

# From foraging to auto-noetic consciousness: The primal self as a consequence of embodied prospective foraging

Thomas T. HILLS<sup>1\*</sup>, Stephen BUTTERFILL<sup>2</sup>

<sup>1</sup> Department of Psychology, University of Warwick, CV4 7AL, Coventry, UK

<sup>2</sup> Department of Philosophy, University of Warwick, CV4 7AL, Coventry, UK

**Abstract** The capacity to adapt to resource distributions by modulating the frequency of exploratory and exploitative behaviors is common across metazoans and is arguably a principal selective force in the evolution of cognition. Here we (1) review recent work investigating behavioral and biological commonalities between external foraging in space and internal foraging over environments specified by cognitive representations, and (2) explore the implications of these commonalities for understanding the origins of the self. Behavioural commonalities include the capacity for what is known as area-restricted search in the ecological literature: this is search focussed around locations where resources have been found in the past, but moving away from locations where few resources are found, and capable of producing movement patterns mimicking Lévy flights. Area-restricted search shares a neural basis across metazoans, and these biological commonalities in vertebrates suggest an evolutionary homology between external and internal foraging. Internal foraging, and in particular a form we call embodied prospective foraging, makes available additional capacities for prediction based on search through a cognitive representation of the external environment, and allows predictions about outcomes of possible future actions. We demonstrate that cognitive systems that use embodied prospective foraging require a primitive sense of self, needed to distinguish actual from simulated action. This relationship has implications for understanding the evolution of auto-noetic consciousness and self-awareness [*Current Zoology* 61 (2): 368–381, 2015].

**Keywords** Self-projection, Mental time travel, Episodic future thinking, Foraging, Consciousness

While studying the abilities of rats to navigate mazes, Tolman and colleagues observed that in many cases rats took routes that suggested they could piece together memories in a way that allowed them to make novel inferences (Tolman, 1948; Tolman and Gleitman, 1949). As opposed to simple cue-response relationships that might be sufficient to learn a maze, the animals demonstrated the ability to make the kind of inferences that would require search over some form of internal representation—akin to realizing that one could go to India by taking a novel route around the earth. Tolman (1948) called these representations “cognitive maps” and argued that they provided a vast repertoire of behaviour not conceivable under the standard behaviorist’s stimulus-response theories. What is meant by a cognitive map and which species, if any, have them is highly debated (Cheung et al., 2014; Cruse and Wehner, 2011; Tsoar et al., 2011). For our purposes it is necessary only that these “maps” be cognitive representations of the world and involve a sense of place such that some things can be nearer to one another than others. The capacity to search through internal representations provides a means

for working out how to get from one place to another. But it also allows another kind of behaviour—deliberation. That is, cognitive representations that can be internally searched should provide organisms with the capacity to explore alternatives, via internal foraging, prior to committing themselves to one course of action. Meunzinger and Fletcher (1938) called behavioural evidence consistent with internal foraging—that is, apparent deliberation at a choice point—*vicarious trial-and-error and learning*, and showed that it correlated with subsequent performance (see also Hu and Amsel, 1995). More recent evidence from neuroscience indicates that animals at choice points activate areas of the brain associated with the outcomes of past decisions at that point. One goal of this article is to provide a basis for understanding why this kind of internal foraging may have been pre-adapted for by the evolution of external foraging.

A second goal is to explore the consequences of internal foraging for the existence of a precursor to self-awareness, or what Tulving (1985) called *auto-noetic consciousness*: “Auto-noetic (self-knowing) conscious-

---

Received Nov. 1, 2014; accepted Feb. 12, 2015.

\* Corresponding author. E-mail: t.t.hills@warwick.ac.uk

© 2015 *Current Zoology*

ness is the name given to the kind of consciousness that mediates an individual's awareness of his or her existence and identity in subjective time extending from the personal past through the present to the personal future" (p. 1, 1985). More specifically, in relation to distinctions made by William James (1890), the "I" is conscious and the "Me" is one of many things the "I" can be conscious of. We are interested in the "Me," the potential object of self-awareness. Our aim will be to describe how a primal version of this self is a consequence of internal foraging. However, readers should not be too worried, as we promise to arrive far short of the kind of self-awareness commonly attributed to humans (Markowitsch and Staniloiu, 2011). Nonetheless, auto-noetic consciousness has been tied to behaviors such as mental time travel (Schacter et al., 2007; Suddendorf and Corballis, 2007), self-projection (Buckner and Carroll, 2007) and episodic future thinking (Atance and O'Neill, 2001), and these have often been proposed as exclusively human capacities. This naturally raises the question of where these abilities may have originated. Our proposal entails that self-awareness and its associated cognitive consequences are unlikely to be exclusively human capacities.


Connecting the dots between foraging and auto-noetic consciousness requires understanding the selective forces most likely to have driven the evolution of cognition, specifically the trade-off between two opposing behavioral motivators—exploiting known resources and exploring for unknown resources. In what follows, we will first discuss the comparative behavioral and biological commonalities underlying adaptations for resolving this trade-off. Then we will provide experimental evidence that this capacity is conserved to search internal representations, and that this is indeed a core function of what cognitive psychologists mean by executive (or effortful) cognition. Finally, we will provide evidence that these capacities are necessary for foresight, or rather, looking ahead, and that, in turn, cognitive systems that use embodied prospective foraging must distinguish between real and simulated events and thus require a precursor to auto-noetic consciousness—recognition of the self.

## 1 External Foraging

Among metazoans, the capacity to forage—that is, to adaptively search for and exploit resources—is ubiquitous. The strategies associated with foraging are varied, ranging from the relatively simple, such as sit-and-wait or random walks, through to more systematic search

strategies (Banks et al., 2008; Stephens and Krebs, 1987). Among these, one strategy for mediating the trade-off between exploration and exploitation—called area-restricted or area-concentrated search—is extremely common. Animals that use area-restricted search localize search around areas where resources have been found in the past (Kareiva and Odell, 1987), but move away from locations where resources are encountered less frequently. A classic example is houseflies increasing their turning angle when they encounter a drop of sucrose, such that they continue searching nearby, but reducing their turning angle as the time since the last encounter grows, so that they move away to explore new locations (Bell, 1991). Area-restricted search is observed throughout metazoans (i.e., animals), including nematodes, moths, leeches, rodents, and humans (see Fig. 1; Hills, 2006). Area-restricted search is also found among animals that are well known for other, more idiosyncratic search behaviors, such as the desert ant *Cataglyphis bicolor*, which is renowned for its path-integration abilities, but quickly engages in area-restricted search when path-integration fails (Wehner and Srinivasan, 1981). Moreover, increasingly research is suggesting that the power-law distributed movement patterns associated with Lévy flights may also be produced by area-restricted search (Hills et al., 2013; Plank and James, 2008).

The commonality of area-restricted search is explained by its adaptation to locating clustered resources. Area-restricted search is a useful strategy for tracking resources that are spatially-autocorrelated, especially when patch boundaries are not well defined. When an encounter with a resource indicates other resources nearby, an animal's turning angle should be a function of resource encounters. This has been demonstrated in numerous studies (Grünbaum, 1998; Hills, 2006; Kareiva and Odell, 1987). Moreover, in direct comparisons with memoryless strategies, such as Lévy flights, memory-based strategies like area-restricted search have a distinct performance advantage when resources are clustered (Ferreira et al., 2012; Plank and James, 2008). Further, species-area distributions are consistently found to be more aggregated than random placement models would predict, indicating spatial-autocorrelation (Harte et al., 2008), and seed dispersal distributions of primary producers also tend to fall-off with distance from the parent (Janzen, 1970); both suggest a strong tendency for resource clustering in the natural world, though whether foragers experience this will depend on the scale at which they forage.



