Animal Cognition
https://doi.org/10.1007/s10071-021-01477-5

ORIGINAL PAPER

Rhesus monkeys (*Macaca mulatta*) monitor evolving decisions to control adaptive information seeking

Ryan J. Brady[1,2] · **Robert R. Hampton**[1,2]

Received: 24 September 2020 / Revised: 6 January 2021 / Accepted: 11 January 2021

Abstract

Adaptive decision making in humans depends on feedback between *monitoring*, which assesses mental states, and *control*, by which cognitive processes are modified. We investigated the extent to which monitoring and control interact iteratively in monkeys. Monkeys classified images as birds, fish, flowers, or people. At the beginning of each trial, to-be-classified images were not visible. Monkeys touched the image area to incrementally brighten the image, referred to as the *brighten response*. The amount by which brightness increased with each brighten response was unpredictable, and the monkeys could choose to classify the images at any time during a trial. We hypothesized that if monkeys monitored the status of their classification decision then they would seek information depending on the amount of information available. In Experiment 1, monkeys rarely used the brighten response when images were bright initially, and they used the brighten response more when earlier uses in a given trial yielded smaller amounts of information. In Experiment 2, monkeys made more brighten responses when the presented image did not belong in any of the trained categories, suggesting monkeys were sensitive to the fact that they could not reach a classification decision despite the image brightening. In Experiment 3, we found that the probability that monkeys used the brighten response correlated with their ability to correctly classify when the brighten response was not available. These findings add to the literature documenting the metacognitive skills of nonhuman primates by demonstrating an iterative feedback loop between cognitive monitoring and cognitive control that allows for adaptive information-seeking behavior.

Keywords Metacognition · Cognition · Decision making · Nonhuman primate

Introduction

Human cognition can be flexible and strategic. One way that human behavior gains flexibility is through metacognition, or thinking about thinking (Benjamin et al. 1998; Dunlosky and Bjork 2008; Flavell 1979; Nelson and Narens 1990; Shimamura and Squire 1986). By assessing what we know, and recognizing when we do not know enough, we can optimize information-seeking before making a decision, balancing the need to gather more information and the cost of acquiring it (Beran and Smith 2011). This dynamic relationship has been described as a feedback loop between *monitoring*, which assesses the mental state, and *control*, which adjusts cognition and behavior to bring about positive change in the monitored mental state (Beran et al. 2012; Metcalfe 2009; Nelson and Narens 1990). Because metacognitive abilities have been linked to aspects of cognition thought to be special in humans, such as consciousness (Koriat 2000), the theory of mind (Proust 2007), and planning, comparative studies of the relationship between metacognitive monitoring and cognitive control may help us better understand the extent to which these processes are unique to humans and how they may have evolved in primates (Tu et al. 2015).

Nonhuman primates sometimes monitor mental states and adjust behavior as a result. Monkeys and apes opt out of difficult tests or seek more information before answering (for review see Kornell 2009; Roberts et al. 2012; Smith 2009). For example, in recognition memory tests monkeys chose to uncover a hidden sample on a computer screen before

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10071-021-01477-5.

✉ Ryan J. Brady
ryanbradyj@gmail.com

1 Emory University, Atlanta, USA

2 Yerkes National Primate Research Center, Atlanta, USA

Published online: 21 January 2021

choosing to see the test options (Beran and Smith 2011). While such studies demonstrate metacognition, they are so-called *one-shot* paradigms, where animals face a binary choice between a single act of information collection or proceeding directly to a test. The sample image is either present or it is not (Beran and Smith 2011); the food is either seen, or not seen (Call and Carpenter 2001; Hampton et al. 2004; Basile et al. 2015), and subjects are only allowed a single opportunity to collect more information. Thus, these studies, as informative as they are about the occurrence of metacognition, did not assess ongoing, iterative feedback between metacognitive monitoring and control in nonhuman primates.

To determine the extent to which monkeys engage in iterative feedback between monitoring and control, Tu et al. (2015) developed a paradigm in which each information-seeking response revealed a small amount of information, and the information-seeking response could be used multiple times in a trial. Monkeys classified images as depicting birds, fish, flowers, or people. At the beginning of the trial, the image was occluded by a grey blocker, and monkeys touched a button which gradually removed parts of the grey blocker with each touch. Monkeys could choose to make their classification response at any time or to continue to reveal more of the image. Thus, this paradigm tested whether monkeys monitored the information accumulated after each reveal response to control when they made their classification response. In one experiment, monkeys made more reveal responses when each response removed smaller blocks. In a second experiment, the image to be classified was shrunken so that it fit under a single critical block, such that when other blocks were removed, they did not reveal any useful information. Monkeys generally made reveal responses until the critical block was removed. In a third experiment, monkeys made more reveal responses when the information provided was insufficient compared to excessive (Tu et al. 2015). These results provide initial evidence that monkeys may dynamically monitor their ongoing decision process and iteratively seek more information when needed.

