

Probabilistic Sensitivity Analysis for Situation Awareness

PI: Adnan Darwiche
Computer Science Department
University of California
Los Angeles, CA 90095
darwiche@cs.ucla.edu

Final Performance Report for FA9550-05-1-0075

Program Manager: David Luginbuhl

June 30, 2008

Abstract

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The 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 the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE			3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)	

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

1 Project Objectives

Probabilistic graphical models are normative tools for managing dynamic, uncertain situations. These models have been finding increased applications as they provide a structured methodology for building complex situation models, and a suite of efficient algorithms for reasoning about modelled situations. The structure of a probabilistic graphical model represents qualitative dependencies among situation attributes, and its parameters (probabilities) represent a quantification of these dependencies. Model parameters represent both the strongest and weakest aspect of probabilistic graphical models. Specifically, it is these parameters that allow one to have fine-grained models of situations (compared to the coarse-grained models afforded by symbolic representations). Yet, it is these parameters that decision makers find the least intuitive, as they worry about their impact on the decision making process (i.e., what would happen if we change a .80 probability to a .85?). *Sensitivity analysis is a remedy to this concern as it focuses on understanding the relationship between the local parameters that quantify a probabilistic model and the global conclusions that are drawn based on the model.* This understanding has far more reaching applications than is suggested above though, as it can be the basis for addressing complex tasks, such as situation control and information-system design.

The objective of this effort is to conduct an investigation into probabilistic sensitivity analysis for aiding the process of situation awareness. As far as technical advancements are concerned, the expected outcomes of this investigation will enable us to (1) assess the impact of parameter changes on the quality of decisions based on probabilistic models, (2) characterize minimal parameter changes that are needed to ensure the conformance of a model with expectations, and (3) quantify and effectively measure the change undergone by a probabilistic model when perturbed through parameter change. These results are expected to significantly improve the effectiveness of probabilistic graphical models as a tool for achieving situation awareness. First, they would allow decision makers to get first hand insights into the sensitivity/robustness of their decisions to the various aspects/assumptions underlying the situation model. Next, they will give them a handle on how to manipulate the situation in order to achieve certain objectives. Finally, they will expand the realm of queries that decision makers can ask about a situation when conducting if-then analysis (What would flip the most likely scenario? What would increase our likelihood of success? Is it worth replacing these sensors by more reliable ones?).

The proposed work involves both a theoretical study, and a correspond-

ing implementation of developed techniques into the SAMIAM system (*reasoning.cs.ucla.edu/samiam*).

2 Project Accomplishments

Most of our initial focus in sensitivity analysis has been on directed graphical models in the form of Bayesian networks. Undirected graphical models (Markov networks), however, are quite important in certain application areas, such as sensor networks and recognition applications in general, where they may provide a more natural modeling tool. One of the main objectives of this effort has been to extend the sensitivity analysis engine developed for Bayesian networks so it can handle undirected models, at least for the simple case of single constraints and single parameter changes.

One of the main accomplishments of this project has been a comprehensive investigation of sensitivity analysis in Markov networks, which resolved most of the pending questions with regard to Markov networks as formulated in the project proposal. Based on these results, which are reported in [1], we now have the basis for extending the sensitivity analysis engine of SAMIAM to Markov networks, making this important type of networks available for the same type of sensitivity analysis that Bayesian networks support. The ability to develop this extension allows us to more naturally handle applications, such as sensor networks, which have been modelled using undirected models, in addition to combinations of directed and undirected models. The majority of these applications center around sensor fusion, with a number of sensors (that measure the same physical quantity such as temperature or proximity) spread according to some physical layout. Model parameters in this case correspond to sensor biases, and other attributes. The sensitivity analysis engine can then be used to reason about the impact of these sensor properties on the decision making process. We should note here that the reading of a sensor can itself be modelled as a model parameter, allowing us to reason about the impact of sensor readings on the decision making process too. For example, we may ask: Would turning this sensor on allow us to prove a certain hypothesis beyond a given threshold?

Sensitivity analysis is usually focused on a particular query type. The most common query, and the one we have focused our initial efforts on, is that of posterior marginals: the probability of some event given some evidence. Here, sensitivity analysis would be concerned with two basic questions: (1) How would the answer to a query get effected by some parameter change? (2) What parameter changes will induce a particular change in a

query? One of the main objectives of this effort has been to extend this type of analysis to one of the most practical types of queries: most likely scenarios. That is, given a most likely scenario in light of some evidence, the goal is to answer two questions: (1) How much change to some parameter would preserve the most likely scenario ? (2) What parameter changes will flip the most likely scenario to some other scenario?

