

IN PRESS: Psychological Science

Running Head: METACOGNITION IN MONKEYS

Transfer of Metacognitive Skills and Hint Seeking in Monkeys

Nate Kornell
University of California, Los Angeles

Lisa K. Son
Barnard College

Herbert S. Terrace
Columbia University and New York State Psychiatric Institute

Abstract

Metacognition is knowledge that can be expressed as confidence judgments about what we know (monitoring) and by strategies for learning what we don't know (control). Although a substantial literature exists on cognitive processes in animals, little is known about their metacognitive abilities. Here we show that rhesus macaques, trained previously to make *retrospective* confidence judgments about their performance on perceptual tasks, transferred that ability immediately to a new perceptual task and to a working memory task. In a second experiment we show that monkeys can also learn to request "hints" when they are given problems that they would otherwise have to solve by trial and error. This shows, for the first time, that non-human primates share with humans the ability to monitor and transfer their metacognitive ability both within and between different cognitive tasks, and to seek new knowledge on a need to know basis.

Transfer of Metacognitive Skills and Hint Seeking in Monkeys

Metacognition, the ability to monitor and control one's knowledge, is a basic feature of human thought (Flavell & Wellman, 1977; Koriat, 1997; Nelson & Narens, 1990). Metacognitive *monitoring* has been studied in experiments in which subjects are asked to state some fact and are then asked to evaluate the accuracy of their response. For example, after responding to a question such as, "What is the capital of Australia?" a subject might be asked, "On a scale of 0-100, rate how confident you are that your answer is correct". In experiments on metacognitive *control*, subjects are typically allowed to determine how much time they spend studying each item during a memory task (Dunlosky & Hertzog, 1998; Metcalfe & Kornell, 2005; Son, 2004). The usual finding, an inverse relationship between accuracy and study time, shows that subjects use their knowledge of what they don't know to determine which items to study.

Although there is a substantial body of research on animal cognition (Shettleworth, 1998), there is very little evidence of metacognitive monitoring and virtually none on metacognitive control processes. This article presents evidence that monkeys trained to make metacognitive judgments about perceptual problems learned to do so in a way that was not task specific. They were able to transfer their metacognitive skills immediately to a qualitatively different working memory task. Experiment 2 shows that monkeys can use metacognitive control to improve their performance on a serial memory task by seeking the information they need while learning that task.

A major difficulty in studying animal metacognition, aside from the need to use non-verbal paradigms, is the paucity of paradigms that can rule out the influence of exteroceptive stimuli in purported instances of metacognition. Some recent psychophysical experiments on

monkeys (Shields, Smith, & Washburn, 1997) and dolphins (Smith, Schull, Strote, McGee, Egnor, & Erb, 1995) illustrate this problem. Subjects were required to make one response (“high”) on trials when one stimulus (S1) was presented and another (“low”) when a second stimulus (S2) was presented. Subjects were also given the option of ending each trial when they were “unsure” of the correct response. Since subjects opted out mainly on trials on which the difference between S1 and S2 was small, opting out was interpreted as evidence of the subject’s uncertainty. However, just as S1 and S2 exerted stimulus control over “high” and “low” responses, S1 and S2 could also exert a similar influence over the opting out response. As noted by Metcalfe (2003) and Shettleworth and Sutton (2003), when a behavior occurs in the presence of the discriminative stimuli to which the subject has been trained, there is no reason to interpret it as metacognitive unless it can be shown to transfer to a new task.

The problem of exteroceptive stimulus control can be obviated by requiring subjects to make *retrospective* judgments of their *memories*. In one experiment, rhesus macaques were given the opportunity to choose between two types of trials on a delayed-matching-sample task *after* the sample was presented but *before* the actual test could be administered (Hampton, 2001). One option was to wait for the test, at which time the subject could earn a desirable reward if it responded correctly, or earn nothing if it made an error. The second option was to end the trial and settle for a less desirable, but certain, reward. Subjects were also given forced-choice trials, on which they could not avoid the test phase. Accuracy was significantly higher on trials on which subjects opted to take the test than on forced choice trials, suggesting that they declined to take the test when their memory of the sample was relatively weak.

