

Scalable Bayesian Network Structure Learning with Splines

Charupriya Sharma

C9SHARMA@UWATERLOO.CA

Peter van Beek

VANBEEK@UWATERLOO.CA

David R. Cheriton School of Computer Science, University of Waterloo

Abstract

The graph structure of a Bayesian network (BN) can be learned from data using the well-known score-and-search approach. Previous work has shown that incorporating structured representations of the conditional probability distributions (CPDs) into the score-and-search approach can improve the accuracy of the learned graph. In this paper, we present a novel approach capable of learning the graph of a BN and simultaneously modelling linear and non-linear local probabilistic relationships between variables. We achieve this by a combination of feature selection to reduce the search space for local relationships and extending the score-and-search approach to incorporate modelling the CPDs over variables as Multivariate Adaptive Regression Splines (MARS). MARS are polynomial regression models represented as piecewise spline functions. We show on a set of discrete and continuous benchmark instances that our proposed approach can improve the accuracy of the learned graph while scaling to instances with a large number of variables.

1. Introduction

Bayesian networks (BNs) are widely used probabilistic graphical models with applications in knowledge discovery, decision support, and prediction (Darwiche, 2009; Koller and Friedman, 2009). A BN represents a joint distribution over random variables as a product of conditional probability distributions (CPDs) for each variable, and it enforces a directed acyclic graph (DAG) structure over the corresponding network. The *score-and-search* approach is often used to learn the structure of a BN from data, where a *scoring function* is used to evaluate the quality of fit of a BN over the data.

The Bayesian Information Criterion (BIC) (Schwarz, 1978) is one of the most widely used scoring functions to learn BNs. It is expressed as the difference of a likelihood and a penalty term for model complexity, and gives theoretical guarantees that enable significant pruning of the search space of DAGs. However, in practice, BIC tends to heavily penalise large parent sets in a DAG, making it less accurate for datasets which need more complex models.

Implementations of score-and-search usually employ simple conditional probability tables (CPTs) to represent the CPD for a variable (e.g., (Bartlett and Cussens, 2013; Chickering, 2002a; van Beek and Hoffmann, 2015)). In settings with limited data or with a large number of variables, such a representation might not be efficient, because CPTs grow exponentially with the size of the parent set. Using CPTs can also result in overfitting. These two disadvantages have led to research of structured representations for CPDs that can model complex relationships between a child and its parents without exponential size and overfitting risks.

One example of a structured representation is context-specific independence (Boutilier et al., 1996), which is a popular method to represent independence given a certain assignment of a subset of variables. Tree structures were extended to be represented as decision graphs (Chickering et al., 1997) with equality constraints. However, the partitioning algorithms used to split the feature space become computationally infeasible for discrete variables with large domains. Other frequently used

representations include the Noisy-OR (Pearl, 1988), which assumes a causal independence within the parent set and allows CPTs to be represented with parameters linear in the size of the parent set, and logistic regression (Buntine, 1991). Most studies of structured representations e.g., (Xiang and Jia, 2007; Zhang and Poole, 1996) do not impose global structural constraints on the DAG.

Friedman and Goldszmidt (1998) were the first to extend the *score-and-search* approach to include structured representations in Bayesian network structure learning (BNSL). They showed this improved results in inference and allowed exploration of more complex network structures. Talvitie et al. (2019) gave a tree-based score to find the optimal BN with some structural constraints. Their experiments with structured CPDs showed an improvement in the ability of the search algorithm to find correct BNs with less data on some real-world datasets, but on other datasets, CPTs performed better. Sharma et al. (2020) extended score-and-search to the Noisy-OR representation and let the search algorithm have the choice between a structured and a Noisy-OR representation. They observed their learned networks to have a mixture of Noisy-ORs and standard CPTs, which improved overall network scores. This shows that there needs to be some flexibility about the structured representation used for the CPD to enable choosing the best fitting representation to find the correct overall structure and maintain the complexity advantage in inference.

Other approaches to scale beyond small network sizes include using neural networks to solve a continuous formulation of the DAG learning problem (Yu et al., 2019; Zheng et al., 2018). Scana-gatta et al. (2015) used approximate search for discrete benchmarks, but did not allow continuous variables or structured representations.

In this paper, we propose the first score-and-search approach for learning Bayesian networks with Multivariate Adaptive Regression Splines (MARS) as possible representations for CPDs and do not make a priori assumptions on the structure of the DAG. MARS are polynomial regression models represented as piecewise spline functions. Our approach can effectively prune most candidate parent sets of a variable by leveraging feature learning algorithms.

