Mental MetaLogic: A New Paradigm in Psychology of Reasoning

Yingrui Yang & Selmer Bringsjord
yangyri@rpi.edu • selmer@rpi.edu
Dept. of Philosophy, Psychology & Cognitive Science (S.B. & Y.Y.)
The Minds & Machines Laboratory (S.B. & Y.Y.)
Department of Computer Science (S.B.)
Rensselaer Polytechnic Institute
Troy NY USA

Abstract

We introduce a new theory of reasoning, Mental MetaLogic (MML), which unifies the mental logic and mental models theories currently deadlocked in opposition to each other. MML, as we explain, is based on the "psychologizing" of *all* of symbolic logic — which includes a syntactic component (proof theory) with which mental logic is associated, a semantic component (model theory and its relatives) with which mental models is associated, and a "meta" component (metatheory or metalogic), hitherto largely ignored in the psychology of reasoning. We set out the seven research questions that MML is designed to enable answers to, present the research agenda that arises from the goal of answering three of these questions, and briefly summarize two samples of MML-based empirical/experimental work, and a sample of MML-based AI R&D.

1 Introduction

Few would dispute the claim that logical reasoning is a fundamental human ability. Smith can't find his car keys, but knows both that he last had them minutes ago in the kitchen, and that he has only been in one other room since then. After scouring every millimeter of the other room, he reenters the kitchen and, with confidence, pokes around and finds them. His confidence is the result of deductive reasoning, that much we know. But what *is* deductive reasoning? In the last three decades, the psychology of reasoning has produced countless papers in psychological journals (such as *Psychological Review*, *Science*, and *Behavioral & Brain Sciences*) aimed at providing an answer to this question, and aimed thereby at providing a scientific explanation of what Smith has just done. Alas, the answer and explanation have never arrived. The reason is that if one had to capture the psychology of reasoning in a word or two, the first choice would doubtless be "The Controversy."

The Controversy is well-known to many of our readers. In brief, it arises from two competing theories for modeling and explaining ordinary[1] human reasoning (especially deductive reasoning). On the one hand, mental logic theory claims that logically untrained humans reason using domain-independent inference rules or schemas (e.g., that from $p \vee q$ and $\neg p$ one can infer to q; proponents of mental logic would say that something like this particular rule is operative in the case of our hypothetical Smith). On the other hand, mental model theory holds that ordinary humans reason using not abstract syntactic rules, but semantic possibilities they in some sense imagine. Scientists on both sides are incontestably seminal and prolific, and the content offered by each is powerful. But the two sides stand vehemently apart, and as a result the science of reasoning is, in a real sense, fundamentally paralyzed.[2] MML will end this paralysis.

2 The 3 Components of Symbolic Logic

MML is a *meta*-theory because it is inspired by modern symbolic logic, which is primarily a *meta*-inquiry, as it is devoted to the mathematical study of mathematical and logical systems.[3] And although MML is a *psychological* theory, to present it we need a quick review of the three main components of symbolic logic.

In broad strokes, modern symbolic logic has three main components: one is syntactic, one is semantic, and one is meta-theoretical. The syntactic component includes specification of the alphabet of a given logical system, the grammar for building well-formed formulas (wffs) from this alphabet, and a proof theory that precisely describes how and when one formula can be proved from a set of formulas. The semantic component includes a precise account of the conditions under which a formula in a given system is true or false. The meta-theoretical component includes theorems, conjectures, and hypotheses concerning the syntactic component, the semantic component, and connections between these two components.

[1] We say 'ordinary' here to exclude those experts who are extensively educated in the ways of formal deductive reasoning, and those humans apparently gifted to the degree that their performance on deductive reasoning tasks is what would be expected from such experts. For a discussion of such experts, see (Bringsjord, Bringsjord & Noel 1998).

[2] There are ways of avoiding the Controversy. For example, one can lump all abstract "analytic" reasoning (deductive reasoning of all kinds, statistical reasoning, and so on) together, and give this reasoning a name, and then give another name to intuitive, concrete, non-symbolic thinking thinking. This is something Stanovich & West (2000) do. But such a move really is avoidance; it makes no progress in the face of the challenge to give a precise scientific account of all the various forms of analytic reasoning, including that species of most concern to us: deductive reasoning.

