

Rhesus Monkeys (*Macaca mulatta*) Adaptively Adjust Information Seeking in Response to Information Accumulated

Hsiao-Wei Tu, Alex A. Pani, and Robert R. Hampton
Emory University

Metacognition consists of monitoring and control processes. Monitoring has been inferred when nonhumans use a “decline test” response to selectively escape difficult test trials. Cognitive control has been inferred from selective information-seeking behavior by nonhumans ignorant of needed knowledge. Here we describe a computerized paradigm that extends previous work and assesses dynamic interactions between monitoring and control. Monkeys classified images as containing birds, fish, flowers, or people. To-be-classified images were initially masked, and monkeys were trained to gradually reveal the images by touching a “reveal button.” Monkeys could choose to classify images at any time or to reveal more of the images. Thus, they had the opportunity to assess when enough of an image had been revealed to support accurate classification. In Experiment 1, monkeys made more reveal responses before classifying when smaller amounts of the image were revealed by each button touch. In Experiment 2, to-be-classified images were shrunk and covered by 1 critical blocker among other blockers that did not provide information when removed. Monkeys made more reveal responses as the critical blocker was removed later in the trial. In Experiment 3, monkeys were presented with previously classified images with either more or fewer blockers obscuring the image than was the case when they chose to classify that image previously. Monkeys made more reveal responses when information was insufficient than when it was excessive. These results indicate that monkeys dynamically monitor evolving decision processes and adaptively collect information as necessary to maintain accuracy.

Keywords: metacognition, information seeking, monitoring, behavioral control, categorization

Humans often engage in metacognition, or thinking about thinking, to adaptively modulate cognitive processing. For example, when preparing for an exam, a student may introspectively evaluate what she knows well already so that she can allot more time to studying material she does not currently know as well. Such metacognition involves feedback between *monitoring*, which assesses the state of ongoing cognition, and *control*, which effects change in cognitive processing (Beran, Brandl, Perner, & Proust, 2012; Metcalfe, 2000; Nelson, 1996; Nelson & Narens, 1990). For example, our student might monitor what she knows and does not know, and use the result of this evaluation to control subsequent study time (e.g., Crystal & Foote, 2011; Flavell, 1979; Kornell, 2009; Metcalfe, 2000). Metacognition has been linked to con-

sciousness (Koriat, 2007), executive control (Fernandez-Duque, Baird, & Posner, 2000; Shimamura, 2000), theory of mind (Proust, 2007), and other challenging phenomena, raising interesting questions for comparative psychologists regarding the extent to which it is a capacity unique to humans or shared with nonhuman animals.

It is comparatively challenging to assess metacognition in nonverbal animals because they are unable to provide verbal commentary on their private cognitive states the way humans do. Nonetheless, experimental designs have been developed to investigate metacognitive abilities in various species (Basile & Hampton, 2014; Beran et al., 2012; Inman & Shettleworth, 1999; Smith, Couchman, & Beran, 2014; Smith et al., 1995; Terrace & Son, 2009). The majority of experiments have emphasized information monitoring and minimized opportunities for cognitive control by providing an option to decline difficult trials without any opportunity to take corrective action to change cognitive state (e.g., Foote & Crystal, 2007; Hampton, 2001; Suda-King, 2008; Templer & Hampton, 2012; Washburn, Smith, & Shields, 2006). Some studies have provided minimal opportunity to exercise control and change cognitive state through information seeking, for example, by choosing to see the location of hidden food before acting (e.g., Basile, Hampton, Suomi, & Murray, 2009; Call & Carpenter, 2001; Hampton, Zivin, & Murray, 2004; Kornell, Son, & Terrace, 2007; Marsh & MacDonald, 2012a). The results of these studies are consistent with metacognitive monitoring and at least minimal cognitive control. However, they provide only “one-shot” opportunities for collecting information and do not examine gradual development of decisions through ongoing feedback between

This article was published Online First August 17, 2015.

Hsiao-Wei Tu, Alex A. Pani, and Robert R. Hampton, Department of Psychology and Yerkes National Primate Research Center, Emory University.

Hsiao-Wei Tu is now at the Neuroscience Division of the Academy of Integrated Science, Virginia Polytechnic Institute and State University.

Supported by National Institute of Mental Health (Grant No. R01MH082819), the National Science Foundation (Grant No. 0745573), the National Center for Research Resources (Grant P51RR165), and the Office of Research Infrastructure Programs (Grant OD P51OD11132). We thank Emily K. Brown, Jessica A. Joiner, and Tara A. Dove-VanWormer for assistance with testing animals and animal care.

Correspondence concerning this article should be addressed to Hsiao-Wei Tu, North End Center, Suite 4300, 300 Turner Street NW, Blacksburg, VA 24061. E-mail: hsiaoweit@vt.edu

monitoring and control. It may be informative to extend the time over which information is collected and begin to examine more gradual and dynamic metacognitive decision making. It is likely that much of the adaptive function of metacognitive monitoring manifests in the role monitoring plays in providing feedback that guides cognitive control.

Recently, the focus of metacognitive research in nonhumans has advanced to the interplay between internal monitoring processes and control of cognitive states. A new paradigm was developed in an attempt to integrate these two aspects of metacognition in a computerized delayed matching-to-sample task, in which the sample and the test were both occluded at the beginning of each trial (Beran & Smith, 2011; Roberts et al., 2009). Monkeys and pigeons were trained to use one cue to reveal the sample and the other to reveal the comparison stimuli for a matching test. If subjects metacognitively evaluated their own knowledge and responded accordingly, they should always uncover the sample before proceeding to the test. To evaluate the alternative explanation that the observed behaviors were due to a series of chained responses learned during training instead of self-monitoring, multiple manipulations were implemented (Beran & Smith, 2011). For example, on some trials, the sample was already uncovered at the beginning, and therefore the subjects were expected to respond immediately to the icon that revealed the matching test. On other trials, both the sample and the test stimuli were presented simultaneously on the screen, in which case subjects should skip the cues and respond to the task directly. Monkeys, but not pigeons, flexibly and adaptively changed their use of the “reveal” option in the different probe conditions (Beran & Smith, 2011; Roberts et al., 2009).

