

How to be a Quantitative Ecologist

The 'A to R' of Green
Mathematics and Statistics

Jason Matthiopoulos

 WILEY



HOW TO BE A QUANTITATIVE ECOLOGIST

HOW TO BE A QUANTITATIVE ECOLOGIST

THE 'A TO R' OF GREEN MATHEMATICS AND STATISTICS

Jason Matthiopoulos

University of St Andrews, Scotland, UK



A John Wiley & Sons, Ltd., Publication

This edition first published 2011

© 2011 John Wiley & Sons, Ltd

Registered office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Matthiopoulos, Jason.

How to be a quantitative ecologist : the 'A to R' of green mathematics and statistics / Jason Matthiopoulos.

p. cm.

Summary: "The book will comprise two equal parts on mathematics and statistics with emphasis on quantitative skills" – Provided by publisher.

Includes bibliographical references and index.

ISBN 978-0-470-69978-2 (hardback) – ISBN 978-0-470-69979-9 (paper)

1. Ecology – Mathematics. 2. Ecology – Vocational guidance. 3. Ecology – Research.
4. Quantitative analysts. 5. Quantitative research. 6. Mathematics – Vocational guidance.

I. Title.

QH323.5.M387 2011

577.0285'5133 – dc22

2010051191

A catalogue record for this book is available from the British Library.

ISBN HB: 978-0-470-69978-2

ISBN PB: 978-0-470-69979-2

ePDF ISBN: 978-1-119-99158-8

eBook ISBN: 978-1-119-99159-5

ePub ISBN: 978-1-119-99172-4

Typeset in 9/11pt PalatinoLinotype by Laserwords Private Limited, Chennai, India

Στους αγαπημένους μου γονείς,
Έφη και Σπύρο


How to be a Quantitative Ecologist

The **A to R** of green
mathematics &
statistics

How I chose to write this book, and why you might choose to
read it xvii
Preface

0. How to start a meaningful relationship with your computer 1
Introduction to R

0.1 What is R?	1	0.8 Basic R usage	8
0.2 Why use R for this book?	2	0.9 Importing data from a spreadsheet	9
0.3 Computing with a scientific package like R	2	0.10 Storing data in data frames	10
0.4 Installing and interacting with R	3	0.11 Exporting data from R	12
0.5 Style conventions	5	0.12 Quitting R	12
0.6 Valuable R accessories	6	Further reading	13
0.7 Getting help	7	References	13



1. How to make mathematical statements	15
<i>Numbers, equations and functions</i>	
1.1 Qualitative and quantitative scales	16
• Habitat classifications	
1.2 Numbers	17
• Observations of spatial abundance	
1.3 Symbols	20
• Population size and carrying capacity	
1.4 Logical operations	21
1.5 Algebraic operations	21
• Size matters in male garter snakes	
1.6 Manipulating numbers	25
1.7 Manipulating units	26
1.8 Manipulating expressions	26
• Energy acquisition in voles	
1.9 Polynomials	29
• The law of mass action in epidemiology	
1.10 Equations	32
1.11 First order polynomial equations	33
• Population size and composition	
1.12 Proportionality and scaling: a special kind of first order polynomial equation	34
• Simple mark-recapture	
• Converting density to population size	
1.13 Second and higher order polynomial equations	36
• Estimating the number of infected animals from the rate of infection	
1.14 Systems of polynomial equations	38
• Deriving population structure from data on population size	
1.15 Inequalities	40
• Minimum energetic requirements in voles	
1.16 Coordinate systems	42
• Non-Cartesian map projections	
1.17 Complex numbers	44
1.18 Relations and functions	47
• Food webs	
• Mating systems in animals	
1.19 The graph of a function	51
• Two aspects of vole energetics	
1.20 First order polynomial functions	52
• Population stability in a time series	
• Population stability and population change	
• Visualising goodness-of-fit	
1.21 Higher order polynomial functions	55
1.22 The relationship between equations and functions	57
• Extent of an epidemic when the transmission rate exceeds a critical value	
1.23 Other useful functions	58
1.24 Inverse functions	60
1.25 Functions of more than one variable	62
• Two aspects of vole energetics	
Further reading	64
References	65

2. How to describe regular shapes and patterns

Geometry and trigonometry

2.1 Primitive elements	68	2.11 Trigonometric identities	88
		• A two-step random walk	
2.2 Axioms of Euclidean geometry	69	2.12 Inverses of trigonometric functions	89
• Suicidal lemmings, parsimony, evidence and proof		• Displacement during a random walk	
2.3 Propositions	70	2.13 Trigonometric equations	91
• Radio-tracking of terrestrial animals		• VHF tracking for terrestrial animals	
2.4 Distance between two points	76	2.14 Modifying the basic trigonometric graphs	93
• Spatial autocorrelation in ecological variables		• Nocturnal flowering in dry climates	
2.5 Areas and volumes	78	2.15 Superimposing trigonometric functions	96
• Hexagonal territories		• More realistic model of nocturnal flowering	
2.6 Measuring angles	80	2.16 Spectral analysis	99
• The bearing of a moving animal		• Dominant frequencies in density fluctuations of Norwegian lemming populations	
2.7 The trigonometric circle	82	• Spectral analysis of oceanographic covariates	
• The position of a seed following dispersal		2.17 Fractal geometry	102
2.8 Trigonometric functions	83	• Availability of coastal habitat	
2.9 Polar coordinates	85	• Fractal dimension of the Koch curve	
• Random walks		Further reading	105
2.10 Graphs of trigonometric functions	86	References	106