[3] Tellingly, mathematical logic, a part of symbolic logic, is sometimes called 'meta-mathematics.'

As an example, consider the familiar logical system known as the propositional calculus (also known as sentential logic). The alphabet in this case is simply an infinite list $p_1, p_2, \ldots, p_n, p_{n+1}, \ldots$ of propositional variables (according to tradition p_1 is p, p_2 is q, and p_3 is r), and the five familiar truth-functional connectives $\neg, \rightarrow, \leftrightarrow, \wedge, \vee$. The grammar for the propositional calculus is straightforward and well-known.

Various proof theories for the propositional calculus are possible. One such theory, a so-called natural deduction theory, includes the inference rule which says that from a pair of formulas having the form $\phi \rightarrow \psi$ and ϕ one can infer ψ; this rule is known as *modus ponens*. Given a full set of rules of inference, one can speak in general about an individual formula ϕ being derivable from a set Φ of formulas that serve as suppositions; this is traditionally written

$$\Phi \vdash \phi.$$

A formula derivable from the null set is said to be a theorem, and where ϕ is such a formula, we follow custom and write

$$\vdash \phi$$

to express such a fact.

The semantic side of the propositional calculus is of course based on the notion of a *truth-value assignment*, an assignment of a truth-value (T or F, or sometimes, instead, 1 or 0) to each propositional variable p_i, and upon the truth-functions expressed by the five connectives alluded to just above; usually these truth-functions are defined via *truth tables*.

Given a truth-value assignment v, the truth or falsity of a given propositional formula, no matter how complicated, can be mechanically calculated. Armed with a truth-value assignment v, we can say that v "makes true" or "models" or "satisfies" a given formula ϕ; this is standardly written

$$v \models \phi.$$

Some formulas are true on all models. For example, the formula $((p \vee q) \wedge \neg q) \rightarrow p$ is in this category. Such formulas are said to be *valid* and are sometimes referred to as *validities*. To indicate that a formula ϕ is valid we write

$$\models \phi.$$

Another important semantic notion is *consequence*. An individual formula ϕ is said to be a consequence of a set Φ of formulas provided that all the truth-value assignments on which all of Φ are true is also one on which ϕ is true; this is customarily written

$$\Phi \models \phi.$$

At this point it's easy enough to describe, through the example of the propositional calculus, the metatheoretical component of modern symbolic logic that is particularly relevant to Mental MetaLogic: Meta-theory would deploy logical and mathematical techniques in order to answer such questions as whether or not provability implies consequence, and whether or not the reverse holds. When the first direction holds, a logical system is said to be *sound*, and this fact can be expressed in the notation we've invoked as

$$\text{If } \Phi \vdash \phi \text{ then } \Phi \models \phi.$$

Roughly put, a logical system is sound if it's guaranteed that true formulas can only yield (through derivations) true formulas; one cannot pass from the true to the false. When the "other direction" is true of a system that system is said to be *complete*; in our notation this is expressed by

$$\text{If } \Phi \models \phi \text{ then } \Phi \vdash \phi.$$

As a matter of fact, the propositional calculus is both provably sound and complete. One consequence of this is that all theorems in the propositional calculus are valid, and all validities are theorems.[4]

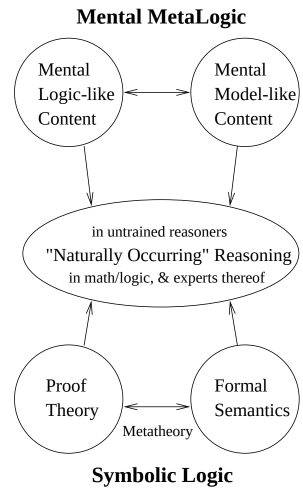

Figure 1: Overview of Mental MetaLogic Symmetry with Symbolic Logic

3 The Components & ML, MM, MML

Now, using the review provided in the previous section, we characterize mental logic, mental models, and Mental MetaLogic. This characterization is summed up in Figure 1.

