Looking to the Past to Conserve Present-Day Diversity



Harri Ravenscroft

St Cross College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity, 2025

Supervised by Erin Saupe

Declaration

I declare that the work presented in this thesis is wholly my own work, except where otherwise stated. The views and opinions expressed herein are mine and not necessarily those of any other person or body unless so attributed.

Harri Ravenscroft

Farewell! I go to find the sun!

–J.R.R. Tolkien

Abstract

This thesis analyses, or develops tools for, the investigation of macroecological patterns of change across varied temporal and spatial scales. I have sought to develop new and innovative methods for the gathering and analysis of a diverse range of ecological data, alongside previously unimplemented strategies within the field of macroecology.

In the first study I developed and then tested the Historical Occurrence Georeferencing System (HOGS). This exploits historical distributional data in printed literature to generate previously untapped distributional occurrence information. HOGS was able to detect and classify large volumes of map-marker occurrence data in a fraction of the time taken for previous manual methods. It provides a rapidly deployable method for the generation of historical data, vital to accurate definition of baseline species conditions.

The second study determined the degree and factorial underpinnings of avian community shift within southern Africa by utilising community weighted mean temperature and precipitation metrics. I found that avian species responded strongly to these climatic changes and the multi-metric approach allowed for a more in-depth analysis than traditionally temperature-only strategies. Indices to measure the rate of change for avian communities employed alongside derivation of important predictive factors is highly useful when designing management strategies going forward.

The third study was conducted in partnership with Kew. I created a voronoi tessellation scheme for the assignment of areas of control in the analysis of socio-economic and wellbeing indicators surrounding Madagascan protected areas. This allowed for a more nuanced understanding of fine scale heterogeneity in poverty indicators that can inform the successful delivery of conservation and sustainable development interventions.

In the final study I developed a modular machine learning pipeline for the automated analysis of mollusc field images. It combines object detection and classification techniques to detect molluscs and drill holes, identify species, and enable biologically meaningful measurements. The system demonstrates strong performance across tasks, achieving high accuracy in mollusc detection, robust species classification across a range of dataset thresholds and reliable localisation and classification of drill holes.

Together, these studies add new methodologies for modern and historical macroecological assessment, which will be crucial in a world experiencing accelerating change.

It is the small everyday deeds of ordinary folk that keep the darkness at bay. Small acts of kindness and love.

–J.R.R. Tolkien

Acknowledgements

The acknowledgements are both the best and the most challenging part to write. They are without the burden of external scrutiny and give a rare opportunity to truly thank the people who have made everything included here possible. However, whether by space on page or an imperfect mind, not everyone can be thanked for all they have done. So before I begin personal acknowledgments, I would like to say thank you to all those who have, in a swirling myriad of ways, helped, encouraged, and even critiqued me throughout this time, it all matters.

Firstly, to the primo supervisor there could be, Erin Saupe. No matter the question, confusion, or issue she has always found the time to help with true kindness and honesty. Her encouragement for me to apply for the PhD and her patience with my constantly tangential mind is not something to be quickly forgotten. She is the epitome of what one hopes an academic could be and I have been exceedingly lucky.

The office of 20.17. There is not enough room even in a thesis to list all the ways in which each person has made the potentially boring office environment truly a place I looked forward to being in (which is impressive considering my opinion of having to be inside!). The quote wall that we kept from our first day together exemplifies the amazing friendships we have all made and the slightly weird people we all are. Lot, Lewis, Ally, Hannah, and Kate have all had a role in keeping me sane and encouraging all aspects of ridiculous plans, for which I am truly grateful.

James Borrell was in theory meant to be a three month supervisor whilst I worked at Kew, it's going to be more like 18 months now. He took me on with boundless enthusiasm and had confidence in my capabilities from the very start.

He also introduced me to wonderful people: Joe, Amber, and Adam being primary of these. Turns out an academic environment surrounded by lovely gardens and lovelier people isn't so bad.

My dear friends who are always kind and patient when I am worn out and constantly snoozy. Albert and Lina, who have been there every step from the first day of undergrad to today, and will be for many more. Emma and Rory, friends for life I couldn't imagine not knowing, and I am determined not to see that reality! John and Zoe for the countless coffee escapes to uni parks and the kindness in many areas of life. Ellie, who even from the far off lands of London continues in support. All the good people of OUTriC, of which there are too many to mention, who have been a fantastic source of sport-based distraction.

My family. Without all of you, now and for the last 28 years, this would never have been possible. To Mum and Dad, the constant kindness and support for everything I have tried across so many years of effort in all areas of life, through a myriad of ups and downs, to now. I would be nowhere without you both and I very much hope you know how very grateful I will always be. For Tom, Rhys, and Dafydd. There have never been a better set of brothers, and I have always known that if I could be anything like you all then I will have done exceedingly well. I have always been lucky to have those that I admire most so close to me from the very start.

Alice, my very greatest fortune and achievement over the last four years is, and always will be, meeting you. Even the hardest, most stressful, days can be made brilliant and ridiculous with you by my side. Here's to more adventures and the wide world ahead.

This PhD has been more than the 4 years of work that will be shown herein. It is a culmination of countless further hours of life, and the basis for more to come.

The road goes ever on and on,
Down from the door where it began.
Now far ahead the road has gone,
And I must follow, if I can,
Pursuing it with eager feet,
Until it joins some larger way
Where many paths and errands meet.
And whither then? I cannot say.

J.R.R. Tolkien

Contents

List of Figures **xiii**

List of Tables **xv**

1 Introduction **1**

1.1 Chapter 2: Release the HOGS 2

1.2 Chapter 3: Southern Africa Bird Shifts 3

1.3 Chapter 4: Madagascan Protected Areas 5

1.4 Chapter 5: Spot the Mollusc 6

1.5 Chapter Integration 7

2 Release the HOGS: An Unsupervised Marker Extraction, Classification, and Georeferencing Approach for Biodiversity Data **10**

2.1 Author Contributions 11

2.2 Abstract 12

2.3 Introduction 12

2.4 Methods 17

2.4.1 Overview of HOGS 17

2.4.2 Shape and Position Classification 23

2.5 Case Study 26

2.5.1 Evaluation of Empirical Model 27

2.6 Results 29

2.6.1 Testing Protocol Accuracy 29

2.7 Discussion 34

2.8 Conclusions 35

3 Evaluating Shifts in Southern African Bird Communities Using Temperature and Precipitation Indices 36

3.1 Author Contributions . . . 37
3.2 Abstract . . . 38
3.3 Introduction . . . 38
3.4 Materials and Methods . . . 41
3.4.1 Bird Data . . . 41
3.4.2 Temperature and Precipitation Data . . . 42
3.4.3 Species Temperature and Precipitation Indices . . . 43
3.4.4 Community Weighted Indices . . . 44
3.4.5 SAR Modelling . . . 46
3.4.6 Jackknife Species Contribution . . . 51
3.5 Results . . . 52
3.5.1 Community Weighted Indices . . . 52
3.5.2 Jackknife Species Contribution . . . 53
3.6 Discussion . . . 54
3.7 Conclusion . . . 59

4 Spatial Determinants of Community Wellbeing Near National Parks in Madagascar 61

4.1 Author Contributions . . . 62
4.2 Abstract . . . 63
4.3 Introduction . . . 64
4.4 Methods . . . 66
4.4.1 Study Area . . . 66
4.4.2 Household Poverty and Wellbeing Indicators . . . 67
4.4.3 Quantifying Community Land Use . . . 69
4.4.4 Landscape Variables . . . 71
4.4.5 Relationships Between Poverty and Wellbeing Indicators . . . 72
4.4.6 Predictors of Spatial Heterogeneity in Socio-economic Outcome Indicators . . . 73

4.5 Results . . . 74
4.5.1 National Scale Drivers of Heterogeneity in Poverty and Well-being . . . 75
4.5.2 Regional Scale Drivers . . . 78
4.6 Discussion . . . 80
4.6.1 Correlations Among Poverty and Wellbeing Metrics . . . 81
4.6.2 National and Local Scale: Reconciling Scale Mismatch . . . 82
4.6.3 Land-use Dynamics and Socio-ecological Patterns . . . 83
4.6.4 Limitations and Future Directions . . . 84
4.7 Conclusion . . . 85
4.8 Ethics Statement and Funding . . . 87

5 TIDAL: Rapid Identification and Classification of Mollusc Assemblages 89
5.1 Author Contributions . . . 90
5.2 Abstract . . . 91
5.3 Introduction . . . 92
5.3.1 Size Categorisation . . . 93
5.3.2 Drill Hole Use . . . 93
5.3.3 Previous Approaches . . . 94
5.3.4 TIDAL . . . 95
5.4 Materials and Methods . . . 97
5.4.1 Image Sourcing . . . 98
5.4.2 Field Collections . . . 99
5.4.3 Specimen Preparation . . . 101
5.4.4 Lab Imaging . . . 101
5.4.5 Labelling for Computer Vision . . . 102
5.4.6 Machine Learning Pipeline . . . 102
5.4.7 Real-world Size Estimation from £1 Coin Calibration . . . 108
5.5 Results . . . 109
5.5.1 Mollusc and £1 Coin Detection . . . 110

5.5.2 Species Classification 114
5.5.3 Drill Hole Detection and Classification 117
5.6 Discussion 118
5.6.1 Limitations and Future Work 119
5.7 Conclusion 122

6 Discussion and Conclusions **123**
6.1 Contribution to Knowledge 124
6.1.1 Chapter 2: Release the HOGS 124
6.1.2 Chapter 3: Southern Africa Bird Shifts 125
6.1.3 Chapter 4: Madagascan Protected Areas 125
6.1.4 Chapter 5: Spot the Mollusc 127
6.2 Limitations and Opportunities 128
6.2.1 Chapter 2: Release the HOGS 128
6.2.2 Chapter 3: Southern Africa Bird Shifts 129
6.2.3 Chapter 4: Madagascan Protected Areas 130
6.2.4 Chapter 5: Spot the Mollusc 131
6.3 Closing Remarks 132

Appendices

A Supporting Information for Chapter 2 **135**
A.1 Additional Information 135
A.1.1 Spatial Point Pattern Test 135
A.2 Additional Tables 137

B Supporting Information for Chapter 3 **139**
B.1 Additional Tables 139
B.2 Additional Plots 140

C Supporting Information for Chapter 4 **145**

C.1 Additional Information 145

C.1.1 Characterising Landscape Occupancy 145

C.2 Additional Tables 146

D Supporting Information for Chapter 5 **148**

D.1 Additional Information 148

D.1.1 Software and Hardware 148

D.2 Additional Plots 149

References **153**

List of Figures

2.1 Marker identification method 18
2.2 Occurrence map marker identification 19
2.3 Compound marker segmentation and splitting 22
2.4 Map alignment 25

3.1 CTI and CPI per pixel trends 53
3.2 Delta CTI and CPI predictors 54
3.3 Jackknife significant species contributions 55

4.1 Protected area map 66
4.2 Voronoi tessellation construction 70
4.3 Livelihood indicators split by protected area 74
4.4 Pairwise plots of economic indicators 75
4.5 Full RDA outcomes - all variables 77
4.6 Partial RDA outcomes - conditioned on ecological factors 80

5.1 Mollusc identification and data extraction protocol 97
5.2 TIDAL methodology flowchart 103
5.3 Multiple mollusc multi-phase detection 111
5.4 Mollusc and £1 multi-phase detection 112
5.5 Validation accuracy of model frameworks with varied threshold image counts 115
5.6 EfficientNetB0 fixed threshold (200) cross-dataset evaluation 116

B.1 Species stability curve 140
B.2 Valid cells given minimum species cutoff 141

B.3 Delta CTI boxplots at varied minimum species cutoffs per cell . . . 142
B.4 Delta CPI boxplots at varied minimum species cutoffs per cell . . . 143
B.5 Change in CPI versus baseline CPI 144
B.6 Change in CPI versus environmental heterogeneity 144

D.1 Pound coin detection - precision and recall 149
D.2 Mollusc detection - precision and recall 150
D.3 Model evaluation scores - image threshold 200 151
D.4 Model evaluation scores - EfficientNetB0 all thresholds 152

List of Tables

2.1 HOGS vs. Manual marker type and occurrence count 30
2.2 Nearest neighbour similarity measures 31
2.3 Decomposed standard distance comparison 32
2.4 Mean centre and standard distance comparison 33

5.1 Field collection localities 100
5.2 Mollusc image source breakdown 102
5.3 Species classification validation accuracy 114
5.4 Evaluation metric results for drill hole classification 117

A.1 Modified-SPPT test, no compound markers 137
A.2 Modified-SPPT test, all markers 137
A.3 HOGS processing times 138

B.1 SAR ΔCTI model terms 139
B.2 SAR ΔCPI model terms 139

C.1 All identified settlement occupancy metrics 146
C.2 Survey village occupancy metrics 146
C.3 RDA predictor variables 147

It's the job that's never started as takes longest to finish.

–J.R.R. Tolkien

1 Introduction

Contents

1.1 Chapter 2: Release the HOGS 2
1.2 Chapter 3: Southern Africa Bird Shifts 3
1.3 Chapter 4: Madagascan Protected Areas 5
1.4 Chapter 5: Spot the Mollusc 6
1.5 Chapter Integration 7

In this thesis I contribute to the development of usable macroecological techniques to identify both climatic and anthropogenic impacts on modern communities and ecosystems. The methods developed and applied herein allow for: rapid access to and use of historical map data for evaluation of species' distributional shifts (Chapter 2: Release the HOGS [2]); assessment of how and why avian communities shift in response to climate change to enable better planning of protected areas (Chapter 3: Southern Africa Bird Shifts [3]); determination of the socio-economic impact of protected areas in ecologically ravaged areas (Chapter 4: Madagascan Protected Areas [4]); and rapid re-surveying and data extraction from shifting marine mollusc communities (Chapter 5: Spot the Mollusc [5]).

Although heterogeneous in the taxonomic groups of study, the chapters included here provide new tool sets that allow for better understanding of the changing macroecological world around us. Using these approaches and by focusing on

quantifying biodiversity change over the last decades, I ask: how faunal communities have shifted in distribution over time in response to natural and anthropogenic factors; and whether our current protected area systems are likely to be effective now and in the future. Each chapter contributes a unique aspect to the puzzle of how we can protect and live alongside the diversity of the natural world.

1.1 Chapter 2: Release the HOGS

Climate change and anthropogenic pressure are expected to continue into the future, prompting changes in the distribution of organisms on Earth and both local and global extinctions across taxonomic levels (Pereira, Leadley et al. 2010; Pereira, Navarro and Martins 2012; Szabo et al. 2012; Pimm et al. 2014; Andermann et al. 2020).

To better contextualise these biodiversity changes, we need historical baselines prior to large-scale anthropogenic impacts (Bush 2023; Toth 2023; Kemp 2023). These historical baselines allow conservation practitioners to better assess the degree of change and to establish effective strategies to reconstitute natural conditions (Kiessling et al. 2019; Huang and Frimpong 2015; Seltmann et al. 2018).

However, anthropogenic activities have been influencing communities and ecosystems prior to the implementation of most modern systematic ecological surveying. The mismatch between the temporal coverage of biodiversity tracking and anthropogenic influence ultimately leads to shifting baseline syndrome (Bonebrake et al. 2010; Papworth et al. 2009; Clavero et al. 2023; Collins, Böhm and Collen 2020), wherein accepted thresholds for environmental and biodiversity standards are continually lowered (Soga and K. J. Gaston 2018).

Combatting this shift is challenging, but historical biodiversity data provide critical ecological insight (Nelson and Ellis 2018; Lister and Climate Change Research Group 2011; Marcer et al. 2022) and an enormous, and largely untapped, reservoir of information about distributional patterns globally (Seltmann et al. 2018; Marcer et al. 2022). To use this biodiversity data in studies of global change, precise knowledge of where specimens were collected, or sightings made, is paramount

(Marcer et al. 2022)—i.e., georeferenced data. Most georeferencing tools for map-based occurrence data are essentially manual, requiring the user to generate a point array over the inputted map image, which necessitates considerable data entry, panning, and on-screen manipulation (Burt et al. 2020) using software packages such as ArcGIS ((*ArcGIS Pro* 2024)) or QGIS ((*QGIS* 2024)). Although the approach can be highly accurate, it incurs a significant time cost and can be prone to user error.

In this chapter, I develop an unsupervised Historical Occurrence Georeferencing System (HOGS). Implemented using the Python programming language, HOGS extracts point occurrence markers from collections of digitised maps and assigns latitude/longitude coordinates to the occurrences. The sole inputs to this algorithm are scans or photographs of the historical maps in png format, and a small number of known coordinate/graticule points for training of a pixel-to-longitude/latitude model.

I test HOGS using two historical distribution atlases for sub-Saharan Passerine and Non-Passerine birds: "An Atlas of Speciation in African Passerine Birds" (Hall and Moreau 1970) and "An Atlas of Speciation in African Non-Passerine Birds" (Snow 1978). These atlases synthesise distributional, ecological, and evolutionary information, including all birds known to breed south of the Sahara ($20°N$) in Africa.

Our approach has the capability to fill temporal and spatial gaps in our knowledge of species' distributions that will allow for more accurate distributional models and a better understanding of pre-anthropogenic-acceleration baseline conditions.

1.2 Chapter 3: Southern Africa Bird Shifts

Climate change is affecting species' distributions, behaviours, and morphologies at a range of spatial and temporal scales (Lehikoinen et al. 2021; Devictor, Julliard et al. 2008; Devictor, Chris van Swaay et al. 2012). Species are generally responding to these climatic changes by shifting their distributions poleward (Pearce-Higgins, R. E. Green and R. Green 2014; Kelly and Goulden 2008; Auer and King 2014) and to higher elevations (Santangeli, Rajasärkkä and Lehikoinen 2017; Parmesan

and Yohe 2003). Some of the strongest signals of these distributional shifts have been apparent at the community level (Morecroft and Speakman 2013), with reorganisation over a variety of spatial scales, leading to unbalanced ecological interactions, local species' extirpations (Peterson and Martínez-Meyer 2008), and biotic homogenisation (Thuiller et al. 2014; Davey et al. 2013; Pierre Gaüzère, Jiguet and Devictor 2015; Montràs-Janer et al. 2024).

The impact of climate change on communities can be assessed using community weighted means (CWMs) (Curley et al. 2022; Sparrius, Top and Swaay 2018; Bowler and Böhning-Gaese 2017; Guetté et al. 2017; Princé and Zuckerberg 2015; Devictor, Julliard et al. 2008; Santangeli and Lehikoinen 2017; Lehikoinen et al. 2021; Hintsanen et al. 2023; Santangeli, Rajasärkkä and Lehikoinen 2017; Oliver et al. 2017). These metrics have been used to summarise the change in community composition over time in response to climate change (Bowler and Böhning-Gaese 2017) and to provide an indication of the occupancy of species within a community (Montràs-Janer et al. 2024).

In this chapter, I utilise spatiotemporal changes in two CWMs, the community temperature index (CTI) (Devictor, Julliard et al. 2008) and the community precipitation index (CPI) (Curley et al. 2022) to detect changes in the composition of communities and potential predictors of those changes.

I focus my analyses on southern Africa (defined here as South Africa, Lesotho, and Eswatini). In addition to ongoing climatic changes, southern Africa has been experiencing significant anthropization in the last century (Niedertscheider, Gingrich and Erb 2012). I utilise species' presence data from the first and second iterations of the South African Bird Atlas Project (SABAP) (gbif.org 2018), which is an extensive citizen science project spanning the years 1987-1992 (SABAP1) and 2007-modern (SABAP2), with over nineteen million occurrence records (Brooks et al. 2022).

Using these data, I seek to determine the trajectory of temperature and precipitation changes, both in magnitude and geographically, alongside understanding how communities are responding to these changes.

I found that temperature and precipitation climate data supported previously observed shifts of increasing temperature and decreasing precipitation (Jury 2013; Wit and Stankiewicz 2006). Additionally, communities appear to be responding to these shifts, reflected by increasing CTI, and decreasing CPI, values. Correlates of CTI and CPI change were found to be unconnected. Microclimate heterogeneity is negatively correlated with CTI change whereas the baseline value of CPI (this is the mean value for SABAP1, 1990-1991) is inversely correlated with delta CPI. This work is important for understanding the magnitude and predictors of community change in an often understudied area of the world where a rapidly increasing human population (Bai and Dent 2007), and accelerating climate change are increasingly altering the landscape (Niedertscheider, Gingrich and Erb 2012; Jury 2013).

1.3 Chapter 4: Madagascan Protected Areas

The protection of species within a changing landscape is often sought through the establishment of protected areas (Ribas et al. 2020; Stolton, Shadie and Dudley 2013). However, effective protected areas must not only conserve biodiversity, but also avoid disadvantaging surrounding communities (Brockington and Wilkie 2015; Gardner, Nicoll, Mbohoahy et al. 2013; Green et al. 2018; Hanauer and Canavire-Bacarreza 2015). Where protected areas result in significant local opportunity costs, positive conservation outcomes may be unsustainable in the long term (Gatiso et al. 2022).

Recognizing the drivers of inequality in the distribution of costs and benefits arising from protected areas is essential to develop equitable interventions that ameliorate these inequalities, and deliver more effective protected area management strategies (Ferraro 2002; Naidoo and Iwamura 2007). While numerous studies have previously sought to assess the overall poverty impacts of protected areas (Andam et al. 2010; Ferraro, Hanauer and Sims 2011; Hanauer and Canavire-Bacarreza 2015; Mammides 2020), particularly in Madagascar (Ferraro 2002; Gardner, Nicoll, Mbohoahy et al. 2013), these have generated varying and sometimes conflicting results.

I contribute to this debate by focusing on fine scale heterogeneity in multiple poverty indicators, sampled across a representative suite of eight protected areas in Madagascar. In addition, I present a novel approach to land-use partitioning among communities utilising a voronoi tessellation procedure. Finally, I consider how a more nuanced understanding of fine scale heterogeneity in poverty indicators can inform the successful delivery of conservation and sustainable development interventions.

I find that at the national scale, ecological gradients—particularly temperature and growing season length—are the dominant drivers of variation in economic outcomes. However, at the regional scale, independent of ecology, outcomes were shaped by land cover change, burned area coverage, and connectivity to major towns. These results suggest that economic outcomes in PA-adjacent communities are more strongly influenced by ecological setting and spatial infrastructure than by proximity to protected areas per se. Policymaking and conservation planning should therefore focus not only on spatial zoning but also on improving rural connectivity, supporting sustainable forest use, and accounting for the socio-economic diversity of communities across Madagascar's varied landscapes.

1.4 Chapter 5: Spot the Mollusc

Climate change and anthropization are powerful forces contributing to change in ecological communities. The dynamics and structure within modern communities are shaped by ecological interactions, predation being key among these. Environmental changes alter existing community dynamics and instigate new interactions that are detrimental to species (Warwick and Turk 2002; Hastings et al. 2020; Elvira S Poloczanska et al. 2013; Chivers, Walne and Hays 2017; García Molinos, Burrows and E S Poloczanska 2017).

Against this backdrop of changes and impacts, repeated, targeted surveys are required to establish baselines against which these dynamic alterations can be measured (Hiscock et al. 2001). Time-averaged death assemblages of shelly molluscs (gastropods and bivalves) are excellent sources of ecological information, as they

retain a strong signal of the rank order of relative species abundances across a large pool of habitat types (Warwick and Turk 2002; Kidwell 2001).

I develop the TIDES protocol, an automated machine learning pipeline capable of detecting molluscs in field images, classifying them by species, sizing them using a predefined scaling object (£1 coin), and identifying and classifying drill holes. The pipeline only requires the input of images to generate this data, thereby minimising the time- and knowledge-intensive aspects of mollusc surveying and dataset creation.

In addition to enabling rapid repeat re-surveying of community composition at fixed locations, the determination of prey and predator size is crucial to investigate predator-prey selection concepts and evaluate age composition within a community. Furthermore, the identification and study of drill holes on molluscs provides the means by which macroecological predation theories can be tested.

1.5 Chapter Integration

The ways in which climate change and anthropogenic environmental degradation are affecting ecological communities are numerous. This thesis cannot detail and combat all the issues the natural world faces, but the chapters described above can facilitate progress in overlapping areas of this endeavour.

Our development of the HOGS protocol (Chapter 2: Release the HOGS [2]) was conceived to provide a tool for simpler and more time-efficient extraction of species' occurrence data from maps and atlases. Application of the method combats shifting baseline syndrome, wherein a progressively lower standard of ecological state is accepted over time (Soga and K. J. Gaston 2018). The potential application of this technique links strongly to the analysis of avian distributional change through time (Chapter 3: Southern Africa Bird Shifts [3]). Our analysis made use of a comprehensive and ongoing citizen science project (SABAP (gbif.org 2018)), and I was able to determine the direction and factors correlating with community shifts over multiple decades (1987-2023). However, climate change and anthropization have been ongoing for a significantly longer period of time (Hockey et al. 2011;

Ziervogel et al. 2014; Warburton Toucher, Schulze and Maharaj 2005; Usman and Reason 2004; Richard et al. 2001; Niedertscheider, Gingrich and Erb 2012). In order to estimate the true scale of changes in communities, the pre-anthropogenic baseline condition would have to be known. If data derived from historical sources could be utilised here, then our knowledge of the true magnitude of change would increase dramatically and potentially allow for conservation innovations, including better protected area placement and conservation practices.

Understanding baseline conditions and ongoing changes is equally important in the marine realm. The fifth chapter (Spot the Mollusc [5]) details the construction and application of the TIDAL methodology. This tool allows users to rapidly generate taxonomic, morphological, and ecological data from mollusc specimens. Importantly, the use of TIDAL is not limited to modern nearshore surveying but can be applied retroactively to institutional collections. This facilitates the creation of historical datasets from which baseline conditions can be better understood. The use of HOGS alongside TIDAL could contribute further historical insight, including more effective designation and migration of marine protected areas in line with ongoing climatic shifts (Hilborn and Kaiser 2022; Arafeh-Dalmau et al. 2021; Fredston-Hermann, Gaines and Halpern 2018).

The implementation of marine and terrestrial protected areas has been shown to be critical to reducing the ongoing impacts experienced by species and communities (UNEP-WCMC and IUCN 2025; Greve et al. 2011; Ribas et al. 2020). However, less often considered in the macroecological context is the viability of these areas when more human-relevant factors are considered. Even with a perfect understanding of species' tolerances, a protected area is unlikely to succeed if anthropogenic pressures mean enforcement is heterogeneous or impossible altogether (Gatiso et al. 2022; Eklund, Blanchet et al. 2016; Eklund, Coad et al. 2019). Our fourth chapter (Madagascan Protected Areas [4]) interrogates the effect of protected area proximity to highly impoverished rural communities with Madagascar. I used an array of socio-economic outputs, derived from extensive surveying conducted by Kew and partner NGO organisations. These determined whether protected areas significantly impact

the wellbeing of local communities and whether they alter the way populations interact with natural capital at a national and regional scale.

The strands of research included are all interwoven. Together, they contribute to both a better understanding of macroecological change and provide some of the tools required to improve our response.

The Eagles are coming! The Eagles are coming!

–J.R.R. Tolkien

2

Release the HOGS: An Unsupervised Marker Extraction, Classification, and Georeferencing Approach for Biodiversity Data

Contents

2.1 Author Contributions 11
2.2 Abstract 12
2.3 Introduction 12
2.4 Methods 17
2.4.1 Overview of HOGS 17
2.4.2 Shape and Position Classification 23
2.5 Case Study 26
2.5.1 Evaluation of Empirical Model 27
2.6 Results 29
2.6.1 Testing Protocol Accuracy 29
2.7 Discussion 34
2.8 Conclusions 35

2.1 Author Contributions

Harri Ravenscroft developed the overarching research ideas and conceptual framework. Methodological design was jointly developed by Harri Ravenscroft and John Gardner. All data collection and data preparation were undertaken by Harri Ravenscroft. Analytical procedures and statistical analyses were performed by Harri Ravenscroft, with advisory input from John Gardner. The initial manuscript was drafted by Harri Ravenscroft, with Professor Erin Saupe contributing to revisions, interpretation of results, and final approval of the chapter.

2.2 Abstract

Distributional data are key for understanding how species and communities respond to abiotic and biotic changes. Large biodiversity occurrence databases have allowed for critical insight into the impact of anthropogenic activities on biodiversity, but these datasets often lack a historical perspective, with most records deriving from the last 25 years. Thus, modern assessments of biodiversity are prone to miss significant historical changes that occurred prior to large-scale record keeping. Developing methodologies to facilitate easy access to untapped historical occurrence data is therefore of pressing concern. To this end, we introduce a new georeferencing algorithm, Historical Occurrence Georeferencing System (HOGS), which exploits historical distributional data in printed literature. This algorithm uses a series of unsupervised steps to identify and isolate occurrence markers from distributional maps, assign each marker to a relevant taxonomic group, and georeference each input image relative to a known baseline coordinate system. Thus, the final output results in the assignment of latitude and longitude coordinates from a series of point occurrences on a map. We test our georeferencing pipeline using two historical atlases containing distributional maps for over 1,000 bird species in sub-Saharan Africa. We georeferenced 153,052 point occurrences from 751 input maps. A test subset of 1,898 occurrences yielded a point count discrepancy of only 5.4% when comparing to labour-intensive hand selection georeferencing approaches. Our new georeferencing protocol enables rapid generation of occurrence data from historical maps, establishing baseline conditions for species' distributions and ecological niche modelling for biogeographical, evolutionary, and conservation studies.

2.3 Introduction

Historical anthropogenic activities have impacted biodiversity through habitat degradation, species extirpation, and range constriction (Kemp 2023). Human impacts on biodiversity are expected to continue into the future and instigate

significant changes to the distribution of organisms on Earth, including widespread extinctions across taxonomic levels (Pereira, Leadley et al. 2010; Pereira, Navarro and Martins 2012; Szabo et al. 2012; Pimm et al. 2014; Andermann et al. 2020). To better contextualise these changes, we need historical baselines prior to large-scale anthropogenic impacts (Bush 2023; Toth 2023; Kemp 2023). Distributional baselines allow conservation practitioners to better assess the degree of change and to establish effective strategies to reconstitute those baseline conditions (Kiessling et al. 2019), manage resources, ensure land preservation and restoration, investigate the effects of climate change, among other applications (Huang and Frimpong 2015; Seltmann et al. 2018).

Most biodiversity data existing in current online repositories (Tobias et al. 2022; Sullivan et al. 2009; *GBIF* 2024a), however, do not provide a historical perspective because they are either too recent or are biassed geographically due to uneven sampling, data storage, and mobilisation efforts (Beck et al. 2014; Graham et al. 2004; Parker et al. 2024). Biodiversity monitoring schemes were largely initiated late in the 20th century, but most of the anthropogenic pressures currently impacting biodiversity have been operating over centuries or even millennia (Mihoub et al. 2017; Marsh and Trenham 2008; Lepetz et al. 2009; Lotze and Worm 2009). The mismatch between the temporal coverage of biodiversity and anthropogenic influence ultimately leads to shifting baseline syndrome (Bonebrake et al. 2010; Papworth et al. 2009; Clavero et al. 2023; Collins, Böhm and Collen 2020), wherein accepted thresholds for environmental and biodiversity standards are continually lowered. This leads to an underestimation of the full impact of anthropogenic activities (Mihoub et al. 2017), which is especially severe in tropical regions where sampling has been limited (Guralnick and Hill 2009) and where environmental shifts are likely to be largest in the coming decades (Sheldon 2019). To address this imbalance, it is increasingly important to utilise historical specimen records (Nowak, Słupecka and Jackowiak 2021).

Historical biodiversity data provide critical ecological insight (Nelson and Ellis 2018; Lister and Climate Change Research Group 2011; Marcer et al. 2022) and

an enormous, and largely untapped, reservoir of information about distributional patterns globally (Seltmann et al. 2018; Marcer et al. 2022). Few modern field surveys have operated for more than a few decades, but historical biodiversity data, either in documents or housed in museums or other academic institutions, can include samples and mapped distributional information that extend at least into the nineteenth century (Boakes et al. 2010). For example, the total number of preserved specimens held in institutions worldwide is on the order of two billion (Ariño 2010; A. Marcer et al. 2021). These historical records provide valuable insight from an interval of accelerating anthropogenic-driven habitat alteration and climate change (Lister and Climate Change Research Group 2011), which is useful for better understanding how biodiversity has been impacted by such changes over decades to centuries (Marcer et al. 2022; Baker et al. 2016; Papworth et al. 2009; McClenachan, Ferretti and Baum 2012; Jachowski et al. 2015; McCollin, Preece and Sparks 2015; Jones et al. 2020) and for making informed projections into the future (Baker et al. 2016; Hallman et al. 2021; Grace, Akçakaya, E. L. Bennett et al. 2021).

To use biodiversity data in studies of global change, precise knowledge of where specimens were collected, or sightings made, is paramount (Marcer et al. 2022)—i.e., georeferenced data. Georeferencing is the process of either verifying the associated geographic coordinates or assigning new ones to locality strings or occurrences on maps (Bloom, Flower and DeChaine 2018; Venceslau and Lopes 2019; A. Marcer et al. 2021; Thorn 2022). The uncertainty associated with historical distributional data usually increases with age due to (i) changing geographic terminology and recording standards over time, (ii) a legacy of inaccurate, imprecise, or vague locality coordinates or text descriptions by field researchers or collectors at the time (Murphy et al. 2004), (iii) imprecise information regarding the historical datum, and (iv) the absence of GPS technology.

Georeferencing protocols and automated procedures for locality string data exist and have been utilised broadly by the scientific community (Thorn 2022; A. Marcer et al. 2021; Marcer et al. 2022; Erp et al. 2015; Bloom, Flower and DeChaine

2018; Grover et al. 2010; Murphy et al. 2004). However, few protocols exist for georeferencing locality information from printed maps or distributional atlases.

Distributional maps represent data from reports of the occurrence of species from a snapshot of time, or can show the results of multi-year repeated data collection efforts (Macías-Duarte and Conway 2015; Peterson 2003; Lerm and Underhill 2019; Underhill and Brooks 2014; Pomeroy 1989). The species' occurrences recorded in high-resolution maps can be crucial in modelling distributions and in examining biodiversity change over time (Huang and Frimpong 2015; Underhill 2016). For example, researchers have been able to use historical presence information for one species, alongside known trophic community associations, to infer the absence of other associated species, allowing for the creation of presence-absence species' distribution models with historical context (Huang and Frimpong 2015).

Most georeferencing tools for map-based occurrence data are essentially manual, requiring the user to generate a point array over the inputted map image, which necessitates considerable data entry, panning, and on-screen data manipulation (Burt et al. 2020) using software packages such as ArcGIS (*ArcGIS Pro* 2024) or QGIS (*QGIS* 2024). Although the approach can be highly accurate, it incurs a significant time cost and can be prone to user error.

To date, and to our knowledge, there has been limited research into automated approaches for georeferencing point occurrences from historical maps that account for the difficulties inherent in early modern-era maps, which are often hand-drawn or lithography-printed (Luft 2020). Previous attempts to automate georeferencing of maps have focused on acquisition of thematic information, which used template matching and Hough line detection to automatically identify grid lines, map boundaries, and coordinate labels (Heitzler et al. 2018; Zatelli, Gabellieri and Besana 2022). Although effective, these methods focus on linear features or character strings, and therefore do not have the flexibility to detect point-level occurrences depicted in distributional atlases, which is often desired for biodiversity analyses.

Other existing approaches for the digitisation and georeferencing of maps are semi-automatic or rely on machine learning models (Burt et al. 2020; Murphy et al.

2004). Both of these methods, however, require a high degree of human input and are therefore time intensive. The supervised machine learning approach developed by Li and Xiao (2020), for example, is highly effective in detecting regional map zones and the projected coordinate system displayed. This supervised method has some potential in automated recognition of occurrence points within maps but would require human input for labelling and training marker (i.e., occurrence) data, with a time requirement perhaps on the same order of magnitude as manual point selection in GIS software.

Here we develop an unsupervised Historical Occurrence Georeferencing System (HOGS). Implemented using the Python programming language, HOGS extracts point occurrence markers from collections of digitised maps and assigns modern latitude/longitude coordinates to the occurrences. The sole inputs to this algorithm are scans or photographs of the historical maps in png format, and a small number of known coordinate/graticule points for the training of a pixel-to-longitude/latitude model. In HOGS, occurrence points are extracted from peripheral background lines and map information by applying a colour thresholding approach that can be tailored to the colour components of any given atlas or map. Alongside simple single occurrence points, complex and overlapping points, which may be discarded or misrepresented by other approaches, are defined by their pixel centroid position and then by their shape. All occurrences are provided as an output with shape, location, and original map information attached. Each stage within the protocol seeks to maximise the accuracy and proportion of data derived from any given input map while minimising human-induced variation and time expenditure from manual selection of occurrence points.

We test HOGS using two historical distribution atlases for sub-Saharan Passerine and Non-Passerine birds: "An Atlas of Speciation in African Passerine Birds" (Hall and Moreau 1970) and "An Atlas of Speciation in African Non-Passerine Birds" (Snow 1978). These atlases synthesise distributional, ecological, and evolutionary information, including all birds known to breed south of the Sahara ($20°N$) in Africa.

Our approach has the capability to fill temporal and spatial gaps in our knowledge of species' distributions that will allow for more accurate distributional models and a better understanding of pre-anthropogenic-acceleration baseline conditions.

2.4 Methods

2.4.1 Overview of HOGS

2.4.1.1 Initial Processing

Map Scans

HOGS is a georeferencing protocol that takes occurrence marker data from maps and automates the assignment of latitude and longitude coordinates. A single digitised image of each of the maps is required as input. Following this, HOGS consists of a series of steps that first isolates markers (i.e, taxon occurrences) of interest, then categorises these markers using an unsupervised clustering approach to identify different populations, species, sexes or other taxon identities, and finally generates georeferenced occurrences (marker positions), given in decimal degree coordinates.

Marker Separation

The first step in HOGS is to digitise the physical map and separate it into two layers: (i) the marker layer with occurrence points (2.2B), and (ii) the background layer with all other features. This is achieved via colour thresholding, which requires markers and background elements to differ in RGB values. Marker pixels are identified by calculating the mean RGB value for each pixel and flagging those below a user-defined threshold, which will be dependent upon the colours in the digital map copy. Common map colours, such as blue (RGB: 0, 0, 255) for aquatic features and green (RGB: 0, 255, 0) for vegetation, can be split easily from common black (RBG: 0, 0, 0) markers in this way.

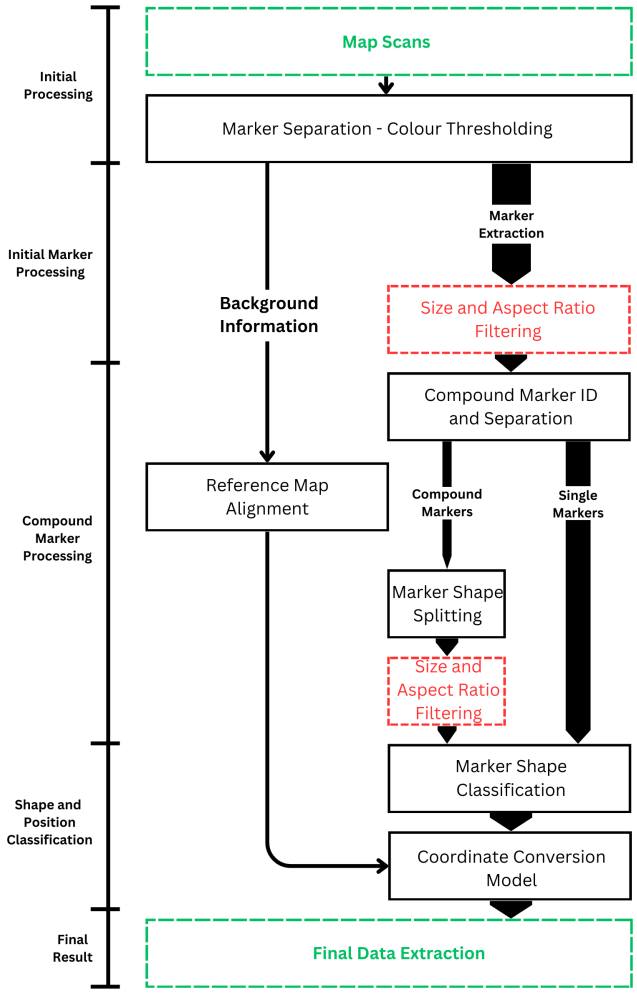


Figure 2.1: Georeferencing protocol steps in HOGS. Green dashed boxes denote inputs and outputs, red dashed boxes indicate data filtering steps, and black lined boxes represent HOGS processing steps. Inputs are high-resolution (typically $\geq$ 300 PPI) digitised historical maps **(Map Scans)**. Occurrence markers are separated from background features (e.g., boundaries, gridlines) and isolated **(Marker Separation)**. Size and aspect ratio thresholds distinguish single markers from complex, overlapping compound markers **(Compound Marker ID and Separation)**. Compound markers are split by identifying the outer contour, locating inflection points, selecting tie-lines, and dividing along these lines **(Marker Shape Splitting)**. Single and separated markers are pooled for PCA, then clustered by morphology using HDBSCAN **(Marker Shape Classification)**, enabling species assignment from legend information. Marker centroid pixel coordinates are recorded relative to each map. **Reference Map Alignment** occurs in parallel, aligning all maps to a reference via automated outline/feature detection and Bayesian-optimised transformation. After manual selection of known geographic features, pixel coordinates are converted to decimal degrees **(Coordinate Conversion Model)**. Final outputs are CSV files containing geographic coordinates (latitude and longitude), shape type, and map ID for each marker **(Final Data Extraction)**.

HOGS retains a background layer in order to later reference map alignment. Background layer creation requires that all the pixels identified as markers are replaced with white (RGB: 255, 255, 255), effectively removing the markers from the image, using our function "separate _markers _and _map". The remaining background, from which the occurrence markers have been removed, is retained.

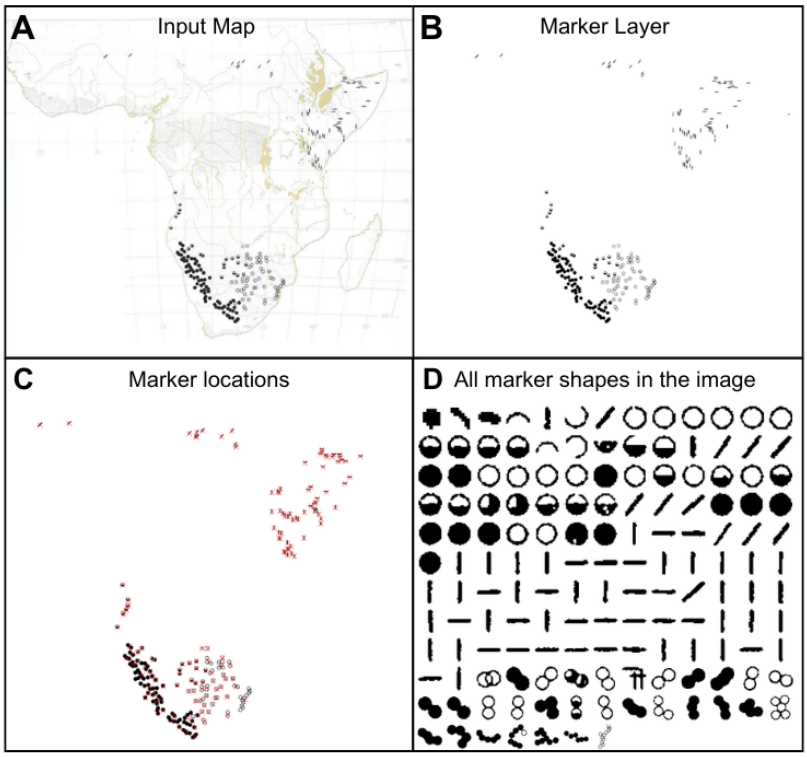


Figure 2.2: **A)** An example of a historical occurrence map prior to any alteration, showing seven different species marker types for sub-Saharan African birds from Hall and Moreau (1970). Markers display members of the *Calendulauda sabota* species group (previously *Mirafra sabota*). **B)** The occurrence marker layer has been isolated from the original map using the "separate_ markers_ and_ map" function in the "Point Separation" step of HOGS. **C)** Marked positions (red crosses) of all single occurrence points in an example from Hall and Moreau (1970), showing four different species and three 'intermediates' (now subspecies) marker types of the *Calendulauda sabota* species group. **D)** Different shapes initially found by the "get_ shapes_ and_ centres" function of HOGS. The intrinsic variability within the marker shapes is due to original printing inaccuracies in the source material. Both single and compound (multiple occurrences) markers were identified. The compound markers need to be split into individual markers in the "Marker Shape Separation" step (see 2.3).

