PROCEEDINGS OF
THE 2007 INTERNATIONAL CONFERENCE ON
ARTIFICIAL INTELLIGENCE

ICAI 2007

Volume I

Editors

Hamid R. Arabnia
Mary Qu Yang, Jack Y. Yang

Associate Editors

Nikolaos G. Bourbakis, Halil Ceylan, Yuehui Chen
Georgios A. Demetriou, David Dodds, Laura L. Elnitski
Okan K. Ersoy, David de la Fuente, Joydeep Ghosh
Peter M. LaMonica, Guo-Zheng Li, Raymond A. Liuzzi
Jun Ni, Andrzej Niemierko, Jose A. Olivas, Bill Rand
Gene Simmons, Ashu M. G. Solo, Andrew Sung
Sule Yildirim, Yanqing Zhang

WORLDCOMP'07
June 25-28, 2007
Las Vegas Nevada, USA
www.world-academy-of-science.org

The Multi-Mind Effect

Selmer Bringsjord[1], Konstantine Arkoudas[2], Deepa Mukherjee[3]
Andrew Shilliday[4], Joshua Taylor[5], Micah Clark[6], Elizabeth Bringsjord[7]
Department of Cognitive Science[1−6]; Department of Computer Science[1,4,5]
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
State University of New York[7]; SUNY Plaza
Albany NY 12246 USA

Abstract *Courtesy of experiments carried out by such thinkers as Wason, Johnson-Laird, and Kahneman & Tversky, there is overwhelming empirical evidence that the vast majority of logically untrained humans are unable to reason in context-independent, normatively correct fashion. However, the* multi-mind effect*, which is predicted by our earlier success at teaching this kind of reasoning, and also by our general theory of human and machine reasoning, shows that while individual persons (with rare exceptions) are unable to solve problems that demand context-independent reasoning,* groups *of persons can often solve such problems.*

Keywords: multi-mind effect, heterogeneous reasoning, multi-agent reasoning, logic-based computational cognitive modeling

1 Introduction

Experimental study of human reasoning has shown that exceedingly few humans can solve problems demanding normatively correct, context-independent reasoning [1]. Two theories in the field of psychology of reasoning, mental logic (ML) [2] and mental models (MM) [3], both predict such failures, giving the same general explanation: humans generally lack the mental machinery required to solve such problems. Predicted failures include phenomena such as *illusory inferences* [4, 5], in which subjects "see" logically valid inferences that simply aren't there.

In earlier work, we have shown that education of a certain kind in the area of formal logic, specifically education in accordance with our theory of human reasoning, *mental meta-logic* (MML) [6, 7, 8, 9], *contra* the claims of some well-known psychologists (e.g., [10]), can produce humans able to negotiate problems demanding context-independent, normatively correct reasoning [11, 7]. We say that such humans acquire *logical minds*, and hold that our work vindicates Piaget to a considerable degree.[1]

Unfortunately, the training required for an individual to reach this level must take place over an extended period of time, and there is no reason to believe that this individual, without ongoing practice, would retain her hard-won ability.

Nonetheless, the possibility remains that normatively correct reasoning might be possible to achieve in a different manner. But how?

While it's indeed true that the vast majority of individuals are unable to solve problems that require context-independent reasoning (unless suitably trained), *groups* of individuals acting together after being stimulated under the right circumstances *can* often solve such problems, even in the absence of extended training. This result is what we call the *multi-mind effect* (MME).

2 Related Research

Of course, it has long been known that groups can out-perform individuals.[2] However, one must distinguish between groups of individuals that can reason in normatively correct, context-independent fashion to solve so-called "unsolvable" problems (as in

[1]Piaget, as is well known, held that in the course of normal development humans would acquire a capacity to think in accordance with first-order logic [12]. Though Piaget's position has fallen out of favor, with sufficient training in formal logic, humans can in fact *exceed* the level of reasoning Piaget called "formal operations," and reach the level in which they can reason in *many* logical systems, as well as about such systems. In this level, humans can also *create* logical systems. In general, this level of reasoning is reached by professionals in the formal sciences.

