

Towards Gaussian Bayesian Network Fusion

Irene Córdoba-Sánchez, Concha Bielza, and Pedro Larrañaga

Departamento de Inteligencia Artificial,
Universidad Politécnica de Madrid, Madrid, Spain
irene.cordoba.sanchez@alumnos.upm.es
{mcbielza,pedro.larranaga}@fi.upm.es

Abstract. Data sets are growing in complexity thanks to the increasing facilities we have nowadays to both generate and store data. This poses many challenges to machine learning that are leading to the proposal of new methods and paradigms, in order to be able to deal with what is nowadays referred to as Big Data. In this paper we propose a method for the aggregation of different Bayesian network structures that have been learned from separate data sets, as a first step towards mining data sets that need to be partitioned in an horizontal way, i.e. with respect to the instances, in order to be processed. Considerations that should be taken into account when dealing with this situation are discussed. Scal-able learning of Bayesian networks is slowly emerging, and our method constitutes one of the first insights into Gaussian Bayesian network aggregation from different sources. Tested on synthetic data it obtains good results that surpass those from individual learning. Future research will be focused on expanding the method and testing more diverse data sets.

1 Introduction

Nowadays, we are entering the *era of Big Data*, as a result of both the generalised trend of massive data collection and the increasing computer capabilities for processing and storage. These data sets are characterized mainly for their huge volume and complexity (they can be noisy, have a fast change rate, etc.). Machine learning methods are rapidly being revised and new paradigms are arising in order to be able to adapt to this kind of data.

One of the main approaches for dealing with high volume of data is to partition it across a cluster, perform some operations and then aggregate the results. This partition can be either horizontal (across the instances) or vertical (across the variables). Horizontal partitions can also naturally arise when we want to jointly analyse information contained at different sources, e.g. records of patients in different hospitals that store the same variables about each of them.

Bayesian networks (BNs) are well-known tools for modelling and dealing with uncertain knowledge and data. Their aggregation has been studied since the days of their conception as belief models from an expert. Martzkevich and

Abramson [7] consider the problem of fusing networks from different experts which shared some variables. They provide an algorithm which seeks to obtain a graph containing all the nodes and arcs from the individual networks, or their reversals. This however may not be the case of interest always when we are thinking about fusing networks that have been learned from data, since the individual networks in this case may contain spurious connections.

The work by del Sagrado and Moral [1] focuses on studying the fusion of DAGs by means of intersection and union of the independence statements represented by each of the involved networks. Richardson and Domingos [11] use knowledge from a group of experts to compute a prior distribution over the BN structures. They motivate their proposal by stating that knowledge elicitation can be facilitated if we allow experts to be noisy on their statements about the BNs, and make up for this flexibility by using multiple different experts. This argument is interesting because it can be compared with the case of huge, noisy data sets, where instead of sub-sampling and learning an individual network, an alternative approach could be to learn multiple networks on different partitions of data and aggregate them afterwards.

Another use case where horizontal partitioning arises naturally is the problem described by López-Cruz *et al.* [5]. In this case a set of experts were asked to classify different neurons, giving rise to one supervised training set from each expert. A cluster process was applied to the set of the individual BNs obtained and a representative BN for each cluster was constructed. These representative networks were then aggregated into a Bayesian multinet.

Regarding the aggregation of parameters, recently Etminani *et al.* [2] propose a method in which they cluster experts' parameters and aggregate only those that correspond to the cluster with the highest number of members, resembling democratic societies. Other popular strategies for parameter fusion in Bayesian networks are Linear Opinion Pools (LinOP) [6] and Logarithmic Opinion Pools (LogOP) [9].

We propose a method for the aggregation of Gaussian BNs (GBNs), which to the best of our knowledge is the first proposal of this kind. It covers both the structure of the network and the parameters of the Gaussian distribution encoded by it. The experiments carried out show promising results for the proposed method.

The paper is organised as follows. Section 2 introduces the necessary background knowledge for the rest of the paper. In Sect. 3 the details of the method are described, whose results from experimental evaluation are discussed in Sect. 4. Finally, the conclusions and future research lines are presented in Sect. 5.

