

Toward Logic-Based Cognitively Robust Synthetic Characters in Digital Environments¹

Selmer BRINGSJORD,
selmer@rpi.edu

Andrew SHILLIDAY, Joshua TAYLOR, Dan WERNER, Micah CLARK,
{shilla, tayloj, werned, clarkm5}@rpi.edu

Ed CHARPENTIER and Alexander BRINGSJORD
{charpe, bringa}@rpi.edu

Rensselaer AI & Reasoning (RAIR) Lab
Department of Cognitive Science
Department of Computer Science
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180

Abstract. With respect to genuine cognitive faculties, present synthetic characters inhabiting online virtual worlds are, to say the least, completely impaired. Current methods aimed at the creation of “immersive” virtual worlds only avatars and NPCs the *illusion* of mentality and, as such, will ultimately fail. Like behaviorism, this doomed approach focuses only on the inputs and outputs of virtual characters and ignores the rich mental structures that are essential for any truly realistic social environment. While this “deceptive” tactic may be suitable so long as a human is in the driver’s seat compensating for the mental deficit, truly convincing autonomous synthetic characters must possess genuine mental states, which can only result from a formal *theory of mind*. We report here on our attempt to invent part of such a theory, one that will enable artificial agents to have and reason about the beliefs of others, resulting in characters that can predict and manipulate the behavior of even *human* players. Furthermore, we present the “embodiment” of our recent successes: Eddie, a four year old child in *Second Life* who can reason about his own beliefs to draw conclusions in a manner that matches human children his age.

Keywords. virtual characters, cognition, logic, theory of mind

¹We are grateful for funding from IARPA and IBM; this funding has partially supported the R&D described herein. We are also indebted to Konstantine Arkoudas for many contributions. Finally, thanks are due to Mike Schoelles for his work on SimBorgs that provided a starting point for automating conversation in *Second Life*.

1. The Problem

Your avatar in the current version of a massively multi-player online (MMO) virtual world (e.g., *Second Life*) is directly tethered to your key strokes, and is nothing more than an automaton controlled directly by what your fingers do. This is great as far as it goes, but the fact remains that your avatar is completely bereft of the cognitive things that make you you, and is in fact devoid of even computational correlates of the things that make you you. He/she doesn't speak or move autonomously, doesn't have any memories, doesn't have any beliefs (none at all, and therefore none of precisely the kind that are distinctive of persons, e.g., beliefs about the beliefs others have about your own beliefs²), doesn't know anything, and therefore certainly can't hold in mind a model of other creatures he/she encounters in such a virtual world.

In laying a foundation for our initial research, we sought to make this problem more concrete: We not only studied the state of the art in AI, but had our laboratory, if you will, enter into it. We thus know that while Rensselaer graduate Joshua Taylor has an avatar in *Second Life*, and also in Nintendo's Wii system (see Figure 1), the fact of the matter is that the real Joshua Taylor has all sorts of beliefs, intentions, goals, and desires that the avatar does not.




Figure 1. Joshua Taylor, Computer Science MS Holder, "Avatared" in Nintendo's Wii System (left) and *Second Life* (right)

The same fundamental problem arises when we are talking not about avatars, but about NPCs (non-player characters). They are just as cognitively empty, and as a result can invariably be spotted in digital environments as mere shells.

If space permitted, we could show, in detail, that all synthetic characters, to this point in time, whether avatars or NPCs, are indeed primitive. (For some examples, see Figure 2.) Even book-length treatments of how to build synthetic characters are impoverished relative to what we are engineering. For example, Figure 3, taken from [2], models synthetic characters using finite state automata. In such models, there is no knowledge and belief, no reasoning, no declarative memories, and no linguistic capacity.

2. Behaviorism Isn't the Answer

In the end, given the stone cold emptiness of avatars and NPCs, simple engineering tricks to make these digital creatures seem as if they have genuine mentality isn't going to work.

²One of the traditional marks of personhood has long been held to be that they can have such so-called *third-order* beliefs. See [1].



Figure 2. Sample Synthetic Characters. Worst to best, in our eyes: Top-left, The Legend of Zelda; SC spits text upon entering room. Top-right, Chrono Trigger; tree-branching conversations. Middle-left, Might & Magic VI (Shopkeepers). Middle-right, Superfly Johnson from Daikatana; behavior scripting, attempts to follow player and act as a sidekick. Bottom-left, Galatea – Interactive Fiction award winner for Best NPC of 2000 (text-based). Bottom-right, Sims 2. Thanks are due to Marc Destefano for these examples and the snapshots.