[2]E.g., in [13] it's shown that cooperative and collaborative learning in mathematics is very effective, and in [14] it's shown that group performance is generally qualitatively and quantitatively superior to the performance of the average individual.

MME), versus groups of individuals who engage in generic problem-solving as a team. Again, there has been extensive research into the working of teams and groups' ability to solve problems effectively and efficiently in a wide variety of fields [15, 16]. The benefits of groups of individuals working together have been documented in areas as diverse as open-source software development [17] and prediction and forecast models [18]. But in all these scenarios, the problems that the groups are working on are readily solvable even without the group, albeit with greater expense and effort. We are unaware of prior results that demonstrate MME.[3]

3 Dearth of C-I Reasoning

Studies of human reasoning have repeatedly shown that logically untrained humans systematically fail to reason in a context-independent manner, even when presented with stimuli that expressly call for this type of reasoning. This failure has been attributed, in part, to the lack of the appropriate reasoning machinery in humans. For example, Rips (1994) has claimed that our mental apparatus is a *partial* selection from the rules and schemas available in standard proof calculi for the propositional calculus. According to this view, humans who have not been extensively trained in logic cannot reason over certain problems accurately, as they simply do not have the inference rules to do so. Empirical evidence for this view has been derived from experiments involving stimuli solvable in normatively correct fashion only if a standard proof theory for propositional calculus is correctly exploited. For example:

> **Problem 1**: Assume that (1) It is false that 'If the square is green, the circle is red'. Given this assumption, can you infer that the square is green?

As Rips (1994) and other proponents of ML have shown, nearly all humans, working as individuals, answer "No" — but in fact the correct response is an affirmative. The reason is that in the propositional calculus, when a (material) conditional of the form 'if ϕ then ψ' is false, ϕ is true while ψ is false. The explanation from Rips is that the inferential rules that would support an affirmative response are simply not part of the reasoning apparatus of humans untrained in formal logic. In other words, ML holds that untrained reasoners reason on the basis of a collection of proof-theoretic rules incomplete from the standpoint of (e.g.) standard extensional formal logic. Rips' (1994) PSYCOP system has a set $\mathcal{R}$ of *incomplete* inference rules for the propositional calculus. For example, the formula $\neg(p \rightarrow q) \rightarrow (\neg p \wedge q)$ isn't provable by $\mathcal{R}$, but this formula, in any standard proof calculus for propositional logic, admits of effortless proof.

Next, consider this somewhat more complicated problem, a slight variant[4] of a puzzle introduced in [19].

> Assume that the following is true:
> 'If there is a king in the hand, then there is an ace in the hand,' or 'If there is not a king in the hand, then there is an ace in the hand,' — but not both of these if-thens are true.
> What can you infer from this assumption? Please provide a careful justification for your answer.

Nearly all untrained subjects, working individually, declare the answer to be: "There is an ace in the hand." Unfortunately, what one *can* infer is that there *isn't* an ace in the hand. This logical illusion is successfully predicted by MM [4, 3], which holds that reasoners conceive of situations containing true premises, but not false premises, and that reasoners are not comfortable reasoning from a false antecedent, and hence do not explicitly represent a false possibility. To correctly solve this illusion one must note the exclusive disjunction. Using obvious symbolization, the given information becomes:

$$((K \rightarrow A) \vee (\neg K \rightarrow A)) \wedge \neg((K \rightarrow A) \wedge (\neg K \rightarrow A))$$

It's not hard to prove in any proof calculus for standard first-order logic that $\neg A$ can be derived from this formula.

4 Basic Experimental Design

The multi-mind effect is predicted by the aforementioned MML theory [6, 7, 8, 9]. MML predicts that groups of agents, appropriately tasked, will often engage in *heterogeneous reasoning*: For example, they will *meta*-reason about the different kinds of reasoning offered by different individuals.