**Fig. 1** Examples of area-restricted search in nematodes and humans

**A.** Paths taken by the nematode *Caenorhabditis elegans* immediately after removal from food (0–5 minutes) and 30 minutes later (30–35 minutes). The white arrow shows a period of roughly straight, forward motion; the dark arrow shows a high-angled turn. The bottom panel of **A** shows the frequency of high-angled turns over the two periods of observation. Reproduced from Hills et al. (2003). **B.** Paths of humans foraging for invisible resources in a 200 m diameter virtual foraging arena with diffuse and clustered resources. The lower panel shows that when encountering resources (indicated by a tone) humans in clustered environments take sharper turns, indicating area-restricted search. Reproduced from Hills, Kalff, and Wiener (2013).

In some cases, the ubiquity of area-restricted search is certainly the outcome of convergent evolution. One can reasonably argue that bacteria engage in area-restricted search, because the second-messenger systems they use to track resources have dephosphorylation times on the order of seconds (Macnab and Koshland, 1972). This allows them to detect differences in resource concentrations larger than their own body length, which in turn allows them to “remember” that a position they were in a moment ago is better than a position they are in now. They can then make a random turn in an effort to return to the general vicinity of their previous location.

In vertebrates, on the other hand, there is strong reason to believe that area-restricted search is symplesiomorphic—a trait shared across a variety of extant taxa that was also possessed by their remote common ancestor. The principal reason for this inference is the shared biological basis of area-restricted search (reviewed in Hills, 2006; Hills and Dukas, 2012). In invertebrates, behaviors modulated by dopaminergic signalling involve feeding or foraging related behaviors, while in vertebrates the behaviors mediated by dopaminergic signalling are more consistent with area-restricted search,

such as spatial foraging and active exploration. Though not without exceptions, the dominant pattern of dopaminergic signalling facilitating persistent or stereotypic behaviors is found across vertebrates (Barron et al., 2010; Winstanley et al., 2012).

## 2 The Evidence for Structure in Internal Foraging

Area-restricted search represents a behavioral solution to the exploration-exploitation trade-off necessary to forage successfully for resources clustered in space. However, the exploration-exploitation trade-off is in the same class as other evolutionary trade-offs known to influence evolution across a variety of domains—such as the surface-area-to-volume trade-off or the size-number trade-off. As a consequence, solutions to the exploration-exploitation trade-off are found across natural living systems, from stress induced mutagenesis at the genomic level to DNA sharing at the population level (Claverys et al., 2006; Galhardo et al., 2007). Importantly for our present purposes, the exploration-exploitation trade-off is also commonplace in cognitive domains, such as in memory search, information search,


problem solving, and visual search (Hills et al., 2015). In each of these domains, cognitive systems must decide between continuing with one course of action that has reasonably certain consequences, versus exploring other possible courses of action for which the outcomes are less certain.

Memory search provides one domain where internal search shares features common to animals foraging in space. In particular, both internal and external search appear to involve dynamic patterns of search over structured domains; items can be “near” or “far” from the position of search. One approach to understanding the structure of memory has been to develop quantitative representations that can then be used to predict and measure behavior. For example, semantic representations of memory derived from statistical models that read through Wikipedia and compute quantitative measures of word similarity have been used with success to predict the order in which people recall lists of items from memory (see Fig. 2A; Hills et al., 2012; Hills et al., 2013; Thompson and Kello, 2014). These structures also allow for testing alternative search algorithms. In model comparisons, the best fitting models of iterated memory search are two-stage models that assume local foraging (using one item in memory to activate the next item in memory) with occasional non-local, long-range transitions to richer regions of the memory representation (Hills et al., 2012; Hills et al., 2013; Hills and Pachur, 2012). This is similar to area-restricted search, and some of the earliest mathematical models of serial recall from memory assume area-restricted search-like dynamics (Raaijmakers and Shiffrin, 1981). Other mod-

els of memory also make the central assumption that memory is a high-dimensional manifold that has structure, such that items in memory can be near or far from one another (Austerweil et al., 2012; Thompson and Kello, 2014).

Problem solving is another domain that provides evidence of internal foraging. Cognitive models of problem solving often use quasi-spatial representations to predict how people will arrive at solutions. For example, in an anagram-like task where people are asked to produce multiple four-letter words from one letter set, such as RNTPOE, people produce solutions that are more similar by string-edit distance to one another the more proximate they are in the series of produced solutions (Hills et al., 2008; 2010). In the above example, a series of solutions might look as follows: ROPE, ROTE, NOTE, RENT, PENT, etc. The first three and last two solutions can each be arrived at in series by changing only one letter, whereas the third and fourth solutions (NOTE to RENT) suggest a larger transition across the solution space. Using a computational model called the executive search process, Hills et al. (2010) demonstrated that search in this task could be explained by assuming a local search process that made iterative changes to the prior solution, but occasionally emptied the memory buffer and started again by sampling from the original letter set.

The Remote Associates Test is another problem solving domain that is well modeled by local search strategies. The task is to find, as quickly as possible, a word that is common to three other words, such as MOON, DEW, and COMB. Here the answer is HONEY. In a



**Fig. 2** Examples of structure in internal foraging

**A.** A typical sequence of items produced in the animal fluency task in which participants are asked to “name all the animals you can think of.” Participants often produce items in clusters of semantically similar items (solid lines) with infrequent transitions between clusters (dotted line). **B.** A typical sequence of items produced in the remote associates task, in which participants are asked to produce a word that is common to each of three other words. Here, the words are COMB, MOON, and DEW. In Smith et al. (2013), participants were also asked to produce prospective solutions that came to mind as they searched for the final solution. Participants appeared to produce prospective solutions in clusters (solid lines) in association with a cue word, and transitioned between these (dotted lines) as they homed in on their final solution.

recent study asking participants to produce the incorrect solutions they thought of while hunting for the correct solution, responses were well characterized by a process that involved sequential dependencies between adjacent responses (see Fig. 2B; Smith et al., 2013).

The evidence of sequential dependencies across these domains reveals two things. First, internal search appears to share with external search the property of involving search in structured spaces. And, second, quantitative representations of internal spaces (e.g., memory or symbolic problem spaces) can be used to predict cognitive navigation of internal spaces in much the same way as Euclidean distances can be used to predict navigation in external spaces.

### 3 Evidence for an Evolutionary Relationship between Internal and External Foraging

Are the same mechanisms used to search internal and external spaces? In humans, there are three kinds of evidence for such a generalized cognitive search process: priming between domains, the general nature of executive control, and the neurobiology of cognitive control.

Evidence of priming between domains is provided in a study aimed at priming search behavior between a visuospatial search task and an internal lexical search task (Hills et al., 2008). People who first foraged in a clustered spatial environment (versus a diffuse spatial environment) subsequently searched as if there were more (or fewer) words in a group of letters in a lexical search task. This priming required search—simply harvesting resources from known locations was insufficient to produce priming (Hills et al., 2010). Thus, the search process used to forage in space appeared to be subsequently used to forage in an abstract lexical space.