Our contributions are the following:

- Our scoring method can model both linear and non-linear interactions between variables without imposing any rigid constraints on their structure while scoring them simultaneously (Section 3). Variables are scored independently of each other.
- Our scoring algorithm can identify high-degree variable interactions without a limit on the number of variables or terms in the model (Section 3).
- Our scoring algorithm can compute structured representations for both discrete and continuous valued variables (Section 3).
- We provide a feature selection method to identify the most promising parents for a node, and scale our score-and-search approach to very large networks with 1000+ nodes (Section 4).
- We show empirically on a set of discrete and continuous benchmark instances that our proposed approach can improve the accuracy of the learned DAG (Section 5).

2. Preliminaries

A BN represents a joint distribution over random variables $\mathbf{V} = \{V_1, \dots, V_n\}$ with a labeled directed acyclic graph (DAG), $G = (\mathbf{V}, E)$ where the edges E represent direct influence of one

random variable on another, and each node V_i is labeled with a conditional probability distribution $P(V_i | \Pi_i)$ that defines the relationship of the variable V_i with its set of parents Π_i in the DAG and $P(V_1, \dots, V_n) = \prod_{i=1}^n P(V_i | \Pi_i)$. Each variable V_i has a state space Ω_i . Each set of parents $\Pi_i = \{V_{i1}, \dots, V_{i|\Pi_i|}\}$ has as state space the set of candidate instantiations of the nodes in Π_i , $\Omega_{\Pi_i} = \Omega_{i1} \times \dots \times \Omega_{i|\Pi_i|}$.

Score-and-search: Given a set of N instances $I = \{I_1, \dots, I_N\}$, where each instance I_i is an n -tuple that is a complete instantiation of the n variables in \mathbf{V} , a *scoring function* $\sigma(G | I)$ assigns a real value to a DAG G to measure the quality of its fit to data I . When the dataset is clear from context, we express this as $\sigma(G)$. The objective of Bayesian network structure learning (BNSL) is to find the network $\operatorname{argmin}_G \sigma(G)$, and a lower score is assumed to represent a better quality fit without loss of generality.

3. MARS: Multivariate Adaptive Regression Splines

Linear models like linear and logistic regression are a simple, intuitive and fast to compute way to represent data. However, they impose linear relations on the data which limits their accuracy if the ground truth contains non-linear relationships. Multivariate Adaptive Regression Splines (MARS) (Friedman, 1991) are an extension of linear models that combines linear basis functions to model non-linearities and variable interactions. MARS models are of the form,


$$M(x) = c_0 + \sum_{j=1}^k c_j B_j(x), \quad (1)$$

where each *model coefficient* c_j for $j \in [0, k]$ is a constant, and each B_j is a *basis function*. Basis functions take one of the two forms:

1. A *hinge function* of the form $(x - t)_+ = \max(0, x - t)$ or $(t - x)_+ = \max(0, t - x)$, where t is a constant, called a *knot*.
2. A product of two or more unique hinge functions. This prevents terms of higher degree which are approximated by the hinge functions instead.

Hinge functions are zero over part of their range, and their product is nonzero only over the region where each of the component hinge functions have nonzero values. Thus, data can be partitioned into disjoint regions and processed independently, enabling processing of high dimensional inputs while still being able to capture complex non-linear relationships with high order interactions (see Figure 1).

Figure 1: MARS (red) and linear regression (blue) models for the same set of points. MARS places a knot at every change of the slope.



Variable modelling: For a variable V_i and a candidate parent set Π_i we write the MARS model for V_i and Π_i as,

$$M(\Pi_i) := c_0 + \sum_{j=1}^k c_j B_j(\Pi_i). \quad (2)$$

For the computation of the model, we are given a set of N instantiations $I = \{I_1, \dots, I_N\}$ for V_i and Π_i , i.e. $I_j = (y_i, x_1, \dots, x_{|\Pi_i|})$ with $y_i \in \Omega_i, x_j \in \Omega_{ij}$.