Alternative non-metacognitive explanations for these findings remain to be addressed (Tu et al. 2015). For instance, in Experiment 1, because different blocker sizes were introduced sequentially, it is possible that monkeys learned that some fixed number of touches was required in order to get trials correct in each blocker size condition, rather than learning to monitor the information available after each reveal response per se. In Experiment 2, because only one block contained critical information, monkeys might have learned to keep using the reveal response until a nonblank blocker was revealed. Finally, in Experiment 3, monkeys tended to "overshoot," making some revelation responses even when they should not have been necessary, raising the concern that use of the reveal response may not have been tightly coupled with monitoring the status of their decision process.

The current study was designed to test whether the findings of Tu et al. (2015) would be obtained when the amount of information revealed by each reveal response is unpredictable. By making the amount of information revealed by each information-seeking response variable, we assessed the extent to which monkeys engaged in on-going monitoring of information as it accumulated, seeking information repeatedly, or making a choice, as the current state of their knowledge dictated. We conducted three experiments in which monkeys used a *brighten response* to brighten to-be-classified images that initially started dark. The extent to which the images brightened after each brighten response was unpredictable, and the monkeys could choose to classify the image at any time during the trial. We hypothesized that if monkeys monitored the ongoing accumulation of information, and information-seeking behavior was controlled by this metacognitive assessment, then they should make more brighten responses on trials where brightness accumulated more slowly. In an important control test, monkeys viewed images that could not be classified, even though they became brighter with each brighten response. If monkeys monitored the evolution of their classification decision, these stimuli should elicit a high number of brighten responses because the monkeys could not make a decision even though the image got brighter and brighter. In a final test, we determined whether monkeys could distinguish between easy and difficult trials.

Experiment 1: monkeys learned to use the brighten response

Experiment 1 tested whether monkeys would learn to use the brighten response to brighten the image until they knew which category to assign a central image. When monkeys touched the location of the sample image, the image would brighten, making it easier to classify. We presented monkeys with two types of trials, *start-dark* trials in which the sample image started completely dark so there was no way of answering accurately without using the brighten response, and *start-bright* trials, in which the sample image started at full brightness. Monkeys received sessions that included start-dark and start-bright trials until they became highly accurate on start-dark trials, by using the brighten response. It is possible that during this training, the monkeys learned the non-metacognitive technique of using the brighten response ballistically, touching it as many times as possible, or a fixed number of times, irrespective of whether or not they had enough information to make a classification response. Alternatively, monkeys could learn to use the brighten response based on their metacognitive assessment of the status of their classification decision, making the image brighter only until they could reliably classification the sample image. We tested two hypotheses. First, if monkeys are metacognitive, then they will use significantly more brighten responses when the image starts dark compared to when the image starts bright. Second, if the use of the brighten response is driven by the amount of information currently available, then monkeys will make significantly fewer brighten responses when the first brighten responses increases the brightness by a large amount, compared to when the first brighten response increases brightness by a small amount.

Subjects and apparatus

We tested five adult male rhesus monkeys (*Macaca mulatta)*, each with a computer attached to the front of their cage. Monkeys were individually housed due to social incompatibility. Laptop computers ran custom programs written in Visual Basic (Microsoft, Redmond, WA), displaying on 15-in color LCD touch-sensitive screens (ELO Touch Systems, Menlo Park, CA) with a resolution of 1024, 768 pixels. Nutritionally balanced primate pellets (Test Diet, St. Louis, MO) were delivered by food dispensers (Med Associates, St. Albans, VT) into food cups below the screen. Monkeys received a full ration of food each day, with ad libitum access to water. Testing took place between 10 am and 5 pm 6 days a week. All monkeys had extensive experience with automated cognitive testing using touch-screen computers.

Stimuli

600 color images, 150 from each of four categories: birds, fish, flowers, people, were used. Images were collected from the online photo downloader Bulkr (Antibody software, 2020) and visually checked afterwards to ensure that each image contained at least one exemplar from only a single category. Duplicates were removed using DupDetector (Prismatic software, 2020) and visual inspection.