2 Preliminaries

2.1 Bayesian Networks

A BN can be defined as a way of representing the factorization of a joint probability distribution over a random vector $\mathbf{X} = (X_1, \dots, X_p)$, where $\mathbf{Pa}(X_i)$ are called the *parents* of X_i ,

$$f(\mathbf{x}) = \prod_{i=1}^p f(x_i | \mathbf{pa}(x_i)). \quad (1)$$

A BN consists on a qualitative part, commonly called the structure, and a quantitative component, the parameters. More formally, it is defined [10] as a pair (G, Θ) , where G is a DAG and Θ are the numerical parameters which define the factorization in Eq. (1). The nodes of G are the components of \mathbf{X} and its arcs represent probabilistic dependencies between the variables, in such a way that the DAG satisfies the *Markov condition*: each variable is conditionally independent of its non-descendants given its parents. Two DAGs are *Markov equivalent* if they represent the same set of conditional independences between the variables. This defines a binary relation which gives rise to equivalence classes and partitions the DAG space.

In order to learn a BN from data it is necessary to learn both the structure (G) and the numerical parameters (Θ). There are two main approaches for BN structure learning: constraint based and score-and-search. Constraint based methods try to find the Bayesian network structure that represents most of the dependence relations present in data, detected by means of statistical tests. The PC algorithm [14], which has as output an equivalence class of DAGs, is a representative example of these types of methods.

On the other hand, score-and-search methods try to find the structure that best fits the data. They are characterized by a representation of the solution space, a search method and a score. The KES algorithm [8] is an example of such methods, which performs the search in the equivalence class space. Searching in this space has several advantages when compared to the DAG space, such as it being a more efficient and robust representation [16], although there is still some controversy regarding this choice. Many search heuristics and scores can be combined and give rise to the different methods appearing in the literature.

2.2 Gaussian Bayesian Networks

A GBN [4] encodes a joint Gaussian distribution over \mathbf{X} , i.e., with joint density function

$$f(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^p |\boldsymbol{\Sigma}|}} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^t \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\},$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_p)$ is the vector of unconditional means and $\boldsymbol{\Sigma}$ is the covariance matrix. Each factor in Eq. (1) corresponds in this case to a univariate normal distribution,

$$f(x_i | \mathbf{pa}(x_i)) \equiv \mathcal{N} \left(\mu_i + \sum_{x_j \in \mathbf{pa}(x_i)} \beta_{ji} (x_j - \mu_j), v_i \right), \quad (2)$$

where β_{ji} reflects the strength of the relationship between X_i and its j -th parent, and v_i is the conditional variance of X_i given its parents, i.e.,

$$v_i = \sigma_i - \boldsymbol{\Sigma}_{i\mathbf{pa}(X_i)} \boldsymbol{\Sigma}_{\mathbf{pa}(X_i)}^{-1} \boldsymbol{\Sigma}_{\mathbf{pa}(X_i) i}^t. \quad (3)$$

In Eq. 3 σ_i is the unconditional variance of X_i , $\Sigma_i \mathbf{Pa}(X_i)$ is the matrix of covariances between X_i and $\mathbf{Pa}(X_i)$, and $\Sigma_{\mathbf{Pa}(X_i)}$ is the covariance matrix of $\mathbf{Pa}(X_i)$.

Thus, the parameters of a GBN are the vector of means $\boldsymbol{\mu}$, the vector of conditional variances \mathbf{v} and the coefficients β_{ji} . Assuming standardized data ($\mu_i = 0$ and $v_i = 1$), the parameter estimation is reduced to solving the linear regression model

$$x_i = \sum_{x_j \in \mathbf{pa}(x_i)} \beta_{ji} x_j + \epsilon_i,$$

with ϵ_i being the Gaussian noise term with zero expectation.

3 Method

Although the aggregation of different individual GBNs is a first step towards the analysis of massive data, where the data set would be split into slices distributed across a cluster, here we will assume that we already have different data sets over the same variables available (i.e., at this stage we are not concerned with the preprocessing and splitting processes).

The structure learning method we have used for learning the individual networks is the score-and-search hill climbing [15] with the Bayesian information criterion (BIC) [12] score on the DAG space. After each network has been learned, they are aggregated using majority vote below a threshold. This procedure is outlined in Algorithm 1.

Algorithm 1. Structure learning

Input: *datasets*. Data sets from where the individual Bayesian networks will be learned.

Input: *threshold*. Threshold for the majority arc voting.

Output: Aggregated Bayesian network structure learned with the specified thresholds.

```

1:  $n\_bn \leftarrow \text{size}(\text{datasets});$ 
2:  $bn\_list \leftarrow \text{list}();$ 
3: for  $i \in \{1, n\_bn\}$  do
4:    $bn\_list[i] \leftarrow \text{learn\_struc}(\text{datasets}[i])$ 
5: end for
6:  $v\_matrix \leftarrow \text{get\_votes}(bn\_list);$ 
7:  $result \leftarrow \text{bn\_aggr}(\text{threshold}, v\_matrix);$ 
8: return  $result;$ 

```

The functions *get_votes* and *bn_aggr* in Algorithm 1 are further detailed in Algorithms 2 and 3 respectively. *get_votes* consists of the process of extracting how many networks contribute to the same arc, i.e., how common across the learned networks an arc is. Thus a matrix containing the sums of the appearances of each arc in the networks is obtained. The threshold for the majority vote is the

Algorithm 2. `get_votes`

Input: `bn_list`. List of BN structures already learned on each data set.

