

What Is Foraging?¹

David L Barack

University of Pennsylvania

Forthcoming in Biology and Philosophy. Please cite the published version.

Abstract

Foraging is a central competence of all mobile organisms. Models and concepts from foraging theory have been applied widely throughout biology to the search for many kinds of external resources, including food, sexual encounters, minerals, water, and the like. In cognitive science and neuroscience, the tools of foraging theory are increasingly applied to a wide range of other types of search, including for abstract resources like information or for internal resources like memories, concepts, and strategies for problem solving. Despite its importance in ecology and increasing relevance for the study of cognition, the concept of foraging is rarely analyzed. Here, I aim to rectify this situation. I outline three desiderata: first, an analysis should differentiate foraging from search and decision making more generally; second, an analysis should unify different types of foraging; and third, an analysis should help ground predictions. I present an analysis of foraging as the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts. Not all search is serial and not all decision making is exclusive, differentiating foraging from search and decision making generally. With the aid of Markov decision processes and directed cyclical models, I show how the analysis implies a cyclical graph. This cyclical graph is embedded in the description of many types of foraging, unifying the different instances. Finally, I argue that the cyclical graph is also embedded in representations of novel task contexts that have not previously been viewed as foraging. I illustrate this novel application of the concept of foraging by arguing that reasoning is a type of foraging.

¹ Acknowledgements: I would like to thank Naomi Rosenkranz, Thomas Hills, Eric Leonardis, and the Future of Foraging Seminar Series for helpful commentary and feedback.

I. Introduction

Foraging, long a focus of study in ecology, is increasingly a hot topic in neuroscience and psychology. Foraging refers to the search for resources in the environment like food or water, and foraging theory (such as Stephens and Krebs 1986) provides a set of models to understand how animals should make such decisions. But models from foraging theory have been extended to cognitive domains (Todd and Hills 2020), including memory (Hills et al. 2012), concepts (Hills et al. 2015), task scheduling (Payne et al. 2007), and more. In addition, foraging models have been applied to the search for abstract resources like information (Pirolli and Card 1999; Pirolli 2007).

Despite this increasing breadth of application, the concept of foraging is rarely analyzed. What exactly is foraging? Is foraging to be identified with search? When animals search for and consume resources, they make a host of decisions, such as where to search; how to search; which food items to search for, which to consume, and which to forego; between different items encountered during search; and many others. Does foraging differ from decisions generally that animals make? And if so, how? What makes foraging different from search on the one hand and decision making on the other?

Clarifying the nature of foraging will help validate and explain this breadth of application. First, foraging is relevant to a range of philosophical work. Sober uses foraging as an illustration of the role and importance of optimality models in evolutionary explanations (Orzack and Sober 1994; Elgin and Sober 2002). Abrams appeals to foraging theory to understand the role of uncertainty in biology (Abrams 2023). But how do we know that optimality models in foraging theory are applied to the right domains? Or, how do we know that a given example of putative foraging behavior is in fact foraging? An analysis of foraging will help justify the use of these examples for diverse philosophical purposes.

Second, foraging theory is increasingly appearing in discussions of the psychological and neural mechanisms of decision making (Pearson et al. 2014; Mobbs et al. 2018). Many choice tasks are said to be foraging tasks, such as temporal discounting (or intertemporal choice) tasks (Stephens 2008). Foraging processes have recently been connected to evidence integration (Davidson and El Hady 2019), central to a wide range of decision phenomena (Gold and Shadlen 2007). Finally, foraging phenomena are being identified in numerous cognitive processes, from searches through memory (Hills et al. 2012) to visuospatial reasoning (Barack et al. 2023) to

scheduling tasks (Payne et al. 2007). The relevance of this diverse body of evidence to foraging remains to be supported, however. My analysis of foraging is, in part, intended to provide a framework for determining whether a given task is a foraging one.

Third, foraging is often cited as crucial to understanding human evolution. Selective pressures from foraging have been argued to give rise to primate intelligence (Genovesio et al. 2014; Passingham 2021). Foraging is also thought to shape social and political ways of life in humans (Kelly 2013). These claims require foraging to be distinct in some way from search and decision making generally; otherwise, they amount to the banal claim that the environment has shaped human intelligence or the social and political lives of humans. How, then, does foraging differ from search and decision making? I will provide an answer to this question herein.

Fourth, an analysis of foraging can also help resolve scientific debate within the foraging community. As illustration, consider a recent debate about foraging through semantic memory, the ability to recall facts about the world such as words, concepts, or numbers (Yee et al. 2014). Hills and colleagues (Hills et al. 2012) reported foraging behavior when participants were instructed to name as many items of a given category, such as animals, as they could in three minutes, concluding that participants forage through memory. Abbott and colleagues (Abbott et al. 2015) undermine this conclusion, showing that behavior of a model that utilizes a random search on a network with a certain structure can match the effects reported by Hills and colleagues. Now, I maintain that not every task is a foraging task; so, one way to resolve the debate is to determine whether such a task is a foraging task, making behavior on the task foraging behavior. My analysis also makes sense of models of decision making as foraging models. Algorithms designed for foraging tasks, such as Hills and colleagues construct, can produce foraging behavior, but so can algorithms designed for any kind of search, such as Abbott and colleagues use. Below, I argue that the task is a foraging task that can be solved with non-foraging algorithms. Furthermore, the equivocal evidence cited by Hills et al. does not undermine the status of the task as a foraging one.

In this essay, I will provide an analysis of foraging. I argue that foraging is the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts. I then defend this argument against some objections. The analysis can be formalized using the framework of Markov decision processes and directed cyclic graphs. After illustrating how such a framework can capture a range of food foraging contexts, I apply it to foraging in memory. I next suggest a novel application of foraging to some reasoning contexts. Foraging is a central capacity of all mobile

organisms, but foraging contexts appear in a wide range of tasks facing many different types of organisms, and the description of foraging mechanisms promises insights into behavior beyond the banal search for alimentary resources.

II. Search, Decision Making, and Foraging

The search for resources is fundamental to biological life. However, not all searches are alike. They differ in the type (material or abstract) and location (internal or external) of the resource, and in the nature (serial or parallel) and specificity (specific or general) of the search. In addition, decision contexts during search vary with respect to the number of options (one, two, or more), exclusivity (options are considered one-at-a-time or many-at-a-time), and persistence of options chosen between. In this section, I present a brief list of desiderata for an analysis of foraging and analyze foraging as a special type of search and decision making.

i. Desiderata for an Analysis of Foraging

What are the constraints that any satisfactory analysis of foraging should fill? I offer three constraints. Note, though, that these desiderata may be met to a greater or lesser degree. They serve as a way to assess and as a guide to constructing analyses of foraging.

First, an analysis of foraging will not simply identify foraging with search or with decision making. Briefly, I understand search to be a type of behavior where some agent selects actions to identify some target under some ignorance about it, and decision making to be a type of behavior where some agent selects an option (which could be an action, item, or some other type) from some set in the absence of complete coercion. A satisfactory analysis of foraging will describe how foraging differs from search or decision making more generally (differentiation desideratum). Why should an analysis of foraging distinguish foraging from search on the one hand and decision making on the other? Separating foraging from search and decision making helps us understand the relevance of different scientific findings to explanations of choice behavior. What should we make of Abbott and colleagues' findings that a random search on a structured network can resemble foraging behavior? Or consider appeals to foraging to understand optimality explanations in biology. If foraging is to be identified with search or decision making in general, then why are

foraging explanations good illustrations of optimality explanations? Why not appeal to any search or choice behavior? Finally, what is special about appeals to foraging contexts to help explain intelligence? Why not appeal to any search or decision? A reason that random search on structured networks doesn't undermine foraging findings, that foraging explanations are good illustrations of optimality explanations, and that foraging may have special explanatory power in evolutionary explanations of intelligence is that foraging is distinct from search and decision making in general. The differentiation desideratum is intended to capture this special role for foraging.

Second, an analysis of foraging should be consistent with a wide range of the uses of the concept in the literature. There are many different types of foraging, including patch foraging (Charnov 1976), traplining (Berger-Tal and Bar-David 2015), and central-place foraging (Houston and McNamara 1985). If they are all foraging, then an analysis of foraging will not only accommodate these different types but also explain why they are all manifestations of the same sort of phenomenon. A satisfactory analysis of foraging, then, will unify these different types (unification desideratum). An account that flies in the face of the unity of foraging theory would fail to be consistent with scientific practice in ecology. The unification desideratum is intended to respect scientific practice and help explain why a wide range of tasks and behaviors are all considered foraging.

