Title
Developmental Accounts of Theory-of-Mind Acquisition: Achieving Clarity via Computational Cognitive Modeling

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 28(28)

ISSN
1069-7977

Authors
Bello, Paul
Cassimatis, Nicholas

Publication Date
2006

Peer reviewed

Developmental Accounts of Theory-of-Mind Acquisition: Achieving Clarity via Computational Cognitive Modeling

Paul Bello (Paul.Bello@rl.af.mil)
Air Force Research Laboratory - Information Directorate
525 Brooks Rd.
Rome, NY 13441 USA

Nicholas Cassimatis (cassin@rpi.edu)
Rensselaer Polytechnic Institute - Dept. of Cognitive Science
110 8th St.
Troy, NY 12180 USA

Abstract

People's ability to predict and explain the beliefs, desires and actions of others, often called their *Theory of Mind* (ToM) is an central component of cognition. There is ample experimental evidence suggesting a major developmental shift in children's ToM somewhere between the ages of three and five (Wellman, Cross & Watson 2001). These results are often interpreted in terms of children acquiring rich understanding of belief and its representational nature. Both of these descriptors are ambiguous and controversial. This makes evaluating the evidence of such a shift, and its implication about the knowledge and mechanisms underlying ToM more difficult. Computational cognitive models are precise by nature, and can help resolve many of these difficulties. We present Polyscheme/ToM, a computational model of human behavior in some important ToM tasks. It uses a domain-general mechanism for representing alternate states of affairs in order to represent the way other people view the world. We demonstrate that the model has several implications: 1) it is possible to succeed in false belief tasks without explicitly representing beliefs; 2) behavior in experiments purporting to show success in false-belief-like tasks at earlier ages require significantly simpler mechanisms than success in the traditional false-belief task, and 3) the shift from younger to older behavior in these tasks can be accounted for by small, but highly consequential changes in cognitive mechanisms.

Problems in interpreting ToM experiments

Developmental psychologists have designed a number of clever and important experiments to chart out the acquisition of so-called *Theory of Mind* (ToM) by young children (Liu & Wellman 2004). The central debate among theoretical frameworks for describing the acquisition and nature of ToM is focused differing views on how children might be able to predict the actions of other agents. For example, one these theories characterizes children as "little scientists" who construct theories of human behavior through connecting stimuli, mental states, and actions taken by observable human agents. This framework is often called *theory-theory* (Gopnik & Meltzoff 1997). Alternatively, proponents of *simulation theory* subscribe to the assumption that children are able to assume the perspective of other people, and use their own cognitive apparatus for dealing with their perspective of the world as a good first-approximation to that of the person to be simulated (Goldman 1989). The results of many different experiments have been used by proponents of both views to buttress support for their respective positions. However, these results are often presented using terms such as representation, belief, knowledge, simulation and model, all of which can be interpreted in variety of ways to suit different purposes. This has several important consequences for the theoretical impact of ToM experiments:

1. **Disagreement among various researchers concerning the nature of mental states such as beliefs, desires, and intentions** have led to any number of interpretations for both performance and developmental data in accounting for them. Philosophers who have studied these problems for millennia cannot even agree upon precise meanings for these terms, and remain at odds over the general properties these states may possess. For example, do children actually represent the beliefs of other agents (i.e., do they have specific beliefs about beliefs), or do they instead have an overall representation of another person's perspective on the world without explicitly thinking about the individual beliefs that make up that representation, as researchers such as Josef Perner seem to suggest (Perner 1991).

2. **Disagreements over definition make experimental data difficult to interpret.** It is often assumed that the false-belief task as presented in (Wimmer & Perner 1983) is the gold standard for determining whether a child has acquired the so-called *representational theory of mind*, and thus a thorough understanding of beliefs. A meta-analysis conducted in (Wellman et al. 2001) strongly corroborates success on these kinds of tasks with a relatively narrow age range (usually between 3.5 and 5.0 years old). However, there are studies that suggest that younger children (e.g., (Wellman & Bartsch 1988)) and even chimpanzees (Call, Hare, Carpenter & Tomasello 2004) have ToM abilities normally only ascribed to older human children.

