

Collective Mining of Bayesian Networks from Distributed Heterogeneous Data

R. Chen (rchen@eecs.wsu.edu) and K. Sivakumar (siva@wsu.edu)

School of EECS, Washington State University, Pullman, WA 99164-2752

H. Kargupta (hillol@cs.umbc.edu)

Department of CSEE, University of Maryland Baltimore County, Baltimore, MD 21250

Abstract. We present a collective approach to learning a Bayesian network from distributed heterogeneous data. We first learn a local Bayesian network at each site using the local data. Then each site identifies the observations that are most likely to be evidence of coupling between local and non-local variables and transmits a subset of these observations to a central site. Another Bayesian network is learnt at the central site using the data transmitted from the local site. The local and central Bayesian networks are combined to obtain a collective Bayesian network, that models the entire data. Experimental results and theoretical justification that demonstrate the feasibility of our approach are presented.

Keywords: Bayesian Network, Collective Data Mining, Distributed Data Mining, Heterogeneous Data

1. Introduction

Raw data is useful only when it is transformed into knowledge or useful information. This involves data analysis and transformation to extract interesting patterns and correlations among the problem variables. In practical applications, such transformations require efficient data access, analysis, and presentation of the outcome in a timely manner. The advent of large distributed environments in both scientific and commercial domains (e.g. the Internet and corporate intranets) introduces a new dimension to this process — a large number of distributed sources of data that can be used for discovering knowledge. Cost of data communication between the distributed databases is a significant factor



© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

in an increasingly mobile and connected world with a large number of distributed data sources. This cost consists of several components like (a) Limited network bandwidth, (b) data security, and (c) existing organizational structure of the applications environment. The field of Distributed Knowledge Discovery and Data Mining (DDM) studies algorithms, systems, and human-computer interaction issues for knowledge discovery applications in distributed environments for minimizing this cost.

In this paper, we consider a Bayesian network (BN) model to represent uncertain knowledge. Specifically, we address the problem of learning a BN from heterogeneous distributed data. It uses a collective data mining (CDM) approach introduced earlier by Kargupta et. al. [30, 32, 33, 35]. Section 2 provides some background and reviews existing literature in this area. Section 3 presents the collective Bayesian learning technique. Experimental results for two datasets — one simulated and one real world — are presented in Section 4. We would like to mention that these experiments are intended mainly to serve as a proof-of-concept. More extensive results for real data mining applications along the web mining line would be published later. Finally, Section 5 provides some concluding remarks and directions for future work.

2. Background, Motivation, and Related Work

In this section, we provide background and motivation to the problem by means of an example. We then review the existing literature in this area.

Distributed data mining (DDM) must deal with different possibilities of data distribution. Different sites may contain data for a common set of features of the problem domain. In case of relational data this would mean a consistent database schema across all the sites. This is the homogeneous case. Tables I and II illustrate this case using an example from a hypothetical credit card transaction domain.¹ There are

¹ Please note that the credit card domain may not always have consistent schema. The domain is used just for illustration.

Table I. Homogeneous case: Site A with a table for credit card transaction records.

Account Number	Amount	Location	Previous record	Unusual transaction
11992346	-42.84	Seattle	Poor	Yes
12993339	2613.33	Seattle	Good	No
45633341	432.42	Portland	Okay	No
55564999	128.32	Spokane	Okay	Yes

Table II. Homogeneous case: Site B with a table for credit card transaction records.

Account Number	Amount	Location	Previous record	Unusual transaction
87992364	446.32	Berkeley	Good	No
67845921	978.24	Orinda	Good	Yes
85621341	719.42	Walnut	Okay	No
95345998	-256.40	Francisco	Bad	Yes

two data sites A and B, connected by a network. The DDM-objective in such a domain may be to find patterns of fraudulent transactions. Note that both the tables have the same schema. The underlying distribution of the data may or may not be identical across different data sites.

In the general case the data sites may be *heterogeneous*. In other words, sites may contain tables with different schemata. Different features are observed at different sites. Let us illustrate this case with relational data. Table III shows two data-tables at site X. The upper table contains weather-related data and the lower one contains demographic data. Table IV shows the content of site Y, which contains holiday toy sales data. The objective of the DDM process may be detecting relations between the toy sales, the demographic and weather

Table III. Heterogeneous case: Site X with two tables, one for weather and the other for demography.

City	Temp.	Humidity	Wind Chill
Boise	20	24%	10
Spokane	32	48%	12
Seattle	63	88%	4
Portland	51	86%	4
Vancouver	47	52%	6

City	State	Size	Average earning	Proportion of small businesses
Boise	ID	Small	Low	0.041
Spokane	WA	Medium	Medium	0.022
Seattle	WA	Large	High	0.014
Portland	OR	Large	High	0.017
Vancouver	BC	Medium	Medium	0.031

related features. In the general heterogeneous case the tables may be related through different sets of key indices. For example, Tables III(upper) and (lower) are related through the key feature *City*; on the other hand Table III (lower) and Table IV are related through key feature *State*. We consider the heterogenous data scenario in this paper.

2.1. MOTIVATION

Bayesian networks offer very useful information about the mutual dependencies among the features in the application domain. Such information can be used for gaining better understanding about the dynamics of the process under observation. Financial data analysis,

Table IV. Heterogeneous case: Site Y with one table holiday toy sales.

State	Best Selling Item	Price (\$)	Number Items Sold (In thousands)
WA	Snarc Action Figure	47.99	23
ID	Power Toads	23.50	2
BC	Light Saber	19.99	5
OR	Super Squirter	24.99	142
CA	Super Fun Ball	9.99	24

manufacturing process monitoring, sensor data analysis, web mining are a few examples where mining Bayesian networks has been quite useful. Bayesian techniques will also be useful for mining distributed data. In this section we discuss one such scenario and explain how the proposed collective Bayesian learning algorithm can be useful in practice.

Consider a subscriber of a wireless network. This person travels frequently and uses her palmtop computer and cell phone to do business and personal transactions. Her transactions go through different servers depending upon her location during the transaction. Now let us say her wireless service provider wants to offer more personalized service to her by paying careful attention to her needs and tastes. This may be useful for choosing the instant messages appropriate for her taste and needs. For example, if she is visiting the Baltimore area the company may choose to send her instant messages regarding the area Sushi and Italian restaurants that she usually prefers. Since too many of such instant messages are likely to be considered a nuisance, accurate personalization is very important. This is indeed quite well appreciated by the business community and use of Bayesian techniques for personalizing web sites has already been reported elsewhere [49, 50, 5, 6].

The scenario described here is however somewhat different from the traditional web personalization applications where web-log data are

centrally located. In the current case the transaction data are logged at different locations since the user is mobile and the wireless service provider is interested in analyzing the overall transactional patterns of the user. As the user moves from one wireless cell to another the servers change and therefore the transactions go through different servers. Today the major wireless service providers support millions of customers. Centralizing this inherently distributed large volume of data may not be scalable. Moreover, these transaction data are heterogeneous. There is no guarantee that the user will perform only a certain type of transactions at every location. The user may choose to perform a wide variety of transactions (e.g. purchasing gifts, money transaction, monitoring the stock portfolio, reading news, and ordering pizza) at different sites. Therefore the features defining the transactions observed at different sites are likely to be different in general although we may have some overlapping characteristic features (e.g. monitoring the stock portfolio everyday!).

