NAG C Library Chapter Introduction

g04 – Analysis of Variance

Contents

1 Scope of the Chapter

This chapter is concerned with methods for analysing the results of designed experiments. The range of experiments covered includes:

(1) single factor designs with equal sized blocks such as randomised complete block and balanced incomplete block designs,

(2) row and column designs such as Latin squares, and

(3) complete factorial designs.

Further designs may be analysed by combining the analyses provided by multiple calls to routines or by using general linear model routines provided in Chapter g02.

2 Background

2.1 Experimental Designs

An experimental design consists of a plan for allocating a set of controlled conditions, the treatments, to subsets of the experimental material, the plots or units. Two examples are:

(a) In an experiment to examine the effects of different diets on the growth of chickens, the chickens were kept in pens and a different diet was fed to the birds in each pen. In this example, the pens are the units and the different diets are the treatments.

(b) In an experiment to compare four materials for wear-loss, a sample from each of the materials is tested in a machine that simulates wear. The machine can take four samples at a time and a number of runs are made. In this experiment, the treatments are the materials and the units are the samples from the materials.

In designing an experiment, the following principles are important.

(i) Randomisation: Given the overall plan of the experiment, the final allocation of treatments to units is performed using a suitable random allocation. This avoids the possibility of a systematic bias in the allocation and gives a basis for the statistical analysis of the experiment.

(ii) Replication: Each treatment should be 'observed' more than once. So, in example (b), more than one sample from each material should be tested. Replication allows for an estimate of the variablity of the treatment effect to be measured.

(iii) Blocking: In many situations the experimental material will not be homogeneous and there may be some form of systematic variation in the experimental material. In order to reduce the effect of systematic variation, the material can be grouped into blocks so that units within a block are similar but there is variation between blocks. For example, in an animal experiment litters may be considered as blocks; in an industrial experiment it may be material from one production batch.

(iv) Factorial designs: If more than one type of treatment is under consideration, for example, the effect of changes in temperature and changes in pressure, a factorial design consists of looking at all combinations of temperature and pressure. The different types of treatment are known as factors and the different values of the factors that are considered in the experiment are known as levels. So, if three temperatures and four different pressures were being considered, then factor 1 (temperature) would have three levels and factor 2 (pressure) would have four levels and the design would be a 3×4 factorial, giving a total of 12 treatment combinations. This design has the advantage of being able to detect the interaction between factors, that is, the effect of the combination of factors.

The following are examples of standard experimental designs; in the descriptions, it is assumed that there are t treatments.

(i) Completely Randomised Design: There are no blocks and the treatments are allocated to units at random.

(ii) Randomised Complete Block Design: The experimental units are grouped into b blocks of t units and each treatment occurs once in each block. The treatments are allocated to units within blocks at random.

(iii) Latin Square Designs: The units can be represented as cells of a $t \times t$ square, classified by rows and columns. The t rows and t columns represent sources of variation in the experimental material. The design allocates the treatments to the units so that each treatment occurs once in each row and each column.

(iv) Balanced Incomplete Block Designs: The experimental units are grouped into b blocks of $k < t$ units. The treatments are allocated so that each treatment is replicated the same number of times and each treatment occurs in the same block with any other treatment the same number of times. The treatments are allocated to units within blocks at random.

(v) Complete Factorial Experiments: If there are t treatment combinations derived from the levels of all factors, then either there are no blocks, or the blocks are of size t units.

Other designs include: partially balanced incomplete block designs, split-plot designs, factorial designs with confounding, and fractional factorial designs. For further information on these designs, see Cochran and Cox (1957), Davis (1978) or John and Quenouille (1977).

2.2 Analysis of Variance

The analysis of a designed experiment usually consists of two stages. The first is the computation of the estimate of variance of the underlying random variation in the experiment along with tests for the overall effect of treatments. This results in an analysis of variance (ANOVA) table. The second stage is a more detailed examination of the effect of different treatments either by comparing the difference in treatment means with an appropriate standard error or by the use of orthogonal contrasts.

The analysis assumes a linear model such as

$$y_{ij} = \mu + \delta_i + \tau_l + e_{ij}$$

where y_{ij} is the observed value for unit j of block i, μ is the overall mean, δ_i is the effect of the ith block, τ_l is the effect of the lth treatment which has been applied to the unit, and e_{ij} is the random error term associated with this unit. The expected value of e_{ij} is zero and its variance is σ^2.

