

Imperial College London
Department of Life Sciences

Evaluating camera trap data for studying spatiotemporal
avoidance and predation between mammal species
across a gradient of habitat degradation

By

Danielle Lisa Norman

A thesis submitted for the degree of Doctor of Philosophy

Statement of Originality and Copyright

Statement of Originality

I declare that the research contained in this thesis is my own work. Ideas and/or research from others have been duly acknowledged following standard referencing procedure. Research carried out for this thesis has led to the following peer-reviewed publication:

Norman, D.L., Bischoff, P.H., Wearn, O.R., Ewers, R.M., Rowcliffe, J.M., Evans, B., Sethi, S., Chapman, P.M. and Freeman, R., 2022. Can CNN-based species classification generalise across variation in habitat within a camera trap survey?. *Methods in Ecology and Evolution*, 14(1), pp.242-251.

This publication is included as Chapter 2 in this thesis, therefore uses 'we/our', while I/my is used in all other chapters. This chapter underwent multiple rounds of comments from reviewers before being published.

Marcus Rowcliffe, Robert Ewers and Robin Freeman provided detailed comments on earlier versions of each chapter.

Copyright

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution – Non Commercial 4.0 International Licence (CC BY-NC). Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Abstract

Ongoing climate change and anthropogenic disturbance negatively impact biodiversity. We must be able to capture, monitor and understand how ecosystem processes, such as, species interactions, are impacted so that we are better positioned to protect against further biodiversity loss.

One inherent challenge is data collection. Camera traps enable us to remotely capture large volumes of data with minimal disturbance to behaviour, but current automated classification methods are unable to generalise well across locations. I investigate the ability of convolutional neural networks to generalise across a gradient of habitat degradation within a camera trap dataset collected in tropical forest. I found generalisability was poor, but was helped by using a detector-classifier combination.

Methods are needed to detect interaction signals from the large volume of camera trap data. Here, I apply statistical methods to test for spatiotemporal avoidance across land-use and disturbance gradients, using the hypothesised avoidance of humans by bearded pigs as a case study. The results did not support the hypothesis, but highlighted the need to understand the data requirements to power the method.

Using an agent-based model to simulate animal movement and generate a camera trap dataset, I test our ability to detect species interactions at varying population and camera trap densities, and interaction strengths. The results showed the difficulty in discerning the type of interaction, and in detecting avoidance behaviour across a range of parameters, despite the large volume of simulated data.

The use of camera trap data for ecological analyses is a growing field, with the potential for transformative analysis, including in our understanding of species interactions. Reliable rapid processing of the images, as well as sensitive methods to detect interactions, are, however, still lacking, and further development is needed to better quantify species' responses to anthropogenic disturbance in order to identify the species most impacted.

Acknowledgements

First, I would like to thank my supervisors – Robin Freeman, Marcus Rowcliffe and Rob Ewers – for their constant faith, support and invaluable discussions.

This research was funded by the Natural Environment Research Council Centre for Doctoral Training in Quantitative and Modelling Skills in Ecology and Evolution [Grant Number: NE/P012345/1]. This work was additionally supported by Sime Darby Foundation funding to the SAFE Project. I'd like to thank Ollie Wearn and the rest of the SAFE project's camera trapping team and taggers for their inputs into the camera trap dataset used throughout.

Undertaking and completing this PhD would not have been possible without the support from my family and friends. First, to my parents, Denise and Graham, thank you for raising me to believe that I can achieve whatever I set my mind to, for providing me with a space to work and for feeding me on Tuesdays! Thank you to my "study buddy" Verity Robinson for being my constant cheerleader, to Emma Bruggenwirth for the pick-me-ups and for keeping me on task, and to Chris Coop for being in my working from home bubble. Thank you to Phil Hayward for the pep talks, to Kirsty Oswald for keeping me sane, and for all the members of my Book Club for very welcome distractions and community: Abigail Hulance, Kristina Bullen, Steph Gautier, Tamsin Marshall, Effie Maguire Ward and Kate Young.

Finally, to my partner, Matthew Udwin: I could not have done this without you. Thank you for your unwavering belief in me, endless support and for picking me back up when it felt overwhelming.

Contents

Statement of Originality and Copyright.....	1
Statement of Originality	1
Copyright.....	1
Abstract.....	2
Acknowledgements.....	3
Contents	4
List of Tables and Figures.....	7
Chapter 1.....	11
1 Introduction.....	11
1.1 Background.....	11
1.1.1 Importance of species interactions.....	12
1.1.2 Methodologies.....	14
1.2 Dataset.....	17
1.2.1 SAFE Project.....	17
1.2.2 Sampling design.....	18
1.2.3 Image processing.....	20
1.2.4 Key findings to date.....	20
1.3 Research Objectives.....	21
1.4 Thesis Outline.....	22
Chapter 2.....	23
2 Can CNN-based species classification generalise across variation in habitat within a camera trap survey?	23
2.1 Abstract.....	23
2.2 Introduction.....	24
2.3 Materials and Methods.....	26
2.3.1 Generalisability.....	27
2.3.2 Bounding boxes.....	28
2.3.3 Metrics.....	29
2.4 Results.....	29
2.4.1 Network and dataset comparison.....	29

2.4.2	Generalisability	31
2.5	Discussion	34
2.5.1	Conclusions	37
Chapter 3	38
3	How are spatiotemporal interactions between mammals in tropical forest impacted by land-use change and disturbance?	38
3.1	Abstract	38
3.2	Introduction	38
3.3	Methods.....	41
3.3.1	Dataset	41
3.3.2	Analysis	41
3.4	Results.....	42
3.4.1	Camera trap events.....	42
3.4.2	Intervals between species.....	44
3.4.3	Land-use category model.....	47
3.4.4	Disturbance level model	48
3.5	Discussion	49
Chapter 4	53
4	Detectability of competitive avoidance and predator-prey interactions in a simulated camera trap survey	53
4.1	Abstract	53
4.2	Introduction	53
4.3	Methods.....	55
4.3.1	Simulation model	55
4.3.2	Interaction scenarios.....	57
4.3.3	Camera traps.....	60
4.3.4	Analysis	61
4.4	Results.....	62
4.4.1	Avoidance and attraction scenarios.....	62
4.4.2	Combined scenario	74
4.5	Discussion	76
Chapter 5	80

5	Discussion and Conclusion	80
5.1	Automated classification of camera trap images	80
5.1.1	Key findings	80
5.1.2	Recent developments	81
5.2	Detection of species interactions.....	81
5.2.1	Key findings	81
5.2.2	Alternative methods for inferring species interactions	83
5.3	Suggested future research.....	86
5.4	Conclusion	87
	Bibliography	88
	Appendices	100
6	Appendix 1: Chapter 2: Supplementary Information	100
6.1	Methods.....	100
6.1.1	Data	100
6.1.2	Machine learning	102
6.1.3	Generalisability – individual disturbance level comparison	107
6.1.4	Bounding boxes.....	109
6.2	Results.....	110
6.2.1	Network comparison.....	110
6.2.2	Generalisability – combinations of disturbance levels	113
7	Appendix 2: Chapter 3 Supplementary Information	114
7.1	Intervals capped at 1 week	114
7.2	Intervals capped at 2 days	115
8	Appendix 3: Chapter 4 Supplementary Information	118
8.1	Influence of interaction zone radius.....	118
8.1.1	Simulated interaction events.....	118
8.1.2	Camera trap events.....	119
8.1.3	Intervals extracted	121
8.2	Applying the Niedballa et al. (2019) threshold	123
8.3	ABM code	124

List of Tables and Figures

Figure 1.1: SAFE camera trap sampling design.	18
Table 1.1: Definitions of disturbance levels within the SAFE project camera trap dataset (Wearn et al., 2017).	19
Table 2.1: All metrics for the best configurations of VGG16, ResNet50 and Inceptionv3 evaluated on the baseline dataset.	29
Figure 2.1: Summary metrics for the Inceptionv3 network.	30
Figure 2.2: Network generalisability.	33
Figure 3.1: Number of camera trap events.	43
Figure 3.2: Distribution of detection times for bearded pigs and humans.	44
Figure 3.3: Number of intervals extracted.	45
Figure 3.4: Distribution of interval durations across land-use categories.	46
Figure 3.5: Distribution of interval durations across habitat disturbance levels.	47
Figure 3.6: Results for the land-use mixed-effect model.	48
Figure 3.7: Results for the habitat disturbance level mixed-effect model.	49
Table 4.1: Parameter values used in simulations, with sources.	55
Figure 4.1: Model schematic for Scenarios 1-3.	60
Table 4.2: Model formulae used in R.	61
Figure 4.2: Interaction rates.	63
Figure 4.3: Number of intervals extracted for Scenarios 1-3.	64
Table 4.3: Median and mean interval duration in days for Scenario 1: CRW.	64
Figure 4.4: Distribution of interval durations for Scenario 1: CRW.	65
Figure 4.5: Distribution of interval durations per pair orientation for Scenarios 1-3.	66
Figure 4.6: Boxplots for Scenarios 2 and 3 at maximum parameter values.	67
Figure 4.7: Boxplots for Scenario 3: one-way attraction.	67
Figure 4.8: Model coefficient estimates for Scenario 1: CRW (control).	68
Figure 4.9: Model coefficient estimates for Scenario 2: one-way avoidance.	69
Figure 4.10: Model coefficient estimates for Scenario 3: one-way attraction.	70
Figure 4.11: Model-predicted interval durations for Scenarios 2 and 3.	71
Figure 4.12: Cap on interval durations: number of intervals extracted.	72
Figure 4.13: Maximum cap on interval durations: interaction zone radii at which the interaction can be detected.	73
Table 4.4: Parameter setups that did not meet the Niedballa et al. requirements of 50 intervals per pair orientation for the analysis.	74
Figure 4.14: Number of intervals extracted for Scenario 4: combined avoidance and attraction.	74

Figure 4.15: Boxplot for Scenario 4: combined attraction and avoidance, with a camera density of 4/km ² .	75
Figure 4.16: Model-predicted interval durations for Scenario 4: combined attraction and avoidance.	76
Table 6.1: Definitions of disturbance levels (Wearn et al., 2017).	100
Table 6.2: Species included within class groups; note all other classes are single species.	101
Table 6.3: Detail of how images for each class were allocated to training, validation and test sets for the network and dataset comparison analyses.	102
Table 6.4: Number of images available and allocated per class to training, validation and test for the baseline dataset.	103
Figure 6.1: Number of (a) events, (b) cameras and (c) images in the training, validation, and test sets for the baseline dataset.	105
Table 6.5: Data augmentation applied during training.	105
Table 6.6: Baseline hyperparameter settings for all networks.	106
Table 6.7: Hyperparameter settings varied during optimisation; note LR: learning rate.	106
Table 6.8: Final learning rates used with optimisation algorithms.	106
Figure 6.2: Training image distribution.	107
Table 6.9: Mean detection rate per camera for each species in each disturbance level for the dataset split at camera level.	108
Table 6.10: Number of images per species class used in the combined disturbance level datasets with and without bounding boxes after passing images through the MegaDetectorv3.	109
Figure 6.3: Top-1 accuracy, Top-5 accuracy and F1-score for settings combinations tried for (a) Inceptionv3, (b) ResNet50 and (c) VGG16.	110
Figure 6.4: Training and validation (a) loss and (b) accuracy after training for 40 epochs for the best setup of each network.	110
Table 6.11: Training and validation metrics for dropout analysis.	111
Figure 6.5: (a) F1-score and (b) recall plotted against number of training images per class in the baseline dataset.	111
Figure 6.6: Generalisability: mean F1-score across the baseline, common only and increased cap datasets. Mean F1-score for the network trained on undisturbed forest (bottom) through to open area (top) and tested on undisturbed forest through to open area (left-right) using the event-level dataset.	112
Figure 6.7: Generalisability: Top-1 accuracy.	113

Table 7.1: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each land-use category, with a maximum cap of 7 days on intervals included in the analysis.	114
Figure 7.1: Results for the land-use mixed-effect model with a cap of 7 days imposed on intervals included.	114
Table 7.2: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each habitat disturbance level, with a maximum cap of 7 days on intervals included in the analysis.	114
Figure 7.2: Results for the habitat disturbance level mixed-effect model with a cap of 7 days imposed on intervals included.....	115
Table 7.3 Median interval duration in days and number of intervals (N) recorded for each pair orientation in each land-use category, with a maximum cap of 2 days on intervals included in the analysis.	115
Figure 7.3: Results for the land-use mixed-effect model with a cap of 2 days imposed on intervals included.	116
Table 7.4: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each habitat disturbance level, with a maximum cap of 2 days on intervals included in the analysis.	116
Figure 7.4: Results for the habitat disturbance level mixed-effect model with a cap of 2 days imposed on intervals included.....	117
Table 8.1: Mean number of avoidance events in Scenario 2, and attraction events in Scenario 3, per simulation for 6, 12 and 24 agents per species at each detection zone radius.	118
Figure 8.1: Total number of camera trap capture events per 100 simulations of the CRW and avoidance scenarios.	119
Figure 8.2: Total number of camera trap capture events per 100 simulations of the CRW and attraction scenarios.	119
Figure 8.3: Total number of camera trap events per 100 simulations of Scenario 4: combined avoidance and attraction.....	120
Figure 8.4: Total number intervals extracted from 100 simulations of the CRW and avoidance scenarios.	121
Figure 8.5: Total number of intervals extracted per 100 simulations of the CRW and attraction scenarios.	121
Figure 8.6: Total number of intervals extracted from 100 simulations of Scenario 4: combined avoidance and attraction.....	122
Figure 8.7: Model coefficient estimates for Scenario 2: one-way avoidance with setups that did not meet the Niedballa et al. threshold of 50 records per pair excluded.	123

Figure 8.8: Model coefficient estimates for Scenario 3: one-way attraction with setups that did not meet the Niedballa et al. threshold of 50 records per pair excluded..... 124

Chapter 1

1 Introduction

1.1 Background

Changes in land use and climate affect the distribution, abundance and behaviour of species globally (Wearn et al., 2017; Gaynor et al., 2018; WWF, 2022). To mitigate negative effects of anthropogenic activity on biodiversity, and to protect vulnerable ecosystems, a good understanding of how species respond to environmental change is required. In particular, species interactions have been identified as a potential mediator of impacts and are therefore vital to monitor (Rahman & Candolin, 2022).

Detailed, observational studies can provide important insight on a small-scale, but effective monitoring over large areas and long time periods is necessary to allow patterns, and any changes in them, to be captured. With their increasing affordability, camera traps provide the opportunity to continuously survey large areas remotely. In addition, they are able to capture mammal species that are often scarce, shy, elusive or nocturnal, so are difficult to observe directly. Analytical techniques have been developed for estimating key community, population and behavioural parameters from camera trap imagery, yielding insight into distribution and abundance, as well as activity patterns, movement and interactions between species (Caravaggi et al., 2017; Frey et al., 2017; Kellner et al., 2022).

One obstacle to the use of camera trap data, however, is in labelling the huge number of images produced. Recent advances in machine learning have opened up methods for automating the time-consuming tagging process (Norouzzadeh et al., 2018), but obstacles remain in having sufficient data to train reliable, transferable classifiers (Beery et al., 2018). With more accessible automated image analysis in the future, the widespread, long term use of these techniques could catalyse a step-change in our understanding of the status of mammal populations and the processes driving change.

In the following sections of this chapter, I will therefore outline the literature pertaining to:

- the importance of species interactions;
- methodologies, including:
 - camera trap surveys and the analytical methods used to infer or to model species interactions;
 - automated classification of camera trap imagery; and
- the dataset used throughout this thesis.

1.1.1 Importance of species interactions

Biodiversity is declining, with more species at risk of global extinction than ever before (IPBES, 2019). The driving forces behind the degradation of ecosystems and decline in biodiversity include the overexploitation of plants and animals, climate change, pollution and invasive species (WWF, 2022). The most important driver of biodiversity loss has been identified as land-use change (WWF, 2022). The effects of anthropogenic activity, such as land-use change, on biodiversity can be monitored via ecological parameters including changes in abundance, species richness and range (Gido et al., 2019; Jetz et al., 2019). Underlying these more readily measurable quantities, are the processes causing them, which are more difficult to observe and quantify (Frey et al., 2017). Anthropogenic activity can trigger changes in animal behaviour, such as changes to daily activity patterns and the occurrence and frequency of species interactions (Frey et al., 2017). Shifts in behaviour, such as in species interactions, can thus be considered “an early-warning system” (Caravaggi et al., 2017) by signifying changes at the process level before they can be seen at the population level. Monitoring changes in behaviour, therefore, may be vital to our understanding of how ecological communities respond and adapt to environmental stressors, and thus how anthropogenic activity results in reduced biodiversity.

Species interactions have been identified as a potential mediator of how species respond to environmental change. Environmental change alters species interactions, species ranges and, consequently, co-occurrence (Morales-Castilla et al., 2015; Rahman & Candolin, 2022). A recent study of terrestrial mammals situated in tropical forests worldwide found that interactions between species were found to mediate distributional dynamics, with ecological similarity being linked to colonisation and extinction dynamics (Beaudrot et al., 2019). In a study of African mammals, Buschke et al. (2015) found that while climate may determine local persistence of species, species interactions drive more general patterns of co-occurrence. This is in agreement with a global study that found predator-prey interactions, in conjunction with environmental effects, to be an important driver of large-scale diversity gradients (Sandom et al., 2013). In particular, predator richness variance in their model was most explained by prey bottom-up effects, while prey richness was most strongly associated with productivity and climate, with predator top-down effects significant (Sandom et al., 2013). Conversely, Greenville et al. (2017) found that under future climate change scenarios, while vegetation cover and productivity are likely to decline, top-down suppression from introduced predators is predicted to have the strongest negative impact on prey populations (Greenville et al., 2017). While it is clear that both environmental effects and species interactions influence the persistence and range of species, as well as diversity gradients, understanding the effect

of habitat composition on species interactions and ecosystem functioning is needed to inform ecosystem management (Casula et al., 2017).

In areas of land use change, species must respond to changes in their habitat and resource availability, which can cause or reduce competition. Studies have found increased nocturnality in animal species in response to human activity and along land-use gradients, including a switch from diurnal to nocturnal activity patterns (Gaynor et al., 2018; Davison et al., 2019). One US study established that avian nest predation was increased in areas of habitat loss; as natural gas development intensity increased, rodent predator abundance and/or activity increased, causing higher nest predation (Hethcoat & Chalfoun, 2015). Human-induced environmental change was therefore shown to be altering the nest predator assemblage, which impacted the demography of local prey communities through predation. Species interactions and forest fragmentation were both found to be important determinants of carnivore occupancy in a recent study into carnivore distribution along an urban gradient (Parsons et al., 2019). The mixture of positive and negative interactions seen, with some dependency on green space availability, led the authors to suggest that fragmentation leads to higher levels of spatial interaction (Parsons et al., 2019). Since species interactions appear to be an important mechanism through which species respond to environmental change, monitoring changes in interactions in response to disturbance could be crucial for our understanding of the wider ecological impacts.

Invasions of non-native species, which are often caused or facilitated by anthropogenic activity (Simberloff et al., 2013), are another threat to biodiversity that acts by altering species interactions, including predation, competition and disease transmission. In particular, invasive mammalian predators can be very destructive, and often endanger island species that are evolutionarily distinct (Doherty et al., 2016; Spatz et al., 2017). It is vitally important to understand and mitigate the impact of these predators to conserve diversity and slow biodiversity decline (Doherty et al., 2016).

Interactions between species have been described as the “architecture of biodiversity” (Bascompte & Jordano, 2007). Communities are made up of populations of species co-existing, so any change to the behaviour of one species can have a knock-on effect on the rest of the community. Similarly, extinction of a species can cause unpredictable secondary extinctions that cascade throughout an entire ecosystem (Ebenman & Jonsson, 2005). This highlights the importance of understanding the interactions between species so as to be in the best position to mitigate further loss of species in the face of current threats to biodiversity.

In mammals in particular, community structure is predominantly shaped by antagonistic interactions, such as competition, harassment and predation (Palomares & Caro, 1999;

Sinclair et al., 2003). Where habitat is fragmented, and populations confined, competitive interactions among species that were previously spatially segregated can be observed (Durant, 1998). Species might then change their behaviour in order to minimise or avoid competitive encounters through altering their diet, daily activity patterns or their use of the habitat (Carothers et al., 1984; Palomares & Caro, 1999; Vanak et al., 2013). Predation can be a critical component of an ecosystem, with apex predators acting as a keystone species that regulates the abundance and distribution of its prey species. In temperate ecosystems, loss of predators has been shown to lead to trophic release, whereby the abundances of prey species subsequently increase (Ripple & Beschta, 2012). In contrast, no evidence of trophic release was found in tropical forest, with predator-prey relationships appearing to be weak or positive (Brodie & Giordano, 2013).

Understanding how and why interactions between species occur will enable us to build a better picture of how they are impacted by other ecological processes and inform mechanism-based predictive models (Poisot et al., 2015). Given the important role mammals play in ecosystems and how influential antagonistic interactions between them are in shaping communities, our ability to capture spatiotemporal avoidance and predation between mammal species is the focus of this thesis. For this, we require robust methods that enable us to identify and quantify interactions.

1.1.2 Methodologies

Survey methods for capturing species interactions include, for example, direct observation, DNA metabarcoding or the use of tagging with GPS collars. DNA metabarcoding to analyse the diets of species can provide important information on predation (Rytönen et al., 2019), but this approach is limited by the necessity to locate and collect samples. Direct observation can be useful for recording pollination or other frequent, short-term interactions that occur over a small survey area. GPS devices have opened up the ability to capture fine-scale movement data for species that were previously difficult to study, such as aquatic species, migratory songbirds and wide-ranging migratory mammals (Hebblewhite & Haydon, 2010). The cost of collars, the necessity to capture and/or sedate individuals to tag them, and difficulties in retrieving the devices can, however, result in weaker study design and small sample sizes (Hebblewhite & Haydon, 2010), as well as short-term effects on activity, behaviour and stress of the animals (Stabach et al., 2020).

1.1.2.1 Camera trap surveys

With advancements in technology and reduction in price, camera traps have become a popular tool for monitoring global biodiversity (Burton et al., 2015; Caravaggi et al., 2017). Camera trap surveys have the advantage of being non-invasive, they can provide measurements at

high resolution (Rowcliffe et al., 2016), and produce vast quantities of data. They have also been identified as a potential tool for creating a global monitoring network (Steenweg et al., 2017). Camera traps have been used in a variety of studies, including wildlife distribution, abundance and community structure (Burton et al., 2015). They can also help quantify animal behaviour, such as foraging, scent marking, daily activity patterns and interactions (Caravaggi et al., 2017). Techniques have been developed to estimate behavioural metrics, such as travel speed and day range, from camera trap data, which inform our understanding of processes such as energy use, foraging success, and disease transmission (Rowcliffe et al., 2016).

Methods that rely on indirectly observed or passive collection of data (e.g. camera traps) to infer species interactions face analytical difficulties. Until recently, camera trap surveys have generally been used to explore hypotheses in a spatial context, with spatial analysis of species interactions involving examining species co-occurrence patterns (Frey et al., 2017). Inference of species interactions from co-occurrence analyses has been criticised by Blanchet et al. (2020) who argue that the complexity of ecosystems blurs the link such that spatial associations alone are poor proxies for interactions; among the challenges outlined, they note the importance of accounting for temporal variations and its impact on species and their interactions (Blanchet et al., 2020).

The use of time-stamped camera trap images to inform temporal analyses, such as variation in activity patterns and partitioning along the temporal niche axis, has been highlighted as critical to developing our understanding of population and community dynamics (Frey et al., 2017). To date, few studies have empirically quantified how external factors may influence temporal niche partitioning (Frey et al., 2017). In their review, Frey et al. (2017) suggest that to investigate how interacting variables influence species segregation along the temporal niche axis, coordinated experiments distributed across a range of stressors and community compositions are needed.

Analyses using combined spatiotemporal data could offer further insight (Cusack et al., 2017). Spatiotemporal analyses assessing coexistence of felids in tropical forest using camera traps have found evidence of temporal partitioning, with spatial habitat use mostly determined by prey availability (Haidir et al., 2018; Santos et al., 2019). Others have identified avoidance behaviour (Ross et al., 2013; Ramesh et al., 2017) or discovered adaptability in behaviour across gradients of resource availability (Karanth et al., 2017). Two recent studies have combined occupancy models with a continuous-time detection process to detect spatial and temporal interactions between predator and prey species (Kellner et al., 2022), and between two competing carnivores along an urbanisation gradient (Parsons et al., 2022).

A framework for exploring spatiotemporal interactions in camera trap data has recently been proposed, alongside a comparison of methods for detecting interactions in camera trap data; linear models were identified as the most suitable approach for detecting spatiotemporal avoidance and a permutation test performed on daily activity patterns for assessing temporal segregation (Niedballa et al., 2019). The application of this framework to real-world data has yet to be explored.

1.1.2.2 Species identification in camera trap images

The increased uptake of camera trap surveys and their voluminous output has resulted in bottlenecks in image processing. It can take many hours for researchers to label each individual image with the species present. Thus citizen scientists have been recruited to ease the burden, with a reported accuracy of 96.6% compared with experts (Swanson et al., 2015). Researchers have also turned to machine learning techniques to improve the speed with which camera trap images can be processed. In a comparison of convolutional neural network architectures in the task of identifying species from the 3.2 million-image Snapshot Serengeti dataset, an overall accuracy of >93.8% was achieved (Norouzzadeh et al., 2018). When restricted to only images the network was confident of having categorised correctly, this rose to 99.3%, which was calculated to equate to a saving of >8.4 years of manual human labelling time (Norouzzadeh et al., 2018). A more recent study has since achieved an accuracy of 97.6% (Tabak et al., 2019). The number of tools available, and their application to species identification tasks, is predicted to continue to increase (Wäldchen & Mäder, 2018).

One current issue with the applicability and impact of these machine learning techniques is their lack of ability to generalise well across background scenes. Previous studies have highlighted a drop in performance when the network is used to classify images from camera trap locations unseen during training (Beery et al., 2018; Tabak et al., 2019). Variation in background scenery and vegetation, lighting, camera position or average distance of subject from camera, for example, can all impact on performance of trained classifiers (Beery et al., 2018). Many applications of machine learning to classification thus far have had a particular geographic focus (Weinstein, 2018), but in order for these techniques to be widely applicable and impactful, we require architectures that can be used by multiple researchers on different datasets, ideally without having to perform the time-consuming training on each new dataset.

Proposed methods for improving generalisability include the use of object detectors to locate animals within images and therefore provide cropped, targeted images to pass to the network (Beery et al., 2018). Another study found that generating artificial images for rare species can improve generalisability for these species, but this approach is not practical for most ecologists since it requires the assistance of a graphic artist (Beery, Liu, et al., 2019). More generally,

incorporating metadata alongside images into a multi-input network has demonstrated improved overall accuracy, and it is suggested this could improve generalisability (Terry et al., 2020).

The impact of habitat degradation on network generalisability has not been considered to date. With ongoing anthropogenic activity leading to habitat loss, fragmentation, and degradation, it is important to understand how these classifiers perform on camera trap datasets from habitats undergoing change. Classifiers need to be robust to changes in image background to facilitate long-term monitoring of these habitats. Otherwise, misclassifications as a result of poor generalisability have the potential to impact on the conclusions drawn from any subsequent ecological analyses.

1.1.2.3 Agent-based models

Species interactions are often difficult to directly study due to logistical difficulty in achieving direct observations and experimental challenges (Smith et al., 2020). The use of computational models to simulate interacting species therefore allows us to generate data that is otherwise challenging to capture. These models also provide the ability to perform in-silico experiments and to test scenarios before they are applied to the real world such that management practices can be optimised, and the potential impacts of disturbances can be best mitigated.

Agent-based models (ABMs) are a popular tool in ecology for modelling the movement and behaviour of individual 'agents' around a study space. They have been used to model ecosystem processes such as seed dispersal distance and the impact of human-induced reductions on mammal movement in tropical forest (Tucker et al., 2021), and the relationship between habitat features and performance on predator-prey event outcomes (Wheatley et al., 2020). ABMs have also been used to study the marine environment including the impact of acoustic disturbance on swimming trajectories, bioenergetics and population size of marine species (Mortensen et al., 2021).

Using ABMs, we can perform experiments in a controlled environment to gain insight into the processes at play, and to generate data that we can use to reliably test methods.

1.2 Dataset

1.2.1 SAFE Project

The data used throughout this thesis is taken from the Stability of Altered Forest Ecosystems (SAFE) Project situated in Sabah, Malaysian Borneo (Ewers et al., 2011). The project was established to generate a broad understanding of the ecological impacts of tropical forest

modification (Ewers et al., 2011). To this aim, the SAFE project study area comprises forest fragments in old-growth forest within the Maliau Basin Conservation Area, logged forest within the Kalabakan Forest Reserve, as well as oil palm plantations straddling the Kalabakan Forest Reserve boundary.

As a result of repeated logging in the study area, there are a variety of habitat types within the logged forest fragments, ranging from grassy open areas or low scrub vegetation to lightly logged forest (Wearn et al., 2013).

1.2.2 Sampling design

For the camera trap surveys at SAFE (Figure 1.1), a clustered hierarchical nested sampling design was used (Wearn et al., 2016). Individual sampling points were clustered into rectangular sampling plots of 48 (4 x 12) points separated by 23 m, each covering an area of 1.75 ha (Wearn et al., 2016). Three to six plots were then clustered into blocks (approximately 25 ha), with three to four blocks per land use category: 13 plots (in 4 blocks) were situated in old-growth forest, 24 plots (in 4 blocks) in logged forest and 9 plots (in 3 blocks) in oil palm plantations (Wearn et al., 2017). Separation distances between plots (170-290 m) and blocks (0.6-3 km) was similar across land uses, although the arrangement of blocks differed in logged forest compared with old-growth forest and oil palm (Wearn et al., 2016).

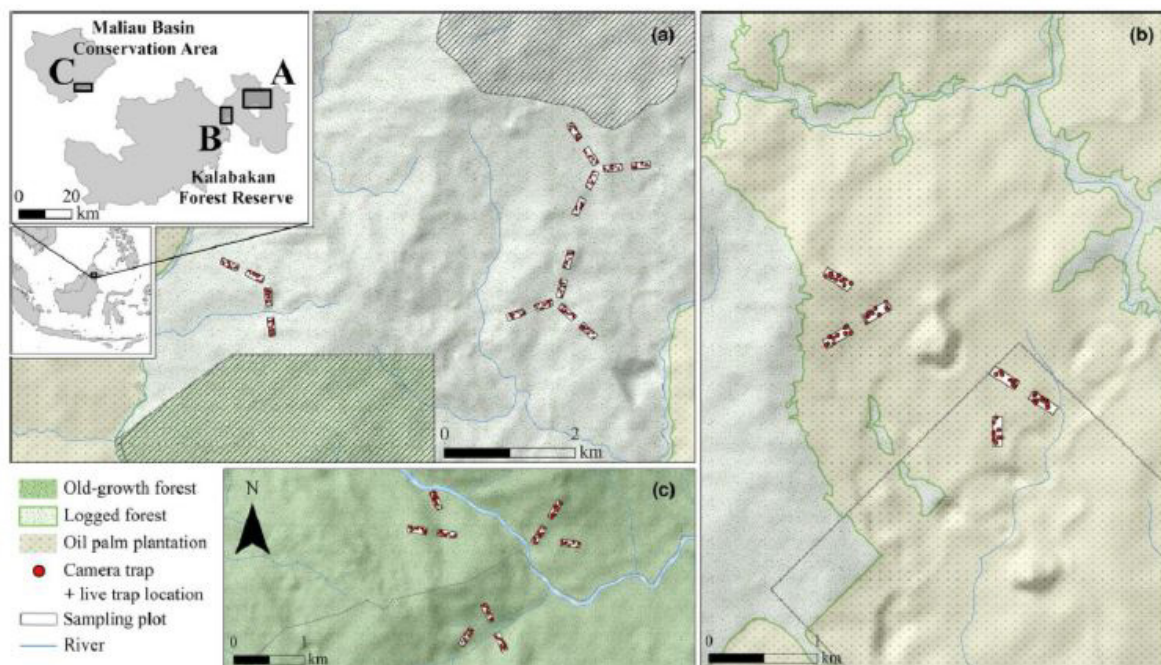


Figure 1.1: SAFE camera trap sampling design. Sampling points are shown across (a) logged forest, (b) oil palm and (c) old-growth forest. Shaded areas (in (a)) lie outside the forest reserve. Taken from (Wearn et al., 2019).

The majority of the camera trap images were collected between May 2011 and April 2014, with most plots (40 of 46) being sampled in multiple years (Wearn et al., 2017). Following April 2014, plots were surveyed once per year, but not all plots were sampled in all years (Davison et al., 2019). Within each plot, a random subset of the 48 sampling points was selected for camera-trapping, giving 681 points sampled in total. Remotely operated digital cameras (Reconyx HC500) were deployed within 5 m of each of these selected points, usually set at a height of 30 cm, with some set higher where the location demanded (Wearn et al., 2013). During each survey, the camera traps were deployed for at least 30 days, culminating in a total of 32, 542 camera trap nights (Davison et al., 2019). The cameras captured 10 photos in quick succession when triggered (Wearn et al., 2013). No lure or bait was used, and disturbance to vegetation was minimised. The survey collected 753, 442 images in total.

For their assessment of mammal abundance within the study area, Wearn et al. (2017) recorded a habitat disturbance score for a 5 m radius around each sampling point. The habitat disturbance was scored on a scale from low to high disturbance, with five categorical levels (Table 1.1). These levels correspond to environmental variation due to the effects of logging and subsequent regeneration.

Table 1.1: Definitions of disturbance levels within the SAFE project camera trap dataset (Wearn et al., 2017).

Disturbance level	Definition
Undisturbed forest	Dominated by old-growth dipterocarps. High, continuous canopy with sparsely-vegetated understorey. Unlogged, with little recent disturbance evident.
Disturbed forest	Mostly pioneer tree species (typically <i>Macaranga</i> species), but some old-growth dipterocarp species may be present. Discontinuous canopy. Lower intensity of logging or natural disturbance.
Heavily disturbed forest	High scrub or dense understorey layer (typically with vines and <i>Dinochloa</i> climbing bamboo species), with a low, heavily-broken canopy layer (< 20 m). Possibly some large, isolated trees (> 20 m). Intensively-logged area or large gap disturbance.
Herbaceous scrub	Dominated by herbs (typically <i>Zingiberaceae</i>), vines and shrubs, with no trees > 3 m in height (except oil palm <i>Elaeis guineensis</i>). Typically representing secondary re-growth from clear-felling, or large gaps due to landslides.
Open area	Open area. Dominated by grasses and small shrubs (< 1 m in height). Typically on logging roads or old log landing areas.

1.2.3 Image processing

Images were reviewed in Adobe Photoshop Lightroom (Adobe Inc., San Jose, California, USA), with keyword tags added to record the species, as well as the capture of a new individual (Wearn et al., 2017). For bearded pigs, tags to denote individual or group, as well as the number and age class of the individuals present were added (Davison et al., 2019). Metadata was extracted using ExifTool (Harvey, 2017).

1.2.4 Key findings to date

Results from SAFE to date highlight the importance of remnant forest patches among agricultural landscapes to native mammal species (Daniel et al., 2022). Some mammals are able to persist in logged forest, increasing in relative abundance by 28% from primary forest, but the opposite is seen in oil palm plantations with a 47% decrease compared to forest (Wearn et al., 2017). Responses among mammal species vary, however, with small mammals shown to increase much more proportionately (169%) than large mammals (13%) from old-growth to logged forest, as well as omnivores and insectivores increasing more than other trophic guilds (Wearn et al., 2017). Similarly, responses to logging in terms of species richness have been shown to vary, with an increase in richness in small mammals at all spatial grains in logged forest compared with old-growth forest, while large mammals showed reduced richness in logged forest at the grain of individual sampling points, but no change at the land-use level (Wearn et al., 2016). Within mammal communities, large mammal communities were found to be more heterogenous at coarse spatial grains in old-growth forest, while small mammal communities were more homogenous, with the reverse being observed in logged forest (Wearn et al., 2016).

