## Aalborg Universitet

## AALBORG UNIVERSITY

## Learning Bayesian Networks with Mixed Variables

Bøttcher, Susanne Gammelgaard

Publication date:
2005

Document Version
Publisher's PDF, also known as Version of record

Link to publication from Aalborg University

Citation for published version (APA):
Bøttcher, S. G. (2005). Learning Bayesian Networks with Mixed Variables.

## General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy
If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

# Learning Bayesian Networks with Mixed Variables 

# Susanne Gammelgaard Bøttcher 

Ph.D. Thesis

2004

## AALBORG UNIVERSITY

Department of Mathematical Sciences

## Preface

This thesis is the result of my Ph.D. study at the Department of Mathematical Sciences, Aalborg University, Denmark. The work has mainly been founded by Aalborg University, but also in parts by the ESPRIT project P29105 (BaKE) and by Novo Nordisk A/S.

The thesis concerns learning Bayesian networks with both discrete and continuous variables and is based on the following four papers:
I. Learning Conditional Gaussian Networks.
II. deal: A Package for Learning Bayesian Networks.
III. Prediction of the Insulin Sensitivity Index using Bayesian Networks.
IV. Learning Dynamic Bayesian Networks with Mixed Variables.

Many of the results in Paper I are published in Bøttcher (2001). Paper II is published in Bøttcher and Dethlefsen (2003a). The developed software package, deal, is written in $R$ ( $R$ Development Core Team 2003) and can be downloaded from the Comprehensive R Archive Network (CRAN) http: //cran. R-project.org/. Paper II and Paper III are written together with Claus Dethlefsen, Aalborg University.

The individual papers are self-contained with an individual bibliography and figure, table and equation numbering. Parts and bits therefore appear in more than one paper. A basic understanding of the results in Paper 1 is though an advantage in reading the other papers. Those who are not familiar with Bayesian networks in general, might consult introductory books such as Jensen (1996) and Cowell, Dawid, Lauritzen and Spiegelhalter (1999).

I am much indebted to my supervisor Steffen L. Lauritzen for inspiring discussions and for many valuable comments. Also thanks to Søren LundbyeChristensen who has always taken the time to help me, when I had some questions.

I owe a special thanks to Claus Dethlefsen for an inspiring collaboration and for sharing his expertise within statistical programming with me.

Part of the work is based on data provided by Torben Hansen, Novo Nordisk $\mathrm{A} / \mathrm{S}$. Thanks for letting us use these data.

From September 1999 to November 1999 I visited the Fields Institute in Toronto, Canada, as a participant in the research program "Causal Interpretation and Identification of Conditional Independence Structures". I wish to thank all the people there for a memorable stay.

Acknowledgements also go to my colleagues at the Department of Mathematical Sciences for providing a nice working environment and especially to E. Susanne Christensen, who has been my mentor and one of my best friends during the years I have worked at the department.

Special thanks go to my colleagues and very dear friends Malene, Kim and Claus. We always have a lot of fun together and we are also their for each other during the hard times.

Both my parents and my in-laws have been very supportive during this study and that has meant a lot to me, so a word of thanks also goes to them.

All my thanks go to my two sons, Johannes and Magnus, for reminding me of the things that are the most important in life.

Finally, I would like to thank my husband Peter for his invaluable support and endless encouragement during my Ph.D. study. Also thanks for caring for Johannes and Magnus during my sometimes long work hours. It was the thought of having more time to spent with the three of you that finally made me finish this thesis.

## Summary

The main topic of this thesis is learning Bayesian networks with discrete and continuous variables.

A Bayesian network is a directed acyclic graph that encodes the joint probability distribution for a set of random variables. The nodes in the graph represent the random variables and missing arrows between the nodes, specify properties of conditional independence between the variables. It consists of two parts, a knowledge base and an inference engine for handling this knowledge. This thesis relies on already developed methods for inference and concentrate on constructing the knowledge base.

When constructing the knowledge base, there are two things to consider, namely learning the graphical structure and learning the parameters in the probability distributions. In this thesis, the focus is on learning Bayesian networks, where the joint probability distribution is conditional Gaussian. To learn the parameters, conjugate Bayesian analysis is used and parameter independence and complete data are assumed. To learn the graphical structure, network scores for the different structures under evaluation, are calculated and used to discriminate between the structures. To calculate these scores, the prior distribution for the parameters for each network under evaluation, must be specified. An automated procedure for doing this is developed. With this procedure, the parameter priors for all possible networks are deduced from marginal priors calculated from an imaginary database.

Bayes factors to be used when searching for structures with high network score, are also studied. To reduce the search complexity, classes of models are identified for which the Bayes factor for testing an arrow between the same two variables, is the same.

To be able to use the methods in practice, a software package called deal, written in $R$, is developed. The package includes procedures for defining priors, estimating parameters, calculating network scores, performing heuristic search as well as simulating data sets with a given dependency structure.

To illustrate the Bayesian learning procedure, a dataset from a study concerning the insulin sensitivity index, is analyzed. The insulin sensitivity index is an index that can be used in assessing the risk of developing type 2 diabetes. Interest is in developing a method to determine the index from measurements of glucose and insulin concentrations in plasma sampled subsequently after an glucose intake. As the dependency relations between the glucose and insulin measurements are complicated, it is proposed to use Bayesian networks. The conclusion is that the insulin sensitivity index for a non-diabetic glucose tolerant subject can be predicted from the glucose and insulin measurements, the gender and the body mass index, using Bayesian networks.

Finally, dynamic Bayesian networks with mixed variables are studied. A dynamic Bayesian network is just a simple extension of an ordinary Bayesian network and is applied in the modeling of time series. It is shown how the methods developed for learning Bayesian networks with mixed variables, can be extended to use for learning dynamic Bayesian networks with mixed variables. As the Markov order of a times series is not always known, it is also shown how to learn this order.

## Summary in Danish - sammendrag

Denne afhandling omhandler konstruktion (indlæring) af bayesianske netværk med diskrete og kontinuerte variable.

Et bayesiansk netværk er en orienteret graf uden kredse, der beskriver den simultane sandsynlighedsfordeling for en mængde af stokastiske variable. Knuderne i grafen repræsenter de stokastiske variable og manglende pile imellem knuderne repræsenterer betingede uafhængighedsantagelser. Et bayesiansk netværk består af to dele, en vidensbase og en inferensmaskine til at håndtere denne viden. Denne afhandling bruger allerede udviklede metoder til inferens og fokuserer på at konstruere vidensbasen.

Konstruktionen af vidensbasen kan deles op i to dele, nemlig selektion af den grafiske struktur og estimation af parametrene i sandsynlighedsfordelingerne. I denne afhandling fokuseres der på bayesianske netværk, hvor den simultane sandsynlighedsfordeling er betinget gaussisk. Til parameter estimation bruges konjugeret bayesiansk analyse og det antages, at parametrene er uafhængige og at data er fuldstændige.

Til selektion af den grafiske struktur beregnes et mål for hvor godt en given struktur beskriver data, i afhandlingen kaldet for en netværksscore. Netværksscoren beregnes for alle de strukturer, der tages i betragtning og bruges således til at diskriminere imellem de forskellige strukturer.

For at kunne beregne disse netværksscorer skal man kende apriori fordelingen for parametrene i alle de betragtede netværk. En automatisk procedure til at deducere disse apriori fordelinger fra marginale apriori fordelinger, beregnet fra en imaginær database, udvikles. Desuden studeres bayes faktorer, da disse bruges i forskellige søge strategier til søgning efter netværk med høj netværksscore. For at reducere søge kompleksiteten identificeres klasser af modeller, hvor bayes
faktoren til at teste en pil mellem de samme to variable, er den samme.
For at kunne bruge de udviklede metoder i praksis, er et software program, kaldet deal, udviklet. Pakken, som er skrevet til R, inkluderer procedurer til at definere apriori fordelinger, estimere parametre, beregne netværksscorer, søge efter netværk med høj netværksscore og simulerer datasæt med en given afhængighedsstruktur.

Til illustration af den bayesianske indlæringsprocedure analyseres et datasæt fra et studie, der omhandler insulin sensitivitets indekset. Insulin sensitivitets indekset er et indeks, der kan bruges til at vurdere risikoen for at udvikle type 2 diabetes. Formålet med studiet er at udvikle en metode, der kan bestemme dette indeks ud fra gentagne målinger af glukose og insulin koncentrationerne i plasma efter et glukose indtag. Da afhængighedsstrukturen mellem glukose og insulin målingerne er kompleks, bruges bayesianske netværk til at repræsentere disse afhængigheder. Konklusionen er at insulin sensitivitets indekset for ikkediabetiske glukose tolerante individer, kan predikteres fra glukose og insulin målingerne, kønnet og body mass indekset, ved at bruge bayesianske netværk.

Til sidst i afhandlingen studeres dynamiske bayesianske netværk med blandede variable. Et dynamisk bayesiansk netværk er en simpel udvidelse af de sædvanlige bayesianske netværk og anvendes til modellering af tidsrække data. Det vises hvordan de metoder, der er udviklet til indlæring af de sædvanlige bayesianske netværk med blandede variable, kan udvides, så de kan anvendes til indlæring af dynamiske bayesianske netværk med blandede variable. Da Markov ordenen af en tidsrække ikke altid er kendt, vises det også, hvordan man kan indlære denne orden.

## Contents

Preface ..... iii
Summary ..... v
Summary in Danish - sammendrag ..... vii
Introduction ..... 1
Paper I. Learning Conditional Gaussian Networks ..... 11
1 Introduction ..... 13
2 Bayesian Networks ..... 14
3 Bayesian Networks for Mixed Variables ..... 15
4 Learning the Parameters in a CG Network ..... 17
5 Learning the Structure of a CG Network ..... 23
6 The Master Prior Procedure ..... 26
7 Local Masters for Mixed Networks ..... 31
8 Model Search ..... 34
9 Example ..... 41
Paper II. deal: A Package for Learning Bayesian Networks ..... 57
1 Introduction ..... 59
2 Bayesian Networks ..... 61
3 Data Structure ..... 62
4 Specification of a Bayesian Network ..... 63
5 Parameter Learning ..... 67
6 Learning the Structure ..... 73
7 Hugin Interface ..... 78
8 Example ..... 79
9 Discussion and Future Work ..... 82
10 Manual Pages for deal ..... 84
Paper III. Prediction of the Insulin Sensitivity Index using Bayesian Networks ..... 107
1 Introduction ..... 109
2 Data ..... 110
3 Bayesian Networks ..... 111
4 Inference ..... 113
5 Results ..... 115
6 Discussion ..... 124
Paper IV. Learning Dynamic Bayesian Networks with Mixed Variables 127
1 Introduction ..... 129
2 Dynamic Bayesian Networks ..... 130
3 Dynamic Bayesian Networks for Mixed Variables ..... 134
4 Examples of DBNs ..... 137
5 Learning DBNs with Mixed Variables ..... 140
6 Specifying Prior Distributions ..... 149
7 Example ..... 152

## Introduction

The main focus of this Ph.D. thesis is to develop statistical methods for learning Bayesian networks with mixed variables. To be able to use these methods in practice, the software package deal is developed. Besides, the methods are extended to use for dynamic Bayesian networks.

## Background

Bayesian networks was developed in the late 80 's by Pearl (1988) and Lauritzen and Spiegelhalter (1988). For terminology and theoretical aspects, see Lauritzen (1996), Jensen (1996) and Cowell et al. (1999) among others.

A Bayesian network is a directed acyclic graph that encodes the joint probability distribution for a set of random variables. The nodes in the graph represent the random variables and missing arrows between the nodes, specify properties of conditional independence between the variables.

A Bayesian network consists of two parts, a knowledge base and an inference engine for handling this knowledge. Generally, inference is computationally heavy as it involves calculating huge joint distributions, especially if there are many variables in the network. Therefore efficient methods of implementing Bayes' theorem are being used. These implementations uses the fact that the the joint probability distribution of all the variables in a network, factorizes according to the structure of the graph. The distributions of interest can then be found by a series of local computations, involving only some of the variables at a time, see $e . g$. Cowell et al. (1999) for a thorough treatment of these methods. The methods are implemented in e.g. Hugin (http://www.hugin.com). Bayesian networks are therefore suitable for problems where the variables exhibit a complicated dependency structure. See Lauritzen (2003) for a recent overview over different applications.

In this thesis, we will rely on already developed methods for inference and concentrate on constructing the knowledge base. The work is documented through four papers, which will be described in the following.

## Paper I. Learning Conditional Gaussian Networks

When constructing the knowledge base there are two things to consider, namely specifying the graphical structure and specifying the probability distributions. Paper I addresses these issues for Bayesian networks with mixed variables.

In this paper, the focus is on learning Bayesian networks, where the joint probability distribution is conditional Gaussian. For an introductory text on learning Bayesian networks, see Heckerman (1999). To learn the parameters in the local probability distributions, conjugate Bayesian analysis is used. As conjugate local priors, the Dirichlet distribution is applied for discrete variables and the Gaussian-inverse gamma distribution is applied for continuous variables, given a configuration of the discrete parents. We assume parameter independence and complete data. To learn the graphical structure, network scores for the different structures under evaluation, are calculated and these scores are used to discriminate between the structures. To calculate these scores, the prior distribution for the parameters, for each network under evaluation, must be specified. In Heckerman, Geiger and Chickering (1995) and Geiger and Heckerman (1994) an automated procedure for doing this in respectively the purely discrete case and the purely continuous case, is developed. Their work is based on principles of likelihood equivalence, parameter modularity, and parameter independence. It leads to a method where the parameter priors for all possible networks, are deduced from one joint prior distribution, in this thesis called a master prior distribution.

In Paper I, we build on their results and develop a method, which can be used on networks with mixed variables. If used on networks with only discrete variables or only continuous variables, it coincides with the methods developed in in respectively Heckerman et al. (1995) and Geiger and Heckerman (1994).

If the number of random variables in a network is large, it is computationally infeasible to calculate the network score for all the possible structures. Therefore different methods for searching for structures with high network score, are being used, see e.g. Cooper and Herskovits (1992). Many of these methods use Bayes factors as a way of comparing the network scores for two different models. We therefore study Bayes factors for mixed networks. To reduce the search complexity, classes of models are identified for which the Bayes factor
for testing an arrow between the same two variables, is the same.
Finally, an analysis of a simple example illustrates the developed methods and is also used for showing how the strength of the prior parameter distribution affects the result of the analysis.

## Paper II. deal: A Package for Learning Bayesian Networks

To be able to use the methods presented in Paper I in practice, we have developed a software package called deal, written in $R$ ( R Development Core Team 2003).

In particular, the package includes procedures for defining priors, estimating parameters, calculating network scores, performing heuristic search as well as simulating data sets with a given dependency structure. The package can be downloaded from the Comprehensive R Archive Network (CRAN) http: //cran.R-project.org/ and may be used under the terms of the GNU General Public License Version 2.

The package supports transfer of the learned network to Hugin (http: / /www . hugin. com). The Hugin graphical user interface (GUI) can then be used for further inference in this network. Besides, deal adds functionality to $R$, so that Bayesian networks can be used in conjunction with other statistical methods available in $R$ for analyzing data. In particular, deal is part of the $g R$ project, which is a newly initiated workgroup with the aim of developing procedures in $R$ for supporting data analysis with graphical models, see http://www. r-project.org/gR.

## Paper III. Prediction of the Insulin Sensitivity Index using Bayesian Networks

To illustrate the Bayesian learning procedure, we have in Paper III analyzed a dataset collected by Torben Hansen, Novo Nordisk A/S.

The insulin sensitivity index, $S_{I}$, is an index that can be used in assessing the risk of developing type 2 diabetes. The index is determined from an intravenous glucose tolerance test (IVGTT), where glucose and insulin concentrations in plasma are subsequently sampled after an intravenous glucose injection. However, an IVGTT is time consuming and expensive and therefore not suitable for large scale epidemiological studies. Therefore interest is in developing a method to assess $S_{I}$ from measurements from an oral glucose tolerance test (OGTT). In an OGTT, glucose and insulin concentrations in plasma are, after an glucose
intake, sampled at a few time points.
In the present study, 187 non-diabetic glucose tolerant subjects underwent both an OGTT and an IVGTT. From the IVGTT, the $S_{I}$ values are determined using Bergmans minimal model (Bergman, Ider, Bowden and Cobelli 1979) as done in Pacini and Bergman (1986). The aim of our analysis is to determine the $S_{I}$ values from the measurements from the OGTT and investigate whether the $S_{I}$ values from the oral study, are correlated to the $S_{I}$ values determined from the intravenous study.

As the dependency relations between the glucose and insulin measurements are complicated, we propose to use Bayesian networks. We learn various Bayesian networks, relating measurements from the OGTT to the $S_{I}$ values determined from the IVGTT. We conclude that the $S_{I}$ values from the oral study, determined using Bayesian networks, are highly correlated to the $S_{I}$ values from the intravenous study, determined using Bergmans minimal model.

## Paper IV. Learning Dynamic Bayesian networks with Mixed Variables

A dynamic Bayesian network is an extension of an ordinary Bayesian network and is applied in the modeling of time series, see Dean and Kanazawa (1989). In Murphy (2002) a thorough treatment of these models for first order Markov time series, is presented and in Friedman, Murphy and Russell (1998), learning these networks in the case with only discrete variables, is described. In Paper IV, methods for learning dynamic Bayesian networks with mixed variables, are developed. These methods are just simple extensions of the methods described in Paper I for learning Bayesian networks with mixed variables. It is therefore also straight forward to use deal to learn dynamic Bayesian networks.

Contrary to previous work, we consider time series with Markov order higher than one and show how the Markov order can be learned.

To illustrate the developed methods, the Wölfer's sunspot numbers are analyzed.

## References

Anderson, T. W. (1971). The Statistical Analysis of Time Series, John Wiley and Sons, New York.

Badsberg, J. H. (1995). An Environment for Graphical Models, PhD thesis, Aalborg University.

Bergman, R. N., Ider, Y. Z., Bowden, C. R. and Cobelli, C. (1979). Quantitative estimation of insulin sensitivity, American Journal of Physiology 236: E667-E677.

Bernardo, J. M. and Smith, A. F. M. (1994). Bayesian Theory, John Wiley \& Sons, Chichester.

Bøttcher, S. G. (2001). Learning Bayesian Networks with Mixed Variables, Artificial Intelligence and Statistics 2001, Morgan Kaufmann, San Francisco, CA, USA, pp. 149-156.

Bøttcher, S. G. (2003). Learning Bayesian networks with mixed variables, Aalborg University, http://www.math. auc.dk/~alma.

Bøttcher, S. G. and Dethlefsen, C. (2003a). deal: A Package for Learning Bayesian Networks, Journal of Statistical Software 8(20): 1-40.

Bøttcher, S. G. and Dethlefsen, C. (2003b). Learning Bayesian networks with R, in K. Hornik, F. Leisch and A. Zeileis (eds), Proceedings of the 3rd international workshop on distributed statistical computing. ISSN 1609395X.

Bøttcher, S. G., Milsgaard, M. B. and Mortensen, R. S. (1995). Monitoring by using dynamic linear models - illustrated by tumour markers for cancer, Master's thesis, Aalborg University.

Chickering, D. M. (1995). A transformational characterization of equivalent Bayesian-network structures, Proceedings of Eleventh Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 87-98.

Chickering, D. M. (1996). Learning Bayesian networks is NP-Complete, in D. Fisher and H. J. Lenz (eds), Learning from Data: Artificial Intelligence and Statistics V, Springer-Verlag, New York, pp. 121-130.

Chickering, D. M. (2002). Optimal structure identification with greedy search, Journal of Machine Learning Research 3: 507-554.

Cooper, G. and Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data, Machine Learning 9: 309-347.

Cowell, R. G., Dawid, A. P., Lauritzen, S. L. and Spiegelhalter, D. J. (1999). Probabilistic Networks and Expert Systems, Springer-Verlag, Berlin-Heidelberg-New York.

Dawid, A. P. (1982). The Well-Calibrated Bayesian, Journal of the American Statistical Association 77(379): 605-610.

Dawid, A. P. and Lauritzen, S. L. (1993). Hyper Markov laws in the statistical analysis of decomposable graphical models, The Annals of Statistics 21(3): 1272-1317.

Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation, Computational Intelligence 5: 142-150.

DeGroot, M. H. (1970). Optimal Statistical Decisions, McGraw-Hill, New York.

Drivsholm, T., Hansen, T., Urhammer, S. A., Palacios, R. T., Vølund, A., BorchJohnsen, K. and Pedersen, O. B. (2003). Assessment of insulin sensitivity index and acute insulin reponse from an oral glucose tolerance test in subjects with normal glucose tolerance, Novo Nordisk A/S.

Edwards, D. (1995). Introduction to Graphical Modelling, Springer-Verlag, New York.

Friedman, N., Murphy, K. P. and Russell, S. (1998). Learning the Structure of Dynamic Probabilistic Networks, Proceedings of Fourteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 139-147.

Frydenberg, M. (1990). Marginalization and collapsibility in graphical interaction models, Annals of Statistics 18: 790-805.

Furnival, G. M. and Wilson, R. W. (1974). Regression by Leaps and Bounds, Technometrics 16(4): 499-511.

Geiger, D. and Heckerman, D. (1994). Learning Gaussian Networks, Proceedings of Tenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 235-243.

Harrison, P. J. and Stevens, C. F. (1976). Bayesian forecasting, Journal of Royal Statistics 38: 205-247.

Haughton, D. M. A. (1988). On The Choice of a Model to fit Data From an Exponential Family, The Annals of Statistics 16(1): 342-355.

Heckerman, D. (1999). A Tutorial on Learning with Bayesian Networks, in M. Jordan (ed.), Learning in Graphical Models, MIT Press, Cambridge, MA.

Heckerman, D., Geiger, D. and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data, Machine Learning 20: 197-243.

Ihaka, R. and Gentleman, R. (1996). R: A language for data analysis and graphics, Journal of Computational and Graphical Statistics 5: 299-314.

Jensen, F. V. (1996). An Introduction to Bayesian Networks, UCL Press, London.

Lauritzen, S. L. (1992). Propagation of probabilities, means and variances in mixed graphical association models, Journal of the American Statistical Association 87(420): 1098-1108.

Lauritzen, S. L. (1996). Graphical Models, Clarendon press, Oxford, New York.
Lauritzen, S. L. (2003). Some modern applications of graphical models, in P. J. Green, N. L. Hjort and S. Richardson (eds), Highly Structured Stochastic Systems, Oxford University Press, Oxford, pp. 13-32.

Lauritzen, S. L. and Jensen, F. (2001). Stable local computation with conditional Gaussian distributions, Statistics and Computing 11: 191-203.

Lauritzen, S. L. and Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems (with discussion), J. Royal Statist. Soc. Series B.

Martin, B. C., Warram, J. H., Krolewski, A. S., Bergman, R. N., Soeldner, J. S. and Kahn, C. R. (1992). Role of glucose and insulin resistance in development of type 2 diabetes mellitus: results of a 25 -year follow-up study, The Lancet 340: 925-929.

Morrison, D. F. (1976). Multivariate Statistical Methods, McGraw-Hill, USA.

Murphy, K. P. (2002). Dynamic Bayesian Networks: Representation, Inference and Learning, PhD thesis, University of California, Berkeley.

Pacini, G. and Bergman, R. N. (1986). MINMOD: a computer program to calculate insulin sensitivity and pancreatic responsivity from the frequently sampled intraveneous glucose tolerance test, Computer Methods and Programs in Biomedicine 23: 113-122.

Pearl, J. (1988). Probabilistic Reasoning in Intelligence Systems: Networks of Plausible Inference, Morgan Kaufmann, Los Altos, CA.

R Development Core Team (2003). R: A language and environment for statistical computing, R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-00-3.

Robinson, R. W. (1977). Counting unlabeled acyclic digraphs, Lecture Notes in Mathematics, 622: Combinatorial Mathematics V pp. 239-273.

Schaerf, M. C. (1964). Estimation of the covariance and autoregressive structure of a stationary time series, Technical report, Department of Statistics, Stanford University.

Seber, G. A. F. (1984). Multivariate Observations, John Wiley and Sons, New York.

Shachter, R. D. and Kenley, C. R. (1989). Gaussian influence diagrams, Management Science 35: 527-550.

Spiegelhalter, D. J. and Lauritzen, S. L. (1990). Sequential updating of conditional probabilities on directed graphical structures, Networks 20: 579605.

Steck, H. and Jaakkola, T. S. (2002). On the Dirichlet Prior and Bayesian Regularization, Conference on Advances in Neural Information Processing Systems, Vol. 15, MIT Press, Cambridge, USA.

Tong, H. (1996). Non-Linear Time Series, Clarendon Press, Oxford.
Venables, W. N. and Ripley, B. D. (1997). Modern Applied Statistics with S$P L U S$, second edn, Springer-Verlag, New York.

West, M. and Harrison, J. (1989). Bayesian Forecasting and Dynamic Models, Springer-Verlag, New York.

Yule, G. U. (1927). On a method for investigating periodicities in disturbed series with special reference to Wölfer's sunspot numbers, Philosophical Transactions of the Royal Society, Series A 226: 267-298.

# Paper I <br> Learning Conditional Gaussian Networks 

Susanne G. Bøttcher

Many of the results in this paper have been published in
Proceedings of the Eight International Workshop on
Artificial Intelligence and Statistics (2001),
pp. 149-156.

# Learning Conditional Gaussian Networks 

Susanne G. Bøttcher<br>Aalborg University, Denmark


#### Abstract

. This paper considers conditional Gaussian networks. The parameters in the network are learned by using conjugate Bayesian analysis. As conjugate local priors, we apply the Dirichlet distribution for discrete variables and the Gaussian-inverse gamma distribution for continuous variables, given a configuration of the discrete parents. We assume parameter independence and complete data. Further, to learn the structure of the network, the network score is deduced. We then develop a local master prior procedure, for deriving parameter priors in these networks. This procedure satisfies parameter independence, parameter modularity and likelihood equivalence. Bayes factors to be used in model search are introduced. Finally the methods derived are illustrated by a simple example.


## 1 Introduction

The aim of this paper is to present a method for learning the parameters and structure of a Bayesian network with discrete and continuous variables. In Heckerman et al. (1995) and Geiger and Heckerman (1994), this was done for respectively discrete networks and Gaussian networks.

We define the local probability distributions such that the joint distribution of the random variables is a conditional Gaussian (CG) distribution. Therefore we do not allow discrete variables to have continuous parents, so the network factorizes into a discrete part and a mixed part. The local conjugate parameter priors are for the discrete part of the network specified as Dirichlet distributions and for the mixed part of the network as Gaussian-inverse gamma distributions, for each configuration of discrete parents.

To learn the structure, $D$, of a network from data, $d$, we use the network score, $p(d, D)$, as a measure of how probable $D$ is. To be able to calculate this score for all possible structures, we derive a method for finding the prior distribution of the parameters in the possible structures, from marginal priors calculated from an imaginary database. The method satisfies parameter independence, parameter modularity and likelihood equivalence. If used on networks with only
discrete or only continuous variables, it coincides with the methods developed in Heckerman et al. (1995) and Geiger and Heckerman (1994).

When many structures are possible, some kind of strategy to search for the structure with the highest score, has to be applied. In Cooper and Herskovits (1992), different search strategies are presented. Many of these strategies use Bayes factors for comparing the network scores of two different networks that differ by the direction of a single arrow or by the presence of a single arrow. We therefore deduce the Bayes factors for these two cases. To reduce the number of comparisons needed, we identify classes of structures for which the corresponding Bayes factor for testing an arrow between the same two variables in a network, is the same.

Finally a simple example is presented to illustrate some of the methods developed.

In this paper, we follow standard convention for drawing a Bayesian network and use shaded nodes to represent discrete variables and clear nodes to represent continuous variables.

The results in Section 2 to Section 7 are also published in Bøttcher (2001).

## 2 Bayesian Networks

A Bayesian network is a graphical model that encodes the joint probability distribution for a set of variables $X$. For terminology and theoretical aspects on graphical models, see Lauritzen (1996). In this paper we define it as consisting of

- A directed acyclic graph (DAG) $D=(V, E)$, where $V$ is a finite set of vertices and $E$ is a finite set of directed edges between the vertices. The DAG defines the structure of the Bayesian network.
- To each vertex $v \in V$ in the graph corresponds a random variable $X_{v}$, with state space $\mathcal{X}_{v}$. The set of variables associated with the graph $D$ is then $X=\left(X_{v}\right)_{v \in V}$. Often we do not distinguish between a variable $X_{v}$ and the corresponding vertex $v$.
- To each vertex $v$ with parents $\mathrm{pa}(v)$, there is attached a local probability distribution, $p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)$. The set of local probability distributions for all variables in the network is denoted $\mathcal{P}$.
- The possible lack of directed edges in $D$ encodes conditional independencies between the random variables $X$ through the factorization of the joint probability distribution,

$$
p(x)=\prod_{v \in V} p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)
$$

A Bayesian network for a set of random variables $X$ is thus the pair $(D, \mathcal{P})$. In order to specify a Bayesian network for $X$, we must therefore specify a DAG $D$ and a set $\mathcal{P}$ of local probability distributions.

## 3 Bayesian Networks for Mixed Variables

In this paper we are interested in specifying networks for random variables $X$ of which some are discrete and some are continuous. So we consider a DAG $D=(V, E)$ with vertices $V=\Delta \cup \Gamma$, where $\Delta$ and $\Gamma$ are the sets of discrete and continuous vertices, respectively. The corresponding random variables $X$ can then be denoted $X=\left(X_{v}\right)_{v \in V}=(I, Y)=\left(\left(I_{\delta}\right)_{\delta \in \Delta},\left(Y_{\gamma}\right)_{\gamma \in \Gamma}\right)$, i.e. we use $I$ and $Y$ for the sets of discrete and continuous variables, respectively. We denote the set of levels for each discrete variable $\delta \in \Delta$ as $\mathcal{I}_{\delta}$.

In this paper we do not allow discrete variables to have continuous parents. This e.g. ensures availability of exact local computation methods, see Lauritzen (1992) and Lauritzen and Jensen (2001). The joint probability distribution then factorizes as follows:

$$
p(x)=p(i, y)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) \prod_{\gamma \in \Gamma} p\left(y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}\right)
$$

where $i_{\mathrm{pa}(\gamma)}$ and $y_{\mathrm{pa}(\gamma)}$ denote observations of the discrete and continuous parents respectively, i.e. $i_{\mathrm{pa}(\gamma)}$ is an abbreviation of $i_{\mathrm{pa}(\gamma) \cap \Delta}$ etc.

We see that the joint probability distribution factorizes into a purely discrete part and a mixed part. First we look at the discrete part.

### 3.1 The Discrete Part of the Network

We assume that the local probability distributions are unrestricted discrete distributions with

$$
p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) \geq 0 \quad \forall \quad \delta \in \Delta
$$

A way to parameterize this is to let

$$
\begin{equation*}
\theta_{i_{\delta}| |_{\mathrm{pa} a}(\delta)}=p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta| |_{\mathrm{pa} a}(\delta)}\right), \tag{1}
\end{equation*}
$$

where $\theta_{\delta| |_{\mathrm{pa}(\delta)}}=\left(\theta_{\left.i_{\delta}\right|_{\mathrm{p}_{\mathrm{pa}}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}$.
Then $\sum_{i_{\delta} \in \mathcal{I}_{\delta}} \theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=1$ and $0 \leq \theta_{i_{\delta} \mid i_{\mathrm{pa}}(\delta)} \leq 1$. All parameters associated with a node $\delta$ is denoted $\theta_{\delta}$, i.e. $\theta_{\delta}=\left(\theta_{\delta| |_{\mathrm{pa}(\delta)}}\right)_{\mathrm{p}_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}}$.

Using this parameterization, the discrete part of the joint probability distribution is given by

$$
p\left(i \mid\left(\theta_{\delta}\right)_{\delta \in \Delta}\right)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta| |_{\mathrm{pa}(\delta)}}\right) .
$$

### 3.2 The Mixed Part of the Network

Now consider the mixed part. We assume that the local probability distributions are Gaussian linear regressions on the continuous parents, with parameters depending on the configuration of the discrete parents. Let the parameters in the distribution be given by $\theta_{\left.\gamma\right|_{\mathrm{pa}}(\gamma)}=\left(m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}, \beta_{\left.\gamma\right|_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)$. Then

$$
\begin{equation*}
\left(Y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}, \theta_{\gamma| |_{\mathrm{pa}}(\gamma)}\right) \sim \mathcal{N}\left(m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}+\beta_{\gamma \mid i_{\mathrm{pa}}(\gamma)} y_{\mathrm{pa}(\gamma)}, \sigma_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{2}\right), \tag{2}
\end{equation*}
$$

where $\beta_{\gamma \mid i_{\mathrm{pa}}(\gamma)}$ are the regression coefficients, $m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}$ is the regression intercept, and $\sigma_{\left.\gamma \mid i_{\mathrm{p}(\gamma)}\right)}^{2}$ is the conditional variance. Thus for each configuration of the discrete parents of $\gamma$, the distribution of $Y_{\gamma}$ is Gaussian with mean and variance given as in (2). There are three special cases of the above situation, namely when $\gamma$ has no discrete parents, when it has no continuous parents and when it has no parents at all. If it has no discrete parents, (2) is just the Gaussian distribution,

$$
\left(Y_{\gamma} \mid y_{\mathrm{pa}(\gamma)}, \theta_{\gamma}\right) \sim \mathcal{N}\left(m_{\gamma}+\beta_{\gamma} y_{\mathrm{pa}(\gamma)}, \sigma_{\gamma}^{2}\right),
$$

and $\theta_{\gamma}=\left(m_{\gamma}, \beta_{\gamma}, \sigma_{\gamma}^{2}\right)$. When $\gamma$ has no continuous parents, we have

$$
\left(Y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) \sim \mathcal{N}\left(m_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}, \sigma_{\gamma \mid \mathrm{p}_{\mathrm{pa}(\gamma)}}^{2}\right),
$$

with $\theta_{\gamma \mid i_{\mathrm{pa}}(\gamma)}=\left(m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}, \sigma_{\left.\gamma\right|_{\mathrm{p} a}(\gamma)}^{2}\right)$, i.e. for each $\gamma$, the mean depends solely on $i_{\mathrm{pa}(\gamma)}$. Finally, when $\gamma$ has no parents at all,

$$
\left(Y_{\gamma} \mid \theta_{\gamma}\right) \sim \mathcal{N}\left(m_{\gamma}, \sigma_{\gamma}^{2}\right),
$$

with $\theta_{\gamma}=\left(m_{\gamma}, \sigma_{\gamma}^{2}\right)$.
With $\theta_{\gamma}=\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}}$, the mixed part of the joint distribution can be written as

$$
p\left(y \mid i,\left(\theta_{\gamma}\right)_{\gamma \in \Gamma}\right)=\prod_{\gamma \in \Gamma} p\left(y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)
$$

### 3.3 The Joint Network

If we let $\theta=\left(\left(\theta_{\delta}\right)_{\delta \in \Delta},\left(\theta_{\gamma}\right)_{\gamma \in \Gamma}\right)$, the joint probability distribution for $X=$ $(I, Y)$ is given by

$$
\begin{equation*}
p(x \mid \theta)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \prod_{\gamma \in \Gamma} p\left(y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}, \theta_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}\right) \tag{3}
\end{equation*}
$$

It can easily be shown by induction that when the local probability distributions are given as defined in (1) and (2), the joint probability distribution for $X$ is a CG distribution with density of the form

$$
p(x \mid \theta)=p(i, y \mid \theta)=p(i)\left|2 \pi \Sigma_{i}\right|^{-\frac{1}{2}} \exp \left\{-\frac{1}{2}\left(y-M_{i}\right)^{\mathrm{T}} \Sigma_{i}^{-1}\left(y-M_{i}\right)\right\}
$$

For each $i, M_{i}$ is the unconditional mean, that is unconditional on continuous variables and $\Sigma_{i}$ is the covariance matrix for all the continuous variables in the network. In Shachter and Kenley (1989) formulas for calculating $\Sigma_{i}$ from the local probability distributions can be found.

A Bayesian network, where the joint probability distribution is a CG distribution is in the following called a CG network.

## 4 Learning the Parameters in a CG Network

When constructing a Bayesian network there is, as mentioned earlier, two things to consider, namely specifying the DAG and specifying the local probability distributions. In this section we assume that the structure of the DAG is known and the distribution type is given as in the previous section and we consider the specification of the parameters in the distributions. For this we need the concept of conjugate Bayesian analysis.

### 4.1 Conjugate Bayesian Analysis

There are several ways of assessing the parameters in probability distributions. An expert could specify them, or they could be estimated from data. In our approach we encode our uncertainty about the parameter $\theta$ in a prior distribution $p(\theta)$, use data to update this distribution, i.e. learn the parameter and hereby, by using Bayes' theorem, obtain the posterior distribution $p(\theta \mid$ data $)$, see DeGroot (1970).

Consider a situation with one random variable $X$. Let $\theta$ be the parameter to be assessed, $\Theta$ the parameter space and $d$ a random sample of size $n$ from the probability distribution $p(x \mid \theta)$. We call $d$ our database and $x^{c} \in d$ a case. Then, according to Bayes' theorem,

$$
\begin{equation*}
p(\theta \mid d)=\frac{p(d \mid \theta) p(\theta)}{p(d)}, \quad \theta \in \Theta, \tag{4}
\end{equation*}
$$

where $p(d \mid \theta)=\prod_{x^{c} \in d} p\left(x^{c} \mid \theta\right)$ is the joint probability distribution of $d$, also called the likelihood of $\theta$. Furthermore the denominator is given by

$$
p(d)=\int_{\Theta} p(d \mid \theta) p(\theta) d \theta
$$

and for fixed $d$ it may be considered as a normalizing constant. Therefore (4) can be expressed as

$$
p(\theta \mid d) \propto p(d \mid \theta) p(\theta),
$$

where the proportionality constant is determined by the relation $\int_{\Theta} p(\theta \mid d) d \theta=$ 1.

When the prior distribution belongs to a given family of distributions and the posterior distribution, after sampling from a specific distribution, belongs to the same family of distributions, then this family is said to be closed under sampling and called a conjugate family of distributions. Further, if a parameter or the distribution of a parameter has a certain property which is preserved under sampling, then this property is said to be a conjugate property.

