A Socially Aware Bayesian Model for Competitive Foraging

Sheeraz Ahmad (sahmad@cs.ucsd.edu)
Department of Computer Science and Engineering
University of California, San Diego

Angela J. Yu (ajyu@ucsd.edu)
Department of Cognitive Science
University of California, San Diego

Abstract

Effectively solving the problem of when and where to forage is critical for survival in many animal species. The task is further complicated when there are other agents, potentially competing for the same limited resources. Previous models of foraging consider agents either in isolation or in groups but without competition. Here, we present a novel Bayesian model for competitive foraging, Socially Aware Bayesian Agent (SABA), that takes into explicit account the presence of other agents for both learning and decision-making. For comparison, we also implement a simple Naive Agent model that completely ignores the presence of other agents. We find that although all models perform well in a stationary environment, converging quickly to the optimal population-level solution, only SABA with the stochastic foraging policy can readily adapt when the environment is non-stationary. These results represent a first step toward cognitively sophisticated representations for learning and decision-making in competitive foraging.

Keywords: competitive foraging; Bayesian learning; decision-making; public information; Ideal Free Distribution (IFD).

Introduction

Foraging is the general problem of deciding how to choose among different reward options (e.g. food patches) adaptively, based on sequentially observed, noisy data, and thus in the presence of uncertainty associated with reward values of the different options. In competitive foraging, there is both benefit, in terms of information, and cost, in terms of competition, due to the presence of other foraging agents. Daily life is rife with such competitive foraging examples, such as picking a restaurant, asking someone out on a date, bidding on a house, picking a vacation location, etc. In the experimental literature, this has most often been studied in the context of animals foraging in an environment containing patches of food (see (Pyke et al., 1977; Perry & Pianka, 1997) for review). The optimal solution for the population is to have the agents distributed proportional to the underlying reward rates at the different food patches, as that ensures maximal efficiency of food intake in the population. This distribution is also known as the Ideal Free Distribution (IFD), and has been observed in humans (Sokolowski et al., 1999; Kraft & Baum, 2001; Madden et al., 2002; Goldstone & Ashpole, 2004; Goldstone et al., 2005) and other species such as ducks (Harper, 1982), guppies (Dill, 1987), sticklebacks (Milinski, 1979), Bumblebees (Dreisig, 1995), etc.

Interestingly, IFD is also the individually optimal Nash equilibrium solution (Nash et al., 1950), in which no agent can do better by moving to another patch. Competitive foraging, therefore, presents an ecologically convenient case study for how individual-based, decentralized learning and decision-making, can lead to coherent and desirable behavioral patterns at the group level. Surprisingly, despite the rich body of empirical literature on the subject, the cognitive and neural mechanisms leading to IFD are not well understood. Most of the models (see (Pyke, 1984) for a detailed review) are descriptive in nature, and make unrealistic assumptions about foraging agents' knowledge about the reward environment and their behavioral constraints. In particular, agents are assumed to have precise and unbiased knowledge about the reward distribution (the *ideal* assumption), and that they can freely travel from one patch to another without any limitation or cost (the *free* assumption). In more natural scenarios, humans and animals often have to forage without (initial) knowledge about the relative reward value of the different options, and there may be costs associated with switching between options (e.g. time, energy, opportunity cost, money). A more complete understanding of competitive foraging behavior needs to incorporate these factors.

Recently, there has been a shift toward models that relax these strong assumptions, such as EPICURE (Roberts & Goldstone, 2006), Polyworld (Griffith & Yaeger, 2011), Avidian agents (Walker, 2011), Omega agents (Seth, 2011), and Dynamic system (Krivan et al., 2008). While these models have demonstrated that IFD can be achieved even when the ideal assumption is relaxed, they still make the free assumption. More importantly, these models are overly simplistic in their assumptions about each agent's representation of the environmental reward distribution and other competing agents. Most of these assume that learning only proceeds based on *private information*, i.e. an agent's own history of choices and reward outcomes (Griffith & Yaeger, 2011; Walker, 2011; Seth, 2011). EPICURE (Roberts & Goldstone, 2006) does consider the effect of competition but it just fits an ad-hoc, descriptive parameter corresponding to agent density, so that the competition has a unidirectional (negative for their data) effect on agent choices, failing to account for the fact that depending on agent's knowledge, at times high agent density might make an option more lucrative. A large body of empirical ecology research shows that animals readily use *public information*, i.e. other agents' choices and possibly also their outcomes, in foraging (e.g. (Krebs, 1973; Boyd, 1988; Templeton & Giraldeau, 1995)), as well as in other bio-

logically important behavior, such as mate selection, predator avoidance, path selection, etc. (see (Galef Jr & Laland, 2005) for review). Thus a psychologically plausible model of foraging should take public information into account.

