6

Search Processes and Hippocampus

A. David Redish

Abstract

Deliberation entails the sequential, serial search through possible options. This means that deliberation requires a mechanism to represent the structure of the world, from which predictions can be generated concerning these options and the expectations of the consequences of taking those options. Deliberation requires a mechanism to move mentally through those predictions as well as a mechanism to evaluate and compare those predictions. Neural signals for each of these factors have been found in the rat.

Introduction

The concept of the cognitive map introduced by Tolman (1938, 1939, 1948) fundamentally entails representations of the structure of the world. In fact, Tolman's original formulation of the "cognitive map" was more "cognitive" than "map." Tolman did not necessarily envision the cognitive map as spatial (for a discussion, see Johnson and Crowe 2009). Nevertheless, the translation of the cognitive map into modern neuroscience (primarily through O'Keefe and Nadel 1978) was fundamentally a spatial vision. Over the subsequent several decades, the concept of the cognitive map was more about whether the hippocampus encoded an actual "map" than about how the map was used (Nadel 1991; Eichenbaum et al. 1992; O'Keefe 1999; Eichenbaum 2000; Redish 2001). Examining the information and computational processes that the cognitive map would provide to the rat allowed an integration of the memory and spatial results into a unified theoretical picture (Redish 1999). In this chapter, I will return to the question of Tolman's original concept: in order to predict, one needs representations of the structure of the world.

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Place Cells and the Structure of the World

The spatial tuning of place cells is known to be derived from internal dead-reckoning representations: representations of spatial location and orientation maintained through self-motion information (Redish 1999). Those dead-reckoning systems appear to lie in the medial entorhinal cortex utilizing the grid-cell representations now known to lie therein (Fyhn et al. 2004; Hafting et al. 2005). These internal coordinate systems are then associated in hippocampus proper with external sensory signals, providing information about spatial position (Knierim et al. 1998; Redish 1999). The hippocampus can, in turn, use these learned associations between representations of external landmarks and internal representations of position to reset the internal coordinate system when the animal becomes lost (Redish and Touretzky 1997; Touretzky and Redish 1996) and to prevent drift during navigation (Samsonovich and McNaughton 1997; Redish et al. 2000). Although hippocampal lesions disrupt the spatial reliability of the grid-cell tuning, they do not seem to disrupt the internal coherence of grid cells (Bonnevie et al. 2010). In contrast, the internal coherence of grid cells depends on theta rhythmicity interactions arising from septal nucleus (Brandon et al. 2010; Koenig et al. 2010). This implies that the grid cells use septal signals to integrate dead-reckoning information and can continue to do so without hippocampus, but that the hippocampus is necessary to prevent drift.

In a variety of conditions, however, place cells also show reliable nonspatial tuning (Redish 1999; Pastalkova et al. 2008). Unlike spatial tuning, which appears to be ubiquitous in hippocampal representations, nonspatial tuning can appear, or not, in an environment depending on the specific path distribution (McNaughton et al. 1983; Muller et al. 1994), task and goal distribution (Markus et al. 1995; Olypher et al. 2002; Jackson and Redish 2007; Ainge et al. 2011), and training within that environment (Wood et al. 2000; Bower et al. 2005). This led several authors to conclude that the nonspatial information was represented as different maps (reference frames: Touretzky and Redish 1996; Redish and Touretzky 1997; Redish 1999; charts: McNaughton et al. 1996; Samsonovich and McNaughton 1997). A more nuanced description may be that the nonspatial representations depend, as proposed by Tolman in his original formulation of the cognitive map, on the structure of the world (Johnson and Crowe 2009).

Several recent experiments have found that hippocampal "place" cells will represent sequences of nonspatial information (Fortin et al. 2002; Agster et al. 2002) and distance (Pastalkova et al. 2008; Takahashi et al. 2009a; Gill et al. 2011) through a waiting period, as originally proposed by Levy (1996; see also Levy et al. 2005). Even on spatial experiments, cells can differentiate overlapping paths that originate from and proceed to different locations (Wood et al. 2000; Ferbinteanu and Shapiro 2003; Ainge et al. 2011). For example, Wood et al. (2000) found that when rats alternate on a simple T-choice with returns,

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

hippocampal place cells encode the two overlapping paths on the central track depending on whether the journey is a left-to-right or right-to-left occurrence. However, Bower et al. (2005) found that such a differentiation occurred only if the rat was initially trained with a separation between directions. Both Griffin et al. (2007) and Ainge et al. (2007) found that even when rats were trained identically to those in Wood et al. (2000), if a delay was imposed (thus making the task hippocampally dependent), the differentiation seen by Wood et al. (2000) vanished.

