

UC Merced

UC Merced Electronic Theses and Dissertations

Title

The Role of Spatial Structure and Memory in Human Foraging

Permalink

<https://escholarship.org/uc/item/7338n813>

Author

Kerster, Bryan Elvis

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, MERCED

The Role of Spatial Structure and Memory in Human Foraging

A dissertation submitted in partial satisfaction of the requirements

for the degree Doctor of Philosophy

in

Cognitive and Information Sciences

by

Bryan Elvis Kerster

Committee in charge:

Professor Christopher Kello, Chair
Professor David Noelle
Professor Jeffrey Yoshimi

Copyright
Bryan Kerster, 2016
All rights reserved

The Dissertation of Bryan Elvis Kerster is approved, and it is acceptable in quality and form for
publication on microfilm and electronically:

Professor Chris Kello, Chair

Professor David Noelle

Professor Jeffrey Yoshimi

University of California, Merced

2016

I dedicate this dissertation to Harold Kerster and Kurt Wille.

Table of Contents

List of Figures	vii
List of Tables	ix
Acknowledgements	x
Curriculum Vitae	xi
Abstract	xiii
Introduction.....	1
Lévy Flights Theory.....	1
Optimal Foraging Theory	3
Area Restricted Search.....	4
Foraging in Cognition.....	4
Experiment 1	9
Methods.....	10
Results.....	13
Discussion.....	19
Experiment 2.....	20
Methods.....	20
Results.....	22
Discussion.....	25
Computational Model	26

Experiment 3.....	30
Results.....	32
Dyadic Search Model.....	35
Results.....	37
Conclusion	42
Future Directions	44
References.....	45
Appendix A.....	48

List of Figures

Figure 1: Three example patterns of movements created by Lévy flights. m in this figure refers to the μ parameter	2
Figure 2: Example movement through an environment with two types of patches.....	5
Figure 3: View of the simulation from (Wolfe, 2013).....	6
Figure 4: Example images from the task in play. The top demonstrates the zoomed out view. The bottom figure shows the zoomed in view with the ship near a resource.....	11
Figure 5: Pseudocode for the algorithm used to generate the resource distributions.	12
Figure 6 Example distributions with varying levels of clustering. .05 (top left), .15 (top right), .35 (bottom left), and .5 (bottom right).	13
Figure 7: Examples of movements produced by participants. Red lines indicate movements, and the pink squares demonstrate the locations where the players zoomed in. The top figure demonstrates a typical sweep strategy, the bottom a more haphazard strategy.	15
Figure 8: Demonstration of how scores varied across different resource distributions. Strong performers improved when resources were more clustered (left), and there is little effect of quantity (right)	16
Figure 9: The effect of resource distributions on mean path lengths. Generally better performers have shorter path lengths, although they increase with resource clustering	17
Figure 10: A typical example of a path length distribution in log/log coordinates, along with its fit line (green). The rounding at the beginning of the distribution occurs across a range of a fairly small number of pixels, and is likely a result of the lack of precision people have when making very small mouse movements.	18
Figure 11: Regression slopes as a function of quantity, clustering, and performance category	19
Figure 12: An example distribution produced by a participant in experiment two. The boxes indicate the area searched on a click. The yellow dots indicate the locations of resources and were not visible to participants.....	22
Figure 13: The path length distributions from all participants sorted by distribution category.	23
Figure 14: Demonstration of how α correlates with score at different levels of clustering (left) and density (right).	24
Figure 15: Normalized time since last click (z score) plotted against log binned path lengths. Path lengths increase with the time since the last resource was found across all conditions.	25
Figure 16: Plot comparing the path length distributions of the model and humans across each resource distribution.	27

Figure 17: Time since last resource found plotted by path length for both humans and the model. The model follows the same trend as humans where path lengths increase as a function of the time since a resource was last found.	28
Figure 18: An example movement created by the model. Red indicates a visited location where a negative trace was left. Green indicates a positive trace. The yellow dots indicate the actual locations of resources.....	29
Figure 19: Demonstration of how α correlates with score at different levels of clustering (left) and density (right) for the model.	30
Figure 20: An example of how experiment 3 appeared during gameplay. Squares indicate where the player has foraged thus far. White squares indicate nothing was found at that location, and red squares indicate that resources were found, with brighter shades indicating higher values.....	31
Figure 21: Path length distributions for each trial. The results are very similar to the earlier experiments and demonstrate the same power law like distribution.....	32
Figure 22: Average α values across each condition.....	34
Figure 23: Demonstration of how α correlates with score at different levels of clustering (left) and density (right). Unlike previous experiments there is no clear correlation across resource density.....	37
Figure 24: The dyadic model is identical in function to the single model with one exception. The models take turns selecting locations, and locations selected by other agents are treated by the non-moving model as having been visited with a trace value equal to that model's attraction value parameter.	35
Figure 25: Examples of model movements. In order the images are examples of the attract, repel, follow, and control conditions. For each image yellow dots indicate the location of resources, and blue and red squares indicate the locations each agent visited. For the follow condition red is the follower and blue is repelled.	38
Figure 26: Logarithmically binned plots of the path length distributions produced by the models. The interaction conditions are in order: attract, repel, follow attracted, follow repelled, and control.....	40
Figure A1: Example model run with high β and moderate-to-high γ values. The movement cost becomes overly strong and forces movements to mostly remain near one location.	48
Figure A2: Example model run with high γ and low-to-moderate β . High gain makes the probability distribution overly peaked, which causes the model to make movements among a small number of overly dense clusters.....	49
Figure A3: Example model run with low γ paired with any β values. Low gain makes the probability distribution too flat, which causes movements to become random and hence not sufficiently clustered.....	50

List of Tables

Table 1: Table comparing search performance and estimated exponents and their standard errors for the humans and model.....	30
Table 2: Performance of the model across levels of density and clustering in terms of the percent of resources found, and degree of clustering as measured by the estimated α value from a pareto distribution.	33
Table 3: Percent of the increase due to each of the two changes made between Experiment 2 and 3, across density and clustering.	34
Table 4: Clustering of each dyadic condition as measured by estimated α	41
Table 5: Average performance and efficiency for each dyadic condition. For the solo condition efficiency is measured over the first 150 steps. (All other conditions only have 150 steps). Values are also normalized for easier interpretation.	41

Acknowledgements

First, I would like to acknowledge my committee, their support and guidance was always appreciated and helpful. A meeting with my committee was always useful and informative, and my abilities as a researcher and scientist owe a great deal to your influence. You have all been a part of my education since before graduate school and the lessons I have learned from you will last the rest of my life.

I would like to Chris Kello, for being my advisor and mentor, for seeing my potential, and giving me this opportunity. You always allowed to find projects that matched my interests and my skills, and found time to offer advice and support, even when you were beyond busy. We would frequently have rousing discussions about the next step in a project or the meaning of a result, and as a result of those discussions and arguments we were able to elevate the quality of the research we did together in a way I don't think either of us could have done on our own. It has been a pleasure to work with you for these last few years and I cannot thank you enough for all you have done for me.

I would like to thank Jeff Yoshimi. I still remember the lecture you gave in CORE when I was a freshman, and without that influence I might never have chosen Cognitive Science. I am also thankful for the time we spent together to guide my papers, and generate experimental ideas. Your passion and insight have given me a new appreciation for philosophy.

I would like to thank Dave Noelle. Your ability to see to the heart of any problem and identify the flaws or weaknesses in an approach was always remarkable, and my research has been much improved by your insight. Your ability to explain complex and technical subjects in a way that was somehow understandable without losing details is also a talent I am thankful for. I often speak to my peers about the usefulness of your courses, and years later the frequency with which I make those recommendations continues to increase.

I would also like to thank the rest of the faculty in Cognitive and Information Sciences. All of you have touched my life, and research in innumerable ways. The success of this program, and my success in it is a result of the quality of the faculty here both as teachers and mentors, and as people.

I would like to thank the graduate students and post docs who have overlapped with my tenure at UC Merced. You've all been tremendous people, and the most amazing friends. There are too many of you who have had an indelible effect on my life to name here, but know that I am incredibly thankful to all of you.

I would also like to thank the National Science Foundation for partial funding of this work through grant BCS 1031903.

Curriculum Vitae

EDUCATION

UC Merced, Merced, CA PhD in Cognitive and Information Sciences	2016
UC Merced, Merced, CA B.S. in Cognitive Science & B.A. in Psychology	2009

AWARDS

2016 Insight Data Science Fellowship	Summer 2016
2015 CIS Summer Fellowship, <i>UC Merced</i>	Summer 2015
2014 Koostama Data Science Fellowship	Winter 2014-2015
2014 Koostama Data Science Fellowship	Summer 2014
2014 CIS Summer Fellowship, <i>UC Merced</i>	Summer 2014
Cognitive and Information Sciences Graduate Student of the Year	2012-2013
2012-2013 CIS Travel Fellowship, <i>UC Merced</i>	May 2013 – August 2013
2013 CIS Summer Fellowship, <i>UC Merced</i>	Summer 2013
2012 GRC Summer Fellowship, <i>UC Merced</i>	Summer 2012
2011 GRC Summer Fellowship, <i>UC Merced</i>	Summer 2011

TEACHING EXPERIENCE

UC Merced, Merced, CA Teaching Assistant – “Methods in Cognitive Science”	Spring 2016
Teaching Assistant – “Introduction to Logic”	Fall 2015
Teaching Assistant – “Methods in Cognitive Science”	Spring 2015
Teaching Assistant – “Cognitive Engineering”	Fall 2014
Teaching Assistant- “Introduction to Cognitive Science”	Spring 2014
Teaching Assistant- “Methods in Cognitive Science”	Spring 2013
Teaching Assistant- “Introduction to Cognitive Science”	Spring 2012
Teaching Assistant- “Methods in Psychology”	Spring 2011

PUBLICATIONS AND PAPERS

- Kerster, B., Kello, C.T.** (2016, August) *Spatial Memory and Foraging: How Perfect Spatial Memory Improves Foraging Performance. Proceedings of the Annual Meeting of the Cognitive Science Society*
- Kerster, B., Kello, C.T., Rhodes, T.** (2016, January) *Spatial Memory in Foraging Games. Cognition.*
- Rhodes, T., Kello, C.T., **Kerster B.** (2014, June) *Intrinsic and Extrinsic Contributions to Heavy Tails in Visual Foraging. Visual Cognition*
- Kerster, B., Kello, C.T., Rhodes, T., Bien-Aimes, R.J.** (2013, July) *Adaptive Foraging: Effects of Resource Conditions on Search Paths in a Web-Based Foraging Game. Proceedings of the Annual Meeting of the Cognitive Science Society.*
- Szary, J., **Kerster, B., Kello, C.** (2011, August) *What Makes a Brain Smart? Reservoir Computing as an Approach for General Intelligence. Talk in the Proceedings of the fourth Conference on Artificial General Intelligence.*
- Rhodes, T., Kello, C. T., **Kerster, B.** (2011, July) *Distributional and Temporal Properties of Eye Movement Trajectories in Scene Perception. Talk at The Annual Meeting of the Cognitive Science Society.*
- Kello C. T., **Kerster, B., Johnson, E.** (2011, July) *Critical Branching Neural Computation, Neural Avalanches, and $1/f$ Scaling. Talk at The Annual Meeting of the Cognitive Science Society.*
- Kello, C. T. & **Kerster, B.** (2011, February). *Power laws, memory capacity, and self-tuned critical branching in an LIF model with binary synapses. Poster session presented at Computational and Systems Neuroscience 2011*

Kello C. T., Rhodes, T., Hollis, G., **Kerster, B.** (2010, August) *Coordination Dynamics in Speech and Lexical Semantics*. Poster session presented at *The Annual Meeting of the Cognitive Science Society*.

CONFERENCE PRESENTATIONS

Drew Abney, **Bryan Kerster**, Chris Kello (July 2014) *The time course of visuospatial information in drawing from memory*, *The Annual Meeting of the Cognitive Science Society* (Poster)

Drew H. Abney, Chris T. Kello, & **Bryan Elvis Kerster**. (July 2014) *Multiscale spatial distributions of eye fixations decouple over time*, *North American Meeting of the International Society for Ecological Psychology* (Talk)

Thompson G. *, **Kerster, B. ***, Dale, R. (2013, August) *Emotiv as an Affordable Means of Extracting Recurrent Brain States*. *Society for Complex Systems in Cognitive Science*. Berlin, Germany (Talk)

*Co-Presenters

FUTHER TRAINING

APA Advanced Training Institute, Exploratory Data Mining
in Behavioral Research
Davis, CA

2013

APA Advanced Training Institute, Nonlinear Methods
for Psychological Science
Cincinnati, Ohio

2012

Tutorial, "Dynamic Field Theory: Conceptual Foundations and
Applications in the Cognitive and Developmental Sciences"
32nd Annual Meeting of the Cognitive Science Society
Portland, Oregon

2010

WORK HISTORY

Electrician's Assistant
Kurt Wille Electric
Provided assistance to master electrician, ran wiring, installed lighting

Summers 2007-2010

MEMBERSHIPS

Cognitive Science Society
American Psychological Association

Abstract

Foraging is an essential process for all mobile organisms. It allows organisms to locate resources such as food and mates. There is a long history of research on animal foraging in the ecology literature and recent work in cognitive science has revealed similarities between cognitive search behaviors and animal foraging behaviors. This gives rise to the possibility of bringing the rich animal foraging literature to bear on cognitive search processes. Historically there have been several major approaches to the study and modeling of foraging animals. One approach is known as optimal foraging theory which is focused on optimizing the amount of time an organism spends foraging a single location before moving on to another location. Another is an expansion on that approach which operationalizes those ideas into a spatial model known as area-restricted search. A third approach is known as the Lévy flight hypothesis. It focuses on the longer term distributional properties of foraging animals, and optimizes coverage of a given search space.

These approaches all make assumptions about environmental conditions faced by foraging organisms. The degree of resource sparsity and clustering in a foraging environment are believed to be important, but it is unclear how they affect foraging behaviors. Additionally spatial memory is a key concept important to animal search strategies, but is frequently ignored in the existing literature. For this dissertation, a series of experiments utilized a web-based foraging game to test how these variables affect foraging behaviors. The first experiment demonstrated that the degree of clustering in the environment had a significant impact on search strategies, and provided qualitative evidence that memory played a role in people's search behaviors. The experiments revealed distributional patterns very similar to those predicted by the Lévy flight approach. The second experiment refined the method and directly tested memory cues and a broader range of resource densities. This experiment revealed results similar to the first with the addition of significant effects of both memory cues and resource density. This dissertation then discusses a model that combined key concepts from both optimal foraging and Lévy foraging to produce results very similar to those produced by human participants, but with significantly higher performance. Experiment 3 examined how human performance changes when specific advantages are provided that can be found in our model, including perfect memory and accuracy.

Finally a continuation of the model is discussed that explores the dynamics of multiple foragers searching the same space. Overall, I demonstrate that people will generate Lévy-like search distributions in a wide variety of environmental conditions, but that search strategies will alter based on the current environment. I also demonstrate that spatial memory is a key factor in foraging, and provide a simple memory-based model that produces foraging behavior very similar to those utilized by people.

Introduction

Search functions are ubiquitous amongst living organisms, with foraging being perhaps the oldest (evolutionarily) and the most common. Foraging is a fundamental ability utilized by organisms to move around an environment in search of resources, such as food, shelter, or other organisms like mates. The study of animal foraging is a topic that has been difficult to study until relatively recently due to technical limitations. The advent of new technologies such as small global positioning systems (GPSs), or harmonic radar have begun to change that over the last 20 years or so, with researchers now able to actually track the movements of animals in their natural environments.

The rapid growth of research into animal foraging has led to the development of a number of competing models and theories that attempt to explain and fit the patterns of movement exhibited by animals as they forage through their environments. Some of these models and theories have started to pick up notice in the cognitive science community. It is easy to imagine that some of the same mechanisms and functions observed in animals may be prevalent in *Homo sapiens* as well. There are also those who have postulated that those same systems may have been adapted into use in other search processes as well, such as memory search and visual search. Or if not the same systems, it is possible that the same sort of search function is utilized.