[3] Some well-read readers may wonder whether an effect studied at the RAND Corporation during the Cold War anticipates MME. Known commonly as the *Delphi effect* (DE), it has been leveraged in a methodology by the same name and has been (and continues to be) used extensively in prediction and forecast models in a wide variety of settings [18]. DE has often been cited in the field of open-source software development [17], where many individuals come together to create sophisticated software that they would not be able to achieve individually. Attempts to exploit DE are apparently ongoing in the field of open-source software development, through organizations like the Creative Commons. This being said, MME is clearly distinct from DE.

[4] The variation arises from disambiguating Johnson-Laird's 's or else s'' as 'either s or s', but not both.'

They will also exploit both the proof-theoretic techniques stressed by ML and the model-based techniques stressed by MM; and they will move back and forth between these techniques. In addition, groups will often reason in logical systems much more powerful than those to which mental models and mental logic are anchored. These, at any rate, are the broad-stroke predictions.[5] Do the predictions hold? Using the following preliminary experimental method, one can seek relevant data.

Suppose that P_{ML} and P_{MM} are (respectively) such that: P_{ML} is a problem subjects cannot solve according to ML, and P_{MM} is a problem subjects cannot solve according to MM.

Next, let us verify the situation in traditional experimental fashion. That is, let us give n (logically untrained) subjects a problem P_{XX} (where this notation ranges over both mental logic and mental models), with the (guaranteed) result being that only a very few of these subjects solve this problem. Next, we remove the subjects who are successful. At this point, we give the subjects who remain, all of whom are now known, by ML and MM, to be inherently incapable of solving P_{XX}, a problem P'_{XX} that is formally isomorphic to P_{XX}. However, we instruct these remaining subjects to work on P'_{XX} in randomly assigned groups. The size of the groups will depend upon how many subjects remain. MML, and our earlier work on de-biasing through logic training, predict the surprising result that P'_{XX} will be solved by some groups, despite the fact that all groups are composed only of individuals who succumbed to bias in the original case. This immediately implies that neither proponents of ML nor proponents of MM can afford to stay silent, since, after all, the"multi-agents" still lack, by these theories, the mental machinery needed to solve P'_{XX}.

Our hypothesis is that the sophisticated reasoning of highly successful and accurate individual reasoners is distinguished by heterogeneous reasoning. Such reasoning is carried out by some individuals working alone in accordance with processes that are multi-agent in nature. Our approach can be contrasted with the emergentist position adumbrated and advocated in [20], which is that collective cognition isn't reducible to individual cognition (collective cognition, to use their term, is *autonomous*), while it is produced out of individuals interacting. In our case, a precise account of the multi-mind effect would enable an individual problem solver to leverage that account to achieve performance on par with the group. In fact, we hypothesize that the very highest levels of individual problem solvers encompass techniques and patterns of thought seen in MME.

[5]Due to space constraints, we leave aside MML's predictions regarding the general failure of *individuals* faced with logical illusions and the like.

5 Initial Supportive Results

Three pilot experiments were conducted by the first author to test for the existence of the predicted MME. The aim of these experiments was to explore MME as a phenomenon occurring in multi-agent reasoning. We also wanted to explore the techniques by which groups of reasoners could leverage the cognitive apparatus of the individual members to come up with performance that was greater than the performance of any (or the best) individual in the group. The problems used were variants of Problems 1 and 2 from above, the famous Wason card selection task (WST) [21], and the Wise Man Puzzle (WMP), which is well-known in logic-based AI.[6] The experiments were carried out on students in the course *Logic and Artificial Intelligence*. These students had taken less than 3 courses in formal logic, and therefore may be considered 'untrained.'[7]

The group size originally was 13. The students were presented with the problems and asked to solve them individually. Once the responses were collected, students who gave a correct response to the problems were separated from the rest of the class. The students were not told whether their responses were accurate or not, to prevent any biases. A cover story was presented to explain the separation of some the students from the rest of the class. Only one subject was removed in the three non-WMP cases. No individual solved WMP (under severe time constraints). The remaining students were randomly assigned to different groups, with four groups of three students per group. The groups were given problems isomorphic to Problems 1 and 2, WST, and WMP. The groups were all asked for justifications for their solutions. These could be in the form of sentences, proofs, or diagrams, and permutations thereof. All the groups reached the correct solution for the isomorphic problem for Problems 1 and 2 and WST. One of the three groups solved WMP. These results, though extremely preliminary, show support for the presence MME.