Output: Matrix containing the votes for each arc.

```
1:  $n\_nodes \leftarrow \text{nodes}(bn)$ ;  
2:  $v\_matrix \leftarrow \text{matrix}(n\_nodes, n\_nodes)$ ;  
3: for  $bn \in bn\_list$  do  
4:    $bn\_arcs \leftarrow \text{arcs}(bn)$ ;  
5:   for  $arc \in bn\_arcs$  do  
6:      $from \leftarrow \text{from}(arc)$ ;  
7:      $to \leftarrow \text{to}(arc)$ ;  
8:      $v\_matrix[from][to] \leftarrow v\_matrix[from][to] + 1$ ;  
9:   end for  
10: end for  
11: return  $v\_matrix$ ;
```

main parameter of this method and we will analyse it further on the experimental section.

In `bn_aggr` the threshold is compared with each of the entries in the matrix of arcs, and the corresponding arc is added to the final network if its value reaches the threshold. In the same algorithm we can notice that when an arc addition causes a cycle in the DAG it is discarded.

Algorithm 3. `bn_aggr`

Input: `threshold`. Threshold for the arc voting.

Input: Matrix containing the votes for each arc.

Output: `bn`. Aggregated Bayesian network.

```
1:  $bn \leftarrow \text{empty\_dag}()$ ;  
2: for  $i \in \text{cols}(v\_matrix)$  do  
3:   for  $j \in \text{rows}(v\_matrix)$  do  
4:     if  $v\_matrix[i][j] \geq \text{threshold}$  then  
5:       if not  $\text{arc\_causes\_cycle}(bn, i, j)$  then  
6:          $\text{add\_arc}(bn, i, j)$ ;  
7:       end if  
8:     end if  
9:   end for  
10: end for  
11: return  $bn$ ;
```

After the aggregation of the structure has finished, the linear regression coefficients of each variable on its parents is learned by maximum likelihood estimation (MLE) from each data set, but this time using the aggregated structure. This is what del Sagrado and Moral [1] call *topological fusion*, that is, obtaining a consensus structure and then estimating the model parameters, as opposed to *graphical representation of consensus*, which consists of aggregating the probability distributions of each network and then obtaining the structure that represents it.

The aggregation of the parameters obtained from each data set is performed using the method explained hereafter. Consider a multiple linear regression model on $\{X_1, \dots, X_n\}$ predictors. Assume that the data is distributed across k slices. Let $\hat{\beta}_j = (\hat{\beta}_{1j}, \dots, \hat{\beta}_{nj})$ be the vector of estimates obtained in slice j . For each predictor X_i , $i \in \{1, \dots, n\}$, let

$$\tilde{\beta}_i = \sum_{j=1}^k w_{ij} \hat{\beta}_{ij}$$

be the aggregated estimate, where $w_{ij} = \sigma_{ij}^{-2} / \sum_{j=1}^k \sigma_{ij}^{-2}$, $\sigma_{ij} = \text{var}(\hat{\beta}_{ij})$. Because we are dealing with GBNs, MLE is equivalent to the least squares (LS) method, and thus $\tilde{\beta}_i$ is the estimator of minimum variance [3] among those with form

$$\sum_{j=1}^k w_{ij} \hat{\beta}_{ij}, \text{ where } \sum_{j=1}^k w_{ij} = 1,$$

Asymptotic normality is also established on Fan *et al.* [3].

The pseudo-code of the outlined procedure for learning the parameters of the linear regression for each variable on its parents can be found in Algorithm 4.

Algorithm 4. Parameter learning

Input: *datasets*. Data sets from where the individual parameters will be learned.