2.4.1.2 Initial Marker Processing

Marker Extraction

Following isolation of the occurrence point layer and conversion to a black-white binary colour scale, the shape and location of each isolated occurrence must be extracted. This can be difficult when the occurrences overlap, which can occur in areas of high marker density (i.e., areas of high population density or regions with high numbers of sightings by collectors). Occurrence data may also be represented by different markers or shapes for different populations, species, or breeding status, which need to be distinguished. The "findContours" algorithm from the Python OpenCV catalogue (Bradski 2000) is therefore used to distinguish the outer contours of each occurrence point.

Our function "get_ shapes_ and_ centers" defines the contours (boundaries) of all the connected components (shapes) in the binary image marker as a list of points representing each shape's boundary. A contour is the curve that joins all the continuous points along a boundary with the same colour or intensity. Only external contours are retrieved, while nested contours are ignored, since any occurrence marker with internal complexity such as white space could generate additional, incorrect, occurrence points.

For each contour found, the bounding box coordinates (i.e, pixel values) are retrieved. The centroid position for the bounding box of each shape is also calculated as the average of the minimum and maximum x and y pixel coordinates. The bounding box coordinates and centroid locations for each shape are stored within an array.

Markers are then filtered by size and aspect ratio. The default aspect ratio threshold (<3) removes line markers, dust, and other artefacts. The minimum marker size is set to 20 pixels; at the high resolution baseline of 300 pixels per inch (PPI), occurrence markers are rarely smaller than 30–50 pixels. This threshold may be adjusted for lower-resolution maps or unusually small markers.

At this stage, the occurrence markers have not been grouped by shape, and overlapping/compound markers have not been identified

2.4.1.3 Compound Marker Processing

Compound Marker Identification and Separation

Occurrence points from the marker layer must be identified as either single or overlapping (i.e., or compound) markers. Compound markers are identified and filtered from single marker occurrences based upon size thresholding:

$$\text{Compound shape} = \text{Shape extent} > (\text{Median extent} \cdot \text{Outlier cutoff}) \quad (2.4.1.1)$$

Equation 1. How a compound shape is defined based upon size thresholding.

In HOGS, the default outlier size cutoff multiplier is 1.5. This value allows for variation from scanning errors and printing flaws, while excluding overlapped markers. The threshold can be adjusted in the "filter_ outliers" function based on marker size in the input maps.

Marker Shape Splitting

The compound markers pose an issue, as they represent a conglomeration of distributional data markers that cannot be easily discriminated (2.3A). Removal of these shapes would reduce the information retrievable from a given data source and produce misleading classifications of distributional patterns, often in areas where data are most dense. Therefore, the compound markers must be isolated into smaller, single occurrence components to return their corresponding shapes and centres. To address this, HOGS automatically splits compound markers into their individual components, identifying their shapes and centres so they can be treated as single occurrences. This is achieved using contour extraction and shape-splitting methods (2.3).

We build upon the contour extraction technique (Zou et al. 2021; Zafari et al. 2015) for the decoupling of compound markers in HOGS (2.3C).

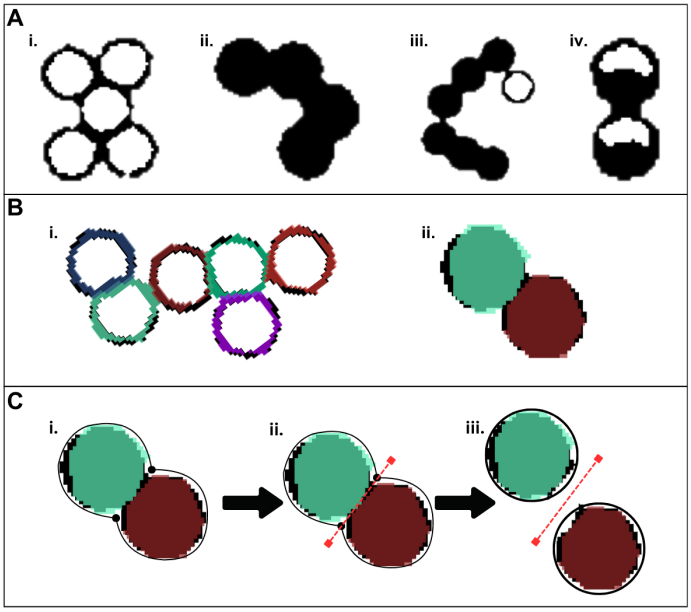


Figure 2.3: A) Examples of overlapping or compound markers derived from a single map input for three species from Hall and Moreau (1970). **i)** Five circular markers with no infill overlap only at their edges. Low levels of overlap means shapes can be more easily distinguished. **ii)** Four infilled circular markers. Some of the components are overlapped to a greater degree, reducing the depth of inflection point between markers. A reduction in inflection depth can lead to difficulty in establishing boundaries between overlapping shapes. **iii)** Six infilled circular markers and a singular circular marker without infill, representing a separate species, are connected to form a compound marker, demonstrating the possibility for cross-species data confusion. **iv)** The overlap of two complex markers with the same shape, exemplifying the internal complexity of some shapes and compounding effects of overlap on marker identification. **B)** Compound markers, each component of this complex shape represents a single occurrence point to be divided. **i)** A compound marker composed of six overlapping unfilled circles. **ii)** A compound marker composed of two overlapping filled circles. **C)** Schematic demonstrating the compound marker splitting process in HOGS, which is used to define clusters of occurrence points. These steps are defined within the function "extract_ single_ shapes_ and_ centres". **i)** The outer shape of the compound point is identified and the concave inflection points between the two components located. **ii)** The cost function identifies the points to join across this pinch point. **iii)** The two individual markers are divided along this tie-line.

The outer shape of each compound marker, defined as an external contour (2.3), is identified after converting the image to grayscale and inverting colours. The contour is extracted using OpenCV v4.10.0 (cv2.findContours) and simplified with the Ramer–Douglas–Peucker algorithm (rdp v0.8.0; (Hirschmann 2016)), which reduces the number of points while preserving shape (Ramer 1972; Douglas and Peucker 1973). Concave points are detected by analysing the turn orientation of

consecutive contour points ("extract_ concave_ points").

The contour is then split into a connected graph using the custom "connect_ the_ dots" function. This links point pairs to form chordless cycles, each representing a single occurrence within the compound marker. The process considers outline geometry and pinch points—narrow sections where a split is likely—using a cost function to decide which points to connect. The cost function, set automatically within "connect_ the_ dots", links two contour points if the chord length (straight-line distance) is below the maximum allowed and the resulting loop length (perimeter of the cycle) is acceptable. Pinch points are prioritised, and if the cost is below the threshold, an edge is added to the graph. This repeats for multiple point pairs.

Single markers are then identified sequentially. For each, the relevant contour section is extracted, its centroid calculated from bounding box coordinates, and its pixels taken from a binary mask. Shapes that are too small, fail the minimum size requirement, or have an unviable aspect ratio (<3.0 X:Y) are discarded, as are shapes whose bounding boxes completely contain others.

The final output is a set of single occurrence markers with centroid pixel coordinates relative to the original map. Processing time is typically under ten seconds per test map (A.3).

2.4.2 Shape and Position Classification

Marker Shape Classification

After deriving markers from the input map, we use clustering to assign shape labels to groups of occurrence markers, accommodating multiple marker types used to represent different taxonomic units, sexes or other taxon identities.

The approach used here groups markers to clusters in a three-stage process. First, all markers are reshaped to an n-by-n grid of the same size using our "resize_ shape" function. The default grid size value in our HOGS protocol is 40-by-40 pixels. This value can be modified based on the marker sizes for individual user maps. Resizing ensures that each marker has the same number of pixels, which is

necessary to flatten and process them with PCA. If markers are initially smaller, then zero-value pixels are padded on the edges to the desired dimensions, otherwise markers are scaled down, again using the "resize_ shape" function.

Second, two-dimensional marker shapes are flattened to a one-dimensional array for use in PCA using the NumPY function v2.1.0 array() in (Harris, Millman et al. 2020). Third, a clustering algorithm (HDBSCAN: Hierarchical Density-Based Spatial Clustering of Applications with Noise, (Campello, Moulavi and Sander 2013)) is used to cluster the markers based on their PCA embedding using the Python hdbscan v0.8.40 package (McInnes, Healy and Astels 2017). This clustering identifies marker types that belong to a particular taxonomic unit (or other object) shown on a map. HDBSCAN does not force every marker into a cluster.Instead, some markers may be classified as noise (outliers), indicated by a label of -1, and are therefore of "unknown" identity.

Reference Map Alignment

This step is relevant if multiple maps derive from the same source, with the same spatial extent. Each map from the same source is aligned to a single reference map from the dataset to automate georeferencing. The reference has markers removed via colour thresholding ("Marker Separation"). This approach assumes all maps in the set are identical, with only minor orientation or scale differences. Highly variable maps (e.g., different regions or projections) must be grouped separately. Alignment corrects small scanning and placement differences that, if uncorrected, could cause marker offsets of tens to hundreds of kilometres (2.4).

A tailored Bayesian optimisation function, implemented with the Optuna library v4.0.0 (Akiba et al. 2019), determines the optimal transformation parameters—translation (x, y), rotation (angle), and scaling—to align each processed map with the reference. The target function (2.4.2.1) minimises the difference between the transformed map and the reference. By default, the optimisation runs 300 trials, exploring translations of ±30 pixels, rotations of ±2°, and scaling from 0.99–1.01. These ranges and the number of trials can be modified within the "get_

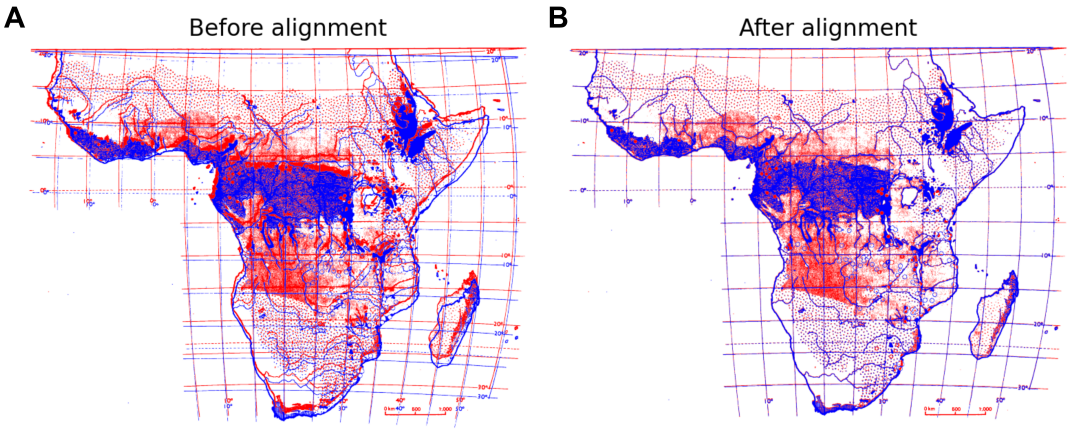


Figure 2.4: A) Disagreement between two historical occurrence maps due to scanning or digitisation inaccuracies. The blue map represents a baseline reference map against which the red map is being aligned. Both maps are sourced from Snow (1978). **B)** The results of the map alignment process, where a tailored Bayesian optimisation function ("get_ transformation") is used to find the optimal transformation parameters (translation, rotation, and scaling) to align all visible features to the reference map.

transformation" function. Each trial tests a unique parameter combination, and the set that produces the closest alignment is selected.

$$1 - \frac{\sum(A \cdot B)}{\sum A} \tag{2.4.2.1}$$

Equation 2. Target function to optimise for alignment of the reference map (A) to the target map (B) in 2.4. In this equation, A and B represent the 2D arrays of values that define the input figures. The target function is at a minimum (0) when B is identical to A, and increases as the maps become more different.

Coordinate Conversion Model

After aligning all input maps to the reference map and transforming the occurrence marker pixel coordinates using the same map-specific transformation, the pixel coordinates must be converted into geographically-relevant coordinates (i.e., latitude and longitude). The default conversion used by HOGS is decimal degrees in the World Geodetic System 1984 (WGS84). The default can be altered based on user preference within the protocol.

First, a dataframe of manually-selected known georeferenced points is created using the Python OpenCV v4.10.0 package (Bradski 2000). Pixel positions can be

chosen at points of interest such as graticule intersections, distinctive geographic features (e.g., mountain peaks, peninsula tips), or other features with known latitude and longitude, obtained from map content or external sources. Each point thus needs and has both geographic and pixel coordinates. A minimum of four points are required as the training input for a 3rd-order polynomial regression model (Rawlings, Pantula and Dickey 2006) to predict geographic coordinates (lat, lon) from pixel coordinates (x, y), although HOGS recommends ≥ 10 for improved accuracy.

Our protocol uses the 'scikit-learn' v1.5.2 package in Python (Pedregosa et al. 2011), which implements a 3rd-order polynomial, to model non-linear relationships between the coordinates. Using our "transform_ points" function, the marker positions with x-y pixel coordinates are transformed to align with a reference map using the Bayesian parameters determined for map alignment ("Reference Map Alignment"). Their geographic coordinates are then predicted using the trained model, and the output is saved in .csv format along with information on marker type and original map identity.

2.5 Case Study

HOGS helps to automate the process of assigning latitude and longitude coordinates to occurrence data represented on historical maps. To showcase its utility, we apply our new pipeline to a total of 751 historical maps from two atlases documenting bird distributions for Passerine (Hall and Moreau 1970) and Non-Passerine (Snow 1978) sub-Saharan African species. These atlases synthesise distributional, ecological, and evolutionary information, including all birds known to breed south of the Sahara ($20°N$) in Africa. The maps were not drawn to show absolute abundances, but do show species' ranges and broad ecological requirements. Distinctions were made for migration and breeding patterns when data were available, but this is not consistent throughout (Hall and Moreau 1970). Although the authorship changes between the two atlases, there is consistency in methodology and presentation. Bird presence

points for each map are unique in their number and location, but species' legend markers are highly variable between- and within- atlases.

Previously, studies have utilised some of the published information to study aspects of the geographic patterns of species' distribution, but, to our knowledge, none have made an attempt at digitisation of the complete works. Studies that utilised the atlases limited their scope to a subset of species, such as the 290 forest passerines (Diamond and Hamilton 1980), or deployed coarse scaling practices to reduce the challenges of manual data analysis, such as gridded quadrats across a range of spatial scales (1" (Burgess et al. 2000), 2.5" (Pomeroy and Ssekabiira 1990), and 5" (Diamond and Hamilton 1980)), with distributional data summarised within these geographical squares. This approach works well to identify continental-scale distributional patterns but lacks the precision to investigate more local or regional patterns.

We use HOGS to rapidly digitise the 751 distributional maps included in the two sub-Saharan bird atlases, which serves to validate our protocol and to demonstrate its capacity to access data for an understudied geographic area in a taxonomic group projected to be impacted by ongoing climatic changes (Jetz, Wilcove and Dobson 2007; Bradshaw, Sodhi and Brook 2009). Before digitisation, each species was taxonomically verified against the Clements Taxonomic Checklist (Clements, Rasmussen et al. 2023) to verify accepted taxonomic position and correct scientific naming (see Supplement: Clements Taxonomic Checklist).

2.5.1 Evaluation of Empirical Model

To evaluate the marker output from HOGS, we manually georeferenced a selection of historical distributional maps from Hall and Moreau (1970) and Snow (1978). Ten maps, five from each of the two atlases, were randomly selected for manual point selection (see Supplement: Test Maps). These maps include variation in marker densities, inter- and intra-map marker shapes, and the number of species displayed (i.e., per-map shape diversity).

2. Release the HOGS: An Unsupervised Marker Extraction, Classification, and Georeferencing Approach for Biodiversity Data

We compare georeferenced output from HOGS against manual georeferencing using ArcGIS Pro v3.2.1 (*ArcGIS Pro* 2024) using (i) The number of markers detected. High similarity between the number of markers detected is indicative of low rates of data omission; (ii) The count of species. Equal counts of species in a given map suggests HOGS is accurately detecting changes to marker shape, meaning taxa are not being misassigned; (iii) The Nearest Neighbour (NN) similarity metric. Low values of NN are indicative of spatial congruence between the manual and HOGS datasets. The similarity metric is measured by finding the occurrence from the manual set that is closest to each occurrence from HOGS, and vice-versa, and taking the mean of these two values. Values are calculated in longitude-latitude degree (WGS84 projection). Each occurrence point in both datasets (HOGS and manual) has a similarity score assigned in this way. (iv) Finally, we use two spatial deviation metrics to evaluate HOGS performance. We quantify the standard distance from the mean centre of the occurrence data to provide a summary measure of point distribution about this centre. The standard distance for a distribution is calculated as the square root of the mean square euclidean distance from the mean centre (2.5.1.1). We additionally quantified the decomposed standard distance, which is a measure of spread wherein latitude and longitude are analysed separately. For each map, highly similar values for these metrics between the manual and HOGS georeferencing indicate highly similar point distributions.

$$\mathrm{SD} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}((x_i - \overline{x})^2 + (y_i - \overline{y})} \qquad (2.5.1.1)$$

Equation 3. Equation used to calculate Standard Distance (SD) metric. n, the total number of points within the distribution, with each point having coordinates x_i, y_i. Coordinates, $\overline{x}, \overline{y}$ are those of the mean centre (centroid).

Calculations of discrepancy provide a way to compare the previously described HOGS-derived point data test metrics against the manually-counted values. Percentage discrepancy between HOGS versus manual latitude-longitude, and spatial deviation metric values for each given map is calculated following 2.5.1.2 below.

A baseline average discrepancy, wherein the sum of the HOGS-to-manual pair discrepancies is summed and divided by the number of pairs (2.5.1.2), provides a measure of relative protocol accuracy.

$$\text{Discrepency}_i = \frac{|x_i - R_i|}{R_i} \cdot 100 \tag{2.5.1.2}$$

Equation 4. Equation used to calculate individual discrepancies for each automated HOGS value x_i and its corresponding manually-counted reference value R_i. Values can be latitude-longitude coordinate differences, or for calculation of differences in spatial deviation metrics.

2.6 Results

A total of 159,852 marker points were identified, categorised, and exported in easily-readable data table format from the 750 maps across the two atlases (Hall and Moreau 1970; Snow 1978). The compound marker splitting step accounted for 38,933 of these occurrence points, or ~25.44% of all data extracted. The required hardware and computing time was low (A.3), making the process highly flexible for those without access to large-scale computing resources. Manual digitisation of each of the historical maps utilised in this study takes an average ~15-20 minutes. Comparatively, marker identification and splitting processes for 751 maps took ~12 minutes, and automated georeferencing and data frame creation adding ~3 minutes. Selection of points for the coordinate conversion model takes ~10 minutes.

2.6.1 Testing Protocol Accuracy

For each map, the manual species count based on different marker types matched the automated species count based on HOGS, indicating there was no omission or commission error (2.1).

Manual sampling across the ten maps detected a mean of 194.0 individual marker points for each map, and HOGS detected and classified an average of 185.3, a discrepancy of only 4.8%. Only three of the ten maps had an absolute difference of more than 10 markers, with two being exact marker number matches (2.1). The negative discrepancy means that HOGS failed to detect a small minority

of occurrence points across the ten maps. Most omissions were due to heavily-overlapping points that were too connected to allow for compound marker splitting. The degree of omission-causing overlap will vary by point shape and their relative shape orientation. In general, an overlap of approximately 90% or more will induce point omission. When two markers are overlaid almost perfectly, it is not possible to delimit their extents using the protocol detailed here, or using any other known approach.

Map		**Occurrence Number**			**Marker Count (i.e., Number of Species)**		
Atlas	**Map Number**	**Manual**	**HOGS**	**Percentage Discrepancy (%: 2d.p)**	**Manual**	**HOGS**	**Discrepancy (%)**
Passerine	94	251	258	2.71	2	2*	0.0
Passerine	100	75	75	0.00	2	2*	0.0
Passerine	180	75	68	9.33	3	3*	0.0
Passerine	192	141	134	4.96	2	2*	0.0
Passerine	273	252	250	0.79	5	5*	0.0
Non-Passerine	7	450	417	7.33	3	3*	0.0
Non-Passerine	28	63	66	4.76	1	1*	0.0
Non-Passerine	117	241	224	7.05	3	3*	0.0
Non-Passerine	178	118	118	0.00	1	1*	0.0
Non-Passerine	344	274	243	11.31	3	3*	0.0
Average	**NA**	**194.0**	**185.3**	**4.82**	**25**	**25**	**0.0**

Table 2.1: The number of occurrences identified per test map by automated (HOGS) and manual georeferencing, and the counts of the number of species (identified by a change in marker shape) for the same maps. Note that values signified with a "*" indicate the "unknown" marker shapes have been excluded in this count. See 2.5.1.1 and 2.5.1.2 for calculation of discrepancy.

The comparison of manually- and HOGS-derived point distribution parameters are indicative of congruent outputs across both atlases. Absolute differences in decomposed standard distance are $0.31°N$ and $0.39°N$ for longitude, and $0.47°N$ and $0.19°N$ for latitude for Passerine and Non-Passerine map groups, respectively (2.3).

Mean centre calculations output a mean absolute discrepancy of $0.14°N$ longitude and $0.37°N$ latitude for Passerine maps, and $0.65°N$ (Longitude) and $0.22°N$ (Latitude) for Non-Passerines (2.4). Aside from the discrepancy in longitude for Non-Passerine birds, these values fall within mean marker sizes.

The distribution, measured by standard distance, around the mean centre is similar between HOGS and manual georeferencing, with only one value exceeding 5% percentage discrepancy, with means of 4.65% (0.58°) and 1.70% (0.31°) across Passerine and Non-Passerine maps, respectively (2.4). The higher discrepancy (13.89%) for Passerine map 180 is a pattern found across all of the metrics, with

Map Type	**Non-Passerine**						**Passerine**					
Map Number	7	28	117	178	344	**Mean**	94	100	180	192	273	**Mean**
NN Similarity (degrees: 2s.f.)	0.47	0.49	0.20	0.54	0.17	**0.37**	0.23	0.39	0.87	0.44	0.23	**0.43**
NN Similarity (km: 1d.p.)	52.0	53.6	21.9	59.0	18.3	**40.96**	24.8	43.9	95.2	47.1	24.7	**47.14**

Table 2.2: Nearest neighbour similarity in degrees, derived from marker longitude-latitude positions (WGS84) for each of 10 test maps using automated (HOGS) and manual georeferencing.

higher discrepancy for spatial standard deviation, NN similarity, and relative marker counts.

Map		Longitude (degrees)			Latitude (degrees)		
Atlas	**Map #**	**Manual**	**HOGS**	**% Δ**	**Manual**	**HOGS**	**% Δ**
Passerine	94	7.17	6.71	13.95	5.94	5.78	2.69
Passerine	100	13.19	13.01	1.37	4.60	4.63	0.65
Passerine	180	6.43	5.58	13.22	11.19	9.62	14.03
Passerine	192	7.35	7.40	0.68	13.80	13.24	4.10
Passerine	273	8.18	8.19	0.12	12.37	12.42	0.40
Passerine	**Mean**	NA	NA	**5.87**	NA	NA	**4.37**
Non-Passerine	7	14.44	13.75	4.78	14.42	13.94	3.33
Non-Passerine	28	3.23	3.87	16.54	7.41	7.61	2.70
Non-Passerine	117	13.42	13.62	1.49	11.68	11.84	1.37
Non-Passerine	178	13.10	13.38	2.14	10.93	11.00	0.64
Non-Passerine	344	16.62	16.46	0.96	5.01	5.04	0.60
Non-Passerine	**Mean**	NA	NA	**5.18**	NA	NA	**1.73**

Table 2.3: Decomposed standard distance derived from marker longitude-latitude positions (WGS84) per test map using automated (HOGS) and manual georeferencing.

Map		Mean Centre and Standard Distance					
Atlas	**Map #**	**Manual Mean Centre (Lon - Lat)**	**HOGS Mean Centre (Lon - Lat)**	**Abs. Discrepancy (Lon - Lat)**	**Manual SD**	**HOGS SD**	**% Δ SD**
Passerine	94	30.15 -12.94	30.25 -13.02	0.10 0.08	9.29	8.84	4.84
Passerine	100	16.85 0.97	17.06 1.05	0.21 0.08	13.88	13.71	1.24
Passerine	180	31.26 -13.18	31.15 -13.77	0.11 0.59	12.82	11.04	13.89
Passerine	192	28.41 -13.54	28.25 -12.61	0.16 0.93	15.58	15.11	3.02
Passerine	273	27.23 -15.74	27.36 -15.56	0.13 0.18	14.80	14.84	0.27
Passerine	**Mean**	NA	NA	**0.14 0.37**	NA	NA	**4.65**
Non-Passerine	7	23.22 -6.52	23.89 -6.56	0.67 0.04	20.40	19.55	4.17
Non-Passerine	28	30.00 0.41	29.24 0.49	0.76 0.08	8.02	8.06	0.50
Non-Passerine	117	25.20 -5.23	25.46 -4.68	0.26 0.55	17.75	18.01	1.47
Non-Passerine	178	15.38 -0.82	14.79 -0.99	0.59 0.17	16.99	17.25	1.53
Non-Passerine	344	17.24 5.75	16.26 5.99	0.98 0.24	17.32	17.18	0.81
Non-Passerine	**Mean**	NA	NA	**0.65 0.22**	NA	NA	**1.70**

Table 2.4: Mean centre and standard distance from centre comparisons for each of the 10 test maps using automated (HOGS) and manual georeferencing.

2.7 Discussion

HOGS allows for rapid generation of usable spatial data from historical map sources. In our empirical case study, the application of this technique represents a time saving of almost three orders of magnitude with minimal data loss or alteration.

Analysis of differences between manually-selected marker points and automated-marker selection suggests our pipeline generates coordinate values that are highly similar both in terms of the number of species and occurrences identified compared to manual georeferencing protocols. The spatial position of the occurrences are also similar between manual georeferencing and HOGS.

These results demonstrate that HOGS produces outputs of comparable accuracy to the time-intensive manual georeferencing methods used previously, and suggests this approach may be useful for rapid generation of marker/occurrence data from historical, non-digitised, and non-evaluated sources.

Implementation and refinement of the methodology presented here can provide access to underutilised data prior to modern, large-scale biodiversity collection efforts. The only initial inputs required to implement this approach are versions of the historical data in the form of scans or photographs in png format, and known coordinate/graticule occurrence points for a limited number of initial training points. This built-in flexibility is important due to the highly-variable nature of historical map documentation.

Although we have attempted to make the protocol as accurate as possible, HOGS could be improved in future versions. First, additional in-depth testing of the compound marker splitting would be useful to determine those combinations of shapes that are most challenging to differentiate. Further testing of the compound-splitting protocol could include factors such as the percentage overlap between shapes, the morphology of the constituent shapes, and the variety of shapes within a single compound marker to assess their impacts on method accuracy. A thorough analysis of all factors relevant to compound marker complexity (marker shape; degree of overlap; number of points involved; resolution of marker; relative orientation etc), however, is beyond the scope of this study.

Second, the manual selection of points with known latitude and longitude coordinates required for the pixel coordinate to geographic coordinate conversion requires user input, slowing the protocol and increasing human expertise requirements. Following Burt et al. (2020), automated location of map boundaries and ground control points to produce georeferenced images without operator assistance could be pursued. HOGS also requires user input of certain parameters if defaults are inappropriate. Automated detection of input map colouration would eliminate the need for manual selection of colour thresholds (Chesneau 2011), and map resolution identification could simplify minimum pixel size thresholding for single versus compound marker discrimination, thereby improving protocol efficiency.

The automated detection and georeferencing of points described here, however, represents a substantial step forward in rapidly digitising historical map data. HOGS greatly reduces the time cost of previous manual selection strategies whilst maintaining a high degree of precision and accuracy. Automation can facilitate the integration of untapped historical data on species' distributions that can be used in myriad applications, such as ecological niche modelling for biogeographical, evolutionary and conservation studies, allowing for more robust baselines and actionable conservation outcomes in line with pre-anthropogenic targets.

2.8 Conclusions

The protocol presented in this study represents an important step in the rapid generation of historical distributional data from often-overlooked but highly valuable biodiversity maps and atlases. Such data will be important if we are to develop reliable species' distribution baselines and combat the expanding impact of anthropogenic activities. Validation of our protocol through analysis of ten historical maps (Hall and Moreau 1970; Snow 1978) demonstrates the viability of generating new data sources from old. Continued improvement of HOGS, and its application to other historical maps and atlases, is needed into the future.

All have their worth and each contributes to the worth of the others.

–J.R.R. Tolkien

3

Evaluating Shifts in Southern African Bird Communities Using Temperature and Precipitation Indices

Contents

3.1 Author Contributions 37
3.2 Abstract 38
3.3 Introduction 38
3.4 Materials and Methods 41
3.4.1 Bird Data 41
3.4.2 Temperature and Precipitation Data 42
3.4.3 Species Temperature and Precipitation Indices 43
3.4.4 Community Weighted Indices 44
3.4.5 SAR Modelling 46
3.4.6 Jackknife Species Contribution 51
3.5 Results 52
3.5.1 Community Weighted Indices 52
3.5.2 Jackknife Species Contribution 53
3.6 Discussion 54
3.7 Conclusion 59

3.1 Author Contributions

Harri Ravenscroft developed the overarching research ideas and conceptual framework. All data collection and data preparation were undertaken by Harri Ravenscroft. Analytical procedures and statistical analyses were performed by Harri Ravenscroft, with advisory input from Professor Erin Saupe. The initial manuscript was drafted by Harri Ravenscroft, with Professor Erin Saupe contributing to revisions, interpretation of results, and final approval of the chapter.

3.2 Abstract

Climate change is rapidly altering the distribution of species, reshuffling ecological communities across spatial scales. Community reorganisations can lead to unbalanced ecological interactions, species' extirpations, and biotic homogenisation. Here we study how bird assemblages have changed across southern Africa between 1987 to 2023 using data from the first and second iterations of the Southern African Bird Atlas Project, characterised by more than 600 species.

We use spatial autoregressive (SAR) models to determine how and why community weighted mean temperature and precipitation indices (CWMs) have changed over this interval. We find that avian assemblages within southern Africa have warmed by ~0.07°C and become drier by ~0.41 mm over the last 36 years. The correlates of these changes differ for precipitation and temperature, such that change in the community temperature index was negatively correlated to increased environmental heterogeneity, whereas change in the community precipitation index was negatively influenced by the baseline index value.

To guide effective conservation and future protected area delineation, it is important to understand how assemblages are changing spatially across southern Africa and the potential drivers behind those changes.

3.3 Introduction

Climate change is impacting species' distributions, behaviours, and morphologies at a range of spatial and temporal scales (Lehikoinen et al. 2021; Devictor, Julliard et al. 2008; Devictor, Chris van Swaay et al. 2012). Species are responding to these climatic changes by shifting their distributions, with general movement both poleward (Pearce-Higgins, R. E. Green and R. Green 2014; Kelly and Goulden 2008; Auer and King 2014) and to higher elevations (Santangeli, Rajasärkkä and Lehikoinen 2017; Parmesan and Yohe 2003). These shifts are apparent at the community level in their "warming" (Morecroft and Speakman 2013), in which the

abundance of warm-adapted species increases over time relative to cool-adapted species under increasing ambient temperatures (Parmesan and Yohe 2003; De Frenne et al. 2013; Khaliq et al. 2024; Bowler and Böhning-Gaese 2017).

Rapid reorganisation of community structures over a variety of spatial scales can lead to unbalanced ecological interactions, species' extirpations (Peterson and Martínez-Meyer 2008), and biotic homogenisation (Thuiller et al. 2014; Davey et al. 2013; Pierre Gaüzère, Jiguet and Devictor 2015; Montràs-Janer et al. 2024). Communities have become increasingly homogenised as generalist species increase in abundance relative to specialist species (Pearce-Higgins, Eglington et al. 2015; Devictor, Julliard et al. 2008; Gaüzère et al. 2020), since the former tend to be better able to withstand abiotic and biotic changes (Curley et al. 2022).

The impact of climate change on communities has been assessed previously using community weighted temperature and precipitation means (CWMs) (Curley et al. 2022; Sparrius, Top and Swaay 2018; Bowler and Böhning-Gaese 2017; Guetté et al. 2017; Princé and Zuckerberg 2015; Devictor, Julliard et al. 2008; Santangeli and Lehikoinen 2017; Lehikoinen et al. 2021; Hintsanen et al. 2023; Santangeli, Rajasärkkä and Lehikoinen 2017; Oliver et al. 2017). CWMs summarise the shift in community composition over time in response to climatic change (Bowler and Böhning-Gaese 2017) and provide an indication of the relative occupancy of species within a community (Montràs-Janer et al. 2024). Previous studies have found that relatively warm-adapted species increase in abundance in regions where temperatures have become favourable to them at the expense of cold-adapted species, the latter of which may experience range contraction and local extinction (Lajeunesse and Fourcade 2023; Curley et al. 2022).

Here, we use two CWMs, the community temperature index (CTI) (Devictor, Julliard et al. 2008) and community precipitation index (CPI) (Curley et al. 2022), to detect spatiotemporal changes in the composition of avifauna assemblages and potential predictors of those changes over the last 26 years. Birds have been shown to be sensitive to both precipitation (Illán et al. 2014; Pearce-Higgins, Eglington et al. 2015; Curley et al. 2022; VanDerWal et al. 2012; Tingley et al. 2009; Bateman

et al. 2016) and temperature (Freeman and Class Freeman 2014; Santangeli and Lehikoinen 2017; Devictor, Julliard et al. 2008; Lehikoinen et al. 2021) in previous work. Previous studies have shown birds have responded both to widespread climatic shifts and to anthropogenic changes by shifting their distributions (La Sorte and Thompson 2007; Montràs-Janer et al. 2024; Hockey et al. 2011; Buchan et al. 2022), potentially affecting community composition (Lindström et al. 2013; Princé and Zuckerberg 2015; Guetté et al. 2017; Ghosh, B. Matthews and Petchey 2024).

We focus our analyses on southern Africa (defined here as South Africa, Lesotho, and Eswatini), as this region may be especially vulnerable to climate change (Lennard et al. 2018). In South Africa, for example, temperatures have increased by at least 1.5 times the observed global average of $0.65°C$ over the past five decades (Engelbrecht et al. 2015; Wright et al. 2021), and the warmest years on record have occurred in the last decade (Blunden and Arndt 2020). Extreme rainfall events have also increased in frequency (Hockey et al. 2011; Ziervogel et al. 2014; Warburton Toucher, Schulze and Maharaj 2005), alongside interannual variability in rainfall (Richard et al. 2001; Usman and Reason 2004). These climatic changes have increasingly altered the landscape (Niedertscheider, Gingrich and Erb 2012). Net primary productivity in South Africa decreased against a background of human population increase of approximately 50% between 1981 and 2003 (Bai and Dent 2007).

We utilise species' presence data from the first and second iterations of the Southern African Bird Atlas Project (SABAP1: 1987–1992; SABAP2: 2007–present) (gbif.org 2018; Brooks et al. 2022). We test how avian assemblages across southern Africa have changed by quantifying shifts in the Community Temperature Index (CTI) and Community Precipitation Index (CPI) between these periods. Specifically, we hypothesise that (i) regional CTI and CPI will correlate positively with long-term trends in temperature and precipitation (Hockey et al. 2011; S. N. Matthews et al. 2011; Rushing et al. 2020; Auer and King 2014; Blumstein 2006; Gillings, Balmer and Fuller 2015); (ii) assemblages with higher baseline CTI or CPI values will exhibit smaller subsequent changes, reflecting reduced scope for further shifts towards

climatic extremes (Şekercioğlu, Primack and Wormworth 2012; Herrando et al. 2019); (iii) areas experiencing greater human modification or population density growth will show larger CWM changes due to elevated ecological disturbance and turnover; and (iv) environmental heterogeneity will buffer assemblages against regional climate trends by providing fine-scale climatic refugia (Curley et al. 2022; Hugo and Rensburg 2008; Rensburg, Chown and Gaston 2002). We then evaluate the relative influence of baseline conditions, anthropogenic pressures, climate trends, and microclimatic heterogeneity on spatial variation in CTI and CPI change.

3.4 Materials and Methods

3.4.1 Bird Data

All statistical analyses and data manipulation were performed using RStudio v4.3.3 (Posit team 2024). We downloaded avian occurrence data at species level from GBIF (*GBIF* 2024a) using the "rgbif" v3.8.1 package (Chamberlain, Barve et al. 2024), specifying the countries of South Africa, Lesotho, and Eswatini for the years 1987 to 2023.

We selected this region due to the availability of occurrence data from the SABAP1 (gbif.org 2018) and SABAP2 (Underhill 2016; gbif.org 2018) citizen science monitoring programmes. The two SABAP (Southern African Bird Atlas Project) programmes ran from 1987-1992 (with most data collected between 1987-1991), and from 2007 to the present day. We used bird occurrence data from GBIF from within two, five-year collection periods of the SABAP citizen science projects: 1987-1991 and 2019-2023. The later SABAP2 interval (i.e., 2019-2023) was selected to maximise the temporal gap between the two time intervals so as to examine change over a longer time interval.

Using "rgbif" v3.8.1 (Chamberlain, Barve et al. 2024) functions, we downloaded SABAP occurrence data for each species in a looped protocol, retaining records based on: (i) "Human Observation" or "Observation" to avoid inclusion of fossil or museum samples, (ii) "present" occurrence status, and (iii) coordinates more

than 2,000 metres from a country centroid, to avoid samples lacking coordinates that have been assigned to the centre of the country. The "CoordinateCleaner" v3.0.1 package (Zizka et al. 2019) was then used to flag and remove records with potential geographical issues, such as zero-coordinates, assignment to a capital city centroid, or any coordinates outside the countries of interest.

Both SABAP initiatives collected occurrence data using a defined grid format and species checklists. The checklist was used by the field recorder to note the presence of taxa (Brooks et al. 2022). Only full protocol checklist data were included, in which the observer is expected to identify at least 90% of species encountered whilst visiting the area for at least two hours and up to five days (Lee, Altwegg and Barnard 2017).

The grid format for SABAP1 used quarter degree grid cells (15'×15') as the sampling unit, corresponding to standard southern African 1:50,000 topographical maps (Harebottle et al. 2007). The SABAP2 spatial sampling unit was a 5'×5' pentad, with nine pentads nested within each quarter degree grid cell (Underhill 2016). We aggregated the data to 0.5°×0.5° resolution, to match that of our climate data (see section 3.3.2), and retained only unique species records per grid cell (i.e., we consider only presence/absence data, not abundance data).

3.4.2 Temperature and Precipitation Data

We use two CWMs, the community temperature index (CTI) (Devictor, Julliard et al. 2008) and community precipitation index (CPI) (Curley et al. 2022), to detect spatiotemporal changes in the composition of avifauna communities and potential predictors of those changes over the last 26 years. We focus on both temperature and precipitation because both have been shown to constrain the distribution of birds in previous work (Illán et al. 2014; Pearce-Higgins, Eglington et al. 2015; Curley et al. 2022; VanDerWal et al. 2012; Tingley et al. 2009; Bateman et al. 2016; Freeman and Class Freeman 2014; Santangeli and Lehikoinen 2017; Devictor, Julliard et al. 2008; Lehikoinen et al. 2021). Highly vagile species, such as birds, are expected to track their climatic niches defined by temperature and precipitation

(Tingley et al. 2009). In topographically-complex environments, temperature is thought to be a primary driver of community change at high elevations, whereas precipitation may better constrain community turnover in lowlands (Enquist 2002; Şekercioğlu, Primack and Wormworth 2012; Parra and Monahan 2008).

Monthly mean temperature ($°C$) and precipitation (mm/month) data were obtained from the Climatic Research Unit gridded Time Series (CRU-TS v4), a global climate dataset on a 0.5° latitude by 0.5° longitude grid (Harris, Osborn et al. 2020). For each year of the study, climate data were available as monthly layers and were aggregated into yearly means for each grid cell.

3.4.3 Species Temperature and Precipitation Indices

Community weight means (CWMs) for temperature and precipitation were derived from the species temperature index (STI) (Devictor, Julliard et al. 2008) and species precipitation index (SPI) (Curley et al. 2022). STI and SPI are calculated here as the average temperature or precipitation experienced by a species across their geographical range during a specified window of time (Devictor, Julliard et al. 2008). These values reflect the realised niche (Soberón 2007) of a species when calculated from all known geographic range information. The realised niche is usually a subset of the fundamental niche in which biological interactions, dispersal limitations, and climatic and habitat availability limit full fundamental niche occupation (Tingley et al. 2009; Soberón 2007).

Species with a high STI or SPI relative to other species are considered warm-adapted, and those with a lower STI are considered cold-adapted by comparison. Similarly, a high SPI value means a species is wet-adapted relative to other species considered, whereas a lower SPI value indicates a preference for drier conditions compared to other species.

Species STI and SPI indices were calculated as the average of mean annual temperature or mean annual precipitation, respectively, experienced by a species across their entire distributional range during the timeframe of the study, from 1987

to 2023 (Devictor, Julliard et al. 2008). Practitioners may choose to use only the study area or a regional zone to establish the STI or SPI of each species (Devictor, Chris van Swaay et al. 2012; Sparrius, Top and Swaay 2018; Curley et al. 2022), but we used the entire distributional range polygons from BirdLife Int (BirdLife 2022) to more accurately reflect their realized climatic niches (Soberón 2007) and to moderate against the effect of distributional change during the study period.

To match each unique species from the SABAP data to their polygon range from BirdLife International, we checked species taxonomy using the "taxize" v0.10.0 R package (Chamberlain and Szocs 2013). In total, 635 species were matched and assigned ranges. The global CRU-TS v4 (Harris, Osborn et al. 2020) data were used to calculate the mean temperature or mean precipitation over the geographic range of a given species, using one occurrence per grid cell, for each year. We then calculated a mean across all years of the study (1987 to 2023) to assign a single STI and SPI value to each of the 635 species. Each occurrence in the dataset was therefore associated with a species ID, latitude and longitude value, and STI and SPI value.

3.4.4 Community Weighted Indices

Community weighted means (CWMs) quantify the average climatic affinity of the species present in a given assemblage. For each grid cell and year, we calculated the CWM as the mean Species Temperature Index (STI) or Species Precipitation Index (SPI) of all species recorded in that cell:

$$\text{CWM} = \frac{1}{S}\sum_{i=1}^{S} x_i, \tag{3.4.4.1}$$

Equation 1. Where x_i is the STI or SPI value of species i, and S is the number of species recorded in that grid cell and year.

Higher CWM values therefore indicate assemblages dominated by species associated with warmer (for CTI) or wetter (for CPI) parts of their climatic niche, whereas lower values indicate assemblages dominated by cooler- or drier-adapted

species. Changes in CWMs through time reflect shifts in the relative prevalence of species with different climatic affinities, providing a functional indicator of community-level responses to climate change.

Community weighted means (CWMs) were derived from STIs and SPIs to calculate two functional indices of community composition: the community temperature index (CTI) (Devictor, Julliard et al. 2008) and the community precipitation index (CPI) (Curley et al. 2022). The CWMs were calculated per grid cell for each year, meaning that each grid cell was characterised by a time series of CTI and CPI values.

Previous studies have often focused on the use of CTI as an easily-calculated indicator of community temperature preference over time (Bowler and Böhning-Gaese 2017). CTI can be used to assess changes in the relative composition of high- (top quartile of STI scores) versus low- (bottom quartile of STI scores) temperature dwellers in local communities (Devictor, Julliard et al. 2008; Bowler and Böhning-Gaese 2017). CTI can be calculated with abundance data (Princé and Zuckerberg 2015; Curley et al. 2022; Santangeli and Lehikoinen 2017) or simple presence-absence information (Devictor, Julliard et al. 2008; Bowler and Böhning-Gaese 2017; Devictor, Chris van Swaay et al. 2012; Pierre Gaüzère, Jiguet and Devictor 2015).

Although CTI is a powerful metric of change on its own, its robustness and biological justification relies on the choice of continuous variable used to characterize species. A single mean climate variable may not reflect the primary climatic influence on species' ranges (Devictor, Julliard et al. 2008). A meta-analysis of species shifts have indicated that responses to climate change are often diverse, and do not always track those of temperature (Chen et al. 2011). For example, Curley et al. (2022) found no correlation between CTI and temperature, but a positive correlation between CPI and precipitation. We therefore consider CPI alongside CTI (Curley et al. 2022).

Unlike some previous studies (Curley et al. 2022; Lehikoinen et al. 2021; Devictor, Chris van Swaay et al. 2012), we do not have predefined study locations or repeat transect routes. Our repeat sampling locations were defined by the $0.5° \times 0.5°$ grid

cells. This means that our data do not represent 'communities' in the ecological sense, but rather assemblages of species that occupy similar regions.