As pointed out elsewhere [36] there are many similar situations where data are distributed among a large number of sites and centralized data mining is not scalable. The technique proposed here will be applicable to many such domains.

2.2. RELATED WORK

The volume of DDM literature is growing fast. There exist a reasonably large body of work on DDM architectures and data mining techniques for the homogeneous and heterogeneous cases. In the following, we review only the existing literature for heterogeneous DDM.

Mining from heterogeneous data constitutes an important class of DDM problems. This issue is discussed in [54] from the perspective of inductive bias. The WoRLD system [2] addressed the problem of concept learning from heterogeneous sites by developing an “activation spreading” approach that is based on first order statistical estimation of the underlying distribution. A novel approach to learn association rules from heterogeneous tables is proposed in [17]. This approach exploits the foreign key relationships for the case of a star schema to develop de-

centralized algorithms that execute concurrently on the separate tables, and subsequently merge the results. An order statistics-based technique for combining high-variance models generated from heterogeneous sites is proposed in [64].

Kargupta and his colleagues [35] also considered the heterogeneous case and proposed the *Collective* framework to address data analysis for heterogeneous environments. They proposed the *Collective Data Mining* (CDM) framework for predictive data modeling that makes use of orthonormal basis functions for correct local analysis. They proposed a technique for distributed decision tree construction [35] and wavelet-based multi-variate regression [30]. Several distributed clustering techniques based on the Collective framework are proposed elsewhere [32, 34]. They also proposed the collective PCA technique [34, 33] and its extension to a distributed clustering application. Additional work on distributed decision tree learning [4], clustering [45, 48, 55], genetic learning [47], DDM design optimization [65], classifier pruning [53], DDM architecture [38], and problem decomposition and local model selection in DDM [43], are also reported.

We now review important literature on learning using Bayesian networks (BN). A BN is a probabilistic graphical model that represents uncertain knowledge [51, 31, 11]. Learning parameters of a Bayesian network from complete data is discussed in [58, 10]. Learning parameters from incomplete data using gradient methods is discussed in [7, 61]. Lauritzen [41] has proposed an EM algorithm to learn Bayesian network parameters, whereas Bauer et. al. [3] describe methods for accelerating convergence of the EM algorithm. Learning using Gibbs sampling is proposed in [63, 25]. The Bayesian score to learn the structure of a Bayesian network is discussed in [16, 10, 27]. Learning the structure of a Bayesian network based on the Minimal Description Length (MDL) principle is presented in [8, 39, 60]. Learning BN structure using greedy hill-climbing and other variants is introduced in [28], whereas Chickering [14] introduced a method based on search over equivalence network classes. Methods for approximating full Bayesian model averaging are presented in [10, 28, 44].

Learning the structure of Bayesian network from incomplete data, is considered in [15, 12, 20, 21, 46, 56, 62]. The relationship between causality and Bayesian networks is discussed in [28, 52, 59, 29]. See [10, 23, 39] for discussion on how to sequentially update the structure of a network as more data is available. Applications of Bayesian network to clustering (AutoClass) and classification is discussed in [12, 19, 22, 57]. Zweig and Russel [66] use Bayesian networks for speech recognition, whereas Breese et. al. [9] discuss collaborative filtering methods that use Bayesian network learning algorithms. Applications to causal learning in social sciences is discussed in [59]. In [40] the authors report a technique to automatically produce a Bayesian belief network from discovered knowledge using a distributed approach.

An important problem is how to learn the Bayesian network from data in distributed sites. The centralized solution to this problem is to download all datasets from distributed sites. Kenji [37] worked on the homogeneous distributed learning scenario. In this case, every distributed site has the same feature but different observations. In this paper, we address the heterogenous case, where each site has data about only a subset of the features. To our knowledge, there is no significant work that addresses the heterogenous case.

3. Collective Bayesian Learning

In the following, we briefly review Bayesian networks and then discuss our collective approach to learning a Bayesian network that is specifically designed for a distributed data scenario.

3.1. BAYESIAN NETWORKS: A REVIEW

A Bayesian network (BN) is a probabilistic graph model. It can be defined as a pair (\mathcal{G}, p) , where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed acyclic graph (DAG) [31, 26]. Here, \mathcal{V} is the vertex set which represents variables in the problem and \mathcal{E} is the edge set which denotes probabilistic relationships among the variables. For a variable $X \in \mathcal{V}$, a parent of X is a

node from which there is a directed link to X . Let $pa(X)$ denote the set of parents of X , then the conditional independence property can be represented as follows:

$$P(X \mid \mathcal{V} \setminus X) = P(X \mid pa(X)). \quad (1)$$

This property can simplify the computations in a Bayesian network model. For example, the joint distribution of the set of all variables in \mathcal{V} can be written as a product of conditional probabilities as follows:

$$P(\mathcal{V}) = \prod_{X \in \mathcal{V}} P(X \mid pa(X)). \quad (2)$$

The conditional independence between variables is either obtained from a priori expert knowledge or discerned from data, or a combination of both [31]. The set of conditional distributions $\{P(X \mid pa(X)), X \in \mathcal{V}\}$ are called the parameters of a Bayesian network. Note that if variable X has no parents, then $P(X \mid pa(X)) = P(X)$ is the marginal distribution of X .

Figure 1 is a Bayesian network called the ASIA model (adapted from [42]). The variables are Dyspnoea, Tuberculosis, Lung cancer, Bronchitis, Asia, X-ray, Either, and Smoking. They are all binary variables. The joint distribution of all variables is

$$\begin{aligned} P(A, S, T, L, B, E, X, D) = & P(A)P(S)P(T \mid A)P(L \mid S) \\ & P(B \mid S)P(E \mid T, L)P(X \mid E)P(D \mid B, E). \end{aligned} \quad (3)$$

The ordering of variables constitutes a constraint on the structure of a Bayesian network. If variable X appears before variable Y , then Y can not be a parent of X . We use the ordering (A, S, T, L, B, E, X, D) as prior knowledge in our example.

Two important issues in using a Bayesian network are : (a) learning a Bayesian network and (b) probabilistic inference. Learning a BN involves learning the structure of the network (the directed graph), and obtaining the conditional probabilities (parameters) associated with the network. Once a Bayesian network is constructed, we usually need to determine various probabilities of interest from the model. This process is referred to as probabilistic inference. For example, in the ASIA

model, a diagnosis application would require finding the probability $P(B \mid D)$ of Bronchitis, given the (observed) symptom Dyspnoea. This probability (usually called posterior probability) can be computed using the Bayes rule.




Figure 1. ASIA Model

3.2. COLLECTIVE BAYESIAN NETWORK LEARNING STRATEGY

We now present a collective strategy to learn a Bayesian network (both structure and parameters) when data is distributed among different sites. The centralized solution to this problem is to download all datasets from distributed sites to a central site. In many applications, this would not be feasible because of the size of the data, available communication bandwidth, or due to security considerations. Learning a BN for the homogeneous case was studied by Kenji [37]. In this case, every distributed site has the same set of features but has different set of observations. We address here the heterogeneous case, where each distributed site has all the observations for only a subset of the features.