Model computation: To compute the basis functions in $M(\Pi_i)$, the MARS algorithm operates in two phases: 1) A forward stage which starts with the intercept c_0 and proceeds to generate many candidate basis functions in pairs $[(t - x)_+, (x - t)_+]$ with $x \in \Pi_i$. The model coefficients are estimated by minimizing the residual sum-of-squares (RSS),

$$RSS(M(\Pi_i)) = \sum_{j=1}^N (y_j - M(I_j))^2. \quad (3)$$

A candidate pair of basis functions is added to the model $M(\Pi_i)$ if it reduces the training error of $M(\Pi_i)$; i.e. if $RSS(M_{new}(\Pi_i)) - RSS(M_{old}(\Pi_i)) < 0$. This process continues until there is no significant improvement. 2) A backward stage that is used to address overfitting in $M(\Pi_i)$. The backward stage iterates over functions in $M(\Pi_i)$ and prunes terms until the subset of the terms that provide the best score remains. The subsets are scored by using a generalized cross-validation (GCV) score,

$$GCV(M(\Pi_i)) = \frac{RSS(M(\Pi_i))}{((1 - \lambda)/N)^2}, \quad (4)$$

where λ is the number of terms in $M(\Pi_i)$. We call $GCV(M(\Pi_i))$ the local score of V_i in a BN which assigns Π_i as the parent set of V_i . GCV provides a computationally fast approximation to leave-one out cross-validation, for linear fitting under squared-error loss. The GCV score can be shown to be similar to the Akaike information criterion (AIC) score, which is frequently used for model selection (Hastie et al., 2009).

Scoring BNs: GCV scores are decomposable, meaning for a BN G with parent set Π_1, \dots, Π_i assigned to variables V_1, \dots, V_n , the total GCV score can be calculated by,

$$GCV(G) = \sum_{i=1}^n GCV(M(\Pi_i)), \quad (5)$$

where $M(\Pi_i)$ is the model computed for representing V_i with candidate parent set Π_i .

4. Scaling with Feature Selection

MARS provides a fast way to score a candidate parent set for a child. However, there are an exponential number of candidate parent sets for each node. To scale the search to large networks, it is crucial to score a small subset of candidate parent sets. In this section, we describe our algorithm BN-FS-MARS (Algorithm 1) for BNSL which is based on MARS and a feature selection mechanism.

BN-FS-MARS operates in two phases: 1) Computing candidate MARS models and their GCV scores for all variables and candidate parent sets. We prune candidate parents for a child V_i by using

feature selection: we compute an ordering for the set of candidate parents $\mathbf{V} \setminus \{V_i\}$ which ranks their importance, and use the candidate parent importance rankings to create candidate parent sets. 2) The first phase provides multiple possible parent sets for each variable with an associated score for each parent set and this gives us an optimisation problem of finding the DAG G that minimizes the total score. We feed the candidate parent sets and scores from the MARS models of phase 1 into a search algorithm to construct a DAG of minimum GCV score. Algorithm 1 shows the pseudocode of the algorithm.

The performance of BN-FS-MARS has significant dependence on the quality of the variable importance ranking used in Step 3 of Algorithm 1. There are a number of feature selection methods such as random forests (Breiman, 2001), decision trees, and the Pearson correlation coefficient that can generate such a ranking (see, e.g., (Guyon and Elisseeff, 2003) and references therein). The accuracy of the feature selection method used to generate this ranking depends significantly on the amount and type of data we have available. In the next section, we will show the performance of BN-FS-MARS with random forests.

Algorithm 1 BN-FS-MARS

Inputs: A set of random variables \mathbf{V} and a set of observations I

Outputs: A DAG G minimizing $GCV(G)$

Phase 1: Compute MARS models and GCV scores for all variables guided by feature selection

- 1: Initialize $R = \emptyset$ to collect scores for Phase 2
- 2: **for** each variable $V_i \in \mathbf{V}$ **do**
- 3: $\Pi_i := \text{FilterCandidatesFS}(V_i)$
- 4: **for** $\ell = 0, \dots, |\Pi_i|$ **do**
- 5: **for** parent sets $\Pi_i \subseteq \Pi_i, |\Pi_i| = \ell$ **do**
- 6: Compute MARS model $M(\Pi_i)$ *see Eqn. (2)*
- 7: Compute $GCV(M(\Pi_i))$ *see Eqn. (4)*
- 8: $R = R \cup \{(V_i, \Pi_i, GCV(M(\Pi_i)))\}$
- 9: **if** time limit for V_i reached **then**
- 10: Continue for-loop in line 2

Phase 2: Choose MARS models for each variable to create a DAG that minimizes $GCV(G)$

- 11: Run a search algorithm on R to compute a DAG G minimizing $GCV(G)$
-

Algorithm 2 FilterCandidatesFS

Inputs: A random variable $V_i \in \mathbf{V}$, a constant $\ell < n$

Outputs: A candidate parent set Π_i

- 1: Compute a variable importance ranking for V_i using $\mathbf{V} \setminus \{V_i\}$
 - 2: Return $\Pi_i := \ell$ -most important variables based on feature selection.
-

5. Experimental Evaluation

In this section, we show the performance of our algorithm in computing DAGs for a given dataset. All experiments are conducted on computers with 2.2 GHz Intel E7-4850V3 CPUs with a memory limit of 32 GB, and each score-and-search experiment is limited to a total of 24 hours.

