STATISTICAL DEVELOPMENTS FOR UNDERSTANDING ANTHROPOGENIC IMPACTS ON MARINE ECOSYSTEMS

Laura Marshall

A Thesis Submitted for the Degree of PhD
at the
University of St. Andrews

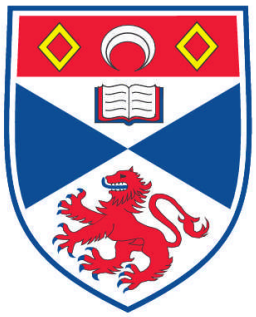

2012

STATISTICAL DEVELOPMENTS FOR UNDERSTANDING ANTHROPOGENIC IMPACTS ON MARINE ECOSYSTEMS

Laura Marshall

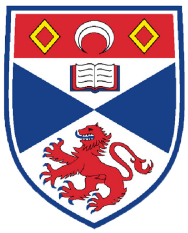

Thesis submitted for the degree of
DOCTOR OF PHILOSOPHY
in the School of Biology
UNIVERSITY OF ST ANDREWS
ST ANDREWS
JANUARY 2011

Abstract

Over the past decades technological developments have both changed and increased human influence on the marine environment. We now have greater potential than ever before to introduce disturbance and deplete marine resources. Two of the issues currently under public scrutiny are the exploitation of fish stocks worldwide and levels of anthropogenic noise in the marine environment. The aim of this thesis is to investigate and develop novel analyses and simulations to provide additional insight into some of the challenges facing the marine ecosystem today. These methodologies will improve the management of these risks to marine ecosystems.

This thesis first addresses the issue of competition between humans and grey seals (*Halichoerus grypus*) for marine resources, providing compelling evidence that a substantial proportion of the sandeels consumed by grey seals in the North Sea are in fact *H. lanceolatus*, which is not commercially exploited, rather than the commercially important *A. marinus*. In addition, we present quantitative results regarding sources of bias when estimating the total biomass of sandeels consumed by grey seals. Secondly, we investigate spatially adaptive 2-dimensional smoothing to improve the prediction of both the presence and density of marine species, information that is often key in the management of marine ecosystems. Particularly, we demonstrate the benefits of such methods in the prediction of sandeel occurrence. Lastly this thesis provides a quantitative assessment of the protocols for real-time monitoring of marine mammal presence, which require that acoustic operations cease when an animal is detected within a certain distance (i.e. the "monitoring zone") of the sound source. We assess monitoring zones of different sizes with regards to their effectiveness in reducing the risks of temporary and permanent damage to the animals' hearing, and demonstrate that a monitoring zone of 2 km is generally recommendable.

Declarations

I, Laura Helen Marshall, hereby certify that this thesis, which is approximately 36000 words in length, has been written by me, that it is the record of work carried out by me and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in January 2006 and as a candidate for the degree of Doctor of Philosophy in Statistics in January 2007; the higher study for which this is a record was carried out in the University of St Andrews between 2006 and 2010.

date:_____ signature of candidate:_____

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in Statistics in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

date:_____ signature of supervisor:_____

date:_____ signature of supervisor:_____

date:_____ signature of supervisor:_____

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. We have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the electronic publication of this thesis:

Embargo on Chapter 4 of printed copy and electronic copy for the same fixed period of five years on the following grounds: publication would be commercially damaging to the researcher, or to the supervisor, or the University; publication would be in breach of law or ethics

date:_____ signature of candidate:_____

date:_____ signature of supervisor:_____

date:_____ signature of supervisor:_____

date:_____ signature of supervisor:_____

Acknowledgments

I feel very lucky to have had the opportunity to work in the wonderfully supportive and friendly environment at CREEM. I have worked with many different people and have learned a lot from each them.

- I thank my supervisors John Harwood, Monique MacKenzie and Carl Donovan for their wonderful support and advice throughout the work on this thesis. I am particularly grateful for their efforts in the final month which coincided with the festive period.

- I would like to thank Steve Buckland. When I completed my undergraduate degree in statistics and computer science I had not planned on continuing in academia. It was on Steve's suggestion that I applied for the masters conversion course to environmental biology. I feel this was a very impotant decision in my life that has led me to wonderful places in both my professional and family life.

- The scope of my thesis has meant that I have been very lucky in that I have had the opportunity to work with many different people. I am particularly grateful to Sophie Smout, Charles Paxton, Len Thomas, David Borchers, Phil Hammond and Kate Grellier for their discussions and suggestions for the work in this thesis.

- I am also grateful to Jason Matthiopoulos for his part in the review process of this thesis and his encouragement to pursue my ideas.

- I thank Rhona Rodgers for her help sorting out various bureaucracy issues and Phil Le Feuvre for his help with the computer issues!

- I have been very lucky to have shared an office for a number of years with two wonderful people Lindesay Scott-Hayward and Joanne Potts. I thank you for

all your hours of discussions as well as your treasured friendship. We have had many adventures together and I hope we will have many more.

- To all the other people who have made CREEM such a wonderful place to work. A special thank you to Leslie New for being a wonderful friend and for all your help and support - see you in Chile soon! Another special thank you to Monique MacKenzie, who as well as being a wonderful supervisor, has also been a good friend and taught me a lot far beyond the scope of this thesis, both professionally and personally. Thank you to Bruno Caneco, Lilliana Martins, Debbie Russell, Danielle Harris, Darren Kidney, Joyce Yuan, Angelika Studney, Glenna Evans and Cornelia Oedekoven for their friendship and discussions. Further thanks to Catriona Harris, Janine Illian, Ruth King, Lorenzo Milazzo, Steve Buckland, Louise Burt and Eric Rexstad for sharing their knowledge and experience.

- Thank you to Maggie Winton, from the Student Support Services, for helping me focus when things seemed out of control

- To all my other friends and family who have helped coax me on when the finish line was no where in sight.

- Thank you to my family in Chile who have welcomed me into their lives. A special thank you to Nairy who in the final months took good care of me at lunch times before chasing me back upstairs to work :-).

- A huge thank you to my mum and dad for all their love, support, advice, patience and wonderful cooking :-) Thank you also to my brother and sister for their support. Further thanks to my family for all the DIY help which gave me a wonderful place to live whilst working in St. Andrews.

- A special thank you to my husband Rodrigo Wiff, especially for understanding my vastly different approach to work than his own. I have learned a great deal from him and I thank him for his love, patience, advice and wonderful cooking :-) A final thank you to my wonderful little boy Tiago, for giving me that bit of extra drive at the end when I really needed it. You are the best reward in the world xxx.

Despite the direct contributions of everyone else to this thesis, the remaining mistakes are my sole responsibility.

Data acknowledgments

With thanks to the following people and institutions for providing data. In addition we acknowledge the work of all participant in the surveys.

- Otolith and fish mass data from G. Engelhard, CEFAS, Lowestoft.
- Grey seal feeding experimental data from K. Grellier, SMRU, University of St. Andrews.
- North Sea scat data from P.Hammond and K. Grellier, SMRU, University of St. Andrews.
- Sandeel acoustic densities in the Firth of Forth and Wee Bankie region from S. Greenstreet, FRS, Marine Laboratory, Aberdeen.
- Sightings data from the Small Cetaceans in the European Atlantic and North Sea II (SCANS-II), SMRU, University of St. Andrews.
- Northern bottlenose whale sightings data, 2001 North Atlantic Sightings Surveys (NASS-2001).

Institutional funding acknowledgments

- With thanks to NERC and BAE Systems Insyte Ltd for funding this PhD thesis.

Table of Contents

List of Figures

List of Tables

Chapter 1

General Introduction

1.1 Anthropogenic Impacts on the Marine Environment

Over the past decades the challenges facing the marine environment due to human influence have changed considerably, and predators towards the top of the food chain, such as marine mammals, are one of the species groups that are potentially the most at risk. The Dolphins, Whales and Porpoises: 2002-2010 Conservation Action Plan for the World's Cetaceans (Reeves et al., 2003) states that "cetacean diversity, like all biodiversity worldwide, is crumbling; we are losing it at a rapid and increasing rate". Species such as the baiji (*Lipotes vexillifer*), vaquita (*Phocoena sinus*), and North Atlantic right whale (*Eubalaena glacialis*) are near extinction and local populations of other species have disappeared or are seriously threatened (Reeves et al., 2003). Although most marine mammal species are now protected from direct harvesting (for example, whaling) increasing technology has led to increased risks for the marine ecosystem as a whole. Humans now have greater potential than ever before to deplete marine resources and introduce new threats to the marine environment. Two of these issues currently under public scrutiny are the exploitation of worldwide fish stocks and increasing levels of anthropogenic noise.

Fisheries and aquaculture provide direct employment to around 35 million people worldwide and in 1999, fish contributed around 15.8% of animal protein consumed by humans (FAO, 2002). However, development of more effective fishing techniques without adequate management has led to the depletion of many fish stocks. According to a 1997 FAO review, 6% of worldwide fisheries are classed as depleted and a further 16% are over-exploited (Botsford et al., 1997).

The decline in many wild fish stocks has increased the demand for farmed fish and caused the aquaculture industry to grow more rapidly than any other sources of animal production (FAO, 2002). In 2002, farmed fish comprised around one third of all fish directly consumed by humans (FAO, 2002). However, the methods currently used to farm carnivorous fish, such as salmon, require large quantities of fishmeal and/or fish oil. Even with increasing efficiency, around 2.6 kg of raw fish are required to produce 1 kg of farmed salmon (Mente et al., 2006).

In Europe, the fish used in the production of the fishmeal and fish oil required for aquaculture are smaller species for which there is no demand for direct human consumption. Such fish include sandeels (*Ammodytidae sp.*, mostly *Ammodytes marinus*), capelin (*Mallotus villosus*), Norway pout (*Trisopterus esmarkii*), blue whiting (*Micromesistius poutassou*), sprat (*Sprattus sprattus*), horse mackerel (*Trachurus trachurus*) and herring (*Clupea harengus*) (Mente et al., 2006). Over-fishing these species has the potential to have devastating effects on marine diversity because many marine species, including larger and commercially important fish species, various species of sea bird and marine mammals, rely on these fish. Myers and Worm (2005) concluded that "industrial fisheries have changed marine ecosystems in fundamental ways", and Pauly et al. (2003, 2005) suggested that, if current fishing trends continue, this industry is likely to collapse in the next decades, and take the supporting ecosystems

with it.

A second threat to the marine environment that has also been the source of increasing public attention is increasing levels of anthropogenic noise (Southall et al., 2007). This is also, in part, linked with the increasing aquaculture industry, because acoustic deterrent devices are deployed to deter marine mammals from predation on farmed fish (Booth, 2010). In addition, there are many other forms of acoustic pollution in operation. Increasing levels of sonar are being used in our oceans and seas, both to locate ever-decreasing fish stocks, and also for various military and scientific applications (Hildebrand, 2004). A number of marine mammal stranding events are suspected to be associated with marine noise (Frantzis, 1998; D'Amico et al., 2009), and this has led to increased concern over the potential consequences of anthropogenic noise for marine mammals. However, as yet, the full extent of such consequences, as well as the effectiveness of current management measures, is far from understood and more research is urgently required (Southall et al., 2007; Dolman et al., 2009).

1.2 Management of the Marine Environment

The term "environmental management" refers to the process by which we aim to supervise or control human activities that may affect the natural environment. With regard to marine ecosystems the objectives of our management processes may have different priorities for different people. One objective may be that of conservation, for example, maintaining the diversity of species. Another may be an economic objective concerned with the sustainable yield of fish and other marine resources. One thing is certain: the management of our marine environment is becoming increasingly challenging as we struggle to better understand the complexities of that environment as well as balance human requirements and desires with the need for economic sustainability and the conservation of the environment. To address such challenges it

is important to develop better biological understanding of the interactions between anthropogenic activities and the species which exist in this environment. One of the key techniques for increasing that understanding is through the implementation of statistical analyses and simulations. However, there are often a number of technical challenges that may be encountered when attempting such analyses. For example, the data may not lend themselves to traditional off-the-shelf analyses or there simply may be no existing technique to deal with the question of interest. The main aim of this thesis is to develop a number of statistical techniques whose application has important consequences for the management of the marine environment.

The concerns over maintaining the diversity of the marine environment has led to large quantities of literature on threatened and endangered species, particularly marine mammals, sea birds and marine turtles (Larkin, 1996). In addition, the National implementations of EU Council Directive 92/43/EEC on the conservation of natural habitats and of wild fauna and flora (generally known as the Habitats Directive) now provides a level of protection for all cetaceans in its waters. However, such species were not considered in the single species stock assessments carried out to try to manage the anthropogenic effects of the fisheries. Therefore, to better assess the potential impacts of human activity on the ecosystem as a whole, increasing attention has been placed on the consideration of multispecies models rather than these simpler single species stock assessments (Larkin, 1996). Such models not only consider the species being subjected to exploitation but also the predators and other species which may be affected by that exploitation. This is an important step towards ensuring that the exploitation of fish stocks does not endanger other species, as well as the economic sustainability of the industry. However, if management measures based on such models are to be successful, it is essential that we have an in depth knowledge of the interactions between species within the ecosystem and try to minimise the number of

caveats in our knowledge.

One example of where such knowledge is lacking is the level of competition between the sandeel fisheries, particularly those directed at the lesser sandeel, *Ammodytes marinus*, and the grey seal (*Halichoerus grypus*, Fabricius). Currently, there is debate about the inclusion of predation by seals on *A. marinus* in North Sea multispecies fisheries models because there is an uncertainty as to how many of the sandeels consumed by the grey seals are in fact this species. There are a further four species of sandeel thought to inhabit the North Sea, although only the largest - *Hyperoplus lanceolatus* - is thought to make up a substantial proportion of the total sandeel biomass. Chapter 2 of this thesis presents a novel method for improving both the estimation of the total biomass of sandeels consumed by grey seals and the proportions of *A. marinus* in the grey seal diet. By doing this, it is now possible to ensure that seal predation on *A. marinus*, one of the key prey species in the North Sea, is included in multispecies fisheries models.

Another commonly implemented conservation technique involves the designation of protected areas, for example permanent marine reserves where all activities are restricted (Larkin, 1996; Myers and Worm, 2005; Booth, 2010) or spatial or temporal closure of areas to commercial fishing - so called "no-take" zones (Myers and Worm, 2005; Greenstreet et al., 2006). Such methods allow populations within areas of importance (e.g. breeding ground or biodiversity hotspots) to recover from the negative impacts of human activities. Pauly et al. (2005) suggests that no-take zones must be viewed as a "legitimate and obvious management tool" to protect species that are now in danger of becoming extinct. However, the reliable implementation of such management techniques requires an accurate understanding of the distribution of the species that is being protected and the factors that determine this.

Spatial modelling is an area of statistics that has been developing rapidly and is helping us to understand and predict the distributions of species. In this thesis we present a novel technique that aims to improve the modelling of species distributions, particularly when the underlying spatial distribution of the species is more variable in some parts of the species' distribution than others. Firstly, we use this technique to investigate the distribution of sandeels off the east coast of Scotland between 1997 and 2002. The results are of particular interest as they can be used to assess the effects of the closure of the sandeel fishery in this region in 2000. Secondly, we use this technique to investigate the distribution of harbour porpoise off the west coast of the United Kingdom. Harbour porpoise (*Phocoena phocoena*) are one of the species of interest in the Environmental Risk Management Capability (ERMC) (Mollett et al., 2009), which is used by the Royal Navy to assess the risks that sonar poses to marine mammals. One of the components of this system is a set of estimates of the density of all marine mammal species in different parts of the world. These estimates are used to aid the selection of test sites so that, whenever possible, sonar operations may be conducted in areas with relatively low densities of marine mammals. It is therefore important that such systems have reliable estimates of animal density through space.

Ideally, sonar operations should avoid areas where there is a high density of marine mammals but, inevitably, there are limitations to our ability to predict animal density. Indeed, marine mammal distribution may vary seasonally, or shift from year to year. In addition, national priorities may make it necessary to operate in areas where marine mammals are abundant. Chapter 4 of this thesis therefore modifies and uses part of the ERMC system to to investigate the effectiveness of monitoring zones in reducing the risks of hearing damage to marine mammals in the vicinity of sonar operations. Monitoring zones are designated areas around the sound source in which observers search for marine mammals, either visually or acoustically, and implement specified

protocols when animals are detected. Although monitoring zones are commonly used to mitigate the effects of anthropogenically generated sound, this is the first detailed evaluation of their effectiveness.

1.3 Computer Intensive Methodology in Statistics

In addition to the development of analytical techniques that are particularly applicable to management of the marine environment, a further theme which permeates this thesis is the use of computer intensive methodologies. The invention of the computer has revolutionised the statistical analysis of biological data. Analyses that would previously have been impossible because they were very time consuming now take a matter of milliseconds on a personal computer.

Since computers came into existence there has been, and continues to be, a rapid development in processing power. Personal computers have increasingly greater processing speed, more processing cores and have much greater capacity for data storage. In addition, it is becoming common place to run complex or time consuming analyses and simulations in parallel across a number of cores or different processors. With these developments comes the ability to explore scenarios and reveal insights into biological problems which were, until now, too complex to consider.

Roff (2006) defines computer intensive methods as those which implement some kind of iterative procedure, and are therefore only possible with the aid of a computer. One of the most widely implemented computer intensive statistical techniques is that of maximum likelihood estimation. In addition, other computer intensive methods, which are often implemented when analysing biological data, include bootstraps (both parametric and non-parametric), cross-validation, smoothing techniques within generalised additive models and Bayesian methods (Roff, 2006).

Such computer intensive statistics have also been widely applied in the context of marine management (for example, Harwood and Stokes (2003); Essington et al. (2006)) and also in environmental risk assessment methodology (for example, Mollett et al. (2009)). The development of such techniques has allowed the management of marine ecosystems to be based on more thorough and scientific methodology (Harwood and Stokes, 2003). One way in which computationally intensive statistics help achieve this is through the incorporation of uncertainty.

This thesis utilises computationally intensive methodology to provide additional insight into some of the challenges facing the marine ecosystem today. It is hoped that these methods will contribute to more successful and sustainable management of marine ecosystems.

In Chapter 2 we investigate the predation of sandeels by grey seals using maximum likelihood techniques to estimate the parameters associated with a novel model representing the process by which otoliths ingested by the seals are digested. We then use a parametric bootstrap to estimate the uncertainty associated with these parameter estimates. These methods not only allow us to quantify sources of potential bias in estimates of the total biomass of sandeels consumed by the seals in comparison with existing methods, but also to quantify what proportion of the sandeels in their diet belong to the commercially important species *A. marinus*.

In Chapter 3 we exploit computational power to apply a branch and bound optimisation routine for selecting specific areas of flexibility across a 2- dimensional surface. Further, we implement cross validation techniques to assess the predictive power of our resulting models and compare them to existing methodologies in the context of predicting the distribution of both sandeels and harbour porpoise.

Finally, in Chapter 4 we use computationally intensive simulations, implemented using parallel processing techniques, to assess the effectiveness of monitoring zones. By simulating the movements of hundreds of thousands of animals through time, we produce a robust analysis of the effectiveness of different sized monitoring zones at reducing the risks of sonar to marine mammal hearing.

Bibliography

C. Booth. *Variation in habitat preference and distribution of harbour porpoises west of Scotland*. PhD thesis, School of Biology, University of St. Andrews, 2010.

L. Botsford, J. Castilla, and C. Peterson. The managmenet of fisheries and marine ecosystems. *Science*, 277:509–514, 1997.

A. D'Amico, R. Gisiner, D. Ketten, J. Hammock, C. Johnson, P. Tyack, and J. Mead. Beaked whale strandings and naval exercises. *Aquatic Mammals*, 35:452–472, 2009.

S. Dolman, C. Weir, and M. Jasney. Comparative review of marine mammal guidance implemented during naval exercises. *Marine Pollution Bulletin*, 58:465–477, 2009.

T. Essington, A. Beaudreau, and J. Wiedenmann. Fishing through marine food webs. *Proceedings of the National Academy of Sciences*, 103:3171–3175, 2006.

FAO. The state of world fisheries and aquaculture, part 1: world review of fisheries and aquaculture. Technical report, Rome: Food and Agriculture Organization, Fisheries Department, 2002.

A. Frantzis. Does acoustic testing strand whales? *Nature*, 392:29, 1998.

S. Greenstreet, E. Armstrong, H. Mosegaard, H. Jensen, I. Gibb, H. Fraser, B. Scott, G. Holland, and J. Sharples. Variation in the abundance of sandeels *Ammodytes marinus* off southeast Scotland: an evaluation of area-closure fisheries management

and stock abundance assessment methods. *Journal of Marine Science*, 63:1530–1550, 2006.

J. Harwood and K. Stokes. Coping with uncertainty in ecological advice: lessons from fisheries. *TRENDS in Ecology and Evolution*, 18:617–622, 2003.

J. Hildebrand. Sources of anthropogenic sound in the marine environment. Technical report, Report to the Policy on Sound and Marine Mammals: An International Workshop. U.S. Marine Mammal Commission and Joint Nature Conservation Committee, UK. London, England., 2004.

P. Larkin. Concepts and issues in marine ecosystem management. *Reviews in Fish Biology and Fisheries*, 6:139–164, 1996.

E. Mente, G. Pierce, M. Santos, and C. Neofitou. Effect of feed and feeding in the culture of salmonids on the environment: a synthesis for European aquaculture. *Aquaculture International*, 14:499–522, 2006.

A. Mollett, C. Schofield, I. Miller, J. Harwood, C. Harris, and C. Donovan. Environmental risk management capability: Advice on minimising the impact of both sonar and seismic offshore operations on marine mammals. Technical report, SPE Offshore Europe Oil & Gas Conference & Exhibition held in Aberdeen, UK, 8th to 11th September, 2009.

R. Myers and B. Worm. Extinction, survival or recovery of large predatory fishes. *Philosophical Transactions of the Royal Society*, 360:13–20, 2005.

D. Pauly, J. Alder, E. Bennet, V. Christensen, P. Tyedmers, and R. Watson. The future for fisheries. *Science*, 302:1359–1361, 2003.

D. Pauly, R. Watson, and J. Alder. Global trends in world fisheries: impacts on marine

ecosystems and food security. *Philosophical Transactions of the Royal Society*, 360: 5–12, 2005.

R. Reeves, B. Smith, E. Crespo, and G. Sciara. *Dolphins, Whales and Porpoises: 2002-2010 conservation action plan for the world's cetaceans.* Union Internationale pour la Conservation de la Nature et de ses Ressources,Switzerland, 2003.

D. Roff. *Introduction to Computer-Intensive Methods of Data Analysis in Biology.* Cambridge University Press, 2006.

A. Southall, B. Bowles, W. Ellison, J. Finneran, R. Gentry, C. Green Jr., D. Kastak, D. Ketten, J. Miller, P. Nachtigall, W. Richardson, J. Thomas, and P. Tyack. Marine mammal exposure criteria: Initial scientific recommendations. *Aquatic Mammals*, 33(4):411–509, 2007.

Chapter 2

Grey Seal (*Halichoerus grypus*) consumption of the Lesser Sandeel (*Ammodytes marinus*) in the North Sea

2.1 Introduction

The lesser sandeel (*Ammodytes marinus*, Raitt) is the focus of increasing attention. Not only did it, until recently, constitute the greatest landed weight of any single species fishery in the North Sea (Wright et al., 2000; Holland et al., 2005) but it is also a crucial component of the North Sea's food web (Greenstreet et al., 2006; ICES, 2008). Sandeels (*Ammodytidae*), of which there are thought to be five species in the North Sea (Macer, 1966), are an important prey for various marine mammals (Santos et al., 2004; Hammond and Grellier, 2006), seabirds (Furness, 2002; Frederiksen et al., 2006) and fish, including mackerel, cod, haddock, whiting and saithe (Bromley et al., 1997; Greenstreet et al., 1998; Engelhard et al., 2008) all of which are valuable to North Sea fisheries (Harwood and Croxall, 1988). Sandeels, therefore, are not only important in sustaining the North Sea ecosystem from a conservation perspective, but their success also has significant economic implications. Despite their importance,

relatively little is known about the ecological consequences of the fisheries that target such large quantities of sandeels (ICES, 2008).

Sandeels are caught for industrial purposes rather than direct human consumption. Historical catches in the North Sea peaked at over 1.1 million t in 1997 (ICES, 2009); between 1973 and 2008, European fisheries caught 24.9 million t of sandeel in the North Sea. Atlantic herring constitute the second heaviest catch in the North Sea, with only 10.3 million t landed in the same time period (ICES, 2009). However, from 2003 onwards sandeel landings have decreased to between 0.2 and 0.4 million t per annum (ICES, 2009), and the spawning biomass of sandeels has been in decline over the past 15 years (ICES, 2008). However, as sandeel recruitment is density-dependent the reasons for such declines are apparently not related to the fisheries (van Deurs et al., 2009). Instead the decline in sandeel recruitment, and therefore numbers, has been associated with climatic changes which have altered the distribution of their prey, e.g. *Calanus finmarchicus* (van Deurs et al., 2009). Although the decline in sandeel numbers has not been associated with over-fishing, such decreases may result in sandeel abundance crossing a critical limit. Once sandeel numbers have fallen below this point, despite their density-dependent recruitment, additional pressure from fishing may cause further reductions in the spawning stock and therefore recruitment. This leaves the sandeel population more vulnerable to environmental changes (Brander, 2005) and therefore with increased potential for a population crash. The result of such a crash would be devastating to the North Sea ecosystem. For this reason in 2000, concern over decreasing sandeel catches and reduced breeding success in seabirds led to the closure of the sandeel fisheries off the east coast of Scotland (Greenstreet et al., 2006).

Sandeels, with such complex interactions within the food web, demonstrate the importance of considering sustainability from an ecosystem perspective rather than a single species stock assessment. In addition, ecosystem and multispecies modelling techniques allow predation rates on sandeels to be updated using time dependent information on their predators' population dynamics. van Deurs et al. (2009) noted a discrepancy between the higher model predictions and the lower observed sandeel recruitment in 2002 which they suggested may be attributed to "unusually high predation pressure" from herring. Such suspicions must lead us to not only monitor the fishing pressure on this species but also the pressures from its predators. However, such multispecies models, which can account for these predation pressures, have an increased complexity and therefore increased risks; that the misspecification of one species' population parameters may not only invalidate results for that species but potentially every other species in the model which interacts with it (Vinther, 2001). The success of using multispecies modelling to aid management decision making (e.g. fishing restrictions) therefore relies on accurate knowledge of the interactions between species.

The interaction that forms the subject of this study is between the sandeel and the grey seal (*Halichoerus grypus*, Fabricius). Grey seal populations increased threefold between 1985 and 2002, incurring a similar increase in their estimated annual consumption of sandeels (Hammond and Grellier, 2006). It is estimated that during 2002 sandeels made up the largest proportion of the grey seal's diet with around 69,000 t being consumed annually in the North Sea, the second most common prey species was cod with an estimated consumption of 8,300 t (Hammond and Grellier, 2006). Although sandeels undoubtedly make up a large proportion of the grey seal diet, the importance of commercial species is a matter of debate (Prime and Hammond, 1990; Hammond et al., 1994a,b). *A. marinus* is the most abundant species of sandeel in

the North Sea and the most accessible to the fisheries due to its preference for sandy substrates. It therefore makes up over 90% of the commercial catch (Jensen et al., 2003). However, some of the sandeels consumed by grey seals are too large to have been *A. marinus* (Prime and Hammond, 1990; Hammond et al., 1994a,b). In addition, there is evidence of bimodality in the size distribution of sandeel otoliths, and therefore the sandeels, consumed by grey seals. Two hypotheses have been proposed to explain this bimodality in size distribution: the sandeels consumed were from two or more distinct species (Prime and Hammond, 1990; Hammond et al., 1994a,b); or the sandeels consumed by grey seals may represent two distinct age groups of *A. marinus* (Macer, 1966; CEFAS, 2007). One objective of this study is to investigate these hypotheses and try to quantify the levels of consumption of *A. marinus*, the species of interest to commercial fisheries, compared to the consumption of other species of sandeel. In doing so it will then be possible to consider the fishing and predation pressures which *A. marinus* is subjected to independently of other species of sandeel. This is vital as such a large proportion of the fisheries catch is *A. marinus*.

In addition to *A. marinus* there are another four species of sandeel that are thought to occur in the North Sea: *A. tobianus* (Raitt), *Gymnammodytes semisquamatus* (Jourdain), *Hyperoplus lanceolatus* (Le Sauvage) and *H. immaculatus* (Corbin). However, the data available on the geographical distribution and size ranges of these species is limited. *A. marinus* is widely accepted as the most abundant species of sandeel in the North Sea (Macer, 1966) with the second most abundant thought to be *H. lanceolatus* (Pearson, 1968; Rijnsdorp et al., 1996). Although the presence of *A. tobianus* (Pearson, 1968), and *G. semisquamatus* (Macer, 1966) has also been documented in the North Sea we do not consider these species in this chapter as there is no evidence to suggest that they make up any more than a very small proportion of the sandeels in the North Sea.

It is thought that *A. marinus* may reach a maximum length of 25 cm (Macer, 1966; Froese and Pauly, 2009); this is consistent with data from the Danish fishery. *H. lanceolatus*, otherwise known as the greater sandeel, can reach a length of up to 40 cm, but is more commonly around 20 cm (Froese and Pauly, 2009). It is especially important for this study to obtain an estimate of the plausible size range for *A. marinus* and, more specifically, a plausible size range for the otoliths belonging to *A. marinus*, which are used to identify the presence of this species in the grey seal diet. Previous diet studies have relied on comparing estimated fish length with historical data on observed lengths of *A. marinus*. However, these estimated fish lengths are based on otolith dimensions which are, in fact, only estimates of original dimensions based on partly digested otoliths. Therefore, there are two sources of potential error involved in these comparisons. Here we compare measurements of the lengths of otoliths from over 2000 specimens of *A. marinus*. These are compared with the estimated undigested sizes of otoliths recovered from the grey seal scats throughout the North Sea.

Although the analysis of hard parts recovered from scats is probably the best single method of diet analysis for grey seals (Prime and Hammond, 1990; Hammond et al., 1994a), it presents some challenges. Estimating the actual number and biomass of prey consumed is difficult, because not all otoliths ingested by the seals are recovered in the scats (Bowen, 2000; Grellier and Hammond, 2006) and the consumed biomass needs to be estimated from the sizes of the recovered otoliths. A number correction factor (NCF) can be used to predict the number of otoliths ingested from the number recovered in the scats. The estimated NCF for *A. marinus* is approximately 3 (Grellier and Hammond, 2006). However, Grellier and Hammond (2006) observed that the proportions of consumed sandeel otoliths that were subsequently recovered varied

from 0.11 to 0.51. Some of this variation could be explained by difference in NCF between otoliths of different size. Ignoring this factor could lead to an under-estimation of the number of smaller otoliths (most likely belonging to *A. marinus*) ingested by the seals. This study therefore includes the estimation of a size-dependent NCF that accounts for the loss of completely digested otoliths and more accurately predicts the size distribution of otoliths ingested.

An estimate of the biomass of prey consumed is more useful for comparing diet composition than an estimate of the number of prey. This biomass, can be estimated using the relation between fish mass and otolith size. In some cases a direct power relation between these variables has been used (Hammond and Grellier, 2006), but otolith size has also been used to predict the length of the fish consumed and this length has then been used to predict their mass (Lewis et al., 2003). A concern with the North Sea scat data is that the wrong relations may be used in these calculations if all the otoliths are treated as if they are from *A. marinus*, when the seals are actually consuming a range of sandeel species. Calculating *A. marinus* biomass from otoliths of fish that are outside the expected range for this species could cause severe bias. This study therefore compares the fish mass to otolith size relations for *A. marinus* and *H. lanceolatus* and highlights the importance of choosing the correct relation.

2.1.1 Chapter Objectives

In summary, the objectives of this chapter are:

- To quantify the amount of *A. marinus* consumed by grey seals so that this species can be considered by itself in multispecies models. This is important as this species alone makes up most of the fisheries catch and it is therefore vital to consider its management separately to the management of other species of sandeel.

- To more accurately quantify the total biomass of all sandeels consumed by the grey seal and investigate potential sources of bias. This is important to assess the impact of the grey seal on on its various prey species. The grey seal is not only a major predator in the North Sea but its population size, and therefore potential impact on the ecosystem, is also increasing.

2.2 Data

2.2.1 North Sea Scat Analysis

Hammond and Grellier (2006) sampled grey seal scats from various haul-out sites throughout the North Sea in 2002 to study grey seal diet composition. The haul-out sites were grouped into five locations: Orkney, Shetland, Moray Firth, Donna Nook and central North Sea (which encompasses haul-out sites on the UK east coast between Abertay and the Farnes). The Orkney, Donna Nook and the central North Sea locations were visited quarterly throughout the year; however, the second quarter was missed at Shetland and only the third quarter was sampled at the Moray Firth. The data of particular interest in this study are the length measurements of the sandeel otoliths recovered from these scats. The otoliths were measured and graded to indicate the extent of digestion: 1 - pristine, 2 - moderately digested, 3 - considerably digested. The data analysed here come solely from grade 2 and 3 otoliths; no otoliths recovered from the feeding experiments in Grellier and Hammond (2006) or those recovered from the North Sea were in pristine condition. To obtain the estimated lengths of the otoliths recovered from the North Sea prior to partial digestion, values which henceforth shall be referred to as x_1, the lengths of the otoliths recovered from the North Sea were multiplied by grade specific digestion coefficients (DC) from Grellier and Hammond (2006) (Table 2.1).

Table 2.1: Grade specific digestion coefficients for grade 2 and 3 otoliths (Grellier and Hammond, 2006). The standard errors of the digestion coefficient estimates are given in parenthesis.

Grade	Otolith Width DC (se)	Otolith Length DC (se)
2	1.22 (0.046)	1.25 (0.041)
3	1.68 (0.030)	1.58 (0.034)
All	1.65 (0.030)	1.56 (0.033)

2.2.2 Otolith Recovery Rates

Grellier and Hammond (2006) and Tollit et al. (1997) performed a series of experiments to estimate the probability of recovering an ingested otolith in a seal scat. Grellier and Hammond (2006) recorded the proportion of otoliths recovered from sandeels which were fed to five captive grey seals (Table 2.2). However, these experiments were not designed to look at how recovery rates varied with sandeel otolith size, and the individual meals fed to each of the seals had very similar mean otolith dimensions. This made it difficult to estimate a recovery function from these data alone. Tollit et al. (1997) carried out a similar feeding experiment with harbour seals (*Phoca vitulin*), in this case the experiment was specifically designed to test for varying recovery probability based on otolith size (Table 2.3). The estimated proportions from these data, from both sets of feeding experiments, are henceforth referred to as x_2. We present the standard errors of these estimated proportions but did not incorporate them into the analyses for reasons discussed in section 2.5.5. Although there is some debate as to whether or not feeding experiments on harbour seals can be extrapolated to grey seals, Bowen (2000) found no significant differences in the recovery rates of otoliths from Atlantic and Pacific herring for these two seal species.

Table 2.2: Proportion of sandeel otoliths recovered and their mean length and width (Grellier and Hammond, 2006). Standard errors of the proportion estimates are given in parenthesis.

Seal Name	Undigested Mean Length (mm)	Undigested Mean Width (mm)	Proportion Recovered (se)
Lola	2.24	1.18	0.42 (0.017)
Q	2.21	1.18	0.27 (0.016)
Tess	2.22	1.18	0.51 (0.022)
Ulrika	2.18	1.17	0.11 (0.014)
Vera	2.78	1.40	0.44 (0.020)

Table 2.3: Proportion of sandeel otoliths recovered and their mean length and width (Tollit et al., 1997). Standard errors of the proportion estimates are given in parenthesis.

Undigested Mean Length (mm)	Undigested Mean Width (mm)	Proportion Recovered (se)
1.70	0.90	0.14 (0.033)
2.59	1.29	0.15 (0.020)
2.93	1.43	0.34 (0.029)
3.26	2.58	0.28 (0.026)

2.2.3 Relation between body size and otolith size for *A. marinus*

Engelhard et al. (2008) sampled *A. marinus* from the Dogger Bank region in the North Sea during April and May of 2004, 2005 and 2006, July 2004, and September/October 2005 and 2006 (CEFAS, 2007). Dredge tows were used to collect the sandeels during the night when they were buried in the sediment. The lengths of the fish were recorded to the nearest 0.5cm and their otoliths were extracted to establish the age of the fish. Fish were chosen to provide similar sample sizes across the range of observed fish lengths. The length and width of 2169 of these otoliths, one from each fish, were measured to the nearest 0.01 mm using an image analysis microscope.

2.3 Methods

2.3.1 Modelling the Consumption of Sandeels

A mixture model combined with a size-related probability of recovery was used to predict the distribution of otoliths ingested by the seals using the sizes of otoliths recovered from the scats. It was hoped that this mixture of distributions would make it easier to classify the otoliths to either the smaller *A. marinus*, or the larger *H. lanceolatus*. The distribution of otolith sizes recovered from the scats appeared to be right-skewed, suggesting that lognormal distributions would be appropriate to model these data. Lognormal distributions are commonly used to represent biological data and are especially useful when modelling size distributions (Limpert et al., 2001). In addition, whereas the size distribution of fish and therefore the size distribution of otoliths for different age classes tend to be normally distributed (Fournier et al., 1998; Rindorf and Lewy, 2001), the sum of these normal distributions across the age classes in the population often gives rise to log-normal distributions. Therefore, this approach is more appropriate for the identification of different species of sandeel rather than different age groups within species. A mixing parameter allowed different weights to be allocated to each log-normal distribution describing the proportion of otoliths thought to have come from each distribution.

Various functions were considered to model the probability of recovery. Given the sparse nature of the data to which the curve was fitted it was especially important to choose a function which was biologically plausible. The first requirement was that the function should approach zero as the otolith size approaches zero. This is due to the belief that the smaller the ingested otolith the less the likelihood of recovery, a phenomenon which has been observed across species (Tollit et al., 1997; Bowen, 2000; Grellier and Hammond, 2006). In addition, the methods used to extract the otoliths

from the scats/tank rely on using a mesh which sieves them out of the scats/water, this means that any otolith which could fit through this mesh would stand very little chance of being recovered. The second requirement was that the curve should approach an upper asymptote between 0 and 1 as otolith size increases. The Gompertz curve meets these criteria and is less restrictive in shape than the logistic because it does not enforce 180 ° symmetry around the point of inflexion. This meant that while the Gompertz curve predicted an upper asymptote which was consistent with what was expected based on assuming a constant probability of recovery, the logistic curve predicted what appeared to be unrealistically high probabilities of recovery for the larger otoliths. Fitting this curve at the same time as the mixture model improves the estimation of the curve's parameters as the observed otolith sizes in the scats provide additional information regarding the shape of the curve. For example, if 100 otoliths recovered from the scats were observed as having a length of 1.8 mm, it is highly unlikely that the probability of recovery is close to zero for otoliths of this size.

Based on the aforementioned considerations, this method makes the following assumptions:

1. The distribution of ingested otolith sizes is well approximated by a mixture of two lognormal distributions. Violation of this assumption would lead to incorrect prediction of the size distribution of otoliths ingested as well as potentially introducing bias in both the estimates of total biomass and the predicted biomass in each of the distributions.