This difference likely depends on how animals are bridging the "gaps" in the task (Redish 2001). There are three ways to bridge a temporal or spatial gap: (a) priming, which is likely dependent on changes in sensory cortices, requiring the repetition of a cue for memory; (b) active maintenance of information and rehearsal, likely dependent on working memory and recurrent circuits in the prefrontal cortex; and (c) recall, dependent on storage and recall of episodic memory situations.

Open Questions:

- When do hippocampal place cells encode nonspatial information and when do they not?
- How does the structure of the world impact those representations?

Episodic Future Thinking and Episodic Past Thinking

Memory is only evolutionarily useful if it affects future actions. An important question, therefore, is: How does episodic memory affect future decision making? Four not necessarily mutually exclusive hypotheses include (a) recognition of individual past events for single-trial learning (Zilli and Hasselmo 2008; Lengyal and Dayan 2007), (b) training up other systems (Marr 1971; Alvarez and Squire 1994; Sutherland and McNaughton 2000), (c) reevaluation of the past (Loftus and Palmer 1974; Schacter 2001; Schacter and Addis 2007; Buckner and Carroll 2007), and (d) episodic future thinking (Buckner and Carroll 2007; Schacter et al. 2008).

The hippocampus has long been identified as a means of "bridging a gap" in tasks that require recognition of individual trials (Rawlins 1985; Redish 1999, 2001). For example, in T-maze alternation, the hippocampus is only necessary if a delay is imposed between the trials (Dember and Richman 1989; Ainge et al. 2007). In an explicit model of this, Zilli and Hasselmo (2008) showed that a one-trial learning memory capable of recognizing past individual events can bridge gaps in hippocampal-dependent tasks. Differentiating between recognition memory, in which one recognizes the familiarity of a situation, and recollection memory, in which one recalls an earlier situation and compares the memory with one's current observations, Eichenbaum et al. (2007) suggest that the hippocampus is necessary for recollection, but not recognition. In

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

spatial tasks, one might identify this as including a "self-localization" process in which the animal identifies the situation it is in and differentiates it from other similar situations (Redish and Touretzky 1997; Redish 1999; Fuhs and Touretzky 2007). In general, recollection will depend on a reconstruction process through which past experiences will need to be rebuilt for comparison with present circumstances (Buckner and Carroll 2007; Schacter et al. 2007).

Open Question: Does the phenomenon of self-localization that is seen in rodents correspond to this more general process in humans?

From almost the very beginning of the place-cell literature, it was noted that place cells fired outside their place fields during nonattentive rest states (Pavlides and Winson 1989; O'Keefe and Nadel 1978). From the first ensemble recordings, it was noted that the ensembles reactivated during sleep states after a task more than before a task (Wilson and McNaughton 1994; Kudrimoti et al. 1999). Bringing these results in line with observations of a limited length retrograde amnesia after hippocampal damage, Squire and colleagues (Squire 1987; Reed and Squire 1998; Squire and Alvarez 1995; cf. Nadel and Moscovitch 1997; Sutherland et al. 2001, 2011) suggest that the hippocampus might be training up other systems during off-line sleep states (Marr 1971; Alvarez and Squire 1994; Redish et al. 1998; Hoffmann and McNaughton 2002; Ji and Wilson 2007; Euston et al. 2007). Although there is good evidence that both the hippocampus and cortex replay representations during sleep that had been experienced during previous wake states (Pavlides and Winson 1989; Buzsáki 1989; Hasselmo 1993), it is still not clear how veridical those replays are. Although there is good evidence that hippocampus plays a role in bridging gaps, particularly contextual ones, it has been well-established that conscious human recall of past events is not veridical. Past memory is "constructed" (Loftus and Palmer 1974; Schacter 2001). The other two hypotheses for the role of hippocampal memory attempt to explain this lack of veridicality.

