A Bayesian model of navigation in squirrels

Anna S. Waisman (annawais@uw.edu)
Institute for Learning and Brain Sciences, Portage Bay Building, Room 450
Seattle, WA 98195 USA

Christopher G. Lucas (cglucas@gmail.com)
Department of Educational Psychology, Baker Hall 342c
Pittsburgh, PA 15213 USA

Thomas L. Griffiths (tom_griffiths@berkeley.edu)
Lucia F. Jacobs (jacobs@berkeley.edu)
Department of Educational Psychology, 3210 Tolman Hall #1650
Berkeley, CA 94720 USA

Abstract

Fox squirrels have an impressive ability to remember the location of stored food. In doing so, they combine information from landmarks of different types. We define a Bayesian model that indicates how an ideal observer would optimally integrate landmark cues, and fit this model to the decisions made by squirrels in a spatial memory task. The resulting model provides a unifying framework for characterizing different strategies to cue integration, and a tool for investigating the circumstances under which particular cues are used. We show that the best fitting models changed depending on the season at testing and the details of the task. These analyses support previous claims that squirrels adopt flexible strategies in landmark use.

Keywords: spatial cognition; Bayesian modeling; spatial memory; animal cognition; navigation

Introduction

Animals of many kinds display remarkable skill at spatial navigation, and formal models of how animals navigate have many potential uses. For example, one can use them to develop robots capable of autonomous movement (Thrun, 2005) and to aid in designing new animal conservation principles (Fevre, Bronsvoort, Hamilton & Cleaveland, 2006; Simons, 2004). In this paper, we analyze the problem of identifying a spatial location from memory as a kind of Bayesian inference. This approach provides a way to quantify degrees of belief and uncertainty, and thus provides a natural framework in which to develop an “ideal observer” model. In cases where multiple kinds of landmarks are available, the Bayesian approach allows us to take into account the perceived reliabilities of each landmark or landmark type. This information can be used to identify the location most consistent with the animal’s memory. Formalizing spatial memory in these terms gives us a tool for identifying which types of landmarks animals use in navigation, what factors influence the use of these landmarks, and what kinds of strategies animals are using based on how they use those landmarks in combination.

Birds and other animals that are either nectar-feeding or store food have excellent spatial memory abilities due to natural selection pressures, and have often been used in spatial memory experiments. Traditionally, such animals have been described as using landmarks in a hierarchical fashion by which the animal works its way down its preference hierarchy of landmark types until it finds the rewarded location (Brodbeck, 1994; Clayton & Krebs, 1994; Herz, Zanette & Sherry 1994). Bats, hummingbirds, and squirrels have shown such preference hierarchies (Healy & Hurly, 1998; Jacobs & Shiflett, 1999; Thiele & Winter, 2005; Vlasak, 2006a).

This traditional hierarchical model has recently been challenged by converging evidence in favor of the plasticity of landmark use in both mammals and birds (Pigeons: Legge, Spetch & Batty, 2009; Chickadees: LaDage, Roth, Fox & Pravosudov, 2009; Flying Squirrel: Gibbs, Lea & Jacobs, 2007; Fox Squirrel: Waisman & Jacobs, 2008). We were interested in exploring in more detail how different combinations of landmark types trade off in guiding search behavior and whether animals might be using Bayesian inference to determine their search strategy. This would explain the flexibility in their search strategies, and predict their strategies in a wide range of novel situations.

The plan of the paper is as follows. In the first section, we briefly explain the general structure of a Bayesian model of landmark use. In the second section, we describe a series of cue combination experiments with squirrels and the specific model that we used to characterize their behavior. In the third section, we demonstrate how the model can be used to examine factors that influence the use of landmarks. The final section concludes the paper.

A Bayesian analysis of cue combination in squirrel spatial navigation

To construct a Bayesian model of squirrel spatial navigation we must first define the problem of squirrel spatial navigation in Bayesian terms. For simplicity, the model presented here will focus on the navigation problem encountered by a single species – the fox squirrel. A squirrel must rely on environmental landmarks and its memory of those landmarks when searching for food. Environments change and a squirrel’s memory has finite precision. This leads to a navigational problem for which the ideal solution requires accounting for these sources of error by determining the probability that the food is hidden in a particular location. To do this, the squirrel must consider the variability of landmarks including global landmarks, local landmarks, and the perceptual features of candidate locations. More formally, we imagine that the squirrel is choosing between a discrete set of locations L, trying to identify the location $l \in L$ that matches a remembered set of landmarks m. Locations have a history of being used for storage, with $p(l)$ being the probability of each location being used. The squirrel seeks to compute $p(l|m)$, the posterior probability of each location given the information provided by the memory m. This can be done by applying Bayes' rule,

