


COMPUTATIONAL MODELING IN COGNITION

P R I N C I P L E S

A N D

P R A C T I C E



STEPHAN LEWANDOWSKY
SIMON FARRELL





30809591

COMPUTATIONAL MODELING IN COGNITION

PRINCIPLES
AND
PRACTICE

STEPHAN LEWANDOWSKY

University of Western Australia

SIMON FARRELL

University of Bristol



 **SAGE**

Los Angeles | London | New Delhi
Singapore | Washington DC

Copyright © 2011 by SAGE Publications, Inc.

All rights reserved. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher.

For information:



SAGE Publications, Inc.
2455 Teller Road
Thousand Oaks, California 91320
E-mail: order@sagepub.com

SAGE Publications India Pvt. Ltd.
B 1/I 1 Mohan Cooperative
Industrial Area
Mathura Road, New Delhi 110 044
India

SAGE Publications Ltd.
1 Oliver's Yard
55 City Road, London, EC1Y 1SP
United Kingdom

SAGE Publications Asia-Pacific
Pte. Ltd.
33 Pekin Street #02-01
Far East Square
Singapore 048763

Printed in the United States of America

Library of Congress Cataloging-in-Publication Data

Lewandowsky, Stephan.

Computational modeling in cognition : principles and practice / Stephan Lewandowsky and Simon Farrell.

p. cm.

Includes bibliographical references and index.

ISBN 978-1-4129-7076-1 (pbk.)

1. Cognition—Mathematical models. I. Farrell, Simon, 1976- II. Title.

BF311.L467 2011

153.01'5118—dc22

2010029246

This book is printed on acid-free paper.

10 11 12 13 14 10 9 8 7 6 5 4 3 2 1

<i>Acquisitions Editor:</i>	Christine Cardone
<i>Editorial Assistant:</i>	Sarita Sarak
<i>Production Editor:</i>	Brittany Bauhaus
<i>Permissions Editor:</i>	Karen Ehrmann
<i>Copy Editor:</i>	Gillian Dickens
<i>Typesetter:</i>	C&M Digital (P) Ltd.
<i>Proofreader:</i>	Sally Jaskold
<i>Cover Designer:</i>	Gail Buschman
<i>Marketing Manager:</i>	Dory Schrader

Preface

The rapid progress in cognitive science during the past decade is intimately linked to three exciting and particularly active areas of research: computational and quantitative modeling of cognition, advances in the neurosciences, and the emphasis on Bayesian techniques as a tool to describe human behavior and to analyze data. Each of those areas is sufficiently broad to fill (at least one) textbook. This volume therefore focuses exclusively on cognitive modeling: We do not consider current research in the neurosciences or Bayesian techniques for data analysis and modeling because we cannot do justice to those additional topics in a single volume. Instead, this book is best considered an introductory stepping stone for further reading: Many of the issues discussed in this book constitute basic knowledge for the Bayesian data analyst and modeler.

What do you need to read this book? We have aimed this volume at an audience with only a limited background in mathematics. For example, we expect you to know the difference between a scalar, a vector, and a matrix, but we do not expect you to be an expert in matrix algebra.

In addition, we rely throughout on MATLAB to illustrate the core concepts with programming examples. If you want to follow those examples—and we strongly recommend that you do—then you will need access to MATLAB, and you need some prior knowledge of how to program in that language. If you have no background at all in programming, then you need to acquire some basic skills before you can tackle this book. Computational modeling, after all, involves computing.

Our intention was to write a book that would allow a junior Ph.D. student or a researcher without background in modeling to begin to acquire the skills necessary for cognitive modeling. Similarly, this book is suitable for use in an advanced undergraduate course on modeling if accompanied by suitable tuition. We also expect that many experts may find it a useful reference guide; however, to do justice to our primary intended target audience, there are many issues that we—reluctantly—had to omit from this volume. Accordingly, we have not covered such topics as Bayesian parameter estimation or multilevel modelling. Applying the

Pareto principle, we believe that 80% of our readership will be interested in 20% of the field—and so we focused on making those 20% particularly accessible.

