Towards a Computational Analogical Theory of Mind

Irina Rabkina (irabkina@u.northwestern.edu)
Clifton McFate (c-mcfate@northwestern.edu)
Kenneth D. Forbus (forbus@northwestern.edu)
Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60208, USA

Christian Hoyos (choyos@u.northwestern.edu)
Department of Psychology, Northwestern University
2029 Sheridan Road, Evanston, IL, 60208, USA

Abstract

Several theories about Theory of Mind (ToM) have been proposed. The most well-known of these are Theory Theory and Simulation Theory, although alternative and hybrid theories do exist. One such theory, proposed by Bach (2011, 2014), is based on the Structure-Mapping theory of analogy, which has been shown to play a key role in cognitive development. There is evidence that children are more likely to pass false belief tasks when trained using stories that are easy to compare via structural alignment, as opposed to stories that are difficult to compare in this way (Hoyos, Horton & Gentner, 2015). This paper shows how a computational model based on Bach's account can provide an explanation for the Hoyos et al. training study and proposes directions for future research on human subjects.

Keywords: analogy; theory of mind; false belief; structure-mapping; cognitive modeling

Introduction

The mechanisms behind Theory of Mind (ToM) have been hotly debated for decades. According to one popular theory, Theory Theory, children are little scientists who develop theories about others' beliefs (e.g. Gopnik & Wellman, 1992). Another theory, Simulation Theory, suggests that children play out scenarios as if they were the agents involved (e.g. Goldman, 1992). Other accounts include hybrid theories (e.g. Bach, 2011), which attempt to combine aspects of Theory Theory and Simulation Theory (see Related Work).

Another important question is how ToM is learned and when. Interestingly, several studies have shown that at least some aspects of ToM can be improved via brief intervention (e.g. Hoyos et al., 2015; Hale & Tager-Flusberg, 2003; Lohman & Tomasello, 2003). This paper considers one such study and models how analogical generalization may lead to improved performance on the false belief tasks tested. The model generates testable predictions for future work.

We begin by discussing the theories that underlie our model, the Structure-Mapping Theory of analogy (SMT; Gentner, 1983) and Bach's (2011) structure-mapping account of ToM, along with our computational models of analogical matching and generalization used in the model. We then summarize a ToM training study (Hoyos et al., 2015) and describe how our model explains the performance

improvements provided by training. We close with related work and future directions.

Background

We base our model on the Structure-Mapping approach to Theory of Mind proposed by Bach (2011, 2014). Because understanding Structure-Mapping Theory (SMT; Gentner, 1983) is essential to understanding this theory and our model, we describe it first. This is followed by a description of Bach's theory. Finally, we describe the computational models of SMT processes that we are using.

Structure-Mapping Theory

Structure-Mapping (Gentner, 1983) is a theory of analogy and similarity. Under SMT, relational/structural similarity is emphasized over similarity based on features alone. Humans' ability to see these structural similarities across dissimilar cases is a key aspect of higher order cognition, which suggests that structural similarity is used in everyday reasoning. SMT proposes that comparison involves the alignment of elements between two cases, called a base and a target.

Consider a common pedagogical analogy: "A cell is like a city. The city government controls the city. The nucleus controls all the cell's activities. A power station provides electricity. A mitochondrion is like the power station." (Chang & Forbus, 2015). In this example, the cell acts as a target and the city acts as a base. Structural representations of the two are aligned to form a mapping. SMT predicts that the cell maps to the city, the nucleus maps to the government, control of the cell maps to control of the city, and the mitochondrion maps to the power station (Fig. 1). What about providing electricity? Because of the match between the mitochondrion and the power station, we can infer that the mitochondrion does something like providing electricity. This conclusion is called a candidate inference.

