

Does Implicit Metacognition Provide a Tool for Self-Guided Learning in Preschool Children?

Frances Balcomb (francesb@u.arizona.edu)

Department of Psychology, 1303 E University Blvd
Tucson, AZ 85719 USA

LouAnn Gerken (gerken@u.arizona.edu)

Department of Psychology, 1303 E University Blvd
Tucson, AZ 85719 USA

Abstract

Many core cognitive processes, such as associative learning, are thought to appear early in infancy and develop rapidly throughout childhood, while more complex processes such as metacognition are thought to appear significantly later. However, there is an apparent paradox in that the rapid growth in cognitive development cannot be explained by passive learning mechanisms only, yet for children to actively guide their own learning, they must be at least somewhat self-referential. In the experiments reported here, metacognition was assessed using non-verbal tasks originally developed for non-human animals. Children aged 3.5 were able to demonstrate metacognition by strategically responding on a task in which they had to access their knowledge states. In addition, results indicate that children whose metacognitive assessments more closely matched their actual memory performance also showed superior overall learning compared to children whose metacognitive assessments were less accurate, suggesting that metamemory skill is correlated with memory itself. The results imply that metacognition may be intrinsically related to core cognitive processes and implicitly emergent in children as the processes themselves emerge, providing an explanation for how they are able to guide learning, even as infants.

Introduction

A key component of learning involves seeking out and encoding of new information and avoiding relearning already known information. In fact, people are active learners, and the ability to guide learning allows selectivity and efficiency; without it, reliance on passive mechanisms would result in redundancy, disorganization, coding of circumstantial but unrelated conjunctions, and slow learning of significant ones. For example, in language paradigms there is evidence that adults are able to use sensitivity to different types of information (such as perceptual salience, co-occurrence, or predictive values) in ongoing stimulus streams to optimize learning depending on the characteristics of the input (Gomez, 2002). We know, too, that infants are especially good at selective learning at very early ages. Indeed, infants' differential preference for what is novel vs. what is familiar in commonly-used testing methods is typically attributed to their seeking an optimal level of new and old information to promote learning (Hunter, Ames, & Koopman, 1983). Rather than being merely associative learners and pattern matchers, they can

select 'good' from 'bad' information sources (Namy & Waxman, 2000), and informationally rich from coincidental co-occurrences as infants (Aslin, Saffran, & Newport, 1998; Newport & Aslin, 2004). How infants accomplish these feats of learning remains a puzzle, and the questions about underlying mechanisms are debated in a variety of fields including computational modeling, linguistics, and cognitive development.

What is proposed here is that executive function skills, including metacognition, may be inherently intertwined with core processes, existing implicitly at very early ages and only emerging explicitly when children are much older. Thus infants can accomplish strategic learning by influencing control over processes like memory encoding through such cognitive behaviors as sustaining attentional focus. These behaviors are influenced by implicit access to internal states of knowledge, with the goal of most efficiently maximizing learning while avoiding redundancy. This implicit monitoring and control of core processes reflects procedural metacognition.

Attribution of metacognition in this way, therefore, does not imply a complex system with two distinct cognitive mechanisms, one to perform a core function and another to evaluate and monitor. What is proposed here is that metacognition may be implicit, and occurs as an integral part of the cognitive process itself. For example, metamemory may be a resultant process of memory, such that during the retrieval process some items are marked as 'more certain' and others as 'less certain' either by familiarity, a sense of recognition, a threshold of recall strength, or other factor. If, indeed, a feeling of knowing is subserved by a simple mechanism integrated into the memory process itself, we might begin to understand how very young children appear well able to capitalize on information in memory to learn rapidly and efficiently, without demonstrating explicit awareness of their states of knowledge.

Metacognition in Children

In general many researchers believe that preschoolers do not have very sophisticated metacognition skills. Young children are not credited with strategic responding based on self-assessments until late preschool at the earliest, despite

having developed an extensively organized knowledge base and cognitive strategies by then. Much research has demonstrated a significant effect of age on a variety of metacognitive tasks in school age children, indicating that by age 3, children are reliably able to show some early metacognitive indicators, like the use of verbs 'think' and 'know' (Kuhn, 2000). However, it is not until at least age four, and more typically at older ages, that children are able to demonstrate success with standard metacognition tasks relying on explicit reporting. Performing in these tasks, children demonstrate continued development until middle school. For example, in one study, kindergarten, first, and third grade children were presented with line drawings, and asked to name them. If a child failed to correctly name an item s/he was asked if s/he would recognize the correct label if s/he heard it. Subsequently, children were presented with an array of nine pictures and one verbal label, and asked to point to the matching picture. Children of all ages were able to predict, to some degree, which picture-label matches they were able to recognize, with a significant increase in accuracy with age (Wellman, 1977).