A separate line of empirical and modeling studies have considered group foraging scenarios in which competition is not a factor (Green, 1980; Lima, 1985; Valone, 1989; Livoreil & Giraldeau, 1997; Olsson & Holmgren, 1998; Sernland et al., 2003), either because rewards are so abundant that there is no need for competition, or because the species under study is by nature highly cooperative. So although these studies take public information into account for learning, they do not, and have no need to, take into account other agents' anticipated behavior in making foraging decisions. There is also the caveat that these models assume that other agents' choice *outcomes* are publicly observable as well, whereas in many natural situations, perhaps only others' choices are observable and not the outcomes. Nevertheless, others' choices are still useful for learning, as they represent crowd-sourced information about the relative reward values for the different options.

Thus, there are two broad families of existing models for foraging, one considering competition but not the use of public information, and the other using public information but sidestepping competition. In this paper, we consider both the cost (competition) and benefit (public information) of foraging among competing agents. We propose a family of models, which we call SABA (Socially Aware Bayesian Agents), that takes into account: (1) knowledge about the fact that more competition at a given patch leads to lower reward, and (2) public information in the form of other agents' choices (and not their outcomes). We also consider two self-consistent decision policies, one that makes deterministic choices (based on the highest "patch profitability ratio") and another that makes stochastic choices (proportional to the inferred "patch profitability ratios" of the different options). We also present a Naive Agent model, a Bayesian analogue of the existing competitive foraging models (like EPICURE, Polyworld, Avidian and Omega agents), which is oblivious of other agents both for learning and decision-making.

The rest of the paper is organized as follows. We describe SABA and the Naive Agent model in the next section, explaining the inference and decision strategies for both. We then present simulation results on a simple two patch foraging problem, both for a scenario in which reward distribution is stationary over time, and another in which it is non-stationary. We conclude with a discussion of implications and limitations of our work, as well as pointers for future research.

Socially Aware Bayesian Agents (SABA)

We consider a scenario in which each agent decides which patch to forage next based on the history of individual reward and public information (crowd-density at different patches). There are n patches, and patch k has a (hidden) abundance level $\theta_k \in \{0,1,2,\ldots,(R-2),(R-1)\}$, where R is the number of discrete abundance levels. Similarly, patch k has (observed) crowd-density $\phi_k \in \{0,1,2,\ldots,(C-2),(C-1)\}$, where C is the number of discrete levels of crowd-density. There are two reasons for considering quantized levels of crowd-density instead of exact number of competitors in the model: (1) numerosity judgment noise scales with magnitude in humans and animals (Dehaene et al., 1998), (2) having arbitrarily precise numerosity information can easily overwhelm private information and lead to informational cascade (Bikhchandani et al., 1992).

Bayesian Inference

We use a Bayesian generative model to capture the agent's knowledge about the statistical nature of how reward is generated at a patch, and how crowd-density serves as an indicator of patch richness. We assume that the agents use Bayesian inference in the recognition model to update their belief about the hidden state of the world (patch richness).

At patch k, reward $x_t \in \{0,1\}$ is generated as:

$$p(x^t|\theta_k,\phi_k^t) \sim Bernoulli(x^t; 1-exp(-\theta_k/\phi_k^t)) \qquad (1)$$

This generative model captures the idea that the probability of getting a reward ($x^t = 1$) increases with increasing richness level and decreases with increasing crowd-density, but the dependence is non-linear in both cases.