Datasets: We test the performance of our learned networks on synthetic networks generated using the software Tetrad ¹. In addition, we show results on real-world Bayesian networks from the Bayesian Network Repository ² for large and very large network sizes. We also show results on 100 node networks in the DREAM 4 challenge benchmarks, which are gene regulation networks using the genenetweaver software ³. Gene regulation networks are a collection of biological regulators that interact with each other, and this interaction can be modelled as a directed graph.

Comparison to other algorithms: We ran our structure learning algorithm, BN-FS-MARS, on the datasets to learn the DAG and compared the results obtained with tabular (CPT-based) BIC and Greedy Equivalent Search (Chickering, 2002b), which is a score-based local search algorithm that searches over the space of equivalence classes of Bayesian-network structures. Tabular BIC results are computed using GOBNILP (Bartlett and Cussens, 2013) (see below) with standard pruning rules applied, and GES results are reported using the R package pcalg (v2.7-4) ⁴ with the score as BIC. For tabular BIC, parent set size was limited to 2, as scoring with larger parent set sizes lead to the vast majority of the experiments failing with an out of memory (OM) error. We refer to this variant of tabular BIC as BIC2. We also experimented with the neural network based DAG-GNN, however as observed in (Yu et al., 2019), we noted that the performance was weaker than that of GOBNILP, so in our experiments, we have reported the GOBNILP results. In addition, we experimented with WINASOBS from Scanagatta et al. (2015), however within our experimental setup it was outperformed by GES, and thus we have not reported the results in detail.

Performance evaluation metrics: We show the performance of BN-FS-MARS using the structural Hamming distance (SHD) and the F1-score. Given a ground truth network and a learned network, SHD measures the distance between the completed partially directed acyclic graphs (CPDAG) representations of the networks, where a CPDAG captures the equivalence class to which a network belongs (see Tsamardinos et al. (2006)). It is the sum of false negatives (FN), which are missing edges; false positives (FP) which are extra edges; and edges with wrong orientations (WD). We use the undirected graph (the skeleton of the DAG) for computing F1-scores. Given the number of edges present (SP) in both the skeleton of the ground truth network and the skeleton of the learned network, i.e. the correct edges, the false positive (FP) and the false negatives (FN), we can compute the F1-score with $\frac{SP}{SP + \frac{1}{2}(FP + FN)}$. The F1-score is between 0 and 1, with 1 being the best value.

5.1 Implementation Details

Computing local scores with MARS: MARS models for representing parent-child relationships were computed using the earth package (v5.3.1) ⁵ in R. This package builds regression models using the techniques proposed by (Friedman, 1991). The earth package limits every term in the model to have at most ten hinge functions in one product. However, for our experiments, we did not find this to be a limitation, as the models we learned did not exceed five hinge functions in a product.

To process discrete data for generating a MARS model, earth splits a categorical variable into l indicator columns of 1s and 0s, where l is the number of unique values the variable can take. For categorical response variables with l values, earth computes l models simultaneously. While the basis functions remain constant across these models, the coefficients can differ. Forward and backward phases are performed with the GCVs and RSSs summed across all l to minimize the sum of the GCV scores across all l models.

¹ <https://github.com/cmu-phil/Tetrad>

² www.bnlearn.com/bnrepository

³ <http://gnw.sourceforge.net>

⁴ <https://cran.r-project.org/web/packages/pcalg/>

⁵ <https://cran.r-project.org/web/packages/earth>

id	network	n	m	id	network	n	m	id	network	n	m
0	pathfinder	109	195	1	munin1	186	273	2	andes	223	338
3	diabetes	413	602	4	pigs	441	592	5	link	724	1125
6	munin2	1003	1244	7	munin4	1038	1306	8	munin3	1041	1388
9	munin	1041	1397								

Table 1: Benchmarks in bnlearn.

