REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	01/26/95	Final Technical 09/15/91 – 11/14/94

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
Reasoning under uncertainty	C F49620-91-C-0080

6. AUTHOR(S)

Dr. Joseph Y. Halpern

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
IBM Corporation Almaden Research Center 650 Harry Road San Jose, CA 95120-6099	AFOSR-TR· 9 ɔ - 0 2 6 8

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
Air Force Office of Scientific Research AFOSR 110 Duncan Avenue Suite B115 Bolling AFB, DC 20322-0001	

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Unlimited

This document has been approved for public release and sale; its distribution is unlimited.

12b. DISTRIBUTION CODE
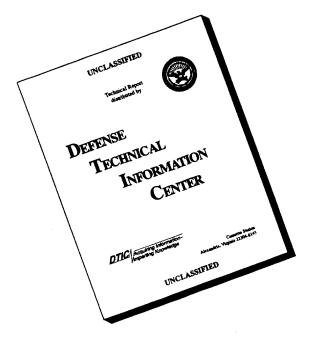
F

DTIC SELECTED
JUN 1 9 1995

19950615 057

13. ABSTRACT (Maximum 200 words)

This report summarizes the research performed under this contract. The research can be divided into four areas: (1) Computing degrees of belief, (2) Reasoning about knowledge and communication, (3) Nonmonotonic reasoning, and (4) Belief change. Topic (1) is concerned with how an agent such as a doctor should assign degrees of belief to events such as "Tom has hepatitis" given a large knowledge base of first-order and statistical information (such as "90% of patients with jaundice have hepatitis"). The research in topic (2) focused on how knowledge changes as a result of communication and finding models for algorithmic (computable) knowledge. The work in topic (3) involved an analysis of the notion of "only knowing" and the introduction of a new tool to model model likelihood, plausibility measures. Finally, the work on topic (4) focused on finding appropriate models that could capture belief, plausibility, and time, and showing how they could be used to provide a deeper understanding of earlier work, as well as providing new results. Other highlights of the contract period are also listed, and a list of papers published under the contract is appended.

DTIC QUALITY INSPECTED 5

14. SUBJECT TERMS	15. NUMBER OF PAGES
Keywords: degrees of belief, nonmonotonic reasoning, belief change, reasoning about knowledge, reasoning about probability, qualitative reasoning, statistical inference.	7
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

DISCLAIMER NOTICE

UNCLASSIFIED

Technical Report
distributed by

DEFENSE
TECHNICAL
INFORMATION
CENTER

DTIC Acquiring Information-
Imparting Knowledge

Cameron Station
Alexandria, Virginia 22304-6145

UNCLASSIFIED

THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE
COPY FURNISHED TO DTIC
CONTAINED A SIGNIFICANT
NUMBER OF PAGES WHICH DO
NOT REPRODUCE LEGIBLY.

GENERAL INSTRUCTIONS FOR COMPLETING SF 298

The Report Documentation Page (RDP) is used in announcing and cataloging reports. It is important that this information be consistent with the rest of the report, particularly the cover and title page. Instructions for filling in each block of the form follow. It is important to *stay within the lines* to meet *optical scanning requirements*.

Block 1. Agency Use Only *(Leave blank)*.

Block 2. Report Date. Full publication date including day, month, and year, if available (e.g. 1 Jan 88). Must cite at least the year.

Block 3. Type of Report and Dates Covered. State whether report is interim, final, etc. If applicable, enter inclusive report dates (e.g. 10 Jun 87 - 30 Jun 88).

Block 4. Title and Subtitle. A title is taken from the part of the report that provides the most meaningful and complete information. When a report is prepared in more than one volume, repeat the primary title, add volume number, and include subtitle for the specific volume. On classified documents enter the title classification in parentheses.

Block 5. Funding Numbers. To include contract and grant numbers; may include program element number(s), project number(s), task number(s), and work unit number(s). Use the following labels:

C	-	Contract	PR - Project	
G	-	Grant	TA - Task	
PE	-	Program Element	WU - Work Unit Accession No.	

Block 6. Author(s). Name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. If editor or compiler, this should follow the name(s).

Block 7. Performing Organization Name(s) and Address(es). Self-explanatory.

Block 8. Performing Organization Report Number. Enter the unique alphanumeric report number(s) assigned by the organization performing the report.

Block 9. Sponsoring/Monitoring Agency Name(s) and Address(es). Self-explanatory.

