THE UNIVERSITY *of* EDINBURGH

Modelling animal movement for behavioural inference

Rebecca Ayodeji Akeresola

Doctor of Philosophy

University of Edinburgh

October 2025

Declaration

I declare that this thesis was composed by myself under the guidance of my supervisors and that the work contained therein is my own, except where explicitly stated otherwise in the text. Chapters 4 and 5 are based on my published paper for which I am the lead and first author:

Akeresola, R. A., Butler, A., Jones, E. L., King, R., Elvira, V., Black, J. Robertson, G. (2024). Validating hidden Markov models for seabird behavioural inference. *Ecology and Evolution,* 14(3).

Available at https://doi.org/10.1002/ece3.11116

The work in this thesis has not been submitted for any other degree or professional qualification except as specified.

(*Rebecca Ayodeji Akeresola*)

To God Almighty, El Roi

Acknowledgements

I extend my sincere appreciation to my supervisors — Prof. Victor Elvira, Prof. Ruth King (University of Edinburgh), Dr. Gail Robertson, Dr. Adam Butler, Dr. Esther Jones (BioSS), and Prof. Andy White (Heriot-Watt University). Their expertise, insightful feedback, and invaluable guidance have been instrumental in the successful completion of this thesis. I am particularly thankful for their unwavering support and the constructive discussions that greatly enhanced the quality of this work.

I am deeply indebted to the Maxwell Institute Graduate School (MAC-MIGS) for providing financial support and the platform to undertake this research.

I am also grateful to the Biomathematics and Statistics Scotland (BioSS) team for their collaboration and support during my PhD. My special thanks go to Ana for her kind assistance with the codes.

My heartfelt gratitude goes to my husband, Olawale, and my daughter, Oreofeoluwa, for their unconditional love, patience, and constant encouragement during the course of my PhD studies.

I am grateful to my parents, Mr. and Mrs. Akeresola, for their love, prayers, and steadfast support, particularly during the writing phase of this thesis.

Finally, I would like to thank my family and friends: Joshua, Mayowa, Tammy, and Aminat for their encouragement and support throughout the PhD journey.

Lay Summary

Imagine standing at the seashore observing a seabird. It is easy to tell where the bird is and what the seabird is doing at each time; however, it becomes impossible to tell once the bird is out of sight. Tracking devices are often attached to animals to know where they have been, but cannot tell what the animal was doing during the tracking period.

Ecological studies are often interested in knowing what an animal is doing (behavior) to help address challenges (e.g., threats faced by animals) that are critical for effective conservation of animals. Since the animal's behaviour influences the way it moves, the movement pattern when flying will differ from that when resting and walking. Therefore, we can say the behaviours are (hidden) underlying factors that influence movement. Ecological statisticians and statistical ecologists often use a popular model on the tracking data to predict animal behaviours from the movement data. This model is known as *hidden Markov models* (HMMs). Researchers often use very large amounts of data on animal movement in these models, but the size of the data can make it difficult to run these models easily. As a result, they first reduce the size of the data before using it in the models for analysis.

This thesis addresses two key problems. First, it is challenging to know if behaviours predicted by HMMs are actually correct. This is because it is difficult to directly observe the behaviour of animals and so do not have true behavioural data to validate the predicted behaviours. Second, reducing the size of movement data can lead to loss of underlying behavioural information that we want to predict, and can limit the analysis.

In this thesis, we address the first problem with a unique data on the observed behaviour of birds and the movement data of the boat used in following the observed bird. We demonstrate that movement tracks of the boat are a close representation of the approximate bird's tracks and use HMMs to predict the bird behaviour from the boat tracks. We assess predicted behaviours with actual behaviour and demonstrate

that HMMs are useful for identifying behaviours from movement data. Results suggest HMMs can be used more confidently to identify behaviours. Finally, we address the second problem by analyzing the movement data in the frequency domain and demonstrate that our approach (the use of *Nyquist-Shannon sampling theorem*) provides a more efficient way of reducing the number of data points but still preserves and minimizes information loss, unlike the traditional thinning technique. Findings allow a large volume of datasets to be used more easily and efficiently.

Abstract

Effective animal conservation is more important than ever in the face of biodiversity and climate crises. Understanding animal movement and behaviour can aid spatial planning and inform conservation management. However, it is difficult to directly observe and record behaviours in remote and hostile terrain such as the marine environment. Researchers often attempt to infer animal behaviour from telemetry data using hidden Markov models (HMMs), as animal movement patterns are influenced by their underlying behaviour. However, these inferred behaviours are not typically validated due to difficulty in obtaining 'ground truth' behavioural information. In practice, ecological researchers often use the inferred behaviours to suggest practical ways of improving conservation efforts. This includes identifying and designating important foraging areas as Marine Protected Areas (MPAs), spatial planning, and construction of offshore wind farms. Additionally, since some of these animals, particularly those in the marine environment, receive less protection at sea, it is crucial that the inferred behaviors are actually identified in the right areas. Furthermore, large-scale animal telemetry datasets are readily available due to advances in biologging technologies, and studies on animal movement and behaviour have been on the increase as a result. The use of these telemetry datasets for statistical modelling can be computationally costly due to the large volume and high-resolution of the datasets. To easily fit these models, the need to reduce the number of data points prior to analysis often arises. In the process of reducing the number of data points, some of the information embedded in the original dataset that may be useful for analysis can be lost.

In the first part of this thesis, we have a dataset obtained from the visual tracking method that provides a ground truth behavioural dataset of terns and the boat GPS track as a proxy for bird tracks. Leveraging on this unique data, we assess whether (i) the visual tracking information from the boat is a good proxy for true bird locations in relation to inferred behaviours of the fitted HMM, and (ii) the inferred behaviours from HMMs fitted to visual tracking data accurately represent true behaviours as

identified by behavioural observations taken from the boat. We demonstrate that visual tracking data can be regarded as a good proxy for true movement data in terms of similarity in inferred behaviours. In the second validation, we assess the validity of HMMs-inferred behaviour by fitting HMMs to boat visual tracks, inferring behaviours from fitted models, and assessing inferred behaviour using ground-truth observed behavioural data. Our results suggest that HMMs fitted to tracking data can accurately identify important conservation-relevant behaviours in seabird species, as demonstrated using visual tracking data.

In the second part of this thesis, we examine a more efficient way of subsampling animal telemetry data in addition to the standard thinning approach. We adopt the Nyquist-Shannon sampling theorem (NSST) to inform the choice of subsampling frequency. While NSST can be used to inform the temporal resolutions for obtaining animal movement trajectories, it also describes how to sub-sample a discrete-time signal (animal movement data) in a way that there is little or no information loss about the underlying latent behavioural process. We explain how to adopt NSST in animal movement modelling to inform the choice of sampling frequency and subsequently reduce the number of data points while minimizing information loss. We initially show that NSST can be useful for determining whether a chosen sampling interval is sufficient for behavioural analysis. We also showed using two scenarios (i) cases where the standard thinning approach alone produces results similar to the NSST approach with respect to behavioural inference and close representation of original movement location, (ii) cases where the NSST approach provides better results than the thinning approach. Lastly, we examine how the accuracy of inferred behaviours is sensitive to the choice of subsampling frequencies informed by NSST. Results suggest that the NSST approach is useful for subsampling as it provides improved behavioural accuracy compared to the standard thinning approach.

Contents

Chapter 1

Introduction

Ecologists are often interested in the three "W" questions that are fundamental to animal movement ecology: the "where", "what", and "why" questions. Where has the animal been? What was the animal doing? Why was the animal in a given location? Answering these three questions is relevant to identifying what environmental variables are needed for effective conservation management. The "where" question tells us about animal movement in space and time, the "what" question provides insights into the animal's behaviour, and the "why" question reveals the interesting things about the area of the environment that attracts the animal. To address these questions, ecologists use the movement tracks of animals, often obtained by attaching Global Positioning System (GPS) devices to individuals, and then use appropriate statistical models to analyze resulting movement tracks and infer behaviour from movement (Patterson et al. 2009, Langrock et al. 2012, McClintock et al. 2017). The purpose of these analyses includes predicting where species might be distributed in the environment, where animals will be in the future, and the distribution and behaviour of animals in areas in which tracking has not taken place. However, there are various problems with using GPS movement data to infer distribution and behaviour: (i) some species may not be able to carry tracking devices such as GPS due to their small size (though this is becoming less of a problem as technological advances allow GPS to become smaller (Cooke et al. 2004)); (ii) it is uncertain if behaviours inferred from movement tracks are correct as results are not often validated using ground truth data and (iii) it is difficult to run statistical models using the very large datasets generated by animal tracking, and these inferential statistical models that tells the ecologists why the animals are doing what they doing. This thesis addresses the three problems by firstly validating an alternative approach (visual tracking) of collecting movement tracks of animals as a proxy to GPS data collection, secondly validating commonly used statistical methods

used to infer behaviours of animals from movement tracks and thirdly explaining how to reduce datapoints by subsampling movement datasets without reducing the amount of information they represent for use in statistical models. The work done in this thesis makes difficult data choices for statisticians and non-statisticians in ecology more possible because inferential statistical models can be used for movement data obtained using either visual tracking or GPS tracking with more validity, confidence (Chapters 4 and 5), and more easily (Chapter 6).

1.1 Animal movement and behaviour

Conservation of biodiversity is vital for the health of the ecosystem. Over the years, anthropogenic activities ranging from the destruction of natural habitats to the introduction of invasive species, hunting, over-consumption, and over-exploitation have contributed to the decline, endangerment, and extinction of species (Ehrlich & Wilson 1991, Lande 1998, Jeff 2019). Although extinction is a natural process (Chapin et al. 2000), the current rate of extinction is alarming (Johnson 2020). In 2020, the International Union for Conservation of Nature (IUCN) Red list reported a total of 15,403 threatened (i.e., critically endangered, endangered, and vulnerable) animal species. The number of threatened species will continue to increase if anthropogenic activities persist as normal and drastic conservation actions are not taken (Jeff 2019). The decline in biodiversity is detrimental to ecosystem function and is a direct consequence of the decline and extinction of species (Larsen et al. 2005). Some extinct or nearly extinct species can be restored through reintroduction, but this can be costly, challenging, and sometimes unsuccessful (Hilbers et al. 2020). The original causes of extinction must be addressed before reintroduction is likely to be successful, further stressing the importance of ecosystem resilience. The increase in the rate of endangered species and extinction can be curbed through other conservation actions such as controlling invasive species, establishing protected areas, and biodiversity programmes (IUCN 2020).

Knowledge of how animals use the environment and areas important for breeding and survival enhances the effectiveness of conservation management and planning, particularly in identifying areas to be protected (Lascelles et al. 2016). Answering questions of where animals move to, how long they spend in a given location (for instance, animals are more likely to stay longer in areas with abundant food compared to areas without food), and what they are doing in a given location is part of the necessary information required to identify areas that need to be protected. Thus, understanding animal movement and behaviour is relevant for conservation purposes.

Several studies have demonstrated the value of animal movement research in informing applied conservation. Movement data have been used to guide infrastructure planning (e.g., road placement and urban development, wildlife crossing design, and wind farm construction), to support the design and evaluation of protected areas (including marine protected areas and special protected areas), to mitigate human–wildlife conflict, and to identify critical habitats such as foraging grounds, breeding sites, and migratory stopover sites (Wakefield et al. 2009, Lascelles et al. 2016, Nathan et al. 2022).

For example, Whittington et al. (2022) quantified the extent to which the development

of towns and roads altered movement behaviour and habitat use of bears and wolves, highlighting how movement data can guide infrastructure planning while minimising anthropogenic disturbance. In efforts to conserve the critically endangered Mojave desert tortoise (*Gopherus agassizii*), Hromada et al. (2020) analysed their movement data at study sites along the California–Nevada border to inform the design of conservation corridors, and their findings indicated that mitigation corridors designated between solar facilities must be sufficiently wide to retain and maintain function of home range. Similarly, movement data from Woodland caribou (*Rangifer tarandus*) have been used to assess the impact of oil and gas well sites on caribou habitat use in Alberta and to inform best management and conservation practices for this threatened species (MacNearney et al. 2021).

In the marine ecosystem, van Zinnicq Bergmann et al. (2022) integrated movement data from sharks and rays with spatial conservation planning tools to design multiple marine protected areas (MPAs) around Bimini, Bahamas and evaluate how effectively existing proposed MPAs protect essential habitats. Likewise, movement data from migratory shorebirds in coastal protected areas in mainland China revealed that existing protected area boundaries failed to encompass the full range of habitats used by birds during tidal cycles and local movements, leading to recommendations for boundary revisions to ensure effective protection of mobile species and key habitat (Choi et al. 2019). More broadly, movement data of seabirds have enabled the identification of key foraging areas and informed candidate MPAs, supporting more effective and sustainable spatial management of marine ecosystems (Le Corre et al. 2012, Lascelles et al. 2012, Thaxter et al. 2012).

Although many studies show that animal movement research is important for conservation, applying these findings directly to policy and management can be challenging. These include mismatches between the spatial scale of animal movements and the scale at which management decisions are made, short-term or incomplete data, limited funding, difficulty of integrating dynamic movement patterns into fixed planning and barriers in knowledge transfer and practical implementation (Lascelles et al. 2016, Fischer 2024). Furthermore, despite significant advances in tag technology to allow large volumes of tracking data to be collected, methods to analyse these data using movement models can be complex due to missing or incomplete information, and translating insights about animal behaviour into policy-relevance metrics remains challenging. Nevertheless, research on animal movement is growing rapidly and continues to provide valuable insights for effective conservation.

Movement is a fundamental characteristic of living organisms and is crucial for

survival and reproduction. While there are various motives for movement, animals move primarily in search of food, shelter, safety, and reproduction. Different movement patterns that can be classified into biologically meaningful animal behavioural states may often exist in animal movement because the behaviours of animals directly influence their movement. For instance, patterns in animals' movement when searching for food may differ from when they are feeding. Data on animal movement can be obtained by tracking their movement. Advances in technology have led to the development of various devices for tracking animal movement. Some of these devices attached to animals include GPS transmitters, acoustic tags, radio transmitters, Argos Doppler tags, biologgers, and harmonic radar trackers (Cooke et al. 2004, Cant et al. 2005, Leos-Barajas et al. 2017). The choice of tracking devices for retrieving animal movement depends on the animal to be tracked and what the research aims to address.

The movement data of animals obtained from tracking devices provides direct information about the movement paths of animals, thus answering the question of where they move to. However, the movement data do not provide direct information about the behaviour of animals, that is, why they move to a given location and, particularly, what they do in that location. The knowledge of what the animals are doing at a given location can help delineate important locations animals use consistently, which is usually of interest in conservation management. Animal behaviour can be classed into different states such as flying, resting, feeding, travelling, searching, etc. Accessing first-hand behaviours of tracked animals is difficult in practice, especially within the marine environment. As a result, behavioural data of tracked animals are rarely available. An exception, however, is when tagged animals are visually tracked so that they can be observed to note down their behaviours (Perrow et al. 2011) or when devices that can be used to obtain relevant information that serve as a proxy for animals' behaviour are attached to tracked animals. For example, the use of accelerometers, visual high-frequency collars, cameras, and heat sensors.

Due to advances in technology, vast movement data retrieved from satellite and GPS transmitters, biologgers, and harmonic radar trackers (Cooke et al. 2004, Cant et al. 2005) are readily available (e.g., `https://www.movebank.org`). As a result, studies on animal movement in movement ecology have received much attention over the years. It has been established in several works of literature (e.g., Morales et al. (2004), Jonsen et al. (2005), Edelhoff et al. (2016)) that unobservable behavioural states characterized by different movement patterns can often be inferred from animal movement data. In addition, environmental data integrated with animal movement data improves the inference of animal behaviour (Patterson et al. 2009, McClintock

et al. 2017), as they may be used as a proxy for behavioural data where the latter are not available (Browning et al. 2018). Ecological researchers have employed various approaches to model animal movement data, including state-space models (e.g., Jonsen et al. (2003), Patterson et al. (2008)), machine learning (e.g., Gutierrez-Galan et al. (2018), Wijeyakulasuriya et al. (2020)), hidden Markov models (e.g., Patterson et al. (2009), McClintock (2021)), to mention a few. The nature of the movement data, the goals to be achieved, and the questions to be answered may help inform the choice of a suitable approach. For instance, state-space models with an infinite number of states (i.e., the state variable is continuous) are good candidates for regular or irregularly spaced movement data with non-negligible measurement error, while hidden Markov models (HMMs) with a finite number of states are primarily suitable for regularly spaced movement data with negligible measurement error (Langrock et al. 2012). In animal movement modelling, where classification of behavioural states is the focus, supervised machine learning approaches are only possible when the output data, that is, animal behavioural data, are available alongside the input data, that is, the movement data. However, animal behavioural data are rarely available due to the difficulty in observing animals over a continuous period of time, for instance, at sea, and the supervised machine learning approach may not be suitable in such cases. Unsupervised machine learning, on the other hand, is possible in the absence of behavioural data since the approach does not require such observed output data.

HMMs are time series models with observation and hidden processes where the latent (unobserved) states describe the underlying behaviour of the individual (Langrock et al. 2012). It has been widely applied in animal movement modelling (Morales et al. 2004, Patterson et al. 2009, McClintock 2021) due to its interpretability, flexibility, and usefulness in separating out the different observational hidden processes. In particular, HMMs are useful for inferring the hidden behaviours of animals from their movement data and can be implemented using readily available software designed for animal movement modelling, such as the R packages moveHMM (Michelot et al. 2016) and momentuHMM (McClintock & Michelot 2018). HMM-inferred behaviours can be used to inform conservation decision-making, for example, the size and location of protected areas. However, very little research has focused on validating these inferred behaviours because contemporaneous behavioural observations on tracked individuals that can serve as a ground-truth dataset can be challenging to collect, particularly in hostile or featureless environments, such as open ocean (Joo et al. 2013).

Some studies have attempted to validate HMM-inferred behaviour from movement data, such as Joo et al. (2013), which validated the behaviour of fishing vessels using ground truth data recorded by onboard observers. Bennison et al. (2018) and Conners

et al. (2021) also validated HMM-inferred behaviours of northern gannet (*Morus bassanus*) and albatross using behaviours from depth-recorder and sensors as ground truth data, respectively. However, depth recorders and sensors are also proxies for ground truth data with their own error structures. Identifying animal behaviours from their movement data is helpful for effective conservation. For example, inferred behaviours can be used to delineate important areas that can be protected from anthropogenic threats. As a result, it is crucial to ensure that the inferred behaviours are actually correct. The work done in the first part of this thesis fills in the gap by addressing the certainty of inferred behaviours. Using a unique dataset described in Chapter 2, we initially examine the validity of the method that provides access to first-hand behavioural data and then assess the validity of HMM-inferred behaviours.

1.2 Animal movement and behaviour: a frequency-domain perspective

Animal movement occurs in continuous time. However, animal locations in space and time are typically sampled at fixed, discrete time intervals (McClintock et al. 2014). As such, the discrete representation of animal movement provides an approximation of the underlying continuous movement trajectories. Patterns in animal movement trajectories are shaped by the underlying behaviours of the animals that are often not observed (Patterson et al. 2009, McClintock et al. 2017). Therefore, the sampling interval at which an animal's location is obtained largely determines the temporal resolution at which behavioural inference can be made (Johnson et al. 2017). The same also applies to movement data subsampled at a chosen sampling interval from already existing data. The sampled locations are used to derive movement metrics such as step length, velocity, speed, and turning angle that are useful for analysis (e.g., behavioural inference) (Joo et al. 2013, Reynolds et al. 2015, Conners et al. 2021). This implies that any underlying issues (e.g., limitations or errors) with the sampled location can impact the usefulness of the derived movement metric for analysis.

Various studies have examined the impact of sampling interval on behavioural inference (Ryan et al. 2004, Gaylord et al. 2016) and movement metrics obtained from movement data (McCann et al. 2021, Schoombie et al. 2024). However, analyses in these studies have focused primarily on the time domain. While the analysis in the time domain is common and informative, they do not fully reveal the reasons why the choice of sampling intervals for collecting data and subsampling already collected data on behavioral inference may be sufficient. For instance, the scale at which a sampled animal's locations are recorded may fail to capture the fine detailed

behavioural dynamics of the animal, and analysis from such data may be limited for making inference on a specific behaviour. Similarly, subsamples drawn from existing movement datasets may fail to preserve the underlying behavioral patterns when behavioural dynamics are not considered prior to subsampling. A possible reason for this is that the sampling rate may have been chosen such that there are frequency components contained in the data to be subsampled that are higher than the sampling rate. However, this is not evident when observing data from the time domain.

The time-domain shows how movement data changes over time. On the other hand, the frequency-domain shows the frequency components of the movement data and the strength of each frequency (Allen & Mills 2004, Boashash 2015). Consequently, the frequency-domain framework can be used to assess insights on how the choice of sampling and subsampling intervals impacts behavioural inference. Since it is difficult to validate some inferred behaviours in practice due to difficulty obtaining ground-truth data (Cooke et al. 2004), it is important to minimize misrepresentations and loss of information from sampled and subsampled data where possible. Therefore, understanding how this can be done from the frequency domain perspective with appropriate frequency analysis is important. Nyquist-Shannon Sampling Theorem, usually used in the frequency-domain framework, can be used to determine how best to subsample enormous movement data that can be difficult to analyse without losing useful behavioural information (Nathan et al. 2022). The Nyquist–Shannon Sampling Theorem (NSST) states that a continuous signal can be perfectly represented by discrete samples if it is sampled at least twice as fast as its highest frequency component (Shannon 1949). The theoretical foundations from the NSST provide conditions under which perfect reconstruction of (i) a continuous signal can be achieved from its samples (ii) the sampled signal can be achieved from the subsampled version. However, an oversight of the conditions which produce misrepresentation or loss of information of the original signal data can be avoided with a technique applied prior to subsampling. NSST is commonly used in fields other than movement ecology. While there are more studies on the use of NSST to inform the choice of sampling interval for collecting accelerometer movement data (Tatler et al. 2018, Hounslow et al. 2019, Yu et al. 2023) than GPS movement data, the principles of NSST also apply to animal movement data (Nathan et al. 2022).

Subsampling movement data is common in animal movement ecology. In most cases, the subsampled movement data is simply assumed to be a faithful representation of the original movement data, with no special care taken to ensure the sampling theorem conditions. In this thesis, we apply the concept of NSST and examine how the theorem can be used to determine whether a sampling interval is sufficient or not,

using simulation studies followed by real examples as a case study. We also show how the technique from NSST can be applied to enhance the standard subsampling method by minimizing information loss and reducing the misrepresentation of movement data.

1.3 Thesis outline

This thesis specifically addresses the challenges highlighted in Sections 1.1 and 1.2 by answering the following questions:

1. Is visual tracking data a good proxy for movement data collected from GPS devices attached to animals?
2. Can we validate HMM-inferred behaviours using ground-truth data?
3. Is it possible to sub-sample existing animal movement data without losing too much information?

The thesis structure that provides an overview of the work done in answering the above questions is as follows.

Chapter 2 contains a description of the data used in this thesis. Specifically, we initially describe the study species, study sites and the visual tracking data used in Chapters 4 and 5 of the thesis, and then describe the simulated and case study data used in Chapter 6 of the thesis.

Chapter 3 provides a comprehensive overview of the main methods used in this thesis. A detailed description of hidden Markov models and commonly used R packages for HMMs is provided.

Chapter 4 is a preliminary to Chapter 5. In this Chapter, we present the first validation study for the first part of this thesis. We describe the visual tracking method that enables access to first-hand observed behavioural data. We then validate the visual tracking data as a proxy for GPS data collection using appropriate validation methods. The results are reported and are also discussed in detail.

Chapter 5 presents the second validation study, which involves the validation of HMM-inferred behaviours. We give a detailed description of the models and validation methods. The results are also reported and discussed in detail.

Chapter 6 presents how we apply the Nyquist-Shannon sampling theorem to subsample animal movement data. We present our approach for the subsampling process, accompanied by a demonstration using simulated data, and then apply it to two real datasets as case study examples.

The final chapter contains the conclusion of the thesis and discusses possible future extensions of the work done. The first part of this thesis, contained in Chapters 4 and 5, has been published in the Journal of Ecology and Evolution (Akeresola et al. 2024).

Chapter 2

Data Structure

In this chapter, we introduce the study species and describe the tern data used for Chapters 4 and 5. We also describe the simulated data and the case study data used in Chapter 6.

2.1 Tern data

HMM-inferred behaviours are rarely validated, as animals that are tagged with instruments such as GPS devices are often species that are difficult to observe, as they live in remote and hostile terrain, such as the marine environment, for example. However, validating HMMs is possible when observation data can be collected contemporaneously with location information, and used to ground-truth. To this end, we consider terns (*Sterna* spp.) as a study species because of the existing observational data on terns behaviours that provides a unique opportunity to validate HMM-inferred behaviours.

2.1.1 Study species

Terns (Figure 2.1) are part of the *Laridae* family, which includes gulls, terns, skimmers, and noddies. There are around forty species of tern worldwide, of which five regularly breed in the UK. Protected under Annex I of the EC Birds Directive, terns are qualifying species for 41 Special Protected Areas with Marine Components in the UK (JNCC 2021, 2020). This study investigates the movement behaviour of four tern species in order of abundance in the British Isles: Arctic (*Sterna paradisaea*; 56,115 pairs), common (*S.hirundo*; 14,323 pairs), Sandwich (*S.sandvicensis*; 14,252 pairs), and roseate (*S.dougallii*; 997 pairs [from 2006 surveys]) (Mavor et al. 2008).

Population estimates from surveys conducted between 1998-2002 and 2015 indicate there have been changes in trends for each species. Whilst 3 species have increased between these periods [roseate (229%), Arctic (17%), and Sandwich (13%)], common terns have declined by an estimated 10%, although all are species of least concern on the IUCN red list of threatened species. Arctic terns tend to breed in coastal areas in the north and west of the UK, with 80% occurring in Shetland, Orkney, and the Outer Hebrides. Common terns have a widespread coastal distribution around the UK and also nest in small colonies inland along rivers and islets. Sandwich terns congregate in large colonies around the UK coast, and the majority of roseate terns breed on Coquet Island (Northumberland) with some pairs occasionally breeding in south-east Scotland, Norfolk, and Hampshire (Wilson et al. 2014).

(a) Arctic (b) Common (c) Roseate (d) Sandwich

Figure 2.1: Study species - Terns. Source: *birdsoftheworld.org*.

Terns are ground-nesting colonial breeders, raising one brood each breeding season (May-June) and laying a clutch of one to three eggs. While breeding adult terns are central-place foragers throughout the breeding season, they are particularly restricted during chick rearing when they must return regularly to provision their chicks, and adults spend up to 80% of their time foraging (Thaxter et al. 2012). Their foraging ranges vary by species, and mean maximum foraging ranges are Arctic (7.1 $\pm$ 2.2 km), Sandwich (11.5 $\pm$ 4.7 km), common (4.5 $\pm$ 3.2km), and roseate (12.2 $\pm$ 12.1 km) (Thaxter et al. 2012). Sandwich terns are specialist predators that can exploit clupeids and sandeels from deeper water, potentially due to their wider foraging range. Likewise, roseate terns are specialists who also forage by plunge diving to depth, catching prey items of predominately sandeels, herring, and sprat. Common terns are generalist predators, and prey items include invertebrates, clupeids, sandeels, and gadoids. Arctic terns forage using several techniques, but are heavily dependent on sandeel, and changes in prey availability can affect their breeding success (Eglington & Perrow 2014).

2.1.2 Data source and study site

The tern datasets used in Chapters 4 and 5 were provided by the Joint Nature Conservation Committee (JNCC). The visual tracking method (see Chapter 4 for a detailed description) was employed by JNCC for data collection. The visual tracking method provided two distinct datasets: (i) the movement data, and (ii) the behavioural data. The latter is represented by the observed behaviour of terns that were visually tracked, while the former is represented by the GPS location (longitude and latitude) of the boat used for the visual tracking. The datasets were collected during incubation and chick-rearing periods (see for example Figure 2.2) in study sites that comprised of 9 breeding colonies across the UK (Figure 2.3): Blue Circle (54°49′N, 5°46′W) and Cockle Island (54°40′N, 5°37′W) in Northern Ireland; Cemlyn Bay (53°24′ N, 4°30′ W) in North Wales; Glas-Eileanan Island (56°49′N, 5°71′W), Forvie (57°18′N, 1°58′W), Isle of May (56°10′N, 2°32′W), Leith (55°96′N, 3°16′W) and South Shian (56°46′N, 5°36′W) in Scotland; and Coquet Island (55°20′N, 1°32′W) in England.

2.2 Simulated data

In this section, we describe how we simulate data used in Chapter 6. Simulations include three-state movement tracks of seabirds. The three states include travel, search, and forage. We initially describe the general approach of simulation for a general number of states and then describe the 3-state specifics. All the simulation is done in R (Allaire 2012) and we discuss the simulation below.

Movement data are simulated from a hidden Markov model - time series models with two stochastic processes (i) observed (movement metrics) and (ii) unobserved (underlying behavioural states). We simulate data by providing an initial location point and then simulate step length and turning angle (conditional on behavioural state). To achieve this, we first describe a generic approach for simulating the behavioural states, followed by the movement metrics, and finally the location of animals.

The sequences of behaviour are first simulated because modelling movement metrics from their respective distribution is dependent on the simulated underlying behavioural states. The sequence of behaviours consists of N states and is simulated to satisfy the Markov property, i.e., given the present latent state, S_t, the probability of being in the future state, S_{t+1}, is independent of the past state, S_{t-1}:

$$Pr(S_{t+1} \mid S_t, \cdots, S_1) = Pr(S_{t+1} \mid S_t).$$

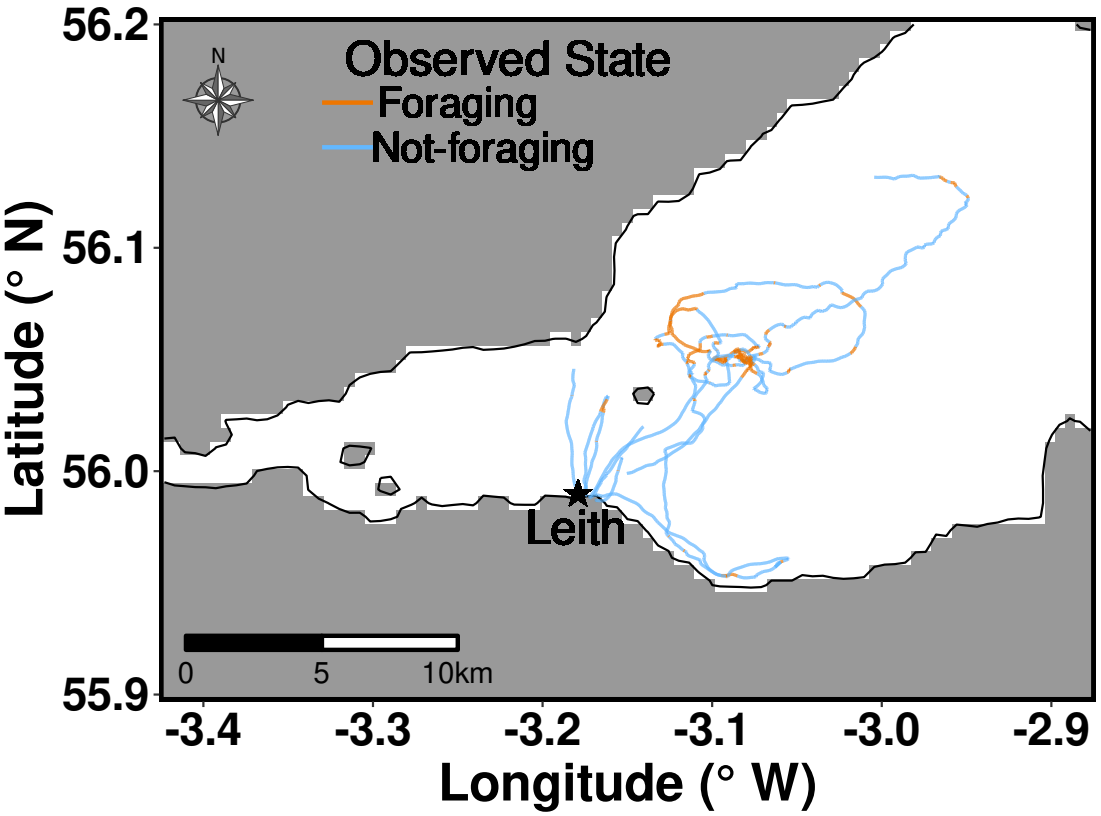

(a)

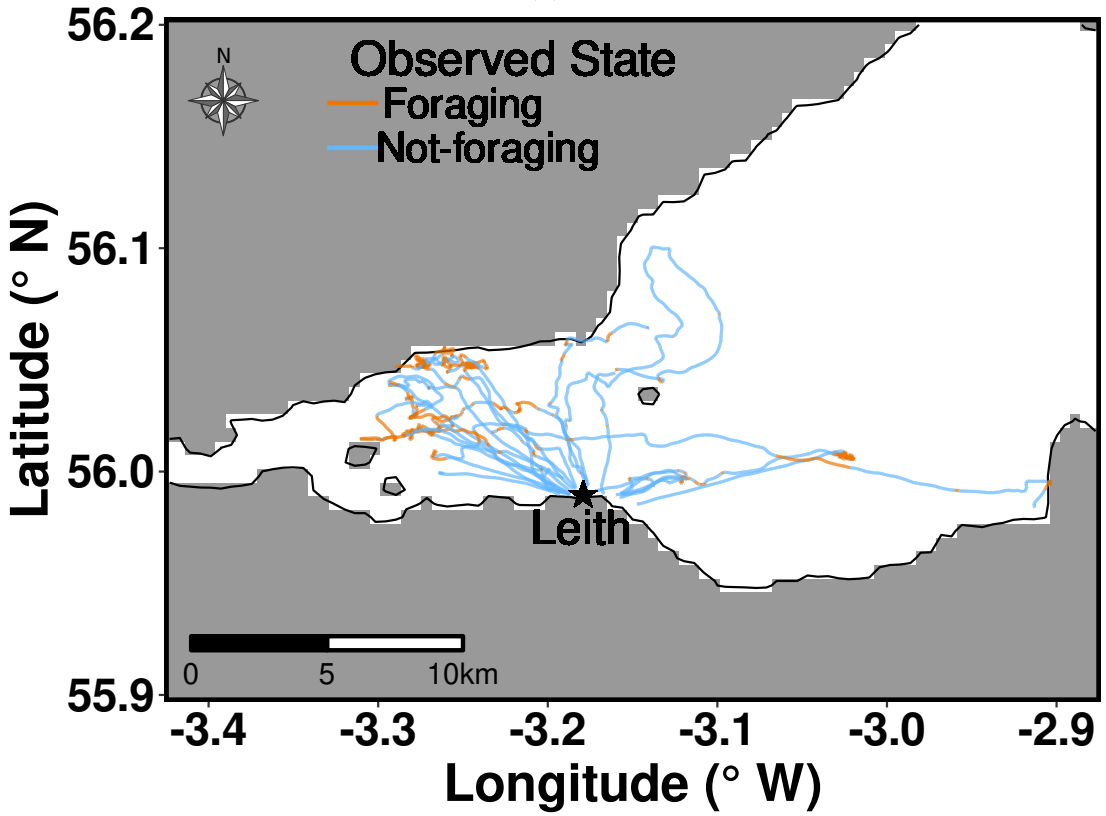

(b)

Figure 2.2: Visual tracks of common terns coloured with observed behavioural states during (a) incubation and (b) chick-rearing period from Leith dock, 2010.

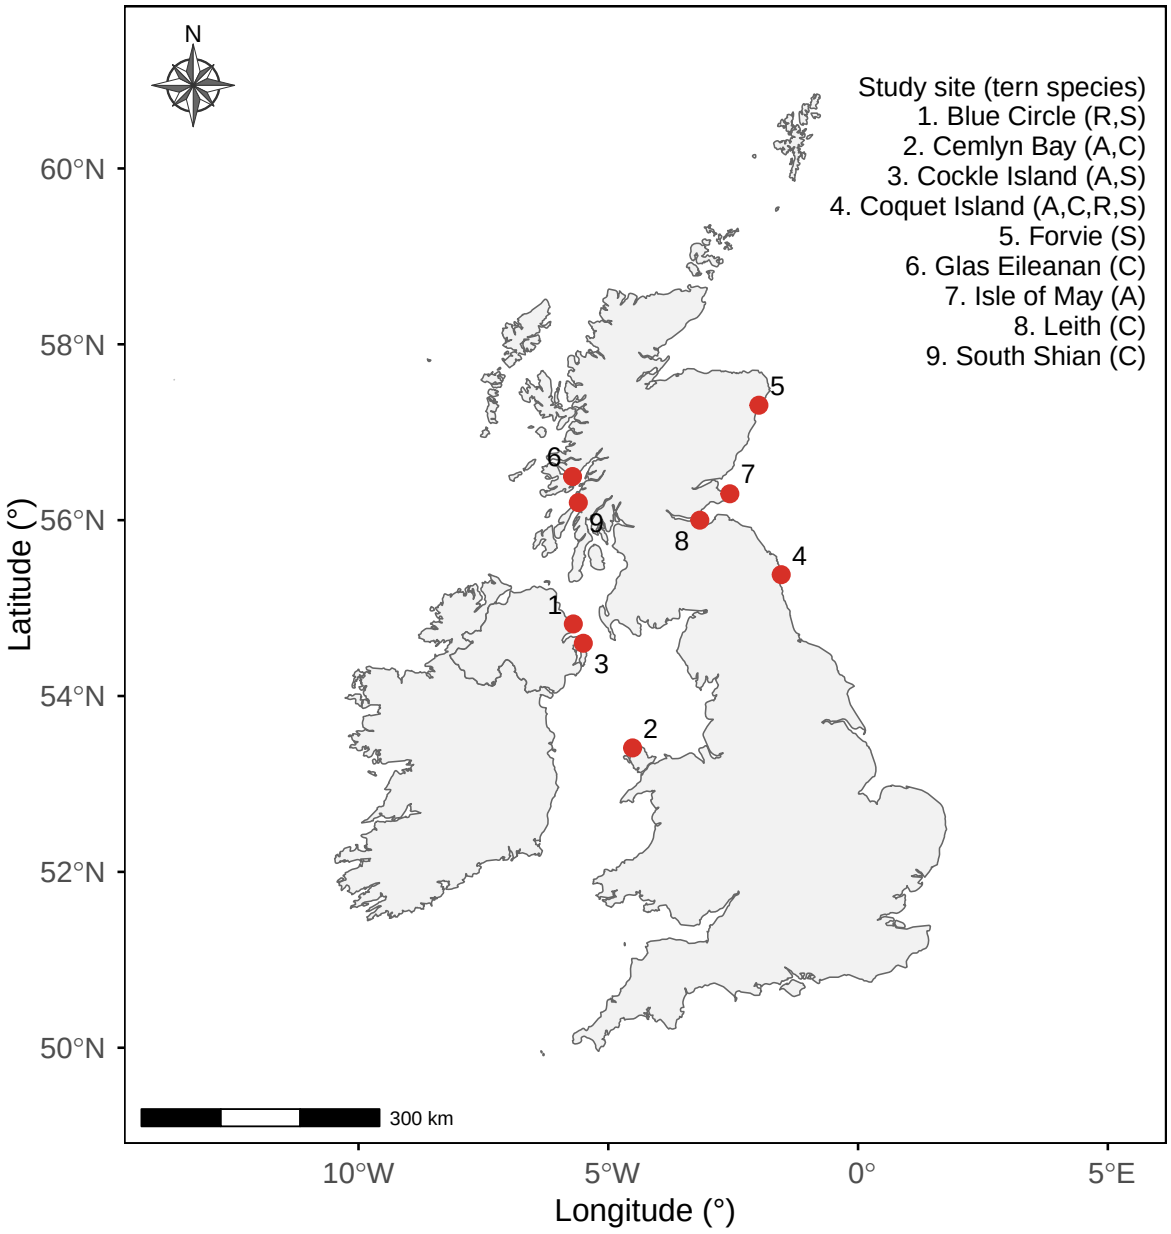

Figure 2.3: Study sites in the United Kingdom consisting of 9 tern breeding colonies. A = Arctic tern, C = common tern, R = roseate tern, and S = Sandwich tern.