SMT can be extended to include analogical generalization (Kuehne et al., 2000). As a person is exposed to alignable cases, generalizations are formed. For example, we can form a generalization between the city and the cell. This would state that "Something like a city or cell has something like a city government or nucleus that controls it and something like

Fig. 1: A visual representation of SMT. Entities are shown as rectangles. Relationships are in diamonds. Dotted lines show correspondences between the two cases. Dashed lines show the candidate inference.

a power station or a mitochondrion that gives it energy." Eventually, generalizations become abstract schemas that can represent, for example, a single type of event. They can be stored in long term or working memory.

SMT Theory of Mind

Bach (2011, 2014) has proposed that ToM is developed via structure-mapping. He proposes that two forms of base domains are used. The first are abstract schemas built up over time. The second are events from autobiographical memory. This provides a hybrid model: Mappings to the schema domain correspond to theories as described in Theory Theory models, and mappings to the autobiographical domain correspond to simulation. For example, to decide whether a person who arrived 15 minutes late to a flight that was delayed by 10 or a person who arrived 15 minutes late to a flight that left on time would be more upset, a person might retrieve an abstract schema that says "people are very upset when they narrowly miss their goal" or they might simulate how they would feel if they were the person in question by mapping to an autobiographical memory. Bach argues that simulation tends to happen when the general heuristic has not yet been formed, and involves complex combinations of cases (see Bach, 2011 for specifics).

Because we do not attempt to model a complete Theory of Mind in this paper, we assume a simplified version of Bach's theory. Our model focuses on the learning aspect, so we assume that heuristic-like abstractions have not yet been formed. Thus, only concrete autobiographical memories are retrieved from long term memory. Generalizations are formed in working memory, which we propose as a mechanism by which schemas are learned.

SME and SAGE-WM

The Structure Mapping Engine (SME; Forbus et al., 2016) implements the analogical mapping process of SMT. SME compares a base and target case, both represented in predicate calculus, and computes one or more mappings that align statements and entities. Each mapped expression receives an initial score, which is propagated to its children. Thus, highly

nested expressions have high scores. The score of a mapping is the sum of the scores of its constituents. Thus, mappings between cases that have high structural similarity receive higher scores. Mappings also include candidate inferences that project missing information from one case to the other.

In this model, we deliberately do not model retrieval from long-term memory, to avoid the cost of providing enough distractors to make this challenging, and instead assume that retrieval finds reasonable autobiographical memories. However, we have proposed (Kandaswamy et al., 2014) that analogical generalization also occurs in working memory, what we call *interim generalizations*. The SAGE-WM model[1] keeps a list of generalizations and recent examples. Given a new example it uses SME to compute a score between the probe and each generalization in turn, ordered by recency. If the score is over a pre-determined threshold, the probe is assimilated into the generalization. If no generalization is above threshold, the new example is compared to each outlier in turn using SME, again ordered by recency. If any mapping is above threshold, a new generalization is formed. Otherwise the probe becomes a new ungeneralized example.

Learning Theory of Mind

Several studies have shown that Theory of Mind can be acquired in part using experimental interventions (e.g. Lohman & Tomasello, 2003; Hale & Tager-Flusberg, 2003; see Hofmann et al., 2016 for meta-analysis). However, most of these studies involve extended training. On the other hand, there is evidence that ToM can be acquired much more quickly when training examples are highly structurally alignable. In particular, a study by Hoyos et al. (2015) showed that structurally alignable unexpected contents-style stories can improve children's performance on false belief tasks, given just three training examples. In this paper, we examine and model the results of this experiment.

Modeling Task

In the Hoyos et al. (2015) study, children were first given a false belief pre-test containing one unexpected contents task (UC), one verbal false belief task (VFB), and one unexpected location task (UL). In the UC task, a container (e.g. a cookie box) is shown to have unexpected contents (e.g. grass) and participants are asked to predict what someone who has never seen inside would think the container contains. In VFB participants are told another child holds a false belief (that they think an item is somewhere it is not) and asked to predict where the child will look for the item. Finally, in UL, participants are told a story where one child places an object in a location and leaves the room. Another child then moves the object, and the participants are asked to predict where the first child will look for the object when they return.