One of the main accomplishments of this project has been the results we obtained on sensitivity analysis with respect to most-likely-scenario queries, where we showed that we can indeed perform this task efficiently. These results, which are reported in [2], have been implemented in our SAMIAM system. We now, therefore, have the nuclear algorithm for developing a very practical and effective engine for aiding decision makers in examining the robustness of most like scenarios with respect to parameter changes. For example, if a critical scenario turns out too sensitive to some parameter, a decision maker may then expend more resources to accurately estimate that parameter, as a measure to enhance their confidence in any decision they make based on the scenario.

Sensitivity analysis in graphical models rely on the ability to do probabilistic inference efficiently. It is well known that the efficiency of this inference is tied to the topology (connectedness) of graphical models, where heavily connected models are known to be hard computationally. We have had a long tradition of work that deals with this problem, which has focused on exploiting the structure exhibited by model parameters (in addition to that exhibited by network topology). Our main accomplishments on this subject have been reported in [4], which extend our previous results in [3]. These results show that by exploiting parametric structure, we can solve (and exactly) models whose topology make them completely outside the scope of classical inference methods. Hence, we are now able to perform sensitivity analysis on these models, were we could not before. The corresponding results have already been implemented and released in our ACE system available at <http://reasoning.cs.ucla.edu/ace/>.

We have also pursued a new direction in approximate inference, during the last phases of this project, which simplifies the network topology by deleting edges from a given model. This is known to produce approximate models, but the goal is then to make the new approximate model as close as possible to the original model in order to minimize the discrepancy between the exact and approximate answers. Our accomplishments on this front have been quite extensive and are reported in [5, 7, 8, 6].

We finally note that we have applied our sensitivity engine to a real-world problem that involves hundreds of sensors from the domain of electrical

power systems used by NASA on a number of space vehicles. The system we worked with is part of the Advanced Diagnostic and Prognostic Testbed (ADAPT) at NASA Ames <http://ti.arc.nasa.gov/projects/adapt/>.

3 Personnel Supported

Hei Chan (Graduate student), Mark Chavira (Graduate Student), Arthur Choi (Graduate student), Keith Cascio (Staff Programmer) and Adnan Darwiche (PI).

4 Honors and Awards

The PI was elected in 2007 as a AAAI Fellow. This is the fellow program of the Association for the Advancement of Artificial Intelligence (AAAI) which was started in 1990 to recognize individuals who have made significant, sustained contributions—usually over at least a ten-year period—to the field of artificial intelligence. The PI citation reads:

For significant contributions to the development and application of both probabilistic and logical methods in automated reasoning.

The PI has also been selected to serve a four year term in the Journal of Artificial Intelligence Research (JAIR), first as an Associate Editor-in-chief for 2007-2008, and then as Editor-in-chief for 2009-2010.

The PI group won the Gold Medal (first place) at the 2007 international competition for satisfiability testing. This is a world-wide, extremely competitive event that takes place every two-years. Our winning SAT solver, RSAT, has been downloaded more than a 1000 times since its release, with most downloads taking place since it was announced as the winner of the SAT'07.

The PI group was the only team to solve all problem instances at the International Evaluation of Exact Probabilistic Reasoning Systems, which took place at the uncertainty in AI conference (UAI) in 2006.

5 Publications

Following are the relevant publications, referenced in the above description. These publications are available at: reasoning.cs.ucla.edu

References

- [1] Hei Chan and Adnan Darwiche. Sensitivity analysis in Markov networks. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1300–1305. Professional Book Center, 2005.
- [2] Hei Chan and Adnan Darwiche. On the robustness of most probable explanations. In *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 63–71, Arlington, Virginia, 2006. AUAI Press.
- [3] Mark Chavira and Adnan Darwiche. Compiling Bayesian networks with local structure. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1306–1312, 2005.
- [4] Mark Chavira and Adnan Darwiche. On probabilistic inference by weighted model counting. *Artificial Intelligence Journal*, 172(6–7):772–799, 2008.
- [5] Arthur Choi, Hei Chan, and Adnan Darwiche. On Bayesian network approximation by edge deletion. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 128–135, Arlington, Virginia, 2005. AUAI Press.
- [6] Arthur Choi, Mark Chavira, and Adnan Darwiche. Node splitting: A scheme for generating upper bounds in bayesian networks. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 57–66, 2007.
- [7] Arthur Choi and Adnan Darwiche. An edge deletion semantics for belief propagation and its practical impact on approximation quality. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2006.
- [8] Arthur Choi and Adnan Darwiche. A variational approach for approximating Bayesian networks by edge deletion. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2006.