2. The size-related probability of recovery can be well approximated by the Gompertz curve. Violation of this assumption would also lead to an incorrect prediction of the distribution of otoliths ingested and therefore also affect biomass estimates. An idea of the effects of violating this assumption can be seen by

comparing the traditional analyses, which assume a constant NCF across otolith size, with the results presented here which allow the NCF to be size dependent.

3. There is independence between all data points. This assumption seems valid for the x_2 data as each set of trials were carried out on a different seal. However, there may be spatial dependence in the x_1 data as it is highly likely that the otoliths from each seal scat were the result of the seal foraging in a specific area. In addition, there may be effects of prey selection (e.g. for the largest sandeels). Violation of this assumption would mean that the shape of the distribution of ingested otoliths would be different to what is expected - a mixture of lognormal distributions - which would have the same implications as the previous two assumptions.

2.3.1.1 Likelihood Specification

2.3.1.1.1 Lognormal Mixture Model

The probability density function (pdf) for a lognormal distribution of recovered otolith measurements x_{1j} for $j = 1, ..., n$ otoliths is shown in equation 2.1; μ and σ are the mean and standard deviation on the log scale for this distribution.

$$Pr\left(X = x_{1j}\ ;\ \mu,\ \sigma\right) = \frac{1}{x_{1j}\sigma\sqrt{2\pi}} exp\left[-\frac{\left(ln\left(x_{1j}\right) - \mu\right)^2}{2\sigma^2}\right] \tag{2.1}$$

Mixture models combine a number of distributions through summation; they are scaled using mixing parameters to ensure the whole distribution integrates to 1. To create a mixture model for the recovered otolith sizes from two lognormal distributions the lognormal pdfs are summed and scaled using a mixing parameter w_r, where $0 \leq w_r \leq 1$. Fitting this mixture model to just the data from the North Sea, without considering the recovery rate curve or the feeding experiment data, implies the

mixing parameter, w_r, represents the proportion of **recovered** otoliths, as opposed to the proportion of **ingested** otoliths, in the first lognormal distribution. As the observations are assumed to be independent, a likelihood function L_{LM} (equation 2.2) can be constructed as the product of the pdfs for each data point.

$$L_{LM} = \prod_{j=1}^{n} \left[w_r \; Pr\left(X = x_{1j} \; ; \mu_1, \sigma_1\right) + \left(1 - w_r\right) \; Pr\left(X = x_{1j} \; ; \mu_2, \sigma_2\right) \right] \tag{2.2}$$

To ensure only one maximum is found during the optimisation μ_2 was modelled as $\mu_1 + \mu_{add}$. Choosing to maximise the log likelihood l_{LM} (equation 2.3), makes the optimisation computationally easier.

$$\begin{aligned} l_{LM} &= ln\left(L_{LM}\right) \\ &= \sum_{j=1}^{n} \ln \left[w_r \; Pr\left(X = x_{1j} \; ; \theta_1\right) + \left(1 - w_r\right) \; Pr\left(X = x_{1j} \; ; \theta_2\right) \right] \\ \theta_1 &= \{\mu_1, \sigma_1\} \\ \theta_2 &= \{\mu_1, \mu_{add}, \sigma_2\} \end{aligned} \tag{2.3}$$

2.3.1.1.2 Binomial Likelihood

A binomial probability mass function (pmf) - equation 2.4 - was used to incorporate the recovery rate data from the feeding experiments. The observations m_i and r_i are the number of otoliths fed and recovered respectively for the i^{th} seal, p is the probability of recovery.

$$Pr\left(R = r_i\right) = \binom{m_i}{r_i} p^{r_i} \left(1 - p\right)^{m_i - r_i} \tag{2.4}$$

This pmf was modified (equation 2.5) to allow p to be a function of an explanatory variable (Welch and Foucher, 1988) x_{2i}, representing the mean size of the otoliths fed to the i^{th} seal.

$$Pr\left(R = r_i\right) = \binom{m_i}{r_i} f\left(x_{2i}\,;\,\theta_3\right)^{r_i}\left(1 - f\left(x_{2i}\,;\,\theta_3\right)\right)^{m_i - r_i} \tag{2.5}$$
$$\theta_3 = \{\alpha,\,\delta,\,\lambda\}$$

Size-related recovery (equation 2.6) was modelled using a Gompertz curve, where α represents the upper asymptote, and δ and λ influence the steepness and location of the inflexion. δ and λ were restricted to be positive to ensure the probability of recovery increases with otolith size.

$$p = f\left(x_{2i};\,\theta_3\right) = \alpha\; exp^{-exp^{\delta - \lambda x_{2i}}} \tag{2.6}$$

Because we assumed that our observations are independent, we again constructed a likelihood function L_{RR} (equation 2.7) as the product of the pmfs for $i = 1, ..., K$ seals.

$$\begin{aligned} L_{RR} &= \prod_{i=1}^{K}\; Pr\left(R = r_i\right) \\ &= \prod_{i=1}^{K}\; \binom{m_i}{r_i} f\left(x_{2i};\,\theta_3\right)^{r_i}\left(1 - f\left(x_{2i};\,\theta_3\right)\right)^{m_i - r_i} \end{aligned} \tag{2.7}$$

To simplify optimisation we again chose to maximise the log likelihood, l_{RR} (equation 2.8). In addition, we ignored the constant $\sum_{i=1}^{K}\; ln\binom{m_i}{r_i}$ because its value does not depend on the parameters of the Gompertz curve and it therefore does not affect the location of the maximum.

$$l_{RR} = ln(L) = \sum_{i=1}^{K}\; r_i ln\left(f\left(x_{2i};\,\theta_3\right)\right) + \left(m_i - r_i\right) ln\left(1 - f\left(x_{2i};\,\theta_3\right)\right) \tag{2.8}$$

2.3.1.1.3 Combined Model

Rather than just fitting a mixture model to the recovered otolith sizes, we can model the process which results in the observed size distribution of recovered otoliths. The process which gives rise to an observed otolith of size $x_1 j$ in the scats is assumed to be a multiplication of the probability that an otolith of size $x_1 j$ was ingested, multiplied by the probability that it was recovered given that it was ingested and of size $x_1 j$. The first of these processes can be described by a lognormal mixture model similar to that in equation 2.2, except now we define the mixing parameter as w_e rather than w_r. This is because we are now fitting a model which describes the whole process which led to an otolith of size $x_1 j$ being recovered in a scat rather than just fitting a mixture model to the recovered otoliths. The second process can be described using the size dependent recovery curve given in equation 2.6, which gives the probability of recovering an otolith dependent on its size given that it was ingested. This combined process is defined in equation 2.9.

$$P(X = x_{1j}) = \left[w_e \; Pr\left(X = x_{1j} \; ; \theta_1\right) + (1 - w_e) \; Pr\left(X = x_{1j} \; ; \theta_2\right)\right] f\left(x_{1j}; \theta_3\right) \quad (2.9)$$

The pdf of this process is then this combined function divided by its integral, and the likelihood and log-likelihood are therefore defined as shown in equations 2.10 and 2.11, respectively.

$$L_{LMRR} = \prod_{j=1}^{n} \frac{P(X = x_{1j})}{\int_0^\infty P(X = x_{1j}) dx} \quad (2.10)$$

$$l_{LMRR} = \sum_{j=1}^{n} ln \left[\frac{P(X = x_{1j})}{\int_0^\infty P(X = x_{1j}) dx}\right] \quad (2.11)$$

The two sets of otolith sizes, x_1 and x_2, are assumed to be independent of each other, and the combined (L_{LMRR}) and binomial (L_{RR}) likelihoods share common parameters. Therefore, we can construct a single likelihood (equation 2.12) as the product of these two likelihoods. Similarly, the overall log-likelihood (equation 2.13) is simply the sum of the two log-likelihood functions.

$$L_{full} = L_{LMRR}\, L_{RR} \tag{2.12}$$

$$l_{full} = l_{LMRR} + l_{RR} \tag{2.13}$$

2.3.1.2 Model Estimation

A separate model was fitted for each location and each season in the North Sea scat analysis dataset. The parameters associated with these models were estimated using maximum likelihood techniques. Although in theory all eight parameters could be estimated simultaneously, in practice to obtain convergence and improve model stability, the parameter λ had to be fixed within the model as a function of δ (equation 2.14). This relation was estimated from the binomial model by fitting solely to the experimental data, and seemed a reasonable approach as the parameter estimates for λ and δ had a correlation of 0.99.

$$\lambda = 0.132 + 0.574 * \delta \tag{2.14}$$

The parameters associated with these models were estimated by optimising the likelihoods using the `optim` routine in the `stats R` library. The variance-covariance matrix was obtained by inverting the Hessian matrix. The L-BFGS-B method (Bryrd et al., 1995) was chosen because it allows upper and lower bounds of the parameters to be specified. It was important to restrict the mixing parameter and the asymptote of the recovery function so they could only take values between 0 and 1. Bounds for

the other parameters were chosen to ensure that the function was finite within their ranges and wide enough so that neither boundary was chosen as a solution.

The standard deviations associated with the lognormal distributions were estimated using the log scale, thereby limiting these parameters to positive values. The estimated values of the shape parameter, δ, were larger than the other parameters and it was therefore also estimated on the log scale to ensure a biologically plausible solution and a valid variance-covariance matrix. Estimating parameters that are too dissimilar in scale can cause negative variances to be obtained when the Hessian matrix is inverted.

The full likelihood approach (equation 2.13) provided unstable and highly variable results for some of the models. Therefore, the mixture model part of the likelihood (equation 2.3) was fitted by itself to the North Sea data. This is henceforth termed the mixture model approach. Due to what was suspected to be artificially reduced variability in the Gompertz parameters estimated assuming λ was a function of δ, these parameters were estimated independently in this approach. Although this approach does not allow information from the North Sea data to influence the estimation of the Gompertz parameters, it gives more stable parameter estimates. However, in this approach the mixing parameter is not the proportion of ingested otoliths in one distribution compared with the other (w_e) as desired, it is the proportion of recovered otoliths in one distribution compared to the other (w_r). We therefore need a method of calculating w_e based on our estimate of w_r, Section 2.3.1.3.

2.3.1.3 Estimation of w_e for the Mixture Model Approach

The mixing parameter w_r estimated from fitting only the mixture model to the recovered otoliths (equation 2.3) refers to the proportion of otoliths estimated to be in

the first distribution of those otoliths recovered in the seal scats. As the two lognormal distributions contain otoliths of different lengths and the probability of recovery varies with otolith length, the proportion of otoliths ingested from each distribution will be slightly different to those recovered. Specifically, the probability an otolith was in the first lognormal distribution given that it was ingested (w_e) is different to the probability an otolith was in the first lognormal distribution given that it was recovered (w_r). However, we can use the probabilities that an otolith was recovered given that it was either in the first or second lognormal distribution to rescale w_r to give w_e. Each of these probabilities are obtained by calculating the integral of each lognormal distribution multiplied by the recovery rate curve. After rescaling both w_r and $(1 - w_r)$ to get the correct ratio of $P\left(dist_1 \mid ingest\right) : P\left(dist_2 \mid ingest\right)$ we must then rescale these values so that they sum to 1, and therefore give us w_e. This process is described in equation 2.15.

$$\begin{aligned}
\widehat{w}_e &= P\left(dist_1 \mid ingest\right) = \frac{\phi_1}{\phi_1 + \phi_2} \\
\phi_1 &= \frac{P\left(dist_1 \mid rec\right)}{P\left(rec \mid dist_1\right)} = \frac{\widehat{w}_r}{\int_0^\infty Pr\left(X = x\,;\widehat{\theta}_1\right)\;\; f\left(x\,;\widehat{\theta}_3\right)\; dx} \\
\phi_2 &= \frac{P\left(dist_2 \mid rec\right)}{P\left(rec \mid dist_2\right)} = \frac{1 - \widehat{w}_r}{\int_0^\infty Pr\left(X = x\,;\widehat{\theta}_2\right)\;\; f\left(x\,;\widehat{\theta}_3\right)\; dx}
\end{aligned} \tag{2.15}$$

2.3.1.4 Estimation of the Number of Otoliths Consumed

The methodology described thus far, only predicts the **distribution** of otoliths ingested by the seals. Further calculations are required to obtain an estimate of the **number** of otoliths ingested. This unknown value will be termed N. The estimates of the number of otoliths from each of the two distributions, N_1 and N_2, can be estimated as $N \times w_e$ and $N \times (1 - w_e)$, respectively.

Two methods were considered for estimating N. The first is based on the standard Horvitz-Thomson estimator (Horvitz and Thompson, 1952; Buckland et al., 2004), which will henceforth be termed the HT standard estimator. If we define $\widehat{P}_j$ as the estimated probability of recovery for the j^{th} otolith x_{1j}, each of these probabilities is obtained using the estimated size-dependent probability of recovery function (equation 2.16). The sum of the inverse of these proportions gives an estimate of N (equation 2.17).

$$\widehat{P}_j = \widehat{\alpha}\ exp^{-exp^{\widehat{\delta}-\widehat{\lambda}x_{1j}}} \tag{2.16}$$

$$\widehat{N} = \sum_{j=1}^{n} \frac{1}{\widehat{P}_j} \tag{2.17}$$

An alternative approach uses the integral of the entire process, rather than the individual probabilities from the recovery function. By definition, the lognormal mixture model integrates to 1. However, when it is multiplied by the proportional recovery function we are multiplying the whole distribution by values less than 1; the integral is therefore reduced and is equal to the proportion of otoliths ingested and subsequently recovered in the seal scats. The inverse of this proportion multiplied by the number of otoliths recovered in the scats gives an alternative Horvitz-Thompson-like estimator (equation 2.18), which will henceforth be termed the HT integral estimator.

$$\widehat{N} = \frac{n}{\int_0^\infty P(X = x_{1j})dx} \tag{2.18}$$

2.3.1.5 Variance Estimation

A parametric bootstrap was used to quantify the variability in the estimated number of ingested otoliths and the corresponding sandeel biomass. A multivariate normal distribution was used to resample the model parameters based on the estimated values

and the variance-covariance matrix. Ninety-five percent confidence intervals where then obtained by extracting the 2.5 and 97.5 percentiles.

2.3.2 Simulation

Simulated data were used to evaluate the effectiveness and robustness of the methods proposed in section 2.3.1. Otolith lengths were randomly drawn from two lognormal distributions to represent the ingested otoliths, some of these otoliths were then selected as recovered using a binomial distribution. The probability of recovery was chosen based on the Gompertz curve fitted to the experimental data. The distributions of otolith lengths in the simulations were located with respect to the probability of recovery as determined by the Gompertz curve. The distribution of ingested otolith lengths in the first set of simulations (A) was chosen so that the whole range of otolith sizes had enough probability of recovery that the entire distribution of otolith lengths was represented in the simulated recovered data (i.e. although the recovery rate varied across the range of otolith sizes, no part of the distribution was lost entirely). In a second set of simulations (B) the distribution of otolith lengths was chosen so that the smallest otoliths had almost no probability of recovery, causing the lower end of the distribution of ingested otoliths to be lost. The simulations involved randomly generating 9999 datasets from each of the two sets of parameters and using both the full likelihood (FL) and mixture model (MM) approach to estimate the two sets of lognormal parameters, as well as the number and biomass of ingested sandeels. The true biomass and biomass estimates were calculated based on the fish mass to otolith length relation from Leopold et al. (2001).

Note that in these simulations we do not reduce the size of the recovered otoliths as if they had been partially digested. This is because the otolith dimensions in the North Sea scat analysis data are corrected for the effects of digestion prior to analysis.

2.3.3 Modelling the relations between otolith dimensions and fish length

The specimens of *A. marinus* collected from the Dogger Bank were used to determine the relation between otolith size and fish length. Otolith lengths and widths were modelled as a function of fish size and age. Generalised additive models (GAMs) were used to test for non linearities and generalised least squares models were used in the case of non-constant variance. The best model was chosen based on the minimum Bayesian information criterion (BIC) (Schwarz, 1978). These models were also compared to previously published relations (Tollit et al., 1997; Grellier and Hammond, 2006).

2.4 Results

2.4.1 Simulation Results

Both the FL and the MM approach produce an accurate fit to the data for both simulation A and simulation B. The fit of the model to the simulated data is indicated by a comparison of the fitted model to a histogram of the simulated data in Figures 2.1-2.4(a).

The Q-Q plots (Figs. 2.1-2.4(b)) allow us to compare data generated from the hypothesized process with that which was observed. In the case of the FL approach this is the fitted model, because the recovery curve is estimated at the same time as the mixture model. However, in the case of the MM approach the hypothesized process combines the fitted mixture model with the adjusted weighting parameter, w_e, and the separately estimated recovery curve to generate the theoretical quantiles. Based on these Q-Q plots it appears that both approaches provide plausible processes which may have given rise to the simulated data in both simulations; only the MM approach

appears to over-estimate the sizes of a few of the largest otoliths in simulation B 2.3(b).

Both approaches are able to accurately predict the shape of the distribution of ingested otoliths for simulation A (Figs 2.1(c) and 2.2(c)). The mean parameter estimates in Table 2.4 confirm this observation. However, the histogram of the mean parameter estimates show that there is some bias in the estimates of μ_1 and σ_1 for the MM approach (Figs. 2.5(a) and 2.5(b)). The FL approach provides less biased estimates of μ_1 and σ_1 for simulation A (Figs. 2.6(a) and 2.6(b)). Both approaches produce unbiased estimates of μ_2 and σ_2 (Figs. 2.5-2.6, (c) and (d)).

The largest discrepancies between truth and prediction for simulation B are in the distribution of smaller otoliths (Figs. 2.3(c) and 2.4(c)). The MM approach over-estimates μ_1 and under-estimates σ_1, whereas the FL approach under-estimates μ_1 (Table 2.4). More serious problems with the FL approach are seen in the confidence intervals and the histograms of the parameter estimates for μ_1 and σ_1 (Figs. 2.8(a) and 2.8(b)). Although, on average, the mean estimates of μ_1 and σ_1 are less biased for the FL approach than the MM approach, the estimates have a skewed distribution. In contrast, although the estimates of μ_1 and σ_1 from the MM approach are biased they have a symmetric distribution (Figs. 2.7(a) and 2.7(b)). Despite these problems, the estimates of μ_2 and σ_2 for both approaches provide unbiased and consistent estimates of the true parameters (Table 2.4, Figs. 2.7 and 2.8, (c) and (d)).

The Gompertz curve, estimated using only the recovery rate data obtained from the feeding experiments, is shown in Figures 2.1(d) and 2.3(d). Figures 2.2(d) and 2.4(d) show the Gompertz curve estimated simultaneously from the simulated data and the recovery rate data. The associated parameter estimates are given in Table 2.5, where λ for the FL approach is estimated as a function of δ (equation 2.14). The FL approach provides accurate estimates of the Gompertz curve parameters.

However, the variability associated with these parameter estimates is smaller than the variability estimated using only the binomial likelihood (Table 2.5).

Table 2.4: A comparison of the true parameters used to generate the data and the parameter estimates from both the mixture model (MM) and the full likelihood (FL) approach. The 2.5 and 97.5 percentiles of the parameter estimates from the 9999 repetitions are given in parenthesis.

	$\boldsymbol{\mu_1}$	$\boldsymbol{\sigma_1}$	$\boldsymbol{\mu_2}$	$\boldsymbol{\sigma_2}$
Truth A	0.90	0.12	1.30	0.15
MM Estimates	0.90 (0.89,0.92)	0.11 (0.11,0.12)	1.30 (1.27,1.32)	0.15 (0.14,0.16)
FL Estimates	0.90 (0.89,0.91)	0.12 (0.11,0.13)	1.30 (1.28,1.32)	0.15 (0.14,0.16)
Truth B	0.50	0.12	1.12	0.20
MM Estimates	0.58 (0.58,0.60)	0.09 (0.08,0.10)	1.12 (1.10,1.13)	0.20 (0.19,0.21)
FL Estimates	0.49 (0.36,0.53)	0.12 (0.10,0.17)	1.12 (1.10,1.14)	0.20 (0.19,0.21)

Figures 2.1 to 2.4: (a) The fitted model plotted over a histogram of the simulated data, (b) Quantile-Quantile plot of the simulated otoliths lengths versus points generated from the hypothesized process, (c) The mean (thick dashed line) and 95% confidence interval (dotted line) of the predicted distribution of ingested otoliths compared with truth (solid line); a histogram of recovered otoliths is given for reference, (d) The binomial recovery curve, the crosses represent the experimental data (Tables 2.2 and 2.3), with the distribution of recovered otoliths for reference.

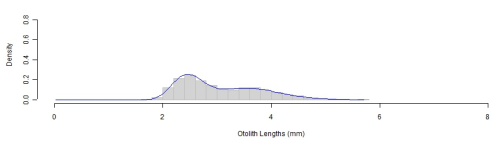

(a) Fitted Model

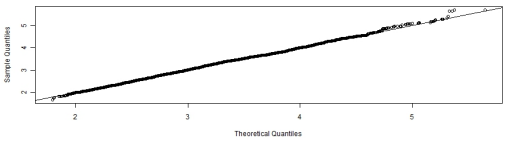

(b) Q-Q Process Plot

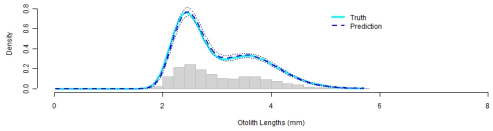

(c) Truth versus Predicted Ingested

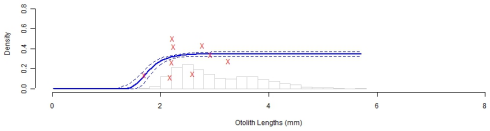

(d) Binomial Recovery Component

Figure 2.1: Example results for simulation A using the mixture model approach.

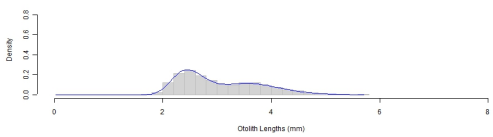

(a) Fitted Model

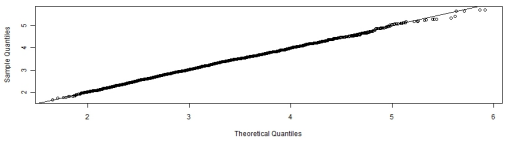

(b) Q-Q Process Plot

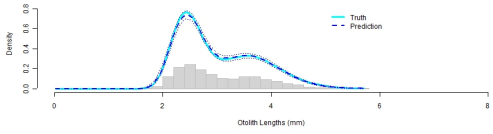

(c) Truth versus Predicted Ingested

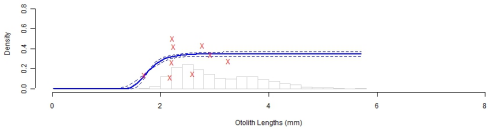

(d) Binomial Recovery Component

Figure 2.2: Example results for simulation A using the full likelihood approach.

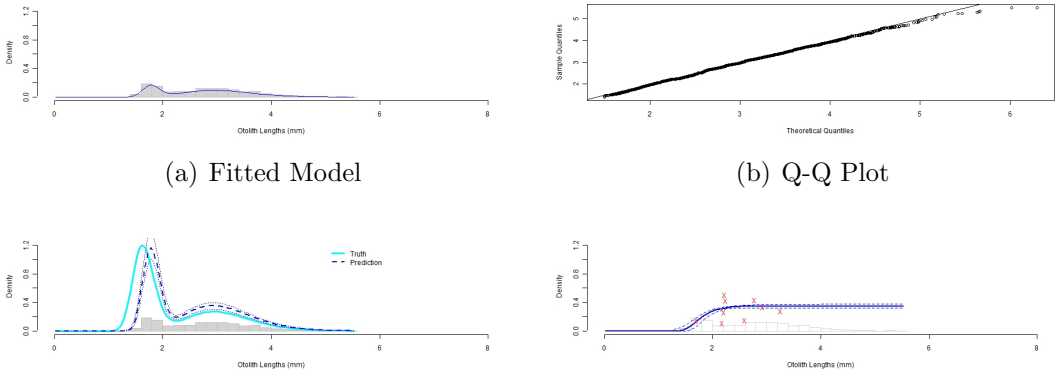

(a) Fitted Model (b) Q-Q Plot

(c) Truth versus Predicted Ingested (d) Binomial Recovery Component

Figure 2.3: Example results for simulation B using the mixture model approach.

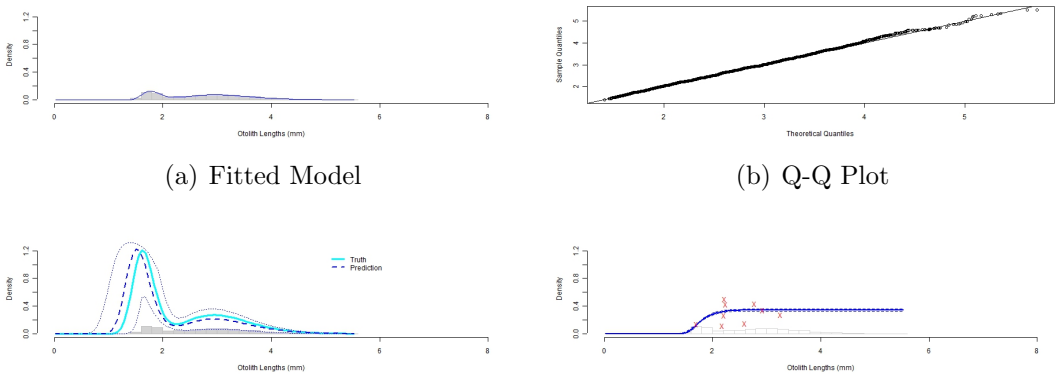

(a) Fitted Model (b) Q-Q Plot

(c) Truth versus Predicted Ingested (d) Binomial Recovery Component

Figure 2.4: Example results for simulation B using the full likelihood approach.

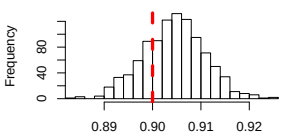

(a) Mean estimates of μ_1

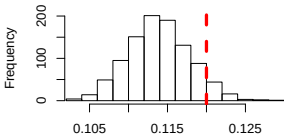

(b) Mean estimates of σ_1

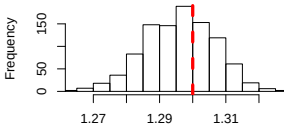

(c) Mean estimates of μ_2

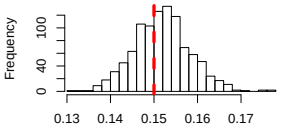

(d) Mean estimates of σ_2

Figure 2.5: Mean parameter estimates for the 9999 replicates of simulation A estimated using the mixture model approach. Truth is indicated by the dashed line.

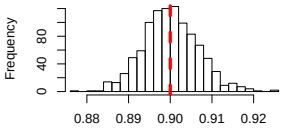

(a) Mean estimates of μ_1

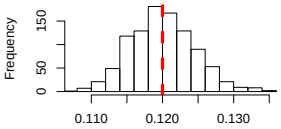

(b) Mean estimates of σ_1

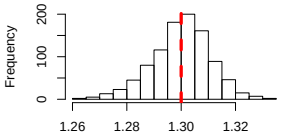

(c) Mean estimates of μ_2

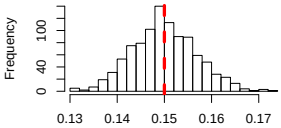

(d) Mean estimates of σ_2

Figure 2.6: Mean parameter estimates for the 9999 replicates of simulation A estimated using the full likelihood approach. Truth is indicated by the dashed line.

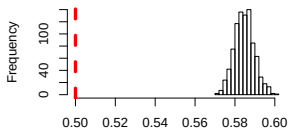

(a) Mean estimates of μ_1

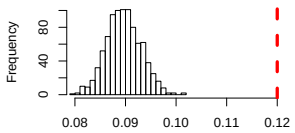

(b) Mean estimates of σ_1

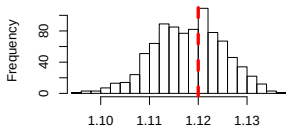

(c) Mean estimates of μ_2

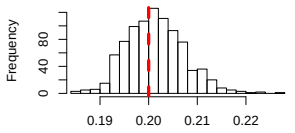

(d) Mean estimates of σ_2

Figure 2.7: Mean parameter estimates for the 9999 replicates of simulation B estimated using the mixture model approach. Truth is indicated by the dashed line.

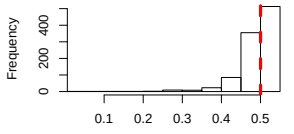

(a) Mean estimates of μ_1

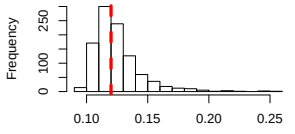

(b) Mean estimates of σ_1

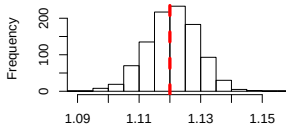

(c) Mean estimates of μ_2

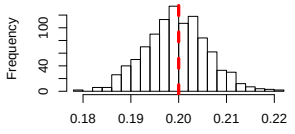

(d) Mean estimates of σ_2

Figure 2.8: Mean parameter estimates for the 9999 replicates of simulation B estimated using the full likelihood approach. Truth is indicated by the dashed line.

Table 2.5: A comparison of the Gompertz curve parameters used to generate the simulated data (those estimated using the binomial likelihood) and those estimated in the full likelihood (FL) approach for simulations A and B. The 2.5 and 97.5 percentiles of the parameter estimates from the 9999 repetitions are given in parenthesis.

	α	δ	λ
Binomial Likelihood	0.35 (0.32,0.37)	8.12 (4.47,11.82)	4.79 (2.74,6.87)
FL estimates (Sim A)	0.35 (0.34,0.35)	8.16 (7.40,9.07)	4.81 (4.38,5.34)
FL estimates (Sim B)	0.35 (0.34,0.36)	8.23 (6.81,10.19)	4.86 (4.04,5.98)

The results for simulation A in Table 2.6 suggest that both the HT integral and the HT standard methods provide accurate estimates of N, N_1 and N_2, for both the MM and the FL approach. In addition, the 95% percentile intervals for all four sets of estimates are largely comparable with one another (Table 2.6) and display a symmetric distribution of estimates (Figs. 2.9 and 2.10). However, the histograms of the parameter estimates indicate that the MM approach for both HT estimators slightly under-estimates N_1 (Figs. 2.9(b), 2.9(e)).

The estimates of N, N_1 and N_2 for simulation B are less accurate and less consistent. The HT standard estimator provides the most accurate estimates of N; there was little difference between the MM and the FL approach (Table 2.6). However, Figures 2.11 and 2.12 indicate that there are some problems with the HT standard estimator because the distributions of the estimates are severely skewed. In the extreme cases, this method provides estimates of N at around 50,000, in the case of the MM approach, and 22,000, with the FL approach (Figs. 2.11(d) and 2.12(d)). The same problem is seen in the HT integral estimates of N for the FL approach (Fig. 12(a)). Only the HT integral estimates of N for the MM approach have a symmetric distribution, but they under-estimate N (Fig. 2.11(a)).

Estimation of N_1 is especially problematic for simulation B. While the HT standard method with the FL approach provides a relatively unbiased mean estimate of N_1, the distribution of the estimates is severely skewed (Table 2.6, Fig. 2.12(e)). In contrast, the HT integral estimates from the MM approach provide a symmetric but biased distribution of estimates (Fig. 2.11(b)).

The HT integral method with the FL approach gives the most accurate estimates of N_2; there is no sign of bias and the estimates have a symmetric distribution (Fig. 2.12(c)).

Table 2.6: A comparison of the true number of otoliths ingested along with the estimated numbers from both the mixture model (MM) and the full likelihood (FL) model. The two methods of estimation presented in section 2.3.1.4 are also compared; HT standard and HT integral. The 2.5 and 97.5 percentiles of the estimates are given in parenthesis.

	Estimator	$\boldsymbol{N}$	$\boldsymbol{N_1}$	$\boldsymbol{N_2}$
Truth A		9000	5000	4000
MM Estimates	HT Integral	8964 (8714,9244)	4886 (4524,5245)	4078 (3727,4426)
	HT Standard	8993 (8745,9270)	4901 (4540,5266)	4091 (3739,4436)
FL Estimates	HT Integral	8983 (8717,9257)	4993 (4598,5354)	3991 (3638,4370)
	HT Standard	8982 (8720,9254)	4992 (4598,5353)	3990 (3637,4368)
Truth B		11000	6500	4500
MM Estimates	HT Integral	8565 (8205,8949)	3976 (3674,4295)	4589 (4377,4803)
	HT Standard	11046 (9965,13305)	5129 (4462,6279)	5917 (5352,7137)
FL Estimates	HT Integral	12097 (9557,20752)	7606 (5097,16176)	4491 (4267,4704)
	HT Standard	11135 (9600,14912)	6843 (5133,11602)	4292 (2962,4863)

The results of the estimation of ingested biomass are directly related to the number and size of the ingested otoliths. Both approaches, as well as both HT estimators, provide accurate estimates of ingested biomass for simulation A (Table 2.7). The results for simulation B are more interesting. Although the estimates of N from the MM approach combined with the HT integral estimator are biased, the estimates of B are less biased and have a symmetric distribution (Fig. 2.13(a)). In fact, the 97.5

percentile corresponds to the true value (Table 2.7). However, there is still a large amount of bias in the estimates of B_1 (Fig. 2.13(b)). The HT standard estimator using the MM approach provides less biased estimates of B_1, however, these are severely skewed (Fig. 2.13(e)). The HT standard estimator using the FL approach provides the best estimation of B_1 (Fig. 2.14(e)), although there is still some bias and skew in the distribution of estimates. Of the other FL estimates of B, B_1 and B_2, only the HT integral estimates of B_2 appear unbiased and without a skewed distribution (Fig 2.14).

Table 2.7: A comparison of the true biomass of sandeels in kg ingested along with the estimated biomass from both the mixture model (MM) and the full likelihood (FL) model. The two methods of estimation presented in section 2.3.1.4 are also compared; HT standard and HT integral. The 2.5 and 97.5 percentiles of the estimates are given in parenthesis.

	Estimation	$\boldsymbol{B_{total}}$	$\boldsymbol{B_1}$	$\boldsymbol{B_2}$
Truth A		65.6 (64.9,66.5)	20.2 (20.0,20.4)	45.5 (44.7,46.2)
MM Estimates	HT Integral	65.8 (63.3,68.4)	19.9 (17.8,22.0)	45.9 (42.9,48.9)
	HT Standard	66.1 (63.6,68.7)	20.0 (17.9,22.2)	46.1 (43.2,49.2)
FL Estimates	HT Integral	65.5 (62.9,68.2)	20.2 (18.1,22.3)	45.3 (42.2,48.4)
	HT Standard	65.5 (63.1,68.1)	20.2 (18.0,22.4)	45.3 (42.3,48.2)
Truth B		43.9 (43.2,44.7)	9.6 (9.5,9.7)	34.3 (33.6,35.1)
MM Estimates	HT Integral	42.0 (40.0,43.9)	7.1 (6.5,7.7)	34.9 (33.0,36.8)
	HT Standard	54.2 (48.8,65.5)	9.2 (8.0,11.3)	45.0 (40.3,54.7)
FL Estimates	HT Integral	44.7 (41.1,53.0)	10.4 (8.0,17.5)	34.3 (32.4,36.3)
	HT Standard	42.4 (34.5,47.9)	9.7 (8.0,12.6)	32.8 (22.9,37.2)

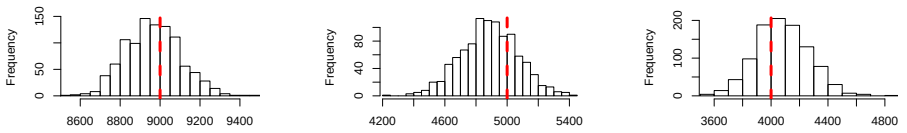

(a) HT integral estimates of N (b) HT integral estimates of N_1 (c) HT integral estimates of N_2

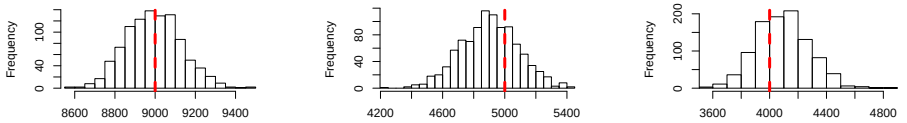

(d) HT standard estimates of N (e) HT standard estimates of N_1 (f) HT standard estimates of N_2

Figure 2.9: Mixture model mean estimates of N, N_1 and N_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation A. Truth is indicated by the dashed line.

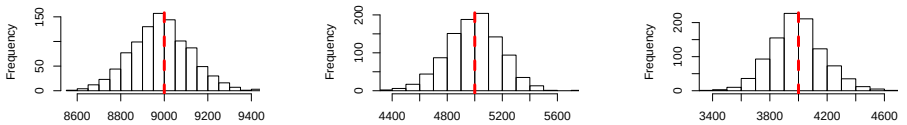

(a) HT integral estimates of N (b) HT integral estimates of N_1 (c) HT integral estimates of N_2

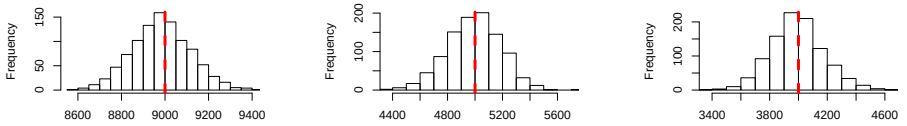

(d) HT standard estimates of N (e) HT standard estimates of N_1 (f) HT standard estimates of N_2

Figure 2.10: Full likelihood mean estimates of N, N_1 and N_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation A. Truth is indicated by the dashed line.

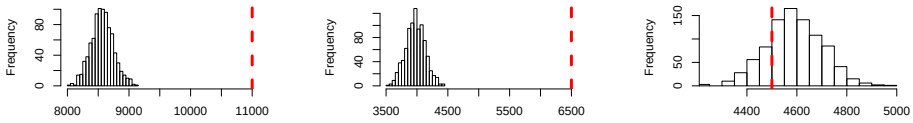

(a) HT integral estimates of N (b) HT integral estimates of N_1 (c) HT integral estimates of N_2

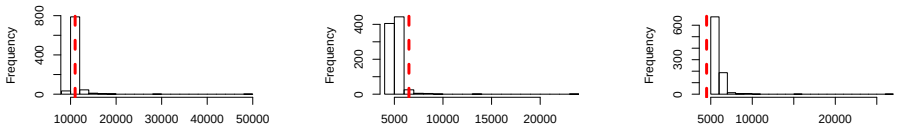

(d) HT standard estimates of N (e) HT standard estimates of N_1 (f) HT standard estimates of N_2

Figure 2.11: Mixture model mean estimates of N, N_1 and N_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation B. Truth is indicated by the dashed line.

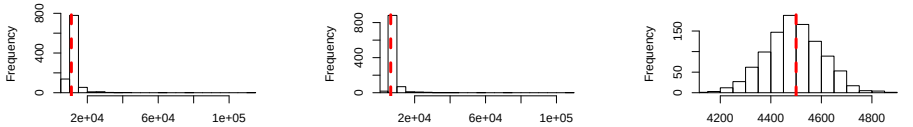

(a) HT integral estimates of N (b) HT integral estimates of N_1 (c) HT integral estimates of N_2

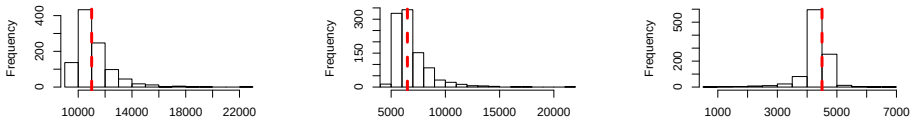

(d) HT standard estimates of N (e) HT standard estimates of N_1 (f) HT standard estimates of N_2

Figure 2.12: Full likelihood mean estimates of N, N_1 and N_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation B. Truth is indicated by the dashed line.

