

# THE TAG LOCATION PROBLEM

by

Michael D. Sumner, B.Sc., B.Ant.Stud. Hons (Tas)

Submitted in fulfilment of the requirements  
for the Degree of Doctor of Philosophy

Institute of Marine and Antarctic Studies  
University of Tasmania  
May, 2011



I declare that this thesis contains no material which has been accepted for a degree or diploma by the University or any other institution, except by way of background information and duly acknowledged in the thesis, and that, to the best of my knowledge and belief, this thesis contains no material previously published or written by another person, except where due acknowledgement is made in the text of the thesis.

Signed: \_\_\_\_\_  
Michael D. Sumner

Date: \_\_\_\_\_

This thesis may be made available for loan and limited copying in accordance with the *Copyright Act 1968*.

Signed: \_\_\_\_\_  
Michael D. Sumner

Date: \_\_\_\_\_

## Statement of Co-Authorship

Chapter 3 has been published in largely the same form as Sumner, M. D., Wotherspoon, S. J., and Hindell, M. A. (2009). Bayesian estimation of animal movement from archival and satellite tags. *PLoS ONE*, 4(10):e7324.

The following people contributed to the publication of work undertaken as part of this thesis:

Michael Sumner (50%), Simon Wotherspoon (25%), Mark Hindell (25%).

Details of the Author's roles:

*Wotherspoon and Sumner outlined the original mathematical and software approach. Wotherspoon was responsible for most of the mathematical content of the manuscript. Sumner completed the software, ran the models and wrote the manuscript. Hindell advised on the ecological goals and context and provided the data.*

**We the undersigned agree with the above stated “proportion of work undertaken” for the above published and peer-reviewed manuscript contributing to this thesis.**

Signed: \_\_\_\_\_

Supervisor  
Institute of Marine and Antarctic Studies  
University of Tasmania

Date: \_\_\_\_\_

Signed: \_\_\_\_\_

Institute of Marine and Antarctic Studies  
University of Tasmania

Date: \_\_\_\_\_

# ABSTRACT

This thesis provides an integrated approach to common problems with location estimation for animal tracking. Many existing techniques confound problems of location accuracy with simplistic track representations and under-utilization of available data.

Traditional techniques such as speed filtering and time spent maps are illustrated with a software package developed by the author with examples of location estimates from southern elephant seals. This software enables the application and exploration of various techniques that have previously not been available in a single solution. These include filtering, temporal gridding, projection transformation and GIS integration.

A novel Bayesian approach is introduced for the more general problems faced by different tagging techniques. This approach integrates all sources of data including movement models, environmental data and prior knowledge. This general framework is illustrated by application to satellite tag data and light-measuring tag data. Examples are used to detail the use of movement models with a powerful track representation model, and the application of raw light data for location estimation. Previously under-utilized sources of data are used to inform location estimates.

A method for applying light level geo-location within the framework is presented. This approach provides a primary location estimate for each twilight and utilizes all of the available data from archival tags.

These model runs result in very large databases of samples from Markov Chain Monte Carlo (MCMC) simulations and techniques for summarizing these for the variety of analysis outputs are illustrated. This system solves issues of location uncertainty with a full path representation and provides spatial maps of residency for multiple animals.

The relation between archival tag data and ocean circulation is used to extend the application of archival tag data for location estimation for diving animals in a manner similar to commonly used SST methods. Diving profiles from elephants seals are compared with 4D oceanographic datasets. Older tags are limited by problems with measurement lags for temperature—this problem is addressed with a proxy model for temperature at depth to ocean height.

This thesis provides a number of important improvements to the derivation of location from various types of tag data by integrating disparate information sources in a systematic way. Location estimates are produced with inherent quantification of errors. The approach provides the variety of metrics and analysis types required with an extensible software package. These contributions help bridge the divides between various analytic techniques traditionally employed for animal tracking.

# ACKNOWLEDGEMENTS

Primarily, thanks to Simon Wotherspoon and Mark Hindell for constant encouragement, guidance and patience. Supervision is crucial to this process and Simon and Mark are a crucial component of this work.

Thanks to Kelvin Michael and Corey Bradshaw for early encouragement and help. (Kelvin also deserves special thanks for ongoing meta-supervision).

Maria Clippingdale for helping start the whole thing in the first place. Greg Lee for lakeside conversations and for pushing me beyond my first programming language. Gareth James ensured that I left Sydney long enough to actually start this work. Toby Patterson for many conversations and emails, and specific guidance on early drafts and ideas. Aleks Terauds for lots of feedback and encouragement from the very beginning.

Jeremy Sumner and Sarah Cupit for constant support and encouragement, and helpful reviews on final drafts. David Pointing was a terrific office mate and remains a very good friend. Dave Procter and Edith O'Shea for being keen to hang out whenever we can.

John Corbett, Brett Muir, Sean Malloy, Steve Harwin and the Eonfusion team for a terrific work environment. Brett Muir provided an invaluable review at the last minute. The timeliness and attention to detail was incredibly helpful.

Mary-Anne Lea for so much enthusiasm and helpful feedback. Matt Dell for constant goodwill and for being a great mate. Michael O'Shea and Sean Malloy for helping build an awesome carpet burner. Students and colleagues at the Antarctic Wildlife Research Unit, WANKERS for life. Especially Judy Horsburgh, Michelle Thums, Virginia Andrews-Goff, Tim Reid, Iain Field, Caitlin Vertigan and Ben Arthur. Students and staff of IASOS, Maths and Physics and CMAR.

Dave Watts provided a whole swag of cool IDL code that helped me learn some of the mysteries of spatial and temporal data. John van den Hoff, Rachael Alderman, Jennifer Kingston Sinclair, Karen Evans, Chris Wilcox, Sophie Bestley, Kevin Redd, Kirk Buttermore, Klaus Engelberg, Roger Hill, Phil Ekstrom, John Gunn, Mark Bravington all provided specific wisdom and guidance.

There are a number of specific contributions of resources, code, data and knowledge that deserve special mention.

Rob Harcourt and the Macquarie University for support at the beginning of my candidature. Thanks to Gwen Beauplet for the use of the fur seal tag data in Chapter 3. Sergeui Sokolov for explaining the usage of the GEM proxy for sea surface height and providing example code. Sebastian Luque for many helpful emails over the years.

Thanks to Andrew Meijers for many enlightening discussions, and for generous assistance in the provision and use of the GEM model. Andrew provided very valuable guidance for the description of ocean processes, ocean models and climatologies.

The following acknowledgements are specific to the published version of Chapter 3.

NCEP Reynolds Optimally Interpolated Sea Surface Temperature Data Sets were obtained from the NASA JPL Physical Oceanography Distributed Active Archive Center, <http://podaac.jpl.nasa.gov>. The 2-Minute Gridded Global Relief Data (ETOPO2) were obtained from the National Geophysical Data Center, National Oceanic and Atmospheric Administration, U.S. Department of Commerce, <http://www.ngdc.noaa.gov>. Solar position algorithms of Meeus (1991) were adapted from the Solar Position Calculator made available by the NOAA Surface Radiation Research Branch, <http://www.srrb.noaa.gov/highlights/sunrise/azel.htm>. Greg Lee and Toby Patterson provided helpful suggestions on an earlier draft. Helpful feedback from two anonymous reviewers has been incorporated.

The R community is a force of nature and I owe a great debt to the creators and core developers of R, and to the very many package contributors. I greatly appreciate the efforts of colleagues and friends who have enthusiastically embraced this tool and helped me on the way. In particular I owe a debt of thanks Edzer Pebesma and Roger Bivand.

Manifold GIS and its awesome users for a lively and exciting community.

Tim Callaghan provided the LaTeX template and styles for the entire structure of the document. This thesis was written mostly in Emacs with Sweave and R, using Subversion, BibTeX and JabRef. To build the document figures and displayed code I have used R 2.12.0 with the following packages: MASS, mgcv, deldir, spatstat, lattice, sp, rgdal, maptools, geosphere, zoo, maps, mapdata, RODBC and ff.

Any errors are of course mine and cannot be attributed to my very generous helpers. When something takes as long as this did there are lot of people to acknowledge. I have surely failed to mention everyone or to fairly pay my thanks so if this applies to you please accept my apologies and gratitude. It was appreciated.

Norm, my best friend. Life is horrendously unfair and you improved mine immensely. I wish desperately that you could share the rest.

My mother and father for love and support, and providing me with everything I've ever needed.

Simone van der Heyden deserves the world for putting up with me during this project: I wouldn't have finished without you.

*We have named our language R—in part to acknowledge the influence of S and in part to celebrate our own efforts. Despite what has been a nearly all-consuming effort, we have managed to remain on the best of terms and retain our interest in computers and computing.*

Ross Ihaka and Robert Gentleman, 1996

*In light of this, I've come to the conclusion that rather than “fixing” R, it would be much more productive to simply start over and build something better.*

Ross Ihaka, 12 September 2010

# TABLE OF CONTENTS

<b>TABLE OF CONTENTS</b>	<b>i</b>
<b>LIST OF TABLES</b>	<b>iv</b>
<b>LIST OF FIGURES</b>	<b>v</b>
<b>1 INTRODUCTION</b>	<b>1</b>
1.1 The thesis . . . . .	2
1.2 Chapter outlines . . . . .	2
<b>2 THE TAG LOCATION PROBLEM</b>	<b>4</b>
2.1 R and the <b>trip</b> package . . . . .	5
2.2 Problems with location estimation . . . . .	5
2.2.1 What is a trip? . . . . .	7
2.2.2 Practical data issues . . . . .	8
2.2.3 Joining the dots . . . . .	10
2.2.4 Summary of problems . . . . .	12
2.3 Summarizing animal behaviour from point-based track data . . . . .	13
2.3.1 Filtering . . . . .	13
2.3.2 Non-destructive filtering . . . . .	18
2.3.3 Spatial smoothing—surfaces from track data . . . . .	18
2.3.4 Partitioning tracks and grids into time periods . . . . .	26
2.3.5 Map projections . . . . .	26
2.3.6 Tools for traditional methods . . . . .	29
2.4 The trip package for animal track data . . . . .	29
2.4.1 Generating trip objects . . . . .	30
2.4.2 Filtering for unlikely movement . . . . .	35

2.4.3	Creating maps of time spent . . . . .	40
2.5	The need for a more general framework . . . . .	43
2.5.1	Location estimation from raw archival data . . . . .	46
2.5.2	Filtering locations . . . . .	47
2.5.3	Intermediate locations . . . . .	48
2.6	Conclusion . . . . .	49
<b>3</b>	<b>BAYESIAN ESTIMATION OF ANIMAL MOVEMENT FROM ARCHIVAL AND SATELLITE TAGS</b>	<b>50</b>
3.1	Introduction . . . . .	51
3.2	Materials and Methods . . . . .	53
3.2.1	Ethics Statement . . . . .	53
3.2.2	Assumptions . . . . .	53
3.2.3	Satellite tags . . . . .	55
3.2.4	Archival tags . . . . .	55
3.2.5	Posterior estimation . . . . .	56
3.2.6	Examples . . . . .	57
3.3	Results . . . . .	59
3.3.1	Argos tag data set . . . . .	61
3.3.2	Archival tag data set . . . . .	64
3.4	Discussion . . . . .	67
3.4.1	Path . . . . .	67
3.4.2	Precision . . . . .	67
3.4.3	Latent estimates . . . . .	67
3.4.4	Combinations . . . . .	69
3.4.5	Updating the models . . . . .	69
<b>4</b>	<b>LIGHT LEVEL GEO-LOCATION FROM ARCHIVAL TAGS</b>	<b>72</b>
4.1	Light-based derivation techniques . . . . .	73
4.1.1	Fixed point methods . . . . .	73
4.1.2	Curve methods . . . . .	74
4.1.3	Aims of this study . . . . .	74
4.2	Methods . . . . .	76
4.2.1	Auxiliary data . . . . .	79

4.2.2	Movement model . . . . .	80
4.2.3	Estimation . . . . .	80
4.3	Examples . . . . .	81
4.3.1	Discussion . . . . .	81
4.4	Conclusion . . . . .	88
<b>5</b>	<b>REPRESENTATION OF TRIP ESTIMATES</b>	<b>89</b>
5.1	Representation of location estimates . . . . .	90
5.2	Primary locations and intermediate locations . . . . .	91
5.3	Features of the representation scheme . . . . .	96
5.4	Examples . . . . .	96
5.5	Conclusion . . . . .	101
<b>6</b>	<b>AUXILIARY DATA FOR LOCATION ESTIMATION</b>	<b>102</b>
6.1	A problem of scale . . . . .	103
6.2	Data availability . . . . .	103
6.3	Slabs versus points . . . . .	105
6.4	Temperature applications . . . . .	105
6.4.1	Temperature profiles . . . . .	106
6.4.2	Ocean height proxy . . . . .	107
6.4.3	Masks . . . . .	109
6.5	Large data set example . . . . .	110
6.6	Conclusion . . . . .	111
<b>7</b>	<b>CONCLUSION</b>	<b>112</b>
<b>A</b>	<b>MARKOV CHAIN MONTE CARLO</b>	<b>114</b>
A.1	Metropolis Hastings . . . . .	114
A.2	Example . . . . .	115
<b>B</b>	<b>ARGOS DATA ESTIMATION EXAMPLE</b>	<b>119</b>
B.1	Argos data example . . . . .	119
	<b>BIBLIOGRAPHY</b>	<b>124</b>

# LIST OF TABLES

3.1	Estimate precision for Argos data set Summary of precision calculated from the posterior for $x$ by original Argos class (km). Each row presents a quantile summary for the CI ranges (95 %) from each Argos class for longitude and latitude. The seven classes are an attribute provided with the original Argos locations (Service Argos, 2004). . . . .	64
3.2	Estimate precision for archival data set. Summary of precision calculated from the posterior for $x$ from the archival tag. A quantile summary for the CI ranges for longitude and latitude. . . . .	67
A.1	Comparison of exact and MCMC approximations for $N = 500$ samples and $N = 5000$ samples for the Binomial example with $n = 20$ and $y = 12$ . . . . .	118

# LIST OF FIGURES

2.1	Raw Argos estimates for a Macquarie Island elephant seal. The line connecting the points is coloured from blue through purple to yellow in order relative to the time of each position. (Macquarie Island can just be seen in black beneath the third dark blue line segment). The outward and inward journeys are much faster than the journey throughout the Ross Sea, as shown by the colour scale change. A graticule is shown for scale, 10 degrees on the main plot, and 20 degrees on the inset. . . . .	6
2.2	Land filter and Argos quality class filter. In the first panel any point occurring on land has been removed, and in the second any point with an Argos class of quality < "B" has been removed. In each panel the state of the track prior to filtering is shown as a dotted line. The coloured filtered line uses the same time-based colouring as in Figure 2.1. . . . .	15
2.3	Recursive speed filter applied using the McConnell method to the track data shown in Figure 2.1. The small panels show the RMS speed calculated for each point for two of the ten iterations of the filter. At each iteration any point with an RMS speed above the required maximum threshold speed (the horizontal line) is flagged (shown as "x") and removed. The RMS speed axis is logarithmic. The resulting filtered track is shown in the larger panel on the right with the same temporal colour scale as used in Figure 2.1. . . . .	17
2.4	Argos track filtered by penalizing by sum of squares speed. The filtered track is shown with the same time-based colour scale as in Figure 2.1, and the unfiltered track is shown as a dashed line. The original data was used without resampling. . . . .	19
2.5	Argos filtered longitude and latitude from Figure 2.4 by penalized sum of squares speed. The coloured lines use the same time-based colour scale as in Figure 2.1. The unfiltered longitude and latitude are shown as a grey line. . . . .	20
2.6	Track lines binned into a coarse grid, using a line segment value to sum into each cell. Cells are coloured from light to dark blue for increasing values. The sum of values resulting from each smaller line segment is shown for each cell. . . . .	23

2.7	Track lines binned into a fine grid. The track and grid origin is the same as in Figure 2.6. The coarser grid outline is shown for occupied cells as thick black lines. . . . .	24
2.8	Side by side plots of line gridding and KDE gridding to a coarse grid. The input line is shown at an arbitrary height above the cells. . . .	25
2.9	Side by side plots of line gridding and KDE gridding to a fine grid. The input line is shown at an arbitrary height above the cells. . . .	26
2.10	Raw Argos estimates and penalized smoothed track plotted in raw longitude and latitude. Great circles are drawn as curves, from Macquarie Island through the most distant raw Argos position (solid) and the most distant smooth position (dashed). . . . .	27
2.11	Raw Argos estimates and penalized smoothed track plotted in an equal-area projection. Great circles (as-the-crow-flies) are shown as lines. These pass through the Macquarie Island site to the most distant raw Argos position (solid) and the most distant smooth position (dashed). . . . .	28
2.12	Plot of Argos track data from a <code>trip</code> as points, with lines showing the speed filtered tracks, coloured for each separate trip. . . . .	37
2.13	Original <code>sdafilter</code> (red) and custom <code>trip</code> speed distance angle filter (thick grey line). The original raw track is shown as a thin grey line. . . . .	41
2.14	Simple image of a time spent grid for the trip object. . . . .	42
2.15	Four variations on a <code>tripGrid</code> . <code>grid.filtered</code> and <code>grid.c026</code> use the simple line-in-cell method for all four seals and for seal <code>c026</code> alone. <code>kde3.filtered</code> and <code>kde1.filtered</code> use the kernel density method for line segments with a sigma of 0.3 and 0.1 respectively. Each grid has dimensions 100x100 and time spent is presented in hours. . . . .	44
2.16	A relatively coarse version of <code>fourgrids</code> with a specific kilometre-based cell size (80x60km), specified within an equal area map projection. . . . .	45
2.17	Light level geo-locations for black marlin in the Coral Sea. This image is taken from Figure 3a in Gunn et al. (2003) and is reproduced with the permission of the authors. . . . .	48
3.1	<b>Satellite tag data and estimates</b> Panel A: The sequence of original Argos estimates for an adult female Weddell seal tagged in the Vestfold Hills, with time scale from red to blue. All location classes are shown. The different length scale bars for north and east represent 10 kilometres. Panel B: Posterior means for $x$ from the Argos data set plotted spatially, with time scale from red to blue as in panel A. The sequence is far more realistic, without the noise and positions on land. Panel C: Map of time spent from full path estimates from the Argos data set. The density represents a measure of time spent per area incorporating the spatial uncertainty inherent in the model. Bin size is 150 m by 140 m. . . . .	60

3.2 **Estimates and time spent for archival data set** Panel A: Posterior means for  $x$  from the archival data set plotted spatially, with time scale from red to blue. The sequence provides a realistic trajectory for an elephant seal. The dashed grey line shows the (approximate) position of the Southern Boundary of the Antarctic Circumpolar Current. Panel B: Map of time spent from full path estimates from the archival data set. Bin size is 5.5 km by 9.3 km at 54 S and 3 km by 9.3 km at 72 S. . . . . 61

3.3 **Individual longitude, latitude estimates for Argos** Posterior means for  $x$  from the Argos data set for longitude and latitude, with time scale from red to blue as in Figure 3.1. The grey line shows the implied sequence of the original Argos estimates. Also shown is the range of the 95 % CI of each estimate (km), determined with the mean by directly summarizing the posterior. . . . . 63

3.4 Intermediate estimates for Argos Posterior means for  $x$  of longitude and latitude for a short period (23-26 Feb 2006) with CI ranges shown. The CI range for intermediate estimates (full path) is shown as a continuous band. . . . . 65

3.5 **Posterior means for archival data set** Posterior means for  $x$  from the archival data set for longitude and latitude, with time scale from red to blue as in Figure 3.2a. Also shown is the range of the 95 % CI of each estimate (km), determined with the mean by directly summarizing the posterior. . . . . 66

3.6 **Intermediate estimates for archival data set** Individual mean estimates of longitude and latitude for a 10 day period in February with CI ranges shown, as well as the CI range for intermediate estimates (full path) shown as a continuous band. . . . . 68

4.1 Locations around Macquarie Island ( $158^{\circ} 57'E$ ,  $54^{\circ} 30' S$ ). Square icons indicate locations along the current twilight line, triangles indicate the corresponding “evening” locations. Macquarie Island is at the intersection of the two curves. . . . . 78

4.2 Morning twilight solar elevations for locations around Macquarie Island from Figure 4.1 (squares). The time at which the elevation is 0 (sunrise) is indicated with a vertical line, and the horizontal lines show solar elevations -5 and 3. . . . . 79

4.3 Morning twilight solar elevations for location around Macquarie Island from Figure 4.1 (triangles). The time at which the sun is at elevation is 0 (sunrise) is indicated with a vertical line, and the horizontal lines show solar elevations -5 and 3 . . . . . 79

4.4 Mean longitude and latitude values from the posterior, plotted with surrounding 95% confidence intervals for elephant seal example. . . 82

4.5 Mean longitude and latitude values from the posterior, plotted with surrounding 95% confidence intervals for fur seal example. . . . . 83

4.6	Mean longitude and latitude values from the posterior, plotted with the 95% confidence intervals (closed circles) and the full binned estimate boundary (open circles) for the fur seal example. . . . .	84
4.7	Mean longitude and latitude values from the posterior, plotted with the 95% confidence intervals (closed circles) and the full binned estimate boundary (open circles) for the elephant seal example. . . . .	85
4.8	Time spent estimate for entire trip derived from all posterior samples for Macquarie Island elephant seal. . . . .	86
4.9	Time spent estimate for entire trip derived from all posterior samples for Amsterdam Island fur seal. The zonally (east/west) oriented lines in the inset are mean locations of Southern Ocean fronts (Orsi et al., 1995). Only the Subantarctic front is seen in the main figure region. . . . .	87
5.1	(A) Simple track of three locations in a traditional representation, moving from left to right. The grey region is included to represent a boundary to the animal's movement, such as a coastline. (B) The points represent unknown locations with some uncertainty. (C) A representation of these location estimates which are disconnected in time. (D) Connecting the measured locations are regions that are again unknowns with some uncertainty. . . . .	92
5.2	(A) Between the first two measured locations is a region of unknown migration connecting the two end points, which can be rather wide. (B) The second connected region, more constrained relative to the first in terms of possible lateral movements. (C) Notation of $x_i$ for primary estimates. (D) Notation of $z_i$ for intermediate estimates. . . . .	93
5.3	Illustration of the grid index scheme. Two child grids are shown within the parent, with indicated index values. . . . .	94
5.4	Example showing a real child grid within a parent grid. The inset window shows the detail of the grid. . . . .	95
5.5	A short region of an Argos track of a ringed seal in Newfoundland. The primary estimates shown as contours are compared to the original Argos estimates. All but three Argos estimates have class "B" and are marked with "+", the others are labelled with their classes "2" and "A". The confidence in the "2" location is reflected by the proximity of the estimate to it and that it is relatively localized. None of the estimates fall on land like some of the Argos locations, and some are tightly bound by the land—but not clipped arbitrarily to it. . . . .	97
5.6	Original Argos track with sequence of modal locations from the posterior. The dark line joins "X" symbols derived as the mode from estimates shown in Figure 5.5. The modal locations are feasible for a ringed seal, but the line joining them introduces a problem by traversing the land area. The original line of the raw Argos track is shown as a dashed line. . . . .	98

5.7	Individual intermediate estimates shown as contours with modal locations from the primary estimates shown as dark “X” symbols (as in Figure 5.6). The intermediate estimates overlap the primary estimates in space as well as connecting them in time. . . . .	99
5.8	Time spent estimate from combination of the nine intermediate estimates. This single surface is constructed by combining the separate intermediate estimates for the time duration, each weighted by their corresponding time interval. . . . .	100
6.1	Temperature and salinity profiles from two regions for an elephant seal track. Seal dive profile values are shown as a red line. The points are all data values from the profiles taken from the GEM model for the corresponding seal depths at the grid of locations shown. The first is near Macquarie Island, the second near Adélie Land. . . . .	108
6.2	Series of location likelihood maps based on SSH proxy from temperature at-depth. . . . .	109
A.1	Two samples drawn by the Metropolis Hastings algorithm for the Binomial example with $n = 20$ and $y = 12$ . The plots show a histogram of $N$ MCMC samples and the exact posterior density, for $N = 500$ and $N = 5000$ . . . . .	117

## CHAPTER 1

# INTRODUCTION

Modern animal tracking analysis consists of a diverse set of techniques used by a wide range of researchers in many disparate fields of science. There are established traditional techniques that range from very simplistic spatial representations of tracks (Samuel et al., 1985; Bradshaw et al., 2002; Croxall et al., 2004), to powerful statistical models that rely on modern computing hardware for their solution (Jonsen et al., 2005; Patterson et al., 2008; Breed, 2008). There are significant divides between fields in terms of the importance of certain techniques and in the access to existing tools. Some of these differences are due to the species being studied or the priorities of the research, and some are due simply to the established habits of a community or their access to the necessary tools.

In terms of the available software tools, there is a divide between the tool users and tool makers in that many of the problems involved with analysis do not come to light until algorithm development is attempted (Calenge et al., 2009). The major problems with estimates of tag location accuracy and reliability are well-understood, but the related issues of data consistency and management are rarely seen except by those who develop the software tools (Coyne and Godley, 2005; Halpin et al., 2006). Somewhat paradoxically, with widespread access of modern computing hardware and software, more researchers are exposed to the implementation details applied to raw data and track analyses and are required to act as practitioners managing and analysing data (Block et al., 2003a; Hartog et al., 2009).

The availability of computing and visualization tools has provided the ability to scrutinize analyses against existing knowledge or intuition. Mapping tools enable realistic representation of animal tracking analyses and can present insights or problems that were not initially apparent. There is great promise for accessible software to enable greater realism, or at least provide the representation of very complex models involving time-developing 3D with multiple variables of interest (Andres et al., 2009).

Animal tracking is a modern field that has seen great advances made with researchers working closely with tag manufacturers to ensure that tag development closely matches the priorities and requirements for data collection (Wilson et al., 2002; Afanasyev, 2004). This same improvement has been occurring with software tools as researchers “get their hands dirty” working with raw track data. The development of techniques like MCMC for the solution of Bayesian models provides

the ability to model very complex problems with spatial and temporal structure, involving large and disparate data sets (Dennis, 1996; Dixon and Ellison, 1996).

## 1.1 The thesis

By the application of Bayesian methods we can utilize all information sources to help with the variety of problems of location estimation for animal tracking. An investigation of traditional track methods shows that the point and line model is too simplistic and reveals a number of inherent problems. From this perspective we see that no matter the method, approaches that deal with tracks as a time series of points or lines are inherently limited, and that these limitations apply to data collected by a variety of tag types. Some process that generates a time series of locations has a deeper rawer source of data behind it and it is this that fits best in an integrated Bayesian approach.

This thesis considers problems with animal tracking data, focussing on the estimation of location. There is some consideration of data handling, mostly from an estimation and quality control perspective. One of the main goals is to take advantage of the large collection of existing data sets, rather than attempt to design best methods for obtaining data. Traditional techniques are presented to illustrate the main problems involved with tag location analysis and provide a toolkit for executing these traditional techniques. Many researchers are aware of cutting-edge analytical techniques that are improving the outputs from tracking data, but access to these is often difficult and there is a lack of understanding of their limitations.

## 1.2 Chapter outlines

Chapter 2 provides an overview of general problems with the analysis of tracking data, assuming that we have location estimates to begin with. We look at the common features of tracking data and some of the common problems faced when using these data. The chapter provides a toolkit for dealing with tracking and associated data, pointing out some of the limitations when applying simple models to quality assurance. The chapter concludes with the need for integration of track analysis with the derivation of location from raw data, such as from archival tags.

Chapter 3 provides a general Bayesian framework for estimating location by applying any available data, prior knowledge and movement models. The theoretical approach is provided with examples of using the framework for two different data sets.

Chapter 4 presents the application of the Bayesian framework to an archival tag data set in full detail. A curve method for light level geo-location is introduced that provides the practitioner with open access to every component of the model in freely available tools.

Chapter 5 presents methods for representing track estimates. This includes a suite of summary tools for the model outputs, and the use of these data for visualization and further diagnostics.

Chapter 6 extends the application of archival tag data to use subsurface temperature for diving animals. A four dimensional ocean data set is used to demonstrate the potential for subsurface archival tag data and applying very large data sets. Alternative methods are presented using surface temperature as an efficient mask, as well as temperature at depth as a proxy for sea surface height.

## CHAPTER 2

# THE TAG LOCATION PROBLEM

This chapter discusses problems faced with tracking data that concern the estimation of location and provides a flexible software environment for exploring data and applying solutions. Examples are used to illustrate the variety of problems and some of the limitations of the traditional techniques by which tracking data are analysed. The **trip** package is a dedicated suite of tools developed by the author in the software language R. This package integrates data handling and analytical techniques for track data with broader GIS and spatial data tools. This makes track data handling tools easily available, even for those without strong programming skills. The chapter concludes by extending the concerns regarding the limitations of traditional techniques to methods for deriving locations from raw data.

This chapter is not intended to be a critique of modern methods of dealing with tracking data, but introduces the variety of issues encountered and tools for applying them. Simple-to-use tools for handling spatial and temporal data are still rare and some of the problems encountered cause difficulties for researchers before they have an opportunity to explore sophisticated methods. The aim here is to illustrate some classical techniques within a software toolkit that provides better control over the details of analysis tools for those without advanced programming skills. Later chapters present solutions for the remaining problems. Work by Patterson et al. (2008) and Breed (2008) provide a more critical review of recent methods.

Aims of this chapter:

1. To introduce existing problems in tracking analyses presented with examples of classical techniques.
2. To illustrate the complexity of problems and areas that require more sophisticated solutions than traditional techniques. The problems presented here illustrate the need for solutions that come later in the thesis.
3. To present a flexible and readily customized software package as a framework for classical analyses and starting point for more sophisticated analyses.
4. To explain the compromises that are often made with regard to data representation and storage, as dictated by traditional systems, rather than an expressive model of the problem represented by the spatial and temporal data.

5. To encourage the use of techniques for automatic data validation, spatial and temporal data storage and integration with database and GIS technologies.

## 2.1 R and the trip package

The software package **trip** developed with this chapter provides an integrated system for data validation and a development framework for track analyses. This can be used as a launching point for further analysis such as validating input to Bayesian methods, or filtering for state-space models (Patterson et al., 2010). As an extension of the R environment, **trip** also provides integration with tools for data access and output, integration with GIS and other data systems, and metadata for projections and coordinate systems. The **trip** package ties together numerous tracking analysis techniques, which previously were only available through a wide variety of disparate tools, each having various requirements and limitations.

The **trip** package was developed within the freely available software platform R, a statistical programming environment consisting of a vast community of contributing developers and users (R Development Core Team, 2010). R is organized into modules known as *packages* which provide the functionality of the language, and also the mechanism by which it is extended<sup>1</sup>. New packages are created using the same tools by which R itself is built and can be contributed to a public repository such as the Comprehensive R Archive Network (CRAN<sup>2</sup>). The repository system for contributed packages is one of the great strengths of R and is part of the reason for its ease of use and subsequent popularity. The spatial and temporal capabilities of R are advanced, including strong integration with other technologies such as databases, GIS software and the wide variety of spatial and other complex data formats.

The tools provided in R for coercion between complex data types are very powerful, allowing rather different packages to share data with tight integration (Bivand et al. (2008); Chambers (2008)). There are some fundamental data representations required for spatial analysis and careful organization is needed to provide coercions between different types to get all the tools to work together. The underlying spatial tools used to create the **trip** package are described in Section 2.4 with examples of using the software.

## 2.2 Problems with location estimation

This section presents actual tag location estimates to illustrate common problems associated with track data. The location data were provided by System Argos, which is a worldwide satellite service that provides location estimates from small mobile transmitters.

The first example is a sequence of Argos position estimates for a single elephant seal in Figure 2.1. All raw estimates provided are shown, with the longitude and latitude values transformed to an equal-area projection and drawn connected as a

---

<sup>1</sup>See [http://en.wikipedia.org/wiki/Package\\_Development\\_Process](http://en.wikipedia.org/wiki/Package_Development_Process).

<sup>2</sup>See [http://en.wikipedia.org/wiki/Software\\_repository](http://en.wikipedia.org/wiki/Software_repository).

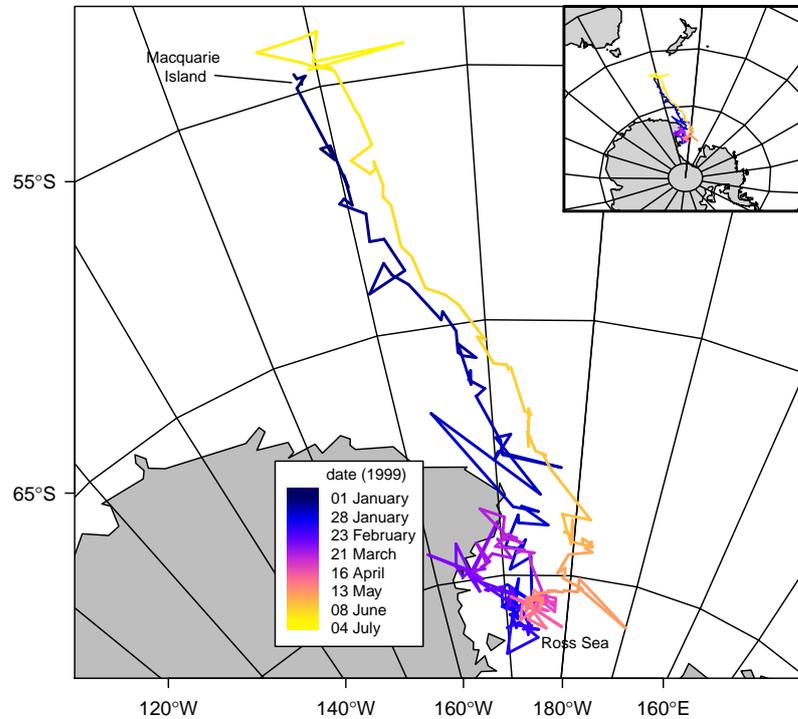


Figure 2.1: Raw Argos estimates for a Macquarie Island elephant seal. The line connecting the points is coloured from blue through purple to yellow in order relative to the time of each position. (Macquarie Island can just be seen in black beneath the third dark blue line segment). The outward and inward journeys are much faster than the journey throughout the Ross Sea, as shown by the colour scale change. A graticule is shown for scale, 10 degrees on the main plot, and 20 degrees on the inset.

line. While there is obvious noise in the track, the general sequence of events is clear: the seal leaves Macquarie Island, swimming predominantly south-east to the Ross Sea where it spends several weeks, and then returns via a similar path in reverse.

There are a number of problems with the location estimates, some that are very obvious, but others that are more subtle. First, some of the dog-legs in the path seem very unlikely. On the outward journey the blue path shows some lateral movement to the west and east, and just before the return to Macquarie Island there is a similar movement to the west, then east. These are obvious problems seen as noise in the track, with positions that cannot be believed that do not otherwise obscure the general pattern of movement. Other dog-legs in the path are less extreme, so are more plausible.

Another problem is that there are locations that are well within the land mass of Antarctica. For a marine animal, these locations are clearly incorrect but as with the track dog-legs there are similar issues with levels of plausibility that are

difficult to evaluate. A location on land may be plausible if it is near enough to the coast, though this can interact with further issues requiring interpretation. These include that the start and end location are not exactly at the known “release” and “re-capture” sight, which was the isthmus at Macquarie Island. This isthmus is quite narrow and is readily crossed by elephant seals, though regions of the island to either side of the isthmus can only be traversed by sea. Another section of the track at the beginning has the path of the animal crossing Macquarie Island itself. At the scale of this image this inaccuracy seems unimportant since the island is such a small area within the study region. However, if the region of land concerned was a much larger peninsula then the problem of choosing which region was actually visited remains.

A scheme that proposes to remove or correct extreme positions faces the problem of defining appropriate thresholds. “Extreme” dog-legs cannot simply be discarded as the question of what is “too-extreme” does not have a simple answer. A simple rule to discard any location on land will not work since these animals do actually visit coastal regions. The distance that a seal might travel inland is not very far but depending on the species studied and the environment the situation may not be so clear-cut.

There are other issues of plausibility. For example, the coastline in the figure is quite coarse and while it may be sufficient for the scale of the image it does not represent the actual coastal boundary available to the seal. The real coastline is far more tortuous, detailed and dynamic—and may be significantly different from a particular data set due to the current fast- or sea-ice cover. This is a general issue with any data set available for informing models of animal location—the assumptions and limitations must be understood and used appropriately.

In terms of the incorrect first and last positions in Figure 2.1, these could be updated to be the actual release and recapture sites, but it might be more correct to actually add those locations and times to the start and end of the sequence of records. This is a data consistency issue that leads to the next family of problems in track data.

### 2.2.1 What is a trip?

There are a number of practical issues associated with the organization of track data that can present more prosaic problems. This section discusses some terminology and suggests the use of a “trip” as the unit of interest and that can be defined with database-like validation restrictions. The idealization of an animal’s trip is a continuous line in four dimensional space-time that is perfectly accurate as a representation of the seal’s position. For practical reasons, this ideal case can only be represented as a series of point samples of uncertain accuracy, in two or three spatial dimensions parameterized by time.

Animal tracking can be carried out in a variety of ways, here restricted to the broad class of data derived from “tagging”. A “tag” is a device attached to an animal that is used to directly sense and record data or that is indirectly detected by a remote sensing system. For the purpose of the current discussion, refer to the first type of tag as “archival” and the second type as “remotely sensed”. Reviews of the practical methods for tagging in the broader context of biotelemetry for marine

and terrestrial species are provided by Cooke et al. (2004), Wilson et al. (2002) and Kenward (1987).

Archival tags record and store data that is later retrieved from the tag, while remotely sensed tags emit an electromagnetic or acoustic signal that is detected by an installed system of sensors, such as a satellite or acoustic array. (This categorization is not always maintained in practice, as archival tags may be satellite-linked in order to upload data in near real-time, but for the purpose of location estimation the distinction holds for the types of available data).

A loose set of definitions then is:

**tag** the device put on the animal.

**track data** any temporally referenced location data resulting from a device attached to an animal.

**trip** a specific tracking “interval” where a tag was attached, the animal is released for a time, and the tag (or just its data) is eventually retrieved. A trip may be represented in some way with track data that has some quality control or modelling applied.

Data resulting from the tagging process are identified by the device ID, as well as by the identity of the individual animal the tag is attached to. The same tag might be put on different animals, or the same animal might have been tracked on two or more separate occasions. For central-place foragers, a single tagging event may involve the same animal leaving the tagging site and returning multiple times. Once the interval of interest is defined it can be regarded as a trip. A single leave / return or a tag deployment / retrieval may be referred to as a trip. Whether multiple leave / return events are considered within a single trip depends on the research question—migratory animals usually won’t have such clear trip boundaries as central-place foragers, for example. The difference is important as the behaviour of interest for central-place foragers is primarily between leave / return events, and the return event is usually the only opportunity to retrieve a tag. Finally, there may not be data for the entirety of the trip of interest due to tag loss or memory restrictions, and so require the inclusion of trip sections where location is uncertain or completely unknown.

For the current discussion, define a trip to coincide with the interval of interest for which there is useable data from the tagging process. “Tracks” or “track data” then are just a set of location and other data, with no particular organization or quality control.

### 2.2.2 Practical data issues

The minimum organization and quality control for trip data involves the ordering and relation between data records. The ordering of records is perhaps inconsequential, as there is the inherent order of the date-time value stored for each location, but this may reveal more basic errors in the data. There must not be occurrences of duplicated date-time records within a single trip, although duplicated locations in

subsequent records are acceptable. Duplicates in time do not make sense since either they are redundant copies of a previous record, or there is an implied infinite speed. These are common from the Argos Service, an example is found in the `seal` data example given by Freitas (2010) which is used in Section 2.4.2.<sup>3</sup> Analytical methods sometimes apply a non-zero time difference arbitrarily to avoid divide-by-zero errors. Less serious is the issue of successive duplicate locations, but care must be taken when calculating metrics such as speed based on inter-point distances. Each of these cases should be investigated carefully in case they hide errors from other causes such as mistaken data handling.