Procedure

Monkeys had previously been trained to classify images into the four categories used in these experiments. Each trial began when the monkey touched a green square (100 × 100 pixels) at the bottom of the screen. After two touches (FR 2), a white outline of a square appeared in the middle of the screen (350 × 450 pixels, Fig. 1) surrounded by four-choice icons that corresponded to the four categories. The four icons always appeared in the same corners. On *start-bright* trials, the to-be-classified image (400 × 300 pixels) appeared at 100% brightness inside the white box outline immediately after the monkeys initiated the trial. On *start-dark* trials, the to-be-classified image was initially completely black. Touching twice inside the white outline constituted a brighten response. Each brighten response on start dark trials caused the image to brighten by a randomly determined amount within the bounds of 0.01 and 0.05. These units correspond to brightness values in Visual Basic. A value of 1.0 is the highest brightness setting, however, an image can be identified at much lower values (see Fig. 4). After a brighten response is made, the white outline surrounding the image flashed off for 500 ms during which time additional touches had no effect. Monkeys were not limited in the number of brighten responses they could make, however, once the image was at a brightness equal to that of the image when collected, it could not be brightened further. Thus, monkeys could make the brighten response on "start bright" trials, and these responses were recorded but had no effect. Monkeys classified images whenever they chose, using the icons in the screen corners. Sessions consisted of 600 trials, with 480 start-dark trials and 120 start-bright trials. Every

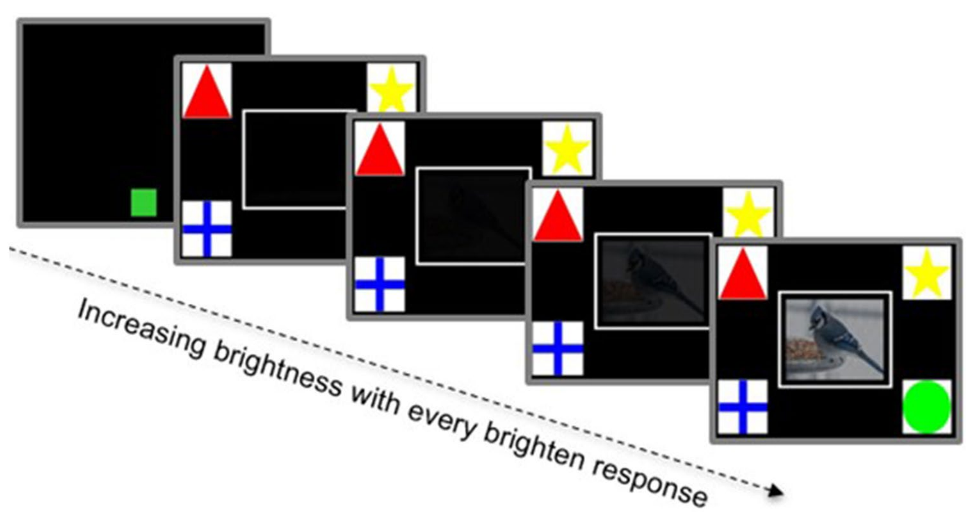

Fig. 1 Information seeking paradigm. Monkeys touched the green square to begin a trial. On start-dark trials, shown here, the sample image started totally dark. On start-bright trials, the image started with full brightness. On start dark trials, touching within the white box outline caused the image to brighten an unpredictable amount. Monkeys could choose to classify the images at any time

block of five trials included one start-bright trial and four start-dark trials with the order determined pseudo-randomly. Monkeys received these sessions until they reached a criterion of 85% correct classifications on start-dark trials in two consecutive sessions.

Data analysis

The primary dependent variable in these experiments was the number of brighten responses. We calculated the average number of brighten responses monkeys made before categorizing on start dark and start bright trials in the last two sessions of training, during which accuracy was above the 85% criterion level. We also calculated the number of brighten responses made after the first response, as a function of how much brighter the image became after this first response. Proportion correct scores were arc sin transformed prior to analysis (Aron and Aron 1994).

Results and discussion

Three monkeys met criterion after three sessions, one after seven and one after ten sessions. Monkeys made significantly more brighten responses before choosing to classify the image on start-dark trials compared to start-bright trials (Fig. 2 left panel; paired t test: t $(4) = -10.1$, $p = 0.001$, $d = 4.5$). Monkeys also made significantly more brighten responses when the first brighten response increased brightness by low amounts compared to high amounts (Fig. 2; repeated measures ANOVA main effect of brightness value: F $(4,16) = 46.5$, $p < 0.001$, $\eta = 0.92$, paired t test between brightness levels 0.01 and 0.05: t $(4) = 12.9$, $p < 0.001$, $d = 5.78$). These results suggest that monkeys used the brighten response based on the status of their decision process, and the information available, and did not adopt a non-metacognitive strategy of touching the screen a particular number of times before responding. If they had employed such a non-metacognitive strategy, the number of brighten responses made would not have varied as a function of how much change in brightness resulted from each response.

