A Cognitive Model of Theory of Mind

Laura M. Hiatt (laura.hiatt.ctr@nrl.navy.mil)
J. Gregory Trafton (greg.trafton@nrl.navy.mil)
Naval Research Laboratory
Washington, DC 20375 USA

Abstract

It is generally well acknowledged that humans are capable of having a theory of mind (ToM) of others. We present here a model which borrows mechanisms from three dissenting explanations of how ToM develops and functions, and show that our model behaves in accordance with various ToM experiments (Wellman, Cross, & Watson, 2001; Leslie, German, & Polizzi, 2005).

Keywords: cognitive architectures; theory of mind

Introduction

The concept of "theory of mind" (ToM) refers to one's ability to infer and understand the beliefs, desires and intentions of others, given the knowledge that one has available; without it, people can be severely impaired in their ability to interact with others (Baron-Cohen, Leslie, & Frith, 1985). A large body of research has tried to explain how this critical ability functions by studying its development in children (Wellman et al., 2001), but has led to many contradictory accounts.

We have built a model that borrows ideas from various explanations of how ToM develops and functions to form a cohesive theory of ToM, and show that it produces behavior in accordance with various ToM experiments (Wellman et al., 2001; Leslie et al., 2005). While the similarities between a model's behavior and data is not a certain indicator of cognitive plausibility (Cassimatis, Bello, & Langley, 2008), it can distinguish between models that show performance and data fit (which, to us, are preferred) and models that do not.

Theories of the Theory of Mind

There are, in general, three competing views for how ToM takes place at a cognitive level. They are typically described in the context of "belief and desire" reasoning: ToM is the understanding that different people can have different beliefs, not all of which may be actually true; people also have internal desires that cause them to act in certain ways, physically, in the world. There is also a distinction between "true-beliefs," or beliefs that are true in the physical world, and "false-beliefs," which others may have but which are not actually true. The ability to understand a false-belief task, then, indicates evidence that a person can appreciate the distinction between the mind and the world (Wellman et al., 2001).

Conceptual change (commonly called theory-theory) is one possible explanation for ToM (Wellman et al., 2001). Theory-theorists believe that children learn a set of causal laws, or theories, about the beliefs and desires of people in general (Gopnik, 1993). Children then use these causal laws to explain behavior observed in others, to predict desires and behaviors, and to perform other related ToM tasks.

Simulation theory is a second view (Gallese & Goldman, 1998). It posits that when a person ("A") tries to understand another ("B"), A simulates what he/she would do in B's place, and attributes the result to B. More specifically, the theory states that humans perform ToM by representing the mental states of others, and then using their own decision-making systems to operate on these foreign mental states to predict others' behavior; similar processes can be used to explain observed behavior, making backward inferences. Gallese and Goldman (1998) describe the distinction between this and theory-theory as, while theory-theory is performed as a "'detached' theoretical activity," simulation theory involves attempting to mimic or impersonate the mental state of another.

A third body of literature posits that the mind has two separate mechanisms that work together to provide ToM (Leslie, Friedman, & German, 2004). The theory of mind mechanism (ToMM) allows people to generate and represent multiple possible beliefs. It is argued that this mechanism is fully functional in even very young children. The second mechanism provides a selection process (SP) that uses inhibition to reason about others' beliefs, such as inhibiting a true-belief to select a false-belief answer; this processing ability, it is argued, develops in children during the pre-school years. To describe how the mechanisms work together as "ToMM-SP" to provide ToM, the authors break it down into four steps: (1) identify candidate belief possibilities; (2) provide *a priori* weights to the candidates, with true-belief receiving the highest weight; (3) adjust the weights given the belief inquiry; and (4) select the highest-weighted candidate as the answer.

A variety of experiments, primarily in developing children, have led to a range of results that supports each of these theories. We describe next some of these experiments, followed by our interpretation of the data and our overall view of ToM.

Experiments in Developing Children

The majority of experiments in this area concerns false-belief tasks. Arguably, the most well-known false-belief task (and the one on which we focus in this paper) is the Sally-Anne task (Baron-Cohen et al., 1985), in which a child is shown a play with two characters, Sally and Anne (Figure 1). The true-belief answer (to where Sally believes the marble is) is that the marble is in Anne's box (the "TB box"), since that is where the marble actually is. In contrast, the correct answer is the false-belief answer, Sally's box (the "FB box").

