

**DEAL: A Package for Learning Bayesian
Networks**

by

Susanne G. Bøttcher and Claus Dethlefsen

February 2003

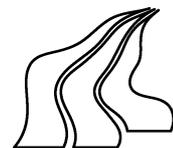
R-2003-03

**DEPARTMENT OF MATHEMATICAL SCIENCES
AALBORG UNIVERSITY**

Fredrik Bajers Vej 7 G ■ DK - 9220 Aalborg Øst ■ Denmark

Phone: +45 96 35 80 80 ■ Telefax: 45 98 15 81 29

URL: www.math.auc.dk/research/reports/reports.htm



DEAL: A Package for Learning Bayesian Networks

Susanne G. Bøttcher

Dept. of Mathematical Sciences
Aalborg University
Fr. Bajers Vej 7G
9220 Aalborg, Denmark
alma@math.auc.dk

Claus Dethlefsen

Dept. of Mathematical Sciences
Aalborg University
Fr. Bajers Vej 7G
9220 Aalborg, Denmark
dethlef@math.auc.dk

February 10, 2003

Abstract

DEAL is a software package freely available for use with R. It includes several methods for analysing 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 (Ihaka and Gentleman, 1996), 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 freely for non-commercial purposes.

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 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.

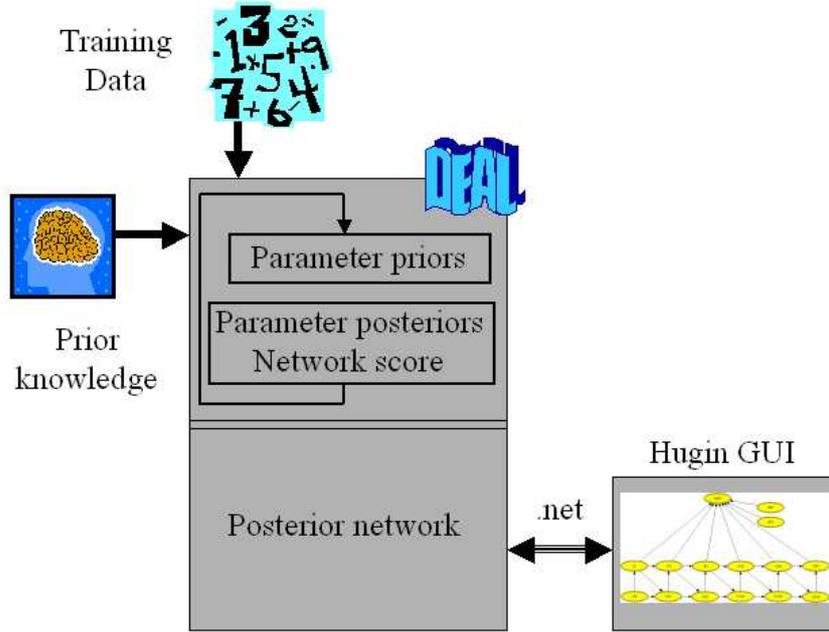


Figure 1: From prior knowledge and training data, a posterior network is produced by *DEAL*. The network may be transferred to Hugin for further inference.

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 $X = (X_v)_{v \in V}$. Often, we do not distinguish between a variable X_v and the corresponding node v . To each node v with parents $\text{pa}(v)$ a local probability distribution, $p(x_v | x_{\text{pa}(v)})$, 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,

$$p(x) = \prod_{v \in V} p(x_v | x_{\text{pa}(v)}).$$

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 Δ and Γ are the sets of discrete and continuous nodes, respectively. The set of variables X can then be denoted $X = (X_v)_{v \in V} = (I, Y) = ((I_\delta)_{\delta \in \Delta}, (Y_\gamma)_{\gamma \in \Gamma})$, where I and Y are the sets of discrete and continuous variables, respectively. For a discrete variable, δ , we let \mathcal{L}_δ 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

$$p(x) = p(i, y) = \prod_{\delta \in \Delta} p(i_\delta | i_{\text{pa}(\delta)}) \prod_{\gamma \in \Gamma} p(y_\gamma | i_{\text{pa}(\gamma)}, y_{\text{pa}(\gamma)}).$$

3 Data Structure

DEAL expects data as specified in a data frame which is a standard data structure in R. For example, the data could be stored in an ASCII file with one column per variable and one line per observation, such as the rats example in Table 1, analysed in Morrison (1976) and Edwards (1995). The first line in the data file should be the names of the variables.

The rats 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
1	1	5	6
1	1	7	6
1	1	9	9
1	1	5	4
1	2	9	12
1	2	7	7
1	2	7	6
1	2	6	8
1	3	14	11
1	3	21	15
1	3	12	10
1	3	17	12
2	1	7	10
2	1	8	10
2	1	6	6
2	1	9	7
2	2	7	6
2	2	10	13
2	2	6	9
2	2	8	7
2	3	14	9
2	3	14	8
2	3	16	12
2	3	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*. Default, all variables in the rats example are set to numeric by `read.table()`. 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 discrete variables and should be altered as

```
rats.df$Sex <- factor(rats.df$Sex)
rats.df$Drug <- factor(rats.df$Drug)
```

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 has several attributes, added or changed by methods described in following sections.

A network is generated by the following command

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

and 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.

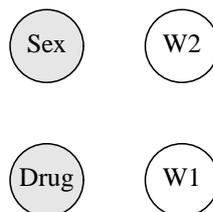


Figure 2: Graphical representation of the rats network.

The primary attribute of a network is the list of nodes, in the example: `rats$nodes`. Each entry in the list is an object of class `node` representing a 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.

