

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

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 Service, 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 August 18, 1997		3. REPORT TYPE AND DATES COVERED Annual, 4/1/96-3/31/97	
4. TITLE AND SUBTITLE Knowledge Representation for an Uncertain World				5. FUNDING NUMBERS N00014-96-1-0718	
6. AUTHOR(S) Daphne Koller					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Stanford University Stanford, CA 94305				8. PERFORMING ORGANIZATION REPORT NUMBER n/a	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) ONR Code 311, Room 607 800 North Quincy St. Arlington, VA 22217-5660				10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release: distribution unlimited.				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Any application where an intelligent agent interacts with the real world must deal with the problem of uncertainty. Bayesian belief networks have become dominant in addressing this issue. This is a framework based on principled probabilistic semantics, which achieves effective knowledge representation and inference capabilities by utilizing the locality structure in the domain: typically, only very few aspects of the situation directly affect each other. Despite their success, belief networks are inadequate as a knowledge representation language for large, complex domains: Their attribute-based nature does not allow us to express general rules that hold in many different circumstances. This prevents knowledge from being shared among applications; the initial knowledge acquisition cost has to be paid for each new domain. It also inhibits the construction of large complex networks. We deal with this issue by presenting a rich knowledge-representation language from which belief networks can be constructed to suit specific circumstances, algorithms for learning the network parameters from data, fast approximate inference algorithms designed to deal with the large networks that result. We show how these techniques can be applied in domains involving continuous variables, in situations where the world changes over time, and in the context of planning under uncertainty.					
14. SUBJECT TERMS reasoning under uncertainty, Bayesian networks, object-oriented models, probabilistic inference, learning				15. NUMBER OF PAGES 13	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT unclassified	20. LIMITATION OF ABSTRACT unlimited		

Progress Report for ONR Contract N00014-96-1-0718

Daphne Koller

August 18, 1997

1 Administrative Information

- a. Title: Knowledge Representation for an Uncertain World
- b. Organization: Stanford University
- c. Subcontractors: None
- d. Principal Investigator:

Name: Daphne Koller
Phone: (650) 723-6598
Email: koller@cs.stanford.edu

e. Team Members / Graduate Students:

- Eric Bauer, Masters student
- Xavier Boyen, PhD student
- Raya Fratkina, Masters student
- Lise Getoor, PhD Student
- Manfred Jaeger, Postdoctoral Researcher.
- Alexander Kozlov, PhD Student
- Uri Lerner, PhD student
- Avi Pfeffer, PhD Student
- Mehran Sahami, PhD Student
- Leon Wong, PhD student.

19970825 080

f. Numerical Productivity Measures:

- Refereed papers submitted but not yet published: 2.
- Refereed papers published: 10.
- Unrefereed reports and articles: 2.
- Books or parts thereof submitted but not yet published: 1.
- Books or parts thereof published: 0.
- Patents filed but not yet granted: 0.
- Patents granted: 0.
- Invited Presentations: 2.
- Contributed Presentations: 6 (+6 by students and co-authors).

- Honors Received: 1.
- Prizes or Awards Received: 1.
- Promotions obtained: 0.
- Graduate Students Supported: 8.
- Post-docs supported: 1.
- Minorities supported: 0.

2 Project Summary

Objective The general goal of this program is to scale up our ability to do uncertain reasoning to domains that are significantly larger and more complex. As part of this effort, we are defining a natural and expressive knowledge representation language which combines the expressive power of Bayesian networks with that of first-order logic and object-oriented models. This language exploits the fact that most complex domains are highly structured, and can therefore be described compactly. Knowledge represented in this language is modular and reusable, so that it can be applied to many situations. The model can be adapted over time, in response to changes in the world or to conclusions reached in the reasoning process. In conjunction with our work on representation, we are also developing inference algorithms for large-scale complex model, and learning algorithms which will allow these models to be learned from data.

Background Over the last decade, there has been growing consensus that the ability to reason and act under uncertainty is fundamental to intelligent behavior. After all, the available sensors are rarely accurate, there are many relevant factors that cannot be observed, the future can never be predicted with certainty, and actions do not always have the desired outcomes.

In recent years, probabilistic representations have come to dominate the field of reasoning under uncertainty. These allow the agent's reasoning process to be based on sound theoretical foundations, and to utilize such well-understood techniques as *conditioning* for incorporating new information, and *expected utility maximization* for making decisions.