Feature selection with random forests: Random forests are an ensemble learning method that operate by constructing a set of decision trees at training time and outputting the class that is the average prediction of the individual trees for regression problems. Training random forests involves applying bagging to tree learners. Given a training set, bagging repeatedly selects a random sample of the training data and fits a decision tree to that sample. Random forests can rank variables by importance in a regression problem as follows: 1) Fit a random forest to the data set and record the out-of-bag error for each observation (averaged over the trees in the forest). 2) The importance of the variable V_i is then determined by perturbing the values of variable V_i in the training data and computing the out-of-bag error again on the perturbed data set. The importance score for variable V_i is then determined by averaging the difference of the out-of-bag error before and after the perturbation over all trees in the random forest and normalizing it by the standard deviation of the differences. Random forests were computed using H2O (v3.32.0.5) ⁶ in Python.

Finding DAGs of minimum score: To solve the optimisation problem of finding the DAGs with the minimum total GCV score, we use GOBNILP, which is a state-of-the-art integer programming method for finding an optimal Bayesian network given a list of candidate parent sets and local scores for each variable V_i , with the objective of solving BNSL to find a DAG G that minimises the total score, $\sigma(G)$. GOBNILP is an anytime algorithm, and it will return the best network at a given time.

5.2 Performance on Discrete Variables

To compare the performance of Algorithm 1 on discrete data, we generated synthetic networks of 100-900 nodes with Tetrad, with the maximum in-degree and out-degree of each node set as 3. We also used real world networks from bnlearn, with sizes between 100-1000+ nodes (see Table 1). From both these sets of networks, we randomly generated datasets with 1000 and 5000 samples, and scored candidate parent sets of all nodes with BN-FS-MARS using the top 5 features ranked by random forests in Step 3 of Algorithm 1. Finally, we computed the DAG structures using GOBNILP (Step 11), and compared the performance of BN-FS-MARS with BIC2 and GES. Due to the anytime nature of the algorithm, we are able to obtain the best network found at any given time, and we use this property to report results on the network found within the time limit of 24 hours for score-and-search.

Performance on F1-scores: Figure 2 reports the F1-score of the synthetic Tetrad networks learned by each of the three methods in column (a). We observe that BN-FS-MARS(BFM) has the highest F1-scores, followed by GES. Figure 3 reports the F1-score of the bnlearn networks learned by each of the three methods in column (a). BN-FS-MARS usually has the best F1-scores, especially as n increases, and the second best F1-scores are from GES. For both Tetrad and bnlearn, we observe that increasing sample size from $N = 1000$ to $N = 5000$ improves F1 scores of BN-

⁶ <https://docs.h2o.ai>

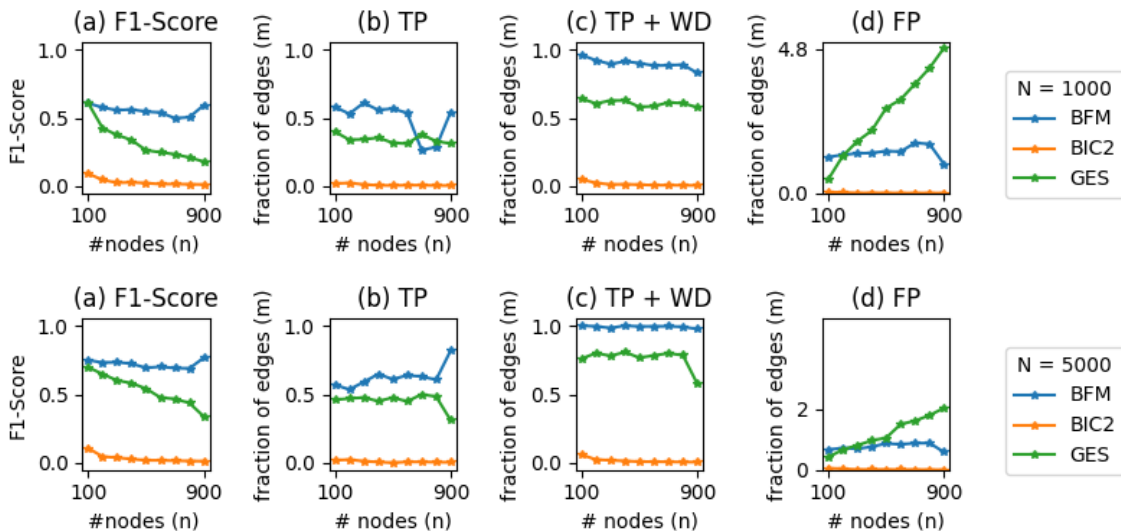


Figure 2: Comparison of structure learning approaches on various networks from Tetrad, where n is the number of variables and the number of edges in the ground truth network: our BN-FS-MARS (BFM), tabular BIC with maximum parent set size as 2 (BIC2), and GES with the BIC score (GES). For each ground truth network, we report the following for corresponding learned networks for sample sizes $N = 1000$ and $N = 5000$: (a) F1-score, (b) TP (true positives), (c) sum of TP and WD (wrong direction), and (d) FP (false positives).