Variations on the Sally-Anne task have also been explored. One is the *avoidance* false-belief task (which we shorten to "avoidance task"). In a sample set-up, the marble is replaced by a kitten that crawls between boxes while Sally is out of

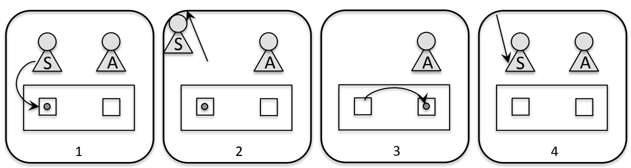

Figure 1: A diagram of the Sally-Anne task. A child watches while: (1) Sally puts a marble in her box; (2) Sally leaves the room; (3) Anne moves the marble to Anne's box; (4) Sally returns to the room. The child is then asked where Sally believes the marble is.

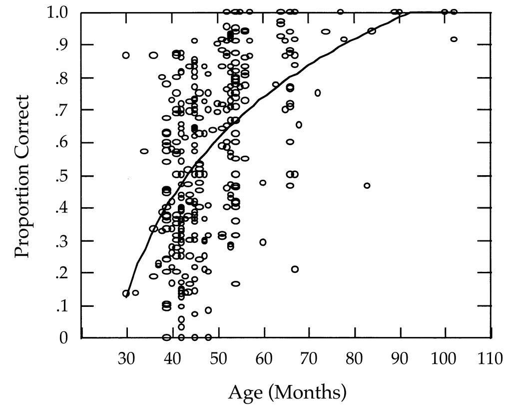

Figure 2: Results from (Wellman et al., 2001) showing a scatterplot of the results and best-fit curve.

the room; when Sally returns, she wants to put a piece of fish under the unoccupied box so that the kitten will not eat the food and get sick. Therefore, the correct answer to the question "where will Sally try to put the fish" is the TB box. This task involves not only identifying Sally's false belief, but also taking into account her avoidance desire to predict her behavior, presumably making the task more difficult.

To individually consider all the experiments in this area is nearly impossible. Instead, we focus on a meta-analysis that compiled a broad range of false-belief experiments (Wellman et al., 2001), and a more detailed experiment performed after the meta-analysis was compiled (Leslie et al., 2005). These two studies involve two developmental shifts that are believed to occur in children. The first is at about 3-4.5 years of age, when children go from being mostly incorrect to mostly correct on the standard false-belief task; this seems to correlate with the ability to recognize and identify beliefs of others. The second developmental shift is at around 4.5-6 years, when children go from having difficulty with the avoidance task to performing it mostly correctly; this seems to correlate with a child's ability to account for both beliefs and desires, and to use them to predict the behavior of others.

The meta-analysis performed by Wellman et al. (2001) provides three results pertinent to this paper. First, it identified several task components that were statistically insignificant, including the exact type of task being performed as well as the phrasing of the false-belief question (*e.g.*, whether it asks where Sally will look, what Sally believes, or what she will say). Other factors such as whether the characters in the task are dolls, photographs, etc., are also inconsequential. Our focus on the Sally-Anne task, then, and the exact experimental set-up we chose should not affect the validity of the results.

Secondly, several task components were identified as main effects, which improve performance but do not interact with age, including whether the child participated in the experiment (*e.g.*, helped to set up props), whether Sally's absence was explicitly emphasized, and in which country the experiment took place. We do not model such task variations.

Thirdly, the compiled results show a significant, if noisy, effect between age and the proportion of children that answered the false-belief query correctly ($p < 0.001$). Figure 2 shows the findings; it plots the results from each individual study considered, as well as the curve that best fits it. They found that at an age of about 44 months, the odds of answering correctly are even, or 1.0; then, the odds of being correct increase 2.94 times for every year. The linear regression model which considers only age is $y = -3.96 + 0.09 \cdot [\textit{age in months}]$, with $r^2 = 0.39$[1]. Their best statistical model, which had six variables (including age, the country in which the experiment took place, and child participation), yielded an R^2 of 0.55. The results clearly document the developmental shift that seems to happen between roughly 3 to 4.5 years of age where children go from being mostly incorrect to mostly correct on the standard false-belief task.

