

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Modeling the Role of Memory Function in Primate Game Play

#### **Permalink**

<https://escholarship.org/uc/item/8vk799c5>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 31(31)

#### **ISSN**

1069-7977

#### **Authors**

Coen, Michael  
Colman, Ricki  
Dassow, Angela  
et al.

#### **Publication Date**

2009

Peer reviewed

# Modeling the Role of Memory Function in Primate Game Play

Michael H. Coen<sup>1,2</sup>, Vandhana Selvaprakash<sup>2</sup>, Angela M. Dassow<sup>3</sup>, Shelley Prudom<sup>5</sup>, Ricki Colman<sup>5</sup>, and Joseph Kennitz<sup>4,5</sup>

Department of Biostatistics and Medical Informatics<sup>1</sup>

Department of Computer Sciences<sup>2</sup>

Department of Zoology<sup>3</sup>

Department of Physiology<sup>4</sup>

Wisconsin National Primate Research Center<sup>5</sup>

University of Wisconsin, Madison, WI 53706

{mhcoen, vandhana}@cs.wisc.edu, {amdassow}@wisc.edu, {sprudom, rcolman, kennitz}@primate.wisc.edu

## Abstract

In this paper, we present a novel interpretation of the role of *forgetfulness* (i.e., memory impairment) in understanding game play in primates. Specifically, we examine how two primate species play a variant of the Wisconsin Card Sorting Task (WCST), a widely used clinical assessment game for measuring neurological and cognitive function. Our goal is to understand the role that memory plays in both learning and subsequently playing this game in humans and rhesus monkeys (*Macaca mulatta*). We are also interested in clustering these two populations based on their forgetfulness. This enables establishing baseline correspondences for cross-species comparison of memory function between different age groups, with the intention of enabling translational clinical treatments for a host of pathologies that involve memory dysfunction, such as senile dementia and Alzheimer's disease. Doing so requires a more in-depth understanding of the role memory plays in cognitive tests like the WCST, which we provide here through computational modeling. We also show this model surprisingly provides a clear indication that learning of an unknown game has actually occurred. It thereby disputes earlier monkey studies on a variant of the WCST by providing evidence their subjects never actually learned to play the game on which they were being evaluated. The model also demonstrates that the effect of memory impairment on game performance is highly non-linear. We find memory degradation has little gradual effect; rather, it shows a steep response past a threshold value, which has strong implications for understanding the dynamics of human aging.

**Keywords:** Game Learning; Memory Function; Cross-Species Cognitive Studies; Translation Medicine; Human Aging.

## Introduction

Few topics evoke a fear of aging as does loss of memory. However, one may ask whether this fear is realistic. While it is often the case that the extraordinarily aged suffer both cognitive and memory related impairments (Boone et al. 1990), what is the impact of a *gradual* impairment in memory function? Is it the case that all young people share near-perfect recall and that this ability deteriorates as a linear function of their age? Does any loss in memory imply a direct loss in cognitive function?

We may also wonder about other species of primates. What age related memory changes do they experience and are these similar to ours? More importantly, if we found some way to treat memory dysfunction in another species of primate, would that treatment translate to people? However, without a basis for comparison, it is difficult to gage “normal” as opposed to “abnormal” memory function across different species.

In this paper, we examine these topics, in the framework of a variant of the Wisconsin Card Sorting Task (WCST). This memory-based clinical test is widely used to assess neurological and cognitive function in people, and its variants are now being explored in the non-human primate world. We introduce a computational model of this game that provides a number of surprising results. Our work confirms the findings of (Fristoe et al. 1997) that memory impairment has a highly non-linear impact on performance. We also found that the variance of *forgetfulness* among healthy young adults is enormous. In other words, many young adults have relatively poor memories but the assays for detecting this are not informative, as these impairments have little detectable functional impact.