Those who passed all three tests were excluded from the study. The remaining children were split into two groups: high alignment and low alignment. Both groups were

[1] Sequential Analogical Generalization Engine, Working Memory

presented with three stories in the style of an UC task, in a repetition-break pattern: the main character in the first two stories held a true belief (e.g. she thought that there was cereal in a cereal box, and there really was cereal inside), while the character in the last held a false belief (e.g. she thought there were crayons in the crayon box, but there were really rocks). The difference was that the stories heard by children in the high alignment condition were very similar, in terms of both structure and linguistic content. The stories heard by children in the low alignment condition, on the other hand, differed on both counts. Following training, all children were tested on the same three tasks (UC, VFB, UL) as before.

Hoyos et al. found that children in both conditions made significant gains from pre- to post-test. Importantly, they found that the children in the high alignment condition made significantly higher gains than those in the low alignment condition. Hoyos et al. concluded that structural alignment aids false belief understanding. Furthermore, they, like Bach (2011, 2014) postulated that analogical comparison is "instrumental in children's understanding of mental states and their relation to the factual world." In this paper, we propose a mechanism for how structural alignment during learning can aid in false belief understanding and forming a complete Analogical Theory of Mind.

Learning Analogical Theory of Mind

The mean performance increase by children in the high alignment group was 0.75 out of 3 possible, with significant gains made in all three of the false belief tests. Yet few children learned more than one. On the other hand, children in the low alignment condition made an average of 0.23 gains. Only gains in the UC task were significant. Since all of the training examples were variants of UC, it is not surprising that this was the easiest task to learn. However, learning ToM requires the ability to transfer to other tasks, as was the case with children in the high alignment condition. The process of making gains in UL and VFB tasks must, then, be different than the process of only gaining UC.

We argue that analogical comparison in working memory alone leads to gains in the UC. That is, immediate recall of the training examples themselves is sufficient to cause gains. In contrast, a generalization between a training example and an autobiographical memory retrieved from long term memory leads to transfer to the other two tasks, VFB and UL. The violation of expectation generated during training causes the child to probe long term memory for a case of similar surprise. What exactly they find surprising about the training—that something other than what they expected was inside the box, that the character in the story was incorrect in her guess, or something else—affects the case that is retrieved from long term memory. This in turn affects which of UL and VFB the child is able to answer.

A Computational Model

Our model, like Bach's theory, is based in SMT, using SAGE-WM for reasoning and learning.

Our Model

A simplified English version of each training and testing example from Hoyos et al. (2015) was semi-automatically encoded using a natural language understanding system (EA NLU; Tomai & Forbus, 2009). Although syntax was simplified, overall structure and word choices were consistent with the original stories. Figure 2 shows a partial representation of a true belief story. Events are represented in the neo-Davidsonian style: a reified event with role relations connecting it to other constituents. The conjunction of statements about an event participates in causal relations. In English, Figure 2 states that because it is not the case that there is a seeing event in the box by Kim, Kim thinks that there is a containment event wherein the box contains cereal.

During training, the appropriate examples were passed into SAGE-WM in the order that the children in the corresponding condition saw them (true belief, true belief, false belief). The threshold for whether or not a probe was generalized was set to 0.01. If the incoming example matched to an example already in working memory with a score greater than 0.01, the model asked whether the match was correct. This corresponds to feedback in the Hoyos et al. (2015) experiment. When told it was correct, the model assimilated the examples into a generalization. Its behavior when told it was incorrect, on the other hand, depended on its calculation of surprise. Surprise occurs when the model encounters an incorrect match whose score is the same order of magnitude as the previous correct match. We propose that this comes out of the repetition break structure of the story order (Hoyos et al., 2015; Loewenstein & Heath, 2009): the high similarity to the interim generalization leads to a strong expectation of sameness, and the violation leads to a search for re-categorization. When surprised, the model probes long term memory for an alternative case to align with.

