

# Reporting On Some Logic-Based Machine Reading Research

Selmer Bringsjord, Konstantine Arkoudas

Micah Clark, Andrew Shilliday, Joshua Taylor, Bettina Schimanski, Yingrui Yang

Department of Cognitive Science

Department of Computer Science

Rensselaer AI & Reasoning (RAIR) Lab:

<http://www.cogsci.rpi.edu/research/rair/>

Troy NY 12180 USA

{selmer,clarkm5,shilla,tayloj,schimb,yangyri}@rpi.edu konstantine@alum.mit.edu

<http://www.rpi.edu/~brings>

## Introduction

Much sponsored research in our lab either falls under or intersects with machine reading. In this short paper we give an encapsulated presentation of some of the research in question, leaving aside, for the most part, the considerable detailed technical information that underlies our work.<sup>1</sup> Demonstrations of our technology will be provided at the symposium itself.

Our machine reading research can be viewed as falling under two categories, viz.,

Fast, Primitive Machine Reading in Real-World Systems.

Here we are interested in building into deployed software a capacity to read text expressed in English. The machine reading in question is primitive because the English is restricted: it's what we call *logically controlled English*. We report herein on how the Slate software system reads logically controlled English, and extracts knowledge from this English to be represented in multi-sorted logic (MSL), the chief native languages of Slate.

Machine Reading of Diagram-Infused Text. Here we are concerned with engineering systems that can read diagram-infused text. Such text, as opposed to text without diagrammatic or pictorial content, is by far the dominant form of text seen in academia, especially in technical areas — such as physics, chemistry, mathematics, computer science, astronomy, and so on, and also in high-stakes standardized testing, for example in the SAT. (The presence of diagrams in domains such as these was noted in (Friedland *et al.* 2004).) Our research in this area is based on a new theory of learning (so-called *poised-for learning by reading*, or just PFLbR), and a new theory of heterogenous reasoning called *mental metalogic*.

In the remainder of the paper, we briefly describe our work under these two categories. Finally, before we begin, we point out that, for better or worse, our approach is a thoroughly logicist one (Bringsjord & Fer-

rucci 1998a; Bringsjord & Ferrucci 1998b; Genesereth & Nilsson 1987; Nilsson 1991).<sup>2</sup>

## Primitive Machine Reading

### Slate, Briefly

Slate is a robust interactive reasoning system. It allows the human “pilot” to harness an ensemble of intelligent agents in order to construct, test, and express proofs and argumentation of various sorts. Slate is designed to empower professionals in the business of producing natural-style<sup>3</sup> argumentation — mathematicians, logicians, analysts, wargamers, designers and producers of standardized reasoning tests, and so on.

### Machine Reading by the Slate System

Slate can translate text expressed in logically controlled English into multi-sorted logic (MSL), build knowledge expressed in MSL, and reason over that knowledge in proof-theoretic and model-based fashion. It can do this both on its own, and under the guidance of a human user of the system. In light of this capability, we say that Slate, in a fixed and confessedly limited sense, can “read.” A conceptualization of the process by which

---

<sup>2</sup>We make use of strength factors, and abductive and inductive inference, but not of probabilistic or statistical formalisms.

<sup>3</sup>There are various ways to understand “natural” argumentation. For us, the hallmark of such argumentation is that it conforms to the kind of reasoning that humans produce, not the kind of inference generally preferred by automated reasoners. The latter are for the most part resolution-based, but resolution is well nigh impenetrable to humans, and certainly logicians, mathematicians, and various other professionals who reason for a living do not use resolution. Instead, their reasoning is driven by suppositions, and has a block-like structure. When the reasoning is deductive, we are thus talking about a Fitch-style natural deduction calculus. A standard presentation of such a calculus is presented in many introductory textbooks; e.g., (Bergmann, Moor, & Nelson 1997; Barwise & Etchemendy 1999). Slate's workspace provides the human with a visual natural calculus of our own invention, in which suppositions are readily identifiable, and arguments are built up in modular fashion.