More recent research has demonstrated that children as young as four years old are able to accurately predict recognition performance to some degree when the stimuli used are more familiar (pictures of other children either well-known, less well-known or unfamiliar) and the task demands are less (Cultice & Somerville, 1983). Furthermore, recent work has indicated that children's metacognition estimates may be confounded by actual memory contents (Butterfield, Nelson, & Peck, 1988), or by confusion between wishful thinking and realistic estimates of performance (Schneider, 2000). Finally, although most developmental research supports the idea that metacognitive skills are more accurate in older children than younger children (Flavell & Wellman, 1977; Kuhn, 2000; Schneider, 1999), there are some exceptions (Butterfield et al, 1988; Lockl & Schneider, 2002)

These results suggest that metacognition may exist at younger ages than has traditionally been thought and maybe highly intertwined with source knowledge in children, rather than emerging separately upon already established structures.

Metacognition in Non-humans

The apparent lack of directed learning or metacognitive ability in human children is surprising, given that many non-human animals demonstrate 'uncertainty' behaviors and respond strategically. This strategic responding may be a demonstration of metacognition, and has been studied with animals in memory tasks. Recent research with rhesus monkeys has indicated that they are able to manifest at least rudimentary metacognition skills.

Metacognition has been studied in a variety of nonhuman animals including dolphins, pigeons and monkeys. In a metacognition study with rhesus monkeys Smith et al. (1998) presented a serial list to adult humans and monkeys and gave them a choice of responding to the task, or

choosing an 'uncertain' response. Humans were more likely to choose the 'uncertain' response for items in the middle of the list (more difficult) than items in the beginning and end of the list, while monkeys were more variable in their success. In a subsequent study animals were trained on a visual matching task using paired polygon shapes (Shields 1999, as cited in Smith, Shields, & Washburn, 2003). At test, one member of a pair was presented, and the monkey was given the option to take the task for the risk of a food reward or a long time out, depending on success or failure, or to skip it for a lesser food reward but no time out. When accuracy on skipped trials was compared to that of trials taken, the monkeys did not show differential performance. In other words, monkeys did not show evidence of optimizing performance by skipping the items they didn't know. In a follow up study, monkeys were given the same task with the difference that on test the monkey saw not only the target, but also the other member of the pair, along with an incorrect foil. In this condition, in which all the information was present, monkeys were successful at choosing to take only the trials on which they were more accurate, indicating at least some ability to monitor their own states of knowledge. The work with non-humans suggests that we need to re-examine our view that metacognitive skill is late emerging in human children

To summarize, the results of metacognition in both young children and non-human animals suggest that they may have implicit access to internal knowledge states. If children are able to use this information to strategically direct learning in optimal trajectories, we could begin to explain their very rapid learning at very young ages. The purpose of the current research was to examine the implicit metamemory-memory relationship in preschool children using methods adapted from comparative literature.

Experiment I

This experiment was designed to determine whether young children would demonstrate the ability to assess their state of knowledge with a visual pairs task in which there were two stimuli sets, a familiar and a less familiar set. In this task children were exposed to some visual pairs multiple times and exposed to others only once. Thus all pairs consisted of familiar animals, but some should have been better learned and thus better known than others. At test, children were shown one item of a pair and chose to either accept or decline taking a trial in which they selected the mate of that item. It was predicted that children would choose to accept more trials with more-familiar versus less-familiar pairs.

Subjects

Subjects were 9 children aged 3:0-4:2 (m=3:6).