The primary steps in our approach are:

- Compute local BNs (local model) involving the variables observed at each site (local variables) based on local data.
- At each site, based on the local BN, identify the observations that are most likely to be evidence of coupling between local and non-local variables. Transmit a subset of these observations to a central site.

- At the central site, a limited number of observations of all the variables are now available. Using this, compute a non-local BN consisting of links between variables across two or more sites.
- Combine the local models with the links discovered at the central site to obtain a collective BN.

The non-local BN thus constructed would be effective in identifying associations between variables across sites, whereas the local BNs would detect associations among local variables at each site. The conditional probabilities can also be estimated in a similar manner. Those probabilities that involve only variables from a single site can be estimated locally, whereas the ones that involve variables from different sites can be estimated at the central site. Same methodology could be used to update the network based on new data. First, the new data is tested for how well it fits with the local model. If there is an acceptable statistical fit, the observation is used to update the local conditional probability estimates. Otherwise, it is also transmitted to the central site to update the appropriate conditional probabilities (of cross terms). Finally, a collective BN can be obtained by taking the union of nodes and edges of the local BNs and the nonlocal BN and using the conditional probabilities from the appropriate BNs. Probabilistic inference can now be performed based on this collective BN. Note that transmitting the local BNs to the central site would involve a significantly lower communication as compared to transmitting the local data.

It is quite evident that learning probabilistic relationships between variables that belong to a single local site is straightforward and does not pose any additional difficulty as compared to a centralized approach. The important objective is to correctly identify the coupling between variables that belong to two (or more) sites. These correspond to the edges in the graph that connect variables between two sites and the conditional probability(ies) at the associated node(s). In the following, we describe our approach to selecting observations at the local sites that are most likely to be evidence of strong coupling between variables at two different sites. The key idea of our approach is that the samples that do not fit well with the local models are likely to be evidence

of coupling between local and non-local variables. We transmit these samples to a central site and use them to learn a collective Bayesian network.

3.3. SELECTION OF SAMPLES FOR TRANSMISSION TO GLOBAL SITE

For simplicity, we will assume that the data is distributed between two sites and will illustrate the approach using the BN in Figure 1. The extension of this approach to more than two sites is straightforward. Let us denote by \mathcal{A} and \mathcal{B} , the variables in the left and right groups, respectively, in Figure 1. We assume that the observations for \mathcal{A} are available at site A, whereas the observations for \mathcal{B} are available at a different site B. Furthermore, we assume that there is a common feature (“key” or index) that can be used to associate a given observation in site A to a corresponding observation in site B. Naturally, $\mathcal{V} = \mathcal{A} \cup \mathcal{B}$.

At each local site, a local Bayesian network can be learned using only samples in this site. This would give a BN structure involving only the local variables at each site and the associated conditional probabilities. Let $p_A(\cdot)$ and $p_B(\cdot)$ denote the estimated probability function involving the local variables. This is the product of the conditional probabilities as indicated by (2). Since $p_A(x)$, $p_B(x)$ denote the probability or likelihood of obtaining observation x at sites A and B, we would call these probability functions the likelihood functions $l_A(\cdot)$ and $l_B(\cdot)$, for the local model obtained at sites A and B, respectively. The observations at each site are ranked based on how well it fits the local model, using the local likelihood functions. The observations at site A with large likelihood under $l_A(\cdot)$ are evidence of “local relationships” between site A variables, whereas those with low likelihoods under $l_A(\cdot)$ are possible evidence of “cross relationships” between variables across sites. Let $S(A)$ denote the set of keys associated with the latter observations (those with low likelihood under $l_A(\cdot)$). In practice, this step can be implemented in different ways. For example, we can set a threshold ρ_A and if $l_A(x) \leq \rho_A$, then $x \in S_A$. The sites A and B transmit the set of keys S_A , S_B , respectively, to a central site, where the intersection $S = S_A \cap S_B$ is computed. The observations corresponding to the set of

keys in S are then obtained from each of the local sites by the central site.

The following argument justifies our selection strategy. Using the rules of probability, and the assumed conditional independence in the BN of Figure 1, it is easy to show that:

$$P(\mathcal{V}) = P(\mathcal{A}, \mathcal{B}) = P(\mathcal{A} | \mathcal{B})P(\mathcal{B}) = P(\mathcal{A} | nb(\mathcal{A}))P(\mathcal{B}), \quad (4)$$

where $nb(\mathcal{A}) = \{B, L\}$ is the set of variables in \mathcal{B} , which have a link connecting it to a variable in \mathcal{A} . In particular,

$$P(\mathcal{A} | nb(\mathcal{A})) = P(A)P(T | A)P(X | E)P(E | T, L)P(D | E, B). \quad (5)$$

Note that, the first three terms in the right-hand side of (5) involve variables local to site A, whereas the last two terms are the so-called *cross terms*, involving variables from sites A and B. Similarly, it can be shown that

$$P(\mathcal{V}) = P(\mathcal{A}, \mathcal{B}) = P(\mathcal{B} | \mathcal{A})P(\mathcal{A}) = P(\mathcal{B} | nb(\mathcal{B}))P(\mathcal{A}), \quad (6)$$

where $nb(\mathcal{B}) = \{E, D\}$ and

$$P(\mathcal{B} | nb(\mathcal{B})) = P(S)P(B | S)P(L | S)P(E | T, L)P(D | E, B). \quad (7)$$

Therefore, an observation $\{A = a, T = t, E = e, X = x, D = d, S = s, L = l, B = b\}$ with low likelihood at both sites A and B; i.e. for which both $P(\mathcal{A}) = P(A = a, T = t, E = e, X = x, D = d)$ and $P(\mathcal{B}) = P(S = s, L = l, B = b)$ are small, is an indication that both $P(\mathcal{A} | nb(\mathcal{A}))$ and $P(\mathcal{B} | nb(\mathcal{B}))$ are large for that observation (since observations with small $P(\mathcal{V})$ are less likely to occur). Notice from (5) and (7) that the terms common to both $P(\mathcal{A} | nb(\mathcal{A}))$ and $P(\mathcal{B} | nb(\mathcal{B}))$ are precisely the conditional probabilities that involve variables from both sites A and B. In other words, this is an observation that indicates a coupling of variables between sites A and B and should hence be transmitted to a central site to identify the specific coupling links and the associated conditional probabilities.

In a sense, our approach to learning the cross terms in the BN involves a selective sampling of the given dataset that is most relevant

to the identification of coupling between the sites. This is a type of *importance sampling*, where we select the observations that have high conditional probabilities corresponding to the terms involving variables from both sites. Naturally, when the values of the different variables (features) from the different sites, corresponding to these selected observations are pooled together at the central site, we can learn the coupling links as well as estimate the associated conditional distributions. These selected observations will, by design, not be useful to identify the links in the BN that are local to the individual sites. This has been verified in our experiments (see Section 4).

3.4. PERFORMANCE ANALYSIS

In the following, we present a brief theoretical analysis of the performance of the proposed collective learning method. We compare the performance of our collective BN with that of a Bayesian network learned using a centralized approach (referred to as centralized BN in the sequel).