The initial state is obtained by sampling and specifying the initial probabilities of being in each of the N states. The subsequent states are obtained by sampling and specifying probabilities of either remaining in the current state or switching to another state from the transition probability matrix, $\mathbf{\Gamma}$. These probabilities are the entries of the $N \times N$ transition probability matrix, $\mathbf{\Gamma}$.

To simulate movement metrics, initial values are first specified for the parameters of the distribution of movement metrics. The number of initial values specified for each parameter is the same as the number of behavioural states selected. For the step length, l_t„ initial values for the mean and standard deviation of the gamma distribution are specified for N behavioural states. On the other hand, for the turn angle, θ_t, the initial concentration parameter of the von Mises distribution is specified for N behavioural states. The corresponding movement metrics are then simulated and are assumed to (i) be regularly sampled at intervals and (ii) satisfy the conditional independence assumption, i.e., at time, t, for the step length, l_t, we have

$$f(l_t \mid S_1, \cdots, S_t, l_1, \cdots, l_{t-1}) = f(l_t \mid S_t),$$

while for the turning angle, θ_t, we have a

$$f(\theta_t \mid S_1, \cdots, S_t, \theta_1, \cdots, \theta_{t-1}) = f(\theta_t \mid S_t).$$

We specify the initial turning angle $\theta_1 = 0$ while the mean of the von Mises distribution for subsequent time, t, is the same as the turning angle at time $t - 1$.

At each time, t, simulations of longitude and latitude locations, (x_t, y_t), are obtained using the simulated movement metrics corresponding to each time. Values are specified for the starting location, (x_1, y_1), while subsequent locations are obtained using the formula below

$$x_{t+1} = x_t + l_t cos\theta_t$$
$$y_{t+1} = y_t + l_t sin\theta_t.$$

We go on to describe the three-state specifics in Section 2.2.1.

2.2.1 Three-state movement track

The three-state movement track simulation we conduct here is designed to replicate more realistic foraging trips of seabirds. Seabirds are known to be central place foragers, i.e., they head out to the sea from their breeding colonies in search of food, forage, and then head back to their breeding colonies either after a successful or unsuccessful foraging trip. We simulate a three-state movement track in the spirit of Michelot et al. (2017), and although the simulation codes adopted from Michelot et al. (2017) were designed for southern elephant seals (*Mirounga leonine*), it can be modified to suit other central place foragers, including seabirds. The three behavioural states include travel, $S_t = 1$, search, $S_t = 2$, and forage, $S_t = 3$. The travel state represents the departure of seabirds from the centre of attraction or breeding colony to foraging regions and return trip to the breeding colony or centre of attraction. The travel state is characterized by larger steps and high directional persistence. The search state is characterized by moderate steps and some directional persistence, while the foraging state has smaller steps and low directional persistence.

The simulation approach, which is an extension of the general approach described earlier, is detailed in (Michelot et al. 2017). We provide the parameter values used in simulating the behavioural state, movement metrics, and the movement locations. The values for the mean, μ, and standard deviation, σ, of step length for the four states are $(\mu_1, \mu_2, \mu_3) = (50, 10, 8)$ and $(\sigma_1, \sigma_2, \sigma_3) = (10.96, 8.00, 3.04)$ respectively. The values of the concentration parameter, κ, of the turning angle for the four states are $(\kappa_1, \kappa_2, \kappa_3) = (6.87, 3.70, 0.08)$. The centre of attraction, which is the breeding colony, is represented with longitude and latitude values specified as $(-2.555586, 56.185728)$. Tracks are then simulated as shown in Figure 2.4 with each track length summarized in Table 2.1.

Table 2.1: Simulated data with sampling interval, Δt = 5-mins.

	Track IDs									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Track length	158	77	191	158	178	132	164	162	151	216

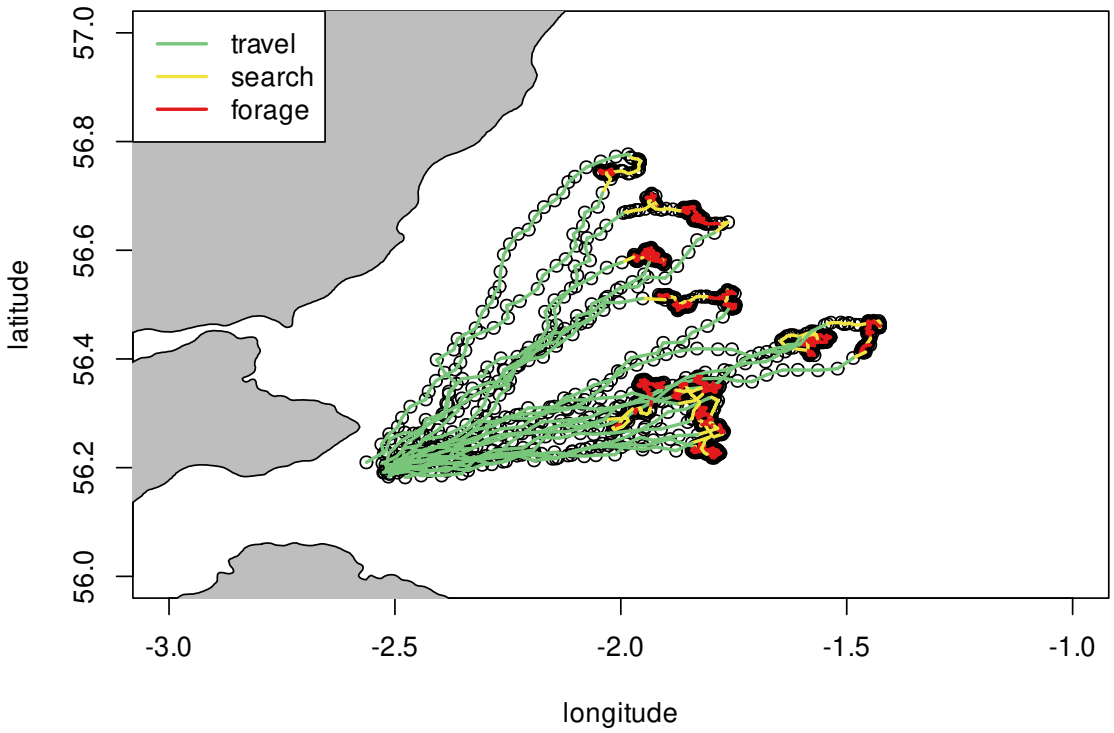

Figure 2.4: Three-state simulated tracks.

2.3 Case study data

In this section, we provide a brief description of the data used as case studies in Chapter 6.

2.3.1 Black-legged Kittiwake data

Black-legged kittiwake (*Rissa tridactyla*, hereafter "kittiwake") data were obtained from the GPS tracking of kittiwakes on Coquet Island, Northeast England (55°20′N, 1°32′W) during chick-rearing in 2012 (Robertson, Bolton, Grecian & Monaghan 2014). Kittiwakes (Figure 2.5) are central place foragers that feed almost exclusively on sandeel during the breeding season (Robertson, Bolton, Grecian & Monaghan 2014) and have an average body mass of 361.4 grams (Chivers et al. 2012). They are also considered to be good indicators of the health of the marine environment (**?**). A summary of the kittiwake tracks considered is presented in Table 2.2 while the GPS positions of tracked kittiwakes (Figure 2.6) were obtained at a sampling interval of 100-secs.

Figure 2.5: Black-legged kittiwake. Source: Wikipedia

Table 2.2: Summary of kittiwake GPS-tracking data obtained at sampling interval, $\Delta t = 100$-secs.

	Track IDs							
	0040	**0044**	**0048**	**0052**	**0059**	**0062**	**0064**	**0066**
Track length	918	1208	625	1873	640	315	1084	2348

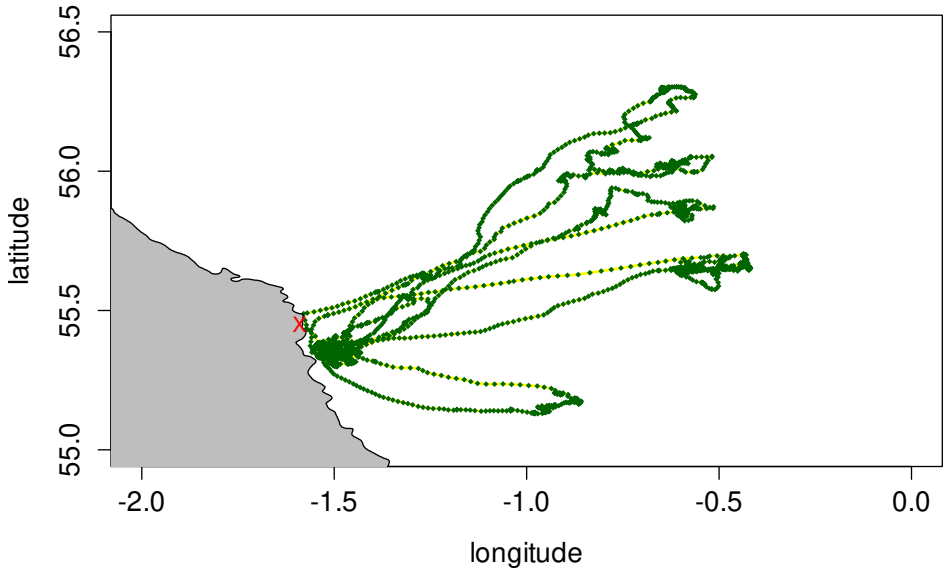

Figure 2.6: Kittiwake tracks during chick-rearing in 2012. Coquet Island (55°20′N, 1°32′W) is represented by "x".

Chapter 3

Hidden Markov Models

Hidden Markov models (HMMs) are time series models with a doubly stochastic process with the observation process driven by an underlying hidden Markov state process (Rabiner & Juang 1986, Zucchini et al. 2016). The hidden state process is assumed to satisfy the first-order Markov property so that at any given time, the next hidden state depends only on the current state and not on any earlier states. HMMs have wide application in fields such as speech recognition (Rabiner 1989), psychology (Visser et al. 2002), health (Kawamoto et al. 2013), agriculture (Leite et al. 2008), finance (Ghysels 1994), and movement ecology (Morales et al. 2004, McKellar et al. 2014). HMMs can be applied to a system of observations generated by an underlying hidden process and can be used to address problems classed as evaluation, decoding and learning. Given the observed sequence and the HMM, the learning problem optimizes the model parameters to best describe the observed sequence. The decoding problem on the other hand finds an optimal sequence of hidden states that best describes the observed sequence while the evaluation problem addresses the likelihood that the observed sequence was generated by the model (Rabiner & Juang 1986).

In the context of animal movement studies, HMMs can be used to find the most likely hidden sequence of behaviour (decoding problem) by drawing inferences about the hidden behaviours of animals from their movement data. HMMs can also be used to learn the movement pattern of each hidden behaviour and probability of switching behaviours (learning problem), and the likelihood of the movement patterns being produced by the model (evaluation problem). HMMs are natural candidates for modelling animal movement data particularly for addressing the decoding problem since access to behavioural data of animals is rare due to the difficulty in obtaining them (Dean et al. 2013, Joo et al. 2013). A HMM framework is often applied to telemetry data arising from tagged individuals with locational GPS data (e.g., Bennison et al.

(2018)) and accelerometer data (e.g., Conners et al. (2021)), which are used to classify hidden behavioural states in an underlying animal movement pattern. Indeed, HMMs have been successfully used to analyse the movement data of woodpeckers (McKellar et al. 2014), albatross (Conners et al. 2021), tuna (Patterson et al. 2009), eagle and shark (Leos-Barajas et al. 2017), elk (Morales et al. 2004), seal (McClintock et al. 2017), shearwater (Dean et al. 2013), and proxy movement data of terns (Akeresola et al. 2024).

The use of HMMs in animal movement modelling often comes with challenges. With telemetry data often serving as the sole information source, selecting a biologically meaningful and interpretable number of states to fit *a priori* can be challenging as animals exhibit different behaviours. Pohle et al. (2017) proposed a pragmatic approach to inform the number of states to fit to HMMs. This involves specifying a range of plausible state numbers based on biological knowledge, fitting corresponding models, and using model selection criteria (e.g., AIC) to guide selection. However, models with more states are often favoured by AIC, despite reduced biological interpretability. Thus, when the objective is behavioural inference, a model with fewer, more interpretable states may be preferred, even at the expense of model fit. Conversely, for forecasting future locations, a model with more states and better fit may be advantageous despite lower interpretability (Conners et al. 2021).

The simplifying assumptions underlying standard HMMs—including a first-order Markov process for the hidden states, negligible measurement error, and regularly spaced observations—are frequently violated in animal movement data as these assumptions are most times unrealistic for animal movement data. Telemetry data are often irregularly sampled, subject to location error, and influenced by underlying behavioural dynamics that may not strictly satisfy the Markov property (Patterson et al. 2017). As a result, direct application of basic HMM formulations can present methodological and inferential challenges. Patterson et al. (2017) outline practical strategies to address such limitations. Despite these limitations, HMMs remain widely used in statistical ecology due to their flexibility and potential capacity for extension to more complex framework (Langrock et al. 2012), as well as their accessibility through dedicated R packages (Michelot et al. 2016, McClintock & Michelot 2018).

3.1 Definition

Let X_t for time $t = 1, \ldots, T$ be an observed or state-dependent process and S_t for time $t = 1, \ldots, T$ be an unobservable or state process. An HMM is a time series model with two components: the observed process, X_t, and the state process, S_t. The observed process, X_t, is said to be driven by an underlying hidden state process, S_t. Figure 3.1 provides a graphical representation of an HMM.

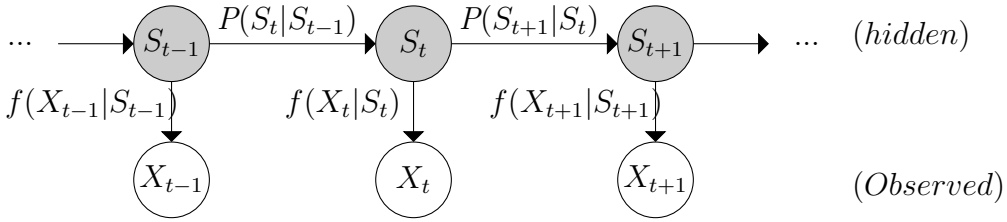

Figure 3.1: Graphical representation of a HMM where S_t denotes the state or hidden process, and X_t denotes the observed component.

The state process, S_t, takes a value on a finite set of N possible values (states) and is assumed to be a first-order Markov chain with the state transition probability, γ, given as

$$\gamma_{ij} = P(S_t = j | S_{t-1} = i), \tag{3.1}$$

that is, conditional on the current state $S_{t-1} = i$, the future state $S_t = j$ is independent of the past state. γ_{ij} are the entries of the $N \times N$ transition probability matrix, $\mathbf{\Gamma}$, defining the probability of state switching.

The observed process, X_t, which can be univariate or multivariate, is assumed to be regularly spaced in time, t, with the associated observation process distribution given as

$$f(X_t | X_{t-1}, \ldots, X_1, S_t, \ldots, S_1) = f(X_t | S_t), \tag{3.2}$$

at any given time t, that is, the distribution that generates an observation at each time depends on the state of the underlying and hidden process at that time.

3.2 Likelihood of a hidden Markov model (HMM)

The corresponding likelihood of an HMM is a function of the following parameters:

(i) $\boldsymbol{\delta}$: $(1 \times N)$ vector of initial state distribution

The initial distribution $\boldsymbol{\delta}$ is a row vector of probabilities with length equal to the number of states N. i.e., $\boldsymbol{\delta} = (\delta_i = P(S_i = i), \ldots, \delta_N = P(S_N = N))$ where $\delta_i \in [0, 1] \; \forall \; i = 1, \ldots, N$ and $\sum_{i=1}^{N} \delta_i = 1$.

(ii) $\boldsymbol{\Gamma}$: $(N \times N)$ matrix of transition probabilities.

The transition probability matrix $\boldsymbol{\Gamma}$ defining the probability of switching from one state at time t to another at time $t + 1$ is given as

$$\boldsymbol{\Gamma} = \begin{pmatrix} \gamma_{11} & \cdots & \gamma_{1N} \\ \vdots & \ddots & \vdots \\ \gamma_{N1} & \cdots & \gamma_{NN} \end{pmatrix}, \tag{3.3}$$

where N is the number of states, $\boldsymbol{\gamma} = \{\gamma_{ij} : i, j \in \{1, \ldots, N\}$,$\gamma_{ij} = \Pr(S_{t+1} = j | S_t = i)$ for $i, j = 1, \ldots, N$. It follows that $\gamma_{ij} \in [0, 1] \quad \forall \quad i, j$ and $\sum_{j=1}^{N} \gamma_{ij} = 1 \quad \forall \quad i$.

(iii) $\mathbf{P}(X_t)$: $(N \times N)$ diagonal matrix corresponding to the observation process given as

$$\mathbf{P}(X_t = x_t) = \begin{pmatrix} f(X_t = x_t | S_t = 1) & \ldots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \ldots & f(X_t = x_t | S_t = N) \end{pmatrix}. \tag{3.4}$$

The general likelihood function, $\mathcal{L}$, of an HMM is then given by

$$\mathcal{L} = P(X_t = x_t \mid \boldsymbol{\theta}) = \boldsymbol{\delta}\boldsymbol{\Gamma}\mathbf{P}(x_1) \ldots \boldsymbol{\Gamma}\mathbf{P}(x_T)\mathbf{1}, \tag{3.5}$$

where $\mathbf{1}$ is a column vector of length N with all entries equal to 1 and $\boldsymbol{\theta} = \{\boldsymbol{\delta}, \boldsymbol{\gamma}, \mathbf{P}(x_t)\}$, are parameters to be estimated.

Estimating the likelihood of an HMM is done via maximum likelihood estimation using numerical optimisation, and the most likely state sequence is obtained using the Viterbi algorithm (Zucchini et al. 2016). Numerical maximization requires specifying initial parameter values to start the maximization. Although there is no rule of thumb in

specifying starting values to use, we followed the guides discussed in Michelot et al. (2016) to specify initial starting values for the model parameters. We used the R package momentuHMM (McClintock & Michelot 2018) for fitting HMMs to the visual tracks since the package is widely used by scientists studying animal movement.

3.2.1 State decoding in HMMs

In addition to estimating the model parameters in HMMs, the underlying hidden states can be decoded either locally or globally. The local state decoding determines the most probable state for each time $t \in \{1, \ldots, T\}$ given an observation and is expressed as

$$i_t^* = _{i=1,\ldots,N} \Pr(S_t = i \mid X_t = x_t), \tag{3.6}$$

where

$$\Pr(S_t = i \mid X_t = x_t) = \frac{\Pr(S_t = i, X_t = x_t)}{\Pr(X_t = x_t)} \quad \text{and} \quad N \quad \text{is the number of states.}$$

The global decoding, on the other hand, determines the most likely sequence of hidden states $s_1, \ldots, s_T$ that maximizes $\Pr(S_t = i \mid X_t)$. This is achieved using the Viterbi algorithm where the most likely state sequence is obtained recursively from (Zucchini et al. 2016)

$$i_T = _{i=1,\ldots,N} \Pr(S_T = i \mid X_T = x_T). \tag{3.7}$$

In this project, the global state decoding is used instead of the local state decoding. This is because in animal movement modelling, the decoded sequence of states is preferred over the decoded state at each separate time t.

3.3 HMMs in animal movement modelling

In animal movement modelling, the observation process, X_t, corresponds to animal movement represented with GPS tracks, and the hidden state process, S_t, is the animal's behavioural states. The observation (tracking) data are often longitudinal and latitudinal locations. The locations can be converted into metrics of step lengths and turning angles that can tell us something about the animal's behaviour. The step length is the distance between two consecutive locations, while the turning angle is the directional angle between three consecutive positions.

Each observed data point, X_t, is bi-dimensional, consisting of the step length (km), r_t, and the turning angle (radians), ψ_t. At each time t, the distribution of X_t is conditional

on the current latent state, S_t, such that

$$f_j(X_t) = \Big(f(r_t \mid s_t = j), f(\psi_t \mid s_t = j) \Big) \quad \text{for} \quad j = 1, \cdots, N, \tag{3.8}$$

where N is the number of states, r_t is modelled from a gamma distribution with parameters mean, μ, and standard deviation, σ i.e., $r_t|(S_t = j) \sim Gamma(\mu_j, \sigma_j)$. The turning angle is from a von Mises distribution with mean, ρ, and concentration, κ as parameters, i.e., $\psi_t|(S_t = j) \sim von - Mises(\rho_j, \kappa_j)$. We assume the distributions are independent for each time t, conditional on the underlying state S_t.

We note that HMMs are suitable for fitting the tern visual tracking data and simulated data described in Chapter 2. The tern data are collected at a regularly spaced interval, while the simulated tracks are obtained with a regular time interval. Additionally, for each time, t, we specify $N = 2$ discrete states corresponding to foraging ($S_t = 1$) and not-foraging ($S_t = 2$) for the tern tracks. Similarly for the simulated tracks, we specify $N = 3$ discrete states corresponding to travel ($S_t = 1$), search ($S_t = 2$), and forage ($S_t = 3$).

3.3.1 momentuHMM R package for HMMs

HMMs have been widely applied in animal movement modelling due to the availability of HMMs software designed for animal movement modelling, such as the R packages moveHMM (Michelot et al. 2016) and momentuHMM (McClintock & Michelot 2018). We use the latter for HMM analysis in this thesis since it is an updated version of the former. Illustrating with the tern data described in Chapter 2, we provide a brief description of how the package is used. The columns of the data considered in the analysis are the track IDs, longitude values, latitude values and species type. The proxy movement data are processed before fitting HMMs to them. In fitting HMMs using momentuHMM, step lengths and turning angles metrics are obtained from the longitude and latitude values using the "preData" feature. In case of missing data, momentuHMM handles them by replacing these points with the closest non-missing values. momentuHMM fits HMMs to the processed data using the "fitHMM" feature. It does this using the numerical maximization of likelihood explained in Section 3.2 with the help of the nlm optimiser. Therefore, starting values are required to be specified for model fitting. For each state considered (e.g., foraging and travel), initial values are specified for the parameters of the distribution of the step length and turning angles. The gamma distribution is chosen for the step length and requires specifying four initial values: 2 corresponding to the mean (μ) in foraging and travel states, and

2 corresponding to the standard deviation (σ) for both states. An additional starting value is specified for the zero mass parameter (zm) (where available) in both states if there is an exact step length of 0 in the data. This needs to be specified because values of 0 cannot be modelled with the gamma distribution. As a result, the zero mass gives the proportion of step lengths that are exactly equal to zero. For the turning angles, the von-Mises distribution requires specifying four initial values: 2 for the angle mean parameter (ρ), one for each state: foraging and travel, and 2 for the angle concentration parameter (κ), one for each state. The gamma and von-Mises distributions are considered because they are commonly specified for the step lengths and turning angles, respectively, in animal movement modelling (Patterson et al. 2009, McClintock et al. 2017, Akeresola et al. 2024). Multiple starting values are specified to find the global maximum likelihood and behavioural state decoded as explained in Sections 3.2 and 3.2.1.

Chapter 4

Validating visual tracking data as a proxy for GPS data collection

4.1 Introduction

HMM-inferred behavioural states from telemetry data have not been validated in many previous studies due to the difficulty in obtaining concurrent observed behavioural data. Most animal movement studies obtain data using tracking devices attached to animals. While some species are large enough to easily carry these tracking devices, the small sizes of others pose a hindrance to attaching them. Tracking devices, such as GPS, can become too heavy to carry and thus cause discomfort or interfere with the normal behaviour of small-sized species, thus preventing their use in these situations, for ethical reasons (Perrow et al. 2011). As mentioned in Chapter 1, the knowledge of animal behaviour is useful for conservation management. GPS trackers attached to animals are ideal for obtaining positional data of where the animals have been, and although they do not provide any context as to the animals behaviour at any point, it is possible to infer the behaviour of animals from their movement data based on characteristics of the observed data and known ecology of the species. In situations where it is important to understand space use and identify the behaviours of smaller species (e.g., seabirds-terns), an alternative approach would be to either conduct an aerial survey or boat-based transect survey (Louzao et al. 2009). However, due to issues such as limited encounters with feeding birds for the boat survey, not being able to identify some birds to species level for the aerial survey, and providing less information on behaviour than tracking data for both aerial and boat-based surveys, both surveys may be ineffective in delineating important foraging areas of birds (Energy 2009, Perrow et al. 2011). These limitations led to the development of the visual tracking approach.

The visual tracking method was developed by Perrow et al. (2011) to determine the foraging movements of breeding terns, understand their at-sea foraging behaviour, and identify foraging areas for conservation purposes. Terns are seabirds that feed mainly on fish, of which Arctic, Common, and Roseate terns weigh between 90-150 grams, while Sandwich terns weigh between 210-250 grams (Winkler et al. 2020). When the method was developed, terns were considered too small to carry GPS devices since the standard GPS trackers used at the time weighed above the 3% maximum body weight recommendation (Phillips et al. 2003, Igual et al. 2005, Reynolds et al. 2015). Whilst smaller GPS trackers were available, they were expensive and, as such, only a limited number would be available for any survey. Furthermore, even with appropriately sized trackers, terns are difficult to track remotely. In Perrow et al. (2011), they attempted to track the biggest of the tern species (Sandwich terns) remotely, but the trackers were either removed by the birds at the colony or by the force of the birds diving for food. The visual tracking method is feasible for terns because they forage relatively close to the colony and can be readily tracked visually over short distances. The visual tracking method detailed in Section 4.2 involves using rigid-hulled inflatable boats (RHIB) with an on-board GPS to follow individual terns as they leave the colonies to forage at sea. The skipper operates the RHIB with the aim of closely replicating the movements of terns so that the on-board GPS track can serve as proxy movement data for the individual terns being tracked, while the observers on the boat note down the behaviours of the terns.

Wilson et al. (2014) in their study visually tracked terns with the aim of identifying appropriate areas to designate as marine SPAs in the UK. Seabirds in general are considerably less protected at sea than they are at their breeding colonies and so faced with increasing threats when they are out in the sea. In an attempt to limit the increasing impact of human activities on seabird populations, important foraging areas are being designated as SPAs (Thaxter et al. 2012) based on their fairly consistent foraging areas (Lascelles et al. 2012). The visual tracking data used in Wilson et al. (2014) study were collected by JNCC and are described in Chapter 2. Given the nature of the visual tracking method, the behavioural data of terns recorded by observers on a boat provides a rare opportunity to validate inferred behaviours from animal movement data using HMMs (see Chapter 5). In practice, the movement data are actual GPS tracking data of the animals, unlike the visual tracking method, where the GPS location of the boat is used as movement data of the birds. Therefore, it is important to assess whether the GPS location of the boat is an adequate proxy for the bird's location, since we go on to use the proxy movement data for analysis in Chapter 5.

This Chapter serves as a preliminary to Chapter 5. We examine the validity of the boat

tracks by investigating how well they replicate the movement of tracked individuals, using additional information on the animal's recorded position in relation to the boat, which includes the distance and bearing of the boat to the bird. The additional information is only available for a subset of the visual tracking dataset. We start by giving a detailed description of the visual tracking method in Section 4.2. In Section 4.3, we provide an explanation of the validation procedure, which includes reconstructing an approximate movement tracks of terns, comparing approximate tern tracks to boat tracks, fitting HMMs to the boat tracks and approximate tern tracks, inferring behaviours from fitted models, and comparing how these are similar for boat tracks and approximate tern tracks. We then present the results in Section 4.4 and implications for wider understanding of ecology and conservation in Section 4.5.

4.2 Visual tracking method

The visual tracking data of terns provided by JNCC was collected using the visual tracking method technique developed by Perrow et al. (2011) and detailed in Wilson et al. (2014). The visual tracking of terns was conducted during chick-rearing (June and July) and incubation (early May to mid-June) between 2009 and 2011. Rigid hull inflatable boats used for the visual tracking were operated by different skippers across the study sites (as shown in Figure 2.3). The boats were kept c.50-200m from the terns whilst an individual was tracked to avoid disturbing the birds and affecting their behaviour. Longitude and latitude of the boats were recorded using an onboard GPS device set to a 1-second sampling frequency. Individuals were tracked on return foraging trips from their breeding colony. One observer maintained constant sight of the tracked individual, while another recorded behavioural information.

An ethogram of continuous flight behaviours and instantaneous foraging events was provided by each observer, and the timing of each behaviour was recorded (Wilson et al. 2014). Flight behaviours were categorised by observers as active search, transit search, and direct flight. Direct flight was defined by observers as a clear and consistent direction with fast flight usually returning to the colony with food. An active search was defined as an erratic flight course actively searching for food, which may include instances of diving and surface feeding. It is hypothesised that for a direct flight, terns have a fixed location in view and fly in a clear and consistent direction, whereas for transit search, they may change direction but not erratically to search for food (Wilson et al. 2014). We classify the direct and transit search as observed not-foraging behaviour since they do not contain behaviours associated with foraging, unlike the active search, which we classify as observed foraging behaviour.

The location of each observed behaviour was calculated from the boat's GPS track log. Unique IDs were assigned to the data of individual terns tracked in each colony. In 2009 and 2011, tracking only took place during chick-rearing. In 2010, tracking was conducted during chick-rearing and incubation periods. Figure 2.2 provides an example of visual tracks for the two breeding seasons. The data combined both complete and incomplete tracks of terns. The track of terns was considered complete if individual terns were tracked leaving and returning to the colony. Incomplete tracks were terns that could not be successfully followed back to the colony. Reasons for incomplete tracks could be individuals flying faster than the boat could follow, flying over a physical obstruction that hindered the boat from following, observers confusing the tracked individuals with other terns, or insufficient fuel in the boat (Perrow et al. 2011).

Visual tracks for which (i) a single observed behaviour was recorded throughout the tracking trip and (ii) total tracking time that did not exceed 1 min are omitted because we consider 2-state HMM models and the sampling interval is 1 secs (see Tables 5.2 and 5.1 for summary information). GPS coordinates of the boat expressed in the form of longitude and latitude are subsequently converted into step length (km) and turning angle (radians) metrics. These calculated metrics of the boat potentially provide information about tern behaviour. For example, foraging behavioural activities are typically characterized by slow and tortuous flight, indicating smaller step lengths and low directional persistence in turnings. In contrast, not-foraging behavioural activities are generally characterized by longer step lengths and high directional persistence in turnings (Morales et al. 2004).

4.3 Validation method

Given that the boat followed at a distance usually between 50 and 200m from the tracked terns, we investigate how well boat tracks replicate the actual movement of tracked individuals using additional information on the animal's recorded position in relation to the boat. For a subset of tracks (see Table 4.1) recorded at the Coquet Island colony during the chick-rearing period in 2009, additional data were also collected corresponding to the distance and bearing of the tern from the boat, thus permitting the reconstruction of the (approximate) longitude and latitude location of the tern.

4.3.1 Reconstructing approximate tern location

We describe how we reconstruct the approximate location of the terns that are visually tracked in Table 4.1. The bearing and distance of the boat to the tern were mostly

Table 4.1: Visual tracking data of terns from Coquet colony, 2009.

Tern species	Number of tracks	Number of locations
Arctic	3	9500
Common	3	6274
Roseate	1	6200
Sandwich	6	8847

provided for the observed behaviours corresponding to foraging behaviour, while very few tracks had information on bearing and distance provided for both foraging and not-foraging observed behaviours. Mathematically, let the boat's longitude and latitude position be denoted by Lon_{boat} and Lat_{boat}, and the bearing and distance of the boat to the tern be indicated by "bearing" and "distance", respectively. Given the start point, which is the boat location (Lon_{boat}, Lat_{boat}), the distance and bearing between the boat and the tern, we obtain the destination point, which is the approximate tern location. Then the corresponding tern longitude and latitude (Lon_{tern}, Lat_{tern}) are given by (MTS 2022):

$$Lat_{tern} \approx \arcsin\left(\sin\left(Lat_{boat}\right) \times \cos\left(\delta\right) + \cos\left(Lat_{boat}\right) \times \sin\left(\delta\right) \times \cos\left(bearing\right)\right),$$

$$Lon_{tern} \approx Lon_{boat} + \arctan\Big(\sin\left(bearing\right) \times \sin\left(\delta\right) \times \cos\left(Lat_{boat}\right), \cos\left(\delta\right) - \sin\left(Lat_{boat}\right) \times \sin\left(Lat_{tern}\right)\Big),$$

where δ is the angular distance and is calculated as $distance/R$. The radius of the earth, R, is 6371km.

4.3.2 Assessing visual tracks

Following the reconstruction of the approximate tern tracks, we then assess whether the visual tracks are good proxies for the reconstructed approximate tern tracks. We describe the three-way approach used for the assessment below.

Firstly, we carry out a graphical assessment by plotting (i) the visual tracks recorded from boats and the reconstructed approximate tern tracks on the same graph and (ii) the distribution of step length and angles corresponding to the visual tracks and approximate tern tracks side by side to allow for direct and easy comparison. We carry out these comparisons to determine whether the former can be used as an

approximation for the movement of individual terns.

Secondly, we conduct a qualitative assessment by quantifying the differences between the visual track and the approximate tern track using the metric of average distance. To compute the average distance (in metres), we first use the Haversine formula below to get the distance between each individual location point within the boat tracks and the tern tracks.

Let

$$\begin{aligned}
\Delta Lat &= (Lat_{tern} - Lat_{boat}),\\
\Delta Lon &= (Lon_{tern} - Lon_{boat}),\\
a &= \sin^2\left(\frac{\Delta Lat}{2}\right) + \cos(Lat_{boat}) \times \cos(Lat_{tern}) \times \sin^2\left(\frac{\Delta Lon}{2}\right),\\
c &= 2 \times \arctan(\sqrt{a}, \sqrt{1-a}) \quad \text{and}\\
distance &= R \times c,
\end{aligned}$$

where Lon_{boat} and Lat_{boat} are the longitude and latitude of the boat used in visually tracking the tern, Lon_{tern} and Lat_{tern} are the longitude and latitude of the approximate tern track, a is the square of half the chord length between the points (Lon_{boat}, Lat_{boat}) and (Lon_{tern}, Lat_{tern}), and c is the angular distance in radians between the points (Lon_{boat}, Lat_{boat}) and (Lon_{tern}, Lat_{tern}).

Then we compute the average distance, D, as

$$D = \frac{1}{n}\sum_{i=1}^{n} distance_i$$

where n is the total number of locations in the boat and the approximate tern track.

Thirdly, we carry out the behavioural inference assessment. To do this, we model the visual tracking data and approximate tern tracking data using hidden Markov models to account for the different movement patterns dependent on the (unknown) underlying behavioural states. We then extract the inferred behavioural states from models fitted to both datasets and create a confusion matrix to assess differences and similarities in inferred behavioural states.

4.3.3 Model specification

As mentioned in the introduction, it is important to assess whether the visual tracks of the boat are good proxies for the actual tern tracks because we want to validate HMMs, often used for inferring behaviours from actual tracks, using visual tracks. To achieve this, we examine possible varieties of HMMs that can be fitted by ecologists to actual GPS tracks of animals to infer behaviours without validation. We start with the simplest form of an HMM and then gradually proceed to the complex ones. As for the number of behavioural states, we focus on 2-state HMMs as they are commonly used in the analyses of telemetry data and usually serve as proxies to the foraging and not-foraging behaviours (Morales et al. 2004, Jonsen et al. 2005, Patterson et al. 2017). In this section, we give a detailed description of the seven hidden Markov models fitted to the visual and approximate tern tracks.

