

UC Davis

UC Davis Electronic Theses and Dissertations

Title

Patterns of Learning in the Navigation of Selectively Foraging Mammals

Permalink

<https://escholarship.org/uc/item/90k3z4nm>

Author

Vining, Alexander Q

Publication Date

2023

Peer reviewed|Thesis/dissertation

**Patterns of Learning in the Navigation of Selectively Foraging  
Mammals**

By

ALEXANDER Q. VINING

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Animal Behavior

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA  
DAVIS

Approved:

---

(Damien Caillaud), Chair

---

(Margaret Crofoot)

---

(Andrew Sih)

Committee in Charge

2023

Patterns of Learning in the Navigation of Selectively Foraging Mammals

Copyright © 2023

by

Alexander Q. Vining

*“To my parents, Jodie, Neil, and Suzanne, for opening the whole world to me. And to the all the queers who learned to navigate that world before me, and bequeathed to me something like a map.”*

## Acknowledgements

First and foremost, I would like to thank all of the co-authors who contributed to the work in this thesis. None of this work was completed on my own, and I am incredibly grateful for the ideas, analysis, support, data, writing, and advice that other contributed toward this work. I am especially grateful to my thesis advisor, Meg Crofoot, the other members of my committee, Damien Caillaud and Andy Sih, and to the academic mentors I had along the way, including Mark Grote, Roland Kays, Seth Frey, and Jeffrey Schenk. I am immensely grateful to my peers in the Animal Behavior Graduate Group, The Crofoot Lab, The Smithsonian Tropical Research Institute, and the Max Planck Institute of Animal Behavior for the communities rich in support and exchange they fostered for me across seven years and three countries. Special thanks are warranted for Grace Davis, Shauhin Alavi, and Amelia Munson, who got me through the day-to-day, read seemingly endless drafts, and taught me more than I can measure. Calixto Rodríguez was a rock star in the field, and endured through absurd conditions with me to collect some of these data. And of course, thank you to my family for all of the support and encouragement, especially my parents to whom this is dedicated. Financial support for this work was provided by the University of California, Davis Animal Behavior Graduate Group and Hemispheric Institute for the Americas, The Smithsonian Tropical Research Institute, The National Science Foundation, the Max Planck Institute of Animal Behavior, and Richard Coss.

## Abstract

Memory and learning distinguish the movement of animals from other things, resulting in trajectories that change over time, but still repeatedly return to specific locations. The frequency and predictability with which animals return to favored locations, and the patterns of change in the paths they use to get there, offer insights into the cognitive systems of learning and memory that guide them. For animals that rely on resources that are concentrated in sparsely distributed, high value patches, these systems are particularly important for avoiding the costs of inefficient random-search foraging. In this thesis, I analyze the trajectories of such animals across multiple contexts and spatial scales, particularly the trajectories of primates and kinkajous (animals that look and behaves much like a primate, but are in fact *Carnivoran*). I find that most of these animals are quick to learn efficient paths between foraging locations, and some are able to generalize strategies for efficient navigation to novel contexts. There is some evidence that more selective foragers rely more on routine ‘traplines’ between multiple known locations, but are faster to deploy strategies that exploit changing resource distributions. Taken together, results of these studies suggest that diverse animal species integrate episodic-like memories into a cognitive map that helps them plan movements over large distances. Evidence that kinkajous flexibly use knowledge of detailed route-networks through complex canopy substrate to quickly find and exploit new resources, is particularly important because it highlights that the advanced development of these cognitive systems is not unique to primates and their socially complex groups, but may evolve readily in response to particular resource distributions and environmental properties.

# Contents

<b>Abstract</b>	<b>v</b>
<b>List of Tables</b>	<b>vii</b>
<b>List of Figures</b>	<b>viii</b>
<b>Introduction</b>	<b>1</b>
<b>1 Mild movement sequence repetition in five primate species and evidence for a taxonomic divide in cognitive mechanisms.</b>	<b>6</b>
Abstract . . . . .	7
Introduction . . . . .	8
Methods . . . . .	13
Results . . . . .	22
Discussion . . . . .	28
Ethics . . . . .	36
Data availability . . . . .	37
<b>2 The Movement and Behavior of Kinkajous in a Dynamic Foraging Hotspot</b>	<b>38</b>
Abstract . . . . .	39
Introduction . . . . .	40
Methods . . . . .	47
Results . . . . .	57
Discussion . . . . .	69

<b>3</b>	<b>Comparing the navigation of mammals on BCI reveals ecological drivers of cognitive evolution</b>	<b>76</b>
	Abstract . . . . .	77
	Introduction . . . . .	77
	Methods and Results . . . . .	81
	Discussion . . . . .	86
<b>4</b>	<b>Evidence of diet-mediated learning rates in the route-formation of four sympatric mammal species</b>	<b>89</b>
	Abstract . . . . .	90
	Introduction . . . . .	91
	Methods . . . . .	95
	Results . . . . .	105
	Discussion . . . . .	107
	<b>Conclusion</b>	<b>119</b>
	Building a Cognitive Map . . . . .	120
	Modelling a Cognitive Map . . . . .	126
	The role of Cognitive Map models in Optimal Foraging . . . . .	129
<b>A</b>	<b>Appendix</b>	<b>132</b>
	Chapter 1 . . . . .	132
	Chapter 2 . . . . .	134
	Chapter 3 . . . . .	135
	Chapter 4 . . . . .	136
	<b>Colophon</b>	<b>137</b>
	<b>References</b>	<b>145</b>



# List of Tables

1.1	Socio-ecological traits of the five primate species included in our sample.	12
1.2	Individuals and trial sample size included in the analysis. . . . .	17
2.1	Hypotheses and Predictions . . . . .	48

# List of Figures

1.1	Design of the navigational arrays used . . . . .	15
1.2	Raw data and model regressions for empirical analyses . . . . .	24
1.3	Posterior credible intervals for species-level slope and intercept parameters	25
1.4	Posterior distributions of alpha estimates . . . . .	29
2.1	Data collection methods . . . . .	50
2.2	Changes in kinkajou residence and flower density. . . . .	59
2.3	Posterior distributions of selected parameters from Model 1 . . . . .	63
2.4	Influence of memory on kinkajou decisions . . . . .	67
2.5	Posterior predictive probabilities of socially informed decisions . . . . .	68
2.6	Posterior distributions of predicted departure rates . . . . .	70
3.1	High Resolution Tracking of a Kinkajou . . . . .	83
3.2	Molly’s GPS tracks by experimental phase . . . . .	84
3.3	Molly retrieves a piece of mango from the Eastern provisioning station.	85
3.4	A Learning Model of Molly’s Inter-Station Travel Distances . . . . .	87
4.1	GPS Tracking Data . . . . .	101
4.2	Telemetry data and segmentation methods . . . . .	102
4.3	Posterior distributions of key model parameters. . . . .	108
4.4	Transit length prediction curves . . . . .	109

# Introduction

In her new book, “Techno-Scientific Practices”, Federica Russo (2022) claims that knowledge has four essential properties: it is relational, embodied, distributed, and material. This means, in brief, that knowledge is not an intangible element of individual minds, but it is the physical inscription of information across a network of actors. This knowledge allows actors in that network to ascribe meaning to the world “out there”. Russo builds on this notion to develop a theory of agency in the co-production of knowledge between scientists and technology, which asks us scientists to recognize that our ability to model the world depends on the knowledge embodied in the technology we use. Technology, then, plays an active role in producing further knowledge, and holds a degree of agency in the process. Russo deems this process of co-producing knowledge “poiêsis”. Esoteric as the idea may seem, it draws on a powerful combination of Philosophy of Science, Philosophy of Technology, and Philosophy of Information to challenge the idea of knowledge as belonging to “psyche”, and posits that knowledge is a physical thing that can be shared among individuals - and studied. I believe this perspective is not only broadly important for scientists and their work, but specifically relevant to the way we study knowledge in animals. We can know what an animal knows when our knowledge is shared. That is - when we ascribe the same meaning to world around us.

Not coincidentally, it is in large part the study of animals that brought us to this un-

## CHAPTER 0. INTRODUCTION

derstanding of knowledge. Anthony Chemero (2013) traces the tradition of embodied cognitive science, a foundation of Russo's work, to Darwin. Darwin's theory on the continuity of life challenged the predominant view in psychology that the existence of thought fundamentally held humans apart from other animals. Foundational Behaviorists that embraced Darwin's theory, such as B.F. Skinner, John Watson, and Ivan Pavlov, built theories of learning by studying the reinforcement of behavior in animals. Yet, in rejecting the "introspective" methods of other psychologists, these animal researchers left a gap between the development of behavior through reinforcement and the complex psychology of humans. Descartes' rationalist view of humans still reigned, and even as the prominent philosopher Martin Heidegger wrote extensively about the nature of "being in the world" as a continuous process of affording meaning to the state of matter around us (Heidegger 1962), he also described animals as "poor-in-world", unable to conceive of things around them as "being" (Heidegger 1946).

There was at least one psychologist, however, that found a middle ground. Edward Tolman's "Purposive Behavior in Animals and Men" (Tolman 1932) describes an extensive series of experiments, mostly based on the navigation of rats through mazes, in which animals 1) latently acquired information about the world around them, 2) behaved in novel situations as though they were considering multiple courses of action, and then 3) pursued courses of action that they could only have inferred (correctly) would most effectively bring about desired outcomes, no reinforcement required. Most famously, he demonstrated that when he changed the structure of a maze, rats would spontaneously take shortcuts to out-of-sight food rewards. Through these experiments, Tolman detailed an elaborate framework of the mind, which intricately connected the behaviorists' notion of reinforcement learning to Heidegger's notion of "being in the world". In his central claim, he made analogy of the mind to a map, which modeled the world outside by

connecting concepts by “distance” and “direction”. The key target of reinforcement, in this framework, was not stimulus-action pairs that yielded (un)desirable outcomes, but distance and direction models that accurately predicted the future. As concepts in an individual’s cognitive map were arranged and built up hierarchically into knowledge of the world, he claimed, increasingly complex thoughts, “raw feels”, and even consciousness would emerge naturally.

Tolman’s took a long time to become popular. It was not until O’Keefe (1976) measured the activity of individual “place neurons” in the brains of rats and claimed that these provided embodied and material evidence of Tolman’s cognitive map that “cognitivism” entered the psychological mainstream. The decades of neuroscience research since have slowly built a rich and compelling picture of how distance and direction can be embodied by linking place neurons to grid neurons, and how cognitive maps can emerge as these neurons are integrated into memory traces (engrams) through reinforcement (Leutgeb et al. 2005, Weber and Sprekeler 2018, Josselyn and Tonegawa 2020). Validating the ideas of both Tolman and Heidegger, the most fundamental feature of this system is increasingly seen not as a the mapping of space, but of time (Dragoi and Tonegawa 2011, Ekstrom and Ranganath 2018). In today’s psychology, the cognitive map concept is tightly connected to the study of episodic memory, and the notion that behavioral flexibility emerges from the ability to map information from the past onto possible futures (Suddendorf and Corballis 2007).

Even now, a tug-of-war exists between the rationalist approach, which identifies elements of human thought that demonstrate our unique capacity for reflection, and the embodied approach, which claims that all thought has material origins traceable through an organisms history. For example, a distinction is often made between “episodic-like” memory, evidenced by animals’ ability to act according to information about when, where, and

## CHAPTER 0. INTRODUCTION

how they experienced something, and episodic memory, which includes the awareness of one's-self "being" in the memory (Clayton 2017, Crystal and Suddendorf 2019). Tulving describes the representation of self in memory as *auto-noesis*, and claims that even if animals do this too, it is not possible to get in their heads in a way by which we could know for sure (Tulving 2005). Perhaps he is right, or perhaps it possible to understand how the process by which *auto-noesis* emerges and identify the signatures of that process in animal behavior, as Tolman once did when he sought evidence that rats make inferences and visualize the future.

If we are to hold the view that knowledge is relational, embodied, distributed, and material (as follows from Tolman's theory of cognitive maps) and that it is co-produced by networks of human *and* non-human actors (as put forth by Russo), then it is possible to see knowledge of the self not as an innate, ineffable aspect of human psychology, but as tangible, and perhaps inevitable, product of a cognitive organism engaging with the world around it. To study this process, however, we must step outside the world of laboratories and rats, and into the complex, poorly defined world in which the human capacity for *auto-noesis* evolved. We must ask how cognitive maps are built, from the ground-up, and put ourselves back into the equation. What is the distance and direction between the way we move through the world and the way an animal moves through world, and what can then be inferred about the knowledge we share?

The rest of this thesis is not about self-awareness or consciousness. I try to limit myself to the marginally less messy concepts of knowledge, learning, and inference. But my background as a cognitive psychologist with an interest in primate theory of mind informs the whole endeavor. My mission, in starting this Ph.D. program, was to take the big questions of psychology and primate evolution to the field, and ask how cognition works when you take away the controls. What I discovered is that in the field of behavioral

ecology, the concept of cognitive maps is used more like a tool than a theory (though it still often traced back to Tolman and the use of shortcuts by rats in mazes). The cognitive map, in ecology, takes the most literal form of Tolman's cognitive map: a mental topological representation of space. Though this may not accurately describe most animals' knowledge, it can be useful for approximating animal knowledge when studying their decision making. And, if the use of shortcuts is demonstrated in wild animals, its proof of a cognitive map's existence can be rightfully hailed as a milestone for the study of animal cognition. In each chapter of this thesis, I make an attempt to integrate the ecological concept of cognitive maps (which I henceforth call topological maps) and the psychological concept of cognitive maps.

What I mean by this is, I make an effort in each chapter to measure the ways different animal's movement reflects an inference at a specific spatial scale, and what the nature of this inference says about the animals' knowledge. In the first chapter, I explore multi-destination route-optimization by primates in local resources arrays. I ask how their movement provides evidence that they have learned either the specific patterns that optimize a given array, or general strategies that yield efficient routes in many arrays. In the second chapter, I bring this controlled, lab test to the messy field, exploring how kinkajous forage for flowers in the analagous but much more complicated crown of a balsa tree. Then I expand the scale, asking how kinkajous learn routes and navigate between feeding stations I experimentally baited with food over nearly two months. Finally, I expand the scope, analyzing the movements of four different species for signs of spatial inference as they forage for a particularly essential and dynamic food resource over a two three-month seasons. In the process, I have repeatedly broken down and rebuilt my own understanding of the nature of cognition. I hope, by the end, the reader might share in some facet of that knowledge.

# Chapter 1

## Mild movement sequence repetition in five primate species and evidence for a taxonomic divide in cognitive mechanisms.

Authors: Alexander Q. Vining<sup>1,2,3,4,†</sup>, L. Tamara Kumpan<sup>5,6,†</sup>, Megan M. Joyce<sup>7</sup>, William D. Aguado<sup>8</sup>, Eve A. Smeltzer<sup>5</sup>, Sarah E. Turner<sup>7</sup> & Julie A. Teichroeb<sup>5</sup>

1. Animal Behavior Graduate Group, University of California, Davis, USA
2. Department for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior, Konstanz, Germany
3. Department of Biology, University of Konstanz, Konstanz, Germany
4. Smithsonian Tropical Research Institute, Balboa, Republic of Panamá
5. Department of Anthropology, University of Toronto Scarborough, Toronto, ON M1C 1A4, Canada.
6. School of the Environment, University of Toronto, Toronto, Canada.
7. Geography, Planning and Environment, Concordia University, Montréal, Canada.
8. Department of Anthropology, Rutgers University, New Brunswick, USA.



## ABSTRACT

†These authors contributed equally: L. Tamara Kumpan and Alexander Q. Vining.

## Abstract

When animals forage, they face complex multi-destination routing problems. Traplining behaviour—the repeated use of the same route—can be used to study how spatial memory might evolve to cope with complex routing problems in ecologically distinct taxa. We analyzed experimental data from multi-destination foraging arrays for five species, two cercopithecine monkeys (vervets, *Chlorocebus pygerythrus*, and Japanese macaques, *Macaca fuscata*) and three strepsirrhines (fat-tailed dwarf lemurs, *Cheirogaleus medius*, grey mouse lemurs, *Microcebus murinus*, and aye-ayes, *Daubentonia madagascariensis*). These species all developed relatively efficient route formations within the arrays but appeared to rely on variable cognitive mechanisms. We found a strong reliance on heuristics in cercopithecoid species, with initial routes that began near optimal and did not improve with experience. In strepsirrhines, we found greater support for reinforcement learning of location-based decisions, such that routes improved with experience. Further, we found evidence of repeated sequences of site visitation in all species, supporting previous suggestions that primates form traplines. However, the recursive use of routes was weak, differing from the strategies seen in well-known traplining animals. Differences between strepsirrhine and cercopithecine strategies may be the result of either ecological or phylogenetic trends, and we discuss future possibilities for disentangling the two.

## Introduction

Foraging animals face complex multi-destination routing problems as they move between food sites. By locating routes that strategically connect biologically meaningful locations across space, animals may benefit from increased route efficiency during travel. However, finding the most efficient path connecting multiple sites requires the cognitive capacity to cope with the classic mathematical problem of the travelling salesperson (TSP), in which a set of fixed locations are each visited once before returning to the point of origin (Anderson 1983, Cook et al. 2011). When returning to the start is not required, the problem is termed an optimal Hamiltonian path problem, open-TSP, or shortest path problem (Janson 2014). The number of possible routes in a TSP-like problem increases exponentially as sites are added, quickly making computation of the most efficient route intractable. Animals with differing computational abilities likely evolved alternative strategies to address these navigational challenges in the wild (Janmaat et al. 2021).

One approach to solving TSP-like foraging problems is to learn, through trial and error, efficient sequences of location transitions within an array; a method sometimes called iterative learning. In bumblebees, an individual-based decision model of this type has been used to explain traplining behaviour (Lihoreau et al. 2012b, Reynolds et al. 2013)—a foraging strategy that involves the repeated use of a circuit to feeding sites in a stable and predictable sequence (Thomson et al. 1997). In simulations and foraging experiments, bumblebees eventually converge on an optimal or near-optimal route through simple arrays via trial and error with iterative improvement, where the likelihood of a given transition between locations increases if that transition was used in a trial that resulted in a final path length decrease, relative to previous trials (Reynolds et al. 2013).

Alternatively, animals may utilize a different cognitive process, using pre-existing heuris-

## INTRODUCTION

tics (Siekman and Crocker 2010), which are simple “rules of thumb,” either learned or innate, for solving multi-destination routes. For example, animals may prefer to always move to the nearest location that has not yet been visited—a strategy called the nearest-neighbour rule (NNR) (Ohashi et al. 2008). This appears to be a solution used by a variety of taxa to solve navigational challenges due to the low cognitive effort it requires (bees (Ohashi et al. 2008); rats (Bureš et al. 1992); non-human primates (Cramer and Gallistel 1997, Teichroeb and Smeltzer 2018, Teichroeb and Vining 2019)). Animals might also rely on other heuristics such as the “convex hull,” where a mental loop is placed around the targets and visits occur in order based on distance from the outer edge (Janson 2000, Teichroeb 2015, Teichroeb and Vining 2019). Importantly, though many heuristic rules are possible (Janmaat et al. 2021), it can be expected that those that are most generalizable to the greatest number of resource distributions should be the most adaptive. The full suite of heuristics that humans use to solve TSP-like problems is still unknown and new possibilities continue to be generated (Ryser-Welch et al. 2015). One approach to understanding human heuristics and how such heuristics emerge is to study how our closest relatives, primates, solve similar challenges within ecological contexts.

Several cognitive processes have been proposed as drivers of trapline foraging, including iterative learning and nearest-neighbour heuristics (Lihoreau et al. 2012a). Thus, traplining behaviour can inform how spatial memory might evolve to cope with complex routing problems in ecologically distinct taxa (Ayers et al. 2018). Further, primate use of traplining has been suggested in the literature (Garber 1988, Janson 1998, Dew and Wright 1998, Di Fiore and Suarez 2007, Noser and Byrne 2010) however, these studies did not explicitly test the hypothesis that primates form traplines.

Investigating whether primates use traplines, how quickly they develop traplines, and how closely their routes resemble standard traplines, would reveal how they manage po-

tentially complex trade-offs between reinforcement learning, cognitively simple heuristics, and the development of more complex, but flexible heuristics.

To examine primate strategies for approximating solutions to multi-destination routes, we analyzed movement data from multi-destination foraging arrays for five species, two cercopithecine monkeys and three strepsirrhines: vervet monkeys (*Chlorocebus pygerythrus*), Japanese macaques (*Macaca fuscata*), fat-tailed dwarf lemurs (*Cheirogaleus medius*), grey mouse lemurs (*Microcebus murinus*), and aye-ayes (*Daubentonia madagascariensis*). We measured the recursive movement characteristic of traplining behavior by quantifying 1) the repetitiveness of foraging sequences in which individuals completed experimental arrays over multiple trials, and 2) the distance travelled in the arrays fitted to simple learning curves. Foraging arrays were originally designed for previous studies that examined different theoretical questions (Teichroeb 2015, Teichroeb and Aguado 2016, Teichroeb and Smeltzer 2018, Teichroeb and Vining 2019, Kumpan et al. 2019, Joyce 2021) and thus they vary for some species, however these data are useful to compare these five primate species in their navigational strategies. Each array required decisions in small-scale space where all sites were visible to one another and exploration between platforms was not required.

We predicted differences in movement patterns based on known dietary and ecological differences between our study species (Table 1) (Rosati 2017). Fat-tailed dwarf lemurs are small nocturnal cheirogaleid strepsirrhines that are frugivorous (Lahann 2007). Since they exhibit a strong reliance on stationary replenishing food items and have been found to improve their accuracy and speed with experience in a multi-destination array (Teichroeb and Vining 2019), we predicted that they would show a decrease in distance traveled with experience consistent with an iterative improvement model. Grey mouse lemurs are also small nocturnal cheirogaleid strepsirrhines but exhibit pronounced feed-

## INTRODUCTION

ing plasticity, being omnivorous and consuming fruit, gum, insect secretions, and small vertebrates (Dammhahn and Kappeler 2008). Aye-ayes are larger nocturnal strepsirrhines and are also known to exhibit diverse diets depending on where they live, relying on insects in primary forest (Erickson et al. 1998, Randimbiharirina et al. 2018) and consuming fruits, coconut, flowers, and flower nectar in secondary forests (Ancrenaz et al. 1994, Andriamasimanana 1994). For these two species that exhibit a stronger reliance on ephemeral and mobile food items, we predicted a minor decrease in distance traveled with experience, and minimally repetitive route sequences. This prediction is supported by results showing more randomness (i.e., rarely repeating foraging paths in experimental set-ups) in the navigation patterns of strepsirrhines relying on ephemeral resources (Teichroeb and Vining 2019). In addition, while vervet monkeys are medium-bodied African monkeys that rely mainly on fruits and flowers and opportunistically prey on insects (Wrangham and Waterman 1981, Bruerton et al. 1991), Japanese macaques are medium-bodied Asian monkeys with a diet that varies as a function of forest type but generally includes new and mature leaves, flowers, fruits, insects, and fungi (Hanya 2004, Go 2010). Since vervets have previously been shown to rely strongly on navigational heuristics (Cramer and Gallistel 1997, Teichroeb 2015, Teichroeb and Aguado 2016, Teichroeb and Smeltzer 2018) and Japanese macaques have similar diets, we predicted that these species would show patterns consistent with heuristic use: distance traveled and route-repetition would both stay consistent throughout trials (because they stick to, for example, a solution guided by the nearest neighbor heuristic), but distance traveled would be significantly lower and route-repetition significantly higher when compared to our simulations of other navigational strategies.

We tested two hypotheses focused on two different cognitive processes to understand how primates approximate solutions to multi-destination routes. H1—Primates use rein-

Table 1.1: Socio-ecological traits of the five primate species included in our sample.

Species	Activity Pattern	Social Structure	Leaves	Insects	Fruit	Gum	Flowers
Strepsirrhines							
Fat-tailed dwarf lemur ( <i>C. medius</i> ) <sup>1,2</sup>	Nocturnal	Small groups	None	Low	High	None	High
Grey mouse lemur ( <i>M. murinus</i> ) <sup>3,4</sup>	Nocturnal	Solitary foragers	Low	High	Mod	High	Mod
Aye-ayes ( <i>D. madagascariensis</i> ) <sup>5</sup>	Nocturnal	Solitary	None	High	High	Low	Mod
Cercopithecines							
Vervet monkeys ( <i>C. pygerythrus</i> ) <sup>6</sup>	Diurnal	Large groups	Mod	Low	High	None	High
Japanese macaques ( <i>M. fuscata</i> ) <sup>7</sup>	Diurnal	Large groups	High	Low	High	None	Mod

Fietz and Ganzhorn (1999)<sup>1</sup>, Fietz et al. (2003)<sup>2</sup>, Génin (2003)<sup>3</sup>, Thorén et al. (2011)<sup>4</sup>, Sefczek et al. (2012)<sup>5</sup>, Cancelliere et al. (2018)<sup>6</sup>, Hill (1997)<sup>7</sup>

forcement learning of location-based decisions that are more likely to lead to shorter and more consistent routes and less travel with experience. This is a slow and cognitively costly process of information acquisition and alteration of behaviour that may require spatial rehearsal within visuo-spatial working memory and mental ‘chunking’ of locations together (Zimmer et al. 2010). H2—Primates could use a less cognitively costly process, heuristics, that can be generalized to approximate solutions to multi-destination routes in many different arrays and if the best heuristics are applied, lead to both distance traveled and route repetition remaining relatively consistent throughout trials. Additionally, under this hypothesis, we would expect animals to immediately use more efficient paths through resource arrays than expected by random decision making. Under the null hypotheses of random decision-making, we would expect to see no decrease in distance traveled with experience and minimally repetitive route sequences. Since traplining is beneficial for renewable and spatially predictable resources, ecological variables may also predict traplining in primates. Specifically, frugivorous species or species that rely on renewable, predictable resources should be more likely to trapline than species relying on insects or other mobile, non-renewable resources (Lihoreau et al. 2011).

We also anticipated that our results might be impacted by living conditions because some species were tested in captivity whereas others were wild. This could lead to many variables altering experimental outcomes, such as internal factors influencing motivation,

## METHODS

unknown variables such as diet, health status, distractions faced by wild but not captive animals, and unnatural territory sizes experienced in captivity. Coincidentally, all strepsirrhines were captive and both cercopithecoids were wild, so these potentially confounding factors may lead to findings of similar navigational strategies for strepsirrhines versus catarrhines.

## Methods

### Study subjects

We conducted foraging experiments on strepsirrhines ( $N_{individuals} = 18$ ) at the Duke Lemur Center (DLC), North Carolina, from February to November 2015<sup>13</sup>. Our sample includes six fat-tailed dwarf lemurs (3–16 years of age, 3 males, 3 females), six gray mouse lemurs (3–7 years of age, all female), and six aye-ayes (17–32 years of age, 2 males, 4 females). Because these species are solitary and nocturnal, most animals were housed singly and were kept on a reversed light cycle such that they were active and could be tested during the day. Housing conditions were similar for all individuals, and they were all fed daily in a similar manner with a diet that included fruits, vegetables, meal worms, and monkey chow (details in Teichroeb and Vining (2019)).

All vervet data were collected on wild animals ( $N_{individuals} = 12$ ) at Lake Nabugabo, Uganda (0°22′–12° S and 31°54′ E) during four separate field seasons (April–June 2013, Double Trapezoid array, M group (Teichroeb 2015); June–September 2013, Pentagon array, M group (Teichroeb and Aguado 2016); August–September 2015, Z-array, M group (Teichroeb and Smeltzer 2018); July–August 2017, Pentagon array, KS group (Kumpan et al. 2019)). M group was composed of between 21–28 individuals, containing 2–3 adult

males, 7–9 adult females, 2 subadult males, 1–3 subadult females, and 9–12 juveniles and infants. KS group was composed of 39–40 individuals including 5 adult males, 11 adult females, 3 sub-adult males, 5 sub-adult females, and 15–16 juveniles and infants. All individuals were reliably identified based on natural features (details in Kumpan et al. (2019); Teichroeb (2015); Teichroeb and Aguado (2016); Teichroeb and Vining (2019)). Outside of foraging experiments, wild vervets were not provision fed.

All Japanese macaque data ( $N_{individuals} = 10$ ) were collected at the Awajishima Monkey Centre (AMC), Awaji Island, Japan (34°14'43.6" N and 134°52'59.9" E) between July and August 2019 (Z-array26). AMC is a privately-run tourist and conservation center visited by a large group of free-ranging Japanese macaques (~400 individuals) called the “Awajishima group” (Kaigaishi et al. 2019). The group is composed of different-aged individuals of both sexes, with bachelor males and bachelor male groups living around the periphery (Nakagawa et al. 2010). The Awajishima group forages on wild foods for much of their dietary requirements but is also provision-fed a combination of wheat and soybeans, supplemented with peanuts, fruits, and vegetables twice daily for ~10 months of the year (details in Kaigaishi et al. (2019); Turner et al. (2008); Turner et al. (2018)).

## Study design

### Navigation arrays

The strepsirrhines and vervets were tested on a “double-trapezoid” shaped multi-destination array with six feeding platforms (Teichroeb 2015, Teichroeb and Vining 2019) modified from Lihoreau et al. (2012a) (Fig. 1a), where there were 720 possible routes (6!). Three different double-trapezoid arrays were built to account for differences in body size: one for the smaller dwarf and mouse lemurs, one for the mid-sized aye-ayes,



## METHODS

and one for the larger, wild vervets. Arrays were scaled such that the distance from platform 1–2 (the shortest distance between targets) was approximately twice the body length of the subject species. This does not necessarily control for the perceptual abilities of the different species, and issue we address in the discussion. Vervets were additionally tested on a Z-shaped array with six feeding platforms (720 possible routes, Fig. 1b (Teichroeb and Smeltzer 2018)), and a pentagon-shaped array with five feeding platforms (120 possible routes, Fig. 1c (Teichroeb and Aguado 2016, Kumpan et al. 2019, Joyce 2021)). Japanese macaques were tested on an identically sized Z-array (Joyce 2021).




Figure 1.1: Design of the navigational arrays used, with (a) the Double Trapezoid array used for *Cheirogaleus medius*, *Microcebus murinus*, *Daubentonia madagascariensis*, and *Chlorocebus pygerythrus*. Three different arrays were built and scaled to the body size of animals (see “Methods”). (b) The Z-array used for *C. pygerythrus* and *Macaca fuscata*. The same size array was used for both species because they are similar in adult body lengths (vervet mean range from four sites: 34.5–42.6 cm (Turner et al. 1997), Japanese macaque mean range from six sites: 48.9–59.7 cm (Fooden and Mitsuru 2006)). (c) The Pentagon used for *C. pygerythrus*. Distances here are unitless but roughly proportional to the body size of each species tested. Created in R version 4.0.4 and ProCreate.

For strepsirrhine trials, DLC staff captured individuals in their enclosures and transported them in padded crates to the testing room. The dwarf and mouse lemur array was set up in a specially designed box ( $0.91 \times 1.83$  m) with a small compartment to contain strepsirrhines for rebaiting between trials. The aye-aye array was set up on the ground in a room measuring  $2.44 \times 4.27$  m, where subjects stayed during the duration of their daily

## CHAPTER 1. SEQUENCE REPETITION

trials Teichroeb and Vining (2019). Vervet and macaque trials occurred when individual monkeys voluntarily left their group to participate in foraging experiments alone. Vervet arrays were set up using wooden feeding platforms (0.75 m long, 0.75 m wide×0.75 m high) placed in an outdoor clearing measuring roughly 10×14 m in the home range of the study group. Japanese macaque arrays were also set up using small wooden feeding tables (0.40 m long, 0.30 m wide, 0.21 m high), covered in green plastic labeled with the platform number. Two identical arrays were built in neighbouring provision-feeding fields at the AMC (Near Lower Field: ~10×35 m, and Far Lower Field: ~15×45 m).

In these studies, all platforms were baited with a single food item. The reward used varied by species (strepsirrhines: grape piece, apple piece, honey, agave nectar, or nut butters, vervets: slice of banana, piece of popcorn; macaques: single peanut or piece of sweet potato). Strepsirrhines have sensory adaptations for using olfaction to locate food (Cunningham et al. 2021), while the cercopithecoids are heavily reliant on vision to locate resources (Teichroeb and Chapman 2014), so we ensured that each platform was baited with identical food items within a trial that smelled and looked the same to avoid biasing where the animals chose to go. Platforms for the wild monkeys were not rebaited between trials until all animals were  $\geq 20m$  away and the entire sequence could be rebaited before their return.

For all species, we started a trial when the tested individual entered the array and took the reward at a platform. We then recorded each successive platform visit (including revisits to empty platforms) until all rewards had been collected ending the trial. In our analyses, we included a total of 852 trials collected over six navigational experiments, completed by 40 unique individuals (18 lemurs, 12 vervets, 10 macaques) (Table 2).

## METHODS

Table 1.2: Individuals and trial sample size included in the analysis.

Species	Array	Individuals	Trials
<i>C. pygerythrus</i>	Pentagon	4	110
	Z-Array	9	269
	Double-Trapezoid	3	70
<i>C. medius</i>	Double-Trapezoid	6	158
<i>M. murinus</i>	Double-Trapezoid	6	88
<i>D. madagascariensis</i>	Double-Trapezoid	6	42
<i>M. fuscata</i>	Z-Array	10	115
Totals		44	852

### Data simulations

In addition to empirically collected data, we simulated agents learning to travel efficiently in the same set of arrays using a simple iterative-reinforcement learning model based on the one used by Reynolds et al. (2013) to test for traplining behavior in bumblebees. In this model, agents move randomly between locations in an array until they visit all locations, then reset for another trial. If the agent completed a trial by travelling less distance than on previous trials, the probability of the agent repeating location-to-location transitions that occurred in that trial increased for future trials by a reinforcement factor. Initial transition probabilities were inversely proportional to the distance between two locations. Unlike Reynolds et al. (2013), our simulated agents started at a random location and were not required to return to that location to complete the trial. This matches the trial structure used in our experiments (open-TSP), and reflects multiple central place foraging patterns in primates (Chapman et al. 1989). Finally, agents could not return to the location they had just come from, using an “avoid the last location” behavioral heuristic observed in nectivores (Pyke 1978a, 1981), which prevented agents from getting stuck in “loops” between two locations (Online Supplement - Simulation Validation).

Within each of the arrays used to collect empirical data, we ran simulations with reinforcement factors of 1 (no reinforcement), 1.2 (mild reinforcement), and 2 (strong reinforcement). These factors indicate the magnitude by which the probability of transitions taken in a new shortest-path trial were multiplied, before re-normalizing all transition probabilities. For each array and reinforcement factor combination, we ran 100 agents that each completed 120 trials, where there was an equal probability of starting each trial at any location. Then, for each array and reinforcement factor combination, we ran 100 additional simulations per species tested in the given array, where the probability of starting a trial at any location was equal to the empirically observed location-starting probabilities of the respective species.

These simulations were designed to help us test predictions of our two hypotheses regarding primate learning and decision making within the arrays. If primates learn to solve navigational arrays efficiently by reinforcing movements between platform pairs, they should exhibit overall greater receptiveness in their sequences of location visits than reinforcement factor 1 simulations, and a greater decrease over time in total distance traveled to complete the arrays. If primates are pre-disposed to navigate arrays using appropriate heuristics, they should exhibit shorter distances travelled on initial trials than in simulations.

## Data analysis

From the raw sequences of locations visited in each trial, we calculated two metrics: minimum distance traveled, and the proportion of platform revisits that occurred within identical 3-platform visit sequences (determinism-DET) Ayers et al. (2018). All calculations were done using R version 4.0.4 (Team 2020) and packages rstan (Stan Development Team 2023) and tidyverse (Wickham et al. 2019). A fully reproducible

## METHODS

data notebook containing this work, as well as all analyzed data, is available at <https://github.com/aqvinig/Do-Primates-Trapline>. All figures were created by AQV in R version 4.0.4 and ProCreate.

### Distance traveled

To calculate minimum distance traveled, we created a distance matrix for each resource array containing the relative linear distance between any two resource locations. These minimum linear distances approximate the distances traveled by the animals, which may not necessarily be linear. We then summed the linear distances for all transitions made in a trial. Because resource arrays were scaled to the subject species' body size, these relative distances were standardized.

### Determinism

Given a sequence of observations, Ayers et al. (2015) defines determinism (*DET*) as the proportion of all matching observation-pairs (recurrences) that occur within matching sub-sequences of observations (repeats) of a given length (*minL*). This metric has been previously used to distinguish sequences of resource visitation generated by trawling behaviour from sequences generated by known processes of random movement within a given resource array (De Groeve et al. 2016, Riotte-Lambert et al. 2017, Ayers et al. 2018). It has several advantages in the analysis of foraging patterns, including the ability to detect repeated sequences between non-consecutive foraging bouts, imperfect repeats in sequences (i.e., omission or addition of a particular site), and distinguishing between forward- and reverse-order sequence repeats (Ayers et al. 2015).