Block 10. Sponsoring/Monitoring Agency Report Number. *(If known)*

Block 11. Supplementary Notes. Enter information not included elsewhere such as: Prepared in cooperation with...; Trans. of...; To be published in.... When a report is revised, include a statement whether the new report supersedes or supplements the older report.

Block 12a. Distribution/Availability Statement. Denotes public availability or limitations. Cite any availability to the public. Enter additional limitations or special markings in all capitals (e.g. NOFORN, REL, ITAR).

DOD - See DoDD 5230.24, "Distribution Statements on Technical Documents."
DOE - See authorities.
NASA - See Handbook NHB 2200.2.
NTIS - Leave blank.

Block 12b. Distribution Code.

DOD - Leave blank.
DOE - Enter DOE distribution categories from the Standard Distribution for Unclassified Scientific and Technical Reports.
NASA - Leave blank.
NTIS - Leave blank.

Block 13. Abstract. Include a brief *(Maximum 200 words)* factual summary of the most significant information contained in the report.

Block 14. Subject Terms. Keywords or phrases identifying major subjects in the report.

Block 15. Number of Pages. Enter the total number of pages.

Block 16. Price Code. Enter appropriate price code *(NTIS only)*.

Blocks 17. - 19. Security Classifications. Self-explanatory. Enter U.S. Security Classification in accordance with U.S. Security Regulations (i.e., UNCLASSIFIED). If form contains classified information, stamp classification on the top and bottom of the page.

Block 20. Limitation of Abstract. This block must be completed to assign a limitation to the abstract. Enter either UL (unlimited) or SAR (same as report). An entry in this block is necessary if the abstract is to be limited. If blank, the abstract is assumed to be unlimited.

Summary of research under AFOSR Contract #F49620-91-C-0080

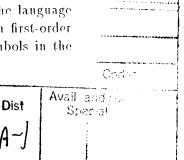**Personnel:** Joe Halpern (PI), Nir Friedman, Adam Grove, Daphne Koller

The research performed under this contract can be roughly divided into four areas:

Computing degrees of belief: Consider an agent (or expert system) with a knowledge base KB that includes statistical information (such as "90% of patients with jaundice have hepatitis") who would like to assign a *degree of belief* to a particular statement φ (such as "Tom, a patient with jaundice, has hepatitis"). For example, a doctor may want to assign a degree of belief that a patient has a particular illness based on the symptoms exhibited by the patient, together with general information about symptoms and diseases. Since the actions the agent takes may depend crucially on this degree of belief, we would like techniques for computing degrees of belief in a principled manner. We have investigated three methods of doing so. All three methods are applications of the "principle of indifference", because they assign equal degree of belief to all basic "situations" consistent with the knowledge base. They differ because there are competing intuitions about what the basic situations are. Various natural patterns of reasoning, such as the preference for the most specific statistical data available, turn out to follow from some or all of the techniques. This is an improvement over earlier theories, such as work on direct inference and reference classes (going back to the 1930's) which arbitrarily postulate these patterns without offering any deeper explanations or guarantees of consistency.

The three methods we investigate have surprising characterizations, that show their relationships to other approaches to obtaining degrees of belief. For example, one approach is strongly related to maximum entropy (and, indeed, can be viewed as a generalization of it), while another can be shown, roughly speaking, to maximize independence among predicates as much as possible. There are also unexpected connections between the three approaches, that help us understand why the specific language chosen (for the knowledge base) is much more critical in inductive reasoning of the sort we consider than it is in traditional deductive reasoning.

One of the applications of this framework is to default reasoning. Suppose we capture a default such as "typically birds fly" as the statistical statement "almost 100% of birds fly". If we include such defaults in the KB, then we automatically get some of the properties that have been most sought-after in default formalisms, including preference for more specific information and ability to ignore irrelevant information. The methods can also be extended to deal with causal reasoning in a natural way. Thus, our approach can be used for a wide spectrum of reasoning: quantitative, qualitative, and causal.

All three methods mentioned above can be viewed as compatible with a Bayesian approach, which says start with a prior probability and condition on whatever information you receive. The problem with the Bayesian approach is that it does not tell us what space to take a prior over, or what that prior should be. Our work can be viewed as answering these questions. Once we fix our *language* (where, technically, the language consists of constant symbols, function symbols, and predicate symbols, as in first-order logic), a *possible world* just defines a possible interpretation for all the symbols in the

language. That is, if the language includes the predicate symbols *Hepatitis* and *Jaundice*, and the constant *Tom*, an interpretation describes which individuals have hepatitis, which ones have jaundice, and which one is Tom. We put a prior on the set of possible worlds. Our focus has been on priors that satisfy some uniformity conditions, and our results show that they have quite attractive properties. However, our results also show that they are not always appropriate. More recently, rather than considering a particular prior, we have considered the question of what properties of priors give us desirable properties of beliefs (such as allowing us to ignore irrelevant information). We are also beginning to consider algorithmic aspects of this approach.