Materials

Stimulus materials consisted of 20 visual pairs presented on a DVD. A pair consisted of colored line drawing of an

animal paired with a picture of a common object like a bicycle. Audio files for training and test stimuli were recorded by a female native English speaker using Sound Studio 2.1.1. Pairs were presented in movie format created with iMovie HD on a Macintosh Mini computer. First the picture of the animal appeared with a sentence giving its name (e.g. 'this is Andy.') Next an object appeared with a sentence describing the animal-object relationship (e.g. 'He likes to drive a fire truck.') Finally, the animal and object appeared together with an additional sentence about the relationship (e.g. 'He drives a fire truck to all the fires.') Each picture appeared for 4 seconds, with a 1 second gap between. Movies were burned to DVD. There were 2 DVD movies. The home DVD contained 10 animal-object pairs repeated three times. The lab DVD showed 20 animal-object pairs; the 10 familiar pairs, and 10 novel pairs, intermixed throughout. At test, children judged their knowledge of all 20 animal-object pairs. Thus, all pairs had been seen equally recently, but some were more familiar and some were less familiar, as defined by amount of viewing time. Test stimuli consisted of the animal-object pairs, and also included a line drawing of a colored arrow. The arrow served as an 'opt out' option, such that if children could touch the arrow to skip trials. Audio stimuli for test consisted of the questions 'Do you know what s/he likes?' and four types of positive (e.g. 'yay!'), one negative ('Uh-oh, that's not it!') and one neutral ('We don't know what s/he likes') feedback sound file.

Testing was conducted on a Dell Inspiron 1100 laptop with a Keytec Magic Touch touch-screen mounted on a 15" LCD upright monitor. The experiment was run using the DMASTR software developed at Monash University and at the University of Arizona by K. Forster and J. Forster

Procedures

The experiment consisted of a learning phase and test phase.

Learning

The DVD with ten items was sent to children's homes and parents were asked to ensure that their children watched all three repetitions. After that, they were allowed to watch it as often as they liked. When children came into the laboratory, they watched the 20-item DVD. Subsequently, they went into a different room for training and testing.

Testing

At test, children had two potential tasks. The judgment of memory (*judgment task*) was performed on every trial. Children were shown a picture of one of the animals from the learning phase along with an arrow. They touched the arrow on the touch screen if they thought they did not remember the item that had been paired with the animal (indicating a 'decline' response), which resulted in a new trial. They touched the animal if they thought they remembered the paired item (indicating 'accept'), which started the second task, the *mate selection task*. On the mate selection task, two pictures appeared, one of the item that

had been paired with the animal from the judgment task (the mate), and one of an item that had been paired with some other animal (the foil). If the child chose the correct response (the mate), s/he received positive feedback, but if s/he selected the incorrect target (the foil), s/he received negative feedback.

There were two training sets, to ensure that children understood the procedure. The first training consisted of 3 common animals and common associates (e.g. a dog and a bone, with eyeglasses as a foil), and 3 novel imaginary animals. The experimenter completed the first training with children, ensuring that they skipped the novel animals, and chose to take the task with familiar animals. To encourage children to transfer the procedures to the animals from the DVD, the second training consisted of two familiar animals from the DVD, and two less familiar. (These items were not used on the subsequent test.) Children completed the second training with the experimenter verbally reviewing the child's responses after each item. The actual test (judgment and mate selection) consisted of 14 pairs of animals and objects, which children completed without help. After each trial, the experimenter echoed the feedback given by the computer.

Results

The design of the experiment ensured that half of the pairs would be better known than the other half. Therefore, it was predicted that children would choose to accept trials involving familiar pairs significantly more often than trials involving less familiar pairs.

Only two children appeared to follow a simple rule of accepting all seven of the familiar items and declining all seven of the less familiar; the rest accepted some items from both sets. The number of 'accept' and number of 'decline' trials was totaled for each pair type (familiar and less familiar.) Children made more accept responses with more familiar animals ($M=5.4$ $SD=1.8$) than with less familiar animals ($M=3.2$, $SD=3.2$; $t(8)=1.86$, $p=0.04$, one-tailed). To determine whether children not only made more accept responses with more familiar animals but also got more of these trials correct, the percent correct was calculated for all familiar trials and compared to percent correct for unfamiliar trials. Results indicated that children were significantly more accurate on more familiar trial items ($M=91$, $SD=19$) than less familiar ($M=55$, $SD=35$; $t(11)=2.11$, $p=0.02$, one-tailed).

Discussion

The purpose of this study was to determine whether children were able to demonstrate metacognition non-verbally on a visual paired associates task, by choosing to accept test trials with items expected to be more accessible in memory.

Results indicate that children made significantly more accept responses with more familiar pairs than with less familiar pairs, supporting the existence of early metacognition. This result is similar to one of the studies done with monkeys using serial list learning, in which

monkeys were expected to skip memory items that were predicted to be more difficult (Smith et al., 1998.) Since monkeys did not show consistent results, but adults did, it was interesting to discover that young children can readily indicate their implicit knowledge state when asked in the right way.