There are three major approaches to modeling foraging which are relevant to my work in the dissertation which I will discuss in detail below. The first is the Lévy flights theory which takes a random walk approach to modeling the movements of animals. Optimal foraging theory which abstracts resource distributions into abstract patches and models movements in terms of how long a forager should exploit a single patch before advancing to another one. Finally, area restricted search is a model of search that focuses on the adjustment of search patterns based on the recency of finding a resource. I will discuss all three of these models in terms of their approach, the data they best represent, and where they fall short. I will also discuss the current literature on the relationship between animal foraging and human cognitive search.

Lévy Flights Theory

There are two major approaches to studying foraging patterns in living organisms. One of these approaches is the Lévy flight approach. A Lévy flight refers to a special class of random walk, where the step lengths of the random walk are drawn from a Lévy distribution (Viswanathan et al., 1996). A Lévy distribution is a class of probability distribution with a power-law tail. The Lévy flight is characterized by the following distribution function: $P(l_j) \sim l_j^{-\mu}$ where l_j represents the flight length, and μ serves as a slope parameter. For a Lévy flight μ falls between 1 and 3. When values are ≥ 3 the distribution becomes equivalent to a normal distribution (Viswanathan et al., 1999), and thus a Brownian walk. Lévy flights can be described as patterns of smaller movements occasionally interspersed by much larger movements. μ values near 2 have been shown to be optimal for foraging, and higher μ values within that range create distributions with more frequent small movements and fewer long movements. Lower values create more frequent longer movements, and fewer short movements. See fig. 1 below as an example.




Figure 1: Three example patterns of movements created by Lévy flights. m in this figure refers to the μ parameter (Viswanathan et al., 1999).

Lévy flights have been observed across a range of physical phenomena including fluid dynamics, and dynamical systems (Shlesinger, Zaslavsky, & Klafter, 1993). The Lévy flight was popularized as a fit for animal foraging data by Viswanathan and colleagues (1996). Viswanathan and colleagues had data collected from five different adult albatrosses in the Atlantic Ocean. Each albatross in the study had a device attached to their leg that simply recorded whether or not it was wet. Being wet indicated that the bird was eating or taking a break. Thus, a collection of durations between foraging sessions was acquired. Long range correlations were found within these durations with a μ value of approximately 2. They hypothesized that long-range correlations, or scale invariance are the result of the animals foraging in a Lévy flight like pattern. They further suggested that such patterns of movement could be advantageous to foraging animals. To provide plausibility to their explanations they also exhibited the similarity between the movements of the albatrosses, and a simple random walk model that pulls movement distances from a Lévy distribution.

This work was followed up a few years later in an article that focused on the optimality of long-tailed power-law distribution based random walks in foraging (Viswanathan et al., 1999). This paper focused on a Lévy walk model as opposed to a Lévy flight model. They are very similar models, with the difference being that a Lévy walk model moves incrementally and may be interrupted by finding resources as it moves, as opposed to Lévy flight that does not observe resources until it arrives at a preselected destination. It is worth noting the specific mechanics of the Lévy walk model used. It had a vision parameter, such that any resources within a specific range were “visible” to the model, resulting in motion towards the nearest visible resource. If none were visible then a direction was chosen at random and a distance was pulled from a

probability distribution. The model then moved incrementally towards that new point, until it either reached it and picked a new direction and distance, or “saw” a resource along the way. With this model, they show that a μ value of 2 may be the optimal parameter for the Lévy exponent. Evidence was presented both by analytically solving a model equation, and then further supported through simulations. According to the simulations different slopes offer only a slight advantage when the distance between foraging targets is small, however the advantage in foraging efficiency at 2 was much greater compared to other slopes when targets were farther apart. It is also worth noting that resources in both the analytical analysis as well the simulations were assumed to be uniformly randomly distributed across the search space. The paper concluded by looking at the path length or path time distributions for bees, deer, amoebas and albatrosses. All four exhibited power law distributions with a slope parameter (μ) of around 2, matching nicely with their analytical and simulated results.

The regularity with which Lévy distributions with a slope of 2 have been observed in a diverse range of foraging animals is quite striking. It was mentioned earlier that bees, deer, amoebas, and albatrosses all exhibited distributions of path lengths when foraging that are fit by heavy-tailed power laws with slopes right around 2. While those four animals cover a broad swath of the animal kingdom, they are far from the only animals to demonstrate this pattern. Reindeer (Mårell, Ball, & Hofgaard, 2002), micro zooplankton (Bartumeus, Peters, Pueyo, Marrasé, & Catalan, 2003), grey seals (Austin, Bowen, & McMillan, 2004), spider monkeys (Ayala-Orozco et al., 2004), 14 different species of fish (Humphries et al., 2010), and even human manned fishing boats (Bertrand, Bertrand, Guevara-Carrasco, & Gerlotto, 2007), are a selection of species that have been observed to follow the same distributions of foraging movements.

However, recent literature has begun to place a great deal of doubt on the Levy flight theory as a realistic model of animal foraging. Benhamou and Collet (2015) demonstrate an alternative model that significantly outperforms Lévy flights in contexts in which it was theorized to be optimal. Additionally, researchers have demonstrated that patterns of movement that are Lévy distributed may result from patterns of movement that are not driven by Lévy flights (Breed, Severns, & Edwards, 2015; Hills, Kalff, & Wiener, 2013).

Optimal Foraging Theory

Another common approach to modeling foraging is known as “patch foraging” or more frequently “optimal foraging theory.” Optimal foraging theory is partly characterized by a tendency to group sets of nearby resources together into single resource “patches” as an abstraction, as well as calculating the optimal amount of time to spend in a given location. This patch-focused approach tends to stay differentiated in the literature from the power law focused approach that characterizes the Lévy flight literature. This patch foraging literature is heavily influenced by two papers that were released in the same year, that lay out the idea of abstracting resources into patches, as well as the idea of trying to calculate how to “optimally” forage (Emlen, 1966; MacArthur & Pianka, 1966)., These ideas were later summarized into a fairly simple equation known as the marginal value theorem (Charnov, 1976).

Optimal foraging theorists treat the resource environment as being a “patchy habitat” Essentially groups of nearby resources are clumped together into individual patches with values assigned to them based on the number of resources contained within the patch. In some environments this simplification makes sense. It is easy to imagine dense packs of prey animals, or schools of fish, or a number of meal worms in a chunk of wood. The abstraction of groups of

resources into patches is critical to the optimal foraging approach. With patches one no longer worries about what pattern of movement to utilize while foraging within a patch, instead it is abstracted into the quantity of resources that can be collected per unit time. It is also assumed that this rate of resource acquisition decreases as resources within the patch are consumed. Charnov (1976) demonstrated that the optimal amount of time to spend in a given patch to constantly maximize net resource intake can be calculated as a function of the value of the current patch, the expected value of patches in the area, and the travel time from one patch to the next. Thus with knowledge about the surrounding environment it is possible to determine exactly how to optimally forage. This ability to mathematically calculate optimal search times in a patch is a key strength to this type of foraging model.

Area Restricted Search

Another approach to studying animal foraging behavior is area-restricted search. Area-restricted search in its most general form is very simple. When a resource is obtained the organism is more likely to stay in the same area for a period time (Kareiva & Odell, 1987). This kind of simple model can perform quite well under clustered or patchy resource conditions. Although, when implemented as a model it does leave one with a number of free parameters such as: How large an area is search restricted to? How long does the organism stay in the area? How should area restriction be realized, as an area of a specific size or more frequent turning? Etc... In implementation the simple area-restricted search can become fairly complex, but serves to solve a key issue with optimal foraging theory. Optimal foraging theory abstracts space into arbitrary patches. Area-restricted search anchors these patches into space, and provides a concrete method for moving within and between them.




Figure 2: Example movement through an environment with two types of patches. (Charnov, 1976)

Models of foraging are important tools for understanding how organisms locate resources in an environment, but in recent years researchers have begun to also apply these models towards certain cognitive functions. Processes such as visual search and memory search are the most frequent targets of such applications. This angle of research has begun to accumulate significant evidence of close similarities between foraging processes and cognitive search processes.

Foraging in Cognition

Visual search is an obvious cognitive domain to apply an animal foraging approach. Basic task constraints between the two are quite similar, beyond them both being search tasks. The goal in each is to move in a manner that best facilitates the location of resources in an environment. In visual search the resources are localized visual features, and the environment is

the visual scene. Typical patterns of movement also contain important similarities between foraging and eye movements. They can both be described as clusters of small movements within an area of interest, interspersed with larger movements between areas of interest. For visual search, eye movements occur as microsaccades interrupted by saccades, and for foraging these are periods of acquisition interspersed by relocations (Liversedge & Findlay, 2000). Both of these domains also frequently restrict the acquisition of resources to only one of these phases of movement. Birds generally do not acquire berries while flying from bush to bush, and eyes do not register features while performing saccades (Henderson, 2003). The similarities between these two tasks have been previously observed by researchers both from a Lévy flight perspective, and an area-restricted search perspective.

Unsurprisingly, researchers working from a Lévy flight perspective focused on the statistical properties of eye movement trajectories. Rhodes, Kello & Kerster (2011) demonstrated that eye movement trajectories follow heavy-tailed distributions with long-range correlations, across several different types of search tasks. Free viewing, counting, and a “find the hidden object” task were used, on a variety of naturalistic images that had various different power law spectra within them. Participants engaged in each of these tasks while wearing an eye-tracker to record the physical movements of their eyes. The tasks had a reliable effect on the parameters of the distributions, but remained power-law, while the distributional properties of the image, such as spatial distribution of luminance, had no reliable effect on eye movements. Power-law distributions in eye movements have also been observed in a visual search task known as the “visual world paradigm” (Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995), although the issue was only indirectly addressed (Stephen & Mirman, 2010). Eye movements in these studies proved to be heavy-tailed, but were not necessarily Lévy distributed. In both studies the distributions were best fit by lognormal distributions. However, this does not rule out the possibility that eye movements during search follow a Lévy flight-like strategy, as lognormal distributions can be viewed as a “constrained” power law (Stephen & Mirman, 2010), which leaves the possibility that the Lévy distribution may have been discernible if the experiment did not occur in a constrained space (like the computer monitors that were used in these eye-tracking studies.)

Visual search has also been approached from a marginal value theorem perspective (Wolfe, 2013). Marginal value theorem provides a very reasonable model of visual search especially if one is most interested in the relationship between saccade frequencies and the distribution of resources in the visual scene. It is a very small theoretical step to imagine the small eye movements an eye makes while it fixates in an area as being like foraging within a patch of resources. Then the important question becomes when do the eyes move their search to another fixation location? Wolfe frames this as a question of “search termination”; when do you terminate the search in one location to move on to the next?

Wolfe attempted to address the usefulness of marginal value theorem models, that is models which attempt to optimize the amount of time spent in a given patch before moving on to the next one, by using a computer based simulation, with a participant controlled mouse as a stand in for direct eye movements. The simulation had an “aerial view” and a “magnified view”. Clicking on a location in the aerial view caused their cursor to “walk” to that location. They could then collect berries within the magnified view by clicking on them (see Figure 3). Wolfe ran a series of six experiments using this general paradigm.

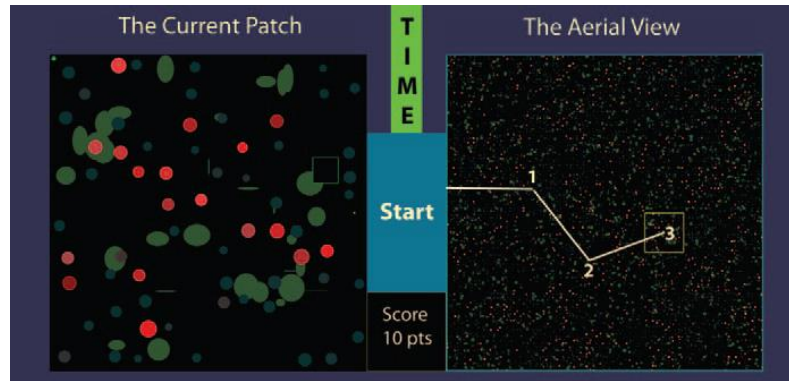


Figure 3: View of the simulation from (Wolfe, 2013)

This series of experiments tested a variety of different parameters to examine how they influenced the search behavior of participants. When resources were uniformly distributed behavior was well fit by a marginal value theorem model with participants becoming more likely to leave a patch as it was depleted such that the rate of yield from the patch dropped below the average rate for the task. For example, when travel time between patches was greater participants stayed longer in individual patches, or when it was more difficult to collect resources within a patch, both behaviors predicted by marginal value theorem. However, he also found that when patches were non uniform, behavior no longer appeared to fit with that predicted by marginal value theorem (Wolfe, 2013). It is also important to note that while this research may raise some interesting possibilities about foraging in human visual search, the experimental design failed to directly test visual search, and paid no attention to where participants were actually gazing at any point.

Optimal foraging models of visual search have been tested using methods that more directly test visual processes, although they have tended to use methods that have only one or no resources available per trial, and they measured the amount of time it took for participants to either spot the resource or give up looking, allowing researchers to look at search termination times (J. M. Wolfe et al., 2007). This line of research is was strongly focused on tasks such as finding the dangerous items in a baggage x-ray, and can be difficult to relate directly to foraging models (which are based around much less constrained environments, and generally contain multiple resources). However, this kind of methodology was used to show that individuals searched longer on each trial if they expected a higher target prevalence, as would be predicted by the marginal value theorem (Cain, Vul, Clark, & Mitroff, 2012). Individuals also demonstrated sensitivity to the resource distribution over the course of the experiment, and changed their strategies appropriately.

Another domain that has been explored from the perspective of animal foraging is memory search. Rhodes & Turvey (2007) performed the first study of Lévy flights in memory. They focused specifically on the domain of semantic memory, using a category recall task. Semantic memory search has fewer obvious direct similarities to animal foraging than visual search, but it is easy to imagine that the resources in semantic memory, that is, words, are distributed across a mental space, in a way similar to how resources are distributed in nature.

Participants were instructed to recite as many different animals as they could within a twenty minute time period. It should not be surprising that participants appeared to generate

animal names in bursts intervened by longer pauses. For example someone might name “cat, leopard, lion” then pause before continuing with “wolf, fox, dog.” Interestingly, these pauses appeared to be quite varied in duration, to the degree that some were exponentially longer than others. In fact, the distribution of the durations between utterances, that is the inter-response intervals, when plotted as distributions demonstrated heavy tails, and matched slopes within the range consistent with the Lévy distribution. They were also well fit by simulated Lévy flights. Duration between recalls functioned as a reasonable proxy for the “distance” between acquisitions used in physical search tasks, and the Lévy distributions produced in this task were striking in their similarity to the distributions produced by foraging animals.

Semantic search has also been studied from the perspective of optimal search theory. It is easy to see how the bursts of utterances produced in these category recall tasks could be viewed as series of patches, rather than as one single time series of utterances. In fact, it is on the basis of observing such clusters that Hills, Jones, and Todd (2012) argued that optimal foraging could be applied to semantic memory. They tested this using an experiment very similar to Rhodes & Turvey (2009) where participants generated words based on the category “animals.” There a couple of key differences between the two experiments. Items were typed on a computer instead of uttered aloud (although this probably did not impact the results see (Thompson, Kello, & Montez, 2013)), and the experiment only lasted three minutes instead of twenty. They also compared the behavioral data to a pair of search models operating over an artificial representation of semantic space based on hand coded categorizations (Troyer, Moscovitch, & Winocur, 1997) and a corpus based semantic space model (Jones & Mewhort, 2007). Using this representation of semantic space they created two models, one which is referred to as a static model that simply moves through adjacent items in the space, and a dynamic model which swaps between categories when they grow sparser (likened more to an area-restricted search approach). They showed that the dynamic model better fits the behavioral data, and argued that the dynamic model is more akin to how individuals explore semantic space (Hills et al., 2012).