We also consider an experiment involving the avoidance task (Leslie et al., 2005). The experiment, performed with 4.75-year-olds on average, supports the belief that this task is more difficult than the standard task, and provides evidence for the second developmental shift. After several children were eliminated for failing the false-belief task, only 25% of 16 children correctly answered the query of "Where will Sally try to put the fish." The experiment showed, however, that by asking the question in terms of where the *first* place Sally will try to put the fish is, almost three times as many children (71%) passed the task; we refer to this as "look-first avoidance." Overall, the results suggest that children gain the ability to understand others' desires and their implications after they gain the capability to understand their beliefs.

Discussion of Experiments

The area of how children develop theory of mind remains controversial. One of the pressing questions that emerges from the literature is whether the various developmental shifts are due to learning concepts and causal laws (for clarity, we refer to this as "learning"), as the theory-theorists strongly posit, or due to increasing capabilities/functionality of mechanisms of the brain (we refer to this as "maturation"), as others argue. There is certainly evidence for both.

[1]This model transformed the proportion correct, p, via a logit transformation, $\ln(p/(1-p))$ where "ln" is the natural logarithm.

Leslie et al. (2004) argues that maturation of processing capabilities and resources, alone, can account for all ToM developments, and have designed reasonable process models (*e.g.*, ToMM-SP) demonstrating the idea's plausibility. Further evidence shows that the capabilities of specific mechanisms in the brain (such as selection processing and inhibition of beliefs) play a crucial role in ToM (German & Hehman, 2006; Carlson, Moses, & Claxton, 2004).

Wellman et al. (2001), however, makes several arguments for learning over maturation based on the results of the meta-analysis; specifically, the strong presence of task manipulations that act as main effects (*e.g.*, child participation). If maturation were true, presumably many task manipulations would interact with age since they should help younger children's processing competence more than older children's; however, they do not. The presence of such manipulations does, however, support conceptual change accounts. Overall, the authors argue that there is a potential interrelation of learning and maturation: children improve as they grow and acquire conceptual understanding of ToM but, within an age group, processing capabilities could be highly correlated with performance and could account for much of the variance.

Many of the above papers argue against simulation theory based on these results; however, much of the arguments are neither substantive nor well supported. Wellman et al. (2001) argues that, since children do not systematically err about their own false beliefs, simulation theory is not as plausible; however, this could easily be explained by children remembering their own past mental states. Leslie et al. (2004) simply says about simulation theory, "it is also hard to see a role for 'simulation' in accounting for this data... the mechanisms of theory of mind might simply figure out what one would do... there is currently no evidence that it is the first-person singular." The opposite argument could just as easily be made. Unfortunately, there are few developmental accounts available for simulation theory; (Harris, 1992) is an exception, and states that a child's inability to perform simulation early on may be due to memory limitations. In general, simulation theorists support their arguments as in (Gallese & Goldman, 1998), with the presence of mirror neurons that fire both when one views an action and when one performs it.

Overall, we agree in part with Wellman et al. (2001), who say that the ability of children to recognize false-beliefs in others is due to both learning and maturation, accounting for the first developmental shift we discussed where children gain the ability to recognize and predict beliefs in others. We argue, however, that the second developmental shift that occurs, which results in children being able to account for both beliefs and desires to predict another's behavior, is due to children gaining the ability to perform simulation. This accounts for 4.75-year-olds' inabilities to reliably answer the avoidance query: they are still in the middle of learning and maturing this ability. Note that this view is not necessarily incompatible, at the process level, with some of the others; *e.g.*, in highly complex situations, there is not much difference between Leslie et al. (2004)'s SP mechanism inhibiting everything that should *not* be used and operating only on what is left, and identifying pertinent beliefs and decision-making processes and subsequently using them in simulation.

Some recent experiments also suggest that very young children (15 months of age) can perform implicit (non-verbal) false-belief tasks (Onishi & Baillargeon, 2005). This supports the theory of processing mechanisms in the brain that work with false-beliefs and, further, suggests that the ability to recognize situations involving false-beliefs develops before the ability to explicitly reason about them. We anticipate further modeling work concerning this would be compelling.