We adapted the methods of Ayers et al. (2015) to calculate the number of recurrences and repeats generated by the sequence of location visits in each trial of our experiments

and simulations. Based on an analysis of the sensitivity of *DET* scores to the parameterization of *minL*, we set *minL* to three for our calculations (Online Supplement - Sensitivity Analysis).

## Statistical analyses

### Learning rates

We modelled distance travelled as a function of trial number, species, and individual. Metrics of animal performance on learned tasks are known to follow power functions over time and experience Jaber (2015), so we a priori applied log transformations to distance travelled and trial number, then fit a linear model. Thus, in the resulting model, the intercept can be interpreted as an estimated distance travelled on the first trial and the slope can be interpreted as the exponent of a learning curve. We modelled species and individual effects on the intercept by summing an estimated grand mean ( $\mu_0$ ), species level deviation ( $\mu_{sp}$ ), and individual level deviation ( $\mu_{id}$ ). We treated species and individual level effects on the learning rate parameter (slope) the same way, summing a grand mean ( $b_0$ ), species level deviation ( $b_{sp}$ ), and individual level deviation ( $b_{id}$ ). We estimated additional parameters for the variance of individual level deviations in intercept and slope ( $\sigma_{\mu_{id}}$  and  $\sigma_{b_{id}}$ , respectively). Finally, after finding residuals in an initial analysis to have variances predicted by trial number and species, we estimated a separate error variance for each species ( $\sigma_{\epsilon_{sp}}$ ) and weighted the standard deviations of the resulting error distributions by dividing them by the square root of one plus the trial number.

We set regularizing priors on the model parameters, assuming distances travelled would remain within one order of magnitude of the most efficient route, but not setting any strict boundaries. For the grand mean of the intercept, we used a normal distribution centered

## METHODS

around twice the minimum possible distance required to visit all platforms in the array, with a variance of one. For the grand mean of the slope and all species and individual level deviations to the slope and intercept, we used normal distributions centered at zero with variance of one. For all error terms, we used half-cauchy priors with a location parameter of zero and a scale parameter of one. The full, hierarchical definition of the model is given in Eq. (1).

$$\begin{aligned} \text{Distance} &\sim \mu_0 + \mu_{sp} + \mu_{id} + (b_0 + b_{sp} + b_{id})\text{Trial} + \epsilon \\ \mu_0 &\sim N(4.78, 1) \\ \mu_{sp}, b_0, b_{sp} &\sim N(0, 1) \\ \mu_{id} &\sim N(0, \sigma_{\mu_{id}}) \\ b_{id} &\sim N(0, \sigma_{b_{id}}) \\ \epsilon &\sim N(0, \sigma_{\epsilon_{sp}}/\sqrt{1 + \text{Trial}}) \\ \sigma_{\mu_{id}}, \sigma_{b_{id}}, \sigma_{\epsilon} &\sim \text{HalfCauchy}(0, 1) \end{aligned}$$

## Determinism

To compare *DET* between species, and between empirical and simulated data, we created a binomial model of expected repeats generated in a trial given the number of recurrences

(Eq. 2).

$$\text{Repeats} \sim \text{binom}(\text{Reursions}, \text{DET})$$

$$\text{DET} = \text{logit}^{-1}(\alpha)$$

$$\alpha = \alpha_0 + \alpha_{sp} + \alpha_{src} + \alpha_{int} + \alpha_{ID}$$

$$\alpha_0, \alpha_{sp}, \alpha_{src}, \alpha_{int} \sim N(0, 1)$$

$$\alpha_{id} \sim N(0, \sigma_{id})$$

$$\sigma_{id} \sim \text{HalfCauchy}(0, 1)$$

where  $\alpha_0$  is the mean intercept,  $\alpha_{sp}$  is one of four coefficients determined by the species (simulations are of the “species” which was used to assign its starting-location probabilities),  $\alpha_{src}$  is one of four coefficients determined by the source (empirical data and each level of simulated reinforcement factor),  $\alpha_{int}$  is one of 16 interaction coefficients (each possible combination of  $\alpha_{sp}$  and  $\alpha_{src}$ ), and  $\alpha_{id}$  is a varying effect of the individual. Because the length of a sequence affects  $DET$ , we limit our analysis of  $DET$  to the sequences generated by a subject’s or an agent’s first ten trials. Subjects that completed fewer than ten trials were excluded from this portion of the analysis.

## Results

### Distance travelled and learning rates

#### Double trapezoid

In the Double Trapezoid array, the estimated minimum distance travelled on first trial (model intercept) was far lower for vervets (95% Credible Interval (CI): 4.06–4.29) than



## RESULTS

for the strepsirrhine species. Between the strepsirrhines, the estimated intercept was lowest for dwarf lemurs, then mouse lemurs, and finally aye-ayes (95% CIs: dwarf lemur: 4.73–5.18, mouse lemur: 4.67–5.37, aye-aye: 4.86–5.52). The mean regression and sample regressions from the posterior are plotted through the data for each species and all arrays in Fig. 2. Because there is considerable overlap in the posterior distributions of intercept parameters for strepsirrhine species, the hypothesis that dwarf lemurs travel the least in initial trials is only weakly supported.

More than 80% of the posterior distribution for the slope parameter for mouse lemurs falls below zero. This rises to >90% for aye-ayes, and 100% for dwarf lemurs (Fig. 3). The mean estimated slope parameter for vervets is positive, though 20% of the posterior sample falls below zero. Thus, there is strong evidence that dwarf lemurs and aye-ayes improve their performance with experience and moderate evidence of the same for mouse lemurs, but no evidence for vervets. However, even with improvement over experience, none of the strepsirrhine species reach the initial performance of vervets, nor does our model predict they would until at least 120 trials.

When the same statistical model used to analyze the empirical data was applied to simulated agents using a reinforcement learning algorithm to transition between locations in the same array, the model did not fit as well. Model residuals were highly structured, containing a strong negative relationship with predicted distance travelled and starkly different variances dependant on trial number, suggesting our statistical model poorly describes the improvement dynamics of the iterative reinforcement algorithm. Nonetheless, estimated learning rates for agents of each level of reinforcement factor showed predicted patterns; the “species” level posterior for agents with reinforcement factor of 1 contains 0 (95% CI: 0.009 to 0.031), it is slightly (but entirely) below zero for those with a reinforcement factor of 1.2 (95% CI: 0.042 to : 0.0003), and well below 0 for individuals

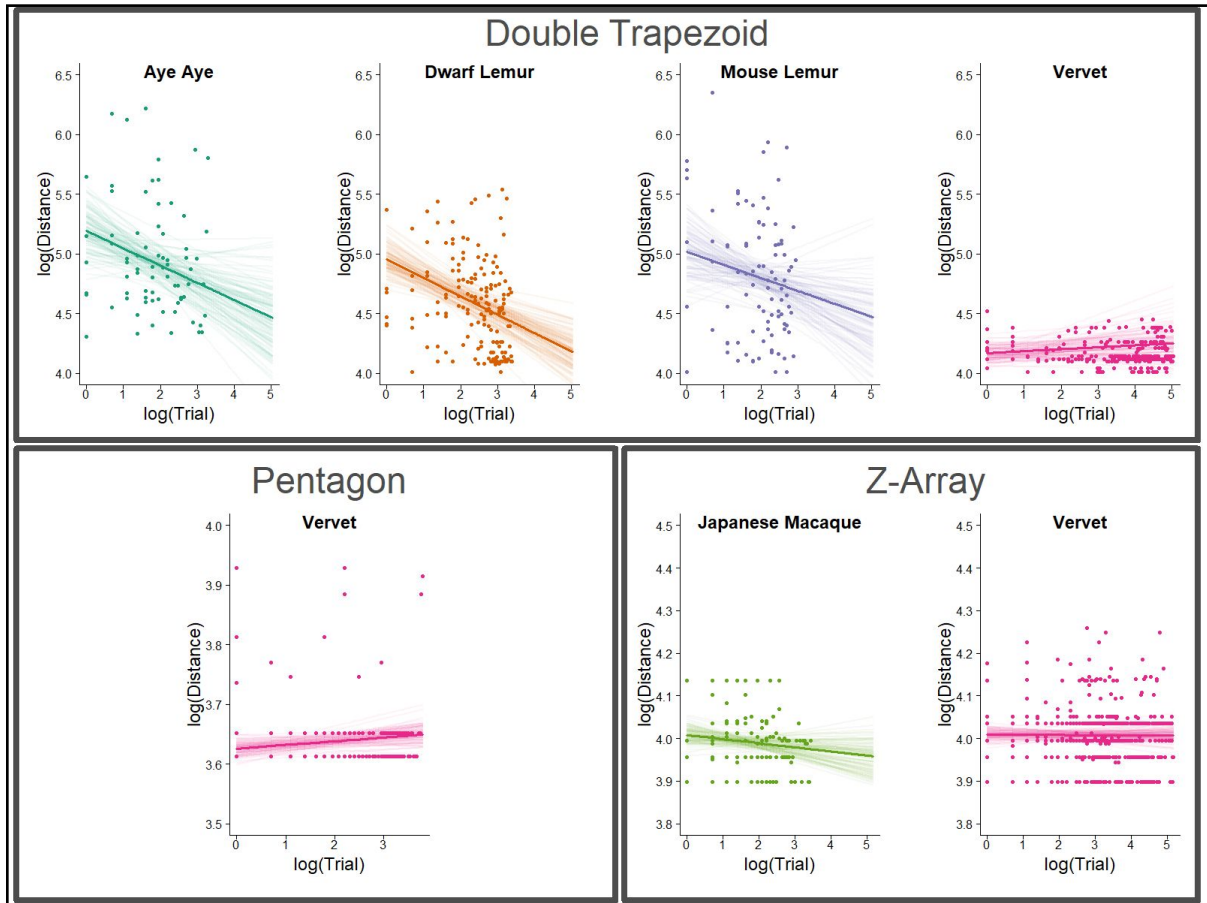


Figure 1.2: Raw data and model regressions for empirical analyses. The mean linear regression of log minimum distance traveled over log trial number for each species (thick lines), plotted through raw data (points) for all empirical observations. Thin lines represent linear regressions of 100 randomly selected samples from the posterior. All regressions are determined by adding the species-specific deviations for the intercept and slope in a given sample to the estimates of the grand mean for the intercept and slope, respectively. Note that axes are fixed within arrays, but not between.

RESULTS

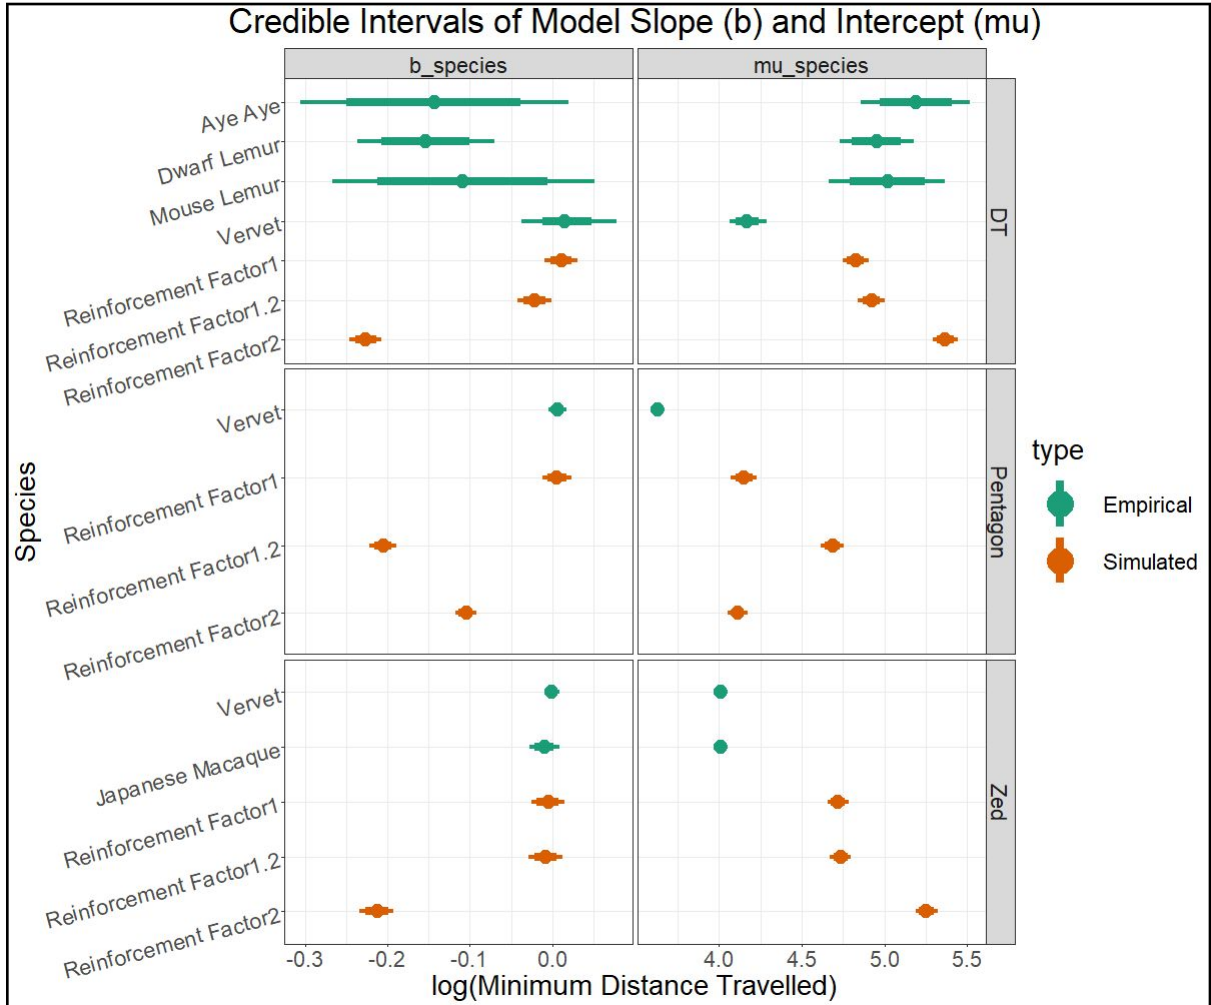


Figure 1.3: Posterior credible intervals for species-level slope and intercept parameters. 80% intervals (wide bars) and 95% intervals (thin bars) of posterior distributions of the model intercept for all species and simulated agents. The only species that show clear improvement with experience are dwarf lemurs in the Double Trapezoid (credible interval does not contain 0, dashed red line), though mouse lemurs and aye ayes also show some evidence of improvement in this array. Across arrays, vervets and Japanese macaques travel significantly less on initial trials than other species, and less than expected given the distance weighted location transitions of the Reinforcement Factor 1 agents. Created in R version 4.0.4.

with a reinforcement factor of 2 (95% CI: 0.246 to : 0.207). Unexpectedly, the estimated intercept (initial performance) is much higher for agents with a reinforcement factor of 2 (95% CI 5.29–5.44) than for those with 1.2 (95% CI: 4.84–5.00) or 1 (95% CI: 4.75–4.90). Given that all agents (regardless of reinforcement factor) are known to behave identically on the first trial, this result is certainly an artifact of the mismatch between the statistical learning model and the actual dynamics of the reinforcement algorithm. Plotting the regressions through the data reveals a floor effect, where agents lock into one of a few short routes well before trials end, which may contribute to this poor fit. Inspection of individual improvement curves over trial number also point to a cause of this mismatch; individuals with reinforcement factors of 2 show initially flat or even increasing distances traveled, followed by sharp decreases. The timing of this drop-off varies greatly by individual (never even appearing for some) and explains both the high inter-individual variance in estimated model parameters for reinforcement factor 2 agents and the inaccurate intercept parameter estimates. Further analysis and visualization of data from simulations can be found in the online Data Notebook.

## Pentagon

Vervets in the Pentagon Array showed results comparable to their performance in the Double Trapezoid. Initial performance started very close to optimal (Fig. 2), with individuals circumnavigating the perimeter of the array on the majority of early trials. This pattern is consistent with a predisposition for a number of route selection heuristics, including nearest-neighbor and convex-hull. In later trials, there was a tendency to select slightly less optimal routes, as indicated by the positive value of  $b_0 + b_{vervet}$  in the majority of our posterior samples (Fig. 3). However, with support from only 89% of the posterior sample, this positive trend is insufficient to convincingly rule out the possibility that

## RESULTS

vervets do not alter their route selection heuristics with experience; instead, less efficient routes in later trials may have occurred by chance.

Simulations in the Pentagon showed a different pattern than in the other arrays. Agents with a reinforcement factor of 1.2 exhibited a strong learning effect, with a posterior distribution of the slope greater in magnitude than agents with reinforcement factor of 2. This can likely be attributed to a highly overestimated intercept, as reinforcement factor 1.2 agents did not achieve short routes as often as reinforcement factor 2 agents.

### Z-Array

Japanese macaques and vervets in the Z-Array performed similarly on initial trials (species intercept 95% CIs: vervets 3.99–4.03, macaques 3.97–4.04; Fig. 3). While vervets remained consistent in their performance across trials (as in other arrays), Japanese macaques showed slight improvement over trial number, with 86% percent of posterior samples estimating a negative slope. This is not sufficient evidence to rule out the hypothesis that decreased distance travelled in later trials occurred by chance, but is notable given how close the initial performance of Japanese macaques was to optimal. Simulations in the Z-array yielded similar results to the Double Trapezoid, though reinforcement factor 1.2 agents appear slightly less able to find efficient routes in this array.

### **Traplining in primates: recursion relative to random transitions in the arrays**

Posterior estimates of parameters in the binomial model of *DET* were well mixed (online supplement “DET Analysis”). In our empirical data, estimates of *DET* in the Double-Trapezoid array were highest for aye ayes and lowest for vervets, but overlap between the

posterior distributions of these estimates (Fig. 4) are too large to draw conclusions. As expected, estimates of *DET* show clear distinctions between simulations with different reinforcement factors, with higher reinforcement factors resulting in higher *DET*. For the most part, lemurs exhibited higher *DET* than species-relevant simulations with a reinforcement factor of one, but less *DET* than species-relevant simulations with a reinforcement factor of 2. The exceptions are vervets, for which about 50% of posterior *DET* estimates fall within or below the posterior distribution of Learning Factor 1 estimates, and aye ayes, for which more than 50% of the posterior *DET* estimates fall above the posterior distribution of Learning Factor 2 estimates.

### **Pentagon and Z-array**

The posterior distributions of estimated *DET* scores for vervets in both the pentagon array and the Z-array were entirely above those for simulated reinforcement learning agents (see online supplement “DET Analysis”). Japanese macaques, which were only tested in the Z-array, exhibited *DET* scores in this array lower than those of the vervets. The posterior distribution of estimated *DET* scores for Japanese macaques in the Z-array fully contained that of reinforcement learning agents with a reinforcement factor of 2.

## **Discussion**

Our results show several key takeaways: (1) Primate movement decisions in our multi-destination arrays were more consistent than would be expected given random transitions with probabilities proportional to distance between targets (i.e. Learning Factor 1); (2) Wild vervets and Japanese macaques, with little to no experience, navigated our arrays with far less travel than captive strepsirrhines and chose paths close to optimal; (3)

DISCUSSION

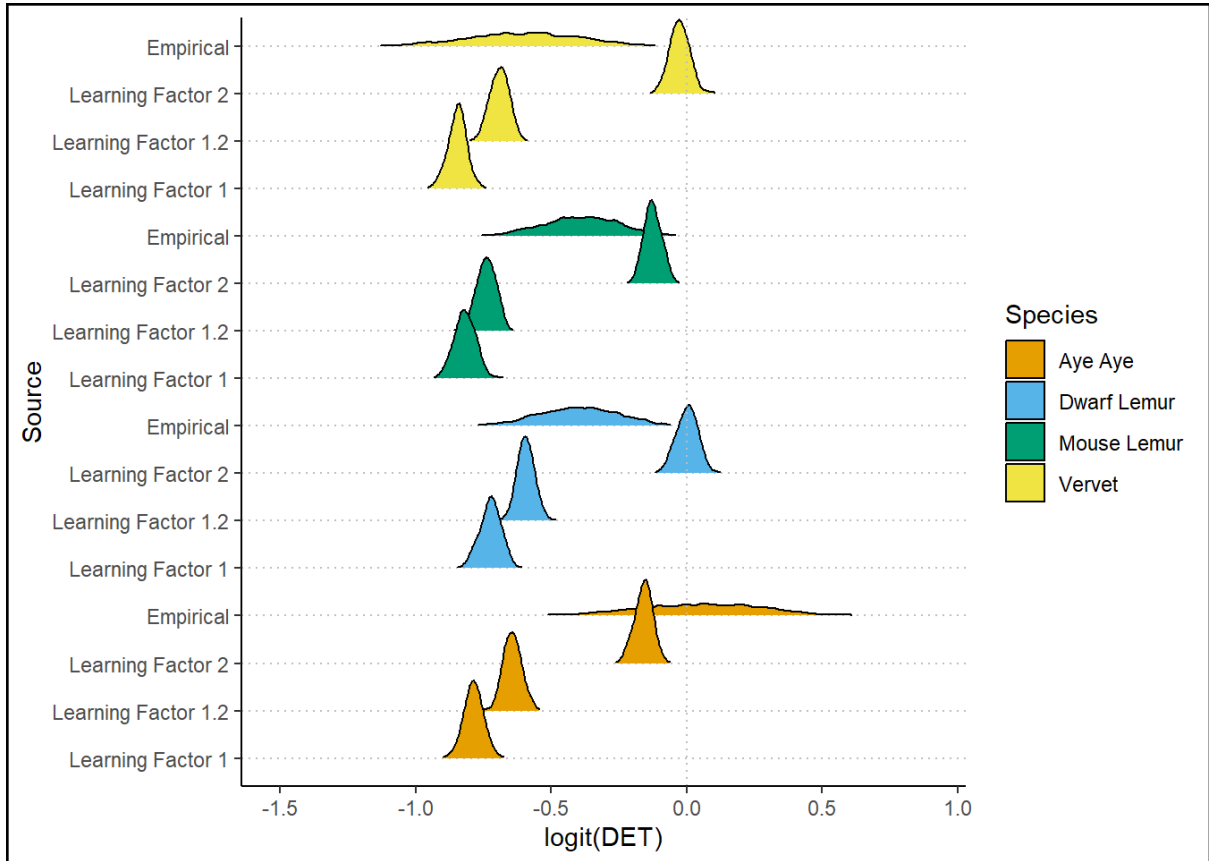


Figure 1.4: Posterior distributions of alpha estimates. Estimates of alpha (related to  $DET$  through a logit link-function) are calculated by summing the mean intercept ( $\alpha_0$ ) with relevant coefficients for species, source, and interaction. Colored plots represent the density of alpha estimates in the posterior for each possible combination of these coefficients, arrayed along the y-axis. The y-axis labels denote the source and plot color denotes the species. The scale of the density axis (height) is not shown, but consistent across all plots.

Captive strepsirrhines exhibit reduced travel distance with more experience, showing improvement rates that are well fit by traditional statistical models of learning rates; (4) These improvement rates correlate with the degree of frugivory in strepsirrhines—being greatest in dwarf lemurs, followed by aye-ayes, and then mouse lemurs. However there is too much uncertainty in our estimates of learning rates to make definitive conclusions; (5) The iterative-reinforcement algorithm proposed by Reynolds et al. (2013) to explain traplining patterns in bumblebees is not sufficient to explain route-finding patterns in the primates we tested, which are generally not central place foragers.

Our hypothesis that primates’ approximate solutions to multi-destination routes by reinforcement learning of location-based decisions was supported most strongly in our strepsirrhine sample. We found strong evidence for improvement with experience in dwarf lemurs and aye-ayes, and notable improvement in mouse lemurs. Vervets finished the Double Trapezoid array faster than the strepsirrhines but the credible intervals for their learning rates contained zero and trended positive (increased distance travelled with experience). Initial performance by vervets was so close to optimal that we may not have observed reinforcement-learning of location-based decisions even if it did occur. Thus, we can rule out the null hypothesis that strepsirrhines do not improve their performance with experience but cannot do the same for vervets. Our data suggest that the improved performance in strepsirrhines results from reinforcement learning of location-based decisions. There are, however, alternative explanations, such as strepsirrhines becoming more motivated as they grew comfortable with the testing environment, thus exploring less and reducing travel distances. Additionally, because the strepsirrhines were raised in captivity, they may simply not have learned spatial heuristics potentially used by their wild counterparts.

Analysis of our simulations also suggests that if strepsirrhines are improving perfor-



## *DISCUSSION*

mance through reinforcement learning, this process is not well modeled by the iterative-reinforcement algorithm of Reynolds et al. (2013). At reinforcement factors that produced learning rates comparable to strepsirrhines, this algorithm also yielded very high inter-individual variation in learning rates and learning curves that were not well fit by traditionally used power law functions. Strepsirrhine learning curves were comparably consistent between individuals and well-fit by power law function. Thus, strepsirrhines may use search strategies that balance exploration with the identification and exploitation of efficient navigation strategies. The simulations, conversely, depend on stochastic transitions to find an efficient route and then quickly exploit that route, to the exclusion of other possibilities.

Our second alternative hypothesis—that primates have either innate or previously learned heuristics that can be generalized to approximate solutions to multi-destination routes in many different arrays—was most strongly supported by our cercopithecoid sample. We found that vervets and Japanese macaques did not significantly reduce their distance traveled over time in any array. Their observed distance traveled was significantly lower than in simulations of the other strategies; they immediately started with efficient routes and maintained them. This evidence suggests that vervets and Japanese macaques depended on previously learned or innate spatial heuristics from the beginning of their experience within the arrays and did not deviate from these. This is reinforced by the finding that vervets in the Z-array strongly relied on the nearest-neighbour heuristic, even though it was not consistent with the shortest possible path (Teichroeb and Smeltzer 2018). Japanese macaques also showed a reliance on heuristics in the Z-array that did not lead to the shortest paths (Joyce 2021). These findings are not surprising—previous work has found that other primates frequently choose routes that are more efficient than chance but less efficient than optimal (Janson 2014). Optimal routes were frequently

observed only in the relatively simple Pentagon array, but less frequently in the more complex Z-array or Double Trapezoid array.

Though our learning simulations eventually reached an optimal path in the Pentagon array (within 70–100 trials), the vervets achieved this optimal path often on the first trial—much quicker than our model simulations. Use of either the convex-hull or the nearest-neighbour rule leads to the shortest path. Thus, the simple heuristics that vervets appear to apply in this array are far more efficient than a reinforcement-learning based approach. This suggests that under certain conditions, these “fast and frugal” heuristics are more adaptive and efficient than the cognitively more demanding learning-based alternative (Gigerenzer and Todd 1999).

## Traplining in primates

Within the first 10 trials in the Double Trapezoid Array, individuals from all species exhibited DET values that were greater than those calculated via simulations with random transitions between targets (Learning Factor 1). This evidence of repeated sequences of site visitation supports previous suggestions that primate food site visitation sequences were reminiscent of traplining (Garber 1988, Janson 1998, Dew and Wright 1998, Di Fiore and Suarez 2007, Noser and Byrne 2010), though these sequences occur in a less stereotyped manner than well-established traplining animals (e.g., bees). The wide posterior distributions of empirical *DET* estimates likely reflect both our small sample sizes and the likely possibility that primates vary between trials in their tendency to repeat paths. Additional studies with more animals could effectively model the effects of experience on *DET* and potentially find narrow estimates that can be more usefully compared to simulations with different learning factors.

The performance of our iterative reinforcement models were also array-dependent. In all

## *DISCUSSION*

arrays, a reinforcement factor of 2 produced strong decreases in distance over trial number, while a lower learning factor of 1.2 produced a strong decrease only in the Pentagon and a minor decrease in the Double Trapezoid. This again reflects the simplicity of the Pentagon array relative to the Double Trapezoid and Z-arrays. Reynolds et al. (2013) also utilized an iterative-reinforcement algorithm to analyze traplining behavior across a variety of arrays, finding fast and convergent learning rates between their algorithm and bumblebees in a pentagon array, and noting the failure of the algorithm in other arrays. By adding a no-backtracking rule to this algorithm, we were able to successfully simulate improvements in array navigation, but these improvements were not well modeled by a power-law learning curve, as often seen in animal learning (Donner and Hardy 2015). In future work, more complex learning models may be able to more accurately reflect the patterns of exploration and improvement that animals use to efficiently exploit resources in environments with complex structure. New ways of measuring patterns of exploration and improvement may complement DET metrics by revealing different aspects of animal decision making.

## **Evidence of a phylogenetic signal in navigational behaviour, ecological determinants, or a wild vs. captive difference?**

It is intriguing that the navigational behaviours of the strepsirrhines in our dataset were closely aligned, while the cercopithecoids were also similar to one another. This suggests that the cognitive skills underlying navigational strategies in primates have some unknown level of phylogenetic signal, consistent with other behavioural traits (e.g., social organization (Ossi and Kamilar 2006); activity pattern (Wu et al. 2017)). However, despite any potential constraints in navigational abilities introduced via evolutionary history, some behavioural variability may still be introduced via specialized adaptations

for ecological niches.

The strepsirrhine species we studied all share a distinct behavioural trait that separates them from vervets and macaques—they are primarily solitary in contrast to the highly social cercopithecoids. This is notable because recent work suggests that spatial cognition brain regions are expanded in strepsirrhines and solitary primates, potentially because animals ranging alone require greater spatial memory to find food and mates (DeCasien and Higham 2019). Conversely, social species rely on each other for increased accessibility to food, including detection and defense. Social species may benefit more from improved sensory perception instead of enhanced spatial abilities, which would be more useful for distinguishing between resources. Our findings may provide some preliminary support for a trade-off between visual processing centers and spatial abilities in social primates, with reduced spatial learning in cercopithecoids and enhanced spatial learning in strepsirrhines. Future research with wild animals under natural conditions is needed to provide further support for a trade-off between spatial abilities and sensory perceptions in social versus solitary species. However, although the strepsirrhines aligned more closely when compared to the cercopithecoids, we found differences between the strepsirrhines which may be caused by their different ecological niches.

Ecological and dietary variability is great within primates, so it is highly unlikely that most species will rely on the same or similar cognitive mechanisms during navigation. Among the strepsirrhines in our dataset, the most frugivorous species (i.e., dwarf lemurs) had the lowest overall credible intervals for initial travel paths and learning curves. While this pattern fits our prediction that more frugivorous animals would adopt strategies to efficiently navigate arrays, the degree of overlap in parameter estimates for the strepsirrhines prevents us from making any strong conclusions. Although dwarf lemurs align more closely with the other strepsirrhines with regards to initial performance and faster

## *DISCUSSION*

learning, we know from previous work that they do appear to utilize some navigational heuristics more often than either aye-ayes or mouse lemurs (Teichroeb and Vining 2019), similar to cercopithecoids. Heuristic use in dwarf lemurs is also weakly supported by our finding that they traveled shorter distances initially relative to other strepsirrhines in the same array. Dwarf lemurs thus show some evidence for both reinforcement-based learning of navigational paths and use of heuristics and may be utilizing both strategies. This could be because dwarf lemurs face the added constraint of spending up to seven months at a time in torpor, where they rely on accumulated fat for subsistence Fietz and Ganzhorn (1999). Thus, they have limited time out of torpor to take in resources and build up their reserves, and likely need to do this as efficiently as possible. Other species heavily reliant on replenishing resources have also been reported to use multiple navigational strategies, for example pollinating bees show iterative learning while also linking nearest-neighbour flowers<sup>5</sup>. Our findings for dwarf lemurs suggest that the tendency to combine navigation strategies may be motivated by a diet of spatially stable, replenishing food sources. These findings support our prediction that dwarf lemurs would show shorter distances with experience and increased determinism.

Our findings of similar navigational strategies for strepsirrhines versus cercopithecoids also conform to the different conditions experienced by our sample (i.e., the strepsirrhines were captive and the monkeys were wild). Many variables could potentially alter experimental outcomes, such as internal factors that influence motivation (Janmaat 2019). Wild monkeys were tested opportunistically at different times during the day and some individuals ran more trials consecutively than others, which may have led to satiation and performance fatigue. We also cannot consider unknown variables such as complete diet and health status. For example, macaques were provision fed whereas vervets were not. Provision feeding in particular may encourage repeated sequences of visits to baited

sites. Further, wild animals experienced distractions including alarm calls, vigilance for predators, and changes in weather conditions. The strepsirrhines in our sample were all raised in captivity; environments much smaller than their natural territories that lack the ecological pressures they evolved with. This may have kept them from learning heuristics they might otherwise use Trapanese et al. (2018). These species are also all nocturnal and may face different perceptual constraints. We cannot be certain of the effect that these or other variables may have had on our results. Exploring whether captive cercopithecoids also exhibit first-trial heuristic use in multi-destination routes is an important future direction. Likewise, investigating a larger sample that includes species with overlapping behavioural traits, such as social strepsirrhines (e.g., ring-tailed lemurs), or strepsirrhines with varying activity patterns (e.g., diurnal or cathemeral), and comparing with our results would be useful future studies.

## Ethics

This work was approved by the Uganda Wildlife Authority (UWA/COD/96/02), the Uganda National Council for Science and Technology (NS 537), the University of Toronto Animal Care Committee (#20011416), the Duke Institution Animal Care and Use Committee (A290-14-12), and the Concordia University Animal Research Ethics Committee (#30009663). These experiments were conducted according to ARRIVE guidelines and were performed in accordance with all institutional and/or national guidelines and regulations.

## *DATA AVAILABILITY*

### **Data availability**

The raw data are included in the manuscript via a public link: <https://github.com/aqvinig/Do-Primates-Trapline>.

## Chapter 2

# The Movement and Behavior of Kinkajous in a Dynamic Foraging Hotspot

Authors: Alexander Q. Vining<sup>1,2,3,4</sup>, Chase L. Núñez<sup>2,3</sup>, Nele Stockmeyer<sup>3</sup>, Margaret C. Crofoot<sup>2,3,4,5,6</sup>

1. Animal Behavior Graduate Group, University of California, Davis, USA
2. Department for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior, Konstanz, Germany
3. Department of Biology, University of Konstanz, Konstanz, Germany
4. Smithsonian Tropical Research Institute, Balboa, Republic of Panamá
5. Center for the Advanced Study of Collective Behavior, University of Konstanz, Konstanz, Germany
6. Department of Anthropology, University of California, Davis, CA, United States



## ABSTRACT

### Abstract

Within the *Procyonidae* family, three lineages (kinkajous, olingos, and ringtails) have independently evolved to become highly arboreal, nocturnal, and fruit dependent (frugivorous) - despite their membership in the Order *Carnivora*. Reliance on fruit is hypothesized to select for behavioral flexibility mediated by cognitive mechanisms such as episodic memory and topological mental mapping. This hypothesis arises because the semi-predictable patterns of fruit production is patchy, and requires extensive environmental knowledge to exploit effectively. Little is known, about how the shift to frugivory has affected the cognitive evolution and foraging patterns of Procyonids, however. In this study, we analyze the movement and behavior of the most strictly frugivorous of the *Procyonidae*, kinkajous (*Potos flavus*), within a large balsa (*Pyramidale ochroma*) crown. Balsa trees produce large flowers that open for one night, generate large quantities of nectar, and then fall off the tree - creating a high value resource patch with internal structure that changes every night. Our goal was to determine the contexts in which kinkajou movements responded to the distribution of balsa flowers, and the types of learning and memory necessary for such a response to be possible. We find that patterns of kinkajou movements between major regions of this tree crown did not vary with the changing distribution of flowers, but kinkajou movements do reflect a tendency to sample different regions and stay longer where flowers are denser. We also find that kinkajous frequently co-forage within the tree, showing increased levels of movement when other kinkajous are present. These results suggest that foraging hotspots may act as social hubs for otherwise solitary kinkajous - despite competitive dynamics. We discuss the implications of these findings for understanding kinkajou cognition at different ecological scales.

## Introduction

Intelligence, by one definition, is the ability to accurately predict the state of the world at some future point in time. Organisms are often deemed more intelligent when they anticipate conditions and outcomes further into the future, or by using knowledge of multiple, interacting processes (Decker et al. 2016, Redshaw and Suddendorf 2020). Social intelligence is often differentiated from ecological intelligence, with ongoing debate as to the consequences natural selection in either domain has on other types of cognitive performance (Rosati 2017, Ashton et al. 2018). Ecological and social hypotheses for the origins of animal intelligence both posit that the ability to understand the nature of multiple, interacting relationships provides fitness benefits to an animal. In the case of ecological intelligence, the relationships in question are physical rather than social (e.g. the spatial relations between patches, the temporal relations between productive periods, the causal relations between rainfall and phenology). Understanding such environmental relationships provides animals a means of directing their movements between highly productive patches, which is more energy efficient than random search foraging and presumed to carry fitness benefits when patches are sparsely distributed (Boyer and Walsh 2010).