Output: Bayesian network parameters aggregated.

```

1:  $n\_bn \leftarrow \text{size}(\text{datasets});$ 
2:  $\text{param\_list} \leftarrow \text{list}();$ 
3: for  $i \in \{1, n\_bn\}$  do
4:    $\text{param\_list}[i] \leftarrow \text{learn\_param}(\text{datasets}[i]);$ 
5: end for
6:  $n\_param \leftarrow \text{size}(\text{param\_list});$ 
7:  $n\_nodes \leftarrow \text{nodes}(\text{param\_list});$ 
8:  $\text{param} \leftarrow \text{matrix}(n\_nodes, n\_nodes);$ 
9: for  $i \in \{1, n\_param\}$  do
10:  for  $\text{node} \in \text{param\_list}[i]$  do
11:    for  $\text{parent} \in \text{parents}(\text{node})$  do
12:       $\text{coef} \leftarrow \text{get\_coeff}(\text{param\_list}[i], \text{node}, \text{parent});$ 
13:       $\text{weight} \leftarrow \text{get\_weight}(\text{param\_list}[i], \text{node}, \text{parent});$ 
14:       $\text{param}[\text{node}][\text{parent}] \leftarrow \text{param}[\text{node}][\text{parent}] + \text{coef} * \text{weight};$ 
15:    end for
16:  end for
17: end for
18: return  $\text{normalize}(\text{param});$ 

```

4 Experiments

We have used some utilities from the R package `bnlearn` [13] and tested the proposed method using synthetic data sets generated from a multivariate Gaussian distribution whose DAG structure is shown in Fig. 1.




Fig. 1. Structure of the Bayesian network used for the experiments.

In a real use case of this method we could have been given a number of separate data sets over the same variables but differing in the number of instances each one contains. On the other hand, if we were to apply it to a huge data set, the different partitions would probably contain a similar amount of instances. In this synthetic experiment we have generated eight different data sets with a sample size of 50 instances each.

Figure 2 shows the different BN structures obtained from each data set. We can notice that a high portion of the original network is learned in most of the cases, being false positives the most common error. We have aggregated the results using all possible values for the threshold parameter, getting as result the networks that appear on Fig. 3.




Fig. 2. Structure learned on each of the data sets. Green arcs are those correctly learned, red arcs are false positives and blue arcs are false negatives (Color figure online).




Fig. 3. Aggregated structures for the different thresholds (1 to 8), increasing from left to right and top to bottom. Green arcs are those correctly aggregated, red arcs are false positives and blue arcs are false negatives (Color figure online).

The metrics we are going to use to evaluate the results obtained with respect to the original structure are the false positive, false negative and true positive rate and the Structural Hamming Distance (SHD) [15]. The latter consists of the number of operations needed to match the Partial DAGs (PDAGs) representing the equivalence classes of each network. The operations considered are arc addition, deletion or reversal and edge addition or deletion. Therefore, the PDAG is extracted from the respective DAGs before calculating this metric. SHD provides a way to compare the two BNs in terms of the conditional independencies encoded by the BN, and thus avoids the penalization of differences in arcs that might be statistically undistinguishable.

In Table 1 we can see the value metrics for the individual structures (left) and the aggregated ones (right) when compared with the original network.

Table 1. Results of the GBN learned on each data set (left) and the aggregated GBN (right) compared with the original network. TP, FP and FN indicate the true positive, false positive and false negative rates (respectively). SHD denotes the Structural Hamming Distance. The networks are numbered according to their order of appearance in Fig. 2

Network	SHD	TP	FP	FN	Threshold	SHD	TP	FP	FN
1	3	7	2	0	1	8	7	8	0
2	3	6	2	1	2	3	7	3	0
3	2	7	2	0	3	0	7	0	0
4	3	6	2	1	4	0	7	0	0
5	2	7	2	0	5	0	7	0	0
6	0	7	0	0	6	1	6	0	1
7	0	7	0	0	7	1	6	0	1
8	5	6	4	1	8	1	6	0	1

Obviously when the threshold is 1 every arc that appears in the individual networks is added to the final one (unless a cycle is caused), so this produces worse results than any individual network when comparing it to the original structure. However, for thresholds above 1 the aggregated result is better than most of the isolated ones, because thanks to the majority threshold false dependences are eliminated. For too restrictive thresholds this can however result in the deletion of a valid arc, so it should be adjusted to an intermediate value for the best results. In a real use case this would depend on the data characteristics, the application domain and the availability of a training set.

Finally, parameter learning is influenced by noise if the structure is not correctly learned because false parents of variables arise, which means that false coefficients are estimated in the linear regression. However it is the case, as one would expect, that the coefficients corresponding to these false parents are very close to zero (e.g., 0.007 mean for the extreme case of threshold 1) and the variations on the value of the other parents are barely noticeable. This is not the case when we learn from a single network. For example, in the last network in Fig. 2, $\beta_{BF} = 1.5$, and B is a false parent of F .