The representational system of current mental logic theory can be viewed as (at least to some degree) a psychological selection from the syntactic components of systems studied in modern symbolic logic. For example, in mental logic *modus ponens* is often selected as a schema while *modus tollens*[5] isn't; yet both are valid inferences in most standard logical systems. Another example can be found right at the heart of Lance Rips' system of mental logic, PSYCOP, set out in (Rips 1994), for this system includes conditional proof (p. 116), but *not* the rule which sanctions passing from the denial of

[4] A fact captured by: $\models \phi$ if and only if $\vdash \phi$.

[5] From $\phi \rightarrow \psi$ and $\neg\psi$ one can infer to $\neg\phi$.

a conditional to the truth of this conditional's antecedent (pp. 125–126).

In parallel to the relationship between mental logic and the syntactic component of symbolic logic, mental model theory is related to the semantic side of symbolic logic. For example, consider the truth table for conditionals in the propositional calculus, shown in Figure 2. The representational system of the mental model theory consists of a number of sets of initial mental models selected from the semantic possibilities. Because each proposition can have two truth-values, true or false, the truth table as a truth function for $\rightarrow$ has four semantic (truth) possibilities. When p is true and q is false, the conditional is false. By the truth principle of mental model theory, reasoners represent only what is true but not what is false, due to limited working memory. So this possibility (row 3) is excluded from the set of possibilities. Proponents of MM also predict that ordinary reasoners will not reason from a false antecedent, thus the two semantic possibilities where p is false (rows 4 and 5) are not represented explicitly, but rather are represented implicitly by making a mental footnote as ". . .." Thus, the only explicit model left is the case in which both p and q are true (row 2). Hence, by using the presence/absence of a mental token to substitute the truth value accordingly, the initial model set for the form of conditionals consists of one explicit and one implicit mental model, which can be represented as

$$\begin{array}{cc} p & q \\ \multicolumn{2}{c}{\ldots} \end{array}$$

As we have seen earlier, the meta-theoretical component of modern symbolic logic covers formal properties (e.g., soundness and completeness) that bridge the syntactic and semantic components of logical systems. In selecting from either but not both of the syntactic or semantic components of logical systems, psychologists have, in a sense, broken the bridges built by logicians and mathematicians. Mental logic theory and mental model theory became incompatible from the perspective of symbolic logic: the former is explicitly and uncompromisingly syntactic, while the latter is explicitly and uncompromisingly semantic. Mental MetaLogic aims to bridge between these two theories.

4 Our Seven Driving Questions

Seven questions drive our setting out and using Mental MetaLogic theory, viz.,

Q1 What is deductive reasoning? More specifically:

- Q1a What is the "essence" of deductive reasoning?
- Q1b What is the structure of deductive reasoning?

Q2 How do people reason deductively?

Q3 How do people (best) learn to reason deductively, and how can we more effectively teach deductive reasoning?

Q4 How can we (best) tell if people are reasoning deductively (well)? More specifically:

- Q4a How good are established "high stakes" tests of deductive reasoning (e.g., AR and LR in GRE and LSAT)? What do these tests test, really?
- Q4b How good are other tests of deductive reasoning (e.g., the Watson-Glazer)?
- Q4c What new and better tests of deductive reasoning can be developed?

Q5 How can we harness MML in such a way as to build appropriate specialized AI reasoning systems that go beyond the purely ML-based BRUTUS system from Bringsjord & Ferrucci (2000)? More specifically:

- Q5a How can we incarnate MML in the form of automated reasoning systems that allow for the full range of deductive reasoning across diagrams, images and various other "semantic" elements, and across proof schemes beyond standard extensional first-order logic, e.g., second-order logic, infinitary logic, self-referential systems, modal logic, etc.?
- Q5b How can we harness MML in this way as to build AI reasoning systems that teach deductive reasoning?
- Q5c How can we harness MML in this way as to build AI reasoning systems that automatically generate tests (e.g., the GRE) of reasoning?

Q6 What is deductive reasoning good for in "real life"?