3. How to change things, one step at a time

Sequences, difference equations and logarithms

3.1 Sequences	108	3.4 Initial conditions and parameters	115
• Reproductive output in social wasps		3.5 Solutions of a difference equation	116
• Unrestricted population growth		3.6 Equilibrium solutions	117
3.2 Difference equations	111	• Harvesting an unconstrained population	
• More realistic models of population growth		• Visualising the equilibria	
3.3 Higher order difference equations	114		
• Delay-difference equations in a biennial plant			

3.7 Stable and unstable equilibria	119	3.10 Exponential function	130
• Parameter sensitivity and ineffective fishing quotas		• Modelling bacterial loads in continuous time	
• Stable and unstable equilibria in a density-dependent population		• A negative blue tit? Using exponential functions to constrain models	
3.8 Investigating stability	122	3.11 Logarithmic function	132
• Cobweb plot for an unconstrained, harvested population		• Log-transforming population time series	
• Conditions for stability under unrestricted growth		3.12 Logarithmic equations	135
3.9 Chaos	126	Further reading	136
• Chaos in a model with density dependence		References	136

4. How to change things, continuously 137

Derivatives and their applications

4.1 Average rate of change	138	• Diurnal rate of change in the attendance of insect pollinators	
• Seasonal tree growth			
• Tree growth			
4.2 Instantaneous rate of change	141	4.8 Higher order derivatives	158
		• Spatial gradients	
4.3 Limits	142	4.9 Derivatives of functions of many variables	160
• Methane concentration around termite mounds		• The slope of the sea-floor	
4.4 The derivative of a function	145	4.10 Optimisation	163
• Plotting change in tree biomass		• Maximum rate of disease transmission	
• Linear tree growth		• The marginal value theorem	
4.5 Differentiating polynomials	148	4.11 Local stability for difference equations	168
• Spatial gradients		• Unconstrained population growth	
		• Density dependence and proportional harvesting	
4.6 Differentiating other functions	152	4.12 Series expansions	171
• Consumption rates of specialist predators		Further reading	174
4.7 The chain rule	155	References	174

5. How to work with accumulated change 177

Integrals and their applications

5.1 Antiderivatives	178	<ul style="list-style-type: none"> • Total number of arrivals and departures at migratory stop-over 	
<ul style="list-style-type: none"> • Invasion fronts • Diving in seals 			
5.2 Indefinite integrals	181		
<ul style="list-style-type: none"> • Allometry 			
5.3 Three analytical methods of integration	182		
<ul style="list-style-type: none"> • Stopping invasion fronts 			
5.4 Summation	187		
<ul style="list-style-type: none"> • Metapopulations 			
5.5 Area under a curve	190		
<ul style="list-style-type: none"> • Swimming speed in seals 			
5.6 Definite integrals	193		
<ul style="list-style-type: none"> • Swimming speed in seals 			
5.7 Some properties of definite integrals	195		
<ul style="list-style-type: none"> • Total reproductive output in social wasps • Net change in number of birds at migratory stop-over 			
		5.8 Improper integrals	198
		<ul style="list-style-type: none"> • Failing to stop invasion fronts 	
		5.9 Differential equations	202
		<ul style="list-style-type: none"> • A differential equation for a plant invasion front 	
		5.10 Solving differential equations	203
		<ul style="list-style-type: none"> • Exponential population growth in continuous time • Constrained growth in continuous time 	
		5.11 Stability analysis for differential equations	209
		<ul style="list-style-type: none"> • Constrained growth in continuous time • The Levins model for metapopulations 	
		Further reading	212
		References	212

6. How to keep stuff organised in tables 213

Matrices and their applications

6.1 Matrices	214	6.6 Eigenvalues and eigenvectors	230
<ul style="list-style-type: none"> • Plant community composition • Inferring diet from fatty acid analysis 		<ul style="list-style-type: none"> • Growth in patchy populations • Metapopulation growth 	
6.2 Matrix operations	217	6.7 Leslie matrix models	233
<ul style="list-style-type: none"> • Movement in metapopulations 		<ul style="list-style-type: none"> • Stage-structured seal populations • Equilibrium of linear Leslie model • Stability in a linear Leslie model • Stable age structure in a linear Leslie model 	
6.3 Geometric interpretation of vectors and square matrices	221		
<ul style="list-style-type: none"> • Random walks as sequences of vectors 			
6.4 Solving systems of equations with matrices	223	6.8 Analysis of linear dynamical systems	237
<ul style="list-style-type: none"> • Plant community composition 		<ul style="list-style-type: none"> • A fragmented population in continuous time • Phase-space for a two-patch metapopulation • Stability analysis of a two-patch metapopulation 	
6.5 Markov chains	227		
<ul style="list-style-type: none"> • Redistribution between population patches 			