While formulating this model, we also wondered about the separate roles of memory in *learning* as opposed to *playing* a novel game. To our surprise, we found that we could reliably detect when a subject has learned a new game by modeling their *forgetfulness* while playing it. This led us to reexamine previous work on rhesus monkeys playing the same game and determine that they never actually learned to play the game on which they were being evaluated. We also examine how to improve the experimental paradigm to enable a far more informative set of cross-species cognitive studies, towards establishing a baseline for translational medical research.

## The Conceptual Set Shifting Task

The Wisconsin Card Sorting Test (WCST) is a standard clinical task used to assess cognitive function, particularly in cases of frontal lobe damage, memory impairment, and senile dementia (Milner 1964; Lezak et al. 2004). The original game is described in Figure 1. In this paper, we use




Figure 1: The Wisconsin Card Sorting Test (WCST). The player's task is to draw cards from a deck and match them with one of four displayed cards corresponding to a secret concept (one of *shape*, *color*, or *number*) of the examiner. So, for example, if the examiner cares about *shapes*, the player must pair correct shapes on cards drawn from the deck to those on the displayed cards. The player receives positive or negative feedback after each card is placed. After some number of correct placements in a row, the examiner secretly changes his target concept. (From Milner 1964.)

a simplified version of the WCST known as the Conceptual Set-Shifting Task (Moore et al. 2005), illustrated in Figure 2. This game removes one of the shapes (the “plus sign”) used in the WCST, and more importantly, changes the structure such that playing the game does not require verbal elucidation of the rules, which is required for evaluating non-human subjects.

In the Conceptual Set-Shifting Task (CSST) (see Figure 2), there are six possible target concepts, namely: triangle, star, circle, red, green, and blue. The opponent selects a single concept without revealing it. The player is then shown a display that contains three objects. Because each object has both a color and a shape, the game guarantees the target concept is always contained in some displayed object. (Namely, the number of concepts (6) = the number of displayed objects (3) x the number of features/object (2).) Each guess receives immediate feedback from the opponent indicating whether it contains the target concept.

The player's goal is to select an object containing the target concept some number of times, e.g., ten, in a row. After this, unbeknownst to the player, the opponent selects some other target concept, in what is known as a concept




Figure 2: The Conceptual Set-Shifting Task (CSST). In this example, the opponent's first target (or secret) concept is *triangle*. The player is shown each board in succession. His goal is to guess a shape – indicated with a dotted circle – containing the target concept. After the player has correctly guessed the target concept some number of times, the opponent changes to a new secret concept, here, *red*. This change is called a *concept drift*. Shapes are labeled with R, G, or B indicating their respective colors.

drift. As a result, the player must first realize his previous answer is no longer correct and then determine what the new concept is. These games are widely studied to determine the presence of perseverative errors (Lezak et al. 2004), namely, the inability of a subject to successfully adapt to a concept drift; these types of errors are indicative of a wide range of clinical pathologies.

However, in this paper, we are interested in using the CSST to study memory and learning in the context of an unknown game. Although some variants of the WCST provide verbal explanation to human players, there is no way to provide such instructions to non-human primates. Thus, no player (human or otherwise) of the CSST in our experiments receives any information about the game in advance. The game therefore has two distinct components: (1) learning its rules; and (2) playing it successfully after rule acquisition occurs.

## A Machine Learning Perspective

This type of game is often considered an *online learning* scenario in the machine learning community. It is typically viewed as an adversarial situation, where the opponent changes the concept as soon as he is convinced the learner has acquired it, which is known as a *concept drift*. Although it is sometimes helpful to view the CSST as an online learning game, in clinical settings test givers will often *simplify* the game if a subject is having difficulty acquiring a concept. This may involve reducing the number of required repetitions, reducing the number of trials, or limiting the number of different concepts. In our scenario, these are all fixed. We believe, however, that the number of required repetitions was selected without an appreciation that it is far larger than necessary. Requiring ten repetitions reduces that amount of collected data without providing any additional benefit. In fact, the expected number of selections for reaching a concept drift by random guessing is 88,572. (This is determined via a recurrence over the expected value:  $E(i) = 3 + 3E(i-1)$ , for  $i=10$ .) We propose future experiments lower this number significantly; five or six correct guesses in a row would certainly be sufficient to insure learning had occurred.