Q7 How do human capacities and skills related to deductive reasoning change across the spectrum of toddlers to the elderly?

5 5-Part MML Agenda for Q1, Q2, Q5

Because of space constraints here, we simply leave aside Q3, Q4, Q6, and Q7.[6] Here is an agenda produced by the desire to answer Q1, Q2, and Q5:

1. *Empirically Investigate Psychological Correlates of Symbolic Logic*. MML must propose and empirically investigate the psychological correlates of symbolic logic. For example, MML must conceptually clarify and redefine meta-theoretical notions such as consistency and completeness *from the psychological perspective*. We briefly described the purely logical concept of completeness above. Corresponding to this concept is the *psychological* notion of completeness: it is for the most part concerned with the *modifiability* of the competing mental logic and mental models theories: if a logical system can be modified to accommodate empirical discoveries about ordinary deductive reasoning, it can in the psychological sense be said to be complete. As to the purely logical concept of soundness, it generally doesn't allow invalid rules (rules that yield false formulas from true ones), but that might not necessarily be the case in "real life:" humans may well routinely make invalid inferences. If so, one of the jobs of MML is to give a precise account of the conditions under which these inferences are made.

2. *Extend the "Representational Reach of ML and MM*. In MML, each of the competing theories (mental logic, mental models) needs to have its "representational reach" extended in order to account for the empirical evidence for the other. In addition, *novel* representational reach, that is, representational power not anticipated in any way by ML and MM, needs to be derived from MML, and instantiated in general-purpose automated reasoning system (see below the fifth part of the agenda).

[6] A paper designed in part to answer Q3 is under review: (Rinella, Bringsjord & Yang under review).

p	q	$p \rightarrow q$
T	T	T
T	F	F
F	T	T
F	F	T

Figure 2: Truth Table for the Conditional in the Propositional Calculus

3. *Predict New Empirical Phenomena*. New empirical phenomena should be predicted by using representational schemes that result from modifications of those schemes employed by mental logic and mental models. For example, illusory schemas in MML should predict the illusions in reasoning with quantified *compound* statements. (Illusions studied through mental models have to this point focussed on propositional statements and quantified *atomic* sentences.)

4. *Develop New Coding Schemes/Techniques for Relevant Verbal Protocols*. The fourth task to be undertaken by MML is that of modeling the interaction of different reasoning skills, such as "schema-skill" and the "mental-model" skill. Thus, the development of a coding method for "think-aloud" protocols should be a part of MML. (The fourth task is tackled in (Yang & Bringsjord under review).)

5. *Implemented Systems*. Finally, MML should give rise to implemented systems. For example, MML should undergird a theorem prover that captures human mathematical reasoning in its full glory. Today's theorem provers are invariably restricted essentially to mental logic: they are based on purely syntactic proofs, and almost always these proofs are in standard first-order logic (Bringsjord 1998). (Arguably the world's most sophisticated automated reasoning system is OSCAR (Pollock 1995), which is based exclusively in natural deduction proofs in first-order logic.) Mental MetaLogic should change all this by producing revolutionary automated reasoning systems able to handle diagrams, images, and other kinds of "semantic" representations.[7] The precursors to these systems should prove useful for teaching deductive reasoning.

6 Point 2 of The Agenda in Action

The following two samples indicate how it's possible for a psychological metatheory to resolve two major recent challenging issues in the domain of quantified predicate reasoning.

Sample 1. People can solve many problems, but of course they don't find all these problems equally difficult. In other words, we can distinguish between two levels of cognitive routines in reasoning: Reasoners can solve certain kinds of inference problems errorlessly, and they are able to perceive systematically the difference between these problems in terms of their difficulty. Mental logic theory proposes a set of inference schemas, and predicts

P1 Reasoners can apply them in reasoning errorlessly, and

P2 Each schema has its particular difficulty.

For example, consider a pair of verbal instantiations of two schemas:

(1) All the children got some red beads. Therefore, all the girls got red beads.

(2) All the beads are either red or blue. The red beads are square. The blue beads are round. Therefore, all the beads are either round or square.