Fruit is one type of resource that is often distributed this way, leading to the hypothesis that more selective foragers (i.e. those that eschew leaves and other hard-to-digest, easy-to-find energy sources in favor of high-energy fruits) have evolved greater intelligence (Milton 1981). This hypothesis is supported by phylogenetic analysis of primate and bird brain sizes with regard to degree of frugivory (DeCasien et al. 2017, Hardie and Cooney 2023, but see also Powell et al. 2017) and by the performance of both primates and birds that rely on sparsely distributed food patches on multiple cognitive tests (Mals-

## INTRODUCTION

burg et al. 2020). Neither of these approaches, however, directly measures an animal's understanding of environmental relationships; they provide only indirect support of the ecological intelligence hypothesis.

Increasingly, behavioral ecologists are using technological advances in animal telemetry and remote environmental sensing to analyze animal movements across different types of landscapes, such as energy, risk, and information (Williams and Safi 2021). This approach allows animal decision making to be studied as a sequential process that is organized hierarchically across multiple spatial and temporal scales - context that is missing from standard laboratory studies of cognition (Stephens 2008). However, determining the relevant scales of animal decision making in natural contexts and collecting appropriate multi-scale data remains a challenging and rarely achieved goal (Kay et al. 2017). Recent examples have highlighted species that both live in complex social groups and depend on dynamic resource distributions like fruit or watering holes (Presotto et al. 2019, elephants, Toledo et al. 2020, fruit bats, Trapanese et al. 2022, macaques and capuchins). There is an opportunity for studies of ecological decision making to enhance our understanding of the relationship between ecological and social intelligence by applying this approach to species that selectively forage on sparsely distributed, high yield resources, but do not live in large, complex social groups.

The *Carnivora* Order is a compelling place to begin this work, given that social carnivores exhibit the same markers of intelligence as social primates, but do not meaningfully diverge from solitary carnivores in terms of relative brain size (Holekamp and Benson-Amram 2017). Kinkajous (*Potos flavus*) are particularly interesting in this context, as they have evolved over the last 24 million years (approximately) from an omnivorous, terrestrial ancestor into a highly specialized arboreal frugivore (Koepfli et al. 2007). We thus predict that they are skilled learners of environmental relationships. Because

kinkajous live in small social groups and typically forage alone (Kays and Gittleman 2001), comparing their understanding of environmental relationships to that of more social animals will help distinguish the roles of diet and sociality in cognitive evolution.

Kinkajous often consume the nectar of balsa flowers (*Pyramidale ochroma*), which are clustered across multiple spatial scales and thus provide an excellent system for studying kinkajou decision making. Balsa flowers typically open around 17:00 hours with already-high volumes of nectar, continue producing diminishing volumes of nectar through the night, and fall off the tree with a day or two (Kays et al. 2012). Thus, within tree crowns, the concentration of nectar in patches of large flowers is analagous to that of the inflorescences exploited by bees and hummingbirds; a system that has been used extensively to study learning mechanisms in these species (van der Kooi et al. 2021). At larger scales, balsa trees are wind-dispersed pioneer specialists that grow best in forest gaps and disturbed areas, resulting in highly patchy spatial distributions (Aikman 1955). In this study, we begin the work of investigating the degree to which kinkajou movement decisions reflect an understanding of the environmental relationships that define the distribution of balsa nectar.

## Kinkajou Socio-Ecology

Kinkajous are highly adapted to a diet of fruit and nectar, which make up 90-100% of a kinkajou's diet (Julien-Laferrière 1999, Kays 1999). Like platyrrhines (New World monkeys), kinkajous have moderate body sizes (2-5 kilograms), prehensile tails, grasping paws, alternating diagonal gaits, and visual sensitivity to brightness contrasts (Chausseil 1992), all of which aid in finding and accessing hard-to-reach fruit in the canopy. Compared to their *Carnivoran* relatives, kinkajous also have shortened snouts, flattened molar dentition, and squared jaws for processing fruit (Ford and Hoffmann 1988). Kinkajous

## INTRODUCTION

do not appear to specialize on any one food source, however; they exploit a rich variety of fruit and nectar sources across their broad range (much of Central and South America) (Julien-Laferrière 1999, Kays 1999, Nascimento et al. 2017), with some evidence of population differences in food preferences (Séguigne et al. 2022). The selectivity and diversity of kinkajous' diets suggests they are promising candidates for the study of ecological intelligence, but as of yet there exists little data that can elucidate their cognitive flexibility or predictive ability.

Another advantage of studying the foraging of kinkajous is that they are largely solitary (Julien-Laferriere 1993). Social groups consisting of one adult female, one or two adult males, and two juveniles have been found associating at sleep sites and in large foraging trees during up to 10% of their active hours, but kinkajous nearly always travel alone (Kays and Gittleman 2001). Kinkajous are, however, known to be quite vocal and frequently scent-mark using specialized glands on their cheeks, throat, and abdomen (Ford and Hoffmann 1988, Kays 1999). It is possible that kinkajous maintain nuanced social relationships with their neighbors mediated through channels of communication that remain invisible to us. Further uncovering the social relationships of kinkajous remains an intriguing line of research. At the same time, the solitary nature of kinkajous' travel between foraging locations suggests decisions about where to forage are made largely independently, meaning it is safer (relative to primates) to assume kinkajous travel where they want to go, and not where other individuals want to go. Kinkajous thus make an excellent species for studying decision making in dynamic, challenging foraging environments.

Recent work revealed that kinkajous' transits between foraging locations tend to exhibit higher "pathiness" (i.e. occur along the same route) than those of other sympatric frugivores (Alavi et al. 2022). Kinkajous' nocturnality may be a contributing factor

here; while diurnal frugivores can use distant landmarks to orient toward foraging locations, kinkajous are likely limited perceptually to local landmarks and thus become more dependent on learned routes. Determining the extent to which kinkajous use their route-networks flexibly to plan visits to highly productive trees, rather than relying on fixed patterns of visitation between foraging sites, will require the study of kinkajou movements relative to the changes in resource distribution in their home ranges. A complete picture of kinkajou decision making requires this work to be done at multiple scales, such that the decision making heuristics used to navigate local, perceptually available resource distributions can be used to contextualize decisions made at home-range scales.

## The Kinkajou-Balsa Mutualism

Balsa trees (*Pyramidale ochroma*) are an important food resource for kinkajous, and one that provides opportunities to study kinkajou decision making at multiple scales. During their flowering period, each of a balsa tree's large flowers can produce around 22mL of nectar a night (Kays et al. 2012). These flowers open around dusk and fall off 24-48 hours later. The timing and quantity of nectar production suggest balsas are specifically soliciting larger bodied, nocturnal pollinators (Lynn Carpenter 1978). Indeed, over 126 nights of observation across three balsa trees, 25 mammalian and bird species were observed visiting. Kinkajous were the most frequently observed species (Kays et al. 2012).

Whether kinkajous and balsa have specifically co-adapted to each other is not known, but we put forth a few hypotheses for why this could be a likely scenario. From the kinkajous' perspective, balsas in lowland tropical rainforests tend to flower during the transition from wet season to dry season, a period on Barro Colorado Island when fruit availability is typically very low (Foster 1982). Thus, balsa nectar may be an essential

## *INTRODUCTION*

“fallback food”, sustaining kinkajous when their primary food source is scarce. From the balsa’s point of view, kinkajous may be especially effective for cross-pollinating between balsa stands. Kinkajous use their long tongues to drink nectar from the flowers, streaking pollen across their faces in the process. Additionally, balsa most often occur in isolated patches where treefalls or other disturbances have occurred ([Aikman 1955](#)); therefore, cross-pollination between these patches requires that pollinators travel long-distances over short period of times. Pollination efficiency would be increased if pollinators moved between stands along direct routes. Pollination efficiency could be further increased if pollinators selectively moved between actively flowering crowns by making flexible, predictive foraging decisions, as we predict that kinkajous do.

Finally, the importance of balsa trees for kinkajous may extend beyond foraging gains. As hotspots of activity, balsa trees may also serve as important social hubs for kinkajous. A clearer picture of kinkajous’ social relationships, and the impacts these may have on their movements, can be gained by observing the nature of kinkajous’ interactions during co-foraging.

## **Defining and Measuring Intelligence Through Foraging Decisions**

Foraging presents animals with a hierarchy of decisions about when to forage, where, and for how long. These decisions require an animal to balance what it can perceive in its local environment, the properties of the environment that it has learned, and the consequences of each decision on future possibilities ([Stephens 2008](#)). Animals are expected to make decisions that minimize their risk of acquiring insufficient energy (often best achieved by maximizing energy intake), but are limited in their ability to make optimal decisions by what they know about the properties of their environment ([Pyke 1984](#)). Thus, the divergence of animal decisions from optimal decisions can reveal what the animal has

learned.

Within individual balsa crowns, the nightly turnover of flowers creates a sort of natural experiment, in which optimal paths through the crown vary each night. Thus, we are able to test whether kinkajous attend to changing resource distributions at a local scale, and flexibly navigate the tree crown by changing their travel route to optimize their foraging success on a night to night basis. Within the balsa crown, it is possible that kinkajous (a) assess flower distributions perceptually (likely by smell), (b) utilize learned behavioral routines to efficiently exploit average flowering patterns, and/or (c) monitor changing flower distributions - using recent experience and knowledge of relationships in flowering patterns between regions of the tree to selectively exploit areas predicted to be most productive. Testing between these possibilities is therefore a good way to begin understanding kinkajou decision making.

## Hypotheses and Predictions

In this study, we conducted observations of a single balsa tree. We collected data on the distribution of flowers within the tree crown each night, and minute-by-minute data on the distribution of kinkajous in the crown over four hour periods. By pairing these data, we aimed to test several hypotheses regarding the ways kinkajous respond to changing resource distributions at a local scale (Table 1).

First, under the hypothesis that animals will selectively forage in areas with above-average energy returns ([Charnov 1976](#)), we predicted that kinkajous would spend more time in regions of the tree crown with greater flower density. We then tested two alternative hypotheses of how kinkajous decide to move between different regions of the tree against a null hypothesis of random movement: A) kinkajous assess, perceptually, which regions have the greatest density of flowers and B) kinkajous predict which regions of the tree



## *METHODS*

will have the greatest density of flowers using information obtained on previous nights.

Next, we considered the role of social bonds on kinkajou behavior while foraging in the balsa crown. We tested the additional hypothesis that kinkajous avoid direct foraging competition by foraging in separate regions of the tree against the hypothesis that kinkajous use time in the balsa tree for physical means of social bonding such as grooming.

Lastly, we considered how working memory might influence kinkajou foraging decisions, hypothesizing that kinkajous will avoid foraging in regions they have recently visited (and likely depleted).

## **Methods**

### **Study Site**

Data were collected on Barro Colorado Island (BCI), Panama. BCI is a 15.6 km<sup>2</sup> island with semi-deciduous lowland tropical forest. The wet and dry seasons on BCI are very distinct, where approximately 90% of an average annual 2600mm rainfall occurs between May and December ([Dietrich 1982](#)). We observed a single balsa tree beginning December 14, 2019 until January 26, 2020. This is a period of transition from the wet to dry season during which there is little fruit production on BCI ([Wright et al. 1999](#)). The tree under observation grows from the forest edge over a clearing made for a residential building within the Smithsonian Tropical Research Institute Biological Research Station. As a result, this tree crown has grown very large and is the only tree crown connecting two forest segments. Consequently, there is direct visual access to the tree crown from above, below, and from three sides. A steep hill to the west of the tree provides an eye-level view into the tree crown.

CHAPTER 2. MOVEMENT IN HOTSPOT

Table 2.1: Hypotheses and Predictions

Hypothesis	Prediction	Statistical Support
1) Kinkajous prefer to forage in regions with greater flower density	Kinkajous spend more time in the tree when there are more flowers	Negative
1) Kinkajous prefer to forage in regions with greater flower density	Kinkajous spend more time in regions of the crown that have greater flower density	Weak
2A) Kinkajous perceptually assess the density of flowers across the tree crown when deciding where to go	When moving between regions, kinkajous tend to select the region with greater flower density	None
2A) Kinkajous perceptually assess the density of flowers across the tree crown when deciding where to go	The difference in flower density between regions is positively associated with the probability a kinkajou moves into the higher density region	None
2B) Kinkajous predict which regions will have more flowers using information obtained from foraging on previous nights	When moving between regions, kinkajous tend to select the region that has historically had greater flower density	None
2B) Kinkajous predict which regions will have more flowers using information obtained from foraging on previous nights	The difference in average previous flower density between regions is positively associated with the probability that a kinkajou moves into the higher density region	None
3) While in the tree, kinkajous avoid direct foraging competition	Kinkajous are more likely to depart a tree region when there are other kinkajous in that region	Strong
3) While in the tree, kinkajous avoid direct foraging competition	When moving between regions, kinkajous are more likely to enter the region with the fewest other kinkajous	None
4) Kinkajous remember where they have been, and avoid foraging in recently depleted regions	The order of minimum conditional entropy in sequences of region visits is $> 1$	Weak

## *METHODS*

### **Behavioral Observations and Thermal Videography**

We conducted behavioral observations and thermal videography for four hour periods over most nights of the study period. Prior to the observation, the tree architecture was diagrammed and each limb out to three branches (forks) from the trunk was given a unique label. During observations, we audio dictated branch-to-branch movements of kinkajous and opportunistically noted other interesting behaviors using the Animal Observer app. In addition, scan samples were recorded in the app every 15 minutes, with data taken on the behavior of each kinkajou in the tree (ethogram: resting, moving, foraging, not visible) and the branch on which it resided.

During each observation period, we used FLIR t1020 cameras to collect thermal video from two perspectives: an East-facing angle into the tree crown from an adjacent hilltop (Figure 1A), and an upward facing angle from below the tree crown. Because it was not possible for a single observer to live-track the movements of all kinkajous in the tree crown, these videos were used to collect data on kinkajou locations within the tree at every minute. However, compression of the three dimensional crown area into two dimensions often made the identification of specific branches impossible. Thus, we used the diagram of the tree architecture to hand-draw three semi-distinct regions (Front, Top, and Back; labelled 1, 2, and 3, respectively) that could be distinguished from both a side view (thermal video, Figure 1A) and an aerial view (drone images of flowers, Figure 1B). Thermal videos were converted from raw .seq format to .mp4 files using FLIR Thermal Studio Pro. All videos were initially converted using a scale minimum of 20.0° C, a scale maximum of 28.0° C, and an isotherm at 28.3° C. In most cases, these settings highlighted kinkajous in green, against greyscale thermal imagery; conversion settings were adjusted for individual videos to achieve this effect where conditions required. Due

to small discrepancies in the frame-rate recorded in video metadata and the actual number of frames collected per second, video conversion resulted in time dilation/expansion that over a full four hour period could add up to several minutes. To account for the discrepancy between converted video length and video length in the raw metadata, four second breaks were inserted at regular intervals of the videos using iMOVIE. This prevented the synchronization of the two video feeds, and only the side view of the full tree was used for further analysis. We then used the ProCreate software to trace region outlines on the first frame of every video, and iMovie to overlay these traces across the full videos. We used the resulting videos to record how many kinkajous were in each region




Figure 2.1: (A) Still frame from thermal video recorded on December 30, 2019. Colored lines denote the three regions, and have been redrawn over the original image for visual clarity in print. An isotherm was used during video conversion to set recorded temperatures above 84 F to green, making kinkajous easily identifiable, as with the kinkajou visible inside Region 3 (Panel D). (B) Orthomosaic image of drone photos from December 30, 2019, with region shapes (black outlines) and geo-referenced upright-flower annotations (circles) added in QGIS. Panel C zooms in on a region of the orthomosaic with two upright flowers.

of the tree crown at each minute of the video. For parts of the videos when there was no kinkajou movement, the video was fast forwarded up to 10 times speed. When there were several kinkajous in the tree moving simultaneously, the video was slowed down to

## *METHODS*

the original speed and reviewed as necessary. Because Region 3 largely occurs within and behind Region 1 in the videos, we used our best judgement to determine when (a) a kinkajou was inside the Region 3 trace, but had not in reality left Region 1 and (b) when kinkajous were hidden from view but had not left the tree.

### **Drone Flights and Flower Counts**

Between December 27, 2019 and January 23, 2020, we piloted a DJI Phantom4 Pro drone to collect aerial photographs of the tree crown each afternoon before the flowers opened and each morning at dawn (weather permitting). Over 100 images were taken in a regular grid on each flight, then stitched into an orthomosaic image using Agisoft. Using QGIS (Version 3.2.9), we geo-referenced all orthomosaics to Google Satellite imagery of the research station and hand-traced outlines of the three regions. For each orthomosaic, we then created a geo-referenced spatial point over each visible flower, annotating whether the flower was upright or fallen over (Figure 1B). Flowers were assigned to regions using the QGIS intersect function. Flower densities were calculated by dividing upright flower counts by the area of corresponding region shapes in QGIS. On nights for which drone flights on both the preceding evening and the next morning produced complete, high quality orthomosaics, the average flower count of the two images was used for analysis.

### **Statistical Analysis**

#### **Overview**

To assess the drivers of kinkajou movement within the balsa crown, we fit hierarchical models examining 1) nightly kinkajou residence times within each region, 2) kinkajou

transitions between regions based on flower densities, and 3) kinkajou transitions between regions based on the presence of other kinkajous. To create a useful reference for understanding how kinkajou behavior deviates from what could be expected if an animal moved at random between regions, we simulated random transitions and fit the same models to the simulated data. Finally, to understand how memory influences kinkajou movement through the tree crown, we extracted sequences of movement between regions during periods when only one kinkajou was in the tree (and thus we could be confident it was the same kinkajou) and analyzed the Shannon’s Entropy and 1st-4th order conditional entropies of the sequences according to protocols described in Riote-Lambert et al. (2017). These measures tell us how random a kinkajou’s transitions appear given information about some number of previous locations it has been. If gaining information about more previous locations makes transitions appear less random, it suggests the kinkajous remember where they have been and modulate their behavior in response. The posterior distributions of key regression parameters, based on our predictions, were compared to null expectations (either 0, or the posterior distribution of the same parameter when the model was fit to simulated data).

All analyses were conducted in R, using the package “rethinking” (McElreath 2020) to interface with STAN (Carpenter et al. 2017). Each model was fitted using four Markov Chain Monte Carlos (MCMCs) of 3000 iterations, with the first 1500 iterations discarded. We used effective sample size and  $\hat{R}$  calculations alongside visual assessment to determine whether MCMCs converged and parameter estimates were adequately mixed.

### Reference Model Simulations

To serve as a reference model for our analysis, we simulated kinkajou movements between tree regions using a first order Markov process with transition probabilities derived from

## METHODS

the rate of kinkajou transitions between regions across the whole study. We maintained the overall timing of kinkajou visits by initializing agents entering the tree at the same times (and to the same regions) as kinkajous in the empirical data. We simulated this process 100 times and use the resulting data throughout subsequent analyses as a reference for what kinkajou movements would look like if they were based purely on fixed, per-minute transition rates and had no dependence on previous behavior, flower counts, or other kinkajous.

### Comparing Kinkajou Residence Time and Regional Flower Densities

To examine the relationship between regional flower densities and kinkajous' residence times, we fit a hierarchical, zero-augmented gamma model (a type of hurdle model) to the data,

$$Residence \sim dZAGamma(p, \mu, \Theta)$$

$$\text{logit}(p) \sim p_0 + p_R + \beta_1\Psi + \beta_2\psi_R + p_t$$

$$\text{log}(\mu) \sim \mu_0 + \mu_R + \beta_3\Psi + \beta_4\psi_R + \mu_t$$

$$p_0, \mu_0 \sim N(0, 2)$$

$$\beta_1, \beta_2, \beta_3, \beta_4 \sim N(0, 2)$$

$$p_R \sim N(0, \sigma_{p_R})$$

$$\mu_R \sim N(0, \sigma_{\mu_R})$$

$$p_t \sim N(0, \sigma_{p_t})$$

$$\mu_t \sim N(0, \sigma_{\mu_t})$$

$$\Theta, \sigma_{p_R}, \sigma_{\mu_R}, \sigma_{p_t}, \sigma_{\mu_t} \sim \text{exp}(1)$$

where *Residence* is the predicted number of minutes kinkajous spent in a tree region per hour of observation on a given night,  $p$  is the probability of a kinkajou entering the region on that night,  $\mu$  is the average residence time per hour given that at least one kinkajou enters the region,  $\Theta$  is the scale parameter of a gamma distribution,  $p_0$  is a grand mean for the intercept of  $p$ ,  $p_R$  gives expected change in  $p$  given the region of interest,  $\Psi$  gives the density of flowers in the full tree crown on the given night,  $\psi_R$  gives the difference between the density of flowers in the region of interest and the density of flowers in the full tree,  $p_t$  and  $\mu_t$  give the effect of the given night on  $p$  and  $\mu$ ,  $\sigma$  parameters give the variance of fixed effects on  $p$  and  $\mu$  across regions ( $R$ ) and nights ( $t$ ), and  $\beta$  coefficients are linear estimators of corresponding flower densities on  $p$  and  $\mu$ . `dZAGamma` is a density function resulting from a mixture of the Bernoulli distribution and the gamma distribution such that  $p$  gives the probability of observing a 0, and the probability of a non-zero value  $x$  is  $(1 - p) * dgamma(x, \mu/\Theta, \Theta)$

We use this model to test the predictions of Hypothesis 1 (Kinkajous forage longer in regions with greater flower density).

### **Kinkajou Decision Making**

While the model of kinkajou residence time, presented above, reveals how kinkajous respond to regional flower densities, it can not elucidate how kinkajous assess multiple regions at once and make decisions. For this purpose, we analyzed the scans where the distribution of kinkajous in the tree crown changed. Because we predicted that kinkajous would choose regions with greater density, and the movement from one of the three regions to another presents a simple binary choice, we encoded each region transition of a kinkajou as either a success (moved to the available region with greater flower density) or a failure (moved to the available region with lesser flower density). We then used



## METHODS

Bayesian estimation to fit a hierarchical model of success to the data (Equation 2)

$$Success \sim \text{binomial}(1, p)$$

$$\text{logit}(p) \sim p_0 + p_{From} + p_{Pref} + p_I + \beta_5 \Delta_\psi$$

$$p_0, \beta_5 \sim N(0, 2)$$

$$p_{From}, \sim N(0, \sigma_{From})$$

$$p_{Pref}, \sim N(0, \sigma_{Pref})$$

$$p_I, \sim N(0, \sigma_I)$$

$$\sigma_{From}, \sigma_{Pref}, \sigma_I \sim \text{exp}(1)$$

where  $p$  is the probability that a particular transition is a success,  $p_0$  is a grand mean for the intercept of  $p$ ,  $p_{From}$  is a conditional adjustment on the probability of success based on the region in which the kinkajou started,  $p_{Pref}$  is a further adjustment to the intercept based on the unoccupied region with the higher flower density,  $p_I$  is an interaction effect given the combination of  $p_{From}$  and  $p_{Pref}$ ,  $\Delta_\psi$  is the difference in flower density between the region selected and the region not selected,  $\beta_5$  is a linear regression coefficient of  $\Delta_\psi$ , and  $\sigma_{trans}$  is the variance of transition effects.

We use this model to test the predictions of Hypothesis 2A (Kinkajous perceptually assess the density of flowers across the tree crown), predicting a mean value of  $\hat{p} > 0.5$ , marginal to regional effects, and that  $\hat{\beta}_5 > 0$ .

### The Role of Memory in Decision Making

In addition to assessing the current distribution of flowers in the tree crown, kinkajous may use information obtained during foraging on previous nights to inform their movement decisions. To test this hypothesis, we fit our data to a third model, identical in

structure to Equation 2 except that  $\beta_5\Delta_\psi$  is replaced by  $\beta_6\Delta_{\bar{\psi}}$ , where  $\Delta_{\bar{\psi}}$  is the difference between (1) the average flower density of the selected region on all previous nights and (2) the average flower density of the unselected region on all previous nights.

To assess the impact of working memory of kinkajou decision making, we isolated periods during which only one kinkajou was in the tree crown, and thus sequences of transitions could be identified. We calculated the Shannon's Entropy and 1st-4th order conditional entropy of these sequences, using the adjustments described in Riotte-Lambert et al. (2017). We used the same calculations on data from Reference Model 1 to determine whether the order and magnitude of the minimum conditional entropy from the empirical sequences differ significantly from expectations given a first-order Markov transition process.

## Social Behavior

We analyzed the role of con-specific presence in the tree crown on kinkajou decision making the same way we analyzed the role of flower density. Thus, our fourth model uses the same structure as Equation 2, except that Success = 1 when a kinkajou chooses the region currently occupied by more con-specifics and  $\beta_5\Delta_\psi$  is replaced by  $\beta_7\Delta_\kappa$ , where  $\Delta_\kappa$  gives the difference in number of present conspecifics between the selected region and the unselected region.

We also explored how the distribution of kinkajous within the tree crown affects the timing of kinkajou transitions. To this end, we fit a fifth model to determine the probability of a kinkajou leaving the region it is in during any given minute. This model takes the

## RESULTS

form given in Equation 3

$$\begin{aligned} \text{Departures} &\sim \text{binomial}(n, p) \\ \text{logit}(p) &\sim p_0 + p_R + p_t + \beta_8 n + \beta_9 k \\ p_0 &\sim N(0, 2) \\ \beta_8, \beta_9 &\sim N(0, 2) \\ p_R &\sim N(0, \sigma_R) \\ p_t &\sim N(0, \sigma_t) \\ \sigma_R, \sigma_t &\sim \text{exp}(1) \end{aligned}$$

where *Departures* is the number of kinkajous that leave a region during a given minute-long scan period,  $n$  is the number of kinkajous in the given region that minute,  $p$  is the probability of a single kinkajou leaving,  $p_0$  is a baseline departure rate,  $p_R$  is a region specific adjustment to the departure rate,  $p_t$  is a night-specific adjustment to the departure rate,  $k$  is the number of kinkajou in *other* regions, and  $\sigma_R$  and  $\sigma_t$  give the variance in conditional adjustments to the intercept for region and night, respectively.

## Results

### Patterns of Flower Density and Kinkajou Residence

Prior to any statistical analysis of the relationship between the distribution of flowers and the movements of kinkajous in the tree crown, we performed some descriptive analysis of these patterns independently. The average density of visible, upright flowers in the tree crown over the course of the study was 1.45 flowers/10m<sup>2</sup> (sd = 0.64). The distribution of

flowers across the regions we studied was relatively even and consistent: Region 1 differed in flower density from the full tree crown by an average of  $-0.04$  flowers/ $10\text{m}^2$  ( $\text{sd} = 0.21$ ); Region 2 differed by an average of  $0.08$  flowers/ $10\text{m}^2$  ( $\text{sd} = 0.15$ ); Region 3 differed by an average of  $-0.09$  flowers/ $10\text{m}^2$  ( $\text{sd} = 0.28$ ). We tracked the residence time of kinkajous in the tree by measure of kinkajou-minutes (e.g. if three kinkajous were observed during a 1-minute scan, we counted three kinkajou-minutes); kinkajous, collectively, spent an average of  $25.23$  kinkajou-minutes/hour in the tree crown during our observations ( $\text{sd} = 22.92$ ). The average times spent in Regions 1, 2, and 3 were  $6.51$  minutes/hour ( $\text{sd} = 7.44$ ),  $9.51$  minutes/hour ( $\text{sd} = 9.54$ ), and  $9.21$  minutes/hour ( $\text{sd} = 10.65$ ), respectively. Flower density trended upward over the course of the study, while kinkajou residence times were highly variable night-to-night, with a drop-off at the end of the study period (Figure 2). There were no immediately apparent differences between regions in these trends.

## Comparing Kinkajou Residence Time and Regional Flower Densities

To statistically analyze the relationship between flower density and kinkajou residence time, we built a model to estimate the residence time in a given region during an observation period in terms of 1) the probability at least one kinkajou entered the region and 2) the expected residence time in the region if at least one kinkajou visited, and as a function of 1) the density of flowers in the whole tree crown during that observation period, 2) the flower density difference of the given region, relative to the full three crown, 3) the region in question, and 4) the night of the observation period (Model 1). Assessment of MCMCs used to estimate parameter values in all models revealed strong mixing and convergence, meaning we are reasonably confident that estimates of our model parameters

## RESULTS

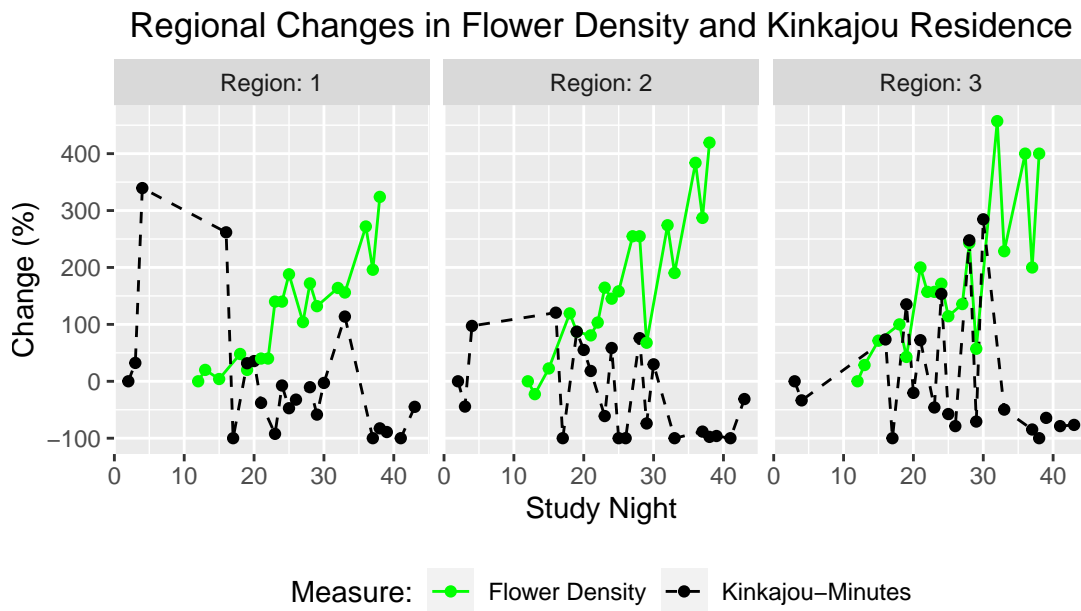


Figure 2.2: Changes in kinkajou residence and flower density across tree crown regions for the full study period. All values are given in terms of percent change from initial values in each region for unbiased visual comparison between variables.

## CHAPTER 2. MOVEMENT IN HOTSPOT

in the posterior distributions accurately represent the range of most-likely values.

Neither the total flower density in the tree crown nor the relative flower density of a region had a clear effect on the probability that at least one kinkajou entered the region in question. The mean estimate of  $p_0$ , or the intercept of the logit-probability of a kinkajou visiting an “average” region, was -2.39 (89% CI [-4.53, -0.19]). The effect of absolute flower density in the full tree crown on this probability was greater than zero in 67% of the posterior samples ( $\bar{\beta}_1 = 0.31$ ; 89% CI [-1.04, 1.55]), offering no clear evidence of an effect. The effect of the difference in flower density between the given region and the tree-crown average on the probability of a kinkajou arriving was greater than zero in 34% of the posterior samples ( $\bar{\beta}_2 = -0.67$ ; 89% CI [-3.27, 1.95]).

Flower density had clearer effects on residence time in a region when at least one kinkajou visited. The mean estimate of  $\mu_0$ , or the intercept of the log-residence-time in a region (when not zero), was 2.58 (89% CI [0.94, 4.05]). In other words, our model predicts that if at least one kinkajou visited a region with no flowers, we could expect a residence time of approximately thirteen kinkajou-minutes in the region that night. The effect of absolute flower density in the full tree crown on expected kinkajou-minutes ( $\beta_3$ ) was greater than zero in only 11% of the posterior samples ( $\bar{\beta}_3 = -0.67$ ; 89% CI [-1.5, 0.22]). This means that, contrary to our initial prediction, if at least one kinkajou visited a region when both the regional flower density and full tree crown density were at the study average density of 1.45 flowers/10m<sup>2</sup>, our model predicts a residence time of 4.95 kinkajou minutes, and we can be moderately confident that there is a true negative relationship between the total density of flowers in the tree and the time that kinkajous spent in any given region. The effect of the difference in flower density between the given region and the tree-crown average on expected residence time ( $\beta_4$ ) was greater than zero in 75% of the posterior samples ( $\bar{\beta}_4 = 0.3$ ; 89% CI [-0.45, 1.02]). We thus have weak-to-moderate evidence that

## RESULTS

kinkajous spend more time in regions of the tree crown that have greater flower density than other regions during that observation period (and this is not simply due to kinkajous spending more time in regions that, on average, have greater flower density.)

Comparisons between the empirical data and simulated data reveal that as the flower density of a region increases relative to the rest of the tree, kinkajous visit that region less frequently - but stay longer - than would be expected if the kinkajous' patterns of movement through the tree were driven by a memory-less process agnostic to current flower distributions. We observe this because posterior distributions of Model 1 coefficients, when fitted to the simulated data, reflect the posterior distributions when fitted to empirical data, except in the cases of  $\beta_2$  and  $\beta_4$  (Figure 3). Whereas the majority of estimates for  $\beta_2$  (effect of relative regional flower density on arrival probability) are negative for the empirical data, most estimates are positive for the simulated data ( $\bar{\beta}_{2sim} = 0.65$ ; 89% CI [0.02, 1.28]). Conversely, the majority of estimates for  $\beta_4$  (effect of relative regional flower density on stay duration) were positive for the empirical data, but most estimates were negative for the simulated data ( $\bar{\beta}_{4sim} = -0.21$ ; 89% CI [-0.3, -0.12]). Interpreting these differences requires careful consideration of why the real data are different from simulations. The simulations, by averaging transition rates between regions across the study, would fail to capture the dynamics of kinkajous moving through the tree based on where they have already been, or switching between states of movement and residency. The simulations maintain, however, the time and location that kinkajous first entered the tree. Thus, the somewhat surprising finding that flower density positively affects the probability of simulated kinkajous entering a region must reflect that real kinkajous tend to initially enter higher density regions and then sample the full tree; simulated kinkajous, by contrast, are more likely to exit the tree without first visiting the lower density regions. At the same time, an overall tendency to move out of the

initial-entry region for sampling, captured by the negative relationship between flower density and residence time for simulated kinkajous, may be opposed in real kinkajous that stay longer only when that region actually has greater flower density. For both parameters, however, the bulk of posterior estimates from the simulated data remain within the credible interval of estimates from the real data, preventing strong conclusions from being made.

## Kinkajou Decision Making

To analyze whether kinkajous chose to move into regions with greater flower density, we built a second model to estimate the probability that a transition between regions was a “success”, meaning the kinkajou moved into the region with greater flower density, as a function of 1) the transition the kinkajou made (e.g. Region 1 to Region 2) and 2) the difference in flower density between the region the kinkajou entered (Region 2, in previous example) and the region it did not (Region 3, in previous example). We observed 54 transitions by kinkajous between tree crown regions on nights for which flower density data were also available. Of these transitions, 35 (65%) were to the available region with greater flower density. We found that success was driven primarily by the transition being made, meaning that - although kinkajous were more often successful than not - we can not say with any confidence whether this was due to a preference for moving into more flower dense regions, or a tendency to move into the central region which, more often than not, happened to be the most flower-dense (Figure 4A).