Missing values must also be handled carefully. Location and time coordinates cannot be used if they are missing or non-finite, even if their record appears in the correct order. Missing values can arise in a number of ways—infinities or undefined numeric values from numeric errors, or out of bounds coordinates, transformation errors, data missing from a regular sequence—and the exact reasons need to be carefully understood.<sup>4</sup> This is a different approach taken to that of Calenge et al. (2009) who explicitly allow missing coordinates as part of “trajectories”. This is most pertinent in the context of tracks of regular time intervals where a missing point can be significant in terms of interpretation. The definitions here are not intended to determine which approach is more appropriate and there is no reason the two rationales cannot co-exist, but the current implementation in the `trip` package disallows missing coordinates.

From a programming perspective, the use of rigid classes (definitions) with validity checking can significantly reduce the time wasted solving these problems (Chambers, 1998). Based on the above, the minimal data-consistency preparation required can be achieved in the following way. Read all records, sort by trip ID then date-time, remove duplicated records or records with missing or non-numeric spatial or temporal coordinates. (The definition of “invalid” for a coordinate may involve out of bounds values such as those for longitude and latitude, but this step only refers to the data values, not their interpretation). Remove or adjust any records with duplicate date-times within a single trip ID. Up to this point no interpretation has been applied to the data—this will provide a useable set of records that can pass minimal validation but each step should be carefully investigated to ensure that automated decisions are not introducing new errors.

One way to adjust duplicate times is to simply modify the values forward or back by a small amount, but this can be problematic depending on the time differences involved. The reason for duplicated times is more likely to be a problem with the data itself and should be investigated.

Other problems in this regard deal with the sensibility of movements in a particular coordinate system. The most commonly used coordinate system for tracking data is longitude and latitude on the WGS84 datum. For animals that traverse

---

<sup>3</sup>Another recent example of duplicated times in a GPS data set is discussed here: <http://lists.faunalia.it/pipermail/animov/2010-August/000635.html>

<sup>4</sup>A natural assumption is that recorded values of date-time are correct beyond question: so there is some information even if one of the spatial coordinate values is missing. This issue is a corollary to the use of filtering techniques that remove locations from track data or otherwise correct spatial positions. If there is a date-time why not interpolate or otherwise estimate missing spatial coordinates?

hemispheres and cross critical meridians such as the Pacific Ocean dateline (longitude 180 W / 180 E) or the prime meridian (longitude 0) a continuous path must be represented appropriately, such as longitudes in  $[-180, 180]$  or  $[0, 360]$  respectively. Many species will cross both these critical boundaries and so representing simple lines requires a smarter choice of map projection. All map projections have these regions of non-optimal usage and so the focus should be on intelligent choice of projection using tools that provide easily applied transformations.

### 2.2.3 Joining the dots

A further problem is the common practice of joining points with “straight lines”. Usually the only available data are temporally referenced point locations, and lines are artefacts introduced for visual purposes. However, using these lines is quite artificial, and can become error prone when used quantitatively. Joining the points imposes a specific model of behaviour, namely that the path is a straight line between points.

This is not correct on several levels. First, the animal is moving in three spatial dimensions not two, and the movement in this third dimension is quite significant for diving animals, though it may be largely ignored for many flying or surface dwelling species. Second, even if the points represent accurate positions for the animal the line joining them most certainly does not represent the intermediate path correctly. The animal could be traversing either side of the line, or taking a far longer, more convoluted path. Thirdly, the coordinate system used to interpolate the intermediate positions can have a large effect on the outcome. “Straight-line” movement is usually assumed, but what is drawn as a straight line on a map has a very different meaning depending on the coordinate system or map projection used. For most coordinate systems shorter step lengths will be closer to the “great circle” path, but the nature of the deviation will also depend on the region traversed (Gudmundsson and Alerstam, 1998).

Joining points with a line guides the eye helpfully to show the sequence of points, and the mind can often overlook problems of inaccuracy to see roughly what actually happened. It is this mental capacity for reducing noise and seeing the overall picture of events that sophisticated models of track analysis and movement aim to replicate in an objective way. When our minds provide these ideas they do so by applying knowledge of the physical system: an animal swimming through water over great distances, an animal that will tend to travel quickly to an area of interest, then spend time in localized regions, an animal that will not venture far from the coastline to areas inland, etc. “An effective EDA [Exploratory Data Analysis] display presents data in a way that will make effective use of the human brain’s abilities as a pattern recognition device” (Maindonald and Braun, 2007).

There is no end to this problem when dealing with points or lines segments themselves as the entities of interest. If a particular position is just slightly wrong, and its neighbouring points also a little inaccurate then any assessment of the distance from one point to another or the intermediate path taken between points is thrown into question.

### Treatment of spatial and temporal data in modern software

The temporal nature of track data stems from the fact that the physical process of animal movement is a continuous path. This continuous process is only measured by discrete samples and so the data are inherently discontinuous. However, treatment of time in software is rarely truly continuous but rather applied as a sequence of “time slices”. This is a legacy limitation that unfortunately matches the way in which track records are usually measured and stored. To choose a high-profile example, animations of tracks in Google Earth (Google, 2010) show sequences of discrete line segments that are progressively revealed or hidden as the slider intersects the time spans of the segments. Why is the line not represented as a continuously changing entity, with extents that match the slider’s extent exactly? Partial line segments could be shown, and the line shown continuously without being so bound to its input points. This is a problem not only for Google Earth but a true limitation in the representation of most spatial data in GIS and GIS-like visualizations.

This must be part of the reason why tracking analysis is rarely tightly coupled with GIS—analytically (if not visually) track data is treated as continuous or near-continuous, with more information than the static line segments joining subsequent points. Also track data is routinely processed based on great circle travel (assuming that’s how the animal would choose to move) but then presented visually in a simple 2D longitude by latitude plot. Map projections provide visualizations that help prevent our brains from giving us the wrong message about distance and area on a simple plot. Ultimately a 4D visualization on a globe may be a “truer” way to visualize track data, but though current tools such as WorldWind and Google Earth will draw lines along great circles they are not well suited to track data that varies continuously in time.

GIS traditionally provides complex shapes such as multi-segment lines with multiple branches, or polygons with branched holes and islands but support for a third coordinate value for elevation is rare, and time is usually non-existent.<sup>5</sup> Though routine computer graphics in games provides complex surfaces and lines composed of primitive elements with incredibly complex representations and interactions, it is rare to find treatment of track data as a multi-part line object, let alone with fine control over the continuous span of a single line. Modern GIS in its most commonly understood terms is not easily extended for temporal data, but provides an excellent platform for dealing with data sources, geometry and gridded data, and map projections.

The availability of data manipulation and analysis tools is a major factor in the effective use of animal tracking data for ecological studies. While there are many analytical techniques and a wide array of software applications, some lack the flexibility required or are restricted by cost or the required expertise to use them effectively. For some purposes track data needs to be represented as points, and for others as lines, or even better as probability density functions. Tools for seamless conversion between these data structures are practically non-existent for everyday research.

---

<sup>5</sup>Polygons are literally incapable of representing continuous 2D topological surfaces in 3(+D) geometric space and the special status of planar polygons that imposes this limitation surely will eventually be transcended by future GIS technology.

An illustrative example of the limitations of GIS data structures is seen when attempting to represent track data. As points, the geometry is stored as X and Y coordinates (and, rarely, with Z coordinates). Time is relegated to the attribute table and even then is not always supported completely by common GIS interchange formats.<sup>6</sup> GIS supports more complex geometry than simple points requiring more than one vertex: lines, polygons and “multipoints”. It should be simple to store a track in either a “point” or “line” version, but for lines each line segment is composed of two vertices so there is no longer a simple match between a point’s date-time (or other) coordinate and those of the line. The line is represented as a single object with multiple X and Y vertices with only a single attribute record, or as a series of line segments composed of two X and Y vertices each. Neither version provides a clean translation of even very simple track data to a GIS construct.

Access to the true continuous nature of track data is simply not provided by common software tools. This is a problem that extends to a general definition of topology versus geometry, for representing objects flexibly in a chosen space but discussion of that is out of scope here. Hebblewhite and Haydon (2010) highlight the need for ecologists to become more adept at matching temporally varying environmental data to animal movement data. There are emerging technologies that allow for a separation of geometry and topology, unlimited coordinate attributes on GIS features, and generalizations of this are permitted by general database theory (Butler et al., 2008; Beegle-Krause et al., 2010; Pauly et al., 2009; Anderson et al., 2010; Fledermaus, 2010).

#### 2.2.4 Summary of problems

The main problems can be described as a set of overlapping issues:

**Inaccurate sampling** Position estimates are inaccurate, with some unknown relation to the true position.

**Irregular and incomplete sampling** Position estimates represent discrete samples from an animal’s continuous path. These may be at irregular time intervals with large gaps in the time series, and no ability to control this because of practical limitations.

**Incomplete paths** Paths composed of too few positions, inconsistent motion and assumptions about straight line movement.

**Unlikely dog-legs** There is no sense in the animal being so erratic.

**Simplistic models of movement and residency** Intermediate locations are shown by joining the dots, using an assumption of direct linear motion between estimates.

Many traditional analyses of modern track data deal with these problems by chaining a series of improvements in an *ad hoc* way, and the need for better approaches is well understood (Breed, 2008; Patterson et al., 2008). Incorrect positions

---

<sup>6</sup>The obscure “measure” value for a fourth coordinate in shapefiles is sometimes used for time, but was not designed for it and is rarely supported by software packages.

are removed by filtering, based on speed, distance, angle, attributes on the location data or spatial masking (McConnell et al., 1992; Austin et al., 2003; Douglas, 2006; Croxall et al., 2004; Freitas et al., 2008). Positions are updated by correction with an independent data source, such as sea surface temperature (SST) or the need to traverse a coastal boundary (Beck and McMillan, 2002; Shaffer et al., 2005). Unlikely dog-legs are removed by filtering, or “corrected” by smoothing the line. Smoothing is also used to augment small samples, by interpolating along a smooth line, or smoothing positions into a 3D grid partitioned by time (Bradshaw et al., 2002; Tremblay et al., 2006; Campagna et al., 2006). There are further requirements for smoothing to estimate latent locations or to match disparate spatial and temporal scales.

Many of these techniques have their own problems, compounded when these operations are chained one after the other. Models of the process may be overly simplistic (linear movement between estimates), or applied inconsistently—positions are removed, then estimates are smoothed, or compared with other data to correct or update them. Later chapters present new methods for incorporating these issues in a more integrated way.

## 2.3 Summarizing animal behaviour from point-based track data

This section revisits some of the problems presented previously and looks at the details of algorithms used. The techniques are useful for first-pass summaries, or exploring ideas, but they rely on simplistic models and are difficult to integrate sensibly.

Putting aside the limitations mentioned earlier and the fact that there is no clear basis for deciding which combination of tests should apply, some of the issues can be illustrated further by proceeding with these simple to more complex filters.

### 2.3.1 Filtering

Filtering is used to remove or modify data in some way based on metrics available or calculated from the track data. Destructive filters categorize locations for removal from the trip. Non-destructive filters update the location data for some positions. Again there is no clear distinction between these two types as a filter can be used to discard some locations entirely, update others and interpolate new locations for various purposes.

At the simplest level, destructive filtering involves categorizing each point for removal or retention. An example is a “land mask” that would deal with the issue of the points on the Antarctic continent as discussed in Section 2.2. A land mask filter would determine any points that fall on land, mark them for removal and discard them. The filter is very simple as it is a basic check for each point individually, with no interaction of its relationship to other points or other data sources. All points that fall on land can be determined and removed in one step, or they could

be checked one after another in any order. The way the filter is applied will have no impact on the filtered result.

A more complex case applies recursion, where once some points are removed the status of the test for remaining points may have changed and so must be determined again. Metrics on successive locations fundamentally rely on the order and relation between points, and so once points are removed the calculations must be repeated until the desired metric is reached for all retained points. Existing filters apply measures such as Argos location quality, distance between successive points, speed of movement, turning angle and land masks. A classic speed filter in marine animal tracking is a recursive rolling root-mean-square speed filter by McConnell et al. (1992). This filter is widely used and widely cited especially in marine applications.

There is a practically endless number of metrics that can be derived from the location data that range from the very simple to complex. However, no matter what combination of decisions are applied, the main limitation of these methods is their arbitrary nature. They are applied to a purely geometric interpretation of the data that largely ignores the physical system being modelled. Much information goes unused, and what data is used is applied independently of other sources.

The use of destructive filters is also problematic because data is discarded and the filter decision is based on the location itself, rather than the process used to estimate the location. It is hardly ever mentioned, but the Argos Service estimation process is not published and therefore not amenable to modelling.

Recursive filters are relatively complicated, but still result in a categorical decision as much simpler filters like a land mask—there is no single number that can be calculated for a given point, and the implications of minor decisions for a given filter can greatly affect the result.

### Destructive filtering

Here the use of two types of destructive filter are demonstrated to remove points based on a land mask, Argos quality and speed calculations. In Section 2.4 the **trip** package is used to create a version of a speed-distance-angle filter.

The Argos Service is a global satellite data collection system that provides an estimate of position based on doppler-shift signals transmitted from tags (Service Argos, 2004). The basic location product is a time series of longitude and latitude values with a categorical measure of location quality that is provided as a label. There is more information available with the service and guidelines for its use, but the scope of the following discussion is restricted to the widely used quality class measure. Location classes take the values “Z”, “B”, “A”, “0”, “1”, “2”, or “3” in order of increasing accuracy. The reported accuracies are > 1000 m for “0”, and that 68% of fixes should be <1000 m for “1”, <350 m for “2”, and <250 m for “3” (Service Argos, 2004; Goulet et al., 1999). No estimate is given for “Z”, “B” or “A” classes although studies have shown that these can have an acceptable accuracy (Vincent et al., 2002).

The first filter removes any points that fall on land, and then any points that have an Argos class of “B” or worse. In Figure 2.2 the two panels show the Argos

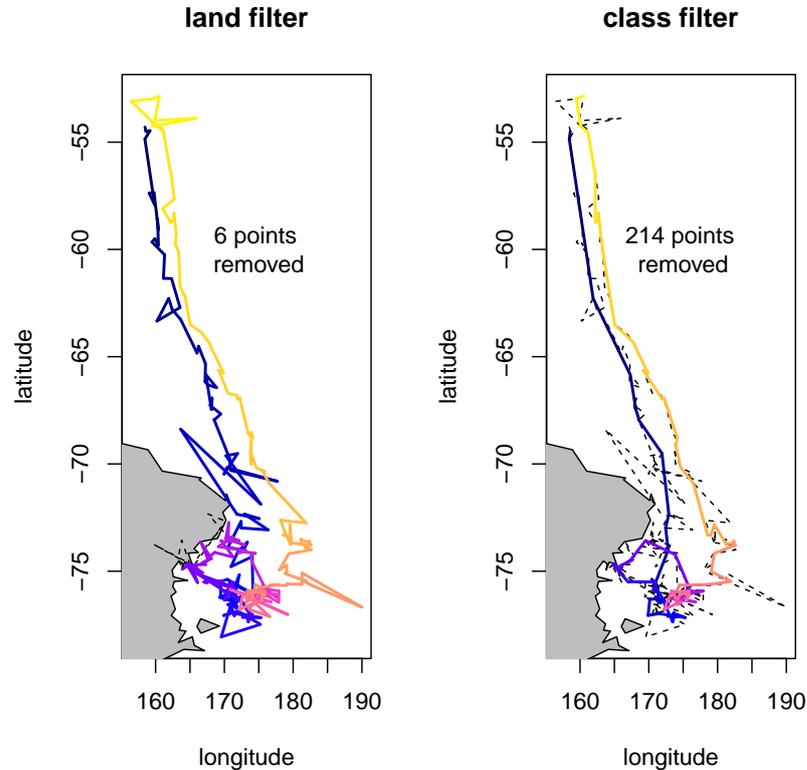


Figure 2.2: Land filter and Argos quality class filter. In the first panel any point occurring on land has been removed, and in the second any point with an Argos class of quality  $< \text{“B”}$  has been removed. In each panel the state of the track prior to filtering is shown as a dotted line. The coloured filtered line uses the same time-based colouring as in Figure 2.1.

track plotted in longitude and latitude coordinates.

In the first panel the original track has been filtered for points on the land area, and the second for points that have a location quality class of “B” or “Z”. The land filter only applies to six points in the track, but the class filter has removed 214, which is a majority of the available 351 points. The effects that these filters have are independent of one another and it would not matter which were performed before the other, although the result may be quite different in combination than in the use of either one alone.

In terms of the land mask, the filter succeeds in removing points from land but there is still a line segment that crosses a portion of the Antarctic continent. A filter that applies to points is quite different to one that applies to a continuous line segment. The class filter provides a track that seems more realistic than the noisy input track, but at the expense of discarding more than half of the available data.

The next example demonstrates a recursive speed filter applying the method of

McConnell et al. (1992) to the Argos track. This technique specifies a running root-mean-square (RMS) summary of speed calculated by determining the instantaneous speeds for each point to each of its two previous and next locations. The RMS is defined as the square root of the mean of the speeds squared. Any sequence of locations with an RMS exceeding the maximum speed value have the peak RMS location discarded and the RMS is recalculated. This continues until all locations have an RMS less than the maximum speed. The threshold applied in this example is 12 km/h. Again the track is presented as longitude/latitude coordinates with distance measurements calculated relative to the WGS84 ellipsoid.

In the left panels of Figure 2.3 are two plots of the RMS values, for the second iteration after some positions are removed and the sixth iteration when only a few nearby positions remain above the threshold. The unfiltered RMS values are shown as a grey line representing the original points, the current points that have RMS values above the maximum are shown as crosses and the points below the maximum are shown in faded black. The threshold speed is shown as a horizontal line. As successive peaks of RMS values above the maximum are removed the categorization for the remaining points changes. This filter took ten iterations to remove all of the 73 locations that imply a speed of movement faster than the threshold, and the resulting track is shown in the right panel.

This result seems reasonable, though there is no end to the complexity and arbitrary nature of the decisions to be made. There are practical uses for techniques like these however. Speed filtering can give a reasonable result and there is a need to quantify the limits here with comparative studies of various algorithms on known data sets.

There are more sophisticated filtering algorithms that apply a suite of tests. A published speed-distance-angle filter first removes a minimum class (such as Z), filters for speed with the McConnell algorithm and then recursively removes further offending points that have a combination of acute angles and long implied distances travelled (Freitas et al., 2008). However, the accompanying software is quite specific and unrelated to other packages. In Section 2.4 a speed-distance-angle filter is defined completely using tools available in the **trip** package.

Related filters applying a combination of metrics have been published by Douglas (2006), Croxall et al. (2004) and Austin et al. (2003). The algorithm published by Austin is available in the R package **DiveMove** (Luque, 2007) and those by Freitas et al. (2008) and McConnell et al. (1992) in **argosfilter**. Freitas et al. (2008) gives an overview of the relative merits of various techniques, and the need for hierarchical techniques to avoid discarding good quality locations.

Another decision process is to choose between two possible locations for each time step provided by the Argos Service in the DIAG format (Service Argos, 2004). The Douglas Argos-Filter Algorithm applies this as part of the filter (Douglas, 2006).

### 2.3.2 Non-destructive filtering

Non-destructive filters aim to update track points by some mechanism without removing corresponding data records. There are many examples such as algorithms

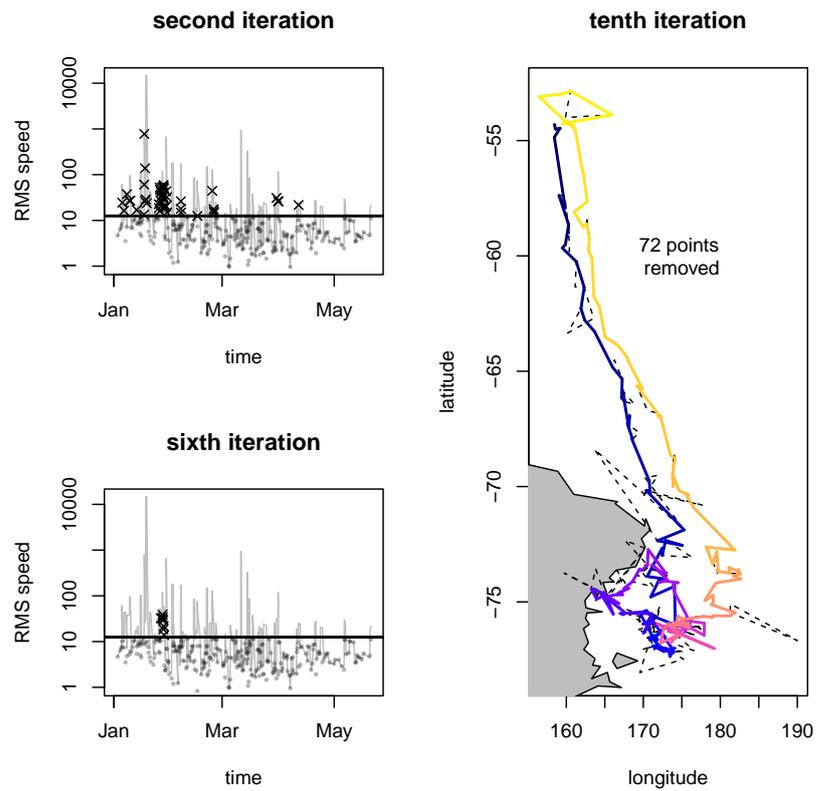


Figure 2.3: Recursive speed filter applied using the McConnell method to the track data shown in Figure 2.1. The small panels show the RMS speed calculated for each point for two of the ten iterations of the filter. At each iteration any point with an RMS speed above the required maximum threshold speed (the horizontal line) is flagged (shown as "x") and removed. The RMS speed axis is logarithmic. The resulting filtered track is shown in the larger panel on the right with the same temporal colour scale as used in Figure 2.1.

to modify light level geo-locations for latitude with SST (Beck and McMillan, 2002) and interpolative techniques to smooth tracks (Tremblay et al., 2006). Many destructive filters can be recast in a non-destructive form using a penalty smoothing approach in the style of Green and Silverman (1994). For example, rather than filter the track to exclude locations that would imply an unrealistic speed of travel, the track can be smoothed using a speed of travel penalty. The smoothed locations are determined by minimizing the functional

$$J_\lambda(\hat{x}) = \sum_{i=1}^n d(x_i, \hat{x}_i)^2 + \lambda \sum_{i=1}^{n-1} v(\hat{x}_i, t_i, \hat{x}_{i+1}, t_{i+1})^2$$

where  $\{x_1, \dots, x_n\}$  represent the raw (unsmoothed) locations,  $\{\hat{x}_1, \dots, \hat{x}_n\}$  their smoothed counterparts,  $d(x, y)$  represents the distance from  $x$  to  $y$ , and  $v(x, t_x, y, t_y)$  the speed required to travel from  $x$  at time  $t_x$  to  $y$  at time  $t_y$ , and  $\lambda$  is the smoothing parameter. The first term is a measure of goodness of fit of the smoothed locations  $\hat{x}_i$  to the raw locations  $x_i$ , while the second is a speed penalty. Minimizing  $J_\lambda$  trades off goodness of fit against speed of travel. When  $\lambda = 0$  the smoothed track reproduces the raw track exactly. Increasing  $\lambda$  favours tracks requiring less extreme speeds, at the expense of reproducing the  $x_i$ .

As for the more traditional application described by Green and Silverman (1994) this process can be interpreted in a Bayesian context. Adopting a penalty based on squared speeds is equivalent to adopting a Gaussian prior on speeds, adopting a penalty based on the absolute values of speed is equivalent to an exponential prior on speed.

An application of this non-destructive filter was applied to the example Argos data set discussed earlier. Figure 2.4 shows the filtered result on a map and Figure 2.5 shows the same result with longitude and latitude plotted against time. The result seems reasonable with a plausibility comparable to that of the recursive speed filter with the advantage of not having removed any data from the trip.

### 2.3.3 Spatial smoothing—surfaces from track data

There are a number of ways of creating bivariate (surface) estimates from track data, the most common involve a surface defined by a grid of cells or connected points. (Regular grids are historically easy to store and to compute, and so are applied most commonly—irregular grids and meshes are not considered here). The most direct methods are variously called “rasterization”, “gridding” or “pixellation”—these effectively generate a bivariate histogram from the input point or line data. Kernel methods apply a bivariate distribution (such as a Gaussian) to the input points or lines. Point methods are not discussed here, but can be achieved in a similar way to the line-based examples shown in Section 2.4.

The following is a very simplified folk history of tracking techniques, but does help explain some of the existing practices that differentiate terrestrial and marine applications. There is an apparent difference between terrestrial and marine applications in the way that surface creation, or gridding methods are applied. Terrestrial applications tend to ignore time or adapt data to avoid temporal auto-correlation. VHF techniques were the original primary tool for wildlife tracking, and the metric

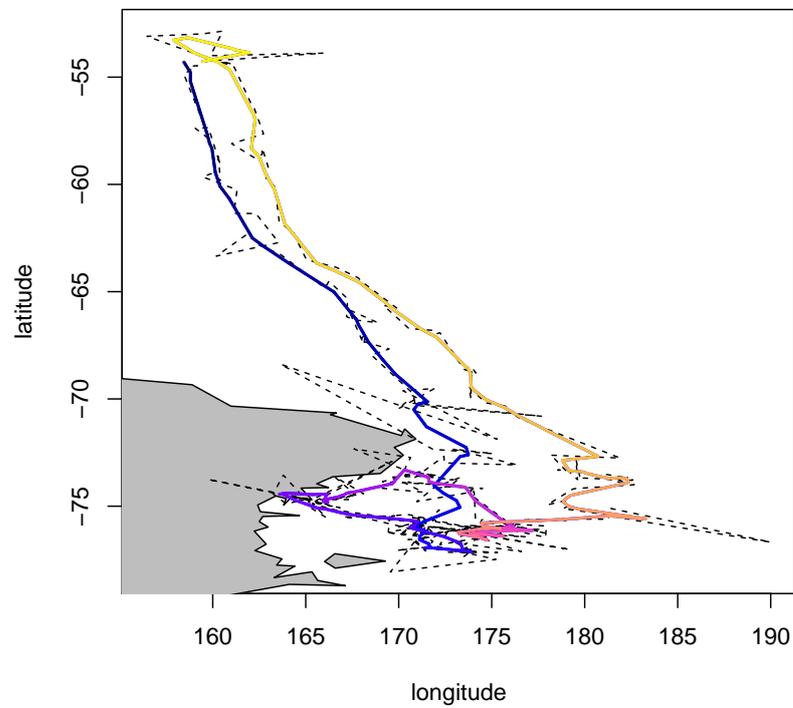


Figure 2.4: Argos track filtered by penalizing by sum of squares speed. The filtered track is shown with the same time-based colour scale as in Figure 2.1, and the unfiltered track is shown as a dashed line. The original data was used without resampling.

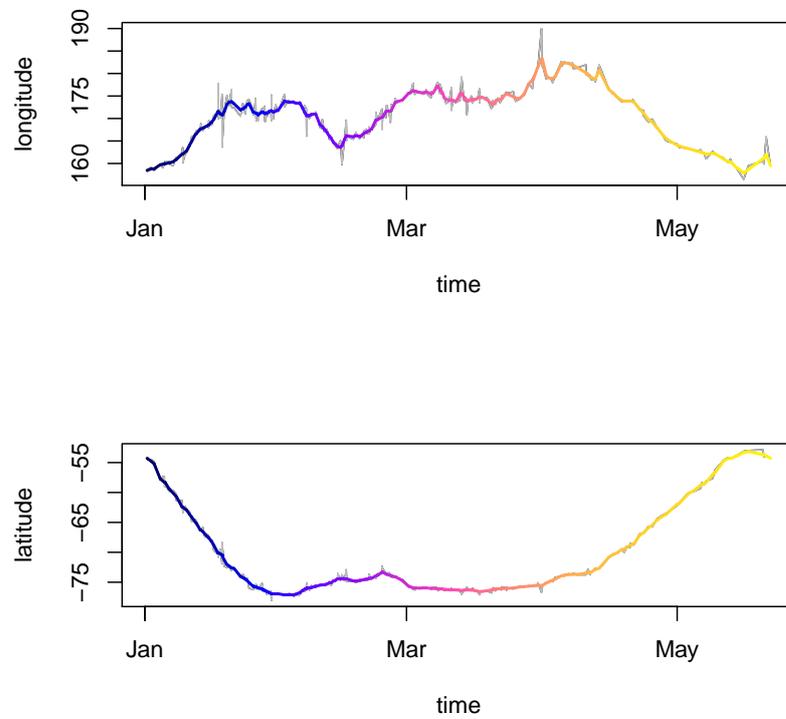


Figure 2.5: Argos filtered longitude and latitude from Figure 2.4 by penalized sum of squares speed. The coloured lines use the same time-based colour scale as in Figure 2.1. The unfiltered longitude and latitude are shown as a grey line.

of interest was pure residency—the minimal region in which the animal is present (home range) and the animal’s core region. In this context actual “tracks” are not the main interest. Marine applications have traditionally been explicitly interested in tracks and the temporal relationships, perhaps because of the real and perceived differences in the dynamics of marine environments. Modern techniques are seeing a far greater cross-over in these originally different fields and the differences are now out-weighed by the common goal of reliable location estimation.

Cell binning and kernel density estimation (KDE) can unproblematically convert a linear geometric track representation into a 2D histogram-like smoothing of residency, time spent map, or other utilization distribution but both must grapple with serious problems of interpretation. Points and lines simply do not represent the movement of an animal completely. Both methods must deal with issues of independence, point or line interpretations and complex boundaries and environmental relationships. If these issues can be dealt with or ignored, both cell gridding or KDE can be used as a convenient smoother for track data.

The following combinations of methods are applied in various ways in many existing publications (for example see Seaman et al. (1999); Wood et al. (2000); Nel et al. (2002); Bradshaw et al. (2002); Croxall et al. (2004)). The grid surface is generated by operating on points or on line segments. Line segments provide a continuity through space and so they provide a natural way to connect regions visited by the animal. Lines present a harder problem to convert to a grid than points and so this is often approximated by providing interpolated points in place of line segments, assuming constant travel speed.

The influence of the points or lines on the surface is calculated by binning into the overlapping grid cells directly, or by calculating the contribution to neighbour cells via a “kernel”. In the case of binning, the coverage provided by overlay with cells is relatively small, and completely dependent on the chosen bin size. This leads to compromises balancing positional accuracy and the need to match the scales of covariate data Bradshaw et al. (2002). These limitations are the same as those for histograms in general—the discontinuous histogram presents analytical difficulties as it quite sensitive to the chosen origin, bin size and orientation (Simonoff, 1996; Silverman, 1998). Kernel density has the advantages that the result is smooth without the blocky, discontinuous nature of a bivariate histogram. Regular or irregular “wireframe” representations give a smoother result and have continuous analytical and visualization counterparts via interpolation, but these data structures are more complicated to calculate and are much less widely supported.

Finally, when the cell value is determined there are various ways in which the contribution of a point or line can be calculated. The point or line can simply be summed into the cell—a point is counted as present in the cell whereas the proportional length of an overlapping line segment is added to the bin. Each bin contribution can be multiplied by a factor, such as a point value or the time interval available to a line segment. This is one of the natural cases for using line segments from track data rather than points, as when the duration of time is of interest the input points cannot represent the duration of time. Many studies approximate this by creating a regular time series by interpolation.

When kernel density methods are applied, a weighting factor can be given for ei-

ther dimension of the kernel, or more generally a two-dimensional correlated weighting is used (Simonoff, 1996).

The following examples show the gridding of a very simple track by some different methods.

### Exact gridding

Gridding (pixellation or rasterization) methods generate a grid of cells that extend completely over the region of input track data.

A very simple track is shown in Figure 2.6 with an underlying grid. Each line segment has a value assigned to it, the time duration between each track point. These values are [2, 3, 1, 7, 1, 1] and the first and last line segments are quite long, so their contribution to the cells is small relative to the middle segments. This is reflected in the grid, which represents the implied “time spent” in each cell. In order to determine the contribution of each line segment’s value to a cell, the segment is split on the boundaries of the cells that it crosses and the value shared proportionally based on the length of the resulting segment-portion. This computation is not simple—to describe it in GIS terms this is an overlay of the line and the cells (“topology overlay”) with a rule to transfer values from the lines to the cell, in this case a proportional sum. GIS can be used to perform these calculations, but as discussed in Section 2.2.3 working with time in GIS is not well supported and this must be done as an attribute on line objects, rather than on inherently continuous lines that vary through time or other dimensions as well as space.

The dependence of the gridding method on bin size is easy to see. Figure 2.7 shows the same gridding process applied to a finer grid, with the same origin. The grid again completely summarizes the contribution of the track, but the actual regions influenced by the grid are quite different—the small grid cells snugly trace the track and cover a much smaller overall area than in the case of the coarse grid. The total time duration represented by the coarse and fine grid is exactly the same, but the results are quite different.

This dependence on scale has been used explicitly to provide a compromise between positional accuracy and spatial coverage with the need to match with environmental data by Bradshaw et al. (2002) and Burns et al. (2004). Bradshaw et al. (2002) also explicitly used grid size to help account for spatial uncertainty. As discussed previously the assumptions made by these analyses become substantial, involving issues such as uniform location accuracy, uniform grid scale and constant straight line travel.

### Kernel density methods

Kernel density methods are applied to overcome the dependence of the probability density function on the origin and bin size, as it is the properties of the kernel that dictate the contribution to each cell. (In theory every input point or line contributes to every cell to some degree). The result is thus less related to the data structure choices, given a sufficiently fine resolution. Figure 2.8 shows the



Figure 2.6: Track lines binned into a coarse grid, using a line segment value to sum into each cell. Cells are coloured from light to dark blue for increasing values. The sum of values resulting from each smaller line segment is shown for each cell.

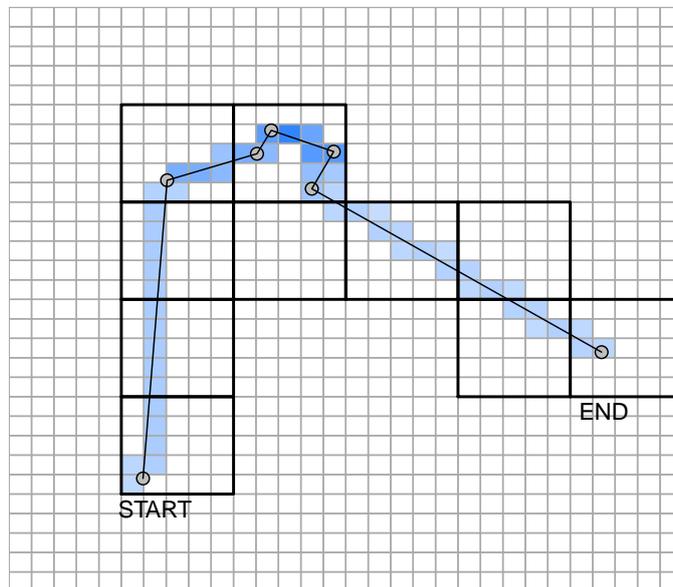


Figure 2.7: Track lines binned into a fine grid. The track and grid origin is the same as in Figure 2.6. The coarser grid outline is shown for occupied cells as thick black lines.

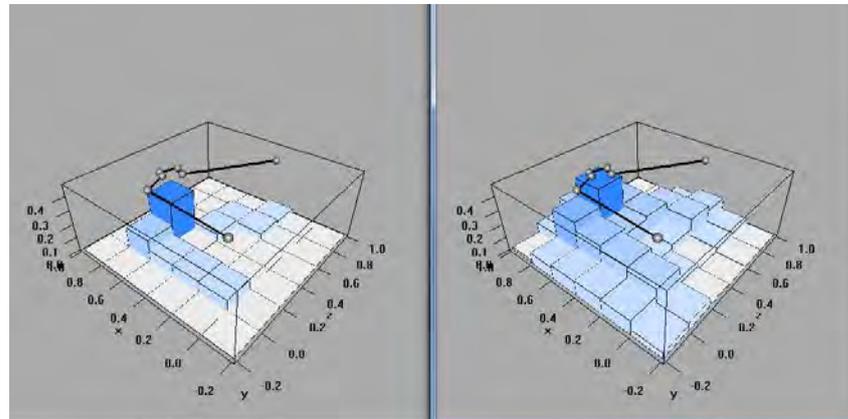


Figure 2.8: Side by side plots of line gridding and KDE gridding to a coarse grid. The input line is shown at an arbitrary height above the cells.

same coarse simple grid with an equivalent KDE grid in a perspective plot. The KDE result is still blocky, but the overall value of the surface is continuous, and not restricted to regions near the line. By increasing the resolution, the continuous nature of the KDE approach becomes more apparent—this is shown in Figure 2.9, again equivalent to the vertical view of Figure 2.7.

Kernel methods are more commonly used in terrestrial home-range applications than marine applications, perhaps as the measure of interest is a spatial region, rather than the actual trajectory of the track. Explicit line-to-cell methods have been used in marine applications because of the focus on time spent based utilization distributions. There is a convergence of aims here that is still developing, and is seeing resolution in modern statistical methods.

Kernel density methods are commonly applied to track data because of these advantages, but perhaps due to implementation difficulties with using line segments published applications use an equal-time interpolation between track points (Croxall et al., 2004).

The **trip** package provides tools to produce simple and KDE grids with line or point interpretations. The distinction here is one between discrete and continuous representations which can be represented in the data itself, or in the visualization technique or analysis used. The lack of clarity for these distinctions in spatial software is one of the problems faced in tracking research.

Other approaches have tended to use the locations more directly, by categorizing them as belonging to “focal foraging areas” (Simmons et al., 2007) or with “first passage time” analysis (Pinaud, 2007).

### Attribution of point or line values

Time duration between measured locations is a common value of interest, but many others may be used such as drift dives (e.g. Thums et al. (2008)), maximum depth, temperatures encountered, etc. The **trip** package provides tools to apply arbitrary

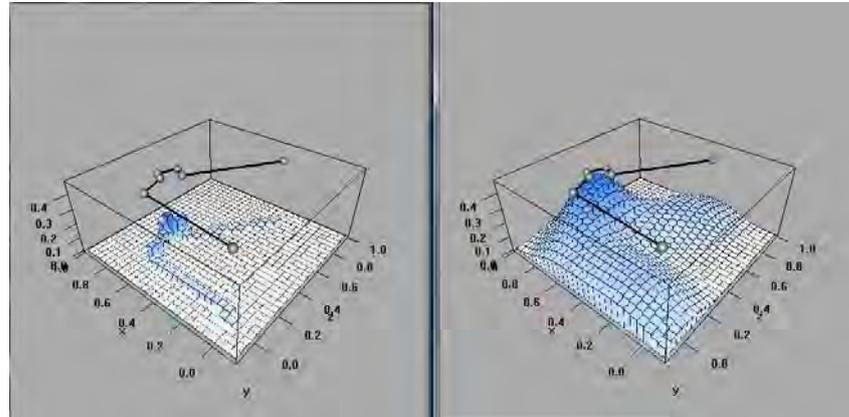


Figure 2.9: Side by side plots of line gridding and KDE gridding to a fine grid. The input line is shown at an arbitrary height above the cells.

values to gridding methods by providing conversions between point and line based interpretations and the ability to store multiple attribute values for points and lines.

### 2.3.4 Partitioning tracks and grids into time periods

Gridding methods are easily applied to entire tracks by binning into a single grid. In order to partition track data into time periods the lines (or other approximation to continuous path) must be cut into time durations. This requires that the path be cut at a point intermediate to the input points as these will rarely coincide with the period boundaries.