Finally, the third constraint for a satisfactory analysis of foraging is that it should help ground predictions that can be made by scientists. It may, for example, help ground predictions that can be tested against acknowledged foraging behavior. Alternatively, it may help ground predictions that foraging behavior will be observed in contexts that are not acknowledged as foraging ones. Regardless of the sort, a satisfactory analysis should help ground the predictions that scientists make about foraging and behavior more generally (empirical desideratum). What motivates the empirical desideratum? One aim for science is to be productive in the sense of extending theories and models to phenomena previously outside the scope of some account. Another aim is to for theories and models in science to be clear enough to be useful—that is, to provide enough detail as to apply to their target phenomena. The empirical desideratum is meant to reflect these general superempirical values in science. To capture these different roles for concepts, I state the desideratum more generally as helping ground predictions.

ii. An Analysis of Foraging

Here, I outline an analysis of foraging that will satisfy the three desiderata. I first consider the nature of search. I distinguish different types of epistemic status that agents can have regarding the target of search and the way the search is performed. I then consider the nature of decision making. I distinguish encounter sets from option sets and discuss decision making given option sets that contain only one or more than one option. I also distinguish exclusive decisions, where agents' choices exclude other choices at the same time, from non-exclusive decisions where making one choice does not exclude making a different choice at the same time. Finally, I distinguish persistent decision contexts from non-persistent ones. This discussion will result in a proposed analysis of foraging.

I will apply my analysis of foraging to a set of central cases in the next section. That each such case is an instance of the concept is explained by the analysis. However, my analysis is not essentialist—it is not intended as a list of necessary and sufficient conditions or some other sense of 'essential'. There may be some central cases of the concept that do not fall under the analysis at all, and some cases that are not taken as central—or even as peripheral—may be classified as an instance of the concept in virtue of falling under the analysis. An analysis of foraging need not cover every use of the concept of foraging as some uses might be specific to some areas of research or fall outside the analysis for other reasons. In addition, I don't maintain that an analysis of foraging must unify every context that is in fact foraging under the same description. Rather, some appropriately broad range of central cases should be covered by the analysis; different such proposals might be judged by the breadth of their coverage. Hence, the fact that some instances of foraging are left out of some analysis does not necessarily undermine the analysis, and the unification desideratum can be fulfilled by degree, with better fulfillment the wider the range of cases that are unified by the analysis.

To begin, consider search. All search involves targets, whatever it is that the agent searches for. Some searches are done with background knowledge about the target: what is being searched for, either the identity of the target (specific knowledge) or the type of target (general knowledge); where the target is located; or the path to get to the target. Some searches are conducted without any information about the target, but agents can't search for something about which they have complete information. Hence, searches are performed under some ignorance.

Variability in the searcher's epistemic state can augment what they search for.² Searchers may know the identity, location, or path to the target. Knowledge of the identity of the target can be general or specific. In general search, the agent searches for some kind of item but not a specific one. In specific search, the agent searches for one particular item. For example, when shopping for a new television, I might engage in general search, trying out different brands until settling on a choice. In specific search (sometimes called 'template matching search', see e.g. Brunelli 2009), I instead have a specific make and model of television in mind. I search for that specific item, either a token of that make and model or, even more specifically, a particular television, when I am at the store. Besides the identity of the target, searchers may know its location. A lion on the savannah searching for the next watering hole may know its location, but if they don't, then they may decide to travel to a lookout point to gather information and reduce uncertainty about the location. The lion searches first for information, then uses that information to help them navigate to the next watering hole. Furthermore, even if the identity and location of the target are known, agents may have to search for a path to get to the target. Having moved to a new town, I may know that I need to get to the grocery store and I may know where it is. My unfamiliarity with the street pattern, however, forces me to search for a route there. The type of knowledge at hand can be representational, as occurs for humans, but it need not be. The 'knowledge' might be implicit, as-if knowledge, such as in the case of foraging by worms in search of resource patches (Calhoun et al. 2014). My use of 'knowledge' should not be read as placing any representational demands on the searcher.

There are different types of targets for which animals can search. In foraging, agents search for resources. They can search for material resources, such as food (Stephens and Krebs 1986), mates (Giraldeau and Caraco 2000), minerals, water, and the like. They can also search for abstract resources, such as information (Pirulli 2007) or reputation. Though the historical study of search regards resources external to the agent, search has also long been thought to be central to cognition (Newell 1994), and increasingly cognitive psychology and related disciplines posit that agents can forage for internal resources (Todd and Hills 2020), such as concepts (Hills et al. 2015), memories

² Here, I often talk about the forager's knowledge of what they are searching for. Knowledge may not be the best description of the epistemic status of foragers. For convenience, however, I will often speak of knowledge. A deeper evaluation of the epistemic status of foragers must await a different venue.

(Rhodes and Turvey 2007; Hills et al. 2012), or ideas. This wide range of resource types suggests any analysis should permit a correspondingly wide range for foraging.

When foraging, the search is for resources of some type and not for a particular item. Hence, foraging is search for general resources. In some cases, the foraging search will be for tokens of a type, such as in cases of so-called diet selection, where only items of one type are accepted (see, e.g., Krebs et al. 1978); others will be for some type or other; but never for a particular token. This leaves lots of flexibility for foraging. In the case of the grocery store, I may forage through different routes even though the identity and location of the target are known. A monkey may forage for the next fruit patch in ignorance of its location. A lion on the hunt may opportunistically forage for some prey or other, attempting to capture the next item it finds.

In addition to variability in the searcher's epistemic state, the way the search is performed can vary. Searches can be serial or parallel. In parallel search, the agent can simultaneously search in multiple places. In serial search, the decision to search in one place excludes simultaneous searches in another place. Most single-agent natural external searches occur in this way; however, artificial or computer agents may be able to engage in parallel external search, and internal searches in the brain may or may not be parallel as well. Foraging searches are a type of serial search.

Foraging, like all search, also involves decision making. Making decisions always requires selecting between items. Encounter sets are the set of items encountered during search. Decision contexts vary in the size of the encounter set. Some decision contexts are relatively simple, where agents encounter a single item; others are more complex, involving encounters with more than one item. In addition to encounter sets, there are also option sets, the set of options among which agents select.

Option sets may or may not match encounter sets. Agents may simultaneously encounter multiple items, such as coming across a bunch of bananas at the same time but make decisions in serial about them one at a time, with the option set containing only a single item, considering first one banana, accepting or rejecting it, then the next, and so on. Alternatively, agents may decide between options simultaneously, where the option set contains multiple items. Upon encountering a bunch of bananas, the agent's option set might contain two bananas at a time, and the agent decides between them by (say) comparing the bananas.

Whether decisions are single-alternative or multi-alternative depends on the nature of the decision process and not just the number of items encountered at a time. When there is a single

encountered item, then necessarily there is a single option. When there are few items in the encounter set, all the items may be included in the option set and their features may be evaluated. Agents can then choose between options based on the evaluation of what is available. These are non-exclusive choices because accepting one item is to simultaneously reject the others; similarly, rejecting one item is to accept others (or, if rejecting all options is an option, to reject others). Some non-exclusive choices can involve multiple simultaneous accept or reject options. Exclusive choices, in contrast, involve accepting or rejecting one item without thereby accepting or rejecting other items in the option set (if there are any). The manager of a sports team may be able to select all the players that are available if there is enough room on the roster. Many exclusive choices are single alternative, where agents make decisions about single options. Agents may still evaluate options in these decisions. However, the absence of a specific alternative confounds any comparison to some such alternative. Instead, these decisions often involve a comparison between the offered option and some generic one, such as the historic average (Charnov 1976) or some expectation (McNamara 1982; Davidson and El Hady 2019).

On my analysis of foraging, foraging decisions are exclusive, accept-or-reject decisions where the decision to accept or reject one option is exclusive to the decision regarding another. In other words, foraging is single-alternative choice. The motivation underlying the focus on single-alternative decisions is that the selective environments of many organisms hypothetically involved serial encounters with single prey items. Because encounters tend to occur one at a time, option sets tended to be singletons. The ancient forager had to make an accept-or-reject choice regarding those items. This selective environment then shaped the decision mechanisms that are present in extant species. While for almost all organisms the study of selective environments is impossible (because they are in the past), this hypothesis motivates the restriction of foraging to single-alternative decisions, those where a single item is in the option set.

Finally, items chosen between can be more or less persistent. A persistent item is one that does not disappear simply in virtue of failing to be chosen. Most natural decision contexts are persistent, as items in the environment tend to persist. Many decision-making tasks, by contrast, are not persistent contexts because items disappear if unchosen. In persistent contexts, items need not persist for eternity if unchosen; predation, competition, and natural processes will often remove items after they are unchosen. Foraging involves persistent choice contexts. In fact, tracking renewal rates and changes due to competition in items available to become options are often central

to understanding animal foraging; for example, hermit hummingbirds are sensitive to both renewal rates for resources and to competition pressures, changing the frequency of visits to resources depending on the competitive context (Gill 1988).