One can try, through surface-level tricks (such as the “chirping” in *The Sims*, used to make it *seem* as if characters are genuinely conversing), to make the observable behavior of synthetic characters such that these characters *seem* to have knowledge, beliefs, memories, and so on, even when they don’t have these things, but this strategy, ultimately, will always fail. Any approach that treats avatars as empty shells, as puppets, will fail for the same reasons that behaviorism failed: if focuses on inputs (stimuli) and outputs (responses) and ignores everything in between. A rich mental life, however, can only be understood, predicted, and explained (indeed, it can only be *had*) on the assumption that there are rich mental structures mediating inputs and outputs. That was the central tenet of the “cognitive revolution,” and unless we want a form of AI as bankrupt as behaviorism, we must subscribe to it when building synthetic characters. We must thus build synthetic characters with rich mental structures capable of representing (at least computational correlates of) beliefs, knowledge, perceptual input, and so on. And we further argue that we need to do that in a rigorous way, one that specifically capitalizes on the logical properties of mental states, as investigated by logicians, technical philosophers, theoretical cognitive scientists, and so on (rather than by building, say, neural networks and other numerical, non-declarative systems). A very deep kind of *logicist* engineering is thus, at least by our lights, needed. For us, the engineering must be in the mold of logic-based AI [3,4,5,6,7] and computational cognitive modeling [8]. This logic-based work stands in what we see as interesting, stark contrast to brain-based approaches in AI and cognitive science (e.g., to [9]).³

³While we don’t discuss the issue herein, it may be that hybrid approaches bringing together top-down logic-based techniques with bottom-up neuroscience-oriented techniques are worth pursuing. Sun’s [10] is directly relevant here.




Figure 3. Impoverished FSA-Based Formalism for Synthetic Characters (from Champandard 2003).

3. Our Long-Term Objective

In order to lay a foundation for engineering cognitively robust synthetic characters in massively multiplayer online (MMO) virtual worlds (e.g., *Second Life*) and other digital environments, we are in the process of codifying, to an unprecedented degree, the principles of “common-sense psychology” (CSP) in the form of an implementable logico-mathematical theory. This theory must include rigorous, declarative definitions of all the concepts central to a theory of mind — concepts like lying, betrayal, evil, and so on. The implementation of this theory will be one that artificial agents can deploy in order to understand and predict psychological aspects of the behavior of other artificial agents, and in order to be genuine stand-ins for human beings. No such understanding or prediction is possible today.

Our methodology can be understood along the lines of Lewis’s prescription:

Collect all the platitudes regarding the causal relations of mental states, sensory stimuli, and motor responses. Add also all the platitudes to the effect that one mental state falls under another... Perhaps there are platitudes of other forms as well. Include only the platitudes which are common knowledge amongst us: everyone knows them, everyone knows that everyone else knows them, and so on. [11, p. 256]

Of course, Lewis here is concerned with the construction of a *complete* theory of CSP, one that would be capable of explaining virtually all aspects of human behavior. It is

We also leave aside the economic side of the present work. Surmounting the impoverished nature of synthetic characters, making them, as we say, *cognitively robust*, would have great economic value. Millions of people are willing today to spend hard-earned money to merely control “puppets” in MMO virtual worlds. Billions would, we suspect, spend such money to play in virtual worlds populated by *bona fide* digital psyches. A small but important part of our efforts includes study of the economic and business dimension of MMO virtual worlds and synthetic characters, because we see our R&D as being bound up with the realities of the marketplace. Our study in this area, led by Alexander Bringsjord, is left aside in the present document.

questionable whether the construction of such a theory is a practical possibility, and whether it could be carried out in the absence of a complete physical theory. Even if such a theory is practically possible, it is doubtful whether practical reasoning could be performed on its basis. Therefore, at least initially, we are building a micromodel of CSP.