Due to the nature of our data, we determined CWMs based on species' presence data, not abundance, and therefore retained only one occurrence per species per grid cell for each year of our CTI and CPI calculations. Previous studies have found no significant difference between using abundance or presence-absence data for CTI and CPI calculations (Devictor, Julliard et al. 2008; Devictor, Chris van Swaay et al. 2012; Bowler and Böhning-Gaese 2017; Pierre Gaüzère, Jiguet and Devictor 2015).

Since sampling effort varied across space and between our two study time intervals, we subsampled each grid cell to a set number of species within each year. To determine the minimum number of species to subsample, we tested values from 10 to 150 in increments of 10, with 100 replicates per size. We only performed these sensitivity analyses on grid cells with ≥ 150 unique species. Each cell was therefore characterized by 100 CTI and CPI values per sample size. The CTI and CPI metrics were then calculated for each subsampling iteration at every sample size by taking the mean of the sampled species' STI or SPI values. To assess the stability of results based on the number of species used, we took the standard deviation of CTI and CPI across sampling replicates and grid cells (see B.1). Based on this sensitivity analysis, we chose 60 unique species for calculation of CTI and CPI for a given grid cell as this balanced cell retention and outcome stability (see B.2, B.3, B.4).

3.4.5 SAR Modelling

Spatial data like ours often exhibit positive spatial autocorrelation, such that adjacent temperature or precipitation values will tend to have similar values (Sokal and Oden 1978). The presence of spatial autocorrelation is problematic for classic statistical tests, as these methods assume independently distributed errors (Legendre 1993) and can return more significant results than the data justify (Dale and Fortin 2002). SAR models use a spatial weights matrix to explicitly account for patterns in the response variable that are not predicted by explanatory variables, but instead

are related to values in neighbouring locations (Kissling and Carl 2008; Ver Hoef et al. 2018). They have been found to outperform ordinary least squares regression models and are expected to yield better predictions for mobile species such as birds, even in coarse grained models with large extent (Bahn, O'Connor and Krohn 2006).

We used spatial autoregressive (SAR) linear lag models to determine how CTI and CPI have changed over the study interval, and to test potential correlates of these changes. We conducted two separate analyses: (i) one that examined change over time in CTI/CPI for each year from 1987-1991 and from 2019-2023 (section 3.3.5), and (ii) one that examined potential correlates of change in CTI/CPI, using a linear SAR model of change in CTI and CPI between two aggregated time intervals (1990-1991 and 2019-2020) (section 3.3.5).

3.4.5.1 Quantifying changes in CTI and CPI over time

To determine whether communities have become more thermophilic or dry-species dominant, we analysed change in CTI and CPI over time. SABAP presence data were characterized for each year, for a total of ten separate years (from 1987 to 1991, and from 2019 to 2023). We generated a SpatRaster (terra v1.8-42) (Hijmans 2025) time series, wherein each layer is a year of CTI or CPI grid cell data. We enforced the ≥ 60 species per grid cell criterion for this year-by-year analysis, resulting in 398 grid cells for analysis. Bootstrapping was used to randomly sample a single occurrence of each unique species up to 60 species. The final CTI or CPI for a grid cell was calculated as the median of the 100 subsampled iterations. We used the median rather than the mean to reduce the effect of outliers on CWM metrics.

We calculated the slope of the CTI and CPI temporal data for each grid cell. The slope for a grid cell is the coefficient of the linear model: lm(C[T or P]I ~year). For both models, we generated the weights matrix from our grid-cell-defined dataset using a rook continuity (k=4) nearest neighbour approach, with weights normalised with row standardisation. This captures direct spatial influence of units sharing

an edge and avoids the potential for over-smoothing compared to queen continuity (k=8). We modelled the slope using only the intercept.

The output of the model indicates the average slope per grid cell for CTI and CPI, and whether there is a meaningful spatial pattern to change, such that nearby cells change at similar rates. Spatial autocorrelation was not significant for precipitation (Rho: 0.098; LR test value: 1.758, p-value: 0.185; Wald statistic: 1.921, p-value: 0.166), and thus we used a linear model that does not account for spatial relationships.

3.4.5.2 Predictors of ΔCTI and ΔCPI

To test for potential correlates of change in CTI and CPI, we aggregated the yearly data into two five-year time bins (1990-1991 and 2019-2020). We then selected 60 species for each grid cell and aggregated time bin to calculate CTI and CPI in a series of 100 bootstrap replicates. The final CTI or CPI for a grid cell was calculated as the median of the 100 subsampled iterations. We only performed analyses for the 359 grid cells shared across the two time intervals (1990-1991 and 2019-2020). The change in CTI and CTI was calculated by subtracting the CTI and CPI from SABAP1 to SABAP2 on a cell-by-cell basis (i.e., ΔCTI and ΔCPI).

We modelled ΔCTI and ΔCPI as a function of six anthropogenic and climatic predictor variables: the baseline CTI or CPI value (1990-1991), human modification index (Theobald et al. 2020), population density change (Liu, X. Cao et al. 2024), change in mean temperature/precipitation (Harris, Osborn et al. 2020), change in maximum temperature/precipitation (Harris, Osborn et al. 2020), and environmental heterogeneity.

Change in CTI and CPI were modelled as outcomes separately with two SAR models. Only mean and maximum temperature change were included for ΔCTI models, and only mean and maximum precipitation change were included for ΔCPI models.

Anthropogenic variables

We included two measures of human pressure: the human modification index (Theobald et al. 2020), which evaluates anthropogenic land use change, and population density (Liu, X. Cao et al. 2024).

Population density change was calculated from ~1 km resolution global data from Liu, X. Cao et al. (2024), resampled to 0.5°×0.5° grid cells using the bilinear resample function in "terra" v1.1-4 (Hijmans 2025). The difference between the mean 1990-1991 value and 2019-2020 value was taken for each grid cell.

Change in the human modification index was extracted from the $0.09km^2$ resolution data of Theobald et al. (2020). Rasters were resampled to the resolution of the CTI and CPI raster (0.5°×0.5°) using the "resample" function in "terra" v1.1-4 (Hijmans 2025). This data was only available for single years at five-year increments between 1990 and 2015. Therefore, the difference between 1990 and 2015 values was calculated for each grid cell.

Climatic variables

We considered four climate-based predictors for changes in CTI and CPI: (i) change in mean temperature or precipitation, (ii) change in maximum temperature or precipitation, (iii) baseline CTI or CPI, and (iv) environmental heterogeneity.

Temperature and precipitation trends were calculated from CRU-TS v4 data (Harris, Osborn et al. 2020) as the difference between mean or maximum values for the SABAP1 period (1990–1991) and the SABAP2 period (2019–2020).

Baseline CTI and CPI values were defined as the mean values within each grid cell during SABAP1. Baseline environmental conditions are known to shape both subsequent community change and modern biodiversity patterns (Antão et al. 2020; Outhwaite, McCann and Newbold 2022).

To estimate microclimatic heterogeneity derived from topography, we combined high-resolution (30 m) DEM data (OpenTopography 2013) with average solar-incidence layers following Montràs-Janer et al. (2024). This metric provides a static spatial proxy for fine-scale thermal and moisture variability rather than a measure

of temporal environmental change.

To test the hypothesis that greater change in the predictor variables will lead to greater change in the CWM metrics, we converted the delta predictor and CWM values to absolute values. By using an absolute difference, we converted the relationship to approximately linear, since we were primarily interested in the magnitude of change.

The "vegan" R package v2.7-0 (Oksanen et al. 2025) was used for predictor variable standardisation, wherein each variable was centred around a mean of 0 and standard deviation of 1. The variance inflation factor (VIF) for each predictor was calculated and any variable >10 is removed following recommendations of Dormann et al. (2013). None of our predictors were found to be autocorrelated.

We used SAR generalized linear model with a weights matrix generated from our grid-cell-defined dataset using a rook continuity (k=4) nearest neighbour approach, with weights normalised with row standardisation. The full models for ΔCTI and ΔCPI each include six predictors:

$$\begin{aligned} |\Delta\text{CTI}| \sim\ & \text{Baseline CTI} + \text{Environmental Heterogeneity} \\ & + \Delta\text{Mean Temperature} + \Delta\text{Max Temperature} \\ & + \Delta\text{Population Density} + \Delta\text{HMI} \end{aligned} \tag{3.4.5.1}$$

Equation 2. Linear model variables used to explain absolute change in CTI.

$$\begin{aligned} |\Delta\text{CPI}| \sim\ & \text{Baseline CPI} + \text{Environmental Heterogeneity} \\ & + \Delta\text{Mean Precipitation} + \Delta\text{Max Precipitation} \\ & + \Delta\text{Population Density} + \Delta\text{HMI} \end{aligned} \tag{3.4.5.2}$$

Equation 3. Linear model variables used to explain absolute change in CPI.

We then used the stepAIC() function in both directions to determine the combination of variables that minimises AIC using the "MASS" v7.3-65 (Venables

and Ripley 2002) R package. For both ΔCTI and ΔCPI models, the application of SAR spatial weighting removed significant spatial autocorrelation prior to modelling. Post-SAR correction, Moran's I for ΔCTI (~-0.004, p ~0.51) and ΔCPI (~0.007, p ~0.39) was small and not statistically significant.

3.4.6 Jackknife Species Contribution

We assessed how individual species influence the CWM relationships using jackknife analysis to identify species that have the greatest impact on assemblage structure (Davey et al. 2013; Curley et al. 2022). Analyses were conducted on an annual basis for each year from 1987-1991 and from 2019-2023, for a total of 10 years. For each year, we calculated baseline CTI and CPI values for all valid grid cells based on the full assemblage composition. We then removed each species in turn and recalculated the CTI and CPI values per cell. The percentage difference between the baseline and the recalculated values quantified each species' influence on CTI and CPI in that year.

To assess the robustness of each species' effect, we implemented a bootstrapping procedure. For species ranked among the top 10 most impactful on CTI or CPI in a given year, we resampled grid cells with replacement 1,000 times and recalculated CTI and CPI differences in each iteration. This allowed us to generate 95% confidence intervals around the average CTI and CPI effect per species. A species' contribution was considered statistically significant if its confidence interval did not include zero.

We used species presence data for our jackknife analysis. This means that a dominance in abundance will not be detected. This limitation is necessary due to the unequal sampling introduced by the procedural changes between SABAP1 and SABAP2.

3.5 Results

3.5.1 Community Weighted Indices

3.5.1.1 Temperature and precipitation change

Based on the CRU-TS climate data (Harris, Osborn et al. 2020), mean temperature across South Africa, Lesotho, and Eswatini (n=321 grid cells) increased by ~1.02°C from 1987 to 2023, while average precipitation declined by ~0.88 mm.

The mean ΔCTI across grid cells (n=321) was ~0.07°C, and the mean ΔCPI across the region was ~-0.41 mm, meaning that overall conditions became drier, supporting previous work (Simmons et al. 2009).

3.5.1.2 CTI and CPI changes

We examined how CTI and CPI changed over the 1987-2023 interval. We found moderate spatial autocorrelation of CTI trends (Rho value 0.234), confirmed by a highly significant likelihood ratio (LR test value 11.336) and Wald test (Wald statistic: 12.155; p-values $p < 0.001$). This autocorrelation was accounted for by the use of the SAR approach, with non-significant spatial autocorrelation of the residuals (Lagrange multiplier test statistic 0.072, p-value 0.788). The CTI SAR model suggests that, on average, CTI increased by ~$0.003° \pm 0.001$ per year between 1987 and 2023 across grid cells (p-value 0.016) (3.1).

Spatial autocorrelation for the CPI SAR model was small and non-significant (Rho 0.098), and likelihood ratio (LR test value 1.758, p-value 0.185) and Wald tests (LR test value 1.758, p-value 0.166) demonstrate that the spatial lag term does not improve the model compared to a non-spatial version. We therefore modelled CPI change using a standard linear model and found CPI declines by -0.028 ± 0.007 mm per year between 1987 and 2023 ($p < 0.001$) (3.1).

3.5.1.3 Correlates of ΔCTI and ΔCPI

The best model for change in CTI based on AIC included three predictors: change in mean temperature, change in population density, and environmental heterogeneity

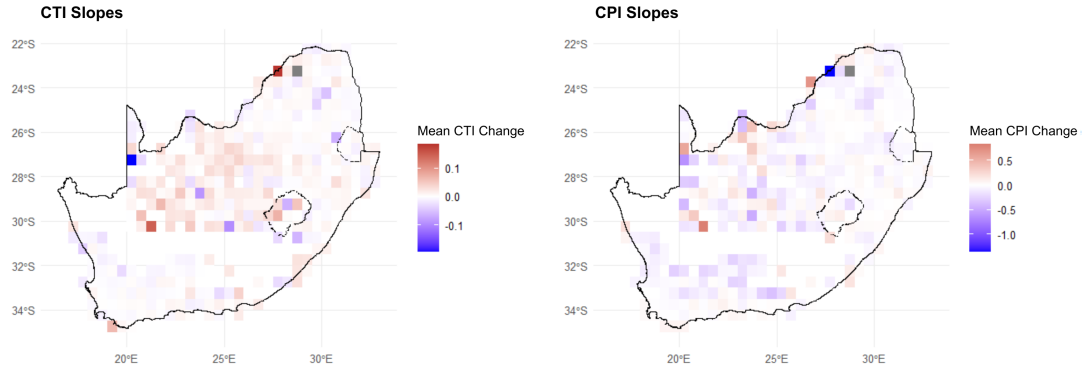


Figure 3.1: Changes in slope for CTI and CPI for each grid cell between 1987 and 2023. Positive slope values indicate a positive trend (warmer or wetter conditions), and negative slopes indicate a negative trend (colder or drier conditions).

(B.1; 3.2). Of these predictors, only environmental heterogeneity was statistically significant and negatively correlated with delta CTI (-0.03 ± 0.01).

The best model for change in CPI based on AIC included only one predictor: baseline CPI (B.2; 3.2). Baseline CPI was statistically significant and negatively correlated with delta CPI (-0.68 ± 0.09).

3.5.2 Jackknife Species Contribution

We calculated the mean effect of each species on CTI and CPI using only the years in which their effect was statistically significant, defined as years where the 95% confidence interval of their estimated contribution did not include zero.

A total of five species showed significant effects on CTI in at least one year, while three species showed significant effects on CPI. For CTI, species were significant for between five and ten years (mean = 7.8 years). *Passer domesticus* (House sparrow) and *Saxicola torquatus* (African stonechat) were each significant in all ten years, while *Motacilla aguimp* (African pied wagtail) was significant in eight (3.3).

In contrast, CPI effects were more limited in duration, with significant species showing effects in two to four years (mean = 3.0 years). *Elanus caeruleus* (Black-winged kite) was the most consistently significant species for CPI, followed by *Crithagra mozambica* (Yellow-fronted canary) and *Emberiza capensis* (Cape bunting). These results suggest that a small subset of species consistently influenced CTI

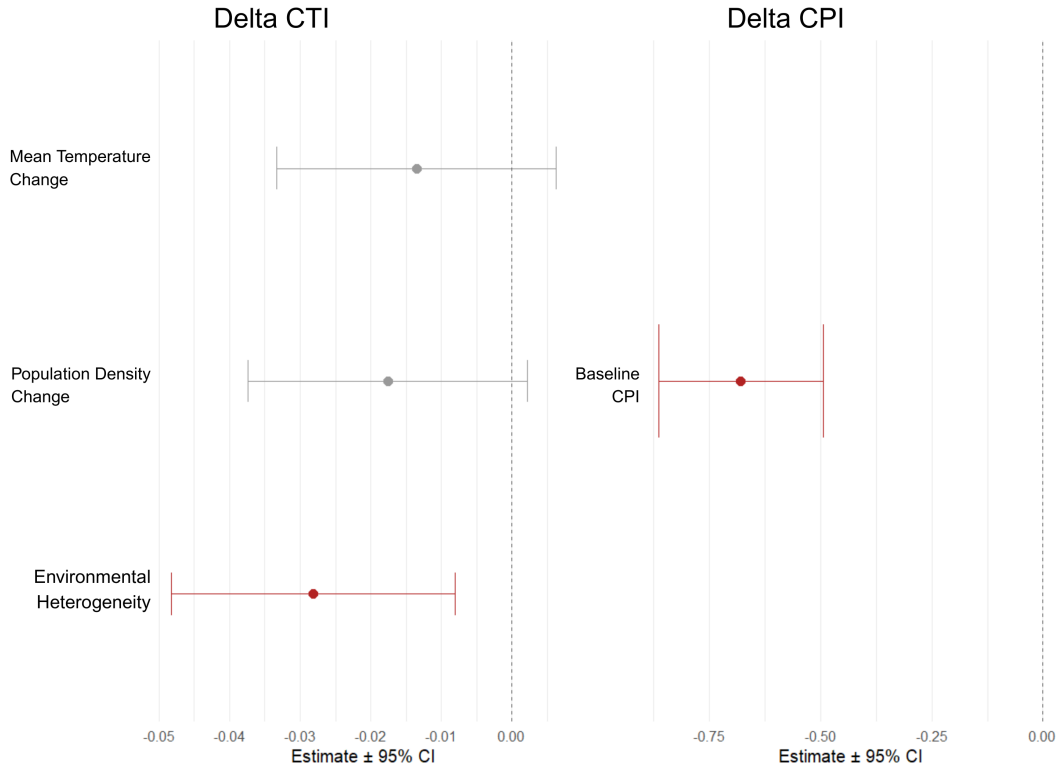


Figure 3.2: Model coefficients with standard errors determined from the ΔCTI and ΔCPI models. Red bars indicate statistical significance at alpha of 0.05.

across years, while CPI responses were driven by fewer species with more temporally restricted effects (3.3).

Figure 3.3 (3.3) shows heatmaps of CTI and CPI jackknife output, with the change in CTI or CPI percentage associated with each species per year displayed. These provide some support to overall positive CTI and negative CPI shifts. The CTI heatmaps show that the statistically significant species change across the period are only those contributing to increased CTI. CPI heatmaps show the opposite trend, significant species years only contribute to a reduction in CPI.

3.6 Discussion

The temperature and precipitation climate data used in this study support previously-observed shifts of increasing temperature and decreasing precipitation in southern

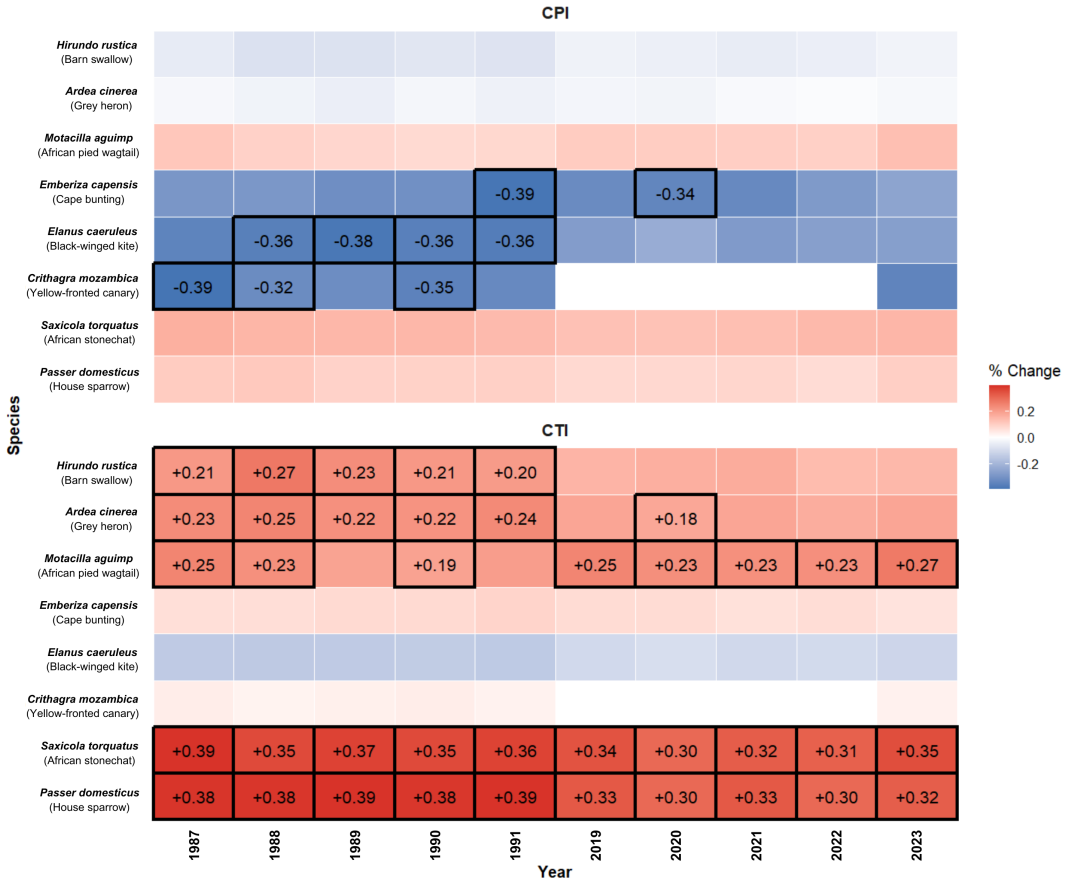


Figure 3.3: Species-level contributions to CTI and CPI across years. Each tile shows the percentage change in CTI or CPI when a species was removed from the assemblage in a given year. Values are displayed only for species whose effect was statistically significant (95% confidence interval did not include zero), indicated by a black outline. Red tones represent positive contributions (species associated with warmer or wetter assemblages), and blue tones represent negative contributions (species associated with cooler or drier assemblages).

Africa (Jury 2013; Wit and Stankiewicz 2006). Our study also supports previous findings that suggest the composition of bird communities globally appear to respond to climate changes by becoming increasingly warm- and dry-adapted (Santangeli, Rajasärkkä and Lehikoinen 2017; Oliver et al. 2017; Lehikoinen et al. 2021; Sparrius, Top and Swaay 2018; Lajeunesse and Fourcade 2023; Curley et al. 2022; Montràs-Janer et al. 2024). However, although mean temperature across southern Africa increased by ~1.02°C and mean precipitation decreased by ~-0.88 mm, ΔCTI only increased by ~0.07°C and ΔCPI decreased by ~-0.41 mm during the same time interval. This lag between ΔCTI and temperature has been found previously, and

suggests birds are not able to effectively track such changes (Devictor, Julliard et al. 2008; Lajeunesse and Fourcade 2023; Lindström et al. 2013).

Our results support microclimate heterogeneity as a regulator of CTI change (B.6). Higher microhabitat heterogeneity is associated with reduced change in ΔCTI. This result is congruent with previous research that found responses to macroclimate warming may be buffered by micro-climatic heterogeneity (De Frenne et al. 2013; Rull 2009; Maclean et al. 2017; Pincebourde and Suppo 2016; Guan et al. 2021; Marta et al. 2023), likely due to the provision of suitable microhabitat conditions for threatened species (Pearson 2006). Indeed, mosaics of microclimates and microhabitats may better allow species to find suitable conditions within relatively confined spatial regions as climate changes, possibly allowing organisms to withstand, at least temporarily, the effects of warming (Pierre Gaüzère, Princé and Devictor 2017; Marta et al. 2023).

Unlike previous studies, however, we do not find a correlation between change in CTI and anthropogenic influence (i.e. population density change and human modification index change) or temperature variables (i.e. change in mean temperature, and change in maximum temperature). Previously, Lajeunesse and Fourcade (2023) identified a significant interaction between temperature trends and human influence index, wherein CTI and temperature trends were weaker with high values of the human influence index. Other studies have also found that within protected areas, where human influence is lower, CTI values are lower than the surrounding zones (Hintsanen et al. 2023; Santangeli and Lehikoinen 2017) and that increased area of intensely managed land stymies CTI change exacerbating declines in cold-adapted bird species whilst also preventing increases in the proportion of warm-adapted birds (Oliver et al. 2017). The lack of correlation in our study could be due to the variables incorporated into our model and the coarse spatial resolution ($0.5°\times0.5°$) of the cell sampling approach. A lack of reliable land-cover and use data for our study period (1987-2023) precluded the analysis of anthropogenic influence through intensification or use alteration. Additionally, spatial resolution, determined by

available gridded CRU-TS climate data (Harris, Osborn et al. 2020) would not have allowed for accurate representation of all but the largest protected areas.

The absence of a CTI–temperature relationship contrasts with findings from Europe (Devictor, Julliard et al. 2008; Santangeli, Rajasärkkä and Lehikoinen 2017; Lehikoinen et al. 2021), North America(Lehikoinen et al. 2021), and New Guinea (Freeman and Class Freeman 2014), where avian CTI shifts closely track warming. However, these studies come from ecological contexts that differ markedly from southern Africa's climate and biogeography, which may partly explain the discrepancy. Several non-mutually exclusive mechanisms could account for the weak CTI–temperature association we observe. First, species' distributions in southern Africa are often constrained as much by precipitation and habitat structure as by temperature (Şekercioğlu, Primack and Wormworth 2012; Illán et al. 2014), meaning that thermal affinities may not be the primary axis along which assemblages reorganise. Second, strong microclimatic and topographic heterogeneity may buffer communities from regional temperature trends (Curley et al. 2022), limiting detectable thermophilisation signals at the scale of our grid cells. Third, assemblages in this region include a mix of resident and partial-migrant species whose responses to warming may be asynchronous or muted (Rushing et al. 2020; Both and Visser 2001).

Consistent with my results, Lajeunesse and Fourcade (2023) found that although thermophilisation was evident overall, local variations in CTI were only weakly correlated with local temperature trends. Their analysis was conducted at a relatively coarse spatial scale (200 km radius sliding windows). By adopting such a scale, one implicitly assumes that 'local' temperature change can be represented by the average trend across a broad area, which may be too coarse to reflect the fine-scale climatic variation that bird communities actually respond to.

Our analysis of changes in CPI suggest that only the baseline value of CPI (i.e., the mean value for SABAP1, 1990-1991) is inversely and significantly correlated with ΔCPI (B.5). The relationship between ΔCPI and baseline CPI has not been found previously. It supports the idea that existing environmental conditions can determine the magnitude of community change (Antão et al. 2020; Outhwaite,

McCann and Newbold 2022). In the case of precipitation, less change was found for assemblages that existed in drier conditions already. This is likely due to initially dry-adapted assemblages having a lower capacity for further downward change in CPI through the introduction of alternative species. If an assemblage already contains the most dry-adapted species in the migratable neighbourhood then further precipitation decreases will not drive species change, except through extirpation.

Unlike previous studies, however, we do not find evidence for a relationship between change in CPI and change in precipitation metrics (Curley et al. 2022). This could be driven by the coarse spatial scale of our analysis ($0.5° \times 0.5°$) and its divergence from more traditional individual repeat recording sites. Aggregation of precipitation and CPI indices at this scale may very well be too large to accurately represent the climate change that communities are responding to as it overlooks microclimatic variation within this area (Lajeunesse and Fourcade 2023). Additionally, previous studies utilising the CPI index are limited in number (Curley et al. 2022), it is therefore challenging to infer the driving factors at varied spatial and temporal scales across different areas of the global landscape.

Whilst this study capitalises on the citizen science work of the SABAP programmes (gbif.org 2018) to elucidate the direction and potential drivers of avian assemblage change in southern Africa, our results must be contextualised based on the nature of the data. SABAP data collection was not uniform temporally and spatially (Bonnevie 2011) and was not collected in standardized transects, as used in similar studies (Devictor, Julliard et al. 2008; Devictor, Chris van Swaay et al. 2012; Curley et al. 2022; Guerin et al. 2024). Although we attempt to correct for any potential sampling biases (see section 3.3), the data collection protocols may impact our results. However, any potential added noise from differing sample protocols would likely make it more difficult to find secular changes in CPI and CTI through time, such that the patterns returned here may represent conservative estimates of change.

Changes in protocol between SABAP1 and SABAP2 meant direct analysis of changes in abundance was not possible (Bonnevie 2011; Underhill 2016; Underhill

and Brooks 2016). Changes in abundance have been shown to be highly influential in determining community structure through increasing population numbers of warm-adapted birds (Princé and Zuckerberg 2015). However, we do show that changes to species occurrence can produce consistent, low magnitude (<0.5%) changes to CTI and CPI over this 36 year timeframe.

Given the nature of the SABAP data, we analysed bird assemblages at a relatively coarse spatial resolution (0.5°×0.5°). This coarse spatial resolution is not necessarily reflective of true ecological communities (Wiens 1989; Stroud et al. 2015; Morin 2009). Analyses at finer spatial resolutions may reveal stronger patterns of change, especially if tied more closely to localized climate change.

We found relatively small changes in CTI (~0.07°C) and CPI (~-0.41 mm) over time. This is likely due to the short temporal span of our study (36 years). The SABAP2 project is ongoing, and therefore future work using these data could elucidate longer-term secular changes, potentially at a higher spatial resolution.

The integration of citizen science outputs with indigenous local ecological knowledge may allow for more accurate understanding of assemblage changes over time. Ehlers Smith et al. (2021) found that people of the grassland communities of southern KwaZulu-Natal, South Africa, who were heavily reliant on natural resources attained a higher local ecological knowledge score, indicating a greater breadth of ecological knowledge. Community members confirmed the otherwise unknown presence of conservation concern species within the area. This highlights the limitations of using only citizen science databases for conservation management.

3.7 Conclusion

Previous studies have shown that highly vagile species such as birds adjust their geographic ranges to track climate change (Gregory et al. 2009; Tingley et al. 2009; Chen et al. 2011; Auer and King 2014; Rubenstein et al. 2023; Virkkala and Lehikoinen 2014; Peterson and Martínez-Meyer 2008; Rushing et al. 2020; La Sorte and Thompson 2007; Taheri, Naimi and Araújo 2016; Gillings, Balmer

and Fuller 2015; Hockey et al. 2011). We support this work, finding that bird community composition has become increasingly dry- and warm-adapted over a period of 36 years. Our analyses suggest that different factors may be responsible for regulating how assemblages respond to these changes across niche axes, such that changes in CTI are negatively correlated only with microhabitat heterogeneity and changes in CPI negatively correlated only with baseline CPI. Community dynamics may therefore be more nuanced than temperature metrics alone suggest (Herrando et al. 2019).

In southern Africa, climate change and anthropogenic encroachment are inextricably linked, occurring simultaneously and compounding pressures on biodiversity (Hockey et al. 2011). Understanding these intertwined drivers is essential for mitigating future species loss. Ensemble models have indicated that a majority of bird species (62%) will lose climatically suitable space in southern Africa before 2100 (Coetzee et al. 2009). To guide effective conservation and future protected area delineation, it is important to understand how assemblages are changing spatially across southern African and the potential drivers behind those changes.

The wide world is all about you: you can fence yourselves in, but you cannot forever fence it out.

–J.R.R. Tolkien

4

Spatial Determinants of Community Wellbeing Near National Parks in Madagascar

Contents

4.1 Author Contributions 62
4.2 Abstract 63
4.3 Introduction 64
4.4 Methods 66
4.4.1 Study Area 66
4.4.2 Household Poverty and Wellbeing Indicators 67
4.4.3 Quantifying Community Land Use 69
4.4.4 Landscape Variables 71
4.4.5 Relationships Between Poverty and Wellbeing Indicators 72
4.4.6 Predictors of Spatial Heterogeneity in Socio-economic Outcome Indicators 73
4.5 Results 74
4.5.1 National Scale Drivers of Heterogeneity in Poverty and Wellbeing 75
4.5.2 Regional Scale Drivers 78
4.6 Discussion 80
4.6.1 Correlations Among Poverty and Wellbeing Metrics . . 81
4.6.2 National and Local Scale: Reconciling Scale Mismatch . 82
4.6.3 Land-use Dynamics and Socio-ecological Patterns 83
4.6.4 Limitations and Future Directions 84
4.7 Conclusion 85
4.8 Ethics Statement and Funding 87

4.1 Author Contributions

Harri Ravenscroft developed the overarching research ideas and conceptual framework, with advisory input from Dr James Borrell. Data collection was conducted in situ by trained Malagasy professionals and overseen by Kew and its partner organisations. Data preparation was primarily undertaken by Harri Ravenscroft, with wellbeing indicator metrics defined by Kew partners. The weighted Voronoi tessellation protocol was developed and implemented by Harri Ravenscroft. All analytical procedures and statistical analyses were conducted by Harri Ravenscroft. The initial manuscript was drafted by Harri Ravenscroft, with Dr James Borrell contributing to revisions, interpretation of results, and final approval of the chapter.

4.2 Abstract

Effective protected areas must balance biodiversity conservation with the livelihoods and wellbeing of surrounding communities. Where conservation imposes high local opportunity costs, desired outcomes may be unsustainable and long-term success unachievable. In Madagascar, a global biodiversity hotspot, millions of rural households depend on ecosystems within and around protected areas for their livelihoods and wellbeing. However, the costs of conservation are not evenly distributed: some communities face far greater restrictions and opportunity costs than others, depending on their location and access to natural resources. Here we analyse spatial heterogeneity in multidimensional poverty around eight protected areas in Madagascar using a representative baseline household survey (n = 2,135). We assess the relationships among alternative wellbeing indicators and examine how ecological, anthropogenic, and land cover variables are associated with disparities in poverty around protected areas. To better reflect patterns of local resource use, we apply a novel land-use partitioning method using Voronoi tessellation.

We find that predictors of wellbeing vary markedly by scale. At the national level, higher ecological productivity and suitability for cash crops correlate with improved socio-economic outcomes. By contrast, at the regional scale, outcomes were shaped by land cover change, burned area coverage, and connectivity to major towns. Increased burned area coverage and shorter distances to towns or ports are more closely associated with lower wellbeing, whereas increased shrub/scrub coverage is associated with higher wellbeing.

Our findings highlight the disconnect between aggregate poverty indicators and the local scale spatial heterogeneity of poverty around protected areas. While aggregate indicators may guide policy at scale, they mask substantial local variation in wellbeing that shapes how communities experience conservation. By integrating spatially explicit poverty assessments into conservation planning, we can support the effective management of protected areas and better target sustainable development interventions to the surrounding communities most impacted.

4.3 Introduction

Effective protected areas must not only conserve biodiversity, but also avoid disadvantaging surrounding communities (Brockington and Wilkie 2015; Gardner, Nicoll, Mbohoahy et al. 2013; Green et al. 2018; Hanauer and Canavire-Bacarreza 2015). The impact of protected areas on local people can be highly heterogeneous and context-dependent (Forest Carbon Partnership Facility 2020; Rodríguez-Rodríguez et al. 2024; Toillier et al. 2011; Andriatsitohaina et al. 2020). Where protected areas result in significant local opportunity costs, positive conservation outcomes may be unsustainable in the long term (Gatiso et al. 2022). Conversely, protected areas can benefit communities through diversified alternative livelihoods (e.g. tourism), and access to or benefits from greater ecosystem service provisioning (Brockington and Wilkie 2015). Nevertheless, protected areas often deliver a greater benefit to a wider national or global community, while imposing more acute costs on those living closest (Cooke et al. 2022; Green et al. 2018). Recognizing the drivers of inequality in the distribution of costs and benefits arising from protected areas is essential to develop equitable interventions that ameliorate these inequalities, and deliver more effective protected area management strategies (Ferraro 2002; Naidoo and Iwamura 2007; Ferraro, Hanauer and Sims 2011).

Madagascar provides a globally important case study for examining the social impacts of conservation. It harbours exceptional biodiversity, much of it endemic and highly threatened (Gardner, Nicoll, Birkinshaw et al. 2018; Ralimanana et al. 2022), yet also ranks among the poorest countries globally, with over two-thirds of the population experiencing multidimensional poverty and widespread food insecurity (World Bank Group 2024; Global Hunger Index 2024). Protected areas have been expanded iteratively to safeguard biodiversity, but they are embedded in densely populated, nature-dependent landscapes. Communities around these areas often rely on forest and ecosystem access to meet basic needs, creating intense pressure on limited natural resources (Toillier et al. 2011; McConnell, Sweeney and Mulley 2004; Andriatsitohaina et al. 2020; Herrera et al. 2021). This creates a

cycle in which further exploitation of a diminishing resource is required, resulting in continued and in some cases accelerating degradation (Ralimanana et al. 2022).

While numerous studies have sought to assess the overall poverty impacts of protected areas (Andam et al. 2010; Ferraro, Hanauer and Sims 2011; Hanauer and Canavire-Bacarreza 2015; Mammides 2020), particularly in Madagascar (Ferraro 2002; Gardner, Nicoll, Mbohoahy et al. 2013), results remain inconsistent. It is widely accepted that local political, ecological and socioeconomic context can substantially impact conservation and poverty outcomes (Waylen et al. 2010), and that this can be further exacerbated by the wide diversity of poverty and conservation effectiveness indicators (Roe et al. 2014). To date, few studies resolve the fine-scale spatial patterns that underpin these dynamics. This limits our ability to target interventions or understand which contextual conditions shape conservation–development trade-offs across different landscapes.

Here, we contribute to this debate by reporting a large household survey dataset from eight protected area landscapes in Madagascar to assess spatial heterogeneity in poverty and wellbeing. The high spatial resolution and sampling depth of our dataset enable us to examine two primary questions. (1) To what extent are poverty and wellbeing indicators correlated? and (2) How do bioclimatic, topographic and socio-economic variables, including proximity to protected areas, influence the fine-scale spatial distribution of poverty metrics? Additionally, to improve spatial inference, we introduce a novel Voronoi-based approach to better capture land-use partitioning among communities,in order to quantify the influence of local land use variables on community poverty and wellbeing indicators.

By identifying the socio-ecological conditions associated with poorer or better-off communities near protected areas, our findings aim to inform more targeted, equitable, and effective conservation planning. Such evidence is needed not only for Madagascar but also for other biodiversity-rich, poverty-affected regions where protected areas are central to national and international environmental strategies.

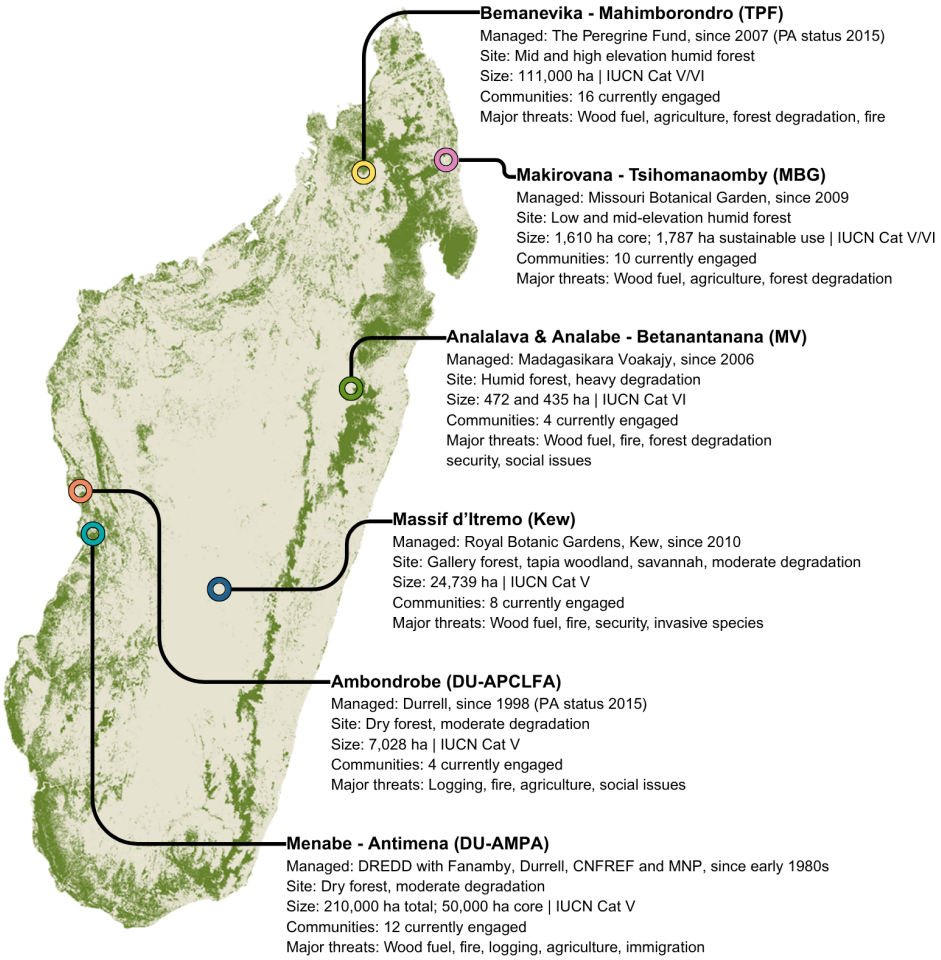


Figure 4.1: Sites used in this study. Fitantanana Maharitra Holovain-jafy (FMH) partner institutions and project sites. Category V protected areas are protected landscapes aimed at safeguarding the integrity of the developed interaction of people and nature. Category VI are protected areas with sustainable use of natural resources. Mainly in natural condition, with only low-level non-industrial natural resource use compatible with nature conservation (Stolton, Shadie and Dudley 2013).

4.4 Methods

4.4.1 Study Area

Our analysis covers eight protected areas across six regions of Madagascar, ranging from eastern rainforest sites, through to central highland grasslands and west coast dry forests (4.1).

These sites are managed by five national and international NGOs (4.1), and all

were established in 2015 or later. Together, these sites form part of a major six-year conservation investment project named "Fitantanana Maharitra Holovain-jafy" (FMH), or "Sustainable Management for Future Generations". The data outlined here forms part of a baseline biodiversity and wellbeing assessment, against which the impact of future interventions can be assessed.

4.4.2 Household Poverty and Wellbeing Indicators

A household survey was conducted as part of the FMH project baseline analysis in May-June 2023, covering 2,135 households across 93 project villages (around 33% of villages in the total project landscape). The household survey instrument was designed to primarily assess the socio-economic impacts of the project, with modules including the household schedule (demographic details of all household members, education levels, and income sources), living conditions (e.g., building materials, fuel and lighting, water and sanitation, savings and credit), health, food security, assets, agriculture, natural resource harvesting, quality of life, shocks and coping strategies, and natural resource management. From this survey we generated three principle socio-economic metrics of household poverty and wellbeing:

1. **Months of adequate household food provisioning (MAHFP):** A food security indicator based on the number of months in the past year that the household did not have enough food to meet their family's needs (Bilinsky and Swindale 2010). For the FMH baseline survey, access to "rice", Madagascar's predominant staple, was used instead of "food".

2. **Multidimensional Poverty Index (MPI):** The MPI is a composite poverty measure based on ten indicators across three dimensions—health, education, and living standards (UNDP-OPHI 2023). The FMH household survey implemented a modified version of this index to reflect locally available data. In the health dimension, the standard nutrition indicator was replaced with the MAHFP food-security measure. Several additional indicators (e.g.,

toilet type, drinking water access, and school attendance) were adjusted to match the FMH survey structure where exact age or technical information was not collected. For each household, deprivation scores across all indicators were calculated and aggregated following standard MPI weighting procedures. In the following analysis, we use a scaled derivative of this composite score, referred to as Household Poverty (HHP), which ranges from 0 to 1 (or 0–100%) and provides a comparable measure of multidimensional deprivation across households..

3. **Global Person Generated Index (GPGI):** Global Person Generated Index (GPGI): A subjective wellbeing measure in which respondents identify the life domains most important to them, rate their current status in each, and assign relative weights to those domains (Rasolofoson, Nielsen and Jones 2018). Each respondent selects five priority domains and rates each on a scale from 0 ("very bad as you fear") to 4 ("very good as you want"). They then allocate ten 'pebbles' across the five domains to indicate relative importance; these pebble allocations act as weights. The household GPGI score is calculated by multiplying each domain's rating by its corresponding pebble weight, summing the weighted scores across the five domains (maximum 40), and rescaling this total to a percentage out of 100. This produces a household-level wellbeing score that is comparable across households despite differing underlying domains.

Development of the survey questionnaire and modules was influenced by (a) the Building Resilience and Adaptation to Climate Change in Malawi household survey (questions on land and agriculture, assets, household demography and schedule, and food coping and security) (Duchoslav et al. 2020), (b) the Ecosystem Services for Poverty Alleviation P4GES household survey (Bidaud et al. 2016) (overall application of the survey to the Malagasy context, and questions on land use and ecosystem services provision, agriculture, social and human capital, and assets), (c) the Fourth Integrated Household Survey from the World Bank

Living Standards Measurement Survey (National Statistical Office (NSO) 2017) (questions on household demography, the household schedule, household shocks and coping strategies, and food security), (d) the CIFOR-PEN project household survey (PEN 2016) (questions on woodland use, ecosystem services, and household response to shocks), and e) the USAID Demographic and Health Surveys Model Household Questionnaire (*The DHS Program* n.d.) (household assets). The final survey instrument is included in the Supplementary Materials, with annotations based on feedback for future iterations of this survey.

