Chapter 10
Mixed Models for the Analysis of Optimization Algorithms

Marco Chiarandini and Yuri Goegebeur

Abstract We review linear statistical models for the analysis of computational experiments on optimization algorithms. The models offer the mathematical framework to separate the effects of algorithmic components and instance features included in the analysis. We regard test instances as drawn from a population and we focus our interest not on those single instances but on the whole population. Hence, instances are treated as a *random factor*. Overall these experimental designs lead to *mixed effects linear models*. We present both the theory to justify these models and a computational example in which we analyze and comment on several possible experimental designs. The example is a component-wise analysis of local search algorithms for the 2-edge-connectivity augmentation problem. We use standard statistical software to perform the analysis and report the R commands. Data sets and the analysis in SAS are available in an online compendium.

10.1 Introduction

Linear statistical models are well-developed mathematical tools for the separation of effects in the observed results of an experiment. Among them, there is the classical analysis of variance (ANOVA). scientific disciplines and also in the field of optimization. In operations research, application examples to test mathematical programming software go back to the late 1970s, see, e.g., Zanakis (1977), Lin and Rardin (1979), Coffin and Saltzman (2000); while in computer science and in test-

Marco Chiarandini
Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark e-mail: marco@imada.sdu.dk

Yuri Goegebeur
Department of Mathematics and Computer Science, University of Southern Denmark, Odense, Denmark, and Research Group Quantitative Psychology and Individual Differences, K.U.Leuven, Belgium e-mail: yuri.goegebeur@stat.sdu.dk

T. Bartz-Beielstein et al. (eds.), *Experimental Methods for the Analysis of Optimization Algorithms*, DOI 10.1007/978-3-642-02538-9_10, © Springer-Verlag Berlin Heidelberg 2010

ing heuristic and evolutionary computation methods their use can be traced back to the late 1990s. Prominent articles in this case are those by Barr et al. (1995), McGeoch (1996), Rardin and Uzsoy (2001), and Czarn et al. (2004). However, only a very small number of articles, relative to those published in these fields, report use of these statistical methods. This fact might be explained by two factors: the need for a background in statistics and experimental design techniques in order to correctly apply and fully understand the results provided; and the presence of underlying assumptions that make the researcher in computer science or operations research sceptical about the real applicability of these methods in the field of optimization. The aim of this chapter is to introduce the reader to the use of linear statistical models in the cases where they can be applied. We aim to present the basic theory behind the methods, their practical application by means of publically available software and the possible outcomes. We go perhaps to a deeper level of detail compared with previous publications in this field, hoping to facilitate future applications. However, we do not aim to remove completely the two barriers above: understanding of statistics and a careful investigation of applicability to each specific case are necessary preconditions. Knowledge of the material in the appendix of this book might be required to follow this chapter.

We emphasize that our intention is to present these tools as complementary to and not substitutes for the current practice of reporting numerical results on benchmark instances with appropriate tables. This practice is indeed helpful to guarantee comparability and verifiability of results. The methods in this chapter are however desirable for scientific experimental analysis, where the interest is in explaining the causes of success of a certain optimization approach rather than in mere comparative studies; see Hooker (1996) for a discussion on these guidelines.

To illustrate the application of the statistical tools we use a case example in which we study heuristic algorithms for a graph problem: finding the cheapest augmentation of arcs that make a network 2-edge-connected (Bang-Jensen et al. 2009). The heuristics are local search algorithms (Michiels et al. 2007) obtained by the combination of some specific components, which may be *qualitative*, like for the presence or not of an algorithmic step or *numerical*, like for parameters that assume real values. Our interest is in understanding the contribution of these components.

In statistical terms, these components are called *factors*. The interest is in the effects of the specific *levels* chosen for these factors. Hence, we say that the levels and consequently the factors are *fixed*. Moreover, when for two factors, every factor level of a factor appears with every factor level of another factor we say that the two factors are *crossed*. We restrict ourselves to analyze the effect of these factors on a univariate measure of performance, namely the quality of the solutions returned by the algorithm at termination. Multivariate analysis are however also possible by extensions of these methods; we refer to Johnson and Wichern (2007) for an overview of these.

Typically, the researcher takes a few instances for the problem at hand and collects the results of some runs of the algorithms on these instances. The instances are treated as *blocks* and all algorithms are run on each single instance. Results are therefore *grouped* per instance. The instances are chosen at random from a large set