The hypothesis that the role of hippocampus is to reevaluate the past suggests that the primary effect of hippocampus on decision making will occur during off-line processes. Presumably, this reevaluation will occur during replay, which suggests that the replay events may include search processes (Samsonovich and Ascoli 2005). It has recently been established that replay is useful for learning (Jackson et al. 2009; Ego-Stengel and Wilson 2010; Wilson et al. 2010), as suggested by early computational models (Buzsáki 1989; Hasselmo 1993). There is evidence that the hippocampus may provide search-like off-line mechanisms during replay in awake states (Gupta et al. 2010; for additional discussion, see Derdikman and Moser 2010); however, whether these same search-like processes occur during sleep states is still unknown.

The three hypotheses above suggest that hippocampus plays no immediate (online) role in decision making, only a supportive (off-line) role. As discussed above, if one has a representation of the structure of the world (a cognitive map), then one could use it to "search through the future" to predict outcomes. This would be particularly useful for one-off critical decisions, like deciding

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

where to go to graduate school or which job to take. In practice, this would become an imagination of what a future would be like: it would be episodic future thinking (Buckner and Carroll 2007; Schacter et al. 2008). Humans with hippocampal damage are severely impaired in the ability to construct imagined situations, including potential future situations (Hassabis et al. 2007). When rats come to choice points, they search through (at least the immediately available) future options (Johnson and Redish 2007) and evaluate them (van der Meer and Redish 2009). A role for the hippocampus in episodic future thinking suggests an explanation for the lack of veridicality of past declarative memories: the brain is using the same episodic imagination process it evolved to imagine future representations to construct past memories (Buckner and Carroll 2007; Schacter and Addis 2007). In a sense, episodic memory may not be memory after all, but rather "episodic past thinking" reconstructed from partial memories stored in cortical systems.

Finally, it is important to remember that the hippocampal activity, at a moment in time, is (usually) a relatively accurate representation of the animal's location within a context. From hippocampal neural ensembles during behavior, it is possible to decode position to an accuracy of better than 1 cm (Wilson and McNaughton 1993) and to decode time within a gap to an accuracy of better than 0.5 second (Itskov et al. 2011). These signals could potentially be used to signal contextual information for conditioning (Holland and Bouton 1999; Rudy 2009) or navigation (Burgess et al. 1994; Redish 1999; Foster et al. 2000).

Of course, all of these hypotheses are potentially viable; they do not necessarily conflict with each other. To determine whether they are incompatible with each other or not, one must first address the question of how they would be implemented computationally in the brain. Computational models have shown that one can, for example, bridge gaps, self-localize spatially, and replay memories all within the same network without interference between them (e.g., Redish et al. 1998; Redish 1999).

Open Questions:

- To what extent does the phenomenon in which place cells represent other places and other experiences correspond to the phenomena of episodic future thinking and recollection (episodic past thinking) seen in humans?
- What role does the hippocampus play in decision making? Does it play an active role online, or only an off-line monitoring role?

Replay

Because place cells carry information about the spatial location of the rat, it is possible to decode spatial location from a population of cells (Wilson and

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

McNaughton 1993; Brown et al. 1998; Zhang et al. 1998). Because place cells also fire spikes outside of their place fields, even decoding attempts based on the tautology of taking both training and test sets from the same data set occasionally decode to a different location from that of the rat (Jensen and Lisman 2000). Because these representations are self-consistent (Johnson et al. 2008), we interpret these as representations of other locations or other times (Johnson et al. 2009).

Replay has historically been interpreted as being related to consolidation of memory from hippocampal (episodic) representations to cortical (semantic) representations (Wilson and McNaughton 1994; Alvarez and Squire 1994; Nadel and Moscovitch 1997; Hoffmann and McNaughton 2002; Euston et al. 2007). However, replay might also support training (Foster and Wilson 2006; Johnson and Redish 2005), exploration (Samsonovich and Ascoli 2005; O'Neill et al. 2008; Csicsvari et al. 2007), or planning (Diba and Buzsáki 2007; Johnson et al. 2008; Singer and Frank 2009).