Bayesian belief networks make probabilistic reasoning tractable by utilizing the *locality structure* in the domain: the fact that, typically, only very few aspects of the situation directly affect each other. They have been used successfully in a variety of applications, ranging from medical diagnosis to real-time monitoring of space shuttle launches.

Despite their great success, Bayesian networks do not address all of the representational needs of a large and complex domain. In particular, a Bayesian network describes the world in terms of a predetermined set of attributes and a fixed set of connections among them. As a consequence, domains that involve many complex entities must be modeled using a single large network, one which is usually too complex to be useful. The static nature of the network is even more problematic in domains where the set of relevant entities is not known in advance. Finally, the fixed nature of the network structure also makes it difficult to share knowledge among applications, so that a high initial knowledge acquisition cost must typically be paid for each new domain.

Approach We are addressing this task using two main tools: the development of a flexible representation language which allows knowledge to be used in a wide variety of situations, and the development of algorithms that support effective inference by tailoring their behavior to the context and to the user's task.

Our representation language augments Bayesian networks with some of the expressive power of first-order logic (predicate calculus) and object-oriented models. It exploits the same basic idea, of viewing the world as a set of domain entities (or objects). We associate properties with these entities, and allow them to be related to each other. In a military context, for example, each military unit is an entity; various geographic locations where troops can be located are also entities, with their own properties. By defining *classes* of entities, we can exploit regularities in the domain, i.e., provide a single representation for complex behavior patterns that recur many times in the domain. For example, we can specify general models for military structure and tactics, for terrain, etc. We can also specify connections between different objects, e.g., the connection between a unit's activity and the properties of the terrain where it's currently located. These general models can then be used to generate specific models for a variety of situations.

In our recent work, we have defined an object-oriented Bayesian network language that supports modeling of this type. We have shown that the use of encapsulation allows the construction of complex entities out of simpler ones, with the ability to ignore the internal structure except where necessary. The use of inheritance hierarchies gives rise to a natural notion of abstraction, supporting reasoning at different levels of granularity. We can ascribe well-defined declarative semantics for this language using ideas from functional programming languages. We are currently integrating this framework with that of logical frame-based systems. The language will thus provide the first integration between two disparate strands of work: frame-based or object-oriented systems and Bayesian networks. As such, it allows a direct integration between sensor to symbol level systems and higher level schema-based representations.

In addition to the work on representation, we are also working on inference algorithms for more complex domains. In our recent work, we have shown that when we represent a complex domain using a more expressive language, we are not increasing the complexity of inference. Rather, the additional structure encoded in the language can be used to make inference significantly more efficient. For example, our language allows one object to be encapsulated within another. This property can be used to modularize the reasoning within objects, providing scaling guarantees for the performance of large domain models specified in this language. Our language also allows a model to explicitly encode the fact that the domain contains many similar entities (e.g., vehicles); an inference algorithm can then combine the inference process for all of them, avoiding unnecessary duplication of work. In some cases, we obtain an exponential decrease in complexity.

We are also continuing our work on inference targetted to a particular context and task. Our language allows us to focus our inference only to the entities that are relevant to the query. We are also developing techniques which allow intelligent approximate inference via a context-sensitive abstraction of parts of the domain.

Finally, we are doing substantial work on learning models from data, with the goal of easing or avoiding the process of knowledge elicitation from human experts. Our focus so far has been on learning the probabilistic parameters, since probabilities are notoriously hard to elicit from people. In the future, we are planning to begin investigating the problem of learning the domain structure: the probabilistic relationships between the properties of an entity, the taxonomic hierarchy over entities, etc. We believe that the additional structure encoded in our representation language will prove beneficial for the design of learning algorithms, as it did for the design of inference algorithms.

Progress There has been substantial progress on the project in the last year. We have completed the design of a stochastic functional language which forms the semantic underpinnings for the rest

of our work. We have defined an object-oriented Bayesian network language, which allows us to explicitly represent entities, inheritance (class) hierarchies, and component (part-of) hierarchies. We have shown how these additional language features can be used to support faster inference. We have used the idea of context-based inference for dealing with the task of reasoning with continuous-valued variables (such as location, velocity, or temperature). We provided an algorithm for discretizing such variables in a way that places the emphasis on those values that are likely to be relevant; e.g., when discretizing a location variable, the algorithm constructs a finer partition in those parts of the space where the object is likely to be. We have also made significant progress on the task of learning probabilistic parameters. We have shown how existing parameter learning algorithms can be made significantly faster at no extra cost. We have also extended these algorithms to the task of parameter learning for the more expressive representations that we are in the process of developing. We have shown that learning algorithms, like inference algorithms, can utilize the structure encoded in our more expressive representation language to achieve performance gains.