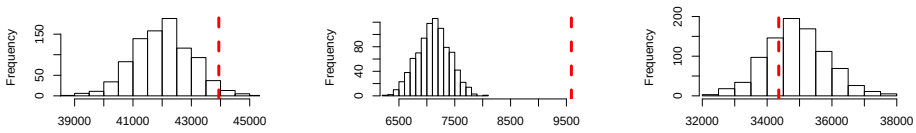

(a) HT integral estimates of B (b) HT integral estimates of B_1 (c) HT integral estimates of B_2

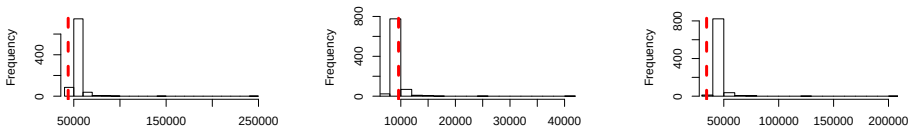

(d) HT standard estimates of B (e) HT standard estimates of B_1 (f) HT standard estimates of B_2

Figure 2.13: Mixture model mean estimates of B, B_1 and B_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation B. Truth is indicated by the dashed line.

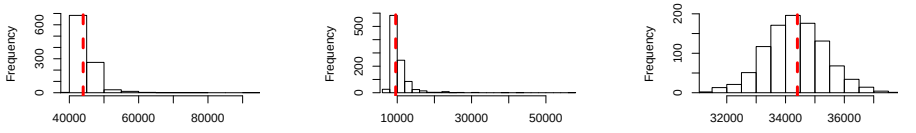

(a) HT integral estimates of B (b) HT integral estimates of B_1 (c) HT integral estimates of B_2

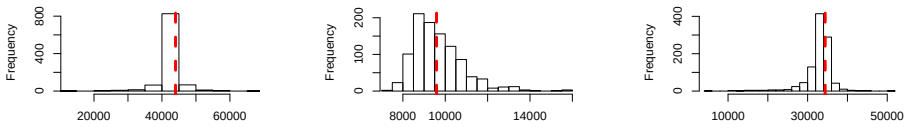

(d) HT standard estimates of B (e) HT standard estimates of B_1 (f) HT standard estimates of B_2

Figure 2.14: Full likelihood mean estimates of B, B_1 and B_2 from the HT integral (panels (a), (b) and (c)) and HT standard methods (panels (d), (e) and (f)) for simulation B. Truth is indicated by the dashed line.

2.4.1.1 Simulation Summary

The following summary can be used to aid interpretation of the results for the North Sea data; the discussion of the simulation results is found in section 2.5.1. Both the MM and FL approach recover the distribution of ingested otoliths accurately, provided the number of small otoliths with almost no probability of recovery is relatively low. The same is true for the estimates of the number of otoliths and biomass of sandeels. On the other hand, there is little hope of accurately predicting the total number of ingested otoliths or sandeel biomass if many of the otoliths are from the size range where the probability of recovery is small. While the MM approach gives more consistent but biased estimates, the FL approach is less biased but produces inconsistent results with high variability. However, by applying and comparing the results from these different approaches and estimators, it is possible to make judgments about the reliability of the estimates of the parameters of interest.

2.4.2 Data Results

2.4.2.1 Size Dependent Probability of Recovery

The FL approach for both HT estimators, as well as the MM approach with the HT standard estimator, gives higher (sometimes implausibly so) estimates of N than the MM HT integral method (Table 2.8). These results are consistent with those from the simulation study (Table 2.6). However, from the simulation results in section 2.4.1, we know that while the MM HT integral estimates are consistent they tend to be negatively biased. In contrast, both of the FL estimates as well as the MM HT standard estimates have the potential to provide more accurate predictions. We can therefore use the comparison between these approaches and estimators to assess the accuracy of the predictions for the North Sea data.

Modelling the probability of recovery as a function of otolith size led to an increase in all mean estimates of the total number of ingested otoliths (N) and almost all mean estimates of total consumed biomass (B) (Tables 2.8 and 2.9) compared with the mean estimates calculated assuming a constant recovery rate (CRR) of 0.35 (Grellier and Hammond, 2006). The largest increase in the estimated number of ingested otoliths was seen in the central North Sea region, quarter 3 (Table 2.8), and this was also where there was the largest increase in estimated consumed biomass (Table 2.9). In general, the confidence intervals for the estimates of total number of consumed otoliths suggest significant increases in estimates of N for the models fitted to the central North Sea, Moray Firth and Orkney regions (Table 2.8). However, the confidence intervals for the estimates of total biomass of consumed sandeels only suggest significant increases for the models fitted to the central North Sea and Moray Firth regions (Table 2.9). In addition, the confidence intervals for the estimates of B are less consistent between the different approaches and estimators, and this makes it difficult to identify significant increases in the estimates of B.

Table 2.8: Estimates of N, calculated using a constant recovery rate (CRR) of 0.35 compared with the HT integral and HT standard estimators for both the mixture model (MM) and full likelihood (FL) approach. The comparison is indicated as a percentage difference compared with the CRR estimate. 95% confidence intervals are given in parenthesis.

				Percentage Difference (%)	
Region	Quarter	CRR $\widehat{N}$	Method	HT Integral	HT Standard
CNS	1	6694	MM	8.22 (3.72,14.5)	8.24e+04 (22.2,5.24e+35)
			FL	16.4 (6.54,27.3)	23.8 (11.1,236)
	2	2517	MM	10.1 (4.57,19.1)	331 (13.9,1.80e+14)
			FL	39.0 (4.41,283)	70.2 (18.6,2.40e+04)
	3	3543	MM	78.8 (49.3,147)	344 (79.4,2.37e+09)
			FL	498 (57.4,2030)	289 (145,638)
	4	6849	MM	7.55 (2.96,14.3)	20.6 (7.34,3.84e+04)
			FL	35.0 (1.43,105)	17.9 (9.59,53.5)
DN	1	803	MM	3.61 (-1.74,12)	13.8 (3.49,1560)
			FL	2.86 (-1.99,9.22)	7.47 (0.87,26.5)
	2	1874	MM	3.09 (-1.55,9.55)	6.14 (5.92,46.7)
			FL	2.56 (-0.80,6.35)	3.84 (0.48,7.68)
	3	1057	MM	1.23 (-4.16,8.51)	2.37 (-3.03,15.6)
			FL	0.76 (-4.35,7.28)	1.61 (-3.88,9.18)
	4	814	MM	1.11 (-4.42,8.35)	1.47 (-3.81,8.60)
			FL	1.72 (-3.69,8.35)	1.60 (-3.56,8.11)
MF	3	1406	MM	10.7 (6.19,17.1)	18.1 (11.1,92.1)
			FL	28.4 (1.49,50.7)	19.7 (14.3,47.2)
ORK	1	15400	MM	10.5 (5.73,17.0)	7.73e+08 (71.2,3.01e+74)
			FL	18.7 (11.3,27.3)	22.1 (13.6,42.9)
	2	429	MM	5.83 (0.70,12.6)	35.0 (9.32,4.26e+04)
			FL	42.0 (0.00,599)	22.6 (10.0,163)
	3	331	MM	10.9 (4.83,20.8)	135 (19.3,1.53e+07)
			FL	43.8 (5.14,113)	53.8 (21.1,574)
	4	17040	MM	8.32 (3.62,15.0)	3.02e+04 (12.8,8.45e+32)
			FL	29.7 (12.5,50.6)	119 (23.5,3.67e+03)
SHET	1	13543	MM	1.15 (-4.58,8.39)	2.79 (-3.34,1.19e+03)
			FL	0.98 (-4.42,7.32)	0.90 (-4.22,9.45)
	3	814	MM	0.61 (-5.41,8.11)	0.74 (-5.28,8.23)
			FL	0.86 (-4.91,7.99)	0.74 (-4.91,7.74)
	4	4580	MM	1.97 (-3.82,9.59)	799 (6.29,4.24e+17)
			FL	12.7 (-7.40,73.8)	4.76 (-3.56,41.5)

Table 2.9: Estimates of B in kg, calculated using a constant recovery rate (CRR) of 0.35 and assuming the fish mass to otolith length relation of Leopold et al. (2001). The differences in the estimates of B for the HT integral and HT standard estimators for both the mixture model (MM) and full likelihood (FL) approach are indicated as a percentage difference compared with the CRR estimate. 95% confidence intervals are given in parenthesis. The MM HT standard estimates are omitted due to the somtimes exceptionally high predicted values of N

				Percentage Difference (%)	
Region	Quarter	CRR $\widehat{B}$	Method	HT Integral	HT Standard
CNS	1	35.1	MM	6.90 (2.31,14.56)	
			FL	2.04 (-2.80,8.95)	8.52 (-0.31,186)
	2	12.7	MM	6.76 (1.03,14.6)	
			FL	7.61 (1.18,1.62e+06)	31.9 (-11.9,3.42e+06)
	3	13.6	MM	34.9 (21.4,64.2)	
			FL	113 (14.5,309)	38.4 (2.79,91.3)
	4	38.7	MM	4.24 (-0.66,11.2)	
			FL	6.14 (1.03,16.4)	-7.57 (-26.5,24.3)
DN	1	6.2	MM	3.42 (-19.3,82.8)	
			FL	0.80 (-9.77,13.6)	5.06 (-7.47,27.3)
	2	15.6	MM	2.16 (-6.72,12.8)	
			FL	-0.45 (-7.05,6.20)	0.93 (-6.04,7.68)
	3	9.3	MM	0.83 (-9.05,11.9)	
			FL	0.55 (-10.1,11.9)	1.20 (-9.79,13.64)
	4	6.2	MM	1.47 (-6.71,15.0)	
			FL	1.05 (-7.41,13.3)	0.69 (-7.62,13.2)
MF	3	5.6	MM	10.5 (4.38,19.3)	
			FL	11.4 (4.32,50.8)	4.14 (-4.97,78.4)
ORK	1	90.9	MM	8.85 (3.87,15.9)	
			FL	-0.97 (-5.70,4.65)	1.70 (-4.24,16.5)
	2	2.2	MM	6.07 (-4.51,21.9)	
			FL	6.94 (-2.71,928)	-8.49 (-67.6,260)
	3	2.3	MM	3.36 (-9.49,21.9)	
			FL	5.39 (-9.67,43.0)	13.5 (-24.4,380)
	4	108.7	MM	3.94 (-1.00,10.8)	
			FL	3.51 (-1.18,9.69)	82.0 (7.90,2670)
SHET	1	110.7	MM	0.99 (-4.75,8.59)	
			FL	-0.59 (-5.61,6.03)	-0.66 (-5.76,7.84)
	3	7.0	MM	0.84 (-6.20,11.7)	
			FL	0.44 (-5.57,15.7)	0.58 (-5.76,15.0)
	4	40.3	MM	1.70 (-4.48,9.95)	
			FL	-3.64 (-8.85,9.99)	-10.7 (-31.8,13.3)

The Gompertz curve fitted to the experimental data (Figure 2.15) indicates that the probability of recovery may vary substantially over the range of data used by Grellier and Hammond (2006) to calculate the digestion coefficients. Based on this relation, which is estimated solely from the experimental data, we conclude that there is a negative bias in the overall estimated digestion coefficient (Table 2.1) of 1.6%. The digestion coefficient is calculated as the mean otolith length fed to the seals divided by the mean otolith length recovered in the scats (Grellier and Hammond, 2005). To obtain an unbiased estimate, each ingested otolith must have an equal probability of being recovered. If the probability of recovery is higher for larger otoliths, the mean length of the recovered otoliths is more representative of these larger ones, and the mean size of the otoliths recovered is therefore over-estimated, causing the digestion coefficient to be under-estimated.

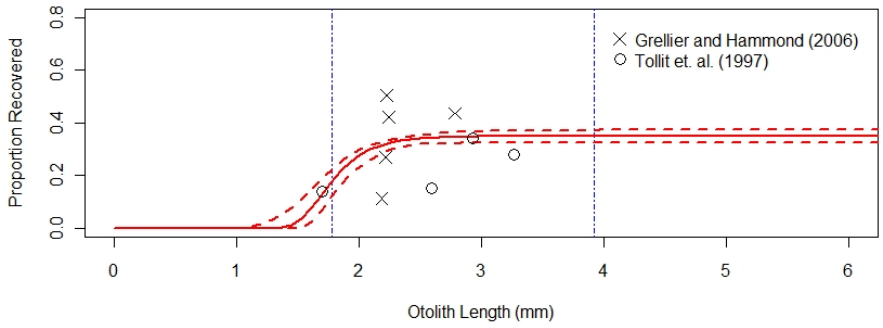

Figure 2.15: Size dependent probability of recovery in sandeel otoliths consumed by grey / harbour seals. 95% confidence intervals are represented by the dashed lines. The vertical dot-dash lines represent the size range of otoliths included in the experiments by Grellier and Hammond (2006). The data are represented by the crosses and open cirlces.

The parameter estimates for the Gompertz curve representing the size dependent probability of recovery are given in Table 2.10. While λ and δ were estimated independently in the binomial model fitted to only the experimental data, the FL approach only estimated δ and assumed a fixed relation between λ and δ. When estimated within the FL approach, we observed some variability in the mean parameter estimates for the Gompertz curve (Table 2.10). The estimates of α from the FL approach are higher on average than those from the binomial model and they are negatively correlated with δ (correlation = -0.96). In addition, the variability associated with the estimates tended to be lower for the FL models than the binomial model.

Table 2.10: Parameter estimates for the Gompertz probability of recovery curve (equation 2.6). The parameter estimates from the binomial model fitted only to the experimental data are compared with the estimates from the FL models for each location and season in the North Sea scat data. 95% confidence intervals are given in parenthesis.

Region	Quarter	α	δ	λ
Experimental Data Only		0.35 (0.32,0.37)	8.12 (4.47,11.82)	4.79 (2.74,6.87)
CNS	1	0.37 (0.35,0.39)	5.16 (4.04,6.56)	3.09 (2.45,3.90)
	2	0.35 (0.33,0.37)	7.23 (5.74,9.69)	4.28 (3.43,5.69)
	3	0.35 (0.33,0.37)	7.86 (6.41,8.82)	4.65 (3.81,5.20)
	4	0.35 (0.33,0.37)	7.76 (6.17,9.78)	4.58 (3.68,5.74)
DN	1	0.36 (0.33,0.38)	6.62 (4.64,9.45)	3.93 (2.80,5.55)
	2	0.36 (0.35,0.37)	6.63 (6.49,6.78)	3.94 (3.85,4.02)
	3	0.35 (0.33,0.38)	7.55 (5.02,11.22)	4.47 (3.01,6.57)
	4	0.35 (0.33,0.37)	8.16 (5.24,12.79)	4.81 (3.14,7.47)
MF	3	0.34 (0.32,0.36)	9.28 (6.23,14.01)	5.46 (3.71,8.17)
ORK	1	0.39 (0.37,0.41)	4.00 (3.42,4.69)	2.43 (2.10,2.82)
	2	0.35 (0.33,0.37)	7.29 (5.3,10.17)	4.31 (3.17,5.97)
	3	0.36 (0.33,0.38)	6.98 (5.18,9.35)	4.14 (3.10,5.50)
	4	0.36 (0.34,0.37)	6.67 (5.84,7.66)	3.96 (3.49,4.53)
SHET	1	0.35 (0.33,0.37)	7.16 (5.31,9.66)	4.24 (3.18,5.68)
	3	0.35 (0.33,0.37)	8.15 (5.24,12.77)	4.81 (3.14,7.46)
	4	0.37 (0.35,0.39)	5.17 (4.15,6.57)	3.10 (2.52,3.90)

2.4.2.2 *A. marinus* Otolith Sizes and their Relation with Fish Length

The *A. marinus* specimens from the Dogger Bank (Section 2.2.3) ranged from 7.0 to 21.0 cm in length; the longest otolith extracted from them was 3.53 mm whilst the widest measured 1.69 mm across. The regression between otolith length (OL) and fish length (FL) (equation 2.19) estimated a smaller slope parameter than previous studies (Fig. 2.16(a)). Although the multiple regression between otolith width (OW) and fish length included age as a factor variable, (equation 2.20), the age factor parameters (Table 2.11) suggest only small differences between age groups. The relation between otolith width and fish length showed a similar slope parameter to that in Grellier and Hammond (2006). However, both parameter values are smaller than the relation presented in Tollit et al. (1997) (Fig. 2.16(b))

Using the regression given in equation 2.19, a mean otolith length of 3.89 mm was predicted for a sandeel of length 25 cm, the largest recorded specimen of *A. marinus*. We used this as an estimate of the largest otolith size which we would expect to belong to *A. marinus*, and investigated the proportion of otoliths predicted to be greater than this value in Section 2.4.2.3.

$$OL = 0.277 + 0.145FL \tag{2.19}$$

$$OW_i = 0.270 + 0.063FL + A_i \tag{2.20}$$

where A_i is the parameter estimate associated with a fish of age $i = 1...6$, Table 2.11.

Table 2.11: Parameter estimates for the age factor variable categories A_i for fish aged $i = 0...6$, in the relation between otolith width and fish length.

Age Category	Estimate
A_0	0.000
A_1	0.000
A_2	0.019
A_3	0.039
A_4	0.041
A_5	-0.034
A_6	0.036

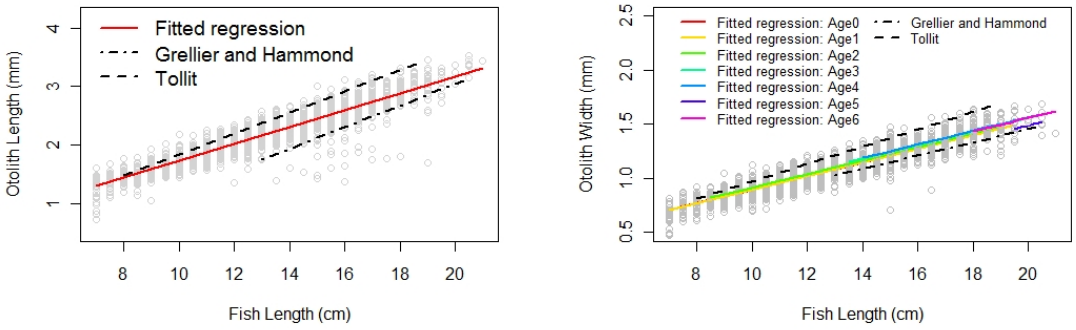

(a) Regression of otolith length and fish length compared with Grellier and Hammond (2005) and Tollit et al. (1997)

(b) Multiple regression of otolith width and fish length compared with Grellier and Hammond (2006) and Tollit et al. (1997)

Figure 2.16: A comparison of the regressions fitted to the *A. marinus* specimens from the Dogger Bank (equations 2.19 and 2.20) and previously published relations.

2.4.2.3 Otolith Lengths from the North Sea

Leopold et al. (2001) suggested that any otolith greater than 4.00 mm in length would always belong to *H. lanceolatus*. The estimated maximum undigested otolith lengths of the North Sea scat data are greater than 4.00 mm in all regions for all quarters (Table 2.12). These otolith lengths, summarised in Table 2.12, are corrected for the effects of digestion using the digestion coefficients in Table 2.1. The Donna

Nook region had the highest proportions of otoliths estimated to have been longer than 3.89 mm, especially during the first three quarters (Table 2.12). Shetland also showed fairly high proportions of large otoliths, especially in quarters 1 and 4. The lowest proportions of large otoliths were seen in the central North Sea and the Moray Firth regions. There did not appear to be any consistent seasonal variation with respect to the size ranges and the proportions of large otoliths observed.

Table 2.12: North Sea otolith lengths corrected for effects of digestion. The range of the otolith sizes and the proportion greater than 3.89 mm. The number of otoliths sampled from each of the location is given in the final column

Region	Quarter	Range	% Otoliths > 3.89mm	n
CNS	1	(1.14,5.73)	1.7%	2,343
	2	(1.26,4.64)	1.2%	881
	3	(1.31,6.07)	2.3%	1,240
	4	(1.39,6.88)	2.0%	2,397
DN	1	(1.46,5.29)	13.9%	281
	2	(1.53,7.49)	16.2%	656
	3	(1.58,6.70)	17.0%	370
	4	(1.88,5.36)	8.1%	285
MF	3	(1.54,4.72)	1.0%	492
ORK	1	(1.03,6.93)	4.6%	5,390
	2	(1.42,4.60)	4.6%	150
	3	(1.36,6.73)	1.3%	116
	4	(1.14,7.68)	5.7%	5,964
SHET	1	(1.42,6.18)	10.1%	4,740
	3	(1.99,4.63)	6.3%	285
	4	(1.23,6.63)	10.7%	1,603

2.4.2.4 Mixture Models

The MM approach provided results which were more robust to different starting values than the FL approach. For the FL approach, the starting values which estimated the larger mean for the lognormal distribution of smaller otoliths were chosen. The simulations indicated that erroneous results obtained from the FL approach were

associated with small estimates of this parameter, this implied that the distribution of smaller otoliths was located in a region with very low probabilities of recovery. Model instability was also seen in the estimates of the mixing parameter w_e, for example the FL approach 95% confidence intervals in the central North Sea, quarter 2, Orkney, quarter 2 and Moray Firth results range from ∼3% to ∼95% (Table 2.13).

The models from both the FL and the MM approach provide a satisfactory fit to the North Sea data (panel (a) of Figures 2.17-2.48). The Q-Q plots in panel (b) of Figures 2.17-2.48 indicate that, in general, the hypothesized process estimated using the FL approach provides a better explanation of the origin of the data than the MM approach. This is especially apparent in Figures 2.35, 2.36, 2.47 and 2.48, where we can see that the theoretical quantiles for the smaller otoliths are greater than those which were observed. However, in the case of the central North Sea, quarter 3 the MM approach appears to provide a better explanation of the origin of the data than the FL approach (Figs. 2.21 and 2.22).

Although these methods provide a good fit to the data and, in most instances, a plausible explanation as to the origin of the data, their ability to distinguish between the larger otoliths, thought to belong to *H. lanceolatus*, and the smaller otoliths, thought to belong mainly to *A. marinus*, is limited. The results for Orkney, quarter 2 and Shetland, for both the MM and the FL approach, show that the range of one lognormal distribution is entirely contained within the range of the other (Table 2.13, Figs. 2.37, 2.38, 2.43-2.48). This phenomenon was also observed in the MM results for Orkney, quarter 1, and the FL results for Donna Nook, quarters 3 and 4, and Moray Firth and Orkney, quarter 3 (Table 2.13, Figs. 2.35, 2.30, 2.32, 2.34 and 2.40). Furthermore, the results for Donna Nook from both approaches predict lognormal distributions all of which contain otoliths too large to be consistent with A. marinus

(Table 2.13, Figs. 2.25-2.32).

The central North Sea results, for both approaches, predict one lognormal distribution which is consistent with *A. marinus* and another distribution of larger otoliths that is more consistent with the expected sizes of *H. lanceolatus* (Table 2.13, Figs. 2.17-2.24). However, the upper limits of the larger distribution for quarters 1, 2 and 4 are not entirely inconsistent with otolith sizes for *A. marinus*. Only the upper limit of the larger distribution of otoliths for quarter 3 provides convincing evidence for the presence of *H. lanceolatus*. The MM results for Orkney data, quarters 2 and 4 as well as the FL results for Orkney, quarters 1 and 4, also predict one lognormal distribution which is consistent with *A. marinus* and another more consistent with *H. lanceolatus* (Table 2.13, Figs. 2.37, 2.41, 2.36 and 2.42). Finally, although the Moray Firth results for the MM approach predict a smaller distribution consistent with *A. marinus* and another larger distribution consistent with *H. lanceolatus*, 99-100% of the otoliths are predicted to have come from the distribution of smaller sized otoliths.

Of the central North Sea results, where both the MM and the FL approach select one lognormal distribution which is consistent with *A. marinus* and another more consistent with *H. lanceolatus*, there is consistency between each pair of parameter estimates associated with the distribution of larger otoliths (Table 2.13). In contrast, the FL parameter estimates associated with the distribution of smaller otoliths tends to be smaller than the MM estimates for each of these models. A larger weighting of biomass was also attributed to this distribution of smaller otoliths in the case of the FL approach in comparison with the MM approach (Table 2.14).

Table 2.13: Means along with 2.5 and 97.5 percentiles for each lognormal distribution fitted to the North Sea data. The mixing parameter, w_e, represents the proportion of ingested otoliths predicted to be in the first distribution; a 95% confidence interval is given in parenthesis. The latter four values are omitted in the case that the model would only converge using a single lognormal distribution.

Reg.	Q.	Method	$mean_1$ (mm)	$lower_1$ (mm)	$upper_1$ (mm)	$mean_2$ (mm)	$lower_2$ (mm)	$upper_2$ (mm)	Mixing Parameter (w_e)
CNS	1	MM	2.43	1.57	3.60	2.92	2.16	3.86	0.49 (0.30,0.67)
		FULL	2.05	1.16	3.34	2.84	2.05	3.83	0.41 (0.30,0.52)
	2	MM	1.78	1.42	2.20	2.71	1.89	3.77	0.10 (0.02,0.17)
		FULL	1.47	0.87	2.35	2.71	1.89	3.77	0.32 (0.03,0.96)
	3	MM	1.71	1.40	2.06	2.82	1.84	4.16	0.70 (0.62,0.78)
		FULL	1.37	0.98	1.87	2.85	1.87	4.16	0.92 (0.76,0.99)
	4	MM	1.90	1.49	2.40	2.87	2.07	3.87	0.13 (0.09,0.17)
		FULL	1.72	0.75	3.41	2.89	2.14	3.82	0.41 (0.24,0.58)
DN	1	MM	2.79	1.73	4.26	3.31	2.17	4.84	0.39 (0.03,0.98)
		FULL	3.06	1.83	4.81	NA	NA	NA	NA
	2	MM	2.86	1.95	4.04	4.85	3.43	6.66	0.86 (0.81,0.92)
		FULL	2.84	1.87	4.14	4.96	3.59	6.69	0.88 (0.87,0.90)
	3	MM	3.16	2.23	4.36	4.91	3.74	6.34	0.93 (0.82,0.99)
		FULL	3.03	2.32	3.90	3.52	2.10	5.54	0.49 (0.15,0.84)
	4	MM	3.14	2.29	4.20	4.94	4.33	5.62	0.99 (0.97,1.00)
		FULL	3.08	1.92	4.68	3.17	2.49	3.98	0.35 (0.04,0.76)
MF	3	MM	2.43	1.74	3.30	4.26	3.69	4.89	0.99 (0.99,1.00)
		FULL	2.08	1.18	3.40	2.47	1.84	3.25	0.51 (0.04,0.96)
ORK	1	MM	2.59	1.52	4.14	3.24	2.53	4.08	0.72 (0.66,0.77)
		FULL	2.22	1.15	3.88	3.16	2.38	4.13	0.66 (0.61,0.72)
	2	MM	2.69	2.06	3.45	2.63	1.52	4.24	0.72 (0.35,0.97)
		FULL	1.54	0.49	3.71	2.69	2.04	3.49	0.40 (0.03,0.94)
	3	MM	2.13	1.54	2.86	3.39	2.36	4.73	0.41 (0.24,0.57)
		FULL	2.20	0.87	4.64	3.42	2.94	3.96	0.78 (0.67,0.89)
	4	MM	1.94	1.47	2.51	3.04	2.12	4.21	0.18 (0.14,0.21)
		FULL	1.90	0.84	3.71	3.08	2.23	4.15	0.46 (0.41,0.50)
SHET	1	MM	2.96	1.75	4.70	3.29	2.49	4.25	0.10 (0.05,0.14)
		FULL	2.51	1.20	4.65	3.28	2.47	4.26	0.07 (0.05,0.09)
	3	MM	3.11	1.97	4.67	3.37	2.83	4.00	0.08 (0.01,0.23)
		FULL	2.97	1.76	4.71	3.37	2.81	4.01	0.06 (0.01,0.18)
	4	MM	2.95	1.41	5.48	3.41	2.76	4.15	0.12 (0.09,0.16)
		FULL	1.52	0.35	4.36	3.40	2.74	4.18	0.23 (0.03,0.48)

The mean proportion of consumed biomass predicted to be in the distribution of otolith lengths, consistent with *A. marinus* varies from 0.04 to 0.68, (Table 2.14). Assuming these do represent the proportion of *A. marinus* consumed, many of these models predict the consumed biomass of *H. lanceolatus* to be greater than that of *A. marinus*(Table 2.14). In addition, these results were calculated using the same otolith length to fish mass relation for both distributions of otoliths. Section 2.4.2.5 suggests that choosing a relation that is specific to *H. lanceolatus* would result in even smaller estimates of the proportion of biomass attributed to *A. marinus*. However, for some models the variability in these estimates is high. This is particularly the case for FL estimates for the central North Sea, quarters 2 and 3 as well as the MM estimate for Orkney, quarter 2, (Table 2.14).

Table 2.14: Percentage of biomass attributed to the distribution of smaller otoliths for the models where one distribution of otoliths was consistent with sizes of *A. marinus* and the other was more consistent with *H. lanceolatus*. Biomass was calculated assuming the fish mass to otolith length relation of Leopold et al. (2001).

		Biomass Weighting between distributions	
Region	Quarter	Mixture Model	Full Likelihood
CNS	1	0.39 (0.20,0.60)	0.25 (0.14,0.40)
	2	0.04 (0.01,0.09)	0.10 (0.01,1.00)
	3	0.38 (0.30,0.49)	0.63 (0.49,0.94)
	4	0.05 (0.03,0.08)	0.19 (0.14,0.28)
ORK	1		0.48 (0.40,0.56)
	2	0.68 (0.34,0.97)	
	4	0.06 (0.05,0.08)	0.24 (0.19,0.29)

Figures 2.17 to 2.48: (a) The fitted model plotted over a histogram of the estimated undigested lengths of the recovered otoliths. (b) Quantile-Quantile plot of the estimated undigested lengths of the recovered otolith versus points generated from the hypothesized process. (c) The mean (solid line) and a 95% confidence interval (short-dash) of the predicted distribution of ingested otoliths, the two lognormal components are shown as long-dash lines and a histogram of the estimated undigested lengths is given for reference, (d) The binomial recovery curve (solid line) with 95% confidence interval (dashed line), the crosses represent the experimental data (Tables 2.2 and 2.3); the distribution of estimated undigested lengths is also shown for reference.

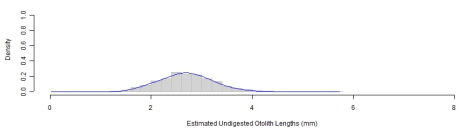

(a) Fitted model.

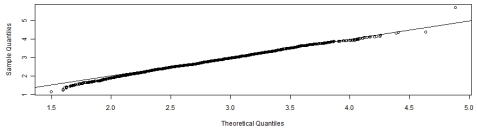

(b) Q-Q plot.

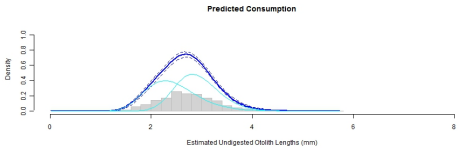

(c) Predicted distribution of ingested otoliths.

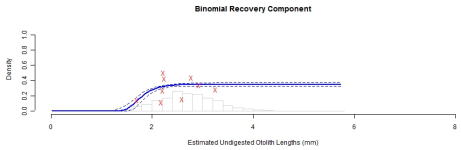

(d) Gompertz probability of recovery curve.

Figure 2.17: Mixture Model Approach: central North Sea, Quarter 1

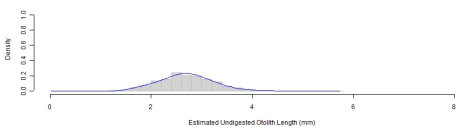

(a) Fitted model.

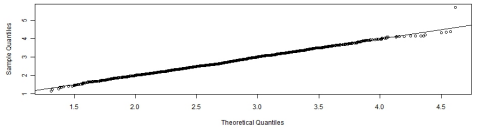

(b) Q-Q plot.

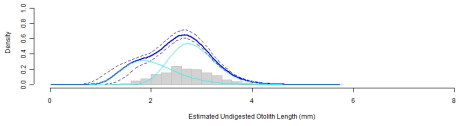

(c) Predicted distribution of ingested otoliths.

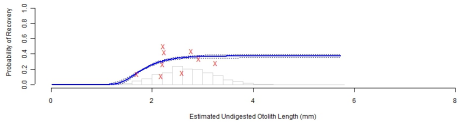

(d) Gompertz probability of recovery curve.

Figure 2.18: Full Likelihood Approach: central North Sea, Quarter 1

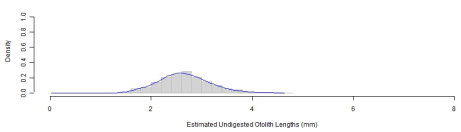

(a) Fitted model.

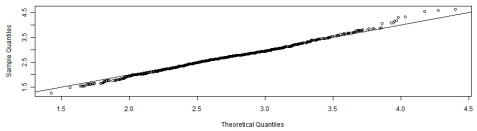

(b) Q-Q plot.

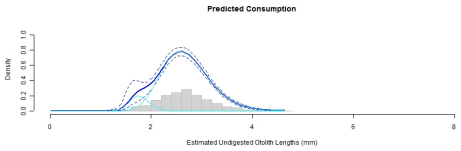

(c) Predicted distribution of ingested otoliths.

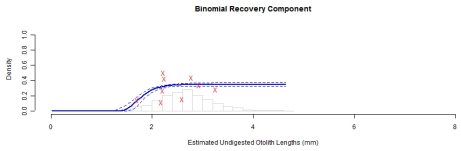

(d) Gompertz probability of recovery curve.

Figure 2.19: Mixture Model Approach: central North Sea, Quarter 2

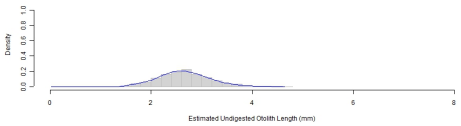

(a) Fitted model.

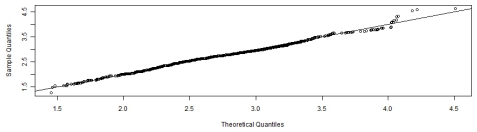

(b) Q-Q plot.

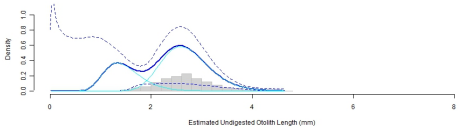

(c) Predicted distribution of ingested otoliths.

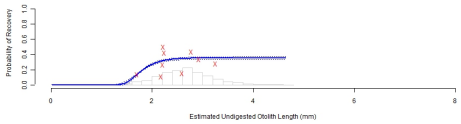

(d) Gompertz probability of recovery curve.

Figure 2.20: Full Likelihood Approach: central North Sea, Quarter 2

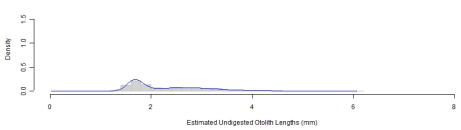

(a) Fitted model.

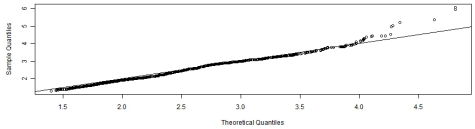

(b) Q-Q plot.

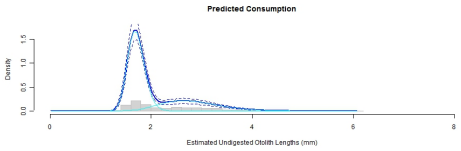

(c) Predicted distribution of ingested otoliths.

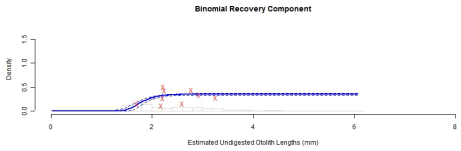

(d) Gompertz probability of recovery curve.

Figure 2.21: Mixture Model Approach: central North Sea, Quarter 3

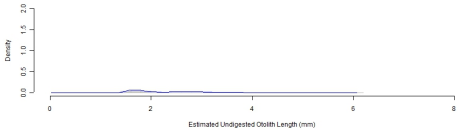

(a) Fitted model.

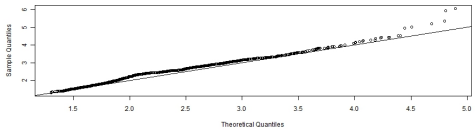

(b) Q-Q plot.

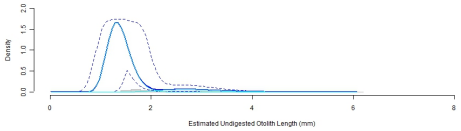

(c) Predicted distribution of ingested otoliths.

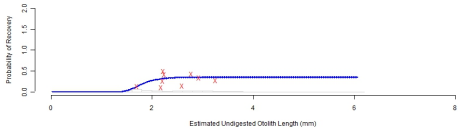

(d) Gompertz probability of recovery curve.

Figure 2.22: Full Likelihood Approach: central North Sea, Quarter 3

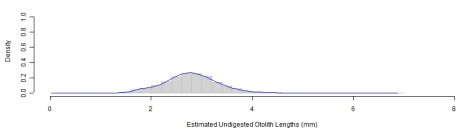

(a) Fitted model.

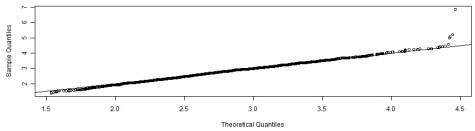

(b) Q-Q plot.

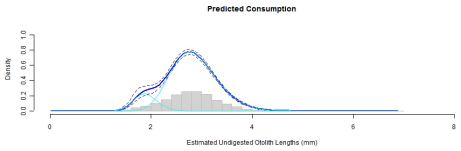

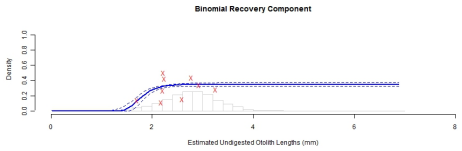

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.23: Mixture Model Approach: central North Sea, Quarter 4

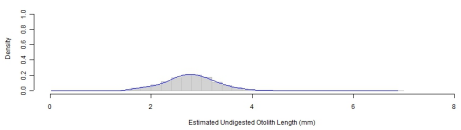

(a) Fitted model.

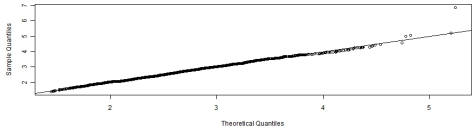

(b) Q-Q plot.

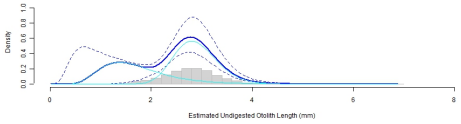

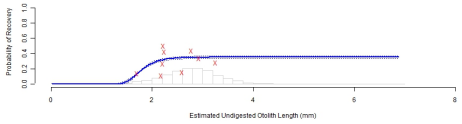

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.24: Full Likelihood Approach: central North Sea, Quarter 4

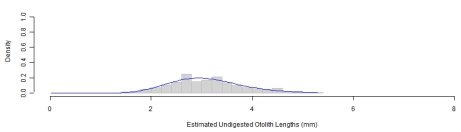

(a) Fitted model.