Replay during sleep states has been reported to be primarily forward in its sequence (Skaggs and McNaughton 1996; Nádasdy et al. 1999), but studies of replay during awake states have found a more complex story emerging. For example, replay during awake states can be reversed (Foster and Wilson 2006), even of remote locations not recently experienced (Davidson et al. 2009; Gupta et al. 2010), and even when animals never experienced that backward sequence (Gupta et al. 2010). Gupta et al. (2010) discovered that one can even find sequences played out during awake sharp waves that the rat has never experienced in either a forward or a backward direction, thus supporting the possibility that the hippocampus is searching through potential paths in the environment (Samsonovich and Ascoli 2005). These these shortcuts were a key component of the original cognitive map proposal (Tolman 1948; O'Keefe and Nadel 1978; Redish 1999).

Although early studies suggested that the amount of time experienced within an environment drove the amount of replay (Kudrimoti et al. 1999; Jackson et al. 2006; O'Neill et al. 2006), more recent studies, which looked at the specifics of what is replayed, find a more complex story. Looking directly at the information played out during awake sharp waves, Gupta et al. (2010) found that it was not the most recent experience that was being played out. On a figure-eight, T-choice maze, Gupta et al. (2010) ran a task that included three reward contingencies: selecting the left side for reward, selecting the right side for reward, or alternating sides for reward. During the critical probe days, the contingency switched halfway through the 40-minute session. In these three reward conditions, the recency with which rats experience the two sides can vary greatly. Gupta et al. (2010) found that rats replayed remote locations (the opposite side of the maze) more often when they had not been recently experienced. This suggests that replay during awake states may also serve to counteract the effect of repeated experiences which could overemphasize certain parts of the map because they were more recently experienced (Gupta et al. 2010;

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Derdikman and Moser 2010). This would be critical to maintaining an accurate representation of the structure of the world. It is also possible that nonhistorical, but spatially valid replay (e.g., as seen by Gupta et al. 2010) could aid in training a more generalized semantic memory and may explain why semantic memory is less autobiographically concrete than episodic memory (Tulving 1983).

Sweeps and Phase Precession

When animals are running, the hippocampal local field potential shows a strong 7 Hz rhythm called “theta” (Vanderwolf 1971; O’Keefe and Nadel 1978). The phase of the spike fired by a given place cell relative to this theta rhythm precesses from late in the theta cycle, when the animal first enters the place field, to earlier and earlier in the cycle, as the animal runs through the place field (Maurer and McNaughton 2007). This is most cleanly seen on linear tracks, where place fields are directional and a simple plot of the phase versus position of spikes fired by a given cell will show a definitive precession (O’Keefe and Recce 1993; Skaggs et al. 1996). However, it is also seen in other tasks in which rats reliably run through place fields, including on open fields (Skaggs et al. 1996; Huxter et al. 2008). It can even be seen during pauses when hippocampal firing divides up those pause times (Pastalkova et al. 2008; Takahashi et al. 2009a; Macdonald et al. 2010).

It is important to recognize that there are two ways to view “phase precession”: as a change in the phase of the cell as an animal runs through the field or as a sequence of firing of cells within a single theta cycle. Although the early studies of phase precession recognized this duality (Skaggs et al. 1996), it was often thought that the purpose of the internally generated phase precession was to construct this internal sequence to allow learning of asymmetric connections for replay (Skaggs et al. 1996; Mehta et al. 1997; Redish and Touretzky 1998). Alternative theories proposed that asymmetric connections drove the internal sequence, producing phase precession (e.g., Tsodyks et al. 1996; Jensen and Lisman 1996). However, an alternate hypothesis is that the sequence of firing is primary and phase precession is an epiphenomenon of a hippocampal generation of the sequence within the theta cycle combined with progress toward a goal (Lisman and Redish 2009). Support for this hypothesis comes from evidence that the sequence of firing within a theta cycle is more stable than the phase of a given cell’s firing (Dragoi and Buzsáki 2006).

Phase precession was first seen on linear tracks where place fields are directional. Skaggs et al. (1996) noted that the phase of spiking provides additional information capable of subdividing a place field (Jensen and Lisman 2000). This means that if phase precession could be seen in nondirectional cells, three observations were possible:

From “Cognitive Search: Evolution, Algorithms, and the Brain,” edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

1. The phase could match at the start of the place field, indicating that cells reflected past history.
2. The phase could match at the peak of the place field, indicating that phase simply reflected firing rate.
3. Or, the phase could match at the end of the place field, indicating that the cells reflected a prediction of the future.

Studies of bidirectional cells on cue-rich linear tracks (Battaglia et al. 2004) and in two-dimensional tasks (Huxter et al. 2008) found definitively that place fields show the third condition, indicating that place fields are representing prediction toward a goal (Lisman and Redish 2009).