In sum, I conceptualize foraging as a type of decision made in the context of a type of search. There are searches that are not foraging, such as parallel search, which violates seriality, or targeted search, which violates generality. What differentiates foraging searches from non-foraging searches is their seriality and their generality. In addition, foraging involves a certain sort of accept-or-reject decision, where foragers choose between accepting an encountered item or rejecting it in search of others but cannot do both. These decisions are exclusive—in accepting or rejecting some option, they do not thereby reject or accept some other. Finally, foraging occurs in decision contexts where items persist for some periods of time. Foraging is the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts.

I will now address three objections to this account of foraging that target the nature of foraging decisions. The first objection maintains that characterizing choice in terms of the selection of items is mistaken. Rather, choice is characterized in terms of actions. When choosing whether to accept or reject an option, the option set contains actions, not items.³ But since it is actions that are selected and not items, the objection continues, the above framing in terms of items is incorrect. For example, what before was characterized as choosing between some encountered item and a generic alternative is now to be described as a choice between two actions, accepting and rejecting a choice. And so, every option set has at least those two actions.

In reply, a shift to selecting actions instead of items violates exclusivity. A shift from items to actions results in at least two options in every choice: accept or reject the item. This reframing highlights a key shortcoming of the alternative action-based frame. In accepting an option, the agent thereby rejects the rejection. But then, they are violating the exclusivity property of foraging, which states that foragers don't simultaneously accept and reject options.

Now, exclusivity itself could be rejected. But then, there is an alternative formulation of the foregoing that treats as the basic foraging decision a binary choice between two actions, accepting or rejecting some option. In short, the account can accommodate the objection by shifting to an action-centered framework with suitable restating where required.

³ Strictly speaking, representations of actions instead of representations of items, but I will skip over this nicety.

A second objection notes that, because all (or most) choice is accept-or-reject, the account fails to differentiate between types of decisions. For example, consider an option set with two items, one banana and a pair of bananas. In choosing between one banana and two, accepting either offer is to reject the other. This holds for larger option sets as well; there is no choice that does not entail accepting or rejecting options. In reply, while it is true that much choice is accept-or-reject, the point is that foraging involves accepting or rejecting single offers, not that foraging is special in focusing on accept-or-reject decisions. Non-exclusive decisions are different because choices are linked: accepting one of the options *eo ipso* is to reject the others. Foraging decisions are not like this—they are exclusive. A forager cannot accept or reject some option and, in accepting or rejecting some option, thereby reject or accept some other.

A third objection notes that the description of the foraging choice is in fact multi-alternative. In other words, the description of foraging choice as a single option made by comparing to some generic alternative is in fact between that option and the generic alternative. But this fundamentally misconstrues the nature of the choice process. The generic alternative, like the average family or mean rainfall, is a mathematical fiction, and so is not an item that can be selected. The average properties of options in the environment can play a role in the decision about whether to accept or reject an offer (see e.g. Charnov 1976; McNamara 1982; Davidson and El Hady 2019), but there is no generic alternative in the option set because there is no generic alternative that can be encountered.

III. Central Cases of Foraging

My analysis of foraging asserts that foraging is the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts. This analysis can be used to provide a description that unifies a range of central cases using the toolkit of directed acyclic graphs and Markov decision processes. The Markov decision process framework provides a method for formalizing tasks, and the graphical framework provides a method for representing them to uncover relations between tasks.

Tasks can be formalized using Markov decision processes. A Markov decision process (MDP) is a 4-tuple $\langle S, A, T, R \rangle$ where S is a set of states called a state space, A is a set of actions called an action space, T is a transition probability matrix that specifies for action A taken in state

S the probability of transition to each state in S , and R is the reward from taking action A in state S and transitioning to state S' . The states in S can be states of the environment or states of agents. Rewards need not refer to alimentary rewards and can include information or any of the other resources that can be gathered. In what follows, I will assume that states, actions, and so forth are discrete, but that assumption can also be relaxed within the framework of MDPs. I will also assume that state transitions are deterministic and occur with unity probability.

This formalism can be used to describe foraging (see also Mangel and Clark 1986). Consider the basic foraging task implied by the discussion above; for convenience, suppose the forager is an animal that encounters prey items. In that context, the forager begins by searching for prey items, encounters them, and then is faced with a decision between accepting the item or rejecting it in search of another. This can be represented as a sequence consisting of an *Explore* state in which the forager can *Search* for prey items or *End* foraging. If foragers *End*, then foraging is over. If foragers *Search*, then they encounter a prey item in state *Encounter*. They then decide and take action *Accept* or action *Reject* in that state. Taking *Accept* leads to some consummatory state *Exploit* with associated reward. After *Exploit*, the forager chooses between *End* and *Search*; if they search, they transition to an *Explore* state and choose again between *Search* and *End* (other actions, such as *Wait*, can be readily imagined; similarly, the generic *Search* could take more specific action forms, such as *Move in d direction* or the like). Taking *Reject* leads straight to the *Explore* state and the choice between *Search* and *End*.

Each of the states and actions in this description refer to types. There are different ways to *Explore*, for example. A similar point holds for actions; there are different ways to *Search*. A particular instance of foraging is representable by the MDP if the states and actions are instances of these types. As illustration, suppose some open water fisherman is trying to catch a sailfish. The *Explore* state here is *Sailing*. The two actions could be to *Return to dock* or *Trawl*, corresponding to *End* and *Search*. As they trawl, suppose one of their fishing lines *Hooks* a fish. This is the *Encounter* state. After reeling in the fish, the fisherman may decide to *Keep* the fish, that is, to *Accept* the offer. They could have also decided to *Throw back* the fish, that is, to *Reject* the offer and return to *Sailing*. If they *Keep* the fish, they have to *Store* it somehow, corresponding to *Exploit*. After storage, the fisherman may continue to *Trawl* and return to the *Sailing* state or to *Return to dock*.

Having characterized tasks in terms of MDPs, the next step is to represent these tasks with a graphical model. The advantage of the graphical model is that it provides a representation that can be used to identify novel cases of foraging by describing graphs that occur as subgraphs of graphs that describe tasks. The graphical model is a type of directed cyclic graph. Briefly, a directed graph G on a set of variables is a set of ordered pairs of those variables, where G contains an arrow from some variable in the set X to another variable Y . A path on a graph starts at some variable, follows arrows from variable to variable, and ends at another variable; a directed path is one in which the arrows all point in the same direction. A directed cyclic graph contains paths from a variable to itself.

To illustrate, consider the MDP that describes the basic foraging context. That process can be represented with the following graph, where boxes correspond to action variables and circles correspond to state variables (Figure 1):

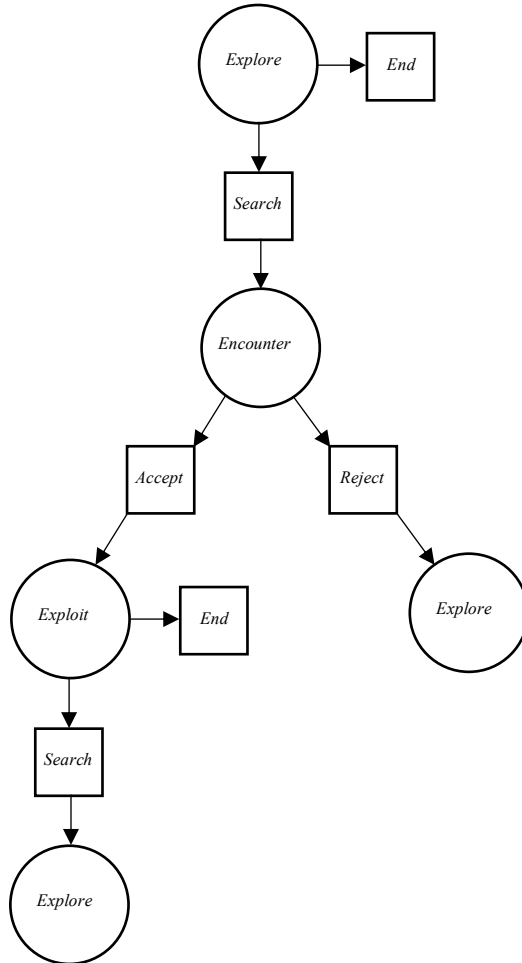


Figure 1. Directed acyclic graph of a foraging decision based on a Markov decision process model of foraging. States are represented as circles and actions as squares. Arrows generally flow in the direction of time from top to bottom (except for *End* choices, which terminate foraging). The foraging decision context starts with the forager in an *Explore* state and choosing to either *Search* for resources or *End* the foraging bout. After deciding to *Search*, the forager enters an *Encounter* state where a prey item is encountered. The forager then must make a choice between *Accept*, entering an *Exploit* state and harvesting the resources, or *Reject*, returning to a new *Explore* state at a later time than the initial state. If the forager decides to *Accept* and enter *Exploit*, then they face another decision to *End* or to *Search*, with subsequent new *Explore* state.

