Gaze-following and awareness of visual perspective in chimpanzees

Anthony M Harrison (anthony.harrison.ctr@nrl.navy.mil)
J. Gregory Trafton (greg.trafton@nrl.navy.mil)
Naval Research Laboratory, Washington DC

Abstract

Recent research suggests that chimpanzees are capable of level 1 perspective taking (Flavell, 1992), but that its expression is limited to situations of increased competition (Brauer, Call, & Tomasello, 2007). We present a model utilizing gaze-following that learns in response to the behavior of a competitor. The model not only learns the proper application of the perspective taking strategy but also the critical spatial characteristics that influence the competitive pressure.

Introduction

Under normal conditions most children will eventually develop a full theory of mind and have full visual perspective taking (Corkum & Moore, 1995,1998; Moll & Tomasello, 2006). Most researchers believe that chimpanzees have neither a full theory of mind nor full visual perspective taking (Povinelli et al., 1994; Tomasello & Call, 1997). Whether chimpanzees have *any* perspective taking ability at all has been subject to some recent debate.

Experimental studies using a variety of paradigms have previously been unable to find strong evidence for perspective taking. In fact, two of the major experimental labs consistently agreed that chimpanzees had no visual perspective taking ability (Povinelli et al., 1994; Tomasello & Call, 1997). However, a novel paradigm suggested that chimpanzees did, in fact, know what others could and could not see (Hare et al., 2000; 2001). In this paradigm a subordinate and dominant chimpanzee competed with each other for two pieces of food, one of which was hidden to the dominant (figure 1, left). Since the subordinate preferred the hidden food, Hare et al. concluded that it was aware of the dominant's visual perspective (2000, 2001).

Unfortunately, in a series of experiments, Karin-D'Arcy and Povinelli (2002) were unable to replicate the original Hare et al. (2000) findings. Karin-D'Arcy and Povinelli used a more stringent coding methodology and suggested that chimpanzees do not understand what others can and cannot see but instead use a variety of competitive strategies to succeed in such scenarios, such as preferring food near barriers.

One difference, however, between the two sets of experiments was the size of the testing area. In the original Hare et al. (2000) experiment, the testing area was 3m x 3m, but Karin-D'Arcy and Povinelli (2002) used a smaller testing area that was 2.6m x 1.8m. It is possible that this size difference could have driven the dynamics and the competitiveness of the situation for the chimpanzees. For example, in a smaller area, it is possible that, since the submissive was released before the dominant, the submissive was able to quickly grab the food, making the use of visual perspective taking less relevant. In the larger area, the competitive aspects of the area could make a quick grab of the food less effective since it would take the submissive longer to approach the food.
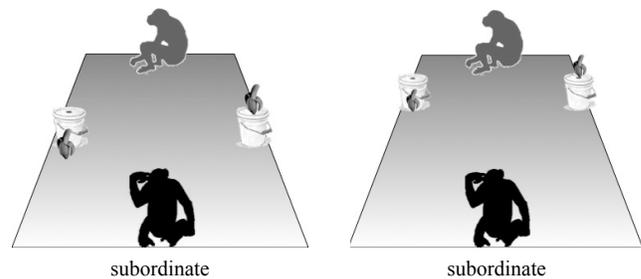

Figure 1. Dual-food layout for Brauer, et al (2007). Visible and hidden food nearer subordinate (left), and further away (right).

Brauer, Call, and Tomasello (2007) tested this idea by making several changes to their experimental paradigm, using the stronger methodology that Karin-D'Arcy and Povinelli (2002) suggested and manipulating the spatial characteristics and therefore the competitive nature of the situation. Specifically, Brauer et al. (2007) manipulated the location of the food to be nearer or farther away from the submissive (figure 1). They found that in the less competitive situation where the food was closer to the submissive, chimps did not seem to use visual perspective-taking. However, in the more competitive situation where the food was further away, chimps did seem to use visual perspective taking, preferring to pursue the hidden food (figure 2).