Experiment 2: monkeys made more brighten responses when the sample image could not be classified

In Experiment 1 we found evidence consistent with monitoring the status of an internal decision process, but we did not rule out the alternative explanation that the use of the brighten response was controlled by brightness per se. The fact that the brightness values at which monkeys chose to take the test was variable suggests that when they chose to classify depended more on the status of their decision than on brightness. However, it is possible that monkeys used the brighten response until brightness fell within a range, and then they classified. By this account, use of the brighten response was controlled by brightness, rather than by the status of a decision process. To further distinguish between the metacognitive and the brightness account of our monkeys' behavior, we replicated Experiment 1 with the addition of a small proportion of start dark trials with an image that did not belong to any of the four categories. These *start-unclassifiable* trials allowed us to distinguish between the use of the brighten response to achieve a certain brightness and use of this response until sufficient information was available to classify. If use of the brighten response was controlled by the accumulation of information guiding classification, then monkeys should generalize their use of the brighten response to a novel condition where their

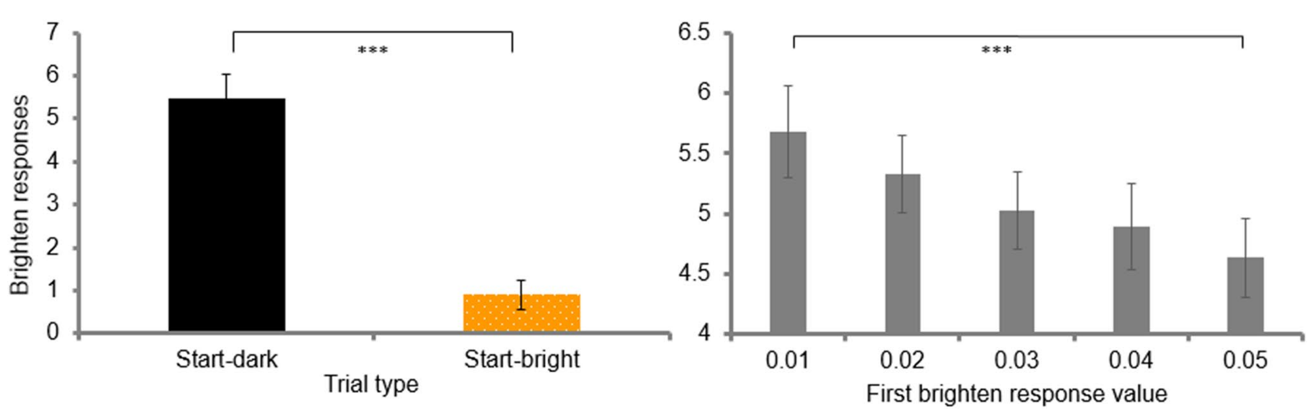

Fig. 2 Monkeys made more brighten responses when the sample started dark or increased in brightness by lesser amounts. When the sample started at full brightness, monkeys rarely chose to make a brighten response (left panel). Monkeys used the brighten response significantly more when the first brighten response resulted in comparatively little brightness (right panel). Error bars represent ± SEM

inability to reach a classification decision is not caused by the image being dim. Metacognitive monkeys should make significantly more brighten responses on unclassifiable trials than on comparable trials with images that can be classified. Alternatively, if the brightness of the image controlled use of the brighten response, monkeys should stop making brighten responses once the image was as bright as on normal trials, then guess at the classification.

Subjects and apparatus and stimuli

Subjects and stimuli were from Experiment 2. For unclassifiable trials, 150 new images were collected using bulkr software. These images were screened by two humans to ensure that they did not include or resemble one of the four categories.

Procedure

Monkeys were tested in five sessions of 600 trials each with 75% start-dark trials, 20% start-bright trials, and 5% start dark unclassifiable trials. Every block of 20 trials contained 15 start-dark, 4 start-bright, and 1 start-unclassifiable trial pseudo-randomly distributed. The unclassifiable images switched to a classifiable image when the monkey had made double the mode number of brighten responses they made on start dark trials in Experiment 1 so that monkeys could potentially earn a reward on each trial. This *switch-value* was determined by taking the last two sessions of each monkey's data from Experiment 1 and doubling the brightness value at which the monkey chose to take the test most often.

Results and discussion

Monkeys used the brighten response significantly more before classifying on start-unclassifiable trials compared to both start-dark and start-bright trials (Fig. 3 left panel; repeated measures ANOVA main effect of condition: F $(2,8)=58.8$, $p \leq 0.001$, $\eta=0.94$; pairwise comparisons Bonferroni corrected alpha = 0.025, start-bright vs start-dark: MD = − 4.94, $p \leq 0.001$, start-dark vs. start-unclassifiable: MD = − 2.88, $p=0.024$; start-bright vs. start-unclassifiable: MD = − 7.81, $p=0.001$). Furthermore, monkeys used the brighten response significantly more on start-dark trials compared to start-bright trials, replicating Experiment 1 (paired t test: $t\,(4)=-10.99$, $p<0.001$, $d=4.93$). Monkeys also used the brighten response significantly more when the first brighten response increased brightness by lesser amounts, also replicating Experiment 1 (Fig. 3 right panel; repeated measures ANOVA main effect of brightness value: $F\,(4,16)=104.3$, $p<0.001$, $\eta=0.96$, paired t-test 0.05 vs 0.01: $t\,(4)=11.26$, $p<0.001$, $d=5.03$). These results indicate that use of the brighten response was controlled by the status of monkeys' decision processes, rather than on the brightness of images. If monkeys had learned to use the brighten response to reach some criterion level of brightness, rather than a criterion level of information, the number of brighten responses made should not have been greater when the sample image was unclassifiable.