People can apply both schemas errorlessly, but we need to make sure that (2) is harder than (1). Now consider a third problem (3):

All the children found some red beads.

The red beads were either round or square

The round beads were plastic.

The square beads were wooden.

- Did all the girls find either plastic or wooden beads?

This problem is soluble by using the appropriate variations of the schemas in (1) and (2). Yang, Braine & O'Brien (1998) conducted a large-scale project to examine the mental predicate logic devised by Braine (1998). About 130 inference problems like (3) were constructed and all the problems were soluble by using 10 direct reasoning schemas like (1) and (2). The participants (N=180, all experiments were individually administrated) were instructed to solve each problem first and then to rate its difficulty on a 7-point scale. As predicted, the overall error rate (of more than 13,000 responses) was lower than 3%. Thus, the introspective difficulty ratings could be used to generate the difficulty-weight for each schema. A parametrical model was used to generate schema-weights by using a least-square method. The theory predicts the mean difficulty of an inference problem by using the sum of difficulty-weights of those schemas included in the proposed solution to solve that problem. As to results, the cross-validation tests (i.e., using the set of weights generated from one data-set to predict the mean difficulty ratings of any other data-sets; total of 9 samples were used) show that the correlation was reliably as high as 0.93, accounting for more than 80% of the variance. This result corrected the illusion that the more schemas required to solve a problem, the more difficult it is. For example, consider some Problem A, whose solution needs three easy schemas and some Problem B, whose solution requires two difficult schemas. The predicted difficulty for Problem A is 1.5 (0.5×3), and that for Problem B is 3 (1.5×2). This result provides direct empirical evidence supporting mental logic theory, and challenges the mental model theory because (i) current mental model theory cannot represent some of the schemas (e.g., a schema involving a quantified conjunction), and (ii) the weights were generated from empirical data. The question is, how could mental model theory account for the schema weights?

Mental MetaLogic responds this challenge in two steps. The first step is to modify the mental model representation in order to represent all the schemas. For example, the form `FOR ALL x (IF Px THEN (Fx OR`

[7] For a description of this great challenge, see (Bringsjord & Bringsjord 1996).

`Gx))` is represented by the model

$$[p]_g^f$$

where [] indicates the universal quantifier, and together the superscript f and the subscript g indicates $f \vee g$. This introduces a logical structure into a single model, which is different from the original version of mental model representation. For the second step, as an example, consider the quantified *modus ponens*:

All the As are Bs. All the Bs are Cs. ∴All the As are Cs.

Its model representation is as follows

[a] b	a [b] c

Here the left-box indicates the model for the first premise, and the next box indicates the result of integrating the information from the second premise with the previous model. Now, one can easily count that the number of the mental tokens used in this model representation is 5. Yang (n.d.) found that the number of the mental tokens included in a consistently modified mental model representation of this type is a very accurate predictor for the schema weights. That is, for each schema examined, its difficulty weight equals the number of the mental tokens included in its minimum mental model representation (divided by 10). The correlation between the two sets of weights (predicted by the mental model method and generated by the mental logical parametrical method) is almost perfect (r = .99).

Sample 2. People can make many correct inferences if they use the schematic representation proposed by mental logic theory; there is simply no questioning this. People can also get many inferences right if they use mental model representations; however, mental model theory also accommodates illusory inferences, which seem very compelling but are fallacious. Consider the following examples:

Only one of the following statements is true:

- Some of the plastic beads are not red, or
- None of the plastic beads are red.

Problem 1 Is it possible that none of the red beads are plastic?

Problem 2 Is it possible that at least some of the red beads are plastic?