[6]For a detailed analysis of WMP, see [22].

[7]In the interests of space, we leave aside discussion of, and evidence for, the view that the number 3 is the dividing line on this issue.

6 Toward Modeling MME

6.1 Declarative CCM

The basic units of declarative computational cognitive modeling (CCM) are declarative in nature, or propositional: they are formal objects naturally associated with those particular sentences or expressions in natural languages (like English, German, Chinese) that are declarative statements (as opposed to expressions in the imperative or inquisitive mode) naturally taking values such as TRUE, FALSE, UNKNOWN, PROBABLE (sometimes to particular numerical degrees), and so on. The basic process over such units is inference, which may be deductive, inductive, probabilistic, abductive, or analogical. Because the basic units of declarative CCM are declarative, a hallmark of this type of modeling is a top-down, rather than bottom-up, approach.[8]

6.2 Logic-Based CCM

As explained in [24], logic-based CCM is the formal kind of modeling that underlies top-down, declarative modeling, and is based on a generalized form of the concept of *logical system* as defined rather narrowly in mathematical logic, where this concept stands at the heart of Lindström's Theorems [25]. Corresponding to every logical system is a logic-based computer program. The execution (or, better, the evaluation) of such a program produces a cognitive model or simulation of the target phenomena. The target in the present case, of course, is the multi-mind effect.

6.3 Prior Logic-Based Modeling

Arkoudas & Bringsjord (2004) have carried out previous work designed to simulate, from an engineering/AI point of view, multi-agent reasoning of the sort required to solve WMP. There is insufficient space to even encapsulate this work, but it should be noted that while it was undertaken outside the perspective of MME, the core formalisms [26, 27] were in line with logic-based CCM. It is also worth mentioning that some prior work seems to us to indicate that MME explains why mathematics involves the interaction of multiple agents, for example see [28].

[8] Modeling carried out by [23] is an exception, since it is at once bottom-up and top-down.

6.4 Toward Modeling MME in Slate

Slate is a system for modeling and facilitating human reasoning and decision-making that has been under development (led by Shilliday, Taylor, and Bringsjord) for the past five years in RPI's RAIR Lab, the research and development made possible by grants from ARDA, DTO, and DARPA. A user of Slate manipulates a number of information-based items, viz., propositions, hypotheticals, sets (i.e., collections of other information-based items), subproofs, models, proofs, and, in the most recent versions of Slate, databases. All of these structures are represented graphically in Slate's workspace within *System S*. System S empowers users of Slate to record and share their inferences and reasoning structures, and allows mechanical treatment of these structures by machine, so that Slate can *automatically* launch searches for proofs and disproofs in connection with arguments under consideration by human reasoners.

6.4.1 On Automated Translation

We are currently working on various tools and frameworks in the RAIR Lab to facilitate interoperability between Slate and other systems. A number of languages have been developed for the purpose of sharing knowledge and expressing the relation between the knowledge representation schemata of various systems. KIF [29] was the first *interlingua* adopted on a (relatively speaking) large scale, but after a few years (during which the internet became much more important), Common Logic [30] was created to address the need for namespaces, URIs and the like. Because of the proliferation of unintegrated knowledge bases and relational databases, DTO launched the IKRIS challenge workshop in April 2005 for the development of semantic interoperability between such systems — a workshop in which we participated. In connection with this work, we have devised a system of *translation graphs* capable of yielding so-called "bridging axioms" to enable semantic interoperation between various divergent representation schemata or ontologies (encoded as *signatures* in Many-Sorted Logic (MSL) [31], Slate's "native" language).

6.4.2 Mental Metalogic Reasoning in Slate

In Slate, items in System S are connected with argument links to graphically depict an argument from some set of premises to a particular conclusion. Arguments may be supported or denied by *witness* objects, viz., models, proofs, or databases.

