

Navigating the Interface Between Learning and Cognition

Jonathon D. Crystal

Indiana University, Bloomington, U.S.A.

The interface of learning and cognition applied to the study of animal behavior represents a target for significant progress if conceptual barriers can be reduced. Is animal behavior exclusively a product of learning or cognition, or are both implicated? Are special (i.e., new) methods required to study cognition or will the enterprise be accomplished by using well-established methods from learning? What types of hypotheses need to be tested to dissociate cognition from learning, and may these hypotheses be profitably tested? This article addresses the above questions by focusing on conceptual, methodological, and hypothesis-testing perspectives for navigating the interface between learning and cognition. Examples from contemporary research are used to develop some suggestions for best practices. The development of a rodent model of episodic memory is used as a case study to feature the validation of an animal model of cognition.

Significant progress in navigating the interface of learning and cognition in animals may be made if conceptual barriers can be reduced. This article explores the role of learning and cognition in animal behavior. Although special (i.e., new) methods may be developed to study cognition, the enterprise may also be accomplished by using well-established methods from learning. Hypotheses that need to be tested to dissociate cognition from learning and approaches for profitably testing these hypotheses are outlined. This article addresses the above issues by focusing on conceptual, methodological, and hypothesis-testing perspectives for navigating the interface between learning and cognition. Some suggestions for best practices are proposed based on examples from contemporary research. A case study (the development of a rodent model of episodic memory) is used to feature the validation of an animal model of cognition.

Conceptual Issues

At the conceptual level, we may ask whether animal behavior is primarily a product of low-level mechanisms or if there is a need to propose other, higher-level mechanisms. At the interface between learning and cognition, learning may be regarded as a lower-level mechanism and cognition may be regarded as a higher-level mechanism. For example, there is a long history of investigation of basic principles of learning in the domains of habituation and association. If these well-established principles of learning explain a particular set of data, then it would not be advisable to propose a cognitive explanation for the data (Shettleworth, 1998, 2010). This application of Morgan's canon (Morgan, 1906) is particularly apt given the array of conceptually complex proposals in contemporary investigations of comparative cognition. Some examples include: metacognition (e.g., Terrace & Son, 2009), episodic memory (e.g., Crystal, 2010), future planning (e.g., Raby & Clayton, 2009), theory of mind (e.g., Call & Tomasello, 2008), and

This work was supported by National Institute of Mental Health grant R01MH080052 to JDC. Correspondence concerning this article should be addressed to Jonathan D. Crystal, Department of Psychological and Brain Sciences, Indiana University, Bloomington, IN, 47405, U.S.A. (jcrystal@indiana.edu).

causal reasoning (e.g., Waldmann, Hagmayer, & Blaisdell, 2006) amongst many others. To take one of the above examples, metacognition is thinking about thinking, and a primary interest in developing animal models of metacognition is to provide a window into examination of awareness or consciousness (Terrace & Metcalfe, 2005). Clearly, the constructs listed in the previous statement are more complex than those which may be listed in a description of associative learning. Hence, this is a dramatic example of the interface of learning and cognition. The existence of cognitive processes in the domain of metacognition may provide fundamental insights about the evolution of cognition but only if a strict standard or criterion can be met, whereby alternative, lower-level explanations may be ruled out. We will return to the case of comparative metacognition in discussing methodological issues below.

At a fundamental level, the decision to assert a learning or cognitive explanation is an empirical issue that will be answered by data. However, it is worth noting that the *quantity* of data that supports a low-level mechanism vs. the quantity that appears to call for a higher-level mechanism is not always helpful for addressing this conceptual issue. By contrast, the issue for higher-level mechanisms is whether there is *any* evidence to require the proposal of higher-level mechanisms.

The types of experiments conducted ultimately determine the quantity of data that lines up in low-level or high-level columns. Thus, extensive programs of research that have been focused on low-level explanations for many years will likely tilt the quantity balance toward low-level mechanisms. If one explores the same topic by conducting many similar experiments, a massive quantity of data will be generated. However, new insights may be obtained by testing other hypotheses, which may ultimately reveal mechanisms that may challenge the relatively large body of research. An example of the quantity issue is the widely held belief that time perception is scalar, which has been taken as evidence for a pacemaker-accumulator mechanism (Gibbon, 1991). Although a scalar representation of time provides an excellent approximation of temporal performance, there are several lines of evidence that suggests that the representation is not scalar, which has been taken as evidence for an oscillator mechanism (e.g., Crystal, 2006a, 2006b, in press).

In this section, an empirical approach was advocated. The interface between learning and cognition will be navigated by conducting experiments that pit low- and high-level explanations against one another. The higher level explanations will only be adopted if some data cannot be explained by well-established lower-level explanations. What types of methods are required to undertake this empirical project?

Methodological Issues

Are new methods needed to dissociate cognition from learning? My perspective on this question is that although new methods may be developed, well-established methods are likely to be more valuable for navigating the interface between learning and cognition. A potential pitfall in the development of new methods will be introduced to advocate for the value of well-established methods.

Research on comparative metacognition has been focused on testing low- and high-level hypotheses from the outset of this research program (Smith, Shields, & Washburn, 2003). A persistent complaint by critics of comparative metacognition has been that perhaps principles of associative learning may explain the putative metacognition data. One approach to addressing this problem has been to develop behavioral testing methods that are outside of the boundaries of associative learning. The rationale for this approach is compelling – if the training protocol is outside the domain of associative learning, then the phenomenon under investigation can be explained by other, in this case cognitive, mechanisms. Although the rationale is compelling, it is important to critically evaluate the hypothesis that the new methods are indeed outside the boundaries of associative learning. The rationale for the new method is undermined if principles of learning may be applied despite the goals at the time that the new methods were developed.

Pervasiveness of Reinforcement in Comparative Metacognition

Experiments on uncertainty monitoring sometimes used direct reinforcement variables to influence the behavior of the subjects; for example, an uncertain response produced a hint or identification of the currently correct response (e.g., Smith, Shields, Allendoerfer, & Washburn, 1998), a guaranteed-reward trial (e.g., Shields, Smith, & Washburn, 1997; Smith, Shields, Schull, & Washburn, 1997), a time-out delay for over-use of the uncertainty response (e.g., Shields et al., 1997; Smith et al., 1997, 1998), or food (e.g., Foote & Crystal, 2007; Hampton, 2001).

Smith and colleagues (Smith, Beran, Couchman, & Coutinho, 2008) argued that a history of reinforcement associated with the uncertainty response is responsible for the applicability of low-level alternative explanations. This observation has led to some creative attempts to circumvent the role of reinforcement (Beran, Smith, Redford, & Washburn, 2006; Smith, Beran, Redford, & Washburn, 2006). However, the functional use of a decline or uncertainty response may be due to the existence of residual reinforcement variables.

Pure uncertainty response. Beran and colleagues (2006) sought to develop a “pure” uncertainty response that would not be contaminated by reinforcement. They trained monkeys in a numerosity discrimination. Up to nine items were presented on a computer screen. When the display had less or more than a designated center value, the monkeys were rewarded for using a joystick to move a cursor to an “L” (less) or “M” (more) on the computer screen, respectively. A wide range of center values was systematically explored, many configurations of items were used across trials, and brightness was controlled. The uncertainty response was a “?” at the bottom-center of the screen. Moving the joystick to this position ended the trial and initiated the next trial. Importantly, the authors emphasize that this method represents a pure uncertainty response in the sense that the uncertainty response was not reinforced by food, information about the correct answer, or the presentation of an easy next trial. Thus, they conclude that this was the purest trial-decline response possible.

However, the attempt to eliminate reinforcement may leave some residual reinforcement in place. The two rhesus monkeys in this study had previous experience with an uncertainty response from an earlier study that did not use the reduced-reinforcement procedure (Shields, Smith, Guttmanova, & Washburn, 2005; Smith et al., 2006). The monkeys had been rewarded in the past for moving the joystick down (which was the response in their study). Training the monkeys to use the joystick involved requiring the monkeys to learn to (1) approach a perch to view the video display, (2) reach through the cage mesh to manipulate the joystick below the monitor, (3) move the joystick so that the cursor on the screen contacted computer-generated stimuli; the joystick response was rewarded with food (Rumbaugh, Richardson, Washburn, Savage-Rumbaugh, & Hopkins, 1989; Shields et al., 2005; Washburn & Rumbaugh, 1992). A history of reinforcement associated with moving the joystick down would presumably be sufficient to generate a low-frequency tendency to select this response. Importantly, these monkeys are relatively task savvy, given that they have a long history of participating in laboratory tasks with joysticks and moving icons to target locations in addition to other laboratory tasks. Task-savvy subjects would be expected to generalize from earlier experiments to the current experiment. This generalization is not surprising given the similarity between earlier experiments and a current experiment (e.g., sitting at the experimental perch, observing the computer display, reaching an arm through the mesh cage, contacting and moving the joystick, receiving reinforcement for joystick movements, etc.); all of these factors promote the use of responses from within their experimental repertoire in new experiments, thereby allowing the experimenters to forgo the extensive training experience that would otherwise be required if new subjects were tested in each experiment.