Figure 3 gives a visual representation of our model. In the high alignment condition (a), the first true belief story is stored in working memory. The second true belief story is then matched to the first, and an interim generalization is formed. When the false belief story comes in, it too matches to the generalization. Due to violated expectations, long term

```
(causes-Underspecified
 (not
  (and
    (inside-UnderspecifiedRegion see85118 box1)
    (perceivedThings see85118
      (InsideOfSpaceRegionFn box1))
    (isa see85118 VisualPerception)
    (doneBy see85118 kim)))
 (opinions kim
  (and
    (containedObject contain84430 cereal84499)
    (containingObject contain84430 box1)
    (isa cereal84499 BreakfastCereal)
    (isa contain84430 ContainingSomething)))))
```

Fig. 2: A partial representation of a true belief story. This statement represents the phrase "Kim thinks that the box contains cereal because Kim has never seen inside the box".

memory (dotted line) is probed. Long term memory is a collection of generalized and specific cases that represent memories formed over time. If a case is retrieved, an interim generalization between the match and the false belief case is created and stored in working memory (b).

In the low alignment condition (c), on the other hand, no generalization is formed between the two true belief cases. This leads to them being stored as separate cases in working memory. When the false belief case comes in, it matches to the first true belief case, but no element of surprise is present when the model is corrected. For this reason, long term memory is never probed, and working memory consists of only the three training examples (d). The contents of working memory during testing predict the questions that the child is able to answer.

Testing proceeded as follows: cases were again encoded semi-automatically using EA NLU. These cases were given to the model which retrieved the most similar case from working memory and generated candidate inferences by analogy. The candidate inferences correspond to what the model predicts is missing from the test cases (e.g. what the agents will do). These candidate inferences were manually inspected to determine whether any could result in correctly answering the test questions.
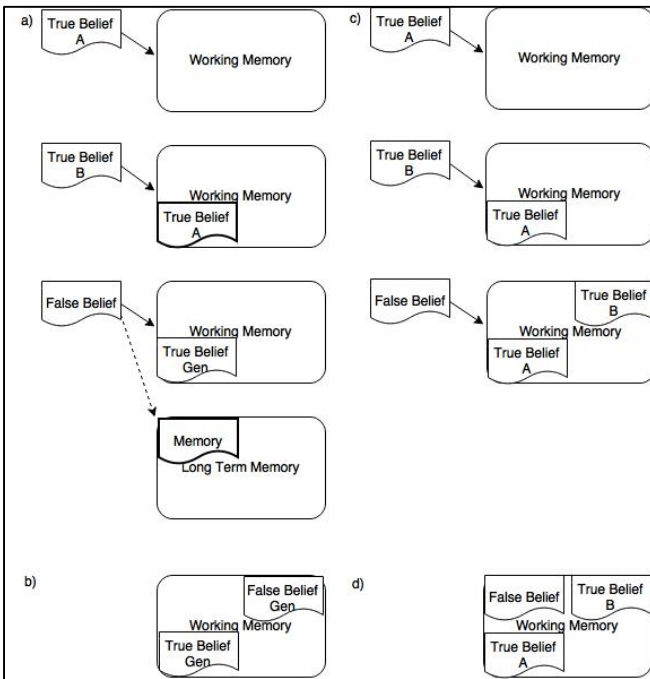

Fig. 3: A visual representation of our model of training in the Hoyos et al. study. (a) shows training in the high alignment condition. (b) is a representation of working memory after high alignment training. (c) and (d) show low alignment training and the consequent working memory, respectively. Cases that are structural matches to the probe are bold.