$$p(l|m) = \frac{p(m|l)\,p(l)}{\sum_{l \in L} p(m|l)\,p(l)} \quad (1)$$

where $p(m|l)$ is the probability of the remembered landmarks m given that l was the true location. In order to simplify the problem, we will restrict our attention to the case where the types of landmarks that comprise the memory m and their associated sources of uncertainty are independent, and no location is more likely a priori. In this case, Equation 1 becomes

$$p(l|m) = \frac{\prod_k p(m_k|l)}{\sum_{l \in L} \prod_k p(m_k|l)} \quad (2)$$

where k indexes specific types of landmarks and m_k is the part of the memory corresponding to landmark type k.

The distribution of m_k given l – the probability of recalling that landmark k takes its remembered value given that l is where the food is currently hidden – depends on that landmark's tendency to change and the accuracy and granularity of the squirrel's memory. For example, if $k = 0$ corresponded to color cues, then $p(m_0|l)$ represents the probability of remembering that the color of the location was red given that the color of the true location was blue. The case where the memory for each landmark is normally distributed was explored in Cheng, Shettleworth, Huttenlocher and Rieser (2007), where it was argued that pigeons might combine such landmarks optimally. The basic prediction produced by this account is that animals should rely on the perceived reliability of either individual landmarks or landmark categories.

Existing evidence suggests that this is the case. For example, animals initially preferring one of two landmarks switched their preference when given evidence that the non-preferred landmark was more reliable (Biegler & Morris, 1996). In terms of broader categories of landmarks, the most reliable landmarks are considered to be the global landmarks, the far away, large landmarks that are not only the most immovable objects in the environment, but also the ones that are least distorted by changes in visual angle or across seasonal changes to the local environment (Shettleworth, 2003). This phenomenon is further supported by evidence that many species of animal prefer to use global landmarks when available. However, what would distinguish the behavior of a Bayesian model from that of other models would be the optimal combination of available landmarks based on their reliabilities. In the remainder of the paper, we explore whether this approach can characterize the strategies that squirrels use in identifying spatial locations from memory.

Applying the model to squirrel navigation

To test the model outlined in the previous section, we analyzed data from a series of experiments in which fox squirrels needed to identify a location based on several different types of landmarks, taken from Waisman and Jacobs (2008). These experiments manipulated the environment in which squirrels were making decisions in order to produce conflict in the information provided by different types of landmarks, and thus provided a good test of our model.

Free-ranging fox squirrels were initially trained to feed from a fixed location within a square array of four feeders (see Figure 1a). Identification of a location could be done using three types of landmarks. Extra-array (EX) landmarks were those that were external to the feeder array, including objects such as large trees and bushes. Intra-array (IN) landmarks were the non-rewarded feeders in the array. Unique feature (UF) landmarks were defined as any features unique to the feeder, including scent and the color and shape of ceramic figurines that were placed on top of the boxes.

Squirrels were then tested in two transformed versions of the array. In hierarchy tests, all three landmark types were in conflict with one another, with no two types being consistent in the location they identified (see Figure 1b). In the majority tests, one landmark type was in conflict with the other two landmark types (see Figure 1c). This provided a test of the majority strategy, examining whether the squirrels always chose the location consistent with the majority of the landmark types.

Experimental setup

Data taken from these experiments were originally analyzed using binomial one-tailed tests. See Waisman and Jacobs (2008) for further discussion of the data. In summary, in Experiment 1, squirrels chose the majority location significantly greater than chance and chose the UF location in the hierarchy test ($p < 0.05$ for all tests). Since fox squirrels are known to prefer EX landmarks in hierarchy tests, the choice of UF in this experiment was perplexing. The contrast suggested that the particular experimental setup had somehow increased the saliency of the UF landmarks. To further explore this possibility, squirrels were tested at the same time the following year, summer, using a different experimental setup and found that, while squirrels continued to choose the majority location in all three majority tests, they now chose the EX location in the hierarchy test ($p < 0.05$ for all tests). From this pattern of data, it was concluded that the experimental setup had indeed increased the saliency of the UF landmarks.