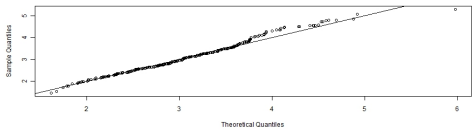

(b) Q-Q plot.

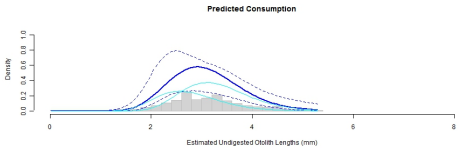

(c) Predicted distribution of ingested otoliths.

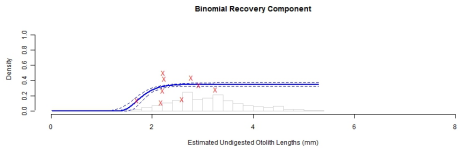

(d) Gompertz probability of recovery curve.

Figure 2.25: Mixture Model Approach: Donna Nook, quarter 1

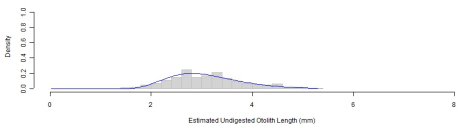

(a) Fitted model.

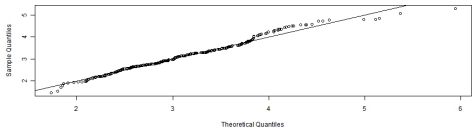

(b) Q-Q plot.

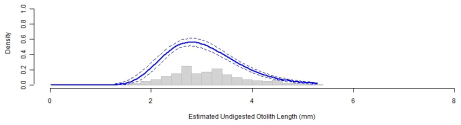

(c) Predicted distribution of ingested otoliths.

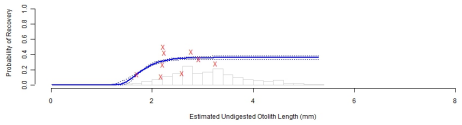

(d) Gompertz probability of recovery curve.

Figure 2.26: Full Likelihood Approach: Donna Nook, quarter 1

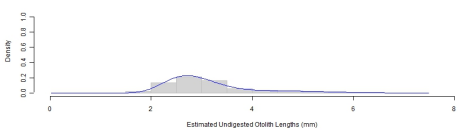

(a) Fitted model.

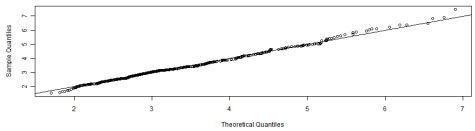

(b) Q-Q plot.

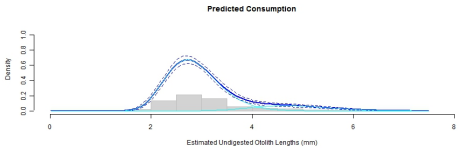

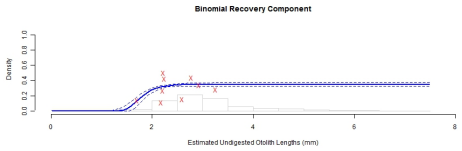

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.27: Mixture Model Approach: Donna Nook, quarter 2

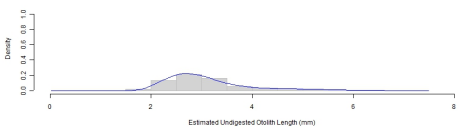

(a) Fitted model.

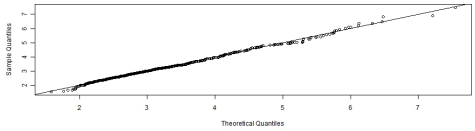

(b) Q-Q plot.

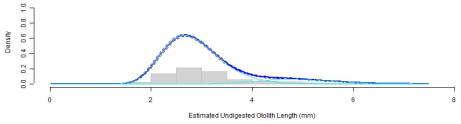

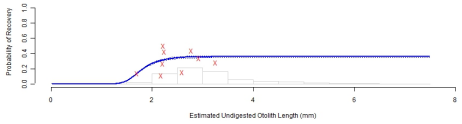

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.28: Full Likelihood Approach: Donna Nook, quarter 2

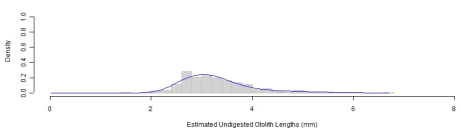

(a) Fitted model.

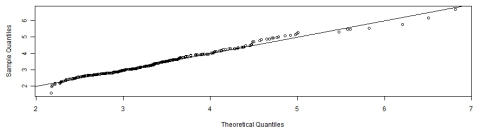

(b) Q-Q plot.

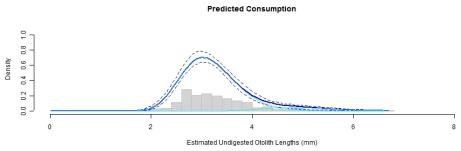

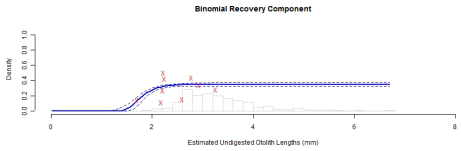

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.29: Mixture Model Approach: Donna Nook, quarter 3

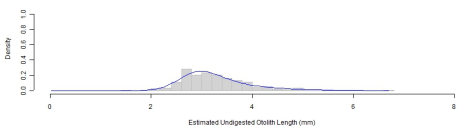

(a) Fitted model.

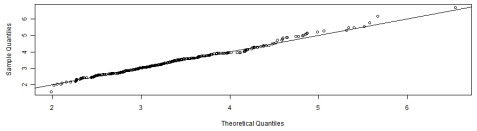

(b) Q-Q plot.

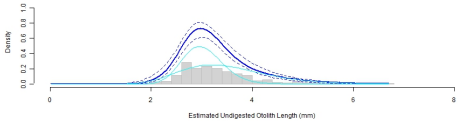

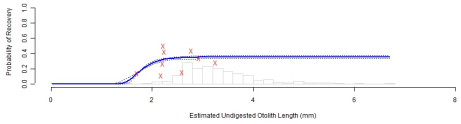

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.30: Full Likelihood Approach: Donna Nook, quarter 3

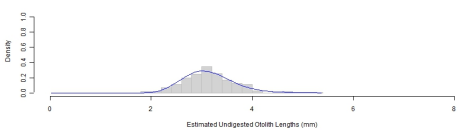

(a) Fitted model.

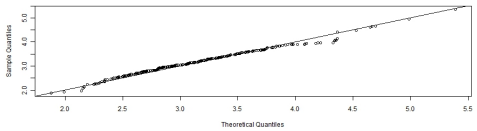

(b) Q-Q plot.

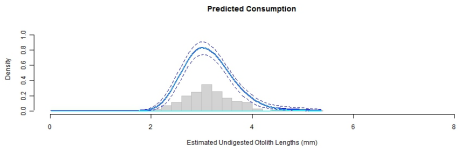

(c) Predicted distribution of ingested otoliths.

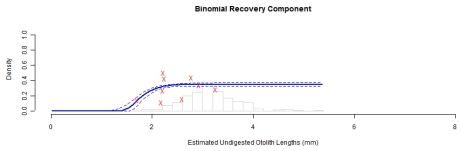

(d) Gompertz probability of recovery curve.

Figure 2.31: Mixture Model Approach: Donna Nook, quarter 4

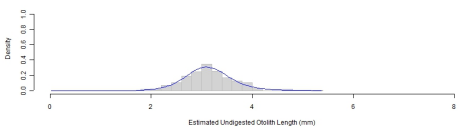

(a) Fitted model.

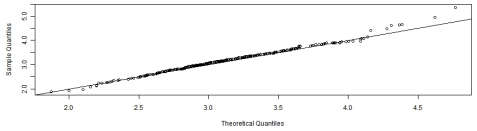

(b) Q-Q plot.

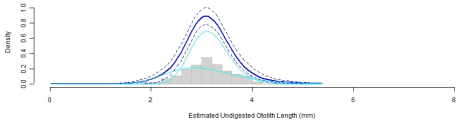

(c) Predicted distribution of ingested otoliths.

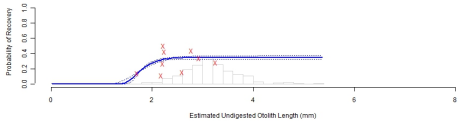

(d) Gompertz probability of recovery curve.

Figure 2.32: Full Likelihood Approach: Donna Nook, quarter 4

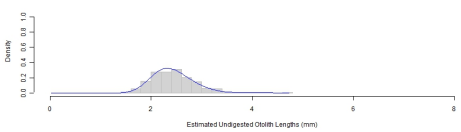

(a) Fitted model.

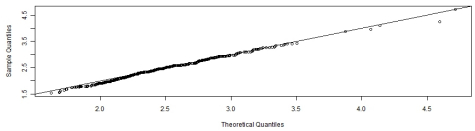

(b) Q-Q plot.

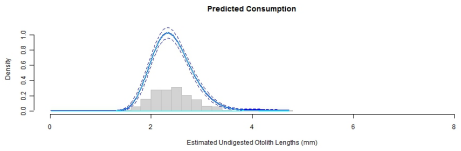

(c) Predicted distribution of ingested otoliths.

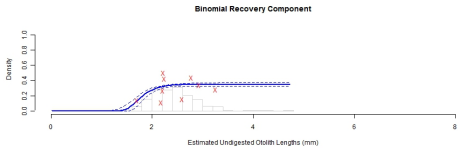

(d) Gompertz probability of recovery curve.

Figure 2.33: Mixture Model Approach: Moray Firth, quarter 3

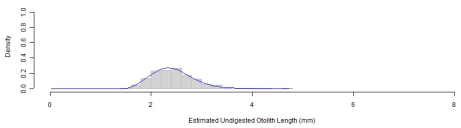

(a) Fitted model.

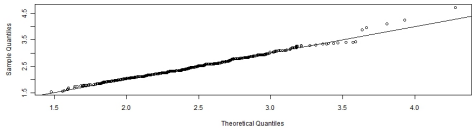

(b) Q-Q plot.

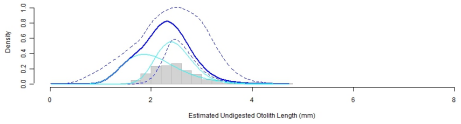

(c) Predicted distribution of ingested otoliths.

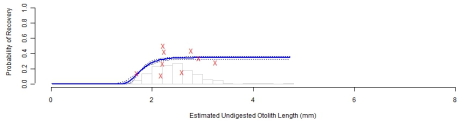

(d) Gompertz probability of recovery curve.

Figure 2.34: Full Likelihood Approach: Moray Firth, quarter 3

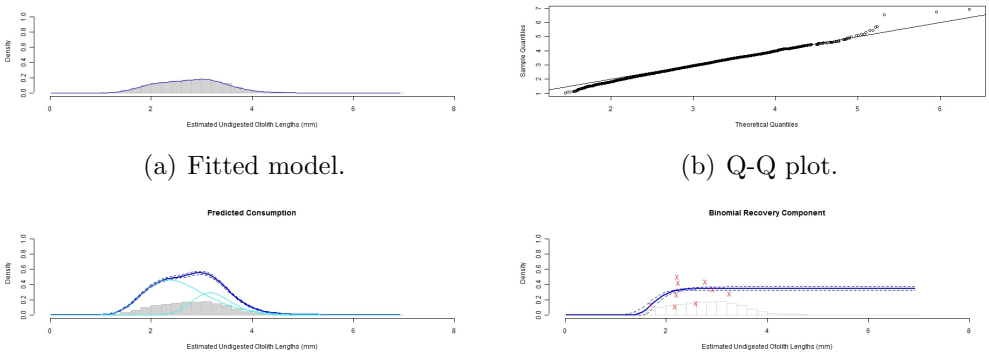

(a) Fitted model.

(b) Q-Q plot.

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.35: Mixture Model Approach: Orkney, quarter 1

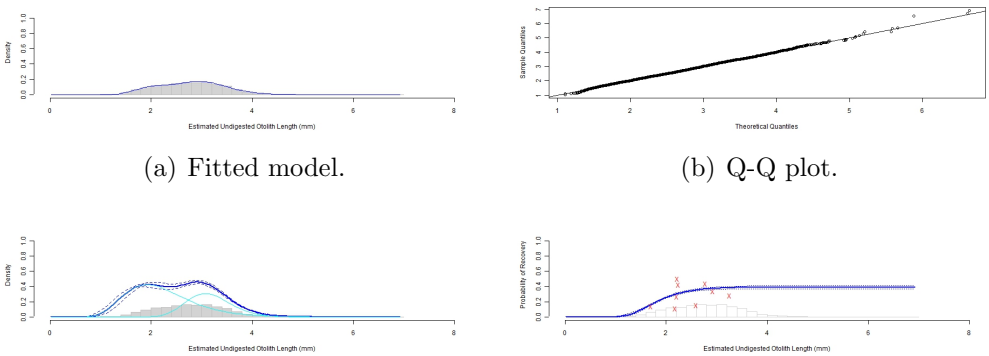

(a) Fitted model.

(b) Q-Q plot.

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.36: Full Likelihood Approach: Orkney, quarter 1

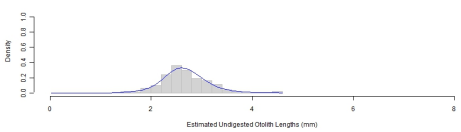

(a) Fitted model.

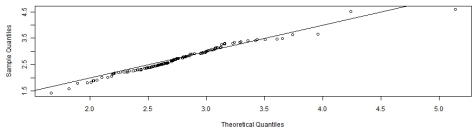

(b) Q-Q plot.

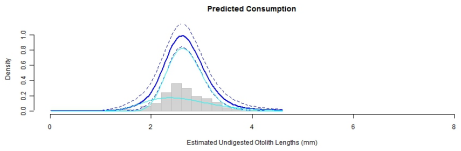

(c) Predicted distribution of ingested otoliths.

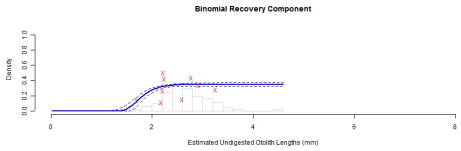

(d) Gompertz probability of recovery curve.

Figure 2.37: Mixture Model Approach: Orkney, quarter 2

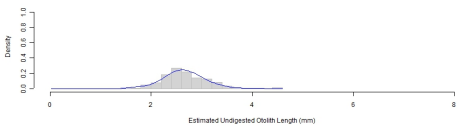

(a) Fitted model.

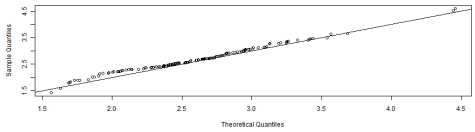

(b) Q-Q plot.

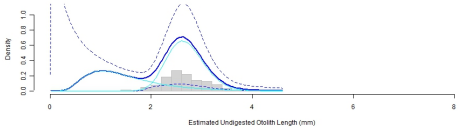

(c) Predicted distribution of ingested otoliths.

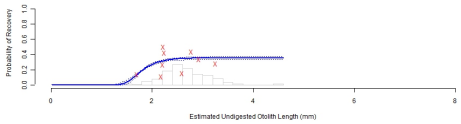

(d) Gompertz probability of recovery curve.

Figure 2.38: Full Likelihood Approach: Orkney, quarter 2

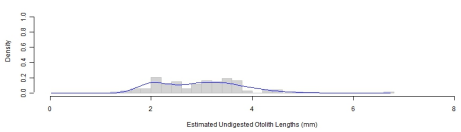

(a) Fitted model.

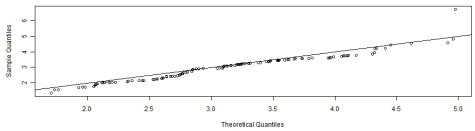

(b) Q-Q plot.

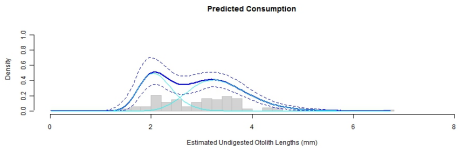

(c) Predicted distribution of ingested otoliths.

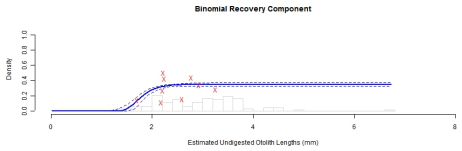

(d) Gompertz probability of recovery curve.

Figure 2.39: Mixture Model Approach: Orkney, quarter 3

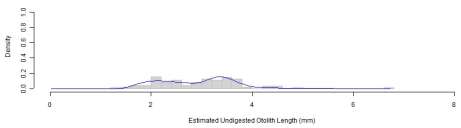

(a) Fitted model.

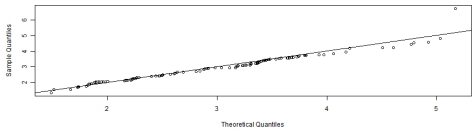

(b) Q-Q plot.

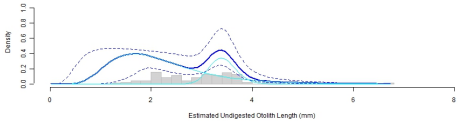

(c) Predicted distribution of ingested otoliths.

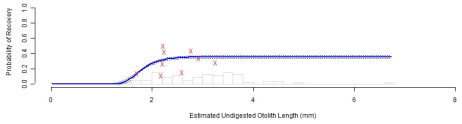

(d) Gompertz probability of recovery curve.

Figure 2.40: Full Likelihood Approach: Orkney, quarter 3

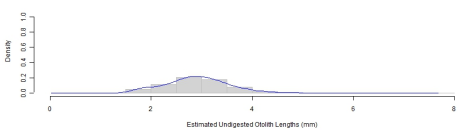

(a) Fitted model.

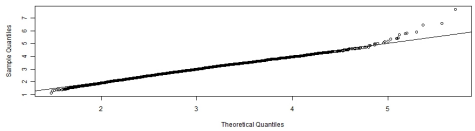

(b) Q-Q plot.

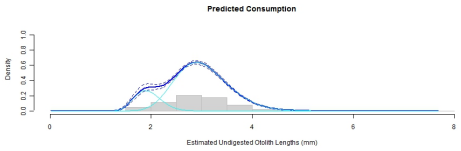

(c) Predicted distribution of ingested otoliths.

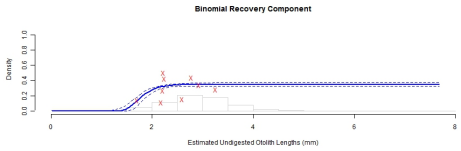

(d) Gompertz probability of recovery curve.

Figure 2.41: Mixture Model Approach: Orkney, quarter 4

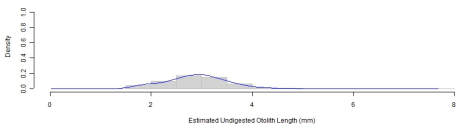

(a) Fitted model.

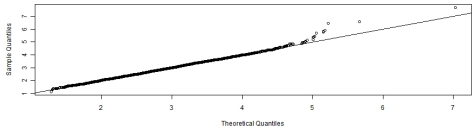

(b) Q-Q plot.

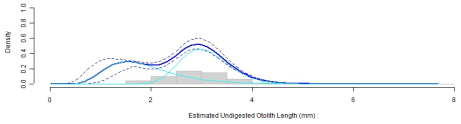

(c) Predicted distribution of ingested otoliths.

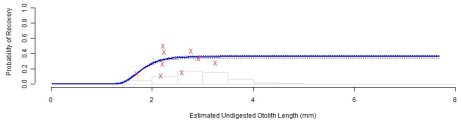

(d) Gompertz probability of recovery curve.

Figure 2.42: Full Likelihood Approach: Orkney, quarter 4

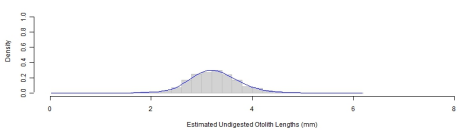

(a) Fitted model.

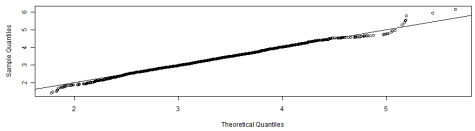

(b) Q-Q plot.

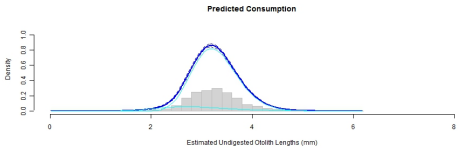

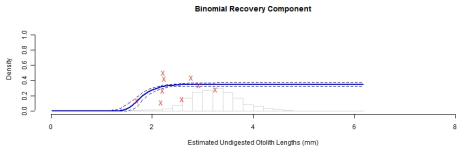

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.43: Mixture Model Approach: Shetland, quarter 1

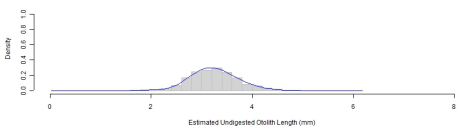

(a) Fitted model.

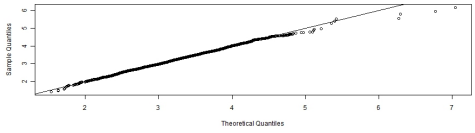

(b) Q-Q plot.

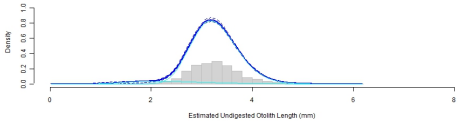

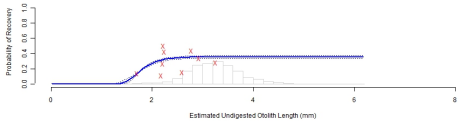

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.44: Full Likelihood Approach: Shetland, quarter 1

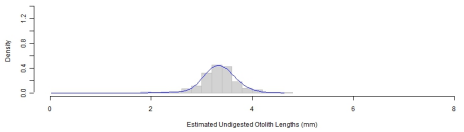

(a) Fitted model.

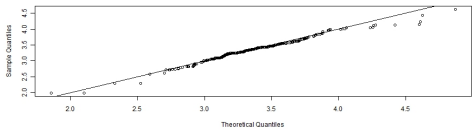

(b) Q-Q plot.

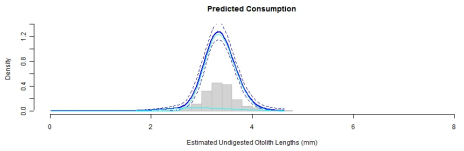

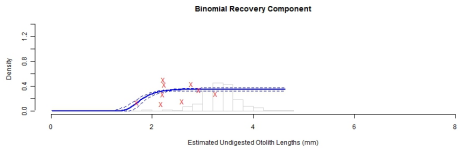

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.45: Mixture Model Approach: Shetland, quarter 2

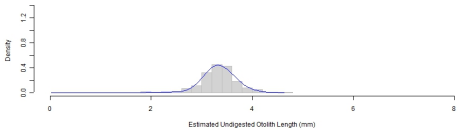

(a) Fitted model.

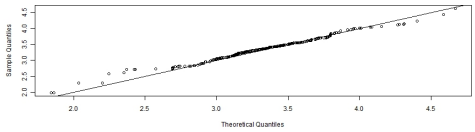

(b) Q-Q plot.

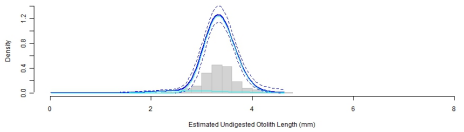

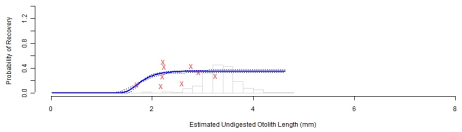

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.46: Full Likelihood Approach: Shetland, quarter 2

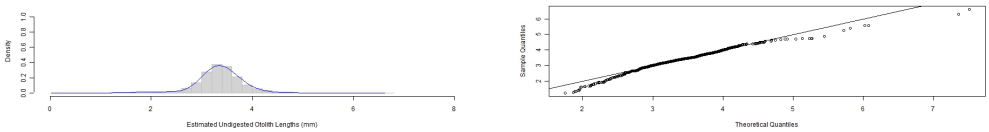

(a) Fitted model.

(b) Q-Q plot.

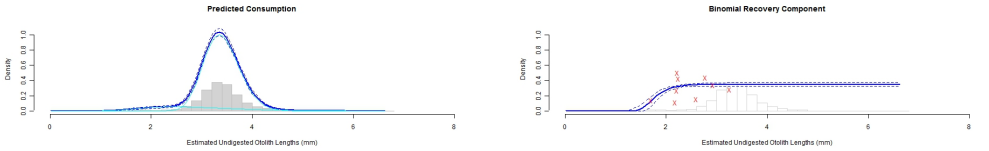

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.47: Mixture Model Approach: Shetland, quarter 4

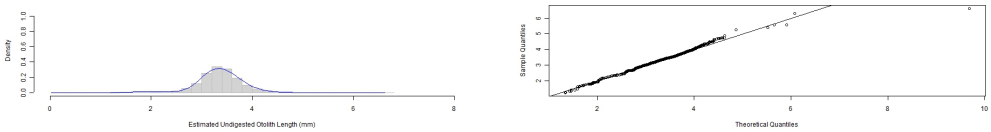

(a) Fitted model.

(b) Q-Q plot.

(c) Predicted distribution of ingested otoliths.

(d) Gompertz probability of recovery curve.

Figure 2.48: Full Likelihood Approach: Shetland, quarter 4

2.4.2.5 Variability in Relations between Fish Mass, Fish Length and Otolith Length for *A. marinus* and *H. lanceolatus*

The two relations between otolith length and fish length for *H. lanceolatus* are similar to those observed for *A. marinus*. However, the variability in the relations within species is large (Fig. 2.49(a)). The mis-identification of an otolith as being from *A. marinus*, when in fact it is *H. lanceolatus*, would probably have no greater impact on estimating fish length than choosing an incorrect relation from the selection available for *A. marinus*.

A large amount of variability was also found in the relations between fish mass and otolith length, both within and between species (Fig. 2.49(b)). On average, these relations predict a greater mass for *H. lanceolatus* than *A. marinus* for otoliths longer than around 2.6 mm (Fig. 2.49(b)). The misidentification of a *H. lanceolatus* otolith as one from *A. marinus* would most likely lead to an under-estimation of consumed biomass. For example, Hammond and Grellier (2006) analysed the otoliths in the North Sea scat data as if they were all *A. marinus* using only Leopold et al. (2001)'s otolith length to fish mass relation. If we assume that all otoliths in the North Sea data greater than 3.89 mm in length were in fact from *H. lanceolatus*, and use the species-specific relations of Leopold et al. (2001), we get a 5.9% increase in estimated consumed biomass. However, if we assume that all otoliths are *A. marinus* and choose the relations used in Lewis et al. (2003) rather than that of Leopold et al. (2001), the increase in estimated consumed biomass is 14.9%.

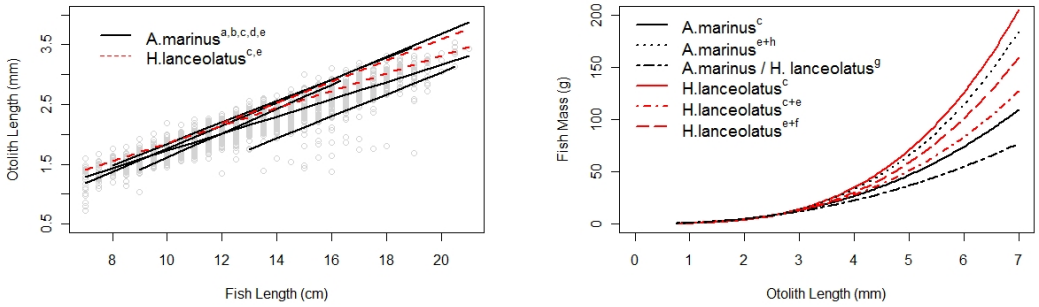

(a) Linear relation between otolith length and fish length.

(b) Power relation between fish mass and otolith length.

Figure 2.49: Comparison of otolith length to fish length and fish mass to otolith length from previously published sources. [a] As presented in section 2.4.2.2, [b] Grellier and Hammond (2005), [c] Leopold et al. (2001), [d] Tollit et al. (1997), [e] Härkönen (1986) in Lewis et al. (2003), [f] Macer (1966), [g] Pierce et al. (2007), [h] Harris and Hislop (1978) in Lewis et al. (2003)

2.5 Discussion

Understanding the predatory relationship between grey seals and sandeels is an important step towards sustainable management of the North Sea ecosystem; such information is essential in the application of effective multispecies modelling. While *A. marinus* is a commercially important species, sandeels are also an important component of the diet of grey seals as well as a number of internationally important seabird colonies. This study provides compelling evidence that a significant proportion of the sandeels consumed by the grey seals are in fact *H. lanceolatus*, which is not commercially exploited, rather than the commercially important *A. marinus*. This result will have significant impacts on the multispecies model of the North Sea and therefore the management of the sandeel fisheries in this area. In addition, it gives quantitative insight into potential sources of bias when estimating the total biomass of sandeels consumed by grey seals using otoliths retrieved from scats. Specifically, there

is bias associated with the chosen forms of the relations between fish mass and otolith length and the effects of otolith size on the probability of recovery. Such sources of bias, if not accounted for, may also lead to mis-management of this precious resource.

2.5.1 Simulation

This chapter not only proposed and implemented a novel method for assessing the quantity of sandeels consumed by grey seals in the North Sea using a size dependent probability of recovery but it also tested the robustness of this method through the use of simulations.

The ingested otolith sizes used in the simulations were drawn from a mixture of lognormal distributions and the probability of recovery was chosen from a Gompertz curve, as assumed in the model. However, it proved difficult to recover the original parameters and therefore the original distribution of otoliths when the probability of recovering smaller otoliths was low. The greater the number of small otoliths ingested, the more of the lower end of the distribution is lost to digestion and the greater the bias when trying to estimate this part of the distribution of otolith sizes. As more of the smallest otoliths are lost to digestion, the MM approach over-estimates the mean of this lower distribution and it appears narrower than it actually is, resulting in under-estimation of the standard deviation. In contrast, as more of the smallest otoliths are lost to digestion the FL approach under-estimates the mean of the lower distribution and the distribution therefore appears wider, resulting in over-estimation of the standard deviation. In addition, the FL approach suffers from high levels of variability and skewed distributions of parameter estimates, which indicate that it is likely to give inconsistent results.

The estimation of the number of otoliths, as well as the biomass of sandeels, ingested is also problematic and less reliable for small otoliths. As more small otoliths are lost,

the HT integral estimator used with the MM approach consistently under-estimates the total number of otoliths ingested whereas the HT standard method can often get closer to the true value (Table 2.6). However, the problem with the MM HT standard estimator is apparent when we consider the spread of the estimates, these are skewed and therefore sometimes are vast over-estimates. The MM HT standard estimator relies solely on the form of the Gompertz curve estimated from the feeding experiment data. Therefore, as otolith size gets smaller and this function approaches zero, any otoliths recovered in this lower tail have extreme implications for the estimates of N and B. This method is, therefore, especially sensitive to mis-specification of the parameters of the recovery function when the otoliths ingested are small. The HT integral estimates are based on the whole process, and are therefore less influenced by the rare recoveries of small otoliths. Because of this, the HT integral method is preferred for estimating the total number of otoliths ingested for the MM approach.

The FL approach re-estimates the Gompertz recovery curve for each set of recovered otoliths, and can therefore adjust its shape if, by chance, small otoliths are recovered, or if the parameters of the recovery curve differ across seasons or locations. However, the simultaneous estimation of the Gompertz curve and the mixture model can lead to model instability, indicated by the skewed distributions of estimates for N and B. In contrast to the MM approach, the FL approach can choose very low estimates of μ_1, moving the distribution of smaller otoliths into the region where very few are recovered. This can happen because the FL approach is looking for a combination of a mixture model and recovery curve which explain the distribution of recovered otoliths whereas the MM approach first fits the lognormal mixture to the recovered otoliths and later corrects the mixing parameter, w_r in order to obtain w_e, based on the recovery curve.

2.5.2 Presence of *H. lanceolatus* in Grey Seal Diet

The central North Sea results from the MM approach provided the most plausible predictions for distinguishing between otoliths from *A. marinus* and *H. lanceolatus*. The MM results for the central North Sea region, quarters 2, 3 and 4 all suggested a distribution of smaller otoliths ranging from around 1.8 to around 2.2 mm in length, which are potentially from *A. marinus*, and a distribution of larger otoliths ranging from around 1.9 to around 4.0mm in length, which are potentially from *H. lanceolatus*. This consistency across seasons, as well as the clear bimodal distribution observed in quarter 3 suggests that the models are distinguishing between these two species. Assuming these models are accurate then they suggest that only a small proportion of the biomass of sandeels consumed by grey seals in this region is *A. marinus*. For quarters 2 and 4 this is predicted to be around 4-5% and for quarter 3 around 38%. However, based on the simulation results these figures are likely to have some negative bias. In addition, the central North Sea region was one of the regions with the lowest percentage of large otoliths (Table 2.12). It is concerning that for the Donna Nook and Shetland areas, which had the largest proportion of otoliths greater than 3.89 mm in length (Table 2.12), the lognormal distributions fitted to the data did not differentiate between *A. marinus* and *H. lanceolatus*.

Previous studies have indicated that *A. marinus* is by far the most abundant species of sandeel in the North Sea (Macer, 1966; Nævdal and Thorkildsen, 2002). In addition, other studies investigating the diet of anglerfish (Laurenson and Priede, 2005) and Northern Gannets (Lewis et al., 2003; Hamer et al., 2007) suggest that *A. marinus* is the predominant sandeel in the diet of these species. However, Pearson (1968) found that although *H. lanceolatus* was scarce in the diet of terns and gulls, they were more frequently found in the diet of auks.

Table 2.12 indicated that many of the otoliths present in the grey seal scats are too large to belong to *A. marinus* and are more likely to belong to *H. lanceolatus*. However, the true proportion of *H. lanceolatus* in the grey seal diet is difficult to estimate. In addition, it must be remembered that there are three other species of sandeel, about which very little is know, in the North Sea (Macer, 1966). Other than the fact that otolith longer than 3.89mm are almost certain to have come from *H. lanceolatus*, the overlap of the otolith sizes of the different species of sandeel make it extremely difficult, if not impossible, to identify otoliths to species based on length alone.

2.5.3 Fish Mass to Otolith Length Relations

This study found that the largest potential source of bias in the estimation of the biomass of sandeels consumed is introduced by the choice of relation used to predict fish mass based on otolith length. Pierce et al. (2007) found that sandeels increased in importance in the diet of harbour porpoises when they used an alternative relation between fish mass and otolith length. They suggest that the chosen relation should be appropriate to the fish in the specific region of interest. This may be especially relevant in the case of *A. marinus*, as recent studies have indicated that the North Sea sandeel stock is likely comprised of several discrete populations (Pedersen et al., 1999; Greenstreet et al., 2006). In addition, more general studies suggest that variation in these relations may not only be regional but also seasonal or annual, depending on environmental conditions, and this can lead to differences in predicted fish mass of 200% or more (Froese, 2006).

Although the presence of *H. lanceolatus* in the diet of grey seals has been previously documented (Prime and Hammond, 1990; Hammond et al., 1994a,b), the effects of using a potentially incorrect fish mass to otolith length relation has not, to our

knowledge, been investigated. Although, it was not possible to confidently quantify the numbers of otoliths which were likely to have come from *H. lanceolatus*, it is clear that identifying which otoliths belong to this species is only half the problem. There is also a requirement for an accurate relation between otolith size and fish mass, and this may need to be specific to region, season and year. Based on the relations presented in Figure 2.49(b) it appears more likely that the identification and specific treatment of *H. lanceolatus* will result in increased estimates of the biomass of sandeels consumed by grey seals. The largest potential biases will most likely be seen in the Donna Nook and Shetland regions, where the large proportions of otoliths greater than 3.89 mm suggest that *H. lanceolatus* is an important component of the seals' diet. In addition, the estimates of undigested otolith length are based on digestion coefficients calculated for *A. marinus*, but the relation for *H. lanceolatus* may be different.

2.5.4 Size Dependent Probability of Recovery

The fact that prey species with larger otoliths tend to have higher otolith recovery rates has led to the implementation of species specific number correction factors (NCF) (Tollit et al., 1997; Bowen, 2000; Grellier and Hammond, 2006). However, despite observations of similar trends within species (Tollit et al., 1997; Bowen, 2000), no one, as far as we are aware, has yet implemented a size dependent NCF function. We incorporated this function by using a Gompertz curve to describe the probability of recovering an otolith, given that it was eaten and of a certain size. Based on this study, we predict that previous estimates of the biomass of sandeels consumed by grey seals are likely to be under-estimates, due to the omission of this size dependent probability of recovery.

The function used here to estimate size-dependent recovery can affect biomass estimates both directly, due to changes in the predicted number of ingested otoliths, and indirectly, through its effects on the digestion coefficient. The most striking increase in estimated biomass due to direct effects ($\sim$36%, based on the MM HT integral and FL HT standard estimates) was seen in the central North Sea model, quarter 3. This estimated increase has implications when assessing seasonal and regional variations in diet. The indirect effects stem from an estimated increase of 1.6% in the digestion coefficient; by itself this would lead to a 4.1% increase in biomass estimates. However, an increase in the digestion coefficient affects all the otolith lengths used to fit the model and, therefore, has other implications. Firstly, it may affect the number of otoliths whose lengths are consistent with *H. lanceolatus*. Secondly, the location of the estimated undigested otolith lengths in relation to the point of inflection in the recovery function will change, which will directly influence the biomass estimates. Finally, the Gompertz curve parameter estimates for the FL approach are based on both the experimental data and the estimated undigested otolith lengths from the North Sea data. As the shape parameter estimates from the FL approach δ (and therefore λ) are significantly correlated with the estimated undigested size of the smallest otolith recovered from the scats (correlation 0.69, p-value 0.003), these are likely to be affected.

Tollit et al. (1997) found that the degree of digestion (and therefore the digestion coefficient) for sandeels was positively related to otolith size. Although we have not explicitly investigated the effects of this, we can use our results to predict its implications. As the digestion coefficients used in this study are an average over most otolith sizes, we would expect size-specific digestion corrections to lead to an increase in the estimated undigested lengths of the largest otoliths. This would lead to increased biomass estimates and also the suggestion of higher levels of *H. lanceolatus*

in the diet. In addition, the estimated undigested lengths of the smaller otoliths would likely become even smaller and this would influence the parameter estimates of the recovery function for the FL approach.

2.5.5 Reliability of Results

The validity of the experimental data is questionable as it was partially based on feeding trials with a different seal species, the harbour seal, while the North Sea scat data is specific to grey seals. Although Bowen (2000) found no significant difference in recovery rates between harbour seals and grey seals fed Atlantic cod (*Gadus morhua*) and Atlantic and Pacific herring (*Clupea harengus/pallasii*), the same may not be true in the case of sandeels. The probability of recovery is highly variable for Atlantic and Pacific herring (Bowen, 2000; Grellier and Hammond, 2006) making differences hard to detect. The recovery rates for Atlantic cod are generally higher and more consistent but their increased robustness compounded with the high variability in experimental conditions, including, seal age, meal sizes and feeding method (Bowen, 2000) would also make differences difficult to detect.