However, unlike natural circumstances in which information may be acquired gradually and the amount of information needed to behave adaptively may vary from case to case, metacognition experiments have defined “information” in an all-or-none fashion. The location of the hidden food was either seen or not seen (Call & Carpenter, 2001; Hampton et al., 2004); the next correct choice was either provided or not provided (Kornell et al., 2007); the sample was either presented or not presented (Beran & Smith, 2011). This dichotomous approach limits the investigation of dynamic interactions between monitoring of information accumulation and seeking of additional information in the development of behavioral decisions. To better understand the extent to which online monitoring of gradually changing cognitive states controls information seeking, we developed a paradigm that begins to allow us to manipulate the amount of information available in a classification task and to examine information seeking and accuracy of classification decisions. Specifically, images that could be classified as birds, fish, flowers, or people were partially or completely occluded at the beginning of each trial. Monkeys were trained to gradually uncover the image by repeatedly touching a “reveal” button. A classification choice could be made at any time, whether or not sufficient information was available to make a correct response. We emphasize that we did not generate this paradigm entirely de novo. Instead, it is a logical extension of the work by Roberts et al. (2009) and Beran and Smith (2011). We conducted three experiments to test the hypothesis that monkeys engage in ongoing monitoring of accumulating information and seek the information necessary for a correct response. Monkeys that monitor and respond adaptively to accumulating information should make many “revelation” responses when information is poor and

few such responses when information is rich. Alternatively, monkeys lacking monitoring or control should show no systematic relation between use of the reveal button and available information.

Experiment 1

Monkeys learned to accurately classify intact images of birds, fish, flowers, and people. After reaching criterion in this training, to-be-classified images were initially concealed behind gray rectangles and monkeys had to touch a button to view the image. Each touch of the reveal button removed a single occluding rectangle. If monkeys monitor the information obtained and use that knowledge to determine whether sufficient information is available for a correct response, we would expect the number of touches of the reveal button before a correct classification decision to be greater when the amount of the image revealed by each touch of the button is lower.

Method

Subjects and apparatus. Six adult male rhesus monkeys (*Macaca mulatta*) were used, each of which was tested with an individual testing rig attached to the front of their home cage. Each testing rig included a personal computer controlling the experiment by custom programs written in Presentation (Neurobehavioral Systems, Albany, CA), a 15-in. color LCD touch-sensitive screen (Elo TouchSystems, Menlo Park, CA) running at a resolution of 1024×768 pixels, and two automatic food dispensers (Med Associates, St. Albans, VT) that delivered nutritionally balanced primate pellets (Bio-Serv, Frenchtown, NJ) into food cups below the screen.

Monkeys were pair-housed, received a full ration of food each day, and had ad libitum access to water. During the 6–7 hr of testing in a day, pairs were separated by dividers that prevented access to the rig in the adjacent cage but allowed limited visual and physical contact. All monkeys had prior experience with automated cognitive tests using touch-screen computers and were previously trained on clip-art visual matching-to-sample.

Stimuli. A total of 600 color photographs, 150 from each of four categories (i.e., fish, flowers, birds, and people), were used in the present study. All images were collected from the online photo repository Flickr (Yahoo!, Sunnyvale, CA) and visually examined to ensure that each photograph depicted one or more exemplars from a single category. Duplicates were eliminated by DupDetector (Prismatic Software, Anaheim, CA) and visual inspection. Images were resized to 300×400 pixels (Height \times Width) by Adobe Photoshop (Adobe, San Jose, CA).

Training procedure. Each subject received one session per day, 6 days a week. Monkeys were trained to classify color photographs into four categories: birds, fish, people, or flowers. A total of 100 photographs per category were randomly chosen from the pool for initial training. Each trial started with a green square (100×100 pixels) at the bottom center of the screen. After touching it twice (FR2), a single photograph was presented at the center of the screen. This to-be-classified image was selected from one of the categories with the restriction that images from the same category did not appear more than four times consecutively. Monkeys touched the image (FR2) and four category icons appeared in

the four corners of the screen (Figure 1a). Icons always appeared in the same corner, such that monkeys could map the categories to either response location or to the identity of the icon. Touching (FR2) the correct icon was rewarded with one pellet together with a positive sound. Incorrect responses resulted in a negative sound, a 5-s timeout, during which the screen stayed black, and a correction trial. The correction trial proceeded as the primary trial with the same to-be-classified image presented again. Responses were reinforced in the same way, except that incorrect responses led to a third trial, during which only the correct classification icon was present at test. Monkeys had to touch it (FR2) to receive positive auditory feedback and a food reward. Only performance on the first iteration of each trial was used in analyses. Consecutive trials were separated by a 3-s intertrial interval. Each to-be-classified image appeared once in a session of 400 trials, and the images were repeatedly used across sessions. When the monkeys performed above 70% correct in a session, an additional 50 new color photographs of each category were intermixed with the original 400 photographs to encourage and evaluate generalization. When the overall accuracy of this set of 600 mixed images was above 70% in a single session, monkeys proceeded to the next stage of training.