The sampling design involved a two-stage clustered randomized approach; first selecting villages from the study area, and then randomly selecting households within the selected villages. While surveyed communities were stratified into control and treatment groups for future activities, here they will be considered as a single cohort. We sampled ~22 households per village. Overall sample size was constrained by budget availability, and future plans to detect outcome variables as a result of subsequent interventions in these landscapes. Surveys were administered using electronic tablets equipped with the Survey CTO application with enumerators fluent in both French and Malagasy and provided with multilingual technical field guidance documents.

4.4.3 Quantifying Community Land Use

To enable quantification of how local landscapes may influence community level poverty and wellbeing metrics, we required a zone of influence (ZOI) for each village within the study area. In the absence of having clearly defined boundaries of community usage, we inferred this zone using a four-step process (4.2). First we defined zones of interest around each protected area (PA) study location. The project villages, from which survey data was collected, were selected due to their proximity to the boundary of a study PA and the expectation of interaction with the natural resources present. At each of the six study locations, the average distance between original project villages was calculated and a convex hull extended

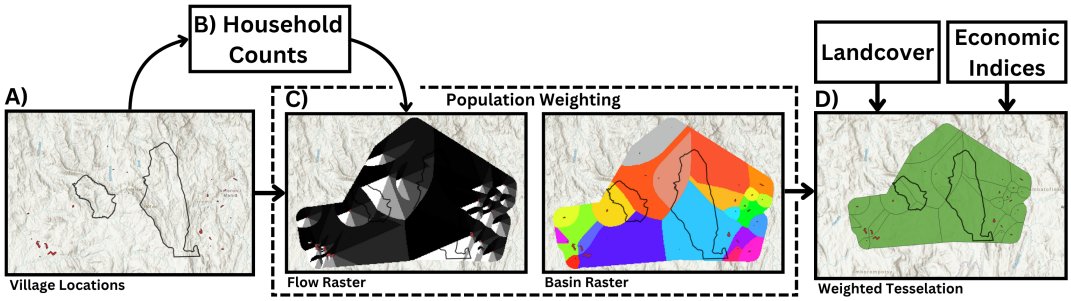


Figure 4.2: Construction of areas of community influence from voronoi tessellation mapping, allowing for assignment of economic and wellbeing indices to communities. **A)** Village locations are identified from survey records, national inventories, and satellite imagery. **B)** Following identification, each building within a settlement is counted, with the total number of buildings being assigned to a settlement acting as a proxy for population size. **C)** Flow and basin raster processes are used to assign areas of influence for each settlement, with population count acting as the weighting. **D)** The final weighted tessellation is reflective of relative settlement positions and size.

by this distance around the outermost village points. This method accounts for the varying distribution of villages and aids in visual identification of potential coordinate outliers, which produce zone shape irregularities.

Second, all settlements surrounding these protected areas were manually identified. We used near real-time ultra high resolution (0.5m×0.5m) satellite imagery to identify communities remotely (Pbc 2024). Identification and delimitation of each settlement within the ZOI is followed by point counts for each building within a community boundary. Building point counts, hereafter referred to as "counts", serve as a fair proxy for the population of a settlement (McConnell, Sweeney and Mulley 2004). Population for each settlement is log_{10} transformed to produce a weighting for Voronoi tessellation determination.

Third, weighted Voronoi polygons within each ZOI were generated in ArcGIS, building on previous applications (Dong 2008; Wilebore and Coomes 2016). Each village acts as a generator point for a Euclidean distance matrix in each ZOI, at 5m×5m resolution. Values within the distance matrices were divided by the weighting of the generator village to produce a weighted output raster for the PA-ZOI. Larger villages will have greater weighting values, therefore producing lower values within this layer.

Finally, all weighted rasters for a given PA-ZOI were merged and minimum matrix values (with each value relating to a single pixel) retained. Thereby giving large villages, with greater populations, more influence. Each pixel value is now equivalent to a pseudo-elevation determined by village weighting and distance to that village. Following Wilebore and Coomes (2016), the spatial analyst "Flow Direction" tool, using D8 flow type, was applied to the merged layer to compute changes in pseudo-elevation between each pixel. Following this, the "Basin" tool was used to delineate drainage basins by analysing the flow direction raster. It identifies pour points, locations at the edge of ZOI where runoff would occur. It then calculates the area that would contribute to each of these pour points. The output raster is then converted to discrete polygons representing village land areas within the PA-ZOI. These polygons were then used to spatially assign landscape variables to each village and to aggregate household-level poverty and wellbeing scores.

4.4.4 Landscape Variables

To investigate potential drivers of variation in poverty and wellbeing, we compiled a suite of 20 spatial variables (C.3). Variables were aggregated at the village level using the Voronoi-derived polygons and grouped into three thematic classes.

Ecological variables included mean annual temperature and precipitation, derived from CHELSA v2.1 (*CHELSA* 2020) at a 30 arc-second (~1 km) resolution, and topographic roughness, calculated from SRTM data collected in 2013 at 1 arc-second (~30 m) resolution (OpenTopography 2013). Agricultural potential was represented by crop suitability and length of the growing period from GAEZ v4 at ~9 km resolution (FAO and IIASA 2024).

Anthropogenic variables comprised population density (c.2020), sourced from WorldPop (2020); Tobler cost path, representing slope-adjusted walking travel time; and Voronoi polygon area, calculated using the spatial distribution and relative

weighting of settlements. Additional variables included travel time by walking to the nearest settlement, town (population > 3,000), and port.

Land cover variables were derived primarily from the Dynamic World v1 dataset (Brown et al. 2022), based on Sentinel-2 imagery at 10 m resolution. Land cover as of 2023 included tree, grass, crops, and shrub/scrub cover. Land cover change between 2016 and 2023 was measured as changes in these four categories. A separate disturbance variable captured the percentage of burned area within each settlement polygon. Burned area was derived using a two-stage detection algorithm at 20 m resolution from Sentinel-2 imagery (Fernández-García, Franquesa and Kull 2024).

All variables were standardised to 1 km resolution and processing was accomplished using the "terra" package v1.8-56 (Hijmans 2025). Prior to modelling we computed the covariance between predictor variables using a Spearman rank correlation coefficient, as not all variables were normally distributed. Any variables with high values (> 0.7) were removed, however none were found to be autocorrelated to this degree. Following this the "vegan" package v2.7-0 (Oksanen et al. 2025) was used for predictor variable standardisation (mean = 0, s.d. = 1).

4.4.5 Relationships Between Poverty and Wellbeing Indicators

Different socio-economic indicators are designed to capture different information about wealth, wellbeing and satisfaction. We first characterised the pairwise relationships between our three socio-economic outcome indicators (HHP, GPGI, and MAHFP) using linear models, from the "stats" package v4.6.0 (R Core Team 2025). The coefficient of determination (R-squared) is then derived as a measure of fit strength between each pairing, in addition to p-value and 95% confidence interval.

4.4.6 Predictors of Spatial Heterogeneity in Socio-economic Outcome Indicators

We used redundancy analysis (RDA) models to evaluate the influence of a combined suite of ecological, anthropogenic, and land cover predictor variables on poverty and wellbeing response variables at national and regional scales. We therefore conducted two separate analyses: one that implemented a full RDA using all twenty available predictors to determine potential correlates of poverty and wellbeing inter-site variance; and a partial RDA to better isolate local drivers by conditioning the analysis on our set of five ecological predictors that naturally vary across the landscape. In both cases, we sought to determine correlates at a broad scale for the combined socio-economic data, and for each component poverty and wellbeing metric in turn (HHP, GPGI, and MAHFP).

The full RDA model utilised all twenty available predictors and the combined outcome suite of HHP, GPGI, and MAHFP. Following this, each socio-economic outcome was modelled individually using an identical approach. This was undertaken so that nuances between these measures could be evaluated and explored.

We then conducted partial RDA modelling. For this, the five ecological variables that generally had low variability within sites, but high variability between sites (mean annual temperature and precipitation, topographic roughness, crop suitability, and length of growing period) were removed from the predictor set and included as conditional variables. As previously, the suite of three socio-economic variables were modelled together and separately.

The direction and significance of variable effects from each model outcome was assessed by extracting the contribution (biplot scores) of each predictor variable to the first RDA axis and merging these with their significance values (p-values) from permutation tests.

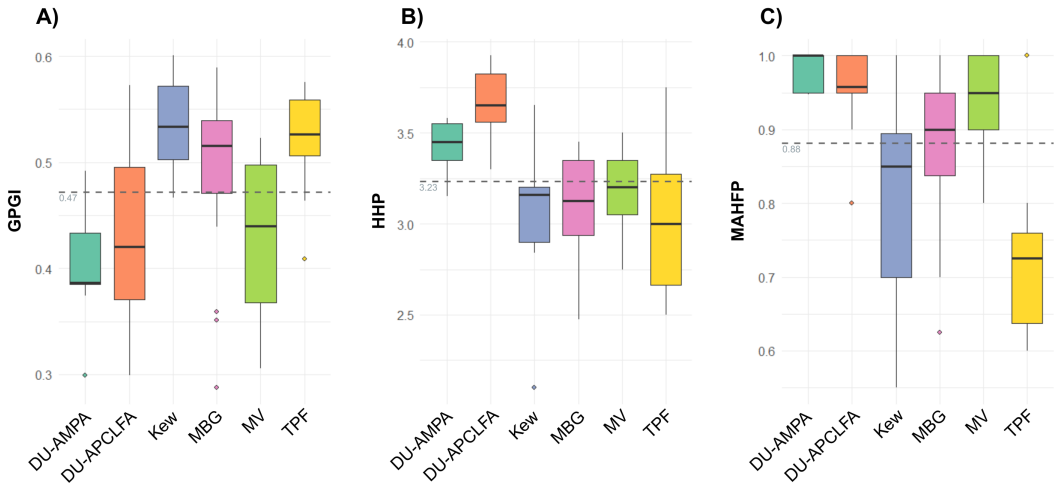


Figure 4.3: Livelihood indicators by protected area. **A)** GPGI, values < 0.5 are indicative of low wellbeing. **B)** HHP, mean values of 3.0 indicate "poor" overall poverty state, 4.0 indicates "severe" state of poverty. **C)** Mean MAHFP values can vary between 0.0 and 1.0. A mean value of 1.0 indicates complete food security.

4.5 Results

We analysed data from 2,135 households across 93 villages in six protected area landscapes. Approximately 90% of surveyed households across the PA network were classified as either 'poor' or 'severely poor' when assessed using the Multidimensional Poverty Index (4.3). The highest level of Household poverty (HHP) was observed at Ambondrobe PA in Western Madagascar. The average MPI across all regions was estimated as 41.9%, with the most widespread deprivations contributing to MPI scores included reliance on solid fuel (e.g., charcoal), inadequate sanitation, poor housing quality, and food insecurity. Notably, 89% of households experienced nutritional deprivation, which peaked during the hot-wet season.

4.5.0.1 Correlation between poverty metrics

Pairwise regressions indicated positive associations between all three core indicators (4.4). The highest correlation was found between household poverty (HHP) and months of adequate household food provisioning (MAHFP), with MAHFP explaining ~50% of the variation in household poverty ($R^2 = 0.501$, $p < 0.001$). This relationship

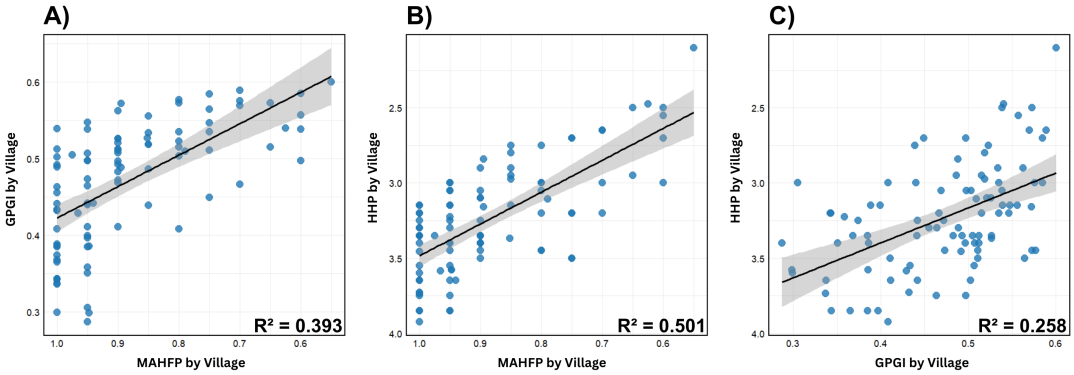


Figure 4.4: Pairwise plots of the three key economic indicators. Positive relationships exist between all three combinations of indicator variables. **GPGI** - Global Person Generated Index, **MAHFP** - Months of Adequate Household Food Provisioning, **HHP** - Household Poverty index.

was strongly positive, with a slope of 2.12 (95% CI: 1.68 to 2.56), indicating that reduced food provisioning was associated with higher levels of poverty. The association between the Global Person Generated Index (GPGI) and MAHFP was also significant ($R^2 = 0.393$, $p < 0.001$), though negative in direction (slope = –0.41, 95% CI: –0.52 to –0.30), suggesting that lower food security was linked with lower self-reported wellbeing. Lastly, household poverty was negatively associated with GPGI ($R^2 = 0.258$, $p < 0.001$), with a slope of –2.33 (95% CI: –3.15 to –1.51), implying that areas with higher wellbeing scores tended to experience lower poverty levels.

However, while metrics were correlated, they captured distinct aspects of household economic experience and wellbeing.

4.5.1 National Scale Drivers of Heterogeneity in Poverty and Wellbeing

The full RDA model assessed inter-site variation in wellbeing outcomes across all regions. The joined suite of economic indices (HHP, GPGI, MAHFP) identified eight statistically significant predictors ($p < 0.05$; 4.5A). The two strongest contributors were length of growing period ($F = 23.89$, $p = 0.001$; variance explained = 0.429) and mean temperature ($F = 21.32$, $p = 0.001$; variance = 0.382), indicating that

regions with longer growing seasons and cooler climates tend to report better economic and wellbeing outcomes. Other significant predictors of negative outcomes included grassland change (F = 5.73, p = 0.008; variance = 0.103), the extent of burned area (F = 3.34, p = 0.040; variance = 0.060), population density (F = 9.07, p = 0.001; var = 0.163), and village area of influence (F = 3.70, p = 0.035; var = 0.066). An increase to the land cover of burned area or grassland correlates with worse socio-economic outcomes in this case. Furthermore, a higher population density and larger village area of influence correlate with overall worse outcomes. Positive contributors also included walking travel time to the nearest town (F = 5.97, p = 0.009; var = 0.107) and port (F = 3.86, p = 0.037; var = 0.069). These suggest that greater travel times to these locations from survey villages are indicative of better socio-economic outcomes. Among the significant terms, the two strongest contributors were ecological, two related to land cover change, and three were directly anthropogenic.

Individual RDA models for each outcome largely reflected patterns observed in the full model. For household poverty (HHP) (4.5B), four predictors remained significant and matched in effect direction: length of growing period (F = 59.88, p = 0.001; var = 0.301), temperature mean (F = 31.12, p = 0.001; var = 0.157), grass cover change (F = 5.98, p = 0.017; var = 0.030), and burned area (F = 5.89, p = 0.016; var = 0.030). Unlike the full model, population density and travel times to the closest port or town were not significant, suggesting outcome-specific sensitivity. Additionally, higher crop suitability was associated with lower material poverty, a result not seen in other socio-economic outcomes.

For the Global Person Generated Index (GPGI) (4.5C), the predictors with highest positive effect strength were travel time to closest towns (F = 8.42, p = 0.010; var = 0.062) and length of growing period (F = 9.21, p = 0.008; var = 0.068). This correlate again suggests that increased growing period and a longer travel time to the nearest town are related to positive outcomes. A higher mean temperature (F = 10.42, p = 0.001; var = 0.077), greater population density (F = 5.95, p =

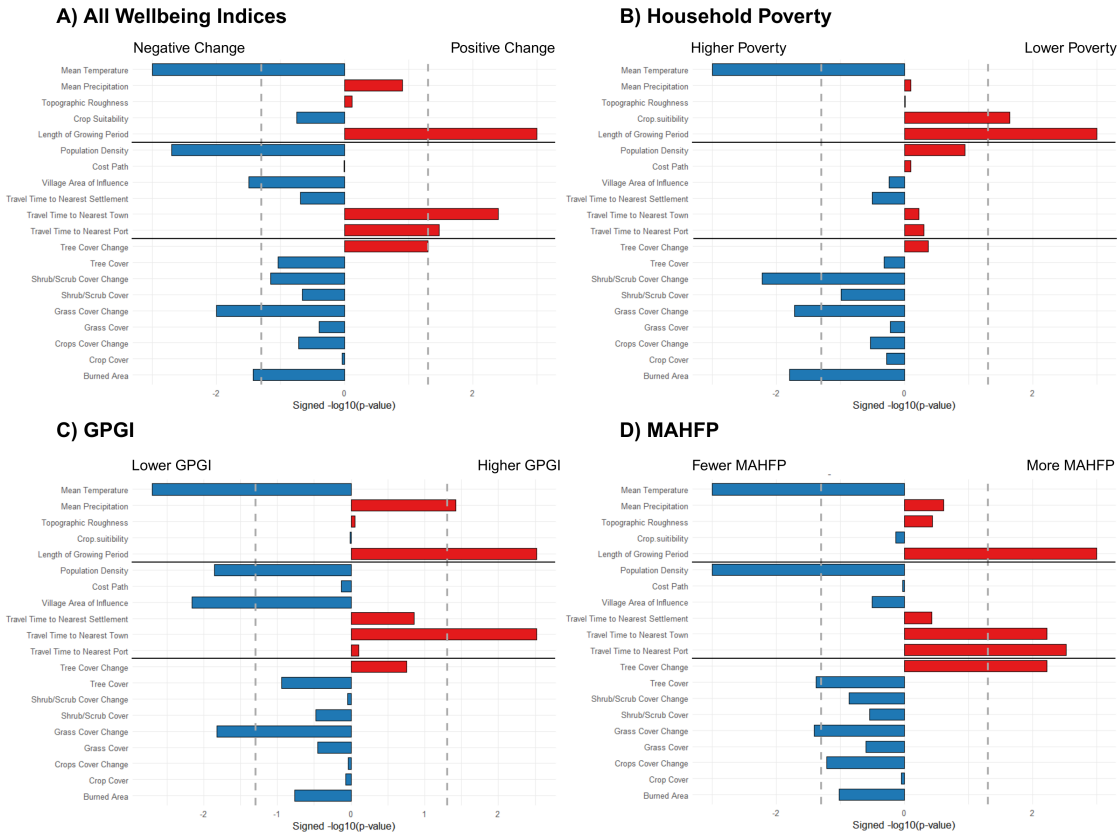


Figure 4.5: RDA outcomes for all input factors. Blue bar colouration indicates negative effect size of the variable, red indicates a positive effect size. Vertical dashed line demonstrates statistical significance of effect size. **A)** All economic indices considered together as full RDA outcome. **B)** Variable effect sizes for household poverty. **C)** Variable effect sizes for GPGI. **D)** Variable effect sizes for MAHFP.

0.011; var = 0.044), and wider village area of influence (F = 8.05, p = 0.008; var = 0.059) were again negatively correlated with better socio-economic outcomes.

The model for MAHFP showed strong similarities compared to the previous models (4.5D). Here, high mean temperature (F = 26.95, p = 0.001; var = 0.149) and greater population density (F = 19.23, p = 0.001; var = 0.106) again had large negative effect sizes. Greater travel time to ports (F = 12.05, p = 0.002; var = 0.067) and to towns (F = 7.87, p = 0.005; var = 0.044), alongside longer growing periods were again positively correlated to positive outcomes. Grassland and tree cover change were also significant ($p < 0.05$), further reinforcing the influence of land cover dynamics on food security.

4.5.2 Regional Scale Drivers

The partial RDA is designed to better determine intra-site at the protected area landscape scale variance through the use of ecological predictors as conditional variables. By conditioning on ecological variables, these models isolate the independent contributions of land cover, accessibility, and anthropogenic variables to wellbeing conditions.

In the combined model incorporating all outcome indices (HHP, GPGI, and MAHFP), two predictors were statistically significant (4.6A). Firstly, shrub/scrub cover had the strongest effect ($F = 9.66$, $p = 0.001$; variance = 0.173), with lower cover being associated with poorer wellbeing outcomes. Secondly, the village area of influence, delimited by Voronoi polygons, was also significant ($F = 3.70$, $p = 0.024$; var = 0.066), indicating that smaller catchment areas were associated with improved wellbeing conditions.

When individual outcome indices were analysed, distinct patterns emerged (4.6B-D). Across all three models, increased burned area (2014–2021) was consistently associated with significantly worse outcomes:

- HHP: $F = 22.04$, $p < 0.01$; var = 0.169
- GPGI: $F = 6.85$, $p < 0.01$; var = 0.058
- MAHFP: $F = 8.39$, $p < 0.01$; var = 0.062

Likewise, shorter walking travel time to towns (>3000 population) had a consistent negative effect:

- HHP: $F = 4.93$, $p < 0.01$; var = 0.038
- GPGI: $F = 13.86$, $p < 0.01$; var = 0.116
- MAHFP: $F = 12.36$, $p < 0.01$; var = 0.092

Beyond these shared influences, each outcome exhibited distinct relationships with landscape and infrastructure variables:

For HHP (4.6B), additional significant predictors associated with higher poverty included loss of grass cover (F = 3.93, p = 0.045; var = 0.030), loss of shrub/scrub (F = 5.04, p = 0.025; var = 0.039), shorter travel time to settlements (F = 6.84, p = 0.014; var = 0.052), and longer travel time to closest town (F = 4.93, p = 0.032; var = 0.038). These results emphasize how accessibility and land cover disturbance influence poverty levels independently of ecological gradients.

For GPGI (4.6C), in addition to shorter travel time to towns, lower GPGI is associated with lower burned area (F = 6.85, p = 0.011; var = 0.058), loss of grass cover (F = 5.49, p = 0.018; var = 0.046) and smaller village area of influence (F = 8.54, p = 0.002; var = 0.072) were significant. This suggests that positive wellbeing perceptions are shaped by access to greater land area and management of landscape vegetation.

For MAHFP (4.6D), there were fewer months of adequate food provisioning associated with higher population density (F = 7.18, p = 0.006; var = 0.053), and greater burned area (F = 8.39, p = 0.007; var = 0.062). Increased provisioning was associated with greater travel time to ports (F = 8.26, p = 0.004; var = 0.061), the travel time to the closest town (F = 12.36, p = 0.002; var = 0.092), and tree cover change (F = 5.74, p = 0.013; var = 0.043).

Overall, while burned area and travel time to towns emerged as consistent predictors across all models, other variables showed outcome-specific relevance. Village area of influence, land cover change, and accessibility to smaller settlements and ports each contributed to distinct aspects of economic wellbeing. These patterns reinforce the value of partial RDA in identifying the independent socio-spatial drivers of livelihood variation beyond ecological constraints.

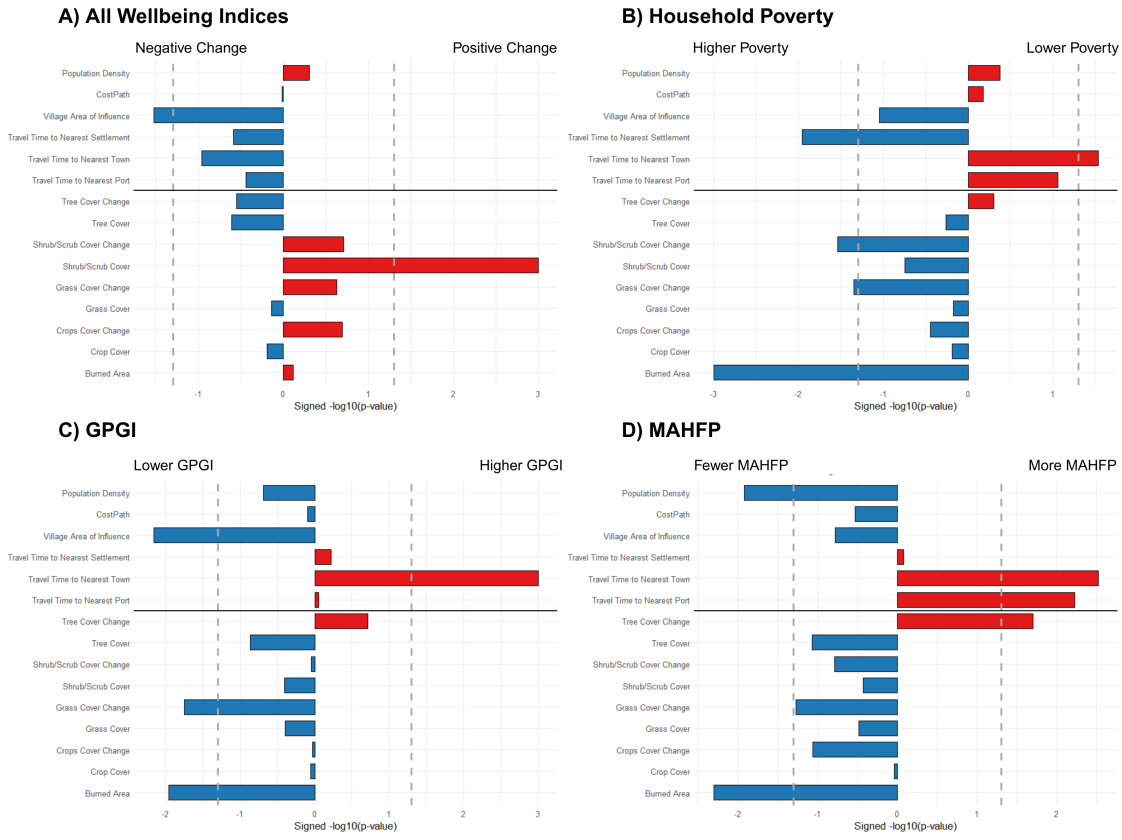


Figure 4.6: Partial RDA outcomes conditioned upon "Ecological" factors. Blue bar colouration indicates negative effect size of the variable, red indicates a positive effect size. Vertical dashed line demonstrates statistical significance of effect size. **A)** All economic indices considered together as partial RDA outcome. **B)** Variable effect sizes for household poverty. **C)** Variable effect sizes for GPGI. **D)** Variable effect sizes for MAHFP.

4.6 Discussion

This study provides a large multi-region assessment of multidimensional poverty and wellbeing around protected areas in Madagascar, a global biodiversity hotspot characterized by severe poverty and intense natural resource dependency. By integrating household-level survey data with ecological, anthropogenic, and land-use variables, we reveal substantial spatial variation in household wellbeing. At broader scales, ecological productivity, specifically temperature and growing-season length, emerges as a dominant determinant of wellbeing, whereas at local scales, wellbeing is more strongly associated with landscape dynamics, including burned area and accessibility to towns and settlements, rather than the cost of proximity to protected areas.

Our findings help inform widely-held assumptions that distance to a protected area influences socio-economic outcomes (Andam et al. 2010; Brockington and Wilkie 2015). Instead, poverty and wellbeing are influenced by a diverse array of ecological, land cover, and anthropogenic factors that vary markedly at different spatial scales. These findings underscore the need to move beyond simplistic distance-based assessment to more nuanced spatial analysis that links ecological productivity, land use patterns and resource use access to conservation planning.

4.6.1 Correlations Among Poverty and Wellbeing Metrics

A priority in conservation-development planning is to reliably measure household wellbeing to inform targeted interventions. Understanding how alternative metrics correlate and diverge is essential for capturing the multidimensional nature of poverty. In our analysis, economic poverty (Household Poverty, HHP) strongly correlated with food insecurity (Months of Adequate Household Food Provisioning, MAHFP), highlighting nutritional adequacy as central to broader economic security. Conversely, subjective wellbeing (Global Person Generated Index, GPGI) correlated less strongly, indicating it reflects distinct aspects of lived experience beyond economic and nutritional measures.

These findings emphasize the necessity of employing complementary wellbeing indicators in socio-ecological assessments. Reliance on single metrics risks obscuring important dimensions of poverty, potentially undermining conservation interventions. Integrating diverse wellbeing indicators thus offers a more comprehensive representation of community vulnerability and resilience, enabling more nuanced and effective conservation planning (McCrea, Walton and Leonard 2014; Johansen, Horney and Tien 2017; Zautra, Hall and Murray 2008; Brousmiche et al. 2021).

4.6.2 National and Local Scale: Reconciling Scale Mismatch

Conservation policy requires having an understanding of how drivers of poverty and wellbeing vary with spatial scale. To do this, it commonly relies on national-scale indicators, such as ecological productivity, climatic suitability, and aggregate socio-economic indices, to allocate resources and design interventions. At this broader scale, we found ecological variables, specifically temperature and length of growing season, were strong and consistent predictors of community wellbeing. Regions with cooler climates and longer growing seasons showed significantly improved outcomes, reinforcing established links between agroecological suitability and livelihood security (Zenda and Rudolph 2024; Tomich et al. 2011; Bezner Kerr et al. 2021). This ecological dependence highlights important constraints that must be accounted for in national-level conservation and development strategies.

However, at finer spatial scales, the level at which conservation interventions directly interface with communities, our findings diverged substantially. Here, local land-use variables such as higher burned cover, shrub cover change, and walking travel time emerged as key determinants of household poverty and food security. Notably, communities closer to protected area boundaries did not experience consistently different economic outcomes compared to more distant communities. This challenges the common assumption within conservation planning that proximity to protected areas intrinsically shapes local wellbeing. Instead, our results clearly demonstrate that local landscape dynamics (e.g., grassland vs. dry forest vs. moist forest) and infrastructural access are far more influential predictors of poverty and wellbeing outcomes than protected-area proximity alone.

This marked divergence between national and local drivers highlights an important scale mismatch often overlooked in conservation and development planning. At national scales, aggregated ecological and climatic indicators provide valuable insights for regional prioritisation. Whereas, at local scales, finer-grained dynamics, land cover changes, resource degradation, and access to markets and infrastructure,

dominate the lived realities of communities. Without explicitly recognising and reconciling these differing scales, conservation efforts risk misallocating resources, exacerbating inequalities, and undermining long-term sustainability and local legitimacy.

Addressing this mismatch requires a paradigm shift in conservation planning toward integrated, spatially explicit strategies that consider both ecological constraints and local socio-economic conditions. Policies and interventions must move beyond simplistic protected-area-centric approaches, prioritising instead targeted support for landscape management practices and sustainable land-use transitions. Ultimately, effective conservation planning demands simultaneous consideration of multiple scales to achieve meaningful, equitable outcomes.

4.6.3 Land-use Dynamics and Socio-ecological Patterns

Our analysis shows that land-use dynamics, particularly the extent of burned area, are consistently correlated with poorer household wellbeing across multiple indicators. However, this relationship should be interpreted cautiously. While fire may be associated with land degradation in some contexts (e.g., tavy - shifting cultivation), it can also reflect livelihood strategies such as pasture renewal. At the spatial resolution of our study, we cannot disentangle the underlying causes or intentions behind fire use. Instead, burned area should be seen as a proxy for broader landscape processes, both ecological and socio-economic, that vary across regions.

Other land-cover changes, including shifts in grass and shrub cover, were also associated with differences in wellbeing, though inconsistently across outcome metrics. Notably, improved wellbeing was consistently linked to longer travel times to towns and smaller village areas of influence, a reversal of what is commonly observed in developing countries (Burger et al. 2020; C. Wang et al. 2020; Lolle 2023). These findings highlight how landscape connectivity and land partitioning shape opportunity structures, often more so than proximity to protected areas. Tree cover change, while often interpreted as a negative ecological trend, showed

mixed associations, possibly reflecting trade-offs between resource access for NTF and forest conversion for agriculture.

These findings suggest that landscape structure and accessibility, rather than proximity to protected areas per se, are key correlates of wellbeing. The counter-intuitive pattern whereby communities closer to towns or ports experience poorer outcomes may reflect several interacting mechanisms: increased connectivity can intensify competition for land, facilitate greater resource extraction, and expose rural households to volatile markets, all of which may disproportionately disadvantage poorer communities. From a conservation perspective, this implies that pressures on biodiversity and protected area integrity may be greatest not in remote regions but in more accessible landscapes where demand for land, forest products, and agricultural expansion is concentrated. However, the mechanisms linking land-use change to poverty remain complex. Understanding these dynamics will require finer-resolution data on fire regimes, forest use, and land management practices, ideally combined with qualitative or longitudinal studies to unpack local motivations and outcomes. For now, these patterns offer important signals for where socio-ecological trade-offs may be most pronounced, and where more targeted investigation is warranted.

4.6.4 Limitations and Future Directions

While this study reveals clear spatial associations between environmental conditions, access, and wellbeing, its cross-sectional design limits causal inference. Longitudinal data are required to more accurately assess how changes in land use, infrastructure, and ecological factors drive or mitigate poverty over time. Moreover, although our household survey was spatially extensive and captured multiple dimensions of wellbeing, potential limitations include sampling biases and the subjectivity inherent in self-reported indicators.

Responses from only 20 households were recorded from each survey village, and many of the villages were those where interventions had been implemented prior to the baseline survey. Although inter-community differences appear to have been

captured well, intra-community scores suggest a lack of variability. This minimised our ability to identify drivers or attributes at a household level.

Future work should integrate longitudinal, spatially explicit surveys with participatory and qualitative approaches to triangulate patterns and validate inferred trade-offs. Incorporating ecosystem service valuation and land-use history would further strengthen understanding of how biophysical changes translate into livelihood impacts. Such approaches would provide a more complete evidence base to inform context-sensitive and equitable conservation strategies.

Although efforts were made to utilise high resolution data derived from Madagascan data collection, this was not possible in almost all cases. Instead, predictor variable data was extracted from global layers at coarser resolution. In addition to potentially overlooking fine-scale dynamics, the data lacked ground truth points in our focal country. This problem is widely seen across the global south (Ortiz-Bobea 2021) and has the potential to curtail fine-scale analysis of this type across broad areas of the globe.

A weighted Voronoi tessellation methodology was used to define the ZOI for each community within the study. Weighting was achieved through the counting of structures in high-resolution satellite imagery (Pbc 2024). However, building counts may be artificially inflated through the inclusion of small "farm houses"/cabins, often called lasy in Malagasy. These lasy are built close to farmers' fields and will be occupied during harvest periods if their fields are too far away from their permanent home but they can be challenging to differentiate as they can look very similar to a typical house in satellite imagery. Future work should determine the commonality of lasy structures throughout Madagascar and whether there is a relationship between their use and agriculture suitability or landscape fragmentation.

4.7 Conclusion

This study sought to assess the spatial drivers of heterogeneity in poverty and wellbeing indicators across communities surrounding Madagascar's National Parks.

Our three outcome metrics—HHP, GPGI, and MAHFP—were highly correlated (4.4), supporting their use as complementary proxies for household economic status. Using a nationally tailored RDA approach and a spatially realistic land assignment protocol, we found that:

- At the national scale, ecological gradients—particularly temperature and growing season length—were the dominant drivers of variation in economic outcomes.
- At the regional scale, independent of ecology, outcomes were shaped by land cover change, burned area coverage, and connectivity to major towns.

Taken together, these results suggest that economic outcomes in PA-adjacent communities are more strongly influenced by ecological setting and spatial infrastructure than by proximity to protected areas per se. Policymaking and conservation planning should therefore focus not only on spatial zoning but also on improving rural connectivity, supporting sustainable forest use, and accounting for the socioeconomic diversity of communities across Madagascar's varied landscapes.

To translate these insights into policy, we outline the following recommendations:

1. **Address the drivers of negative outcomes associated with connectivity:** Investigate why communities closer to major towns exhibit poorer economic outcomes (e.g., market saturation, resource competition, land-use displacement) and design targeted support to mitigate these effects. Connectivity investments should be coupled with safeguards that prevent rural households from being disadvantaged.
2. **Tailor development interventions to ecological gradients:** Recognise that temperature regimes, growing season length, and agroecological constraints fundamentally shape livelihood opportunities. Poverty-alleviation strategies should therefore be differentiated by ecological zone rather than applied uniformly.

3. **Prioritise sustainable land-use and fire management in high-pressure regions:** Where land cover change or burned-area prevalence is strongly associated with lower wellbeing, invest in community-based fire management, sustainable forest product value chains, and land-restoration initiatives.

4. **Embed multidimensional wellbeing metrics into conservation decision-making:** Adopt locally appropriate indicators such as the modified MPI (HHP), GPGI, and MAHFP for routinely assessing the social impacts of conservation policy and for identifying communities most at risk.

Together, these recommendations emphasise that effective conservation and development policy must recognise the complex interactions between ecology, land use, and spatial access—and that simply increasing connectivity is not inherently beneficial without complementary socio-economic safeguards.

Although this analysis focuses on Madagascar, the patterns we identify are likely to hold in other parts of the world where protected areas are embedded within heterogeneous socio-ecological landscapes. In many regions of the global South, ecological productivity, land-use change, and infrastructural access similarly shape livelihood outcomes more strongly than simple proximity to protected areas. Our findings therefore reinforce a need for conservation policies that recognise socio-ecological variation, address local drivers of vulnerability, and avoid assuming uniform impacts of protected areas on neighbouring communities. Integrating multidimensional wellbeing metrics into conservation planning can help ensure that protected areas contribute to equitable development outcomes globally.

4.8 Ethics Statement and Funding

The baseline methodology and instruments were approved prior to data collection by the Durrell Wildlife Conservation Trust Ethics Committee, ensuring adherence with best practice in research with human subjects. In line with this data ethics plan, all baseline data was encrypted and transmitted via a secure internet connection to

NIRAS, where it was anonymised and stored in password-protected cloud storage. Only the principal investigator and co-investigators with the requisite ethical training have access to the data file that links the household identification information to the survey responses.

Funding provided by UK International Development to RBG Kew as lead delivery partner through the Biodiverse Landscapes Fund project 'Achieving Sustainable Forest Management through Community Protected Areas in Madagascar' (ecm 62237).

To the sea, to the sea! The white gulls are crying, The wind is blowing, and the white foam is flying.

–J.R.R. Tolkien

5

TIDAL: Rapid Identification and Classification of Mollusc Assemblages

Contents

5.1 Author Contributions 90
5.2 Abstract 91
5.3 Introduction 92
5.3.1 Size Categorisation 93
5.3.2 Drill Hole Use 93
5.3.3 Previous Approaches 94
5.3.4 TIDAL 95
5.4 Materials and Methods 97
5.4.1 Image Sourcing 98
5.4.2 Field Collections 99
5.4.3 Specimen Preparation 101
5.4.4 Lab Imaging 101
5.4.5 Labelling for Computer Vision 102
5.4.6 Machine Learning Pipeline 102
5.4.7 Real-world Size Estimation from £1 Coin Calibration . 108
5.5 Results 109
5.5.1 Mollusc and £1 Coin Detection 110
5.5.2 Species Classification 114
5.5.3 Drill Hole Detection and Classification 117
5.6 Discussion 118
5.6.1 Limitations and Future Work 119
5.7 Conclusion 122

5.1 Author Contributions

Harri Ravenscroft developed the overarching research ideas and conceptual framework. Methodological design was jointly developed by Harri Ravenscroft and Dr Dafydd Ravenscroft. All data collection and data preparation were undertaken by Harri Ravenscroft. Analytical procedures and statistical analyses were performed by Harri Ravenscroft, with advisory input from Dr Dafydd Ravenscroft. The manuscript was drafted by Harri Ravenscroft, with Professor Erin Saupe contributing to revisions. Harri Ravenscroft approved the final version of the chapter.

5.2 Abstract

Marine molluscs play a crucial role in nearshore marine ecosystems, and have high utility as ecological indicators of changing biodiversity and community structures. Our capacity to use mollusc community assessments is currently limited by the ease and rapidity of necessary ecological evaluations: species identification and morphological measurements are often challenging and time-consuming, requiring knowledge from a limited number of highly-skilled professionals.

Developing end-to-end methodologies to facilitate easy assessment of this data source is therefore needed. Here we present TIDAL (Taxonomic Identification and Dimensional Analysis via Learning), a new supervised machine-learning pipeline for the rapid and accurate generation of ecological data from shoreline UK mollusc specimens, alongside a fully-labelled UK database of approximately 13,000 images covering 382 species of gastropods and bivalves. This pipeline can derive species-level taxonomic classifications, high-precision shell size measurements, and other specimen features, such as the presence, type, and size of predatory drill holes from a single image.

TIDAL's modular three-stage design – object detection, species classification and drill hole identification – proved effective and highly flexible. Our bi-phasal object detection protocol greatly improved detection of mollusc specimens and scaling features (£1 coins) within images. We achieved mollusc species validation accuracy of up to 96.5% using ≥ 400 training images per species, and 83.8% with a sparse dataset of ≥ 50 training images per species. Finally, drill hole classification achieved perfect precision and recall.

The TIDAL pipeline is an effective mechanism for rapid generation of modern ecological data for the analysis of community change and distributional shifts in response to sea surface temperature conditions and anthropogenic disturbances.

5.3 Introduction

The dynamics and structure of modern faunal communities are shaped by ecological interactions, predation being key among these. Studies of extant systems have yielded important prey selection theories, with widespread testing of such ideas critical to determining the nature of community dynamics (Kelley and Hansen 2003; Kingsley-Smith, Richardson and Seed 2003; Palmer 1988). These dynamics will be affected by ongoing warming and hypersalinity brought on by climate change (Buschbaum et al. 2024; Chiba and Sato 2016; Warwick and Turk 2002; Coleman, Byrne and Davis 2014).

Temperature and salinity are two of the most important physical factors that affect intertidal gastropod and bivalve molluscs and play a significant role in their abundance and distribution (Rato et al. 2022; Gosling 2008; Broekhuysen 1940; Segal 1961; Verbrugge et al. 2012; Verween, Vincx and Degraer 2007; Berger and Kharazova 1997). Both factors are primed to impact UK waters, with the UK straddling a biogeographic boundary between cold boreal 'northern' waters and warmer Lusitanian 'southern' waters (Lewis 1953). Many intertidal species reach their northern or southern limits of distribution in or close to the British Isles (Mieszkowska et al. 2006; Hiscock et al. 2001) and multi-decadal time series of intertidal species in the UK and Ireland found that some 'southern' species, including *Osilinus lineatus* and *Gibbula umbilicalis*, have undergone north and north-eastern range extensions that are synchronous around the UK, suggesting the effects of climate change are to blame (Mieszkowska et al. 2006). Range extensions and shifts in response to changing sea surface temperature are a commonly noted feature, however, many species are not able to keep pace with shifting temperatures (Moore and Smale 2020), leading to species being subjected to warmer conditions, potentially outside of their ecological niche (Buschbaum et al. 2024).

Against these plethora of changes and impacts, repeated, targeted surveys are required now to establish baselines against which dynamic alterations can be measured (Hiscock et al. 2001). Surveys of nearshore species assemblages are amenable due to their accessibility, feasibility of replication in experimental

design, and ease of identification. The use of time-averaged death assemblages of shelly molluscs (gastropods and bivalves) for surveying is viable as they retain a strong signal of the rank order of relative species abundances in the original source community across a large pool of habitat types (Warwick and Turk 2002; Kidwell 2001; Mieszkowska et al. 2006)

5.3.1 Size Categorisation

Determination of prey and predator size is crucial to investigate predator-prey selection concepts and evaluate the relative age composition of a community (Kingsley-Smith, Richardson and Seed 2003; Palmer 1988; Chiba and Sato 2012). Prey size and shape selection is evidenced across geographic and taxonomic space. Modern rhynchonelliform brachiopods from around New Zealand (Harper, Robinson and Lee 2011); bivalve death assemblages of Cuban beaches (Diez et al. 2023), and the Red Sea (Chattopadhyay, Zuschin and Tomašových 2015); and the limpet *Lottia pelta* (Yanes and Tyler 2009), all show statistically significant evidence of preferential size and shape selection. The size-refuge concept suggests that bivalve prey at the extremes of the observed size range will show lower predation intensities and an increasingly random pattern of drill hole siting due to predator manipulation limits. Manipulation limits being the result of the mechanical limits of the attack strategy wherein predators cannot firmly envelope prey, with their foot, when the item is relatively very small or very large (Chiba and Sato 2012; Chattopadhyay and Dutta 2013).