Each agent also needs a model for predicting how other agents will move, but an exact model of such dynamics is inherently intractable due to the infinite recursion when an agent, a_i, tries to model other agents' beliefs about the world including a_i. Thus we need a model of crowd-dynamics that is tractable and captures the behavior approximately. We propose:

$$p(\phi_k^t|\theta_k,\phi_k^{t-1}) = \alpha \mathbf{1}_{\{\phi_k^t=\phi_k^{t-1}\}} + (1-\alpha)p'(\phi_k^t|\theta_k) \qquad (2)$$

which asserts that at each time instant, the agent distribution can remain the same with probability α, or a redistribution around θ_k can occur with probability $(1-\alpha)$. The distribution p' is peaked at the hidden reward abundance levels θ_k; in the following simulations, we assume p' to be a discretized Gaussian with peak at θ_k. Eq. 2 captures our assumption that each agent assumes that the population distribution gradually tends toward the IFD from the current configuration, with the parameter α parameterizing how fast this convergence is expected to happen. Note that the equation does not imply that agents have any additional information about the underlying abundance level θ_k, only that agents assume a certain behavior of how other agents move, specifically that agent-density is likely to be (noisily) matched to the hidden abundance level.

In the recognition model, based on the observed reward history $\mathbf{x}_t$ at the chosen patch and the observed crowd density history Φ_k at all patches, each agent updates its belief about patch richness θ_k by applying Bayes' rule. Thus at the chosen

patch:

$$p(\theta_k|\mathbf{x}^t,\Phi_k^t) \propto p(x^t|\theta_k,\phi_k^t)p(\phi_k^t|\theta_k,\phi_k^{t-1})p(\theta_k|\mathbf{x}^{t-1},\Phi_k^{t-1}) \tag{3}$$

And at all other patches:

$$p(\theta_k|\Phi_k^t) \propto p(\phi_k^t|\theta_k,\phi_k^{t-1})p(\theta_k|\Phi_k^{t-1}) \tag{4}$$

Decision Strategies

Since the goal of each agent is to maximize its individual gain, a behaviorally relevant strategy is to move to a patch which has the highest expected ratio of patch-richness to crowd density. Furthermore, relaxing the free assumption would imply that moving between patches has some cost. With these considerations, a deterministic strategy for next patch selection is:

$$k^{t+1} = \underset{k}{\text{argmax}} \left\{ \frac{\mathbb{E}(\theta_k)}{\mathbb{E}(\phi_k^{t+1})} - c_s \mathbf{1}_{\{k \neq k^t\}} \right\} \tag{5}$$

The deterministic policy bases the decision on not only the agent's belief about patch richness but also about its belief about how other agents would move (eq. 2). Since the crowd dynamics model is only approximate, a disagreement between true dynamics and assumed dynamics could potentially lead to misinformed decisions. To avoid this, an agent can also employ a stochastic policy that only depends on its belief about patch richness, so that:

$$p(k^{t+1} = k) \propto \mathbb{E}(\theta_k) \tag{6}$$

Another motivation for using a stochastic strategy comes from the animal behavior data (Harper, 1982) where the equilibrium is dynamic in nature, i.e. agents (ducks) move between patches so that the number of ducks on any patch still approximates the IFD. A stochastic strategy does lead to a dynamic equilibrium as agents choose a location randomly and hence switch often, whereas in a deterministic strategy the choice of location becomes fixed once the system is in equilibrium. As we will see in the result section, a deterministic strategy can get stuck in a suboptimal solution.

Naive Agents

Several existing models of foraging (like (Roberts & Goldstone, 2006; Griffith & Yaeger, 2011; Walker, 2011; Seth, 2011)) propose agents that maintain an estimate of patch-richness based on their individual experience and employ a simple strategy for selecting the next patch. In order to compare SABA to this class of model, we propose agents within our Bayesian framework, that are agnostic about social information and crowd's effect on the reward. Furthermore, these agents greedily choose the next patch as the one which they believe to be the best (has highest expected value of patch richness). Such agents have a wrong statistical model of the environment, attributing the reward outcome solely to (fixed) patch-richness instead of attributing it to the (fluctuating) presence of competitors:

$$\hat{p}(x^t|\theta_k) \sim Bernoulli(x^t;\theta_k) \tag{7}$$

Furthermore, since individual reward is the only source of information, these agents only update belief about the attended patch where reward can be observed. After observing reward x_t, the belief is updated via Bayes rule:

$$p(\theta_k|\mathbf{x}^t) \propto \hat{p}(x^t|\theta_k)p(\theta_k|\mathbf{x}^{t-1}) \tag{8}$$