In a different paradigm, which we also used in the current experiment, monkeys were trained to make retrospective judgments of their accuracy on perceptual tasks (Son & Kornell,

2005; Son, Kornell, Terrace, Sussan, & Flaherty, 2004). Subjects were differentially rewarded for selecting one of two confidence icons (high or low) after each response on a given perceptual task. Reward was maximized when subjects chose the high confidence following correct responses and low confidence following errors. Subjects trained on this paradigm learned to respond to the confidence icons appropriately¹.

Task-specificity is another issue that needs to be addressed in experiments on metacognition in animals. In contrast to judgments of confidence and uncertainty by human subjects, which have been observed in many contexts (e.g. Schwartz, 2002), it has yet to be shown that purported instances of metacognitive behavior in animals transfers to novel types of problems (Shettleworth & Sutton, 2003). Experiment 1 shows, for the first time, that a monkey can transfer its ability to judge the accuracy of its performance from a perceptual task to a working memory task.

Experiment 1

Method

Subjects. The subjects were two male rhesus macaques (*Macaca mulatta*), Lashley and Ebbinghaus. Both were approximately 6 years of age at the start of the experiment and had been trained previously in experiments on simultaneous chaining (Terrace, 2005a) and numerical discrimination (Brannon & Terrace, 1998) and in pilot studies on metacognition (Son & Kornell, 2005; Son et al., 2004). Subjects were housed individually in adjoining cages in a colony of 20 rhesus macaques at the New York Psychiatric Institute (NYSPI). The colony was maintained in accordance with NIH guidelines and those issued by the Institutional Animal Care and Use Committees at NYSPI and Columbia University.

Apparatus. The apparatus was identical to that described in earlier experiments on monkey cognition (Subiaul, Cantlon, Holloway & Terrace, 2004). All training and testing sessions took place in chambers that were housed in sound attenuated booths. Each chamber was equipped with a touch-sensitive video monitor that was used to present stimuli and to detect subjects' responses.

Pilot experiment. Experiment 1 is an extension of a pilot experiment in which we determined optimal values of the parameters needed to train reliable metacognitive performance (Son & Kornell, 2005). During the pilot experiment, subjects were trained on two perceptual tasks: a line-discrimination task, on which they were required to select the longest of nine lines, and a numerical discrimination task, on which one subject (Lashley) was required to select the stimulus containing the largest number of geometrical objects, and one (Ebbinghaus) the fewest. The metacognitive paradigm (described below) was added after subjects learned the line discrimination task. Both subjects learned to respond metacognitively. Subjects were subsequently trained on the numerical discrimination task. After that task was mastered, the metacognitive paradigm was added without further metacognitive training. Both subjects responded metacognitively immediately after the metacognitive paradigm was added.

Metacognitive paradigm. The metacognitive paradigm used in this experiment is similar to that used in the pilot experiment. As shown in Figure 1, two confidence icons were presented at the end of each trial, immediately after the subject responded on a perceptual task. One icon signified high confidence; the other, low confidence. Reward was contingent upon the *relation* between the accuracy of the subject's response on the perceptual task and its choice of confidence icon. A response to the high confidence icon was a "risky" bet. The subject won 3

tokens if its response on the perceptual task was correct, but lost 3 tokens if it was incorrect.

Responses to the “safe” low confidence icon were always rewarded with 1 token.

Tokens were deposited in a “bank” located in the lower right hand corner of the video monitor. The bank contained 9 tokens at the start of each session. Tokens were added to and subtracted from the bank following each metacognitive response. Food reward (a 190-mg banana pellet) was provided when 12 or more tokens had accumulated in the bank. The number of tokens was then reset to 9. If the balance dropped to zero, no further tokens could be lost. We note that this appears to be the first experiment with animals in which tokens were *subtracted* following incorrect responses. In earlier experiments with token economies, incorrect responses did not incur any penalties (Wolfe, 1936; Jackson & Hackenberg, 1996).

Design of Experiment 1. In experiments on human metacognition, subjects don’t require any pre-training when asked to make metacognitive judgments on a new task. Having established the efficacy of the metacognitive paradigm in our pilot study, we sought to determine whether a monkey could transfer its metacognitive skills to a new perceptual problem and to a qualitatively different memory task. Our measure of transfer was based on subjects’ *immediate* performance during the first 4 sessions in which the metacognitive paradigm was added to each new task.