<sup>1</sup>For example, we speak below both of diagrammatic knowledge, and of Denotational Proof Languages (DPLs), but we don't discuss in this short paper the somewhat complicated DPL designed to represent such knowledge, which is known as Vivid.

Slate reads, shown in Figure 1, is described by three distinct phases:

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. At the present time Slate makes use of Attempto Controlled English (ACE) (Fuchs, Schwertel, & Schwitter 1999; Hoefler 2004), a logically controlled English with a fixed, definite clause grammar and a user-defined vocabulary.<sup>4</sup>

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 (Kamp & Reyle 1993; Kamp & Reyle 1996).

Phase 3 The DRSs are automatically translated into MSL, the chief native language of Slate. (Slate has built-in translators for going from MSL to straight first-order logic (FOL), using long-established theorems (Manzano 1996).) 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 Forthcoming, among others.




Figure 1: Slate's Reading Process

To demonstrate Slate's reading ability, we turn to the Intelligence Analysis case study of "Well-Logger #1"<sup>5</sup>. In this factually-based<sup>6</sup> hypothetical scenario about the potential possession of radiological dispersion bombs by terrorists, the analyst is given (i) 14 Premises — explicitly set off for the analyst, and (ii) a table containing additional necessary information. From these two sets of givens, the analyst is challenged to determine and justify which one of twelve possible recommendations should be issued to superiors in position to launch aggressive law enforcement activity. Slate's reading ability enables the direct and automatic formalization of the textual premises from the given English. Perhaps

<sup>4</sup>Phase 1 is currently a manual operation, but techniques developed by Mollá & Schwitter (Mollá & Schwitter 2001) may allow for at least partial automation of this phase.

<sup>5</sup>This is a rather tricky case study in intelligence analysis created by Selmer Bringsjord for ARDA (now DTO), and makes for a good test of human reasoners, machine reasoners, and systems that assist the interaction between the two. The case study is available at

<http://kryten.mm.rpi.edu/SB.LOGGER.CASESTUDY.tar.gz>

<sup>6</sup>The case is based on the real-life theft of well-loggers, many of which contain enough raw material to fashion so-called "dirty bombs."

the simplest of the 14 premises is "If  $x$  has some sufficient amount ( $\geq k$  curies) of iridium,  $x$  has suitable raw material" which, as a result of Phase 1, is rephrased as "If a person  $x$  has some sufficient iridium then  $x$  has some raw material." By passing the sentence through the remaining phases, the following formula is obtained.

$$\begin{aligned} \forall A, \dots, E (( & \text{object}(A, \text{person}, \text{person}) \wedge \\ & \text{quantity}(A, \text{cardinality}, \text{count\_unit}, B, \text{eq}, 1) \wedge \\ & \text{structure}(A, \text{atomic}) \wedge \\ & \text{object}(C, \text{iridium}, \text{object}) \wedge \\ & \text{quantity}(C, \text{dimension}, \text{unit}, D, \text{eq}, \text{unspecified}) \wedge \\ & \text{structure}(C, \text{mass}) \wedge \\ & \text{property}(C, \text{sufficient}) \wedge \\ & \text{predicate}(E, \text{unspecified}, \text{have}, A, C)) \\ \Rightarrow & \\ \exists F, G, H ( & \text{object}(F, \text{material}, \text{object}) \wedge \\ & \text{quantity}(F, \text{dimension}, \text{unit}, G, \text{eq}, \text{unspecified}) \wedge \\ & \text{structure}(F, \text{mass}) \wedge \\ & \text{property}(F, \text{raw}) \wedge \\ & \text{predicate}(H, \text{unspecified}, \text{have}, A, F))) \end{aligned}$$

Note that the complexity of the above formula is a reflection of the encoding strategy and micro-ontology employed by ACE. Through the application of an inverse encoding scheme, as part of Phase 3, a succinct statement is achieved; namely,

$$\begin{aligned} \forall A, B (( & \text{person}(A) \wedge \text{iridium}(B) \wedge \text{have}(A, B)) \Rightarrow \\ \exists C ( & \text{material}(C) \wedge \text{raw}(C) \wedge \text{have}(A, C))) \end{aligned}$$

This final formula is a correct formalization of the initial premise. Of course, Slate can do likewise for the remaining 13 premises of the case study, and for a sentential expression of the information contained in the table provided to the analyst. Once all this knowledge is represented in Slate's workspace, an argument can be constructed by the human, and then validated by her using argument-checking facilities built into Slage. Once that happens, the case study is solved.