While the empirical data suggests that chimpanzees do have some form of visual perspective taking, it is unclear what degree of visual perspective taking is needed. Other researchers have suggested different levels of visual perspective-taking, mostly focused around the development of human children (Flavell, 1992). This work suggests that human infants, by one year of age, can follow another's gaze to targets (Corkum & Moore, 1995; 1998). By 12-15 months, a child knows a great deal about what others can and can not see, including (a) that an adult's line of sight is blocked by a screen unless it is transparent or has a window in it (Caron et al. 2002; Dunphy-Lelii & Wellman, 2004); (b) that an adult will not be able to see a target while their eyes are closed (Brooks & Meltzof, 2002); and (c) that an

adult can see something that the child can not when the adult looks to locations behind them or behind barriers (Moll & Tomasello, 2004).

Most researchers interpret these findings as evidence of level 1 visual perspective-taking (Flavell, 1992): understanding the content of what a child sees may differ from what another may see. Level 2 visual perspective taking is achieved when a child understands that people can see the same view from different perspectives. After level 1 and 2 visual perspective taking, normally developing human children also achieve a full theory of mind (knowing that others can have different thoughts and beliefs).

Hare, Call, & Tomasello (2001) suggested that chimpanzees are able to engage in level 1 visual perspective taking but not level 2. We modeled level 1 visual perspective taking to determine if it is sufficient to match the data from Brauer et al. (2007). We embed our simulation within a learning framework as well to explore how different competitive strategies can be learned.

Specifically, a model of chimpanzee competitive food foraging was developed within ACT-R (Anderson, Bothell, Byrne, Douglass, Lebiere, & Qin, 2004) utilizing the architecture's procedural learning mechanisms and a new gaze-following capability to support level 1 perspective taking.

Experiment

The refined methodology of Brauer et al. (2007) used a testing environment that was 2.5m x 2.6m, with barriers placed at the extreme sides of the cage. In the *near* condition, the barriers were equidistant between the two entrances. For the *far* condition they were moved 0.5m closer to the dominant's entrance. Food pieces were either placed behind the barrier (visible to the subordinate only) or on top (visible to both). On each trial, there could be two pieces of food (one hidden and one visible), one visible or one hidden.

The trial began when the subordinate's door was opened allowing it into the environment. After the subordinate entered the cage, the dominant's door was opened (usually within 2s). The subordinate's food preference was recorded when it made a reaching gesture in the direction of a piece of food before the dominant had approached any barrier.

The single food trials were control conditions testing the possibility that the subordinate might simply prefer food located near barriers (Karin-D'Arcy & Povinelli, 2002). The critical comparison is between the two distance conditions. When the pieces of food were *near* the subordinate, it chose indiscriminately. Because of its head start (~2s), the subordinate could pursue either piece, and was often able to acquire both. However, when the food was closer to the dominant, the subordinate preferred the hidden food almost 2:1 (figure 2).
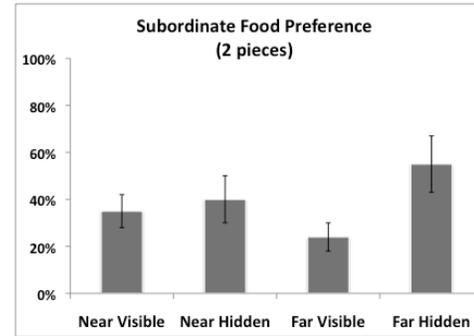

Figure 2. Subordinates prefer hidden food when competitive pressures are greatest (right). Error bars are SE (Brauer, et al, 2007).

Model

Models of both the dominant and subordinate chimpanzees were built in ACT-R (Anderson, et al., 2004). These models were run within the Player/Stage environment (Collett, MacDonald, & Gerkey, 2005) that mimicked the structure of the actual experiment.