3 Work Planned for Next Year

- Exploit the paradigm of context-based inference for a much wider range of tasks, including context-specific abstraction in standard Bayesian networks, and automatic choice of the appropriate level of granularity in object-oriented Bayesian networks.
- Implement a prototype system for object-oriented Bayesian networks.
- Extend the object-oriented Bayesian network framework to deal with domains where there may be uncertainty over the domain structure.
- Extend the framework to deal with domains involving multiple objects whose state evolves over time.
- Develop inference algorithms for such temporal domains, exploiting the partition of the domain into different entities.
- Explore the problem of learning the structure of complex hierarchical models from data.
- Utilize hierarchical structure in representations for planning under uncertainty, and explore planning algorithms that utilize this structure.

4 Technical Transitions

- The object-oriented Bayesian representation language, whose development was initiated under this contract, will be used as the core representation of uncertainty in DARPA's DMIF program. The complexity of this application requires a modular structured representation language with the functionality described above, as well as the ability to generate specific domain models tailored to different regional conflicts and military exercises. The technology transfer is done in collaboration with IET (Information Extraction and Transport), Inc., one of the prime DMIF contractors. The technology transfer started a year ago, and will proceed throughout the DMIF contract; it is supported by a Stanford subcontract from IET's DMIF contract.
- Our object-oriented Bayesian representation language is being integrated into logical frame-based systems. The resulting system will utilize a recently stabilized Generic Frame Protocol

(GFP), which will allow it to interface with most of the existing large-scale knowledge-based systems (including Stanford University's Ontolingua system and ISI's Loom system). This work is being supported by a recently awarded DARPA contract which is part of DARPA's HPKB program.

- As part of the work on this contract, we have defined a probabilistic extension to *description logic* (a tractable sublanguage of first-order logic). A sublanguage of this language has been utilized by AT&T in their Information Manifold project. This project maintains a high-level description of the contents of various information sources on the web; the description is used to guide information gathering in response to user queries. The project currently utilizes standard description logic as its primary representation language, an approach which is limited since the contents of a data source can never be described with complete accuracy. The use of our language allows the Information Manifold to explicitly represent the extent to which a data source contains information relevant to the query, thereby intelligently prioritizing access to various data sources.
- Elements of our 1996 work on context-specific structure are in the process of being implemented into Microsoft's Bayesian Network reasoning tool.

5 Significant Accomplishments of Last Year

- The design and specification of an object-oriented Bayesian network language, and the specification of associated inference algorithms that utilize the additional domain structure in our more expressive representation language.
- The transfer of this technology to two major DARPA programs (DMIF and HPKB).
- Specification of an inference algorithm for arbitrary models with continuous attributes.
- Development of effective parameter learning algorithms for Bayesian networks and for models in our more expressive language.

6 List of Publications

The following is a list of publications completed, accepted for publication, or published over the last year. Those involving work (partially) supported by this contract are 1, 3, 5, 6, 7, 8, 9, 10, 13, and 14.

1. "Representations and Solutions for Game-Theoretic Problems," D. Koller and A.J. Pfeffer. *Artificial Intelligence*, **94**(1), July 1997, pages 167–215.
2. "From statistical knowledge bases to degrees of belief," F. Bacchus, A.J. Grove, J.Y. Halpern, and D. Koller. *Artificial Intelligence*, **87**, February 1997, pages 75–143.
3. "Adaptive probabilistic networks with hidden variables," J. Binder, D. Koller, S.J. Russell, and K. Kanazawa. *Machine Learning Journal*, to appear.
4. "(De)randomized construction of small sample spaces in NC ," D.R. Karger and D. Koller. *Journal of Computer and System Sciences*, special issue of selected papers from FOCS '94 (by invitation), accepted for publication.