Core Cognitive Architecture

As our core cognitive architecture we use ACT-R, a hybrid symbolic/sub-symbolic production-based system (Anderson, 2007). ACT-R consists of a number of modules, buffers and a central pattern matcher. Modules contain a relatively specific cognitive faculty typically associated with a specific region of the brain. For each module, there are one or more buffers that communicate directly with that module as an interface to the rest of ACT-R. At any point in time, there may be at most one symbolic item, or "chunk," in any individual buffer; the module's job is to decide when to put chunks into a buffer. Chunks are used to represent knowledge or memories related to any of the modules/buffers, and, in addition to symbolic information, contain subsymbolic information (*e.g.*, activation). The pattern matcher uses the contents of the buffers, if any, to match specific productions which, when fired, can modify the current contents of the buffers. Ties between competing productions are broken based on the productions' expected utilities, which can be initially set and adjusted via a reinforcement learning process; random noise can also be added in during execution to affect production selection.

The relevant modules of ACT-R to this paper are the intentional and declarative modules. In addition, ACT-R interfaces with the world through the visual, vocal, motor and aural modules. The open-source, robotic simulation environment Stage (Collett, MacDonald, & Gerkey, 2005) was used as the "world" of the model in order to enable fast model development and data collection.

ACT-R is able not only to learn new facts and rules, but also to learn which rule should fire (called utility learning in ACT-R). It accomplishes this by learning which rule or set of rules lead to the highest reward. ACT-R uses an elaboration of the Rescorla-Wagner learning rule and the temporal-difference algorithm (Fu & Anderson, 2006). This algorithm has been shown to be related to animal and human learning theory.

Any time a reward is given (*e.g.*, children being told they responded with the correct answer), a reward is propagated back in time through the rules that had an impact on the model getting that reward. Punishments are performed similarly.

Model Description

As stated above, our model is based on the conjecture that, as children grow, they learn and mature simultaneously; *i.e.*, as they develop, they learn to take advantage of their maturing ability to select between competing beliefs. Further, we believe that being able to select between beliefs acts as a precurser for simulation, which allows people to use the beliefs and desires of others to predict and understand their behavior, and is ultimately what provides full-fledged ToM.

In our model, the Sally-Anne task takes place in the Stage simulator, which feeds the model visual information; *i.e.*, it passes the model visual locations to fixate on and, when attended to, what is at that location. This allows the model to "watch" the Sally-Anne play unfold. As the story unfolds, the model explicitly notes what happened (*e.g.*, Sally moved the marble into her box), and who saw it happen (*e.g.*, only Anne saw herself move the marble into her box). After the play completes, the model is asked several false-belief questions. If the model answers a question correctly, the model is rewarded; otherwise, it is punished.

We first describe the core mechanisms that enable ToM. Then, we describe how the model learns to effectively use these mechanisms (as well as develops the ability to use them). Although much of the description is in the context of the Sally-Anne task, as are our experiments, recall that this acts as a proxy for false-belief tasks in general and our results are not specific to this task (Wellman et al., 2001).

Theory of Mind Mechanisms

When its goal is to answer a query about someone's belief, a fully-developed model will answer the question similar to Leslie et al. (2004)'s ToMM-SP. As the story unfolds, the model generates possible beliefs for the marble's location; for the standard Sally-Anne task, then, this set is $\{$sallys-box, annes-box$\}$. The model first retrieves the TB answer because it has the highest activation. It realizes, however, that the answer is not correct since Sally does not know about it. To address this, it considers the various possible beliefs of the marble's location and, from these, it selects the most salient belief that Sally was known to be privy to, the FB box.

When faced with an avoidance task, a fully-developed model will first use the above process to select knowledge to use as input to its simulation. For the Sally-Anne avoidance task variant, the simulation's input would be the different boxes, as well as Sally's belief of the location of the kitten. All subsymbolic information of the knowledge, including activation levels, is preserved. The model next performs simulation by spawning a sub-model with: this input; access to the model's productions and cognitive resources; and the goal of deciding where to put the kitten (Kennedy, Bugajska, Harrison, & Trafton, 2009). Then, the sub-model can infer that, if Sally wants to put the fish under a box without the kitten, she will put it under the TB box.