Transfer to a new perceptual task. Subjects were trained on a perceptual task based on circle size. The level of difficulty varied randomly on each trial. One subject (Ebbinghaus) was trained to select the circle that had the largest area; the other (Lashley), the smallest. Trials were initiated by the subject when it touched a “start” stimulus on the monitor. Correct responses were followed by the simultaneous delivery of primary reward (two 190-mg banana-flavored Bio-Serv© food pellets) and visual and auditory feedback that informed the subject that its response

was detected by the monitor. Each session lasted 20 minutes. The metacognitive paradigm was introduced once each subject completed at least 65% of the trials in one session correctly.

Transfer to a serial working memory task. Following metacognitive testing on the size discrimination problem, subjects were trained on a working memory task that was similar to those used in experiments on serial working memory (SWM) with human subjects (Sternberg, 1969) and monkeys (Wright, Santiago, Sands, Kendrick, & Cook, 1985). The trial structure of the SWM task is shown in Figure 2. A sequence of 6 trial-unique photographs was presented during the study portion of each trial. To maintain attention to the sample, subjects were required to touch each photograph before the next one was presented. Responses to each of the first 5 samples were followed by a 2-s inter-stimulus-interval. The test portion of each trial began .5 s after the subject responded to the 6th sample. Nine photographs were displayed simultaneously during test, one of which had appeared previously as a sample. The remaining 8 photographs were distractors. The subject's task was to select the sample. Correct responses were rewarded with two banana pellets. A response to any of the distractors ended the trial without reward. The samples and distractors were selected at random from a library of over 2,500 photographs. A given photograph was never repeated during the same session. When subjects satisfied the accuracy criterion of 65% correctly completed trials/session, the two metacognitive icons were introduced at the end of each trial, without any additional metacognitive training. The contingencies for responding to the high- and the low-confidence icons were the same as those used on the size task.

Results

Evidence of metacognition in a non-verbal animal requires a correlation between task accuracy and a metacognitive judgment. The measure we used was the phi-correlation

coefficient (ϕ) between accuracy and confidence. The value of ϕ is the same as the value of Pearson's r when r is calculated from data arranged in a 2 x 2 table, in this instance, correct vs. incorrect X high vs. low risk. The value of ϕ approaches 1.0 to the extent that subjects chose the high confidence icon after responding correctly and the low confidence icon after making an error.

Perceptual task. Both subjects responded metacognitively as soon as the risk icons were introduced to the perceptual task. This means that they responded more frequently to the high confidence icon after correct responses than after errors, and that they responded more frequently to the low confidence icon after errors than after correct responses. The value of ϕ for the first four days of responding on the new perceptual task is shown in Figure 3, along with the 2 x 2 tables from which each value of ϕ was derived. In each instance, the value of ϕ was significant beyond the .01 level (as computed by Fisher's exact test)². Given that no additional metacognitive training was provided during the area tasks, these data provide clear evidence that an animal can transfer its metacognitive ability from one perceptual task to another.

Memory task. Metacognitive responding also transferred immediately when the confidence icons were added to the SWM task. As shown in Figure 3, there was no decrement in the value of ϕ , as compared to its value on the area discrimination task, the most recent perceptual task. For both subjects, the value of ϕ was significant beyond the .001 level³. This result is especially striking because of the long delays that occurred between the subject's responses to sample items and the presentation of the confidence icons. In the SWM task, the median duration of the interval between Ebbinghaus and Lashley's responses to sample 1 and the appearance of the confidence icons was 22 and 23 s, respectively. For sample 6, it was 4 and 5 s, respectively.

Serial position of probe. On most serial tasks, the difficulty of a list item is correlated with its serial position (primacy and the recency effects). Only a recency effect was obtained here. Accuracy increased significantly with serial position for both monkeys (Ebbinghaus, $r = .14$, $p < .001$; Lashley $r = .12$, $p < .01$). By contrast, the position of the sample had no effect on the value of ϕ . Indeed, the values of the ϕ differed significantly from zero at each serial position ($p < .05$), with the sole exception of Ebbinghaus at serial position 6. These data show that difficulty, as defined by an item's serial position, cannot account for subjects' confidence in the accuracy of their responses.

Influence of reaction time. It might be argued that subjects' responses to the confidence icons were based on the difference between their reaction times (RT) on correct and incorrect responses. In experiments with human subjects (e.g. Benjamin, Bjork & Schwartz, 1994), RTs of correct responses are typically shorter than RTs of errors. The same was true on the SWM task: There was a negative correlation (r) between RT and accuracy during the first ten sessions of metacognitive training on the SWM task (Ebbinghaus $r = -.11$, $t(9) = -2.05$, $p = .07$; Lashley $r = -.40$, $t(9) = -9.77$, $p < .0001$), confirming that RT decreased as accuracy increased.