A second kind of evidence for shared processes is provided by a cognitive construct called executive capacity, which has been shown to be involved in both external and internal search (Longstaffe et al., 2014; Hills and Pachur, 2012). One aspect of executive capacity can be measured by working memory span. Working memory span is a measure of the amount of independent pieces of information that can be held in short term memory, which is hypothesized to be a measure of people's ability to maintain top-down ('executive') focus on specific goals while inhibiting distracting information (Engle and Kane, 1999). Working memory span is seen as a general cognitive capacity associated with top-down control of behavior, and associated with a variety

of performance measures such as SAT scores and measures of general fluid intelligence (Unsworth and Engle, 2007). People with higher working memory spans also tend show longer periods of local foraging in internal memory representations before making long-distance jumps to new regions of the memory space (Hills et al., 2013; Hills and Pachur, 2012; see also Rosen and Engle, 1997). Thus, features of internal memory search appear to share general control processes known to be involved in other behaviors associated with the control of attention.

Finally, the capacity for behavioral focus (i.e., attention), like the capacity for spatial foraging, appears to be a conserved neuromolecular process across vertebrates. At the synaptic level, this is characterized by the dopaminergic modulation of G-protein linked receptors, inducing a cascade of intracellular events that modulate cellular responses to other neurotransmitters—most commonly, glutamate (Nicola et al., 2000). Dopamine is both excitatory and inhibitory, depending on the receptor subtype found at the membrane: the D1 subtype is typically excitatory while the D2 subtype is typically inhibitory. Together, activation of the dopamine receptor subtypes enhances the signal to noise ratio of a neural message. This facilitates activity locally while inhibiting activity more globally (Floresco et al., 1996; Seamans and Floresco, 1998). These neuromodulatory microcircuits are found throughout the midbrain (Dani and Zhou, 2004)—a region associated with the control of attention and directly connected with the prefrontal cortex, a region typically associated with goal maintenance (Miller and Cohen, 2001). In vertebrates, dopaminergic neurons in the midbrain are required for learning the relations between unconditioned and conditioned stimuli, and thus for forming associative chains between actions and outcomes (Daw and Doya, 2006). They are also involved in enhancing hippocampus-dependent memory persistence (McNamara et al., 2014), which we will show in the next section is a central component of embodied prospective foraging by vicarious trial-and-error.

The above provides evidence at three different levels—behavioral priming, cognitive modeling, and neurobiology—for a relation between external and internal foraging. Though not decisive, the evidence provides substantial support for the view that goal-directed cognition and cognitive control more generally are evolutionary descendants of spatial foraging (Hills, 2006). Indeed, dopaminergic modulation facilitates persistent or stereotypic behaviors across species (Barron et al., 2010), for example, modulating area-restricted search in

*Caenorhabditis elegans* (Hills et al., 2004), response to cocaine in *Drosophila melanogaster* (Bainton et al., 2000), exploratory behavior in the rat (Fink and Smith, 1980), and cognitive flexibility in humans (Chermahini and Hommel, 2010; Cools and D'Esposito, 2009; Winstanley et al., 2012). Furthermore, both external and internal foraging require balancing the trade-off between exploration and exploitation, and appear to share similar behavioral features and biological mechanisms for doing so (e.g., area-restricted search, (Hills et al., 2015). The capacity to use these mechanisms for more general goal-directed behavior is likely to be a derived trait borrowed (i.e., exapted) from its more spatial, ancestral trait; this would explain observations that dopamine is strongly associated with “reward,” “novelty,” and “information” (e.g., Costa et al., 2014; Bromberg-Martin and Hikosaka, 2009). This may also explain dopamine’s involvement in a wide range of goal-directed behaviors (especially in vertebrates), but involvement in only a narrower subset of foraging-related behaviors across species (e.g., Barron et al., 2010). From an evolutionary perspective, few behaviors are as widely shared as foraging with as clear an indicator of common descent as that provided by dopaminergic modulation. Though the evidence that internal and external search share a common evolutionary pathway is not incontrovertible, we know of no more plausible hypothesis nor of any proposed alternatives.

#### 4 From Internal Foraging to Embodied Prospective Foraging

Thus far, we have provided evidence that foraging in internal and external spaces shares common behavioral and biological features. However, internal foraging provides something that external foraging does not: a capacity to deliberate. As noted in the introduction, even short range, stimulus-response relationships provide animals with the ability to predict future outcomes that increase the probability of short-term rewards. Richer internal representations, such as those that support internal foraging, should provide better predictive capacities. In this section we focus on a kind of internal foraging for which there is evidence in nonhuman animals, which we call *embodied prospective foraging*. This is the form of internal foraging in which searching involves much the same processes, and draws on much the same abilities, that would be involved in actually searching on the ground, with memory standing in for experience. In identifying the evidence for embodied prospective foraging, we focus on the hippocampus, an

area of the midbrain purported to be involved in episodic memories in humans and requiring dopaminergic modulation for both flexibility and encoding of long-term memories (Shohamy and Adcock, 2010; see also McNamara et al., 2014). The hippocampus is also well understood as the vertebrate brain’s spatial representation system (e.g., at the level of place cells and grid cells, see Moser, Kropff, and Moser, 2008).

Animals replay neural activation associated with past events, at least in part by reactivating neurons in the hippocampus. This appears to be associated with both memory consolidation and goal-directed planning (reviewed in Pezzulo et al., 2014). This was initially observed during periods of sleep, when rats exposed to mazes during the day were later found to replay sequences of activation in their hippocampus concordant with those observed during their waking experience (Skaggs and McNaughton, 1996; Wilson and McNaughton, 1994). These sequences of replay could cover distances greater than 10-m (Davidson et al., 2009), and disruption of them impaired spatial learning (Ego-Stengel and Wilson, 2010). Similar patterns of activation have subsequently been found in animals during awake states (Carr et al., 2011; Diba and Buzsáki, 2007; Foster and Wilson, 2006), and these are associated with choice points prior to decision making, consistent with vicarious trial-and-error learning. Here, rats show activation patterns in their hippocampus that are associated with points in the maze in front of their current position (Johnson and Redish, 2007). These patterns or “sweeps” progress in series ahead of the animal’s current location, and terminate with activation in reward centers of the brain, such as the ventral striatum (Johnson et al., 2007; Pennartz et al., 2011; Pezzulo et al., 2014).

The most parsimonious explanation for the observed data prior to decision making is that animals use their current context to activate sequences of memory that provide information necessary for future planning. For example, in an experiment by Tolman and Gleitman (1949), a rat was allowed to freely explore a T-maze, which had a light and a dark region (both containing food) at the goal boxes at the end of each passage. After exploring, the animal was moved to a dark chamber, separate from the original T-maze, where it experienced electric foot shocks. Finally, the animal was placed once again in the T-maze, where it moved to the light goal box. One explanation for this behavior is that the rat could look ahead in the T-maze down the arm associated with the dark goal box, recall the dark goal box that then activated memories of the foot shock, and avoid

that arm (Hesslow, 2012). Recent evidence from Gupta et al. (2010) found that animals who experienced two routes, one with high frequency and another with low frequency, did not preferentially recall the more experienced sequence, but showed activation more evenly across all possible future routes prior to making a decision—as if they were exploring alternatives.