To investigate possible seasonal effects on landmark use, the experiment was also run in the spring, using the second experimental setup, with less salient UF landmarks. In this experiment, squirrels no longer chose the majority location when it was the combination of IN and UF landmarks, and chose the EX location in the hierarchy test ($p < 0.05$). The spring data suggest that squirrels were no longer taking into account UF landmarks when making spatial decisions. Taken altogether, the data are consistent with the proposal that squirrels are using a majority strategy when possible, but that there are both seasonal and salience effects on the use of UF landmarks. The results suggested that rather than being limited to a strict hierarchy, as has been proposed in earlier studies, squirrels used a more flexible majority strategy when possible.

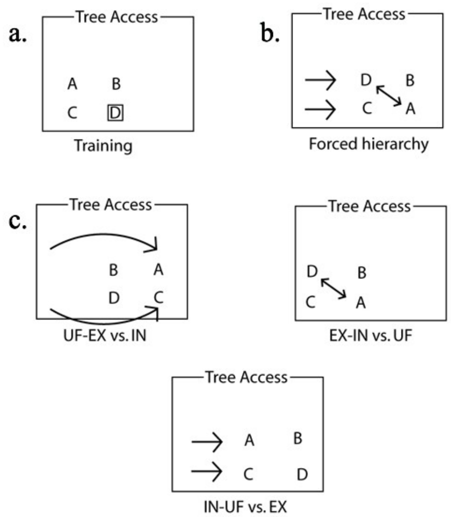

Figure 1: (a) A training trial in which feeder D is rewarded. (b) Hierarchy test: the entire array is moved horizontally and feeder D is switched with feeder A. Unique feature (UF) cues indicate search at D, intra-array (IN) landmarks indicate search at A, and extra-array (EX) landmarks indicate search at C. (c) Majority tests based on having been trained to feeder D.

Applying the model to the experiment

The Bayesian model introduced in the previous section gives us a sophisticated tool for characterizing the strategies that squirrels used in these experiments. Again, the squirrel's problem is to determine which location has the highest posterior probability of being correct given its memory of the location. We will use l to denote candidate locations and their associated landmarks. These landmarks include all three types of landmarks: extra-array, array position, and unique features. The squirrel's memory of the location m contains recalled values for extra-array (m_{EX}), intra-array (m_{IN}), and feature (m_{UF}) cues. The squirrels were naive and had no information besides the landmarks leading them to prefer one particular location over another, and local landmark and feature cues were novel and independently varying, so our earlier formulation (Equation 2) applies, yielding

$$p(l|m) \propto p(m_{EX}|l)\, p(m_{IN}|l)\, p(m_{UF}|l) \qquad (3)$$

Our model is thus defined by specifying the likelihood terms for each of these types of landmarks (e.g. $p(m_{EX}|l)$).

We adopted a slightly different probabilistic model for each landmark type. For global landmarks, EX, we assumed that recalled locations are distributed normally around the true location, in accordance with previous cue integration models: $p(m_{EX}|l) = N(m_{EX}|l, \sigma^2 I)$, where σ^2 is a parameter representing the variance associated with location. The landmark's value is represented as two-dimensional vectors (where the components give a location in centimeters from the true location) and I is the two-dimensional identity matrix. For intra-array landmarks, IN, we characterized the difference between a candidate true location and the recalled location as a perturbation in grid positioning in either of two independent, orthogonal directions. If we let the probability of a perturbation in each direction be ϵ_{IN}, then $p(m_{IN}|l) = (1-\epsilon_{IN})^{2-p}\epsilon_{IN}^{p}$, where p is the number of perturbations out of a possible two. For feature-based cues, UF, if the probability that a feature differs between the hypothesized location and the squirrel's memory is ϵ_{UF} (e.g the probability that the squirrel remembers red but the true location is blue), so $p(m_{UF}|l) = 1-\epsilon_{UF}$ if the recalled and candidate features are identical, else $\epsilon_{UF}/3$ given that there are three other locations with their own distinct features.

Under this specification, the three parameters σ^2, ϵ_{IN}, and ϵ_{UF} capture the reliability of each type of landmark, and by extension the strength of the evidence that that landmark type provides and the likelihood that the squirrel will prefer that type of landmark when all three are in conflict. The results of the experiment indicate the percentage of squirrels choosing a particular observed location, which can be used to estimate these parameters from the data. We compute the log-likelihood of the choices made by the squirrels by assuming that they used a standard probability-matching decision rule, with the probability of choosing a location being equal to the posterior probability of that location.