It is possible that monkeys learned that unclassifiable images switched to classifiable images if the brightness level exceeded a certain amount. Using this strategy would require that monkeys detected an unclassifiable image, then used the brighten response until the image changed. While this would not indicate that monkeys made all their brighten responses based on when they knew the answer, this explanation would

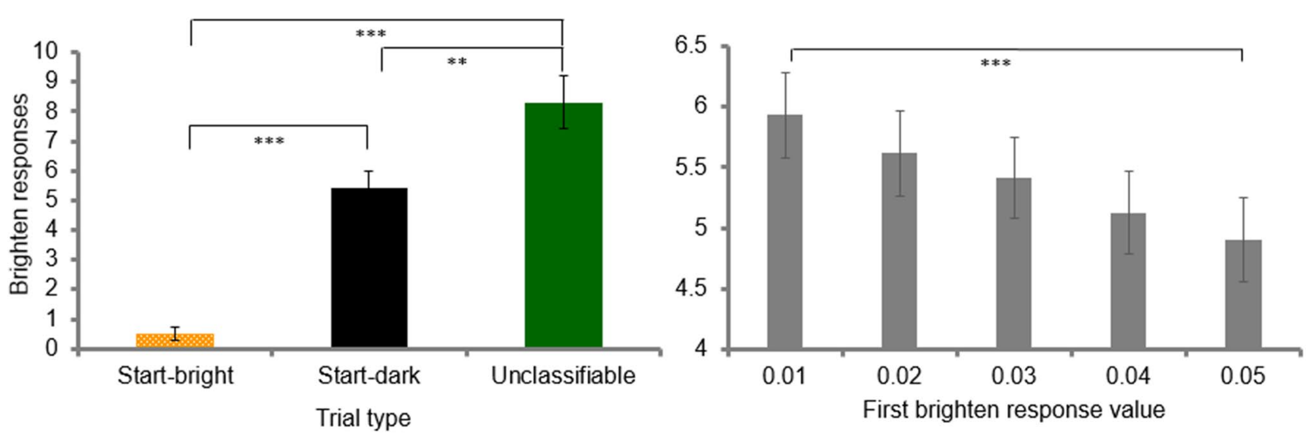

Fig. 3 Monkeys made more brighten responses when the sample was unclassifiable and replicated findings from Experiment 1. Brighten responses on start-bright, start dark, and unclassifiable trials (left panel). Average number of brighten responses as a function of brightness achieved by the first brighten response (right panel). Error bars represent ± SEM

still be consistent with some kind of metacognitive monitoring and control because monkeys would have to recognize the image as unclassifiable, inhibit making a classification choice, and use the brighten response instead. To further test whether the use of the brighten response was controlled by the status of monkeys' classification decision, we empirically determined how accurate monkeys were at each brightness level on average, while concurrently measuring how often they made brighten responses at these levels.

Experiment 3: monkeys made fewer brighten responses when their ability to accurately classify was high

In Experiments 1 and 2, we found evidence suggesting that the monkeys' use of the brighten response was not controlled by the brightness of the image, but rather by an internal assessment of whether or not they could classify the image. In Experiment 3, we directly tested whether the use of the brighten response was controlled by the status of the classification decision. We added *fixed-value* probe trials where the first use of the brighten response would set the brightness of the image to a predetermined brightness value that could not be changed by additional brighten responses. These probe trials allowed us to objectively measure the monkeys' ability to classify at each brightness level. We measured their subjective assessment of their ability concurrently through their use of the brighten response. We hypothesized that if monkeys chose to make their classification response because they knew they knew the answer, then the likelihood that monkeys used the brighten response should be low for brightness levels for which accuracy was high.

Subjects and apparatus and stimuli

The same subjects, equipment, and images used in Experiment 1 were used in Experiment 3.

Procedure

Monkeys received five sessions of 600 trials each with 120 start-bright trials, 120 start-dark fixed-value trials, and 360 start-dark regular trials. Every block of five trials contained one start bright trial, one start-dark fixed value probe trial, and three start-dark regular trials, presented in a pseudorandom order. On fixed-value trials, the trial resembled a start-dark trial, but when monkeys made the first brighten response, the brightness increased to one of four fixed values: 0.02, 0.04, 0.08, 0.16 (Fig. 4). The order in which these values occurred was pseudo-randomly chosen such that each value occurred once before repeating. If monkeys used the brighten response after the image had brightened to the fixed value, the white box outline disappeared and reappeared as it always did, and the brighten responses was recorded, however, the brightness remained constant until the monkey made a classification response. On start-dark regular trials each brighten response increased the brightness according to the fixed-value trial brightness levels. The first brighten response increased the brightness to 0.02, the second to 0.04, the third to 0.08, and the fourth to 0.16. If the monkey made a fifth brighten response, the brightness increased to 0.32 and further brighten responses had no effect.