However, it is possible that children's responses on the judgment task was based only on association of a particular response (accept or decline) with a type of stimulus presented (more familiar vs. less-familiar, or seen at home vs. not seen at home.) For example, when declining to take the trial for the animal named 'Andy,' children might not be able to search their memory for the appropriate paired associate and find it lacking, but rather might simply respond based on the feeling that they had not seen Andy very frequently. Given that children did choose a mix of both familiar and unfamiliar pairs, and their overall accuracy was quite high, this explanation is unlikely to provide a complete description of the results. However, this explanation could account for the main effect of more accept responses for more familiar pairs. Therefore, a follow up study was conducted to examine whether children's choice to decline or accept the task was based on judgments of knowing specific pairs, as opposed to pairs from different sets.

Experiment 2

The purpose of Experiment 2 was to further examine the extent to which children are able to indicate metacognitive competence by accepting or declining trials based on memory of matched pairs. In this study, memory of the paired associates was examined in a subsequent memory task in which children had to select one of two objects (the mate vs. a foil) on every trial. The second task was added to determine if children had poorer memory for the trials that they had declined in the metamemory task. In addition, Experiment 2 explored the possible relation between metamemory and memory by asking whether those children who made the most accurate assessments of their memory also showed better memory for the visual pairs than children who either overestimated or underestimated their knowledge of the pairs.

Subjects

Subjects were 25 children aged 3:5-3:7 (m=3:6).

Materials

Fifteen of the visual matching pairs used in Experiment 1 were selected and made into one movie in which all pairs appeared once. This movie was viewed for the first time when children arrived at the laboratory. Therefore, at test, all items had equal exposure. Test apparatus and the first training were identical to that of Experiment 1. The second training consisted of 4 animals; the first and last animals of the DVD (which, according to list learning effects should be better accessible to memory) and two animals from the middle of the movie, which should similarly be less

accessible. These items were not used in the subsequent tests. There were 11 test trials.

Procedures

The experiment consisted of a learning phase and two separate tests; the first test was the same two-task test as in Experiment 1 (the judgment task, and if the child made an accept response, the mate selection task). The second test was a separate recognition memory test (memory test), which assessed children's memory for the items they had declined in the combined test. (See Figure 1.)

Learning: In the instruction phase children watched the DVD once.

Testing: Children completed the metamemory test first, followed by the memory test.

Metamemory- test: Training and test procedures were the same as in Experiment 1 with the exception that rather than seeing only the animal and arrow, children saw a picture of the animal, an arrow, the object mate of the animal, and a foil (which was an associate for a different animal). The child was asked, 'do you know what s/he likes?' Immediately after the question, the mate and foil disappeared, and the child selected the animal or arrow with only those two images present. This design was modeled after the task with which rhesus monkeys were successful at indicating metacognitive competence on visual matching tasks (Shields et al, 1999). As previously discussed, monkeys were able to decline less well known items and accept more well known items only when they saw the target, the match and the foil at the time of choice.

Memory test: After the metamemory test, children were presented with an 11-trial recognition memory test. On each trial, children saw one of the 11 animals, the animal's object mate and another object that was the mate of a different animal. They were simply asked to find select which object had been paired with the animal.

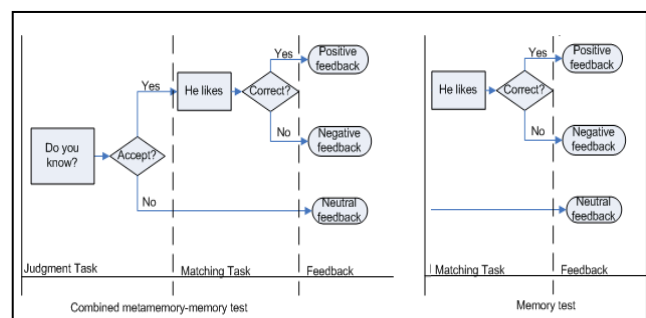


Figure 1: Test Procedures for Experiment 2

For some children, some trials of the memory test involved being tested on an animal a second time if they had accepted the trial involving that animal on the metacognitive test. The same positive and negative feedback sound files were used as in the metacognition test trials.

Results

One child declined all trials. Sixteen children accepted and declined at least one trial. Eight children accepted all trials.