Model 0: Complete pooling

Model 0 is considered the base model. The model specifies that the state and observed processes are pooled across the individual visual tracks so that the model parameters, which include the estimate of distribution parameters, the transition probability matrix, and the initial state distribution, are assumed to be the same for all individuals. This implies that there is no individual-level effect on the model parameters and the model does not account for individual variation. The disadvantage of this implication is that the base model is then considered biologically unrealistic since individuals are likely to exhibit differences. However, it is still worth considering the base model (i) because there may be a need to draw out patterns common to all individuals in order to scale individual movement up to the population-level (ii) in case the data do not contain enough information to estimate the parameters in the more complicated models with any precision, and (iii) for easy interpretability with less parameterization. The likelihood of the model is described as

$$\mathcal{L} = \prod_{m=1}^{M} \boldsymbol{\delta}\boldsymbol{\Gamma}\mathbf{P}(x_{m,1}) \dots \boldsymbol{\Gamma}\mathbf{P}(x_{m,T_m})\mathbf{1}, \tag{4.1}$$

where M is the total number of individual tracks, $x_m (m = 1, \dots, M)$ is the observation for each individual track, T_m $(m = 1, \dots, M)$ is the total number of time point for each individual track, $\mathbf{1}$ is the column vector of ones, $\boldsymbol{\delta}$ is the initial distribution, and transition probability matrix $\boldsymbol{\Gamma} = \begin{pmatrix} \gamma_{11} & \gamma_{12} \\ \gamma_{21} & \gamma_{22} \end{pmatrix}$.

Model 1: Partial pooling - observed process

Model 1 specifies a pooling effect on the observed process only and accounts for individual variability within the state process by removing the pooling effect on the state process. This implies that the model parameter— the step length distribution parameters — are assumed to be the same for all individuals. The transition probability matrix parameter, Γ, on the other hand, is assumed to be unique for each individual. The likelihood of the model is described as

$$\mathcal{L} = \prod_{m=1}^{M} \boldsymbol{\delta}\boldsymbol{\Gamma}_m\mathbf{P}(x_{m,1})\dots\boldsymbol{\Gamma}_m\mathbf{P}(x_{m,T_m})\mathbf{1}, \tag{4.2}$$

where $\mathbf{\Gamma_m}$, the transition probability matrix for each individual track, is now dependent on individual m.

Model 2: Partial pooling - state process

Model 2 accounts for individual variability within the observed process by specifying a pooling effect on the state process while removing the pooling effect on the observed process. This implies that the model parameter—transition probability matrix— is assumed to be the same for all individuals, while the step length parameters differ across individuals. The likelihood of the model is described as

$$\mathcal{L} = \prod_{m=1}^{M} \boldsymbol{\delta_m}\boldsymbol{\Gamma}\mathbf{P}_m(x_{m,1})\dots\boldsymbol{\Gamma}\mathbf{P}_m(x_{m,T_m})\mathbf{1}, \tag{4.3}$$

where the initial distribution, $\boldsymbol{\delta_m}$, and the observation process, $\mathbf{P_m}$, are now dependent on individual m.

Model 3: No pooling

Model 3 accounts for individual variation in both state and observed processes by removing the pooling effect on these processes, so that the model parameters are assumed to be different for all individuals. This implies that the transition probability matrix parameter, Γ, and the step length distribution parameters are unique for each individual. Unlike Model 1, the base model, Model 3 recognizes the differences that individual animals may exhibit, and although highly parameterized, it avoids any distributional assumptions about the individual effects (McClintock 2021). The likelihood of the model is described as

$$\mathcal{L} = \prod_{m=1}^{M} \boldsymbol{\delta}_m\boldsymbol{\Gamma}_m\mathbf{P}_m(x_{m,1})\dots\boldsymbol{\Gamma}_m\mathbf{P}_m(x_{m,T_m})\mathbf{1} \tag{4.4}$$

where $\mathbf{\Gamma}_m$, $\boldsymbol{\delta}_m$ and $\mathbf{P}_m$ are now unique for each individual m track.

Model 4: Covariate effect on the state process.

In Model 4, the parameters associated with the observed and state processes are pooled across individual visual tracks while including the Euclidean distance of the boat to the colony as an environmental covariate effect on the state process. The likelihood is similar to (4.1) with the inclusion of the covariate in the transition probability matrix.

In HMMs, covariates $c_1, \ldots, c_p$ are included as an effect on the state transition probabilities using a multinomial logit link (Patterson et al. 2017) given as

$$\gamma_{ij} = \frac{\exp(\eta_{ij})}{\sum_{k=1}^{N} \exp(\eta_{ik})}, \tag{4.5}$$

where

$$\eta_{ij} = \begin{cases} \beta_0^{(ij)} + \sum_{l=1}^{p} \beta_l^{(ij)} c_l & \text{if } i \neq j; \\ 0 & \text{if } i = j \end{cases} \quad \text{for } i, j = 1, \ldots, N. \tag{4.6}$$

For our $N = 2$-state model, with a single covariate, c_l, which is the Euclidean distance of the boat to the colony, we have that

$$\begin{cases} \eta_{12} = \beta_0^{(12)} + \beta_1^{(12)} c_{12}, \\ \eta_{21} = \beta_0^{(21)} + \beta_1^{(21)} c_{21}, \\ \eta_{11} = \eta_{22} = 0, \end{cases} \tag{4.7}$$

where $\beta_0^{(12)}$, $\beta_1^{(12)}$ correspond to the intercept and regression parameter of the covariate when switching from state 1 (foraging) to state 2 (not-foraging). On the other hand, $\beta_0^{(21)}$, $\beta_1^{(21)}$ correspond to the intercept and regression parameter of the covariate when switching from state 2 (not-foraging) to state 1 (foraging).

Model 5: Covariate effect on the observed process.

Similar to Model 4, the parameters associated with the observed and state process in Model 5 are pooled across individual visual tracks while including the Euclidean distance of the boat to the colony as an environmental covariate effect on the observed process. The likelihood is similar to equation (4.1) with the inclusion of the covariate in the step length distribution parameter of the observed process as described in equations (4.5), (4.6), and (4.7). In this model, η_{ij} from (4.6) is now the mean of the step length distribution. As mentioned earlier, the GPS coordinates used to obtain the movement metrics of step length and turning angle are of the boat rather than the bird. As such,

there will be limitations in using this to infer bird behaviours, in particular, that the turning angle of the boat is unlikely to match the turning angle of the bird. So, we restrict the covariate effect on the step length only.

Model 6: Covariate effect on both state and observed process.

In Model 6, the parameters associated with the observed and state processes are unique across individual visual tracks, while including the Euclidean distance of the boat to the colony as an environmental covariate effect on both state and observed processes. For seabird species in the breeding season, distance to colony will typically be a dominant factor in determining movement and spatial distributions at-sea (e.g., Wakefield et al. (2017)), because adults need to return to the colony regularly in order to feed their chicks. It is therefore natural to include the Euclidean distance of the boat to the colony as a potential covariate to describe both the probability of transition between states and the distribution of step length within each state, and the energetic cost of travelling to a particular location from the breeding colony (Wilson et al. 2014). The likelihood is similar to equation (4.4) with the inclusion of the covariate in the transition probability matrix, and step length distribution parameter of the observed process as described in equations (4.5), (4.6), and (4.7).

4.3.4 Confusion matrix

A confusion matrix is an $N \times N$ table with four elements that can be used to evaluate the performance of classification models, where N is the number of classes (Hossin & Sulaiman 2015). In constructing the confusion matrix in Figure 4.1, the ground truth data and the predicted data are needed. The ground truth data are represented by the observed behavioural states (validation data) while the predicted data are represented by the behavioural states decoded from HMMs fitted to the visual and approximate tern tracks. In addition, these states are of two classifications: foraging and not-foraging behavioural (travelling) states.

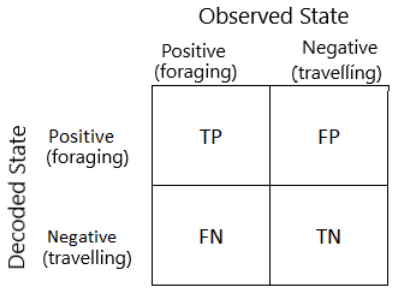

Figure 4.1: Representation of confusion matrix.

The four elements that make up the confusion matrix table are the true positives (TP), true negatives (TN), false positives (FP), also known as type-I error, and false negatives (FN), also known as type-II error. TP are the correctly predicted positive instances, implying that the model predicts a foraging state when the observed state is foraging. FP, on the other hand, are the incorrectly predicted positive instances and imply that the model predicts a foraging state when the observed state is a not-foraging state. TN are the correctly predicted negative instances, meaning the model predicts a not-foraging state when the observed state is not-foraging. FN, on the other hand, are the incorrectly predicted negative instances, which translates to the model predicting not-foraging state when the observed state is foraging.

While the overall accuracy is calculated from these four elements as

$$\text{accuracy} = \frac{\text{Number of TP} + \text{Number of TN}}{\text{Number of TP} + \text{Number of FP} + \text{Number of FN} + \text{Number of TN}}, \tag{4.8}$$

We use the metrics of the true positives and negatives, false positives and negatives to assess the similarity between behaviours inferred from the visual tracks and approximate tern tracks.

4.4 Results

The results reported here are based on visual tracking conducted at the Coquet Island colony during the chick-rearing period in 2009 and summarised in Table 4.1. We compare the boat locations to the associated inferred movement track of nine terns and distributions of the derived step lengths and turning angles. Typical foraging movement patterns generated by the boat and tern for six tracks are provided in Figures 4.2-4.4. The left panels provide the recorded and approximated locations, and the right panels show the associated histograms of the step lengths and turning angles. There are strong similarities between the locations (as would be expected given the boats were following the birds) and step length distributions (though with a slightly larger right-hand tail for the inferred true location). However, there appear to be more substantial differences in relation to the turning angle distributions (lower right panels, Figures 4.2-4.4). The latter difference can be explained by the bird making quicker turns compared to the boat, which has smoother turning movements.

The average distance, which is used to quantify the difference between boats and approximate tern tracks, is reported in Table 4.2. The average distance computed lies within the range of 50-200m for most tracks, except for two of the Sandwich tracks, which have an average distance slightly larger than 200m. The results of the average

distance between the boat and approximate tern track falling within the range of 50-200m are expected because the boat followed at a distance usually between 50 and 200m from the tracked terns as reported in Perrow et al. (2011).

Table 4.2: Quantifying differences between visual track and approximate tern track

Tern species	Track ID	Average distance (m)
Arctic	25	107.42
	82	68.94
	117	78.02
Common	41	112.07
	91	100.65
	107	97.82
Roseate	84	190.36
Sandwich	20	93.24
	21	149.14
	22	128.20
	35	121.96
	37	249.79
	87	206.26

The results of the behavioural inference assessment further validate the boat tracks as an adequate proxy to approximate tern tracks. We fit seven 2-state HMMs described in Section 4.3.3 to the boat and approximate tern location data of 1 Arctic, 2 Common, and 5 Sandwich terns. We then infer foraging and not-foraging behaviour from the fitted models, and observed little difference in the inferred behavioural states when using boat location to approximate the tern location. In particular, the confusion matrix metrics in Figure 4.5 indicate a slightly higher proportion of false negatives for Common tern in Models 2 and 3. However, generally, there is a very low proportion of false positives and false negatives evident across HMMs fitted to both boat and inferred tern location data. Thus, further reassuring that the boat locations are good proxies for the approximate tern location.

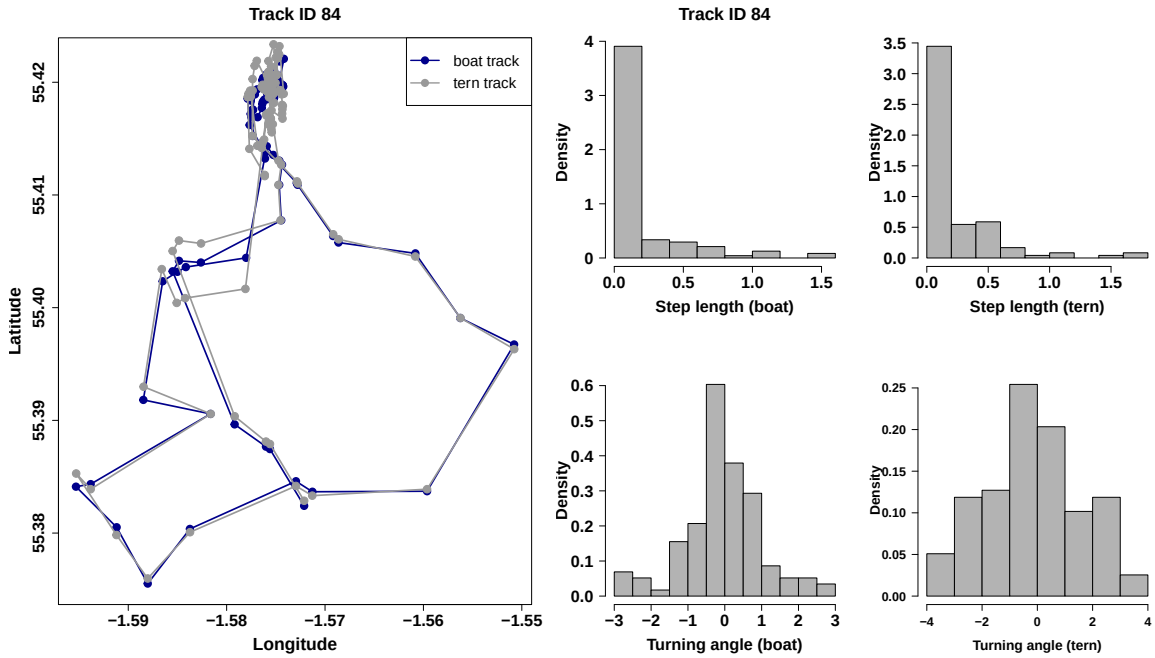

(a)

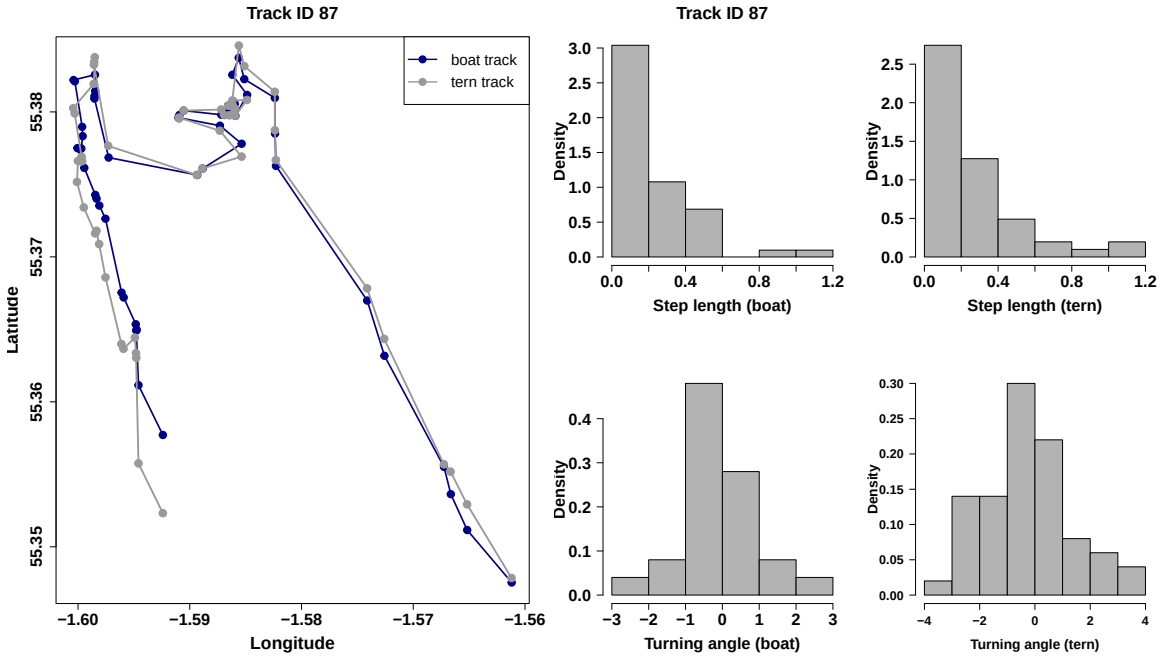

(b)

Figure 4.2: Left panel: Approximate tern tracks and boat tracks of roseate tern (Track ID84) and Sandwich tern (Track ID87). Right panel: Histogram showing the distribution of step length (km) from boat tracks (top left) and approximate tern tracks (top right), and histogram showing the distribution of turning angle (radians) from boat tracks (bottom left) and approximate tern tracks (bottom right) of roseate tern (Track ID84) and Sandwich tern (Track ID87) from Coquet Island, 2009.

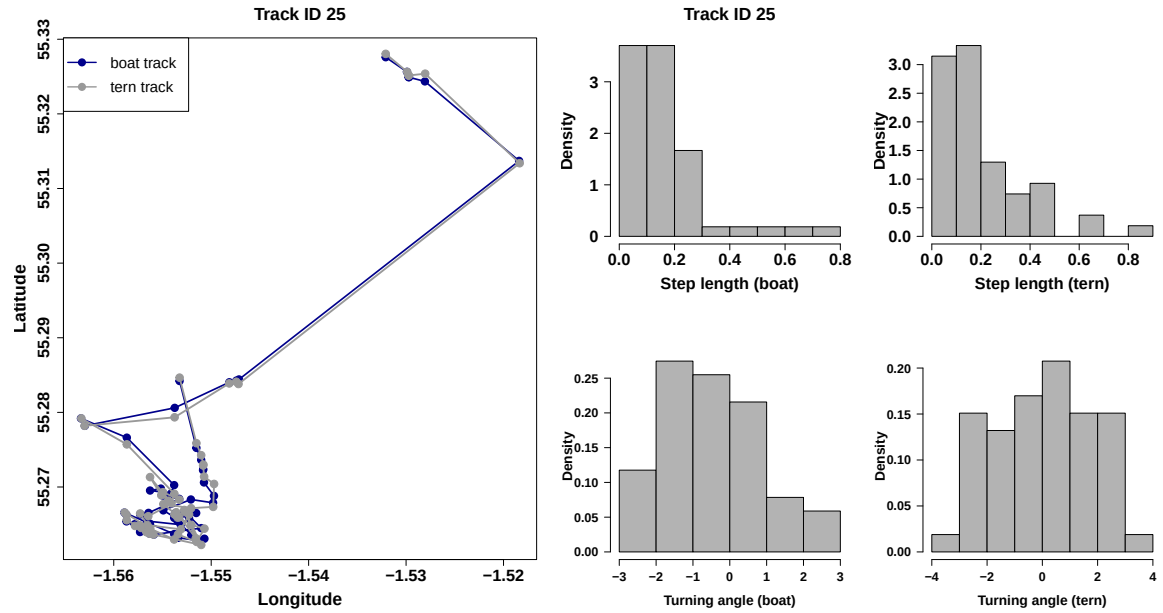

(a)

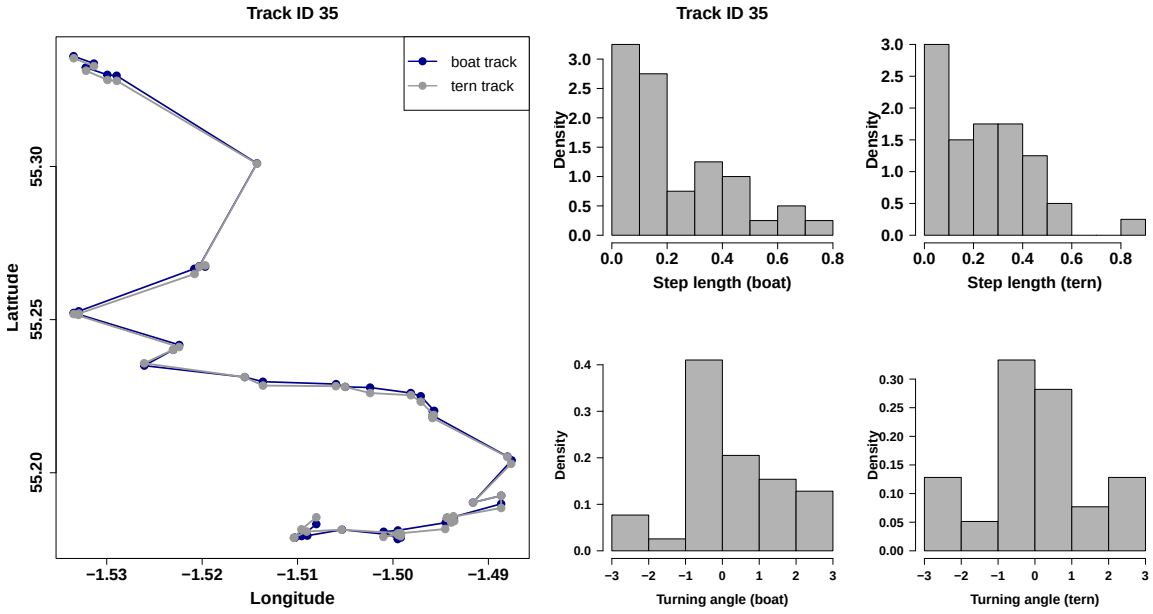

(b)

Figure 4.3: Left panel: Approximate tern tracks and boat tracks of Arctic tern (Track ID25) and Sandwich tern (Track ID35). Right panel: Histogram showing the distribution of step length (km) from boat tracks (top left) and approximate tern tracks (top right), and histogram showing the distribution of turning angle (radians) from boat tracks (bottom left) and approximate tern tracks (bottom right) of Arctic tern (Track ID25) and Sandwich tern (Track ID35) from Coquet Island, 2009.

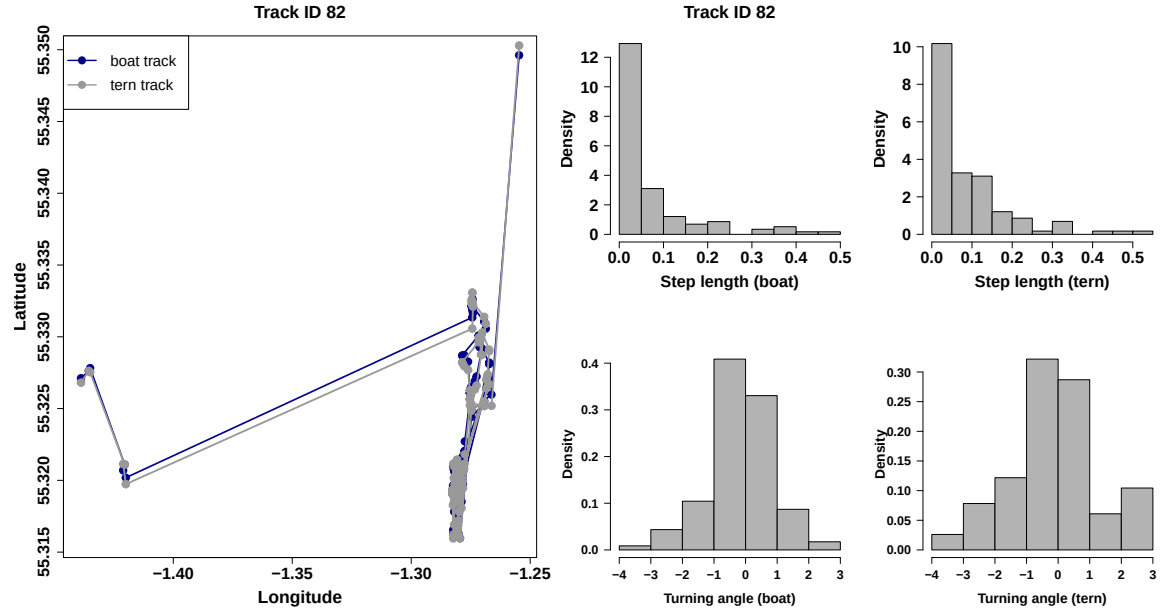

(a)

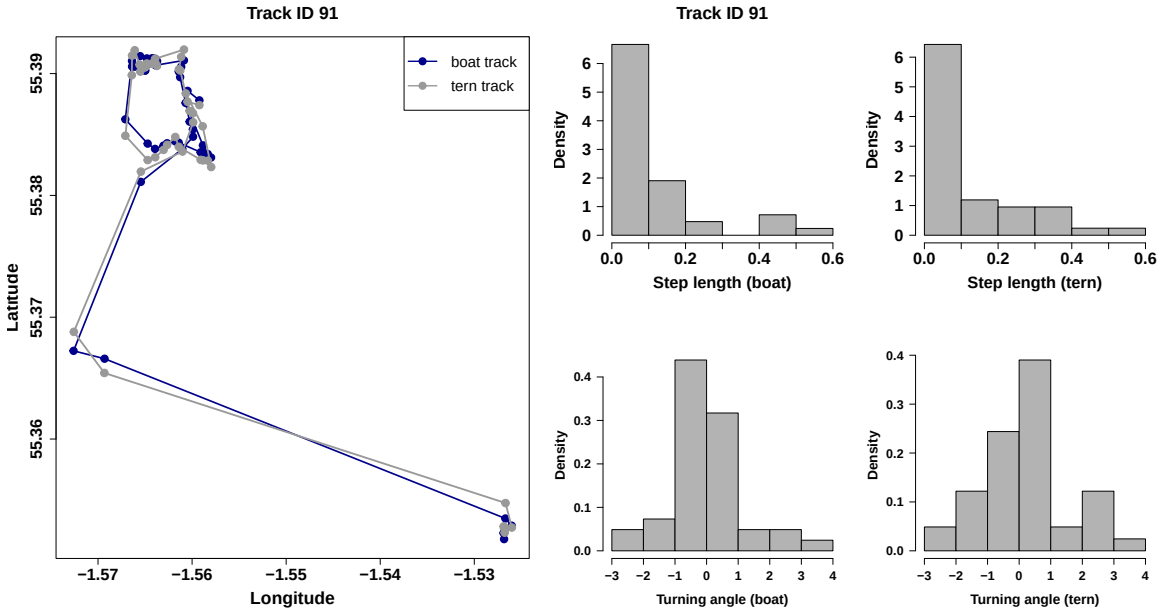

(b)

Figure 4.4: Left panel: Approximate tern tracks and boat tracks of Arctic tern (Track ID82) and Common tern (Track ID91). Right panel: Histogram showing the distribution of step length (km) from boat tracks (top left) and approximate tern tracks (top right), and histogram showing the distribution of turning angle (radians) from boat tracks (bottom left) and approximate tern tracks (bottom right) of Arctic tern (Track ID82) and Common tern (Track ID91) from Coquet Island, 2009.

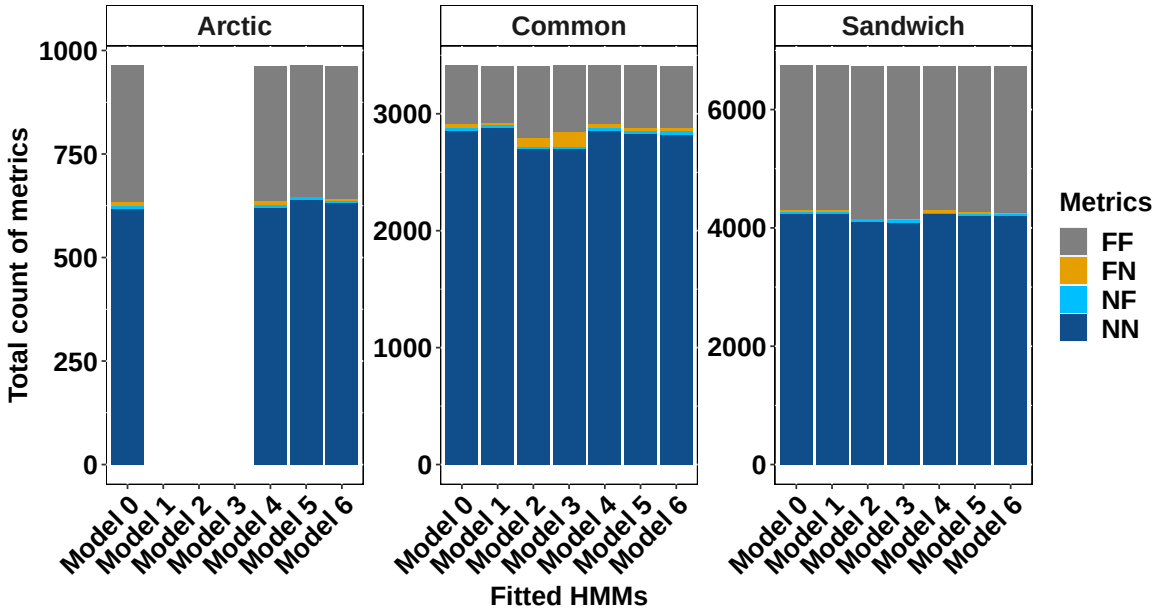

Figure 4.5: Confusion matrix metrics of inferred behavioural states from HMMs fitted to visual tracking and approximate location data of 1 Arctic, 2 common, and 5 Sandwich terns, from Coquet colony during chick-rearing, 2009. FF = true positive, FN = false negative, NF = false positive, and NN = true negative. Positive is represented by foraging behaviour, and negative is represented by not-foraging behaviour. Models 1, 2, and 3 for Arctic tern have no bar because a minimum of 2 terns is required to fit such models.

4.5 Discussion

HMM-inferred behavioural states from telemetry data have not been validated in many previous studies due to the difficulty in obtaining concurrent observed behavioural data. However, these inferred behaviours are used in ecology to delineate important areas, such as those used for foraging, and effective conservation planning and management decisions are taken based on the location of these behaviours. Given the current climate and appetite for increasing the number of protected areas on land and sea globally (e.g., protecting 30% of the earth by 2030 target from the UN Biodiversity COP 15), it is crucial to assess the validity of behaviors inferred from HMMs used in identifying the size and location of essential areas to be protected.

In practice, behavioural states of seabirds are mostly inferred from HMMs fitted to telemetry data (Langrock et al. 2012). We acknowledge that there may be potential effects of the boat following the seabirds on their behaviour, the inferred states, and the validation process itself. However, previous studies have shown that the visual tracking method does not unduly affect bird behaviour due to a reasonable distance

maintained between the individuals and the boat; moreover, most birds appear to ignore the boat. The distance between the boat and the bird was, however, increased when there was a noticeable change in behaviour, such as evasive flight, observed for a few birds (Perrow et al. 2011, Wilson et al. 2014, Robertson, Bolton, Grecian, Wilson, Davies & Monaghan 2014). These previous studies did not investigate the extent to which boat-based tracks replicate the path taken by the birds.

In our study, we show that movement data from the boat being used to visually track terns closely replicated those from the estimated location data of the terns being tracked, particularly for movement tracks corresponding to the foraging behavioural states of terns. Additionally, similar behavioural states of terns were inferred from HMMs fitted to the boat tracks and the corresponding actual (estimated) tern location data. We acknowledge that the boat and approximate tern position were compared for a small number of tern species and restricted to a single colony (Coquet Island) and breeding period (chick-rearing), which may impose limitations on how representative the data and how generalised the interpretation of the results can be. However, there are previous studies where individual terns were tracked visually using a boat with tracks obtained from the onboard GPS as proxies for foraging tracks have been used to successfully identify foraging behaviours and areas of tern species (Perrow et al. 2011, Wilson et al. 2014).

The unique approach of the visual tracking method provides telemetry data for the boat, a proxy for the tracks of the terns they are following, combined with additional behavioural observation data, which are difficult to access in terrains such as the marine environment. Consequently, it allows HMM-inferred behaviours of seabirds to be validated using behavioural observations. From our findings, we can conclude that visual tracking is a suitable method to identify foraging movement and at-sea behaviour of terns, consistent with Perrow et al. (2011). Furthermore, we show that visual tracking provides an effective alternative to telemetry in contexts where attaching biologging devices may not be feasible or appropriate (e.g., in species particularly susceptible to behavioural impacts from the attachment process or devices(Gillies et al. 2020)).

Chapter 5

Validating hidden Markov models for animal behavioural inference

5.1 Introduction

Seabirds are key indicators of marine and coastal environmental health (Parsons et al. 2008, Lascelles et al. 2012) but are the most threatened and anthropogenically pressured group of birds globally (Croxall et al. 2012). Threats, including invasive species at breeding colonies, climate change, over-fishing, and offshore renewable developments, have resulted in a global decline in seabird populations of 70% over the last five decades (Vulcano et al. 2021). In the UK, some species of seabirds (e.g. Northern fulmar (*Fulmarus glacialis*), little tern (*Sternula albifrons*), European shag (*Phalacrocorax aristotelis*), Arctic skua (*Stercorarius parasiticus*), and Black-legged kittiwake (*Rissa tridactyla*)) have continued to decline (JNCC 2021, Stanbury et al. 2021). Of the 25 seabird species that regularly breed in the UK, 24 are listed as Red or Amber on the UK's Birds of Conservation Concern (Stanbury et al. 2021). In addition to threats already being faced by seabirds, the Americas, Europe, and Asia are undergoing outbreaks of High Pathogenicity Avian Influenza (HPAI) sub-type H5N1 in wild birds, which has affected many important breeding colonies of Great Skua (*Stercorarius skua*) and Northern Gannets (*Morus bassanus*) (Ashworth 2022). Coquet Island, Northumberland, currently the UK's only breeding colony of red-listed roseate tern (*Sterna dougallii*), was also reported to have been severely affected by the outbreak (RSPB 2022). Under the Habitats Directive (EC/92/43) and Birds Directive (EC/79/409), Special Protection Areas (SPA) are established to form the Natura 2000 network, which protects species and habitats (European-Commission et al. 2008). Although SPAs have historically been restricted to small areas focused

on seabird breeding colonies, recent extensions and new classifications in the marine environment have expanded the SPA network across the UK (Jim 2020, JNCC 2020). Seabirds are restricted to central-place foraging during the breeding season. Therefore, understanding their at-sea behaviour, including characterising important foraging areas, is vital to ensure adequate protection measures are in place to prevent further population decline.

Seabird tracking studies, where individuals are tagged using biologging technology, are an effective way to understand their space use and behaviour (Lascelles et al. 2012, Wakefield et al. 2017, Bennison et al. 2018). Technological advances over the last 30 years have accelerated the availability of biologging information from devices such as Global Positioning System (GPS) transmitters, accelerometers, conductivity-temperature-depth (CTD) tags, and harmonic radar trackers (Cooke et al. 2004, Cant et al. 2005). Telemetry data provides information on animal locations at discrete intervals but does not provide direct information about the underlying behaviour of the tagged animals. To infer behavioural states such as foraging, flying, and resting from these movement data, hidden Markov Models (HMMs) have been widely used (Morales et al. 2004, Patterson et al. 2009, Langrock et al. 2012, McKellar et al. 2014, Patterson et al. 2017, McClintock 2021). HMMs are time series models with observation and processes where the latent (unobserved) states describe the underlying behaviour of the individual (Langrock et al. 2012). Ancillary data such as environmental information about the landscape the animal is moving through or accelerometery information from instruments on a tagged individual can be integrated with the animal movement data to improve inference of these behavioural patterns (Patterson et al. 2009, McClintock et al. 2017). HMM-inferred behaviours can be used to inform conservation decision-making, for example, the size and location of protected areas.

One limitation of using HMM-inferred behaviours to inform conservation-relevant decision-making is the difficulty in validating models using ground truth data. Some studies have attempted to validate HMM-inferred behaviour from movement data, such as Joo et al. (2013), which validated the behaviour of fishing vessels using ground truth data recorded by onboard observers. Bennison et al. (2018) and Conners et al. (2021) also validated HMM-inferred behaviours of northern gannet (*Morus bassanus*) and albatross using behaviours from depth-recorder and sensors as ground truth data, respectively. However, depth recorders and sensors are also proxies for ground truth data with their own error structures. Overall, little research has focused on evaluating the model performance of HMMs fitted to animal movement data through data validation because contemporaneous behavioural observations on

tracked individuals can be challenging to collect, particularly in hostile or featureless environments, such as open ocean (Joo et al. 2013). To examine the performance of HMMs fitted to movement data, we consider a unique dataset provided by the Joint Nature Conservation Committee (JNCC) described in Chapter 2 and obtained via the visual tracking method of terns (*Sterna* spp.) described in Chapter 4.

First-hand behavioural data of seabirds such as that collected by Wilson et al. (2014) is generally not feasible to collect directly alongside GPS tracking location data. Given that terns do not usually forage far offshore, and direct visual observations are possible and feasible to collect, this provides a unique opportunity where behavioural observations of individuals were collected concurrently with GPS locations of the boat to give proxy locations of the observed individuals. As a result, we consider terns as a case study to examine the performance of HMMs for behavioural inference. To the best of our knowledge, this is the first study to validate HMM-inferred behaviour from movement data using observed behavioural data of seabirds. Our study aims to leverage the rare opportunity provided by the unique JNCC dataset to validate inferred behaviours of seabirds from HMMs using observed seabird behavioural data.

5.2 Data preprocessing

Table 5.1: Summary of visual tracking data of terns across the study sites during chick-rearing period.

Colony, Year	Tern species	Number of tracks	Number of locations
Cemlyn, 2009	Arctic	2	2861
	common	15	23328
Blue Circle, 2010	Sandwich	9	13167
Cockle, 2010	Sandwich	5	13014
Coquet, 2010	Arctic	5	10793
	common	2	2384
	roseate	1	2984
	Sandwich	3	6684
Isle of May, 2010	Arctic	5	3292
Leith, 2010	common	17	34681
Eilean Glas, 2011	common	15	25991
Forvie, 2011	Sandwich	15	38747
South Shian, 2011	common	6	14272

Table 5.2: Summary of visual tracking data of terns across the study sites during incubation period.

Colony, Year	Tern species	Number of tracks	Number of locations
Blue Circle, 2010	roseate	1	9236
Cockle, 2010	Arctic	1	1793
	Sandwich	3	16076
Isle of May, 2010	Arctic	2	8997
Leith, 2010	common	6	15077

5.3 Model specification

GPS locations of the boat were recorded at 1s intervals and do not have missing data (as would be normal in animal-tagged data). The completeness of the data means it is possible to use the recorded positions directly without the need to standardize the recording frequency by interpolating in time and space. Seabirds have been shown to vary their behaviour and area use at different breeding stages, travelling further from the colony to rich foraging grounds during incubation and remaining closer to the colony to feed chicks during chick-rearing (Robertson, Bolton, Grecian, Wilson, Davies & Monaghan 2014). As behaviour is expected to differ between the two time periods, we expect model parameters to differ, reflecting this. Therefore, for each tern species at each colony during the two breeding seasons, we consider a set of different models by varying model parameters, described in Section 4.3.3 and summarised below and in Table 5.3.