Using labeled edges, this graph can be simplified (Figure 2):

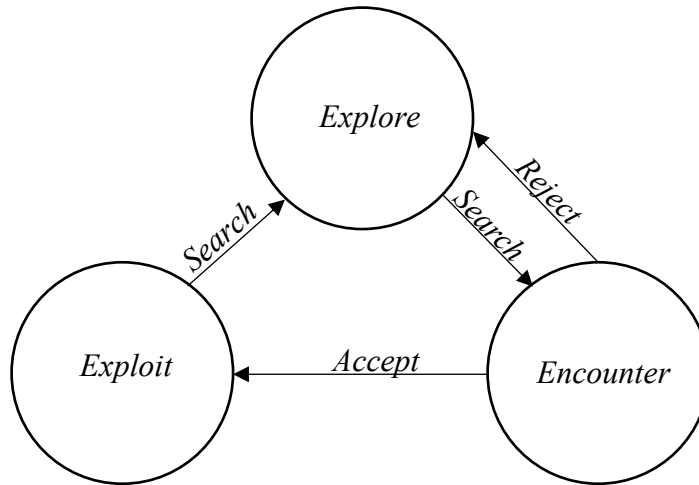


Figure 2. Condensed cyclical graph of foraging decisions.

Here, the actions are represented by labeled edges that connect starting and ending states. Note that this graph maps on to the analysis of foraging as the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts. Search is one of the actions represented in the cyclical model in either the *Exploit* or *Explore* state. The seriality of search is represented in two parts. First, the model is cyclical, meant to capture how serial searches consider sequences of locations in the search for resources. Second, there is only one each of *Explore*, *Exploit*, and *Encounter*. In contrast, a parallel process will have different subprocesses in distinct sets of states. The accept-or-reject nature of foraging decisions is represented by the two labeled edges leading from the *Encounter* state, one for the *Accept* action and one for *Reject*. The exclusivity of this choice is represented by a single labeled edge for each action type. The nature of the search target, a general resource, is left out of the representation, as is the persistence of items.⁴

This basic foraging graph can be identified in a range of different task contexts, allowing for the classification of tasks as foraging or non-foraging. A task embeds a foraging task if a subgraph of the graph is the foraging graph. By subgraph, I mean two things: first, the pattern of states and connections is present in the graph for the containing task; second, the types of states (i.e., *Explore*, etc.) and the types of actions (i.e., *Search*, etc.) in the subgraph are instances of the kinds that occur in the foraging graph. This simple thesis provides a formal method for identifying foraging tasks up to the generality and persistence of the resource being sought.

⁴ Persistence is essentially temporal. Graphs can be time-resolved as well, but shifting to time-resolved graphs is a complication beyond the scope of this discussion.

Representations of different cases of foraging, including patch foraging, traplining, and central place foraging, embed the cyclical graph (Figure 3). Take patch foraging first. Examples of patch foragers include long-line fishing trawlers, who tend to group their lines in patches to track schools of fish (O'Farrell et al. 2019), and Nahua mushroom gatherers, who follow classic foraging laws in their search for fungi (Pacheco-Cobos et al. 2019). In patch foraging, the forager is exploring, searching the environment for resources that are gathered into clumps or 'patches'. Upon encountering a patch of resources, the forager must decide between staying at the patch and harvesting prey items and leaving the patch to search for new resources. Deciding to harvest resources is a type of accept decision and leaving the patch to search for a new one is a type of reject decision. Patch foragers cannot both harvest and search. Finally, patch foragers often do not know the location of the patches or the specific prey items therein, instead searching for a general type of prey. Within the patch, the forager also faces a similar sort of foraging decision. Prey items are often encountered serially and foragers must decide between accepting the prey item or rejecting the prey item and searching for a new one. In both decisions to accept or reject a patch and to accept or reject a prey item encountered in a patch, patch foragers face single-alternative, accept-or-reject decisions in the serial search for a type of resource. These properties of foraging contexts are evident in the graphical model for patch foraging (Figure 3B).

In trapline foraging (traplining), animals search for resources at a set of locations in sequence (Freeman 1968; Berger-Tal and Bar-David 2015). Whereas in patch foraging, the locations of resources are unknown, in trapline foraging the resource locations are known. For example, baboon troops are known to trapline forage through a set range that contains fig trees, visiting them in linear succession until feeding is complete (Noser and Byrne 2010). The trapper faces a version of the basic foraging problem. Traplining animals do not know which location contains prey, though they know the location of the 'traps'; they visit the sites in sequence to find the ones with prey. For sites with prey, trappers can choose to consume the prey item, a type of accept decision, or to skip the prey item, a type of rejection. Trappers cannot both consume and skip prey items, however. After this choice, they can continue to search at the next location on the trapline or return home. Like all foraging, then, traplining involves the serial search for prey items and subsequent accept-or-reject decisions when prey items are encountered (Figure 3C).

Finally, consider central place foraging, where foragers start foraging from some central location, search for resource patches, harvest from those patches, and then leave those patches to

return home. Many species of birds are classic central place foragers (Houston and McNamara 1985), such as the great tit which brings back a single prey item to the nest (Royama 1970). Central place foraging is straightforwardly a case of patch foraging. Foragers leave their home, travel to a resource patch, make an initial accept decision to search the patch, make a series of accept-or-reject decisions about prey within the patch, and a reject decision when they leave the patch to return home. Consequently, central place foraging also satisfies the cyclical model (Figure 3D). Some central place foraging may involve search for some type of prey that does not inhabit a patch. This sort still satisfies the cyclical model because foragers explore for the prey, make an accept-or-reject decision upon encountering an instance of the prey type, and at some point cease exploring to return home.

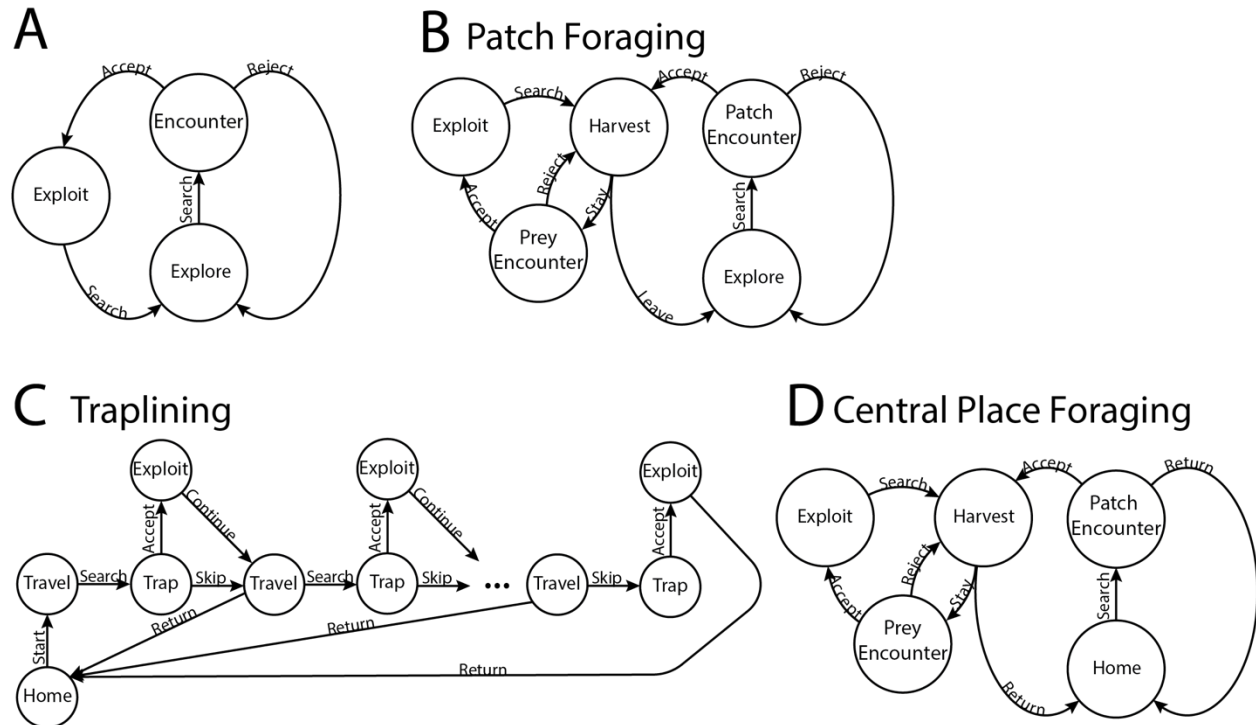


Figure 3. The basic structure of foraging decisions appears across a range of types of foraging. In these graphs, states are points and actions correspond to labeled edges, to condense presentation. **A.** The basic structure of foraging decisions can be represented as a cyclic graph. Though visually different from the display in Fig. 2, the condensed graph is the same. **B.** Patch foraging displays this same basic structure, both within patch as well as for deciding whether or not to forage in a patch. **C.** Trapline foraging displays this same basic structure, except search decisions are limited to decisions to reject the current option to travel on to the next predetermined location ('Skip'). **D.** Central place foraging displays this same basic structure, essentially matching patch foraging except search reflects a decision to return to the central place.