## A Computational Model of Game Play

Writing an algorithm to play the CSST optimally as an online learning game is straightforward and algorithms are presented in (Zhu et al. 2008; Coen et al. 2009). However, people and monkeys, whether healthy or infirmed, do not play the CSST optimally. We assume this is due to boredom or distraction in all players and organic brain disorders in players with underlying pathologies. We therefore model a notion of *forgetfulness* by modifying the algorithms presented in (ibid.). This algorithm is presented in Figure 3 and we now examine it in more detail.

Let  $h = (f_1, \dots, f_d)$  be an object with  $d$  Boolean features, representing our hypothesis about the target object. In the CSST,  $d = 6$ , representing the six possible target

1. Initialize hypothesis  $h = (1,1,1,1,1)$ .
2. Randomly select object  $x \in \{x_1, x_2, x_3\}$   
where  $x = \underset{x}{\operatorname{argmax}} |h \wedge x|$  (The maximally informative  $x$ )
3. If  $x$  is correct,  $h = h \wedge x$
4. If  $x$  is wrong,  $h = h \wedge \neg x$ .
5. If  $h = 0$ , go to step 1, as there has been a concept drift.
6. If  $\operatorname{rand}() < \rho$ , set some 0 in  $h$  to 1. (Forget something.)

Figure 3. Our algorithm for playing the CSST with forgetfulness. (1) In the beginning, the player initializes his current hypothesis  $h$  to include all possible target concepts, entertaining the notion they are all equally valid. (2) At each subsequent step, he selects an object that provides maximal information, based on what has been so far deduced. Suppose the selected object is *green triangle*. (3) If that guess is correct, he knows the target must either contain *green* or *triangle*, so his hypothesis can immediately be set to (0,1,0,1,0,0). (4) If the guess is wrong, he knows that neither of these is the target feature, so he can remove them from his hypothesis. Once the player has narrowed his hypothesis to a single feature, it is necessarily the target concept of his adversary, so he may repeatedly guess it. (5) However, following a concept drift, he will likely need to reset his space of possible guesses to once include all features. (6) Also, the player will occasionally forget that some feature has been previously ruled out in this round.

concepts of *red*, *green*, *blue*, *triangle*, *star*, and *circle*. For example, a red triangle has the feature vector  $x = (100100)$ . We define a *concept* as a feature vector with exactly one non-zero value, namely  $\sum f_i = 1$ ; for example green has a concept vector (010000). In our game, the opponent picks a target concept and the player's goal is to guess it.

At each stage of the game, the player is presented with  $d/2$  objects, each of which contains 2 features corresponding to color and shape. Therefore, for each object,  $\sum f_i = 2$ , as an object may only have one shape and one color. By the definition of the CSST, the three objects' features are mutually exclusive, implying  $\sum f_i = 6$  over all three displayed objects. In other words, the target concept is always shown to the user on the display; it is therefore always possible to guess the correct answer.

In step (6) of the algorithm, we incorporate a forgetfulness parameter  $\rho$ , which indicates the likelihood of resetting some previously ruled out feature in our candidate hypothesis  $h$  from 0 back to 1, making it once again a plausible choice. This captures our intuitive notion that players will forget concepts that have been ruled out due to a variety of reasons. We represent forgetfulness stochastically, using a function  $\operatorname{rand}()$  that uniformly draws at random from  $[0,1]$ . If the value of  $\operatorname{rand}() < \rho$ , the algorithm forgets something it has previously learned.

The CSST proceeds in rounds. The player must guess the correct concept some number of times in a row; here, this number is ten, corresponding to the length of a round. After this, the opponent selects a new target in a concept drift, and

the game repeats. The game ends after four concept drifts, corresponding to five rounds of play.