In a conjugate family of distributions it is generally straightforward to calculate the posterior distribution.

### 4.2 Some Simplifying Properties

In the previous section we showed how to update a prior distribution for a single parameter $\theta$. In a Bayesian network with more than one variable, we also have to look at the relationship between the different parameters for the different variables in the network. In this paper we assume that the parameters associated with one variable is independent of the parameters associated with the other variables. This assumption was introduced by Spiegelhalter and Lauritzen (1990) and we denote it global parameter independence. In addition to this, we will assume that the parameters are independent for each configuration of the discrete parents, which we denote as local parameter independence. So if the parameters have the property of global parameter independence and local parameter independence, then

$$
\begin{equation*}
p(\theta)=\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\left.\delta \mid i_{\mathrm{pa}(\delta)}\right)} \prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}\right)\right. \tag{5}
\end{equation*}
$$

and we will refer to (5) simply as parameter independence.
A consequence of parameter independence is that, for each configuration of the discrete parents, we can update the parameters in the local distributions independently. This also means that if we have local conjugacy, i.e. the distributions of $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ and $\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ belongs to a conjugate family, then because of parameter independence, we have global conjugacy, i.e. the joint distribution of $\theta$ belongs to a conjugate family.

Further, we will assume that the database $d$ is complete, that is, in each case it contains at least one instance of every random variable in the network. With this we can show that parameter independence is a conjugate property.

Due to the factorization (3) and the assumption of complete data,

$$
\begin{aligned}
p(d \mid \theta) & =\prod_{c \in d} p\left(x^{c} \mid \theta\right) \\
& =\prod_{c \in d}\left(\prod_{\delta \in \Delta} p\left(i_{\delta}^{c} \mid i_{\mathrm{pa}(\delta)}^{c}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \prod_{\gamma \in \Gamma} p\left(y_{\gamma}^{c} \mid y_{\mathrm{pa}(\gamma)}^{c}, i_{\mathrm{pa}(\gamma)}^{c}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)\right)
\end{aligned}
$$

where $i^{c}$ and $y^{c}$ respectively denotes the discrete part and the continuous part of
a case $x^{c}$. Another way of writing the above equation is

$$
\begin{align*}
p(d \mid \theta)= & \prod_{\delta \in \Delta i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \prod_{c: i i_{\mathrm{pa}(\delta)}^{c}=i_{\mathrm{pa}(\delta)}} p\left(i_{\delta}^{c} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \\
& \times \prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \prod_{\mathrm{pa}(\gamma)}^{c \cdot i_{\mathrm{pa}}^{c}} \sum_{\mathrm{pa}(\gamma)} p\left(y_{\gamma}^{c} \mid y_{\mathrm{pa}(\gamma)}^{c}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right), \tag{6}
\end{align*}
$$

where the product over cases is split up into a product over the configurations of the discrete parents and a product over those cases, where the configuration of the discrete parents is the same as the currently processed configuration. Notice however that some of the parent configurations might not be represented in the database, in which case the product over cases with this parent configuration just adds nothing to the overall product.

By combining (5) and (6) it is seen that

$$
p(\theta \mid d)=\prod_{\delta \in \Delta i_{\mathrm{pa}}(\delta) \in \mathcal{I}_{\mathrm{pa}}(\delta)} \prod_{\delta \mid} p\left(\theta_{\left.\delta\right|_{\mathrm{pa} a}(\delta)} \mid d\right) \prod_{\gamma \in \Gamma i_{\mathrm{pa}}(\gamma) \in \mathcal{I}_{\mathrm{pa}}(\gamma)} p\left(\theta_{\gamma| |_{\mathrm{pa}}(\gamma)} \mid d\right),
$$

i.e. the parameters remain independent given data. We call this property posterior parameter independence. In other words, the properties of local and global independence are conjugate.

Notice that the posterior distribution, $p(\theta \mid d)$, can be found using batch learning or sequential learning. In batch learning, $p(\theta \mid d)$ is found by updating $p(\theta)$ with all cases in $d$ at the same time, i.e. in a batch. In sequential learning, $p(\theta)$ is updated one case at a time, using the previous posterior distribution as the prior distribution for the next case to be considered. When the database $d$ is complete, batch learning and sequential learning leads to the same posterior distribution and the final result is independent of the order in which the cases in $d$ are processed. It is of course also possible to process some of the cases in a batch and the rest sequentially, which could be done if e.g. a new case is added to an already processed database, see Bernardo and Smith (1994).

### 4.3 Learning in the Discrete Case

We now consider batch learning of the parameters in the discrete part of the network. Recall that the local probability distributions are unrestricted discrete distributions defined as in (1). As pointed out in the previous section we can,
because of the assumption of parameter independence, find the posterior distribution of $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ for each $\delta$ and each configuration of $\mathrm{pa}(\delta)$ independently.

So given a specific configuration of $i_{\mathrm{pa}(\delta)}$, we need to find $p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid d\right)$. From Bayes' theorem, Equation (4), we have that

$$
\begin{equation*}
p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid d\right) \propto \prod_{c: i_{\mathrm{pa}(\delta)}^{c}=i_{\mathrm{pa}(\delta)}} p\left(i_{\delta}^{c} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \tag{7}
\end{equation*}
$$

A conjugate family for multinomial observations is the family of Dirichlet distributions. So let the prior distribution of $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ be a Dirichlet distribution $\mathcal{D}$ with hyperparameters $\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}$, also written as

$$
\begin{equation*}
\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid \alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \sim \mathcal{D}\left(\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \tag{8}
\end{equation*}
$$

The probability function for this Dirichlet distribution is given by

$$
p\left(\theta_{\delta \mid i_{\mathrm{pa}}(\delta)} \mid \alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)=\frac{\Gamma\left(\alpha_{+\delta \mid i_{\mathrm{pa}}(\delta)}\right)}{\prod_{i_{\delta} \in \mathcal{I}_{\delta}} \Gamma\left(\alpha_{\left.i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right)}\right)} \prod_{i_{\delta} \in \mathcal{I}_{\delta}}\left(\theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)^{\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}-1}
$$

where $\alpha_{+\delta \mid i_{\mathrm{pa}(\delta)}}=\sum_{i_{\delta} \in \mathcal{I}_{\delta}} \alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}$ and $\Gamma(\cdot)$ is the gamma function. Because of notational convenience, we do not in what follows write the hyperparameters explicitly in the conditioning.

It then follows from (7) and (8) that the posterior distribution is given as

$$
\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid d\right) \sim \mathcal{D}\left(\alpha_{\delta \mid i_{\mathrm{pa}}(\delta)}+n_{\delta \mid i_{\mathrm{pa}}(\delta)}\right)
$$

where the vector $n_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta \in \mathcal{I}_{\delta}}}$, also called the counts, denotes the number of observations in $d$ where $\delta$ and $\mathrm{pa}(\delta)$ have that specific configuration. Notice that, for at given parent configuration, the number of observations in a batch, $|b|$, is the same as $n_{+_{\delta} \mid i_{\mathrm{pa}(\delta)}}$, where $n_{+_{\delta} \mid i_{\mathrm{pa}(\delta)}}=\sum_{i_{\delta} \in \mathcal{I}_{\delta}} n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}$.
Because of parameter independence, the joint prior distribution of all the parameters for the discrete variables in the network, is given by the product of the local parameter priors.

The above learning procedure can also be used for sequential learning by applying the above formulas one case at a time, using the previous posterior distribution as the prior distribution for the next case to be processed.

### 4.4 Learning in the Mixed Case

In the mixed case we write the local probability distributions as

$$
\left(Y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) \sim \mathcal{N}\left(z_{\mathrm{pa}(\gamma)}\left(m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)^{\mathrm{T}}, \sigma_{\gamma \mid \mathrm{p}_{\mathrm{pa}(\gamma)}}^{2}\right),
$$

where $z_{\mathrm{pa}(\gamma)}=\left(1, y_{\mathrm{pa}(\gamma)}\right)$. This vector has dimension $k+1$, where $k$ is the number of continuous parents to $\gamma$.

As in the discrete case we can because of parameter independence update the parameters for each $\gamma$ and each configuration of the discrete parents independently. By Bayes' theorem,

$$
p\left(\theta_{\gamma\left|i_{\mathrm{pa}(\gamma)}\right|} \mid d\right) \propto \prod_{c: i_{\mathrm{pa}}^{c}(\gamma)=i_{\mathrm{pa}(\gamma)}} p\left(y_{\gamma}^{c} \mid y_{\mathrm{pa}(\gamma)}^{c}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) p\left(\theta_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}\right) .
$$

We now join all the observations $y_{\gamma}^{c}$ for which $i_{\mathrm{pa}(\gamma)}^{c}=i_{\mathrm{pa}(\gamma)}$ in a vector $y_{\gamma}^{b}$, i.e. $y_{\gamma}^{b}=\left(y_{\gamma}^{c} i_{\mathrm{pa}(\gamma)}^{c}=i_{\mathrm{pa}(\gamma)}\right.$. The same is done with the observations of the continuous parents of $\gamma$, i.e. $y_{\mathrm{pa}(\gamma)}^{b}=\left(y_{\mathrm{pa}(\gamma)}^{c}\right) i_{\mathrm{pa}(\gamma)}^{c}=i_{\mathrm{pa}(\gamma)}$. As the observations in $d$ are independent, $p\left(y_{\gamma}^{b} \mid y_{\mathrm{pa}(\gamma)}^{b}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma| |_{\mathrm{pa}}(\gamma)}\right)$ is the likelihood function for a multivariate normal distribution with mean vector $z_{\mathrm{pa}(\gamma)}^{b}\left(m_{\gamma \mid i_{\mathrm{pa}}(\gamma)}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)^{\mathrm{T}}$ and covariance matrix $\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2} I$, where $I$ is the identity matrix and $z_{\mathrm{pa}(\gamma)}^{b}$ is defined through $y_{\mathrm{pa}(\gamma)}^{b}$.
The posterior distribution of $\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ can now be written as

$$
p\left(\theta_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)} \mid d\right) \propto p\left(y_{\gamma}^{b} \mid y_{\mathrm{pa}(\gamma)}^{b}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}}(\gamma)}\right) p\left(\theta_{\left.\gamma| |_{\mathrm{pa}(\gamma)}\right)}\right.
$$

A standard conjugate family for these observations is the family of Gaussianinverse gamma distributions. Let the prior joint distribution of $\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)$ and $\sigma_{\gamma \mid \dot{p}_{\mathrm{p}}(\gamma)}^{2}$ be as follows.

$$
\begin{aligned}
\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid \sigma_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}( }^{2}\right) & \sim \mathcal{N}_{k+1}\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}^{2} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\right) \\
\left(\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right) & \sim \mathcal{I} \Gamma\left(\frac{\rho_{\gamma \mid i_{\mathrm{pa}}(\gamma)}}{2}, \frac{\phi_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}^{2}}{2}\right)
\end{aligned}
$$

The posterior distribution is then

$$
\begin{aligned}
\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}, d\right) & \sim \mathcal{N}_{k+1}\left(\mu_{\gamma| |_{\mathrm{pa}(\gamma)}^{\prime}}^{\prime}, \sigma_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{2}\left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\right)^{\prime}\right) \\
\left(\sigma_{\gamma \mid \mathrm{p}_{\mathrm{pa}(\gamma)}}^{2} \mid d\right) & \sim \mathcal{I} \Gamma\left(\frac{\rho_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{\prime}}{2}, \frac{\phi_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}^{2}}{2}\right),
\end{aligned}
$$

where

$$
\begin{aligned}
& \tau_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{\prime}=\tau_{\gamma| |_{\mathrm{pa}}(\gamma)}+\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}} z_{\mathrm{pa}(\gamma)}^{b} \\
& \mu_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{\prime}=\left(\tau_{\gamma \mid \mathrm{p}_{\mathrm{pa}}(\gamma)}^{\prime}\right)^{-1}\left(\tau_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}} \mu_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}+\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}} y_{\gamma}^{b}\right) \\
& \rho_{\gamma \mid i_{\mathrm{pa}}(\gamma)}^{\prime}=\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+|b| \\
& \phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}=\phi_{\gamma| |_{\mathrm{pa}}(\gamma)}+\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid \mathrm{p}_{\mathrm{pa}}(\gamma)}^{\prime}\right)^{\mathrm{T}} y_{\gamma}^{b} \\
& +\left(\mu_{\gamma \mid i_{\mathrm{p} a}(\gamma)}-\mu_{\gamma \mid i_{\mathrm{p} a}(\gamma)}^{\prime}\right)^{\mathrm{T}} \tau_{\gamma| |_{\mathrm{pa}(\gamma)}} \mu_{\gamma| |_{\mathrm{pa}(\gamma)}},
\end{aligned}
$$

where $|b|$ denotes the number of observations in $b$.
As for the discrete variables, we can with these formulas also use the sequential approach and update the parameters one case at a time.

Further, because of parameter independence, the joint prior distribution is given as the product of the local prior distributions for all parameters in the network.

## 5 Learning the Structure of a CG Network

In this section we consider how to learn the structure of a CG network.

### 5.1 The Network Score

There are basically two ways of determining which DAG should represent the conditional independencies between a set of random variables. First, if the relations between the variables are well understood by an expert, then he could specify the DAG, using a causal interpretation of the arrows. Second, we could learn the DAG from data. That is, we could find out how well a DAG $D$ represents the conditional independencies, by measuring how probable $D$ is, given that we have observed data $d$. Different approaches use different measures. An
often used measure is the posterior probability of the DAG, $p(D \mid d)$, which from Bayes' theorem is given by

$$
p(D \mid d) \propto p(d \mid D) p(D)
$$

where $p(d \mid D)$ is the likelihood of $D$ and $p(D)$ is the prior probability. As the normalizing constant does not depend upon structure, another measure, which gives the relative probability, is

$$
p(D, d)=p(d \mid D) p(D)
$$

We refer to the above measures as network scores. So learning the DAG from data, we can in principle first calculate the network scores for all possible DAGs and then select the DAG with the highest network score. If many DAGs are possible, it is computationally infeasible to calculate the network score for all these DAGs. In this situation it is necessary to use some kind of search strategy to find the DAG with the highest score, see e.g. Cooper and Herskovits (1992).

In some cases it can be more accurate to average over the possible DAGs for prediction, instead of just selecting a single DAG. So if $x$ is the quantity we are interested in, we can use the weighted average,

$$
p(x \mid d)=\sum_{D \in D A G} p(x \mid d, D) p(D \mid d)
$$

where $D A G$ is the set of all DAGs and $p(D \mid d)$ is the weight.
Again, if many DAGs are possible, this sum is to heavy to compute, so instead, by using a search strategy, we can find a few DAGs with high score and average over these.

### 5.2 The Network Score for a CG Network

In order to calculate the network score for a specific DAG $D$, we need to know the prior probability and the likelihood of the DAG. For simplicity, we could for example choose to let all DAGs be equally likely, then

$$
p(D \mid d) \propto p(d \mid D)
$$

In a CG network, the likelihood of the DAG $D$ is given by

$$
\begin{aligned}
& p(d \mid D)=\int_{\theta \in \Theta} p(d \mid \theta, D) p(\theta \mid D) d \theta \\
& =\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \int_{c: i_{\mathrm{pa}(\delta)}^{c}} \prod_{\mathrm{pa}(\delta)} p\left(i_{\delta}^{c} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}, D\right) p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid D\right) d \theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \\
& \times \prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \int_{c: i_{\mathrm{pa}(\gamma)}^{c}} \prod_{=i_{\mathrm{pa}(\gamma)}} p\left(y_{\gamma}^{c} \mid y_{\mathrm{pa}(\gamma)}^{c}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, D\right) p\left(\theta_{\gamma\left|i_{\mathrm{pa}(\gamma)}\right|} \mid D\right) d \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} .
\end{aligned}
$$

Again we see that we can consider the problem for the discrete part and the mixed part of the network separately.

The discrete part is from the formulas in Section 4.3 found to be

$$
\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \frac{\Gamma\left(\alpha_{+\delta} \mid i_{\mathrm{pa}(\delta)}\right)}{\Gamma\left(\alpha_{+\delta}\left|i_{\mathrm{pa}(\delta)}+n_{+\delta}\right| i_{\mathrm{pa}(\delta)}\right)} \prod_{i_{\delta} \in \mathcal{I}_{\delta}} \frac{\Gamma\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}+n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)}{\Gamma\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)} .
$$

In the mixed part of the network, the local marginal likelihoods are non-central $t$ distributions with $\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ degrees of freedom, location vector $z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ and scale parameter $s_{\gamma \mid i_{\mathrm{pa}(\gamma)}}=\frac{\phi_{\gamma \mid i_{\mathrm{pa}}(\gamma)}}{\rho_{\gamma \mid i_{\mathrm{pa}}(\gamma)}}\left(I+\left(z_{\mathrm{pa}(\gamma)}^{b}\right) \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}}\right)$. The index $b$ is defined as in Section 4.4.

So the mixed part is given by

$$
\begin{gathered}
\prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \frac{\Gamma\left(\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+|b|\right) / 2\right)}{\Gamma\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} / 2\right)\left[\operatorname{det}\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} s_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \pi\right)\right]^{\frac{1}{2}}} \times \\
{\left[1+\frac{1}{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) s_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)^{\mathrm{T}}\right]^{\frac{-\left(\rho_{\gamma\left|i_{\mathrm{pa}(\gamma)}+|b|\right)}^{2}\right.}{2}} .}
\end{gathered}
$$

The network score for a CG network is thus the product of the prior probability for the DAG $D$, the term for the discrete part and the term for the mixed part. Notice that the network score has the property that it factorizes into a product over terms involving only one node and its parents. This property is called decomposability.

To evaluate which DAG or possible several DAGs that represent the conditional independencies in a Bayesian network well, we want to find the DAG or DAGs
with the highest network scores. To calculate these scores, we must specify the local probability distributions and the local prior distributions for the parameters for each network under evaluation. In the next section, a method for doing this is developed.

## 6 The Master Prior Procedure

The papers Heckerman et al. (1995) and Geiger and Heckerman (1994) develops a method for finding the prior distributions for the parameters in respectively the purely discrete case and the purely continuous case. The work is based on principles of likelihood equivalence, parameter modularity, and parameter independence. It leads to a method where the parameter priors for all possible networks are deduced from one joint prior distribution, in the following called a master prior distribution.

In this paper we will build on this idea, which can be used on networks with mixed variables. We will therefore in the following describe their method for the pure cases.

### 6.1 The Master Prior in the Discrete Case

In the purely discrete case, or the discrete part of a mixed network, the following is a well known classical result.

Let $A$ be a subset of $\Delta$ and let $B=\Delta \backslash A$. Let the discrete variables $i$ have the joint distribution

$$
p(i \mid \Psi)=\Psi_{i}
$$

Notice here, that the set $\Psi=\left(\Psi_{i}\right)_{i \in \mathcal{I}}$ contains the parameters for the joint distribution, contrary to $\theta$ in Section 3, which contains the parameters for the conditional local distributions.

In the following we use the notation $z_{i_{A}}=\sum_{j: j_{A}=i_{A}} z_{j}$, where $z$ is any parameter. Then the marginal distribution of $i_{A}$ is given by

$$
p\left(i_{A} \mid \Psi\right)=\Psi_{i_{A}}
$$

and the conditional distribution of $i_{B}$ given $i_{A}$ is

$$
p\left(i_{B} \mid i_{A}, \Psi\right)=\frac{\Psi_{i}}{\Psi_{i_{A}}}=\Psi_{i_{B} \mid i_{A}}
$$

Further if the joint prior distribution for the parameters $\Psi$ is Dirichlet, that is

$$
(\Psi) \sim \mathcal{D}(\alpha)
$$

where $\alpha=\left(\alpha_{i}\right)_{i \in \mathcal{I}}$, then the marginal distribution of $\Psi_{A}$ is Dirichlet, i.e.

$$
\left(\Psi_{A}\right) \sim \mathcal{D}\left(\alpha_{A}\right)
$$

with $\alpha_{A}=\left(\alpha_{i_{A}}\right)_{i_{A} \in \mathcal{I}_{A}}$. The conditional distribution of $\Psi_{B \mid i_{A}}$ is

$$
\left(\Psi_{B \mid i_{A}}\right) \sim \mathcal{D}\left(\alpha_{B \mid i_{A}}\right)
$$

with $\alpha_{B \mid i_{A}}=\left(\alpha_{i_{B} \mid i_{A}}\right)_{i_{B} \in \mathcal{I}_{B}}$ and $\alpha_{i_{B} \mid i_{A}}=\alpha_{i}$. Furthermore the parameters are independent, that is

$$
\begin{equation*}
p(\Psi)=\prod_{i_{A} \in \mathcal{I}_{A}} p\left(\Psi_{B \mid i_{A}}\right) p\left(\Psi_{A}\right) \tag{9}
\end{equation*}
$$

From the above result we see, that for each possible parent/child relationship, we can find the marginal parameter prior $p\left(\Psi_{\delta \cup p a(\delta)}\right)$. Further, from this marginal distribution we can, for each configuration of the parents, find the conditional local prior distribution $p\left(\Psi_{\left.\delta \mid i_{\mathrm{pa}(\delta)}\right)}\right.$. Notice that $\Psi_{\delta \mid i_{\mathrm{pa}(\delta)}}=\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$, where $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ was specified for the conditional distributions in Section (3.1). Further, because of parameter independence, given by (9), we can find the joint parameter prior for any network as the product of the local priors involved.

To use this method, we must therefore specify the joint Dirichlet distribution, i.e. the master Dirichlet prior. This was first done in Heckerman et al. (1995) and here we follow their method. We start by specifying a prior Bayesian network $(D, \mathcal{P})$. From this we calculate the joint distribution $p(i \mid \Psi)=\Psi_{i}$. To specify a master Dirichlet distribution, we must specify the parameters $\alpha=\left(\alpha_{i_{\delta}}\right)_{i \in \mathcal{I}}$ and for this we use the following relation for the Dirichlet distribution,

$$
p(i)=\mathbb{E}\left(\Psi_{i}\right)=\frac{\alpha_{i}}{n}
$$

with $n=\sum_{i \in \mathcal{I}} \alpha_{i}$. Now we let the probabilities in the prior network be an estimate of $\mathbb{E}\left(\Psi_{i}\right)$, so we only need to determine $n$ in order to calculate the parameters $\alpha_{i}$. We determine $n$ by using the notion of an imaginary database. We imagine that we have a database of cases, from which we from total ignorance have updated the distribution of $\Psi$. The sample size of this imaginary database is thus $n$. Therefore we refer to the estimate of $n$ as the imaginary sample size and it expresses how much confidence we have in the dependency structure expressed in the prior network.

### 6.2 The Master Prior in the Gaussian Case

For the Gaussian case, the following result is used, see e.g. Dawid and Lauritzen (1993). Let $A$ be a subset of $\Gamma$ and let $B=\Gamma \backslash A$. If

$$
(y \mid m, \Sigma) \sim \mathcal{N}(m, \Sigma)
$$

then

$$
\left(y_{A} \mid m, \Sigma\right) \sim \mathcal{N}\left(m_{A}, \Sigma_{A A}\right)
$$

and

$$
\left(y_{B} \mid y_{A}, m_{B \mid A}, \beta_{B \mid A}, \Sigma_{B \mid A}\right) \sim \mathcal{N}\left(m_{B \mid A}+\beta_{B \mid A} y_{A}, \Sigma_{B \mid A}\right),
$$

where

$$
\begin{gathered}
\Sigma=\left(\begin{array}{cc}
\Sigma_{A A} & \Sigma_{A B} \\
\Sigma_{B A} & \Sigma_{B B}
\end{array}\right), \Sigma_{B \mid A}=\Sigma_{B B}-\Sigma_{B A} \Sigma_{A A}^{-1} \Sigma_{A B}, \\
m_{B \mid A}=m_{B}-\beta_{B \mid A} m_{A} \text { and } \beta_{B \mid A}=\Sigma_{B A} \Sigma_{A A}^{-1} .
\end{gathered}
$$

Further, if

$$
(m \mid \Sigma) \sim \mathcal{N}\left(\mu, \frac{1}{\nu} \Sigma\right) \text { and }(\Sigma) \sim \mathcal{I} W(\rho, \Phi)
$$

where the scale matrix $\Phi$ is partitioned as $\Sigma$, then

- $\left(m_{A} \mid \Sigma_{A A}\right) \sim \mathcal{N}\left(\mu_{A}, \frac{1}{\nu} \Sigma_{A A}\right)$
- $\left(\Sigma_{A A}\right) \sim \mathcal{I} W\left(\rho, \Phi_{A A}\right)$
- $\left(\Sigma_{B \mid A}\right) \sim \mathcal{I} W\left(\rho+|A|, \Phi_{B \mid A}\right)$
- $\left(m_{B \mid A}, \beta_{B \mid A} \mid \Sigma_{B \mid A}\right) \sim \mathcal{N}\left(\mu_{B \mid A}, \Sigma_{B \mid A} \otimes \tau_{B \mid A}^{-1}\right)$
- $m_{A}, \Sigma_{A A} \Perp m_{B \mid A}, \beta_{B \mid A} \Sigma_{B \mid A}$
where

$$
\mu_{B \mid A}=\left(\mu_{B}-\Phi_{B A} \Phi_{A A}^{-1} \mu_{A}, \Phi_{B A} \Phi_{A A}^{-1}\right)
$$

and

$$
\tau_{B \mid A}^{-1}=\left(\begin{array}{cc}
\frac{1}{\nu}+\mu_{A}^{\mathrm{T}} \Phi_{A A}^{-1} \mu_{A} & -\mu_{A}^{\mathrm{T}} \Phi_{A A}^{-1} \\
-\Phi_{A A}^{-1} \mu_{A} & \Phi_{A A}^{-1}
\end{array}\right)
$$

and $\otimes$ denotes the Kronecker product. Notice that the dimension of $\mu_{B \mid A}$ is given as $(|B|,|B| \times|A|)$.

As in the discrete case, this result shows us how to deduce the local probability distributions and the local prior distributions from the joint distributions.

Further, because of parameter independence, the joint parameter prior for any Gaussian network can be specified as the product of the local priors. Notice that the parameters found here for a node given its parents, coincides with the parameters specified in Section 3.2.

Before we show how to construct the master prior, we need the following result. The Gaussian-inverse Wishart prior is conjugate to observations from a Gaussian distribution (DeGroot 1970). So let the probability distribution and the prior distribution be given as above. Then, given the database $d=\left\{y^{1}, \ldots, y^{n}\right\}$, the posterior distributions are

$$
(m \mid \Sigma, d) \sim \mathcal{N}\left(\mu^{\prime}, \frac{1}{\nu^{\prime}} \Sigma\right) \text { and }(\Sigma \mid d) \sim \mathcal{I} W\left(\rho^{\prime}, \Phi^{\prime}\right)
$$

where

$$
\begin{align*}
\nu^{\prime} & =\nu+n \\
\mu^{\prime} & =\frac{\nu \mu+n \bar{y}}{\nu+n}  \tag{10}\\
\rho^{\prime} & =\rho+n \\
\Phi^{\prime} & =\Phi+s s d+\frac{\nu n}{\nu+n}(\mu-\bar{y})(\mu-\bar{y})^{\mathrm{T}}
\end{align*}
$$

with

$$
\bar{y}=\frac{1}{n} \sum_{i=1}^{n} y_{i} \quad \text { and } \quad s s d=\sum_{i=1}^{n}\left(y_{i}-\bar{y}\right)\left(y_{i}-\bar{y}\right)^{\mathrm{T}} .
$$

From these updating formulas we see that $\nu^{\prime}$ and $\rho^{\prime}$ are updated with the number of cases in the database. Further $\mu^{\prime}$ is a weighted average of the prior mean and the sample mean, each weighted by their sample sizes. Finally $\Phi$ is updated with the $s s d$, which expresses how much each observation differs from the sample mean, and an expression for how much the prior mean differs from the sample mean.

To specify the master prior, we need to specify the four parameters $\nu, \mu, \rho$ and $\Phi$. As for the discrete variables we start by specifying a prior Bayesian network, $(D, \mathcal{P})$. From this, a prior joint probability distribution $p(y \mid m, \Sigma)=\mathcal{N}(m, \Sigma)$ can be deduced. Now imagine that the mean $m$ and the variance $\Sigma$ were calculated from an imaginary database, so that they actually are the sample mean and the sample variance. Further, assume that before this imaginary database was observed, we were totally ignorant about the parameters. The formulas in (10) can now be used to "update" the parameters on the basis of the imaginary
database. As we have not seen any cases before, $\nu$ and $\rho$ are estimated by the size of the imaginary database. Further

$$
\mu=m \quad \text { and } \quad \Phi=s s d=(\nu-1) \Sigma .
$$

In Geiger and Heckerman (1994), $\mu$ and $\Phi$ are found in a slightly different way. They use the fact that the marginal likelihood $p(y)$ is a multivariate non-central $t$ distribution with $\rho$ degrees of freedom, location vector $\mu$ and scale matrix $S=\frac{\nu+1}{\nu \rho} \Phi$. Now the mean and covariance matrix in the $t$ distribution is given by

$$
\mathbb{E}(y)=\mu \text { and } \operatorname{Cov}(y)=\frac{\rho}{\rho-2} S .
$$

They then let the mean and covariance matrix from the prior network estimate the mean and covariance matrix in the $t$ distribution, which implies that

$$
\mu=m \quad \text { and } \quad \Phi=\frac{\nu(\rho-2)}{\nu+1} \Sigma .
$$

Experimental results have not shown noticeable differences between the two approaches.

### 6.3 Properties of the Master Prior Procedure

The method for finding prior parameter distributions described in the previous section has some properties, which we will describe here. In this section we use $\Psi$ as a parameter defined for a joint distribution, i.e. $\Psi$ can be the parameter for the discrete variables or in the continuous case, $\Psi=(m, \Sigma)$.

Clearly a consequence of using the above method is that the parameters are independent. Further it can be seen, that if a node $v$ has the same parents in two DAGs $D$ and $D^{*}$, then

$$
p\left(\Psi_{v \mid \mathrm{pa}(v)} \mid D\right)=p\left(\Psi_{v \mid \operatorname{pa}(v)} \mid D^{*}\right)
$$

This property is referred to as parameter modularity. Now both the discrete and the Gaussian distribution has the property that if the joint probability distribution $p(x)$ can be factorized according to a DAG $D$, then it can also be factorized according to all other DAGs, which represents the same set of conditional independencies as $D$. A set of DAGs, $D^{e}$, which represents the same independence constraints is referred to as independence equivalent DAGs. So let $D$ and $D^{*}$ be independence equivalent DAGs, then

$$
p(x \mid \Psi, D)=p\left(x \mid \Psi, D^{*}\right)
$$

This means, that from observations alone we can not distinguish between different DAGs in an equivalence class. In the papers Heckerman et al. (1995) and Geiger and Heckerman (1994) it is for respectively the discrete and the Gaussian case shown, that when using the master prior procedure for the construction of parameter priors, the marginal likelihood for data is also the same for independence equivalent networks, i.e.

$$
p(d \mid D)=p\left(d \mid D^{*}\right)
$$

This equivalence is referred to as likelihood equivalence. Note that likelihood equivalence imply that if $D$ and $D^{*}$ are independence equivalent networks, then they have the same joint prior for the parameters, i.e.

$$
p(\Psi \mid D)=p\left(\Psi \mid D^{*}\right)
$$

## 7 Local Masters for Mixed Networks

In this section we will show how to specify prior distributions for the parameters in a CG network. In the mixed case, the marginal of a CG distribution is not always a CG distribution. In fact it is only a CG distribution if we marginalize over continuous variables or if we marginalize over a set $B$ of discrete variable, where $(B \Perp \Gamma) \mid(\Delta \backslash B)$, see Frydenberg (1990). Consider the following example. We have a network of two variables, $i$ and $y$, and the joint distribution is given by

$$
p(i, y)=p(i) \mathcal{N}\left(m_{i}, \sigma_{i}^{2}\right)
$$

Then the marginal distribution of $y$ is given as a mixture of normal distributions

$$
p(y)=\sum_{i \in \mathcal{I}} p(i) \mathcal{N}\left(m_{i}, \sigma_{i}^{2}\right)
$$

so there is no simple way of using this directly for finding the local priors.

### 7.1 The Suggested Solution

The suggested solution is very similar to the solution for the pure cases. We start by specifying a prior Bayesian network $(D, \mathcal{P})$ and calculate the joint probability distribution

$$
p(i, y \mid H)=p(i \mid \Psi) \mathcal{N}\left(m_{i}, \Sigma_{i}\right)
$$

with $H=\left(\Psi,\left(m_{i}\right)_{i \in \mathcal{I}},\left(\Sigma_{i}\right)_{i \in \mathcal{I}}\right)$. So from the conditional parameters in the local distributions in the prior network, we calculate the parameters for the joint distribution. Then we translate this prior network into an imaginary database, with imaginary sample size $n$. From the probabilities in the discrete part of the network, we can, as in the pure discrete case, calculate $\alpha_{i}$ for all configurations of $i$. Now $\alpha_{i}$ represents how many times we have observed $I=i$ in the imaginary database. We can assume that each time we have observed the discrete variables $I$, we have observed the continuous variables $Y$ and therefore set $\nu_{i}=\rho_{i}=\alpha_{i}$. Now for each configuration of $i$, we let $m_{i}$ be the sample mean in the imaginary database, and $\Sigma_{i}$ the sample variance. Further, as for the pure Gaussian case, we use $m_{i}=\mu_{i}$ and $\Phi_{i}=\left(\nu_{i}-1\right) \Sigma_{i}$. However, for $\Phi_{i}$ to be positive, $\nu_{i}$ has to larger than 1 , for all configurations $i$ and this has an impact on how small we can choose $n$ to be, as $n=\sum_{i} \nu_{i}$. If the number of discrete variables is large, and/or the number of configurations of the discrete variables is large, then we might have to let $n$ be larger than the value, that really reflects our confidence in the prior network. For these situations it might therefore be better to e.g. let $\Phi_{i}=\nu_{i} \Sigma_{i}$ as we then can choose the value of $n$ any way we want. Or, we can just choose $\nu_{i}$ and $\rho_{i}$ independently of $n$.

All the parameters needed to define the joint prior distributions for the parameters are now specified, so

$$
\begin{aligned}
p(\Psi) & =\mathcal{D}(\alpha), \\
p\left(M_{i} \mid \Sigma_{i}\right) & =\mathcal{N}\left(\mu_{i}, \frac{1}{\nu_{i}} \Sigma_{i}\right), \\
p\left(\Sigma_{i}\right) & =\mathcal{I} W\left(\rho_{i}, \Phi_{i}\right) .
\end{aligned}
$$

But we can not use these distributions to derive priors for other networks, so instead we use the imaginary database to derive local master distributions.

Let, for each family $A=v \cup \mathrm{pa}(v)$, the marginal CG distribution of $X_{a}$ given $H_{A}$ be given by

$$
\left(X_{A} \mid H_{A}\right) \sim C G\left(\Psi_{i_{A \cap \Delta}}, m_{A \cap \Gamma \mid i_{A \cap \Delta}}, \Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}}\right) .
$$

Then we suggest that the marginal prior distributions, also called the local masters, are found in the following way:

Let, for any variable $z, z_{i_{A \cap \Delta}}=\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} z_{j}$. Then

$$
\begin{aligned}
\left(\Psi_{A \cap \Delta}\right) & \sim \mathcal{D}\left(\alpha_{A \cap \Delta}\right) \\
\left(\Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}}\right) & \sim \mathcal{I} W\left(\rho_{i_{A \cap \Delta}},\left(\tilde{\Phi}_{A \cap \Gamma \mid i_{A \cap \Delta}}\right)\right. \\
\left(m_{A \cap \Gamma \mid i_{A \cap \Delta}} \mid \Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}}\right) & \sim \mathcal{N}\left(\bar{\mu}_{A \cap \Gamma \mid i_{A \cap \Delta}}, \frac{1}{\nu_{i_{A \cap \Delta}}} \Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}}\right)
\end{aligned}
$$

where

$$
\bar{\mu}_{i_{A \cap \Delta}}=\frac{\left(\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} \mu_{j} \nu_{j}\right)}{\nu_{A \cap \Delta}}
$$

and

$$
\tilde{\Phi}_{i_{A \cap \Delta}}=\Phi_{i_{A \cap \Delta}}+\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} \nu_{j}\left(\mu_{j}-\bar{\mu}_{i_{A \cap \Delta}}\right)\left(\mu_{j}-\bar{\mu}_{i_{A \cap \Delta}}\right)^{\mathrm{T}}
$$

The equations in the above result are well known from the analysis of variance theory, see e.g. Seber (1984). The marginal mean is found as a weighted average of the mean in every group, where a group here is given as a configuration of the discrete parents we marginalize over. The weights are the number of observations in each group. The marginal $s s d$ is given as the within group variation plus the between group variation. Notice that with this method, it is possible to specify mixed networks, where the mean in the mixed part of the network depends on the discrete parents, but the variance does not.

From the local masters we can now, by conditioning as in the pure cases, derive the local priors needed to specify the prior parameter distribution for a CG network. So the only difference between the master procedure and the local master procedure is in the way the marginal distributions are found.

### 7.2 Properties of the Local Master Procedure

The local master procedure coincides with the master procedure in the pure cases. Further, the properties of the local master procedure in the mixed case, are the same as of the master prior procedure in the pure cases.

Parameter independence and parameter modularity follows immediately from the definition of the procedure. To show likelihood equivalence, we need the following result from Chickering (1995). Let $D$ and $D^{*}$ be two DAGs and let $R_{D, D^{*}}$ be the set of edges by which $D$ and $D^{*}$ differ in directionality. Then, $D$ and $D^{*}$ are independence equivalent if and only if there exists a sequence of $\left|R_{D, D^{*}}\right|$ distinct arc reversals applied to $D$ with the following properties:

- After each reversal, the resulting network structure is a DAG, i.e. it contains no directed cycles and it is independence equivalent to $D^{*}$.
- After all reversals, the resulting DAG is identical to $D^{*}$.
- If $w \rightarrow v$ is the next arc to be reversed in the current DAG, then $w$ and $v$ have the same parents in both DAGs, with the exception that $w$ is also a parent of $v$ in $D$.