Specifically, kinkajous typically moved from Regions 1 and 3 into Region 2 and therefore success rates are high when Region 2 has greater flower density, but lower when it is not the more flower dense region. In a small majority of posterior estimates, the probability of success also increases as the magnitude of difference in flower density between regions



## RESULTS

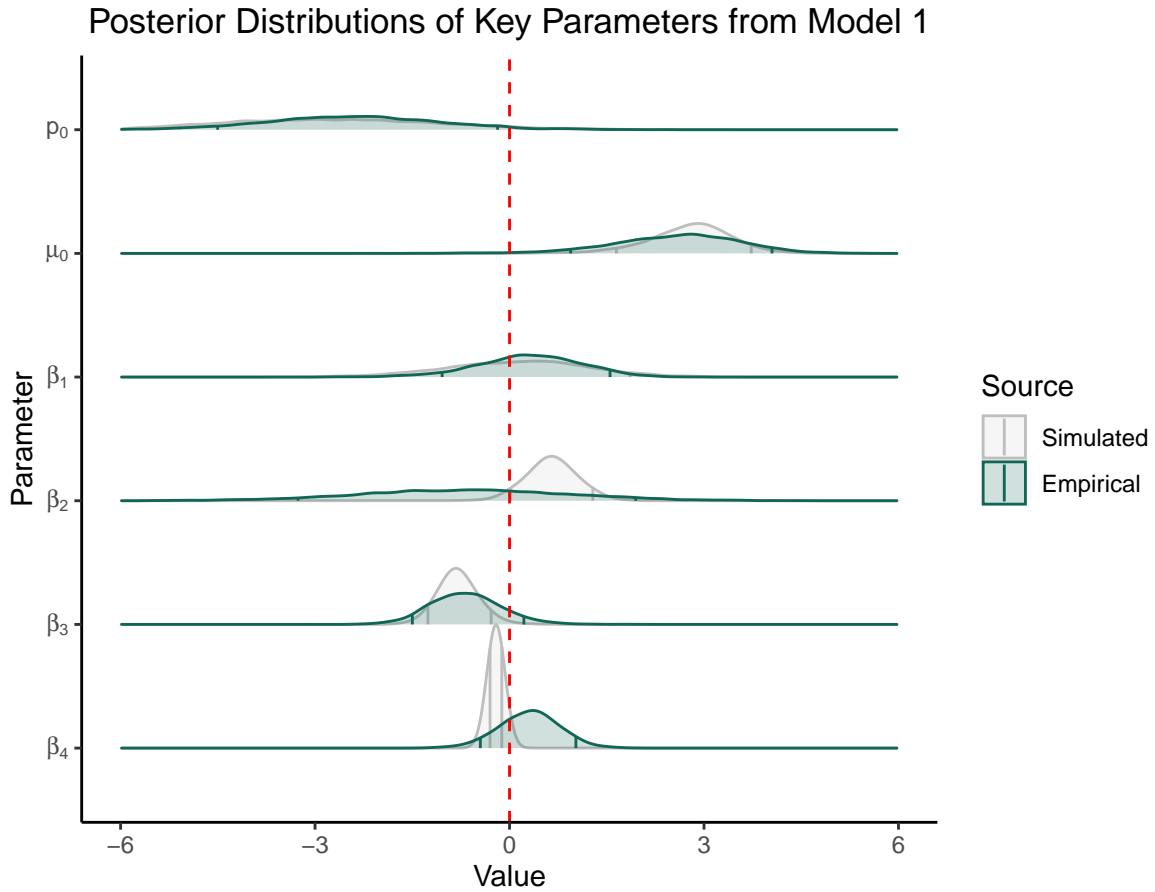


Figure 2.3: Posterior distributions of selected parameters from Model 1, which identifies the impact of regional flower densities on kinkajou residence time.  $\beta_1$  and  $\beta_2$  give the effect of one flower/10m<sup>2</sup> in the whole tree and a given region, respectively, on the logit-probability that a kinkajou enters the given region during the observation period.  $\beta_3$  and  $\beta_4$  give the effect of one flower/10m<sup>2</sup> in the whole tree and a given region, respectively, on the expected log(minutes/hour) that kinkajous will spend in the given region during the observation period.

under consideration increases, but the range of this posterior distribution is wide ( $\bar{\beta}_5 = 0.48$ , 89% CI = [-1.83, 2.81]). This amounts to very weak evidence that if choosing between regions with a greater differential in flower density, kinkajous are more likely to go to the more dense region. Overall, however, when the probability of success is marginalized over conditions (i.e. we control for repeated measures across nights and specific transitions), the credible interval ranges from nearly zero to nearly one across all levels of flower density difference (Figure 4C); i.e. we have no ability to determine how likely a kinkajou is to choose the more flower dense region without knowing specifically where it is, where it ought to go, and how frequently it makes that transition overall. Comparisons of empirical and simulated data offer little additional insight. When Model 2 was fit to the simulated data, in which kinkajous moved between regions with probabilities equal to full-study averages in the empirical data, the posterior distribution of estimated success probability was nearly identical to estimates from the empirical model fit (Figure 4B). The posterior distribution of estimated effects of flower density difference on kinkajou success rate for the simulated data was much narrower than the empirical model fit, but still centered near zero ( $\bar{\beta}_{5sim} = -0.02$ , 89% CI = [-0.31, 0.27]). Comparing the posterior distributions of empirically fit model parameter estimates to the posterior distributions of model parameters fit to simulated data (rather than comparing to zero) accounts for potential contingencies between observed flower distributions and apparent patterns of decision making that might arise if kinkajous' typical rates of movement through the tree were driven by a memory-less process agnostic to current flower distributions. In this case, such contingencies do not appear to be a concern.

## RESULTS

### Analyzing the Role of Memory in Decision Making

When evaluating kinkajous' transitions with regard to the average density of flowers on all previous nights (Model 3), the probability of success also varied significantly by the transition being made, and along the same lines as Model 2 (Figure 4A). Of the 54 transition observed, 35 (65%) were to the region that had a higher average density of flowers on previously monitored nights. Though this success rate is coincidentally identical to that of Model 2, it was not the same set of transitions that counted as success in the two models: only 21 transitions were to a region that had both the higher density of flowers that night, and on average over previous nights. There was very little variation in the magnitude of average historical density difference between regions (mean = 0.02, sd = 0.06), which likely accounts for the greater degree of uncertainty for estimates of the effect of average previous density on success, relative to that of current density. What variation did exist did not have a notable impact on the decisions of kinkajous. ( $\bar{\beta}_6 = -0.74$ , 89% CI = [-2.23, 0.78]). This holds true when Model 3 was fit to the simulated data as well (Figure 4B). As with the difference in current flower density, the estimated probability of a kinkajou moving to the historically more flower dense region varies considerably by conditions; without knowing which region the kinkajou is currently in and which of the other regions was historically more flower dense, we are unable to predict the kinkajous' probability of moving to the historically more flower dense region (Figure 4C).

### Recursion Analysis

The adjusted conditional entropies of solo-movement sequences offer no evidence that the region a kinkajou came from influences the region it will go to next. We predicted that

the empirical data will have a minimum conditional entropy at an order greater than or equal to 2, meaning that information about its last location will improve our ability to predict where it goes next, relative to having information only about its currently location. The second order conditional entropy (0.725) is indeed slightly smaller than the first order conditional entropy (0.737), but this reduction is not outside the range of variation in the simulated data, which uses a memory-less process (Figure 4D).

## Social Behavior

### Region Transitions

As with flower density, we found no evidence that the distribution of other kinkajous in the tree crown influences the region transition choices of kinkajous. We observed 16 transitions by kinkajous between tree crown regions during which the number of kinkajous in the available regions differed, 5 (31%) of which were to the region with more kinkajous. Unlike the distribution of flowers, the distribution of other kinkajous during observed decision making events was not regionally biased, and thus the transition made had relatively little impact on our “success” metric (moving to the region with more kinkajous). Unsurprisingly, given the small number of social-decision-making events observed, the posterior distributions of estimated success probability and its underlying coefficients are wide (Figure 5). Overall, kinkajous moved more frequently into the region with fewer kinkajous, but not more frequently than could be expected given random variation ( $\bar{p}_0 = -0.42$ ; 80% CI [-2.8, 1.99]).

RESULTS

Asocial Predictors of Region Transitions by Kinkajous

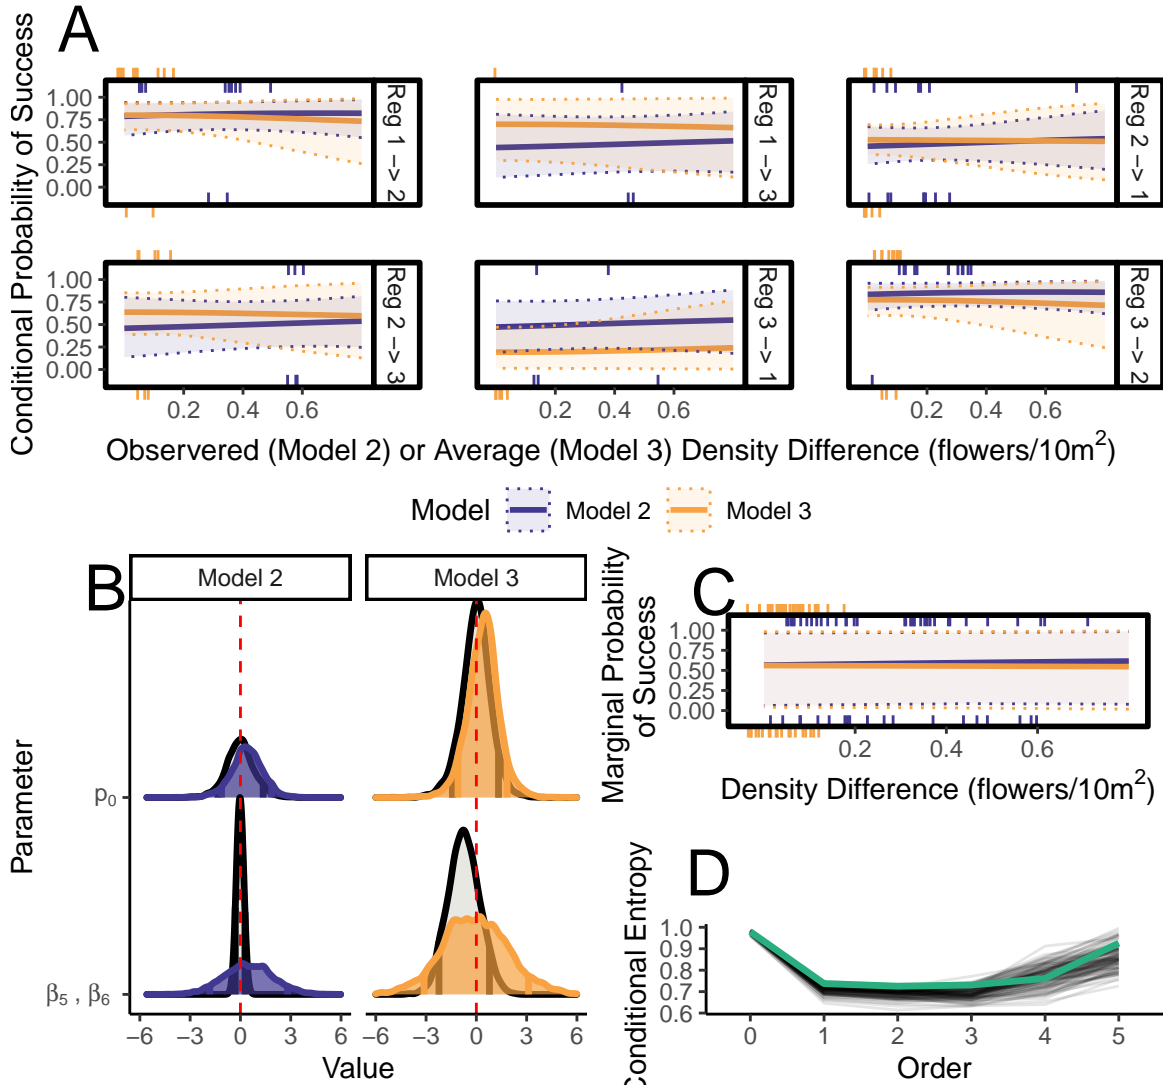


Figure 2.4: A) Posterior predictive probability of a kinkajou choosing the region with greater flower density, by success condition. Solid lines give the mean posterior estimate and dotted lines the 89% credible interval. B) Posterior distributions of select Model 2 and 3 parameters when fit to empirical (colored) and simulated (black) data.  $p_0$  is the grand mean of the success probability,  $\beta_5$  and  $\beta_6$  are the linear regression coefficients for the flower density difference between regions at the moment of choice (Observed, Model 2), and on previous nights (Average, Model 3). C) Posterior predictive probability of a kinkajou choosing the region with greater flower density, marginalized across the estimated variances in condition-based adjustments to the intercept. D) Adjusted conditional entropy of location sequences in empirical (green) and simulated (black) data. 'Order' denotes the number of previous locations used to calculate entropy.

Social Predictors of Region Transitions by Kinkajous

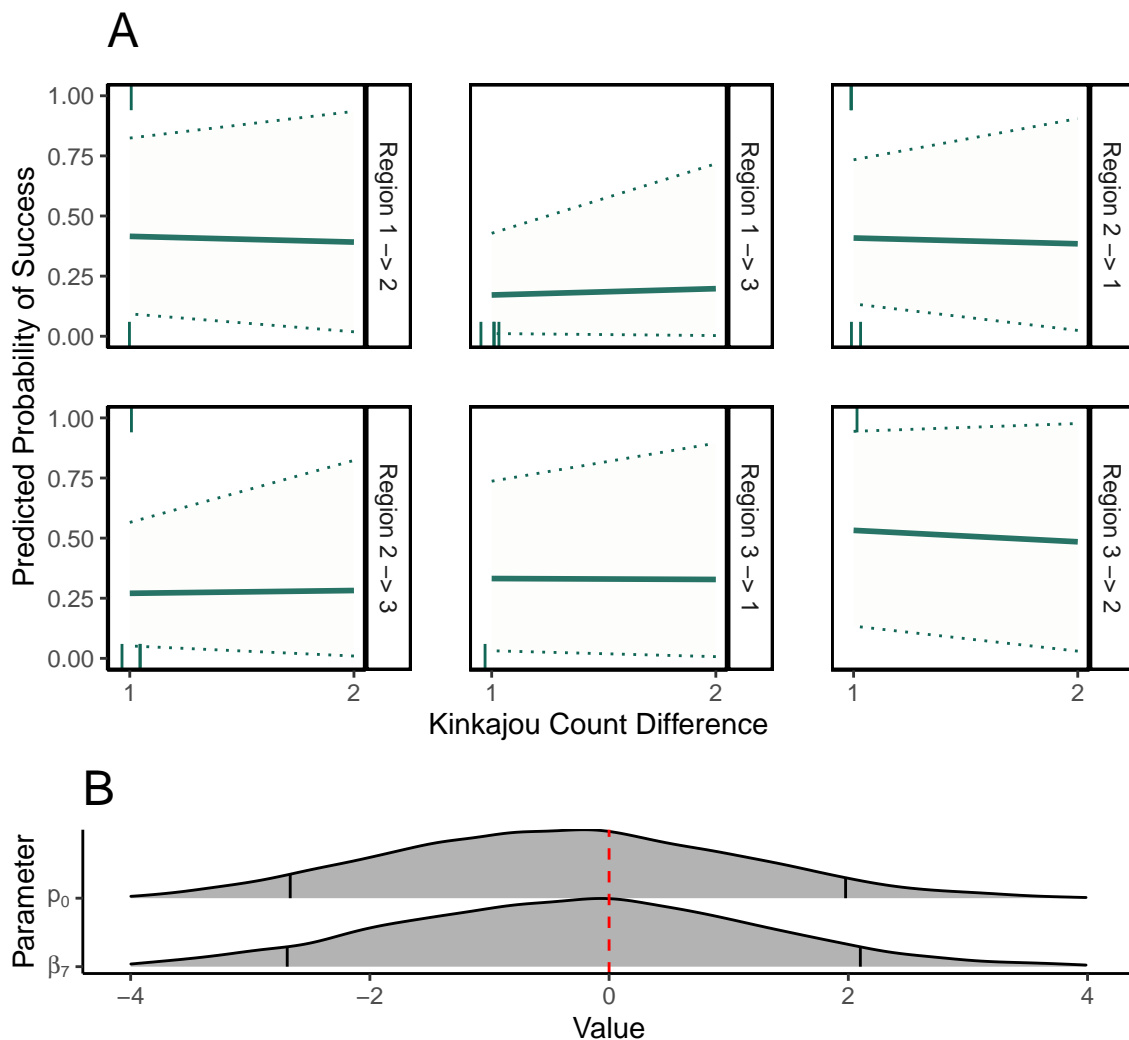


Figure 2.5: A) Posterior predictive probability of a kinkajou choosing the region with more kinkajous, by success condition (e.g. the condition 'Region 1 -> 2' represents all decisions for which the kinkajou was in Region 1 and Region 2 had more kinkajous than Region 3). Solid lines give the mean posterior estimate and dotted lines the 89% credible interval. B) Posterior distributions of select Model 4 parameters when fit to empirical data.  $p_0$  is the grand mean of the success probability,  $\beta_7$  is the linear regression coefficient for the kinkajou count difference between regions at the moment of choice.

## DISCUSSION

### Departure Rates

Though we did not find evidence that the presence of kinkajous affects which region a kinkajou moves to, we did find evidence that the presence and distribution of other kinkajous in the tree affects how often kinkajous move between regions. We observed at least one kinkajous in the balsa tree during 1128 minute-long scans. Of these scans, 212 (18.8%) included at least one kinkajou moving to a different region or exiting the tree. Overall, kinkajous were more likely to move between regions as the number of kinkajous increased (Figure 6). Kinkajous were significantly more likely to depart a region as the number of other kinkajous in that region increased ( $\bar{\beta}_8 = 0.3$ ; 89% CI [0.07, 0.52]). Kinkajous were also more likely to depart a region as the number of kinkajous in other regions increased ( $\bar{\beta}_9 = 0.14$ ; 89% CI [-0.04, 0.33]). We thus see strong evidence that kinkajous move around more in tree when other kinkajous are present, and weak evidence that this effect is stronger when the kinkajous are in the same region, relative to when they are in other regions.

## Discussion

### Conclusions

When examining the presence of kinkajous in the balsa crown, relative to the distribution of flowers, three clear patterns stand out.

First, while the overall density of flowers in the crown increased over the course of the study, the duration of kinkajou visits to the tree fell off sharply toward the study's end. This stands in clear contrast to our prediction that kinkajous would spend more time in the tree crown when flower density was greater. This contradiction may be explained

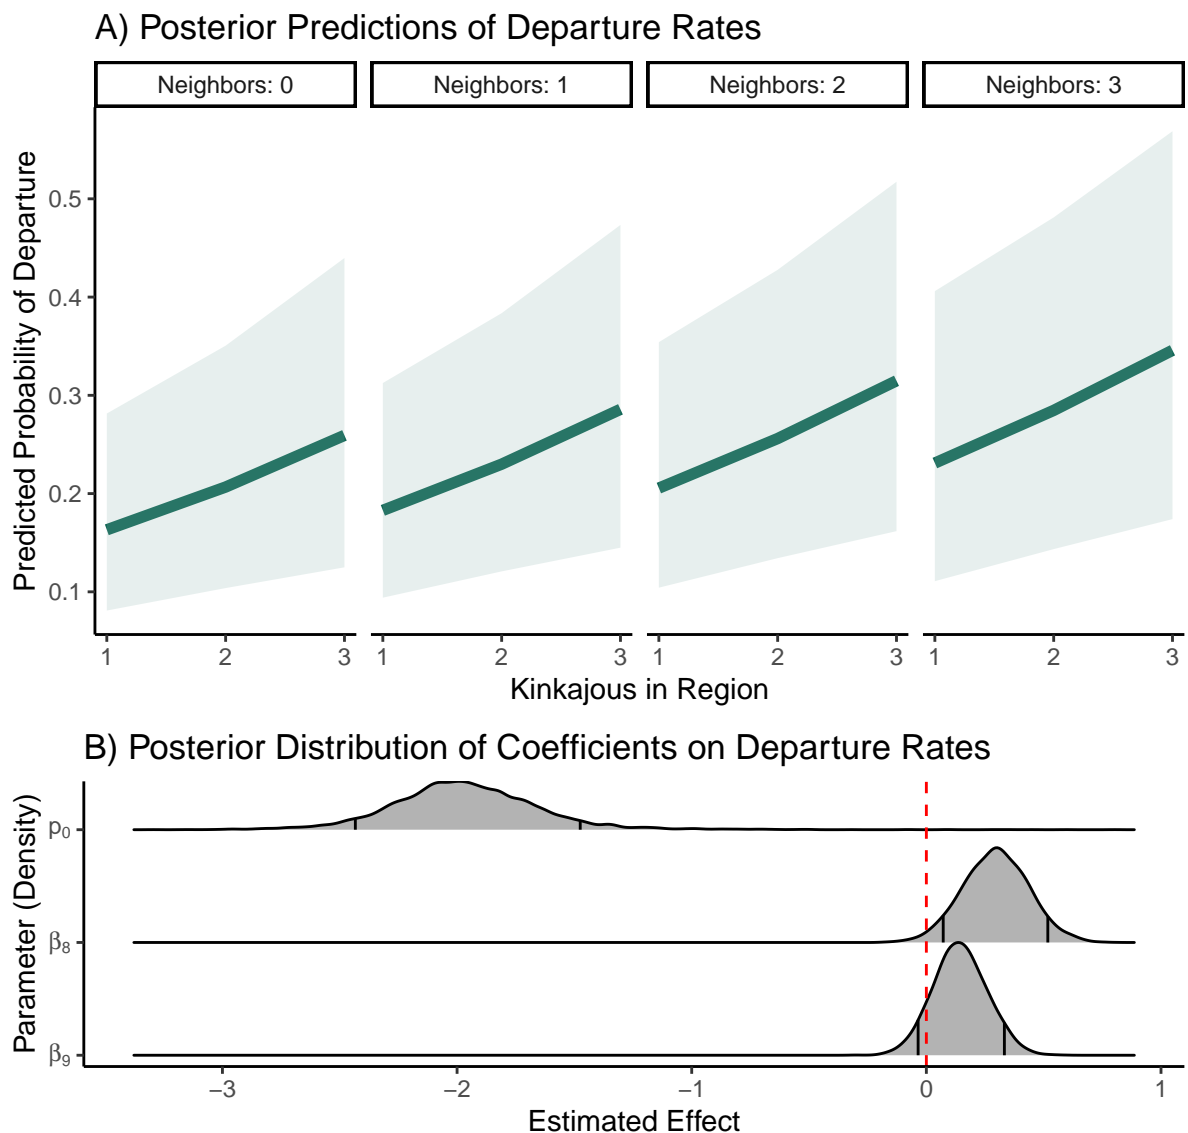


Figure 2.6: Posterior distributions of predicted departure rates of kinkajous from a balsa crown region. A) Departure rates conditioned on the number of kinkajous present in the region of interest and the number of kinkajous in other regions of the tree for a typical region (region-specific adjustment to intercept equal to zero). B) Posterior distributions of key parameters in the Departure Model (Model 5), where  $p_0$  is the grand mean of the departure probability intercept,  $\beta_8$  is the linear effect of the number of kinkajous in the region of interest, and  $\beta_9$  is the linear effect of the number of kinkajous in other regions.



## DISCUSSION

by the increasing availability of another food source, *Dipteryx oleifera* around the time kinkajous visit durations decreased, suggesting balsa nectar may serve primarily as a fall-back food during a resource scarce time-period.

Second, kinkajous move between regions more frequently when there are other kinkajous in the tree, and especially when those kinkajous are in the same region of the tree crown. This appears to support our hypothesis that kinkajous engage in scramble competition for balsa flowers, but it is hard to rule out any number of other factors that could increase overall movement activity when multiple kinkajous are present in the tree.

Third, the movement of kinkajous between regions is highly structured, with certain transitions made much more often than others. This underlying structure appears to be the primary driver of whether kinkajous move into a more flower-dense region or not, with no evidence to suggest that between-night variance in regional flower densities affects these decisions.

A fourth pattern is not strong enough to merit strong conclusions, but none-the-less warrants attention. As the density of flowers in a specific region increases relative to the other regions, kinkajous were less likely to visit that region but stayed longer when they did. This is not likely to be a result of flower processing time, as the number of flowers in the whole tree does not increase the duration of stay. If not simply a product of noise, this pattern would suggest kinkajous are not making movement decisions based on overall flower density, but do match foraging effort to local resource density - in keeping with the basic tenets of Optimal Foraging Theory. For such behavior to occur, kinkajous must be using information about the density of flowers in regions of the tree they do not currently occupy. One way they could gain this information is directly through perception. It would be somewhat surprising if they did this, given that they do not chose to move into more dense regions, but could be explained if the size of our regions

does not capture the scale of kinkajou decision making, or if kinkajous visit and deplete patches of flowers systematically. This latter explanation would indicate kinkajous have pre-existing knowledge of patch locations within the tree. Alternatively, kinkajous may acquire information about the density of flowers in each region on a given night through sampling. This would require the kinkajous to visit each region of the tree for at least a short period each night (consistent with the data) and maintain within their working memory an estimate of the flower density in previously visited regions.

Though we found no evidence that kinkajous assess the current distribution of flowers across the tree crown when making decisions, or that their movements through the tree crown were shaped by previous experience with the distribution of flowers, we are not able to rule these hypotheses out. The number of relevant “decisions” we observed is small, and the credible intervals on the parameters of interest in our models are large. In short, we lack the power to rule anything out. Additionally, it is quite possible - likely even - that kinkajous make decisions that are not captured by the boundaries and scale of our regional distinctions. Exploring how kinkajous make decisions at individual branches in the tree crown structure may yield different results. Unfortunately, our efforts to collect additional data with better protocols for resolving the three-dimensional distribution of kinkajous and flowers as a part of this project were disrupted by the COVID-19 pandemic.

## **The need to track individuals**

Our ability to extend the conclusions of this study to the way kinkajous learn during a foraging task is greatly limited by our lack of individually identified kinkajous. For example, we use average flower counts on previous nights to estimate what information a kinkajou *might* have; the power of this analysis would increase considerably if we could assess the distribution of flowers when we knew a kinkajou to have previously visited the

## *DISCUSSION*

tree. For example, we could ask whether the distribution of flowers during one visit by a kinkajou predicts that kinkajous initial trajectory through the tree on its next visit - and whether this effect decreases the longer apart the two visits are. Such analyses could reveal rates of learning and forgetting in kinkajous and whether they match the rates at which flower density in the tree crown changes. Additionally, we could explore variation in the patterns which kinkajous move through the tree. For example, if individual kinkajous tend to move through the tree in a consistent pattern, but these patterns differ between kinkajous, it would suggest kinkajou are more prone to reinforce decisions that are ‘good enough’, rather than to explore at random or sample different strategies. Analyses such as these are integral to studying the roles of reinforcement and inference in animal learning, central themes of other chapters in this dissertation.

## **What kinkajou-balsa interactions can tell us about animal learning**

Nectar is often patchily distributed at multiple spatial scales, and thus nectar foraging has been used extensively to study animal learning. When foragers can perceive multiple options, they typically choose the options that are nearest and from which energy can most efficiently be extracted (Pyke 1981, Marden and Waddington 1981). Often, the nectar concentrations of flowers are not immediately apparent, however, and thus nectar foragers’ decisions usually change with time and experience as they learn to associate cues such as location or honest scent signals with nectar quality (Knauer and Schiestl 2015, Burdon et al. 2020). Learning by iterative reinforcement is sufficient to explain patterns of route optimization between multiple, spatially stable foraging locations in honeybees (Reynolds et al. 2013). In accord with the ecological intelligence hypothesis, the speed and consistency of such route formation is strongest when foraging locations are further

apart (Buatois and Lihoreau 2016). There were clear patterns in how kinkajous moved between regions of the tree crown in our study, which may reflect learned routines or the way kinkajous intrinsically interact with the structure of the tree. Further efforts to keep track of kinkajous' identities over the course of a similar study would help reveal how individual behavior changes over time, and thus how they learn to associate certain crown regions with higher overall nectar productivity.

In changing environments, foragers must balance the exploitation of learned environmental properties with continued exploration and sampling. Nectar foraging bees, for example, vary predictably in rates of exploratory behavior and behavioral extinction (the abandonment of reinforced behaviors following loss of reinforcing stimuli) based on species level differences in physiology, social structure, and central-place constraints (Townsend-Mehler 2010, Al Toufalia et al. 2013). The species in most nectar foraging studies (i.e. bees and hummingbirds), however, often fail to exhibit types of learning associated with cognitive flexibility that are seen in mammals. Honeybees, for example, fail to use novel shortcuts between known foraging locations, as has been demonstrated in bats (Harten et al. 2020), indicating the bees can learn the behavioral links between foraging locations, but do not learn the spatial relations (Dyer 1991). Conversely, some primates adopt near-optimal routes through multi-destination arrays without needing experience, perhaps by generalizing heuristics learned previously in similar situations (Chapter 1). In a laboratory testing paradigm that switches the reward contingencies back-and-forth between behavioral options, mammals - including one kinkajou (Braveman and Katz 1971) - learn to switch to previously unrewarded behaviors faster with subsequent reversals in reward outcomes, while invertebrates maintain the same rates of behavior switching no matter how many times the reversal occurs (Bitterman 1965).

Though the production of nectar in balsa crowns appears to offer a natural occurrence

## *DISCUSSION*

of such an experiment, the spatial scale of our study may also be too small for any benefits of flexible decision making to carry much consequence for kinkajous, potentially explaining why we do not find evidence of such behavior. Additionally, the variation in flower distribution across our selected regions is both minimal and fast; kinkajous would have little opportunity to learn how to exploit a recent change before another change occurred. Should kinkajous truly fail to learn new foraging strategies in response to changing resource distributions, it would challenge the existing phylogenetic and ecological theories of cognitive evolution. Thus, our results point to a need for inquiry at smaller scales (e.g. branch level decision making) and larger scales (e.g. across trees in their home-ranges) in which a lack of flexible decision-making and reversal-learning would have greater fitness consequences for the kinkajous.

# Chapter 3

## Comparing the navigation of mammals on BCI reveals ecological drivers of cognitive evolution

Authors: Alexander Q. Vining<sup>1,2,3,4</sup>, Roland Kays<sup>4,5,6</sup>, Margaret C. Crofoot<sup>2,3,4,7,8</sup>

1. Animal Behavior Graduate Group, University of California, Davis, USA
2. Department for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior, Konstanz, Germany
3. Department of Biology, University of Konstanz, Konstanz, Germany
4. Smithsonian Tropical Research Institute, Balboa, Republic of Panamá
5. North Carolina Museum of Natural Sciences, Raleigh, NC USA
6. Department of Forestry and Environmental Resources, North Carolina State University, Raleigh, NC USA
7. Center for the Advanced Study of Collective Behavior, University of Konstanz, Konstanz, Germany
8. Department of Anthropology, University of California, Davis, CA, United States

## ABSTRACT

### Abstract

Research on the behavior of Barro Colorado Island's primates has yielded insight into how the distribution of food resources has shaped the evolution of human intelligence. In this chapter, we first synthesize primate research on BCI that has helped define navigational behaviors representative of intelligence, such as route-use and directed movements. We then explore how measuring these behaviors in non-primate frugivores, such as kinkajous (*Potos flavus*) supplements previous research by enabling comparisons to solitarily foraging animals. Finally, we present new data from a study of kinkajou navigation, offering preliminary evidence that kinkajous, like primates, flexibly navigate networks of learned routes to efficiently exploit changing resource distributions.

### Introduction

#### A history of ecological intelligence research on BCI

Based on observations of black-handed spider monkeys (*Ateles geoffroyi*) and mantled howler monkeys (*Alouatta palliata*), on Barro Colorado Island (BCI), Milton (1981) hypothesized that selective foragers evolve greater intelligence. The supposition is thus: if higher energy foods are concentrated at more distant points in time and space, choosy animals must be able to learn when and where those points will be to avoid fruitlessly searching at random. Howler monkeys, Milton noted, were somewhat choosy and appeared somewhat intelligent, while spider monkeys were quite choosy and appeared quite intelligent. The interrogation of these observations and subsequent development of Milton's Ecological Intelligence Hypothesis have remained a centerpiece of primate research

### CHAPTER 3. NAVIGATION AND COGNITION

on BCI, and throughout the world.

Spider monkeys and mantled howler monkeys make for a good comparative analysis; they both weigh an average of 7-9 kg as adults, live in outsider-averse groups of variable sizes (approx. 15-45 individuals (Milton 1981, Chapman 1990)), and they (or closely related species) range in canopies through much of Central and South America. The diet of these species differs substantially, however. While Milton observed spider monkeys on BCI to spend 72% of their annual feeding time eating fruit, this number is as low as 42% for the island's howler monkeys – which spent much more time eating young leaves.

Such dietary differences are likely a proximate cause of other dissimilarities. The average BCI howler monkey home-range is 31 hectares; that of the spider monkeys is 800. Howler monkeys travel and forage cohesively as a group; spider monkeys employ a fission-fusion approach. Milton (1981) argued that spider monkeys' choosiness requires them travel much further between resource patches, yielding larger home ranges and the need to reduce inter-patch travel by foraging independently. The direct relationship between sparsely distributed resources, larger home ranges, and fusion-fusion group structure, captured by the ecological constraints model, is now well established (Chapman and Chapman 2000).

Milton's Ecological Intelligence Hypothesis, however, relies on a series of further arguments that remain less agreed upon. Route-learning, such as Milton observed in howler monkeys, is common in animals and requires only the reinforcement of behavioral responses to specific landmarks (Collett et al. 2003). For example, Chaib et al. (2021) experimentally demonstrate the capacity of BCI's *Megalopta genalis*, a sweat bee, to navigate home by using patterns of the canopy against the night sky as landmarks. Yet, more recent research found that the routes of BCI's howler monkeys are more likely to occur in areas with higher annual food production (Hopkins 2011), and are used flexi-



## INTRODUCTION

bly to exploit emerging resources beyond the monkeys' perceptual range ([Hopkins 2016](#)). This mix of flexible and routine behavior highlights how gradations in a primate's dietary composition can shape the balance of multiple, interconnected cognitive faculties.

By contrast, spider monkeys appear able to find their own way as individuals and to use their knowledge of the environment more flexibly as they respond to rapid changes in resource availability. This is seen in their less concentrated space-use overall and, subsequent to the discovery of a newly available type of fruit, rapid integration of newly productive trees into daily travel paths. Novel use of known landmarks to achieve different goals, as spider monkeys seem capable, suggests extensive spatial knowledge and flexible decision making. The information processing required for these faculties may also explain why spider monkeys exhibit long periods of infant dependency, large repertoires of communicative signals, and (maybe) parental teaching ([Milton 1981](#)). It is difficult to quantify, however, degrees of flexible decision making or landmark use from broader patterns of movement.

New technology is helping researchers on BCI link animal movements to environment features, and thus assess decision making. Analyzing primate movements in conjunction with canopy structure data obtained by airborne light detection and ranging (LiDAR), [McLean et al. \(2016\)](#) found that a spider monkey and a capuchin, but not howler monkeys, preferred to move toward areas with a higher canopy. This behavior could be associated with seeking vistas for visual wayfinding ([Harel et al. 2022](#)). High resolution GPS tracking of white-faced capuchins, known for their tool-use and extractive foraging skill ([Oppenheimer 1968](#)), and spider monkeys revealed that neither species relied extensively on specific routes to move through their home ranges ([Alavi et al. 2022](#)). Taken together, these results emphasize the complex interplay of locomotory capacity, dietary need, social structure, and cognition in determining the movement of primates.

## Kinkajou navigation can differentiate ecological and social intelligence

Disentangling the roles of diet and social structure in the evolution of intelligence requires comparisons to animals outside the primate lineage, as frugivorous diets and complex social structure are highly correlated in primates (Powell et al. 2017). Kinkajous (*Potos flavus*) are an ideal study species for this effort. Fruit makes up 90-100% of the kinkajou diet (Julien-Laferrière 1999), and many physical adaptations reflect this exclusive diet combined with strict arboreality. A prehensile tail, grasping digits and hooked hallux, flexible spine, sharp claws, and reversible hind paws help kinkajous to climb, cling, and leap. A short snout, flattened dentition, and squared jaw help with fruit processing (Ford and Hoffmann 1988). And finally, large, forward facing eyes and sensitivity to brightness contrasts likely help kinkajous visually locate fruit, despite their nocturnal activity patterns (Ford and Hoffmann 1988, Chausseil 1992). Most of these features evolved convergently with primates, and kinkajous' morphology more closely resembles that of a lemur than most of their *Carnivoran* relatives.

However, kinkajous spend 80-99% of their active hours alone (Kays 1999). Though kinkajou decision making is likely influenced by the need to coordinate sleep sites with small family units (up to two adult males, one adult female, and two offspring, (Kays and Gittleman 2001)), defend territorial boundaries from neighboring conspecifics, and seek out extra-pair copulations (Kays 2003), kinkajous do not need to achieve group consensus through their movement. This begs the question: has kinkajous' cognition evolved convergently with primates, like their morphology, or is primate intelligence specific to their social contexts?

A recent study on BCI tracked kinkajous, spider monkeys, capuchins, and coatis using

## *METHODS AND RESULTS*

GPS collars. Compared to the other species, kinkajou movements were more constrained to a network of routes within their home ranges (Kays et al. [in review](#)). This pattern of movement resembles that of howler monkeys, prompting several further questions: how do kinkajous learn routes, given that they don't have access to the same social information as howler monkeys? How precise are kinkajou routes, exactly – do they travel along the same branches? Do kinkajous, like howler monkeys, flexibly use their route networks to exploit ephemeral fruit resources?

Because kinkajous forage alone and mostly on fruit (like spider monkeys) but in a smaller range (like howler and capuchin monkeys), answers to these questions will help identify the mixture of selective pressures that shape the evolution of cognition across lineages.

### **Do BCI kinkajous utilize route-networks flexibly?**

We recently piloted two new approaches to better understand how kinkajous on BCI use routes and navigate between resources in the canopy. First, we equipped one kinkajou with a high-resolution tracking collar that recorded a GPS fix every second for six hours each night over twelve nights. Second, we tracked a different kinkajou (at lower resolution) for two months, including a period of experimental food provisioning.