Model results for all experiments

We can use this Bayesian model as the basis for defining a nested hierarchy of models that differ in the assumptions they make about the landmarks that fox squirrels use in choosing locations (see Figure 2). The simplest model sets the parameters to values that make all landmark type values equally likely: $\sigma^2 = \infty$, $\epsilon_{IN} = 0.5$ and $\epsilon_{UF} = 0.75$. This model corresponds to completely uniform choices of location. The fit between this model and the data is represented by its log likelihood value: -257.92. We can then examine the effects of estimating the parameters associated with each landmark type. The model that only

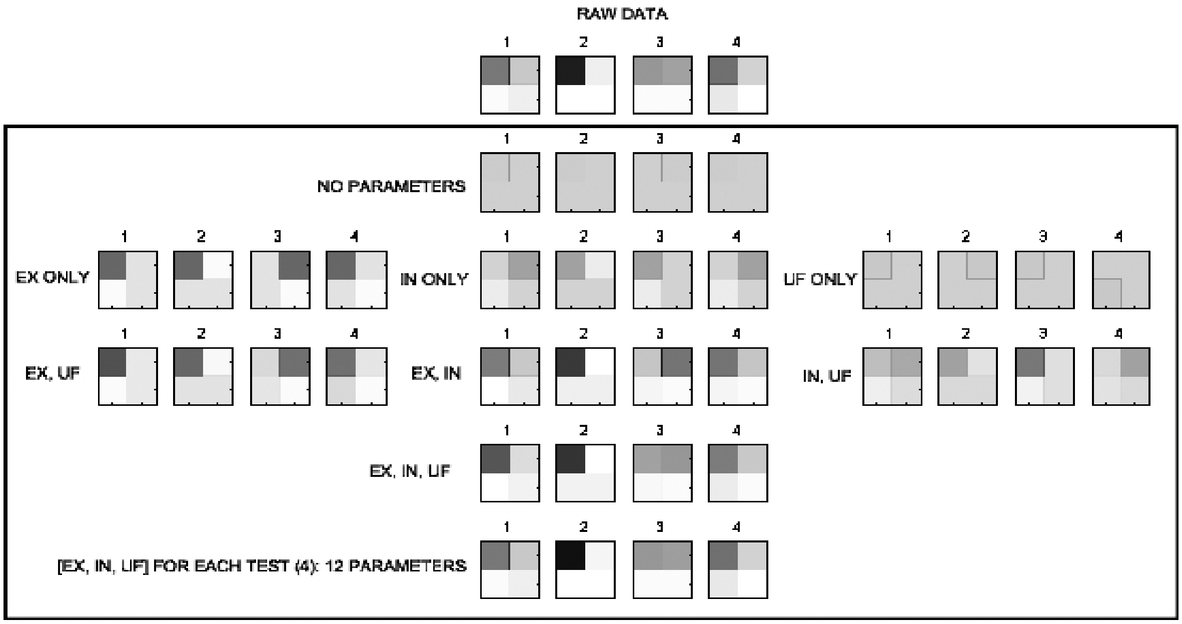

Figure 2: Predictions of each model in the nested hierarchy analysis, labeled with which parameters were estimated and the test type. Plot 1 represents the hierarchy test, while 2, 3, and 4 represent the majority tests. Each quadrant represents the posterior probability of an observed location, corresponding to the four locations shown in Figure 1. Darker indicates higher probability, while lighter indicates lower probability.

estimated the UF parameter, ϵ_{UF}, gave a similar log-likelihood of -257.68. Thus, the UF landmark on its own does not provide a significant amount of information for explaining the data. In contrast, the log likelihood values yielded by estimating either the EX parameter σ^2, or the IN parameter, ϵ_{IN}, on their own were significantly different from that of the chance model (-173.79 and -235.37 respectively; $p < 0.01$ for both, by a likelihood-ratio test).

As one can see in Figure 2, of the two estimated parameter models, the model that estimated both the EX and IN parameters most closely fit with the raw data, with a log-likelihood of -160.11. However, most of this fit seems to be captured by the EX parameter, since the model with EX and UF parameters also does well at fitting the data, with a log-likelihood of -172.48. Even so, the log-likelihood of the model that estimated parameters for all three landmark types, -155.32, was significantly better at predicting the data than any combination of only two parameters or one parameter on its own ($p < 0.01$ for all comparisons). This model estimated the values of the parameters to be 239.01 for σ^2, 0.30 for ϵ_{IN} and 0.60 for ϵ_{UF}. So while it seems that EX has the greatest influence on the choices of the squirrels, their behavior is consistent with that of a model that uses all three landmark types.