## Experimental Methodology

### Humans

The CSST was administered to 55 young adults, between 18-30 years old. These were comprised of a population of graduate students from multiple institutions and young professionals. The players interacted with the game via a web browser under supervision, intended to eliminate distractions from external influences such as e-mail and instant messages. The humans receive a minimum of verbal instruction, namely asking them to sit down and interact with the browser. Humans have no difficulty latching onto the notion that the experiment involves clicking on displayed objects and receiving feedback. Selections are immediately rewarded with visual feedback in the form of emoticons.

### Monkeys

While humans are quite familiar with interacting with icons displayed on a computer screen, this is very much a new experience for the macaques. A group of eight elderly macaques were therefore trained to use a touch screen apparatus, similar to the system described in (Voytko 2002). These monkeys were on a calorie-restricted diet, as part of an effort to gauge the effects of reduced caloric intake on their cognition. Rewards were therefore given in the form of highly desired banana pellets, and negative feedback consisted of a dissonant sound.

The monkeys went through an involved series of training steps that taught them to: (1) touch objects on the screen; (2) disregard the position of objects on the screen; and (3) learn to associate particular objects with a reward when multiple candidates were displayed. In this sense, the monkeys had to acquire an inductive bias for playing the CSST that was already assumed by people familiar with computer interaction.

We note the monkeys only played four rounds of the CSST while the human played five. This distinction has no impact on our results, but as we discovered humans play faster, we decided to gather more data from them.

### The Role of Memory in Game Performance

We note the algorithm in Figure 3 *already assumes the player knows the rules of the game*. However, this is not the case with either our human or monkey population. Neither species knows the game's rules in advance and part of the problem is simply figuring out what is going on during play. The question of modeling learning using an informed algorithm is examined in depth in (Coen et al. 2009). Here, we simply note some basic conclusions of that work.

First, to examine the effects of different values of  $\rho$ , we ran 10,000 simulations of this algorithm playing the CSST for 100 different values of  $\rho$  ranging between 0 and 0.8. The results are shown in Figure 4. From these simulations,




Figure 4: Mistake per trial as a function of  $\rho$ . We ran 10,000 simulations of our algorithm playing the CSST for 100 different values of  $\rho$  to determine how forgetfulness affects the number of errors per trial. This allows us to construct a reverse lookup table, so we can determine an individual's forgetfulness  $\rho$  by observing how mistakes he makes. This curve fits the function  $y = 1222x^3 - 310x^2 + 35.2x + 2.1$ , with  $R^2 = .999$ .

we constructed a reverse lookup table, whereby observing the number of errors per trial made by a player, we can determine the forgetfulness  $\rho$  corresponding to this number of mistakes. For example, if  $\rho = 0$ , we expect the player to make 2.1 errors on average, corresponding to optimal play. If  $\rho = .25$ , namely,  $\frac{1}{4}$  of the time we expect the player to forget something he has learned, there will be approximately 12 errors per round. Thus, we can determine a subject's forgetfulness simply by looking up his number of mistakes in our table of  $\rho$  values constructed from these simulations.

For each subject, we recorded all data corresponding to their play over all rounds of the CSST, including their selections and timing information. Using each player's number of mistakes, we computed his derived forgetfulness value  $\rho$ , using our lookup table. An aggregate view of the distributions of forgetfulness values  $\rho$  for the human players is contained in Figure 5.

We note a very interesting cognitive feature displayed in this graph; namely, relatively large increases in forgetfulness  $\rho$  (represented by circles) have little effect on observed performance (represented by triangles) below a




Figure 5: Contrasting  $\rho$  and mean errors per round. The aggregate results for the 43 human subjects who were able to learn the CSST. The  $\rho$  values were computed from the lookup table constructed above. The two horizontal lines show the average and worst case performance of the algorithm above, assuming  $\rho = 0$ , corresponding to perfect play.

threshold. In fact, only when features are forgotten 15% of the time does the number of errors surpass the worst case performance of the optimal algorithm. When  $\rho$  surpasses 20%, performance deteriorates markedly. This sudden deterioration is functionally due to  $\rho$  fitting a cubic polynomial, but it mirrors human age-related impairment quite well. It also suggests that *medium grade memory impairment has little impact on performance in these types of games*. We discuss these implications further below.