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$nodes          # the list of nodes
rats$nodes[[1]]    # the first node
rats$nodes$Drug     # the node ``Drug``

```

A collection of networks can be represented in an object of class `networkfamily`. Operations on a network family is thus applied to all the networks in the family.

4.2 Specification of the Probability Distributions

The joint distribution of the random variables in a network in *DEAL* is a CG distribution.

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

$$\theta_{i_\delta|i_{\text{pa}(\delta)}} = p(i_\delta|i_{\text{pa}(\delta)}, \theta_{\delta|i_{\text{pa}(\delta)}}),$$

where $\theta_{\delta|i_{\text{pa}(\delta)}} = (\theta_{i_\delta|i_{\text{pa}(\delta)}})_{i_\delta \in \mathcal{I}_\delta}$. The parameters fulfil $\sum_{i_\delta \in \mathcal{I}_\delta} \theta_{i_\delta|i_{\text{pa}(\delta)}} = 1$ and $0 \leq \theta_{i_\delta|i_{\text{pa}(\delta)}} \leq 1$.

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

$$\theta_{\gamma|i_{\text{pa}(\gamma)}} = (m_{\gamma|i_{\text{pa}(\gamma)}}, \beta_{\gamma|i_{\text{pa}(\gamma)}}, \sigma_{\gamma|i_{\text{pa}(\gamma)}}^2),$$

so that

$$(Y_\delta|i_{\text{pa}(\gamma)}, y_{\text{pa}(\gamma)}, \theta_{\gamma|i_{\text{pa}(\gamma)}}) \sim \mathcal{N}(m_{\gamma|i_{\text{pa}(\gamma)}} + y_{\text{pa}(\gamma)}\beta_{\gamma|i_{\text{pa}(\gamma)}}, \sigma_{\gamma|i_{\text{pa}(\gamma)}}^2).$$

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

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

$$p(i_\delta|i_{\text{pa}(\delta)}) = 1/\mathcal{I}_\delta.$$

Define $z_{\text{pa}(\gamma)} = (1, y_{\text{pa}(\gamma)})$ and let $\eta_{\gamma|i_{\text{pa}(\gamma)}} = (m_{\gamma|i_{\text{pa}(\gamma)}}, \beta_{\gamma|i_{\text{pa}(\gamma)}})$, where $m_{\gamma|i_{\text{pa}(\gamma)}}$ is the intercept and $\beta_{\gamma|i_{\text{pa}(\gamma)}}$ is the vector of coefficients. For a continuous variable γ , the suggested local probability distribution $\mathcal{N}(z_{\text{pa}(\gamma)}\eta_{\gamma|i_{\text{pa}(\gamma)}}, \sigma_{\gamma|i_{\text{pa}(\gamma)}}^2)$ is determined as a regression on the continuous parents for each configuration of the discrete parents.

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

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.*

```

rats$nodes$Sex$prob <- c(0.6, 0.4)
rats$nodes$W1$prob  <- c(10, 0)
rats$nodes$W2$prob  <- 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

$$p(i) = \prod_{\delta \in \Delta} p(i_\delta | i_{\text{pa}(\delta)}).$$

For continuous variables, the joint distribution $\mathcal{N}(M_i, \Sigma_i)$ 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 γ and all previously processed nodes. From the prior network, we have given $\eta_{\gamma | \text{pa}(\gamma)} = (m_{\gamma | \text{pa}(\gamma)}, \beta_{\gamma | \text{pa}(\gamma)})$ and $\sigma_{\gamma | \text{pa}(\gamma)}^2$. Previously evaluated are $M_{\text{pa}(\gamma)}$ and $\Sigma_{\text{pa}(\gamma)}$. Now, the covariance is given by

$$\Sigma_{\gamma, \text{pa}(\gamma)} = \Sigma_{\text{pa}(\gamma)} \beta_{\gamma | \text{pa}(\gamma)}.$$

The variance and mean are then given by

$$\begin{aligned} \Sigma_\gamma &= \sigma_{\gamma | \text{pa}(\gamma)}^2 + \Sigma_{\gamma, \text{pa}(\gamma)} \beta_{\gamma | \text{pa}(\gamma)} \\ M_\gamma &= m_{\gamma | \text{pa}(\gamma)} + \beta_{\gamma | \text{pa}(\gamma)} M_{\text{pa}(\gamma)}. \end{aligned}$$

In *DEAL*, we can assess these quantities by

```
rats.j <- jointprior(rats)
rats.j$jointmu
rats.j$jointsigma
rats.j$jointalpha
```

The discrete part, $p(i)$, is not returned directly but may be deduced from `rats.j$jointalpha` by division by `sum(rats.j$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 the Bayesian approach. We encode our uncertainty about parameters θ in a prior distribution $p(\theta)$, use data d to update this distribution, and hereby obtain the posterior distribution $p(\theta|d)$ by using Bayes' theorem,

$$p(\theta|d) = \frac{p(d|\theta)p(\theta)}{p(d)}, \quad \theta \in \Theta. \quad (1)$$

Here Θ is the parameter space, d is a random sample from the probability distribution $p(x|\theta)$ and $p(d|\theta)$ is the joint probability distribution of d , also called the likelihood of θ . 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.*

$$p(\theta) = \prod_{\delta \in \Delta} \prod_{i_{\text{pa}(\delta)} \in \mathcal{I}_{\text{pa}(\delta)}} p(\theta_{\delta} | i_{\text{pa}(\delta)}) \prod_{\gamma \in \Gamma} \prod_{i_{\text{pa}(\gamma)} \in \mathcal{I}_{\text{pa}(\gamma)}} p(\theta_{\gamma} | i_{\text{pa}(\gamma)}), \quad (2)$$

We refer to (2) as *parameter independence*. Further, as we have assumed complete data, the parameters stay independent given data, see Böttcher (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* $p(\theta_{v|i_{\text{pa}(v)}})$ 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 Böttcher (2001), for the pure discrete and the pure 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.

5.3.1 Master Prior

Let $\Psi = (\Psi_i)_{i \in \mathcal{I}}$ 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 $\alpha = (\alpha_i)_{i \in \mathcal{I}}$. To specify this Dirichlet distribution, we need to specify these hyperparameters.