In addition, forward looking activation patterns in the hippocampus appear to take one route at a time (Johnson and Redish, 2007). In other words, the patterns of activation do not indicate a general spreading of parallel activation across all possible future routes, but are isolated to plausible routes, each in turn (Pezzulo et al., 2014). These forward looking activation patterns are most frequent when the animals have limited experience with a particular choice; as they gain experience with a choice, and especially when one of the choices is favored over another, the rats engage in fewer forward looking patterns of activation (Hu and Amsel, 1995; Johnson and Redish, 2007). The forward looking activation patterns also predict goal-directed future action in environments with multiple alternative routes, and they can generate novel routes to known goals (Pfeiffer and Foster, 2013). This further indicates that these forward sweeps of activation are not simply recall of past experience, but patterns of activation that foreshadow actions in a goal-directed manner.

The above evidence suggests that the animals studied can predict the outcomes of possible future actions by making use of an internal representation to search for exploitable resources in embodied prospective foraging. It also suggests that animals actively explore internal representations with continuous movement through a metric ‘cognitive’ space in much the same way that we described for external and internal foraging in the previous sections.

Though more difficult to study, there is considerable evidence that humans use a similar process involving hippocampal-cortical interactions—along with a number of other areas (see work on imagination networks in Hassabis et al., 2007)—to imagine future events (Buckner, 2010; Conway et al., 2003). For example, patients with damage to the hippocampus, like H.M. (Scoville, 1968) and D.B. (Klein, Loftus and Kihlstrom, 2002), have difficulty imagining future events even though they may retain general knowledge about the future. Further, episodic future thinking in humans is associated with the hippocampus—and other brain regions—and directly influences decision making (Peters and Büchel, 2010). Importantly for what we discuss in the

next section, damage to the hippocampus is also associated with a degraded sense of self (Hassabis et al., 2007; Corkin, 2002).

The observation that humans and non-human animals can, without actually taking an action, fire-up neural activation patterns that resemble the patterns of activation they would have when actually taking that action is the basis of the simulation theory of cognition (reviewed in Hesslow, 2012). This theory is based on two key observations, for which there is ample experimental support: 1) That activation of the motor and sensory cortex during non-action resembles that during real action, except that the execution of motor actions are suppressed, and 2) that this activation tends to produce representations of outcomes in proportion to the real likelihood of occurrence. Evidence that the action-like processes associated with non-action have much in common with real actions has been observed both with respect to the relative time it takes to imagine actions—for example, in mental rotation tasks (Shepard and Metzler, 1971)—as well as in relation to the neural activation patterns associated with particular actions—for example, when piano players play or imagine playing music (Meister et al., 2004).

In humans, research on the ability to predict the future often focuses on episodic future thinking (Atance and O’Neill, 2001), which is understood as the future-directed counterpart of episodic memory and sometimes also referred to as ‘mental time travel’ (Schacter et al., 2007; Suddendorf and Corballis, 2007) or ‘self-projection’ (Buckner and Carroll, 2007). Episodic future thinking is held to require metacognition (i.e., thinking about thinking), auto-noetic consciousness, concepts of self and time, abilities to identify with one’s future self, and mindreading in the form of an understanding that your future self may have needs and mental states different from those you now have (Suddendorf and Corballis, 2007; see also Perner et al., 2007). Perhaps unsurprisingly, evidence for episodic future thinking so conceived is largely or entirely confined to humans. Some have argued that there are forms of recall that, although demanding less than full blown episodic memory, do involve re-experiencing events and places rather than merely retrieving information about them (Clayton and Russell, 2009; see also Conway, 2005). The existence of embodied prospective foraging in human and non-human animals implies, relatedly, that capacities to use memory to guide actions and to predict their outcomes do not always require episodic future thinking at the level observed in humans.

In sum, evidence from vertebrate brain activity suggests a clear relationship between internal foraging and future action. The evidence further suggests that this internal foraging can involve simulation of alternative routes prior to motor execution of choice. In the next section we investigate potential consequences of this relationship for understanding the origin of the self as it features in auto-noetic consciousness.

## 5 The Primal Self as A Consequence of Embodied Prospective Foraging


We have just seen evidence that a variety of animals are capable of embodied prospective foraging. As we stressed, embodied prospective foraging is likely to demand far less cognitive sophistication than full blown episodic future thinking, at least as often understood. What embodied prospective foraging does demand, however, are the ingredients for the construction of a primal self. Or so we shall argue in this section.

Because embodied prospective foraging involves simulating actions not on the basis of actual perceptual inputs but rather on the basis of memories associated with places, objects, and events, this kind of internal foraging involves re-activating experiences of things which enable at least some of the actions that actually experiencing those things would enable. We assume for now that these memories must specify a point of view from which to act in much the way that perceptual experiences do. (We consider how our view would change if this assumption were false in the next section.) Embodied prospective foraging therefore involves an ability to switch between the animal's actual point of view and points of view associated with the animal's memories.

In addition, embodied prospective foraging involves an ability to internally generate temporally extended sequences of points of view. To see why, consider an individual who uses embodied prospective foraging to decide which of two routes to take by mentally exploring (or, as we shall sometimes say, 'querying') the first and then mentally exploring the second. The decision of which route to actually take depends on which route yields the greatest subjective gain during embodied prospective foraging. It thus involves distinguishing between two sequences of simulated actions, each involving a point of view that moves continuously through space. In principle it would be possible to combine arbitrary pieces from the two prospective sequences of points of view irrespective of whether their combination


specifies a route that could plausibly be taken and use the subjective gains associated with such combinations in deciding which route to take. But this, we assume, does not happen because—in this case—it would be adaptively useless. Rather, each prospective foraging trip is treated as a single unit because it represents a coherent set of possible actions and the associated likely subjective gains. Deciding between two routes by means of embodied prospective foraging therefore involves an ability to segment sequences of points of view in such a way that each sequence represents an interval of a possible future life. Each of these coherent sequences of points of view we will call a *q-self*, because each represents a query in the internal foraging space that could represent a future self (see Fig. 3).

How could embodied prospective foraging enable an animal to decide which of several routes to take? Consider what we will call *the naive inhibition model* (Fig. 4): When the animal is unable to select a route, the last phases of action preparation are inhibited so that cognition no longer leads to bodily movements, but only to their simulation. A decision is made at random and embodied prospective foraging occurs. This leads to reactivation of past associations in response to stimuli in the current context, which can then be input back into the cognitive system to generate further associations (e.g., Hesslow, 2002) until a plausible sequence of actions and outcomes is generated. After some time the simulation terminates. If the animal remains unable to select a route, another decision is made at random and simulation occurs again—further sampling possible sequences from memory (Pezzulo et al., 2014). Once sufficient information is accumulated to make a decision, motor and sensory areas are disinhibited and an action is taken.



**Fig. 3** During internal foraging, the animal inhibits action and queries the outcome of potential choices by activating memories in a way that reveals structure

These queries appear to be associated with coherent points of view, labelled here as *q-selves*, which each involve experiencing a stream of events (dotted lines). Here, the *p-self* is located at the actual position of the animal, as further explained in the text.