Once the trips are cut into exact time boundaries summaries derived from them represent some form of 3D grid in order to represent each time period separately. The track line must be cut exactly at the time boundaries, assuming constant travel speed and the contribution of each part summed separately into two grids. The **trip** package provides tools to partition sets of tracks in this way and generate collections of spatial grids. Using the simple methods above the trip data can be partitioned in time to give temporal-based utilization distributions (Keating and Cherry, 2009).

A more general method of representing arbitrary time periods from modelled tracks is discussed in Chapter 5.

### 2.3.5 Map projections

A paper published by Gudmundsson and Alerstam (1998) details the importance of the choice of map projection for analysing migration routes. Despite the wide availability of software tools for working with map projections<sup>7</sup> their use is still virtually non-existent in modern tracking studies.

<sup>7</sup>See <http://trac.osgeo.org/proj/>, Keitt et al. (2010), <http://www.eos.ubc.ca/~rich/map.html>, <http://www.manifold.net>, <http://gmt.soest.hawaii.edu/> for general projection transformations, and Hijmans et al. (2010) for working with orthodromic and loxodromic paths.

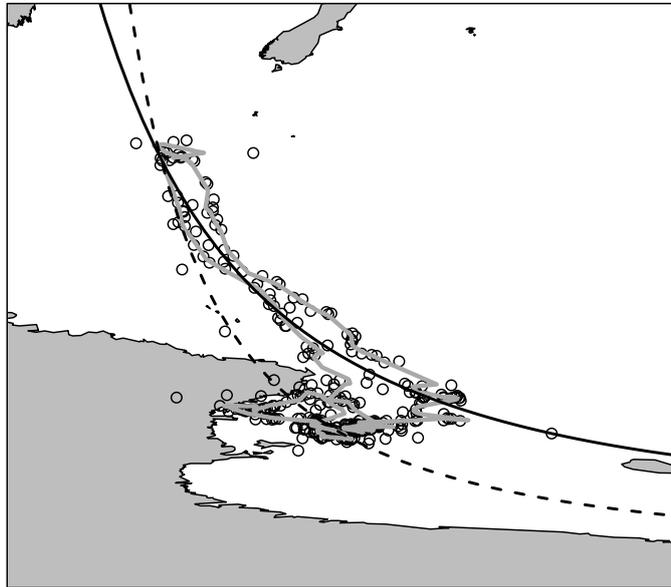


Figure 2.10: Raw Argos estimates and penalized smoothed track plotted in raw longitude and latitude. Great circles are drawn as curves, from Macquarie Island through the most distant raw Argos position (solid) and the most distant smooth position (dashed).

Distance and direction are the most obvious metrics that are calculated and many studies only use spherical approximations which are accurate except in polar regions (Banerjee, 2004). These approximations will be accurate for most purposes but visualization is still usually done using an equal-angle projection—plotting longitude and latitude directly on the x and y axes. Other studies provide ample justification and explanation for working with map projections, so only the importance of their use for the line-interpretation of a track is presented here. The **trip** package provides access to a full suite of projection transformations via the spatial support of the **rgdal** package in R (Keitt et al., 2010). A widely used transformation library with string codes for map projections, PROJ.4, is used to specify the required metadata and calculations (Evenden, 1990).

Figure 2.10 shows that the perceived error in the Argos positions in the east/west direction is far less and it is clear that intuitive ideas about straight line travel are not easily conveyed in longitude and latitude plots. The same data in Figure 2.11 uses a projection that more accurately represent great circles as straight-lines.

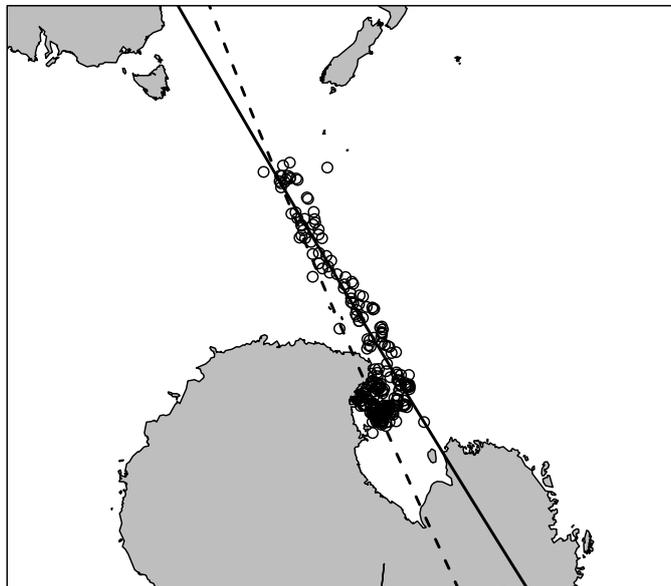


Figure 2.11: Raw Argos estimates and penalized smoothed track plotted in an equal-area projection. Great circles (as-the-crow-flies) are shown as lines. These pass through the Macquarie Island site to the most distant raw Argos position (solid) and the most distant smooth position (dashed).

### 2.3.6 Tools for traditional methods

Despite these limitations there is still a use for these filtering and gridding methods for exploratory analyses. They are also used as part of more sophisticated modelling approaches (Jonsen et al., 2005; Patterson et al., 2010). The **trip** package provides a framework for running these algorithms on sets of animal tracks within the validation requirements discussed above. As it builds on the existing spatial infrastructure the package provides simple access to arbitrary map projections and to input and output for commonly used GIS vector and raster formats. Extra capability is provided by linkages to the **raster** package to simplify access to imagery, raster data and array formats like NetCDF and to the **spatstat** package

## 2.4 The trip package for animal track data

The **trip** package developed by the author provides programming definitions and functions for working with animal track data in the R programming environment. The package provides convenient access to commonly used methods discussed in this chapter. The **trip** package is available on CRAN: <http://cran.r-project.org><sup>8</sup>.

The **trip** package builds on the spatial infrastructure of the **sp** package. The **sp** package uses the S4 programming classes and methods of Chambers (2008) to provide a coherent system for the usual spatial data types: points, lines, polygons and grids. This system provides visualization and manipulation of spatial data, a consistent conversion for data between the variety of spatial statistics packages already in R, and interfaces to GIS and map projections (Pebesma and Bivand, 2005). The following description of **sp** objects given here is derived from Pebesma and Bivand (2005) and Bivand et al. (2008), which should be consulted for a more complete description.

All **sp** (and therefore **trip**) objects share the basic property **Spatial** which stores only the bounding box and map projection metadata. Each new data type then extends this to the variety of spatial types with increasing specialization: coordinates and levels of organization are added to provide **SpatialPoints**, **SpatialLines**, **SpatialPolygons** and **SpatialGrid**. **Trip** objects are specializations of the **SpatialPoints** class and so the other types are not considered further here.

The class **SpatialPoints** consists of a matrix of coordinates and the **Spatial** metadata for their bounding box and map projection—this is sufficient to locate each point on a map, but applying more information for each point requires a matching attribute table. For example data such as the name, size, colour, direction and quality for each point may be stored as well as its location. The basic table component in R is a **data.frame**, and this is extended by the class **SpatialPointsDataFrame**. A **SpatialPointsDataFrame** can be considered as a table of records (a **data.frame**) containing X and Y coordinates and other attributes. The coordinates are treated specially so that they may be used to visualize or manipulate the object in a spatial context<sup>9</sup>. Objects of class **trip** extend this class directly, by applying another

<sup>8</sup>**Trip** replaced the experimental package **timeTrack** version 1.1-6, which provided similar functionality but was not integrated with **sp**.

<sup>9</sup>This is rather simple for points since there is only one coordinate for each record: the separation

level of information about the identity of each trip, the temporal coordinate of each point and ensuring that the records can be validated by the rules laid out in Section 2.2.2. This is done with a new class `TimeOrderedRecords` that stores the names of the date-time and ID attribute in the records. Checks are applied for any object aspiring to be of class `Spatial` and of class `TimeOrderedRecords` by ensuring that time values are in order within ID and that there are no duplicate times, no missing coordinates etc. If any of these constraints is broken then validation cannot occur and the object creation will fail. Further discussion and illustration of the `trip` class definition in the context of `sp` is given in Bivand et al. (2008).

By using the `sp` classes `trip` automatically gains all the available methods for `Spatial` objects including `plot`, `subset`, coordinate system metadata and reprojection, `summary` and `print`. A `trip` object is considered as a set of “TimeOrdered” points, rather than multi-segment lines. Some functions do assume line-interpretion of trips, but in general the storage of attribute data on ordered points provides greater flexibility than storage of line objects for reasons discussed in Section 2.2.3.

The `trip` package depends entirely on the `methods` and `sp` packages, and in part on the `spatstat` and `mapproj` packages (Baddeley and Turner, 2005; Lewin-Koh and Bivand, 2010).

The next section demonstrates some examples using the `trip` package for importing data, dealing with common problems and “filtering” and “gridding” track data.

### 2.4.1 Generating trip objects

The following examples show the simplest means of generating a trip object from a table of data. The function `readArgos` does all of this internally, as shown below. The `trip` class is derived from classes in the package `sp` and so the first step is to create a `SpatialPointsDataFrame`.

The following R code reads data from a text file, promotes the resulting table to a Spatial object and converts dates and times from text.

```
> dat <- read.csv("trackfile.csv")
> names(dat)

[1] "long" "lat" "seal" "date" "local" "lq"

> library(sp)
> coordinates(dat) <- c("long", "lat")
> dat$gmt <- as.POSIXct(strptime(paste(dat$date,
  dat$local), "%d-%b-%y %H:%M:%S"), tz = "GMT") -
  10 * 3600
```

A data frame in R is read in from text file using `read.csv`. The column names are printed by the `names` function. The `sp` function `coordinates` is used to promote of coordinates and attribute data becomes important for more complicated spatial objects.

a data frame to a `Spatial` object by specifying which columns contain the spatial coordinates.

The dates and times in this object are still just text from the file, so these are converted to the `DateTimeClasses` provided by R. An offset is applied to ensure the correct timezone interpretation. The `strptime` function provides all the required templates for parsing date-times from text—here the date stamp is a parochial format with day, short month and short year.

The examples here use longitude and latitude data, but `trip` objects can be generated using any valid coordinates. The projection metadata may be stored as per the definition of the `Spatial` classes in `sp`.

The following code attempts to generate a `trip` object by assigning the time and ID attributes. This step often presents some problems that need to be dealt with.

```
> library(trip)

> tr <- trip(dat, c("gmt", "seal"))
```

```
Error in validityMethod(object) : duplicated records within data
```

The `trip` function takes the `SpatialPointsDataFrame` and the names of the date-time and ID columns. In this case the attempt fails as the data will not validate in its current form. This is important as it ensures that these simple problems are dealt with upfront so they are not propagated further to cause problems in subsequent analyses.

The next section shows an alternative method for reading `trip` data from an Argos format.

### Reading data from Argos records

Argos (PRV/DAT) files can be read directly using the function `readArgos` and if the defaults for basic quality control are successful this will return a `trip` object.<sup>10</sup>

```
> library(trip)
> argosfiles <- list.files(path = "G:/DATA/tracks/blackBrowed/",
  pattern = ".dat", full.names = TRUE, ignore.case = TRUE)
> argosdata <- readArgos(argosfiles[1:3])
```

```
Longitudes contain values greater than 180, assuming proj.4 +over
```

```
Data fully validated: returning object of class trip
```

<sup>10</sup>These data were provided by the DPIWE Macquarie Island Albatross Project (Terauds et al., 2006). Only three of the available Argos files are imported for this example.

```
> summary(argosdata)
```

```
Object of class trip
```

	tripID ("ptt")	No.Records	startTime ("gmt")
1	14257	445	2001-12-06 01:35:31
2	14403	479	2001-12-02 04:03:06
3	14418	684	2001-12-02 05:46:51

	endTime ("gmt")	tripDuration
1	2001-12-27 04:40:19	21.12833 days
2	2001-12-18 20:16:06	16.67569 days
3	2001-12-27 06:18:30	25.02198 days

Total trip duration: 5428167 seconds (1507 hours, 2967 seconds)

Derived from Spatial data:

```
Object of class SpatialPointsDataFrame
```

```
Coordinates:
```

	min	max
longitude	147.872	189.025
latitude	-61.207	-37.800

```
Is projected: FALSE
```

```
proj4string : [+proj=longlat +ellps=WGS84 +over]
```

```
Number of points: 1608
```

```
Data attributes:
```

	prognum	ptt	nlines
Min.	:1807	Min. :14257	Min. : 2.000
1st Qu.:	:1807	1st Qu.:14257	1st Qu.: 4.000
Median	:1807	Median :14403	Median : 6.000
Mean	:1807	Mean :14369	Mean : 6.342
3rd Qu.:	:1807	3rd Qu.:14418	3rd Qu.: 8.000
Max.	:1807	Max. :14418	Max. :14.000

	nsensor	satname	class	date
Min.	:4	D:245	Z: 0	2001-12-09: 89
1st Qu.:	:4	H:316	B:234	2001-12-10: 87
Median	:4	J:310	A:202	2001-12-06: 84
Mean	:4	K:358	0:669	2001-12-11: 82
3rd Qu.:	:4	L:379	1:346	2001-12-17: 77
Max.	:4		2:135	2001-12-07: 75
			3: 22	(Other) :1114

	time	altitude	transfreq
05:44:16:	2	Min. :0	Min. :401653551
06:37:59:	2	1st Qu.:0	1st Qu.:401653710
07:28:45:	2	Median :0	Median :401653830
08:38:14:	2	Mean :0	Mean :401653849
11:34:27:	2	3rd Qu.:0	3rd Qu.:401653970
17:20:03:	2	Max. :0	Max. :401654168
(Other)	:1596		

```

      gmt
Min.   :2001-12-02 04:03:06
1st Qu.:2001-12-07 20:19:05
Median :2001-12-12 20:09:02
Mean   :2001-12-13 13:20:02
3rd Qu.:2001-12-18 13:46:50
Max.   :2001-12-27 06:18:30

```

In Argos PRV (DAT) files the fields `longitude` and `latitude` contain the spatial coordinates (these have been extracted from the other data in the `SpatialPointsDataFrame` in the usual way), `date` and `time` the temporal information (these have been combined into an R date-time column called `gmt`), and `ptt` is the ID for individual instruments that is used as the trip ID. The function `readArgos` will perform some sensible quality control corrections by default. The `summary` function returns a listing of the individual trips, their ID, start and end times, and number of locations. The remaining data are summarized in the usual way for a `SpatialPointsDataFrame`.

When reading data from PRV files the following coordinate system is assumed:

longitude / latitude on the WGS84 datum.

This is specified using the following PROJ.4 string:

```
+proj=longlat +ellps=WGS84.
```

If longitude values greater than 180 are present the “+over” element is applied for the “Pacific view [0,360]” longitude convention. No further checking is done.

The function `readDiag` reads Argos DIAG (diagnostic) format that provides two sets of location coordinates. This function returns a data frame with the attributes from the files.

### Dealing with common problems in track data

This section begins with the raw data frame of track data from 2.4.1.

```
> dat <- as.data.frame(dat)
```

There are a number of simple problems at this stage that the `trip` class automatically provides validation for. Duplicated rows, such as those from overlapping Argos files, can be safely dropped.

```

> dat <- dat[!duplicated(dat), ]
> head(dat)
      long      lat seal      date      local lq
1 158.9467 -54.49333 b284 21-Jan-99 13:43:51  A
2 155.2533 -54.42833 b284 21-Jan-99 15:18:53  B

```

```

3 155.8400 -56.30833 b284 23-Jan-99 16:45:51 B
4 159.3850 -56.10167 b284 23-Jan-99 13:03:35 B
5 159.9650 -58.35333 b284 25-Jan-99 16:19:17 B
6 158.8917 -57.99000 b284 25-Jan-99  4:45:15 B
      gm
1 1999-01-21 03:43:51
2 1999-01-21 05:18:53
3 1999-01-23 06:45:51
4 1999-01-23 03:03:35
5 1999-01-25 06:19:17
6 1999-01-24 18:45:15

```

The discarded rows removed are exact duplicate rows, matched on every data value for each column.

The ordering of rows in the data is assumed to follow the order of the date-time values with a trip ID. The date-time values are not used automatically to order the data, as this could hide deeper problems in a data set.

```
> dat <- dat[order(dat$seal, dat$gmt), ]
```

A final problem is that subsequent date-times within a single trip ID can be duplicates. Removing duplicate rows and ordering the rows still leaves the problem of what that can mean. If the location coordinates are different, the implication is that the animal moved a certain distance in no time at all. If the locations are the same then it raises the question of why the tracking system distinguishes the records at all. A zero duration time difference results in meaningless metrics of speed of movement, for example. There is no obvious solution for this and the issue must be investigated in the context of the actual data used. A simplistic solution that allows us to move on is to adjust these duplicate times by a very small amount, so that the time difference is not zero.

```
> dat$gmt <- adjust.duplicateTimes(dat$gmt, dat$seal)
```

Further data problems are more fundamental and do not have easy fixes. The `trip` class will fail to validate in the following cases and these must be dealt with individually as appropriate.

**Insufficient records for a given trip ID** Each set of records must have three or more locations to qualify as a `trip`. Sometimes the only available data is a start and end location but this case is deemed inappropriate for `trip`.

**Invalid values for critical data** Missing values or otherwise non-numeric values for locations and date-times cannot be included. The `sp` classes ensure this for location coordinates, and `trip` adds the limitation for date-times and IDs.

**Non-existent date and ID data** This is somewhat obvious, but `trip` will not assume a simple date-time value from the order of records or provide a default ID for single-trip data. These must be explicitly provided.

Once all of the fixes required have been applied, `trip` validation is possible.

```
> coordinates(dat) <- c("long", "lat")
> tr <- trip(dat, c("gmt", "seal"))
```

The location quality class is converted to an ordered factor, and the appropriate PROJ.4 string is applied.

```
> tr$class <- ordered(tr$lq, c("Z", "B", "A", "0",
  "1", "2", "3"))
> proj4string(tr) <- CRS("+proj=longlat +ellps=WGS84 +over")
```

The location quality class provided for Argos data is automatically converted to an “ordered factor” by `readArgos`, here shown manually. This can be used for simple selection of a range of records from a data set, even though the tokens used as text have no inherent order. This allows the data to be used directly without creating a numeric proxy for the class values.

The resulting `trip` object read from CSV text is now validated and equivalent to that returned by `readArgos`.

### 2.4.2 Filtering for unlikely movement

The infrastructure provided by the `trip` classes allows efficient implementation of custom filters based on a variety of metrics.

The data are from southern elephant seals from Macquarie Island. These animals can swim up to 12 km/hr (Bradshaw et al., 2002) so that value is used to calculate a filter, which is added as a column in the data frame. The filtering algorithm is that of McConnell et al. (1992).

```
> tr$ok <- speedfilter(tr, max.speed = 12)
> summary(tr)
```

```
Object of class trip
  tripID ("seal") No.Records  startTime ("gmt")
1          b284         281 1999-01-21 03:43:51
2          b290         255 1999-01-29 19:30:59
3          c026         351 1999-01-01 02:27:28
4          c993         258 1999-01-07 02:59:53
      endTime ("gmt")  tripDuration
1 1999-06-09 04:34:22 139.0351 days
2 1999-07-02 22:52:25 154.1399 days
3 1999-05-20 18:51:27 139.6833 days
4 1999-04-24 03:04:44 107.0034 days
```

```
Total trip duration: 46644047 seconds (12956 hours, 2447 seconds)
```

Derived from Spatial data:

Object of class SpatialPointsDataFrame

Coordinates:

```

          min      max
long 128.27167 198.99167
lat  -78.06333 -52.85667

```

Is projected: FALSE

proj4string : [+proj=longlat +ellps=WGS84 +over]

Number of points: 1145

Data attributes:

seal	date	local	lq
b284:281	5-Feb-99 : 20	15:29:50: 3	0:114
b290:255	10-Mar-99: 17	14:42:21: 2	1: 31
c026:351	11-Feb-99: 17	14:48:07: 2	2: 18
c993:258	21-Mar-99: 17	16:08:14: 2	3: 1
	23-Feb-99: 16	16:18:30: 2	A:281
	10-Apr-99: 15	16:39:55: 2	B:700
	(Other) :1043	(Other) :1132	

gmt	class	ok
Min. :1999-01-01 02:27:28	Z: 0	Mode :logical
1st Qu.:1999-02-11 06:03:36	B:700	FALSE:224
Median :1999-03-10 19:43:05	A:281	TRUE :921
Mean :1999-03-15 13:10:03	0:114	NA's :0
3rd Qu.:1999-04-10 06:18:01	1: 31	
Max. :1999-07-02 22:52:25	2: 18	
	3: 1	

The `speedfilter` function assumes that speed is specified in km/hr and distance is calculated based on the projection metadata. If the projection is specified and not longitude and latitude then Euclidean methods are applied. If the projection is not specified longitude and latitude is assumed and ellipsoid methods for WGS84 are used.

The summary shows that a number of locations are now classified by a boolean value in a new “ok” column. Although the speed filter has not removed many locations, a customized subset can be defined based on other data. Using a minimum Argos location quality class of “A” the raw data is plotted with a default point symbol and reduced size, with lines connecting the remaining filtered points. The plot is shown in figure 2.12.

If this were calculated with coordinates in a different unit or projection the `speedfilter` function takes this into account and calculates distance accordingly. This integration of spatial metadata and calculation provides for a very flexible working environment for exploratory analysis.

```
> plot(tr, axes = TRUE, cex = 0.4)
> plot(world, add = TRUE, col = "grey")
> lines(tr[tr$ok & tr$class > "B", ], lwd = 2, col = bpy.colors()[seq(10,
  90, length = 4)])
```

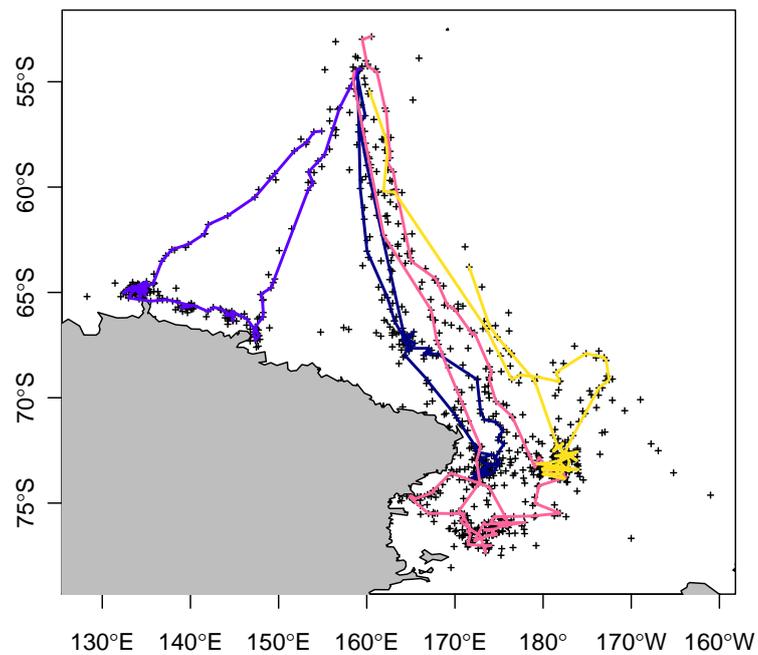


Figure 2.12: Plot of Argos track data from a `trip` as points, with lines showing the speed filtered tracks, coloured for each separate trip.

### Extending the trip package

One of the advantages of the infrastructure provided by the **trip** classes is that simple tasks can be strung together efficiently. This example shows how the speed-distance-angle filter of Freitas et al. (2008) can be built up from the basic tool kit.

The speed-distance-angle filter is distributed with an easily run example, and the following is reproduced almost exactly from the documentation for **argosfilter** in Freitas (2010).

```
> library(argosfilter)
> data(seal)
> lat <- seal$lat
> lon <- seal$lon
> dtime <- seal$dtime
> lc <- seal$lc
> cfilter <- sdafilter(lat, lon, dtime, lc)
> seal$sda <- !(cfilter == "removed")
```

This filter results in a column of values that specify whether the filter retains or discards the location.

The following steps create a new version of this filter using the **trip** package. Load the example data and validate as a single event trip object.

```
> library(argosfilter)
> library(sp)
> library(trip)
> library(maptools)
> trackAngle <- function(xy) {
  angles <- abs(c(trackAzimuth(xy), 0) - c(0,
    rev(trackAzimuth(xy[nrow(xy):1, ]))))
  angles <- ifelse(angles > 180, 360 - angles,
    angles)
  angles[is.na(angles)] <- 180
  angles
}
> vmax <- 2
> ang <- c(15, 25)
> distlim <- c(2500, 5000)
> coordinates(seal) <- ~lon + lat
> proj4string(seal) <- CRS("+proj=longlat +ellps=WGS84")
> seal$id <- "seal"
```

Perform some simple sanity checks before proceeding.

```
> range(diff(seal$dtime))

Time differences in secs
[1] 0 362737
```

```
> which(duplicated(seal$dttime))

[1] 18 117 123 1009 1159 1232 1294 1301
```

There are some preliminary requirements to load packages and functions for earth distances and track turning angles. The `seal` object is a data frame with columns “lon”, “lat”, “dttime” and “lc”. The date-time values are already in POSIXct format, and lc uses a numeric code for the Argos class.

The records are in order of date-time, but some of the date-time values are duplicated, so for illustration these are dropped and then a new `trip` object is created.

```
> seal <- seal[!duplicated(seal$dttime), ]
> seal.tr <- trip(seal, c("dttime", "id"))
```

First, add the filter from the `sdafilter` function to the trip object, and also the basic speed filter available in the `trip` package (using km/hr rather than m/s).

```
> seal.tr$speed.ok <- speedfilter(seal.tr, max.speed = vmax *
  3.6)
```

Perform the initial simple processing for the speed distance angle filter.

```
> dsts <- trackDistance(coordinates(seal.tr)) *
  1000
> angs <- trackAngle(coordinates(seal.tr))
> dprev <- c(0, dsts)
> dnext <- c(dsts, 0)
> ok <- (seal.tr$speed.ok | dprev <= 5000) & (seal.tr$lc >
  -9)
```

This creates a logical vector where the speedfilter and distance-previous pass, and discard any location class of “Z” (coded as “-9” in this example). This is effectively the first pass of the filter, distinguishing distance previous and distance next for each point.

Now the remaining parts of the filter can be run—testing for angle and distance combinations over the specified limit. Some housekeeping—create an index to keep track while running the filter on the two distance/angle limits. A temporary copy of the `trip` is used to make matching the filter easy as points are removed.

```
> seal.tr$filt.row <- 1:nrow(seal.tr)
> seal.tr$ok <- rep(FALSE, nrow(seal.tr))
> df <- seal.tr
> df <- df[ok, ]
> for (i in 1:length(distlim)) {
  dsts <- trackDistance(coordinates(df)) * 1000
```

```

angs <- trackAngle(coordinates(df))
dprev <- c(0, dsts)
dnext <- c(dsts, 0)
ok <- (dprev <= distlim[i] | dnext <= distlim[i]) |
      angs > ang[i]
ok[c(1:2, (length(ok) - 1):length(ok))] <- TRUE
df <- df[ok, ]
ok <- rep(TRUE, nrow(df))
}

```

The result is now a reduced `trip` with missing rows discarded by the filter. Using the row index created earlier, match the filter result to the original rows and tabulate the filter values for the original `sdafilter` and the `trip` version. The number of accepted points is nearly the same.

```

> seal.tr$ok[match(df$filt.row, seal.tr$filt.row)] <- ok
> sum(seal.tr$sda)

[1] 1135

> table(seal.tr$sda, seal.tr$lc)
      -9  -2  -1   0   1   2   3
FALSE 26 302  48  17  17   6   2
TRUE   0 374 374  61 150 116  60

> sum(seal.tr$ok)

[1] 1142

> table(seal.tr$ok, seal.tr$lc)
      -9  -2  -1   0   1   2   3
FALSE 26 299  49  15  16   5   1
TRUE   0 377 373  63 151 117  61

```

Plot the result to see that the new `trip` version is comparable. There are differences since the distance and angle calculations are ellipsoid based, rather than spherical as used by `argosfilter` and some locations were first discarded due to impossible duplicate times in the original data. There are probably also differences in the detail of the recursive speed filter and the way that peaks are assessed.

The examples above consist of a working prototype that can be wrapped as a function for the `trip` package to efficiently apply this filter to sets of tracks within a `trip` object. The `argosfilter` example above takes at least three times as long to complete as the code developed here. The ability to extend the functionality in this way mirrors the development of the `trip` package itself, from the `sp` package, and in turn the R platform.

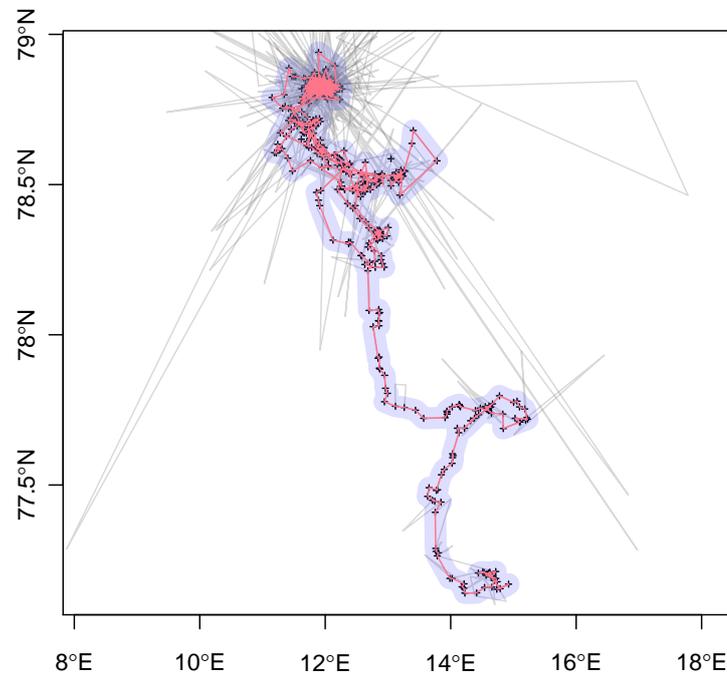


Figure 2.13: Original `sdafilter` (red) and custom `trip` speed distance angle filter (thick grey line). The original raw track is shown as a thin grey line.

```
> trg <- tripGrid(tr[tr$ok, ])
> image(trg, col = oc.colors(100), axes = TRUE)
```

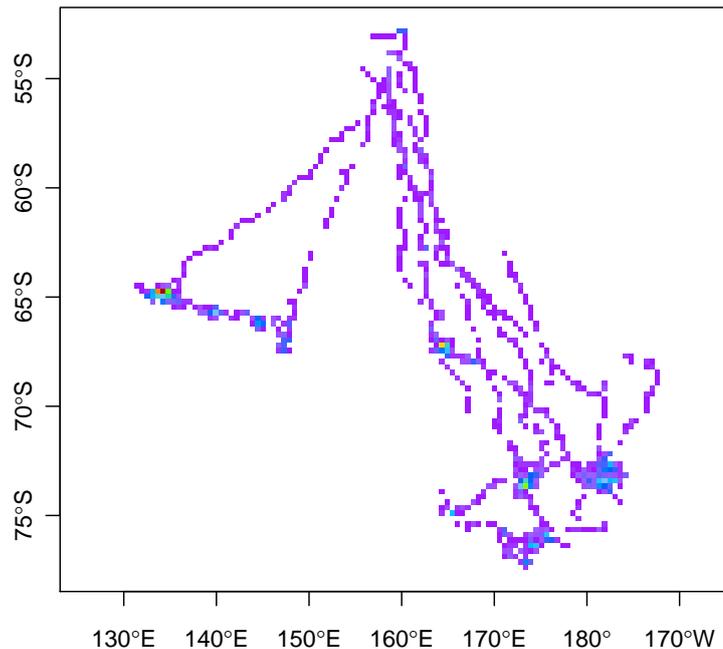


Figure 2.14: Simple image of a time spent grid for the trip object.

### 2.4.3 Creating maps of time spent

Assuming that the filtered locations give realistic information about position for the animal and that motion between these positions is constant and straight, a map of time spent, or residency can be created. The choice of grid cell size might reflect the confidence in the accuracy of the location data, or require a specific cell size for comparison with another data set.

Using the trip locations accepted by the speed filter attribute generate a grid of time spent, shown in Figure 2.14.

The function `tripGrid` is used with the subset of the trip object accepted by the speed filter to create a grid of time spent. The algorithm used cuts the line segments exactly as described in section 2.3.3 by the `spatstat` function `pixellate`. If any zero-length line segments are present a warning is issued and their time contribution is included as points. This is because the `pixellate` algorithm relies on weightings that are relative to the line length, which is another example of the subtle implications

of point versus line interpretations.

The gridded object `trg` is a `SpatialGridDataFrame`, whose class definition is provided by the `sp` package. These grids support multiple attributes and so, conveniently, variations on the map created can be stored in a single object.

Add three new attributes to the grid object by isolating a single animal's trip, and using the kernel density method with two different sigma values.

```
> names(trg) <- "grid.filtered"
> gt <- getGridTopology(trg)
> trg$grid.c026 <- tripGrid(tr[tr$ok & tr$seal ==
  "c026", ], grid = gt)$z
> trg$kde1.filtered <- tripGrid(tr[tr$ok, ], grid = gt,
  method = "density", sigma = 0.1)$z
> trg$kde3.filtered <- tripGrid(tr[tr$ok, ], grid = gt,
  method = "density", sigma = 0.3)$z
> for (col.name in names(trg)) trg[[col.name]] <- trg[[col.name]]/3600
```

The default name for an attribute from `tripGrid` is “z” so first this is renamed to “grid.filtered”. Calculation for the extra attributes requires that they share the same origin and scale so this is stored in the `GridTopology` object `gt` and used again. The `trip` object is subset on the speed filter attribute and the seal ID, and the resulting grid attribute “z” is extracted and assigned to the grid object in one step. For the kernel density versions two values of sigma are passed onto the `density` function for each grid. The default output is in seconds, so this is converted to hours for each grid column by division. Finally, the multi-panel `spplot` function in package `sp` provides a conveniently scaled image for each of the four grids shown in Figure 2.15.

By default, `tripGrid` will provide a grid with dimensions 100x100 cells. This can be controlled exactly using a `GridTopology` passed in as the grid argument to the function. The convenience function `makeGridTopology` allows the user to define a specific grid from the trip object itself.

The next examples use `trip` object to create grids with a different scale.

The first example shows the creation of a grid topology with dimensions of 50x50, then another is created using a given cell size. The resulting plot from this coarser version in an equal area map projection is shown in Figure 2.16. The grid generation will assume kilometres for cell size as at the centre of the grid for longitude and latitude coordinates.

For approximate methods `tripGrid.interp` will interpolate between positions based on a specified time duration. A shorter period will result in a closer approximation to the total time spent, but will take longer to complete. The approximate method is similar to that published by Croxall et al. (2004) and Bradshaw et al. (2002).

```
> require(lattice)
> library(trip)
> trellis.par.set("regions", list(col = oc.colors(256)))
> print(spplot(trg))
```

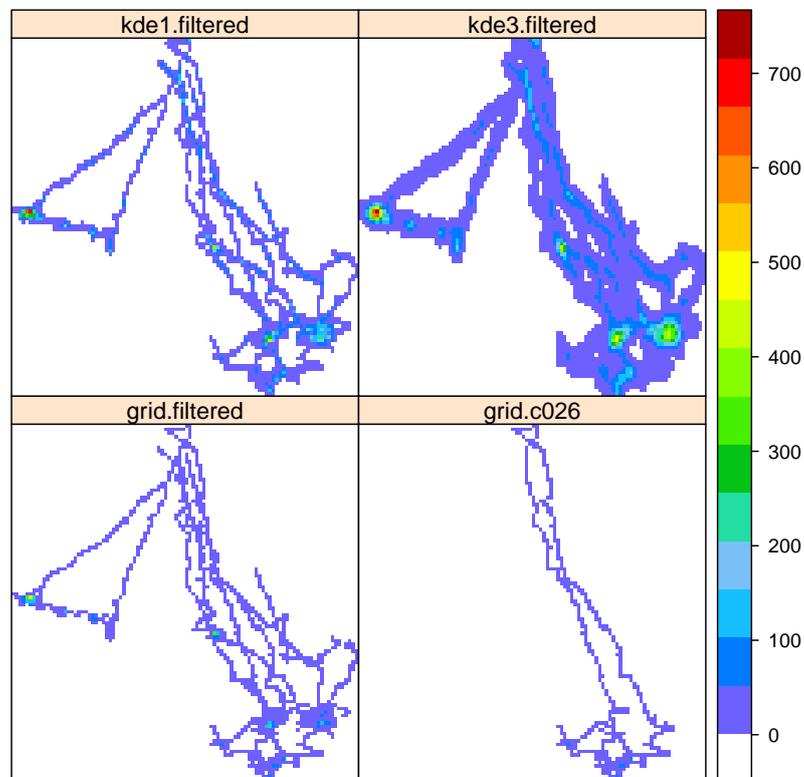


Figure 2.15: Four variations on a tripGrid. `grid.filtered` and `grid.c026` use the simple line-in-cell method for all four seals and for seal c026 alone. `kde3.filtered` and `kde1.filtered` use the kernel density method for line segments with a sigma of 0.3 and 0.1 respectively. Each grid has dimensions 100x100 and time spent is presented in hours.

```
> proj4string(tr) <- CRS("+proj=longlat +ellps=WGS84")
> p4 <- CRS("+proj=laea +lon_0=174.5 +lat_0=-65.5 +units=km")
> ptr <- tripTransform(tr, p4)
> gt <- makeGridTopology(ptr, c(50, 50))
> gt1 <- makeGridTopology(ptr, cellsize = c(80,
  60))
> grd2 <- tripGrid(ptr, grid = gt1)
> image(grd2, col = oc.colors(256), axes = TRUE)
> library(maptools)
> data(wrld_simpl)
> plot(spTransform(wrld_simpl, p4), add = TRUE,
  col = "grey")
```

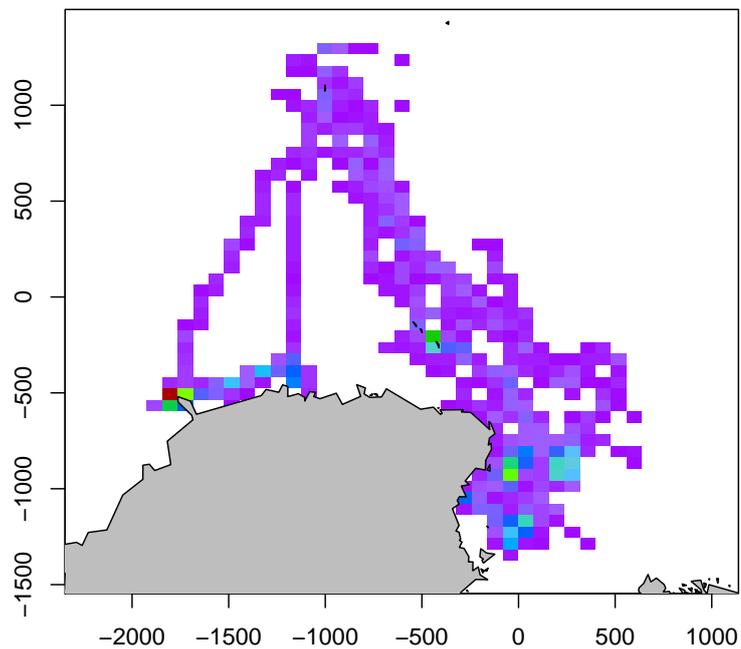


Figure 2.16: A relatively coarse version of `fourgrids` with a specific kilometre-based cell size (80x60km), specified within an equal area map projection.

## 2.5 The need for a more general framework

Previous sections presented tools for a more systematic handling of animal track data and a variety of methods for improving track estimates and deriving spatial summaries such as time spent. Some of the problems with animal tracking are relatively simple and have reasonable solutions. The traditional methods above illustrate ways of dealing with them in an extensible software toolkit. Most of these solutions however are just “single-targets”—the easiest aspects are cherry-picked for a first-pass answer that solves a small aspect of the larger problem.