Moreover, there may have been concurrent reinforcement because the uncertainty response reduced the delay to reinforcement in subsequent trials. Reducing delay to reward is a reinforcement variable (Carlson, 1970; Kaufman & Baron, 1968; Richardson & Baron, 2008). To examine the role of delay to reinforcement in these types of experiments Crystal and Foote (2009) conducted a simulation of reinforcement rate. The simulation used the feedback described by Beran et al. (2006) for their purest trial-decline response. On the primary task, a correct response produced one food pellet, and an incorrect response did not produce food. Critically, in their procedure, an incorrect response produced a time out of 20 sec. An uncertainty response did not produce food and did not produce a time out. We used a threshold to model the selection of the uncertainty response, as proposed by Smith et al. (2008). In the simulations, we varied a response strength for the uncertainty response from 0 to 1 using many intervening values and held all other aspects of the simulation constant (for details see Crystal & Foote, 2009). If delay to reinforcement is not a reward variable in these studies, then the amount of food per unit time will be constant as a function of the threshold values in the simulations. By contrast, if delay to reinforcement functions as a reward variable, then there will be some threshold parameter for the uncertainty response that maximizes food per unit time.

Figure 1 shows the results of the simulation. Importantly, there is a peak in food per unit time. Thus, it is possible that a subject in these types of experiments could adjust its threshold level to maximize food per unit time, and this adjustment

of the “non-reinforced” uncertainty response is reinforced by reduced delay to reinforcement in the overall procedure. This simulation shows that despite the lack of direct reward for use of the uncertainty response, there are residual reinforcement variables at work in these types of experiments.




Figure 1. Results of a simulation of reinforcement density as a function of variation in threshold for the uncertainty response. The simulation used the generalization and constant-threshold concepts from Smith et al. (2008). Reinforcement and delays were based on Beran et al. (2006). Although no food was delivered upon selecting the uncertainty response, the simulation shows that the value of the threshold for selecting the uncertainty response influences the amount of food obtained per unit time in the primary discrimination. Thus, the uncertainty response was indirectly reinforced despite efforts to eliminate reinforcement. Reproduced with permission from Crystal & Foote (2009).

Trial-by-trial feedback. Another example of the difficulty encountered in curtailing reinforcement comes from a recent study by Smith and colleagues (2006). In this uncertainty monitoring study, trial-by-trial feedback was delayed and uncoupled from the responses that earned the feedback. In particular, the monkeys were presented with a computer display that had a variable number of randomly placed pixels. Some displays were sparse and others were dense. The monkeys were required to use a joystick to move a cursor to an “S” (sparse) or “D” (dense) on the screen. A “?” appeared at the bottom-center, which was used as the uncertainty response. As the trials progressed, the monkeys earned food rewards or 20 s penalties (i.e., a time-out with a buzzer sound) based on correct or incorrect responses, respectively. However, the earned rewards or penalties were *not* delivered at the end of each trial. Instead, these consequences were *delayed* until the completion of a block of four trials. To further uncouple consequences from the responses that earned them, the feedback was *not* presented in the order in which they were earned. Instead, when the block ended, all rewards were presented first, followed by all time outs. The proportion of “S” and “D” responses tracked the density of the stimuli and declined toward the central value. The use of the “?” response peaked near the central value for one of the monkeys.

This is an innovative method to uncouple feedback in the density discrimination from the specific stimuli that were present when the feedback was earned. It is impressive that monkeys learned the task contingencies when responses and outcomes were uncoupled. The central question for our purpose is whether this new procedure eliminates the potential application of a low-level reinforcement explanation.

The study employed monkeys with previous reinforcement of the joystick response. A history of reinforcement with moving the joystick down may be sufficient to generate a low-frequency tendency to use this response in the future as discussed above. The rest of the work is done by response strength generalization for the “S” and “D” options. Because the proportion of “S” and “D” responses tracked the density of the stimuli, we may conclude that the animals had lower response tendencies for “S” and “D” near the central value based on learning the density discrimination. Note that they had these response tendencies despite the lack of transparent feedback, but density-discrimination performance is presumably based on a generalization decrement for the “S” and “D” responses. It is not necessary to hypothesize a secondary representation (i.e., “knowing that you know”) to explain the use of the three responses. In addition to the history of reinforcement of the joystick response described above, there also may have been a residual source of concurrent reinforcement that would maintain the tendency to select “?” at a relatively low frequency. The selection of the uncertainty response would reduce delay to reinforcement in the next block of trials as outlined in the simulation above. Consequently, reinforcements per unit time would be higher when the monkey selected the uncertain option compared to the scenario of not using the uncertain option. Therefore, the analysis from Figure 1 may apply even to this study.

It is worth noting that although this is an impressive procedure that made significant progress in making feedback on the primary task opaque, these features of the experiment do not eliminate response strengths for the primary task. The monkeys in Smith and colleagues’ (2006) study responded with higher accuracy on easy problems near the end of the stimulus continuum compared to difficult problems near the middle of the stimulus continuum. Thus, we could trace out a psychophysical function for the sparse-dense continuum. This function is consistent with high response tendencies to respond “sparse” for the least-dense stimuli and to respond “dense” for the most-dense stimuli. Thus, response strengths appear to be much like what would be observed if feedback was transparent.

Summary

Despite creative attempts to curtail reinforcement, the functional use of an uncertainty response may be due to the existence of residual reward variables. Indeed, if an uncertainty response was never reinforced, it seems unlikely that it would be produced by the subject, and it seems virtually impossible that it would be used functionally to express uncertainty or escape a difficult trial.

The development of new methods that are outside the domain of reinforcement learning is fraught with problems, which may be difficult, if not

impossible, to overcome. It is recommended that established methods be used to test hypotheses that pit learning and cognitive explanations against one another, which is developed in the next section.

Hypothesis Testing

What types of hypotheses need to be tested to dissociate cognition from learning? Navigating the interface of learning and cognition requires careful examination of predictions. Early work in the development of an animal model of cognition often takes the form of developing proof of concept. This work involves demonstrating that animals can produce a pattern of behavior predicted by a cognitive process. Although this is a minimal and necessary initial step in the development of an animal model, it is generally not sufficient. It is often necessary to follow-up with experiments that validate the model. Validation of the model requires eliminating alternative explanations. These studies take the form of critical experiments in that alternative proposals (e.g., low-level and high-level mechanisms) make different predictions for the study. Importantly, if an experiment can be explained by both low- and high-level explanations, additional experiments are required to decide between these alternatives. Until such experiments are conducted, it is unknown which explanation should be applied, and in the interim it is advisable to not accept complex explanations when simpler ones are adequate to explain existing data.

Validating an Animal Model of Cognition: The Case of Episodic-like Memory in Rats

Episodic memory is distinguished from semantic memory by the source of information that is represented. Semantic memory stores representations of facts about the world, without information about the context in which the memories were stored. By contrast, episodic memory includes contextual information that occurred at the time of memory storage. Tulving (1983, 1985, 2001, 2005) has argued that the subjective experiences that accompany retrieval, such as conscious recollection or the experience of the event re-occurring, are critical aspects of episodic memory. Studies of human memory exploit both behavioral and subjective sources of information, and both sources may be used to generate hypotheses about human memory. However, the development of animal models of human memory focuses exclusively on behavioral sources of evidence because subjective sources cannot be evaluated in non-verbal animals. Consequently, Clayton and colleagues (Clayton, Bussey, & Dickinson, 2003) developed behavioral criteria for studying episodic memory that focus on Tulving's (1972) classic definition of episodic memory: *what*, *where*, and *when* an event occurred during an earlier episode. It is important to note that *what-where-when* memory focuses on the content of episodic memory, which can be evaluated in animals. Because behavioral criteria do not assess subjective experiences, Clayton and Dickinson (1998) introduced the term *episodic-like* memory to describe representations of the content of episodic memories in animals.