In terms of species interactions, Wearn et al. (2019) found that the prevalence of spatial avoidance decreased along the land-use gradient; their co-occurrence analysis detected only a few instances of spatial avoidance in any land-use, however, with 13 instances detected in old-growth forest. Of these, three were among congeneric pairs: *Maxomys* rats (*M. surifer* and *M. rajah*), the greater and lesser (*T. kanchil*) mouse-deer, and the thick-spined and Malay porcupines (*Hystrix crassispinis* and *H. brachyura*) (Wearn et al., 2019). The authors concluded competition may be occurring, but it is unlikely to be a dominant assembly process in the study system. It was suggested that since the study site consists of a trophically diverse mammal assemblage, species are able to avoid competition through resource-partitioning. The species could also be segregated vertically in space, or temporally, such that the co-occurrence analyses would be unable to detect such avoidance (Wearn et al., 2019).

In a separate investigation of mammal behaviour at SAFE using camera trap data, shifts in activity pattern between primary and logged forest were observed in six species (Chapman et

al., 2018). Sambar deer and bearded pigs were found to shift from diurnal in old growth forest to more nocturnal in logged forest. The long-tailed porcupine showed a significant shift in activity peak from the early hours of the morning in primary forest to late evening in logged forest (Chapman et al., 2018). Activity in greater mouse-deer was strongly crepuscular in old growth forest, but more constant throughout the day in logged forest. An earlier decline in activity during dusk in logged than in old growth forest was seen in red muntjac, while a slightly increased activity level before dawn, which increased more slowly after dawn in logged forest was seen in yellow muntjac (Chapman et al., 2018).

Bearded pig behaviour between old-growth and logged forest, and oil palm plantations, was further analysed by Davison et al. (2019). In a study of activity patterns using camera trap data, the authors found bearded pigs shift from diurnal in old-growth forest, to crepuscular in logged forest and nocturnal in oil palm plantations. The shift to crepuscular in logged forest was attributed to both an avoidance of humans and the higher midday temperatures in the low canopy cover, less sheltered logged forests (Davison et al., 2019). The nocturnal activity in oil palm probably represented foraging raids from nearby forests (Davison et al., 2019). The lack of early morning activity was suggested to be likely avoidance of human workers in oil palm plantations, who are consistently active in the mornings. Bearded pigs could also be avoiding free-roaming dogs, whose core activity patterns overlapped minimally with bearded pigs (Davison et al., 2019).

The SAFE Project provides valuable insight into the impact of ongoing human-mediated disturbance in tropical forests. Comprising communities that are particularly vulnerable to extinction, it is especially important to understand the processes underlying the biodiversity changes in these areas.

1.3 Research Objectives

1. To assess the performance of established convolutional neural network architectures, identifying the best network for classifying images within the SAFE camera trap dataset, and to investigate the generalisability of the chosen network within the dataset.
2. To test our ability to detect spatiotemporal avoidance of one species by another in a real-world camera trap dataset used for long-term monitoring, and to assess whether this behaviour varies across land-use categories and disturbance levels. The hypothesised avoidance of humans by bearded pigs in oil palm plantations will be used as a case study.
3. To investigate the impact of interaction strength, population density and camera trap density on our ability to detect spatiotemporal avoidance and predatory stalking, with the aim of

identifying thresholds in these parameters, beyond which competitive avoidance and/or predation can be captured.

1.4 Thesis Outline

In Chapter 2, I focus on one of the ongoing issues with the application of automated classification for species identification in camera trap images: generalisability. I train and evaluate performance of three convolutional neural networks (CNNs) on the SAFE camera trap dataset, then assess the ability of the best-performing CNN to generalise across the habitat disturbance levels within the dataset.

In Chapter 3, I move to considering mammalian species interactions, specifically whether I can detect suspected spatiotemporal avoidance of humans by bearded pigs. I apply a method that was tested on simulated data and identified as optimal for detecting spatiotemporal avoidance of one species by another within camera trap data. I use the results to both test my hypotheses and to assess the application of this method to real-world, as opposed to simulated, data.

In Chapter 4, I construct an agent-based model comprising individuals from two species. The model allows me to simulate different interaction behaviours: competitive or predator avoidance, as well as stalking of prey. A camera trap survey is then simulated by modelling camera traps placed at random within the study area such that they are triggered by nearby agents. Using a control simulation, I evaluate our ability to detect the species interactions using the method applied in Chapter 2.

Chapter 5 contains a discussion of the key findings in context, suggestions for future research and final conclusions.

Chapter 2

2 Can CNN-based species classification generalise across variation in habitat within a camera trap survey?

Note: this work has been published as: Norman, D.L., Bischoff, P.H., Wearn, O.R., Ewers, R.M., Rowcliffe, J.M., Evans, B., Sethi, S., Chapman, P.M. and Freeman, R., 2022. Can CNN-based species classification generalise across variation in habitat within a camera trap survey?. *Methods in Ecology and Evolution*, 14(1), pp.242-251.

2.1 Abstract

Camera trap surveys are a popular ecological monitoring tool that produce vast numbers of images making their annotation extremely time-consuming. Advances in machine learning, in the form of convolutional neural networks, have demonstrated potential for automated image classification, reducing processing time. These networks often have a poor ability to generalise, however, which could impact assessments of species in habitats undergoing change.

Here, we: (i) compare the performance of three network architectures in identifying species in camera trap images taken from tropical forest of varying disturbance intensities; (ii) explore the impacts of training dataset configuration; (iii) use habitat disturbance categories to investigate network generalisability; and (iv) test whether classification performance and generalisability improve when using images cropped to bounding boxes.

Overall accuracy (72.8%) was improved by excluding the rarest species and by adding extra training images (76.3% and 82.8%, respectively). Generalisability to new camera locations within a disturbance level was poor (mean F1-score: 0.32). Performance across unseen habitat disturbance levels was worse (mean F1-score: 0.27). Training the network on multiple disturbance levels improved generalisability (mean F1-score on unseen disturbance levels: 0.41). Cropping images to bounding boxes improved overall performance (F1-score: 0.77 vs 0.47) and generalisability (mean F1-score on unseen disturbance levels: 0.73), but at a cost of losing images that contained animals which the detector failed to detect.

These results suggest researchers should consider using an object detector before passing images to a classifier, and an improvement in classification might be seen if labelled images from other studies are added to their training data. Composition of training data was shown to be influential, but including rarer classes did not compromise performance on common classes, providing support for the inclusion of rare species to inform conservation efforts.

These findings have important implications for use of these methods for long-term monitoring of habitats undergoing change, as they highlight the potential for misclassifications due to poor generalisability to impact subsequent ecological analyses. These methods therefore need to be considered as dynamic, in that changes to the study site would need to be reflected in the updated training of the network.

2.2 Introduction

Camera traps have become an increasingly popular survey tool among ecologists and conservationists, being used in a variety of studies, including of wildlife distribution, abundance, occupancy, behaviour and community structure (Burton et al., 2015). Their two biggest advantages are that they sample a relatively broad-spectrum of wildlife, making them effective for monitoring species richness, and that they can operate night-and-day for months at a time, meaning that they can produce useful data on even the rarest species (Wearn et al. 2019). They also produce thousands, or in some cases, millions, of images for analysis.

Sifting out empty images and tagging images of animals can be a very time-consuming task for researchers. Although workflow efficiency and task complexity are probably hugely variable in 'real-world' settings, our experience is that an operator can process on the order of 1,000-5,000 images per day (assuming a basic task of tagging species and counting individuals). Recent advances in machine learning have seen the application of neural networks to this task to reduce the burden on researchers and reduce processing time (Swanson et al., 2015; Beery et al., 2018; Norouzzadeh et al., 2018; Tabak et al., 2019; Willi et al., 2019). In the largest comparison of machine learning architectures for the task of identifying species to date – based on the 3.2 million-image Snapshot Serengeti dataset – an overall accuracy of 93.8 % was achieved (Norouzzadeh et al., 2018). When restricted to only images the network was confident of having categorised correctly, this rose to 99.3%. Overall, automating the task of identifying species could have saved over 8.4 years of manual human labelling time if implemented from the outset (Norouzzadeh et al., 2018). More recent studies have achieved even higher accuracies of 95.6% (Schneider et al., 2020) and 97.6% (Tabak et al., 2019). In their review, Waldchen and Mader (2018) predicted that the number of tools available, and their application to species identification tasks, will continue to increase in the future.

These high accuracy results are impressive, but do not provide the full picture since they represent performance when the network is trained and tested on images from the same camera trap locations. When networks are tested on images from camera locations unseen during training, performance invariably drops; the networks do not generalise well. Previous studies have reported varying accuracies in this case: 68.7% (Schneider et al., 2020) and 59% (Beery et al., 2018) when tested on unseen camera locations from within the camera trap

dataset, and 82% (Tabak et al., 2019) when tested on camera locations from an alternative dataset. This drop in performance could be due to variables such as changes to the background scenery, lighting, camera position or average distance of subject from camera. Performance can also be impacted by variation in the distribution and density of species recorded by each camera (Wei Koh et al., 2021). The issue of poor generalisability is not unique to automated classification of camera trap images, however. In the related context of acoustic detection in birds, networks generalised poorly to new conditions including differing species balances, noise conditions, or recording equipment (Stowell et al., 2019). Similarly, a 14.4 % drop in marine mammal classifier accuracy occurred when testing on whistle data from a different region than that trained on (Erbs et al., 2017).

Many applications of machine learning to classification thus far have had a particular geographic focus (Weinstein, 2018), but in order for these techniques to be widely applicable and impactful, architectures are required that can be used by multiple researchers on different datasets, ideally without having to perform the time-consuming network training at each new location (Wearn et al., 2019). In a world increasingly impacted by anthropogenic activity resulting in habitat degradation and fragmentation, it is also important that we have classifiers that are robust to changes in image background to facilitate long-term monitoring of habitats undergoing change. Otherwise, when conducting analyses using images classified by a network trained on images from pristine habitat, we risk drawing wrong conclusions if the new habitat has been altered to such an extent that the image backgrounds have changed. The impact of habitat degradation on network generalisability has not been considered to date.

Here, we compare the performance of three established architectures to identify species in camera trap images taken from undisturbed and disturbed tropical forests in Borneo. We specifically aim to explore the extent to which a network is able to generalise, which we achieve by splitting the dataset into an environmental gradient of varying levels of habitat disturbance generated by historical logging. Our goals are to: (1) Assess the performance of established architectures and identify the best network for classifying images within our dataset; (2) Explore the impacts of training dataset configuration on overall performance, specifically restricting the data to common species only or increasing the number of images per species class included; (3) Use the disturbance level categories attributed to the camera trap locations to investigate the generalisability of the chosen network within our dataset; and (4) Compare generalisability performance when images are cropped to bounding boxes. The results of this study will inform the robust application of automated image classification for monitoring biodiversity in habitats undergoing change.

2.3 Materials and Methods

The dataset comprises camera trap images taken from the Stability of Altered Forest Ecosystem (SAFE) Project (Ewers et al., 2011) in Malaysian Borneo, a subset of which form the open access BorneoCam dataset. Camera trapping took place between May 2011 and March 2018, following the sampling procedure laid out in Wearn et al., 2013. Data was originally collected under approval from the Government of Malaysia, with the following permit numbers: Economic Planning Unit 40/200/19/2656; Maliau Basin Management Committee MBMC/2010/15, and Sabah Biodiversity Council JKM/MBS.1000-2/3 (84), JKM/MBS.1000-2/2 JLD.7 (51), JKM/MBS.1000-2/2 JLD.5 (142), JKM/MBS.1000-2/2 JLD.4 (192) and JKM/MBS.1000-2/2 JLD.3 (125). This dataset makes an ideal case study since it represents a realistic ecological dataset, in terms of size and in level of imbalance between classes, and it comprises a variety of habitat disturbance levels. These images have previously been used to inform analyses of mammalian species abundance (Wearn et al., 2017), diversity (Wearn et al., 2016) and behaviour (Davison et al., 2019) across a gradient of land-use comprising unlogged forest, logged forest and oil palm plantations. Forest quality at the locations of individual camera traps has been quantified into a five-step disturbance scale: 1: undisturbed forest; 2: disturbed forest; 3: heavily-disturbed forest; 4: herbaceous scrub; and 5: open area (Wearn et al., 2017). (Full descriptions of disturbance categories are provided in Table 1.1.)

The total raw data consisted of 753,442 images from 681 camera deployments. To construct a dataset of labelled images, untagged images were removed, as well as images captured during the setup process or a camera malfunction, or containing non-target (reptile or invertebrate) or multiple species. Empty images were also removed since we were interested in classification of species and so made the explicit assumption that the step of separating images into empty and non-empty had previously taken place. Camera traps (Reconyx HC500) were programmed to take a rapid burst of 10 images, termed a capture event. Image labelling for this dataset is at the level of images, rather than events. All non-empty images from a given event were allocated together to either the training or test dataset. A small proportion (0.05 % of images), where the event grouping was not recorded in the metadata, were discarded. This reduced the dataset to 378,000 images from 640 deployments. Both day and night images were included since previous studies have found this had little effect on performance (Norouzzadeh et al., 2018; Tabak et al., 2019).

A minimum of 40 images or four capture events, per species class was imposed. To include as many species as possible, species that fell below this threshold were grouped together with related species, e.g. Hose's civet *Diplogale hosei* images were included within the banded civet *Hemigalus derbyanus* class. These group classes comprised between 2 and 15 species

(detail provided in Table 6.2). To limit imbalance within the dataset, and to reduce computation time and resources required, the maximum number of images per class was restricted to 5,000. We investigated an increased maximum per class (Figure 2.1) and found that it improved Top-1 accuracy while having a small impact on mean F1-score, suggesting that it resulted in more bias towards common classes. A 90:5:5 split for training, validation and test sets was used following Willi et al. (2019), and to ensure matching distributions across classes within the three sets (Appendix 1: Table 6.4, Figure 6.1: Number of (a) events, (b) cameras and (c) images in the training, validation, and test sets for the baseline dataset. Figure 6.1). This resulted in training, validation and test sets consisting of 76,637, 4,290 and 4,309 images, respectively, each containing images from 51 classes. Images were resized to 256 x 256 pixels before passing to the neural networks. Data augmentation was also performed, consisting of random shearing, horizontal flipping, cropping and brightness modification (Appendix 1: Table 6.5). This is commonly carried out in image classification problems to bolster training data and prevent overfitting (Krizhevsky et al., 2012; Beery et al., 2018).

To identify the best network for our dataset we compared performance of three architectures: VGG16 (Simonyan & Zisserman, 2015), Inceptionv3 (Szegedy et al., 2016) and ResNet50 (He et al., 2016). In each case, the network was pre-trained on ImageNet, which is a large database of quality-controlled, human-annotated images, including animal classes, and is commonly used to pre-train networks for image classification tasks (Deng et al., 2009). Our baseline hyperparameter settings were based on those used by Norouzzadeh et al. (2018). All models were trained for 40 epochs – more epochs and early stopping were also assessed with small changes in validation loss as the stopping criteria, but no difference was found in resulting models.

As well as optimising hyperparameter settings for our dataset, we also investigated the impact of altered dataset configurations. We created a second dataset consisting of only the most common species by restricting the baseline dataset to classes that had a minimum number of 1,000 images (rather than 40), which left 21 of the original 51 classes. We also created a third dataset in which the cap on the number of images per class was increased from 5,000 to 10,000, which affected 11 of the 51 classes. In each case, performance was evaluated on the baseline and common-only test sets.

2.3.1 Generalisability

To form the datasets for comparing performance across individual habitat disturbance levels, we removed all images from locations without a disturbance score (26,546 images). We then formed two datasets: one following the same procedure as above, where all images from a single event were allocated to either the training or test set following a 90-10 split (event-level),

and one where 10 % of the cameras within each disturbance level were withheld to form a pool for the test set, with the remaining 90 % forming the pool of images for the training set (camera-level). For the event-level dataset, we imposed a minimum of four capture events per class per habitat type, leaving 14 classes. For the camera-level dataset, we restricted the data to classes which were captured on at least two cameras in all five disturbance levels, leaving 14 classes. In both cases, we imposed a cap of 5,000 images per class. Only the best performing network, Inceptionv3, as identified from the initial network comparison, was used for the generalisability analysis.

To assess the effect of increasing the number of disturbance levels included in the training set on generalisability, we trained the network on images from every possible combination of disturbance level. To negate the impact of varying numbers of images across the disturbance level combinations, for each dataset configuration we fixed the total number of training images per combination to the smallest individual disturbance level training set size, and randomly sampled images evenly across the included disturbance levels to meet this, ensuring all classes were captured. A test set was similarly formed for each individual disturbance level, ensuring consistent size and all classes included, and used to assess the performance of each combination on images from disturbance levels both seen and unseen during training.

2.3.2 Bounding boxes

One suggested method for improving generalisability is to use an object detector to locate animals within the image and pass the image cropped to the resulting bounding box to the network for training (Beery et al., 2018). Here, we passed the images used for the disturbance level combinations datasets through the Microsoft ‘MegaDetector’v3 (Beery, Morris, & Yang, 2019). In most cases, a single object was identified, and the image was cropped to this bounding box and resized to 256 x 256 as above. Where more than one object was detected, we used the bounding box with highest confidence. In some cases, no object was detected despite being manually labelled as containing an animal. This was caused overwhelmingly by false negatives on the part of the MegaDetector, especially when animals were entering or exiting the field of view and were only partially visible (e.g. only parts of the legs or tail visible). These images were excluded from the generalisability analysis to create a fair comparison. To replicate performance in a ‘real-world’ scenario, however, a comparison of accuracy with and without these images included in the test sets is provided in Appendix 1 (Figure 6.7).

The combined disturbance level datasets were then replicated using these bounding boxes in place of the whole images and the networks trained. We tested the networks on both the original test set for the disturbance level combinations, and the corresponding test set with images cropped to bounding boxes, for comparison.

2.3.3 Metrics

Model performance was assessed against a test dataset which contained images distinct from those used to train the classifier. The performance metrics used are in line with those used in similar studies (Norouzzadeh et al., 2017; Beery et al., 2018; Tabak et al., 2019): (1) Top-1 and Top-5 accuracy: the proportion of all individual images in the test set that were correctly classified within the top 1 or 5 predictions, respectively; (2) F1-score: the harmonic mean of precision and recall, where (a) Precision is the proportion of predictions per class that were correct, i.e. an indication of how reliable the predictions are for a given class; and (b) Recall is the proportion of images per class that were correctly identified, i.e. how fully detected a given class is; and (3) Top-1 accuracy on an event basis: the proportion of capture events that contain at least one correctly classified image.

All metrics were evaluated for the initial network comparison, but generalisability was assessed using F1-score only. Since Top-1 accuracy is heavily influenced by the most common species, we consider F1-score to be a better metric to assess overall performance on an imbalanced dataset.

2.4 Results

2.4.1 Network and dataset comparison

Performance of the three network architectures was comparable (Table 2.1), with the same pattern seen across the four metrics. All networks achieve higher Top-5 accuracy (mean 87%) and Top-1 event accuracy (mean 81%) than Top-1 accuracy on individual images (mean 73%). F1-score is consistently lowest (mean 0.62). Following optimisation (Appendix 1, Figure 6.3), we chose to proceed with the Inceptionv3 network on the basis of F1-score.

Table 2.1: All metrics for the best configurations of VGG16, ResNet50 and Inceptionv3 evaluated on the baseline dataset.

Architecture	Top-1 accuracy	Top-5 accuracy	Top-1/event	F1-score
Inceptionv3	72.8%	86.5%	79.8%	0.63
ResNet50	73.5%	88.6%	81.2%	0.62
VGG16	73.2%	87.1%	80.8%	0.61

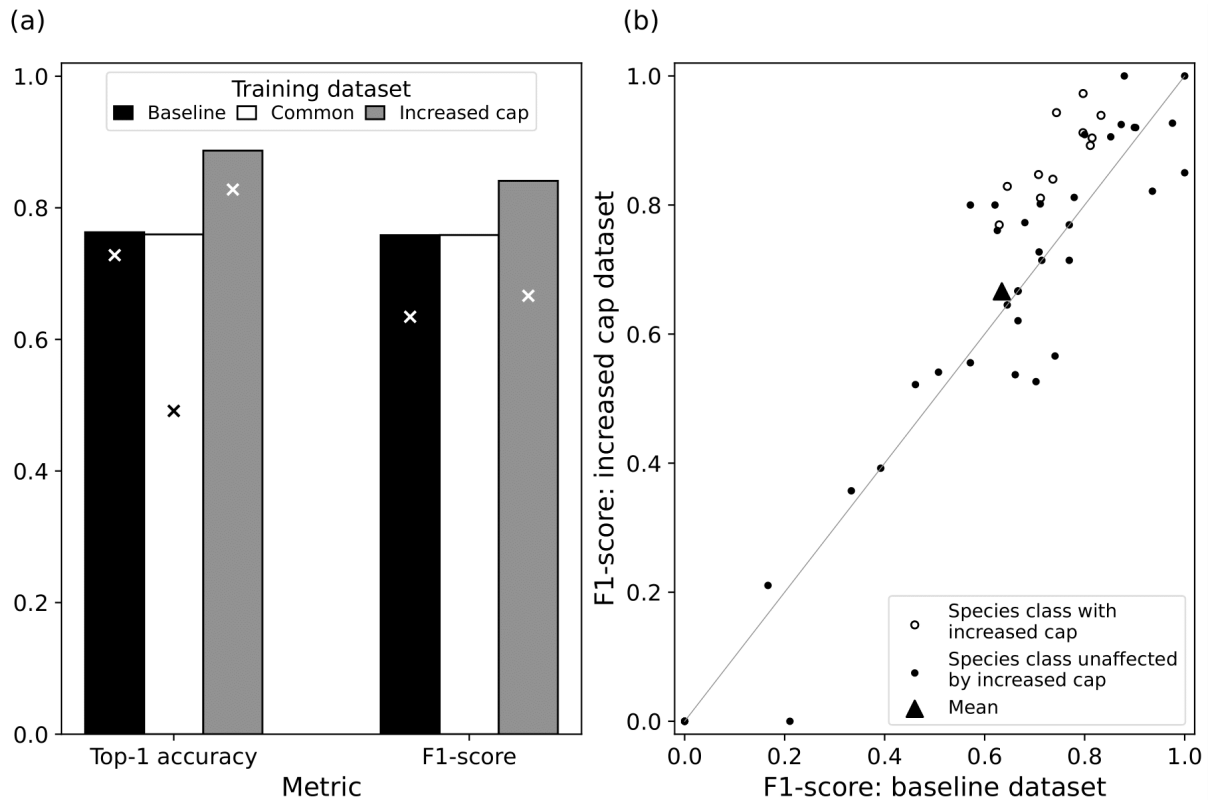


Figure 2.1: Summary metrics for the Inceptionv3 network. (a): Top-1 accuracy and F1-score for the Inceptionv3 network trained on the baseline dataset (max. 5,000 images per class, black), common species only dataset (white) and increased cap dataset (max. 10,000 images per class, grey) when tested on common species only (bars) and when tested on the baseline test dataset (x). Note: it is not possible to calculate F1-score for the network trained on the common species only and evaluated on the full baseline test set due to model structure. (b): F1-score per species class when evaluated on the baseline test set for the network trained on the dataset with an increased cap against trained on the baseline dataset, with a 1:1 line for reference.

When evaluated on only the common classes, overall performance was comparable for the network trained on only the common species and the network trained on all species in the baseline dataset (Top-1: 76% and F1: 0.76 in both cases) (Figure 2.1(a)). Overall Top-1 accuracy and F1-score were improved by increasing the cap on the number of images per class in the training data (Top-1: 89%, F1: 0.84) (Figure 2.1(a)). Including the rarer species in the test set, that is, evaluating performance on the full baseline dataset, saw lower scores for both networks trained on the baseline dataset and on the increased cap dataset (Top-1: 73% and 83%, F1: 0.63 and 0.67, respectively). There was a bigger loss of performance in terms of F1-score than Top-1 accuracy with the inclusion of the rarer species (Figure 2.1(a)), reflecting the bias towards common species in the Top-1 accuracy metric. Including all of the rarer species' test images in the evaluation of the network trained on common species only results in an absolute reduction in Top-1 accuracy of 27% (Figure 2.1(a)). Species-level F1-score and recall tended to increase with a greater number of training images available in the baseline dataset (Appendix 1: Figure 6.5). This was again demonstrated in the increased cap

dataset where all classes that benefitted from extra images saw an increase in F1-score, although the overall mean was only slightly higher than that trained on the baseline dataset (mean F1: 0.67 and 0.63, respectively) (Figure 2.1(b)). This highlights the trade-off with increasing the imbalance within the training data, where some of the classes that did not have any additional training images saw a decrease in F1-score.

2.4.2 Generalisability

For the dataset split at event-level only, peaks in F1-score occurred where the network was trained and tested on images from the same habitat disturbance level (mean 0.76), while performance dropped substantially on disturbance levels not present in the training data (mean: 0.30) (Figure 2.2(a)). Although the distribution across classes for each disturbance level was roughly even, heavily disturbed forest had the greatest number of cameras and images (Appendix 1: Figure 6.2), which may have contributed to it achieving the highest F1-score (0.46) on an unseen disturbance level.



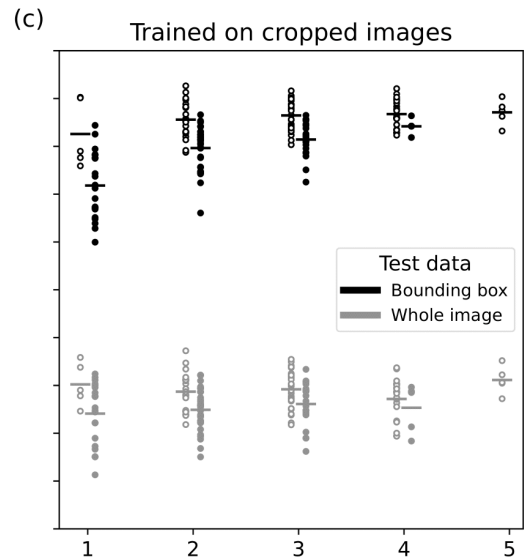
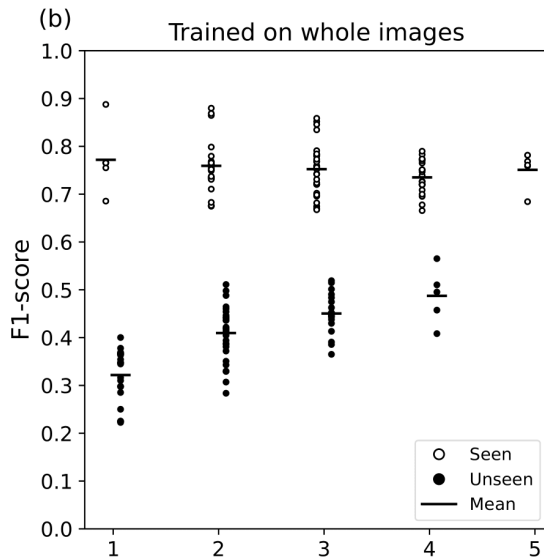
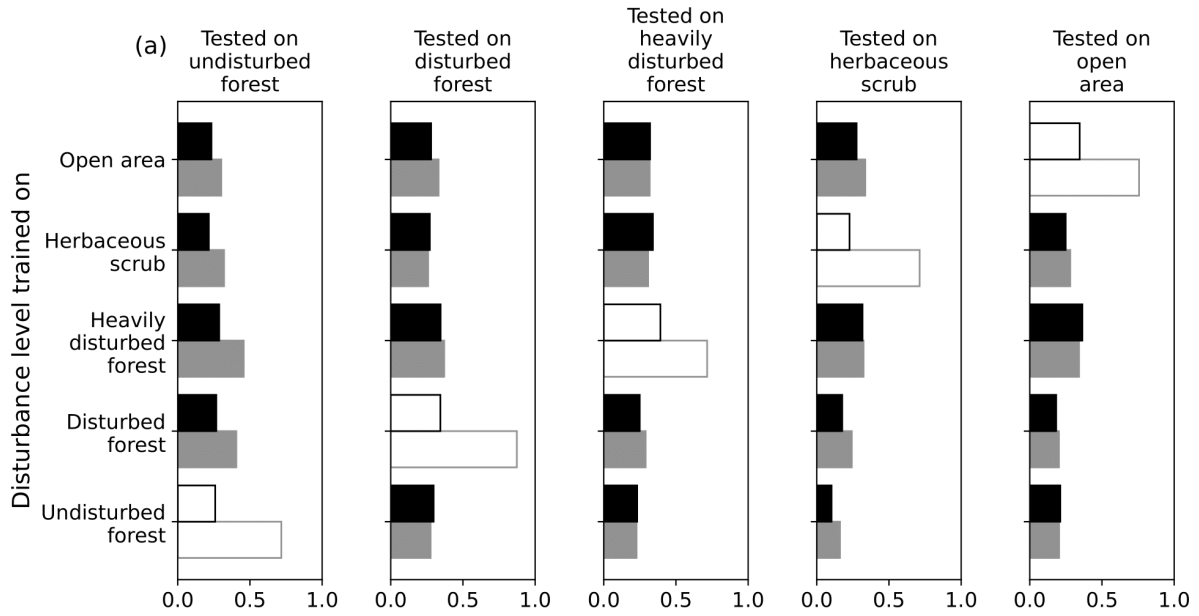


Figure 2.2: Network generalisability. (a) Disturbance level comparison: F1-score per individual disturbance levels. White bars denote F1-score for the same disturbance level used in both training and testing. Filled bars denote F1-score for disturbance levels not seen during training. Results for the dataset split at event level are shown in black and at camera level in grey. (b) and (d): Disturbance level combinations for the dataset split at event-level (circles) and camera-level (diamonds): F1-score for each disturbance level tested following training on a combination of disturbance levels. Every combination of disturbance level was included. Disturbance levels seen during training are denoted by a white marker, while those unseen are denoted by a black marker. The mean F1-score for each number of combinations is also marked. (c) and (e) Disturbance level combinations using bounding boxes: as for (b) and (d), respectively, but training was performed on images cropped to bounding boxes. Black symbols denote testing on images cropped to bounding boxes, while grey represent performance when tested on whole images.

As the number of disturbance levels included within the training dataset was increased from one to four, the mean F1-score on the unseen habitats also increased (0.32, 0.41, 0.45, 0.49, respectively), that is, the network generalised better (Figure 2.2(b)). Conversely, F1-score for the disturbance levels seen during training tended to decrease (0.77, 0.76, 0.75, 0.74, respectively) (Figure 2.2(b)). Performance on unseen habitats was still relatively poor, however, when only one disturbance level was omitted from the training dataset (Figure 2.2(b)).

Using images cropped to bounding boxes in both the training and test sets improved both the overall mean F1-score (0.87) and generalisability (mean F1 score on unseen disturbance levels when trained on a combination of 1, 2, 3 and 4 disturbance levels, respectively: 0.72, 0.80, 0.81, 0.84) (Figure 2.2(c)). Training on bounding boxes and testing on whole images showed a large drop in performance (Figure 2.2(c)).

For the dataset split at camera-level, and network trained on whole images, performance on seen disturbance levels was slightly better than on unseen disturbance levels (Figure 2.2(d)). As with the event-level dataset, an improvement was seen when the number of disturbance levels included within the training data was increased (mean F1-score on seen disturbance level combinations of one, two, three and four levels, respectively: 0.32, 0.40, 0.42, 0.44; unseen: 0.27, 0.36, 0.41, 0.41, Figure 2.2(d)). Performance was best when the network was trained on all disturbance levels (mean F1-score: 0.47).

Using bounding boxes on this dataset again improved performance, but overall F1-score was lower than that achieved with the event-level dataset (mean F1 on all five disturbance levels when trained and tested on cropped images: 0.77 vs 0.87, Figure 2.2(e)).

In a 'real-world' scenario, in which the images containing an animal undetected by the MegaDetector are included in the test set, we can see that Top-1 accuracy drops by 5 % for the network trained on cropped images, since all of these images are deemed to be wrongly classified (Appendix 1: Figure 6.7(d)). For the network trained on whole images, performance does decrease, but to a lesser extent (1%) since the network has the opportunity to classify these images (Appendix 1: Figure 6.7(c)).

2.5 Discussion

This is the first study to assess the application of automated image classification methods and, more specifically, the implications of poor generalisability of CNNs, when considered across a gradient of habitat disturbance in tropical rainforest. Our results highlight the ongoing issues with poor generalisability to unseen camera locations in camera trap image classification, as well as the additional problem of generalisability to changes in background associated with varying levels of habitat disturbance within a single camera trap dataset (Figure 2.2). Training across multiple disturbance levels improved generalisability, suggesting that these differences can be mitigated. Our results demonstrate that an awareness of variation in habitat backgrounds is required when planning a camera trap survey intended for automated classification, and when training the classifier.

We found that in addition to classification accuracy being lower at unseen camera locations within a disturbance level, performance was worse in unseen disturbance levels. One important implication of a lack of generalisability across levels of habitat disturbance, particularly in the context of increasing levels of habitat change, is that classifiers should not be considered 'static'. If a habitat changes over time, naturally or through anthropogenic impacts, new data and additional computer power may be needed to ensure derived classifications and ecological estimates are correct. Alternatively, a dataset comprising images from across the range of possible disturbances should be sought for training at the outset. Field ecologists wanting to use automated classification should therefore consider the generalisability issue when designing future camera trap surveys by stratifying their sites, *a priori*, by broad background types.

Although we fixed the overall number of training images per disturbance level combination, the number of images per class was allowed to vary – by sampling in this way we aimed to replicate the abundance distribution of species within each habitat. Our analysis evaluated performance across all possible combinations of disturbance levels, which should mitigate the impacts of particularly distinct distributions of species in some disturbance levels. We additionally explored how our results changed when we only applied our classifier to the most abundant species within our study (Figure 2.1(a), Appendix 1: Figure 6.6), and found little difference. We also ensured that every class included occurred in both training and test sets to avoid differences in class distribution. Our results suggest that researchers working on smaller camera trap studies might see an improvement in classification performance if labelled images from other studies from similar habitat were to be added to their training data. Further, these results support the aggregation of images from across studies on platforms such as

Wildlife Insights (Ahumada et al., 2020) to enhance available training data and improve classification.

Other researchers have found that a detector-classifier combination was more generalisable than a classifier alone when applied to their dataset (Beery et al., 2018). Our results from the bounding box analysis support this, showing that by focussing images on the animals present and reducing the amount of background, the network was better able to identify species across all disturbance levels (Figure 2.2(c) and (e)). The results also highlight the need to test on bounding boxes rather than the whole image (Figure 2.2(c) and (e)). This is important from a practical perspective, since all images would need to be passed through a detector before being classified, including any new test images from ongoing projects; this adds computational time. We note, however, that even with images cropped to bounding boxes, the impact of background differences across disturbance levels is still evident, with mean F1-score on unseen disturbance levels rising from 0.62, when a classifier is trained on a single disturbance level, to 0.73 for a classifier trained on four levels (Figure 2.2(e)).

The use of an object detector highlighted some discrepancy between the expert labellers and the detector in identifying images as being empty. As a result, object detection may miss subjects that could have been classified correctly. In our data, the missed detections equated to a 5 % reduction in Top-1 accuracy (Appendix 1: Figure 6.7(d)). Although not the focus of this study, a review highlighted that in cases where only a very small part of the animal is visible, or the animal is mostly obscured by vegetation, a human has been able to identify that an animal is present using the visual aid of the whole event sequence, whereas a detector, without that context, could not. Image metadata has been found to improve both automated classification (Terry et al., 2020) and per-species detection performance (Beery et al., 2020) thus could similarly improve detection here.