3. **Imprecision makes it difficult to assess the magnitude of development.** Without a specific and precise account of cognition, both pre- and post- success on the false belief task, it is difficult to assess the magnitude of the development in children's thinking that the behavior shift in these tasks represents. Theory-theorists are faced with a fundamental problem of accounting for the induction and adaptation of a large knowledge-base corresponding to a näive theory of human behavior. It is difficult to asses how large of a problem this is without a more precise characterization of these theories in both younger and older children. Simulation theorists aren't faced with the problem of a large theory, they face a more fundamental set of difficulties entailed by their own commitments, partially outlined

in (Saxe 2005). Is it reasonable to assume that children acquire this large theory in a short period of time, or is the theory supplemented by mechanisms which allow for use of the child's already-compiled knowledge? Precise instantiations of theory-theory and simulation theory would help clarify these issues.

We have set out to develop a computational model of children's behavior in ToM tasks in hopes of shedding light on these apparent confusions.

Common Themes in ToM Tasks

We developed our model by first identifying some common (and commonly accepted) themes among some of the now classic "metarepresention" tasks in the literature. The more well-known of these tasks are the *appearance-reality distinction* (Flavell 1986), the *smarties task* in (Hogrefe, Wimmer & Perner, 1986), and infamous *false-belief task* (Wimmer & Perner 1983).

- **Children must represent separate and potentially contradictory states of affairs.** In the case of appearance-reality tasks, this corresponds to how objects appear versus how they actually are. In the false belief task children must separate the state of affairs they know to be true from the state of affairs which another person holds to be true, regardless if there is a mismatch between the two. In the smarties task, children must distinguish between the state of the affairs they currently hold to be the case from the state of affairs which they formerly held to be the case.
- **Theories relating people to alternate states of affairs**. The false belief and smarties tasks both require that these alternate states of affairs must be associated with people.
- **Causal theories**, which define the relationships between the multiple states of affairs. For example, perceptual events have a causal impact on the state of affairs associated with another person's perspective.

Two features of these common themes motivate our choice of cognitive modeling framework. First is the need to represent multiple states of affairs and the second is the need to combine perception and reasoning. For these reasons we chose the Polyscheme cognitive architecture (Cassimatis 2005) to implement our ToM model. Polyscheme is rare among cognitive architectures in that it has rich facilities for representing alternate worlds and perhaps is unique in being designed to combine this ability with mechanisms for naturally integrating reasoning and perception.

Polyscheme

Polyscheme is a cognitive architecture which has been used to model phenomena as wide-ranging as reasoning with näive theories of physics to syntactic parsing mechanisms (for English) using the same set of common functions (Cassimatis 2005). A Polyscheme model consists of a sequence of interacting modules called specialists that represent and make inferences about aspects of the world. Specialists are based on their specialized data structures and algorithms. Specialists are all able to represent and make inferences about alternate states of worlds. For the purpose of this discussion, we will only deal with two specialists, the "rule specialist" which represents propositions symbolically and reasons about them using standard rule-based techniques (*modus ponens* and *modus tollens*), and the "temporal-perception specialist," which keeps track of when things are seen. Specialists in Polyscheme communicate through a propositional language. These propositions have *truth-values* on the following scale: <Maybe, Likely, Very Likely, Certainly>. The truth value of a proposition is a tuple (x,y), where x represents negative evidence in favor of the truth of the proposition, and y represents the contrasting positive evidence. Specialists vote on propositions by simultaneously focusing on them, and forming opinions on their truth value. Propositions are of the form $R(x,y,t,w)$ and state that relation R holds over entities x and y over temporal interval t in world w. When $t = E$, that relation holds over all time (eternity). The "real world," or the world as it is seen by Polyscheme in a first person sense, is denoted as *R*, so in the default case, $w = R$. Unless there is reason (i.e. perception or inference) to believe otherwise, specialists assume what is true in *R* is true in other worlds. For a more thorough overview, refer to (Cassimatis 2005).