Note that as we only reverse $\left|R_{D, D^{*}}\right|$ distinct arcs, we only reverse arcs in $R_{D, D^{*}}$. For mixed networks this means that we only reverse arcs between discrete variables or between continuous variables, as the only arcs that can differ in directionality are these. So we can use the above result for mixed networks.

From the above we see that we can show likelihood equivalence by showing that $p(d \mid D)=p\left(d \mid D^{*}\right)$ for two independence equivalent DAGs $D$ and $D^{*}$ that differ only by the direction of a single arc. As $p(x \mid H, D)=p\left(x \mid H, D^{*}\right)$ in CG networks, we can show likelihood equivalence by showing that $p(H \mid D)=$ $p\left(H \mid D^{*}\right)$.

In the following let $v \rightarrow w$ in $D$ and $w \rightarrow v$ in $D^{*}$. Further let $\nabla$ be the set of common discrete and continuous parents for $v$ and $w$. Of course, if $v$ and $w$ are discrete variables, then $\nabla$ only contains discrete variables. The relation between $p(H \mid D)$ and $p\left(H \mid D^{*}\right)$ is given by:

$$
\begin{align*}
\frac{p(H \mid D)}{p\left(H \mid D^{*}\right)} & =\frac{p\left(H_{v \mid w \cup \nabla}, D\right) p\left(H_{w \mid \nabla}, D\right)}{p\left(H_{w \mid v \cup \nabla}, D^{*}\right) p\left(H_{v \mid \nabla}, D^{*}\right)} \\
& =\frac{p\left(H_{v \cup w \mid \nabla}, D\right)}{p\left(H_{v \cup w \mid \nabla}, D^{*}\right)} \tag{11}
\end{align*}
$$

When using the local master procedure, the terms in (11) are equal. This is evident, as we find the conditional priors from distributions over families $A$, in this case $A=v \cup w \cup \nabla$, which is the same for both networks. Therefore likelihood equivalence follows.

## 8 Model Search

In the search for Bayesian networks with high network score, we can, in theory, calculate the network score for all possible DAGs and then choose the DAG or DAGs with the highest score.

In Robinson (1977), a recursive formula for the number of possible DAGs that contains $n$ nodes, is found to be

$$
f(n)=\sum_{i=1}^{n}(-1)^{i+1}\binom{n}{i} 2^{i(n-i)} f(n-i)
$$

where $\binom{n}{i}$ are the binomial coefficient. As we in mixed networks do not allow discrete nodes to have continuous parents, the number of possible mixed DAGs is given by

$$
f(|\Delta|, \mid \Gamma) \mid)=f(|\Delta|) \times f(|\Gamma|) \times 2^{|\Delta| \times|\Gamma|}
$$

where $f(|\Delta|)$ and $f(|\Gamma|)$ are the numbers of DAGs for respectively the discrete and the continuous nodes, and $2^{|\Delta| \times|\Gamma|}$ denotes the number of different combinations of arrows from discrete to continuous nodes. If the number of random variables in the network is large, it is computationally infeasible to calculate the network score for all the possible DAGs. Therefore different methods for searching for DAGs with high network score have been tried, see e.g. Cooper and Herskovits (1992). In Section 8.3 we will describe one of these methods, namely greedy search with random restarts. This method, like many others, make use of Bayes factors as a way of comparing the network scores for two different DAGs. In the next section we will therefore consider Bayes factors for mixed networks.

### 8.1 Bayes Factors

A way to compare the network score for two different networks, $D$ and $D^{*}$, is to calculate the posterior odds, given by

$$
\frac{p(D \mid d)}{p\left(D^{*} \mid d\right)}=\frac{p(D, d)}{p\left(D^{*}, d\right)}=\frac{p(D)}{p\left(D^{*}\right)} \times \frac{p(d \mid D)}{p\left(d \mid D^{*}\right)}
$$

where $p(D) / p\left(D^{*}\right)$ is the prior odds and $p(d \mid D) / p\left(d \mid D^{*}\right)$ is the Bayes factor.
The posterior odds is for numerical reasons often calculated using the logarithm,

$$
\log \left(\frac{p(D \mid d)}{p\left(D^{*} \mid d\right)}\right)=\log (p(D \mid d))-\log \left(p\left(D^{*} \mid d\right)\right)
$$

For two models that differ only by a single arrow, the Bayes factor is, because of decomposability, especially simple. In this section, we will specify the Bayes
factor in the case where two DAGs differ by the direction of a single arrow and in the case where two DAGs differ by the presence of a single arrow.

First we look at the former case. As discrete nodes can not have continuous parents, we only look at reversing an arrow between two discrete variables or two continuous variables. In the following let $v \leftarrow w$ in $D$ and $v \rightarrow w$ in $D^{*}$. Further let $\nabla_{w}$ be the parents of $w$ in $D$ and $\nabla_{v}$ the parents of $v$ in $D^{*}$. As $D$ and $D^{*}$ only differ by the direction of the arrow between $v$ and $w$, the parents of $w$ in $D^{*}$ are $\nabla_{w}$ and $v$ and the parents of $v$ in $D$ are $\nabla_{v}$ and $w$. Notice that if $v$ and $w$ are discrete nodes, then the nodes in $\nabla_{v}$ and $\nabla_{w}$ can only be discrete, whereas if $v$ and $w$ are continuous nodes, they can be both discrete and continuous.

To simplify, we let the database consist of just one case, so $d=\{x\}$. As the likelihood terms are decomposable, the Bayes factor is given by

$$
\begin{aligned}
\frac{p(x \mid D)}{p\left(x \mid D^{*}\right)} & =\frac{p\left(v \mid \nabla_{v}, w, D\right) p\left(w \mid \nabla_{w}, D\right)}{p\left(w \mid \nabla_{w}, v, D^{*}\right) p\left(v \mid \nabla_{v}, D^{*}\right)} \\
& =\frac{\int p\left(x_{v} \mid x_{w \cup \nabla_{v}}, H_{v \mid w \cup \nabla_{v}}, D\right) p\left(H_{v \mid w \cup \nabla_{v}} \mid D\right) d H_{v \mid w \cup \nabla_{v}}}{\int p\left(x_{w} \mid x_{v \cup \nabla_{w}}, H_{w \mid v \cup \nabla_{w}}, D^{*}\right) p\left(H_{w \mid v \cup \nabla_{w}} \mid D^{*}\right) d H_{w \mid v \cup \nabla_{w}}} \\
& \times \frac{\int p\left(x_{w} \mid x_{\nabla_{w}}, H_{w \mid \nabla_{w}}, D\right) p\left(H_{w\left|\nabla_{w}\right|} \mid D\right) d H_{w \mid \nabla_{w}}}{\int p\left(x_{v} \mid x_{\nabla_{v}}, H_{v \mid \nabla_{v}}, D^{*}\right) p\left(H_{v \mid \nabla_{v}} \mid D^{*}\right) d H_{v \mid \nabla_{v}}} .
\end{aligned}
$$

So to calculate the Bayes factor between $D$ and $D^{*}$, we only need to consider the terms involving the conditional distributions of $v$ and of $w$.

Notice that if $\nabla_{v}=\nabla_{w}$, then $D$ and $D^{*}$ are independence equivalent networks and the Bayes factor is equal to one.

Now let $D$ and $D^{*}$ be two different networks, that differ by a single arrow between the nodes $v$ and $w$, with $v \leftarrow w$ in $D$ and $v \nleftarrow w$ in $D^{*}$. Here $v$ and $w$ can be either both discrete variables, both continuous or $v$ continuous and $w$ discrete. Again, let $\nabla_{v}$ be the set of variables that are parents of $v$ in $D^{*}$, so in $D$ the parents of $v$ are $\nabla_{v}$ and $w$. As the likelihood terms are decomposable, the Bayes factor is given by

$$
\begin{aligned}
\frac{p(x \mid D)}{p\left(x \mid D^{*}\right)} & =\frac{p\left(x_{v} \mid x_{w \cup \nabla_{v}}, D\right)}{p\left(x_{v} \mid x_{\nabla_{v}}, D^{*}\right)} \\
& =\frac{\int p\left(x_{v} \mid x_{w \cup \nabla_{v}}, H_{v \mid w \cup \nabla_{v}}, D\right) p\left(H_{v \mid w \cup \nabla_{v}} \mid D\right) d H_{v \mid w \cup \nabla_{v}}}{\int p\left(x_{v} \mid x_{\nabla_{D}}, H_{v \mid \nabla_{v}}, D^{*}\right) p\left(H_{v \mid \nabla_{v}} \mid D^{*}\right) d H_{v \mid \nabla_{v}}} .
\end{aligned}
$$

### 8.2 Equivalent Bayes Factors

To compare network scores for all networks which differ by only one arrow, is computationally inefficient. When using the local master procedure, we can reduce the number of comparisons needed.

Our goal is to identify classes of DAGs for which the corresponding Bayes factors for testing an arrow between the same two variables in the network, are the same. So let $D_{1}$ and $D_{1}^{*}$ be two different networks that differ by a single arrow between the nodes $v$ and $w$, with $v \leftarrow w$ in $D_{1}$ and $v \nleftarrow w$ in $D_{1}^{*}$. Further, let $\nabla_{v_{1}}$ be the set of variables that are parents of $v$ in both $D_{1}$ and $D_{1}^{*}$, i.e. in $D_{1}$ the parents of $v$ are $\nabla_{v_{1}}$ and $w$ and in $D_{1}^{*}$ just $\nabla_{D_{1}}$.

Further let $D_{2}$ and $D_{2}^{*}$ be another two networks different from $D_{1}$ and $D_{1}^{*}$ that differ by an arrow between $v$ and $w$ and let $\nabla_{v_{2}}$ be the set of variables that are parents of $v$ in both $D_{2}$ and $D_{2}^{*}$. There are two situations to consider, namely when $v \leftarrow w$ in $D_{2}$ and when $v \rightarrow w$ in $D_{2}$.

Consider first the former situation. The Bayes factor for testing $D_{1}$ against $D_{1}^{*}$ was in the previous section found to be

$$
\begin{equation*}
\frac{p\left(x \mid D_{1}\right)}{p\left(x \mid D_{1}^{*}\right)}=\frac{\int p\left(x_{v} \mid x_{w \cup \nabla_{v_{1}}}, H_{v \mid w \cup \nabla_{v_{1}}}, D_{1}\right) p\left(H_{v \mid w \cup \nabla_{v_{1}}} \mid D_{1}\right) d H_{v \mid w \cup \nabla_{v_{1}}}}{\int p\left(x_{v} \mid x_{\nabla_{v_{1}}}, H_{v \mid \nabla_{v_{1}}}, D_{1}^{*}\right) p\left(H_{v \mid \nabla_{v_{1}}} \mid D_{1}^{*}\right) d H_{v \mid \nabla_{v_{1}}}} . \tag{12}
\end{equation*}
$$

Likewise the Bayes factor for testing $D_{2}$ against $D_{2}^{*}$ is

$$
\frac{p\left(x \mid D_{2}\right)}{p\left(x \mid D_{2}^{*}\right)}=\frac{\int p\left(x_{v} \mid x_{w \cup \nabla_{v_{2}}}, H_{v \mid w \cup \nabla_{v_{2}}}, D_{2}\right) p\left(H_{v \mid w \cup \nabla_{v_{2}}} \mid D_{2}\right) d H_{v \mid w \cup \nabla_{v_{2}}}}{\int p\left(x_{v} \mid x_{\nabla_{v_{2}}}, H_{v \mid \nabla_{v_{2}}}, D_{2}^{*}\right) p\left(H_{v \mid \nabla_{v_{2}}} \mid D_{2}^{*}\right) d H_{v \mid \nabla_{v_{2}}}} .
$$

As the local master procedure has the property of parameter modularity, then if $\nabla_{v_{1}}=\nabla_{v_{2}}$ it follows that

$$
p\left(H_{v \mid w \cup \nabla_{v_{1}}} \mid D_{1}\right)=p\left(H_{v \mid w \cup \nabla_{v_{2}}} \mid D_{2}\right)
$$

and

$$
p\left(x_{v} \mid x_{w \cup \nabla_{v_{1}}}, H_{v \mid w \cup \nabla_{v_{1}}}, D_{1}\right)=p\left(x_{v} \mid x_{w \cup \nabla_{v_{2}}}, H_{v \mid w \cup \nabla_{v_{2}}}, D_{2}\right) .
$$

So the Bayes factor for testing the arrow from $v$ to $w$ is equivalent to testing this arrow in any other network, where $v$ has the same parents as in $D_{1}$, i.e. if $\nabla_{v_{1}}=\nabla_{v_{2}}$. This is illustrated in Figure 1.


Figure 1: Equivalence due to parameter modularity.


Figure 2: Equivalence due to property of local master procedure.

Consider now the situation where $v \rightarrow w$ in $D_{2}$. Let $\nabla_{w_{2}}$ be the set of variables, that are parents of $w$ in both $D_{2}$ and $D_{2}^{*}$. The Bayes factor is given as

$$
\begin{aligned}
\frac{p\left(x \mid D_{2}\right)}{p\left(x \mid D_{2}^{*}\right)} & =\frac{p\left(x_{w} \mid x_{v \cup \nabla_{w_{2}}}, D_{2}\right)}{p\left(x_{w} \mid x_{\nabla_{w_{2}}}, D_{2}^{*}\right)} \\
& =\frac{\int p\left(x_{w} \mid x_{v \cup \nabla_{w_{2}}}, H_{w \mid v \cup \nabla_{w_{2}}}, D_{2}\right) p\left(H_{w \mid v \cup \nabla_{w_{2}}} \mid D_{2}\right) d H_{w \mid v \cup \nabla_{w_{2}}}}{\int p\left(x_{w} \mid x_{\nabla_{w_{2}}}, H_{w \mid \nabla_{w_{2}}}, D_{2}^{*}\right) p\left(H_{w \mid \nabla_{w_{2}}} \mid D_{2}^{*}\right) d H_{w \mid \nabla_{w_{2}}}} .
\end{aligned}
$$

Again we see that because of parameter modularity, this Bayes factor is the same as the Bayes factor given in (12), if $\nabla_{v_{1}}=\nabla_{w_{2}}$, i.e. if $w$ in $D_{2}$ has the same parents as $v$ does in $D_{1}$, with the exception that $v$ is a parent of $w$ in $D_{2}$. For an illustration, see Figure 2.

To show that these situations are the only ones where the Bayes factors always are the same, it is easy to find an example where $\nabla_{v_{1}} \neq \nabla_{v_{2}}$ and the Bayes factors are not same.

The above result is summarized in the following theorem.

## Theorem 8.1

The Bayes factor for testing the arrow $v \leftarrow w$ in a $D A G D_{1}$ is equivalent to the Bayes factor for testing the same arrow in any other network $D_{2}$ if and only if the following two criteria are met:
(1) $v \leftarrow w$ and $v$ in $D_{2}$ has the same parents as in $D_{1}$.
(2) $v \rightarrow w$ and $w$ in $D_{2}$ has the same parents as $v$ does in $D_{1}$, with the exception that $v$ is a parent of $w$ in $D_{2}$.

Although using the two criteria reduces the number of comparisons, there will still, for large networks, be too many comparisons needed for finding the most likely DAG. Therefore it is still necessary to use some kind of search strategy.

### 8.3 Greedy search with random restarts

As mentioned earlier, many search strategies use Bayes factors as a way to compare the network score for two different networks. In the following we will describe one such strategy called greedy search.

Greedy search is initialized by choosing a network $D$ from which to start the search. Let $\Delta e$ be the posterior odds between two networks that differ by an arrow. Calculate then $\Delta e$ for all DAGs $D^{*}$ that differ from $D$ by a single arrow $e$, either added, removed or reversed. Make the change $e$ for which $\Delta e$ is a minimum, that is where $p\left(D^{*} \mid d\right)$ is a maximum and continue the search from this new network. The search is terminated when there is no $e$ with $\Delta e$ smaller than 1 . As shown in the previous section, the posterior odds is because of decomposability especially simple, as $D$ and $D^{*}$ only differ by one arrow. Further, it is possible to reduce the time complexity by using the equivalence criteria developed in Section 8.2.

As this search is local in the sense that it only evaluates local changes to the network, there is a chance that the found maximum is only a local maximum. A way to overcome this problem is to randomly perturb the structure of the start network $D$ and restart the greedy search from this new network. This can be repeated a manageable number of times and between the networks found by the search strategy, the network with the highest score is chosen.

### 8.4 Priors on DAGs

In this section we will consider how to assign prior probabilities to the possible DAGs in a given problem. As shown in various papers, there are different ways of doing this. The Bayesian way would be to assess the prior belief in each DAG, but as the number of different DAGs grow, this is not manageable. Instead automated methods is being used.


Figure 3: Models for which the Bayes factors are equivalent.

An often used approach is to assume that all DAGs are equally likely, thus letting the prior probability distribution over DAGs be uniform. This approach is mostly used only for simplicity and can be refined in various ways. For example, if we know that some of the DAGs are not possible, then we can assign probability zero to these and equal probabilities to the rest. Because of likelihood equivalence, DAGs within the same equivalence class will, with this approach, be assigned the same network score.

One argument against letting the prior over DAGs be uniform is that the number of different DAGs in an equivalence class varies between equivalence classes. This means that the conditional independencies represented in an equivalence class with many DAGs, a priori are more probable than those represented in an equivalence class with fewer DAGs. When using model averaging, this is a problem because it involves a sum over all the different DAGs. The conditional independencies represented by a large equivalence class, therefore influence the result more than those represented by a small equivalence class. A way to handle this problem is to either include only one DAG from each equivalence class or instead let all equivalence classes be equally likely and assign to each DAG a prior probability inversely proportional to the number of DAGs in the equivalence class it belongs to.

This last approach has, however, an affect on the posterior odds. Consider the following example, illustrated in Figure 3.

According to criteria one in Theorem 8.1, the Bayes factor for testing the presence of the arrow $v \leftarrow w$ in $D_{1}$ is equivalent to testing $v \leftarrow w$ in $D_{2}$, i.e.

$$
\frac{p\left(v \mid w, D_{1}\right)}{p\left(v \mid D_{1}^{*}\right)}=\frac{p\left(v \mid w, D_{2}\right)}{p\left(v \mid D_{2}^{*}\right)}
$$

If we assign equal priors to all DAGs, the posterior odds are the same as the Bayes factors and they will therefore also be equivalent in the above example. However, if we let all equivalence classes be equally likely and assign to each DAG a prior probability inversely proportional to the number of DAGs in the
equivalence class it belongs to, the posterior odds are no longer the same as the Bayes factors. In the above example, the number of DAGs in the equivalence classes for $D_{1}, D_{1}^{*}, D_{2}$ and $D_{2}^{*}$ are respectively $3,2,2$ and 1 . So the prior odds are not equivalent, i.e.

$$
\frac{p\left(D_{1}\right)}{p\left(D_{1}^{*}\right)}=\frac{2}{3} \neq \frac{1}{2}=\frac{p\left(D_{2}\right)}{p\left(D_{2}^{*}\right)}
$$

and therefore the posterior odds are not equivalent either. So this approach should not be used if we in a search strategy want to utilize that some of the Bayes factors are equivalent.

## 9 Example

In the following, some of the methods derived are illustrated by a simple example. This example was constructed by Morrison (1976) and also studied in Edwards (1995).

### 9.1 The Dataset

The dataset is from a hypothetical drug trial, where the weight losses of male and female rats under three different drug treatments have been measured after one and two weeks. Thus we have the discrete variables $I_{\text {sex }}$ and $I_{d r u g}$ with states

$$
\begin{aligned}
& I_{\text {sex }}=\{\text { male }=1, \text { female }=2\} \\
& I_{d r u g}=\{1,2,3\},
\end{aligned}
$$

and the continuous variables $Y_{w 1}$ and $Y_{w 2}$ which respectively represents the weight losses after one and two weeks. For every drug, four rats of each sex have been treated, which gives a total of 24 observations. The observations are shown in Table 1.

### 9.2 Specifying the Prior Network

We start by specifying a prior Bayesian network $(D, \mathcal{P})$. To simplify the specification of the joint parameter prior, we choose to let all the variables be inde-

| sex | drug | $w 1$ | $w 2$ | sex | $d r u g$ | $w 1$ | $w 2$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 1 | 5 | 6 | 2 | 1 | 7 | 10 |
| 1 | 1 | 7 | 6 | 2 | 1 | 8 | 10 |
| 1 | 1 | 9 | 9 | 2 | 1 | 6 | 6 |
| 1 | 1 | 5 | 4 | 2 | 1 | 9 | 7 |
| 1 | 2 | 9 | 12 | 2 | 2 | 7 | 6 |
| 1 | 2 | 7 | 7 | 2 | 2 | 10 | 13 |
| 1 | 2 | 7 | 6 | 2 | 2 | 6 | 9 |
| 1 | 2 | 6 | 8 | 2 | 2 | 8 | 7 |
| 1 | 3 | 14 | 11 | 2 | 3 | 14 | 9 |
| 1 | 3 | 21 | 15 | 2 | 3 | 14 | 8 |
| 1 | 3 | 12 | 10 | 2 | 3 | 16 | 12 |
| 1 | 3 | 17 | 12 | 2 | 3 | 10 | 5 |

Table 1: Observations of weight loss of male and female rats under three different drug treatments.
pendent, so the local probability distribution for each node only depends on the node itself, and we can specify them as follows.

For each discrete variable, we let each state be equally likely, so

$$
p\left(i_{\operatorname{sex}}=1\right)=p\left(i_{\operatorname{sex}}=2\right)=\frac{1}{2}
$$

and

$$
p\left(i_{d r u g}=1\right)=p\left(i_{d r u g}=2\right)=p\left(i_{d r u g}=3\right)=\frac{1}{3}
$$

This in fact is true by design.
For the continuous variables we use the sample mean and the sample variance as an initial estimate of the mean and the variance. Using this approach, the position and scale of the parameters are determined. We find that

$$
p\left(y_{w 1}\right)=\mathcal{N}(9.6,17.1)
$$

and

$$
p\left(y_{w 2}\right)=\mathcal{N}(8.7,7.6)
$$

So jointly

$$
p(i, y)=p(i) \mathcal{N}\left(m_{i}, \Sigma_{i}\right)
$$

with

$$
p(i)=\frac{1}{6}, \quad m_{i}=\binom{9.6}{8.7} \quad \text { and } \quad \Sigma_{i}=\left(\begin{array}{cc}
17.1 & 0 \\
0 & 7.6
\end{array}\right)
$$

for all possible configurations of $i$.
Be aware that in this way the dataset is used twice, namely both to initially specify the local probability distributions and later to find the posterior parameter distributions. This could result in parameter values that are overfitted to data.

### 9.3 Specifying Parameter Priors

In order to specify parameter priors for all possible networks, we use the local master procedure.

First we translate the prior network into an imaginary database. The parameters needed to represent this imaginary database are $n, \alpha_{i}, \nu_{i}, \rho_{i}, \mu_{i}$ and $\Phi_{i}$.

Here we let $\Phi_{i}=\left(\nu_{i}-1\right) \Sigma_{i}$, so $\nu_{i}$ must be larger than 1 . This means in this example that $n$ must be larger than 6 . We choose $n=12$ and find that

$$
\alpha_{i}=\nu_{i}=\rho_{i}=p(i) n=\frac{1}{6} 12=2 .
$$

Further

$$
\mu_{i}=m_{i}=\binom{9.6}{8.7} \quad \text { and } \quad \Phi_{i}=\left(\nu_{i}-1\right) \Sigma_{i}=\left(\begin{array}{cc}
17.0 & 0 \\
0 & 7.6
\end{array}\right)
$$

for all configurations of $i$.
We can now specify parameter priors for all possible networks. As an illustration, consider the parameter prior for the network in Figure 4.

We need to find the local masters for the following four families

$$
\begin{aligned}
& A_{1}=\{\text { sex }\} \\
& A_{2}=\{d r u g\} \\
& A_{3}=\{w 1\} \\
& A_{4}=\{s e x, w 1, w 2\} .
\end{aligned}
$$



Figure 4: The DAG in the example for specification of local parameter priors.

As the variables in $A_{1}, A_{2}$ and $A_{3}$ do not have any parents, the local masters for these families are also the local parameter priors. Thus the local parameter prior for $I_{\text {sex }}$ is given by

$$
\Psi_{s e x} \sim \mathcal{D}\left(\alpha_{s e x}\right)
$$

with

$$
\alpha_{i_{s e x}=1}=\sum_{j: j_{\text {sex }}=1} \alpha_{j}=6 \text { and } \alpha_{i_{s e x}=2}=\sum_{j: j_{s e x}=2} \alpha_{j}=6 .
$$

Similarly the local parameter prior for $I_{d r u g}$ is

$$
\Psi_{d r u g} \sim \mathcal{D}\left(\alpha_{d r u g}\right)
$$

with

$$
\alpha_{i_{d r u g}=1}=\alpha_{i_{d r u g}}=2=\alpha_{i_{d r u g}=3}=4
$$

For $Y_{w 1}$ we find the local parameter prior to be

$$
\begin{aligned}
\Sigma_{w 1} & \sim \mathcal{I} W\left(\rho, \tilde{\Phi}_{w 1}\right) \\
m_{w 1} \mid \Sigma_{w 1} & \sim \mathcal{N}\left(\bar{\mu}_{w 1}, \frac{1}{\nu} \Sigma_{w 1}\right)
\end{aligned}
$$

with

$$
\rho=\sum_{j} \rho_{j}=12 \text { and } \nu=\sum_{j} \nu_{j}=12
$$

and

$$
\begin{aligned}
\bar{\mu} & =\frac{\sum_{i} \mu_{i} \nu_{i}}{\nu}=\binom{9.6}{8.7}, \\
\tilde{\Phi} & =\sum_{i} \Phi_{i}+\sum_{i} \nu_{i}\left(\mu_{i}-\bar{\mu}\right)\left(\mu_{i}-\bar{\mu}\right)^{\mathrm{T}}=\left(\begin{array}{cc}
102.6 & 0 \\
0 & 45.6
\end{array}\right),
\end{aligned}
$$

SO

$$
\bar{\mu}_{w 1}=9.6 \text { and } \tilde{\Phi}_{w 1}=102.6 .
$$

The local master for the family $A_{4}$ is given as

$$
\begin{aligned}
\left(\sum_{i_{\text {sex }}}\right) & \sim \mathcal{I} W\left(\rho_{i_{\text {sex }}},\left(\tilde{\Phi}_{i_{\text {sex }}}\right)\right) \\
\left(m_{i_{\text {sex }}}\right) \mid\left(\sum_{i_{\text {sex }}}\right) & \sim \mathcal{N}\left(\left(\bar{\mu}_{i_{\text {sex }}}\right), \frac{1}{\nu_{i_{\text {sex }}}}\left(\Sigma_{i_{\text {sex }}}\right)\right)
\end{aligned}
$$

with

$$
\rho_{i_{\text {sex }}=1}=\sum_{j: j_{s e x}=1} \rho_{j}=6 \text { and } \rho_{i_{s e x}=2}=\sum_{j: j_{s e x}=2} \rho_{j}=6 .
$$

Likewise for $\nu_{i_{s e x}}$. Further

$$
\bar{\mu}_{i_{s e x}=1}=\frac{\sum_{j: j_{s e x}=1} \mu_{j} \nu_{j}}{\nu_{i_{s e x}=1}}=\binom{9.6}{8.7}
$$

and

$$
\begin{aligned}
\tilde{\Phi}_{i_{s e x}=1} & =\sum_{j: j_{\text {sex }}=1} \Phi_{j}+\sum_{j: j_{\text {sex }}=1} \nu_{j}\left(\mu_{j}-\bar{\mu}_{i_{\text {sex }}=1}\right)\left(\mu_{j}-\bar{\mu}_{i_{s e x}=1}\right)^{\mathrm{T}} \\
& =\left(\begin{array}{cc}
51.3 & 0 \\
0 & 22.8
\end{array}\right)
\end{aligned}
$$

and the same for $i_{\text {sex }}=2$.
The local parameter prior for $Y_{w 2}$ given $Y_{w 1}$ and $I_{s e x}$ can now be found by conditioning in this local master distribution.

We have now specified the parameters needed to calculate the likelihood of a DAG, $p(d \mid D)$. To calculate the network score of $D$, we also need to specify the prior probability of $D$. In this example we just choose to let all DAGs be equally likely and thus use the likelihood $p(d \mid D)$ as the network score.

### 9.4 Result

Using the formula on page 35, we find that for a network with two discrete and two continuous nodes, there are 144 possible DAGs. So in this example, there are no computational problems in calculating the network score for all these DAGs. Further, if we only calculate the score for DAGs that are not independence equivalent, the number of different DAGs are reduced to 88 .

| Prior network |  | $\begin{array}{ll} 0 & 0 \\ 0 & 0 \end{array}$ | Imaginary sample size |  | 12 |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | $\underset{\substack{9 \\ 0.30}}{\substack{0 \\ 0}}$ | $\begin{gathered} \underset{\sim}{\circ \rightarrow 0} \\ 0.20 \end{gathered}$ |  |  |
|  |  |  |  |  |  |
| $\begin{aligned} & \substack{\rightarrow \infty \\ 0 \\ 0.023 \\ 0} \end{aligned}$ |  |  |  |  |  |
|  |  | $\begin{aligned} & \text { on } \\ & 0.0069 \end{aligned}$ |  |  |  |
|  | $\xrightarrow[\rightarrow 0]{\substack{0}}$ | $\begin{aligned} & 0 \quad 0 \\ & \mathrm{O} \longrightarrow \mathrm{O} \\ & 0.0019 \end{aligned}$ |  |  | $\begin{aligned} & \stackrel{\longrightarrow}{\longrightarrow} \\ & \underset{9.6 \cdot 10^{-4}}{\longrightarrow} \end{aligned}$ |
|  |  | $5.2 \cdot 10^{-4}$ |  | $\begin{aligned} & \stackrel{\leftrightarrow}{\circ} \mathrm{O} \\ & 0.9 \cdot 10^{-4} \end{aligned}$ |  |
|  |  |  |  |  |  |
|  |  |  | $\begin{aligned} & \circ \longrightarrow 0 \\ & \circ \\ & 0 \quad \text { o } \\ & 8.0 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & \leftrightarrow 0 \\ & \rightarrow 0 \\ & 7.2 \cdot 10^{-5} \end{aligned}$ |  |
| $5.2 \cdot 10^{-5}$ |  |  |  |  | $\begin{aligned} & \text { on } \\ & 4.2 \cdot 10^{-5} \end{aligned}$ |
| continued on next page |  |  |  |  |  |


| continued from previous page |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Prior network |  | $\begin{array}{ll} 0 & 0 \\ 0 & 0 \end{array}$ | Imaginary sample size |  | 12 |
| $\begin{aligned} & \text { of of } \\ & \text { of } \\ & 3.9 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & 0 \text { on } \\ & 0 \text { o } \\ & 3.8 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & 0 \\ & 0.4 \cdot 10^{-5} \\ & 3.0 \end{aligned}$ |  |  |  |
|  | $\begin{aligned} & 0 \text { P } \\ & 0 \text { o } \\ & 2.0 \cdot 10^{-5} \end{aligned}$ | $1.6 \cdot 10^{-5}$ |  |  |  |
| $\begin{aligned} & 98 \\ & 0 \text { o } \\ & 5.9 \cdot 10^{-6} \end{aligned}$ | $\begin{aligned} & 818 \\ & 4.9 \cdot 10^{-6} \end{aligned}$ | $\begin{aligned} & 0.0 \text { on } \\ & 0.0 \\ & 3.6 \cdot 10^{-6} \end{aligned}$ |  | $\begin{aligned} & \text { os } \\ & \text { of } \\ & 1.1 \cdot 10^{-6} \end{aligned}$ |  |
| $\begin{aligned} & \circ \text { م } 0 \\ & 0 \text { o } \\ & 3.2 \cdot 10^{-7} \end{aligned}$ | $\begin{array}{ll} \hline \circ & \circ \\ \circ & \circ \\ 3.0 \cdot 10^{-7} \\ \hline \end{array}$ |  | $\begin{aligned} & 90 \\ & 0 \quad 0 \\ & 2.5 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \hline \bigcirc \bigcirc \\ & \circ \circ \\ & 1.5 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \text {-o } \\ & 1.3 \cdot 10^{-7} \end{aligned}$ |
| $\begin{aligned} & 0 \text { o } \\ & 0 \text { o } \\ & 9.4 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \text { if } \circ \\ & \text { o } \circ \\ & 8.9 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \text { of } \\ & 8.8 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \text { oo } \\ & 0 \text { o } \\ & 7.4 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \rightarrow 0 \\ & 7.2 \cdot 10^{-8} \end{aligned}$ |  |
|  | $\begin{aligned} & \xrightarrow[0]{\circ} \\ & 0.8 \\ & 3.8 \cdot 10^{-8} \end{aligned}$ |  |  |  |  |

Table 2: The DAGs in the reduced search space, listed in decreasing order of probability. The number below each DAG is the Bayes factor between the given DAG and the DAG with the highest network score.

In Table 2 the result of the learning procedure is given. The DAGs are listed in decreasing order of probability, and the number below each DAG is the posterior odds between the given DAG and the DAG with the highest network score. This number expresses the relative probability of a DAG, that is, relative to the DAG with the highest network score. As we have chosen a uniform prior over DAGs, the posterior odds is in this example equal to the Bayes factor.

Before analyzing the result, we can discard some of the networks in Table 2. By design, the discrete variables sex and drug are independent, so there should not be an arrow between sex and drug. Further, there is a time restriction between $w 1$ and $w 2$, as $w 1$ is observed before $w 2$. So if $w 1$ and $w 2$ are dependent, the arrow between $w 1$ and $w 2$ must go from $w 1$ to $w 2$. Taking these restrictions into account, we only consider the 32 different DAGs listed in Table 3.

In the most probable DAG, we see that $w 2$ depends on $w 1$ and $w 1$ depends on

| Prior network |  | $\begin{array}{ll} 0 & 0 \\ 0 & 0 \end{array}$ | Imaginary sample size |  | 12 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{lr} 0 & 0 \\ 0 & 0 \\ 1 \end{array}$ | $\begin{gathered} \underset{\sim}{\sim \rightarrow 0} \\ 0 \rightarrow 0 \\ 0.68 \end{gathered}$ | $\begin{aligned} & 0 \underset{\sim}{0} \underset{0}{0} 0 \\ & 0.12 \end{aligned}$ | $\begin{aligned} & 9 \\ & 0 \rightarrow 0 \\ & 0.075 \end{aligned}$ | $\begin{aligned} & \rightarrow 0 \rightarrow 0 \\ & 0.051 \\ & 0.0 \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \rightarrow 0 \\ & 0.023 \end{aligned}$ |
| $\begin{gathered} 0_{0}^{\text {P }} \\ 0.0093 \\ 0 . \end{gathered}$ | $\begin{aligned} & 0 \underset{0}{0} \\ & 0.0020 \\ & 0.0 \end{aligned}$ | $\begin{aligned} & \circ \quad \circ \\ & \circ \rightarrow 0 \\ & 0.0019 \end{aligned}$ | $\begin{aligned} & \substack{0 \\ \rightarrow \infty \\ 0.0017} \end{aligned}$ | $\begin{aligned} & \stackrel{\longrightarrow}{\longrightarrow} \\ & \underset{9.6 \cdot 10^{-4}}{ } \end{aligned}$ | $\begin{aligned} & \stackrel{\infty}{\infty \rightarrow \infty} \\ & 4.5 \cdot 10^{-4} \end{aligned}$ |
| $\begin{array}{ll} \circ & \AA \\ 0 & 0 \\ 1.6 \cdot 10^{-4} \end{array}$ | $\underbrace{\substack{0}}_{1.5 \cdot 10^{-4}}$ | $\begin{aligned} & 0.0 \\ & 0 \rightarrow 0 \\ & 1.4 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & 1.3 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & \circ \longrightarrow \\ & \circ \\ & 0 \\ & 1.1 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & \infty \rightarrow \infty \\ & 8.9 \cdot 10^{-5} \end{aligned}$ |
|  | $\begin{aligned} & \rightarrow 0 \\ & 0.4 \cdot 10^{-5} \\ & 0 \rightarrow 0 \end{aligned}$ |  | $1.6 \cdot 10^{-5}$ | $3.6 \cdot 10^{-6}$ |  |
| $\begin{aligned} & 0 \text { o } \\ & 0 \text { o } \\ & 3.2 \cdot 10^{-7} \end{aligned}$ | $\begin{array}{ll} \circ & \circ \\ \circ & \circ \\ 3.0 \cdot 10^{-7} \end{array}$ | $\begin{aligned} & 9 \\ & 00_{0}^{9.6} 10^{-7} \\ & 2.0 \end{aligned}$ | $\begin{aligned} & 0 \\ & 2.5 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \circ \rightarrow 0 \\ & \circ \circ \\ & 1.5 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \stackrel{\infty}{\infty} \\ & 1.3 \cdot 10^{-7} \end{aligned}$ |
| $\begin{aligned} & \text { ole } \\ & 0 \\ & 7.2 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \underbrace{0}_{0} \\ & 5.9 \cdot 10^{-8} \end{aligned}$ |  |  |  |  |

Table 3: The DAGs in the reduced search space, listed in decreasing order of probability. The number below each DAG is the Bayes factor between the given DAG and the DAG with the highest network score.
$d r u g$. Further $w 2$ and $d r u g$ are conditionally independent given $w 1$ and both $w 1$ and $w 2$ are independent on sex.

Almost the same dependency structure is seen in the second and third best DAG, except that here $w 2$ also depends on respectively sex and $d r u g$.

Generally we see that in the first 12 DAGs, $w 1$ depends on $d r u g$. The first DAG that does not show this dependency relation is only 0.00016 times as probable as the best DAG. Likewise we see that in the first 7 DAGs, $w 2$ depends on $w 1$ and the first DAG that does not contain this dependency relation is only 0.0020 as probable as the best DAG. Therefore we should not consider any model that does not include these dependencies.