There is no simple method for controlling for possible influences of RT duration on confidence judgments. It is, however, possible to partial out the contribution of RT ($r_{RT,acc}$) to the correlation between accuracy and confidence choice ($r_{a,c}$) and to then compare the value of $r_{a,c}$ with zero during each session. The values of $r_{a,c}$ during the first ten sessions of metacognitive training on the SWM task were significantly greater than zero for both subjects (Ebbinghaus $r = .28$, $t(9) = 5.74$, $p < .001$; Lashley $r = .24$, $t(9) = 6.62$, $p < .0001$). It follows that

the duration of RTs on correct and incorrect trials is not the sole determinant of subjects' metacognitive performance.

Experiment 2

Experiment 1 showed that a monkey can *monitor* its accuracy on perceptual tasks, and transfer that ability to monitoring its memory. In Experiment 2 we asked, can a monkey acquire a second component of human metacognition, *control* of one's knowledge? Previous studies (Call & Carpenter, 2001; Hampton, Zivin, & Murray, 2004) showed that apes and monkeys sought information by looking for food when they were uncertain about its location, but that they reached for it directly if they saw where it was hidden. However, similar behavior occurs naturally during foraging. In certain contexts, animals search for food by default unless they can acquire it directly. It is therefore possible that the animals in question simply used a familiar food-seeking strategy in a new context. They didn't have to know what they knew. Instead they simply needed to know where the food was. If not, they searched for it.

The task used in Experiment 2 required subjects to learn novel 4-item sequences composed of arbitrary photographs. Subjects were given the opportunity to ask for "hints" as to the next response in the sequence. Without "hints", subjects would have to learn each sequence by trial and error, a skill they acquired in a previous experiment (Terrace, Son, & Brannon, 2003). Thus, Experiment 2 addressed two questions: would subjects learn to request hints when starting to learn a new list, and would hint seeking decrease as they became more proficient at executing the list?

Method

Subjects. The subjects were two 10-year old male rhesus macaques (Macduff and Oberon). Both monkeys had extensive experience at learning 3-, 4- and 7-item lists as simultaneous chains (Terrace et al, 2003).

Apparatus. The apparatus was identical to that used in Experiment 1.

Simultaneous chaining paradigm. Under the simultaneous chaining paradigm, all list items are displayed simultaneously throughout each trial, typically on a touch-sensitive video monitor (Figure 4a). The subject's task is to respond to each item in a particular order, regardless of its spatial position on the monitor (Figure 4b). The physical location of list items is randomized on each trial to prevent subjects from performing the required sequence as a fixed-motor pattern. All of the lists used in Experiment 2 were composed of photographs (1.5" x 2") of natural or man-made objects (e.g., animals, people, scenery, flowers, cars, bridges, etc.). All of the stimuli were novel and none were repeated when composing new lists.

Simultaneous chains are typically learned by trial and error from feedback that follows each response Terrace (2005a). Correct responses were followed by brief (0.5-s) visual and auditory feedback; errors, by a 4-s time out during which the screen was darkened. A trial ended either when the subject responded to an item incorrectly or when the subject responded to all of the items in the correct order. Food reward (a 190-mg banana pellet) was provided only after a correct response to the last list item.

Training simultaneous chains with a hint option. As shown in Figures 5c and 5d, a hint could be obtained by responding to an icon to the right of the list items. The hint was four blinking lines that appeared around the item to which the subject should next respond. If the subject completed a trial correctly *without* requesting any hints, it was rewarded with a highly desirable M&M candy. If the subject requested one or more hints during the course of

completing a trial correctly, it received the usual banana pellet as a reward. Subjects could obtain as many as 4 hints on each trial, one for each list item.

Subjects were pre-trained on a familiar list for 6 days with the hint-icon present on 50% of the trials⁴. Each subject was then trained on 10 new lists, 4 sessions on each list, 60 trials/session. The familiar list was presented on half of the trials, and the new list on the other half. During each session, the subject could request a hint on half of the trials. Thus, out of 60 trials, a hint was available on 15 trials on which a new list was presented and on 15 trials on which the familiar list was presented. No hint was available on the remaining 30 trials: 15 with the new list and 15 with the familiar list. Trial types alternated randomly.