4. The Theory of Mind Theory

The ability to ascribe mental states to others and to reason about such mental states is indispensable for social communication. All social transactions — from lying and detecting lying, to engaging in commerce and negotiating, to making jokes and empathizing with other people’s pain or joy — require at least a rudimentary mastery of CSP. Today’s synthetic characters lack such a facility, and hence essentially suffer from autism, which is literally the disease afflicting that part of the brain responsible for reasoning about the mental states of others. An inability to attribute mental states to other characters is “tantamount to not differentiating between the world of objects (with physical states) and the world of persons (with mental states)” [12, p. 65]. A grasp of CSP is particularly important for agents who are trying to manipulate the behavior of other agents. The ability to manipulate through lying, for example, requires sophisticated reasoning about the beliefs of other agents. More generally, social interaction has been described as a form of “social chess” [13], whereby one agent, for instance, “may wish by his own behavior to change the behavior of another; but since the social animal is himself reactive and intelligent, the interaction soon becomes a two-way argument where each ‘player’ must be ready to change his tactics — and maybe his goals — as the game proceeds.” The principles and techniques that neurobiologically normal humans deploy in order to understand, predict, and manipulate the behavior of other humans are collectively referred to as a *theory of mind* [14]. What we are now starting to do is to engineer part of the theory that artificial agents could use to understand, predict, and manipulate the behavior of other agents, and in order to be veridical avatars for human beings, or to be autonomous intellects in their own right.

One quick way to make our project clearer is to briefly turn specifically to lying, a speck in the overall theory we seek, but nonetheless informative to consider.⁴ One well-known definition of lying within analytic philosophy is

L lies to $D =_{def} \exists p$ such that

1. either L believes that p is not true or L believes that p is false; and
2. L asserts p to D [16, 152].

where

L asserts p to $D =_{def}$ L states p to D and does so under conditions which, he believes, justify D in believing that he, L , not only accepts p , but also intends to contribute causally to D ’s believing that he, L , accepts p [16, 152].

As can be clearly seen by inspection, even if these definitions are only approximately correct, any human D who believes that he or she is being lied to, must have beliefs about the beliefs the liar has about her (i.e., D ’s) beliefs, and at the heart of the matter is propositional content.

⁴Clark has extended, formalized, and implemented this definition in what he calls “the lying machine[15].”

5. Near-Term Goal: The Crucial Pair of Demos

In the near term, we are working toward progressively more elaborate versions of two key demonstrations (both of which will be given at AGI2008). The first (the first two versions of which we have already achieved; see below) will subject our system to the litmus test of theories of mind: passing a false-belief test. Experiments with false beliefs were first carried out by Wimmer and Perner[17]. In a typical scenario, a child is presented with a story in which a character A places an object (e.g., a marble, teddy bear, cookie, etc.) in a certain location l_1 , say in a particular kitchen cabinet. Then A leaves, and during his absence a character B removes the object from its original location l_1 and puts it in a different location l_2 (say, a different cabinet). The child is then asked to predict where A will look for the object when he gets back — the right answer, of course, being the original location, l_1 . (Very young children ($\approx \leq 4$) don't produce this answer; they do not have a theory of the mind of others.) We are putting artificial agents to a similar test. To pass the test successfully, the agent will be required to make the correct prediction, by deploying our CSP model.

Though we have much work to do, some primitive versions of this first demo have been implemented in connection with *Second Life*. A snapshot from one of our demos is shown in Figure 4. To take a look at a couple of the full videos, visit the link at the bottom of http://www.cogsci.rpi.edu/research/rair/asc_rca.




Figure 4. Snapshot of Demo in SL.

Our demo in *Second Life* is achieved with help from an automated theorem prover, Snark, coupled with procedures for generating and executing AppleScript (a scripting language from Apple Inc.). When Snark discovers a proof that agent s will perform action a , a procedure written in Common Lisp produces the code that, when executed, simulates keystrokes in *Second Life*, thereby enabling control of an avatar. Figure 5 shows the proof associated with “Ed’s” immature response in the false-belief task carried out in *Second Life*. In this case, s is the autonomous synthetic character Eddie, and the action a is the selection of the container that the bear has been moved to (so this is an example of failure in the task).