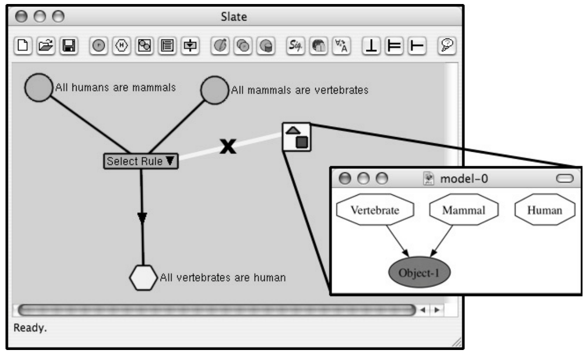

Figure 1: Slate simulates mental model-based reasoning.

This mechanism can be used to model model-based reasoning in Slate. For example, consider the following argument:

	(P1)	All humans are mammals.
	(P2)	All mammals are vertebrates.
∴	C	All vertebrates are human.

The easiest way to realize that this argument is invalid is to imagine a *model* in which the premises are true, and yet the conclusion is false. Consider a model in which there is a single object that is a vertebrate, a mammal, and yet not a human. (P1) is satisfied: All humans are mammals (as there are no humans at all in the model). And (P2) is satisfied: All mammals are vertebrates (as the one object that is a mammal is also a vertebrate.) However, (C) is false: There is some vertebrate which is not a human (the single object in the model is a vertebrate, and yet not a human.)

This entire process, which some subjects engage in, can be simulated by Slate; the process is summed up in Figure 1. Each premise and conclusion is represented iconically in System S, and the connection between these objects reflects the inferential structure of the argument given above. The (mental) model can be inspected graphically (the inset image), as well as manipulated iconically in the System S. The iconic model is connected to the argument with an X, denoting that the model shows the argument to be invalid.

6.4.3 Multi-Agent Reasoning in Slate

Slate can be used to model multi-agent reasoning analogous to the interactions between actual human reasoners. To get a glimpse of how this is possible, consider the following scenario.

Four telemarketing companies, A, B, C, and D, each of which we can regard to be an agent, have decided, mostly due to their highly developed sense of morality, that they should collaborate with each other to ensure that the potential customers are contacted no more than by two of the companies within a two month period. Each firm keeps a database of those people they have contacted, and would like to provide this data, while revealing neither sensitive data nor the structure of their own records, to the other companies. Each firm also wants to receive such data from the other marketers. Adding to the difficulty, not only are sensitive data and record structures to be kept private, but the information which *can* be shared is stored *differently* from firm to firm:

- Firm A maintains records only on whether they have called a particular household within a given month.
- Business B considers a contact made when they have telephoned a household and the household has made a call *back* to them.
- Company C considers having made a contact with a household if and only if they called the household and the phone call lasted at least fifteen minutes.
- D considers a contact made when *either* they have called a household *or* the household has called them.

Though the systems have different notions of contacting households, an intermediate office I works in conjunction with the four companies to help them determine who they can call, as follows:

I records the initiating party, the receiving party, the date, and the duration of phone calls. Each firm is able to give this information to I, and when a firm wants to call a household, they first ask I whether they are permitted to call the household. I has at its disposal information from each system and so can determine whether it is permissible to call the household. Note that no system has to reveal its internal structures or sensitive information, but that each system, with the help of I, can behave appropriately.

Given *translation graphs*, the relationships between the representations used by the different com-

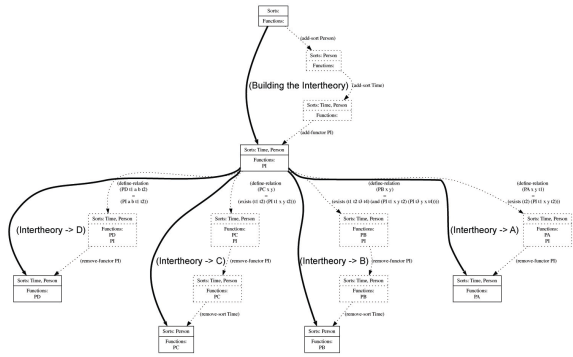

Figure 2: Slate models four agents interacting through an *interlingua*.

panies (and more generally, agents), can be explored in Slate, and a process for reconciling the representations constructed. A set of *bridging axioms* can be extracted automatically from this translation graph. Such a graph is shown in Figure 2.