There are several ways in which this book can be used and perused. The order of our chapters is dictated by logic, and we thus present basic modeling tools before turning to model selection and so on. However, the chapters can be read in a number of different orders, depending on one's background and intentions.

For example, readers with very little background in modeling may wish to begin by reading Chapters 1, 2, and 3, followed by the first part of Chapter 7 and all of Chapter 8. Then, you may wish to go back and read Chapters 4, 5, and 6. In contrast, if this book is used for formal tuition in a course, then we suggest that the chapters be assigned in the order in which they are presented; in our experience, this order follows the most logical progression in which the knowledge is presented.

This project would not have been possible without support and assistance from many sources. We are particularly grateful to Klaus Oberauer, John Dunn, E.-J. Wagenmakers, Jeff Rouder, Lael Schooler, and Roger Ratcliff for comments and clarifications on parts of this book.

—*Stephan Lewandowsky*

—*Simon Farrell*

Perth and Bristol,

April 2010

Contents

Preface	ix
1. Introduction	1
1.1 Models and Theories in Science	1
1.2 Why Quantitative Modeling?	3
1.3 Quantitative Modeling in Cognition	5
1.3.1 Models and Data	5
1.3.2 From Ideas to Models	8
1.3.3 Summary	10
1.4 The Ideas Underlying Modeling and Its Distinct Applications	10
1.4.1 Elements of Models	10
1.4.2 Data Description	11
1.4.3 Process Characterization	16
1.4.4 Process Explanation	19
1.4.5 Classes of Models	25
1.5 What Can We Expect From Models?	25
1.5.1 Classification of Phenomena	25
1.5.2 Emergence of Understanding	26
1.5.3 Exploration of Implications	27
1.6 Potential Problems	28
1.6.1 Scope and Testability	28
1.6.2 Identification and Truth	32
2. From Words to Models: Building a Toolkit	35
2.1 Working Memory	35
2.1.1 The Phonological Loop	36
2.1.2 Theoretical Accounts of the Word Length Effect	37
2.2 The Phonological Loop: 144 Models of Working Memory	39
2.2.1 Decay Process	40
2.2.2 Recall Process	41
2.2.3 Rehearsal	41

2.3	Building a Simulation	42
2.3.1	MATLAB	42
2.3.2	The Target Methodology	43
2.3.3	Setting Up the Simulation	44
2.3.4	Some Unanticipated Results	49
2.3.5	Variability in Decay	51
2.4	What Can We Learn From These Simulations?	53
2.4.1	The Phonological Loop Model Instantiated	53
2.4.2	Testability and Falsifiability	54
2.4.3	Revisiting Model Design: Foundational Decisions	55
2.5	The Basic Toolkit	56
2.5.1	Parameters	57
2.5.2	Discrepancy Function	61
2.5.3	Parameter Estimation and Model-Fitting Techniques	63
2.6	Models and Data: Sufficiency and Explanation	64
2.6.1	Sufficiency of a Model	65
2.6.2	Verisimilitude Versus Truth	67
2.6.3	The Nature of Explanations	68
3.	Basic Parameter Estimation Techniques	71
3.1	Fitting Models to Data: Parameter Estimation	71
3.1.1	Visualizing Modeling	72
3.1.2	An Example	74
3.1.3	Inside the Box: Parameter Estimation Techniques	82
3.2	Considering the Data: What Level of Analysis?	96
3.2.1	Implications of Averaging	97
3.2.2	Fitting Individual Participants	100
3.2.3	Fitting Subgroups of Data	103
3.2.4	Fitting Aggregate Data	103
3.2.5	Having Your Cake and Eating It: Multilevel Modeling	104
3.2.6	Recommendations	105
4.	Maximum Likelihood Estimation	109
4.1	Basics of Probabilities	110
4.1.1	Defining Probability	110
4.1.2	Properties of Probabilities	111
4.1.3	Probability Functions	112
4.2	What Is a Likelihood?	117
4.2.1	Inverse Probability and Likelihood	124
4.3	Defining a Probability Function	125
4.3.1	Probability Functions Specified by the Psychological Model	126
4.3.2	Probability Functions via Data Models	127