Recall from Section 4.3.3:

- Model 0: the base model, specifies that the state and observed processes are completely pooled across the individual visual tracks so that the model parameters are assumed to be the same for all individuals.
- Model 1: assumes a unique transition probability matrix parameter, Γ, for each individual by removing the pooling effect on the state process.
- Model 2: assumes unique step length parameters for each individual track by removing the pooling effect on the observed process across tracks.
- Model 3: Models 1 and 2 are nested in Model 3 since the pooling effect is not specified on both state and observed process.
- Model 4: is an extension of the base model with the Euclidean distance of the

boat to the colony specified as an environmental covariate on the state process.

- Model 5: is an extension of the base model with the Euclidean distance of the boat to the colony specified as an environmental covariate on the observed process in Model 5
- Model 6: assumes unique parameters for both state and observed process, with the Euclidean distance of the boat to the colony specified as an environmental covariate on both processes.

Table 5.3: HMMs (Models 0 - 6) fitted to the boat tracking data of visually tracked terns across study sites during incubation and chick-rearing. Covariate = Euclidean distance of the boat to the study site.

	Pooling effect		Covariate effect	
Models	State process	Observed process	State process	Observed process
0	✓	✓	✗	✗
1	✗	✓	✗	✗
2	✓	✗	✗	✗
3	✗	✗	✗	✗
4	✓	✓	✓	✗
5	✓	✓	✗	✓
6	✗	✗	✓	✓

Model selection was performed using the Akaike information criterion (AIC) (Burnham & Anderson 2002). The AIC value is expressed as

$$\text{AIC} = -2\ln(\widehat{\mathcal{L}}) + 2\text{p}, \tag{5.1}$$

where $\widehat{\mathcal{L}}$ is the likelihood evaluated at the MLE of the model parameters and p is the number of model parameters. We define ΔAIC_i as

$$\Delta\text{AIC}_i = \text{AIC}_i - \text{AIC}_{min} \quad \text{for} \quad i = 0, \ldots, 6, \tag{5.2}$$

such that $\Delta\text{AIC}_i = 0$ for the model deemed optimal.

5.4 Model validation

The validation data consist of the observed behaviours of visually tracked terns. The inferred behavioural states from HMMs and validation data are assumed to be binary classifications: foraging and not-foraging. Common evaluation metrics for binary classification tasks includes accuracy and F1-score from the confusion matrix, area under a ROC curve, and logarithmic loss (Hossin & Sulaiman 2015, Gaurav 2020). Accuracy from the confusion matrix is more appropriate to measure the overall performance when the classification is balanced, as accuracy results can be misleading when one of the classes is more dominant than the other. The F1-score, however, mitigates the effect of imbalanced classification by punishing the extremes values (Hossin & Sulaiman 2015). The area under a ROC curve uses predicted probabilities to determine how well a model can differentiate between classes, while the log-loss metric measures not only the correctness of classification but also accounts for the model's confidence in its predictions (Hossin & Sulaiman 2015, Gaurav 2020). In Chapter 4 of this thesis, we use the metrics of true positives and negatives, false positives and negatives from the confusion matrix in order to assess the similarity between behaviours inferred from the visual tracks and approximate tern tracks. However, in this chapter, we use both log-loss and F1-score metrics to validate behavioural states of visually tracked terns inferred from HMMs, since there is an uneven distribution of behavioural states. In particular, we identified an unbalanced classification of the observed behavioural states of terns for some breeding colonies such as Cemlyn, Isle of May, and Leith. Additionally, the foraging behavioural state is typically of more interest as this helps to identify tern foraging areas. Therefore, a false negative, which fails to identify a foraging behaviour, is equally as important as a false positive, which wrongly identifies a foraging behaviour.

5.4.1 F1-score metrics

To calculate the F1-score metric, we first obtain the precision, also known as the positive predictive value (PPV), which is the proportion of correct positives identified from all the predicted positives, calculated as

$$\text{Precision} = \frac{\text{number of true positive}}{\text{number of true positive} + \text{number of false positive}} \tag{5.3}$$

Similarly, the recall, or true positive rate (TPR), is the proportion of the positives that

are predicted correctly and expressed as

$$\text{Recall} = \frac{\text{number of true positive}}{\text{number of true positive} + \text{number of false negative}} \tag{5.4}$$

A false positive is identified when a foraging behavioural state is decoded by the HMM, and a not-foraging behavioural state is recorded by the observers on the boat. A false negative translates to HMMs decoding a not-foraging behavioural state when the observers on the boat record a foraging behavioural state. Using Equations (5.3) and (5.4), the F1-score is calculated as the harmonic mean of the precision and recall given as

$$\text{F1-score} = 2\left(\frac{\text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}}\right) \tag{5.5}$$

We also report the negative predictive value (NPV), which is the percentage of correct not-foraging behavioural states of all the decoded not-foraging states expressed as

$$\text{NPV} = \frac{\text{number of true negative}}{\text{number of true negative} + \text{number of false negative}} \tag{5.6}$$

Although the F1-score is a good validation metric, it does not account for how close the predicted classification (i.e., the decoded behavioural state) is to the observed behavioural state. To account for this, we consider the logarithmic loss metrics.

5.4.2 Log-loss metrics

The logarithmic loss metric based on probability accounts for the uncertainty in predicted classification (Gaurav 2020). We use the logarithmic loss to account for the uncertainty of the behavioural states inferred from fitted HMMs. At each point, the observed behavioural states y_i, and the predicted probabilities of inferred behavioural state, q_i, are used to calculate the logarithmic loss metric as

$$\text{Log loss}(y, q) = -\frac{1}{n}\sum_{i=1}^{n}\Big[y_i \log(q_i) + (1 - y_i)\log(1 - q_i)\Big] \tag{5.7}$$

where n is the number of observations. The fitted model with the lowest log-loss value is deemed optimal for this criteria, and we report the F1-score, PPV, and TPR corresponding to optimal HMMs.

5.4.3 Foraging event metric

In addition to the validation metrics explained above, we obtain the total number of foraging events identified within each observed behavioural data for visual tracking trips conducted across breeding colonies for each tern species. We define a foraging event as a bout within which only foraging behavioural states are recorded in the observed behavioural data of the individual tracked tern species. We then calculate the proportion of observed foraging events where optimal HMMs correctly infer (i) less than 25% (0% exclusive), (ii) $25 - 49\%$, (iii) $50 - 74\%$, and (iv) at least 75% of foraging behavioural states. Also, we obtain the proportion of observed foraging events completely missed from the foraging behavioural states inferred from optimal HMMs (i.e., observed foraging events where the model infers foraging at 0% of the time points).

5.5 Results

Tables 5.4 and 5.5 present the summarised results of 2-state HMMs fitted to the visual tracking data during incubation and chick-rearing, respectively (see Tables 5.6, 5.8 and 5.9 for detailed results). HMMs with the lowest log-loss value are deemed optimal. The result of the positive predictive value (PPV) shows that the correctly decoded foraging states relative to total decoded foraging states ranged from 65% to 98% during chick-rearing. Additionally, from the result of the true positive rate (TPR), we note that the correctly decoded foraging states relative to total observed foraging states ranged from 70% to 91% except for Arctic terns from Cemlyn and Coquet study sites with 58% and 61%, respectively. Overall, the performance of HMMs in correctly inferring behavioural states during chick-rearing is at least 70% across study sites except for Arctic terns in Coquet, with a percentage of 64%. Validation of HMM results for incubation data shows a low performance compared to models fitted to chick-rearing data in inferring behavioural states. For example, we recorded a minimum of 70% for only one roseate tern visually tracked at the Blue Circle colony during incubation. The overall low performance during this breeding season may be due to the small sample sizes.

Examining the corresponding observed behavioural data for each movement track of the boat, we identified and defined a foraging bout within each track where observed foraging behaviours were recorded as a foraging event. Out of the 626 observed foraging events, foraging behavioural states decoded by optimal HMMs across study sites make up 525 foraging events. 16% of the decoded foraging events captured $< 25\%$ and 25-49% of the observed foraging behaviours within the foraging event, with most foraging

Table 5.4: Summarised validation results of 2-state HMMs fitted to visual tracking data of terns during incubation across 4 study sites. A = Arctic tern, C = common tern, R = roseate tern, S = Sandwich tern, Γ = transition probability matrix, covariate = Euclidean boat distance to the colony.

Incubation		HMM deemed optimal (i.e. based on least log-loss value)		Additional validation metrics (%)		
Colony	Species	Model	Model Description	PPV	TPR	F1-score
Leith	C	0	complete pool	61	60	60
Blue Circle	R	4	covariate on Γ	83	60	70
Cockle	A	5	covariate on step	60	68	64
Cockle	S	6	covariate on Γ and step	59	49	54
Isle of May	A			64	32	43

Table 5.5: Summarised validation results of 2-state HMMs fitted to visual tracking data of terns during chick-rearing in 2010 across the 9 study sites. A = Arctic tern, C = common tern, R = roseate tern, S = Sandwich tern, Γ = transition probability matrix, covariate = Euclidean distance of boat to colony.

Chick-rearing		HMM deemed optimal (i.e. based on least log-loss value)		Additional validation metrics (%)		
Colony	Species	Model	Model Description	PPV	TPR	F1-score
Coquet	C	0	complete pool	88	79	83
Glas Eileanan				84	73	78
Coquet	A	1	no pool on Γ	67	62	64
Isle of May				87	70	78
Blue Circle	S	2	no pool on step	81	78	79
Leith	C			74	70	72
Cockle	S	3	no pool on Γ and step	85	91	88
Coquet				87	75	80
Forvie				65	80	72
Cemlyn	A	4	covariate on Γ	99	59	74
Coquet	R			68	86	76
South Shian	C			70	90	79
Cemlyn	C	6	covariate on Γ and step	72	81	76

events during incubation categorised within this range (Figure 5.1, lower panel). About 24% and 44% of the decoded foraging events accounted for 50-74% and 75-100% of the observed foraging behaviours within the foraging event, respectively, with results during chick-rearing categorised most times in this range (Figure 5.1, upper panel).

The number of observed foraging events completely missed by optimal HMMs (i.e., observed foraging events where the model infers foraging at 0% of the time points) sum to 101, with an average time of 18 secs (Figure 5.2).

The visual tracks coloured with behavioural states (see, for example, Figure 5.3) reveals similarity in the inferred and observed behavioural states across time points within visual tracking trips conducted across breeding colonies. Figure 5.4 provides histograms of the step length and turning angle overlaid with the density curves of the inferred behavioural states for a given track. The inferred states assigned to foraging show shorter step lengths and lower directional persistence in turnings than the not-foraging states, which exhibit larger step lengths and high directional persistence in turnings.

All models fitted appeared to have similar inferred states so that the inferred states were largely insensitive to the set of models considered. However, AIC identified the same, relatively complex model, see Tables 5.6, 5.8 and 5.9, (e.g., an HMM with a relatively large number of model parameters) across many species and breeding colonies, whilst the validation metrics identified much simpler models. During incubation, we observe that the HMM accounting for the Euclidean distance of the boat to the colony as a covariate effect is mostly considered optimal compared to the chick-rearing period. Furthermore, since there are no young terns to look after at the colony during incubation, terns are likely to forage further from the colony during this period compared to chick-rearing. Thus, accounting for the distance of the terns to the colony in HMMs may provide better behavioural inference.

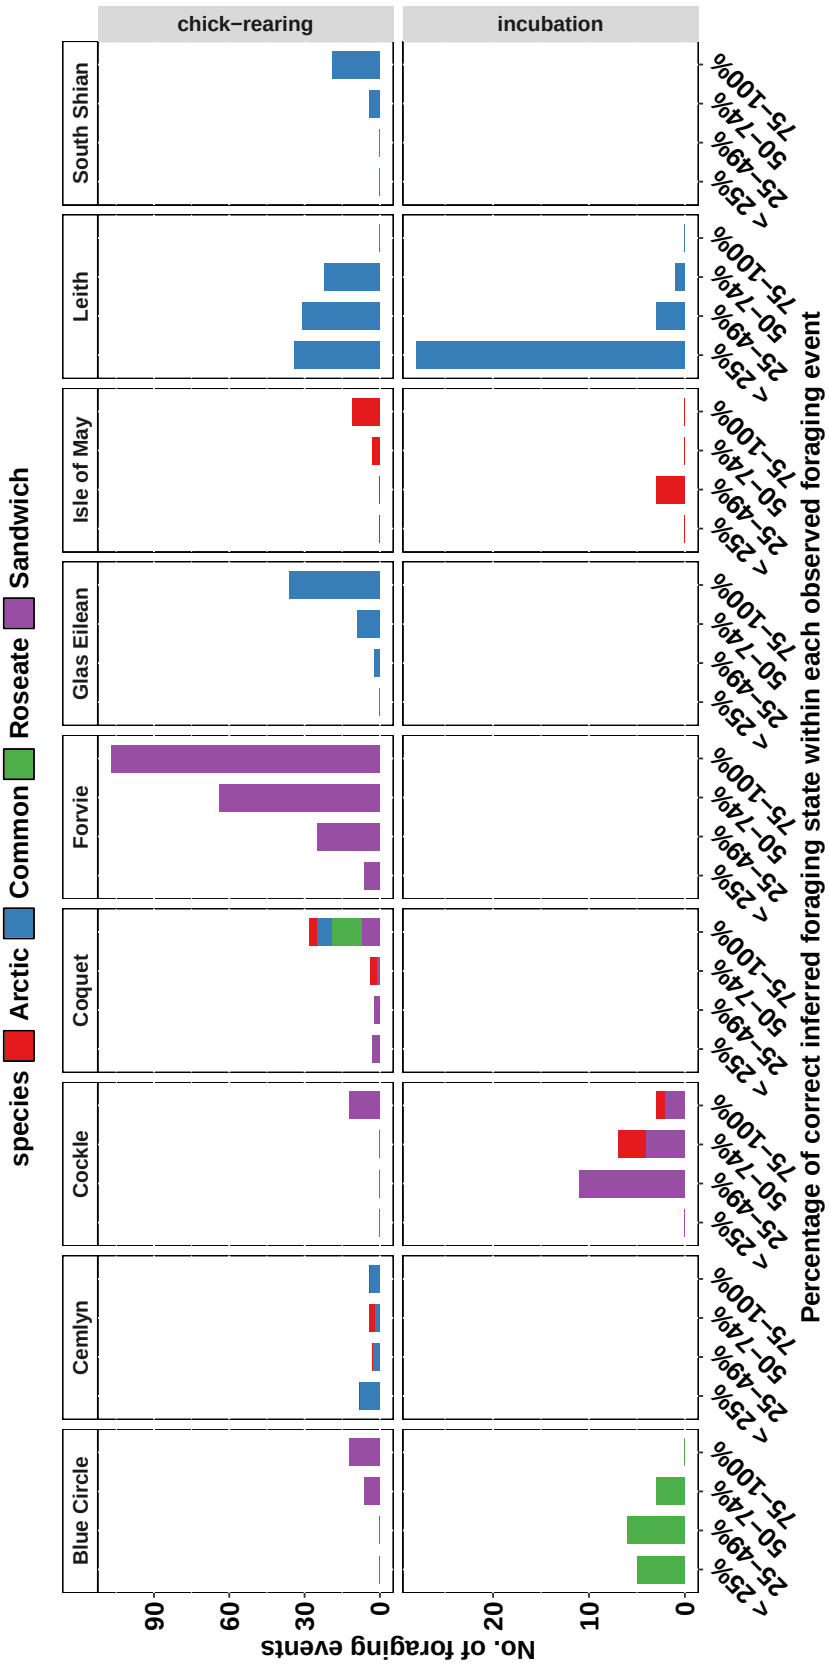

Figure 5.1: Proportion of correctly inferred foraging states within observed foraging events across study sites during chick-rearing and incubation, based on optimal HMMs (lowest log-loss).

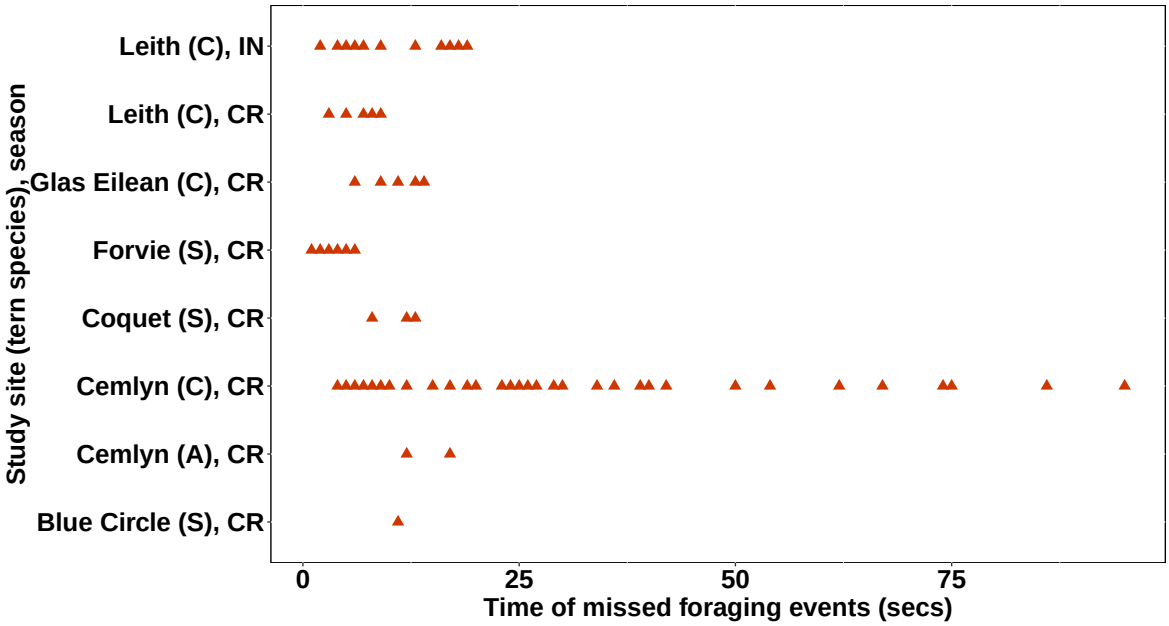

Figure 5.2: Observed foraging events completely missed from inferred foraging events across the study sites during chick-rearing (CR) and incubation (IN). Inferred foraging states are from HMMs deemed optimal based on validation (i.e., HMMs with the lowest log-loss value).

5.6 Discussion

Our study investigated the accuracy of HMMs fitted to visual tracking data from different tern species across breeding colonies in the UK during the breeding season, using behavioural observation data recorded by observers on the boats. Results suggest that HMMs can accurately infer behavioural states from tracking data. A similar observation has been shown for inferred behavioural states from HMMs using additional accelerometer and magnetometer data from four species of albatross (Conners et al. 2021) and fishermen's movement data with frequency differing from the observed behaviours (Joo et al. 2013). The use of additional accelerometer and magnetometer data to validate inferred behaviours are subject to the accuracy of the measurement devices. Our study is the first to validate HMMs using observations of behaviours taken concurrently as the tracking data in the same spatial and temporal context. Generally, HMMs performed reasonably well at decoding behavioural states. However, the performance during incubation was poor compared to chick-rearing, particularly for Arctic terns at the Isle of May (42% see Table 5.4). Terns on the Isle of May had reduced breeding success in 2010. Therefore, terns that were tracked may have included failed or non-breeders which are not required to return to the colony regularly to attend to eggs or chicks, and so the data for this colony and year may be potentially

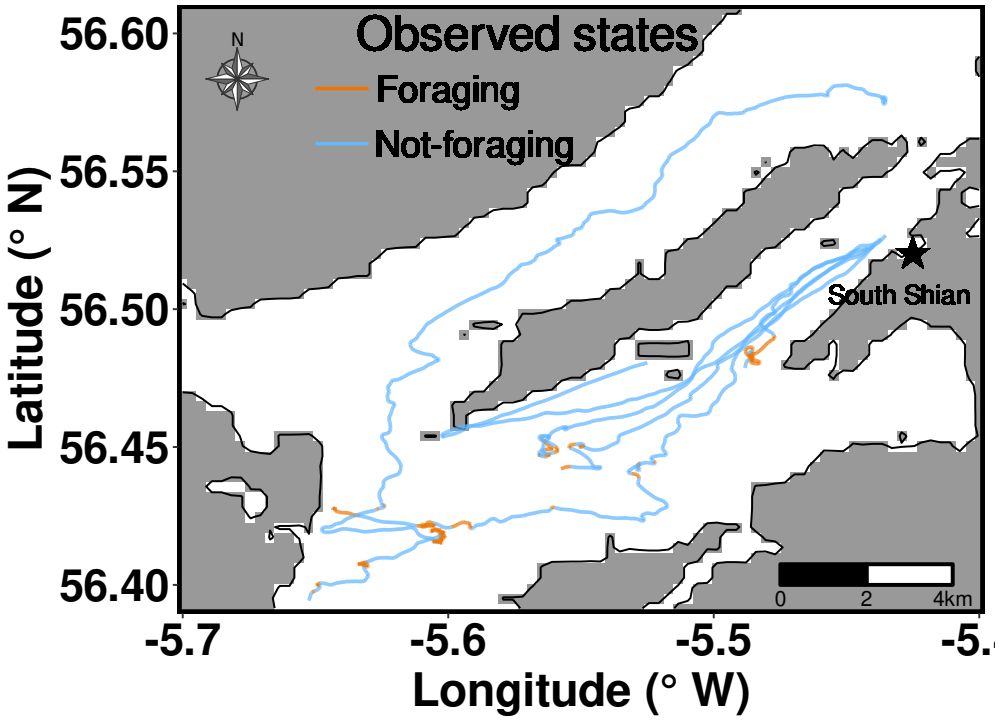

(a)

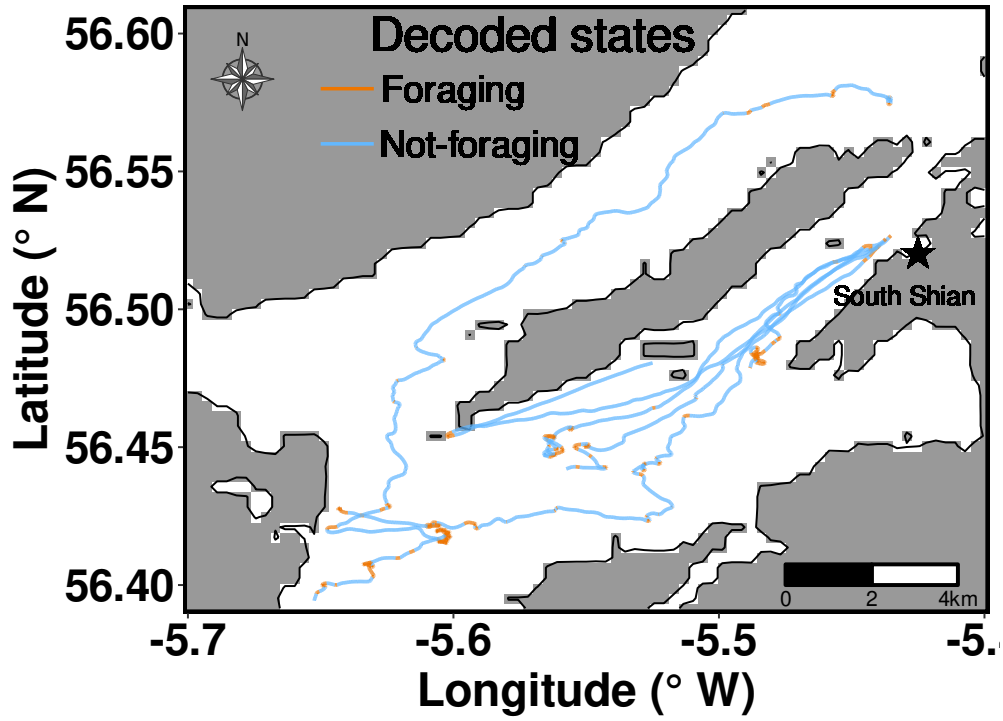

(b)

Figure 5.3: Visual tracks of 6 common terns coloured with (a) observed and (b) decoded behavioural states from South Shian colony during the chick-rearing period, 2011.

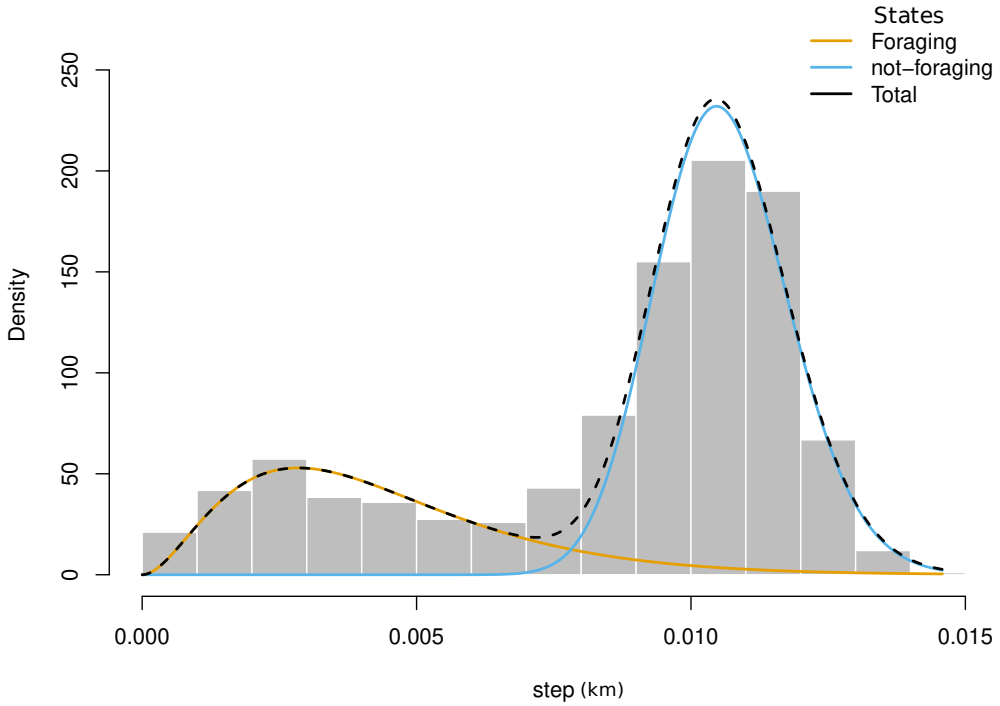

(a)

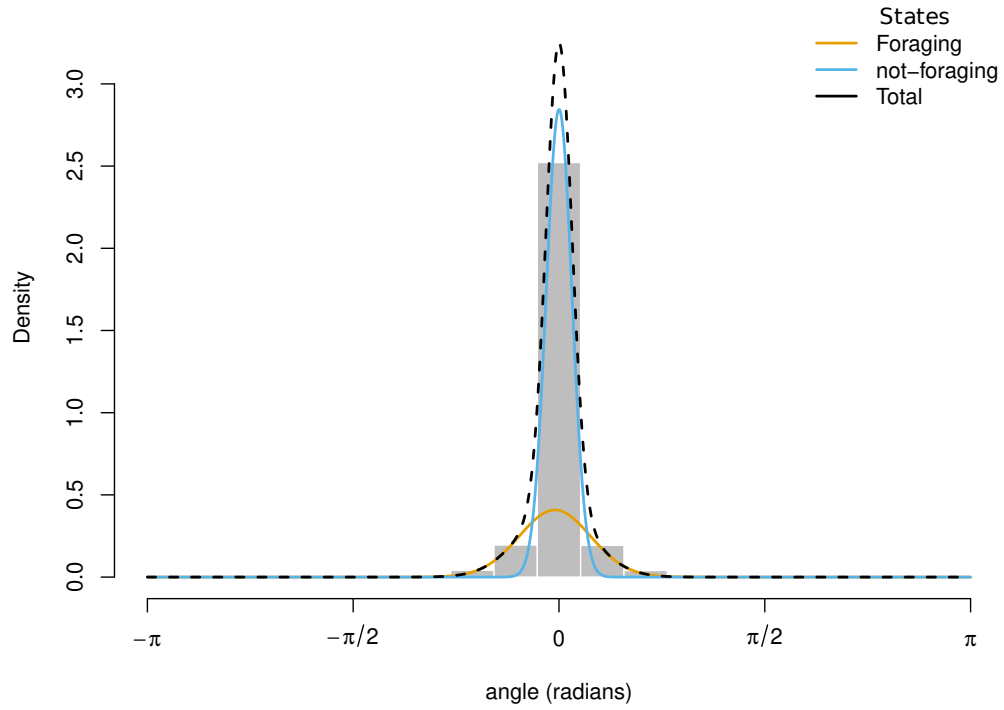

(b)

Figure 5.4: Histograms showing the distribution of (a) step length and (b) turning angle of 6 visually (boat) tracked common terns from South Shian colony during chick-rearing period, 2011. Lines represent HMM-fitted state-dependent distributions and are coloured according to the decoded behavioural states: foraging (orange) and not-foraging (blue).

unrepresentative of breeding adults (Wilson et al. 2014).

The capacity of HMMs in identifying and capturing most foraging behavioural activities within a foraging bout was low for roseate terns at Blue Circle and common terns at Leith during incubation in 2010 (Figure 5.1, lower panel) and common terns at Cemlyn and Leith during chick-rearing (Figure 5.1, upper panel). The visual tracking method was aimed at chick-rearing (2009-2011) but was extended to incubation in 2010, resulting in a reduction in the frequency of data collection (through survey effort being split between time periods) (Wilson et al. 2014), which may be a potential reason why the HMMs fitted less well during the incubation period. Observed behavioural data showed that common terns at Leith colony foraged closer to the colony during chick-rearing, 2010 (Figure 2.2b). The Leith common tern colony is in a port, so there may have been speed restrictions on the boat as well as limitations to how well the boat could closely replicate the movement of the terns. It is unclear from our study the exact reason why fitted HMMs did not identify most foraging behavioural states of common terns within foraging events at Leith and Cemlyn. However, 70% (Leith) and 81% (Cemlyn) of the foraging behavioural states were decoded correctly from HMMs.

HMMs inferred foraging behavioural states 0% of the time for some observed foraging events that lasted for an average of 18 seconds. These missed foraging events were mostly common in chick-rearing. Terns forage close to the colony during chick-rearing and do not travel for long distances (as they do in incubation) (Eglington & Perrow 2014). Also, observers noted short sessions of foraging behavioural activities of some tracked terns in some colonies (JNCC personal communication). As a result, the track of the boat may not capture tern movement corresponding to these short observed foraging events. Consequently, boat tracks may not have represented the tern's track correctly within those short phases of foraging events. As such, the HMMs fitted to boat tracks from such a scenario could not have decoded foraging states within the foraging bout from the boat tracks.

The choice of the number of behavioural states to fit in HMMs is a major challenge in animal movement modelling, particularly when the goal is to infer behavioural states from telemetry data. AIC tends to select HMMs with a higher number of states but may not correspond to or have a meaningful biological interpretation of the animal being studied. Pohle et al. (2017) provides practical guides in selecting the number of states to fit HMMs. Given a fixed number of states, an additional model selection process may include covariates or consider pooling across individual tracks. However, our study showed that these different models did not lead to any substantial differences between the inferred behavioural states, as identified by McClintock (2021). Therefore,

fitting less complex HMMs may likely outperform complex models in inferring hidden behavioural states from movement data. As such, when behavioural inference is the study's goal, it may be preferable to consider simpler models (i.e., including a smaller number of model parameters) when choosing an appropriate HMM to fit after selecting the desired number of states.

Our findings are informative for conservation management and planning. Seabird colonies are more likely to be included as part of protected area networks due to their aggregated nature and relative ease of delineation than areas used by seabirds at sea, especially for species with large foraging ranges from the colony. Foraging areas are considered important habitats to include within seabird-protected area networks (Lascelles et al. 2016). Thus, foraging behavioural activities can be a focus for future studies looking at using behavioural states to inform conservation and management, such as identifying the optimal size and location of foraging areas around seabird colonies. In addition, our study could be extended to assess how temporal validation translates to spatial validation. The visual tracking data could be used to compare the spatial distribution of behaviours inferred from HMMs with the spatial distribution of observed behaviours to determine the accuracy of foraging areas detected using HMMs with real-world implications for conservation and management.

Our study shows that using HMMs to infer foraging behavioural states can help identify most foraging events correctly, as HMMs decoded foraging activities within observed foraging events. Furthermore, missed foraging events or bouts may be less frequent from HMMs fitted to telemetry data of seabirds as GPS devices attached to seabirds are more likely to capture movement patterns influenced by short foraging behavioural activities that last a short time than HMMs fitted to visual tracking data. Therefore, using HMMs for behavioural inference, particularly the foraging behaviour of seabirds, can aid spatial planning and inform conservation decisions, hence providing a tool for the effective management of the impact of human activities on seabirds and other species.

In summary, using HMMs to infer important conservation-relevant behaviours from telemetry data appears defensible based on our results and can inform the design of designated protected areas. Lastly, there is evidence from our validation study that, given the same number of behavioural states, there may be no substantial differences in the performance of simpler and complex HMMs in inferring behavioural states, even in situations where standard model selection approaches, such as AIC, strongly suggest the use of more complex models.

Table 5.6: Validation results of HMMs fitted to terns visual tracks from Cemlyn, Cockle, and Coquet colonies during the chick-rearing period in 2010. Sp = species, A = Arctic tern, C = common tern, R = roseate tern. nPar = number of model parameter, PPV = positive predictive value (foraging), NPV = negative predictive value (not-foraging), TPR = true positive rate. HMM deemed optimal based on *AIC and ** validation (least log-loss value).

2010 Chick-rearing					Validation metrics				
Colony	Sp	Model	nPar	ΔAIC	PPV	NPV	TPR	F1-score	log-loss
Cemlyn	A	0	13	1066	0.98	0.26	0.58	0.73	12.40
		1	15	1062	0.99	0.26	0.58	0.73	12.41
		2*	19	**0**	0.97	0.25	0.59	0.73	12.75
		3	21	1	0.97	0.25	0.59	0.73	12.73
		4**	15	1120	0.98	0.26	0.58	0.73	**12.38**
		5	19	1599	0.98	0.25	0.55	0.70	13.59
		6	21	991	0.98	0.26	0.58	0.73	12.75
	C	0	13	7308	0.71	0.86	0.79	0.75	6.76
		1	41	7191	0.71	0.86	0.79	0.75	6.74
		2	97	58	0.68	0.86	0.79	0.73	7.60
		3*	125	**0**	0.68	0.86	0.80	0.73	7.56
		4	15	8824	0.71	0.86	0.79	0.75	6.74
		5	19	13169	0.71	0.85	0.78	0.74	6.92
		6**	21	8688	0.71	0.87	0.81	0.76	**6.66**
Coquet	A	0	13	3286	0.66	0.83	0.62	0.64	7.57
		1**	21	3089	0.66	0.82	0.61	0.64	**7.52**
		2	37	233	0.59	0.78	0.49	0.54	9.12
		3*	45	**0**	0.63	0.79	0.50	0.56	8.67
		4	15	3449	0.66	0.82	0.61	0.64	7.63
		5	19	4693	0.66	0.82	0.62	0.63	7.77
		6	21	4676	0.66	0.82	0.62	0.63	7.77
	C	0**	13	140	0.88	0.87	0.79	0.83	**4.18**
		1	15	110	0.88	0.87	0.79	0.83	4.25
		2	19	37	0.88	0.87	0.80	0.83	4.27
		3*	21	**0**	0.87	0.87	0.79	0.83	4.33
		4	15	244	0.88	0.87	0.79	0.83	4.22
		5	19	596	0.88	0.87	0.79	0.83	4.27
		6	21	595	0.88	0.87	0.79	0.83	4.26
	R	0*	11	**0**	0.68	0.90	0.83	0.75	6.67
		4**	13	99	0.68	0.91	0.86	0.76	**6.63**
		5	15	547	0.70	0.89	0.80	0.75	6.90
		6	17	550	0.70	0.88	0.80	0.75	6.90

Table 5.7: Validation results of HMMs fitted to terns visual tracks from Blue Circle and Cockle colonies in 2010, and Forvie colony in 2011 during the chick-rearing period. Sp = species, S = Sandwich tern. nPar = number of model parameter, PPV = positive predictive value (foraging), NPV = negative predictive value (not-foraging), TPR = true positive rate. HMM deemed optimal based on *AIC and ** validation (least log-loss value).