Polyscheme/ToM

Polyscheme/ToM is a Polyscheme model augmented to account for behavior in theory-of-mind tasks. It is motivated by the common themes in the metarepresentation tasks which we highlighted in the last section. The following outlines our model of children who normally succeed in these tasks, roughly corresponding to four-year-old behavior:

- **Alternate worlds represent states of affairs:** Polyscheme/ToM can reason about alternate worlds. In our model, we use such worlds to represent the mind of other agents.
- **The `MindOf` predicate captures the relation between states of affairs and people.** Polyscheme/ToM implements a special relation called `MindOf`, which holds over a person and a world which represents their mind. For example, if John thinks that chair is in front of the piano, we say that `MindOf w john t R + FrontOf chair piano t w .` Notice that while `MindOf` defines *w* as being related to John, the proposition `FrontOf` is indexed to John's mind, which is being represented by *w*. This captures the notion that other agents may view the world differently than Polyscheme/ToM does. While perception is always veridical for Polyscheme/ToM (that is, everything perceived in the real world is assumed to be certainly the case, this is not so for other agents. To account for the discrepancy, Polyscheme implements two different ways of creating alternate worlds which we describe later.
- **Causal theories:** Polyscheme/ToM contains a theory of perception, and its effect on producing representations in other minds. Particularly, Polyscheme/ToM knows that certain kinds of occlusion make certain perceptual modalities unavailable. If Polyscheme/ToM is not blindfolded, and John happens to be, then Polyscheme/ToM can modify the status of propositions in alternate world representing John's mind by reasoning about them in

terms of how (lack of) perception of unique objects and their attributes impute selected mental states to John. The rule corresponding to this is `NOT CanPerceive ?person ?Attribute ?t R + MindOf ?w ?person E R ==> , Blocked perceptionRule ?Attribute ?t ?w .`, stating that if a person can't perceive a particular type of attribute, and the person's mind is represented by the world *w*, then we block perception of any proposition corresponding to the attribute in question in the mind of this person. If John is blindfolded, he therefore doesn't have any solid conception concerning the location of his coat if someone moves it out of the room while he can't see.

`MindOf` works in the following way: by default, when Polyscheme/ToM instantiates a world *w*, it inherits all propositions in certain in *R* as being certain in *w*. We call this *ontological inheritance*. When *w* represents the mental state of a person, e.g., $MindOf(w,Susan)$, only inherits propositions into *w* as being likely. We call this *epistemic inheritance*. By allowing things to be different in *w* than they are in the real world, we allow Polyscheme/ToM to have a representation about an external agent's view of the world. `MindOf` allows Polyscheme/ToM to keep these representations independent, yet connected through causal theories.

Polyscheme/ToM: Action Prediction

So far we have only described how our model decides to infer what another person believes to be the case. Our model uses the following set of rules to predict what a person will do:

- `Wants ?cat ?person ?tWanting ?w + Category ?object ?cat E ?w + Location ?object ?loc ?tWanting ?w ==> , SearchesAt ?person ?loc ?tWanting ?w .`
- `NOT CanPerceive ?person ?Attribute ?t R + MindOf ?w ?person E R ==> , Blocked perceptionRule ?Attribute ?t ?w .`

Which amounts to stating "IF P wants OBJECT (of a certain category) and in his mind the object is at LOCATION, THEN P will search for OBJECT at LOCATION." This *desire-action rule* is preliminary and will need to be expanded because it only deals with perception and location for now. However, it outlines how we can predict what a person will do based on the state of affairs that THEY, and not Polyscheme/ToM, finds to be the case.

Three-year old behavior

Our model of four-year old behavior explains how people can succeed in metarepresentation tasks by relating the state of affairs that someone else believes in to the actions they take (using the world mechanism and `MindOf` predicate). We can account for three-year old behavior by simply removing these from the model. This would lead to the following rule:

```
Wants ?cat ?p ?tWanting R + Category ?o
?cat E R + Location ?o ?loc ?tWanting R +
Perceive ?p ?loc ?tSeeing R + Before ?tSeeing
?tWanting E R ==> , SearchesAt ?p ?loc
?tWanting R .
```

Stating that "IF P wants OBJECT and P sees object at LOCATION, THEN P will search for OBJECT at LOCATION."

The difference in three- and four-year old models

To summarize, our model accounts for four-year old success with two mechanisms: the ability to represent the state of affairs a person believes in with a world that is associated with that person using the `MindOf` predicate; and an action rule that takes this state of affairs into account. That is the three-year old action rule merely makes reference to what the three-year old (not other people) believes is true about the world and what other people have seen. The four-year old desire-action rule makes reference to an alternate state of affairs associated with another person:

- **3-year-old:** "IF P wants OBJECT and P **perceives** object at LOCATION, THEN P will search for OBJECT at LOCATION.
- **4-year-old:** "IF P wants OBJECT and **in his mind** the object is at LOCATION, THEN P will search for OBJECT at LOCATION."

