Linköping Electronic Articles in
Computer and Information Science
Vol. 5(1999): nr 28

On the Role of Partial Differentiation in Probabilistic Inference

Adnan Darwiche

Linköping University Electronic Press
Linköping, Sweden

http://www.ep.liu.se/ea/cis/1999/028/

Posted on April 6,
formally published on December 21, 1999
by Linköping University Electronic Press
581 83 Linköping, Sweden

Linköping Electronic Articles in
Computer and Information Science
ISSN 1401-9841
Series editor: Erik Sandewall

©*1999 Adnan Darwiche*
Typeset by the author

Recommended citation:
<Author>. <Title>. Linköping Electronic Articles in Computer and Information Science, Vol. 5(1999): nr 28. http://www.ep.liu.se/ea/cis/1999/028/. *December 21, 1999.*

This URL will if possible also contain a link to the author's home page.

The publishers will keep this article on-line on the Internet (or its possible replacement network in the future) for a period of 25 years from the date of publication, unless the author requests its withdrawal and the article becomes available from another source.

The on-line availability of the article implies a permission for anyone to read the article on-line, to print out single copies of it, and to use it unchanged for any non-commercial research and educational purpose, including making copies for classroom use, for as long as the article's webpage at the URL specified below does not indicate that this permission has been revoked.

The publication of the article on the date stated above included also the production of a limited number of copies on paper, which were archived in Swedish university libraries like all other written works published in Sweden. The publisher has taken technical and administrative measures to assure that the on-line version of the article is permanently equal to the archived printed copies.

For additional information about the Linköping University Electronic Press and its procedures for publication and for assurance of document integrity, please refer to its WWW home page: http://www.ep.liu.se/ *or by conventional mail to the address stated above.*

Abstract

We present in this paper one of the simplest, yet most comprehensive frameworks for inference in Bayesian networks. According to this framework, one compiles a Bayesian network into a polynomial – in which variables correspond to potential evidence and network parameters – and then computes the partial derivatives of this polynomial with respect to each variable. Once such derivatives are made available, one can compute in constant-time answers to a large class of probabilistic queries, which are central to classical inference, parameter estimation, model validation and sensitivity analysis. We show a number of key results relating to this framework.

First, given a Bayesian network of size n and an elimination order of width w, we present an elimination algorithm for compiling the polynomial in $O(n \star exp(w))$ time and space.

Next, given some evidence and parameter setting, we show that the compiled polynomial can be evaluated, and all its first partial derivatives computed simultaneously, in $O(n \star exp(w))$ time and space. Once these derivatives are made available, we show that one can compute in constant-time: the posterior marginal of any network variable or family, the probability of evidence after having retracted the value of any evidence variable, and the sensitivity of evidence probability to change in any network parameter.

Finally, we show that second partial derivatives can all be computed simultaneously in $O(n^2 \star exp(w))$ time and space. Moreover, once such derivatives are made available, one can compute in constant-time a variety of sophisticated queries, such as context-specific independence between any pair of variables, and the amount of change to some network parameter which is needed to ensure a particular ranking on some hypotheses.

The proposed framework provides new insights into the role of partial differentiation in probabilistic inference. Moreover, its combined simplicity, comprehensiveness and computational complexity appear to be unique among existing frameworks for inference in Bayesian networks.

For technical reasons, this article is stored as two separate computer files, labelled "Part I" and "Part II", respectively. Please replace the present page by a printout of the "Part II" file.