Chick-rearing					Validation metrics				
Colony	Sp	Model	nPar	ΔAIC	PPV	NPV	TPR	F1-score	log-loss
Blue Circle	S	0	13	4139	0.79	0.81	0.77	0.78	6.85
		1	29	4102	0.79	0.81	0.77	0.78	6.85
		2**	61	47	0.80	0.81	0.78	0.79	**6.49**
		3*	77	**0**	0.80	0.81	0.78	0.79	6.50
		4	15	4090	0.79	0.81	0.78	0.78	6.84
		5	19	3860	0.78	0.80	0.77	0.77	7.00
		6	21	3818	0.78	0.80	0.77	0.78	6.98
Cockle	S	0	13	6739	0.87	0.95	0.86	0.86	2.30
		1	21	6709	0.87	0.95	0.86	0.86	2.31
		2*	37	**0**	0.84	0.96	0.91	0.87	2.25
		3**	45	136	0.84	0.96	0.91	0.87	**2.18**
		4	15	6737	0.87	0.95	0.86	0.86	2.30
		5	19	58	0.88	0.94	0.84	0.86	2.34
		6	21	58	0.88	0.94	0.84	0.86	2.34
Forvie	S	0	13	8483	0.63	0.89	0.80	0.71	7.35
		1	41	8205	0.64	0.89	0.80	0.71	7.34
		2	97	146	0.65	0.89	0.80	0.72	7.15
		3**,*	125	**0**	0.65	0.88	0.79	0.71	**7.13**
		4	15	10766	0.64	0.89	0.80	0.71	7.34
		5	19	15588	0.64	0.88	0.79	0.71	7.38
		6	21	10760	0.64	0.88	0.79	0.71	7.26

Table 5.8: Validation results of HMMs fitted to terns visual tracks from Isle of May, and Leith in 2010, and Glas Eileanan, and South Shian colonies in 2011 during chick-rearing. Sp = species, A = Arctic tern, C = Common tern. nPar = number of parameter, PPV = positive predictive value (foraging), NPV = negative predictive value (not-foraging), TPR = true positive rate. HMM deemed optimal based on *AIC and ** validation (least log-loss value)

Chick-rearing					Validation metrics				
Colony	Sp	Model	nPar	ΔAIC	PPV	NPV	TPR	F1-score	log-loss
Isle of May	A	0	13	776	0.86	0.72	0.71	0.77	7.51
		1**	21	761	0.86	0.71	0.70	0.77	**7.50**
		2*	37	**0**	0.84	0.70	0.69	0.76	8.08
		3	45	5	0.84	0.70	0.69	0.76	8.08
		4	15	779	0.86	0.72	0.71	0.77	7.51
		5	19	750	0.86	0.71	0.70	0.77	7.54
		6	21	753	0.86	0.71	0.70	0.77	7.54
Leith	C	0	13	9960	0.74	0.80	0.65	0.69	7.58
		1	45	9812	0.74	0.80	0.64	0.69	7.60
		2**	109	50	0.74	0.82	0.70	0.72	**7.21**
		3*	141	**0**	0.73	0.81	0.69	0.71	7.23
		4	15	9904	0.74	0.80	0.65	0.69	7.58
		5	19	8039	0.74	0.79	0.63	0.68	7.70
		6	21	7988	0.74	0.79	0.63	0.68	7.70
Glas Eileanan	C	0**	13	8929	0.84	0.81	0.73	0.78	**6.04**
		1	41	8338	0.84	0.81	0.72	0.78	6.09
		2	97	182	0.79	0.82	0.76	0.78	6.40
		3*	125	**0**	0.81	0.82	0.75	0.78	6.25
		4	15	9662	0.84	0.81	0.73	0.78	6.04
		5	19	14241	0.83	0.80	0.71	0.77	6.36
		6	21	9579	0.84	0.81	0.72	0.78	6.04
South Shian	C	0	13	4459	0.70	0.96	0.89	0.78	3.91
		1	23	4432	0.69	0.96	0.89	0.78	3.93
		2	43	2	0.68	0.96	0.88	0.77	4.24
		3*	53	**0**	0.68	0.96	0.88	0.77	4.23
		4**	15	4796	0.70	0.96	0.90	0.79	**3.87**
		5	19	8286	0.70	0.96	0.88	0.78	4.04
		6	21	8286	0.70	0.96	0.88	0.78	4.03

Table 5.9: Validation results of HMMs fitted to terns visual tracks from Blue circle, Cockle, Isle of May, and Leith colonies during incubation breeding season. Sp = species. A = Arctic, C = Common, R = Roseate, and S = Sandwich tern. nPar = number of model parameter, PPV = positive predictive value (foraging), NPV = negative predictive value (not-foraging), TPR = true positive rate. HMM deemed optimal based on *AIC and ** validation (lowest log-loss value).

2010 Incubation					Validation metrics				
Colony	Sp	Model	nPar	ΔAIC	PPV	NPV	TPR	F1-score	log-loss
Blue Circle	R	0	13	811	0.83	0.41	0.60	0.70	12.80
		4**	15	782	0.83	0.41	0.60	0.70	**12.78**
		5	19	14	0.82	0.40	0.59	0.68	13.15
		6*	21	**0**	0.82	0.40	0.58	0.68	13.15
Cockle	A	0	13	332	0.60	0.79	0.66	0.63	10.09
		4	15	329	0.61	0.79	0.67	0.64	10.01
		5**	19	11	0.60	0.79	0.68	0.64	**9.77**
		6*	21	**0**	0.60	0.79	0.68	0.64	9.79
	S	0	13	2293	0.59	0.70	0.49	0.53	11.35
		1	17	2093	0.59	0.70	0.49	0.53	11.37
		2	25	224	0.59	0.70	0.47	0.52	11.47
		3*	29	**0**	0.58	0.70	0.47	0.52	11.48
		4	15	2294	0.59	0.70	0.49	0.53	11.35
		5	19	1437	0.59	0.70	0.49	0.54	11.34
		6**	21	1433	0.59	0.70	0.49	0.54	**11.33**
Isle of May	A	0	13	611	0.63	0.65	0.32	0.43	12.12
		1	15	598	0.63	0.65	0.32	0.43	12.12
		2	19	312	0.62	0.65	0.32	0.42	12.12
		3	21	299	0.62	0.65	0.32	0.43	12.12
		4	15	544	0.63	0.65	0.32	0.43	12.12
		5	19	82	0.63	0.65	0.31	0.42	12.06
		6**,*	21	**0**	0.63	0.65	0.32	0.42	**12.05**
Leith	C	0**	13	1071	0.61	0.83	0.60	0.60	**7.94**
		1	23	1045	0.61	0.83	0.60	0.60	7.95
		2	43	22	0.60	0.83	0.60	0.60	8.04
		3*	53	**0**	0.60	0.83	0.60	0.60	8.04
		4	15	994	0.61	0.83	0.60	0.60	7.95
		5	19	555	0.61	0.84	0.61	0.61	7.95
		6	21	476	0.61	0.84	0.61	0.61	7.96

Chapter 6

Subsampling animal movement data

6.1 Introduction

As discussed in earlier chapters, tracking devices are often used to obtain direct information about animals movement in space and time as well as proxy information about animal movement (see Chapter 4). Movement data obtained from tracking devices provide a discrete representation of animal movements since the devices are set to record animals' position at specified time intervals. Modelling animal movement data can be done in discrete-time framework but requires that movement data are either collected at regular intervals, or are subsampled, aggregated, or interpolated to obtain tracks at regular intervals (Johnson et al. 2008). While discrete-time models for animal telemetry data are more intuitive to interpret, inference from discrete-time models is not time-invariant, and so can be misleading when the time-scale of observations does not match the behaviours of animals of interest (McClintock et al. 2014).

The behaviour of animals is an underlying latent component that is assumed to influence the movement of animals. This implies that when sampling the animal's position, these underlying latent behaviours are captured within the sampled movement data. Sampling an animal's movement at a coarse temporal scale may not capture underlying latent behaviours that occur only at a finer scale resolution and vice versa (in cases where the battery life of the telemetry device limits longer tracking duration). Given that the behaviours are inferred from movement data, it is crucial to look into the question of how well the animal's movement data can be collected using tracking devices and also how well already collected movement data can be subsampled, so that little or no information is lost about the underlying latent process that is of interest to the ecological community.

Tracking technologies for recording animal movement data have trade-offs based on tag longevity (e.g., battery power), location sampling frequency, hardware costs, size, and receiver or network range. Depending on the research question and approach, some trade-offs can be balanced or mitigated against. For example, a lower sampling frequency that preserves battery life, allowing for a tag to transmit over long periods, may be adequate for investigating long-term flight patterns of seabirds in the nonbreeding season when short-term provisioning behaviours are not of interest. Fine-scale resolution data requiring a high sampling rate may be preferable for understanding short-term home ranges in central-place foragers during the breeding season (Murgatroyd et al. 2016). However, data collection parameters cannot always be controlled, so subjective decisions taken in processing the data, combined with intrinsic properties within the data, can lead to various issues, such as information loss in terms of behavioural information, and spatial and temporal correlation. Unless these issues are addressed through advanced statistical methodology, results and interpretation of analyses may be limited. For example, missing out on foraging areas due to animal telemetry data lacking resolution suitable for inferring behaviours relevant to foraging.

The definition of low and high-resolution movement data is context-dependent and is determined by the specification of a sampling interval, Δt. For example, the definition of 'low' and 'high' resolution data may vary between species or taxa and by the question of interest, such as 1 location every second (e.g., albatross and petrels, Schoombie et al. (2024)), 1 location every 30 min (e.g., swamp wallabies *Wallabia bicolor*, Fischer et al. (2018)), 1 location every hour or more (e.g., African elephants, Wittemyer et al. (2008)). Since there is no rule of thumb to select Δt and specifying a general standard Δt is not feasible, the choice of sampling intervals should be informed by specific research or management questions, goals, and the type of animal being tracked (Gaylord et al. 2016). Previous studies have examined the effect of GPS sampling interval on examining the social interaction of birds (He et al. 2023), estimating animal distance travel (Marcus et al. 2012, McCann et al. 2021), foraging tracks of African Penguins (Ryan et al. 2004), inferring movement metrics (Schoombie et al. 2024), and inferring behaviours (Gaylord et al. 2016). The latter study in particular showed that a sampling interval of 64-secs was able to distinguish short bouts of running, while a 5-mins interval accurately classified different feeding behaviors (e.g., graze, browse, hay) of Rocky Mountain elk (*Cervus elaphus nelsoni*). The availability of modern data collection technologies has made the generation of a high volume of movement data possible. When the need to reduce the volume of data for analysis arises, ecologists tend to first regularise their data into specific units (e.g., one location every 5 mins) and then

sub-sample it using regularised thinning (e.g., choosing 1 point in 10) (Marcus et al. 2012). Behaviours that may be inferred from movement data largely depend on the extent to which the movement data captures these underlying behaviours. However, with thinning, it is possible to lose information that may be helpful for behavioural inference.

In this chapter, we consider an improved approach to experimental design and subsampling data to infer behaviours that are helpful to delineate important areas for organisms that need to be conserved. The approach, examined from the optic of the Nyquist-Shannon sampling theorem (NSST) (Shannon 1949), helps to inform the choice of sampling frequency for data collection. Additionally, NSST can be used to subsample existing datasets to minimize the loss of information in animal telemetry data. NSST originates from the field of electrical engineering, in particular, information theory and signal processing. The theorem establishes the conditions under which (i) sampling a continuous-time signal to obtain a discrete-time signal and (ii) subsampling a discrete-time signal can be done without losing any information contained in the original signal (Shannon 1949). This means that, under such assumptions, the continuous-time signal can be perfectly reconstructed from the discrete-time signals. Similarly, the latter can also be perfectly reconstructed from the subsampled signal. This theorem also quantifies and minimizes the information loss in cases where the assumptions are not met.

NSST has widespread practical application in engineering and science, such as digital audio and video processing (Lévesque 2014, Devi & Pugazhenthi 2016), digital communication systems (Gabor 1946), medical imaging (Cardoso et al. 2015), and biomedical engineering (Hong & Coté 2024). However, the application of NSST in ecology is minimal, with only a few applications to ecological data such as accelerometer data (Hounslow et al. 2019, Yu et al. 2023) and no current application to GPS data to the best of our knowledge. In the context of animal movement ecology, the theorem can be useful for selecting temporal resolutions in the animal movement trajectories such that there is little or no loss of information about the underlying latent behavioural process. In this work, we show how to apply NSST using simulations of GPS movement data and behavioural data as examples, then apply it to real GPS-based movement data to give guidance on how to sub-sample these trajectories. Furthermore, we then assess how the approach taken to reducing the sampling frequency influences the ability to identify underlying behavioural latent states using HMMs.

6.2 Time and frequency domain representation of animal movement data

Temporal data (e.g., animal movement data), which are data collected with respect to time, can be viewed from two different perspectives: the time and frequency domains (Boashash 2015). The time-domain provides insight into the behaviour of data over time. The frequency-domain, on the other hand, shows how the energy of the temporal data is distributed over a range of frequencies. The frequency domain is the domain where the temporal data are represented as a sum of sinusoidal components, revealing their frequency component and subsequent spectral content (Smith 1999). The representation of temporal data in the frequency domain is useful for manipulating data, such as applying functions, for example (i) filtering, which involves designing filters to pass (i.e., retain) or block (i.e., remove) the information associated to certain frequencies, and (ii) sampling a continuous-time signal to obtain its discrete-time version (Allen & Mills 2004).

Temporal data can be transformed from the time-domain to the frequency-domain using a Fourier transform, and back to the time-domain using an inverse Fourier transform. The Fourier transformation decomposes the signal in the time-domain into real (cosine) and imaginary (sine) parts to yield the frequency-domain representation (Smith 1999, Cerna & Harvey 2000). There are four classes of these transforms, which include the Fourier series, the Fourier transform, the Discrete Time Fourier Transform, and the Discrete Fourier Transform. For a periodic signal, the Fourier series is used if continuous, while the Discrete Time Fourier Transform is used if discrete. On the other hand, when the signal is not periodic, the Fourier Transform is used if continuous, and the Discrete Fourier Transform is used if discrete (Smith 1999). In this thesis, we employ the latter because the animal movement data considered and simulated in Chapter 2, particularly the latitudinal and longitudinal positions of animals (referred to as signals henceforth) are discrete and non-periodic. Implementing the Discrete Fourier Transform (DFT) of a time series with N data points is time-consuming, as it requires $O(N^2)$ steps. However, a more computationally efficient algorithm for implementing the DFT is the Fast Fourier Transform (FFT), which requires just $O(N \log N)$ steps (Smith 1999).

The signals are typically obtained from telemetry devices attached to the animals. These telemetry devices are usually set to record the locations of animals at specified time intervals. The time-domain representation of the signal is typically a graphical plot of the positions over time. For example, Figure 6.1 (left panel) shows the time-domain representation of the longitudinal locations of a kittiwake obtained from the dataset

described in Chapter 2. Section 6.2.1, provides the details of how the frequency-domain representation of the signals (see, for example, Figure 6.1, right panel), are obtained in this chapter.

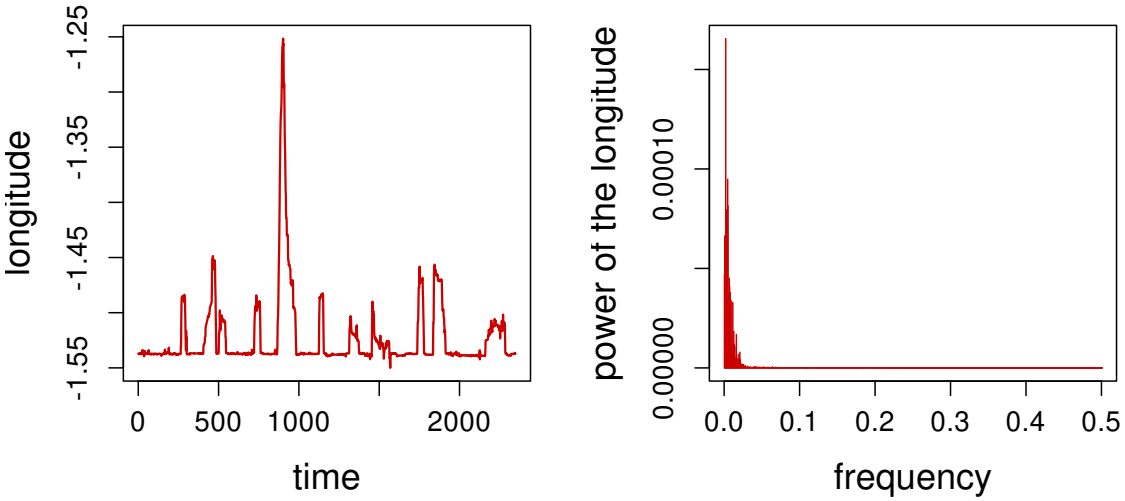

Figure 6.1: Left: Longitudinal movement trajectories of a black-legged kittiwake (ID0066) over time. Right: Corresponding power spectra shown against normalised frequency (0–0.5).

6.2.1 Graphical representation of the frequency domain

Let l_n for time $n = 1, \ldots, N$ be the signal in the time domain with finite length N i.e., N observed locations. As mentioned earlier, we use the DFT to calculate the frequency domain representation of the signal from the time domain, and is described in the equation below (Smith 1999):

$$L_k = \sum_{n=0}^{N-1} l_n \exp\left(- i\frac{2\pi}{N}kn\right), \quad k = 0, 1, 2, \ldots, N-1, \tag{6.1}$$

where

- N is the total number of samples in the signal, i.e., the number of observed locations.
- k is the index of the frequency bin
- $i = \sqrt{-1}$ is the imaginary number
- $\exp\left(- i\frac{2\pi}{N}kn\right)$ is a complex sinusoid that can be expressed as $\cos\left(\frac{2\pi kn}{N}\right) -$

$i \sin\left(\frac{2\pi kn}{N}\right)$

- L_k is the DFT output with the same length as the signal in the time-domain. The DFT output is usually in the form of a complex signal, $X_k + iY_k$. The real part, X_k, and imaginary part, Y_k, are the basis functions of the DFT that correspond to the signal's magnitude and phase, respectively. The real part is expressed as:

$$x_k = \sum_{n=0}^{N-1} l_n \cos\left(\frac{2\pi kn}{N}\right), \quad k = 0, 1, 2, \ldots, N-1, \tag{6.2}$$

while the imaginary part is expressed as:

$$y_k = \sum_{n=0}^{N-1} l_n \sin\left(\frac{2\pi kn}{N}\right), \quad k = 0, 1, 2, \ldots, N-1. \tag{6.3}$$

The real part (6.2) and imaginary part (6.3) are further useful for the graphical representation of the frequency via the power spectrum. The power spectrum is a plot of the portion of a signal's power falling within given frequency bins. In our case, the frequencies are normalised and calculated as:

$$f_k = \frac{k}{N}, \quad 0 \leq f_k < 1, \quad k = 0, 1, 2, \ldots, N-1. \tag{6.4}$$

The signal's power is the energy or strength of the signal per unit time, and is calculated by using the magnitude (i.e., the real part, x) and phase (i.e., the imaginary part, y) of the transformed signal. Let us define the modulus of the transformed signal as

$$z_k = \sqrt{(x_k^2 + y_k^2)}. \tag{6.5}$$

For each frequency, f_k, we can then compute the power of the signal as

$$\text{Power} = \frac{1}{N} z_k^2, \tag{6.6}$$

where N is the number of observations of the transformed signal, i.e., the number of the observed locations.

For a real-valued signal as we have in our case, the DFT output, L_k, satisfies the symmetry property. i.e. only first half of the output ($k = 0, \ldots, N/2$) contains unique information. As a result, the power spectrum produced is also symmetric (Cerna & Harvey 2000). To visualize power spectrum which is the frequency-domain representation of the signal, the plot of the power spectrum is restricted to the unique power values within [0, 0.5] frequencies (see, for example, Figure 6.2, right panel) since

those within [0.5, 1] is symmetric (see, for example, Figure 6.2, left panel).

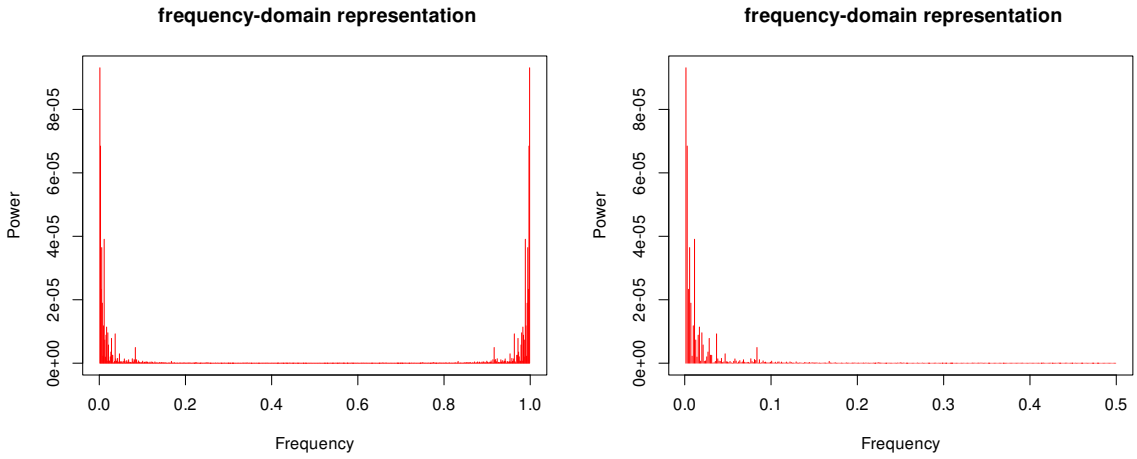

Figure 6.2: Power spectrum of longitudinal movement trajectories shown against normalised frequency: full range (0–1, left) and up to the standard Nyquist range (0–0.5, right).

The fast discrete Fourier transform (fft) R package (R Core Team 2024) is used to obtain the frequency-domain representation of the animal's signal in this chapter. The fft package inputs the signals, and transforms it to produce complex numbers as outputs where the real and imaginary parts correspond to the signal's magnitude and phase, respectively, as discussed earlier. It is also possible to recover the signal, in its respective time-domain, from the frequency-domain using the inverse Fourier transform described as:

$$l_n = \frac{1}{N}\sum_{k=0}^{N-1} L_k \exp\left(i\frac{2\pi}{N}kn\right), \quad k = 0, 1, 2, \ldots, N-1, \tag{6.7}$$

where l_n is the signal in time-domain, L_k is the signal in frequency-domain (i.e., equation (6.1)), and N is the total number of samples in the signal, i.e., the number of observed locations.

Analysis can be carried out in each of these domains depending on the context in which the analysis is being performed. For instance, switching between the frequency and time domains allows data manipulation to be carried out more easily. In the context of animal movement modelling, this work focuses on the valuable insights that can also be derived from frequency-domain analysis, in addition to those from the time-domain. Particularly, we consider the thinning of animal movement dataset.

6.2.2 Standard thinning of animal movement data

Data thinning reduces the number of data points by selecting a subset of the (temporal) data. In animal movement data, the subset is usually obtained by specifying a parameter k, such that every kth location point in the original movement data is extracted. The sampling interval for the extracted subset, Δt_{subset}, is such that

$$\Delta t_{\text{subset}} = k * \Delta t_{\text{original}},$$

where $\Delta t_{\text{original}}$ is the sampling interval of the original movement data to be thinned. We consider these subsets to what might have been obtained had the GPS or tracking device attached to animals recorded a location every Δt_{subset} unit instead of every $\Delta t_{\text{original}}$ unit.

Reducing the number of data points by thinning is sometimes necessary for various reasons. For example, fitting species distribution models is difficult with large datasets, and some movement data collected at high frequency may include redundant information. However, thinning data can lead to loss of information that may be useful for statistical inference (Marcus et al. 2012, McCann et al. 2021). Thinning animal movement data may have consequences on behavioral inference. However, there are some issues that may arise with thinning, particularly when the original movement data contains short durations of certain behaviours. When data are thinned indiscriminantly (i.e., assuming that each data point contributes the same level of information), information critical to the data can be removed so that the thinned data (i) may fail to capture characteristics of interest, or (ii) may remove underlying relevant behaviours within the original movement data.

As mentioned earlier, the time and frequency domains are the two perspectives from which the animal movement data can be viewed, such that whatever is done within one of the domains has consequences for what happens in the other domain. Just as the behaviour of the thinned data differs from the original data in the time-domain representation, its frequency component will also differ from those of the original data when the thinned data is viewed from the frequency domain perspective. Understanding why the thinning of animal movement data may or may not have any negative impact on behavioral inference cannot be achieved from the time-domain representation alone. In this chapter, we show that a look into the frequency-domain representation and analysis within the frequency-domain can help provide a better insight into why this is the case.

6.3 Nyquist-Shannon Sampling Theorem in animal movement modelling

When sampling signals, it is possible to lose information. The amount of information loss largely depends on the choice of the frequency at which sampling is done. The Nyquist-Shannon sampling theorem (NSST) is usually applied to guide how frequently a signal is to be sampled in order to avoid information loss. The theorem states that "If a continuous function x(t) contains no frequencies higher than B hertz, then it can be completely determined from its samples at a sequence of points spaced less than $1/(2B)$ seconds apart" (Shannon 1949), i.e., given a sample rate, f_s, we can reconstruct the signal perfectly when the highest frequency of the signal is less than the Nyquist frequency (i.e., $B < f_s/2$).

In the light of animal movement ecology, the theorem implies that for a given sampling rate, f_s, we can reconstruct the movement data perfectly when the highest frequency in the data is less than the Nyquist frequency, $f_s/2$ (Nathan et al. 2022). In this context, frequency refers to the temporal rate of variation in the movement trajectory, such as changes in latitude and longitude over time. High-frequency components correspond to rapid changes in the trajectory (e.g., sharp turns or abrupt displacements), whereas low-frequency components reflect more gradual or sustained movement patterns. Accordingly, the Nyquist theorem implies that the sampling rate should be at least twice the highest frequency present in the trajectory to ensure accurate reconstruction of movement data (Tatler et al. 2018, Yu et al. 2023). However, identifying this highest frequency in advance is challenging in animal movement studies because it depends on the temporal resolution chosen for data collection. Consequently, NSST cannot be directly used to determine the sampling interval for telemetry data collection, although it can be applied retrospectively when subsampling animal movement data.

When movement data are sampled at a rate lower than twice the highest frequency, important features of the trajectory may be lost or misrepresented, which can in turn distort the underlying behavioural information — a phenomenon known as aliasing (Nathan et al. 2022). To minimise aliasing, signals are often low-pass filtered, whereby frequencies above a chosen cut-off value are removed and lower frequencies are retained. We discuss our approach to apply the NSST framework to animal movement modelling from the perspective of the hidden process (the behavioural information) in Section 6.3.1 and how we use the theorem to subsample animal movement data from the perspective of the observed process (the movement data) in Section 6.3.2.

6.3.1 Exploiting NSST from the hidden process

We use both simulated behavioural states and observed tern behavioural data (collected via visual tracking and used in Chapters 4 and 5) to understand how the choice of sampling interval impacts behavioural inference from the NSST perspective.

Animals that switch behaviours frequently require short sampling intervals (i.e., high sampling frequency) to accurately capture behavioural transitions. Conversely, species with infrequent behavioural changes can be adequately characterised using longer sampling intervals (i.e., low sampling frequency). Selecting an appropriate sampling interval therefore depends on the rate at which behaviours change. The NSST provides a framework for assessing whether a given sampling interval is sufficient to capture these behavioural dynamics. To assess this, we examine the behavioural data (simulated and observed) with the steps highlighted below

1. For each sequence of behaviours, we define a bout as the the period an animal remains in a behavioural state before switching. For each behaviour, we identify all the possible bouts from the behavioural data.

2. For each behavioural state, $i = 1, 2, \ldots, n$, we note the minimum bout duration, t_i.

3. Motivated by the NSST principle, and as a rule of thumb, we choose a new sampling interval, Δt_{new}, such that

$$\Delta t_{new} = min\left(\frac{1}{2}\left(t_1, t_2, \ldots, t_n\right)\right),$$

where n is the total number of underlying behavioural states.

The sampling intervals of the simulated and observed behavioural data are concurrent with the sampling intervals of their movement data. Consequently, we compare the new sampling interval, Δt_{new}, informed by NSST, to the sampling interval, Δt, at which the movement data were originally obtained. We make the comparison with respect to how well the underlying behaviours inferred from the movement data of the two sampling intervals capture the observed behavioural information. We highlight how the results can be described below:

- If $\Delta t_{new} = \Delta t$, then it implies that the sampling interval, Δt, at which the movement data was originally sampled is sufficient.

- However, when $\Delta t_{new} > \Delta t$, it implies that the original sampling interval does not necessarily have to be small and that a less finer sampling interval, Δt_{new}, is

sufficient. We therefore resample both movement and behavioural data at Δt_{new} and compare the behavioural inferences with those obtained using the original interval Δt.

- On the other hand, if $\Delta t_{new} < \Delta t$, then the sampling interval, Δt, at which the data were originally collected is inadequate.

6.3.2 Exploiting NSST from the observed process

As mentioned earlier in Sections 6.2 and 6.2.2, analysis in the frequency domain can provide insight that may not be apparent in the time domain. This section explains how we resample movement trajectories (collected at a fine scale) from the frequency domain in a way that preserves most underlying unobservable behavioural information within these trajectories when there is no access to the first-hand behavioural information of animals being tracked. To achieve this, we examine the movement trajectories from the NSST perspective.

Recall, NSST explains that we can correctly reconstruct a signal from its samples at a sequence of points spaced less than $1/(2B)$ seconds apart, provided there are no frequencies higher than B in the signal. We want to show that NSST approach of thinning is beneficial for minimizing information loss and obtaining a better representation of the original movement trajectories. When we thin, we can understand from the frequency domain perspective via NSST why it is possible to

1. show that thinning with no prior filtering and thinning with prior filtering produce similar results if NSST condition is satisfied (i.e., no high-frequency component in the signal)

2. lose much information if NSST requirement is violated (i.e., the remaining sampling rate of thinned data is below twice the highest frequency in the original data) and no prior filtering is done before thinning, and

3. lose minimal information unlike in (2) due to the low-pass filter applied before thinning.

Using the simulated data, we consider two scenarios to further explain the possibilities highlighted above. Scenario 1 assumes that the movement trajectories contain no high-frequency components. In this scenario, we consider frequencies, f, above 0.25 (which is half of 0.5 in the frequency domain) as high frequencies (i.e., $0.25 < f \leq 0.5$). Scenario 2, on the other hand, acknowledges the presence of high-frequency components in the movement trajectories. It is uncommon for animal movement data (longitude and latitude, hereafter, signals) to lack a high-frequency component when viewed in the

frequency domain. Recall from Section 6.2 we mentioned that the frequency domain shows the distribution of the signal's energy across its frequency components. In order for the signals to satisfy the assumption for Scenario 1, we get a representation of the signals in the frequency domain, and modify its frequency components such that we discard high frequencies higher than 0.25. We do this by setting the energy values corresponding to the frequency values above the 0.25 cut-off to be zero. Following the modification of the signal, we then investigate and show that possibilities (i), (ii) and (iii) holds for Scenario 1. In Scenario 2, we examine signals with some high-frequency component, as we have for most animal movement data. That is, we do not modify the signals prior to analysis. The aim of Scenario 2 is to show that possibility (i) does not hold and that filtering prior to thinning results in minimal information loss compared to thinning without filtering (i.e., (ii) and (iii) hold). We investigate whether filtering the signals in the frequency domain before thinning provides a closer representation of the actual signals than the standard thinning approach, which does not involve filtering.

In both scenarios, filtering before thinning requires that we first filter by selecting a cut-off frequency value (e.g., 0.25, which is half of 0.5 in the frequency domain) to low-pass filter the signal as described in the earlier paragraph, and then recover the filtered signal in the time domain using the inverse Fourier transform. The recovered filtered signal has a length equal to that of the original signal. We then proceed to thin the filtered signal (e.g., by a factor of 2 because a cut-off frequency value of 0.25 is used). On the other hand, the thinning without prior filtering requires that we thin the signals directly. In order to compare how close the filtered-thinned version and the thinned version are to the original signal, we first interpolate both versions (e.g., by a factor of 2) to obtain the same length as the original signal and then quantify the difference between (i) the filtered-thinned-interpolated signal and the original signal, and (ii) the thinned-interpolated signal and the original signal.

So, we compute the average distance (in metres) for (i) and (ii) using the formula below (Kabir et al. 2021):

$$
\begin{aligned}
\Delta\text{Lat} &= (\text{Lat}_{\text{thin}} - \text{Lat}_{\text{original}}),\\
\Delta\text{Lon} &= (\text{Lon}_{\text{thin}} - \text{Lon}_{\text{original}}),\\
\text{a} &= \sin^2\left(\frac{\Delta\text{Lat}}{2}\right) + \cos(\text{Lat}_{\text{original}}) \times \cos(\text{Lat}_{\text{thin}}) \times \sin^2\left(\frac{\Delta\text{Lon}}{2}\right),\\
\text{c} &= 2 \times \arctan(\sqrt{a}, \sqrt{1-a}),\\
\text{distance} &= \text{R} \times \text{c},
\end{aligned}
$$

where

- R = 6371000 metres, is the radius of the Earth
- $\text{Lon}_{\text{original}}$ and $\text{Lat}_{\text{original}}$ are the original longitude and latitude
- Lon_{thin} and Lat_{thin} are either the filtered-thinned-interpolated versions of the longitude and latitude or thinned-interpolated versions of the longitude and latitude.
- ΔLat and ΔLon are the change in latitude and change in longitude, respectively.
- a is the square of half the chord length between the points ($\text{Lon}_{\text{original}}$, $\text{Lat}_{\text{original}}$) and ($\text{Lon}_{\text{thin}}$, Lat_{thin}).
- c is the angular distance in radians between the points ($\text{Lon}_{\text{original}}$, $\text{Lat}_{\text{original}}$) and ($\text{Lon}_{\text{thin}}$, Lat_{thin}).

Following the explanation of the scenarios used to highlight the essence of filtering for the standard thinning approach, we proceed to the second part and examine the impact of NSST on behavioural inference. NSST describes how to subsample a discrete-time signal in a way that there is little or no information loss. The theorem, stated in Section 6.3, explains that we can correctly reconstruct a signal from its samples at a sequence of points spaced less than $1/(2B)$ seconds apart, provided there are no frequencies higher than B in the signal. (Shannon 1949). This implies that B is essential to guide for subsampling. While in most signals, B equals the Nyquist frequency, $f_s/2$, or $B = 0.5$ in the normalized spectrum, in practice, most of the power or energy of the signal is at low frequencies (see Figure 6.3).

A low-pass filter can be applied to the signal to discard high frequencies at a $B' < f_s/2$, where B' is the cut-off frequency value. Then, it is possible to subsample, working at a $f'_s = 2B' < f_s$, i.e., less sample frequency than the original signal. Since the value of the cut-off frequency, B', determines how much information we discard from the signal, we also examine how the choice of B' impacts animal behavioural inference. In particular, how HMM inference is impacted in distinguishing between behavioural states. To achieve this, we (i) calculate high quantiles of the power distribution, which is the same as the percentage of energy distribution preserved, to inform the choice of B' (ii) execute the filtering process with chosen B' and infer behaviours from the filtered signal using HMMs, and then (iii) investigate how sensitive the HMM results are to which of these high quantiles is used for calculating the cut-off frequency value, B'.

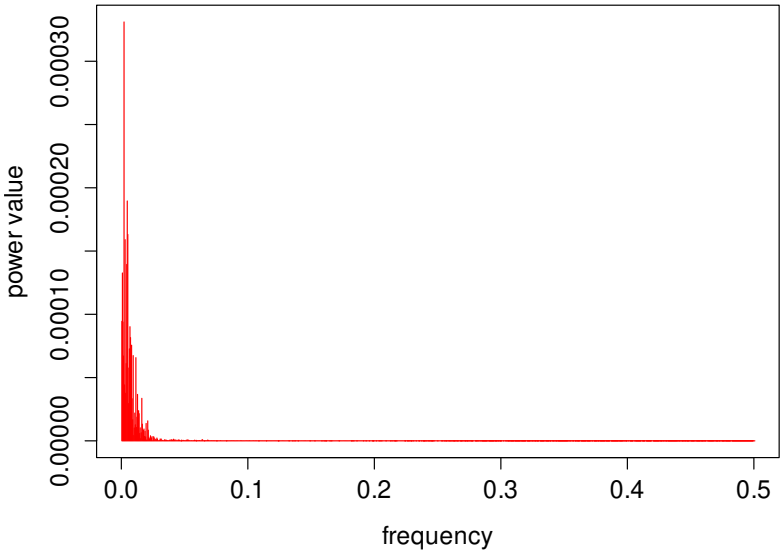

Figure 6.3: Power spectrum of longitudinal movement trajectories of a black-legged kittiwake (ID0066) showing variation across normalised frequencies (0–0.5).

Exploiting NSST: summary of the approach

In summary, we have described how insights from NSST approach, examined from the hidden perspective (i.e., prior knowledge of the underlying behaviour that is often unavailable in practice), can help inform the choice of a sampling interval. We also examined how NSST is used for subsampling discrete-time signals while retaining information, and how filtering prior to thinning provides a better representation of the signals with some high-frequency components. The filtering is particularly useful for discarding high-frequency components from the frequency domain and determines the amount of information retained in the data. In the remaining part of this chapter, we go on to present the simulation study, followed by examples of real movement data used for the case study.

6.4 Simulation study

Movement patterns in tracking data are influenced by underlying behavioural states that are not directly observable. These tracks can be plotted from geographical coordinates retrieved by attaching GPS devices to the animals being studied. Assessing the effect of GPS sampling intervals on behavioural inference via regular thinning and NSST requires that the inferred behavioural states are validated. Since it is rare in practice to access first-hand behavioural data for validation, movement tracks alongside behavioural states corresponding to the movement tracks are simulated. In practice, common behavioural state classifications are travel, search, and forage, with the latter being of more importance to conservation purposes. Therefore, we consider simulation studies using the simulated data described in Chapter 2, and then discuss each category of the simulation. In particular, we present the result of the simulation studies with respect to the use of NSST for the simulated behavioural information in Section 6.4.1 and the use of NSST for the simulated tracks in Section 6.4.2.

6.4.1 NSST from the perspective of the hidden process

This section presents the application of NSST to the behavioural information of the simulated movement tracks described in Section 2.2 and follows the steps highlighted in Section 6.3.1 to inform the choice of a sampling interval for data collection. We present the total number of bouts corresponding to each behavioural activity (i.e., travel, search, and forage) in Table 6.1 for all the simulated tracks. The travel behavioural bout is the period the bird spends in a travel behavioural state before changing to another behaviour, similarly for the forage and search behavioural bouts. Additionally, for each behaviour, the length and time of the bout with the minimum activity are also reported in Table 6.1 for all simulated tracks.

Recall from the NSST approach that sampling at a time interval, Δt, is sufficient to characterize animal behaviour that lasts longer than $2\Delta t$. So, we choose a sampling interval Δt_{new} such that

$$\Delta t_{new} = min\left(\frac{1}{2}\left(t_1, t_2, \ldots, t_n\right)\right)$$

where n is the total number of underlying behavioural states and $t_1, \ldots, t_n$ is the time length of behavioural bout with the minimum activity.

Table 6.1 suggests that a sampling interval, Δt_{new}, required to capture all the behavioural bout will be

$$\Delta t_{new} = 2.5\text{mins}$$

Table 6.1: Observed behavioural bouts from simulated data with Δt = 5-mins. n = number of observed bout for each behaviour.