As we said earlier, Problem 1 is predicted by mental model theory as an illusory problem (i.e., the correct answer is "No" but most people would say "Yes"), and Problem 2 is predicted to be a control problem (i.e., most people can get the correct answer, "Yes"). Johnson-Laird and his colleagues have published, at last count, a dozen of papers on this topic, including the recent *Science* paper (P. N. Johnson-Laird, Girotto & Legrenzi 2000). These results have caused something arguably approaching a crisis for mental logic theory, because current mental logic theory cannot account for both the controls and the illusions (Yang & Johnson-Laird 2000). But here is how MML works in a series of steps toward resolving the crisis. First, MML analyzes the difficulties facing mental logic theory: (1) Johnson-Laird's illusory phenomenon requires an extension of mental logic theory to deal with possibility, which the current theory does not handle. (2) The phenomenon requires mental logic theory to modify the Principle of Truth via the mental model claim that people often represent only what is true but not what is false. This is difficult because the principle concerns the truth-values (true or false), but any current schema is used only "at its face value." (3) The phenomenon requires allowing some invalid schemas. It is conceptually difficult to swallow this from the classical metatheoretical point of view, because in classical logic and mathematics everything follows from a contradiction.

The second step involves a revision of the current version of the theory. For difficulty (3), we say invalid schemas should be allowed, because as psychologists spanning The Controversy, we have after all agreed that, in (untrained) human reasoning, nothing follows from contradiction. For (2), we revise the original definition of the inference schema due to Braine (1998) by adding one clause as italicized just below.

> "An inference schema specifies the form of an inference: Given information whose semantic representation, *or the representation of its semantic consequence*, has the form specified in the schema, one can infer the conclusion whose form is also specified. A semantic consequence does not have to be given."

Notice that Braine formulated inference schemas with certain *semantic elements*. For example, because his predicate logic is domain-specific, each quantified statement in a schema is with a given domain, which belongs to formal semantics in classical logic. Then we can define the **Understanding Conditional** as:

> The formula $p \Rightarrow x$ is read as p understandingly implies x, meaning x is a semantic consequence of p.

Accordingly, $p \Rightarrow x$ should be read as If-Then when x is given. For the third step, we can now define two illusory schemas as below. First we have

$$(p \rightarrow x) \text{ ori } (q \rightarrow x); p \text{ ore } q; \therefore x$$

where "ori" denotes inclusive disjunction and "ore" exclusive disjunction, The variable x may include negation or some modal operator, or both. For instance, here is a specific version of the schema accounting for illusory inferences of impossibility:

$$(p \Rightarrow \neg\Diamond s) \text{ ori } (q \Rightarrow \neg\Diamond s); p \text{ ore } q; \therefore \neg\Diamond s$$

Handled in this way, illusory schemas accommodate the Principle of Truth, and it is general enough for the proof-theoretic approach to account for the published illusory inferences caused mainly by exclusive disjunction, and thus it should also predict the same type of illusory inferences with compound statements. Some other illusory schemas would of course need to be developed to account for other possible type of illusory inferences (and corresponding work is required for the full adaptation of mental model theory as well).