There are two types of errors involved in learning a BN: (a) Error in BN structure and (b) Error in parameters (probabilities) of the BN. The structure error is defined as the sum of the number of correct edges missed and the number of incorrect edges detected. For parameter error, we need to quantify the “distance” between two probability distributions. We only consider learning error in the parameters, assuming that the structure of the BN has been correctly determined (or is given). A widely used metric is the Kullback-Leibler (KL) distance (cross-entropy measure) $d_{KL}(p, q)$ between two discrete probabilities, $\{p_i\}$, $\{q_i\}$, $i = 1, 2, \dots, N$

$$d_{KL}(p, q) = \sum_{i=1}^N p_i \ln\left(\frac{p_i}{q_i}\right) \quad (8)$$

where N is the number of possible outcomes.

Indeed, if p^* is the empirically observed distribution for data samples $\{s_i, 1 \leq i \leq M\}$ and h is a hypothesis (candidate probability

distribution for the underlying true distribution), then [1]

$$\begin{aligned} d_{KL}(p^*, h) &= \sum_{i=1}^M p^*(s_i) \ln\left(\frac{p^*(s_i)}{h(s_i)}\right) = \sum_{i=1}^M \frac{1}{M} \ln \frac{1}{M} - \sum_{i=1}^M \frac{1}{M} \ln(h(s_i)) \\ &= \ln \frac{1}{M} - \frac{1}{M} \sum_{i=1}^M \ln(h(s_i)). \end{aligned} \quad (9)$$

Therefore, minimizing the KL distance with respect to the empirically observed distribution is equivalent to finding the maximum likelihood solution h^* of $\sum_{i=1}^M \ln(h(s_i))$.

Since the BN provides a natural factorization of the joint probability in terms of the conditional probabilities at each node (see (2)), it is convenient to express the KL distance between two joint distributions in terms of the corresponding conditional distributions. Let h and c be two possible (joint) distributions of the variables in a BN. For $i = 1, 2, \dots, n$, let $h_i(x_i \mid \pi_i)$, $c_i(x_i \mid \pi_i)$ be the corresponding conditional distribution at node i , where x_i is the variable at node i and π_i is the set of parents of node i . Following [18], define a distance $d_{CP}(P, c_i, h_i)$ between h_i and c_i with respect to the true distribution P :

$$d_{CP}(P, c_i, h_i) = \sum_{\pi_i} P(\pi_i) \sum_{x_i} P(x_i \mid \pi_i) \ln\left(\frac{c_i(x_i \mid \pi_i)}{h_i(x_i \mid \pi_i)}\right). \quad (10)$$

It is then easy to show that

$$d_{KL}(P, h) - d_{KL}(P, c) = \sum_{i=1}^n d_{CP}(P, c_i, h_i). \quad (11)$$

Equations (10) and (11) provide a useful decomposition of the KL distance between the true distribution P and two different hypotheses c , h . This will be useful in our analysis of sample complexity in the following sub-section.

3.5. SAMPLE COMPLEXITY

We now derive a relationship between the accuracy of collective BN and the number of samples transmitted to the central site. We consider

the unrestricted multinomial class BN, where all the node variables are Boolean. The hypothesis class H is determined by the set of possible conditional distributions for the different nodes. Given a BN of n variables and a hypothesis class H , we need to choose a hypothesis $h \in H$ which is close to a unknown distribution P . Given an error threshold ϵ and a confidence threshold δ , we are interested in constructing a function $N(\epsilon, \delta)$, such that if the number of samples M is larger than $N(\epsilon, \delta)$

$$\text{Prob}(d_{KL}(P, h) < d_{KL}(P, h_{opt}) + \epsilon) > 1 - \delta, \quad (12)$$

where $h_{opt} \in H$ is the hypothesis that minimizes $d_{KL}(P, h)$. If smallest value of $N(\epsilon, \delta)$ that satisfies this requirement is called the sample complexity. This is usually referred to as the probably approximately correct (PAC) framework. Friedman and Yakhini [24] have examined the sample complexity of the maximum description length principle (MDL) based learning procedure for BNs.

Dasgupta [18] gave a thorough analysis for the multinomial model with Boolean variables. Suppose the BN has n nodes and each node has at most k parents. Given ϵ and δ , an upper bound of sample complexity is

$$N(\epsilon, \delta) = \frac{288n^2 2^k}{\epsilon^2} \ln^2\left(1 + \frac{3n}{\epsilon} \ln \frac{18n^2 2^k \ln(1 + 3n/\epsilon)}{\epsilon\delta}\right). \quad (13)$$

Equation (13) gives a relation between the sample size and the (ϵ, δ) bound. For the conditional probability $h_i(x_i | \pi_i) = P(X_i = x_i | \Pi_i = \pi_i)$, we have (see (10))

$$d_{CP}(P, h_{opt}, h) \leq \frac{\epsilon}{n} \quad (14)$$

We now use the above ideas to compare the performance of the collective learning method with the centralized method. We fix the confidence δ and suppose that an ϵ^{cen} can be found for the centralized method, for a given sample size M using (13). Then, following the analysis in [18, Section 5],

$$d_{CP}(P, h_{opt}^{cen}, h^{cen}) \leq \frac{\epsilon^{cen}}{n}, \quad (15)$$

where h_{opt}^{cen} is the optimal hypothesis and h^{cen} is the hypothesis obtained based on a centralized approach. Then from (11)

$$\begin{aligned} d_{KL}(P, h^{cen}) - d_{KL}(P, h_{opt}^{cen}) &= \sum_{i=1}^n d_{CP}(P, h_{i,opt}^{cen}, h_i^{cen}) \\ &\leq \sum_{i=1}^n \frac{\epsilon^{cen}}{n} = \epsilon^{cen}. \end{aligned} \quad (16)$$

For the collective BN learning method, the set of nodes can be split into two parts. Let V_l be the set of nodes, which have all their parent nodes at the same local site, and V_c be the set of nodes, which have at least one parent node belonging to a site different than the node itself. For ASIA model, $V_l = \{A, S, T, L, B, X\}$ and $V_c = \{E, D\}$. We use n_l and n_c to denote the cardinality of the sets V_l and V_c . If a node $x \in V_l$, the collective method can learn the conditional probability $P(x \mid pa(x))$ using all data because this depends only on the local variables. Therefore, for $x \in V_l$,

$$d_{CP}(P, h_{opt}^{col}, h^{col}) \leq \frac{\epsilon_1^{col}}{n} = \frac{\epsilon^{cen}}{n}, \quad (17)$$

where, for the local terms, $\epsilon_1^{col} = \epsilon^{cen}$. For the nodes in V_c , only the data transmitted to the central site can be used to learn its conditional probability. Suppose M_c data samples are transmitted to the central site, and the error threshold ϵ_2^{col} satisfies (13), for the same fixed confidence $1 - \delta$. Therefore, for $x \in V_c$, we have from (14) that $d_{CP}(P, h_{opt}^{col}, h^{col}) \leq \frac{\epsilon_2^{col}}{n}$, where $\epsilon_2^{col} \geq \epsilon^{cen}$, in general, since in the collective learning method, only $M_c \leq M$ samples are available at the central site. Then from (11) and (17)

$$\begin{aligned} d_{KL}(P, h_{opt}^{col}) - d_{KL}(P, h^{col}) &= \sum_{i=1}^n d_{CP}(P, h_{i,opt}^{col}, h_i^{col}) \\ &= \sum_{i \in V_l} d_{CP}(P, h_{i,opt}^{col}, h_i^{col}) + \sum_{i \in V_c} d_{CP}(P, h_{i,opt}^{col}, h_i^{col}) \\ &= \frac{n_l}{n} \epsilon^{cen} + \frac{n_c}{n} \epsilon_2^{col} \end{aligned} \quad (18)$$

