

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Adaptive Foraging: Effects of Resource Conditions on Search Paths in a Web-Based Foraging Game

Permalink

<https://escholarship.org/uc/item/7vh330tw>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Kerster, Bryan
Kello, Christopher
Rhodes, Theo
et al.

Publication Date

2013

Peer reviewed

Adaptive Foraging: Effects of Resource Conditions on Search Paths in a Web-Based Foraging Game

Bryan Elvis Kerster (bkerster@ucmerced.edu)
Christopher T. Kello (ckello@ucmerced.edu)
Cognitive and Information Science, 5200 N. Lake Road
Merced, CA 95343 USA

Theo Rhodes (theorhodes@gmail.edu)
Cognitive Science and Psychology, 7060 Route 104
Oswego, NY 13126 USA

Ralph Jerry Bien-Aime (rbien-aime@ucmerced.edu)
Cognitive and Information Science, 5200 N. Lake Road
Merced, CA 95343 USA

Abstract

Foraging is a search process common to mobile organisms, and foraging paths commonly exhibit statistical patterns akin to Lévy walks. There may be common factors and benefits underlying these patterns, but investigations are hindered by difficulty in assessing and manipulating search environments and task conditions. In the present study, a simple foraging game was developed to isolate and manipulate two factors hypothesized to make Lévy walks adaptive search strategies—sparsity, and spatial clustering of targets in the search environment. Players navigated a fuel-limited ship over a 2D grid to find as many targets as possible, rendered as asteroids in outer space. Over 1800 participants were recruited to play using Amazon’s Mechanical Turk, in order to widely sample the parameter space defined by degrees of target sparsity and clustering. Observed search paths resembled Lévy walks with memory, and those of high performers were found to vary adaptively with clustering, but not sparsity. Results indicate that Lévy-like walks can emerge from search strategies and algorithms adapted to environments with clustered resources.

Keywords: Foraging; Lévy walks; crowdsourcing; adaptive search.

Introduction

Perhaps the most ancient kind of search function in biological organisms, in terms of evolutionary history, is foraging—moving about one’s environment in search of resources like food, locales like shelter, or other organisms like mates. Studies of animal foraging have found that foraging paths tend to resemble Lévy walks (Viswanathan et al., 1996). Paths are clustered such that most path segments are relatively short, but they are interspersed with longer segments, occasionally much longer. Paths resemble Lévy walks in that distributions of path lengths follow an inverse power law, $P(l) \sim 1/l^\alpha$, where $\alpha \sim 2$. Lévy-like paths are observed for foragers from bacteria (Berg, 1993) to humans (Rhee, Shin, Hong, Lee, & Kim, 2011).

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 maximize the rate of finding targets compared with Gaussian-distributed random walks.

The potential benefits of Lévy walks recently have led cognitive scientists to investigate whether they occur in perceptual, memory, and decision-making search tasks. First were Rhodes and Turvey (2007), who 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. Then, Rhodes, Kello, and Kerster (2011) found that saccade lengths in visual foraging tasks also followed a heavy-tailed distribution resembling efficient Lévy walks, 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 & Buldyrev, 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 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.

In the present 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 us to know and manipulate search conditions. We made the game web-based to collect data from large numbers of participants on Amazon's Mechanical Turk. Recent studies show that Turk yields data comparable to university participant pools (Germiné et al., 2012; Snow, O'Connor, Jurafsky, & Ng, 2008), and we confirmed this in a pilot study 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.

Comparisons of high versus low scorers are critical because they test whether adaptive search is associated with closer-to-optimal (i.e. higher scoring) performance. That is, do high performers adapt their search strategies, as measured by path length distributions, to changes in sparsity and clustering? Do high scoring foraging paths more closely resemble Lévy flights with the theoretically optimal exponent of two? Addressing these questions will provide evidence on 1) whether Lévy-like foraging paths can emerge from searcher-environment interactions on timescales no longer than minutes, and 2) whether adapting paths to sparsity or clustering is associated with better performance.

Methods

The foraging game was framed as a task of exploring outer space to find resources on asteroids (see Figure 1, and <http://cogmech.ucmerced.edu/downloads.html> to play). Participants used a mouse (or functionally equivalent device) to move a spaceship over a 1280x1024 grid of space. Movement was controlled at two scales, *zoomed in* and *zoomed out*. When zoomed out, the entire space was visible at once, and participants clicked on a location to "fly" the ship to that spot (shown by animation). Participants pressed the space bar to zoom in 15X at a given location, at which point they again could navigate the ship via point-and-click. Hubble images were used as background to help engage players by giving the sense of outer space, and to provide environmental cues that are, in general, ubiquitous to natural search conditions. These cues may encourage use of memory in navigation (e.g. Vinson, 1999), which shall be discussed later.