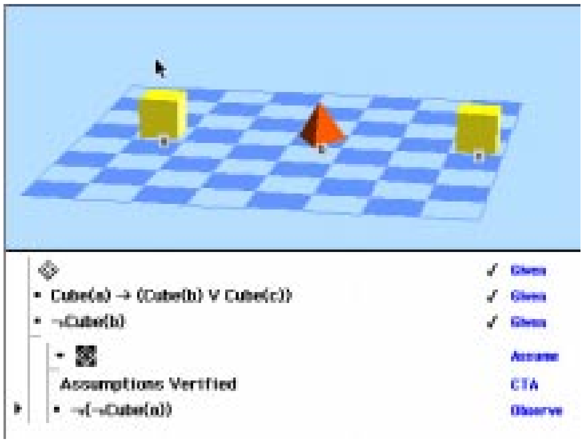

Figure 3: MML-Style Proof (Nonconsequence) in HYPERPROOF

7 Point 5 of The Agenda in (Brief!) Action

Recall from our framing discussion earlier in the paper that proofs establishing $\Phi \vdash \phi$, where Φ is a set of formulas in some logical system, and ϕ is one such formula, are entirely syntactic entities: they consist, essentially, of sequences of formulas chained together by rules of inference. However, there is a radically different sort of proof (rarely taught in elementary logic classes) which by their very nature embody part of the spirit of Mental MetaLogic. These are proofs of non-consequence; that is, to use the enabling machinery introduced earlier in this paper, they are proofs that establish facts of the form of $\Phi \nvdash \phi$. Now problems in which one must determine whether or not $\Phi \vdash \phi$ are thus rather interesting. On the one hand, if in fact ϕ can be derived from Φ by schemas and rules, then (in the first-order case) a computing machine can find a proof by churning away "at the level of mental logic." But on the other hand, if ϕ *can't* be derived from Φ in "mental logic fashion," then, in general, by definition there can't exist a mental logic-style proof of this fact. In addition, it's hard to see how a mental model style justification for this fact can be given, because such a justification must feature the (syntactic) form of the formulas in question. The answer is a proof of the sort that falls within Mental MetaLogic. Such a proof is approximated in Figure 3. This proof shows that from the first two formulas shown in the figure, one can't derive that ¬`Cube(a)`. The proof is shown in the system HYPERPROOF, which is well-suited to proofs having an MML flavor, and whose underlying mathematics is in line with MML as well (see Barwise & Etchemendy 1995). Unfortunately, HYPERPROOF is really only courseware, and so, as good as it is, it doesn't come close to realizing MML. We believe that the development of our general-purpose MML-based theorem prover (intended to be suitable for representing, solving, and (in part) generating reasoning problems of the sort seen on the GRE), now underway at RPI, will mature to the point of answering Q5a and Q5b.

8 Acknowledgments

Our MML research is currently supported by the National Science Foundation and the Educational Testing Service. We're grateful for this indispensable support. We're also greatly indebted to Phil Johnson-Laird for countless objections, insights, and suggestions.

References

Barwise, J. & Etchemendy, J. (1995), Heterogeneous logic, *in* J. Glasgow, N. Narayanan & B. Chandrasekaran, eds, 'Diagrammatic Reasoning: Cognitive and Computational Perspectives', MIT Press, Cambridge, UK, pp. 211–234.

Braine, M. (1998), Steps toward a mental predicate-logic, *in* M. Braine & D. O'Brien, eds, 'Mental Logic', Lawrence Erlbaum Associates, Mahwah, NJ, pp. 273–331.

Bringsjord, S. (1998), 'Is Gödelian model-based deductive reasoning computational?', *Philosophica* **61**, 51–76.

Bringsjord, S. & Bringsjord, E. (1996), 'The case against ai from imagistic expertise', *Journal of Experimental and Theoretical Artificial Intelligence* **8**, 383–397.

Bringsjord, S., Bringsjord, E. & Noel, R. (1998), In defense of logical minds, *in* 'Proceedings of the 20^{th} Annual Conference of the Cognitive Science Society', Lawrence Erlbaum, Mahwah, NJ, pp. 173–178.

Bringsjord, S. & Ferrucci, D. (2000), *Artificial Intelligence and Literary Creativity: Inside the Mind of Brutus, a Storytelling Machine*, Lawrence Erlbaum, Mahwah, NJ.

P. N. Johnson-Laird, P. L., Girotto, V. & Legrenzi, M. S. (2000), 'Illusions in reasoning about consistency', *Science* **288**, 531–532.

Pollock, J. (1995), *Cognitive Carpentry: A Blueprint for How to Build a Person*, MIT Press, Cambridge, MA.

Rinella, K., Bringsjord, S. & Yang, Y. (under review), 'Efficacious logic instruction: People are not irremediably poor deductive reasoners'.

Rips, L. (1994), *The Psychology of Prooof*, MIT Press, Cambridge, MA.

Stanovich, K. E. & West, R. F. (2000), 'Individual differences in reasoning: Implications for the rationality debate', *Behavioral and Brain Sciences* **23**(5), 645–665.

Yang, Y. (n.d.), How mental model theory can predict quantified schema-weights.

Yang, Y., Braine, M. & O'Brien, D. (1998), Some empirical justification of one predicate-logic model, *in* M. Braine & D. O'Brien, eds, 'Mental Logic', Lawrence Erlbaum Associates, Mahwah, NJ, pp. 333–365.

Yang, Y. & Bringsjord, S. (under review), 'The mental possible worlds mechanism: A new method for analyzing logical reasoning problems on the gre'.

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