Results

Hint seeking was inversely related to accuracy on each of the 10 novel lists, for Macduff ($r = -.74, p < .0001$) and Oberon ($r = -.87, p < .0001$). As can be seen in Figure 5a, requests for hints, when a hint was available (dashed function), decreased across the four days of training from 78% to 53%. At the same time, the proportion of correct responses increased from 18% to 37% when a hint was not available (solid function).

The inverse relation between hint seeking and accuracy was also evident in the subjects' performance on new and familiar lists, for Macduff ($r = -.74, p < .0001$) and Oberon ($r = -.94, p < .0001$). As can be seen in Figure 5b, the requests for hints when a hint was available were less frequent on the familiar list (13%) than on new lists (65%). In contrast, proportion correct was higher on the familiar list (83%) than on the new lists (30%).

General Discussion

Because experiments on metacognition rely on verbal paradigms and because metacognitive abilities have not been observed in children until the age of 4, some researchers

have argued that metacognition requires verbal ability (Metcalf & Shimamura, 1994; Tulving, 1994). Using behavioral measures, Experiment 1 showed that monkeys can learn to make metacognitive judgments that are functionally similar to those obtained from human subjects. Monkeys that learned to monitor their performance on a series of perceptual tasks transferred that ability to a qualitatively different task based on working memory. That result demonstrates that a monkey's metacognitive ability is not task specific. It is also the first demonstration that an animal can learn to respond metacognitively about the match (or lack thereof) between its memory of a sample and the item it selected during test. Because the monkey's response was based on its memory of stimuli that were no longer present, its choice of confidence icon cannot be attributed to the physical characteristics of any exteroceptive stimulus. Experiment 2 provided the first evidence to date that a monkey can apply appropriate control strategies to correct deficiencies in its knowledge. When it lacks the information it needs to execute new lists, it acquired that information by asking for a hint. Taken together, the results of Experiments 1 and 2 show that a monkey's metacognitive behavior provides analogues of the two basic features of human metacognition: monitoring and controlling knowledge.

In the absence of verbal ability, there is no reason to assume that monkeys are conscious of their metacognitive judgments, or for that matter, of the metacognitive states of other individuals (Povinelli & Vonk, 2003). Although humans are conscious of many of their metacognitive processes, they often make metacognitive judgments before they can be processed consciously (Cary & Reder, 2002; Reder & Schunn, 1996; Son & Metcalfe, 2005). Thus, non-verbal primates appear to share with humans the ability to make unconscious metacognitive judgments. Rhesus macaques can, across different tasks, differentiate between the cognitive states of knowing and not knowing. That ability adds to the growing list of cognitive processes

that evolved without language in which animals learn to solve problems by relying on representations of previously experienced events (Kinsbourne, 2005; Terrace, 2002; Terrace, 2005b).

References

- Benjamin, A. S, Bjork, R. A., & Schwartz, B. L. (1998). The mismeasure of memory: When retrieval fluency is misleading as a metamnemonic index. *Journal of Experimental Psychology: General*, 127, 55-68.
- Brannon, E. M., & Terrace, H. S. (1998). Ordering of the numerosities 1 to 9 by monkeys. *Science*, 282, 746-749.
- Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, 4, 207-220.
- Cary, M. & Reder, L. M. (2002). Metacognition in Strategy Selection: Giving Consciousness Too Much Credit. In M. Izaute, P. Chambres, & P.J. Marescaux (Eds.), *Metacognition: Process, Function, and Use* (pp. 63-78). New York, NY: Kluwer.
- Dunlosky, J., & Hertzog, C. (1998). Training programs to improve learning in later adulthood: Helping older adults educate themselves. In D. J. Hacker, J. Dunlosky, & A. C. Graesser (Eds.), *Metacognition in educational theory and practice* (pp. 249–275). Mahwah, NJ: Lawrence Erlbaum Associates.
- Dunlosky, J., & Nelson, T. O. (1992). Importance of the kind of cue for judgments of learning (JOL) and the delayed-JOL effect. *Memory & Cognition*, 20, 373-380.
- Flavell, J. & Wellman, H. (1977). Metamemory. In R. Kail, & J. Hagen (Eds.), *Perspectives on the Development of Memory and Cognition* (pp. 3-33). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, 98, 5359-5362.