The upper asymptote of the size dependent probability of recovery function is estimated at between 0.34 and 0.39, which is very similar to the constant recovery rate of 0.35 reported in Grellier and Hammond (2006). However, due to the high uncertainty in the estimates of the shape parameters λ and δ, we conclude that the quantity and quality of experimental data is insufficient to evaluate whether or not the Gompertz curve is an adequate function for modelling the recovering rate. This high level of uncertainty surrounding these parameter estimates, along with the FL model instability, led to the decision not to include any additional uncertainty from the estimated proportions of the feeding experiment data. It is suggested that if better data on the otolith recovery rates become available then these uncertainties

may be incorporated using a bootstrap which resamples the proportion estimates for each set of trials. Doing this would result in greater uncertainty in the estimation of the Gompertz curve parameters which would lead to greater uncertainty in the estimates of consumed biomass of sandeels. In addition, we should be cautious of the lower uncertainty in the parameter estimates from the FL approach compared with the binomial model. This is most likely an artifact of the fixed relation between λ and δ specified in the FL approach, in contrast the binomial model in the MM approach allowed these two parameters to be estimated independently of one another. In addition, a more appropriate likelihood for combining these recovery rate data may be an overdispersed binomial distribution due to the variability between the trials which were carried out on different seals. It was not possible to test for overdispersion in the feeding experiment data due to time constraints but again this could be an additional source of variability which has not been accounted for.

The reliability of the FL and MM approaches for predicting the number and biomass of consumed sandeels is difficult to assess. Based on the simulations, when the four estimates of N and B (given by the HT standard and HT integral estimators from the FL and MM approaches) agree with each other, we can have more confidence in the estimates. However, this is not often the case and when it is, it only leads to small, non-significant increases in N and B, for example Shetland, quarter 1, Table 2.9.

The prediction of 3.89 mm as the largest expected mean otolith length for *A. marinus*, for a fish of 25 cm, was based on data from fish with lengths between 7 and 21 cm. This prediction is, therefore, an extrapolation outside the range of the data. To our knowledge, there are no published relations or observations of otolith lengths for fish greater than 21 cm. As we have already observed, these relations are highly

variable and all other relations presented in this study would lead to predictions greater than 3.89 mm. However, our value is consistent with the maximum otolith length of 4.00 mm given in Leopold et al. (2001). Because 3.89 mm is an estimate of the mean it implies that there will be some variability around this value, so the observation of slightly larger otoliths is not surprising.

2.6 Conclusions

The simulations showed that provided the number of small otoliths with almost no probability of recovery is relatively low the estimates of the number of otoliths and biomass of sandeels consumed can be predicted accurately. However, if there are many otoliths ingested which are in the size ranges of those with little probability of recovery then we cannot hope to be able to accurately predict the quantity consumed.

Although we found that the largest source of bias in the estimation of the biomass of sandeels consumed by grey seals is associated with the choice of relation used to predict fish mass from otolith length, we did not consider all sources of bias. Hammond and Grellier (2006) suggested that the largest potential source of bias is the estimation of seal population size. Choosing a different model to predict seal population size led to an increase of 40% in the estimated consumed biomass in the North Sea.

To improve the estimation of the total biomass of sandeels consumed by grey seals, effort is probably better spent obtaining accurate fish mass to otolith length relations for *A. marinus* than identifying *H. lanceolatus* otoliths. These relations should, as far as possible be specific to the sandeels in the study, and conform to the recommendations in Froese (2006). An alternative approach, given the scale of this data, would be to incorporate more variability into the model by using a range of fish mass/otolith length relations.

Estimation of the probability of recovery function would greatly benefit from improved experimental data. We therefore recommend further experiments, specific to grey seals, which would assess the effects of otolith size on both the probability of

recovery and the degree of digestion. This would allow more accurate and precise estimation of the digestion coefficient as well as improving the estimates of the biomass of sandeels consumed by the grey seal.

There may be potential to try to identify sandeel otoliths recovered from seals scats using multivariate shape analysis techniques (Stransky et al., 2008). However, Leopold et al. (2001) indicate that the otolith shapes for these two species are very similar. Nævdal and Thorkildsen (2002) have identified genetic markers which successfully differentiate between different species of sandeel. DNA testing of faecal samples may therefore provide a promising alternative to hard part analysis, when trying to quantify the importance of these species in the grey seals diet (Deagle and Tollit, 2007; Deagle et al., 2009).

There is a lack of up-to-date information in the literature on the abundance and distribution of *H. lanceolatus* and other species of sandeel in the North Sea. Such information is vital in assessing the prevalence of *H. lanceolatus* in grey seal diet and whether or not more than a negligible amount of other sandeel species, which were not considered in this chapter, are also being consumed. Given the number of species for which *H. lanceolatus* is a prey item, this information would also be valuable for understanding the North Sea ecosystem and ensuring sustainable management of sandeel stocks.

Bibliography

W. Bowen. Reconstruction of pinniped diets: accounting for complete digestion of otoliths and cephalopod beaks. *Can. J. Fish. Aquat. Sci.*, 57:898–905, 2000.

K. Brander. Cod recruitment is strongly affected by climate when stock biomas is low. *Journal of Marine Science*, 62:339–343, 2005.

P. Bromley, T. Watson, and J. Hislop. Diel feeding partterns and the development of food webs in pelagic 0-group cod (*Gadus morhua* L.), haddock (*Melanogrammus aeglefinus* L.), whiting (*Merlangius merlangus* L.), saithe (*Pollachius virens* L.), and Norway pout (*Trisopterus esmarkii* Nilsson) in the northern North Sea. *ICES Journal of Marine Science*, 54:846–853, 1997.

R. Bryrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM J. Scientific Computing*, 16:1190–1208, 1995.

S. Buckland, D. Anderson, K. Burnham, J. Laake, D. Borchers, and L. Thomas. *Advanced Distance Sampling: Estimating abundance of biological populations*. Oxford University Press, 2004.

CEFAS. Multispecies fisheries management: A comprehensive impact assessment of the sandeel fishery along the English east coast. Annex 1 of final report to DEFRA on project MF0323, 83pp. Technical report, CEFAS, 2007.

B. Deagle and D. Tollit. Quantitative analysis of prey DNA in pinniped faeces: potential to estimate diet composition? *Conservation Genetics*, 8:743–747, 2007.

B. Deagle, R. Kirkwood, and S. Jarman. Analysis of Australian fur seal diet by pyrosequencing prey DNA in faeces. *Molecular Ecology*, 18:2022–2038, 2009.

G. Engelhard, J. van der Kooij, E. Bell, J. Pinnegar, J. Blanchard, D. Mackinson, and D. Righton. Fishing mortality versus natural predation on diurnally migrating sandeels *Ammodytes marinus*. *Marine Ecology Progress Series*, 369:213–227, 2008.

D. Fournier, J. Hampton, and J. Sibert. MULTIFAN-CL: a length based, age-structured model for fisheries stock assessment, with application to south Pacific albacore, *Thunnus alalunga*. *Can. J. Fish. Aquat. Sci.*, 55:2105–2116, 1998.

M. Frederiksen, M. Edwards, A. Richardson, N. Halliday, and S. Wanless. From plankton to top predators: bottom-up control of a marine food web across four trophic levels. *Journal of Animal Ecology*, 75:1259–1268, 2006.

R. Froese. Cube law, condition factor and weight-length relationships: history, meta-analysis and recommendations. *Journal of Applied Ichthyology*, 22:241–253, 2006.

R. Froese and D. E. Pauly. Fishbase. World Wide Web electronic publication. www.fishbase.org, version (09/2009) [acceded 29 october 2009]. 2009.

R. Furness. Management implications of interactions between fisheries and sandeel-dependent seabirds and seals in the North Sea. *ICES Journal of Marine Science*, 59:261–269, 2002.

S. Greenstreet, J. McMillan, and E. Armstrong. Seasonal variation in the importance of pelagic fish in the diet of piscivorous fish in the Moray Firth, NE Scotland: a response to variation in prey abundance? *ICES Journal of Marine Science*, 55: 121–133, 1998.

S. Greenstreet, E. Armstrong, H. Mosegaard, H. Jensen, I. Gibb, H. Fraser, B. Scott, G. Holland, and J. Sharples. Variation in the abundance of sandeels *Ammodytes marinus* off southeast Scotland: an evaluation of area-closure fisheries management and stock abundance assessment methods. *Journal of Marine Science*, 63:1530–1550, 2006.

K. Grellier and P. Hammond. Feeding method affects otolith digestion in captive gray seals: Implications for diet composition estimation. *Marine Mammal Science*, 21 (2):296–306, 2005.

K. Grellier and P. Hammond. Robust digestion and passage rate estimates for hard parts of grey seal (*Halichoerus grypus*) prey. *Can. J. Fish. Aquat. Sci.*, 63:1982–1998, 2006.

K. Hamer, E. Humphreys, S. Garthe, J. Hennicke, G. Peters, D. Gémillet, R. Phillips, M. Harris, and S. Wanless. Annual variation in diets, feeding location and foraging behaviour of gannets in the North Sea: flexibility, consistency and constraint. *Marine Ecology Progress Series*, 388:295–305, 2007.

P. Hammond and K. Grellier. Grey seal diet composition and prey consumption in the North Sea. Technical report, Executive Summary project MF0319, SMRU, University of St. Andrews, 2006.

P. Hammond, A. Hall, and J. Prime. The diet of grey seals around Orkney and other island and mainland sites in north-eastern Scotland. *The Journal of Applied Ecology*, 31(2):340–350, 1994a.

P. Hammond, A. Hall, and J. Prime. The diet of grey seals in the Inner and Outer Hebrides. *Journal of Applied Ecology*, 31(4):737–746, 1994b.

J. Harwood and J. P. Croxall. The assessment of competition between seals and commercial fisheries in the North Sea and the Antarctic. *Marine Mammal Science*, 4(1):13–33, 1988.

G. Holland, S. Greenstreet, I. Gibb, H. Fraser, and M. Robertson. Identifying sandeel *Ammodytes marinus* sediment habitat preferences in the marine ecosystem. *Marine Ecology Progress Series*, 303:269–282, 2005.

D. Horvitz and D. Thompson. A generalization of sampling without replacement from a finite universe. *J. Am. Stat. Assoc.*, 47:663–685, 1952.

ICES. Report of the ICES advisory committee 2008. Technical Report Book 6, pp 326, ICES Advice, 2008.

ICES. Ices catch by species, area and year (1973-2008). Eurostat/ICES database on catch statistics - ICES 2007, 2009.

H. Jensen, P. Wright, and P. Munk. Vertical distribution of pre-settled sandeel (*Ammodytes marinus*) in the North Sea in relation to size and environmental variables. *ICES Journal of Marine Science*, 60:1342–1351, 2003.

C. Laurenson and I. Priede. The diet and trophic ecology of anglerfish *Lophius piscatorius* at the Shetland Islands, UK. *J. Mar. Biol. Ass. U.K.*, 85:419–424, 2005.

M. Leopold, C. van Damme, C. Philippart, and C. Winter. Otoliths of North Sea fish - fish identification key by means of otoliths and other hard parts. version 1.0. Technical report, [CD-ROM] ETI (Expert Centre for Taxonomic Identification), University of Amsterdam, Amsterdam, the Netherlands., 2001.

S. Lewis, T. Sherratt, K. Hamer, M. Harris, and S. Wanless. Contrasting diet quality of northern gannets *Morus bassanus* at two colonies. *Ardea*, 91:167–176, 2003.

E. Limpert, W. Stahel, and M. Abbt. Log-normal distributions across the sciences: Keys and clues. *BioScience*, 51(5):341–352, 2001.

Macer. Sand eels (*Ammodytidae*) in the south-western North Sea; their biology and fishery. *Ministry of Agriculture, Fisheries and Food, Fishery Investigation Series*, 24(6):55, 1966.

G. Nævdal and S. Thorkildsen. Genetic studies on species composition and population structure of sand eels (genera: *Ammodytes, Hyperoplus and Gymnammodytes*) in Norwegian waters. *Journal of Applied Ichthyology*, 18 (2):124–126, 2002.

T. Pearson. The feeding biology of sea-bird species breeding on the Farne Islands, Northumberland. *Journal of Animal Ecology*, 37(3):521–552, 1968.

S. Pedersen, P. Lewy, and P. Wright. Assessments of the lesser sandeel (*Ammodytes marinus*) in the North Sea based on revised stock divisions. *Fisheries Research*, 41:221–241, 1999.

G. Pierce, M. Santos, and S. Cerviño. Assessing sources of variation underlying estimated of cetacean diet composition: a simulation study on analysis of harbour porpoise diet in Scottish (UK) waters. *Journal of the Marine Biological Association*, 87:213–221, 2007.

J. Prime and P. Hammond. The diet of grey seals from the south-western North Sea assessed from analysis of hard parts found in faeces. *The Journal of Applied Ecology*, 27(2):435–447, 1990.

A. Rijnsdorp, P. van Leeuwen, N. Daan, and H. Heessen. Changes in abundance of demersal fish species in the North Sea between 1906-1909 and 1990-1995. *ICES Journal of Marine Science*, 53:1054–1062, 1996.

A. Rindorf and P. Lewy. Analyses of length and age distributions using continuation-ratio logits. *Can. J. Fish. Aquat. Sci.*, 58:1141–1152, 2001.

M. Santos, G. Pierce, and J. Learmonth. Variability in the diet of harbour porpoises (*Phocoena phocoena*) in Scottish waters 1992-2003. *Marine Mammal Science*, 20 (1):1–27, 2004.

G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2): 461–464, 1978.

C. Stransky, H. Baumann, S. Fevolden, A. Harbitz, H. Høie, K. Nedreaas, A. Salberg, and T. Skarstein. Separation of Norweigian coastal cod and Northeast Arctic cod by outer otolith shape analysis. *Fisheries Research*, 90:26–35, 2008.

D. Tollit, M. Stewart, P. Thompson, G. Pierce, M. Santos, and S. Hughes. Species and size differences in the digestion of otoliths and beaks: implications for estimates of pinniped diet composition. *Can. J. Fish. Aquat. Sci.*, 54:105–119, 1997.

M. van Deurs, R. van Hall, T. Tomczak, S. Jónasdóttir, and P. Dolmer. Recruitment of lesser sandeel *Ammodytes marinus* in relation to density dependence and zooplankton composition. *Marine Ecology Progress Series*, 381:249–258, 2009.

M. Vinther. Ad hoc multispecies VPA tuning applied for the Baltis and North Sea fish stocks. *Journal of Marine Science*, 58:311–320, 2001.

D. Welch and R. Foucher. A maximum likelihood methodology for estimation length-at-maturity with application to Pacific cod (*Gadus macrocephalus*) population dynamics. *Can. J. Fish. Aquat. Sci.*, 45:333–343, 1988.

P. Wright, H. Jensen, and I. Tuck. The influence of sediment type on the distribution of the lesser sandeel, *Ammodytes marinus*. *Journal of Sea Research*, 44:243–256, 2000.

Chapter 3

Spatially Adaptive Multidimensional Smoothing: investigating knot placement using a branch and bound algorithm

3.1 Introduction

Spatial modelling describes methods which recognises spatial dependence between data points. This dependence may either be explicitly modelled through the use of spatial coordinates or through the use of explanatory variables which themselves vary across space. In this chapter we investigate a two-dimensional spatially adaptive smoothing technique for modelling species distributions. Understanding where species occur and why they occur in particular places is of great ecological importance and the prediction of the occurrence of species is fundamental for conservation biology and wildlife management (Williams et al., 2009). In a world where the distribution of the human race is expanding and our effects even more so, modelling the distribution of species is of increasing importance.

This chapter presents two case studies which illustrate how spatial modelling may be used to aid the management of the marine environment. We develop a novel

spatial modelling technique which attempts to address some of the limitations of existing techniques. The methods in this chapter, unlike many existing methods, can be employed in any modelling framework in which linear explanatory terms can be incorporated. The first case study uses these methods to investigate the distribution of both the density and presence of sandeels off the east coast of Scotland. These data were collected over six years, the first three years were collected while the fishery was still in operation and the latter three years were collected following the closure of the sandeel fishery, they can therefore be used to assess the impacts of closing this region to the sandeel fisheries. This case study is particularly useful from a statistical viewpoint as it illustrates how these methods can be implemented both within a logistic regression, to assess the spatial distribution of sandeel presence, and within a zero-inflated model framework, to assess the spatial distribution of sandeel density. It is hoped that such knowledge will lead to a better understanding of sandeel distribution and help ensure that the fisheries do not contribute to any further decline of this key species. The second case study applies these novel techniques within a Poisson count model to investigate the density of Harbour Porpoise off the west coast of Scotland. The aim of modelling marine mammal distributions, such as Harbour Porpoise, is so that naval exercises involving sonar may attempt to avoid the areas where animal density is predicted to be highest.

Although there has been much development recently in the spatial modelling field there still remains room for further development. Of the variety of techniques available, the parametric generalised linear models (GLM) and the semi-parametric generalised additive models (GAM) (Hastie and Tibshirani, 1990; Wood, 2006) are the most widely used for species distribution modelling (Guisan et al., 2002; Rushton et al., 2004). The development of the GAM was an important step allowing the assumptions of linearity associated with the GLM to be relaxed. However, controlling

these flexible relations to capture the underlying function whilst not modelling the noise can prove challenging. The most commonly used implementation of the GAM lies within the `MGCV` package in `R` (Wood, 2006). `MGCV` by default uses penalized regression splines, a form of thin plate regression spline whereby the degree of smoothing can be selected using various criteria or alternatively supplied by the user. However, there is concern that these methods can over-fit to the data in some situations (Baker, 2008).

A further concern and topic of current research is the ability of existing semi-parametric models to be able to accommodate surfaces with varying levels of smoothness. For example, it may be the case that the underlying distribution of a species varies a great deal more in one area of a survey region than another. In this situation we would like some areas of the surface to be more flexible than others to accommodate these features. The ability to accommodate this local flexibility is dictated by the method used to smooth the surface. Most current spline-based methods, such as the penalized thin plate regression splines implemented in `MGCV`, smooth the whole surface to the same degree (Wood, 2006). Specifically, this method does not require the user to define knot positions explicitly and instead uses a penalty term to control the "wiggliness" of the surface (Wood, 2006). In contrast, pure regression splines control the flexibility of a surface solely using the number and location of knots, which are often chosen according to a regular grid or using space filling algorithms (Ruppert et al., 2003; Royle and Nychka, 1998). However, carefully positioning knots across a surface can be a valuable tool in controlling the flexibility of a surface. Increasingly, spatially adaptive smoothing techniques based on knot selection routines are being investigated to achieve improved fit and prediction (Mammen and van de Geer, 1997; Zhou and Shen, 2001). More recently, Walker et al. (2010) published a spatially adaptive local smoothing algorithm (SALSA) based on iteratively adding, removing and

relocating knot positions in an attempt to reduce the largest residual. In addition, Cox (2008) demonstrated the potential of a branch and bound (global optimisation) algorithm (section 3.2.3) to find optimal knot locations for modelling missing data from acoustic krill surveys. This chapter further investigates the potential of the branch and bound algorithm for knot placement in the context of modelling species distributions.

Pure regression splines, with requisite knot locations, return basis functions which decay with distance from the knot location. There are a variety of basis functions available; one of the most commonly implemented is the thin plate spline (TPS) basis. However, the TPS basis has three drawbacks, it is not scale invariant (Wood, 2006), it is global in its effects (Wood, 2006) and its shape is such that it can predict very large values just outside the range of the data (Austin and Meyers, 1996). Its global nature can be especially problematic for spatially adaptive modelling as the bases have an effect on all points across a surface. Therefore, in additions to TPS bases we also investigate local Gaussian spline (LGS) bases which are effectively local since they tend to a zero as distance from the knot position increases, thus reducing the risk of predicting very large values outside the range of the data. In addition, the rate of decay of these bases can be chosen to be specific for each dataset by altering the value of the parameter r.

One final consideration addressed in this chapter is the ongoing debate about whether to include spatial coordinates (e.g. latitude and longitude) in a model with environmental covariates (Austin, 2002). One potential downside of including both spatial and environmental covariates in a model is the collinearity that can result due to overlap between the information contained in the spatial location and environmental variables. This collinearity can result in very large estimates of precision

for some model parameters (Graham, 2003) and correspondingly large p-values. This can be exacerbated when local changes in the flexibility of a spatial surface happen to coincide with one or more of the environmental covariates. Even when the overlap between the spatial co-ordinates and environmental covariates is not prohibitively severe, fitting them together in a model may still cause incorrect conclusions to be drawn about the relation between the response and environmental covariates. For example, it is suspected that in some circumstances performing model selection while considering both spatial coordinates and environmental covariates simultaneously will cause important environmental covariates to be excluded from the model. For this reason, many biologists resist including spatial coordinates in their models.

Despite these concerns, modelling species distribution using spatial co-ordinates is often valuable. Typically many environmental factors which affect species distribution are not measured and spatial variability remains, which can be captured using models which include spatial coordinates. We therefore follow a three stage modelling approach. Firstly, a model is fitted to the environmental variables alone. Secondly, the residuals from this model are modelled using spatial co-ordinates and a spatially adaptive smoothing technique. Finally, the spatial basis functions associated with the knot positions identified in stage 2 are included alongside the important environmental covariates identified in stage 1.

This chapter assesses the above methods using two case studies. We compare the branch and bound algorithm for knot placement using both the thin plate spline and the local Gaussian spline basis functions with the industry standard, penalised thin plate regression splines implemented in `MGCV`. Although the case studies presented here involve modelling species distributions, the methods proposed are applicable in a range of other situations (e.g. image analysis, spatial epidemiology) and could in

theory be extended to modelling in higher dimensions.

3.2 Methods

This section includes an overview of the general method we propose for fitting a 2-dimensional smooth, followed by details of pure regression splines and the branch and bound algorithm which are generic across case studies. Those parts of the procedure which need to be customised dependent on the dataset, as well as standard statistical modelling decisions such as testing for collinearity and non-linearity, are detailed separately for each case study in sections 3.3.3 and 3.4.3.

3.2.1 Overview

This chapter demonstrates a novel method for selecting and including a 2-dimensional spatial smooth in a variety of modelling frameworks. Although the zero-inflated and generalised additive models, both logistic and Poisson, are used in this chapter as examples, these methods may be implemented in any modelling framework that allows the inclusion of linear explanatory terms. The main stages in implementing this method are as follows:

1. Choose a modelling framework suitable for your data, e.g. generalised linear model, generalised additive model, generalised estimating equations, zero-inflated model etc.

2. Perform model selection on any covariates you wish to include in the model but excluding the 2-dimensional spatial coordinates. We shall refer to this model as the "environmental covariate" model.

3. Extract the most Normally distributed residuals from your environmental covariate model. These by definition should be the deviance residuals but some

modelling frameworks (e.g. zero-inflated models) may not provide deviance residuals. In this case you should consider other types of residual, e.g. Pearson or response residuals. These residuals will be used to try to find the 2-dimensional spatial smooth which best explains the remaining variability in the data.

4. Select all data points you would like to consider as candidate knot locations for the pure regression splines (Section 3.2.2). If you have data which are sparsely distributed throughout the spatial range of the data, and not too many data points, then you may wish to consider all data points as candidate knot locations. If you have data over several years implying that some of the data points have the same spatial coordinates, or if the data are too densely distributed within the spatial range you may need to choose a subset of the data. Candidate knot locations which are too close together may cause problems in the branch and bound routine. The candidate knot locations should be chosen using the spatial coordinates in the dataset (so that knots may only be located where there is a data point) and should give as even a coverage of the study region as possible.

5. Divide the candidate knot locations into starting groups of no more than 30 (due to computational restrictions). Each group of candidate knot locations should give as even a coverage of the survey region as possible. Other methods, such as dividing the candidate knots into groups that were close together in space were tested but led to less optimal models.

6. Construct basis functions for each of the candidate knot locations (Section 3.2.2). In this chapter we consider two basis functions, those associated with the thin plate regression spline and the local Gaussian regression spline. These are discussed and compared for each case study in Sections 3.3.5.4, 3.4.5.1 and 3.5.2.

If you are uncertain of which to use, cross validation techniques are advised for spline selection.

7. Use the branch and bound algorithm (section 3.2.3) to perform an exhaustive search on the set of basis functions associated with each starting group of knots. The `regsubsets` routine in the `leaps R` library (Lumley, 2009) is used in the analyses presented in this chapter. This routine takes in a response variable (the residuals from the environmental covariate model) and a number of explanatory variables (the basis functions associated with the candidate knot locations). This routine returns the best possible subset of explanatory variables (basis functions) for each size of subset and provides a selection of criteria for choosing between them. To choose the best number of basis functions to model the residuals we recommend the BIC criterion (Schwarz, 1978), see Section 3.5.1 for further discussion.

8. Group the resulting selections of basis functions for each starting group with another set of selected basis functions. This was done by grouping the basis functions chosen from the first and second groups, third and fourth groups etc. Steps 7 and 8 are then repeated until only one group of selected basis functions remains.

9. Add the 2-dimensional smooth associated with this final selection of basis functions to the environmental covariate model to form the "full model". This is done by simply including the selected basis functions in the environmental covariate model as linear terms.

Note that the requirement for steps 5 and 8 occurs due to computational limitations. Ideally we would simply supply the `regsubsets` routine with all candidate knots at once and perform an exhaustive search over the entire set.

3.2.2 Pure Regression Splines

Generalised additive models are an extension of generalised linear models which relax the assumption of linear relations between the response and the explanatory covariates on the link scale (Hastie and Tibshirani, 1990; Wood, 2006). One of the ways this can be done is through the use of splines; these replace the previously linear relations with data-driven non-linear functions. The implementation of the pure regression spline essentially involves creating additional covariates (i.e. additional columns in the model design matrix), usually one for each knot location in the covariate space. The values of these "covariates" are calculated as some type of smooth basis function ($f(d)$) across the range of the explanatory variable. The coefficients are then estimated in the same way as those associated with linear terms, since the function is linear in its parameters. The addition of these basis functions multiplied by their respective parameter estimates leads to a non-linear smooth function representing the relation between the explanatory and response. The general form of this model is seen in equation 3.1.

$$E[Y_i] = \beta_0 + \beta_1 x_{1i} + ... + \beta_m x_{mi} + \sum_{k=1}^{K} \beta_{k+m} f(d_{ik}) \tag{3.1}$$

where $Y_i \sim$ some exponential family distribution.

Y_i is the i-th observation, the x_i's are the m linear explanatory covariates, $f(d_{ik})$ is some basis function of the Euclidean distance d to knot location k for k in 1,...K knots and the β's are the regression coefficients. In the case of a two dimensional smooth across latitude (lat) and longitude (lon) this distance would be calculated as

$$d_{ik}^2 = \|\boldsymbol{x}_i - \boldsymbol{x}_k^*\| = (lat_i - lat_k^*)^2 + (lon_i - lon_k^*)^2 \tag{3.2}$$

where $\boldsymbol{x}_i = [lat_i, lon_i]^T$ is the i-th coordinate vector and $\boldsymbol{x}_k^* = [lat_k^*, lon_k^*]$ is the k-th knot position vector (Hastie et al., 2001).

Thin Plate Regression Splines (TPRS)

Thin plate regression splines (TPRS) are constructed using the thin plate spline (TPS) basis function, equation 3.3 (Harder and Desmarais, 1972). They are one of the most well known basis functions (Buhmann, 2000) and are commonly used in species distribution modelling, for example shorebirds (Granadeiro et al., 2004) and marine mammals (Panigada et al., 2008). However, despite their popularity there is some debate as to their optimality. It can be seen from equation 3.3 that the value of the TPS basis function increases with distance d (where $d > 1$) from the knot location and these splines are therefore global in their effects.

$$f\left(d_{ik}\right) = d_{ik}^{2} log\left(d_{ik}\right) \tag{3.3}$$

Figure 3.1 provides an illustration of a 2-dimensional smooth using TPRS.

Local Gaussian Regression Splines (LGRS)

Local Gaussian splines (LGS) are less commonly associated with GAMs, however they have some useful properties which may improve the modelling of species distributions. As their basis functions are constructed using a Gaussian decay (equation 3.4) they tend to a flat plane with increasing distance d from the knot position.

$$f\left(d_{ik}\right) = exp\left(\frac{-d_{ik}^{2}}{r^{2}}\right) \tag{3.4}$$

In contrast to the global nature of the TPS, the LGS can be more local in its effects with the basis function decaying to effectively zero beyond $3r$. In addition, how fast they approach zero can be controlled based on the value of r which can be chosen specifically for each dataset. The smaller the value of r the more local the effects of the basis function. Figure 3.2 provides an illustration of a 2-dimensional smooth

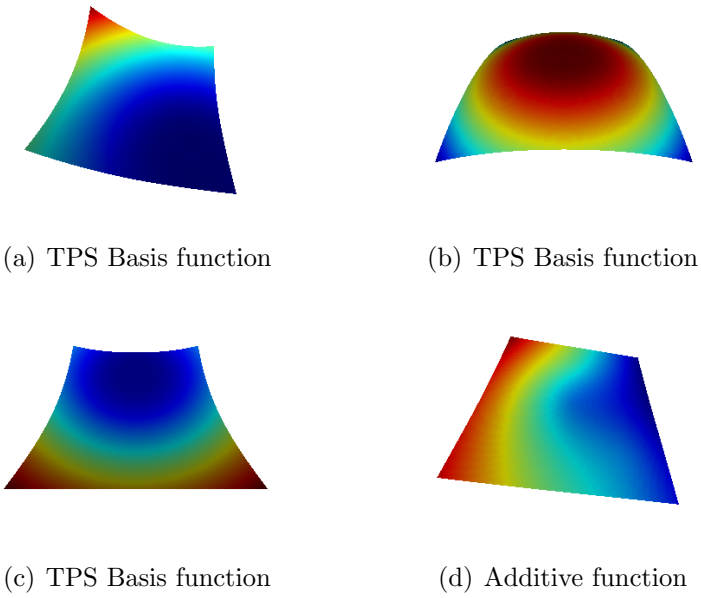

(a) TPS Basis function (b) TPS Basis function

(c) TPS Basis function (d) Additive function

Figure 3.1: Thin plate regression spline example. The basis functions are colour coded from blue for low values to red for high values. Panels (a) to (c) show three basis functions which have been multiplied by their respective coefficients. Note that the inverted shape of panel (b) indicates a negative coefficient. Panel (d) shows the resulting function based on the summation of these three basis functions.

using LGS.

3.2.3 Branch and Bound Algorithm

Through the implementation of pure regression splines we have turned the problem of smoothing into one of variable selection. The same as we would choose from a number of explanatory covariates we can now perform variable selection on a number of spline basis functions each associated with a different knot location. Currently the only method of variable selection which ensures a global optimum is reached, rather than the local optimums which may be reached using methods such as stepwise selection, is an exhaustive search (Miller, 2002). The branch and bound is an algorithm which can be used to perform this exhaustive search and find this global optimal

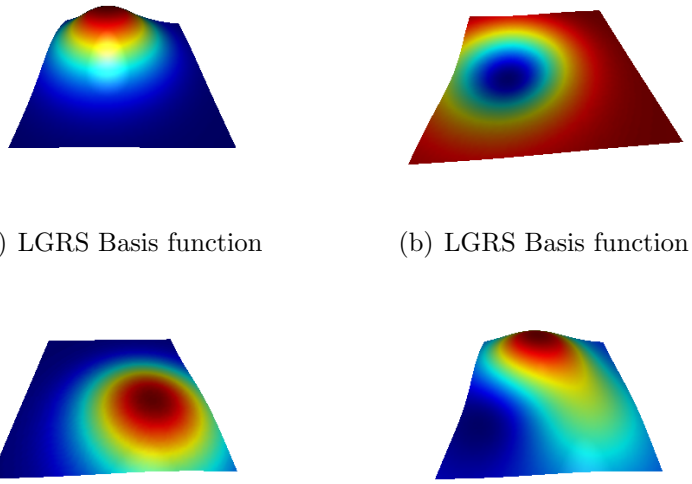

(a) LGRS Basis function (b) LGRS Basis function

(c) LGRS Basis function (d) Additive function

Figure 3.2: Local Gaussian regression spline example. The basis functions are colour coded from blue for low values to red for high values. Panels (a) to (c) show three basis functions which have been multiplied by their respective coefficients. Note that the inverted shape of panel (b) indicates a negative coefficient. Panel (d) shows the resulting function based on the summation of these three basis functions.

solution without actually calculating all possible outcomes. One of the earliest implementations of this algorithm for the purposes of variable selection can be found in Furnival and Wilson (1974). In this chapter, we apply the `regsubsets` function in the `R` library `leaps` (Lumley, 2009).

The branch and bound works by trying to minimise the residual sums of squares (RSS). It does this by iteratively dividing explanatory variables (or in our case basis functions) into two or more smaller sets, which depending on the implementation may or may not overlap. The recursive application of this procedure creates a tree structure, sometimes termed the 'search tree', with each of the subsets of basis functions as the nodes. The minimum RSS is then simply the smallest RSS over all possible

subsets / child nodes. The bounding part of the algorithm involves calculating these upper and lower bounds for the minimum RSS for a given subset. When comparing these bounds, if a lower bound of one subset is higher than upper bound of another subset the former subset may be safely removed from the search tree, a process known as pruning. The branch and bound terminates when either the current set of basis functions only contains one element or alternatively when the upper bound of the current candidate set matches the lower bound.

3.3 Case Study 1: Sandeel distribution in the Firth of Forth and Wee Bankie region

3.3.1 Introduction

The lesser sandeel (*Ammodytes marinus*), introduced in Chapter 2, is a vital component of the North Sea's food web (Greenstreet et al., 2006). Not only is it an important food source for many seabirds (Furness, 2002; Frederiksen et al., 2006) and marine mammals (Santos et al., 2004; Hammond and Grellier, 2006), it is also predated by many commercial fish species (Bromley et al., 1997; Greenstreet et al., 1998; Engelhard et al., 2008) and, in addition, is itself subject to large-scale industrial fishing. However, the dramatic drop in sandeel landings from over 1.1 million t in 1997 to less than 0.5 million t per annum since 2003 (ICES, 2009) has triggered concern over the ability of sandeel stocks to withstand this intense exploitation. In 2000, this concern, coupled with reduced breeding success of seabirds led to the closure of the sandeel fisheries off the east coast of Scotland in the Firth of Forth and Wee Bankie region (Greenstreet et al., 2006). The data in this study were collected between 1997 and 2002 and therefore allow us to assess the impacts of this closure on the distribution of sandeels in the region.

Being such an important species, both ecologically and economically, sustainable management of *A. marinus* is a priority. To achieve this, we must be able to accurately assess the population size of the sandeel stock and understand the factors which affect its distribution. In addition, Matthiopoulos et al. (2008) highlights the importance of understanding the spatial availability of prey so that this can be related to predator consumption. Although there have been a number of studies relating to sandeels in the North Sea looking at population dynamics (Arnott et al., 2002), stock assessment (Pedersen et al., 1999), habitat preference (Wright et al., 2000) and evaluation of management strategies (Greenstreet et al., 2006), there are still caveats in our knowledge. For example, Wright et al. (2000) suggests that sandeel aggregations could contract and expand around preferred sediment type, a behaviour which may result in fisheries maintaining high catch levels even when the stock has become depleted. In addition, Pedersen et al. (1999) discuss the importance of choosing the correct stock assessment regions for the North Sea based on evidence which suggests the existence of a number of smaller self-sustained sub stocks which need to be assessed individually. They also point out that there is a lack of information on the abundance of sandeels outside the fishing grounds. This study uses novel spatially adaptive modelling techniques to provide valuable insights into some of the issues associated with assessing the distribution of the sandeel stock.

The acoustic data reported in Greenstreet et al. (2006) presents two main challenges; they contains a large number of zero's and they have a patchy distribution. Firstly, this study investigates the potential of spatially adaptive modelling to better predict the patchy distribution of the sandeels' presence and absence. Secondly, we demonstrate how our spatially adaptive smoothing can be implemented within a zero-inflated framework, something which to our knowledge has not previously been

achieved. We also consider whether including the additional spatial covariate information in this model affects the interpretation of the relations between the response variate and the environmental covariates and whether a three stage modelling approach reduces the associated problems of collinearity. Finally, we assess if the more local Gaussian spline basis functions are superior to the globally acting thin plate spline for modelling the patchy distribution of the sandeel.

3.3.2 Data

The survey region (Figure 3.3) covers an area off the east coast of Scotland from the Firth of Forth to the Tay estuary (56;0N to 56;30N, 03;0W to 01;0W) and includes the main fishing grounds in the area, the Wee Bankie, Marr Bank, and Berwick's Bank (Greenstreet et al., 2006). The data were collected using acoustic sampling techniques over 3 to 8 consecutive days in either June or July of each year from 1997 to 2002. From 1997 to 1999 the sandeel fishery in the area was active, but was closed from 2000 onwards. There were some changes to sampling effort: in 1998 the two most northerly transects were omitted from the survey due to adverse weather conditions, while in 2000 and 2001 additional effort was expended across the major sandbanks and around the Isle of May where a major seabird colony breeds.

Acoustic surveys can detect sandeels in the water column but cannot detect those which are buried in the sediment. For this reason, the surveys were carried out between 03:00 and 15:00 hours, when the majority of sandeels emerge from the sediment to feed. These acoustic measurements had been processed prior to analysis and converted into measurements of sandeel density (g/m^2) (Greenstreet et al., 2006). The transects were divided into between 40 and 126 segments and the sandeel density was integrated over each segment to give a single value. The midpoint coordinates of

each segment, along with the date and time, were also recorded. The spatial coordinates were converted from latitude and longitude to east and north nautical miles to ensure equivalent scales in both dimensions. In addition, various environmental covariates were recorded, including depth, salinity, density, temperature, and various measures relating to sediment type. Further information on survey methods and data processing can be found in Greenstreet et al. (2006).

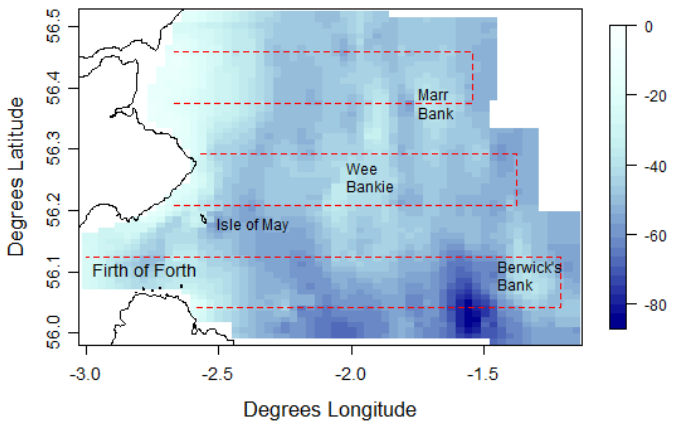

Figure 3.3: Survey region off the east coast of Scotland. The depth (m) of the seabed is indicated by the blue shading and the main sand banks / fishing grounds (Wee Bankie, Marr Bank, Berwick's Bank) are labelled. The planned main survey effort is indicated by the red dashed lines.