Developmental Mechanisms

As stated, our model both learns and matures as it develops ToM. The learning mechanism is similar to standard ACT-R learning. The model begins with a production that answers false-belief queries simply by retrieving the belief chunk with the highest activation, and returning it. It can learn, however, to consider an alternate competing production that, upon the retrieval of the belief, considers whether the person the query is about knows about the belief. This production acts as the gateway to the selection process. Learning over time can teach the model to exclusively favor this production, as it ultimately leads to the correct answer. A similar process occurs when learning to perform simulation.

ACT-R does not normally model increasing functionality in the brain. In order to model maturation, therefore, we introduce the notion of a "maturation parameter." This parameter determines whether a model has the ability to fire certain sets of productions (*i.e.*, whether the model is mature enough to have that functionality). Since maturation is not an "all or nothing" concept, and happens gradually, the parameter acts as a guideline for how strong the model's abilities are at that moment. Any time the model attempts to fire a maturing set of productions, their availability is random according to the parameter (*e.g.*, if a randomly selected number is less than the parameter, the productions will be able to fire). Intuitively, maturation parameters should be correlated with age: the older the child, the higher the parameter.

In the case of selecting between different beliefs, the maturation parameter is called the "selection parameter" and determines the availability of the productions that select between beliefs. A model with a selection parameter of 0 would never be able to correctly select a false-belief as the involved productions would be unable to fire; a model with a selection parameter of 0.5 would be able to do so on half of its attempts; and a model would a selection parameter of 1 will always be able to fire the involved productions.

In the case of simulation, the model should be able to perform larger and larger simulations as it ages. This is in accordance with Harris (1992)'s view that children have difficulty performing simulation early on due to memory limitations. The "simulation parameter" determines the availability of the productions that perform simulation, given the size of simulation that is being attempted; for low sizes, the model is more likely to be able to do it, but at high sizes the model becomes overwhelmed and cannot process all the data, and so simulation is less likely. Specifically, any time a simulation is attempted, the probability that the simulation productions will be available is $\min(1, sp/s)$, where sp is the simulation parameter and s is the size of the attempted simulation. The size of the simulation is discussed in the subsequent section.

Modeling Developmental Progress

The model begins at approximately 2 years of age with the ability to generate multiple possible beliefs (Leslie et al., 2004). Model development mirrored the two ToM developmental phases. With respect to the first phase and the standard false-belief task, the model has a selection parameter of 0.5, but does not yet know to do the selection; *i.e.*, when it initially retrieves the most salient belief, it does not know to check whether Sally saw it and simply returns the belief. Of course, the most salient belief is likely the true-belief, and so the model will be incorrect, leading to a punishment. This causes the model to begin to explore using the selection process. If the model is able to access that functionality (*i.e.*, if a number randomly selected at the time of the attempt is less than the selection parameter), it will attain the correct answer and receive reward; otherwise, it will default to returning the initially-retrieved belief, likely leading to punishment. Note that if this occurs, the productions leading to the selection attempt will incur lower expected utility, making it less likely that the model will attempt selection during the next trial.

Experience is simulated by engaging the model in false-belief trials and by slowly increasing the selection parameter. Therefore, as the model grows more experienced, it concurrently learns to utilize its selection mechanism and is able to more reliably perform selection: by the age of about 44 months (3.7 yrs), the selection parameter is up to 0.8, and by the age of 68 months (5.7 yrs) that parameter equals 0.95. Note that, as the selection parameter increases, so does the efficacy of learning, since more trials that attempt to select the false-belief do so successfully and receive positive reward. Learning was concentrated such that about 2 trials approximates 12 months of experience; the function relating learning trials to age was determined post hoc after comparing our results with those of (Wellman et al., 2001).

The second developmental component (concerning the avoidance task) occurs in an analogous way. Whenever the child successfully answers the standard false-belief task, it is queried about the look-first avoidance task (and, upon successfully answering that, is further queried on the standard avoidance task). The model first tries to calculate Sally's belief exactly as in the standard false-belief task; note that, especially at early ages, it may or may not be able to do so and may end up thinking about either the TB or FB box. Once a belief is in hand, the model initially does not know what to do with it; so it defaults to where it would put the kitten, the FB box, resulting in punishment. Over time, the model will start using the initial belief as input to simulation. If the model is able to simulate, it will return the box other than the belief; otherwise, it will again default to returning the FB box.