Consider the following relation for the Dirichlet distribution,

$$p(i) = \mathbb{E}(\Psi_i) = \frac{\alpha_i}{N},$$

with $N = \sum_{i \in \mathcal{I}} \alpha_i$. Now we use the probabilities in the prior network as an estimate of $\mathbb{E}(\Psi_i)$, so we only need to determine N in order to calculate the parameters α_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 Ψ 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 $A = \delta \cup \text{pa}(\delta)$. Let

$$\alpha_{i_A} = \sum_{j: j_A = i_A} \alpha_j,$$

and let $\alpha_A = (\alpha_{i_A})_{i_A \in \mathcal{I}_A}$. Then the marginal distribution of Ψ_A is Dirichlet, $p(\Psi_A) \sim \mathcal{D}(\alpha_A)$. This is the master prior in the discrete case.

5.3.2 Local Parameter Prior

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

$$\begin{aligned} \alpha_{i_{\delta|i_{\text{pa}(\delta)}}} &= \alpha_{i_A}, \\ \alpha_{\delta|i_{\text{pa}(\delta)}} &= (\alpha_{i_{\delta|i_{\text{pa}(\delta)}}})_{i_{\delta} \in \mathcal{I}_{\delta}}, \\ \theta_{\delta|i_{\text{pa}(\delta)}} | \alpha_{i_{\delta|i_{\text{pa}(\delta)}}} &\sim \mathcal{D}(\alpha_{\delta|i_{\text{pa}(\delta)}}). \end{aligned}$$

5.3.3 Local Parameter Posterior

Let $n_{\delta|i_{\text{pa}(\delta)}}$ 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 $\alpha'_{\delta|i_{\text{pa}(\delta)}}$ are given by

$$\alpha'_{\delta|i_{\text{pa}(\delta)}} = \alpha_{\delta|i_{\text{pa}(\delta)}} + n_{\delta|i_{\text{pa}(\delta)}}.$$

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.

5.4.1 Master Prior

Böttcher (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 α_i was determined in Section 5.3.1. Also, $\Phi_i = (\nu_i - 1)\Sigma_i$.

The joint parameter priors are assumed to be distributed as

$$\begin{aligned} p(M_i|\Sigma_i) &= \mathcal{N}\left(\mu_i, \frac{1}{\nu_i}\Sigma_i\right) \\ p(\Sigma_i) &= \mathcal{IW}(\rho_i, \Phi_i). \end{aligned}$$

We cannot use these distributions to derive priors for other networks, so 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 $A = \gamma \cup \text{pa}(\gamma)$, the local master prior is then found as

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

where

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

5.4.2 Local Parameter Prior

The local prior parameters, given as

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

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 $i_{A\cap\Delta} = i_{\text{pa}(\gamma)}$ and assume that $A \cap \Gamma$ is ordered with γ as the first entry. Write $\text{pa}(\gamma)$ for the continuous parents $\{A \cap \Gamma\} \setminus \{\gamma\}$. Define the partitioning

$$\tilde{\Phi}_{A\cap\Gamma|i_{A\cap\Delta}} = \begin{bmatrix} \tilde{\phi}_\gamma & \tilde{\Phi}_{\gamma, \text{pa}(\gamma)} \\ \tilde{\Phi}_{\text{pa}(\gamma), \gamma} & \tilde{\Phi}_{\text{pa}(\gamma)} \end{bmatrix}$$

$$\bar{\mu}_{i_{A \cap \Delta}} = \left(\bar{\mu}_\gamma, \bar{\mu}_{\text{pa}(\gamma)} \right).$$

Then

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

5.4.3 Local Parameter Posterior

Define $z_{\text{pa}(\gamma)|i_{\text{pa}(\gamma)}}^b$ 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 $y_{\gamma|i_{\text{pa}(\gamma)}}^b$ be the vector of observations of the node γ for a configuration of the discrete parents.

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

$$\begin{aligned} \tau'_{\gamma|i_{\text{pa}(\gamma)}} &= \tau_{\gamma|i_{\text{pa}(\gamma)}} + (z_{\text{pa}(\gamma)|i_{\text{pa}(\gamma)}}^b)^\top z_{\text{pa}(\gamma)|i_{\text{pa}(\gamma)}}^b \\ \mu'_{\gamma|i_{\text{pa}(\gamma)}} &= (\tau'_{\gamma|i_{\text{pa}(\gamma)}})^{-1} \times \left(\tau_{\gamma|i_{\text{pa}(\gamma)}} \mu_{\gamma|i_{\text{pa}(\gamma)}} + (z_{\text{pa}(\gamma)|i_{\text{pa}(\gamma)}}^b)^\top y_{\gamma|i_{\text{pa}(\gamma)}}^b \right) \\ \rho'_{\gamma|i_{\text{pa}(\gamma)}} &= \rho_{\gamma|i_{\text{pa}(\gamma)}} + n \\ \phi'_{\gamma|i_{\text{pa}(\gamma)}} &= \phi_{\gamma|i_{\text{pa}(\gamma)}} \\ &\quad + (y_{\gamma|i_{\text{pa}(\gamma)}}^b - z_{\text{pa}(\gamma)}^b \mu'_{\gamma|i_{\text{pa}(\gamma)}})^\top y_{\gamma|i_{\text{pa}(\gamma)}}^b \\ &\quad + (\mu_{\gamma|i_{\text{pa}(\gamma)}} - \mu'_{\gamma|i_{\text{pa}(\gamma)}})^\top \tau_{\gamma|i_{\text{pa}(\gamma)}} \mu_{\gamma|i_{\text{pa}(\gamma)}}. \end{aligned}$$

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,

```

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

```

The result is attached to each learned node as two attributes. These contain the parameters in the local prior distribution and the parameters in the local posterior distribution, respectively. For the node `Sex`, the attributes are assessed as

```

rats$nodes$Sex$condprior
rats$nodes$Sex$condposterior

```

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

$$S(D) = p(D, d) = p(d|D)p(D),$$

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 δ and $S_{\gamma}(D)$ is the contribution from the continuous node γ .