Monkeys were trained to use the “reveal” button to uncover to-be-classified images. Trials proceeded as before, except that, after monkeys touched the green start square, a gray rectangle appeared with the image “concealed” underneath, and a purple reveal button appeared below the covered image (see Figure 1b). Monkeys could touch one of the four category symbols at any time, including while the image was completely concealed. One touch (FR1) to the reveal button uncovered the entire to-be-classified image. There was no restriction on how many times the reveal button could be touched, but extra touches after the image had been uncovered had no effect. Monkeys advanced to the test session when classification accuracy was greater than 80% correct for each of three consecutive 600-trial sessions.

Test procedure. There were four testing stages that varied in the amount of the gray rectangle that was removed by each touch of the reveal button. Specifically, the to-be-classified image was segmented into grids of 150×200 , 100×100 , 50×50 , or 30×40 pixels in Stage 1, 2, 3, and 4, respectively. Each segment of the image was covered by a gray rectangular blocker. Each touch (FR1) to the reveal button removed one randomly selected blocker, revealing part of the image. With smaller blockers, more touches

were required to reveal the same amount of the image. There were 160 trials in each session, including 40 images from each category randomly chosen from the training session without replacement. Monkeys had to perform better than 80% correct on each of three consecutive sessions before moving to the next stage.

Data analyses. The primary dependent variable in this experiment was the average number of reveal button touches before classification in the last three sessions with each of the four different blocker sizes. Data from all trials were compared across the four stages of different blocker size by repeated-measures analysis of variance. Post hoc analyses were conducted by paired *t* tests. We report standardized Cohen’s d_z for within-subjects designs (Lakens, 2013) and the 95% confidence interval (CI) for this measure of effect size for each *t* value larger than 1.

Results and Discussion

On average, 7.17 sessions were needed for classification performance to reach 70%, and all monkeys transferred immediately to the 50 new images from each category without decrease in accuracy ($M = 77\%$ with only the 200 new images and $M = 81\%$ across all images vs. the 70% criterion for moving to this generalization test). This indicated that the monkeys had learned the classification task. Monkeys were trained for an additional 5.17 sessions on average to use the reveal button. At the end, all monkeys revealed at least part of the image before making a classification decision at least 97% of the time ($M = .99$). Monkeys stayed in Stage 1, 2, 3, and 4 for 4.33, 6.83, 13.67, and 9.17 sessions, respectively, before meeting criterion.

Monkeys made more touches to the reveal button when more were needed to reveal sufficient information for classification. The average number of reveal button touches in the last three sessions differed significantly across different blocker sizes, $F(3, 15) = 35.66$, mean squared error (MSE) = 61.43, $p < .001$, $\eta_p^2 = 0.88$ (see Figure 2). Post hoc analyses showed that monkeys made significantly more touches to the reveal button with each decrease in the size of the blocks removed by the reveal button, Stage 1 versus Stage 2: $t(5) = 6.98$, $p = .001$, $d_z = 2.85$, 95% CI [0.93, 4.66]; Stage 2 versus Stage 3: $t(5) = 5.69$, $p = .002$, $d_z = 2.32$, 95% CI [0.70, 3.91]; Stage 3 versus Stage 4: $t(5) = 2.79$, $p = .04$, $d_z = 1.14$, 95% CI [0.06, 2.16]. These results suggest that monkeys monitored their evolving decision process and collected more information as necessary to reach a reliable decision.




Figure 1. An example of a regular classification trial (a). There was no reveal button, and the monkeys had to touch the image (FR2) for the category symbols to appear for a decision. Each category symbol was always presented at the same corner. An example of a trial in Experiment 1 (b). The image was completely occluded by gray blockers at the beginning of each trial, but the size of these blockers and the “information” given per reveal button touch decreased across sessions. An example of a probe trial in Experiment 2 (c). The image was shrunk to the size of a 60×80 blocker, and the timing for it to be revealed was manipulated. See the online article for the color version of this figure.




Figure 2. The average number of reveal button touches before correct and incorrect classification in the last three sessions of each testing stage in Experiment 1. As the blockers became smaller, the monkeys revealed more before making a classification decision. Error bars indicate standard errors.

In the last three sessions of Stage 1 where the blockers had the largest size, the average number of reveal button touches on incorrect trials ($M = 2.38$) was significantly lower than that on correct trials ($M = 2.71$), $t(5) = 7.46$, $p = .001$, $d_z = 3.05$, 95% CI [1.04, 5.03]. This finding is consistent with more information promoting higher accuracy, but also raises the question why monkeys did not make more reveal responses on incorrect trials so that they could respond correctly. However, the difference in reveal responses between correct and incorrect trials was not significant in the following three stages, where the area uncovered by each touch became smaller (see Figure 2), Stage 2: $t(5) = 2.25$, $p = .074$, $d_z = 0.92$, 95% CI [-0.09, 1.86]; Stage 3: $t(5) = 0.70$, $p = .52$; Stage 4: $t(5) = 0.78$, $p = .47$. The lack of a difference in the number of reveal responses on correct and incorrect trials in the last three stages of testing raises the possibility that monkeys used the reveal response without sensitivity to the amount of information they had collected. One might expect errors to be associated with lower numbers of reveal responses. However, a metacognitive account of use of the reveal button is premised on monkeys making a subjective assessment of the adequacy of the information available, and this assessment will not always be accurate. Monkeys may make classification responses before sufficient information is available, either out of misjudgment of the information available or as a result of poor impulse control. With the current design, the objective state of the image and the subjective state of the decision process could not be rigidly linked, although they must be related. Because we do not know precisely when sufficient information is available for monkeys to respond correctly, we designed Experiment 2 to clearly link use of the reveal response to the availability of a sufficient, and in fact maximal, amount of information.