The use of MDPs and graphical formalisms also illustrates how certain task contexts, especially those that involve simultaneous consideration of multiple options, are not foraging contexts. Consider a classic two-alternative choice context. In this context, the option set contains two items A and B. This choice context can be represented (Figure 4):

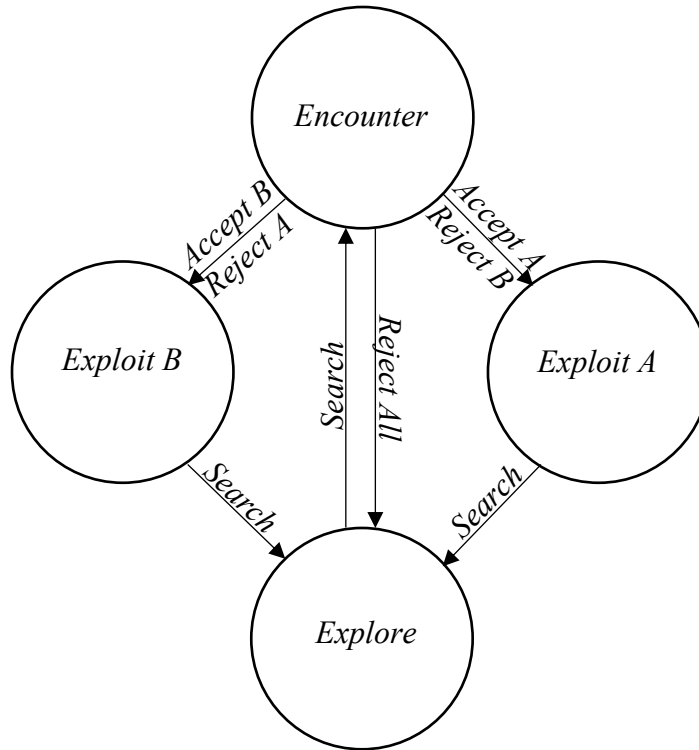


Figure 4. Cyclical graph for two-alternative simultaneous encounters.

There are two relevant differences between this graph and the previous foraging graph. First, note the two labels, *Accept* and *Reject*, in the above graph for each choice—reflecting the nature of simultaneous multi-alternative choice. This violates exclusivity, because accepting (rejecting) one option is to reject (accept) the other. Second, there are two distinct *Exploit* states depending on the agent’s selection. This implies that there are two distinct directed paths to the *Explore* state after choosing to *Accept*, in contrast to the single directed path in the foraging decision context.

The graphical model for simultaneous encounters provides grounds for an objection to the foregoing analysis. The model for simultaneous encounters embeds two instances of the foraging cyclic graph. First, upon encountering A and B, the agent may *Accept A / Reject B*. Then they enter an *Exploit A* state, followed by an *Explore* state and then another *Encounter*. This just is the foraging cyclic graph. Second, the agent may *Accept B / Reject A*, resulting in a second foraging cyclic graph. Recall that one of the desiderata of an analysis of foraging is that foraging be

differentiated from search and decision making more generally. The previous argument that foraging is distinct from decision making more generally focused on demonstrating that foraging was a subset of the class of decision contexts. But now, we have demonstrated that foraging underlies all choice. If foraging underlies all choice, however, then foraging is not distinct from decision making more generally.

I admit that this is a tantalizing thesis. Instead of defending it herein, however, I will stipulate a response. Let the outdegree of a variable be the number of paths that lead away from the variable. In a foraging task, each *Encounter* variable has an outdegree of two: either *Accept* the offer and enter *Exploit* or *Reject* the offer and return to *Explore*. In contrast, in simultaneous encounters, each *Encounter* variable has an outdegree greater than two: *Accept / Reject* choices for each item in the option set and *Reject* all items in the set. The final constraint on a task graph embedding a foraging subgraph is that the outdegree of *Encounter* is no greater than two.

IV. Other Accounts of Foraging

Stephens and Krebs (1986) offer a detailed account of foraging theory in their classic textbook. While they are primarily concerned with foraging models, I will infer an analysis of foraging from their framing of the topic (page numbers in the following refer to Stephens and Krebs 1986). In the introduction to the book, they describe three elements of foraging models (p. 5): decision assumptions that specify the variables used to make those decisions, currency assumptions that specify the evaluation of those choices, and constraint assumptions that limit the first two types of assumption. The decision assumptions include “the type of choice... the animal is assumed to make... rather than a specific choice” (p. 6). Consequently, like my analysis, foraging is a type of general search. The constraint assumptions limit the nature of the decision problem and can help characterize foraging. First, foraging assumes that search and exploitation are either/or: “the predator cannot exploit... items such as prey or patches while searching for new ones” (p. 11). Second, that items are encountered sequentially: “items are encountered one at a time” (p. 11). They emphasize that “...the two most important assumptions of the conventional foraging models are long-term average-rate maximization... and the exclusivity of searching and exploiting” (p. 11). On Stephens and Krebs analysis, foraging is the general decision to search or exploit in serial encounters to maximize long-term average rates of return.

Though similar, I will discuss one main difference between my analysis and Stephens and Krebs. Recall my analysis above: foraging is the serial search for general resources in accept-or-reject, exclusive, persistent decision contexts. The difference lies in their focus on average-rate maximization. There are obvious reasons to focus on such maximization—most importantly, if foragers are as efficient as possible in the sense of maximizing their average intake rate, then their overall fitness is likely to increase (however that is measured). However, I eschew a similar focus as I am concerned with foraging generally. Organisms may or may not maximize average rates of return when foraging, and performance during foraging cannot distinguish foraging from non-foraging: such a measure assumes some distinction between foraging and other phenomena. In short, Stephens and Krebs’ normative focus threatens to subvert the descriptive goal of analyzing foraging. I agree however that from a normative perspective, a focus on average rates of return is important.

Other foraging researchers have offered state-based approaches (Mangel and Clark 1986; Houston and McNamara 1999). Consider the account of Universal Foraging Theory (UFT) from Mangel and Clark (1986). UFT is based on three core concepts: first, “a state variable (or set of variables)... that characterizes the current physiological state of the forager... and that changes in a stochastic fashion, depending upon the state of the environment, the state of the forager, and its decisions”; second, “a concept of fitness that is directly related to the long-term contribution to the gene pool”; and third, “a methodology for determining the behavioral strategy (as a function of the state variable) that optimizes fitness over a long time interval, using stochastic dynamic programming... and Markov decision processes” (Mangel and Clark 1986, p. 1128). The key to their approach is “...the introduction of a state variable, $X(t)$, which characterizes the condition of the forager at time t ” (Mangel and Clark 1986, p. 1128). On UFT, foraging refers to decisions based on the state of the agent to maximize long-term average rates of return.

The UFT account emphasizes how decisions made when foraging change both the environment and the state of the forager. Many models of foraging emphasize these changes. Such accounts suffer from being both too general and too focused. First, the focus on the state of the agent limits the applicability of the analysis. There may be foraging decisions that do not turn on the state of the agent, especially in its extended application outside food foraging (discussed below). In addition, the focus on any decisions based on the agent’s state to maximize long-term return applies to far too many types of decision. This is a deliberate choice: “[w]e choose

[‘foraging’] rather than Unified Theory of Animal Behavior or Mathematical Ethology because foraging behavior must always be taken into consideration by an animal, unless it is sleeping or hibernating” (Mangel and Clark 1986, p. 1128). But then, the approach fails to differentiate foraging from search or decision making. In addition, as with the Stephens and Krebs approach, the normative focus threatens the descriptive aim of providing an analysis of foraging.