In the domains I discussed prior to semantic memory search most of the experiments, models, and simulations are tested in very different environments between those interested in Lévy flight models, and those interested in optimal search models. These models are generally uniformly distributed individual resources for Lévy, and patchy distributions for optimal search. In the domain of semantic memory search both foraging approaches are now testing their ideas in effectively the same resource space, human semantic memory. This really begins to highlight some of the similarities and compatibilities between the two approaches. It is especially notable that both approaches have successfully fit themselves to what are effectively the same data sets by focusing on different aspects of the search process. Lévy focuses on the general shape of the distribution, and how it might be produced, while optimal value approaches focus on the larger transitions. This leads to both also having significant gaps in their explanations. Lévy approaches, while elegantly simple, also fail to account for the likely possibility that foraging decisions are probably not random pulls from a distribution, but are likely effected in some way by knowledge of the local environment. Optimal value approaches on the other hand ignore spatial and other relationships among patches, and the details of how foraging occurs within somewhat arbitrarily defined patches.

The importance of the underlying distributions in the foraging environments remains an important question in semantic memory space, and one that may actually be partially addressable by looking at how these foraging models fit the behavioral data. It is still not entirely clear whether the underlying distributions are what drive patterns seen in behavioral data, or whether

foragers just engage in strategies that best fit their current environment. That is, do foragers use one general strategy, or do they engage in different strategies depending on the current environment. Either way, existing evidence makes it clear that foraging behavior changes, and has to be modeled differently under different resource environments. For example in Lévy research it is believed Lévy flights are most optimal under sparse resource environments (Viswanathan et al., 1999), and it has been observed that animal foraging patterns in animals no longer appear Lévy-like when resources are densely distributed (Humphries et al., 2010). Similarly optimal search theory relies entirely on being able to categorize the underlying resource environment in clusters (Charnov, 1976). Thus, given the success of these models in the domain of semantic memory search, it would seem reasonable to make certain inferences about the underlying “shape” of semantic memory. The strong fit exhibited by the Lévy flight model implies that semantic memory consists of randomly and sparsely distributed items (Rhodes & Turvey, 2007). The fit by optimal search theory models seems to contradict this, and instead implies that semantic memory is patchy and uniformly distributed (Hills et al., 2012). This will be one of the key questions addressed in the work to follow.

The random walk approach, the patch foraging approach, and area-restricted search have led to findings that appear to reveal important characteristics of animal foraging. However, none serve as complete descriptions of observed behavior. It is likely that a more complete model of animal foraging could be generated by combining aspects of all three approaches. Area-restricted search has started this process by implementing the concepts of patch foraging, and maximizing the ratio of time spent and expected value in a location. My work integrates what has been learned from existing theories by incorporating Lévy Flights. My work will also address another issue that affects this body of research, which is a lack of clear understanding of the relationship between foraging distributions, and the underlying resource distributions that are foraged. Both approaches tend to make unrealistic assumptions regarding resource distributions, and both acknowledge that significant variations in behavior may occur as resource distributions change. The work in this dissertation focuses on providing several key new findings to the literature: demonstrating that people produce distributional search patterns that are well fit by a Lévy distribution, but that a random walk model of foraging is insufficient to capture the patterns displayed in human foraging behavior, showing that human foraging behavior is altered by the underlying resource distributions being foraged, and that spatial memory is a crucial feature for any model that attempts to accurately model human foraging.

This dissertation consists of five elements including three empirical experiments and two models which serve to demonstrate and test findings listed above. All three experiments were built as online foraging games that allow for the easy collection of large quantities of data. The first experiment allowed users to forage a variety of different environmental resource distributions using a two-phase movement heavily inspired by ocular movements. The second experiment is similar but greatly simplifies the constraints of the task and explores a few questions raised by the first, including the role of memory and time costs. The final experiment is very similar to the second but provides aids to the player that should allow for near optimal performance in the task. I will also describe a simple computational model that provides a very close fit to the human data, while performing very well at the task. The final element of my dissertation expands the model to explore dyadic foraging. Many of the findings presented here have been previously published in other forms. Experiment 1 was published in Kerster, Kello, Rhodes, & Bien-Aime (2013), experiment 2 and the computational model were discussed in Kerster, Rhodes, & Kello (2016), and experiment 3 was discussed in Kerster & Kello (2016).

Experiment 1

Lévy foraging paths can be modeled simply as random walks with path lengths sampled from a power law, although path directions may be correlated over time (Viswanathan et al., 2001), and mechanisms of navigation are left unspecified. Despite their simplicity, Lévy walks have proven influential because they suggest that search benefits are conferred by power law path lengths, at least under certain foraging conditions (Viswanathan & Buldyrev, 1999). In particular, when $\alpha \sim 2$ and targets are sparsely and randomly distributed, Lévy walks are able to search spaces more completely compared with Gaussian distributed random walks, and find targets at a maximal rate.

The potential benefits of Lévy walks have led cognitive scientists to investigate whether they occur in perceptual, memory, and decision-making search tasks. Rhodes and Turvey (2007), investigated Lévy walks in a classic category recall paradigm (Bousfield & Sedgewick, 1944). Participants recalled as many animals as they could from long-term memory, for twenty minutes. Inter-response intervals were used as indirect measures of memory “path lengths”, and they were found to be best fit by inverse power law functions with exponents near two. Rhodes, Kello, and Kerster (2011) found that saccade lengths in visual foraging tasks also followed a heavy-tailed distribution resembling the optimal Lévy walk, although the lognormal function provided the best fit to data. A lognormal can be viewed as a constrained power law (Stephen & Mirman, 2010), which should be expected when search is constrained to a relatively small space (a computer monitor). Most recently, Radicchi and Baronchelli (2012) found search intervals to be Lévy-like when buyers searched the bid space in online auctions, and observed exponents were shown to maximize economic gains.

These and other similar studies raise the question of what mechanisms and factors give rise to Lévy-like search paths across so many different species and foraging conditions. Theoretical analyses suggest that sparsity of targets is a factor, but it is prohibitively difficult to test this hypothesis in natural foraging conditions, including visual and memory foraging of natural scenes and categories. Also, most theoretical analyses have assumed randomly distributed targets (Viswanathan et al., 1999), but food and other resources may instead tend to be clustered in nature, as is the case with plankton distributions, for instance (Mackas & Boyd, 1979). In terms of mechanism, Lévy-like foraging may be intrinsic to both biological and cognitive search functions, in which case Lévy walks would occur regardless of search conditions. Alternatively, Lévy walks instead may emerge as a result of interactions between search processes and their environments. These interactions may unfold over the course of minutes and even faster timescales.

Much of the data used in the study of animal foraging is observational, and in an ecological setting, it is frequently impossible to perform experiments. For example, you cannot change the distribution of fish available in the ocean to observe if an albatross changes its feeding habits. However, exactly these kinds of constraints can be experimentally tested in a simulated laboratory setting using human participants. There is sufficient evidence to indicate that humans seem to engage in foraging behaviors very like those observed in animals. This is true both in naturalistic foraging behaviors such as hunting, deep sea fishing and exploring theme parks (Bertrand et al., 2007; Raichlen, Wood, Gordon, Mabulla, & Marlowe, 2013; Rhee, Shin, Hong, Lee, & Kim, 2011) and in cognitive behaviors such as visual search and memory search (Rhodes, Kello, & Kerster, 2011a; Rhodes & Turvey, 2007).

In this experiment, we examined the roles of sparsity and clustering in a web-based video game designed to mimic canonical foraging. We used a video game because it allowed me to know and explicitly manipulate search conditions. We made the game web-based so that we could collect data from large numbers of participants on Amazon's Mechanical Turk. Recent studies have shown that Turk yields data comparable to university participant pools (Germine et al., 2012; Snow, O'Connor, Jurafsky, & Ng, 2008), and we confirmed this in a pilot study of my video game run through the UC Merced participant pool (not reported). Turk allowed us to robustly sample the parameter space created by factorial manipulation of sparsity and clustering. Turk also allowed us to collect enough data to compare players who find greater versus fewer numbers of targets, i.e. high versus low performers. The main goal of this experiment was to determine how people's foraging behavior is affected by the environment, specifically if people's behavior are responsive to variations in resource clustering or density. Additionally, this experiment served to establish whether such variations in behavior would interact with foraging performance.

Methods

For this experiment we designed a two dimensional, web-based foraging game, focused on the manipulation of two resource distribution variables: resource density, and resource clustering. The design of this game was heavily influenced by the movement of flying animals such as the albatross flying over the ocean, and by how the eyes of humans move, as in a visual search task.

The game was framed as a space exploration task, with the goal of finding resources on asteroids. Participants used their mouse cursor to move a spaceship over a 1280x1024 grid of space. Movement occurred at two discrete scales: *zoomed in*, and *zoomed out*. While zoomed out the entire space was visible to the participant. Clicking anywhere on the space would "fly" the ship to that location. Participants could then press the space bar to zoom in by 15X and explore the local vicinity by clicking. Images from the Hubble space telescope were used as backgrounds to help engage players by being visually interesting, and "outer spacey." See figure 4 for an example of what the game looked like in play.




Figure 4: Example images from the task in play. The top demonstrates the zoomed out view. The bottom figure shows the zoomed in view with the ship near a resource.

Asteroids were only visible, and could only be collected while zoomed in. Resources were collected for points by moving the ship graphic to collide with the asteroid graphic. Successful collections were indicated with visual and audio feedback, a point for each collected

asteroid was added to the score. A set amount of fuel was provided for each trial of the game, such that each trial would take approximately five minutes to complete based on pilot work. Fuel expenditure (which was indicated with a fuel bar in the top right of the screen) was a linear function of distance traveled, as well as a small constant for each zoom in and zoom out action. Fuel could never be increased, only consumed. Play continued until all fuel was consumed.

Resource distributions were generated using an algorithm we designed that can generate distributions with varying levels of clustering. The algorithm takes a parameter which can vary from 0.0 (all resources on a single point) to 0.5 (a uniform random distribution). The algorithm functions by splitting the total area in half and then probabilistically placing the resources onto a given half based on the supplied parameter. For example with a parameter of 0.25, 25% of the resources would be assigned to one half, and 75% to the other half. Each half would then be split again, with assigned resources being again split in the same manner. This process continues recursively until either a single resource remains in a rectangle, in which case it will be assigned to a random location in the current rectangle, or until a rectangle less than the size of a pixel, in which case the all remaining resources are placed on that pixel. See the pseudocode (figure 5) for additional details on the algorithm. This algorithm created clusters whose sizes followed a nested scaling relation to varying degrees, consistent with findings of scaling law clustering of natural resource (Humphries et al., 2010; Mackas & Boyd, 1979). Conditions were generated using clustering algorithm parameters of 0.05, 0.15, 0.25, 0.35, and 0.5 (see figure 6 for examples). This was crossed with the following varying levels of resource density: 25, 50, 100, 150. These twenty different resource distributions were randomly assigned to participants.

```
function distributeResources(rectangle, prob_split)
{
  if (rectangle.stars_remaining < 1) return
  else if (rectangle.size < 1 pixel) {
    place remaining stars at pixel
    return
  }
  else if (rectangle.stars_remaining == 1) {
    place star randomly in rectangle
    return
  }

  // alternate between vertical and horizontal splitting
  (rectangle1,rectangle2) = splitRectangle(rectangle,alternate)

  for each star {
    if (random_prob() < prob_split) star in rectangle1
    else star in rectangle2
  }

  // bias direction is randomized each split
  if (random_prob() < 0.5) prob_split = 1 - prob_split

  distributeResources(rectangle1, prob_split)
  distributeResources(rectangle2, prob_split)
}
```

Figure 5: Pseudocode for the algorithm used to generate the resource distributions.




Figure 6: Example distributions with varying levels of clustering. .05 (top left), .15 (top right), .35 (bottom left), and .5 (bottom right).

The game was coded in Adobe Flash so it could be distributed via the web, with game data collected on a local server. We used Amazon's Mechanical Turk to find people willing to play the game twice (plus a 1 minute practice session to learn) for 75 cents in compensation. Pilot work indicated that the availability and quality of Turk workers fell off precipitously for tasks lasting more than 10-15 minutes. A step-by-step demo, along with instructions in English, was presented at the beginning of play, and each play was set to last about 5 minutes (assuming no breaks). Two plays were required for each paid work session.

Results

A total of 1,825 play sessions were administered on Mechanical Turk. Participants who did not produce more than 80 zoom in actions per play were excluded from analysis (603 participants). Pilot work indicated that participants who simply expended fuel to complete the task, rather than endeavored to find asteroids, were usually revealed as making mostly large zoom out movements to expend fuel quickly. Of the remaining 1,222 participants, it was determined that 393 of them were participants who played in two or more Turk sessions.

Analyses with and without these repeats indicated no qualitative change in results, so all plays were included in the reported results. Analyses combine zoomed in and zoomed out path lengths.

Visual inspection of zoomed out flight paths revealed directional movements that ranged in their temporal correlations, which are a basic expression of memory in search paths. Two example paths at the two ends of this range are shown in Figure 7. Paths that consisted of highly regular directional movements were seen as “sweep” strategies designed to systematically cover the space in left-right, top-down, spiral, and other search patterns. Paths at the other end of the range consisted of seemingly haphazard directional movements, akin to random walks. In the middle were mixtures of the two, plus directional movements that followed irregular contours of Hubble images (despite instructions that distributions were independent of images).

To minimize effects of practice and learning, only the second of two plays per Turk session was analyzed (unless specified otherwise). Performance was measured as the proportion of available targets harvested, and plays were divided into three categories of performance, for each of the 16 game conditions: Top 20, middle 20, and bottom 20 scores. More than 60 Turk sessions were randomly assigned to each game condition, so any additional plays were excluded from reported analyses. The constant of 20 plays per cell simplified statistical analyses, and excluding plays in between performance categories helped to make those categories distinct.




Figure 7: Examples of movements produced by participants. Red lines indicate movements, and the pink squares demonstrate the locations where the players zoomed in. The top figure demonstrates a typical sweep strategy, the bottom a more haphazard strategy.

All results are graphed and analyzed as a function of sparsity, clustering, and performance category. A three-way analysis of variance was conducted using the three categories as predictors for each dependent measure, and only main effects and two-way interactions are reported because these are the statistical tests relevant to my research questions and hypotheses. First, we examined score as function of these three factors, shown in Figure 8. The main effect of performance category is itself based on score, and is so large throughout our analyses that reporting its reliability was unnecessary.

For the other two main effects, sparsity was not reliable, $F(3,18) = 2.1$, $p > 0.1$, but clustering was, $F(3,18) = 38.42$, $p < .05$. The interaction of performance category with sparsity was also not reliable, $F(6,18) = 0.84$, $p > 0.5$, but it did interact with clustering, $F(6,18) = 55.41$, $p < .05$. Visual inspection shows that scores improved with clustering for high performers, but the opposite effect occurred for low performers. These results show that foragers adapted to clustering but not sparsity, and low performers appeared to adapt counterproductive strategies in terms of score. This provides initial evidence that high performers took advantage of the spatial correlations in clustering, suggesting that foraging paths emerged from interactions between search processes and game conditions.




Figure 8: Demonstration of how scores varied across different resource distributions. Strong performers improved when resources were more clustered (left), and there is little effect of quantity (right)

However, one might argue that each forager had an intrinsic strategy, or a strategy chosen without regard to conditions. Effects of clustering may merely show that the consequences of good versus bad strategies increase with clustering. We tested this possibility by examining the change in score from first to second play. We found that score increased over time for high performers (+19%), but decreased for low performers (-14%), $t(638) = 17.7$, $p < .01$. This

difference suggests that strategies changed over the course of play, for better or worse, indicative of interactions between search processes and game conditions.

Looking at mean path lengths high performers had shorter path lengths overall, which indicates that shorter path lengths allowed for greater coverage of the space (confirmed by coverage analyses not reported). This main effect would be expected to diminish, and possibly even reverse, if fuel cost more per zoom.

There were no main effects of sparsity on mean path length (see figure 9), $F(3,18) = 2.19$, $p > .1$, or clustering, $F(3,18) = 1.7$, $p > .2$. The interaction of performance category with sparsity again was not reliable, $F(6,18) = 1.66$, $p > .15$, but it interacted reliably with clustering once again, $F(6,18) = 4.61$, $p > .05$. Visual inspection shows that path lengths for high performers increased with clustering, whereas they decreased for low performers. The increase for high performers presumably reflected the increased need for larger jumps as clusters became more tight and scattered. Again, low performers appeared to adjust strategies as well, but in counterproductive ways.