Further uses of the Bayesian framework

In this section, we consider how these results relate to squirrels' choice of strategy, and how the model can be used to explore variation in the importance of landmark types due to differences in salience and season.

The majority strategy and Bayes

In the study from which we took our data, when fox squirrels were presented with situations in which each landmark type indicated a different location, they predominantly relied on external, global landmarks. However, when presented with situations in which two landmark types were consistent with one another and in conflict with the preferred, global landmarks, they chose to search in the location consistent with the greatest number of landmark types. Waisman and Jacobs (2008) called this the majority strategy. These results support the idea that squirrels are able to adapt their search strategy to the particular spatial context.

This type of decision strategy can be modeled using a Bayesian approach. Under the same distributional assumptions as before and a maximizing or probability matching decision rule, the squirrel is most likely to visit the location for which $p(m|l)$ is maximal. For this location l_i,

$$p(m_{EX}|l_i)\,p(m_{IN}|l_i)\,p(m_{UF}|l_i) > p(m_{EX}|l_j)\,p(m_{IN}|l_j)\,p(m_{UF}|l_j)$$

for all $j \neq i$. Equivalently, we can represent this relationship in terms of the evidence provided by each feature, i.e., the log odds of the recalled landmark values for location i versus j. For instance, the evidence provided by the intra-array landmarks is $e_{IN}(i,j) = \log\left(\frac{p(m_{IN}|l_i)}{p(m_{IN}|l_j)}\right)$. For the

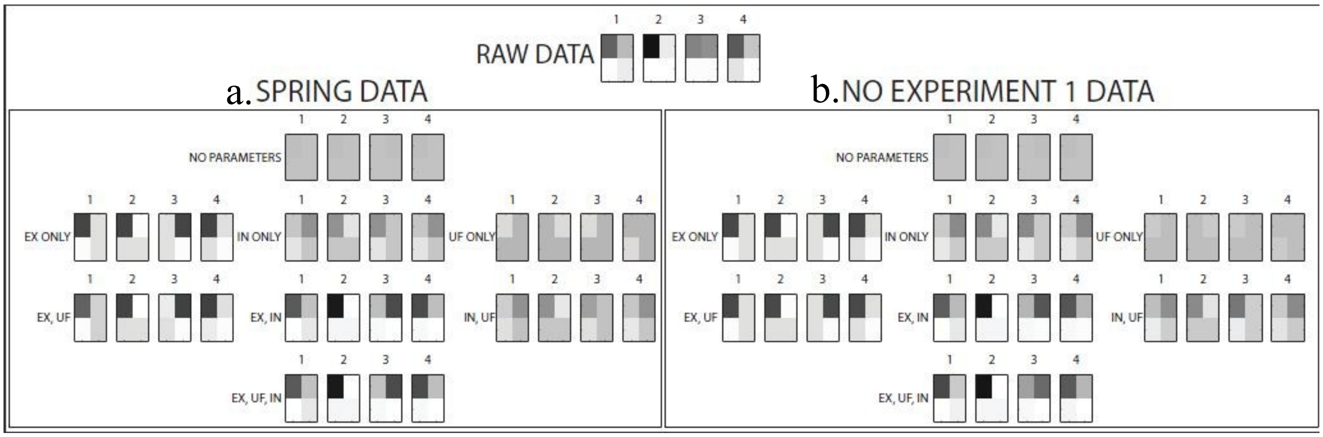

Figure 3: Model predictions for the seasonal and salience analyses. Format is the same as Experiment 1. (a) are the analyses using the subset of the data that were collected only in the spring, Experiments 2-4. (b) are the analyses using the subset of the data that were collected using the second experimental setup, Experiments 2-5.

majority strategy to apply, no one type of landmark may provide evidence dominating the combined evidence from the other two: the maximum possible evidence for one landmark type cannot exceed the the sum of the maximum negative evidence provided by the other two. We can establish some bounds on σ^2 (in terms of a maximal distance d between candidate and recalled locations), ϵ_{IN} and ϵ_{UF} that determine when a majority strategy will no longer apply: violations of

$$max\ e_a(i,j) < max\ e_b(j,i) + max\ e_c(j,i)$$

for any combination of cues a, b, and c imply that, as a group, the squirrels' behavior deviates from a majority strategy.