Experiment 2

In Experiment 1, monkeys were familiarized with the use of the reveal button and used it in an apparently metacognitive manner, matching effort spent seeking information to information required. Specifically, the number of touches to the reveal button increased as the area uncovered by each touch decreased. However, because

we introduced the different blocker sizes sequentially, it is possible that monkeys learned at each stage that a specific average number of touches was required to secure a reward, rather than learning to monitor their developing decision process and the information available. Such learning is consistent with the lack of significant differences in the number of reveal responses on incorrect and correct trials in most phases of Experiment 1. Experiment 2 was designed to test whether reveal button touches were controlled by an association with a given blocker size or stage of training, or instead was controlled by the status of the classification decision process or the available information.

With the method used in Experiment 1, it was not possible to know how much of a particular image monkeys had to reveal to successfully classify it. For this reason, it was not possible to determine whether they only revealed enough of the image to support classification. To provide a specific test of whether monkeys regulated use of the reveal button in response to information, we shrank the to-be-classified image and placed it under a single blocker among others on probe trials. Only revealing this critical blocker would offer the information necessary for correct classification. If monkeys learned in Experiment 1 to emit a certain number of touches to the reveal button before classifying, rather than monitoring the amount of information accumulated, they should emit that number of responses whether or not the critical blocker had been removed.

Method

The same 600 color photographs used in training were used here. Blockers were 60×80 pixels (Height \times Width), resulting in a gray rectangle consisting of 25 blockers. Each touch to the reveal button (FR1) removed one blocker. Monkeys were allowed to touch the reveal button as many times as they wished and to classify at any time. Due to a programming error, we did not record extra touches after all 25 blockers were removed. This bounded the number of recorded touches at 25.

In each session, 15 trials per category were randomly selected as probe trials, and the remaining 540 trials (135 trials per category) were presented as regular trials. On regular trials, the to-be-classified image was 300×400 pixels, covered by 25 blockers. On probe trials, a rescaled to-be-classified image (60×80 pixels) was positioned behind a randomly selected "critical blocker." Removing this blocker revealed the shrunken image, whereas removing the other blockers revealed only an uninformative light yellow background (see Figure 1c). The critical blocker was programmed to be removed on the 1st, 7th, 13th, 19th, or 25th touch to the reveal button. Each revealing schedule was applied to three probe trials in each of the four categories. Probe trials were distributed pseudorandomly in a session so that no two probe trials occurred consecutively. Contingencies of reinforcement stayed the same as in Experiment 1 for both regular trials and probe trials, but there was no correction trial for incorrect responses. Each monkey was tested for five sessions. Trials from each of the revelation schedules were pooled across image categories for analysis, yielding a total of 60 probe trials per revealing schedule. Accuracy in each schedule and average number of blockers removed before classifying were analyzed in the same way as in Experiment 1.

Results and Discussion

Classification accuracy did not differ significantly across changes in the number of reveal responses required to remove the critical blocker, $F(4, 20) = 2.18$, $MSE = 0.01$, $p = .11$, $\eta_p^2 = 0.30$, showing that monkeys used the reveal button at least enough to uncover the critical blocker. In contrast, the average number of reveal button touches prior to classification did vary significantly as a function of the number required to remove the critical blocker, $F(4, 20) = 61.82$, $MSE = 2.64$, $p < .001$, $\eta_p^2 = 0.93$ (see Figure 3). The average number of touches increased significantly with each increase in the number of responses required, 1 touch versus 7 touches: $t(5) = 3.26$, $p = .02$, $d_z = 1.33$, 95% CI [0.17, 4.46]; 7 touches versus 13 touches: $t(5) = 9.20$, $p < .001$, $d_z = 3.76$, 95% CI [1.37, 6.14]; 13 touches versus 19 touches: $t(5) = 6.51$, $p = .001$, $d_z = 2.66$, 95% CI [0.86, 4.43]; 19 touches versus 25 touches: $t(5) = 3.72$, $p = .01$, $d_z = 1.52$, 95% CI [0.28, 2.70]. Thus, the monkeys adjusted use of the reveal button adaptively to ensure they acquired the information necessary for accurate classification, and did not touch the reveal button according to any fixed schedule.

However, monkeys did not immediately stop touching the reveal button when the critical blocker was removed, especially when the removal happened early in a trial. There were significantly more touches to the reveal button than were needed on probe trials when the shrunken image was revealed after 1, 7, and 13 reveal button touches, 1 touch: $M = 9.19$, $t(5) = 7.69$, $p = .001$, $d_z = 3.14$, 95% CI [1.09, 5.17]; 7 touches: $M = 11.87$, $t(5) = 8.19$, $p < .001$, $d_z = 3.34$, 95% CI [1.18, 5.49]; 13 touches: $M = 16.11$, $t(5) = 4.09$, $p = .01$, $d_z = 1.67$, 95% CI [0.36, 2.93], but not when 19 or 25 touches were required, 19 touches: $M = 19.39$, $t(5) = 0.37$, $p = .73$; 25 touches: $M = 21.87$, $t(5) = 2.02$, $p = .10$, $d_z = 0.82$, 95% CI [-0.15, 1.74]. In addition, two of the six monkeys may have shown a bimodal distribution, such that they sometimes touched 25 times when not necessary. This result indicates that monkeys tend to “overshoot” use of the reveal button, uncovering more blockers even after the to-be-classified image had already appeared. Be-