Since classification performance could have a significant impact on the outcomes of ecological studies, the choice of metric used is important. Our Top-1 accuracy scores (overall: 73%; generalisability analysis: 68%) were slightly lower than one other similar sized study (79% on seen locations (Beery et al., 2018)), but much lower than reported by others (>93.8% (Norouzzadeh et al., 2018), 98% (Tabak et al., 2019)). Part of these differences might be due to the difficulty of the classification in different datasets, for example caused by the extent of the background noise that needs to be overcome in order to detect animals. Images from dense forest environments, as used here, might be expected to present a harder task than open grassland environments (e.g. Snapshot Serengeti as used in Norouzzadeh et al. (2017)). In addition, Top-1 accuracy is naturally dominated by the more common species and can therefore be a misleading metric. We chose to mostly report results as F1-scores here as this

combines precision and recall, and can more accurately reflect classification performance across classes. In practice, Top-5 accuracy can be useful for shortlisting possible species for manual classification, and Top-1 event accuracy can identify events for manual review, both resulting in time-savings. Future work might look into how network performance is assessed in the context of the ecological questions we wish to ask with the data. In particular, an important extension could explore the impact of biases arising from poor generalisability across disturbance levels on resulting ecological studies (e.g. on bias and precision of state variable estimates, such as animal density and occupancy, or on the statistical power to detect differences between experimental treatments).

Although the composition of training data was shown to influence image classification accuracy, including rarer classes did not compromise performance on more common classes, which is important for the continued inclusion of data on rare species to inform conservation efforts. The increased performance achieved through additional training images supports existing studies that have concluded that, although good results can be achieved on smaller training datasets, classification accuracy is generally improved by a larger training dataset (Willi et al., 2019). Here, we saw increased classification performance when the cap on common species was increased, which, through the use of the F1-score metric, we can confidently say was not driven by a simple numerical increase of correctly identified common species.

Class imbalance is common within camera trap datasets, and has been shown to impair the performance of neural networks (Buda et al., 2018). The resulting difficulty in training neural networks on rare species is a known problem (Beery, Liu, et al., 2019). In practice, if a network is able to satisfactorily classify and remove common species such that manual classification is reduced to the rare species, this would still result in substantial time savings. For conservation projects, however, where rare species are the main interest, we might require better performance, especially in recall (since false negatives likely have a higher conservation cost than false positives). Specifically improving classification on rare species was not the focus of this study, but oversampling and weighted loss methods have been tried elsewhere with some success (Norouzzadeh et al., 2018; Terry et al., 2020). Others have tried generating artificial images containing the rare species or incorporating images from other datasets (Schneider et al., 2018; Beery, Liu, et al., 2019). Future work could therefore include applying these methods to this dataset and assessing performance.

Despite the inherent difficulties with training on rarer species, and the general trend seen for increased F1-score with number of training images, we did see instances of high F1-score for relatively rare species (Appendix 1: Figure 6.5). There was also variation in performance on

the most common species. Future work will need to consider the degree of morphological variation within and among species as a possible contributing factor as to why networks are able to learn some species better than others.

2.5.1 Conclusions

This study highlights the ongoing issue of poor performance of automated species classifiers across unseen locations in camera trap studies. Importantly, it also demonstrates that unseen backgrounds (here disturbance levels) can further impair classification performance. Unseen locations in novel habitat disturbance-levels had poorer classification performance than those from unseen locations in habitat levels seen during training. Generalisability can be improved by the use of bounding-box object detection prior to species classification, but the use of bounding boxes did not completely eliminate the problem. As camera trap datasets become more abundant, and the use of machine learning for automated classification becomes more commonplace, it will be critically important to ensure that estimation of changes in ecosystem function and composition are not biased by methodological choices in detection and identification of species. This is particularly important in the context of current global biodiversity loss, for monitoring the impacts of anthropogenic activities on ecosystems and mitigating further declines.

Chapter 3

3 How are spatiotemporal interactions between mammals in tropical forest impacted by land-use change and disturbance?

3.1 Abstract

Species interactions provide a link between biodiversity and ecosystem functioning, and have been identified as a potential mediator of species responses to environmental change. The impact of disturbance on species interactions is difficult to capture and has been poorly studied to date.

Here, bearded pigs and humans are used as a case study to test for spatiotemporal avoidance in disturbed tropical forest using camera trap data. It has previously been shown that bearded pigs shift from diurnal activity in old-growth forest, to crepuscular in logged forest and to nocturnal in oil palm plantations (Davison et al., 2019). It was hypothesised that this could be due to avoidance of oil palm plantation workers, who are mostly active in the morning, or of the higher midday temperatures that occur in oil palm plantations.

Linear mixed-effect models were constructed to model the duration between captures of the two species as a function of which species is recorded first. A separate model was constructed for each of two measures of environmental variation: land-use categories (old-growth forest, logged forest and oil palm) and habitat degradation due to logging, ranging from undisturbed forest to open area, to test for spatiotemporal avoidance across categories.

No evidence of spatiotemporal avoidance of humans by bearded pigs was found. It was not possible to conclude whether this was due to the absence of the interaction occurring, or the inability to detect it due to insufficient sensitivity in the method, or insufficient data. It is possible that spatiotemporal avoidance was occurring, but that it was happening at a temporal scale too fine for the method to detect. This study highlights the challenge in being able to extract enough datapoints to power the method, such that reliable conclusions could be drawn, from a camera trap dataset representative of a long-term monitoring program.

3.2 Introduction

Environmental changes have both direct and indirect effects on ecosystem structures and processes (Rahman & Candolin, 2022). Species interactions provide a link between biodiversity and ecosystem functioning, and have thus been identified as a potential mediator

of species responses to environmental change (Morales-Castilla et al., 2015; Rahman & Candolin, 2022). The impact of habitat disturbance on mammalian wildlife has been demonstrated and measured through changes in abundance, richness and diversity (Rustam et al., 2012; Brodie et al., 2015; Wearn et al., 2016). The impact of disturbance on species interactions is less straightforward to capture and has been poorly studied to date (Meijaard & Sheil, 2008; Denis et al., 2019; Rahman & Candolin, 2022).

One area that has seen, and continues to experience, considerable human-mediated disturbance are tropical rainforests undergoing logging and conversion to plantations. Comprising communities that are particularly vulnerable to extinction as a result of climate change (Wiens, 2016), it is especially important to understand the processes underlying the biodiversity changes in these areas. Some mammals have been previously shown to persist in logged forest, increasing in mean relative abundance by 28% from old-growth forest, but the opposite was seen in oil palm plantations with a 47% decrease compared to forest (Wearn et al., 2017). Species that are able to survive in disturbed landscapes are ecologically important but understudied.

Bearded pigs, which are in decline and have been classified as vulnerable (IUCN), have been shown to persist in disturbed forest, despite a decline in relative abundance of 87% from old-growth forest (Wearn et al., 2017). Through foraging, soil rooting and nest building, bearded pigs provide ecosystem services such as seed dispersal and are an ecologically important species (Malhi et al., 2022). In a study as part of the Stability of Altered Forest Ecosystems (SAFE) Project in Borneo, bearded pigs were found to shift from diurnal activity in old-growth forest, to crepuscular in logged forest and to nocturnal in oil palm plantations (Davison et al., 2019). It was thought that this change in activity pattern could be due to avoidance of humans in the plantations, with workers there being mostly active between 7 am and midday (Davison et al., 2019).

As well as, or as an alternative to, temporal partitioning to avoid competitors or predators, some species avoid others in space. Some evidence of spatial avoidance has already been discovered within the SAFE project area. In a study examining co-occurrence patterns of mammal species at the SAFE study site, 13 instances of spatial avoidance were found in old growth forest, three of which were in congeneric pairs: *Maxomys* rats (*M. surifer* and *M. rajah*), the greater and lesser (*T. kanchil*) mouse-deer, and the thick-spined and Malay porcupines (*Hystrix crassispinis* and *H. brachyura*) (Wearn et al., 2019). It was suggested that competition could be driving spatial niche separation in these species pairs (Wearn et al., 2019).

The studies outlined above used data collected via camera trap surveys. The integration of spatial and temporal analyses of camera trap data has been highlighted as critical to improve

our understanding of the impact of anthropogenic disturbance and land-use changes on species interactions and community dynamics (Frey et al., 2017). Previous spatiotemporal analyses have found evidence of temporal partitioning among tropical forest mammals (Haidir et al., 2018; Santos et al., 2019). In their study of neotropical forest carnivores, Santos et al. (2019) found apparent spatial and temporal partitioning for most of the species pairs analysed using a combination of occupancy models and temporal activity overlap assessment. Comparatively, Haidir et al. (2018) found no evidence of spatial avoidance by two mesopredators in Sumatran tropical forest, but did find evidence for temporal niche separation, whereby clouded leopard appeared more nocturnal and thus had higher temporal overlap with more nocturnal prey species: porcupine and mouse-deer. Again, these authors used occupancy models, as well as Bayesian modelling and kernel density estimate associations. Limitations in using co-occurrence to infer species interactions have been raised by Blanchet et al. (2020), who argued that the complexity of ecological systems blurs the link between interactions and co-occurrence, and other methods and data besides presence-absence are needed.

Following the suggestion in Davison et al. (2019) that the observed change in bearded pig behaviour between forest and oil palm could be avoidance of either humans or the higher temperatures in oil palm, I test the former here. Using the framework set-out by Niedballa et al. (2019) in their analysis of methods for detecting spatiotemporal interactions in camera trap data, I apply their best-powered method to the SAFE dataset to test for spatiotemporal avoidance of humans by bearded pigs as indicated by their temporal shift in activity. Since the dataset used contains measures of environmental change in both land use and habitat disturbance level, I consider a potential change in behaviour across both separately. Here, habitat disturbance corresponds to environmental variation due to the effects of logging and subsequent regeneration, ranging from undisturbed forest to open area (Wearn et al., 2017). I will test three hypotheses: (i) bearded pigs avoid humans in oil palm, (ii) bearded pigs avoid humans in open areas, and (iii) interactions between humans and bearded pigs vary across land-use categories and disturbance levels. Following Niedballa et al. (2019), in each case of hypothesised spatiotemporal avoidance of humans by bearded pigs, I would expect to see significantly longer durations between a capture of a human and the subsequent capture of a bearded pig than between a capture of a bearded pig and subsequent human. Where no interaction is occurring, I would expect to see no significant difference in interval durations.

3.3 Methods

3.3.1 Dataset

The dataset comprises camera trap data collected as part of the Stability of Altered Forest Ecosystems (SAFE) Project situated in tropical forest in Malaysian Borneo (Ewers et al., 2011), and detailed in Chapter 1.2. This dataset comprises images from three land-use categories: old-growth forest, logged forest and oil palm plantations. Individual camera trap sites sampled during the initial survey were also given a score of habitat disturbance for a 5 m radius around the sampling point. The disturbance levels are categorised on a scale comprising 5 levels from undisturbed forest to open area (see Table 1.1 for details). Note that camera trap sites not included within the initial survey have not been given a disturbance level score.

For this analysis, all images used during camera trap setup, or tagged with malfunction or an incorrect time stamp, were removed from the dataset. To ensure independence of capture events, consecutive triggers of the same species occurring within 30 minutes of each other were grouped to one event (Burton et al., 2015; Davison et al., 2019; Easter et al., 2020).

3.3.2 Analysis

In their assessment of methods for detecting spatiotemporal interactions using simulated camera trap data, Niedballa et al. (2019) considered four methods: linear models (using both log-transformed and untransformed response variables), the Mann–Whitney U-test, a permutation test and a test based on randomly generated records. All tests were assessed as valid, above a minimum of 10 records per species, but a linear model using a log-transformed response was found to be the most suitable approach for spatiotemporal avoidance (Niedballa et al., 2019). The response variable used was time intervals between recordings of two species: AB and BA where the notation represents the order of the events, e.g. AB is the time interval between an event capturing species A and the next event of species B, and vice versa. These intervals have been used in previous analyses of camera trap data: Harmsen et al. (2009) found evidence of spatiotemporal avoidance between jaguars and pumas using a similar linear model approach, while Karanth et al. (2017) found minimal spatiotemporal overlap of three sympatric carnivores using a permutations approach.

3.3.2.1 Interval extraction

The duration of intervals between capture events of the two species were extracted from the data using the sequence of events at each camera for each deployment as a separate time series. Here, a “human-bearded pig” interval represents the time between a capture of a human and the subsequent capture of a bearded pig, and a “bearded pig-human” capture is

the reverse. It was assumed for this analysis that the presence of other species did not impact the behaviour of the two species of interest, so detections for all other species were ignored. For bearded pig-human intervals, where multiple detections of a bearded pig occurred before a subsequent detection of a human, only the interval between the last detection of a bearded pig in the sequence and the first subsequent human was recorded, and similarly for human-bearded pig intervals.

3.3.2.2 Linear model

For this analysis, a linear model was constructed with the \log_{10} of the interval duration as the response variable and pair orientation as a fixed effect explanatory variable, following Niedballa et al. (2019). The pair orientation variable is a factor consisting of two levels and refers to whether the interval is between a capture of a bearded pig followed by a human (bearded pig-human), or a human followed by a bearded pig (human-bearded pig). In their study, Niedballa et al. (2019) considered a single camera, whereas the SAFE dataset comprises images from multiple cameras, so I therefore also added camera site as a random effect. Spatiotemporal avoidance of humans by bearded pigs is expected to result in significantly longer human-bearded pig intervals compared with bearded pig-human intervals.

3.3.2.3 Environmental variation

To investigate whether the interaction between bearded pigs and humans changes with habitat disturbance, two measures of environmental variation were considered within the model. The interactions are first considered across land-use categories including old-growth forest, logged forest and oil palm. These are incorporated into the model as an additional explanatory factor variable. The interactions are then separately evaluated across habitat disturbance levels, consisting of undisturbed forest, disturbed forest, heavily-disturbed forest, herbaceous scrub and open area. These levels are informed by the score given to each camera trap site at the time of deployment. Similarly to the land-use model, a disturbance level factor variable is included as an additional explanatory variable.

3.4 Results

3.4.1 Camera trap events

Following grouping of detections occurring within 30 minutes, the dataset contained 2,377 bearded pig events and 1,627 human events. These events were unevenly distributed across both land-use categories and habitat disturbance levels (Figure 3.1). The majority of bearded pig images were captured in logged forest (2,006), with only a small number (39) in oil palm (Figure 3.1(a)). Human captures in logged forest and oil palm were similar in number (716 and 715, respectively), but much lower in old-growth forest (196).

Capture events for bearded pigs across disturbance levels were similar (range: 235-283), except for a higher number (1,000) seen in heavily-disturbed forest (Figure 3.1(b)). Human captures varied more across disturbance levels, with 557 capture events in open area, but only 69 in disturbed forest.

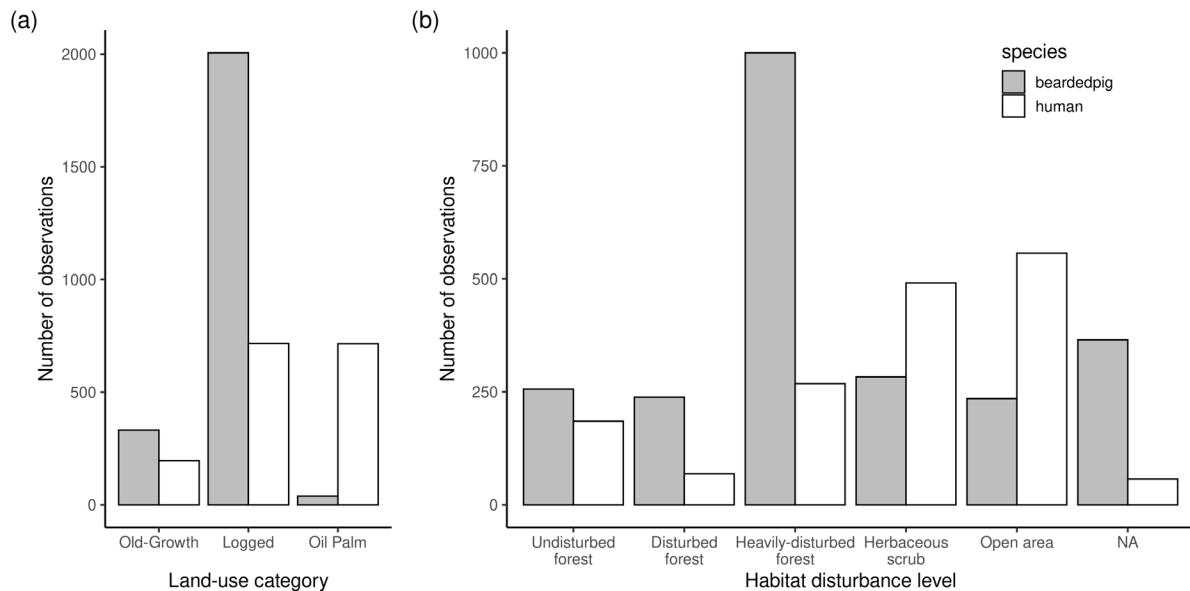


Figure 3.1: Number of camera trap events. Number of independent camera trap capture events per (a) land-use category and (b) habitat disturbance level. Note that "NA" represents detections from camera traps without a habitat disturbance score.

Bearded pig captures occurred during the day - predominantly in the morning - in old-growth forest, around dawn and dusk in logged forest, and during the night in oil palm, as expected (Figure 3.2(a)). The distribution of bearded pig captures in undisturbed forest closely follows that of old-growth forest, while in all other disturbance levels, there are peaks around dawn and dusk, with smaller peaks during the night (Figure 3.2(c)). Human captures consistently occurred during the daytime (between 6 am and 6 pm) in all land-use categories and habitat disturbance levels (Figure 3.2(b) and (d)).

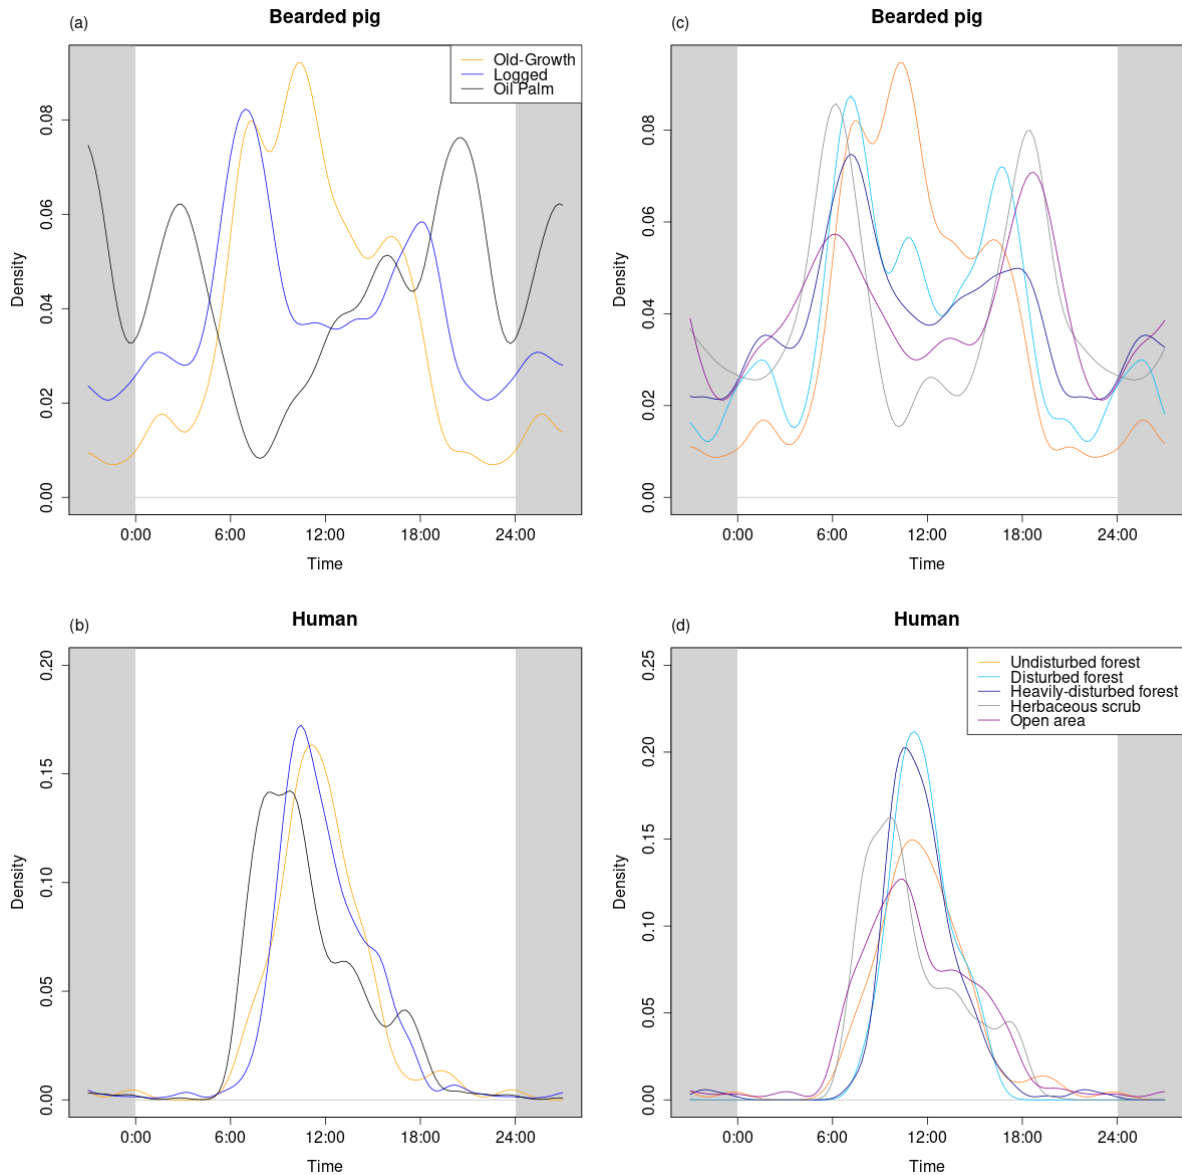


Figure 3.2: Distribution of detection times for bearded pigs and humans. Density of detection times for independent capture events of (a) bearded pigs and (b) humans in land-use categories, and of (c) bearded pigs and (d) humans in habitat disturbance levels.

3.4.2 Intervals between species

After extracting the interval durations between species from the camera trap time series, the dataset comprised 265 intervals in total from logged forest, 27 in old-growth forest and 28 in oil palm (Figure 3.3(a)). Intervals across disturbance levels were also unevenly split with approximately 100 in each of heavily-disturbed forest and open area, but only 22 in disturbed forest and 27 in undisturbed forest (Figure 3.3(b)).

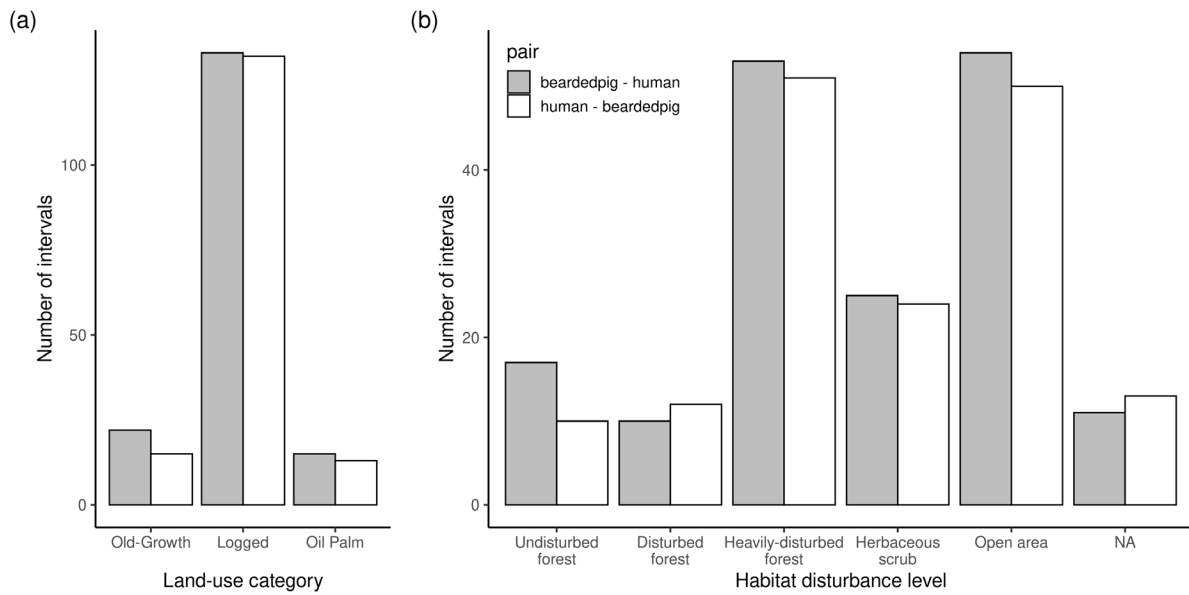


Figure 3.3: Number of intervals extracted. Number of intervals between captures of bearded pigs and humans (grey) and humans and bearded pigs (white) extracted per (a) land-use category and (b) habitat disturbance level. Note that “NA” represents detections from camera traps without a habitat disturbance score.

The interval durations were longer in old-growth forest (median 10 days for bearded pig-human intervals and 9 days for human-bearded pig intervals) compared to the other land-use categories (1 day for bearded pig-human intervals in both logged forest and oil palm; 3.5 and 4.5 days for human-bearded pig intervals in logged forest and oil palm, respectively) (Figure 3.4).

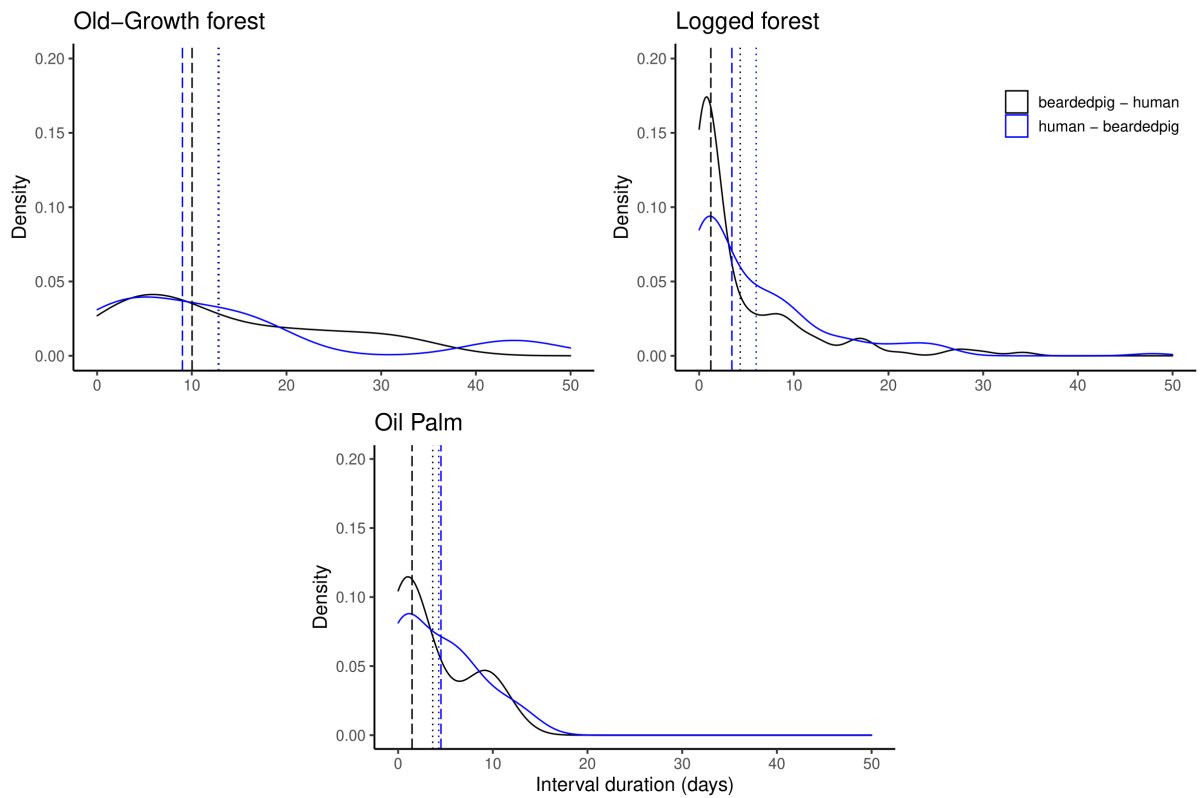


Figure 3.4: Distribution of interval durations across land-use categories. Shown for bearded pig-human intervals (black) and human-bearded pig intervals (blue). The median (dashed line) and mean (dotted line) interval durations are also plotted.

Similarly, the interval durations were longer in undisturbed forest (median 10 days for bearded pig-human intervals and 9.5 days for human-bearded pig intervals) compared to the other disturbance level categories (Figure 3.5). A reduction in median interval duration with increased disturbance was seen, with the shortest interval durations recorded in herbaceous scrub and open area (median 0.9 days for bearded pig-human intervals and 1.2 days for human-bearded pig intervals in both) (Figure 3.5).

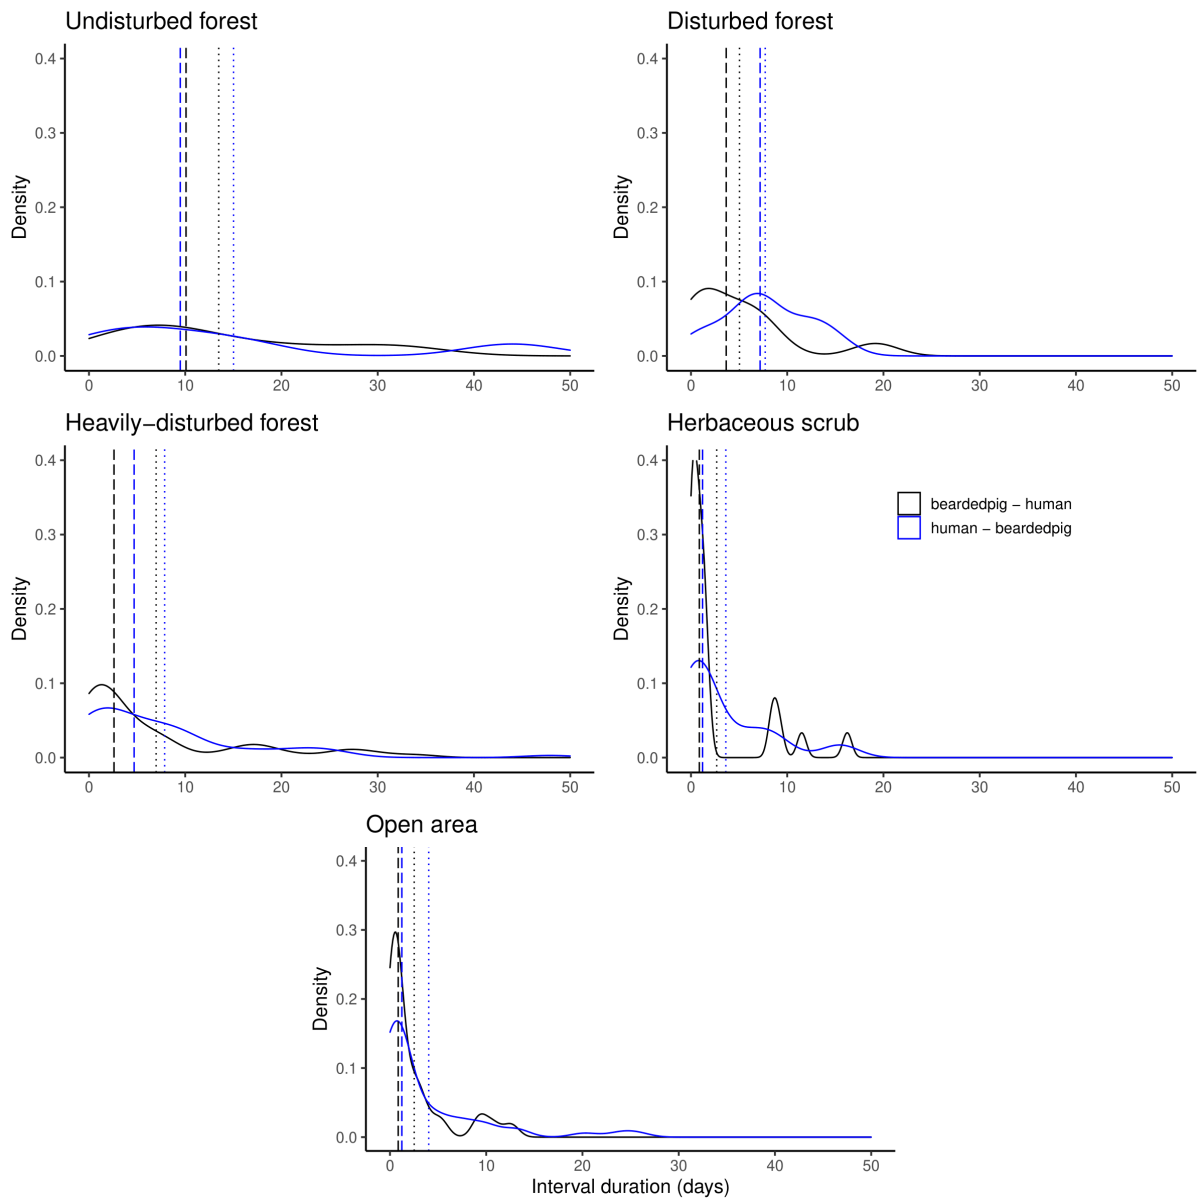


Figure 3.5: Distribution of interval durations across habitat disturbance levels. Shown for bearded pig-human intervals (black) and human-bearded pig intervals (blue). The median (dashed line) and mean (dotted line) interval durations are also plotted.

3.4.3 Land-use category model

Modelling interval duration against land-use category showed no evidence of spatiotemporal avoidance of humans by bearded pigs in any land-use category (Figure 3.6(a)). The interval durations were longer in old-growth forest compared to the other land-use categories. There was no significant difference in predicted interval durations between the pair orientations (Figure 3.6(b)).

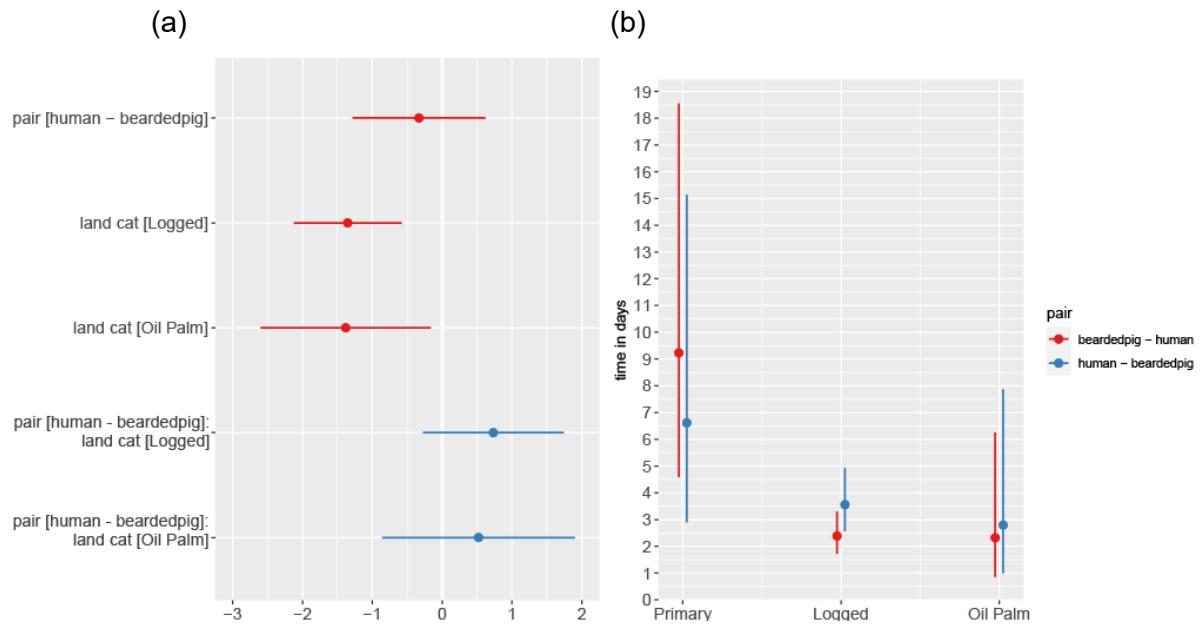


Figure 3.6: Results for the land-use mixed-effect model. (a) Model coefficients and (b) predicted model intervals. This model takes the duration between detections of bearded pigs and humans as the response variable, with pair (the order in which the species occur) and land-use category as explanatory variables, including the interaction between these explanatory variables. The intercept (not shown) comprises interval durations between a bearded pig and subsequent human in old-growth forest.