5.3.2 Drill Hole Use

The identification and study of drill holes on molluscs provides the means by which macroecological predation theories can be tested (Palmer 1988; Harper, Robinson and Lee 2011; Harper, Crame and Sogot 2018; Chattopadhyay and Baumiller 2007). Large drill holes (>0.2 mm in diameter) are predominantly produced by

Naticidae and Muricidae gastropods. The drill holes produced by muricid and naticid gastropods are a predation trace created during feeding, and can be easily recognised due to their circular outline, axis of drill hole perpendicular to the outer valve surface, narrow outer borehole diameter range, and limited size distribution (Harper 2003). The strength of interaction between predator and prey species can be quantified by the predation intensity metric: the proportion of drill holes in any given assemblage. Additionally, the size of the drill hole can also inform on the size of the attacking predator. The inner borehole diameter (IBD) and outer borehole diameter (OBD) have been shown to correlate significantly to causal predator length (Chiba and Sato 2012; Kingsley-Smith, Richardson and Seed 2003)

5.3.3 Previous Approaches

A common computer vision task is to identify 2D images using a set of previously-labelled, or identified, images. This set of images is termed the "training set" and allows the system to build an understanding of how an input (an image) maps to a categorical output (classifier or class) (Hsiang et al. 2019). It is often the case that the production of training sets is limited by access to images of sufficient quality (Liu, Jiang et al. 2022), or that time constraints prevent the formation of sets large enough to train the most complex, deepest, neural networks. Two methods can be employed to combat this problem. The first mechanism is to pre-train a convolutional neural network (CNN) to learn very general features of objects (Kiel 2021; Liu, Jiang et al. 2022). This can be done using image datasets containing millions of general images, the most commonly used of these being ImageNet (Deng et al. 2009). Fine-tuning such pre-trained models can dramatically reduce the number of training examples required to reach an acceptable accuracy than when training from scratch. The second method for generating a larger training set is to perform data augmentation procedures, including random contrast adjustments, flipping and rotations (Perez and Wang 2017).

Existing approaches for the automated determination of marine fauna attributes have typically only focused on resolving a single attribute of the data available in an image (e.g. species). Tests using the EfficientNetV2 (Tan and Le 2019) architecture to categorise genera of the Bivalve and Brachiopod Fossil Image Dataset (BBFID) (870 total genera, with 16,596 labels) output a best identification accuracy of >80% for 22 genera and ~64% accuracy for 343 species (Sun et al. 2023). A similar approach that analysed colour, shape, and texture in 134 shell species with >10 training samples achieved 77.39% shell recognition accuracy initially (Zhang et al. 2019), before a pre-trained deep NN was able to achieve 96% accuracy (Wijaya and Town 2022). The incorporation of pre-trained models was crucial to the higher accuracy outcome of the second taxonomic assessment. A greater number of training images will also positively impact accuracy (Kumar, Mileo et al. 2023; Bianchi and Gonçalves 2021; Wilson et al. 2021). Maravillas, Feliscuzo and Nogra (2022) achieve classification accuracy of 97.29% by deploying >800 images for each of the twelve species within their study. Similar studies with varied marine groups (81% species accuracy using a dataset of >34,000 images of 113 graptolite species (Niu and Xu 2022), 91.85% genus-level classification of 9 radiolarian genera (Dionisio et al. 2020), and ~90% accuracy classifying large foraminifera image sets (Marchant et al. 2020) have been undertaken, but they all fail to determine metrics of interest (size, drill hole type/presence) other than taxonomy.

5.3.4 TIDAL

To address the challenges of large-scale ecological data collection and annotation, we propose TIDAL, an automated machine learning pipeline capable of detecting molluscs in field images, classifying them by species, and identifying and classifying drill holes. This pipeline is designed to support high-throughput ecological analysis by increasing the scalability, objectivity and reproducibility of mollusc-based field studies.

TIDAL is composed of three distinct but integrated model stages, each tailored to a specific sub-task in the analysis:

1. **Object detection** using a YOLOv8 model (Jocher, Chaurasia and Qiu 2023) to localise molluscs and accompanying scale indicators (£1 coins) within raw field images.

2. **Species classification** of each detected mollusc using transfer learning with a convolutional neural network (CNN), applied to individual cropped mollusc regions.

3. **Drill hole detection and classification**, again using YOLOv8 (Jocher, Chaurasia and Qiu 2023), trained on cropped mollusc regions to localise and categorise drill holes into predefined biological types (Bevelled, Straight or Unfinished).

Each model stage was deliberately designed to separate and specialise the computational tasks, improving flexibility and interpretability of the overall system. This modular approach also provides the ability to retrain or swap individual components without affecting the rest of the pipeline.

Existing tools for mollusc image analysis are limited in several important ways. Most pre-existing approaches focus on a single task—typically species identification (Sun et al. 2023; Wijaya and Town 2022; Maravillas, Feliscuzo and Nogra 2022)—and therefore cannot support full predation analysis workflows. None provide an integrated system capable of detecting specimens in the field, identifying species, and locating and classifying drill holes within the same pipeline. In addition, many existing methods rely on rigid architectures, making them difficult to adapt across sites, taxa, or imaging conditions. TIDAL addresses these limitations by combining high-performance deep learning models in a flexible, modular framework that can be retrained or extended as new taxa or ecological questions arise. As such, it represents the first end-to-end, scalable solution designed specifically for mollusc-based predation studies.

The overall objective of TIDAL is to provide a reliable, reproducible and scalable framework for mollusc identification and predation analysis, thereby reducing the bottleneck of manual annotation and enabling more comprehensive ecological data collection at scale.

5.4 Materials and Methods

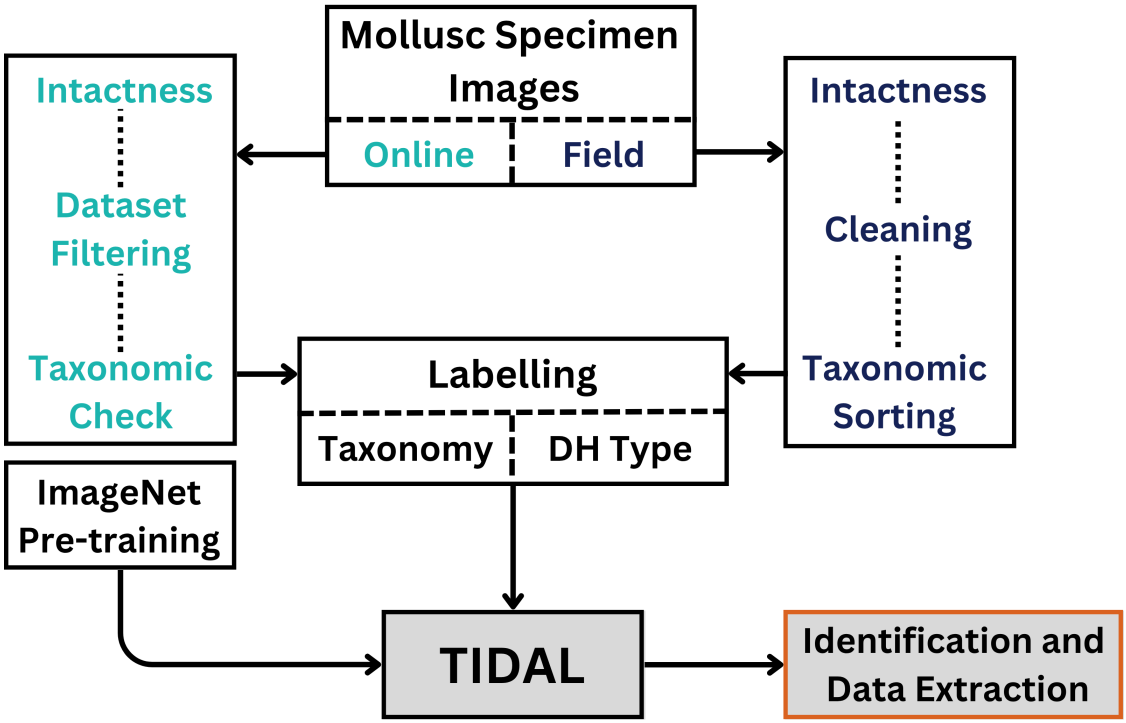


Figure 5.1: The steps in our mollusc identification and data extraction protocol. The primary inputs of this protocol are image files of UK mollusc specimens (Mollusc Specimen Images). Online images, already assigned to a species group, may progress immediately to the labelling stage of the protocol. Field-collected specimens are processed prior to imaging in order to ensure: i) damaged specimens are removed (<100), ii) sand, detritus, and algae have been removed, and iii) taxonomic identification has been conducted to a species level. All images are labelled, using Roboflow online software (Dwyer, Nelson and Hansen 2024), with taxonomic and drill hole (DH) information. Alongside ImageNet pretraining (Deng et al. 2009), this new, labelled, image dataset is used to train our model. The trained model can then be used to identify and extract readily implementable ecological information from field images across the UK

Bivalve and gastropod images are the basis on which TIDAL is constructed.

Both training and test images must clearly show the mollusc, or group of molluscs to allow for accurate taxonomic and drill hole labelling. There is a strong general relationship between the number of images available for training of a neural network in a given class, and the classification score of that network (Kumar, Mileo et al. 2023; Bianchi and Gonçalves 2021; Wilson et al. 2021). Therefore, collection of an image dataset, for later labelling, must be comprehensive in both taxonomic coverage and number of examples per species, or drill hole type.

The three primary sources for the compilation of a comprehensive dataset are: online open-source repositories; taxonomic expert and museum collections; and field studies. Online repository images are those that have been uploaded to an open-source web database, such as the Global Biodiversity Information Facility (GBIF) (*GBIF* 2024a), and have undergone location and taxonomic checks. Similarly, expert and museum collections will have undergone similar checks during digitalisation. Images sourced in these manners will also have dates associated with their collection, although most are collected within the last two decades (Hedrick et al. 2020; Miller et al. 2020; Lendemer et al. 2020). Field studies are also a valuable tool for data collection. Imaging may be conducted with either the target in-situ, at a beach or other shoreline locality, or in a controlled lab set-up, following cleaning and identification procedures.

5.4.1 Image Sourcing

High taxonomic coverage of UK specimens is crucial in meeting the stated target of a model that can rapidly identify and extract morphological information from all UK specimens input. Ideally, all species would be represented by $\gg$10 specimens, in line with previous studies (Zhang et al. 2019; Wijaya and Town 2022), with different orientations present.

Image sourcing from online repositories and sites represent the logistically-easiest mechanism for generation of a training dataset. To generate the TIDAL

training database >11,000 images were sourced from online collections, generating >15,000 labelled annotations.

The primary image source was GBIF (*GBIF* 2024a) (~9,000 images). All image data relating to marine molluscs in the UK region were initially downloaded (~50,000 images). This database was filtered to remove those images with high coordinate uncertainty and any terrestrial or freshwater species. Each image was checked for taxonomic discrepancy against the GBIF assigned label.

Other images were collected from Taxonomically-labelled Internet Sources (TLIS) (~2,000 images). These are images wherein the source has each image checked by taxonomists for categorisation accuracy. This is a highly useful image source due to the high initial confidence of correct species allocation.

5.4.2 Field Collections

To create an in-situ mollusc image dataset for model development and validation, we conducted field excursions across ten coastal sites in England and Wales between February and October 2024 (Table 5.1). These sites were selected to span a broad range of shoreline environments—including mud–silt, sandy beaches, pebble shores, and rocky substrates—to maximise species diversity and ensure that the training dataset was not biased toward any single deposition environment. Each site was visited once, providing a geographically broad but non-repeated sampling of coastal habitats.

Field visits were timed to coincide with low tide and calm sea states to maximise specimen visibility and minimise bias caused by wave-driven transport. Although sampling occurred across multiple months, all sites were surveyed under comparable tidal and weather conditions. This approach ensured consistency in sampling effort and reduced potential confounding effects of seasonally varying storm intensity, which can redistribute shells higher up the beach or fragment them more readily.

At each site, we conducted five equally spaced beach transects (10 m width) extending from the low to high tide marks. These transects were designed to

Site Name	Coordinates	Substrate Type	Sampling Month
Rhossili Bay	51.5732, -4.2919	Sand	August
Blackpill, Swansea Bay	51.5968, -3.9937	Sand/Pebble Mix	February
Broad Haven Beach	51.7814, -5.1044	Sand	October
Marloes Sands	51.7200, -5.2070	Sand	October
Traeth Llanddwyn	53.1446, -4.3919	Sand	September
Traeth Lleiniog	53.2919, -4.0684	Mud/Boulder Mix	September
Seahouses Beach	55.5909, -1.6678	Sand	May
Reculver Beach	51.3798, 1.2073	Sand/Pebble Mix	August
Bowleaze Cove	50.6352, -2.4253	Pebble	October
Charmouth Beach	50.7332, -2.8979	Sand	October

Table 5.1: Field collection sampling localities.

capture variation in the hydrodynamic sorting of shells, as species differ in shell robustness and hydrodynamic behaviour, which influences their deposition zone on the shore (Chattopadhyay 2009). By sampling across the entire intertidal gradient, we minimised the risk of selectively collecting only easily detectable or recently deposited specimens.

Within each transect, all material meeting the $\geq 85\%$ completeness criterion (Klompmaker 2009) and containing no living organism was collected. This threshold reduces selection bias towards visually intact or aesthetically attractive shells (Battista and Schultz 2024; Meineke and Daru 2021; Tang 2000) and ensures that specimens remain identifiable for species classification and drill-hole assessment. All samples from a transect were placed into labelled, sealable bags noting site name, coordinates, date, and transect number.

Although these field samples provide high-quality, known-provenance material, they were not intended to serve as the primary source of training data. Instead, they supplement the larger labelled image dataset by contributing standardised, well-documented examples of shell morphology and predation traces collected under controlled field conditions.

5.4.3 Specimen Preparation

All collected specimens underwent preparation for lab-imaging in the same three step process: completeness check, deep cleaning, and taxonomic sorting.

Initially, all specimens for each transect within a locality were checked for completeness. Following (Klompmaker 2009) all specimens must conform to ≥85% shell completeness. If gastropods are missing a significant portion of their spire (≥30%) they were rejected regardless of the remaining completeness of their shell. Heavily damaged specimens were excluded as they lack features necessary for accurate taxonomic designation (Stribling, Moulton and Lester 2003).

All specimens were cleaned by transect batch for a given locality. A 50-50 mixture of bleach and water was used and all specimens were fully exposed to this for 15 continuous minutes to kill any bacterial colonies growing on the material without reducing shell integrity or colouration. Upon removal, each specimen was washed thoroughly in water to remove any residual cleaning solution and dried thoroughly.

Taxonomic sorting was conducted following the WoRMS (Ahyong et al. 2024) taxonomy. Each specimen was checked for signature morphological features using a variety of sources (WoRMS (Ahyong et al. 2024), MarLIN (*Marine Life Information Network* 2020), (*Mollusca Types in Britain and Ireland* 2023)).

5.4.4 Lab Imaging

Controlled lab imaging was conducted for all specimens collected across the ten field sites with a consistent set up, including: fixed imaging table with vertically translatable camera stand; white-tone ring light; black velvet backing material, for specimen-background contrast; Nikon D7200 DSLR camera with macro lens; and £1 coin for scale. The £1 coin was used as scale due to its commonality and ease of access, when compared to typical binary colour scale bars, and the uncommon shape of the coin (12-sided) meant that it can easily be identified in any specimen image. Additionally, diameter is consistent across all £1 coins (22.5 mm).

Two images were captured for each individual mollusc specimen. For bivalves, we imaged both the outer and inner valve surface. For gastropods, we captured one image of the specimen with aperture facing towards the camera, and the other was spire-directed towards the camera. The orientations were chosen to maximise the morphological features for the training dataset (Callomon 2019).

5.4.5 Labelling for Computer Vision

To create the train-test-validation dataset for TIDAL the free online image labelling and dataset testing software "Roboflow" (Dwyer, Nelson and Hansen 2024) was used. Initially, an image segmentation project is created, and images are uploaded into this in png or jpg format. For the purposes of this project the uploaded images were split into three categories, dependent upon image source (GBIF (*GBIF* 2024a), taxonomically-labelled internet sources (TLIS), own imagery (Field Collection)).

	Image Source		
	GBIF	Field Collection	TLIS
Image Count	9,152	1,785	1,950
Annotation Count	11,774	2,158	3,330
Species Classes	150	37	360
DH Classes	2	3	2

Table 5.2: Breakdown of labelling for each component image dataset. Annotation count is higher than image count as there can be multiple annotations per image, due to multiple specimens or the presence of a drill hole (DH). **GBIF** - Global Biodiversity Information Facility, **Field Collection** - images taken by authors from their own fieldwork collections, **TLIS** - Taxonomically-labelled Internet Sources.

5.4.6 Machine Learning Pipeline

The pipeline comprised three main sequential stages:

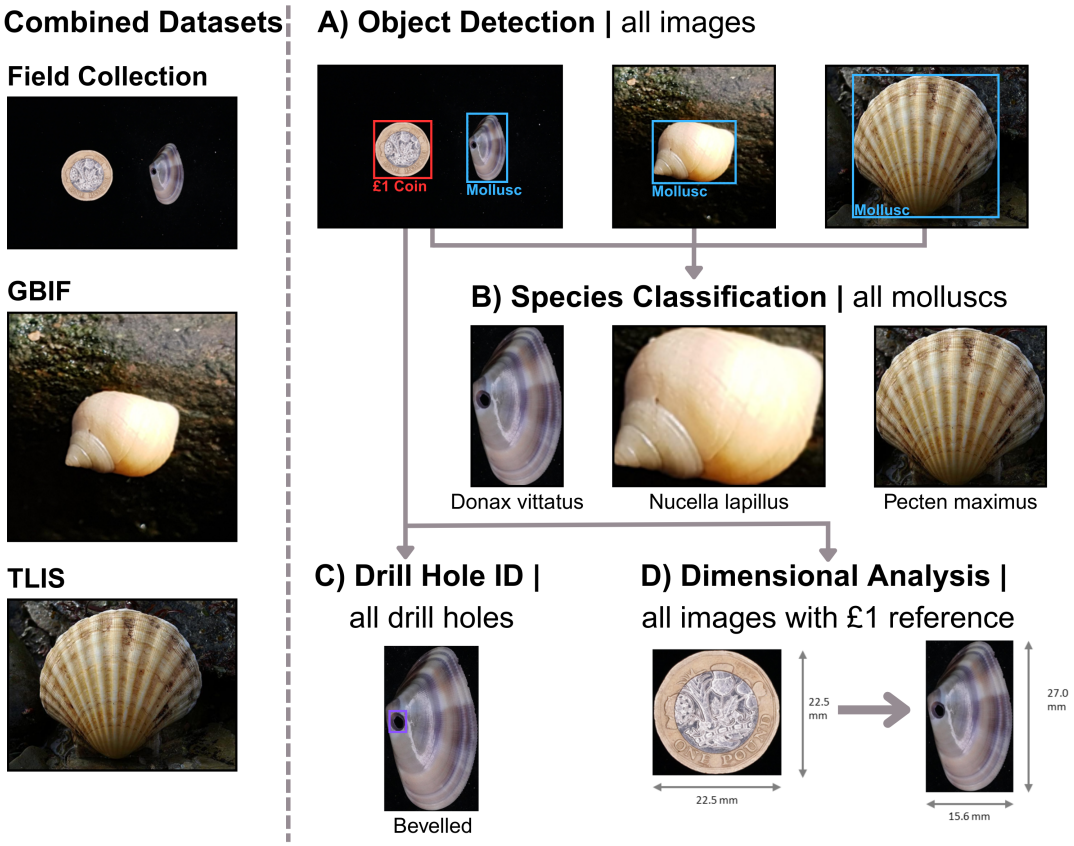


Figure 5.2: Flowchart showing the stages of the TIDAL pipeline. All images from the three datasets are used. **A)** Object detection using a YOLOv8 model (Jocher, Chaurasia and Qiu 2023) is performed in which molluscs and £1 coins are identified and bounding boxes applied. **B)** The identified molluscs are cropped and fed to a machine learning (ML) classifier which predicts their species class. **C)** Images with drill holes are fed through another YOLOv8 model (Jocher, Chaurasia and Qiu 2023) to create bounding boxes for the drill hole and categorise them as straight or bevelled. **D)** Images containing £1 coins undergo dimensional analysis – using the known size reference the size of the mollusc, and any drill holes it contains can be calculated.

5.4.6.1 Stage 1: mollusc and £1 coin detection

The first stage of the pipeline is dedicated to detecting molluscs and reference objects (£1 coins) within raw field images. This step is important for the downstream tasks in the pipeline, namely, species classification, drill hole identification and size estimation, which all rely on the accuracy of the localisation of regions of interests (ROIs) at this stage. Even small localisation errors could impact cropping, measurement or classification in later stages, so reliable performance here supports the pipeline's overall consistency and utility.

To solve this task, we adopted YOLOv8 (Jocher, Chaurasia and Qiu 2023), a state-of-the-art object detection architecture developed by Ultralytics. YOLOv8 was selected for its strong trade-off between speed and accuracy, its ability to generalise well with limited data, and its streamlined training process and support for deployment in varied environments. The model was trained to detect two object classes: mollusc and £1 coin. To maximise robustness while leveraging high-quality labels, we implemented a two-phase training approach:

Phase 1: Clean data pretraining

The model was initially trained on the "Field Collection" dataset, a carefully curated collection of images featuring single molluscs photographed under controlled conditions with a visible £1 coin placed nearby for scale. Each image in this dataset was fully annotated with bounding boxes and segmentation masks for both molluscs and coins, providing highly reliable training labels.

Images were resized to a fixed resolution of 640×640 pixels and augmented using standard image transformations, including random contrast adjustments, flipping and rotations, to promote generalisability (Perez and Wang 2017). This phase focused on helping the model develop precise detection capability, particularly for £1 coins, which are essential for later size analysis.

Phase 2: Adaptation to field conditions

Following pretraining, the model was retrained on a combined dataset comprising all three datasets (Field Collection, TLIS, and GBIF).

In this phase, £1 coin annotations were only available in "Field Collection". To extend coin detection coverage, the pretrained model from Phase 1 was applied to the other two datasets to generate candidate £1 coin detections. Predictions exceeding a defined confidence threshold of 0.5 were extracted and subsequently reviewed to identify high-confidence, visually verified detections. These verified predictions were then used as pseudo-labels to augment training data in Phase 2, enabling coin localisation in previously unlabelled datasets. All mollusc annotations

were retained as ground truth. Images were resized to 640×640 and bounding boxes scaled accordingly.

This two-phase process is analogous to curriculum learning, starting with simple, high-confidence examples and gradually introducing more complex, noisy data (Bengio et al. 2009). It also reflects semi-supervised learning, using model predictions to extend sparse annotations (Engelen and Hoos 2020).

At inference time, the trained model processes raw input images and returns a list of detections, each comprising a bounding box, class label, and confidence score. Detected mollusc ROIs are cropped for subsequent species classification and drill hole detection, while £1 coin detections enable pixel-to-millimetre conversion via the known coin diameter (22.5 mm). This enables accurate estimation of absolute mollusc and drill hole sizes.

5.4.6.2 Stage 2: species classification

The second stage of the pipeline focuses on identifying the species of each individual mollusc. This is treated as a multi-class image classification task, where the goal is to assign one of many possible species labels to each mollusc detected in the previous stage (Joshi, Porikli and Papanikolopoulos 2009).

Species classification is performed using cropped mollusc ROIs, extracted from the original images using the bounding boxes labels. By isolating the mollusc from the surrounding environment and other distracting elements, this approach provides a focused input for classification and reduces background noise, improving both model accuracy and interpretability.

Several well-established convolutional neural network (CNN) architectures were tested for this task: MobileNetV2 (Sandler et al. 2018), ResNet50 (He et al. 2015), EfficientNetB0 and EfficientNetB1 (Tan and Le 2019). These models were selected based on their proven performance in image classification tasks, their widespread use in ecological and biological imaging contexts (Sun et al. 2023; X. Liu and Song 2020), and the availability of ImageNet-pretrained weights for each (Deng et al.

2009), which facilitates transfer learning. Each architecture represents a different point on the trade-off curve between computational efficiency and model capacity:

- **MobileNetV2** offers fast inference with low memory usage, suitable for lightweight applications or deployment (Sandler et al. 2018).
- **ResNet50** is a deeper, residual network known for stable performance across domains (He et al. 2015).
- **EfficientNetB0** and **B1** use compound scaling to balance depth, width and resolution, often achieving high accuracy with fewer parameters (Tan and Le 2019).

To address the challenge of limited labelled data for many species, transfer learning was used in all models. This approach leverages general-purpose feature representations learned from large-scale datasets and adapts them to the specific mollusc classification task (Zhuang et al. 2021).

Training was performed in two phases:

1. **Feature extraction phase:** the pretrained base layers of each model were frozen and only the top classification layers were trained using the mollusc images. This allows the model to quickly adapt its output layer to the target classes without disrupting well-established low-level features.
2. **Fine-tuning phase:** all layers were unfrozen and trained end-to-end at a reduced learning rate, allowing the model to refine both general and task-specific features for improved performance.

Each model was trained and evaluated at multiple class thresholds (minimum number of samples per species): 50, 100, 200 and 400. This enabled analysis of how performance scaled with class diversity and training set size. This trade-off is important for ecological studies where rare species may be underrepresented

but biologically important (Y. Cao, D. D. Williams and N. E. Williams 1998; Dee et al. 2019).

Input images were resized to 224×224 pixels and pixel values were normalised using architecture-specific preprocessing functions to match the original training conditions of each model. The model was trained using the categorical cross-entropy loss function and the Adam optimiser (Kingma and Ba 2014). To promote generalisation and prevent overfitting, early stopping was applied with a patience of 5 epochs, restoring the best weights based on validation loss. Additionally, learning rate scheduling was implemented via a ReduceLROnPlateau (TensorFlow v2.16.1 (Martín Abadi et al. 2015)) callback, which halved the learning rate if validation loss did not improve for 3 consecutive epochs.

5.4.6.3 Stage 3: drill hole detection and classification

The final stage of the pipeline involves detecting and categorising predatory drill holes within mollusc shells. Drill holes offer valuable ecological insight into predator-prey interactions and shell usage patterns, and identifying their presence and morphology can aid in understanding the behaviour of mollusc species (Harper, Robinson and Lee 2011; Chattopadhyay 2009; Chattopadhyay and Baumiller 2007).

This task was approached using a second YOLOv8 model, trained separately from the mollusc detector and applied specifically to the cropped mollusc ROIs. This allows the classification model to focus on fine-grained features of the shell surface, avoiding the complexity of interpreting full-scene field images.

The classification model was trained to detect and classify two drill hole types: *Bevel* and *Straight*. A third label, *Unfinished*, was present in the original annotations but was excluded from training and evaluation due to its extremely limited representation in the dataset. Removing this class allowed the model to focus on reliably classifiable examples and avoid training instability or biased predictions due to imbalance.

Training data consisted of bounding box annotations for drill holes located within mollusc shells. During preprocessing, each drill hole's bounding box was transformed to coordinates relative to its enclosing mollusc ROI, ensuring consistency in scale and orientation. Each image was resized to 640×640 pixels and data augmentation techniques such as flipping and contrast adjustment were used to improve robustness (Perez and Wang 2017).

As with the earlier detection model, YOLOv8 was selected for its ability to perform detection and classification in a single pass, with high efficiency and accuracy. The output of the model consists of bounding boxes indicating the predicted location of each drill hole within the mollusc, along with the predicted class (*Bevel* or *Straight*) and a confidence score.

Drill hole detection is designed to operate independently from species classification. This enables the pipeline to make drill hole predictions for any detected mollusc, regardless of species, making the approach applicable to broader ecological datasets where species identity may not always be available or confidently classified.

5.4.7 Real-world Size Estimation from £1 Coin Calibration

To enable size-based analysis of mollusc morphology and drill hole characteristics, the pipeline incorporates a pixel-to-millimetre scaling step based on the detection of £1 coins within the original field images. £1 coins serve as a convenient and reliable reference object, with a known real-world diameter of 22.5 mm.

During the mollusc and coin detection stage, any detected £1 coins are localised using bounding boxes predicted by the YOLOv8 model. When a £1 coin is detected within the same image as a mollusc, its bounding box width (in pixels) is used to compute the pixel-to-millimetre conversion factor. Specifically, the average pixel width of the coin's bounding box is divided by 22.5 mm, yielding an image-specific pixels-per-millimetre ratio.

This ratio is then applied to the bounding box of the corresponding mollusc or drill hole to estimate its real-world width, height and diagonal length in millimetres.

In the case where no coin is detected in a particular image, no size estimation is performed for that sample.

By incorporating a simple and consistent real-world scaling mechanism, this step enhances the interpretability of the results and enables biologically meaningful comparisons across samples.

5.5 Results

Each phase of modelling was evaluated using a tailored set of performance metrics, appropriate to the nature of the task - classification or object detection. The following metrics were employed:

- **Accuracy:** measures the proportion of correctly classified instances out of the total number of predictions.
- **Top-5 accuracy:** considers a prediction correct if the true class is among the model's five most probable predictions.
- **Precision:** measures the proportion of predicted objects that were correct, i.e., true positives out of all predicted positives.
- **Recall:** measures the proportion of actual objects that were correctly detected, i.e., true positives out of all ground truth objects.
- **F1 score:** the harmonic mean of precision and recall, this provides a single measure that balances both false positives and false negatives.
- **Intersection over Union (IoU):** For bounding box comparison, IoU measures the overlap between predicted and ground truth boxes.
- **mAP@0.5:** (mean Average Precision at IoU ≥ 0.5) averages the precision over all classes and detection thresholds, counting a detection as correct if its bounding box overlaps the ground truth by at least 50%.

- **mAP@0.5:0.95:** is a stricter metric that averages mAP across multiple IoU thresholds (from 0.5 to 0.95 in steps of 0.05), providing a more nuanced measure of localisation accuracy.

5.5.1 Mollusc and £1 Coin Detection

The effectiveness of the object detection module was evaluated across the three datasets using both the Phase 1 and Phase 2 YOLOv8 models. Precision, recall, mean Intersection over Union (IoU), and mean Average Precision (mAP) were computed for mollusc and pound coin detection. Evaluations were conducted using a confidence threshold of 0.5 and an IoU threshold of 0.5, providing a balance between tolerance for detection variance and discrimination of false positives.

On the "Field Collection" dataset, which contains high-quality, curated images with consistent lighting and well-annotated molluscs and £1 coins, the Phase 1 model already achieved excellent performance (Figures D.1 and D.2). Mollusc detection yielded a precision of 0.979 and recall of 0.956, with a mean IoU of 0.961 and mAP@0.5 of 0.984. The detection of £1 coins in this dataset was perfect (precision and recall of 1.000). These results indicate that under controlled conditions, the model is highly reliable for both object classes.

Following the two-phase training strategy, the Phase 2 model retained this high performance on the "Field Collection" data, further improving mollusc precision to 1.000 and recall to 0.997, and increasing mean IoU to 0.987. This indicates that the additional training data introduced in Phase 2 did not degrade performance in curated conditions and actually offered slight gains likely due to improved generalisation.

The generalisation capabilities of the model were more substantially tested on the TLIS and GBIF datasets, both of which contain images with more complex visual environments. These datasets introduce challenges such as natural beach settings, cluttered backgrounds, varying lighting conditions and multiple or partially visible molluscs within a single frame. Importantly, these datasets do not contain labelled £1 coins, and their presence is sparse or absent entirely, meaning pound coin evaluation is either omitted or based on model-generated pseudo-labelling where available.

On the TLIS dataset, the Phase 1 model struggled to generalise, achieving a recall of just 0.321 for mollusc detection and a precision of 0.773. The mAP@0.5 stood at 0.866. These results reflect a high false-negative rate (many molluscs were simply missed) while the moderate precision suggests some degree of false positives. However, the Phase 2 model significantly improved performance on this dataset, with mollusc recall increasing to 0.875 and mAP@0.5 reaching 0.970. The increase in mean IoU to 0.943 reflects much stronger bounding box alignment, highlighting the importance of exposing the model to more diverse training examples (Figures D.1 and D.2).

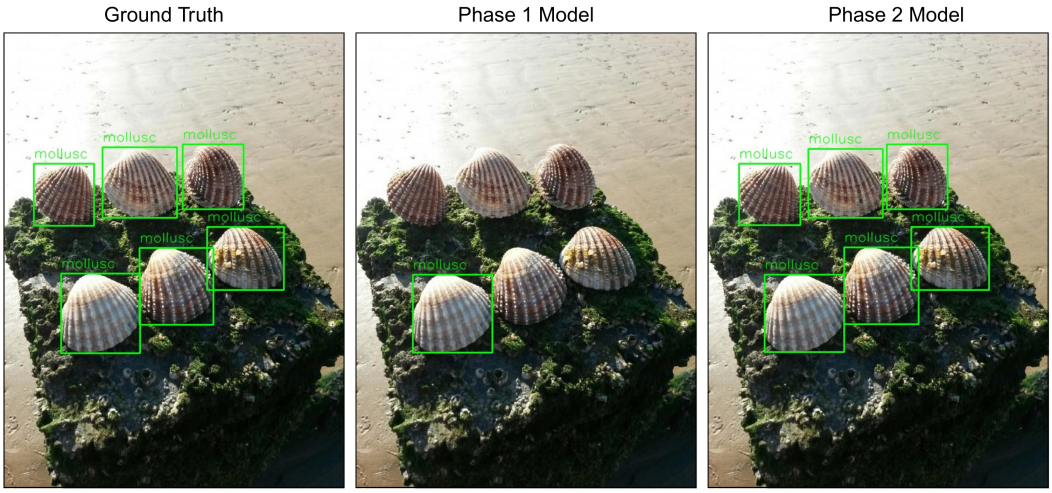


Figure 5.3: Detection of multiple mollusc specimens in a natural environment image. Ground truth has been initially defined by labelling.

Similar trends were observed on the more challenging GBIF dataset. The Phase 1 model exhibited poor mollusc recall at 0.203 and a low precision of 0.402, indicating both frequent misses and false detections. With Phase 2 training, recall rose sharply to 0.923 and precision improved to 0.755 (Figures D.1 and D.2). The mean IoU increased from 0.848 to 0.956, and mAP@0.5 climbed from 0.741 to 0.963. These results confirm that the Phase 2 model was far better equipped to handle real-world complexity, including occlusions, overlapping specimens, and variable capture conditions. Although some false positives remained, particularly

for molluscs in noisy backgrounds, the Phase 2 model successfully mitigated the weaknesses observed in Phase 1.

Notably, pound coin detection remained perfect in "Field Collection" across both phases, validating the use of pseudo-labels generated from the Phase 1 model in Phase 2 training. Although pound coins are rarely present in the other datasets, the limited number of detections in GBIF showed strong agreement with the ground truth, suggesting the model was capable of recognising £1 coins even in novel environments.

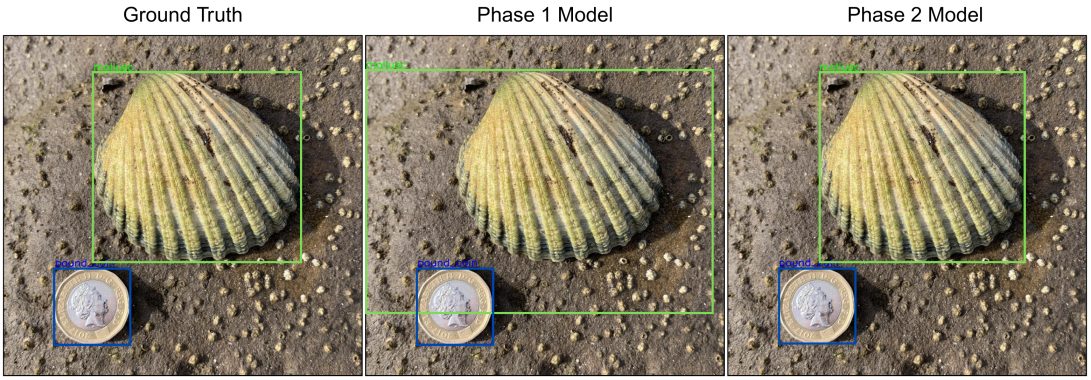


Figure 5.4: Detection of a single mollusc specimen alongside a pound coin measurement item. Ground truth has been initially defined by labelling.

While the second phase YOLO model exhibited high overall accuracy and strong generalisation across diverse datasets, a number of important caveats emerged when evaluating performance, particularly on the GBIF dataset. These reflect not only model limitations but also inherent challenges in working with large, heterogeneous, and partially annotated datasets.

One notable source of false positives in the "pound_coin" class was the presence of circular objects that were not £1 coins but were used for scale in some field photographs. These included other denominations or tokens, which despite their visual similarity to a £1 coin were not explicitly annotated in the dataset and were never encountered during training. As a result, the model occasionally misclassified these distractors as pound coins reducing precision scores. Future work could mitigate this issue through targeted negative mining, incorporating examples of non-£1 coins explicitly labelled as background. Alternatively, a multi-class setup

distinguishing true £1 coins from other reference objects may improve robustness in visually ambiguous cases.

A second challenge was posed by inconsistent or incomplete labelling of molluscs in crowded or noisy scenes, particularly in GBIF and TLIS datasets. Unlike the curated "Field Collection" dataset where all molluscs were reliably annotated, many field images in GBIF and TLIS included multiple individuals, some of which were partially occluded, overlapping or relegated to the background. Annotation efforts appeared to prioritise clearer, foreground molluscs, leading to false positive detections by the model in cases where unlabelled molluscs were nonetheless correctly localised. This issue likely contributed to inflated false positive rates and underestimation of model precision during evaluation. Addressing this would require a more exhaustive annotation process or the use of weak supervision to augment labels via model-generated proposals verified by human reviewers.

These findings highlight the limitations of relying solely on standard evaluation metrics in noisy real-world datasets. The model may generalise better than metrics suggest when accounting for ambiguities in labelling or incomplete ground truth. Nevertheless, these results reinforce the importance of high-quality annotations in evaluation and suggest that detection-based pipelines may benefit from additional curation or post-hoc filtering mechanisms when deployed in uncontrolled environments.

The two-phase training approach proved important for generalisation. Training first on curated data established robust foundational detection for structured images and reliable pseudo-labelling of £1 coins. Subsequent training on a wider set of real-world examples dramatically improved recall, bounding box accuracy, and overall robustness. The results confirm that mollusc and coin detection in diverse field settings is not only achievable with YOLOv8 but can be effectively adapted with a staged training strategy designed to balance accuracy and generalisability.

5.5.2 Species Classification

The central task in this study is the classification of mollusc species based on visual morphology. Following initial object detection, cropped mollusc regions were passed to a series of CNN classifiers to assign species-level labels. This stage is designed to replicate and ultimately replace the manual species identification process traditionally carried out by field researchers and taxonomists.

To evaluate the impact of sample size and class imbalance on performance, classification models were trained using data subsets filtered by four sample-count thresholds: 50, 100, 200 and 400 images per species. These thresholds determined which species were retained in each training run, allowing for controlled experiments that balance species diversity with the depth of representation per class.

Four commonly used CNN architectures were compared: EfficientNetB0 (Tan and Le 2019), EfficientNetB1 (Tan and Le 2019), ResNet50 (He et al. 2015), and MobileNetV2 (Sandler et al. 2018). These models vary in depth, capacity and computational requirements, and all were fine-tuned from pretrained ImageNet (Deng et al. 2009) weights using the two-phase transfer learning strategy as detailed in the Methods section. Validation accuracy for each model and threshold combination are shown in 5.3 and 5.5.

Image Count Threshold	Number of Classes (Unique Species)	Model Framework Accuracy			
		EfficientNetB0	EfficientNetB1	ResNet50	MobileNetV2
50	63	83.8%	82.5%	82.5%	76.4%
100	47	85.5%	86.5%	86.5%	80.8%
200	32	89.5%	88.1%	89.5%	89.5%
400	8	95.6%	96.5%	96.5%	94.1%

Table 5.3: Species classification validation accuracy percentage. Accuracy scores are shown for each model at the specified image count thresholds.

As expected, validation accuracy increases with higher thresholds, reflecting the trade-off between the number of species represented and the volume of training data available for each class. EfficientNet models consistently performed well across thresholds, particularly at lower thresholds where they maintained relatively high accuracy despite the increased number of classes. ResNet50 matched or exceeded

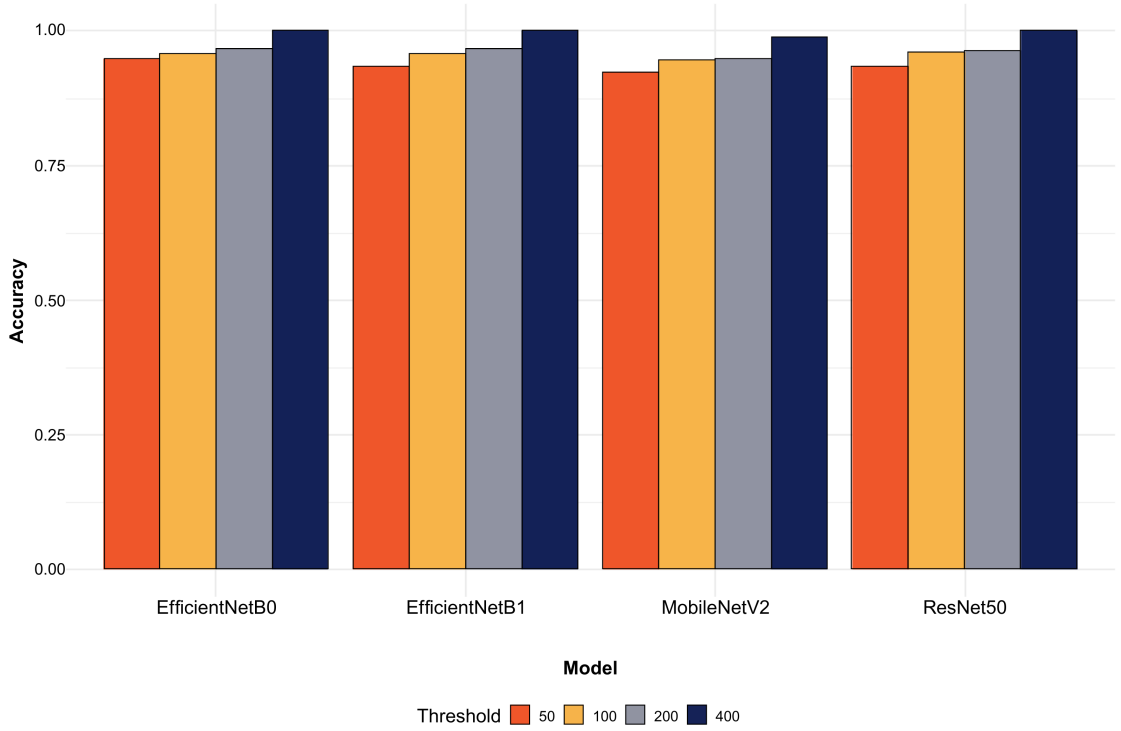


Figure 5.5: Validation accuracy of model frameworks with varied threshold image counts.

EfficientNet performance at higher thresholds. While MobileNetV2 may be the fastest and most lightweight model, its reduced capacity limits its performance on fine-grained species discrimination and it underperformed across all thresholds.

As well as accuracy other metrics were explored to better understand the performance of the model on a dataset with unbalanced classes (D.3 and D.4). Precision and recall were generally well balanced across all experiments, suggesting that models were neither overly conservative nor prone to over-classification. High precision indicates few false positives, while high recall reflects success in capturing the correct class when it's present. The similarity between these two metrics across models suggests stable classification behaviour without extreme bias toward precision or recall.

Top-5 accuracy remained consistently high across all runs indicating that even when the top prediction was incorrect, the correct species was often among the top five suggestions. This can be valuable in ecological workflows, where automated systems might assist rather than replace human experts. A shortlist of probable

species can greatly reduce the verification burden, especially when distinguishing between visually similar taxa.

We also assessed how performance varied across datasets using a selected model (EfficientNetB0) at a threshold of 200. As shown in figure 6, the "Field Collection" dataset demonstrated the strongest performance across all metrics, followed closely by TLIS. The GBIF dataset consistently underperformed, especially in precision and recall. This is likely due to its higher visual complexity as images are often captured in situ, with cluttered natural backgrounds, multiple overlapping molluscs, partial occlusions and diverse lighting conditions. These factors introduce noise and ambiguity that are absent in the more controlled photographic conditions of the other datasets. The same pattern is seen at other thresholds, as in (D.4).