Interpreting Conflicting Results

We will contrast the study of false-belief made famous in (Wimmer & Perner 1983) with the claim that children younger than predicted by the age-range reported in the 1983 experiment display some understanding of false belief, particularly as it is related to predicting action (Wellman & Bartsch 1988).

Wimmer and Perner's Experiment

The paradigmatic false-belief task involves two subjects, a child and another agent who are in a room with the experimenter. There are two cookie jars in the room, with both child and the agent knowing that there is a cookie in jar 1, but not in jar 2. The agent is asked to leave the room for a moment. While the agent is gone, the cookie is switched from jar 1 to jar 2. The agent is then reintroduced into the room, and the experimenter asks the child which jar the agent will look in to find the cookie. A task of this type requires a notion of common knowledge, but also that the child understands that the agent can be in possession of faulty information (or false belief), which ultimately will have consequences for its behavior. As we've stated previously, a transition from failure to success on this task is in the narrow range of 3.5 to 4.5 years of age.

Bartsch and Wellman's Experiment

Bartsch and Wellman describe a number of experimental results in their 1988 paper which seem to indicate that children have a rich understanding of belief before the age predicted by Wimmer and Perner's 1983 experiments. We focus on the third experiment presented in the paper. In the experiment, children were presented with Susan, who is leaving her house for work in the morning. She sees a black magic marker on the table in the kitchen on her way out. There is another marker in the house on the shelf in the living room. The child is told that Susan hasn't seen this particular marker. Susan comes home, and wants to get the marker in order to write something down. The children are asked where Susan will look for the marker. This is not quite the same as the false belief task in Wimmer and Perner, but it is very close. In the false belief task, the object in question (the cookie) changes status. In this case, the marker on the shelf acts as a possible distracter for the child, giving the experimenter some confidence concerning whether or not the child is linking Susan's perception to her beliefs to her actions.

Model: False-Belief

The task proceeds according to the following temporal sequence (from the standpoint of the subject):

- t_1: perceive the locations of the two cookie jars, the location of the cookie, and the status/location of susan, who is the agent that Polyscheme/ToM will make a judgment about.
- t_2: perceive that susan leaves the room. To simplify modeling, we say that susan is blindfolded.
- t_3: perceive that the location of the cookie is changed from its original placement.
- t_4: perceive that susan has her blindfold removed.
- t_5: infer what susan thinks about the location of the cookie.
- t_5: on the basis of this inference, predict what susan will do if she wants the cookie.

To model this scenario, we make three simplifying assumptions: first, that instead of leaving the room, susan is blindfolded; secondly that both susan and Polyscheme/ToM are told the original location of the cookie; and finally, that having cookie-jars or other occluders aside from the blindfold doesn't do anything but overcomplicate the essence of the task. The Polyscheme input file describing the task looks like:

```
###RULES

Wants ?cat ?person ?tWanting ?w + Category
?object ?cat E ?w + Location ?object ?loc
?tWanting ?w ==> , SearchesAt ?person ?loc
?tWanting ?w .

NOT CanPerceive ?person ?Attribute ?t R
+ MindOf ?w ?person E R ==> , Blocked
perceptionRule ?Attribute ?t ?w .

Blindfolded ?person ?t R + MindOf ?w ?person E
R ==> , NOT CanPerceive ?person Location ?t R .

Location ?object ?person ?t w + Before ?t1 ?t2
E R ~~> , Location ?object ?person ?t2 w .

NOT Blocked perceptionRule Location t1 ?w +
Perceive ?object t1 ?w ==> , Location ?object
p3 t1 ?w .

NOT Blocked perceptionRule Location t4 ?w +
Perceive ?object t4 ?w ==> , Location ?object
p4 t4 ?w .

###INPUTS
```

<u>Time: 1</u>

```
NOT Blocked perceptionRule Location t1 R .
NOT Blocked perceptionRule Location t4 R .
NOT Blocked perceptionRule Location t6 R .

Before t1 t4 E R .
Before t1 t6 E R .
Before t4 t6 E R .

Perceive cookie t1 R .
NOT Blindfolded susan t1 R .
```

<u>Time: 10</u>

```
Blindfolded susan t2 R .
Blindfolded susan t3 R .
Blindfolded susan t4 R .
Blindfolded susan t5 R .
```