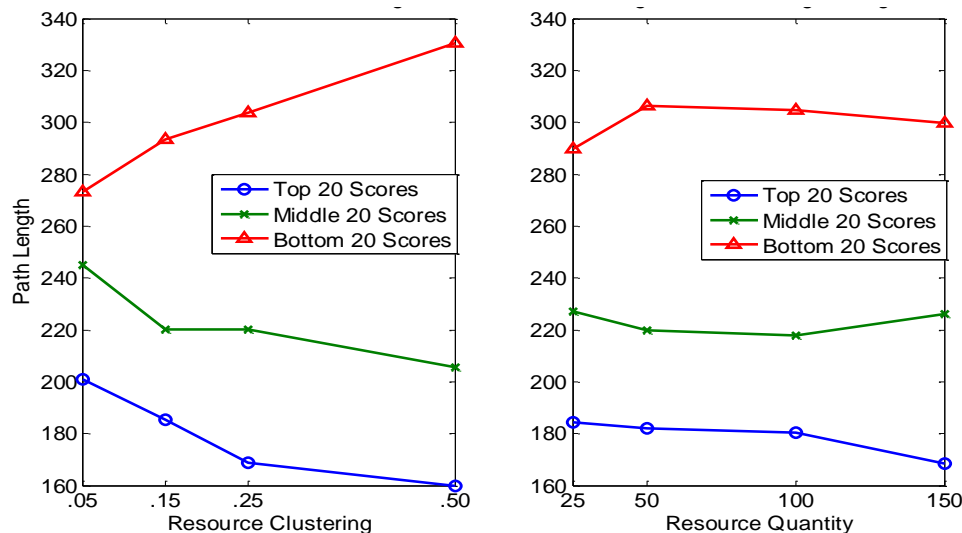


Figure 9: The effect of resource distributions on mean path lengths. Generally better performers have shorter path lengths, although they increase with resource clustering

Next we examined whether foraging paths resembled Lévy walks, in the sense that path length distributions were power law distributed with estimated exponents near two. We used multi-model inference (Symonds & Moussalli, 2010) to test which of four different functions provided the best fit to the distribution of path lengths for each participant (mean of 217 path segments per participant): Normal, exponential, lognormal, and Pareto. Only the latter two are heavy-tailed and Lévy-like, and the method uses Akaike's information criterion (AIC) to find the function with the shortest information-theoretic distance to the data.

The lognormal function provided the best fit for 68% of the participants, with the remaining trials roughly evenly split between normal and exponential fits (Pareto never provided the best fit). As mentioned earlier, the lognormal is akin to a constrained power law, and the foraging game constrained movements in terms of a limited amount of space. An example distribution from one participant is plotted in Figure 10 in logarithmic coordinates, which is

representative of the majority of participants. The constrained, normal-like portion of the distribution is seen as a slight hump on the left side, and the power law-like tail is seen as a negatively sloped line on the right.




Figure 10: A typical example of a path length distribution in log/log coordinates, along with its fit line (green). The rounding at the beginning of the distribution occurs across a range of a fairly small number of pixels, and is likely a result of the lack of precision people have when making very small mouse movements.

We further examined the tails of path length distributions in two ways. First, the fitted lognormal functions have μ and σ parameters, where the latter roughly corresponds with the heaviness of the tail. Heavier tails indicate more Lévy-like distributions. Best-fitting lognormal σ 's are used in the following analyses for all participants, including those whose data were better fit by normal or exponential distributions.

The overall pattern of results were similar to previous ones. There was no main effect of sparsity, $F(3,18) = 1.17$, $p > .35$, but clustering was again reliable, $F(3,18) = 10.52$, $p < .05$. Visual inspection revealed the possibility of an effect in the high/mid performers which was supported by a reliable interaction of performance category with clustering, $F(6,18) = 12.29$, $p < .05$. Once again, there was no reliable interaction with sparsity, $F(6,18) = 1.77$, $p > .15$. Visual inspection shows that the tails of path length distributions were heavier overall for low performers, but they *became* heavier with greater clustering *only* for high performers.

To gauge whether distributions were becoming more similar to the theoretical power law exponent of two, we fit regression lines to the right half of distributions in logarithmic coordinates, and results are shown in Figure 11. None of the main effects (excluding performance category) or interactions were reliable, but slopes were generally in the neighborhood of the theoretical optimum of -2. Moreover, slopes for high performers were closest in their approach towards -2 with greater clustering.




Figure 11: Regression slopes as a function of quantity, clustering, and performance category

To summarize, search paths generally resembled the optimal Lévy walks predicted to occur, to some degree. The majority of path length distributions were heavy-tailed, and tails resembled power laws with exponents near two, to some degree. Distributions most closely resembled Lévy walks for high performers in the most clustered resource conditions. Foraging paths also departed from Lévy walks, in that directions were never drawn purely at random. Example paths showed that directions tended to be correlated in time. This tendency can be quantified simply by computing the proportion of times that next steps went in the same direction, within some threshold. Angular changes were between 0 and 180 degrees (collapsing left versus right turns) and divided evenly into 45 bins. The proportion of movements falling into the smallest angular bin was 25%, about ten times greater than chance. Thus, foraging paths had memory in terms of directional correlations.

Discussion

This experiment revealed a number of interesting findings. First of all search movements appeared to follow a constrained power-law distribution in player search movements, with a slope near -2 as would be predicted by a Lévy-walk theoretical approach. Furthermore, better performers in the search task tended to be closer to the theoretical optimum than poor performers. These findings provided support for the idea that a Lévy distributed random walk model is an accurate representation of organisms forage in an environment.

However, we also find several effects that cannot be accounted for by a Lévy-walk model approach. We found that people may be responsive to their foraging environment, however that finding is possibly confounded by the fact that participants were grouped by their performance. We also found that search patterns alter significantly across varying resource distributions. This suggests that people were also not using a random walk strategy. They were instead using some kind of more sophisticated search process that is sensitive, at least to some extent, to the

underlying resource distribution. Additionally the distributions of angular change in people's movements were inconsistent with a random walk which would produce an even distribution, where participants demonstrated a tendency to prefer forward movement.

It is also interesting that we repeatedly found evidence for strong interactions between foraging performance (score) and different degrees of resource clustering in the environment, but not for the density of the resources. This is surprising because existing modeling work using Lévy-flight models has demonstrated strong relationships between resource density and search performance (Viswanathan et al., 1999). One possible explanation for this is the range of densities used was relatively small. This is one of the ideas that will be explored further in experiment 2.

Experiment 2

There were some concerns in experiment 1 regarding having two stages of movement baked directly into the task. One important aspect of the first experiment was that movement was multi-scaled by the very nature of how the ship moved in two separate scales. This was done with the intent of capturing a system akin to eye-movements or a foraging bird. However, this also gives rise to some concerns about how this two-scale movement translated into path length distributions. Perhaps multi-scale movement patterns emerged solely as the result of the movements being inherently multi-scale. To address this concern this second experiment further simplified the foraging game, and removed the “zoom” style movement and resource collection.

Experiment one also made it clear that the strategies users chose to utilize in the game were clearly not random walks, as one might expect from the Lévy-flight foraging theory. In fact, many users used a “mowing” or “sweeping” strategy to cover the space, and in general users were much more likely to continue moving forward than any other direction. This suggests that spatial memory may be an important part of foraging strategy for organisms capable of utilizing it. At the very least, this appears to be the case for humans. Experiment two will attempt to test the importance of spatial memory.

Methods

For experiment two a second foraging game was created. This game attempted to simplify the task as much as possible to try to capture very basic foraging. This time the game was framed as a task to explore space by sending out “probes” to search for hidden resources. Participants were presented with a 1280 by 1024 pixel screen, where each pixel on the screen represented a location where resources could be hidden. Again, no visual cues were provided to indicate where resources might be hidden. Instead of moving a ship, and zooming in and out the game was played using a computer mouse (or functionally equivalent pointing device) to click on any area of the screen to “send a probe” to that location. This caused a 15 by 15 pixel area centered on the mouse cursor to be explored. Participants were notified of any found resources in an explored location through auditory and visual cues. Resources were removed when found and never regenerated. For example clicking on the same location with a resource twice would indicate a reward the first time, and indicate no resource the second time. A 15x15 pixel square outline was also placed around the mouse cursor at all times so that participants could see the size of the area that would be searched if they chose to click at that location. Instead of removing fuel based on distance travel, participants were allowed to search for resources 300 times per trial no matter what. This created a trial length of approximately two to three minutes. This change was made for two reasons, the first being to simplify the game further, and the second being to

discourage players from attempting to choose strategies that finish the game faster as opposed to maximize score. Such behavior was relatively common (although easy to detect) in experiment one. This change made the experiment take approximately the same amount of time regardless; a feature we discovered is especially useful for experiments that will be run online. A meter in the top right corner of the screen indicated what percentage of probes remained; the raw number of remaining probes was also displayed adjacent to the meter. In the top left corner of the screen, as in experiment one, the number of found resources for the current trial was displayed.

At the beginning of each experiment, participants viewed a text briefing that described how to play the game and informed them that there would be a leader board. Participants then viewed an example distribution of resources that had the same density and clustering as would be in their trials. This meant that participants did not have to spend time during the trials to learn how resources were distributed, and could attempt to start with a strategy tailored to their particular distribution in their trials.

This game was also written in Adobe Flash and Amazon's Mechanical Turk service was used to find participants. Each participant was presented with the briefing, two trials, and a debriefing. Each trial was also preceded by a short practice trial. The entire experiment took participants approximately eight minutes, and they were compensated \$0.50 for their effort. Data was collected from 2,040 participants.

The experiment had four factors; three of them were manipulated between subjects: the density of the resource distribution, the clustering of the resource distribution, and image background. All three of these manipulations were randomly assigned.

Resource density levels were chosen that were much greater than in experiment one, and cover a much greater range. In hindsight I realized that the range of density values in experiment one were much smaller than those used to test many models e.g. (Viswanathan et al., 1999). There were three levels: 100 resources, 600 resources, and 1,100 resources.

Resource clustering was also manipulated across three levels. The parameters used were 0.1, which produces a very heavily clustered distribution, 0.3, which produces a moderately clustered distribution, and 0.5, which produces a uniform random distribution.

Half the participants were shown a background image during their trials. These images were randomly chosen from a small selection of Hubble space telescope images. The other half of participants saw a plain black background during their trials. This condition tested the effects of memory on foraging behavior. As in experiment 1 the content of the images was completely uncorrelated with the placement of resources, so any effects of this manipulation could be attributed to participants' ability to use features in the images as landmarks.

The within subjects manipulation adjusted the limitations on how frequently participants could click to use a probe. Each participant saw both levels of this manipulation in one trial with the order randomly determined. One version of this manipulation was that participants could click as frequently as they liked without any limitation until they ran out of probes. The other version created a variable delay linearly correlated with the distance since their last click. For example if a participant clicked somewhere on screen, and then somewhere far away there would be a significant delay before they could click again, however if they clicked somewhere nearby they would be able to click again with almost no delay. This delay was animated with a dot moving from the previous location clicked to the new one, with the delay ending when the dot reached the newly clicked location. This manipulation added an embodied cost to moving large distances, in an effort to see how this would change foraging strategies, similar to that experienced by participants in experiment one. We had questions from experiment 1 about the

impact this kind of embodied time cost might have on participants' strategies. This ultimately had little effect on the results, so this manipulation is ignored in the analyses.

Results

The Euclidean distance between each successive pair of clicks was calculated, and distances were integrated into summed path lengths for consecutive segments that formed an angle of less than 10° . Each game yielded 165 integrated path lengths on average, and path length distributions were analyzed along with percent targets found per player, as well as path lengths as a function of number of consecutive locations foraged without targets. An example of a foraging trial can be seen in figure 12. The effects of our foraging manipulations on these dependent measures were investigated to test the adaptive nature of search trajectories.

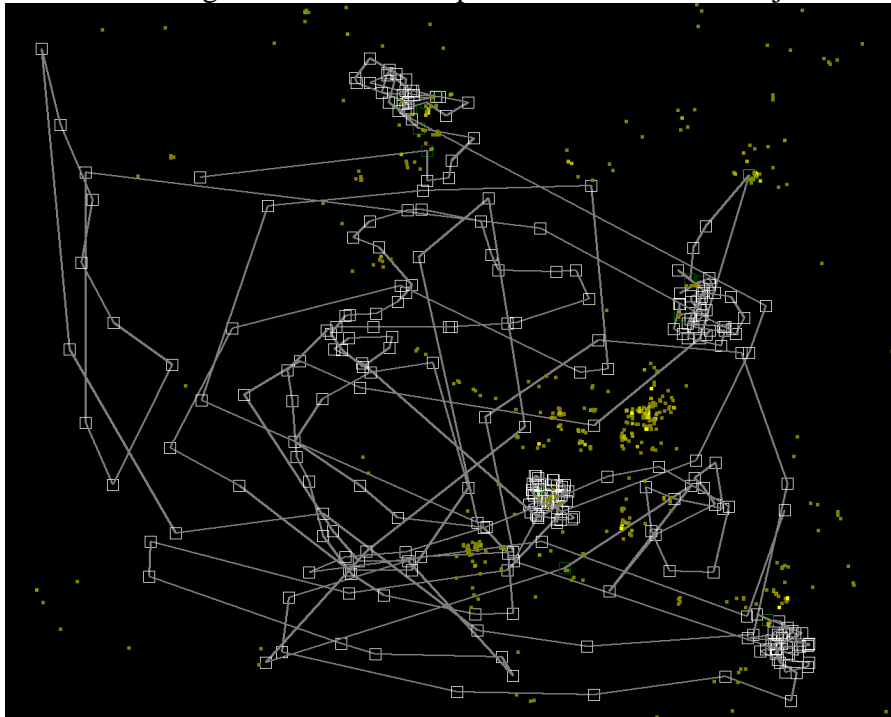


Figure 12: An example distribution produced by a participant in experiment two. The boxes indicate the area searched on a click. The yellow dots indicate the locations of resources and were not visible to participants.

For each path length distribution produced by each subject, the same maximum likelihood method from experiment 1 was applied to estimate the best fit estimates for lognormal, exponential, pareto (power-law), truncated power law and Gaussian functions. The large majority of distributions were best fit by the lognormal (96%), whose heavy tail is indicative of clustered search trajectories. While the distributions clearly were not strictly power law functions as predicted by Lévy walks, we could nonetheless estimate the amount of clustering by fitting a power law function to the tail of distribution, i.e. path lengths > 42 pixels as a minimum cutoff determined using the method developed by Clauset and Shalizi (2009). These distributions can be viewed in figure 13. The estimated exponent (referred to as α in this experiment) served to quantify the heavy tail (Rhee, Shin, Hong, Lee, Kim, & Chong, 2011), and the cutoff separated out smaller path lengths that were more uniformly distributed, presumably

due to noise from screen, mouse, and motor control limitations. Interestingly, estimated exponents (α values) were near the theoretic optimum of 2 for Lévy walks, even though distributions were lognormal rather than power law.




Figure 13: The path length distributions from all participants sorted by distribution category.

Our finding of heavy-tailed (lognormal) path length distributions in all conditions indicated that search trajectories were inherently clustered. However, degrees of clustering changed as a function of target distributions and visual landmarks. In particular, clustering of search trajectories increased with target density, decreased when targets were randomly distributed, and increased with visual landmarks. These effects were tested using mixed effects analyses of variance (ANOVA) models, with subject as a random variable. Each effect was tested using likelihood ratios to compare the full mixed effects model against a null model that excluded the effect in question. Chi-square statistics showed reliable effects of target density, $\chi^2(2) = 90.18, p < .0001$ and target clustering, $\chi^2(2) = 48.70, p < .0001$, and more clustering with visual landmarks ($M = 1.92, SE = 0.007$) than without ($M = 1.86, SE = 0.007$), $\chi^2(1) = 18.54, p < .0001$.