Location estimates from methods such as archival tagging and acoustic tagging can also be dealt with in this way, and there are many studies that apply these techniques as well as more sophisticated models. There is an important opportunity here though since the normal “data product” for archival tags is not location estimates over time but dense temporal records of environmental values, such as light, temperature and depth.

Importantly, different satellite methods such as the Argos Service and GPS will be handled uniformly by a general approach. These methods ultimately rely on data as raw as archival tags and acoustic arrays, with doppler shifts or ephemeris data used to determine location. The practical difference of these methods from those of archival tags is that the raw data are simply not available for research purposes, for a variety of reasons.

Another great opportunity presented by raw data is that the concept of integration of all data sources comes very naturally. For example, determining position by light level is plagued by environmental conditions that attenuate the ambient light, and the movement of animals complicates the relation of environmental measurements to independent data sets. The scale of measurement is another issue when relating values such as dive depth or water temperature to synoptic data sets.

To turn attention to “raw data” methods such as those required for archival tags, the aim is to provide a more complete solution that integrates solar information, environmental values such as temperature and depth and applies constraints on movement. This approach contrasts with the filters in this chapter by applying as much information to each location estimate as possible. Also it aims to prevent the practice of discarding “bad” data—the influence of unreliable data should be downplayed but not ignored completely.

There are of course many existing location estimates for which there is no rawer data. From this perspective an Argos location can be treated as a mere data point—not an absolute position to be retained or discarded and then smoothed by some mechanism, but simply a piece of data to help inform our models. Even completely invalid positions have a date-time value and so at the very least it can be inferred that the animal was “visible” to the remote sensing system. In conjunction with other data about the environment this tiny piece of data can be valuable.

The following describes the problems specific to the two types of tags that originally motivated this work. By considering the location estimates that come from archival tag methods we see that dealing with the points without reference to the other available information is not going to be good enough.

### 2.5.1 Location estimation from raw archival data

Earlier methods of light level geo-location rely on the determination of critical times during the day from the sequence of raw light levels. Longitude is determined directly from the time of local noon which, based on the symmetry between dawn and dusk, can be easily measured. Latitude is determined by the length of the day, measured by choosing representative times such as dawn or dusk that are distinctive in the light record. While this is a very simplified explanation given the range of traditional methods the thrust of the argument is basically correct, see Chapters 3 and 4 and Metcalfe (2001) for more detail.

In effect, this approach reduces the data set of light values to a more abstract summary with local peaks at noon and inflection points at dawn and dusk. The majority of the light data is not used directly and only one location can be determined per day. This approach to estimation is susceptible to the movement of the animal between twilights, to poor choices for the critical times, and to day length at equinox periods. Hill and Braun (2001), Ekstrom (2004) and Welch and Eveson (1999) review these methods and provide quantifications for their accuracy.

There is an opportunity for delving more deeply into the raw data available for location estimation provided by archival tags. Considering the problems for quality testing of derived locations, studies can utilize the raw data from the archival tag and provide an estimation that takes into account more information such as the conditions at twilight, diving depth, water temperature, and movement in an integrated way. The raw data is not just a single point estimate to be discarded or corrected, but a sequence of light, a sequence of temperature and a sequence of depths. These are all linked temporally within the main data set. In theory, this is no different for satellite tags except in that case the raw data are not available. Raw data in this case would be the doppler shift signals and satellite orbital details for Argos, and satellite ephemeris and signal timings for GPS. Acoustic tagging applications similarly have a wealth of raw data for informing locations, and potential for improving on existing techniques.

Figure 2.17 presents another example of the “obvious” problem, with archival tracks that are very erratic and have some estimates well inland. The sorts of inaccuracies shown tend to be easily understood by researchers, but the public perception of research, so important to biological programs, can be easily undermined by figures like this. This figure also highlights the need for uncertainty in estimates to be integrated and represented as part of track visualizations and other summaries.

### 2.5.2 Filtering locations

Destructive filters that remove “erroneous” locations are susceptible to the following problems. Data is lost, and this may be significant for a diving animal or other situation where the opportunity for obtaining fixes is rare as was shown earlier in Section 2.3.1. If nothing else, an “erroneous” location at least has useful information about the time at which a fix was available. In practice techniques tend to smooth out some filter metric across multiple locations, with no clear reason to choose one technique over another. Speed, distance and direction all require at least two

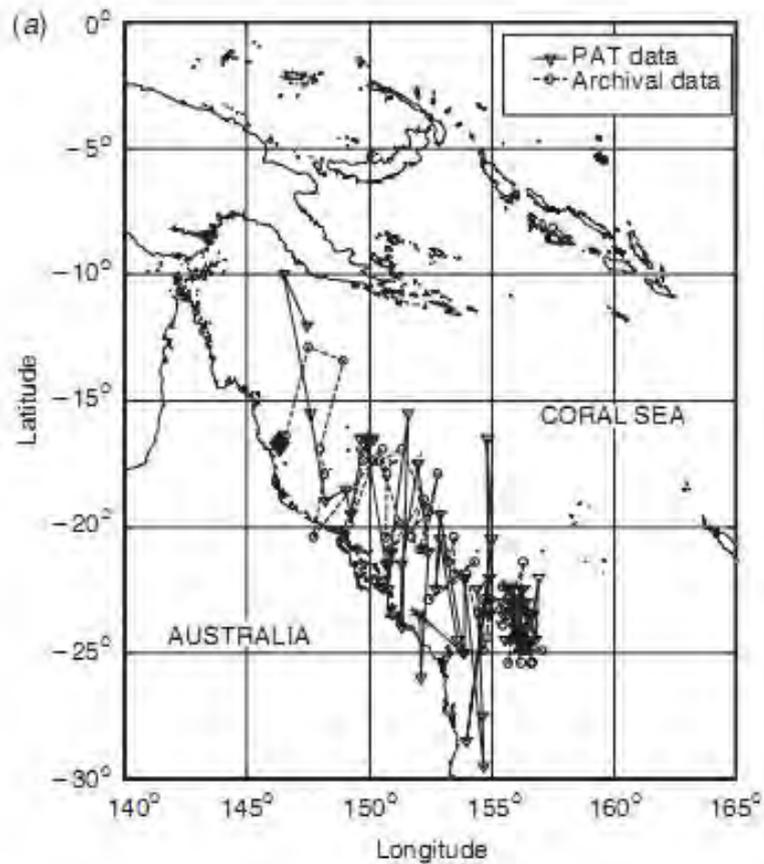


Figure 2.17: Light level geo-locations for black marlin in the Coral Sea. This image is taken from Figure 3a in Gunn et al. (2003) and is reproduced with the permission of the authors.

locations for their calculation, and so this forces the consideration of a point or line interpretation, or perhaps a hybrid of the two. The removal of a data point changes the relevant filter decision for its neighbours and so begs the question of whether the first location removed is not more valid, and perhaps better than others that should be removed in preference. An example of this can be seen with Argos diagnostic (DIAG) data in which there are paired solutions for each locations (Service Argos, 2004). Usually the choice is obvious, but for many positions it depends on the choice made for neighbouring locations. Updating a location to an estimate nearer its neighbours based on some constraint may be more sensible as in Section 2.3.2, but there is still only a point estimate. Preferably there should be a probability distribution for the location, something that gives an indication of “how correct”.

### 2.5.3 Intermediate locations

There is a need for modelling the location of the animal at times that are not accompanied by data. As discussed in Section 2.3.3 the goals of cell gridding and density analysis attempt to integrate this with location uncertainty, but there are limits to the use of tracks represented by simple points and lines. In Chapter 3 an approach that distinguishes “primary” locations from “intermediate” locations is presented. The importance of this for track representations is illustrated in greater detail in Chapter 5.

## 2.6 Conclusion

This chapter has presented traditional solutions to many of the issues faced by tracking analyses with a readily available and extensible software toolkit. The **trip** package developed by the author provides an integrated environment for applying traditional algorithms in the context of widely used spatial data tools. Seemingly isolated problems with track data were shown to be inter-related in complicated ways that preclude an otherwise simple chaining together of individual solutions. The scope of these problems can be extended to include archival tag data and raw methods like light level geo-location. This highlights the need for treating any data, even ostensibly accurate location estimates, as only part of a larger suite of available information. The next chapter provides an integrated statistical approach to modelling location using these disparate types of tag data.

## CHAPTER 3

# BAYESIAN ESTIMATION OF ANIMAL MOVEMENT FROM ARCHIVAL AND SATELLITE TAGS

### Preface

Many of the issues highlighted in the the previous chapter motivated the following work. The combination of requirements for satellite-linked and light-measuring archival tags led to the Bayesian framework presented here which was published as Sumner et al. (2009).<sup>1</sup> That document is reproduced here from the original versions with only minor modifications, including a fix to Table 3.2, the addition of this preface and the addition of a conclusion.

Originally, the goal of this work was to develop an improved approach to archival tag location estimation. However, consideration of the practical needs for correcting Argos data provided insights into the shared requirements for the modelling process of both methods. The early estimation code for light level geo-location focussed on the solar calculations for primary estimates with a speed constraint placed directly on the twilight locations, and the first attempts at movement modelling for Argos estimates resulted in a distinction between the primary locations and their intermediate (latent) intervals. Merging these goals in a single approach required a model of the data collection process and identification of the appropriate data sources to inform models of the true locations. This chapter presents a general Bayesian approach that can be applied to any tracking location method.

### Abstract

The reliable estimation of animal location, and its associated error is fundamental to animal ecology. There are many existing techniques for handling location error, but these are often *ad hoc* or are used in isolation from each other. In this study we present a Bayesian framework for determining location that uses all the

---

<sup>1</sup>Available in original form at <http://dx.doi.org/10.1371/journal.pone.0007324>.

data available, is flexible to all tagging techniques, and provides location estimates with built-in measures of uncertainty. Bayesian methods allow the contributions of multiple data sources to be decomposed into manageable components.

We illustrate with two examples for two different location methods: satellite tracking and light level geo-location. We show that many of the problems with uncertainty involved are reduced and quantified by our approach.

This approach can use any available information, such as existing knowledge of the animal's potential range, light levels or direct location estimates, auxiliary data, and movement models. The approach provides a substantial contribution to the handling of uncertainty in archival tag and satellite tracking data using readily available tools.

### 3.1 Introduction

Estimating the movements of animals is a fundamental requirement for many ecological questions. These include elucidating migratory patterns, quantifying behaviour in terms of the physical environment and understanding the determinants of foraging success, all of which can influence larger population processes (Nel et al., 2002; Xavier et al., 2004; Hindell et al., 2003). Types of movement data can range from simple mapping of positions to behavioural models that attempt to account for unlikely estimates, provide estimates of behavioural states and predict latent variables.

There are two common methods for obtaining position estimates, which can be broadly categorized as remote and archival. Remote methods use techniques such as radio or satellite telemetry to locate a tag attached to an animal. Archival methods require the tag to record aspects of the animal's environment over time (such as light levels and water temperature) which are then processed to infer location (Smith and Goodman, 1986; Hill, 1994; Nel et al., 2002).

Before any analysis can be done, position estimates require some quantification of precision and accuracy to provide statistical confidence in results (Hays et al., 2001; White and Sjöberg, 2002; Phillips et al., 2004). Quantification of location precision, and crucially, the incorporation of these into synoptic spatial representations of animal movement, is an important problem common to both methods that many authors have attempted to address in recent studies (Matthiopoulos et al., 2005; Royer et al., 2005; Jonsen et al., 2005; Teo et al., 2004; Nielsen et al., 2006; Bradshaw et al., 2002).

Location precision is generally lower in archival methods due both to the theoretical basis and practical problems of the location estimation (Hill and Braun, 2001; Ekstrom, 2002). To overcome this limitation, archival methods routinely integrate primary location estimation with auxiliary data sets (DeLong et al., 1992; Smith and Goodman, 1986; Teo et al., 2004; Domeier et al., 2005). In principle this enables the integration of the estimation and error estimation processes but this remains an under-utilized opportunity: published uses of archival methods usually separate the estimation of the quality of position estimates from their derivation. Satellite-derived estimates provide less opportunity in this regard, as the process is

proprietary and information regarding error is minimal. However, satellite locations still require a modelling framework to incorporate auxiliary information and provide the best possible estimates (Jonsen et al., 2005) including a quantification of precision.

The simplest analysis of movement data is to visualize the sequence of locations visited by the animal. It is slightly more complex to provide a path estimate of the animal, which requires the ability to determine position both from available data as well as for latent times where no data were measured. An obvious simple model is to “join the dots”, assuming that movement is both linear and regular between measured positions. A more realistic approach demands that estimates of an animal’s path consider both primary and latent location estimates, because in general there are open-ended scenarios that could occur between primary estimates. There are a multitude of methods for achieving this (Turchin, 1998; Wentz et al., 2003; Bradshaw et al., 2002; Tremblay et al., 2006; Ovaskainen et al., 2008), but none have been directly integrated with the estimation process from raw data.

Once an estimate of an animal’s path is obtained biologists often need to calculate speed of and distance of travel, generate spatial representations of an animal’s use of space in terms of time spent in geographic regions, metabolic effort or other measure of resource allocation. More sophisticated analyses aim to determine behavioural states more exactly (Matthiopoulos et al., 2004; Jonsen et al., 2005), or to differentiate migration from foraging behaviour. These aims are beyond the present work, where we will be focusing on the first step in the process—description of an animal’s path and the precision with which this can be estimated.

Earlier work has attempted to account for spatial uncertainty by choosing a scale for interpreting location data (Bradshaw et al., 2002), or spatial smoothing (Wood et al., 2000). These techniques fail to estimate statistical uncertainty for individual estimates, and provide only an overall average of precision. Other techniques are used to estimate latent position by interpolation or similar technique (Tremblay et al., 2006), but these must assume that positions are known.

Given the diversity of questions asked of movement data, there are understandably many approaches to data analysis. Many existing techniques are specific to particular questions and species and have little scope outside the given application. Further, each application has its own problems of scale, location error, data quality and summarizing of behaviour. In this context, sophisticated model approaches are seeing greater use in tracking studies (Nielsen et al., 2006), but these have only been applied to pre-derived positions and leave the problem of location estimation from raw data unaddressed. No study has yet provided a general approach to dealing with the twin issues of estimate precision and accuracy for both archival and satellite location data. There is a growing need for just such an approach as more large multi-species studies are being undertaken (Croxall et al., 2004; Block et al., 2003b; Halpin et al., 2006). Such multi-species studies inevitably utilize a range of tracking techniques as no one method is suitable for all species. For example, fish which rarely come to the surface are not usually suitable for satellite tracking (Hunter et al., 1986).

Here we present a Bayesian framework for the analysis of movement data that directly addresses the estimation of location from raw data collected by archival

tags and can also be applied to other data sets of pre-derived position estimates such as Argos locations. We apply the approach to both an archival tag data set and a satellite tag data set. Our primary goal is to integrate all available sources of information for estimating location. Using all available information may sound obvious, but it is a missed feature of many applications. Secondly, we aim to integrate the location estimation and the estimation of location precision. The approach should also be able to provide all of the desired end-uses of tracking data as mentioned above. In the Bayesian context, each of these measures, including appropriate confidence intervals (CI) (Gelman et al., 2004; Gilks et al., 1995), can be determined by specifying appropriate priors and distributions for each data source and calculating the posterior.

## 3.2 Materials and Methods

### 3.2.1 Ethics Statement

Data were collected under permits from the University of Tasmania Animal Ethics Committee (A6790 and A6711).

### 3.2.2 Assumptions

We propose a Bayesian approach to the tag location problem that uses Markov Chain Monte Carlo methods to approximate the posterior.

There are three main elements to the process of Bayesian estimation; the prior, the likelihood and the posterior. The *prior* distribution  $p(\theta)$  represents our knowledge of the parameters  $\theta$  before any data is observed. The *likelihood*  $p(y|\theta)$  is the probability of observing data  $y$  for a given set of parameters  $\theta$ , and represents our knowledge of the data collection process. From these we calculate the posterior distribution  $p(\theta|y)$  via Bayes' rule

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{\int p(y|\theta)p(\theta) d\theta}. \quad (3.1)$$

The posterior  $p(\theta|y)$  represents our knowledge of the parameters after the data  $y$  have been observed. In essence, Bayes' rule provides a consistent mechanism for updating our knowledge based on observed data.

The data available for forming location estimates can be classified into four broad types.

**Prior knowledge of the animal's movements** Invariably something is known of an animal's home range, migratory pattern or habitat preference, and any location estimate should be consistent with this information. This information can range from being quite specific such as the species generally stays over the continental shelf (e.g. shy albatross (Brothers et al., 1997)) or more vague such as the species often heads south (e.g. southern elephant seals (Bradshaw et al., 2002)).

**Primary location data** The primary location data  $y$  is data collected primarily for the purposes of location estimation, and directly inform about the locations  $x = \{x_1, x_2, \dots, x_n\}$  of the tag at a sequence of (possibly irregular) times  $t = \{t_1, t_2, \dots, t_n\}$ . Examples include the light levels recorded by an archival tag, or for an Argos tag the locations provided by the Argos service.

**Auxiliary environmental data** Many tags also record additional environmental data  $q$ , and this data may be compared to external databases to further constrain location estimates (Smith and Goodman, 1986; Delong et al., 1992; Teo et al., 2004; Domeier et al., 2005; Nielsen et al., 2006). For example, in the marine context depth and temperature measurements can be compared to remotely sensed or modelled sea surface temperature (SST) data to confine locations to regions where SST is consistent with the temperatures observed by the tag.

**Movement models** Movement models constrain the trajectory of the animal, reducing or removing the occurrence of location estimates that correspond to improbable or impossible trajectories. Several forms of movement models appear in the literature; at the simplest level is speed filtering which prohibits estimates that imply impossible speeds of travel (McConnell et al., 1992; Austin et al., 2003), while other authors propose more complex state space approaches that model correlation between successive legs of the trajectory (Matthiopoulos et al., 2004; Jonsen et al., 2005).

Several authors have noted the advantages of Bayesian methods in complex problems in ecological research (Dixon and Ellison, 1996; Dorazio and Johnson, 2003; Roberts et al., 2004; Wintle et al., 2003; Ellison, 2004); for the tag location problem one principal advantage is that four disparate data sources can be systematically incorporated into a single unified estimator of location.

The novel aspect of the method we propose is the adoption of a simple yet powerful representation of the movement model that not only constrains the animal's trajectory, but also allows this trajectory to be estimated. Between each pair of successive locations  $x_i$  and  $x_{i+1}$ , introduce a new latent point  $z_i$  representing the location of the tag at a time  $\tau_i$  uniformly distributed in the interval  $[t_i, t_{i+1}]$ , and let  $d_i$  be the length of the dog-leg path from  $x_i$  through  $z_i$  to  $x_{i+1}$ . The movement model then simply prescribes the joint distribution  $p(d|t)$  of the dog-leg distances  $d = \{d_1, d_2, \dots, d_{n-1}\}$ . For example, adopting a model where the  $d_i$  are independently uniformly distributed

$$d_i \sim U(0, s(t_{i+1} - t_i))$$

implements a simple speed filter that limits the maximum speed of travel to  $s$ . Alternately, migration and large scale consistency of motion can be modelled by adopting a distribution that allows for more complex patterns of dependence between the successive  $d_i$ .

Note there is no explicit expression for the  $z_i$ , they are defined implicitly through the dog-leg distances  $d_i$ . However, any choice of  $p(d|t)$  that places realistic bounds on each  $d_i$  is sufficient to ensure that the  $z_i$  are estimable (in a Bayesian sense), while also constraining location estimates. Most importantly, as  $\tau_i$  is uniformly distributed in the interval  $[t_i, t_{i+1}]$ , the posterior distribution for  $z_i$  describes the

possible paths between  $x_i$  and  $x_{i+1}$ . In a sense,  $z_i$  is not intended to refer to the tag location at one particular time in the interval  $[t_i, t_{i+1}]$ , but all times in the interval  $[t_i, t_{i+1}]$ .

The second key assumption of the method is that the primary location data, the auxiliary environmental data and the behavioural model are all independent, and so the likelihood  $p(y, q, d | x, t, E)$  reduces to a product of contributions from each of these three sources

$$p(y, q, d | x, t, E) = p(y | x, t)p(q | x, t, E)p(d | t).$$

Here  $p(y | x, t)$  is the likelihood of observing the primary location data  $y$  given locations  $x$  at times  $t$ ,  $p(q | x, t)$  is the likelihood of observing the environmental data  $q$  given locations  $x$  at times  $t$  and a database  $E$  of known environmental data, and  $p(d | t)$  is the distribution of dog-leg distances between the successive locations described above. The exact form of  $p(y | x, t)$  and  $p(q | x, t)$  will depend on the precise nature of the data collected by the tag, and several common examples are discussed below.

The prior for  $x$  and  $z$  reflects knowledge of the animal's home range, habitat preference, migratory patterns or other fundamental environmental considerations. For example, a known home range can be modelled by adopting a prior of the form

$$p(x, z) \propto \prod_i^n I(x_i \in \Omega) \prod_i^{n-1} I(z_i \in \Omega)$$

where  $\Omega$  is the known home range and  $I$  is the indicator function

$$I(x) = \begin{cases} 1 & \text{if } A \text{ is true} \\ 0 & \text{if } A \text{ is false.} \end{cases}$$

Migration can be accommodated by allowing  $\Omega$  to vary with season, while habitat preference can be incorporated by assigning greater probability density to more favourable habitat. We must also supply a prior for  $\tau$  that simply reflects our assumption that  $\tau_i \sim U(t_i, t_{i+1})$ . The form of  $p(y | x, t)$  as the contribution of the primary location data to the total likelihood depends on the nature of the tag in question.

### 3.2.3 Satellite tags

For satellite tracked tags, the primary location data  $y$  consists of primary estimates  $X = \{X_1, X_2, \dots, X_n\}$  of the true tag locations  $x = \{x_1, x_2, \dots, x_n\}$  at times  $t = \{t_1, t_2, \dots, t_n\}$  provided by a remote sensing service, possibly augmented with some indicators of location reliability  $\{r_1, r_2, \dots, r_n\}$ . In this case the contribution  $p(y | x, t)$  to the total likelihood is determined by assuming the observed locations  $X_i$  are bivariate Normally distributed about the true locations  $x_i$ ,

$$X_i \sim N(x_i, \sigma^2(r_i))$$

with a variance  $\sigma^2$  that is a function of the reliabilities  $r_i$ . For less consistent services, longer tailed distributions such as the bivariate  $t$  can be used to accommodate the occasional erroneous location (Gelman et al., 2004).

### 3.2.4 Archival tags

For archival tags there are no initial estimates of tag location; the primary location data consists of light intensities recorded by the tag at regular intervals over the day. The tags' location can be estimated from the light level data by the methods of Ekstrom (2004) and Hill and Braun (2001). We use a version of the template-fitting method (Ekstrom, 2004) to provide a location estimate for each twilight. The full computational details are complex and will be the subject of a future publication, but in essence the method is as follows. The time series of light levels corresponding to each twilight recorded by the tag is extracted, and for marine applications, corrected for attenuation due to depth. This yields a sequence of time series; one time series  $l_i = \{l_{i1}, l_{i2}, \dots, l_{im}\}$  for each twilight, where  $l_{ik}$  is the corrected light level recorded at time  $t_{ik}$ . A function  $l(\theta)$  that maps solar elevation  $\theta$  to the (unattenuated) log light level  $l$  recorded by the tag is determined by laboratory calibration. The contribution  $p(y|x, t)$  to the total likelihood is determined by assuming the log corrected light levels are distributed as

$$\log l_{ik} \sim N(\log l(\theta(x_i, t_{ik})) + k_i, \sigma^2),$$

where  $\theta(x, t)$  is the Sun's elevation at location  $x$  and time  $t$ , and  $k_i$  is a constant to allow for attenuation due to cloud. The variance  $\sigma^2$  is determined by the recording error in the tag.

Similarly, the contribution  $p(q|x, t, E)$  the auxiliary environmental data  $q$  makes to the total likelihood will depend on the nature of the data recorded by the tag and the availability of a suitable reference database  $E$  with which to compare.

For example, for marine tags that record both water temperature and depth, for each  $x_i$  an estimate  $s_i$  of the SST can be derived from the temperature and depth data recorded by the tag in some small time interval  $[t_i - \Delta t, t_i + \Delta t]$  surrounding  $t_i$ . This estimate might then be assumed to be Normally distributed about a reference temperature  $S(x_i)$  determined from a remotely sensed SST database  $E$ ,

$$s_i \sim N(S(x_i), \sigma_s^2)$$

where the variance  $\sigma_s^2$  is determined by the accuracy of both the tag and the remotely sensed database. Alternately, a more conservative approach similar to that employed by Hindell et al. (1991) is to suppose that the temperature  $s_i$  measured by the tag is a very poor indicator of average SST, but could be no greater than an upper limit  $S(x_i) + \Delta S$  and no lower than  $S(x_i) - \Delta S$  and assume  $s_i$  is uniformly distributed in this interval

$$s_i \sim U(S(x_i) - \Delta S, S(x_i) + \Delta S).$$

Again  $\Delta S$  is determined by both the accuracy of the tag and database.

As a second example, for marine applications the depth data recorded by a tag can be exploited by noting that the maximum depth recorded in a time interval  $[t_i - \Delta t, t_i + \Delta t]$  surrounding  $t_i$  provides a lower bound  $h_i$  for the depth of the water column at  $x_i$ . We can then refine the estimate of  $x_i$  comparing  $h_i(x)$  to a high resolution topography database  $E$  and excluding regions that are too shallow by

including in the likelihood a factor of the form

$$\prod_{i=0}^n I(h_i < H_h(x_i))$$

where  $H_h$  is the bottom depth determined from the database and  $I$  is again the indicator function.

### 3.2.5 Posterior estimation

Once the prior and likelihood have been defined, the posterior  $p(x, z, \tau | y, q, t, E)$  is determined by Bayes' rule

$$p(x, z, \tau | y, q, t, E) = \frac{p(y, q, d | x, t, E)p(x, z)p(\tau)}{\int p(y, q, d | x, t, E)p(x, z)p(\tau) dx dz d\tau}.$$

Typically however, the integral in the denominator is computationally intractable, and instead we resort to Markov Chain Monte Carlo (MCMC) to approximate the posterior.

Appendix A provides an explanation of MCMC methods with an example.

MCMC (Gilks et al., 1995) is a family of methods that allows us to draw random samples from the posterior distribution. Summarizing these samples approximates the properties of the posterior, in the same way that a sample mean is an approximation to a population mean. In principle, the approximation can be made arbitrarily accurate by increasing the number of samples drawn.

For the tag location problem we use a block update Metropolis algorithm based on a multivariate Normal proposal distribution (Gilks et al., 1995). The Metropolis algorithm was chosen for its simplicity and genericity – it is easily implemented and the implementation is not strongly tied to particular choices of likelihood and prior. We have used a block update variant of the algorithm, where each  $x_i$  and each  $z_i$  are updated separately. Using a block update improves computational efficiency provided parameters from separate blocks are not strongly correlated. For the time intervals between locations typical of satellite and geolocation data and reasonable choices of movement model  $p(d | t)$ , we have not found the correlation between successive locations estimates to be so great as to greatly impede the mixing of the chain.

### 3.2.6 Examples

To illustrate this basic framework, we present two simple examples.

The first example is a Weddell seal tagged at the Vestfold Hills (78°E, 68°S) tracked with a satellite tag (9000X SRDL; Sea Mammal Research Unit, St. Andrews, Scotland) with locations provided by the Argos service (Service Argos, 2004).

The Argos service provides approximate locations  $X = \{X_1, X_2, \dots, X_n\}$  and corresponding location qualities  $\{r_1, r_2, \dots, r_n\}$  for a sequence of times  $t = \{t_1, t_2, \dots, t_n\}$ . This forms the primary location data. Each  $r_i$  categorizes the corresponding  $X_i$  into

one of seven quality classes based on the number of satellites used in its determination (Service Argos, 2004). We translate the  $r_i$  into approximate positional variances  $\sigma^2(r_i)$  based on the results of Vincent et al. (2002) and assume

$$X_i \sim N(x_i, \sigma^2(r_i)).$$

So that the contribution to the likelihood from the primary location data is

$$p(y | x, t) = \prod_{i=1}^n (2\pi\sigma^2(r_i))^{-1} \exp\left(\frac{-(X_i - x_i)^T(X_i - x_i)}{2\sigma^2(r_i)}\right).$$

This particular tag recorded no environmental data, and so the corresponding contribution to the likelihood is  $p(q | x, t, E) = 1$ .

For this example a very simple movement model was adopted. We choose  $p(d | t)$  so that the mean speeds  $d_i/(t_{i+1} - t_i)$  between successive locations are independently log Normally distributed

$$p(d | t) = \prod_i^{n-1} (2\pi\sigma_s^2)^{-1/2} \exp\left(\frac{-(\log(d_i/(t_{i+1} - t_i)) - \mu_s)^2}{2\sigma_s^2}\right)$$

with  $\mu_s = 0.25ms^{-1}$  and  $\sigma_s = 0.8ms^{-1}$ , where these figures were chosen conservatively based on an examination of Argos data of the highest quality class.

Finally, we adopted a prior  $p(x, z)$  for  $x$  and  $z$  that was uniform over the ocean, that is

$$p(x, z) \propto \prod_i^n I(x_i \in \Omega) \prod_i^{n-1} I(z_i \in \Omega)$$

where  $\Omega$  is the ocean. This was implemented by comparing  $x$  and  $z$  to a high resolution land/sea raster mask generated from A Global Self-consistent, Hierarchical, High-resolution Shoreline Database (Wessel and Smith, 1996). Creating a raster mask to indicate sea/land allows the prior to be computed very efficiently by avoiding complicated point-in-polygon tests.

The second example is a mature southern elephant seal (*Mirounga leonina*) tagged at Macquarie Island (158° 57'E, 54° 30' S), with data from a time-depth-recorder (Mk9 TDR; Wildlife Computers, Seattle, WA, USA). The data were collected using methods described by Bradshaw et al. (2006). This tag provides regular time series of measurements of depth, water temperature, and ambient light level.

In this case the primary location data consist of the time series of depth and ambient light level. As outlined above, the depth adjusted light level is assumed to be log Normally distributed about the log expected light level for the sun elevation adjusted for cloud cover so that

$$p(y | x, t) = \prod_{i=1}^n \prod_{k=1}^{n_i} (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(\log l_{ik} - \log l(\theta(x_i, t_{ik})) + k_i)^2}{2\sigma^2}\right).$$

For this example, the depth and water temperatures recorded by the tag were used to estimate sea surface temperatures that were then compared to NCEP

Reynolds Optimally Interpolated SST. For each twilight, estimates of minimum  $L_i$  and maximum  $U_i$  SST observed in the surrounding 12 hour period were derived from the depth and water temperature records. These estimates form the auxiliary environmental data  $q$ , and  $p(q | x, t, E)$  was then chosen as

$$p(q | x, t, E) = \prod_{i=1}^n p(L_i, U_i | x_i, t_i, E)$$

where

$$p(L_i, U_i | x_i, t_i, E) = \begin{cases} 1 & \text{if } L_i \leq S(x_i, t_i) \leq U_i \\ 0 & \text{otherwise} \end{cases}$$

and  $S(x, t)$  is the NCEP Reynolds Optimally Interpolated SST. This example shows the great difficulty in choosing  $p(q | x, t, E)$  – typically the data from the tag and the data from the reference database are recorded on wildly disparate spatial and temporal scales, making it very difficult to make any reasonable comparison of the two.

Again the movement model  $p(d | t)$  is chosen so that the mean speeds  $d_i/(t_{i+1} - t_i)$  between successive locations are independently log Normally distributed

$$p(d | t) = \prod_i^{n-1} (2\pi\sigma_s^2)^{-1/2} \exp\left(\frac{-(\log(d_i/(t_{i+1} - t_i)) - \mu_s)^2}{2\sigma_s^2}\right)$$

In this case we use  $\mu_s = 1.4ms^{-1}$  and  $\sigma_s = 0.8ms^{-1}$ , and these figures were chosen conservatively based on knowledge of elephant seal behaviour.

Finally, just as for the satellite tag example a prior  $p(x, z)$  uniform on the ocean was adopted  $x$  and  $z$ , but in this case the land/sea raster mask generated from the 2-Minute Gridded Global Relief Data (ETOPO2).

The primary rationale behind our choices for examples was to show the application of our approach to both satellite locations and archival tag data. Further to this, for the satellite example we wish to demonstrate the use of our approach for a situation involving a complex inshore coastline and the handling of existing estimates that occur on land. We are not attempting to show the best possible application for our examples, but demonstrating a consistent approach that is able to use all available sources of data.

### 3.3 Results

For the satellite tag example an initial 10,000 samples were drawn and discarded to allow for both burn-in and tuning of the proposal distribution (Gilks et al., 1995). A further 300,000 samples were then drawn, and standard convergence tests applied (Best et al., 1995). The same strategy was adopted for the archival tag example, with 30,000 samples drawn for burn-in, and a further 800,000 samples drawn. In neither case was there any evidence that the chains had failed to converge, but it must be realized that these are problems of extremely high dimension, and as such a subtle convergence problem may be difficult to detect.

The provided Argos Service locations for the satellite tag example are displayed in Figure 3.1a, showing the primary location data. This includes all raw positions from Argos, including every location quality class. The time-series of locations, is quite noisy and many of the positions fall on land. The sequence suggests that the animal has begun in the southern region of the area, with excursions into and out of various inlets, travelling to the north overall, but with an excursion returning to the south somewhat offshore. The record ends in the northern region. From this plot it is clear that there are many unlikely locations given the presence on land and the implied tortuous path. The outputs of our modelled estimates for this data set are discussed below. Posterior mean locations for  $x$  from the archival tag data set may be seen in Figure 3.2a. Unlike the Argos example, there are no 'raw locations' to present as the primary location data are light level measurements. The range of the track estimate has no local topographic features (coastline or bathymetry) that constrains the locations, as the area visited is for the most part deeper than -2000m (Wessel and Smith, 1996). However, we know that these locations are consistent with the matching sea surface temperature data, under the assumptions of our model.

### 3.3.1 Argos tag data set

In Figure 3.3 the posterior means for  $x$  are plotted separately for longitude and latitude with the sequence of original Argos Service positions overplotted as a line. Also shown are the individual confidence interval (CI) estimates (95% level, presented as a range in kilometres). The sequence of estimates is clearly more realistic than the original Argos locations in terms of likely movement, even though no time steps have been discarded. The confidence intervals in Figure 3.3 are summarized from their 2-dimensional versions and plotted here with longitude and latitude separated to easily show the relative precision of each. Most of the estimates have a range of less than 5 km, with a maximum above 30 km. This simple plotting of individual parameters with CIs leaves out a lot more information than exists in two dimensions. A supporting information file (Figure S1 in Sumner et al. (2009)) provides an animation of the full path with the implied path of the original Argos locations to illustrate the improvement provided by our approach.

The posterior means for  $x$  longitude and latitude are presented spatially in Figure 3.1b. The main differences with the raw estimates is that there are now no estimates that fall on land, and the sequence of positions is far more realistic in terms of likely movement. The 1124 original Argos locations included 179 that fell within the bounds of the coastline data used. The overall travel to the north can be seen in more detail, with an excursion into the main large inlet and then movement around the bay into the region of islands to the north. There are two large excursions when the animal has returned briefly to the southern region, first to the large inlet, then to an island further south, but the more extreme outliers are no longer present. This journey is typical for these seals, as shown by Lake et al. (2005). (We do not present the points connected by lines as this would be visually messy and also imply impossible trajectories based on the simplistic "join the dots" model. The connectivity, or full-path, of estimates is provided by the intermediate estimates.)

A map of time spent per unit area is shown in Figure 3.1c. This density plot shows the "full path" estimate using the intermediate locations, summarized by

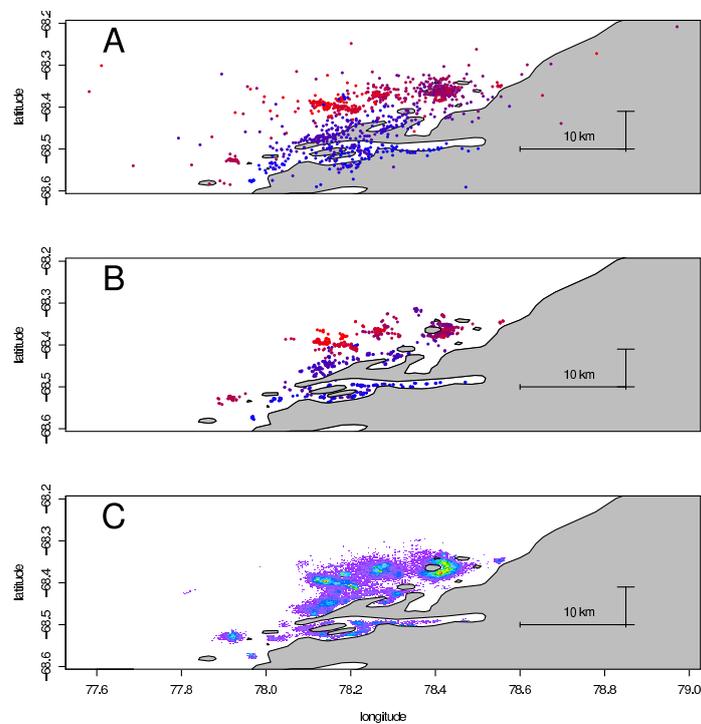


Figure 3.1: **Satellite tag data and estimates** Panel A: The sequence of original Argos estimates for an adult female Weddell seal tagged in the Vestfold Hills, with time scale from red to blue. All location classes are shown. The different length scale bars for north and east represent 10 kilometres. Panel B: Posterior means for  $x$  from the Argos data set plotted spatially, with time scale from red to blue as in panel A. The sequence is far more realistic, without the noise and positions on land. Panel C: Map of time spent from full path estimates from the Argos data set. The density represents a measure of time spent per area incorporating the spatial uncertainty inherent in the model. Bin size is 150 m by 140 m.

binning the posterior and weighting each segment by the time difference between each original Argos time step. The full track estimate is shown here providing

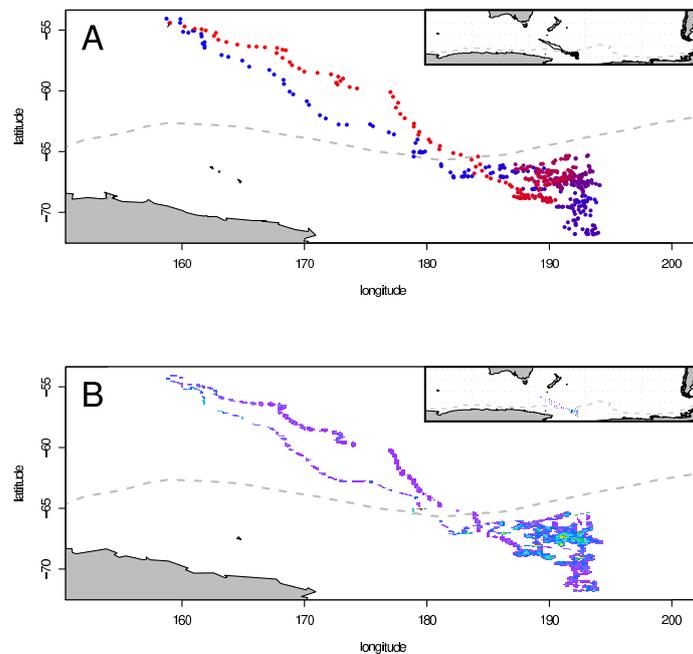


Figure 3.2: **Estimates and time spent for archival data set** Panel A: Posterior means for  $x$  from the archival data set plotted spatially, with time scale from red to blue. The sequence provides a realistic trajectory for an elephant seal. The dashed grey line shows the (approximate) position of the Southern Boundary of the Antarctic Circumpolar Current. Panel B: Map of time spent from full path estimates from the archival data set. Bin size is 5.5 km by 9.3 km at 54 S and 3 km by 9.3 km at 72 S.

a single view of the entire trip. Again, this neglects a lot of information that is available from the posterior, as any segment of the path may be interrogated, down to the level of individual estimates. The bin size here is 150 m by 140 m, simply chosen for convenience given the image plot size. This image portrays the areas of

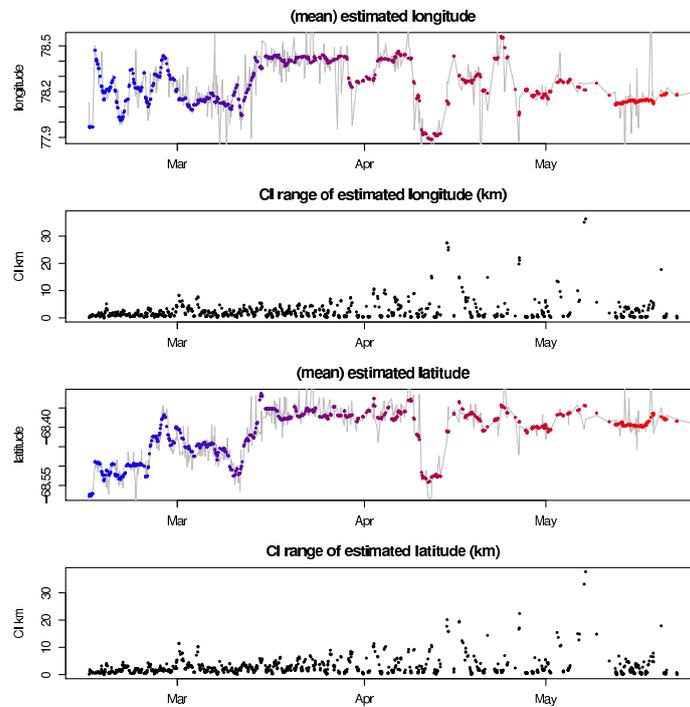


Figure 3.3: **Individual longitude, latitude estimates for Argos** Posterior means for  $x$  from the Argos data set for longitude and latitude, with time scale from red to blue as in Figure 3.1. The grey line shows the implied sequence of the original Argos estimates. Also shown is the range of the 95 % CI of each estimate (km), determined with the mean by directly summarizing the posterior.

most time spent by the animal, with the spatial precision of estimates implicit in the spread of time spent density. Importantly, the transition between time in the water and the position of land is smooth as the estimation takes the presence of land into account as it proceeds. There is no artificial clipping of the distribution as would be required if a simple spatial smoother was used on raw estimates. This achieves the shared goals of smoothing techniques such as kernel density (Matthiopoulos, 2003)

and cell gridding.