3.3.3 Modelling Methods

3.3.3.1 Zero-Inflated Model

A common problem when modelling species distributions, and one which is present in this dataset, is an excess of zero observations. Specifically, the data contain a larger proportion of zeros than expected given the distribution that would typically have been used to model them. Failure to model these zeros correctly can lead to biased estimation of ecological effects and their variability, which would result in incorrect

model inference and prediction (Martin et al., 2005). Therefore, we must use models specifically designed for this type of data. Two commonly used types of model for this situation, are the hurdle model (Cragg, 1971; Mullahy, 1986) and the zero-inflated model (Lambert, 1992).

Zero observations in the context of modelling species distributions are usually classified as either true or false zeros. True zeros, also termed structural zeros, arise when the species never occupies a site. False zeros are the random zeros which occur by chance, even though the species does occupy the site. False zeros can either arise due to the species not being present at the time of the survey, or the species is present but the observer fails to detect it. The difference between the hurdle model and the zero-inflated model is their approach to the zeros. The hurdle model makes no distinction between true and false zeros; firstly, it uses a binomial model and associated probability "hurdle" to determine whether a data point will have a zero outcome or a positive realisation. If a data point has a probability of presence greater than the hurdle value (which is not necessarily zero) its fitted value is determined by the conditional distribution, generally a count distribution truncated at zero. In comparison, the zero-inflated model approaches the problem using a mixture model, consisting of a count component and a point mass at zero component. The difference being, zeros could now either have arisen from the point mass at zero or from the count component. This specification means that the point mass at zero component, is now only modelling the probability of a true (or structural) zero.

This distinction is an important one when modelling species distributions, however if the aim is to simply quantify the instantaneous distribution of a species (e.g. as a part of biomass estimation) then the distinction between the different types of zero can be ignored. If the aim is to model habitat preference then the difference between

a true zero and a false zero is important. In addition, Martin et al. (2005) points out the importance of considering which zeros are the cause of the zero-inflation problem when choosing a modelling technique. While the zero-inflated model can be used for an excess of true or false zeros, the hurdle model generally only has a satisfactory application when the excess of zeros are true zeros. It should be noted, however, that there is currently very little literature which discusses preferred methods when the excess zeros are both true and false. The source of the zeros in this dataset is up for debate. While many of the zeros could be false zeros as sandeels form highly aggregated groups when they are in the water column and therefore have a patchy distribution throughout the range of the habitat they occupy, there is also an argument that many of the zeros could be true zeros. Sandeels have very specific sediment preferences and therefore some of the habitat within the range of the survey region would be unsuitable for the sandeels to bury themselves in at night (Wright et al., 2000). Such sediment preferences may also limit the spatial distribution of the sandeel in the water column during the day.

The aim of this investigation was to assess if including a spatial smooth gave a better description of the the underlying spatial distribution of sandeels. Our interest in predicting sandeel density as well as how the environmental covariates are related to sandeel distribution meant the treatment of zeros was important. For this reason, and considering that many of zeros were suspected to be false zeros, a zero-inflated model was implemented.

The zero-inflated models were fitted using the `zeroinfl` routine in the `pscl` library within `R` (Zeileis et al., 2008). The zero-inflated model implemented within `pscl`, unlike the `gamlss` (Stasinopoulos and Rigby, 2007) and `VGAM` (Yee, 2008), allow different sets of regressors for the zero component to the count component. Although

generalised additive modelling techniques have been previously implemented within a zero-inflated framework (e.g. Barry and Welsh (2002); Yee (2008)), to our knowledge no one has yet implemented spatially adaptive generalised additive modelling techniques within a zero-inflated modelling framework.

The available data comprise of measurements of sandeel density, collected via acoustic sampling. However, the distributions implemented in `pscl` for the non-zero part of the model are the Poisson and Negative Binomial, and therefore the response variable must be integer values. As the majority of the densities were less than 1, they were multiplied by 100 to preserve the information which was recorded to 2 decimal places. The predictions were then back-transformed later to the original scale.

3.3.3.2 Presence / Absence Models

When the probability of a species being present is of interest rather than its abundance or density, logistic regression is commonly used. This study also considered the use of the branch and bound method, in the context of a Binomial Generalized Additive Model used to predict the sites where sandeels are most likely to occur.

3.3.3.3 Environmental Covariate Selection

Throughout this case study, for ease of description, we refer to all covariates which are not spatial coordinates (i.e. North and East) as environmental covariates, and the models which are fitted to these alone, as environmental covariate models. Under this definition we include temperature, salinity, density, depth, time of day, year and substrate measures RoxAnn e1 and e2. The year factor variable was transformed into a two level factor variable relating to those years before the closure of the fishery (1997-1999) and those years after the closure of the fishery (2000-2002). Those models which include spatial coordinates in addition to the environmental covariates, will henceforth be referred to as the full models.

Variance inflation factors (VIFs) were used to assess collinearity (Fox, 1997) since including a number of collinear variables in a model may result in incorrect inference with regard to those covariates. Generalized additive models (GAM) were fitted using the default options in the `MGCV` library (Wood, 2006) to assess whether the relations between the response and the environmental covariates were significantly non-linear. These were assessed based on the associated p-values and verified by a visual assessment of the partial regression plots. This process was carried out for both the presence/absence data using a binomial distribution, and for the count data using a Poisson distribution. Any non-linear relations were included in the zero-inflated model or binomial GAM through the implementation of cubic B-splines with a knot placed at the median value (Hastie and Tibshirani, 1990). A forward stepwise procedure was used to select the environmental covariates because including all environmental covariates in the zero-inflated model resulted in non-convergence of the model. For the zero-inflated model, forwards stepwise selection was performed on the binomial component first and then the count component. The final models were selected based on the minimum Bayesian information criterion (BIC, equation 3.5) (Schwarz, 1978).

$$BIC = -2lnL + \theta log(n) \tag{3.5}$$

where lnL is the log-likelihood of the model, θ represents the number of parameters in the model and n is the number of data points used to fit the model.

3.3.3.4 Including Spatial Coordinates

Residuals from the environmental covariate models were used to select the specific knot locations for the 2-dimensional spline basis functions across the spatial covariates (East and North). As the branch and bound attempts to find the "best" subset of knot locations based on minimising the residual sum of squares, ideally, we would

like residuals which are normally distributed. In the case of the binomial model, the deviance residuals were used since these are approximately normal (Pierce and Schafer, 1986). The zero-inflated model only provides Pearson and response residuals. On inspection, the response residuals were more normally distributed and were used to govern the branch and bound algorithm.

Due to a large number of (almost identical) spatial co-ordinates in the data set, a subset of these were used as candidate knot locations. A grid was created over the range of the spatial coordinates to ensure the selected locations were evenly spaced throughout the range of the data, and the data points with coordinates closest to the centre of each grid cell were used as the candidate locations (Figure 3.4(a)). The grid scale was selected so that it was fine enough to give good spatial coverage but not too fine as to make the fitting process unacceptably time consuming.

In theory, the branch and bound has no limit on the number of candidate knot locations. However, in practice with current computing power, the branch and bound algorithm implemented here can only deal with up to approximately 30 candidate knots. Manually, limiting the candidate knot locations to 30 runs the risk of poor spatial coverage. To reduce this risk, an iterative routine was used to allow all candidate knot locations indicated in Figure 3.4 an opportunity to be selected. This method therefore distinguishes itself from Cox (2008) in that it allows more than 30 candidate knot locations to be considered. Initially, the candidate knot locations were divided up into groups of 30 coordinates, which were chosen so that each subset maintained good spatial coverage. This was done by moving along the transects and assigning the first candidate knot location to the first group, the second to the second etc. until the last group had been assigned a knot location and then the next knot location along the transect was assigned to the first group and so on. This was

repeated until all the candidate knot locations had been assigned starting groups. An example of these starting groups, each represented by a different colour, is shown in Figure 3.4(b).

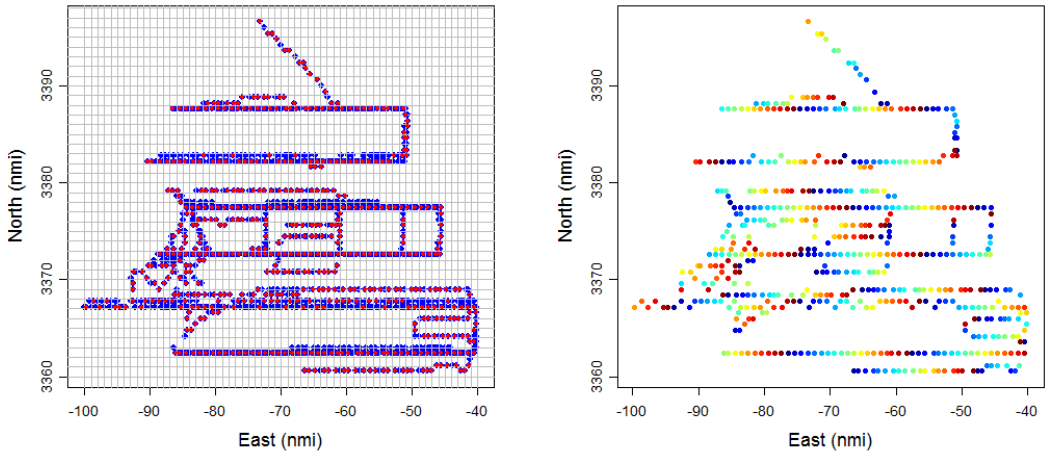

(a) Candidate Knot Selection

(b) Groups for Initial Selection

Figure 3.4: Initial selection process of candidate knot locations. (a) Indicates the grid laid across all data points. The points closest to the centre of each grid cell are selected and coloured red; the unselected data points are blue. (b) Indicates only the selected candidate knots. These are colour coded to indicate the initial grouping. Note that due to the number of knots sets, some colours for adjacent points only differ slightly in shade.

The branch and bound knot selection process is carried out by the `regsubsets` routine provided by the `leaps` library in R. The regsubsets routine returns the best possible subset of basis functions, for all possible sizes of subset. In addition, it provides various criteria for choosing between these; again the BIC was used to choose the best sized subset of knot locations. Initially, the best subsets from each of the start groups are obtained (Figure 3.5(a)) and each group of knots selected at iteration 1 is then combined with an adjacent group of selected knots, and this process is repeated until only one set of selected knots remain. This process is illustrated in Figure 3.5.

The final set of knot locations are then used to generate the spatial basis functions which are included in the environmental covariate models to give the full models.

3.3.3.5 Model Assessment

The application of the branch and bound iterative knot selection technique was explored in a number of ways. Firstly, its ability to model the spatial variability remaining in the residuals, once the environmental covariates were accounted for, was investigated for both the presence/absence GAM as well as the zero-inflated model. Secondly, the full GAM and zero-inflated models, which included both the spatial information and the environmental covariates, were compared with environmental covariate models. In addition, the models of the residuals were compared to intercept only models and GAMs using the penalised thin plate regression splines implemented in MGCV. Finally, the full presence/absence GAM constructed using the branch and bound was compared with a full presence/absence GAM using the penalized thin plate regression splines in MGCV.

The predictive power of the models to data unseen by the model was assessed using 10 fold cross validation (Kohavi, 1995). The data were first divided into independent blocks based on the autocorrelation detected in the auto-correlation function plots (Venables and Ripley, 2002) of the residuals. Next, these blocks were divided into 10 non-overlapping subsets of equal size. Each of these 10% subsets were removed in turn and the remaining 90% of the data was used to fit the respective models. These models were then used to make predictions to the 10% of the data which were removed before model fitting. The ability of the model to predict this unseen data is measured using the mean square error (MSE) (equation 3.6), where n_v is the number of data points removed.

$$MSE = \frac{\sum(predicted - observed)^2}{n_v} \tag{3.6}$$

The fit to the data was also assessed using a pseudo R^2 score (equation 3.7), as there are no comparable R^2 values across these different modelling techniques. AIC and BIC statistics were also assessed for their consistency with the MSE results.

$$pseudoR^2 = (correlation(fitted, observed))^2 \tag{3.7}$$

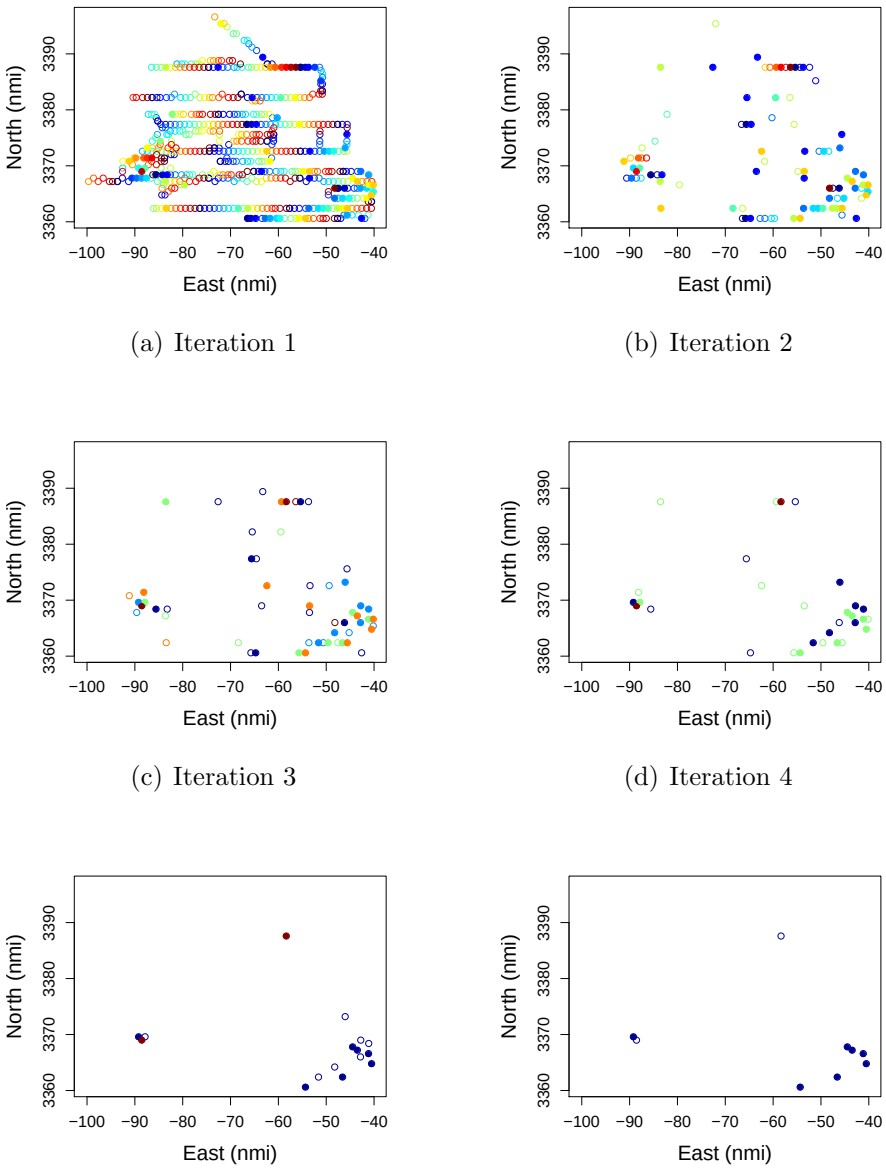

(a) Iteration 1

(b) Iteration 2

(c) Iteration 3

(d) Iteration 4

(e) Iteration 5

(f) Iteration 6

Figure 3.5: Iterative knot selection process: Open circles represent candidate knot locations, closed circles represent selected knot locations and the colours allow each of the groups to be distinguished. At each iteration the circles correspond to the closed circles of the previous iteration.

3.3.4 Results

3.3.4.1 Model Selection

High levels of collinearity were detected between salinity and water density (VIF values of 21 and 27, respectively.) and salinity, rather than water density, was retained as a possible explanatory covariate based on the highest pseudo R^2 values. Collinearity was also evident between year, day of year and temperature. While year and temperature were retained as potential explanatory covariates, day of year was not found to be significant at the 5% level once all other covariates were included, and was omitted from the models. The final models considered the following as candidate covariates, salinity, year, temperature, depth, time of day and substrate measures RoxAnn e1 and RoxAnn e2, as well as a factor representing pre and post closure of the fishery.

Presence / Absence Model

Time of day, temperature, RoxAnn e1 and RoxAnn e2 were significant predictors of sandeel presence in the environmental covariate model (Table 3.1). While the relations between sandeel presence and both time of day and RoxAnn e1 were significantly non-linear at the 5% level (based on the MGCV p-values), the relations between sandeel presence, and temperature and RoxAnn e2 were not found to be significantly non-linear. There was also a significant relation (at the 5% level) between sandeel presence and the pre/post closure factor variable

The branch and bound method selected seven local Gaussian spline (LGS) basis functions to model the remaining spatial variability found in the residuals of the environmental covariate model (Table 3.1). These seven basis functions, as well as linear terms for North and East, formed the spatial smooth which was found to be a highly significant predictor of sandeel presence (Table 3.1). Despite the additional

coefficients the BIC decreased from 2878 to 2822 when the LGS branch and bound spatial smooth was included in the model. In contrast, the BIC increased to 2907 when the global MGCV smooth was used ($df = 11.2$). In addition, the pseudo R^2 increased from 0.16 to 0.21 when the LGS branch and bound spatial smooth was included in the model, but only increased to 0.18 when the global MGCV smooth was used.

Table 3.1: Presence/Absence model selection results. The degrees of freedom (df) used for each covariate are indicated in the second column. The significance of the relation with each covariate is indicated by the p-values, for both the environmental covariate (EC) model and the full model. The p-values for the linear terms are taken from the model summary while those for the smooths are calculated using likelihood ratio tests. The full model was fitted using local Gaussian splines with an r value of 10.

Covariate	df	p-value (EC model)	p-value (Full model)
s(Time of Day)	4	< 2.2e-16 ***	3.616e-14 ***
Temperature	1	1.332e-11 ***	4.191e-07 ***
Pre/Post Closure	1	0.0009248 ***	1.656e-05 ***
s(RoxAnn e1)	4	2.417e-08 ***	0.002462 **
RoxAnn e2	1	0.0109324 *	0.173971
s(North, East)	9		< 2.2e-16 ***

Including the branch and bound-based 2-dimensional smooth had consequences for the RoxAnn e2 covariate (Table 3.1); this term was no longer significant at the 5% level In addition, model selection based on the minimum BIC would no longer choose to retain this covariate in the branch and bound based model. This change in significance (and BIC-related selection) was also apparent when the global MGCV-based smooth was fitted.

Apart from the change in interpretation regarding RoxAnn e2, the inclusion of the branch and bound spatial smooth has not vastly changed the interpretation of the

model regarding the relations between the other environmental covariates and sandeel presence. In both the environmental covariate model and the full model, the sandeels are more likely to be present in the water column at around 8 or 9 am and then later at around 7pm (Figs. 3.6(a) and 3.6(b)). There also appears to be a higher probability of sandeel presence for values of the RoxAnn e1 sediment between 0.5 and 1.1 (Figs. 3.6(c) and 3.6(d)). These similarities are due to no substantial (or statistically significant) differences in the parameter estimates for the full model and the environmental covariate model. These consistencies were also observed when the global MGCV smoothing approach was used (Fig. 3.7).

Such consistencies are also seen in the liner relations between the environmental covariates and sandeel presence (table 3.2). Under the EC model, the probability of sandeel presence is predicted to decrease linearly with increasing temperature (between 7.5 and 11.5 degrees Celsius) and increase linearly with the RoxAnn e2 sediment measure (between 0.5 and 1.9). In addition, all models indicate that the probability of sandeel presence is significantly higher post-closure of the sandeel fisheries than pre-closure.

Table 3.2: Linear parameter estimates from the presence/absence models on the scale of the link function. Results are presented for the environmental covariate (EC) model and the full models fitted using both the branch and bound (B&B) and MGCV. The standard errors of the estimates are given in parenthesis

Covariate	EC Model	Full Model (B&B)	Full Model (MGCV)
Temperature	-0.60 (0.09)	-0.49 (0.10)	-0.54 (0.09)
Pre Closure	-0.53 (0.16)	-0.74 (0.17)	-0.62 (0.17)
RoxAnn e2	0.60 (0.24)	0.48 (0.35)	0.48 (0.35)

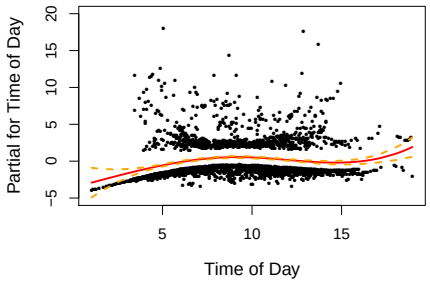

(a) Partial residual plot for the time of day covariate in the environmental covariate model

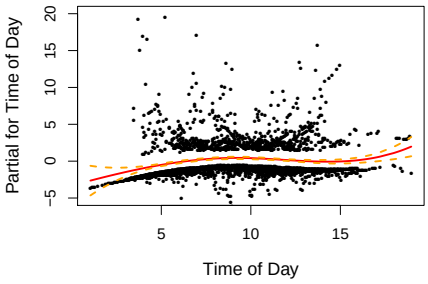

(b) Partial residual plot for the time of day covariate in the full model

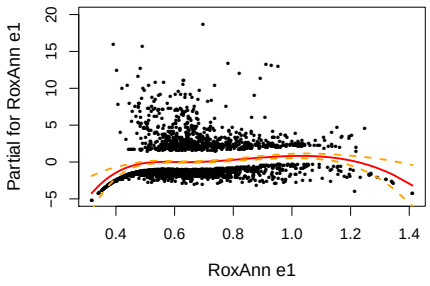

(c) Partial residual plot for the RoxAnn e1 covariate in the environmental covariate model

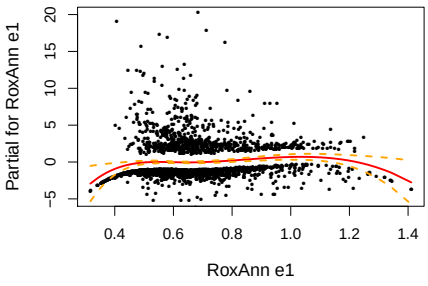

(d) Partial residual plot for the RoxAnn e1 covariate in the full model

Figure 3.6: Comparison of the partial residual plots for the environmental covariate model and the branch and bound LGS GAM modelling sandeel presence/absence. Panels (a) and (c) show the relations between sandeel presence and the covariates in the environmental covariate model. Panels (b) and (d) show the relations between sandeel presence and the covariates in the full model.

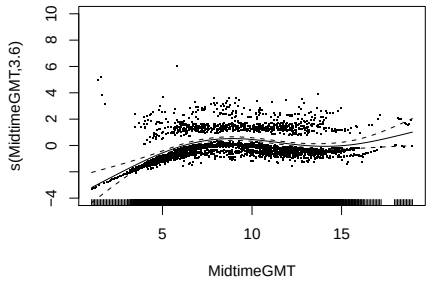

(a) Partial residual plot for the time of day covariate in the environmental covariate model

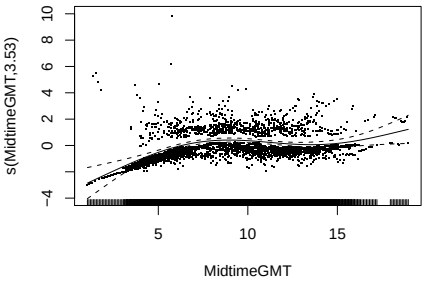

(b) Partial residual plot for the time of day covariate in the full model

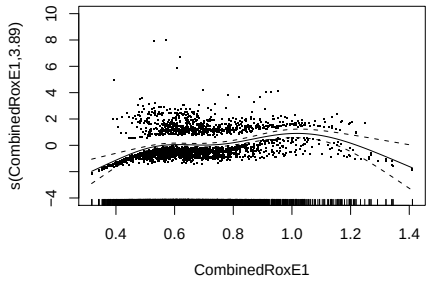

(c) Partial residual plot for the RoxAnn e1 covariate in the environmental covariate model

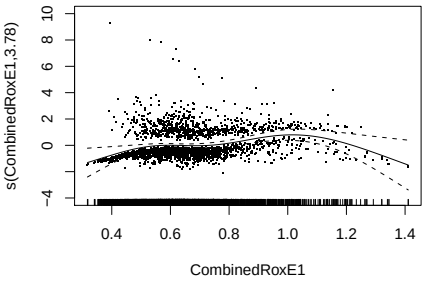

(d) Partial residual plot for the RoxAnn e1 covariate in the full model

Figure 3.7: Comparison of the partial residual plots for the environmental covariate model and the MGCV GAM modelling sandeel presence/absence. Panels (a) and (c) show the relations between sandeel presence and the covariates in the environmental covariate model. Panels (b) and (d) show the relations between sandeel presence and the covariates in the full model.

3.3.4.2 Modelling Sandeel Density

A Negative Binomial distribution was selected to model the count component of the zero-inflated model, based on the model with the minimum BIC, and a significant value for the scale parameter (θ) associated with the Negative Binomial distribution (p-value: < 0.0001). The same set of environmental covariates were selected by the BIC to predict sandeel presence in the zero mass component of the zero-inflated model as the binomial model; Time of day, temperature, pre/post closure, RoxAnn e1 and RoxAnn e2. In addition, time of day and year were selected to model the count component of the zero-inflated model.

The branch and bound method selected four local Gaussian spline (LGS) basis functions to model the remaining spatial variability found in the residuals of the environmental covariate model. The spatial smooth was included in the count component of the zero-inflated model, in order to evaluate the branch and bounds' potential for improving the estimation of the density surface. This smooth consisted of four basis functions, as well as East and North, and was found to be a highly significant predictor of sandeel density (Table 3.3). The BIC decreased from 18191 to 18132 when the spatial smooth chosen by the branch and bound was included in the model, and the pseudo R^2 value increased from 0.032 to 0.055. However, it is apparent from the pseudo R^2 value and Figure 3.8 that some observations are severely under-estimated.

There was negligible change in the parameter estimates (and associated p-values) for the zero mass component of the model when the spatial smooth was included. Unlike the binomial model, the RoxAnn e2 covariate was still significant at the 5% level (Table 3.3).

Table 3.3: Zero Inflated Model model selection results. The degrees of freedom (*df*) used for each covariate are indicated in the second column. The significance of the relation with each covariate is indicated by the *p*-values, for both the environmental covariate (EC) model and the full model. The *p*-values for the linear terms are taken from the model summary while those for the smooths are calculated using likelihood ratio tests. The full model was fitted using local Gaussian splines with an r value of 15

Zero Mass Component	*df*	*p*-value (EC model)	*p*-value (Full model)
s(Time of Day)	4	< 2.2e-16***	< 2.2e-16***
Temperature	1	2.66e-11 ***	2.70e-11 ***
Pre/Post Closure	1	0.001113 **	0.001025 **
s(RoxAnn e1)	4	2.635e-08***	2.635e-08***
RoxAnn e2	1	0.011690 *	0.011457 *
Neg Bin Component	***df***	**P-value (EC model)**	**P-value (Full model)**
s(Time of Day)	4	1.488e-08***	0.004426 **
Year	5	< 2.2e-16***	< 2.2e-16***
s(North, East)	6		< 2.2e-16***

There were some changes, in the the time of day covariate for the negative binomial count component of the zero-inflated model. The *p*-value associated with time of day increased from < 0.0001 to 0.0044 (Table 3.3) and the BIC-based selection would no longer retain this covariate. The year covariate still remained highly significant (Table 3.3). Interestingly, the shape of the predicted smooth across time of day appears to have only changed slightly (Fig. 3.9(a)). More pronounced changes can be seen in the predicted differences between years. The biggest difference is associated with the 1998 parameter estimate (Fig. 3.9(b)). The second largest difference is seen in 2000, while other parameter estimates appear largely unchanged.

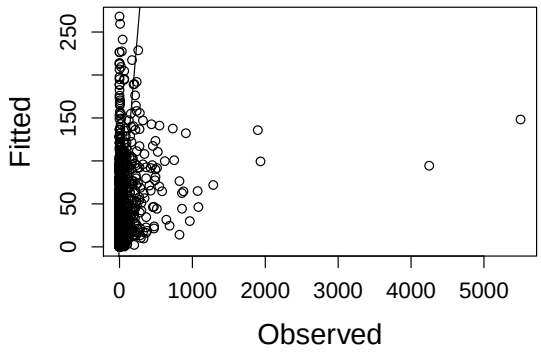

Figure 3.8: Fitted values plotted against the observed values for the full zero-inflated model, including both environmental covariates and the spatial smooth.

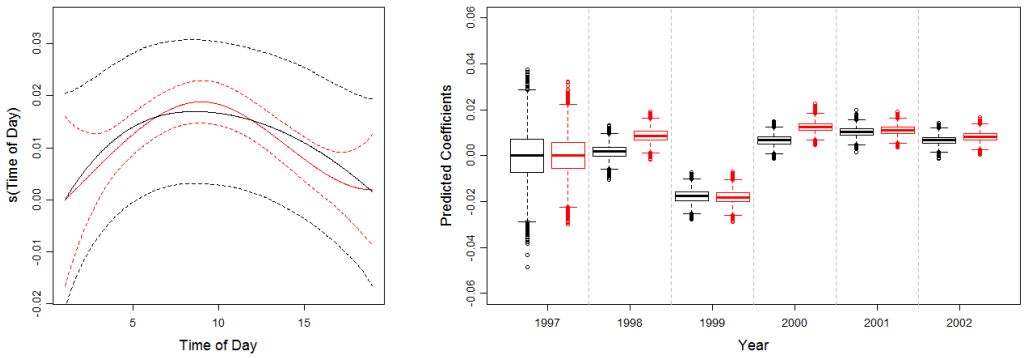

(a) Partial relation between sandeel density and time of day.

(b) Partial relation between sandeel density and year

Figure 3.9: Zero-inflated model partial relations with time of day and year. The relations predicted by the full model are shown in black and those predicted by the environmental covariate model are shown in red. These relations are based on parameters on the log scale.

3.3.4.3 Cross Validation

Presence/Absence Residual Model

The branch and bound method performed well to data unseen by the model. On average, the best (i.e. lowest) cross-validation based mean squared error (MSE), corresponds to the branch and bound model using local Gaussian splines (LGS) with an r value of between 6 to 8 (Table 3.4). However, the median MSE score might suggest that an r value of 4 would typically perform better.

The branch and bound outperforms MGCV with respect to the mean MSE scores in 6 of the 8 models presented. The intercept only model (i.e. no covariates) performs worst on average with respect to the mean MSE score, closely followed by the branch and bound LGS model with an r value of approximately 2.

Table 3.4: Cross validation results for the residuals from the presence/absence environmental covariate model. The branch and bound results are compared with an intercept only model and an MGVC GAM. r is the parameter associated with the local Gaussian splines and the standard deviation of the MSE score is given in parenthesis. All values are averages across the cross-validation folds.

Method	Spline	r	Pseudo R^2	MSE scores			AIC	BIC	Knots
				mean	median	(sd)			
Intercept	NA	NA	NA	1.016	1.03	(0.012)	6961	6972	NA
B&B	LGS	2	0.075	1.015	1.02	(0.127)	6793	6874	9.9
B&B	LGS	4	0.065	0.993	0.99	(0.124)	6817	6884	7.5
B&B	LGS	6	0.055	0.991	1.01	(0.129)	6841	6902	6.5
B&B	LGS	8	0.050	0.991	1.02	(0.117)	6853	6911	6.1
B&B	LGS	10	0.046	0.995	1.01	(0.121)	6862	6920	5.9
B&B	LGS	15	0.039	1.008	1.02	(0.127)	6881	6941	6.3
B&B	LGS	20	0.030	1.013	1.03	(0.127)	6901	6955	5.4
B&B	TPS	NA	0.057	0.999	1.02	(0.114)	6839	6916	9.3
MGCV	TPS	NA	0.023	1.010	1.05	(0.129)	6926	6997	10.3

The fit to the data (indicated by the pseudo R^2; Table 3.4) is higher for all branch and bound models compared with MGCV. Unsurprisingly, the branch and bound LGS model with the smallest r value trialled ($r = 2$) provides the best fit to the data, but overfits to the data and provides poor prediction of the data unseen by the model. No pseudo R^2 is available for the intercept only model since all predicted values are equal to the mean of the data and therefore cannot have a correlation with the observations.

The lowest AIC and BIC values are also associated with the branch and bound LGS model with an r value of 2, which is expected when you consider that r is not reflected in the model degrees of freedom. Under this scheme, there is no penalty incurred by the extra flexibility in the model surface that is possible when r decreases. One of the consequences of considering smaller r values however, is that more knots tend to be required which was also the case here (Table 3.4), and this is reflected in the model df. The branch and bound model with thin plate splines (TPS) and MGCV also use a larger number of degrees of freedom to fit the smooth than most of the LGS models and despite this incur a worse fit to the data and worse prediction of data unseen by the model.

Presence / Absence Model

The cross validation based MSE results for the presence/absence models, presented in Table 3.5, are similar to the results from the presence/absence residual models. The best MSE scores are again found using the branch and bound LGS model with r values of between 4 and 10. In addition, the branch and bound TPS model also provides a low mean MSE score based on cross-validation results.

Consistent with the presence/absence residual model results for the branch and bound LGS models, the smaller the value of r the better the fit for reasons stated

Table 3.5: Cross validation results for the the presence/absence models. The branch and bound full model results are compared with the environmental covariate model and an MGVC GAM. r is the parameter associated with the local Gaussian splines and the standard deviation of the MSE score is given in parenthesis. All values are averages across the cross-validation folds.

Method	Spline	r	Pseudo R^2	MSE scores			AIC	BIC	Knots
				mean	median	(sd)			
Env. Cov	NA	NA	0.164	0.173	0.18	(0.020)	2526	2596	NA
B&B	LGS	2	0.236	0.171	0.17	(0.020)	2379	2518	9.9
B&B	LGS	4	0.225	0.168	0.17	(0.022)	2401	2524	7.5
B&B	LGS	6	0.219	0.169	0.17	(0.022)	2415	2538	6.5
B&B	LGS	8	0.211	0.168	0.18	(0.020)	2431	2545	6.1
B&B	LGS	10	0.210	0.168	0.17	(0.020)	2434	2551	5.9
B&B	LGS	15	0.201	0.172	0.18	(0.022)	2453	2570	6.3
B&B	LGS	20	0.193	0.173	0.18	(0.021)	2474	2586	5.4
B&B	TPS	NA	0.219	0.169	0.18	(0.019)	2417	2552	9.3
MGCV	TPS	NA	0.184	0.172	0.18	(0.022)	2500	2625	9.6

earlier. In addition, the branch and bound TPS method improves the fit in comparison with penalized thin plate splines implemented in MGCV.

The following observations are also consistent with the results from the presence/ absence residual models: firstly, the lowest AIC and BIC values are associated with the branch and bound LGS model with an r value of 2, secondly the smaller the value or r for the LGS models, the higher the number of knots selected and, lastly, the branch and bound and MGCV also use a relatively high number of degrees of freedom to fit the spatial smooth in comparison with many of the LGS models.

Zero-Inflated Residual Model

In contrast to model results obtained thus far, the zero-inflated residual model results indicate that a more global smoothing approach to the model is appropriate. The best cross-validation based MSE scores are associated with the branch and bound LGS model with an r value of 15, the intercept only model and the globally

operating MGCV model (Table 3.6). There is however, great uncertainty about the CV-based MSE scores; the standard deviations are invariably larger than the MSE score themselves.

Table 3.6: Cross validation results for the residuals from the zero-inflated environmental covariate model. The branch and bound results are compared with an intercept only model and an MGVC GAM. r is the parameter associated with the local Gaussian splines and the standard deviation of the MSE score is given in parenthesis. All values are averages across the cross-validation folds.

Method	Spline	r	Pseudo R^2	MSE scores			AIC	BIC	Knots
				mean	median	(sd)			
Intercept	NA	NA	NA	274	99.3	(406)	31803	31815	NA
B&B	LGS	2	0.104	306	146.8	(389)	31550	31611	6.5
B&B	LGS	4	0.047	279	101.3	(399)	31700	31750	4.6
B&B	LGS	6	0.026	276	101.9	(402)	31747	31785	2.5
B&B	LGS	10	0.019	275	100.4	(402)	31766	31804	2.5
B&B	LGS	15	0.016	274	99.8	(402)	31774	31817	3.4
B&B	LGS	20	0.011	275	100.5	(402)	31784	31822	2.5
B&B	TPS	NA	0.045	282	106.2	(401)	31710	31774	7.0
MGCV	TPS	NA	0.011	274	99.9	(405)	31795	31860	9.1

Consistent with previous analyses, the smaller the value of r the better the fit and the branch and bound LGS based model with $r = 2$ provides the best fit to the data. In addition, the AIC and BIC also both select the branch and bound LGS model with the smallest r value trialled.

Zero-Inflated Model

The cross validation results for the zero-inflated models also suggest a more global smoothing approach is preferable (Table 3.7); on average, the best MSE score is found using the branch and bound LGS model with $r = 15$. However, the mean MSE score for the environmental covariate model (without a spatial surface) is also low, and the median MSE score indicates that typically this model would be expected to perform better. As in Table 3.6, there is great uncertainty about the CV-based MSE scores.

Table 3.7: Cross validation results for the the zero-inflated models. The branch and bound full model results are compared with the environmental covariate model. r is the parameter associated with the local Gaussian splines and the standard deviation of the MSE score is given in parenthesis. All values are averages across the cross-validation folds.

Method	Spline	r	Pseudo R^2	MSE scores			AIC	BIC	Knots
				mean	median	(sd)			
Env. Cov	NA	NA	0.033	276	100.0	(408)	16249	16382	NA
B&B	LGS	2	0.105	299	148.9	(393)	16141	16323	6.5
B&B	LGS	4	0.077	280	106.6	(397)	16142	16313	4.6
B&B	LGS	6	0.059	278	105.6	(404)	16167	16326	2.5
B&B	LGS	10	0.055	276	103.3	(404)	16166	16325	2.5
B&B	LGS	15	0.056	275	104.2	(402)	16167	16331	3.4
B&B	LGS	20	0.045	277	109.0	(400)	16182	16341	2.5
B&B	TPS	NA	0.070	279	108.4	(406)	16141	16325	7.0

Consistent with all the previous results for the branch and bound LGS models, the smaller the value of r the better the fit and the branch and bound LGS model with $r = 2$ provides the best fit to the data. However, in contrast, the BIC selects the branch and bound LGS model with $r = 4$. The lowest AIC value is now found in both the branch and bound LGS model with an r value of 2 and the branch and bound TPS model.

3.3.4.4 Prediction

We can use the predictions from our models to better assess the benefits of including the spatially adaptive smooth for assessing the effects of the sandeel fishery closure for both the presence and density of sandeels throughout the survey region.

Predictions were made from both the environmental covariate models and the full models detailed in Tables 3.1 and 3.3. Predictions for each year from 1997 to 2002 were computed using year-specific temperature values. The time of day was fixed across predictions as the median value in the dataset (9.125). In addition, predictions were made using the models of the residuals from both the presence / absence and zero-inflated environmental covariate models. These aided the assessment of the spatially adaptive smooth at picking up residual spatial variation once the effects of the environmental covariates had been modelled.