Hills and colleagues have offered an account of foraging in terms of area-restricted search (ARS) (see, e.g., (Todd and Hills 2020)). As they state, “patterns of extensive and intensive foraging in response to resource absence or presence, respectively, have also been widely observed across species... This pattern of movement is called area-restricted... search” (Hills et al. 2013, p. 1). Area-restricted searches are those where the agent makes smaller movements with greater turning radii to concentrate on a local area. On ARS, foraging is the area-restricted search for resources.

While many foraging contexts do involve ARS, not all do. Consider again trapline foraging or central-place foraging. Trapline foraging is explicitly not area-restricted because the locations can range far apart. Central-place foraging is also not area restricted as the search for prey items can range quite widely. The ARS defender might object that trapline foraging and central place foraging are area restricted. Trapline foraging is restricted to the subset of locations where resources are known to be present. In reply, these locations may be spread quite widely, so it is unclear in what sense the area is restricted. Nonetheless, at each location, the search can be very restricted, including to a single location (e.g., where a hunter has placed a trap). Central place foraging may also be area restricted as foragers venture out into patches and, once they have harvested resources, then return home. The foraging is restricted to patches, however, which occupy a restricted area of the environment. In reply, while I maintain that such foraging need not be area restricted, a second argument against ARS is that not every instance of an ARS is foraging. For example, targeted search may be area restricted but is not foraging. Consequently, even if every case, or some subset of central cases, of foraging is area-restricted, the proposal would still be insufficient as an analysis of foraging. Similarly, the area-restricted search approach does not distinguish foraging from decision making more generally.

Finally, consider one last approach to an analysis of foraging, which focuses only on the algorithms and processes for making choices, and does not consider the properties of contexts in which decisions are made. While not explicit advocates, Stephens and Krebs (1986) can also be

interpreted as implicitly adopting this way of understanding foraging. On this behavior account, foraging is defined in terms of the way that some decisions are made: foraging is a decision made in a foraging way while searching. A foraging way of making a decision will vary depending on the account of such decisions on offer. On optimal foraging theory, for example, one sort of foraging decision is a decision made by comparing instantaneous to average rates of resource intake. Foragers must track the average values of items in the environment, the value of current offers, and perform a comparison between them to make a decision. This account disregards the context in which a search is made, identifying foraging purely in terms of decision processes.

The main problem with the behavior account is that it fails the differentiation criterion. What makes a way (algorithm or process) of making a decision a foraging way? The defender of the behavior account can't merely say that some ways of making a decision just are foraging ways; that either begs the question or leaves foraging as a brute analysandum, depending on how this reply is interpreted. Further, on this account, any way of making a decision could be seen as a foraging way. But then, foraging is no longer any different from other types of decisions, violating the differentiation criterion. In addition, the sole focus on decision processes results in a failure to distinguish foraging from search more generally.

In sum, alternative approaches to the analysis of foraging are either unduly normative (Stephens and Krebs account, UFT), fail to sufficiently analyze central cases of foraging (UFT, ARS), or fail the differentiation criterion (behavior account). In contrast, the analysis of foraging as serial search for general resources in accept-or-reject, exclusive, persistent decision contexts distinguishes foraging from search or decision making more generally as well as capturing important central cases of foraging. But does the approach help ground predictions?

V. Extending Foraging Beyond Food

The discussion thus far has focused on examples of foraging for external resources. However, as mentioned in the introduction, foraging models and analyses are increasingly used in cognitive science and cognitive neuroscience to explain search for internal resources. Can the analysis accommodate these recent developments? I will use an example to illustrate how the analysis does cover those cases. In addition, I will suggest the analysis helps ground predictions like that foraging behavior may be observed certain classes of task that may not intuitively

correspond to foraging tasks but in fact may be considered foraging because the task representation embeds the cyclical graph.

There are two main advantages to such an extension. First, by analyzing tasks as foraging tasks, the structure of the decision environment becomes clearer. These decision environments might then elicit foraging behavior. The general idea is that the pressures present in the selective environments of many organisms in their evolutionary history select for choice processes that reflect those pressures. For example, if the ancestral choice environment is a foraging one, then one has grounds to predict that when faced with choice problems in their environment, animals may exhibit foraging choice behavior—behavior that has been selected for success at foraging. So, if cognitive problems also exhibit foraging-context-like properties, then we might be able to explain cognitive behavior in foraging terms. Second, getting clear on the structure of foraging decision contexts helps us specify what good or even optimal performance looks like. Optimal foraging theory, such as the presentation in Stephens and Krebs mentioned above, contains a wealth of formal tools and concepts for understanding how foragers should behave. The application of these tools to unobvious foraging contexts can suggest new norms for behavior, norms that say how agents should act in those contexts to behave optimally. While I won't have the space herein to fully flesh out this idea, below I argue that reasoning embeds a foraging context. Because it embeds a foraging context, there may be novel reasoning norms that can be stated that define optimal or successful reasoning.

i. Other Resources

As noted in the introduction, foraging theory has been applied to the search for a wide range of non-alimentary resources, including information (Pirulli and Card 1999; Fu and Pirulli 2007), concepts (Hills et al. 2015), and memory (Hills et al. 2012). While the desideratum stated above that an analysis of foraging should unify a range of different types of foraging did not explicitly limit those types to searches for food (or water and the like), the examples adumbrated were searches for alimentary resources. This raises the question of whether the analysis can be extended to searches for abstract (e.g., information) or internal (e.g., concepts) resources.

Consider the semantic search task used by Hills and colleagues (Hills et al. 2012). The design of this task was very simple: participants were instructed to type as many items as possible

of a given category in three minutes. Participants were first presented with the category title (e.g., ‘animals’). They then typed out the names of as many items from the category as they could. These responses were then analyzed in terms of their distance in semantic space.

Internal search of the sort utilized on this task can be challenging to represent in the MDP framework. What is a state on this framework? States need not refer to the environment. On an internal search, the states may refer to states of the mind or brain or subsystems of the mind or brain. What about actions? Here, I appeal to the not-uncontentious notion of a mental action (see, e.g., Proust 2013). During internal search, some mental process occurs such that an item is generated, rises to the awareness of the agent, and the agent makes a mental accept or reject decision regarding the item. The task can be described as follows (where I have deliberately chosen different terms for the states and actions). (Note that this may not reflect the phenomenology of performing the task because the states and actions described may not rise to the level of conscious awareness of the participants.) During semantic search, participants are in a *Query* state, where they can *Generate* new items or *End* their efforts. If participants choose to *Generate*, some mental or neural process will generate a possible item (such as possible animal names like ‘paw’ (incorrect) or ‘tiger’ (correct)). Participants then enter an *Evaluate* state where the items are assessed, which they can choose to *Type* or *Skip*. If they choose to *Skip* an item, they return to *Query*; if they choose to *Type* an item, they transition to a *Response* state, after which they can *Continue* the task (or *End*, not represented below). The following condensed task diagram illustrates the semantic search task (Figure 5):




Figure 5. Condensed cyclical graph for semantic search task reported in Hills et al. 2012.

This task representation shares the same structures as the cyclical model of foraging above. There are three states (*Query*, *Evaluate*, and *Response*), two actions in one of the states (*Skip* or *Type*), and the graph is cyclical. However, this is not yet sufficient to make this a foraging task. A directed cyclic graph is identified not only by the number of variables and their connections but also the identity of the variables. To use the graphical representation to identify foraging tasks, the variables in this semantic search task and other such tasks must be instances of the state and action types presented above.

First, consider the state variables, *Query*, *Evaluate*, and *Response*. The *Query* state in the task is when participants must generate items of the relevant category but have not yet done so. From this state, participants can search their memory stores to generate a possible reply or they can end the task. This state is a type of exploration. In *Explore* states above, foragers search for resources. In the semantic search task, participants search for words. Both types of search involve actions on the part of the searcher—movement in the physical world in the classic foraging context and mental actions in the semantic search task (like searching through memory). And both involve responses to outcomes in some sense outside the searcher: in foraging, coming across a prey item and entering the *Encounter* state, and in semantic search, bringing to mind a word and entering the *Evaluate* state. The nature of the search that occurs during the semantic search task implies that the *Query* state is a sort of exploration. *Evaluate* and *Response* states are more clearly instances of their relevant state types. *Evaluate* and *Encounter* both follow on exploration and pose a sort of accept-or-reject decision regarding some option, a prey item in the case of classic foraging and a word in the case of semantic search. Hence, the *Evaluate* state is one type of *Encounter*. Finally, *Exploit* and *Respond* are also relevantly similar; they both involve accruing the benefits of accepting some offer. Hence, the *Respond* state is a sort of *Exploit*.