For a discrete node, δ , the score contribution is given by

$$S_{\delta}(D) = \prod_{i_{\text{pa}(\delta)} \in \mathcal{I}_{\text{pa}(\delta)}} \frac{\Gamma(\alpha_{+\delta|i_{\text{pa}(\delta)}})}{\Gamma(\alpha_{+\delta|i_{\text{pa}(\delta)}} + n_{+\delta|i_{\text{pa}(\delta)}})} \times \prod_{i_{\delta} \in \mathcal{I}_{\delta}} \frac{\Gamma(\alpha_{i_{\delta}|i_{\text{pa}(\delta)}} + n_{i_{\delta}|i_{\text{pa}(\delta)}})}{\Gamma(\alpha_{i_{\delta}|i_{\text{pa}(\delta)}})},$$

where $\alpha_{+\delta|i_{\text{pa}(\delta)}} = \sum_{i_{\delta} \in \mathcal{I}_{\delta}} \alpha_{i_{\delta}|i_{\text{pa}(\delta)}}$ and $n_{+\delta|i_{\text{pa}(\delta)}} = \sum_{i_{\delta} \in \mathcal{I}_{\delta}} n_{i_{\delta}|i_{\text{pa}(\delta)}}$.

For a continuous node, γ ,

$$S_\gamma(D) = \prod_{i_{\text{pa}(\gamma)} \in \mathcal{I}_{\text{pa}(\gamma)}} \frac{\Gamma\left(\frac{\rho_\gamma | i_{\text{pa}(\gamma)} + n}{2}\right)}{\Gamma\left(\frac{\rho_\gamma | i_{\text{pa}(\gamma)}}{2}\right) \sqrt{\det(\rho_\gamma | i_{\text{pa}(\gamma)} S_\gamma | i_{\text{pa}(\gamma)} \pi)}} \times \left[1 + \frac{1}{\rho_\gamma | i_{\text{pa}(\gamma)}} a_{\gamma | i_{\text{pa}(\gamma)}} S_{\gamma | i_{\text{pa}(\gamma)}}^{-1} a_{\gamma | i_{\text{pa}(\gamma)}}^\top \right]^{-\frac{\rho_\gamma | i_{\text{pa}(\gamma)} + n}{2}},$$

where

$$s_{\gamma | i_{\text{pa}(\gamma)}} = \frac{\phi_{\gamma | i_{\text{pa}(\gamma)}}}{\rho_{\gamma | i_{\text{pa}(\gamma)}}} \left(I + z_{\text{pa}(\gamma)}^b T_{\gamma | i_{\text{pa}(\gamma)}}^{-1} (z_{\text{pa}(\gamma)}^b)^\top \right)$$

$$a_{\gamma | i_{\text{pa}(\gamma)}} = y_{\gamma | i_{\text{pa}(\gamma)}}^b - z_{\text{pa}(\gamma) | i_{\text{pa}(\gamma)}}^b \mu_{\gamma | i_{\text{pa}(\gamma)}}.$$

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 DAG's 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 <- learn(rats,rats.df,rats.prior)$nw
rats$nodes$Sex$loglik
rats$score # log network score
```

6.2 Model Search

In principle, we could evaluate the network score for all possible DAG's and indeed this is provided in *DEAL*

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

However, the number of possible DAG's grows more than exponential 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 the number of random variables in a network is large, it is not computationally possible to calculate the network score for all the possible DAG's. For these situations a strategy for searching for DAG's 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 DAG's, D and D^* , we use the posterior odds,

$$\frac{p(D|d)}{p(D^*|d)} = \frac{p(D,d)}{p(D^*,d)} = \frac{p(D)}{p(D^*)} \times \frac{p(d|D)}{p(d|D^*)},$$

where $p(D)/p(D^*)$ is the prior odds and $p(d|D)/p(d|D^*)$ is the Bayes' factor. At the moment, the only option in *DEAL* for specifying prior distribution over DAG's is to let all DAG's be equally likely, so

# 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 prior odds are always equal to one. Therefore, we use the Bayes’ factor for comparing two different DAG’s.

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 <- autosearch(rats,rats.df,rats.prior)$nw
```

returns all tried networks in a greedy search from the initial network `rats`, which may be constructed using `drawnetwork()`.

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

```
rats <- drawnetwork(rats,rats.df,rats.prior)$nw
```

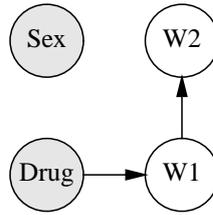


Figure 3: The network with the highest score in the rats example.

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

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 calling `perturb()` directly

```
rats.rn <- 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 $v \leftrightarrow w$ 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 <- heuristic(rats,rats.df,rats.prior,trylist=rats.tl)$nw
```

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 networks but cannot learn either pure 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 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 file 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 analyse a large data set which includes both discrete and continuous variables. The *ksl* 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.

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

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.

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 $\log(\text{BMI})$. 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)

## transform BMI to log(BMI)
ksl$BMI <- log(ksl$BMI)
names(ksl)[4] <- "logBMI"

## create factors for discrete variables
ksl$Sex <- factor(ksl$Sex)
ksl$Year <- factor(ksl$Year)
ksl$Smok <- factor(ksl$Smok)
ksl$Alc <- factor(ksl$Alc)
ksl$Work <- factor(ksl$Work)
```

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
banlist <- matrix(c(5,5,6,6,7,7,9,
                   8,9,8,9,8,9,8),ncol=2)
ksl.nw$banlist <- banlist
```

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 <- learn(ksl.nw,ksl,ksl.prior)$nw

## Do structural search
thebest <- autosearch(ksl.nw,ksl,ksl.prior,saveall=FALSE,trace=TRUE)$nw[[1]]

## perturb 'thebest' and rerun search twice.
hiscorelist <- heuristic(thebest,ksl,ksl.prior,restart=2,degree=10,trace=TRUE)

thebest2 <- hiscorelist$nw[[1]]