<u>Time: 20</u>

```
Perceive cookie t4 R .
```

<u>Time: 30</u>

```
NOT Blindfolded susan t6 R .
```

<u>Time: 40</u>

```
MindOf w susan E R .
```

<u>Time: 50</u>

```
WONDER Location cookie p3 t1 w .
WONDER Location cookie p3 t6 w .
WONDER Location cookie p4 t6 w .
```

<u>Time: 60</u>

```
Category cookie Cookie E R .
Wants Cookie Sue t6 w .
```

While we have previously given explanations for the first two rules in our model, it behooves us to quickly describe the operation of the rest of the model describing Polyscheme/ToM's näive theory of perception. The third rule in the model simply states that blindfolded people can't perceive locations, while the fourth rule states that if an object is perceived as being at a particular place at time t, it is likely to be there at time $t+1$. The remaining rules are what we called *perceptual rules*, and defeasibly mediate between perception and knowledge through the use of the relation `Blocked`. When `Blocked` is true, perception doesn't lead to knowledge, however in the default case (where a person isn't blindfolded), `Blocked` is assumed to not be the case.

False-Belief Task: Inference Trace

Extrapolating the inferential chain in the false-belief task allows us to present how our model is able to predict the behavior of another agent within the task. All propositions referenced to the real-world *R* in our input file are assumed to be certainly true by the Polyscheme/ToM model. As soon as the Polyscheme/ToM model simulates the mind of the other agent (by invoking MindOf with a world argument *w*), some interesting inferences begin to be made:

- By invoking `MindOf`, using it to bind *w* to susan, and wondering about the location of the cookie in *w*; Polyscheme/ToM creates an alternate world *w* corresponding to susan's mind using epistemic inheritance, resulting in `Location cookie p3 t1 w` being considered "likely true" rather than certain.
- Rule 3, which tells us that blindfolded people cannot perceive locations, matches, and given that Polyscheme/ToM knows that the other person is blindfolded from t2 through t5, it infers that the other person cannot perceive location at those times.
- Now, given the fact that the other person cannot see from t2 through t5, Polyscheme/ToM infers that any perception-Rule referencing any time in the interval t2 through t5 is blocked.
- Recall that since Polyscheme/ToM assumes ontological inheritance by default, Perceive cookie t4 R is considered to be true in *w*. But since the perceptionRule at t4 is blocked, no inference is made such that `Location cookie p4 t4 w` is certainly true.
- When Polyscheme/ToM wonders about the location of the cookie from susan's perspective at time t6, it applies the 3rd rule to `Location cookie p4 t4 w`, but due to contradiction with `Location cookie p3 t4 w` it retracts its' belief, and re-assumes the last piece of true information available, which corresponds to before susan was blindfolded.

Model: Action-Prediction

The action-prediction task in (Wellman & Bartsch 1988) proceeds according to the following temporal sequence (from the standpoint of the subject):

- t_1: perceive the locations of the markers on both the table and the shelf, and notice that susan only perceived the marker on the table.
- t_2: perceive that susan wants the marker.
- t_3: wonder where susan will search to find it.

The only difference between this model and the previous model (aside from obvious differences in task structure), is that this task doesn't make any reference at all to worlds representing alternate states of affairs or to `MindOf` or its associated machinery. This task also doesn't have a fully cashed out theory of perception (how blindfolding would affect the perception of location-changing, et cetera). It doesn't need one, although it would certainly be simple enough to recast this problem in terms of perceptual limitations.

```
###RULES
```

```
Wants ?cat ?p ?tWanting R + Category ?o ?cat E
R + Location ?o ?loc ?tWanting R + Perceive ?p
?loc ?tSeeing R + Before ?tSeeing ?tWanting E R
==> , SearchesAt ?p ?loc ?tWanting R .
```

```
###INPUTS
```

<u>Time: 1</u>

```
Category markerTable Marker E R .
Category markerShelf Marker E R .

Location markerTable pl:1-0-1 t1 R .
Location markerTable pl:1-0-1 t2 R .
Location markerTable pl:1-0-1 t3 R .
Location markerShelf pl:2-0-1 t1 R .
Location markerShelf pl:2-0-1 t2 R .
Location markerShelf pl:2-0-1 t3 R .

Perceive susan pl:1-0-1 t1 R .
NOT Perceive susan pl:2-0-1 t1 R .
```

<u>Time: 10</u>

```
Wants Marker susan t3 R .
Before t1 t3 E R .
```

<u>Time: 20</u>

```
WONDER SearchesAt susan ?loc t3 R .
```

It should be clear that performance can be modeled in this task with surprising less complexity. In fact, it is just a single application of rule-matching. This is a perfect example of misconstruing a simple behavioral rule with a rule making reference to Susan's mind.