Second, consider the action variables, *Generate*, *Skip*, and *Type*. *Skip* and *Type* are obviously ways to *Reject* and *Accept* respectively. When in the *Query* state, participants *Generate* possible responses. As noted, this action involves the searcher but also the search environment, just as *Search* does in foraging. *Generate* then is a type of *Search*. When in the *Response* state, participants can continue to *Generate* responses, leading to the *Query* state while response generation goes on. From there, participants can continue to *Generate*. Hence, *Generate* initiates a search process, much as *Search* leading from *Exploit* is to commit to looking for the next item. In sum, the graph of the semantic search task embeds the foraging graph as the graph structure of

the task matches the foraging graph and the task states and actions are tokens of the respective types of foraging states and actions.

Granted that the semantic search task embeds the cyclical foraging graph, the analysis classifies the semantic search task as a foraging task. In light of selective considerations on choice mechanisms, the classification helps ground predictions of foraging behavior. Many models of foraging decisions (such as those described in Stephens and Krebs 1986) predict that decisions to accept or reject offers will be dictated by current offers or instantaneous outcomes from choices dropping below the average outcome across the environment. In these models, foragers compare a current offer to a threshold set by the value of a generic option. If the current offer drops below that threshold, then they reject the offer. The value of the generic option is often set to the average across the environment. In that case, the current offer is compared to the average offer and, if the current offer is less than the average, the current offer is rejected.

Such effects of averages are indeed reported by Hills and colleagues. Specifically, participants tended to switch between ‘word patches’, groups of animals that share semantic similarity, when the similarity between items dropped to the average similarity across all items. This effect is akin to the depletion of a patch of resources in the environment: harvest resources at a patch until it is depleted below the average on offer in the environment, and then decide to explore for another patch or go home (i.e., end the task). Semantic similarity is a good currency to make these decisions because response times are correlated with it. More similar responses are made more quickly. By ensuring that responses are more similar than the average semantic similarity, responses will be quicker and subjects will be able to provide more responses in the allotted time. In opposition to this evidence, Abbott and colleagues present evidence that the behavior reported by Hill and colleagues is equivocal: that behavior could result from a foraging choice algorithm that compares the value of offers to average values to make a decision or from a choice algorithm that selects a random direction on a graph with a certain structure. This debate, however, is about the choice process. On my analysis of foraging, animals can make foraging decisions in many ways so long as choices are accept-or-reject and exclusive; the disagreement between Abbott and Hills is about how foraging decisions are made, not about whether agents are foraging. In both algorithms, decisions are accept-or-reject, exclusive choices: options are accepted or rejected (not accepted and rejected) and decisions between options are made singly. My analysis, then, partly

helps resolve the disagreement by refining the debate to be about how foraging decisions are made and not whether participants are foraging.

This illustration is intended to hold generally. The analysis applies to other types of tasks, specifically those that regard the search for internal resources. Those tasks embed the cyclical foraging graph, implying that they are a type of foraging. Consequently, the analysis, which implies the cyclical model, can be extended to cases of foraging for internal and other types of resources.

ii. Reasoning

Having illustrated how the cyclical model can be extended to tasks where foragers search for resources other than food, water, and the like, what is the breadth of tasks that the model can encompass? This is related to the third desideratum, that an analysis of foraging be able to be used by scientists to help ground predictions. I believe that many but not all tasks embed the cyclical foraging graph. Here, I will make the case that some instances of reasoning, in the sense of transitioning from a set of premise sentences (premises) to a conclusion sentence (conclusion), are instances of foraging.

Reasoners start with a set of premises and transition to some conclusion or other. Consider the classic argument:

All humans are mortal.

Socrates is human.

Therefore, Socrates is mortal.

This argument consists in the transition from the set of premises {‘*All humans are mortal*’, ‘*Socrates is human*’} to the conclusion {‘*Socrates is mortal*’} (or, alternatively, from the belief that all humans are mortal and the belief that Socrates is human to the concluding belief that Socrates is mortal; I won’t distinguish between reasoning as operating over attitudes and reasoning as operating over sentences herein). This example obscures the large number of other possible conclusions that a reasoner could draw. For example, other possible classically valid conclusions include ‘*All humans are mortal and Socrates is human*’, ‘*Socrates is human or the moon is made of cheese*’, and many more. Besides these valid conclusions, there are innumerable invalid ones as

well, such as ‘*Socrates is not mortal*’. And, of course, reasoners make other kinds of transitions including inductive, abductive, and more.

What is the nature of the consideration of the set of possible conclusions? During explicit conscious reasoning, reasoners will sometimes start with a set of premises, consider what those premises support, and then transition to one of the conclusions under consideration. This consideration process starts with some possible conclusion and evaluates the justification for the conclusion in light of the premises. Hence, explicit conscious reasoning imposes a serial restriction on the reasoning process. Reasoning is sometimes a process in which, given a set of premises, reasoners serially consider possible conclusions before transitioning to one of them.

This description of the task confronting the reasoner is sufficient to establish its foraging pedigree. We can construct a MDP for reasoning and represent this process as a cyclical graph. Reasoners start out in some *Premise* state corresponding to having in mind some set of premises. Reasoners then *Generate* some possible conclusion, whereby they enter a state of *Consideration* and face the decision of whether to *Adopt* it and transition to a new *Update* state (where they possibly change their view; cf. (Harman 1986)) or to *Discard* it and return to the *Premise* state to *Generate* new possible conclusions in turn. After *Update*, they *Select* the premises for the next move in reasoning (such as a new set of premises that includes the newly adopted sentence). This context can be represented (Figure 6):

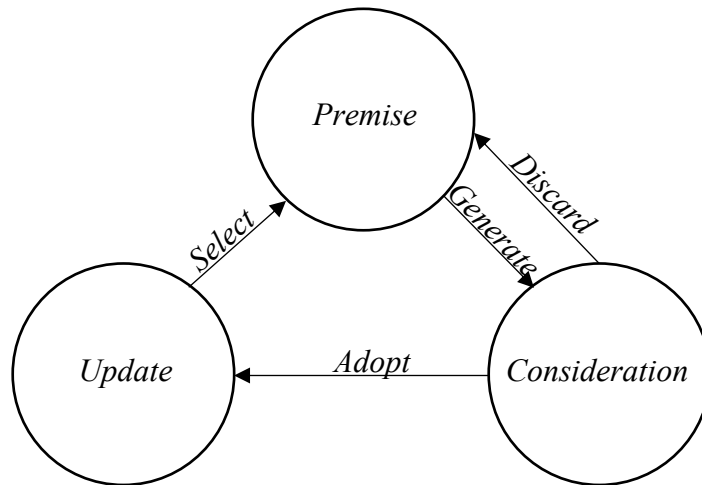


Figure 6. Condensed cyclical graph for reasoning from a set of premises to some conclusion.

This proposal is less obviously an instance of the cyclical foraging model. Take the states first. *Consideration* is a type of *Encounter* insofar as the reasoner considers what action to take regarding some item, the conclusion under examination. *Update* is also a type of *Exploit* insofar

as (say) changes in the reasoner's view are the outcomes of adopting the conclusion, just as a forager might gain energy as the outcome of accepting a prey item. But in what sense is *Premise* a type of exploration? Reasoners come up with conclusions given some set of premises and the aims of their inquiry. *Premise* is the state of the reasoner before a conclusion has been generated but after premises have been selected. Reasoners are in such a state during reasoning. But is this a state of exploration? Being engaged in reasoning is to engage in inquiry, to determine the answer to a question, say. But then, reasoners are searching for an answer to the question. In addition, the state involves action on the part of the reasoner, viz. generating possible answers to their question. Finally, the reasoner responds to the outcomes of this process by entering the *Consideration* state and making a decision with regard to the answer. *Premise*, then, is a type of *Exploration* state.

The (mental) actions undertaken by reasoners are also instances of the types of actions taken by foragers. *Adopt* and *Discard* share obvious properties with *Accept* and *Reject*; adopting some conclusion results in (say) the update of the reasoner's view, which is akin to gathering the calories from a prey item, and discarding a conclusion means reasoners must search for a new one, akin to rejecting a prey item in search of others. The other two actions are *Select* and *Generate*. Unlike the other instances of foraging tasks discussed above where the same action *Search* leads from the *Exploit* state to the *Explore* state and from the *Explore* state to the *Encounter* state, these two actions are not the same. However, they are both types of *Search*, albeit with different search areas. *Select* is the action taken by reasoners to determine the premises for the next move in reasoning after (say) updating their view. Reasoners must search through the numerous possible premise sets to use to further their inquiry, and *select* is the action taken to commit to that premise set. In contrast, *Generate* is the action taken after the premises have been determined and that produces possible conclusions. As illustrated above, for any given set of premises, there are innumerable conclusions, and reasoners must generate conclusions to evaluate and potentially *Adopt* or *Discard*. Both *Select* and *Generate* are types of *Search*.