The first analysis was conducted to examine if children demonstrated better accuracy on items they accepted than on those they declined on the metamemory test, indicating accurate judgment of their knowledge. For the 16 children who both accepted and declined at least one trial, the percent correct on the metacognitive test for accepted trials was compared to the percent correct on the memory test for items children declined on the metamemory. As predicted, children were significantly more accurate on 'accept' items ($M=80$, $SD=22$) than 'decline' ($M=64$, $SD=40$; $t(15)=1.83$, $p=0.04$, one-tailed).

A second analysis was conducted to examine whether better metamemory performance was related to better memory performance. Each child took all 11 memory trials, therefore the total number of correct responses was known. For each child, the number correct on the memory test was compared to the number of items the child expected to get correct, as indicated by the number of items s/he accepted, to determine overall metamemory performance. A difference score was calculated (total correct minus total accepted) and converted into z scores. Children were ranked as to how far their performance estimates differed from their actual performance, with zero (no difference) being optimal. Children were divided at the median in two groups, 'higher accuracy' (HA) and 'lower accuracy' (LA).

The accuracy on the memory test of HA versus LA children was compared. Results indicate that children in the HA group were significantly more accurate on the memory test ($M=83$ $SD=14$) than children in the LA group ($M=64$, $SD=24$; $t(20)=2.09$, $p=0.03$, two-tailed). To ensure that this difference was not an effect of the HA children learning during the metamemory task and applying what they learned to the subsequent memory task, a comparison was made of HA versus LA children for 'accept' trials on the metamemory task. Results indicate that children in the HA group were significantly more accurate on 'accept' trials ($M=84$ $SD=15$) than children in the LA group ($M=68$, $SD=27$; $t(17)=1.74$, $p=0.04$, two-tailed), indicating that children whose overall estimates of memory were better also had better recognition memory performance. Children who either over or underestimated their memory abilities showed poorer overall memory.

Discussion

These experiments indicate that young children are able to demonstrate evidence of implicit metamemory skills before the age at which they can verbalize about their knowledge,

and in addition, that metamemory may be linked to source memory itself.

The results are a first step toward addressing the paradox of excellent learning abilities in infants and young children with their here-to-fore assumed poor metacognitive skills. The picture that is beginning to emerge is one in which learners are able to keep track of what they do and do not know from very early in life. Some aspects of this ability can be tapped in non-verbal metacognition tasks such as ours by the pre-school years (and perhaps before). However, the ability to verbalize about the contents of one's memory develops later. Thus, knowing and being able to make use of the contents of one's mind represents a continuum of task-particular abilities.

There are alternate explanations for the performance seen here. Perhaps, in generating a memory search, children used different markers than a feeling of knowing per se as decision criteria. For example, children may have chosen to accept trials based on a sense of 'liking' for certain items over others. An alternate explanation is a simple threshold account. By this account, when a child generates a memory search that results in an activation of a memory representation above a certain threshold the child experiences successful recall and chooses to 'accept,' otherwise s/he experiences recall failure and responds 'decline.'

However, perhaps it is these very 'alternate' explanations that comprise metacognition. A feeling of knowing can arise from a feeling of liking, or a feeling of liking may be what marks an item as being 'known', and disliking what marks it as 'not-known.' What is interesting is that these mechanisms appear to be successful, as evidenced by the accuracy on accepted versus declined items. Also, the fact that such judgments, resulting from simple processes or not, accord with but do not perfectly match actual memory suggests that they reflect at the least, differential selection processes for the qualitatively different judgment and memory tasks. If memory responses and judgment processes were one and the same, then children's judgment and memory performance would be expected to more exactly correspond. For example, if judgment response was set by a simple threshold such that pairs exceeding a certain strength were given the value 'accept' and pairs below given the value 'decline,' children would be expected to do one of two things; either they would have 100% accuracy on accepted items and a mix on decline, or 100% inaccuracy on decline and a mix of accuracy on accepted. If one makes the model more complicated by adding an area in the middle, in which any pair with intermediate strength can be accepted or declined, then that must be an area of uncertainty, in which a judgment response is selected by both the memory search and an additional, different process.