## Machine Reading Diagram-Infused Text Poised-For Learning by Reading (PFLbR)

Put informally, the core idea behind poised-for learning by reading (PFLbR) is this: An agent<sup>7</sup> can be said to have p.-f. learned some text  $\mathcal{I}$  if, in the absence of any output from the system that would normally justify assertions that the system had learned this text, by virtue of having on hand not just declarative knowledge of a sort that can be represented as formulas in a logic, but also

- a certain class of algorithms ready to produce correct output, and
- diagrammatic knowledge produced by reading  $\mathcal{I}$ ,

the agent is ready (poised) to produce such output. The algorithms in question are specifically designed to produce theorems, and proofs that establish theorems. Knowledge that includes such algorithms, and the diagrammatic knowledge to which we have alluded, is

<sup>7</sup>In the now-orthodox sense of 'agent' set out in (Russell & Norvig 2002).

called *p.-f. knowledge*. PFLbR is consistent with results in cognitive science indicating that certain human learners, when reading, are able to self-test and self-explain, which puts them in position to deliver superior performance when actual testing comes (Chi *et al.* 1994; VanLehn, Jones, & Chi 1992). PFLbR is also consistent with the possibility that, in the future, whether an agent had in fact learned in some desirable way could be determined by simply inspecting the “brain” of the agent, obviating the need for carrying out testing.

In what follows we say a few words about the general structure of PFLbR, and the two above-bulleted features of p.-f. knowledge.

### Overall Structure of PFLbR

Let us assume that we are concerned with the kind of sophisticated learning that comes through reading, but also that we are specifically talking about the domain of (elementary) astronomy. The context is thus one in which an agent — who we will call ‘Hugh,’ or sometimes just ‘ $\mathcal{H}$ ’ — is charged with learning about this subject (one quite new to him) from a group of relevant books.<sup>8</sup> Let’s refer to the collective input coming from these books as  $\mathcal{I}$ , and let’s use  $\mathcal{O}$  to refer to Hugh’s output, produced in response to a test (or, as we soon say, to a *query*). At this point the situation is quite generic; it coincides with what some of us have called *Psychometric AI* (Bringsjord & Schimanski 2003); and the situation can be summed up by Figure 2.




Figure 2: Highest-Level View of the Overall Structure

Our context includes that Hugh would ordinarily be said to have learned if he was able to answer sufficiently difficult questions about astronomy correctly, with accompanying justifications of those answers. Accordingly, we assume that a query  $\mathcal{Q}$  is given to Hugh, and that he would be asked to provide an answer  $\mathcal{A}$  to it, supported by justification  $\mathcal{J}$ ; and we assume that the pair  $(\mathcal{J}, \mathcal{A})$ , which comprises  $\mathcal{O}$ , is of very high quality.

### Diagrammatic Knowledge and Additional Structure of PFLbR

It’s tempting to say that the elements of  $\mathcal{I}$  are characters, words, sentences, paragraphs, and so on. This response is inaccurate. As you may remember from your grade school days, or perhaps as you can guess, the stimuli in the case at hand, that which appears on the pages of the books Hugh studies, includes both linguistic and pictorial information. Consider for example any of the books on astronomy cited earlier. Each of them, *on each and every page*, includes both textual *and* diagrammatic information. As an example, consider that

<sup>8</sup>A nice group of such books is: (Simon 1992; Awan 2004; Lippincott 2004; Dickinson 1998).

constellations are picked out and remembered with help from diagrams superimposed on stars and planets seen when looking (save for Sagittarius) on a line of sight beyond the Milky Way. Figure 3 indicates how the trick works for Scorpio.