**Fig. 4 The naive inhibition model of embodied prospective foraging**


The q-self represents a query into memory, leading to internally generated sequences that provide predictions about the outcomes of future actions. Inhibition prevents sensory input and motor output, but still allows for updating during embodied prospective foraging.

This naive inhibition model makes few assumptions in that it requires no novel representational capacities on the part of the animal, only simulated action guided by selective activation of memories associated with the present context. This is sufficient to explain the data described above (e.g., Johnson and Redish, 2007; Gupta et al., 2014; Pfeiffer and Foster, 2013; Pezzulo et al., 2014). However, the naive inhibition model is fundamentally maladaptive for behavior. This is because, according to the naive inhibition model, whenever an action results in an outcome, the animal's behaviours and representations of its needs and environment are updated irrespective of whether the action is simulated or actual. But outcomes can occur in simulated action without it being beneficial for the animal to update its actual behaviors and representations. This is true even though the same outcomes occurring in actual action would require updating behaviors and representations. For example, resources can become depleted, the animal's position can change, predators can appear or disappear, mating can occur, and so on. If the naive inhibition model were correct, embodied prospective foraging would cause the animal to acquire incorrect representations about its environment.

The failure of the naive inhibition model suggests that if embodied prospective foraging is to enable an animal to decide which of several routes to take, then the animal must, in updating behaviours and representations, distinguish simulated actions and outcomes from their actual counterparts. Let us call any model involving such a distinction a *self-actuating* model. Although nothing in our account hinges on which self-actuating model is correct, it is helpful to consider a particular model in order to make it clear that distinguishing simulated actions and outcomes from actual ones does not require sophisticated perspective taking, temporal concepts or theory of mind.

Figure 5 introduces the *crude self-actuating model*, which is an attempt to improve on the naive inhibition model. This improved model states that embodied prospective foraging involves three processes. First, there is a clone process that produces some form of copy or marking of the animal's current mental states, action capacities, and memories. The clone process need not be taken literally—rather, it is a process that marks a separation between the simulated and actual animal. Second, there is a simulation process in which the animal internally forages, creating a coherent sequence of possible experiences. As noted above, this sequence is a q-self. During the simulation, experiences accrete to the q-self: the animal itself does not become sated, or falsely believe, as a consequence of the process of internal foraging, that resources have been depleted at one location relative to another. Third, there is a merge process that occurs when a q-self terminates. This involves the animal updating some beliefs and preferences, but not others, based on the outcomes of the actions associated with the q-self. For example, the animal should update with information about the availability of resources, but should not update with information that may have changed during the internal foraging process (such as resource depletion, mating encounters, or the presence of an aerial versus a ground predator).

The crude self-actuating model may be incorrect in many respects, but it exemplifies something we take to be a necessary feature of any adequate model: it requires the animal to distinguish between those sequences of points of view associated with embodied prospective foraging (the q-selves) and those sequences of points of view associated with actually foraging. As we stressed above, this requirement appears to be necessary.



**Fig. 5 The crude self-actuating model of embodied prospective foraging**

Clone indicates a process where some separation is made between simulated and actual events, such that the animal can simulate action-event sequences to inform decision making. This need not be a literal copying. The merge process indicates the place where information from the simulation is used to update the state of the actual animal's beliefs. The p-self does not necessarily change its state during the simulation, but is used to guide appropriate updating in the merge process.

ssary given that some but not all of what happens in embodied prospective foraging should result in updates to the animal's behaviours and representations of its needs and environment.

Where this requirement is met, as in the crude self-actuating model, embodied prospective foraging entails two things. First, the animal, either simultaneously or in close temporal proximity, activates multiple coherent sequences of points of view. Some of these sequences are q-selves, which are associated with simulated actions and experiences. But one point of view is anchored in reality in a way that the others are not: it specifies the actual animal. This matters: the animal's actually eating is linked to its survival in ways that its simulated eating is not. Second, in updating its behaviours and representations of its needs and environment, the animal distinguishes one of these sequences of points of view as its actual point of view from those that are merely q-selves.

In treating the sequences of points of view that are merely q-selves differently from the way it treats the sequence of points of view associated with the actual animal, the animal is manifesting a precursor of self-awareness. The animal is treating one of the sequences of points of view as itself: distinguishing the points of view associated with actual actions and actual consequences from other points of view just is distinguishing the self from things that are not the self. This gives rise to a primal notion of a self or self-like entity, which we shall call the *p-self*. The p-self is created when the animal determines which sequence of points of view should be treated as the actual animal.

The existence of the crude self-actuating model of embodied prospective foraging as a hypothetical possibility offers additional insight into what is meant by the p-self. The p-self is unlike the notion of self often thought to be necessary for mental time travel (Suddendorf and Corballis, 2007), nor is it a pre-reflective account of the self often proposed to lie at the heart of first person experience (Block, 1995; Ghallager, 2000; Metzinger, 2004; Panksepp, 1998; Vandekerckhove and Panksepp, 2009; see also Tagini and Raffone, 2010). In particular, the p-self is temporally short-lived relative to the animal's lifespan, because it is only required during activation of q-selves. The p-selves created across internal foraging events separated in time need not be cognitively represented as the same p-self. Invoking a p-self therefore does not require that the animal suddenly envision itself as having a life.

Second, the p-self is in some minimal sense a bearer

of subjective rewards or a subject of experiences. This follows from the fact that, as noted earlier, effectively merging the animal's actual and simulated behaviors cannot simply be a matter of copying information associated with the q-self back onto the p-self, for this would produce errors in the animal's representation of the state of the surrounding world. The merging operation therefore involves distinguishing either what the q-self experiences from what the p-self experiences or the q-self's rewards from the p-self's rewards. The p-self—in that it has short-lived persistence, a point of view, and mental states (in some minimal sense)—provides the primal version of the self, Self 0.1.

This distinction points to a further interesting problem in identifying what the q-self actually returns to the animal, via the merge process. Numerous cognitive models of decision making hypothesize that exploratory processes return values that summarize the world along some set of dimensions (e.g., risk and reward, Hau et al., 2008; Tversky and Shafir, 1992). For cognitive level explanations in the laboratory, this is perfectly plausible and is a powerful predictive tool. However, natural environments vary along countless dimensions (e.g., effort required, predator presence, kinds of predators, escape routes, kinds of resources, shortcuts to other resources, number and quality of competitors, and so on) and may require more sophisticated expectations. A merge process that returned a single value would not prepare an animal for possible action-event outcomes in ways that a more richly valued system would—nor would it appear to fit with the evidence. For example, imagined actions can influence mental and physiological states such as anxiety and heart rate (Taylor et al., 1998; Decety, 1996). However, taken to its extreme, a merge process that included all relevant information would not be any different than a process that returns the simulation. Here, we cannot offer a solution to this problem, but simply note that the results of the merge process determine what the q-self returns to the animal. It cannot simply replace the p-self—as it would in the naive inhibition model—but it may also be maladaptive to reduce the q-self experiences to a single value.