It is not clear which independencies should be included in the model, except for those introduced when we reduced the search space. The second DAG is
for example 0.68 times as probable as the first DAG, and the third to the sixth DAG is between 0.12 and 0.023 as probable as the best DAG. This suggest that there is some unexplained variation not accounted for in the best DAG and it might therefore be more accurate to select e.g. the first six models and use model averaging.

In Edwards (1995) the dataset is analyzed using undirected graphical models. He uses the software MIM for maximum likelihood estimation and likelihood ratio test. The result is displayed in Figure 5 and we see that it is not in conflict with our result.


Figure 5: Previous result.

### 9.5 Sensitivity to Prior Information

In this section we will explore how the size of the imaginary database and the choice of the prior network influences the result. The findings agree with findings for a purely discrete case described in Steck and Jaakkola (2002).

Recall that the prior network ideally expresses which dependency structure we believe there is between the variables in the network and the size of the imaginary database expresses how much confidence we have in this dependency structure.

In the previous section we used the empty network as the prior network and set the size $n$ of the imaginary database to 12 . This is less than the number of real observations in the example, which is 24 . We will therefore also learn the networks using a larger value of $n$ and to see the difference clearly, we use $n=2000$. The result is given in Table 4.

If we look at the three best networks from the previous result, we see that the relative probabilities for these networks in this result, are between 0.94 and 0.97 . They are no longer the most probable networks, but they are still very probable. Actually all the networks are very probable and the relative probability of the least probable network is as much as 0.78 .

| Prior network |  | $\begin{array}{ll} 0 & 0 \\ 0 & 0 \end{array}$ | Imaginary sample size |  | 2000 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \rightarrow 0 \\ & 1 \end{aligned}$ | $\begin{aligned} & \rightarrow \infty \\ & 0 \rightarrow 0 \\ & 0.99 \end{aligned}$ | $\begin{aligned} & 0 \text { P } \\ & 0 \rightarrow 0 \\ & 0.97 \end{aligned}$ | $\begin{aligned} & \substack{9 \\ 0 \rightarrow 0 \\ 0.96} \end{aligned}$ | $\begin{aligned} & \underset{\sim}{0 \rightarrow 0} \\ & 0 \rightarrow 0 \\ & 0.96 \end{aligned}$ | $\begin{aligned} & \rightarrow \infty \\ & 0 \rightarrow 0 \\ & 0.95 \end{aligned}$ |
| $\begin{aligned} & 0 \\ & 0 \rightarrow 0 \\ & 0.94 \end{aligned}$ | $\begin{aligned} & 0 \\ & 0 \\ & 0.93 \\ & 0.9 \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 1 \\ & 0 \\ & 0.89 \\ & 0.89 \end{aligned}$ | $\begin{aligned} & 0 \\ & 00 \\ & 0.89 \\ & 0.0 \end{aligned}$ | $\begin{aligned} & \underset{\rightarrow 0}{0 \rightarrow 0} \\ & 0.89 \\ & 0.89 \end{aligned}$ | $\xrightarrow[0]{\circ}$ |
| $\begin{array}{lr} \circ \quad 0 \\ 0 \rightarrow 0 \\ 0.88 \end{array}$ | $\begin{aligned} & \underset{=0}{\rightarrow 0} \\ & 0.88 \end{aligned}$ | $\begin{aligned} & 9 \\ & 0 \\ & 0.88 \end{aligned}$ | $\begin{aligned} & \stackrel{\longrightarrow}{\circ} \\ & \underset{0}{\longrightarrow} \\ & 0.87 \end{aligned}$ | $\begin{aligned} & 00 \\ & 0.0 \\ & 0.87 \end{aligned}$ | $\begin{aligned} & \stackrel{\rightarrow 0}{\rightarrow 0} \\ & 0.87 \end{aligned}$ |
| $\begin{aligned} & 0 \\ & 0 \\ & 0 \\ & 0.86 \\ & 0.86 \end{aligned}$ | $\begin{aligned} & 88 \\ & 0.86 \\ & 0.8 \end{aligned}$ | $\begin{aligned} & \circ \rightarrow 1 \\ & \circ \\ & 0.0 \\ & 0.85 \end{aligned}$ | $\begin{aligned} & \infty \rightarrow 0 \\ & 0 \\ & 0.85 \end{aligned}$ | $\begin{array}{ll} \circ & \circ \\ \circ & 0 \\ 0.84 \end{array}$ | $\begin{aligned} & 9 \\ & 0 \\ & 0.83 \\ & 0.83 \end{aligned}$ |
| $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \\ & 0.79 \\ & 0.79 \end{aligned}$ | $\begin{aligned} & 0 \\ & 0 \\ & 0 \\ & 0.79 \end{aligned}$ | oos | $\begin{aligned} & 0 \\ & 0.79 \\ & 0.79 \end{aligned}$ | $\begin{array}{lc} \circ & \circ \\ \circ & \circ \\ 0.78 \end{array}$ | $\begin{aligned} & 0 \\ & 0 \\ & 0 \\ & 0.78 \end{aligned}$ |
| $\begin{aligned} & \circ \longrightarrow \\ & \circ \\ & 0 \\ & 0.78 \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \\ & 0.78 \end{aligned}$ |  |  |  |  |

Table 4: The revised result with the prior network and the imaginary sample size specified as in the first line of this table.

The reason for this is that the prior network is the empty network, which represents that all the variables are independent. This model is therefore a submodel of all other models. When $n$ is large, we have much confidence in these independencies, so all networks will a priori be very probable. As the real database only contains few observations, we have not enough information to differentiate between these networks and all the networks are therefore almost equally likely.

We will now explore what happens if we change the prior network. First we will learn the structure using the most probable structure from Table 3 as the prior network. The results with $n=12$ and $n=2000$ are given in respectively Table 5 and Table 6.

For $n=12$ we see almost the same result as when using the empty network. The best networks are, not surprisingly, the same, only the order between them are a little different. To some extent, this also applies for $n=2000$.

| Prior network |  | $\begin{array}{cc} 0 & 0 \\ 0 & 0 \end{array}$ | Imaginary sample size |  | 12 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{aligned} & 0 \\ & 0 \rightarrow 0 \\ & 1 \end{aligned}$ |  |  | $\begin{array}{lr} 0 & 0 \\ 0 \rightarrow 0 \\ 0.34 \end{array}$ | $\underbrace{\infty}_{0 \rightarrow 0}$ $0.17$ | $\begin{aligned} & \rightarrow \infty \\ & \rightarrow \rightarrow+\infty \\ & 0.10 \end{aligned}$ |
| $\begin{gathered} 0,0 \\ 0 \\ 0.064 \\ 0 \end{gathered}$ | $\begin{aligned} & 0 \\ & 0 \\ & 0.056 \end{aligned}$ | $\begin{aligned} & \circ \stackrel{\circ}{\circ} \mathrm{O} \\ & \underset{2.7 \cdot 10^{-4}}{\circ} \end{aligned}$ | $\begin{aligned} & \stackrel{\longrightarrow}{\circ} \\ & \underset{\sim}{\circ} \\ & 1.3 \cdot 10^{-4} \end{aligned}$ | $\xrightarrow[\substack{0}]{\substack{0 \\ 1.2 \cdot 10^{-4}}}$ | $\begin{aligned} & \stackrel{\sim}{\infty \rightarrow 0} \\ & 4.8 \cdot 10^{-5} \end{aligned}$ |
| $\begin{aligned} & 0.0 \\ & 4.5 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & \underset{\infty}{\rightarrow 0} \\ & 2.2 \cdot 10^{-5} \end{aligned}$ |  | $\underset{\substack{0 \rightarrow 6}}{\substack{0 \\ 7.9 \cdot 10^{-6}}}$ | $7.5 \cdot 10^{-8}$ |  |
|  | $\begin{aligned} & \text { aro } \\ & 0.9 \cdot 10^{-8} \\ & 2.0 \end{aligned}$ |  | $\begin{aligned} & \circ \\ & \circ \\ & \circ \\ & \hline \\ & 2.6 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \infty \rightarrow \infty \\ & 1.9 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & 9.0 \\ & 1.7 \cdot 10^{-8} \end{aligned}$ |
| $\begin{aligned} & \circ \quad \circ \\ & \circ \\ & \circ \\ & 2.1 \cdot 10^{-11} \end{aligned}$ | $\begin{aligned} & 90 \\ & 1.4 \cdot 10^{-11} \end{aligned}$ | $\begin{aligned} & \circ \longrightarrow \\ & \circ \\ & \circ \\ & 9.9 \cdot 10^{-12} \end{aligned}$ | $\begin{aligned} & \text { or o } \\ & \text { oo o } \\ & 8.9 \cdot 10^{-12} \end{aligned}$ | $\begin{aligned} & \leftrightarrow \infty \\ & \circ \rightarrow 0 \\ & 6.5 \cdot 10^{-12} \end{aligned}$ | $\begin{aligned} & 9 \\ & 0.8 \\ & 5.8 \cdot 10^{-12} \end{aligned}$ |
| $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \\ & 0.6 \cdot 10^{-12} \end{aligned}$ |  |  |  |  |  |

Table 5: The revised result with the prior network and the imaginary sample size specified as in the first line of this table.

Further we see that for both $n=12$ and $n=2000$, the 32 networks categorize as follows. The 8 networks with both arrows $d r u g \rightarrow w 1$ and $w 1 \rightarrow w 2$ are the 8 most probable networks. In the succeeding 8 networks we have $d r u g \rightarrow w 1$ and $w 1 \nrightarrow w 2$, after that the 8 networks with $d r u g \nrightarrow w 1$ and $w 1 \rightarrow w 2$. In the last 8 networks we have $d r u g \nrightarrow w 1$ and $w 1 \nrightarrow w 2$. Also we see that within each category, the networks are almost equally likely, mostly pronounced for $n=2000$. These finding are what we expected. The arrows included in the prior network are all represented in the most probable networks and these networks are all almost equally likely, as the prior network is a submodel of these. Further there is a large difference in relative score between the different categories, which shows that networks which include the arrows drug $\rightarrow w 1$ and $w 1 \rightarrow w 2$, are much more likely than those that do not. As this is valid for both $n=12$ and $n=2000$, it is not only due to the influence of the prior network, but also because the dataset supports these dependencies.

We will now explore what happens if we choose the prior network to be the least

| Prior network |  |  | Imaginary sample size |  | 2000 |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | $\begin{aligned} & \rightarrow \infty \\ & 1 \\ & \rightarrow 0 \end{aligned}$ |  |  | $\begin{aligned} & \mathrm{O} \longrightarrow \mathrm{O} \\ & 0 \rightarrow 0 \\ & 0.98 \end{aligned}$ |  |
| $\begin{aligned} & \rightarrow 0 \\ & 0 \rightarrow 0 \\ & 0.97 \end{aligned}$ | $\underset{\substack{9}}{\substack{0 \\ 0.96}}$ |  |  | $\begin{aligned} & \stackrel{\circ}{\circ} \underset{\sim}{\circ} \\ & 6.4 \cdot 10^{-4} \end{aligned}$ |  |
| $\begin{aligned} & \circ \stackrel{\circ}{\circ} \\ & \underset{\rightarrow}{\circ} \\ & 1.9 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & \underset{\longrightarrow}{\mathrm{o}} \mathrm{O} \\ & \underset{\longrightarrow}{\longrightarrow} \\ & 1.8 \cdot 10^{-4} \end{aligned}$ |  |  | $3.5 \cdot 10^{-9}$ |  |
|  |  |  |  | $\begin{aligned} & \circ \text { o } \\ & \circ \text { O } \\ & 3.4 \cdot 10^{-9} \end{aligned}$ |  |
|  |  | $\begin{aligned} & \text { o م o } \\ & \text { o } 0.10^{-12} \\ & 2.2 \cdot 10^{2} \end{aligned}$ |  | $\begin{aligned} & \circ \circ \\ & \circ \circ \\ & 6.4 \cdot 10^{-13} \end{aligned}$ |  |
| $\begin{aligned} & \circ \longrightarrow 0 \\ & 0 \quad 0 \\ & 6.4 \cdot 10^{-13} \end{aligned}$ |  |  |  |  |  |

Table 6: The revised result with the prior network and the imaginary sample size specified as in the first line of this table.
probable network from Table 3. The results are for $n=12$ and $n=2000$ given in respectively Table 7 and Table 8.

For $n=12$ we see almost the same result as with the other prior networks. For $n=2000$ we see that the 8 most probable models actually are the 8 models that are possible with both the arrows sex $\rightarrow w 1$ and sex $\rightarrow w 2$. Further we see that all networks are almost equally likely and there is not, as would be expected, a large difference in score between networks with both arrows and the others. Actually for both $n=12$ and $n=2000$ the result is very similar to the result with the empty network as the prior networks. The reason for this is that the probability distribution of the prior network is estimated from data, i.e. we use the sample mean and sample variance as the mean and variance in the prior network. If data does not support a dependence between sex and respectively $w 1$ and $w 2$, then this prior network will be almost the same as the empty prior network and so will the result of the learning procedure. However, it can be seen that even small differences from the empty prior network have an impact

| Prior network |  | $0$ | Imaginary sample size |  | 12 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{array}{ll} 0 & 0 \\ 0 & 0 \\ 1 \end{array}$ | $\begin{aligned} & \underset{\sim}{a \rightarrow 0} \\ & \underset{0.25}{a} \end{aligned}$ | $\begin{aligned} & 0 \xrightarrow{0} \underset{0}{0} \\ & 0.13 \end{aligned}$ | $\begin{gathered} 0 \\ 0 \rightarrow 0 \\ 0.094 \end{gathered}$ | $\begin{aligned} & \substack{\rightarrow 0 \\ 0 \rightarrow 0 \\ 0.023} \end{aligned}$ | $\begin{aligned} & 9 \\ & 0.012 \\ & 0.0 \end{aligned}$ |
| $\begin{aligned} & \stackrel{\text { PO}}{0 \rightarrow 0} \\ & 0.0079 \end{aligned}$ | $\begin{aligned} & \text { 0 } \\ & \underset{0}{\infty} \\ & 0.0020 \end{aligned}$ | $\begin{aligned} & 0 \quad 0 \\ & 0 \rightarrow 0 \\ & 0.0018 \end{aligned}$ | $\begin{aligned} & \stackrel{\rightarrow}{\circ} \\ & \underset{9.6 \cdot 10^{-4}}{\circ} \end{aligned}$ | $7.4 \cdot 10^{-4}$ | $\begin{aligned} & \stackrel{\infty}{\infty \rightarrow} \\ & 3.9 \cdot 10^{-4} \end{aligned}$ |
| $\begin{aligned} & \text { O. } \\ & 0 \rightarrow 0 \\ & 1.8 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & \substack{\circ \\ 1.7 \cdot 10^{-4}} \end{aligned}$ | $\begin{aligned} & \circ \\ & \circ \\ & \circ \\ & \circ \\ & 1.6 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & 9.1 \\ & 1.4 \cdot 10^{-4} \end{aligned}$ | $\begin{aligned} & \underset{\sim}{\rightarrow \rightarrow 0} \\ & 9.0 \cdot 10^{-5} \end{aligned}$ |  |
| $\underset{\substack{0 \rightarrow 0}}{\substack{0.7 \cdot 10^{-5}}}$ | $\begin{aligned} & \text { ob } \\ & 0 \text { o } \\ & 3.5 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & \circ \text { op } \\ & 0 \text { o } \\ & 2.0 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & 980 \\ & 1.8 \cdot 10^{-5} \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 1.2 \cdot 10^{-6} \end{aligned}$ | $\begin{aligned} & 1.1 \cdot 10^{-6} \\ & 0801 \\ & 0 \end{aligned}$ |
| $\begin{aligned} & \circ \text { م } 0 \\ & 0 \text { o } \\ & 3.1 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \circ \quad \circ \\ & \circ \\ & \circ \\ & 2.9 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & 9 \\ & 080 \\ & 2.8 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & 0.0 \\ & 0.6 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \circ \longrightarrow 0 \\ & \circ \circ \\ & 1.5 \cdot 10^{-7} \end{aligned}$ | $\begin{aligned} & \infty \\ & 1.4 \cdot 10^{-7} \end{aligned}$ |
| $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \rightarrow 0 \\ & 6.2 \cdot 10^{-8} \end{aligned}$ | $\begin{aligned} & \text { os } \\ & 0<0 \\ & 5.5 \cdot 10^{-8} \end{aligned}$ |  |  |  |  |

Table 7: The revised result with the prior network and the imaginary sample size specified as in the first line of this table.
when $n$ is large, as the 8 most probable networks actually are the ones with both sex $\rightarrow w 1$ and sex $\rightarrow w 2$.

| Prior network |  | $0$ | Imaginary sample size |  | 2000 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\begin{gathered} \infty \rightarrow 0 \\ 0 \rightarrow 0 \end{gathered}$ <br> 1 | $\begin{gathered} \rightarrow 0 \\ 0 \rightarrow 0 \\ 0.99 \end{gathered}$ | $\underset{\substack{0}}{\substack{0 \\ 0.94}}$ | $\begin{aligned} & \underset{\sim \rightarrow 0}{\rightarrow 0} \\ & 0.94 \end{aligned}$ | $\begin{aligned} & 0,0 \\ & 000 \\ & 0.91 \end{aligned}$ | $\begin{aligned} & \longrightarrow 8 \\ & 0 \\ & 0.90 \end{aligned}$ |
| oros | $\begin{aligned} & \substack{0 \\ 0 \\ 0.86} \end{aligned}$ | $\begin{aligned} & \substack{8 \\ 0 \rightarrow 0 \\ 0.67} \end{aligned}$ | $\begin{gathered} 9 \\ 0.0 \\ 0.65 \end{gathered}$ | $\begin{aligned} & 9 \\ & 0.1 \\ & 0.61 \end{aligned}$ | $\begin{aligned} & 9 \\ & 0 \\ & 0.59 \\ & 0.59 \end{aligned}$ |
| $\begin{aligned} & 9 \\ & 0.0 \\ & 0.59 \end{aligned}$ | $\begin{aligned} & 00 \\ & 0 \rightarrow 0 \\ & 0.59 \end{aligned}$ | $\begin{aligned} & 9 \\ & 0 \\ & 0.54 \end{aligned}$ | $\begin{aligned} & 90 \\ & 0 \text { o } \\ & 0.53 \end{aligned}$ | $\begin{aligned} & \substack{0 \rightarrow 0 \\ 0 \rightarrow 0 \\ 0.38} \end{aligned}$ | $\begin{aligned} & \stackrel{\rightarrow}{\circ \rightarrow+} \\ & 0 \rightarrow 0 \\ & 0.37 \end{aligned}$ |
|  | $\begin{aligned} & \underset{\rightarrow 0}{\rightarrow 0} \\ & 0.35 \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \\ & 0.34 \\ & 0.34 \end{aligned}$ | $\begin{aligned} & \mathrm{O} \rightarrow \mathrm{o} \\ & \circ \\ & 0.0 \\ & 0.33 \end{aligned}$ | $\begin{aligned} & 0 \rightarrow 0 \\ & 0 \\ & 0.32 \end{aligned}$ | $\begin{aligned} & \circ \longrightarrow 0 \\ & \circ \\ & 0 \\ & 0.31 \end{aligned}$ |
| $\begin{aligned} & \circ \xrightarrow{\circ}+ \\ & 0.05 \end{aligned}$ | $\begin{array}{lr} 0 & 0 \\ 0 \rightarrow 0 \\ 0.25 \end{array}$ | $\begin{gathered} 0 \\ 0 \\ 0 \\ 0.22 \\ 0.22 \end{gathered}$ | $\begin{aligned} & 0 \\ & 0.22 \\ & 0.0 \\ & 0.0 \end{aligned}$ | $\begin{array}{lr} \circ \quad 0 \\ 0 \rightarrow 0 \\ 0.22 \end{array}$ | $\begin{array}{ll} \circ & 0 \\ 0 & 0 \\ 0.22 \end{array}$ |
| $\begin{aligned} & \text { o م } \\ & \text { o o } \\ & 0.20 \end{aligned}$ | $\begin{array}{lr} \circ & \circ \\ \circ & \circ \\ 0.20 \end{array}$ |  |  |  |  |

Table 8: The revised result with the prior network and the imaginary sample size specified as in the first line of this table.

## Acknowledgements

This research was supported by the ESPRIT project P29105 (BaKE). Also I would like to thank my supervisor Steffen L. Lauritzen for his dedicated help and support.

## References

Bernardo, J. M. and Smith, A. F. M. (1994). Bayesian Theory, John Wiley \& Sons, Chichester.
Bøttcher, S. G. (2001). Learning Bayesian Networks with Mixed Variables, Artificial Intelligence and Statistics 2001, Morgan Kaufmann, San Francisco, CA, USA, pp. 149-156.
Chickering, D. M. (1995). A transformational characterization of equivalent Bayesian-network structures, Proceedings of Eleventh Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 87-98.
Cooper, G. and Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data, Machine Learning 9: 309-347.
Dawid, A. P. and Lauritzen, S. L. (1993). Hyper Markov laws in the statistical analysis of decomposable graphical models, The Annals of Statistics 21(3): 1272-1317.
DeGroot, M. H. (1970). Optimal Statistical Decisions, McGraw-Hill, New York.
Edwards, D. (1995). Introduction to Graphical Modelling, Springer-Verlag, New York.

Frydenberg, M. (1990). Marginalization and collapsibility in graphical interaction models, Annals of Statistics 18: 790-805.
Geiger, D. and Heckerman, D. (1994). Learning Gaussian Networks, Proceedings of Tenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 235-243.
Heckerman, D., Geiger, D. and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data, Machine Learning 20: 197-243.
Lauritzen, S. L. (1992). Propagation of probabilities, means and variances in mixed graphical association models, Journal of the American Statistical Association 87(420): 1098-1108.
Lauritzen, S. L. (1996). Graphical Models, Clarendon press, Oxford, New York.
Lauritzen, S. L. and Jensen, F. (2001). Stable local computation with conditional Gaussian distributions, Statistics and Computing 11: 191-203.
Morrison, D. F. (1976). Multivariate Statistical Methods, McGraw-Hill, USA.
Robinson, R. W. (1977). Counting unlabeled acyclic digraphs, Lecture Notes in Mathematics, 622: Combinatorial Mathematics V pp. 239-273.

Seber, G. A. F. (1984). Multivariate Observations, John Wiley and Sons, New York.

Shachter, R. D. and Kenley, C. R. (1989). Gaussian influence diagrams, Management Science 35: 527-550.

Spiegelhalter, D. J. and Lauritzen, S. L. (1990). Sequential updating of conditional probabilities on directed graphical structures, Networks 20: 579605.

Steck, H. and Jaakkola, T. S. (2002). On the Dirichlet Prior and Bayesian Regularization, Conference on Advances in Neural Information Processing Systems, Vol. 15, MIT Press, Cambridge, USA.

# Paper II <br> deal: A Package for Learning <br> Bayesian Networks 

Susanne G. Bøttcher and Claus Dethlefsen

Published in:
Journal of Statistical Software (2003), 8(20):1-40.

# deal: A Package for Learning Bayesian Networks 

Susanne G. Bøttcher<br>Aalborg University, Denmark<br>Claus Dethlefsen<br>Aalborg University, Denmark


#### Abstract

. deal is a software package for use with $R$. It includes several methods for analyzing data using Bayesian networks with variables of discrete and/or continuous types but restricted to conditionally Gaussian networks. Construction of priors for network parameters is supported and their parameters can be learned from data using conjugate updating. The network score is used as a metric to learn the structure of the network and forms the basis of a heuristic search strategy. deal has an interface to Hugin.


## 1 Introduction

A Bayesian network is a graphical model that encodes the joint probability distribution for a set of random variables. Bayesian networks are treated in e.g. Cowell et al. (1999) and have found application within many fields, see Lauritzen (2003) for a recent overview.

Here we consider Bayesian networks with mixed variables, i.e. the random variables in a network can be of both discrete and continuous types. A method for learning the parameters and structure of such Bayesian networks has recently been described by Bøttcher (2001). We have developed a package called deal, written in R (R Development Core Team 2003), which provides these methods for learning Bayesian networks. In particular, the package includes procedures for defining priors, estimating parameters, calculating network scores, performing heuristic search as well as simulating data sets with a given dependency structure. Figure 1 gives an overview of the functionality in deal. The package can be downloaded from the Comprehensive R Archive Network (CRAN) http://cran.R-project.org/ and may be used under the terms of the

GNU General Public License Version 2.


Figure 1: From prior knowledge and training data, a posterior network is produced by deal. The network may be transferred to Hugin for further inference.

In Section 2 we define Bayesian networks for mixed variables. To learn a Bayesian network, the user needs to supply a training data set and represent any prior knowledge available as a Bayesian network. Section 3 shows how to specify the training data set in deal and Section 4 discusses how to specify a Bayesian network in terms of a Directed Acyclic Graph (DAG) and the local probability distributions.
deal uses the prior Bayesian network to deduce prior distributions for all parameters in the model. Then, this is combined with the training data to yield posterior distributions of the parameters. The parameter learning procedure is treated in Section 5.

Section 6 describes how to learn the structure of the network. A network score is calculated and a search strategy is employed to find the network with the highest score. This network gives the best representation of data and we call it the posterior network.

Section 7 describes how to transfer the posterior network to Hugin (http: //www.hugin.com). The Hugin graphical user interface (GUI) can then be
used for further inference in the posterior network.
In the appendix we provide manual pages for the main functions in deal.

## 2 Bayesian Networks

Let $D=(V, E)$ be a Directed Acyclic Graph (DAG), where $V$ is a finite set of nodes and $E$ is a finite set of directed edges (arrows) between the nodes. The DAG defines the structure of the Bayesian network.

To each node $v \in V$ in the graph corresponds a random variable $X_{v}$. The set of variables associated with the graph $D$ is then $X=\left(X_{v}\right)_{v \in V}$. Often, we do not distinguish between a variable $X_{v}$ and the corresponding node $v$. To each node $v$ with parents $\mathrm{pa}(v)$ a local probability distribution, $p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)$, is attached. The set of local probability distributions for all variables in the network is $\mathcal{P}$.

A Bayesian network for a set of random variables $X$ is the pair $(D, \mathcal{P})$.
The possible lack of directed edges in $D$ encodes conditional independencies between the random variables $X$ through the factorization of the joint probability distribution,

$$
p(x)=\prod_{v \in V} p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)
$$

Here, we allow Bayesian networks with both discrete and continuous variables, as treated in Lauritzen (1992), so the set of nodes $V$ is given by $V=\Delta \cup \Gamma$, where $\Delta$ and $\Gamma$ are the sets of discrete and continuous nodes, respectively. The set of variables $X$ can then be denoted $X=\left(X_{v}\right)_{v \in V}=(I, Y)=$ $\left(\left(I_{\delta}\right)_{\delta \in \Delta},\left(Y_{\gamma}\right)_{\gamma \in \Gamma}\right)$, where $I$ and $Y$ are the sets of discrete and continuous variables, respectively. For a discrete variable, $\delta$, we let $\mathcal{I}_{\delta}$ denote the set of levels.

To ensure e.g. availability of exact local computation methods, we do not allow discrete variables to have continuous parents. The joint probability distribution then factorizes into a discrete part and a mixed part, so

$$
p(x)=p(i, y)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) \prod_{\gamma \in \Gamma} p\left(y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}\right) .
$$

## 3 Data Structure

deal expects data as specified in a data frame which is a standard data structure in R. For example, standard ASCII data files with one column per variable and one line per observation can be read using read. table () which returns a data frame.

The rats example in Table 1 was constructed by Morrison (1976) and also studied in Edwards (1995). The data set is from a hypothetical drug trial, where the weight losses of male and female rats under three different drug treatments have been measured after one and two weeks.

| Sex | Drug | W1 | W2 |
| ---: | ---: | ---: | ---: |
| M | D1 | 5 | 6 |
| M | D1 | 7 | 6 |
| M | D1 | 9 | 9 |
| M | D1 | 5 | 4 |
| M | D2 | 9 | 12 |
| M | D2 | 7 | 7 |
| M | D2 | 7 | 6 |
| M | D2 | 6 | 8 |
| M | D3 | 14 | 11 |
| M | D3 | 21 | 15 |
| M | D3 | 12 | 10 |
| M | D3 | 17 | 12 |
| F | D1 | 7 | 10 |
| F | D1 | 8 | 10 |
| F | D1 | 6 | 6 |
| F | D1 | 9 | 7 |
| F | D2 | 7 | 6 |
| F | D2 | 10 | 13 |
| F | D2 | 6 | 9 |
| F | D2 | 8 | 7 |
| F | D3 | 14 | 9 |
| F | D3 | 14 | 8 |
| F | D3 | 16 | 12 |
| F | D3 | 10 | 5 |

Table 1: An example data file, rats.dat.

The data are loaded into a data frame rats. df by the following command

```
rats.df <- read.table("rats.dat",header=TRUE)
```

Before continuing, it is essential that the column variables have the correct types. Discrete variables should be specified as factors and continuous variables as numeric. To alter the type of a variable so that it is regarded as a discrete variable, use the factor () function (standard in R). In the rats example, Sex and

Drug are interpreted to be factors by read.table, and thus no changes are necessary.

We assume that we have observed complete data which means that no NA's are present in the data frame.

## 4 Specification of a Bayesian Network

As described in Section 2, a Bayesian network is specified by a DAG and a set of local probability distributions. In this section we will show how to specify these terms in deal.

### 4.1 The Network Class and Associated Methods

In deal, a Bayesian network is represented as an object of class network. The network object is a list of properties that are added or changed by the methods described in following sections.

A network is generated by the following command

```
rats <- network(rats.df)
```

and by default it is set to the empty network (the network without any arrows).
If the option specifygraph is set, a point and click graphical interface allows the user to insert and delete arrows until the requested DAG is obtained.

```
rats <- network(rats.df,specifygraph=TRUE)
```

A plot of the network (see Figure 2) is generated by

```
plot(rats)
```

Note that discrete nodes are grey and continuous nodes are white.
The primary property of a network is the list of nodes, in the example the list is: nodes (rats). Each entry in the list is an object of class node representing a


Figure 2: Graphical representation of the rats network.
node in the graph, which includes information associated with the node. Several methods for the network class operate by applying an appropriate method for one or more nodes in the list of nodes. The nodes appear in the node list in the same order as in the data frame used to create the network object.

It is possible to access the individual nodes in a network by referring either to their index (the column number in the data frame) or to their name:

```
rats.nd <- nodes(rats)# the list of nodes
rats.nd[[1]] # the first node
rats.nd$Drug # the node '`Drug''
```

A collection of networks is represented in an object of class networkfamily, which has associated print () and plot () functions.

### 4.2 Specification of the Probability Distributions

The joint distribution of the random variables in a network in deal is a conditional Gaussian (CG) distribution.

For discrete nodes, this means that the local probability distributions are unrestricted discrete distributions. We parameterize this as

$$
\theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)
$$

where $\theta_{\delta| |_{\mathrm{paq}}(\delta)}=\left(\theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}$. The parameters fulfill $\sum_{i_{\delta} \in \mathcal{I}_{\delta}} \theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=1$ and $0 \leq \theta_{\left.i_{\delta}\right|_{\mathrm{pa}(\delta)}} \leq 1$.

For continuous nodes, the local probability distributions are Gaussian linear regressions on the continuous parents with parameters depending on the configuration of the discrete parents. We parameterize this as

$$
\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}=\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma| |_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right),
$$

so that

$$
\left(Y_{\gamma} \mid i_{\mathrm{pa}(\gamma)}, y_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) \sim \mathcal{N}\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+y_{\mathrm{pa}(\gamma)} \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)
$$

A suggestion for the local probability distributions is generated and attached to each node as the property prob. The suggestion can then be edited afterwards.

For a discrete variable $\delta$, the suggested local probability distribution $p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right)$ is taken to be uniform over the levels for each parent configuration, i.e.

$$
p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right)=1 / \mathcal{I}_{\delta}
$$

Define $z_{\mathrm{pa}(\gamma)}=\left(1, y_{\mathrm{pa}(\gamma)}\right)$ and let $\eta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}=\left(m_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)$, where $m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ is the intercept and $\beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ is the vector of coefficients. For a continuous variable $\gamma$, the suggested local probability distribution

$$
\mathcal{N}\left(z_{\mathrm{pa}(\gamma)} \eta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)
$$

is determined as a regression on the continuous parents for each configuration of the discrete parents.

The prob property for discrete nodes is a multi-way array with the node itself occupying the first dimension and the parents each occupying one dimension. For continuous nodes, $\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}$ and $\eta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ are stored in a matrix with one row for each configuration of the discrete variables. The first column contains $\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}$ and the remaining columns $\eta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$.

It is possible to inspect the suggested local probability distributions by setting the option inspectprob as

```
rats <- network(rats.df,inspectprob=TRUE)
```

This gives a graphical way of inspecting the local probability distribution by clicking on the nodes. Then, it is possible to adjust the local distributions, e.g.

```
localprob(rats,"Sex") <- c(0.6, 0.4)
localprob(rats,"W1") <- c(10, 0)
localprob(rats,"W2") <- c(10, 0, 1) # if eg. W2|W1
```


### 4.3 The Joint Distribution

We now show how the joint probability distribution of a network can be calculated from the local probability distributions.

For the discrete part of the network, the joint probability distribution is found as

$$
p(i)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) .
$$

For continuous variables, the joint distribution $\mathcal{N}\left(M_{i}, \Sigma_{i}\right)$ is determined for each configuration of the discrete variables by applying the following sequential algorithm, see Shachter and Kenley (1989).

The order is determined so that the joint distribution of the parents have already been determined for the current node. For notational convenience, we skip the index $i$ and determine the joint distribution of node $\gamma$ and all previously processed nodes, $p$. From the prior network, we have given $\eta_{\gamma \mid \mathrm{pa}(\gamma)}=$ $\left(m_{\gamma \mid \mathrm{pa}(\gamma)}, \beta_{\gamma \mid \mathrm{pa}(\gamma)}\right)$ and $\sigma_{\gamma \mid \mathrm{pa}(\gamma)}^{2}$. Previously evaluated are $M_{p}$ and $\Sigma_{p}$. Now, the covariance is given by

$$
\Sigma_{\gamma, p}=\Sigma_{p} \beta_{\gamma \mid p}
$$

where $\beta_{\gamma \mid p}$ is a column vector of the regression coefficients given all previously evaluated nodes. All coefficients are zero, except the coefficients, $\beta_{\gamma \mid \mathrm{pa}(\gamma)}$, corresponding to the parents of the node. The variance and mean are then given by

$$
\begin{aligned}
\Sigma_{\gamma} & =\sigma_{\gamma \mid \mathrm{pa}(\gamma)}^{2}+\Sigma_{\gamma, p} \beta_{\gamma \mid p} \\
M_{\gamma} & =m_{\gamma \mid \mathrm{pa}(\gamma)}+\beta_{\gamma \mid p}^{\top} M_{p} .
\end{aligned}
$$

In deal, we can assess these quantities by
rats.j <- jointprior(rats)
and inspect the properties jointmu, containing $M_{i}$, jointsigma, containing $\Sigma_{i}$, and jointalpha. The discrete part, $p(i)$, is not returned directly, but is found by dividing rats.j\$jointalpha with sum(rats.j\$jointalpha).

## 5 Parameter Learning

In the previous section we showed how to specify a Bayesian network, i.e. a DAG and the local probability distributions. In this section we will show how to estimate the parameters in the local probability distributions from data. The first sections present the theory behind the learning procedure and the last section shows how it is done in deal.

### 5.1 The Bayesian Approach

To estimate the parameters in the network, we use a Bayesian approach. We encode our uncertainty about parameters $\theta$ in a prior distribution $p(\theta)$, use data $d$ to update this distribution, and hereby obtain the posterior distribution $p(\theta \mid d)$ by using Bayes' theorem,

$$
\begin{equation*}
p(\theta \mid d)=\frac{p(d \mid \theta) p(\theta)}{p(d)}, \quad \theta \in \Theta \tag{1}
\end{equation*}
$$

Here $\Theta$ is the parameter space, $d$ is a random sample from the probability distribution $p(x \mid \theta)$ and $p(d \mid \theta)$ is the joint probability distribution of $d$, also called the likelihood of $\theta$. We refer to this as parameter learning or just learning.

In deal, we assume that the parameters associated with one variable are independent of the parameters associated with the other variables and, in addition, that the parameters are independent for each configuration of the discrete parents, i.e.

$$
\begin{equation*}
p(\theta)=\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) . \tag{2}
\end{equation*}
$$

We refer to (2) as parameter independence. Further, as we have assumed complete data, the parameters stay independent given data, see Bøttcher (2001). This means that we can learn the parameters of a node independently of the parameters of the other nodes, i.e. we update the local parameter prior $p\left(\theta_{v \mid i_{\mathrm{pa}(v)}}\right)$ for each node $v$ and each configuration of the discrete parents.

As local prior parameter distributions, we use the Dirichlet distribution for the discrete variables and the Gaussian-inverse gamma distribution for the continuous variables. These distributions are conjugate to observations from the respective distributions and this ensures simple calculations of the posterior distributions.

In the next section we present an automated procedure for specifying the local parameter priors associated with any possible DAG. The procedure is called the master prior procedure. For the mixed case it is treated in Bøttcher (2001), for the purely discrete and the purely continuous cases it is treated in Heckerman et al. (1995) and Geiger and Heckerman (1994), respectively.

### 5.2 The Master Prior Procedure

In the following sections we will show how to deduce and update the local prior parameter distributions for discrete and continuous nodes, respectively. Here, we will summarize the steps in the master prior procedure.

The idea is that from a given Bayesian network, we can deduce parameter priors for any possible DAG. The user just has to specify the Bayesian network as he believes it to be. We call this network a prior Bayesian network.

1. Specify a prior Bayesian network, i.e. a prior DAG (Section 4.1) and prior local probability distributions (Section 4.2). Calculate the joint prior distribution (Section 4.3).
2. From this joint prior distribution, the marginal distribution of all parameters in the family consisting of the node and its parents can be determined. We call this the master prior.
3. The local parameter priors are now determined by conditioning in the master prior distribution.

This procedure ensures parameter independence. Further, it has the property that if a node has the same set of parents in two different networks, then the local parameter prior for this node will be the same in the two networks. Therefore, we only have to deduce the local parameter prior for a node given the same set of parents once. This property is called parameter modularity.