savenet(thebest2, "ksl.net")
```

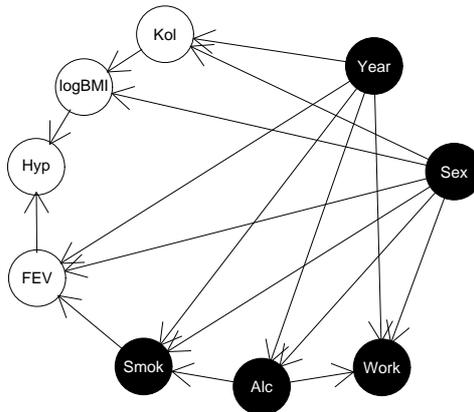


Figure 4: The network with the highest score, $\log(\text{score}) = -15957.91$.

The resulting network `thebest2` 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

$\text{Smok} \leftarrow \text{Alc} \rightarrow \text{Work}$ in the given DAG represents the same probability distribution as the relations $\text{Smok} \leftarrow \text{Alc} \leftarrow \text{Work}$ and $\text{Smok} \rightarrow \text{Alc} \rightarrow \text{Work}$, *i.e.* the three DAG's 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 conditional 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 analysing data. In particular, *DEAL* is part of the gR 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 analysing 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 BDe criteria. We intend to also incorporate the Bayesian Information Criteria (BIC) and Akaike's 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.

References

- J.H. Badsberg. *An Environment for Graphical Models*. PhD thesis, Aalborg University, 1995.
- S.G. Bøttcher. Learning Bayesian networks with mixed variables. In *Proceedings of the Eighth International Workshop in Artificial Intelligence and Statistics*, 2001.
- S.G. Bøttcher. Learning Bayesian networks with mixed variables. URL <http://www.math.auc.dk/~alma>. Aalborg University, 2003.
- D.M. Chickering. *Learning from Data: Artificial Intelligence and Statistics V*, chapter Learning Bayesian networks is NP-complete, pages 121–130. Springer-Verlag, 1996.
- D.M. Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, November 2002.
- R.G. Cowell, A.P. Dawid, S.L. Lauritzen, and D.J. Spiegelhalter. *Probabilistic Networks and Expert Systems*. Springer-Verlag, 1999.
- D. Edwards. *Introduction to Graphical Modelling*. Springer-Verlag, New York, 1995.
- D. Geiger and D. Heckerman. Learning Gaussian networks. Technical Report MSR-TR-94-10, Microsoft Research, 1994.
- D. Heckerman, D. Geiger, and D. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. *Machine Learning*, 1995.
- R. Ihaka and R. Gentleman. R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics*, 5:299–314, 1996.
- S.L. Lauritzen. Propagation of probabilities, means and variances in mixed graphical association models. *Journal of the American Statistical Association*, 82:1082–1108, 1992.
- S.L. Lauritzen. Some modern applications of graphical models. In P.J. Green, N.L. Hjort, and S. Richardson, editors, *Highly Structured Stochastic Systems*. Oxford University Press, 2003.
- D.F. Morrison. *Multivariate Statistical Methods*. McGraw-Hill, USA, 1976.
- R.D. Shachter and C.R. Kenley. Gaussian influence diagrams. *Management Science*, 35:527–550, 1989.

A Manual Pages for *DEAL*

autosearch *Greedy search*

Description

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

Usage

```
autosearch(initnw,data,prior=jointprior(network(data)),maxiter=50,
           trylist= rep(list(NULL),initnw$n),trace=FALSE,
           timetrace=FALSE,showban=FALSE,saveall=TRUE)
```

```
heuristic(initnw,data,prior=jointprior(network(data)),
           maxiter=100,restart=10,degree=initnw$n,
           trylist= rep(list(NULL),initnw$n),trace=FALSE,
           timetrace=TRUE,saveall=TRUE)
```

Arguments

initnw	The initial network from which the search is started.
data	Data frame used for learning the network, see <code>network</code> .
prior	A joint prior, see <code>jointprior</code> .
maxiter	Maximum search steps in the search algorithm.
restart	The number of times to perturb <code>initnw</code> and rerun the search.
degree	Degree of perturbation, see <code>perturb</code> .
trylist	Used internally for reusing learning of nodes, see <code>maketrylist</code> .
trace	If TRUE, plots the accepted networks during search.
timetrace	If TRUE, prints some timing information on the screen.
showban	Passed to <code>plot.network</code> .
saveall	If TRUE, all networks explored in the search is returned (memory consuming). Otherwise, only the best is returned.

Details

In `autosearch`, each step, a list of networks is 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 networks has a higher network score than the previous network, the search is terminated. The returned network list contains all tried networks (depending on the value of `saveall`).

`heuristics` restarts by perturbing `initnw` degree times and calling `autosearch` again. The number of restarts is given by the option `restart`.

Value

A list with two elements

<code>nw</code>	A network family with all tried networks or only the network with the highest score if <code>saveall</code> is <code>FALSE</code> . Duplicate networks have not been removed (use <code>nwfunique</code>).
<code>trylist</code>	Updated, learn.

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 <- learn(fit,rats,fit.prior)$nw
fit <- insert(fit,2,1,rats,fit.prior)$nw
fit <- insert(fit,1,3,rats,fit.prior)$nw
hisc <- autosearch(fit,rats,fit.prior)$nw
hisc <- nwfsort(hisc)
```

Description

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

Usage

```
drawnetwork(nw,df,prior,trylist=rep(list(NULL),nw$n),
            scale=10,unitscale=20,cexscale=8,
            arrowlength=.25,nocalc=FALSE,
            yr=c(0,350),xr=yr,...)
```

```
inspectprob(nw,scale=10,unitscale=20,cexscale=8,
            arrowlength=.25,xr=c(0,350),yr=xr,...)
```

Arguments

nw	The initial network, see network.
df	Data frame used for learning the network, see network.
prior	A joint prior, see jointprior.
trylist	Used internally for reusing learning of nodes, see maketrylist.
cexscale	Passed to plot.network.
arrowlength	Passed to plot.network.
nocalc	Used when no learning is wanted, see eg. simulation.
unitscale	Passed to plot.network.
scale	Passed to plot.network.
xr	Range on x-axis.
yr	Range on y-axis.
...	Passed to plot.network.

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 attribute 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`).

`inspectprob` draws the network and makes it possible to inspect the `prob` attributes of the nodes by clicking on them. The result is shown in the output window.

Value

A list with two elements

`nw` A network family with all created networks.
`trylist` Updated, see `learn`.