4.3.3 Two Types of Probability Functions	132
4.3.4 Extending the Data Model	133
4.3.5 Extension to Multiple Data Points and Multiple Parameters	136
4.4 Finding the Maximum Likelihood	138
4.5 Maximum Likelihood Estimation for Multiple Participants	142
4.5.1 Estimation for Individual Participants	143
4.5.2 Estimating a Single Set of Parameters	143
4.6 Properties of Maximum Likelihood Estimators	145
5. Parameter Uncertainty and Model Comparison	149
5.1 Error on Maximum Likelihood Estimates	150
5.1.1 Standard Errors Across Participants	150
5.1.2 Curvature of the Likelihood Surface	153
5.1.3 Bootstrapping	163
5.1.4 What Do Confidence Limits Tell Us?	170
5.2 Introduction to Model Selection	170
5.3 The Likelihood Ratio Test	177
5.4 Information Criteria and Model Comparison	179
5.4.1 Kullback-Leibler Distance and Akaike's Information Criterion	179
5.4.2 Bayesian Information Criterion	183
5.4.3 Model Comparison With the AIC and BIC	184
5.4.4 An Example of Model Comparison Using AIC and BIC	189
5.4.5 Choosing Between AIC and BIC	191
5.5 Conclusion	192
6. Not Everything That Fits Is Gold: Interpreting the Modeling	197
6.1 Psychological Data and the Very Bad Good Fit	198
6.1.1 Overfitting	198
6.1.2 Generalizability Versus Goodness of Fit	200
6.2 Parameter Identifiability and Model Testability	204
6.2.1 Identifiability	205
6.2.2 Testability	211
6.2.3 Putting It All Together: Deciding the Identifiability and Testability of a Model	212
6.2.4 Models That Are Identifiable but Not Testable	213
6.3 Drawing Lessons and Conclusions From Modeling	219
6.3.1 Explorations of a Model: Effects of Parameters	220
6.3.2 Demonstrations of Sufficiency	222
6.3.3 Demonstrations of Necessity	226
6.3.4 Summary	232

7. Drawing It All Together: Two Examples	235
7.1 WITNESS: Simulating Eyewitness Identification	235
7.1.1 WITNESS: Architecture	236
7.1.2 WITNESS and Verbal Overshadowing	238
7.1.3 WITNESS in MATLAB	240
7.1.4 WITNESS Simulation Results	252
7.2 Exemplar Versus Boundary Models: Choosing Between Candidates	257
7.2.1 Models of Categorization	258
7.2.2 A Probabilistic Feedback Experiment	263
7.2.3 MATLAB Code for ML Parameter Estimation for GCM, DEM, and GRT	265
7.2.4 Fitting the Models to Data	269
7.2.5 What Have We Learned About Categorization?	279
7.3 Conclusion	281
8. Modeling in a Broader Context	283
8.1 Bayesian Theories of Cognition	284
8.2 Neural Networks	288
8.2.1 Basic Architecture and Operation	288
8.2.2 Hebbian Models	290
8.2.3 Backpropagation Models	293
8.2.4 Summary	297
8.3 Neuroscientific Modeling	298
8.3.1 The Allure of Neuroscience	299
8.3.2 The True Promise of Neuroscience	303
8.4 Cognitive Architectures	306
8.4.1 Cognitive Architectures: Convergence to a Standard	307
8.4.2 Cognitive Architectures: Successes, Problems, and Solutions	308
8.4.3 Marrying Architectures and Neuroscience	310
8.4.4 Architectures and Models	311
8.5 Conclusion	313
8.5.1 Memory	314
8.5.2 Language	314
8.5.3 Perception and Action	315
8.5.4 Choice and Decision Making	315
8.5.5 Identification and Categorization	316
References	319
Author Index	345
Subject Index	353
About the Authors	359

Introduction

1.1 Models and Theories in Science

Cognitive scientists seek to understand how the mind works. That is, we want to *describe* and *predict* people's behavior, and we ultimately wish to *explain* it, in the same way that physicists predict the motion of an apple that is dislodged from its tree (and can accurately describe its downward path) and explain its trajectory (by appealing to gravity). For example, if you forget someone's name when you are distracted seconds after being introduced to her, we would like to know what cognitive process is responsible for this failure. Was it lack of attention? Forgetting over time? Can we know ahead of time whether or not you will remember that person's name?