3.4.4 Disturbance level model

In the disturbance level model, a trend can be seen for shorter interval durations with increasing disturbance (Figure 3.7). As with the land-use category model, no evidence of spatiotemporal avoidance of humans by bearded pigs was found. For all disturbance levels except undisturbed forest, the human-bearded pig intervals are predicted to be longer than the bearded pig-human intervals, but not significantly so (Figure 3.7(b)).

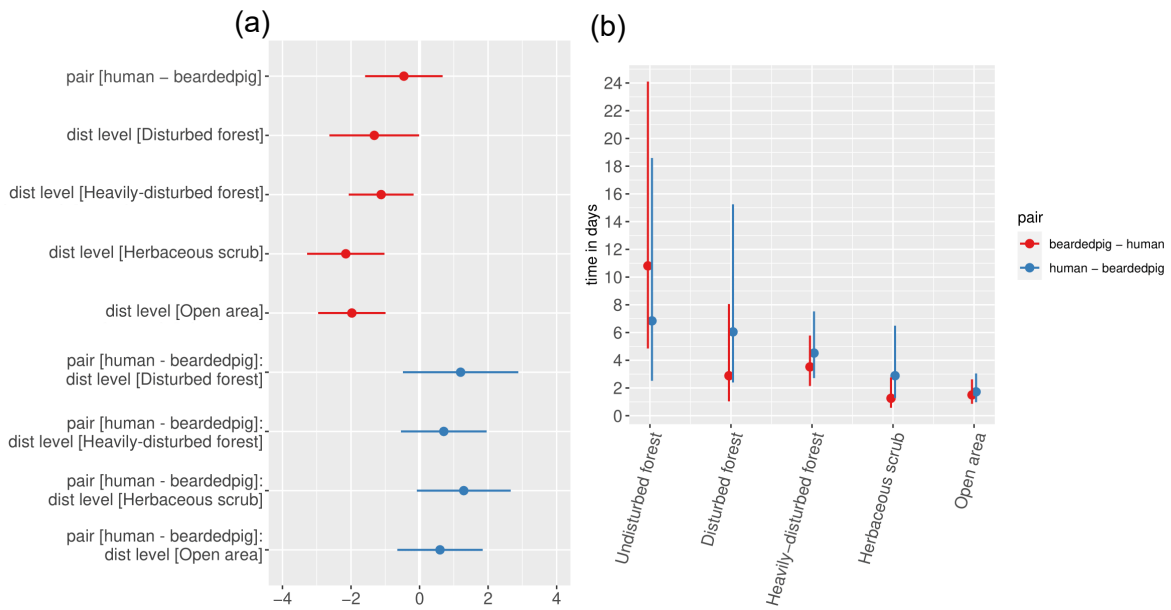


Figure 3.7: Results for the habitat disturbance level mixed-effect model. (a) Model coefficients and (b) predicted model interval durations. This model takes the duration between detections of bearded pigs and humans as the response variable, with pair (the order in which the species occur) and disturbance level as explanatory variables, including the interaction between these explanatory variables. The intercept (not shown) comprises interval durations between a bearded pig and subsequent human in undisturbed forest.

3.5 Discussion

In this study, I have investigated whether spatiotemporal avoidance of humans by bearded pigs can be detected in a camera trap dataset taken from disturbed tropical forest, where a shift in bearded pig activity patterns has previously been observed (Davison et al., 2019). Using a method previously identified as optimal (Niedballa et al., 2019), I found no evidence to support the hypothesis that bearded pigs avoid humans in oil palm plantations. It is not possible to conclude whether this was due to the absence of the interaction occurring, or the inability to detect it due to insufficient sensitivity in the method, or insufficient data.

It has previously been shown that bearded pigs shift their daily activity pattern from diurnal in old-growth forest to nocturnal in oil palm (Davison et al., 2019). In their study, the authors found that human and bearded pig activity periods did not overlap in oil palm, and their ranges overlapped minimally. It was suggested that the shift in activity pattern in bearded pigs could be an avoidance of oil palm plantation workers, or avoidance of the higher midday temperatures that occur in oil palm compared to forest. The authors noted that a lack of early morning activity by bearded pigs when microclimates would be favourable lent weight to the avoidance of humans hypothesis (Davison et al., 2019). The results here do not support this hypothesis.

Further, the distribution of human detections shows that human activity consistently occurs during the daytime across land-use categories (Figure 3.4(c)). The majority of human

detections were recorded between 6 am and 12 pm in oil palm, and slightly later in logged and old-growth forest, with a peak around 11 am and 12 pm, respectively. With consistent diurnal human activity, we might expect more crepuscular or nocturnal bearded pig behaviour in other land-use categories. Davison et al. (2019), however, note that old-growth forest offers more cover than the other categories, so experiences less extreme temperatures and provides a buffer between the species.

It is possible that there is spatiotemporal avoidance occurring, but that it is happening at a temporal scale that is too fine for the method to detect. The average interval times extracted from the data ranged from 1 to 10 days, so it might be unrealistic to expect to detect any responses happening on scales of minutes or hours. Re-runs of the analysis with a focus on shorter timescales of 7 days and 2 days also did not provide evidence in support of the hypothesis, however, and were under-powered due to a reduction in data (Appendix 2). The process of data extraction for this method means that while bearded pig activity patterns in oil palm appear nocturnal, there are detections during the day, and these “outliers” could be influencing the results here and masking any underlying avoidance process, especially for the smaller number of intervals extracted in this land-use. A cross-validation analysis to assess the influence of individual data points might provide insight into whether this is the case, but again is challenging with the paucity of data.

These results showed a trend for shorter interval durations, for both interval orientations, with increased disturbance (Figure 3.6(b), Figure 3.7(b)). This is an unexpected result since with less overlap in activity, we might expect longer durations between records of the two species. This trend does not follow the abundance pattern seen for bearded pigs, which show an increase from old growth forest to logged forest then a decrease to oil palm (Wearn et al., 2017) so is not a consequence of higher abundance resulting in more frequent camera trap triggers. It could be a result of increased human activity with disturbance such that human detections occur more frequently, thus reducing interval durations between captures of the two species.

The camera trap dataset used in this analysis is taken from a long-term project designed to capture biodiversity and ecosystem function change as forests are impacted by human disturbance. The survey was consequently designed to cover a range of habitat types, and camera traps were placed at random within the sampling design. Other species interactions studies have designed camera trap surveys specifically to capture their species of interest, by placing them along roads or known trails (Karanth et al., 2017; Haidir et al., 2018). In these targeted camera trap surveys, we might expect to capture the species of interest more frequently and to build a more accurate representation of their behaviour. It is possible that

the method used here might perform better in these datasets, which more closely represent the simulated camera trap data, than on a more general monitoring camera trap survey as used here.

Although the dataset met the minimum requirements for this method to be valid (>10 records), the requirement for the analysis to be reliable (>50 records) (Niedballa et al., 2019) was not met for all land-use categories or disturbance levels. This dataset was taken from an extensive survey, which captured thousands of bearded pig and human events. It resulted in small numbers of datapoints (intervals between species) for the analysis, however. This method relies on both species of interest being captured on the same camera, such that intervals between the two can be extracted. Where species have a spatial avoidance interaction, they may not be captured on the same cameras, thus providing no or limited data to power the analysis. In my dataset, where only one of the species was captured, no interval could be recorded, and the data discarded. Further, during the interval extraction process, where a sequence of detections of one species was captured before a detection of the other species, only one interval is recorded, that is, the shortest duration between the two species so results in only one datapoint.

This study has highlighted that a method that works well on simulated data does not necessarily translate across to real data. In their analysis, Niedballa et al. (2019) simulated strong spatiotemporal avoidance in just two species on one camera. Using a linear model, the authors were able to detect this avoidance. In my dataset, I have more than 50 species and in focussing on just one species pair, the detections of all other species were excluded. The analysis therefore assumed that the behaviour of the two species of interest was unimpacted by the presence and behaviour of any other species. This assumption would hold for most human behaviour, where the humans captured are researchers or plantation workers, but hunting has also been recorded in the study area and in these instances, the humans are likely tracking targeted species.

Another weakness in this method is that results could misrepresent any interaction as a result of the relative timing of activity peaks. Considering the temporal activity patterns of humans and bearded pigs reported by Davison et al. (2019), it is possible that using this method, any spatiotemporal avoidance could be masked by the peaks in activity impacting the durations observed. For instance, species that do not interact but have activity peaks that are 6 hours apart might appear to interact via this method due to the intervals for species A-species B being approximately 6 hours, while the intervals for species B-species A would be approximately 18 hours, which could present as a significant difference. It is consequently difficult to set clear expectations for detecting an interaction. An alternative interval, such as

AA vs ABA, as used by Zalewska et al (2021), would provide a “control”. In this case, the influence of the presence of humans could be assessed by comparing the interval durations between bearded pig captures when a human is present compared with when the human is absent.

Many interaction studies using camera trap data have used co-occurrence analyses to infer interactions between species (Ramesh et al., 2017; Davis et al., 2018; Parsons et al., 2019). Challenges with this approach have been highlighted by Blanchet et al (2020), who argue that the complexity of ecosystems blurs the link between species interactions and cooccurrence. In their analyses of mammal species at the SAFE project, Wearn et al (2019) agree that while little evidence of spatial avoidance was found in their study, it is possible that the species were segregated in time or vertically in space (e.g. in the forest canopy), which the cooccurrence analyses were not able to detect. The method used here utilises the detection of a species in both space and time so is a step forward, although it also cannot account for vertically segregated species. This study demonstrates the importance of considering both spatial and temporal data to gain a more complete understanding of behavioural processes.

Species interactions remain an important topic of study, and particularly in species such as the bearded pig, which modify the habitat around them and are able to persist in tropical forests that have been impacted by human disturbance. As habitats continue to undergo anthropogenic change, capturing the subsequent changes in the interactions between the species present could provide vital insight into the impacts on biodiversity. The challenges in being able to power the method used in this study such that reliable conclusions can be drawn from a camera trap dataset that is representative of a long-term monitoring program, however, have been highlighted here. We need methods that are able to detect species interactions, but the complexity of ecosystems and the corresponding network of interacting species continues to present challenges to the analysis.

Chapter 4

4 Detectability of competitive avoidance and predator-prey interactions in a simulated camera trap survey

4.1 Abstract

Species interactions, including competition and predation, can play a dominant role in shaping ecological communities, particularly in mammals. They are, however, difficult to directly study due to logistical challenges in capturing the behaviour and performing experiments. Species interactions can be inferred from indirect observations, such as through camera traps.

An evaluation of methods for detecting spatiotemporal avoidance in simulated camera trap data previously identified a linear model of time intervals between species detections as effective. This analysis featured only one camera, however, thus it did not represent a typical camera trap survey performed in the field.

Here, I construct an agent-based model (ABM) to simulate the movement of multiple individuals from two species around a camera trap 'grid' for one week. Each simulation is performed 100 times. This forms a camera trap dataset to which I apply a linear mixed-effect model to detect simulated interactions between species. I further consider the thresholds required in species population density and camera trap density, as well as interaction zone radius as a proxy for interaction strength, for the interaction to be detected.

I have shown that this method can detect avoidance and attraction behaviour. Attraction behaviour could be detected at all but the weakest strength, while avoidance could only be detected at maximum strength, and with a cap of 72 hours applied to the maximum interval duration. Since both interaction behaviours produced the same qualitative results, difficulty in discerning the specific interaction at play could be challenging in real-world data analyses using this method. Applying a previously identified threshold in data volume required to power the analysis suggested that for the parameter values investigated here, a camera trap density of 2/km² would be required for a population density of ~0.1/km², while a camera trap density of 1/km² is sufficient for larger populations.

4.2 Introduction

Species interactions have been described as the "architecture of biodiversity" (Bascompte & Jordano, 2007). Interactions such as competition and predation can play a dominant role in shaping ecological communities, particularly in mammals (Palomares & Caro, 1999; Sinclair et al., 2003). Where we are able to capture data on species interactions, we are able to glean

insights into ecosystem processes. Further, if we are able to monitor species interactions in the long-term such that we are able to identify any changes in frequency or spatiotemporal occurrence, for example, we may be able to mitigate any subsequent impacts on biodiversity (Caravaggi et al., 2017).

Species interactions are difficult to directly study, but can be inferred from indirect observations, such as through camera traps (Amir et al., 2022). Previous spatiotemporal analyses of animal behaviour using camera trap data have discovered adaptability in behaviour across gradients of resource availability (Karanth et al., 2017), avoidance behaviour (Harmsen et al., 2009; Ross et al., 2013; Ramesh et al., 2017), and temporal partitioning (Haidir et al., 2018; Santos et al., 2019).

Niedballa et al. (2019) identified that a linear model of the time interval between captures of individuals from two different species was the most powerful method for detecting spatiotemporal avoidance between two species using camera trap data. One shortcoming of this study was the analysis simulated camera trap data from two individuals on only one camera. In practice, a camera trap survey commonly comprises multiple camera placements and can capture many individuals of the same species, so any analytical methods applied need to take this into account.

In Chapter 3, I applied a recommended method for detecting spatiotemporal avoidance between species in camera trap data to the SAFE dataset, but was unable to detect an interaction. It was not clear whether the inability to detect the interaction was due to the absence of the interaction or a weak interaction strength, or insufficiencies in the method or volume of data. Therefore, here I explore the volume of data and interaction strength required to power the method to detect modelled interactions. I construct an agent-based model (ABM) to simulate the movement of multiple individuals from two species around a camera trap 'grid' to form a dataset. I then apply the linear model suggested in Niedballa et al. (2019) to assess the ability of this method to detect interactions between species.

ABMs are an ideal framework for this simulation since they allow a 'bottom-up' approach, in which individuals, rather than the population as a whole, are modelled (Mortensen et al., 2021). In ABMs, rules are applied at each time-step to update the status of each individual agent. ABMs have previously been applied to ecological problems such as predicting population trajectories for reintroduced species (Andersen et al., 2022), assessing movement of mammals in conservation priority areas to highlight areas requiring restoration or mitigation (Jayadevan et al., 2020), and the evaluation of effective predator control methods (Latham et al., 2019; Pacioni et al., 2021).

In this application, I create a set of rules tailored to each species that allow me to simulate movement according to different interactions. I use these to investigate the impact of interaction strength, population density and camera trap density on our ability to detect spatiotemporal avoidance and predatory stalking using the method outlined by Niedballa et al. (2019). I aim to identify thresholds in these parameters, beyond which competitive avoidance and/or predation can be captured. Three interaction scenarios are considered, in addition to a control scenario in which all agents follow a Correlated Random Walk (CRW). I use the SAFE Project (Ewers et al., 2011) situated in Borneo as a case study to inform the study site size and the population densities of the species within the model.

4.3 Methods

4.3.1 Simulation model

An agent-based model (ABM) was constructed in Python to simulate the movement of individuals around an 8 km x 8 km study site. The model comprised two species, A and B, both of which consisted of a population of either 6, 12 or 24 individual agents. These population sizes were calculated from estimated mammal densities from a previous camera trap study, which suggested that most mammal species have densities less than 0.25/km² in tropical forest (Wearn et al., 2022). Movement speed (0.15 m/s) was similarly taken from a previous camera trap study (Rowcliffe et al., 2016). Each individual was assigned a randomly generated starting location within the study site, and a randomly generated direction of travel. A time-step size of 10 seconds was chosen to allow for time between camera trap triggers, assuming that when a camera trap is triggered, it takes a rapid burst of 10 images over approximately 5 seconds (Wearn et al., 2022). At each time-step, the agents' locations were updated according to the interaction scenario (Figure 4.1). 100 simulations were performed for each combination of parameter values, with each simulation lasting for 1 week (60,480 time steps).

Interaction strength is modelled as the distance within which the presence of an individual from the other species triggers the interaction behaviour. A range of interaction zone radii are considered between 15 m and 150 m. These values were chosen based on the vision distance of 15 m used in the predator-prey ABM constructed by Wheatley et al. (2020), and the recorded flight initiation distances of black-tailed deer of up to 150m, with most data points between 25 and 100 m (Stankowich & Coss, 2006).

Table 4.1: Parameter values used in simulations, with sources.

Parameter	Value(s)	Source/Notes
No. simulations	100	-

Study site size	8 km x 8 km	SAFE project size
No. camera traps	64, 128, 256	Represent densities of 1, 2 and 4 per km ² . Camera trap studies vary greatly in camera trap spacing and density, from 20 per km ² (Caravaggi et al., 2016) to 1 per 2 km ² (Santos et al., 2019). The densities used here were chosen to balance representing data captured at different spatial scales with the computational power required.
No. species	2	-
No. agents per species	6, 12, 24	Most mammal species have densities <0.25 km ² at SAFE (Wearn et al., 2022). Here, 6 and 12 represent densities ~0.1 and 0.2 per km ² . 24 provides an additional higher density for comparison.
Time step size	10 seconds	To allow for camera being triggered to take 10 images.
Length of simulation	1 week	-
Velocities	0.15 m/s	Speed trap CT paper (Rowcliffe et al., 2016) data for Red brocket deer.
Standard deviation for normal distribution of theta	$\pi/24$	Chosen to tune moving angle in CRW.
Avoidance/attraction zone radius	15 m, 25 m, 50 m, 75 m, 100 m, 150 m	Range of values between 15 m (Wheatley et al., 2020) and 150 m (Stankowich & Coss, 2006): - Wheatley et al., (2020) used a 15 m vision distance in their predator-prey ABM - Stankowich & Coss (2006) (Fig2b): black-tailed deer, flight initiation distance up to 150m, but most data points between 25 and 100 m
Kill zone radius	1 m	Wheatley et al., 2020

Camera trap trigger detection distance	6 m	Based on red deer (Table 3, Hofmeester et al., 2017). Wearn et al., (2022) also suggests a range of 2-10 m, depending on species.
Grouping into capture event	30 mins	To ensure independence of capture events, consecutive triggers of the same species occurring within 30 minutes of each other were grouped to one event (Burton et al., 2015; Davison et al., 2019; Easter et al., 2020).

4.3.2 Interaction scenarios

Scenario 1: Correlated random walk (CRW): control

In the first scenario, used as a control, both species followed a correlated random walk (CRW) with no interactions occurring (Figure 4.1(a)). At each time step, each agent received a new direction of travel (θ) sampled from a normal distribution with their previous direction as the mean and with a constant standard deviation of $\pi/24$. The agent's location was then updated using this new θ and their speed.

Scenario 2: One-way avoidance: weak competitor avoiding a strong competitor, or prey avoiding predator

In this scenario, individuals of species A avoid individuals of species B, while species B follows a CRW as in Scenario 1. At each time step, for each agent from species A, a scan is performed to identify the locations of any agents from species B that are within their avoidance zone (Figure 4.1(b)). If any are found, the vector between the two agents is calculated and used to update the trajectory of the species A agent such that they move in the opposite direction to where the species B agent is located (Figure 4.1(d)). Where more than one species B agent is located within the avoidance zone of a species A agent, an average is taken such that the species A agent is moving away from all of the species B agents equally (Figure 4.1(c)).

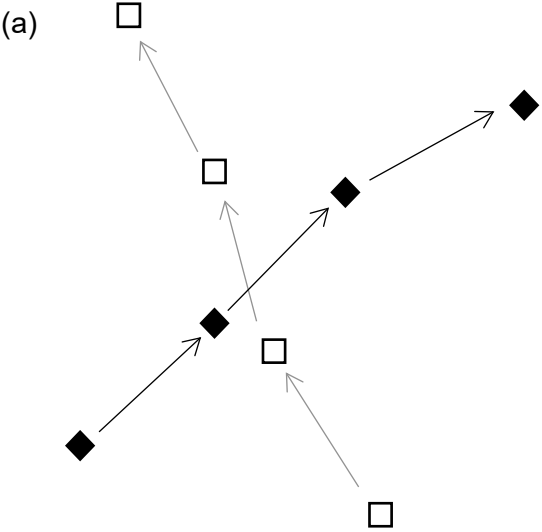
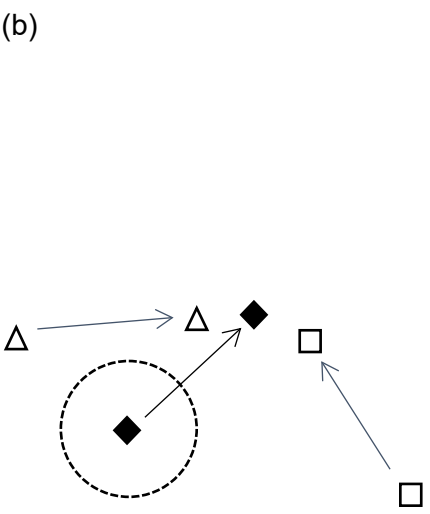
Scenario 3: One-way attraction: predator stalking prey

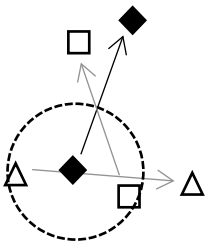
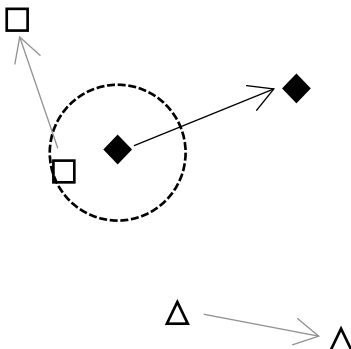
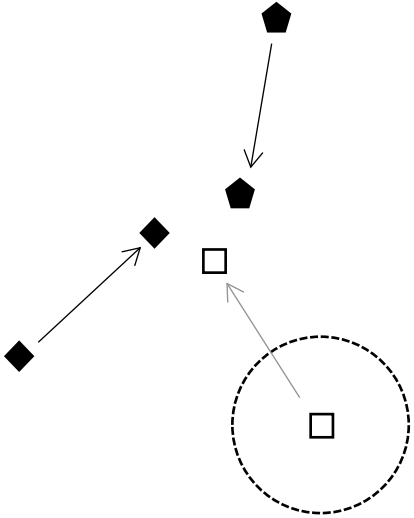
Here, species A continues on a CRW while species B is attracted towards individuals from species A. Similarly to Scenario 2, at each time step, for individuals from species B, a scan is performed to identify the locations of any agents from species A within their attraction zone (Figure 4.1(e)). Again, the vector between the agents is calculated and this is used to move the agent from species B towards an agent from species A. In contrast to Scenario 2, where there are multiple species A agents within the interaction zone, the agent from species B

orients towards the closest agent from species A (Figure 4.1(g)). Once the agent from species B is within a kill zone radius from the species A agent, that species A agent is killed and regenerates randomly within the study space to maintain the population densities (Figure 4.1(f)).

Scenario 4: Combined avoidance and attraction: predator-prey relationship

In this combination, species A avoids species B (as in Scenario 2) while species B is attracted towards species A (as in Scenario 3). The same mechanisms are used as described above.

Scenario 1: Correlated random walk	
<p>(a)</p> 	<p>At each time step, agents from both species continue in a direction similar to their previous step, unaffected by the presence or proximity of other agents.</p>
Scenario 2: One-way avoidance	
<p>(b)</p> 	<p>No agents from species B (white) are in the avoidance detection zone (dashed black circle) for the species A agent (black) so all agents move as per CRW.</p>

<p>(c)</p> 	<p>Both species B agents (white) are within the avoidance detection zone of the species A agent (black).</p> <p>The species A agent moves in a direction that is equally opposite to the two species B agents, while the species B agents continue on a CRW.</p>
<p>(d)</p> 	<p>One species B agent (white square) is within the avoidance detection zone of the species A agent (black).</p> <p>The species A agent moves in the opposite direction to where the species B agent is located, while both species B agents continue on a CRW.</p>
<p>Scenario 3: One-way attraction</p>	
<p>(e)</p> 	<p>No agents from species A (black) are in the attraction detection zone (black circle) for the species B agent (white) so all agents move as per CRW.</p>

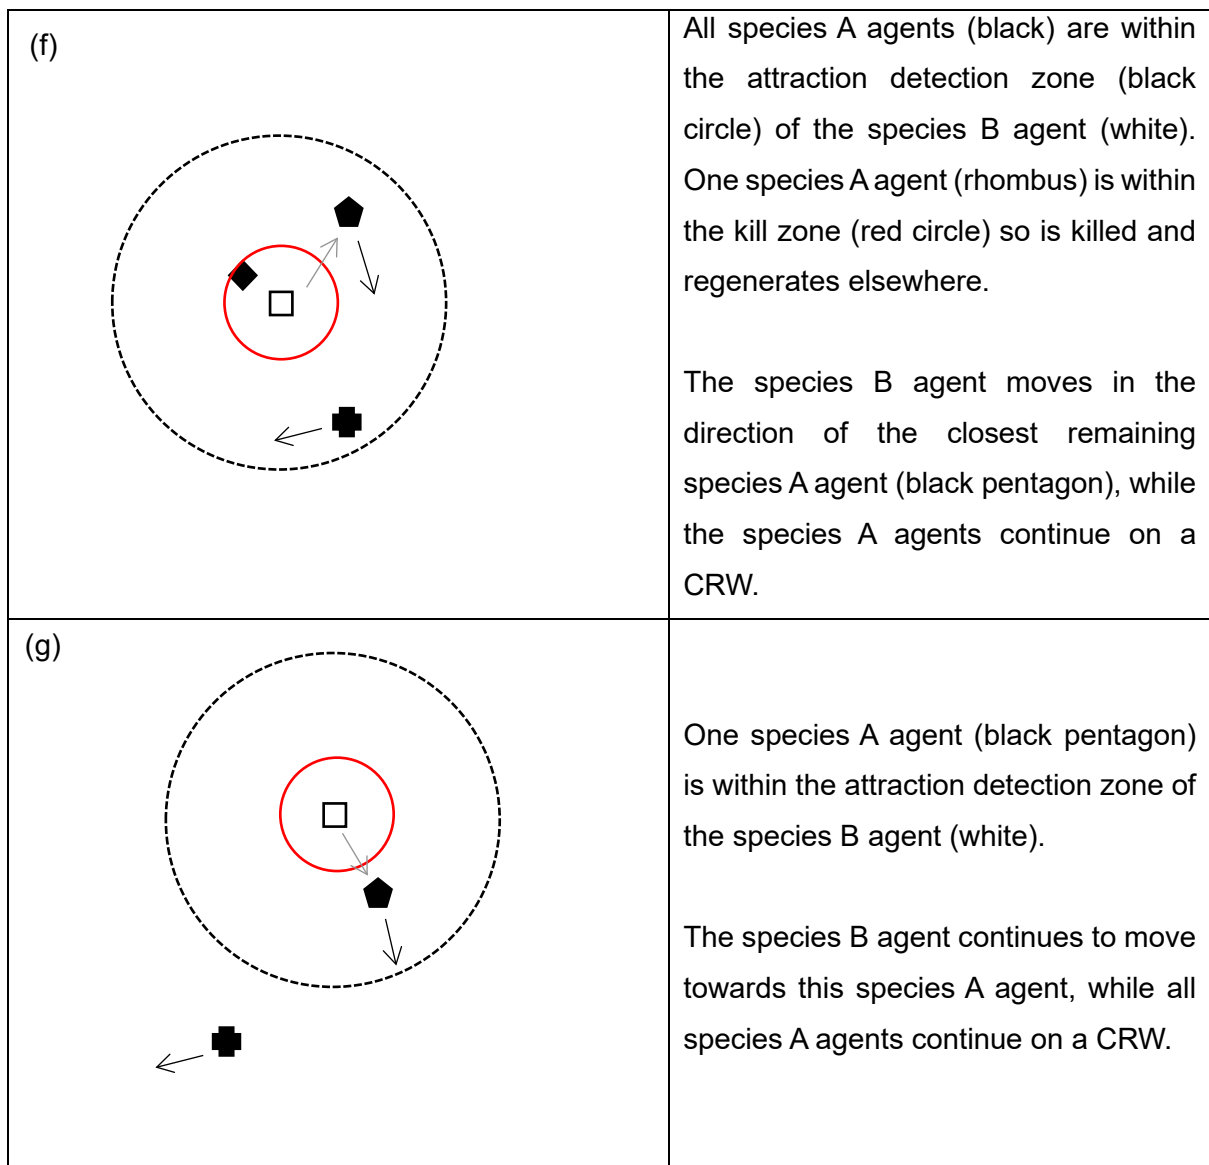


Figure 4.1: Model schematic for Scenarios 1-3. Figure (a) depicts a correlated random walk, Figures (b)-(d) illustrate movement within the avoidance scenario and Figures (e)-(g) show movement within the attraction scenario.

4.3.3 Camera traps

Camera trap locations were randomly generated using camera-trap densities of 1/km², 2/km² and 4/km². Each agent trajectory was compared with the camera trap locations, and ‘triggers’ recorded for instances where an individual was situated within the camera trap detection distance of 6 m (Hofmeester et al., 2017). Consecutive triggers for individuals of the same species that occurred within 30 minutes were grouped together to form one capture event, following common procedure for analysing observational camera trap data to ensure independent events (Burton et al., 2015; Davison et al., 2019; Easter et al., 2020).

4.3.4 Analysis

4.3.4.1 Interactions and data collection

I first measured the number of interaction events occurring within the ABM simulations and compared across scenarios and parameter values to quantify the impact of increasing the camera or agent density, as well as interaction strength. I repeated this for the camera trap events recorded across the simulations.

4.3.4.2 Intervals analysis

The duration of intervals between a capture event of species A and the next capture event of species B, and vice versa, were then extracted. The subsequent analyses were based on comparisons of the distribution, mean and median of these interval durations for each scenario.

Following the approach highlighted by Niedballa et al. (2019), a linear mixed model was constructed for each scenario with the \log_{10} of the interval durations as the response variable, the pair orientation (see note below) as a fixed-effect explanatory variable and camera location as a random effect (Table 4.2). For Scenarios 2-4, the interaction zone radius was also included as a fixed-effect explanatory variable via an interaction term with the pair orientation (Table 4.2). Results were deemed significant where the absolute value of the associated t-value was greater than 2.

Note that the pair variable is a factor with two levels: AB and BA, where AB represents an interval between a camera trap capture of an individual from species A and the subsequent capture of an individual from species B, while BA represents the reverse.

In the CRW scenario, where no interactions are occurring, no significant difference is expected between the durations of AB intervals compared with BA intervals. In the avoidance scenario, where species A agents avoid species B agents, significantly longer BA intervals than AB intervals are expected. Similarly, for the attraction scenario, where species B agents move towards species A agents, AB intervals are expected to be shorter than BA.

Table 4.2: Model formulae used in R.

Scenario	Model Formulae
1. CRW	$\log(\text{duration}) \sim \text{pair} * n_agents + \text{pair} * n_cams + (1 \text{camera})$
2. One-way avoidance	$\log(\text{duration}) \sim \text{pair} * \text{avoid_radius} + \text{pair} * n_agents + \text{pair} * n_cams + (1 \text{camera})$
3. One-way attraction	$\log(\text{duration}) \sim \text{pair} * \text{attract_radius} + \text{pair} * n_agents + \text{pair} * n_cams + (1 \text{camera})$
4. Combined	$\log(\text{duration}) \sim \text{pair} * \text{avoid_radius} * \text{attract_radius} + n_agents + n_cams + (1 \text{camera})$

4.3.4.3 Maximum interval duration cap

To investigate the ability to detect interactions when focussed on shorter timescales, I imposed caps on the maximum interval duration included within the analysis for the avoidance and attraction scenarios. For each cap, the dataset was reduced to only intervals with durations less than or equal to the cap, and the linear model re-run using this restricted data. The timescale on which species interactions occur can vary greatly; average interval durations from the analysis in Chapter 3 ranged from 0.9 to 10 days so caps of 1 to 6 days were applied, as well as shorter caps of 6 and 12 hours. The number of intervals available for the analysis at each cap is considered alongside the linear model results to explore the trade-off between a shorter timescale and data availability.

4.3.4.4 Applying the Niedballa et al. (2019) threshold

In their study, Niedballa et al. (2019) found that more than 10 records per species were required for their method to be valid, and more than 50 records per species were required for their method to produce reliable results, that is, for it to have power > 0.8. Here I apply this threshold as the requirement to extract 50 of both AB and BA intervals for each parameter setup to identify the combination of population and camera trap densities, and interaction zone radius required.

4.4 Results

4.4.1 Avoidance and attraction scenarios

4.4.1.1 Interactions and data collection

An increase in the interaction zone radius predictably resulted in more interaction events occurring in the simulations (Figure 4.2), with the largest radii of 150 m generating between 22 and 36 times the number of interaction events per simulation seen at the smallest radii of 15 m. This did not impact the number of camera trap triggers or intervals extracted in the avoidance scenario (Appendix 3: Figure 8.1 and Figure 8.4), but an increase in both with interaction zone radius was seen in the attraction scenario (Appendix 3: Figure 8.2 and Figure 8.5). Both increasing the number of agents and increasing the camera trap density resulted in an increase in the number of intervals between species extracted for the analyses for all scenarios (Figure 4.3). Since more capture trap events were recorded in the attraction scenario than for the CRW or avoidance scenarios, particularly at the highest camera trap density, more intervals between species were extracted (4193 vs 3078 and 3084 for the CRW and avoidance scenarios, respectively, at a camera trap density of 4/km², Figure 4.3).

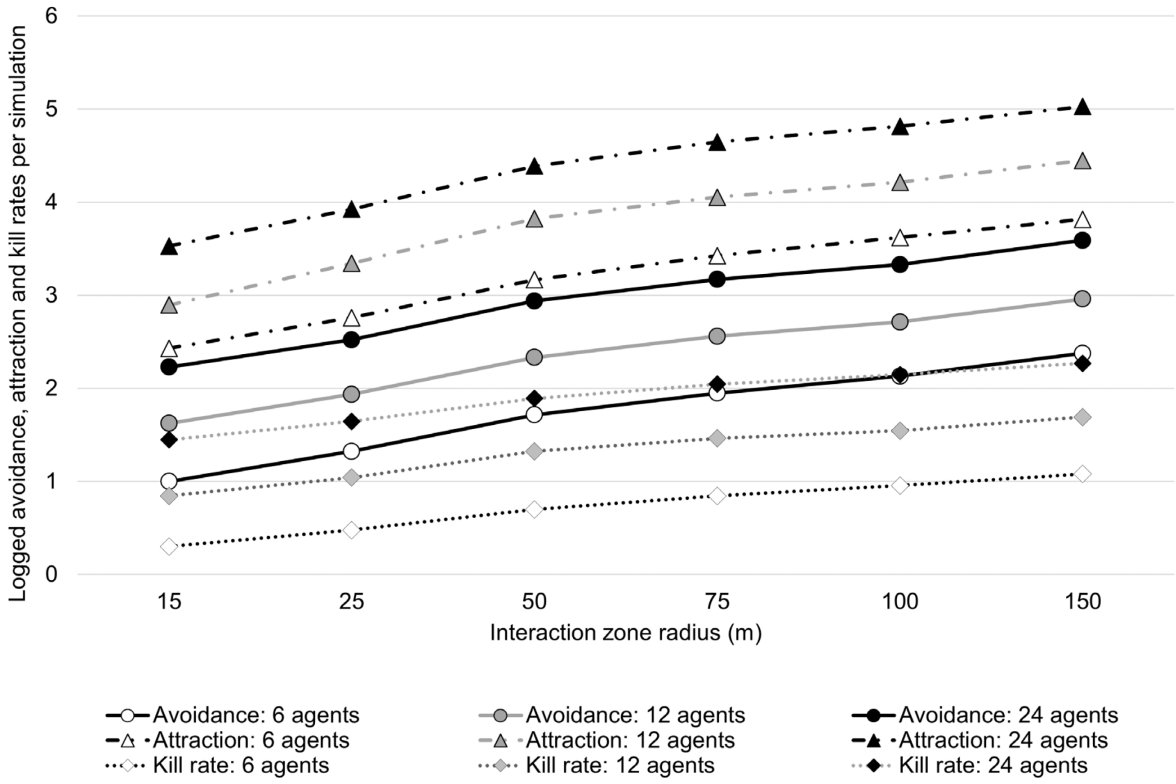


Figure 4.2: Interaction rates. Avoidance (circles), attraction (triangles) and kill (rhombus) rates per simulation for avoidance and attraction scenarios, respectively, comprising 6 (white), 12 (grey) and 24 (black) agents per species. Note that a \log_{10} scale has been used for the y-axis.