A summary of the precision of estimates for longitude and latitude for each original Argos class estimate is presented in Table 3.1. This summary shows that our estimates are consistent with and often better than the expected precision given by the Argos class and, while that point is slightly circular given our use of the class information in the model, our approach is able to combine the contribution of the Argos class with other information and show that the precision of estimates is not necessarily directly related to the class assigned.

Longitude					
<i>class</i>	<i>Min.</i>	<i>1stQu.</i>	<i>Median</i>	<i>3rdQu</i>	<i>Max</i>
Z	0.27	1.09	1.90	2.99	22.05
B	0.27	0.95	1.77	3.95	36.20
A	0.27	1.09	2.18	3.78	15.38
0	0.13	1.36	2.30	4.08	25.86
1	0.27	0.82	1.23	2.04	5.99
2	0.14	0.41	0.61	0.95	2.31
3	0.14	0.27	0.41	0.54	1.50
Latitude					
<i>class</i>	<i>Min.</i>	<i>1stQu.</i>	<i>Median</i>	<i>3rdQu</i>	<i>Max</i>
Z	0.45	1.21	1.97	3.79	17.13
B	0.15	1.21	2.12	4.40	37.75
A	0.30	1.52	2.27	4.40	13.64
0	0.15	1.52	2.50	4.66	19.56
1	0.15	1.06	1.67	2.73	14.86
2	0.15	0.60	0.99	1.67	5.00
3	0.15	0.45	0.61	1.06	3.03

Table 3.1: Estimate precision for Argos data set Summary of precision calculated from the posterior for  $x$  by original Argos class (km). Each row presents a quantile summary for the CI ranges (95 %) from each Argos class for longitude and latitude. The seven classes are an attribute provided with the original Argos locations (Service Argos, 2004).

Finally in Figure 3.4 we can see the relationship between the direct estimates (plotted individually with CI ranges) and CI range of intermediate estimates (plotted as a continuous band) for a short period between 23-26 February 2006. The intermediate estimates provide a continuous path estimate, with latent times of no data “filled in” with estimates constrained only by the movement model and the environmental data. This figure also shows the utility of the method in terms of providing overall full path estimates, as well as individual point estimates with a measure of precision.

Figure 3.4 also shows a deficiency of the assumed movement model - the estimated path at each  $t_i$  tends to be more variable than the corresponding  $x_i$ . This is because there is no constraint on the individual legs of the dog-leg path from  $x_i$  to  $x_{i+1}$ . So it is possible for  $z_i$  to be a great distance from  $x_i$  an instant after  $t_i$

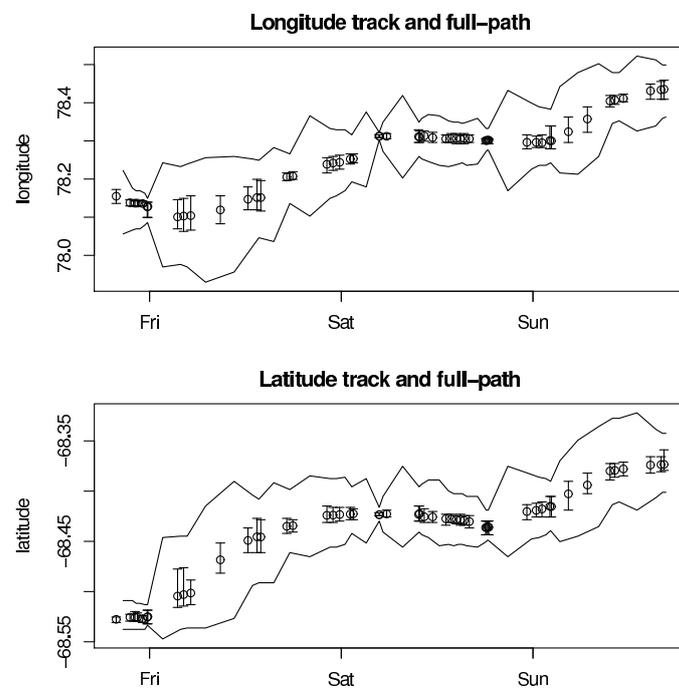


Figure 3.4: Intermediate estimates for Argos Posterior means for  $x$  of longitude and latitude for a short period (23-26 Feb 2006) with CI ranges shown. The CI range for intermediate estimates (full path) is shown as a continuous band.

or from  $x_{i+1}$  an instant before  $t_{i+1}$ , provided the total distance traversed over the dog-leg path is reasonable. It is difficult to resolve this issue without requiring a much more detailed understanding of the animal's behaviour.

### 3.3.2 Archival tag data set

Posterior means for  $x$  longitude and latitude are plotted separately with accompanying confidence intervals Figure 3.5. This includes a location for every local twilight, as seen in the raw light data. The sequence seems consistent with the time steps involved (12 hourly, on average), with no extreme or obviously problematic movements. The confidence interval of each estimate is also plotted, with a spatial range that is usually less than 30 km for longitude and 40 km for latitude. A summary of the precision of estimates for longitude and latitude is presented in Table 3.2.

These estimated location are plotted spatially in Figure 3.2a. This animal has left Macquarie Island (1 February, 2005) and travelled directly to the southeast to a region north of the Ross Sea. Here it spends the period from early March to mid September with a short excursion to the south during April. Finally the animal reverses its outward journey, returning to Macquarie Island on 8 October 2005. The sequence of locations seems reasonable, with no obviously extreme estimates, and this is a fairly typical journey for these seals (Bradshaw et al., 2002).

In Figure 3.2b a density map shows more clearly the spatial precision of the estimates and the areas where most time has been spent. It is clear that this region south of the Southern Boundary of the Antarctic Circumpolar Current (Orsi et al., 1995) is an important feeding area for this animal.

A summary of the precision of estimates for longitude and latitude is presented in Table 3.2. We can see the distinction between the primary and intermediate estimates plotted in Figure 3.6. This time the difference between the primary and intermediate estimates is less than with the satellite tag example.

Longitude				
<i>Min.</i>	<i>1stQu.</i>	<i>Median</i>	<i>3rdQu</i>	<i>Max</i>
3.74	15.52	18.51	21.42	57.03
Latitude				
<i>Min.</i>	<i>1stQu.</i>	<i>Median</i>	<i>3rdQu</i>	<i>Max</i>
4.68	23.38	28.06	37.41	135.60

Table 3.2: Estimate precision for archival data set. Summary of precision calculated from the posterior for  $x$  from the archival tag. A quantile summary for the CI ranges for longitude and latitude.

## 3.4 Discussion

The flexibility provided by Bayesian methods for complex problems (Ellison, 1996; Dorazio and Johnson, 2003; Wintle et al., 2003) proved fruitful in this study. We have demonstrated a general approach for estimating true locations from both archival tag data and satellite fixes, accepting either source as raw data. This approach handles erroneous existing location estimates and other problems by incorporating all available sources of information in one unified process. We have shown how this approach can be used to obtain all of the common measures of interest in

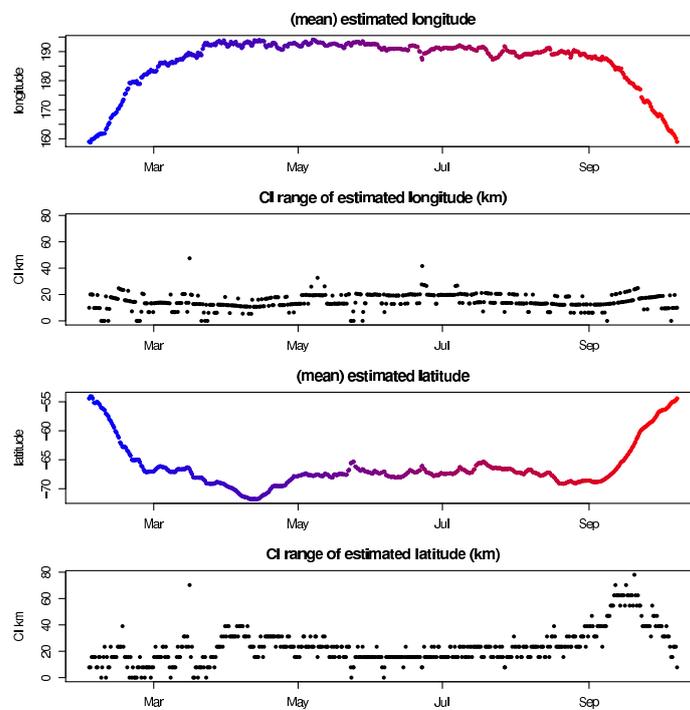


Figure 3.5: **Posterior means for archival data set** Posterior means for  $x$  from the archival data set for longitude and latitude, with time scale from red to blue as in Figure 3.2a. Also shown is the range of the 95 % CI of each estimate (km), determined with the mean by directly summarizing the posterior.

tracking studies by summarizing the posterior. These are path estimates, estimate precision, latent estimates, combinations and diagnostics of location estimates.

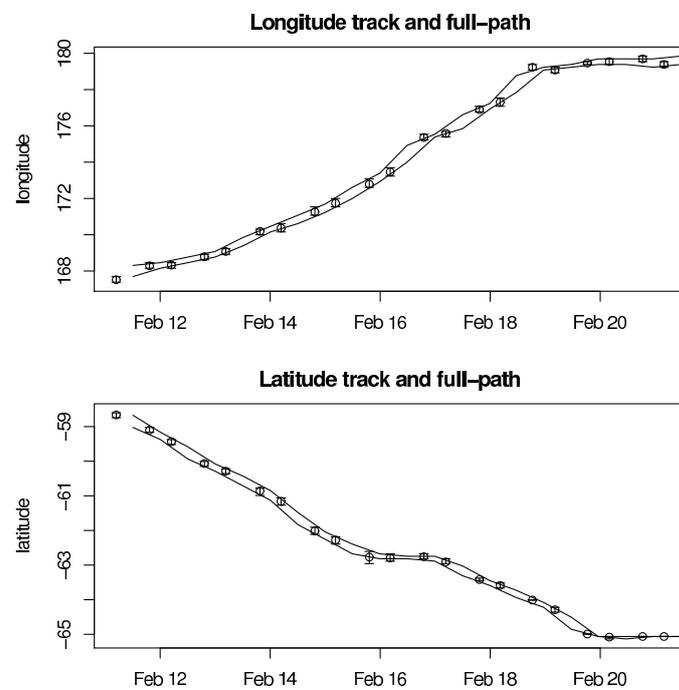


Figure 3.6: **Intermediate estimates for archival data set** Individual mean estimates of longitude and latitude for a 10 day period in February with CI ranges shown, as well as the CI range for intermediate estimates (full path) shown as a continuous band.

### 3.4.1 Path

The likely (posterior mean) path for a basic representation of position over time. These can be used to plot simple tracks, or to query other data sets (such as productivity measures) for corresponding information at that location and time.

### 3.4.2 Precision

For each estimate we can obtain precision estimates (CI). These probability densities are bivariate and can be obtained separately for each time step in the sequence, or for combined durations as required. This information can be used for more nuanced interrogation of other data sets to obtain representative values based on the spatial precision of the estimate.

### 3.4.3 Latent estimates

Estimates of latent locations can be obtained, representing the intermediate positions between those explicitly measured. These represent each period between Argos locations or times between each twilight for archival tags: in general they represent periods between those of (primary) data collection relevant to location estimation. Latent estimates may also be summarized as a mean and CI, and used to provide estimates of the full path between individual time steps. The density of intermediate locations provides a model of the possible range of the track, similar in intention to the spatial smoothing mechanisms employed in other studies.

While primary estimates are constrained by likely movement regimes as well as the available data, the latent estimates represent the residual possible movement in-between.

Unlike some studies using techniques that require subsequent clipping (Bradshaw et al., 2002; Croxall et al., 2004), time spent estimates can be made without spurious presence on land or other out-of-bounds areas. Also, there is a more realistic probability transition from land to marine areas even for complexly shaped coastlines.

The use of latent estimates utilization distributions is better than either cell gridding or kernel density as there is no dependence on the choice of grain size or kernel. The final step to quantize values into a density grid can be done directly from the posterior, without intermediate processing.

### 3.4.4 Combinations

The structure of our estimates enables us to combine estimates from different animals for spatial measures of resource usage. This may be done for arbitrary time periods and groups of individuals. Also raw coordinates may be projected for summaries based on an appropriate coordinate system for particular groups or areas of interest.

### 3.4.5 Updating the models

Time spent maps and track summaries (mean and CI values) were generated by summarizing the posterior for each example. The intermediate locations represent the 'full path' and hence are appropriate for time spent maps and similar spatial summaries. The primary locations are estimates for each time step from the raw location data - individual twilights for the archival tag, Argos times for the satellite

tag. Interrogating individual  $x$  or  $z$  estimates provides feedback on the performance of the model run that may be used to identify problems or areas that require improvement. An example of this feedback was discussed with Figure 3.4 where we see how the movement model requires an improved implementation for the satellite tag. This is one of the most powerful aspects of our approach, more important than the results presented here as it provides a foundation from which remaining problems with location estimates may be identified and related to deficiencies in source data, model specification or model assumptions.

Other studies have successfully applied Bayesian methods to tracking problems with similar success (Jonsen et al., 2003, 2005), but applied only to pre-derived location estimates, and it is not clear how archival tag data could be incorporated in such an approach. The quantities of data involved and the non-linear complexity of the models involved are difficult to implement with more efficient statistical sampling regimes such as Gibb's sampling. Our approach enables the use of the raw archival tag data and incorporation of independent environmental databases. High quality location methods such as satellite tracking can also benefit from our approach. For example: similar to the satellite example presented here, Thompson et al. (2003) also report dealing with large numbers of Argos locations that were clearly deficient as they place marine animals on the land. Our approach allows the systematic use of the appropriate coastline to data account for this inconsistency.

The advantages of our approach are relevant to all users of tracking data including tag manufacturers, ecological researchers and environmental decision makers. The key benefits are:

1. A convenient mechanism for separating large complex problems into manageable components, enabling the use of all available information sources.
2. Obviously incorrect locations are avoided, and when data are absent or of poor quality the estimates will have a lower precision.
3. Estimates are continuous in the posterior and may be summarized as required, rather than being discretized or otherwise simplified.

While we have illustrated our approach using seals, these techniques clearly have broader implications for the tracking of other species and other tagging methods. This approach to location estimation better enables multi-species ecosystems comparisons irrespective of the methods used to collect data. A particularly important area of application is in fishery studies, which have large quantities of archival tag data e.g. Gunn et al. (2003) and Teo et al. (2004), or satellite data e.g. Croxall et al. (2004); Block et al. (2003b); Halpin et al. (2006). The improvement of location estimation will enable further research aimed at relating fisheries management to that of other marine species and processes.

While our approach can provide location estimates with confidence intervals based on the data model, there remains the need for independent validation of the techniques with known locations. The assessment of accuracy of these techniques is crucial to their use, and opportunities exist with double-tagging experiments, recapture studies and experimental validation.

The relationship between tag-measured temperatures in near-surface waters and remotely sensed surface temperature remains largely unexplored in animal tracking studies (Sumner et al., 2003). This is due to the discrepancy between traditional physical oceanographic interests and those of biological studies. Access to hierarchical data sets of SST (Domeier et al., 2005), models of surface and at-depth water temperature and sources of higher quality local environmental data will improve the contributions from this auxiliary information. A more detailed approach would match auxiliary data values in a probabilistic sense similar to methods employed by Teo et al. (2004), enabling the application of distributions to account for error in all measurements.

The use of depth and temperature at depth also remains a largely unexplored aspect, no further work has been published since Smith and Goodman (1986) and Hindell et al. (1991). The utility of this data source obviously depends on the environment visited and the animal's diving behaviour, but also highlights the breadth of opportunities that are available for various species.

Many of our implementation decisions have been deliberately based on simplistic, first-pass practicalities in order to demonstrate the generality of our approach to a wide range of problems. The application of MCMC demands careful diagnosis of model convergence (Plummer et al., 2006) and we have omitted this important but onerous aspect from the present work in order to focus on the primary goal of integrating all the available data. While our movement model is flexible it does not account for movement regimes that are auto-correlated or seasonal. Auto-correlation of speed is recognized as an important aspect of modelling movement, also missing from our initial implementation. For example, in both examples we have assumed that the successive  $d_i$  are independent. However, we can model serial correlation in the track by choosing the joint distribution of distances so that successive  $d_i$  are correlated. The impact of a variety of correlation models could be explored (Jonsen et al., 2005; Viswanathan et al., 2000).

In this study we applied a single scheme to the derivation of location estimates from two very different tracking data sets. Each data set was composed of separate sources of information integrated using our four-part approach. This was used to derive location estimates from raw archival tag data, as well as from pre-derived location estimates from a satellite service. In each case, where limitations from a particular source could have produced problematic estimates, this was augmented by the strengths of others.

This method is clearly practically applicable to the real-world problem of analysing behaviour from many large archival tag data sets employed by marine animal studies, and is appropriate for the tracking data from many species. It is also useful for applying behavioural constraints to the latent aspects of nearly error-free location estimation such as GPS.

## Conclusion

The framework presented in this chapter provides a general Bayesian approach to location estimation that can incorporate prior knowledge, primary location data

from different tag types, large environmental data sets and movement models. The work in this chapter forms the basis of a novel approach to estimating location from disparate tagging methods. An implementation of the approach is freely available in easily extensible form in the **tripEstimation** package developed by the author (Sumner and Wotherspoon, 2010). No previous study has provided as broad integration of the many issues faced by the modelling of animal movement.

This chapter has focussed on the broad structure of a general approach to location estimation, applied to a satellite tag data set and an archival tag data set. The following chapters deal with various aspects of these issues in greater detail. Chapter 4 describes the specifics required for light level geo-location and the requirements for tag calibration and inclusion of solar elevation information. Chapter 5 describes the details of the track representation model given by the separation between primary locations and the estimation of intermediate intervals. The implications of this are outlined in detail with examples to demonstrate the importance of commonly neglected issues such as the use of a hard boundary to inform location estimates.

## CHAPTER 4

# LIGHT LEVEL GEO-LOCATION FROM ARCHIVAL TAGS

Determining location from archival tag data, such as light level and temperature, is an important method for the study of marine animals, particularly those that spend time at depth and migrate over long distances. In general, locations derived from archival tags are not as accurate as those from satellite tags but satellite methods do not work underwater and the tags are considerably more expensive. The major source of inaccuracy is simply that observed light levels are such an indirect indicator of global location, but the problem is compounded by analytical approaches that are statistically inefficient and under-utilize the tag-measured data. For example, the evaluation of estimate accuracy is often performed as a separate process to the estimation, with reference to secondary sources, such as maps of sea surface temperature. These problems were described by Sibert and Nielsen (2007), who provided a state-space approach for geo-location from raw data with a random walk movement model. By applying the general approach adopted in Chapter 3 (Sumner et al., 2009), we present a methodology that provides two locations per day, each with quantified precision and estimates for intermediate locations from an integrated movement model.

The technique of determining location from light levels relies on the relationship of tag-measurable light to solar elevation, at particular times of day. Solar elevation is the angle of the position of the sun relative to the horizon. “Light level” is not an absolute measure and the actual relation between solar elevation and measured light depends on the particular tag construction and calibration. However, the daily pattern of change in solar elevation has a predictable effect on measured light levels and methods of determining location from this pattern are referred to as “light level geo-location” (Hill, 1994). There are two main approaches: “threshold methods” which detect signature patterns that mark particular solar events, and “curve methods” (or “template methods”) which use patterns in the rate of change of solar elevation (Welch and Eveson, 1999; Hill and Braun, 2001; Musyl et al., 2001; Ekstrom, 2004, 2007). Both approaches rely on twilight as the most informative time of day when solar motion has the greatest influence on changes in measurable light level.

Direct astronomical measurements can be used to derive location with an accuracy of a few kilometres or better (Bowditch, 2002). The “longitude problem” for eighteenth century navigation was due to the lack of precise and reliable time-keeping

machines (Sobel, 1998). With an accurate reference time, longitude is easily derived from the sun's azimuth and latitude from its elevation. Time-keeping in miniature tags is now cheap and simple but given that tags cannot directly measure solar elevation their successful use is largely hindered by the "latitude problem".

Light data from archival tags provide only an indirect indication of angular solar position. Light level geo-location is only possible during times when the sun actually rises and sets, and the factors that degrade the estimation process tend to mostly affect the estimation of latitude. Longitude is generally much simpler to determine from light levels than latitude (Hill and Braun, 2001), but both are susceptible to a variety of problems. In terms of modelling the data collection process they are two dependent parameters that arguably cannot be cleanly separated.

## 4.1 Light-based derivation techniques

Light level geo-location has been used since the mid 1980's, with the first work from a collaborative study by Northwest Pacific and NOAA for tagging tuna (Smith and Goodman, 1986; Hunter et al., 1986). Smith and Goodman (1986) gave a detailed analysis on the limits on latitude determination by depth-specific temperatures. Wilson et al. (1992) provided the first general application of light level geo-location, shortly followed by Hill (1994). Together these works established the problems of latitude determination during equinox and the use of ocean temperatures as a way to improve otherwise problematic estimates. This largely set the context for applications of geo-location: longitude and latitude are estimated separately, then further work is done to correct poor estimates. Innovations in techniques for deriving location from light data have been relatively rare and many publications state that the limits of light derivation have been met and must be augmented by auxiliary data (Beck and McMillan, 2002; Block et al., 2003c; Teo et al., 2004; Shaffer et al., 2005). Published works that explore the application of light methods in detail are Wilson et al. (1992); Welch and Eveson (1999); Musyl et al. (2001). Recently methods to extend the use of light data in new ways have been introduced by Sibert and Nielsen (2007); Ekstrom (2007) and Evans and Arnold (2008).

The major difference between existing methods for light-level geo-location is that of "fixed point" methods that choose critical times and "curve" methods that use the rate of change in light level during twilight. The classic and well-known description is a form of threshold method that requires the definition of reference angles for solar elevation (corresponding to noon, sunrise and sunset) inferred from identifiable light level values. Based on the same principles used by eighteenth century mariners, the time of local noon is determined from the mid-point of an assumed symmetric light signal (giving longitude), and the length of day is determined from the time between sunrise and local noon (giving latitude). The classic method effectively assumes a fixed point that is constant over a whole day, then solves for that in the solar equations. This threshold approach is particularly susceptible to asymmetric weather conditions between twilight events, Solar equinoxes, the movement of the animal, and to slight errors in the determination of sunrise, sunset and noon (Hill and Braun, 2001; Ekstrom, 2004).

### 4.1.1 Fixed point methods

In fixed point, or threshold, methods there are a number of ways that the times of twilights or zenith values are determined. The major distinctions are between “fixed reference”, “variable reference” and “reflection” methods (Metcalf, 2001), with the difference between these concerning the estimate of latitude. A fixed reference light level is chosen as representative of the time of the definitive zenith angle, and latitude is found by using these times in standard solar algorithms. Variable reference methods choose a level as a relative value for each day. Reflection methods rely on comparison of subsequent dawn/dusk light curves to define these times by finding the best fit or match between them (Hill, 2005). In essence, these methods estimate latitude based on the apparent day length, and so latitudinal estimates are particularly susceptible to error near an equinox, when the day length is the same at every latitude.

### 4.1.2 Curve methods

The precise rate at which the sun appears to rise or set is a function of latitude. Where fixed point methods effectively estimate latitude from the apparent day length, curve methods estimate latitude from the apparent rate of change of light level.

Published accounts of curve methods are confined to the works of Ekstrom (2004), Musyl et al. (2001) and Sibert and Nielsen (2007). Importantly, these methods provide a means for better utilizing the available light data and allowing for atmospheric effects. The advantages of the curve approach depends on three major facets: it uses the rate of change of solar elevation over time, prescribes a limit to be applied to the range of angles used based on physical principles, and easily accounts for effects caused by slow variations in weather. Curve methods use the actual series of recorded light levels and its relationship to solar angles. This uses more of the actual light data, and is less susceptible to movement of the tag between the twilight periods and the equinox problem described above (Ekstrom, 2004).

Curve methods are still prone to physical problems that disturb the relationship between solar elevation and measured light that would otherwise provide a solution built purely on principles of astronomy and atmospheric physics (Ekstrom, 2007). The animal’s movement and behaviour during twilight, and attenuation by clouds and water depth disturb the relation between ambient light level and solar elevation that are difficult or impossible to account for directly. Further problems such as moonlight, atmospheric effects and depth attenuation remain incompletely explored, with the work of Ekstrom (2007); Sibert and Nielsen (2007) and Evans and Arnold (2008)<sup>1</sup> providing directions for further work. The actual measured response of the tag to light, data discretization, or further processing depend on the particular construction and design of any given tag, presenting even more variations that need to be accounted for in models used to derive location. This complex of problems requires a very general and flexible approach that can be applied to a wide range of tags and data collection scenarios.

---

<sup>1</sup>See the abstract by Hartog et al. in Evans and Arnold (2008).

### 4.1.3 Aims of this study

There is considerable room for improvement in existing light-level geo-location estimates, and in particular, from methods that separate the derivation of location from the estimation of model uncertainty fail to fully utilize the available information. We apply a version of a curve method within the Bayesian framework introduced in Chapter 3 (Sumner et al., 2009).

The key component provides the likelihood of predicted light levels given the measured light levels. A given location  $x$  at time  $t$  has a known sequence of solar elevation providing the predicted light levels via a tag calibration. As in the previous chapter we augment the light data with auxiliary environmental data to limit the range of locations by comparison to independent environmental databases. A behavioural model links the estimates temporally in a physically realistic way.

The following assumptions apply to our application of light level geo-location.

1. The construction of a relationship between solar elevation and light depends upon the existence of the calibration data.
2. This works on (log) light measured at or near the surface, allowing for an additive offset to account for general cloud conditions and ignoring the need for depth correction.
3. Twilight periods are identified from the raw data (ensuring that twilight occurs within the period is sufficient) and are assumed as inputs to the model.

These assumptions will not work for all tags, but they provide a relatively simple application to illustrate the overall approach. Indeed, implementations must be very tag-specific as there are a wide variety of models, with different light responses, quantization errors and onboard processing. The key component is a likelihood function that can be defined for the data set and alternative methods, such as those by Ekstrom (2007) or Sibert and Nielsen (2007), or methods relying on measurement of the length of twilight could be used. We want to encourage a modular approach to the problem in which various components can be deployed from related studies as appropriate for particular species, environments or tags.

We present the required components of the approach: a function to calculate solar elevation from position and time, a calibration function relating measured light-level to solar elevation, a lookup function to provide a time-specific mask location and a behavioural model. We present two examples of archival tags from different species: a southern elephant seal and a Subantarctic fur seal.

The key advantages of the approach presented here are as follows.

1. The method provides two independent locations per day
2. Each location has an individual error estimate.
3. Estimation is less susceptible to equinox problems due to independence from calculated daylength and ability to incorporate non-light data.

4. The framework can admit any information that is available so that periods in which the light data are poor or missing are automatically augmented by the other data sources.
5. The approach is open and fairly simple so remaining problems can be addressed by inspecting the contribution of each component to the result and modifying the application as required.

## 4.2 Methods

We illustrate our method with data from two archival tags. The first from a Subantarctic fur seal (*Arctocephalus tropicalus*) from Amsterdam Island ( $37^{\circ} 55'S$ ,  $77^{\circ} 30'E$ ) using data collected with a Wildlife Computers Time-Depth-Recorder (TDR; MK7) from 10 June 1999 to 19 July 1999. The data for this tag was kindly provided by Gwen Beauplet and its collection is described in Beauplet et al. (2004). The second tag is from a southern elephant seal (*Mirounga leonina*) from Macquarie Island ( $158^{\circ} 57'E$ ,  $54^{\circ} 30'S$ ) using data collected with a MK7 TDR from 28 January 2002 to 26 September 2002. Data collection methods for this tag are described in Bradshaw et al. (2002).

Our method used four broad types of data for determining location.

1. Prior knowledge which prescribes a maximum range for each animal's trip.
2. The primary location data based on the light level time series for each twilight.
3. Auxiliary environmental data which define the ocean region based on factors such as the presence of land or known sea surface temperature.
4. A movement model applied to the representation of primary locations  $x$  and intermediate locations  $z$  links the estimates temporally (Sumner et al., 2009).

Incorporating the contribution of light data to the likelihood requires the following components.

**Solar Elevation Function** The solar elevation function  $\theta(x, t)$  determines the elevation  $\theta$  of the sun when observed from location  $x$  at time  $t$ , and is determined by standard astronomical formulae (Meeus, 1991).

**Calibration Function** The calibration function  $l(\theta)$  relates the light level  $l$  recorded by the tag for a given solar elevation  $\theta$ . The calibration function can be provided by the tag manufacturer, or determined by applying semi-parametric regression techniques to light levels recorded by the tag at a fixed location.

**Likelihood Function** The likelihood function  $p(l|x, t)$  is the probability of observing a light level  $l$  at a given location  $x$  and time  $t$ . In combination, the solar elevation and calibration functions allow an expected light level  $\hat{l}$  to be determined for a given location and time. Given the expected light level, the probability of the observed light level is then determined based on the properties of the tag.

The primary location data  $y$  is tag light data informing the locations  $x = \{x_1, x_2, \dots, x_n\}$  of the tag at a sequence of times of twilight  $t = \{t_1, t_2, \dots, t_n\}$ . For each twilight, the light data is extracted for the period in which the sun is rising or setting. This yields a sequence of time series, one time series  $l_i = \{l_{i1}, l_{i2}, \dots, l_{im}\}$  for each twilight, where  $l_{ij}$  is the light level (corrected for depth) recorded at time  $t_{ij}$  within the twilight period.

The contribution  $p(l_{ik}|x_i, t_{ik})$  to the total likelihood is determined by assuming the log corrected light levels are Normally distributed

$$\log l_{ij} \sim N(\log l(\theta(x_i, t_{ij})) + k_i, \sigma^2 f),$$

where  $\theta(x, t)$  is the Sun’s elevation at location  $x$  and time  $t$ , and  $k_i$  is a constant to allow for attenuation due to cloud. The solar position algorithms of Meeus (1991) used are available in the R package **tripEstimation** (Sumner and Wotherspoon, 2010). For efficiency within the estimation the location-independent components of solar declination and hour angle are precomputed from  $t$ . Following Ekstrom (2004), a limit is applied to the elevations used by weighting the distribution for values outside of a given range.

The posterior is approximated with Markov Chain Monte Carlo (MCMC) using a block update Metropolis algorithm as in Chapter 3. An individual proposal provides solar elevations for each  $t_{ij}$  using  $\sigma(x)$ , and from these the calibration function provides predicted light levels offset by attenuation  $k_i$ . The likelihood of predicted light levels are then calculated according to the equation above, providing the basis for estimation.

An illustration of the effects of different locations on the expected solar elevation (and thus expected light) is shown in Figure 4.1. The matching solar elevation plots for locations along both curves at particular locations during a morning are shown in Figure 4.2 (squares) and Figure 4.3 (triangles) respectively. These figures illustrate the variation in solar elevation for different locations at a given time, and how subtle the changes can be.

The relative line thickness in Figures 4.2 and 4.3 matches the relative point sizes for the locations in Figure 4.1. There are very small difference in the pattern of solar elevations for relatively large geographical differences along the axis of the morning twilight between light and dark. This axis, and its positions later in the day are represented by gradations in the background colour.

The band of variation for the pattern of solar elevation around sunrise<sup>2</sup> is very narrow, matching the relative “twilight band” that is represented by coloured regions in Figure 4.1. Different choices of location along the sunrise great circle provide very little variation in expected pattern of solar elevation. The pattern for locations along the corresponding but diametrically opposite “evening twilight band” provide a lot more differentiation. This is an alternative way of viewing the narrow “twilight region” of information that is available during the morning and evening. The orientation for each of the twilight bands switches daily, which for this time match the

---

<sup>2</sup>Here and elsewhere “sunrise” and “sunset” simply refer to their everyday meanings of “general time of day” at which the sun rises and sets. The “twilight periods” are defined as the time of day when there is a reliably recognizable change in the light levels—the input data do not need to be exactly classified as such as long as the period between dark and light is captured.

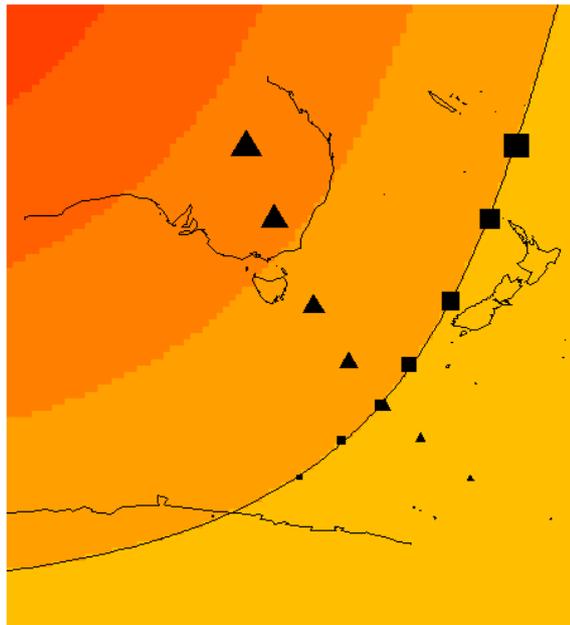


Figure 4.1: Locations around Macquarie Island ( $158^{\circ} 57'E$ ,  $54^{\circ} 30' S$ ). Square icons indicate locations along the current twilight line, triangles indicate the corresponding “evening” locations. Macquarie Island is at the intersection of the two curves.

intersecting lines of triangles and squares in Figure 4.1.

This provides the basis for calculating the likelihood for a given location, by comparing the difference between measured and predicted light levels. There are many ways to determine this measure of “difference”, we use the lognormal density as described in Section 3.2.6. Other methods could include simply the absolute difference or a correlation measure, and the data itself could be subjected to smoothing or curve-fitting of various kinds as necessary.

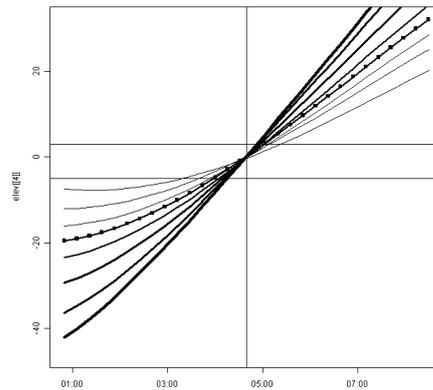


Figure 4.2: Morning twilight solar elevations for locations around Macquarie Island from Figure 4.1 (squares). The time at which the elevation is 0 (sunrise) is indicated with a vertical line, and the horizontal lines show solar elevations -5 and 3.

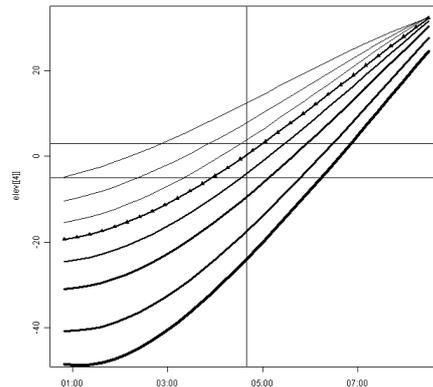


Figure 4.3: Morning twilight solar elevations for location around Macquarie Island from Figure 4.1 (triangles). The time at which the sun is at elevation is 0 (sunrise) is indicated with a vertical line, and the horizontal lines show solar elevations -5 and 3

### 4.2.1 Auxiliary data

For each segment, any MCMC proposal locations that do not fall within the prescribed range given by environmental data are simply excluded. For efficiency, these ranges are pre-calculated and made available as a simple lookup function for location during a given segment. This masking approach provides a very conservative usage of these data, to avoid potentially erroneous assumptions. We used a topographic data set to limit the estimation to areas of ocean for both examples, as well as Reynolds OISST for the elephant seal. More discussion on the application of these environmental data sets was provided in Section 3.2.6, and the use of masks is revisited in Chapter 6.

### 4.2.2 Movement model

The model also includes a behavioural constraint. The movement model is chosen so that the mean speeds between successive locations are independently log Normally distributed.

### 4.2.3 Estimation

Initialization of the MCMC estimation requires a starting location for each primary location  $x$  along the track. This is done using a simplistic grid search to determine the approximate maximum likelihood estimates. The likelihood is calculated at each point of a coarse grid of locations, using a fixed value for the attenuation  $k_i$ . Locations that are not consistent with the prior region or masks are ignored. Taking the intersection of the current, previous and next regions for each grid provides a location for each segment that is consistent with the light data and the movement of the animal. Note that this result will vary for different values of  $k_i$  and ignores the behavioural model.

While this is a very coarse method for initialization, this gives rough “first pass” locations that can seem quite reasonable. MCMC is in general not dependent upon the actual starting points used (Gilks et al., 1995) and in theory this only affects the time taken for eventual burn-in<sup>3</sup>. However, this initialization can be problematic as the behavioural model does not account for the topology of the prior space where accepted dog-legs paths may imply an otherwise impossible journey due to the wide variety of coastline shapes, or other irregular spaces implied by the masks. Also, these initial points are of course not the final estimates, which must be derived from samples from the posterior after running the full estimation model and ensuring good mixing of the chains.