The effects of target density and clustering indicate that players modified search trajectories based on information in the foraging environment. Analyses of search performance indicate that these modifications were adaptive to the search environment because effects on exponents corresponded to similar effects on percentages of targets found: Target density and clustering again had reliable effects on percent targets found, $\chi^2(2) = 79.44, p < .0001$ and

$\chi^2(2) = 575.8, p < .0001$, where performance increased with density and was lowest when target distribution was random. Also in line with exponents, performance was marginally better with visual landmarks ($M = 7.51\%$, $SE = 0.19$) than without ($M = 7.03\%$, $SE = 0.17$), $\chi^2(1) = 3.39, p = .066$.

These additional results indicated that clustering behavior in search trajectories adapted to clustering structure in target distributions in order to enhance performance. We examined the relationship between search and target clustering directly by dividing the range of observed exponent values into 10 bins with equal numbers of data points, and plotting the mean exponent of each bin against its corresponding mean percent targets found (see figure 14). Graphs showed positive linear relationships that were steeper with greater target clustering, and no discernable effect of target density. Using linear regression, estimated exponents were reliably correlated with percent found at each of the three levels of target density, and in both clustered conditions ($p < .0001$ in each case). Surprisingly, score appeared to increase as α increased, where prior research would predict a peak at 2.0. There was no such relationship in the random condition because finding targets provided no information about locations of other targets.

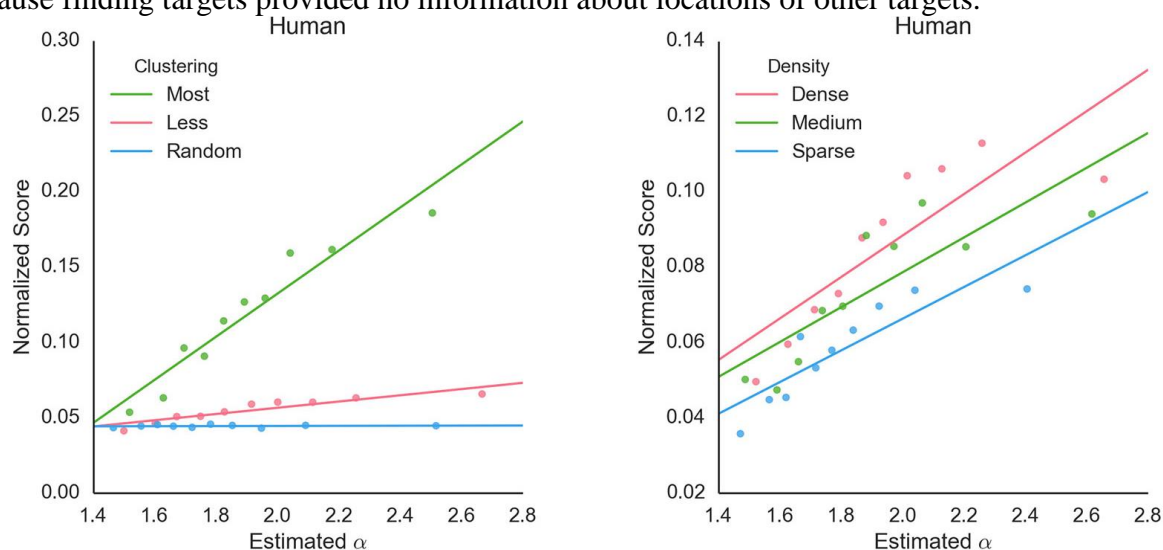


Figure 14: Demonstration of how α correlates with score at different levels of clustering (left) and density (right). A higher α value indicates a more clustered pattern of movement.

Finally, correlations between search and target clustering provided indirect evidence that foraging movements were adapted as search unfolded and specific target clusters were found. We also tested the dependence of search movements on learned information about target distributions more directly by examining the relationship between each individual path length and the corresponding number of prior clicks without finding a target (figure 15). These analyses showed that players made shorter movements soon after finding targets, which became increasingly longer as the number of consecutive unsuccessful clicks increased.




Figure 15: Normalized time since last click (z score) plotted against log binned path lengths. Path lengths increase with the time since the last resource was found across all conditions.

Discussion

This experiment confirmed several of the key findings from experiment 1. Despite significant differences in how players physically explored the space in the two experiments, both revealed that people generate Lévy-like lognormal distributions with estimated power-law exponents near the theoretical optimum of 2, much like the distributions produced by foraging animals. This occurred even on trials where participants failed to find any resources which

implies that Lévy-like movement distributions are a feature inherent to how people forage and not solely a product of the search environment. Furthermore, higher power-law exponents correlated with better search performance. These findings would be predicted by a Lévy-flight modeling approach.

However, we also found that people are sensitive to both the density and clustering of resources in their search environment and make significant adjustments to their foraging patterns in order to compensate, and unlike in experiment 1, we can be confident in this effect because no post hoc grouping of participants was performed for those analyses. This result would not be predicted by a Lévy-flight model. The fact that people clearly used memory to help guide their searches also makes it clear that people were not searching using canonical Lévy-flights, as they are a memory-less process. The fact that participants used shorter paths after finding resources would instead suggest people were performing an area-restricted search. However, that model fails to explain the Lévy-like distribution of path lengths, especially in the case where no resources are ever discovered. A new model is needed to be able to account for all of these findings.

Computational Model

The first two experiments revealed several important features that as a group are not yet accounted for by any single model of foraging behavior. People produced inherently clustered movements with distributional properties that resemble a Levy walk. But they also adjusted their movements to match the search environment, and produced shorter path lengths shortly after finding resources. Those are both behaviors that would be predicted by an area-restricted search model. However, it also appears that people were utilizing memory to a degree that neither model would predict. In order to address the shortcomings of these existing models we have designed a memory-based computational model which combines important aspects of both area-restricted search and Lévy-flight search. This model assumed complete spatial and temporal memory as a simplification, rather than a necessity. The foraging space is represented as 2D grid of squares, where the search agent chose each move by sampling from a probability distribution of unvisited squares indexed by xy ,

$$P \left[\exp \left(\gamma \frac{E_{xy} + 1}{\beta M_{xy} + 1} \right) \right]$$

where E_{xy} is the expected number of targets at square xy , M_{xy} is the Euclidean distance from the agent's current location. The model has two free parameters, one to balance randomness versus determinism in selecting squares to visit (γ), and one to weigh the effect of distance on movements (β). The parameter space was searched and two values were found ($\gamma = 5.5$ and $\beta = 0.125$) that maximized overall fit of the model to the human data. More details on the effects on these parameters can be found in Appendix A. For the findings described here the values were held constant across all simulations and conditions, which means that patterns of differential effects across conditions reflect how model behavior changes solely as a function of the target distribution.




Figure 16: Plot comparing the path length distributions of the model and humans across each resource distribution.

Expected numbers of targets were based on previously visited locations weighted by distance, i.e. $E_{xy} = \sum V_i/D_i$, where V_i is the trace value at each previously visited square i , D_i is the distance between xy and each visited square, and the sum is over all previous visits. Traces encoded numbers of targets as they were found, i.e. $V_i = N_i$ for $N > 0$ upon visiting each square i . Thus expected numbers of targets were partly extrapolated from previously found targets, and movements were drawn to locations nearby. However, when no traces were found, an opposite, *repelling* force is applied, i.e. $V_i = -\ln(E_{xy})$ for $N = 0$. These negative values played an essential role in offsetting unmet expectations, and thereby driving movements away from previously found targets as searching around them failed to yield more targets (the logarithm was necessary to attenuate repulsion and thereby ensure clustering). The dynamic can be conceptualized as a spatially continuous version of patch foraging, with the added feature of spatial memory to guide movements towards or away from specific locations based on search history.

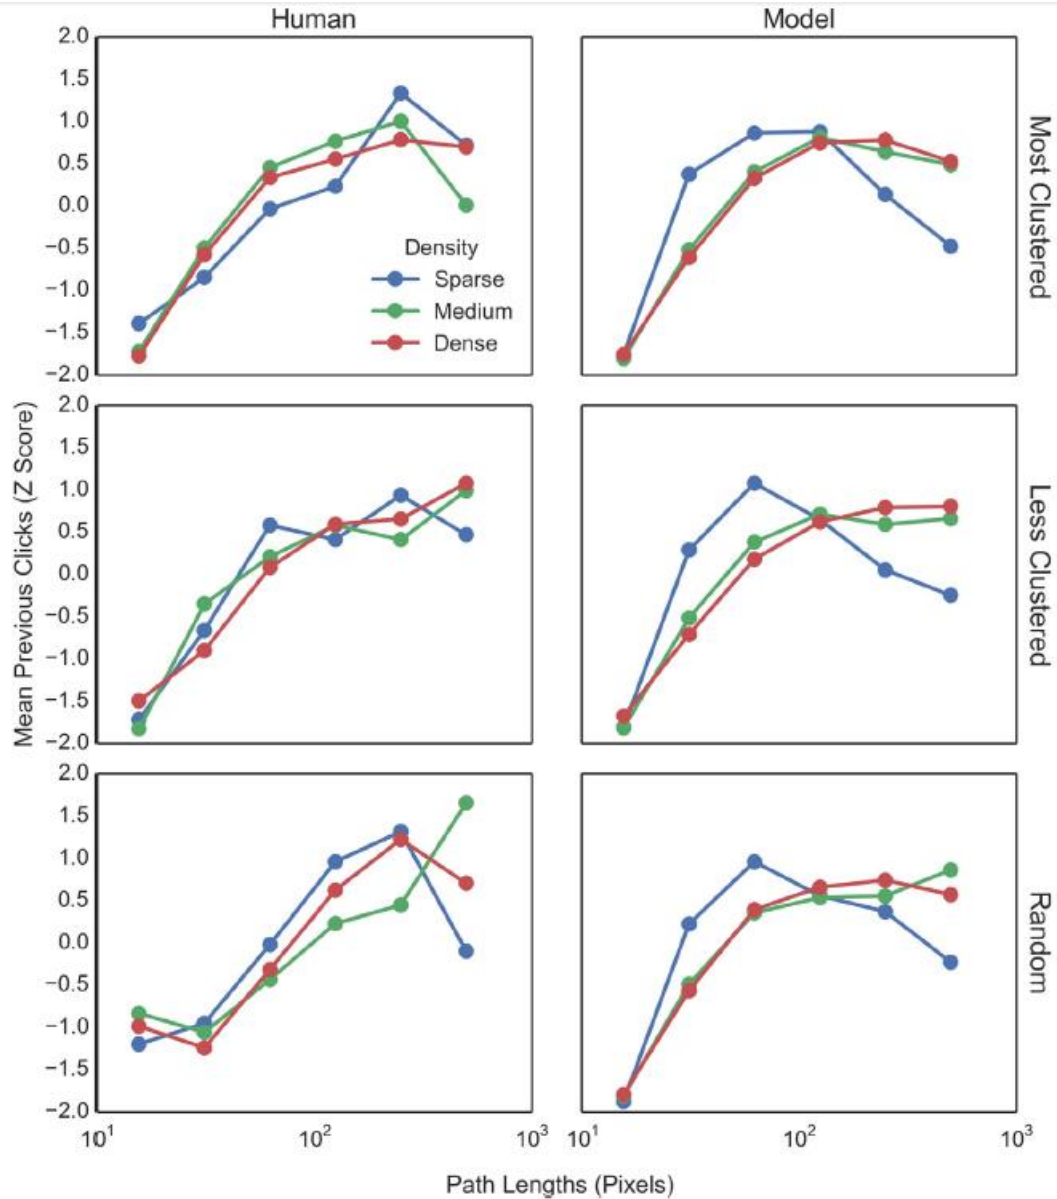


Figure 17: Time since last resource found plotted by path length for both humans and the model. The model follows the same trend as humans where path lengths increase as a function of the time since a resource was last found.

The model simulated each game play by each subject, and model trajectories were analyzed in the same manner as the human trajectories were. The model exhibited the same qualitative patterns of results as humans, shown throughout the figures and tables. Trajectories were inherently clustered (figure 16) with similar exponents fit to heavy-tailed path length distributions, although model trajectories were all best fit by power law functions instead of lognormal. The power law means that search trajectories resembled Lévy walks, and thus the

model may help to explain evidence for environmental influences on Lévy walks in animal foraging (Ferreira, Raposo, Viswanathan, & da Luz, 2012). The discrepancy with human players requires further study, but may be due to the empirical limitations noted earlier.

Like humans, the model generated clustered movements around targets when first discovered, and then made longer movements away as locations were increasingly barren (figure 17). This dynamic was an adaptive feature of the model, but the same underlying mechanism also caused inherent clustering when no targets are found. The reason is that, when target expectations drop below one, $-\ln(E_{xy})$ becomes positive and a trace is encoded *as if* a target was found. The search process then clusters around this internally generated trace, until enough empty locations are visited to trigger a long excursion, and the cycle may be repeated. This dynamic of target expectations was responsible for the intrinsic power law path length distributions, and it distinguished our model from previous foraging models with spatial memory (e.g. see Gautestad & Mysterud, 2010).

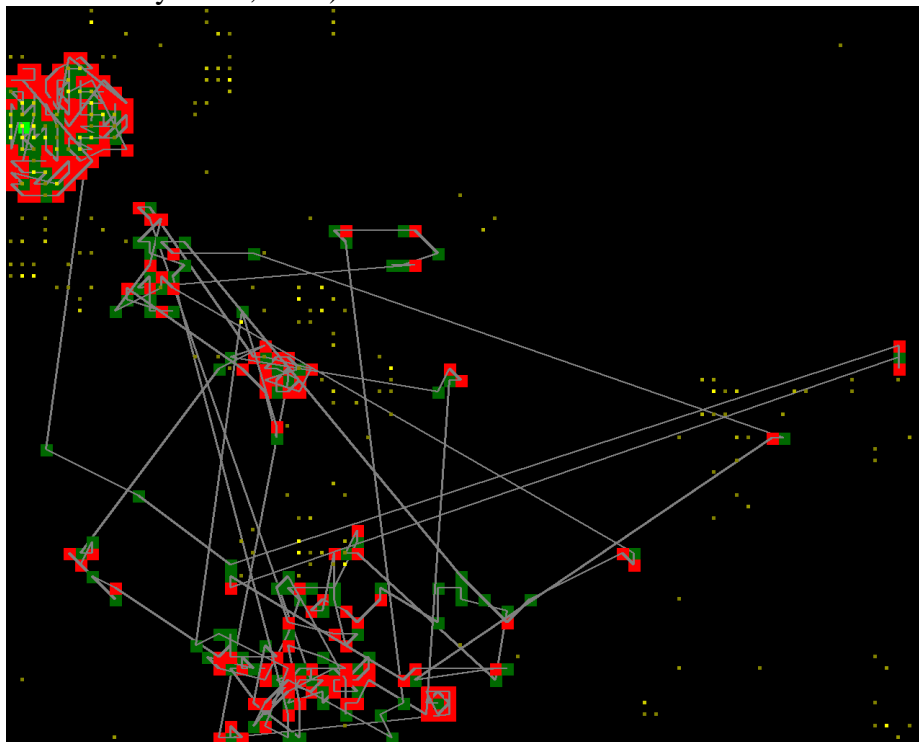


Figure 18: An example movement created by the model. Red indicates a visited location where a negative trace was left. Green indicates a positive trace. The yellow dots indicate the actual locations of resources.




Figure 19: Demonstration of how α correlates with score at different levels of clustering (left) and density (right) for the model.

Model performance and estimated exponents followed the same patterns as those observed for human players, including the variability in scores, as seen in similar coefficients of variation (see table 1 and figure 19). The main differences between human and model foraging derived from the model's complete memory—performance was about double that of humans when target distributions contained information about target locations, and path length distributions for the model were less variable than for human players, as seen in smaller coefficients of variation for estimated exponents.