As an integrated architecture, ACT-R provides multiple mechanisms for representation and learning. These particular models rely upon ACT-R's procedural memory and learning. At any given time there is a set of productions (if-then rules) that may fire because their conditions match the current external state of the environment or internal state of the model. From this set of competing productions, a single one is selected and fired, ultimately modifying the environment or internal state. ACT-R uses the predefined or learned utilities of productions to determine which will be fired.

To learn production utilities, ACT-R uses an elaboration of the temporal-difference (TD) algorithm (Sutton & Barto, 1998). The elaboration in ACT-R is more applicable for human learning and allows it to be more easily incorporated into a production-system framework (Fu & Anderson, 2006). Briefly, any time reinforcement is given (e.g., a banana eaten or physical punishment) the reinforcement value is propagated back in time through the rules that had an impact on the model receiving that reinforcement. Reinforcements (either positive or negative) gradually shift utility values and therefore the relative probability that a particular production will be selected over others within a set of competitors.

The application of ACT-R to non-human cognition presents many challenges. Even though chimpanzee cognition shares many similarities to that of humans, the architecture may still provide too much capability. Because of this we intentionally used the least-common-denominator in these models. The chimpanzee models make no use of declarative encoding or retrievals, nor does it engage in any imaginal operations. The models are driven predominantly by reactive productions and rely upon an impoverished goal representation (merely storing what target to pursue).

Gaze-following

To implement gaze-following in ACT-R, a new set of optional constraints were introduced to the visual search mechanism. ACT-R's basic visual search mechanism takes a request to find a percept matching some set of features (e.g. where is a red object?). The possibly features include both visual properties (i.e. color, size) and limited spatial information (e.g. nearest the current focus of attention). The location of the first matching object is returned to the model allowing it to attend to that location and encode the actual visual representation of that percept.

Within this mechanism, gaze-following was implemented as a directed visual search along a retinotopic vector. Specifically, instead of returning the first matching location in search, the full set of matches is passed through a secondary filter. This filter merely sorts the locations by their distance from the retinotopic vector. Given a starting point and either an angle or an end point, the visual search returns the location on an object somewhere along that line within a specified tolerance. Knowing the visual location of the dominant chimp (A in figure 3) and the food (C in figure 3), the subordinate performs a visual search for any object along the line segment AC. Finding the barrier (B), the subordinate can (generally) assume that the food is not directly visible to the dominant.
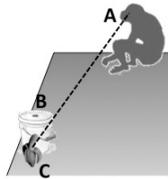

Figure 3. Retinotopic searches to find objects 1) between A and C or 2) along the ray starting at A.

This simple mechanism allows the visual system to find objects along a gaze line, or any potential obstructions between two points. While this mechanism is not accurate for all gaze-directions (particularly as the ray approaches the viewer), they are adequate for basic searches. More advanced gaze-following is addressable by having the model perform more detailed processing of the returned visual locations and the actual visual percepts at those locations, such as testing the distance, size, or opacity of an obstruction. Given the nature of the experimental environment, these higher-level strategies were not implemented.

Model Structure

The dominant and subordinate models are composed of the same constituent parts. Each model performs a full environment scan from its current position, looking not only for the food, but also the other chimpanzee and the buckets. The targets are evaluated to determine which should be pursued.

Environmental Scan The environmental scan is a rapid visual search of the environment that attends to all visible objects. If the object is a piece of food, a bucket, or another chimpanzee, the first occurrence is retained in the model's limited goal representation. If no objects are found, the model physically rotates its body to get a different view of the environment.