Asteroids were visible only when zoomed in, and resources were harvested by moving to them such that collision occurred between the asteroid and ship graphics. A set amount of fuel was provided for each play of the game,

and fuel usage (shown by a fuel bar) was a linear function of distance traveled, plus a small constant for each zoom in/out. The amount of fuel provided was determined based on pilot work to allow for about 5 minutes per play, and to enable players to find some but not all asteroids. Each successful harvest was indicated by sight and sound, and asteroids could only be harvested once (the un/harvested status of asteroids was not displayed). Each harvest added one point to the score (no fuel was added), and play continued until all fuel was expended.




Figure 1: Example game space shown zoomed out (above), along with the corresponding zoomed in view (below).

Current score is shown in upper left corner, and fuel bar with remaining fuel in red is shown in upper right corner.

The game was designed to mimic foraging as exemplified by aquatic birds hunting for fish, or the eyes scanning a scene to gather visual information. Relatively short movements are made during resource acquisition (while in the water and close to the surface, or during fixations), interspersed with longer-scale movements when no resources are acquired (while flying high above the water, or during saccades when visual information uptake is attenuated; Ross, Morrone, Goldberg, & Burr, 2001). Foraging costs in natural searches (e.g. risk and energy expenditure) were lumped into the fuel cost of travel, although time costs were also a factor, given the natural tendency to minimize time spent foraging.

The game was coded in 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.




Figure 2: Example of 150 asteroids clustered at 0.05, 0.15, 0.25, and 0.5, left to right, top to bottom.

The number of asteroids per play was set at four different levels: 25, 50, 100, and 150. Pilot work indicated that 25 asteroids meant that players occasionally found only a few of them (or even none), and 150 meant that players found asteroids nearly every time the zoomed in. Clustering of asteroids was manipulated at four different levels of a probabilistic parameter: 0.05, 0.15, 0.25, and 0.5. This parameter controlled the probability of dividing asteroids evenly (0.5) or entirely to one side (0.0) in an algorithm that divided a given set of asteroids recursively into alternating horizontal and vertical splits of a given 2D space. Asteroids were placed when only one remained in a given recursively split section of the space (placed at random in the section), and/or when the space could be split no further (see Figure 2 for example asteroid distributions, and Figure 3 for pseudocode). This algorithm created clusters whose sizes followed a nested scaling relation to varying degrees, consistent with findings of scaling law clustering of natural resources (Humphries et al., 2010; Mackas & Boyd, 1979). The algorithm also created asteroid distributions that were independent of Hubble image backgrounds, and participants were informed of this independence during the demo/instruction period.

The full 4x4 factorial of sparsity and clustering levels was tested. Each participant played twice in only one of the 16 possible conditions, chosen at random at the start of each Turk session. The demo and instructions included an

example asteroid distribution for the condition the participant was in, to help them formulate an informed foraging strategy. Players were also encouraged to achieve the best score possible by maintaining a high score board, and allowing high scorers to enter their initials for display to other players.

```
function distributeResources(rectangle, prob_split)
{
  // stop when no stars or no space left
  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
  }

  // split rect in half, alternate between vert and horiz
  (rectangle1,rectangle2) = splitRectangle(rectangle,alternate)

  // randomize bias for placing stars in each half
  if (random_prob() < 0.5) prob_split = 1 - prob_split

  for each star {
    if (random_prob() < prob_split) place star in rectangle1
    else place star in rectangle2
  }
  distributeResources(rectangle1, prob_split)
  distributeResources(rectangle2, prob_split)
}
```

Figure 3: Pseudocode for asteroid distribution algorithm. Note that a smaller parameter leads to increased clustering.

Results

A total of 1,825 game sessions were administered on 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 revealed by low numbers of zoom-in actions. Of the remaining 1,222 participants, 393 played in two or more Turk sessions. Analyses with and without these repeats indicated no qualitative change in results, so both 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 express a very simple memory (i.e. effect of history) in search paths. Two example paths at the two ends of this range are shown in Figure 4. 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. Other paths consisted of apparently 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. Performance was measured as the proportion of available targets found, 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 intermediate performances helped to further distinguish our three categories.