Embodied prospective foraging and the crude self-actuating model offer an explanation of how foraging leads to the requirement for a primal self, thus potentially providing clues to the origins of auto-noetic consciousness or self-awareness. Moreover, consistent with the evidence provided in previous sections, embodied prospective foraging allows the animal to look ahead of its current position, to navigate over past experiences

and potentially chain these together in novel ways, and to deliberate over possible outcomes in a way that offers behavioral flexibility. In addition, the p-self/q-self distinction allows the animal to do this in a way that does not confuse imagined experience with real experience (e.g., see Hassabis et al., 2007). If the q-self finds and eats all the strawberries, the p-self allows the animal to forage as if the strawberries are still there.

Since the p-self has minimal cognitive requirements, it is reasonable to expect many kinds of agents, including artificial agents such as robots, could have a p-self (see Holland and Goodman, 2003). The core ingredients are an ability for embodied prospective foraging and an ability, in making decisions based on such foraging, to distinguish sequences of points of view which are merely q-selves from its actual points of view. Any cognitive system with these properties has a p-self, whatever the mechanism used to make the distinction.

Finally, we note that the notion of self produced during embodied prospective foraging solves a problem in relation to the adaptive function of the self. In particular, if we are willing to entertain the idea that the earliest versions of self-awareness, however humble, arise out of a process of embodied prospective foraging, then this suggests that *the primal function of the self was to mark the distinction between actual and simulated action*.

## 6 Discussion

The aim of the present work has to been to explore the relationship between external foraging and internal foraging, and the potential consequences of the latter for understanding the evolution of auto-noetic consciousness or self-awareness. Notably, we have arrived at something like a self, the *p-self*, which is not derived from many of the elegant and rich accounts of consciousness (e.g., Metzinger, 2004; Markowitsch and Staniloiu, 2011; Panksepp, 1998). Rather, the p-self solves a particular kind of problem created by cognitive systems capable of internal foraging, and specifically embodied prospective foraging.

The p-self as a consequence of embodied prospective foraging is also consistent with evidence that damage to the hippocampus leads to problems in both future thinking and self representation (Hassabis et al., 2007; Corkin, 2002). Individuals with hippocampal damage appear to show degraded auto-noetic consciousness, indicating that some features of neural processing associated with embodied prospective foraging are potentially also features required for self-awareness. This is particularly telling because long-term memories do not suffer from hippo-

campal damage. Thus, auto-noetic consciousness might be no more than an evolutionary consequence of embodied prospective foraging: whatever mechanisms enable animals to distinguish actual from simulated actions in prospective foraging may also be needed for auto-noetic consciousness.

The present work also raises a number of questions. First, to say that an animal is capable of embodied prospective foraging implies that it can recall past memories sufficient to guide (simulated) actions, which suggests that the recalled information may specify a point of view. Progress in determining whether, for various animals, recalled information can specify a point of view may be made by measuring their abilities in perspective taking tasks (e.g. Vauclair et al., 1993; Mauck and Denhardt, 1997; Burmann et al, 2005). Though we see the reactivation of points of view as parsimonious with past research, the degree to which this is true is an open question. Animals could encode information in ways that are degraded or symbolic. However, we do not think this would alter the need for a p-self. Even if embodied prospective foraging were entirely symbolic, if the symbolic updating that occurs during internal foraging alters the animal's correct representation of its current state, then it must distinguish the real from the imagined. How it makes this distinction is not relevant to our argument, but it is nonetheless an open question. For example, *reality monitoring* is the capacity to distinguish internally generated information from externally derived information, and has been shown to involve comparisons between the quality of the information produced from the different sources (Johnson et al., 1993). Even though this system occasionally makes errors, it nonetheless solves the same problem that we have outlined here and therefore distinguishes the self in distinguishing the actual from the imagined. Indeed, the presence of occasional errors in real biological systems is evidence for the need to distinguish real from imagined actions: Their confusion (source memory errors) leads to false memories and hallucinations (Brainerd and Reyna, 2005; Mitchell and Johnson, 2009; Morrison, 2001).

The current work also asks when it is beneficial for an animal, or human, to engage in embodied prospective foraging, or internal foraging of any kind. Uncertainty is one kind of answer, but we may also ask in what kinds of environments is embodied prospective foraging adaptive, and why dispositions to engage in it differ, for example, across individuals. Moreover, we may be interested in further understanding how internal

representations lead to activation of q-selves—that is, what kinds of memories and experiences are most associated with embodied prospective foraging. Furthermore, it is interesting to ask how many q-selves are appropriate and what stopping rules animals use in shifting from internal foraging to actually acting. We hope that the present work offers a challenge to understand these questions as well as to understand what cognitive capacities may arise out of internal foraging more generally.