A defining feature of episodic memory is that the memories are for *unique events or episodes*. Therefore, research that seeks to evaluate evidence for episodic-like memory in animals must document that the memory is about a specific earlier event. It is necessary to rule out alternative explanations that may exploit rules or strategies that do not require memory for a unique event.

The next sections describe the development of a rodent model of episodic memory in rats. The review is divided into proof of concept and validation studies.

Proof of Concept Studies

We developed our approach to ask if rats have what-where-when memories based on related studies by Clayton and colleagues (Clayton, Dally, Gilbert, & Dickinson, 2005; Clayton & Dickinson, 1998, 1999a, 1999b, 1999c; Clayton, Yu, & Dickinson, 2001, 2003; de Kort, Dickinson, & Clayton, 2005) with scrub jays. Our experiments (Babb & Crystal, 2005, 2006a, b; Zhou & Crystal, 2009, 2011) demonstrate that rats have a detailed representation of what, where, and when specific events occurred. In this section, initial evidence for what-where-when memory is described. In the next section, studies that identify the temporal processes that support what-where-when memory are described. In a later section, our initial attempt to validate the model through convergent lines of evidence is described.

The rats were trained in an eight-arm radial maze, which has eight runways radiating from a central hub. In a standard radial maze experiment, a small piece of food is consumed in its entirety when encountered at the distal end of each location. The rat is permitted to explore the maze and consume the available food. Because the food is consumed in its entirety when encountered and no additional food is available after consumption in a standard radial maze experiment, the optimal strategy is to visit each location once. Extensive evidence indicates that rats perform at a near-optimal level and that they do so based on memory for recently visited locations (Olton & Samuelson, 1976; Roberts, 1998).

Babb and Crystal (2006b) used a modified version of the standard radial maze experiment as follows. The trial was divided into three parts: a *study* phase, a retention interval, and a *test* phase. The retention interval interrupted the trial, which otherwise continued from study to test phases. We use the terms *study* and *test* because the locations of food in the test phase depended on memory for locations of food in the study phase. In the study phase, four randomly selected locations provided food (the other arms were blocked by closed guillotine doors). Two of the study locations had standard chow-flavored food. One of the study locations, randomly selected each trial, had a distinctive grape-flavored food, and one study location, also randomly selected, had a distinctive raspberry-flavored food. After eating its first helpings of food in the study phase, the rat was removed from the maze for a retention interval which was either 1 hr (short) or 6 hrs (long). The rat ate its second helpings of food in the test phase, at which point the trial was a continuation from the study phase. All doors were open in the test phase, but food was not available at the recently visited chow locations from the study phase. By contrast, the locations that were closed in the study phase provided chow-flavored food in the test phase. Thus the test phase is a test of memory for recently

presented information from the study phase (i.e., memory about the study phase was required to avoid visiting now-depleted chow locations). The locations that recently had grape and raspberry replenished grape and raspberry, respectively, after a long retention interval, but not after a short retention interval.

Optimal performance in the task described above required what-where-when memory. The rats received replenishment of distinctive flavors after one, but not the other, retention interval. Thus, a temporal component was required to efficiently revisit the locations that were about to replenish, and reduce this tendency at other times. The rats revisited the grape and raspberry locations at a higher rate when these locations were about to replenish relative to other occasions when these locations were not about to replenish. This could be accomplished if they remembered the locations that recently had a distinctive flavor and temporal information about the study and test phases. The rats accomplished this differential rate of revisiting grape and raspberry locations while maintaining high accuracy in avoiding revisits to chow locations (which never replenished).

Next, we tested the hypothesis that rats remember the *specific* flavors (grape vs. raspberry) that were encountered. To accomplish this objective, we devalued one of the distinctive flavors, while leaving the other intact. In one experiment, we satiated the rats to one of the flavors (grape or raspberry) during the long retention interval while leaving the other flavor intact; our experiment used a flavor-specific devaluation (Balleine & Dickinson, 1998; Colwill, 1994). When the trial continued after access to one of the flavors, the rats selectively reduced revisits to the devalued flavor while continuing to revisit the location with the non-devalued flavor. In another experiment we used banana and chocolate distinctive flavors. During the long retention interval, we devalued chocolate by pairing it with lithium chloride; our experiment used a learned taste aversion manipulation (Batson, Best, Phillips, Patel, & Gilleland, 1986; Melcer & Timberlake, 1985). The rats selectively eliminated revisits to the chocolate location without reducing revisits to the banana location. These data suggest that the rats remember the specific flavors that were encountered at the locations, in addition to temporal information.

Temporal Processing in What-where-when Memory

In the experiments described above, the rats adjusted revisit rates after different retention intervals at the appropriate locations, which implicates what-where-when memory. But what type of temporal mechanism was used? And is memory of a *specific episode* required to accomplish the adjustment in revisit rates (i.e., is it *episodic* memory)? There are three types of temporal information that may support what-where-when memory. First, because the trials always started at an approximately constant time of day, the test phases also occurred at approximately constant times of day (i.e., 1 hr or 6 hrs later). Therefore, time of day at the *test* phase could be used as a cue to adjust revisit rates. Second, the *interval* between study and test phases could be used to adjust revisit rates to the distinctively baited locations. Third, the rats might have remembered the specific study episode, including *when* (i.e., the time of day at which) the *study* event occurred. Importantly, the third proposal, but not the first two proposals, requires

memory of the study episode. Thus, discrimination of time of day at test or timing an interval with respect to an earlier event represent alternative explanations of what-where-when memory that would not require episodic memory.

A number of experiments have investigated the temporal information that rats used for the *when* component of what-where-when memory. Three investigations will be reviewed.

First, Babb and Crystal (2006a) tested the hypothesis that rats were using time of day at the test phase as a cue to solve the what-where-when task. To test this hypothesis, the time of test after both short and long retention intervals was held constant. We did this by using 1 hr and 25 hr retention intervals, and starting the study phases at an approximately constant time of day. Chocolate replenished at its trial-unique location after a 25 hr retention interval, but not after a 1-hr retention interval. For example, consider a study phase that starts at 1200. The test phase began 1 hr or 25 hrs later. If the retention interval was 1 hr, then the test phase began at 1300 on the *same day* as the study phase. By contrast, if the retention interval was 25 hr, then the test phase began at 1300 on the *next day*. Note that it was 1300 at test after either short or long retention intervals. Thus, time of day at test could not be used to adjust revisit rates at the chocolate location. Yet the rats revisited the distinctively baited location at a higher rate when it was about to replenish relative to times when it was not about to replenish. The differential revisit rate clearly could not be based on time of day at test because it was the same time of day in both retention interval conditions. Note that the time of study was also constant (1200 in the example above). Thus, a potential solution to the task is to estimate the time since the study phase. This could be accomplished by timing an interval between study and test. Alternatively, the rat could encode the time of occurrence of the study phase and current time of occurrence and obtain the interval by subtraction (Gallistel, 1990). Using an independent method, we showed that rats can discriminate alternate days (Pizzo & Crystal, 2007). Because 1 hr and 25 hr retention intervals produced tests that occurred at the same time of day on alternate days, the discrimination of alternate days (e.g., noon today vs. noon yesterday) is another mechanism by which the animals may have solved this what-where-when task. Recently, Naqshbandi and colleagues (Naqshbandi, Feeney, McKenzie, & Roberts, 2007) replicated our study using different methods, which provides additional evidence that rats use what-where-when memories and that rats did not solve the discrimination by using time of day at the test phase.

In a second investigation on temporal information in what-where-when memory, Roberts and colleagues (2008) carefully selected the times at which study and test phases occurred to eliminate the correlation between time of study and interval between study and test. In the Roberts et al. study, some trials had study phases that started at a constant time of day (thereby having test phases at varying times of day); other trials had the test phases occur at a constant time of day (thereby having study phases start at varying times of day). The time of day at which study and test phases occurred and the retention interval between study and test were arranged so that some rats received a *consistent replenishment pattern* with respect to time of study (referred to as the *when* group), retention interval (referred to as the *how-long-ago* group), or both (the *when + how-long-ago* group).

The when group failed to learn the replenishment contingency whereas the other two groups adjusted revisit rates to correspond with the replenishment contingency. It should be noted that rats in the when group received *inconsistent* feedback (i.e., replenishment) after short and long retention intervals. A potential explanation of these data is the hypothesis that when both when and how-long-ago information are available, rats appear to rely on how-long-ago (or learn about it more rapidly); this may be a form of overshadowing, which occurs under conditions of cue competition (De Houwer, Beckers, & Vandorpe, 2005). This hypothesis does not preclude the possibility that time of study may be encoded, but different experimental techniques might be necessary to reveal remembering of when the study episode occurred.

