

Preface

This book is about modeling psychophysical data with modern statistical methods using the software language **R** [146]. **R** is an interactive environment, and the novice user can just enter commands one by one at the command line prompt. It is also a powerful programming language in its own right. **R** is easily interfaced with other languages, such as C, Fortran, C++, Java, Perl, and Python. Another reason we recommend using **R** is that it permits access to a growing repository of new and powerful statistical methods, many of which are potentially useful to the psychophysical researcher. In the last chapter, for example, we describe methods that allow the researcher to model individual differences in observers, using mixed-effects models.

Many packages exist to supplement the functions of **R**. Over 3,000 are currently available on the Comprehensive R Archive Network (or CRAN) at <http://www.r-project.org/>. Further reasons for using **R** are that it is free and is supported by a large and enthusiastic online community.

Psychophysics begins with a question, posed to a *psychophysical observer*, human, animal, or neural.¹ The question may be as simple as judging which of two lights is brighter. The experimenter selects the question, constrains the observer's possible responses, and manipulates stimuli. The observer's response is effectively stochastic, and the experimenter's proximate goal is to model the observer's response and the mechanisms underlying that response. The final goal, of course, is to learn how neural systems process sensory information and provide access to the world around us.

The term “psychophysics” is due to Fechner [59], and there are several excellent introductions to psychophysical methods [91, 157–159, 164], the history of psychophysics [17, 115], and psychophysical theory [57, 151]. We note that psychophysics has no connection with psychometric theory. For an excellent introduction to psychometric theory using **R** we recommend Boeck and Wilson [16].

¹We can treat any device that categorizes stimuli as a psychophysical observer. See [22, 138, 174] for examples of psychophysical models of neural response.

This book is *not* an introduction to psychophysical methods, history, or theory. It is written for those who have a good idea what psychophysics entails and who regularly model psychophysical data or would like to do so—typically established or aspiring psychologists and neuroscientists studying perception. Some of the material is pertinent to researchers in memory using signal detection theory.

The focus of the book is on modern statistical methods readily available in **R** and how they can be used to analyze psychophysical data. Throughout the book, we illustrate how traditional and recently developed models of psychophysical performance can be viewed as special cases of statistical models of great power and generality, notably the generalized linear model (GLM). In Chap. 1, we use **R** to explore and model a celebrated set of psychophysical data due to Hecht, Shlaer, and Pirenne [79]. When possible, as in this chapter, we use a running example within the chapter so that readers can appreciate how different methods might be applied to the same data set. Thus, packages and data sets loaded earlier in the chapter as well as variables will be assumed to be defined from the point in the chapter in which they are invoked, but not across the chapters.

The statistical methods employed in most psychophysical laboratories were developed for the most part over the past 50 years and until recently, commonly available statistical packages were of limited use in fitting psychophysical data. That has now changed and much of what the psychophysicist wishes to learn from data and experiment can be learned using off-the-shelf statistical packages developed in the last few decades. Moreover, the statistical methods and packages now available can serve to inspire the psychophysicist to try more sophisticated analyses and consider better experimental designs. If, for example, the experimenter has run an experiment with a factorial design, then he might like to do the nonlinear equivalent of an analysis of variance (ANOVA) and determine whether variation of any of the factors led to changes in observers' performance and whether variations in the factors interact. We illustrate how to carry out ANOVA-like analyses where the data are appropriate for fitting by signal detection models or psychometric functions.

The book should be accessible to upperlevel undergraduate and graduate students and could be used as a textbook for a one-semester course. It presupposes only some knowledge of elementary statistics, some facility with vector and matrix notation, and some familiarity with programming. We do not assume that the reader is familiar with **R** or is particularly fluent in any other programming language. Almost all of the examples are drawn from studies of human vision only because that is what we know best. The methods presented are readily adaptable to model data in other disciplines ranging from memory to experimental economics, and in testing this book with students we have found that they could readily adapt the code we include as examples to their own data without any extensive programming.

Chapter 2 describes the principal statistical methods we will use: the linear model (LM), the generalized linear model (GLM)² the (generalized) additive model

²The generalized linear model is *not* the same as the general linear model commonly used in cognitive neuroscience. The near-coincidence in terminology is unfortunate. We will not make use of the general linear model at any point.

(GAM), and a brief discussion of methods for fitting nonlinear models in general. The reader familiar with LM in the form of multiple regression or analysis of variance can become familiar with how models are specified and analyzed in R. The remainder of this chapter introduces GLM and GAM by analogy with LM. The treatment of GLM and GAM in this book is self-contained. We provide references to texts on LM, GLM, and GAM using R for readers who want to pursue these topics further. If already familiar with R and its approach to modeling, the reader may choose to skip this chapter on an initial reading and to start instead with one of the later chapters describing particular psychophysical methods. The reader can refer back to Chap. 2 as needed.

We also include two appendices that summarize a few useful terms and results from mathematical statistics (Appendix B) and a brief description of data structures in R as well as basic information about input and output of data in R (Appendix A). We reference particular sections of these appendices in the main text, and they need not be read through if one is already familiar with their content.