6.9 Analysis of nonlinear dynamical systems	243	Further reading	249
• The Lotka–Volterra, predator–prey model		References	249
• Stability analysis of the Lotka–Volterra model			

7. How to visualise and summarise data 251

Descriptive statistics

7.1 Overview of statistics	252	7.8 Skewness and kurtosis	266
7.2 Statistical variables	253	7.9 Graphical summaries	267
• Activity budgets in honey bees		7.10 Data sets with more than one variable	268
7.3 Populations and samples	255	7.11 Association between qualitative variables	268
• Production of gannet chicks		• Community recovery in abandoned fields	
7.4 Single-variable samples	255	7.12 Association between quantitative variables	270
7.5 Frequency distributions	256	• Height and root depth of tree ferns	
• Activity budgets in honey bees		7.13 Joint frequency distributions	273
• Activity budgets from different studies		• Mosaics of abandoned fields	
• Visualising activity budgets		• Joint distribution of tree height and root depth	
• Height of tree ferns		• Joint and marginal distributions of tree size	
• Gannets on Bass rock		Further reading	276
7.6 Measures of centrality	260	References	276
• Chick rearing in red grouse			
• Swimming speed in grey seals			
• Median of chicks reared by red grouse			
7.7 Measures of spread	263		
• Gannet foraging			

8. How to put a value on uncertainty 279

Probability

8.1 Random experiments and event spaces	280	8.4 Equally likely events	285
• Assumptions of random experiments		• Something is certain to occur (even if it is nothing)	
8.2 Events	281	• Undirected movement	
• Plant occurrence in survey quadrats		8.5 The union of events	286
• Overlapping events		• Seed germination on a gridded landscape	
• Mutually exclusive events		• Coexisting sparrows	
8.3 Frequentist probability	284	8.6 Conditional probability	288
• Fluctuating frequency of newborn male wildebeest		• Territoriality and survival in red grouse	

8.7 Independent events	291	<ul style="list-style-type: none"> • Null and alternative hypotheses for wildebeest sex ratio • Bayesian updating for wildebeest sex ratio
8.8 Total probability	292	
<ul style="list-style-type: none"> • Seed germination in a heterogeneous environment 		Further reading
		296
8.9 Bayesian probability	293	References
<ul style="list-style-type: none"> • Does the sex ratio at birth deviate from 1:1? 		297

9. How to identify different kinds of randomness 299

Probability distributions

9.1 Probability distributions	300	9.10 Number of occurrences in a unit of time or space: the Poisson distribution	317
<ul style="list-style-type: none"> • Probability distribution of nominal variables 		<ul style="list-style-type: none"> • Data on plant abundance 	
9.2 Discrete probability distributions	301	9.11 The gentle art of waiting: geometric, negative binomial, exponential and gamma distributions	319
<ul style="list-style-type: none"> • Beetle eggs per cluster: a count • PMF for egg cluster size • CDF for egg cluster size 		<ul style="list-style-type: none"> • Catastrophic events and species extinctions 	
9.3 Continuous probability distributions	304	9.12 Assigning probabilities to probabilities: the beta and Dirichlet distributions	322
<ul style="list-style-type: none"> • Are all fern heights equally likely? • CDF for tree fern heights • PDF for tree fern heights 			
9.4 Expectation	306	9.13 Perfect symmetry: the normal distribution	324
<ul style="list-style-type: none"> • Mean egg cluster size • Expected tree fern heights 		<ul style="list-style-type: none"> • Weight distribution in voles 	
9.5 Named distributions	309	9.14 Because it looks right: using probability distributions empirically	326
		<ul style="list-style-type: none"> • A beta prior for wildebeest sex at-birth 	
9.6 Equally likely events: the uniform distribution	310	9.15 Mixtures, outliers and the t-distribution	328
<ul style="list-style-type: none"> • River otter home ranges 		<ul style="list-style-type: none"> • Bimodal weight distribution in voles 	
9.7 Hit or miss: the Bernoulli distribution	312	9.16 Joint, conditional and marginal probability distributions	329
<ul style="list-style-type: none"> • Bernoulli births 		<ul style="list-style-type: none"> • Feeding site fidelity in kittiwakes 	
9.8 Count of occurrences in a given number of trials: the binomial distribution	313	9.17 The bivariate normal distribution	333
<ul style="list-style-type: none"> • Wildebeest reproductive histories • Wildebeest reproductive output 		<ul style="list-style-type: none"> • Height and root depth of tree ferns 	
9.9 Counting different types of occurrences: the multinomial distribution	316	9.18 Sums of random variables: the central limit theorem	335
<ul style="list-style-type: none"> • Metapopulation transitions 		<ul style="list-style-type: none"> • Food provisioning in starlings • Mixed-diet provisioning in starlings 	

9.19 Products of random variables: the log-normal distribution	337	9.21 Stochastic simulation	341
• Stochastic exponential growth		• Population viability analysis	
9.20 Modelling residuals: the chi-square distribution	340	Further reading	344
• Position of limpets relative to water		References	344