Results and discussion

Monkeys were less likely to seek additional information on easy trials. The brighter the image that appeared after the first brighten response on start-dark fixed value trials, the less likely monkeys were to make a brighten response (Fig. 5; repeated measures ANOVA main effect of brightness: F (3,12) = 15.6, $p < 0.001$, $\eta = 0.80$). Monkeys were also more accurate on brighter start-dark fixed value trials (Fig. 5; repeated measures ANOVA main effect of brightness: F (3,12) = 40.5, $p < 0.001$, $\eta = 0.91$). Furthermore, monkeys used the brighten response more often on start-dark trials compared to start-bright trials, replicating Experiments 1 and 2 (paired t test: t (4) = 12.04, $p < 0.001$, $d = 5.39$). These results suggest that monkeys adjust their information seeking based on whether they can make an accurate classification response with the information

Fig. 4 Progression of the fixed brightness values used in Experiment 3. On fixed-value trials, the trial started dark, and the first brighten response set the brightness to one of the fixed values. Further brighten responses did not change the brightness. On start-dark regular trials, each brighten response increased the brightness according to the fixed values

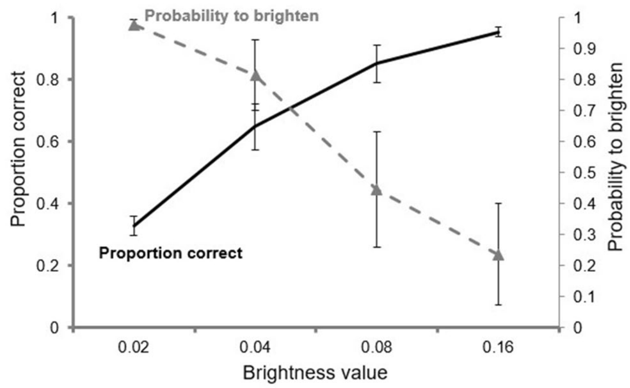

Fig. 5 Monkeys were less likely to use the brighten response when their ability to accurately classify was high. Shown are accuracy (solid line) and probability of making at least one brighten response (dashed line) on fixed-value trials. Error bars represent ± SEM

available. Combined with the results of Experiments 1 and 2, these results provide strong evidence against alternative non-metacognitive explanations such as the monkeys learning to use the brighten response a fixed number of times or using the brighten response until a threshold of brightness was reached.

Because we matched the brightness values on regular trials to the values of the fixed value trials, we were able to compare the probability that monkeys used the brighten response at a given brightness value when (1) the monkey had to use the brighten response multiple times already (regular trials) and when (2) the monkey had used the brighten response just once (fixed-value trials). We found that the probability that monkeys used the brighten response was greater on regular trials when the monkey had already used the brighten response multiple times, compared to the same brightness value on fixed-value trials (Supplemental Fig. 1; repeated measures ANOVA (condition × difficulty interaction: F (3,12) = 3.79, $p = 0.04$ $\eta = 0.49$; main effect of condition: F (1,4) = 8.33, $p = 0.045$, $\eta = 0.68$, main effect of difficulty: F (3,12) = 21.3, $p < 0.001$, $\eta = 0.84$). This comparison between the regular and fixed value trials suggests that while our monkeys did use the brighten response based on whether they knew the answer, they also showed some "overshooting" or behavioral momentum, as observed by Tu et al. (2015). Particularly on trials on which the monkeys had already made multiple brighten responses, they were likely to make more additional responses than they would when confronted with the same image on a trial on which they had not made multiple brighten responses. From these results, we cannot determine whether this reflects a failure of monitoring, such that monkeys did not always attend to the effect of each brighten response, or a lack of inhibitory control such that once they got started making brighten responses, they could not easily stop.

General discussion

Across three experiments, we found that monkeys used the information-seeking response adaptively, in a way that was contingent on the amount of information already available. In Experiment 1, monkeys made more brighten responses the dimmer the initial image and made fewer brighten responses the more brightness had accumulated after their first response. In Experiment 2, monkeys used the brighten response significantly more when the image was unclassifiable, compared to regular trials. This indicates that use of the brighten response was controlled by monitoring evolving decisions, rather than by external cues such as brightness alone. In Experiment 3, monkeys sought more information at brightness values where they could not accurately classify the image. This provides converging evidence that use of the brightness response was controlled by whether the monkeys could accurately classify the image, such that when their accuracy was high, the likelihood they chose the brighten response was low. Taken together, these results indicate that monkeys iteratively monitored the ongoing accumulation of information or their decision process. The results of this monitoring engaged cognitive control to delay making a classification response when they were not confident. This resulted in them seeking more information until they were able to answer accurately.