### 5.3 Discrete Nodes

We will now show how to find the local parameter priors for the discrete nodes. Recall that the local probability distributions are unrestricted discrete distributions defined as in Section 4.2.

## Master Prior

Let $\Psi=\left(\Psi_{i}\right)_{i \in \mathcal{I}}$ be the parameters for the joint distribution of the discrete variables. The joint prior parameter distribution is assumed to be a Dirichlet distribution

$$
p(\Psi) \sim \mathcal{D}(\alpha)
$$

with hyperparameters $\alpha=\left(\alpha_{i}\right)_{i \in \mathcal{I}}$. To specify this Dirichlet distribution, we need to specify these hyperparameters.

Consider the following relation for the Dirichlet distribution,

$$
p(i)=\mathbb{E}\left(\Psi_{i}\right)=\frac{\alpha_{i}}{N}
$$

with $N=\sum_{i \in \mathcal{I}} \alpha_{i}$. Now we use the probabilities in the prior network as an estimate of $\mathbb{E}\left(\Psi_{i}\right)$, so we only need to determine $N$ in order to calculate the parameters $\alpha_{i}$.

We determine $N$ by using the notion of an imaginary data base. We imagine that we have a data base of cases, from which we have updated the distribution of $\Psi$ out of total ignorance. The imaginary sample size of this imaginary data base is thus $N$. It expresses how much confidence we have in the (in)dependencies expressed in the prior network, see Heckerman et al. (1995).

We use this joint distribution to deduce the master prior distribution of the family $A=\delta \cup \mathrm{pa}(\delta)$. Let

$$
\alpha_{i_{A}}=\sum_{j: j_{A}=i_{A}} \alpha_{j}
$$

and let $\alpha_{A}=\left(\alpha_{i_{A}}\right)_{i_{A} \in \mathcal{I}_{A}}$. Then the marginal distribution of $\Psi_{A}$ is Dirichlet, $p\left(\Psi_{A}\right) \sim \mathcal{D}\left(\alpha_{A}\right)$. This is the master prior in the discrete case.

## Local Parameter Prior

From the master prior, we calculate the conditional distribution $\Psi_{\delta \mid i_{\mathrm{pa}(\delta)}}=$ $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ which is the local parameter prior in the discrete case. Then,

$$
\begin{gathered}
\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=\alpha_{i_{A}}, \\
\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(\alpha_{i \delta \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}, \\
\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid \alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}} \sim \mathcal{D}\left(\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}} .\right.
\end{gathered}
$$

## Local Parameter Posterior

Let $n_{\delta \mid i_{\mathrm{pa}(\delta)}}$ be the number of cases observed with the particular parent configuration in the data base and let $n$ be the total number of observations.

Then, the posterior parameters $\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}^{\prime}$ are given by

$$
\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}^{\prime}}^{\prime}=\alpha_{\delta \mid i_{\mathrm{pa}}(\delta)}+n_{\delta \mid i_{\mathrm{pa}(\delta)}}
$$

### 5.4 Continuous Nodes

We now show how to find the local parameter priors for the continuous nodes. Recall that the local probability distributions are normal distributions defined as in Section 4.2.

## Master Prior

Bøttcher (2001) derived this procedure in the mixed case. For a configuration $i$ of the discrete variables we let $\nu_{i}=\rho_{i}=\alpha_{i}$, where $\alpha_{i}$ was determined in Section 5.3. Also, $\Phi_{i}=\left(\nu_{i}-1\right) \Sigma_{i}$.

The joint parameter priors are assumed to be distributed as

$$
\begin{aligned}
p\left(M_{i} \mid \Sigma_{i}\right) & =\mathcal{N}\left(\mu_{i}, \frac{1}{\nu_{i}} \Sigma_{i}\right) \\
p\left(\Sigma_{i}\right) & =\mathcal{I} W\left(\rho_{i}, \Phi_{i}\right)
\end{aligned}
$$

where $\mathcal{I} W$ is the inverse Wishart distribution.

However, since the marginal distribution of a CG distribution is not necessarily a CG distribution, there is no simple way to derive priors for other networks. Instead we use the imaginary data base to derive local master priors.

Define the notation

$$
\rho_{i_{A \cap \Delta}}=\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} \rho_{j}
$$

and similarly for $\nu_{i_{A \cap \Delta}}$ and $\Phi_{i_{A \cap \Delta}}$. For the family $A=\gamma \cup \mathrm{pa}(\gamma)$, the local
master prior is then found as

$$
\begin{aligned}
\Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}} & \sim \mathcal{I} W\left(\rho_{i_{A \cap \Delta}}, \tilde{\Phi}_{A \cap \Gamma \mid i_{A \cap \Delta}}\right) \\
M_{A \cap \Gamma \mid i_{A \cap \Delta}} \mid \Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}} & \sim \mathcal{N}\left(\bar{\mu}_{A \cap \Gamma \mid i_{A \cap \Delta}}, \frac{1}{\nu_{i_{A \cap \Delta}}} \Sigma_{A \cap \Gamma \mid i_{A \cap \Delta}}\right)
\end{aligned}
$$

where

$$
\begin{aligned}
\bar{\mu}_{i_{A \cap \Delta}} & =\frac{\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} \mu_{j} \nu_{j}}{\nu_{i_{A \cap \Delta}}} \\
\tilde{\Phi}_{A \cap \Gamma \mid i_{A \cap \Delta}} & =\Phi_{i_{A \cap \Delta}}+\sum_{j: j_{A \cap \Delta}=i_{A \cap \Delta}} \nu_{j}\left(\mu_{j}-\bar{\mu}_{i_{A \cap \Delta}}\right)\left(\mu_{j}-\bar{\mu}_{i_{A \cap \Delta}}\right)^{\top} .
\end{aligned}
$$

## Local Parameter Prior

Using $\mathcal{I} \Gamma$ for the inverse gamma distribution, the local prior parameters, given as

$$
\begin{aligned}
\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right) & \sim \mathcal{N}\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\right), \\
\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2} & \sim \mathcal{I} \Gamma\left(\frac{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}{2}, \frac{\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}{2}\right),
\end{aligned}
$$

are deduced from the local master prior by conditioning as follows. To simplify notation, we ignore all subscripts in the master prior and thus consider the configuration $i_{A \cap \Delta}=i_{\mathrm{pa}(\gamma)}$ and assume that $A \cap \Gamma$ is ordered with $\gamma$ as the first entry. Write $\mathrm{pa}(\gamma)$ for the continuous parents $\{A \cap \Gamma\} \backslash\{\gamma\}$. Define the partitioning

$$
\begin{aligned}
\tilde{\Phi}_{A \cap \Gamma \mid i_{A \cap \Delta}} & =\left[\begin{array}{cc}
\tilde{\phi}_{\gamma} & \tilde{\Phi}_{\gamma, \mathrm{pa}(\gamma)} \\
\tilde{\Phi}_{\mathrm{pa}(\gamma), \gamma} & \tilde{\Phi}_{\mathrm{pa}(\gamma)}
\end{array}\right] \\
\bar{\mu}_{i_{A \cap \Delta}} & =\left(\bar{\mu}_{\gamma}, \bar{\mu}_{\mathrm{pa}(\gamma)}\right)
\end{aligned}
$$

Then

$$
\begin{aligned}
\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & \left.=\left(\bar{\mu}_{\gamma}-\tilde{\Phi}_{\gamma, \mathrm{pa}(\gamma)}\right) \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1} \bar{\mu}_{\mathrm{pa}(\gamma)}, \tilde{\Phi}_{\gamma, \mathrm{pa}(\gamma)} \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1}\right) \\
\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & =\tilde{\phi}_{\gamma}-\tilde{\Phi}_{\gamma, \mathrm{pa}(\gamma)} \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1} \tilde{\Phi}_{\mathrm{pa}(\gamma), \gamma}, \\
\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & =\rho_{i_{A \cap \Delta}}+|\mathrm{pa}(\gamma)|, \\
\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & =\left(\begin{array}{cc}
1 / \nu_{i_{A \cap \Delta}}+\bar{\mu}_{\mathrm{pa}(\gamma)}^{\top} \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1} \bar{\mu}_{\mathrm{pa}(\gamma)} & -\bar{\mu}_{\mathrm{pa}(\gamma)}^{\top} \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1}, \\
-\tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1} \bar{\mu}_{\mathrm{pa}(\gamma)} & \tilde{\Phi}_{\mathrm{pa}(\gamma)}^{-1}
\end{array}\right)^{-1}
\end{aligned}
$$

## Local Parameter Posterior

Define $z_{\mathrm{pa}(\gamma) \mid i_{\mathrm{pa}(\gamma)}}^{b}$ as the matrix with $n$ rows and with a column of ones and columns of the observed continuous parents for a given configuration of the discrete parents. Let $y_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{b}$ be the vector of observations of the node $\gamma$ for a configuration of the discrete parents.

Then, the prior parameters $\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ are updated to posterior parameters (denoted with a prime) by the following relations

$$
\begin{aligned}
\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\left(z_{\mathrm{pa}(\gamma) \mid i_{\mathrm{pa}(\gamma)}}^{b}\right)^{\top} z_{\mathrm{pa}(\gamma) \mid i_{\mathrm{pa}(\gamma)}}^{b}, \\
\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{-1} \times\left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\left(z_{\mathrm{pa}(\gamma) \mid i_{\mathrm{pa}(\gamma)}}^{b}\right)^{\top} y_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{b}\right) \\
\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+n, \\
\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \\
& +\left(y_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{\top} y_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{b} \\
& +\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}-\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{\top} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}} .
\end{aligned}
$$

### 5.5 The Learning Procedure in deal

Assume that the training data are available in a data frame, rats.df, as described in Section 3. Also, assume that the user has specified a Bayesian network to be used as prior network, called rats, see Section 4.

The parameters of the joint distribution of the variables in the network are then determined by the function jointprior () with the size of the imaginary data base as optional argument. If the size is not specified, deal sets the size to a reasonably small value.

```
rats.prior <- jointprior(rats)
## auto set size of imaginary data base
rats.prior <- jointprior(rats,12)
## set size of imaginary data base to 12
```

The parameters in the object rats. prior may be assessed as

```
rats.prior$jointalpha
rats.prior$jointnu
rats.prior$jointrho
rats.prior$jointphi
```

The procedure learn () determines the master prior, local parameter priors and local parameter posteriors and may be called on all nodes or just a single node. The result is accessed using the getnetwork () extractor function.

```
rats <- getnetwork(learn(rats,rats.df,rats.prior))
## all nodes
rats <- getnetwork(learn(rats,rats.df,rats.prior, 2))
## only node 2
```

In the result, each learned node has now attached two properties. These contain the parameters in the local prior distribution and the parameters in the local posterior distribution, respectively. For the node $S$ ex, the properties are assessed as
localprior (nodes (rats) \$Sex)
localposterior (nodes (rats) \$Sex)

## 6 Learning the Structure

In this section we will show how to learn the structure of the DAG from data. The section is based on Bøttcher (2001), Heckerman et al. (1995) and Geiger and Heckerman (1994).

### 6.1 Network Score

As a measure of how well a DAG $D$ represents the conditional independencies between the random variables, we use the relative probability

$$
S(D)=p(D, d)=p(d \mid D) p(D)
$$

and refer to it as a network score.

The network score factorizes into a discrete part and a mixed part as

$$
S(D)=\prod_{\delta \in \Delta} S_{\delta}(D) \prod_{\gamma \in \Gamma} S_{\gamma}(D)
$$

where $S_{\delta}(D)$ is the contribution from the discrete node $\delta$ and $S_{\gamma}(D)$ is the contribution from the continuous node $\gamma$.

For a discrete node, $\delta$, the score contribution is given by
where $\alpha_{+\delta \mid i_{\mathrm{pa}(\delta)}}=\sum_{i_{\delta} \in \mathcal{I}_{\delta}} \alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}$ and $n_{+\delta \mid i_{\mathrm{pa}(\delta)}}=\sum_{i_{\delta} \in \mathcal{I}_{\delta}} n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}$.
For a continuous node, $\gamma$,

$$
\begin{aligned}
S_{\gamma}(D)= & \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \frac{\Gamma\left(\frac{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+n}{2}\right)}{\Gamma\left(\frac{\left.\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)}{} \sqrt{\operatorname{det}\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} s_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \pi\right)}\right.} \times \\
& {\left[1+\frac{1}{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}} a_{\gamma \mid i_{\mathrm{pa}(\gamma)}} s_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1} a_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\top}\right]^{-\frac{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}}{+n}} }
\end{aligned}
$$

where

$$
\begin{aligned}
s_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & =\frac{\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}\left(I+z_{\mathrm{pa}(\gamma)}^{b} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\top}\right), \\
a_{\gamma \mid i_{\mathrm{pa}(\gamma)}} & =y_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{b}-z_{\mathrm{pa}(\gamma) \mid i_{\mathrm{pa}(\gamma)}}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}} .
\end{aligned}
$$

Note that the network score factorizes into a product over terms involving only one node and its parents. This property is called decomposability.

It can be shown that the network scores for two independence equivalent DAGs are equal. This property is called likelihood equivalence and it is a property of the master prior procedure.

In deal we use, for computational reasons, the logarithm of the network score. The log network score contribution of a node is evaluated whenever the node is learned and the $\log$ network score is updated. The results are inspected as

```
rats <- getnetwork(learn(rats,rats.df,rats.prior))
score (nodes (rats) $Sex)
score(rats) # log network score
```


### 6.2 Model Search

In principle, we could evaluate the network score for all possible DAGs and indeed this is provided in deal.

```
allrats <- networkfamily(rats.df,rats,rats.prior)
allrats <- nwfsort(getnetwork(allrats))
```

However, the number of possible DAGs grows more than exponentially with the number of nodes (see Table 2) and, in general, the problem of identifying the network with the highest score is NP-complete (see Chickering (1996)). If

| \# nodes | \# networks |
| ---: | ---: |
| 1 | 1 |
| 2 | $2-3$ |
| 3 | $12-25$ |
| 4 | $144-543$ |
| 5 | $4800-29281$ |
| 6 | $320000-3781503$ |
| 7 | $\approx 56 \cdot 10^{6}-10^{9}$ |
| 8 | $\approx 10^{10}-10^{11}$ |
| 9 | $\approx 10^{13}-10^{15}$ |
| 10 | $\approx 10^{16}-10^{18}$ |

Table 2: The (approximate) number of networks for a given number of nodes. Since we do not allow arrows from continuous to discrete nodes, the number of networks for a given number of nodes is given as a lower and upper bound.
the number of random variables in a network is large, it is not computationally possible to calculate the network score for all the possible DAGs. For these situations a strategy for searching for DAGs with high score is needed. In deal, the search strategy greedy search with random restarts, see Heckerman et al. (1995), is implemented. As a way of comparing the network scores for two different DAGs, $D$ and $D^{*}$, we use the posterior odds,

$$
\frac{p(D \mid d)}{p\left(D^{*} \mid d\right)}=\frac{p(D, d)}{p\left(D^{*}, d\right)}=\frac{p(D)}{p\left(D^{*}\right)} \times \frac{p(d \mid D)}{p\left(d \mid D^{*}\right)},
$$

where $p(D) / p\left(D^{*}\right)$ is the prior odds and $p(d \mid D) / p\left(d \mid D^{*}\right)$ is the Bayes factor. At the moment, the only option in deal for specifying prior distribution over

DAGs is to let all DAGs be equally likely, so the prior odds are always equal to one. Therefore, we use the Bayes factor for comparing two different DAGs.

In greedy search we compare models that differ only by a single arrow, either added, removed or reversed. In these cases, the Bayes factor is especially simple, because of decomposability of the network score.

Greedy search works as follows.

1. Select an initial DAG $D_{0}$, from which to start the search.
2. Calculate Bayes factors between $D_{0}$ and all possible networks, which differ by only one arrow, that is
(a) One arrow is added to $D_{0}$.
(b) One arrow in $D_{0}$ is deleted.
(c) One arrow in $D_{0}$ is turned.
3. Among all these networks, select the one that increases the Bayes factor the most.
4. If the Bayes factor is not increased, stop the search. Otherwise, let the chosen network be $D_{0}$ and repeat from 2 .

In deal
rats.s <- getnetwork(autosearch(rats, rats.df,rats.prior))
returns all tried networks in a greedy search from the initial network rats, which may be constructed using drawnetwork ().


Figure 3: The network with the highest score in the rats example.

To manually assess the network score of a network (e.g. to use as initial network in a search), use

```
rats <- getnetwork(drawnetwork(rats,rats.df,rats.prior))
```

In the drawnetwork () procedure, it is possible to mark (ban) some of the arrows. In the search, deal then disregards any DAG which contains any of these arrows, and this reduces the search space.

The search algorithm may also be used with restarts which is implemented in the function heuristic (). The initial network is then perturbed according to the parameter degree and the search is performed starting with the perturbed network. The process is restarted the number of times specified by the option restart. A network family of all visited networks is returned.

```
rats.h <- getnetwork(heuristic(rats,rats.df,rats.prior,
    restart=10, degree=5))
```

The perturbation of the initial network is done as follows

1. Randomly choose between one of three actions
(a) Insert an arrow.
(b) Delete an arrow.
(c) Turn an arrow.
2. After selection of the action, perform the action according to Insert Choose randomly between all possible insertions of one arrow.
Delete Choose randomly between all possible deletions of one arrow.
Turn Choose randomly between all possible turns of one arrow.
If the action is not possible, return the unchanged network.
Perturbation is done automatically in heuristic() by calling the function perturb (). However, a random graph may also be generated by directly calling perturb()
```
rats.rn <- getnetwork(perturb(rats,rats.df,rats.prior,
    degree=10))
```


### 6.3 Using Equivalence Relations to Speed up Model Search

In Bøttcher (2003) two types of equivalences are identified and it is shown that no other equivalences exist. Let $D_{1}$ and $D_{1}^{*}$ be two different networks that differ by a single arrow between the nodes $v$ and $w$, with $v \leftarrow w$ in $D_{1}$ and $v \leftarrow w$ in $D_{1}^{*}$. Further, let $D_{2}$ and $D_{2}^{*}$ be another two networks that differ by an arrow between $v$ and $w$.

1. The Bayes factor for testing the arrow from $v$ to $w$ is equivalent to testing this arrow in any other network, where $v$ has the same parents as in $D_{1}$.
2. The Bayes factor for testing the arrow from $v$ to $w$ is equivalent to this arrow in the network, where $w$ has the same parents in $D_{2}$ as $v$ has in $D_{1}$, with the exception that $v$ is also a parent of $w$ in $D_{2}$.

We use the first equivalence in all functions that call the learning procedure, including heuristic(), learn(), drawnetwork(), networkfamily() and perturb (), by maintaining a so-called trylist. The trylist may be given as input to the functions and is returned in an updated version.

The trylist contains a list for each node in the network. The list for a node consists of the result after learning the node for all parent configurations that has previously been tried. When a node is learned after a change in its parent structure, we first look in the trylist to see if the node has been learned before with the same parent configuration (Equivalence 1). If the equivalence cannot be used, the node is learned and the result is inserted in the trylist. Utilization of Equivalence 2 is not yet implemented in deal.

The cost of looking in the trylist is smaller than learning a node. Note, however, that the trylist must be recalculated if the imaginary data base size is changed or if the data base is changed.

In deal, there is support for generating the complete trylist, that is, all nodes are learned with all possible parent configurations.

```
rats.tl <- maketrylist(rats,rats.df,rats.prior)
rats.h <- getnetwork(heuristic(rats,rats.df,rats.prior,
    trylist=rats.tl))
```


## 7 Hugin Interface

A network object may be written to a file in the Hugin . net language. Hugin (http://www.hugin.com) is commercial software for inference in Bayesian networks. Hugin has the ability to learn networks with only discrete variables, but cannot learn either purely continuous or mixed networks. deal may therefore be used for this purpose and the result can then be transferred to Hugin.

The procedure savenet () saves a network to a connection (for example a file).

For each node, we use point estimates of the parameters in the local probability distributions.

The readnet () procedure reads the network structure from a connection but does not, however, read the probability distributions. This is planned to be included in a future version of deal.

## 8 Example

In this section, deal is used to analyze a large data set which includes both discrete and continuous variables. The $k s l$ data set, included in Badsberg (1995), is from a study measuring health and social characteristics of representative samples of Danish 70-year old people, taken in 1967 and 1984. In total, 1083 cases have been recorded and each case contains observations on nine different variables, see Table 3.

| Node index | Variable | Explanation |
| :---: | :--- | :--- |
| 1 | Fev | Forced ejection volume - lung function |
| 2 | Kol | Cholesterol |
| 3 | Hyp | Hypertension (no/yes) |
| 4 | BMI | Body Mass Index |
| 5 | Smok | Smoking (no/yes) |
| 6 | Alc | Alcohol consumption (seldom/frequently) |
| 7 | Work | Working (yes/no) |
| 8 | Sex | Gender (male/female) |
| 9 | Year | Survey year (1967/1984) |

Table 3: Variables in the ksl data set. The variables Fev, Kol, BMI are continuous variables and the rest are discrete variables.

The purpose of our analysis is to find dependency relations between the variables. One interest is to determine which variables influence the presence or absence of hypertension. From a medical viewpoint, it is possible that hypertension is influenced by some of the continuous variables Fev, Kol and BMI. However, in deal we do not allow continuous parents of discrete nodes, so we cannot describe such a relation. A way to overcome this problem is to treat Hyp as a continuous variable, even though this is obviously not most natural. This is done in the analysis below.

Further, the initial data analysis indicates a transformation of BMI into $\log$ (BMI). With these adjustments, the data set is ready for analysis in deal.

First, deal is activated and the data are read into a data frame and prepared for analysis.

```
library(deal) ## invoke DEAL
data(ksl) ## read data (included in DEAL)
```

The next step in the analysis is to specify a prior Bayesian network. We have no prior knowledge about specific dependency relations, so for simplicity we use the empty DAG as the prior DAG and let the probability distribution of the discrete variables be uniform. The assessment of the probability distribution for the continuous variables is based on data, as described in Section 4.2.

```
## specify prior network
ksl.nw <- network(ksl)
## make joint prior distribution
ksl.prior <- jointprior(ksl.nw)
```

We do not allow arrows into Sex and Year, as none of the other variables can influence these variables. So we create a ban list which is attached to the network. The ban list is a matrix with two columns. Each row contains the directed edge that is not allowed. The ban list could also have been created interactively using the function drawnetwork ().

```
## ban arrows towards Sex and Year
mybanlist <- matrix(c(5,5,6,6,7,7,9,
    8,9,8,9,8,9,8),ncol=2)
banlist(ksl.nw) <- mybanlist
```

Finally, the parameters in the network are learned and structural learning is initiated using autosearch() and heuristic(). We use the prior DAG as starting point for the structural search.

```
## learn the initial network
ksl.nw <- getnetwork(learn(ksl.nw,ksl,ksl.prior))
## Do structural search
ksl.search <- autosearch(ksl.nw,ksl,ksl.prior,trace=TRUE)
## perturb 'thebest' and rerun search twice.
ksl.heuristic <- heuristic(getnetwork(ksl.search),
    ksl,
    ksl.prior,
    restart=2, degree=10,
    trace=TRUE,
    trylist=gettrylist(ksl.search))
thebest2 <- getnetwork(ksl.heuristic)
savenet(thebest2, file("ksl.net"))
```



Figure 4: The network with the highest score, $\log ($ score $)=-15957.91$.
The resulting network thebest 2 is shown in Figure 4 and it is the network with the highest network score among those networks that have been tried through the search.

In the result we see for the discrete variables that Alc, Smok and Work depend directly on Sex and Year. In addition, Smok and Work also depend on Alc.

These two arrows are, however, not causal arrows, as Smok $\leftarrow$ Alc $\rightarrow$ Work in the given DAG represents the same probability distribution as the relations Smok $\leftarrow$ Alc $\leftarrow$ Work and Smok $\rightarrow$ Alc $\rightarrow$ Work, i.e. the three DAGs are independence equivalent.

Year and Sex are independent on all variables, as specified in the ban list.
For the continuous variables all the arrows are causal arrows. We see that Fev depends directly on Year, Sex and Smok. So given these variables, Fev is conditionally independent on the rest of the variables. Kol depends directly on Year and Sex, and logBMI depends directly on Kol and Sex.

Given logBMI and Fev, the variable Hyp is conditionally independent on the rest of the variables. So according to this study, hypertension can be determined by the body mass index and the lung function forced ejection volume. However, as Hyp is not continuous by nature, other analyses should be performed with Hyp as a discrete variable, e.g. a logistic regression with Hyp as a response and the remaining as explanatory variables. Such an analysis indicates that, in addition, Sex and Smok may influence Hyp, but otherwise identifies logBMI as the main predictor.

## 9 Discussion and Future Work

deal is a tool box that adds functionality to $R$ so that Bayesian networks may be used in conjunction with other statistical methods available in R for analyzing data. In particular, deal is part of the $g R$ project, which is a newly initiated workgroup with the aim of developing procedures in R for supporting data analysis with graphical models, see http://www.r-project.org/gR.

In addition to methods for analyzing networks with either discrete or continuous variables, deal handles networks with mixed variables.
deal has some limitations and we plan to extend the package with the procedures described below. Also, it is the intention that the procedures in deal will eventually be adjusted to the other procedures developed under the gR project.

The methods in deal are only applicable on complete data sets and in the future, we would like to incorporate procedures for handling data with missing values and networks with latent variables.

The criteria for comparing the different network structures in deal, is the relative probability $S(D)$. We intend to also incorporate the Bayesian Information Criteria (BIC) and Akaikes Information Criteria (AIC) and let it be up to the user to decide which criteria to use.

Another possible extension of deal is to incorporate procedures for specifying mixed networks, where the variance in the mixed part of the network does not depend on the discrete parents, but the mean does.

Finally, we are working on an implementation of the greedy equivalence search (GES) algorithm, see Chickering (2002), which is an algorithm for search between equivalence classes. Asymptotically, for the size of the database tending to infinity, this algorithm guarantees that the search terminates with the network with the highest network score.

## Acknowledgements

The work has been supported by Novo Nordisk A/S. We thank the reviewer for many helpful comments.

## References

Badsberg, J. H. (1995). An Environment for Graphical Models, PhD thesis, Aalborg University.
Bøttcher, S. G. (2001). Learning Bayesian Networks with Mixed Variables, Artificial Intelligence and Statistics 2001, Morgan Kaufmann, San Francisco, CA, USA, pp. 149-156.
Bøttcher, S. G. (2003). Learning Bayesian networks with mixed variables, Aalborg University, http: / /www.math. auc.dk/~alma.
Chickering, D. M. (1996). Learning Bayesian networks is NP-Complete, in D. Fisher and H. J. Lenz (eds), Learning from Data: Artificial Intelligence and Statistics $V$, Springer-Verlag, New York, pp. 121-130.
Chickering, D. M. (2002). Optimal structure identification with greedy search, Journal of Machine Learning Research 3: 507-554.
Cowell, R. G., Dawid, A. P., Lauritzen, S. L. and Spiegelhalter, D. J. (1999). Probabilistic Networks and Expert Systems, Springer-Verlag, Berlin-Heidelberg-New York.

Edwards, D. (1995). Introduction to Graphical Modelling, Springer-Verlag, New York.

Geiger, D. and Heckerman, D. (1994). Learning Gaussian Networks, Proceedings of Tenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 235-243.
Heckerman, D., Geiger, D. and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data, Machine Learning 20: 197-243.
Lauritzen, S. L. (1992). Propagation of probabilities, means and variances in mixed graphical association models, Journal of the American Statistical Association 87(420): 1098-1108.
Lauritzen, S. L. (2003). Some modern applications of graphical models, in P. J. Green, N. L. Hjort and S. Richardson (eds), Highly Structured Stochastic Systems, Oxford University Press, Oxford, pp. 13-32.
Morrison, D. F. (1976). Multivariate Statistical Methods, McGraw-Hill, USA.
R Development Core Team (2003). R: A language and environment for statistical computing, R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-00-3.
Shachter, R. D. and Kenley, C. R. (1989). Gaussian influence diagrams, Management Science 35: 527-550.

## 10 Manual Pages for deal

## Description

From initial network, does local perturbations to increase network score.

## Usage

```
autosearch(initnw, data,prior=jointprior(network(data)),maxiter=50,
    trylist= vector("list",size(initnw)),trace=TRUE,
    timetrace=TRUE,showban=FALSE,removecycles=FALSE)
heuristic(initnw,data, prior=jointprior(network(data)),
    maxiter=100,restart=10,degree=size(initnw),
```

```
    trylist= vector("list",size(initnw)),trace=TRUE,
    timetrace=TRUE,removecycles=FALSE)
gettable(x)
```


## Arguments

| initnw | an object of class network，from which the search is started． |
| :---: | :---: |
| data | a data frame used for learning the network，see network． |
| prior | a list containing parameter priors，generated by jointprior． |
| maxiter | an integer，which gives the maximum number of steps in the search algorithm． |
| restart | an integer，which gives the number of times to perturb initnw and rerun the search． |
| degree | an integer，which gives the degree of perturbation，see perturb． |
| trylist | a list used internally for reusing learning of nodes，see maketrylist． |
| trace | a logical．If TRUE，plots the accepted networks during search． |
| timetrace | a logical．If TRUE，prints some timing information on the screen． |
| showban | a logical passed to the plot method for network objects．If FALSE，the banned arrows are not shown in the plots（if trace is TRUE）． |
| removecycles | a logical．If TRUE，all networks explored in the search is returned，except for networks containing a cycle．If FALSE，all networks are returned，including cyclic networks． |
| x | an output object from a search． |

## Details

In autosearch，a list of networks is in each step created with either one arrow added，one arrow deleted or one arrow turned（if a cycle is not generated）．The network scores of all the proposal networks are calculated and the network with the highest score is chosen for the next step in the search．If no proposed network has a higher network score than the previous network，the search is terminated．The network with the highest network score is returned，along with a list containing all tried networks （depending on the value of removecycles）．
heuristic restarts by perturbing initnw degree times and calling autosearch again．The number of restarts is given by the option restart．

## Value

autosearch and heuristic returns a list with three elements，that may be accessed using the functions getnetwork，gettable and gettrylist．The elements are
nw an object of class network，which gives the network with the highest score．
table a table with all tried networks．If removecycles is FALSE，the networks may contain cycles．The table contains two columns：model with a string repre－ sentation of the model and score with the corresponding log network score． The table can be translated to a networkfamily using makenw．
trylist an updated list used internally for reusing learning of nodes，see maketrylist．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## See Also

```
perturb
```


## Examples

```
data(rats)
fit <- network(rats)
fit.prior <- jointprior(fit,12)
fit <- getnetwork(learn(fit,rats,fit.prior))
fit <- getnetwork(insert(fit,2,1,rats,fit.prior))
fit <- getnetwork(insert(fit,1,3,rats,fit.prior))
hisc <- autosearch(fit,rats,fit.prior,trace=FALSE)
hisc <- autosearch(fit,rats,fit.prior,trace=FALSE,
    removecycles=TRUE) # slower
plot(getnetwork(hisc))
hisc2 <- heuristic(fit,rats,fit.prior,restart=10,trace=FALSE)
plot(getnetwork(hisc2))
print(modelstring(getnetwork(hisc2)))
plot (makenw(gettable(hisc2),fit))
```


## Description

drawnetwork allows the user to specify a Bayesian network through a point and click interface.

## Usage

```
drawnetwork(nw,df,prior,trylist=vector("list",size(nw)),
    unitscale=20,cexscale=8,
    arrowlength=.25, nocalc=FALSE,
    yr=c(0,350),xr=yr,...)
```


## Arguments

| nw | an object of class network to be edited. |
| :--- | :--- |
| df | a data frame used for learning the network, see network. |
| prior | a list containing parameter priors, generated by jointprior. |
| trylist | a list used internally for reusing learning of nodes, see maketrylist. <br> cexscale |
| arrowlength numeric passed to the plot method for network objects. Measures the scaled |  |
| size of text and symbols. |  |$\quad$| a numeric passed to the plot method for network objects. Measures the length |
| :--- |
| of the edges of the arrowheads. |


| nocalc | a logical．If TRUE，no learning procedure is called，see eg．rnetwork． |
| :--- | :--- |
| unitscale | a numeric passed to the plot method for network objects．Scale parameter for <br> chopping off arrow heads． |
| xr | a numeric vector with two components containing the range on x－axis． |
| yr | a numeric vector with two components containing the range on y－axis． |
| $\ldots$ | additional plot arguments，passed to the plot method for network objects． |

## Details

To insert an arrow from node＇ A ＇to node＇ B ＇，first click node＇ A ＇and then click node＇ B ＇．When the graph is finished，click＇stop＇．
To specify that an arrow must not be present，press＇ban＇（a toggle）and draw the arrow．This is shown as a red dashed arrow．It is possible to ban both directions between nodes．The ban list is stored with the network in the property banlist．It is a matrix with two columns．Each row is the＇from＇node index and the＇to＇node index，where the indices are the column number in the data frame．
Note that the network score changes as the network is re－learned whenever a change is made（unless nocalc is TRUE）．

## Value

A list with two elements that may be accessed using getnetwork and gettrylist．The elements are
nw an object of class network with the final network．
trylist
an updated list used internally for reusing learning of nodes，see maket rylist．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## See Also

```
network
```


## Examples

```
data(rats)
rats.nw <- network(rats)
rats.prior <- jointprior(rats.nw,12)
rats.nw <- getnetwork(learn(rats.nw,rats,
    rats.prior))
## Don't run: newrat <- getnetwork(drawnetwork(rats.nw,rats,
    rats.prior))
```


## Description

Given a network with a prob property for each node, derives the joint probability distribution. Then the quantities needed in the local master procedure for finding the local parameter priors are deduced.

## Usage

jointprior(nw, N=NA, phiprior="bottcher",timetrace=FALSE)

## Arguments

| nw | an object of class network. Each node must have a prob property to describe <br> the local probability distribution. The prob property is created using prob <br> method for network objects, which is called by the network function. |
| :--- | :--- |
| an integer, which gives the size of the imaginary data base. If this is too small, |  |
| NA's may be created in the output, resulting in errors in learn. If no N is |  |
| given, the procedure tries to set a value as low as possible. |  |

## Details

For the discrete part of the network, the joint probability distribution is calculated by multiplying together the local probability distributions. Then, jointalpha is determined by multiplying each entry in the joint probability distribution by the size of the imaginary data base N .
For the mixed part of the network, for each configuration of the discrete variables, the joint Gaussian distribution of the continuous variables is constructed and represented by jointmu (one row for each configuration of the discrete parents) and joint sigma (a list of matrices - one for each configuration of the discrete parents). The configurations of the discrete parents are ordered according to findex. The algorithm for constructing the joint distribution of the continuous variables is described in Shachter and Kenley (1989).
Then, jointalpha, jointnu, jointrho, mu and jointphi are deduced. These quantities are later used for deriving local parameter priors.
For each configuration $i$ of the discrete variables,

$$
\nu_{i}=\rho_{i}=\alpha_{i}
$$

and

$$
\phi_{i}=\left(\nu_{i}-1\right) \Sigma_{i}
$$

if phiprior="bottcher", see Bøttcher(2001) and

$$
\phi_{i}=\nu_{i}\left(\rho_{i}-2\right) \Sigma_{i} /\left(\nu_{i}+1\right)
$$

if phiprior="heckerman", see Heckerman, Geiger and Chickering (1995).

## Value

A list with the following elements，

| jointalpha | a table used in the local master procedure for discrete variables． |
| :--- | :--- |
| jointnu | a table used in the local master procedure for continuous variables． |
| jointrho | a table used in the local master procedure for continuous variables． |
| jointmu | a numeric matrix used in the local master procedure for continuous variables． <br> jointsigma |
| a list of numeric matrices（not used in further calculations）． |  |
| jointphi | a list of numeric matrices used in the local master procedure for continuous <br> variables． |

## Author（s）

Susanne Gammelgaard Bøttcher $\langle$ alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## References

Bøttcher，S．G．（2001）．Learning Bayesian Networks with Mixed Variables，Artificial Intelligence and Statistics 2001，Morgan Kaufmann，San Francisco，CA，USA，pp．149－156．
Heckerman，D．，Geiger，D．and Chickering，D．（1995）．Learning Bayesian networks：The combination of knowledge and statistical data，Machine Learning 20：197－243．
Shachter，R．D．and Kenley，C．R．（1989）．Gaussian influence diagrams，Management Science 35：527－ 550.

## See Also

```
network, prob
```


## Examples

```
data(rats)
rats.nw <- network(rats)
rats.prior <- jointprior(rats.nw,12)
## Don't run: savenet(rats.nw,file("rats.net"))
## Don't run: rats.nw <- readnet(file("rats.net"))
## Don't run: rats.nw <- prob(rats.nw,rats)
## Don't run: rats.prior <- jointprior(rats.nw,12)
```


## Description

Updates the distributions of the parameters in the network, based on a prior network and data. Also, the network score is calculated.

## Usage

```
learn (nw, df, prior=jointprior(nw),
    nodelist=1:size(nw),
    trylist=vector("list",size(nw)),
    timetrace=FALSE)
```


## Arguments

| nw | an object of class network. |
| :--- | :--- |
| df | a data frame used for learning the network, see network. |
| prior | a list containing parameter priors, generated by jointprior. |
| nodelist | a numeric vector of indices of nodes to be learned. |
| trylist | a list used internally for reusing learning of nodes, see maketrylist. |
| timetrace | a logical. If TRUE, prints some timing information on the screen. |

## Details

The procedure learn determines the master prior, local parameter priors and local parameter posteriors, see Bøttcher (2001). It may be called on all nodes (default) or just a single node.
From the joint prior distribution, the marginal distribution of all parameters in the family consisting of the node and its parents can be determined. This is the master prior, see localmaster. The local parameter priors are now determined by conditioning in the master prior distribution, see conditional. The hyperparameters associated with the local parameter prior distribution is attached to each node in the property condprior.
Finally, the local parameter posterior distributions are calculated (see post) and attached to each node in the property condposterior.
A so-called trylist is maintained to speedup the learning process. The trylist consists of a list of matrices for each node. The matrix for a given node holds previously evaluated parent configurations and the corresponding log-likelihood contribution. If a node with a certain parent configuration needs to be learned, it is checked, whether the node has already been learned. The previously learned nodes are given as input in the trylist parameter and is updated in the learning procedure.
When one or more nodes in a network have been learned, the network score is updated and attached to the network in the property score.
The learning procedure is called from various functions using the principle, that networks should always be updated with their score. Thus, e.g. drawnetwork keeps the network updated when the graph is altered.