In a third investigation on temporal-information processing in what-where-when memory, Zhou and Crystal (2009) sought to evaluate what-where-when memories under conditions in which how-long-ago cues were irrelevant to predicting replenishment. Because the data of Roberts et al. (2008) suggest that how-long-ago dominates when multiple temporal cues are available, we made how-long-ago cues *irrelevant* to predicting replenishment. When how-long-ago was rendered irrelevant, we found that rats remember *when* (i.e., time of day) an earlier study episode occurred, in addition to *what* occurred and *where* it took place. In the Zhou and Crystal study, rats were tested in the morning or afternoon (but not both) on separate days (see Fig. 2a). Chocolate replenished at a daily unique location at only one of these times (morning for half of the rats; afternoon for the other rats). The interval between the study and test phases was constant (approximately 2 min). Because the location of chocolate varied randomly across days and the morning and afternoon sessions were presented in random order, what-where-when memory would be needed to seek out the chocolate location selectively on occasions when chocolate was about to replenish. When the chocolate location was about to replenish, the rats revisited that location at a higher rate relative to non-replenishment trials (Fig. 3a). These data suggest that rats used what-where-when memories to adjust revisit rates to the daily-unique chocolate location. Importantly, what-where-when in this study could not be based on the delay between study and test. Consequently, memory performance cannot be based on judging relative familiarity of the study items or timing an interval between study and test).

Next Zhou and Crystal (2009) sought to determine the type of timing mechanism used in what-where-when memory. There are two remaining hypotheses. According to the *circadian time-of-day* hypothesis, the rats used a circadian signal (i.e., morning vs. afternoon) (Gallistel, 1990; Takahashi, Turek, & Moore, 2001) to adjust revisit rates at the daily-unique chocolate location. Alternatively, according to the *interval-timing* hypothesis, the rats timed the interval from light onset in the colony to the morning and afternoon sessions. Morning sessions occurred 1 hr after light onset in the colony, and afternoon sessions occurred 7 hr after light onset. Consequently, we used a phase shift of 6 hr to test these hypotheses. The lights in the colony were turned on 6 hr early and the probe session was conducted at the usual time in the morning (see Fig. 2b). According to the circadian time-of-day hypothesis, the rats would treat the probe as a morning session because an endogenous circadian oscillator is not expected to

adjust immediately to a phase shift (Gallistel, 1990; Gibbon, Fairhurst, & Goldberg, 1997; Takahashi et al., 2001). Alternatively, according to the interval-timing hypothesis, the rats would treat the probe as an afternoon session because afternoon sessions typically occur 7 hr after light onset in the colony; there is independent evidence that rats can time long intervals with respect to colony-light onset in the range of hours (Crystal, 2001). However, the rats did not use the interval between light onset and the session, suggesting that they used circadian time of day (Fig. 3b). Next, we sought to determine if it was the *time of day* at *study* or at *test* that was responsible for the different rates of revisiting the chocolate location. Because a 2-min delay between study and test is too small for rats to discriminate based on a circadian oscillator (Pizzo & Crystal, 2004), we increased the delay to 7 hr (see Fig. 3c). Importantly, the time of day at study was *familiar* from prior training, but the time of day at test was *unfamiliar* (approximately 7 hr later than usual). Consequently, if the rats used time of day at study, then they should continue to differentially revisit the chocolate locations. Alternatively, if the rats used time of day at test, then there is no basis for them to revisit chocolate locations at different rates in the morning and afternoon because the test times were unfamiliar. When tested with *novel* test times of day after *familiar* morning or afternoon study times of day, we observed complete transfer (i.e., the differential rates of revisiting occurred on the very first trial in the morning and the very first trial in the afternoon; Fig. 3c and d). These data suggest that at the time of memory assessment, the rats remembered the time of day at which the study episode occurred. We obtained additional evidence for the same conclusion by conducting a conflict test. Because the 7-hr delays between study and test phases produced a 1-hr overlap between the two types of trials, it was possible to start a trial with a late study phase and end the trial with an early test phase (see Fig. 2d). Again we sought to determine if the rats were adjusting revisit rates in the test phase based on the time of day at test (we refer to this proposal as the *test-time* hypothesis) or based on *memory* of the time of day at which the study phase occurred (referred to as the *study-time* hypothesis). According to the *test-time* hypothesis, the rats should revisit at the usual baseline rate that typically occurred on tests at that time of day. Alternatively, according to the *study-time* hypothesis, the rats should revisit at the usual time of day that occurred after a later study time (which usually is followed by a test 7 hrs later rather than 1 hr later). The rats adjusted chocolate revisits based on the time of day at study rather than the time of day at test (Fig. 3e). These data also suggest that rats remembered the study episode, and the time of day at which the study episode occurred.

At the time of memory assessment, the rats remembered the time of day at which the study episode occurred. Importantly, these experiments suggest that rats remember what-where-when under conditions in which how-long-ago cues were made irrelevant to performance. Thus, the *relative strength* of memories that decay over time (i.e., *relative familiarity* of study items) *cannot* explain our results because the delay between study and test phases was constant in each experiment.

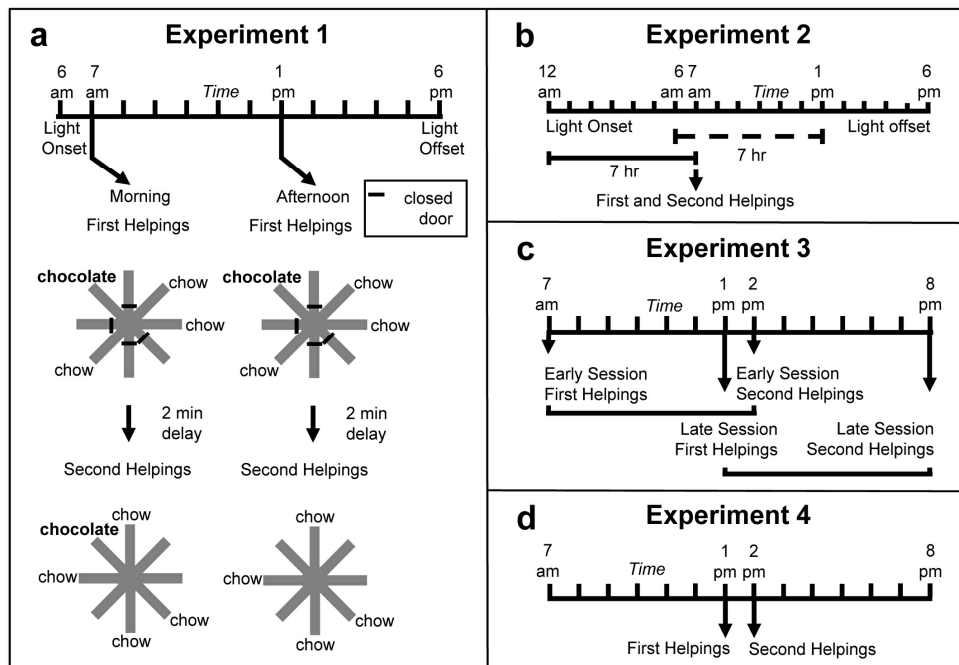


Figure 2. Experimental design of Zhou and Crystal's (2009) study. **A)** Design of Experiment 1. The morning or afternoon was randomly selected for presentation of first helpings (study phase; encoding) and second helpings (test phase; memory assessment) of food. An example of the accessible arms and flavors in a study and test phases is shown. Chocolate or chow flavored pellets were available at four arms in the study phase (randomly selected); closed doors prevented access to the other arms. After a 2-min retention interval, chow-flavored pellets were available at previously inaccessible locations in the test phase. Chocolate replenished at the location that had chocolate in the study phase in the replenishment condition (shown for the morning session). Chocolate did not replenish at the other time of day in the non-replenishment condition (shown in the afternoon session). Chocolate replenished in the test phase conducted in the morning (7 a.m.) but not in the afternoon (1 p.m.) for half of the rats; these contingencies were reversed for the remaining rats (not shown). For each rat, one session (i.e., study and test phases) was conducted per day. The same study-phase baiting pattern was used in the figure, but these arms were randomly selected in each session for each rat. **b)** Phase-shift design of Experiment 2. Light onset occurred at 12 a.m. (i.e., 6 hr earlier than in Experiment 1) and the study and test phases occurred at the time of a typical morning session (i.e., starting at 7 a.m.). Note that 7 hrs elapsed between light onset and the study-test sequence (solid horizontal line), which is comparable to the time between the typical light onset and a typical afternoon session (dashed horizontal line) in Experiment 1. The design of the experiment puts predictions for time-of-day and how-long-ago cues in conflict. Thus, a rat would be expected to behave as in its morning baseline (based on time of day) or as in its afternoon baseline (based on how long ago). **c)** Transfer-test design of Experiment 3. The time of day at which the study phase occurred was the same as in Experiment 1 (i.e., 7 a.m. in early sessions or 1 p.m. in late sessions). The introduction of 7-hr retention intervals in Experiment 3 produced test phases that occurred at novel times of day (2 p.m. in early sessions and 8 p.m. in late sessions). Early and late sessions had study times (but not test times) that corresponded to those in Experiment 1. The first two sessions in Experiment 3 consisted of one replenishment and one non-replenishment condition. An early or late session was randomly selected on subsequent days. Differential revisits to the chocolate location is expected if the rats were adjusting revisit rates based on the time of day at which the study episode occurred; revisit rates are expected to be equal in early and late sessions if the rats used time of day at which the test phase occurred. Study and test phases were as in Experiment 1, except that they were separated by 7-hr delays (shown by horizontal brackets). **d)** Conflict-test design of Experiment 4. The study and test phases occurred at 1 p.m. and 2 p.m., respectively. These times correspond to the typical time of day at which a late-session study phase and early-session test phase occurred in