The evidence presented here indicates that monkeys iteratively monitor their metacognitive state over the gradual accumulation of information, resembling the dynamic relationship between metacognitive monitoring and cognitive control as it exists sometimes in humans. The critical difference between our paradigm and previous information-seeking paradigms (Basile et al. 2009; Beran and Smith 2011; Call and Carpenter 2001; Castro and Wasserman 2013; Hampton 2009; Hampton et al. 2004; Tu et al. 2015) was that each information-seeking response added an unpredictable amount of information. This was an important addition for at least two reasons. First, monkeys typically had to seek information more than once to accurately respond, moving away from typical "one-shot" information-seeking paradigms. Second, the unpredictable schedule of information increase prevented monkeys from learning a rote strategy based on making a particular number of responses. Thus, we ruled out leading alternative non-metacognitive explanations of our results and those of Tu et al. (2015).

Understanding why and how metacognitive monitoring evolved begins by asking what a metacognitive organism can do that one without metacognition cannot do (Hampton et al. 2020). Metacognitive monitoring may be of little value until it is used to guide cognitive control. After all, there is not much value in the assessment of knowledge, without the means to do something based on that assessment. It is likely that monitoring allows discrimination between knowing and not knowing (Hampton et al. 2020), and the relations of monitoring with cognitive control networks allow for the flexible choice of adaptive responses. How closely monitoring and control interact may therefore be critical in explaining similarities and differences in behavioral flexibility between species. For instance, while our results demonstrate a relationship between monitoring and control in monkeys, we also observed that monkeys in one condition tended to *overshoot*, or use the brighten response more than necessary, in line with previous findings (Tu et al. 2015). This finding perhaps illustrates how the control processes are not necessarily engaged immediately upon receiving input from the monitoring system.

The overshooting we observed may be explained by the low cost in choosing to make an extra brighten response to ensure accuracy (Tu et al. 2015). Apes and children also show excessive information seeking when the cost of doing so is low, or when the reward at risk is high (Call and Carpenter 2001; Marsh and MacDonald 2012). Interestingly, however, in the current experiment we observed more overshooting behavior on trials in which the monkey had to use more brighten responses to reach a given brightness value, compared to when the first brighten response gave that same value. Apparently, once the monkeys "got going" using the brighten response, it was more difficult to stop, compared to if they had not started using it already. This observation may suggest that the overshooting behavior was due to a lack of inhibitory control, rather than a strategy related to the ease of the response or the value of the reward. It is possible that the monkeys' metacognitive monitoring system recognized that a classification decision had been reached, but the control system took extra time to change the behavioral response from information seeking to classification. Capuchin monkeys tested on similar paradigms do not perform as flexibly as rhesus monkeys and rarely show adaptive information seeking behaviors (Basile et al. 2009; Beran and Smith 2011; Fujita 2009). This may be explained by a relative lack of inhibitory control in this species. It could also be due to differences in monitoring abilities (Beran and Smith 2011). Behavior may fail to reflect metacognition immediately either because the monitoring system is not updating the assessment of knowledge, the monitoring system is updating knowledge but is not relaying this to the control system fast enough, or the control system receives the signal to change behavior but is not changing behavior fast enough. Any of these explanations would result in overshoot in brighten responses. Thus, future studies should more closely investigate the relationship between metacognitive monitoring and inhibitory control abilities.

Comparative studies so far have drawn many parallels between human and nonhuman primate metacognitive abilities. Monkeys show metamemory, or the ability to monitor the presence or absence of a memory (Brown et al. 2017, 2019; Hampton 2001; Templer et al. 2019). Monkeys and apes both seek information when they do not know where a valued item is hidden (Basile et al. 2015; Call and Carpenter 2001; Hampton et al. 2004). Rhesus monkeys but not capuchin monkeys show information seeking in various computerized situations (Beran and Smith 2011) and monkeys spontaneously show memory monitoring without training (Hampton and Hampstead 2006; Rosati and Santos 2016). Furthermore, monkeys are subject to metacognitive illusions similar to humans (Ferrigno et al. 2017). Our findings add to this growing literature, demonstrating an iterative feedback loop between cognitive monitoring and cognitive control that allows for adaptive information-seeking behavior.

Acknowledgements We thank Tara A. Dove-VanWormer for help with monkey testing.

Funding This work was supported by National Science Foundation Grants BCS-1632477 and BCS-1946767, and by ORIP/OD P51OD011132.

Compliance with ethical standards

Conflicts of Interest Authors declare no conflicts of interest.

Availability of data and material Available upon request.