Presence/Absence Model

Both the environmental covariate models and the full models suggest an increase in probability of sandeel presence after the closure of the fishery (Figs. 3.10 and 3.11), with the highest predicted probabilities of sandeel presence seen in 2001 (Figs. 3.11(c) and 3.11(d)). For example, the environmental covariate model predicts an average probability of sandeel presence of between 0.15 and 0.28 (with 95% confidence) in 1997 - 1999 when the fishery was open, compared with between 0.23 and 0.40 (with 95% confidence) in 2000-2002 when the fishery was closed. These predictions are based on the median values of time of day, temperature and substrate measures.

Including the spatially adaptive smooth in the model, led to increases in the predictions for the probability of sandeel presence around two of the three main fishing

areas (Marr Bank and Berwick's Bank) in comparison with the environmental covariate model. The full model also predicts a higher probability of sandeel presence in the area around the Isle of May, where sandeels are thought to be an important component of the sea birds' diet. In addition, these differences appear to be supported by the data, particularly in the years following the closure of the sandeel fishery (Fig. 3.11).

Further differences between the environmental covariate and full models, which are also supported by the data, are higher predictions to the north east of the Marr Bank and lower predictions running roughly north-south between around -70 and -80 East nautical miles. However, although the data in to the north east of the Marr Bank indicate a large number of presences (mostly in the years post closure), the highest predictions in this area are out with the range of the data (extrapolation) and give some cause for concern.

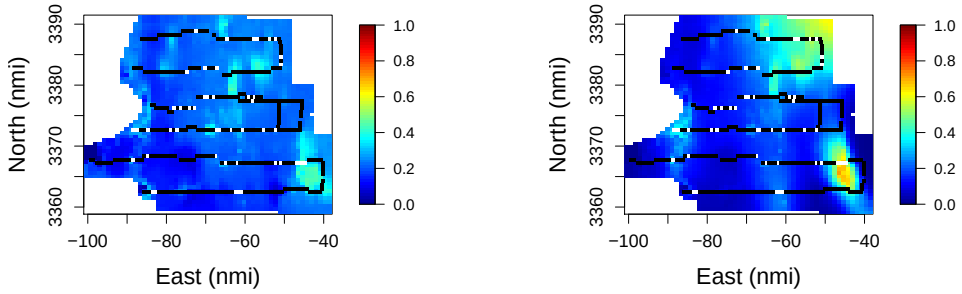

(a) 1997, environmental covariate model

(b) 1997, full model

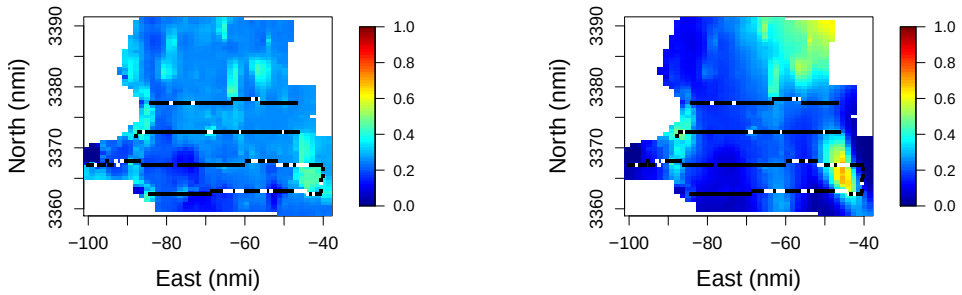

(c) 1998, environmental covariate model3

(d) 1998, full model

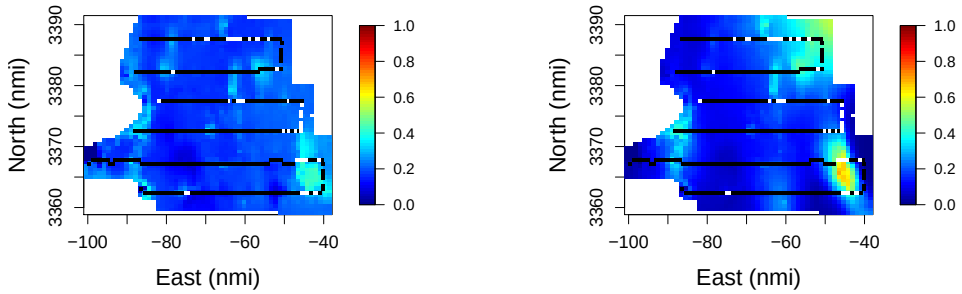

(e) 1999, environmental covariate model

(f) 1999, full model

Figure 3.10: Predicted probability of sandeel presence in the years when the sandeel fishery was active. The observations are shown by black squares for an absence and white squares for a presence.

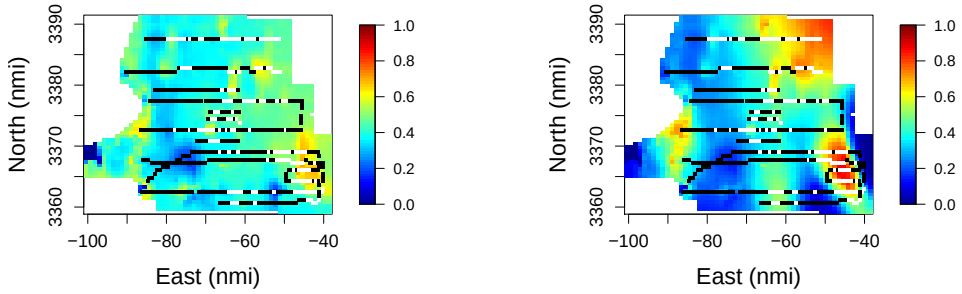

(a) 2000, environmental covariate model

(b) 2000, full model

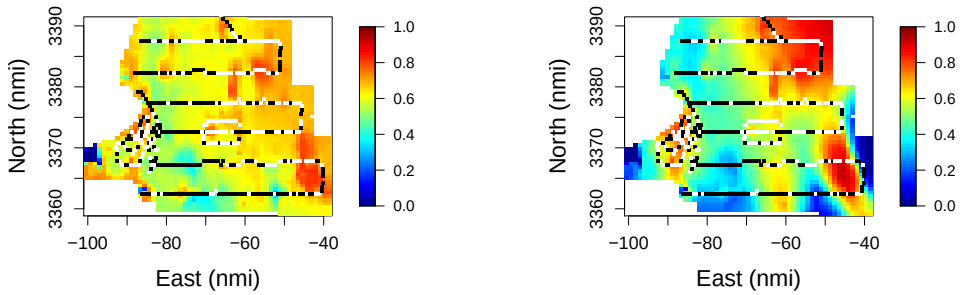

(c) 2001, environmental covariate model

(d) 2001, full model

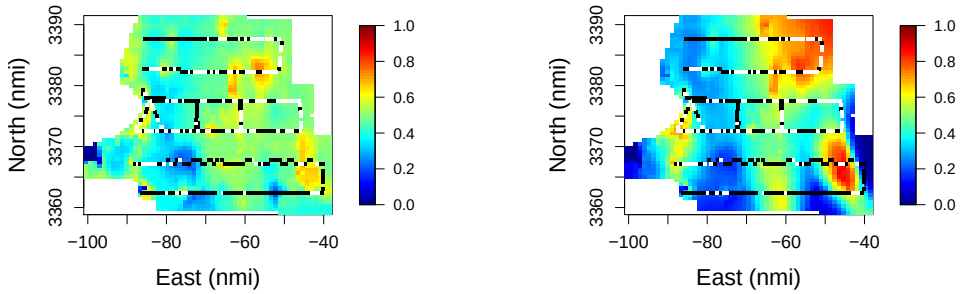

(e) 2002, environmental covariate model

(f) 2002, full model

Figure 3.11: Predicted probability of sandeel presence in the years after the sandeel fishery was closed. The observations are shown by black squares for an absence and white squares for a presence.

Presence / Absence Residuals Model

The model fitted to the residuals of the environmental covariate presence / absence model indicates evidence of remaining spatial pattern after all environmental covariates (as selected by the BIC) had been included in the model. The fit of the spatially adaptive smooth to the residuals gave a pseudo R^2 value of 0.049.

Similar patterns to those detected in the differences between the environmental covariate and full models are also seen in the model fitted to the residuals (Figure 3.12). These include higher predictions around the Marr Bank, Berwick's Bank and the Isle of May, as well as slightly lower predictions running north-south between around -70 and -80 East nautical miles.

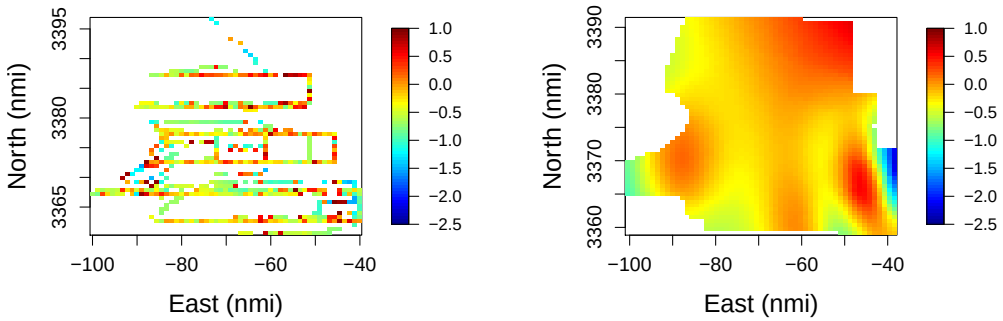

(a) Observed Residuals

(b) Predictions

Figure 3.12: Observed and predicted values for the residuals from the environmental covariate presence/absence model. Predictions were made from a model fitted using the branch and bound LGS method with an r values of 10.

In addition, the prediction plot of the residuals allows us to more easily assess the differences between the environmental covariate models and the full models. It becomes more apparent that there are lower predictions in the Firth of Forth and to the east of Berwick's Bank than would be made by the environmental covariate model

(Figs. 3.10-3.12). There is also another area of slightly higher predicted values to the west of Berwick's Bank. The latter two observations appear to be more strongly supported by the data than the former. In addition, there was very little survey effort within the Firth of Forth.

Zero-Inflated Model

The largest predicted values for the full models were much higher than the environmental covariate models. In addition, the scale on the plots from 2000-2002 (Fig. 3.14) have a much greater range of predicted densities than 1997 to 1999 (Fig. 3.13). The differences were so great for the predictions from the full models for 2000 to 2002 that only those values below 200 were plotted to allow comparisons with the environmental covariate predictions. The greatest predicted density for the full model was in 2001 at a value of almost 500.

In general, both the environmental covariate models and the full models suggest an increase in average sandeel density after the closure of the fishery, with 2001 having the highest sandeel densities predicted. The lowest sandeel densities were predicted in 1999. Naturally, those models that include spatial terms capture local changes in sandeel densities and more adequately represent the "hot-spots" in these areas.

The differences between the environmental covariate models and the full models are largely in the same locations as for the binomial models. The full models predict sometimes dramatically higher sandeel densities in the Berwick's Bank and Marr Bank regions. In addition, there are small increases observed in the Isle of May region and also lower sandeel densities running roughly North-South between around -70 and -80 East nautical miles.

There are some differences between the environmental covariate model and the full model which give some cause for concern. Although there are fewer problems of high values in the top right corner (as seen in the full presence / absence models), the predicted hot spot near the Berwick's Bank region is centred quite far outside the range of the data.

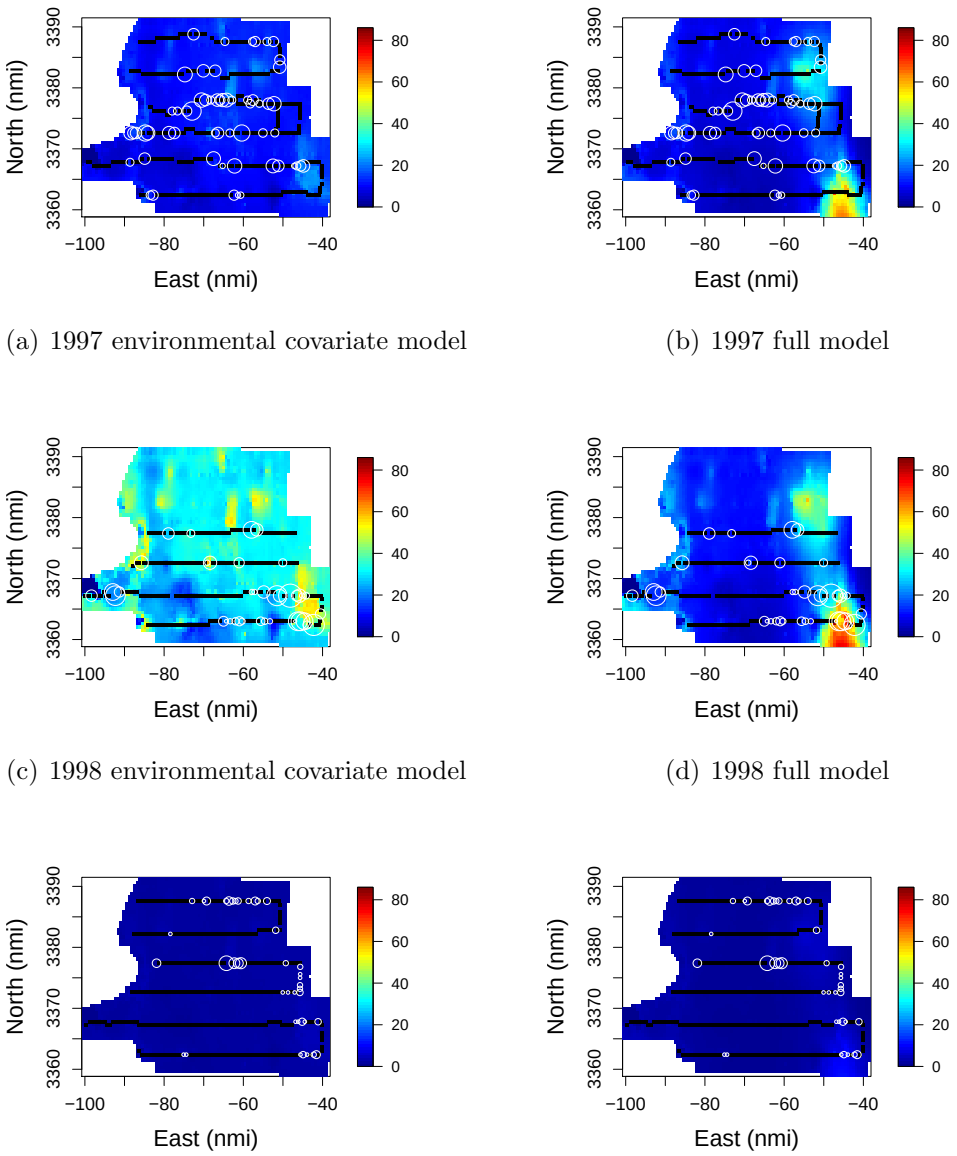

(a) 1997 environmental covariate model

(b) 1997 full model

(c) 1998 environmental covariate model

(d) 1998 full model

(e) 1999 environmental covariate model

(f) 1999 full model

Figure 3.13: Predicted density of sandeels in the years when the sandeel fishery was active. The observed sandeel densities are indicated by the white circles, the bigger the circle the higher the observed density. The zero observations are shown by black squares.

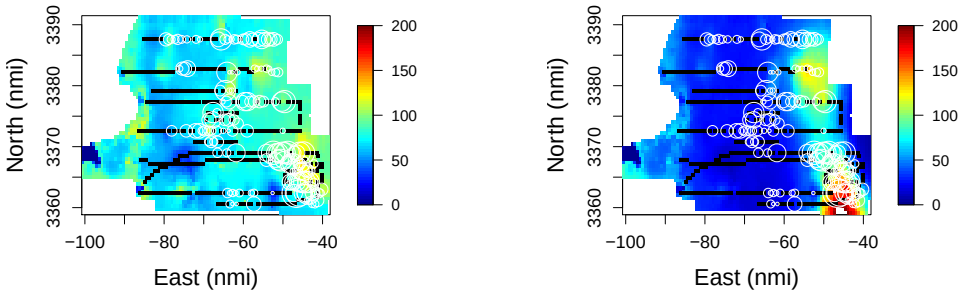

(a) 2000 environmental covariate model

(b) 2000 full model

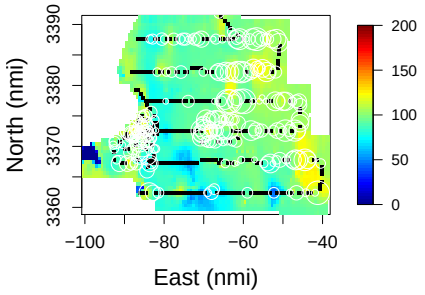

(c) 2001 environmental covariate model

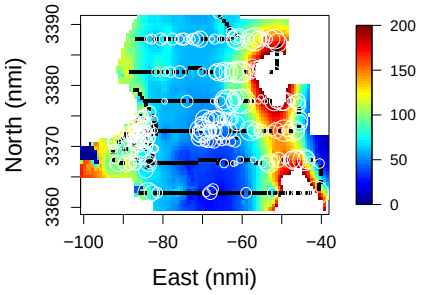

(d) 2001 full model

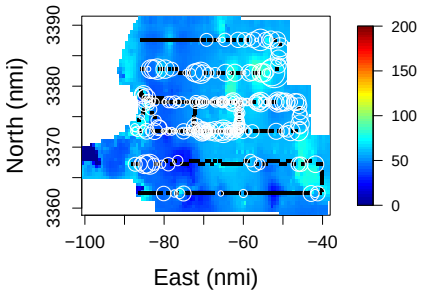

(e) 2002 environmental covariate model

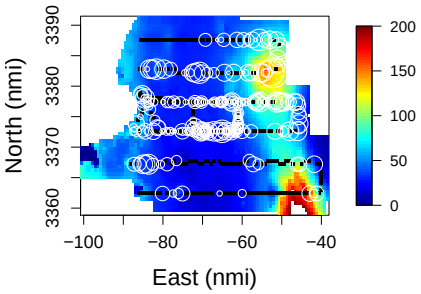

(f) 2002 full model

Figure 3.14: Predicted density of sandeels in the years after the sandeel fishery was closed. The observed sandeel densities are indicated by the white circles, the bigger the circle the higher the observed density. The zero observations are shown by black squares. Only predicted values <= 200 are displayed to allow comparison between the models.

Zero-Inflated Residual Model

Spatial variability was also found in the residuals from the zero-inflated model. However, the spatially adapted smooth fitted to these residuals only had a pseudo R^2 value of 0.017.

The residuals from the zero inflated model are characterised by a few very high values, indicating that the environmental covariate model vastly under-predicted these values (Figures 3.8 and 3.15(a)).

These predictions highlight the same patterns as the differences between the environmental covariate and full zero-inflated models. The most dramatic hotspots are indicated in the Berwick's Bank and Marr Bank regions (Figure 3.15(b)).

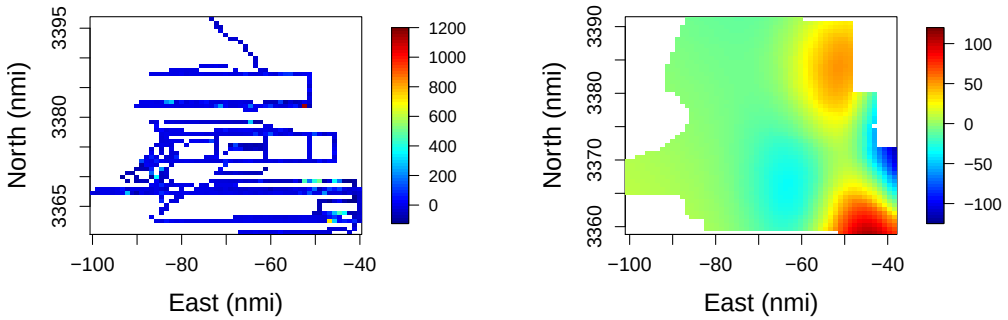

(a) Observed Residuals

(b) Predictions

Figure 3.15: Observed and predicted values for the residuals from the environmental covariate zero-inflated model. Predictions were made from a model fitted using the branch and bound LGS method with an r values of 15.

3.3.5 Discussion

Understanding the distribution of the sandeel stocks and how it changes in relation to fishing pressure is a necessary tool in the sustainable management of this valuable resource. This study has demonstrated the potential of a novel spatially adaptive smoothing method both within a binomial generalised additive model to improve the prediction of sandeel presence, as well as within a zero-inflated model.

This study also provides an assessment of the implications of including a spatial smooth as well as environmental covariates in a single model. For the first time a three stage modelling technique was investigated, and for the binomial data, was also compared with including a more flexible thin plate spline using MGCV.

3.3.5.1 Three Stage Modelling Approach

Presence/Absence Model

The results from the three stage modelling approach for including the spatially adaptive smooth were largely comparable to including the MGCV penalised thin plate regression splines, for this dataset. In fact, if anything the MGCV spatial smooth had a slightly lower impact on model interpretation than the branch and bound method. The differences in parameter estimates in Table 2 between the environmental covariate model and MGCV are smaller than those with the branch and bound model. This is an unexpected result as the surface fitted by MGCV had the potential to be more flexible than that of the branch and bound method. While the MGCV smooth could take on any shape within the full model, the branch and bound smooth should have been more restricted due to the specific choice of knot locations based on the spatial variability remaining in the residuals once the environmental covariates had been accounted for.

Overall the impact on model interpretation from including the spatial smooth was not extreme; only the interpretation of the significance of the RoxAnn e2 covariate changed. However, the shape of the predicted relations did not change substantially. The covariate RoxAnn e2 which was no longer significant at the 5% level after the inclusion of the smooth had a p-value of 0.011 prior to its inclusion and was not therefore highly significant initially. As is generally the case with the inclusion of any additional covariates, the variability of the parameter estimates for the existing covariates increases. This could explain some of the reduction in significance. In addition, because RoxAnn e2 was only significant at the 5% level it is possible that it was acting as a proxy for another unmeasured covariate.

Zero-Inflated Model

Small changes in model interpretation were also observed when the spatially adaptive smooth was included in the zero-inflated model. However, while time of day would no longer be retained in the model as chosen by the BIC, this covariate was still significant at the 0.5% level consistent with the environmental covariate model. Perhaps this gives us concern that the BIC may not provide a good model selection criterion for negative binomial zero-inflated models.

3.3.5.2 Selecting the Appropriate Spatial Smooth

The AIC and the BIC are two commonly used criteria for model selection (Johnson and Omland, 2004; Kuha, 2004). However, although the best subset of knots was selected within the iterative branch and bound using the BIC, neither this nor the AIC statistic appear to be useful when assessing which spatial smooth to use. Both the AIC and the BIC almost always select the branch and bound LGS method with an r value of 2, however, the MSE scores suggested that these models over-fit to the data. It is likely that this happens as the data are very peaked and have a few very

high values. The LGS models with smaller r values allow a much better fit to these data points which in turn dramatically reduces the loglikelihood value associated with the model. Although there are generally more knot positions (and therefore more parameter) selected for these models the loglikelihood is improved sufficiently that the AIC and BIC almost always select these models.

3.3.5.3 Predicting Sandeel Presence and Density

Models of the Residuals

The fitting of a spatial smooth to the residuals from an environmental covariate model is a valuable exercise for improving the understanding of the distribution of the species of interest, and potentially in the design of future surveys. These models tell us where there is a higher probability of presence or density of sandeels observed than was predicted based on the environmental covariate models. It is possible that such spatial patterns may point to additional previously unmeasured covariates which could then be recorded in subsequent surveys and included used to improve model predictive power. Alternatively, it is possible that autocorrelation is down to some other spatial process such as extinction, speciation, dispersal etc. Identifying the source of such spatial correlation, can also be a useful tool for management purposes.

Sandeel Presence

A better fit to the data, as well as improved prediction power to unseen data, was achieved by including the branch and bound spatial smooth in the presence / absence model. Cox (2008) also found that the branch and bound performed well for modelling missing pings in acoustic krill data, particularly for complex highly variable surfaces.

The predictions from these models also highlighted areas which are known fishing ground for sandeels. This suggests that these are locations with higher levels of

sandeel presence which were not fully predicted based on the environmental covariates alone.

These differences in predicted sandeel presence between the environmental covariate model and the full model provide insights which may aid the management of the sandeel stocks in this region. Pedersen et al. (1999) discussed the importance of choosing the correct division of sandeel stocks for management purposes as there is evidence that sandeels exist in smaller self-contained stocks. The full models predict a lower probability of sandeel presence running north-south between around -70 and -80 East nautical miles. This suggests that there is some geographic separation between the sandeels around the Isle of May region, which are an important food source for the seabirds, and those around two of the main fishery areas, the Berwick's Bank and Marr Bank. This separation is less distinctive in the years following the closure of the sandeel fishery which may also support the theory of Wright et al. (2000), that the distribution of sandeels expand and contract around their preferred habitat. However, the full models fitted in this chapter assume the same form of spatial smooth across years. To investigate these hypotheses further, models should be fitted to subsets of the data, either pre and post closure or every year individually.

Sandeel Density

Including the spatial smooth in the zero inflated models did not improve the predictive power over the environmental covariate model. The reason for this is likely due to the extreme right skewed distribution of the response (sandeel densities). Sandeels form aggregations in the water column and therefore may only appear in one or two adjacent segments of a transect. In addition, it appears that some of these aggregations were of particularly high density (Figure 3.8). This was also exhibited in the low pseudo R^2 values. As these aggregations only appear in very few of the consecutive

data points, the cross validation MSE scores are very sensitive to which data points are removed. Therefore, it is likely that the spatial smooths will overfit to the data and have worse MSE scores than the environmental covariate model.

However, the differences generated by the inclusion of the spatial smooth appear to provide genuine improvements in the prediction of where high densities of sandeels occur. Similar to the presence / absence results, the enhanced hotspots (by the inclusion of the spatial smooth) largely coincide with the sandbanks where the fisheries operate. Similar implications for management are indicated to those found for the presence / absence models; a potential division in the sandeel stocks around the Isle of May and those in the Berwick's Bank and Marr Bank region and potentially expanding and contracting distributions. Again further investigations specific to subsets of this dataset would be advised, either by pre and post closure or individual year.

3.3.5.4 Comparison of Splines

The LGS models provide a better fit as well as improved predictive power for the presence / absence models. In contrast, due to the severly right skewed distribution of the sandeel density variate, a more global smooth was preferred for the zero-inflated models.

Although it was hoped that the LGS models would provide more stable results at the edges of the survey region and not predict high values just outside the range of the data (as can be the case with TPS), this was not found to be the case. In both the presence / absence models and the zero-inflated models the LGS models predicted hotspots which were centered out with the range of the data. This was because two knot positions were selected which were relatively close to one another as well as the edge of the survey region. The knot closer to the centre of the region had a negative parameter estimate while the knot close to the edge of the survey region

had a positive parameter estimate. The sum of these two basis functions, which must have provided a better fit to the observed data points, then resulted in the prediction of a peak whose centre lay outwith the spatial range of the data and was therefore unsupported by the data. If such predictions are believed to be unrealistic, it may be possible to lessen these effects by implementing a buffer zone around the edge of the data in which candidate knot locations cannot be selected. However, as this dataset is collected along transects, this would result in many candidate knot locations being excluded.

3.3.5.5 Concluding Remarks and Future Investigations

The inclusion of the LGS branch and bound smooth in the presence / absence model enhances the prediction of where sandeels are likely to occur. These results suggested that further investigations of the Marr Bank and Berwick's Bank regions may reveal additional information about habitat preference in sandeels. In addition, insight into the variations in sandeel presence pre and post closure may be found by fitting individual models to separate subsets of the data.

These investigations have not dealt with any remaining temporal or spatial autocorrelation. Although the zero-inflated models do not provide any options to account for such problems, it would be possible to fit a presence / absence model which allows for temporal autocorrelation using generalised estimating equations (Hardin and Hilbe, 2003).

3.4 Case Study 2: Harbour Porpoise distribution off the west coast of the UK

3.4.1 Introduction

In section 3.3 we investigated the potential of the branch and bound algorithm to select optimal knot locations in order to implement a spatially adaptive smooth to better model both the presence and density of sandeels in space. The sandeel data provided two main challenges including an excess of zeros and a response which changed on more local scale in some areas than others. In this case study, we investigate the same spatially adaptive smoothing technique for modelling the distribution of harbour porpoise (*Phocoena phocoena*). These data are different to the sandeel data in that their spatial variability is more uniform across the surface and the areas of high density are less peaked.

Harbour porpoises are the most common species of cetacean along the West coast of Scotland (Evans, 2003). In addition to all species of dolphin, porpoise and whale, harbour porpoise are named in the European protected species list in Schedule 2 of The Conservation of Habitats and Species Regulations 2010[1]. There have been a number of studies carried out to predict the abundance and / or distribution of these animals, for example SCANS /SCANS II (Burt et al., 2006), Embling et al. (2010), Booth (2010), as monitoring the abundance of and understanding the spatial extent of this species is vital in its conservation.

Increasing levels of anthropogenic noise in the marine environment is of concern along the west coast of the United Kingdom. The sources of such noise include military sonar activities (Parsons et al., 2000), seismic exploration (Stone and Tasker, 2006) and acoustic deterrent devices used by fish farms (Booth, 2010). The effects of

[1]http://www.legislation.gov.uk/uksi/2010/490/contents/made

acoustic pollution and disturbance can have a number of potential implications for marine mammals and as awareness of these risks increases, so does the pressure on the organisations responsible to implement better management practices.

One such response to these pressures is the Environmental Risk Management Capability (ERMC) (Mollett et al., 2009). The ERMC project provides a quantitative risk assessment system for assessing the impact of sonar on marine mammal hearing. One component of this system involves density maps which describe the spatial distribution of the various species of marine mammal. Carrying out operations in areas which are least used by marine mammals is one obvious technique to try to minimise the risk to the animals' hearing and disturbance effects. Such techniques are long standing and have been widely implemented in the form of marine reserves e.g. Larkin (1996); Myers and Worm (2005); Booth (2010).

However, the effectiveness of these marine reserves and the predictive power of risk assessment systems such as ERMC, are limited by the quality of the models describing the spatial density of the species of interest. Here we investigate the potential of the branch and bound smoothing algorithm to model the instantaneous distribution of harbour porpoise in comparison with the off-the-shelf standard penalised thin plate spline methods implemented in MGCV. Consistent with section 3.3, we compare the use of both local Gaussian spline basis functions as well as those of the thin plate spline. Such models of the instantaneous distribution of harbour porpoise may then be used to validate the beliefs of existing management measures and risk assessment systems such as ERMC about the spatial distributions of the species of interest.

3.4.2 Data

Line transect surveys were undertaken throughout European waters as part of the second Small Cetaceans in the European Atlantic and North Sea and Adjacent waters survey (SCANS II) (Burt et al., 2006). Here we consider the data collected in those strata to the west and north of the United Kingdom. The effort and sightings of harbour porpoise groups are displayed in Figure 3.20. Strata P, Q and T were surveyed using ships, while strata J, N1 and R were surveyed using light aircraft. Further information on data collection can be found in (Burt et al., 2006).

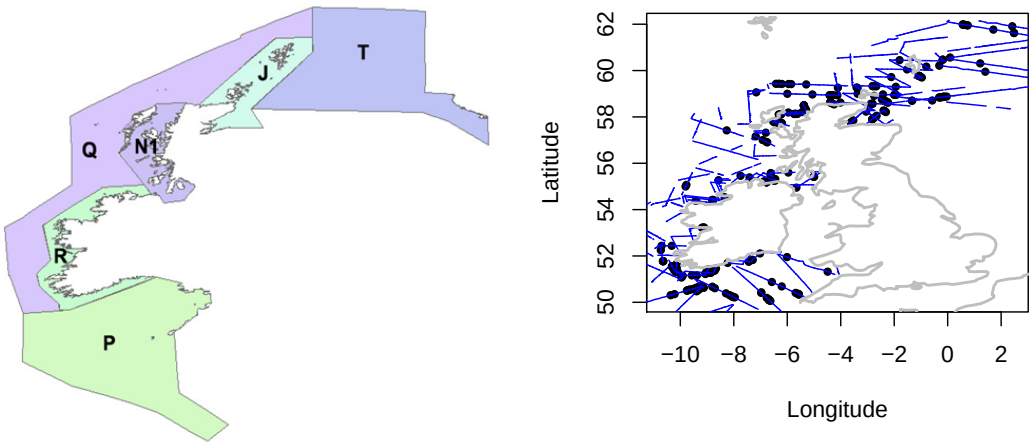

(a) SCANS-II strata

(b) Harbour porpoise sightings and surveyed transects.

Figure 3.16: (a)SCANS-II survey strata (P, Q, T, J, N1 and R) to the west and north of the UK. (b) Effort is displayed by the blue line while sightings of harbour porpoise schools are indicated by the black points, the coastline is given in grey.

Consistent with the scale of the density estimates in the ERMC system, the data were divided into half degree grid cells which were referenced by their centre coordinates. The latitude and longitude of each sighting was calculated based on the vessels

heading, along with the angle and radial distance associated with each sighting. The ships heading was calculated using the end points of each transect and assumed the ship followed a "rhumb line" or loxodrome; meaning that the ship would be following a constant bearing. The code used to perform these calculations, was adapted from equations obtained from http://www.movable-type.co.uk/scripts/LatLong.html. Once the locations of the sightings were established the number in each cell were counted and recorded.

Each of the zigs or zags of the survey are classified as distinct transects, and in distance sampling are assumed to be independent of one another. In the situation more than one transect intersected the half degree two or more separate entries were included in the data.

3.4.3 Modelling Methods

3.4.3.1 Estimation of Effective Effort

The effort associated with each cell was calculated as the length of transect within the cell multiplied by twice the effective strip half width. Line transect surveys differ from strip transects in that not all animals are counted within a certain distance of the transect. Instead, the distances to the sightings are recorded. A function is then fitted to these distances which describes the probability of observing an animal or group of animals given their distance from the transect. Assuming that the transects are randomly located throughout the survey regions means that on average the same number of animals are available for detection at each distance from the transect. The effective strip half width (μ) is then the distance at which as many animals are thought to be missed between distance 0 and μ as are observed beyond distance μ, Figure 3.17. This distance give us an effective strip (half) width, i.e. the equivalent width of a strip transect which would have generated the same number of sightings if we were to

have counted everything within this distance of the transect. In order to estimate the effective strip widths, detection functions were fitted to the perpendicular distances from both the ship and aerial surveys.

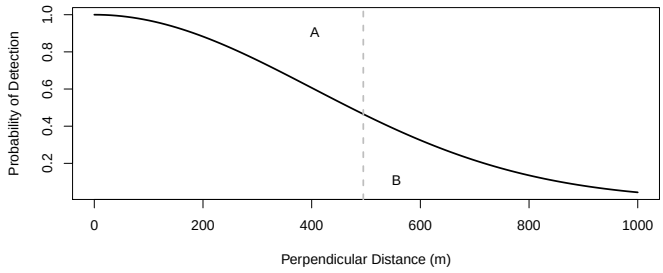

Figure 3.17: Half normal detection function, describing the probability of detecting an animal or group of animals given perpendicular distance from the transect. The dashed grey line represents the effective strip half width (μ) calculated as the distance at which area A is equal to area B.

3.4.3.2 General Model

A Poisson generalised additive model was used to model the instantaneous distribution of harbour porpoise. The counts of harbour porpoise sightings in each cell were offset by the effort in each cell (equation 3.8).

$$log\left(y_i/e_i\right) = \beta_0 + \sum_{k=1}^{K} \beta_k f\left(d_{ik}\right) + \epsilon_i \tag{3.8}$$

$$\text{where } \epsilon_i \sim Poisson(\lambda)$$

y_i and e_i denotes the count and effort associated with the i^{th} cell, respectively. β_0 represents the intercept term, $f(d_{ik})$ the value of the spline basis function for the i^{th} cell and k^{th} knot location and β_k represents the corresponding coefficient. ϵ_i represents the error term associated with the i^{th} cell.

3.4.3.3 Knot Selection

All data points were considered as candidate knot locations. Consistent with the sandeel case study, the candidate knot locations were divided into starting groups which were evenly spaced throughout the survey region, Figure 3.18. As these data were in a grid format is was possible to choose each subset based on a grid with a wider spacing. For example, if we consider the starting group of candidate knot locations coloured yellow then it can be seen that there are two candidate knot locations of other colours (belonging to other starting groups) between each yellow point in both the horizontal and vertical directions. The same iterative process as described in Section 3.3.3, (Figure 3.5) was used to select the final set of knot locations. Also consistent with the sandeel case study, the BIC statistic from the `regsubsets` routine was used to select the optimum size of subset at each iteration.

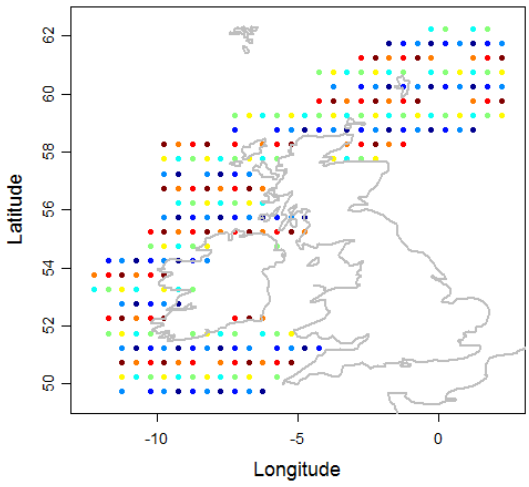

Figure 3.18: Candidate knot sets for the branch and bound method. The different colours represent the different starting sets.

3.4.3.4 Model Assessment

The harbour porpoise sightings were modelled using both local Gaussian splines and thin plate regression splines with the branch and bound method used for knot selection. In addition, penalised thin plate spline methods within MGCV were used for comparison. This model was restricted to use 15 degrees of freedom or less to model this surface.

The predictive power of the models to data unseen by the model was assessed using cross validation. Distance sampling assumes independence between transects, therefore, each validation test was carried out by removing an individual transect from the dataset. The model was then fitted to the remaining data and used to predict the data points which were removed before fitting. The ability of the model to predict this unseen data is measured using the mean square error (MSE) (equation 3.9), where n_v is the number of data points removed.

$$MSE = \frac{\sum(predicted - observed)^2}{n_v} \tag{3.9}$$

Consistent with the sandeel case study, the fit to the data was assessed using a pseudo R^2 score (equation 3.10) and additional statistics for model selection, including the AIC, BIC were recorded.

$$pseudoR^2 = (correlation(fitted, observed))^2 \tag{3.10}$$

The data used to fit the final models were ordered through time and the residuals were checked for temporal autocorrelation using both a Wald-Wolfowitz (`runs.test`) in the `lawstat` library, as well as the `acf` function in R.

3.4.4 Results

3.4.4.1 Effective Strip Half-Widths

The effective strip (half) widths were assumed to be 400 m and 190 m for the shipboard and aerial surveys, respectively. The detection functions used to estimate these values are provided in Figure 3.19.