FS-MARS and GES, and BN-FS-MARS continues to have the highest scores. However, adding more data does not improve the F1-scores for BIC2.

Performance on SHD: Figure 2 shows results for Tetrad on networks learned by each of the three methods in columns (b)-(d), and the numbers are normalised by the total number of edges m of the corresponding dataset. We see that BN-FS-MARS (BFM) almost always has the highest number of TPs as well as the highest sums of TP and WD. GES reports the second highest numbers, and BIC reports the lowest numbers, which are close to 0. This is in line with what we observed in F1-scores, showing that networks learned by BN-FS-MARS have a higher portion of edges present in the ground truth network, though some edges may not have the correct orientation. Note that FNs would correspond to the complement of the TP+WD plot, as they would show the number of edges that were not discovered at all.

Figure 2 (d) shows numbers for FP. GES almost always has the highest number of FPs, indicating a tendency to learn extra edges not present in the ground truth. BN-FS-MARS has the second highest FP, but these are usually significantly lower than those for GES. We observe that increasing sample size from $N = 1000$ to $N = 5000$ reduces the false positives significantly for both these methods. BIC2 has close to zero FPs. However, we can see that BIC2 had also failed to learn most of the edges in the network.

Figure 3 shows SHD results for bnlearn networks. It reports the SHD of the networks learned by each of the three methods in columns (b)-(d), and the numbers are normalised by the total number of edges m , of the corresponding dataset. Here we see that GES almost always has the highest number of TPs as well as the highest sums of TP and WD. BN-FS-MARS reports the second highest

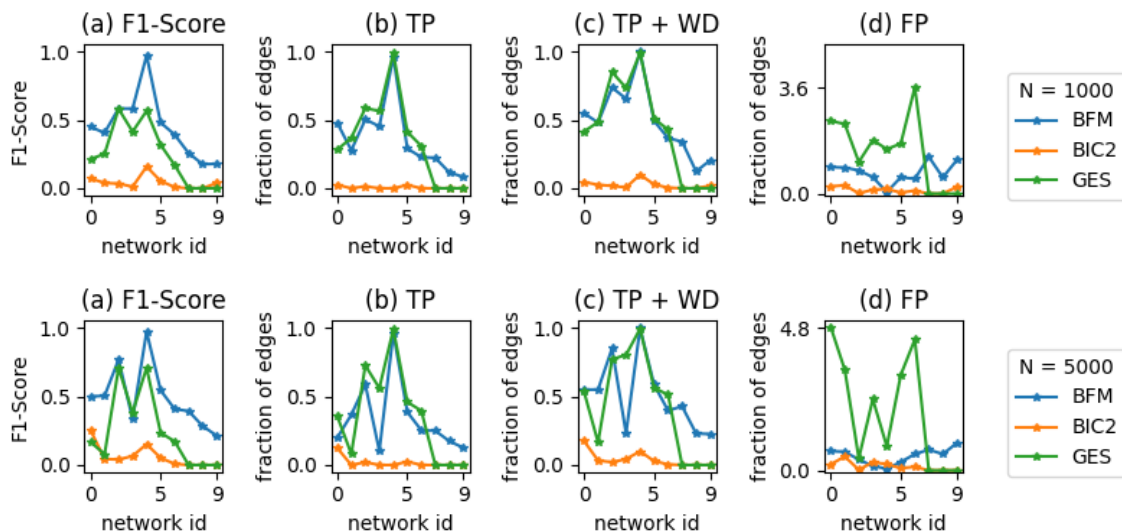


Figure 3: Comparison of structure learning approaches on various networks from bnlearn: our BN-FS-MARS (BFM), tabular BIC with maximum parent set size as 2 (BIC2), and GES with the BIC score (GES). For each ground truth network, we report the following for corresponding learned networks for sample sizes $N = 1000$ and $N = 5000$: (a) F1-score, (b) TP (true positives), (c) sum of TP and WD (wrong direction), and (d) FP (false positives).

numbers, which are close to the GES numbers, and BIC2 reports the lowest numbers, which are often close to 0.