Track ID	Behavioural bout (n)	bout with minimum behavioral activity		
		Bout length	Time of bout length, t (mins)	t/2 (mins)
ID 1	Travel (2)	8	40	20
	Search (4)	1	5	**2.5**
	Forage (3)	7	35	17.5
ID 2	Travel (2)	22	110	55
	Search (2)	3	15	7.5
	Forage (1)	2	10	**5**
ID 3	Travel (2)	11	55	27.5
	Search (4)	1	5	**2.5**
	Forage (3)	7	35	17.5
ID 4	Travel (2)	16	80	40
	Search (7)	1	5	**2.5**
	Forage (6)	7	35	17.5
ID 5	Travel (2)	32	160	80
	Search (5)	1	5	**2.5**
	Forage (4)	1	5	**2.5**
ID 6	Travel (2)	22	110	55
	Search (5)	1	5	**2.5**
	Forage (4)	7	35	17.5
ID 7	Travel (2)	26	130	65
	Search (4)	1	5	**2.5**
	Forage (3)	13	65	32.5
ID 8	Travel (2)	18	90	45
	Search (3)	1	5	**2.5**
	Forage (2)	45	225	112.5
ID 9	Travel (2)	19	95	47.5
	Search (6)	1	5	**2.5**
	Forage (5)	6	30	15
ID 10	Travel (2)	14	70	35
	Search (4)	1	5	**2.5**
	Forage (3)	18	90	45

In this case, it appears the NSST suggests that the sampling interval, Δt = 5-mins that has actually been used to simulate the tracking data is insufficient and that a higher frequency (i.e., < 5-mins) is needed. Given the insufficiency, it is possible for the simulated tracks to miss out on the fine details of some behaviours. Consequently,

we fit a complete pooling HMM (see Section 4.3.3) to all the simulated tracks to assess whether there are behavioural bouts that the tracks fail to capture. We choose the complete pooling HMM since the work in Chapter 5 shows that simpler HMMs are sufficient for inferring behavioural states. The inferred behaviours are compared with the simulated behaviours to identify behavioural bouts the decoded behaviours fail to capture in the simulated behaviours. From Table 6.2 we observe that for all simulated tracks, behavioural bouts that last for 5-mins (Δt) and 10-mins ($2\Delta t$) are completely missed as the decoded states fail to capture them. The missed behavioural bouts are mostly search behaviours and include the foraging behaviour as well. The results confirm what the NSST suggests that we may not be able to characterize behaviours that last $\leq 2\Delta t$ with a sampling interval of Δt. Although the performance (F1-score, see Section 5.4) of the HMM in decoding the traveling, foraging, and search behavioural states is 97.57%, 95.82 %, and 87.85%, respectively, there is room for improving the accuracy if the missed bout can be captured, particularly for the search behavioural state. Using one of the simulated tracks as an example, Table 6.3 gives a further breakdown of all the simulated behavioural bouts alongside the total number of simulated behaviours captured and missed by the decoded behaviours (see Figure 6.4 for the decoded behaviours).

Table 6.2: Number (n) of missed behavioural bouts from 3-state HMMs fitted to simulated data with Δt = 5-mins. s= search, f = forage.

	Track IDs									
bout length	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
5-mins	s(1)	-	s(1)	s(1)	s(1) f(1)	s(2)	s(1)	s(1)	s(1)	s(1)
10-mins	-	f(1)	-	-	s(1)	-	-	-	s(1)	-

Table 6.3: Identifying behavioural activity in bouts that are captured and missed by decoded states from a 3-state HMM fitted to simulated data-ID10 with an original sampling interval, Δt = 5-mins. The numbers in the Table translates as: **no. of simulated behaviours within bout**(no. of behaviours identified by decoded state, no. of behaviours missed by decoded state)

	Bouts			
	Bout 1	**Bout 2**	**Bout 3**	**Bout 4**
Behaviour				
Travel	**18**(17,1)	**14**(14,0)	-	-
Search	**6**(6,0)	**14**(13,1)	**4**(3,1)	**1**(0,1)
Forage	**18**(14,4)	**23**(23,0)	**118**(117,1)	-

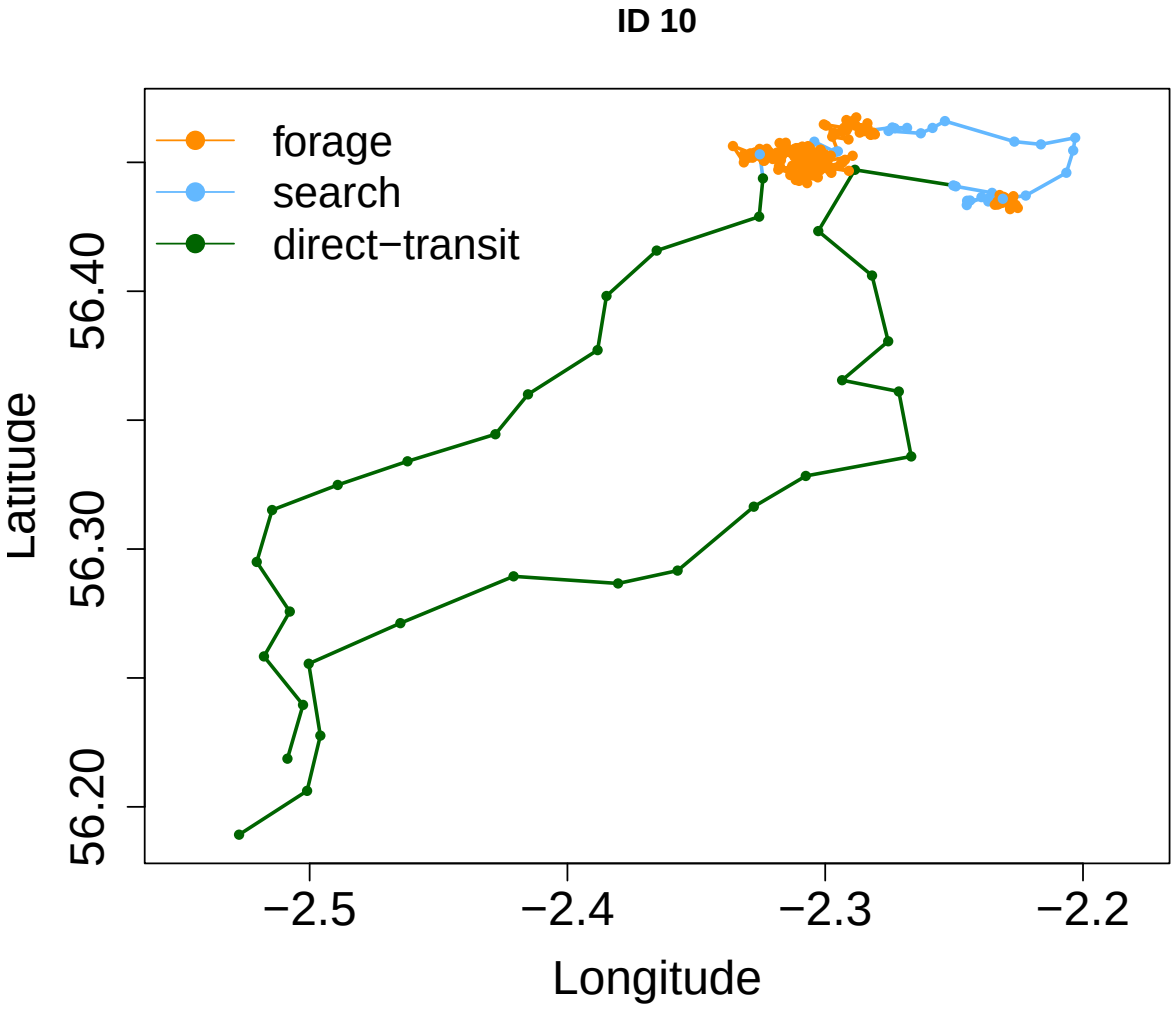

Figure 6.4: Simulated track-ID10 coloured by decoded behavioural states: forage, search and travel (direct-transit), at 5-mins sampling interval.

6.4.2 NSST from the perspective of the observed process

Recall from Subsection 6.3.2 that the aim of Scenario 1 (i.e., signals with no high-frequency components) is to show that thinning signals with and without filtering produce similar results if and only if the signal to be thinned contains no high-frequency components when viewed from the frequency-domain perspective, while the aim of Scenario 2 (i.e., signals with some high-frequency components) is to show that there is minimal loss of information when filtering is done before thinning compared to thinning without filtering. We first provide a graphical illustration of both scenarios, then compare how close both versions of the simulated tracks (i.e., thinning only and filtering prior to thinning) are to the simulated tracks under Scenarios 1 and 2, and finally assess the accuracy of behavioural inference of both versions under the two scenarios. The simulated tracks considered in this section are described in Section 2.2.

Graphical illustration of scenarios 1 and 2

We explain the scenarios in detail with a graphical illustration using one of the simulated longitude locations, $\text{lon}_{original}$, in Section 2.2. In Figure 6.5, the frequency-domain representation of the longitude location, $\text{lon}_{original}$, contains some high-frequency components because the power value (Y-axis) corresponding to the higher frequencies (X-axis) is small but never zero. For Scenario 2, which is a signal with some high-frequency component, we use the simulated locations as it is without any modification in the frequency-domain. However, for Scenario 1, we modify the simulated data to ensure it has no high-frequency components by setting the power values of frequencies within the range of 0.25 and 0.5 to zero. Then, we show that filtering prior to thinning and the standard thinning techniques in scenario 1 provide the same results when compared with the simulated locations. However, the results are different in Scenario 2. The results are quantified by the average error between the simulated locations and the interpolated versions of the filtering before thinning, and the thinning without filtering.

The filtering is done in the frequency domain. We use a low-pass filter by selecting a cut-off frequency value and then discarding frequencies above the cut-off value. The choice of the cut-off value determines and quantifies how much information is discarded. For thinning, we reduce the data points by dropping every kth sample and then interpolate the thinned version to obtain the same number of data points as the simulated data. In this simulation study, we thin by a factor of 2, 3, 4, 5, and 6, and filter with cut-off frequency values of 0.25, 0.16, 0.125, 0.1, and 0.08. We use one of

the simulated tracks, a thinning factor of 2, and a cut-off frequency value of 0.25 as an example to explain the interpolation process. A cut-off value of 0.25 implies that we discard half of the information of the simulated data by setting the power values corresponding to frequencies within the range of 0.25 to 0.5 to zero. Thinning by a factor of 2 implies dropping one out of every two samples to obtain half of the length of the simulated data. A cut-off value of 0.25 is paired with a thinning factor of 2 because we want to be able to directly compare the outcome of (i) filter-thin-interpolate to (ii) thin-interpolate versions

Figure 6.5 (row 2) shows the time and frequency domain representation of the thinned version. For the thinned version to have the same length as the original simulated data, we need to first upsample the thinned version and then interpolate. Upsampling by an integer factor, m, implies that we insert $m-1$ zeros between every data point in the time domain to represent new sample positions that have no data yet. We then apply a low-pass filter to the upsampled version to smooth out the inserted gaps (i.e., the empty samples) by interpolating new, realistic values between the longitude locations. Since we thinned the simulated data by a factor of 2, we will need to upsample by a factor of 2, meaning that we insert just a single zero between every data point (i.e., the location). The frequency-domain representation of the upsampled version of the thinned location in Figure 6.5 has a mirror image of the thinned location's frequency spectrum. This is because of the zeros that we have included (see Liu et al. (2023) for a detailed explanation). To eliminate this, we need to interpolate and apply a low-pass filter to the upsampled version so that the frequency domain plot can be similar to the original simulated location. Therefore, we select a cut-off frequency for the low-pass filter to be 0.25.

After the interpolation, the new frequency domain plot is shown in the right panel of row 4, Figure 6.5. The recovered longitudinal location after interpolation is shown in the left panel of row 4, Figure 6.5. The same process is repeated for the simulated longitudinal and latitudinal location using cut-off frequencies of 0.16, 0.125, 0.1, and 0.08, and thinning by factors of 3, 4, 5, and 6, respectively. We then compute the average distance between the (i) simulated locations and the filter-thin-interpolate versions, and (ii) simulated locations and the thin-interpolate versions, with results presented in the Figures 6.6, 6.7, and Table 6.4.

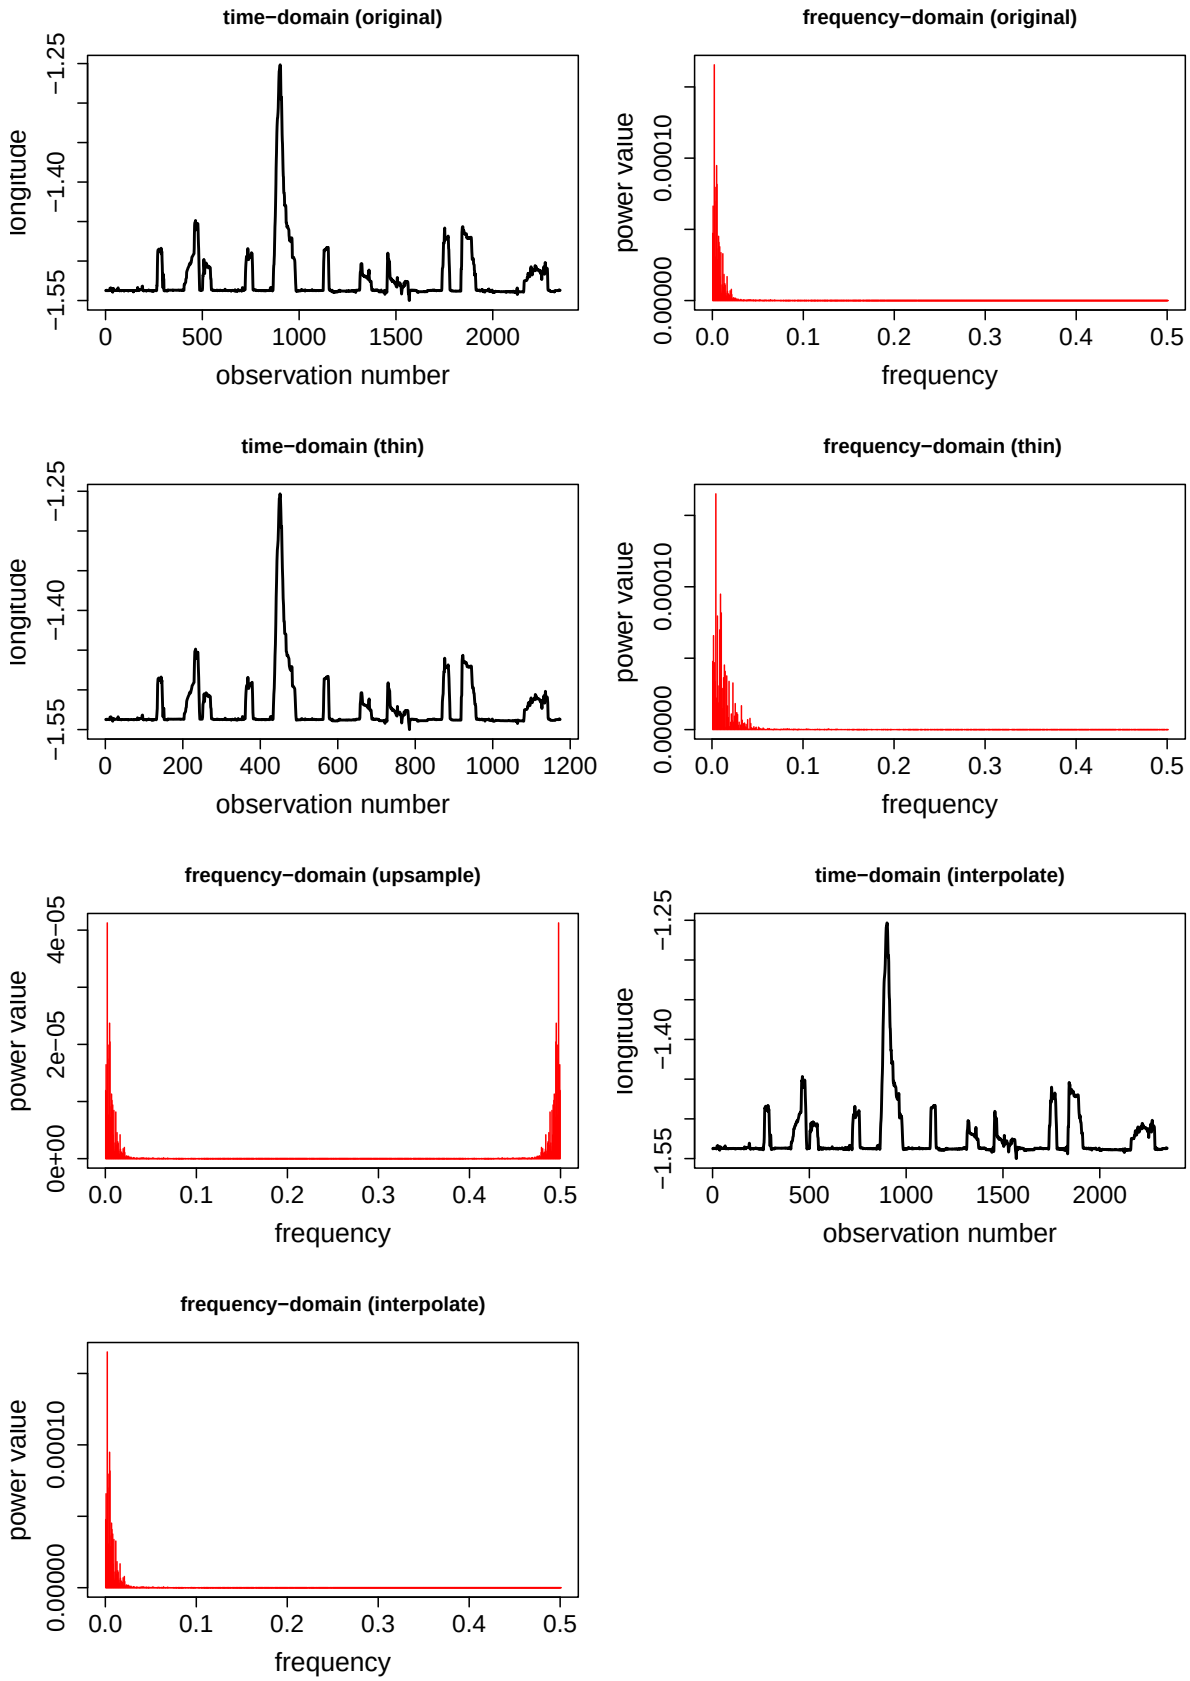

Figure 6.5: Processing of longitudinal movement trajectories in time and frequency domains: Original signal (Row 1), thinned by factor 2 (Row 2), upsampled (Row 3, left), and interpolated in time (Row 3, right) and frequency domains (Row 4).

Comparing closeness of filter-thin-interpolate and thin-interpolate versions of the simulated tracks to the original simulated tracks.

In Scenario 1 (i.e., simulated tracks with no high frequency component), Figure 6.6 shows that at 0.25 cut-off frequency and thinning factor of 2, the filter-thin-interpolate version of the simulated data produces the same average distance as the thin-interpolate version when both versions are compared to the simulated data. In some cases, an average distance of 0 is seen, indicating that the filter-thin-interpolate and thin-interpolate versions are exactly the same as the simulated tracks. A noticeable increase in the average distance for both versions is observed as the thinning factor increases and the cut-off frequency decreases. However, the average distance of the filter-thin-interpolate version is slightly lower than that of the thin-interpolate version at each cut-off frequency and thinning factor. In contrast to Scenario 1, the average distance for both filter-thin-interpolate and thin-interpolate versions is not exactly the same but close at 0.25 cut-off frequency and thinning factor of 2 in Scenario 2 (Figure 6.7). While there is an increase in the average distance for both versions as the thinning factor increases and the cut-off frequency decreases, the average distance of the filter-thin-interpolate version is much lower than that of the thin-interpolate version at each cut-off frequency and thinning factor.

We also consider the filtered versions of the simulated tracks at different cut-off frequencies and then compare with the simulated tracks for Scenarios 1 and 2 in Table 6.4. Results show that in Scenario 1, the filtered versions of the simulated tracks at 0.25 cut-off frequency have an average distance of 0 when compared with the simulated tracks. This implies that the filtered versions are exactly the same as the simulated tracks at a cut-off frequency of 0.25. However, the same does not hold for Scenario 2. An overall increase in the average distance is observed as the cut-off frequency decreases for each individual track in both scenarios, although the average distance for Scenario 1 is much lower than that of Scenario 2 at each cut-off frequency. This implies that the filtered versions are much closer to the simulated tracks in Scenario 1 than in Scenario 2.

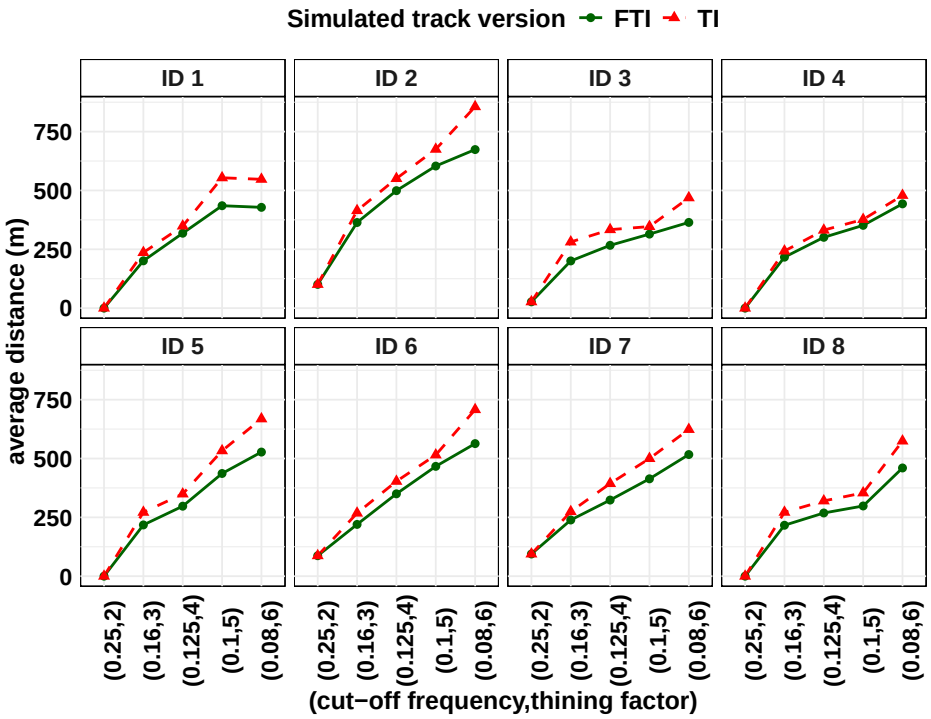

Figure 6.6: Comparison of average distances from the simulated track without high-frequency components (Scenario 1) to the filter–thin–interpolate (FTI; green solid line) and thin–interpolate (TI; red dashed line) versions.

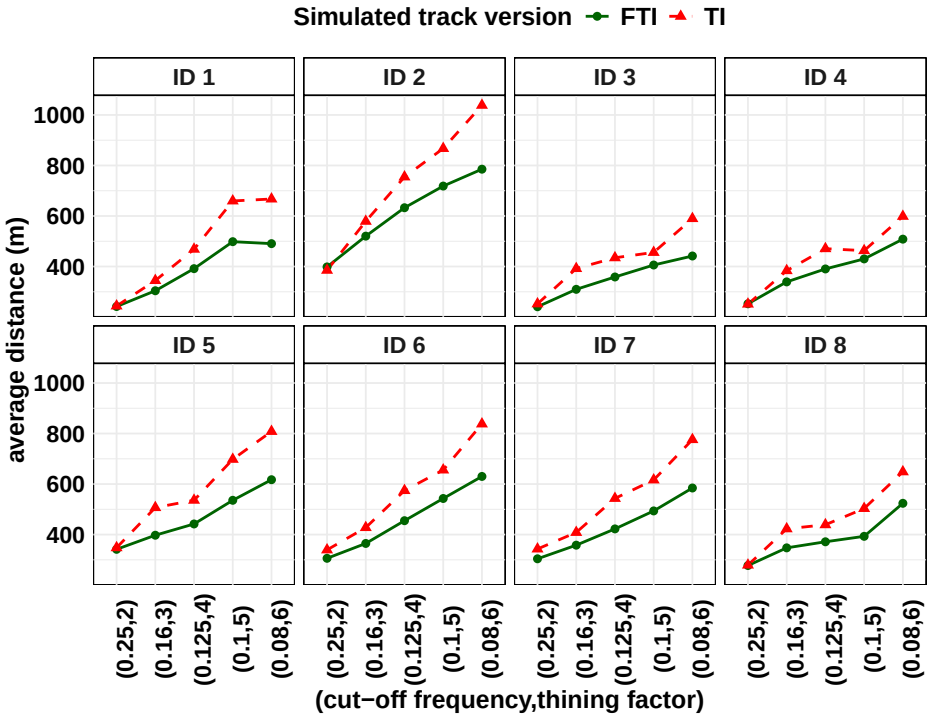

Figure 6.7: Comparison of average distances from the simulated track with some high-frequency components (Scenario 2) to the filter–thin–interpolate (FTI; green solid line) and thin–interpolate (TI; red dashed line) versions.

Table 6.4: Average distance (in metres) between locations of simulated data and (i) filtered (F), (ii) filter-thin-interpolate (FTI), (iii) thin-interpolate (TI) versions. cf = cut-off frequency for filtering and tf = thinning factor

		Scenario 1			Scenario 2		
Track ID	(cf, tf)	F	FTI	TI	F	FTI	TI
ID 1	(0.25,2)	0.00	0.00	0.00	241.62	241.62	243.76
	(0.16,3)	202.66	201.01	236.20	304.94	304.04	344.50
	(0.12,4)	286.78	318.32	348.82	372.39	391.53	468.41
	(0.10,5)	421.42	435.15	553.81	488.44	498.35	659.54
	(0.08,6)	481.46	428.23	547.40	529.82	490.52	667.24
ID 2	(0.25,2)	0.00	100.83	100.83	409.98	398.39	385.10
	(0.16,3)	356.59	363.77	414.40	522.32	520.26	578.97
	(0.12,4)	484.00	499.15	550.59	617.48	632.70	754.61
	(0.10,5)	628.16	603.85	674.98	734.24	718.48	867.23
	(0.08,6)	684.89	673.74	855.98	785.83	785.32	1038.19
ID 3	(0.25,2)	0.00	26.80	26.80	236.90	240.71	252.09
	(0.16,3)	200.74	200.59	281.17	310.35	309.67	391.94
	(0.12,4)	257.39	266.7	333.49	354.34	358.57	434.60
	(0.10,5)	339.66	314.52	346.53	423.35	405.93	455.81
	(0.08,6)	376.08	364.17	468.83	452.57	441.63	589.74
ID 4	(0.25,2)	0.00	0.00	0.00	252.52	252.52	251.19
	(0.16,3)	218.95	216.57	242.63	340.25	339.18	383.62
	(0.12,4)	266.29	300.81	331.27	365.88	390.38	470.59
	(0.10,5)	310.20	351.725	376.85	393.00	430.12	462.85
	(0.08,6)	506.10	443.13	479.01	565.42	508.33	598.62
ID 5	(0.25,2)	0.00	0.00	0.00	341.63	341.63	347.94
	(0.16,3)	207.44	217.79	271.50	398.71	397.40	506.81
	(0.12,4)	312.27	297.42	349.40	445.14	442.00	535.89
	(0.10,5)	427.41	436.40	532.94	531.31	535.70	697.86
	(0.08,6)	501.56	527.40	668.50	594.71	617.18	808.29
ID 6	(0.25,2)	0.00	87.01	87.01	293.75	306.03	339.36
	(0.16,3)	220.01	220.01	268.04	364.86	364.86	427.72
	(0.12,4)	350.00	350.00	403.09	454.95	454.95	574.25
	(0.10,5)	479.28	466.71	515.12	556.60	542.86	655.67
	(0.08,6)	563.35	563.35	707.78	630.14	630.14	837.90
ID 7	(0.25,2)	0.00	94.35	94.35	286.75	304.14	342.95
	(0.16,3)	240.81	238.83	274.29	358.99	358.02	408.51
	(0.12,4)	323.47	323.47	393.63	422.48	422.48	543.13
	(0.10,5)	399.10	413.72	499.95	483.96	493.65	616.23
	(0.08,6)	505.73	516.88	623.49	567.83	584.44	775.94
ID 8	(0.25,2)	0.00	0.00	0.00	277.43	277.43	279.21
	(0.16,3)	216.42	216.42	271.58	347.41	347.41	422.78
	(0.12,4)	290.50	268.60	319.90	388.58	371.47	438.92
	(0.10,5)	306.06	298.36	353.88	401.46	393.06	503.26
	(0.08,6)	459.85	459.85	574.36	523.81	523.81	648.41

Comparing accuracy of behavioural inference from filter-thin-interpolate and thin-interpolate versions under Scenarios 1 and 2

We fit a 3-state complete pooling HMM described in Section 4.3.3 to both the filter-thin-interpolate and thin-interpolate versions of the simulated track IDs under Scenarios 1 and 2. The model's overall accuracy (see Section 4.3.4) in inferring behavioural state for Scenarios 1 and 2 across both versions of the simulated data tracks is further compared. The model's results in Table 6.5 reveal that the overall accuracy of behaviours inferred from the filter-thin-interpolate version at all the selected cut-off frequencies and thinning factors is the same in both Scenarios 1 and 2. This suggests that applying the low-pass filter before thinning helps retain behaviourally meaningful features even after thinning and interpolating. However, this is not the case for the thin-interpolate version, as we observe a slight decrease in accuracy for Scenario 2. Most of the time, the accuracy of behaviours inferred from the filter-thin-interpolate version is higher than that of the thin-interpolate version under both scenarios, except for the case with the cut-off frequency of 0.25 and thinning factor of 2 in Scenario 2 and cut-off frequency of 0.12 and thinning factor of 4 in Scenario 1. This suggests that by applying the low-pass filter before thinning, NSST reduces the aliasing effect on behavioural inference.

Table 6.5: Classification accuracy of behaviours inferred from 3-state complete pooling HMM fitted to all the simulated tracks (ID 1 - ID 8). FTI = filter-thin-interpolate, TI = thin-interpolate versions.

	Overall accuracy (%)			
	Scenario 1		Scenario 2	
(cutoff frequency, thinning factor)	FTI	TI	FTI	TI
(0.25,2)	88.52	88.52	88.52	89.92
(0.16,3)	89.92	88.03	89.92	82.79
(0.12,4)	86.97	88.28	86.97	86.80
(0.10,5)	84.10	82.62	84.10	82.70
(0.08,6)	82.46	74.75	82.46	72.87

Selecting the choice of cut-off frequency values, B', for the low-pass filter

Recall from Section 6.3.2 that the cut-off frequency value, B', makes it possible to subsample working at a lower sample frequency than the original signal (i.e., $f_s' = 2B' < f_s$). We present the sensitivity analysis to assess the impact of the choice of

B' on the performance of 3-state HMMs in inferring behavioural states using one of the simulated tracks. A very low cut-off frequency, B', may lead to loss of meaningful behavioural detail (e.g., fast turns). Similarly, selecting a cut-off frequency that is too high may result in retaining unnecessary or unwanted information. Therefore, we select a more informed choice of cut-off frequency by using the high quantiles (i.e, 99%, 95%, 90%, 85%, and 80%) of the signal's power value. To achieve this, we find the frequency below which 99%, 95%, 90%, 85%, and 80% of the signal's total power lies. We then set the frequency as the cut-off frequency. Since the cut-off frequency determines how much information is lost, the high quantiles we selected also correspond to losing 1%, 5%, 10%, 15%, and 20% of information, respectively. The cut-off frequencies are then used to low-pass filter the trajectories in the frequency domain. The filtered trajectories are then converted back to the time domain (see, for example, Figures 6.8 and 6.9). Accuracy of behaviours inferred from the 3-state complete pooling HMMs fitted to the filtered simulated tracks is compared with behaviours inferred from the unfiltered (i.e., no low-pass filter) simulated tracks in Table 6.6. The F1-score is used as a metric for assessing the accuracy of behaviours (see Section 5.4).

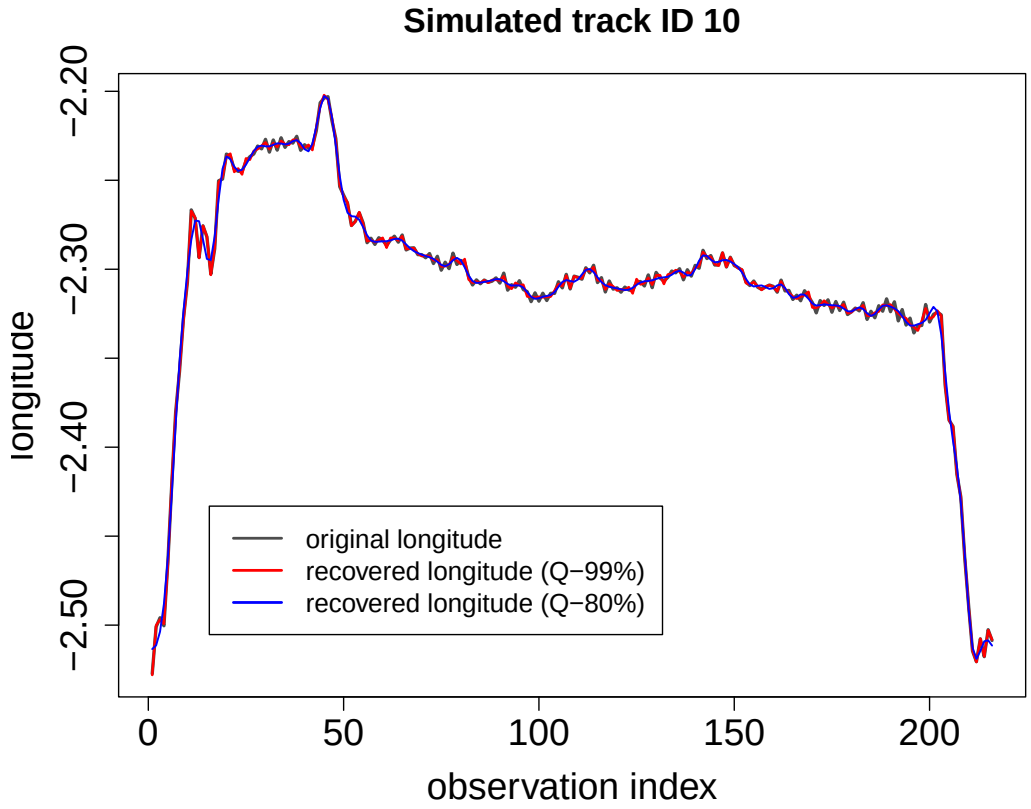

Figure 6.8: Original longitudinal trajectory of simulated track ID 10 compared with its recovered versions after low pass filter with cut-off frequencies chosen at 99% (red line) and 80% (blue line) quantiles.

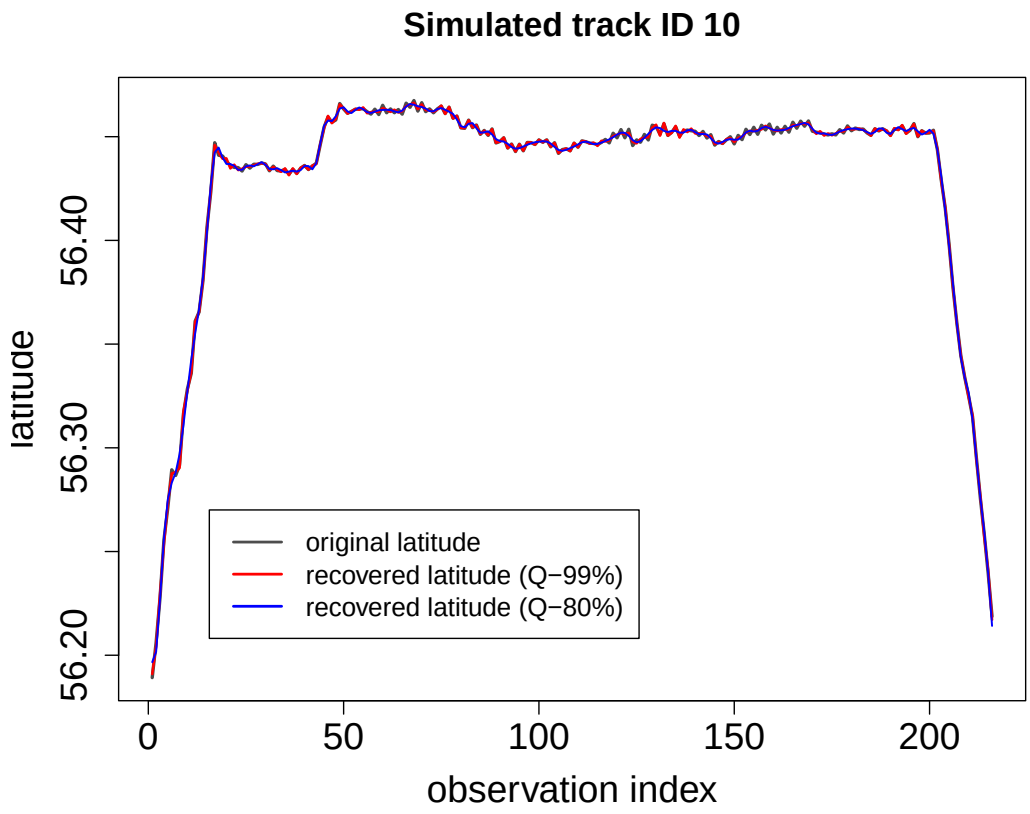

Figure 6.9: Original latitudinal trajectory of simulated track ID 10 compared with its recovered versions after low pass filter with cut-off frequencies chosen at 99% (red line) and 80% (blue line) quantiles.