Figure 3: The superimposition of a scorpion to produce Scorpio. The input here involves diagrammatic/visual information, as well as textual information. Taken from (Awan 2004).

This implies that we can provide a bit more structure in our overview of p.-f. learning: We can say that the input  $\mathcal{I}$  is composed of textual information  $\Theta$  and diagrammatic information  $\Delta$ . At this point p.-f.-learning can be summed up by Figure 4.




Figure 4: High-Level View with Basic Structure of Input  $\mathcal{I}$  and Output  $\mathcal{O}$

But we can uncover additional structure in p.-f.-learning. We draw your attention to Figure 5. Notice that we now assume that the input, courtesy of help from a natural language understanding module, leads to the representation of this input (in some logical system; see e.g., (Ebbinghaus, Flum, & Thomas 1984; Bringsjord & Ferrucci 1998a; Bringsjord & Yang 2003)), augmented by a system for handling pictorial input  $\Delta$ . In addition, we include now  $\Psi$ , knowledge Hugh was in command of prior to his study of astronomy. P.f. knowledge is denoted by  $\Pi$ , which is produced from: the representation of both text and diagrams; from  $\Psi$ ; and from queries  $Q_1, \dots, Q_m$  that the agent produces internally, in a “self-reflective” mode that helps anticipate actual queries  $Q$ . Once  $\Pi$  is constructed, a query  $Q$  leads to a representation of the output (in some logical system), and this representation, with help from a natural language generation module, yields the final answer and corresponding justification in natural language. Again, the overall process is summed up in Figure 5.

The *representation* of the input is itself a mixture of the syntactic and semantic. There is now overwhelm-




Figure 5: Additional Structure in the Sequence of P.f.-Learning

ing empirical evidence for the view that while some human knowledge does seem to be accurately modeled in purely syntactic or symbolic form (the theory of *mental logic* proposes such knowledge; see e.g. (Rips 1994; Yang, Braine, & O'Brien 1998)), some knowledge is represented in irreducibly semantic form, or in what we call, using Johnson-Laird's descriptor, *mental models* (e.g., see (Johnson-Laird 1983; Johnson-Laird et al. 2000)). Mental models can be pictorial or imagistic in form, rather than symbolic, syntactic, or linguistic. P.-f. knowledge of astronomy includes *both* types of knowledge. The theory within cognitive science that posits, explains, and empirically justifies (among other things) this mixed mode is *mental meta-logic* ('MML,' for short), and is due to Bringsjord and Yingrui Yang (Yang & Johnson-Laird 2000a; Yang & Johnson-Laird 2000b; Yang & Bringsjord forthcoming; Rinella, Bringsjord, & Yang 2001; Yang & Bringsjord 2001a; Yang & Bringsjord 2001b; Yang, Braine, & O'Brien 1998).<sup>9</sup>

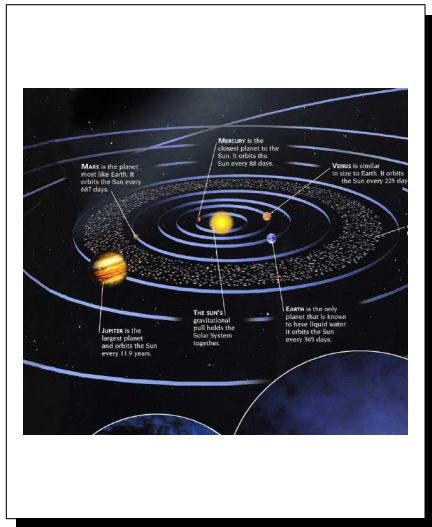


Figure 6: Standard Overview of our Solar System (from Awan 2004)

In order to make this more concrete, let's turn to

<sup>9</sup>MML holds as well that another hallmark of human reasoning is *meta-reasoning*. This means, for example, that humans are capable of reasoning about patterns of reasoning. Meta-reasoning is mechanized in PFLbR through *methods*, discussed below.