## Value

A list with two elements that may be accessed using getnetwork and gettrylist. The elements are
nw an object of class network，with the condposterior properties updated for the nodes．Also，the property score is updated and contains the network score．The contribution to the network score for each node is contained in the property loglik for each node．
trylist an updated list used internally for reusing learning of nodes，see maketrylist．

## Author（s）

Susanne Gammelgaard Bøttcher $\langle$ alma＠math．auc．dk〉，
Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## References

Bøttcher，S．G．（2001）．Learning Bayesian Networks with Mixed Variables，Artificial Intelligence and Statistics 2001，Morgan Kaufmann，San Francisco，CA，USA，pp．149－156．

## See Also

networkfamily，jointprior，maketrylist，network

## Examples

```
data(rats)
fit <- network(rats)
fit.prior <- jointprior(fit,12)
fit.learn <- learn(fit,rats,fit.prior,timetrace=TRUE)
fit.nw <- getnetwork(fit.learn)
fit.learn2<- learn(fit,rats,fit.prior,trylist=gettrylist(fit.learn),
    timetrace=TRUE)
```

maketrylist Creates the full trylist

## Description

For faster learning，a trylist is maintained as a lookup table for a given parent configuration of a node．

## Usage

```
maketrylist(initnw,data, prior=jointprior(network(data)),
    timetrace=FALSE)
```


## Arguments

| initnw | an object of class network，from which the search is started． |
| :--- | :--- |
| data | a data frame used for learning the network，see network． |
| prior | a list containing parameter priors，generated by jointprior． |
| timetrace | a logical．If TRUE，prints some timing information on the screen． |

## Details

This procedure is included for illustrative purposes．For each node in the network，all possible parent configurations are created and learned．The result is called a trylist．To create the full trylist is very time－ consuming，and a better choice is to maintain a trylist while searching and indeed this is automatically done．The trylist is given as output to all functions that call the learning procedure and can be given as an argument．

## Value

A list with one element per node in the network．In the list，element $i$ is a matrix with two columns： a string with the indices of the parent nodes，separated by＂：＂，and a numeric with the log－likelihood contribution of the node given the parent configuration．Whenever learning is performed of a node given a parent configuration，the trylist is consulted to yield faster learning，especially useful when using autosearch or heuristic．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## See Also

networkfamily，autosearch heuristic

## Examples

```
data(rats)
rats.nw <- network(rats)
rats.pr <- jointprior(rats.nw,12)
rats.nw <- getnetwork(learn(rats.nw,rats,rats.pr))
rats.tr <- maketrylist(rats.nw, rats,rats.pr)
rats.hi <- getnetwork(heuristic(rats.nw,rats,rats.pr,
    trylist=rats.tr))
```


## Description

A Bayesian network is represented as an object of class network. Methods for printing and plotting are defined.

## Usage

```
network(df,specifygraph=FALSE,inspectprob=FALSE,
    doprob=TRUE,yr=c (0,350), xr=yr)
## S3 method for class 'network':
print(x,filename=NA,condposterior=FALSE,
                                    condprior=FALSE,...)
## S3 method for class 'network':
plot (x,arrowlength=.25,
    notext=FALSE,
    sscale=7,showban=TRUE, yr=c (0,350),xr=yr,
    unitscale=20,cexscale=8,...)
```


## Arguments

| df | a data frame, where the columns define the variables. A continuous variable <br> should have type numeric and discrete varibles should have type factor. |
| :--- | :--- |
| specifygraph | a logical. If TRUE, provides a call to drawnetwork to interactively specify a <br> directed acyclic graph and possibly a ban list (see below). |
| inspectprob | a logical. If TRUE, provides a plot of the graph and possibility to inspect the <br> calculated probability distribution by clicking on the nodes. <br> a logical. If TRUE, do not calculate a probability distribution. Used for example <br> in rnetwork. <br> an object of class network. |
| doprob | a string or NA. If not NA, output is printed to a file. |
| x a logical. If TRUE, the conditional prior is printed, see conditional. |  |
| filename | a logical. If TRUE, the conditional posterior is printed, see learn. |
| condposterior |  |
| sscale | a numeric. The nodes are initially placed on a circle with radius sscale. <br> anitscale numeric. Scale parameter for chopping off arrow heads. |
| cexscale | a numeric. Scale parameter to set the size of the nodes. |
| arrowlength | a numeric containing the length of the arrow heads. <br> a numeric vector with two components containing the range on x-axis. |
| xr | a numeric vector with two components containing the range on y-axis. |
| yr | a logical. If TRUE, no text is displayed in the nodes on the plot. |
| notext | a logical. If TRUE, banned arrows are shown in red. <br> additional plot arguments, passed to plot. node. |
| showban |  |

## Value

The netork creator function returns an object of class network，which is a list with the following elements（properties），

| nodes | a list of objects of class node．If doprob is TRUE，the nodes are given the <br> property prob which is the initial probability distribution used by <br> jointprior． <br> an integer containing the number of nodes in the network． |
| :--- | :--- |
| n | a numeric vector of indices of discrete nodes． |
| discrete | a numeric vector of indices of continuous nodes． |
| continuous | a numeric matrix with two columns．Each row contains the indices i <br> arrows that may not be allowed in the directed acyclic graph． |
| banlist of |  |
| score | a numeric added by learn and is the log network score． |
| relscore | a numeric added by nwfsort and is the relative network score－compared <br> with the best network in a network family． |

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen $\langle\operatorname{dethlef}$＠math．auc．dk〉．

## See Also

networkfamily，node，rnetwork，learn，drawnetwork，jointprior，heuristic， nwequal

## Examples

```
A <- factor(rep(c("A1","A2"),50))
B <- factor(rep (rep(c("B1","B2"), 25),2))
thisnet <- network( data.frame(A,B) )
set.seed(109)
sex <- gl(2,4,label=c("male","female"))
age <- gl (2,2,8)
yield <- rnorm(length(sex))
weight <- rnorm(length(sex))
mydata <- data.frame(sex,age,yield,weight)
mynw <- network(mydata)
# adjust prior probability distribution
localprob(mynw,"sex") <- c(0.4,0.6)
localprob(mynw,"age") <- c(0.6,0.4)
localprob(mynw,"yield") <- c(2,0)
localprob(mynw,"weight")<- c(1,0)
print (mynw)
plot (mynw)
prior <- jointprior(mynw)
```

mynw <- getnetwork(learn(mynw, mydata, prior))
thebest <- getnetwork (autosearch (mynw, mydata, prior))
print (mynw, condposterior=TRUE)
\#\# Don't run: savenet (mynw,file("yield.net"))
networkfamily Generates and learns all networks for a set of variables.

## Description

Method for generating and learning all networks that are possible for a given set of variables. These may be plotted or printed. Also, functions for sorting according to the network score (see nwfsort) and for making a network family unique (see the unique method for networkfamily objects) are available.

## Usage

```
networkfamily(data,nw=network(data), prior=jointprior(nw),
    trylist=vector("list",size(nw)), timetrace=TRUE)
## S3 method for class 'networkfamily':
print(x,...)
## S3 method for class 'networkfamily':
plot(x,layout=<<see below>>,
    cexscale=5,arrowlength=0.1,sscale=7,...)
```


## Arguments

| nw | an object of class network. This should be the empty network for the set of variables. |
| :---: | :---: |
| data | a data frame used for learning the network, see network. |
| prior | a list containing parameter priors, generated by jointprior. |
| trylist | a list used internally for reusing learning of nodes, see maketrylist. |
| timetrace | a logical. If TRUE, prints some timing information on the screen. |
| x | an object of class networkfamily. |
| layout | a numeric two dimensional vector with the number of plots in the rows and columns of each plotting page. <br> Default set to rep (min(1+floor (sqrt (length(x))),5), 2). |
| cexscale | a numeric. A scaling parameter to set the size of the nodes. |
| arrowlength | a numeric, which gives the length of the arrow heads. |
| sscale | a numeric. The nodes are initially placed on a circle with radius sscale. |
|  | additional plot arguments passed to the plot method for network objects. |

## Details

networkfamily generates and learns all possible networks with the nodes given as in the initial network nw．This is done by successively trying to generate the networks with all possible arrows to／from each node（see addarrows）．If there is a ban list present in nw（see network），then this is respected，as are the restrictions described in insert
After generation of all possible networks，a test for cycles（see cycletest）is performed and only networks with directed acyclic graphs are returned．

## Value

The function networkfamily returns a list with two components，
nw an object of class networkfamily．
trylist an updated list used internally for reusing learning of nodes，see maketrylist．

## Note

Generating all possible networks can be very time consuming！

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉，
Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## See Also

network，genlatex，heuristic，nwfsort，unique．networkfamily，elementin， addarrows，cycletest

## Examples

```
data(rats)
allrats <- getnetwork(networkfamily(rats))
plot(allrats)
print(allrats)
```


## Description

Various extraction／replacement functions for networks

## Usage

```
modelstring(x)
makenw(tb,template)
as.network(nwstring,template)
size(x)
banlist(x)
banlist(x) <- value
getnetwork(x)
gettrylist(x)
```


## Arguments

| x | an object of class network. |
| :--- | :--- |
| tb | a table output from autosearch or heuristic in the list property table. <br> Can be translated into a networkfamily. <br> an object of class network with the same nodes as the networks described in <br> the table tb. |
| template | a string representing the network. <br> a numeric matrix with two columns. Each row contains the indices i <br> arrows that may not be allowed in the directed acyclic graph. |
| nwstring of |  |
| value |  |

## Details

The string representation of a network is a minimal size representation to speed up calculations. The functions modelstring, as.network and makenw converts between the string represention and network objects.
size extracts the number of nodes in a network object
banlist extracts the banlist from a network object.
get network and gettrylist are accessor function that extracts a network object or trylist from the result from autosearch, heuristic, learn, perturb, networkfamily, drawnetwork.
node Representation of nodes

## Description

An important part of a network is the list of nodes. The nodes summarize the local properties of a node, given the parents of the node.

Usage

```
node (idx,parents,type="discrete",name=paste(idx),
    levels=2,levelnames=paste(1:levels),position=c(0,0))
## S3 method for class 'node':
print (x,filename=NA,condposterior=TRUE,condprior=TRUE,...)
```

```
## S3 method for class 'node':
plot (x,cexscale=10, notext=FALSE,...)
nodes(nw)
value <- nodes(nw)
```


## Arguments

\(\left.$$
\begin{array}{ll}\mathrm{x} \\
\text { parents } \\
\text { idx } & \begin{array}{l}\text { an object of class node. } \\
\text { a numeric vector with indices of the parents of the node. } \\
\text { an integer, which gives the index of the node (the column number of the corre- } \\
\text { sponding data frame). }\end{array}
$$ <br>
type <br>
a string, which gives the type of the node. Either "discrete" (for factors) <br>
or "cont inuous" (for numeric). <br>
a string, which gives the name used when plotting and printing. Defaults to the <br>
column name in the data frame. <br>
an integer. If type is "discrete", this is the number of levels for the dis- <br>
crete variable. <br>

if type is "discrete", this is a vector of strings (same length as levels)\end{array}\right]\)| with the names of the levels. If type is "cont inuous", the argument is |
| :--- |
| ignored. |

## Details

The operations on a node are typically done when operating on a network, so these functions are not to be called directly.
When a network is created with network, the nodes in the nodelist are created using the node procedure.
Local probability distributions are added as the property prob to each node using prob. node. If the node is continuous, this is a numeric vector with the conditional variance and the conditional regression coefficients arising from a regression on the continuous parents, using data. If the node has discrete parents, prob is a matrix with a row for each configuration of the discrete parents. If the node is discrete, prob is a multiway array which gives the conditional probability distribution for each configuration of the discrete parents. The generated prob can be replaced to match the prior information available. nodes gives the list of nodes of a network. localprob gives the probability distribution for each node in the network.

## Value

The node creator function returns an object of class node, which is a list with the following elements (properties),
\(\left.$$
\begin{array}{ll}\text { idx } & \begin{array}{l}\text { an integer．A unique index for this node．It MUST correspond to the column } \\
\text { index of the variable in the data frame．}\end{array}
$$ <br>

a string．The printed name of the node．\end{array}\right]\)| a string．Either＂continuous＂or＂discrete＂． |
| :--- |
| an integer．If the node is of type＂discrete＂，this integer is the number of |
| levels of the node． |
| levels |
| levpe is＂discrete＂，this is a vector of strings（same length as levels） |
| with the names of the levels．If type is＂cont inuous＂，the node does not |
| have this property． |

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉，
Claus Dethlefsen $\langle$ dethlef＠math．auc．dk〉．
numbermixed The number of possible networks

## Description

Calculates the number of different directed acyclic graphs for a set of discrete and continuous nodes．

## Usage

numbermixed（nd，nc）

## Arguments

$\begin{array}{ll}\text { nd } & \text { an integer，which gives the number of discrete nodes．} \\ \text { nc } & \text { an integer，which gives the number of continuous nodes }\end{array}$

## Details

No arrows are allowed from continuous nodes to discrete nodes．Cycles are not allowed．The number of networks is given by Bøttcher（2003），using the result in Robinson（1977）．
When $n d+n c>15$ ，the procedure is quite slow．

## Value

A numeric containing the number of directed acyclic graphs with the given node configuration．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## References

Bøttcher，S．G．（2003）．Learning Conditional Gaussian Networks．
http：／／www．math．auc．dk／～alma．Aalborg University， 2003.
Robinson，R．W．（1977）．Counting unlabeled acyclic digraphs，Lecture Notes in Mathematics，622： Combinatorial Mathematics V pp．239－273．

## Examples

numbermixed $(2,2)$
\＃\＃Don＇t run：numbermixed $(5,10)$
nwfsort Sorts a list of networks

## Description

According to the score property of the networks in a network family，the networks are sorted and the relative score，i．e．the score of a network relative to the highest score，is attached to each network as the relscore property．

## Usage

nwfsort（nwf）

## Arguments

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．
perturb Perturbs a network

## Description

Randomly insert／delete／turn arrows to obtain another network．

## Usage

```
perturb(nw, data,prior, degree=size(nw),trylist=vector("list",size(nw)),
    nocalc=FALSE,timetrace=TRUE)
```


## Arguments

| nw | an object of class network，from which arrows are added／removed／turned． |
| :--- | :--- |
| data | a data frame used for learning the network，see network． |
| prior | a list containing parameter priors，generated by jointprior． |
| degree | an integer，which gives the number of attempts to randomly insert／remove／turn <br> an arrow． |
| trylist | a list used internally for reusing learning of nodes，see maketrylist． |
| nocalc | a logical．If TRUE no learning procedure is called，see eg．rnetwork． |
| timetrace | a logical．If TRUE，prints some timing information on the screen． |

## Details

Given the initial network，a new network is constructed by randomly choosing an action：remove，turn， add．After the action is chosen，we choose randomly among all possibilities of that action．If there are no possibilites，the unchanged network is returned．

## Value

A list with two elements that may be accessed using getnetwork and gettrylist．The elements are
nw an object of class network with the generated network．
trylist
an updated list used internally for reusing learning of nodes，see maketrylist．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## Examples

```
set.seed(200)
data(rats)
fit
fit.prior <- jointprior(fit)
fit <- getnetwork(learn(fit,rats,fit.prior))
fit.new <- getnetwork(perturb(fit,rats,fit.prior,degree=10))
data(ksl)
ksl.nw <- network(ksl)
ksl.rand <- getnetwork(perturb(ksl.nw,nocalc=TRUE,degree=10))
plot(ksl.rand)
```

prob Local probability distributions

## Description

Methods for accessing or changing the local probability distributions and for accessing the local prior and posterior distributions

## Usage

```
prob(x,df,...)
## S3 method for class 'node':
prob (x,df,nw,...)
## S3 method for class 'network':
prob (x,df,...)
localprob(nw)
value <- localprob(nw,name)
localprior(node)
localposterior(node)
```


## Arguments

| df | a data frame，where the columns define the variables．A continuous variable <br> should have type numeric and discrete varibles should have type factor． |
| :--- | :--- |
| nw | an object of class network． |
| node | an object of class node． |
| name | a string，which gives the node name． |
| $\ldots$ | additional arguments for specific methods． |

## Details

The prob methods add local probability distributions to each node．If the node is continuous，this is a numeric vector with the conditional variance and the conditional regression coefficients arising from a regression on the continuous parents，using data．If the node has discrete parents，prob is a matrix with a row for each configuration of the discrete parents．If the node is discrete，prob is a multiway array which gives the conditional probability distribution for each configuration of the discrete parents． The generated prob can be replaced to match the prior information available．
localprob returns the probability distribution for each node in the network．
In a learned network，the local prior and posterior can be accessed for each node using localprior and localposterior．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉，
Claus Dethlefsen 〈dethlef＠math．auc．dk〉．
readnet Reads／saves ．net file

## Description

Reads／saves a Bayesian network specification in the ．net language（see http：／／developer． hugin．com／documentation／net／）．

## Usage

```
readnet(con=file("default.net"))
savenet(nw, con=file("default.net"))
```


## Arguments

| con | a connection． |
| :--- | :--- |
| nw | an object of class network． |

## Details

readnet reads only the structure of a network，i．e．the directed acyclic graph． savenet exports the prob property for each node in the network object along with the network structure defined by the parents of each node．

## Value

readnet creates an object of class network with the nodes specified as in the ．net connection． The network has not been learned and the nodes do not have prob properties（see prob．network）． savenet writes the object to the connection．

## Note

The call to readnet（savenet（network））is not the identity function as information is thrown away in both savenet and readnet．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## See Also

network

## Examples

```
data(rats)
nw <- network(rats)
## Don't run: savenet(nw,file("default.net"))
## Don't run: nw2 <- readnet(file("default.net"))
## Don't run: nw2 <- prob(nw2,rats)
```


## Description

## Usage

```
score(x,...)
## S3 method for class 'node':
score (x,...)
## S3 method for class 'network':
score (x,...)
```


## Arguments

```
x an object of class node or network.
... additional arguments for specific methods.
```


## Value

For networks，the $\log$ network score is returned．For nodes，the contribution to the log network score is returned．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．
rnetwork Simulation of data sets with a given dependency structure

## Description

Given a network with nodes having the simprob property，rnetwork simulates a data set．

## Usage

```
rnetwork(nw, n=24, file="")
```


## Arguments

| nw | an object of class network，where each node has the property simprob（see <br> makesimprob）． |
| :--- | :--- |
| n | an integer，which gives the number of cases to simulate． <br> a string．If non－empty，the data set is stored there． |

## Details

The variables are simulated one at a time in an order that ensures that the parents of the node have already been simulated．For discrete variables a multinomial distribution is used and for continuous variables，a Gaussian distribution is used，according to the simprob property in each node．

## Value

A data frame with one row per case．If a file name is given，a file is created with the data set．

## Author（s）

Susanne Gammelgaard Bøttcher 〈alma＠math．auc．dk〉， Claus Dethlefsen 〈dethlef＠math．auc．dk〉．

## Examples

```
A <- factor(NA, levels=paste("A",1:2, sep=""))
B <- factor(NA, levels=paste("B",1:3,sep=""))
c1 <- NA
c2 <- NA
df <- data.frame(A,B,C1,c2)
nw <- network(df,doprob=FALSE) # doprob must be FALSE
nw <- makesimprob(nw) # create simprob properties
set.seed(944)
sim <- rnetwork(nw, n=100) # create simulated data frame
```


## Paper III

# Prediction of the Insulin Sensitivity Index using Bayesian Networks 

Susanne G. Bøttcher and Claus Dethlefsen

# Prediction of the Insulin Sensitivity Index using Bayesian Networks 

Susanne G. Bøttcher<br>Aalborg University, Denmark<br>Claus Dethlefsen<br>Aalborg University, Denmark


#### Abstract

. The insulin sensitivity index $\left(S_{I}\right)$ can be used in assessing the risk of developing type 2 diabetes. An intravenous study is used to determine $S_{I}$ using Bergmans minimal model. However, an intravenous study is time consuming and expensive and therefore not suitable for large scale epidemiological studies. In this paper we learn the parameters and structure of several Bayesian networks relating measurements from an oral glucose tolerance test to the insulin sensitivity index determined from an intravenous study on the same individuals. The networks can then be used in prediction of $S_{I}$ from an oral glucose tolerance test instead of an intravenous study. The methodology is applied to a dataset with 187 patients. We find that the $S_{I}$ values from this study are highly correlated to the $S_{I}$ values determined from the intravenous study.


## 1 Introduction

Type 2 diabetes is a clinical syndrome that can result from several disorders that interfere with insulin secretion and/or the ability of the target tissues to respond to insulin. Martin, Warram, Krolewski, Bergman, Soeldner and Kahn (1992) found evidence in a 25 year follow-up study that insulin sensitivity index $\left(S_{I}\right)$ can be used to predict the development of type 2 diabetes up to a decade before diagnosis. Assessment of $S_{I}$ is by Bergmans minimal model, see Bergman et al. (1979), which is based on data from an intravenous glucose tolerance test (IVGTT). In the minimal model, the glucose and insulin kinetics are separately described by two sets of differential equations. The parameters in the model are traditionally estimated by a non-linear weighted least squares estimation technique, see for example Pacini and Bergman (1986). From these parameters, $S_{I}$ can be determined.

However, an IVGTT is time consuming and expensive and therefore not suitable for large scale epidemiological studies. Interest is therefore in developing a
method to assess the insulin sensitivity index from an oral glucose tolerance test (OGTT).

In Drivsholm, Hansen, Urhammer, Palacios, Vølund, Borch-Johnsen and Pedersen (2003), multiple linear regression is used to derive predictive values of $S_{I}$ from measurements from an OGTT. These are compared with the values of $S_{I}$ obtained from an IVGTT and calculated using Bergmans minimal model. The results show that it is possible to predict estimates of $S_{I}$, which are highly correlated to IVGTT-derived $S_{I}$ for subjects with normal glucose tolerance.

In this paper, we express the relation between the observed variables in a Bayesian network. We try different approaches of establishing a Bayesian network, which can be used to predict $S_{I}$ from measurements from an OGTT. We learn the parameters and structure of a Bayesian network from a training data set, where all patients underwent both an IVGTT and an OGTT. Bergmans minimal model were used to determine $S_{I}$ from the IVGTT. We then calculate the predictive value of $S_{I}$ from the Bayesian network and compare it with the value of $S_{I}$ obtained from the IVGTT.

Like the multiple linear regression approach, the Bayesian network approach gives predictions of $S_{I}$ that are highly correlated to IVGTT-derived $S_{I}$ for subjects with normal glucose tolerance. In addition, the complex dependency structure between the variables is modeled adequately. Further, using Bayesian networks makes it possible to incorporate any prior information available, e.g. the physiological understanding of the problem or results from previous studies.

## 2 Data

In this paper we consider 187 non-diabetic glucose tolerant subjects, with one parent having diabetes. All the subjects underwent a 75 gram frequently sampled OGTT. In such a test, the subject drinks 75 gram fluent glucose, after a 12 hour overnight fast. Venous blood samples are then drawn at 10,5 and 0 minutes before the OGTT and after the start of the OGTT, at $10,20,30,40,50$, $60,75,90,105,120,140,160,180,210$ and 240 minutes. From these blood samples, the glucose and insulin concentrations are determined.

Within one week after the OGTT examination, all subjects underwent a tolbutamide modified frequently sampled IVGTT. In an IVGTT, glucose is injected directly into the venous. Blood samples are drawn at 10, 5 and 0 minutes before
the injection and frequently up until 180 minutes after the injection. At 20 minutes, a bolus of tolbutamide is injected to elicit secondary pancreatic beta cell response. In the time between the two examinations, the subjects were asked not to change their lifestyle. The insulin sensitivity index $\left(S_{I}\right)$ was for each subject calculated from the observations in the IVGTT using Bergmans minimal model and estimated by a non-linear weighted least squares estimation technique, as described Pacini and Bergman (1986).

Other variables in the study are age, sex, weight, height, waist circumference, hip circumference, fat mass and information on physical activity. From the weight and height, the body mass index (BMI) can be calculated.

## 3 Bayesian Networks

We perform the analysis using Bayesian networks for discrete and continuous variables in which the joint distribution of all the variables are conditional Gaussian (CG), see Lauritzen (1992).

### 3.1 Bayesian Networks with Mixed Variables

Let $D=(V, E)$ be a Directed Acyclic Graph (DAG), where $V$ is a finite set of nodes and $E$ is a finite set of directed edges (arrows) between the nodes. The DAG defines the structure of the Bayesian network. To each node $v \in V$ in the graph corresponds a random variable $X_{v}$. The set of variables associated with the graph $D$ is then $X=\left(X_{v}\right)_{v \in V}$. Often, we do not distinguish between a variable $X_{v}$ and the corresponding node $v$. To each node $v$ with parents $\mathrm{pa}(v)$, a local probability distribution, $p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)$ is attached. The set of local probability distributions for all variables in the network is $\mathcal{P}$. A Bayesian network for a set of random variables $X$ is then the pair $(D, \mathcal{P})$.

The possible lack of directed edges in $D$ encodes conditional independencies between the random variables $X$ through the factorization of the joint probability distribution,

$$
\begin{equation*}
p(x)=\prod_{v \in V} p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right) \tag{1}
\end{equation*}
$$

Here, we allow Bayesian networks with both discrete and continuous variables, as treated in Lauritzen (1992), so the set of nodes $V$ is given by $V=\Delta \cup \Gamma$,
where $\Delta$ and $\Gamma$ are the sets of discrete and continuous nodes, respectively. The set of variables $X$ can then be denoted $X=\left(X_{v}\right)_{v \in V}=(I, Y)=$ $\left(\left(I_{\delta}\right)_{\delta \in \Delta},\left(Y_{\gamma}\right)_{\gamma \in \Gamma}\right)$, where $I$ and $Y$ are the sets of discrete and continuous variables, respectively. For a discrete variable, $\delta$, we let $\mathcal{I}_{\delta}$ denote the set of levels.

To ensure availability of exact local computation methods, we do not allow discrete variables to have continuous parents. The joint probability distribution then factorizes into a discrete part and a mixed part, so

$$
p(x)=p(i, y)=\prod_{\delta \in \Delta} p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) \prod_{\gamma \in \Gamma} p\left(y_{\gamma} \mid y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}\right) .
$$

A method for estimating the parameters and learning the dependency structure of a conditional Gaussian networks with mixed variables is presented in Bøttcher (2001) and implemented in the software package deal, see Bøttcher and Dethlefsen (2003).

### 3.2 Parameter and Structure Learning

To estimate the parameters in the network and to find the structure of the network, we use a Bayesian approach. So, considering the parameters, we encode our uncertainty about $\theta$ in a prior distribution $p(\theta)$, use data $d$ to update this distribution, i.e. learn the parameters, and hereby obtain the posterior distribution $p(\theta \mid d)$ by using Bayes' theorem,

$$
\begin{equation*}
p(\theta \mid d)=\frac{p(d \mid \theta) p(\theta)}{p(d)}, \quad \theta \in \Theta . \tag{2}
\end{equation*}
$$

Here, $\Theta$ is the parameter space, $d$ is a random sample of size $n$ from the probability distribution $p(x \mid \theta)$ and $p(d \mid \theta)$ is the joint probability distribution of $d$, also called the likelihood of $\theta$. As prior parameter distributions we use the Dirichlet distribution for the discrete variables and the Gaussian inverse-Gamma distribution for the continuous variables. These distributions are conjugate to observations from the respective distributions and this ensures simple calculations of the posterior distributions.

Now, to learn the structure of the network, we calculate the posterior probability of the DAG, $p(D \mid d)$, which from Bayes' theorem is given by

$$
p(D \mid d)=\frac{p(d \mid D) p(D)}{p(d)},
$$

where $p(d \mid D)$ is the likelihood of $D$ and $p(D)$ is the prior probability of $D$. As the normalizing constant $p(d)$ does not depend upon structure, another measure, which gives the relative probability, is

$$
p(D, d)=p(d \mid D) p(D)
$$

We use the above measure and refer to it as the network score. For simplicity, we choose to let $p(D)$ be the same for all DAGs, so we are only interested in calculating the likelihood $p(d \mid D)$. It is given as

$$
p(d \mid D)=\int_{\theta \in \Theta} p(d \mid \theta, D) p(\theta \mid D) d \theta
$$

and we see that it, besides the likelihood of the parameters, also involves the prior distribution over the parameters, $p(\theta \mid D)$. This means that we for each possible DAG have to specify a prior distribution for the parameters. In the papers Heckerman et al. (1995) and Geiger and Heckerman (1994) an automated method for doing this in respectively the purely discrete and the purely Gaussian case is developed. In Bøttcher (2001) this method is extended to the mixed case. With this method, the parameter priors for all possible networks can be deduced from one joint parameter prior, called a master prior. To specify this master prior, we only have to specify a prior Bayesian network, i.e. a prior DAG and a prior probability distribution, together with a measure of how confident we are in the prior network. With a few assumptions, the network score is obtainable in closed form.

If many DAGs are possible, it is computational infeasible to calculate the network score for all DAGs. In this situation it is necessary to use some kind of search strategy to find the DAG with the highest score, see e.g. Cooper and Herskovits (1992). In this paper we use a search strategy called greedy search. In greedy search we compare DAGs that differ only by a single arrow, either added, removed or reversed. The change that increases the network score the most is selected and the search is continued from this new DAG.

## 4 Inference

Having established a Bayesian network for a set of random variables, this represents the knowledge we, at this stage, have about these variables. When information on some or all of the variables becomes available, we can use this
"knowledge base" to make inference about the unobserved variables in the network.

Inference in Bayesian networks is performed using Bayes' theorem. Consider a network for a set of random variables $X$ and assume that some of the variables, $B$, are observed and the rest, $A$, are not. We can then, by using Bayes' theorem, calculate the conditional distribution of $A$ given $B$ as

$$
p(A \mid B) \propto p(B \mid A) p(A)
$$

Thus $p(A)$ is the prior distribution of $A$, i.e. the distribution of $A$ before we observe $B, p(B \mid A)$ is the likelihood of $A$ and $p(A \mid B)$ is the posterior distribution of $A$, i.e. the distribution of $A$, when we have observed $B$. Generally, finding these distributions is computationally heavy as it involves calculating huge joint distributions, especially if there are many variables in the network. Therefore efficient methods of implementing Bayes' theorem are being used. These implementations uses the fact that the the joint probability distribution of all the variables in a network factorizes according to (1). The marginal or conditional distributions of interest can then be found by a series of local computations, involving only some of the variables at a time, see e.g. Cowell et al. (1999) for a thorough treatment of these methods.

So having observed some of the variables in a network, we can use this new evidence to calculate the posterior distribution of any unobserved variable $X_{v}$, given the evidence. Notice that we do not need to observe all the other variables before calculating the posterior distribution, as we can update the prior distribution of $X_{v}$ with any information available. Of course, the more information we have, the better the posterior distribution is determined. However, not all information will have an impact on the posterior distribution of a variable $X_{v}$. Consider the following result. A node $v$ is conditional independent on the rest of the nodes in the network, given the Markov blanket of $v, \mathrm{bl}(v)$, i.e.

$$
v \Perp V \backslash v \mid \mathrm{bl}(v) .
$$

The Markov blanket of $v$ is the set of $v$ 's parents, children and children's parents, i.e.

$$
\operatorname{bl}(v)=\operatorname{pa}(v) \cup \operatorname{ch}(v) \cup\{w: \operatorname{ch}(w) \cap \operatorname{ch}(v)=\emptyset\}
$$

where $\mathrm{pa}(v)$ is the parents of $v$ and $\operatorname{ch}(v)$ is the children of $v$, see Cowell et al. (1999). So if all the variables in the Markov blanket are observed, we do not get further information about the distribution of $X_{v}$ by observing the variables outside the Markov blanket. But if we have not observed all the variables in the Markov blanket, then observing some variable outside the Markov blanket, can influence the posterior distribution of $X_{v}$.

## 5 Results

We will now present the results obtained.

### 5.1 Preliminaries

In the present study, 187 subjects without known diabetes underwent both an OGTT and an IVGTT. In the OGTT, measurements were recorded of plasma glucose (G) and serum insulin levels (I) at time points 10,5 and 0 before intake of 75 gram glucose and at $10,20,30,40,50,60,75,90,105,120,140,160$, 180,210 and 240 minutes after the intake.

In this analysis, the observations to time 10,5 and 0 before the glucose intake are, for both insulin and glucose, averaged and represented by the corresponding observation to time 0 . Further, based on previous results, see Drivsholm et al. (2003), we use the logarithm of the insulin sensitivity index $\log S_{I}$ instead of $S_{I}$ and we also include the sex of the patient and the body mass index (BMI) in the models. Sex is a binary variable, but we choose to treat it as a continuous variable. This has the effect that the variance is assumed equal for male and female observations, whereas the means can differ. If sex is treated as a discrete variable, the data is split into two groups with a parameter set for each group. and we have found that we do not have enough data to support this. Consider for example the simple case, where the only parent to $\log S_{I}$ is sex. If sex is treated as a continuous variable, the distribution of $\log S_{I}$ is given as

$$
\left(\log S_{I} \mid \operatorname{sex}\right) \sim \mathcal{N}\left(m+\beta \operatorname{sex}, \sigma^{2}\right)
$$

None of the parameters $m, \beta$ and $\sigma^{2}$ depend on sex, but the mean is $m$ if sex is 0 and $m+\beta$ if sex is 1 . If sex is treated as a discrete variable, the distribution of $\log S_{I}$ is

$$
\left(\log S_{I} \mid \operatorname{sex}\right) \sim \mathcal{N}\left(m_{\mathrm{sex}}, \sigma_{\mathrm{sex}}^{2}\right)
$$

i.e. both the mean and the variance depends on sex.

In the following we will try different ways of establishing a Bayesian network, which can be used to predict $\log S_{I}$ from measurements from an OGTT and from BMI and sex. So the networks we will consider in the following, only contain continuous variables. Notice that, when using the theory presented for mixed networks on networks with only continuous variables, it coincides with theory developed for purely continuous networks, see B $\emptyset$ ttcher (2001). To learn
the parameters and structure of a Bayesian network, we use the software package deal, see Bøttcher and Dethlefsen (2003). The package is written for R, see Ihaka and Gentleman (1996).

To validate the models, we split the dataset into a subset with 140 subjects, used as training data, and a subset with 47 subjects, used as validation data. For each model, we use deal with the training data to learn the parameters and structure of the Bayesian network. The posterior parameter distribution of $\log S_{I}$ is used to derive point estimates of the parameters. For the Gaussian parameters, we use the mean of the posterior and for the gamma distributed parameter, we use the mode of the posterior. These point estimates are then transfered to Hugin (www. hugin. com). For each subject in both the training data and the validation data, the conditional distribution of $\log S_{I}$ is calculated given the observations from the OGTT using Hugin. In the following, we call this distribution the predictive distribution of $\log S_{I}$. Notice, however, that if a fully Bayesian approach had been used, the predictive distribution for one subject is

$$
p\left(\log S_{I} \mid d\right)=\int_{\theta \in \Theta} p\left(\log S_{I} \mid d, \theta\right) p(\theta) d \theta
$$

where $d$ denotes the subjects OGTT measurements and $\theta$ are the parameters. This distribution is a $t$ distribution with degrees of freedom increasing with, among other numbers, the number of subjects in the training dataset. In this study we have 140 subjects and we find that the error using a Gaussian distribution instead, is very small.

The predictive distribution is then, for each subject, compared with the corresponding $\log S_{I}$ value determined from the IVGTT in the following way. For each subject we use the predictive distribution to calculate the $95 \%$ 's credibility intervals $\mu \pm 1.96 \cdot \sigma$, where 1.96 is the $97.5 \%$ 's quantile in the Gaussian distribution. So if a Bayesian network can predict the value of $\log S_{I}$, we expect that $95 \%$ of the corresponding $\log S_{I}$ values found in the IVGTT study, will lie within this interval. If this is the case, we say that the predictive distribution of $\log S_{I}$ is well calibrated, see Dawid (1982).

Further, we perform an ordinary linear regression of the IVGTT obtained $S_{I}$ on the predicted $S_{I}$ and calculate the residual standard deviation, $S D$, and the correlation coefficient, $R^{2}$, obtained from this regression. To show that there is no systematic bias in these regressions, we report the intercept and slope of these regressions lines.

### 5.2 The Different Models

In the following we will present different approaches for finding a Bayesian network, that can model the dependency relations between the variables in the problem. Further, we will present the results of a previous approach, where multiple linear regression is used and also the results of using the leaps and bounds algorithm for best subset selection.

## Bayesian regression network

Previous results have shown that predictions of $\log S_{I}$ from a multiple regression on OGTT plasma glucose and serum insulin levels, BMI and sex, are highly correlated to the corresponding IVGTT-derived $S_{I}$ estimates, see e.g. Drivsholm et al. (2003). We will therefore learn the parameters and the structure of a network, where $\log S_{I}$ can depend on these variables, and these variables are marginally independent, i.e. the only arrows that are allowed in the model, are arrows into $\log S_{I}$. This network represents a regression model, so we will refer to it as the Bayesian regression network. To learn this network, we need to specify a prior network, i.e. a prior DAG and a prior probability distribution. As prior DAG we use, for simplicity, the empty DAG, i.e. the one without any arrows. This DAG represent that all the variables are independent, so the local probability distribution for each node only depends on the node itself. To specify the prior probability distribution, we use the sample mean and the sample variance as an initial estimate of the mean and variance. As a measure of our confidence in this network, we use $N=100$ for the size of the imaginary database. Figure 1 shows the result of the structural learning procedure. We see that $\log S_{I}$ depends on almost all of the insulin measurements, except for I10, and a few of the glucose measurements.