Comparing (16) and (18), it is easy to see that the error threshold of the collective method is $\epsilon^{col} = \frac{n_l}{n} \epsilon^{cen} + \frac{n_c}{n} \epsilon_2^{col}$. The difference of the

error threshold between the collective and the centralized method is

$$\epsilon^{col} - \epsilon^{cen} = \frac{n_c}{n}(\epsilon_2^{col} - \epsilon^{cen}) \quad (19)$$

Equation (19) shows two important properties of the collective method. First, the difference in performance is independent of the variables in V_l . This means the performance of the collective method for the parameters of local variables is same as that of the centralized method. Second, the collective method is a tradeoff between accuracy and the communication overhead. The more data we communicate, more closely ϵ_2^{col} will be to ϵ^{cen} . When $M_c = M$, $\epsilon_2^{col} = \epsilon^{cen}$, and $\epsilon^{col} - \epsilon^{cen} = 0$.

4. Experimental Results

We tested our approach on two datasets — ASIA model and Web log data. We present our results for the two cases in the following subsections.

4.1. ASIA MODEL

Our experiments were performed on a dataset that was generated from the BN depicted in Figure 1 (ASIA Model). The conditional probability of a variable is a multidimensional array, where the dimensions are arranged in the same order as ordering of the variables, viz. $\{A, S, T, L, B, E, X, D\}$. Table V (top) depicts the conditional probability of node E . It is laid out such that the first dimension toggles fastest. From Table V, we can write the conditional probability of node E as a single vector as follows: $[0.9, 0.1, 0.1, 0.01, 0.1, 0.9, 0.9, 0.99]$. The conditional probabilities (parameters) of ASIA model are given in Table V (bottom) following this ordering scheme. We generated $n = 6000$ observations from this model, which were split into two sites as illustrated in Figure 1 (site A with variables A, T, E, X, D and site B with variables S, L, B). Note that there are two edges ($L \rightarrow E$ and $B \rightarrow D$) that connect variables from site A to site B, the rest of the six edges being local.

Table V. (Top) The conditional probability of node E and (Bottom) All conditional probabilities for the ASIA model

No.	T	L	E	Probability
1	F	F	F	0.9
2	T	F	F	0.1
3	F	T	F	0.1
4	T	T	F	0.01
5	F	F	T	0.1
6	T	F	T	0.9
7	F	T	T	0.9
8	T	T	T	0.99

A	0.99	0.01						
S	0.5	0.5						
T	0.1	0.9	0.9	0.1				
L	0.3	0.6	0.7	0.4				
B	0.1	0.8	0.9	0.2				
E	0.9	0.1	0.1	0.01	0.1	0.9	0.9	0.99
X	0.2	0.6	0.8	0.4				
D	0.9	0.1	0.1	0.01	0.1	0.9	0.9	0.99

Local Bayesian networks were constructed using a conditional independence test based algorithm [13], for learning the BN structure and a maximum likelihood based method for estimating the conditional probabilities. The local networks were exact as far as the edges involving only the local variables. We then tested the ability of the collective approach to detect the two non-local edges. The estimated parameters of these two local Bayesian network is depicted in Table VI. Clearly, the estimated probabilities at all nodes, except nodes E and D , are close to the true probabilities given in Table V. In other words, the parameters

Table VI. The conditional probabilities
of local site A and local site B

Local A				
A	0.99	0.01		
T	0.10	0.84	0.90	0.16
E	0.50	0.05	0.50	0.95
X	0.20	0.60	0.80	0.40
D	0.55	0.05	0.45	0.95
Local B				
S	0.49	0.51		
L	0.30	0.59	0.70	0.41
B	0.10	0.81	0.90	0.19

that involve only local variables have been successfully learnt at the local sites.

A fraction of the samples, whose likelihood are smaller than a selected threshold T , were identified at each site. In our experiments, we set

$$T_i = \mu_i + \alpha\sigma_i, \quad i \in \{A, B\}, \quad (20)$$

for some constant α , where μ_i is the (empirical) mean of the local likelihood values and σ_i is the (empirical) standard deviation of the local likelihood values. The samples with likelihood less than the threshold (T_A at site A T_B at site B) at both sites were sent to a central site. The central learns a global BN based on these samples. Finally, a collective BN is formed by taking the union of edges detected locally and those detected at the central site. The error in structure learning of the collective Bayesian network is defined as the sum of the number of correct edges missed and the number of incorrect edges detected. This is done for different values of α . Figure 2 (left) depicts this error as a function of the number of samples communicated (which is determined by α). It




Figure 2. Performance of collective BN: (left) structure learning error (right) parameter learning error.

is clear that the exact structure can be obtained by transmitting about 5% of the total samples.

Next we assessed the accuracy of the estimated conditional probabilities. For the collective BN, we used the conditional probabilities from local BN for the local terms and the ones estimated at the global site for the cross terms. This was compared with the performance of a BN learnt using a centralized approach (by aggregating all data at a single site). Figure 2 (right) depicts the KL distance $d(p_{cntr}(\mathcal{V}), p_{coll}(\mathcal{V}))$ between the joint probabilities computed using our collective approach and the one computed using a centralized approach. Clearly, even with a small communication overhead, the estimated conditional probabilities based on our collective approach is quite close that obtained from a centralized approach.

A more important test of our approach is the error in estimating the conditional probabilities at nodes E and D , since these are the cross terms, estimated at the global site, based on a selective transmission of data. The KL distance between the conditional probabilities was computed based on our collective BN and a BN obtained using a centralized approach (by transmitting all data to one site), for the cross terms: $p(E | T, L)$ and $p(D | E, B)$. Given that these are conditional probabilities, we compute the sum over all the possible values of

$\{T, L\}$, of the KL distance between $p_{\text{coll}}(E | T, L)$ and $p_{\text{cntr}}(E | T, L)$, estimated using our collective approach and the centralized approach, respectively. Figure 3 (top left) depicts the KL distance $\sum_{T,L} d(p_{\text{cntr}}(E | T, L), p_{\text{coll}}(E | T, L))$, between the two estimates. Figure 3 (top right) depicts the sum $\sum_{B,E} d(p_{\text{cntr}}(D | B, E), p_{\text{coll}}(D | B, E))$, over all the possible values of $\{E, B\}$, of the KL distance between the two estimates. Clearly, even with a small data communication, the estimates of the conditional probabilities of the cross-terms, based on our collective approach, is quite close to that obtained by the centralized approach. To further verify the validity of our approach, the transmitted data at the central site was used to estimate two local conditional probabilities, $p(X | E)$ and $p(L | S)$. The corresponding KL distances are depicted in the bottom row of Figure 3 (left: node L and right: node: X). It is clear that the estimates of these probabilities is quite poor, unless a substantial fraction of the data is transmitted. This clearly demonstrates that our technique can be used to perform a biased sampling for discovering relationships between variables across sites.