As mentioned, the model's ability to perform simulation is dependent on a simulation parameter, which in turn is dependent on the "size" of the simulation. For the look-first avoidance query, the simulation size is 1, as the child is being asked to predict Sally's actions only one step in the future. For the standard avoidance query, the simulation size is set to 3[2]. When the model begins at age 2, the simulation parameter is 0 and so no simulation is possible; by age 56 months (4.7 yrs), it is 1, and by age 72 months (6 yrs), it is 5.

For all models, we kept most of the ACT-R parameter defaults. We did change the utility noise parameter (set at a moderate 1.0) to allow low-use productions to occasionally fire. Because the rate of learning is dependent entirely on the utility learning rate parameter (set at the default of 0.2), learning occurred quite quickly in this model. Utility learning rate could be scaled down substantially to match actual development and learning time. In order to do this correctly, it would be important to know approximately how often children encounter false-belief and avoidance tasks, and learn from them.

[2]Although this is ad hoc, with such limited data to match, a more pleasing parameter choice and justification is not possible.

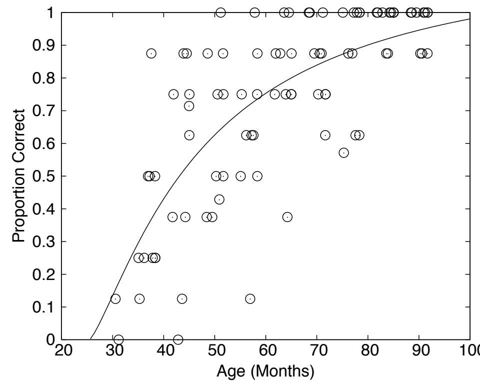

Figure 3: Model results showing a scatterplot of the standard false-belief results and best-fit curve.

Model Results

In our first experiment, corresponding to the first phase of model development, we started testing the model at age 32 months (2.7 yrs), and test roughly every 7 months until the model reaches around age 92 months (7.7 yrs), for a total of 10 tests. Each test period consisted of 8 repetitions of the Sally-Anne task, including all three queries. During these tests, learning is turned off in order to reliably test the model's abilities at that age. To simulate the variability of children's development, we randomly perturbed the models' starting ages around their *a priori* value of 2 years, selecting uniformly in the range [17, 31] months. This made the age of the models in our experiment comparable to the ages of the children in the meta-analysis (Wellman et al., 2001).

Figure 3 shows the results for the false-belief task, and plots each model's age during a test period against the proportion of correct answers the model gave during the test. The graph appears very similar, visually, to that of Figure 2, and shows a clear learning trend as well as noise which presumably stems from different maturation levels. Using Wellman et al. (2001)'s linear regression model (which considers only age) on this data, $r^2 = 0.51$ with a residual standard error of 1.73. This is considerably higher than their $r^2 = 0.39$. It also approaches the R^2 of their multi-variate model, 0.55. We argue, then, that our model is stronger since it is both a process model that learns to perform this task, as compared to a statistical model, and depends on fewer parameters.

Note that this curve is due to an interaction between the selection parameter increasing, and the model learning that attempting to select between beliefs often leads to the correct answer. We expect, therefore, that if the selection parameter increased more slowly, learning would be impeded and models' performance would not improve as quickly.

Our avoidance false-belief results were also compared to those of (Leslie et al., 2005), which showed that 71% of children around the age of 4.75 years could answer the look-first avoidance query but only 25% could answer the standard avoidance query. We were able to match these results, but further experimental data is needed in order to distinguish our parameterization from other valid possibilities.

Discussion

We have shown in this paper a cognitive model for theory of mind. Our model borrows ideas from all three main postulates of ToM to develop a cohesive explanation for how ToM functions. The model uses a selection process to identify the beliefs and knowledge others may have; then, to predict the desires and behaviors of others, it uses the identified concepts as input to its own decision-making mechanisms, simulating what the model would do in the other's place. This ToM functionality develops by concurrent learning and maturation of the required functional capabilities. The model was found to be a good match to existing data from developing children.