Target & Strategy Evaluation Once a target has been attended to it must be evaluated. For the dominant model this is simple: if it's food, pursue it, otherwise keep looking. The subordinate has more to consider. First, the subordinate must determine whether the food is *near* or *far*. Once classified, the subordinate can then choose which strategy to use. It can either try to make a mad-dash for the food (grab-and-go), or use gaze-following to ensure that the coast is clear. If the subordinate chooses grab-and-go, it runs the risk of contention with the dominant, particularly if the food is far away. For gaze-following, the subordinate will use the location of the dominant's head and the target to find any intermediate object that may be a visual barrier. If a visual barrier is found, the subordinate assumes the dominant cannot see the target and will pursue it. If no barrier is found, the subordinate rescans the environment ignoring the rejected target.
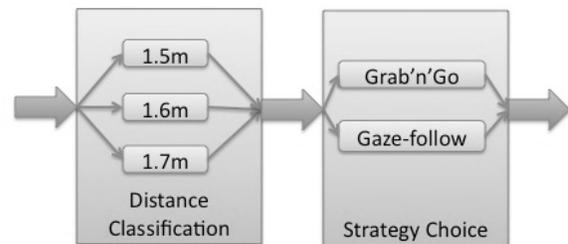

Figure 4. Two choice points for the subordinate model. The model must learn which distance threshold to use for classification and then which strategy to use.

Target Pursuit Since the Brauer, et al. experiment recorded food preference based on the initial reaching behavior, models' food preferences were recorded immediately after evaluation. The full models, however, are able to navigate in the environment, grab food and even strike each other.

Model Assumptions and Parameter Selection

At their heart all models are simplified abstractions of their respective phenomenon. Simplifications can be for reasons of computational tractability, interpretability, or theoretical relevance. The models described here must operate at a high-level of fidelity in order to capture the embodied nature of the task. The computational costs of the embodied simulations required a handful of simplifying assumptions.

Environmental Assumptions

In the actual experiment, doors into the experiment cage were opened allowing the chimpanzees to enter the space. After the subordinate entered, the dominant's door was opened, typically after around 2 seconds. Lacking doors in the simulation, each model was "beamed" into the experiment space. The delay between the subordinate and

the dominant was fixed at 2 seconds. Since the subordinate's food preference is only recorded if it is made before the dominant makes one, this delay acts as a scalar for the food preference measure. Increasing the delay allows the subordinate more time to choose, increasing the absolute food preference scores.

Model Assumptions

Learning Brauer, et al (2007), Hare et al. (2000; 2001) and Karin-D'Arcy & Povinelli (2002) all noted a lack of learning within their studies. All concluded that the preferences and skills exhibited had developed prior to testing. For the models to exhibit these behaviors they either have to be hand tuned by the modeler or they must be given sufficient training prior to testing. Having an architecture that can learn allows us to avoid the problem of custom tuned models. Each model was run through a series of learning trials, which consisted of ten sets of the full factorial design of the experiment (e.g. single & dual pieces of food at both the near & far distances), for a total of 60 trials. This was a rough surrogate for the individual's life experience with competitive food foraging.

Additionally, since gaze-following is learned over time in humans (Corkum & Moore, 1995), initial utilities of the gaze-following productions were lowered below those of the grab-and-go productions (to -1.5). This provides an early bias towards grab-and-go, delaying the onset of gaze-following, potentially providing the model with the time necessary to learn the distance classifications.

Reinforcement Probabilities In order to learn from these trials, the models must receive some reinforcement based upon their target choices. However, since the trials terminate after target choices are made they normally wouldn't receive any reinforcement. One alternative would be to run each trial to completion (after either has actually consumed the food or been hit). Unfortunately, full trials, with the possibility of the dominant chasing the subordinate around the cage, are extremely costly computationally (by almost an order of magnitude).

Reinforcements were provided based on the model target choices. When either chooses an uncontested piece of food, it is rewarded. When both the dominant and subordinate decide to pursue the same target there is some chance that the dominant will charge and strike the subordinate. Naturally, as the distance between the target and dominant decreases, the probability that the subordinate will be punished for pursuing that same target increases. All other things being equal, when the distance to the target is equivalent, there is roughly a 50% chance that the subordinate will be able to reach the target first. The chance of being hit is further reduced by the subordinate's two-second head start in the experiment design. The qualitative behavioral pattern (i.e. subordinate preferring hidden food when both pieces are closer to the dominant) holds through probability values where $P(hit|near) < 0.5 \leq P(hit|far) \leq 1$.