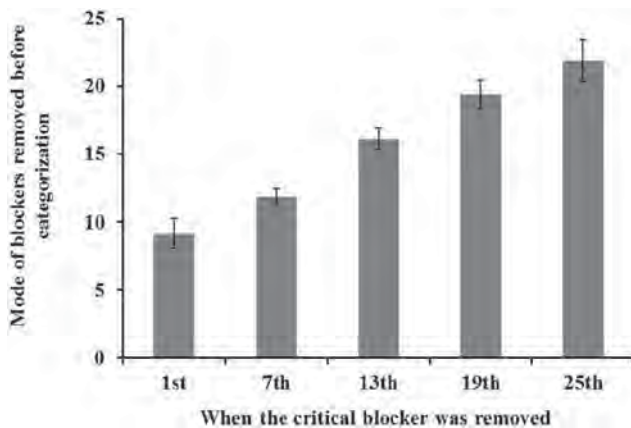


Figure 3. In Experiment 2, monkeys made significantly more touches to the reveal button as the number of blockers needed to be removed before the to-be-classified image appeared increased. The accuracy did not change significantly among different revelation schedules. Error bars indicate standard errors.

cause the reveal button was presented at approximately the same location as the green start square, which the monkeys had to touch twice to start the trial, they may have often continued touching that location before examining the to-be-classified image and reacting metacognitively. These touches were very easy to make because their hand was already in position and tapping. Also, because nearly all trials required some touches before sufficient information was available, monkeys may have emitted several touches before checking whether sufficient information was available. By this account, we would expect to observe the most overshoot when only one touch was required to remove the critical blocker and progressively less as the number of responses required was higher, and that is what we observed.

We also found that there were significantly fewer touches to the reveal button in the last session than in the first session when the critical blocker was removed after 1, 7, and 13 reveal button touches, 1 touch: $t(5) = 3.87$, $p = .01$, $d_z = 1.58$, 95% CI [0.31, 2.79]; 7 touches: $t(5) = 2.88$, $p = .03$, $d_z = 1.18$, 95% CI [0.08, 2.21]; 13 touches: $t(5) = 3.31$, $p = .02$, $d_z = 1.35$, 95% CI [0.18, 2.46], but, in the last session, monkeys still revealed significantly more than needed when the to-be-classified image appeared after 1 and 7 reveal button touches, 1 touch: $M = 8.22$, $t(5) = 5.98$, $p = .002$, $d_z = 2.44$, 95% CI [0.76, 4.16]; 7 touches: $M = 11.39$, $t(5) = 3.82$, $p = .012$, $d_z = 1.56$, 95% CI [0.30, 2.76]. Although this indicates that monkeys might have learned to better control their button touch behavior with practice, it is difficult to distinguish whether the monkeys had learned to behave metacognitively or had learned to “keep revealing until a nonblank blocker is revealed,” which may not involve metacognitive processes.

Experiment 3

In Experiment 2, we found that monkeys made fewer revelation responses when fewer responses were required to remove the critical blocker, suggesting that they monitored whether the information needed for accurate classification was available and flexibly controlled use of the reveal button. However, a viable nonmetacognitive explanation is that monkeys had learned to “keep revealing until a nonblank blocker is revealed.” To further evaluate the extent to which monkeys match use of the reveal button to the amount of information accumulated, three types of manipulations were used in Experiment 3. Using only correct regular trials from Experiment 2, we reconstructed the images monkeys saw at the point they chose to make the classification response. Given that the monkeys responded correctly, we inferred that most of the time these displays contained sufficient information for accurate classification. These partially revealed images were presented in probe trials with the exact number of blockers in exactly the same positions as when each monkey had chosen to classify in Experiment 2. We hypothesized that the monkeys would reveal no more and respond immediately because there should be enough information for them to classify correctly. In a second condition, we put some blockers back on these partially revealed images before they were presented to the monkeys, so that the images were more occluded than they had been when the monkey chose to classify in Experiment 2. If the behavior of monkeys is controlled metacognitively, they should touch the

reveal button more times on these trials in order to receive more information before classifying. Finally, we also removed additional blockers from these partially revealed images so that more information than required was presented for making a decision. If use of the reveal button is controlled entirely by the information available, monkeys should classify immediately without touching the reveal button.

Method

The same 600 color photographs used in Experiment 2 were used here. Blockers of 60×80 pixels (Height \times Width) were used to cover the entire image on regular trials and partial images on probe trials. Each touch to the reveal button randomly removed one blocker. Monkeys were allowed to touch one of the category icons at any time.

A total of 100 images (25 of each category) that had seven blockers revealed (i.e., 18 blockers remained on the image) at the time of correct classification were randomly selected among all correct regular trials in Experiment 2. Each of these 100 images was manipulated in three ways, resulting in 300 probe trials. First, in the “unchanged” condition, exactly the same seven blockers were removed, leaving 18 blockers cover exactly the same areas on the images as when they were classified in Experiment 2. Second, in the “more revealed” condition, five additional blockers were randomly removed from each image so that only 13 blockers remained. Third, in the “less revealed” condition, five extra blockers were randomly added to each image so that 23 blockers obscured the image. Twenty probe trials from each condition were selected to include in a session with the constraint that the same image only appeared once in each session.

There were 540 regular trials and 60 probe trials in a session. On regular trials, the entire to-be-classified image was covered by 25 blockers at the beginning, whereas, on probe trials, the manipulated images were presented immediately after the monkeys touched the green start square at the beginning of each trial. Each monkey was tested for five sessions yielding 100 of each type of probe trial. Test procedures and data analyses remained the same as in Experiment 2.