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 <- learn(rats.nw,rats,rats.prior)$nw

newrat <- drawnetwork(rats.nw,rats,rats.prior)$nw
```

`jointprior` *Calculates the joint prior and the quantities needed to specify local parameter priors*

Description

Given a network with a `prob` attribute for each node, the joint probability distribution is derived. 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)
jointdisc(nw,timetrace=FALSE)
jointcont(nw,timetrace=FALSE)
```

Arguments

<code>nw</code>	A network. Each node must have a <code>prob</code> attribute to describe the local probability distribution, see <code>network</code> .
<code>N</code>	The size of the imaginary data base. If this is too small, errors might occur. If no <code>N</code> is given, the procedure tries to set a value as low as possible.
<code>phiprior</code>	The way of specifying the prior for ϕ . Either <code>phiprior="bottcher"</code> or <code>phiprior="heckerman"</code> can be used.
<code>timetrace</code>	If <code>TRUE</code> , prints some timing information on the screen.

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 `jointsigma` (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 e.g. Shachter and Kenley (1989).

Then, `jointalpha`, `jointnu`, `jointrho`, `mu` and `jointphi` are deduced. These quantities are later used for deriving local parameter priors.

For the configuration `i` of the discrete variables,

$$\nu_i = \rho_i = \alpha_i$$

and

$$\phi_i = (\nu_i - 1)\Sigma_i$$

if `phiprior="bottcher"` and

$$\phi_i = \nu_i(\rho_i - 2)\Sigma_i/(\nu_i + 1)$$

if `phiprior="heckerman"`.

Value

A list with the following elements

<code>jointalpha</code>	Used in local master procedure for discrete variables.
-------------------------	--

<code>jointnu</code>	Used in local master procedure for continuous variables.
<code>jointrho</code>	Used in local master procedure for continuous variables.
<code>jointmu</code>	Used in local master procedure for continuous variables.
<code>jointsigma</code>	Sigma matrices (not used in further calculations).
<code>jointphi</code>	Used in local master procedure for continuous variables.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
 Claus Dethlefsen (dethlef@math.auc.dk).

References

Shachter and Kenley (1989), Gaussian influence diagrams. *Management Science* 35:527–550.

See Also

`network`

Examples

```
data(rats)
rats.nw <- network(rats)
rats.prior <- jointprior(rats.nw,12)
```

<code>learn</code>	<i>Estimates the parameters in the local probability distributions from data</i>
--------------------	--

Description

Update 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:nw$n, trylist=
      rep(list(NULL),nw$n),
      timetrace=FALSE,
      usetrylist = TRUE
      )
learnnode(node,nw,df,prior=jointprior(nw),timetrace=FALSE)
```

```
udisclik(node,nw,df)
reuselearn(node,trylist)
```

Arguments

<code>nw</code>	The network.
<code>df</code>	Data frame used for learning the network, see <code>network</code> .
<code>prior</code>	A joint prior, see <code>jointprior</code> .
<code>odelist</code>	A vector of indices of nodes to be learned.
<code>trylist</code>	Used internally for reusing learning of nodes, see <code>maketrylist</code> .
<code>timetrace</code>	If <code>TRUE</code> , prints some timing information on the screen.
<code>node</code>	Node to be learned.
<code>usetrylist</code>	Use <code>trylist</code> or not. If the number of nodes is large, the look up in the <code>trylist</code> may be slower than re-learning the node.

Details

The procedure `learn` determines the master prior, local parameter priors and local parameter posteriors. It may be called on all nodes 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 attribute `condprior`.

Finally, the local parameter posterior distributions are calculated (see `post`) and attached to each node in the attribute `condposterior`.

A so-called `trylist` is maintained. It consists of learned nodes with particular parent configurations. 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. The learning procedure calls `reuselearn` which traverses through the `trylist` to see if it is possible to reuse previously learned nodes.

When one or more nodes in a network have been learned, the network score is updated and attached to the network in the attribute `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

<code>nw</code>	The network with the <code>condposterior</code> attributes updated for the nodes. Also, the attribute <code>score</code> is updated and contains the network score. The contribution to the network score for each node is contained in the attribute <code>loglik</code> for each node.
<code>trylist</code>	Updated.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
 Claus Dethlefsen (dethlef@math.auc.dk).

See Also

`networkfamily`, `jointprior`, `maketrylist`, `network`, `post`

Examples

```
data(rats)
fit      <- network(rats)
fit.prior <- jointprior(fit,12)
fit      <- learn(fit,rats,fit.prior)$nw
```

<code>maketrylist</code>	<i>Creates the full trylist</i>
--------------------------	---------------------------------

Description

For faster learning, a trylist is maintained.

Usage

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

Arguments

<code>initnw</code>	The network
<code>data</code>	Data Frame
<code>prior</code>	Joint prior
<code>timetrace</code>	For timing purposes

Details

For each node in the network, all possible parent configurations are created and learned. The result is called a trylist. This is very time-consuming, and a better choice is to maintain a trylist while searching and indeed this is automatically done.

Value

A list with one element per node in the network. Element i is a list of nodes. These nodes are copies of node i in the network, except that all possible parent configurations are present. For each copy, the node is learned so that the parameter posteriors and log-likelihood contribution are calculated. These nodes can then be copied into networks where the node appears.

Note

The full trylist is also output from `networkfamily`.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
Claus Dethlefsen (dethlef@math.auc.dk).

See Also

`reuselearn`, `networkfamily`, `heuristic`

Examples

```
data(rats)
rats.nw <- network(rats)
rats.pr <- jointprior(rats.nw,12)
rats.nw <- learn(rats.nw,rats,rats.pr)$nw
rats.tr <- maketrylist(rats.nw,rats,rats.pr)

rats.hi <- heuristic(rats.nw,rats,rats.pr,trylist=rats.tr)$nw
```

network

Bayesian network data structure

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)
print(x, filename=NA, master=FALSE, condposterior=FALSE,
      condprior=FALSE, ...)
plot(x, scale=10, arrowlength=.25,
     notext=FALSE,
     sscale=.7*scale, showban=TRUE, yr=c(0, 350), xr=yr,
     unitscale=20, cexscale=8, ...)

prob.network(x, df)
```