The MCMC estimation proceeds by sampling from these starting locations. Results are generated by direct summarization of the posterior, by binning the samples for  $x$  and  $z$ . From these binned estimates, quantiles were calculated for the binned primary  $x$  estimates to give “most probable” locations and precision, and time spent

---

<sup>3</sup>“Burn-in” is the early period of MCMC samples that are transient and unrepresentative of the equilibrium distribution.

maps generated from the binned intermediate  $z$  estimates. A simple density estimate on each segment allow us to determine a percentile range. For calculating the location precision of estimates we first project the longitude-latitude coordinates to a local instance of the Lambert Azimuthal Equal Area projection (Evenden, 1990) to simplify the distance calculations.

The software used in this work is available on the R repository CRAN in the package **tripEstimation** (Sumner and Wotherspoon, 2010).

### 4.3 Examples

For the fur seal example, only one twilight period is available during the afternoon on the 10 June 1999 at the release site at Amsterdam Island. For the elephant seal example, there was one appropriate twilight at the start of the trip (28 January 2002), and three at the end of the of the trip at Macquarie Island (25-26 September 2002). The solar elevation range for both examples was set at  $[-8, 5]$  chosen from short test runs. To exclude light data outside of this range, the  $\sigma$  term is increased by an order of magnitude.

The fur seal example was run for  $1.4 \times 10^5$  iterations, the elephant seal example was run for  $7.8 \times 10^5$  iterations. For both the first  $1 \times 10^4$  samples were discarded to allow for burn-in.

The estimates over time for longitude and latitude are shown in Figure 4.5 for the Subantarctic fur seal and in Figure 4.4 for the southern elephant seal. These estimates consist of two independent locations per day, with confidence intervals for each. The posterior mean value is plotted (red) with the 95% confidence interval. There is no change in the overall value of the location precision at times of equinox for either example.

The range (in kms) of each estimate is plotted for longitude and latitude in Figure 4.7 and Figure 4.6.

Figure 4.8 and Figure 4.9 show the time spent maps derived from the entire posterior for the intermediate locations. The fur seal has travelled directly to the east, taking time to slowly move to the north. No data is available for this after 19 July 1999. The fur seal travels from Macquarie Island directly to the south-west to the coast of Antarctica at  $140^\circ$  E where it remains for several weeks when it moves further west. It then returned back to the east and north of Macquarie Island to the shallower waters of the Campbell Plateau, remaining for several weeks and then returning directly to Macquarie Island.

#### 4.3.1 Discussion

In this work we demonstrate a systematic method for determining location from archival tag data, with light levels as the primary data source. These results present rich location estimates for two different diving animals from data recorded by attached archival tags. The variance of location estimates is readily determined from the posterior by calculating appropriate confidence intervals. The posterior esti-

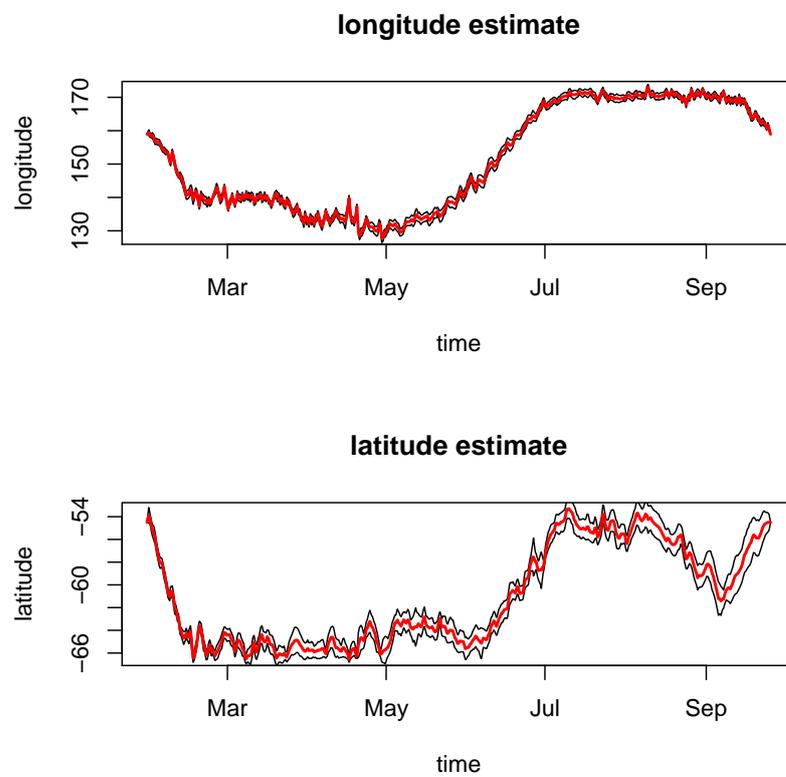


Figure 4.4: Mean longitude and latitude values from the posterior, plotted with surrounding 95% confidence intervals for elephant seal example.

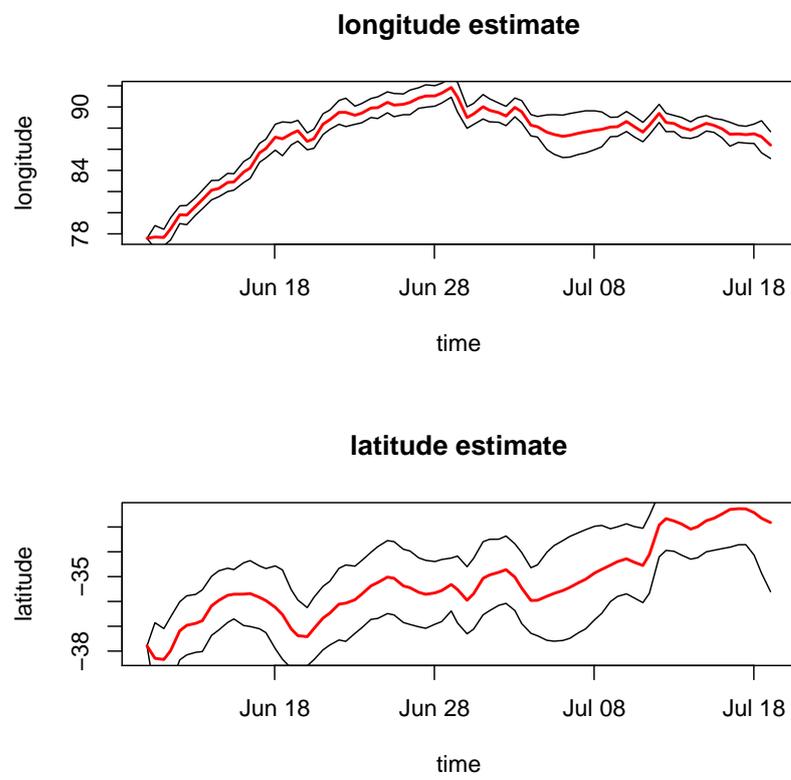


Figure 4.5: Mean longitude and latitude values from the posterior, plotted with surrounding 95% confidence intervals for fur seal example.

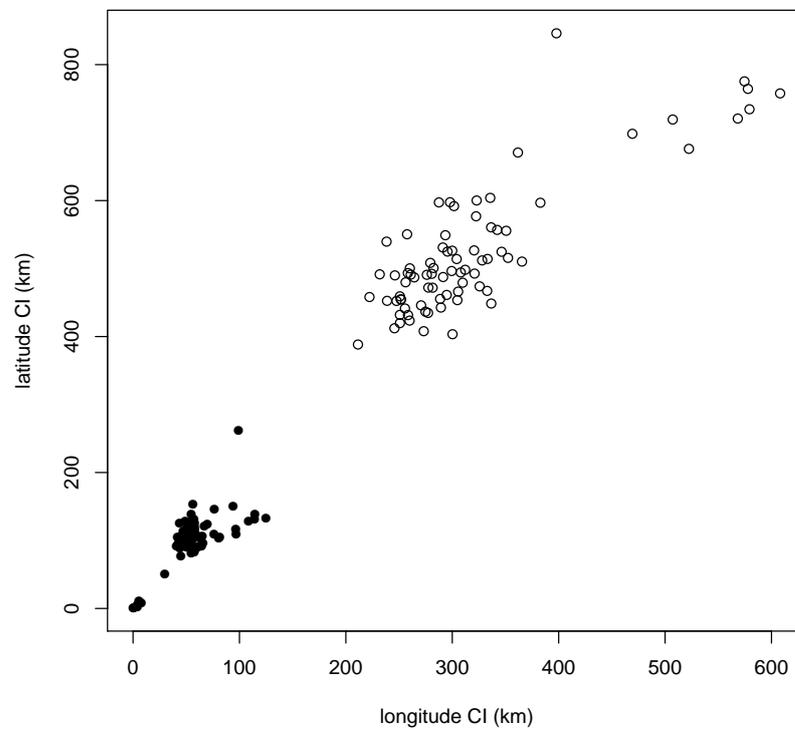


Figure 4.6: Mean longitude and latitude values from the posterior, plotted with the 95% confidence intervals (closed circles) and the full binned estimate boundary (open circles) for the fur seal example.

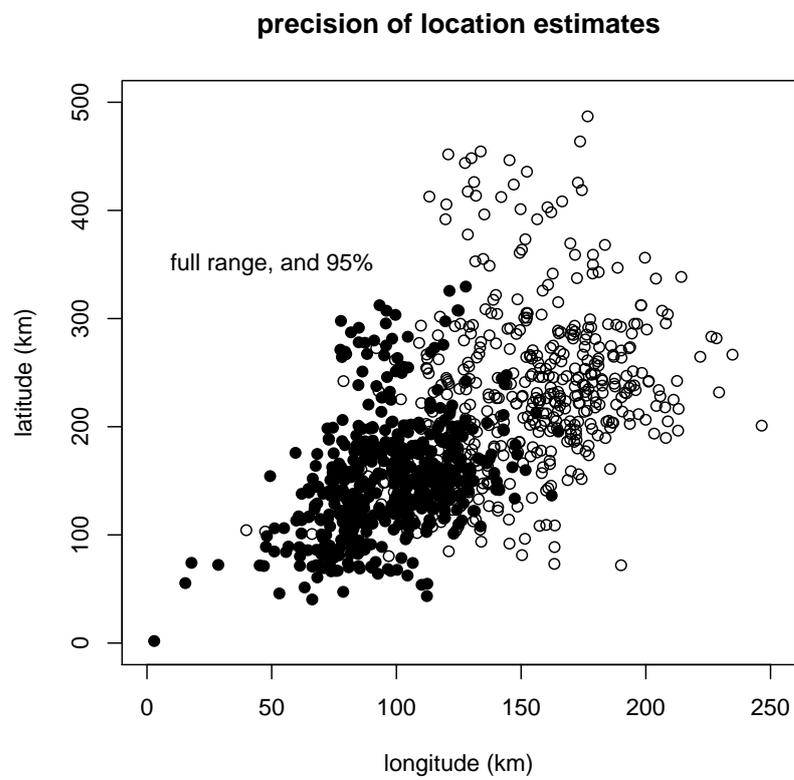


Figure 4.7: Mean longitude and latitude values from the posterior, plotted with the 95% confidence intervals (closed circles) and the full binned estimate boundary (open circles) for the elephant seal example.

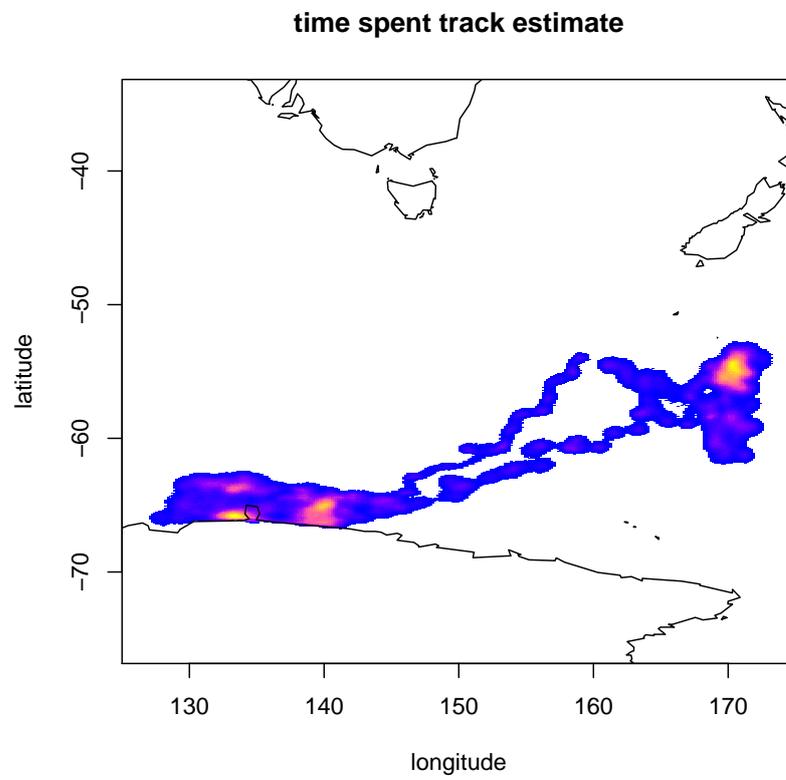


Figure 4.8: Time spent estimate for entire trip derived from all posterior samples for Macquarie Island elephant seal.

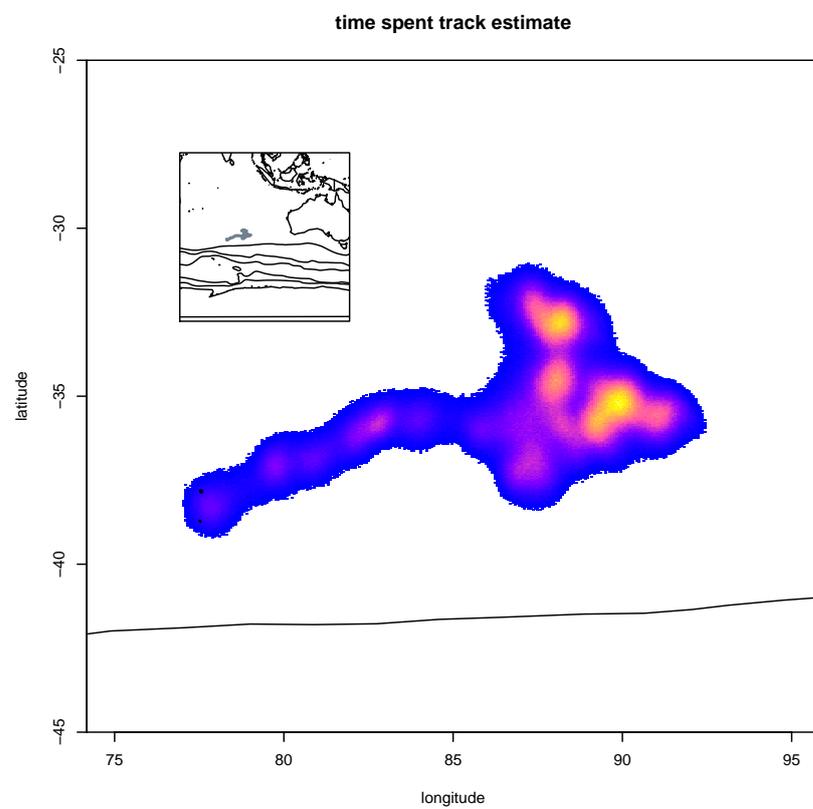


Figure 4.9: Time spent estimate for entire trip derived from all posterior samples for Amsterdam Island fur seal. The zonally (east/west) oriented lines in the inset are mean locations of Southern Ocean fronts (Orsi et al., 1995). Only the Subantarctic front is seen in the main figure region.

mates represent a starting point for analysis in that once calculated, the required percentile for confidence intervals are chosen as appropriate for a given model. This approach represents a significant improvement for the use of disparate sources of data.

In addition, the estimation of intermediate as well as primary locations provides a way of applying measures of behavioural resources to location estimates. Here we have presented full-trip time spent estimates as summaries, in the next chapter we explore the distinction between primary and intermediate locations in greater detail.

For the Subantarctic fur seal example we did not apply any environmental data, as the animal never dives to a depth that would provide any limit in the surrounding bathymetry, and the difficulties with temperature matching tag and existing SST data sets overwhelmed any clear benefit. These difficulties are the focus of ongoing work on temperature validation and use of ocean models which is presented in Chapter 6.

Future improvements could include a more topological constraint for complicated coastlines and environmental time series that can account for non-sensible paths that are otherwise unable to be detected by the methods used here. The approach presented here provides a strong foundation for improving existing methods and developing new methods for light-based geo-location techniques.

## 4.4 Conclusion

We have demonstrated that light data can be used to greater potential than existing methods by incorporating it in a Bayesian approach applying environmental data sets and integrated movement models. This chapter presented the detail and requirements for performing light level geo-location as a component of the general framework presented in Chapter 3. The approach was demonstrated with two example tag data sets, one from a Subantarctic fur seal and one from a southern elephant seal. The software used for these examples is freely available as the **tripEstimation** package developed by the author (Sumner and Wotherspoon, 2010). The examples presented here show time spent estimates of entire trips from intermediate locations as well as primary location estimates with integrated measures of accuracy derived from the posterior. The posterior is represented by large databases of MCMC samples for each primary and intermediate estimate.

In the next chapter we explore the track representation behind the primary and intermediate location model, and show how it can be used to model individual locations as well as the full path, and be used for combining estimates from multiple trips.

## CHAPTER 5

# REPRESENTATION OF TRIP ESTIMATES

Previous chapters introduced a general framework for modelling track locations from tag data. The results from these models provide a rich resource for generating reliable tracks and other summaries such as time spent maps. In previous examples only the full path has been illustrated but the posterior can be queried arbitrarily for different metrics for the whole track or portions of the track, and can be combined in multi-animal estimates. This chapter illustrates the track representation of primary and intermediate estimates, and details a system for efficiently dealing with these model outputs.

Ecological studies involving animal tracking require metrics and summary outputs derived from spatial and temporal movement data. Methods used are varied, depending on the species, region, tag method or available software and no universal or recommended scheme exists. Location methods inherently provide estimates at instants in time, sampled from the otherwise continuous movement of the animal. This chapter presents a representation of tracks that distinguishes between discrete locations and the full path that connects them. The representation is illustrated with practical examples using a public data sample with software in the R package **tripEstimation** developed by the author (Sumner and Wotherspoon, 2010). The representation model of track data presented here provides a new perspective that admits modern statistical methods for location estimation and track summaries.

There are many new modern statistical techniques for estimating and analysing animal track data, but freely available software tools for efficient handling of sets of track data and implementing modern techniques are still rare, as discussed in Section 2.2.3. More traditional methods such as speed filters, cell binning and kernel density are easily applied, but this is rarely done in a way that utilizes the temporal aspect of the track in a flexible or interactive way. The wealth of multi-disciplinary studies highlights an increasing need for access to the tools and techniques for multi-dimensional analysis, including abstractions from real world space. The limited support for continuous variables in GIS vector was discussed in Chapter 2. The crossover of these research domains is rare partly because of this divide in data representation, but is increasingly relevant due to multi-disciplinary studies (de La Beaujardière et al., 2009; Veness, 2009; Beegle-Krause et al., 2010).

This chapter presents methods for representing track data that provide greater emphasis on the discrete versus continuous nature of the estimation involved. The Markov Chain Monte Carlo (MCMC) model runs for the Bayesian approach presented in Chapter 3 produce long chains of samples from the posterior. A data structure scheme for storing binned summaries of these posterior samples is presented.

The challenges for handling large databases of MCMC samples include the simultaneous handling of multiple individuals (even populations and species) over variable time frames. A simple approach demands no more than to bin the required time range of samples into a grid, but this can unnecessarily limit the analyses performed and disregards the importance of quantifying and representing error for individual location estimates. Supporting a flexible range of analyses requires the ability to choose sets of tags, time periods and spatial extent in near real-time for analysis.

Here we describe an approach, implemented in **tripEstimation** that performs the binning of samples for each animal trip into a spatio-temporal grid. These may be thought of as 3D arrays stored in a sparse manner as only the spatial bins that are visited are generated. We present a scheme for representing track estimates as a time series of density grids derived from MCMC samples and illustrate its use for generating metrics for track data. We also apply the scheme to more traditional track representations and smoothing methods.

No existing study has made an explicit distinction between primary locations that represent purpose-measured data and intermediate times between these. By making this distinction a track estimate is more clearly representative of the data collection process. The intermediate locations are truly continuous in that each individual element represents the entire interval between each subsequent primary location. Combining these estimates inherently provides a full-path estimate of the animal, and the design easily handles sets of locations from multiple animals.

## 5.1 Representation of location estimates

In Chapter 2 several problems with simple representations of track data as points and lines were identified. Points do not clearly indicate the temporal aspect of a trip, and while lines improve on this they do not readily add clarity in terms of location uncertainty or time spent. There are a number of requirements for improving on this such as estimates that incorporate a measure of their uncertainty, estimates that readily yield metrics such as mean and standard error, line summaries for path representation, and spatial measures of residency or “time spent”.

The following design provides an efficient mechanism for handling long time series of probabilistic estimates of track locations. Each “location” element is a small “raster”—a matrix of values that is spatially and temporally registered. Raster support for geo-registered 2D grids is easily available in many software packages, and they are easily contoured, quantized, visualized and summed. Using a modern software language, the ability to work with matrices as indexable objects provides an efficient mechanism to store long time series of estimate densities as a sparse array. Individual time steps are easily combined using array index functions, whether to

visualize individual time steps as an animation, or to sum time steps from longer intervals of a track, from multiple animals or across seasons. With quantization functions and contouring these can be quickly summarized as vector data or other related GIS formats.

A “parent” grid encompassing the entire region is defined with a given grain size and offset. This is then treated as a virtual 3D array, without requiring that the parent matrix be duplicated for every time step. (See Figure 5.3). A spatial “child” window of each time step is stored to encompass only the samples for each estimate. The extents in the third dimension are set by the number of time steps in the trip estimation. The 3rd axis is defined implicitly by the tag-times and hence may be irregular, but the first 2D axes are set by the parent.

## 5.2 Primary locations and intermediate locations

Chapter 3 introduced a representation of track data as a series of interleaved “primary”  $x$  and “intermediate”  $z$  location estimates. The  $x$  locations are for instants or short intervals at which the primary location data were recorded. These correspond to the times provided by a GPS or the Argos service, or to the periods of light data collected during twilight for an archival tag. The  $z$  locations represent the latent intervals that are intermediate between each  $x$ . These are times at which no primary location data were collected and represent the entire interval between each  $x$ . Distinguishing between these primary and intermediate locations provides a number of benefits discussed below.

The design of this scheme is too complicated to show in a single diagram, so the following illustration introduces the relevant concepts. Consider a very simple example in Figure 5.1A, with a track of three measured locations shown in the traditional way. As shown in Figure 5.1B each of these location estimates has some underlying uncertainty, depending on the methods used to produce them. No matter the location error or what methods were used these points represent the time at which data specific to determining location was measured.

By using these data and models of the animal’s movement the aim is to improve on this uncertainty to pin down a location to a known region, with as great a confidence as possible. This is represented in Figure 5.1C as three densities giving a relative measure estimate of location for the points. These densities could be practically any shape, possibly with multiple peaks and be as steep or as flat as required to represent what is known about the location. There is another aspect to the uncertainty: the time periods that are intermediate to each point. No matter how well a position estimate is known there remains a region between each estimate that could be very narrow and directed or very sparse and wide. The regions shown with question marks in Figure 5.1D are potentially long durations in time for which there is no location-specific data.

Based on a model of the animal’s movement, with a distribution or limit on likely speeds these intermediate uncertainties can be defined to more or less directed regions. In Figure 5.2 A and B representations of the two intermediate zones are shown as lines, indicating the width of region possible between the end points. Note

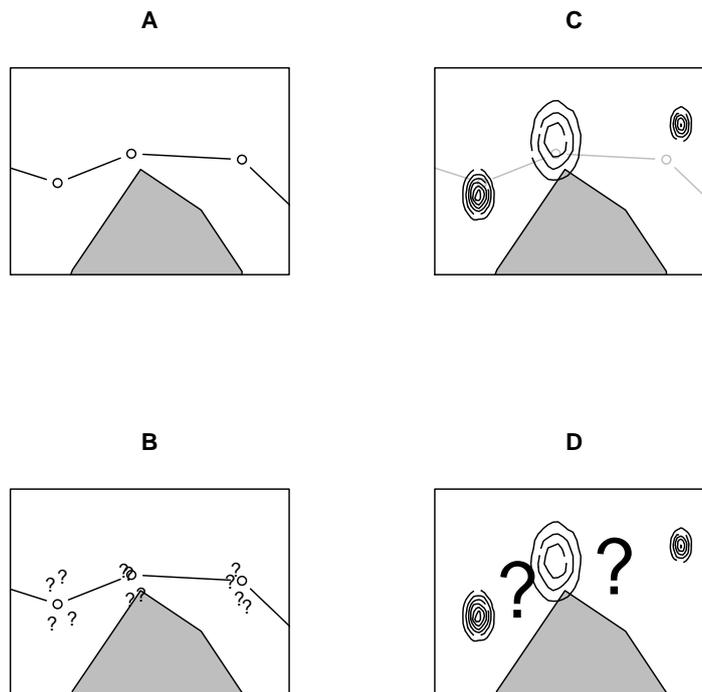


Figure 5.1: (A) Simple track of three locations in a traditional representation, moving from left to right. The grey region is included to represent a boundary to the animal's movement, such as a coastline. (B) The points represent unknown locations with some uncertainty. (C) A representation of these location estimates which are disconnected in time. (D) Connecting the measured locations are regions that are again unknowns with some uncertainty.

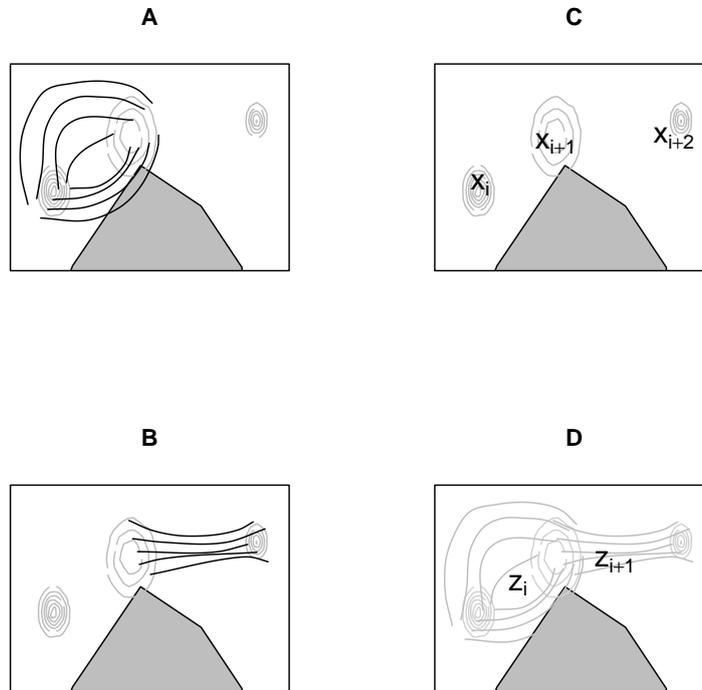


Figure 5.2: (A) Between the first two measured locations is a region of unknown migration connecting the two end points, which can be rather wide. (B) The second connected region, more constrained relative to the first in terms of possible lateral movements. (C) Notation of  $x_i$  for primary estimates. (D) Notation of  $z_i$  for intermediate estimates.

that this is less about location measurement *per se*, and more a measure of how far the animal may have ranged—this is mostly constrained by the model of movement and any remaining auxiliary data, such as the coastline or environmental data sets that could exclude certain regions. In Figure 5.2 is corresponding notation introduced in Chapter 3 for the primary estimates (C) and the intermediate estimates (D).

Separating these primary locations (corresponding to the times at of recording primary data) and intermediate locations (corresponding to latent times) allows for representing the full path of the animal, as well as a series of discrete estimates. Note that depending on the actual path of the animal subsequent  $x$  estimates may overlap in space, and that  $z$  estimates should always overlap with the corresponding  $x$  estimates.

The MCMC methods described in Chapters 3 and 4 produces samples from the posterior. The samples for the primary locations  $x$  and intermediate locations  $z$  must be binned spatially. For efficiency, this scheme defines a parent grid to encompass

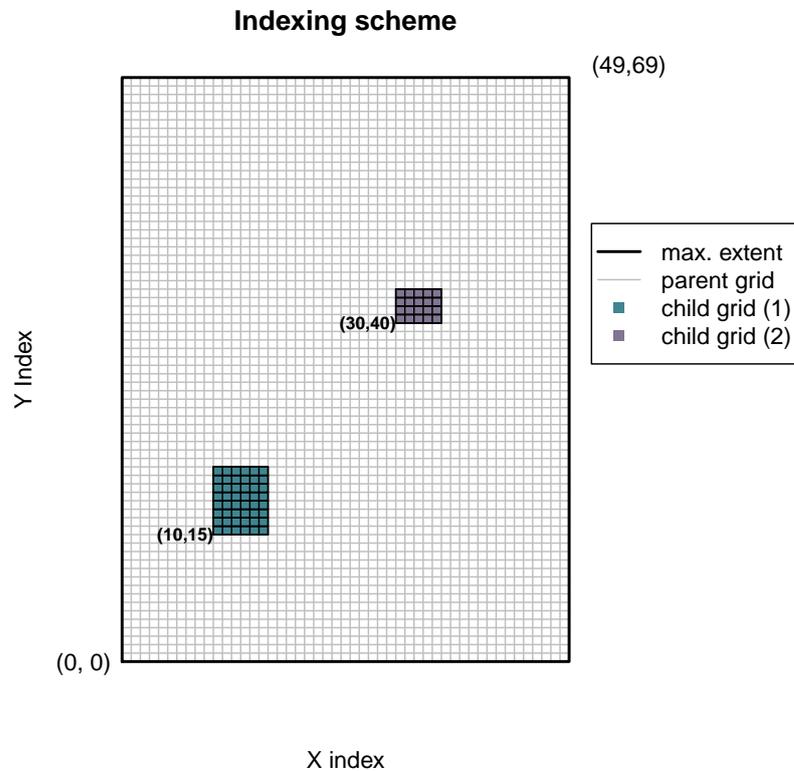


Figure 5.3: Illustration of the grid index scheme. Two child grids are shown within the parent, with indicated index values.

all accepted samples at a chosen resolution and each separate time step is binned into a window cut from the parent grid. This is illustrated in Figure 5.3 with the parent grid displayed in full, with two “child” grids shown. This simple grid shows clearly the full parent grid and the relationship of the child grid as a simple index position within.

Figure 5.4 shows a more practical grid, with a much finer resolution. The detail of the grid itself is shown in the inset.

At each time step, we store only as much of the parent grid as required to hold the samples. That is, for each  $x$  and  $z$ , we record a rectangular subgrid just large enough to contain the samples for that  $x/z$  together with an offset that locates the subset within the large parent grid. This is in effect a three-dimensional sparse array, where the X and Y dimensions are regularly spaced and the third dimension corresponds to the times  $t_i$ , which may represent instants or short intervals. For  $x$  these are the times of the primary location data, such as Argos times or twilights for archival tags, for  $z$  they represent the interval between each  $x$ .

Only the smallest required subset is stored, and so the binning is fast as we are not handling redundant empty cells. Before the actual tabulation is performed a

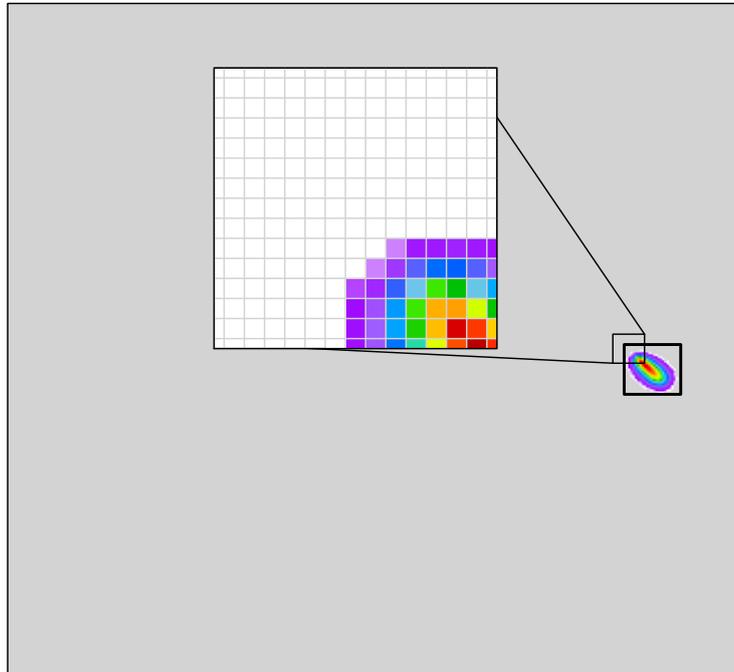


Figure 5.4: Example showing a real child grid within a parent grid. The inset window shows the detail of the grid.

test for out-of-bounds coordinates invokes any required expansion of the child grid to encompass samples outside the range yet seen. The parent will never be less than is required as the prior bounds applied by the estimation already ensures this. This model can also be applied to other forms of track representation such as those discussed in Chapter 2.

### 5.3 Features of the representation scheme

This section describes the overall features that apply to any track represented with this method, followed by examples.

#### Binning the posterior

The posterior is represented as a multi-dimensional array of samples that can be rather large and unwieldy. There will be at least two model parameters (X and Y, or longitude and latitude coordinates) for each time step, multiplied by the number of samples ( $1 \times 10^5$  is a rough rule-of-thumb (Gilks et al., 1995)) and this nearly doubles for the intermediate locations. With archival or satellite records of several months this can amount to hundreds of millions of values which, while in the realm of current day desktop computers to handle in memory, still presents difficulties for creating summaries and visualization (Unwin, 2006).

This system we have presented is very flexible and allows for many options that are otherwise difficult to apply or restricted by certain constraints in other methods. These include arbitrary map projections and choice of origin and scale for the final grid, temporal partitions, and multi-trip and multi-animal combinations. The system can also be used for handling sets of grids from traditional track data like that discussed in Chapter 2.

### 5.4 Examples

This section presents more specific outputs from a trip estimate represented with the scheme.

A publicly available data set is used with code that can be run with some simple preparation. The data set is of Argos satellite estimates from a ringed seal (*Phoca hispida*) caught and released in New York (Luque, 2007). The description to follow visits a cached posterior after the estimation has been run. The details for running the estimation are found in Appendix B.

Figure 5.5 shows several Argos locations with corresponding estimated versions. The resulting estimates are individually rich in information and have unique distributions. These are discrete estimates representing the position of the animal at the time instant of the Argos fix. The full richness of the estimate is in the raw samples from the posterior, and these contoured estimates provide a useful summary. Each estimate can be interrogated for a mean or mode, standard errors or confidence intervals.

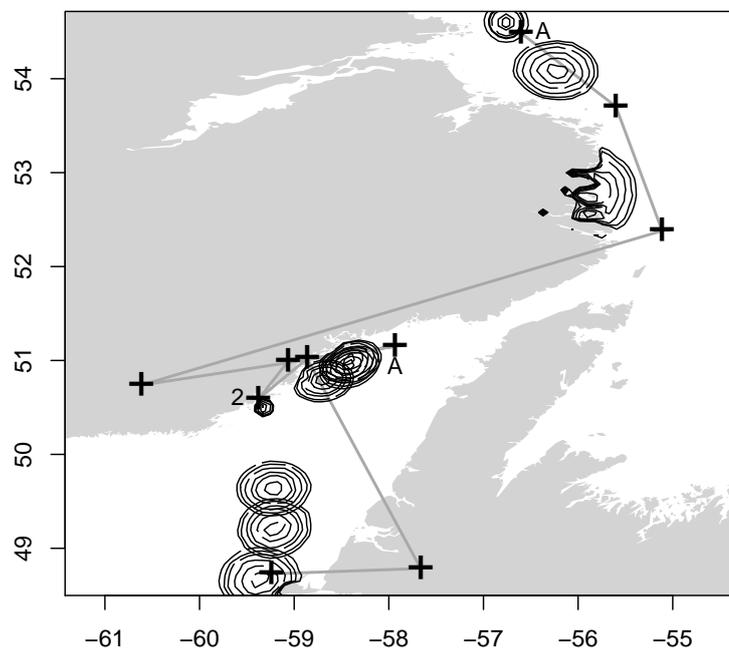


Figure 5.5: A short region of an Argos track of a ringed seal in Newfoundland. The primary estimates shown as contours are compared to the original Argos estimates. All but three Argos estimates have class “B” and are marked with “+”, the others are labelled with their classes “2” and “A”. The confidence in the “2” location is reflected by the proximity of the estimate to it and that it is relatively localized. None of the estimates fall on land like some of the Argos locations, and some are tightly bound by the land—but not clipped arbitrarily to it.

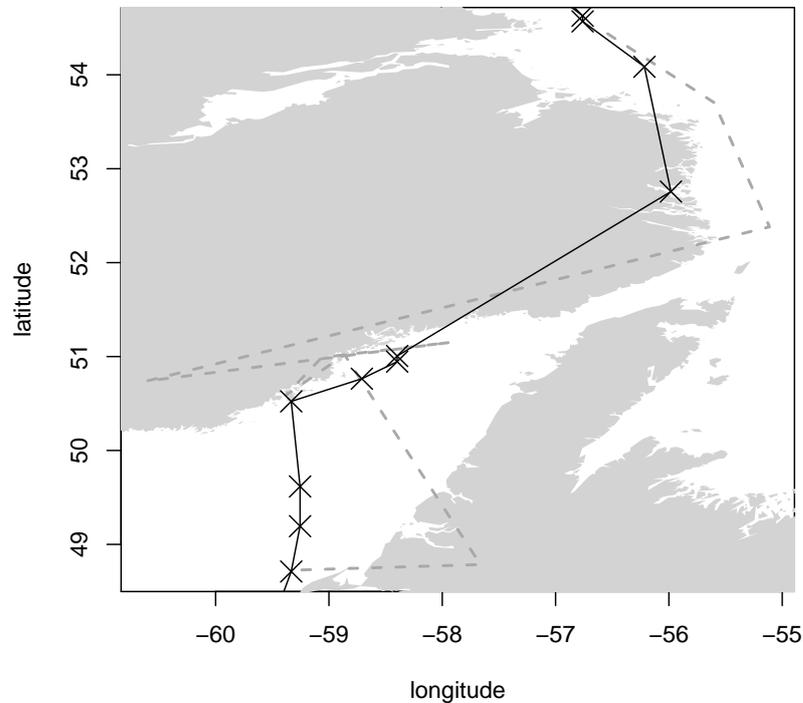


Figure 5.6: Original Argos track with sequence of modal locations from the posterior. The dark line joins “X” symbols derived as the mode from estimates shown in Figure 5.5. The modal locations are feasible for a ringed seal, but the line joining them introduces a problem by traversing the land area. The original line of the raw Argos track is shown as a dashed line.

A very simple summary is to calculate the modal point of each estimate—on the face of it this provides the “most probable track” but as can be seen in Figure 5.6 this is not necessarily a useful output as it simply ignores the continuous nature of a track, as discussed in Section 2.2.3. In Figure 5.6 the original Argos sequence is shown, as well as the modal locations for each primary estimate also joined by a line. This line crosses the land area, even though the estimates do not, which illustrates another problem with the practice of joining dots.

To more accurately represent these connected parts of the “full path” we interrogate samples from the posterior for the intermediate estimates. Individual intermediate estimates are shown in Figure 5.7 with the primary estimates shown as dark crosses. This figure is quite complicated but it does show the nature of the intermediate estimates in connecting the otherwise discrete primary estimates. The elongated intermediate through the strait shows the flexibility of representing complex estimates that can be multi-modal. The MCMC for such estimates may need to be run longer in order to ensure stability of the chain.