For both Tetrad and bnlearn, we observe that increasing sample size from $N = 1000$ to $N = 5000$ usually improves TP and WD of networks learned using BN-FS-MARS and GES, and BN-FS-MARS is able to recover almost all of the edges in one of the instances (pigs). However, adding more data does not improve results for BIC2. For bnlearn benchmarks exceeding 1003 nodes, i.e. some of the munin sub-networks GES and BIC2 did not converge at the end of the time limit. GES stopped with an out of time error, and BIC2 ran out of the memory.

Figure 3 (d) shows numbers for FP for bnlearn. As with Tetrad networks, GES almost always has the highest number of FPs, indicating a tendency to learn extra edges not present in the ground truth. BN-FS-MARS has the second highest FP, but these are usually significantly lower than those for GES. We observe that increasing sample size from $N = 1000$ to $N = 5000$ sometimes leads to an increase in FPs for GES, but almost always a reduction for BN-FS-MARS. BIC2 has the close to zero FPs. However, as with Tetrad, BIC2 also failed to recover most ground truth edges.

We note that for one of the benchmarks (diabetes), BN-FS-MARS has decreased F1-scores and TPs with $N = 5000$. A limitation of using the MARS method on discrete values is that it attempts to optimise the set of basis functions for all models for the child variable. If the child’s domain is of large size, the results for the overall aggregated model will suffer in performance. Random forests are also known to perform poorly for such variables. The benchmark diabetes has several variables with domains of size 20+, leading to the observed degradation of performance of BN-FS-MARS. The more erratic nature of the plots for bnlearn is also because of the very different ranges of variable domain sizes across the benchmarks.

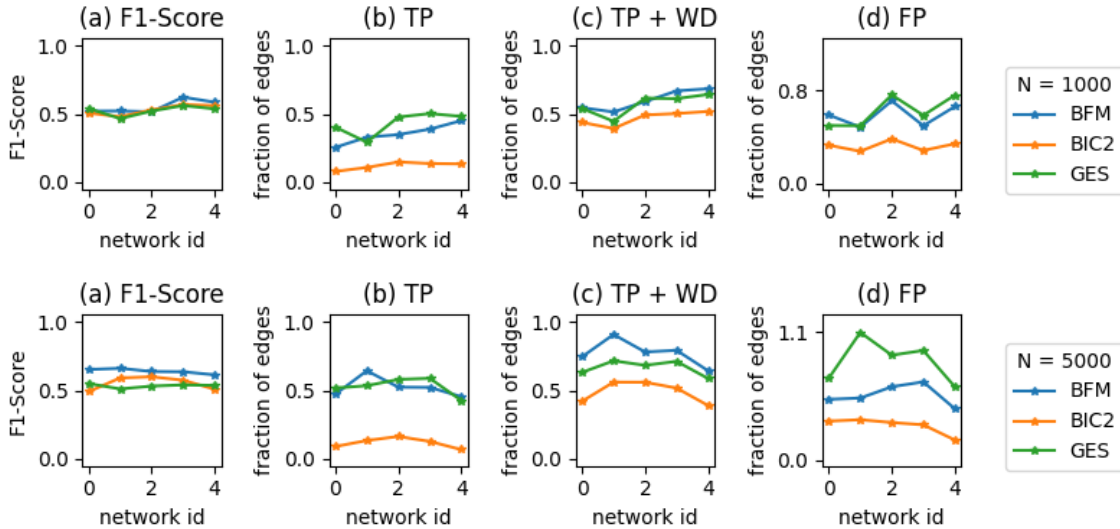


Figure 4: Comparison of structure learning approaches on various $n = 100$ networks from Dream4: our BN-FS-MARS (BFM), tabular BIC with maximum parent set size as 2 (BIC2), and GES with the BIC score (GES). For each ground truth network, we report the following for corresponding learned networks for sample sizes $N = 1000$ and $N = 5000$: (a) F1-score, (b) TP (true positives), (c) sum of TP and WD (wrong direction), and (d) FP (false positives).

5.3 Performance on Continuous Variables

To test the performance of Algorithm 1 on continuous data, we used the 5 Dream4 networks, each of which have 100 nodes and between 176 – 249 edges. We generated datasets using simulated steady-state measurements with the *genenetworker* software, with 1000 and 5000 samples. We scored candidate parent sets of all nodes with BN-FS-MARS with random forests as the feature selection method. For scoring, we used the top 5 scoring candidate parent sets in Step 3 of Algorithm 1. Once the scoring was complete, we computed the DAG structures using GOBNILP (Step 11). Due to the anytime nature of the algorithm, we are able to obtain the best network found at any given time, and we use this property to report results on the network found within the time limit of 24 hours for score-and-search. We compare the performance of BN-FS-MARS with tabular BIC (using *pyGOBNILP* because of the ability of continuous data processing) and BIC based GES.