Arguments

<code>df</code>	A data frame, where the columns define the variables. A continuous variable should have type <code>numeric</code> and discrete variables should have type <code>factor</code> .
<code>specifygraph</code>	Interactively specify a directed acyclic graph and possibly a ban list (see below).
<code>inspectprob</code>	Inspect the calculated probability distribution using a graphical interface.
<code>doprob</code>	If <code>TRUE</code> , do not calculate a probability distribution. Used for example in <code>simulation</code> .
<code>x</code>	The network.
<code>filename</code>	If a string argument is given, output is printed to a file.
<code>condprior</code>	Print <code>condprior</code> ? <code>TRUE/FALSE</code>
<code>condposterior</code>	Print <code>condposterior</code> ? <code>TRUE/FALSE</code>
<code>master</code>	Print <code>master</code> ? <code>TRUE/FALSE</code> .
<code>scale</code>	The axes in network plots are from <code>-scale</code> to <code>scale</code> .
<code>sscale</code>	The nodes are placed on a circle with radius <code>sscale</code> .
<code>unitscale</code>	Scale parameter for chopping off arrow heads.
<code>cexscale</code>	The size of the nodes.
<code>arrowlength</code>	The length of the arrow heads.
<code>xr</code>	Range of x-axis.
<code>yr</code>	Range of y-axis.
<code>notext</code>	No text is displayed in the nodes on the plot if <code>notext</code> is <code>TRUE</code>
<code>rotate</code>	The nodes are placed on a circle and they may be rotated by this number (in radians).
<code>showban</code>	If banned arrows should be shown (they are shown in red).
<code>...</code>	Passed to <code>plot.node</code> .

Value

A network object with attributes

<code>nodes</code>	A list of nodes. If <code>doprob==TRUE</code> , the nodes are given the attribute <code>prob</code> which is the initial probability distribution used by <code>jointprior</code> .
<code>n</code>	The number of nodes.
<code>discrete</code>	A vector of indices of discrete nodes.
<code>continuous</code>	A vector of indices of continuous nodes.
<code>banlist</code>	A matrix with two columns. Each row contains the indices <code>i -> j</code> of arrows that may not be allowed in the directed acyclic graph.
<code>score</code>	Added by <code>learn</code> and is the log network score.
<code>relscore</code>	Added by <code>nwfsort</code> and is the relative log network score – compared with the best network in a network family.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
Claus Dethlefsen (dethlef@math.auc.dk).

See Also

`networkfamily`, `node`, `simulation`, `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) )
```

`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 `nwfunique`) are available.

Usage

```
networkfamily(data,nw=network(data), prior=jointprior(nw),
              trylist=rep(list(NULL),nw$n), timetrace=TRUE)

print(x,...)
plot(x,layout=rep(min(1+floor(sqrt(length(nwf))),5),2),
      cexscale=5,arrowlength=0.1,scale=10,sscale=.7*scale,...)
```

Arguments

<code>nw</code>	The initial network from which the full network family is generated.
<code>data</code>	Data frame used for learning the parameters in each network.
<code>prior</code>	A joint prior. See <code>jointprior</code> .
<code>trylist</code>	Used internally for reusing learning of nodes. See <code>maketrylist</code> .
<code>timetrace</code>	For timing purposes.
<code>x</code>	A network family.
<code>layout</code>	2-dimensional vector of the number of plots in the rows and columns of each plotting page.
<code>cexscale</code>	See <code>plot.network</code>
<code>arrowlength</code>	See <code>plot.network</code>
<code>scale</code>	See <code>plot.network</code>
<code>sscale</code>	See <code>plot.network</code>
<code>...</code>	Other plot parameters passed to <code>plot.network</code> .

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.

After generation of all possible networks, a test for cycles (see `cycletest`) is performed and only networks with directed acyclic graphs are returned and learned.

Value

The function `networkfamily` returns list with two components

<code>nw</code>	The network family.
<code>trylist</code>	An updated trylist (see <code>maketrylist</code>).

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, nwfunique, elementin, addarrows, cycletest

Examples

```
data(rats)
allrats <- networkfamily(rats)$nw
plot(allrats)
print(allrats)
```

node	<i>Nodes</i>
------	--------------

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, name=paste(idx),
      levels=2, levelnames=paste(1:levels), position=c(0,0))
print (x, filename=NA, master=FALSE, condposterior=TRUE, condprior=TRUE, ...)
plot (x, cexscale=10, notext=FALSE, scale=10, ...)
prob.node (x, nw, df)
```

Arguments

x	Node.
parents	A vector with indices of the parents of the node.
idx	The index of the node (the column number of the data frame).
type	"discrete" or "continuous".
name	The name used when plotting and printing. Defaults to the column name in the data frame.
levels	If type="discrete", this is the number of levels for the discrete variable.
levelnames	If type="discrete" this is a vector of strings (same size as levels) with the names of the levels.

<code>position</code>	The coordinates where the node should appear in the plot. Usually set by <code>network</code> and <code>drawnetwork</code> .
<code>df</code>	A data frame, where the columns define the variables. A continuous variable should have type "numeric" and discrete variables should have type "factor".
<code>nw</code>	The network.
<code>filename</code>	If a string argument is given, output is printed to a file.
<code>condprior</code>	Print condprior? TRUE/FALSE
<code>condposterior</code>	Print condposterior? TRUE/FALSE
<code>master</code>	Print master? TRUE/FALSE.
<code>scale</code>	Used to position node.
<code>cexscale</code>	The size of the nodes.
<code>notext</code>	If FALSE, do not display text on plots.
<code>...</code>	Passed to plot.

Value

A node object may have the following attributes

<code>idx</code>	A unique index for this node. It MUST correspond to the column index of the variable in the data frame.
<code>name</code>	The printed name of the node.
<code>type</code>	Either <code>continuous</code> or <code>discrete</code> .
<code>levels</code>	If the node is of type <code>discrete</code> , this integer is the number of levels of the node.
<code>levelnames</code>	If the node is of type <code>discrete</code> , this vector of strings is the names of the levels. This must have length equal to <code>levels</code> .
<code>parents</code>	A vector of indices of the parents to this node. It is best to manage this vector using the <code>insert</code> function.
<code>prob</code>	The initial probability distribution, set by <code>network</code> .
<code>condprior</code>	The parameter priors deduced from <code>jointprior</code> using the master prior procedure (see <code>localmaster</code>).
<code>condposterior</code>	The parameter posteriors obtained from <code>learnnode</code> .
<code>loglik</code>	The log likelihood contribution for this node, calculated in <code>learnnode</code> .
<code>simprob</code>	Added by <code>makesimprob</code> and used by <code>simulation</code> .