Results

The model behaved as predicted. In the high alignment condition, the model generalized the true belief cases with a normalized match score of 0.075. It then matched the false belief to the generalization with a score of 0.066, which corresponds to the child incorrectly predicting that the character in the story knows what is in the box. The model was then informed that this match was incorrect. Because the similarity scores it had encountered were within the same order of magnitude, it searched long term memory for another match. It then retrieved one of two memory cases that matched with a normalized score of 0.083 or 0.066, and created an interim generalization between it and the false belief case. We used stories intended to approximate a memory a child might have (e.g. thinking that a magician put a ball inside of a hat, only to find the hat empty) to model what might plausibly be retrieved. Depending on the case retrieved, the model was then able to answer VFB or UL. Correctness was evaluated based on the candidate inferences generated from the best mapping between the test case and the contents of working memory. For example, to correctly answer "Where is Nora going to look for her ball?" (UL) the mapping must produce a candidate inference stating that there might be a looking event, in which Nora looks for her ball in the appropriate location.

In the low alignment condition, on the other hand, the second true belief case matched to the first with a very low similarity score of 0.0014, well below threshold. For this reason, the model did not form a generalization between them. When the false belief case was compared, it had a match score of 0.066 with the first true belief case. Similar to the high alignment condition, the model was informed that this was not a correct match.

Because the previous match score was of a different order of magnitude, the model did not look into long term memory, and instead stored the false belief case alongside the two true belief cases. When the UC case came in, the false belief case was retrieved. The mapping generated a candidate inference that would allow the model to properly answer "What does she think is in the box?" This candidate inference stated that not having looked inside the cookie box would cause the agent to believe that it contained something analogous to crayons in the crayon box from the training example. That is, cookies.

Note that this retrieval is due to recency in working memory: the UC test case lacks the explanation present in the training cases about why a person holds a certain belief (e.g. "Kim thinks that the cereal box contains cereal because Kim has never looked inside the box."), so the first true belief case had the same match score. If that case had been retrieved, the model would not have been able to answer UC correctly.

Discussion

Our model gives one explanation for the results of the Theory of Mind training study presented in Hoyos et al. (2015). It also suggests that an important step in ToM development is generalizing belief-state cases in long term memory. In the

training studies, understanding that the training cases can, and indeed should, be assimilated to long term memory with belief-state interpretation cases is crucial. In other words, children may be accumulating experiences that require reasoning about belief states in long term memory, but these memories remain inert until a surprising event—such the one experienced by the high alignment participants in the Hoyos et al. study—stimulates their retrieval and begins the process of creating schemas that can be used in future ToM reasoning. This predicts that children in the high alignment condition of Hoyos et al. (2015) are more likely to retain what they have learned than the children in the low alignment condition: the children in the high alignment condition were more likely to access those experiences from long term memory and form a generalization with them.

In addition, our model predicts that reversing the order of training examples would cause children in both conditions to fail. In the low alignment case, when the most recent training example is retrieved, children would match the UC task to a true belief scenario, and answer incorrectly. Children in the high alignment case would similarly fall back on retrieval of the most recent case, as they would not experience the surprise caused by the repetition break structure.

Previous studies (e.g. Hale & Tager-Flusberg, 2003; Lohmann & Tomasello, 2003) have suggested that experience plays a role in ToM development. Our model provides a concrete explanation for how these experiences might lead to ToM and provides further suggestions for human subject experiments.

Related Work

Theories of Theory of Mind

Here, we summarize the best-known ToM theories.

Theory Theory One of the most popular takes on ToM is Theory Theory, which views the child as a scientist with regard to interpreting other people's mental states (e.g. Gopnik & Wellman, 1994). The child begins with a naïve theory about others, sometimes referred to as a folk psychology, which she modifies and adapts as evidence that supports or refutes the theory is observed. The theory gradually develops from only understanding desire states, to belief states, to how belief and desire states influence each other and behavior (Bartsch & Wellman, 1995).