4.2. WEBSERVER LOG DATA

In the second set of experiments, we used data from real world domain — a web server log data. Web server log contains records of user interactions when request for the resources in the servers is received. Web log mining can provide useful information about different user profiles. In our application, the raw web log file was obtained from the web server of the School of EECS at Washington State University — <http://www.eecs.wsu.edu>. There are three steps in our processing. First we preprocess the raw web log file to transform it to a session form which is useful to our application. Each session corresponds to the logs from a single user in a single web session. We consider each session as a data sample. Then we categorize the resource (html, video, audio etc.) requested from the server into eight categories: E-EE Faculty, C-CS Faculty, L-Lab and facilities, T-Contact Information, A-Admission Information, U-Course Information, H-EECS Home, and R-Research. These categories are our features. Each feature value in a session is

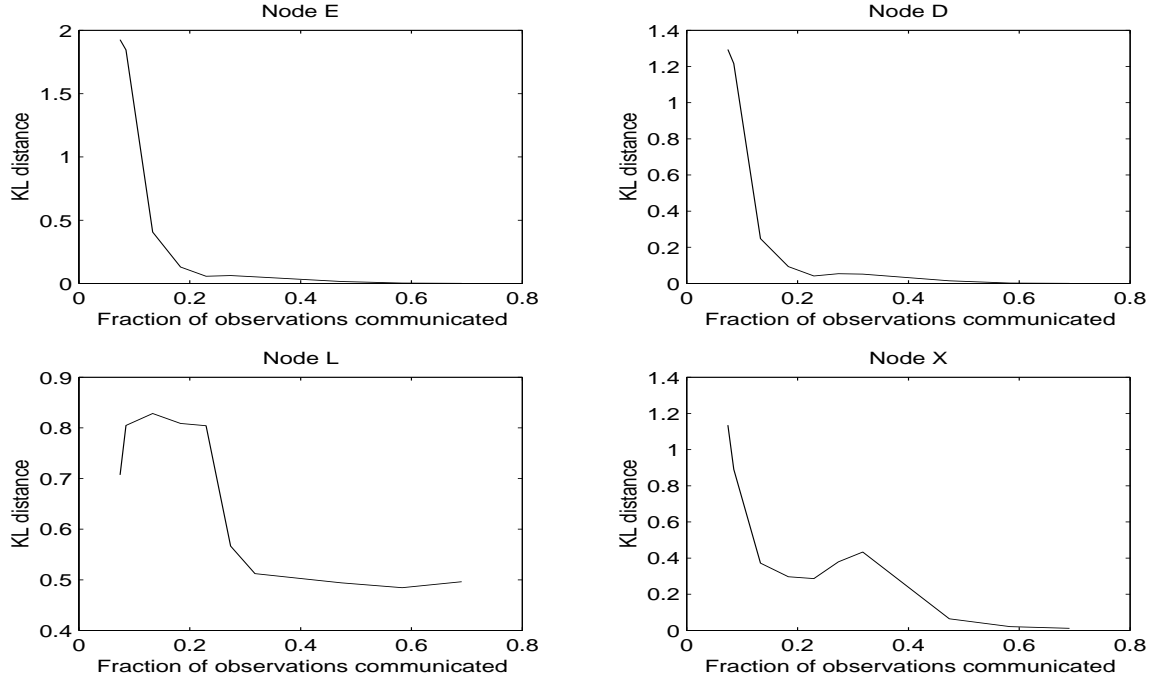


Figure 3. KL distance between conditional probabilities

set to one or zero, depending on whether the user requested resources corresponding to that category. An 8-feature, binary dataset was thus obtained, which was used to learn a BN.

A central BN was first obtained using the whole dataset. Figure 4 depicts the structure of this centralized BN. We then split the features into two sets, corresponding to a scenario where the resources are split into two different web servers. Site A has features E, C, T, and U and site B has features L, A, H, and R. We assumed that the BN structure was known, and estimated the parameters (probability distribution) of the BN using our collective BN learning approach. Figure 5 shows the KL distance between the central BN and the collective BN as a function of the fraction of observations communicated. Clearly the parameters of collective BN is close to that of central BN even with a small fraction of data communication.




Figure 4. KL distance between conditional probabilities

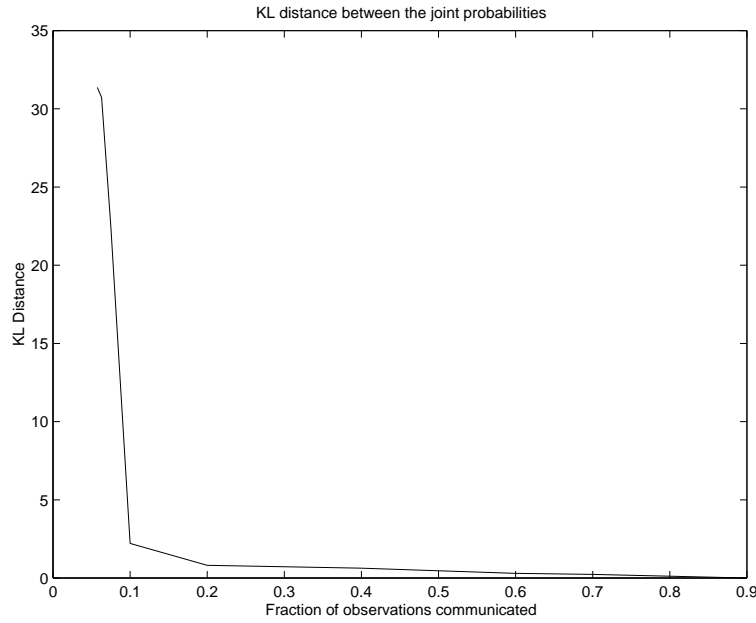


Figure 5. KL distance between conditional probabilities

5. Conclusions

We have presented an approach to learn Bayesian networks from distributed heterogeneous data. This is based on a collective learning strategy, where a local model is obtained at each site and the global associations are determined by a selective transmission of data to a central site. In our experiments, the performance of the collective Bayesian network was quite comparable to that obtained from a centralized approach,

even for a small data communication. To our knowledge, this is the first approach to learning Bayesian networks from distributed heterogeneous data.

A number of theoretical and practical issues remains to be resolved. In particular, we need to obtain tight bounds for the performance of our collective BN as compared to that obtained from a centralized approach, as a function of the data communication involved. Some preliminary analysis of this problem was presented in Section 3. A thorough understanding of this problem would provide performance guarantees for our approach (in terms of (ϵ, δ) error bounds used in Section 3), in terms of the communication overhead. Other important problems include (a) scalability of the approach with respect to number of sites, and number of variables and (b) effect of the choice of existing learning algorithms on performance of the collective Bayesian network. We are actively pursuing these issues and would report this in a future publication.

Acknowledgements

The third author would like to acknowledge support from an NSF Career Award IIS-0093353.