The central thesis of this book is that to answer questions such as these, cognitive scientists must rely on quantitative mathematical models, just like physicists who research gravity. We suggest that to expand our knowledge of the human mind, consideration of the data and verbal theorizing are insufficient on their own.

This thesis is best illustrated by considering something that is (just a little) simpler and more readily understood than the mind. Have a look at the data shown in Figure 1.1, which represent the position of planets in the night sky over time.

How might one describe this peculiar pattern of motion? How would you explain it? The strange loops in the otherwise consistently curvilinear paths describe the famous "retrograde motion" of the planets—that is, their propensity to suddenly reverse direction (viewed against the fixed background of stars) for some time before resuming their initial path. What explains retrograde motion? It took more than a thousand years for a satisfactory answer to that question to become




Figure 1.1 An example of data that defy easy description and explanation without a quantitative model.

available, when Copernicus replaced the geocentric Ptolemaic system with a heliocentric model: Today, we know that retrograde motion arises from the fact that the planets travel at different speeds along their orbits; hence, as Earth “overtakes” Mars, for example, the red planet will appear to reverse direction as it falls behind the speeding Earth.

This example permits several conclusions that will be relevant throughout the remainder of this book. First, the pattern of data shown in Figure 1.1 defies description and explanation unless one has a *model* of the underlying process. It is only with the aid of a model that one can describe and explain planetary motion, even at a verbal level (readers who doubt this conclusion may wish to invite friends or colleagues to make sense of the data without knowing their source).

Second, any model that explains the data is itself unobservable. That is, although the Copernican model is readily communicated and represented (so readily, in fact, that we decided to omit the standard figure showing a set of concentric circles), it cannot be directly observed. Instead, the model is an abstract explanatory device that “exists” primarily in the minds of the people who use it to describe, predict, and explain the data.

Third, there nearly always are *several* possible models that can explain a given data set. This point is worth exploring in a bit more detail. The overwhelming




Figure 1.2 The geocentric model of the solar system developed by Ptolemy. It was the predominant model for some 1,300 years.

success of the heliocentric model often obscures the fact that, at the time of Copernicus's discovery, there existed a moderately successful alternative—namely, the geocentric model of Ptolemy shown in Figure 1.2. The model explained retrograde motion by postulating that while orbiting around the Earth, the planets also circle around a point along their orbit. On the additional, arguably somewhat inelegant, assumption that the Earth is slightly offset from the center of the planets' orbit, this model provides a reasonable account of the data, limiting the positional discrepancies between predicted and actual locations of, say, Mars to about 1° (Hoyle, 1974). Why, then, did the heliocentric model so rapidly and thoroughly replace the Ptolemaic system?¹

The answer to this question is quite fascinating and requires that we move toward a *quantitative* level of modeling.

1.2 Why Quantitative Modeling?

Conventional wisdom holds that the Copernican model replaced geocentric notions of the solar system because it provided a better account of the data.

4 Computational Modeling in Cognition

But what does “better” mean? Surely it means that the Copernican system predicted the motion of planets with less quantitative error—that is, less than the 1° error for Mars just mentioned—than its Ptolemaic counterpart? Intriguingly, this conventional wisdom is only partially correct: Yes, the Copernican model predicted the planets’ motion in latitude better than the Ptolemaic theory, but this difference was slight compared to the overall success of both models in predicting motion in longitude (Hoyle, 1974). What gave Copernicus the edge, then, was not “goodness of fit” alone² but also the intrinsic elegance and simplicity of his model—compare the Copernican account by a set of concentric circles with the complexity of Figure 1.2, which only describes the motion of a single planet.