5 Conclusions and Future Work

We have considered the problem of horizontal partitioning in the context of Big Data and proposed a method for aggregating several GBNs learned from different data sets as a first step towards scalable GBN learning. The method obtains good results both in the case of structure and parameter learning on synthetic data. The aggregated results surpass in most cases those derived from learning from a single data set by taking into account all the data available without the need of analysing it as a single block. This is specially useful for its potential applications when analysing partitions of massive data sets.

As a future line of research, when learning the aggregated structure the treatment of cycles will be refined and will involve more sophisticated techniques such as arc reversal, checking the strength of the connection in each of the individual networks (coefficients of the regression), establishing a suitable ordering of arc consideration, etc.

In the case of applying the proposed method on a distributed setting, where each data set is in a computer within a cluster, it would be interesting to define some communication protocol during the learning process. This would be useful for gathering stepwise information that could be used later on for example when aggregating the individual networks (e.g. cycles) but also for developing more sophisticated voting schemes which could depend on the adequateness of each data set for the learning process.

Finally, we will also focus on performing more testing with diverse (noise, missing values) and real data sets.

Acknowledgements. The authors thank the reviewers for comments and critics which significantly contributed to improve the paper; and also J.M. Peña, J. Nielsen,

J. Mengin and M. Serrurier for the valuable help. This work has been partially supported by the Spanish Ministry of Economy and Competitiveness through the Cajal Blue Brain (C080020-09; the Spanish partner of the Blue Brain initiative from EPFL) and TIN2013-41592-P projects, and by the Regional Government of Madrid through the S2013/ICE-2845-CASI-CAM-CM project.

References

1. del Sagrado, J., Moral, S.: Qualitative combination of Bayesian networks. *Int. J. Intell. Syst.* **18**(2), 237–249 (2003)
2. Etminani, K., Naghibzadeh, M., Peña, J.M.: DemocraticOP: A democratic way of aggregating Bayesian network parameters. *Int. J. Approximate Reasoning* **54**(5), 602–614 (2013)
3. Fan, T.H., Lin, D.K., Cheng, K.F.: Regression analysis for massive datasets. *Data Knowl. Eng.* **61**(3), 554–562 (2007)
4. Geiger, D., Heckerman, D.: Learning Gaussian networks. In: *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, pp. 235–243. Morgan Kaufmann Publishers Inc. (1994)
5. López-Cruz, P.L., Larrañaga, P., DeFelipe, J., Bielza, C.: Bayesian network modeling of the consensus between experts: an application to neuron classification. *Int. J. Approximate Reasoning* **55**(1), 3–22 (2014)
6. Maynard-Reid II, P., Chajewska, U.: Aggregating learned probabilistic beliefs. In: *Proceedings of the Eighteenth International Conference on Uncertainty in Artificial Intelligence*, pp. 354–361. Morgan Kaufmann Publishers Inc. (2001)
7. Matzkevich, I., Abramson, B.: The topological fusion of Bayes nets. In: *Proceedings of the Eighth International Conference on Uncertainty in Artificial Intelligence*, pp. 191–198. Morgan Kaufmann Publishers Inc. (1992)
8. Nielsen, J.D., Kočka, T., Peña, J.M.: On local optima in learning Bayesian networks. In: *Proceedings of the Nineteenth International Conference on Uncertainty in Artificial Intelligence*, pp. 435–442. Morgan Kaufmann Publishers Inc. (2002)
9. Pennock, D.M., Wellman, M.P.: Graphical representation of consensus belief. In: *Proceedings of the Fifteenth International Conference on Uncertainty in Artificial Intelligence*, pp. 531–540. Morgan Kaufmann Publishers Inc. (1999)
10. Pearl, J.: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc, California (1988)
11. Richardson, M., Domingos, P.: Learning with knowledge from multiple experts. In: *Proceedings of the Twentieth International Conference on Machine Learning*, pp. 624–631. AAAI Press (2003)
12. Schwarz, G.: Estimating the dimension of a model. *Ann. Stat.* **6**(2), 461–464 (1978)
13. Scutari, M.: Learning Bayesian networks with the bnlearn R package. *J. Stat. Softw.* **35**(3), 1–22 (2010)
14. Spirtes, P., Glymour, C.N., Scheines, R.: *Causation, Prediction, and Search*. MIT press, Cambridge (2000)
15. Tsamardinos, I., Brown, L.E., Aliferis, C.F.: The max-min hill-climbing Bayesian network structure learning algorithm. *Mach. Learn.* **65**(1), 31–78 (2006)
16. Vidaurre, D., Bielza, C., Larrañaga, P.: Learning an L1-regularized Gaussian Bayesian network in the equivalence class space. *IEEE Trans. Syst. Man Cybern. Part B Cybern.* **40**(5), 1231–1242 (2010)