## **Methods and Results**

### **High Resolution Tracking**

We fitted one adult male kinkajou, Tony Stark, with a high-resolution telemetry collar. The collar was programmed to collect GPS fixes every second from 11PM to 5AM, for twelve nights. Plotting Tony Stark's movement data from each night reveals multiple

regions in which he traveled along highly specific paths over multiple nights (Figure 1). In one case, travel paths from four different nights are so tightly interwoven as to be indistinguishable on the map. The degree of precision across paths through this route ( $<5$  m) relative to the route's length ( $\sim 200$ m) suggests either that his movement through the canopy in this region is highly constrained, or that he is orienting to a highly localized series of landmarks such as individual branches or previously laid scent markings. Our on-the-ground observations that Tony Stark and other kinkajous move nimbly on a variety of canopy substrates, including small branches and lianas, lead us away from the former hypothesis. Notably, Tony Stark's movement through both routes highlighted in Figure 1 occurs in either direction, indicating that he is able to use a single set of landmarks to orient toward multiple goals.

## Food Provisioning Experiment

Under the hypothesis that kinkajous flexibly use learned route networks to efficiently exploit novel and ephemeral resources, we predicted that kinkajous would rapidly integrate new, reliable food sources into their nightly movement. To test this prediction, we equipped a single kinkajou (Molly) with the same GPS collar described previously, but programmed to collect GPS points every four minutes over two months. We conducted a food provisioning experiment during these months, divided into three phases: Baseline (11 days), Learning (21 days), and Randomization (19 days). The location and frequency of food provisioning varied between these phases, allowing us to measure how Molly's movement changed in response to the introduction of new resources.

During the Baseline Phase, we used incoming telemetry data to delineate Molly's core usage area and find locations for a "wagon-wheel" array of seven arboreal provisioning stations, each separated from its closest neighbors by about 150m (Figure 2). During the

METHODS AND RESULTS

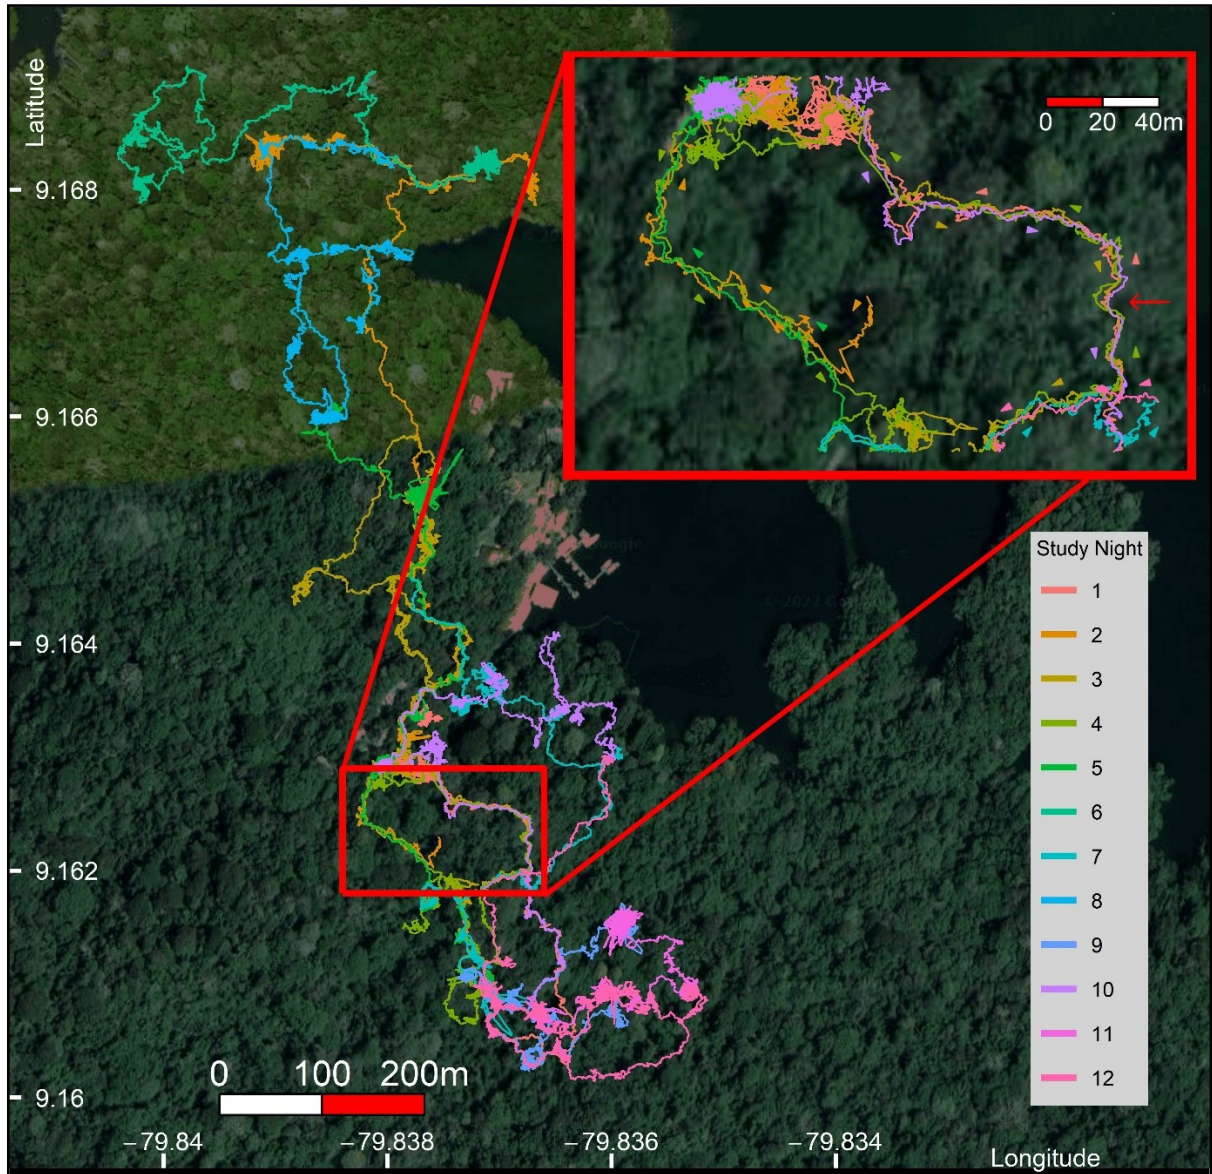


Figure 3.1: High Resolution Tracking of a Kinkajou. 1 Hz GPS tracks from a single male kinkajou, Tony Stark, tracked for six hours each night over twelve nights. Movements in the inset panel are highly correlated between nights, indicative of route use. Elongated triangles mark the direction of along each path within the inset window. The non-linear section of the route present in all four paths and indicated by the red arrow is illustrative of the degree of precision in movement that can be captured with high-resolution GPS tracking, and could reflect the use of a learned route or the avoidance of a tree-fall gap.

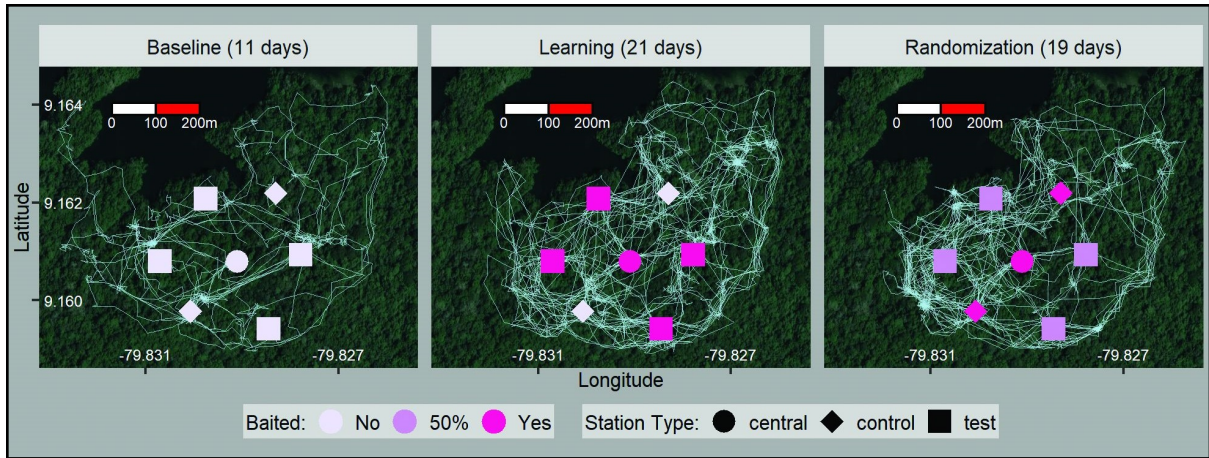


Figure 3.2: Molly's GPS Tracks by Experimental Phase. The schedule of food availability at seven provisioning stations changed through a series of three phases. The four types of stations are indicated by shape, and the frequency with which they contained food by color. Neighboring stations were separated by approximately 150 meters. New routes (light blue tracks) appeared to form between stations over the course of the study.

Learning Phase, we baited the central station and four perimeter (test) stations each night by placing a portion of fruit ( $1/3$  banana or  $1/4$  mango, both in the central station) into a bucket and hoisting it into the canopy with a pulley (Figure 3). The two remaining stations (control) were not baited during the Learning Phase. In the Randomization Phase, the central and control stations were baited each night, along with two of the four test stations, selected randomly. The total available portions available each night was held constant. Based on kinkajous' presumed ability to smell and locate ripe fruit from non-trivial distances, we predicted that Molly would visit the central and test provisioning stations sooner each night in both the Learning and Randomization Phases. Then, according to our primary hypothesis that kinkajous rapidly learn to exploit new resources, we predicted the time spent in each of these phases would have an additional effect on the speed with which Molly visited the central and test stations each night. In the Learning phase, when food was always present at these stations, we predicted Molly would visit them sooner as the phase progressed. In the Randomization phase, when

## METHODS AND RESULTS



Figure 3.3: Molly retrieves a piece of mango from the Eastern provisioning station.

food was only sometimes available at test stations but newly and consistently available at control stations, we predicted a gradual increase in the distance traveled between test stations.

To test these predictions, we calculated the distance Molly traveled each night between visits to the test and central stations. We fit custom statistical models to the distance Molly traveled each night for her respective 1st, 2nd, and 3rd station visits. These models estimated an overall effect of each phase (Condition Effects) and an accumulated, logarithmic effect of time spent in each phase (Learning Effects) on distances traveled between stations. A complete model description and its analysis can be found in the online supplement to this chapter (data notebook).

Molly quickly located and used the feeding stations (Figure 4). Because the distance Molly traveled before visiting the first station was highly dependent on where she slept, and thus not well predicted by our model, we focus on the distances Molly traveled

between stations for her 2nd and 3rd station visits each night. Relative to the Baseline Phase, the distance Molly traveled for her 2nd visit was substantially reduced in the Learning Phase (Condition Effect mean difference = -612m, 89% CI = [-1162m, -183m]) and the Randomization Phase (Condition Effect mean difference = -461m, 89% CI = [-1043m, 43m]). Our model suggests there is moderate evidence that Molly’s 2nd visit travel distances were further reduced with time spent in the Learning Phase (Learning Effect 89% CI = [-0.34,0.06]). Molly never visited three stations in a single night during the Baseline Phase, but did so frequently in the Learning and Randomization Phases. Both her 2nd and 3rd visit travel distances subsequently increased over time spent in the Randomization Phase (2nd visit Learning Effect 89% CI = [0.02,0.47]; 3rd visit Learning Effect 89% CI = [0.00, 0.50]).

## Discussion

We discovered that a kinkajou used multiple specific and bi-directional routes between commonly visited locations. This behavior indicates knowledge of many different landmarks within its home range, as well as the ability to use these landmarks for multiple navigational purposes. This same repeated narrow route use is also seen in all kinkajous tracked by Kays et al. ([in review](#)), especially males.

We also experimentally demonstrated that a different kinkajou rapidly integrated new resources into its nightly travel paths. In addition to immediately reducing travel distances between feeding stations when provisioning began, this kinkajou continued to reduce travel distances over time, evidence that the kinkajou is learning to optimize travel paths to exploit new resources. Likewise, when reliable provisioning shifted to control stations, a gradual increase in distance traveled between test station occurred.



DISCUSSION

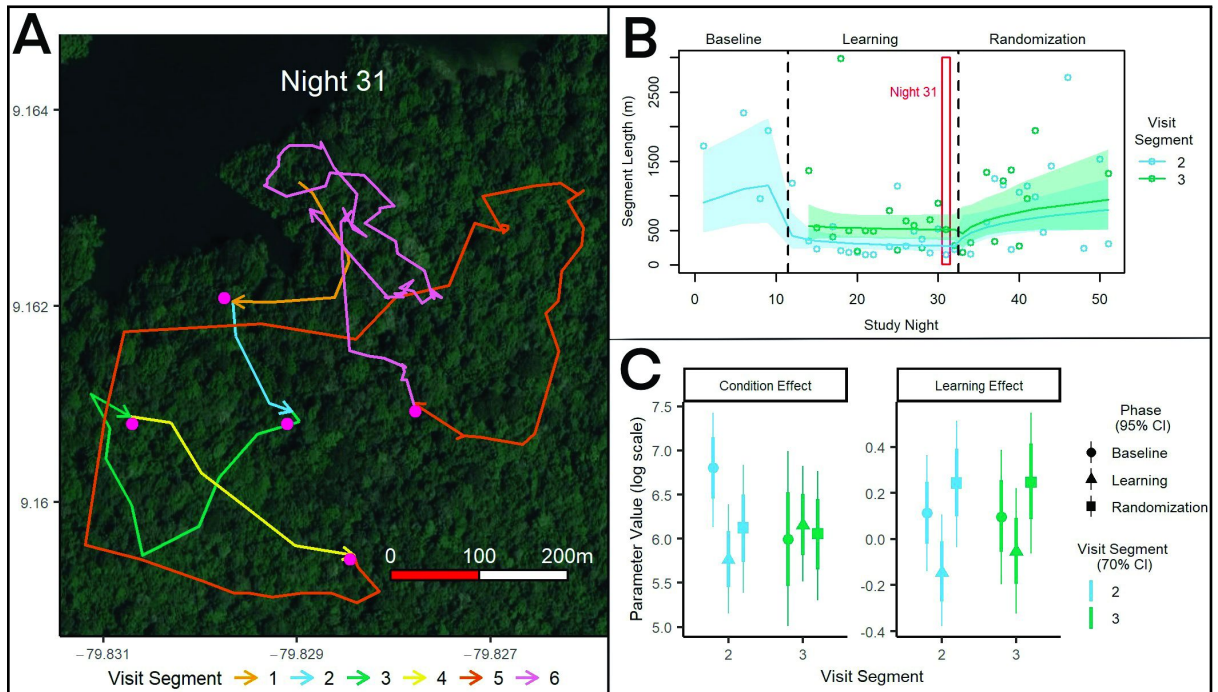


Figure 3.4: A Learning Model of Molly's Inter-Station Travel Distances. (A) Molly's movement on night 31 of the study, illustrating how the travel path on each night is broken into visitation segments. (B) The mean predicted distance for the second and third segments (differentiated by color) across nights (solid lines). Empty circles represent the observed distances traveled. Shaded areas indicate the 95% posterior predictive credible interval. (C) Credible intervals of the Condition Effect and Learning Effect parameters in the learning models.

### CHAPTER 3. NAVIGATION AND COGNITION

Taken together, these results suggest kinkajou route-use is not the result of fixed behavioral routines, but reflects the flexible use of landmarks to efficiently exploit dynamic resource distributions.

Our studies of kinkajou navigation offer preliminary evidence that kinkajous, like primates, use extensive knowledge of their environments to move efficiently between resources, and that they can rapidly integrate new information to flexibly and efficiently exploit ephemeral resources. In kinkajous, route-use may be a necessity of nocturnal travel, in which perceptual limits require greater dependence on orientation to local landmarks rather than distant ones. Our experimental results indicate kinkajous use routes flexibly to connect many different resource locations and that they rapidly alter their movement patterns in response to new resources. Kinkajous' clear use of route-learning and at least some degree of flexible decision-making is in accord with the Ecological Intelligence Hypothesis, given kinkajous high-fruit diet. Though the difficulty of directly comparing these kinkajou data to previously collected primate data prevents us from drawing conclusions about the relative roles of sociality and diet, our methodology creates the potential for controlled, experimental comparisons if applied to other species. Thus, further research that uses high resolution tracking and simple manipulations of resources can be used to directly compare mammalian navigation on BCI. This work will help elucidate the relationship between diet and sociality in cognitive evolution, continuing a decades long research tradition on Barro Colorado Island.

# Chapter 4

## Evidence of diet-mediated learning rates in the route-formation of four sympatric mammal species

Authors: Alexander Q. Vining<sup>1,2,3,4</sup>, Damien Caillaud<sup>4,5</sup>, Mark Grote<sup>5</sup>, Shauhin Alavi<sup>2,3</sup>, Roland Kays<sup>4,6,7</sup>, Ben Hirsch<sup>4,5,8</sup>, Margaret C. Crofoot<sup>2,3,4,5,9</sup>

1. Animal Behavior Graduate Group, University of California, Davis, USA
2. Department for the Ecology of Animal Societies, Max Planck Institute of Animal Behavior, Konstanz, Germany
3. Department of Biology, University of Konstanz, Konstanz, Germany
4. Smithsonian Tropical Research Institute, Balboa, Republic of Panamá
5. Department of Anthropology, University of California, Davis, CA, United States
6. North Carolina Museum of Natural Sciences, Raleigh, NC USA
7. Department of Forestry and Environmental Resources, North Carolina State University, Raleigh, NC USA
8. Center for Tropical Environmental and Sustainability Science, James Cook University, Queensland, Australia
9. Center for the Advanced Study of Collective Behavior, University of Konstanz, Konstanz, Germany

## Abstract

For animals that selectively forage on resources distributed in spatially distant patches with short periods of high productivity, such as fruit, the development of efficient routes can greatly improve foraging efficiency. A number of cognitive mechanisms, including reinforcement learning, episodic memory, and cognitive mapping, likely play a role in such route-formation. Though route-use is commonly observed in free-ranging animals, few studies have quantified the development of such routes. Such efforts can elucidate the learning mechanisms an animal uses for navigation and, when compared across many species, the contexts in which those mechanisms are likely to evolve. We used high-resolution GPS collars to track forty-six individual animals of four different mammal species foraging for fruit in a lowland tropical rainforest. Due to seasonal scarcity, we were able to map all potentially productive trees during the periods of study. Combining a novel combination of recursive sequence analysis and a hierarchical Bayesian model, we identify highly revisited points of interest (POIs) within each animal's territory and describe an asymptotically curved relationship between the distance separating two POIs and the length of animal transits between them. We found that an animal's previous experience with a given transit significantly reduced its expected transit length, with variation in learning rates mostly predicted by degree of frugivory. We suggest that our results offer insight into how different animals learn about and map their environments, and that our methods open the door to broader, analytically rigorous comparative analyses.

## Introduction

Route-formation reflects an animal's ability to build knowledge of the relationships between places and desired outcomes, such as food or mates (Reid and Staddon 1998). The learning mechanisms of route-formation have been well studied through navigational experiments, but most such studies occur in captive settings, such as the radial arm maze and the Morris water maze (Knierim and Hamilton 2011). Janmaat et al. (2021) argue that by studying the formation of routes (among other patterns) in animals' natural movement trajectories, it is possible to expand the scope of learning research in many ways, including but not limited to: identifying learning patterns at larger spatio-temporal scales (which may or may not differ from patterns at captive scales as a function of, e.g. energetic constraints or memory (Fagan et al. 2013)); comparing learning patterns across a broader scope of species; and exploring how learning shapes the behavior and fitness of animals in a complex world. Measuring the formation of routes in the wild, however, has been historically limited by the difficulty of collecting and analyzing data on animal movements and resource distributions at appropriate resolutions and scales. In this study - with the aid of new GPS tracking technology, remote environmental sensing, and developing statistical methods - we analytically compare rates of route formation in 46 individual animals of four different species.

There is a long history of using patterns of routine movement to make inferences about the cognition of animals, including those with relatively small brains. The formation of traplines - repeated visitation sequences between foraging locations - in honeybees (Buatois and Lihoreau 2016), bumblebees (Lihoreau et al. 2012b), and hummingbirds (Garrison and Gass 1999) has been studied extensively using artificial feeding arrays, with compelling evidence that these animals learn traplines by reinforcing movement responses

## CHAPTER 4. ROUTE FORMATION

to specific landmarks that, on average, yield greater foraging gains during individual bouts from a central place (Pyke 1978a, Reynolds et al. 2013). Evidence that honey bees also make novel shortcuts between known foraging locations (Menzel et al. 2012) has sparked debate about the roles of planning, path integration, and spatial mapping in insect navigation (Collett et al. 2013, Menzel 2021). Overall, many insects appear to exhibit site-fidelity and repeated travel sequences, indicating the storage of foraging memory for route-development even at large spatial scales (Janzen 1971, Ackerman et al. 1982, Collett et al. 2003, Wehner 2003, Chaib et al. 2021), with some variation in the extent to which knowledge of learned routes is integrated into what is often called a place-map (Collett and Collett 2006, Najera 2009).

For larger brained animals, ubiquitous patterns of site fidelity and routine space-use suggest more wide-spread use of place-maps, which is widely hypothesized to be a primary benefit in the selection for larger brains (DeCasien et al. 2017, Powell et al. 2017, Rosati 2017). Routes are used, for example, by spider monkeys and howler monkeys to monitor large numbers of fruiting trees (Di Fiore and Suarez 2007, Hopkins 2016) and by sea turtles to make long distance migrations (Krochmal et al. 2021). However, large-brained animals also exhibit varying degrees of behavioral flexibility (Riotte-Lambert et al. 2017) and possess neurological systems that appear capable of integrating place-maps into near-euclidean spatial representations (Jacobs and Schenk 2003) and/or hierarchical systems of knowledge (i.e. cognitive maps) (Buzsáki 2005, Behrens et al. 2018, Ekstrom and Ranganath 2018). A small number of recent studies have used long-term tracking to reveal the use of novel shortcuts in the natural foraging of, for example, elephants (Presotto et al. 2019) and fruit bats (Harten et al. 2020), evidence of euclidean representations. Some experiments with primates have also revealed flexible use of spatial knowledge, suggestive of cognitive maps, to make foraging decisions (Janson 2016,

## INTRODUCTION

Trapanese et al. 2022). Understanding the evolutionary and ontogenetic development of reinforcement-based place-maps into more sophisticated mental representations, however, requires continuous metrics that can be compared across species and developmental periods. Path linearity has been used as one measure of spatial knowledge (e.g. Harten et al. (2020); Jang et al. (2019)), but Janmaat et al. (2021) note that path linearity can result from many mechanisms other than spatial representations, and that a full picture of what wild animals know requires measures of how path linearity (and other metrics) change over time, experience, and context. As of yet, no study has quantitatively compared the development of path linearity in multiple species of free ranging animals. This is the primary goal of our study.

For a number of reasons, we focus our study on the fruit-eating mammals of Barro Colorado Island, Panama. Firstly, this lowland tropical rainforest experiences a period of food scarcity from November to December, alleviated somewhat by the fruiting of a single tree species, *Dipteryx oleifera*, in January and February. The fruit-producing crowns of *D. oleifera* are usually large, sparsely distributed, and visible from above the canopy. They also exhibit significant inter-tree variation in the timing of peak fruit productivity across the fruiting season (Caillaud et al. 2010). Therefore, we predicted that rapid formation of efficient routes between productive trees will be important if fruit-dependent species are to competitively exploit this resource.

Additionally, the diverse mix of fruit-eating species on the island lends itself to interesting comparative analyses. Spider monkeys (*Ateles geoffroyi*) and kinkajous (*Potos flavus*) are both exclusively arboreal, highly selective foragers that depend extensively on a diet of high-quality fruit (Julien-Laferriere 1993, Kays 1999, Di Fiore et al. 2008, González-Zamora et al. 2009, Lambert et al. 2014). Fruit is also an important dietary resource for coatis (*Nasau narica*) and capuchins (*Cebus capucinus*) on the island, but these

## CHAPTER 4. ROUTE FORMATION

species additionally forage for insects, seeds, and other protein sources (coatis primarily-but-not-always on the ground and capuchins primarily-but-not-always in the canopy) (Alves-Costa et al. 2004, Hirsch 2009, Mosdossy et al. 2015). Coatis and capuchins also both live in relatively stable social groups (Gompper 1997, Fragaszy et al. 2004, Crofoot et al. 2011); spider monkey social groups exhibit fission-fusion dynamics (Asensio et al. 2008); kinkajous den with small family units but typically forage alone (Kays and Gittleman 2001). Kinkajous are the only one of the four species primarily active at night. Kinkajous and coatis are both in the *Procyonidae* Family of *Carnivora*; spider monkeys and capuchins belong to different Families of *Primates*. Because the variation in some key life history traits in these species are not taxonomically linked, a comparison can offer preliminary insights into potential convergent evolution of the cognitive mechanisms for route-formation.

For example, when high quality resources are concentrated in sparse but spatially stable patches with semi-predictable patterns of depletion and renewal (as with *D. oleifera* fruit), the Ecological Intelligence Hypothesis predicts an energetic trade-off associated with selective foraging: selective foragers need not invest in specialized digestive systems for processing low-quality foods, but must travel further between food sources and maintain energetically expensive cognitive systems to remember where food can be found, predict when it will be available, and find efficient travel paths between productive foraging patches (Milton 1981). For example, howler monkeys (which are closely related to and broadly sympatric with spider monkeys) invest a lot of energy in a specialized digestive tract for processing young leaves; compared to spider monkeys, however, they have much shorter daily travel lengths and more stable (i.e. routine) travel patterns (Hopkins 2011). For coatis and capuchins, the case may be that they have a more generalized diet, but are not necessarily less selective. Instead they exhibit extractive foraging strategies for



## METHODS

exploiting quality food sources that are more randomly distributed and difficult to find or access than fruit. We predict that quickly learning efficient routes between highly productive resource patches is less important for these extractive foragers.

In summary, the goal of our research is to identify the socio-ecological factors that drive route-formation, keeping in mind 1) the inferences that can be drawn from rates of route-formation regarding animal learning and 2) the trajectories of cognitive evolution across different taxa.

## Methods

### Study Site

Data were collected on Barro Colorado Island (BCI), Panama. BCI is a 15.6 km<sup>2</sup> island with semi-deciduous lowland tropical forest. The wet and dry seasons on BCI are very distinct, where approximately 90% of an average annual 2600mm rainfall occurs between May and December (Dietrich et al., 1982). Fruit availability fluctuates with these seasons, with a period of scarcity in the late wet season (October and November) broken by the fruiting of *D. oleifera* at the start of the dry season (mid-December) (Foster 1982).

### Mapping *Dipteryx oleifera*

All trees on the island were mapped during 2015 and 2017 using either a custom designed fixed wing (Penguin, FinWing, China) or multi-rotor UAV (Phantom 4 Pro, DJI, Shenzhen China). Two sets of flights were conducted in each year of the study, the first at or a bit before peak flowering and the second approximately two weeks later. Due to weather conditions and time constraints, flight altitude differed across the 4 sets of flights (but

was held constant within each set), yielding imagery that ranged in resolution from 4-15 cm/pixel. Photos were processed using the Agisoft Metashape program (Agisoft LLC, St. Petersburg, Russia), to create four georeferenced orthomosaic images of BCI, which were imported into ArcGIS (ESRI, Redlands, CA, USA). Dipteryx trees were identified by their distinctive pink/purple color and polygons were manually delineated to cover the maximum visible extent of each tree crown. We merged contiguous tree crowns into larger polygons, thus some tree patch visits by our animals represented feeding at more than one tree in a single patch visit.

## Animal Capture and Tracking

We fit GPS/3-D accelerometer collars (e-Obs Digital Telemetry, Gruenwald, Germany) to individuals from four species, two primates, capuchins (*Cebus capucinus*), spider monkeys (*Ateles geoffroyi*), and two procyonid carnivores, kinkajous (*Potos flavus*) and coatis (*Nasua narica*). Collars were programmed to collect a burst of six consecutive (1 Hz) GPS locations every 4 min during the animal's active periods: 06:00–18:00 for capuchins and spider monkeys, 06:00–18:30 for coatis, and 23:00–6:30 for the nocturnal kinkajous. 3D acceleration was recorded at 1-min intervals to determine activity profiles. Collaring occurred in 2015 and in 2017, with 20 individuals tagged the first field season and 26 individuals tagged the second field season. 8 spider monkeys, 7 capuchin monkeys, 16 coatis, and 14 kinkajous were tagged in total. From December 2015 to March 2016, the GPS sampling regime of collars on kinkajous and coatis was ACC-informed, with collars collecting data as described above when accelerometer readings were above a specified threshold (1,000 mV). ACC-informed sampling was not used in the second field season, from December 2017 to March 2018. All collars were programmed to timeout if they did not acquire a fix after 90 s.

## METHODS

### Data Preparation

The last fix of each GPS burst consistently had the best horizontal accuracy measurement, therefore only the last fix of each burst was used for all analyses. All data were uploaded to Movebank, an online repository for animal movement data. Duplicate and outlier fixes were removed using Movebank’s data filters, filtering fixes by the height above ellipsoid. All fixes with height above ellipsoid values less than or equal to 21 or greater than 244 were marked as outliers. This corresponds to the first quartile minus twice the interquartile range and the third quartile plus twice the interquartile range, respectively. Subsequent outlier detection was done using the *ctmm* package in R (Calabrese et al., 2016), using error information, straight line speeds, and distances from the median latitude and longitude to manually identify outliers via the *outlie* function. Further, obviously impossible locations, such as location estimates in the water and clearly outside the boundaries of the island, were marked as outliers. For ACC informed collars, GPS locations were interpolated for times when the animals were below their ACC thresholds. The error on the interpolated positions was modeled to replicate the observed GPS error of a stationary collar in a tree, and was drawn from a negative binomial distribution with a mean of 5.46m and a dispersion parameter of 2.4m.

### Identifying Points of Interest

In order to study how animals learn to move between important locations, it was first essential to identify where those important locations are. Though one option would be to simply use *D. oleifera* tree crowns as points of interest (POIs), this misses two important features of animal behavior. First, clusters of nearby tree crowns may be viewed by an animal as a single foraging patch. Treating each individual crown as a POI would

similarly strain the computational limits of our analysis and potentially reduce our ability to identify the role of experience if an animal uses the same route to arrive at multiple, adjacent crowns. Second, there are many reasons other than food why an animal may frequently return to a location, including sleep sites, mating opportunities, territory defense, and social coordination. Failing to consider these points could result in biased analyses if, for example, animals frequently detour to these locations between foraging bouts. We thus use each animal’s own movement to identify its individual POIs. We do this by combining a recursion analysis (Bracis et al. 2018) with a clustering analysis (Hahsler et al. 2019).

The recursion analysis works by identifying, for each GPS fix in an animal’s tracking data, the number of times that location was revisited. It requires the user to set a radius that defines the visitation region around each point, and a temporal threshold that defines the time outside the visitation region after which the animal is considered to have left the area. Any time the animal’s movement trajectory re-enters the visitation region after a period of time outside the region greater than the temporal threshold, it is counted as a re-visit, or recursion. We visualized heat-maps of recursion values using different parameters, and used these heatmaps in combination with our own intuition about the spatio-temporal scale of animal foraging decisions and the error of our GPS collars to set a radius of ten meters and a temporal threshold of thirty minutes (Supplement 1: Path Clustering and Segmentation).

In the second stage of POI identification, we selected an upper percentile of recursively visited GPS fixes, performed a spatial cluster analysis of these points using the DBSCAN algorithm, and drew minimum convex polygons around the resulting clusters. We call the percentile of GPS points used the revisit cut-off (e.g. a revisit cut-off of 0.5 means we cluster on the GPS points with revisit counts in the upper 50 percentile of an indi-

## METHODS

vidual animal’s track). The DBSCAN algorithm detects clusters by identifying sets of “density connected” points in the data. Points are said to be density connected when at least one point in the cluster (a core point) contains at least some minimum number (MinPts) of other points within a neighborhood of radius *epsilon*, and all points in the cluster are within the neighborhood of at least one (other) core point. Given a value for MinPts, the algorithm will select an optimal value of *epsilon* based on the distribution of distances between each point and its MinPts nearest neighbor. This algorithm has several advantages for us over other clustering approaches such as k-means: it does not require a specific number of clusters to be pre-specified, it can detect density clusters of any shape, and it will label points that do not meet the criteria for being a cluster “outliers” rather than assigning them a cluster. We ran the full recursion-and-clustering algorithm across all animal tracks individually (producing unique POIs for each animal ID) with several values of revisit cutoff and MinPts. For each combination of values, we visually assessed resulting POI polygons across all animals tracks and chose values for our final analysis that 1) minimized the total number of polygons and 2) maximized the distance between polygons, but 3) minimized the number of polygons that included large areas of low animal use (Supplement 1: Path Clustering and Segmentation). These values were 0.5 for the revisit cutoff and 50 for MinPts.

This amounts to a (mostly) quantitatively rigorous way of defining our POIs as regions of continuous space 1) which an animal revisited more often than average, 2) in which the same animal spent approximately four or more hours total, and 3) which are separated by a meaningful distance from other such spaces.

## Hierarchical Modelling

Following the identification of POIs, we split each animal’s movement trajectory into transits, defined as the movement segments between two different POIs. We discarded transits between POIs less than 10 meters apart, as we are primarily interested in how animals move between regions they do not have immediate perceptual access to. We treated the remaining transits as our primary unit of analysis, using as a dependent variable  $\Delta L$ , the difference between the transit length (assuming straight line travel between GPS fixes) and the linear distance (D) separating the start and end POIs. We built a hierarchical model of  $\Delta L$  iteratively, first (1) exploring various non-linear models and error distributions to best fit the relationship between D and  $\Delta L$ . Based on the appearance of a humped curve in plots of transit  $\Delta L$ ’s against D, we primarily focused on non-linear models with this feature, including the Jentsch curve, the Shephard Curve, and a generalized Ricker curve. Ultimately, we chose to use the generalized Ricker curve in further analysis for several practical reasons: 1) it provided a better fit to the data, 2) it has easily interpretable parameters, 3) it has a log-linear form which simplifies Bayesian parameter estimation, and 4) it can flexibly fit initial curves that are concave-up or concave-down. Additionally, Persson et al. (1998) used the curve to model trade-offs in predator attack rates as predator body size increases (e.g. increasing attack rates due to development of perceptual systems vs. decreasing attack rates due to reduced responsiveness). This has a rough analogue to the notion that animal transits may be the result of either random search or directed travel, resulting in competing effects on  $\Delta L$  as D increases. This is because the rate of increase in  $\Delta L$  relative to increases in D would be much greater in transits resulting from random search than those from directed travel. However, the probability that a transit resulted from random search decreases as D increases. We believed such a dynamic may have been producing the patterns we

METHODS

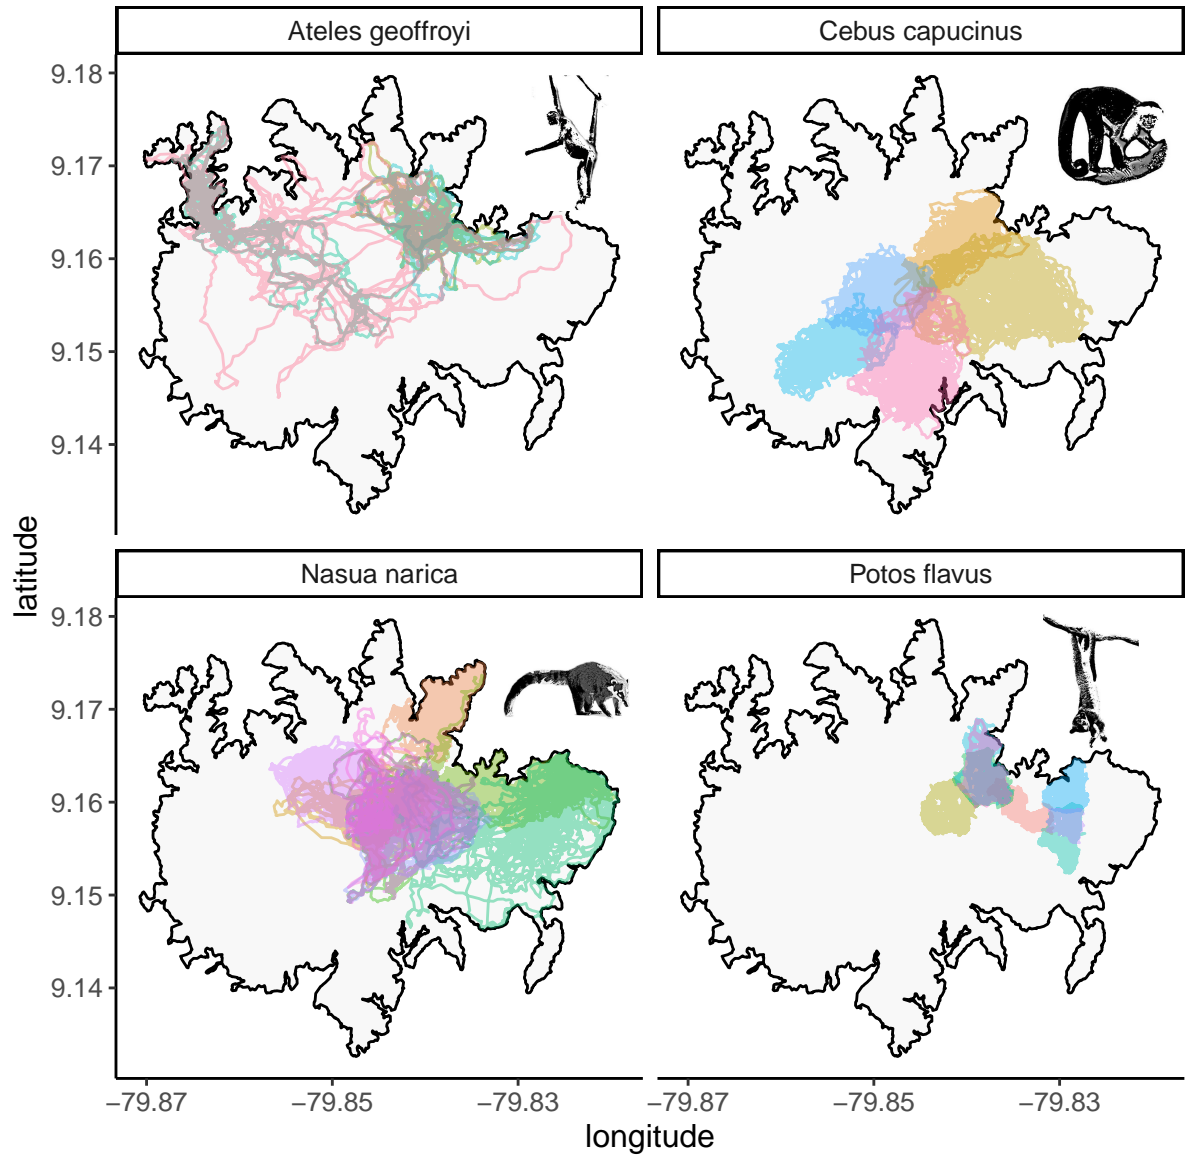


Figure 4.1: GPS trajectories for animals tracked during the 2017-2018 season. Each color is a different individual.