References

Aron A, Aron E (1999) Statistics for psychology. Prentice Hall, Upper Saddle River, NJ

Basile BM, Hampton RR, Suomi SJ, Murray EA (2009) An assessment of memory awareness in tufted capuchin monkeys (*Cebus apella*). Anim Cogn 12(1):169–180

Basile BM, Schroeder GR, Brown EK, Templer VL, Hampton RR (2015) Evaluation of seven hypotheses for metamemory performance in rhesus monkeys. J Exp Psychol Gen 144(1):85–102

Benjamin AS, Bjork RA, Schwartz BL (1998) The mismeasure of memory: when retrieval fluency is misleading as a metamnemonic index. J Exp Psychol Gen 127(1):55–68

Beran MJ, Smith JD (2011) Information seeking by rhesus monkeys (*Macaca mulatta*) and capuchin monkeys (*Cebus apella*). Cognition 120(1):90–105

Beran MJ, Brandl J, Perner J, Proust J (eds) (2012) Foundations of metacognition. Oxford University Press, Oxford

Brown EK, Templer VL, Hampton RR (2017) An assessment of domain-general metacognitive responding in rhesus monkeys. Behav Proc 135:132–144

Brown EK, Basile BM, Templer VL, Hampton RR (2019) Dissociation of memory signals for metamemory in rhesus monkeys (*Macaca mulatta*). Anim Cogn 22(3):331–341

Call J, Carpenter M (2001) Do apes and children know what they have seen? Anim Cogn 3(4):207–220

Castro L, Wasserman EA (2013) Information-seeking behavior: exploring metacognitive control in pigeons. Anim Cogn 16(2):241–254
Dunlosky J, Bjork RA (2008) The integrated nature ofmetamemory and memory. In: Dunlosky J, Bjork RA (eds) Handbook of metamemory and memory. Psychology Press, New York, NY, pp 11–28
Ferrigno S, Kornell N, Cantlon JF (2017) A metacognitive illusion in monkeys. Proc R Soc B Biol Sci 284(1862):20171541
Flavell JH (1979) Metacognition and cognitive monitoring: a new area of cognitive–developmental inquiry. Am Psychol 34(10):906–911
Fujita K (2009) Metamemory in tufted capuchin monkeys (*Cebus apella*). Anim Cogn 12(4):575–585
Hampton RR (2001) Rhesus monkeys know when they remember. Proc Natl Acad Sci 98(9):5359–5362
Hampton RR (2009) Multiple demonstrations of metacognition in nonhumans: converging evidence or multiple mechanisms? Comp Cogn Behav Rev 4:17–28
Hampton RR, Hampstead BM (2006) Spontaneous behavior of a rhesus monkey (*Macaca mulatta*) during memory tests suggests memory awareness. Behav Proc 72(2):184–189
Hampton RR, Zivin A, Murray EA (2004) Rhesus monkeys (*Macaca mulatta*) discriminate between knowing and not knowing and collect information as needed before acting. Anim Cogn 7(4):239–246
Hampton RR, Engelberg JW, Brady RJ (2020) Explicit memory and cognition in monkeys. Neuropsychologia 138:107326
Koriat A (2000) The feeling of knowing: Some metatheoretical implications for consciousness and control. Conscious Cogn 9(2):149–171
Kornell N (2009) Metacognition in humans and animals. Curr Dir Psychol Sci 18(1):11–15
Marsh HL, MacDonald SE (2012) Orangutans (*Pongo abelii*) “play the odds”: Information-seeking strategies in relation to cost, risk, and benefit. J Comp Psychol 126(3):263
Metcalfe J (2009) Metacognitive judgments and control of study. Curr Dir Psychol Sci 18(3):159–163
Nelson, Narens (1990) Metamemory: a theoretical framework and new findings. Psychology of learning and motivation, vol 26. Academic Press, pp 125–173
Proust J (2007) Metacognition and metarepresentation: is a self-directed theory of mind a precondition for metacognition? Synthese 159(2):271–295
Roberts WA, McMillan N, Musolino E, Cole M (2012) Information seeking in animals: metacognition? Comp Cogn Behav Rev 7:85–109
Rosati AG, Santos LR (2016) Spontaneous metacognition in rhesus monkeys. Psychol Sci 27(9):1181–1191
Shimamura AP, Squire LR (1986) Memory and metamemory: a study of the feeling-of-knowing phenomenon in amnesic patients. J Exp Psychol Learn Mem Cogn 12(3):452–460
Smith JD (2009) The study of animal metacognition. Trends Cogn Sci 13(9):389–396
Templer VL, Gazes RP, Hampton RR (2019) Co-operation of long-term and working memory representations in simultaneous chaining by rhesus monkeys (*Macaca mulatta*). Q J Exp Psychol 72(9):2208–2224
Tu HW, Pani AA, Hampton RR (2015) Rhesus monkeys (*Macaca mulatta*) adaptively adjust information seeking in response to information accumulated. J Comp Psychol 129(4):347–355