- Hampton, R. R., Zivin, A., & Murray, E. A. (2004). Rhesus Monkeys (*Macaca mulatta*) Discriminate Between Knowing and Not Knowing and Collect Information As Needed Before Acting. *Animal Cognition*, 7, 239-254.
- Jackson, K., & Hackenberg, T. D. (1996). Token reinforcement, choice, and self-control in pigeons. *Journal of the Experimental Analysis of Behavior*, 66, 29-49.
- Kinsbourne, M. (2005) A continuum of self-consciousness that emerges in phylogeny and ontogeny. In H. S. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness* (pp. 142-156). New York: Oxford University Press.
- Koriat, A. (1997). Monitoring one's own knowledge during study: A cue-utilization approach to judgments of learning. *Journal of Experimental Psychology: General*, 126, 349-370.
- Metcalfe, J. (2003) Drawing the line on metacognition. *Behavioral and Brain Sciences*, 26, 350-351.
- Metcalfe, J., & Kornell, N. (2005). A region of proximal learning model of study time allocation. *Journal of Memory and Language*, 52, 463-477.
- Metcalfe, J., & Shimamura, A. P. (1994). *Metacognition: knowing about knowing*. Cambridge, MA: MIT Press.
- Nelson, T. O. (1984). A comparison of current measures of the accuracy of feeling-of-knowing predictions. *Psychological Bulletin*, 95, 109-133.
- Nelson, T. O., & Dunlosky, J. (1991). When people's judgments of learning (JOLs) are extremely accurate at predicting subsequent recall: The "delayed-JOL effect." *Psychological Science*, 2, 267-270.

- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 26, pp. 125–141). New York: Academic Press.
- Povinelli, D. J., & Vonk, J. (2003) Chimpanzee minds: Suspiciously human? *Trends in Cognitive Sciences*, 7, 157-160.
- Reder, L. M., & Schunn, C. (1996). Metacognition does not imply awareness: Strategy choice is governed by implicit learning and memory. In L.M. Reder (Ed.), *Implicit memory and metacognition* (pp. 45-77). Mahwah, NJ: Erlbaum.
- Schwartz, B. L. (2002). *Tip-of-the-tongue states: Phenomenology, mechanism, and lexical retrieval*. NJ: Erlbaum.
- Shettleworth, S. J., & Sutton, J. E. (2003). Metacognition in animals: It's all in the methods. *Behavioral and Brain Sciences*, 23, 353-354.
- Shettleworth, S. J. (1998) *Cognition, evolution, and behavior*. New York: Oxford University Press.
- 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.
- Smith, J. D., Schull, J., Strote, J., McGee, K., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (*Tursiops truncatus*). *Journal of Experimental Psychology: General*, 124, 391-408.

- Son, L. K. (2004). Spacing one's study: evidence for a metacognitive control strategy. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 601-604.
- Son, L. K., & Kornell, N. (2005). Meta-confidence judgments in rhesus macaques: Explicit versus implicit mechanisms. In H. S. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness* (pp. 296-320) New York: Oxford University Press.
- Son, L. K., Kornell, N., Terrace, H. S., Sussan, D., & Flaherty, M. (2004). Measuring confidence judgments non-verbally by using a betting paradigm. Paper presented at the *Annual Meeting of Comparative Cognition*. Melbourne Beach, Florida, March 25-27.
- Son, L. K., & Metcalfe, J. (2005). Judgments of Learning: Evidence for a Two-Stage Model. *Memory & Cognition*, 33, 1116-1129.
- Sternberg, S. (1969). Memory scanning: Mental processes revealed by reaction time experiments. *American Scientist*, 57, 421-457.
- Subiaul, F. X., Cantlon, J. F., Holloway, R., & Terrace, H. S. (2004). Cognitive imitation in rhesus macaques. *Science*, 305, 407-410.
- Terrace, H. S. (2002). Serial Expertise and the Evolution of Language. In J. H. A. Wray, & F. J. Newmeyer (Ed.), *The Transition to Language* (pp. 64-90). New York: Oxford University Press.
- Terrace, H. S. (2005a) The simultaneous chain: A new approach to serial learning. *Trends in Cognitive Sciences*, 9, 202-210.
- Terrace, H. S. (2005b) Metacognition and the Evolution of Language. In H. S. Terrace & J. Metcalfe (Eds.), *The Missing Link in Cognition: Origins of Self-Reflective Consciousness* (pp. 84-115) New York: Oxford University Press.