10. How to see the forest from the trees

345

Estimation and testing

10.1 Estimators and their properties	345	• Maximum likelihood estimator for wildebeest sex ratio	
• Sampling gannet egg weights			
10.2 Normal theory	347	10.10 Bayesian estimation	364
10.3 Estimating the population mean	348	• Bayesian estimator for wildebeest sex ratio	
10.4 Estimating the variance of a normal population	349	10.11 Link between maximum likelihood and Bayesian estimation	368
10.5 Confidence intervals	350	• Comparing ML and Bayes estimates of sex ratio at birth	
• Estimating the unknown parameters of the gannet egg weight distribution		10.12 Hypothesis testing: rationale	369
10.6 Inference by bootstrapping	354	10.13 Tests for the population mean	370
• Bootstrap inference for gannet age distribution		10.14 Tests comparing two different means	373
10.7 More general estimation methods	356	• Paired samples of gannet condition	
• Estimating parameters for animal movement		• How to rank observations	
10.8 Estimation by least squares	358	10.15 Hypotheses about qualitative data	376
• Weight distribution in voles of different ages		• Tree selection in Peruvian ants	
• Mean weight in a cohort of voles		10.16 Hypothesis testing debunked	379
10.9 Estimation by maximum likelihood	359	Further reading	379
• Maximum likelihood estimation of mean and variance of vole weight distribution		References	379

11. How to separate the signal from the noise**381***Statistical modelling*

11.1 Comparing the means of several populations	382	<ul style="list-style-type: none"> • Seeing the linear model as a special case of the GLM • Likelihood for a log-linear GLM • Modelling senescence in turtles • Estimating survival as a function of age
<ul style="list-style-type: none"> • Samples of gannet condition • Estimates of multiple population averages 		
11.2 Simple linear regression	386	11.8 Evaluation, diagnostics and model selection for GLMs
<ul style="list-style-type: none"> • Density dependence of gannet condition • Modelling density dependence of gannet condition 		407
11.3 Prediction	390	11.9 Modelling dispersion
<ul style="list-style-type: none"> • Predicting density-dependent changes in gannet condition 		409
11.4 How good is the best-fit line?	391	11.10 Fitting more complicated models to data: polynomials, interactions, nonlinear regression
<ul style="list-style-type: none"> • Diagnostics for the density dependence model 		410
11.5 Multiple linear regression	394	<ul style="list-style-type: none"> • Density dependence and polynomial terms • Predation, immigration and interaction terms
<ul style="list-style-type: none"> • Combined effects of density and the environment 		
11.6 Model selection	396	11.11 Letting the data suggest more complicated models: smoothing
<ul style="list-style-type: none"> • Should the beating of a butterfly's wings be used to explain gannet condition? • Model selection by adjusted r^2 and AIC • Collinearity in covariates of gannet condition 		415
11.7 Generalised linear models	400	<ul style="list-style-type: none"> • Distribution along a linear habitat • Other density estimation by kernel smoothing • Other density estimation by GAM
<ul style="list-style-type: none"> • Many response variables are constrained • Log-transforming fecundity data • Modelling count data 		
		11.12 Partitioning variation: mixed effects models
		419
		<ul style="list-style-type: none"> • Samples of gannet condition • Modelling individual and colony variation • Why not use colony as a factor?
		Further reading
		422
		References
		423

12. How to measure similarity**425***Multivariate methods*

12.1 The problem with multivariate data	426	12.3 Principal components analysis	428
<ul style="list-style-type: none"> • Characterising environmental similarity • Characterising patterns of occurrence 		<ul style="list-style-type: none"> • Collinearities between four environmental variables • PCA for four environmental variables 	
12.2 Ordination in general	427	12.4 Clustering in general	432
<ul style="list-style-type: none"> • Correlations represent redundancy 			

• Identifying functional groups in ecological communities		12.8 Logistic regression: two classes	436
• Clustering data frame for Antarctic species		• Characterising fern habitat	
12.5 Agglomerative hierarchical clustering	433	12.9 Logistic regression: many classes	438
• Dendrogram for Antarctic species		• Classification of whale vocalisations	
12.6 Nonhierarchical clustering: k means analysis	434	Further reading	440
12.7 Classification in general	435	References	440
 Appendix: Formulae			441
 R Index			453
 Index			457

How I chose to write this book, and why you might choose to read it

(Preface)

*'Many of our biology students are refugees from high-school mathematics'
John Ollason, thinker, cynic and the first quantitative ecologist I ever met.*

The evolution of most languages is driven by ease of use and the need for fast information exchange. In this sense, phrases like the 'language of mathematics' and 'computer programming language' are cruel euphemisms (arguably, even seasoned mathematicians find it harder to read through an unfamiliar equation than an unfamiliar piece of text). These languages are driven by the desire to eliminate ambiguity, at the expense of user-friendliness. Because of our excellent ability to perceive context in everyday situations, few of us feel the need to communicate unambiguously. However, scientific research happens outside the comfort zone of well-understood phenomena. For this reason, modern ecologists need to be trained in quantitative methods but find the process painful. Although there is also a great deal of pleasure in using mathematics and computers in science, it is sometimes hard to keep sight of it during the early period of training. Assuming you have decided that you need quantitative skills for your scientific career (a good call), there are five tricks that can make your learning experience less arduous:

- ❶ Begin from material that you already know well and work your way up to the harder stuff. Try not to be impatient: often, you will discover that you didn't know the basic stuff as well as you thought.
- ❷ Pick material that contains a good mix of equations and words. The idea of using text around equations is to give an intuitive understanding of their meaning, but you cannot implement the concepts with words alone. So, do not skip the equations. It defeats the purpose.