		% Found				Est. Exponent			
		<u>Human</u>		<u>Model</u>		<u>Human</u>		<u>Model</u>	
Density	<u>Sparse</u>	9.9%	(0.81)	14.1%	(1.00)	1.84	(0.16)	1.74	(0.03)
	<u>Medium</u>	11.3%	(0.90)	22.5%	(0.94)	1.95	(0.19)	1.89	(0.05)
	<u>Dense</u>	13.2%	(0.88)	28.1%	(0.80)	2.00	(0.16)	2.08	(0.08)
Clustering	<u>Most</u>	21.7%	(0.57)	42.8%	(0.53)	1.97	(0.16)	1.92	(0.09)
	<u>Less</u>	7.4%	(0.23)	14.5%	(0.35)	2.00	(0.17)	1.98	(0.11)
	<u>Random</u>	5.5%	(0.17)	7.1%	(0.17)	1.81	(0.18)	1.84	(0.05)

Table 1: Table comparing search performance and estimated exponents and their standard errors for the humans and model.

Experiment 3

The model demonstrated the importance of accuracy and spatial memory in foraging tasks. The difference in performance between participants in experiment 2 and the model raise

the possibility that improving people's spatial memory and accuracy would improve their performance to a level comparable to the model. Experiment 3 empirically tested the predictions of the model by utilizing a method very similar to experiment 2 while externally boosting participants spatial memory and accuracy.

This experiment used the same methods as experiment 2 with some minor changes to explore these questions of memory and accuracy. Rather than allow users to click anywhere on the screen, the screen was explicitly divided into 15x15 pixel squares. This created a search environment that is effectively identical to the one used by the model, and it will not be possible for users to accidentally overlap portions of their search attempts. It was also indicated to players locations that they had visited, using hue to indicate the magnitude of a find at the location. This provided them perfect memory similar to what the model enjoys. These adjustments created an experiment with results that could be more validly compared directly to the model, and ultimately confirm the importance of spatial memory on foraging performance (see figure 20 for an example).




Figure 20: An example of how experiment 3 appeared during gameplay. Squares indicate where the player has foraged thus far. White squares indicate nothing was found at that location, and red squares indicate that resources were found, with brighter shades indicating higher values.

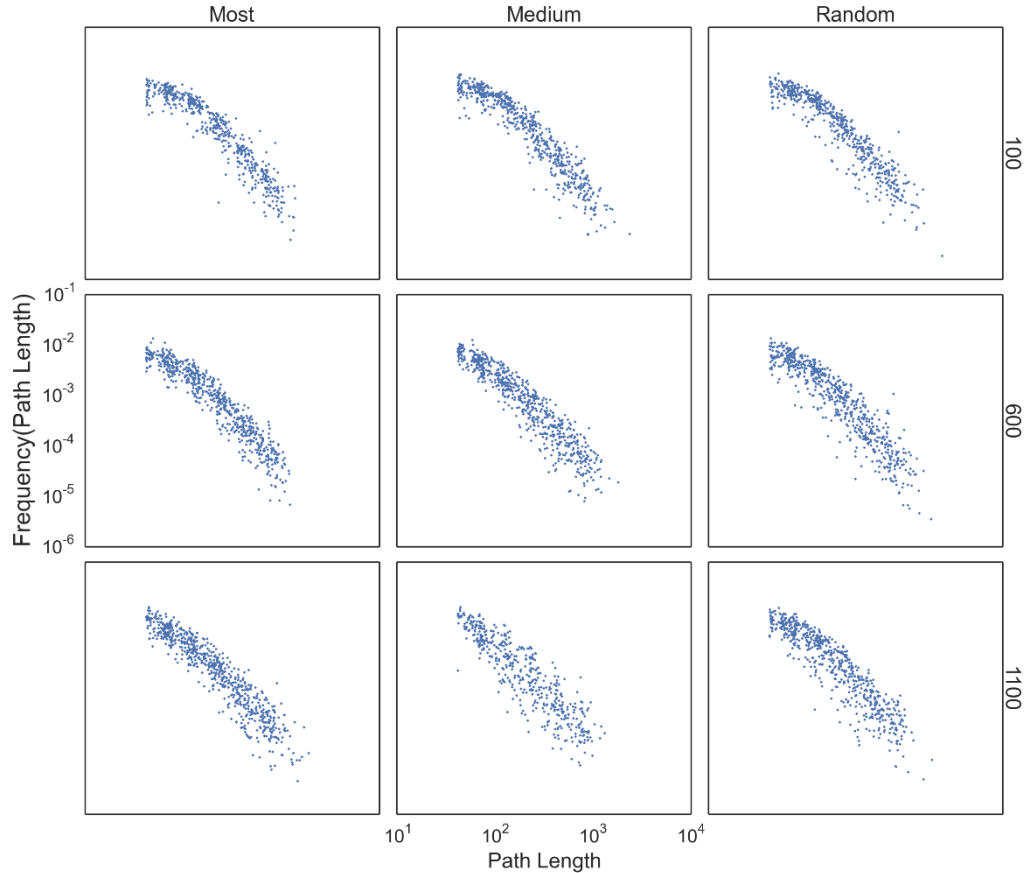


Figure 21: Path length distributions for each trial. The results are very similar to the earlier experiments and demonstrate the same power law like distribution.

Results

The data was analyzed using the same methods as in experiment two. The movement patterns observed in this experiment proved to be very similar to those observed in previous experiments (figure 21). The distributions of movements demonstrate the same tendency towards inherent clustering, with most of the distributions being lognormally distributed. This was determined using the same maximum likelihood method described earlier. 939 out of 1033 (90.1%) trials were determined to be lognormal.

One of the key goals in this experiment was to see if giving participants the same benefits we afforded the model, perfect memory, and perfect accuracy (that is the ability to never overlap previous search locations) would boost scores into ranges similar to those observed in the model. On this front, the experiment was wildly successful. Scores were boosted substantially across all conditions (table 2). In fact, participant's scores in this experiment surpass those produced by the model. This experiment demonstrates empirically how search performance can be boosted drastically simply by improving memory.

Density	% found		Est exponent	
	Mean	Variation	Mean	Variation
100	22.8%	0.86	1.93	0.11
600	28.2%	0.80	1.97	0.12
1100	32.6%	0.75	2.06	0.15
Clustering				
1	55.6%	0.20	1.95	0.10
3	15.9%	0.28	2.04	0.14
5	7.0%	0.14	1.98	0.14

Table 2: Performance of the model across levels of density and clustering in terms of the percent of resources found, and degree of clustering as measured by the estimated α value from a pareto distribution.

As mentioned earlier, two key changes were made between this experiment and experiment 2. Without further controls it is difficult to precisely determine how much of the increases in score were due to improvements in spatial memory, and how much were due to increased search efficiency because search locations could no longer overlap. Nevertheless, we developed a simple method to control for search efficiency by scaling normalized score by the number of unique pixels visited. This allows for the relative comparison of scores accounting for inefficiency due to repeatedly searching the same area. This analysis revealed that, for each condition except the random conditions, the majority of observed performance increases (about 70-80%) were due to improvements in spatial memory (see Table 3).

	% increase due to search efficiency	% increase due to perfect memory
<i>Density</i>		
Sparse	23.2%	76.8%
Medium	24.5%	75.5%
Dense	26.5%	73.5%
<i>Clustering</i>		
Most	20.3%	79.7%
Less	30.7%	69.3%
Random	77.9%	22.1%

Table 3: Percent of the increase due to each of the two changes made between Experiment 2 and 3, across density and clustering.

The next question is whether the degree of clustering in the search movements is modulated by the underlying resource distributions in the same way that was observed in experiment 2. The pattern observed here is similar to that in experiment 2, but has a few differences. The most striking of these is that the α values in the most clustered condition are actually lower than the values in both of the other conditions (figure 23). In experiment 2 this condition fell between the values of the medium and random clustering conditions. It is also worth noting that the observed values fall within a narrower range than in the previous experiment, so there are fewer overall differences between search strategies in this experiment (see figure 22). The reason for these differences is unclear and remains a potential topic for future research.

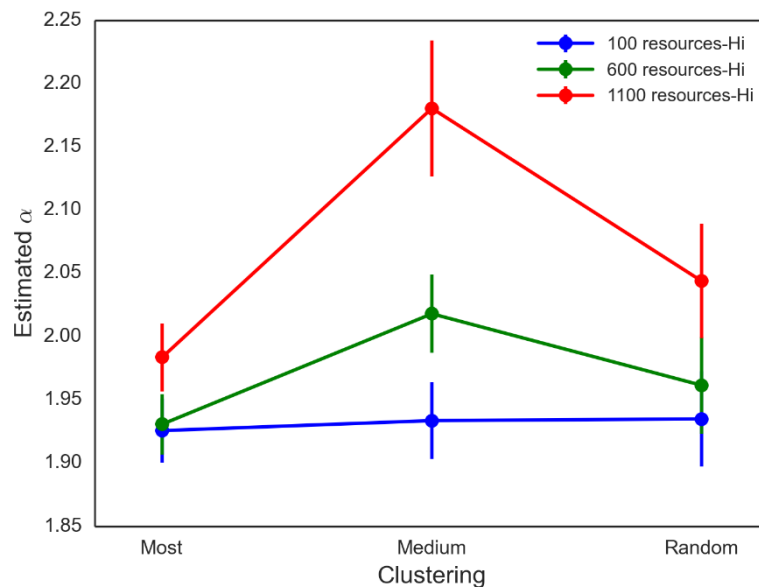


Figure 22: Average α values across each condition.

While there is less overall variation in the search strategies utilized by participants in this experiment, there is still a clear effect of search clustering on score. For both of the clustered conditions higher α values correlate with higher search performance. The pattern observed across clustering remains the same as before, however when measured across density there is no longer a correlation.

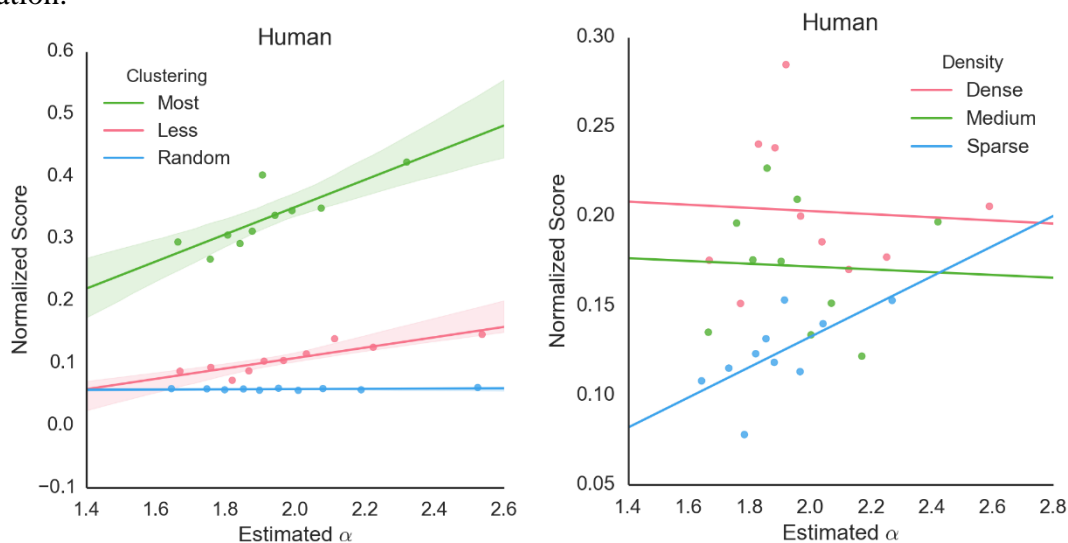


Figure 23: Demonstration of how α correlates with score at different levels of clustering (left) and density (right). Unlike previous experiments there is no clear correlation across resource density.

This experiment succeeded in replicating many of the findings from both experiments 1 and 2. Most notably the path lengths remained Levy distributed with exponents near 2. In fact, granting improved memory slightly increased clustering as measured by the power law exponent value, and brought it closer to the theoretical optimum of 2. This experiment also confirmed several of the assumptions tested by the model, by demonstrating the importance of spatial memory in foraging performance. Foraging performance was nearly doubled in this experiment largely due to the effects of improving spatial memory.

This experiment also revealed how task dependent some of the findings from experiment 2 were. Both the effect of clustering on estimated α values and the interaction between estimated α and score are quite different. Given the size of the samples involved in both experiments it is reasonable to believe these are real differences in behavior driven by the subtle differences in the task constraints. Unfortunately, teasing out how these constraints are interacting with these effects will have to remain a task for future research.

Dyadic Search Model

Organisms rarely have the opportunity to forage a physical space in isolation. Typically, there will be other organisms competing or cooperating with them. This extends beyond physical search and foraging as well. People frequently engage in collaborative cognitive search tasks together. Examples of collaborative search tasks can range from problem solving to memory recall to visual search. Interestingly, while people regularly engage in collaborative search tasks,


it has been shown that groups are often outperformed by individuals (Rajaram & Pereira-Pasarin, 2010). The model presented earlier offers the opportunity to explore questions of collaborative search from a foraging perspective. The model is based on new findings, as well as findings from two competing areas of foraging research and so should be a good tool for exploring collaborative search behaviors.

The work presented in this section is both preliminary and exploratory. Multi-agent foraging and search are rich and important areas of future study, and a natural extension of the work presented here thus far. However, moving to multiple agents also amplifies many of the difficulties of studying foraging through increased complexity. Multi-agent collaborative search, is already an area of study for many fields, but adding a foraging perspective would provide a new and valuable approach.

In this exploratory work the model was used to look at how a variety of different dyadic interaction styles affect several key foraging behaviors. Specifically, how do clustering behaviors change as measured by power law exponents, and how does search efficacy and efficiency change?

Minor modifications were made to the foraging model to allow for the exploration of three different dyadic foraging approaches: attraction, avoidance, and following. The model was adjusted such that two agents were able to act on the same space by taking turns. Each model in the space “knows” the locations that have been searched by the other model in the space. This prevents either model from moving to **any** previously visited location. In order to allow interaction between the two models, each model had an “attraction” value that they associated with the other agent. That value was placed as a trace value for any location that the other agent visited. For the control condition in which two models forage the same space without interacting the trace values were set to 0. This meant that the locations visited by the other agent had no effect on the probability of the current agent visiting any given location. The only exception to this being that the current agent no longer had visited locations available.

To generate interactive patterns of behavior these values were adjusted. Setting the values positive for both agents led to an attraction dynamic. The visited locations of the other agent served as small attractors to the current agent, in much the same way that finding a resource would. Generally this created a dynamic where if one agent spent a large amount of time in one location it is likely that the other agent would also visit that general location. Effectively the agents are able to share locations where they have been successful.



$$\begin{array}{ccc}
 \mathbf{P} \left[\exp \left(\gamma \frac{E_{xy} + 1}{\beta M_{xy} + 1} \right) \right] & & \mathbf{P} \left[\exp \left(\gamma \frac{E_{xy} + 1}{\beta M_{xy} + 1} \right) \right] \\
 E_{xy} = \sum \frac{V_i}{D_i} & & E_{xy} = \sum \frac{V_i}{D_i} \\
 V_i = \begin{cases} N_{xy}, & N_{xy} > 0 \\ -\ln(E_{xy}), & N_{xy} = 0 \end{cases} & & V_i = \begin{cases} N_{xy}, & N_{xy} > 0 \\ -\ln(E_{xy}), & N_{xy} = 0 \end{cases}
 \end{array}$$

Fig 24: The dyadic model is identical in function to the single model with one exception. The models take turns selecting locations, and locations selected by other agents are treated by the non-moving model as having been visited with a trace value equal to that model's attraction value parameter.

If the values for both agents are set to negative numbers then a repelling dynamic was generated. Each agent was slightly repelled by the locations that the other agent has visited. This creates a dynamic that is very much the opposite of the attraction dynamic. Instead of sharing locations where the agents have been successful they avoid them. One could imagine that this would be a useful strategy as it should lead to the agents exploring separate areas.

Finally, a kind of following dynamic was created by setting the values for one agent to be positive, and the values for the other agent to be negative. This led to a kind of following dynamic in which one agent is attracted to the other, who in turn is repelled away from the first. This had the possibility of leading to a dynamic in which agent effectively acted as a “searcher”, while the other agent acted as a “collector.”