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
 Claus Dethlefsen (dethlef@math.auc.dk).

Examples

```
A <- factor(rep(c("A1", "A2"), 50))
B <- factor(rep(rep(c("B1", "B2"), 25), 2))
thisnet <- network( data.frame(A,B) )
```

numbermixed	<i>The number of possible networks</i>
-------------	--

Description

Calculates the number of different DAG's for a set of discrete and continuous nodes.

Usage

```
numbermixed(nd, nc)
```

Arguments

nd	The number of discrete nodes.
nc	The number of continuous nodes.

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 $nd+nc>15$, the procedure is quite slow.

Value

The number of DAG's with the given node configuration.

Author(s)

Susanne Gammelgaard Bøttcher <alma@math.auc.dk>,
Claus Dethlefsen <dethlef@math.auc.dk>.

References

Bøttcher, S.G. (2002). Learning Bayesian Networks with Mixed Variables.
Robinson, R.W. (1977). Counting unlabeled acyclic digraphs, Lecture Notes in Mathematics, 622: Combinatorial Mathematics.

Examples

```
numbermixed(2,2)
numbermixed(5,10)
```

nwfsort

Sorts a list of networks

Description

According to the `score` attribute 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` attribute.

Usage

```
nwfsort(nwf)
```

Arguments

`nwf` Network family.

Author(s)

Susanne Gammelgaard Bøttcher <alma@math.auc.dk>,
Claus Dethlefsen <dethlef@math.auc.dk>.

nwfunique

Makes a network family unique.

Description

Makes a network family unique, i.e. removes networks that are equal or equivalent to networks already in the network family.

Usage

```
nwfunique(nwf, equi=FALSE, timetrace=FALSE, epsilon=1e-12)
```

Arguments

nwf	Network family.
equi	If TRUE, also equivalent networks are thrown out (i.e. if their score is within epsilon from another network).
timetrace	For timing purposes.
epsilon	Closeness of equivalent networks.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
Claus Dethlefsen (dethlef@math.auc.dk).

perturb	<i>Perturbs a network</i>
---------	---------------------------

Description

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

Usage

```
perturb(nw, data, prior, degree=nw$n, trylist=rep(list(NULL), nw$n),  
        nocalc=FALSE, timetrace=TRUE)
```

Arguments

nw	The network from which arrows are added/removed/turned.
data	Data frame used for learning the network.
prior	A joint prior. See jointprior().
degree	The number of attempts to randomly insert/remove/turn an arrow.
trylist	Used internally for reusing learning of nodes.
nocalc	If TRUE, the network score is not updated
timetrace	Timing information available.

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 possibilities, the unchanged network is returned.

<code>phi</code>	The prior parameter ϕ . An integer.
<code>y</code>	Data vector of observations of the current node.
<code>z</code>	Matrix with a column of ones and columns with the observations of the continuous parents.
<code>timetrace</code>	For timing purposes.

Details

These functions are called by the learning routines (`learn`, `learnnode`).

`post0`: posterior for continuous node with 0 parents as batch learning.

`postc0c`: as `post0`, but using sequential learning in C.

`postc`: posterior for continuous node with continuous parents. Sequential learning.

`post`: as `postc`, but as batch learning.

`postM`: as `post`, but using the Matrix library.

`postcc`: as `postc`, but using C.

Value

Updated parameters and log network score contribution of current node.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
Claus Dethlefsen (dethlef@math.auc.dk).

See Also

`learnnode`

Examples

```
data(rats)
fit      <- network(rats, specifygraph=FALSE, inspect=FALSE)
fit.prior <- jointprior(fit, 12)
W1       <- fit$nodes$W1
W1       <- cond.node(W1, fit, fit.prior)
W1.post  <- postc0c(W1$condprior[[1]]$mu,
                    W1$condprior[[1]]$tau,
                    W1$condprior[[1]]$rho,
                    W1$condprior[[1]]$phi,
                    rats[, W1$idx])
```

savenet

Saves Bayesian network as .net file

Description

Exports a network object in the standard format .net (used by eg. Hugin).

Usage

```
savenet(nw, filename = "default.net")
```

Arguments

nw	Network object, see <code>network</code> .
filename	Filename (should end with .net).

Details

The `prob` attribute for each node in the network object is exported along with the network structure defined by the parents of each node.

Value

A file is created.

Author(s)

Susanne Gammelgaard Bøttcher (alma@math.auc.dk),
Claus Dethlefsen (dethlef@math.auc.dk).

See Also

`network`

Examples

```
data(rats)
nw <- network(rats) # the empty network
savenet(nw)
```

readnet

Reads .net file

Description

Reads a Bayesian network specification in the .net language (see e.g. Hugin).

Usage

```
readnet(fn)
```

Arguments

fn File name to be read.

Details

Reads only the structure of a network, i.e. the DAG.

Value

A network with the nodes specified as in the .net file. The network has not been learned and does not contain 'prob' attributes.

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, savenet

Examples

```
data(rats)
nw <- network(rats)
savenet(nw)
nw2 <- readnet("default.net")
```

Description

Given a network with nodes having the 'simprob'-attribute, simulates a data set.

Usage

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

Arguments

nw	A network where each node have the attribute <code>simprob</code> (see <code>makesimprob</code>).
n	The number of cases to simulate.
file	If a file name is given, 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 normal distribution is used, according to the `simprob` attribute 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, specifygraph=FALSE, doprob=FALSE) # doprob must be False
nw <- makesimprob(nw) # create simprob attributes

set.seed(944)
sim <- simulation(nw, n=100) # create simulated data frame
```