Results and Discussion

Classification accuracy remained similar across different conditions ($M = .99$) regardless of how much of the to-be-classified image was revealed at the beginning of a trial, $F(2, 10) = 0.24$, $MSE = 0.004$, $p = .79$, $\eta_p^2 = 0.05$, indicating that monkeys normally revealed at least enough information for accurate classification. More importantly, monkeys adaptively adjusted their use of the reveal button in response to different initial states of the to-be-classified image. The number of additional blockers revealed at test differed significantly across the three conditions (see Figure 4), $F(2, 10) = 63.94$, $MSE = 0.17$, $p < .001$, $\eta_p^2 = 0.93$. Specifically, significantly more blockers were removed when less of the image was revealed initially, less revealed condition versus unchanged condition: $t(5) = 16.82$, $p < .001$, $d_z = 6.87$, 95% CI [2.71, 11.07]; unchanged condition versus more revealed condition: $t(5) = 2.57$, $p = .05$, $d_z = 1.05$, 95% CI [0.00, 2.04]. Monkeys thus acquired more information when less was revealed than had been when they classified images in Experiment 2, and

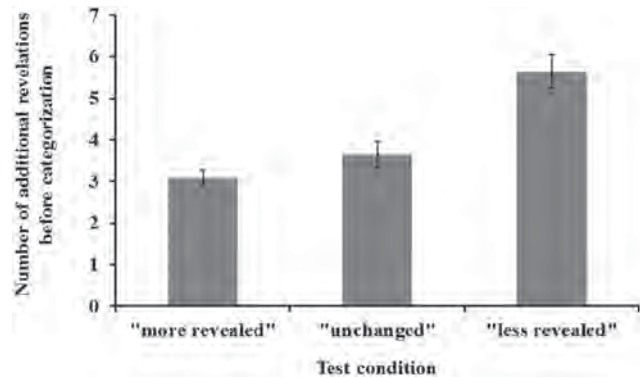


Figure 4. In Experiment 3, the number of additional blockers revealed at test differed significantly across the three conditions. Specifically, more blockers were removed before classification when the image was largely occluded, whereas fewer blockers were removed before classification when the image was almost completely shown. Error bars indicate standard errors.

tended to acquire less information if more was available than had been at the time they chose to classify in Experiment 2.

In the less revealed condition, five additional blockers were added to conceal more of the occluded image from Experiment 2. Therefore, on those probe trials, one might predict that the monkeys would, on average, make about five touches to the reveal button before classification. This is close to what we found. The number of revelations did not differ significantly from five in the “less revealed” condition, $t(5) = 1.61$, $p = .17$, $d_z = 0.66$, 95% CI [−0.26, 1.53], consistent with our hypothesis that monkeys should stop touching the revelation cue when enough information was obtained for a correct classification decision. However, on trials in which the image was revealed sufficiently at the beginning, monkeys still removed some blockers instead of choosing one of the category symbols immediately, unchanged condition $M = 3.65$, $t(5) = 11.40$, $p < .001$, $d_z = 4.65$, 95% CI [1.76, 7.55]; more revealed condition: $M = 3.08$, $t(5) = 16.36$, $p < .001$, $d_z = 6.68$, 95% CI [2.63, 10.77].

We found that the number of these unnecessary button touches was significantly lower in the last session compared with the first session, unchanged condition: $M_{\text{first}} = 4.23$, $M_{\text{last}} = 3.12$, $t(5) = 4.11$, $p = .01$, $d_z = 1.68$, 95% CI [0.36, 2.94]; more revealed condition: $M_{\text{first}} = 3.63$, $M_{\text{last}} = 2.85$, $t(5) = 2.81$, $p = .04$, $d_z = 1.15$, 95% CI [0.06, 2.18], indicating that monkeys did learn to better control excessive responding with experience. But again, we cannot determine whether this reflects improved cognitive control, better attention to the to-be-classified image early in trials, or trial-and-error learning that fewer responses were required.

General Discussion

We allowed monkeys to incrementally reveal occluded images and freely determine when to classify these images in Experiments 1 and 3. Monkeys rarely classified images immediately, but instead revealed information that increased the probability of correct classification. This suggests that monkeys privately evaluated an evolving decision process and collected more information until they could confidently classify the images. Thus, monkeys appear

to dynamically monitor their knowledge states, and this monitoring exerts feedback to regulate information seeking, constituting a type of cognitive control. The “revelation” paradigm used here represents an advance in nonhuman metacognition research because it extends previous work to provide an opportunity to study sustained interactions between cognitive monitoring and cognitive control. Previous research on animal metacognition has focused mostly on monitoring, using decline test paradigms in which subjects can avoid difficult tests but have no opportunity to alter the course of cognitive processing by collecting more information (e.g., Smith, Redford, Beran, & Washburn, 2010; Smith & Washburn, 2005; Suda-King, 2008; Templer & Hampton, 2012). Other studies have addressed cognitive control in a minimal way by including an active “one-shot” information-seeking component (e.g., Basile et al., 2009; Beran & Smith, 2011; Call & Carpenter, 2001; Castro & Wasserman, 2013; Hampton et al., 2004; Marsh & MacDonald, 2012b; Roberts, McMillan, Musolino, & Cole, 2012). Although cognitive monitoring alone can be useful in avoiding situations where needed information is not available, it is likely that the critical role cognitive monitoring plays in adaptively regulating cognitive control is a major reason for the evolution of cognitive monitoring. Further studies of the dynamic interplay between cognitive monitoring and cognitive control will allow us to better understand how the feedback from monitoring modulates cognitive control.