References

1. Abe, N., J. Takeuchi, and M. Warmuth: 1991, 'Polynomial learnability of probabilistic concepts with respect to the Kullback-Leibler divergence'. In: *Proceedings of the 1991 Workshop on Computational Learning Theory*. pp. 277–289.
2. Aronis, J., V. Kulluri, F. Provost, and B. Buchanan: 1997, 'The WoRLD: Knowledge Discovery and Multiple Distributed Databases'. In: *Proceedings of the Florida Artificial Intelligence Research Symposium (FLAIRS-97)*. pp. 11–14. Also available as Technical Report ISL-96-6, Intelligent Systems Laboratory, Department of Computer Science, University of Pittsburgh.

3. Bauer, E., D. Koller, and Y. Singer: 1997, 'Update rules for parameter estimation in Bayesian networks'. In: D. Geiger and P. Shanoy (eds.): *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*. pp. 3–13.
4. Bhatnagar, R. and S. Srinivasan: 1997, 'Pattern Discovery in Distributed Databases'. In: *Proceedings of the AAAI-97 Conference*. Providence, pp. 503–508.
5. Billsus, D. and M. Pazzani: 1996, 'Revising User Profiles: The Search for Interesting Web Sites'. In: *Proceedings of the Third International Workshop on Multistrategy Learning*. AAAI Press.
6. Billsus, D. and M. Pazzani: 1997, 'Learning Probabilistic User Models'. In: *Workshop notes of Machine Learning for User Modeling — Sixth International Conference on User Modeling*. Chia Laguna, Sardinia.
7. Binder, J., D. Koller, S. Russel, and K. Kanazawa: 1997, 'Adaptive probabilistic networks with hidden variables'. *Machine Learning* **29**, 213–244.
8. Bouckaert, R. R.: 1994, 'Properties of Bayesian network learning algorithms'. In: R. L. de Mantaras and D. Poole (eds.): *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*. pp. 102–109.
9. Breese, J. S., D. Heckerman, and C. Kadie: 1998, 'Empirical analysis of predictive algorithms for collaborative filtering'. In: G. F. Cooper and S. Moral (eds.): *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*.
10. Buntine, W.: 1991, 'Theory refinement on Bayesian networks'. In: B. D. D'Ambrosio and P. S. and P. P. Bonissone (eds.): *Proceedings of the Seventh Annual Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, pp. 52–60.
11. Charniak, E.: 1991, 'Bayesian networks without tears'. *AI Magazine* **12**, 50–63.
12. Cheeseman, P. and J. Stutz: 1996, 'Bayesian Classification (AutoClass): Theory and Results'. In: U. Fayyad, G. P. Shapiro, P. Smyth, and R. S. Uthurasamy (eds.): *Advances in Knowledge Discovery and Data Mining*. AAAI Press.
13. Cheng, J., D. A. Bell, and W. Liu: 1997, 'Learning Belief Networks from Data: An Information Theory Based Approach'. In: *Proceedings of the Sixth ACM International Conference on Information and Knowledge Management*.
14. Chickering, D. M.: 1996, 'Learning equivalence classes of Bayesian network structure'. In: E. Horvitz and F. Jensen (eds.): *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*.
15. Chickering, D. M. and D. Heckerman: 1997, 'Efficient approximation for the marginal likelihood of incomplete data given a Bayesian network'. *Machine Learning* **29**, 181–212.
16. Cooper, G. F. and E. Herskovits: 1992, 'A Bayesian method for the induction of probabilistic networks from data'. *Machine Learning* **9**, 309–347.

17. Crestana, V. and N. Soparkar: 1999, 'Mining Decentralized Data Repositories'. Technical Report CSE-TR-385-99, University of Michigan, Ann Arbor, MI.
18. Dasgupta, S.: 1997, 'The sample complexity of learning fixed-structure Bayesian networks'. *Machine Learning* **29**, 165–180.
19. Ezawa, K. J. and S. T.: 1995, 'Fraud/uncollectable debt detection using Bayesian network based learning system: A rare binary outcome with mixed data structures'. In: P. Besnard and S. Hanks (eds.): *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. pp. 157–166.
20. Friedman, N.: 1997, 'Learning Bayesian networks in the presence of missing values and hidden variables'. In: D. Fisher (ed.): *Proceedings of the Fourteenth International Conference on Machine Learning*.
21. Friedman, N.: 1998, 'The Bayesian structural EM algorithm'. In: G. F. Cooper and S. Moral (eds.): *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*.
22. Friedman, N., D. Geiger, and M. Goldszmidt: 1997, 'Bayesian network classifiers'. *Machine Learning* **29**, 131–163.
23. Friedman, N. and M. Goldszmidt: 1997, 'Sequential update of Bayesian network structure'. In: D. Geiger and P. Shanoy (eds.): *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*.
24. Friedman, N. and Z. Yakhini: 1996, 'On the sample complexity of learning Bayesian networks'. In: *Proceedings of the twelfth conference on uncertainty in artificial intelligence*.
25. Gilks, W., S. Richardson, and D. Spiegelhalter: 1996, *Markov chain Monte Carlo in practice*. Chapman and Hall.
26. Heckerman, D.: 1998, 'A Tutorial on Learning with Bayesian Networks'. In: M. I. Jordan (ed.): *Proceedings of the NATO Advanced Study Institute on Learning in Graphical Models*. Kluwer Academic Publishers.
27. Heckerman, D., D. Geiger, and D. M. Chickering: 1995, 'Learning Bayesian networks: The combination of knowledge and statistical data'. *Machine Learning* **20**, 197–243.
28. Heckerman, D. and D. Gieger: 1995, 'Learning Bayesian networks: A unification for discrete and Gaussian domains'. In: P. Besnard and S. Hanks (eds.): *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence*. pp. 274–284.
29. Heckerman, D., C. Meek, and G. Cooper: 1997, 'A Bayesian approach to causal discovery'. Technical Report MSR-TR-97-05, Microsoft Research.
30. Hershberger, D. and H. Kargupta: 1999, 'Distributed Multivariate Regression Using Wavelet-based Collective Data Mining'. Technical Report EECS-99-02, School of EECS, Washington State University. To be published in the Special Issue on Parallel and Distributed Data Mining of the Journal of Parallel