## References

- Atance CM, O'Neill DK, 2001. Episodic future thinking. *Trends in Cognitive Sciences* 5: 533–539.
- Austerweil JL, Abbott JT, Griffiths TL, 2012. Human memory search as a random walk in a semantic network. In: Pereira F, Burges CJC, Bottou L, Weinberger KQ ed. *Advances in Neural Information Processing Systems Proceedings of Neural Information Processing Systems 2012, Lake Tahoe, Nevada*.
- Bainton R, Tsai L, Singh C, Moore M, Neckameyer WS et al., 2000. Dopamine modulates acute responses to cocaine, nicotine and ethanol in *Drosophila*. *Current Biology* 10: 187–194.
- Banks A, Vincent J, Phalp K, 2008. Natural strategies for search. *Natural Computing* 8: 547–570.
- Barron AB, Søvik E, Cornish JL, 2010. The roles of dopamine and related compounds in reward-seeking behavior across animal phyla. *Frontiers in Behavioral Neuroscience* 4: 1–9.
- Bell WJ, 1991. *Searching Behaviour: The Behavioural Ecology of Finding Resources*. London: Chapman and Hall Ltd.
- Block N, 1995. How many concepts of consciousness? *Behavioral and Brain Sciences* 18: 272–287.
- Brainerd CJ, Reyna VF, 2005. *The Science of False Memory*. New York: Oxford University Press.
- Bromberg-Martin ES, Hikosaka O, 2009. Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63: 119–126.
- Buckner RL, 2010. The role of the hippocampus in prediction and imagination. *Annual Review of Psychology* 61: 27–48.
- Buckner RL, Carroll DC, 2007. Self-projection and the brain. *Trends in Cognitive Sciences* 11: 49–57.
- Burmann B, Dehnhardt G, Mauck B, 2005. Visual information processing in the lion-tailed macaque *Macaca Silenus*: Mental rotation or rotational invariance? *Brain, Behavior and Evolution* 65: 168–76.
- Carr MF, Jadhav SP, Frank LM, 2011. Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience* 14: 147–153.
- Chermahini SA, Hommel B, 2010. The (b) link between creativity and dopamine: Spontaneous eye blink rates predict and dissociate divergent and convergent thinking. *Cognition* 115: 458–465.
- Cheung A, Collett M, Colle TS, Dewar A, Dyer F et al., 2014. Still no convincing evidence for cognitive map use by honeybees. *Proceedings of the National Academy of Sciences* 111: E4396–E4397.
- Claverys JP, Prudhomme M, Martin B, 2006. Induction of competence regulons as a general response to stress in gram-positive bacteria. *Annual Review of Microbiology* 60: 451–475.
- Clayton NS, Russell J, 2009. Looking for episodic memory in animals and young children: Prospects for a new minimalism. *Neuropsychologia* 47: 2330–2340.
- Conway MA, 2005. Memory and the self. *Journal of Memory and Language* 53: 594–628.
- Conway MA, Pleydell-Pearce CW, 2000. The construction of autobiographical memories in the self-memory system. *Psychological Review* 107: 261.
- Conway MA, Pleydell-Pearce CW, Whitecross SE, Sharpe H, 2003. Neurophysiological correlates of memory for experienced and imagined events. *Neuropsychologia* 41: 334–340.
- Cools R, D'Esposito M, 2009. Dopaminergic modulation of flexible cognitive control in humans. In: Björklund A, Dunnett S, Iversen L, Iversen S eds. *Dopamine Handbook*. Oxford: Oxford University Press, 249–260.
- Corkin S, 2002. What's new with the amnesic patient HM? *Nature Reviews Neuroscience* 3: 153–160.
- Costa VD, Tran VL, Turchi J, Averbeck BB, 2014. Dopamine modulates novelty seeking behavior during decision making. *Behavioral Neuroscience* 28: 556.
- Cruse H, Wehner R, 2011. No need for a cognitive map: Decentralized memory for insect navigation. *PLoS Computational Biology* 7: e1002009.
- Dani JA, Zhou F-M, 2004. Selective dopamine filter of glutamate striatal afferents. *Neuron* 42: 522–524.
- Davidson TJ, Kloosterman F, Wilson MA, 2009. Hippocampal replay of extended experience. *Neuron* 63: 497–507.
- Daw ND, Doya K, 2006. The computational neurobiology of learning and reward. *Current Opinion in Neurobiology* 16: 199–204.
- Decety J, 1996. Do imagined and executed actions share the same neural substrate? *Cognitive Brain Research* 3: 87–93.
- Diba K, Buzsáki G, 2007. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience* 10: 1241–1242.
- Ego-Stengel V, Wilson M A, 2010. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus* 1: 1–10.
- Engle RW, Kane MJ, Tuoholski SW, 1999. Individual differences in working memory capacity and what they tell us about controlled attention, general fluid intelligence, and functions of the prefrontal cortex. In: Miyake A, Shah P eds. *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. New York, NY: Cambridge University Press, 102–134.
- Ferreira AS, Raposo EP, Viswanathan GM, da Luz MGE, 2012. The influence of the environment on Lévy random search efficiency: Fractality and memory effects. *Physica A: Statistical Mechanics and Its Applications* 391: 3234–3246.
- Fink J, Smith G, 1980. Mesolimbicocortical dopamine terminal fields are necessary for normal locomotor and investigatory exploration in rats. *Brain Research* 199: 359–385.
- Floresco S, Seamans J, Phillips A, 1996. A selective role for dopamine in the nucleus accumbens of the rat in random foraging but not delayed spatial win-shift-based foraging. *Behavioural Brain Research* 80: 161–168.
- Foster DJ, Wilson MA, 2006. Reverse replay of behavioural sequences in hippocampal place cells during the awake state.

- Nature 440: 680–683.
- Gallagher S, 2000. Philosophical conceptions of the self: Implications for cognitive science. *Trends in Cognitive Science* 4: 14–21.
- Galhardo RS, Hastings PJ, Rosenberg SM, 2007. Mutation as a stress response and the regulation of evolvability. *Critical Reviews in Biochemistry and Molecular Biology* 42: 399–435.
- Grünbaum D, 1998. Using spatially explicit models to characterize foraging performance in heterogeneous landscapes. *The American Naturalist* 151: 97–113.
- Gupta AS, van der Meer MAA, Touretzky DS, Redish AD, 2010. Hippocampal replay is not a simple function of experience. *Neuron* 65: 695–705.
- Harte J, Zillio T, Conlisk E, Smith AB, 2008. Maximum entropy and the state-variable approach to macroecology. *Ecology* 89: 2700–2711.
- Hassabis D, Kumaran D, Vann SD, Maguire EA, 2007. Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences* 104: 1726–1731.
- Hassabis D, Kumaran D, Maguire EA, 2007. Using imagination to understand the neural basis of episodic memory. *The Journal of Neuroscience* 27: 14365–14374.
- Hau R, Pleskac TJ, Kiefer J, Hertwig R, 2008. The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making* 21: 493–518.
- Hesslow G, 2002. Conscious thought as simulation of behaviour and perception. *Trends in Cognitive Sciences* 6: 242–247.
- Hesslow G, 2012. The current status of the simulation theory of cognition. *Brain Research* 1428: 71–79.
- Hills TT, 2006. Animal foraging and the evolution of goal-directed cognition. *Cognitive Science* 30: 3–41.
- Hills T, Dukas R, 2012. The evolution of cognitive search. In: Todd PM, Hills T, Robbins T eds. *Cognitive Search: Evolution, Algorithms, and the Brain*. Strüngmann Forum Reports, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press.
- Hills TT, Pachur T, 2012. Dynamic search and working memory in social recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38: 218–228.
- Hills TT, Jones MN, Todd PM, 2012. Optimal foraging in semantic memory. *Psychological Review* 119: 431–440.
- Hills TT, Kalf C, Wiener JM, 2013. Adaptive Lévy processes and area-restricted search in human foraging. *PLoS ONE* 8: e60488.
- Hills TT, Todd PM, Goldstone RL, 2008. Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science* 19: 802–808.
- Hills TT, Todd PM, Goldstone RL, 2010. The central executive as a search process: Priming exploration and exploitation across domains. *Journal of Experimental Psychology: General* 139: 590–609.
- Hills TT, Todd PM, Lazer D, Redish AM, Couzin I et al., 2015. Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences* 19: 46–54.
- Hills TT, Brockie PJ, Maricq AV, 2004. Dopamine and glutamate control area-restricted search behavior in *Caenorhabditis elegans*. *The Journal of Neuroscience* 24: 1217–1225.
- Hills TT, Mata R, Wilke A, Samanez-Larkin G, 2013. Mechanisms of age-related decline in memory search across the adult life span. *Developmental Psychology* 49: 2396–2404.
- Holland O, Goodman R, 2003. Robots with internal models: A route to machine consciousness? *Journal of Consciousness Studies* 10: 4–5.
- Hu D AN, Amsel A, 1995. A simple test of the vicarious trial-and-error hypothesis of hippocampal function. *Proceedings of the National Academy of Sciences* 92: 5506–5509.
- Janzen DH, 1970. Herbivores and the number of tree species in tropical forests. *The American Naturalist* 104: 501–528.
- Johnson A, Redish AD, 2007. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience* 27: 12176–12189.
- Johnson A, van der Meer MA, Redish AD, 2007. Integrating hippocampus and striatum in decision-making. *Current Opinion in Neurobiology* 17: 692–697.
- Kareiva P, Odell G, 1987. Swarms of predators exhibit “preytaxis” if individual predators use area-restricted search. *The American Naturalist* 130: 233–270.
- Klein SB, Loftus J, Kihlstrom JF, 2002. Memory and temporal experience: The effects of episodic memory loss on an amnesic patient’s ability to remember the past and imagine the future. *Social Cognition* 20: 353–379.
- Longstaffe KA, Hood BM, Gilchrist ID, 2014. The influence of cognitive load on spatial search performance. *Attention, Perception, and Psychophysics* 76: 49–63.
- Macnab RM, Koshland DE, 1972. The gradient-sensing mechanism in bacterial chemotaxis. *Proceedings of the National Academy of Sciences of the United States of America* 69: 2509–2512.
- Markowitsch HJ, Staniloiu A, 2011. Memory, autoegetic consciousness, and the self. *Consciousness and Cognition* 20: 16–39.
- Mauck B, Dehnhardt G, 1997. Mental rotation in a California sea lion *Zalophus Californianus*. *Journal of Experimental Biology* 200: 1309–1316.
- McNamara CG, Tejero-Cantero Á, Trouche S, Campo-Urriza N, Dupret D, 2014. Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neuroscience* 17: 1658–1660.
- Meister IG, Krings T, Foltys H, Boroojerdi B, Müller M et al., 2004. Playing piano in the mind: An fMRI study on music imagery and performance in pianists. *Cognitive Brain Research* 19: 219–228.
- Metzinger T, 2004. *Being no one: The self-model theory of subjectivity*. Cambridge, MA: MIT Press.
- Miller EK, Cohen JD, 2001. An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience* 24: 167–202.
- Mitchell KJ, Johnson MK, 2009. Source monitoring 15 years later: What have we learned from fMRI about the neural mechanisms of source memory? *Psychological bulletin* 135: 638.
- Morrison AP, 2001. The interpretation of intrusions in psychosis: An integrative cognitive approach to hallucinations and delusions. *Behavioural and Cognitive Psychotherapy* 29: 257–276.
- Moser EI, Kropff E, Moser MB, 2008. Place cells, grid cells, and the brain’s spatial representation system. *Neuroscience* 31: 69.
- Muenzinger KF, Fletcher FM, 1936. Motivation in learning: VI. Escape from electric shock compared with hunger-food tension in the visual discrimination task. *Journal of Comparative Psychology* 22: 79–91.