Further support for this hypothesis has come from examinations of hippocampal "place" cells during nonspatial running. Hirase et al. (1999) had rats run on a running wheel within a given cell's place field. They found that cells did not show phase precession; instead, cells fired at constant phase. Pastalkova et al. (2008) trained rats to run on a running wheel for a given time and found (as predicted by Levy et al. 2005) that cells divided up the time on the running wheel. In contrast to Hirase et al. (1999), Pastalkova et al. (2008) found that their cells showed phase precession—the difference is that Pastalkova et al.'s rats had a goal toward which they were running.

When rats come to difficult decision points on spatial tracks, they pause and look back and forth. This behavioral observation has been termed "vicarious trial and error" (Muenzinger and Gentry 1931; Muenzinger 1938; Tolman 1938). During these attentive-pausing behaviors, rats remain in theta (Vanderwolf 1971). Decoding neural ensembles during these decision processes revealed theta phase-coupled sweeps of representation far ahead of the animal (Johnson and Redish 2007). Like phase precession, these sweeps were initiated from the location of the animal or slightly behind it, and they proceeded ahead of the animal within a single theta cycle. However, the sweeps proceeded much farther ahead than the sequential firing within a theta cycle typically seen on simple tasks. In addition, sweeps occurred first in one direction and then in the other, changing direction completely on each theta cycle.

These far-reaching sweeps occurred only on passes through the decision point during which animals showed vicarious trial and error. On the tasks used by Johnson and Redish (2007), animals eventually learned to run through the decision point without stopping, having made their decision earlier or having transferred the decision-making processes into a different nonhippocampal system (Schmitzer-Torbert and Redish 2002; van der Meer et al. 2010). As the behavior changed, decoded hippocampal representations first swept in both directions, then in only one direction; then they became indistinguishable from phase precession, only going a short distance ahead of the rat (Johnson and Redish 2007). The suggestion that phase precession is actually an epiphenomenon of the within-theta cycle sequence (Dragoi and Buzsáki 2006) and progress toward a goal (Lisman and Redish 2009) suggests that phase precession

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

and the sweeps of representation seen by Johnson and Redish (2007) may reflect a single process of prediction.

Open Questions:

- Is phase precession a special case of sweeps?
- What is the relationship between preplay before movement that is reported as occurring during sharp waves (e.g., Diba and Buzsáki 2007; Singer and Frank 2009) and sweeps that are reported as occurring during theta (e.g., Johnson and Redish 2007)?

Evaluation and Action Selection

The purpose of cognitive search is to identify the best action to take in a given situation. Thus, when performing a cognitive search, it is not enough to predict the future, one must also evaluate those future expectations to determine their relative value to one's goals and then select between them. These evaluation processes have been suggested to exist within the ventral striatum or nucleus accumbens. Even as far back as the 1980s, it was suggested that the ventral striatum may serve as the "interface between the limbic system and actions" (Mogenson et al. 1980). In part, this was because manipulations of ventral striatum were known to affect actions, but it was not clear whether they were affecting the action-selection process itself (Mogenson 1984) or the evaluation process (Berridge 1996, 2007), which would lead to action-selection changes (Atallah et al. 2007).

The hippocampal system projects to the ventral striatum through the CA1 and subiculum regions (Groenewegen et al. 1987; Voorn et al. 2004). Functionally, hippocampal firing and ventral striatal firing are correlationally coupled (Martin 2001). For example, sharp waves in the hippocampus (during which replay occurs) precede the firing of ventral striatal cells (Pennartz et al. 2004). In fact, the reactivated sequence in the hippocampus leads to specific reactivated reward-related firing in the ventral striatum (Lansink et al. 2008, 2009).