- ③ If you are (or training to be) a research scientist, then quantitative techniques are just a means to an end. So, look for a book that contains lots of examples from your own area of expertise. The mathematical concepts may be the same, but there is no doubting the motivational power of examples that are interesting to you.
- ④ You are unlikely to know, in advance, which techniques will come in handy later on in your research. Who knows? Maybe you will need to combine two apparently unrelated techniques to achieve your objective. So, early on, look for breadth rather than depth.
- ⑤ Don't be fazed by notation and terminology. Most terms and symbols have a common-sense, plain-language interpretation. I would hazard that only about 1 in 10 new terms requires repeated readings to digest. If you are struggling with every single term you encounter, you need to go back to a simpler text.

This book is geared towards these requirements. It's the sort of textbook that I wish I had as an ecology student. I have tried to give it a logical structure that nevertheless doesn't make the narrative so linear as to be boring. Each section prepares you for future sections, each chapter builds on previous ones and the entire book prepares you for more advanced texts that are certainly out there. Almost all chapters follow a classical entry to their subject matter and develop to more contemporary themes towards the end. This should re-animate faint high-school skills and then smoothly carry you to what you need to know for your research today.

Focusing on ecology was a selfish choice but it has allowed a happier co-existence between elementary and more advanced material. Basic maths and stats should not be seen as the bitter pill that has to be swallowed in anticipation of the good stuff. Therefore, while the later examples aim to elucidate abstract mathematical concepts by couching them in an ecological context, the earlier examples hopefully show that even basic high-school techniques can be used to address some interesting ecological questions. Indeed, the priority throughout is to convince you that quite a lot of useful quantitative ecology can be done with a modicum of technical knowledge and that some rather fuzzy ecological concepts can readily be recast and understood in formal mathematical language.

Having said that, I have not tried to pretend that quantitative ecology is easy – it is not. Making this admission means that I do not have to skip over the easier bits by pretending that they are trivial and neither do I need to hide the harder bits under the carpet. There are, therefore, few black boxes in the presentation of the theory. I realise that this is a risky decision because most ecologists don't particularly want to know what goes into the methods they are using, but perhaps this is the mentality that we should be working hard to change. As a result, this text contains hundreds of equations, but one of the earlier ones is $1 + 1 = 2$.

The inclusion of computing with R is a double blessing: from your point of view as a reader, it is useful to see how to implement complicated numerical solutions in practice with minimum effort. From my point of view, the extensive R libraries meant that I could include advanced techniques by explaining their theoretical basis but not their exact implementation. This trick makes for a lower page count and stretches some of the chapters towards the cutting edge of the discipline.

It is not the aim of this book to be an exhaustive guide to either R or the science of ecology, but it is most definitely intended as a comprehensive introduction to maths and stats for green scientists. Similarly, I am not hoping to convert ecologists into modellers (although using this book for a structured, two-semester course would go a long way towards this objective). Quite honestly, this material represents the minimum level of quantitative skills currently required of practising ecologists. The choice of topics is broad for two reasons: first, modern ecology is a melting pot of different quantitative concepts and techniques. The best advances in the primary literature seem to come from cross-fertilisation of ideas. Second, the biggest obstacle faced by a neophyte theoretician is psychological: lack of familiarity with the

basic terminology and scope of applications means that ecologists lack the confidence to tackle more technical papers and they are left looking for a 'way in' among textbooks that are either far too basic or way too advanced. As a result, many good hunches remain unsatisfactorily verbal (and unpublished) because they are conceived by colleagues who lack the overview of quantitative theory needed to either formalise their own notions or contact the correct specialist. If this book helps a few of these ideas make it into the primary literature, then writing it will have been worth my while. I hope you will find reading it just as rewarding, especially if one of those ideas belongs to you.

Supplementary material for this book (Exercises, Computer projects, R code, etc.) can be found on the online resource : www.wiley.com/go/quantitative_ecologist

Jason Matthiopoulos
St Andrews
20 May 2010

Thank you to . . .

. . .those who toiled

Popi Gkikopoulou, John Harwood, Helen Heyes, Debbie Russell, Gayatri Shanker, Debbie Steele,
Students of St Andrews MRes in Environmental Biology and in Marine Mammal Science (classes
2001–2010), Steve Smart.