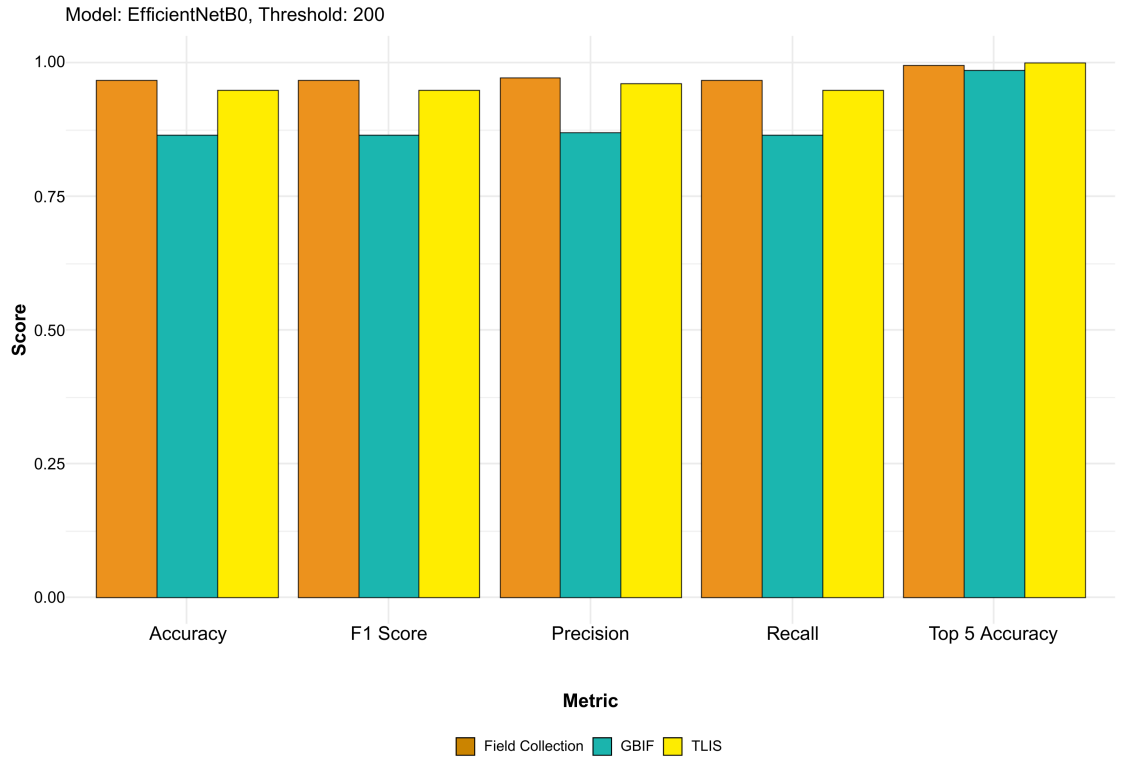


Figure 5.6: Comparison of datasets across evaluation metrics. Model structure (EfficientNetB0) and image threshold (200) has been held constant.

5.5.3 Drill Hole Detection and Classification

The final stage of the pipeline focuses on detecting and classifying predatory drill holes on mollusc shells using a specialised YOLOv8 model. This model was trained on cropped mollusc ROIs to detect two types of drill holes, *"Bevel"* and *"Straight"*, based on visual morphology. A third category, *"Unfinished"*, was excluded from model training and evaluation due to an insufficient number of labelled instances, which would have limited the model's ability to learn meaningful patterns for that class.

Table 3 (5.4) summarises the model's performance on a validation set of 33 images with drill holes across the two retained classes. The model achieved high precision and recall for both classes, with mAP@0.5 values close to 1.0 and strong localisation performance under the more stringent mAP@0.5:0.95 metric.

Class	Precision	Recall	mAP@0.5	mAP@0.5:0.95
Bevel	1	1	0.995	0.935
Straight	1	1	0.995	0.929

Table 5.4: Evaluation metric results for drill hole classification.

The model exhibited perfect precision for both drill hole types, indicating no false positive detections in the validation set. Recall values were also perfect, suggesting reliable coverage of ground-truth annotations. These results demonstrate the model's capability to detect subtle and visually localised features, even with relatively modest training data volumes.

Drill hole detection is implemented independently of species classification, allowing predictions to be made on any identified mollusc, regardless of species label. This improves the pipeline's flexibility and enables application to broader datasets, including those where species classification may be uncertain or incomplete.

5.6 Discussion

TIDAL provides a practical, modular framework for automating the analysis of mollusc field images. It combines object detection and classification techniques to detect molluscs and drill holes, identify species and enable biologically meaningful measurements. Overall, the system demonstrates strong performance across tasks, achieving high accuracy in mollusc detection, robust species classification across a range of dataset thresholds and reliable localisation and classification of drill holes.

The use of a modular three-stage design - object detection, species classification and drill hole identification - proved effective. It allowed each model to be optimised for a specific task and made the pipeline more interpretable and easier to adapt or extend. For instance, species classification can be retrained independently on new species subsets without needing to modify the detection or drill hole stages. Similarly, the reuse of mollusc ROIs across stages ensures that outputs are spatially aligned and biologically meaningful.

Despite these strengths, there are several important limitations to acknowledge. In the species classification stage, performance was highly dependent on the number of classes and sample availability per class. Models performed well at higher thresholds (with fewer, better-sampled species), but accuracy dropped when trying to classify larger numbers of species with limited representation. This trade-off between ecological coverage and predictive accuracy is a key challenge in biodiversity modelling (Dionisio et al. 2020; Maravillas, Feliscuzo and Nogra 2022), and future work may explore class balancing strategies, data augmentation (Alomar, Aysel and Cai 2023; Kumar, Brennan et al. 2024) or ensemble models (Mahadevkar et al. 2022) to address it.

Finally, while this pipeline was tested on a controlled and annotated dataset, its generalisability to other imaging conditions, field settings or mollusc assemblages remains to be evaluated. Applying the system to a broader set of field images, potentially with more diverse lighting, substrate or shell presentation, would provide a more complete picture of its robustness in practical ecological workflows.

Overall, the results presented here demonstrate the value of applying machine learning to ecological image analysis and establish a strong foundation for future enhancements. It provides a scalable, extensible framework that can support a range of ecological investigations into mollusc communities, predator-prey interactions and shell morphology, reducing reliance on manual image interpretation and enabling larger-scale and more consistent data collection.

5.6.1 Limitations and Future Work

While the TIDAL pipeline demonstrates strong performance across mollusc and pound coin detection, species classification, and drill hole identification, several limitations remain. These highlight valuable opportunities for improvement in future iterations of the system.

A central challenge lies in the diversity and completeness of the datasets used. The first phase of detection training relied primarily on the Field Collection dataset, which consists of high-quality, well-lit images with clearly annotated molluscs and pound coins. While this allowed for strong initial model performance, especially in controlled conditions, generalisation to less curated images was more challenging. The GBIF and TLIS datasets presented more difficult scenarios, with molluscs photographed in situ, complex natural backgrounds, varying lighting, and partial occlusions. These real-world conditions led to reduced performance, particularly in the earlier Phase 1 model. While Phase 2 training substantially improved generalisation by incorporating more diverse examples, future work would benefit from a larger and more balanced training dataset that better captures the range of ecological imaging scenarios.

Expanding the dataset to include more examples of underrepresented species and environments would improve model robustness. This could be achieved by directly encouraging broader image collection, for instance through citizen science platforms where users capture mollusc photos in the wild using mobile devices. These citizen science projects have proved to be highly effective for other

taxonomic groups (Mäder et al. 2021; stagmerlinapp n.d.) If paired with automated annotation tools (Zhuravlyov and Grigorian 2024; Sager, Janiesch and Zschech 2021; Beijbom et al. 2012), such a framework could greatly accelerate dataset growth and improve generalisation.

Species classification performance was also impacted by class imbalance and the fine-grained nature of mollusc taxonomy. Many species exhibit only subtle morphological differences and rarer species had to be excluded from training due to limited sample size. While this improved accuracy for the well-represented species, it reduced ecological coverage. Future efforts should aim to include more rare species by collecting additional images and exploring classification strategies that better handle sparse classes. Hierarchical classification (Ho et al. 2023), in which genus or family are predicted before species, could be explored. If suitable image metadata were available, models that incorporate contextual cues such as geographic location or co-occurring species (Terry, Roy and August 2020), could improve accuracy in difficult cases.

The current system uses a top-1 classification prediction, but top-5 accuracy was demonstrated to remain high even when the most confident prediction was incorrect. This suggests the model often included the correct label among its top candidates and future versions could take advantage of this by presenting a shortlist to human users in an assisted classification workflow. Additionally, class probability scores already produced by the model could be used to communicate prediction certainty, filter low-confidence output or enable semi-automated review (Mahadevkar et al. 2022).

A further limitation stems from inconsistencies in ground truth annotations, especially in the GBIF dataset. Many images contain additional molluscs beyond those explicitly annotated, particularly in the background or at the edges of images, leading to false positives when the model correctly detects these unlabelled individuals. In such cases, measured precision is artificially reduced despite accurate model behaviour. To address this, future versions of the pipeline could include optional human review of predictions, semi-automated labelling tools (Mahadevkar et al. 2022) to refine incomplete annotations or uncertainty-aware metrics that

tolerate ambiguous labels during evaluation (Tabarisaadi et al. 2024; Mendes, Romano and Garlan 2024; Chandu et al. 2024).

Drill hole classification, though highly accurate on the validation set, was limited by a small number of positive training examples. The dataset contained far more molluscs without drill holes than with them, and only two types of drill holes, *Bevel* and *Straight*, were included. The third category, *Unfinished*, had too few labelled examples to be reliably modelled. More balanced datasets, with additional annotated examples of all three classes and explicit inclusion of negative samples would help reduce this bias. Further to this, the current model only detects drill hole presence and type. It does not yet incorporate information about hole placement or shape, which may be biologically relevant (Harper, Crame and Sogot 2018; Kingsley-Smith, Richardson and Seed 2003). Future iterations could integrate segmentation models or geometric analysis to support more detailed morphological classification.

Additionally, the use of bounding boxes to localise drill holes may be too coarse given their small, irregular shape. Future work could explore segmentation-based approaches to better capture their structure (Minaee et al. 2022). Including geometric features such as shape or position relative to the shell may also support more nuanced biological interpretation.

Across all three stages, the current pipeline is modular with each model operating independently and sequentially. While this offers flexibility and transparency, it also means errors can propagate from one stage to the next. For example, if a mollusc is not detected in the first stage, it will not be passed to the species classifier or drill hole detector. One potential future direction is to explore end-to-end models that learn all tasks jointly, potentially improving consistency and robustness. Additionally, integrating model interpretability techniques such as saliency maps (Gupta et al. 2020; G. Li and Yu 2015) or Grad-CAM visualisations (Selvaraju et al. 2017) could help diagnose errors, build trust, and provide insight into the biological basis for predictions.

While the TIDAL pipeline already shows strong potential as an automated tool for analysing mollusc imagery, future improvements should focus on expanding

and diversifying the training dataset, improving annotation quality, incorporating more sophisticated classification strategies, and refining detection granularity. With these developments, the system could support a wider range of ecological research questions and enable scalable, field-ready deployment in biodiversity monitoring and citizen science initiatives.

5.7 Conclusion

This study presents a modular machine learning pipeline, TIDAL, that successfully automates the analysis of mollusc field images by integrating object detection, species classification, and drill hole identification. The system shows strong performance across all tasks and offers a flexible, scalable framework for ecological image analysis.

While challenges remain, particularly around species classification with limited data and the need for more comprehensive drill hole annotations, the pipeline lays a solid foundation for future improvements. Its design enables easy adaptation and integration into broader ecological workflows, with the potential to enhance research on mollusc communities, predation patterns, and shell morphology at scale. Future work should focus on improving model generalisability, expanding annotated datasets, and incorporating additional biologically relevant features to maximise its utility in diverse field conditions.

All's well that ends better.

–J.R.R. Tolkien

6 Discussion and Conclusions

Contents

6.1 Contribution to Knowledge 124
6.1.1 Chapter 2: Release the HOGS 124
6.1.2 Chapter 3: Southern Africa Bird Shifts 125
6.1.3 Chapter 4: Madagascan Protected Areas 125
6.1.4 Chapter 5: Spot the Mollusc 127
6.2 Limitations and Opportunities 128
6.2.1 Chapter 2: Release the HOGS 128
6.2.2 Chapter 3: Southern Africa Bird Shifts 129
6.2.3 Chapter 4: Madagascan Protected Areas 130
6.2.4 Chapter 5: Spot the Mollusc 131
6.3 Closing Remarks 132

In this thesis I developed new protocols for the study of macroecology across a range of spatial and temporal scales and use these tools to investigate climatic and anthropogenic impacts on modern communities and ecosystems. I began with the development and testing of the Historical Occurrence Georeferencing System (HOGS): a workflow designed to extract historical occurrence data previously trapped behind the wall of inordinate time costs. I followed this work with an investigation of climate-modulated avian shifts in southern Africa. Utilising multi-decadal citizen science occurrence data as input, I applied spatial autoregressive modelling to assess how and why community weighted precipitation and temperature

values change. In the fourth chapter, I applied an established strategy to a different field of research: I used weighted voronoi tessellations alongside extensive survey data and bioclimatic variables to better understand the complex, two-way relationship between rural communities and protected areas in Madagascar. My final chapter details the creation of a new way to rapidly analyse marine mollusc communities using supervised machine learning strategies (TIDAL). Each chapter seeks to explore new, more efficient or realistic, techniques for the study of the macroecological world. Here I discuss my major findings and elaborate on remaining unknowns and drawbacks that merit further study.

6.1 Contribution to Knowledge

6.1.1 Chapter 2: Release the HOGS

In chapter 2, I introduced and applied a new workflow for the rapid, automated extraction of occurrence information that may otherwise sit unused in historical maps and atlases. The aim was to provide a readily usable tool that could be deployed across a range of ecological disciplines for the benefit of greater understanding of ecological baselines in often understudied areas or groups.

The Historical Occurrence Georeferencing System (HOGS) uses a series of unsupervised steps to identify and isolate occurrence markers from distributional maps, assign each marker to a relevant taxonomic group, and georeference each input image relative to a known baseline coordinate system. It does not require encyclopedic coding knowledge to be applied successfully and is fast: for my use case, HOGS georeferenced over 150,000 points in approximately 20 minutes. The primary application for HOGS is the redefinition of shifting ecological baselines to more closely reflect reality. The degree to which true baseline conditions for a species or community can be discerned will be limited by the availability of data. HOGS cannot solve all data availability issues but it can help in some cases.

6.1.2 Chapter 3: Southern Africa Bird Shifts

My third chapter investigated distributional shifts in southern African avian species communities between 1987-2023. I tested the degree to which communities could and do track changes in temperature and precipitation, and the effect of anthropogenic and ecological predictor variables in explaining these shifts.

I indirectly determined the shifting constituents of communities under the influence of changing climatic conditions using the community weighted means: "Community Temperature Index" (CTI) and "Community Precipitation Index" (CPI). The implementation of temperature and precipitation community indices has been used sparingly in previous studies and never before in this region. The results showed that avian communities track the direction of both changing temperature and precipitation conditions, but that a combination of baseline conditions and environmental heterogeneity are significantly correlated with changes in CTI and CPI.

The implementation of temperature and precipitation indices together helps paint a fuller and more applicable picture of the changing dynamics of a specious and under-threat area. Accurate determination of predictors and trends give a better opportunity to intervene effectively both now and in the future.

6.1.3 Chapter 4: Madagascan Protected Areas

My fourth chapter was formulated in conjunction with Kew. I utilised a comprehensive household survey undertaken in six zones surrounding eight protected areas (PAs). The survey gathered data on a large array of socio-economic factors, from land use and ownership, to family structure, and reactions to natural perturbations and shocks to normal subsistence practices. The chapter focused on developing a better understanding of the predictors of three key socio-economic outcomes: Months of Adequate Household Food Provisioning (MAHFP), Global Person Generated Index (GPGI), and Household Poverty (HHP) derived from the Multidimensional Poverty Index (MPI). From a large pool of explanatory variables, I utilised both full

and partial Redundancy Analyses (RDA) to determine the factors that best explain variance in the socio-economic outcomes between village sites in each PA zone.

A key contribution to the analysis was my application of weighted voronoi tessellations to more realistically define areas of village resource influence in and around each PA. The use of weighted voronoi tessellations has significant advantages over unweighted tessellation and commonly utilised strategies of circular buffer zones around settlements. Circular buffers do not account for relative proximity of village sites and can therefore overlap, incorrectly suggesting a co-use scenario that is not commonly found. However, buffers can be weighted by the size of a settlement. Unweighted tessellations take relative proximity into account and thus prevent co-use scenarios. However, all sites are assumed to have the same influence. Again this is not realistic, as a village of 1,000 people will exert more pressure on the surrounding area than a settlement of 50 people. My application of weighted voronoi tessellations marries the positive aspects of the above strategies to determine areas where settlement size and relative position are accounted for.

The use of a two-pronged RDA approach, coupled with a determination of predictors into source groups: anthropogenic, ecological, and landcover, allowed for a more nuanced understanding of factors linking rural Malagasy populations with their environment. The outcome – that ecological factors dominate variance at a national scale, and burned area coverage coupled with travel times are significant at a regional scale – is new and provides exciting information for more effective intervention schemes implemented by the NGO partners in each zone.

Taken together, these results suggest that economic outcomes in PA-adjacent communities are shaped more by ecological setting and spatial infrastructure than by proximity to protected areas. Achieving socially just and ecologically effective outcomes requires moving beyond static zoning models. Embedding high-resolution, multidimensional wellbeing assessments into conservation planning offers a key pathway to reconcile biodiversity goals with poverty alleviation—not only in Madagascar, but across similarly vulnerable landscapes globally.

6.1.4 Chapter 5: Spot the Mollusc

To address the challenges of large-scale ecological data collection and annotation, we developed TIDAL, an automated machine learning pipeline capable of detecting molluscs in field images, classifying them by species, sizing them using a predefined scaling object (£1 coin), and identifying and classifying drill holes.

The implementation of the modular three-stage design - object detection, species classification, and drill hole identification - proved effective. The first stage of the pipeline, object detection of molluscs and £1 coins, achieved extremely high performance (>0.99) across all standard object detection metrics (precision, recall, and mean Average Precision (mAP) at both the lenient ($IoU \geq 0.5$) and more stringent (IoU 0.5:0.95) thresholds. Classification of mollusc species based on visual morphology was undertaken using four sample-count thresholds: 50, 100, 200, and 400 images per species. Even at the lowest threshold (50 images per species), species validation accuracy was >80% for the different models tested. As expected, validation accuracy increased with higher image thresholds (>94% for all models with 400 species images). Using a specialised computer vision model, we identified two types of drill holes, "Bevel" and "Straight", based on visual morphology. The model achieved high precision and recall for both classes, suggesting a balanced and effective modelling strategy.

Modularisation allowed each model to be optimised for a specific task and made the pipeline more interpretable and easier to adapt or extend. For instance, species classification can be retrained independently on new species subsets without needing to modify the detection or drill hole stages. Similarly, the reuse of mollusc regions of interest across stages ensures that outputs are spatially aligned and biologically meaningful.

The accurate identification of species demonstrates the significant utility of this tool in rapid identification of specimens. Applications can be found in the accelerated digitisation push of museums and other institutions. By adopting this tool, significant time, cost, and specialist expertise could be saved. Additionally, drill hole size is a well-authenticated (Palmer 1988; A. A. Klompmaker et al. 2017) proxy

for predator size, and the size of the prey itself can be directly measured. Rapid generation of datasets containing these metrics may be helpful in testing ecological concepts of predator-prey interaction and determining whether shifts have/are/will occur in community structure (Clements and Comeau 2019; Cintra-Buenrostro 2012; Leonard-Pingel and Jackson 2016; Watson et al. 2014; Parmesan 2006; Chiba and Sato 2016).

6.2 Limitations and Opportunities

6.2.1 Chapter 2: Release the HOGS

HOGS provides a workflow for the extraction and georeferencing of occurrence data from historical map-based sources. Although it is a leap forward, there are still three primary steps that could be taken to improve ease of use and detection accuracy.

HOGS utilises a colour thresholding approach to initially split background information, such as political or geographic boundaries, from occurrence marker points of interest. This introduces the need for a user-defined threshold for separation. Ideally, colour thresholding would be completed through an automatic process, wherein colour bands within an input are detected and an appropriate colour split designation is made. Although the lack of this aspect does not affect the functionality of HOGS, it does add a human step that could be given over to automation, thus further increasing the speed of the process and reducing the capacity for input error.

An additional opportunity to remove dispensable user input is the development and implementation of automated graticule or feature selection for the alignment of multiple input maps of the same area. Currently, the user is required to select graticule intersections, or notable basemap features for HOGS to then conduct reorientation of all matching inputs. Manual alignment can be completed on GIS or Python systems and can take between ~5 to 20 minutes to complete, or the same amount of time as for all automated steps of the algorithm to process the data. As with automation of colour thresholding, implementation of this would significantly reduce user time investment and diminish the chance of user introduced error.

The final limitation relates to the capacity to accurately distinguish shapes that are very complex or with high marker overlap. In its current iteration, HOGS was found to yield a point count discrepancy of 5.4% in the ten map test case comparison. Although this is an accuracy of 94.6%, improvements could be made. Many of these inaccuracies were errors of omission. Markers that extensively overlap each other (>90% overlap) were not readily splittable, even with the inflection point approach applied. An improvement to the current system would require the detection of high degree overlaps and accurate discernment of contributing initial shapes. In the test case, the imprecision of marker sizings and a wide array of potential marker shapes made this untenable. Potentially, this problem would be less pronounced and more readily solvable on maps where shape is less variable, both within and without type.

6.2.2 Chapter 3: Southern Africa Bird Shifts

Appropriate implementation of the above HOGS protocol may be useful in addressing potential limitations relating to the degradation of species occurrence data from the historical record. This was a limitation encountered in the third chapter concerning southern African shifting bird distributions. The study period encompassed a temporal range from 1987-1991 and 2019-2023, extracted from GBIF ((*GBIF* 2024b)). These periods were selected as they aligned with the Southern Africa Bird Atlas Project (SABAP) citizen science observation periods, which provided extensive sightings data. However, data were scarce before and between these intervals. Integration of pre-21st century and older historical data from a variety of sources (including but not limited to maps, journals, and museum collections) would alleviate this issue and allow for greater confidence when performing multi-decadal analyses of change in the natural world.

Utilising this, and similar, data, there are many other questions to ask of understudied zones in the Southern Hemisphere. Do communities change in similar ways across comparative latitudinal or climatic belts between continents? Or, are there additional factors intrinsic to the landscape that I have not explored,

or to the communities and their biotic and abiotic interactions with the world around them? It should also be mentioned that such an endeavour would be well served in the sensitive integration of local knowledge structures and their own historical information (Gardner, Nicoll, Birkinshaw et al. 2018; Fischer et al. 2023; Ehlers Smith et al. 2021; Kemp 2023; Steinmetz, Chutipong and Seuaturien 2006; Jessen et al. 2022; Selin 2010; Bohensky and Maru 2011; Robinson et al. 2021).

6.2.3 Chapter 4: Madagascan Protected Areas

The Madagascan household survey, facilitated by Kew and their other NGO partner organisations, is an example of respectfully integrating local views and knowledge into research for the country. One opportunity for the future is repeated surveying of the same groups over time to determine whether interventions in the NGO zones are having a positive impact and what can be done to improve. This work is being planned, and repeated measurement and engagement with communities in this way is critical to making long-term useful changes that benefit the local people and environment.

Although efforts were made to utilise high resolution data derived from Madagascan data collection, this was not possible in almost all cases. Instead, predictor variable data was extracted from global layers at coarser resolution. In addition to potentially overlooking fine-scale dynamics, the data lacked ground truth points in our focal country. This problem is widely seen across the global south (Ortiz-Bobea 2021) and has the potential to curtail fine-scale analysis of this type across broad areas of the globe.

A weighted Voronoi tessellation methodology was used to define the ZOI for each community within the study. Weighting was achieved through the counting of structures in high-resolution satellite imagery (Pbc 2024). However, building counts may be artificially inflated through the inclusion of small "farm houses"/cabins, often called lasy in Malagasy. These lasy are built close to farmers' fields and will be occupied during harvest periods if their fields are too far away from their permanent

home but they can be challenging to differentiate as they can look very similar to a typical house in satellite imagery. Future work should determine the commonality of lasy structures throughout Madagascar and whether there is a relationship between their use and agriculture suitability or landscape fragmentation.

My chapter made use of only a small subset of the survey data available. This was primarily due to the sheer volume of data. Therefore, in the future, I would like to use these data to analyse the impacts of family structure, land and resource ownership, and generational effects. These community-intrinsic factors are likely to influence the socio-economic status of each household present and are often not considered in broad-scale studies and interventions.

6.2.4 Chapter 5: Spot the Mollusc

My final content chapter seeks to address the data scarcity that undermines macroecological interpretation and understanding. I applied the TIDAL method to allow for the rapid taxonomic classification of mollusc specimens from images. The workflow can size any identified individual within an image; check for and classify any drill holes present; and measure the size of any drill hole present. TIDAL makes the task of repeat community measurement fast using collected specimens. It additionally allows for easier species identification in the absence of taxonomic experts. Despite the strengths of the protocol, there are several important limitations to acknowledge.

Currently, TIDAL is limited to only UK gastropod and bivalve specimens. This limitation is due to the time required to manually collect (from physical sites or online image repositories) specimens and then label them with the correct species identification for later use as training, validation, or testing data. I assigned over 15,000 labels, with all major species within the UK gastropod and bivalve communities included. In order to progress TIDAL further, there would need to be a significant initial input of time and expertise for the retrieval and taxonomic assignment of training data.

Drill hole detection and classification also performed well overall, but remains limited by the availability and balance of training data. In particular, more comprehensive annotation across all drill hole types is needed so that incomplete drill holes or other damage patterns, like those of serpulid worms, can be included in the categorisation. Additionally, the current model only detects drill hole presence and type and does not yet incorporate information about hole placement. Future iterations could integrate segmentation models or geometric analysis to support more detailed morphological classification.

6.3 Closing Remarks

Throughout this thesis, I have attempted to contribute novel tools and strategies for the purpose of understanding ecological dynamics and the impact of human activities on the natural world. The implementation of automated and time-rationing solutions across the fields of palaeo-, historical-, and modern-ecology is critical in the development of effective conservation strategies that are not based solely on present-day knowledge. Greater cross-field data usage could be integral in enabling us to accurately determine accurate baseline conditions for species, communities, and wider geographic areas.

Modern ecological practices have led to huge steps forward in understanding the natural world around us and have been invaluable in generating the evidence-base for protecting species and environments globally. However, both palaeo-ecology and historical-ecology have significant contributions to make that are often overlooked (McClenachan, Ferretti and Baum 2012; Bonebrake et al. 2010; Grace, Akçakaya, E. Bennett et al. 2019; Rodrigues et al. 2019; Rick and Lockwood 2013; Ojaveer et al. 2018). Palaeo-ecology can utilise a greater timespan of data than any other discipline. Although these data have their own biases with respect to distribution, preservation, and taxonomy, amongst others (Sansom 2015; Crees, Collen and Turvey 2019; Vilhena and Smith 2013; Ye 2024; Kiessling et al. 2019), omitting these data would also incur bias. Historical-ecology is sometimes viewed with scepticism and

as fragmented or selective. However, with the creation of time-saving tools and the widespread push for collection digitisation by a plethora of organisations, the availability and utility of this data source is growing. Integration of historical data is critical for many modern applications, such as species distribution models (SDMs) and baseline condition assessments. Any SDM or baseline that does not make every effort to integrate available historical information is limiting its own accuracy and utility, and potentially limiting the conservation of any species involved (Baker et al. 2016; Labay et al. 2011; Huang and Frimpong 2015).

Appendices

A

Supporting Information for Chapter 2

Contents

A.1 Additional Information 135
A.1.1 Spatial Point Pattern Test 135
A.2 Additional Tables 137

A.1 Additional Information

A.1.1 Spatial Point Pattern Test

To further investigate statistical similarity between manually-assigned and automated coordinates, the Spatial Point Pattern Test (SPPT), developed by Andresen (2016), was used with modifications to implement multinomial logistic regression by Wheeler, Steenbeek and Andresen (2018). The approach splits a study area into areal units and compares the number and location of occurrences identified with the manual and automated georeferencing protocols. Local (dis)similarities from each areal unit are aggregated to one global measure of pattern similarity. The areal units are defined by the user based on an equal area grid system, or any number of political or social boundaries, such as country borders or localised administrative locales (see below). Differences in point proportions are calculated

using Chi-square, and p-values are adjusted for multiple comparisons using the false discovery rate (Benjamini and Yekutieli 2001).

To define the areal units for this test, two spatial area segmentations were applied. We used the country boundaries of the African continent in the first areal segmentation. Within this boundary, the area of each unit is highly variable, and the boundaries of country units are arbitrary in terms of land cover change. The second areal unit coverage was generated using a hexagonal grid with 100 km spacing between hexagons based on the Icosahedral Snyder Equal Area Aperture 3 Hexagonal Grid projection (Barnes 2018). The grid covers the box extent of the African continent and results in a much finer grid structure that is of equal area across the study region.

$$S = \frac{\sum_{i=1}^{n} S_i}{n} \tag{A.1.1.1}$$

Equation S1. n = number of areas, $S_i = 1$ if spatial pattern of two datasets is considered similar for areal unit i, 0 otherwise.

We use both the "standard" and "robust" indices of similarity, which range from 0 (no similarity) to 1 (perfect similarity). Both indices are defined by the same base equation (A.1.1.1), but the "Standard" index includes all areal units (defined by the spatial area segmentations), whereas the "Robust" index considers only the number of areal units that actually have events occurring in them in either dataset. In the current usage, the "Robust" output is more likely to reflect reality, since the large number of empty cells will artificially increase the similarity value. The bivariate correlation context represents the most similar application of this spatial point test, from which coefficients ranging from 0.80 to 0.90 are considered "strong". Cohen (1988) considers lower levels of correlation to be considered "strong", but these higher values are considered to be more conservative. It should be noted that this threshold is not, and should not, be considered in a binary context such as statistical significance. Modified-SPPT was run for all ten maps considered, with both the country-boundary areal zones and the hexagonally-gridded equal-area areal zones.

The modified-SPPT test was applied to assess point pattern similarity between manually-selected occurrence points and HOGs identified occurrence points for the ten maps (A.2). Modified-SPPT was run for all ten maps considered, with both the country-boundary areal zones and the hexagonally-gridded equal-area areal zones. In all tests, using both the "Standard" and "Robust" indices, there is a 1.0 (perfect similarity) agreement in point pattern similarity.

A.2 Additional Tables

Map		Country Boundary Areal Units		Equal Area Hexagonal Units	
Group	Map Number	Standard	Robust	Standard	Robust
Passerine	94	1	1	1	1
Passerine	100	1	1	1	1
Passerine	180	1	1	1	1
Passerine	192	0.9824561	0.9333333	1	1
Passerine	273	1	1	1	1
Non-Passerine	7	1	1	1	1
Non-Passerine	28	1	1	1	1
Non-Passerine	117	1	1	1	1
Non-Passerine	178	1	1	1	1
Non-Passerine	344	1	1	1	1

Table A.1: Outcomes for the modified-SPPT tests for all ten maps. Indices of similarity range from 0 (no similarity) to 1 (perfect similarity). No compound HOGS markers are considered.

Map		Country Boundary Areal Units		Equal Area Hexagonal Units	
Atlas	Map Number	Standard	Robust	Standard	Robust
Passerine	94	1	1	1	1
Passerine	100	1	1	1	1
Passerine	180	1	1	1	1
Passerine	192	1	1	1	1
Passerine	273	1	1	1	1
Non-Passerine	7	1	1	1	1
Non-Passerine	28	1	1	1	1
Non-Passerine	117	1	1	1	1
Non-Passerine	178	1	1	1	1
Non-Passerine	344	1	1	1	1

Table A.2: Outcomes for the modified-SPPT tests for all ten maps. Indices of similarity range from 0 (no similarity) to 1 (perfect similarity). All HOGS markers are considered, including those generated from the compound marker splitting protocol.

Processing Step	CPU Time
Marker-background separation	total: 55 s
Identify shapes and centres	total: 7.32 s
Find compound shapes	total: 55.7 ms
Discriminate single and compound shapes	total: 470 ms
Split and define compound shape components	total: 10min 3s
Total	∼11min 6s

Table A.3: Processing times for each protocol stage.

B

Supporting Information for Chapter 3

Contents

B.1 Additional Tables 139
B.2 Additional Plots 140

B.1 Additional Tables

Term	Coefficient	Std. Error	z-value	p-value
Intercept	+0.19	0.02	9.06	<0.05
Population Density Change	-0.02	0.01	-1.75	0.08
Mean Temperature Change	-0.01	0.01	-1.34	0.18
Environmental Heterogeneity	-0.03	0.01	-1.34	<0.05

Table B.1: Model terms from the SAR ΔCTI weighted lag model.

Term	Effect Size	Std. Error	z-value	p-value
Intercept	+1.60	0.16	9.87	<0.05
Baseline CPI	-0.68	0.09	-7.21	<0.05

Table B.2: Model terms from the SAR ΔCPI weighted lag model

B.2 Additional Plots

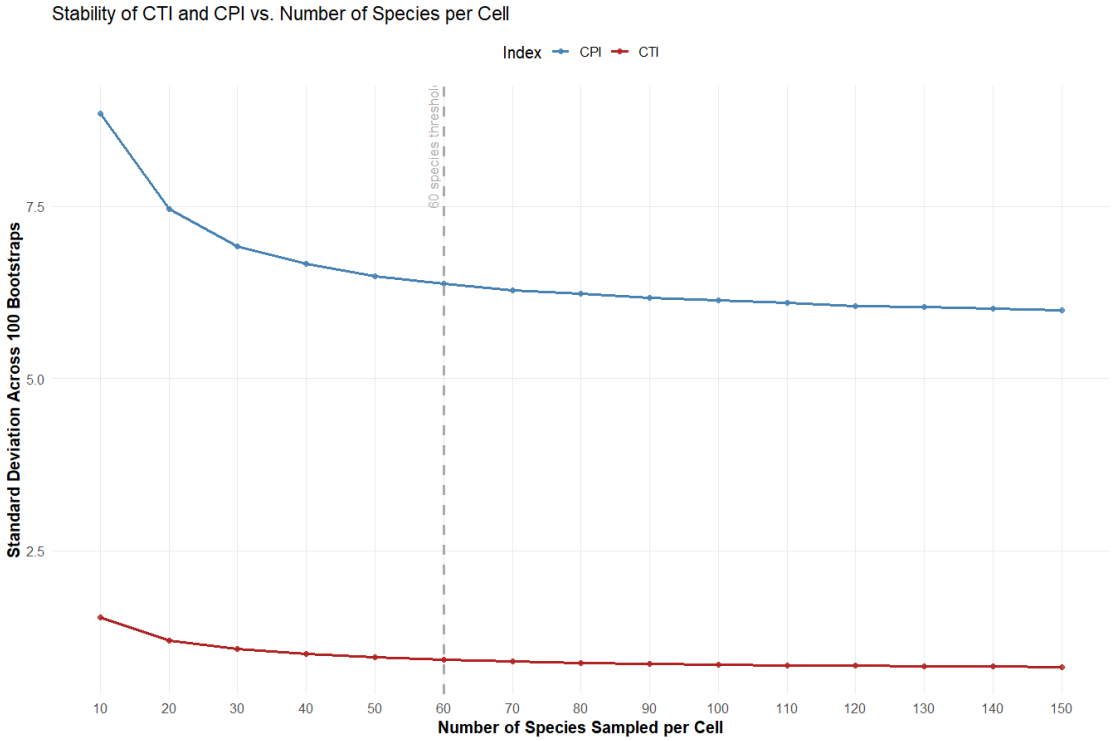


Figure B.1: The number of unique species per grid cell and the standard deviation from CTI and CPI calculation across 100 subsampling iterations.

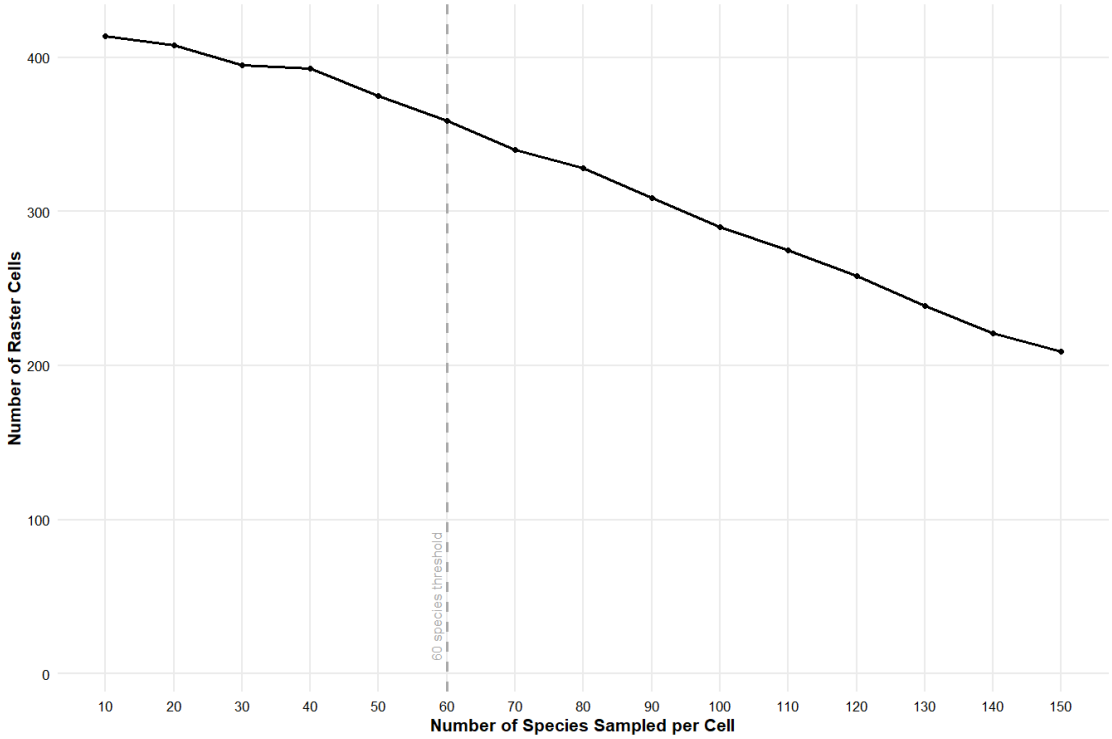


Figure B.2: Number of valid raster cells within the study region as determined by minimum species sampling requirement per cell. The grey dashed line indicates our optimal tradeoff between result stability and proportion of data cells that can be included in analyses.

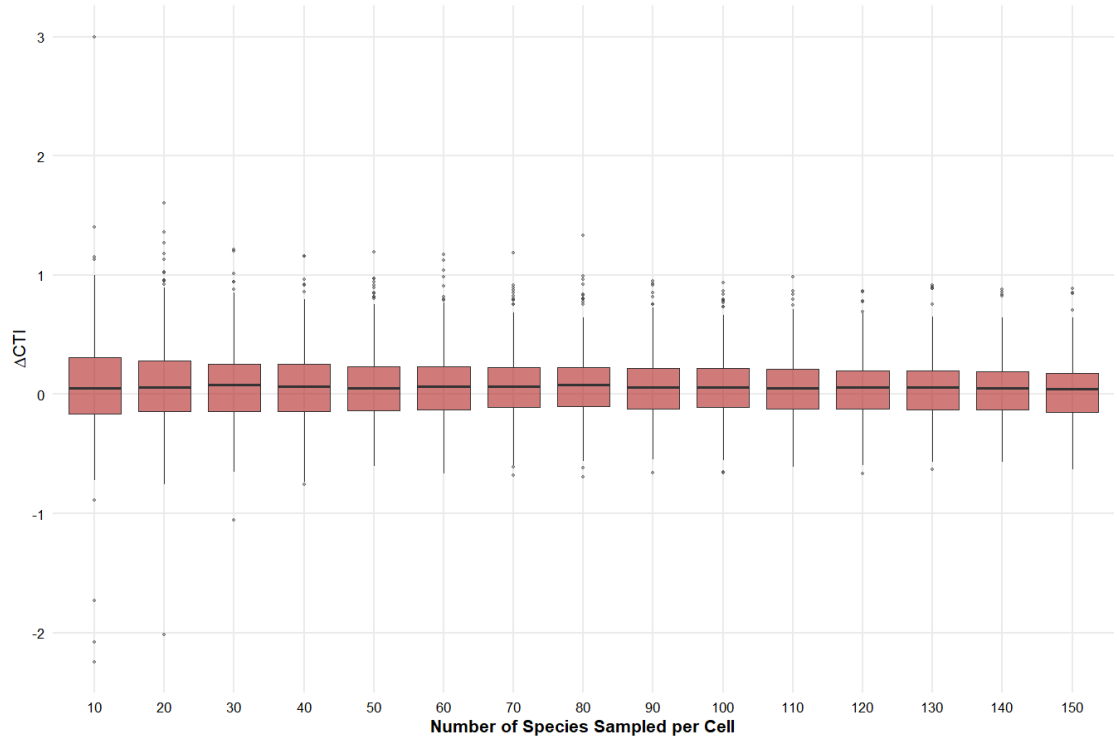


Figure B.3: Boxplots showing the distribution of ΔCTI across raster cells for each species-per-cell threshold. Thresholds represent the minimum number of species required in a cell to be included in the analysis (from 10 to 150 in increments of 10). As the threshold increases, the number of cells decreases, but the consistency of ΔCTI values may improve by reducing noise from sparsely sampled cells.

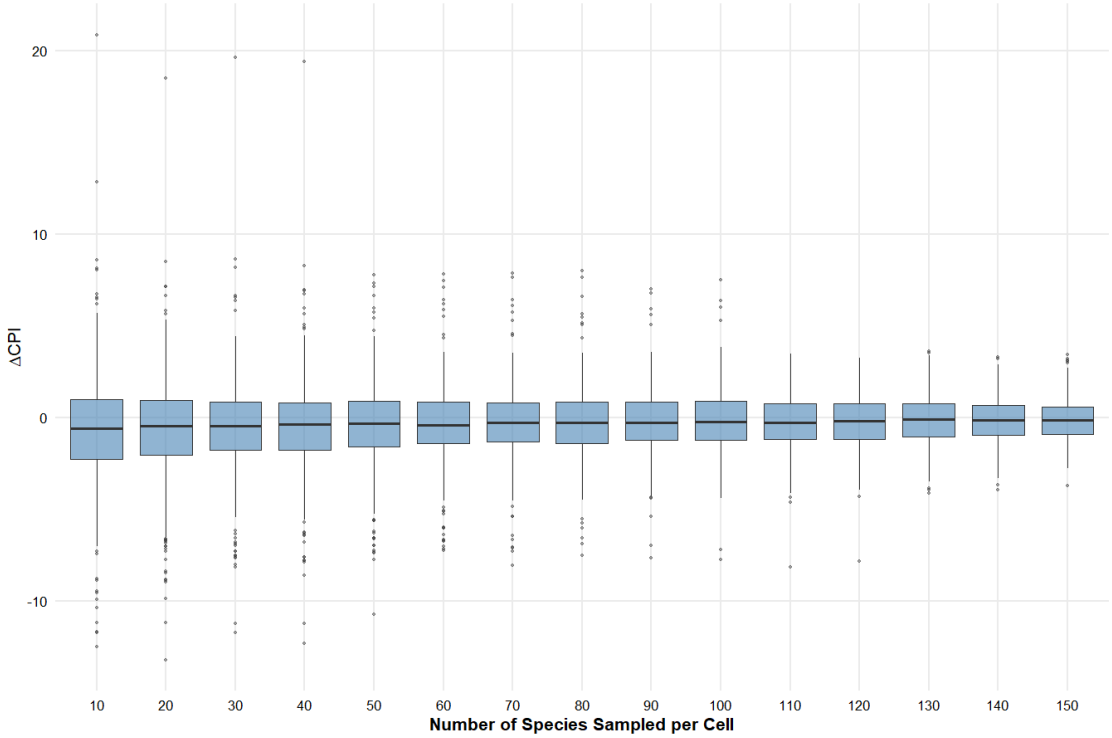


Figure B.4: Boxplots showing the distribution of ΔCPI across raster cells for each species-per-cell threshold. Thresholds represent the minimum number of species required in a cell to be included in the analysis (from 10 to 150 in increments of 10). As the threshold increases, the number of cells decreases, but the consistency of ΔCPI values may improve by reducing noise from sparsely sampled cells.