Papers A2, A3, A4, A6, B1, B2, B3, B4, B5, B6, B9, B16, B17, C3, and C5 summarize our work in this area.

Reasoning about knowledge and communication: A recent trend in analyzing complicated systems has been to think in terms of one agent knows about what other agents know, and how this knowledge changes as a result of communication. Under this contract, we have (a) shown that this approach (and the formal model of systems that goes along with it) can be used to capture aspects of knowledge bases more elegantly and concisely than the traditional axiomatic approach, (b) provided a formal logic for analyzing communication and knowledge change, and (c) provided a logic that captures the more computational aspects of reasoning about knowledge. This latter problem is a longstanding one, since it is well known that the standard models for reasoning about knowledge, using possible-world structures, suffer from the problem of *logical omniscience*: agents know all tautologies, and the logical consequences of their knowledge. Our model captures the algorithms used by agents to compute their knowledge in a direct way; most of the previous attempts to solve the problem can be embedded in our solution in a natural way. Thus, our approach can be viewed as a generalization of them all. The work in this area is summarized in papers A4, B5, B10, and B15.

Nonmonotonic reasoning: While a great deal of work has been done on nonmonotonic reasoning, almost all of it has assumed that there is only a single agent. It has been a longstanding problem to extend some of the standard approaches to nonmonotonic reasoning to multiple agents. This problem has now been solved for two approaches to "only knowing", due to Levesque and to Halpern and Moses. The approach should extend naturally to Moore's autoepistemic logic. In addition, by going back to the original definitions of various approaches to nonmonotonic reasoning, we were able to acquire a deeper understanding of sources of potential problems.

More recently, work was done trying to provide a general model for nonmonotonic reasoning. Many nonmonotonic logics have been proposed, and there is little agreement as to what is the "right" one, or whether there is any "right" one. There is some agreement, on the other hand, on some basic "core" properties for nonmonotonic reasoning. These properties have been called the KLM properties (for Kraus, Lehmann, and Magidor). Many different approaches to nonmonotonic reasoning have been proposed, all of which have been shown to be characterized by the KLM properties in the propositional case: including preferential logic, ϵ-semantics, possibility theory, and κ rankings. This has been viewed as quite surprising. Recently, we have proposed a qualitative way of capturing

2

likelihood that we call *plausibility measures*. It is straightforward to show that all the approaches to giving semantics to nonmonotonic reasoning that can be characterized by the KLM properties can be easily mapped into plausibility spaces. Moreover, a single property of plausibility is provided that is necessary and sufficient to guarantee that the KLM properties are sound, and another is weak property is provided that guarantees that they are complete. Since all the known approaches are easily seen to have both properties, it is no longer so mysterious that they can be characterized by the KLM properties.

Things change significantly once we move to the first order case. Using the insights afforded by thinking in terms of plausibility, we can show that all the approaches alluded to above other than ϵ-semantics lack significant expressive power once we introduce first-order quantification. These results emphasize the naturalness and appropriateness of an approach based on plausibility.

The work in this area is summarized in papers A7, B7, B8, C1, and C6.

Belief change: The study of *belief change* has been an active area in philosophy and in artificial intelligence and, more recently, in game theory. The focus of this research is to understand how an agent should revise his beliefs as a result of getting new information. In the literature, two instances of this general phenomenon have been studied in detail: *Belief revision* attempts to describe how an agent should accommodate a new belief (possibly inconsistent with his other beliefs) about a static world. *Belief update*, on the other hand, attempts to describe how an agent should change his beliefs as a result of learning about a change in the world.

Belief revision and belief update describe only two of the many scenarios in which beliefs change. Under this contract, we have constructed a framework to reason about belief change in general. The framework starts starts with a framework introduced earlier by the PI and his collaborators for reasoning about knowledge in multi-agent systems. The framework explicitly includes time, so lets us investigate how knowledge changes over time. We extend it to include beliefs and plausibilities, using the notion of plausibility measure discussed above. Since our framework allows us to talk about knowledge, belief, and time, it gives us a powerful tool for capturing various notions of belief change. Many situations previously studied in the literature, such as *diagnostic reasoning* and the *prisoner's dilemma* from game theory, can be easily captured in the framework.