Ventral striatal medium spiny neurons show a variety of responses on tasks. Some cells show "ramps" or "anticipatory" activity, increasing their activity as the animal approaches a reward, usually spatially (Lavoie and Mizumori 1994; Miyazaki et al. 1998; Mulder et al. 1999; van der Meer and Redish 2009, 2011; but also temporally, Carelli and Deadwyler 1994; Nicola et al. 2004a). Other cells show responses that reflect value-related motor actions (Mulder et al. 2004; Roesch et al. 2009; Nicola 2010) and cues that signal impending reward (Carelli and Deadwyler 1994; Nicola et al. 2004a; Roitman et al. 2005; van der Meer and Redish 2009). Finally, some cells show firing in response to reward-receipt (Carelli and Deadwyler 1994; Lavoie and Mizumori 1994; Miyazaki et al. 1998; Nicola et al. 2004b; van der Meer and Redish 2009).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

At task points in which sweeps are known to occur, ventral striatal reward-related cells show excess activity, such that a decoding algorithm applied to the data decodes to times of reward-receipt rather than to the location of the animal (Carelli and Deadwyler 1994; Lavoie and Mizumori 1994; Miyazaki et al. 1998; Nicola et al. 2004b; van der Meer and Redish 2009). Although not explicitly studied this way, reward-related cells fire just before movement toward a goal (German and Fields 2007; Nicola 2010), when hippocampal ensembles show a pre-play representation of expected future paths (Diba and Buzsáki 2007; Singer and Frank 2009). The obvious hypothesis is that these cells represent a covert expectation or evaluation of reward (Johnson et al. 2009; van der Meer and Redish 2009, 2010). Importantly, they occur before an animal turns around when correcting the final decision in a vicarious trial and error (VTE) event (van der Meer and Redish 2009).

Recently, we found that ventral striatal anticipatory "ramp" cells phase precess relative to the hippocampal theta rhythm (van der Meer and Redish 2011). It has long been suggested that these cells could play a role related to value representations (Daw 2003), because they increase activity as they approach the goal; upon reaching the goal, firing drops dramatically. In the cognitive map literature, these predicted cell types have been referred to as "goal" cells because they encode distance to a goal (Burgess et al. 1993). If phase precession in hippocampus actually reflects the combination of a sweep-like sequence within a given theta cycle and progress toward a goal, then phase of a ventral striatal distance-to-goal or value-of-the-current-situation cell may reflect the evaluative step of this sweep-like sequence.

Other structures have also been identified as being involved in search, planning, and evaluation, including both the prefrontal cortex (Jones and Wilson 2005; Hyman et al. 2010; Peters and Büchel 2010; DeVito and Eichenbaum 2011) and the orbitofrontal cortex (Gallagher et al. 1999; Tremblay and Schultz 1999; Padoa-Schioppa and Assad 2006). In particular, the hippocampal-prefrontal interaction coupling identified by cellular and local field potential interactions improves during successful search-based tasks, but not during task failures (Jones and Wilson 2005; Hyman et al. 2010). Representations preceding expected outcomes in the orbitofrontal cortex depend on hippocampal integrity (Ramus et al. 2007).

We have recently found that orbitofrontal cortex reward-related neurons also show excess activity during the same VTE events as hippocampal sweeps and ventral striatal covert-reward activity (Steiner and Redish 2010). These orbitofrontal representations, however, occur after the turn around when correcting the final decision in a VTE event, which suggests that the orbitofrontal cortex is not part of the evaluation step in the decision-making process. This is consistent with recently hypothesized roles of the orbitofrontal cortex in signaling information about expectancies more than evaluation (Schoenbaum and Roesch 2005; Murray et al. 2007; Takahashi et al. 2010; Wilson et al. 2010; McDannald et al. 2011).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Open Questions:

- What role does ventral striatum play in action selection? Is it only evaluative? Or does it include action-selection components?
- What is the function of phase precession in ventral striatal ramp cells?
- How does the relationship between the hippocampus and the prefrontal cortex change during cognitive search-based processes?
- What role does the orbitofrontal cortex play in action selection? Is it evaluative or does it only encode expectations?

Automaticity

In situations with repeated, reliable reward contingencies, rats automate behaviors, switching from deliberative, flexible map-based decision systems to habit-based, inflexible, situation-response decision systems (Restle 1957; O'Keefe and Nadel 1978; Packard and McGaugh 1996; Balleine and Dickinson 1998a; Killcross and Coutureau 2003; Redish et al. 2008). Although the mechanism of these habit-based decision systems is beyond the scope of this chapter, theoretical suggestions differentiate deliberative from habit-based decisions by their search processes (Daw et al. 2005) and their representations of future expected outcomes (Balleine and Dickinson 1998a; Niv et al. 2006; van der Meer and Redish 2010).