Simulation Theory Under the Simulation Theory view, a child mentally simulates events in order to predict others' actions and beliefs (Goldman, 2006), and develops by improvement in simulation abilities (Flavell, 2004). Criticisms of Simulation Theory include that errors made by both children and adults are not consistent with those predicted by Simulation Theory accounts (Saxe, 2005) and that simulation is not sufficient for describing observed developmental patterns (Perner & Howes, 1992).

Modular Theories Another common account is that ToM can be explained as a single cognitive module. Scholl and Leslie (1999) list six characteristics of modules: they are domain-specific, their behavior is, at least in part non-

voluntary, their processing is fast, their outputs are shallow and highly constrained, they are often located in a particular region of the brain, and their processes may be impaired—and selectively impaired—by neural damage. Importantly, according to Scholl and Leslie, modularity theories "intend to capture only the origin of the basic ToM abilities" (1999). In this sense, modularity theories do not necessarily compete with other theories of ToM discussed here.

Hybrid Theories Several hybrid theories have been proposed to bridge the gap between Theory Theory and Simulation Theory. Some, which Bach (2011) calls *divided-hybrid models*, alternately assign aspects of Theory of Mind to simulation or theorizing, depending on which is better supported by empirical data (e.g. Heal, 1996). This approach, as Bach notes, avoids discussion of acquisition. It is unclear how a child learns to use simulation for some tasks and theory for others, and how simulation and theory develop concurrently. Other hybrid theories, which Bach (2011) calls *dynamic-hybrid models*, focus on continued development. Bach's model falls under this category. Like other dynamic-hybrid theories, Bach's allows for development and changes to ToM not only throughout childhood, but into adulthood. This includes switching between theorizing and simulating to complete the same tasks at different points in development. As psychologists continue to find evidence of ToM shifts throughout adulthood (e.g. Hess, 2006), dynamic-hybrid theories become more and more plausible.

Computation Models of Theory of Mind

Hiatt and Trafton (2010) implemented a model of Theory of Mind using the ACT-R cognitive architecture (Anderson, 2007) that learned to perform the Sally-Ann task. It extracted facts out of the scenario and was asked several false belief questions about what it saw. It was rewarded for answering correctly and punished for answering incorrectly, leading it over time to inhibit true belief responses, producing a learning curve consistent with developmental data. However, unlike our model, the training they used did not follow from an empirical training study. We note that the children in the Hoyos et al. (2015) study were able to learn aspects of false belief after seeing just three examples, only one of which actually was a false belief situation.

Goodman et al. (2006) modeled ToM via two Bayesian networks that respectively represent a naïve and expert theory in a Theory Theory account. They propose the models as competing hypotheses in the Sally-Anne task, and show how, during training, the expert theory becomes preferred over the naïve theory. The need to hand-code both theories in the system's starting endowment makes it more of a computational level model (Marr, 1982), whereas we provide a process-level model of learning. Furthermore, our model is consistent with the evidence from the training study presented by Goodman et al. (2006), which shows that surprise can improve children's ToM performance.

Future Directions

Our results provide evidence that structure-mapping is indeed a plausible process-level mechanism (Marr, 1982) for ToM and how it is learned. As such, our future work will look toward developing a complete computational Theory of Mind, including both the theory and simulation aspects of Bach's theory, using SAGE as the underlying mechanism.

Acknowledgments

References

Anderson, J. R. (2007). Using brain imaging to guide the development of a cognitive architecture. *Integrated models of cognitive systems*, 49-62.

Bach, T. (2011). Structure-mapping: Directions from simulation to theory. *Philosophical Psychology*, 24(1), 23-51.

Bach, T. (2014). A Unified Account of General Learning Mechanisms and Theory- of- Mind Development. *Mind & Language*, 29(3), 351-381.

Bartsch, K., & Wellman, H. M. (1995*). Children talk about the mind*. Oxford university press.