The remaining chapters each focus on traditional and not so traditional methods of fitting data to psychophysical models using modern statistical methods. With the exception of Chap. 5 which builds on Chap. 4, the chapters are independent of one another and may be read in any order. We recommend that, wherever you start, you type in and execute the examples in the text as you read them. All of the code is available in scripts in a companion R package **MPDiR** available on CRAN that also contains the data sets not available elsewhere. There are exercises at the end of each chapter that invite the reader to modify the example code or to explore further the R commands used. If you are motivated to go further, you can apply the example code to your own data. Even without extensive experience in programming or programming in R, you should be able to adapt your own data to our examples. All of the data analyses and graphs in this book were generated using R, and we encourage the reader to execute the examples in the book as they occur. In our experience, using R is the best way to learn R.

In Chap. 3 we describe how to use GLM to fit data from signal detection experiments. Signal detection theory (SDT) was introduced in the 1950s to address the problem of observer bias. In a simple Yes–No task, a psychometric function summarizes the probability of a “Yes” response. However, this probability confounds the observer’s ability to detect the presence or absence of a signal with any preference the observer may have to respond “Yes.” Even when the signal intensity is zero (no signal is presented), the observer may still respond “Yes.” In a signal detection experiment, in simplest form, there is only one intensity of signal that is presented or not presented, at random, on each trial. On each trial, the observer is obliged to judge whether a signal was present. The two measures of interest are the rate of responding “Present” when the signal is present (referred to as the Hit rate) and the rate of responding “Present” when the signal is absent, the False Alarm rate. These two separate response measures are combined to allow separate estimates of a measure of the observer’s sensitivity to the signal and the observer’s bias (tendency to say “Present” independent of the presence or absence of the signal).

In Chap. 4 we illustrate how to fit psychometric functions to data using GLM. The psychometric function describes the relation between the probability of making a specific classification such as “Present” and stimulus intensity. Its earliest formulation dates from at least the beginning of the twentieth century [176]. One long-established use of psychophysical methods is to measure the precision of the human sensory apparatus in detection or discrimination. The researcher’s goal nowadays is often to use such data to test models of sensory, perceptual, or cognitive function, and one of the most venerable is the *psychometric function* relating the observer’s ability to judge a stimulus to stimulus intensity. There are many variant forms of psychometric function each associated with a particular experimental method. For now we assume that the observer responds “Yes” or “No,” judging whether the stimulus was present.

In Chap. 5 we extend the analysis of the psychometric function to more complex situations, for example, when there are multiple choices. We, also, consider here methods for obtaining standard errors and confidence intervals on the parameters estimated for the fitted functions. Finally, we demonstrate some nonparametric approaches to estimating psychometric functions.

In Chap. 6 we describe classification image methods. These methods are a natural extension of signal detection theory and have proven to be of great value in modeling psychophysical data. They are also easy to work with using GLM and GAM.

In the next two chapters, we describe two relatively new methods for representing perception of stimulus intensities well above threshold. That is, these methods are primarily useful in measuring *appearance*. The first method, maximum likelihood difference scaling, is described in Chap. 7, and the second, maximum likelihood conjoint measurement, in Chap. 8. These methods and corresponding psychophysical tasks complement traditional methods. We will illustrate their use.

In many statistical analyses the experimenter wishes to estimate how variable the outcomes of psychophysical measurements are across observers, generalized from results measured for a specific set of observers in an experiment to the population from which they were drawn. In the last chapter, Chap. 9, we focus on how to use newly available R packages to fit mixed-effects models of psychophysical performance. These models take us to the edge of research in statistics and hint at what is to come.

The more advanced programmer will want to know that R is an object-oriented language. Most of the functions that carry out statistical analyses return not just vectors or matrices, but objects summarizing the analyses carried out. Generic functions that print, plot, summarize, or do model comparisons accept these objects and choose the appropriate methods to do what is appropriate. In general, we simply use these functions as needed to illustrate how intuitive they are.

R is a language that is in constant development, and although the core functionality evolves only slowly over time, usually with an eye toward backward compatibility, this is less so for the code in packages. In time, some of the code in this book will be affected by such changes and may not run. We have opted to choose package versions that seem stable but in the event that future changes break some of the code, we will make available supplemental materials on a web

page associated with the book that can be accessed at <http://extras.springer.com/>. The current project home page for the book can be found at <http://mpdir.r-forge.r-project.org/>.

We thank Angela Brown, Frédéric Devinck, Roland Fleming, Jonathan Pierce, Hao Sun, François Vital-Durand, and Li Zhaoping for generously providing data. We thank Rune H. B. Christensen for extended discussions of ordinal models and Aries Arditi, Patrick Bennett, Reinhold Kleigl, Richard Murray, and, above all, Michael Kubovy for reading and critiquing several chapters. We thank our contacts at Springer, the editors John Kimmel and Marc Strauss, and Hannah Bracken. We also thank all of the R Core and R Help List contributors who responded to many questions along the way. Any remaining errors are, of course, our own.

We especially thank Geneviève Inchauspé and Maria Felicita Dal Martello for infinite patience with our preoccupation with R and this book.

Lyon, France
New York
Padova, Italy

Kenneth Knoblauch
Laurence T. Maloney