One of the strengths of this model is that it generalizes to many other types of false-belief and ToM tasks. The maturation parameters are very general, and can be applied with little change to other tasks. The same holds true for simulation; the cognitive mechanism which enables it can accept, and work with, any input. The learning of ToM in this paper is not as general, as it chooses between productions which are relatively task-specific; however, if the model were to have experience on a variety of ToM tasks, we expect that it would generalize what it learns into a broader concept.

Our work is also distinguished from previous work in cognitive architectures. Laird (2001)'s QuakeBot performs mental simulation of opponents to predict their behavior, for example, but to our knowledge their approach has not been matched against human cognitive data.

A future step is to explicitly address other observed ToM phenomena. One experiment added a third "neutral" box to the avoidance task, introducing a second correct answer, and had both children and adults as subjects (Leslie et al., 2004). The study showed that children have a bias towards the TB box, whereas adults have a bias towards the new neutral box. Our model does predict this phenomena for children, since the TB box is the correct box with the highest activation (it is the last box to receive a kitten, and it is identified as the true-belief of the kitten's location before the selection of beliefs begins), and so it is the answer that simulation will select. As far as the results for adults, we believe that with further learning, simple simulations can be avoided in favor of general, learned inference rules. In this case, therefore, adults are simply returning an answer that is true from anyone's perspective. The paper describes further experiments that our model can predict, but that is outside the scope of this paper.

Acknowledgements

This work was supported by the Office of Naval Research under funding document N0001409WX20173 to JGT. The views and conclusions contained in this document should not be interpreted as necessarily representing the official policies of the U. S. Navy.

References

Anderson, J. R. (2007). *How can the human mind occur in the physical universe?* Oxford University Press.

Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a "theory of mind"? *Cognition*, *21*.

Carlson, S. M., Moses, L. J., & Claxton, L. J. (2004). Individual differences in executive functioning and theory of mind: An investigation of inhibitory control and planning ability. *Journal of Experimental Child Psychology*, *87*, 299-319.

Cassimatis, N. L., Bello, P., & Langley, P. (2008). Ability, breadth, and parsimony in computational models of higher-order cognition. *Cognitive Science*, *32*(8), 1304-1322.

Collett, T. H. J., MacDonald, B. A., & Gerkey, B. P. (2005). Player 2.0: Toward a practical robot programming framework. In *Proceedings of the Australasian Conference on Robotics and Automation (ACRA '05).*

Fu, W., & Anderson, J. (2006). From recurrent choice to skill learning: A model of reinforcement learning. *Journal of Experimental Psychology: General*.

Gallese, V., & Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences*, *2*(12).

German, T. P., & Hehman, J. A. (2006). Representational and executive selection resources in 'theory of mind': Evidence from compromised belief-desire reasoning in old age. *Cognition*, *101*, 129-152.

Gopnik, A. (1993). How we know our minds: The illusions of first person knowledge of intentionality. *Behavioral and Brain Sciences*, *16*, 1-14.

Harris, P. L. (1992). From simulation to folk psychology: The case for development. *Mind and Language*, *7*.

Kennedy, W. G., Bugajska, M. D., Harrison, A. M., & Trafton, J. G. (2009). "Like-me" simulation as an effective and cognitively plausible basis for social robotics. *International Journal of Social Robotics*, *1*(2), 181-194.

Laird, J. E. (2001). It knows what you're going to do: Adding anticipation to a quakebot. In *Proceedings of the International Conference on Automomous Agents.*

Leslie, A. M., Friedman, O., & German, T. P. (2004). Core mechanisms in 'theory of mind'. *Trends in Cognitive Sciences*, *8*(12), 528-533.

Leslie, A. M., German, T. P., & Polizzi, P. (2005). Belief-desire reasoning as a process of selection. *Cognitive Psychology*, *50*, 45-85.

Onishi, K. E., & Baillargeon, R. (2005). Do 15-month-old infants understand false beliefs? *Science*, *308*.

Wellman, H. W., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, *72*(3), 655-684.