Generally speaking, the higher the probability of being hit for any given distance, the more likely the subordinate will select the more conservative gaze-following strategy. The values $P(hit|near)=0.1$ and $P(hit|far)=0.9$ were settled upon after a high-level exploration of the parameter space. Simulations testing the validity of these assumptions using the full trial protocol are ongoing.

Hit Probabilities Reinforcement Values ACT-R's reinforcement learning mechanism relies ultimately on time as its metric (Fu & Anderson, 2006). This forces the modeler to map physical rewards and punishments into a temporal reference frame. For this experiment, the reward for getting a piece of food was set at the average maximum time to complete the task using the gaze-following strategy (4 seconds). The punishment for being hit needs to be greater in magnitude than the food reward in order to pull apart the two primary strategies. Parameter explorations yielded good convergence rates for punishments around 8 seconds.

ACT-R's default utility learning rate of 0.2 was used. The only other parameter modified was the utility noise (0.1), which permits weaker productions to occasionally be selected over their stronger competitors.

Simulation Results

For this model to be a viable account for the subordinate chimpanzee's behavior not only must it fit the aggregate food preference measure, but it must also be able to correctly classify the target distances and prefer the gaze-following strategy for far targets. Because the individual learning histories result in greater downstream behavioral variability, large numbers of models had to be run to arrive at stable results. The results presented here are the derived from 1000 individual model runs.

Distance classification

The key factor in the results presented by Brauer, et al (2007) is that the preference for choosing the hidden piece of food is dependent upon how close the food is to the dominant chimpanzee. While they did not do a full parametric exploration of the factor, the simple difference of half a meter was sufficient to tease apart the behaviors.

Similarly the model had to be able to correctly classify the target distances as near or far. At the distance choice-point (figure 4), three productions are in competition, setting the distance threshold to 1.5, 1.6, or 1.7m. Subsequent productions then classify the target's distance using that threshold. In the simulation, target distances $\geq$ 1.6m correspond to the *far* condition. Within each model we can simply examine the relative utilities of the distance threshold productions; 41% of the models converged upon the correct threshold of 1.6m, 21% at 1.5m and 14% at 1.7m. The remaining 24% of the models showed no clear preference as the threshold utilities were all within the model's utility noise.

Strategy Selection

When the food is *near*, it is perfectly rational for the subordinate to make a mad-dash for either piece. With the two-second head start, there is little chance that it will be punished. On the other hand, when the food is further away

(and closer to the dominant), it makes sense to use the gaze-following strategy even though it takes longer and requires waiting for the dominant to enter the experiment space. If the subordinate were to use grab-and-go for far targets, it would run an increased risk of contention with the dominant, even with its head start. On average gaze-following took 0.75 – 1.5 seconds longer than grab-and-go. While this increase in execution time ultimately reduces the temporally discounted reward, it effectively avoids the much more costly punishment when conflict does occur. Figure 5 shows the percentage of model strategic preferences. The majority of the models preferred grab-and-go when *near* and gaze-following when *far*.
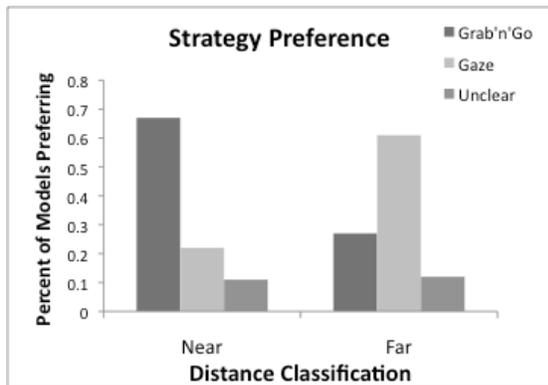

Figure 5. Percentage of models preferring a given strategy for both *near* and *far* target classification.