Experiment 3. The design of the experiment put predictions for time of day at study and time of day at test in conflict. A rat would be expected to behave as in its early-session, second-helpings baseline (based on test time of day) or as in its late-session, second-helpings baseline (based on study time of day). Reproduced with permission from Zhou, W., & Crystal, J. D. (2009). Evidence for remembering when events occurred in a rodent model of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, 106(23), 9525-9529. © 2009 National Academy of Sciences, U.S.A.




Figure 3. **a)** Rats preferentially revisit the chocolate location when it is about to replenish in Experiment 1. The probability of a revisit to the chocolate location in the first four choices of a test phase is shown for replenishment and non-replenishment conditions; replenish and non-replenish sessions were presented in random order. **b)** Rats used time of day, rather than an interval, to adjust revisit rates in Experiment 2. Rats treated the study-test sequence as a morning session, suggesting the use of a time-of-day rather than an interval-timing mechanism. The figure plots the difference between observed and baseline revisit rates. For the bar labeled interval, the baseline was the probability of revisiting chocolate in the afternoon; thus, the significant elevation above baseline shown in the figure suggests that the rats did not use an interval mechanism. For the bar labeled time of day, the baseline was the probability of revisiting chocolate in the morning; thus, the absence of a significant elevation above baseline is consistent with the use of time of day. The horizontal line corresponds to the baseline revisit rate to the chocolate location from Experiment 1. Positive difference scores correspond to evidence against the hypothesis indicated on the horizontal axis. **c)** and **d)** Rats preferentially revisited the chocolate location when it was about to replenish when the study, but not the test, time of day was familiar in Experiment 3. The probability of a revisit to the chocolate location in the first four choices of a test phase is shown for first replenishment and first non-replenishment conditions (**c**; initial) and for subsequent sessions (**d**; terminal). **e)** Rats remembered the time of day at which the study episode occurred in Experiment 4. Rats treated the novel study-test sequence as a late-session test phase, suggesting memory of the time of day at study rather than discriminating time of day at test. The figure plots the difference between observed and baseline revisit rates. For the bar labeled test time, the baseline was the probability of revisiting chocolate in the test phase of the early session in Experiment 3; thus, the significant elevation above baseline suggests that the rats did not use the time of day at test to adjust revisit rates. For the bar labeled study time, the baseline was the probability of revisiting chocolate in the test phase of the late session in Experiment 3; thus, the absence of a significant elevation above baseline is consistent with memory of the time of day at study. The horizontal line corresponds to the baseline revisit rate to the chocolate location from Experiment 3 (terminal). Positive difference scores correspond to evidence against the hypothesis indicated on the horizontal axis. **a-e.** Error bars indicate SEM. **a, c, and d.** The probability expected by chance is 0.41. Repl = replenishment condition. Non-repl = non-replenishment condition. **a)** * $p < 0.001$ difference between conditions. **b)** * $p < 0.04$ different from baseline. **c)** and **d)** * $p < 0.04$ and ** $p < 0.0001$ difference between conditions. **e)** * $p < 0.001$ different from baseline. Reproduced with permission from Zhou, W., & Crystal, J. D. (2009). Evidence for remembering when events occurred in a rodent model of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, 106(23), 9525-9529. ©2009 National Academy of Sciences, U.S.A.

We argued that rats, at the time of memory assessment, remember when the study episode occurred, in addition to what happened and where it took place based on a single, brief encoding episode (Zhou & Crystal, 2009). Recent evidence

suggests that male meadow voles anticipate the sexual receptivity of female voles based on a brief encoding episode that includes time of day, the stage of postpartum estrus, and location of the encounter (Ferkin, Combs, delBarco-Trillo, Pierce, & Franklin, 2008). Honeybees also integrate time of day with visual and place information based on a single brief encoding episode (Pahl, Zhu, Pix, Tautz, & Zhang, 2007; Zhang, Schwarz, Pahl, Zhu, & Tautz, 2006). Thus, the ability to use what, where, and when based on time of day may be quite widespread.

The goal of the research reviewed above was to test the hypothesis that rats have episodic-like memory (i.e., they remember a specific episode including specific details of what, where, and when the event occurred). The validation of an animal model of human memory holds great potential for gaining insight into the neurobiology of human memory and disorders of memory. However, this effort is likely to require convergent lines of evidence (Crystal, 2009; Shettleworth, 1998). It may be argued that any single approach is likely to be limited by a set of competing, alternative explanations. For example, perhaps animals learn about the compound or configural what-where-when cues. Similarly, perhaps the animals have learned a set of semantic rules. Thus, a careful selection of multiple approaches may overcome weaknesses that may exist if each approach were treated separately. Next, we review some initial validation studies.

Validation Studies

Convergent lines of evidence may play an important role in the assessment of psychological processes in animals (Crystal, 2009; Shettleworth, 1998). Any single approach may be limited by a set of competing, alternative explanations. However, a careful selection of multiple approaches may overcome weaknesses that would exist if each approach were treated separately. We consider two convergent lines of evidence as our initial attempt to validate a rodent model of episodic memory.

Encoding Failure Hypothesis

The study by Zhou and Crystal (2009) suggests that rats have specific knowledge about earlier episodes, including when the episode occurred, what happened, and where it took place. However, the rats might have learned that encoding the chocolate location was not required in some time-of-day conditions, which we refer to as the *encoding failure hypothesis*. For example, a rat might solve the task by encoding the location of chocolate at one time of day (e.g., when chocolate replenishes in the morning) but not encoding the location of chocolate at the other time of day (e.g., when chocolate does not replenish in the afternoon); notice that differential rates of revisiting chocolate would occur in this situation without remembering the episode. The encoding failure hypothesis could also explain data from other studies (Naqshbandi et al., 2007).

To address this potential problem, Zhou and Crystal (2011) provided rats with daily information about a preferred food type (chocolate) that replenished or failed to replenish at its previously encountered location (Fig. 4). Another flavor (regular chow) was available at all other locations but never replenished.

Importantly, although some of the information needed to predict replenishment was available at the time of encoding (location, time of day, food flavor), one critical piece of information needed to predict replenishment was not presented until immediately before the memory assessment. The presence or absence of additional chocolate pellets in a central location could be used to predict replenishment when combined with time of day. Critically, although time of day was known at study, the subsequent baiting of the central location could not be predicted at study. Thus, to solve this task, it was necessary to encode the location of chocolate and time of day at study, but decoding of replenishment was delayed until immediately prior to memory assessment. To preferentially revisit the chocolate location when it was about to replenish at the memory assessment phase, the rats needed to remember where they found it during their earlier encoding phase. For example, for some rats, the presence of chocolate in the hub in the morning (left columns of Fig. 4) and the absence of chocolate in the hub in the afternoon (right columns of Fig. 4) allowed the rat to predict the forthcoming replenishment of chocolate. For other rats, the role played by presence and absence of food in the hub was reversed to counterbalance assignment of conditions across the rats. Because it was impossible to predict whether chocolate would be replenished later, rats had to encode the episode on each study occasion. If the *encoding failure hypothesis* explained the results from previous studies, it would be impossible for rats to solve the task. By contrast, if rats encoded detailed information about chocolate, they should show a significantly higher revisit rate to chocolate when it was about to replenish relative to when it was not about to replenish. We conducted two tests of the encoding-failure hypothesis. Initially, we used a constant, minimal (2-min) retention interval. Next, we conducted a transfer test with a longer retention interval (approximately 1 hr).