. . .those who advised

Geert Aarts, Christian Asseburg, Nicole Augustin, Mike Begon, Tim Benton, Luca Borger,
Steve Buckland, Peter Corrigan, Will Cresswell, Carl Donovan, Ann Farrow, John Fieberg, Marie Guilpin,
John Halley, Sonja Heinrich, Monique Mackenzie, Marc Mangel, Juan Morales, Dave Moretti, Robert
Moss, Leslie New, Theoni Photopoulou, Sophie Smout, Matthew Spencer, Simon Wood, Mark
Woolhouse.

. . .those who provided

Susan Barclay, Richard Davies, Heather Kay, Ilaria Meliconi, Sheila Russell, Prachi Sinha-Sahay,
University of St Andrews, Wiley publishing house, the R development team.


. . .my loved ones

Spyros Matthiopoulos, Spyros Phevos Matthiopoulos, Effie Matthiopoulos, Valia Tavoularie-
Matthiopoulos



0

How to start a meaningful relationship with your computer (Introduction to R)



'Part of the inhumanity of the computer is that, once it is competently programmed and working smoothly, it is completely honest'

Isaac Asimov (1920–1992), author of science and fiction

This chapter looks and feels different to the rest of the book. It is short, contains no ecology and simply aims to familiarise you with the language used by scientific programmers and the particular conventions of R. It is not exhaustive, so all further R skills will be presented as needed in later chapters, in their appropriate mathematical, statistical and ecological context. The essential questions of what R is, why I chose to burden you with it and what it feels like to use it are covered in Sections 0.1–0.3. In Sections 0.4–0.7 you will find out where to obtain R and some of its valuable accessories, how to set them up in your computer and where to find help when you need it. I also outline the typesetting conventions that I will use to explain R code in this book. The last three sections (0.8–0.10) explain the basics of R usage and how to import data from other software into **data frames**.

0.1. What is R?

R is an open-source software package developed by a core team of academics and continually augmented by a large list of contributors. It is a numerical environment with a particular

bias towards statistical analysis and modelling. To some extent, it is what you make of it. It may be used interactively to interrogate a data set or as a programming language to construct simulations and automate complicated tasks. Despite being free to academic users, R compares favourably with other data-analysis and modelling software. For example, it can do considerably more than basic proprietary software such as SPSS or MiniTab and it competes well with very expensive software such as SAS, MATHEMATICA and MATLAB.

0.2. *Why use R for this book?*

It is generally better to teach scientific computing using real rather than pseudo-code. It is much better to understand the lofty concepts of programming through a particular language, any language. It is then easier to cross over to another if it is better suited to your purposes. There are several accomplished environments for data analysis and scientific programming but there are several reasons why the choice of R for this book is particularly sound.

The first is its overall suitability to the workflow of ecologists. In the field of quantitative software, packages have historically belonged to one of three camps:

- ❶ traditional programming languages like Pascal, Fortran and C with basic numerical libraries;
- ❷ mathematical software like MATHEMATICA and MAPLE with extensive libraries for symbolic analysis;
- ❸ statistical software such as S, R and SAS with extensive libraries for data analysis.

Any one of these would be suitable, but since the majority of quantitative ecologists spend most of their careers analysing data sets and running simulation models (or solving analytical models numerically), software from the third category seems to fit best.

Another important reason for choosing a software tool, particularly considering the time and effort required to become proficient in it, is longevity. R has a respectable pedigree (its foundations were laid in the software S that has existed since 1976) and it also has considerable potential. Currently, the momentum behind R shows no signs of abating and this augurs well for its future. This momentum guarantees the continuous supply of contributed packages to do almost any imaginable task, books at various levels of specialisation, online resources and text editors for programming (see Section 0.6). Many of these tools and textbooks are aimed at, or motivated by, ecological applications.

Finally, R is freely distributed under the Gnu public licence. There are great ideological reasons for supporting a piece of software developed by publicly-funded academics who then freely distribute their work, placing it at the service of the worldwide academic community. The fact that it is also open-source means that the number of good brains working to improve it is likely to exceed those employed by a private software company.

0.3. *Computing with a scientific package like R*

Most of the tasks that an ecologist would care to do on other specialised packages (e.g. geographical information systems, spreadsheets, databases etc.) can also be done in R, with one crucial difference: because it is a programming language, R is considerably more flexible and customisable. Using the built-in commands and the additional packages that can readily be downloaded from the CRAN website, you can write computer code for any imaginable task. This comes at a price to user-friendliness: as with any large tool-box, you need to know what the tools are for, how to use them and in what order. For larger tasks that need to be done

several times over, you will need to bundle together several tools in a well-defined sequence. This is called **programming**. If, in your career so far, you have only dealt interactively with a computer (ask a question, get back an answer, then ask another question, possibly based on the previous answer, and so on) then you might find that you need to shift your way of thinking about computers somewhat. Specifying complex tasks for a computer to do is an unforgiving and frequently frustrating job. Not only do you first need to perform the task manually (at least once) to make sure you know what you want done, but you then need to explain it to the computer unambiguously, in a language that looks nothing like written speech. Once this is done, you will often spend long hours looking at the screen wondering why on earth your apparently perfect piece of code comes up with an incomprehensible error message. The problem may be a tiny typo, a missing bracket or a fundamental logical inconsistency. Invariably, you will have to swallow the humiliation that, whatever the mistake was, it was yours and not the computer's. Even when these errors (or **bugs**) have been detected and fixed, there is always the possibility that the computer flawlessly performs a task other than the one you want. For example, a computer program may obligingly allow biological populations to recover from extinction long after their size has become negative. So, the process of **debugging** requires you to be untrusting and critical towards your own creation. This is probably one of the best life-lessons that your computer can teach you.