While this example relates to the economic marketplace on its surface, the underlying structures and processes, we believe, are subtle, and are part of what is required to model the multi-mind effect, in which individual humans interact through an *interlingua* to problem-solve.

7 M-M Effect & Education

The ultimate dream of passionate educators would presumably be nothing short of teaching students how to surmount context *de*pendent reasoning. Our basic strategy would be to teach *individual* students how to engage in the efficacious forms of reasoning seen when MME is produced. This strategy, arguably, would stretch back to Plato, who defines thinking as inner (silent) dialogue with oneself.[9] We specifically hold that education in logic is the key to developing context-independent deductive reasoning in students [11]. Such education should include instruction in the following, which we have observed to be in play in MME:

1. *Dis*proving an incorrect answer, and proving that a purported disproof fails, which reinstates the original, targeted would-be proof.
2. Rigorous and general-purpose methods for transforming a natural language (e.g., English) word problem into a formal representation in a logical system.
3. Using diagrammatic techniques to model proofs and disproofs, and meta-reasoning over these meta-representations.

[9]E.g., in *Theaetaus* (189E-190A). He gives similar accounts of thought in the *Sophist* (263E), and in *Philebus* (38E).

8 Next Steps

The empirical evidence presented in this paper is a result of some preliminary studies conducted to test our theory-driven prediction of MME. Further experiments are planned in which this effect will be investigated in detail, in a controlled manner.[10] Our research program is also aimed at precisely modeling MME in Slate, under the paradigm of logic-based CCM, a kind of modeling we have achieved previously, for example in [22].[11] Progress toward these goals will be reported and demonstrated at ECM_MAI 2007.

References

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

[2] L. Rips. *The Psychology of Proof.* MIT Press, Cambridge, MA, 1994.

[3] P. N. Johnson-Laird. *Mental Models.* Harvard University Press, Cambridge, MA, 1983.

[4] P. N. Johnson-Laird, P. Legrenzi, V. Girotto, and M. S. Legrenzi. Illusions in reasoning about consistency. *Science*, 288:531–532, 2000.

[10]However, the only empirical issue open to dispute is whether the multi-mind effect ensures de-biasing. We know the effect occurs, for the simple reason that verbal protocols we have collected invariably show it in action when groups, as opposed to individuals, grapple with relevant problems.

[11]We have stressed reasoning in the foregoing, but the multi-mind effect should surface in connection with decision-making. Put simply, the main reason is that we believe can shortly show, formally and computationally, that the seminal experiments in [32], which reveal a failure of normatively correct decision making, are circumvented by artificial agents that reason in the normatively correct fashion underlying MME.

[5] S. Bringsjord and Y. Yang. Logical illusions and the welcome psychologism of logicist artificial intelligence. In D. Jacquette, editor, *Philosophy, Psychology, and Psychologism: Critical and Historical Essays on the Psychological Turn in Philosophy*, pages 289–312. Kluwer, Dordrecht, The Netherlands, 2003.

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

[7] K. Rinella, S. Bringsjord, and Y. Yang. Efficacious logic instruction: People are not irremediably poor deductive reasoners. In J. D. Moore and K. Stenning, editors, *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, pages 851–856. Lawrence Erlbaum Associates, Mahwah, NJ, 2001.

[8] Y. Yang and S. Bringsjord. Mental metalogic: A new paradigm for psychology of reasoning. In *Proceedings of the Third International Conference on Cognitive Science (ICCS 2001)*, pages 199–204. Press of the University of Science and Technology of China, Hefei, China, 2001.