- Nicola SM, Surmeier J, Malenka RC, 2000. Dopaminergic modulation of neuronal excitability in the striatum and nucleus accumbens. *Annual Review of Neuroscience* 23: 185–215.
- Panksepp J, 1998. The periconscious substrates of consciousness: Affective states and the evolutionary origins of the self. *Journal of Consciousness Studies* 5: 566–82.
- Pennartz CMA, Ito R, Verschure PFMJ, Battaglia FP, Robbins TW, 2011. The hippocampal-striatal axis in learning, prediction and goal-directed behavior. *Trends in Neurosciences* 34: 548–559.
- Perner J, Kloos D, Stöttinger E, 2007. Introspection and remembering. *Synthese* 159: 253–270.
- Peters J, Büchel C, 2010. Episodic future thinking reduces reward delay discounting through an enhancement of prefrontal-midtemporal interactions. *Neuron* 66: 138–148.
- Pezzulo G, van der Meer MA, Lansink CS, Pennartz CM, 2014. Internally generated sequences in learning and executing goal-directed behavior. *Trends in Cognitive Sciences* 18: 647–657.
- Pfeiffer BE, Foster DJ, 2013. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497: 74–79.
- Plank M, James A, 2008. Optimal foraging: Lévy pattern or process? *Journal of the Royal Society Interface* 5: 1077–1086.
- Raaijmakers JGW, Shiffrin RM, 1981. Search of associative memory. *Psychological Review* 88: 93–134.
- Rosen VM, Engle RW, 1997. The role of working memory capacity in retrieval. *Journal of Experimental Psychology: General* 126: 211–227.
- Schacter DL, Addis DR, Buckner RL, 2007. Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience* 8: 657–661.
- Scoville WB, 1968. Amnesia after bilateral mesial temporal-lobe excision: Introduction to the case HM. *Neuropsychologia* 6: 211–213.
- Seamans J, Floresco S, 1998. D1 receptor modulation of hippocampal-prefrontal cortical circuits integrating spatial memory with executive functions in the rat. *The Journal of Neuroscience* 18: 1613–1621.
- Shepard RN, Metzler J, 1971. Mental rotation of three-dimensional objects. *Science* 171: 701–703.
- Shohamy D, Adcock RA, 2010. Dopamine and adaptive memory. *Trends in Cognitive Sciences* 14: 464–472.
- Skaggs WE, McNaughton BL, 1996. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science* 271: 1870–1873.
- Smith KA, Huber DE, Vul E, 2013. Multiply-constrained semantic search in the Remote Associates Test. *Cognition* 128: 64–75.
- Stephens DW, Krebs JR, 1987. *Foraging Theory*. Princeton: Princeton University Press.
- Suddendorf T, Corballis MC, 2007. The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and Brain Sciences* 30: 299–313.
- Tagini A, Raffone A, 2010. The ‘I’ and the ‘Me’ in self-referential awareness:  $\alpha$  neurocognitive hypothesis. *Cognitive Processing* 11: 9–20.
- Taylor SE, Pham LB, Rivkin ID, Armor DA, 1998. Harnessing the imagination: Mental simulation, self-regulation, and coping. *American Psychologist* 53: 429.
- Thompson GW, Kello CT, 2014. Walking across Wikipedia: A scale-free network model of semantic memory retrieval. *Frontiers in Psychology* 5: 86.
- Tolman EC, 1948. Cognitive maps in rats and men. *Psychological Review* 55: 189–208.
- Tolman EC, Gleitman H, 1949. Studies in learning and motivation: I. Equal reinforcements in both end-boxes, followed by shock in one end-box. *Journal of Experimental Psychology* 39: 810–819.
- Tsoar A, Nathan R, Bartan Y, Vyssotski A, Dell’Omo G et al., 2011. Large-scale navigational map in a mammal. *Proceedings of the National Academy of Sciences* 108: E718–E724.
- Tulving E, 1972. *Episodic and Semantic Memory*. Tulving E, Donaldson W ed. *Organization of Memory*. New York, London: Academic Press, 381–402.
- Tulving E, 1985. Memory and consciousness. *Canadian Psychology/Psychologie Canadienne* 26: 1–12.
- Tversky A, Shafir E, 1992. Choice under conflict: The dynamics of deferred decision. *Psychological Science* 3: 358–361.
- Unsworth N, Engle RW, 2007. On the division of short-term and working memory: An examination of simple and complex span and their relation to higher order abilities. *Psychological Bulletin* 133: 1038.
- Wehner R, Srinivasan MV, 1981. Searching behaviour of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera). *Journal of Comparative Physiology* 142: 315–338.
- Wilson MA, McNaughton BL, 1994. Reactivation of hippocampal ensemble memories during sleep. *Science* 265: 676–679.
- Winstanley CA, Robbins TW, Balleine BW, Brown JW, Büchel C et al., 2012. Search, goals, and the brain. In: Todd PM, Hills TT, Robbins TW eds. *Cognitive Search: Evolution, Algorithms, and the Brain*. Cambridge, MA: The MIT Press.
- Vandekerckhove M, Panksepp J, 2009. The flow of anoetic to noetic and auto-noetic consciousness: A vision of unknowing (anoetic) and knowing (noetic) consciousness in the remembrance of things past and imagined futures. *Consciousness and Cognition* 18: 1018–1028.
- Vauclair J, Fagot J, Hopkins WD, 1993. Rotation of mental images in baboons when the visual input is directed to the left cerebral hemisphere. *Psychological Science* 4: 99–103.