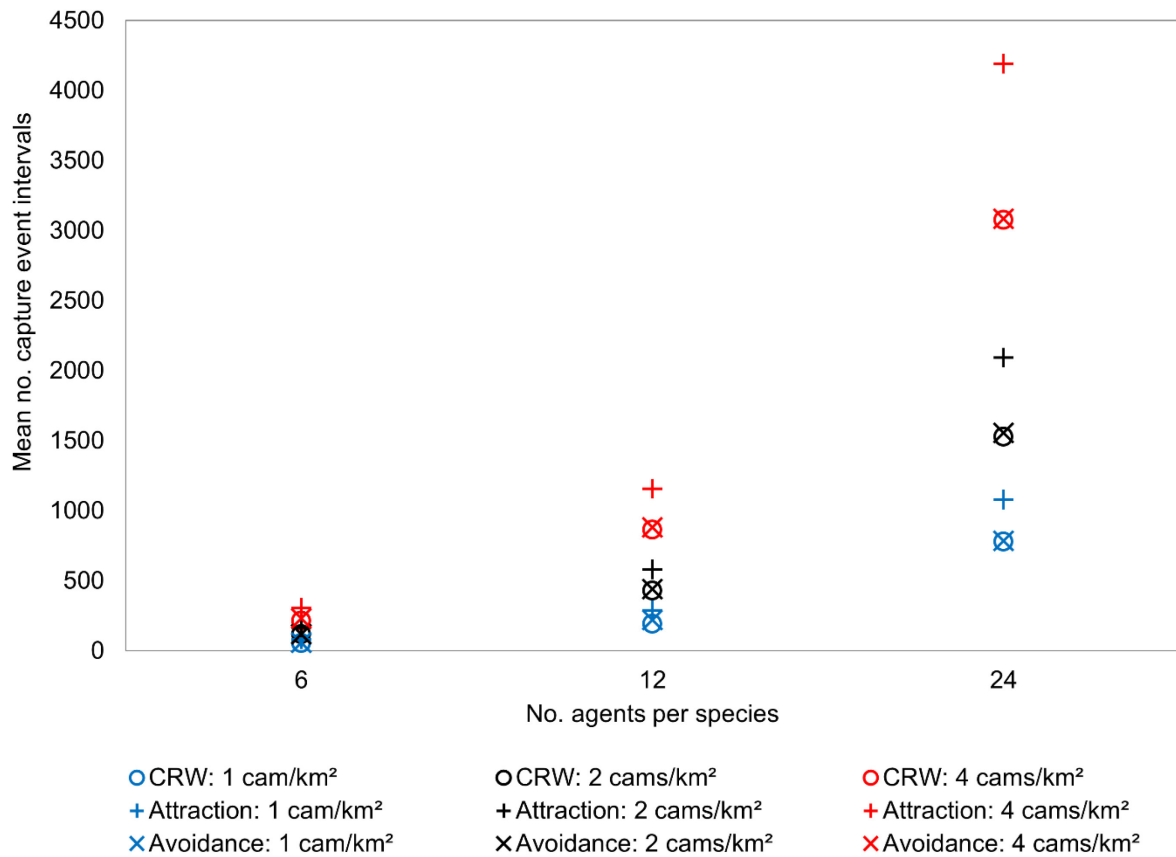


Figure 4.3: Number of intervals extracted for Scenarios 1-3. Mean number of intervals extracted from the simulated camera trap data at camera densities of 1/km² (blue), 2/km² (black) and 4/km² (red) for the CRW (circles), avoidance (cross) and attraction (plus sign) scenarios.

4.4.1.2 Intervals analysis: distribution

Increasing both the number of agents per species and the camera density reduced the overall mean and median interval duration for Scenario 1: CRW (Table 4.3, Figure 4.4).

Table 4.3: Median and mean interval duration in days for Scenario 1: CRW.

	Population size (no. agents per species)			Camera density (per km ²)		
	6	12	24	1	2	4
Median (days)	1.97	1.82	1.73	1.84	1.80	1.72
Mean (days)	2.21	2.19	2.03	2.11	2.11	2.04

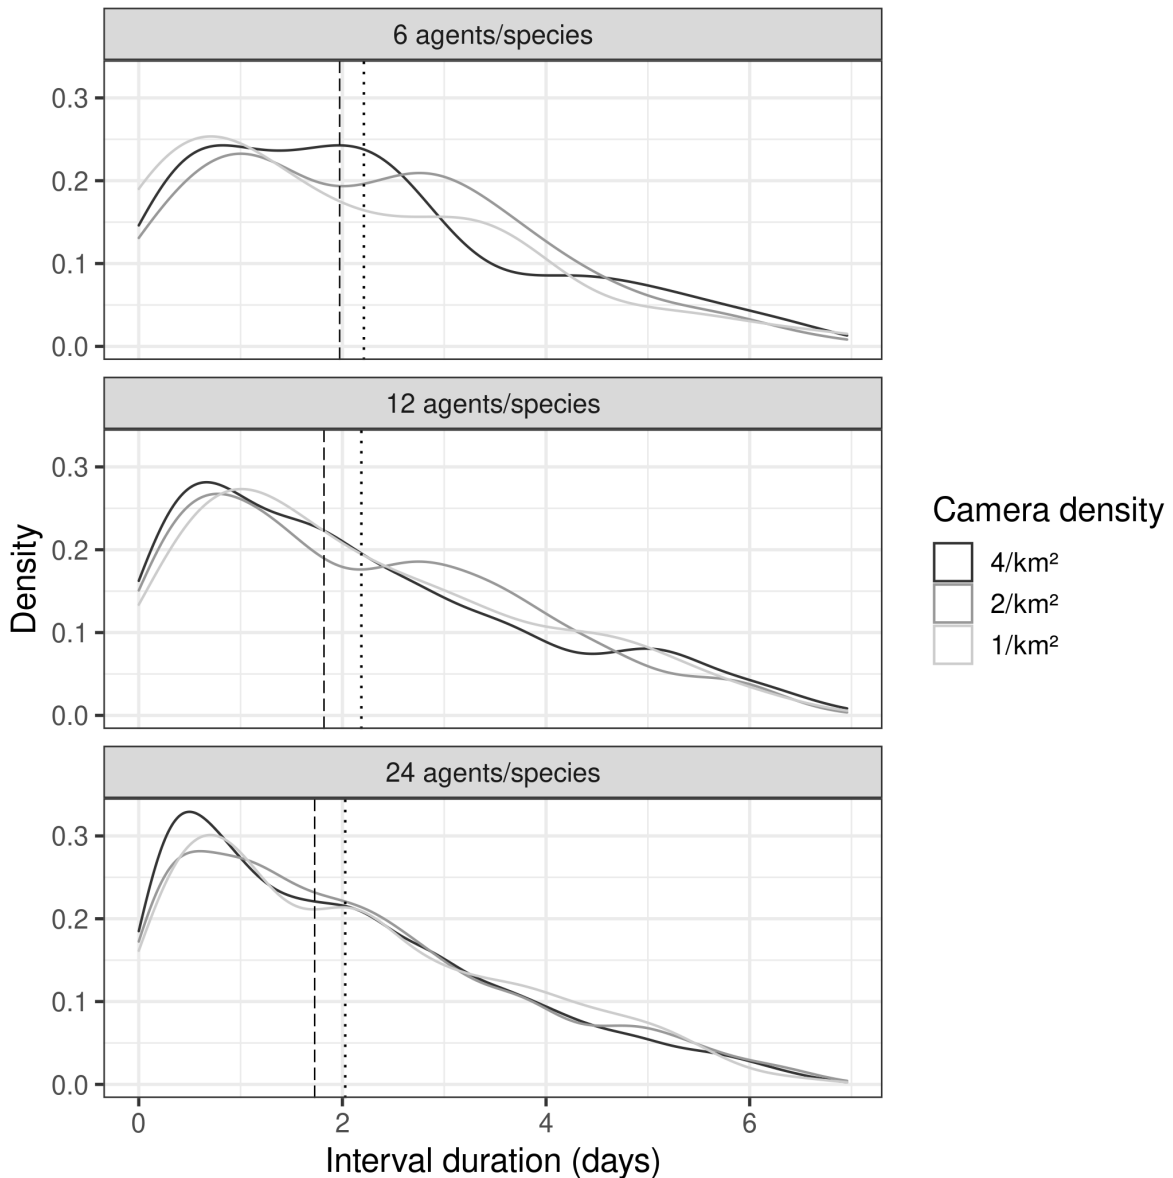


Figure 4.4: Distribution of interval durations for Scenario 1: CRW. Density is shown for population sizes of (top) 6, (middle) 12 and (bottom) 24 agents per species at camera densities of $1/\text{km}^2$ (black), $2/\text{km}^2$ (dark grey) and $4/\text{km}^2$ (light grey). The mean (dotted line) and median (dashed line) are also plotted.

Comparing across interval pairs, the distributions for AB intervals for CRW and avoidance are similar, but for intervals BA we see fewer instances of the shortest durations in the avoidance scenario than for the CRW scenario (Figure 4.5). The mean and median interval duration for BA in the avoidance scenario are also greater than for the AB intervals (mean: 2.07 vs 1.97; median 1.76 vs 1.57) (Figure 4.5). Intervals for the attraction scenario are concentrated close to zero, with a higher peak for AB intervals than for BA (AB median: 0.62 days vs BA median 1.05 days) (Figure 4.5). The mean and median for both pair orientations are shorter than for the CRW and avoidance scenarios (Figure 4.5).




Figure 4.5: Distribution of interval durations per pair orientation for Scenarios 1-3. Interval durations for Scenario 1: CRW (blue), Scenario 2: one-way avoidance (red) and Scenario 3: one-way attraction (black) for intervals between an individual from species A followed by a capture of an individual from species B (top) and the reverse (bottom). The median interval duration (dashed line) and mean (dotted line) for each scenario and pair orientation are also shown. In all scenarios presented, data was taken from the simulations with 24 agents per species and a camera density of 4/km² for maximum data input. The avoidance scenario here illustrates maximum avoidance (zone radius 150 m). The attraction scenario is represented by moderate attraction (zone radius 75 m).

The size of the avoidance zone radius had no visible effect on interval duration length, even with maximum interaction strength, that is, radius 150 m, and maximum population and camera trap densities (Figure 4.6(a)). Conversely, for the attraction scenario, increasing the interaction strength, that is the attraction radius zone, reduced the median interval durations for 24 agents at all camera trap densities, and for 12 agents at the highest camera trap density (Figure 4.7). For maximum population and camera densities, and interaction zone radius, the interval durations are very short in the attraction scenario with a median of 40 seconds (Figure 4.6(b)) (vs 1.69 and 1.67 days for CRW and avoidance, respectively), and a mean of 1.03 days (vs 1.99 and 2.02 days for CRW and avoidance, respectively).

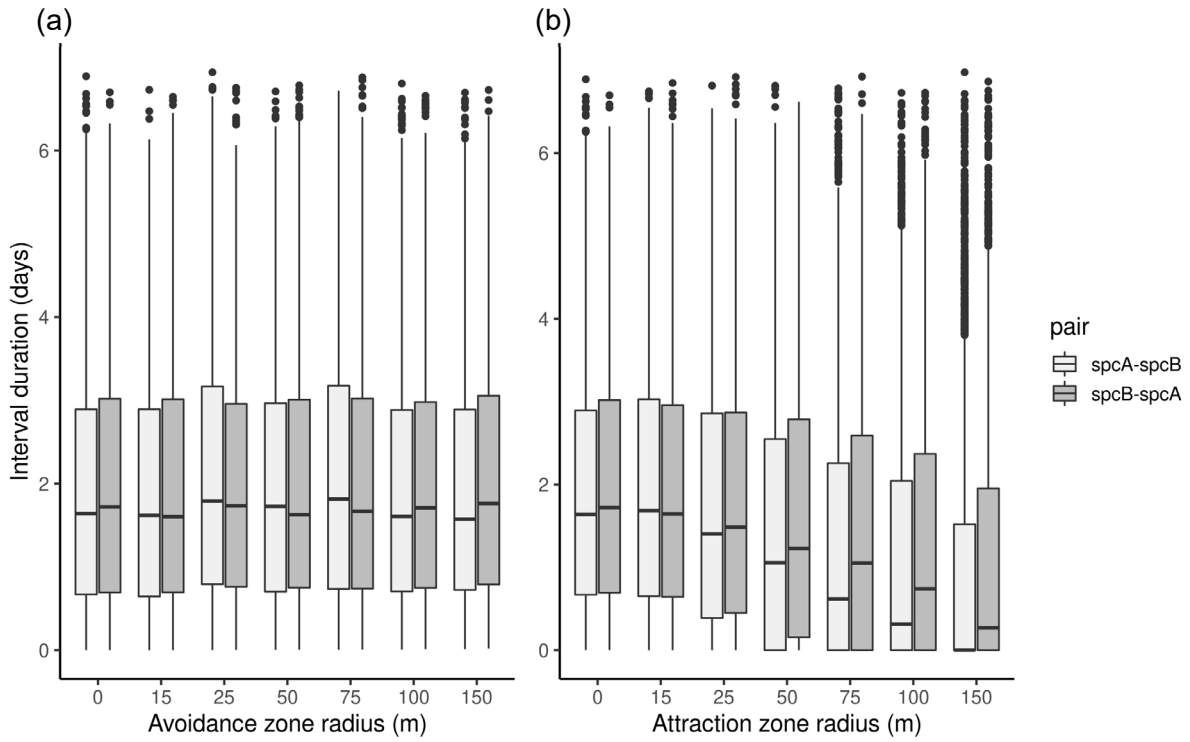


Figure 4.6: Boxplots for Scenarios 2 and 3 at maximum parameter values. (a) Scenario 2: one-way avoidance and (b) Scenario 3: one-way attraction, with maximum population size (24 agents) and camera trap density (4/km²). The x-axis displays interaction zone radius with 0 for a CRW.




Figure 4.7: Boxplots for Scenario 3: one-way attraction. The x-axis for each sub-plot shows the attraction zone radius. The plots show increasing camera density from 1/km² (left column), 2/km² (middle column) to 4/km² (right column), and increasing population size for each species from 6 (bottom row), 12 (middle row) to 24 (top row) agents per species.

4.4.1.3 Intervals analysis: interaction-detection model

For the CRW scenario, a linear model of interval durations following Niedballa et al. (2019) showed no significant difference in the interval duration for intervals BA and AB, as expected. There was also no significant effect on interval duration with increasing number of agents or camera density (Figure 4.8).




Figure 4.8: Model coefficient estimates for Scenario 1: CRW (control). Note that the intercept (not shown) contains AB intervals for 6 agents per species at a camera density of 1/km².

For the avoidance scenario, only maximum population and camera trap densities were found to be significant in reducing interval durations, with no significant difference between AB and BA intervals (Figure 4.9).




Figure 4.9: Model coefficient estimates for Scenario 2: one-way avoidance. Note that the intercept (not shown) contains AB intervals for a CRW with 6 agents per species and a camera density of 1/km².

In Scenario 3: one-way attraction, BA intervals were found to be significantly longer than AB intervals for attraction zone radii greater than or equal to 25 m (Figure 4.10). AB interval durations were also shown to significantly decrease with increased attraction zone radius (Figure 4.10).




Figure 4.10: Model coefficient estimates for Scenario 3: one-way attraction. Note that the intercept (not shown) contains AB intervals for a CRW with 6 agents per species and a camera density of 1/km².

Model-predicted interval durations show a trend for longer BA intervals than AB intervals in the avoidance scenario, but with overlap at all avoidance zone radii (Figure 4.11(a)). In the attraction scenario, the model-predicted interval durations tend to zero with increasing attraction zone radius (Figure 4.11(b)). At each radius, the durations for BA are longer than for AB, and with no overlap in values from an attraction zone radius of 25 m (Figure 4.11).

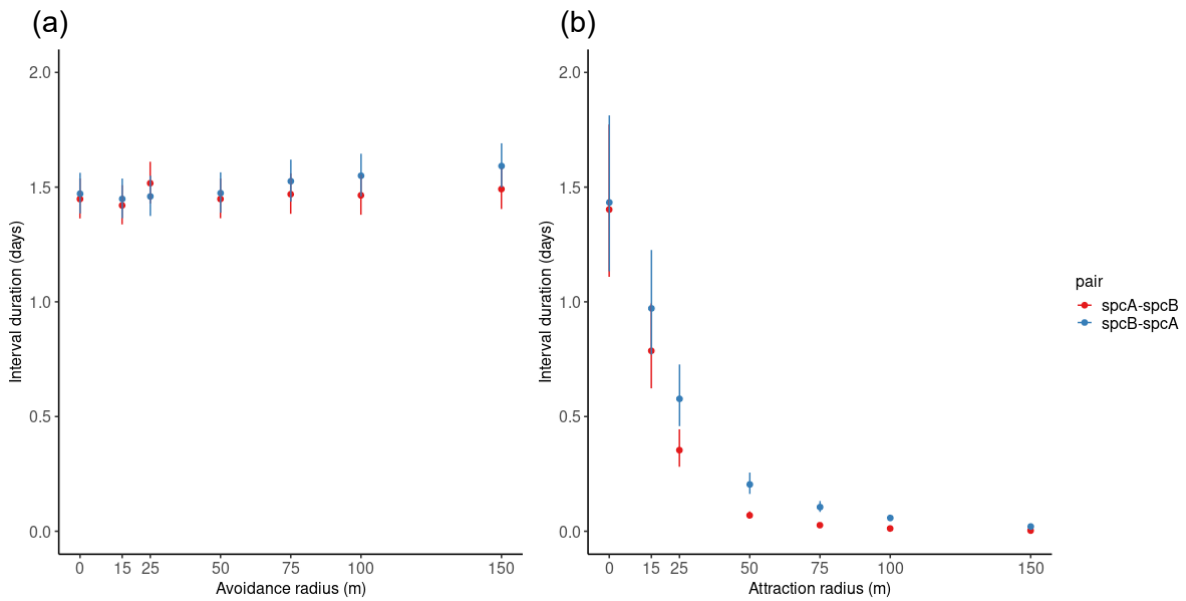


Figure 4.11: Model-predicted interval durations for Scenarios 2 and 3. (a) Scenario 2: one-way avoidance and (b) Scenario 3: one-way attraction. The x-axis shows the avoidance and attraction zone radius, respectively.

4.4.1.4 Maximum interval duration cap

Applying a cap on the interval duration included in the analysis reduces the number of datapoints. For setups with maximum parameter values, that is, maximum population size (24 agents per species), camera density ($4/\text{km}^2$) and interaction zone radius (150 m), the number of intervals extracted reduces from 3,007 and 5,619 for avoidance and attraction, respectively, to 232 and 3,116 for a cap of 6 hours (Figure 4.12(b)). For setups with minimum parameter values, that is the smallest population size (6 agents per species), camera density ($1/\text{km}^2$) and interaction zone radius (15 m), the number of intervals included in the analysis reduces from 62 and 67 for avoidance and attraction, respectively, to 3 and 16 at a cap of 6 hours (Figure 4.12(a))

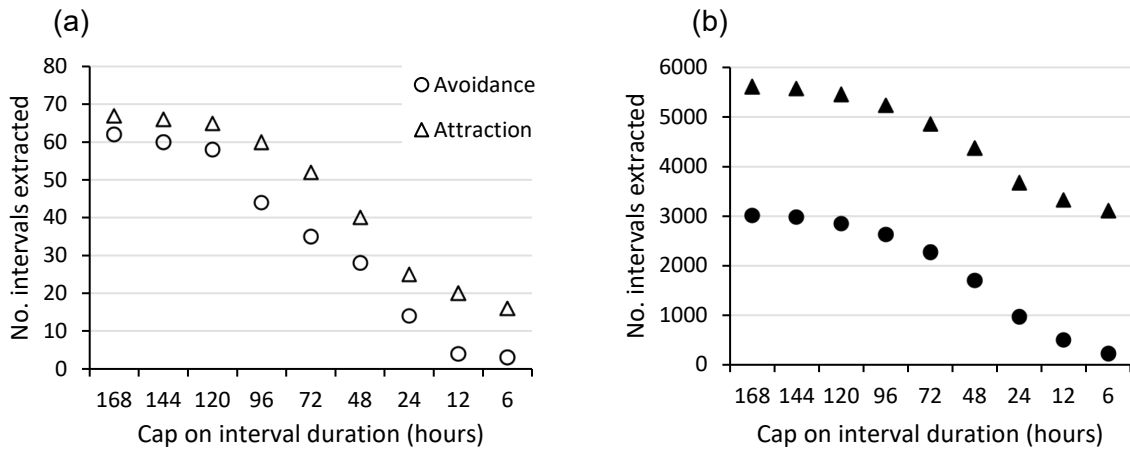


Figure 4.12: Cap on interval durations: number of intervals extracted. Number of intervals extracted for the avoidance (circles) and attraction (triangles) scenarios from setups with (a) minimum parameter values, i.e. 6 agents per species, camera density of 1/km² and an interaction zone radius 15 m, and (b) maximum parameter values, i.e. 24 agents per species, camera density of 4/km² and interaction zone radius 150 m. The x-axis shows the maximum interval duration permitted in hours.

With a cap imposed on the interval durations included in the model data, interactions can be detected at smaller radii (Figure 4.13). For the attraction scenario, the interaction could be detected for radii of 25 m or greater with no cap applied, but with a cap of 6 days or less, the interaction can also be detected at the smallest radius of 15 m (Figure 4.13(b)). For the avoidance scenario, the interaction can be detected for the largest radius of 150 m with a cap of 24 or 72 hours imposed, while interactions at a radius of 75 m or larger could be detected for a cap of 6 or 12 hours (Figure 4.13(a)).

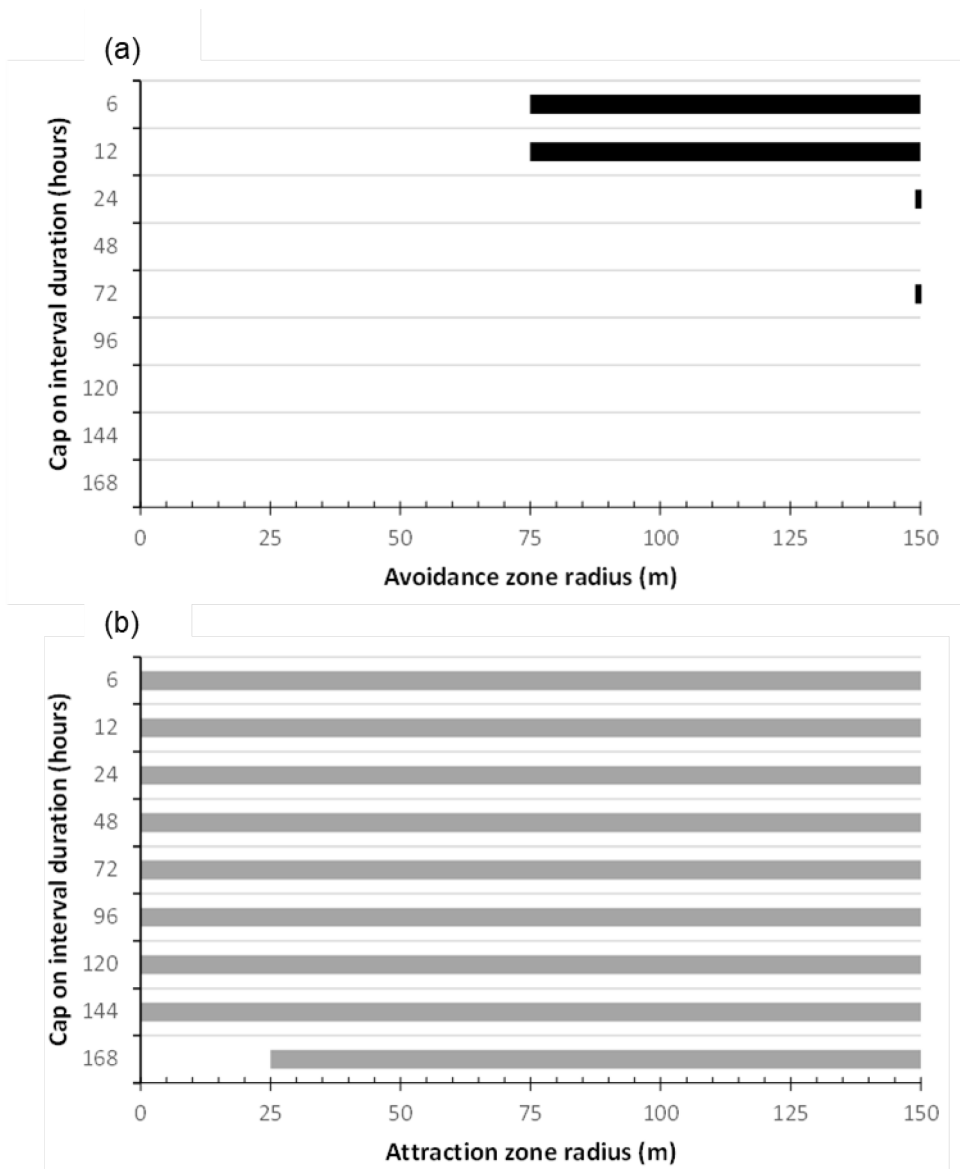


Figure 4.13: Maximum cap on interval durations: interaction zone radii at which the interaction can be detected. Interaction zone radii at which the (a) avoidance and (b) attraction interaction can be detected, that is, radii for which a significant difference between BA and AB interval coefficients is found in the corresponding model. Bars depict that the interaction is detectable at that radius, while blank space depicts radii at which the interaction is not detected in the model. The x-axis shows interaction zone radii and the y-axis shows the maximum interval duration in hours.

4.4.1.5 Applying the Niedballa et al. (2019) threshold

Applying a threshold of 50 intervals for each pair orientation found that a population size of 12 agents was sufficient to generate enough datapoints for all camera densities (Table 4.4). For the smallest population size of 6 agents and camera density of 1/km², only the setup with maximum attraction zone radius (150 m) met the requirements. Removing the corresponding data from these setups did not alter the results (Appendix 3: Figure 8.7 and Figure 8.8).

Table 4.4: Parameter setups that did not meet the Niedballa et al. requirements of 50 intervals per pair orientation for the analysis.

Scenario	No. agents	Camera density (/km ²)	Interaction zone radius (m)
1. CRW	6	1	NA
2. Avoidance	6	1	All
	6	2	75
3. Attraction	6	1	15 - 100

4.4.2 Combined scenario

4.4.2.1 Interactions and data collection

In Scenario 4: combined avoidance and attraction, the number of camera trap events generated, and intervals extracted, was impacted by both the camera density and number of agents, increasing with both, as per the previous scenarios (Figure 4.14).




Figure 4.14: Number of intervals extracted for Scenario 4: combined avoidance and attraction. Mean number of intervals extracted from the simulated camera trap data at camera densities of 1/km² (blue), 2/km² (black) and 4/km² (red).

4.4.2.2 Intervals analysis: distribution

Interval durations tended to decrease with increased attraction zone radius in the combined scenario (Figure 4.15). With avoidance radius greater than attraction radius, results are similar to the avoidance scenario, while for attraction radius greater than or equal to avoidance radius,

results are similar to the attraction scenario (Figure 4.15). For a medium-length attraction zone radius (75 m), interval durations for BA are consistently longer than for AB, but both increase with increased avoidance (Figure 4.15).

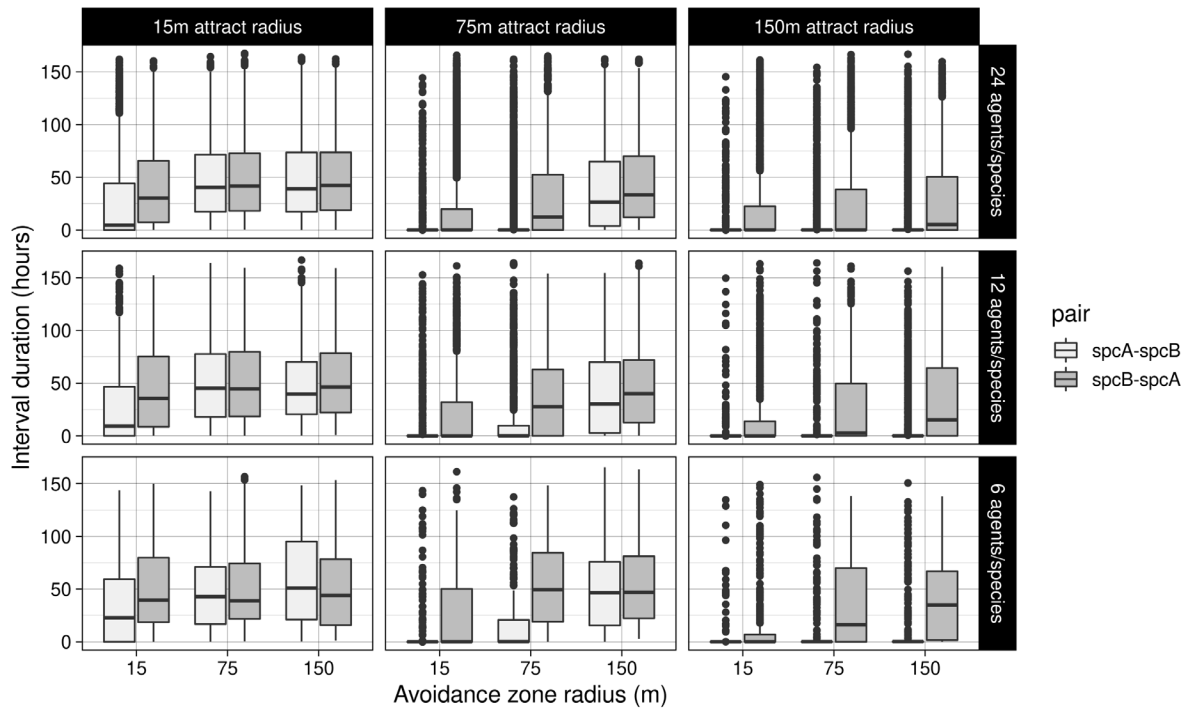


Figure 4.15: Boxplot for Scenario 4: combined attraction and avoidance, with a camera density of $4/\text{km}^2$. The x-axis for each sub-plot shows the avoidance zone radius. The plots show increasing attraction zone radius from 15 m (left column), 75 m (middle column) to 150 m (right column), and increasing population size for each species from 6 (bottom row), 12 (middle row) and 24 (top row) agents per species.

4.4.2.3 Intervals analysis: interaction-detection model

In the combined scenario, for small attraction zone radii, an increase in interval duration with increased avoidance radius for both pair orientations can be seen (Figure 4.16). As the attraction radius increases, it dominates the output, with shorter interval durations seen as a result. In all instances, the model-predicted results show longer intervals for BA than AB (Figure 4.16), without overlap where attraction strength is greater than or equal to avoidance strength.




Figure 4.16: Model-predicted interval durations for Scenario 4: combined attraction and avoidance. The x-axis for each subplot shows the avoidance zone radius, while the subplots from left-right show attraction zone radii 15m, 75 m and 150 m.

4.5 Discussion

This study has extended the approach previously identified as the best-powered for detecting spatiotemporal avoidance in a time series on one camera trap, to consider performance when considering triggers across a camera trap grid and with multiple individuals per species in the landscape. Using an ABM to model different interaction scenarios, I have shown that this method can detect avoidance and attraction behaviour. For the parameter values used, attraction could be detected at all but the weakest strength, that is, from an interaction zone radius equal to or greater than 25 m. Avoidance, however, could only be detected with a cap of 72 hours applied on the maximum interval durations included. This study has highlighted the difficulty in discerning the specific interaction at play, since all interaction scenarios produced the same qualitative results.

Increasing the avoidance zone radius had no impact on the number of camera trap triggers (Appendix 3: Figure 8.1), but an increase in attraction zone radius resulted in a higher number of camera trap triggers (Appendix 3: Figure 8.2). In the avoidance scenario, a single time-step can be enough to encompass the full avoidance interaction event, meaning the majority of time-steps are capturing a simple CRW rather than a continuous series of avoidance interactions. By contrast, in the attraction scenario, the predator moves closer to the prey at each time step until a kill occurs, meaning a trigger by a prey agent that is being stalked will be followed by a trigger from the stalking predator. It is also possible for a prey agent to be stalked by multiple predators since each predator will move towards whichever prey agent is closest within their attraction zone. This could result in clumping or herding where multiple

agents are close together. In this instance, if they were to be within the detection zone of a camera trap, many images would be captured.

For one-way avoidance, fewer instances of the shortest durations were seen than for the CRW scenario, particularly for the BA intervals, but the interaction could only be picked up by the linear model when the interaction strength was strongest, i.e. avoidance zone radius 150 m, and a cap of 72 hours was applied. This is in agreement with Niedballa et al. (2019) who found that the reliability of the method is dependent on interaction strength, with a strong interaction required. This suggests that small-scale avoidance could be difficult to detect in camera trap data. For detecting differences in species interactions, between habitat types for example, we would need a big shift in behaviour in order to detect it using this method. The attraction interaction in the model was stronger than for avoidance since the interaction persisted through multiple time-steps. Consequently, the distribution of interval durations was skewed towards lower values (Figure 4.5) and the model was able to detect the interaction from the smallest attraction zone radius with a cap of 6 days. Attraction was therefore dominant in the combined scenario, with avoidance strength only having an impact for weak attraction (15 m radius). In reality, avoidance behaviour can be triggered by scent, sound or other visual cues, besides the presence of an individual competitor or predator. Avoidance of a site might therefore be triggered by, and sustained due to, the scent of an individual after it has moved on, for example, which is not captured in this model.

Although this method is able to detect that an interaction is occurring, discerning the underlying mechanism is challenging since the three interaction scenarios present the same qualitative results: both avoidance of species B by species A, and attraction towards species A by species B, resulted in longer BA intervals than AB intervals. The complexity of ecological systems has previously been argued to blur the link between species interactions and measurable parameters, such as, co-occurrence (Blanchet et al., 2020). Similarly, the blur here between following and leading behaviour results in attraction appearing as avoidance behaviour. Predator-prey relationships in tropical forest have previously been shown to be weak (Brodie & Giordano, 2013) so we might need an alternative method or metric that is better able to detect these weaker interaction effects.

An avoidance or attraction zone radius was used here as a proxy for interaction strength. An alternative probabilistic approach to modelling interaction strength could be used, whereby interaction strength is represented as a probability of the interaction response occurring. The results here support the finding of Niedballa et al. (2019) that statistical power is strongly affected by interaction strength. In my analysis, an interaction strength of 1, represented by an avoidance zone radius of 150 m, could not provide reliable detection of the interactions for

the avoidance scenario. The analysis would subsequently be unable to detect a reduction in strength. Responses to land-use change or anthropogenic disturbance could include changes in interaction strength, so methods need to be robust to this.