These “interaction values” were effectively free parameters that I chose to set to +0.5 or -0.5 depending on the condition. These values were chosen by performing a parameter search and selecting values that were high enough to lead to interactive behaviors between the two agents while also low enough that they did not completely dominate the behavior of the agents. Choosing fixed values for these parameters also led to interesting behavior, where when nothing has been found or the probability distribution is relatively flat for an agent (each agent generates separate probability values based on the resources they find) the interaction component became more important. When an agent had more local information about the space (i.e. has recently found resources) then the behavior of the other agent had less of an effect.

These four conditions (control, attract, avoid, follow) were each tested by running the model 100 times on each density and clustering condition used by both experiment two and the solo model. This means there were 900 runs of the model per dyadic interaction condition. Each individual agent in the dyadic conditions moved 150 steps, for a total of 300 steps between the two agents. This allowed for direct comparisons between the interaction conditions of the dyadic model, as well as with the solo model.

Results

Before discussing the results in more detail there are several qualitative differences worth noting between the interaction conditions. These can be easily observed in the example images provided in figure 25. The major difference worth noting is in how movements cluster in the

repel condition. In this condition movements appeared more tightly clustered, with many small tight clusters of movements. This was an effect of the models having many small repellers placed throughout the space. This created many “hills” that made it more difficult for the model to leave a location, and more likely that it would explore only directly next to the location it was currently at before leaving (assuming no new resources are found.) Essentially these repellers increased the “cost” of visiting new locations. This same effect can also be observed in the follower condition, but only on the agent that is repelled.

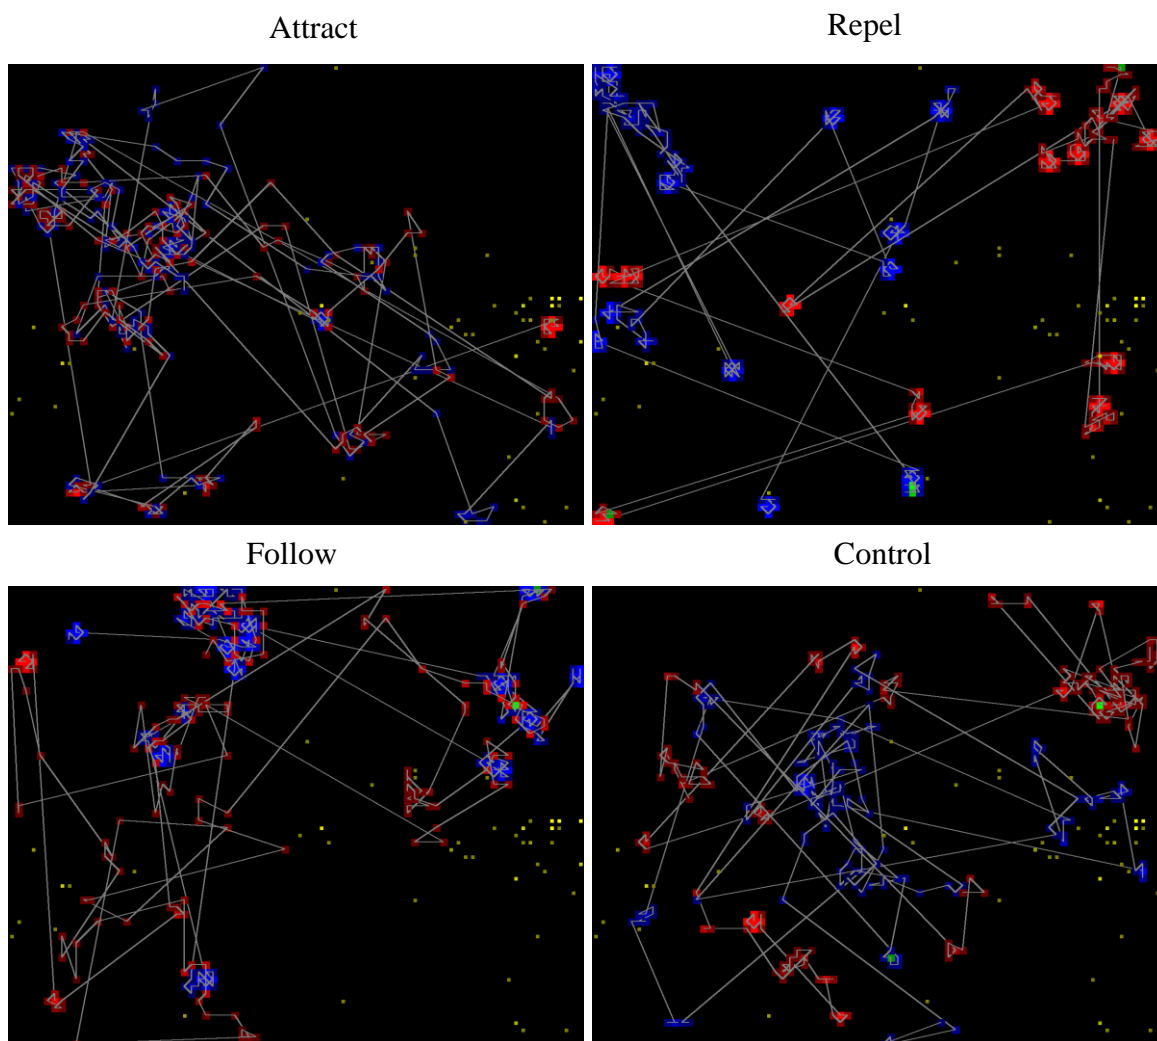


Figure 25: Examples of model movements. In order the images are examples of the attract, repel, follow, and control conditions. For each image yellow dots indicate the location of resources, and blue and red squares indicate the locations each agent visited. For the follow condition red is the follower and blue is repelled.

Foraging path length distributions were generated for each condition as described in experiment 2. In prior examples of path length distributions presented earlier in this document distributions generally appeared mostly flat in logarithmic coordinates with slopes near -2. Some of the dyadic interaction conditions demonstrated some surprisingly different trends. See figure 26. Visually the attract condition, and the attract agent for the follow condition appeared as expected. The repel condition differed in that many of the distributions had longer paths in

greater frequency leading to a shift in the top of the distribution. Much of this effect can be attributed to the difference in how the agents in the repel condition move that was described earlier. Surprisingly, the control condition also demonstrated a similar effect although to a smaller magnitude. This was caused by each agent blocking locations that the other agent would have normally visited. This caused the model to pick longer movements than it otherwise would have thus shifting the distribution. One can be confident that this was the cause for the control condition because that is the only difference between the solo model and the dyadic model in the control condition. The attract condition appeared “normal” because the extra attractors created a tendency for shorter path lengths to pull the distribution towards the right compared to the control condition. The repelled agent in the follow condition provided yet another changed behavior. Rather than shifting as in the repel condition, it instead bowed outwards with a flatter distribution towards the midsized to longer path lengths.

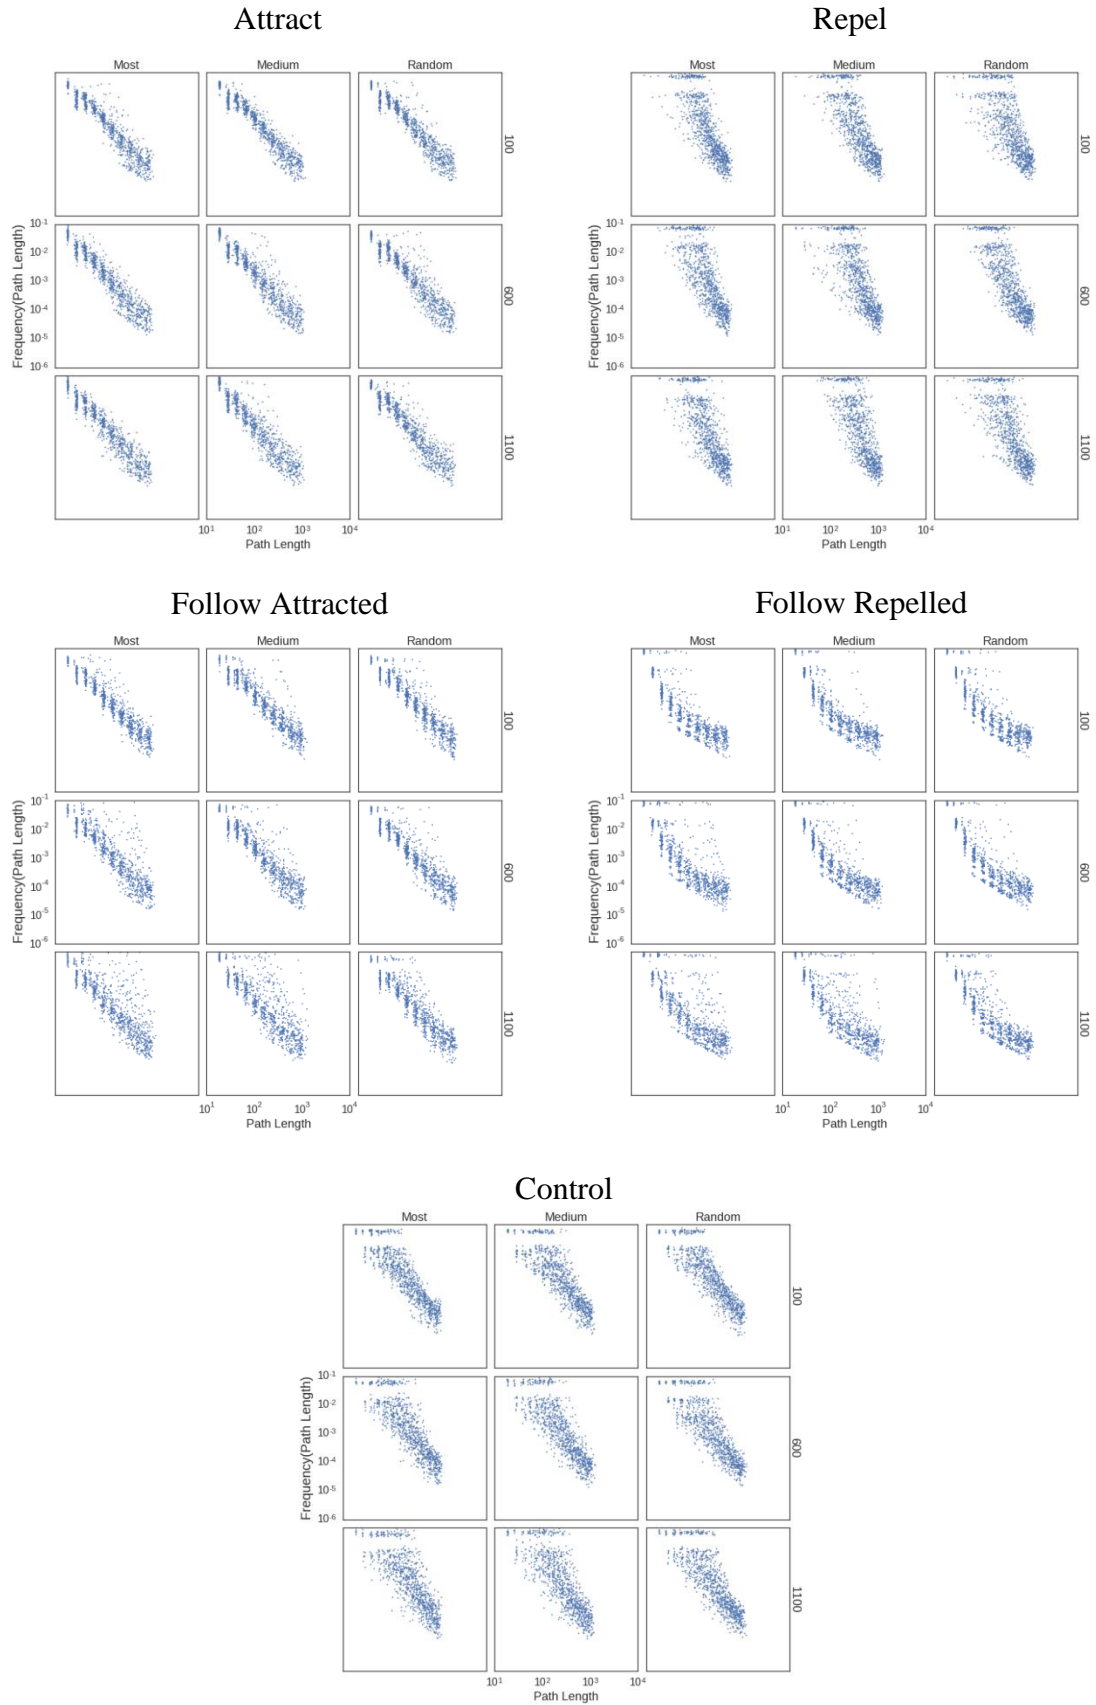


Figure 26: Logarithmically binned plots of the path length distributions produced by the models. The interaction conditions are in order: attract, repel, follow attracted, follow repelled, and control.

Differences in search behavior were further explored by looking at the clustering in the search movements. This was measured in the same way as before by using the method developed by Clauset, et al. (2009) to estimate the power-law slope of the distribution. Table 4 provides a breakdown of estimated α values and their variance. This view of the data revealed that the attract condition is more tightly clustered than any of the other conditions which were much closer to the expected value of 2. This view of the data also revealed that there was much more variance in path length distributions with the dyadic model compared to the original solo model. Many of the differences that can be seen in the distribution plots are at least partially driven by the increased variance.

	Estimated α	Standard Deviation
attract	2.419	0.2653
follow	2.105	0.2394
independent	2.060	0.2608
repel	2.080	0.3132
solo	1.849	0.1567

Table 4: Clustering of each dyadic condition as measured by estimated α .

So far we can see that the dyadic model led to remarkable shifts in model behavior. Changes in behavior were based largely on how agents were interacting however, increased variation occurred regardless. But how did these shifts in behavior affect model performance in finding resources? Based on existing work on collaborative search one would have expected a drop in performance from the solo model. Based on Viswanathan's theories one might have expected the conditions with α values closest to 2 to perform best.

	% Found	Efficiency
attract	0.1086	7.2257
follow	0.1064	6.9719
independent	0.1175	7.5726
repel	0.0944	6.7469
solo	0.1278	7.8484

Table 5: Average performance and efficiency for each dyadic condition. For the solo condition efficiency is measured over the first 150 steps. (All other conditions only have 150 steps). Values are also normalized for easier interpretation.

In this case the results fell in line with what would be predicted based on behavioral research on collaborative search. All of the dyadic conditions performed worse in terms of resources found than the solo model. The fact that the model's performance matches the existing literature may be a sign that this is a valid approach to studying more complicated dyadic search problems.

However, these results also raised a question that often comes in collaborative search research. If collaborative search is inferior why do people do it? Some theories suggest that more people allows for work in parallel saving time overall, or that having more people may also increase the chance of finding something difficult (although this is more applicable to larger groups than pairs) (Panait & Luke, 2005). The results from these simulations may offer an alternative theory. Collaborative foraging may be more efficient, even if you are finding fewer overall resources per unit time. Efficiency is operationalized here in terms of more resources per unit of energy expended, in this case energy is measured as the distance traveled.

Efficiency was measured for each of the dyadic conditions by taking the percentage of resources found and dividing by the total distance moved by the agent. The values presented in table 5 were also normalized for the sake of easy viewing. For the solo model efficiency was only measured by the first 150 steps of movement so as to have an even comparison between the single dyadic agents and the solo agents. Even using this measure which should favor the dyadic models the solo condition remained the best performing.

This set of dyadic models served as an early exploration of how the findings presented earlier could be applied to foraging situations involving multiple agents. It also showed that multi agent foraging can lead to different search behaviors. These findings are limited because interaction between agents was implemented in a very basic and simplified way, and changes in how interaction is implemented will likely lead to further changes in behavior.

Conclusion

Foraging is an ancient and nearly universal search process. Prior work by several researchers has demonstrated probable connections between the processes by which animals forage for food, and various forms of cognitive search including memory and visual search. This dissertation has attempted to utilize the power of crowd sourcing in order to gain a better understanding of how people forage, and answer questions about what leads to some of the previously observed foraging behaviors.