some simple information about our solar system. The basic overview of the system is traditionally provided to readers by pictures like that shown in Figure 6, from (Awan 2004). Let's suppose, then, that Hugh has specifically read such material. In addition, let's suppose that Hugh successfully answers and justifies the query: Is it true that all the planets inside the asteroid belt are smaller than the sun? If he has truly learned from his reading, then he has p.-f. learned from that reading; and this in turn implies that his answer and justification flow from p.-f. knowledge. One possibility for this p.-f. knowledge is shown in Figure 7, which makes use of the kind of blocks world often used in AI for expository purposes. (This particular blocks world is that of Hyperproof (Barwise & Etchemendy 1994).) Here, Hugh has a mental model abstracted from the picture shown in Figure 6; this model corresponds to the first blocks world image. In this image, the sun is shown on the far left, and then the planets move to the right diagonally out to the lineup of tetrahedrons; this lineup represents the asteroid belt. The large dodecahedron after the belt is Jupiter (the remaining four planets aren't shown). In addition, we assume that Hugh knows (syntactically) that the sun is quite large, and a disjunction: that Earth is either roughly the same size as Venus or Mars. Given this, he has knowledge poised to produce an affirmative response in response to the query, as well as a corresponding justification. The affirmative response corresponds to the last formula in the sequence, and overall the sequence is poised for an argument by cases. The two cases are the two disjuncts, and each leads to the situation presented in the second image, in which the relative sizes of the sun and Earth are pinned down. It's a matter of direct mental observation to infer from this second image that all four interior planets are indeed smaller than the sun. Notice that the p.-f. knowledge in this case is *not* a proof. Rather, it's knowledge that is merely *poised for* providing an argument that in turn yields an affirmative response to the query.

### Remarks on "Poised-For Proving"

PFLbR is based on denotational proof languages (DPLs) (Arkoudas 2000). DPLs integrate computation and deduction. They can be used as regular programming languages, as languages for presenting and checking formal proofs in natural-deduction style, and as languages for expressing trusted proof-search algorithms — so-called *methods*. Here we will focus on methods, which are a key element of DPLs for PFLbR.

Put simply, a method is an algorithm for constructing a proof; some methods are allowed to be higher-order: they take methods as arguments. Hugh's p.-f. knowledge can be represented as a higher-order method, that is, as an algorithm for producing a justification for the answer to the query, in the form of a proof, when supplied with appropriate lower-level methods as arguments in order to fill in certain gaps within the higher-order method's reasoning. In a nutshell, if Hugh truly