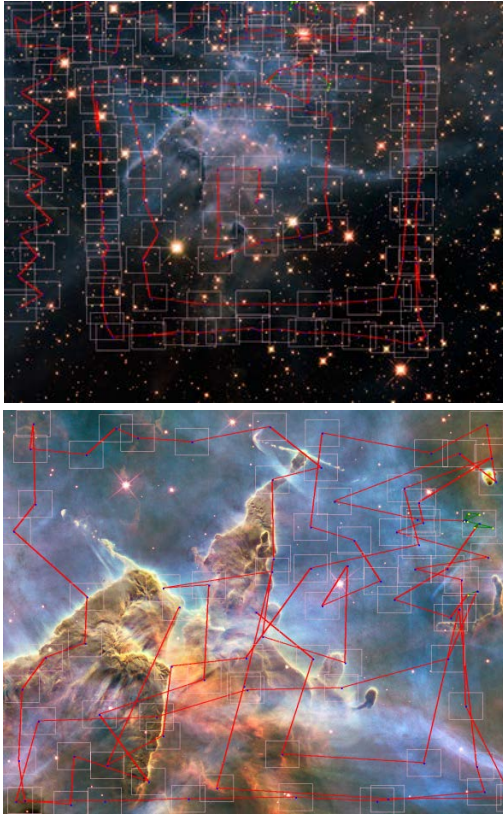


Figure 4: Two example flight paths of more versus less directionally correlated movements (top versus bottom). Red and green lines indicate zoomed out and zoomed in movements, respectively. Blue dots indicate clicked locations, and yellow dots indicate points of harvest. White boxes indicate areas of zoom in.

All results are graphed and analyzed as a function of sparsity, clustering, and performance category. A three-way analysis of variance was conducted for each dependent measure, but we report only main effects and two-way interactions relevant to our research questions and hypotheses. First, we examined score as function of sparsity, clustering, and performance levels (Figure 5). The main effect of performance category is itself based on score, and is so large throughout our analyses that reporting its reliability was unnecessary.

As 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 level 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. Thus we have initial evidence that high performers took advantage of the spatial correlations in clustering, suggesting that foraging paths adapted based on interactions between search processes and game conditions.

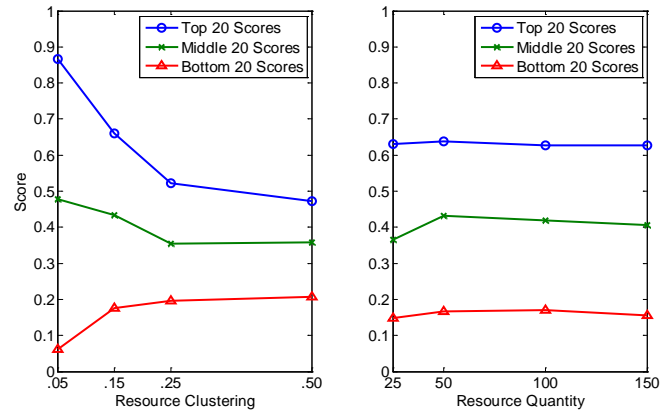


Figure 5: Proportional score as a function of sparsity, clustering, and performance category.

However, one might argue for an alternate interpretation of the data. It may be that each forager chooses a strategy a priori without regard to conditions, and effects of clustering merely show that strategies matter more for greater 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, indicating that strategies were indicative of interactions between search processes and game conditions.




Figure 6: Path Length as a function of sparsity, clustering, and performance category.

Next we examine mean path length, shown in Figure 6. High performers had shorter path lengths overall, reflecting the fact that shorter path lengths allowed for greater coverage of the space by reducing zoom costs. This main

effect would be expected to diminish, and possibly even reverse, if fuel costs were greater per zoom.

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

Next we examine 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 7 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.

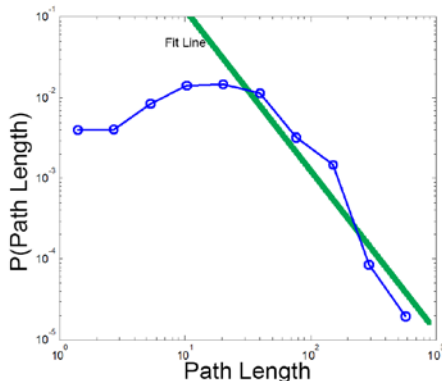


Figure 7: An representative path length distribution for one participant, plotted in logarithmic coordinates

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 σ 's are shown in Figure 8 for all

participants, including those whose data were better fit by normal or exponential distributions.

The overall pattern of results was similar to those reported earlier. 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 reveals 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 (see Figure 7) to the right half of distributions in logarithmic coordinates, and results are shown in Figure 9. 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 (negative of the optimal exponent). Moreover, slopes for high performers were closest in their approach towards -2 with greater clustering.