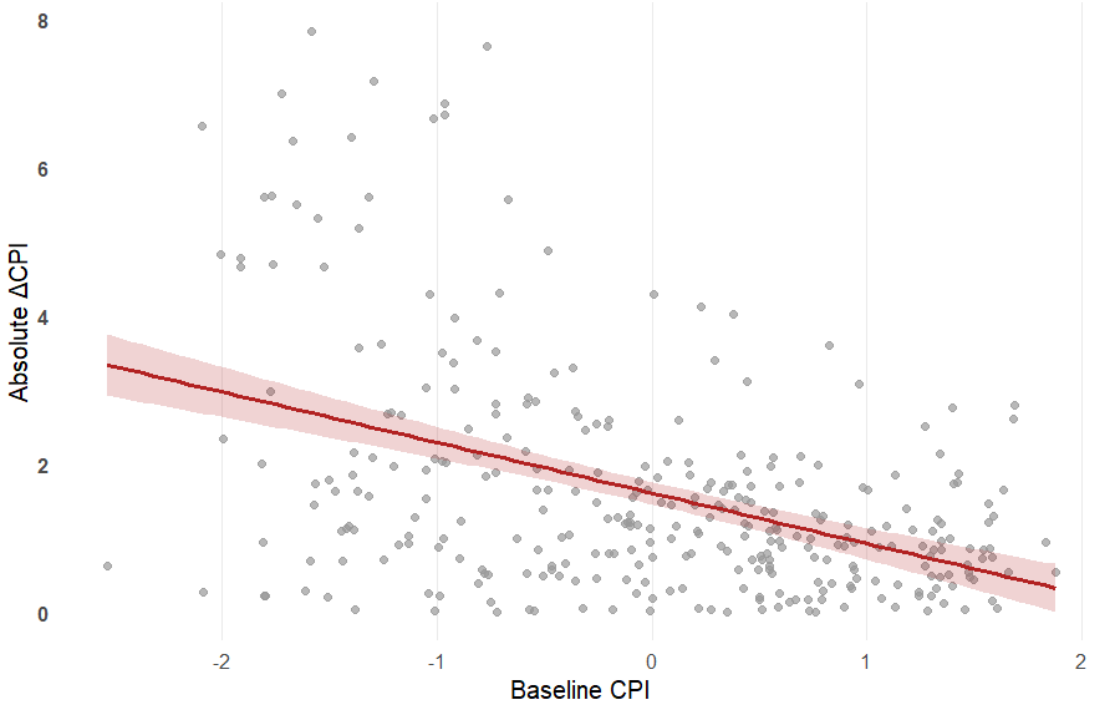


Figure B.5: The inverse relationship between absolute values of ΔCPI and baseline CPI values.

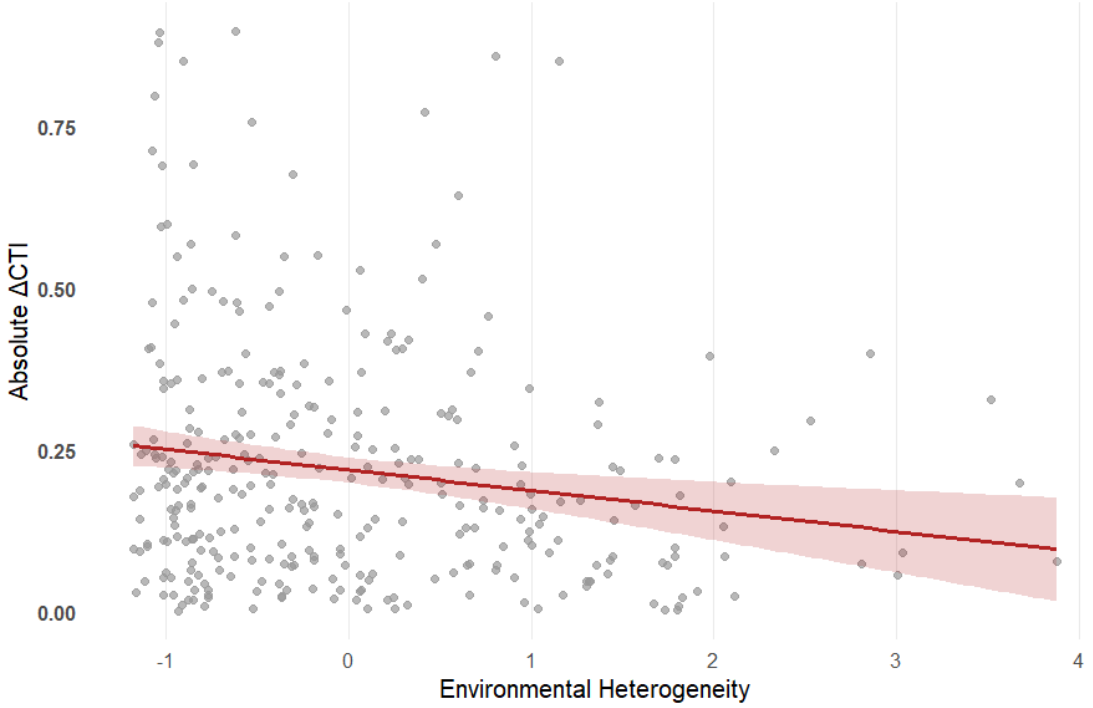


Figure B.6: The inverse relationship between absolute values of ΔCTI and environmental heterogeneity.

C

Supporting Information for Chapter 4

Contents

C.1 Additional Information 145
C.1.1 Characterising Landscape Occupancy 145
C.2 Additional Tables 146

C.1 Additional Information

C.1.1 Characterising Landscape Occupancy

Approximately 13,000 buildings were identified from satellite imagery across all site areas of interest, with an average of ~58 in each of the 227 settlements. There was a range of 468 between the lowest and highest buildings counts. The mean voronoi-defined polygon area was $25.4km^2$ and the range extended between $0.2km^2$ and $420.0km^2$. Site-specific indices of landscape occupancy are summarised in table S2 (C.1. The TPF site contained the most villages (56) and DU-APCLFA the least (24). MBG had a far greater average point count of ~105.8, compared to all other sites having means per settlement below 52.

In the 93 survey villages considered, ~6,500 buildings were identified, with a mean of 70.2 for each settlement. Polygon areas also varied considerably, between

$0.7km^2$ and $206.6km^2$. Furthermore, the mean polygon area for surveyed villages ($17.5km^2$) was lower than the whole dataset figure ($25.4km^2$). Similar to the whole settlement dataset, MBG building count of ~142 was approximately double the average per settlement (C.2).

C.2 Additional Tables

Site	Settlement Count	Mean Building Count (1 d.p.)	Mean Polygon Area (km^2) (1 d.p.)
DU-AMPA	28	36.4	39.9
DU-APCLFA	24	37.7	50.5
KEW	25	37.9	54.6
MBG	53	105.8	10.1
MV	41	45.3	6.1
TPF	56	51.9	22.9

Table C.1: Details of all identified settlement occupancy metrics.

Site	Village Count	Mean Building Count (1 d.p.)	Mean Polygon Area (km^2) (1 d.p.)
DU-AMPA	9	53.2	49.3
DU-APCLFA	18	16.3	25.5
KEW	9	39.4	23.1
MBG	32	142.3	12.2
MV	13	21.5	3.9
TPF	12	46.8	6.5

Table C.2: Details of survey village occupancy metrics.

Data Layer	Predictor Class	Layer Type	Data Source	Timeframe	Spatial resolution
Average Annual Temperature	Ecological	Raster	CHELSA ("CHELSA" 2020)	2020	30"
Average Annual Precipitation	Ecological	Raster	CHELSA ("CHELSA" 2020)	2020	30"
Topographic Roughness	Ecological	Raster	SRTM (OpenTopography 2013)	2013	1"
Crop Suitability	Ecological	Raster	Global Agro-Ecological Zoning (GAEZ) (FAO and IIASA 2024)	2020	9 km
Length of Growing Period	Ecological	Raster	(FAO and IIASA 2024)	2020	9 km
Population Density	Anthropogenic	Raster	WorldPop (WorldPop 2020)	2020	30"
Tobler Costpath	Anthropogenic	Line	Derived	2024	-
Voronoi Polygon Area	Anthropogenic	Polygon	Derived	2024	-
Travel Time: Settlement	Anthropogenic	Integer	Derived	2024	-
Travel Time: Town	Anthropogenic	Integer	Derived	2024	-
Travel Time: Port	Anthropogenic	Integer	Derived	2024	-
Tree Cover	Landcover	Raster	Dynamic World V1 and Sentinel-2 base layer (Brown et al. 2022)	2023	10 m
Grass Cover	Landcover	Raster		2023	10 m
Crops Cover	Landcover	Raster		2023	10 m
Shrub/Scrub Cover	Landcover	Raster		2023	10 m
Tree Cover Change	Landcover	Raster	Dynamic World V1 and Sentinel-2 base layer (Brown et al. 2022)	2016 - 2023	10 m
Grass Cover Change	Landcover	Raster		2016 - 2023	10 m
Crops Cover Change	Landcover	Raster		2016 - 2023	10 m
Shrub/Scrub Cover Change	Landcover	Raster		2016 - 2023	10 m
Burned Area	Landcover	Raster	(Fernández-García, Franquesa, and Kull 2024)	2014 - 2021	20 m

Table C.3: Predictor variables for RDA analysis.

D

Supporting Information for Chapter 5

Contents

D.1 Additional Information 148
D.1.1 Software and Hardware 148
D.2 Additional Plots 149

D.1 Additional Information

D.1.1 Software and Hardware

All experiments were conducted using Python 3.10 within Google Colab, leveraging GPU acceleration. Primary libraries included PyTorch (via Ultralytics' YOLOv8 (Jocher, Chaurasia and Qiu 2023), TensorFlow/Keras (Martín Abadi et al. 2015; Chollet et al. 2015) for CNN model training and OpenCV (Bradski 2000) for image processing. Scripts and experimental logs were managed and version-controlled using Google Drive for reproducibility and ease of access.

D.2 Additional Plots

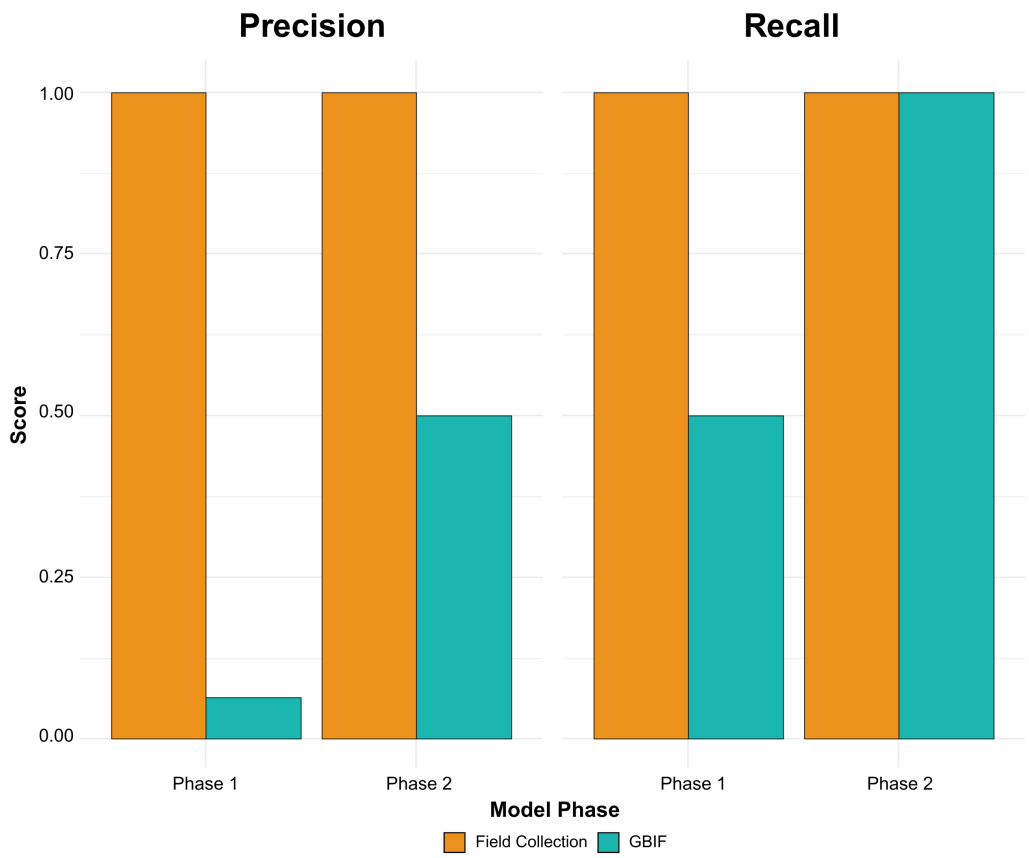


Figure D.1: Precision and recall scores for pound coin detection. Progression of scores between phase 1 and phase 2 of model identification. TLIS dataset is not included as it does not contain any labelled £1 coin features.

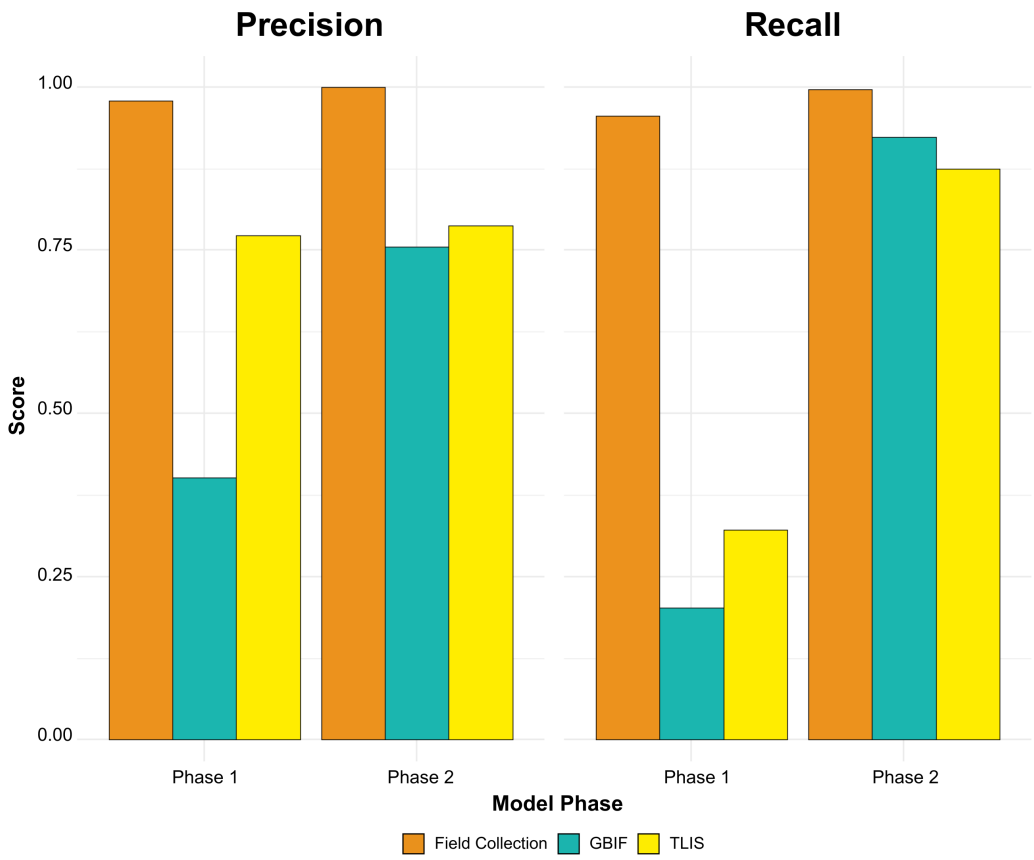


Figure D.2: Precision and recall scores for mollusc detection. Progression of scores between phase 1 and phase 2 of model identification.

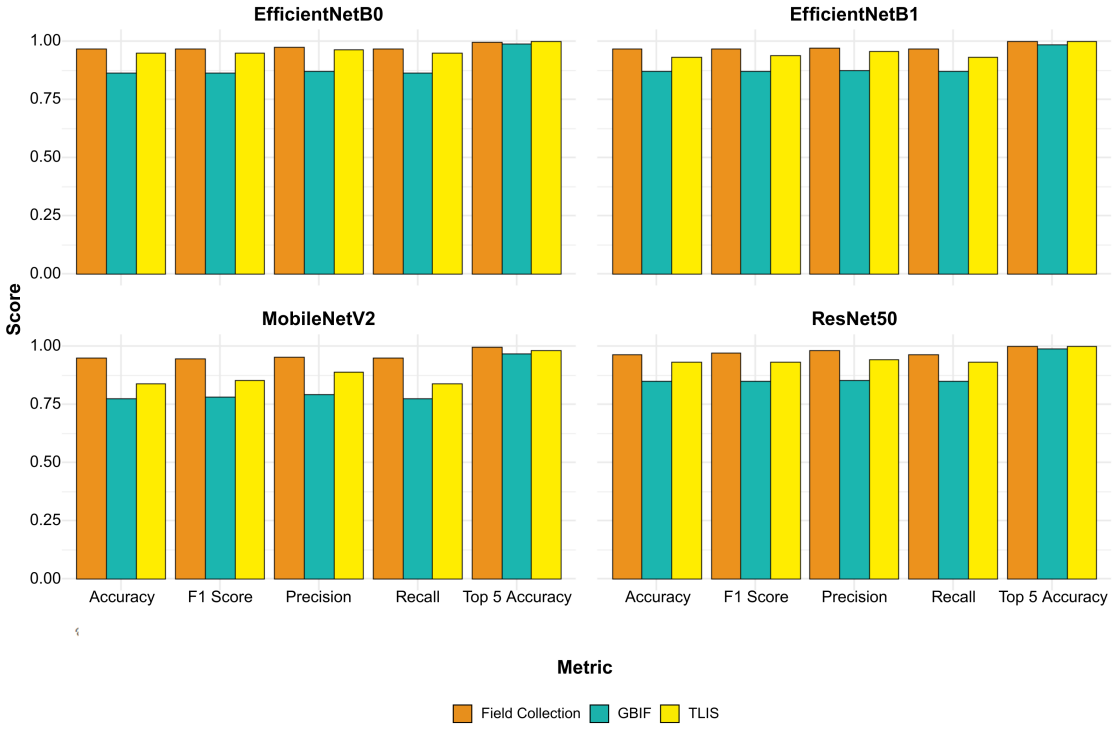


Figure D.3: Model evaluation scores across the four models tested. Image threshold is held constant at 200.

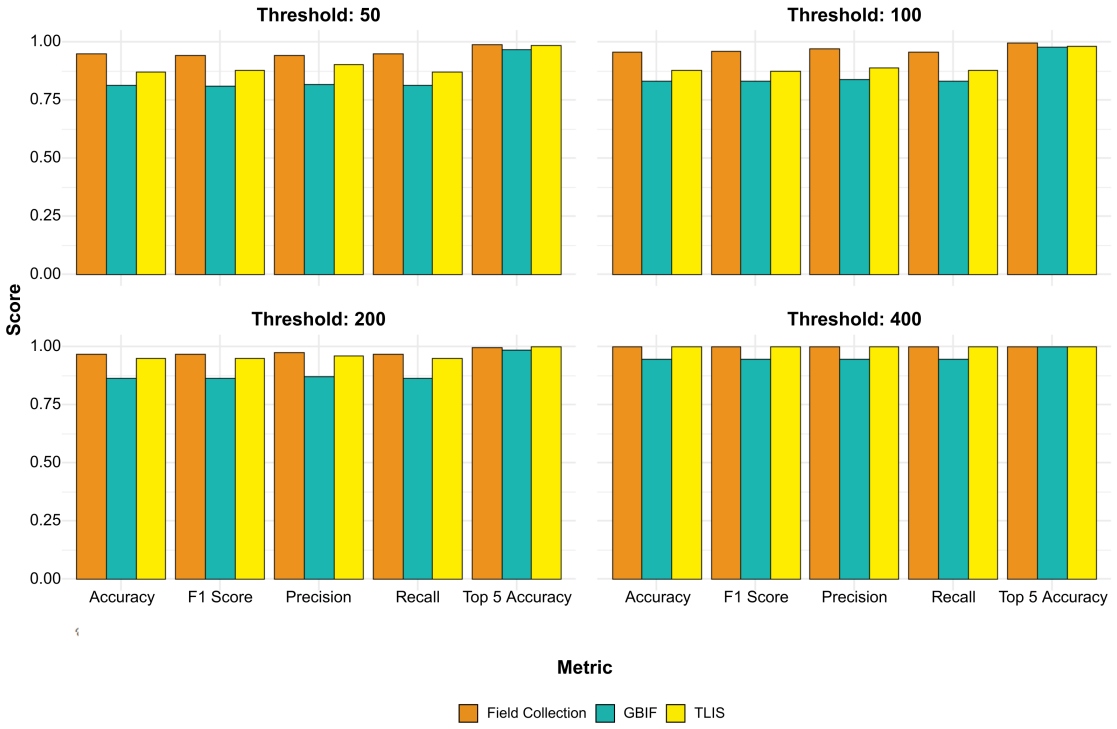


Figure D.4: Model evaluation scores for the EfficientNetB0 (Tan and Le 2019) model at each threshold.

References

Ahyong, S et al. (2024). *World Register of Marine Species (WoRMS).* =https://www.marinespecies.org.

Akiba, Takuya et al. (2019). 'Optuna: A Next-generation Hyperparameter Optimization Framework'. In: *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.*

Alomar, Khaled, Halil Ibrahim Aysel and Xiaohao Cai (Feb. 2023). 'Data augmentation in classification and segmentation: A survey and new strategies'. en. In: *J. Imaging* 9.2, p. 46.

Andam, Kwaw S et al. (June 2010). 'Protected areas reduced poverty in Costa Rica and Thailand'. In: *Proc. Natl. Acad. Sci. U. S. A.* 107.22, pp. 9996–10001.

Andermann, Tobias et al. (2020). 'The past and future human impact on mammalian diversity'. In: *Sci. Adv.* 6.36.

Andresen, Martin A (Jan. 2016). 'An area-based nonparametric spatial point pattern test: The test, its applications, and the future'. In: *Method. Innov.* 9, p. 205979911663065.

Andriatsitohaina, R Ntsiva N et al. (Jan. 2020). 'Participatory Bayesian network modeling to understand driving factors of land-use change decisions: insights from two case studies in northeast Madagascar'. In: *J. Land Use Sci.* 15.1, pp. 69–90.

Antão, Laura H et al. (July 2020). 'Temperature-related biodiversity change across temperate marine and terrestrial systems'. In: *Nat. Ecol. Evol.* 4.7, pp. 927–933.

Arafeh-Dalmau, Nur et al. (Oct. 2021). 'Incorporating climate velocity into the design of climate-smart networks of marine protected areas'. In: *Methods Ecol. Evol.* 12.10, pp. 1969–1983.

ArcGIS Pro (2024). `https://www.esri.com/en-us/arcgis/products/arcgis-pro/overview`. Accessed: 2024-11-5.

Ariño, Arturo H (Oct. 2010). 'Approaches to estimating the universe of natural history collections data'. In: *Biodivers. Inf.* 7.2.

Auer, Sonya K and David I King (Aug. 2014). 'Ecological and life-history traits explain recent boundary shifts in elevation and latitude of western North American songbirds: Range boundary shifts in songbirds'. In: *Glob. Ecol. Biogeogr.* 23.8, pp. 867–875.

Bahn, Volker, Raymond O'Connor and William Krohn (Dec. 2006). 'Importance of spatial autocorrelation in modeling bird distributions at a continental scale: Autocorrelation in bird distributions'. In: *Ecography (Cop.)* 29.6, pp. 835–844.

Bai, Z G and D L Dent (2007). *ISRIC Report 2007/03, (Glada report 1a): Land degradation and improvement in South Africa. 1. Identification by remote sensing.* `https://www.isric.org/documents/document-type/isric-report-200703-glada-report-1a-land-degradation-and-improvement-south`. Accessed: 2025-5-9.

Baker, David J et al. (2016). 'Choice of baseline climate data impacts projected species' responses to climate change'. In: *Glob. Chang. Biol.* 22.7, pp. 2392–2404.

Barnes, Richard (2018). *r-barnes/dggridR: v2.0.3*.

Bateman, Brooke L et al. (Mar. 2016). 'The pace of past climate change vs. potential bird distributions and land use in the United States'. In: *Glob. Chang. Biol.* 22.3, pp. 1130–1144.

Battista, Francesco and Cesar L Schultz (Nov. 2024). 'Sampling and collector biases as taphonomic filters: an overview'. en. In: *An. Acad. Bras. Cienc.* 96.4, e20231242.

Beck, Jan et al. (Jan. 2014). 'Spatial bias in the GBIF database and its effect on modeling species' geographic distributions'. In: *Ecol. Inform.* 19, pp. 10–15.

Beijbom, O et al. (June 2012). 'Automated annotation of coral reef survey images'. en. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 1170–1177.

Bengio, Yoshua et al. (June 2009). 'Curriculum learning'. In: *Proceedings of the 26th Annual International Conference on Machine Learning*. New York, NY, USA: ACM.

Benjamini, Yoav and Daniel Yekutieli (Aug. 2001). 'The control of the false discovery rate in multiple testing under dependency'. In: *Ann. Stat.* 29.4, pp. 1165–1188.

Berger, V J and A D Kharazova (1997). 'Mechanisms of salinity adaptations in marine molluscs'. In: *Interactions and Adaptation Strategies of Marine Organisms*. Dordrecht: Springer Netherlands, pp. 115–126.

Bezner Kerr, Rachel et al. (June 2021). 'Can agroecology improve food security and nutrition? A review'. In: *Glob. Food Sec.* 29.100540, p. 100540.

Bianchi, Filipe Michels and Leonardo Tresoldi Gonçalves (Apr. 2021). 'Getting science priorities straight: how to increase the reliability of specimen identification?' In: *Biol. Lett.* 17.4, p. 20200874.

Bidaud, Cecile et al. (2016). *Household survey investigating the social impact of biodiversity offset: a case study from Madagascar*.

Bilinsky, P and A Swindale (2010). *Months of Adequate Household Food Provisioning (MAHFP) for Measurement of Household Food Access: Indicator Guide (Version 4)*. USAID.

BirdLife (2022). *BirdLife international and handbook of the birds of the world (2022) bird species distribution maps of the world*. `http://datazone.birdlife.org/species/requestdis.`. Accessed: 2023-11-22.

Bloom, Trevor D S, Aquila Flower and Eric G DeChaine (Jan. 2018). 'Why georeferencing matters: Introducing a practical protocol to prepare species occurrence records for spatial analysis'. In: *Ecol. Evol.* 8.1, pp. 765–777.

Blumstein, Daniel T (Feb. 2006). 'Developing an evolutionary ecology of fear: how life history and natural history traits affect disturbance tolerance in birds'. In: *Anim. Behav.* 71.2, pp. 389–399.

'State of the Climate in 2019' (Aug. 2020). In: *Bull. Am. Meteorol. Soc.* 101.8. Ed. by J Blunden and D S Arndt, S1–S429.

Boakes, Elizabeth H et al. (June 2010). 'Distorted views of biodiversity: spatial and temporal bias in species occurrence data'. In: *PLoS Biol.* 8.6, e1000385.

Bohensky, Erin L and Yiheyis Maru (2011). 'What have we learned from a decade of international literature on ?Integration??' In: *Ecol. Soc.* 16.4.

Bonebrake, Timothy C et al. (Dec. 2010). 'Population decline assessment, historical baselines, and conservation'. In: *Conserv. Lett.* 3.6, pp. 371–378.

Bonnevie, Bo T (Aug. 2011). 'Some considerations when comparing SABAP 1 with SABAP 2 data'. In: *Ostrich* 82.2, pp. 161–162.

Both, C and M E Visser (May 2001). 'Adjustment to climate change is constrained by arrival date in a long-distance migrant bird'. en. In: *Nature* 411.6835, pp. 296–298.

Bowler, Diana and Katrin Böhning-Gaese (Sept. 2017). 'Improving the community-temperature index as a climate change indicator'. In: *PLoS One* 12.9, e0184275.

Bradshaw, Corey J A, Navjot S Sodhi and Barry W Brook (2009). 'Tropical turmoil: a biodiversity tragedy in progress'. In: *Front. Ecol. Environ.* 7.2, pp. 79–87.

Bradski, G (2000). 'The OpenCV Library'. In: *Dr. Dobb's J. Softw. Tools Prof. Program.*

Brockington, Daniel and David Wilkie (Nov. 2015). 'Protected areas and poverty'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 370.1681, p. 20140271.

Broekhuysen, G J (Jan. 1940). 'A preliminary investigation of the importance of desiccation, temperature, and salinity as factors controlling the vertical distribution of certain intertidal marine gastropods in False Bay, South Africa'. In: *Trans. R. Soc. S. Afr.* 28.3, pp. 255–292.

Brooks, Michael et al. (Nov. 2022). 'The African Bird Atlas Project: a description of the project and BirdMap data-collection protocol'. In: *Ostrich*, pp. 1–10.

Brousmiche, Delphine et al. (Apr. 2021). 'How can we analyze environmental health resilience and vulnerability? A joint analysis with composite indices applied to the north of France'. In: *Sci. Total Environ.* 763.142983, p. 142983.

Brown, Christopher F et al. (June 2022). 'Dynamic World, Near real-time global 10 m land use land cover mapping'. In: *Scientific Data* 9.1, pp. 1–17.

Buchan, Claire et al. (Sept. 2022). 'Spatially explicit risk mapping reveals direct anthropogenic impacts on migratory birds'. In: *Glob. Ecol. Biogeogr.* 31.9, pp. 1707–1725.

Burger, Martijn et al. (2020). *Urban-Rural Happiness Differentials Across the World*. `https://worldhappiness.report/ed/2020/urban-rural-happiness-differentials-across-the-world/`. Accessed: 2025-6-25.

Burgess, Neil et al. (Jan. 2000). 'A preliminary assessment of congruence between biodiversity patterns in Afrotropical forest birds and forest mammals'. In: *Ostrich* 71, pp. 286–290.

Burt, James E et al. (Jan. 2020). 'Automated and semi-automated map georeferencing'. In: *Cartogr. Geogr. Inf. Sci.* 47.1, pp. 46–66.

Buschbaum, Christian et al. (Aug. 2024). 'Climate change impacts on a sedimentary coast—a regional synthesis from genes to ecosystems'. In: *Mar. Biodivers.* 54.4, pp. 1–33.

Bush, Mark B (Feb. 2023). 'Bridging the divide from paleoecology to neoecology'. In: *Bull. Fla. Mus. Nat. Hist.* 60.2, pp. 57–57.

Callomon, Paul (Feb. 2019). 'Standard views for imaging mollusk shells'. In: *ResearchGate*.

Campello, Ricardo J G B, Davoud Moulavi and Joerg Sander (2013). 'Density-based clustering based on hierarchical density estimates'. In: *Advances in Knowledge Discovery and Data Mining*. Lecture notes in computer science. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 160–172.

Cao, Yong, Dudley D Williams and Nancy E Williams (Nov. 1998). 'How important are rare species in aquatic community ecology and bioassessment?' en. In: *Limnol. Oceanogr.* 43.7, pp. 1403–1409.

Chamberlain, Scott, Vijay Barve et al. (2024). *rgbif: Interface to the Global Biodiversity Information Facility API*.

Chamberlain, Scott and Eduard Szocs (2013). *taxize - taxonomic search and retrieval in R.*

Chandu, Khyathi Raghavi et al. (July 2024). 'Certainly uncertain: A benchmark and metric for multimodal epistemic and aleatoric awareness'. In: *arXiv [cs.AI]*.

Chattopadhyay, Devapriya (2009). 'Predation in molluscs: A multi-taxon approach using neontological and paleontological data'. PhD thesis. University of Michigan.

Chattopadhyay, Devapriya and Tomasz K Baumiller (Nov. 2007). 'Drilling under threat: An experimental assessment of the drilling behavior of Nucella lamellosa in the presence of a predator'. en. In: *J. Exp. Mar. Bio. Ecol.* 352.1, pp. 257–266.

Chattopadhyay, Devapriya and Saurav Dutta (Mar. 2013). 'Prey selection by drilling predators: A case study from Miocene of Kutch, India'. In: *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 374, pp. 187–196.

Chattopadhyay, Devapriya, Martin Zuschin and Adam Tomašových (Dec. 2015). 'How effective are ecological traits against drilling predation? Insights from recent bivalve assemblages of the northern Red Sea'. In: *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 440, pp. 659–670.

CHELSA (July 2020). `https://chelsa-climate.org/downloads/`. Accessed: 2024-8-7.

Chen, I-Ching et al. (Aug. 2011). 'Rapid range shifts of species associated with high levels of climate warming'. In: *Science* 333.6045, pp. 1024–1026.

Chesneau, Elisabeth (Feb. 2011). 'A model for the automatic improvement of colour contrasts in maps: application to risk maps'. In: *Geogr. Inf. Syst.* 25.1, pp. 89–111.

Chiba, Tomoki and Shin'ichi Sato (May 2012). 'Size-selective predation and drillhole-site selectivity in Euspira fortunei (Gastropoda: Naticidae): implications for ecological and palaeoecological studies'. In: *J. Molluscan Stud.* 78.2, pp. 205–212.

— (May 2016). 'Climate-mediated changes in predator–prey interactions in the fossil record: a case study using shell-drilling gastropods from the Pleistocene Japan Sea'. In: *Paleobiology* 42.2, pp. 257–268.

Chivers, William J, Anthony W Walne and Graeme C Hays (Feb. 2017). 'Mismatch between marine plankton range movements and the velocity of climate change'. In: *Nat. Commun.* 8.1, p. 14434.

Chollet, François et al. (2015). *Keras*. `https://keras.io`.

Cintra-Buenrostro, C (2012). 'Would predatory drillhole frequency on Chione spp. Increase under the suggested climate change scenario? Comparing Pleistocene and modern rhodolith beds'. In: *International Journal of Modern Botany* 2, pp. 103–107.

Clavero, M et al. (Apr. 2023). 'Where wolves were: setting historical baselines for wolf recovery in Spain'. In: *Anim. Conserv.* 26.2, pp. 239–249.

Clements and Comeau (Dec. 2019). 'Behavioral Defenses of Shellfish Prey under Ocean Acidification'. In: *shre* 38.3, pp. 725–742.

Clements, P C Rasmussen et al. (2023). 'The eBird/Clements checklist of Birds of the World'. In.

Coetzee, Bernard W T et al. (Nov. 2009). 'Ensemble models predict Important Bird Areas in southern Africa will become less effective for conserving endemic birds under climate change'. In: *Glob. Ecol. Biogeogr.* 18.6, pp. 701–710.

Cohen, Jacob (1988). *Statistical power analysis for the behavioral sciences*. 2nd ed. London, England: Routledge.

Coleman, D W, M Byrne and A R Davis (Aug. 2014). 'Molluscs on acid: gastropod shell repair and strength in acidifying oceans'. en. In: *Mar. Ecol. Prog. Ser.* 509, pp. 203–211.

Collins, Amy C, Monika Böhm and Ben Collen (Feb. 2020). 'Choice of baseline affects historical population trends in hunted mammals of North America'. In: *Biol. Conserv.* 242.108421, p. 108421.

Cooke, A et al. (2022). 'The economic contribution of Madagascar's protected areas – a review of the evidence'. In.

Crees, Jennifer J, Ben Collen and Samuel T Turvey (Dec. 2019). 'Bias, incompleteness and the 'known unknowns' in the Holocene faunal record'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 374.1788, p. 20190216.

Curley, Shannon R et al. (Apr. 2022). 'Evaluating compositional changes in the avian communities of eastern North America using temperature and precipitation indices'. In: *J. Biogeogr.* 49.4, pp. 739–752.

Dale, Mark R T and Marie-Josée Fortin (Jan. 2002). 'Spatial autocorrelation and statistical tests in ecology'. In: *Écoscience* 9.2, pp. 162–167.

Davey, Catherine M et al. (May 2013). 'Impact of climate change on communities: revealing species' contribution'. In: *J. Anim. Ecol.* 82.3, pp. 551–561.

De Frenne, Pieter et al. (Nov. 2013). 'Microclimate moderates plant responses to macroclimate warming'. In: *Proc. Natl. Acad. Sci. U. S. A.* 110.46, pp. 18561–18565.

Dee, Laura E et al. (Aug. 2019). 'When do ecosystem services depend on rare species?' en. In: *Trends Ecol. Evol.* 34.8, pp. 746–758.

Deng, Jia et al. (June 2009). 'ImageNet: A large-scale hierarchical image database'. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 248–255.

Devictor, Vincent, Romain Julliard et al. (Dec. 2008). 'Birds are tracking climate warming, but not fast enough'. In: *Proc. Biol. Sci.* 275.1652, pp. 2743–2748.

Devictor, Vincent, Chris van Swaay et al. (Jan. 2012). 'Differences in the climatic debts of birds and butterflies at a continental scale'. In: *Nature Climate Change* 2.2, pp. 121–124.

Diamond, A W and A C Hamilton (1980). 'The distribution of forest passerine birds and Quaternary climatic change in tropical Africa'. In: *J. Zool. (1987)* 191.3, pp. 379–402.

Diez, Yander L et al. (Sept. 2023). *Drilling predation on bivalve shell assemblages at Playa Guardalavaca, Cuba.*

Dionisio, Abigaile et al. (Jan. 2020). 'A Radiolarian Classifier using Convolutional Neural Networks'. In: *2020 International Conference on Artificial Intelligence and Signal Processing (AISP)*. IEEE, pp. 1–5.

Dong, Pinliang (Apr. 2008). 'Generating and updating multiplicatively weighted Voronoi diagrams for point, line and polygon features in GIS'. In: *Comput. Geosci.* 34.4, pp. 411–421.

Dormann, Carsten F et al. (Jan. 2013). 'Collinearity: a review of methods to deal with it and a simulation study evaluating their performance'. en. In: *Ecography (Cop.)* 36.1, pp. 27–46.

Douglas, David H and Thomas K Peucker (1973). 'ALGORITHMS FOR THE REDUCTION OF THE NUMBER OF POINTS REQUIRED TO REPRESENT A DIGITIZED LINE OR ITS CARICATURE'. In: *Cartographica* 10.2, pp. 112–122.

Duchoslav, Jan et al. (2020). *Building resilience and adaptation to climate change in Malawi: Quantitative baseline report.* Tech. rep. Washington, DC.

Dwyer, Nelson and Hansen (2024). *Roboflow.* `https://roboflow.com.`. Accessed: 2024-11-18.

Ehlers Smith, Yvette C et al. (Apr. 2021). 'Indigenous knowledge of South African bird and rangeland ecology is effective for informing conservation science'. In: *J. Environ. Manage.* 284, p. 112041.

Eklund, Johanna, F Guillaume Blanchet et al. (Nov. 2016). 'Contrasting spatial and temporal trends of protected area effectiveness in mitigating deforestation in Madagascar'. In: *Biol. Conserv.* 203, pp. 290–297.

Eklund, Johanna, Lauren Coad et al. (Oct. 2019). 'What constitutes a useful measure of protected area effectiveness? A case study of management inputs and protected area impacts in Madagascar'. In: *Conserv. Sci. Pract.* 1.10.

Engelbrecht, Francois et al. (Aug. 2015). 'Projections of rapidly rising surface temperatures over Africa under low mitigation'. In: *Environ. Res. Lett.* 10.8, p. 085004.

Engelen, Jesper E van and Holger H Hoos (Feb. 2020). 'A survey on semi-supervised learning'. en. In: *Mach. Learn.* 109.2, pp. 373–440.

Enquist, C A F (Apr. 2002). 'Predicted regional impacts of climate change on the geographical distribution and diversity of tropical forests in Costa Rica'. In: *J. Biogeogr.* 29.4, pp. 519–534.

Erp, Marieke van et al. (Aug. 2015). 'Georeferencing animal specimen datasets: Georeferencing animal specimen datasets'. In: *Trans. GIS* 19.4, pp. 563–581.

FAO and IIASA (2024). *Global Agro Ecological Zones version 4 (GAEZ v4).* `http://www.fao.org/gaez/`. Accessed: 2024–.

Fernández-García, V, M Franquesa and C A Kull (Mar. 2024). 'Madagascar's burned area from Sentinel-2 imagery (2016-2022): Four times higher than from lower resolution sensors'. In: *Sci. Total Environ.* 914, p. 169929.

Ferraro, Paul J (Dec. 2002). 'The local costs of establishing protected areas in low-income nations: Ranomafana National Park, Madagascar'. In: *Ecol. Econ.* 43.2-3, pp. 261–275.

Ferraro, Paul J, Merlin M Hanauer and Katharine R E Sims (Aug. 2011). 'Conditions associated with protected area success in conservation and poverty reduction'. In: *Proc. Natl. Acad. Sci. U. S. A.* 108.34, pp. 13913–13918.

Fischer, Harry W et al. (Nov. 2023). 'Community forest governance and synergies among carbon, biodiversity and livelihoods'. In: *Nat. Clim. Chang.*, pp. 1–8.

Forest Carbon Partnership Facility (2020). 'ER Monitoring Report (ER-MR)'. In.

Fredston-Hermann, Alexa, Steven D Gaines and Benjamin S Halpern (Oct. 2018). 'Biogeographic constraints to marine conservation in a changing climate: Biogeography informs MPAs under climate change'. In: *Ann. N. Y. Acad. Sci.* 1429.1, pp. 5–17.

Freeman, Benjamin G and Alexandra M Class Freeman (Mar. 2014). 'Rapid upslope shifts in New Guinean birds illustrate strong distributional responses of tropical montane species to global warming'. In: *Proc. Natl. Acad. Sci. U. S. A.* 111.12, pp. 4490–4494.

García Molinos, J, M T Burrows and E S Poloczanska (May 2017). 'Ocean currents modify the coupling between climate change and biogeographical shifts'. In: *Sci. Rep.* 7.1, p. 1332.

Gardner, Charlie J, Martin E Nicoll, Christopher Birkinshaw et al. (Apr. 2018). 'The rapid expansion of Madagascar's protected area system'. In: *Biol. Conserv.* 220, pp. 29–36.

Gardner, Charlie J, Martin E Nicoll, Tsibara Mbohoahy et al. (Dec. 2013). 'Protected areas for conservation and poverty alleviation: experiences from Madagascar'. In: *J. Appl. Ecol.* 50.6, pp. 1289–1294.

Gatiso, Tsegaye T et al. (Aug. 2022). 'Effectiveness of protected areas influenced by socio-economic context'. In: *Nat. Sustain.* 5.10, pp. 861–868.

Gaüzère, P et al. (Feb. 2020). 'Long-term effects of combined land-use and climate changes on local bird communities in mosaic agricultural landscapes'. In: *Agric. Ecosyst. Environ.* 289.106722, p. 106722.

Gaüzère, Pierre, Frédéric Jiguet and Vincent Devictor (Sept. 2015). 'Rapid adjustment of bird community compositions to local climatic variations and its functional consequences'. In: *Glob. Chang. Biol.* 21.9, pp. 3367–3378.

Gaüzère, Pierre, Karine Princé and Vincent Devictor (June 2017). 'Where do they go? The effects of topography and habitat diversity on reducing climatic debt in birds'. In: *Glob. Chang. Biol.* 23.6, pp. 2218–2229.

GBIF (2024a). `https://www.gbif.org`. Accessed: 2024-11-13.

GBIF (2024b). `https://www.gbif.org`. Accessed: 2024-11-13.

gbif.org, Registry-Migration (2018). *Southern African Bird Atlas Project.*

Ghosh, Shyamolina, Blake Matthews and Owen L Petchey (Oct. 2024). 'Temperature and biodiversity influence community stability differently in birds and fishes'. en. In: *Nat. Ecol. Evol.* 8.10, pp. 1835–1846.

Gillings, Simon, Dawn E Balmer and Robert J Fuller (June 2015). 'Directionality of recent bird distribution shifts and climate change in Great Britain'. In: *Glob. Chang. Biol.* 21.6, pp. 2155–2168.

Global Hunger Index (2024). *2024 Global Hunger Index: How Gender Justice Can Advance Climate Resilience and Zero Hunger.*

Gosling, Elizabeth (Apr. 2008). *Bivalve molluscs: Biology, ecology and culture.* Chichester, England: Wiley-Blackwell.

Grace, H Resit Akçakaya, Elizabeth Bennett et al. (Dec. 2019). 'Using historical and palaeoecological data to inform ambitious species recovery targets'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 374.1788, p. 20190297.

Grace, H Resit Akçakaya, Elizabeth L Bennett et al. (Dec. 2021). 'Testing a global standard for quantifying species recovery and assessing conservation impact'. In: *Conserv. Biol.* 35.6, pp. 1833–1849.

Graham, Edward et al. (Sept. 2004). 'Site selection for OWL using past, present, and future climate information'. In: *Ground-based Telescopes.* Ed. by Jacobus M Oschmann Jr. SPIE.

Green et al. (Apr. 2018). 'Local costs of conservation exceed those borne by the global majority'. In: *Glob. Ecol. Conserv.* 14.e00385, e00385.