Model Fit

Even with the model complexity and resulting downstream behavior variability, the fits were strong (RMSE=7.2%, R^2=0.96). The qualitative pattern (i.e. preference for hidden food when far and equivalence for near) holds across the majority of the hit probability ranges discussed earlier.
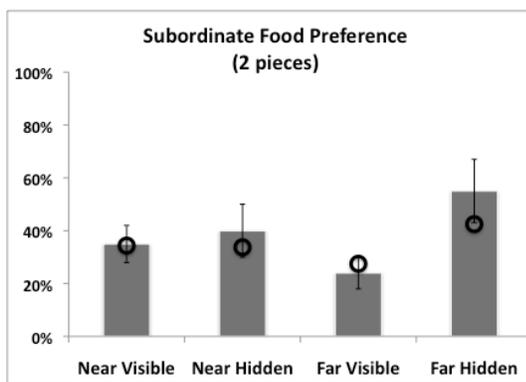

Figure 6. Model (circles) fit to Brauer, et al (2007) data. RMSE=7.2%, R^2=0.96

Distance & Strategy Interactions

The variability in the behavior of any given subordinate model is a direct result of its experiences with the dominant model. That some learned the wrong distance threshold or frequently choose the wrong strategy is hardly surprising. Looking more closely at these models is particularly informative from a rational analysis perspective. All of the models that settled on the 1.5m distance threshold used the gaze-following strategy exclusively for *far* targets (which would have been virtually all of the them). Similarly, over half the models that settled on 1.7m as the distance threshold preferred gaze-following when targets were both *far* and *near*. These overly conservative models were able to stabilize in their patterns because there was no disincentive for misclassifying targets as *far* only *near*, particularly since they could rely upon gaze-following to compensate for incorrect distance classifications.

Discussion

The simulation presented provides a process model of chimpanzee competitive food foraging that combines the awareness that individual visual experiences are different (i.e. Flavell, level 1) and a simple gaze-following mechanism. Leveraging the existing reinforcement-learning component in ACT-R, the model learns to prefer the more conservative gaze-following strategy when the risk of punishment is increased (i.e. when the food is closer to the dominant). The model shows that its "awareness of the other's visual experience" need not entail full visual perspective taking (Hare, Call, & Tomasello, Animal Behaviour, 2001). Knowledge of the particular spatial relationships that the dominant is experiencing are also unnecessary.

Obviously this does not preclude the possibility that chimpanzees possess level 2 skills. It is worth considering how a model of full perspective taking would perform in this situation. Such a model was actually developed before the one reported here. It performed egocentric transformations of its own perspective, aligning them with the perceived position and orientation of the dominant (e.g. Hegarty & Waller, 2004). This model was able to learn the same qualitative behavioral pattern, but at an increased cost. Perspective transformations are particularly costly in terms of time; often taking 2-4x longer than gaze-following depending on assumptions of representational capacity and mental transformation rates.

What is perhaps more interesting is that if full perspective taking and gaze following are allowed to compete, gaze following is consistently preferred. While gaze following isn't as accurate at assessing visibility, it is accurate enough within the confines of the task and significantly faster. Given this, it is unlikely that one could find evidence of full perspective taking in the current experimental paradigm.

These models arose out of our growing interest in embodied cognition. While fully situating a model in an environment makes some tasks quite simple (i.e. inferring intent based on gaze direction), it comes at the cost of requiring higher fidelity models and simulations. This higher fidelity brings with it increasingly complex dynamic interactions between the model and environment (including other intelligent agents). Our work with human-robot interaction has shown us that these dynamic interactions cannot be ignored.