Chang, M.D. and Forbus, K.D. (2015). Towards Interpretation Strategies for Multimodal Instructional Analogies. *Proceedings of the 28th International Workshop on Qualitative Reasoning (QR2015)*. Minneapolis, MN.

Flavell, J. H. (2004). Theory-of-mind development: Retrospect and prospect. *Merrill-Palmer Quarterly*, 50(3), 274-290.

Forbus, K. D., Ferguson, R. W., Lovett, A., and Gentner, D. (2016). Extending SME to handle large-scale cognitive modeling. *Cognitive Science*, DOI: 10.1111/cogs.12377, pp 1-50.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155–170.

Goldman, A. I. (1992). In defense of the simulation theory. *Mind & Language*, 7(1- 2), 104-119.

Goldman, A. I. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. Oxford University Press.

Goodman, N. D., Baker, C. L., Bonawitz, E. B., Mansinghka, V. K., Gopnik, A., Wellman, H., ... & Tenenbaum, J. B. (2006). Intuitive theories of mind: A rational approach to false belief. *In Proceedings of the twenty-eighth annual conference of the cognitive science society* (pp. 1382-1387).

Gopnik, A. & Wellman H. (1992). Why the child's theory of mind really is a theory. *Mind and Language*, 7, 145-171.

Gopnik, A., & Wellman, H. (1994). The theory theory. In L. Hirschfeld & S. Gelman (Eds.), *Domain specificity in culture and cognition*. NY: Cambridge University Press.

Hale, C. M., & Tager- Flusberg, H. (2003). The influence of language on theory of mind: A training study. *Developmental science*, 6(3), 346-359.

Heal, J. (1996) Simulation, Theory, and Content. In Carruthers, P. & Smith, K. (Eds.), *Theories of theories of mind*. Cambridge: Cambridge University Press.

Hess, T. (2006). Adaptive aspects of social-cognitive functioning in adulthood: Age-related goal and knowledge influences. *Social Cognition*, 24(3), 279-309.

Hiatt, L. M., & Trafton, J. G. (2010, August). A cognitive model of theory of mind. In *Proceedings of the 10th International Conference on Cognitive Modeling* (pp. 91-96).

Hofmann, S. G., Doan, S. N., Sprung, M., Wilson, A., Ebesutani, C., Andrews, L. A., ... & Harris, P. L. (2016). Training children's theory-of-mind: A meta-analysis of controlled studies. *Cognition*, 150, 200-212.

Hoyos, C., Horton, W., & Gentner, D. (2015). Analogical comparison aids false belief understanding in preschoolers. In *CogSci*.

Kandaswamy, S., Forbus, K., and Gentner, D. (2014). Modeling Learning via Progressive Alignment using Interim Generalizations. *Proceedings of the Cognitive Science Society*.

Kuehne, S., Forbus, K., Gentner, D. and Quinn, B. (2000). SEQL: Category learning as progressive abstraction using structure-mapping. *Proceedings of CogSci 2000*, August.

Loewenstein, J., & Heath, C. (2009). The Repetition-Break plot structure: A cognitive influence on selection in the marketplace of ideas. *Cognitive Science*, 33, 1-19.

Lohmann, H., & Tomasello, M. (2003). The role of language in the development of false belief understanding. *Child Development*, 74(4), 1130 –1144.

Marr, D. (1982). *Vision*. Freeman Publishers.

Perner, J., & Howes, D. (1992). 'He Thinks He Knows': And More Developmental Evidence Against the Simulation (Role Taking) Theory. *Mind & Language*, 7(1- 2), 72-86.

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

Scholl, B. J., & Leslie, A. M. (1999). Modularity, development and 'theory of mind'. *Mind & Language*, 14(1), 131-153.

Tomai, E. and Forbus, K. (2009). EA NLU: Practical Language Understanding for Cognitive Modeling. *Proceedings of the 22nd International Florida Artificial Intelligence Research Society Conference*. Sanibel Island, Florida.