In generating the camera trap data, both increasing the number of agents or the camera density resulted in an increase in the number of camera trap triggers. This is expected since increasing the density of either will result in any individual coming into proximity of a camera trap more frequently. In their study, Niedballa et al. (2019) found that more than 50 records per species were required for their method to produce reliable results, that is, for it to have power > 0.8. Interpreting that threshold in terms of the number of intervals here suggests that a camera trap density of 2/km² would be required for population sizes of 6 agents, representative of a population density of ~0.1/km², while a camera trap density of 1/km² is sufficient for larger populations (Table 4.4). Removing the corresponding data from these setups did not alter the results here as all setups were combined in the model (Appendix 3: Figure 8.7 and Figure 8.8). Applying this threshold in this context provides guidance to researchers on the suitability of this approach for any population density of interest and what camera density might be required for future interactions studies using camera trap surveys. However, these results should only be interpreted within the range of parameter values used.

The camera traps simulated in this study are always triggered if an animal is within the detection zone, and feature a 360° “lens”. In reality, camera traps can only be triggered and take images in the direction they are facing, and they can malfunction so that they fail to trigger. These aspects have not been included within our study but could be considered in future work. Camera traps can also be triggered by changes in light or movement of foliage by the wind, for example, but these “blank” images would be removed from analyses so this missing aspect would have no impact on our results. Animals can be seen to investigate camera traps in photos, and species sensitive to human disturbance might avoid areas with repeated human visitation (George & Crooks, 2006). Both of these would impact animal behaviour, and could consequently impact species interactions, but the general impact on animal behaviour due to camera traps is minimal and not the focus of this study.

The main focus of this study was on the ability to detect species interactions within camera trap data, but translating our results to real-world camera trap data relies on the assumptions that were built into our basic animal movement simulations. We note in particular that the ABM we constructed necessarily simplified some aspects of biological realism. These include absence of pauses in the animal movement; instead, each agent’s location was updated at each time-step with no possibility of staying in the same location. The velocity was chosen to replicate the speed at which an individual may cover a certain area so incorporated stalled

foraging behaviour to a degree, but we would expect individuals to pause to sleep for a period of time over the week in which the model ran. Since the purpose of the ABM was to generate camera trap triggers for the subsequent evaluation of the power needed to detect the modelled interactions, and animals would not be generating triggers while sleeping, sleep was not modelled. In future work, to develop the model such that it could more realistically represent the 24-hour activity pattern of specific species, periods of rest could be incorporated.

Another simplification in the ABM was the use of a constant movement speed, yet in the real world, prey that detect a nearby predator might either freeze or flee. Consequently, the addition of a freeze state and/or acceleration for prey fleeing attack, as well as a slow stalking pace and/or acceleration for predators giving chase, might better replicate predator-prey events. A time or distance-limit for predators following a prey individual would also allow the prey to escape predation. Including a maximum speed, which can be used while fleeing or giving chase, and applying a limit of 30% (Wheatley 2020), for example, for foraging or stalking would be one way to incorporate speed variation.

Finally, future work might want to consider the effect of environmental variation within the model. Animal movement has been shown to vary according to habitat, with individuals moving more slowly through complex habitats (Wheatley et al., 2020) and generally moving less in highly productive habitat areas as they are able to gather resources in a small area (Mueller et al., 2011). The inclusion in the model of refuges for prey to shelter in and obstacles for both predators and prey to navigate (Wheatley et al., 2020), or waterholes and foraging hotspots where individuals aggregate, would allow for more biological realism, such that interaction behaviour could be more accurately modelled and explored. The incorporation of a vegetation index to represent different habitat types (Tucker et al., 2021) would further allow behaviour to be tailored to the environment and for the impact of concurrent interactions and environmental effects to be investigated.

Chapter 5

5 Discussion and Conclusion

Species interactions are a vital component of ecosystems, yet are poorly understood. With ongoing climate change and anthropogenic disturbance, we must be able to capture, monitor and understand how ecosystem processes, such as species interactions, are impacted so that we are better positioned to mitigate against negative outcomes. Camera trap surveys provide us with the ability to capture large volumes of data with minimal disturbance to natural behaviour. The use of camera trap data for ecological analyses is a growing field, with the potential for transformative analysis, including in our understanding of species interactions. Reliable rapid processing of the images, as well as sensitive methods to detect interactions, are, however, still lacking.

5.1 Automated classification of camera trap images

5.1.1 Key findings

In Chapter 2, I applied automated image classification, in the form of convolutional neural networks, to a large, species-rich camera trap dataset to explore the performance of three established architectures and to investigate generalisability across a gradient of habitat disturbance in tropical rainforest. Classification performance was found to be worse on both unseen camera locations, and unseen disturbance levels. These results highlighted the ongoing issues with poor generalisability to unseen camera locations, as well as the additional problem of poor generalisability to changes in background associated with varying levels of habitat disturbance within a single camera trap dataset.

Training across multiple disturbance levels improved generalisability, suggesting that an awareness of variation in habitat backgrounds is required when planning a camera trap survey intended for automated classification, and when training the classifier. Using a bounding-box object detector in conjunction with the classifier also improved both overall performance and generalisability, but did not completely eliminate the problem. These results support the aggregation of images from multiple camera trap surveys to enhance available training data and subsequently improve classification. Further, as habitats change over time, whether naturally or through anthropogenic disturbance, classifiers similarly need to be considered as dynamic and retrained with images representative of the modified backgrounds.

5.1.2 Recent developments

A recent step-forward in the state of the art is the use of an ensemble of Data-efficient image Transformers (DeiT) (Kyathanahally et al., 2022). DeITs differ from CNNs in that they consider image patches, rather than individual pixels, and employ an attention mechanism that identifies the most relevant part of the image. Ensemble models take an average of predictions from several individually trained models. Kyathanally et al. (2022) validated an Ensemble of DeITs (EDeITs) on ten ecological image datasets, including images of plankton, coral reefs, insects, birds and large animals, in both colour and black-and white, with and without backgrounds. Single-model performance of DeITs matched alternative CNN models, while the ensembles of DeITs significantly outperformed the previous state-of-the-art ensemble of CNN models, in both accuracy and better classification of rare species (Kyathanahally et al., 2022). The increase in performance seen in the study was attributed to higher disagreement between individual models, compared with CNNs, which tend to produce similar predictions.

The DeITs used by Kyathanally et al. (2022) had a similar number of parameters and required similar computational power to the equivalent CNNs. DeITs also have the advantage of being more transparent than CNNs in that they identify the part of the image on which the model is focussed, and they are more robust to perturbations such as occlusions (Naseer et al., 2021). They have additionally been shown to be more robust to data shift (Paul & Chen, 2022), whereby test images differ from the training images, due to sampling method, instrument degradation, environmental variation, etc. (Moreno-Torres et al., 2012). This robustness to data shift could result in increased generalisability across camera locations and disturbance levels.

The application of DeITs to an aggregation of camera trap data collected across multiple locations, habitats and disturbances, such as that on Wildlife Insights (Ahumada et al., 2020) could be a critical step towards making robust classifiers.

5.2 Detection of species interactions

5.2.1 Key findings

In Chapter 3, I set out to test for spatiotemporal avoidance of humans by bearded pigs. Davison et al. (2019) previously showed that bearded pigs shift their activity pattern from diurnal in primary forest, to crepuscular in logged forest and nocturnal in oil palm plantations. In explanation of this result, the authors hypothesised that this behavioural shift could be due to bearded pigs avoiding plantation workers, who were active in the morning, or avoiding the higher daytime temperatures that occur in oil palm (Davison et al., 2019). My analysis did not

find evidence to support the avoidance of humans hypothesis, but it was not possible to conclude whether this was due to the absence of the interaction occurring, or the inability to detect it due to insufficient sensitivity in the method, or insufficient data.

It is possible that any behavioural responses that may be occurring are happening at a temporal scale that is too fine for the method to detect. The average interval times extracted from the data ranged from 1 to 10 days, so it might be unrealistic to expect to detect any responses happening on the scale of minutes, for example. To focus in on shorter timescales is at the cost of exclusion of data. This analysis, therefore, highlighted challenges in being able to power methods such that reliable conclusions can be drawn about species interactions from a camera trap dataset representative of a long-term monitoring project. These interactions remain important to understand and capture, particularly for species such as the bearded pig, which modifies the habitat and is able to persist in disturbed habitat. This study also highlighted the challenge in assessing the effectiveness of methods for detecting interactions in real data. Previous research had highlighted a linear model of time intervals between species detections as best-powered for detecting spatiotemporal avoidance between two species (Niedballa et al., 2019), but in evaluating performance on a simulated dataset for one camera trap, many of the challenges in real-life camera trap datasets were overlooked.

To explore how the results from Chapter 3 are impacted by data availability, as well as population size, camera trap density, and interaction strength, in Chapter 4 I constructed an agent-based model (ABM) to further evaluate our ability to detect avoidance, as well as predatory stalking, in a simulated camera trap survey. This analysis highlighted the difficulty in discerning the type of behaviour detected since both avoidance and attraction presented the same qualitative result via the method used; knowledge of the species' ecology and ecosystem would be required to help inform this.

This analysis also demonstrated the need for a strong interaction. Interaction strength was modelled here as the distance from which the avoidance or attraction behaviour is triggered, with the weakest and strongest interactions represented by an interaction zone radius of 15 m and 150 m, respectively. For attraction behaviour, whereby the interaction occurred over multiple sequential time-steps until a kill occurred, the interaction was detected at all but the weakest strength with no cap imposed. For avoidance behaviour, however, where the interaction could be resolved in just one time-step, the strongest interaction could only be detected with a cap imposed on the interval durations considered. This suggests that any similarly small-scale changes in behaviour, as a result of habitat disturbance, for example, would be similarly difficult to detect due to the challenges in capturing enough data to perform the analysis. These behaviours could therefore be missed using existing methods.

5.2.2 Alternative methods for inferring species interactions

5.2.2.1 Extended occupancy models

Other methodological approaches to investigating species interactions using camera trap data include the use of occupancy models. In their critique of this approach, Blanchet et al. (2020) argued that species interactions could be wrongly inferred from co-occurrence signals that arise simply due to habitat preferences, rather than an interspecific interaction. Recent approaches have consequently extended basic occupancy models to incorporate a continuous-time detection process (Kellner et al., 2022; Parsons et al., 2022). This extension allows for simultaneous analysis of species interactions in space and time, through site occupancy and activity patterns, respectively, while accounting for imperfect detection (Kellner et al., 2022). The addition of the temporal component to the model lends support to any consequent inference of interspecific interactions by allowing for corroboration of spatial findings (Parsons et al., 2022).

Applied to a camera trap survey of coyote and white-tailed deer in the USA, Kellner et al. (2022) found evidence of both spatial and temporal interactions between the two species. At sites where coyotes occurred, deer were found to have greater overall detection intensity. This co-occurrence pattern was thought to be a result of either predation, whereby coyotes seek sites in which deer are present, or a shared preference for edge habitat. There was also a shift in deer detection distribution to proportionally more diurnal in the presence of coyotes, suggesting the deer were avoiding the primarily nocturnal coyotes (Kellner et al., 2022). The authors further found that coyotes were less likely to occur at hunting sites; this was in direct contrast with previous analyses of the same data set (Kays et al., 2017), but was consistent with the earlier study's a priori predictions, that hunting would locally reduce coyote abundance. The inclusion of the temporal data in the approach by Kellner was argued to result in a more intuitive finding, and highlights the importance of accounting for species interactions in occupancy models (Kellner et al., 2022).

In their application to carnivore competition along an urbanisation gradient, Parsons et al. (2022) also found evidence of spatiotemporal avoidance of coyote by gray foxes using Kellner's approach. In a previous analysis, which used spatial data only, Parsons et al. (2019) found no support for their hypothesis that smaller competitors (gray fox) use humans as shields against larger competitors (coyotes) at higher levels of urbanisation. By incorporating temporal data, however, they were able to uncover a behavioural mechanism, that is, gray foxes avoided rural areas with sparse forest after a coyote had passed by. The spatiotemporal analysis provides evidence for an interaction occurring between the two species, whereas a pattern identified using spatial data only could be the result of an interaction occurring or a

response to unmodelled covariates. The results of this study re-emphasise the need to consider temporal behaviour patterns when investigating species interactions mediated by human disturbance (Parsons et al., 2022).

In highlighting the limitations of previous analyses of species activity patterns and interactions using passive data, Kellner et al. (2022) state that these analyses either ignored imperfect detection, were limited to sites where all species cooccur, or required binning data such that fine-scale changes in activity patterns, which indicate temporal partitioning in site use, were lost. The linear model approach used in this thesis made use of the continuous detection data camera trap surveys collect, and has the advantages of being quick and straightforward to implement, as well as allowing for the addition of covariates to be modelled. A significant limitation in this approach, however, is the use of the intervals between species as the response variable; even with a large data set, it is difficult to obtain enough data to power the method. Conversely, the approach taken by Parsons and Kellner is able to make use of all independent detections in the data (Kellner et al., 2022), thus no data is lost and it has more wide-ranging applicability since it allows for analysis of smaller data sets.

5.2.2.2 Multilayer ecological networks

Studies conducted on a species level are important for understanding individual species' responses to habitat disturbance and the consequences for their survival. They can also inform predicted responses of similar, unstudied species. These studies collectively report a variety of responses, however, which cannot inform generalised management or conservation decisions (Meijaard & Sheil, 2008). It is therefore necessary to capture and measure general trends and broader, more consistent responses. These can, in turn, be used to predict community responses to disturbance. Identifying patterns in community responses will inform our knowledge of interactions between functional or trophic groups. From this, we can gain a more mechanistic understanding of how changes in land use influence populations, communities and ecosystems (Bruckerhoff et al., 2020).

Ecological communities, comprising species interacting dynamically, can be naturally represented by ecological networks (Hagen et al., 2012; García-Callejas et al., 2018). Ecological networks consist of nodes, typically representing species or populations, connected by links that represent interactions or movement. The occurrence and intensity of species interactions varies according to the environment, species abundance and occurrence or absence of other interactions (Poisot et al., 2015). Similarly, the structure of ecological networks has been shown to vary across gradients of habitat modification (Tylianakis et al., 2007) and among habitat fragments (Mclaughlin et al., 2010). Several ecological applications

that could particularly benefit from the use of multilayer networks are adaptation in space, niche partitioning and network stability (Hutchinson et al., 2019).

Network studies have often focussed on one type of interaction, such as, trophic or mutualistic, but this excludes other interactions that could be significantly impacting community structure; consequently, their impact is not captured (García-Callejas et al., 2018). Multilayer networks offer a more comprehensive approach since they can incorporate multiple types of interaction, and offer the potential to model system dynamics (García-Callejas et al., 2018). Multilayer networks also offer the opportunity to model how disturbances might percolate through the network, to inform management practice (Hutchinson et al., 2019). This could extend to an evaluation of connectivity scenarios or the identification of keystone species for conservation purposes (Pilosof et al., 2017). The identification of patterns across layers could also inform our understanding of the consistency of species roles and interactions (Mora et al., 2018).

In best practice, values attributed to links between species and between layers should be measured directly, but they can also be inferred (Pilosof et al., 2017). Since comparable data across multiple interactions is challenging to acquire, studies can also be aggregated *a posteriori* (García-Callejas et al., 2018). Alternatively, values can be evaluated systematically to explore how the relative compositions of the two interaction types affects the system overall (Pilosof et al., 2017). One particularly challenging element to quantify is interaction strength (García-Callejas et al., 2018). The probability of an interaction occurring can be approximated from the relative abundances of the species populations involved, which can be extended to approximate interaction strength by incorporating an estimate of the effect of the interaction on each species (P. Vázquez et al., 2007). Camera traps have been identified as a potential data source to parameterise multilayer ecological networks (Hutchinson et al., 2019).

Given the importance of species interactions and how they link to ecosystem function, there is a need for more studies to consider the behaviour of animals when studying the impact of human disturbance on communities and ecosystems (Rahman & Candolin, 2022). Modelling ecosystems as a network of interacting species facilitates investigations into how the impacts of disturbance can permeate throughout the ecosystem (Hutchinson et al., 2019). Using the SAFE Project camera trap dataset, with its span across land-use categories and disturbance levels, to parameterise multilayer networks, could help provide mechanistic insight into the processes and changes occurring as a result of the anthropogenic disturbance in tropical forest.

5.3 Suggested future research

Although automated classification of species in camera trap images is possible for large volumes of data, challenges remain in the applicability of automated classification across habitats. Future work might want to investigate the impact of biases arising from poor generalisability on the subsequent ecological studies undertaken. These misclassifications could bias state variable estimates, such as population density or species richness, as well as the statistical power to detect difference between environmental treatments. One such study into the impact of misclassifications on occupancy modelling found that including false positives at a rate of 3% or more of the data resulted in overestimation of species distributions and reduced estimates of occurrence associations (Clare et al., 2019). Similar studies could provide guidance on the required performance of automated classification methods in order to mitigate resultant errors in analyses of species interactions using camera trap data.

Analysis of the performance of the CNNs used in this thesis on individual species showed that high performance was achieved for some relatively rare species, while there was variation in performance on common species. An investigation into to what degree morphological variation within and between species contributes to the ability of networks to better learn some species over others would also be insightful.

The uptake of automated classification methods will allow for larger camera trap datasets to be collected and collated. This will be an important step forward given the challenge in extracting enough data points for the analysis here, despite the large size of the dataset. Even with large volumes of simulated data, interaction behaviours were not consistently detected. Alternative methods, such as occupancy models with an incorporated continuous-time detection process (Kellner et al., 2022; Parsons et al., 2022) offer the potential to gain insight into spatiotemporal interactions, without such data limitations. More complex methods, such as the use of networks, allow for insight into how any changes to these species interactions impact other elements of the ecosystem. More work is needed to better quantify species' responses to anthropogenic disturbance so that we are able parametrise these methods and use them to identify the species most affected. Studies that focus on keystone species, ecosystem engineers and dominant species are particularly important, as they have a large influence on the processes and other species in the ecosystem (Rahman & Candolin, 2022).

Future work could look to construct multilayer ecological networks that span the habitat disturbance levels in the SAFE dataset. Connections between land-use fragments would allow for evaluation of fragment composition that allow species to persist and inform management strategies for tropical forests undergoing land-use change. Constructing a network such that connections between species in each layer represent competition and links between layers

represent predation, we could vary the balance of weights between the two interaction types to assess their relative influences on community structure. Alternatively, constructing a network with layers representing time points would allow an analysis of how the community interactions change over time.

A community-based approach to the simulated camera trap survey in this thesis could also improve aspects of biological realism in the model. Although the ABM used to simulate the animal movement was realistic in terms of animal and camera trap density, the inclusion of noise in terms of additional species would more closely represent realistic camera trap datasets, and the possible dilution of an interaction by intervening species. The model could also be tailored to specific species within a chosen habitat. This would enable closer realism in terms of the species' behaviour when in the presence of a competitor, predator or prey. Similarly, the addition of environmental variation, such as watering or foraging hotspots, and refuges, would allow for modelling of more realistic predator-prey interactions by including prey-aggregation zones, and opportunities to escape, respectively. Analyses could then evaluate the influence of environmental variation on direct interactions, to help inform discernment of interaction behaviour patterns in real-world data.

5.4 Conclusion

As camera trap datasets become more abundant, and the use of machine learning for automated classification becomes more commonplace, it will be critically important to ensure that estimation of changes in ecosystem function and composition are not biased by methodological choices in identification of species or their interactions. This is particularly important in the context of current global biodiversity loss, for monitoring the impacts of anthropogenic activities on ecosystems and mitigating further declines. Increasing volumes of data and improved methodologies for analysing camera trap data will facilitate further studies into behavioural shifts and species interactions, such that their potential as an early warning system can be achieved.

Bibliography

- Ahumada, J. A., Fegraus, E., Birch, T., Flores, N., Kays, R., O'Brien, T. G., ... Dancer, A. (2020). Wildlife Insights: A Platform to Maximize the Potential of Camera Trap and Other Passive Sensor Wildlife Data for the Planet. *Environmental Conservation*, 47(1), 1–6. doi:10.1017/S0376892919000298
- Amir, Z., Sovie, A., & Luskin, M. S. (2022). Inferring predator–prey interactions from camera traps: A Bayesian co-abundance modeling approach. *Ecology and Evolution*, 12(12), 1–15. doi:10.1002/ece3.9627
- Andersen, D., Yi, Y., Borzée, A., Kim, K., Moon, K. S., Kim, J. J., ... Jang, Y. (2022). Use of a spatially explicit individual-based model to predict population trajectories and habitat connectivity for a reintroduced ursid. *Oryx*, 56(2), 298–307. doi:10.1017/S0030605320000447
- Bascompte, J., & Jordano, P. (2007). Plant-Animal Mutualistic Networks: The Architecture of Biodiversity. *Annual Review of Ecology, Evolution, and Systematics*, 38(1), 567–593. doi:10.1146/annurev.ecolsys.38.091206.095818
- Beaudrot, L., Acevedo, M. A., Lessard, J. P., Zvoleff, A., Jansen, P. A., Sheil, D., ... Ahumada, J. (2019). Local temperature and ecological similarity drive distributional dynamics of tropical mammals worldwide. *Global Ecology and Biogeography*, 28(7), 976–991. doi:10.1111/geb.12908
- Beery, S., Liu, Y., Morris, D., Piavis, J., Kapoor, A., Meister, M., ... Perona, P. (2019). Synthetic Examples Improve Generalization for Rare Classes. *ArXiv Preprint ArXiv:1904.05916*. Retrieved from <http://arxiv.org/abs/1904.05916>
- Beery, S., Morris, D., & Yang, S. (2019). Efficient Pipeline for Camera Trap Image Review. Retrieved from <http://arxiv.org/abs/1907.06772>
- Beery, S., Van Horn, G., & Perona, P. (2018). Recognition in Terra Incognita. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11220 LNCS, 472–489. doi:10.1007/978-3-030-01270-0_28
- Beery, S., Wu, G., Rathod, V., Votel, R., & Huang, J. (2020). Context R-CNN: Long Term Temporal Context for Per-Camera Object Detection. *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13072–13082. doi:10.1109/CVPR42600.2020.01309.
- Blanchet, F. G., Cazelles, K., & Gravel, D. (2020). Co-occurrence is not evidence of ecological

- interactions. *Ecology Letters*, 23(7), 1050–1063. doi:10.1111/ele.13525
- Brodie, J. F., & Giordano, A. (2013). Lack of trophic release with large mammal predators and prey in Borneo. *Biological Conservation*, 163, 58–67. doi:10.1016/j.biocon.2013.01.003
- Brodie, J. F., Giordano, A. J., & Ambu, L. (2015). Differential responses of large mammals to logging and edge effects. *Mammalian Biology*, 80(1), 7–13. doi:10.1016/j.mambio.2014.06.001
- Bruckerhoff, L. A., Connell, R. K., Guinnip, J. P., Adhikari, E., Godar, A., Gido, K. B., ... Welti, E. (2020). Harmony on the prairie? Grassland plant and animal community responses to variation in climate across land-use gradients. *Ecology*, 0(0), 1–17. doi:10.1002/ecy.2986
- Buda, M., Maki, A., & Mazurowski, M. A. (2018). A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106, 249–259. doi:10.1016/j.neunet.2018.07.011
- Burton, A. C., Neilson, E., Moreira, D., Ladle, A., Steenweg, R., Fisher, J. T., ... Boutin, S. (2015). Wildlife camera trapping: A review and recommendations for linking surveys to ecological processes. *Journal of Applied Ecology*, 52(3), 675–685. doi:10.1111/1365-2664.12432
- Buschke, F. T., Brendonck, L., & Vanschoenwinkel, B. (2015). Simple mechanistic models can partially explain local but not range-wide co-occurrence of African mammals. *Global Ecology and Biogeography*, 24(7), 762–773. doi:10.1111/geb.12316
- Caravaggi, A., Banks, P. B., Burton, A. C., Finlay, C. M. V., Haswell, P. M., Hayward, M. W., ... Wood, M. D. (2017). A review of camera trapping for conservation behaviour research. *Remote Sensing in Ecology and Conservation*, 3(3), 109–122. doi:10.1002/rse2.48
- Caravaggi, A., Zaccaroni, M., Riga, F., Schai-Braun, S. C., Dick, J. T. A., Montgomery, W. I., & Reid, N. (2016). An invasive-native mammalian species replacement process captured by camera trap survey random encounter models. *Remote Sensing in Ecology and Conservation*, 2(1), 45–58. doi:10.1002/rse2.11
- Carothers, J. H., Jaksić, F. M., & Jaksic, F. M. (1984). Time as a Niche Difference: The Role of Interference Competition. *Oikos*, 42(3), 403. doi:10.2307/3544413
- Casula, P., Luiselli, L., Milana, G., & Amori, G. (2017). Habitat structure and disturbance affect small mammal populations in Mediterranean forests. *Basic and Applied Ecology*, 19, 76–83. doi:10.1016/j.baae.2016.11.003
- Chapman, P. M., Wearn, O. R., Riutta, T., Carbone, C., Rowcliffe, J. M., Bernard, H., & Ewers,

- R. M. (2018). Inter-annual dynamics and persistence of small mammal communities in a selectively logged tropical forest in Borneo. *Biodiversity and Conservation*, 27(12), 3155–3169. doi:10.1007/s10531-018-1594-y
- Clare, J. D. J., Frett, S., Townsend, P. A., Singh, A., Anhalt-Depies, C., Zuckerberg, B., ... Stenglein, J. L. (2019). Making inference with messy (citizen science) data: when are data accurate enough and how can they be improved? *Ecological Applications*, 29(October 2018), e01849. doi:10.1002/eap.1849
- Cusack, J. J., Dickman, A. J., Kalyahe, M., Rowcliffe, J. M., Carbone, C., MacDonald, D. W., & Coulson, T. (2017). Revealing kleptoparasitic and predatory tendencies in an African mammal community using camera traps: a comparison of spatiotemporal approaches. *Oikos*, 126(6), 812–822. doi:10.1111/oik.03403
- Daniel, O. Z., Heon, S. P., Bernard, H., Norman, D. L., Wearn, O. R., Deere, N. J., ... Ewers, R. M. (2022). Mammals in an oil palm dominated landscape. *The Planter, Kuala Lumpur*, 98(1151), 97–116.
- Davis, C. L., Rich, L. N., Farris, Z. J., Kelly, M. J., Di Bitetti, M. S., Blanco, Y. Di, ... Miller, D. A. W. (2018). Ecological correlates of the spatial co-occurrence of sympatric mammalian carnivores worldwide. *Ecology Letters*, 21(9), 1401–1412. doi:10.1111/ele.13124
- Davison, C. W., Chapman, P. M., Bernard, H., Ewers, R. M., & Wearn, O. R. (2019). Shifts in the demographics and behavior of bearded pigs (*Sus barbatus*) across a land-use gradient. *Biotropica*, 00(July), 1–11. doi:10.1111/btp.12724
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Denis, T., Richard-Hansen, C., Brunaux, O., Guitet, S., & Hérault, B. (2019). Birds of a feather flock together: Functionally similar vertebrates positively co-occur in Guianan forests. *Ecosphere*, 10(3), e02566. doi:10.1002/ecs2.2566
- Doherty, T. S., Glen, A. S., Nimmo, D. G., Ritchie, E. G., & Dickman, C. R. (2016). Invasive predators and global biodiversity loss. *Proceedings of the National Academy of Sciences*, 113(40), 11261–11265. doi:10.1073/pnas.1602480113
- Durant, S. M. (1998). Competition refuges and coexistence: An example from Serengeti carnivores. *Journal of Animal Ecology*, 67(3), 370–386. doi:10.1046/j.1365-2656.1998.00202.x

- Easter, T., Bouley, P., & Carter, N. (2020). Intraguild dynamics of understudied carnivores in a human-altered landscape. *Ecology and Evolution*, 10(12), 5476–5488. doi:10.1002/ece3.6290
- Ebenman, B., & Jonsson, T. (2005). Using community viability analysis to identify fragile systems and keystone species. *Trends in Ecology and Evolution*, 20(10), 568–575. doi:10.1016/j.tree.2005.06.011
- Erbs, F., Elwen, S. H., & Gridley, T. (2017). Automatic classification of whistles from coastal dolphins of the southern African subregion. *The Journal of the Acoustical Society of America*, 141(4), 2489–2500. doi:10.1121/1.4978000
- Ewers, R. M., Didham, R. K., Fahrig, L., Ferraz, G., Hector, A., Holt, R. D., ... Turner, E. C. (2011). A large-scale forest fragmentation experiment: The stability of altered forest ecosystems project. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1582), 3292–3302. doi:10.1098/rstb.2011.0049
- Frey, S., Fisher, J. T., Burton, A. C., & Volpe, J. P. (2017). Investigating animal activity patterns and temporal niche partitioning using camera-trap data: challenges and opportunities. *Remote Sensing in Ecology and Conservation*, 3(3), 123–132. doi:10.1002/rse2.60
- García-Callejas, D., Molowny-Horas, R., & Araújo, M. B. (2018). Multiple interactions networks: towards more realistic descriptions of the web of life. *Oikos*, 127(1), 5–22. doi:10.1111/oik.04428
- Gaynor, K. M., Hohnowski, C. E., Carter, N. H., & Brashares, J. S. (2018). The influence of human disturbance on wildlife nocturnality. *Science*, 360(6394), 1232–1235. doi:10.1126/science.aar7121
- George, S. L., & Crooks, K. R. (2006). Recreation and large mammal activity in an urban nature reserve. *Biological Conservation*, 133(1), 107–117. doi:10.1016/j.biocon.2006.05.024
- Gido, K. B., Propst, D. L., Whitney, J. E., Hedden, S. C., Turner, T. F., & Pilger, T. J. (2019). Pockets of resistance: Response of arid-land fish communities to climate, hydrology, and wildfire. *Freshwater Biology*, 64(4), 761–777. doi:10.1111/fwb.13260
- Greenville, A. C., Wardle, G. M., & Dickman, C. R. (2017). Desert mammal populations are limited by introduced predators rather than future climate change. *Royal Society Open Science*, 4(11). doi:10.1098/rsos.170384
- Hagen, M., Kissling, W. D., Rasmussen, C., De Aguiar, M. A. M., Brown, L. E., Carstensen,

- D. W., ... Olesen, J. M. (2012). *Biodiversity, Species Interactions and Ecological Networks in a Fragmented World. Advances in Ecological Research* (1st ed., Vol. 46). Elsevier Ltd. doi:10.1016/B978-0-12-396992-7.00002-2
- Haidir, I. A., Macdonald, D. W., & Linkie, M. (2018). Assessing the spatiotemporal interactions of mesopredators in Sumatra's tropical rainforest. *PLoS ONE*, 13(9), 1–18. doi:10.1371/journal.pone.0202876
- Harmsen, B. J., Foster, R. J., Silver, S. C., Ostro, L. E. T., & Doncaster, C. P. (2009). Spatial and temporal interactions of sympatric jaguars (*panthera onca*) and pumas (*puma concolor*) in a neotropical forest. *Journal of Mammalogy*, 90(3), 612–620. doi:10.1644/08-MAMM-A-140R.1
- Harvey, P. (2017). ExifTool: Read, write and edit meta information. Retrieved from <http://www.sno.phy.queensu.ca/~phil/exiftool>
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2016-Decem*, 770–778. doi:10.1109/CVPR.2016.90
- Hebblewhite, M., & Haydon, D. T. (2010). Distinguishing technology from biology: A critical review of the use of GPS telemetry data in ecology. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1550), 2303–2312. doi:10.1098/rstb.2010.0087
- Hethcoat, M. G., & Chalfoun, A. D. (2015). Towards a mechanistic understanding of human-induced rapid environmental change: A case study linking energy development, nest predation and predators. *Journal of Applied Ecology*, 52(6), 1492–1499. doi:10.1111/1365-2664.12513
- Hofmeester, T. R., Rowcliffe, J. M., & Jansen, P. A. (2017). A simple method for estimating the effective detection distance of camera traps. *Remote Sensing in Ecology and Conservation*, 3(2), 81–89. doi:10.1002/rse2.25
- Hutchinson, M. C., Bramon Mora, B., Pilosof, S., Barner, A. K., Kéfi, S., Thébault, E., ... Stouffer, D. B. (2019). Seeing the forest for the trees: Putting multilayer networks to work for community ecology. *Functional Ecology*, 33(2), 206–217. doi:10.1111/1365-2435.13237
- IPBES. (2019). Global assessment report of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services. *Brondízio, Eduardo Sonnewend Settele, Josef Díaz, Sandra Ngo, Hien Thu (Eds)*, 1082.