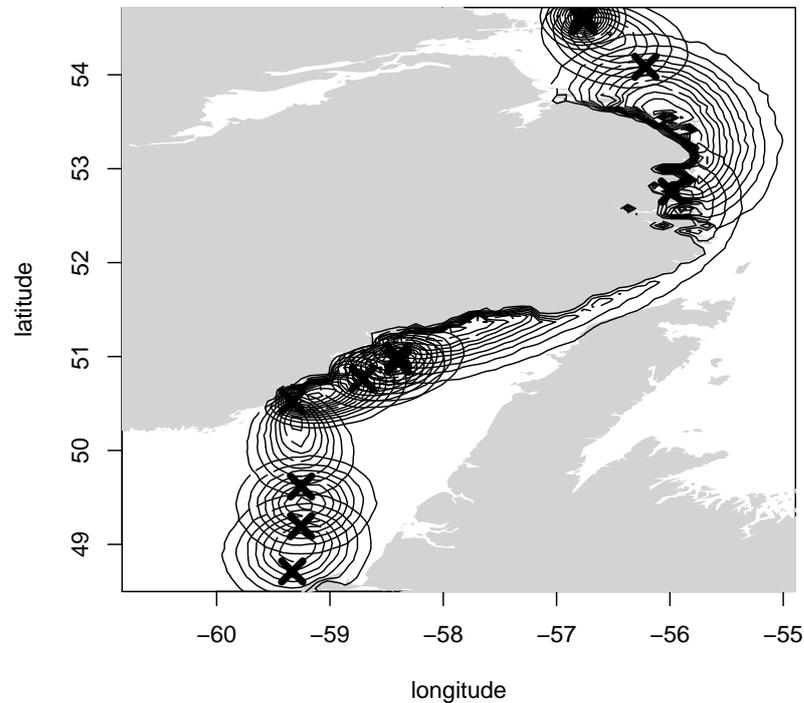


Figure 5.7: Individual intermediate estimates shown as contours with modal locations from the primary estimates shown as dark “X” symbols (as in Figure 5.6). The intermediate estimates overlap the primary estimates in space as well as connecting them in time.

The modal location of the primary estimate and the time spent estimate suggests the animal travelled very closely to the coastline of the Newfoundland and Labrador region north of the Island of Newfoundland—this may not be realistic and highlights that the input land boundary data may be too detailed. A very smooth boundary that less tightly hugs the real coastline could provide a more realistic constraint.

These individual estimates are treated discretely here, though they represent a continuous time-based process. A more natural way to present the intermediate estimates is as a surface estimate of time spent, which is shown in Figure 5.8. The likely path of the animal is seen through the more light-coloured region of the surface. The plot shows the entire distribution for each intermediate, these could be trimmed based on confidence intervals prior to the combination into a single surface. This figure is weighted by the time duration for each intermediate estimate in a way analogous to the time spent estimates presented in Chapter 2. For this particular example the only covariate data available is the time difference between fixes, and this has been used to calculate relative time spent per area. For different applications on other species this could be replaced by a different variable such as the

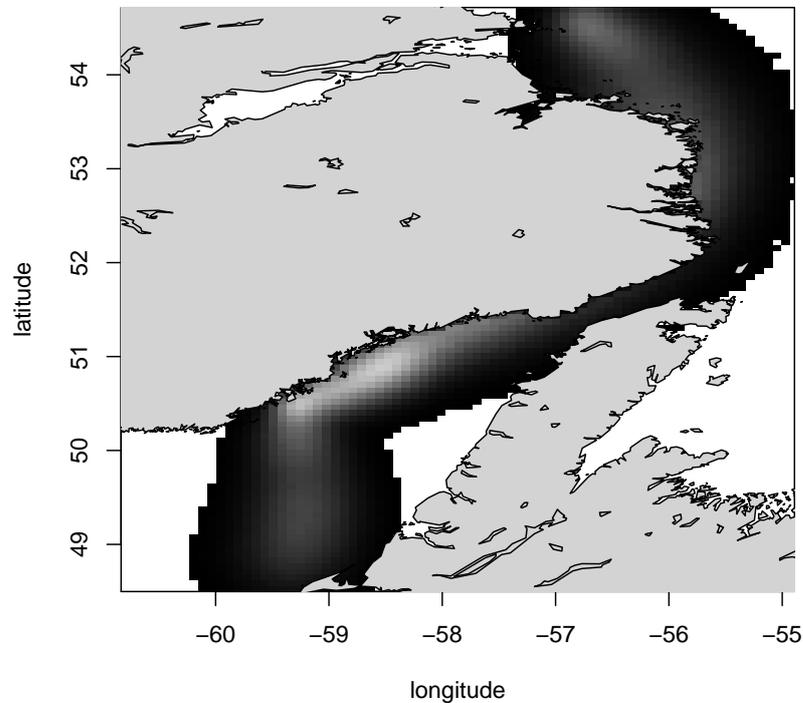


Figure 5.8: Time spent estimate from combination of the nine intermediate estimates. This single surface is constructed by combining the separate intermediate estimates for the time duration, each weighted by their corresponding time interval.

rate of drift dives (Thums et al., 2008), maximum depth, temperatures encountered or any relevant covariate of interest.

#### **Limitations of the primary and intermediate estimates**

The intermediate locations represent the interval between each primary location as a single block, with limited knowledge of their continuous connection through the primary estimates. They provide continuity in terms of the full path of the track, but not for the fine details between individual primary estimates. A more complex version, perhaps by introducing multiple intermediate locations, could reflect that at times close to the primary estimates the intermediates are closely matched to their primary counterparts. This limitation was also discussed in regards to Figure 3.4.

## 5.5 Conclusion

We have demonstrated techniques for efficient handling of summaries from a Bayesian approach to location estimation that provide the variety of outputs required by wildlife ecology. These include the most likely path, measures of uncertainty for individual locations, models of intermediate locations and estimates of residency that incorporate behavioural metrics such as time spent.

This chapter presented efficient methods for summarizing samples from the posterior of location estimation models. A new model for the representation of track data was demonstrated that is integrated with a general Bayesian approach to location estimation. The distinction between primary and intermediate location was illustrated with simulated and real-world data using a scheme for efficient and flexible summaries.

The next chapter deals with the use of environmental data as an integrated component of location models, presenting methods for applying data as a part of Bayesian models, and some examples demonstrating the potential for the application of subsurface data that have not previously been exploited.

## CHAPTER 6

# AUXILIARY DATA FOR LOCATION ESTIMATION

Recent collaborative studies have seen the use of diving animals as oceanographic platforms for measuring ocean properties (Thys et al., 2001; Boehlert et al., 2001; Sokolov et al., 2006). The location of the animal as it migrates is obviously a crucial requirement for oceanographic applications of this kind and for investigating the animal's habitat. However, as discussed in Chapters 3 and 4, these data have been used for location estimation for many years. Since the earliest marine tracking studies, data from archival tags have been used to help locate animals based on the relation of tag measured data to independent environmental data sets and models. Using these data for location estimation remains an important component of archival tag studies (Nielsen et al., 2006; Lam et al., 2008) and increasing knowledge of biophysical processes in the ocean provides opportunities for improvement and the development of new methods.

There is potential for a convergence of the investigation into the navigation and foraging strategies employed by marine animals with the understanding of the relation of movement to many properties of ocean circulation of interest to predators—current speed and direction, frontal activity, water bodies, basin topography and more. This work has concentrated on the estimation of location, but models of “where” need to integrate with models of “why” in a much tighter way, as discussed by Brillinger and Stewart (1998); Alerstam et al. (2001); Patterson et al. (2008) and Schick et al. (2008). While data collected using animals cannot be used both for oceanographic platforms and for estimating animal positions, exploratory analyses with at-depth data can provide insights and improvements for other applications.

Here we discuss the issues faced when matching tag measurements to synoptic environmental data and discuss some of the practical data access problems. We present two examples showing the potential for location information to be derived from subsurface temperatures and present the detail of the masking technique used in Chapters 3 and 4. As a function of expected data based on location, this component can be added to the estimation model. This is a feature of the general framework presented in Chapter 3, as a model of the data collection process.

## 6.1 A problem of scale

The main difficulty for exploiting auxiliary environmental data as a location data source is the sheer mismatch in scales—synoptic measures are aggregated and modelled based on relatively sparse samples by satellite or roving platforms. Synoptic products tend to be provided on time scales of a week (finer temporal scales do not have as complete coverage) and spatial scales of the order of many kilometres. These scales are much larger than the sampling scale of the tag itself, which can measure environmental properties directly within minutes and metres.

In Chapter 3, a method for integrating auxiliary data such as ocean temperature or maximum depth to the modelling framework was introduced. This method differs from more traditional approaches in that the restriction is not applied to existing location estimates by filtering or updating point estimates, but to provide an overall constraint that can augment deficiencies in primary location data. There are many ways in which this information can be incorporated using a Bayesian approach, but unfortunately these scale disparities become very important. As well as questions of how to match the underlying scales of the synoptic and tag measurements, it is not obvious how the natural time series of short twilight durations should be lined up with continuous temperature measurements. Decisions made to match one time scale (tag temperatures at seconds or hours to synoptic data on weekly or daily scales) are compounded with the need to incorporate the primary time scale of (for example) diurnal light variation.

## 6.2 Data availability

Oceanographic data are generally readily available, but can present difficulties in terms of accessibility from some platforms and in the handling of large, multi-dimensional data sets. The use of these data by animal movement models may be direct, with a large array of data stored directly in memory, or otherwise queried from database or file caches on demand. Both approaches are ideal for directly accessing “slabs” of data from an array that are then used for masking areas, querying model likelihoods, or specifying environmental correlates. However, for arbitrary queries of individual coordinates the commonly used tools can be inefficient. To query a 3D array database of environmental data with an animal track requires large data transfers of the entire section of the array that intersect. This problem is increased for 4D arrays and again for multiple variable data. This may not be problematic for a given track, but the methods like those used in Chapter 3 must deal with large numbers of proposal tracks used to sample the oceanographic data.

The use of oceanographic data and models as data sources for location estimation requires fast access to large data sets. Data interfaces generally use one of two types: direct manipulation of “slabs” of array data, or database-like lookup of exact samples queried by 3D or 4D track coordinates. In general terms, most applications have used slab-like access to do a small number of heavy computations to limit locations to a particular region, or to seek out a particular cell that has a better data match than an existing estimate. Sampling methods used in MCMC require a large number of small computations to be carried out in order to extract the data required for a

single iteration.

These problems are exacerbated by further requirements such as:

1. Individual tracks that span large distances and durations with many samples.
2. Studies involving multiple tracks over multiple seasons.
3. Model likelihoods involving several environmental variables.
4. The need for interpolation to exact coordinates within an array.

If the amount of data required is too large to be held in memory in a simple fashion the slab approach can break the analysis into a number of smaller chunks that can be carried out sequentially. In the database approach, the data access cannot be carried out sequentially as there may be temporal dependence in the model. For each sample, every section of data must be available so each individual time-step can find its matching data value. Traditionally, this problem is solved by using database techniques but, for a variety of reasons, large environmental data sets are rarely available in this way.

Most published examples provide their own specific method or workflow for accessing the environmental data, and matching it to the point or line scales used for the application. There are some notable projects that aim to cross this divide or provide database-linkages, such as the Rasdaman project, MGET and Spatial Analyst and STAT (Coyne and Godley, 2005)<sup>1</sup>.

There are three main types of access methods that are considered here, and each one assumes that accessing arbitrary portions of the data and that sampling or interpolating point values can be done efficiently.

**Access data as a slab.** This assumes that there is enough memory to work with the data as a whole.

**Compression for masks.** Large slabs of data can be compressed into bits for very efficient storage. This is useful when the data can be pre-processed into minimum/maximum ranges, as we discuss in Section 6.4.3.

**Database system.** This is the ideal, providing the entire data set as a slab with efficient overlay and access functions.

In practice the database system is simply not available, without significant effort in multiple software environments. In Section 6.5 we present an example that provides a powerful version of this approach using memory-mapped files in a single, readily accessible software environment.

Online servers providing access to data are an important component, but they still rely on client-side management of data as large slabs and do not provide any arbitrary access for point samples or interpolation. These include Reynolds online

---

<sup>1</sup>For Rasdaman see <http://rasdaman.com/>, for MGET see <http://code.env.duke.edu/projects/mget> and for Spatial Analyst see <http://www.spatial-analyst.net>

via **kfsst** (Nielsen et al., 2006), the PO.DAAC Ocean ESIP Tool (<http://poet.jpl.nasa.gov>) and various OpenDAP and related services (<http://www.opendap.org>, <http://www.tpac.org.au/main/>).

### 6.3 Slabs versus points

An obvious efficiency for data access is to only require the handling of the “region of interest” or, more generally, the specific overlap in time and space for the model being used. This is obvious in general terms, but has specific technology-based issues that keep it at the forefront of data usage for most practitioners. Even if this aspect is handled seamlessly, the next level provides a serious challenge that can be at odds with the optimizations applied for the “slab” problem.

Point samples provide a direct overlay of single coordinates with an arrayed data set. In its own right this is very simple, but include requirements for interpolation to exact coordinates and dynamic interaction and the common slab implementations are lacking. This is analogous to the deficiencies of topological data identified in Chapter 2—decisions or habits established at one level have serious ramifications for the simplest next level of generalization. The marriage of spatial and temporal data analysis is still plagued by these problems and it is a topic of ongoing development (Beegle-Krause et al., 2010)<sup>2</sup>.

### 6.4 Temperature applications

Here we present a number of issues and future directions for the use of temperature as a location source.

Surface temperature has been used by many studies to improve the latitude component of light level geo-location estimates, as discussed in Section 4.1. Many papers now cite a common method of determining location from traditional light level geo-location, with subsequent correction of the latitude component by temperature matching (see Nielsen et al. (2006) for an overview). The framework presented in Chapters 3 and 4 showed that a more systematic approach can utilize more of the light data while simultaneously incorporating environmental data. Much of the data collected by archival tags is currently under-utilized and could be enhanced by comparison with ocean subsurface properties. The prospects for subsurface applications have previous been discussed by Nielsen et al. (2006), Anderson et al. (2007) and Lam et al. (2008).

Many studies have used surface temperature-matching for improving locations, but there is very little standardization. As discussed by Lam et al. (2008) the first question is “Which product is most appropriate?”, and besides the availability of data for the region of interest there are many, often arbitrary, choices required for matching the spatial and temporal scales involved. Some studies have explored more sophisticated relations between animal migration and gradients in ocean properties

---

<sup>2</sup>See relevant, albeit very different, projects at <http://opengeostatistics.org/> and <http://www.eonfusion.com>

(Gaspar et al., 2006; Campagna et al., 2006) and with bottom topography and tidal patterns (Gröger et al., 2007; Pedersen et al., 2008). There is a need for generalization of the problems involved and exploration of new models and data sets as they become available.

Subsurface temperature data have been little used though some early works have published examples. Smith and Goodman (1986) discussed the prospects for latitude determination by depth-specific temperatures and Hindell et al. (1991) published estimates of the minimum regions that could have been visited based on temperature at-depth data. The work by Hindell et al. (1991) manually compared potential locations to maps of temperature at-depth, similar to the masking approach described later in Section 6.4.3. The scale mismatch for tag data versus synoptic environmental data means that a masking approach can provide some arguably safer assumptions. Below we present examples to show the use of environmental data as masks and the potential for using subsurface temperature profiles.

Archival tag temperatures also tend to suffer from a slow response time in which the sensor must stabilize at a constant temperature for some time before the value can be considered reliable (Boehlert et al., 2001). To account for this problem Sokolov et al. (2006) apply a correction by modelling the temperature response of the unit and Shaffer et al. (2005) select water temperatures that are stable for at least 20 minutes. Modern CTD-level tags do not suffer from this issue, and the availability of these modern tags improves the opportunity for marine animals to act as oceanographic-sensing platforms. The issue of matching spatial and temporal scales remains.

Below we present an example of using sea surface height (SSH) as a location source from a proxy relationship of at-depth temperatures from archival tags.

#### 6.4.1 Temperature profiles

The vertical temperature and salinity structure of water masses in the Southern Ocean is very stable, with a very coherent spatial pattern in the meridional (north/south) axis (Sokolov and Rintoul, 2002). Very generally, the geographic pattern at the surface is of strong meridional temperature and salinity gradients. The Southern Ocean is dominated by zonally (east/west) oriented fronts which are relatively narrow regions of fast eastward flowing currents. At each of these fronts are strong meridional gradients of temperature and salinity that extend to great depths. These fronts, identified as sharp increases in SSH towards the north, meander north and southward by up to 100 km periodically spinning off eddy rings.

Climatologies and models of the ocean can provide full four-dimensional representations of ocean currents with a resolution that intersects the spatial and temporal scales of interest to the location estimation of wide ranging marine animals. Elephant seals can dive to depths of at least 2000m, and the majority of their time is spent in water masses deeper than the mixed layer. For species of tuna the deepest dive is not as extreme, but the time at depth is more consistent.

The 4-D Gravest Empirical Mode (GEM) model takes advantage of the highly coherent vertical structure of water masses in the Antarctic Circumpolar Current

(ACC) to produce a very strong empirical relationship between subsurface temperature/salinity profiles and their surface expression of SSH at any given longitude (Meijers et al., 2010). As the SSH is observable via satellite altimetry at relatively high spatial and temporal resolutions, the observed SSH values can be used to construct highly accurate estimates of the subsurface temperature and salinity. The altimetry is provided on a  $1/3^\circ$  Mercator grid at weekly intervals, allowing the resolution of oceanic features around 50-100 km in size. This resolution limit, along with other factors, means GEM temperature estimates have RMS errors of around  $1^\circ\text{C}$  in the mixed layer, decreasing to below  $0.2^\circ\text{C}$  RMS below about 1000 m. This error approaches the *a priori* noise in the Southern Ocean, as determined from *in situ* hydrography.

Figure 6.1 presents two small regions from a single elephant seal track with nearby cell points from the 4-D GEM model (Meijers et al., 2010). The corresponding temperature and salinity plots are the model temperatures over depth for these cells and the line is the summary CTD temperatures measured by the diving seal. The region for each plot is based on the available Argos location data. We can be confident that the GEM sample points cover the region containing the true position of the seal for the profile as this region is significantly larger than the underlying Argos error. The correspondence between the GEM profiles and the CTD temperature and salinity values is quite strong, with quite different patterns in the two regions. Within localized regions of the Southern Ocean the pattern of temperature and salinity over depth is consistent.

The combination of temperature and salinity provides a region restricted in longitude as well as latitude, provide a potentially far more informative location source than SST alone. However, the sheer size of the GEM data set presents substantial challenges for its use in geo-location studies. See Section 6.5 for a description of accessing this large data set on low-end desktop computers.

### 6.4.2 Ocean height proxy

As discussed in Section 6.4, traditional archival tags can suffer measurement lags that obscure the temperature signal. Newer tags offer CTD-quality measurements (Boehme et al., 2009), but there is a rich long-term record of traditional archival tags for many marine species. Also, subsurface temperatures from older archival tags will often have durations at which the animal stays at a constant depth, allowing for the sensor to stabilize to a constant temperature. Relating these more reliable temperatures to other synoptic products provides another promising source of location information.

The GEM parameterization provides a relationship between the surface ocean height and the underlying temperature profile of the Southern Ocean. This is based on the relation of temperature at a given depth and an integral of the density field (Sokolov et al., 2006). By applying this relationship in reverse it is possible to predict the surface height from temperature at-depth<sup>3</sup>. This was used to generate spatial fields of the difference between predicted ocean height and satellite-derived ocean height. These provide a source of location, primarily for latitude but with some

---

<sup>3</sup>Example code to apply this was generously provided by Sergei Sokolov from the CSIRO.

deldir 0.0-13

Please note: The process for determining duplicated points has changed from that used in version 0.0-9 (and previously).

spatstat 1.22-0

Type 'help(spatstat)' for an overview of spatstat

'latest.news()' for news on latest version

'licence.polygons()' for licence information on polygon calculations

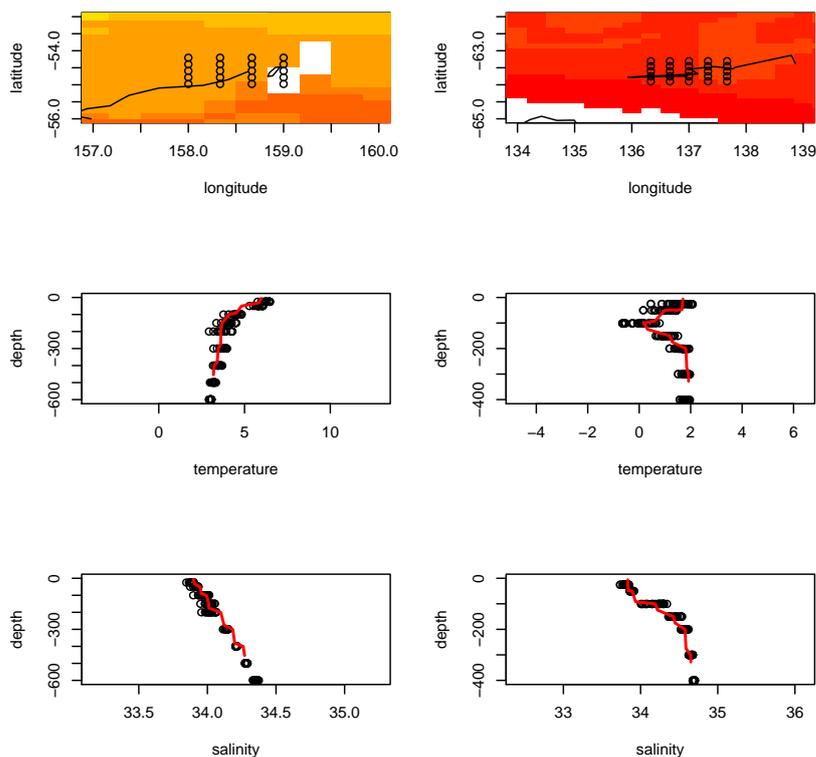


Figure 6.1: Temperature and salinity profiles from two regions for an elephant seal track. Seal dive profile values are shown as a red line. The points are all data values from the profiles taken from the GEM model for the corresponding seal depths at the grid of locations shown. The first is near Macquarie Island, the second near Adélie Land.

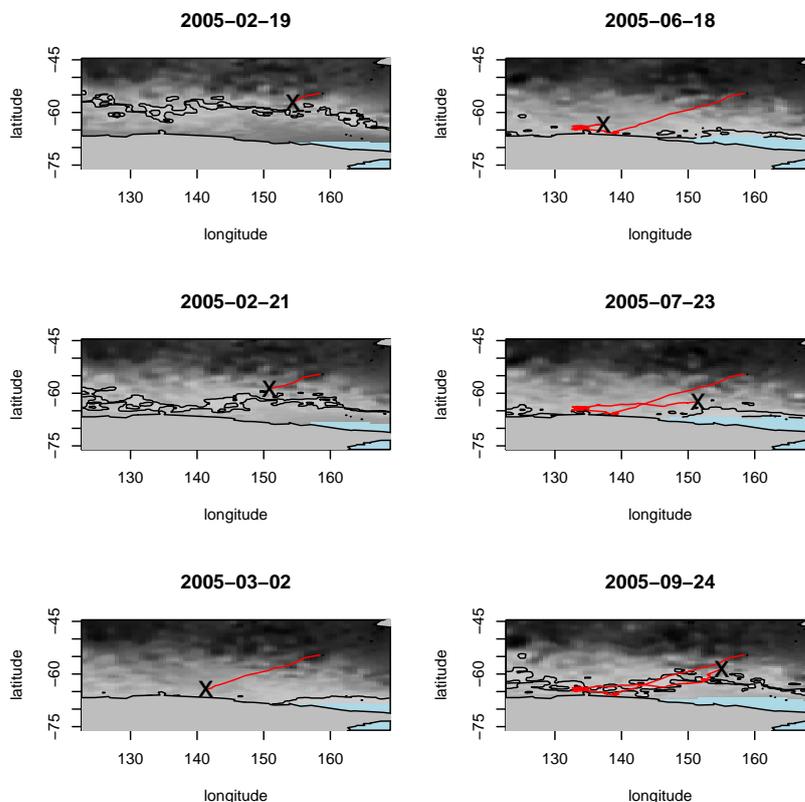


Figure 6.2: Series of location likelihood maps based on SSH proxy from temperature at-depth.

meridional variation based on frontal meanderings as described in Section 6.4.1.

A Mean Sea Level Anomaly gridded data set is available as a complete-coverage product, is not susceptible to cloud cover and is available at weekly intervals since 1996 on a  $1/3^\circ$  Mercator grid<sup>4</sup>. Maps of absolute SSH were produced using the same methods as Sokolov et al. (2006).

In Figure 6.2 a selection of these fields is shown as the absolute difference between predicted and measured ocean height with the corresponding Argos track and current Argos position. The likely regions of the field are highlighted with contours which neatly line up with the concurrent Argos location, and the general movement of the trip.

### 6.4.3 Masks

Here we outline the process for generating a time series of masks for archival tag temperatures compared to an SST data set for input to a location estimation model.

<sup>4</sup>This is available from [http://www.marine.csiro.au/dods/nph-dods/dods-data/sat\\_alt\\_msla/](http://www.marine.csiro.au/dods/nph-dods/dods-data/sat_alt_msla/)

This requires the mask to be organized to match the time intervals of the model locations and since it is a binary mask, the final data object can be substantially compressed.

The first step requires that the temporal scales of the SST data be matched to that of the archival data. Archival tag temperatures restricted to near-surface values are summarized into a minimum and maximum value for all temperatures measured for six hours either side of the primary time. (There are many possible strategies here.)

By converting the original temperature data to a series of masks, each pixel can be stored as a single bit, providing substantial compression from the temperature data set as a whole. This is true even though many more masks are created than there are time slices in the original data. This technique can then be used on even low-performance computers. An example is presented in the manual page for the function `get.mask` in the **tripEstimation** package (Sumner and Wotherspoon, 2010).

The final step is to create a lookup that can query the mask for sampled points for the primary times. This is a function that takes a set of locations, a coordinate at each time step, and returns a mask value for each. This function is the input mechanism for applying auxiliary environmental data that was outlined in Chapter 3.

A lookup function for the mask returns a boolean value for each location, which is readily incorporated into the modelling framework presented in Chapter 3. A mask lookup function that had to deal with the entire set of SST data for region would need to do the comparison to minimum/maximum values each time in order to return the appropriate decision for a given location. This scheme provides a very fast lookup that uses a very small memory footprint. Another advantage of the lookup function is that it can perform other necessary tasks, for example reprojecting a sample of points to match the environmental data potentially avoiding the need to warp a large array to the model coordinate system.

These masks could be extended to use more of the subsurface data, either for improved temperatures for older archival tags or to integrate the information at multiple depths for reliable temperature profiles from CTD tags. No matter what information a practitioner chooses to include in the creation of the mask, the final product for the simulation model is a lookup function, the details of which can be hidden from the running of the model itself.

Finally, in general terms these functions may sample actual data values from the original data for more sophisticated masks or likelihood models on environmental data. The memory requirements for this kind of application is much greater than a mask and this is increased for samples required by simulation models, for interpolating into cells and for higher dimensional applications.

## 6.5 Large data set example

We describe a simple approach implemented in freely available tools that can be used for fast arbitrary access to very large time series of oceanographic data. The

example uses the 4-D GEM climatology described in Section 6.4.1.

The GEM data set consists of four 4-D (longitude, latitude, depth, time) arrays for temperature, salinity, and current vectors totalling about 150Gb of double-precision floating point values. The data traverses the entire meridional extent of the Southern Ocean, between latitudes -70 and -34.74. Depth extends from the surface to 5400m in 36 steps. The temporal range of the data is from 4 October 1992 to 10 September 2006 in weekly steps. Within the ocean region the data is relatively complete with few gaps, although horizontal coverage reduces at the deepest areas due to the ocean floor topography, and changes over time. The latitude and vertical axes are not regular, though the geographic space can be interpreted in a regular grid by using the Mercator projection. The vertical step length increases with increasing depth. Land areas and missing data are identified with missing values.

These data were converted to a generic form to provide arbitrary access tools that could overcome some memory limitations. The files were converted in Octave from Matlab export format to generic binary arrays with a reversed index order<sup>5</sup>. These were then used to populate memory-mapped files controlled by the R package `ff` (Adler et al., 2010).

The final system consists of three memory-mapped objects for each variable, a read function to generate new objects as a single array from subsets to the entire 4D data and lookup functions to sample point values from the field. Using this system the entire GEM data set can be made accessible and large subsets of it (up to 16Gb each) manipulated in a single object for efficient and simple use with track data. The data for Figure 6.1 was prepared using this system, and via the lookup functions these data can be made available to complex location models such as the examples presented in Chapters 3 and 4.

## 6.6 Conclusion

We have demonstrated that there is far greater potential for the incorporation of large environmental data sets to location models than has been previously explored. Computing tools and environments must be carefully chosen to accommodate the requirements for these large-scale data sets. We have shown methods for accessing a very large oceanographic model that can be deployed on relatively low-end desktop computers, demonstrating that modern statistical techniques for location estimation can incorporate large and complex data sets.

---

<sup>5</sup>This orientation is not necessary for the final memory-mapped system, but allowed for simple georeferencing of the raw files for simple and efficient conversion to GIS formats via the GDAL virtual raster format (Warmerdam and the GDAL development team, 2010). Also, there are R packages that provide reading of Matlab files, but installation details and machine limitations meant that this was not feasible.

## CHAPTER 7

# CONCLUSION

Location estimation for animal tracking is a developing field with a long history of contributions for many different applications and techniques. It is impossible to classify all of the contributions that have been made, and the great variety of issues dealt with by different applications. This thesis presents a unification of many disparate issues in animal tracking analysis focussed on the use of large raw data sets.

Chapter 2 surveyed a range of problems in tracking data regarding the accuracy of location estimates and the representation of track data. Uncertainties in location are compounded by inconsistencies with track representations and the lack of integrated software tools for analysing and exploring different methods. Seemingly simple things such as track via points or lines introduce assumptions and requirements that are not commonly dealt with explicitly, such as the way track data are manipulated to produce summaries of residency or time spent. The need for more systematic representations of track data and location error was demonstrated with examples, including easily available software. The **trip** package provides an integrated software environment for employing traditional methods. Track data are automatically validated as part of a formal system in order to avoid common problems, providing access to a huge range of tools required for dealing with spatial and temporal data. Also demonstrated was the need for wider use of geographic map projections for the representation of tracks and for simulation studies. Access to software tools providing efficient conversion between map projections is also made available by **trip**. Chapter 2 also presented traditional methods for track data with modern software tools, providing much needed access to exploratory analyses in a single environment. The issues discussed regarding traditional methods were then used to put a new perspective on track estimation from data sources that are not inherently spatial, providing the context for a general framework for location estimation.

Chapter 3 provided a general framework for location estimation using Bayesian methods. The approach provides a broad classification of sources of location information that can be used for a variety of location methods. These sources are prior knowledge, primary location data, auxiliary location data and movement models. The framework includes a model for track representation that explicitly differentiates primary locations from intermediate locations. This distinction bridges the point and line track representations for traditional analyses and provides all the

required metrics and measures for modelling animal movement. The framework was applied to two example data sets: Argos Service locations and measurements from an archival tag, demonstrating the generality of the approach. The full detail of the light level geo-location example was omitted here in order to focus on the general applicability of the framework.

Chapter 4 focussed on the full detail for the light level geo-location example of Chapter 3. This provides a novel method for determining location from light level by relating measured light to solar elevation. Extensible and freely available software for running light level geo-location is provided to enable the application of this approach in the R package **tripEstimation**.

Chapter 5 expanded on the track representation model presented in Chapter 3. This includes an integrated data structure design and binning mechanism for efficient visualization and analysis of individual location estimates or full-path estimates with in-built measures of uncertainty. The data structures can be easily used for generating residency or time spent estimates from groups of individuals and visualizing, or otherwise quantifying, changing patterns of spatial residency over time. These provide a powerful mechanism for comparing spatial estimates consisting of probability distributions with data sets of environmental covariates. Common problems with requirements for data resampling by disparate scales or different map projections is reduced. For example, by binning samples from the posterior to match an environmental data set, otherwise required warping of gridded data can be avoided.

Chapter 6 illustrated the potential for using subsurface ocean properties for estimating location for marine animals. There is a largely unrecognized and unrealized potential for the use of both surface and subsurface data and there are challenging issues with scale and sampling disparities, data error problems and efficient access methods to large data sets. We have demonstrated methods for exploiting subsurface temperatures and salinity with large multi-dimensional ocean models for inclusion in statistical models.

For the first time the wide range of data analysis techniques required for animal tracking have been unified. This approach provides consistent data models for spatial and temporal data, location estimation, track representation and flexibility for exploring new methods.

## APPENDIX A

# MARKOV CHAIN MONTE CARLO

The computational challenge in applying Bayesian techniques is the evaluation of complex integrals required for practical inference. Markov Chain Monte Carlo (MCMC) techniques provide a simple and generic solution to this problem.

It is trivial to determine the posterior density to within a multiplicative constant, as by Bayes' rule (3.1)

$$p(\theta | y) \propto p(y | \theta)p(\theta).$$

But any practical inference, such as the calculation of expected values and quantiles, requires the calculation of the normalization constant

$$\int p(y | \theta)p(\theta)d\theta,$$

and other associated integrals of the posterior distribution. Unfortunately, the evaluation of the required integrals is often computationally demanding.

MCMC techniques provide a generic and simple approach to this problem through simulation. MCMC methods allow samples from the posterior distribution when the posterior density is known only to within a multiplicative constant. Any properties of the posterior can then be approximated by the properties of the sample—in essence, evaluating the complex integrals required for inference by Monte Carlo quadrature.

### A.1 Metropolis Hastings

The Metropolis Hastings (MH) algorithm is the most generic form of MCMC. Given a target distribution  $p$  from which we wish to draw samples, the Metropolis Hastings algorithm constructs a Markov Chain that has  $p$  as its stationary distribution.

The algorithm is surprisingly simple. To generate a sequence  $X_1, X_2, X_3 \dots$  of draws from  $p$ , at each stage  $i$  a candidate point  $Y$  is drawn from a proposal distribution  $q(\cdot | X_i)$ . The candidate is accepted with probability  $\alpha(Y, X_i)$ , where

$$\alpha(Y, X) = \min \left( 1, \frac{p(Y)q(X | Y)}{p(X)q(Y | X)} \right).$$

If the candidate is accepted, then  $X_{i+1} = Y$ , otherwise  $X_{i+1} = X_i$ .

The sample  $\{X_1, X_2, X_3 \dots\}$  constructed in this way will have distribution  $p$ , and so any property of  $p$  can be approximated from the sample.

The key feature of the Metropolis Hastings algorithm is that  $p$  occurs in both the numerator and denominator of  $\alpha$ , and so  $p$  need only be known to within a multiplicative constant. This is the reason the MH algorithm is so useful for Bayesian inference.

The choices of the initial point  $X_1$  and proposal distribution  $q(\cdot | \cdot)$  are effectively arbitrary, but poor choices impact efficiency. A poor choice of  $X_1$  can result in points in the neighbourhood of  $X_1$  being over-represented in the sample. For this reason it is common to discard the initial segment of the chain to reduce the dependence on  $X_1$ . A poor choice of  $q(\cdot | \cdot)$  will result in fewer candidates being accepted, and it may become necessary to draw a very large sample to obtain accurate results.

More details of the properties of the Metropolis Hastings algorithm can be found in Gilks et al. (1995).

## A.2 Example

To illustrate the Metropolis Hastings algorithm, consider a simple Binomial example. Suppose  $y$  heads are observed from  $n$  tosses of a coin, and interest lies in estimating  $\theta$ , the probability of throwing a head in a single toss.

Adopt a uniform prior for  $\theta$

$$p(\theta) = 1 \quad 0 \leq \theta \leq 1,$$

so that before any data is observed, all values of  $\theta$  in the interval  $[0, 1]$  are believed equally reasonable.

The number of heads  $y$  is Binomially distributed

$$y \sim \text{Bin}(n, \theta)$$

so the likelihood  $p(y | \theta)$  is simply the Binomial mass function

$$p(y | \theta) = \binom{n}{y} \theta^y (1 - \theta)^{n-y}.$$

By Bayes's rule, the posterior density is

$$\begin{aligned} p(\theta | y) &= \frac{\binom{n}{y} \theta^y (1 - \theta)^{n-y}}{\int_0^1 \binom{n}{y} \theta^y (1 - \theta)^{n-y} d\theta} \\ &= \frac{\Gamma(n + 2) \theta^y (1 - \theta)^{n-y}}{\Gamma(y + 1) \Gamma(n - y + 1)}. \end{aligned}$$

Aficionados will recognise this as the density of a Beta distribution, and that the posterior distribution is

$$\theta | y \sim \text{Beta}(y + 1, n - y + 1).$$

The posterior distribution encapsulates all knowledge of  $\theta$  given the observed data  $y$ , and can be summarized to yield a number of quantities of interest. For example, the posterior mean and variance of  $\theta$ —the mean and variance of  $\theta$  given the observations—are given by

$$\begin{aligned} \text{E}(\theta | y) &= \int_0^1 \theta p(\theta | y) d\theta \\ &= \frac{\Gamma(n + 2)}{\Gamma(y + 1)\Gamma(n - y + 1)} \int_0^1 \theta^{y+1}(1 - \theta)^{n-y} d\theta \\ &= \frac{y + 1}{n + 2}, \\ \text{Var}(\theta | y) &= \int_0^1 (\theta - \text{E}(\theta | y))^2 p(\theta | y) d\theta \\ &= \frac{\Gamma(n + 2)}{\Gamma(y + 1)\Gamma(n - y + 1)} \int_0^1 \left(\theta - \frac{y + 1}{n + 2}\right)^2 \theta^y(1 - \theta)^{n-y} d\theta \\ &= \frac{(y + 1)(n - y + 1)}{(n + 2)^2(n + 3)}. \end{aligned}$$

The posterior quantile  $\theta_\alpha$  is given by the solution of the equation

$$\begin{aligned} \alpha = \text{Pr}(\theta < \theta_\alpha | y) &= \int_0^{\theta_\alpha} p(\theta | y) d\theta \\ &= \int_0^{\theta_\alpha} \frac{\Gamma(n + 2)\theta^y(1 - \theta)^{n-y}}{\Gamma(y + 1)\Gamma(n - y + 1)} d\theta \end{aligned}$$

and  $\theta_{0.025}$  and  $\theta_{0.975}$  form an exact confidence interval for  $\theta$ .

To determine the probability that  $\theta$  lies within the interval  $[0.4, 0.6]$ , that is, the probability that the coin is *roughly* fair, we calculate

$$\text{Pr}(0.4 < \theta < 0.6 | y) = \int_{0.4}^{0.6} p(\theta | y) d\theta = \int_{0.4}^{0.6} \frac{\Gamma(n + 2)\theta^y(1 - \theta)^{n-y}}{\Gamma(y + 1)\Gamma(n - y + 1)} d\theta.$$

To apply the Metropolis Hastings algorithm to this problem, it suffices to determine the posterior density to within a constant of proportionality independent of  $\theta$

$$p(\theta | y) \propto p(y | \theta)p(\theta) \propto \theta^y(1 - \theta)^{n-y}.$$

Choosing a Normal distribution centred on the current point as the proposal distribution

$$q(Y | X) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-(X-Y)^2/(2\sigma^2)}$$

the Metropolis Hastings acceptance probability becomes

$$\alpha(X, Y) = \begin{cases} \min\left(1, \frac{Y^y(1-Y)^{n-y}}{X^y(1-X)^{n-y}}\right) & \text{if } 0 \leq Y \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

To generate a sequence of draws  $X_1, X_2, \dots$  from the posterior  $p(\theta | y)$ , we simply draw proposal points from  $q(Y | X_i)$ , accepting  $Y$  with probability  $\alpha(X_i, Y)$ .