Gregory, Richard D et al. (Mar. 2009). 'An indicator of the impact of climatic change on European bird populations'. In: *PLoS One* 4.3, e4678.

Greve, M et al. (June 2011). 'The ecological effectiveness of protected areas: a case study for South African birds: The effectiveness of protected areas'. In: *Anim. Conserv.* 14.3, pp. 295–305.

Grover, Claire et al. (Aug. 2010). 'Use of the Edinburgh geoparser for georeferencing digitized historical collections'. In: *Philos. Trans. A Math. Phys. Eng. Sci.* 368.1925, pp. 3875–3889.

Guan, Yanlong et al. (Nov. 2021). 'Changes in global climate heterogeneity under the 21st century global warming'. en. In: *Ecol. Indic.* 130.108075, p. 108075.

Guerin, Greg R et al. (Oct. 2024). 'Detecting ecosystem trends in response to climate and disturbance across continental plot networks: a power analysis'. In: *bioRxiv*, p. 2024.10.14.618336.

Guetté, Adrien et al. (Aug. 2017). 'Measuring the synanthropy of species and communities to monitor the effects of urbanization on biodiversity'. In: *Ecol. Indic.* 79, pp. 139–154.

Gupta, Ashish Kumar et al. (Oct. 2020). 'Salient object detection techniques in computer vision-A survey'. en. In: *Entropy (Basel)* 22.10, p. 1174.

Guralnick, Robert and Andrew Hill (Feb. 2009). 'Biodiversity informatics: automated approaches for documenting global biodiversity patterns and processes'. In: *Bioinformatics* 25.4, pp. 421–428.

Hall and Reginald Ernest Moreau (1970). *An atlas of speciation in African passerine birds*. Trustees of the British museum (Natural history).

Hallman, Tyler A et al. (2021). 'Building a better baseline to estimate 160 years of avian population change and create historically informed conservation targets'. In: *Conserv. Biol.* 35.4, pp. 1256–1267.

Hanauer, Merlin M and Gustavo Canavire-Bacarreza (Nov. 2015). 'Implications of heterogeneous impacts of protected areas on deforestation and poverty'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 370.1681, p. 20140272.

Harebottle, D M et al. (2007). 'Southern African bird atlas project 2: Instruction manual'. In: *Cape Town: Animal Demography Unit, University of Cape Town*.

Harper, Elizabeth M (Dec. 2003). 'Assessing the importance of drilling predation over the Palaeozoic and Mesozoic'. In: *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 201.3, pp. 185–198.

Harper, Elizabeth M, J Alistair Crame and Caroline E Sogot (June 2018). '"Business as usual": Drilling predation across the K-Pg mass extinction event in Antarctica'. In: *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 498, pp. 115–126.

Harper, Elizabeth M, Robinson and Lee (May 2011). 'Drill hole analysis reveals evidence of targeted predation on modern brachiopods'. In: *Palaeogeogr. Palaeoclimatol. Palaeoecol.* 305.1-4, pp. 162–171.

Harris, K Jarrod Millman et al. (2020). 'Array programming with NumPy'. In: *Nature* 585.7825, pp. 357–362.

Harris, Timothy J Osborn et al. (Apr. 2020). 'Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset'. In: *Sci. Data* 7.1, p. 109.

Hastings, Reuben A et al. (Apr. 2020). 'Climate change drives poleward increases and equatorward declines in marine species'. In: *Curr. Biol.* 30.8, 1572–1577.e2.

He, Kaiming et al. (Dec. 2015). 'Deep Residual Learning for Image Recognition'. In: *arXiv [cs.CV]*.

Hedrick, Brandon P et al. (Mar. 2020). 'Digitization and the future of natural history collections'. In: *Bioscience* 70.3, pp. 243–251.

Heitzler, Magnus et al. (2018). *A modular process to improve the georeferencing of the Siegfried map*. `http://www.e-perimetron.org/Vol_13_2/Heitzler_et_al.pdf`. Accessed: 2023-8-30.

Herrando, Sergi et al. (Apr. 2019). 'Contrasting impacts of precipitation on Mediterranean birds and butterflies'. en. In: *Sci. Rep.* 9.1, p. 5680.

Herrera, James P et al. (June 2021). 'Food insecurity related to agricultural practices and household characteristics in rural communities of northeast Madagascar'. In: *Food Secur* 13.6, pp. 1393–1405.

Hijmans, Robert J (2025). *terra: Spatial Data Analysis*.

Hilborn, Ray and Michel J Kaiser (July 2022). 'A path forward for analysing the impacts of marine protected areas'. In: *Nature* 607.7917, E1–E2.

Hintsanen, Leena et al. (Oct. 2023). 'Temperature niche composition change inside and outside protected areas under climate warming'. In: *Conserv. Biol.* 37.5, e14134.

Hirschmann, Fabian (2016). *rdp: Python/Numpy implementation of the Ramer-Douglas-Peucker algorithm*.

Hiscock, Keith et al. (Jan. 2001). *The impact of climate change on subtidal and intertidal benthic species in Scotland*.

Ho, Madison et al. (Jan. 2023). 'Hierarchical multi-label taxonomic classification of carbonate skeletal grains with deep learning'. In: *Sediment. Geol.* 443, p. 106298.

Hockey, Philip A R et al. (Mar. 2011). 'Interrogating recent range changes in South African birds: confounding signals from land use and climate change present a challenge for attribution: Range changes in South African birds'. In: *Divers. Distrib.* 17.2, pp. 254–261.

Hsiang, Allison Y et al. (July 2019). 'Endless forams: >34,000 modern planktonic foraminiferal images for taxonomic training and automated species recognition using convolutional neural networks'. In: *Paleoceanogr. paleoclimatology* 34.7, pp. 1157–1177.

Huang, Jian and Emmanuel A Frimpong (June 2015). 'Using historical atlas data to develop high-resolution distribution models of freshwater fishes'. In: *PLoS One* 10.6, e0129995.

Hugo, Sanet and van Rensburg (Sept. 2008). 'The maintenance of a positive spatial correlation between South African bird species richness and human population density'. In: *Glob. Ecol. Biogeogr.* 17.5, pp. 611–621.

Illán, Javier Gutiérrez et al. (Nov. 2014). 'Precipitation and winter temperature predict long-term range-scale abundance changes in Western North American birds'. In: *Glob. Chang. Biol.* 20.11, pp. 3351–3364.

Jachowski, David S et al. (2015). 'Redefining baselines in endangered species recovery'. In: *J. Wildl. Manage.* 79.1, pp. 3–9.

Jessen, Tyler D et al. (Mar. 2022). 'Contributions of Indigenous Knowledge to ecological and evolutionary understanding'. In: *Front. Ecol. Environ.* 20.2, pp. 93–101.

Jetz, Walter, David S Wilcove and Andrew P Dobson (2007). 'Projected impacts of climate and land-use change on the global diversity of birds'. In: *PLoS Biol.* 5.6, e157.

Jocher, Glenn, Ayush Chaurasia and Jing Qiu (2023). *Ultralytics YOLOv8*. Version 8.0.0.

Johansen, Chloe, Jennifer Horney and Iris Tien (June 2017). 'Metrics for evaluating and improving community resilience'. In: *J. Infrastruct. Syst.* 23.2, p. 04016032.

Jones et al. (Dec. 2020). 'Investigating the implications of shifting baseline syndrome on conservation'. In: *People Nat. (Hoboken)* 2.4, pp. 1131–1144.

Joshi, Ajay J, Fatih Porikli and Nikolaos Papanikolopoulos (June 2009). 'Multi-class active learning for image classification'. en. In: *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 2372–2379.

Jury, Mark R (2013). 'Climate trends in southern Africa'. In: *S. Afr. J. Sci.*

Kelley and Hansen (2003). *The Fossil Record of Drilling Predation on Bivalves and Gastropods*.

Kelly, Anne E and Michael L Goulden (Aug. 2008). 'Rapid shifts in plant distribution with recent climate change'. In: *Proc. Natl. Acad. Sci. U. S. A.* 105.33, pp. 11823–11826.

Kemp, Melissa E (Feb. 2023). 'Characterizing past communities to build future ones: Lessons from Caribbean conservation paleobiology'. In: *Bull. Fla. Mus. Nat. Hist.* 60.2, pp. 58–58.

Khaliq, Imran et al. (Mar. 2024). 'Warming underpins community turnover in temperate freshwater and terrestrial communities'. In: *Nat. Commun.* 15.1, p. 1921.

Kidwell, S M (Nov. 2001). 'Preservation of species abundance in marine death assemblages'. In: *Science* 294.5544, pp. 1091–1094.

Kiel, Steffen (Apr. 2021). 'Assessing bivalve phylogeny using Deep Learning and Computer Vision approaches'. In: *bioRxiv*, p. 2021.04.08.438943.

Kiessling, Wolfgang et al. (Dec. 2019). 'Addressing priority questions of conservation science with palaeontological data'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 374.1788, p. 20190222.

Kingma, Diederik P and Jimmy Ba (Dec. 2014). 'Adam: A method for stochastic optimization'. In: *arXiv [cs.LG]*.

Kingsley-Smith, Peter R, Christopher A Richardson and Raymond Seed (Nov. 2003). 'Stereotypic and size-selective predation in Polinices pulchellus (Gastropoda: Naticidae) Risso 1826'. In: *J. Exp. Mar. Bio. Ecol.* 295.2, pp. 173–190.

Kissling, W Daniel and Gudrun Carl (Jan. 2008). 'Spatial autocorrelation and the selection of simultaneous autoregressive models'. In: *Glob. Ecol. Biogeogr.* 17.1, pp. 59–71.

Klompmaker (Nov. 2009). 'Taphonomic bias on drill-hole predation intensities and paleoecology of Pliocene mollusks from langenboom (mill), the Netherlands'. In: *Palaios* 24.11, pp. 772–779.

Klompmaker, Adiël A et al. (June 2017). 'Increase in predator-prey size ratios throughout the Phanerozoic history of marine ecosystems'. In: *Science* 356.6343, pp. 1178–1180.

Kumar, Teerath, Rob Brennan et al. (2024). 'Image data augmentation approaches: A comprehensive survey and future directions'. en. In: *IEEE Access* 12, pp. 187536–187571.

Kumar, Teerath, Alessandra Mileo et al. (Jan. 2023). 'Image Data Augmentation Approaches: A Comprehensive Survey and Future directions'. In: *arXiv [cs.CV]*.

La Sorte, Frank A and Frank R Thompson 3rd (July 2007). 'Poleward shifts in winter ranges of North American birds'. In: *Ecology* 88.7, pp. 1803–1812.

Labay, Ben et al. (Sept. 2011). 'Assessing historical fish community composition using surveys, historical collection data, and species distribution models'. In: *PLoS One* 6.9, e25145.

Lajeunesse, Armelle and Yoan Fourcade (Feb. 2023). 'Temporal analysis of GBIF data reveals the restructuring of communities following climate change'. In: *J. Anim. Ecol.* 92.2, pp. 391–402.

Lee, Res Altwegg and Phoebe Barnard (Sept. 2017). 'Estimating conservation metrics from atlas data: the case of southern African endemic birds'. In: *Bird Conserv. Int.* 27.3, pp. 323–336.

Legendre, Pierre (Sept. 1993). 'Spatial Autocorrelation: Trouble or New Paradigm?' In: *Ecology* 74.6, pp. 1659–1673.

Lehikoinen, Aleksi et al. (May 2021). 'Wintering bird communities are tracking climate change faster than breeding communities'. In: *J. Anim. Ecol.* 90.5, pp. 1085–1095.

Lendemer, James et al. (Jan. 2020). 'The extended specimen network: A strategy to enhance US biodiversity collections, promote research and education'. In: *Bioscience* 70.1, pp. 23–30.

Lennard, C J et al. (June 2018). 'On the need for regional climate information over Africa under varying levels of global warming'. In: *Environ. Res. Lett.* 13.6, p. 060401.

Leonard-Pingel, Jill S and Jeremy B C Jackson (Aug. 2016). 'Drilling predation increased in response to changing environments in the Caribbean Neogene'. In: *Paleobiology* 42.3, pp. 394–409.

Lepetz, Virginie et al. (Nov. 2009). 'Biodiversity monitoring: some proposals to adequately study species' responses to climate change'. In: *Biodivers. Conserv.* 18.12, pp. 3185–3203.

Lerm, Rion Elias and Les G Underhill (Mar. 2019). 'Bird Distribution Dynamics - African black oystercatcher in South Africa'. In: *Biodiv. Obs.* 10, 10.3:1–7.

Lewis, John R (Nov. 1953). 'The ecology of rooky shores around Anglesey'. In: *Proc. Zool. Soc. Lond. (1944)* 123.3, pp. 481–550.

Li, Guanbin and Yizhou Yu (June 2015). 'Visual Saliency Based on Multiscale Deep Features'. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Li and Ningchuan Xiao (Nov. 2020). 'Using Machine Learning Methods to Identify and Classify the Regions and Projections of Online Maps'. In: *Geocomputation 2019*. unknown.

Lindström, Åke et al. (Mar. 2013). 'Rapid changes in bird community composition at multiple temporal and spatial scales in response to recent climate change'. en. In: *Ecography (Cop.)* 36.3, pp. 313–322.

Lister, Adrian M and Climate Change Research Group (Apr. 2011). 'Natural history collections as sources of long-term datasets'. In: *Trends Ecol. Evol.* 26.4, pp. 153–154.

Liu, Xin Cao et al. (2024). *GlobPOP: A 33-year (1990-2022) global gridded population dataset generated by cluster analysis and statistical learning*.

Liu, Shouyi Jiang et al. (Mar. 2022). 'Automatic taxonomic identification based on the Fossil Image Dataset (>415,000 images) and deep convolutional neural networks'. In: *Paleobiology*, pp. 1–22.

Liu, Xiaokang and Haijun Song (Dec. 2020). 'Automatic identification of fossils and abiotic grains during carbonate microfacies analysis using deep convolutional neural networks'. In: *Sediment. Geol.* 410, p. 105790.

Lolle, Henrik Lauridsen (Jan. 2023). 'Differences in subjective well-being between rural and urban areas in Denmark'. In: *Rural quality of life*. Manchester University Press, pp. 348–369.

Lotze, Heike K and Boris Worm (2009). 'Historical baselines for large marine animals'. In: *Trends Ecol. Evol.* 24.5, pp. 254–262.

Luft, Jonas (Jan. 2020). 'Automatic Georeferencing of Historical Maps by Geocoding'. In: *International Workshop on Automatic Vectorisation of Historical Maps*. unknown, pp. 77–89.

Macías-Duarte, Alberto and Courtney J Conway (Mar. 2015). 'Distributional Changes in the Western Burrowing Owl (Athene cunicularia hypugaea) in North America from 1967 to 2008'. In: *J. Raptor Res.* 49.1, pp. 75–83.

Maclean, Ilya M D et al. (Jan. 2017). 'Fine-scale climate change: modelling spatial variation in biologically meaningful rates of warming'. en. In: *Glob. Chang. Biol.* 23.1, pp. 256–268.

Mäder, P et al. (2021). 'The flora incognita app-interactive plant species identification'. In: *Methods in Ecology and Evolution*.

Mahadevkar, Supriya V et al. (2022). 'A Review on Machine Learning Styles in Computer Vision—Techniques and Future Directions'. en. In: *IEEE Access* 10, pp. 107293–107329.

Mammides, Christos (Jan. 2020). 'Evidence from eleven countries in four continents suggests that protected areas are not associated with higher poverty rates'. In: *Biol. Conserv.* 241.108353, p. 108353.

Maravillas, Alme B, Larmie S Feliscuzo and James Arnold E Nogra (2022). *Neural network approach for bivalves classification.*

Marcer, Arnald et al. (Mar. 2021). 'Quality issues in georeferencing: From physical collections to digital data repositories for ecological research'. In: *Divers. Distrib.* 27.3, pp. 564–567.

Marcer et al. (Sept. 2022). 'Uncertainty matters: ascertaining where specimens in natural history collections come from and its implications for predicting species distributions'. In: *Ecography (Cop.)* 2022.9, e06025.

Marchant, Ross et al. (2020). 'Automated analysis of foraminifera fossil records by image classification using a convolutional neural network'. In: *Journal of Micropalaeontology* 39.2, pp. 183–202.

Marine Life Information Network (2020). `https://www.marlin.ac.uk`. Accessed: 2024–.

Marsh, David M and Peter C Trenham (2008). 'Current trends in plant and animal population monitoring'. In: *Conserv. Biol.* 22.3, pp. 647–655.

Marta, Silvio et al. (Aug. 2023). 'Heterogeneous changes of soil microclimate in high mountains and glacier forelands'. en. In: *Nat. Commun.* 14.1, p. 5306.

Martín Abadi et al. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems.*

Matthews, Stephen N et al. (2011). 'Changes in potential habitat of 147 North American breeding bird species in response to redistribution of trees and climate following predicted climate change'. In: *Ecography* 34.6, pp. 933–945.

McClenachan, Loren, Francesco Ferretti and Julia K Baum (Oct. 2012). 'From archives to conservation: why historical data are needed to set baselines for marine animals and ecosystems'. In: *Conserv. Lett.* 5.5, pp. 349–359.

McCollin, Duncan, Richard C Preece and Tim H Sparks (Dec. 2015). 'Pushing back the baseline: a novel approach to detect long-term changes in terrestrial faunal abundance using historical qualitative descriptions'. In: *Eur. J. Ecol.* 1.1, pp. 32–42.

McConnell, William J, Sean P Sweeney and Bradley Mulley (Feb. 2004). 'Physical and social access to land: spatio-temporal patterns of agricultural expansion in Madagascar'. In: *Agric. Ecosyst. Environ.* 101.2, pp. 171–184.

McCrea, Rod, Andrea Walton and Rosemary Leonard (Oct. 2014). 'A conceptual framework for investigating community wellbeing and resilience'. In: *Rural Soc.* 23.3, pp. 270–282.

McInnes, Leland, John Healy and Steve Astels (2017). 'hdbscan: Hierarchical density based clustering'. In: *J. Open Source Softw.* 2.11, p. 205.

Meineke, Emily K and Barnabas H Daru (Dec. 2021). 'Bias assessments to expand research harnessing biological collections'. en. In: *Trends Ecol. Evol.* 36.12, pp. 1071–1082.

Mendes, Pedro, Paolo Romano and David Garlan (May 2024). 'Error-Driven Uncertainty Aware Training'. In: *arXiv [cs.LG].*

Mieszkowska, N et al. (Feb. 2006). 'Changes in the range of some common rocky shore species in Britain – A response to climate change?' In: *Hydrobiologia* 555.1, pp. 241–251.

Mihoub, Jean-Baptiste et al. (Jan. 2017). 'Setting temporal baselines for biodiversity: the limits of available monitoring data for capturing the full impact of anthropogenic pressures'. In: *Sci. Rep.* 7.1, p. 41591.

Miller, Sara E et al. (July 2020). 'Building Natural history collections for the twenty-first century and beyond'. In: *Bioscience* 70.8, pp. 674–687.

Minaee, Shervin et al. (July 2022). 'Image segmentation using deep learning: A survey'. en. In: *IEEE Trans. Pattern Anal. Mach. Intell.* 44.7, pp. 3523–3542.

Mollusca Types in Britain and Ireland (2023). `https://gbmolluscatypes.ac.uk/`. Accessed: 2024–.

Montràs-Janer, Teresa et al. (Apr. 2024). 'Anthropogenic climate and land-use change drive short- and long-term biodiversity shifts across taxa'. In: *Nat. Ecol. Evol.* 8.4, pp. 739–751.

Moore, P J and D A Smale (Jan. 2020). 'Impacts of climate change on shallow and shelf subtidal habitats, relevant to the coastal and marine environment around the UK'. In: *MCCIP Science Review*, pp. 272–292.

Morecroft, Michael D and L Speakman (May 2013). *Terrestrial biodiversity climate change impacts - report card 2012 -13*. LWEC Climate Change Impact Report Cards. Swindon, England: Living With Environmental Change.

Morin, Peter J (2009). *Community ecology*. John Wiley & Sons.

Murphy, Paul C et al. (Sept. 2004). 'Georeferencing of museum collections: A review of problems and automated tools, and the methodology developed by the mountain and plains spatio-temporal database-informatics initiative (MaPSTeDI)'. In: 3, pp. 1–29.

Naidoo, Robin and Takuya Iwamura (Nov. 2007). 'Global-scale mapping of economic benefits from agricultural lands: Implications for conservation priorities'. In: *Biol. Conserv.* 140.1-2, pp. 40–49.

National Statistical Office (NSO) (2017). *Fourth integrated household survey 2016-2017*.

Nelson and Shari Ellis (Nov. 2018). 'The history and impact of digitization and digital data mobilization on biodiversity research'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 374.1763, p. 20170391.

Niedertscheider, Maria, Simone Gingrich and Karl-Heinz Erb (Dec. 2012). 'Changes in land use in South Africa between 1961 and 2006: an integrated socio-ecological analysis based on the human appropriation of net primary production framework'. In: *Reg. Environ. Change* 12.4, pp. 715–727.

Niu, Zhi-Bin and Hong-He Xu (Jan. 2022). 'AI-based graptolite identification improves shale gas exploration'. In: *bioRxiv*, p. 2022.01.17.476477.

Nowak, Maciej M, Katarzyna Słupecka and Bogdan Jackowiak (Nov. 2021). 'Geotagging of natural history collections for reuse in environmental research'. In: *Ecol. Indic.* 131.108131, p. 108131.

Ojaveer, Henn et al. (Aug. 2018). 'Historical baselines in marine bioinvasions: Implications for policy and management'. In: *PLoS One* 13.8, e0202383.

Oksanen, Jari et al. (2025). *vegan: Community Ecology Package*.

Oliver, Tom H et al. (June 2017). 'Large extents of intensive land use limit community reorganization during climate warming'. In: *Glob. Chang. Biol.* 23.6, pp. 2272–2283.

OpenTopography (2013). *Shuttle Radar Topography Mission (SRTM) Global*.

Ortiz-Bobea, Ariel (May 2021). 'Climate, Agriculture and Food'. In: *arXiv [econ.GN]*.

Outhwaite, Charlotte L, Peter McCann and Tim Newbold (May 2022). 'Agriculture and climate change are reshaping insect biodiversity worldwide'. In: *Nature* 605.7908, pp. 97–102.

Palmer, A R (1988). 'Feeding biology of Ocenebra-lurida (prosobranchia, muricacea) - diet, predator-prey size relations, and attack behavior'. In: *Veliger* 31, pp. 192–203.

Papworth, S K et al. (2009). 'Evidence for shifting baseline syndrome in conservation'. In: *Conserv. Lett.*

Parker, Evan J et al. (July 2024). 'Insufficient and biased representation of species geographic responses to climate change'. In: *Glob. Chang. Biol.* 30.7, e17408.

Parmesan, Camille (Dec. 2006). 'Ecological and evolutionary responses to recent climate change'. In: *Annu. Rev. Ecol. Evol. Syst.* 37.1, pp. 637–669.

Parmesan, Camille and Gary Yohe (Jan. 2003). 'A globally coherent fingerprint of climate change impacts across natural systems'. In: *Nature* 421.6918, pp. 37–42.

Parra, Juan L and William B Monahan (Oct. 2008). 'Variability in 20th century climate change reconstructions and its consequences for predicting geographic responses of California mammals: 20TH CENTURY CLIMATE CHANGE'. In: *Glob. Chang. Biol.* 14.10, pp. 2215–2231.

Pbc, Planet Labs (2024). *Planet Application Program Interface: In Space for Life on Earth*.

Pearce-Higgins, James W, Sarah M Eglington et al. (July 2015). 'Drivers of climate change impacts on bird communities'. In: *J. Anim. Ecol.* 84.4, pp. 943–954.

Pearce-Higgins, James W, Rhys E Green and Rhys Green (June 2014). *Birds and Climate Change: Impacts and Conservation Responses*. Cambridge University Press.

Pearson, Richard G (Mar. 2006). 'Climate change and the migration capacity of species'. en. In: *Trends Ecol. Evol.* 21.3, pp. 111–113.

Pedregosa, F et al. (2011). 'Scikit-learn: Machine Learning in Python'. In: *J. Mach. Learn. Res.* 12, pp. 2825–2830.

PEN (2016). *CIFOR's Poverty and Environment Network (PEN) global dataset*.

Pereira, Paul W Leadley et al. (Dec. 2010). 'Scenarios for global biodiversity in the 21st century'. In: *Science* 330.6010, pp. 1496–1501.

Pereira, Laetitia Marie Navarro and Inês Santos Martins (Nov. 2012). 'Global biodiversity change: The bad, the good, and the unknown'. In: *Annu. Rev. Environ. Resour.* 37.1, pp. 25–50.

Perez and Wang (Dec. 2017). 'The Effectiveness of Data Augmentation in Image Classification using Deep Learning'. In: *arXiv [cs.CV]*.

Peterson, A Townsend (2003). 'Subtle Recent Distributional Shifts in Great Plains Bird Species'. In: *Southwest. Nat.* 48.2, pp. 289–292.

Peterson, A Townsend and Enrique Martínez-Meyer (Nov. 2008). 'Pervasive poleward shifts among North American bird species'. In: *Biodiversity* 9.3-4, pp. 114–116.

Pimm, S L et al. (May 2014). 'The biodiversity of species and their rates of extinction, distribution, and protection'. In: *Science* 344.6187, p. 1246752.

Pincebourde, Sylvain and Christelle Suppo (July 2016). 'The vulnerability of tropical ectotherms to warming is modulated by the microclimatic heterogeneity'. en. In: *Integr. Comp. Biol.* 56.1, pp. 85–97.

Poloczanska, Elvira S et al. (Oct. 2013). 'Global imprint of climate change on marine life'. In: *Nat. Clim. Chang.* 3.10, pp. 919–925.

Pomeroy, Derek (1989). 'Using East African bird atlas data for ecological studies'. In: *Ann. Zool. Fennici* 26.3, pp. 309–314.

Pomeroy, Derek and Dorothy Ssekabiira (1990). 'An analysis of the distributions of terrestrial birds in Africa'. In: *Afr. J. Ecol.* 28.1, pp. 1–13.

Posit team (2024). *RStudio: Integrated Development Environment for R*. Boston, MA.

Princé, Karine and Benjamin Zuckerberg (Feb. 2015). 'Climate change in our backyards: the reshuffling of North America's winter bird communities'. In: *Glob. Chang. Biol.* 21.2, pp. 572–585.

QGIS (2024). `http://www.qgis.org`. Accessed: 2024-11-13.

R Core Team (2025). *R: A Language and Environment for Statistical Computing*. Vienna, Austria.

Ralimanana, Hélène et al. (Dec. 2022). 'Madagascar's extraordinary biodiversity: Threats and opportunities'. In: *Science* 378.6623, eadf1466.

Ramer, Urs (Nov. 1972). 'An iterative procedure for the polygonal approximation of plane curves'. In: *Comput. Graph. Image Process.* 1.3, pp. 244–256.

Rasolofoson, Ranaivo A, Martin R Nielsen and Jones (May 2018). 'The potential of the Global Person Generated Index for evaluating the perceived impacts of conservation interventions on subjective well-being'. In: *World Dev.* 105, pp. 107–118.

Rato, Ana et al. (July 2022). 'The impact of climate change on bivalve farming: Combined effect of temperature and salinity on survival and feeding behavior of clams Ruditapes decussatus'. In: *Front. Mar. Sci.* 9, p. 932310.

'Polynomial Regression' (2006). In: *Springer Texts in Statistics*. Ed. by J O Rawlings, S G Pantula and D A Dickey. New York: Springer-Verlag, pp. 235–268.

Rensburg, van, Chown and Gaston (May 2002). 'Species richness, environmental correlates, and spatial scale: a test using South African birds'. In: *Am. Nat.* 159.5, pp. 566–577.

Ribas, Luiz Guilherme dos Santos et al. (June 2020). 'A global comparative analysis of impact evaluation methods in estimating the effectiveness of protected areas'. In: *Biol. Conserv.* 246, p. 108595.

Richard, Yves et al. (June 2001). '20th century droughts in southern Africa: spatial and temporal variability, teleconnections with oceanic and atmospheric conditions'. In: *Int. J. Climatol.* 21.7, pp. 873–885.

Rick, Torben C and Rowan Lockwood (Feb. 2013). 'Integrating paleobiology, archeology, and history to inform biological conservation'. In: *Conserv. Biol.* 27.1, pp. 45–54.

Robinson et al. (May 2021). 'Traditional ecological knowledge in restoration ecology: a call to listen deeply, to engage with, and respect Indigenous voices'. In: *Restor. Ecol.* 29.4, e13381.

Rodrigues, Ana S L et al. (Dec. 2019). 'Unshifting the baseline: a framework for documenting historical population changes and assessing long-term anthropogenic impacts'. In: *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 374.1788, p. 20190220.

Rodríguez-Rodríguez, D et al. (Mar. 2024). 'Socioeconomic impacts of small conserved sites on rural communities in Madagascar'. In: *Environmental Development* 49, p. 100965.

Roe, Dilys et al. (Feb. 2014). 'Which components or attributes of biodiversity influence which dimensions of poverty?' In: *Environ. Evid.* 3.1, p. 3.

Rubenstein, Madeleine A et al. (Apr. 2023). 'Climate change and the global redistribution of biodiversity: substantial variation in empirical support for expected range shifts'. In: *Environmental Evidence* 12.1, pp. 1–21.

Rull, Valentí (Mar. 2009). 'Microrefugia'. en. In: *J. Biogeogr.* 36.3, pp. 481–484.

Rushing, Clark S et al. (June 2020). 'Migratory behavior and winter geography drive differential range shifts of eastern birds in response to recent climate change'. In: *Proc. Natl. Acad. Sci. U. S. A.* 117.23, pp. 12897–12903.

Sager, Christoph, Christian Janiesch and Patrick Zschech (July 2021). 'A survey of image labelling for computer vision applications'. en. In: *J. Bus. Anal.* 4.2, pp. 91–110.

Sandler, Mark et al. (Jan. 2018). 'MobileNetV2: Inverted residuals and linear bottlenecks'. In: *arXiv [cs.CV]*.

Sansom, Robert S (Mar. 2015). 'Bias and sensitivity in the placement of fossil taxa resulting from interpretations of missing data'. In: *Syst. Biol.* 64.2, pp. 256–266.

Santangeli, Andrea and Aleksi Lehikoinen (Mar. 2017). 'Are winter and breeding bird communities able to track rapid climate change? Lessons from the high North'. In: *Divers. Distrib.* 23.3, pp. 308–316.

Santangeli, Andrea, Ari Rajasärkkä and Aleksi Lehikoinen (June 2017). 'Effects of high latitude protected areas on bird communities under rapid climate change'. In: *Glob. Chang. Biol.* 23.6, pp. 2241–2249.

Segal, E (May 1961). 'ACCLIMATION IN MOLLUSCS'. In: *Integrative and Comparative Biology* 1.2, pp. 235–244.

Şekercioğlu, Çağan H, Richard B Primack and Janice Wormworth (Apr. 2012). 'The effects of climate change on tropical birds'. In: *Biol. Conserv.* 148.1, pp. 1–18.

Selin, Helaine (Dec. 2010). *Nature across cultures: Views of nature and the environment in non-western cultures*. Ed. by Helaine Selin. Science Across Cultures: The History of Non-Western Science. Dordrecht, Netherlands: Springer.

Seltmann, Katja et al. (Dec. 2018). 'Georeferencing for Research Use (GRU): An integrated geospatial training paradigm for biocollections researchers and data providers'. In: *Res. Ideas Outcomes* 4.

Selvaraju, Ramprasaath R et al. (Oct. 2017). 'Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization'. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Sheldon, Kimberly S (Nov. 2019). 'Climate change in the tropics: Ecological and evolutionary responses at low latitudes'. In: *Annu. Rev. Ecol. Evol. Syst.* 50.1, pp. 303–333.

Simmons, Robert E et al. (Nov. 2009). 'Climate change and birds: Perspectives and prospects from southern Africa'. In: *Ostrich - Journal of African Ornithology* 75.4.

Snow, David William (1978). *An atlas of speciation in African non-passerine birds*. London, England: British Museum Press.

Soberón, Jorge (Dec. 2007). 'Grinnellian and Eltonian niches and geographic distributions of species'. In: *Ecol. Lett.* 10.12, pp. 1115–1123.

Soga, Masashi and Kevin J Gaston (May 2018). 'Shifting baseline syndrome: causes, consequences, and implications'. In: *Front. Ecol. Environ.* 16.4, pp. 222–230.

Sokal, Robert R and Neal L Oden (June 1978). 'Spatial autocorrelation in biology: 1. Methodology'. In: *Biol. J. Linn. Soc. Lond.* 10.2, pp. 199–228.

Sparrius, L, G Top and C Swaay (Jan. 2018). 'An approach to calculate a Species Temperature Index for flora based on open data'. In: *Gorteria* 40.1, pp. 73–78.

stagmerlinapp (n.d.). *Merlin Bird ID - Home*. en. `https://merlin.allaboutbirds.org/`. Accessed: 2025-7-11.

Steinmetz, Robert, Wanlop Chutipong and Naret Seuaturien (Oct. 2006). 'Collaborating to conserve large mammals in Southeast Asia'. In: *Conserv. Biol.* 20.5, pp. 1391–1401.

Stolton, S, P Shadie and N Dudley (2013). *IUCN WCPA Best Practice Guidance on Recognising Protected Areas and Assigning Management Categories and Governance Types*.

Stribling, James B, Stephen R Moulton II and Gary T Lester (Dec. 2003). 'Determining the quality of taxonomic data'. en. In: *J. North Am. Benthol. Soc.* 22.4, pp. 621–631.

Stroud, James T et al. (Nov. 2015). 'Is a community still a community? Reviewing definitions of key terms in community ecology'. en. In: *Ecol. Evol.* 5.21, pp. 4757–4765.

Sullivan, Brian L et al. (Oct. 2009). 'eBird: A citizen-based bird observation network in the biological sciences'. In: *Biol. Conserv.* 142.10, pp. 2282–2292.

Sun, Jiarui et al. (Oct. 2023). 'Automatic identification and morphological comparison of bivalve and brachiopod fossils based on deep learning'. In: *PeerJ* 11, e16200.

Szabo, Judit K et al. (Oct. 2012). 'Global patterns and drivers of avian extinctions at the species and subspecies level'. In: *PLoS One* 7.10, e47080.

Tabarisaadi, Pegah et al. (May 2024). 'An optimized uncertainty-aware training framework for neural networks'. en. In: *IEEE Trans. Neural Netw. Learn. Syst.* 35.5, pp. 6928–6935.

Taheri, Shirin, Babak Naimi and Miguel B Araújo (Aug. 2016). 'Did British breeding birds move north in the late 20th century?' In: *Climate Change Responses* 3.1, pp. 1–5.

Tan, Mingxing and Quoc V Le (May 2019). 'EfficientNet: Rethinking model scaling for convolutional Neural Networks'. In: *arXiv [cs.LG]*.

Tang, C M (June 2000). 'Ugly Fossil Syndrome'. In: *Palaios* 15.3, pp. 175–176.

Terry, J Christopher D, Helen E Roy and Tom A August (Feb. 2020). 'Thinking like a naturalist: Enhancing computer vision of citizen science images by harnessing contextual data'. en. In: *Methods Ecol. Evol.* 11.2, pp. 303–315.

The DHS Program (n.d.). `https://www.dhsprogram.com/`. Accessed: 2025-6-20.

Theobald, David M et al. (Sept. 2020). 'Earth transformed: detailed mapping of global human modification from 1990 to 2017'. In: *Earth Syst. Sci. Data* 12.3, pp. 1953–1972.

Thorn, Caitlin (Aug. 2022). 'Georeferencing Historic Collection Data'. In: *Biodivers. Inf. Sci. Stand.* 6.

Thuiller, Wilfried et al. (2014). 'The European functional tree of bird life in the face of global change'. In: *Nat. Commun.* 5, p. 3118.

Tingley, Morgan W et al. (Nov. 2009). 'Birds track their Grinnellian niche through a century of climate change'. In: *Proc. Natl. Acad. Sci. U. S. A.* 106 Suppl 2.supplement_2, pp. 19637–19643.

Tobias, Joseph A et al. (2022). 'AVONET: morphological, ecological and geographical data for all birds'. In: *Ecol. Lett.* 25.3, pp. 581–597.

Toillier, Aurélie et al. (Jan. 2011). 'Livelihood Strategies and Land Use Changes in Response to Conservation: Pitfalls of Community-Based Forest Management in Madagascar'. In: *J. Sustainable For.* 30.1-2, pp. 20–56.

Tomich, Thomas P et al. (Nov. 2011). 'Agroecology: A review from a global-change perspective'. In: *Annu. Rev. Environ. Resour.* 36.1, pp. 193–222.

Toth, Lauren T (Feb. 2023). 'Geological perspectives on the degradation and restoration of coral reefs'. In: *Bull. Fla. Mus. Nat. Hist.* 60.2, pp. 60–60.

Underhill, Les G (Aug. 2016). 'The fundamentals of the SABAP2 protocol'. In: *Biodiv. Obs.*, pp. 1–12.

Underhill, Les G and Michael Brooks (June 2014). 'Preliminary summary of changes in bird distributions between the First and Second Southern African Bird Atlas Projects (SABAP1 and SABAP2)'. In: *Biodiv. Obs.*, pp. 258–293.

— (Sept. 2016). 'Displaying changes in bird distributions between SABAP1 and SABAP2'. In: *Biodiv. Obs.*, pp. 1–13.

UNDP-OPHI (2023). *Global multidimensional poverty index 2023. Unpacking global poverty: Data for high impact action.* `https://hdr.undp.org/system/files/documents/hdp-document/2023mpireportenpdf.pdf`. Accessed: –.

UNEP-WCMC and IUCN (2025). *Protected Planet: The World Database on Protected Areas (WDPA) and World Database on Other Effective Area-based Conservation Measures (WD-OECM)*. Cambridge, UK.

Usman, M T and C J C Reason (June 2004). 'Dry spell frequencies and their variability over southern Africa'. In: *Clim. Res.* 26.3, pp. 199–211.

VanDerWal, Jeremy et al. (Sept. 2012). 'Focus on poleward shifts in species' distribution underestimates the fingerprint of climate change'. In: *Nat. Clim. Chang.* 3.3, pp. 239–243.

Venables, W N and B D Ripley (2002). *Modern Applied Statistics with S*. New York.

Venceslau, Leonor and Luis Lopes (July 2019). 'Comparison of automated georeferencing tools using insect collection data'. In: *Biodivers. Inf. Sci. Stand.* 3.

Ver Hoef, Jay M et al. (Feb. 2018). 'Spatial autoregressive models for statistical inference from ecological data'. In: *Ecol. Monogr.* 88.1, pp. 36–59.

Verbrugge, Laura N H et al. (June 2012). 'Sensitivity of native and non-native mollusc species to changing river water temperature and salinity'. In: *Biol. Invasions* 14.6, pp. 1187–1199.

Verween, A, M Vincx and S Degraer (Sept. 2007). 'The effect of temperature and salinity on the survival of Mytilopsis leucophaeata larvae (Mollusca, Bivalvia): The search for environmental limits'. In: *J. Exp. Mar. Bio. Ecol.* 348.1-2, pp. 111–120.

Vilhena, Daril A and Andrew B Smith (Oct. 2013). 'Spatial bias in the marine fossil record'. In: *PLoS One* 8.10, e74470.

Virkkala, Raimo and Aleksi Lehikoinen (Oct. 2014). 'Patterns of climate-induced density shifts of species: poleward shifts faster in northern boreal birds than in southern birds'. In: *Glob. Chang. Biol.* 20.10, pp. 2995–3003.

Wang, Chao et al. (Jan. 2020). 'Determinants of subjective health, happiness, and life satisfaction among young adults (18-24 years) in Guyana'. In: *Biomed Res. Int.* 2020.1, p. 9063808.

Warburton Toucher, Michele, Roland Schulze and M Maharaj (2005). 'Is South Africa's temperature changing? An analysis of trends from daily records, 1950-2000'. In: *Climate change and water resources in southern Africa: studies on scenarios, impacts, vulnerabilities and adaptation*, pp. 275–295.

Warwick, R M and S M Turk (Oct. 2002). 'Predicting climate change effects on marine biodiversity: comparison of recent and fossil molluscan death assemblages'. In: *J. Mar. Biol. Assoc. U. K.* 82.5, pp. 847–850.

Watson, Sue-Ann et al. (Jan. 2014). 'Marine mollusc predator-escape behaviour altered by near-future carbon dioxide levels'. In: *Proc. Biol. Sci.* 281.1774, p. 20132377.

Waylen, Kerry A et al. (Aug. 2010). 'Effect of local cultural context on the success of community-based conservation interventions'. In: *Conserv. Biol.* 24.4, pp. 1119–1129.

Wheeler, Andrew, Wouter Steenbeek and Martin A Andresen (Jan. 2018). 'Testing for similarity in area-based spatial patterns: Alternative methods to andresen's spatial point pattern test'. In: *SSRN Electron. J.*

Wiens, J A (1989). 'Spatial Scaling in Ecology'. en. In: *Funct. Ecol.* 3.4, p. 385.

Wijaya, Maleakhi A and Christopher Town (June 2022). 'Deep learning vs keypoint: Which feature set is better for shell recognition?' In: *2022 12th International Conference on Pattern Recognition Systems (ICPRS)*. IEEE, pp. 1–7.

Wilebore, Beccy and David Coomes (Dec. 2016). 'Combining spatial data with survey data improves predictions of boundaries between settlements'. In: *Appl. Geogr.* 77, pp. 1–7.

Wilson, Rebecca J et al. (Dec. 2021). 'Applying computer vision to digitised natural history collections for climate change research: temperature-size responses in British butterflies'. In: *bioRxiv*, p. 2021.12.21.473511.

Wit, Maarten de and Jacek Stankiewicz (Mar. 2006). 'Changes in surface water supply across Africa with predicted climate change'. In: *Science* 311.5769, pp. 1917–1921.

World Bank Group (2024). *Madagascar country climate and development report (CCDR series)*.

WorldPop (2020). *Global 100m Population total adjusted to match the corresponding UNPD estimate*.

Wright, Caradee Y et al. (May 2021). 'Major climate change-induced risks to human health in South Africa'. In: *Environ. Res.* 196.110973, p. 110973.

Yanes, Y and C L Tyler (May 2009). 'Drilling predation intensity and feeding preferences by Nucella (muricidae) on limpets inferred from a dead-shell assemblage'. In: *Palaios* 24.5, pp. 280–289.

Ye, Shan (Apr. 2024). 'Investigating the role of contemporary climate on fossil collecting bias'. In: *Paleontol. Res.* 28.4, pp. 413–425.

Zafari, Sahar et al. (2015). 'Segmentation of Partially Overlapping Nanoparticles Using Concave Points'. In: *Advances in Visual Computing*. Springer International Publishing, pp. 187–197.

Zatelli, Paolo, Nicola Gabellieri and Angelo Besana (Apr. 2022). 'Digitalization and classification of Cesare Battisti's atlas of 1915'. In: *ISPRS Int. J. Geoinf.* 11.4, p. 238.

Zautra, Hall and Kate Murray (July 2008). 'Community development and community resilience: An integrative approach'. In: *Community Dev. (Columb.)* 39.3, pp. 130–147.

Zenda, Mashford and Michael Rudolph (Feb. 2024). 'A systematic review of agroecology strategies for adapting to climate change impacts on smallholder crop farmers' livelihoods in South Africa'. In: *Climate* 12.3, p. 33.

Zhang, Qi et al. (Oct. 2019). 'A shell dataset, for shell features extraction and recognition'. In: *Sci. Data* 6.1, p. 226.

Zhuang, Fuzhen et al. (Jan. 2021). 'A comprehensive survey on transfer learning'. en. In: *Proc. IEEE Inst. Electr. Electron. Eng.* 109.1, pp. 43–76.

Zhuravlyov, A K and K A Grigorian (Dec. 2024). 'Automatic annotation of training datasets in computer vision using machine learning methods'. en. In: *Autom. Doc. Math. Linguist.* 58.S5, S279–S282.

Ziervogel, Gina et al. (Sept. 2014). 'Climate change impacts and adaptation in South Africa: Climate change impacts in South Africa'. In: *Wiley Interdiscip. Rev. Clim. Change* 5.5, pp. 605–620.

Zizka, Alexander et al. (2019). *CoordinateCleaner: standardized cleaning of occurrence records from biological collection databases.*

Zou, Tong et al. (Aug. 2021). 'Recognition of overlapping elliptical objects in a binary image'. In: *Pattern Anal. Appl.* 24.3, pp. 1193–1206.