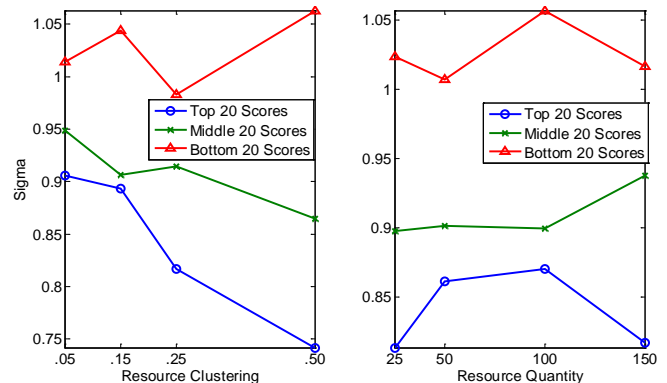


Figure 8: Lognormal σ as a function of sparsity, clustering, and performance category.

To summarize, search paths generally resembled efficient Lévy walks as predicted, at least to some degree. The majority of path length distributions were heavy-tailed, and tails resembled truncated power laws with exponents near two. Distributions most closely resembled Lévy walks for high performers in the most clustered resource conditions. Foraging paths were not like random walks, in that path directions were never drawn purely at random. Instead, path directions tended to be correlated over 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 that the direction of each step was sensitive to previous steps.




Figure 9: Regression slopes as a function of sparsity, clustering, and performance category.

Discussion

In the present study, analyses of data from a web-based foraging game showed that Lévy-like search paths emerge from search processes that change depending on the clustering of resources. This was true for both high and low performers, although only high performers changed their search processes adaptively. The lack of a sparsity effect was conspicuous because prior theoretical analyses suggest that the benefits of Lévy-like search paths are most prevalent when items are rare to be found. However, prior analyses focused on random Lévy walks, whereas observed foraging paths clearly had memory. Memory sometimes manifested as “sweep” searches, but more generally, directions of next steps depended on previous steps.

Our results suggest that the prevalence of Lévy-like walks in both animal and cognitive searches can be better modeled by processes with memory that attend to target and task conditions, as opposed to random walk processes. The observed effects of clustering indicate that search processes take advantage of spatial correlations in resource distributions when they exist and are known or learned. A simple approach to modeling an effect of spatial correlations is for search processes to follow a gradient of resource density. These and related modeling ideas (e.g. Ferreira, Raposo, Viswanathan, & da Luz, 2012) are potentially interesting topics for future research.

Acknowledgments

This work was supported by a grant from the National Science Foundation, BCS 1031903 (PI Kello). The authors thank members of the Cognitive and Information Sciences graduate group for helpful comments and feedback.

References

Berg, H. C. (1993). *Random walks in biology*. Princeton University Press.

Bousfield, W. A., & Sedgewick, C. H. W. (1944). An analysis of sequences of restricted associative responses. *Journal of General Psychology*.

Ferreira, A. S., Raposo, E. P., Viswanathan, G. M., & da Luz, M. G. E. (2012). The influence of the environment

on Lévy random search efficiency: Fractality and memory effects. *Physica A: Statistical Mechanics and its Applications*, 391(11), 3234–3246.

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*.

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.

Mackas, D. L., & Boyd, C. M. (1979). Spectral analysis of zooplankton spatial heterogeneity. *Science*, 204(4388), 62.

Radicchi, F., & Baronchelli, A. (2012). Evolution of optimal Lévy-flight strategies in human mental searches. *Physical Review E*, 85(6), 1–8.

Rhee, I., Shin, M., Hong, S., Lee, K., & Kim, S. (2011). On the levy-walk nature of human mobility. *ACM Transactions on*, 19(3), 630–643.

Rhodes, T., & Turvey, M. (2007). Human memory retrieval as Lévy foraging. *Physica A: Statistical Mechanics and its Applications*, 385(1), 255–260.

Ross, J., Morrone, M., Goldberg, M., & Burr, D. (2001). Changes in visual perception at the time of saccades. *Rev. Neurosci*, 24(2), 113–21.

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), 154–165.

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.

Vinson, N. (1999). Design guidelines for landmarks to support navigation in virtual environments. *Proceedings of the SIGCHI conference on Human ...*

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., Afanasyev, V., Buldyrev, S. V., Havlin, S., Luz, M. G. E. da, Raposo, E. P., & Stanley, H. E. (2001). Lévy flights search patterns of biological organisms. *Physica A: Statistical Mechanics and its Applications*, 295, 85–88.

Viswanathan, G., & Buldyrev, S. V. (1999). Optimizing the success of random searches. *Nature*, 401(6756), 911.