Our second demonstration is in the area of natural language processing; here again an early version is available at the url given above. (See Figure 6 for a snapshot from

```

SNARK-USER 14 >
(in-immature-scenario
 (prove '(t-retrieve subject
          teddybear
          ?z)
        :answer '(looks-in ?c)))

(Refutation
 (Row 1
  (or (not (person ?x)) (not (object ?
y)) (not (container ?z)) (not (in ?y ?
z)) (bel-in ?x ?y ?z))
  assertion)
 (Row 2
  (or (not (person ?x))
      (not (container ?y))
      (not (object ?z))
      (not (w-retrieve ?x ?z))
      (not (bel-in ?x ?z ?y)))
      (t-retrieve ?x ?z ?y))
  assertion)
 (Row 4
  (person subject)
  assertion)
 (Row 6
  (container c2)
  assertion)
 (Row 7
  (object teddybear)
  assertion)
 (Row 8
  (in teddybear c2)
  assertion)
 (Row 9
  (w-retrieve subject teddybear)
  assertion)
 (Row 10
  (not (t-retrieve subject teddybear ?
x))
  negated_conjecture
  Answer (looks-in ?x))
 (Row 11
  (or (not (person ?x)) (bel-in ?x
teddybear c2))
  (rewrite (resolve 1 8) 6 7))
 (Row 25
  (bel-in subject teddybear c2)
  (resolve 11 4))
 (Row 28
  (t-retrieve subject teddybear c2)
  (rewrite (resolve 2 25) 9 7 6 4))
 (Row 30
  false
  (resolve 10 28)
  Answer (looks-in c2)))
:PROOF-FOUND
SNARK-USER 15 > (answer t)
LOOKS-IN C2)

```

Figure 5. Snapshot of False-Belief Task Demo in *Second Life*.

the video.) The goal in the second demonstration is to engineer a virtual character in an MMO virtual world capable of carrying out a conversation based upon certain knowledge and belief. In the demo, this character (BinDistrib) is in conversation with an avatar (CrispyNoodle). When the avatar (CrispyNoodle) speaks, the text in question is read by BinDistrib. For our work here we are leveraging our prior work in the machine reading area [18], an overview of which we now provide.



Figure 6. Snapshot of Conversational Demo in *Second Life*.

5.1. Machine Reading of Text in *Second Life*

In our approach, text to be read is expressed in logically controlled English, which is translated into multi-sorted logic (MSL). The resulting information in MSL, in conjunction with other background knowledge expressed in MSL, is automatically reasoned over in proof-theoretic and model-based fashion. Our approach to machine reading is a three-phase one:

Phase 1 English texts are rephrased in logically controlled English — i.e., a proper subset of full English that can be unambiguously translated into a formal logic. In the past, we have

made use of Attempto Controlled English (ACE) [19,20], a logically controlled English with a fixed, definite clause grammar and a user-defined vocabulary.⁵ However, we are now moving to CELT [22,23,24].

Phase 2 Discourse representation structures (DRSs) are automatically generated from the controlled English. DRSs are a syntactic variant of first-order logic for the resolution of unbounded anaphora. Their use in the interpretation of text is a central element of discourse representation theory [25,26].

Phase 3 The DRSs are automatically translated into MSL. (We have a number of translators in our lab for going from MSL to straight first-order logic (FOL), using long-established theorems [27].) As a DRS is equivalent to a quantified first-order formula, the translation to FOL is not conceptually difficult. Algorithms for performing such translations are provided by Blackburn [28], among others.

We have used this three-phase process in developing machine reading capability in systems we have built with sponsorship from the DoD. In connection with Slate, one of these systems, the three phases can be viewed from a high-level perspective as shown in Figure 7.




Figure 7. Slate's Reading Process

As we have indicated, in prior logic-based NLP R&D, we have used Attempto Controlled English (ACE), but we are now moving to CELT. CELT is potentially more powerful than ACE because it leverages the Suggested Upper Merged Ontology (SUMO), which incorporates useful knowledge such as geometric axioms. In the current version of our demo, a sentence like

Micah puts the teddy bear in the box.

becomes in CELT's translation to FOL the following formula:

```
(exists
  (?box ?event ?teddy_bear)
  (and
    (instance Micah Human)
    (attribute ?box Box)
    (agent ?event Micah)
    (instance ?event Putting)
    (instance ?teddy_bear Artifact)
    (destination ?event ?box)
    (patient ?event ?teddy_bear)))
```

⁵Phase 1 is currently a manual operation, but techniques developed by Mollá & Schwitter [21] may allow for at least partial automation of this phase.