Monkeys adaptively adjusted information seeking in response to variations in information available, but their use of the reveal button was not optimal. Monkeys tended to make more revelation responses than necessary when information was easy to get (e.g., when the critical blocker was the first one to be revealed in Experiment 2) or already given (e.g., when the to-be-classified image was already more revealed than when it was correctly classified earlier in Experiment 3). One nonmetacognitive interpretation is that the monkeys learned to touch the reveal button numerous times before emitting a classification response, rather than monitoring the accumulations of information necessary for a decision. Based on this account, the number of button touches before classification should not differ as a function of the size of the blockers or the amount of information available. In contrast, we observed that monkeys used the reveal button more when each response revealed smaller parts of the to-be-classified image and that use of the reveal button decreased immediately when the blocker size changed from 30×40 pixels in the last manipulation of Experiment 1 ($M = 9.77$) to 60×80 pixels in the first session of Experiment 2 ($M = 7.94$). These findings suggest that the monkeys dynamically monitored the information available to them and changed their information-seeking behavior accordingly. We consider this monitoring dynamic because different numbers of reveal responses were adaptive depending on how rapidly information was revealed. It would be of interest to extend these tests to a situation in which the amount of information revealed by each reveal response is unpredictable, putting more demand on ongoing monitoring of the information accumulated.

While modulation of use of the reveal button suggests information monitoring, excessive use of the reveal button still indicates that this behavior is not always tightly controlled by metacognition. There are likely other sources of stimulus control for use of the reveal button, and they may operate in parallel

with metacognitive monitoring. On the nonprobe trials that made up most trials in each session, monkeys always had to remove multiple blockers before images could be reliably classified correctly. Thus, they may have learned to always make at least several reveal responses before checking whether enough information was available to classify. The effect of such an expectation, or habit, would be exacerbated by the fact that the reveal button appeared near the green start square that initiated trials (see Figures 1b and 1c), making it very easy to switch from touching the start square to touching the reveal button at the beginning of a trial. Future studies might reduce this by delaying the presentation of the reveal button to ensure that monkeys inspect the to-be-classified image before contacting the reveal button, and by moving the reveal button away from the green start square. In contrast to this rapid and simple exercise of habitual touching of the reveal button, metacognition may require hierarchical cognitive processing (Metcalfe, 2000; Nelson, 1996; Nelson & Narens, 1990) that develops over a period of time longer than that needed to initiate a simple motor response (McLaughlin, Simon, & Gillan, 2010). Therefore, although on some probe trials an immediate classification response was optimal, unnecessary responses to the reveal button may have occurred before metacognitive control of responding became active. This new “revelation” paradigm may offer an opportunity for future study of the interactions between stereotypical, habitual motor movements and higher level cognitive control that can modulate such behaviors.

A third possible explanation for the excessive button touches is that our monkeys “double checked” before making classification responses, even though they already had enough information to make a decision. Similar behaviors may be common in humans (Call, 2010; Call & Carpenter, 2001). For example, we may repeatedly confirm that we have our passport in possession when traveling, even when we are quite certain it is in our pocket, because the cost of forgetting to bring the passport is extremely high. Our monkeys appear to have revealed more blockers than necessary only when information was easy to obtain (e.g., when the critical blocker was the first one to be revealed in Experiment 2) and rarely spent the time necessary to reveal the entire image, suggesting that they were able to evaluate the cost and benefit of seeking more information and metacognitively to determine whether more information was needed.

We have argued that use of the reveal response in these experiments is consistent with metacognitive monitoring of a mental state reflecting the status of a decision process. However, we recognize that use of the reveal response is subject to other explanations, some of which were discussed above. One major interpretive issue that is very difficult to address directly with this paradigm, or any information-seeking paradigm, is that the metacognitive response brings about a change in the physical stimulus controlling the primary cognitive response. In the present case, the putatively metacognitive response of pressing the reveal button changed the visibility of the image and therefore the ability of the image to control a correct classification response. So we are left without a certain answer whether cessation of reveal responses was controlled “directly” by the change in the image, or was mediated by a change in a mental state reflecting the status of the classification decision process. Because the status of the decision process is surely determined by that status of the to-be-classified image, it is difficult or impossible to entirely discriminate between the

alternatives of control of behavior by the physical image or by a cognitive state. At some level, attempting to make this distinction is incoherent given that cognitive states arise from physical stimuli. The two cannot be entirely decoupled.

Conclusion

We have presented converging data indicating that rhesus monkeys flexibly adjust information seeking based on monitoring of the amount of knowledge accumulated over time. By allowing monkeys to gradually uncover to-be-classified images in a visual classification test, we have begun to study the dynamic interactions of cognitive monitoring and cognitive control in a simple decision-making process in monkeys. It is likely that cognitive monitoring plays much broader roles in cognitive function than simply triggering “one-shot” avoidance of difficult situations. Studies of how cognitive monitoring provides feedback for ongoing cognitive control of decision making promise to better document the adaptive value of cognitive monitoring.