In addition, single mechanism accounts would fail to describe what differentiates animals and age groups that appear aware of uncertainty but fail to respond adaptively from those who appear both aware of and able to respond strategically to uncertainty. This difference would seem to

imply that at some level of cognitive development, the ability to recognize uncertainty becomes coupled with an ability to then respond to it with a response that is qualitatively different, and results in greater outcome utility per trial. To selectively indicate ‘not-knowing’ requires realizing both a value for the knowledge itself (select the correct target), and a value for that value (accept or decline the task.) A simple retrieve-accept fail to retrieve-decline model, or a model based on a feeling of liking would require more elaboration to describe what differentiates random responses to uncertainty from more strategic responses.

A second finding from this research was that children who demonstrated better metamemory performance also demonstrated better overall recognition memory performance than children who either over- or underestimated their memory. This intriguing finding suggests that metacognition is linked to the process on which it acts, such that performance in one domain is inherently linked to performance in the other. This linkage further suggests that individuals who differ in core processes like memory may also differ in the control and monitoring processes that act upon them. Recent focus on individual differences reveals that adults, also, vary in their ability to optimally respond to uncertain tasks, and that this difference may be correlated with other cognitive and personality variables (Washburn, Smith, & Tagliatela, 2005.) The current results suggest that metacognition and cognition are inherently interrelated, and that metacognition may emerge implicitly very early, in fact as early as when core processes emerge, providing a mechanism by which infants and children can actively direct learning.

Future Directions

Current research in our lab is underway to examine whether children's early metacognition skills rely crucially on recognition memory, or whether they are able to generate accurate metacognition assessments from retrieval processes only. A further important test of this hypothesis is to examine whether children do, indeed, strategically guide their learning based on such implicit metacognition. Future work is planned to look at the relationship of metacognition and other cognitive skills that would seem to rely on the ability to have some awareness of one's own mental state, for example decision-making abilities.

References

Aslin, R., Saffran, J., & Newport, E. (1998). Computation of Conditional Probability Statistics by 8-Month-Old Infants. *Psychological Science, 9*, 321-324.

- Butterfield, E., Nelson, T., & Peck, V. (1988). Developmental Aspects of the Feeling of Knowing. *Developmental Psychology, 24*, 654-663.
- Cultice, J., & Somerville, S. (1983). Preschoolers' Memory Monitoring: Feeling-of-Knowing Judgments. *Child Development, 54*, 1480-1486.
- Flavell, J., & Wellman, H. (1977). Metamemory. In R. Kail & J. Hagen (Eds.), *Memory in Children*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Gomez, R. (2002). Variability and Detection of Invariant Structure. *Psychological Science, 13*, 431-436.
- Kuhn, D. (2000). Metacognitive Development. *Current Directions in Psychological Science, 9*, 178-181.
- Hunter, M., Ames, E., & Koopman, R. (1983). Effects of Stimulus Complexity and Familiarization Time on Infant Preferences for Novel and Familiar Stimuli. *Developmental Psychology, 19*, 338-352.
- Lockl, K., & Schneider, W. (2002). Development Trends in Children's Feeling-of-Knowing Judgments. *International Journal of Behavioral Development, 26*, 327-333.
- Namy, L. & Waxman, S. (2000). Naming and Exclaiming: Infants' Sensitivity to Naming Contexts. *Journal of Cognition and Development, 1*, 405-428.
- Newport, E., & Aslin, R. (2004). Learning at a Distance I. Statistical Learning of Non-Adjacent Dependencies. *Cognitive Psychology, 48*, 127-162.
- Schneider, W. (1999). The Development of Metamemory in Children. In D. Gopher & A. Koriat (Eds.), *Attention and Performance XVII: Cognitive Regulation of Performance; Interaction of Theory and Application*. Cambridge, MA: MIT Press.
- Schneider, W. (2000). Performance prediction in Young Children: Effects of Skill, Metacognition and Wishful Thinking. *Developmental Science, 1*, 291-297.
- Smith, D., Shields, W., & Washburn, D. (2003). The Comparative Psychology of Uncertainty Monitoring and Metacognition. *Behavioral and Brain Sciences, 26*, 317-373.
- Smith, D., Shields, W., & Washburn, D. (1998). Memory Monitoring by Animals and Humans. *Journal of Experimental Psychology: General, 12*, 227-250.
- Washburn, D., Smith, D., & Tagliatela, L. (2005). Individual Differences in Metacognitive Responsiveness: Cognitive and Personality Correlates. *Journal of General Psychology, 132*, 446-461.
- Wellman, H. (1977). Tip of the Tongue and feeling of Knowing Experiences: A Developmental Study of Memory Monitoring. *Child Development, 48*, 13-21.