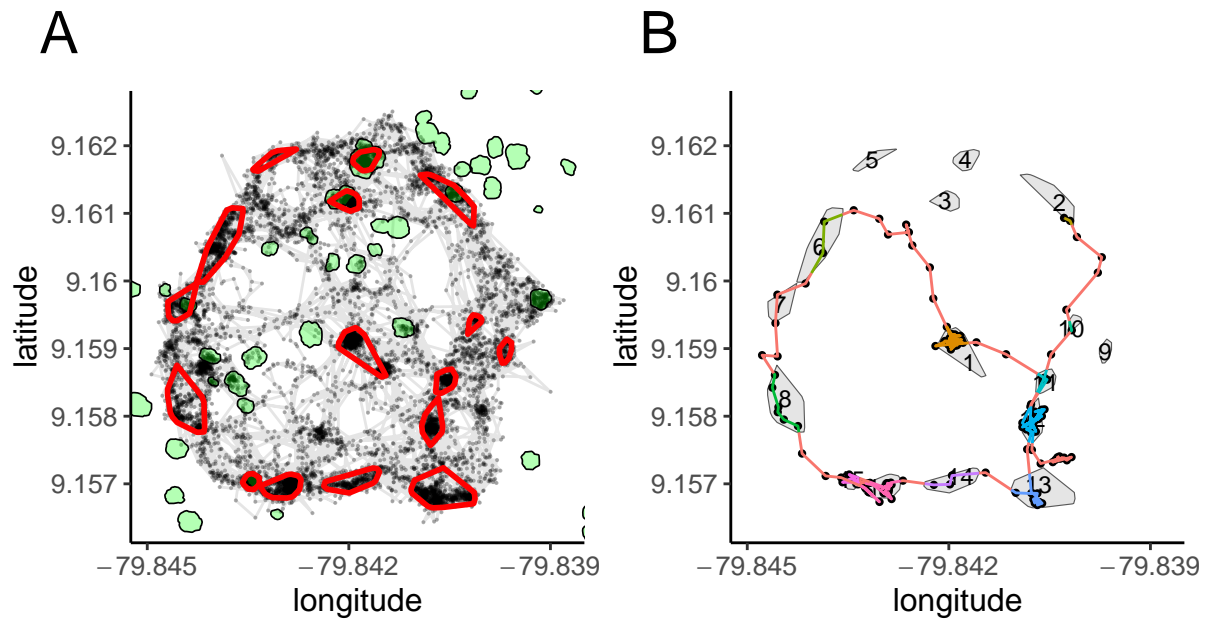


Figure 4.2: Breaking individual GPS trajectories into transits. A) The GPS fixes of a single kinkajou, with POI polygons super-imposed. B) The track from a single day of the same kinkajou in figure B, with transits between POIs colored red.



## METHODS

observed in the data. Finally, we chose to fit the curve using a gamma distribution, as this distribution (a) satisfies the physical constraint that  $\Delta L$  is positive real and (b) accommodates correlation between the expectation (mean) and dispersion (variance) in the data, which we noted during initial data visualizations.

Next, (2) we added co-efficients for species-levels adjustments to each of the three parameters in the generalized Ricker curve, allowing our model to fit an independent curve for each species (and us to compare species level differences in these curves). We then (3) included experience as a predictor of  $\Delta L$  by multiplying the maximum of the Ricker curve by the number of previous times the animal had previously made that same transit, raised to the power of an experience coefficient. In effect, this stretches the entire curve up or down (based on the sign of the coefficient) in a manner that scales with the breadth of previous experience (based on the magnitude of the coefficient). Power laws are frequently used to model learning curves ([Donner and Hardy 2015](#)). Next, we (4) added species-level adjustments to the learning co-efficient, allowing us to model species level differences in learning rates. Finally, we (5) accounted for individual difference and repeated measures by building in individual-level adjustments to the maximum of the generalized Ricker curve and the learning-rate co-efficient. At each stage, we fit the model using Bayesian estimation with STAN ([Carpenter et al. 2017](#)) and RStan ([Stan Development Team 2023](#)) and analyzed both the behavior of the Markov sampler and the outcome of the model to gain a better understanding of the patterns in our data.

However, we primarily present the results from the final model, which can be written

$$\begin{aligned} \Delta L &\sim \Gamma(\alpha, \alpha/\mu) \\ \log(\mu) &= a_{sp} + \beta_{aid} + \log(k)(r_{sp} + \beta_{rid}) + c_{sp}(\log(D) - \log(b_{sp}) + (1 - D/b_{sp})) \\ \alpha &\sim \text{exponential}(5) \\ a_{sp} &\sim N(1, 5) \\ r_{sp} &\sim N(0, 1) \\ \beta_{rid} &\sim N(0, 1) \\ \beta_{aid} &\sim N(0, 1) \\ b_{sp} &\sim N(1, 20) \\ c_{sp} &\sim N(1, 1) \end{aligned}$$

where  $D$  is the linear distance in kilometers between the start and end POIs of a given transit,  $\Delta L$  is the difference in kilometers between the length of the transit and  $D$ ,  $\alpha$  is the shape parameter of a gamma distribution,  $\mu$  is the expectation of the gamma distribution,  $a_{sp}$  is the log of  $\Delta L$  at the critical point in the Ricker curve (i.e. it's max) for the transiting species,  $\beta_{aid}$  is an adjustment to  $a_{sp}$  based on the transiting individual,  $k$  is the number of times the transiting individual has previously made the same transit,  $r_{sp}$  is the learning rate of the transiting species,  $\beta_{rid}$  is an adjustment to  $r_{sp}$  based on the transiting individual,  $b_{sp}$  is the value of  $D$  at the critical point in the Ricker curve, and  $c_{sp}$  modifies the shape of the Ricker curve, such that values close to zero result in a wide peak, and larger values a narrower peak.

## RESULTS

# Results

## Fitting a Generalized Ricker Model

We began our analysis of animal transits between POIs by fitting the differences of the transit length and the linear distance between POIs ( $\Delta L$ ) to a generalized Ricker curve, dependent on the linear distance between POIs ( $D$ ). We set bounds on the parameters to fix the intercept at 0 and ensure the following conditions: 1) predictions of  $\Delta L$  initially increase as  $D$  increases, up until some critical point defined by the first two parameters,  $a$  and  $b$ , which give, respectively,  $\Delta L$  and  $D$  at the critical point, and 2) the shape of the resultant “humped curve” is defined by the third parameter  $c$ , such that the curve widens toward the flat line  $\Delta L = a$  as  $c \rightarrow 0$ , and narrows toward a spike as  $c \rightarrow \text{inf}$ . We predicted that at low values of  $b$ , this model would capture what, from initial visualizations, appeared to be a rapid increase in  $\Delta L$  with increasing  $D$ , followed by a leveling off or decrease in  $\Delta L$  growth rates with  $D$ .

This was not the case: estimates of  $b$  from our baseline model fell well outside observed values of  $D$  (89% CI  $b$  [37.01, 140.07] kilometers). This means that the best fit of the generalized Ricker curve to our data did not use the “humped-curved” features of the model, but instead used just the front part of the curve that increases toward an asymptote at  $\Delta L = a$ . One possible reason for this, statistically, is that our initial assessments of the shape of the data were biased (e.g. by the density of points or the shape of the variance) and an asymptomatic curve would generally provide a better fit to our data than a humped curve. Another is that the expectation of  $\Delta L$  given  $D$  does have a humped curve shape, but this shape cannot be accommodated by the constraints of the generalized Ricker. We will elaborate further on these possibilities

and their meaning in the discussion; for practical purposes, these results influenced the way we further developed the model to account for learning rates. Namely, we noted for values of  $b$  in the posterior distribution, changes to  $a$  would have a stronger and more clearly interpretable effect on predictions, and so chose to model all learning effects as modifications to  $a$ . We opted to continue the analysis using the generalized Ricker instead of switching to a simpler asymptotic curve because the full features of the Ricker curve may have been useful for describing species-specific patterns (they were not), and the failure of these features to prove useful in this context is still informative.

## Rates of Route-Formation

In our final model, there is clear evidence that the transits of all four species become more direct with increasing experience (Figure 3). Species-level posterior distributions of the learning effect, marginalized across individual effects, were entirely below zero and highly overlapping for spider monkeys (mean  $\hat{r} = -0.21$ , 89% CI [-0.26, -0.16]), coatis (mean  $\hat{r} = -0.17$ , 89% CI [-0.2, -0.14]), and kinkajous (mean  $\hat{r} = -0.16$ , 89% CI [-0.19, -0.12]). This means that given our model of transit length, we can be nearly certain these species travel more linearly as they repeat a specific transit over the course of our study. The posterior distribution of the learning effect in capuchins (mean  $\hat{r} = -0.09$ , 89% CI [-0.13, -0.05]) is significantly lower than that of the other species (spider monkeys, 99.88% of posterior samples; coatis 99.25% of posterior samples; kinkajous 97.38% of posterior samples), but still significantly below zero (1% of posterior samples), meaning we can be highly confident capuchin transits become more linear with experience, but at a slower rate than the other species. Recall that  $r$  determines  $\Delta L$  at the critical point of our model according to the term  $\exp(a + r \log(k))$ , where  $k$  represents transit experience, and that this critical point has been estimated at values of  $D$  far outside our data. Thus, the

## DISCUSSION

best way to get a sense of the magnitude of the learning effect at relevant values of  $D$  is visually (Figure 4).

## Discussion

### Route Formation and Implications for Animal Cognition

We find clear evidence that transit lengths between highly revisited locations decrease as a given individual repeats a given transit over the course of our study. There are several possible explanations for this.

The most likely explanation, we believe, is that as the spatial distribution of *D. oleifera* changes - due to the inter-tree variance in peak fruiting period - individual animals learn new routes between currently productive trees.

Such learning could be the result of procedural memory combined with some degree of movement stochasticity, as with the formation of traplines in bees (Reynolds et al. 2013). In other words, when an animal makes a transit that results in the encounter of a productive fruit tree, it is more likely to follow that same sequence of landmarks in the future (via reinforcement learning), with minor variations due to chance. When a more efficient transit is made as a result of such chance, the resultant stimulus-action (landmark-movement) pairs are further reinforced. However, this system *per se* is not particularly flexible, and should produce only gradual changes in daily travel patterns of individuals (Collett et al. 2003). We generally find, in looking at day-to-day paths of animals in our study, that they are highly variable. This suggests that increases in transit linearity over time are made by inferring the direction of a goal relative to a known landmark and orienting accordingly (as described in macaques by Wirth et al. (2017)). At small scales,

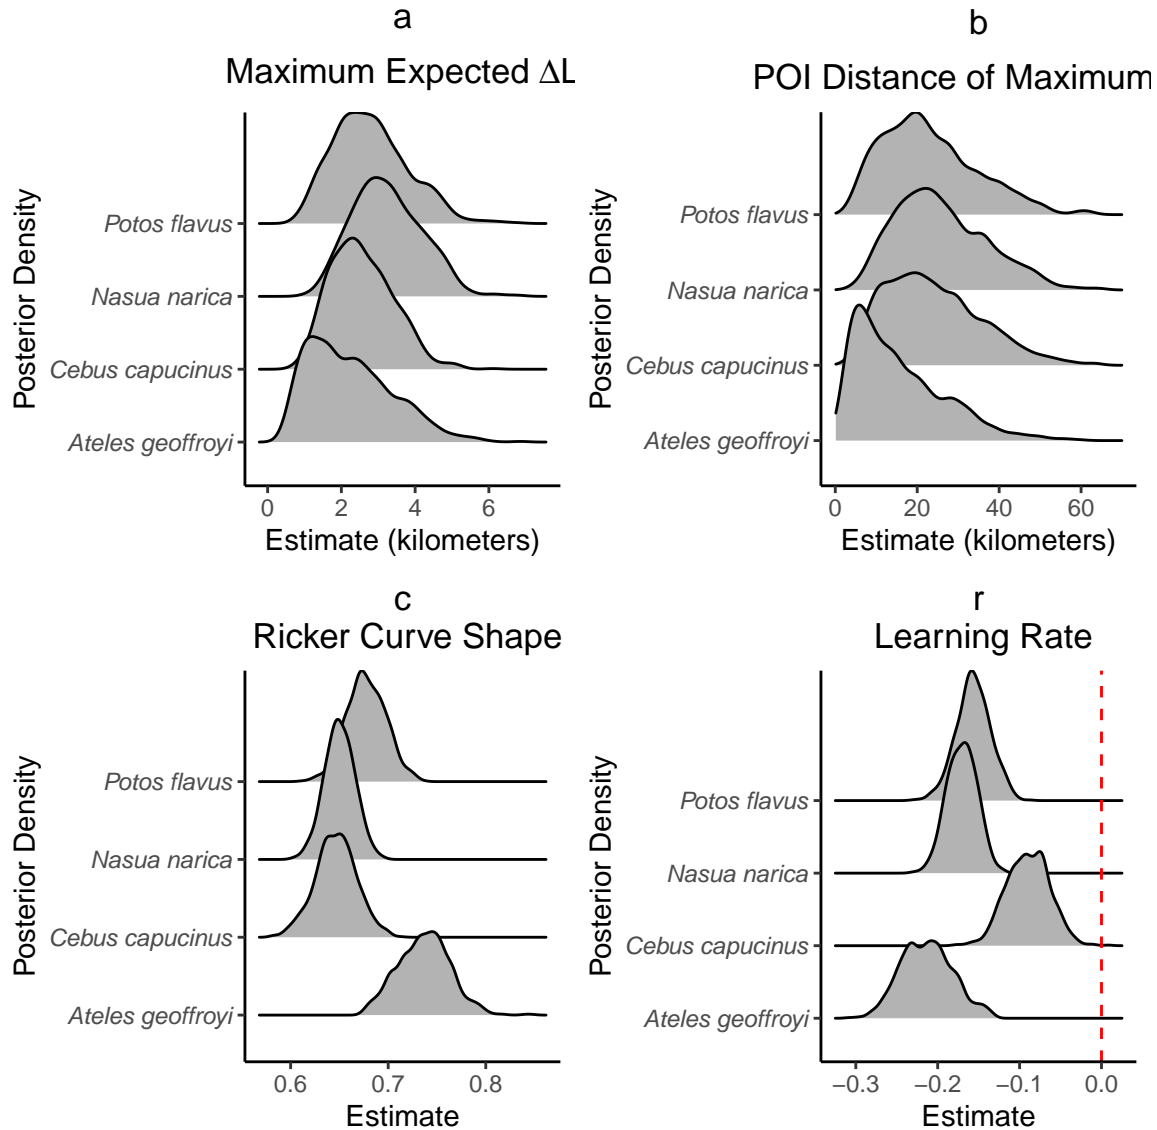


Figure 4.3: Posterior distributions of key model parameters  $a$ ,  $b$ ,  $c$ , and  $r$ , giving (respectively) the maximum expected difference between the straight line path between two POIs and the actual path travelled, the straight line distance between two POIs at which this maximum occurs, the shape of the curve (where higher values mean greater 'peakiness'), and the learning rate, where negative values indicate a reduction in expected travel length with increasing experience.

DISCUSSION




Figure 4.4: Prediction curves of transit length minus POI distance using mean posterior estimates of model parameters, marginalized over individual effects, plotted through the empirical data. We have used a square root transformation on distances and truncated some empirical data points with large observed values in y to better visualize the data.

this could be done through spatial processing of a perceptual field, i.e. an animal finding a shorter path from one landmark to a nearby, intermediary landmark. At a larger scale, this would require an animal to build a topological mental map, for example by using path integration to keep track of its absolute distance and direction from a landmark and remembering this vector when it arrives at a goal (Wang 2016). Though our work does not clearly distinguish the scale at which animals are inferring more linear transits, we posit that the combination of generally flexible space-use with increasingly direct transits is evidence that these animals are using spatial reasoning to optimize their movement during goal-oriented travel. We thus suggest that, at a minimum, coatis, kinkajous, spider monkeys, and capuchins maintain knowledge of multiple important locations in the environments and learn to navigate between them by orienting to a dense array of local landmarks. Whether they optimize their travel paths by inferring direct transits between distant landmarks and goals is an important question for future research.

The most common counter-argument to claims that path linearity is evidence of environmental knowledge is that animals may adopt straight-line travel as a generally efficient strategy for avoiding recently visited foraging patches, without necessarily knowing to which location they are travelling or the path they are taking (Janmaat et al. 2021). Under this counter-hypothesis, our observations that linearity increases over experience would require that motivation for straight-line-travel increases over subsequent transits. Our study occurred across a seasonal transition from food scarcity to greater abundance, which we propose would have the opposite effect: as food becomes more abundant, diverse, and uniformly distributed, the need to avoid recently visited areas decreases and the probability of interspersing straight line travel with bouts of opportunistic foraging increases. It is possible, however, that our analysis has captured an initial increase in motivation to travel in a straight line as *D. oleifera* becomes available, which out-scales the



## *DISCUSSION*

predicted trend in the opposite direction at the end of the season. There may also exist other motivational reasons for straight-line-travel to increase across our study period we have not considered. Thus, though we go one step further than past research in demonstrating the production of environmental knowledge of free-ranging animal movement, we have not entirely removed the confound of motivation from our conclusions.

The shape of our best-fit models offers further insights into the relative roles of knowledge and motivation in determining the length of transits. We find that the relationship between POI distance and transit length takes the shape of an asymptotic curve, such that relative increases to transit length diminish with increasing POI distances. This offers further evidence that animals anticipate or plan their destinations, and increase the linearity of their travel when moving to a far-away goal. We must also consider, however, that animals might alternate between strategies of local-search and linear travel without a specific goal destination (an effective strategy for escaping depleted regions, ala a Lévy walk ([Viswanathan et al. 1999](#))). In this case, bouts of non-linear random-search are more likely to occur between nearby POIs than distant ones. The distribution of error in our data speaks to this reality: variance in transit length is much higher at small POI distances than large ones, a statistical impossibility for a single, positive-real process forced through zero (such as our model). Thus, we believe our data are produced by a mixture of processes, local-search and goal-directed travel among them. The prevalence of these processes could be better disentangled by fitting the data to a function of path linearity over POI distance derived from a mixture of diffusion processes, for example by adding together Brownian diffusion processes that differ in the way their diffusion rates relate to POI distance.

Often, evidence of route-use and path inference are related to the question of whether or not animals have cognitive maps. A theory of cognitive maps, first proposed by Edward

## CHAPTER 4. ROUTE FORMATION

Tolman (Tolman 1948), is well formed in the fields of psychology and neuroscience (see Behrens et al. (2018); Ekstrom and Ranganath (2018); O’Keefe and Nadel (1978); Schiller et al. (2015)). Though novel shortcuts have been used to study the cognitive maps of rats in maze-based experiments, this field generally holds that cognitive maps are not explicitly spatial, but use abstract measures of “distance” and “direction” to connect memories and knowledge. Given this view, the most important evidence of cognitive maps is the use of what has variably been called “vicarious trial and error”, “mental replay and pre-play”, and “episodic memory”. In other words, cognitive maps describe the way an animal uses past information to model and choose between possible futures. Because this model of the cognitive map is so well supported by neurological research into the function of the hippocampus and related brain regions, we hold the view that it is counter-productive to refer to topological mental representations of space as cognitive maps, as is common in behavioral ecology (e.g. Bennett (1996); Dyer (1991); Presotto et al. (2019)), and suggest there is more nuance to studying cognitive maps through movement than demonstrating the use of novel shortcuts.

Instead, we propose that the scale and context of route optimization can be used to quantify the flexibility of animal navigation and compare cognitive maps between individuals and species (see Mikhalevich et al. (2017) for more on this argument). This requires a shift from the thinking that cognitive maps are rare, necessarily high-order systems and toward the view that cognitive maps are common in animals and can develop, with sufficient experience and memory capacity, into complex systems that enable topological inference at large scales. Our analysis suggests that selectively foraging mammals use spatial inference to optimize transits between frequently re-used locations at a time when the distribution of a key fruit resource is shifting dynamically across highly concentrated patches. The useful question, in this context, is not whether these animals have cogni-

## *DISCUSSION*

tive maps, but at what scale their cognitive maps are enabling spatial inference and what causes variation in this process.

### **Inter-Species Differences**

The clearest trend when comparing our model predictions between species is that spider-monkeys increase transit linearity at the fastest rate, followed by coatis and kinkajous, and then capuchins. It is not clear from our analysis whether these differences in learning rates result from making spatial inferences more frequently, or at larger spatial scales. We also note that coatis have longer (i.e. less linear) initial transits than the other species - their faster learning rate, relative to capuchins, resulting in similar transit linearity only when individuals from both species have made the transit multiple times. These results generally comport with the prediction that more fruit-dependent species (spider monkeys and kinkajous) will optimize transit between high value POI's faster, but with only four species there are many other ecological factors that must be considered. Coatis are the most terrestrial of the four species, which may enable greater variance in path linearity - from more tortuous random search to more linear directed travel. This could explain why coatis have initially less linear transits, but are similar in linearity to capuchins when both have high levels of transit experience. Kinkajous are the only nocturnal species in this study; in general, we expect foraging at night would limit an animal's ability to orient using distant landmarks, and thus reduce the scale at which they can make spatial inference. This could explain why kinkajous exhibit route optimization rates less than spider monkeys and similar to coatis, despite their high degree of frugivory. Kinkajous and spider monkeys, in addition to eating the highest percentages of fruit, also travel in smaller groups. Higher rates of route-optimization in these species could reflect an individual's increased ability to use spatial knowledge when there is less need to build

group consensus in movement direction.

This narrative offers a contrast to the social brain hypothesis (Byrne and Whiten 1988, Reader and Laland 2002), which has been popular for decades and posits that primate intelligence evolved largely to keep track of complex social relationships. The species in our study with the least socially complex groups, kinkajous, exhibited rates of route optimization equal or greater than coatis and capuchins, which live in moderate to large, highly cohesive groups. Capuchins, known for maintaining cultural traditions and their use of tools for extractive foraging (Ottoni and Izar 2008, Perry 2011, Barrett et al. 2018), exhibited the lowest rates of path optimization. Our point here is not to refute the idea that capuchins are intelligent, but to note that a less social, non-primate species appears to more flexibly adjust its goal-directed travel to changing resource distributions. Given that many group living animals do not evolve complex social relationships, and diet has been shown to be a better predictor of brain size than social structure in primates (DeCasien et al. 2017), it may be time to flip the thesis of the social brain hypothesis on its head: maybe primates did not evolve intelligence because they were social, but they became social because they evolved intelligence (i.e. they evolved the ability to respond flexibly to changing environments at multiple spatio-temporal scales, and this mechanism also enables them to learn about individual differences and keep track of relationships). It is worth noting that although kinkajous do not live in complex social groups, they are long-lived and may very well develop individual relationships with their neighbors through vocalization, scent-markings, and other means that are largely invisible to us (Kays and Gittleman 2001).

Of course, it is frequently possible to craft a compelling evolutionary narrative when studying fewer species than ecological degrees of freedom. The coherence of our findings with *a priori* predictions is not a strong case for concluding that selectively foraging on

## DISCUSSION

fruit leads to natural selection for cognitive maps that enable spatial inference at large scales, but offers compelling evidence that it *might*. To our knowledge, this is the first study to rigorously compare rates of route optimization between species in an ecological context with large spatial scale. We thus offer a toolkit for analyzing animal cognition in a comparable manner through a combination of telemetry, environmental sensing, and hierarchical modelling. Applying these methods across a broad enough array of species would enable a phylogenetic analysis that more rigorously assesses the roles of diet, sociality, and spatio-temporal niche in the evolution of large scale spatial inference. This work would complement existing strategies such as comparisons of brain size and cognitive test batteries (Malsburg et al. 2020) to paint a fuller picture of how interactions between social structure and resource distribution result in different strategies for storing memories, consolidating knowledge, and predicting the future.

### **Suggestions for future comparative analyses of ecological cognition.**

Janmaat et al. (2021) propose a compelling system of analysis for studying animal cognition through natural movement, which includes 1) comparisons of changes in path linearity with experience, 2) decision analysis as animals move between locations with known, relevant qualities, and 3) informational analysis to quantify patterns of location revisitation. In each case, Janmaat et al. (2021) call for the development of hierarchical statistical methods to compare metrics across species. We have tackled the first case, with applicable lessons for moving forward.

First, we suggest a small adjustment to the linearity index recommended by Janmaat et al. (2021) for quantifying route directedness. This measure (the POI distance divided by the length of the transit) is useful for a cursory comparison of transits with different

straight-line-distances. This, however, washes out any potential effects of POI distance on transit length, which we have noted are both present in our data and a logical product of differing motivations (e.g. random foraging, social cohesion, directed travel) underlying animal movement. Though it is possible to account for these differences by including POI distance as a predictor of the linearity index, properly modelling a fraction as an outcome of its numerator is not straightforward. We use, instead, the difference between straight-line-distance and transit length ( $\Delta L$ ), and suggest that for the statistical modeler this is both simpler and more interpretable.

Second, our work calls for further exploration of the relationship between POI distance and transit length, both to ensure that the transits of animals that tend to travel between distant locations are properly compared to those with shorter transits, and to better identify the underlying factors that result in a non-linear relationship. Because we visually observed a pattern in our data where  $\Delta L$  initially increased, as a function of POI distance, then decreased, and finally leveled out, we focused on exploring “humped-curve” functions to model this relationship. We settled on the generalized Ricker used by Persson et al. (1998) because its parameters usefully describe a critical point that could result from the balancing of a positive effect between predictor and outcome (e.g. the scaling of random-search motivated path lengths as POI distance increases) and a negative effect (e.g. the increasing probability of a transit being goal-directed). However, this model did not fit our data as expected, and the critical point was estimated well outside the range of observed POI distances. This could be because 1) reduction in path linearity at larger POI distances resulting from an increased rate of directed travel does not outweigh the exponentially scaling tortuosity of the large POI distance transits that do occur from random-search, 2) the generalized Ricker does a poor job of modelling the dynamics between goal-directed travel and non-directed foraging (i.e. the humped curve

## *DISCUSSION*

shape of the data is real, but not captured by the constraints of the Ricker curve), or 3) our conceptualization of the processes underlying the relationship between POI distance and  $\Delta L$  is wrong. Though using the generalized Ricker turned out to be useful to us for exploring what the relationship between POI distance and  $\Delta L$  is not, we were unable to make full use of its multiple parameters to derive meaning from the patterns in our data; for those with smaller data sets or the desire to build more appropriate models, we suggest starting with functions that increase positively toward an asymptote.

Third, we suggest there is a need for further theoretical work to understand how different movement processes shape the relationship between POI distance and  $\Delta L$ . This need is motivated most strongly by the strange pattern of error in our data, in which the variance decreases as the POI distance (and predicted  $\Delta L$ ) increase. This inverse relationship between  $\Delta L$  and its variance may explain why visualizing the data gives the impression of a humped curve. It is also anathema to the gamma distribution, or any other distribution that results from a single, positive-real process that runs through the origin. From an applied statistics perspective, one way to address this problem would be to develop mixed models to describe the balance of multiple processes producing transits (e.g. random foraging and directed travel). A promising avenue for determining what such a statistical model might look like is the study of Brownian Bridges. This method calculates probability densities of possible transits given the time and place of the start and end points and the assumption that the underlying movement process is a Brownian random walk. If the same derivations were made for an underlying process that combines multiple Brownian random walks (representing, e.g., deviations from straight-line-directed-travel made due to alternative motivations and deviations made due imperfect knowledge) with magnitudes dependent on start- and end-point distances, the resulting probability densities would describe predicted distributions of  $\Delta L$  relative

to POI distance.

Finally, there are many interesting ways in which the model we have developed can be used as a foundation, and expanded upon to further identify when and where path optimization occurs across animals. Building in parameters to adjust the learning rate based on the properties of a target POI, for example, could illuminate how animal learning is motivated differently by specific outcomes. Similarly, we apply the same learning function to model path optimization across all POI distances, but (like  $\Delta L$ ) learning rates could vary depending on how far apart the POIs are. Appropriate parameters to test this could also be built into our model, and could reveal the scale at which animals make spatial inferences to optimize their routes. These are just a few suggestions - the flexibility of our modelling approach opens the door to comparatively assessing any number of environmental variables that might determine how an animal learns and infers optimal paths between important locations. As these methods develop, and are integrated into existing approaches for comparing animal cognition, a clearer picture will emerge of how animals in complex, natural environments process information to make effective, future oriented decisions.



# Conclusion

I began my graduate research program with the intention to study how a cognitive system evolves such that it is capable of understanding complex societies and building sophisticated technologies, as seen in human behavior. My foundational hypothesis was that the foraging behavior of primates and other species that share their ecology would reveal the basic features of such cognitive systems. The most useful idea I found in the pursuit of this goal is that of the cognitive map. In this conclusion, I will outline the contours of how I understand cognitive maps now, which differs considerably from how I understood cognitive maps when I designed each of my chapters, and which itself differs from how I understood cognitive maps as I wrote the chapters. Once my personal theory of cognitive maps is briefly laid out, I will highlight some of the ways my research helped my views to evolve and discuss some ways I believe that this theory could shape the future of ecological research. Finally, I will consider some of the emerging bridges between the research fields of artificial intelligence and animal cognition that I think have the most potential, especially those that might lead to mathematical models necessary for formally testing conceptual theories of cognitive maps.

## Building a Cognitive Map

The concept of the cognitive map is, at its core, an analogy. In the conclusion to *\*Purposive Behavior in Animals and Men*, Tolman writes

“All science necessarily presents, it seems to us, but a map and picture of reality. If I were to present reality in its whole concreteness, science would be not a map, but a complete replica of reality. And then it would lose its usefulness. For it would have to cover as many pages as does life itself; it would no longer serve as a brief and a handbook. One of the first requisites of a science is, in short, that it be a map, i.e. A short-hand for finding one’s way about from moment of reality to the next—that it be a symbolic compendium by means of which to predict and control. Our account of the mind, we hold, is a map account.”

The central question Tolman attempts to answer in his book is “how do animals make this map?” The framework he presents is rigorous and insightful, but also sprawling, complex, and somewhat archaic - no doubt a reason that its influence in psychology was slow to emerge and often misunderstood. Here, I draw on this framework, as well as more modern accountings and my own insights, to present a simpler outline of the cognitive map’s primary mechanics (as I understand it).

The simplest definition of a cognitive map I can offer is “a hierarchical representation of the state of the world, built from memory by the combined processes of inference and reinforcement.” At the lowest level, inference occurs when an animal determines that one perceptual scene is “like” another. It can then use the outcomes of its behavior in that previous scenario to guide its decisions in the current moment. If the outcomes of

## *BUILDING A COGNITIVE MAP*

the two scenarios are congruous, the shared elements of the scenarios can be stored in a *scheme*. Schema reduce the information in a perceptual scene to a subset of important features, and they are the building blocks of a cognitive map. Learning occurs when an outcome causes the connections between the schema and motor pathways activated prior to that event to become reinforced. In this way, animals learn not only how to respond to specific perceptual cues, but also how to make inferences across increasingly disparate and diverse scenarios by attending to their most important features. The network of associations between schema that an animal builds over time - its semantic knowledge - allows the animal to translate a given perceptual scene into a representation of possible actions and their outcomes - a cognitive map.

My hope, in presenting this highly simplified explanation of a cognitive map, is to highlight how reinforcement and inference are highly integrated processes that build on each other to produce intelligent behavior. This view contrasts with the one I held entering this dissertation that reinforcement and inference form a dichotomy. On one side of such a dichotomy, reinforcement produces stable patterns of action that can only be changed by repeated, countervailing experience. On the other side, inference is a separate cognitive faculty that enables some animals to behave optimally without the need for reinforcement. The research I have presented offers several opportunities to consider why the idea of a cognitive map depicts nuance in the relationship between reinforcement and inference that is necessary for studying the behavior of intelligent animals in complex environments.

Take the differing responses of primates placed into small foraging arrays in Chapter 1, for example. The vervets and Japanese macaques needed no experience to adopt efficient movements between resource platforms. Despite the novelty of the platforms and their arrangement, these primates inferred an efficient way to acquire all of the

## CHAPTER 4. ROUTE FORMATION

food. The strepsirrhines did not. Instead, the strepsirrhines spent many of their initial trials exploring. Only with repeated exposure did strepsirrhines begin showing behaviors that indicated they had inferred, from previous trials, properties of the array: they focused more on the platforms, avoided returning to the same platforms, and in a few cases began to consistently follow efficient routes between all platforms. Because this project lacked an experimental component and there were confounding differences in the nature of the strepsirrhine environments and those with vervets and Japanese macaques, it is impossible to say what drives the observed differences. That the strepsirrhines eventually prove themselves capable of navigating the arrays as efficiently as vervets and Japanese macaques, however, emphasizes the question of why they do not do it in the first place. Perhaps their different sensory organs perceive the scene differently and the food is less salient to them; perhaps they have less prior experience with arrays of depleting resources than the vervets and macaques. Or, perhaps the reinforcement of schema in their cognitive maps is more modular, and less likely to bridge between disparate contexts, even if they share some salient features. Future work could resolve these alternative hypotheses by sequentially introducing individuals of each species to arrays of different scales (eg one-half or twice the sizes used in Chapter 1), and in different contexts (eg. light vs dark, different colored rooms). The differently sized arrays would enable measurement of how learning rates change with the scale of the array, which makes for more robust comparisons between species. Introducing the same individual to arrays in different contexts would allow some control over experience, allowing to measure how individuals generalize their experiences, instead of just guessing.

In Chapter 2, kinkajous face the opposite problem: a familiar environment where the resource array changes. In this case, kinkajous may still benefit from learning movement patterns through the tree crown that optimize foraging success on average (direct

## *BUILDING A COGNITIVE MAP*

reinforcement of cues such as crown structure to motor action). In addition, kinkajous might hypothetically learn strategies that help them efficiently exploit the specific array of flowers on a given night (reinforcing the use of particular schema, e.g. the appearance of flowers on one branch that commonly blooms synchronously with a branch in another region). However, at the scale of our analysis we found neither patterns that could be exploited as such, nor much variation in flower distribution at all. Such absence was reflected in the apparent randomness of kinkajous' movements within the crown.

Analyzing the results of the first two chapters from the perspective of reinforcement and inference illustrates how animals, if they are to exploit the patterns of an ever-changing world, must move beyond simple stimulus-action reinforcement. They need to use schema to make inferences - learning occurs as they reinforce the use of schema that enable better inferences. This type of learning usefully describes the process of developing semantic knowledge from episodic memory. A sense of the cognitive map begins to emerge when considering how an animal combines its knowledge of multiple schema to conceptualize its current context and the actions that would be most useful. The nature and value of the cognitive map becomes clearer when considering it in an explicitly spatial context, as done in Chapters 3 and 4.