Once you have adjusted your expectations of how long it takes to develop a piece of code, things can only get better. You may start to enjoy the hunt for bugs, the creative process of constructing a functioning tool out of nothing, the rewarding feeling of uncovering the secrets of your data. Crucially, you will get better as a scientific programmer. You may even savour the rare occasion when code works perfectly the very first time you run it.

0.4. *Installing and interacting with R*

Day-to-day work with R involves the **R base package**, additional R packages as required, a good text editor and a quick-reference document of your liking. I explain what each of these is and how to obtain them.

The R base package contains the functionality required by most users including the basic user interface, mathematical, graphical and programming functions and all essential statistical tests and models. It can be downloaded from the Comprehensive R Archive Network (or CRAN for short) at <http://cran.r-project.org/>. You need to select the appropriate version for your operating system (Linux, Windows or Mac). You then need to follow the link to the base package and download the current version (v2.9.2 at the time of writing this book). Once prompted by a dialogue box, ask for the executable to be run and follow the default options in the various prompts of the setup program. When the program installs, start it up (e.g. by clicking at the desktop shortcut), you should see a screen like the one in Figure 0.1.

R uses a **command-line interface**, meaning that all the interesting stuff isn't done through the drop-down menu in the RGui window, but by typing commands next to the prompt (`>`) in the **R Console window**. When using it interactively, you type something at the prompt which can cause R to give you an output. Try typing something, say a numerical calculation, and then press return:

```
> 1+1
```

The response from R is

```
[1] 2
```

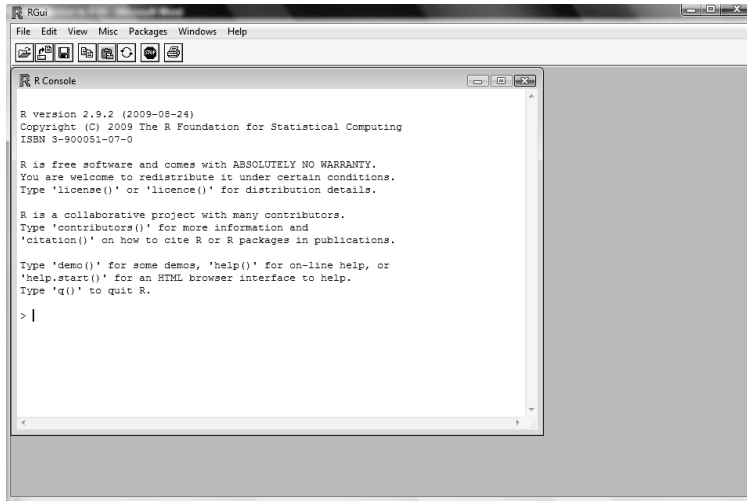


Figure 0.1: Start-up screen of the R command-line interface.

The serial number in square brackets indicates the order of a particular output line resulting from the previous user input. For example, to get R to print a list of all the years from my birth until writing this paragraph, the input and output would be:

```
> 1970:2009
[1] 1970 1971 1972 1973 1974 1975 1976 1977 1978 1979 1980
[12] 1981 1982 1983 1984 1985 1986 1987 1988 1989 1990 1991
[23] 1992 1993 1994 1995 1996 1997 1998 1999 2000 2001 2002
[34] 2003 2004 2005 2006 2007 2008 2009
```

The colon (:) always indicates a range of values. Depending on the width of your screen, R will generate the list in as few lines as possible, quoting the serial number of the next item within square brackets, at the start of each line.

When your input is not understood by R, you will get an error message:

```
> this_input_is_rubbish
Error: object 'this_input_is_rubbish' not found
```

The flow of printed information in the R Console is always downwards. Although you can scroll up to view your previous workings, and you can use your mouse to highlight and copy bits of printed input/output, you cannot navigate up to edit any part of your previous inputs. You may, however, use the up arrow key on your keyboard to quickly copy a previously typed line of code onto the currently active prompt line. Working with such one-liners is not a problem within the R Console but it can get cumbersome when you need to input several lines of code together. In these cases, a better alternative is to type your code in a text editor (Section 0.6) and copy/paste it into the R Console.

Assignments in R can be done with the arrow symbols (\leftarrow , \rightarrow). For example, to give a name (say, `years`) to the above list of years, you would type

```
> years<-1970:2009
```

An assignment prompts no response from R. The information is simply stored under the name `years`, ready for later use. If you want to inspect the information, type `years` and press return.