- Distributed Computing, Guest Eds: Vipin Kumar, Sanjay Ranka, and Vineet Singh.
31. Jensen, F.: 1996, *An Introduction to Bayesian Networks*. Springer.
 32. Johnson, E. and H. Kargupta: 1999, 'Collective, Hierarchical Clustering from Distributed, Heterogeneous Data'. In: *Lecture Notes in Computer Science*, Vol. 1759. pp. 221–244.
 33. Kargupta, H., W. Huang, S. Krishnamrthy, H. Park, and S. Wang: 2000a, 'Collective Principal Component Analysis from Distributed, Heterogeneous Data'. In: D. Zighed, J. Komorowski, and J. Zytkow (eds.): *Proceedings of the Principles of Data Mining and Knowledge Discovery Conference*, Vol. 1910. Berlin, pp. 452–457. Lecture Notes in Computer Science.
 34. Kargupta, H., W. Huang, S. Krishnamurthy, and E. Johnson: 2000b, 'Distributed Clustering Using Collective Principle Component Analysis'. In: *Proceedings of the ACM SIGKDD Workshop on Distributed and Parallel Knowledge Discovery in Databases*. pp. 8–19.
 35. Kargupta, H., B. Park, D. Hersherberger, and E. Johnson: 2000c, 'Collective data mining: A new perspective toward distributed data mining'. In: H. Kargupta and P. Chan (eds.): *Advances in Distributed and Parallel Knowledge Discovery*. Menlo Park, California, USA: AAAI/ MIT Press, pp. 133–184.
 36. Kargupta, H., K. Sivakumar, W. Huang, R. Ayyagari, R. Chen, B.-H. Park, and E. Johnson: 2001, 'Towards Ubiquitous Mining of Distributed Data'. In: R. Grossman, C. Kamath, P. Kegelmeyer, V. Kumar, and R. Namburu (eds.): *Data Mining for Scientific and Engineering Applications*. Kluwer Academic Publishers.
 37. Kenji, Y.: 1997, 'Distributed Cooperative Bayesian Learning Strategies'. In: *Proceedings of the Tenth Annual Conference on Computational Learning Theory*. Nashville, Tennessee, pp. 250–262.
 38. King, R. and M. Novak: 1996, 'Supporting Information Infrastructure for Distributed, Heterogeneous Knowledge Discovery'. In: *Proceedings of SIGMOD 96 Workshop on Research Issues on Data Mining and Knowledge Discovery*. Montreal, Canada.
 39. Lam, W. and F. Bacchus: 1994, 'Learning Bayesian belief networks: An Approach based on the MDL Principle'. *Computational Intelligence* **10**, 262–293.
 40. Lam, W. and A. M. Segre: 1997, 'Distributed Data Mining of Probabilistic Knowledge'. In: *Proceedings of the 17th International Conference on Distributed Computing Systems*. Washington, pp. 178–185.
 41. Lauritzen, S. L.: 1995, 'The EM algorithm for graphical association models with missing data'. *Computational Statistics and Data Analysis* **19**, 191–201.

42. Lauritzen, S. L. and D. J. Spiegelhalter: 1988, 'Local Computations with Probabilities on Graphical Structures and their Application to Expert Systems (with discussion)'. *Journal of the Royal Statistical Society, series B* **50**, 157–224.
43. Lazarevic, A., D. Pokrajac, and Z. Obradovic: 2000, 'Distributed Clustering and Local Regression for Knowledge Discovery in Multiple Spatial Databases'. In: *Proc. 8th European Symposium on Artificial Neural Networks*. Bruges, Belgium.
44. Madigan, D. and A. Raftery: 1994, 'Model selection and accounting for model uncertainty in graphical models using Occam's window'. *Journal of the American Statistical Association* **89**, 1535–1546.
45. McClean, S., B. Scotney, and K. Greer: 2000, 'Clustering Heterogeneous Distributed Databases'. In: *Workshop on Distributed and Parallel Knowledge Discovery at KDD-2000*. Boston, pp. 20–29.
46. Meila, M. and M. I. Jordan: 1998, 'Estimating dependency structure as a hidden variable'. In: *NIPS*.
47. Park, H., H. Kargupta, E. Johnson, R. E. Sanseverino, L. D. Silvestre, and D. Hersherberger: 2000, 'Distributed, collaborative data analysis from heterogeneous sites using a scalable evolutionary technique'. Technical Report EECS98-001, School of Electrical Engineering and Computer Science, Washington State University.
48. Parthasarathy, S. and M. Ogihara: 2000, 'Clustering Distributed Homogeneous Datasets'. In: D. Zighed, J. Komorowski, and J. Zytkow (eds.): *Proceedings of the Principles of Data Mining and Knowledge Discovery Conference*. pp. 566–574.
49. Pazzani, M. and D. Billsus: 1997, 'Learning and revising user profiles: The identification of interesting web sites'. *Machine Learning* **27**, 313–331.
50. Pazzani, M., J. Muramatsu, and D. Billsus: 1996, 'Syskill & Webert: Identifying Interesting Web Sites'. In: *Proceedings of the National Conference on Artificial Intelligence*.
51. Pearl, J.: 1988, *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
52. Pearl, J.: 1993, 'Graphical models. causality and intervention'. *Statistical Science* **8**, 266–273.
53. Prodromidis, A. and S. Stolfo: 2000, 'Cost complexity-based pruning of ensemble classifiers'. In: *Workshop on Distributed and Parallel Knowledge Discovery at KDD-2000*. Boston, pp. 30–40.
54. Provost, F. J. and B. Buchanan: 1995, 'Inductive Policy: The Pragmatics of Bias Selection'. *Machine Learning* **20**, 35–61.
55. Sayal, M. and P. Scheuermann: 2000, 'A Distributed Clustering Algorithm for Web-Based Access Patterns'. In: *Workshop on Distributed and Parallel Knowledge Discovery at KDD-2000*. Boston, pp. 41–48.

56. Singh, M.: 1997, 'Learning Bayesian networks from incomplete data'. In: *Proceedings of the National Conference on Artificial Intelligence*. pp. 27–31.
57. Singh, M. and G. M. Provan: 1995, 'A comparison of induction algorithms for selective and non-selective Bayesian classifiers'. In: A. Prieditis and S. Russel (eds.): *Proceedings of the Twelfth International Conference on Machine Learning*. pp. 497–505.
58. Spiegelhalter, D. J. and S. L. Lauritzen: 1990, 'Sequential updating of conditional probabilities on directed graphical structures'. *Networks* **20**, 570–605.
59. Spirtes, P., C. Glymour, and R. Scheines: 1993, *Causation, Prediction and Search*, No. 81 in Lecture Notes in Statistics. Springer-Verlag.
60. Suzuki, J.: 1993, 'A construction of Bayesian networks from databases based on an MDL scheme'. In: D. Heckerman and A. Mamdani (eds.): *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence*. pp. 266–273.
61. Thiesson, B.: 1995, 'Accelerated quantification of Bayesian networks with incomplete data'. In: *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*. pp. 306–311.
62. Thiesson, B., C. Meek, D. M. Chickering, and D. Heckerman: 1998, 'Learning mixtures of Bayesian networks'. In: *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*.
63. Thomas, A., D. Spiegelhalter, and W. Gilks: 1992, 'Bugs: A program to perform Bayesian inference using Gibbs sampling'. In: J. Bernardo, J. Berger, A. Dawid, and A. Smith (eds.): *Bayesian Statistics*. Oxford University Press, pp. 837–842.
64. Tumer, K. and J. Ghosh: 2000, 'Robust Order Statistics based Ensemble for Distributed Data Mining'. In: *Advances in Distributed and Parallel Knowledge Discovery*. MIT, pp. 185–210.
65. Turinsky, A. and R. Grossman: 2000, 'A Framework for Finding Distributed Data Mining Strategies That are Intermediate Between Centralized Strategies and In-Place Strategies'. In: *Workshop on Distributed and Parallel Knowledge Discovery at KDD-2000*. Boston, pp. 1–7.
66. Zweig, G. and S. J. Russel: 1998, 'Speech recognition with dynamic Bayesian networks'. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence*.