5. "Update rules for parameter estimation in Bayesian networks," E. Bauer, D. Koller, and Y. Singer. *Proceedings of the 13th Annual Conference on Uncertainty in AI (UAI)*, Providence, Rhode Island, August 1997, pages 3–13.
6. "Object-Oriented Bayesian Networks," D. Koller and A. Pfeffer. *Proceedings of the 13th Annual Conference on Uncertainty in AI (UAI)*, Providence, Rhode Island, August 1997, pages 302–313. Winner of the UAI '97 best student paper award.
7. "Nonuniform dynamic discretization in hybrid networks," A.V. Kozlov and D. Koller. *Proceedings of the 13th Annual Conference on Uncertainty in AI (UAI)*, Providence, Rhode Island, August 1997, pages 314–325.
8. "Effective Bayesian Inference for Stochastic Programs," D. Koller, D. McAllester, and A. Pfeffer. *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI)*, Providence, Rhode Island, August 1997, pages 740–747.
9. "P-Classic: A tractable probabilistic description logic," D. Koller, A. Levy, and A. Pfeffer. *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI)*, Providence, Rhode Island, August 1997, pages 390–397.
10. "Learning probabilities for noisy first-order rules," D. Koller and A. Pfeffer. *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI)*, Nagoya, Japan, August 1997, to appear.
11. "Hierarchically classifying documents using very few words," D. Koller and M. Sahami. *Proceedings of the 14th International Conference on Machine Learning (ML)*, Nashville, Tennessee, July 1997, pages 170–178.
12. "Using probabilistic information in data integration," D. Florescu, D. Koller, and A. Levy. *Conference on Very Large Databases (VLDB)*, Athens, Greece, 1997, to appear.
13. "Learning the parameters of first-order probabilistic rules," D. Koller and A. Pfeffer. *Working Notes of the 1996 AAAI Fall Symposium on Learning Complex Behaviors in Adaptive Intelligent Systems*.
14. "Evidence-directed belief network simplification," D. Koller. *Working Notes of the 1996 AAAI Fall Symposium on Flexible Computation in Intelligent Systems*.


7 Online Information

All of my recent papers are available online via my webpage. The URL is:
<http://robotics.stanford.edu/~koller/>.

8 Quad Charts

In separate powerpoint file.

(a) System design



(b) Principal milestones

- December '97:
 - ◆ complete implementation of of prototype system for object-oriented and frame-based Bayesian networks.
 - ◆ demonstrate ability to generate models for variety of situations.
- June '98:
 - ◆ extend framework to temporal domains;
 - ◆ specify associated inference algorithms.
- September '98:
 - ◆ Develop algorithms for learning structure of complex hierarchical models from data.
- April '99:
 - ◆ Demonstrate ability to automatically generate models at appropriate level of abstraction.

(c) Rationale

- Uncertainty is unavoidable in the real world:
 - ◆ partial & noisy sensor data;
 - ◆ indeterminate action effects;
 - ◆ unpredictable exogenous events.
- Bayesian networks (BNs) are highly successful representation of uncertain knowledge.
- Problem: BNs do not support large-scale modeling:
 - ◆ no notion of entities & relations between them;
 - ◆ network structure static & inflexible;
 - ◆ no ability to represent & reuse common substructures;
 - ◆ no representation at different levels of granularity.
- **Our goal:** scale BNs to larger more complex domains.

Challenges & key ideas

- Key challenge: provide the additional expressive power without sacrificing effective reasoning.
 - ◆ models for complex domains involving multiple entities can easily grow intractably large.
- Key ideas:
 - ◆ exploit domain structure and its organization into fairly independent entities;
 - ◆ exploit repetition of common patterns for multiple similar entities;
 - ◆ allow generation of models at different levels of abstraction;
 - ◆ tailor model and inference algorithm to current context & task.

(d) Technical approach

- Develop object-oriented probabilistic language that allows specification of entities, their properties, and relations between them.
- Utilize structure of object-oriented framework---encapsulation and inheritance---to support description of domains at different levels of abstraction.
- Construct algorithms that use this structured language for building domain models tailored to current situation.
- Extend the algorithms to adapt the domain model in response to user needs or to changes in current situation.

(e) Goals

Benefits

- Develop structured representation language for uncertain domains with multiple entities.
- Develop algorithms that generate compact domain models tailored to given situation.
- Develop capability for generating domain models at different levels of abstraction.
- Utilize framework as basis for developing planning algorithms.

- Reduction & amortization of knowledge acquisition effort for complex domains.
- Utilization of general domain model for many qualitatively different situations.
- Support for reasoning under uncertainty in time-critical settings.
- Provide capability for planning in realistic (uncertain) domains.

(f) Exciting significant results

- Development of representation language that integrates two main knowledge-representation paradigms:
 - ◆ frame-based/object-oriented systems;
 - ◆ probabilistic reasoning & Bayesian networks.
- Use of language to integrate sensor-to-symbol level systems and schema-based representations.
- Selection of language as primary representation of uncertainty within DARPA's DMIF & HPKB projects.