And the control strategy is:

$$k^{t+1} = \underset{k}{\text{argmax}} \{\mathbb{E}(\theta_k)\} \tag{9}$$

Results

For ease of visualization, we consider a simple foraging problem over 2 patches, with 100 agents. We address more complex scenarios in the Discussion. The (hidden) patch richness $\theta_k \in \{0,1,2,\ldots,9,10\}$, and the (observed) crowd-density $\phi_k \in \{0,1,2,\ldots,9,10\}$. The α parameter in the approximate crowd-density distribution, eq. 2, is 0.6. Fig. 1 shows the behavior of SABA (Deterministic and Stochastic strategies), as well as Naive Agents for a stationary environment [7 3], i.e. the richness level of patch 1 (blue) is 7, and that of patch 2 (green) is 3. The dotted lines show the IFD, which is matched to patch richness, and is therefore also [7 3]. The solid lines show the simulated distributions when all the agents follow a given strategy. We note that agents following any of the three strategies achieve IFD. However, only the stochastic policy achieves a dynamic equilibrium whereas for the other two, the number of agents on a given patch remains constant once equilibrium is attained. Note that although it seems like that both Deterministic and Naive strategies under- or over-match, this is not the case, since the discretization of richness and crowd density enforces that the observed crowd-density is always an integer. For example, for the deterministic strategy, although there are 68 agents at blue patch, as far as the agents are considered it is crowd-density level 7, which is matched to the patch richness. Thus, just by looking at the aggregate agent behavior for this environment, it may seem that all three strategies are equally good, leading to the optimal solution.

In order to further differentiate the models and investigate the scenarios where considering social information may potentially be useful, we look at the beliefs that agents on different patches infer. In other words, we wish to understand whether the agents, following these different strategies, also correctly figure out the underlying richness of different patches, in addition to just the observed IFD behavior. This is important, because an agent operating with incorrect belief (e.g. believing that all other patches are very low on resources) may never adapt to changes in the environment (e.g. always stays at the settled patch even when its own patch has deteriorated). Fig. 2 shows the average beliefs of agents in a

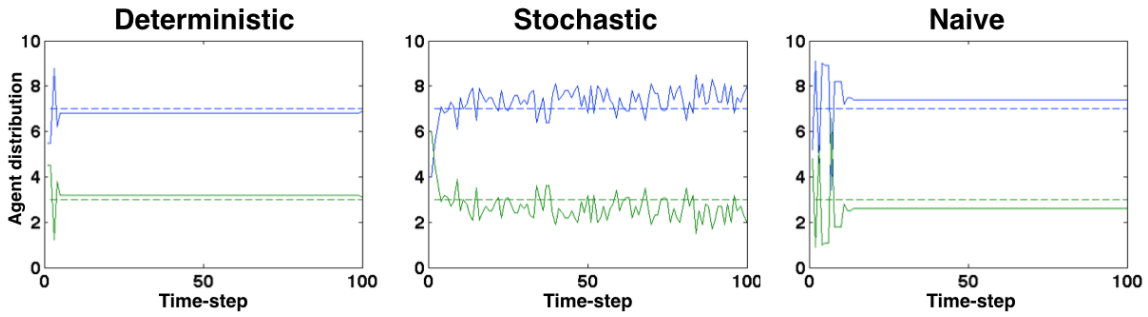

Figure 1: Stationary Environment: [7 3]. Dotted lines indicate the patch richness level (hidden). Solid lines indicate the normalized agent density on each patch, blue: patch 1 and green: patch 2.

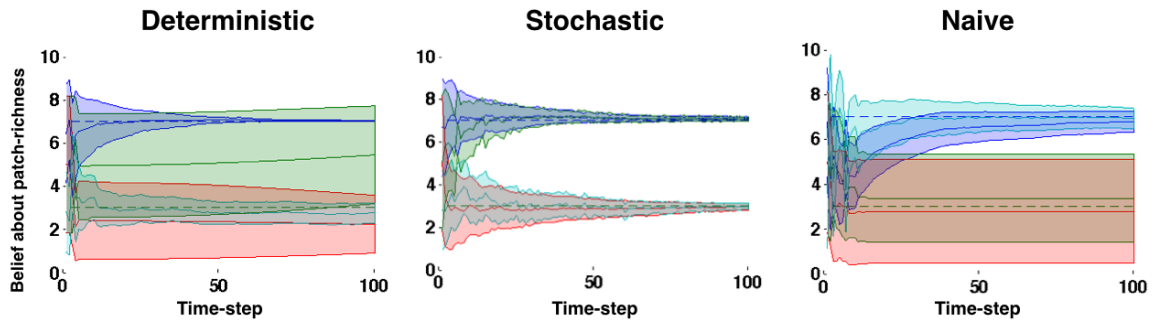

Figure 2: Inferred beliefs of agents at different patches. Dotted lines show the true richness level of the two patches. Blue: Belief about patch 1, of agent at patch 1. Green: Belief about patch 1, of agent at patch 2. Red: Belief about patch 2, of agent at patch 1. Cyan: Belief about patch 2, of agent at patch 2. The boundaries indicate the standard deviation of beliefs.