The framework can easily capture revision and update as special cases, and gives us a much better understanding of them. For example, it brings out how update prefers to defer "abnormalities" to as late a time as possible, an approach which is rarely appropriate in AI applications. More recently, it was shown that making a qualitative *Markov assumption* gives us a notion of belief change that is perhaps closer to our intuitions than others proposed in the literature.

This work is summarized in papers B11, B12, B13, B14, C1, and C2.

Other highlights during the contract period include:

- Adam Grove received his Ph.D. in the summer of 1992, and began working at NEC in September, 1992.

- Daphne Koller received her Ph.D. in December, 1993. She is currently a postdoctoral fellow at U.C. Berkeley, holding both an NSF Postdoctoral Fellowship and a University of California President's Postdoctoral Fellowship. Her thesis won the Arthur Samuel award (for best thesis in the Stanford Comptuer Science Department).

- Joe Halpern was elected a fellow of the AAAI in 1993. In addition, he was invited to speak at the 3rd Kurt Gödel Colloquium in Brno, Czechoslovakia, in August, 1993, at the 13th Conference on Foundations of Software Technology and Theoretical Computer Science, in Bombay, India, in December, 1993, at the International Conference on Epistemic Logic and the Theory of Games and Decision, in Marseille, in January, 1994, at the Third CSLI Workshop on Logic, Language, and Computation, at Stanford, in June, 1994, and at the Fourth International Colloquium on Cognitive Science, to be held in San Sebastian, Spain, in May, 1995. He has also served on the program committees for the 12th and 14th Annual ACM Symposium on Principles of Distributed Computing, the 3rd Bar-Ilan Symposium on Foundations of AI, AAAI '94 (as area chair for knowledge representation), and the Third Israeli Symposium on Theory of Computing and Systems. Finally, he was appointed area editor of the *Journal of the ACM* in 1992 and consulting editor for the new online *Chicago Journal of Theoretical Science* in 1994, and continued to serve as editor of *Information and Computation* and the *Journal of Logic and Computation*, and as editor of the special issue of the *Journal of Computer and Systems Science* devoted to the best papers of the 1991 ACM Symposium on Theoretical Computer Science.

List of papers written under the contract:

A. Journal Publications

1. A. J. Grove and J. Y. Halpern, Naming and identity in propositional logics. I: The propositional case, *Journal of Logic and Computation*, 3:1, 1993, pp. 345-378.

2. F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller, A response to "Believing on the basis of evidence", *Computational Intelligence*, 10:1, 1994, pp. 21-25.

3. A. J. Grove, J. Y. Halpern, and D. Koller, Random worlds and maximum entropy, *Journal of AI Research*, 2, 1994, pp. 33-88.

4. A. J. Grove, J. Y. Halpern, and D. Koller, Asymptotic conditional probabilities: The unary case, accepted for publication, *SIAM Journal on Computing*, 1994.

5. J. Y. Halpern, Should knowledge entail belief, accepted for publication, *Journal of Philosophical Logic*, 1994.

6. A. J. Grove, J. Y. Halpern, and D. Koller, Asymptotic conditional probabilities: The nonunary case, accepted for publication, *Journal of Symbolic Logic*, 1994.

7. J. Y. Halpern and G. Lakemeyer, Levesque's axiomatization of only knowing is incomplete, accepted for publication, *Artificial Intelligence*, 1994.

B. Papers in Conference Procedings

1. A. Grove, J.Y. Halpern, and D. Koller, Asymptotic conditional probabilities for first-order logic, *Proc. of 24th ACM Symposium on Theory of Computing*, 1992, pp. 294–305.

2. A. Grove, J.Y. Halpern, and D. Koller, Random worlds and maximum entropy, *Proc. of 7th Annual IEEE Symposium on Logic in Computer Science*, 1992, pp. 22–33.

3. F. Bacchus, A. Grove, J.Y. Halpern, and D. Koller, From statistics to beliefs, *Proc. of AAAI '92*, 1992, pp. 602–608.

4. D. Koller and J.Y. Halpern, A logic for approximate reasoning, *Proc. of Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992, pp. 153–164.