The extension of the cyclical model to reasoning contexts provides grounds for making predictions. Much like the semantic search task, the analysis classifies reasoning as a foraging task. And, just as with the semantic search task, this classification can help ground predictions by applying foraging models to reasoning behavior. Applying these models is a non-trivial task, however; as the discussion of Stephens and Krebs above suggested, a number of assumptions must be made to apply such models, such as the specification of a currency, how that currency is

transformed by the decision variable, and the rule for making decisions. These models can be applied to the search for premises or the search for conclusions only once these preliminary issues are settled. But the analysis helps ground predictions that these models will describe reasoning behavior better than other search or decision models that don't match the analysis, such as observing that reasoners keep track of their current progress in solving some problem, are sensitive to the average performance or progress in reasoning toward a goal, and make decisions about changing reasoning strategies when their current progress drops below that average.

VI. Conclusion

In this paper, I have presented an analysis of foraging. Foraging is the serial search under constraints for general resources in accept-or-reject, exclusive, persistent decision contexts.

At the outset, I listed three desiderata that a satisfactory analysis of foraging should satisfy: first, it should differentiate foraging from search and decision making more generally (differentiation desideratum); second, it should unify what different types of foraging have in common (unification desideratum); third, it should make predictions (empirical desideratum). How does my analysis stack up against these desiderata?

First, foraging is distinct from search or decision making more generally. Not all search is foraging; parallel or targeted search for example are not foraging. And not all decision making is foraging; simultaneous choices made between multiple items are not foraging. The analysis satisfies the first desideratum.

Second, the analysis of foraging implies that foraging contexts have a certain structure. This structure can be formalized with MDPs and directed graphs. This formalization reveals a cyclical foraging graph. Various foraging tasks can then be formalized to see if they embed the cyclical foraging graph. The result of this process is the unification of a wide range of foraging contexts, including both external and internal searches.

Third, the cyclical graph can be embedded in other contexts that *prima facie* do not resemble foraging decision contexts. I illustrated this extension of foraging theory by arguing that reasoning is a type of foraging and that this helps ground predictions that can be assessed by analyzing reasoning behavior with models of foraging. The extension of the analysis to reasoning helps ground predictions, satisfying the third desideratum.

References

- Abbott, J. T., J. L. Austerweil and T. L. Griffiths (2015). "Random walks on semantic networks can resemble optimal foraging." Psychological review **122**(3): 558-569.
- Abrams, M. (2023). "Random foraging and perceived randomness." Philosophy of Science.
- Barack, D. L., A. Bakkour, D. Shohamy and C. D. Salzman (2023). "Visuospatial information foraging describes search behavior in learning latent environmental features." Scientific Reports **13**(1): 1126.
- Berger-Tal, O. and S. Bar-David (2015). "Recursive movement patterns: review and synthesis across species." Ecosphere **6**(9): art149.
- Brunelli, R. (2009). Template matching techniques in computer vision: theory and practice, John Wiley & Sons.
- Calhoun, A. J., S. H. Chalasani and T. O. Sharpee (2014). "Maximally informative foraging by *Caenorhabditis elegans*." Elife **3**: e04220.
- Charnov, E. L. (1976). "Optimal foraging, the marginal value theorem." Theor Popul Biol **9**(2): 129-136.
- Davidson, J. D. and A. El Hady (2019). "Foraging as an evidence accumulation process." PLoS computational biology **15**(7): e1007060.
- Elgin, M. and E. Sober (2002). "Cartwright On Explanation And Idealization." Erkenntnis **57**(3): 441-450.
- Freeman, R. (1968). "Charles Darwin on the routes of male humble bees." Bulletin of the British Museum (Natural History) Historical Series **3**(6): 177-189.
- Fu, W.-T. and P. Pirolli (2007). "SNIF-ACT: A cognitive model of user navigation on the World Wide Web." Human-Computer Interaction **22**(4): 355-412.
- Genovesio, A., S. P. Wise and R. E. Passingham (2014). "Prefrontal-parietal function: from foraging to foresight." Trends in Cognitive Sciences **18**(2): 72-81.
- Gill, F. B. (1988). "Trapline foraging by hermit hummingbirds: competition for an undefended, renewable resource." Ecology: 1933-1942.
- Giraldeau, L.-A. and T. Caraco (2000). Social foraging theory. Social Foraging Theory, Princeton University Press.
- Gold, J. I. and M. N. Shadlen (2007). "The neural basis of decision making." Annu Rev Neurosci **30**: 535-574.
- Harman, G. (1986). Change in view: Principles of reasoning. Cambridge, MA, MIT Press.
- Hills, T. T., M. N. Jones and P. M. Todd (2012). "Optimal foraging in semantic memory." Psychological review **119**(2): 431.
- Hills, T. T., P. M. Todd and M. N. Jones (2015). "Foraging in semantic fields: How we search through memory." Topics in cognitive science **7**(3): 513-534.
- Houston, A. I. and J. M. McNamara (1985). "A general theory of central place foraging for single-prey loaders." Theoretical Population Biology **28**(3): 233-262.
- Houston, A. I. and J. M. McNamara (1999). Models of adaptive behaviour: an approach based on state, Cambridge University Press.
- Kelly, R. L. (2013). The lifeways of hunter-gatherers: the foraging spectrum, Cambridge University Press.
- Krebs, J. R., A. Kacelnik and P. Taylor (1978). "Test of optimal sampling by foraging great tits." Nature **275**(5675): 27-31.

- Mangel, M. and C. W. Clark (1986). "Towards a unified foraging theory." Ecology **67**(5): 1127-1138.
- McNamara, J. (1982). "Optimal patch use in a stochastic environment." Theoretical Population Biology **21**(2): 269-288.
- Mobbs, D., P. C. Trimmer, D. T. Blumstein and P. Dayan (2018). "Foraging for foundations in decision neuroscience: insights from ethology." Nature Reviews Neuroscience **19**(7): 419-427.
- Newell, A. (1994). Unified Theories of Cognition, Harvard University Press.
- Noser, R. and R. W. Byrne (2010). "How do wild baboons (*Papio ursinus*) plan their routes? Travel among multiple high-quality food sources with inter-group competition." Animal cognition **13**(1): 145-155.
- O'Farrell, S., J. N. Sanchirico, O. Spiegel, M. Depalle, A. C. Haynie, S. A. Murawski, L. Perruso and A. Strelcheck (2019). "Disturbance modifies payoffs in the explore-exploit trade-off." Nature communications **10**(1): 1-9.
- Orzack, S. H. and E. Sober (1994). "Optimality models and the test of adaptationism." The American Naturalist **143**(3): 361-380.
- Pacheco-Cobos, L., B. Winterhalder, C. Cuatianquiz-Lima, M. F. Rosetti, R. Hudson and C. T. Ross (2019). "Nahua mushroom gatherers use area-restricted search strategies that conform to marginal value theorem predictions." Proceedings of the National Academy of Sciences **116**(21): 10339-10347.
- Passingham, R. (2021). Understanding the prefrontal cortex: selective advantage, connectivity, and neural operations, Oxford University Press.
- Payne, S. J., G. B. Duggan and H. Neth (2007). "Discretionary task interleaving: heuristics for time allocation in cognitive foraging." Journal of Experimental Psychology: General **136**(3): 370.
- Pearson, J. M., K. K. Watson and M. L. Platt (2014). "Decision making: the neuroethological turn." Neuron **82**(5): 950-965.
- Pirolli, P. and S. Card (1999). "Information foraging." Psychological review **106**(4): 643.
- Pirolli, P. L. T. (2007). Information foraging theory: Adaptive interaction with information, Oxford University Press.
- Proust, J. (2013). The philosophy of metacognition: Mental agency and self-awareness, OUP Oxford.
- Rhodes, T. and M. T. Turvey (2007). "Human memory retrieval as Lévy foraging." Physica A: Statistical Mechanics and its Applications **385**(1): 255-260.
- Royama, T. (1970). "Factors governing the hunting behaviour and selection of food by the great tit (*Parus major* L.)." The Journal of animal ecology: 619-668.
- Stephens, D. (2008). "Decision ecology: Foraging and the ecology of animal decision making." Cognitive, Affective, & Behavioral Neuroscience **8**(4): 475-484.
- Stephens, D. W. and J. R. Krebs (1986). Foraging Theory. Princeton, NJ, Princeton University Press.
- Todd, P. M. and T. T. Hills (2020). "Foraging in mind." Current Directions in Psychological Science **29**(3): 309-315.
- Yee, E., E. G. Chrysikou and S. L. Thompson-Schill (2014). "Semantic memory."