The results rule out the encoding failure hypothesis. The rats revisited the chocolate location when it was about to replenish relative to the nonreplenishment condition (see Fig. 5a and b). Rats were more likely to revisit the chocolate location in the replenishment conditions compared to the non-replenishment conditions. Revisit probabilities were similar for both retrieval cues and the effect of replenishment condition did not depend on the retrieval cue. Differential rates of revisiting chocolate-flavored locations were accomplished while rats accurately avoided revisits to depleted chow-flavored locations. To successfully solve this task, rats had to encode the episode at study, because the critical information about whether or not chocolate would be replenished at the recently visited location was not available until the presence or absence of chocolate pellets in the central hub immediately before the memory assessment.




Figure 4. Schematic representation of Zhou and Crystal's (2011) study. The morning or afternoon was randomly selected for presentation of study and test phases. The figure shows an example of the accessible arms and flavors in encoding and the corresponding memory assessment phases that would occur after a 2-min retention interval. The presence or absence of chocolate pellets in the central hub immediately prior to memory assessment was needed to predict the replenishment of chocolate in the test phase. In the replenishment conditions, chocolate replenished at the location that recently delivered chocolate, which was predicted by the presence or absence of food (e.g., presence of chocolate in the central hub immediately prior to second helpings memory assessment in the morning but absence of chocolate in the hub in the afternoon); these contingencies were reversed in the non-replenishment conditions. These conditions were counterbalanced across rats (not shown). For each rat, one session (i.e., study phase, hub-baiting retrieval cue, and test phase) was conducted per day. The same study-phase baiting pattern was used to illustrate morning and afternoon sessions in the figure to facilitate inspection of presence and absence of chow and chocolate, but these arms were randomly selected in each session for each rat. Reproduced from Zhou W., & Crystal, J. D.(2011). Validation of a rodent model of episodic memory. *Animal Cognition*, 14(3), 325-340. © 2011 Springer-Verlag.

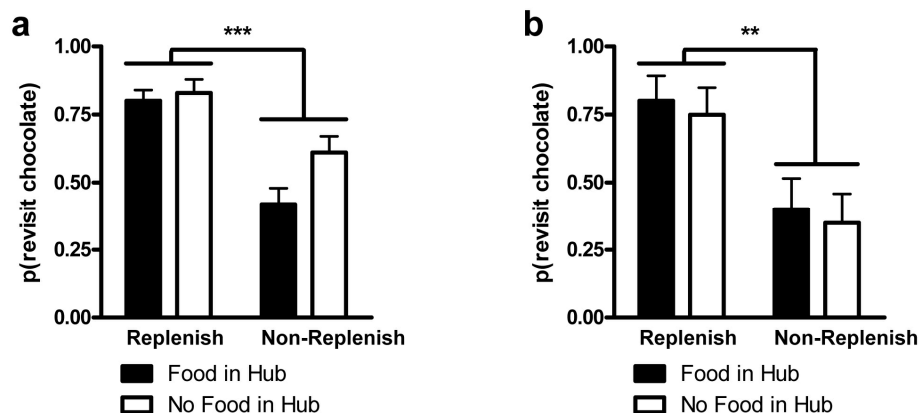


Figure 5. Rats preferentially revisit the chocolate location when it is about to replenish when the retention interval was approximately 2 min (**a**) and 1hr (**b**). The probability of a revisit to the chocolate location in the first four choices of a test phase is shown for replenishment and non-replenishment conditions; replenishment and non-replenishment conditions were presented in random order. The presence or absence of food in the hub, immediately prior to memory assessment, served as a cue that could be used to predict the replenishment or non-replenishment of chocolate. Error bars represent 1 SEM. **a)** *** $p < 0.001$ difference between replenishment and non-replenishment conditions. **b)** Each condition was tested once, in random order. ** $p = 0.009$ difference between replenishment and non-replenishment conditions. Reproduced from Zhou, W., & Crystal, J. D.(2011). Validation of a rodent model of episodic memory. *Animal Cognition*, 14(3), 325-340. © 2011 Springer-Verlag.

Unexpected Question

A fundamental aspect of episodic memory is that it can be used to report information when the test of memory is unexpected. One problem with many of the paradigms used to evaluate episodic-like memory is that extensive training is required (Singer & Zentall, 2007; Zentall, 2005, 2006; Zentall, Clement, Bhatt, & Allen, 2001). Zentall and colleagues (2001) have argued that it is not possible to preclude semantic-like knowledge in the discrimination of what-where-when because the contingencies of food availability are explicitly trained; the explicit training might foster the development of semantic knowledge about experimental contingencies. Therefore, they proposed that it is preferable to assess the capacity for episodic-like memory in animals by using an unexpected question about a recent event. Accordingly, they proposed that documenting episodic-like memory requires a demonstration that the animal can report on a recent event when there was no expectation that such a report would be required (i.e., answer an unexpected question).

Zhou and Crystal (2011) used two rooms to construct an unexpected question. In initial training, the rats were familiarized with two rooms. In Room 1, chocolate was encountered at study and replenished at test (as described above). In Room 2, they were never exposed to chocolate or its replenishment. After demonstrating that the rats discriminate the two rooms, we used a *probe* test to assess performance with an unexpected question. The study phase with chocolate began in Room 1, as in previous training. We may assume that the rats encoded the

location of chocolate within Room 1 and planned to revisit that location or not depending on the identity of the retrieval cue to be provided immediately prior to the test phase. Importantly, the expectation to revisit the location in Room 1 would be defined by the global environmental cues in that room (e.g., Roberts, 1998;). The trial continued with a test in Room 2. We used the retrieval cue that would prompt revisiting of the chocolate location. At the time of encoding in the probe, the rat is expected to store memories of Room 1 locations with respect to Room 1 global environmental cues. It is not possible for the rat to store information with respect to Room 2 cues at the time of encoding because the rat cannot know that it will be unexpectedly tested in Room 2. When the rat is tested in Room 2, it must retrieve a memory of Room 1 and use it in Room 2. They were unexpectedly asked to retrieve information about the expected location of chocolate in a novel context. If rats can answer a question when they were unexpectedly asked to retrieve information about the expected location of chocolate in a novel context, then they should revisit the corresponding chocolate location in Room 2 at a relatively high rate on the first trial in which the unexpected question occurred.

In a second unexpected question probe, we replaced the chocolate arm bait with a novel flavor, sucrose. If rats are able to report about the earlier event of a novel flavor in a novel context, we would expect a relatively high revisit rate to the corresponding arm in Room 2 that was recently baited with sucrose in Room 1. If rats were not able to report the earlier experience of a novel flavor at a novel context, we would not expect to obtain a relatively high revisit rate to the corresponding sucrose arm.

In both chocolate and sucrose unexpected question probes (Fig 6a and b), the rats revisited the corresponding location in Room 2 at a high rate, as expected for a replenishment condition (i.e., significantly higher than the rate in the non-replenishment condition and significantly higher than the rate expected by chance). Differential rates of revisiting the chocolate flavored location were accomplished while rats moderately avoided revisits to depleted chow-flavored locations. We propose that the rats retrieved a memory of an earlier episode when they were unexpectedly asked to do so.