## Bayesian network with empty prior network

In situations where not all the variables are observed, information is gained by modeling the possible correlations between the explanatory variables. So we will now learn a network, where these correlations are allowed. We only consider networks, where arrows between the glucose and insulin measurements point forward in time, where BMI and sex can not have any parents and where $\log S_{I}$ can not have any children. Again we use the empty DAG as prior DAG,


Figure 1: The result of the structural learning procedure for the Bayesian regression network.
the sample mean and sample variance to specify the prior probability distribution and $N=100$ as a measure of our confidence in this network. The result of the structural learning procedure reveals a complicated dependency structure between the variables, see Figure 2.


Figure 2: The Bayesian network with the empty network as prior.
The Markov blanket for $\log S_{I}$ in this network, is the same as the Bayesian regression network, see Figure 1. The reason for this is that $\log S_{I}$, in both networks, is not allowed to have any children and because we in both approaches have used the same prior network. So when all the variables in the Markov blanket are observed, as it is in our study, the prediction results are exactly the same as for the Bayesian regression network.

## Bayesian network with physiological prior network

In the previous two networks, we have for simplicity used the empty DAG as prior DAG. We will now use a prior DAG, called the physiological network, where the knowledge we have about the physiological relations between the
variables is incorporated. In this network, insulin measurements and glucose measurements are assumed to be Markov processes. They are coupled so that the current glucose measurement depends on the previous insulin measurement and the current insulin measurement depends on the current glucose measurement, see Figure 3. This structure is consistent with the physiological model used in Bergmans minimal model to determine $S_{I}$ from an IVGTT. In addition, we let the initial glucose and insulin measurements depend on BMI and sex.


Figure 3: The physiological network.
Like before, we estimate the prior probability distribution from data. However, contrary to the empty network, the variables in the physiological network depends on other variables, so we perform a linear regression on the parents and use the sample mean and sample variance from these regressions as the mean and variance in the local prior probability distributions. Again we use $N=100$ and we only consider networks where arrows between the glucose and insulin measurements point forward in time, where BMI and sex can not have any parents and where $\log S_{I}$ can not have any children. The result of the structural learning procedure is shown in Figure 4. As before, we see a complicated de-


Figure 4: The Bayesian network with the physiological network as prior.
pendency structure between the variables. In Figure 5, the Markov blanket of $\log S_{I}$ is shown and we see that is quite different than with the empty prior,
shown in Figure 1. Only 6 of the insulin measurements and 5 of the glucose measurements are included in the present blanket.


Figure 5: The Markov blanket for the Bayesian network with the physiological network as prior.

## Results using multiple linear regression

In Drivsholm et al. (2003), multiple linear regression is used to derive predictive equations of $\log S_{I}$ using OGTT plasma glucose and serum insulin levels, BMI and gender. To limit the amount of blood samples drawn from the patients, they constrain the models to include glucose and insulin observations to the same time point. By a combination of backwards elimination and forward selection, they find the optimal model to be with sample time points $0,30,60,105,180$, and 240. Notice, though, that they have found their model on the basis of a different training dataset than ours, as the partition of the dataset into training data and validation data is done randomly in both cases.

## Results using the leaps and bound algorithm

Further, we have tried the leaps and bound algorithm by Furnival and Wilson (1974), using the Bayesian information criteria to find the best subset of the explanatory variables. With this approach, the optimal model is with I50, I90, G160, BMI and sex as explanatory variables. In theory, when the size of the database approaches infinity, using the Bayesian information criteria will result in the same subset of explanatory variables as when using the network score as selection criteria, see Haughton (1988).

### 5.3 Evaluation

To compare the different models, we first consider the network score. Notice that we can only compare network scores for networks that are learned using the same prior network.

To be able to evaluate all models using the network score, we have also calculated the $\log$ scores for the results found in the multiple linear regression approach and the leaps and bounds approach. This is done by formulating these results as Bayesian networks and calculating the log scores using respectively the empty network and the physiological network as prior network. Likewise, for the Bayesian regression network found by using the empty network as prior network, we have calculated the log score using the physiological network as prior network.

| Model | Empty prior | Physiological prior |
| :--- | :--- | ---: |
| BR | -17878.30 | -17848.33 |
| BN | -16528.39 | -14851.44 |
| MLR | -17886.17 | -17849.06 |
| L\&B | -17894.95 | -17846.12 |

Table 1: Network scores for the different models.

The results are reported in Table 1. The Bayesian network model (BN) has the lowest $\log$ score, i.e. the highest network score, both when the empty network and the physiological network are used as prior network. This is obvious as the BN is selected using the network score as selection criteria and because the Bayesian regression (BR), the multiple linear regression (MLR) and the leaps and bounds (L\&B) networks are included in the search space, when searching for the BN with the highest score. So unless we have only found a local maximum, instead of a global maximum, the score for the BN must be higher than the score for the other networks.

When comparing the scores found using the empty prior, we see that the network scores for the BR network, the MLR network and the L\&B network are almost all the same. The network score for the BN is over a thousand times higher than for any of the other networks, indicating that the BN provides a much better fit to data. Recall, however, that the Markov blanket for the BR network and the BN are the same, so when all the variables in the Markov blanket are observed, the BR network and the BN will predict the same $\log S_{I}$ values. So the higher network score is not important when data are complete, but can have an impact
when data are incomplete.
When using the physiological network as prior network, we see almost the same result. The network score for the BR, MLR and L\&B networks are almost all the same, whereas the network score for the BN is over 3000 times higher than for any of the other networks.

| Model | $\operatorname{Tr}$. data $R^{2}(S D)$ | Val. data $R^{2}(S D)$ | Tr. outside | Val. outside |
| :--- | ---: | ---: | ---: | ---: |
| BR with empty prior | $0.76(0.31)$ | $0.73(0.35)$ | $1(1 \%)$ | $1(2 \%)$ |
| BN with empty prior | $0.76(0.31)$ | $0.73(0.35)$ | $1(1 \%)$ | $1(2 \%)$ |
| BN with physiological prior | $0.77(0.30)$ | $0.73(0.36)$ | $7(5 \%)$ | $3(6 \%)$ |
| MLR | $0.76(0.31)$ | $0.66(0.40)$ | $3(2 \%)$ | $3(6 \%)$ |
| L\&B | $0.75(0.31)$ | $0.73(0.36)$ | $6(4 \%)$ | $4(9 \%)$ |

Table 2: The table lists the $R^{2}$ and $S D$ values from the linear regressions of the IVGTT obtained $\log S_{I}$ on the predicted $\log S_{I}$ for both the training dataset and the validation dataset. Also listed are how many $\log S_{I}$ values that fall outside the credibility interval $\mu \pm 1.96 \cdot \sigma$.

In Table 2 the $R^{2}$ and $S D$ values from the linear regression of the IVGTT obtained $\log S_{I}$ on the predicted $\log S_{I}$ are reported. The $R^{2}$ and $S D$ values are for all five models acceptable and they are almost the same for all models, except for the multiple regression model, which on the validation dataset does not perform as well as the others. Table 3 shows the intercept and slope of the estimated regression lines and there are no evidence of any systematic bias. We therefore conclude that an OGTT can be used to determine the insulin sensitivity index.

| Model | Tr. data (intercept, slope) | Val. data (intercept, slope) |
| :--- | ---: | ---: |
| BR empty prior | $(-0.19,1.09)$ | $(-0.05,1.01)$ |
| BN empty prior | $(-0.03,1.01)$ | $(0.25,0.87)$ |
| BN phsyiological prior | $(-0.06,1.03)$ | $(0.14,0.92)$ |
| MLR | $(0,1)$ | $(0.06,0.96)$ |
| L\&B | $(0,1)$ | $(0.11,0.93)$ |

Table 3: The intercept and slope of the regressions lines from the regressions of the IVGTT obtained $S_{I}$ on the predicted $S_{I}$. Reported to show that there is no evidence of systematic bias.

In Table 2 we have also listed how many $\log S_{I}$ values that fall outside the credibility interval $\mu \pm 1.96 \cdot \sigma$. Approximatively $5 \%$ of these predictions should
lie outside and $95 \%$ inside the interval for the predictive distributions to be well calibrated. This is clearly fulfilled for the BN with the physiological network as prior network, so the predictive distribution for $\log S_{I}$ is, when using this network, well calibrated. With the MLR approach and the L\&B approach it is almost fulfilled that $5 \%$ of the predictions lie outside the intervals. We will therefore conclude that the predictive distributions are also well calibrated in these cases. For the BR and the BN with the empty network as prior network, very few values lie outside the intervals, indicating that the variance is probably estimated to large. Figure 6 shows the predicted $\log S_{I}$ values and the intervals for the BN with the empty prior and for the BN with the physiological prior. We see that for the two models, the predicted $\log S_{I}$ values are almost the same, but the intervals are much wider for the BN with the empty prior, meaning that the variance in this model is larger.


Figure 6: The predicted $\log S_{I}$ values and the credibility intervals for the Bayesian network with empty prior (dark and disks) versus the Bayesian network with physiological prior (light and triangles).

So to summarize, all the models give adequate predictions of the $\log S_{I}$ values. Evaluating the models using the different validation approaches all together, the BN with the physiological prior model gives a more precise predictive distribution of $\log S_{I}$ compared to the other models. We therefore suggest that this model should be used to derive the predictive values of $\log S_{I}$.

## 6 Discussion

We have established a promising way of determining the insulin sensitivity index from an oral glucose tolerance test rather than from an intravenous glucose tolerance test. All approaches give adequate predictions of $S_{I}$. The Bayesian network with the physiological prior estimates the most precise predictive distribution of $S_{I}$, so we claim that this is the best model. There are also other advantages by using a Bayesian network instead of an ordinary regression model. In a Bayesian network, we can use any prior knowledge available from e.g. previous studies or from the physiological understanding of the problem. Further, we can calculate the predictive distribution of $\log S_{I}$ in situations, where some of the observations are missing. This can be used when a single or a few observations are missing for a specific subject. It can also be used when certain time points are not observed at all, which could be the case if a dataset from another study, using fewer time points, is analyzed.

## Acknowledgements

This work has been supported by Novo Nordisk A/S. The collection of data were done by Torben Hansen, Novo Nordisk A/S.

## References

Bergman, R. N., Ider, Y. Z., Bowden, C. R. and Cobelli, C. (1979). Quantitative estimation of insulin sensitivity, American Journal of Physiology 236: E667-E677.

Bøttcher, S. G. (2001). Learning Bayesian Networks with Mixed Variables, Artificial Intelligence and Statistics 2001, Morgan Kaufmann, San Francisco, CA, USA, pp. 149-156.

Bøttcher, S. G. and Dethlefsen, C. (2003). deal: A Package for Learning Bayesian Networks, Journal of Statistical Software 8(20): 1-40.

Cooper, G. and Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data, Machine Learning 9: 309-347.

Cowell, R. G., Dawid, A. P., Lauritzen, S. L. and Spiegelhalter, D. J. (1999). Probabilistic Networks and Expert Systems, Springer-Verlag, Berlin-Heidelberg-New York.

Dawid, A. P. (1982). The Well-Calibrated Bayesian, Journal of the American Statistical Association 77(379): 605-610.

Drivsholm, T., Hansen, T., Urhammer, S. A., Palacios, R. T., Vølund, A., BorchJohnsen, K. and Pedersen, O. B. (2003). Assessment of insulin sensitivity index and acute insulin reponse from an oral glucose tolerance test in subjects with normal glucose tolerance, Novo Nordisk A/S.

Furnival, G. M. and Wilson, R. W. (1974). Regression by Leaps and Bounds, Technometrics 16(4): 499-511.

Geiger, D. and Heckerman, D. (1994). Learning Gaussian Networks, Proceedings of Tenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 235-243.

Haughton, D. M. A. (1988). On The Choice of a Model to fit Data From an Exponential Family, The Annals of Statistics 16(1): 342-355.

Heckerman, D., Geiger, D. and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data, Machine Learning 20: 197-243.

Ihaka, R. and Gentleman, R. (1996). R: A language for data analysis and graphics, Journal of Computational and Graphical Statistics 5: 299-314.

Lauritzen, S. L. (1992). Propagation of probabilities, means and variances in mixed graphical association models, Journal of the American Statistical Association 87(420): 1098-1108.

Martin, B. C., Warram, J. H., Krolewski, A. S., Bergman, R. N., Soeldner, J. S. and Kahn, C. R. (1992). Role of glucose and insulin resistance in development of type 2 diabetes mellitus: results of a 25 -year follow-up study, The Lancet 340: 925-929.

Pacini, G. and Bergman, R. N. (1986). MINMOD: a computer program to calculate insulin sensitivity and pancreatic responsivity from the frequently sampled intraveneous glucose tolerance test, Computer Methods and Programs in Biomedicine 23: 113-122.

# Paper IV Learning Dynamic Bayesian Networks with Mixed Variables 

Susanne G. Bøttcher

# Learning Dynamic Bayesian Networks with Mixed Variables 

Susanne G. Bøttcher<br>Aalborg University, Denmark


#### Abstract

. This paper considers dynamic Bayesian networks for discrete and continuous variables. We only treat the case, where the distribution of the variables is conditional Gaussian. We show how to learn the parameters and structure of a dynamic Bayesian network and also how the Markov order can be learned. An automated procedure for specifying prior distributions for the parameters in a dynamic Bayesian network is presented. It is a simple extension of the procedure for the ordinary Bayesian networks. Finally the Wölfer's sunspot numbers are analyzed.


## 1 Introduction

In this paper we consider dynamic Bayesian networks (DBNs) for discrete and continuous variables. A DBN is an extension of an ordinary Bayesian network and is applied in the modeling of time series.

DBNs for first order Markov time series are described in Dean and Kanazawa (1989). In Murphy (2002), a thorough treatment of these models is presented and in Friedman et al. (1998) learning these networks in the case with only discrete variables is described.

Here we consider DBNs with both discrete and continuous variables. In these networks we also allow some of the variables to be static, i.e. some of the variables do not change over time. We only treat the case where the distribution of the variables is conditional Gaussian (CG) and show how to learn the parameters and structure of the DBN when data is complete. Further we present an automated method for specifying prior parameter distributions for the parameters in a DBN. These methods are simple extensions of the ones used for ordinary Bayesian networks with mixed variables, described in Bøttcher (2001).

We consider time series, where the Markov order can be higher than one and show how the Markov order can be learned.

In Section 2, DBNs with static and time varying variables are defined. Section 3
presents these DBNs for the mixed case and Section 4 gives some examples of some well known models that can be represented as DBNs. Section 5 shows how to learn the parameters and structure of a DBN with mixed variables. Further, it shows how the Markov order can be learned. Section 6 presents a method for specifying prior distributions of the parameters in the DBN. In Section 7 Wölfer's sunspot numbers are analyzed using a DBN.

## 2 Dynamic Bayesian Networks

A Bayesian network is a graphical model that encodes the joint probability distribution for a set of variables. For terminology and theoretical aspects on graphical models, see Lauritzen (1996). We define it as a directed acyclic graph (DAG) $D=(V, E)$, where $V$ is a finite set of nodes and $E$ is a finite set of directed edges between the nodes. The DAG defines the structure of the Bayesian network. To each node $v \in V$ in the graph corresponds a random variable $X_{v}$. The set of variables associated with the graph $D$ is then $X=\left(X_{v}\right)_{v \in V}$.

To each vertex $v$ with parents $\mathrm{pa}(v)$, there is attached a local probability distribution, $p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)$. The possible lack of directed edges in $D$ encodes conditional independencies between the random variables $X$ through the factorization of the joint probability distribution,

$$
p(x)=\prod_{v \in V} p\left(x_{v} \mid x_{\mathrm{pa}(v)}\right)
$$

In a Bayesian network, the set of random variables $X$ is fixed. To model a multivariate time series we need a framework, where we allow the set of random variables to vary with time. For this we use dynamic Bayesian networks, defined as below. This definition is consistent with the exposition in Murphy (2002), but here we also allow for static variables and Markov orders higher than one.

Let $X^{t}$ be a set of time varying random variables, that is $X^{t}$ can take on the values $X^{0}, X^{1}, \ldots, X^{T}$. We index the time varying variables by the non-negative integers to indicate that the observations are taken at discrete time points. The corresponding nodes in the graph are denoted $V_{t}$, so $X^{t}=\left(X_{v}^{t}\right)_{v \in V_{t}}$ for each time point $t$. Note however that $V_{t}$ is "the same" for all time points $t$, so formally $V_{t}=\{(v, t), v \in V\}$. Further, let $X^{s}$ be a set of static random variables, i.e. variables that do not change over time. The nodes corresponding
to $X^{s}$ are denoted $V_{s}$. The set of variables associated with a DBN is then $X=\left(\left(X^{t}\right)_{t=0}^{T}, X^{s}\right)$ and the set of nodes is $V=\left(\left(V_{t}\right)_{t=0}^{T}, V_{s}\right)$.

We refer to the time varying variables at one time point as a time slice or just a slice. We let the static variables $X^{s}$ belong to the time slice at time $t=0$ and refer to this as the initial time slice. So the initial time slice includes the variables $X^{0}$ and $X^{s}$ and, for $t=1, \ldots, T$, the time slice at time $t$ includes the variables $X^{t}$.

We will mostly consider the variables in the initial time slice jointly, so to ease later notation we define $X^{\tilde{0}}=\left(X^{0}, X^{s}\right)$ and $V_{\tilde{0}}=\left(V_{0}, V_{s}\right)$.

The joint probability distribution of the variables in a dynamic Bayesian network can be very complex, as the number of variables grows over time. Therefore we assume that the time series we are dealing with, is mth order Markov, i.e.

$$
p\left(x^{t} \mid x^{t-1}, \ldots, x^{0}\right)=p\left(x^{t} \mid x^{t-1}, \ldots, x^{t-m}\right)
$$

for all time points $t=m, \ldots, T$.
Further, we assume that the time series has stationary dynamics, so

$$
p\left(x^{t} \mid x^{t-1}, \ldots, x^{t-m}\right)=p\left(x^{m} \mid x^{m-1}, \ldots, x^{0}\right)
$$

for all $t=m, \ldots, T$. Stationary dynamics refers to the fact that the conditional distributions are time independent, while the marginal distributions may be time dependent.

We will first introduce DBNs for time series that are first order Markov. With the above assumptions, a DBN for a first order Markov time series can be defined to be the pair $\left(B_{\tilde{0}}, B_{\rightarrow}\right)$, where $B_{\tilde{0}}$ is a Bayesian network defining the probability distribution of $X^{\tilde{0}}$ as

$$
p\left(x^{\tilde{0}}\right)=\prod_{v \in V_{\tilde{0}}} p\left(x_{v}^{\tilde{o}} \mid x_{\mathrm{pa}(v)}^{\tilde{0}}\right)
$$

and $B_{\rightarrow}$ is a 2-slice temporal Bayesian network defining the conditional distribution of $X^{t}$ as

$$
p\left(x^{t} \mid x^{t-1}, x^{s}\right)=\prod_{v \in V_{t}} p\left(x_{v}^{t} \mid x_{\mathrm{pa}(v)}^{t}, x_{\mathrm{pa}(v)}^{t-1}, x_{\mathrm{pa}(v)}^{s}\right)
$$

The joint probability distribution for a DBN with $T+1$ time points is given as

$$
p\left(x^{0}, \ldots, x^{T}, x^{s}\right)=p\left(x^{\tilde{0}}\right) \prod_{t=1}^{T} p\left(x^{t} \mid x^{t-1}, x^{s}\right)
$$

As we assumed that the time series has stationary dynamics, the DBN is completely specified through $B_{\overline{0}}$ and $B_{\rightarrow}$.

For the dependency relations between the time slices we assume that arrows point forward in time, so the variables in time slice $t$ can have parents in the time slices to time $t$ and $t-1$. Further, they can have parents from $X^{s}$. Due to stationary dynamics, the dependency relations between the time slices are the same for all time points. This also means that if a time varying variable $X_{v}^{t}$ has a static variable $X_{w}^{s}$ as a parent, then $X_{w}^{s}$ is also a parent of $X_{v}^{1}, \ldots, X_{v}^{T}$. The variables in the initial time slice can have parents from the initial time slice and therefore also from $X^{s}$, as $X^{s}$ is included in the initial time slice.

Within a time slice, there are no restrictions of the dependency relations between the variables, as long as the structure is a DAG. Due to stationary dynamics, the dependency relations within a time slice are the same for the time slices to time $t=1, \ldots, T$. They are however not necessarily the same as for the time varying variables in the initial time slice.

So the structure of the DBN repeats itself over time, except for $B_{\tilde{0}}$, where the time series is initialized.

Figure 1 shows an example of the structure of a a first order Markov DBN, $\left(B_{\tilde{0}}, B_{\rightarrow}\right)$, with two time varying variables $Y^{t}$ and $Z^{t}$ and one static variable $X^{s}$. Because of the first order Markov property, the structure is completely specified through the first two time points and the structure of the DBN can therefore be represented by the DAG in Figure 2.


Figure 1: Example of a first order Markov DBN $\left(B_{\tilde{0}}, B_{\rightarrow}\right)$.
For time series with higher Markov order properties, we need to extend the definition.

Consider an $m$ th order Markov time series. The joint probability distribution


Figure 2: A first order Markov DBN $\left(B_{\tilde{0}}, B_{\rightarrow}\right)$ represented by the first two time points.
for $T+1$ time points can be written as

$$
\begin{aligned}
p\left(x^{0}, \ldots, x^{T}, x^{s}\right) & =p\left(x^{\tilde{0}}, x^{1}, \ldots, x^{m-1}\right) \prod_{t=m}^{T} p\left(x^{t} \mid x^{t-1}, \ldots, x^{t-m}, x^{s}\right) \\
& =p\left(x^{\tilde{0}}\right) p\left(x^{1} \mid x^{\tilde{0}}\right) \cdots, p\left(x^{m-1} \mid x^{m-2}, \ldots, x^{\tilde{0}}\right) \\
& \times \prod_{t=m}^{T} p\left(x^{t} \mid x^{t-1}, \ldots, x^{t-m}, x^{s}\right)
\end{aligned}
$$

Following the definition for first order Markov time series, we let $B_{\rightarrow}$ be a $m+1$-slice temporal Bayesian network defining the conditional distribution of $X^{t}$,

$$
p\left(x^{t} \mid x^{t-1}, \ldots, x^{t-m}, x^{s}\right)=\prod_{v \in V_{t}} p\left(x_{v}^{t} \mid x_{\mathrm{pa}(v)}^{t}, \ldots, x_{\mathrm{pa}(v)}^{t-m}, x_{\mathrm{pa}(v)}^{s}\right),
$$

for $t=m, \ldots, T$.
The variables in time slice $t$ can have parents in the time slices to times $t, \ldots, t-$ $m$ and they can have parents from $X^{s}$. Again, due to stationary dynamics, the dependency relations between and within the time slices are the same for all time points $t=m, \ldots, T$. Further, if a time varying variable $X_{v}^{t}$ has a static variable $X_{w}^{s}$ as a parent, then $X_{w}^{s}$ is also a parent of $X_{v}^{m}, \ldots, X_{v}^{T}$.

The question is now how to initialize the time series. The probability distribution $p\left(x^{\tilde{0}}, x^{1}, \ldots, x^{m-1}\right)$ can be written as

$$
\begin{equation*}
p\left(x^{\tilde{0}}, x^{1}, \ldots, x^{m-1}\right)=p\left(x^{\tilde{0}}\right) p\left(x^{1} \mid x^{\tilde{0}}\right) \cdots p\left(x^{m-1} \mid x^{m-2}, \ldots, x^{\tilde{0}}\right) \tag{1}
\end{equation*}
$$

As arrows point forward in time, this factorization defines the possible dependency relations between the variables $X^{\tilde{0}}, \ldots, X^{m-1}$. As before we let $B_{\tilde{0}}$ be a

Bayesian network defining the probability distribution of $X^{\tilde{0}}$ as

$$
p\left(x^{\tilde{0}}\right)=\prod_{v \in V_{\tilde{0}}} p\left(x_{v}^{\tilde{0}} \mid x_{\mathrm{pa}(v)}^{\tilde{0}}\right) .
$$

Now we also define Bayesian networks for the rest of the conditional distributions in (1). We let $B_{1}$ be a 2 -slice Bayesian network defining the conditional distribution of $X^{1}$ given $X^{\tilde{0}}$ as

$$
p\left(x^{1} \mid x^{\tilde{0}}\right)=\prod_{v \in V_{1}} p\left(x_{v}^{1} \mid x_{\mathrm{pa}(v)}^{1}, x_{\mathrm{pa}(v)}^{\tilde{0}}\right),
$$

and likewise for $B_{2}, \ldots, B_{m-1}$, where $B_{m-1}$ is an $m$-slice Bayesian network defining the conditional distribution of $X^{m-1}$ given $X^{m-2}, \ldots, X^{\tilde{0}}$ as

$$
p\left(x^{m-1} \mid x^{m-2}, \ldots, x^{\tilde{0}}\right)=\prod_{v \in V_{m-1}} p\left(x_{v}^{m-1} \mid x_{\mathrm{pa}(v)}^{m-1}, \ldots, x_{\mathrm{pa}(v)}^{\tilde{0}}\right)
$$

So the variables in the time slice to time $t=1$ can have parents from the time slice to time $t=1$ and $t=0$. The variables in time slice $m-1$ can have parents from the time slices to time $t=0, \ldots, m-1$. The dependency relations between the time slices to time $t=0, \ldots, m-1$ are obviously not the same and the dependency relations within these time slices are not necessarily the same.

The tuple $\left(B_{\tilde{0}}, B_{1}, \ldots, B_{m-1}, B_{\rightarrow}\right)$ is thus a DBN for an $m$ th order Markov time series, where the different Bayesian networks in the tuple defines the corresponding probability distributions as above. Notice that we could also just have specified the networks $B_{\tilde{0}}, B_{1}, \ldots, B_{m-1}$ as one large network, with the necessary restrictions on the arrows.

## 3 Dynamic Bayesian Networks for Mixed Variables

In this section we consider DBNs with mixed variables, i.e. the variables in the network can be of discrete and continuous type. We let $V=\Delta \cup \Gamma$, where $\Delta$ and $\Gamma$ are the sets of discrete and continuous variables, respectively. The corresponding random variables $X$ can then be denoted $X=\left(X_{v}\right)_{v \in V}=(I, Y)=$ $\left(\left(I_{\delta}\right)_{\delta \in \Delta},\left(Y_{\gamma}\right)_{\gamma \in \Gamma}\right)$. Again, we index the sets of nodes and the random variables with $t$ for time varying variables, $s$ for static variables and $\tilde{0}$ for the variables in the initial time slice.

To ensure availability of exact local computation methods, we do not allow continuous parents of discrete nodes, so the probability distributions factorize into a discrete part and a mixed part as presented below. To simplify notation, we present the theory for first order Markov time series and comment on how to extend it to higher order Markov assumptions by following the definitions introduced in the previous section.

We consider $B_{\tilde{0}}$ and $B_{\rightarrow}$ separately, and the joint probability distribution is obtained as specified in the previous section.

For $B_{\tilde{0}}$ we have that

$$
\begin{align*}
p\left(x^{\tilde{0}}\right) & =\prod_{v \in V_{\tilde{0}}} p\left(x_{v}^{\tilde{0}} \mid x_{\mathrm{pa}(v)}^{\tilde{0}}\right) \\
& =\prod_{\delta \in \Delta_{\tilde{0}}} p\left(i_{\delta}^{\tilde{0}} \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}\right) \prod_{\gamma \in \Gamma_{\tilde{0}}} p\left(y_{\gamma}^{\tilde{0}} \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, y_{\mathrm{pa}(\gamma)}^{\tilde{0}}\right) \tag{2}
\end{align*}
$$

and for $B_{\rightarrow}$

$$
\begin{align*}
p\left(x^{t} \mid x^{t-1}, x^{s}\right) & =\prod_{v \in V_{t}} p\left(x_{v}^{t} \mid x_{\mathrm{pa}(v)}^{t}, x_{\mathrm{pa}(v)}^{t-1}, x_{\mathrm{pa}(v)}^{s}\right) \\
& =\prod_{\delta \in \Delta_{t}} p\left(i_{\delta}^{t} \mid i_{\mathrm{pa}(\delta)}^{t}, i_{\mathrm{pa}(\delta)}^{t-1}, i_{\mathrm{pa}(\delta)}^{s}\right)  \tag{3}\\
& \times \prod_{\gamma \in \Gamma_{t}} p\left(y_{\gamma}^{t} \mid i_{\mathrm{pa}(\gamma)}^{t}, i_{\mathrm{pa}(\gamma)}^{t-1}, i_{\mathrm{pa}(\gamma)}^{s}, y_{\mathrm{pa}(\gamma)}^{t}, y_{\mathrm{pa}(\gamma)}^{t-1}, y_{\mathrm{pa}(\gamma)}^{s}\right)
\end{align*}
$$

To account for higher order Markov assumptions, we would just have to specify the probability distributions for the intervening networks accordingly.

To simplify notation for $B_{\rightarrow}$, we use the following notation, where the possible parent configurations are not explicitly defined. They must be specified in the given context and according to (3).

$$
\begin{aligned}
p\left(x^{t} \mid x^{t-1}, x^{s}\right) & =\prod_{v \in V_{t}} p\left(x_{v}^{t} \mid x_{\mathbf{p a}(v)}\right) \\
& =\prod_{\delta \in \Delta_{t}} p\left(i_{\delta}^{t} \mid i_{\mathrm{pa}(\delta)}^{\rightarrow}\right) \prod_{\gamma \in \Gamma_{t}} p\left(y_{\gamma}^{t} \mid i_{\mathrm{pa}(\gamma)}^{\overrightarrow{2}}, y_{\mathrm{pa}(\gamma)}^{\overrightarrow{2}}\right)
\end{aligned}
$$

So for example, $i_{\mathrm{pa}(\delta)}^{\rightarrow}$ contains the variables $i_{\mathrm{pa}(\delta)}^{t}, i_{\mathrm{pa}(\delta)}^{t-1}$ and $i_{\mathrm{pa}(\delta)}^{s}$.
In this paper we only consider networks, where the joint distribution of the variables is conditional Gaussian. The local probability distributions are therefore
defined as in the following two sections. In these sections, we do not distinguish between the variables in $B_{\tilde{0}}$ and $B_{\rightarrow}$, as the distribution of these variables is of the same type. The possible parent set differ however between variables in $B_{\tilde{0}}$ and variables in $B_{\rightarrow}$. In the following we therefore just denote the parents of a variable $x_{v}$ by $x_{\mathrm{pa}(v)}$ and $x_{\mathrm{pa}(v)}$ must be specified according to (2) or (3).

### 3.1 Distribution for Discrete Variables

When the joint distribution is conditional Gaussian, the local probability distributions for the discrete variables are just unrestricted discrete distributions with

$$
p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}\right) \geq 0 \quad \forall \quad \delta \in \Delta
$$

We parameterize this as

$$
\theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=p\left(i_{\delta} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)
$$

where $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(\theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}$.
Furthermore $\sum_{i_{\delta} \in \mathcal{I}_{\delta}} \theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=1$ and $0 \leq \theta_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}} \leq 1$. All parameters associated with a node $\delta$ is denoted by $\theta_{\delta}$, so $\theta_{\delta}=\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}}$.

### 3.2 Distribution for Continuous Variables

For the continuous variables, the local probability distributions are Gaussian linear regressions with parameters depending on the configuration of the discrete parents. So let the parameters be given by $\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}=\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}^{2}\right)$. Then

$$
\begin{equation*}
\left(Y_{\gamma} \mid y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}, \theta_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}\right) \sim \mathcal{N}\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} y_{\mathrm{pa}(\gamma)}, \sigma_{\left.\gamma \mid i_{\mathrm{pa}(\gamma)}\right)}^{2}\right) \tag{4}
\end{equation*}
$$

where $\beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ are the regression coefficients, $m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ is the regression intercept, and $\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}$ is the conditional variance. Thus for each configuration of the discrete parents of $\gamma$ the distribution of $Y_{\gamma}$ is Gaussian with mean and variance given as in (4). The parameters associated with a node $\gamma$ is then $\theta_{\gamma}=\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}}$.

### 3.3 The Parameterized Distributions

With the above distributional assumptions, we can specify the parameterized DBN as follows.

Let $\theta^{\tilde{0}}=\left(\left(\theta_{\delta}^{\tilde{0}}\right)_{\delta \in \Delta_{\tilde{0}}},\left(\theta_{\gamma}^{\tilde{0}}\right)_{\gamma \in \Gamma_{\tilde{0}}}\right)$ and $\theta^{\rightarrow}=\left(\left(\theta_{\delta}\right)_{\delta \in \Delta_{t}},\left(\theta_{\gamma}\right)_{\gamma \in \Gamma_{t}}\right)$. Further, let $\theta=\left(\theta^{\tilde{0}}, \theta^{\rightarrow}\right)$. Then $B_{\tilde{0}}$ is given as

$$
p\left(x^{\tilde{0}} \mid \theta^{\tilde{o}}\right)=\prod_{\delta \in \Delta_{\tilde{o}}} p\left(i_{\delta}^{\tilde{0}} \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}}^{\tilde{0}}\right) \prod_{\gamma \in \Gamma_{\tilde{\mathrm{o}}}} p\left(y_{\gamma}^{\tilde{0}} \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\tilde{0}}\right)
$$

and $B_{\rightarrow}$ as

$$
\begin{aligned}
p\left(x^{t} \mid x^{t-1}, x^{s}, \theta^{\rightarrow}\right) & =\prod_{\delta \in \Delta_{t}} p\left(i_{\delta}^{t} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \\
& \times \prod_{\gamma \in \Gamma_{t}} p\left(y_{\gamma}^{t} \mid y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)
\end{aligned}
$$

The joint distribution for $T+1$ time points is given as

$$
p\left(x^{0}, \ldots, x^{T}, x^{s}, \theta\right)=p\left(x^{\tilde{0}} \mid \theta^{\tilde{0}}\right) \prod_{t=1}^{T} p\left(x^{t} \mid x^{t-1}, x^{s}, \theta^{\rightarrow}\right)
$$

Notice that, due to stationarity, $\theta^{\rightarrow}$ is the parameter in the conditional distribution of $x^{t}$ for every time point $t=1, \ldots, T$.

## 4 Examples of DBNs

We will now give some examples of some well known models that can be represented as DBNs. In the figures, shaded nodes represent discrete variables and clear nodes represent continuous variables.

### 4.1 Hidden Markov Models

A Hidden Markov Model (HMM) is a stochastic automaton, where each state generates an observation. Figure 3 shows a HMM, where the hidden states are
first order Markov.


Figure 3: A Hidden Markov Model.

The hidden states, i.e. the discrete hidden variables, are denoted by $I$ and the observations by $Y$. We have represented the observed variables as continuous, but they can also all be discrete. In this HMM, $I^{t+1}$ is conditionally independent of $I^{t-1}$, given $I^{t}$. Further, $Y^{t}$ is conditionally independent of the rest of the variables in the network, given $I^{t}$. A model like this is used in situations, where the observations do not follow the same model all the time, but can follow different models at different times. This gives for example the possibility to account for outliers.

When a HMM is represented as a DBN, we assume that the time series has stationary dynamics. So, together with the first order Markov property, we can specify the joint probability distribution for the variables in this network by just specifying the initial prior probabilities $p\left(i^{0}\right)$, the transition probabilities $p\left(i^{t} \mid i^{t-1}\right)$ and the conditional Gaussian distributions $p\left(y^{t} \mid i^{t}\right)$ (or, if the observed variables are discrete, the conditional multinomial distributions $p\left(j^{t} \mid i^{t}\right)$ ).

There are many variants of this basic HMM, e.g. Buried Markov Model, Mixedmemory Markov Model and Hierarchical HMM, see Murphy (2002) for a presentation of these models represented as DBNs and their application within speech recognition.

### 4.2 Kalman Filter Models

A Kalman Filter Model (KFM), introduced by Harrison and Stevens (1976) as a state space model, models the dynamic behavior of a time series. In such a model, the continuous observations $Y$ are indirect measurements of a latent Markov process $Z$.

In Figure 4, a KFM is shown. The structure is the same as for the HMM, since the two models assume the same set of conditional independencies. The probability distributions to be specified is the Gaussian distribution $p\left(z^{0}\right)$, the Gaus-


Figure 4: A Kalman Filter Model.
sian linear regression $p\left(z^{t} \mid z^{t-1}\right)$ and the Gaussian linear regression $p\left(y^{t} \mid z^{t}\right)$. For a comprehensive treatment of KFMs and their applications, see West and Harrison (1989).

### 4.3 Multiprocess Kalman Filter Models

Multiprocess Kalman Filter Models (MKFMs), also known as switching state space Markov models, are an extension of the KFMs, see Harrison and Stevens (1976), where the aim is to discriminate between different KFMs.


Figure 5: A Multiprocess Kalman Filter Model.
Figure 5 shows a MKFM. Again we see that the continuous observations $Y$ are indirect measurements of a latent continuous Markov process $Z$, i.e. this part of the network represents a KFM. In addition, the process $Z$ depends on the hidden states $I$, which in our example are first order Markov. Like the HMM, this model can be used in situations, where the observations do not follow the same model all the time, but can follow different models at different times, but here the models are KFMs. Applications include modeling piece-wise linear time series, which for example can be used for monitoring purposes, see e.g. Bøttcher, Milsgaard and Mortensen (1995).

Notice that because of the first order Markov property assumed for HMMs, KFMs and MKFMs, these models could have been represented by using only the first two time points, as the structure repeats over time.

### 4.4 Vector Autoregressive processes

Another classical time series model is the Vector Autoregressive process (VAR) of Markov order $p$. This model is equivalent to a DBN of Markov order $p$, in which all the variables are continuous and observed. So the local probability distributions in this model are Gaussian linear regressions on the continuous parents.


Figure 6: A Vector Autoregressive process.
In Figure 6, an example of a VAR process of order 2 is given. Because of the second order Markov property, this model can be represented by the first three time points.

In the next section, we will developed a method for learning the parameters and structure of a DBN. In this paper we assume that data are complete, so we can not learn networks with hidden variables. Therefore, the HMM, the KFM and the MKFM can only be learned with these methods, if a training dataset with complete data is available.