There is an important lesson to be drawn from this fact: The choice among competing models—and remember, there are always several to choose from—inevitably involves an *intellectual judgment* in addition to quantitative examination. Of course, the quantitative performance of a model is at least as important as are its intellectual attributes. Copernicus would not be commemorated today had the predictions of his model been *inferior* to those of Ptolemy; it was only because the two competing models were on an essentially equal quantitative footing that other intellectual judgments, such as a preference for simplicity over complexity, came into play.

If the Ptolemaic and Copernican models were quantitatively comparable, why do we use them to illustrate our central thesis that a purely verbal level of explanation for natural phenomena is insufficient and that all sciences must seek explanations at a quantitative level? The answer is contained in the crucial modification to the heliocentric model offered by Johannes Kepler nearly a century later. Kepler replaced the circular orbits in the Copernican model by ellipses with differing eccentricities (or “egg-shapedness”) for the various planets. By this straightforward mathematical modification, Kepler achieved a virtually perfect fit of the heliocentric model with near-zero quantitative error. There no longer was any appreciable quantitative discrepancy between the model’s predictions and the observed paths of planets. Kepler’s model has remained in force essentially unchanged for more than four centuries.

The acceptance of Kepler’s model permits two related conclusions, one that is obvious and one that is equally important but perhaps less obvious. First, if two models are equally simple and elegant (or nearly so), the one that provides the better quantitative account will be preferred. Second, the predictions of the Copernican and Keplerian models cannot be differentiated by verbal interpretation alone. Both models explain retrograde motion by the fact that Earth “overtakes” some planets during its orbit, and the differentiating feature of the two models—whether orbits are presumed to be circular or elliptical—does not entail any differences in predictions that can be appreciated by purely verbal analysis.

That is, although one can talk about circles and ellipses (e.g., “one is round, the other one egg shaped”), those verbalizations cannot be turned into testable predictions: Remember, Kepler reduced the error for Mars from 1° to virtually zero, and we challenge you to achieve this by verbal means alone.

Let us summarize the points we have made so far:

1. Data never speak for themselves but require a model to be understood and to be explained.
2. Verbal theorizing alone ultimately cannot substitute for quantitative analysis.
3. There are always several alternative models that vie for explanation of data, and we must select among them.
4. Model selection rests on both quantitative evaluation and intellectual and scholarly judgment.

All of these points will be explored in the remainder of this book. We next turn our attention from the night sky to the inner workings of our mind, first by showing that the preceding conclusions apply in full force to cognitive scientists and then by considering an additional issue that is of particular concern to scholars of the human mind.

1.3 Quantitative Modeling in Cognition

1.3.1 Models and Data

Let's try this again: Have a look at the data in Figure 1.3. Does it remind you of planetary motion? Probably not, but it should be at least equally challenging to discern a meaningful pattern in this case as it was in the earlier example. Perhaps the pattern will become recognizable if we tell you about the experiment conducted by Nosofsky (1991) from which these data are taken. In that experiment, people were trained to classify a small set of cartoon faces into two arbitrary categories (we might call them the Campbells and the MacDonalds, and members of the two categories might differ on a set of facial features such as length of nose and eye separation).

On a subsequent transfer test, people were presented with a larger set of faces, including those used at training plus a set of new ones. For each face, people had to make two decisions: which category the face belonged to and the confidence of that decision (called “classification” in the figure, shown on the x -axis), and whether or not it had been shown during training (“recognition,” on




Figure 1.3 Observed recognition scores as a function of observed classification confidence for the same stimuli (each number identifies a unique stimulus). See text for details. Figure reprinted from Nosofsky, R. M. (1991). Tests of an exemplar mode for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 3–27. Published by the American Psychological Association; reprinted with permission.

the y-axis). Each data point in the figure, then, represents those two responses, averaged across participants, for a given face (identified by ID number, which can be safely ignored). The correlation between those two measures was found to be $r = .36$.

Before we move on, see if you can draw some conclusions from the pattern in Figure 1.3. Do you think that the two tasks have much to do with each other? Or would you think that classification and recognition are largely unrelated and that knowledge of one response would tell you very little about what response to expect on the other task? After all, if $r = .36$, then knowledge of one response reduces uncertainty about the other one by only 13%, leaving a full 87% unexplained, right?