Surprisingly, we found we can also use these  $\rho$  values to determine when an individual has learned to play the game.

## The Significance of $\rho$ in Game Learning and Game Play

As we mentioned above, there are two stages to playing the CSST in our scenario. First, subjects must acquire the rules of the game simply by playing it. This has two distinct components: (1) figuring out how to get through a round, namely, determining how to consistently receive rewards; and (2) adapting to concept drifts, where the previously rewarded concept must be abandoned in favor of a new one. Subsequently, they must play the game according to these rules. We found that the transition from *learning* to *playing* is detectable, even without knowing how the subjects are learning the rules.

We note that the aggregate view in Figure 5 combines two distinct stages: (1) that of learning the game and (2) that of playing the game. Instead of viewing the data en masse, let us examine the mean  $\rho$  values for the human players over the individual rounds of the CSST, as in Figure 6. We see that  $\rho$  undergoes a sudden transition between rounds two and three, marking a noticeable increase in performance. *It is at this point we consider the subject has learned the rules of the game*. Subsequent decreases in  $\rho$  in rounds four and five reflect improved asymptotic performance.

Although we intended  $\rho$  to represent forgetfulness in the algorithm in Figure 3 – which models an expert player with imperfect memory – it is difficult to view it this way before the novice player even knows what needs to be remembered. Thus, we find that  $\rho$  has two distinct interpretations. When the player is just learning the game, assuming  $\rho$  represents forgetfulness is not meaningful, as the player has not yet learned what is important to remember. Instead, we take  $\rho$  to represent a *learning rate*.<sup>1</sup> In essence, a person with no knowledge of the game acts like a player who forgets 15% of the time; but this player is not forgetting, he is learning. However, once the player has acquired the rules of the game,  $\rho$  drops substantially. We can establish this further by superimposing the average time taken per move in each round, as shown in Figure 8. The

<sup>1</sup> We note the term “learning rate” is used by a variety of disciplines. For example, in machine learning, it is a parameter that determines the rate of algorithmic convergence. Its use here appears justified, as it captures the empirically observed amount of time and effort necessary for a subject to acquire the game's rules.




Figure 6: How  $\rho$  changes by round. Examining the mean value of  $\rho$  over our 43 human subjects, we note the sudden transition in  $\rho$  between rounds two and three. We associate this with the subject having acquired the rules of the game. From this perspective, the pre-rule version of  $\rho$  reflects an initial learning rate and the post-rule version of  $\rho$  reflects the intended measure of forgetfulness.

players spend the greatest amount of time deciding moves in round 2, after which the time per move drops precipitously along with their error rate. We believe this demonstrates that the game's rules have been acquired. At this point,  $\rho$  takes on its intended value representing memory function, as in the algorithm above. Simultaneously, players can now move very quickly, as they understand the rules governing correct move selection. This transition in  $\rho$  therefore mirrors the transition from a subject *learning* a game to the subject *playing* that game. We are thereby able to determine when a player has actually learned the CSST.

We may contrast this with data obtained from Zhu et al. (2008) describing rhesus monkeys playing the CSST. The poor performance of the monkey population described in that paper was attributed to a high number of perseverative errors due to age-related cognitive impairment. However, it was also hypothesized that perhaps the monkeys were not given a sufficient chance to actually learn the game. By examining the transitions of  $\rho$  over the rounds, shown in Figure 7, we see the complete absence of any transition in  $\rho$ , indicating that learning had not occurred yet in the monkeys. Thus, the data appear to describe the monkeys *learning* rather than *playing* the game.