## 5 Learning DBNs with Mixed Variables

Learning first order Markov DBNs in the purely discrete case with no static variables is described in Friedman et al. (1998). Here we will consider learning DBNs with mixed variables for the case with both time varying and static variables. Further, we will also illustrate how to learn DBNs with higher Markov order and how to learn this order.

As noted in Murphy (2002), learning DBNs is, because of the way DBNs are defined, just a simple extension of learning BNs. This also applies for DBNs
with mixed variables, so we will use the theory for learning Bayesian networks with mixed variables, described in Bøttcher (2001).

### 5.1 Parameter Learning

To learn the parameters for a given DAG, we use a Bayesian approach. We specify a prior distribution of a parameter $\theta$, use a random sample $d$ from the probability distribution $p(x \mid \theta)$ and obtain the posterior distribution by using Bayes' theorem

$$
p(\theta \mid d) \propto p(d \mid \theta) p(\theta)
$$

The proportionality constant is determined by the relation $\int_{\Theta} p(\theta \mid d) d \theta=1$, where $\Theta$ is the parameter space.

To obtain closed formed expressions, we use conjugate distributions of the parameters.

We assume that the parameters associated with $B_{\tilde{0}}$ and $B_{\rightarrow}$ are independent. Further, for the parameters in respectively $B_{\tilde{0}}$ and $B_{\rightarrow}$, we assume that the parameters associated with one variable is independent of the parameters associated with the other variables and that the parameters are independent for each configuration of the discrete parents, i.e.

$$
\begin{align*}
p(\theta) & =p\left(\theta^{\tilde{0}}\right) p\left(\theta^{\rightarrow}\right) \\
& =\prod_{\delta \in \Delta_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\left.\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}\right)} \prod_{\gamma \in \Gamma_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{0}}^{\tilde{0}}\right)\right.  \tag{5}\\
& \times \prod_{\delta \in \Delta_{t}} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \prod_{\gamma \in \Gamma_{t}} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)
\end{align*}
$$

We refer to this as parameter independence. Notice though that it is slightly different than parameter independence for ordinary Bayesian networks, as we here assume that the parameters in $B_{\rightarrow}$ are the same for each time point $t=$ $1, \ldots, T$.

In the case with higher order Markov properties, parameter independence is also valid for the parameters in the networks $B_{1}, \ldots, B_{m-1}$.

We also assume complete data, i.e. each case ${ }^{c} x$ in a dataset $d$ contains one instance of every random variable in the network. With this we can show posterior parameter independence. The likelihood $p(d \mid \theta)$ can be written as follows.

$$
\begin{aligned}
p(d \mid \theta) & =\prod_{c \in d} p\left({ }^{c} x^{0}, \ldots,{ }^{c} x^{T},{ }^{c} x^{s} \mid \theta\right) \\
& =\prod_{c \in d}\left(p\left({ }^{c} x^{\tilde{0}} \mid \theta^{\tilde{0}}\right) \prod_{t=1}^{T} p\left({ }^{c} x^{t} \mid{ }^{c} x^{t-1},{ }^{c} x^{s}, \theta^{\rightarrow}\right)\right)
\end{aligned}
$$

As the time series has stationary dynamics, we see that for each observations of the variables in $B_{0}$, there are $T$ observations of the variables in $B_{\rightarrow}$.

To simplify the expressions, we consider the likelihood terms for $B_{\tilde{0}}$ and $B_{\rightarrow}$ separately. For $B_{\tilde{0}}$ we have that

$$
\prod_{c \in d} p\left({ }^{c} x^{\tilde{0}} \mid \theta^{\tilde{0}}\right)=\prod_{c \in d} \prod_{\delta \in \Delta_{\tilde{0}}} p\left(c_{i} i_{\delta}^{\tilde{0}} \mid{ }_{i}^{c} i_{\mathrm{pa}(\delta)}^{\tilde{0}}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}}^{\tilde{0}}\right) \prod_{\gamma \in \Gamma_{\tilde{0}}} p\left({ }^{c} y_{\gamma}^{\tilde{o}} \mid{ }^{c} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, c_{i}^{\tilde{0}}{ }_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\tilde{0}}\right)
$$

where ${ }^{c} i$ and ${ }^{c} y$ respectively denotes the discrete part and the continuous part of a case ${ }^{c} x$. Our goal is to show posterior parameter independence, so we must show that the likelihood, like the parameters, factorizes into a product over nodes and a product over the configuration of the discrete parents of a node. Therefore we write this part of the likelihood as

$$
\begin{align*}
\prod_{c \in d} p\left({ }^{c} x^{\tilde{0}} \mid \theta^{\tilde{0}}\right) & =\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \prod_{c: c i_{\mathrm{pa}(\delta)}^{\tilde{0}}=i_{\mathrm{pa}(\delta)}^{\tilde{0}}} p\left({ }^{c} i_{\delta}^{\tilde{0}} \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}}\right) \\
& \times \prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \prod_{c: c i_{\mathrm{pa}(\gamma)}^{\tilde{0}}=i_{\mathrm{pa}(\gamma)}^{\tilde{0}}} p\left(\left.y_{\gamma}^{\tilde{0}}\right|^{c} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{0}}^{\tilde{0}}\right) \tag{6}
\end{align*}
$$

We see that the product over cases is split up into a product over the configurations of the discrete parents and a product over those cases, where the configuration of the discrete parents is the same as the currently processed configuration. Notice however that some of the parent configurations might not be represented in the database, in which case the product over cases with this parent configuration just adds nothing to the overall product.

In the case with $m$ th order Markov properties, the likelihood terms for all the networks $B_{1}, \ldots, B_{m-1}$, can be written as in (6).

The likelihood part from $B_{\rightarrow}$ is given as,

$$
\begin{align*}
& \prod_{c \in d} \prod_{t=1}^{T} p\left({ }^{c} x^{t} \mid{ }^{c} x^{t-1},{ }^{c} x^{s}, \theta^{\rightarrow}\right) \\
& =\prod_{c \in d} \prod_{t=1}^{T}\left(\prod_{\delta \in \Delta_{t}} p\left({ }^{c} i_{\delta}^{t} \mid{ }^{c} \vec{i}_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}^{\rightarrow}\right) \prod_{\gamma \in \Gamma_{t}} p\left({ }^{c} y_{\gamma}^{t} \mid{ }^{c} y_{\mathrm{pa}(\gamma)},{ }^{c} \dot{\mathrm{a}}_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)\right) \\
& =\prod_{\delta \in \Delta_{t}} \prod_{\underset{\mathrm{pa}(\delta)}{ } \in \prod_{\overrightarrow{\mathrm{I} a}(\delta)}} \prod_{t=1}^{T} \prod_{c: c_{\underset{\mathrm{pa}(\delta)}{ }}=i_{\mathrm{pa}(\delta)}} p\left({ }^{c} i_{\delta}^{t} \mid i_{\mathrm{pa}(\delta)}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)  \tag{7}\\
& \times \prod_{\gamma \in \Gamma} \prod_{\underset{\mathrm{pa}(\gamma)}{\rightarrow} \in \underset{\mathrm{I} \overrightarrow{\mathrm{pa}(\gamma)}}{ } \prod_{t=1}^{T} \prod_{c: c}^{i_{\mathrm{pa}(\gamma)}^{\vec{~}}=i_{\mathrm{pa}(\gamma)}^{\rightarrow}}} p\left(\left.{ }^{c} y_{\gamma}^{t}\right|^{c} y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)
\end{align*}
$$

The product over cases is split up as before. Further, this is also a product over time points, so for each time point $t$, we take the product over cases with a specific configuration of the discrete parents.

Posterior parameter independence now follows from (5), (6) and (7),

$$
\begin{aligned}
p(\theta \mid d) & =p\left(\theta^{\tilde{0}} \mid d\right) p\left(\theta^{\rightarrow} \mid d\right) \\
& =\prod_{\delta \in \Delta_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}} \mid d\right) \prod_{\gamma \in \Gamma_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{o}}} \mid d\right) \\
& \times \prod_{\delta \in \Delta_{t}} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid d\right) \prod_{\gamma \in \Gamma_{t}} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid d\right) .
\end{aligned}
$$

So due to parameter independence and complete data, the parameters stay independent given data. This means that we can learn the parameters in the local distributions independently and also that the parameters in $B_{\tilde{0}}$ and $B_{\rightarrow}$ can be learned independently. Again, if the time series is $m$ th order Markov, posterior parameter independence also follows and we can learn the parameters in $B_{\tilde{0}}, \ldots, B_{m-1}$ and $B_{\rightarrow}$ independently.

Consider for example in $B_{\tilde{0}}$ a parameter for a discrete node $\delta$, with a specific configuration of the discrete parents, $i_{\mathrm{pa}(\delta)}$. The posterior distribution of $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}}$
is by Bayes' theorem found as

Thus $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}}^{\tilde{0}}$ is updated with the cases in the database, where the configuration of the parents of $\delta$ is $i_{\mathrm{pa}(\delta)}^{\tilde{0}}$.
Likewise with a parameter $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ in $B_{\rightarrow}$,

$$
p\left(\theta_{\delta \mid i_{\mathrm{ia}(\delta)}} \mid d\right) \propto \prod_{t=1}^{T} \prod_{c: c_{\mathrm{ia}(\delta)}=i_{\mathrm{pa}(\delta)}} p\left(i_{i}^{t} \mid i_{\mathrm{pa}(\delta)}^{\overrightarrow{ }}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)
$$

Here $\theta_{\delta \mid i_{\mathrm{pa}(\delta)}}$ is, for each time point $t$, updated with the cases in the database for which the configuration of the parents of $\delta$ is $i_{\mathrm{pa}(\delta)}^{\overrightarrow{2}}$.

In the next sections we will introduce the conjugate distributions of the parameters and show how these are learned. The only difference in how the parameters in $B_{\tilde{0}}$ and $B_{\rightarrow}$ are learned, is the set of cases used to learn them. So in the following we do not differentiate between the parameters in $B_{\tilde{0}}$ and $B_{\rightarrow}$.

### 5.2 Learning the Discrete Variables

As described in DeGroot (1970), a conjugate family for multinomial observations is the family of Dirichlet distributions. Let the prior distribution of $\theta_{\delta \mid i_{\mathrm{pa}}(\delta)}$ be a Dirichlet distribution, $\mathcal{D}$, with hyperparameters $\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta} \in \mathcal{I}_{\delta}}$, also written as

$$
\left(\theta_{\delta \mid i_{\mathrm{pa}}(\delta)} \mid \alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}\right) \sim \mathcal{D}\left(\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)
$$

The posterior distribution is then given as

$$
\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}} \mid d\right) \sim \mathcal{D}\left(\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}+n_{\delta \mid i_{\mathrm{pa}(\delta)}}\right)
$$

where the vector $n_{\delta \mid i_{\mathrm{pa}(\delta)}}=\left(n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)_{i_{\delta \in \mathcal{I}_{\delta}}}$, also called the counts, denotes the number of observations in $d$ where $\delta$ and $\mathrm{pa}(\delta)$ have that specific configuration.

Again $\alpha_{\delta \mid i_{\mathrm{pa}(\delta)}}$ and $n_{\delta \mid i_{\mathrm{pa}(\delta)}}$ can be indexed by $\tilde{0}$ and $\rightarrow$, according to $B_{\tilde{0}}$ and $B_{\rightarrow}$. So for $B_{\tilde{0}}$ we have that $n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}^{\tilde{0}}$ is the number of cases in $d$ with a given configuration of $\delta$ and $\mathrm{pa}(\delta)$. Likewise for $B_{\rightarrow}$, where $n_{i \delta \mid i_{\mathrm{pa}(\delta)}}$ is the number of cases in $d$ and for every time point $t=1, \ldots, T$, with this configuration of $\delta$ and $\mathrm{pa}(\delta)$.

### 5.3 Learning the Continuous Variables

For the continuous variables we can write the local probability distributions as

$$
\left(Y_{\gamma} \mid y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) \sim \mathcal{N}\left(z_{\mathrm{pa}(\gamma)}\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)^{\mathrm{T}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)
$$

where $z_{\mathrm{pa}(\gamma)}=\left(1, y_{\mathrm{pa}(\gamma)}\right)$. A standard conjugate family for these observations is the family of Gaussian-inverse gamma distributions. Let the prior joint distribution of $\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)$ and $\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}$ be as follows.

$$
\begin{aligned}
\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right) & \sim \mathcal{N}_{k+1}\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\right) \\
\left(\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right) & \sim \mathcal{I} \Gamma\left(\frac{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}{2}, \frac{\left.\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\right)}{2}\right)
\end{aligned}
$$

If $\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ is a parameter in $B_{\tilde{0}}$, the posterior distribution is found by

$$
p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}} \mid d\right) \propto \prod_{c: c i_{\mathrm{pa}(\gamma)}^{\tilde{0}}=i_{\mathrm{pa}(\gamma)}^{\tilde{0}}} p\left({ }^{c} y_{\gamma}^{\tilde{0}} \mid{ }^{c}{ }_{y}^{\mathrm{pa}(\gamma)} \tilde{0}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{0}}^{\tilde{0}}\right) p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}}\right)
$$

We now join all the observations ${ }^{c} y_{\gamma}^{\tilde{0}}$ for which $i_{\mathrm{pa}(\gamma)}^{\tilde{0}}=i_{\mathrm{pa}(\gamma)}^{\tilde{0}}$ in a vector ${ }^{b} y_{\gamma}^{\tilde{0}}$, i.e. ${ }^{b} y_{\gamma}^{\tilde{0}}=\left({ }^{c} y_{\gamma}^{\tilde{0}}\right)_{c_{\mathrm{pa}(\gamma)}^{\tilde{0}}=}=i_{\mathrm{pa}(\gamma)}^{\tilde{0}}$.

The same is done with the observations of the continuous parents of $\gamma$, i.e. ${ }^{b} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}=\left({ }^{( } y_{\mathrm{pa}(\gamma)}^{\tilde{0}}\right)_{i_{\mathrm{pa}(\gamma)}^{\tilde{0}}=i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}$. The posterior distribution of $\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ can now be written as

$$
p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\tilde{0}} \mid d\right) \propto p\left({ }^{b} y_{\gamma}^{\tilde{0}} \mid{ }^{b} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid \mathrm{i}_{\mathrm{pa}(\gamma)}}^{\tilde{0}}\right) p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\tilde{0}}\right) .
$$

As the distribution, $p\left(c_{y_{\gamma}^{\tilde{0}}}^{\tilde{0}} \mid y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}}\right)$, is a Gaussian distribution, then $p\left({ }^{b} y_{\gamma}^{\tilde{0}} \mid{ }_{y}^{b} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}}\right)$ is a multivariate Gaussian distribution. The covariance matrix is diagonal as all the cases in the database are independent. This way we consider all the cases in a batch.

The same formulation applies for parameters in $B_{\rightarrow \text {. }}$. Notice that the observations included in ${ }^{b} y_{\gamma}$ and ${ }^{b} y_{\mathrm{pa}(\gamma)}$ are taken for each time point $t=1, \ldots, T$.

The posterior distribution is found to be

$$
\begin{aligned}
\left(m_{\gamma \mid i_{\mathrm{pa}(\gamma)},}, \beta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}, d\right) & \sim \mathcal{N}_{k+1}\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}, \sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}\left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\right)^{\prime}\right) \\
\left(\sigma_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2} \mid d\right) & \sim \mathcal{I} \Gamma\left(\frac{\left.\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{2}, \frac{\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}}{2}\right)}{2}\right.
\end{aligned}
$$

where

$$
\begin{aligned}
\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}} z_{\mathrm{pa}(\gamma)}^{b} \\
\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{-1}\left(\tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}} y_{\gamma}^{b}\right) \\
\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+|b| \\
\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}= & \phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{\mathrm{T}} y_{\gamma}^{b} \\
& +\left(\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}-\mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{\prime}\right)^{\mathrm{T}} \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}
\end{aligned}
$$

where $|b|$ denotes the number of observations in $y_{\gamma}^{b}$.

### 5.4 Structure Learning

To learn the structure of a DBN, we again use a Bayesian approach and calculate the posterior probability of a DAG $D$ given data $d$,

$$
\begin{equation*}
p(D \mid d) \propto p(d \mid D) p(D) \tag{8}
\end{equation*}
$$

where $p(d \mid D)$ is the marginal likelihood of $D$ and $p(D)$ is the prior probability of $D$.

In this paper we choose, for simplicity, to let all DAGs be equally likely a priori and therefore we use the measure

$$
p(D \mid d) \propto p(d \mid D)
$$

We refer to the above measure as a network score. We can, in principle, calculate the network score for all possible DAGs and then select the one with the highest score (or, if using model averaging, select a few with high score). In most situations however, there are too many different DAGs to evaluate and some kind of search strategy must be employed, see e.g. Cooper and Herskovits (1992).

The marginal likelihood $p(d \mid D)$ is given as follows.

$$
\begin{aligned}
& p(d \mid D)=\int_{\theta \in \Theta} p(d \mid \theta, D) p(\theta \mid D) d \theta \\
& =\prod_{\delta \in \Delta_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \int_{c: c i_{\mathrm{pa}(\delta)}^{\tilde{0}}=i_{\mathrm{pa}(\delta)}^{0}} p\left(i_{\delta}^{\tilde{0}} \mid i_{\mathrm{pa}(\delta)}^{\tilde{0}}, \theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{0}}^{\tilde{0}}, D\right) p\left(\theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{0}}^{\tilde{0}} \mid D\right) d \theta_{\delta \mid i_{\mathrm{pa}(\delta)}^{0}}^{\tilde{0}} \times
\end{aligned}
$$

$$
\prod_{\gamma \in \Gamma_{\tilde{0}}} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \int_{c: c i_{\mathrm{pa}(\gamma)}^{\tilde{0}}} \prod_{\mathrm{pa}(\gamma)} p\left(\left.{ }^{c} y_{\gamma}^{\tilde{0}}\right|^{c} y_{\mathrm{pa}(\gamma)}^{\tilde{0}}, i_{\mathrm{pa}(\gamma)}^{\tilde{0}}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}}, D\right) p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}} \mid D\right) d \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}^{\tilde{0}}}^{\tilde{0}} \times
$$

$$
\prod_{\gamma \in \Gamma_{t}} \prod_{i_{\mathrm{pa}(\gamma)}^{\rightarrow} \in \underset{\mathrm{I}}{\overrightarrow{\mathrm{pa}(\gamma)}}} \int \prod_{t=1}^{T} \prod_{c: c_{\mathrm{pa}(\gamma)} \rightarrow} p\left(\left.i_{\mathrm{pa}(\gamma)}^{\vec{c}} y_{\gamma}^{t}\right|^{c} y_{\mathrm{pa}(\gamma)}, i_{\mathrm{pa}(\gamma)}, \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}, D\right) p\left(\theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \mid D\right) d \theta_{\gamma \mid i_{\mathrm{pa}(\gamma)}}
$$

We see that the marginal likelihood $p(d \mid D)$ factorizes into a product over terms involving only one node and its parents, called local marginal likelihoods, so the network score is decomposable. This also means that the likelihood factorizes into terms related to $B_{\tilde{0}}$ and terms related to $B_{\rightarrow}$. For $m$ th order Markov time series, the likelihood factorizes in a similar manner into terms related to $B_{\tilde{0}}, \ldots, B_{m-1}$ and $B_{\rightarrow}$.

Because of the way we specified the possible parent sets of variables in $B_{\tilde{0}}$ and in $B_{\rightarrow}$, we can find the best DAG (the one with the highest network score) by finding the best DAG for $B_{\tilde{0}}$ and the best DAG for $B_{\rightarrow \text {. }}$. So we can learn the structure of $B_{\tilde{0}}$ and $B_{\rightarrow}$ independently and we can learn them just as we learn ordinary Bayesian networks with mixed variables as described in Bøttcher (2001). This also applies for $m$ th order Markov time series in which we can learn the structure of $B_{\tilde{0}}, \ldots, B_{m-1}$ and $B_{\rightarrow}$ independently.

In the following we do not distinguish between variables in $B_{\tilde{0}}$ and $B_{\rightarrow}$, as the terms presented apply for both $B_{\tilde{0}}$ and $B_{\rightarrow}$.

The network score contribution from the discrete variables in a network is given by

$$
\begin{equation*}
\prod_{\delta \in \Delta} \prod_{i_{\mathrm{pa}(\delta)} \in \mathcal{I}_{\mathrm{pa}(\delta)}} \frac{\Gamma\left(\alpha_{+\delta} \mid i_{\mathrm{pa}(\delta)}\right)}{\Gamma\left(\alpha_{+\delta}\left|i_{\mathrm{pa}(\delta)}+n_{+\delta}\right|_{\mathrm{pa}(\delta)}\right)} \prod_{i_{\delta} \in \mathcal{I}_{\delta}} \frac{\Gamma\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}+n_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)}{\Gamma\left(\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}\right)} \tag{9}
\end{equation*}
$$

For the continuous variables, the local marginal likelihoods are non-central $t$
distributions with $\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ degrees of freedom, location vector $z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}$ and scale parameter $s_{\gamma \mid i_{\mathrm{pa}(\gamma)}}=\frac{\phi_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}\left(I+\left(z_{\mathrm{pa}(\gamma)}^{b}\right) \tau_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\left(z_{\mathrm{pa}(\gamma)}^{b}\right)^{\mathrm{T}}\right)$. The index $b$ is defined as in Section 5.3.

The network score contribution from the continuous variables is given by

$$
\begin{gather*}
\prod_{\gamma \in \Gamma} \prod_{i_{\mathrm{pa}(\gamma)} \in \mathcal{I}_{\mathrm{pa}(\gamma)}} \frac{\Gamma\left(\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}+|b|\right) / 2\right)}{\Gamma\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} / 2\right)\left[\operatorname{det}\left(\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}} s_{\gamma \mid i_{\mathrm{pa}(\gamma)}} \pi\right)\right]^{\frac{1}{2}}} \times \\
{\left[1+\frac{1}{\rho_{\gamma \mid i_{\mathrm{pa}(\gamma)}}}\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right) s_{\gamma \mid i_{\mathrm{pa}(\gamma)}}^{-1}\left(y_{\gamma}^{b}-z_{\mathrm{pa}(\gamma)}^{b} \mu_{\gamma \mid i_{\mathrm{pa}(\gamma)}}\right)^{\mathrm{T}}\right]^{\frac{-\left(\rho_{\gamma\left|i_{\mathrm{pa}(\gamma)}+|b|\right)}^{2}\right.}{2}} .} \tag{10}
\end{gather*}
$$

The network score is thus the product of (9) and (10).
So if the time series is first order Markov, we can find the best DAG by finding the best DAG for $B_{\tilde{0}}$ and the best DAG for $B_{\rightarrow}$. If it is $m$ th order Markov, we find the best DAGs for $B_{\tilde{0}}, \ldots, B_{m-1}$ and $B_{\rightarrow}$.

### 5.5 Learning the Markov Order

If the Markov order of the time series is unknown, we can learn it by choosing a "prior" order and learn the DBN with this order. The learned order can then be read from the best DAG for $B_{\rightarrow}$, by determining which time slices $X^{t}$ has parents from. The slice furthest back in time will give the order.

It is important that the prior order is chosen high enough to ensure that no order higher than this is better in describing the time series. How high this prior order in practice should be chosen, depends on any prior information available on the time series, but also of how large a dataset the network is learned from. The higher we choose the order, the more complex the possible DAGs are, with more parameters to estimate and fewer cases to learn them from.

To increase the stability of the search procedure, it could therefore be better to start by learning a DBN with a low Markov order. If the best DAG for $B_{\rightarrow}$ include dependencies up to the chosen order, a network with a higher order should be tried and this should be repeated until no dependencies of higher order reveal themselves. However, with this procedure there is a chance that the
best Markov order will not be learned. If e.g. a prior order of three is chosen and the learned network only reveals second order Markov properties, we would with this procedure conclude that the time series is second order Markov, even though the best order could be higher than three. An example of this is shown in Section 7.

Situations can arise, where the Markov order in the initial DAGs is higher than in $B_{\rightarrow}$. For example, if we have assumed that the time series is third order Markov, we need to learn the structure of $B_{\tilde{0}}, B_{1}, B_{2}$ and $B_{\rightarrow}$. Consider now a situation where $B_{\rightarrow}$ is learned to be first order Markov, i.e. $X^{t}$ has only parents in $X^{t}$ and $X^{t-1}$, while $B_{2}$ is learned to be second order Markov, i.e. to have time varying parents from $B_{\tilde{0}}$. This is not necessarily a problem, but it should be noted that if we had assumed the first order Markov property, then there would have been more cases to learn the parameters in $B_{\rightarrow}$ by. In such situations, the importance of specifying the initialization of the time series correctly, must be compared to the loss of precision in the distribution of the parameters in $B_{\rightarrow}$.

## 6 Specifying Prior Distributions

To learn the structure of the DAG we need to specify prior parameter distributions for all possible DAGs under evaluation. An automated procedure for doing this has been developed for ordinary Bayesian networks. We call it the master prior procedure. The procedure is for the purely discrete case treated in Heckerman et al. (1995), for the purely continuous case in Geiger and Heckerman (1994) and for the mixed case in Bøttcher (2001).

We will here give an outline of the procedure and show how it can be used for specifying prior parameter distributions for DBNs.

### 6.1 The Master Prior Procedure

The idea in the master prior procedure is that from a given Bayesian network, we can deduce parameter priors for any possible DAG. The user just has to specify a prior Bayesian network, which is the Bayesian network as he believes it to be. Also, he has to specify an imaginary sample size, $N$, which is a measure of how much confidence he has in the prior network. The procedure works as follows.

1. Specify an imaginary sample size.
2. Specify a prior Bayesian network, i.e. a prior DAG and prior local probability distributions. Calculate the joint prior distribution.
3. From the joint prior distribution and the imaginary sample size, the marginal distribution of all parameters in the family consisting of a node and its parents can be determined. We call this a master prior.
4. The local parameter priors are now determined by conditioning in these master prior distributions.

This procedure ensures parameter independence. Further, it has the property that if a node has the same set of parents in two different networks, then the local parameter prior for this node will be the same in the two networks. Therefore, we only have to deduce the local parameter prior for a node, given the same set of parents, once. This property is called parameter modularity. Finally, the procedure ensures likelihood equivalence, that is, if two DAGs represent the same set of conditional independencies, the network score for these two DAGs will be the same.

As an example, we will show how to deduce parameter priors for the discrete nodes.

Let $\Psi=\left(\Psi_{i}\right)_{i \in \mathcal{I}}$ be the parameters for the joint distribution of the discrete variables. The joint prior parameter distribution is assumed to be a Dirichlet distribution

$$
p(\Psi) \sim \mathcal{D}(\alpha)
$$

with hyperparameters $\alpha=\left(\alpha_{i}\right)_{i \in \mathcal{I}}$. To specify this Dirichlet distribution, we need to specify these hyperparameters. Consider the following relation for the Dirichlet distribution,

$$
p(i)=\mathbb{E}\left(\Psi_{i}\right)=\frac{\alpha_{i}}{N},
$$

with $N=\sum_{i \in \mathcal{I}} \alpha_{i}$. Now we let the probabilities in the prior network be an estimate of $\mathbb{E}\left(\Psi_{i}\right)$, so we only need to determine $N$ in order to calculate the parameters $\alpha_{i}$.

We determine $N$ by using the notion of an imaginary data base. We imagine that we have a database of cases, from which we have updated the distribution of $\Psi$ out of total ignorance. The imaginary sample size of this imaginary data base is thus $N$. It expresses how much confidence we have in the dependency structure expressed in the prior network, see Heckerman et al. (1995).

We use this joint distribution to deduce the master prior distribution of the family $A=\delta \cup \mathrm{pa}(\delta)$. Let

$$
\alpha_{i_{A}}=\sum_{j: j_{A}=i_{A}} \alpha_{j}
$$

and let $\alpha_{A}=\left(\alpha_{i_{A}}\right)_{i_{A} \in \mathcal{I}_{A}}$. Then the marginal distribution of $\Psi_{A}$ is Dirichlet, $p\left(\Psi_{A}\right) \sim \mathcal{D}\left(\alpha_{A}\right)$. This is the master prior in the discrete case. Notice that the parameters in the master prior can also be found as

$$
\alpha_{i_{A}}=N p\left(i_{A}\right)
$$

where $p\left(i_{A}\right)=\sum_{j: j_{A}=i_{A}} p(i)$.
The local parameter priors can now be found by conditioning in these master prior distributions. The conditional distribution of $\Psi_{\delta \mid i_{\mathrm{pa}(\delta)}}$ is

$$
p\left(\Psi_{\delta \mid i_{\mathbf{p a}(\delta)}}\right) \sim \mathcal{D}\left(\alpha_{\left.\delta \mid i_{\mathbf{p}(\delta)}\right)}\right)
$$

with $\alpha_{i_{\delta} \mid i_{\mathrm{pa}(\delta)}}=\alpha_{i_{A}}$.

### 6.2 The Master Prior Procedure for DBNs

For DBNs , the parameter priors can also be found by using the above procedure. Consider a DBN for a first order Markov time series (the procedure is directly extendible to time series with higher order Markov properties). As the DAG from time $t=1$ and forward repeats itself, the structure of the overall DAG is completely specified by the structure of the first two time slices. So we can specify all the parameter priors we need from a prior network consisting of the variables $X^{\tilde{0}}$ and $X^{1}$. Notice that the parameter priors for $B_{\rightarrow}$ are the same as the parameter priors for the parameters in $X^{1}$, as this is the first time point in the time series.

We will also allow for different imaginary sample sizes for the parameters in $B_{\tilde{0}}$ and the parameters in $B_{\rightarrow}$. One reason for this is that the parameters in $B_{\rightarrow}$ are updated with more cases than the parameters in $B_{\tilde{0}}$ and therefore might need a stronger prior distribution.

The procedure works almost as the procedure for ordinary Bayesian networks, the only difference being the different imaginary sample sizes.

1. Specify an imaginary sample size, $N^{\tilde{0}}$, for $B_{\tilde{0}}$, and an imaginary sample size, $N \rightarrow$, for $B \rightarrow$.
2. Specify a prior Bayesian network for the first two time slices. Calculate the joint prior distribution.
3. From the joint prior distribution and the imaginary sample size, the master prior for all parameters in a family can be determined. For families including only variables from $X^{\tilde{0}}$, the imaginary sample size for $B_{\tilde{0}}$ is used and for the other families, the imaginary sample size for $B_{\rightarrow}$ is used.
4. The local parameter priors are now determined by conditioning in the appropriate master prior distribution.

It is obvious that parameter independence and parameter modularity still applies as these properties are not influenced by the use of different imaginary sample sizes. Neither is likelihood equivalence, as variables in $X^{\tilde{0}}$ can not have parents from $X^{1}$. This means that parameter priors for two DAGs that represent the same set of conditional independencies, are calculated using the same imaginary sample sizes. So likelihood equivalence also still applies.

As a simple example of the master prior procedure for DBNs, consider a time series for a single discrete variable $I^{0}, \ldots, I^{T}$. Assume that the time series is first order Markov. The parameter priors for the DAG in Figure 7 are deduced as follows

$$
\begin{aligned}
\alpha_{i^{0}}^{0} & =N^{0} p\left(i^{0}\right), \\
\alpha_{i^{\dagger} \mid i^{t-1}} & =N^{\rightarrow} p\left(i^{0}, i^{1}\right) .
\end{aligned}
$$



Figure 7: DAG for first order Markov time series.

## 7 Example

In this section, we will analyze the Wölfer's sunspot numbers using a dynamic Bayesian network. The Wölfer's sunspot numbers are annual measures of sunspot activity, collected from 1700 to 1988. In statistical terms, the sunspot numbers
is a univariate continuous time series $Y^{0}, \ldots, Y^{288}$. The dataset we use is from Tong (1996).

The sunspot numbers are shown in Figure 8.


Figure 8: Wölfer's sunspot numbers.

Many statistical investigations of these numbers have been made. Anderson (1971) gives a short review of some of these studies. For example, for annual measures of sunspot activity from 1749 to 1924 , Yule (1927) proposed the autoregressive process as a statistical model. He calculated the $A R(p)$ for $p=2$ and $p=5$ and found that an $\operatorname{AR}(2)$ was sufficient, i.e. he estimated the sequence to be second order Markov. Another example is found in Schaerf (1964). She fits an autoregressive model with lags 1,2 , and 9.

Here we will use a DBN as the statistical model and learn the Markov order by structural learning of the DBN. The software package deal, see Bøttcher and Dethlefsen (2003), is used for the analysis.

Our aim is to learn the Markov order, so we are only interested in learning the structure of $B_{\rightarrow}$. The structure of the initial networks is not of interest and


Figure 9: The learned network, $B_{\rightarrow}$, when an Markov order of 30 is assumed. The variables that do not influence $Y^{t}$, have been omitted.
are actually not likely to be determined by learning from the sunspot numbers. These numbers are namely represented by one time series, meaning that for the initial networks there are only one observation of each variable.

As the prior network we use the empty network, i.e. the one without any arrows. In order to get the right location and scale of the parameters, we estimate the prior probability distribution for the empty network from data, i.e. we use the sample mean and the sample variance as the mean and variance in the prior probability distribution.

As the number of observations in the sunspot series is relatively large, we can choose a rather high Markov order for the DBN. Anderson (1971) concludes that the order is not higher than 18. But to be absolutely sure that we capture the best order, we choose an order of 30 . The result of the structural learning of $B_{\rightarrow}$ is shown in Figure 9. The variables that do not influence $Y^{t}$, have been omitted in the figure. From the result we see that the sunspot numbers can be described by a Markov process of order 9 with lags 1, 2 and 9 , i.e.

$$
Y^{t}=m+\beta_{1} Y^{t-1}+\beta_{2} Y^{t-2}+\beta_{9} Y^{t-9}+\epsilon_{t}, \quad \epsilon_{t} \sim \mathcal{N}\left(0, \sigma^{2}\right)
$$

with parameter estimates $m=5.06, \beta_{1}=1.21, \beta_{2}=0.51, \beta_{9}=0.21$ and $\sigma^{2}=267.5$.

The result is in accordance with some of the previous studies, e.g. Schaerf (1964) as mentioned earlier. Other studies determine that an second order Markov process is sufficient, e.g. Yule (1927). But as mentioned, he only examines an order as high as 5 .

We have also tried to learn $B_{\rightarrow}$ using lower Markov order properties. If we e.g. use a Markov order of 3, we reach the conclusion that the sunspot numbers are 2 . order Markov, with lags 1 and 2. This result is shown in Figure 10. Similarly, if we learn $B_{\rightarrow}$ using the order $2, \ldots, 7$ or 8 , we still reach the conclusion that the sunspot numbers are second order Markov, with lags 1 and 2 . This is therefore an example of the importance of choosing the prior Markov order high enough.


Figure 10: The learned network, $B_{\rightarrow}$, when the 3 . order Markov property is assumed. The variable $Y^{t-3}$ have been omitted as it does not influence $Y^{t}$.

As can be seen from Figure 8, the sunspot numbers are periodical with a period of between 10 and 11 years. To determine the period more precisely, we calculate the spectrum,

$$
f(\omega)=\sigma^{2}\left(1-\sum_{t} \beta_{t} e^{-i t \omega}\right)^{-2}
$$

see Venables and Ripley (1997), using the parameter estimates obtained from deal.

The spectrum is shown in Figure 11.


Figure 11: Spectrum of Wölfer's sunspot numbers.

There is a peak at frequency 0.096 , which corresponds to a period of $1 / 0.096=$ 10.40 years. This result is also in accordance with previous studies.

## Acknowledgements

This research was supported by Novo Nordisk A/S. Also I would like to thank Claus Dethlefsen for useful discussions and his help with the implementation of the example. Finally, I thank my supervisor Steffen L. Lauritzen for many valuable comments.

## References

Anderson, T. W. (1971). The Statistical Analysis of Time Series, John Wiley and Sons, New York.

Bøttcher, S. G. (2001). Learning Bayesian Networks with Mixed Variables, Artificial Intelligence and Statistics 2001, Morgan Kaufmann, San Francisco, CA, USA, pp. 149-156.

Bøttcher, S. G. and Dethlefsen, C. (2003). deal: A Package for Learning Bayesian Networks, Journal of Statistical Software 8(20): 1-40.

Bøttcher, S. G., Milsgaard, M. B. and Mortensen, R. S. (1995). Monitoring by using dynamic linear models - illustrated by tumour markers for cancer, Master's thesis, Aalborg University.

Cooper, G. and Herskovits, E. (1992). A Bayesian method for the induction of probabilistic networks from data, Machine Learning 9: 309-347.

Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation, Computational Intelligence 5: 142-150.

DeGroot, M. H. (1970). Optimal Statistical Decisions, McGraw-Hill, New York.

Friedman, N., Murphy, K. P. and Russell, S. (1998). Learning the Structure of Dynamic Probabilistic Networks, Proceedings of Fourteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 139-147.

Geiger, D. and Heckerman, D. (1994). Learning Gaussian Networks, Proceedings of Tenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, CA, USA, pp. 235-243.

Harrison, P. J. and Stevens, C. F. (1976). Bayesian forecasting, Journal of Royal Statistics 38: 205-247.

Heckerman, D., Geiger, D. and Chickering, D. M. (1995). Learning Bayesian networks: The combination of knowledge and statistical data, Machine Learning 20: 197-243.

Lauritzen, S. L. (1996). Graphical Models, Clarendon press, Oxford, New York.
Murphy, K. P. (2002). Dynamic Bayesian Networks: Representation, Inference and Learning, PhD thesis, University of California, Berkeley.

Schaerf, M. C. (1964). Estimation of the covariance and autoregressive structure of a stationary time series, Technical report, Department of Statistics, Stanford University.

Tong, H. (1996). Non-Linear Time Series, Clarendon Press, Oxford.
Venables, W. N. and Ripley, B. D. (1997). Modern Applied Statistics with SPLUS, second edn, Springer-Verlag, New York.

West, M. and Harrison, J. (1989). Bayesian Forecasting and Dynamic Models, Springer-Verlag, New York.

Yule, G. U. (1927). On a method for investigating periodicities in disturbed series with special reference to Wölfer's sunspot numbers, Philosophical Transactions of the Royal Society, Series A 226: 267-298.