- IUCN. (n.d.). Red List of Threatened Species. Retrieved 17 January 2023, from <https://www.iucnredlist.org/species/41772/123793370>
- Jayadevan, A., Nayak, R., Karanth, K. K., Krishnaswamy, J., DeFries, R., Karanth, K. U., & Vaidyanathan, S. (2020). Navigating paved paradise: Evaluating landscape permeability to movement for large mammals in two conservation priority landscapes in India. *Biological Conservation*, 247(November 2019), 108613. doi:10.1016/j.biocon.2020.108613
- Jetz, W., McGeoch, M. A., Guralnick, R., Ferrier, S., Beck, J., Costello, M. J., ... Turak, E. (2019). Essential biodiversity variables for mapping and monitoring species populations. *Nature Ecology and Evolution*, 3(4), 539–551. doi:10.1038/s41559-019-0826-1
- Karanth, K. U., Srivathsa, A., Vasudev, D., Puri, M., Parameshwaran, R., & Kumar, N. S. (2017). Spatio-temporal interactions facilitate large carnivore sympatry across a resource gradient. *Proceedings of the Royal Society B: Biological Sciences*, 284(1848), 20161860. doi:10.1098/rspb.2016.1860
- Kays, R., Parsons, A. W., Baker, M. C., Kalies, E. L., Forrester, T., Costello, R., ... McShea, W. J. (2017). Does hunting or hiking affect wildlife communities in protected areas? *Journal of Applied Ecology*, 54(1), 242–252. doi:10.1111/1365-2664.12700
- Kellner, K. F., Parsons, A. W., Kays, R., Millsbaugh, J. J., & Rota, C. T. (2022). A Two-Species Occupancy Model with a Continuous-Time Detection Process Reveals Spatial and Temporal Interactions. *Journal of Agricultural, Biological, and Environmental Statistics*, 27(2), 321–338. doi:10.1007/s13253-021-00482-y
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 1097–1105. doi:10.1201/9781420010749
- Kyathanahally, S. P., Hardeman, T., Reyes, M., Merz, E., Bulas, T., Brun, P., ... Baity-Jesi, M. (2022). Ensembles of data-efficient vision transformers as a new paradigm for automated classification in ecology. *Scientific Reports*, 12(1), 1–11. doi:10.1038/s41598-022-21910-0
- Latham, M. C., Anderson, D. P., Norbury, G., Price, C. J., Banks, P. B., & Latham, A. D. M. (2019). Modeling habituation of introduced predators to unrewarding bird odors for conservation of ground-nesting shorebirds. *Ecological Applications*, 29(1), 1–14. doi:10.1002/eap.1814
- Li, Z., Kamnitsas, K., & Glocker, B. (2019). Overfitting of Neural Nets Under Class Imbalance:

- Analysis and Improvements for Segmentation. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 11766 LNCS, 402–410. doi:10.1007/978-3-030-32248-9_45
- Malhi, Y., Riutta, T., Wearn, O. R., Deere, N. J., Mitchell, S. L., Bernard, H., ... Struebig, M. J. (2022). Logged tropical forests have amplified and diverse ecosystem energetics. *Nature*, 612(December). doi:10.1038/s41586-022-05523-1
- Mclaughlin, Ó., Jonsson, T., & Emmerson, M. C. (2010). Temporal variability in predator-prey relationships of a forest floor food web. *Advances in Ecological Research*, 42, 171–264.
- Meijaard, E., & Sheil, D. (2008). The persistence and conservation of Borneo's mammals in lowland rain forests managed for timber: Observations, overviews and opportunities. *Ecological Research*, 23(1), 21–34. doi:10.1007/s11284-007-0342-7
- Mora, B. B., Gravel, D., Gilarranz, L. J., Poisot, T., & Stouffer, D. B. (2018). Identifying a common backbone of interactions underlying food webs from different ecosystems. *Nature Communications*, 9(1). doi:10.1038/s41467-018-05056-0
- Morales-Castilla, I., Matias, M. G., Gravel, D., & Araújo, M. B. (2015). Inferring biotic interactions from proxies. *Trends in Ecology and Evolution*, 30(6), 347–356. doi:10.1016/j.tree.2015.03.014
- Moreno-Torres, J. G., Raeder, T., Alaiz-Rodríguez, R., Chawla, N. V., & Herrera, F. (2012). A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1), 521–530. doi:10.1016/j.patcog.2011.06.019
- Mortensen, L. O., Chudzinska, M. E., Slabbekoorn, H., & Thomsen, F. (2021). Agent-based models to investigate sound impact on marine animals: bridging the gap between effects on individual behaviour and population level consequences. *Oikos*, 130(7), 1074–1086. doi:10.1111/oik.08078
- Mueller, T., Olson, K. A., Dressler, G., Leimgruber, P., Fuller, T. K., Nicolson, C., ... Fagan, W. F. (2011). How landscape dynamics link individual- to population-level movement patterns: A multispecies comparison of ungulate relocation data. *Global Ecology and Biogeography*, 20(5), 683–694. doi:10.1111/j.1466-8238.2010.00638.x
- Naseer, M., Ranasinghe, K., Khan, S., Hayat, M., Khan, F. S., & Yang, M. H. (2021). Intriguing Properties of Vision Transformers. *Advances in Neural Information Processing Systems*, 28(NeurIPS), 23296–23308.
- Niedballa, J., Wilting, A., Sollmann, R., Hofer, H., & Courtiol, A. (2019). Assessing analytical

- methods for detecting spatiotemporal interactions between species from camera trapping data. *Remote Sensing in Ecology and Conservation*, 1–14. doi:10.1002/rse2.107
- Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M., Packer, C., & Clune, J. (2017). Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning, *115*(25). doi:10.1073/pnas.1719367115
- Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., & Clune, J. (2018). Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences*, *115*(25), E5716–E5725. doi:10.1073/pnas.1719367115
- P. Vázquez, D., J. Melián, C., M. Williams, N., Blüthgen, N., R. Krasnov, B., & Poulin, R. (2007). Species abundance and asymmetric interaction strength in ecological networks. *Oikos*, *116*(7), 1120–1127. doi:10.1111/j.2007.0030-1299.15828.x
- Pacioni, C., Ramsey, D. S. L., Schumaker, N. H., Kreplins, T., & Kennedy, M. S. (2021). A novel modelling framework to explicitly simulate predator interaction with poison baits. *Wildlife Research*, *48*(1), 64–75. doi:10.1071/WR19193
- Palomares, F., & Caro, T. M. (1999). Interspecific killing among mammalian carnivores. *American Naturalist*, *153*(5), 492–508. doi:10.1086/303189
- Parsons, A. W., Kellner, K. F., Rota, C. T., Schuttler, S. G., Millsbaugh, J. J., & Kays, R. W. (2022). The effect of urbanization on spatiotemporal interactions between gray foxes and coyotes. *Ecosphere*, *13*(3), 1–13. doi:10.1002/ecs2.3993
- Parsons, A. W., Rota, C. T., Forrester, T., Baker-Whatton, M. C., McShea, W. J., Schuttler, S. G., ... Kays, R. (2019). Urbanization focuses carnivore activity in remaining natural habitats, increasing species interactions. *Journal of Applied Ecology*, *56*(8), 1894–1904. doi:10.1111/1365-2664.13385
- Paul, S., & Chen, P.-Y. (2022). Vision Transformers Are Robust Learners. *Proceedings of the AAAI Conference on Artificial Intelligence*, *36*(2), 2071–2081. doi:10.1609/aaai.v36i2.20103
- Pilosof, S., Porter, M. A., Pascual, M., & Kéfi, S. (2017). The multilayer nature of ecological networks. *Nature Ecology & Evolution*, *1*(4), 1–9.
- Poisot, T., Stouffer, D. B., & Gravel, D. (2015). Beyond species: Why ecological interaction networks vary through space and time. *Oikos*, *124*(3), 243–251. doi:10.1111/oik.01719
- Rahman, T., & Candolin, U. (2022). Linking animal behavior to ecosystem change in disturbed

- environments. *Frontiers in Ecology and Evolution*, 10(July), 1–13. doi:10.3389/fevo.2022.893453
- Ramesh, T., Kalle, R., & Downs, C. T. (2017). Staying safe from top predators: patterns of co-occurrence and inter-predator interactions. *Behavioral Ecology and Sociobiology*, 71(2), 1. doi:10.1007/s00265-017-2271-y
- Ripple, W. J., & Beschta, R. L. (2012). Large predators limit herbivore densities in northern forest ecosystems. *European Journal of Wildlife Research*, 58(4), 733–742. doi:10.1007/s10344-012-0623-5
- Ross, J., Hearn, A. J., Johnson, P. J., & Macdonald, D. W. (2013). Activity patterns and temporal avoidance by prey in response to Sunda clouded leopard predation risk. *Journal of Zoology*, 290(2), 96–106. doi:10.1111/jzo.12018
- Rowcliffe, J. M., Jansen, P. A., Kays, R., Kranstauber, B., & Carbone, C. (2016). Wildlife speed cameras: measuring animal travel speed and day range using camera traps. *Remote Sensing in Ecology and Conservation*, 2(2), 84–94. doi:10.1002/rse2.17
- Rustam, Yasuda, M., & Tsuyuki, S. (2012). Comparison of Mammalian Communities in a Human-Disturbed Tropical Landscape in East Kalimantan, Indonesia. *Mammal Study*, 37(4), 299–311. doi:10.3106/041.037.0404
- Rytkönen, S., Vesterinen, E. J., Westerduin, C., Leviäkangas, T., Vatka, E., Mutanen, M., ... Orell, M. (2019). From feces to data: A metabarcoding method for analyzing consumed and available prey in a bird-insect food web. *Ecology and Evolution*, 9(1), 631–639. doi:10.1002/ece3.4787
- Sandom, C., Dalby, L., Flojgaard, C., Kissling, W. D., Lenoir, J., Sandel, B., ... Svenning, J. C. (2013). Mammal predator and prey species richness are strongly linked at macroscales. *Ecology*, 94(5), 1112–1122. doi:10.1890/12-1342.1
- Santos, F., Carbone, C., Wearn, O. R., Rowcliffe, J. M., Espinosa, S., Moreira, M. G., ... Peres, C. A. (2019). Prey availability and temporal partitioning modulate felid coexistence in Neotropical forests. *PLoS ONE*, 14(3), 1–23. doi:10.1371/journal.pone.0213671
- Schneider, S., Greenberg, S., Taylor, G. W., & Kremer, S. C. (2020). Three critical factors affecting automated image species recognition performance for camera traps. *Ecology and Evolution*, 10(7), 3503–3517. doi:10.1002/ece3.6147
- Schneider, S., Taylor, G. W., & Kremer, S. (2018). Deep learning object detection methods for ecological camera trap data. *Proceedings - 2018 15th Conference on Computer and*

Robot Vision, CRV 2018, 321–328. doi:10.1109/CRV.2018.00052

- Simberloff, D., Martin, J. L., Genovesi, P., Maris, V., Wardle, D. A., Aronson, J., ... Vilà, M. (2013). Impacts of biological invasions: What's what and the way forward. *Trends in Ecology and Evolution*, 28(1), 58–66. doi:10.1016/j.tree.2012.07.013
- Simonyan, K., & Zisserman, A. (2015). Very deep convolutional networks for large-scale image recognition. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, 1–14.
- Sinclair, A. R. E., Mduma, S., & Brashares, J. S. (2003). Patterns of predation in a diverse predator-prey system. *Nature*, 425(6955), 288–290. doi:10.1038/nature01934
- Smith, J. A., Suraci, J. P., Hunter, J. S., Gaynor, K. M., Keller, C. B., Palmer, M. S., ... Beaudrot, L. (2020). Zooming in on mechanistic predator–prey ecology: Integrating camera traps with experimental methods to reveal the drivers of ecological interactions. *Journal of Animal Ecology*, 89(9), 1997–2012. doi:10.1111/1365-2656.13264
- Spatz, D. R., Zilliacus, K. M., Holmes, N. D., Butchart, S. H. M., Genovesi, P., Ceballos, G., ... Croll, D. A. (2017). Globally threatened vertebrates on islands with invasive species. *Science Advances*, 3(10). doi:10.1126/sciadv.1603080
- Stabach, J. A., Cunningham, S. A., Connette, G., Mota, J. L., Reed, D., Byron, M., ... Leimgruber, P. (2020). Short-term effects of GPS collars on the activity, behavior, and adrenal response of scimitar-horned oryx (*Oryx dammah*). *PLoS ONE*, 15(2), 1–22. doi:10.1371/journal.pone.0221843
- Stankowich, T., & Coss, R. G. (2006). Effects of predator behavior and proximity on risk assessment by Columbian black-tailed deer. *Behavioral Ecology*, 17(2), 246–254. doi:10.1093/beheco/arj020
- Steenweg, R., Hebblewhite, M., Kays, R., Ahumada, J., Fisher, J. T., Burton, C., ... Rich, L. N. (2017). Scaling-up camera traps: monitoring the planet's biodiversity with networks of remote sensors. *Frontiers in Ecology and the Environment*, 15(1), 26–34. doi:10.1002/fee.1448
- Stowell, D., Wood, M. D., Pamuła, H., Stylianou, Y., & Glotin, H. (2019). Automatic acoustic detection of birds through deep learning: The first Bird Audio Detection challenge. *Methods in Ecology and Evolution*, 10(3), 368–380. doi:10.1111/2041-210X.13103
- Swanson, A., Kosmala, M., Lintott, C., Simpson, R., Smith, A., Packer, C., ... Park, N. (2015). Snapshot Serengeti, high-frequency annotated camera trap images of 40 mammalian

- species in an African savanna. *Scientific Data*, 2. doi:10.1038/sdata.2015.26
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the Inception Architecture for Computer Vision. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2016-Decem*, 2818–2826. doi:10.1109/CVPR.2016.308
- Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., Sweeney, S. J., Vercauteren, K. C., Snow, N. P., ... Miller, R. S. (2019). Machine learning to classify animal species in camera trap images: Applications in ecology. *Methods in Ecology and Evolution*, 10(4), 585–590. doi:10.1111/2041-210X.13120
- Terry, J. C. D., Roy, H. E., & August, T. A. (2020). Thinking like a naturalist: Enhancing computer vision of citizen science images by harnessing contextual data. *Methods in Ecology and Evolution*, 11(2), 303–315. doi:10.1111/2041-210X.13335
- Tucker, M. A., Busana, M., Huijbregts, M. A. J., & Ford, A. T. (2021). Human-induced reduction in mammalian movements impacts seed dispersal in the tropics. *Ecography*, 44(6), 897–906. doi:10.1111/ecog.05210
- Tylianakis, J. M., Tscharntke, T., & Lewis, O. T. (2007). Habitat modification alters the structure of tropical host-parasitoid food webs. *Nature*, 445(7124), 202–205. doi:10.1038/nature05429
- Vanak, A. T., Fortin, D., Thaker, M., Ogden, M., Owen, C., Greatwood, S., & Slotow, R. (2013). Moving to stay in place: behavioral mechanisms for coexistence of African large carnivores. *Ecology*, 94(11), 2619–2631.
- Wäldchen, J., & Mäder, P. (2018). Machine learning for image based species identification. *Methods in Ecology and Evolution*, 9(11), 2216–2225. doi:10.1111/2041-210X.13075
- Wearn, O. R., Bell, T. E. M., Bolitho, A., Durrant, J., Haysom, J. K., Nijhawan, S., ... Rowcliffe, J. M. (2022). Estimating animal density for a community of species using information obtained only from camera-traps. *Methods in Ecology and Evolution*, 2022(May), 1–14. doi:10.1111/2041-210X.13930
- Wearn, O. R., Carbone, C., Rowcliffe, J. M., Bernard, H., & Ewers, R. M. (2016). Grain-dependent responses of mammalian diversity to land use and the implications for conservation set-aside. *Ecological Applications*, 26(5), 1409–1420. doi:10.1890/15-1363
- Wearn, O. R., Carbone, C., Rowcliffe, J. M., Pfeifer, M., Bernard, H., & Ewers, R. M. (2019). Land-use change alters the mechanisms assembling rainforest mammal communities in

- Borneo. *Journal of Animal Ecology*, 88(1), 125–137. doi:10.1111/1365-2656.12903
- Wearn, O. R., Freeman, R., & Jacoby, D. M. P. (2019). Responsible AI for conservation. *Nature Machine Intelligence*, 1(2), 72–73. doi:10.1038/s42256-019-0022-7
- Wearn, O. R., Rowcliffe, J. M., Carbone, C., Bernard, H., & Ewers, R. M. (2013). Assessing the status of wild felids in a highly-disturbed commercial forest reserve in Borneo and the implications for camera trap survey design. *PLoS ONE*, 8(11). doi:10.1371/journal.pone.0077598
- Wearn, O. R., Rowcliffe, J. M., Carbone, C., Pfeifer, M., Bernard, H., & Ewers, R. M. (2017). Mammalian species abundance across a gradient of tropical land-use intensity: A hierarchical multi-species modelling approach. *Biological Conservation*, 212(October 2016), 162–171. doi:10.1016/j.biocon.2017.05.007
- Wei Koh, P., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., ... Liang, P. (2021). WILDS: a benchmark of in-the-wild distribution shifts. *Proceedings - International Conference on Machine Learning*, PMLR 139, 5637–5664.
- Weinstein, B. G. (2018). A computer vision for animal ecology. *Journal of Animal Ecology*, 87(3), 533–545. doi:10.1111/1365-2656.12780
- Wheatley, R., Pavlic, T. P., Levy, O., & Wilson, R. S. (2020). Habitat features and performance interact to determine the outcomes of terrestrial predator–prey pursuits. *Journal of Animal Ecology*, 89(12), 2958–2971. doi:10.1111/1365-2656.13353
- Wiens, J. J. (2016). Climate-Related Local Extinctions Are Already Widespread among Plant and Animal Species. *PLoS Biology*, 14(12), 1–18. doi:10.1371/journal.pbio.2001104
- Willi, M., Pitman, R. T., Cardoso, A. W., Locke, C., Swanson, A., Boyer, A., ... Fortson, L. (2019). Identifying animal species in camera trap images using deep learning and citizen science. *Methods in Ecology and Evolution*, 10(1), 80–91. doi:10.1111/2041-210X.13099
- WWF. (2022). *Living Planet Report 2022 – Building a nature-positive society*. Almond, R.E.A., Grooten, M., Juffe Bignoli, D. & Petersen, T. (Eds). WWF. Gland, Switzerland. Retrieved from www.livingplanetindex.org
- Zalewska, K., Wagnershauser, C. N., Kortland, K., & Lambin, X. (2021). The best defence is not being there: avoidance of larger carnivores is not driven by risk intensity. *Journal of Zoology*, 315(2), 110–122. doi:10.1111/jzo.12910

Appendices

6 Appendix 1: Chapter 2: Supplementary Information

6.1 Methods

6.1.1 Data

Table 6.1: Definitions of disturbance levels (Wearn et al., 2017).

Disturbance level	Definition
Undisturbed forest	Dominated by old-growth dipterocarps. High, continuous canopy with sparsely-vegetated understorey. Unlogged, with little recent disturbance evident.
Disturbed forest	Mostly pioneer tree species (typically <i>Macaranga</i> species), but some old-growth dipterocarp species may be present. Discontinuous canopy. Lower intensity of logging or natural disturbance.
Heavily disturbed forest	High scrub or dense understorey layer (typically with vines and <i>Dinochloa</i> climbing bamboo species), with a low, heavily-broken canopy layer (< 20 m). Possibly some large, isolated trees (> 20 m). Intensively-logged area or large gap disturbance.
Herbaceous scrub	Dominated by herbs (typically <i>Zingiberaceae</i>), vines and shrubs, with no trees > 3 m in height (except oil palm <i>Elaeis guineensis</i>). Typically representing secondary re-growth from clear-felling, or large gaps due to landslides.
Open area	Open area. Dominated by grasses and small shrubs (< 1 m in height). Typically on logging roads or old log landing areas.

Table 6.2: Species included within class groups; note all other classes are single species.

Group	Species
Banded Palm Civet	Banded palm civet, <i>Hemigalus derbyanus</i> ; Hose's palm civet, <i>Diplogale hosei</i>
Beautiful Squirrels	Ear-spot squirrel, <i>Callosciurus adamsi</i> ; Plantain squirrel, <i>Callosciurus notatus</i> ; Prevost's squirrel, <i>Callosciurus prevostii</i>
Plain Treeshrew	Pen-tailed treeshrew, <i>Ptilocercus lowii</i> ; Long-footed treeshrew, <i>Tupaia longipes</i>
Crane	Red-legged crane, <i>Rallina fasciata</i> ; White-breasted waterhen, <i>Amauornis phoenicurus</i> ; Great egret, <i>Ardea alba</i>
Cuckoo	Greater coucal, <i>Centropus sinensis</i> ; Bornean ground cuckoo, <i>Carpococcyx radiceus</i> ; Crested serpent eagle, <i>Spilornis cheela</i>
Dove	Emerald dove, <i>Chalcophaps indica</i> ; Spotted dove, <i>Spilopelia chinensis</i>
Landfowl	Crested fireback, <i>Lophura ignita</i> ; Great argus, <i>Argusianus argus</i> ; Bulwer's pheasant, <i>Lophura bulweri</i> ; Chestnut-necklaced partridge, <i>Arborophila charltonii</i> ; Blue-breasted quail, <i>Excalfactoria chinensis</i> ; Red junglefowl, <i>Gallus gallus</i> ; Crested partridge, <i>Rollulus rouloul</i>
Langur	Maroon langur, <i>Presbytis rubicunda</i> ; Hose's langur, <i>Presbytis hosei</i>
Low's Squirrel	Low's squirrel, <i>Sundasciurus lowii</i> ; Slender squirrel, <i>Sundasciurus tenuis</i>
Mueller's Giant Sunda Rat	Mueller's giant Sunda rat, <i>Sundamys muelleri</i> ; Small spiny rat, <i>Maxomys baeodon</i> ; Black rat, <i>Rattus rattus</i> ; Whitehead's spiny rat, <i>Maxomys whiteheadi</i>
Songbird	Black-capped babbler, <i>Pellorneum capistratum</i> ; Little spiderhunter, <i>Arachnothera longirostra</i> ; White-crowned shama, <i>Copsychus stricklandii</i> ; Short-tailed babbler, <i>Malacocincla malaccensis</i> ; White-crowned forktail, <i>Enicurus leschenaultia</i> ; Oriental magpie-robin, <i>Copsychus saularis</i> ; Yellow-vented bulbul, <i>Pycnonotus goiavier</i> ; Chestnut munia, <i>Lonchura atricapilla</i> ; Moustached babbler, <i>Malacopteron magnirostre</i> ; Bornean ground-babbler, <i>Ptilocichla leucogrammica</i> ; Giant pitta, <i>Hydrornis caeruleus</i> ; Blue-headed pitta, <i>Hydrornis baudii</i> ; Banded pitta, <i>Hydrornis irena</i> ; Black-and-crimson pitta, <i>Pitta ussheri</i> ; Hooded pitta, <i>Pitta sordida</i>
Striped Ground Squirrel	Four-striped ground squirrel, <i>Lariscus hosei</i> ; Least pygmy squirrel, <i>Exilisciurus exilis</i> ; Tufted ground squirrel, <i>Rheithrosciurus macrotis</i> ; Flying squirrel, <i>Aeromys thomasi</i>

Western tarsier	Western tarsier, <i>Cephalopachus bancanus</i> ; Bornean gibbon, <i>Hylobates muelleri</i>
Yellow-throated Marten	Yellow-throated marten, <i>Martes flavigula</i> ; Malayan weasel, <i>Mustela nudipes</i>

6.1.2 Machine learning

Networks were constructed using the Keras library for Python and were trained on Imperial College London's HPC. Each run was performed once due to computational time.

Table 6.3: Detail of how images for each class were allocated to training, validation and test sets for the network and dataset comparison analyses.

No. of images (N)	No. of events (n)	Training	Validation (V)	Test (T)
N>5000	N/A	Max 90%	Max 5%	Max 5%
N<5000	n≥6	N-V-T	Max (5%, 2 events)	Max (5%, 2 events)
	4≤n≤5	N-T	None	2 events

Note: the minimum number of events at 6 was chosen to ensure at least two events in each of the three sets; and the lower minimum at 4 selected to ensure inclusion of as many species as possible.

Table 6.4: Number of images available and allocated per class to training, validation and test for the baseline dataset.

Class	Total images	Training	Validation	Test
Red muntjac, <i>Muntiacus muntjak</i>	94,927	4,500	250	246
Bearded pig, <i>Sus barbatus</i>	78,888	4,491	243	250
Human, <i>Homo sapiens</i>	49,852	4,497	243	244
Sambar deer, <i>Rusa unicolor</i>	44,165	4,491	247	250
Landfowl [group]	22,077	4,499	241	241
Pig-tailed macaque, <i>Macaca nemestrina</i>	21,378	4,492	250	242
Yellow muntjac, <i>Muntiacus atherodes</i>	14,682	4,498	246	246
Greater mouse-deer, <i>Tragulus napu</i>	14,637	4,500	242	244
Malayan porcupine, <i>Hystrix brachyura</i>	9,270	4,500	246	250
Malayan civet, <i>Viverra zangalunga</i>	8,912	4,500	249	249
Lesser mouse-deer, <i>Tragulus kanchil</i>	7,016	4,492	241	249
Banded palm civet [group]	4,673	4,212	229	232
Long-tailed porcupine, <i>Trichys fasciculata</i>	3,209	2,897	157	155
Sun bear, <i>Helarctos malayanus</i>	2,005	1,814	99	92
Leopard cat, <i>Prionailurus bengalensis</i>	1,974	1,785	95	94
Domestic dog, <i>Canis familiaris</i>	1,694	1,532	85	77
Songbird [group]	1,427	1,287	71	69
Long-tailed giant rat, <i>Leopoldamys sabanus</i>	1,287	1,165	63	59
Dove [group]	1,218	1,101	60	57
Banteng, <i>Bos javanicus</i>	1,164	1,055	54	55
Orangutan, <i>Pongo pygmaeus</i>	1,028	927	50	51
Large treeshrew, <i>Tupaia tana</i>	981	893	42	46
Low's squirrel [group]	934	845	46	43
Plain treeshrew [group]	921	831	45	45

Masked palm civet, <i>Paguma larvata</i>	822	751	33	38
Short-tailed mongoose, <i>Herpestes brachyurus</i>	701	644	31	26
Thick-spined porcupine, <i>Hystrix crassispinis</i>	663	603	27	33
Crane [group]	552	504	25	23
Yellow-throated marten [group]	552	502	28	22
Striped ground squirrel [group]	538	486	27	25
Pangolin, <i>Manis javanica</i>	451	416	19	16
Slender treeshrew, <i>Tupaia gracilis</i>	358	326	19	13
Stink badger, <i>Mydaus javanensis</i>	323	289	14	20
Sunda clouded leopard, <i>Neofelis diardi</i>	318	278	20	20
Striped treeshrew, <i>Tupaia dorsalis</i>	302	262	20	20
Marbled cat, <i>Pardofelis marmorata</i>	272	236	16	20
Horse-tailed squirrel, <i>Sundasciurus hippurus</i>	244	216	13	15
Cuckoo [group]	229	202	15	12
Mueller's giant Sunda rat [group]	198	167	11	20
Borneo elephant, <i>Elephas maximus</i>	194	154	20	20
Binturong, <i>Arctictis binturong</i>	188	158	10	20
Collared mongoose, <i>Herpestes semitorquatus</i>	184	156	12	16
Bay cat, <i>Catopuma badia</i>	168	135	18	15
Common palm civet, <i>Paradoxurus hermaphroditus</i>	99	59	20	20
Langur [group]	98	66	12	20
Long-tailed macaque, <i>Macaca fascicularis</i>	96	65	13	18
Moonrat, <i>Echinosorex gymnura</i>	83	59	12	12
Beautiful squirrels [group]	75	43	12	20
Oriental small-clawed otter, <i>Aonyx cinereus</i>	59	27	20	12
Western tarsier [group]	43	12	20	11

Banded linsang, Prionodon linsang	42	17	9	16
-----------------------------------	----	----	---	----

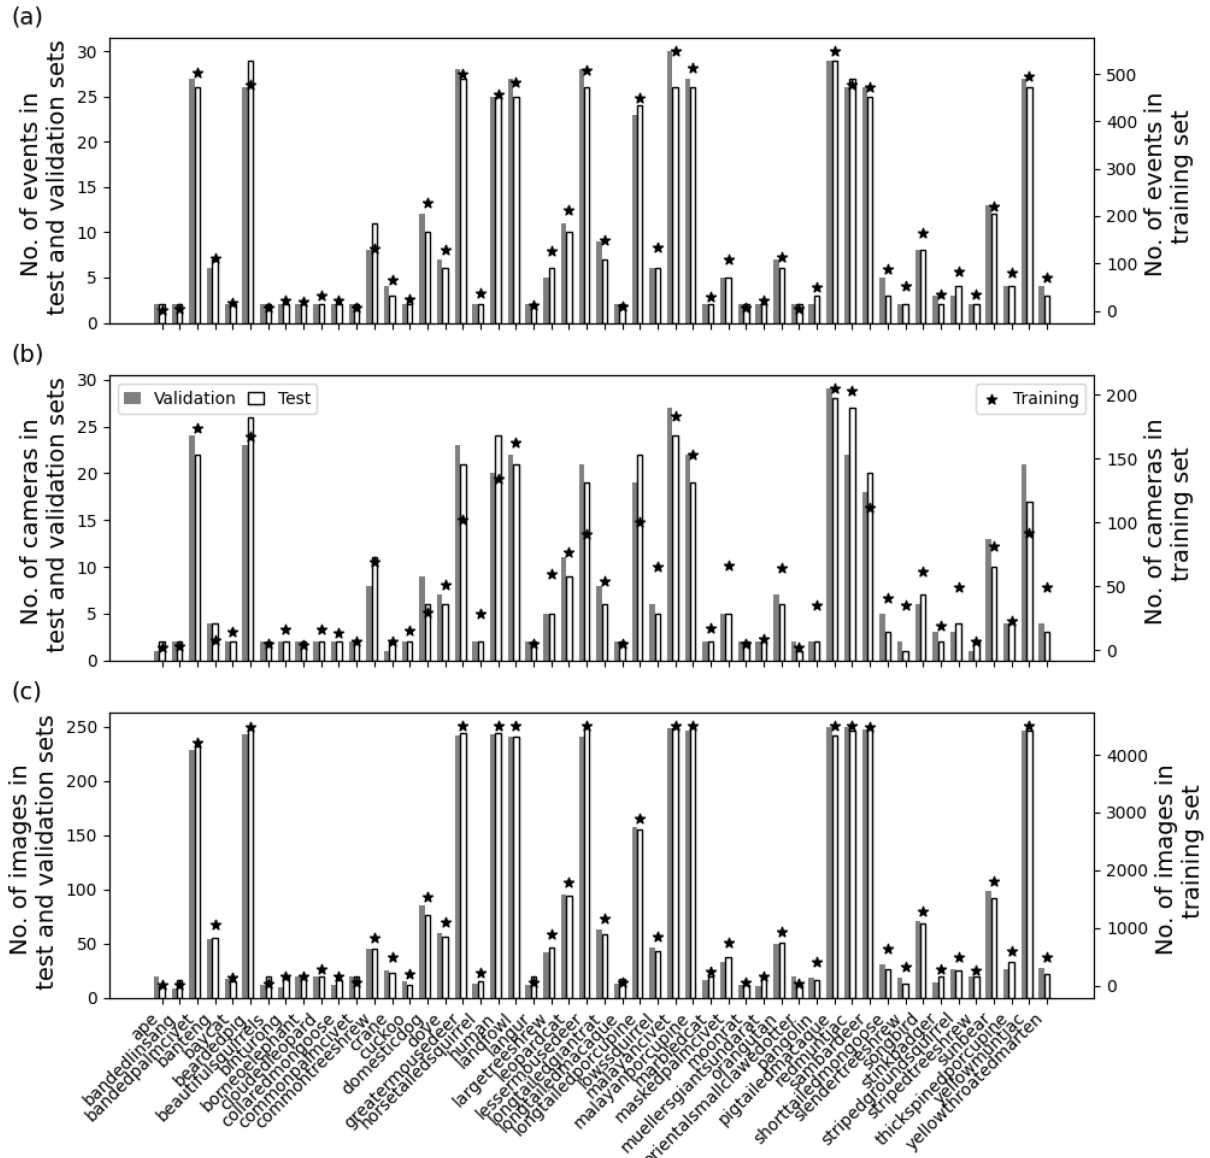


Figure 6.1: Number of (a) events, (b) cameras and (c) images in the training, validation, and test sets for the baseline dataset.

Table 6.5: Data augmentation applied during training.

Data augmentation	Specification
Random shearing	0.2 radians
Flip	Horizontal
Crop size	244 x 244
Brightness modification	+/-20%

Table 6.6: Baseline hyperparameter settings for all networks.

Hyperparameter	Setting
Batch size	128
No. epochs	40
Optimiser	SGD with momentum (0.9)
Learning rate	Dynamic (see Appendix 1, Table 6.8)
Transfer learning	Weights initialised
Weight decay	L2(0.005/2)
Dropout probability	0.5

Table 6.7: Hyperparameter settings varied during optimisation; note LR: learning rate.

Run name	Hyperparameter	Parameter value/setting
dropout prob: 0.75	dropout probability	0.75
dropout prob: 0.25	dropout probability	0.25
frozen conv layers	fully connected layers trained only	-
optimiser: RMSprop	optimising algorithm	RMS prop (LR in Table 6.8)
optimiser: Adam	optimising algorithm	Adam (LR in Table 6.8)
batch64	batch size	64

Table 6.8: Final learning rates used with optimisation algorithms.

Epoch	SGD with momentum (baseline)	RMSprop/Adam
1-5	0.005	0.0001
6-10	0.001	0.0001
11-20	0.005	0.00005
21-30	0.001	0.00001
31-35	0.0005	0.000005
36-40	0.0001	0.000001

6.1.3 Generalisability – individual disturbance level comparison




Figure 6.2: Training image distribution. Shown across classes per disturbance level from undisturbed forest (left) to open area (right) for the dataset split at event. level

Table 6.9: Mean detection rate per camera for each species in each disturbance level for the dataset split at camera level.

	Undisturbed forest		Disturbed forest		Heavily disturbed forest		Herbaceous scrub		Open area	
	Train	Test	Train	Test	Train	Test	Train	Test	Train	Test
Bearded pig	195.65	166.67	300.00	132.67	39.82	35.71	91.84	55.56	77.59	55.56
Human	166.67	125.00	109.76	71.43	84.91	54.00	116.52	10.00	112.50	71.43
Landfowl	141.67	33.50	87.36	117.67	52.94	50.60	90.35	26.75	55.34	50.00
Leopard cat	24.55	28.00	33.27	10.25	14.73	11.00	7.50	8.00	23.63	37.67
Lesser mousedeer	58.50	21.00	41.20	23.50	41.77	30.00	50.00	135.00	94.94	37.00
Malayan civet	65.54	66.67	42.96	37.00	43.79	74.00	25.17	31.00	16.48	23.50
Malayan porcupine	47.45	17.67	79.56	24.00	41.71	26.70	77.36	118.00	31.37	24.63
Masked palm civet	5.00	10.00	14.00	10.00	8.61	21.00	8.17	10.00	15.68	10.00
Orangutan	12.33	1.00	9.00	4.00	14.14	21.57	16.64	10.00	23.75	6.00
Pangolin	13.50	10.00	13.00	10.00	13.27	7.67	7.75	11.50	16.67	9.00
Pig-tailed macaque	126.88	61.00	47.75	1.00	40.54	35.71	70.34	81.50	50.00	41.67
Red muntjac	173.08	47.50	115.38	68.80	32.85	38.46	72.58	62.50	44.12	38.46
Sambar deer	225.00	12.00	321.43	500.00	54.88	62.50	146.58	1.00	38.22	19.00
Sun bear	14.00	10.00	27.50	10.00	26.54	26.50	24.94	18.50	20.57	15.67

6.1.4 Bounding boxes

Table 6.10: Number of images per species class used in the combined disturbance level datasets with and without bounding boxes after passing images through the MegaDetectorV3.

	Total no. images with bounding boxes	Total no. images without bounding boxes	% images without bounding boxes
Bearded pig	31,980	1,306	4%
Human	24,312	4,515	16%
Landfowl	17,006	388	2%
Leopard cat	1,803	81	4%
Lesser mouse-deer	6,105	209	3%
Long-tailed giant rat	1,017	216	18%
Malayan civet	8,053	370	4%
Malayan porcupine	8,079	565	7%
Masked palm civet	716	54	7%
Orangutan	672	13	2%
Pangolin	350	76	18%
Pig-tailed macaque	16,824	756	4%
Red muntjac	35,579	706	2%
Sambar deer	25,888	1,018	4%
Sun bear	1,759	68	4%
Total/total/mean	180,143	10,341	5%

6.2 Results

6.2.1 Network comparison

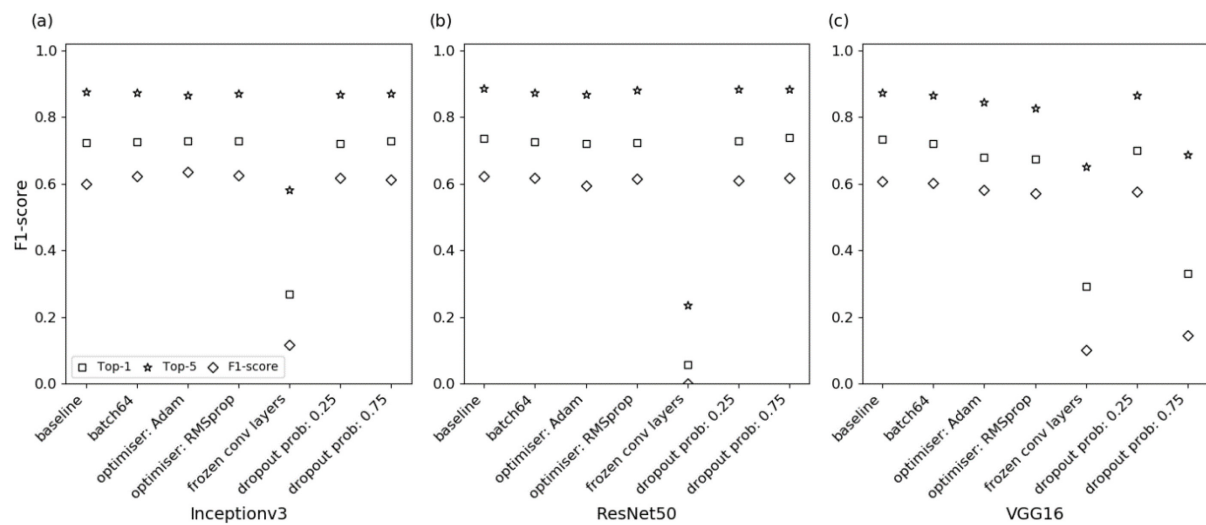


Figure 6.3: Top-1 accuracy, Top-5 accuracy and F1-score for settings combinations tried for (a) Inceptionv3, (b) ResNet50 and (c) VGG16.

Figure 6.3 shows that networks consistently achieve higher Top-5 accuracy than Top-1, which is to be expected since the network has more attempts to make the correct prediction. These metrics are also consistently higher than mean F1-score. Top-1 and Top-5 accuracy are biased towards large classes, whereas F1-score is a better representation of performance across all the classes as they have equal weighting. The final chosen network is Inceptionv3 with Adam optimiser.