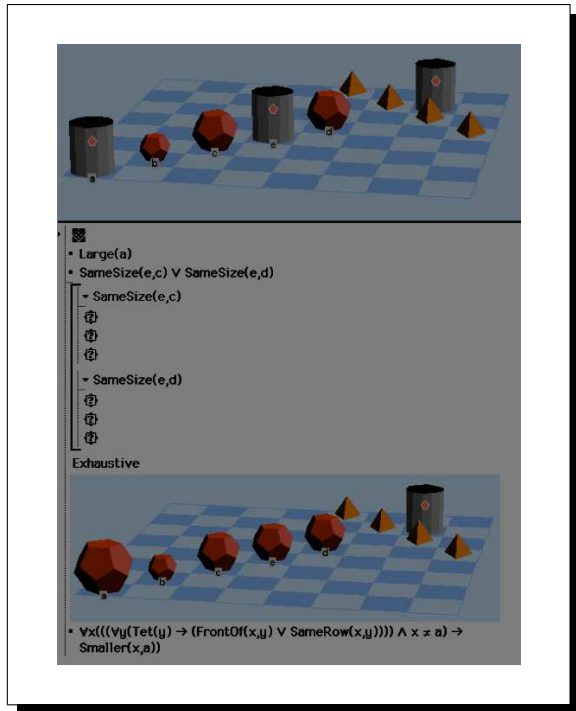


Figure 7: Some Possible P.f.-Knowledge of Our Solar System. The p.f.-knowledge here is shown in Hyperproof. In the first picture (a representation of a diagram Hugh read), the leftmost object, *a*, is the sun. Since its size at this point is unknown, the object depicting it is a cylinder (cylinders indicate that the actual shape of the object is **unknown**; Hyperproof, and our own diagrammatic DPLs (e.g., the aforementioned Vivid, allow for this third truth value), not a dodecahedron. Moving from left to right, *b* is Mercury, *c* Venus, *e* Earth, *d* Mars. The asteroid belt is represented by the line of tetrahedrons, and then the first planet beyond this belt is a representation of Jupiter (whose size is also unclear to Hugh). The second picture shows that sizes resolved, on the basis of an argument by cases that is poised to be completed.

learns by reading, then, *before* he is tested, he stands ready with algorithms which, when fired in response to an actual test, will efficiently produce correct answers and justifications. Such a method goes beyond those in the well-established DPL known as Athena (Arkoudas) by allowing reference to visual models or diagrams — or as we simply say in the sample code (Figure 8), to *diagrams*. A method corresponding to the p.f. knowledge in Figure 7 can be formulated as the `show-relative-size-of-sun` method shown in Figure 8.<sup>10</sup> Note again that the methods here incorporate visual information and reasoning. For instance, the assertion that the sun is large is obtained via the “visual” inference rule `inspect`, which consults a stored diagram

<sup>10</sup>Due to lack of space, we don’t explain the syntax and semantics of every Athena construct appearing in Figure 8. A succinct reference describing Athena’s syntax and formal semantics can be found elsewhere (Arvizo).

```
(define show-relative-size-of-sun
  (method (M1 M2)
    (dlet ((premise1 ((isLarge sun) BY (!inspect diagram1)))
          (premise2 ((forall ?x (iff (within-asteroid-belt ?x)
                                   (or (= ?x mercury)
                                       (= ?x venus)
                                       (= ?x earth)
                                       (= ?x mars))))
                    BY (!inspect diagram1)))
          (premise3 ((and (smallerSize mercury sun)
                          (smallerSize venus sun)
                          (smallerSize mars sun))
                    BY (!inspect diagram1)))
          (premise4 (!claim (forall ?x ?y ?z
                              (if (and (smallerSize ?x ?z)
                                      (roughlySameSize ?x ?y))
                                  (smallerSize ?y ?z))))))
          (premise5 (!claim
                    (or (roughlySameSize earth venus)
                        (roughlySameSize earth mercury))))))
      (premises [premise1 premise2 premise3
                premise4 premise5])
      (case1 (assume (sameSize earth venus)
                    (!M1 (add (sameSize earth venus) premises))))
      (case2 (assume (sameSize earth mercury)
                    (!M2 (add (sameSize earth mercury)
                              premises))))))
    (!by-cases case1 case2)))
```

Figure 8: An Athena Method Representing Part of Hugh’s P.f.-knowledge.

in order to verify the claimed conclusion. Premises 1, 2 and 3 in Figure 8 are obtained in that way from stored visual information (referred to in the code as `diagram1`) derived from  $\Delta$  within the input Hugh has been supplied with via his reading. Premises 4 and 5 come from prior and background knowledge. Prior knowledge is already stored in Hugh’s knowledge base and would be recalled for the purposes of running the method in order to construct the justification/answer pair in response to a query.




Figure 9: Full Structure in the Sequence of P.f.-Learning

## Acknowledgements

We are indebted to NSF, DARPA, and AFOSR for funding that has enabled the majority of the machine-reading research described herein. With respect to the Slate system, and our giving it the capacity to read logically controlled English, with much gratitude, we acknowledge the financial support provided by DTO (formerly Advanced Research and Development Activity (ARDA)) in the past, through two programs: contract # MDA-904-03-C-0408 to Booz Allen Hamilton (general contractor) and RPI (sub-contractor), Novel Intelligence from Massive Data Program (NIMD); and IKRIS. We are also indebted to DTO for just-arrived grants

(AQAIN and ARIVA programs) that will allow us to apply and enhance our machine reading theory and technology in connection with the Solomon and Slate systems.