given patch about the two patches. We note that the agents following the deterministic strategy infer the richness that is qualitatively accurate, i.e. both agents at patch 1 and patch 2 believe that patch 1 is richer, and similarly both believe that patch 2 is poorer, although the beliefs are not always strong. Agents following the stochastic strategy infer the quantitatively accurate richness, i.e. agents at both patches become confident that patch 1 has a richness level 7, and patch 2 has a richness level 3. On the other hand, agents following the Naive strategy fail to infer even the qualitative trends, specifically agents at patch 2 settle on a belief that patch 2 is better than patch 1, opposite of the real richness levels. These observations motivated us to investigate whether the different strategies can perform well when the environment changes.

Fig. 3 shows the behavior of different algorithms for a non-stationary environment $[6\,4] \rightarrow [7\,3]$, the change occurring after 150 time-steps. For this case, the stochastic strategy is the only one that adapts, and achieves IFD in the changed environment, owing to its ability to infer the hidden richness correctly. The deterministic strategy suffers because it puts too much trust on other agents, both during inference and decision making, which leads to self-information being inundated by less accurate public information. The naive strategy suffers because it does not incorporate public information at all. The results shown suggest that a strategy like the stochastic one that uses public information, without letting it overwhelm self information, works the best.

Discussion

We introduced a Bayesian competitive foraging model, SABA, as well as a Naive Agent model. SABA takes into account both the self and public information, which is in accordance with the findings that animals use both of these sources for learning and decision-making while foraging. Introduction of stochasticity leads to an adaptive policy with dynamic equilibrium, as observed for foraging animals. Our model is principled in that it only makes assumptions about the statistical nature of the environment, in contrast to the ad hoc parameter fitting needed by existing approaches. In a comprehensive review on the use of public information (Valone & Templeton, 2002), the authors pointed out two shortcomings of the state-of-the-art research in the field stemming from these assumptions: (1) each agent can not only observe all the other foragers but also whether they received a reward, (2) public information serves as an additional sample information and is given equal weighting as the self-information. We relax assumption 1 by allowing each agent to only observe the crowd-density at other patches instead of the quantity of food available or each forager's individual choice outcome. We relax assumption 2 by modeling public information as having a constant level of informativeness (discretization) regardless of the absolute number of competing agents. These also alleviate the problem of informational cascade since self information contributes more to the inference. Although, we show results for 2 patches, the model extends naturally to

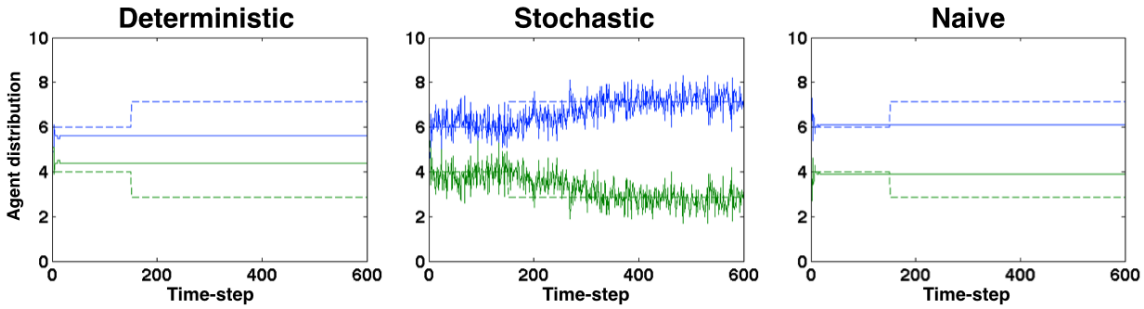

Figure 3: Non-stationary Environment: $[6\ 4] \rightarrow [7\ 3]$. Dotted lines indicate the patch richness level (hidden) which changes at time-step 150. Solid lines indicate the normalized agent density on each patch, blue: patch 1 and green: patch 2.

more patches, and the increase in complexity is linear owing to the inference on patches being decoupled. Simulations for 3 and 4 patches show that the model achieves IFD without high computational demands.