The results of Chapter 3 illustrate how kinkajous use both reinforcement and inference to behave efficiently and flexibly throughout space. In some places, the high resolution tracking of Tony Stark reveals incredibly consistent paths through the canopy, indicating the power of reinforcement to develop efficient behavior in a complex environment. Yet, these paths are used in two directions. This could not be achieved by simply reinforcing stimulus-action responses, as the stimuli moving the opposing directions are likely very different. Instead, an animal learning to use a route bi-directionally might reinforce schema along the route that includes schema of where the animal has been. A couple

## CHAPTER 4. ROUTE FORMATION

interesting factors become apparent here. First, the hierarchical nature of a cognitive map; while lower order schema might store the salient features of a simple concept, like a certain type of object, higher order schema organize lower order schema into increasingly complex concepts, such as a place. Second, the use of bi-directional routes demonstrates the role of cognitive maps in allowing an animal to represent the state of the world outside of its current perceptual context by building associations. As Tony Stark learns to use a route in one direction, if the schema he builds are associated with schema of where he has been, he will be better able to infer the utility of using the route in the other direction.

This same idea applies to the intersections in a kinkajou's route network, but with added complexity. These points in space are not just used for two navigational goals, but three or more. In the food provisioning experiment with Molly, she rapidly integrated provisioning locations into her nightly movements and developed more direct routes between these locations. Though route intersections could emerge if different behaviors are reinforced according to different motivations (e.g. food vs. water), this is not sufficient to explain Molly's behavior. Instead, Molly's behavior suggests she has previously reinforced associations between movement decisions at important intersections and particular place schema. If her scheme for a place changes, say because she has recently found food there, it can influence her decision at an intersection previously associated with the place, but not the food. I have taken some liberty in describing the results of this chapter, as the limited data do not rule out many alternative hypotheses, but what I hope to have illustrated is that the flexible navigation of kinkajous can not be explained by magnitudes of reinforcement learning and/or inference alone; it requires a theory of how reinforcement and inference create a feedback loop that can build an increasingly complex system of knowledge like a cognitive map.

If the behavior of the kinkajous tracked in Chapter 3 point toward an integrated role of

## *BUILDING A COGNITIVE MAP*

reinforcement and inference in the development of a cognitive map, Chapter 4 provides evidence of this process in action. The analysis takes the assumption that animals know of a relationship between two locations, but the connections may be indirect or through weakly reinforced schema. If a change occurs in the animal's cognitive map, e.g. the appearance of new food at one of these two locations, and that change motivates the need for a more direct route, the animal may follow the known connections, but use learned strategies for orientation to build more direct and reinforced schema along the way. Examples of this could include learning new landmarks on the physical path that aid in orientation, or cutting landmarks out of the route by finding a direct connections between the landmarks that come before and after. As an animal's route shifts through space, it will need to infer the utility of new places it visits, then reinforce those schema that are most useful. By measuring the rate at which animals build direct routes between locations separated by varying distances, we are able to assess how they balance reinforcement of paths that work with learned strategies for inferring new choices that may produce better results.

In summary, each chapter of this dissertation finds animals attempting to navigate ever-changing worlds, but the factors that change differ from chapter to chapter. In each case, the changing features render simple stimulus-action reinforcement ineffective for producing optimal behavior. Instead, optimal behaviors could be produced by reinforcing the use and associations of schema that mediate between perception and action, allowing for inference. In the first two chapters, these inferences involve strategies for serial decision-making in local, multi-destination arrays. In chapters 3 and 4, the problem expands to environments that are only partially observable, and thus call for a more complex cognitive map that, through hierarchical organization, can build concepts of places and their spatial properties. Across all chapters, the data reveal variance between

individuals and species in how quickly they learn optimal behaviors, and the breadth of contexts between which they can make inferences.

One piece that is largely missing from these studies is the development of cognitive maps over an individual’s lifetime, or across different experimental contexts. The methods of analysis developed in this work opens the door to rigorously conduct such research. These efforts would be greatly supplemented by the formalization of a cognitive map model, as is happening in the field of artificial intelligence, enabling the simulation of learning (and emergent behaviors) across parameter values that balance inference and reinforcement. Taken together, these efforts could integrate theories of learning with ecological theories such as optimal foraging and the ecological constraints model. In the following sections, I outline further details about the directions this research could take.

## Modelling a Cognitive Map

Inference is of primary interest to AI scientists, and is generally studied within the field under the moniker “first-trial learning”. First-trial learning refers to the ability of an agent to behave more optimally upon its first exposure to a new environment if it has previously had experience in a different environment with some shared properties (relative to agents that have no previous experience).

One of the algorithms most proficient in first-trial learning is called Projective Simulation (PS; Briegel and De Las Cuevas (2012)), which draws explicitly on the ideas of episodic memory, schematization, and reinforcement outlined in the previous section. In the infinite color game, for example, a PS agent is presented with a simple perceptual scene (percept) that contains two elements, an arrow with one of a discrete set of  $n$  directions and a random color drawn from the continuous color space. The actions the agent



## MODELLING A COGNITIVE MAP

can take correspond to the directions of the arrows, and the goal is for agents to learn to choose the action that matches the arrow direction in a percept. The first step for the algorithm to achieve this is to create a network of encountered percepts (episodic memory), ‘wildcard generalizations’ (nodes in the graph that replace at least one element of percept vector with a ‘#’ symbol, i.e. schema), and actions. Second, when an agent encounters a new percept (which occurs at every step because every color is unique), it creates edges between that percept’s node in the memory network and any wildcard nodes to which it can be generalized. It then conducts a random walk through the network, starting at the current percept node and ending when it arrives at an action. Finally, when the agent successfully matches an action to the observed arrow direction, the weight of all the network edges it traversed prior to that decision are reinforced, increasing their likelihood of being followed again in the future (Melnikov et al. 2017). Over time, the links between wildcard schema that contain only arrow direction information become strongly linked to the correct actions, and this architecture thus provides a computationally simple means for the agent to derive the relevant properties of a scene, even when no two scenes are ever the same.

Though it is unsurprising to see that principles of psychology, developed in large part through the study of animal behavior, have been influential in the field of Artificial Intelligence (AI), the insights of AI research have been slower to make their way back to behavioral ecology. The most obvious reason for this is that applying these insights to the behavior of real animals in complex environments is among the most difficult class of AI problems. Russell and Norvig (2020) identify seven binaries that categorize the difficulty of an environment in terms of its difficulty for AI to solve, noting that “the hardest case is *partially observable, multiagent, nondeterministic, sequential, dynamic, continuous, and unknown*. Without going into detail about each of these qualities, I will

## CHAPTER 4. ROUTE FORMATION

suffice it to say they clearly describe the conditions of Chapters 3 and 4 (and also the first two chapters, if one considers the role of the external environment in the animals' behavior - e.g. how changing abundance of other food sources affected the time kinkajous spent in the balsa tree studied in Chapter 2). The types of algorithms that can handle these types of environments are new, and require specialized study to implement (Russell and Norvig (2020)'s 1200 page volume is an *introductory* text to the study of these algorithms). The Projective Simulation algorithm in the infinite color game addresses only the 'continuous' problem of challenging environments. Nonetheless, emerging AI algorithms have enormous potential for informing the development of new models relating individual learning to the emergence of group- and population-level phenomena in behavioral ecology.

I believe that the cognitive map concept will be essential for filling this gap. Consider a hypothetical version of Projective Simulation that includes a 'working memory' element of each percept. This working memory might contain a vector for each possible action the agent can take (memory vector), with a magnitude equal to the number of times that action was taken since some event (e.g. a successful foraging action). Like the colors in the infinite color game, the possible vectors that could be held in this element are infinite, but need only be added to the memory network as they are experienced, with links to the wildcard nodes that generalize the vectors away. Such functionality would enable agents to learn about the 'distance' and 'direction' between experiences and schema, the key features of a cognitive map. If the structure of nodes representing memory vectors is further constrained such that the connection between a node in vector  $x$  of magnitude  $i$  and a node in vector  $y$  of magnitude  $j$  is equal to  $r_{xy}|i - j|$ , where the value of  $r_{xy}$  that minimizes an agent's surprise can be learned, I believe agents in a spatial foraging task would begin to use novel shortcuts. The true power of such an approach, however,

is that it need not be applied only to spatial domains; it would allow an agent to learn generally about how its actions change its environment and the relationships between different actions and their consequences.

## **The role of Cognitive Map models in Optimal Foraging**

The notion of an intelligent agent that can solve foraging tasks in the most challenging environmental cases begs the question of how such agents will behave relative to real animals. For example, the Marginal Value Theorem predicts that animals will stop foraging in a patch when the rate of return in that patch falls below the average rate of return in the environment, including the time costs of travel (Charnov 1976). This theorem assumes that an animal has perfect knowledge of the average return rate and distance of foraging patches in its environment, but no knowledge of the locations or values of specific patches. Perfect knowledge of the average environment seems unlikely for animals, and Chapter 4 adds to a large body of evidence that animals have some knowledge of future destinations. None-the-less, early research using relatively simple systems found that foraging behavior matched the predictions of the Marginal Value Theorem in species as diverse as black-capped chickadees (Krebs et al. 1974), great tits (Cowie 1977), hummingbirds (Pyke 1978b), chipmunks (Giraldeau and Kramer 1982), and bumblebees (Pyke 1978a, Cibula and Zimmerman 1984). However, subsequent work found that animals generally overstay in patches (Nonacs 2001), but sometimes leave energy rich patches too early (Alonso et al. 1995). The ability of animals to match Marginal Value Theorem predictions can be explained by learning models that closely approximate the solution using limited memory (Belisle and Cresswell 1997), while deviations from

these predictions have been attributed to animals learning to leave a patch based on multi-scale foraging contexts (e.g. they leave a quality patch sooner if they know it to be in a high quality region) (Watanabe et al. 2014). Deviations for optimal foraging may also result from animals learning to satisfice multiple competing demands, rather than maximize energy intake (Alonso et al. 1995). A formal cognitive map model of learning could explain such higher order learning, and lead to the derivation of new marginal value predictions based on the multi-scale foraging challenges faced by most animals in real environments.

I will conclude by (1) re-iterating the distinction posed in my introduction between cognitive map models of *knowledge* often used in movement ecology and cognitive map models of *learning* used in psychology, (2) emphasizing the reasons my research motivates a need to integrate cognitive map models of learning into ecological research, and (3) briefly speculating on the power of cognitive map models of learning to extend beyond foraging behavior and into higher level processes like social learning and theory of mind.

- (1) In ecological research, cognitive models generally have a spatially explicit structure; they are useful for assessing how animals make decisions under the assumption they already know relevant features of the environment are location. In psychology, cognitive map models the the way a network of ideas is constructed to hold information about the distance and direction seperating these ideas. These models are useful for understanding how an animal can develop spatially explicit knowledge in the first place, and have the added benefit of being useful for explaining learning in non-spatial domains.
- (2) In my dissertation, I have illustrated how movement decisions and patterns of navigation differ in contexts with small and large spatial scales. Yet, I also observe there is a link between the ways animals learn the strategies effective for each

## THE ROLE OF COGNITIVE MAP MODELS IN OPTIMAL FORAGING

context. Cognitive map models of learning can unify accounts of how animals learn heuristics for navigating arrays within local, perceptually available scenes, and how animals learn routes that traverse the continuous space between distant patches. They provide a means of explaining (and predicting) the inferences animals make (e.g. the relevant context of a novel scene, or a shortcut between patches) that are typically baked into cognitive map models of knowledge.

- (3) Finally, cognitive map models of learning should, in theory, allow for learning beyond spatial domains, such as social domains - so long as the agent or animal can perceive the other relevant actors and learn how sequences of its own actions intervene between one state of the world and the next. If true, these models could be integrated into our understanding of how animals respond to each other in their movements, offering insights into how animals achieve an ideal free distribution (Křivan et al. 2008) or modulate group size according to ecological constraints (Chapman 1990).

To achieve these ends, behavioral ecologists and computer scientists will need to bring their domains of knowledge together. Many behavioral ecologists conducting research like my own are building substantial knowledge of what inferences animals make, and when. Computer scientists are increasingly able to explain how intelligent systems learn to make sophisticated inferences by integrating episodic memory and reinforcement learning into hierarchical networks of knowledge. These systems relate perceived world-states according distances and directions along an agent's action space, much like Tolman described cognitive maps. Somewhere between the leading edges of these two fields, there is a predictive model of animal foraging in stochastic, multi-scale environments that does not need to assume animals possess *a-priori* representations of topological space. The route through this space just needs to be mapped.

# Appendix A

## Appendix

This dissertation is supported by several hundred pages of supplementary code, analysis, and figures. These documents, along with a completely reproducible workflow, can be found in an online repository at [github.com/aqvinig/Thesis-patterns\\_of\\_learning](https://github.com/aqvinig/Thesis-patterns_of_learning). Here, we offer a guide to important supplements (included those cited in the preceding chapters) and where to find them within the main repository.

### Chapter 1

#### 1.1 - Chapter 1 Markdown

Location: “Do Primates Trapline/DOCS/Movement Sequence Repetition.Rmd”

Description: The document used to compile Chapter 1. Primarily text and figure embedding. The bulk of analysis for Chapter 1 was completed in other documents (see below).

#### 1.2 - Chapter 1 Data Notebook

## CHAPTER 1

Location: “Do Primates Trapline/CODE/Data\_Notebook.pdf”

Description: A reproducible report on all analyses included in chapter 1. Includes 1) pointers to the code used to clean and prepare data used for the chapter, 2) Detailed written and visual assessment of model validity for the learning rate models used in chapter 1, 3) Code to reproduce all Chapter 1 figures, and 4) detailed analysis of model posterior distributions. The markdown version of this PDF is also available using the .rmd extension.

### 1.3 - Chapter 1 Simulation Validation

Location: “Do Primates Trapline/CODE/Simulation-Validation.html”

Description: Simulations in Chapter 1 were based on the descriptions of a model from another paper (Reynolds et al, 2013), with modifications. In this report, I reproduce the results of the Reynolds paper, validating that differences in our model were the result of our intentional changes and not due to errors or differences in code. I also note how the Reynolds model performs in the novel arrays used in Chapter 1, discussing the importance of our model modifications for these more complex environments. The markdown version is also available using the .rmd extension.

### 1.4 - Initial analysis of DET metric

Location: “Do Primates Trapline/CODE/Primate DET Report.html”

Description: This report details how the properties of the DET metric were explored and understood, leading the selection of analysis parameters used and described in Chapter 1. It includes several insights into the use of the DET metric for studying primate behavior not included in the main body due to their complexity and tangential relevance. Some caution should be used, however, as the models used for analysis in this report are not as comprehensive as those in the final analysis, and the document has not been updated

or modified from its original form.

### 1.5 - Final Analysis of primate sequence DET

Location: “Do Primates Trapline/CODE/DET\_Analysis.html”

Description: A companion to the Data Notebook, this documents provides a complete and reproducible description of the models used to analyze the DET of primate movement sequences through experimental arrays. This includes code, model diagnostics, detailed descriptions of posterior distributions, and written analysis.

## Chapter 2

### 2.1 Chapter 2 Markdown

Location: “Balsa\_Spatial\_Analysis/DOCS/The-movement-and-behavior-of-kinkajous-within-a-dynamic-foraging-hotspot.Rmd”

Description: The analysis and writing of this chapter were integrated. Thus, this document contains all of the code necessary to produce Chapter 2, including the text markdown, analysis, statistical reports, and figures. This also include figures and statistics for model diagnostics and in-depth analysis that are hidden in the final PDF version.

### 2.2 Chapter 2 Reference Simulations

Location: “Balsa\_Spatial\_Analysis/CODE/Reference\_Model1\_Analysis.html”

Description: A detailed description of the simulations used to create reference data, and the code used. The markdown version is also available using the .Rmd extension

### 2.3 Initial Analysis Short Report

Location: “Balsa\_Spatial\_Analysis/CODE/Kinkajou\_Scan\_Analysis.html”



## CHAPTER 3

Description: An initial report on the raw data data used in chapter 2 along-side an early analytical model. Includes many figures included in the main markdown code, but hidden and not described in the output of that document. This version is thus easier to access and read for anybody interested in a more detailed breakdown of the data.

# Chapter 3

## 3.1 Chapter 3 Markdown

Location: “BCI Anniversary/DOCS/Navigation on BCI markdown.Rmd”

Description: Markdown of text and figure embedding for Chapter 3. Analysis for this chapter was not integrated into the markdown body, and can be found in supplementary documents (see below).

## 3.2 Chapter 3 Data Preparation

Location: “BCI Anniversary/CODE/Path\_Segmentation.html”

Description: Detailed descriptions, figures, and code for processing raw data for chapter 3. Also available in markdown using .Rmd extension

## 3.3 Chapter 3 Reproducible Analysis

Location: “BCI Anniversary/DOCS/Data-Notebook.html”

Description: A more thorough analysis of the models and statistics presented in Chapter 3, including model diagnostics and code used to analyze the data and create figures. The markdown version is available in the CODE folder using the .RMD extension.

## Chapter 4

### 4.1 Chapter 4 Markdown

Location: “Segment Linearity/DOCS/manuscript.Rmd”

Description: Integrated markdown of Chapter 4 text, code, and analysis, sufficient to reproduce the full chapter. Many of the more detailed figures are hidden in the final output.

### 4.2 Chapter 4 Data preparation

Location: Segment Linearity/CODE/Path-Clustering-and-Segmentation.Rmd

Description: Code and figures used to process raw data, with descriptions and explanations of parameter settings used for each step in the data processing pipeline.

### 4.3 Chapter 4 Model Development

Location: Segment Linearity/CODE/Gamma-Generalized-Ricker-Model-Development

Description: The final model presented in Chapter 4 is complex. It was developed iteratively, testing each new parameter and its prior distributions to determine its effect on the model's health and the properties of the model when fit to the data. Each model tested is contained in this document, as well as a brief written description of what was learned at each stage.

# Colophon

This document is set in [EB Garamond](#), [Source Code Pro](#) and [Lato](#). The body text is set at 11pt with *lmr*.

It was written in R Markdown and  $\LaTeX$ , and rendered into PDF using [aggiedown](#) and [bookdown](#).

This document was typeset using the XeTeX typesetting system, and the University of California Thesis class. Under the hood, the elements of the document formatting source code have been taken from the [Latex, Knitr, and RMarkdown templates for UC Berkeley's graduate thesis](#), and [Dissertate: a LaTeX dissertation template to support the production and typesetting of a PhD dissertation at Harvard, Princeton, and NYU](#)

The source files for this thesis, along with all the data files, have been organised into a git repository, which is available at <https://github.com/aqvining/Thesis-Patterns-of-Learning>. A hard copy of the thesis can be found in the University of California, Davis library.

This version of the thesis was generated on 2023-08-25 16:47:26. The repository is currently at this commit:

Commit: 98ee595ec3674312339f1743190e841431bcf7ee

Author: Alexander Vining <aqvining@ucdavis.edu>

APPENDIX A. APPENDIX

When: 2023-08-25 13:49:06 GMT

thesis submission

The computational environment that was used to generate this version is as follows:

```
- Session info -----  
setting value  
version R version 4.2.2 (2022-10-31 ucrt)  
os Windows 10 x64 (build 19044)  
system x86_64, mingw32  
ui RTerm  
language (EN)  
collate English_Germany.utf8  
ctype English_Germany.utf8  
tz Europe/Berlin  
date 2023-08-25  
pandoc 2.19.2 @ C:/Program Files/RStudio/resources/app/bin/quarto/bin/tools/ (via  
  
- Packages -----  
! package * version date (UTC) lib source  
abind 1.4-5 2016-07-21 [1] CRAN (R 4.2.0)  
aggiedown * 1.0 2023-06-14 [1] Github (ryanpeek/aggiedown@ae9930  
arrayhelpers 1.1-0 2020-02-04 [1] CRAN (R 4.2.1)  
backports 1.4.1 2021-12-13 [1] CRAN (R 4.2.0)  
bayesplot * 1.10.0 2022-11-16 [1] CRAN (R 4.2.2)
```

## CHAPTER 4

bitops	1.0-7	2021-04-24	[1]	CRAN	(R 4.2.0)
bookdown	0.34	2023-05-09	[1]	CRAN	(R 4.2.3)
broom	1.0.4	2023-03-11	[1]	CRAN	(R 4.2.3)
cachem	1.0.8	2023-05-01	[1]	CRAN	(R 4.2.3)
callr	3.7.3	2022-11-02	[1]	CRAN	(R 4.2.2)
caTools	1.18.2	2021-03-28	[1]	CRAN	(R 4.2.1)
checkmate	2.2.0	2023-04-27	[1]	CRAN	(R 4.2.3)
class	7.3-22	2023-05-03	[2]	CRAN	(R 4.2.3)
classInt	0.4-9	2023-02-28	[1]	CRAN	(R 4.2.3)
cli	3.4.1	2022-09-23	[1]	CRAN	(R 4.2.2)
cmdstanr	* 0.5.3	2022-08-01	[1]	local	
coda	0.19-4	2020-09-30	[1]	CRAN	(R 4.2.1)
codetools	0.2-19	2023-02-01	[2]	CRAN	(R 4.2.2)
colorspace	2.1-0	2023-01-23	[1]	CRAN	(R 4.2.3)
corrplot	* 0.92	2021-11-18	[1]	CRAN	(R 4.2.1)
crayon	1.5.2	2022-09-29	[1]	CRAN	(R 4.2.2)
curl	5.0.0	2023-01-12	[1]	CRAN	(R 4.2.3)
DBI	1.1.3	2022-06-18	[1]	CRAN	(R 4.2.1)
devtools	* 2.4.5	2022-10-11	[1]	CRAN	(R 4.2.2)
digest	0.6.31	2022-12-11	[1]	CRAN	(R 4.2.3)
distributional	0.3.2	2023-03-22	[1]	CRAN	(R 4.2.3)
dplyr	* 1.1.2	2023-04-20	[1]	CRAN	(R 4.2.3)
e1071	1.7-13	2023-02-01	[1]	CRAN	(R 4.2.3)
ellipsis	0.3.2	2021-04-29	[1]	CRAN	(R 4.2.1)
evaluate	0.21	2023-05-05	[1]	CRAN	(R 4.2.3)
fansi	1.0.4	2023-01-22	[1]	CRAN	(R 4.2.3)

APPENDIX A. APPENDIX

farver	2.1.1	2022-07-06 [1] CRAN (R 4.2.1)
fastmap	1.1.1	2023-02-24 [1] CRAN (R 4.2.3)
foreign	0.8-84	2022-12-06 [2] CRAN (R 4.2.2)
fs	1.6.2	2023-04-25 [1] CRAN (R 4.2.3)
generics	0.1.3	2022-07-05 [1] CRAN (R 4.2.1)
geosphere	* 1.5-18	2022-11-15 [1] CRAN (R 4.2.1)
ggdist	3.3.0	2023-05-13 [1] CRAN (R 4.2.3)
ggforce	* 0.4.1	2022-10-04 [1] CRAN (R 4.2.2)
ggplot2	* 3.4.2	2023-04-03 [1] CRAN (R 4.2.3)
ggridges	* 0.5.4	2022-09-26 [1] CRAN (R 4.2.2)
git2r	* 0.32.0	2023-04-12 [1] CRAN (R 4.2.3)
glue	1.6.2	2022-02-24 [1] CRAN (R 4.2.1)
gplots	* 3.1.3	2022-04-25 [1] CRAN (R 4.2.1)
gridExtra	2.3	2017-09-09 [1] CRAN (R 4.2.1)
gtable	0.3.3	2023-03-21 [1] CRAN (R 4.2.3)
gtools	3.9.4	2022-11-27 [1] CRAN (R 4.2.3)
highr	0.10	2022-12-22 [1] CRAN (R 4.2.3)
hms	1.1.3	2023-03-21 [1] CRAN (R 4.2.3)
htmltools	0.5.5	2023-03-23 [1] CRAN (R 4.2.3)
htmlwidgets	1.6.2	2023-03-17 [1] CRAN (R 4.2.3)
httpuv	1.6.11	2023-05-11 [1] CRAN (R 4.2.3)
httr	1.4.6	2023-05-08 [1] CRAN (R 4.2.3)
inline	0.3.19	2021-05-31 [1] CRAN (R 4.2.1)
jpeg	* 0.1-10	2022-11-29 [1] CRAN (R 4.2.2)
jsonlite	1.8.4	2022-12-06 [1] CRAN (R 4.2.3)
kableExtra	* 1.3.4	2021-02-20 [1] CRAN (R 4.2.2)

## CHAPTER 4

KernSmooth	2.23-21	2023-05-03	[2]	CRAN	(R 4.2.3)
knitr	1.43	2023-05-25	[1]	CRAN	(R 4.2.3)
labeling	0.4.2	2020-10-20	[1]	CRAN	(R 4.2.0)
later	1.3.1	2023-05-02	[1]	CRAN	(R 4.2.3)
lattice	0.21-8	2023-04-05	[2]	CRAN	(R 4.2.3)
lifecycle	1.0.3	2022-10-07	[1]	CRAN	(R 4.2.2)
loo	2.6.0	2023-03-31	[1]	CRAN	(R 4.2.3)
lubridate	* 1.9.2	2023-02-10	[1]	CRAN	(R 4.2.3)
lwgeom	* 0.2-13	2023-05-22	[1]	CRAN	(R 4.2.3)
magrittr	2.0.3	2022-03-30	[1]	CRAN	(R 4.2.1)
maptools	1.1-6	2022-12-14	[1]	CRAN	(R 4.2.3)
MASS	7.3-60	2023-05-04	[2]	CRAN	(R 4.2.3)
Matrix	1.5-4.1	2023-05-18	[1]	CRAN	(R 4.2.3)
matrixStats	0.63.0	2022-11-18	[1]	CRAN	(R 4.2.3)
memoise	2.0.1	2021-11-26	[1]	CRAN	(R 4.2.1)
mime	0.12	2021-09-28	[1]	CRAN	(R 4.2.0)
miniUI	0.1.1.1	2018-05-18	[1]	CRAN	(R 4.2.1)
modelr	* 0.1.11	2023-03-22	[1]	CRAN	(R 4.2.3)
move	* 4.1.12	2023-01-20	[1]	CRAN	(R 4.2.3)
munsell	0.5.0	2018-06-12	[1]	CRAN	(R 4.2.1)
mvtnorm	1.1-3	2021-10-08	[1]	CRAN	(R 4.2.0)
pillar	1.9.0	2023-03-22	[1]	CRAN	(R 4.2.3)
pkgbuild	1.4.1	2023-06-14	[1]	CRAN	(R 4.2.2)
pkgconfig	2.0.3	2019-09-22	[1]	CRAN	(R 4.2.1)
pkgload	1.3.2	2022-11-16	[1]	CRAN	(R 4.2.1)
polyclip	1.10-4	2022-10-20	[1]	CRAN	(R 4.2.1)

APPENDIX A. APPENDIX

posterior	1.4.1	2023-03-14	[1]	CRAN	(R 4.2.3)
prettyunits	1.1.1	2020-01-24	[1]	CRAN	(R 4.2.1)
processx	3.8.1	2023-04-18	[1]	CRAN	(R 4.2.3)
profvis	0.3.8	2023-05-02	[1]	CRAN	(R 4.2.3)
promises	1.2.0.1	2021-02-11	[1]	CRAN	(R 4.2.1)
proxy	0.4-27	2022-06-09	[1]	CRAN	(R 4.2.1)
ps	1.7.5	2023-04-18	[1]	CRAN	(R 4.2.3)
purrr	* 1.0.1	2023-01-10	[1]	CRAN	(R 4.2.3)
R6	2.5.1	2021-08-19	[1]	CRAN	(R 4.2.1)
raster	* 3.6-20	2023-03-06	[1]	CRAN	(R 4.2.3)
rbbt	* 0.0.0.9000	2023-05-26	[1]	Github	(paleolimbot/rbbt@a7053f7)
RColorBrewer	* 1.1-3	2022-04-03	[1]	CRAN	(R 4.2.0)
Rcpp	1.0.10	2023-01-22	[1]	CRAN	(R 4.2.3)
D RcppParallel	5.1.7	2023-02-27	[1]	CRAN	(R 4.2.3)
readr	* 2.1.4	2023-02-10	[1]	CRAN	(R 4.2.3)
remotes	2.4.2	2021-11-30	[1]	CRAN	(R 4.2.1)
rethinking	* 2.21	2022-08-29	[1]	Github	(rmcelreath/rethinking@783)
rgdal	* 1.6-6	2023-04-18	[1]	CRAN	(R 4.2.3)
rgeos	0.6-3	2023-05-24	[1]	CRAN	(R 4.2.3)
rlang	1.1.1	2023-04-28	[1]	CRAN	(R 4.2.3)
rmarkdown	2.22	2023-06-01	[1]	CRAN	(R 4.2.3)
rstan	* 2.26.13	2022-06-25	[1]	local	
rstudioapi	0.14	2022-08-22	[1]	CRAN	(R 4.2.2)
rvest	1.0.3	2022-08-19	[1]	CRAN	(R 4.2.2)
scales	* 1.2.1	2022-08-20	[1]	CRAN	(R 4.2.1)
sessioninfo	1.2.2	2021-12-06	[1]	CRAN	(R 4.2.1)



## CHAPTER 4

sf	* 1.0-13	2023-05-24	[1]	CRAN	(R 4.2.3)
shape	1.4.6	2021-05-19	[1]	CRAN	(R 4.2.0)
shiny	1.7.4	2022-12-15	[1]	CRAN	(R 4.2.3)
sp	* 1.6-0	2023-01-19	[1]	CRAN	(R 4.2.3)
StanHeaders	* 2.26.25	2023-05-17	[1]	CRAN	(R 4.2.3)
stringi	1.7.12	2023-01-11	[1]	CRAN	(R 4.2.2)
stringr	* 1.5.0	2022-12-02	[1]	CRAN	(R 4.2.3)
svglite	2.1.1	2023-01-10	[1]	CRAN	(R 4.2.3)
svUnit	1.0.6	2021-04-19	[1]	CRAN	(R 4.2.1)
systemfonts	1.0.4	2022-02-11	[1]	CRAN	(R 4.2.1)
tensorA	0.36.2	2020-11-19	[1]	CRAN	(R 4.2.0)
terra	1.7-29	2023-04-22	[1]	CRAN	(R 4.2.3)
tibble	* 3.2.1	2023-03-20	[1]	CRAN	(R 4.2.3)
tidybayes	* 3.0.4	2023-03-14	[1]	CRAN	(R 4.2.3)
tidybayes.rethinking	* 3.0.0	2022-08-30	[1]	Github	(mjskay/tidybayes.rethinki
tidyr	* 1.3.0	2023-01-24	[1]	CRAN	(R 4.2.3)
tidyselect	1.2.0	2022-10-10	[1]	CRAN	(R 4.2.2)
timechange	0.2.0	2023-01-11	[1]	CRAN	(R 4.2.3)
tweenr	2.0.2	2022-09-06	[1]	CRAN	(R 4.2.1)
tzdb	0.4.0	2023-05-12	[1]	CRAN	(R 4.2.3)
units	* 0.8-2	2023-04-27	[1]	CRAN	(R 4.2.3)
urlchecker	1.0.1	2021-11-30	[1]	CRAN	(R 4.2.1)
usethis	* 2.1.6	2022-05-25	[1]	CRAN	(R 4.2.1)
utf8	1.2.3	2023-01-31	[1]	CRAN	(R 4.2.3)
V8	4.3.0	2023-04-08	[1]	CRAN	(R 4.2.3)
vctrs	0.6.2	2023-04-19	[1]	CRAN	(R 4.2.3)

APPENDIX A. APPENDIX

viridisLite	0.4.2	2023-05-02	[1] CRAN (R 4.2.3)
webshot	0.5.4	2022-09-26	[1] CRAN (R 4.2.2)
withr	2.5.0	2022-03-03	[1] CRAN (R 4.2.1)
xfun	0.39	2023-04-20	[1] CRAN (R 4.2.3)
xml2	1.3.4	2023-04-27	[1] CRAN (R 4.2.3)
xtable	1.8-4	2019-04-21	[1] CRAN (R 4.2.1)
yaml	2.3.7	2023-01-23	[1] CRAN (R 4.2.3)

[1] C:/Users/avining/AppData/Local/R/win-library/4.2

[2] C:/Program Files/R/R-4.2.2/library

D -- DLL MD5 mismatch, broken installation.

-----

# References

- 10 Ackerman, J. D., M. R. Mesler, K. L. Lu, and A. M. Montalvo. 1982. Food-Foraging Behavior of Male Euglossini (Hymenoptera : Apidae ): Vagabonds or Trapliners? *Biotropica* 14:241–248.
- Aikman, J. M. 1955. The Ecology of Balsa (*Ochroma lagopus* Swartz) in Ecuador 62.
- Al Toufailia, H., C. Grüter, and F. L. W. Ratnieks. 2013. **Persistence to Unrewarding Feeding Locations by Honeybee Foragers (*Apis mellifera*): the Effects of Experience, Resource Profitability and Season.** *Ethology* 119:1096–1106.
- Alavi, S. E., A. Q. Vining, D. Caillaud, B. T. Hirsch, R. W. Havmøller, L. W. Havmøller, and R. Kays. 2022. **A Quantitative Framework for Identifying Patterns of Route-Use in Animal Movement Data** 9:1–16.
- Alonso, J. C., J. A. Alonso, L. M. Bautista, and R. Muñoz-Pulido. 1995. **Patch use in cranes: a field test of optimal foraging predictions.** *Animal Behaviour* 49:1367–1379.
- Alves-Costa, C. P., G. A. Da Fonseca, and C. Christófaro. 2004. Variation in the diet of the brown-nosed coati (*Nasua nasua*) in Southeastern Brazil. *Journal of Mammalogy* 85:478–482.
- Ancrenaz, M., I. Lackman-Ancrenaz, and N. Mundy. 1994. **Field Observations of Aye-**

- Ayes (Daubentonia madagascariensis) in Madagascar*. Folia Primatologica 62:22–36.
- Anderson, D. J. 1983. **Optimal foraging and the traveling salesman**. Theoretical Population Biology 24:145–159.
- Andriamasimanana, M. 1994. Ecoethological Study of Free-Ranging Aye-Ayes (*Daubentonia madagascariensis*) in Madagascar. Folia Primatologica:37–45.
- Asensio, N., A. Korstjens, C. Schaffner, and F. Aureli. 2008. **Intragroup aggression, fission–fusion dynamics and feeding competition in spider monkeys**. Behaviour 145:983–1001.
- Ashton, B. J., A. Thornton, and A. R. Ridley. 2018. **An intraspecific appraisal of the social intelligence hypothesis**. Philosophical Transactions of the Royal Society B: Biological Sciences 373.
- Ayers, C. A., P. R. Armsworth, and B. J. Brosi. 2015. **Determinism as a statistical metric for ecologically important recurrent behaviors with trapline foraging as a case study**. Behavioral Ecology and Sociobiology 69:1395–1404.
- Ayers, C. A., P. R. Armsworth, and B. J. Brosi. 2018. **Statistically testing the role of individual learning and decision-making in trapline foraging**. Behavioral Ecology 29:885–893.
- Barrett, B. J., C. M. Monteza-Moreno, T. Dogandžić, N. Zwyns, A. Ibáñez, and M. C. Crofoot. 2018. **Habitual stone-tool-aided extractive foraging in white-faced capuchins, *Cebus capucinus***. Royal Society Open Science 5.
- Behrens, T. E. J., T. H. Muller, J. C. R. Whittington, S. Mark, A. B. Baram, K. L. Stachenfeld, and Z. Kurth-Nelson. 2018. **What Is a Cognitive Map? Organizing Knowledge for Flexible Behavior**. Neuron 100:490–509.