[9] Y. Yang and S. Bringsjord. The mental possible worlds mechanism and the lobster problem: an analysis of a complex gre logical reasoning task. *Journal of Experimental and Theoretical Artificial Intelligence*, 18(2):157–168, 2006.

[10] P. W. Cheng and K. J. Holyoak. Pragmatic versus syntactic aproaches to training deductive reasoning. *Cognitive Psychology*, 17:391–416, 1985.

[11] S. Bringsjord, E. Bringsjord, and R. Noel. In defense of logical minds. In *Proceedings of the 20^{th} Annual Conference of the Cognitive Science Society*, pages 173–178. Lawrence Erlbaum, Mahwah, NJ, 1998.

[12] B. Inhelder and J. Piaget. *The Growth of Logical Thinking from Childhood to Adolescence*. Basic Books, New York, NY, 1958.

[13] N. Davidson and T. Worsham. *Enhancing Thinking Through Cooperative Learning*. Teachers College Press, New York, NY, 1992.

[14] G. Hill. Group versus individual performance: Are $n + 1$ heads better than one? *Psychological Bulletin*, 91(3):5517–539, 1982.

[15] J. Richard Hackman, Ruth Wageman, Thomas M. Ruddy, and Charles L. Ray. *Industrial and Organizational Psychology*, chapter Team Effectiveness in Theory and Practice, pages 109–119. Blackwell Publishers Inc, Malden, MA, 2000.

[16] Eduardo Salas, C. Shawn Burke, and Janis A Cannon-Bowers. Teamwork: Emerging principles. *International Journal of Management Reviews*, 2(4):339–356, 2000.

[17] Eric Raymond. The cathedral and the bazaar. *Knowledge, Technology & Policy*, 12, 1999.

[18] Harold A. Linstone and Murray Turoff. *The Delphi Method: Techniques and Applications*. Addison Wesley Publishing Company, Reading, MA, 1975.

[19] P.N. Johnson-Laird. Rules and illusions: A criticial study of Rips's *The Psychology of Proof*. *Minds and Machines*, 7(3):387–407, 1997.

[20] Pietro Panzarasa and Nicholas Jennings. Collective cognition and emergence in multi-agent systems. In Ron Sun, editor, *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*, pages 401–408. Cambridge University Press, Cambridge, UK, 2006.

[21] P. Wason. Reasoning. In *New Horizons in Psychology*. Penguin, Hammondsworth, UK, 1966.

[22] Konstantine Arkoudas and Selmer Bringsjord. Metareasoning for multi-agent epistemic logics. In *Fifth International Conference on Computational Logic In Multi-Agent Systems (CLIMA 2004)*, volume 3487 of *Lecture Notes in Artificial Intelligence (LNAI)*, pages 111–125. Springer-Verlag, New York, 2005.

[23] Ron Sun. Robust reasoning: Integrating rule-based and similarity-based reasoning. *Artificial Intelligence*, 75:241–295, 1995.

[24] Selmer Bringsjord. Declarative/logic-based computational cognitive modeling. In Ron Sun, editor, *The Handbook of Computational Cognitive Modeling*. Cambridge University Press, forthcoming.

[25] H. D. Ebbinghaus, J. Flum, and W. Thomas. *Mathematical Logic (second edition)*. Springer-Verlag, New York, NY, 1994.

[26] K. Arkoudas. Athena. http://www.cag.csail.mit.edu/∼kostas/dpls/athena.

[27] Lou Goble, editor. *The Blackwell Guide to Philosophical Logic*. Blackwell Publishing, Oxford, UK, 2001.

[28] S. Colton, A. Bundy, and T. Walsh. Agent based cooperative theory formation in pure mathematics, 2000.

[29] Michael R. Genesereth and Richard E. Fikes. *Knowledge Interchange Format Version 3 Reference Manual*, December 1997.

[30] Christopher Menzel. Common Logic Standard. In *Santa Fe 2003 Metatdata Forum Symposium on Ontologies*, 2003. Presentation of the case for Common Logic.

[31] Maria Manzano. *Extensions of First Order Logic*. Cambridge University Press, 1996.

[32] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 17(2):263–292, 1979.