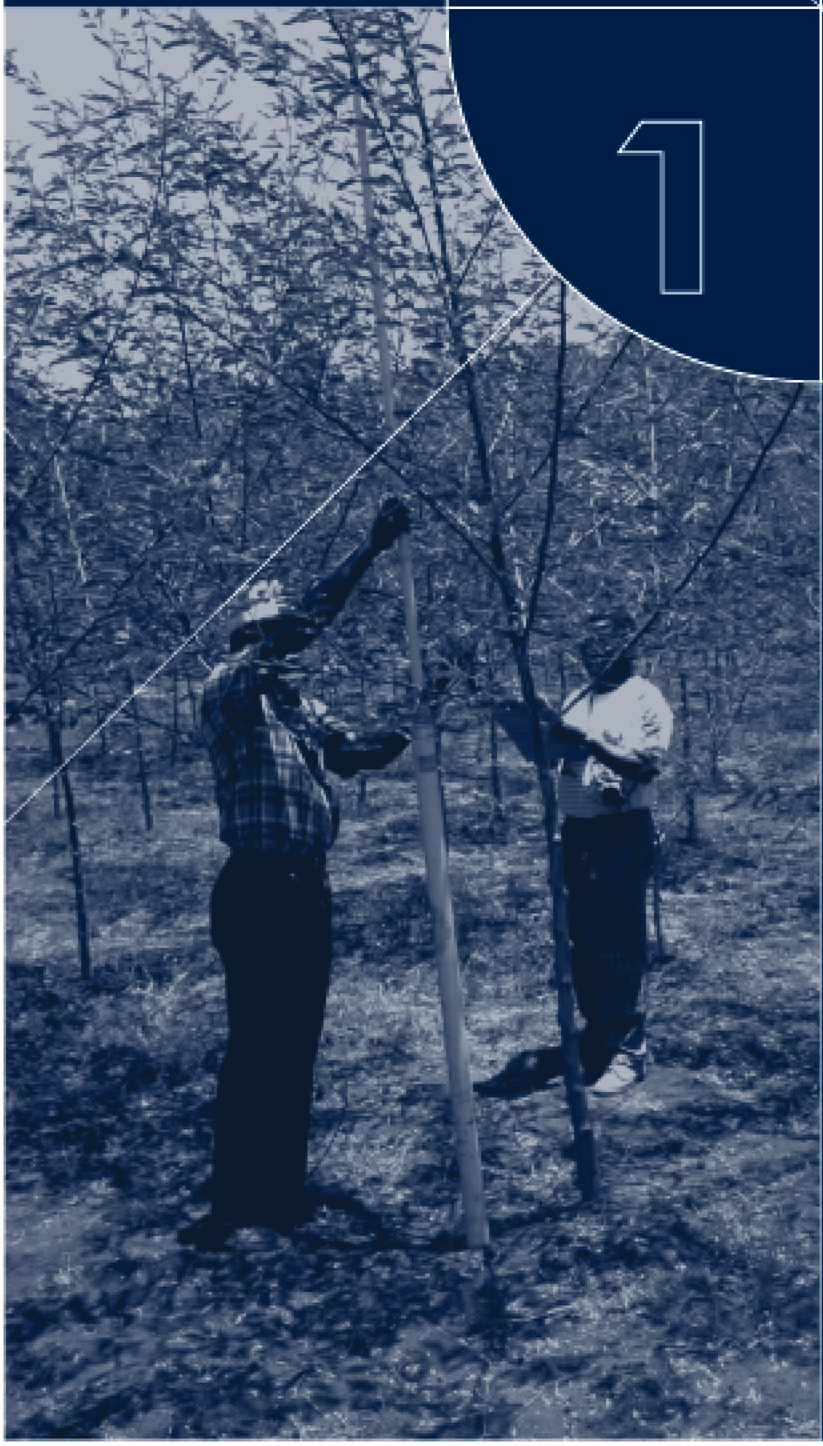

World Agroforestry Centre

TRANSFORMING LIVES AND LANDSCAPES