Determining whether a set of parameters is consistent with the majority strategy requires computing the maximal evidence that can be provided by each kind of cue. In our model, this is $(2\sigma^2)^{-1} d^2$ for EX, $\log((1-\epsilon_{IN})^2 \epsilon_{IN}^{-2})$ for IN, and $\log(\frac{1-\epsilon_{UF}}{\epsilon_{UF}/3})$ for UF. Plugging these values into the constraints identified in the previous paragraph, we obtain the following inequalities:

$$(2\sigma^2)^{-1} d^2 < \log((1-\epsilon_{IN})^2 \epsilon_{IN}^{-2}) + \log((1-\epsilon_{UF}) 3\epsilon_{UF}^{-1})$$

$$\log((1-\epsilon_{IN})^2 \epsilon_{IN}^{-2}) < (2\sigma^2)^{-1} d^2 + \log((1-\epsilon_{UF}) 3\epsilon_{UF}^{-1})$$

$$\log((1-\epsilon_{UF}) 3\epsilon_{UF}^{-1}) < \log((1-\epsilon_{IN})^2 \epsilon_{IN}^{-2}) + (2\sigma^2)^{-1} d^2$$

If these inequalities are satisfied, then our Bayesian model will produce behavior consistent with the majority strategy. Waisman and Jacobs (2008) concluded that squirrels only used the UF cues and thus a majority strategy during the summer season. Accordingly, we evaluated these criteria using parameter estimates from only the summer experiment data. The parameter estimates computed from the summer experiments were: 318.98 for σ^2 0.30 for ϵ_{IN} and 0.40 for ϵ_{UF}. Using these values, the inequalities for the majority strategy were satisfied, suggesting that squirrels' cue combination behavior in the summer season can be explained by the use of a majority strategy.

Capturing variations in salience and season

We analyzed the possible salience and season effects by repeating the analyses from the previous section on two different subsets of the data: one that included only experiments run in the spring, and another excluding Experiment 1 in which the experimental setup included stimuli that seemed to increase the saliency of UF landmarks. Figure 3 summarizes these analyses.

Seasonal effects

Using data from only the spring experiments (Figure 3a), the inclusion of ϵ_{UF} with a log likelihood of -158.99, while significantly better than having no parameters, a log likelihood of -162.20 ($p < 0.05$), did little to enhance any model in which it was included. The inclusion of both IN and EX parameters was, however, significantly better than including either parameter alone (IN: -150.94; EX: -98.70; EX & IN: -89.53; $p < 0.01$ for both comparisons). Unlike the models using the full data set, the log likelihoods of both the model that included only the IN and EX parameters and the model that optimized all three parameters were identical. The model including the EX and IN parameters estimated the parameters to be 209.09 for σ^2 and 0.30 for ϵ_{IN}. These results suggests that the squirrels were not taking UF landmarks into account when making spatial decisions at this time, as concluded in Waisman and Jacobs (2008).

Salience effects

In the analyses that excluded Experiment 1, the experiment that used the setup with more salient UF landmarks, the log likelihood of the model that optimized the IN and EX parameters was once again near identical to that of the model that included all three parameters, with values of -121.8 and -120.7 respectively (Figure 3b). The model that included the IN and EX parameters estimated the parameters to be 210.25 for σ^2 and 0.26 for ϵ_{IN}. Since, once again, the model that best matched the squirrels' behavior was the one that did not optimize the parameter for the UF landmark, these analyses support the idea that they were less salient in the experimental setup.

Summary and Conclusion

Taken together, the results of these analyses illustrate how our Bayesian framework can be used to characterize the landmarks used by animals in navigation. Across all data sets that we ran, the extra- array parameter, σ^2, resulted in a statistically significant improvement in fit whenever it was added to a model ($p < 0.01$ for all comparisons). This agrees with previous research stating that for squirrels in the field the global landmarks are the most salient when navigating (Jacobs & Shiflett, 1999; Vlasak, 2006a, b). The fact that for both the spring data set and the data set excluding the first experiment, optimizing the parameter for the unique feature landmarks, ϵ_{UF}, yielded no predictive power beyond that of the other landmark types, corroborates the conclusion that UF landmarks were less salient to the squirrels both in the spring and when using the second experimental setup. From these analyses we can conclude that a Bayesian model is a useful tool for exploring the spatial strategies. The pattern of choices exhibited by the squirrels matched that of a rational model taking into account all three available landmark types. Additionally, model comparisons provided a tool for investigating both seasonal and salience effects in the data.