R is **vectorised**, meaning that operations can be applied to entire collections of things with the same ease as applying them to single items. For example, my entire list of birthday anniversaries can be calculated (and filed under the name `ages`) as

```
> ages<-years-1970
```

While the R Console handles the alphanumeric input and output, the **R Graphics Device** deals with images. This window appears separately within the RGui window. For example, a plot of `ages` versus `years` can be created by typing

```
> plot(years, ages)
```

Pressing return brings up the graphics window on screen (Figure 0.2). Feel free to rearrange and reposition this within your workspace.

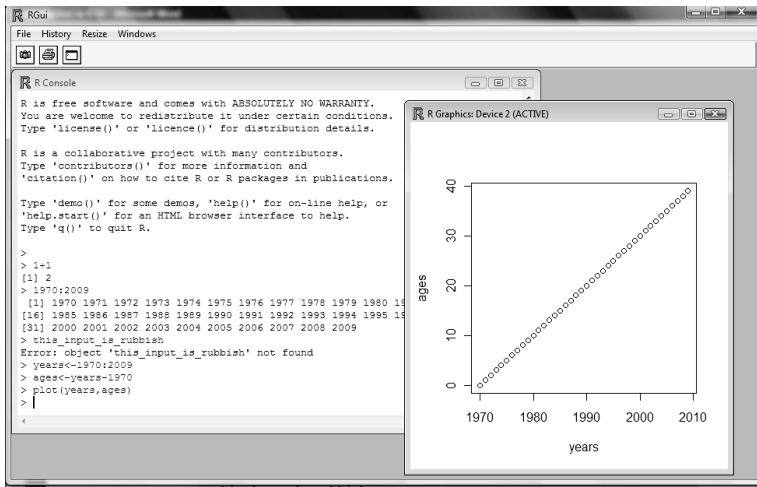


Figure 0.2: R Gui, R Console and R Graphics Device: the three main components of the R interface.

0.5. Style conventions

When editing your code outside of R, it is a good idea to use monospaced fonts (such as Courier), because, unlike proportional fonts (like Times), they allocate equal space to all characters. This retains the spacing of some tabular forms of output and makes code easier to debug. In this chapter, I have placed R input and output on a grey background. In the rest of the book, I use a grey background for the entire R boxes to make the computing sections more obvious among the rest of the text. When describing interactive use of R, I precede user input by the prompt (`>`) and use boldface for the output:

```
> 1+1
[1] 2
```

Larger pieces of code, of a length that might be typed up in a text editor and then pasted into R, are presented without the prompts, accompanied by short comments in English for most lines. Such detailed annotation is good practice when programming and not just for the benefit of others: I am always surprised by how hard it is to understand my own uncommented code a mere few weeks after writing it. The special character `#` tells R to ignore the remainder of that line so that when copy/pasting code, the comments do not interfere with the computation. In the book, such comments are shown in italics. For example:

```
# This plots age v calendar year for a person born in 1970
years<-1970:2009      # A list of years since 1970
ages<-years-1970      # A list of ages since birth
plot(years,ages)      # Generates the plot
```

0.6. Valuable R accessories

The functionality of R can be expanded by installing additional packages. A **package** is a collection of additional functions, example data sets and documentation. Whenever this happens in this book, you will be informed which package to get, but you need to know how. There are two types of packages: those that only need to be loaded into R and those that require a full install (i.e. downloading from the CRAN site and then loading into R). Both can be done via the ‘Packages’ drop-down menu in the RGui window. Upon selecting ‘Load package...’ you will be presented with a selection of about 28 packages. Simply select the one you need and press OK. The package is then loaded and can be used by your current R session. Alternatively, packages not on this list can be obtained by selecting ‘Install package(s)...’, again from the ‘Packages’ drop-down menu. You may be asked to select a CRAN mirror site. Just pick the one that is closest geographically to you. This list of packages is considerably larger. Pick your package and wait until it downloads and expands. Installed packages are saved on the hard disc and stay with your computer even after your R session ends. You do, however, need to load them into R when you start a fresh session. Go through ‘Packages->Load package...’ as before. You will notice that your recently installed package has made its appearance in this list. An alternative to using the drop-down menu to load a package is to do it using the commands `require()` or `library()`. For example, `require(MASS)` will load the package MASS. If your code requires a package, then place the `require()` command right at the beginning, so that the package is loaded before the rest of the code is executed.

The constant introduction of new R commands throughout the book might leave you overwhelmed by the apparent arbitrariness of their names. Rest assured you are not alone. Each computer language may use different names for different purposes and no programmer can remember more than a small vocabulary. Navigating help files is therefore an essential skill (see Section 0.7) but equally important is a good quick-reference guide of the names and syntax of the most frequently used R commands. My personal favourite was created by Tom Short and can be found at <http://cran.r-project.org/doc/contrib/Short-refcard.pdf>.

Finally, there is the delicate issue of the text editor to be used for developing longer pieces of code. You do need one, but the choice is a matter of taste (Figure 0.3). Word-processors are to be avoided, because their spell-checking and slow searching facilities tend to get in the way. You may, instead, use a fast and simple text editor such as TextPad (freely available from <http://www.textpad.com/products/textpad/index.html>). Better still, you can download a text editor that has been developed specifically for R programming. More information on editors can be found at http://www.sciviews.org/_rgui/projects/Editors.html.