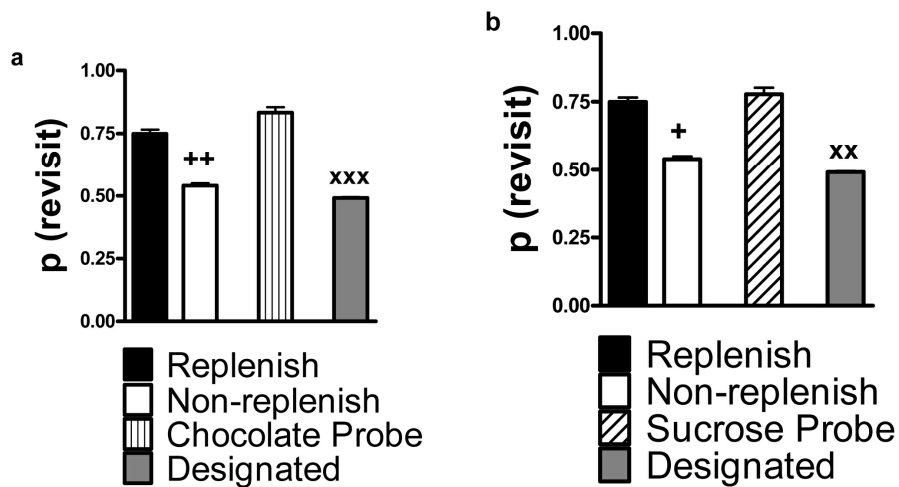


Figure 6. a) In the chocolate probe, rats revisited the corresponding chocolate location in Maze 2 at a higher rate when chocolate was unexpectedly replenished compared with the non-replenishment *mixed baseline* in Maze 1 (labeled Replenish and Non-replenish) and *designated baseline* in Maze 2 (labeled Designated; a measure of randomly selected arm entries). The probability of revisiting corresponding chocolate location in Maze 2 is labeled Chocolate Probe. ++ $p = 0.005$ difference between chocolate probe and non-replenishment mixed baseline; xxx $p = 0.001$ difference between chocolate probe and designated baseline. *b)* In the sucrose probe, rats revisited the corresponding sucrose location in Maze 2 when sucrose was unexpectedly replenished compared with the non-replenishment mixed baseline in Maze 1 and designated baseline in Maze 2. + $p = 0.03$ (one-tailed) difference between sucrose probe and non-replenishment mixed baseline; xx $p = 0.01$ difference between sucrose probe and designated baseline. Reproduced from Zhou, W., & Crystal, J. D.(2011). Validation of a rodent model of episodic memory. *Animal Cognition*, 14(3), 325-340. © 2011 Springer-Verlag.

Summary

The experiments represent our initial attempt to validate a rodent model of episodic memory by providing convergent lines of evidence. Our findings rule out the encoding failure hypothesis, suggesting that failing to encode the content of an episode is not used as an alternative strategy to solve what-where-when tasks. Moreover, rats demonstrated the ability to retrieve a memory of an earlier episode when they were unexpectedly asked to do so. The ability to answer an “unexpected” question captures a feature of human episodic memory – at the time of memory retrieval, information is used that could not be anticipated at an earlier point. Importantly, answering an unexpected question rules out the use of expectations derived from well-learned semantic rules established by extensive training. A further test of an unexpected question in rats would examine incidental encoding given that people are able to report about information that was incidentally encoded.

The ability to find chocolate and sucrose in a novel context in the probes may also be interpreted as evidence of flexible use of memory of the study episode. The flexibility was likely afforded by the rats’ storage of a memory of the *study* episode rather than by storing the future response at the time of study. Documenting flexible use of study-episode memory is an important feature of

episodic-like memory (Clayton, Bussey, et al., 2003).

Conclusions About Best Practices

In our efforts to develop and validate a rodent model of episodic memory, we used established methods in the service of testing cognitive hypotheses. For example, we used well-established flavor-specific devaluation techniques to isolate evidence for cognitive hypotheses, namely flexibility and a detailed representation of the "what" content in what-where-when memory. We used well-established techniques from interval and circadian research domains to explore the "when" component in what-where-when memory.

Our efforts to develop a rodent model of episodic memory have focused on proof of concept and validation. Validation has focused on two alternative low-level hypotheses. Yet, additional tests of low-level hypotheses will enhance confidence that a memory of a previous episode is implicated. Some aspects of human episodic memory will remain unexplored, namely subjective experiences that accompany episodic memory in people. However, other aspects of human episodic memory provide valuable targets for modeling in animals. For example, episodic memories in people can be assessed after incidental encoding; this feature was not captured in our unexpected question (Zhou & Crystal, 2011), although it can be explored in future experiments. Moreover, episodic memories in people include rich source monitoring (Johnson, Hashtroudi, & Lindsay, 1993; Mitchell & Johnson, 2009). Although what-where-when memory captures some source information, the limits of source monitoring in rats remains to be explored in future research.

Although each approach may have limitations, the use of multiple, independent approaches may be used to test the hypothesis that rats and other animals have episodic-like memory. One advantage of a multi-method approach includes the ability to obtain a comprehensive description of the elements of episodic-like memory in rats (Crystal, 2009, 2010). It is possible that rats possess some aspects of episodic-like memory, but in some significant ways, aspects of memory differ from those observed in people or other animals. For example, none of our methods targeted subjective experiences; instead, we focus on the content of episodic memories. Overall, the use of multiple criteria and assessment methods enables a more complete picture of an animal's representational systems than could be obtained if only a single method were used.

References

- Babb, S. J., & Crystal, J. D. (2005). Discrimination of what, when, and where: Implications for episodic-like memory in rats. *Learning & Motivation, 36*, 177-189.
- Babb, S. J., & Crystal, J. D. (2006a). Discrimination of what, when, and where is not based on time of day. *Learning & Behavior, 34*, 124-130.
- Babb, S. J., & Crystal, J. D. (2006b). Episodic-like memory in the rat. *Current Biology, 16*, 1317-1321.

- Balleine, B. W., & Dickinson, A. (1998). The role of incentive learning in instrumental outcome revaluation by sensory-specific satiety. *Animal Learning & Behavior*, *26*, 46-59.
- Batson, J. D., Best, M. R., Phillips, D., Patel, H., & Gilleland, K. R. (1986). Foraging on the radial-arm maze: Effects of altering the reward at a target location. *Animal Learning & Behavior*, *14*, 241-248.
- Beran, M. J., Smith, J. D., Redford, J. S., & Washburn, D. A. (2006). Rhesus macaques (*Macaca mulatta*) monitor uncertainty during numerosity judgments. *Journal of Experimental Psychology: Animal Behavior Processes*, *32*, 111-119.
- Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, *12*, 187-192. doi: 10.1016/j.tics.2008.02.010
- Carlson, J. G. (1970). Delay of primary reinforcement in effects of two forms of response-contingent time-out. *Journal of Comparative and Physiological Psychology*, *70*, 148-153.
- Clayton, N. S., Bussey, T. J., & Dickinson, A. (2003). Can animals recall the past and plan for the future? *Nature Reviews Neuroscience*, *4*, 685-691.
- Clayton, N. S., Dally, J., Gilbert, J., & Dickinson, A. (2005). Food caching by Western scrub-jays (*Aphelocoma californica*) is sensitive to the conditions at recovery. *Journal of Experimental Psychology: Animal Behavior Processes*, *31*, 115-124.
- Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*, 272-274.
- Clayton, N. S., & Dickinson, A. (1999a). Memory for the content of caches by scrub jays (*Aphelocoma coerulescens*). *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 82-91.
- Clayton, N. S., & Dickinson, A. (1999b). Motivational control of caching behaviour in the scrub jay, *Aphelocoma coerulescens*. *Animal Behaviour*, *57*, 435-444.
- Clayton, N. S., & Dickinson, A. (1999c). Scrub jays (*Aphelocoma coerulescens*) remember the relative time of caching as well as the location and content of their caches. *Journal of Comparative Psychology*, *113*, 403-416.
- Clayton, N. S., Yu, K. S., & Dickinson, A. (2001). Scrub jays (*Aphelocoma coerulescens*) form integrated memories of the multiple features of caching episodes. *Journal of Experimental Psychology: Animal Behavior Processes*, *27*, 17-29.
- Clayton, N. S., Yu, K. S., & Dickinson, A. (2003). Interacting cache memories: Evidence for flexible memory use by Western scrub-jays (*Aphelocoma californica*). *Journal of Experimental Psychology: Animal Behavior Processes*, *29*, 14-22.
- Colwill, R. M. (1994). Associative representations of instrumental contingencies. *The Psychology of Learning and Motivation*, *31*, 1-72.
- Crystal, J. D. (2001). Circadian time perception. *Journal of Experimental Psychology: Animal Behavior Processes*, *27*, 68-78.
- Crystal, J. D. (2006a). Sensitivity to time: Implications for the representation of time. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 270-284). New York: Oxford University Press.
- Crystal, J. D. (2006b). Time, place, and content. *Comparative Cognition & Behavior Reviews*, *1*, 53-76.
- Crystal, J. D. (2009). Elements of episodic-like memory in animal models. *Behavioural Processes*, *80*, 269-277.
- Crystal, J. D. (2010). Episodic-like memory in animals. *Behavioural Brain Research*, *215*, 235-243. doi: 10.1016/j.bbr.2010.03.005
- Crystal, J. D. (in press). Sensitivity to time: Implications for the representation of time. In E. A. Wasserman & T. R. Zentall (Eds.), *Handbook of comparative cognition*. Oxford: Oxford University Press.