Figure 7. How  $\rho$  changes by round for rhesus monkeys. The mean values for all monkeys is shown by the bold, blue line. Note that there is no sudden decrease in  $\rho$ , as there is for people in Figure 6 above.




Figure 8: Contrasting time per move with  $\rho$ . Human subjects on average spend the most time per move in round two, as they come to understand the game they are playing. We call this the *primary learning stage*. Once this occurs, the delay drops precipitously, while the accuracy simultaneously increases. We believe this demonstrates the subject has learned the unknown game. (Performance improvements beyond this point are asymptotic.)

## Variance of Memory Function in Young Adults

We have proposed a method to estimate memory impairment in a subject, using the forgetfulness parameter  $\rho$  as a proxy in our computational model. Thus, even in cases where the impact of the impairment on performance is slight, and perhaps even close to negligible, we can estimate  $\rho$  for a given player.

Previous work on the WCST found that while age effects are detectable, they are generally inconsequential for subjects less than 70 years of age (Boone et al. 1993). Similarly, a group comprised of octogenarians performed significantly worse than younger subjects (Haaland et al. 1987). Causally, age related differences in WCST performance have been attributed to impaired working memory for adults 60 to 86 years of age (Fristoe et al. 1997).

We are unaware of previous work on the characterization of memory impairments in young adults; to the best of our knowledge, this is the first work to suggest that young adults also suffer from varying levels of memory impairment. However, this variance is quite difficult to detect. This is because the assay is insufficiently sensitive in the range of relevant  $\rho$  values due to its non-linearity.

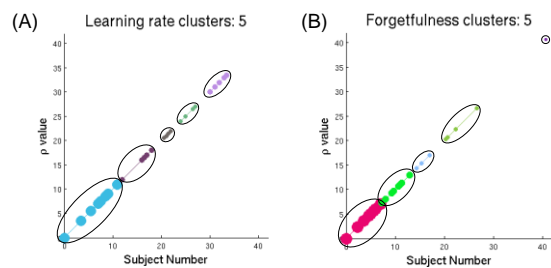


Figure 9. Viewing the variance in human learning rates and forgetfulness. (A) Clustering subjects based on their learning rate. (B) Clustering subjects based on  $\rho$  measuring forgetfulness after learning the game. Identically colored points represent each cluster, where the point size is indicative of the cluster size.



However, our computational model allows us to extract individual  $\rho$  values, which we then clustered using the Affinity Propagation algorithm (Frey and Dueck 1997) as shown in Figure 9.

The benefit of clustering is that it allows us to characterize variations in  $\rho$  and determine its variance across our sample population. Thus, we see that in both learning and playing the game, we can clearly distinguish different degrees of memory impairment in a sample of young adults. This is possible even though their performance on the CSST is quite similar and near optimal in many cases.

One can therefore say that our model provides a more sensitive measure of memory dysfunction than the clinical assessment it is modeling. This is demonstrated by the clustering in Figure (9b), where  $\rho$  corresponds to forgetfulness, as discussed above. In the clustering on the left in Figure (9a),  $\rho$  corresponds to learning rate, which is not measured by the CSST or the WCST. However, we can also see differential ability here, but whether it is attributable to memory or some other cognitive function(s) is unknown.

We note in the memory-based clustering (9b) that the first and largest cluster (red and leftmost) contains subjects who play near optimally – they have an average  $\rho$  value of 5%, and made 2.67 mistakes per trial on average. The fifth cluster on the far right consists of a lone subject who made 94.2 mistakes per trial on average.

Of the 55 human subjects, six were unable to reach the first concept drift, which we attributed to disinterest and removed them from further consideration. Of the remaining 49 subjects, only 43 were able to finish the CSST. To our great surprise, six appeared unable to learn this game in the number of rounds allotted. Several of them made hundreds of moves per concept, yet were unable to latch onto the correct set of rules leading to a concept drift.