These theories suggest that the flexible, map-based decision system includes a model of the state-transition structure of the world (Daw et al. 2005), allowing a prediction of the outcomes of actions (Daw et al. 2005; Johnson et al. 2007) and an online evaluation of the outcomes (Balleine and Dickinson 1998a; Niv et al. 2006; van der Meer and Redish 2010), whereas the habit-based processes simply associate a value with a situation-action pair (Daw et al. 2005; Sutton and Barto 1998). Unfortunately, the names that seem to have stuck are "model-based" and "model-free" because the former requires knowledge of the transition structure of the world (Daw et al. 2005; Niv et al. 2006). However, the latter also requires categorizing situations, entailing some knowledge of the state structure of the world (Redish and Johnson 2007; Gershman et al. 2010).

These theories suggest that deliberation-based systems will learn to recognize situation-situation and situation-action-situation transitions and use those to provide a "cognitive map," consistent with hippocampal learning on hippocampal-dependent tasks (Hirsh et al. 1978; O'Keefe and Nadel 1978; Tse et al. 2007) and the increased hippocampal-prefrontal coupling on these tasks (Jones and Wilson 2005; Hyman et al. 2010; Peters and Büchel 2010; DeVito and Eichenbaum 2011). This "cognitive map" learning will be ubiquitous, and cells will show reliable contextual (spatial) representations on any task, regardless of whether the task is hippocampally dependent or not. Consistent with

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

this hypothesis, place cells are observed on both hippocampally dependent and independent tasks (for a review, see Redish 1999).

In contrast, habit-based systems will learn to associate stimuli (situation/state representations) with actions only in contexts in which those stimuli (situation/state representations) reliably produce reward. This is the whole point of reinforcement learning, which is solving the credit-assignment problem (Sutton and Barto 1998). Lesion studies suggest that the dorsal striatum is likely a key structure in the habit learning system (Packard and McGaugh 1996; Graybiel 1990, 1998). Consistent with these hypotheses, dorsal striatal cells develop representations of key components of the task (Jog et al. 1999; Barnes et al. 2005, 2011; Thorn et al. 2010; van der Meer et al. 2010). These cells only represent information when that information is reliably rewarded (Schmitzer-Torbert and Redish 2008; Berke and Eichenbaum 2009).

Upon explicit examination of dorsal striatal neural ensembles on the same task in which hippocampal sweeps and ventral striatal covert representations of reward were seen, van der Meer et al. (2010) found that dorsal striatal ensembles showed neither effect. Even though dorsal striatum eventually developed representations of spatial location more reliably than the hippocampus on this task (for this task, space carries information about appropriate action selection; Schmitzer-Torbert and Redish 2004, 2008), those dorsal striatal representations never represented the future over the past. Similarly, although dorsal striatum contained reward-related cells (Schmitzer-Torbert and Redish 2004), those reward-related cells never showed covert reactivations. These data strongly support the view that the difference between the deliberative ("model-based") and habit-based ("model-free") systems is the presence of search and expectancy processes in the deliberative system and the lack of such processes in the habit-based system.

Open Questions:

- What controls which system to drive behavior?
- Is this related to the role of the prefrontal cortex in deliberation and evaluation?

Computational Pathologies

The fact that the decision-making process is a mechanical and algorithmic process (even if a complex one) implies that there are potential vulnerabilities or failure modes that can occur within the process which will lead to mismade decisions (Redish et al. 2008). Several potential vulnerabilities can be identified within the deliberative search-and-evaluate process discussed above. The first two errors reviewed below entail errors in the results of the search. The final two errors entail errors in the process itself.

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Errors in Expectations

The simplest error in prediction-based decision-making systems is that the system may predict the wrong outcome, either through misrecognition of situations, through mislearning of outcomes, or through misinformation. Gamblers, for example, are often cited as falling victim to the illusion of control, in which they believe that their own actions can control random effects (Wagenaar 1988; Langer and Roth 1975; Ladouceur and Sévigny 2005). Similarly, a person who believes in the positive effects of alcohol is more likely to drink than someone with a negative belief, independent of the actual effect of alcohol on the subject in question (Goldman et al. 1987; Jones et al. 2001).