6. On Relevant Prior Work

It is important to distinguish our approach from other related research efforts:

1. We are not be doing research on BDI (belief-desire-intention) logics in the standard knowledge-representation tradition [29,30]. Most of that work has focused on codifying *normative* reasoning (e.g., about knowledge) in modal logics, with little attention paid to practical concerns such as computational efficiency. We are more interested in a *descriptive* model of reasoning about mental states, rather than a prescriptive model, and one that is computationally tractable.
2. We are not doing cognitive science. While our insights will be obviously coming from human CSP, we will not be particularly concerned with arriving at a cognitively plausible theory. Our aim is not to construct a computational theory which explains and predicts actual human behavior. Rather, our aim is to build artificial agents which are more interesting and useful by enabling them to ascribe mental states to other agents, to reason about such states, and to have, as avatars, states that are correlates to those experienced by corresponding humans. This might result in an interesting computational simulation of a fragment of human CSP, but our primary concern is engineering in the service of entertainment and gaming, not science. In particular, our effort should not be seen as an endorsement of the “theory-theory” position in the current debate between simulation theorists and folk-psychology theorists in cognitive science [31].

Some prior research and development in our own case provides a foundation for the proposed work. For example, with support from NSF, we designed and built synthetic characters (named Ralph and JR9000) that interacted with Bringsjord to teach AI In addition, the primitive synthetic character known simply as E, was demoed at the 2005 *GameOn!* conference. The associated paper won the “Best Paper” award: [32] In the style of the definition of lying provided above, E is based upon a fully declarative, and fully formal definition of evil.⁶

7. The Main Obstacle: Meaning Mathematized

It’s hard to miss the fact that we are sanguine about our attempt to engineer cognitively robust synthetic characters. In the sober light of AI’s failure to live up to predictions made by its founders, do we really maintain that we can meet with success? Do we not at least see some *seemingly* insuperable barriers out there ahead of us? Actually, we do see a serious hurdle ahead, one that we currently confess we do not have the means to overcome. We do not have the space to carefully describe this hurdle, which is quite technical, but we give here a synopsis of both the challenge, and our approach to meeting it.

The problem is that, to this point, formal logic and formal methods in computer science have provided inadequate mathematical frameworks for pinning down the *meaning* of the propositional attitudes central to our objectives. What does it mean for a person to believe, know, desire, fear, hope, etc. — where this meaning is provided in fully formal fashion that can be rendered in computational form? This question has not been answered. There have been *attempted* answers, but they are clearly inadequate, and logic-based AI is only fooling itself when it takes the current answers to be anywhere near ac-

⁶A copy of the paper can be found at <http://kryten.mm.rpi.edu/GameOnpaper.pdf>.

ceptable. For example, Bringsjord and Konstantine Arkoudas have long been convinced that the appropriation of possible-worlds semantics for supposedly specifying the meaning of *Knows* and *Believes* has been more a marriage of convenience than anything else. For *Necessarily* and *Possibly*, the two operators used in standard modal logic [33], this sort of semantics makes very good sense. For example, it seems eminently sensible to understand $\Box\phi$ (a formula in standard modal logic traditionally used to express that ϕ is logically necessary) to be saying that, in all possible worlds, or in all possible histories, ϕ holds. The standard correlate for knowing, though, is almost laughably forced. What intuitive sense does it make to say (e.g., with [34]) that an agent knows ϕ just in case in all possible worlds epistemically accessible for this agent, ϕ holds?

In the course of the project described herein we are going to need to invent a new logical system for the propositional attitudes at the heart of the robust characters we seek. Some of the first steps in this direction have been taken, and we now say a few words about our system.

Our approach provides a formal semantics for epistemic operators like *Believes* and *Knows* on the basis of *epistemic maps*. Such maps are a particular kind of graph in which the objects of such operators are propositions situated within a web of interconnected inference. Inference is allowed to be deductive, inductive, abductive, analogical, or (a mode that readers are unlikely to be familiar with) creative. Figure 8 shows a very simple epistemic map \mathcal{M} corresponding to a situation in which agent *A* believes, on the strength of a number of other, subsidiary beliefs (all of which, appearing within rectangles, are perceptually based), that Smith’s cat is paralyzed with fear up in an oak tree on Smith’s property. In this case, only one inference has taken place, an abductive one. Where ϕ represents in first-order logic that Smith’s cat is stymied in an oak tree, the formula $\psi = \text{Believes}_A\phi$ is true on this map (written $\mathcal{M} \models \psi$), but because the strength of ϕ is only at level 1 (“probable”) in the continuum from -4 to 4 of strength factors in our approach (a continuum derived directly from [35]), and this level doesn’t provide a sufficiently strong justification to move *A*’s belief into the category of knowledge, $\mathcal{M} \not\models \text{Knows}_A\phi$.