We believe our work is a step toward more sophisticated models of decision making in social scenarios. However, it also has limitations that leave scope for future inquiries, some of which we discuss here. The model of crowd-dynamics, eq. 2, is only approximate and needs to be revisited for more complex problems. This also means that we need ways to manage how much trust to place in public information, since not only the model of crowd-dynamics is approximate, an agent cannot trust other agents to be behaving optimally. Towards this trust management goal, we tried an inference method, called Perturbed Bayes (Jadbabaie et al., 2012), which provides a way to combine high quality self-information with the low-quality public information. However, the technique is ad-hoc and there is no principled way of deciding how to combine the differential-quality information sources, making it only moderately useful for the problem. There are few other investigations in this domain, but it is largely an open problem. Another approach we tried, is to start doubtful about public information, and put incrementally more trust in the crowd, since the crowd is expected to have learned more about the environment as time passes. We implement this idea by decreasing the variance of the distribution p' in eq. 2 with time, but it is not clear how to set the rate of this decrease in a principled manner.

Another way to improve the performance in a non-stationary environment, is to use a Dynamic Belief Model (DBM) (Yu & Cohen, 2009), which assumes non-stationarity of the environment in the generative model, making agent more *aware* during inference that resource levels can change abruptly. Using a model such as DBM would improve not only the Naive Agents but also SABA, especially when combined with a deterministic decision policy, although the parameter specifying the rate of change would have to be higher in Naive Agent compared to SABA, as it would need to account for reward variability not only due to true reward non-stationarity but also net profitability due to the presence/absence of other agents.

Another direction for extension is to consider a scenario in which the reward options are located (embedded) over a terrain, and the hidden reward abundance level is an unknown continuous function over that terrain – this would also allow the possibility of including patch-switching cost that depends on the distance traveled between two options. However, this scenario poses serious tractability issues since the number of patches is not finite (or even countable) any more. One way to get around this is to use a parametric reward surface or a Bayesian non-parametric formulation such as a Gaussian Process (Williams & Rasmussen, 1996). Then each sample can be used to update the agent's beliefs about the entire terrain.

Yet another direction for future enquiry is to consider a collection of heterogenous agents who may differ in capability, costs, and goals. Indeed, heterogeneity has been observed in animals, in which more dominant members may "despotically" monopolize some food patches and push the population distribution away from IFD (Parker & Sutherland, 1986).

To conclude, the problem of competitive foraging in its entirety is highly complex. This model represents a first step toward combining self and public information rationally to infer the hidden richness of different patches and using this inferred knowledge, along with beliefs about other agents' intended actions, to inform future decisions. And this individually optimal, decentralized foraging strategy naturally leads to IFD, the socially optimal solution for competitive foraging.

Acknowledgments

This project was supported by Qualcomm's Fellow Mentor Advisor Fellowship to S. Ahmad and A. J. Yu; we thank Qualcomm mentors Yinyin Liu and Blythe Towal for helpful feedback. This material is based upon work supported in part by the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF1110391.

References

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 992–1026.

Boyd, R. (1988). *Culture and the evolutionary process*. University of Chicago Press.

Dehaene, S., Dehaene-Lambertz, G., & Cohen, L. (1998). Abstract representations of numbers in the animal and human brain. *Trends in neurosciences*, *21*(8), 355–361.

Dill, L. M. (1987). Animal decision making and its ecological consequences: the future of aquatic ecology and behaviour. *Canadian Journal of Zoology*, *65*(4), 803–811.

Dreisig, H. (1995). Ideal free distributions of nectar foraging bumblebees. *Oikos*, 161–172.

Galef Jr, B. G., & Laland, K. N. (2005). Social learning in animals: empirical studies and theoretical models. *Bioscience*, *55*(6), 489–499.

Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, *11*(3), 508–514.

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*(1), 81–87.

Green, R. F. (1980). Bayesian birds: a simple example of oaten's stochastic model of optimal foraging. *Theoretical Population Biology*, *18*(2), 244–256.