## References

- [Arkoudas ] Arkoudas, K. Athena. <http://www.cag.csail.mit.edu/~kostas/dpls/athena>.
- [Arkoudas 2000] Arkoudas, K. 2000. *Denotational Proof Languages*. Ph.D. Dissertation, MIT, Department of Computer Science, Cambridge, USA.
- [Arvizo ] Arvizo, T. A virtual machine for a type- $\omega$  denotational proof language. Masters thesis, MIT, June 2002.
- [Awan 2004] Awan, S. 2004. *Night Sky Atlas*. London, UK: Dorling Kindersley Limited.
- [Barwise & Etchemendy 1994] Barwise, J., and Etchemendy, J. 1994. *Hyperproof*. Stanford, CA: CSLI.
- [Barwise & Etchemendy 1999] Barwise, J., and Etchemendy, J. 1999. *Language, Proof, and Logic*. New York, NY: Seven Bridges.
- [Bergmann, Moor, & Nelson 1997] Bergmann, M.; Moor, J.; and Nelson, J. 1997. *The Logic Book*. New York, NY: McGraw Hill.
- [Blackburn & Bos Forthcoming] Blackburn, P., and Bos, J. Forthcoming. Working with Discourse Representation Theory: An Advanced Course in Computational Semantics.
- [Bringsjord & Ferrucci 1998a] Bringsjord, S., and Ferrucci, D. 1998a. Logic and artificial intelligence: Divorced, still married, separated...? *Minds and Machines* 8:273–308.
- [Bringsjord & Ferrucci 1998b] Bringsjord, S., and Ferrucci, D. 1998b. Reply to Thayse and Glymour on logic and artificial intelligence. *Minds and Machines* 8:313–315.
- [Bringsjord & Schimanski 2003] Bringsjord, S., and Schimanski, B. 2003. What is artificial intelligence? Psychometric AI as an answer. In *Proceedings of the 18<sup>th</sup> International Joint Conference on Artificial Intelligence (IJCAI-03)*, 887–893.
- [Bringsjord & Yang 2003] Bringsjord, S., and Yang, Y. 2003. Representations using formal logics. In Nadel, L., ed., *Encyclopedia of Cognitive Science Vol 3*. London, UK: Nature Publishing Group. 940–950.
- [Chi et al. 1994] Chi, M.; Leeuw, N.; Chiu, M.; and Lavancher, C. 1994. Eliciting self-explanations improves understanding. *Cognitive Science* 18:439–477.
- [Dickinson 1998] Dickinson, T. 1998. *Nightwatch: A Practical Guide to Viewing the Universe*. Willowdale, Canada: Firefly Books.
- [Ebbinghaus, Flum, & Thomas 1984] Ebbinghaus, H. D.; Flum, J.; and Thomas, W. 1984. *Mathematical Logic*. New York, NY: Springer-Verlag.
- [Friedland et al. 2004] Friedland, N.; Allen, P.; Matthews, G.; Witbrock, M.; Baxter, D.; Curtis, J.; Shepard, B.; Miraglia, P.; Angele, J.; Staab, S.; Moench, E.; Oppermann, H.; Wenke, D.; Israel, D.; Chaudhri, V.; Porter, B.; Barker, K.; Fan, J.; Chaw, S. Y.; Yeh, P.; Tecuci, D.; and Clark, P. 2004. Project halo: Towards a digital aristotle. *AI Magazine* 29–47.
- [Fuchs, Schwertel, & Schwitter 1999] Fuchs, N. E.; Schwertel, U.; and Schwitter, R. 1999. Attempto Controlled English (ACE) Language Manual, Version 3.0. Technical Report 99.03, Department of Computer Science, University of Zurich, Zurich, Switzerland.
- [Genesereth & Nilsson 1987] Genesereth, M., and Nilsson, N. 1987. *Logical Foundations of Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann.
- [Hoeffler 2004] Hoeffler, S. 2004. The Syntax of Attempto Controlled English: An Abstract Grammar for ACE 4.0. Technical Report ifi-2004.03, Department of Informatics, University of Zurich, Zurich, Switzerland.