Inspired by the Levy flight approach to foraging, experiment 1 tested the effect of a variety of different resource distributions on foraging behavior in a task with constraints inspired by both the movement of the eye, and limitations of the flying albatross. This experiment revealed that people will produce a power-law like clustering of search movements pretty much regardless of underlying resource distribution, and that they will adjust their search distribution in response to current environmental conditions. This experiment also revealed that people are not moving in Levy flights, as their movements were clearly not memory-less, as would be predicted by a random walk model.

Experiment 2 simplified the constraints of the task and attempted to explore in greater detail how people respond to resource distributions, and whether memory aids would boost search efficacy. This experiment confirmed the results of the first one, and demonstrated that people's search patterns could be receptive to both clustering and sparsity in the environment.

Based on the findings of the first two experiments a model was developed informed by both the Levy flight, and area restricted search approaches to foraging. This model had several key features: it utilized spatial memory, chose locations based both on prior experience, and distance from the current location, and produced inherently clustered movements. This model proved to be a good fit for the behavioral data with one key exception. It proved to be much better at foraging than people in experiment 2.

Experiment 3 served as both a follow up to experiment 2 and the model, by attempting to provide people with two of the key benefits of the model: perfect spatial memory, and perfect accuracy. This experiment confirmed the results of the model when people demonstrated performance at or above the level of the model. Specifically, this experiment demonstrates how powerful a tool spatial memory can be in increasing search performance.

Finally, some explorative work was performed to begin to expand this line of research to foraging situations involving multiple searchers. The model was modified to allow for two agents to forage and interact in the same space. This early work demonstrated the value of future research in this direction, by matching some basic predictions of current collaborative search literature, and allowed for testing hypotheses about why people may choose to engage in collaborative search even though it frequently leads to poorer results.

Overall this line of research has answered many basic questions about how people forage. I have demonstrated that people will produce power-law clustered patterns of movements, almost regardless of their environment, but that they will also adjust those distributions to match their current context. Further I have demonstrated the importance of spatial memory in basic foraging processes. All of these findings were made in the specific context of foraging, but I believe that all of these findings are also applicable to a variety of cognitive search tasks with sufficiently similar constraints. To some extent this findings could be useful for the study of visual search, memory search, information search, and problem solving, as all of these are search problems of some kind. I have also begun the expansion of this research to collaborative search between multiple agents, which I think could eventually add to our understanding of a great number of collaborative tasks.

Future Directions

Search processes are important to a wide variety of human cognitive tasks. My approach to studying human search is based on findings in animal foraging, and has focused thus far on a very basic and simplified foraging type task in order to more closely understand the basics of how people forage in an environment. There is already a variety of evidence for similarity between the foraging patterns observed in wild animals, those observed in my series of experiments, and various cognitive domains such as visual and memory search. One area of future work would be to look again at these areas now that we have a greater understanding of how humans search, and how different search environments affect search patterns. For example, it would be interesting to compare eye-tracking data from a visual search task with how my computational model performs searching the same image.

Another promising area of research is infant language development. Pilot data from Anne Warlaumont's lab indicates that infant vocalizations may also be Lévy-distributed similarly to animal foraging distributions. Examining this in detail may provide an important new understanding of how infants learn to vocalize. A very ambitious study could possibly even examine the vocal environment around the child to see if it affects the distribution produced by the infant. It may even be possible to correlate background "resource" distribution, and the infant's distribution with the speed at which they reach speech milestones.

Another possible area for future research is an applied approach. I see two interesting possible directions this line of research this could take. The first of this is to use what I have learned about human foraging in the form of our computational model to use in Artificial Intelligence (AI) foraging tasks. One could modify the model for use in unmanned drones in a variety of foraging tasks such as locating water, or search and rescue applications.

Another applied application could be to create tools that will help human searchers (like search and rescue workers, or someone just trying to find their keys). For example, a phone app that simply marks where a forager has been may be sufficient to boost human searching, or alternatively a program that aids human searchers by suggesting locations to search based on where they have been, and where they have found "resources." Such an approach would take out some of the difficulties associated with a more pure AI approach, like not having to have a direct method of identifying resources.

References

- Austin, D., Bowen, W., & McMillan, J. (2004). Intraspecific variation in movement patterns: modeling individual behaviour in a large marine predator. *Oikos*, 1(September 2003).
- Ayala-Orozco, B., Cocho, G., Larralde, H., Ramos-Fernandez, G., Mateos, J. L., & Miramontes, O. (2004). Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*). *Behavioral Ecology and Sociobiology*, 55(3), 223–230.
- Bartumeus, F., Peters, F., Pueyo, S., Marrasé, C., & Catalan, J. (2003). Helical Lévy walks: adjusting searching statistics to resource availability in microzooplankton. *Proceedings of the National Academy of Sciences of the United States of America*, 100(22), 12771–5.
- Benhamou, S., & Collet, J. (2015). Ultimate failure of the Lévy foraging hypothesis: two-scale searching strategies outperform scale-free ones even when prey are scarce and cryptic. *Journal of Theoretical Biology*.
- Bertrand, S., Bertrand, A., Guevara-Carrasco, R., & Gerlotto, F. (2007). Scale-invariant movements of fishermen: the same foraging strategy as natural predators. *Ecological applications : a publication of the Ecological Society of America*, 17(2), 331–7.
- Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *Journal of General Psychology*.
- Breed, G. A., Severns, P. M., & Edwards, A. M. (2015). Apparent power-law distributions in animal movements can arise from intraspecific interactions. *Journal of The Royal Society Interface*, 12(103), 20140927.
- Cain, M. S., Vul, E., Clark, K., & Mitroff, S. R. (2012). A Bayesian Optimal Foraging Model of Human Visual Search. *Psychological Science*, (August).
- Charnov, E. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology*, 9(2), 129–136.
- Clauset, A., Shalizi, C., & Newman, M. (2009). Power-law distributions in empirical data. *SIAM review*.
- Emlen, J. (1966). The role of time and energy in food preference. *American Naturalist*, 100(916), 611–617.
- Germine, L., Nakayama, K., Duchaine, B. C., Chabris, C. F., Chatterjee, G., & Wilmer, J. B. (2012). Is the Web as good as the lab? Comparable performance from Web and lab in cognitive/perceptual experiments. *Psychonomic bulletin & review*.
- Henderson, J. (2003). Human gaze control during real-world scene perception. *Trends in Cognitive Sciences*, 7(11), 498–504.
- Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal Foraging in Semantic Memory. *Psychological review*.
- Hills, T. T., Kalf, C., & Wiener, J. M. (2013). Levy Processes and Area-Restricted Search in Human Foraging. *PLoS ONE*, 8(4), 1–7.
- Humphries, N. E., Queiroz, N., Dyer, J. R. M., Pade, N. G., Musyl, M. K., Schaefer, K. M., ... Sims, D. W. (2010). Environmental context explains Lévy and Brownian movement patterns of marine predators. *Nature*, 465(7301), 1066–9.

- Jones, M. N., & Mewhort, D. J. K. (2007). Representing word meaning and order information in a composite holographic lexicon. *Psychological review*, *114*(1), 1–37.
- Kareiva, P., & Odell, G. (1987). Swarms of predators exhibit “preytaxis” if individual predators use area-restricted search. *American Naturalist*, *130*(2), 233–270.
- Kerster, B. E., Kello, C. T., Rhodes, T., & Bien-Aime, R. J. (2013). Adaptive Foraging : Effects of Resource Conditions on Search Paths in a Web-Based Foraging Game. *The Annual Meeting of the Cognitive Science Society*.
- Liversedge, S., & Findlay, J. (2000). Saccadic eye movements and cognition. *Trends in cognitive sciences*, *4*(1), 6–14.
- MacArthur, R., & Pianka, E. (1966). On optimal use of a patchy environment. *American Naturalist*, *100*(916), 603–609.
- Mackas, D. L., & Boyd, C. M. (1979). Spectral analysis of zooplankton spatial heterogeneity. *Science*, *204*(4388), 62.
- Mårell, A., Ball, J., & Hofgaard, A. (2002). Foraging and movement paths of female reindeer: insights from fractal analysis, correlated random walks, and Lévy flights. *Canadian Journal of Zoology*, *86*(5), 854–865.
- Panait, L., & Luke, S. (2005). Cooperative Multi-Agent Learning : The State of the Art. *Autonomous Agents and Multi-Agent Systems*, 387–434.
- Radicchi, F., & Baronchelli, A. (2012). Evolution of optimal Lévy-flight strategies in human mental searches. *Physical Review E*, *85*(6), 1–8.
- Raichlen, D. a, Wood, B. M., Gordon, A. D., Mabulla, A. Z. P., & Marlowe, F. W. (2013). Evidence of Lévy walk foraging patterns in human hunter-gatherers. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(2).
- Rajaram, S., & Pereira-Pasarin, L. P. (2010). Collaborative memory: Cognitive research and theory. *Perspectives on Psychological Science*, *5*(6), 649–663.
- Rhee, I., Shin, M., Hong, S., Lee, K., & Kim, S. (2011). On the levy-walk nature of human mobility. *IEEE/ACM transactions on networking*, *19*(3), 630–643.
- Rhodes, T., Kello, C., & Kerster, B. E. (2011a). Distributional and temporal properties of eye movement trajectories in scene perception. *Proceedings of the 33rd ...*, 178–183.
- Rhodes, T., Kello, C. T., & Kerster, B. E. (2011b). Distributional and Temporal Properties of Eye Movement Trajectories in Scene Perception. *The Annual Meeting of the Cognitive Science Society*.
- Rhodes, T., & Turvey, M. (2007). Human memory retrieval as Lévy foraging. *Physica A: Statistical Mechanics and its Applications*, *385*(1), 255–260.
- Shlesinger, M., Zaslavsky, G., & Klafter, J. (1993). Strange kinetics. *Nature*.
- Snow, R., O’Connor, B., Jurafsky, D., & Ng, A. (2008). Cheap and fast---but is it good?: evaluating non-expert annotations for natural language tasks. ... *Methods in Natural Language ...*, (October), 254–263.
- Stephen, D. G., & Mirman, D. (2010). Interactions dominate the dynamics of visual cognition. *Cognition*, *115*(1), 154–65.

- Symonds, M. R. E., & Moussalli, A. (2010). A brief guide to model selection, multimodel inference and model averaging in behavioural ecology using Akaike's information criterion. *Behavioral Ecology and Sociobiology*, *65*(1), 13–21.
- Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, *268*(5217), 1632–1634.
- Thompson, G., Kello, C., & Montez, P. (2013). Searching Semantic Memory as a Scale-Free Network: Evidence from Category Recall and a Wikipedia Model of Semantics. *Proceedings of the Cognitive Science Society*, 3533–3538.
- Troyer, A. K., Moscovitch, M., & Winocur, G. (1997). Clustering and switching as two components of verbal fluency: evidence from younger and older healthy adults. *Neuropsychology*, *11*(1).
- Viswanathan, G., Afanasyev, V., Buldyrev, S., Murphy, E., Prince, P., & Stanley, H. E. (1996). Lévy flight search patterns of wandering albatrosses. *Nature*, *381*(6581), 413–415.
- Viswanathan, G., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da, Raposo, E. P., & Stanley, H. E. (1999). Optimizing the success of random searches. *Nature*, *401*(6756), 911.
- Wolfe, J. (2013). When is it time to move to the next raspberry bush? Foraging rules in human visual search. *Journal of Vision*, *13*, 1–17.
- Wolfe, J. M., Horowitz, T. S., Van Wert, M. J., Kenner, N. M., Place, S. S., & Kibbi, N. (2007). Low target prevalence is a stubborn source of errors in visual search tasks. *Journal of experimental psychology. General*, *136*(4), 623–38.

Appendix A

The model has two free parameters, γ and β , and thus four combinations of extreme values that lead to aberrant model behaviors. When γ is very low, β has no appreciable effect on model behavior, which means there are actually only three types of aberrant behaviors at the parameter extremes. These types are demonstrated in the following three figures.




Figure A1. Example model run with high β and moderate-to-high γ values. The movement cost becomes overly strong and forces movements to mostly remain near one location.

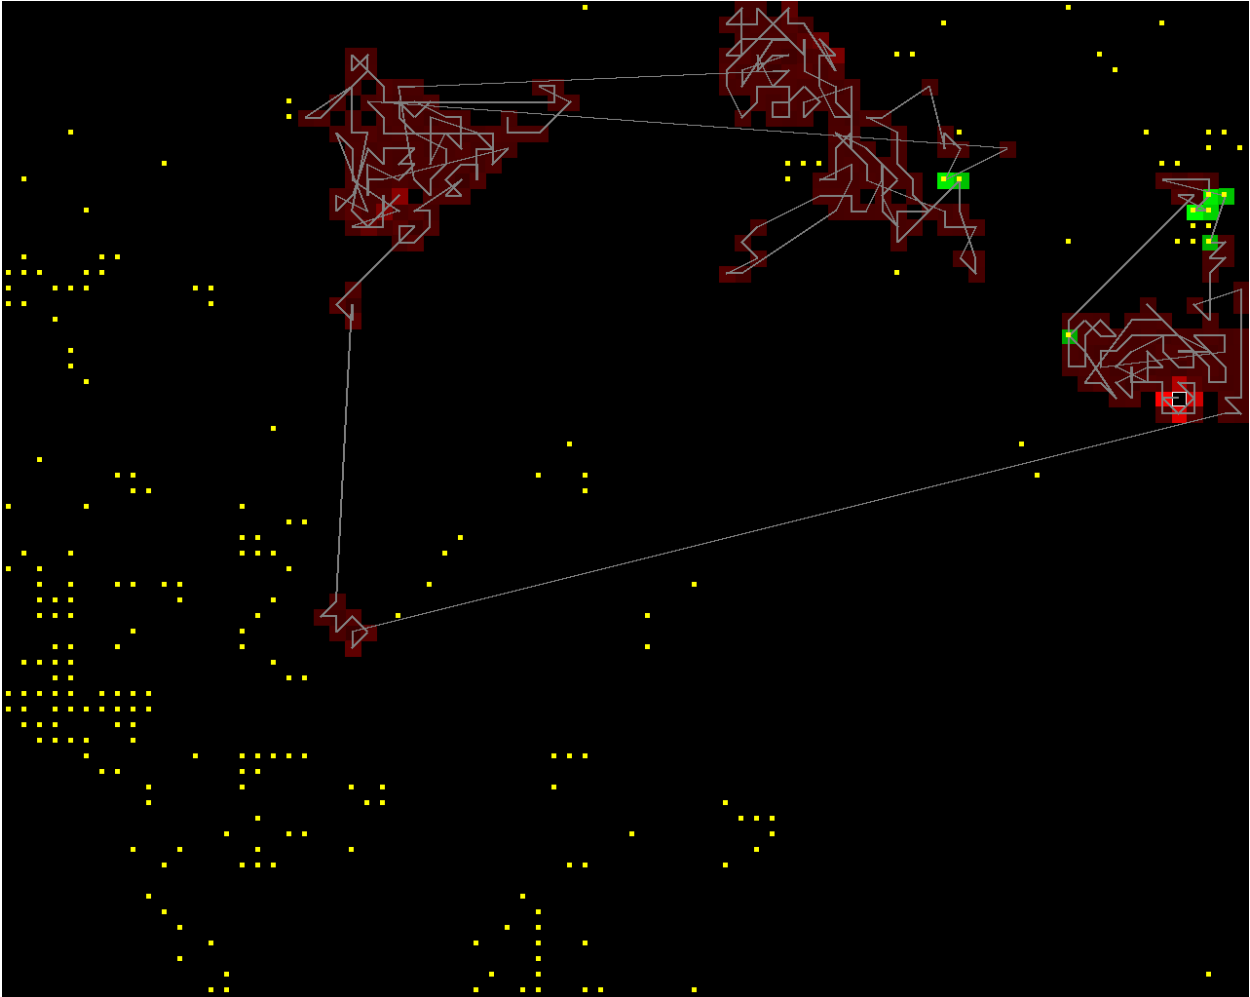


Figure A2. Example model run with high γ and low-to-moderate β . High gain makes the probability distribution overly peaked, which causes the model to make movements among a small number of overly dense clusters.




Figure A3. **Example model run with low γ paired with any β values.** Low gain makes the probability distribution too flat, which causes movements to become random and hence not sufficiently clustered.