- Crystal, J. D., & Foote, A. L. (2009). Metacognition in animals. *Comparative Cognition & Behavior Reviews*, 4, 1-16.
- De Houwer, J., Beckers, T., & Vandorpe, S. (2005). Evidence for the role of higher-order reasoning processes in cue competition and other learning phenomena. *Learning & Behavior*, 33, 239-249.
- de Kort, S. R., Dickinson, A., & Clayton, N. S. (2005). Retrospective cognition by food-caching Western scrub-jays. *Learning & Motivation*, 36, 159-176.
- Ferkin, M. H., Combs, A., delBarco-Trillo, J., Pierce, A. A., & Franklin, S. (2008). Meadow voles, *Microtus pennsylvanicus*, have the capacity to recall the 'what', 'where', and 'when' of a single past event. *Animal Cognition*, 11, 147-159.
- Foote, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, 17, 551-555.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gibbon, J. (1991). Origins of scalar timing. *Learning & Motivation*, 22, 3-38.
- Gibbon, J., Fairhurst, S., & Goldberg, B. (1997). Cooperation, conflict and compromise between circadian and interval clocks in pigeons. In C. M. Bradshaw & E. Szabadi (Eds.), *Time and behaviour: Psychological and neurobehavioural analyses* (pp. 329-384). New York: Elsevier.
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences of the United States of America*, 98, 5359-5362.
- Johnson, M. K., Hashtroudi, S., & Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114, 3-28. doi: 10.1037/0033-2909.114.1.3
- Kaufman, A., & Baron, A. (1968). Suppression of behavior by timeout punishment when suppression results in loss of positive reinforcement. *Journal of the Experimental Analysis of Behavior*, 11, 595-607.
- Melcer, T., & Timberlake, W. (1985). Poison avoidance and patch (location) selection in rats. *Animal Learning & Behavior*, 13, 60-68.
- Mitchell, K. J., & Johnson, M. K. (2009). Source monitoring 15 years later: What have we learned from fMRI about the neural mechanisms of source memory? *Psychological Bulletin*, 135, 638-677. doi: 10.1037/a0015849
- Morgan, C. L. (1906). *An introduction to comparative psychology*. London: W. Scott.
- Naqshbandi, M., Feeney, M. C., McKenzie, T. L. B., & Roberts, W. A. (2007). Testing for episodic-like memory in rats in the absence of time of day cues: Replication of Babb and Crystal. *Behavioural Processes*, 74, 217-225.
- Olton, D. S., & Samuelson, R. J. (1976). Remembrance of places passed: Spatial memory in rats. *Journal of Experimental Psychology: Animal Behavior Processes*, 2, 97-116.
- Pahl, M., Zhu, H., Pix, W., Tautz, J., & Zhang, S. (2007). Circadian timed episodic-like memory - A bee knows what to do when, and also where. *Journal of Experimental Biology*, 210, 3559-3567.
- Pizzo, M. J., & Crystal, J. D. (2004). Time-place learning in the eight-arm radial maze. *Learning & Behavior*, 32, 240-255.
- Pizzo, M. J., & Crystal, J. D. (2007). Temporal discrimination of alternate days in rats. *Learning & Behavior*, 35, 163-168.
- Raby, C. R., & Clayton, N. S. (2009). Prospective cognition in animals. *Behavioural Processes*, 80, 314-324. doi: 10.1016/j.beproc.2008.12.005
- Richardson, J. V., & Baron, A. (2008). Avoidance of timeout from response-independent food: Effects of delivery rate and quality. *Journal of the Experimental Analysis of Behavior*, 89, 169-181.
- Roberts, W. A. (1998). *Principles of animal cognition*. Boston: McGraw-Hill.

- Roberts, W. A., Feeney, M. C., MacPherson, K., Petter, M., McMillan, N., & Musolino, E. (2008). Episodic-like memory in rats: Is it based on when or how long ago? *Science*, *320*, 113-115. doi: 10.1126/science.1152709
- Rumbaugh, D. M., Richardson, W. K., Washburn, D. A., Savage-Rumbaugh, E. S., & Hopkins, W. D. (1989). Rhesus monkeys (*Macaca mulatta*), video tasks, and implications for stimulus-response spatial contiguity. *Journal of Comparative Psychology*, *103*, 32-38.
- Shettleworth, S. J. (1998). *Cognition, evolution, and behavior*. New York: Oxford University Press.
- Shettleworth, S. J. (2010). Clever animals and killjoy explanations in comparative psychology. *Trends in Cognitive Sciences*, *14*, 477-481. doi: 10.1016/j.tics.2010.07.002
- Shields, W. E., Smith, J. D., Guttmanova, K., & Washburn, D. A. (2005). Confidence judgments by humans and rhesus monkeys. *Journal of General Psychology*, *132*, 165-186.
- Shields, W. E., Smith, J. D., & Washburn, D. A. (1997). Uncertain responses by humans and rhesus monkeys (*Macaca mulatta*) in a psychophysical same-different task. *Journal of Experimental Psychology: General*, *126*, 147-164.
- Singer, R. A., & Zentall, T. R. (2007). Pigeons learn to answer the question 'where did you just peck?' and can report peck location when unexpectedly asked. *Learning & Behavior*, *35*, 184-189.
- Smith, J. D., Beran, M. J., Couchman, J. J., & Coutinho, M. V. C. (2008). The comparative study of metacognition: Sharper paradigms, safer inferences. *Psychonomic Bulletin & Review*, *15*, 679-691.
- Smith, J. D., Beran, M. J., Redford, J. S., & Washburn, D. A. (2006). Dissociating uncertainty responses and reinforcement signals in the comparative study of uncertainty monitoring. *Journal of Experimental Psychology: General*, *135*, 282-297.
- Smith, J. D., Shields, W. E., Allendoerfer, K. R., & Washburn, D. A. (1998). Memory monitoring by animals and humans. *Journal of Experimental Psychology: General*, *127*, 227-250.
- Smith, J. D., Shields, W. E., Schull, J., & Washburn, D. A. (1997). The uncertain response in humans and animals. *Cognition*, *62*, 75-97.
- Smith, J. D., Shields, W. E., & Washburn, D. A. (2003). The comparative psychology of uncertainty monitoring and metacognition. *Behavioral and Brain Sciences*, *26*, 317-373.
- Takahashi, J. S., Turek, F. W., & Moore, R. Y. (2001). *Handbook of behavioral neurobiology: Circadian clocks* (Vol. 12). New York: Plenum.
- Terrace, H. S., & Metcalfe, J. (2005). *The missing link in cognition: Origins of self-reflective consciousness*. New York: Oxford University Press.
- Terrace, H. S., & Son, L. K. (2009). Comparative metacognition. *Current Opinion in Neurobiology*, *19*, 67-74.
- Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of memory* (pp. 381-403). New York: Academic Press.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.
- Tulving, E. (1985). How many memory systems are there? *American Psychologist*, *40*, 385-398.
- Tulving, E. (2001). Chronesthesia: Awareness of subjective time. In D. T. Stuss & R. C. Knight (Eds.), *The age of the frontal lobes* (pp. 311-325). New York: Oxford University Press.

- Tulving, E. (2005). Episodic memory and autoeogenesis: Uniquely human? In H. Terrace & J. Metcalfe (Eds.), *The missing link in cognition: Origins of self-reflective consciousness* (pp. 3-56). New York: Oxford University Press.
- Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science*, *15*, 307-311. doi: 10.1111/j.1467-8721.2006.00458.x
- Washburn, D. A., & Rumbaugh, D. M. (1992). Testing primates with joystick-based automated apparatus: Lessons from the Language Research Center's computerized test system. *Behavior Research Methods, Instruments & Computers*, *24*, 157-164.
- Zentall, T. R. (2005). Animals may not be stuck in time. *Learning & Motivation*, *36*, 208-225.
- Zentall, T. R. (2006). Mental time travel in animals: A challenging question. *Behavioural Processes*, *72*, 173-183.
- Zentall, T. R., Clement, T. S., Bhatt, R. S., & Allen, J. (2001). Episodic-like memory in pigeons. *Psychonomic Bulletin & Review*, *8*, 685-690.
- Zhang, S., Schwarz, S., Pahl, M., Zhu, H., & Tautz, J. (2006). Honeybee memory: A honeybee knows what to do and when. *Journal of Experimental Biology*, *209*, 4420-4428.
- Zhou, W., & Crystal, J. (2011). Validation of a rodent model of episodic memory. *Animal Cognition*, *14*, 325-340. doi: 10.1007/s10071-010-0367-0
- Zhou, W., & Crystal, J. D. (2009). Evidence for remembering when events occurred in a rodent model of episodic memory. *Proceedings of the National Academy of Sciences of the United States of America*, *106*, 9525-9529. doi: 10.1073/pnas.0904360106