## References

- Belisle, C., and J. Cresswell. 1997. The Effects of a Limited Memory Capacity on Foraging Behavior. *Theoretical Population Biology* 52:78–90.
- Bennett, A. T. D. 1996. **Do animals have cognitive maps?** *The Journal of Experimental Biology* 199:219–224.
- Bitterman, M. E. 1965. Phyletic Differences in Learning. *American Psychologist* 20:396–410.
- Boyer, D., and P. D. Walsh. 2010. **Modelling the mobility of living organisms in heterogeneous landscapes: Does memory improve foraging success?** *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 368:5645–5659.
- Bracis, C., K. L. Bildstein, and T. Mueller. 2018. **Revisitation analysis uncovers spatio-temporal patterns in animal movement data.** *Ecography* 41:1801–1811.
- Braveman, N. S., and L. Katz. 1971. **Spatial and Visual Probability Learning in the Kinkajou.** *Perceptual and Motor Skills* 33:1115–1121.
- Briegel, H. J., and G. De Las Cuevas. 2012. **Projective simulation for artificial intelligence.** *Scientific Reports* 2.
- Bruorton, M. R., C. L. Davis, and M. R. Perrin. 1991. **Gut microflora of vervet and samango monkeys in relation to diet.** *Applied and Environmental Microbiology* 57:573–578.
- Buatois, A., and M. Lihoreau. 2016. **Evidence of trapline foraging in honeybees.** *Journal of Experimental Biology* 219:2426–2429.
- Burdon, R. C. F., R. A. Raguso, R. J. Gegear, E. C. Pierce, A. Kessler, and A. L. Parachnowitsch. 2020. **Scented nectar and the challenge of measuring honest signals**

- in pollination. *Journal of Ecology* 108:2132–2144.
- Bureš, J., O. Burešová, and L. Nerad. 1992. *Can rats solve a simple version of the traveling salesman problem?* *Behavioural Brain Research* 52:133–142.
- Buzsáki, G. 2005. *Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory.* *Hippocampus* 15:827–840.
- Byrne, R., and A. Whiten. 1988. *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans.* Clarendon Press, Oxford, United Kingdom.
- Caillaud, D., M. C. Crofoot, S. V. Scarpino, P. A. Jansen, C. X. Garzon-Lopez, A. J. S. Winkelhagen, S. A. Bohlman, and P. D. Walsh. 2010. *Modeling the spatial distribution and fruiting pattern of a key tree species in a neotropical forest: Methodology and potential applications.* *PLoS ONE* 5.
- Cancelliere, E. C., C. A. Chapman, D. Twinomugisha, and J. M. Rothman. 2018. *The nutritional value of feeding on crops: Diets of vervet monkeys in a humanized landscape.* *African Journal of Ecology* 56:160–167.
- Carpenter, B., A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell. 2017. *Stan: A probabilistic programming language.* *Journal of statistical software* 76.
- Chaib, S., M. Dacke, W. Wcislo, E. Warrant, S. Chaib, M. Dacke, W. Wcislo, and E. Warrant. 2021. *Dorsal landmark navigation in a Neotropical nocturnal bee.* *Current Biology* 31:3601–3605.e3.
- Chapman, C. A. 1990. *Ecological Constraints on Group Size in Three Species of Neotropical Primates.* *Folia Primatologica* 55:1–9.

## References

- Chapman, C. A., and L. J. Chapman. 2000. Determinants of Group Size in Primates: The Importance of Travel Costs. Pages 24–42 *in* S. Boinski and P. A. Garber, editors. *On the Move: How and Why Animals Travel in Groups*. The University of Chicago Press, Chicago.
- Chapman, C. A., L. J. Chapman, and R. L. McLaughlin. 1989. **Multiple central place foraging by spider monkeys: travel consequences of using many sleeping sites**. *Oecologia* 79:506–511.
- Charnov, E. L. 1976. **Optimal foraging: The marginal value theorem**. *Theoretical Population Biology* 9:129–136.
- Chausseil, M. 1992. **Evidence for color vision in procyonides: Comparison between diurnal coatis (*Nasua*) and nocturnal kinkajous (*Potos flavus*)**. *Animal Learning & Behavior* 20:259–265.
- Chemero, A. 2013. **Radical Embodied Cognitive Science**. *Review of General Psychology* 17:145–150.
- Cibula, D. A., and M. Zimmerman. 1984. **The Effect of Plant Density on Departure Decisions: Testing the Marginal Value Theorem Using Bumblebees and *Delphinium Nelsonii***. *Oikos* 43:154.
- Clayton, N. S. 2017. Episodic-like memory and mental time travel in animals. Pages 227–243 *in* J. Call, G. M. Burghardt, I. M. Pepperberg, C. T. Snowdon, and T. Zentall, editors. *APA handbook of comparative psychology: Perception, learning, and cognition*. American Psychological Association.
- Collett, M., L. Chittka, and Thomas S. Collett. 2013. **Spatial Memory in Insect Navigation**. *Current Biology* 23:R789–R800.

- Collett, M., and T. S. Collett. 2006. [Insect Navigation: No Map at the End of the Trail?](#) *Current Biology* 16:R48–R51.
- Collett, T. S., P. Graham, and V. Durier. 2003. [Route learning by insects.](#) *Current Opinion in Neurobiology* 13:718–725.
- Cook, W. J., D. L. Applegate, R. E. Bixby, and V. Chvatal. 2011. *The traveling salesman problem: a computation study.* Princeton University Press.
- Cowie, R. J. 1977. [Optimal foraging in great tits \(\*Parus major\*\).](#) *Nature* 268:137–139.
- Cramer, A. E., and C. R. Gallistel. 1997. [Vervet monkeys as travelling salesmen.](#) *Nature* 387:464–464.
- Crofoot, M. C., D. I. Rubenstein, A. S. Maiya, and T. Y. Berger-Wolf. 2011. [Aggression, grooming and group-level cooperation in white-faced capuchins \(\*Cebus capucinus\*\): insights from social networks.](#) *American Journal of Primatology* 73:821–833.
- Crystal, D. J., and T. Suddendorf. 2019. [Episodic memory in nonhuman animals?](#) *Current Biology* 29:R1291–R1295.
- Cunningham, E. P., D. Edmonds, L. Stalter, and J. N. Malvin. 2021. [Ring-tailed lemurs \(\*Lemur catta\*\) use olfaction to locate distant fruit.](#) *American Journal of Physical Anthropology* 175:300–307.
- Dammhahn, M., and P. M. Kappeler. 2008. [Comparative Feeding Ecology of Sympatric \*Microcebus berthae\* and \*M. murinus\*.](#) *International Journal of Primatology* 29:1567–1589.
- De Groeve, J., N. Van de Weghe, N. Ranc, T. Neutens, L. Ometto, O. Rota-Stabelli, and F. Cagnacci. 2016. [Extracting spatio-temporal patterns in animal trajectories: An ecological application of sequence analysis methods.](#) *Methods in Ecology and*



## References

- Evolution 7:369–379.
- DeCasien, A. R., and J. P. Higham. 2019. **Primate mosaic brain evolution reflects selection on sensory and cognitive specialization.** *Nature Ecology and Evolution* 3:1483–1493.
- DeCasien, A. R., S. A. Williams, and J. P. Higham. 2017. **Primate brain size is predicted by diet but not sociality.** *Nature Ecology and Evolution* 1:1–7.
- Decker, J. H., A. R. Otto, N. D. Daw, and C. A. Hartley. 2016. **From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning.** *Psychological Science* 27:848–858.
- Dew, J. L., and P. Wright. 1998. **Frugivory and Seed Dispersal by Four Species of Primates in Madagascar’s Eastern Rain Forest<sup>1</sup>.** *Biotropica* 30:425–437.
- Di Fiore, A. N., A. N. Link, and J. L. Dew. 2008. Diets of wild spider monkeys. Pages 81–137 *in* C. J. Campbell, editor. *Spider monkeys: Behavior, ecology, and evolution of the genus Ateles.* Cambridge University Press, Cambridge.
- Di Fiore, A., and S. A. Suarez. 2007. **Route-based travel and shared routes in sympatric spider and woolly monkeys: Cognitive and evolutionary implications.** *Animal Cognition* 10:317–329.
- Dietrich, W. E. 1982. Geology, climate and hydrology of Barro Colorado island. Pages 21–46 *in* A. S. Rand, D. M. Windsor, and E. G. Leigh, editors. *The ecology of tropical forest: seasonal rhythms and long-term changes.* Smithsonian Institution, Washington D.C.
- Donner, Y., and J. L. Hardy. 2015. **Piecewise power laws in individual learning curves.** *Psychonomic Bulletin and Review* 22:1308–1319.

- Dragoi, G., and S. Tonegawa. 2011. **Preplay of future place cell sequences by hippocampal cellular assemblies**. *Nature* 469:397–401.
- Dyer, F. C. 1991. **Bees acquire route-based memories but not cognitive maps in a familiar landscape**. *Animal Behaviour* 41:239–246.
- Ekstrom, A. D., and C. Ranganath. 2018. **Space, time, and episodic memory: The hippocampus is all over the cognitive map**. *Hippocampus* 28:680–687.
- Erickson, C. J., S. Nowicki, L. Dollar, and N. Goehring. 1998. Percussive Foraging: Stimuli for Prey Location by Aye-Ayes (*Daubentonia madagascariensis*). *International Journal of Primatology* 19:111–122.
- Fagan, W. F., M. A. Lewis, M. Auger-Méthé, T. Avgar, S. Benhamou, G. Breed, L. Ladage, U. E. Schlägel, W. W. Tang, Y. P. Papastamatiou, J. Forester, and T. Mueller. 2013. **Spatial memory and animal movement**. *Ecology Letters* 16:1316–1329.
- Fietz, J., and J. U. Ganzhorn. 1999. **Feeding ecology of the hibernating primate *Cheirogaleus medius*: how does it get so fat?** *Oecologia* 121:157–164.
- Fietz, J., F. Tataruch, K. Dausmann, and J. Ganzhorn. 2003. **White adipose tissue composition in the free-ranging fat-tailed dwarf lemur (*Cheirogaleus medius*; Primates), a tropical hibernator**. *Journal of Comparative Physiology B* 173:1–10.
- Fooden, J., and A. Mitsuru. 2006. **Systematic Review of Japanese Macaques, *Macaca fuscata* (Gray, 1870)**. *Fieldiana Zoology* 104:1–198.
- Ford, L. S., and R. S. Hoffmann. 1988. *Potos flavus*. *Mammalian Species*:1–9.
- Foster, R. 1982. The seasonal rhythm of fruitfall on Barro Colorado Island. Pages 151–172 *in* E. Leigh, A. Rand, and D. Windsor, editors. *The ecology of a tropical forest: seasonal rhythms and long-term changes*. Smithsonian Institution Press.

## References

- Fragaszy, D. M., E. Visalberghi, and L. M. Fedigan. 2004. The complete capuchin: the biology of the genus *Cebus*. Cambridge University Press, Cambridge.
- Garber, P. A. 1988. **Foraging Decisions During Nectar Feeding by Tamarin Monkeys (*Saguinus mystax* and *Saguinus fuscicollis*, Callitrichidae, Primates) in Amazonian Peru**. *Biotropica* 20:100.
- Garrison, J. S. E., and C. L. Gass. 1999. **Response of a traplining hummingbird to changes in nectar availability**. *Behavioral Ecology* 10:714–725.
- Génin, F. 2003. **Female dominance in competition for gum trees in the Grey Mouse Lemur *Microcebus murinus***. *Revue d'Écologie (La Terre et La Vie)* 58:397–410.
- Gigerenzer, G., and P. M. Todd. 1999. Fast and frugal heuristics: The adaptive toolbox. Pages 3–34 *Simple Heuristics That Make Us Smart*. Oxford University Press.
- Giraldeau, L.-A., and D. L. Kramer. 1982. **The marginal value theorem: A quantitative test using load size variation in a central place forager, the Eastern chipmunk, *Tamias striatus***. *Animal Behaviour* 30:1036–1042.
- Go, M. 2010. **Seasonal changes in food resource distribution and feeding sites selected by Japanese macaques on Koshima Islet, Japan**. *Primates* 51:149–158.
- Gompper, M. E. 1997. Population ecology of the white-nosed coati (*Nasua narica*) on Barro Colorado Island, Panama. *Journal of Zoology* 241:441–455.
- González-Zamora, A., V. Arroyo-Rodríguez, Ó. M. Chaves, S. Sánchez-López, K. E. Stoner, and P. Riba-Hernández. 2009. **Diet of spider monkeys (*Ateles geoffroyi*) in Mesoamerica: current knowledge and future directions**. *American Journal of Primatology* 71:8–20.
- Hahsler, M., M. Piekenbrock, and D. Doran. 2019. **dbscan: Fast density-based clustering**

- with R. *Journal of Statistical Software* 91:1–30.
- Hanya, G. 2004. *Diet of a Japanese Macaque Troop in the Coniferous Forest of Yakushima*. *International Journal of Primatology* 25:55–71.
- Hardie, J. L., and C. R. Cooney. 2023. *Sociality, ecology and developmental constraints predict variation in brain size across birds*. *Journal of Evolutionary Biology* 36:144–155.
- Harel, R., S. Alavi, A. M. Ashbury, J. Aurisano, T. Berger-Wolf, G. H. Davis, B. T. Hirsch, U. Kalbitzer, R. Kays, K. Mclean, C. L. Núñez, A. Vining, Z. Walton, R. W. Havmøller, and M. C. Crofoot. 2022. *Life in 2.5D: Animal Movement in the Trees*. *Frontiers in Ecology and Evolution* 10.
- Harten, L., A. Katz, A. Goldshtein, M. Handel, and Y. Yovel. 2020. The ontogeny of a mammalian cognitive map in the real world. *Science* 369:194–197.
- Heidegger, M. 1946. Letter on humanism. *Basic writings* 204:189–242.
- Heidegger, M. 1962. *Being and time*. (J. Macquarrie and E. Robinson, Trans.). Seventh. New York.
- Hill, D. A. 1997. *Seasonal variation in the feeding behavior and diet of Japanese macaques (Macaca fuscata yakui) in lowland forest of Yakushima*. *American Journal of Primatology* 43:305–320.
- Hirsch, B. T. 2009. *Seasonal Variation in the Diet of Ring-Tailed Coatis (Nasua nasua) in Iguazu, Argentina*. *Journal of Mammalogy* 90:136–143.
- Holekamp, K. E., and S. Benson-Amram. 2017. *The evolution of intelligence in mammalian carnivores*. *Interface Focus* 7.

## References

- Hopkins, M. E. 2011. [Mantled Howler \(Alouatta palliata\) Arboreal Pathway Networks: Relative Impacts of Resource Availability and Forest Structure](#). *International Journal of Primatology* 32:238–258.
- Hopkins, M. E. 2016. [Mantled howler monkey spatial foraging decisions reflect spatial and temporal knowledge of resource distributions](#). *Animal Cognition* 19:387–403.
- Jaber, M. Y. 2015. [Learning and Forgetting Models and Their Applications](#). Pages 535–566 *in* A. B. Badiru, editor. *Handbook of Industrial and Systems Engineering*. 2nd edition. CRC Press.
- Jacobs, L. F., and F. Schenk. 2003. [Unpacking the Cognitive Map: The Parallel Map Theory of Hippocampal Function](#). *Psychological Review* 110:285–315.
- Jang, H., C. Boesch, R. Mundry, S. D. Ban, and K. R. L. Janmaat. 2019. [Travel linearity and speed of human foragers and chimpanzees during their daily search for food in tropical rainforests](#). *Scientific Reports* 9:1–13.
- Janmaat, K. R. L. 2019. [What animals do not do or fail to find: A novel observational approach for studying cognition in the wild](#). *Evolutionary Anthropology: Issues, News, and Reviews* 8:1–18.
- Janmaat, K. R. L., J. Collet, R. W. Byrne, and B. Robira. 2021. [Using natural travel paths to infer and compare primate cognition in the wild](#). *iScience*.
- Janson, C. 2000. Spatial movement strategies : theory , evidence , and challenges. Pages 165–203 *in* S. Boinski and P. A. Garber, editors. *On the Move: How and Why Animals Travel in Groups*.
- Janson, C. H. 1998. [Experimental evidence for spatial memory in foraging wild capuchin monkeys, Cebus apella](#). *Animal Behaviour* 55:1229–1243.

- Janson, C. H. 2014. **Death of the traveling salesman Primates do not show clear evidence of multi-step Route Planning.** American Journal of Primatology 76:410–420.
- Janson, C. H. 2016. **Capuchins, space, time and memory: An experimental test of what-where-when memory in wild monkeys.** Proceedings of the Royal Society B: Biological Sciences 283.
- Janzen, D. H. 1971. **Euglossine bees as long-distance pollinators of tropical plants.** Science 171:203–205.
- Josselyn, S. A., and S. Tonegawa. 2020. **Memory engrams: Recalling the past and imagining the future.** Science (New York, N.Y.) 367.
- Joyce, M. M. 2021. **Spatial Movement Decisions Among Foraging Japanese Monkeys (*Macaca fuscata*).** MSc, Concordia University, Montréal, Québec, Canada.
- Julien-Laferrriere, D. 1993. **Radio-tracking observations on ranging and foraging patterns by kinkajous (*Potos flavus*) in French Guiana.** Journal of Tropical Ecology 9:19–32.
- Julien-Laferrrière, D. 1999. **Foraging strategies and food partitioning in the neotropical frugivorous mammals *Caluromys philander* and *Potos flavus*.** Journal of Zoology 247:71–80.
- Kaigaishi, Y., M. Nakamichi, and K. Yamada. 2019. **High but not low tolerance populations of Japanese macaques solve a novel cooperative task.** Primates 60:421–430.
- Kay, S. L., J. W. Fischer, A. J. Monaghan, J. C. Beasley, R. Boughton, T. A. Campbell, S. M. Cooper, S. S. Ditchkoff, S. B. Hartley, J. C. Kilgo, S. M. Wisely, A. C. Wyckoff, K. C. VerCauteren, and K. M. Pepin. 2017. **Quantifying drivers of wild pig movement across multiple spatial and temporal scales.** Movement Ecology 5:1–15.
- Kays, R. 1999. **Food preferences of kinkajous (*Potos flavus*): A frugivorous carnivore.**

## References

- Journal of Mammalogy 80:589–599.
- Kays, R. 2003. Social polyandry and promiscuous mating in a primate-like carnivore: the kinkajou (*Potos flavus*). Pages 125–136 in U. H. Reichard and C. Boesch, editors. Monogamy: mating strategies and partnerships in birds, humans and other mammals. Cambridge University Press.
- Kays, R., and J. L. Gittleman. 2001. **The social organization of the kinkajou *Potos flavus* (Procyonidae)**. Journal of the Zoological Society of London 253:491–504.
- Kays, R., B. T. Hirsch, D. Caillaud, R. Mares, S. Alavi, R. W. Havmøller, and M. C. Crofoot. in review. Multi-scale movement syndromes for comparative of animal patterns.
- Kays, R., M. E. Rodríguez, L. M. Valencia, R. Horan, R. Adam, and C. Ziegler. 2012. Animal Visitation and Pollination of Flowering Balsa Trees (*Ochroma pyramidale*) in Panama 16:66–68.
- Knauer, A. C., and F. P. Schiestl. 2015. **Bees use honest floral signals as indicators of reward when visiting flowers**. Ecology Letters 18:135–143.
- Knierim, J. J., and D. A. Hamilton. 2011. **Framing spatial cognition: Neural representations of proximal and distal frames of reference and their roles in navigation**. Physiological Reviews 91:1245–1279.
- Koepfli, K.-P., M. E. Gompper, E. Eizirik, C.-C. Ho, L. Linden, J. E. Maldonado, and R. K. Wayne. 2007. **Phylogeny of the Procyonidae (Mammalia: Carnivora): Molecules, morphology and the Great American Interchange**. Molecular Phylogenetics and Evolution 43:1076–1095.
- Krebs, J. R., J. C. Ryan, and E. L. Charnov. 1974. **Hunting by expectation or optimal**

- foraging? A study of patch use by chickadees. *Animal Behaviour* 22:953–IN3.
- Křivan, V., R. Cressman, and C. Schneider. 2008. *The ideal free distribution: A review and synthesis of the game-theoretic perspective*. *Theoretical Population Biology* 73:403–425.
- Krochmal, A. R., T. C. Roth, and N. T. Simmons. 2021. *My way is the highway: the role of plasticity in learning complex migration routes*. *Animal Behaviour* 174:161–167.
- Kumpan, L. T., J. M. Rothman, C. A. Chapman, and J. A. Teichroeb. 2019. *Playing it safe? Solitary vervet monkeys (Chlorocebus pygerythrus) choose high-quality foods more than those in competition*. *American Journal of Primatology* 81.
- Lahann, P. 2007. *Feeding ecology and seed dispersal of sympatric cheirogaleid lemurs (Microcebus murinus, Cheirogaleus medius, Cheirogaleus major) in the littoral rain-forest of south-east Madagascar*. *Journal of Zoology* 271:88–98.
- Lambert, J. E., V. Fellner, E. McKenney, and A. Hartstone-Rose. 2014. *Binturong (Arctictis binturong) and kinkajou (Potos flavus) digestive strategy: Implications for interpreting frugivory in carnivora and primates*. *PLoS ONE* 9.
- Leutgeb, S., J. K. Leutgeb, M. B. Moser, and E. I. Moser. 2005. *Place cells, spatial maps and the population code for memory*. *Current Opinion in Neurobiology* 15:738–746.
- Lihoreau, M., L. Chittka, S. C. Le Comber, and N. E. Raine. 2012a. *Bees do not use nearest-neighbour rules for optimization of multi-location routes*. *Biology Letters* 8:13–16.
- Lihoreau, M., L. Chittka, and N. E. Raine. 2011. *Trade-off between travel distance and prioritization of high-reward sites in traplining bumblebees*. *Functional Ecology* 25:1284–1292.



## References

- Lihoreau, M., N. E. Raine, A. M. Reynolds, R. J. Stelzer, K. S. Lim, A. D. Smith, J. L. Osborne, and L. Chittka. 2012b. [Radar Tracking and Motion-Sensitive Cameras on Flowers Reveal the Development of Pollinator Multi-Destination Routes over Large Spatial Scales](#). *PLoS Biology* 10:19–21.
- Lynn Carpenter, F. 1978. [Hooks for mammal pollination?](#) *Oecologia* 35:123–132.
- Malsburg, J. H. D., P. M. Kappeler, and C. Fichtel. 2020. Linking ecology and cognition : does ecological specialisation predict cognitive test performance ? *Behavioral Ecology* 74.
- Marden, J. H., and K. D. Waddington. 1981. [Floral choices by honeybees in relation to the relative distances to flowers](#). *Physiological Entomology* 6:431–435.
- McElreath, R. 2020. *Rethinking: An R package for fitting and manipulating Bayesian Models*.
- McLean, K. A., A. M. Trainor, G. P. Asner, M. C. Crofoot, M. E. Hopkins, C. J. Campbell, R. E. Martin, D. E. Knapp, and P. A. Jansen. 2016. [Movement patterns of three arboreal primates in a Neotropical moist forest explained by LiDAR-estimated canopy structure](#). *Landscape Ecology* 31:1849–1862.
- Melnikov, A. A., A. Makmal, V. Dunjko, and H. J. Briegel. 2017. [Projective simulation with generalization](#). *Scientific Reports* 7:14430.
- Menzel, R. 2021. [A short history of studies on intelligence and brain in honeybees](#). *Apidologie* 52:23–34.
- Menzel, R., K. Lehmann, G. Manz, J. Fuchs, M. Koblöfsky, and U. Greggers. 2012. [Vector integration and novel shortcutting in honeybee navigation](#). *Apidologie* 43:229–243.

## References

- Mikhalevich, I., R. Powell, and C. Logan. 2017. **Is behavioural flexibility evidence of cognitive complexity? How evolution can inform comparative cognition.** *Interface Focus* 7.
- Milton, K. 1981. **Distribution Patterns of Tropical Plant Foods as an Evolutionary Stimulus to Primate Mental Development.** *American Anthropologist* 83:534–548.
- Mosdosy, K. N., A. D. Melin, and L. M. Fedigan. 2015. **Quantifying seasonal fallback on invertebrates, pith, and bromeliad leaves by white-faced capuchin monkeys (*Cebus capucinus*) in a tropical dry forest: Capuchin Fallback Foods in a Seasonal Dry Forest.** *American Journal of Physical Anthropology* 158:67–77.
- Najera, D. A. 2009. **Redefining Honeybee Foraging Cognition.** Doctor of Philosophy, University of Kansas.
- Nakagawa, N., Nakamichi, Masayuki, and H. Sugiura, editors. 2010. **The Japanese Macaques.** Springer.
- Nascimento, F. F., M. Oliveira-Silva, G. Veron, J. Salazar-Bravo, P. R. Gonçalves, A. Langguth, C. R. Silva, and C. R. Bonvicino. 2017. **The Evolutionary History and Genetic Diversity of Kinkajous, *Potos flavus* (Carnivora, Procyonidae).** *Journal of Mammalian Evolution* 24:439–451.
- Nonacs, P. 2001. **State dependent behavior and the Marginal Value Theorem.** *Behavioral Ecology* 12:71–83.
- Noser, R., and R. W. Byrne. 2010. **How do wild baboons (*Papio ursinus*) plan their routes? Travel among multiple high-quality food sources with inter-group competition.** *Animal Cognition* 13:145–155.
- O’Keefe, J. 1976. **Place units in the hippocampus of the freely moving rat.** *Experimental*

## References

- Neurology 51:78–109.
- O’Keefe, J., and L. Nadel. 1978. *The Hippocampus as a Cognitive Map*. Oxford University Press.
- Ohashi, K., A. Leslie, and J. D. Thomson. 2008. *Trapline foraging by bumble bees: V. Effects of experience and priority on competitive performance*. Behavioral Ecology 19:936–948.
- Oppenheimer, J. R. 1968. Behavior and ecology of the white-faced monkey, *cebus capucinus* on Barro Colorado Island. University of Illinois.
- Ossi, K., and J. M. Kamilar. 2006. *Environmental and phylogenetic correlates of Eulemur behavior and ecology (Primates: Lemuridae)*. Behavioral Ecology and Sociobiology 61:53–64.
- Ottoni, E. B., and P. Izar. 2008. *Capuchin monkey tool use: Overview and implications*. Evolutionary Anthropology: Issues, News, and Reviews 17:171–178.
- Perry, S. 2011. *Social traditions and social learning in capuchin monkeys (Cebus)*. Philosophical Transactions of the Royal Society B: Biological Sciences 366:988–996.
- Persson, L., K. Leonardsson, A. M. De Roos, M. Gyllenberg, and B. Christensen. 1998. *Ontogenetic Scaling of Foraging Rates and the Dynamics of a Size-Structured Consumer-Resource Model*. Theoretical Population Biology 54:270–293.
- Powell, L. E., K. Isler, and R. A. Barton. 2017. *Re-evaluating the link between brain size and behavioural ecology in primates*. Proceedings of the Royal Society B: Biological Sciences 284:1–8.
- Presotto, A., R. Fayrer-Hosken, C. Curry, and M. Madden. 2019. *Spatial mapping shows that some African elephants use cognitive maps to navigate the core but not*

- the periphery of their home ranges. *Animal Cognition* 22:251–263.
- Pyke, G. H. 1978a. **Optimal foraging: Movement patterns of bumblebees between inflorescences**. *Theoretical Population Biology* 13:72–98.
- Pyke, G. H. 1978b. **Optimal Foraging in Hummingbirds: Testing the Marginal Value Theorem**. *American Zoologist* 18:739–752.
- Pyke, G. H. 1981. **Optimal foraging in hummingbirds: Rule of movement between inflorescences**. *Animal Behaviour* 29:889–896.
- Pyke, G. H. 1984. **Optimal Foraging Theory**. *Annual Review of Ecological Systems* 15:523–575.
- Randimbiharirinirina, D. R., B. M. Raharivololona, M. T. R. Hawkins, C. L. Frasier, R. R. Culligan, T. M. Sefczek, R. Randriamampionona, and E. E. Louis. 2018. **Behaviour and Ecology of Male Aye-Ayes (*Daubentonia madagascariensis*) in the Kianjavato Classified Forest, South-Eastern Madagascar**. *Folia Primatologica* 89:123–137.
- Reader, S. M., and K. N. Laland. 2002. **Social intelligence, innovation, and enhanced brain size in primates**. *Proceedings of the National Academy of Sciences* 99:4436–4441.
- Redshaw, J., and T. Suddendorf. 2020. **Temporal Junctures in the Mind**. *Trends in Cognitive Sciences* 24:52–64.
- Reid, A. K., and J. E. R. Staddon. 1998. **A Dynamic Route Finder for the Cognitive Map**. *Psychological Review* 105:585–601.
- Reynolds, A. M., M. Lihoreau, and L. Chittka. 2013. **A Simple Iterative Model Accurately Captures Complex Trapline Formation by Bumblebees Across Spatial Scales and Flower Arrangements**. *PLoS Computational Biology* 9:1–10.

## References

- Riotte-Lambert, L., S. Benhamou, and S. Chamaillé-Jammes. 2017. [From randomness to traplining: a framework for the study of routine movement behavior](#). *Behavioral Ecology* 28:280–287.
- Rosati, A. G. 2017. [Foraging Cognition: Reviving the Ecological Intelligence Hypothesis](#). *Trends in Cognitive Sciences* 21:691–702.
- Russell, S., and P. Norvig. 2020. *Artificial Intelligence: A Modern Approach*. 4th edition. Pearson.
- Russo, F. 2022. *Technoscientific Practices: An Informational Approach*. Rowman & Littlefield.
- Ryser-Welch, P., J. F. Miller, and S. Asta. 2015. [Generating Human-readable Algorithms for the Travelling Salesman Problem using Hyper-Heuristics](#). Pages 1067–1074 *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation*. ACM, Madrid Spain.
- Schiller, D., H. Eichenbaum, E. A. Buffalo, L. Davachi, D. J. Foster, and S. Leutgeb. 2015. [Memory and Space : Towards an Understanding of the Cognitive Map](#). *The Journal of Neuroscience* 35:13904–13911.
- Sefczek, T. M., Z. J. Farris, and P. C. Wright. 2012. [Aye-Aye \(\*Daubentonia madagascariensis\*\) Feeding Strategies at Ranomafana National Park, Madagascar: An Indirect Sampling Method](#). *Folia Primatologica* 83:1–10.
- Séguigne, M., O. Coutant, B. Bouton, L. Picart, É. Guilbert, and P.-M. Forget. 2022. [Arboreal camera trap reveals the frequent occurrence of a frugivore-carnivore in neotropical nutmeg trees](#). *Scientific Reports* 12:7513.
- Siekman, J., and M. W. Crocker. 2010. [Resource-adaptive cognitive processes](#). Pages

- 1–10 in M. W. Crocker and J. Siekmann, editors. Resource-Adaptive Cognitive Processes. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Stan Development Team. 2023. [RStan: the R interface to Stan](#).
- Stephens, D. W. 2008. [Decision ecology: Foraging and the ecology of animal decision making](#). Cognitive, Affective, & Behavioral Neuroscience 8:475–484.
- Suddendorf, T., and M. C. Corballis. 2007. [The Evolution of Foresight What Is Mental Time Travel and Is It Unique To Humans?](#) Behavioral and Brain Sciences 30:299–351.
- Team, R. C. 2020. R: A Language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.
- Teichroeb, J. A. 2015. [Vervet monkeys use paths consistent with context-specific spatial movement heuristics](#). Ecology and Evolution 5:4705–4716.
- Teichroeb, J. A., and W. D. Aguado. 2016. [Foraging vervet monkeys optimize travel distance when alone but prioritize high-reward food sites when in competition](#). Animal Behaviour 115.
- Teichroeb, J. A., and C. A. Chapman. 2014. [Sensory information and associative cues used in food detection by wild vervet monkeys](#). Animal Cognition 17:517–528.
- Teichroeb, J. A., and E. A. Smeltzer. 2018. [Vervet monkey \(\*Chlorocebus pygerythrus\*\) behavior in a multi-destination route: Evidence for planning ahead when heuristics fail](#). PLOS ONE 13:e0198076.
- Teichroeb, J. A., and A. Q. Vining. 2019. [Navigation strategies in three nocturnal lemur species: diet predicts heuristic use and degree of exploratory behavior](#). Animal Cognition 22:343–354.

## References

- Thomson, J. D., M. Slatkin, and B. A. Thomson. 1997. **Trapline foraging by bumble bees: II. Definition and detection from sequence data.** *Behavioral Ecology* 8:199–210.
- Thorén, S., F. Quietzsch, D. Schwochow, L. Sehen, C. Meusel, K. Meares, and U. Radespiel. 2011. **Seasonal Changes in Feeding Ecology and Activity Patterns of Two Sympatric Mouse Lemur Species, the Gray Mouse Lemur (*Microcebus murinus*) and the Golden-brown Mouse Lemur (*M. ravelobensis*), in Northwestern Madagascar.** *International Journal of Primatology* 32:566–586.
- Toledo, S., D. Shohami, I. Schiffner, E. Lourie, Y. Orchan, Y. Bartan, and R. Nathan. 2020. **Cognitive map-based navigation in wild bats revealed by a new high-throughput wildlife tracking system.** *Science* 193:188–193.
- Tolman, E. C. 1932. *Purposive behavior in animals and men.* Univ of California Press.
- Tolman, E. C. 1948. **Cognitive maps in rats and men.** *Psychological Review* 55:189–208.
- Townsend-Mehler, J. M. 2010. *Decision-Making in a Changing Environment: A Look at the Foraging Behavior of Honeybees and Bumblebees as they Respond to Shifts in Resource Availability.* Doctor of Philosophy, Michigan State University.
- Trapanese, C., H. Meunier, and S. Masi. 2018. **What, where and when: spatial foraging decisions in primates.** *Biological Reviews*.
- Trapanese, C., H. Meunier, and S. Masi. 2022. **Do primates flexibly use spatio-temporal cues when foraging?** *Quarterly Journal of Experimental Psychology* 75:232–244.
- Tulving, E. 2005. **Episodic Memory and Autonoesis: Uniquely Human?** Pages 3–56 *in* H. S. Terrace and J. Metcalfe, editors. *The missing link in performance.*
- Turner, S. E., L. M. Fedigan, H. Nobuhara, T. Nobuhara, H. D. Matthews, and M. Nakamichi. 2008. **Monkeys with disabilities: prevalence and severity of congenital**

- limb malformations in *Macaca fuscata* on Awaji Island. *Primates* 49:223–226.
- Turner, S. E., T. Nobuhara, H. Nobuhara, Nakamichi, Masayuki, and S. M. Reader. 2018. Disability and dominance rank in adult female and male Japanese macaques (*Macaca fuscata*). Pages 133–155 in U. Kalbitzer and K. M. Jack, editors. *Primate Life Histories, Sex Roles, and Adaptability: Essays in Honour of Linda M. Fedigan*. Springer International Publishing, Cham.
- Turner, T. R., F. Anapol, and C. J. Jolly. 1997. Growth, development, and sexual dimorphism in vervet monkeys (*Cercopithecus aethiops*) at four sites in Kenya. *American Journal of Physical Anthropology* 103:19–35.
- van der Kooi, C. J., M. Vallejo-Marín, and S. D. Leonhardt. 2021. Mutualisms and (A)symmetry in Plant–Pollinator Interactions. *Current Biology* 31:R91–R99.
- Viswanathan, G. M., S. V. Buldyrev, S. Havlin, M. G. E. Da Luz, E. P. Raposo, and H. E. Stanley. 1999. Optimizing the success of random searches. *Nature* 401:911–914.
- Wang, R. F. 2016. Building a cognitive map by assembling multiple path integration systems. *Psychonomic Bulletin and Review* 23:692–702.
- Watanabe, Y. Y., M. Ito, and A. Takahashi. 2014. Testing optimal foraging theory in a penguin–krill system. *Proceedings of the Royal Society B: Biological Sciences* 281:20132376.
- Weber, S. N., and H. Sprekeler. 2018. Learning place cells, grid cells and invariances with excitatory and inhibitory plasticity:1–41.
- Wehner, R. 2003. Desert ant navigation: how miniature brains solve complex tasks. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology* 189:579–588.



- Wickham, H., M. Averick, J. Bryan, W. Chang, L. McGowan, R. François, G. Grolemund, A. Hayes, L. Henry, J. Hester, M. Kuhn, T. Pedersen, E. Miller, S. Bache, K. Müller, J. Ooms, D. Robinson, D. Seidel, V. Spinu, K. Takahashi, D. Vaughan, C. Wilke, K. Woo, and H. Yutani. 2019. [Welcome to the Tidyverse](#). *Journal of Open Source Software* 4:1686.
- Williams, H. J., and K. Safi. 2021. [Certainty and integration of options in animal movement](#). *Trends in Ecology & Evolution* 36:990–999.
- Wirth, S., P. Baraduc, A. Planté, S. Pinède, and J.-R. Duhamel. 2017. [Gaze-informed, task-situated representations of space in primate hippocampus during virtual navigation](#). *PLOS Biology* 15:1–28.
- Wrangham, R. W., and P. G. Waterman. 1981. [Feeding Behaviour of Vervet Monkeys on \*Acacia tortilis\* and \*Acacia xanthophloea\*: With Special Reference to Reproductive Strategies and Tannin Production](#). *The Journal of Animal Ecology* 50:715.
- Wright, S. J., C. Carrasco, O. Calder, and S. Paton. 1999. The El Niño Southern Oscillation, Variable Fruit Production, and Famine in a Tropical Forest. *Ecological Society of America* 80:1632–1674.
- Wu, Y., H. Wang, H. Wang, and E. A. Hadly. 2017. [Rethinking the Origin of Primates by Reconstructing Their Diel Activity Patterns Using Genetics and Morphology](#). *Scientific Reports* 7:11837.
- Zimmer, H. D., S. Münzer, and K. Umla-Runge. 2010. [Visuo-spatial Working Memory as a Limited Resource of Cognitive Processing](#). Pages 13–34 *in* M. W. Crocker and J. Siekmann, editors. *Resource-Adaptive Cognitive Processes*. Springer Berlin Heidelberg, Berlin, Heidelberg.