Misevaluation

Even if the expectation leads to the correct answer, the evaluation process may misvalue the outcome. In the deliberative system, valuation is a dynamic process that depends on a combination of motivational, prediction, and memory processes (Balleine 2001; Niv et al. 2006). Several addiction theories are based on the misevaluation of expected outcomes (e.g., Robinson and Berridge 2001, 2003). These valuation processes depend on emotional processes (Damasio 1994) and lead to observable irrationalities (Andrade and Ariely 2009). The valuation process itself is a complex calculation, depending on memory (Balleine 2001), set points (Kahneman et al. 1982; Koob and Le Moal 2006), differences in risk sensitivity to gains and losses (Kahneman and Tversky 1979; Glimcher et al. 2008), and framing effects (Kahneman and Tversky 2000), including incompatibilities in valuation as a function of how these values are measured (Ahmed 2010). For example, rats will work harder for self-administered cocaine or for heroin than for sweetened water, but will prefer sweetened water when given the choice (Cantin et al. 2009, 2010). Similarly, addicts are highly inelastic when faced with small increases in drug costs (Carroll 1993; Grossman and Chaloupka 1998; Bickel and Marsch 2001), but will remain drug-free for very small, but tangible rewards (Higgins et al. 2002).

Obsession

The hypothesis that deliberation entails an actual search through potential future possibilities opens up the possibility that the search may repeat itself. The search process is a memory process (Johnson et al. 2007), and thus retrieving a potential path through the structure of the world entails recall and reconstruction of past episodic memories (Buckner and Carroll 2007). If the representation of the structure of the world is not balanced, the agent may be more likely to retrieve one potential path over others. A memory process that repeatedly retrieves a single path through potential futures may be clinically identifiable as a form of obsession (Redish and Johnson 2007).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Craving

The hypothesis that deliberation entails a representation of future outcomes for evaluation (Balleine and Dickinson 1998a; Niv et al. 2006; van der Meer and Redish 2010) implies a potential model for craving (Redish and Johnson 2007). Craving is an explicit, intense desire for a specific thing (Halikas 1997; Tiffany 1999). This implies that craving must include an expectation of that specific thing, whether it be as a goal to be achieved (Tiffany 1990) or an identification of a potential path to that thing (Goldman et al. 1987). Craving should not appear in habit-based relapse (Tiffany and Wray 2009; Redish 2009), where paths to drug use are reliable and often nonconscious (Tiffany 1990; Robbins and Everitt 1999; Altman et al. 1996; Sayette et al. 2000; Oei and Baldwin 2002; Everitt and Robbins 2005; Dickerson and O'Connor 2006; Redish et al. 2008). As one example, craving appears in alcoholics only when the path to a goal is thwarted (Sinha and O'Malley 1999; Addolorato et al. 2005), presumably leading to a switch from habit-based to deliberative systems (Redish et al. 2008).

Open Questions:

- How repetitive can search be? Is this related to obsession?
- How is the evaluation actually accomplished? Can we explain the irrationalities mechanistically?
- How do the search and evaluation processes interact? How do these interactions change in pharmacological and behavioral addictions?

Summary

The hippocampus represents the spatial and contextual information necessary for decision making (O'Keefe and Nadel 1978; Cohen and Eichenbaum 1993; Redish 1999). In particular, it is critical for the successful integration of those cues in terms of the construction of future expectations (Hassabis et al. 2007; Buckner and Carroll 2007), presumably due to its auto-associative properties (McNaughton and Nadel 1990). Under the hypothesis that decision making is separable into search-based processes and automated or cached processes, the hippocampus is implicated in search-based processes, particularly in the construction of future expectations (Johnson and Redish 2007), and can be contrasted with dorsal striatum, which is implicated in the development of look-up tables for cached-action (non-search-based) decisions (van der Meer et al. 2010). Open questions remain, however, as to the specific role played by the hippocampus in active (search-based) decision making, the hippocampal relation to structures usually associated with evaluation process (such as ventral striatum and orbitofrontal cortex), and the role played by the hippocampus in clinical search failures (such as errors in expectations and craving).

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

Acknowledgments

This work was supported by NIH grants R01 MH080318 and DA024080.

From "Cognitive Search: Evolution, Algorithms, and the Brain," edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.

From “Cognitive Search: Evolution, Algorithms, and the Brain,” edited by Peter M. Todd, Thomas T. Hills, and Trevor W. Robbins. 2012. Strüngmann Forum Report, vol. 9, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01809-8.