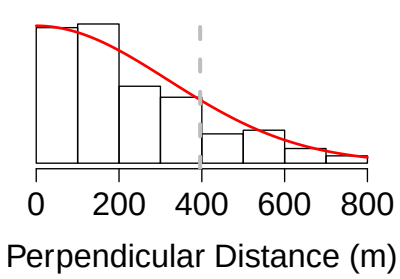

(a) Aerial sightings

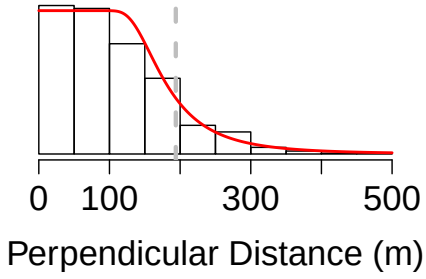

(b) Ship sightings

Figure 3.19: Detection functions fitted to the perpendicular sighting distances from the SCANS II surveys. The effective strip (half) widths are indicated by vertical the dashed line. (a) Half normal detection function fitted to the shipboard sightings from strata T, Q and P. (b) Hazard rate detection fitted to all aerial sightings.

3.4.4.2 Cross Validation

On average, the thin plate spline (TPS) basis functions provide better predictive power to unseen data as well as a better fit to the observed data than the local Gaussian splines (LGS). Not only are the cross validation MSE scores lower on average for the TPS models than the LGS models, the mean pseudo R^2 values are also higher (Table 3.8).

There is no clear choice of modelling technique when comparing the MGCV method and the branch and bound method for choosing the spatial smoothing fitted using

Table 3.8: Harbour porpoise cross validation results for the branch and bound method with various splines compared with MGCV.

Method	Spline	r	Pseudo R^2	MSE scores			AIC	BIC	Knots
				mean	median	(sd)			
B&B	LGS	1	0.081	6.52	0.97	(15.22)	1101	1135	5.0
B&B	LGS	2	0.097	7.14	0.81	(15.94)	1066	1101	5.1
B&B	LGS	3	0.060	6.99	1.18	(15.40)	1068	1101	4.8
B&B	LGS	4	0.025	7.15	1.20	(15.74)	1147	1168	2.0
B&B	LGS	5	0.026	7.08	1.39	(15.58)	1145	1167	2.4
B&B	TPS	NA	0.164	6.38	0.86	(14.69)	958	992	5.0
MGCV	TPS	NA	0.172	6.23	1.02	(14.89)	951	1013	13.8

thin plate splines. However, the results do appear to slightly favour MGCV. The pseudo R^2 values indicate that MGCV provides a better fit to the data than the branch and bound on average. In addition, while the median MSE scores indicate that we would typically expect the branch and bound to do slightly better, the mean MSE scores indicate that on average we would expect MGCV to provide the better MSE score. Also, when the MSE scores are compared in a pair wise manner across these two methods MGCV has the lower MSE score 51% of the time. However, comparing the worst MSE scores (98.04 and 86.57 for MGCV and the branch and bound, respectively) we see than in the worst case scenario the branch and bound outperforms MGCV.

Although the branch and bound LGS method with an r value of 2 also looks like it may provide promising predictive power based on the lowest median MSE score, MGCV has lower MSE scores 61% of the time, and the branch and bound TPS models have lower MSE scores 62% of the time.

Reassuringly, the AIC and BIC scores are roughly consistent with the MSE results. On average, the AIC would select the MGCV model and the BIC would select the

branch and bound model with TPS.

3.4.4.3 Prediction

The predictions from the MGCV model were compared with those from the branch and bound TPS model. Predictions were also obtained from the branch and bound model with LGS ($r = 2$), Figure 3.20.

Consistent with the cross validation results, the MGCV model provided a slightly better fit to the data than the branch and bound TPS model, Table 3.9. Interestingly, the branch and bound LGS model also gave a pseudo R^2 value of 0.17 , matching the fit of MGCV.

In contrast to the cross validation results, the AIC and the BIC would both select the branch and bound LGS model. Second choice of the AIC would be the MGCV model, while more consistent with the cross validation results, the BIC would select the branch and bound TPS model, Table 3.9.

Table 3.9: Comparison of final model statistics for the MGCV GAM as well as the branch and bound models with TPS and LGS with an r value of 2.

Method	Spline	r	pseudo R^2	AIC	BIC	*df*
B&B	LGS	2	0.17	954	987	5
B&B	TPS	NA	0.16	963	996	5
MGCV	TPS	NA	0.17	958	1021	13.83

Although the autocorrelation plots didn't suggest any problems of remaining temporal correlation in the residuals, the Wald-Wolfowitz tests indicated significant correlation for the MGCV and branch and bound LGS models (p-values 0.001 and <0.001, respectively). In contrast the Wald-Wolfowitz test on the residuals from the branch and bound TPS model showed no significant temporal autocorrelation (p-value 0.789).

It is difficult to assess the accuracy of the predicted values. While all models predict hot spots (which are also observed in the data) to the west of Scotland and the south west of Ireland they differ slightly in their predictions of the extent of these hot spots, Figure 3.20. Generally, the two branch and bound models are more consistent with each other in their predictions than the MGCV model. The branch and bound models appear to better predict the high data points to the south west of Scotland / north of Ireland than the MGCV model. However, the MGCV model appears to better predict the band of higher densities (which look to be apparent in the dataset) from the north Scotland up to Shetland and beyond. Finally, the MGCV model predicts a higher density of harbour porpoise further up the west coast of Ireland than the branch and bound models, however it is difficult to tell from the data if this is a more accurate representation of the data or not.

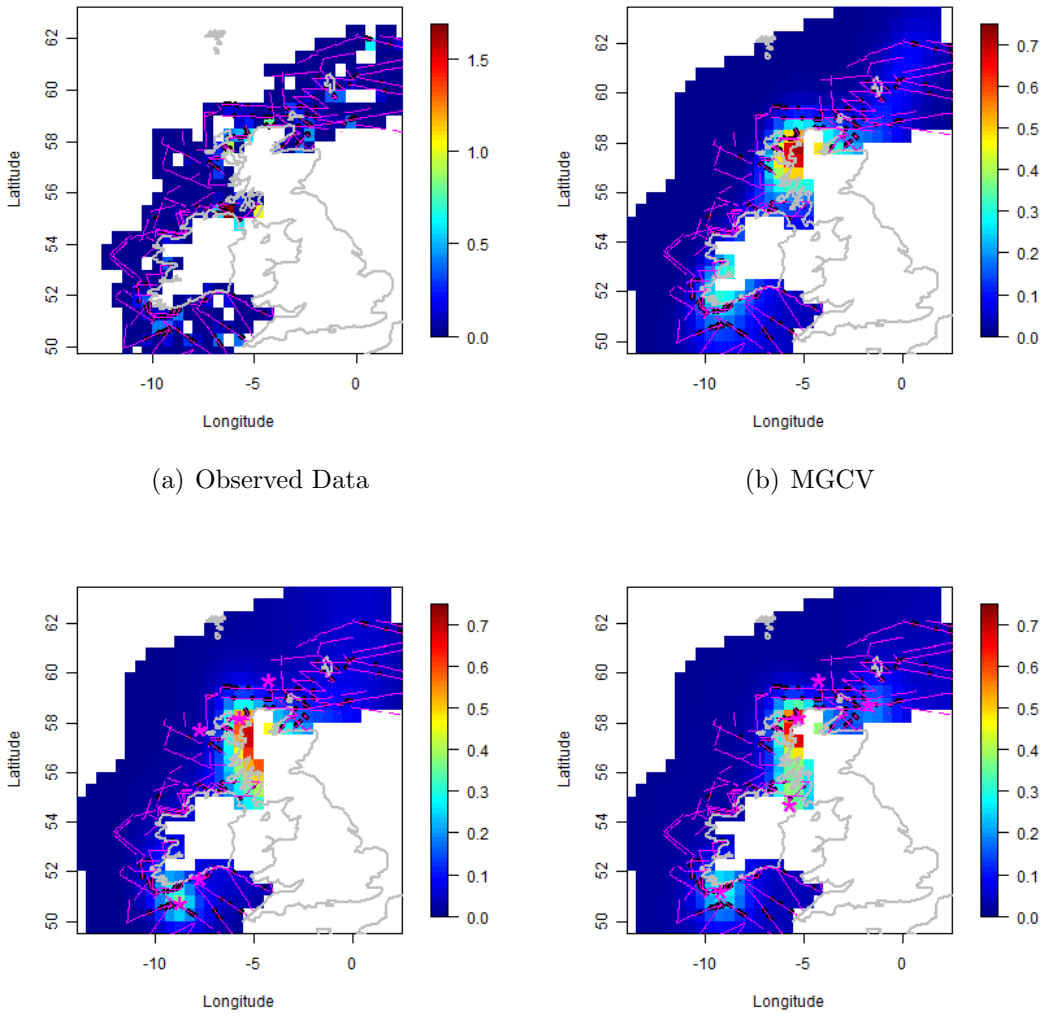

(a) Observed Data

(b) MGCV

(c) Branch and Bound TPS

(d) Branch and Bound LGS

Figure 3.20: Predicted number of harbour porpoise sightings per unit effort compared with the data. (a) Observed values: The average number of sightings per unit effort is indicated by the colour scale. The transects surveyed are show in pink and the sightings as black circles. (b) Predicted surface from the MGCV model using thin plate spline basis functions. (c) Predictions from the branch and bound model using thin plate spline basis functions. (d) Predictions from the branch and bound model using local Gaussian spline basis functions.

3.4.5 Discussion

In agreement with the findings of Cox (2008), the branch and bound method appears less effective in modelling more uniformly complex data. The distribution of harbour porpoise is more homogeneous in its variability and less peaked than the sandeel data.

However, while on average the MGCV GAM method appears to provide marginally better predictive power, the median value indicates that typically the branch and bound TPS model may be preferred and in the worst case scenario the branch and bound out performs MGCV. In addition, the almost 50:50 "win" rate when the MSE scores are compared pair-wise, make it very difficult to choose one model over the other.

3.4.5.1 Model Selection

In contrast to the sandeel case study, the cross validation results indicate that the AIC and BIC provided a reasonable criterion for the selection of the spatial smooth for the harbour porpoise data. This is thought to be related to higher levels of homogeneity in spatial variability across the surface. The LGS model with small r was no longer required to provide a much better fit (and therefore small loglikelihood). However, strangely the LGS model with an r value of 2 had the lowest AIC and BIC values when fitted to the entire dataset.

3.4.5.2 Accuracy of Predictions

The accuracy of the predictions from the different models is difficult to assess. Previous studies indicate that there is a high density of harbour porpoise to the west of Scotland (Booth, 2010). In addition, Booth (2010) found that the distribution of harbour porpoise appeared to be more northerly in 2005 (the year of this survey) than

in other years which may be more consistent with the MGCV predictions. However, the differing scales of this analysis with that of Booth (2010) make comparisons difficult.

3.4.5.3 Concluding remarks and Future Directions

Dependent on which modelling technique is implemented there may or may not be temporal autocorrelation in the residuals. However, the branch and bound spatially adaptive smoothing technique could easily be implemented within a generalised estimating equations framework to relax the assumption of independence.

Due to the difficulties associated with model selection in this instance a possible solution would be to average model predictions across models, in line with AIC or BIC weights (or similar) (e.g. Hjort and Claeskens (2003)). Alternatively, predictions from all models could be used to incorporate model uncertainty; this is an important consideration if management decisions are to be based on such results.

3.5 General Discussion and Conclusions

This chapter has investigated the potential of the branch and bound algorithm for implementing a spatially adaptive smooth. We have demonstrated the improved predictive power to unseen data achieved by including the branch and bound smooth in the presence / absence models. We have also demonstrated how this technique can be used within a zero-inflated framework. In addition, we have shown that this method performs similarly to the global off-the-shelf method (MGCV) for modelling the more spatially homogeneous distribution of harbour porpoise.

3.5.1 Model Selection Criteria

BIC versus Cp mallows

In contrast to the study of Cox (2008) these investigations found that the Cp mallows (similar to the AIC statistic) led to severe problems of over-fitting to the data. The results of the cross validation indicated that the predictive power of the resulting models for the branch and bound method, were vastly improved when the BIC was used to select the size of the appropriate subset of knots.

AIC and BIC in Smoothing Selection

These results suggest that the AIC or BIC may be acceptable for the purposes of spline selection for more uniformly variable surfaces but are less effective at selecting splines for peaked data and heterogeneously variable data. There was reassuring agreement between the AIC, BIC and cross validation MSE scores for the harbour porpoise data but this was not the case for the sandeel data.

3.5.2 Comparison of Splines

These investigations demonstrated that the choice of spline can affect the fit and thus the predictive power of a model. While the local Gaussian splines are more

suited to highly peaked and spatially variable data, the thin plate splines appear to be preferred when modelling more uniformly variable data.

In addition, these investigations found that the local Gaussian and thin plate splines splines suffer similarly from predicting high values just outside the range of the data. Both the full sandeel presence / absence model as well as the full zero-inflated model predicted peaks out with the spatial range of the data.

3.5.3 Future Directions

Developments in computing power and the efficiency of algorithms will make this method more practical in the future. Given that there are no statistics which can be used to consistently and easily choose between splines, cross validation is recommended as the most appropriate method for smooth selection. However, this branch and bound can take a large amount of time to perform knot selection ($\sim$10 minutes) which makes cross validation time consuming.

It is likely that in the near future the number of potential knot locations which the branch and bound algorithm will be able to practically consider will no longer be restricted to 30. Such developments should lead to improved results from this method. Although the method presented in this chapter attempts to work around this restriction by choosing a large number of potential knot locations and implementing an iterative algorithm, we do not achieve the "all possible subsets" ideal. However, the more knots the branch and bound algorithm can consider, the closer we will get to this ideal

Given the successful results of the LGS models, it would be interesting to investigate other locally acting basis functions. In addition, fitting LGS surfaces using difference r values within the same model may improve the ability of the model to capture

variability at different scales. However, this latter idea would lead to a very large number of potential basis functions for the branch and bound to consider.

Generally when you include interactions within a model you should also include the main (linear) terms. Based on this idea, linear terms for both north and east were included in the model as well as the spatial smooth. The effects of doing this were not investigated but may have contributed to the problems observed in the top right corner of the presence / absence predictions for the sandeel data. It would be worth investigating the effects of excluding these main terms from the models.

In this chapter, we assessed the ability of the models to predict to data unseen via cross validation techniques. However, the true assessment of any modelling technique is with reference to the true underlying function. In order to assess this we require an appropriate suite of benchmark functions so that methods such as these can be compared across these standard functions. As yet there are a very limited number of such functions available and none of which are realistic enough to require spatially adaptive modelling techniques to best capture the underlying function. This is an area which requires urgent development.

Bibliography

S. Arnott, G. Ruxton, and E. Poloczanska. Stochastic dynamic population model of North Sea sandeels, and its application to precautionary management procedures. *Marine Ecology Progress Series*, 235:223–234, 2002.

M. Austin. Spatial prediction of species distribution: an interface between ecological theory and statistical modelling. *Ecological Modelling*, 157:101–118, 2002.

M. Austin and J. Meyers. Current approaches to modelling the environmental niche of eucalypts: implications for management of forest biodiversity. *Forest Ecology and Management*, 85:95–106, 1996.

J. D. Baker. Variation in the relationship between offspring size and survival provides insight into causes of mortality in Hawaiian monk seals. *Endangered Species Research*, 5:55–64, 2008.

S. Barry and A. Welsh. Generalized additive modelling and zero inflated count data. *Ecological Modelling*, 157:179–188, 2002.

C. Booth. *Variation in habitat preference and distribution of harbour porpoises west of Scotland*. PhD thesis, School of Biology, University of St. Andrews, 2010.

P. Bromley, T. Watson, and J. Hislop. Diel feeding partterns and the development of food webs in pelagic 0-group cod (*Gadus morhua* L.), haddock (*Melanogrammus aeglefinus* L.), whiting (*Merlangius merlangus* L.), saithe (*Pollachius virens* L.),

and Norway pout (*Trisopterus esmarkii* Nilsson) in the northern North Sea. *ICES Journal of Marine Science*, 54:846–853, 1997.

M. Buhmann. Radial basis functions. *Acta Numerica*, pages 1–38, 2000.

M. Burt, D. Borchers, F. Samarra, and Others. Preliminary abundance estimates from SCANS II. To be published 2007, 2006.

M. Cox. *Acoustic and ecological investigations into predator-prey interactions between Antarctic krill (*Euphausia superba*) and seal and bird predators.* PhD thesis, Biology / Mathematics and Statistics, 2008.

J. Cragg. Some statistical models for limited dependent variables with application to the demand for durable goods. *Econometrica*, 39:829–844, 1971.

C. Embling, P. Gillibrand, J. Gordon, J. Shrimpton, P. Stevick, and P. Hammond. Using habitat models to identify suitable sites for marine protected areas for harbour porpoises (*Phocena phocena*). *Biological Conservation*, 143:267–279, 2010.

G. Engelhard, J. van der Kooij, E. Bell, J. Pinnegar, J. Blanchard, D. Mackinson, and D. Righton. Fishing mortality versus natural predation on diurnally migrating sandeels *Ammodytes marinus*. *Marine Ecology Progress Series*, 369:213–227, 2008.

P. Evans. Shipping as a possible source of disturbance to cetaceans. Technical report, ASCOBANS, 2003.

J. Fox. *Applied Regression, Linear Models, and Related Methods.* Sage, 1997.

M. Frederiksen, M. Edwards, A. Richardson, N. Halliday, and S. Wanless. From plankton to top predators: bottom-up control of a marine food web across four trophic levels. *Journal of Animal Ecology*, 75:1259–1268, 2006.

R. Furness. Management implications of interactions between fisheries and sandeel-dependent seabirds and seals in the North Sea. *ICES Journal of Marine Science*, 59:261–269, 2002.

G. Furnival and R. Wilson. Regressions by leaps and bounds. *Technometrics*, 16: 499–511, 1974.

M. Graham. Confronting multicollinearity in ecological multiple regression. *Ecology*, 84:2809–2815, 2003.

J. Granadeiro, J. Andrade, and J. Palmeirim. Modelling the distribution of shorebirds in estuarine areas using generalised additive models. *Journal of Sea Research*, 52: 227–240, 2004.

S. Greenstreet, J. McMillan, and E. Armstrong. Seasonal variation in the importance of pelagic fish in the diet of piscivorous fish in the Moray Firth, NE Scotland: a response to variation in prey abundance? *ICES Journal of Marine Science*, 55: 121–133, 1998.

S. Greenstreet, E. Armstrong, H. Mosegaard, H. Jensen, I. Gibb, H. Fraser, B. Scott, G. Holland, and J. Sharples. Variation in the abundance of sandeels *Ammodytes marinus* off southeast Scotland: an evaluation of area-closure fisheries management and stock abundance assessment methods. *Journal of Marine Science*, 63:1530–1550, 2006.

A. Guisan, J. Edwards, and T. Hastie. Generalized linear and generalized additive models in studies of species distributions: setting the scene. *Ecological Modelling*, 157:89–100, 2002.

P. Hammond and K. Grellier. Grey seal diet composition and prey consumption in

the North Sea. Technical report, Executive Summary project MF0319, SMRU, University of St. Andrews, 2006.

R. Harder and R. Desmarais. Interpolation using surface splines. *Journal of Aircraft 9*, 1972.

J. Hardin and J. Hilbe. *Generalized Estimating Equations*. Chapman and Hall/CRC, 2003.

T. Hastie and R. Tibshirani. *Generalized Additive Models*. Chapman and Hall, 1990.

T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer-Verlag, New York., 2001.

N. Hjort and G. Claeskens. Frequentist model average estimators. *J. Am. Stat. Assoc.*, 98:879–899, 2003.

ICES. Ices catch by species, area and year (1973-2008). Eurostat/ICES database on catch statistics - ICES 2007, 2009.

J. Johnson and K. Omland. Model selection in ecology and evolution. *TRENDS in Ecology and Evolution*, 19:101–108, 2004.

R. Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *International Joint Conference on Artificial Intelligence*, 1995.

J. Kuha. AIC and BIC comparisons of assumptions and performance. *Sociological Methods and Research*, 33:188–229, 2004.

D. Lambert. Zero-inflated poisson regression, with an application to defects in manufacturing. *Technometrics*, 34:1–14, 1992.

P. Larkin. Concepts and issues in marine ecosystem management. *Reviews in Fish Biology and Fisheries*, 6:139–164, 1996.

T. Lumley. Regression subset selection including exhaustive search, version 2.9. 2009.

E. Mammen and S. van de Geer. Locally adaptive regression splines. *The Annals of Statistics*, 25:387–413, 1997.

T. Martin, B. Wintle, J. Rhodes, P. Kuhnert, S. Field, S. Low-Choy, A. Tyre, and H. Possingham. Zero tolerance ecology: improving ecological inference by modelling the source of zero observations. *Ecology Letters*, 8:1235–1246, 2005.

J. Matthiopoulos, S. Smout, A. Winship, D. Thompson, I. L. Boyd, and J. Harwood. Getting beneath the surface of marine mammal - fisheries competition. *Mammal Review*, 38(2&3):167–188, 2008.

A. Miller. *Subset selection in regression*. Chapman and Hall, 2002.

A. Mollett, C. Schofield, I. Miller, J. Harwood, C. Harris, and C. Donovan. Environmental risk management capability: Advice on minimising the impact of both sonar and seismic offshore operations on marine mammals. Technical report, SPE Offshore Europe Oil & Gas Conference & Exhibition held in Aberdeen, UK, 8th to 11th September, 2009.

J. Mullahy. Specification and testing of some modified count data models. *Journal of Econometrics*, 33:341–365, 1986.

R. Myers and B. Worm. Extinction, survival or recovery of large predatory fishes. *Philosophical Transactions of the Royal Society*, 360:13–20, 2005.

S. Panigada, M. Zanardelli, M. MacKenzie, C. Donovan, F. Mélin, and P. Hammond. Modelling habitat preference for fin whales and striped dolphins in the Pelagos Sanctuary (western Mediterranean Sea) with physiographic and remote sensing variables. *Remote Sensing of Envrionment*, 112:3400–3412, 2008.

E. Parsons, I. Birks, J. Evans, P.G.H.and Gordon, J. Shrimpton, and S. Pooley. The possible impacts of military activity on cetaceans in west Scotland. *European Research on Cetaceans*, 14:185–190, 2000.

S. Pedersen, P. Lewy, and P. Wright. Assessments of the lesser sandeel (*Ammodytes marinus*) in the North Sea based on revised stock divisions. *Fisheries Research*, 41:221–241, 1999.

D. Pierce and D. Schafer. Residuals in generalised linear models. *J. Am. Stat. Assoc.*, 81:977–986, 1986.

J. Royle and D. Nychka. An algorithm for the construction of spatial coverage designs with implementation in SPLUS. *Computers & Geosciences*, 24:479–488, 1998.

D. Ruppert, M. Wand, and R. Carroll. *Semiparametric Regression*. Cambridge University Press, 2003.

S. Rushton, S. Ormerod, and G. Kerby. New paradigms for modelling species distributions? *Journal of Applied Ecology*, 41:193–200, 2004.

M. Santos, G. Pierce, and J. Learmonth. Variability in the diet of harbour porpoises (*Phocoena phocoena*) in Scottish waters 1992-2003. *Marine Mammal Science*, 20 (1):1–27, 2004.

G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2): 461–464, 1978.

D. Stasinopoulos and R. Rigby. Generalized additive models for location scale and shape (GAMLSS) in R. *Journal of Statistical Software*, 23, 2007. URL `http://www.jstatsoft.org/v23/i07/`.

C. Stone and M. Tasker. The effects of seismic airguns on cetaceans in UK waters. *J. Cetacean Res. Manage.*, 8:255–263, 2006.

W. N. Venables and B. D. Ripley. *Modern Applied Statistics with S. Fourth Edition*. Springer-Verlag, 2002.

C. Walker, M. Mackenzie, C. Donovan, and M. O'Sullivan. SALSA - a spatially adaptive local smoothing algorithm. *Journals of Statistical Computation and Simulation*, 00:1–13, 2010.

R. Williams, D. Lusseau, and P. Hammond. The role of social aggregations and protected areas in killer whale conservation: The mixed blessing of critical habitat. *Biological Conservation*, 142:709–719, 2009.

S. Wood. *Generalized Additive Models*. Chapman and Hall/CRC, 2006.

P. Wright, H. Jensen, and I. Tuck. The influence of sediment type on the distribution of the lesser sandeel, *Ammodytes marinus*. *Journal of Sea Research*, 44:243–256, 2000.

T. Yee. *VGAM: Vector Generalized Linear and Additive Models. R package version 0.7-7*, 2008. URL `http://CRAN.R-project.org/package=VGAM`.

A. Zeileis, C. Kleiber, and S. Jackman. Regression models for count data in R. *Journal of Statistical Software*, 27(8), 2008.

S. Zhou and X. Shen. Spatially adaptive regression splines and accurate knot selection schemes. *J. Am. Stat. Assoc.*, 96:247–259, 2001.

Chapter 4

Anthropogenic Noise and Marine Mammals: Assessing real-time monitoring

This chapter has been temporarily embargoed

Chapter 5

General Discussion

5.1 General Conclusions

This thesis has developed a number of statistical techniques whose application has important implications for the management of the marine environment. It has addressed two key concerns in the North Sea at present: the fishing and predation pressure on sandeels, a key prey species, and the management of acoustic noise in the marine environment. To do this we have utilised computationally intensive methods in the development of novel analyses and simulations. While Chapters 2 and 4 have developed techniques to deal with specific management issues which were not addressed in existing analyses or simulations, Chapter 3 developed a novel spatial modelling technique which is widely applicable in any 2-dimensional modelling context.

5.1.1 Management of North Sea Sandeels

Firstly, this thesis provided compelling evidence that a substantial proportion of the sandeels consumed by grey seals are *H. lanceolatus*, a species which is not commercially exploited, rather than the commercially important *A. marinus*. Previous analyses documented the suspected presence of *H. lanceolatus* in addition to *A. marinus* but did not attempt to quantify the biomass of each of these species of sandeel (Prime

and Hammond, 1990; Hammond et al., 1994a,b). The novel methods in Chapter 1 of this thesis attempt to estimate the consumed biomass of these species separately to allow seal predation specifically on *A. marinus*, one of the key prey species in the North Sea, to be included in multispecies fisheries models. Not only did this model attempt to identify two separate distributions of otoliths sizes thought to belong to each of these species using a mixture model, it also implemented a size dependent number correction factor (NCF). Although species specific NCF values have been used in previous hard part diet analyses, a NCF function dependent otolith size has never before been implemented within species. The inclusion of this was important as not only does it allow more accurate predictions of the total consumed biomass of sandeels but also what quantity may be *A. marinus*. A general trend of increasing probability of recovery with increasing otolith size has been observed both within and across species (Tollit et al., 1997; Bowen, 2000; Grellier and Hammond, 2006), therefore it is likely that more of the smaller otoliths belonging to *A. marinus* will be lost to digestion in comparison with the larger otoliths thought to belong to *H. lanceolatus*.

Chapter 1 also presented quantitative results regarding a number of sources of bias in the estimation of the total biomass of sandeels consumed by grey seals. Accounting for the size dependent NCF, a consideration which was not included in previous analyses, suggested that the total biomass of sandeels may have been previously underestimated. In addition, we also considered the implications of using the wrong relation for estimating sandeel mass from otolith dimension. In previous analyses of the North Sea data, the total biomass was calculated assuming a fish mass to otolith dimension relation based solely on *A. marinus*, we found that accounting for the presence of *H. lanceolatus* and using species specific relations would most likely lead to an increase in the estimated total biomass consumed. However, it is likely

that the variability in relations within species would have even more influence on the estimates of consumed biomass than the variability between species, similar to findings by Pierce et al. (2007) and Froese (2006).

In addition to investigating the predatory pressures of grey seals on sandeels, particularly *A. marinus*, in Chapter 3 we also investigated the effects of the closure of the sandeel fishery on the distribution of sandeels in the Firth of Forth and Wee Bankie region off the east coast of Scotland. To do this we developed a novel spatial modelling technique to investigate both the spatial distribution of sandeel presence and density. Unlike many existing models our technique allows the 2-dimensional spatial smooth to vary more in some regions of the surface than others. Although a few other techniques apply the same idea (Mammen and van de Geer, 1997; Zhou and Shen, 2001), none, as far as we know, have been implemented within a zero-inflated model. In addition, our technique is easy to implement in any model which allows the inclusion of linear terms. In contrast, the other methods often require more complex methods and are less generic across different types of model (Mammen and van de Geer, 1997; Zhou and Shen, 2001).

Applying the spatial modelling technique developed in Chapter 3, improved both the fit and the predictive power of the model describing sandeel presence in comparison with both the commonly used `MGCV` package and also the model where no spatial information was used (the "environmental covariate" model). Although the fit of the model to sandeel density was improved by applying this technique there were less clear results regarding the predictive power. However, the addition of our spatially adaptive smooth in both the presence / absence and animal density models suggest that there may be more of a division between the sandeels around the Isle of May region, which is an important feeding ground for seabirds, and the sandeels in

the main fishing areas, Berwick's Bank and Marr Bank, than would be suggested by the environmental covariate models. In addition, our models suggest higher probabilities of presence and higher densities of sandeel in the Berwick's Bank and Marr Bank regions than the environmental covariate models. Such observations lead us to consider whether the sandeels around the Isle of May, and therefore the seabirds which feed on them, may be less affected by fisheries operating on Berwick's Bank and Marr Bank than previously thought. There also appears to be something which makes the Berwick's Bank and Marr Bank especially appealing to sandeels which is not currently included in the environmental covariate models, further investigations into what these factors are will aid understanding of sandeel distribution. Our models also support the theory that the distribution of sandeels contracts and expands around areas of preferred habitat (Wright et al., 2000), therefore it is important to be aware that although fisheries catches may not decrease the sandeel population may well be in decline.

5.1.2 Management of Acoustic Noise in the Marine Environment

When the Royal Navy operates sonar, one consideration is where can they do so with minimal disturbance to marine mammals. Such concerns led to the development of the Environmental Risk Management Capability (ERMC) (Mollett et al., 2009), which amongst other things, uses the predicted densities of various species of marine mammal to try to choose locations which are less likely to have high densities of animals. It is therefore important that the predicted densities of marine mammals in different regions is accurate. In Chapter 3, we apply the same novel spatially adaptive smoothing technique to model the distribution of harbour porpoise to the west of the United Kingdom. The cross validation results suggested that a number of different models fitted using our spatially adaptive smoothing, as well as an `MGCV`

model, performed equally well for predicting the distribution of this species. However, the predictions from each of these models differed somewhat, an uncertainty that as yet is not accounted for in these risk assessment methods. Such results suggest that in such instances we must move to methods which incorporate predictions from a range of models rather than just relying on a single model, perhaps using techniques such as model averaging.

Such sources of uncertainty in our prediction of where species are likely to occur at a given time, as well as operational priorities, mean that sonar may be operated in areas with higher abundance of marine mammals. This has led to the implementation of additional management measures such as monitoring zones, however, currently the size of these zones is chosen arbitrarily (Compton et al., 2008). Chapter 4 of this thesis, incorporated parts of the ERMC software to create sophisticated simulations which provided the first scientific evaluation into the effectiveness of monitoring zones. We have shown that for cetacean species, a monitoring zone of at least 2 km is generally advisable for reducing the risk of TTS and PTS for any animal sighted from the vessel. However, in contrast, monitoring for pinniped presence did not appear to be successful in reducing the risks of TTS and PTS due to the higher sensitivity in the hearing of these animals and the fact that they are harder to detect; a demonstration of the importance of improving our ability to detect marine mammals if this technique is to be effective. A further consideration highlighted in this chapter, is the rapid rate of increase in cumulative sound exposure level at the beginning of the simulation for any animal in the vicinity of the vessel, such observations indicate the importance of monitoring for marine mammal presence for a period of time prior to the operation of the sonar.

5.2 Making the Most of Computer Intensive Methodology

The application of all methods in this thesis has relied on computer intensive methodology. In Chapter 2, we implemented maximum likelihood techniques to generate parameter estimates as well as a bootstrap to estimate uncertainty. Chapter 3, relied on the computationally intensive branch and bound algorithm to attempt to find the optimal spatial smooth and the simulations in Chapter 4 pushed even the most up-to-date personal computer at the time to its limits. However, one point that became obvious, particularly when working on the sandeel recovery data in Chapter 2, is that no amount of computing power will compensate for a lack of data or knowledge. In this instance, we simply did not have enough information to accurately quantify the biomass of *A. marinus* consumed by grey seals, when they were also consuming other species of sandeel. Although increased computational power allows us to more easily test the assumptions and sensitivity of our models to uncertainty, the accuracy and precision of our results ultimately relies on the quality of the data.

Computer intensive methodology can provide assistance in the collection of good quality data. Currently, software such as `WiSP` (Zucchini et al., 2007) and the automated survey design engine within `DISTANCE`, versions 6.0 (Thomas et al., 2010) provide methodologies that aid survey design. This software uses computationally intensive methodologies to assess the effectiveness of different survey designs and sampling protocols prior to commencing the survey. As part of this process, it is possible to test the robustness of different techniques under various scenarios and choose the one that performs the best, e.g. the one that will give results with the least bias or lowest variance.

Such testing of survey or experimental methods via intensive computational methods prior to commencing data collection can also be used to generate more cost effective surveys, as well as providing better quality data . The cost of running computer simulations is vastly lower than the costs of data collection. In addition, Roff (2006) points out that, where possible, experiments or surveys should be designed so that traditional analyses may be applied. It is only when such methods are not applicable that we must resort to analyses that can better address the challenges encountered. Traditional analyses have the advantages that they have been well tested, possess well studied properties and have a large number of associated diagnostic techniques.

5.3 Future Directions

This thesis has investigated a variety of data and computer intensive methodologies. However, within each chapter there is still potential for future research and methodological development.

Firstly, we believe that one of the most promising techniques for accurately quantifying the biomass of *A. marinus* consumed by grey seals would be DNA analysis of faecal samples. Indeed, during seal scat sampling in 2010 samples were collected for DNA analysis in addition to the collection of hard parts (Hammond pers. comm.). However, DNA analysis is expensive, and the required amounts of funding is one of the limitations of this technique. Although there is some potential to use sophisticated shape analysis techniques to distinguish the otoliths of *A. marinus* from those of *H. lanceolatus*, they are difficult to differentiate even prior to digestion.

A further consideration of the otolith analyses is ensuring that the total consumed biomass of sandeels is estimated accurately. This requires further studies to investigate the relation between fish mass and otolith length for both *A. marinus* and *H.*

lanceolatus. In addition, further feeding experiments are required to more accurately determine the relationship between otolith size and the probability of recovery.

The most important development for furthering spatial modelling methodology would be to create a suite of 2-dimensional benchmark functions. The only currently-available 2-dimensional benchmark function that we are aware of is the horseshoe (e.g. Wood et al. (2008)), which lacks a realistic complexity. Such functions can be used to assess the effectiveness of spatial modelling techniques, such as the spatially adaptive smoothing considered in Chapter 3, by assessing the fit to the underlying function, rather than relying on cross validation techniques where the true function is unknown. These benchmark functions can also be used to assess the effectiveness of the modelling techniques under different scenarios, e.g. different levels of observational error.

Finally, there is potential for the monitoring range simulations to be applied to other scenarios, e.g. different species and locations, and different sound characteristics (frequencies, source strength, duty cycle, duration). In addition, the collection of more accurate data on marine mammal response to sound would allow the effectiveness of "soft starts" to be assessed. However, before such investigations are undertaken it is essential to develop more efficient code for the simulations. Currently, the simulations for one species, based on a single detection function, take around two weeks to complete.

5.4 Concluding Remarks

With increasing technology has come increasing risks for the marine environment. However, this thesis has demonstrated a number of ways in which technological and statistical developments can also be used to aid more environmentally sustainable, and

potentially more economically productive, management of these valuable resources. We must therefore ensure that our management techniques keep pace with the rapidly advancing technology of our society.

Bibliography

W. Bowen. Reconstruction of pinniped diets: accounting for complete digestion of otoliths and cephalopod beaks. *Can. J. Fish. Aquat. Sci.*, 57:898–905, 2000.

R. Compton, L. Goodwin, R. Handy, and V. Abbott. A critical examination of worldwide guidelines for minimising the disturbance to marine mammals during seismic surveys. *Marine Policy*, 32:255–262, 2008.

R. Froese. Cube law, condition factor and weight-length relationships: history, meta-analysis and recommendations. *Journal of Applied Ichthyology*, 22:241–253, 2006.

K. Grellier and P. Hammond. Robust digestion and passage rate estimates for hard parts of grey seal (*Halichoerus grypus*) prey. *Can. J. Fish. Aquat. Sci.*, 63:1982–1998, 2006.

P. Hammond, A. Hall, and J. Prime. The diet of grey seals around Orkney and other island and mainland sites in north-eastern Scotland. *The Journal of Applied Ecology*, 31(2):340–350, 1994a.

P. Hammond, A. Hall, and J. Prime. The diet of grey seals in the Inner and Outer Hebrides. *Journal of Applied Ecology*, 31(4):737–746, 1994b.

E. Mammen and S. van de Geer. Locally adaptive regression splines. *The Annals of Statistics*, 25:387–413, 1997.

A. Mollett, C. Schofield, I. Miller, J. Harwood, C. Harris, and C. Donovan. Environmental risk management capability: Advice on minimising the impact of both sonar and seismic offshore operations on marine mammals. Technical report, SPE Offshore Europe Oil & Gas Conference & Exhibition held in Aberdeen, UK, 8th to 11th September, 2009.

G. Pierce, M. Santos, and S. Cerviño. Assessing sources of variation underlying estimated of cetacean diet composition: a simulation study on analysis of harbour porpoise diet in Scottish (UK) waters. *Journal of the Marine Biological Association*, 87:213–221, 2007.

J. Prime and P. Hammond. The diet of grey seals from the south-western North Sea assessed from analysis of hard parts found in faeces. *The Journal of Applied Ecology*, 27(2):435–447, 1990.

D. Roff. *Introduction to Computer-Intensive Methods of Data Analysis in Biology*. Cambridge University Press, 2006.

L. Thomas, S. Buckland, E. Rexstad, J. L. Laake, S. Strindberg, S. L. Hedley, J. R. Bishop, T. A. Marques, and K. P. Burnham. Distance software: design and analysis of distance sampling surveys for estimating population size. *Journal of Applied Ecology*, 47:5–14, 2010. URL `http://www.ruwpa.st-and.ac.uk/distance/`.

D. Tollit, M. Stewart, P. Thompson, G. Pierce, M. Santos, and S. Hughes. Species and size differences in the digestion of otoliths and beaks: implications for estimates of pinniped diet composition. *Can. J. Fish. Aquat. Sci.*, 54:105–119, 1997.

S. Wood, M. Bravington, and S. L. Hedley. Soap film smoothing. *Journal of the Royal Statistical Society. Sers. B*, 70:931–955, 2008.

P. Wright, H. Jensen, and I. Tuck. The influence of sediment type on the distribution of the lesser sandeel, *Ammodytes marinus*. *Journal of Sea Research*, 44:243–256, 2000.

S. Zhou and X. Shen. Spatially adaptive regression splines and accurate knot selection schemes. *J. Am. Stat. Assoc.*, 96:247–259, 2001.

W. Zucchini, D. Borchers, M. Erdelmeier, E. Rexstad, and J. Bishop. WiSP 1.2.4. Technical report, Institut fur Statistik und Okonometrie, Geror-August-Universitat Gottingen, Platz der Gottinger Seiben 5, Gottingen, Germany., 2007. URL `http://www.ruwpa.st-and.ac.uk/estimating.abundance/WiSP/index.html`.