In the analysis of variance, the total variation, measured by the sum of squares of observations about the overall mean, is partitioned into the sum of squares due to blocks, the sum of squares due to treatments, and a residual or error sum of squares. This partition corresponds to the parameters β, τ and σ. In parallel to the partition of the sum of squares there is a partition of the degrees of freedom associated with the sums of squares. The total degrees of freedom is $n - 1$, where n is the number of observations. This is partitioned into $b - 1$ degrees of freedom for blocks, $t - 1$ degrees of freedom for treatments, and $n - t - b + 1$ degrees of freedom for the residual sum of squares. From these the mean squares can be computed as the sums of squares divided by their degrees of freedom. The residual mean square is an estimate of σ^2. An F-test for an overall effect of the treatments can be calculated as the ratio of the treatment mean square to the residual mean square.

For row and column designs the model is extended to give

$$y_{ij} = \mu + \rho_i + \gamma_j + \tau_l + e_{ij},$$

where ρ_i is the effect of the ith row and γ_j is the effect of the jth column. Usually the rows and columns are orthogonal. In the analysis of variance, the total variation is partitioned into rows, columns, treatments and residual.

In the case of factorial experiments, the treatment sum of squares and degrees of freedom may be partitioned into main effects for the factors, and interactions between factors. The main effect of a factor is the effect of the factor averaged over all other factors. The interaction between two factors is the additional effect of the combination of the two factors, over and above the additive effects of the two factors, averaged over all other factors. For a factorial experiment in blocks with two factors, A and B, in which the jth unit of the ith block received level l of factor A and level k of factor B the model is

$$y_{ij} = \mu + \delta_i + (\alpha_l + \beta_k + \alpha\beta_{lk}) + e_{ij},$$

where α_l is the main effect of level l of factor a, β_k is the main effect of level k of factor B, and $\alpha\beta_{lk}$ is the interaction between level l of A and level k of B. Higher-order interactions can be defined in a similar way.

Once the signifcant treatment effects have been uncovered they can be further investigated by comparing the differences between the means with the appropriate standard error. Some of the assumptions of the analysis can be checked by examining the residuals.

3 References

Cochran W G and Cox G M (1957) *Experimental Designs* Wiley

Davis O L (1978) *The Design and Analysis of Industrial Experiments* Longman

John J A (1987) *Cyclic Designs* Chapman and Hall

John J A and Quenouille M H (1977) *Experiments: Design and Analysis* Griffin

Searle S R (1971) *Linear Models* Wiley

4 Available Functions

This chapter contains functions that can handle a wide range of experimental designs

To compute the analysis of variance along with treatment means and their standard errors for any block design with equal sized blocks use nag_anova_random (g04bbc). This function will handle both complete block designs and balanced and partially balanced incomplete block designs. for row and column designs, such as Latin squares, use nag_anova_random (g04bbc).

Having computed an analysis of variance, if there is no structure to the treatments then simultaneous confidence intervals for the differences between pairs of treatments can be computed using nag_anova_confid_interval (g04dbc).

In the case where the treatments have a factorial structure, nag_anova_factorial (g04cac) computes the analysis of variance for a complete factorial design (optionally with blocks) along with treatment means and their standard errors.

Other designs can be analysed by combinations of calls to nag_anova_random (g04bbc), nag_anova_r-ow_col (g04bcc) and nag_anova_factorial (g04cac). The analysis from the functions can be combined by using the residuals computed from one function as the response variable for another. For example, if the treatments in a Latin Square had a factorial structure, the row and column effects could be estimated by nag_anova_random (g04bbc) and then the residuals passed to nag_anova_factorial (g04cac) to compute the analysis of variance for the treatments.

For experiments with missing values, these values can be estimated by using the Healy and Westmacott procedure, (see John and Quenouille (1977)). This procedure involves starting with initial estimates for the missing values and then making adjustments based on the residuals from the analysis. The improved estimates are then used in further iterations of the process.

For designs that cannot be analysed using the functions described above, a general linear model approach may be used (see Searle (1971)). This involves creating a set of dummy (0,1) variables that define the factors and any interactions using nag_dummy_vars (g04eac). The model containing these dummy variables can be fitted to the response variable y_i by nag_regsn_mult_linear (g02dac).