Discussion

Our work on other ToM tasks lends support to the generality of our claims here. We have used Polyscheme/ToM to model the appearance-reality distinction, the so-called "Smarties task," and the false belief task (Flavell 1986) using Polyscheme/ToM (Bello & Cassimatis, under review). Our model demonstrates how precise computational implementations can limit and help to resolve confusions in the interpretation of behavior in ToM tasks. Specifically, Polyscheme/ToM makes the following contributions:

- **Explicit reasoning about beliefs is not necessary for success in false belief tasks.** In our model, children can succeed at false belief tasks about, say, a cookie's location, without reifying (making explicit) the other person's belief about that cookie and associating it with the cookie in the world. In our model, children represent a state of affairs associated with Susan, but in that state of affairs the cookie is a cookie, not a representation or belief about a cookie. It is just a cookie in a different location. In terms of theory-theory, this means that these experiments do not require that a four year old's ontology includes beliefs. It merely requires that it includes states of affairs, objects in the world denoted as persons, and a way of relating them.

- **Precision, in the form of computational models, can serve to clarify ambiguous experimental results.** This is demonstrated by comparing our two models of action-prediction in both the marker-finding (Wellman & Bartsch 1988) and cookie-finding scenarios (Wimmer & Perner 1983), and noting that the marker-finding task doesn't require the special functionality that `MindOf` offers, whereas the cookie-finding task does.

- **More precise assessment of the magnitude of the developmental shift.** Although the differences in our three-year old and four-year old models correspond to a significant shift in inferential power, they do so with surprisingly few changes in knowledge or mechanism. In addition to not requiring specific cognition about beliefs, desires and representing, the "theory" component in our model which explains performance on the false belief task is only one rule about how what one perceives affects his mental state and another rule about how such a mental state together with desire leads to action. The four-year-old innovation is to apply the alternate world mechanism (for which there is independent evidence (Dias & Harris, 1990) to representing other minds. No more sophisticated theoretical or conceptual apparatus is required.

ToM has been relatively unexplored in cognitive modeling. This work demonstrates that precise computable models can illuminate important issues in this literature. By precisely specifying what exactly constitutes children ToM, we have therefore been able to reduce the problem of explaining a broad shift in children's behavior between three and four years to the question of how they aquire/learn/develop one rule, modify a second rule, and begin to apply an alternate world mechanism they already possess toward representing other minds.

References

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

Wimmer H., and Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in children's understanding of deception. *Cognition*, 13, 103-128.

Wellman, H.M., Bartsch, K. (1988). Young children's reasoning about beliefs. Cognition, 30(2), 239-277.

Perner, J. (1991). Understanding the Representational Mind. Cambridge, MA: MIT Press.

Gopnik, A. and Meltzoff, A. (1997). Words, Thoughts and Theories. Cambridge, MA: MIT Press.

Goldman, A. (1989). Interpretation psychologized. *Mind and Language*, *4*, 161-185

Call, J., Hare, B., Carpenter, M., and Tomasello, M. (2004). Unwilling versus unable: chimpanzees understanding of human intentional action. Developmental Science, 7(4), 488-498.

Saxe, R.,(2005). Against simulation: the argument from error. Trends in Cognitive Science, 9(4), 174-179.

Hogrefe, G. J., Wimmer, H., & Perner, J. (1986). Ignorance versus false belief: A developmental lag in attribution of epistemic states. *Child Development*, 57, 567-582

Cassimatis, N. (2005a). Integrating cognitive models based on different computational methods. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Flavell, J. (1986). The development of children's knowledge of the appearance-reality distinction. *American Psychologist*, 41, 418-425.

Bello, P., and Cassimatis, N. (under review). Some unifying principles for computational models of theory-of-mind. *Submitted to the International Conference on Cognitive Modeling*.

Dias, M. & Harris, P. (1990). The influence of the imagination on reasoning by young children. *British Journal of Developmental Psychology*, 8, 305318.