Acknowledgments

We would like to thank José Lopez and Ted Claire for constructing the feeders. The research was supported by a grant to L.J. from the University of California at Berkeley Committee on Research, a sabbatical appointment to L.J. at the Santa Fe Institute and funding from the James S. McDonnell Foundation Causal Learning Collaborative. The research followed federal and university animal care and use rules and guidelines and complied with APA ethical standards in the treatment of animals.

References

Biegler, R., & Morris, R. (1996). Landmark stability: Studies exploring whether the perceived stability of the environment influences spatial representation. *Journal of Experimental Biology, 199*(1), 187-193.

Brodbeck, D. R. (1994). Memory for Spatial and Local Cues - a Comparison of a Storing and a Nonstoring Species. *Animal Learning & Behavior, 22*(2), 119-133.

Cheng, K., Shettleworth, S. J., Huttenlocher, J., & Rieser, J. J. (2007). Bayesian integration of spatial information. *Psychological Bulletin, 133*(4), 625-637.

Clayton, N. S., & Krebs, J. R. (1994). Memory for Spatial and Object-Specific Cues in Food-Storing and Nonstoring Birds. *Journal of Comparative Physiology a-Sensory Neural and Behavioral Physiology, 174*(3), 371-379.

Fevre, E. M., Bronsvoort, B., Hamilton, K. A., & Cleaveland, S. (2006). Animal movements and the spread of infectious diseases. *Trends in Microbiology, 14*(3), 125-131.

Gibbs, S., Lea, S., & Jacobs, L. F. (2007). Flexible use of spatial cues in the southern flying squirrel (*Glaucomys volans*). *Animal Cognition, 10*(2), 203-209.

Healy, S. D., & Hurly, T. A. (1998). Rufous hummingbirds' (*Selasphorus rufus*) memory for flowers: Patterns or actual spatial locations? *Journal of Experimental Psychology-Animal Behavior Processes, 24*(4), 396-404.

Herz, R. S., Zanette, L., & Sherry, D. F. (1994). Spatial Cues for Cache Retrieval by Black-Capped Chickadees. *Animal Behaviour, 48*(2), 343-351.

Jacobs, L. F., & Shiflett, M. W. (1999). Spatial orientation on a vertical maze in free-ranging fox squirrels (*Sciurus niger*). *Journal of Comparative Psychology, 113*(2), 116-127.

LaDage, L. D., Roth, T. C., Fox, R. A., & Pravosudov, V. V. (2009). Flexible cue use in food-caching birds. *Animal Cognition, 12*(3), 419-426.

Legge, E. L. G., Spetch, M. L., & Batty, E. R. (2009). Pigeons' (*Columba livia*) hierarchical organization of local and global cues in touch screen tasks. *Behavioural Processes, 80*(2), 128-139.

Shettleworth, S. J. (2003). Memory and hippocampal specialization in food-storing birds: Challenges for research on comparative cognition. *Brain Behavior and Evolution, 62*(2), 108-116.

Simons, A. M. (2004). Many wrongs: the advantage of group navigation. *Trends in Ecology & Evolution, 19*(9), 453-455.

Thiele, J., & Winter, Y. (2005). Hierarchical strategy for relocating food targets in flower bats: spatial memory versus cue-directed search. *Animal Behaviour, 69*, 315-327.

Thrun, S., (2005). *Probabilistic robotics / Sebastian Thrun, Wolfram Burgard, Dieter Fox*. Cambridge, Mass. :: MIT Press.

Vlasak, A. N. (2006a). Global and local spatial landmarks: their role during foraging by Columbian ground squirrels (*Spermophilus columbianus*). *Animal Cognition, 9*(1), 71-80.

Vlasak, A. N. (2006b). The relative importance of global and local landmarks in navigation by Columbian ground squirrels (*Spermophilus columbianus*). *Journal of Comparative Psychology, 120*(2), 131-138.

Waisman, A. S., & Jacobs, L. F. (2008). Flexibility of cue use in the fox squirrel (*Sciurus niger*). *Animal Cognition, 11*(4), 625-636.