5. A. Grove, Semantics for knowledge and communication, *Proc. of Third International Conference on Principles of Knowledge Representation and Reasoning*, 1992, pp. 213–224.

6. F. Bacchus, A. Grove, J.Y. Halpern, and D. Koller, Statistical foundations for default reasoning, *Proc. of 13th International Joint Conference on Artificial Intelligence (IJCAI '93)*, 1993, pp. 563–569.

7. J.Y. Halpern, Reasoning about only knowing with many agents, *Proc. of AAAI '93*, 1993, pp. 655–661.

8. J.Y. Halpern, A critical reexamination of default logic, autoepistemic logic, and only knowing, *Computational Logic and Proof Theory* (Proceedings of the Kurt Gödel Symposium, KGC '93), Lecture Notes in Computer Science, vol. 713, Springer, 1993, pp. 43–60.

9. F. Bacchus, A. Grove, J.Y. Halpern, and D. Koller, Generating degrees of belief from statistical information: an overview, *Proceedings of the 13th Conference on Foundations of Software Technology and Theoretical Computer Science*, Lecture Notes in Computer Science, Vol. 761, Springer, 1993, pp. 318–325.

10. J. Y. Halpern, Y. Moses, and M. Y. Vardi, Algorithmic knowledge, *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning About Knowledge*, 1994, pp. 255–266.

11. N. Friedman and J. Y. Halpern, A knowledge-based framework for belief change, Part I: Foundations, *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning About Knowledge*, 1994, pp. 44–64.

12. N. Friedman and J. Y. Halpern, A knowledge-based framework for belief change, Part II: Revision and update, *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, 1994, pp. 190–201.

13. N. Friedman and J. Y. Halpern, On the complexity of conditional logics, *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, 1994, pp. 202–213.

14. N. Friedman and J. Y. Halpern, Conditional logics of belief change, *AAAI-94 (Proceedings of the Twelfth National Conference on Artificial Intelligence)*, 1994, pp. 915–921.

15. R. Fagin, J. Y. Halpern, Y. Moses, and M. Vardi, An operational semantics for knowledge bases, *AAAI-94 (Proceedings of the Twelfth National Conference on Artificial Intelligence)*, pp. 1142–1147, 1994.

16. F. Bacchus, A. Grove, J.Y. Halpern, and D. Koller, Forming beliefs about a changing world, *AAAI-94 (Proceedings of the Twelfth National Conference on Artificial Intelligence)*, 1994, pp. 222–229.

17. F. Bacchus, A. Grove, J.Y. Halpern, and D. Koller, Generating new beliefs from old, *Proceedings of the Tenth Conference on Uncertainty in AI*, 1994, pp. 37–45.

C. Other Publications

1. F. Bacchus, J. Y. Halpern, and H. J. Levesque, Reasoning about noisy sensors in the situation calculus, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

2. N. Friedman and J. Y. Halpern, A qualitative Markov assumption and its implications for belief change, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

3. J. Y. Halpern and D. Koller, Irrelevance and conditioning in first-order probabilistic logic, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

4. N. Friedman, J. Y. Halpern, and D. Koller, First-order conditional logic revisited, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

5. J. Y. Halpern and D. Koller, Representation dependence in probabilistic inference, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

6. N. Friedman and J. Y. Halpern, Plausibility measures and default reasoning, submitted to *14th International Joint Conference on Artificial Intelligence (IJCAI '95)*.

Additional publications

Journal Publications

8. D. Koller and N. Megiddo, Constructing small sample spaces satisfying given constraints, *SIAM Journal on Discrete Mathematics*, 7:2, 1994.

9. D. Koller and N. Megiddo, Finding small mixed strategies, accepted for publication, *International Journal of Game Theory*, 1994.

B. Papers in Conference Proceedings

18. D. Koller and N. Megiddo, Constructing small sample spaces satisfying given constraints," *Proceedings of the 25th ACM Symposium on Theory of Computing*, 1993.

19. D. Karger and D. Koller, (De)randomized construction of small sample spaces in $\mathcal{NC}$, *Proceedings of the 35nd Annual Symposium on the Foundations of Computer Science (FOCS)*, 1994.

20. D. Koller, N. Megiddo, and B. von Stengel, Fast algorithms for finding randomized strategies in game trees, *Proceedings of the 26th ACM Symposium on Theory of Computing (STOC)*, 1994.

C. Other Publications

7 . Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller, From statistical knowledge bases to degrees of belief, IBM RJ 9855, 1994, submitted for publication.