The use of a higher quantile to select the cut-off frequency for filtering produces a closer representation of the original longitudinal and latitudinal trajectories after filtered trajectories are recovered back in the time-domain as seen in Figures 6.8 and 6.9. Furthermore, results in Table 6.6 show that higher quantiles of the power signal are preferable for behavioural inference.

Table 6.6: Accuracy of behavioural state inferred from 3-state HMM fitted to simulated data and recovered data after low-pass filter (LPF) with cutoff frequency value selected at different quantiles for simulated TrackID 10

Accuracy(%) - F1-score				
Filter	**Overall**	**travel**	**forage**	**search**
No filter	95.83	96.88	96.86	88.00
LPF: Q-99%	91.67	96.88	98.74	40.00
LPF: Q-95%	92.13	96.88	98.74	44.00
LPF: Q-90%	90.28	90.62	98.74	36.00
LPF: Q-85%	91.20	96.88	98.74	36.00
LPF: Q-80%	89.35	84.38	98.11	40.00

6.5 Case Studies

Following the simulation study in Section 6.4, we present the application of the NSST to some examples of real tracking data described in Chapter 2. We begin with the application of NSST on the behavioural data in Section 6.5.1 and conclude with the application of NSST on the GPS-movement tracks in Sections 6.5.2

6.5.1 NSST from the perspective of the hidden process: a case study of Roseate tern

We consider the roseate tern (*Sterna dougallii*) visual track (Δt = 1-sec, n = 9236) as an example because the behavioural data of the roseate tern, which are available, are necessary to assess the use of NSST from the perspective of the underlying (hidden) process. Table 6.7 reports the behavioural bouts that are present in the behavioural data obtained by following the steps highlighted in Section 6.3.1. There are two behavioural states of foraging and not-foraging noted in the roseate tern data. The latter behavioural state records a total of 15 bouts (i.e., the roseate tern was observed to be consistently performing a not-foraging behaviour without switching to a foraging behaviour, at 15 different times). The former, on the other hand, has a total number of 14 bouts (i.e., the roseate tern was observed to be consistently performing a foraging behaviour without switching to a not-foraging behaviour, at 14 different times). The information about the length and time length of the bout with minimum activity for each behaviour is also reported in Table 6.7. This is because we are interested in assessing whether the sampling interval, Δt = 1-sec, at which the movement tracks were originally collected, is sufficient for behavioural inference. From Table 6.7, the time length of the bout with the minimum number of foraging and not-foraging behavioural activity is 60-secs and 40-secs, respectively. Recall, a sampling interval Δt_{new} informed by the NSST is chosen such that

$$\Delta t_{new} = min\left(\frac{1}{2}\left(t_1, t_2, \ldots, t_n\right)\right)$$

where n is the total number of underlying behavioural states and $t_1, \ldots, t_n$ is the time length of behavioural bout with the minimum activity.

From Table 6.7 we have that Δt_{new} = 20-secs. This suggests that the original sampling interval of Δt = 1-sec is sufficient. However, the original sampling interval does not necessarily need to be as fine as 1-sec and sampling intervals $\leq$ 20-secs will be sufficient as well. Additionally, sampling intervals $\leq$ 30-secs will be adequate if the foraging behavioural state is of interest.

Table 6.7: Observed behavioural bouts from Roseate tern data with Δt = 1-sec. n = number of observed bout for each behaviour.

	bout with minimum behavioral activity		
Behavioural bout (n)	Bout length	Time of bout length, t (secs)	t/2 (secs)
Foraging (14)	60	60	30
Not-foraging (15)	40	40	**20**

The Roseate tern visual tracking data originally obtained at Δt = 1-sec is resampled at 20-secs and 30-secs to compare the accuracy of inferred behaviours among the sampling intervals. A 2-state HMM (state 1= foraging, state 2 = not-foraging) is fitted to the tracks (i.e., tracks Δt = 1-sec, Δt = 20-secs, and Δt = 30-secs) to assess the accuracy of behavioural states inferred using the F1-score and log-loss metrics (see Section 5.4). Figures 6.10 and 6.11 shows the tracks at each sampling interval coloured by the decoded behaviours.

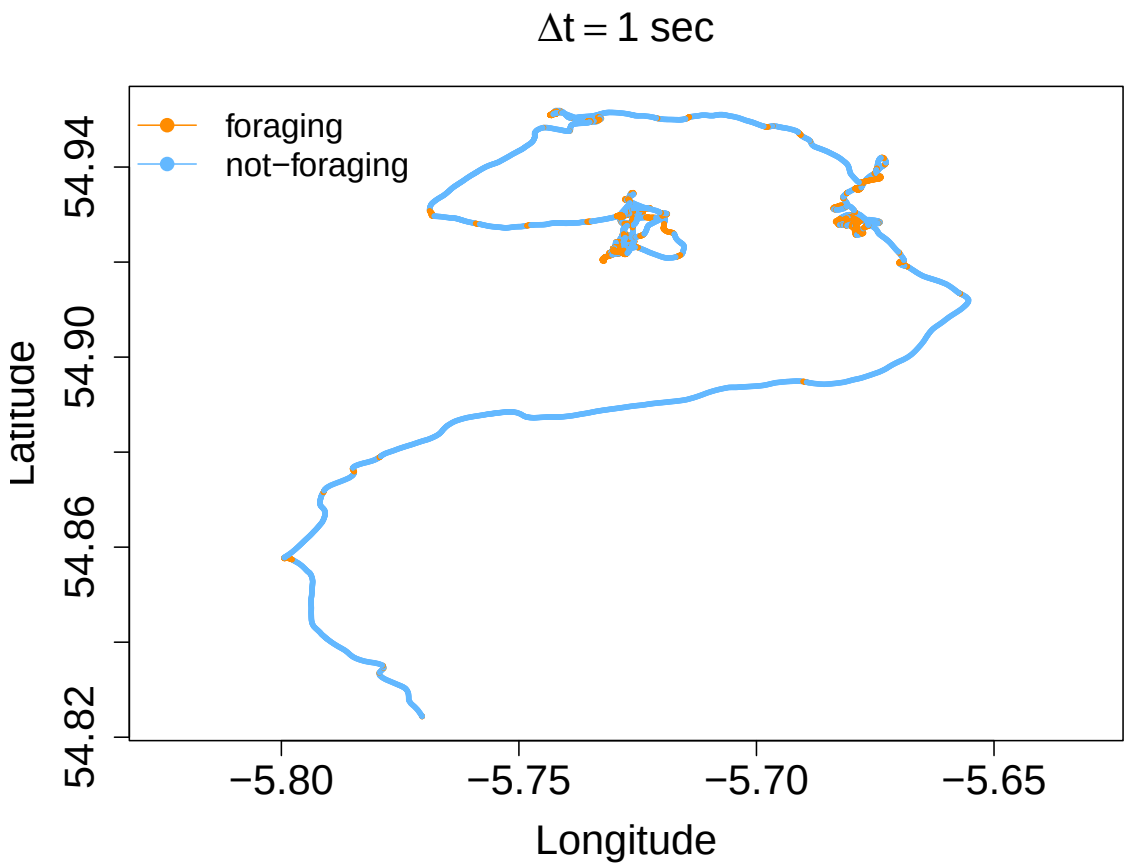

Figure 6.10: Roseate tern track (ID1), coloured by behavioural state decoded from a two-state HMM fitted at sampling intervals, Δt = 1-sec.

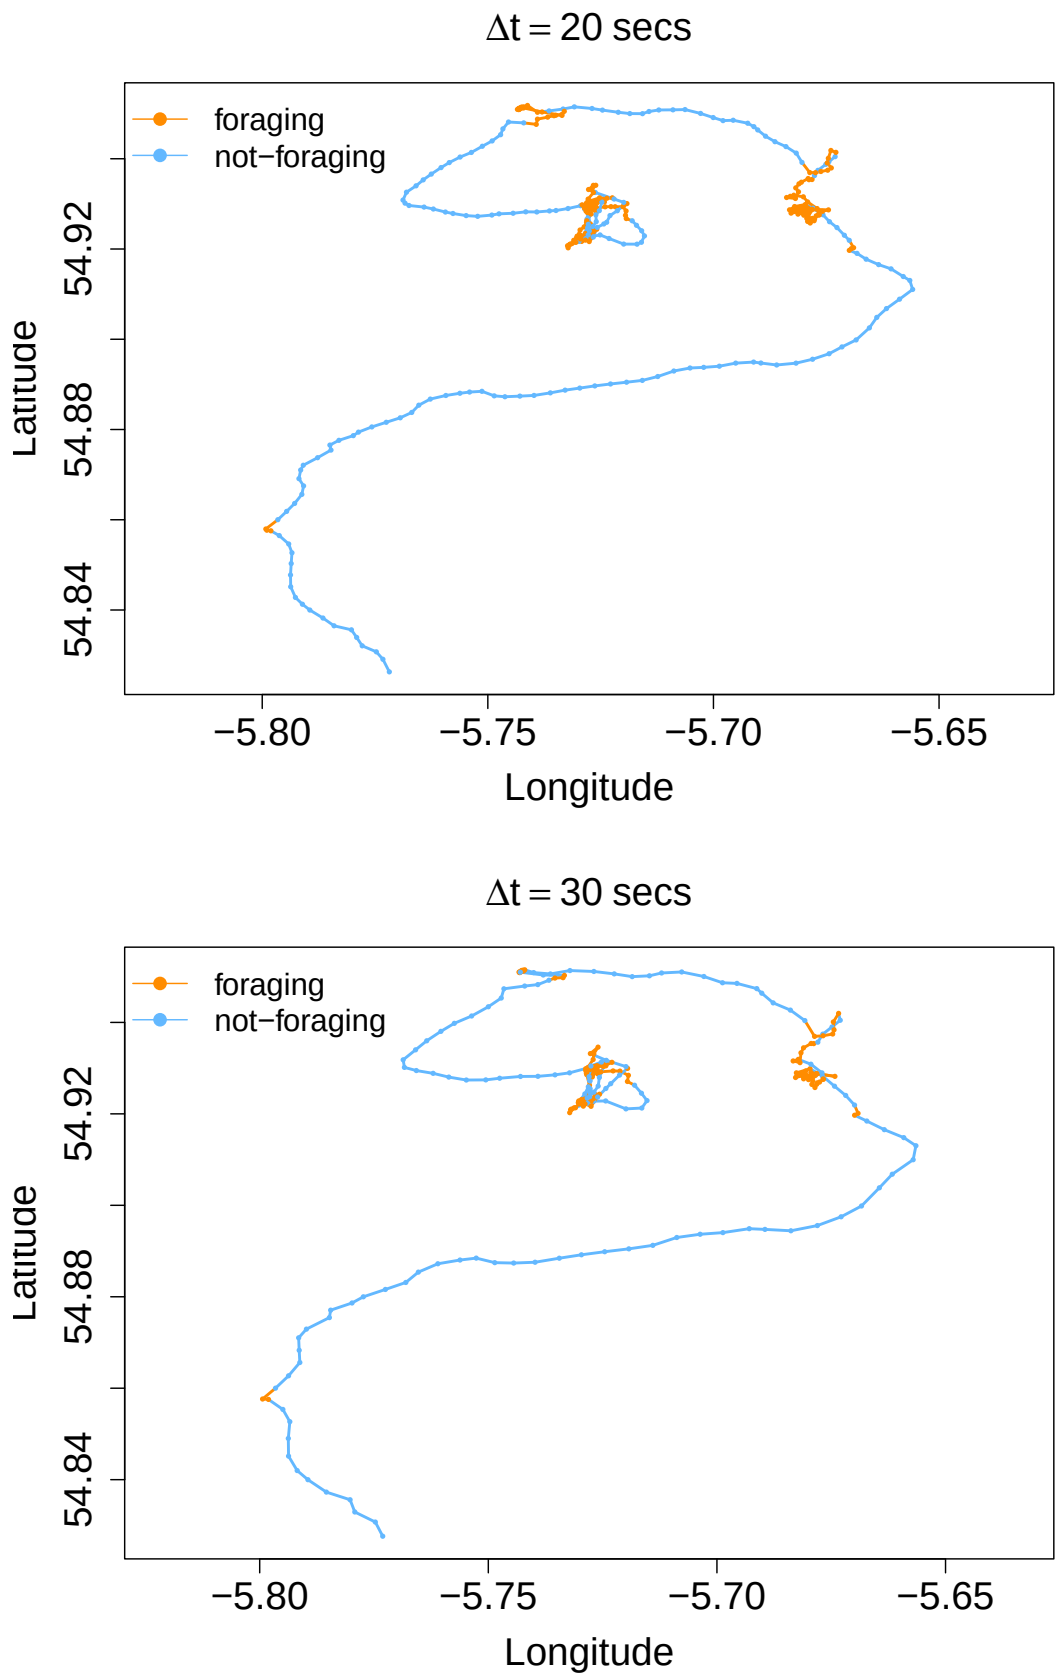

Figure 6.11: Roseate tern track (ID1), coloured by behavioural state decoded from a two-state HMM fitted at sampling intervals, Δt = 20-secs (upper panel) and Δt = 30-secs (lower panel).

Results in Table 6.8 show that the accuracy of behaviours inferred from tracks sampled at $\Delta t = 20$-secs outperforms those of $\Delta t = 1$-sec and $\Delta t = 30$-secs with an F1-score of 75% and log-loss value of 11.65. We also observe that the performance of the resampled track at $\Delta t = 30$-secs also outperforms that of the original track sampled at $\Delta t = 1$-sec. Similarly, Figures 6.10 and 6.11 also confirm the results reported in the Table 6.8, as we see that tracks resampled at $\Delta t = 20$-secs decode the behavioural state better, particularly the foraging behaviours.

Table 6.8: Validation result of 2-state HMM fitted to Roseate tern visual tracking data originally collected at $\Delta t = 1$-sec and subsampled at $\Delta t = 20$-secs and $\Delta t = 30$-secs.

	Sampling intervals		
Accuracy metrics	**1-sec**	**20-secs**	**30-secs**
F1-score (%)	70	75	73
Log-loss	12.81	11.65	12.07

Since NSST suggest that sampling intervals, $\Delta t \leq 20$-secs and $\Delta t \leq 30$-secs will be sufficient for the not-foraging and foraging behavioural state respectively, the tracks with $\Delta t = 1$-sec is also resampled at $\Delta t = 2, 3, 4, \ldots, 30$-secs to assess accuracy of behaviours inferred from the 2-state HMM fitted to the resampled versions. We report only the log-loss metric in Figure 6.12 because it accounts for how close the inferred behaviours are to the actual behaviours. From Figure 6.12, a sharp decrease in the log-loss is seen from $\Delta t = 3$-secs, while the log-loss of sampling intervals ranging from 3-secs to 25-secs is much lower than that of $\Delta t = 1$ to 2-secs compared to sampling intervals from 26-30 secs. This suggests that sampling intervals informed by the NSST perform better than the original sampling interval of $1-$sec in terms of behavioural inference.

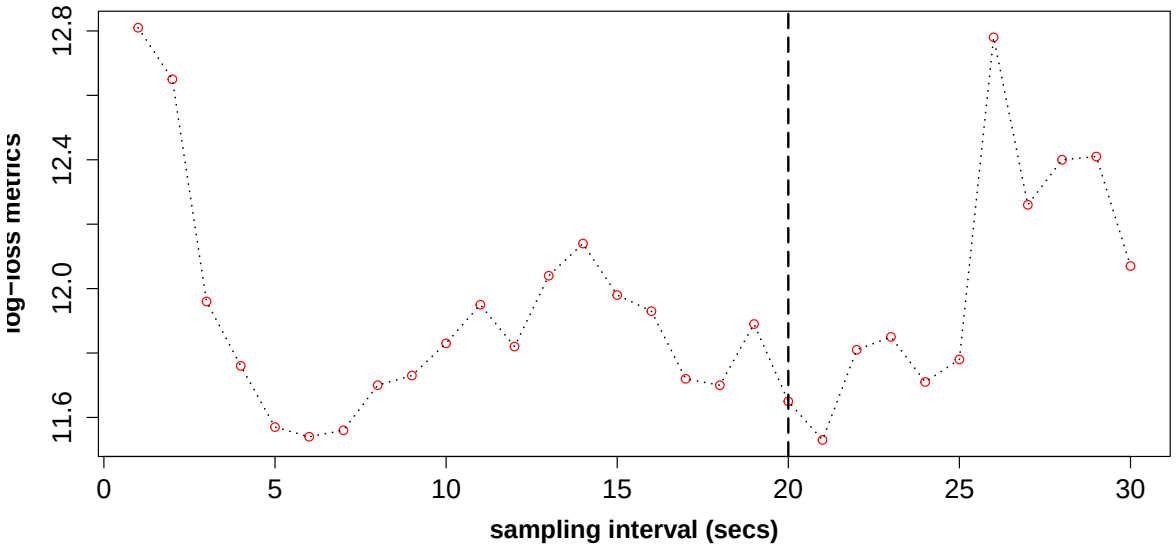

Figure 6.12: Assessing the validity of behaviours inferred from a 2-state HMM fitted to the roseate track at different sampling intervals. The vertical long-dashed line represents the sampling interval (Δt = 20-secs) informed by NSST.

6.5.2 NSST from the perspective of the observed process: a case study of Kittiwake

We consider the kittiwake data described in Chapter 2 as an example to assess how NSST can be used to inform the choice of subsampling frequency and improve the standard thinning approach of reducing the number of data points for animal movement data. The kittiwake tracks are obtained at a sampling interval of 100-secs. The tracks with the least (n = 351) and most (n = 2348) number of data points (Figure 6.13) are used as examples to initially describe the data from the perspective of the time-domain and frequency-domain representations before presenting findings of the case study results for all the tracks considered. The left panel of Figure 6.14 shows how the longitude and latitude positions of the kittiwakes change with respect to time. Frequent fluctuations in longitude and latitude positions are more noticeable for the track with a longer time length (ID0066) compared to the track with a shorter time length (ID0062). This suggests that the kittiwake with longer time length has better exploration of the surrounding locations in space, as seen in Figure 6.14. Consequently, we observe that the track ID0066 has most of the power (i.e., strength of each longitude and latitude location), concentrated in the lower frequencies than those of track ID0062.

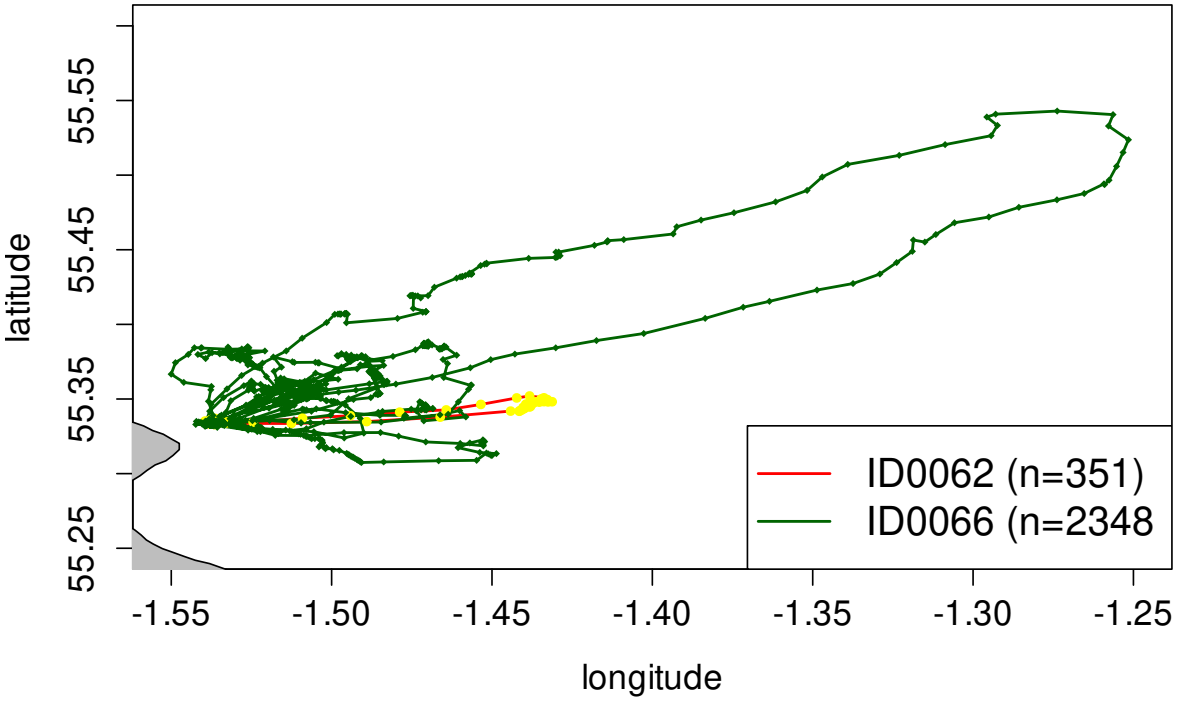

Figure 6.13: An example of the kittiwake movement tracks for ID0062 and ID0066.

We compared filtered versions of the kittiwake tracks at varying cut-off frequencies with the original tracks under Scenarios 1 and 2 (Table 6.9). In Scenario 1, filtering at a 0.25 cut-off frequency resulted in an average distance of 0, indicating that the filtered tracks were identical to the original kittiwake tracks. This pattern was not observed in Scenario 2. In both scenarios, the average distance increased as the cut-off frequency decreased; however, distances were consistently smaller in Scenario 1 than in Scenario 2, indicating closer agreement with the original kittiwake tracks. Findings from the cases study results are also similar to the findings from the simulation study. This suggests that analysis of movement data in the frequency-domain can be useful to gain insights into what may not be obtainable in the time-domain. In particular, results from this case study example show that NSST (analysis in the frequency-domain) when applied prior to the standard thinning approach, can be helpful for minimizing information loss that may be useful for behavioural inference.

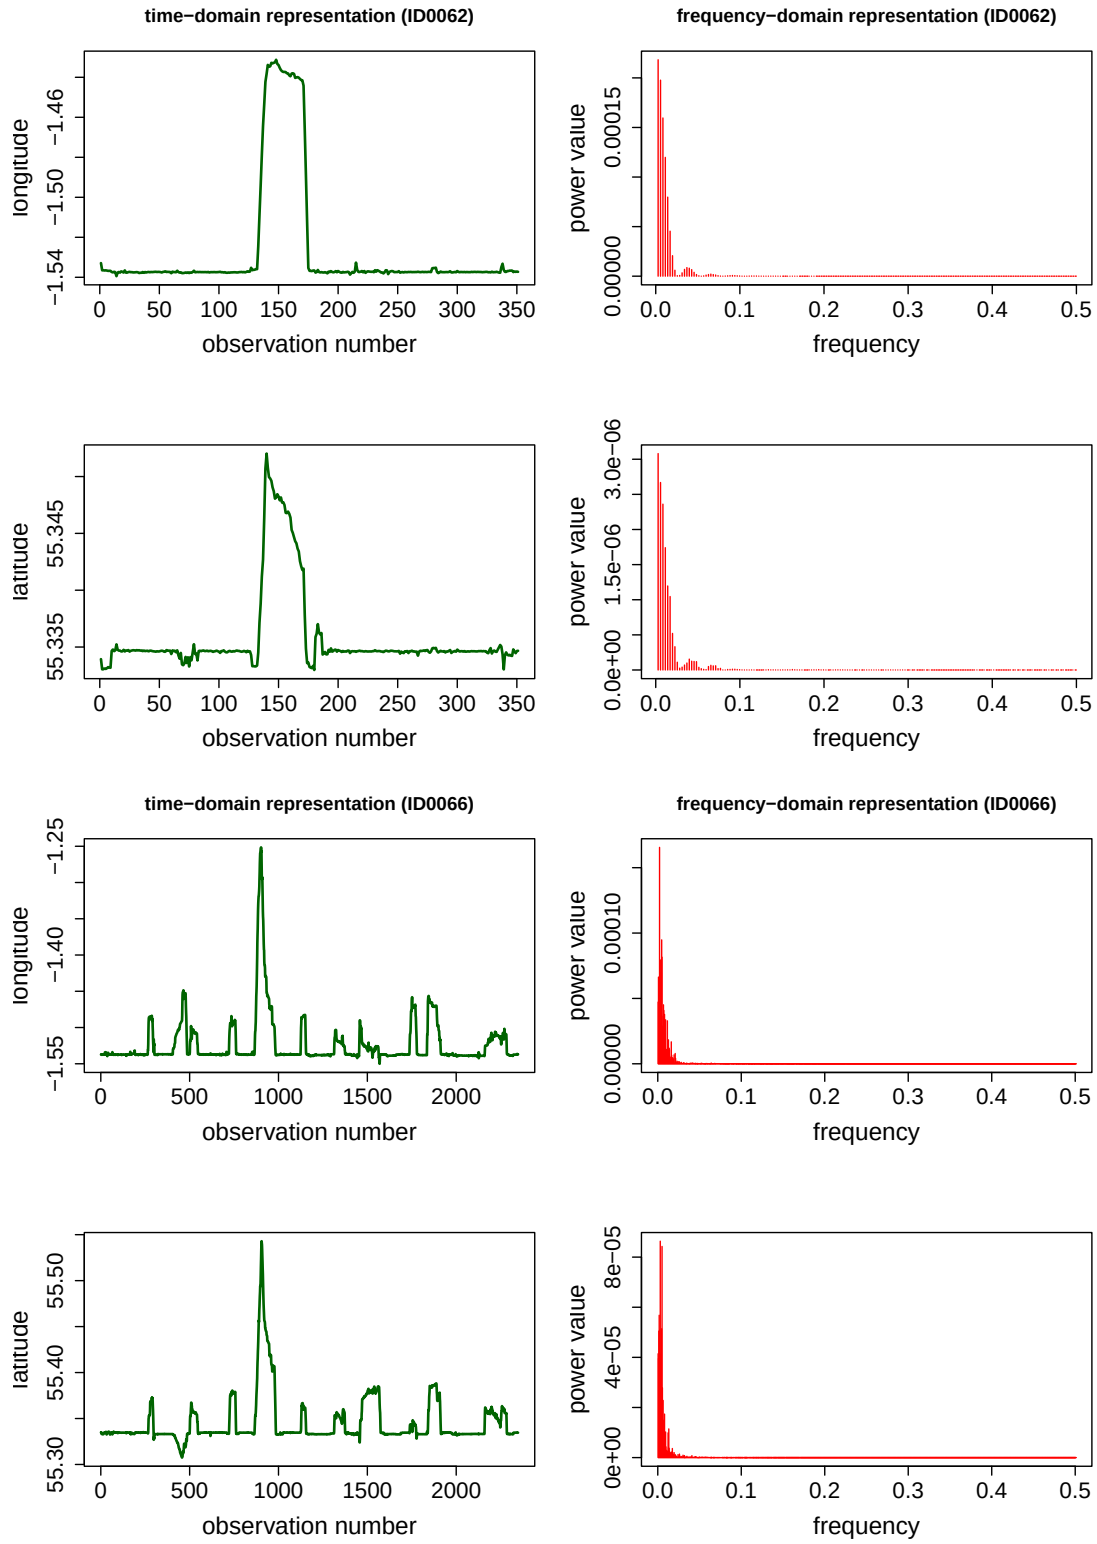

Figure 6.14: Longitudinal and latitudinal movement trajectories of kittiwakes (ID0062: Rows 1–2; ID0066: Rows 3–4) shown in the time domain (left) and frequency domain (right). Time-domain panels illustrate changes over time, while frequency-domain panels show the distribution of power across normalised frequencies (0–0.5).

Table 6.9: Average distance (in metres) between kittiwake data and (i) filtered (F), (ii) filter-thin-interpolate (FTI), (iii) thin-interpolate (TI) versions. cf = cut-off frequency for filtering and tf = thinning factor

		Scenario 1			Scenario 2		
Track ID	(cf, tf)	F	FTI	TI	F	FTI	TI
ID 0040	(0.25,2)	0.00	0.00	0.00	59.08	59.08	57.37
	(0.16,3)	81.97	81.97	101.51	98.35	98.35	120.17
	(0.12,4)	113.22	116.90	142.03	125.92	129.31	157.49
	(0.10,5)	156.36	156.15	181.60	166.53	166.16	196.15
	(0.08,6)	201.28	201.28	250.06	210.86	210.86	259.40
ID 0044	(0.25,2)	0.00	5.93	5.93	286.75	65.07	68.07
	(0.16,3)	240.81	81.26	96.88	358.99	101.39	122.81
	(0.12,4)	323.47	121.99	146.44	422.48	137.30	165.83
	(0.10,5)	399.10	150.72	184.29	483.96	162.92	207.23
	(0.08,6)	505.73	196.05	238.42	567.83	205.09	262.12
ID 0048	(0.25,2)	0.00	0.00	0.00	277.43	16.36	16.35
	(0.16,3)	216.42	19.75	23.22	347.41	25.20	29.55
	(0.12,4)	290.50	27.31	30.18	388.58	31.13	37.97
	(0.10,5)	306.06	30.72	36.67	401.46	34.08	43.03
	(0.08,6)	459.85	54.80	72.91	523.81	57.54	73.87
ID 0052	(0.25,2)	0.00	0.00	0.00	295.83	45.34	45.63
	(0.16,3)	244.23	56.46	65.74	382.53	70.25	83.58
	(0.12,4)	355.85	81.81	96.15	459.71	91.94	110.06
	(0.10,5)	376.53	103.58	125.44	484.91	111.24	135.93
	(0.08,6)	444.77	125.13	162.23	540.28	131.94	176.38
ID 0059	(0.25,2)	0.00	5.81	5.81	254.12	60.45	62.31
	(0.16,3)	168.04	71.80	85.65	307.69	89.85	109.94
	(0.12,4)	218.56	104.38	129.22	332.93	117.45	147.64
	(0.10,5)	335.05	127.04	153.08	423.37	137.59	170.59
	(0.08,6)	477.16	163.48	212.14	540.01	171.82	237.79
ID 0062	(0.25,2)	0.00	0.00	0.00	254.12	22.16	24.47
	(0.16,3)	168.04	27.19	34.48	307.69	34.24	42.70
	(0.12,4)	218.56	42.65	44.82	332.93	46.90	49.48
	(0.10,5)	335.05	57.21	70.27	423.37	61.03	78.76
	(0.08,6)	477.16	74.89	105.96	540.01	76.36	119.76
ID 0064	(0.25,2)	0.00	8.13	8.13	55.87	56.36	57.48
	(0.16,3)	64.05	64.34	77.41	81.29	81.59	97.07
	(0.12,4)	97.32	97.32	117.68	110.77	110.77	133.47
	(0.10,5)	114.70	117.11	141.11	125.08	127.33	161.86
	(0.08,6)	131.26	131.74	162.92	140.17	140.68	180.10
ID 0066	(0.25,2)	0.00	2.86	2.86	254.12	39.33	42.12
	(0.16,3)	168.04	51.54	62.74	307.69	61.42	73.85
	(0.12,4)	218.56	72.80	87.50	332.93	80.09	100.35
	(0.10,5)	335.05	93.37	121.82	423.37	99.12	132.09
	(0.08,6)	477.16	115.54	145.75	540.01	120.52	157.18

6.6 Discussion

This study aimed to apply the Nyquist-Shannon sampling theorem to animal movement data for behavioural inference. We discuss the implications of our results with respect to how NSST can be beneficial for behavioural inference when subsampling existing animal movement data, and also informing sampling frequency for data collection. The choice of sampling frequency for collecting animal movement data depends on the aim of the study (Yu et al. 2023). While a fine-scale temporal resolution may be helpful for capturing more behaviours, it comes at a cost of battery life of the telemetry device. Sampling at a coarser scale prolongs battery life but may miss out on behaviours obtainable at a finer scale (Marcus et al. 2012, Schoombie et al. 2024). Tracking groups of individual animals with telemetry devices set to different sampling frequencies can be expensive, especially when reasonable sample sizes are needed for each category of the specified sampling frequency (He et al. 2023). As such, it is important to decide *a priori* on what sampling frequency to use. Although NSST can be used to inform the choice of sampling frequency in other fields, it cannot be used directly to set sampling frequency for collecting animal data. However, if the frequency range of behaviour of interest is known or can be estimated, then NSST can be used to inform the choice of sampling frequency.

Using simulated and observed behavioral data, we showed how to use NSST to get valuable insights into how the choice of sampling frequency impacts inference of behaviours from GPS movement trajectories. When the sampling frequency selected is at a coarse scale, it is possible that the chosen sampling frequency may be insufficient to capture short bursts of behavioural bout even when models may perform generally well in classifying behavioural states, as suggested by our results. In particular, our result showed that when inferring behaviours, sampling intervals of 5−mins failed to capture bouts that lasted for 5−mins and 10−mins respectively. As indicated in this work, the NSST approach can tell us that the sampling interval should be lower than the actual interval used. However, it cannot tell us exactly how much lower. Therefore, we suggest that a fine-scale sampling frequency may be used (where possible and necessary) to collect movement data and improve the performance of models in inferring behaviours when there is a likelihood of short-burst of behavioural bout. Most importantly, the behavioral understanding of the animals being tracked should be factored in if the aim of the study is to infer behaviors from telemetry data, since ground-truthing data to validate inferred behaviours is difficult to obtain. Furthermore, consistent with previous studies (e.g., Tatler et al. (2018), Schoombie et al. (2024)), our case study results suggested that a high sampling frequency (1−sec in our case) may not always be necessary for the benefit of classifying behavioral

states. Our example with the casestudy data showed that the accuracy of behavioural states inferred from tracks resampled at 20−secs and 30−secs (informed by NSST) outperformed that of the original sampling interval of 1−sec. Additionally, inferred behaviours from the 20−secs and 30−secs were much closer to the observed behaviours than those of the 1−sec sampling interval, as indicated by the log-loss metrics.

Previous studies have shown how NSST can be used to inform the choice of sampling frequency for collecting animal data using an accelerometer and how these impact behavioural inference (Tatler et al. 2018, Hounslow et al. 2019, Yu et al. 2023). However, very few studies have applied NSST to inform the choice of subsampling frequency for existing datasets (Nathan et al. 2022). To the best of our knowledge, this work is the first to use NSST to carry out thinning in a better way by applying low-pass filtering prior to thinning within the context of animal movement ecology. In practice, ecologists often record movement data of animals at high temporal resolutions and gradually down-sample by thinning. Thinning is equivalent to subsampling without filtering. This violates the NSST requirement if the remaining sampling rate is below twice the highest frequency in the data and thus leads to loss of behavioural information.

We showed that in scenarios where there are no high-frequency components in the movement trajectories, the standard thinning approach yields similar results as the NSST approach. However, in practice, animal movement trajectories will often contain some high-frequency components when viewed in the frequency-domain. When subsampling such movement trajectories without using an informed choice of subsampling frequency, it is possible for these high frequencies to be misrepresented as low frequencies in the frequency domain. The misrepresentation, also termed as aliasing, may negatively impact behavioural inference from such subsampled trajectories. Therefore, removing high frequency before subsampling is needed to avoid aliasing. Our results showed that applying the filtering before standard thinning provided a closer representation of the original trajectory compared to thinning without prefiltering. Additionally, the behavioural inference from these two different datasets (i.e., trajectories with and without prefiltering) also indicated that the NSST approach helps retain information contained in the trajectory, as there was improved accuracy of behaviours inferred from the NSST approach compared to the standard thinning approach.

In practice, most of the power of the signal is at low frequencies, including animal movement data. This implies that the behaviourally relevant motion is at the lower-frequencies, while the higher frequencies contains the less relevant ones. The

high frequencies can be discarded at selected cut-off frequency. However, investigating how behavioural inference is sensitive to the choice of cut-off frequency for filtering is important. Our results suggest that a cut-off frequency value chosen based on high quantiles of the signal's power is preferable for improved behavioural inference from filtered movement data. Conclusively, the NSST approach is useful for improving subsampling results when filtering is done prior to thinning with the aim of inferring behaviours from subsampled movement trajectories.

Chapter 7

Conclusions and future work

Studying animal movement contributes to our understanding of animal behaviour and how this understanding can be used to improve conservation actions. While mitigating human activities that serve as threats to animal populations is one of the necessary actions needed to aid conservation efforts (Croxall et al. 2012, Lascelles et al. 2016, Wakefield et al. 2017), identifying and protecting areas which animals use consistently (e.g., for foraging) can help mitigate these threats by reducing activities in conservation-sensitive areas. Since direct observation of animal behaviour can be challenging, especially in the marine environment, behaviours are often inferred from animal tracking data using appropriate movement models (Jonsen et al. 2005, Patterson et al. 2009, Browning et al. 2018, McKellar et al. 2014). GPS devices have been used to track various animals, including seabird species (Burger & Shaffer 2008, Wakefield et al. 2009), and previous studies have inferred behaviours from animal GPS tracking data using HMMs (Morales et al. 2004, Joo et al. 2013, Dean et al. 2013, McClintock et al. 2017). However, despite the large number of tracking studies carried out to date, the validation of these HMM-inferred behaviours using first-hand observed behavioural data of animals has not been done to the best of our knowledge. As such, there are uncertainties associated with these inferred behaviours that may impact conservation efforts if they are in the wrong location. The impact may be more pronounced in the marine environment as seabirds, which are good indicators of marine health, are offered less protection at sea than at their breeding colonies, though at-sea foraging locations are under increasing pressure from human activities (Croxall et al. 2012).

In this thesis, I have validated HMMs commonly used in inferring behaviors by utilising data collected using a little used tracking methodology - visual tracking - specifically designed for terns and developed by Perrow et al. (2011). The method allows detailed information on foraging areas utilised by terns to be examined with two distinct data:

(i) the ground truth behavioural data and (ii) the boat GPS tracks used in visually tracking the birds. In this thesis, I make use of the visual tracking data because they provide access to ground-truth behavioural data, and so provide a unique basis for validating HMMs fitted to movement data to infer behaviours. While the visual tracking data provides ground-truth data, locations were recorded for the boat following birds rather than the birds themselves. Consequently, the boat tracks were initially assessed as a proxy for the GPS track of the birds before fitting HMMs to the boat track to infer behaviours, and finally validating inferred behaviours using ground-truth observational data. The initial step was important because it is common practice to fit models to GPS tracks of animals rather than a proxy (e.g., boat) track. As such, it was necessary to check whether the boat tracks were a suitable representation of the bird track.