Terrace, H. S., Son, L. K., & Brannon, E. M. (2003). Serial expertise of rhesus macaques.

Psychological Science, 14, 66-73.

Tulving, E. (1994). Foreword of J. Metcalfe & A. P. Shimamura (Eds.), *Metacognition: knowing*

about knowing. Cambridge, MA: MIT Press.

Vesonder, G., and J. Voss. (1985). On the ability to predict one's own responses while learning.

Journal of Memory and Language, 24, 363-376.

Wolfe, J. B. (1936). Effectiveness of token-rewards for chimpanzees. *Comparative Psychology*

Monographs, 12, 1-72.

Wright, A. A., Santiago, H. C., Sands, S. F., Kendrick, D. F., and Cook, R. G. (1985). Memory

processing of serial lists by pigeons, monkeys, and people. *Science, 229*, 287-289.

Author Note

Nate Kornell, Department of Psychology, University of California, Los Angeles, Lisa K. Son, Department of Psychology, Barnard College, and Herbert S. Terrace, Department of Psychology, Columbia University and New York State Psychiatric Institute.

This research was supported by a grant from NIMH to HST (MH 40462). We thank Jessica Cantlon, Tamar Kornblum, Dustin Merritt and Tammy Moscrip for their assistance in this project, and Peter Balsam, Janet Metcalfe, Allen Neuringer and Robert Remez for their helpful comments about earlier drafts of this article.

The authors declare no competing financial interests. Correspondence and requests for materials should be addressed to Herbert S. Terrace, 406 Schermerhorn Hall, Columbia University, 1190 Amsterdam Avenue, New York, NY 10027 (terrace@columbia.edu).

Notes

¹ A similar experiment was subsequently performed with monkeys by Shields, Smith, Guttmanova, & Washburn (2005).

² The values of ϕ we obtained compare favorably with those reported in experiments in which adult human subjects were asked to make metacognitive judgments immediately after learning a set of paired associates (Nelson & Dunlosky, 1991). In those studies, in which metacognition was evaluated by the magnitude of Gamma correlations, the values of Gamma were typically low; mean = .38 (Dunlosky & Nelson, 1992), range; .09-.48 (Vesonder & Voss, 1985). ϕ is a more conservative estimate of metacognitive performance than Gamma because Gamma can return values that are higher, but never lower, than the value of ϕ (Nelson, 1984).

³ Video clips of Ebbinghaus performing the SWM task can be found at <http://www.columbia.edu/cu/psychology/primatecognitionlab/videos/MetaSPR.mov>.

⁴ The percentage of trials completed correctly when the hint was not available during the last day of training on the familiar list was 70% and 87% for Macduff and Oberon, respectively.

Figure Captions

Figure 1. Metacognitive task. The high- and low-confidence icons appeared immediately after the subject made its selection on the perceptual task. A response to the high confidence icon resulted in a gain of 3 tokens after a correct response (top), but a loss of 3 tokens after an error (middle). A response to the low-confidence icon always resulted in a gain of one token (bottom).

Figure 2. Trial structure of the serial working memory task. Six sample photographs were displayed successively, followed by a test during which one of the 6 samples was presented with 8 distractors. The subject's task was to select the sample and then select the appropriate confidence icon.


Figure 3. ϕ -correlations. The value of ϕ for the first 4 days of training on the perceptual and serial working memory task. The absolute frequencies based on which each value of ϕ was computed are shown in the 2 x 2 tables shown in each bar.

Figure 4. Examples of displays used in Experiment 2. Subjects were rewarded for responding to all items in a particular order ($A \rightarrow B \rightarrow C \rightarrow D$). On 50% of trials, the hint was not available [a & b]. Panels a & b illustrate that the position of the items varied randomly on each trial. Hints were available on trials on which an icon appeared on the right side of the monitor, above a red circle [c & d]. Touching the hint icon resulted in the appearance of 4 blinking lines around the correct item for that position in the sequence. The red circle below the hint icon signalled that the subject could earn an M&M, instead of a less desirable banana pellet, by completing the trial correctly. The first request for a hint eliminated the red circle.

Figure 5. Proportion of trials completed correctly when no hint was available, and proportion of trials on which a hint was requested when the hint was available. [a] Across the four days of training on new lists. [b] On new lists and on familiar list.

Figure 1

a



b




Figure 2




Figure 3




Figure 4




Figure 5