References

- Basile, B. M., & Hampton, R. R. (2014). Metacognition as discrimination: Commentary on Smith et al. *Journal of Comparative Psychology*, *128*, 135–137. <http://dx.doi.org/10.1037/a0034412>
- Basile, B. M., Hampton, R. R., Suomi, S. J., & Murray, E. A. (2009). An assessment of memory awareness in tufted capuchin monkeys (*Cebus apella*). *Animal Cognition*, *12*, 169–180. <http://dx.doi.org/10.1007/s10071-008-0180-1>
- Beran, M. J., Brandl, J., Perner, J., & Proust, J. (Eds.). (2012). *Foundations of metacognition*. New York, NY: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199646739.001.0001>
- Beran, M. J., & Smith, J. D. (2011). Information seeking by rhesus monkeys (*Macaca mulatta*) and capuchin monkeys (*Cebus apella*). *Cognition*, *120*, 90–105. <http://dx.doi.org/10.1016/j.cognition.2011.02.016>
- Call, J. (2010). Do apes know that they could be wrong? *Animal Cognition*, *13*, 689–700. <http://dx.doi.org/10.1007/s10071-010-0317-x>
- Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, *3*, 207–220. <http://dx.doi.org/10.1007/s100710100078>
- Castro, L., & Wasserman, E. A. (2013). Information-seeking behavior: Exploring metacognitive control in pigeons. *Animal Cognition*, *16*, 241–254. <http://dx.doi.org/10.1007/s10071-012-0569-8>
- Crystal, J. D., & Foote, A. L. (2011). Evaluating information-seeking approaches to metacognition. *Current Zoology*, *57*, 531–542. Retrieved from <http://www.currentzoology.org/temp/%7B465ACF1E-5363-460E-B074-6EAC743C38E8%7D.pdf>
- Fernandez-Duque, D., Baird, J. A., & Posner, M. I. (2000). Awareness and metacognition. *Consciousness and Cognition*, *9*, 324–326. <http://dx.doi.org/10.1006/ccog.2000.0449>
- Flavell, J. H. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, *34*, 906–911. <http://dx.doi.org/10.1037/0003-066X.34.10.906>
- Foote, A. L., & Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, *17*, 551–555. <http://dx.doi.org/10.1016/j.cub.2007.01.061>
- Hampton, R. R. (2001). Rhesus monkeys know when they remember. *PNAS: Proceedings of the National Academy of Sciences of the United States of America*, *98*, 5359–5362. <http://dx.doi.org/10.1073/pnas.071600998>
- 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*, 239–246. <http://dx.doi.org/10.1007/s10071-004-0215-1>
- Inman, A., & Shettleworth, S. J. (1999). Detecting metamemory in non-verbal subjects: A test with pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, *25*, 389–395. <http://dx.doi.org/10.1037/0097-7403.25.3.389>
- Koriat, A. (2007). Metacognition and consciousness. In P. D. Zelazo, M. Moscovitch, & E. Thompson (Eds.), *Cambridge handbook of consciousness* (pp. 289–326). New York, NY: Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511816789.6012>
- Kornell, N. (2009). Metacognition in humans and animals. *Current Directions in Psychological Science*, *18*, 11–15. <http://dx.doi.org/10.1111/j.1467-8721.2009.01597.x>
- Kornell, N., Son, L. K., & Terrace, H. S. (2007). Transfer of metacognitive skills and hint seeking in monkeys. *Psychological Science*, *18*, 64–71. <http://dx.doi.org/10.1111/j.1467-9280.2007.01850.x>
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for *t*-tests and ANOVAs. *Frontiers in Psychology*, *4*, 863. <http://dx.doi.org/10.3389/fpsyg.2013.00863>
- Marsh, H. L., & MacDonald, S. E. (2012a). Information seeking by orangutans: A generalized search strategy? *Animal Cognition*, *15*, 293–304. <http://dx.doi.org/10.1007/s10071-011-0453-y>
- Marsh, H. L., & MacDonald, S. E. (2012b). Orangutans (*Pongo abelii*) “play the odds”: Information-seeking strategies in relation to cost, risk, and benefit. *Journal of Comparative Psychology*, *126*, 263–278. <http://dx.doi.org/10.1037/a0025906>
- McLaughlin, A. C., Simon, D. A., & Gillan, D. J. (2010). From intention to input: Motor cognition, motor performance, and the control of technology. *Reviews of Human Factors and Ergonomics*, *6*, 123–171. <http://dx.doi.org/10.1518/155723410X12849346788741>
- Metcalf, J. (2000). Metamemory: Theory and data. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 197–211). New York, NY: Oxford University Press.
- Nelson, T. O. (1996). Consciousness and metacognition. *American Psychologist*, *51*, 102–116. <http://dx.doi.org/10.1037/0003-066X.51.2.102>
- Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. *Psychology of Learning and Motivation*, *26*, 125–173. [http://dx.doi.org/10.1016/S0079-7421\(08\)60053-5](http://dx.doi.org/10.1016/S0079-7421(08)60053-5)
- Proust, J. (2007). Metacognition and metarepresentation: Is a self-directed theory of mind a precondition for metacognition? *Synthese*, *159*, 271–295. <http://dx.doi.org/10.1007/s11229-007-9208-3>
- Roberts, W. A., Feeney, M. C., McMillan, N., MacPherson, K., Musolino, E., & Petter, M. (2009). Do pigeons (*Columba livia*) study for a test? *Journal of Experimental Psychology: Animal Behavior Processes*, *35*, 129–142. <http://dx.doi.org/10.1037/a0013722>
- Roberts, W. A., McMillan, N., Musolino, E., & Cole, M. (2012). Information seeking in animals: Metacognition? *Comparative Cognition & Behavior Reviews*, *7*, 85–109. <http://dx.doi.org/10.3819/ccbr.2012.70005>
- Shimamura, A. P. (2000). Toward a cognitive neuroscience of metacognition. *Consciousness and Cognition*, *9*, 313–323. <http://dx.doi.org/10.1006/ccog.2000.0450>
- Smith, J. D., Couchman, J. J., & Beran, M. J. (2014). Animal metacognition: A tale of two comparative psychologies. *Journal of Comparative Psychology*, *128*, 115–131. <http://dx.doi.org/10.1037/a0033105>
- Smith, J. D., Redford, J. S., Beran, M. J., & Washburn, D. A. (2010). Rhesus monkeys (*Macaca mulatta*) adaptively monitor uncertainty while multi-tasking. *Animal Cognition*, *13*, 93–101. <http://dx.doi.org/10.1007/s10071-009-0249-5>
- 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. <http://dx.doi.org/10.1037/0096-3445.124.4.391>
- Smith, J. D., & Washburn, D. A. (2005). Uncertainty monitoring and metacognition by animals. *Current Directions in Psychological Science*, *14*, 19–24. <http://dx.doi.org/10.1111/j.0963-7214.2005.00327.x>