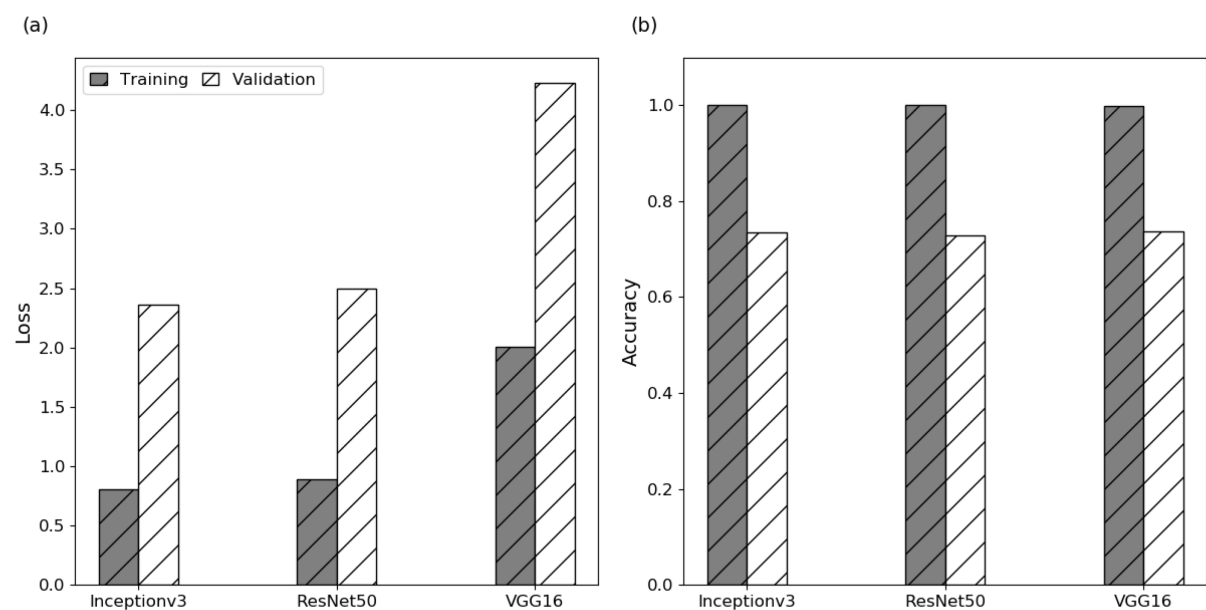


Figure 6.4: Training and validation (a) loss and (b) accuracy after training for 40 epochs for the best setup of each network.

Figure 6.4 illustrates the overfitting present across all three networks. This is a known issue, particularly for imbalanced datasets (Li, Kamnitsas, & Glocker, 2019). More dropout probability values were tried for the Inceptionv3 network to try to improve this but to little effect (Table 6.11).

Table 6.11: Training and validation metrics for dropout analysis.

Dropout probability	Validation loss	Validation accuracy	Training loss	Training accuracy
0.1	2.189	0.733	0.734	1.000
0.2	2.243	0.723	0.760	1.000
0.3	2.321	0.723	0.755	1.000
0.4	2.368	0.723	0.793	1.000
0.5	2.399	0.721	0.805	1.000
0.6	2.470	0.731	0.827	1.000
0.7	2.496	0.731	0.866	1.000
0.8	2.606	0.727	0.936	1.000
0.9	2.794	0.733	1.086	1.000




Figure 6.5: (a) F1-score and (b) recall plotted against number of training images per class in the baseline dataset.

Figure 6.5 illustrates the class imbalance within our training dataset, and the variability in F1-score and recall for rarer classes while more common classes perform consistently above 0.5.




Figure 6.6: Generalisability: mean F1-score across the baseline, common only and increased cap datasets. Mean F1-score for the network trained on undisturbed forest (bottom) through to open area (top) and tested on undisturbed forest through to open area (left-right) using the event-level dataset. Performance when tested on common species only (classes with >1,000 images available) is shown by bars, for the network trained on the baseline dataset (max. 5,000 images per class, black), common species only dataset (white) and increased cap dataset (max. 10,000 images per class, grey). Performance when tested on the baseline test dataset is denoted by x.

6.2.2 Generalisability – combinations of disturbance levels

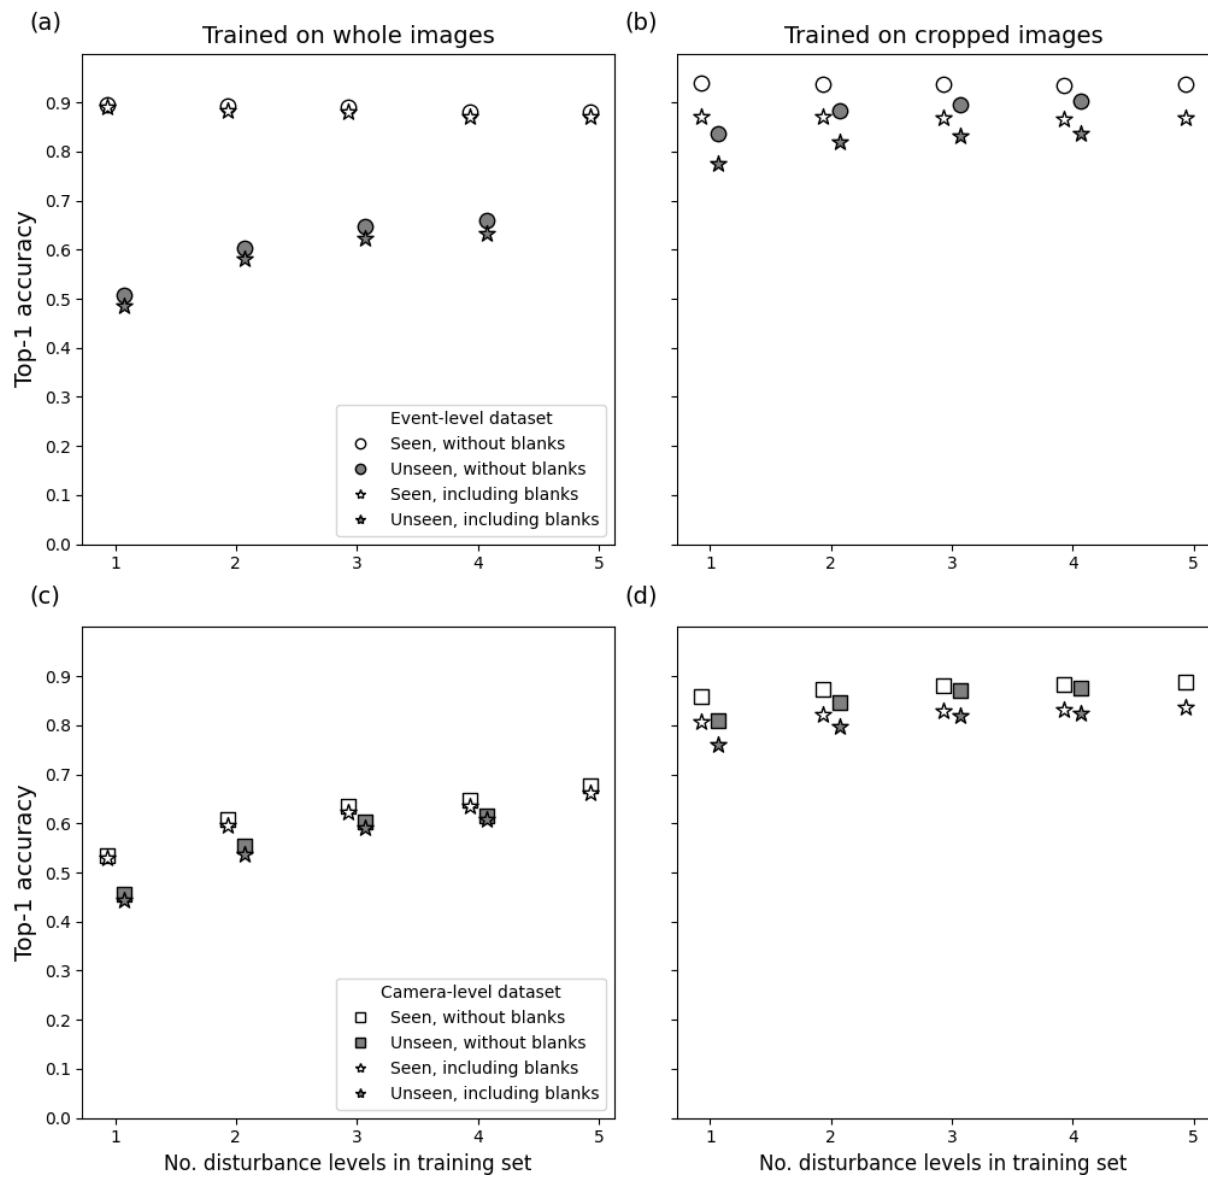


Figure 6.7: Generalisability: Top-1 accuracy. (a) and (b): Mean Top-1 accuracy when trained on combined disturbed levels taken from the event-level dataset with and without the 'blank' images in which the MegaDetector was unable to detect an animal, for the network trained and tested on (a) whole images and (b) cropped images. (c) and (d): as for (a) and (b) but using the camera-level dataset.

7 Appendix 2: Chapter 3 Supplementary Information

7.1 Intervals capped at 1 week

Table 7.1: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each land-use category, with a maximum cap of 7 days on intervals included in the analysis.

	Bearded pig – human Median (N)	Human – bearded pig Median (N)
Primary forest	4.2 days (10)	1.4 days (6)
Logged forest	0.8 days (103)	1.1 days (88)
Oil palm	0.9 days (11)	1.1 days (10)




Figure 7.1: Results for the land-use mixed-effect model with a cap of 7 days imposed on intervals included. Model coefficients (left) and predicted interval durations (right).

Table 7.2: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each habitat disturbance level, with a maximum cap of 7 days on intervals included in the analysis.

	Bearded pig – human Median (N)	Human – bearded pig Median (N)
Undisturbed forest	5.8 days (7)	2.5 days (4)
Disturbed forest	1.8 days (8)	4.4 days (5)
Heavily-disturbed forest	1.0 days (36)	1.4 days (30)
Herbaceous scrub	0.6 days (20)	0.6 days (19)
Open area	0.8 days (47)	0.8 days (39)

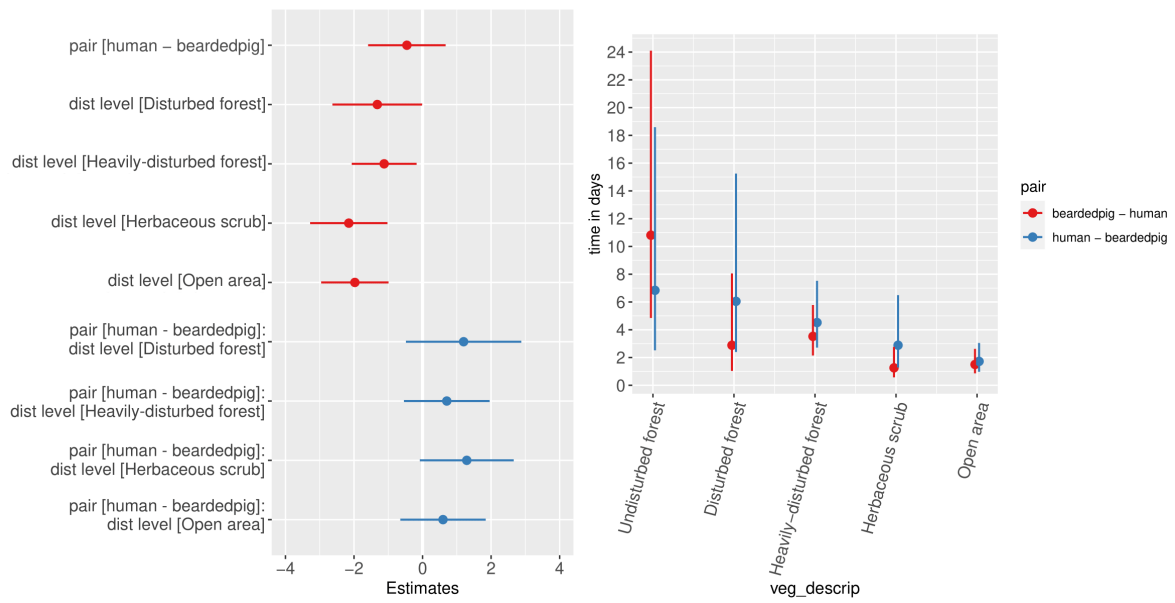


Figure 7.2: Results for the habitat disturbance level mixed-effect model with a cap of 7 days imposed on intervals included. Model coefficients (left) and predicted interval durations (right).

7.2 Intervals capped at 2 days

Table 7.3 Median interval duration in days and number of intervals (N) recorded for each pair orientation in each land-use category, with a maximum cap of 2 days on intervals included in the analysis.

	Bearded pig – human	Human – bearded pig
	Median (N)	Median (N)
Primary forest	1.1 days (2)	0.9 days (4)
Logged forest	0.6 days (74)	0.5 days (56)
Oil palm	0.5 days (8)	0.5 days (6)




Figure 7.3: Results for the land-use mixed-effect model with a cap of 2 days imposed on intervals included. Model coefficients (left) and predicted interval durations (right).

Table 7.4: Median interval duration in days and number of intervals (N) recorded for each pair orientation in each habitat disturbance level, with a maximum cap of 2 days on intervals included in the analysis.

	Bearded pig – human	Human – bearded pig
	Median (N)	Median (N)
Undisturbed forest	1.1 days (1)	0.6 days (2)
Disturbed forest	0.8 days (4)	1.0 days (2)
Heavily-disturbed forest	0.7 days (21)	0.3 days (17)
Herbaceous scrub	0.6 days (20)	0.5 days (15)
Open area	0.6 days (35)	0.6 days (30)




Figure 7.4: Results for the habitat disturbance level mixed-effect model with a cap of 2 days imposed on intervals included. Model coefficients (left) and predicted interval durations (right).

8 Appendix 3: Chapter 4 Supplementary Information

8.1 Influence of interaction zone radius

8.1.1 Simulated interaction events

Table 8.1: Mean number of avoidance events in Scenario 2, and attraction events in Scenario 3, per simulation for 6, 12 and 24 agents per species at each detection zone radius.

Interaction zone radius (m)	6 agents per species	12 agents per species	24 agents per species
Scenario 2: mean no. avoidance events per simulation			
15	10	42	169
25	21	86	332
50	52	215	868
75	89	362	1482
100	135	516	2131
150	238	911	3894
Scenario 3: mean no. attraction events per simulation			
15	269	785	3369
25	574	2198	8413
50	1468	6652	24404
75	2657	11287	44399
100	4177	16364	65426
150	6542	27963	106670

8.1.2 Camera trap events




Figure 8.1: Total number of camera trap capture events per 100 simulations of the CRW and avoidance scenarios. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).




Figure 8.2: Total number of camera trap capture events per 100 simulations of the CRW and attraction scenarios. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).

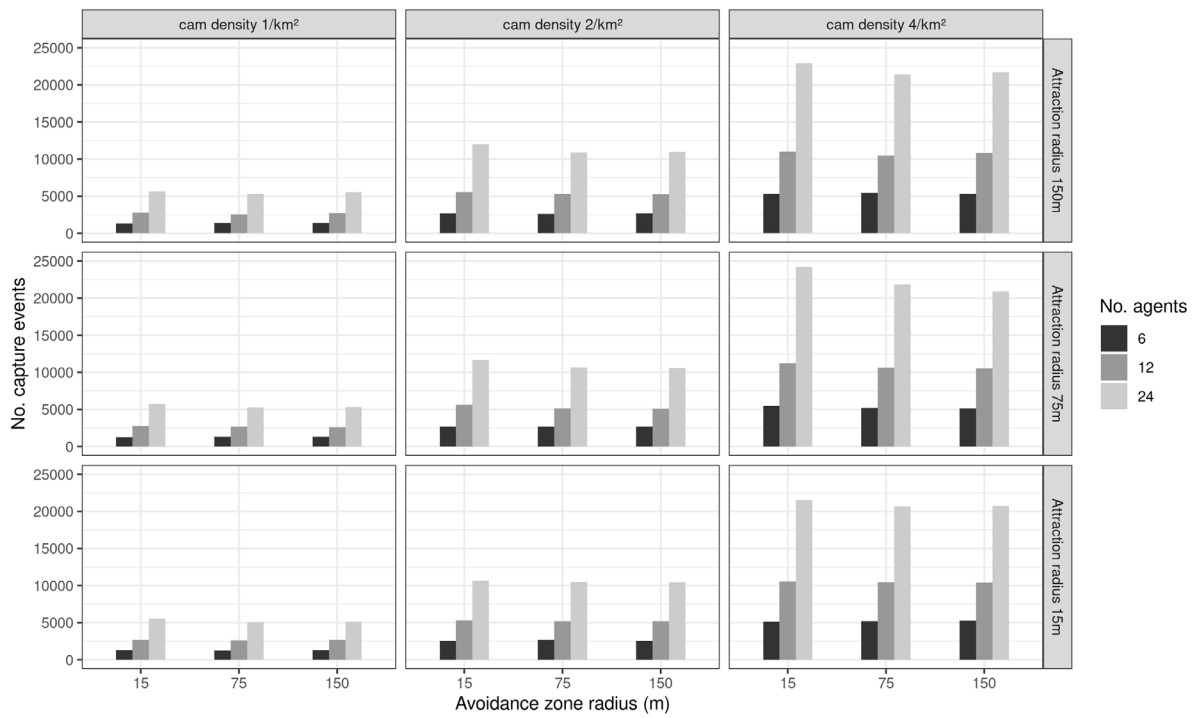


Figure 8.3: Total number of camera trap events per 100 simulations of Scenario 4: combined avoidance and attraction. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).

8.1.3 Intervals extracted

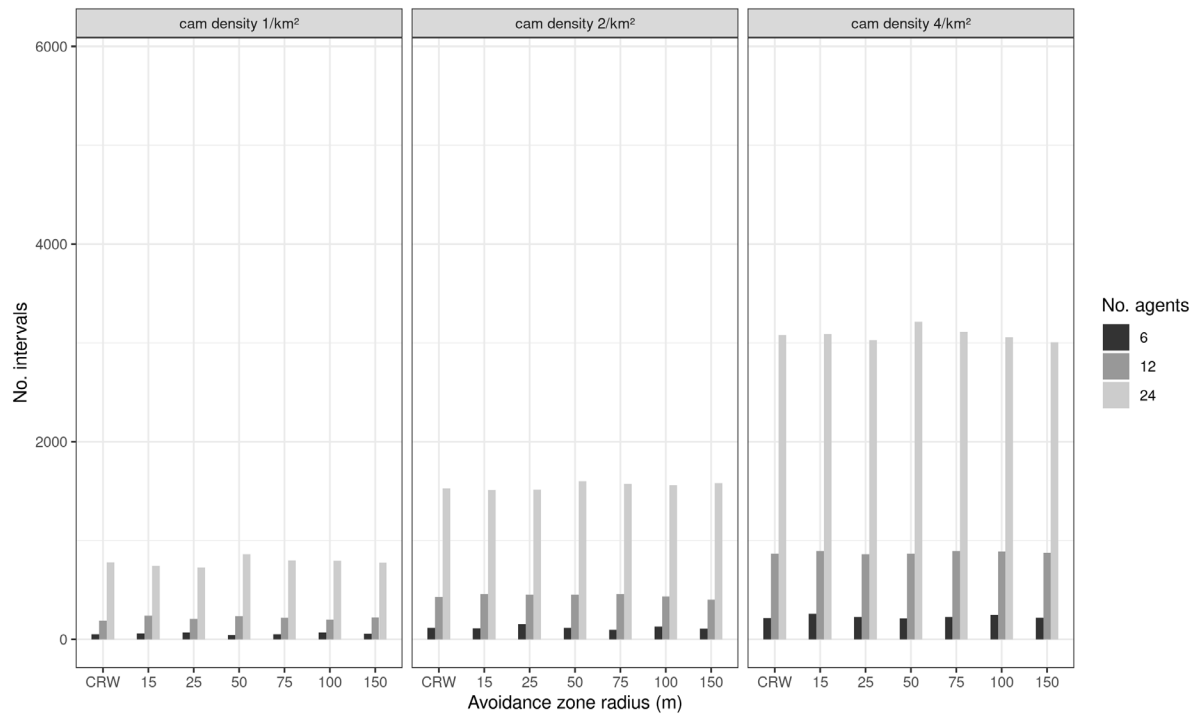


Figure 8.4: Total number intervals extracted from 100 simulations of the CRW and avoidance scenarios. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).




Figure 8.5: Total number of intervals extracted per 100 simulations of the CRW and attraction scenarios. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).

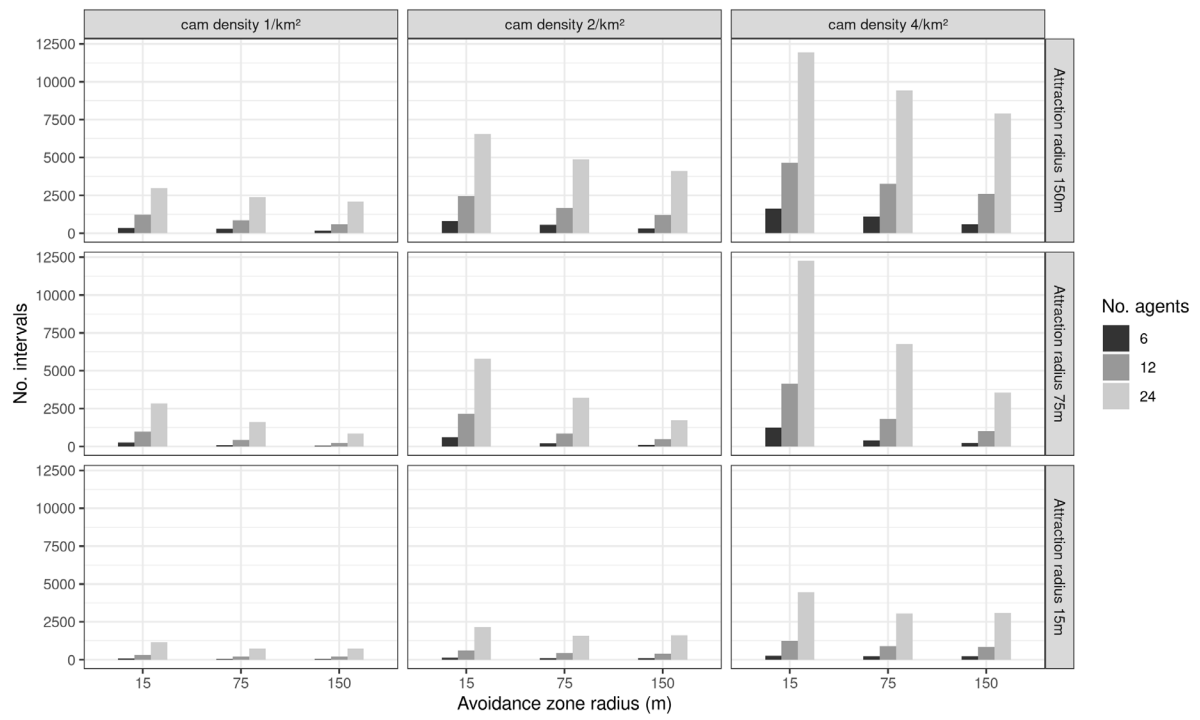


Figure 8.6: Total number of intervals extracted from 100 simulations of Scenario 4: combined avoidance and attraction. Shown for a camera trap density of 1/km² (left), 2/km² (middle) and 4/km² (right).

8.2 Applying the Niedballa et al. (2019) threshold




Figure 8.7: Model coefficient estimates for Scenario 2: one-way avoidance with setups that did not meet the Niedballa et al. threshold of 50 records per pair excluded. Note that the intercept (not shown) contains AB intervals for a CRW with 6 agents per species and a camera density of $1/\text{km}^2$.




Figure 8.8: Model coefficient estimates for Scenario 3: one-way attraction with setups that did not meet the Niedballa et al. threshold of 50 records per pair excluded. Note that the intercept (not shown) contains AB intervals for a CRW with 6 agents per species and a camera density of 1/km².

8.3 ABM code

```
#!/usr/bin/env python3
import random
import numpy as np
import math
import sys
import timeit
import csv
import pickle
import bz2
```

```
## Start timer
```

```

tic=timeit.default_timer()

##### Retrieve settings from bash inputs #####
run_cat = sys.argv[1] # 'CRW', 'avoid1', 'avoid2', 'attract', 'pred_pre'
outputpath = sys.argv[2]
avoid_strength = int(sys.argv[3])
avoid_ddist = int(sys.argv[4])
n_ag = int(sys.argv[5])
n_agents = [n_ag,n_ag]
n_sims = int(sys.argv[6])
sim_start = int(sys.argv[7])
attract_ddist = int(sys.argv[8])
kill_ddist = int(sys.argv[9])

##### Define functions #####

# Agent class with default values
class Agent:
    def __init__(self, id, type = "spcA", x = 0, y = 0, theta = 0, vx = 0, vy = 0, sdev =
math.pi/12, mag = 1.5):
        self.id = id
        self.type = type
        self.x = x
        self.y = y
        self.theta = theta
        self.vx = vx
        self.vy = vy
        self.sdev = sdev
        self.mag = mag
        self.avoided = 0
        self.attracted = 0
        self.caught = 0
    #
    def __str__(self):
        return f"Agent x:{self.x}, y:{self.y}, theta:{self.theta}, vx:{self.vx}, vy:{self.vy},
sdev:{self.sdev}, mag:{self.mag}"
    #
    # Update agent (passing list of other agents)
    def update(self, agents, avoid_ddist, avoid, attract_ddist, kill_ddist, attract, width, height):
        #
        # Store avoid_ddist and attract_ddist squared
        avoid_ddist_sq = avoid_ddist**2
        attract_ddist_sq = attract_ddist**2
        kill_ddist_sq = kill_ddist**2
        #
        if self.type == "spcA":
            # Check for proximity to B agents for being predated on
            if attract > 0:
                for other in agents:
                    # If the agent is of other type
                    if (other.type == "spcB"):
                        ## Calculate distance between this agent and other agent from other species
                        x_dist = other.x - self.x
                        y_dist = other.y - self.y
                        sq_dist = np.square(x_dist) + np.square(y_dist)

```

```

        #print(sq_dist)
        ## If agent is within kill detection radius of agent from species B, regenerate
randomly within the grid
        if sq_dist <= kill_ddist_sq:
            self.x = np.random.random()*width
            self.y = np.random.random()*height
            self.theta = np.random.random()*2*math.pi
            self.caught = 1
            break
        else:
            self.caught = 0
    if avoid > 0:
        # List of agents to avoid
        beta_list = []
        # Loop through all agents
        for other in agents:
            # If the agent is of other type
            if (other.type == "spcB"):
                ## Calculate distance between this agent and other agent from other species
                x_dist = other.x - self.x
                y_dist = other.y - self.y
                sq_dist = np.square(x_dist) + np.square(y_dist)
                #print(sq_dist)
                # If we're doing avoidance and this is a type A species
                if avoid > 0:
                    ## If within detection radius, save angle to agent
                    if (sq_dist <= avoid_ddist_sq):
                        # Calculate angle of avoidance
                        beta_list.append(math.atan2((self.y - other.y), (self.x -
other.x))%(2*math.pi))
                # If there were any agents to avoid:
                if len(beta_list) > 0:
                    # print(len(beta_list))
                    # Take mean angle of avoidance
                    self.theta = sum(beta_list)/len(beta_list)
                    self.avoided = 1
                    if len(beta_list)>1:
                        print('Number of agents being avoided: ' + str(len(beta_list)))
                else:
                    # Otherwise..sample theta as for CRW
                    self.theta = (np.random.normal(self.theta, self.sdev))%(2*np.pi)
                    self.avoided = 0
            else:
                # Otherwise..sample theta as for CRW
                self.theta = (np.random.normal(self.theta, self.sdev))%(2*np.pi)
                self.avoided = 0
        #
        # Re-calculate vx, vy with new theta
        self.vx = (self.mag * np.cos(self.theta))
        self.vy = (self.mag * np.sin(self.theta))
#
else: # for species B
    sq_list=[]
    gamma_list=[]
    if attract > 0:

```

```

for other in agents:
    # If the agent is of other type
    if (other.type == "spcA"):
        ## Calculate distance between this agent and other agent from other species
        x_dist = other.x - self.x
        y_dist = other.y - self.y
        sq_dist = np.square(x_dist) + np.square(y_dist)
        if sq_dist <= attract_ddist_sq:
            sq_list.append(sq_dist)
            gamma_list.append(math.atan2((other.y - self.y), (other.x -
self.x))%(2*math.pi))
        # If there were any agents to follow:
        if len(gamma_list) > 0:
            # Find closest
            ind = sq_list.index(min(sq_list))
            self.theta = gamma_list[ind]
            self.attracted = 1
            # Calculate new velocity
            if min(sq_list)<self.mag:
                self.vx = min(sq_list) * np.cos(self.theta)
                self.vy = min(sq_list) * np.sin(self.theta)
            else:
                self.vx = self.mag * np.cos(self.theta)
                self.vy = self.mag * np.sin(self.theta)
        else:
            # Otherwise..sample theta as for CRW
            # self.theta=math.pi/2
            self.theta = (np.random.normal(self.theta, self.sdev))%(2*np.pi)
            self.attracted = 0
            # Calculate new velocity
            self.vx = self.mag * np.cos(self.theta)
            self.vy = self.mag * np.sin(self.theta)
        else:
            # CRW
            # Calculate new theta
            self.theta = (np.random.normal(self.theta, self.sdev))%(2*np.pi)
            # self.theta=0
            # Calculate new velocity
            self.vx = self.mag * np.cos(self.theta)
            self.vy = self.mag * np.sin(self.theta)
        # update agent location - with torus boundary conditions
        newx = np.add(self.x, self.vx)%width
        newy = np.add(self.y, self.vy)%height
        self.x = newx
        self.y = newy
        #print("%d, %f, %f" % (a.id, a.x, a.y))

def
sim_abm(spc_labs,n_agents,n_steps,sdevs,mag,avoid,avoid_ddist,attract,attract_ddist,kill_d
dist):
    ## Function to simulate correlated random walks for agents from multiple species
    ## Includes switch for 1-way avoidance
    #
    ## Variables
    # spc_labs: species names (list)

```



```

# n_agents: number of agents per species (list)
# n_steps: number of time steps
# sdevs: standard deviations for sampling directional angles per species (list)
# mag: speed per species (list)
# avoid: switch: 0 - none; 1 - 1-way (A avoids B); 2 - 2-way (A avoids B and B avoids A)
# avoid_ddist: detection radius to trigger avoidance
# attract: switch: 0 - none; 1 - B attracted to A
# attract_ddist: detection distance to trigger attraction
# kill_ddist: distance within which agent is killed and regenerates
#
# Generate random agents of each species with random starting positions and directions
positions = []
agents = [] # List of agents
j = 1
for s in range(0, n_species):
    spc = spc_labs[s]
    for i in range(0, n_agents[s]):
        x = np.random.random()*width
        y = np.random.random()*height
        theta = np.random.random()*2*math.pi
        # a = Agent((s+1)*(i+1), spc, x, y, theta, 0, 0, sdevs[s], mag[s])
        a = Agent(j, spc, x, y, theta, 0, 0, sdevs[s], mag[s])
        #print(a)
        agents.append(a)
        j = j+1
#
# Update each agent at each time step
a = agents[1]
for i in range(1, n_steps+1):
    # Shuffle
    # random.shuffle(agents)
    for k in range(0, len(agents)):
        a = agents[k]
        others = agents[:k] + agents[k+1:]
        a.update(others, avoid_ddist, avoid, attract_ddist, kill_ddist, attract, width, height)
        #print("%d, %d, %s, %f, %f" % (i, a.id, a.type, a.x, a.y))
        positions.append((i, a.id, a.type, a.x, a.y, a.avoided, a.attracted, a.caught))
return positions

```

```
##### Space constraints #####
```

```

xmin=0
ymin=0
xmax=8000
ymax=8000

```

```

width = xmax - xmin
height = ymax - ymin

```

```

# n_steps = 30*24*60*6 # 30 days at 10-second time-steps
n_steps = 7*24*60*6

```

```
##### Species traits #####
```

```

## set number of species, agents per species and time steps; and standard deviation for
velocity distributions #####
n_species = 2
spc_labs = ['spcA','spcB']
# mag = [0.15,0.15] # velocity magnitude for each species (m/s)
mag = [1.5,1.5] # m/10seconds
sdevs = [math.pi/24, math.pi/24]# [math.pi/3,math.pi/4] # standard deviation for each species
(radians)
CT_ddist = 6

# Default attract/avoid parameter values
avoid = 0 # no avoidance
# avoid_ddist = 15 # detection distance for triggering avoidance
attract = 0
# attract_ddist = 0
# kill_ddist = 1

##### Update attract/avoid parameter values according to run category
if run_cat == 'CRW':
    avoid = 0
    attract = 0
    store_string = 'CRW'
elif run_cat == 'avoid1':
    avoid = 1
    store_string = 'avoid_1'
elif run_cat == 'attract1':
    attract = 1
    store_string = 'attract_1'
elif run_cat=='combi1':
    attract=1
    avoid=1
    store_string = 'combi1'

##### Simulations
#####

## Loop over n_sims for each test value and storing captures per simulation/species/camera
avoid_list=0
attract_list=0
caught_list=0
combi_list=0
for m in range(sim_start,sim_start + n_sims):
    print('Starting simulation ' + str(m) + ' for ' + store_string)
    ## Run simulation of animal movement
    positions =
sim_abm(spc_labs,n_agents,n_steps,sdevs,mag,avoid,avoid_ddist,attract,attract_ddist,kill_d
dist)
    print(run_cat)
    if run_cat == 'CRW':
        outfile = outpath + 'positions_time_step_10secs_nsteps_' + str(n_steps) + '_sim_' +
str(m) + '_nsims_' + str(n_sims) + '_nagents_' + str(n_agents[0]) + '_' + store_string
    elif run_cat == 'avoid1':
        outfile = outpath + 'positions_time_step_10secs_nsteps_' + str(n_steps) + '_sim_' +
str(m) + '_nsims_' + str(n_sims) + '_avoiddist_' + str(avoid_ddist) + '_nagents_' +
str(n_agents[0]) + '_' + store_string

```

```

elif run_cat == 'attract1':
    outfile = outpath + 'positions_time_step_10secs_nsteps_' + str(n_steps) + '_sim_' +
str(m) + '_nsims_' + str(n_sims) + '_attractdist_' + str(attract_ddist) + '_killddist_' +
str(kill_ddist) + '_nagents_' + str(n_agents[0]) + '_' + store_string
    elif run_cat == 'combi1':
        outfile = outpath + 'positions_time_step_10secs_nsteps_' + str(n_steps) + '_sim_' +
str(m) + '_nsims_' + str(n_sims) + '_dists_' + str(attract_ddist) + '_killddist_' + str(kill_ddist) +
'_nagents_' + str(n_agents[0]) + '_' + store_string
        # print(outfile)
        # Compressing Data
        # pickle.dump(positions,bz2.BZ2File(outfile,"wb"))
        with open(outfile + '.csv','w') as out:
            csv_out=csv.writer(out)
            csv_out.writerow(['time_step', 'id', 'type', 'x', 'y', 'avoided','attracted','caught'])
            csv_out.writerows(positions)
            avoid_list = avoid_list + sum(positions[i][-3] for i in range(len(positions)))
            attract_list = attract_list + sum(positions[i][-2] for i in range(len(positions)))
            caught_list = caught_list + sum(positions[i][-1] for i in range(len(positions)))
            combi_temp = [1 if positions[i][-2]+positions[i][-3]==2 else 0 for i in range(len(positions))]
            combi_list = combi_list + sum(combi_temp)
            print('avoidance events: ' + str(sum(positions[i][-3] for i in range(len(positions)))))
            print('attraction events: ' + str(sum(positions[i][-2] for i in range(len(positions)))))
            print('kill events: ' + str(sum(positions[i][-1] for i in range(len(positions)))))
            print('combination events: ' + str(sum(combi_temp)))

print("Total avoidance events: " + str(avoid_list))
print("Total attraction events: " + str(attract_list))
print("Total kill events: " + str(caught_list))
print("Total combination events: " + str(combi_list))

toc=timeit.default_timer()
print(toc-tic)

```