Griffith, V., & Yaeger, L. S. (2011). Ideal free distribution in agents with evolved neural architectures. *arXiv preprint arXiv:1112.3574*.

Harper, D. (1982). Competitive foraging in mallards:ideal freeducks. *Animal Behaviour*, *30*(2), 575–584.

Jadbabaie, A., Molavi, P., Sandroni, A., & Tahbaz-Salehi, A. (2012). Non-bayesian social learning. *Games and Economic Behavior*.

Kraft, J. R., & Baum, W. M. (2001). Group choice: The ideal free distribution of human social behavior. *Journal of the Experimental Analysis of Behavior*, *76*(1), 21–42.

Krebs, J. R. (1973). Social learning and the significance of mixed-species flocks of chickadees (parus spp.). *Canadian Journal of Zoology*, *51*(12), 1275–1288.

Krivan, V., Cressman, R., Schneider, C., et al. (2008). The ideal free distribution: a review and synthesis of the game-theoretic perspective. *Theoretical population biology*, *73*(3), 403–425.

Lima, S. L. (1985). Sampling behavior of starlings foraging in simple patchy environments. *Behavioral Ecology and Sociobiology*, *16*(2), 135–142.

Livoreil, B., & Giraldeau, L.-A. (1997). Patch departure decisions by spice finches foraging singly or in groups. *Animal behaviour*, *54*(4), 967–977.

Madden, G. J., Peden, B. F., & Yamaguchi, T. (2002). Human group choice: Discrete-trial and free-operant tests of the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, *78*(1), 1–15.

Milinski, M. (1979). An evolutionarily stable feeding strategy in sticklebacks1. *Zeitschrift für Tierpsychologie*, *51*(1), 36–40.

Nash, J. F., et al. (1950). Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, *36*(1), 48–49.

Olsson, O., & Holmgren, N. M. (1998). The survival-rate-maximizing policy for bayesian foragers: wait for good news. *Behavioral Ecology*, *9*(4), 345–353.

Parker, G., & Sutherland, W. (1986). Ideal free distributions when individuals differ in competitive ability: phenotype-limited ideal free models. *Animal Behaviour*, *34*(4), 1222–1242.

Perry, G., & Pianka, E. R. (1997). Animal foraging: past, present and future. *Trends in Ecology & Evolution*, *12*(9), 360–364.

Pyke, G. H. (1984). Optimal foraging theory: a critical review. *Annual review of ecology and systematics*, *15*, 523–575.

Pyke, G. H., Pulliam, H. R., & Charnov, E. L. (1977). Optimal foraging: a selective review of theory and tests. *Quarterly Review of Biology*, 137–154.

Roberts, M. E., & Goldstone, R. L. (2006). Epicure: Spatial and knowledge limitations in group foraging. *Adaptive Behavior*, *14*(4), 291–314.

Sernland, E., Olsson, O., & Holmgren, N. M. (2003). Does information sharing promote group foraging? *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *270*(1520), 1137–1141.

Seth, A. K. (2011). *Optimal agent-based models of action selection.* Modeling natural action selection, Cambridge University Press.

Sokolowski, M., Tonneau, F., & Baqué, E. F. i. (1999). The ideal free distribution in humans: An experimental test. *Psychonomic bulletin & review*, *6*(1), 157–161.

Templeton, J. J., & Giraldeau, L.-A. (1995). Patch assessment in foraging flocks of european starlings: evidence for the use of public information. *Behavioral Ecology*, *6*(1), 65–72.

Valone, T. J. (1989). Group foraging, public information, and patch estimation. *Oikos*, 357–363.

Valone, T. J., & Templeton, J. J. (2002). Public information for the assessment of quality: a widespread social phenomenon. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1427), 1549–1557.

Walker, J. C. (2011). The evolution of optimal foraging strategies in populations of digital organisms. In *Proceedings of the 13th annual conference on genetic and evolutionary computation* (pp. 203–210).

Williams, C. K. I., & Rasmussen, C. E. (1996). Gaussian processes for regression. In M. M. D. S. Touretzky & M. E. Hasselmo (Eds.), *Advances in neural information processing systems 8* (p. 514-20). Cambridge, MA: MIT Press.

Yu, A. J., & Cohen, J. D. (2009). Sequential effects: Superstition or rational behavior? *Advances in Neural Information Processing Systems*, *21*, 1873-80.