Conclusions

A computational learning model was developed that is able to effectively reason about what another can and cannot see. This embodied model is able to learn and exploit regularities in the environment (target distances) to adapt to a competitor's behavior. The model is able to do this with only a basic gaze-following mechanism instead of relying upon full visual perspective taking (Hare, Call, & Tomasello, 2001). This mechanism, implemented as a general directed visual search, provides an important developmental step towards the development of theory-of-mind (Baron-Cohen, 1995; Butterworth, 1991).

Acknowledgments

References

Anderson, J., Bothell, D., Byrne, M., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111* (4), 1036-1060.

Baron-Cohen, S. (1995). The eye direction detector (EDD) and the shared attention mechanism (SAM): Two cases for. *Joint attention: Its origins and role in development*.

Bräuer, J., Call, J., & Tomasello, M. (2007). Chimpanzees really know what others can see in a competitive situation. *Animal Cognition, 10* (4), 439-448.

Butterworth, G. (1991). The ontogeny and phylogeny of joint visual attention.

Call, J., Hare, B., & Tomasello, M. (1998). Chimpanzee gaze following in an object-choice task. *Animal Cognition, 1* (2), 89-99.

Caron, A.J., Kiel, E.J., Dayton, M., & Butler, S.C. (2002). Comprehension of the referential intent of looking and pointing between 12 and 15 months. *Journal of Cognition and Development*, 3(4), 445–464.

Collett, T., MacDonald, B., & Gerkey, B. (2005). Player 2.0: Toward a practical robot programming framework. *Proceedings of the Australasian Conference on Robotics and Automation (ACRA 2005)* .

Corkum, V., & Moore, C. (1995). Development of joint visual attention in infants. In C. Moore & P. J. Dunham(Eds.), *Joint attention: Its origins and role in development* (pp. 61–83). Hillsdale, NJ: Erlbaum.

Corkum, V., &Moore, C. (1998). The origins of joint visual attention in infants. *Developmental Psychology*, 34(1), 28–38.

Dunphy-Lelii, S., & Wellman, H. M. (2004). Infants' understanding of occlusion of others' line-of-sight: Implications for an emerging theory of mind. *European Journal of Developmental Psychology*, 1(1), 49–66.

Flavell, J. H. (1992). Perspectives on perspective-taking. In H. Beilin & P. B. Pufall (Eds.), *Piaget's theory: Prospects and possibilities. The Jean Piaget symposium series* (Vol.14, pp.107–139). Hillsdale, NJ: Erlbaum.

Fu, W.-T., & Anderson, J. (2006). From Recurrent Choice to Skill Learning: A Reinforcement-Learning Model. *Journal of experimental psychology: general , 135* (2), 184-206.

Hare, B., Call, J., & Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal Behaviour , 61* (1), 139-151.

Hare, B., Call, J., & Tomasello, M. (2006). Chimpanzees deceive a human competitor by hiding. *Cognition , 101* (3), 495-514.

Hare, B., Call, J., Agnetta, B., & Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour , 59* (4), 771-785.

Hegarty and Waller. A dissociation between mental rotation and perspective-taking spatial abilities. *Intelligence* (2004) vol. 32 pp. 175-191

Karin-D'Arcy, M., & Povinelli, D. (2002). Do chimpanzees know what each other see? A closer look. *International Journal of Comparative Psychology , 15*, 21–54.

Moll, H., & Tomasello, M. (2004). 12- and 18-month-old infants follow gaze to spaces behind barriers. *Developmental Science*, 7(1), F1–F9.

Moll, H., & Tomasello, M. (2006). Level 1 perspective-taking at 24 months of age. *British Journal of Developmental Psychology*, 24, 603-613.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: MIT Press.

Tomasello, M. (2003). Chimpanzees understand psychological states – the question is which ones and to what extent. *Trends in cognitive sciences , 7* (4), 153-156.

Tomasello, M., Call, J., & Hare, B. (1998). Five primate species follow the visual gaze of conspecifics. *Animal Behaviour , 55* (4), 1063-1069.