Figure A.1 show histograms of a sample of 500 draws and a sample of 5000 draws from the posterior density when  $n = 20$  and  $y = 12$ , together with the true posterior density. As the number of draws increases, the histogram estimates increasingly approximate the posterior density.

Similarly, any posterior quantity can be approximated directly from the sample. Table A.1 show both the exact values and sample estimates for the posterior mean, variance, median, 2.5% and 97.5% quantiles and the probability that  $\theta$  lies within the interval  $[0.4, 0.6]$ . Each of these approximations improve as the number of draws  $N$  from the posterior increases.

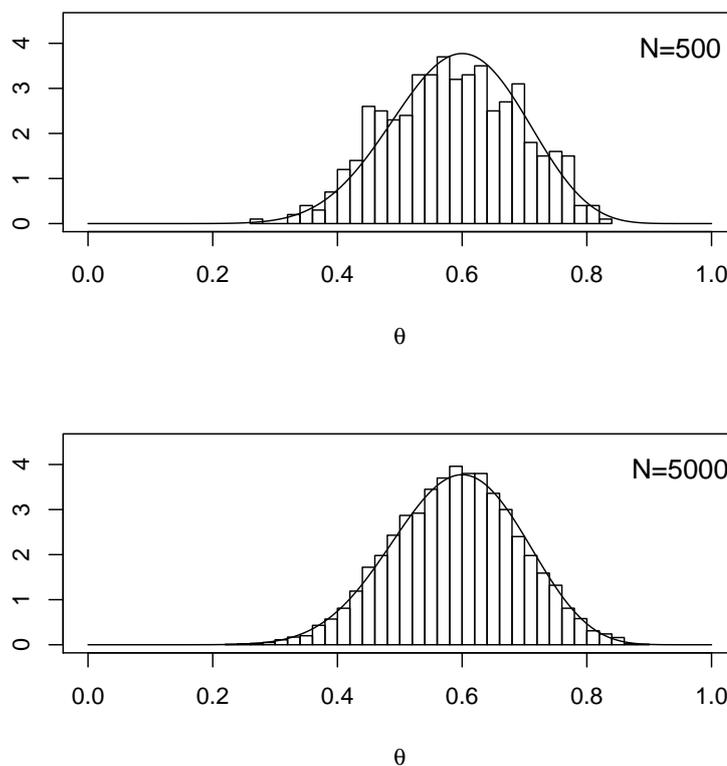


Figure A.1: Two samples drawn by the Metropolis Hastings algorithm for the Binomial example with  $n = 20$  and  $y = 12$ . The plots show a histogram of  $N$  MCMC samples and the exact posterior density, for  $N = 500$  and  $N = 5000$ .

	Exact	N=500	N=5000
mean	0.5909	0.584	0.5894
variance	0.01051	0.01111	0.01023
median	0.5937	0.5827	0.5916
$\theta_{0.025}$	0.3844	0.3899	0.3892
$\theta_{0.975}$	0.7818	0.769	0.782
$\Pr(0.4 < \theta < 0.6   y)$	0.4885	0.518	0.5006

Table A.1: Comparison of exact and MCMC approximations for  $N = 500$  samples and  $N = 5000$  samples for the Binomial example with  $n = 20$  and  $y = 12$ .

## APPENDIX B

# ARGOS DATA ESTIMATION EXAMPLE

### B.1 Argos data example

```
library(tripEstimation)
library(diveMove)
library(trip)
##library(mapdata)
library(rgdal)
library(maptools)
## prepare source data
##locs <- readLocs(system.file(file.path("data", "sealLocs.csv"),
##                             package="diveMove"), idCol=1, dateCol=2,
##                             dtformat="%Y-%m-%d %H:%M:%S", classCol=3,
##                             lonCol=4, latCol=5)

locs <- readLocs(system.file(file.path("data", "sealLocs.csv"),
                             package="diveMove"), idCol=1, dateCol=2,
                             dtformat="%Y-%m-%d %H:%M:%S", classCol=3,
                             lonCol=4, latCol=5, sep=";")
ringy <- subset(locs, id == "ringy" & !is.na(lon) & !is.na(lat))
coordinates(ringy) <- ~lon+lat
tr <- trip(ringy, c("time", "id"))
##mm <- map("worldHires", xlim = bbox(tr)[1,], ylim = bbox(tr)[2,])

## Auxiliary environmental data

## create a simple mask for land/sea from polygon data set
data(wrld_simpl)
coast <- c(grep("Canada", wrld_simpl$NAME),
           grep("United States", wrld_simpl$NAME))
## US and Canada only
```

```

wrlld_simpl <- wrld_simpl[coast, ]
##plot(mm, xlim = bbox(tr)[1,], ylim = bbox(tr)[2,])

bb <- bbox(tr)
xx <- seq(bb[1,1] - 1, bb[1,2]+ 1, length = 200)
yy <- seq(bb[2,1] - 1, bb[2,2]+ 1, length = 250)
res <- overlay(SpatialPoints(expand.grid(x = xx, y = yy)), wrld_simpl)
## the topo mask
topo <- list(x = xx,
             y = yy,
             z = matrix(is.na(res), length(xx), length(yy)))
rm(res)
## lookup function
lookup <- mkLookup(topo)
## set up the proposal functions
xy <- coordinates(tr)
m <- nrow(xy)
## Choose parameters for proposals appropriately
## m proposals one for each time, 2 params (lon, lat), sigma for each
proposal.x <- norm.proposal(m, 2, c(0.1, 0.1))
## m-1 proposals for each intermediate, 2 params, sigma for each
proposal.z <- norm.proposal(m-1, 2, c(0.2, 0.2))
## interactive session to see if choices are "reasonable" - we want to
## jump around the neighbourhood, but not too much - later we will
## tune the proposals to aid mixing, but the choice now helps the MCMC
## initialize efficiently, especially with respect to the mask The Zs
## need greater proposals as they are generally more varied

plot(xy)
## repeat this line a few times to see what the spread is like
points(proposal.x$proposal(coordinates(xy)))
plot((xy[-m,1:2]+xy[-1,1:2])/2)
## repeat this line a few times to see what the spread is like
points(proposal.x$proposal((xy[-m,1:2]+xy[-1,1:2])/2))
## the choice of behav.mean/sd here is based on building this
## distribution, I'm assuming that the Argos locations are reasonably
## representative (in the absence of more information)

## km/hr
spd <- trackDistance(coordinates(tr), longlat = TRUE) /
      diff(unclass(tr$time) / 3600)
x <- seq(0, max(spd) * 1.5, by = 0.1)
op <- par(mfrow = c(2, 1))
plot(x, dnorm(x, mad(spd), sd(spd)),
      main = "likelihood of speed drops off fast with higher values")
plot(x, dnorm(x, mad(spd), sd(spd), log = TRUE),
      main = "log normal as applied in the model")
par(op)

```

```

## choose these values as required, or rewrite the behavioural
## function in the model to suit

speed.mu <- mean(spdx)
speed.sd <- sd(spdx)
tr$class <- ordered(tr$class,
                    levels = c("Z", "B", "A", "0", "1", "2", "3"))
argos.sd <- argos.sigma(tr$class, sigma = c(100, 80, 50, 20, 10, 8, 4),
                      adjust = 111.12)
## build the model
## Argos times and locations
## proposal functions for X and Z
## mask functions, and whether release/recapture are fixed
##   (which helps pin down the start and end)
## starting positions (Zs are intermediate to the Xs)
## Argos position sigma (can pass in a single value, or a value
##   based on the class)
## behavioural model function (defined for lognormal or gamma)
## mean and sd for behavioural model

d.model <- satellite.model(tr$time, xy,
                          proposal.x$proposal, proposal.z$proposal,
                          mask.x = lookup, mask.z = lookup,
                          fix.release = FALSE, fix.recapture = FALSE,
                          start.x = xy, start.z = (xy[-m,1:2]+xy[-1,1:2])/2,
                          posn.sigma = argos.sd, behav = "log",
                          behav.mean = speed.mu, behav.sd = speed.sd)
## Run initial chain to obtain values that meet mask constraints
ch <- metropolis0(d.model, iters=10, thin=10, start.x = d.model$start.x,
                 start.z = d.model$start.z)
while(!(all(ch$mask.x) && all(ch$mask.z))) {
  ch <- metropolis0(d.model, iters=100, thin=10,
                  start.x = ch$last.x, start.z = ch$last.z)
  plot(ch$last.x, type = "l")
  plot(wrld_simpl, add = TRUE)
}
## run a short period, and check
ch <- metropolis(d.model, iters=200, thin=10,
                start.x=ch$last.x,
                start.z=ch$last.z)
plot(tr, pch = 1, col = "grey")
lines(coordinates(tr), col = "grey")
plot(wrld_simpl, add = TRUE)
lines(ch$last.x)
points(ch$last.x, col = "red", pch = 21, cex = 0.5)
## run for a while to settle in
for (i in 1:3) {

```

```

    ch <- metropolis(d.model, iters=2000, thin=10,
                    start.x=ch$last.x,
                    start.z=ch$last.z)

}
plot(tr, pch = 1, col = "grey")
plot(wrld_simpl, add = TRUE, col = "grey")
apply(ch$z, 3, points, pch = ".")
lines(coordinates(tr), col = "grey")
## run and tune the proposals (hopefully enough to allow for burn-in)
for (i in 1:5) {
  ch <- metropolis(d.model, iters=2000, thin=10,
                  start.x=ch$last.x,
                  start.z=ch$last.z)

  proposal.x$tune(ch$x, scale = 0.3)
  proposal.z$tune(ch$z, scale = 0.3)

}
## run and save the results to disk, tuning as we go

xfile <- "X0.bin"
zfile <- "Z0.bin"
for (i in 1:15) {
  ch <- metropolis(d.model, iters=2000, thin=10,
                  start.x=ch$last.x,
                  start.z=ch$last.z)

  proposal.x$tune(ch$x, scale = 0.3)
  proposal.z$tune(ch$z, scale = 0.3)
  chain.write(xfile, ch$x, append = !i == 1)
  chain.write(zfile, ch$z, append = !i == 1)
}
xfile <- "X1.bin"
zfile <- "Z1.bin"
for (i in 1:15) {
  ch <- metropolis(d.model, iters=2000, thin=10,
                  start.x=ch$last.x,
                  start.z=ch$last.z)

  proposal.x$tune(ch$x, scale = 0.3)
  proposal.z$tune(ch$z, scale = 0.3)
  chain.write(xfile, ch$x, append = !i == 1)
  chain.write(zfile, ch$z, append = !i == 1)
}
## etc.
xfile <- "X2.bin"
zfile <- "Z2.bin"

```

```
for (i in 1:15) {  
  ch <- metropolis(d.model, iters=2000, thin=10,  
                  start.x=ch$last.x,  
                  start.z=ch$last.z)  
  
  proposal.x$tune(ch$x, scale = 0.3)  
  proposal.z$tune(ch$z, scale = 0.3)  
  chain.write(xfile, ch$x, append = !i == 1)  
  chain.write(zfile, ch$z, append = !i == 1)  
}
```

# BIBLIOGRAPHY

- Adler, D., Gläser, C., Nenadic, O., Oehlschlägel, J., and Zucchini, W. (2010). *ff: memory-efficient storage of large data on disk and fast access functions*. R package version 2.1-2.
- Afanasyev, V. (2004). A miniature daylight level and activity data recorder for tracking animals over long periods. Technical report, British Antarctic Survey.
- Alerstam, T., Gudmundsson, G. A., and Hedenström, M. G. A. (2001). Migration along orthodromic sun compass routes by arctic birds. *SCI*, 291:300–303.
- Anderson, J., Andres, J., Davis, M., Fujiwara, K., Fang, T., and Nedbal, M. (2010). Voyager: An Interactive Software for Visualizing Large, Geospatial Data Sets. *Marine Technology Society Journal*, 44(4):8–19.
- Anderson, K., Nielsen, A., Thygesen, U., Hinrichsen, H.-H., and Neuenfeldt, S. (2007). Using the particle filter to geolocate Atlantic cod (*Gadus morhua*) in the Baltic Sea, with special emphasis on determining uncertainty. *Canadian Journal of Fisheries and Aquatic Sciences*, 64:618–627.
- Andres, J., Davis, M., Fujiwara, K., Anderson, J., Fang, T., and Nedbal, M. (2009). A Geospatially Enabled, PC-Based, Software to Fuse and Interactively Visualize Large 4D/5D Data Sets.
- Austin, D. A., McMillan, J. I., and Bowen, W. D. (2003). A three-stage algorithm for filtering erroneous argos satellite locations. *Marine Mammal Science*, 19:371–383.
- Baddeley, A. and Turner, R. (2005). Spatstat: an R package for analyzing spatial point patterns. *Journal of Statistical Software*, 12(6):1–42. ISSN 1548-7660.
- Banerjee, S. (2004). On geodetic distance computations in spatial modelling. *Biometry*, 61:617–625.
- Beauplet, G., Dubroca, L., Guinet, C., Cherel, Y., Dablin, W., Gagne, C., and Hindell, M. (2004). Foraging ecology of subantarctic fur seals *Arctocephalus tropicalis* breeding on Amsterdam Island: seasonal changes in relation to maternal characteristics and pup growth. *Marine Ecology Progress Series*, 273:211–225.
- Beck, C. A. and McMillan, J. I. (2002). An algorithm to improve geolocation positions using sea surface temperature and diving depth. *Marine Mammal Science*, 18:940–951.
- Beegle-Krause, C. J., Vance, T., Reusser, D., Stuebe, D., and Howlett, E. (2010). *Pelagic Habitat Visualization: The Need for a Third (and Fourth) Dimension: HabitatSpace*. <http://www.research4d.org/publications/HabitatSpace.pdf>.
- Best, N., Cowles, M., and Vines, K. (1995). CODA: Convergence diagnosis and output analysis software for Gibbs sampling output. *Version 0.3. MRC Biostatistic Unit, Cambridge, UK*.

- Bivand, R., Pebesma, E., and Gómez-Rubio, V. (2008). *Applied spatial data analysis with R*. Springer.
- Block, B., Costa, D., Boehlert, G., and Kochevar, R. (2003a). Revealing pelagic habitat use: the tagging of Pacific pelagics program Idées sur l'utilisation de l'habitat pélagique: le programme de marquage de pélagiques dans le Pacifique. *Oceanologica Acta*, 25:255–266.
- Block, B., Costa, D., Boehlert, G., and Kochevar, R. (2003b). Revealing pelagic habitat use: the tagging of Pacific pelagics program Idées sur l'utilisation de l'habitat pélagique: le programme de marquage de pélagiques dans le Pacifique. *Oceanologica Acta*, 25:255–266.
- Block, B. A., Boustany, A., Teo, S., Walli, A., Farwell, C. J., Williams, T., Prince, E. D., Stokesbury, M., Dewar, H., Seitz, A., and Weng, K. (2003c). Distribution of western tagged atlantic bluefin tuna determined from archival and pop-up satellite tags. *International Commission for the Conservation of Atlantic Tunas Collective Volume of Scientific Papers*, 55(3):1127–1139.
- Boehlert, G., Costa, D., Crocker, D., Green, P., O'Brien, T., Levitus, S., and Le Boeuf, B. (2001). Autonomous pinniped environmental samplers: Using instrumental animals as oceanographic data collectors. *Journal of Atmospheric and Oceanic Technology*, 18:1882–1893.
- Boehme, L., Lovell, P., Biuw, M., Roquet, F., Nicholson, J., Thorpe, S., Meredith, M., and Fedak, M. (2009). Technical Note: Animal-borne CTD-Satellite Relay Data Loggers for real-time oceanographic data collection. *Ocean Sci. Discuss*, 6.
- Bowditch, N. (2002). *The American Practical Navigator*. Paradise Cay Pubns.
- Bradshaw, C. J. A., Hindell, M. A., Littnan, C., and Harcourt, R. G. (2006). Determining Marine Movements of Australasian Pinnipeds. In Merrick, J. R., Archer, M., Hickey, G., and Lee, M., editors, *Evolution and Biogeography of Australasian Vertebrates*, pages 889–911. Australian Scientific Publishers, Sydney.
- Bradshaw, C. J. A., Hindell, M. A., Michael, K. J., and Sumner, M. D. (2002). The optimal spatial scale for the analysis of elephant seal foraging as determined by geo-location in relation to sea surface temperatures. *ICES Journal of Marine Science*, 59:770–781.
- Breed, G. (2008). *State-space analyses indicate experience, prey availability, competition, and reproductive status drive foraging behaviour in a top marine predator*. PhD thesis, Dalhousie University, Halifax, NS, Canada.
- Brillinger, D. R. and Stewart, B. S. (1998). Elephant-seal movements: Modelling migration. *The Canadian Journal of Statistics*, 26:431–443.
- Brothers, N., Reid, T., and Gales, R. (1997). At-sea distribution of shy albatrosses *Diomedea cauta cauta* derived from records of band recoveries and colour-marked birds. *Emu*, 97(3):231–239.
- Burns, J., Costa, D., Fedak, M., Hindell, M., Bradshaw, C., Gales, N., McDonald, B., Trumble, S., and Crocker, D. (2004). Winter habitat use and foraging behavior of crabeater seals along the Western Antarctic Peninsula. *Deep Sea Research Part II: Topical Studies in Oceanography*, 51(17-19):2279–2303.
- Butler, H., Daly, M., Doyle, A., Gillies, S., Schaub, T., and Schmidt, C. (2008). *The GeoJSON Format Specification*, 1.0 edition.

- Calenge, C., Dray, S., and Royer-Carenzi, M. (2009). The concept of animals' trajectories from a data analysis perspective. *Ecological Informatics*, 4(1):34–41.
- Campagna, C., Piola, A., Rosa Marin, M., Lewis, M., and Fernández, T. (2006). Southern elephant seal trajectories, fronts and eddies in the Brazil/Malvinas Confluence. *Deep-Sea Research Part I*, 53(12):1907–1924.
- Chambers, J. (1998). *Programming with data: A guide to the S language*. Springer Verlag.
- Chambers, J. (2008). *Software for data analysis: Programming with R*. Springer Verlag.
- Cooke, S. J., Hinch, S. G., Wikelski, M., Andrews, R. D., Kuchel, L. J., Wolcott, T. G., and Butler, P. J. (2004). Biotelemetry: a mechanistic approach to ecology. *Trends in Ecology and Evolution*, 19:334–343.
- Coyne, M. and Godley, B. (2005). Satellite Tracking and Analysis Tool (STAT): an integrated system for archiving, analyzing and mapping animal tracking data. *Marine Ecology Progress Series*, 301:1–7.
- Croxall, J., Taylor, F., and Silk, J., editors (2004). *Tracking ocean wanderers: the global distribution of albatrosses and petrels. Results from the Global Procellariiform Tracking Workshop*. BirdLife International, Cambridge, UK. 1-5 September, 2003, Gordon's Bay, South Africa.
- de La Beaujardière, J., Beegle-Krause, C., Bermudez, L., Hankin, S., Hazard, L., Howlett, E., Le, S., Proctor, R., Signell, R., Snowden, D., and Thomas, J. (2009). Ocean and Coastal Data Management. *Community White Paper for OceanObs*, 9.
- Delong, R., Stewart, B., and Hill, R. (1992). Documenting migrations of northern elephant seals using day length. *Marine Mammal Science*, 8:155–159.
- Dennis, B. (1996). Discussion: Should ecologists become bayesians? *Ecological Applications*, 6(4):1095–1103.
- Dixon, P. and Ellison, A. M. (1996). Introduction: Ecological Applications of Bayesian Inference. *Ecological Applications*, 6:1034–1035.
- Domeier, M. L., Kiefer, D., Nasby-Lucas, N., Wagschal, A., and O'Brien, F. (2005). Tracking pacific bluefin tuna (*Thunnus thynnus orientalis*) in the northeastern pacific with an automated algorithm that estimates latitude by matching sea-surface-temperature data from satellites with temperature data from tags on fish. *Fishery Bulletin*, 108:292–306.
- Dorazio, R. and Johnson, F. (2003). Bayesian inference and decision theory—a framework for decision making in natural resource management. *Ecological Applications*, 13:556–563.
- Douglas, D. (2006). *The Douglas Argos-Filter Algorithm*. USGS Alaska Science Center, 3100 National Park Road Juneau, AK 99801, v7.02 edition.
- Ekstrom, P. (2002). Blue twilight in a simple atmosphere. *Proceedings of The International Society for Optical Engineering (SPIE)*, 4815:paper 14.
- Ekstrom, P. (2004). An advance in geolocation by light. *Memoirs of the National Institute of Polar Research (Special Issue)*, 58:210–226.
- Ekstrom, P. (2007). Error measures for template-fit geolocation based on light. *Deep Sea Research Part II: Topical Studies in Oceanography*, 54:392–403. Bio-logging Science: Logging and Relaying Physical and Biological Data Using Animal-Attached Tags - Proceedings of the 2005 International Symposium on Bio-logging Science, Second International Conference on Bio-logging Science.

- Ellison, A. M. (1996). An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecological Applications*, 6(4):1036–1046.
- Ellison, A. M. (2004). Bayesian inference in ecology. *Ecology Letters*, 7:509–520.
- Evans, K. and Arnold, G. (2008). Report on a geolocation methods workshop convened by the SCOR Panel on New Technologies for Observing Marine Life. Technical report, SCOR.
- Evenden, G. (1990). Cartographic Projection Procedures for the UNIX Environment. A User's Manual. *USGS Open-file report*, pages 90–284.
- Fledermaus (2010). *Fledermaus*. IVS 3D, Fredericton, New Brunswick.
- Freitas, C. (2010). *argosfilter: Argos locations filter*. R package version 0.62.
- Freitas, C., Lydersen, C., Fedak, M., and Kovacs, K. (2008). A simple new algorithm to filter marine mammal Argos locations. *Marine Mammal Science*, 24(2):315–325.
- Gaspar, P., Georges, J., Fossette, S., Lenoble, A., Ferraroli, S., and Le Maho, Y. (2006). Marine animal behaviour: neglecting ocean currents can lead us up the wrong track. *Proceedings of the Royal Society B*, 273(1602):2697.
- Gelman, A., Carlin, J., and Stern, H. (2004). *Bayesian data analysis*. CRC press.
- Gilks, W. R., Richardson, S., and Spiegelhalter, D. J. (1995). Introducing Markov Chain Monte Carlo. In *Markov Chain Monte Carlo in Practice*, pages 1–19. Chapman & Hall/CRC, Boca Raton, FL.
- Google (2010). *Google Earth User Guide*, 5.0 edition.
- Goulet, A., Hammill, M., and Barrette, C. (1999). Quality of satellite telemetry locations of gray seals (*Halichoerus grypus*). *Marine Mammal Science*, 15(2):589–594.
- Green, P. J. and Silverman, B. W. (1994). *Nonparametric regression and generalized linear models: a roughness penalty approach*. CRC Press.
- Gröger, J., Rountree, R., Thygesen, U., Jones, D., Martins, D., Xu, Q., and Rothschild, B. (2007). Geolocation of Atlantic cod (*Gadus morhua*) movements in the Gulf of Maine using tidal information. *Fisheries Oceanography*, 16(4):317–335.
- Gudmundsson, G. A. and Alerstam, T. (1998). Optimal map projections for analysing long-distance migration routes. *Journal of Avian Biology*, 29:597–605.
- Gunn, J. S., Patterson, T. A., and Pepperell, J. G. (2003). Short-term movement and behaviour of black marlin *Makaira indica* in the coral sea as determined through a pop-up satellite archival tagging experiment. *Marine and Freshwater Research*, 54(4):515–525.
- Halpin, P., Read, A., Best, B., and Hyrenbach, K. (2006). OBIS-SEAMAP: developing a biogeographic research data commons for the ecological studies of marine mammals, seabirds, and sea turtles. *Marine Ecology Progress Series*, 316:239–246.
- Hartog, J., Patterson, T., Hartmann, K., Jumppanen, P., Cooper, S., and Bradford, R. (2009). Developing integrated database systems for the management of electronic tagging data. *Tagging and Tracking of Marine Animals with Electronic Devices*, pages 367–380.
- Hays, G. C., Åkesson, S., Godley, B. J., Luschi, P., and Santidiran, P. (2001). The implications of location accuracy for the interpretation of satellite-tracking data. *Animal Behaviour*, 61:1035–1040.

- Hebblewhite, M. and Haydon, D. (2010). Distinguishing technology from biology: a critical review of the use of GPS telemetry data in ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550):2303.
- Hijmans, R. J., Williams, E., and Vennes, C. (2010). *geosphere: Spherical Trigonometry*. R package version 1.2-4.
- Hill, R. (2005). *WC-GPE: Global Position Estimator Program Suite User's Manual*. Wildlife Computers, Redmond, WA.
- Hill, R. D. (1994). Theory of geolocation by light levels. In Boeuf, B. J. L. and Laws, R. M., editors, *Elephant seals: population ecology, behaviour and physiology*, pages 227–236. University of California Press, Berkeley.
- Hill, R. D. and Braun, M. (2001). Geolocation by light level - the next step: Latitude. In Sibert, J. and Nielsen, J., editors, *Electronic Tagging and Tracking in Marine Fisheries*, pages 315–330. Kluwer, Boston.
- Hindell, M. A., Bradshaw, C. J. A., Sumner, M. D., Michael, K. J., and Burton, H. R. (2003). Dispersal of female southern elephant seals and their prey consumption during the austral summer: relevance to management and oceanographic zones. *Journal of Applied Ecology*, 40:703–715.
- Hindell, M. A., Burton, H. R., and Slip, D. J. (1991). Foraging areas of southern elephant seals, *Mirounga leonina*, as inferred from water temperature data. *Australian Journal of Marine Freshwater Research*, 42:115–128.
- Hunter, J., Argue, A. W., Bayliff, W. H., Dizon, A. E., Fonteneau, A., Goodman, D., and Seckel, G. R. (1986). The dynamics of tuna movements: an evaluation of past and future research. Technical report, FAO Fish.
- Jonsen, I., Flemming, J., and Myers, R. (2005). Robust state-space modeling of animal movement data. *Ecology*, 86(11):2874–2880.
- Jonsen, I. D., Myers, R. A., and Flemming, J. M. (2003). Meta-analysis of animal movement using state-space models. *Ecology*, 84(11):3055–3063.
- Keating, K. and Cherry, S. (2009). Modeling utilization distributions in space and time. *Ecology*, 90(7):1971–1980.
- Keitt, T. H., Bivand, R., Pebesma, E., and Rowlingson, B. (2010). *rgdal: Bindings for the Geospatial Data Abstraction Library*. R package version 0.6-27.
- Kenward, R. E. (1987). *Wildlife Radio Tagging*. Academic Press, San Diego, CA.
- Lake, S., Wotherspoon, S., and Burton, H. (2005). Spatial utilization of fast-ice by Weddell seals *Leptonychotes weddelli* during winter. *Ecography*, 28(3):295–306.
- Lam, C., Nielsen, A., and Sibert, J. (2008). Improving light and temperature based geolocation by unscented Kalman filtering. *Fisheries Research*, 91(1):15–25.
- Lewin-Koh, N. J. and Bivand, R. (2010). *mapproj: Tools for reading and handling spatial objects*. R package version 0.7-34.
- Luque, S. P. (2007). Diving behaviour analysis in R. *R News*, 7(3):8–14.
- Maindonald, J. and Braun, J. (2007). *Data analysis and graphics using R: an example-based approach*. Cambridge Univ Pr.

- Matthiopoulos, J. (2003). Model-supervised kernel smoothing for the estimation of spatial usage. *Oikos*, 102:367–377.
- Matthiopoulos, J., Harwood, J., and Thomas, L. (2005). Metapopulation consequences of site fidelity for colonially breeding mammals and birds. *Journal of Animal Ecology*, 74:716–727.
- Matthiopoulos, J., McConnell, B., Duck, C., and Fedak, M. (2004). Using satellite telemetry and aerial counts to estimate space use by grey seals around the British Isles. *Journal of Applied Ecology*, 41:476–491.
- McConnell, B. J., Chambers, C., and Fedak, M. A. (1992). Foraging ecology of southern elephant seals in relation to bathymetry and productivity of the Southern Ocean. *Antarctic Science*, 4:393–398.
- Meeus, J. (1991). *Astronomical Algorithms*. Willmann-Bell, VA.
- Meijers, A., Bindoff, N., and Rintoul, S. (2010). Estimating the 4-dimensional structure of the Southern Ocean using satellite altimetry. *Journal of Atmospheric and Oceanic Technology*, in press.
- Metcalf, J. D. (2001). Summary report of a workshop on daylight measurements for geolocation in animal telemetry. In Sibert, J. R. and Nielsen, J. L., editors, *Electronic Tagging and Tracking in Marine Fisheries*, pages 331–342.
- Musyl, M. K., Brill, R. W., Curran, D. S., Gunn, J. S., Hartog, J. R., Hill, R. D., Welch, D. W., Eveson, J. P., Boggs, C. H., and Brainard, R. E. (2001). Ability of archival tags to provide estimates of geographical position based on light intensity. In *Electronic Tagging and Tracking in Marine Fisheries*, pages 343–367, Dordrecht. Kluwer.
- Nel, D. C., Ryan, P. G., Nel, J. L., Klages, N. T. W., Wilson, R. P., Robertson, G., and Tuck, G. N. (2002). Foraging interactions between Wandering Albatrosses *Diomedea exulans* breeding on Marion Island and long-line fisheries in the southern Indian Ocean. *Ibis*, 144:141–154.
- Nielsen, A., Bigelow, K. A., Musyl, M. K., and Sibert, J. R. (2006). Improving light-based geolocation by including sea surface temperature. *Fisheries Oceanography*, 15:314–325.
- Orsi, A., Whitworth, T., and Nowlin, W. (1995). On the meridional extent and fronts of the Antarctic Circumpolar Current. *Deep Sea Research Part I: Oceanographic Research Papers*, 42(5):641–673.
- Ovaskainen, O., Rekola, H., Meyke, E., and Arjas, E. (2008). Bayesian methods for analyzing movements in heterogeneous landscapes from mark-recapture data. *Ecology*, 89(2):542–554.
- Patterson, T., McConnell, B., Fedak, M., Bravington, M., and Hindell, M. (2010). Using GPS data to evaluate the accuracy of state-space methods for correction of Argos satellite telemetry error. *Ecology*, 91(1):273–285.
- Patterson, T., Thomas, L., Wilcox, C., Ovaskainen, O., and Matthiopoulos, J. (2008). State-space models of individual animal movement. *Trends in Ecology & Evolution*, 23(2):87–94.
- Pauly, T., Higginbottom, I., Pederson, H., Malzone, C., Corbett, J., and Wilson, M. (2009). Keeping Pace with Technology Through the Development of an Intuitive Data Fusion, Management, Analysis & Visualization Software Solution. In *OCEANS 2009-EUROPE, 2009.*, pages 1–8.

- Pebesma, E. and Bivand, R. (2005). Classes and methods for spatial data in R. *R News*, 5(2):9–13.
- Pedersen, M., Righton, D., Thygesen, U., Andersen, K., and Madsen, H. (2008). Geolocation of North Sea cod (*Gadus morhua*) using hidden Markov models and behavioural switching. *Canadian Journal of Fisheries and Aquatic Sciences*, 65(11):2367–2377.
- Phillips, R. A., Silk, J. R. D., Croxall, J. P., Afanasyev, V., and Briggs, D. R. (2004). Accuracy of geolocation estimates for flying seabirds. *Marine Ecology Progress Series*, 266:265–272.
- Pinaud, D. (2007). Quantifying search effort of moving animals at several spatial scales using first-passage time analysis: effect of the structure of environment and tracking systems. *Journal of Applied Ecology*, 45:91–99.
- Plummer, M., Best, N., Cowles, K., and Vines, K. (2006). CODA: Convergence diagnosis and output analysis for MCMC. *R News*, 6(1):7–11.
- R Development Core Team (2010). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
- Roberts, S., Guilford, T., Rezek, I., and Biro, D. (2004). Positional entropy during pigeon homing I: application of Bayesian latent state modelling. *Journal of Theoretical Biology*, 227:39–50.
- Royer, F., Fromentin, J.-M., and Gaspar, P. (2005). A state-space model to derive bluefin tuna movement and habitat from archival tags. *Oikos*, 109:473–484.
- Samuel, M. D., Pierce, D. J., and Garton, E. O. (1985). Identifying areas of concentrated use within the home range. *JAnE*, 54:711–719.
- Schick, R., Loarie, S., Colchero, F., Best, B., Boustany, A., Conde, D., Halpin, P., Joppa, L., McClellan, C., and Clark, J. (2008). Understanding movement data and movement processes: current and emerging directions. *Ecology letters*, 11(12):1338–1350.
- Seaman, D. E., Millspaugh, J. J., Kernohan, B. J., Brundige, G. C., Raedeke, K. J., and Gitzen, R. A. (1999). Effects of sample size on kernel home range estimates. *Journal of Wildlife Management*, 63:739–747.
- Service Argos (2004). *User's Manual*. Collecte Localisation Satellites (CLS), France.
- Shaffer, S. A., Tremblay, Y., Awkerman, J. A., Henry, R. W., Teo, S. L. H., Anderson, D. J., Croll, D. A., Block, B. A., and Costa, D. P. (2005). Comparison of light- and SST-based geolocation with satellite telemetry in free-ranging albatrosses. *Marine Biology*, 147:833–843.
- Sibert, J. and Nielsen, A. (2007). State-space model for light-based tracking of marine animals. *Canadian Journal of Fisheries and Aquatic Sciences*, 64(8):1055–1068.
- Silverman, B. (1998). *Density estimation for statistics and data analysis*. Chapman & Hall/CRC.
- Simmons, S., Crocker, D., Kudela, R., and Costa, D. (2007). Linking foraging behaviour of the northern elephant seal with oceanography and bathymetry at mesoscales. *Marine Ecology Progress Series*, 346:265–275.
- Simonoff, J. (1996). *Smoothing methods in statistics*. Springer Verlag.

- Smith, P. and Goodman, D. (1986). Determining fish movements from an “archival tag”: Precision of geographical positions made from a time series of swimming temperature and depth. Technical report, National Marine Fisheries Service, National Oceanic and Atmospheric Administration (NOAA), Springfield, VA.
- Sobel, D. (1998). *Longitude. The true story of a lone genius who solved the greatest scientific problem of his time*. Fourth Estate.
- Sokolov, S. and Rintoul, S. (2002). Structure of Southern Ocean fronts at 140 E. *Journal of Marine Systems*, 37(1-3):151–184.
- Sokolov, S., Rintoul, S. R., and Wienecke, B. (2006). Tracking the Polar Front south of New Zealand using penguin dive data. *Deep Sea Research Part I: Oceanographic Research Papers*, 53(4):591–607.
- Sumner, M. and Wotherspoon, S. (2010). *TripEstimation: Metropolis sampler and supporting functions for estimating animal movement from archival tags and satellite fixes*. R package version 0.0-33.
- Sumner, M. D., Michael, K. J., Bradshaw, C. J. A., and Hindell, M. A. (2003). Remote sensing of Southern Ocean sea surface temperature: implications for marine biophysical models. *Remote Sensing of Environment*, 84:161–173.
- Sumner, M. D., Wotherspoon, S. J., and Hindell, M. A. (2009). Bayesian estimation of animal movement from archival and satellite tags. *PLoS ONE*, 4(10):e7324.
- Teo, S. L. H., Boustany, A., Blackwell, S., Walli, A., Weng, K. C., and Block, B. A. (2004). Validation of geolocation estimates based on light level and sea surface temperature from electronic tags. *Marine Ecology Progress Series*, 283:81–98.
- Terauds, A., Gales, R., Baker, G., and Alderman, R. (2006). Foraging areas of black-browed and grey-headed albatrosses breeding on Macquarie Island in relation to marine protected areas. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 16(2):133–146.
- Thompson, D., Moss, S. E. W., and Lovell, P. (2003). Foraging behaviour of South American fur seals *Arctocephalus australis*: extracting fine scale foraging behaviour from satellite tracks. *Marine Ecology Progress Series*, 260:285–296.
- Thums, M., Bradshaw, C., and Hindell, M. (2008). Tracking changes in relative body composition of southern elephant seals using swim speed data. *Marine Ecology Progress Series*, 370:249–261.
- Thys, T. M., Hobson, B. W., and Dewar, H. (2001). Marine animals: the next generation of autonomous underwater vehicle? *Oceans*.
- Tremblay, Y., Shaffer, S. A., Fowler, S. L., Kuhn, C. E., McDonald, B. I., Weise, M. J., Bost, C.-A., Weimerskirch, H., Crocker, D. E., Goebel, M. E., and Costa, D. P. (2006). Interpolation of animal tracking data in a fluid environment. *The Journal of Experimental Biology*, 209:128–140.
- Turchin, P. (1998). *Quantitative Analysis of Movement: measuring and modeling population redistribution in plants and animals*. Sinauer Associates, Inc.
- Unwin, A. (2006). *Interacting with Graphics: Visualizing a Million*. Springer.
- Veness, A. (2009). A real-time spatio-temporal data exploration tool for marine research. Master’s thesis, Geography and Environmental Studies, UTAS.

- Vincent, C., McConnell, B. J., Ridoux, V., and Fedak, M. A. (2002). Assessment of Argos Location Accuracy from Satellite Tags Deployed on Captive Gray Seals. *Marine Mammal Science*, 18:156–166.
- Viswanathan, G., Afanasyev, V., Buldyrev, S., Havlin, S., da Luz, M., Raposo, E., and Stanley, H. (2000). Lévy flights in random searches. *Physica A: Statistical Mechanics and its Applications*, 282(1-2):1–12.
- Warmerdam, F. and the GDAL development team (2010). *GDAL—Geospatial Data Abstraction Library*. Open Source Geospatial Foundation.
- Welch, D. W. and Eveson, J. P. (1999). An assessment of light-based geolocation estimates from archival tags. *Canadian Journal of Fisheries and Aquatic Sciences*, 56:1317–1327.
- Wentz, E. A., Campbell, A. F., and Houston, R. (2003). A comparison of two methods to create tracks of moving objects: linear weighted distance and constrained random walk. *International Journal of Geographic Information Science*, 17:623–645.
- Wessel, P. and Smith, W. (1996). A global, self-consistent, hierarchical, high-resolution shoreline database. *Journal of Geophysical Research*, 101(B4):8741–8743.
- White, N. A. and Sjöberg, M. (2002). Accuracy of satellite positions from free-ranging grey seals using ARGOS. *Polar Biology*, 25:629–631.
- Wilson, R. P., Ducamp, J. J., Rees, W. G., Culik, B. M., and Niekamp, K. (1992). Estimation of location: global coverage using light intensity. In Priede, I. G. and Swift, S. M., editors, *Wildlife Telemetry: Remote Monitoring and Tracking of Animals*, pages 131–135. Ellis Horwood, London.
- Wilson, R. P., Grémillet, D., Syder, J., Kierspel, M. A. M., Garthe, S., Weimerskirch, H., Schäfer-Neth, C., Scolaro, J. A., Bost, C., Plötz, J., and Nel, D. (2002). Remote-sensing systems and seabirds: their use, abuse and potential for measuring marine environmental variables. *Marine Ecology Progress Series*, 228:241–261.
- Wintle, B. A., McCarthy, M. A., Volinsky, C. T., and Kavanagh, R. P. (2003). The use of Bayesian model averaging to better represent uncertainty in ecological models. *Conservation Biology*, 17:1579–1590.
- Wood, A. G., Naef-Daenzer, B., Prince, P. A., and Croxall, J. P. (2000). Quantifying habitat use in satellite-tracked pelagic seabirds: application of kernel estimation to albatross locations. *Journal of Avian Biology*, 31:278–286.
- Xavier, J. C., Trathan, P. N., Croxall, J. P., Wood, A. G., Podesta, G., and Rodhouse, P. G. (2004). Foraging ecology and interactions with fisheries of wandering albatrosses (*Diomedea exulans*) breeding at South Georgia. *Fisheries Oceanography*, 13:324–344.