In assessing the use of boat tracks as a proxy for GPS tracks, the bird movement path was reconstructed for the subset of tracks for which further information was available and compared to the boat tracks. Results showed that the boat tracks close replicated the reconstructed path of the bird. Additionally, behaviours inferred from the boat tracking data do not differ substantially from those of the birds' reconstructed path for the four species of terns. Our results suggest that the visual tracking data (GPS tracks of the boat) is a good proxy for true bird locations in relation to inferred behaviours and that the visual tracking method can be used as a potential alternative to GPS-tracking in situations where the loss of transmitters limits or prevents the use of telemetry. However, a potential limitation of the visual tracking method worth highlighting is that the boat tracks may not capture the turning angle of the birds well enough, especially when they make sharp, unexpected turnings. Our results confirmed this limitation, as the main difference observed in the data between true (predicted) bird locations and boat locations was noted in their respective turning angles, which implies that the turning angle may not be particularly informative. This may impact the ability of the fitted models to decode the foraging areas, as the GPS location of the boat is used as a proxy for the location of the birds, and it is also likely to limit our ability to estimate more than two behaviours using HMMs.

This thesis also focused on validating HMM-inferred behaviour, which is very rarely feasible in practice due to the difficulty of accessing observed behavioural data. Using the visual tracking method, the ground-truth behavioural data were used to assess the performance of HMMs in inferring behavioural states from boat tracks of the four tern species across the nine breeding colonies considered during chick-rearing and incubation. My results with respect to the overall performance of HMMs in inferring behaviours showed that HMMs were able to decode behaviours with a high degree of

accuracy when compared with the ground-truth behavioral data for most colonies and species. Additionally, for the foraging behavioral state, which is of primary interest for conservation purposes, HMMs were also able to capture a substantial amount of the total observed foraging bouts (i.e., a period of observed foraging activity within the boat tracks). Lastly, my results showed that the accuracy of the fitted HMMs was generally insensitive to model choice, even when AIC values ranged widely across different models. This is also supported by the findings from McClintock (2021), where complex HMMs do not significantly improve behavioural state classification when compared to the simpler HMMs, with the latter being preferred for the goal of behavioural inference. Results suggest that HMMs fitted to tracking data can accurately identify important conservation-relevant behaviours, demonstrated using visual tracking data. Furthermore, even with the limitation of the visual tracking (boat) data mentioned earlier, my results showed that HMMs performed well in decoding behavioural states. As such, our findings increase confidence in the use of HMMs for inferring behaviours from GPS tracks, even without validation data, which can be challenging to obtain. This has important implications for animal conservation, such as evaluating the impacts of anthropogenic activities (e.g., offshore windfarms) on seabirds, where understanding how and why animals use the marine environment is critical (Thaxter et al. 2012, Lascelles et al. 2012). Although HMMs have been widely applied to GPS tracking data of animals (Joo et al. 2013, Dean et al. 2013, McClintock et al. 2017) in animal movement ecology, my study indicates that HMMs can also be successfully applied to visual tracking data (where available) for the purpose of behavioural inference, as the visual tracking method is useful for determining the foraging movement of birds (e.g., terns) confirmed in a previous study by Perrow et al. (2011). The limitations of this work, for the validation data, is that the accuracy of fitted HMMs in inferring behaviours relies on the assumption that there are no errors in the observed behavioural data and secondly for the visual tracking method is that it is impossible to determine certain characteristics from the boat (Lascelles et al. 2012), making it impossible to model birds based on sex or age, which may (where available) help improve the accuracy of behavioural inference.

Animal behaviors are the underlying hidden information in movement data, which are usually inferred using appropriate statistical models to inform conservation actions, such as spatial planning, identifying special protected areas, and marine protected areas (Thaxter et al. 2012, Lascelles et al. 2012). As such, it is important to minimize loss of useful behavioural information embedded within animal movement data. This thesis presented an application of NSST to animal movement with the aim of showing the consequences of violating NSST (i.e., information loss) on behavioural inference

when collecting animal data. Most importantly, this thesis presented an application of NSST to animal movement data to enable data to be sub-sampled with minimal information loss resulting from a reduced number of data points, unlike the standard thinning approach.

In conclusion, my work has addressed the question of how well animal behaviour can be inferred from movement data for effective conservation purposes. In particular, validating visual tracking data proxy for GPS tracks, validating HMMs for behavioural inference, and applying NSST to inform the choice of (sub)sampling frequency of movement data for behavioural inference. Based on this thesis, I propose the following potential future research:

1. *Exploring and validating HMM extensions for animal behavioural inference.* This thesis focuses on HMMs for behavioural inference because they are popularly used due to an appealing combination of model flexibility, clear interpretability, and ease of computational management (McClintock & Michelot 2018). Various extensions that relax assumptions of HMMs, such as hidden semi-Markov models (HSMMs) have been widely studied as well (Joo et al. 2013). HSMMs are designed to ease the geometric constraint on the dwell-time distributions (the probability of leaving a behavioural state is the same at every time step) by explicitly modelling the dwell-time distributions (Langrock et al. 2012). A particular study by Joo et al. (2013) applied HSMMs to foraging tracks using ground-truth data on behavioural modes. Through a simulation experiment, it was shown that increasing time resolution from a coarse to a finer scale may decrease the accuracy of HMMs and conversely, increase the accuracy of HSMM inference. However, HSMMs slightly outperform HMMs for the subsampled dataset with a coarse scale (60 minutes). Future research could involve adopting the NSST approach before subsampling to determine if similar results still hold and to ascertain if there is a temporal resolution at which the performance of HMMs matches or surpasses that of HSMMs.

2. *Extending the NSST approach.*

    - Data type: The NSST approach assumes that the signals are uniformly sampled. In this thesis, the theorem has been applied to movement data with regular intervals. However, irregularly spaced movement data is very common in practice. Non-regular data also has practical interest. For instance, there are tags that use variable duty cycles (record at variable frequencies, where something triggers the tag to record at a higher frequency), and there is a whole set of interesting questions around how

best to set the frequencies to use with such tags. Exploring how NSST can be extended to apply to a non-regular movement dataset commonly encountered in practice is an interesting avenue for future research.

- Filter type: In the context of animal movement, NSST requires that the maximum frequency (i.e., fastest body movement) is known *a priori* to inform the choice of sampling interval for collecting movement data. In practice, the maximum frequency is hard to know. As such, the low-pass filter helps to remove aliasing that may be caused by sampling and ensures that the theorem is satisfied. There are various types of low-pass filters, such as Butterworth, windowed-sinc filters (Smith 1999). However, in this thesis, we used the basic form of the low-pass anti-aliasing filters (ideal low-pass filter). A possible area of future research is exploring the various types of low-pass filters to assess their suitability in the context of animal movement modelling. In particular, assessing how the various types of low-pass filters help in improving the minimal loss of information compared to the ideal low-pass filter.

Bibliography

Akeresola, R. A., Butler, A., Jones, E. L., King, R., Elvira, V., Black, J. & Robertson, G. (2024), 'Validating hidden Markov models for seabird behavioural inference', *Ecology and Evolution* **14**(3), e11116. https://doi.org/10.1002/ece3.11116.

Allaire, J. (2012), 'RStudio: integrated development environment for R', *Boston, MA* **770**(394), 165–171.

Allen, R. L. & Mills, D. W. (2004), *Signal analysis: time, frequency, scale, and structure*, John Wiley & Sons, New Jersey, USA. https://dsp-book.narod.ru/SATFSS1.pdf.

Ashworth, J. (2022), *Bird flu outbreak devastates UK seabird colonies*, Natural History Museum. https://nhm.ac.uk/discover/news/2022/june/bird-flu-outbreak-devastates-uk-seabird-colonies.

Bennison, A., Bearhop, S., Bodey, T., Votier, S., Grecian, W., Wakefield, E., Hamer, K. & Jessopp, M. (2018), 'Search and foraging behaviors from movement data: A comparison of methods', *Methods in Ecology and Evolution* **8**, 13–24. https://doi.org/10.1002/ece3.3593.

Boashash, B. (2015), *Time-frequency signal analysis and processing: a comprehensive reference*, Academic Press, Australia. https://pdfs.semanticscholar.org/d82c/0a17f328e30e861cb7e77507ed60c5aeabad.pdf.

Browning, E., Bolton, M., Owen, E., Shoji, A., Guilford, T. & Freeman, R. (2018), 'Predicting animal behaviour using deep learning: GPS data alone accurately predict diving in seabirds', *Methods in Ecology and Evolution* **9**(3), 681–692. https://doi.org/10.1111/2041-210X.12926.

Burger, A. E. & Shaffer, S. A. (2008), 'Perspectives in ornithology application of tracking and data-logging technology in research and conservation of seabirds', *The Auk* **125**(2), 253–264. https://doi.org/10.1525/auk.2008.1408.

Burnham, K. P. & Anderson, D. R. (2002), *Model selection and multimodel inference:*

a practical information-theoretic approach, Springer, New York, NY. https://doi.org/10.1007/978-0-387-22456-5_6.

Cant, E., Smith, A., Reynolds, D. & Osborne, J. (2005), 'Tracking butterfly flight paths across the landscape with harmonic radar', *Proceedings of the Royal Society B: Biological Sciences* **272**(1565), 785–790. https://doi.org/10.1098/rspb.2004.3002.

Cardoso, M. J., Modat, M., Vercauteren, T. & Ourselin, S. (2015), Scale factor point spread function matching: Beyond aliasing in image resampling, *in* 'Medical image computing and computer-assisted intervention – MICCAI 2015', Springer International Publishing, pp. 675–683. https://doi.org/10.1007/978-3-319-24571-3_81.

Cerna, M. & Harvey, A. F. (2000), 'The fundamentals of FFT-based signal analysis and measurement', *National Instruments, Junho* **54**. https://sjsu.edu/people/burford.furman/docs/me120/FFT_tutorial_NI.pdf.

Chapin, F., Zavaleta, E., Eviner, V., Naylor, R., Vitousek, P., Reynolds, H., Hooper, D., Lavorel, S., Sala, O., Hobbie, S., Mack, M. & Diaz, S. (2000), 'Consequences of changing biodiversity', *Nature* **405**, 234–242. https://doi.org/10.1038/35012241.

Chivers, L. S., Lundy, M. G., Colhoun, K., Newton, S. F., Houghton, J. D. & Reid, N. (2012), 'Foraging trip time-activity budgets and reproductive success in the black-legged kittiwake', *Marine Ecology Progress Series* **456**, 269–277. https://doi.org/10.3354/meps.

Choi, C., Peng, H., He, P., Ren, X., Zhang, S., Jackson, M. V., Gan, X., Chen, Y., Jia, Y., Christie, M., Flaherty, T., Leung, K.-S. K., Yu, C., Murray, N. J., Piersma, T., Fuller, R. A. & Ma, Z. (2019), 'Where to draw the line? using movement data to inform protected area design and conserve mobile species', *Biological Conservation* **234**, 64–71. https://doi.org/10.1016/j.biocon.2019.03.025.

Conners, M. G., Michelot, T., Heywood, E. I., Orben, R. A., Phillips, R. A., Vyssotski, A. L., Shaffer, S. A. & Thorne, L. H. (2021), 'Hidden Markov models identify major movement modes in accelerometer and magnetometer data from four albatross species', *Movement Ecology* **9**(1), 7. https://doi.org/10.1186/s40462-021-00243-z.

Cooke, S., Hinch, S., Wikelski, M., Andrews, R., Kuchel, L., Wolcott, T. & Butler, P. (2004), 'Biotelemetry: a mechanistic approach to ecology', *Trends in Evolution and Ecology* **19**, 334–343. https://doi.org/10.1016/j.tree.2004.04.003.

Croxall, J. P., Butchart, S. H., Lascelles, B., Stattersfield, A. J., Sullivan, B., Symes, A. & Taylor, P. (2012), 'Seabird conservation status, threats and priority actions: a global assessment', *Bird Conservation International* **22**(1), 1–34. https://doi.org/10.1017/S0959270912000020.

Dean, B., Freeman, R., Kirk, H., Leonard, K., Phillips, R. A., Perrins, C. M. & Guilford, T. (2013), 'Behavioural mapping of a pelagic seabird: combining multiple sensors and a hidden Markov model reveals the distribution of at-sea behaviour', *Journal of the Royal Society Interface* **10**(78). https://doi.org/10.1098/rsif.2012.0570.

Devi, R. & Pugazhenthi, D. (2016), 'Ideal sampling rate to reduce distortion in audio steganography', *Procedia Computer Science* **85**, 418–424. https://doi.org/10.1016/j.procs.2016.05.185.

Edelhoff, H., Signer, J. & Balkenhol, N. (2016), 'Path segmentation for beginners: an overview of current methods for detecting changes in animal movement patterns', *Movement Ecology* **4**(1), 21. https://doi.org/10.1186/s40462-016-0086-5.

Eglington, S. M. & Perrow, M. R. (2014), *Literature review of tern (Sterna & Sternula spp. foraging ecology. Report to JNCC)*, ECON Ecological Consultancy Ltd., Norwich.

Ehrlich, P. & Wilson, E. (1991), 'Biodiversity studies: science and policy', *Science* **253**(5021), 758–762. https://doi.org/10.1126/science.253.5021.758.

Energy, C. (2009), *Race Bank Offshore Wind Farm—Environmental Statement*, The Crown Estate, London.

European-Commission, Mézard, N., Sundseth, K. & Wegefelt, S. (2008), *Natura 2000 – Protecting Europe's biodiversity*, European Commission. https://data.europa.eu/doi/10.2779/45963.

Fischer, C. (2024), 'Challenges and opportunities when studying movement ecology in science and practical conservation', *Basic and Applied Ecology* **81**, 59–65. https://doi.org/10.1016/j.baae.2024.10.006.

Fischer, M., Parkins, K., Maizels, K., Sutherland, D., Allan, B., Coulson, G. & Di Stefano, J. (2018), 'Biotelemetry marches on: A cost-effective GPS device for monitoring terrestrial wildlife', *PLoS ONE* **13**(7), e0199617. https://doi.org/10.1371/journal.pone.0199617.

Gabor, D. (1946), 'Theory of communication. Part 3: Frequency compression and expansion', *Journal of the Institution of Electrical Engineers - Part III: Radio*

and Communication Engineering **93**, 445–457. https://doi.org/10.1049/ji-3-2.1946.007.

Gaurav, D. (2020), *Intuition behind log-loss score*, Towards Data Science. https://towardsdatascience.com/intuition-behind-log-loss-score-4e0c9979680a.

Gaylord, A. J., Sanchez, D. M. & Van Sickle, J. (2016), 'Choosing sampling interval durations for remotely classifying Rocky Mountain Elk behavior', *Journal of Fish and Wildlife Management* **7**(1), 213–221. https://doi.org/10.3996/042015-JFWM-034.

Ghysels, E. (1994), 'On the periodic structure of the business cycle', *Journal of Business & Economic Statistics* **12**(3), 289–298. https://doi.org/10.1080/07350015.1994.10524544.

Gillies, N., Fayet, A. L., Padget, O., Syposz, M., Wynn, J., Bond, S., Evry, J., Kirk, H., Shoji, A., Dean, B. et al. (2020), 'Short-term behavioural impact contrasts with long-term fitness consequences of biologging in a long-lived seabird', *Scientific Reports* **10**(1), 15056. https://doi.org/10.1038/s41598-020-72199-w.

Gutierrez-Galan, D., Dominguez-Morales, J. P., Cerezuela-Escudero, E., Rios-Navarro, A., Tapiador-Morales, R., Rivas-Perez, M., Dominguez-Morales, M., Jimenez-Fernandez, A. & Linares-Barranco, A. (2018), 'Embedded neural network for real-time animal behavior classification', *Neurocomputing* **272**, 17–26. https://doi.org/10.1016/j.neucom.2017.03.090.

He, P., Klarevas-Irby, J. A., Papageorgiou, D., Christensen, C., Strauss, E. D. & Farine, D. R. (2023), 'A guide to sampling design for GPS-based studies of animal societies', *Methods in Ecology and Evolution* **14**(8), 1887–1905. https://doi.org/10.1111/2041-210X.13999.

Hilbers, J. P., Huijbregts, M. A. & Schipper, A. M. (2020), 'Predicting reintroduction costs for wildlife populations under anthropogenic stress', *Journal of Applied Ecology* **57**(1), 192–201. https://doi.org/10.1111/1365-2664.13523.

Hong, S. & Coté, G. (2024), 'Arterial pulse wave velocity signal reconstruction using low sampling rates', *Biosensors* **14**(2), 92. https://doi.org/10.3390/bios14020092.

Hossin, M. & Sulaiman, M. N. (2015), 'A review on evaluation metrics for data classification evaluations', *International Journal of Data Mining & Knowledge Management Process* **5**(2), 1. https:/doi.org/10.5121/ijdkp.2015.5201.

Hounslow, J., Brewster, L., Lear, K., Guttridge, T., Daly, R., Whitney, N. & Gleiss, A. (2019), 'Assessing the effects of sampling frequency on behavioural

classification of accelerometer data', *Journal of Experimental Marine Biology and Ecology* **512**, 22–30. https://doi.org/10.1016/j.jembe.2018.12.003.

Hromada, S. J., Esque, T. C., Vandergast, A. G., Dutcher, K. E., Mitchell, C. I., Gray, M. E., Chang, T., Dickson, B. G. & Nussear, K. E. (2020), 'Using movement to inform conservation corridor design for mojave desert tortoise', *Movement Ecology* **8**(1), 38. https://doi.org/10.1186/s40462-020-00224-8.

Igual, J. M., Forero, M. G., Tavecchia, G., González-Solís, J., Martínez-Abraín, A., Hobson, K. A., Ruiz, X. & Oro, D. (2005), 'Short-term effects of data-loggers on Cory's shearwater (Calonectris diomedea)', *Marine Biology* **146**(3), 619–624. https://doi.org/10.1007/s00227-004-1461-0.

IUCN (2020), 'Red list summary statistics'. https://www.iucnredlist.org/resources/summary-statistics#Summary%20Tables.

Jeff, T. (2019), 'Humans are driving one million species to extinction', *Nature* **569**, 171. https://doi.org/10.1038/d41586-019-01448-4.

Jim, R. (2020), *Marine protection sites: How did we get here and where do we go from here?*, Natural England. https://naturalengland.blog.gov.uk/author/jim-robinson/.

JNCC (2020), *SPAs with marine components*, Joint Nature Conservation Committee. https://jncc.gov.uk/our-work/spas-with-marine-components/.

JNCC (2021), *Seabird population trends and causes of change: 1986–2019 Report*, Joint Nature Conservation Committee, Peterborough. https://jncc.gov.uk/our-work/smp-report-1986-2019.

Johnson, C. N. (2020), 'Past and future decline and extinction of species', *The Royal Society* . https://royalsociety.org/topics-policy/projects/biodiversity/decline-and-extinction/.

Johnson, D. S., London, J. M., Lea, M.-A. & Durban, J. W. (2008), 'Continuous-time correlated random walk model for animal telemetry data', *Ecology* **89**(5), 1208–1215. https://doi.org/10.1890/07-1032.1.

Johnson, L. R., Boersch-Supan, P. H., Phillips, R. A. & Ryan, S. J. (2017), 'Changing measurements or changing movements? Sampling scale and movement model identifiability across generations of biologging technology', *Ecology and Evolution* **7**(22), 9257–9266. https://doi.org/10.1002/ece3.3461.

Jonsen, I. D., Flemming, J. M. & Myers, R. A. (2005), 'Robust state-space modeling

of animal movement data', *Ecology* **86**(11), 2874–2880. https://doi.org/10.1890/04-1852.

Jonsen, I. D., Myers, R. A. & Flemming, J. M. (2003), 'Meta-analysis of animal movement using state-space models', *Ecology* **84**(11), 3055–3063. https://doi.org/10.1890/02-0670.

Joo, R., Bertrand, S., Tam, J. & Fablet, R. (2013), 'Hidden Markov models: the best models for forager movements?', *PLoS ONE* **8**(8), e71246. https://doi.org/10.1371/journal.pone.0071246.

Kabir, M. H., Ruman, U., Alam, S., Islam, S., Sultana, J. & Islam, M. M. (2021), 'Approximate shortest distance and direction between two places on the spherical earth and the oblate spherical earth', *GUB Journal of Science and Engineering* **8**(1), 49–56. https://doi.org/10.3329/gubjse.v8i1.62332.

Kawamoto, R., Nazir, A., Kameyama, A., Ichinomiya, T., Yamamoto, K., Tamura, S., Yamamoto, M., Hayamizu, S. & Kinosada, Y. (2013), 'Hidden Markov model for analyzing time-series health checkup data.', *MedInfo* **192**, 491–495. https://doi.org/10.3233/978-1-61499-289-9-491.

Lande, R. (1998), 'Anthropogenic, ecological and genetic factors in extinction and conservation', *Researches on Population Ecology* **40**(3), 259–269. https://doi.org/10.1007/BF02763457.

Langrock, R., King, R., Matthiopoulos, J., Thomas, L., Fortin, D. & Morales, J. M. (2012), 'Flexible and practical modeling of animal telemetry data: hidden Markov models and extensions', *Ecology* **93**(11), 2336–2342. https://doi.org/10.1890/11-2241.1.

Larsen, T. H., Williams, N. M. & Kremen, C. (2005), 'Extinction order and altered community structure rapidly disrupt ecosystem functioning', *Ecology Letters* **8**(5), 538–547. https://doi.org/10.1111/j.1461-0248.2005.00749.x.

Lascelles, B., Langham, G., Ronconi, R. & Reid, J. (2012), 'From hotspots to site protection: Identifying Marine Protected Areas for seabirds around the globe', *Biological Conservation* **156**, 5–14. https://doi.org/10.1016/j.biocon.2011.12.008.

Lascelles, B., Taylor, P., Miller, M., Dias, M., Oppel, S.and Torres, L., Hedd, A., Le Corre, M., Phillips, R., Shaffer, S. & Weimerskirch, H. (2016), 'Applying global criteria to tracking data to define important areas for marine conservation', *Diversity and Distributions* **22**(4), 422–431. https://doi.org/10.1111/ddi.12411.

Le Corre, M., Jaeger, A., Pinet, P., Kappes, M. A., Weimerskirch, H., Catry, T., Ramos, J. A., Russell, J. C., Shah, N. & Jaquemet, S. (2012), 'Tracking seabirds to identify potential Marine Protected Areas in the tropical western Indian Ocean', *Biological Conservation* **156**, 83–93. https://doi.org/10.1016/j.biocon.2011.11.015.

Leite, P. B. C., Feitosa, R. Q., Formaggio, A. R., Costa, G., Pakzad, K. & Sanches, I. D. (2008), Crop type recognition based on Hidden Markov Models of plant phenology, *in* '2008 XXI Brazilian Symposium on Computer Graphics and Image Processing', IEEE, pp. 27–34. https://doi.org/10.1109/SIBGRAPI.2008.26.

Leos-Barajas, V., Photopoulou, T., Langrock, R., Patterson, T. A., Watanabe, Y. Y., Murgatroyd, M. & Papastamatiou, Y. P. (2017), 'Analysis of animal accelerometer data using hidden Markov models', *Methods in Ecology and Evolution* **8**(2), 161–173. https://doi.org/10.1111/2041-210X.12657.

Lévesque, L. (2014), 'Nyquist sampling theorem: understanding the illusion of a spinning wheel captured with a video camera', *Physics Education* **49**(6), 697–705. https://doi.org/10.1088/0031-9120/49/6/697.

Liu, Y., Li, J., Pang, Y., Nie, D. & Yap, P.-T. (2023), The devil is in the upsampling: Architectural decisions made simpler for denoising with deep image prior, *in* 'Proceedings of the IEEE/CVF International Conference on Computer Vision', pp. 12408–12417.

Louzao, M., Bécares, J., Rodríguez, B., Hyrenbach, K. D., Ruiz, A. & Arcos, J. (2009), 'Combining vessel-based surveys and tracking data to identify key marine areas for seabirds', *Marine Ecology Progress Series* **391**, 183–197. https://doi.org/10.3354/meps08124.

MacNearney, D., Nobert, B. & Finnegan, L. (2021), 'Woodland caribou (rangifer tarandus) avoid wellsite activity during winter', *Global Ecology and Conservation* **29**, e01737. https://doi.org/10.1016/j.gecco.2021.e01737.

Marcus, R. J., Carbone, C., Kays, R., Kranstauber, B. & Jansen, P. A. (2012), 'Bias in estimating animal travel distance: the effect of sampling frequency', *Methods in Ecology and Evolution* **3**(4), 653–662. https://doi.org/10.1111/j.2041-210X.2012.00197.x.

Mavor, R., Heubeck, M., Schmitt, S. & Parsons, M. (2008), *Seabird numbers and breeding success in Britain and Ireland, 2006. (UK Nature Conservation, No. 31.)*, Joint Nature Conservation Committee, Peterborough.

McCann, R., Bracken, A., Christensen, C., Fürtbauer, I. & King, A. (2021), 'The relationship between GPS sampling interval and estimated daily travel distances in chacma baboons (Papio ursinus)', *International Journal of Primatology* **42**(4), 589–599. https://doi.org/10.1007/s10764-021-00220-8.

McClintock, B. T. (2021), 'Worth the effort? A practical examination of random effects in hidden Markov models for animal telemetry data', *Methods in Ecology and Evolution* **12**(8), 1475–1497. https://doi.org/10.1111/2041-210X.13619.

McClintock, B. T., Johnson, D. S., Hooten, M. B., Ver Hoef, J. M. & Morales, J. M. (2014), 'When to be discrete: the importance of time formulation in understanding animal movement', *Movement Ecology* **2**(1), 21. https://doi.org/10.1186/s40462-014-0021-6.

McClintock, B. T., London, J. M., Cameron, M. F. & Boveng, P. L. (2017), 'Bridging the gaps in animal movement: hidden behaviors and ecological relationships revealed by integrated data streams', *Ecosphere* **8**(3), e01751. https://doi.org/10.1002/ecs2.1751.

McClintock, B. T. & Michelot, T. (2018), 'momentuHMM: R package for generalized hidden Markov models of animal movement', *Methods in Ecology and Evolution* **9**(6), 1518–1530. https://doi.org/10.1111/2041-210X.12995.

McKellar, A. E., Langrock, R., Walters, J. R. & Kesler, D. C. (2014), 'Using mixed hidden Markov models to examine behavioral states in a cooperatively breeding bird', *Behavioral Ecology* **26**(1), 148–157. https://doi.org/10.1093/beheco/aru171.

Michelot, T., Langrock, R., Bestley, S., Jonsen, I. D., Photopoulou, T. & Patterson, T. A. (2017), 'Estimation and simulation of foraging trips in land-based marine predators', *Ecology* **98**(7), 1932–1944. https://doi.org/10.1002/ecy.1880.

Michelot, T., Langrock, R. & Patterson, T. A. (2016), 'moveHMM: an R package for the statistical modelling of animal movement data using hidden Markov models', *Methods in Ecology and Evolution* **7**(11), 1308–1315. https://doi.org/10.1111/2041-210X.12578.

Morales, J. M., Haydon, D. T., Frair, J., Holsinger, K. E. & Fryxell, J. M. (2004), 'Extracting more out of relocation data: building movement models as mixtures of random walks', *Ecology* **85**(9), 2436–2445. https://doi.org/10.1890/03-0269.

MTS (2022), *Calculate distance, bearing and more between Latitude and Longitude points*, Movable Type Scripts. http://www.movable-type.co.uk/scripts/latlong.html.

Murgatroyd, M., Underhill, L. G., Bouten, W. & Amar, A. (2016), 'Ranging behaviour of Verreaux's Eagles during the pre-breeding period determined through the use of high temporal resolution tracking', *PLoS ONE* **11**(10), e0163378. https://doi.org/10.1371/journal.pone.0163378.

Nathan, R., Monk, C. T., Arlinghaus, R., Adam, T., Alós, J., Assaf, M., Baktoft, H., Beardsworth, C. E., Bertram, M. G., Bijleveld, A. I., , Brodin, T., Brooks, J. L., Campos-Candela, A., Cooke, S. J., Gjelland, K. , Gupte, P. R., Harel, R., Hellström, G., Jeltsch, F., Killen, S. S., Klefoth, T., Langrock, R., Lennox, R. J., Lourie, E., Madden, J. R., Orchan, Y., Pauwels, I. S., Říha, M., Roeleke, M., Schlägel, U. E., Shohami, D., Signer, J., Toledo, S., Vilk, O., Westrelin, S., Whiteside, M. A. & Jarić, I. (2022), 'Big-data approaches lead to an increased understanding of the ecology of animal movement', *Science* **375**(6582), eabg1780. https://doi.org/10.1126/science.abg1780.

Parsons, M., Mitchell, I., Butler, A., Ratcliffe, N., Frederiksen, M., Foster, S. & Reid, J. B. (2008), 'Seabirds as indicators of the marine environment', *ICES Journal of Marine Science* **65**(8), 1520–1526. https://doi.org/10.1093/icesjms/fsn155.

Patterson, T. A., Basson, M., Bravington, M. V. & Gunn, J. S. (2009), 'Classifying movement behaviour in relation to environmental conditions using hidden Markov models', *Journal of Animal Ecology* **78**(6), 1113–1123. https://doi.org/10.1111/j.1365-2656.2009.01583.x.

Patterson, T. A., Parton, A., Langrock, R., Blackwell, P. G., Thomas, L. & King, R. (2017), 'Statistical modelling of individual animal movement: an overview of key methods and a discussion of practical challenges', *AStA Advances in Statistical Analysis* **101**(4), 399–438. https://doi.org/10.1007/s10182-017-0302-7.

Patterson, T. A., Thomas, L., Wilcox, C., Ovaskainen, O. & Matthiopoulos, J. (2008), 'State-space models of individual animal movement', *Trends in Ecology & Evolution* **23**(2), 87–94. https://doi.org/10.1016/j.tree.2007.10.009.

Perrow, M. R., Skeate, E. R. & Gilroy, J. J. (2011), 'Visual tracking from a rigid-hulled inflatable boat to determine foraging movements of breeding terns', *Journal of Field Ornithology* **82**(1), 68–79. https://doi.org/10.1111/j.1557-9263.2010.00309.x.

Phillips, R. A., Xavier, J. C. & Croxall, J. P. (2003), 'Effects of satellite transmitters on Albatrosses and Petrels', *The Auk* **120**(4), 1082–1090. https://doi.org/10.1093/auk/120.4.1082.

Pohle, J., Langrock, R., Van Beest, F. M. & Schmidt, N. M. (2017), 'Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement', *Journal of Agricultural, Biological and Environmental Statistics* **22**(3), 270–293. https://doi.org/10.1007/s13253-017-0283-8.

R Core Team (2024), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/.

Rabiner, L. (1989), 'A tutorial on hidden Markov models and selected applications in speech recognition', *Proceedings of the IEEE* **77**(2), 257–286. https://doi.org/10.1109/5.18626.

Rabiner, L. & Juang, B. (1986), 'An introduction to hidden Markov models', *IEEE ASSP Magazine* **3**(1), 4–16. https://doi.org/10.1109/MASSP.1986.1165342.

Reynolds, A. M., Cecere, J. G., Paiva, V. H., Ramos, J. A. & Focardi, S. (2015), 'Pelagic seabird flight patterns are consistent with a reliance on olfactory maps for oceanic navigation', *Proceedings of the Royal Society B: Biological Sciences* **282**(1811), 20150468. https://doi.org/10.1098/rspb.2015.0468.

Robertson, G., Bolton, M., Grecian, W. & Monaghan, P. (2014), 'Inter-and intra-year variation in foraging areas of breeding kittiwakes (Rissa tridactyla)', *Marine biology* **161**(9), 1973–1986. https://doi.org/10.1007/s00227-014-2477-8.

Robertson, G., Bolton, M., Grecian, W., Wilson, L., Davies, W. & Monaghan, P. (2014), 'Resource partitioning in three congeneric sympatrically breeding seabirds: foraging areas and prey utilization', *The Auk: Ornithological Advances* **131**(3), 434–446. https://doi.org/10.1642/AUK-13-243.1.

RSPB (2022), *Coquet Island seabird sanctuary*, The Royal Society for the Protection of Birds (RSPB). https://www.rspb.org.uk/helping-nature/what-we-do/protecting-species-and-habitats/coquet-island-seabird-reserve.

Ryan, P., Petersen, S., Peters, G. & Grémillet, D. (2004), 'GPS tracking a marine predator: the effects of precision, resolution and sampling rate on foraging tracks of African Penguins', *Marine Biology* **145**, 215–223. https://doi.org/10.1007/s00227-004-1328-4.

Schoombie, S., Wilson, R., Ropert-Coudert, Y., Dilley, B. & Ryan, P. (2024), 'The efficiency of detecting seabird behaviour from movement patterns: the effect of sampling frequency on inferring movement metrics in Procellariiformes', *Movement Ecology* **12**(1), 59. https://doi.org/10.1186/s40462-024-00499-1.

Shannon, C. E. (1949), 'Communication in the presence of noise', *Proceedings of the IRE* **37**(1), 10–21. https://doi.org/10.1109/JRPROC.1949.232969.

Smith, S. W. (1999), *The scientist and engineer's guide to digital signal processing (Second Edition)*, California Technical Publishing, San Diego, California. https://www.dspguide.com/pdfbook.htm.

Stanbury, A., Eaton, M., Aebischer, N., Balmer, D., Brown, A., Douse, A., Lindley, P., McCulloch, N., Noble, D. & Win, I. (2021), 'The status of our bird', *British Birds* **114**, 723–747. https://britishbirds.co.uk/sites/default/files/BB_Dec21-BoCC5-IUCN2.pdf.

Tatler, J., Cassey, P. & Prowse, T. A. A. (2018), 'High accuracy at low frequency: detailed behavioural classification from accelerometer data', *Journal of Experimental Biology* **221**(23), jeb184085. https://doi.org/10.1242/jeb.184085.

Thaxter, C. B., Lascelles, B., Sugar, K., Cook, A. S., Roos, S., Bolton, M., Langston, R. H. & Burton, N. H. (2012), 'Seabird foraging ranges as a preliminary tool for identifying candidate Marine Protected Areas', *Biological Conservation* **156**, 53–61. https://doi.org/10.1016/j.biocon.2011.12.009.

van Zinnicq Bergmann, M. P., Guttridge, T. L., Smukall, M. J., Adams, V. M., Bond, M. E., Burke, P. J., Fuentes, M. M., Heinrich, D. D., Huveneers, C., Gruber, S. H. et al. (2022), 'Using movement models and systematic conservation planning to inform Marine Protected Area design for a multi-species predator community', *Biological Conservation* **266**, 109469. https://doi.org/10.1016/j.biocon.2022.109469.

Visser, I., Raijmakers, M. E. & Molenaar, P. C. (2002), 'Fitting hidden Markov models to psychological data', *Scientific Programming* **10**(3), 185–199. https://doi.org/10.1155/2002/874560.

Vulcano, A., Rutherford, C., Staneva, A. & Mitchell, D. (2021), *How are seabirds doing in the EU and UK?*, BirdLife International. https://www.birdlife.org/news/2021/04/02/how-are-seabirds-doing-eu-and-uk/.

Wakefield, E. D., Owen, E., Baer, J., Carroll, M. J., Daunt, F., Dodd, S. G., Green, J. A., Guilford, T., Mavor, R. A., Miller, P. I., Newell, M. A., Newton, S. F., Robertson, G. S., Shoji, A., Soanes, L. M., Votier, S. C., Wanless, S. & Bolton, M. (2017), 'Breeding density, fine-scale tracking, and large-scale modeling reveal the regional distribution of four seabird species', *Ecological Applications* **27**(7), 2074–2091. https://doi.org/10.1002/eap.1591.

Wakefield, E. D., Phillips, R. A. & Matthiopoulos, J. (2009), 'Quantifying habitat use and preferences of pelagic seabirds using individual movement data: a review', *Marine Ecology Progress Series* **391**, 165–182. https://doi.org/10.3354/meps08203.

Whittington, J., Hebblewhite, M., Baron, R. W., Ford, A. T. & Paczkowski, J. (2022), 'Towns and trails drive carnivore movement behaviour, resource selection, and connectivity', *Movement Ecology* **10**(1), 17. https://doi.org/10.1186/s40462-022-00318-5.

Wijeyakulasuriya, D. A., Eisenhauer, E. W., Shaby, B. A. & Hanks, E. M. (2020), 'Machine learning for modeling animal movement', *PLoS ONE* **15**(7), e0235750. https://doi.org/10.1371/journal.pone.0235750.

Wilson, L. J., Black, J., Brewer, M. J., Potts, J. M., Kuepfer, A., Win, I., Kober, K., Bingham, C., Mavor, R. & Webb, A. (2014), *Quantifying usage of the marine environment by terns Sterna sp. around their breeding colony SPAs. JNCC Report No. 500*, Joint Nature Conservation Committee, Peterborough.

Winkler, D. W., Billerman, S. M. & Lovette, I. J. (2020), *Gulls, Terns, and Skimmers (Laridae), version 1.0. In Birds of the World.*, Cornell Lab of Ornithology, Ithaca, NY, USA. https://doi.org/10.2173/bow.larida1.01.

Wittemyer, G., Polansky, L., Douglas-Hamilton, I. & Getz, W. (2008), 'Disentangling the effects of forage, social rank, and risk on movement autocorrelation of elephants using Fourier and wavelet analyses', *Proceedings of the National Academy of Sciences of the United States of America* **105**(49), 19108—19113. https://doi.org/10.1073/pnas.0801744105.

Yu, H., Muijres, F., te Lindert, J., Hedenström, A. & Henningsson, P. (2023), 'Accelerometer sampling requirements for animal behaviour classification and estimation of energy expenditure', *Animal Biotelemetry* **11**(28). https://doi.org/10.1186/s40317-023-00339-w.

Zucchini, W., MacDonald, I. L. & Langrock, R. (2016), *Hidden Markov models for Time Series: An introduction using R, Second Edition (2nd ed.)*, Chapman and Hall, London. https://doi.org/10.1201/b20790.