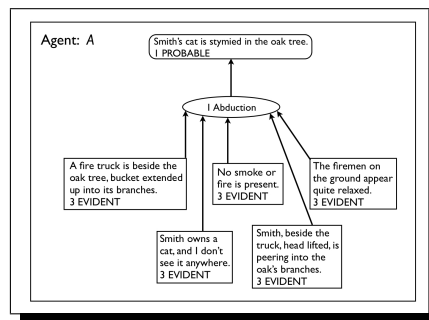


Figure 8. Simple Epistemic Map Re. Smith’s Cat

We conclude with a final point about our approach. As is well known, ordinary substitution into extensional contexts leads to error. For example, though agent *B* may know that John Le Carré is the author of *A Perfect Spy*, it doesn’t follow from this and the fact that Le Carré and David Cornwall are one and the same that *B* knows that David Cornwall wrote this novel. (After all, *B* may not have heard or seen the string ‘David

Cornwall' in his entire life.) Any acceptable epistemic logic must respect such blocked inference. In our approach, as shown in Figure 9, an agent *A* wouldn't believe that *B* knows that David Cornwall wrote *A Perfect Spy*, even if *A* believes that *B* believes that John Le Carré wrote this book. The mechanical reason for this is that epistemic maps can themselves contain epistemic maps, but, in general, propositions holding in a map \mathcal{M} that contains another map \mathcal{M}' ($\text{cal}\mathcal{M} \supset \mathcal{M}'$) do not get "injected" into \mathcal{M}' . (On the other hand, for all such pairs of epistemic maps, and all formulae δ , if $\mathcal{M}' \models \delta$, then $\mathcal{M} \models \phi$.) Put in terms of Figure 9, the maps shown there don't satisfy the proposition that *A* believes that *B* believes that Cornwall wrote the novel in question.




Figure 9. Opaque Context in Our Approach

References

- [1] S. Bringsjord, *Abortion: A Dialogue*. Indianapolis, IN: Hackett, 1997.
- [2] A. Champandard, *AI Game Development*. Berkeley, CA: New Riders, 2003.
- [3] S. Bringsjord, K. Arkoudas, and P. Bello, "Toward a general logicist methodology for engineering ethically correct robots," *IEEE Intelligent Systems*, vol. 21, no. 4, pp. 38–44, 2006.
- [4] S. Bringsjord and D. Ferrucci, *Artificial Intelligence and Literary Creativity: Inside the Mind of Brutus, a Storytelling Machine*. Mahwah, NJ: Lawrence Erlbaum, 2000.
- [5] N. Nilsson, "Logic and Artificial Intelligence," *Artificial Intelligence*, vol. 47, pp. 31–56, 1991.
- [6] S. Bringsjord and D. Ferrucci, "Logic and artificial intelligence: Divorced, still married, separated...?," *Minds and Machines*, vol. 8, pp. 273–308, 1998.
- [7] S. Bringsjord and D. Ferrucci, "Reply to Thayne and Glymour on logic and artificial intelligence," *Minds and Machines*, vol. 8, pp. 313–315, 1998.
- [8] S. Bringsjord, "Declarative/logic-based cognitive modeling," in *The Handbook of Computational Psychology* (R. Sun, ed.), Cambridge, UK: Cambridge University Press, forthcoming.
- [9] J. Hawkins, *On Intelligence*. New York, NY: Holt Paperbacks, 2005.
- [10] R. Sun, *Duality of the Mind*. Mahwah, NJ: Lawrence Erlbaum Associates, 2001.
- [11] D. Lewis, "Psychophysical and Theoretical Identifications," *Australasian Journal of Philosophy*, vol. 50, pp. 249–258, 1972.
- [12] U. Frith, "Cognitive explanations of autism," *Acta Paediatrica Supplement*, vol. 416, pp. 63–68, 1996.
- [13] N. Humphrey, *Consciousness regained*. Oxford University Press, 1984.
- [14] D. Premack and G. Woodruff, "Does the chimpanzee have a theory of mind?," *Behavioral and Brain Sciences*, vol. 4, pp. 515–526, 1978.
- [15] M. Clark, *Cognitive Illusions and the Lying Machine*. PhD thesis, Rensselaer Polytechnic Institute (RPI), 2008.