## Conclusions

The goal of this work has been to establish cognitive baselines between human and rhesus monkey performance in the Conceptual Set Shifting Task. Our hope is by understanding how performance in one species relates to that of the other, we can enable translational medical approaches to the wide range of human disease pathologies affecting memory. Doing so required a far more in depth understanding of the role of memory in the CSST and other cognitive assays. We therefore developed a computational model to describe the role of memory in playing these games at a fine level of detail.

This work has lead to unexpected conclusions about how humans learn and play unknown games, the high degree of memory variance in young adult populations, and the need for simpler experimental frameworks to derive meaningful data from non-human primates. Regarding the latter point, using a smaller number of repetitions per trial would likely enable the monkeys to transition from learning to playing the CSST.

Our results agree with predictions about gradual degradation in human performance as a result of memory impairment, accompanied by a sudden transition reflecting a significant impact on problem solving. However, our finding that a high degree of variation in memory function is common in younger populations presents a new framework for examining memory and understanding how it changes over time.

## Acknowledgements

This work was supported by the School of Medicine and Public Health, the Wisconsin Alumni Research Foundation, the Department of Biostatistics and Medical Informatics, and the Department of Computer Sciences at the University of Wisconsin-Madison. This research was conducted in part at a facility constructed with support from Research Facilities Improvement Program grant numbers RR15459-01 and RR020141-01 from NCRR. The authors gratefully acknowledge the technical assistance provided by S. Baum, J. Christensen, J. A. Adriansjach, C. E. Armstrong, and the Animal Care, Veterinary and Pathology Staff of the Wisconsin National Primate Research Center. Thanks to S. Asthana, S. Johnson, and M.H. Ansari for helpful comments. Finally, we thank the anonymous reviewers for their thoughtful and insightful suggestions.

## References

- Boone, K. B., Ghaffarian, S., Lesser, I. M., Hill-Gutierrez, E., & Berman, N. G. (1993). Wisconsin Card Sorting Test performance in healthy, older adults: Relationship to age, sex, education, and IQ. *Journal of Clinical Psychology*, 49: pp54-60.
- Boone, K.B., Miller, B.L., Lesser, I.M., and E. Hill. (1990). Performance on frontal lobe tests in healthy, older individuals. *Developmental Neuropsychology*, pp. 215–223.
- Coen, M., Selvakrishnan, V., Prudom, S., Colman, R., and Kemnitz, J. (2009). Learning Unknown Games in Humans and Other Primates. In submission.
- Frey, B. J., and Dueck, D. (2007). Clustering by Passing Messages Between Data Points. *Science*, 315(5814), 972–976.
- Fristoe, N.M., Salthouse, T.A., and Woodard, J.L. (1997). Examination of age-related deficits on the Wisconsin Card Sorting Test. *Neuropsychology*, 11: 428.
- Haaland, K.Y., Vranes, L.F., Goodwin, J.S., Garry, P.J. (1987) Wisconsin card sort test performance in a healthy elderly population. *J Gerontol* 39: pp166–169.
- Lezak, M., Howieson, D.B., Loring, D.W. (2004). *Neurological Assessment*, Oxford University Press, Oxford.
- Milner, B. (1964). Some effects of frontal lobectomy in man. Warren and Akert (eds.). *The Frontal Granular Cortex and Behavior*. New York: McGraw-Hill.
- Moore, T.L., Schettler, S.P., Killiany, R.J., and Herndon, J.G. (2005). Cognitive impairment in aged rhesus monkeys associated with monoamine receptors in the prefrontal cortex. *Behavioural Brain Research*. Vol 160:2, pp208-221.
- Voytko, M.L., (2002). Estrogen and the cholinergic system modulate visuospatial attention in monkeys (Macaca fascicularis). *Behavioural Neuroscience*. 116, pp187–197.
- Zhu, X., Coen, M., Prudom, S., Colman, R., and Kemnitz, J. (2008). Online learning in monkeys. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI-08)*.