- [Johnson-Laird et al. 2000] Johnson-Laird, P. N.; Legrenzi, P.; Girotto, V.; and Legrenzi, M. S. 2000. Illusions in reasoning about consistency. *Science* 288:531–532.
- [Johnson-Laird 1983] Johnson-Laird, P. N. 1983. *Mental Models*. Cambridge, MA: Harvard University Press.
- [Kamp & Reyle 1993] Kamp, H., and Reyle, U. 1993. *From Discourse to Logic: Introduction to Model-theoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Springer, 1 edition.
- [Kamp & Reyle 1996] Kamp, H., and Reyle, U. 1996. A Calculus for First Order Discourse Representation Structures. *Journal of Logic, Language and Information* 5:297–348.
- [Lippincott 2004] Lippincott, K. 2004. *Astronomy*. New York, NY: DK Eyewitness Books.
- [Manzano 1996] Manzano, M. 1996. *Extensions of First Order Logic*. Cambridge, UK: Cambridge University Press.
- [Mollá & Schwitter 2001] Mollá, D., and Schwitter, R. 2001. From Plain English to Controlled English. In *Proceedings of the 2001 Australasian Natural Language Processing Workshop*, 77–83.
- [Nilsson 1991] Nilsson, N. 1991. Logic and Artificial Intelligence. *Artificial Intelligence* 47:31–56.
- [Rinella, Bringsjord, & Yang 2001] Rinella, K.; Bringsjord, S.; and Yang, Y. 2001. Efficacious logic instruction: People are not irremediably poor deductive reasoners. In Moore, J. D., and Stenning, K., eds., *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Lawrence Erlbaum Associates. 851–856.
- [Rips 1994] Rips, L. 1994. *The Psychology of Proof*. Cambridge, MA: MIT Press.
- [Russell & Norvig 2002] Russell, S., and Norvig, P. 2002. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice Hall.
- [Simon 1992] Simon, S. 1992. *Our Solar System*. New York, NY: Morrow Junior Books.
- [VanLehn, Jones, & Chi 1992] VanLehn, K.; Jones, R.; and Chi, M. 1992. A model of the self-explanation effect. *Journal of the Learning Sciences* 2(1):1–60.
- [Yang & Bringsjord 2001a] Yang, Y., and Bringsjord, S. 2001a. Mental metalogic: A new paradigm for psychology of reasoning. In *Proceedings of the Third International Conference on Cognitive Science (ICCS 2001)*. Hefei, China: Press of the University of Science and Technology of China. 199–204.
- [Yang & Bringsjord 2001b] Yang, Y., and Bringsjord, S. 2001b. The mental possible worlds mechanism: A new method for analyzing logical reasoning problems on the gre. In *Proceedings of the Third International Conference on Cognitive Science (ICCS 2001)*. Hefei, China: Press of

the University of Science and Technology of China. 205–210.

[Yang & Bringsjord forthcoming] Yang, Y., and Bringsjord, S. forthcoming. *Mental Metalogic: A New, Unifying Theory of Human and Machine Reasoning*. Mahway, NJ: Erlbaum.

[Yang & Johnson-Laird 2000a] Yang, Y., and Johnson-Laird, P. N. 2000a. How to eliminate illusions in quantified reasoning. *Memory and Cognition* 28(6):1050–1059.

[Yang & Johnson-Laird 2000b] Yang, Y., and Johnson-Laird, P. N. 2000b. Illusory inferences with quantified assertions. *Memory and Cognition* 28(3):452–465.

[Yang, Braine, & O'Brien 1998] Yang, Y.; Braine, M.; and O'Brien, D. 1998. Some empirical justification of one predicate-logic model. In Braine, M., and O'Brien, D., eds., *Mental Logic*. Mahwah, NJ: Lawrence Erlbaum Associates. 333–365.