- [16] R. M. Chisholm and T. D. Feehan, "The Intent to Deceive," *Journal of Philosophy*, vol. 74, pp. 143–159, March 1977.
- [17] H. Wimmer and J. Perner, "Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception," *Cognition*, vol. 13, pp. 103–128, 1983.
- [18] S. Bringsjord, K. Arkoudas, M. Clark, A. Shilliday, J. Taylor, B. Schimanski, and Y. Yang, "Reporting on some logic-based machine reading research," in *Proceedings of the 2007 AAAI Spring Symposium: Machine Reading (SS-07-06)*, (Menlo Park, CA), pp. 23–28, AAAI Press, 2007.
- [19] N. E. Fuchs, U. Schwertel, and R. Schwitter, "Attempto Controlled English (ACE) Language Manual, Version 3.0," Tech. Rep. 99.03, Department of Computer Science, University of Zurich, Zurich, Switzerland, 1999.
- [20] S. Hoefler, "The Syntax of Attempto Controlled English: An Abstract Grammar for ACE 4.0," Tech. Rep. ifi-2004.03, Department of Informatics, University of Zurich, Zurich, Switzerland, 2004.
- [21] D. Mollá and R. Schwitter, "From Plain English to Controlled English," in *Proceedings of the 2001 Australasian Natural Language Processing Workshop*, (Macquarie University, Sydney, Australia), pp. 77–83, 2001.
- [22] A. Pease and J. Li, "Controlled english to logic translation," forthcoming.
- [23] A. Pease and C. Fellbaum, "Language to logic translation with phrasebank," in *Proceedings of the Second International WordNet Conference (GWC 2004)* (P. Sojka, K. Pala, P. Smrz, C. Fellbaum, and P. Vossen, eds.), (Masaryk University, Brno, Czech Republic), pp. 187–192, 2004.
- [24] A. Pease and C. Fellbaum, "An english to logic translator for ontology-based knowledge representation languages," in *Proceedings of the 2003 IEEE International Conference on Natural Language Processing and Knowledge Engineering*, (Beijing, China), pp. 777–783, 2003.
- [25] H. Kamp and U. Reyle, *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Springer, 1 ed., 1993.
- [26] H. Kamp and U. Reyle, "A Calculus for First Order Discourse Representation Structures," *Journal of Logic, Language and Information*, vol. 5, pp. 297–348, 1996.
- [27] M. Manzano, *Extensions of First Order Logic*. Cambridge, UK: Cambridge University Press, 1996.
- [28] P. Blackburn and J. Bos, "Working with Discourse Representation Theory: An Advanced Course in Computational Semantics." Forthcoming.
- [29] M. E. Bratman, *Intention, Plans, and Practical Reason*. CSLI Publications, 1999.
- [30] M. Wooldridge, *Reasoning About Rational Agents*. MIT Press, 2000.
- [31] M. Davies and T. Stone, eds., *Folk Psychology: The Theory of Mind Debate*. Blackwell Publishers, 1995.
- [32] S. Bringsjord, S. Khemlani, K. Arkoudas, C. McEvoy, M. Destefano, and M. Daigle, "Advanced synthetic characters, evil, and E," in *Game-On 2005, 6th International Conference on Intelligent Games and Simulation* (M. Al-Akaidi and A. E. Rhalibi, eds.), pp. 31–39, Ghent-Zwijnaarde, Belgium: European Simulation Society, 2005.
- [33] B. F. Chellas, *Modal Logic: An Introduction*. Cambridge, UK: Cambridge University Press, 1980.
- [34] R. Fagin, J. Halpern, Y. Moses, and M. Vardi, *Reasoning About Knowledge*. Cambridge, MA: MIT Press, 2004.
- [35] R. Chisholm, *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice-Hall, 1966.