Wrong. There is at least one quantitative cognitive model (called the GCM and described a little later), which can relate those two types of responses with considerable certainty. This is shown in Figure 1.4, which separates classification and recognition judgments into two separate panels, each showing the




Figure 1.4 Observed and predicted classification (left panel) and recognition (right panel). Predictions are provided by the GCM; see text for details. Perfect prediction is represented by the diagonal lines. Figure reprinted from Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 3–27. Published by the American Psychological Association; reprinted with permission.

relationship between observed responses (on the y-axis) and the predictions of the GCM (x-axis). To clarify, each point in Figure 1.3 is shown twice in Figure 1.4—once in each panel and in each instance plotted as a function of the *predicted* response obtained from the model.

The precision of predictions in each panel is remarkable: If the model's predictions were absolutely 100% perfect, then all points would fall on the diagonal. They do not, but they come close (accounting for 96% and 91% of the variance in classification and recognition, respectively). The fact that these accurate predictions were provided by the same model tells us that classification and recognition can be understood and related to each other within a common psychological theory. Thus, notwithstanding the low correlation between the two measures, there is an underlying model that explains how both tasks are related and permits accurate prediction of one response from knowledge of the other. This model will be presented in detail later in this chapter (Section 1.4.4); for now, it suffices to acknowledge that the model relies on the comparison between each test stimulus and all previously encountered exemplars in memory.

The two figures enforce a compelling conclusion: “The initial scatterplot . . . revealed little relation between classification and recognition performance. At that limited level of analysis, one might have concluded that there was little in common between the fundamental processes of classification and recognition. Under

the guidance of the formal model, however, a unified account of these processes is achieved” (Nosofsky, 1991, p. 9). Exactly paralleling the developments in 16th-century astronomy, data in contemporary psychology are ultimately only fully interpretable with the aid of a quantitative model. We can thus reiterate our first two conclusions from above and confirm that they apply to cognitive psychology in full force—namely, that *data never speak for themselves but require a model to be understood and to be explained* and that *verbal theorizing alone cannot substitute for quantitative analysis*. But what about the remaining earlier conclusions concerning model selection?

Nosofsky’s (1991) modeling included a comparison between his favored exemplar model, whose predictions are shown in Figure 1.4, and an alternative “prototype” model. The details of the two models are not relevant here; it suffices to note that the prototype model compares a test stimulus to the *average* of all previously encountered exemplars, whereas the exemplar model performs the comparison one by one between the test stimulus and each exemplar and sums the result.³ Nosofsky found that the prototype model provided a less satisfactory account of the data, explaining only 92% and 87% of the classification and recognition variance, respectively, or about 5% less than the exemplar model. Hence, the earlier conclusions about model selection apply in this instance as well: There were several alternative models, and the choice between them was based on clear quantitative criteria.

1.3.2 From Ideas to Models

So far, we initiated our discussions with the data and we then . . . poof! . . . revealed a quantitative model that spectacularly turned an empirical mystery or mess into theoretical currency. Let us now invert this process and begin with an idea, that is, some psychological process that you think might be worthy of exploration and perhaps even empirical test. Needless to say, we expect you to convert this idea into a quantitative model. This raises at least two obvious questions: First, how would one do this? Second, does this process have implications concerning the role of modeling other than those we have already discussed? These questions are sufficiently complex to warrant their own chapter (Chapter 2), although we briefly survey the latter here.

Consider the simple and elegant notion of rehearsal, which is at the heart of much theorizing in cognition (e.g., A. D. Baddeley, 2003). We have all engaged in rehearsal, for example, when we try to retain a phone number long enough to enter it into our SIM cards. Several theorists believe that such subvocal—or sometimes overt—rehearsal can prevent the “decay” of verbal short-term memory traces, and introspection suggests that repeated recitation of a phone number is a good means to avoid forgetting. Perhaps because of the overwhelming intuitive appeal of the notion and its introspective reality, there have been few if any attempts