Performance on F1-scores: Figure 4 reports the F1-score of networks learned by each of the three methods in column (a). For sample size $N = 1000$, the methods have similar scores. Increasing sample size to $N = 5000$ improves F1-scores of all methods slightly, and BN-FS-MARS shows the best scores.

Performance on SHD: Figure 4 shows SHD results for the networks learned by each of the three methods in columns (b)-(d), and the numbers are normalised by the total number of edges m of the corresponding dataset. GES almost always has the highest number of TPs, but it is close to BN-FS-MARS in sums of TP and WD. BIC2 reports the lowest numbers for both of these cases, but unlike with the discrete benchmarks, the numbers are not close to zero. For both Tetrad and *bnlearn* increasing sample size to $N = 5000$ improves TP and WD of networks learned by all methods, and BN-FS-MARS has the highest numbers for TP+WD.

The final column (d) shows numbers for FP. As with Tetrad and bnlearn networks, GES almost always has the highest number of FPs, indicating a tendency to learn extra edges not present in the ground truth. BN-FS-MARS has the second highest FP. We observe that increasing sample size to $N = 5000$ sometimes leads to an increase in FPs for GES, but always a reduction for BN-FS-MARS. BIC2 has the lowest number of FPs. A change in the behaviour of these three methods can be because of the smaller network size ($n = 100$). With a much smaller search space, BIC2 gives competitive results for F1-scores, in part because of its ability to learn fewer FPs. However, the strength of BN-FS-MARS comes from its feature selection mechanism, which is harder to solve for continuous data (Guyon and Elisseeff, 2003).

5.4 Model Complexity

One of the strengths of BN-FS-MARS is its ability to identify large parent sets. Consider a variable in the benchmark pathfinder with a domain of size 4. It has five parents of domains sizes between 2 – 3. Tabular BIC would need almost 300 parameters to represent it in CPT form, and it would give it a large penalty, making it difficult to be able to detect it with score-and-search. BN-FS-MARS assigns it a model with 32 terms, with none of the basis functions having a product of more than two terms. The lowered penalty makes such a parent set more likely to be selected, especially if it has a good score. In our experiments we found that modeling parent sets with BN-FS-MARS enabled us to score very large parent sets easily.

6. Conclusion

In this paper, we propose a novel approach to score-and-search for learning Bayesian networks with MARS relations as possible representations for CPDs. This algorithm lets us model non-linear relationships between nodes with low complexity, enabling us to learn large parent set sizes, and it does not place any constraints on the global network structure. Our approach can effectively prune most candidate parent sets of a variable by leveraging variable importance results from feature learning algorithms. We show empirically that this algorithm can solve both discrete and continuous instances with a very large number of nodes.

BIBLIOGRAPHY

- M. Bartlett and J. Cussens. Advances in Bayesian network learning using integer programming. In *Proc. of UAI*, pages 182–191, 2013.
- C. Boutilier, N. Friedman, M. Goldszmidt, and D. Koller. Context-specific independence in Bayesian networks. In *Proc. of UAI*, pages 115–123, 1996.
- L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- W. Buntine. Theory refinement on Bayesian networks. In *Proc. of UAI*, pages 52–60, 1991.
- D. M. Chickering. Learning equivalence classes of Bayesian network structures. *J. Mach. Learn. Res.*, 2:445–498, 2002a.
- D. M. Chickering. Optimal structure identification with greedy search. *J. Mach. Learn. Res.*, 3: 507–554, 2002b.

- D. M. Chickering, D. Heckerman, and C. Meek. A Bayesian approach to learning Bayesian networks with local structure. In *Proc. of UAI*, pages 80–89, 1997.
- A. Darwiche. *Modeling and Reasoning with Bayesian Networks*. Cambridge University Press, 2009.
- J. H. Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, pages 1–67, 1991.
- N. Friedman and M. Goldszmidt. Learning Bayesian networks with local structure. In *Learning in Graphical Models*, pages 421–459. Springer, 1998.
- I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *J. Mach. Learn. Res.*, 3:1157–1182, 2003.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2009.
- D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009.
- J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- M. Scanagatta, C. P. de Campos, G. Corani, and M. Zaffalon. Learning Bayesian networks with thousands of variables. In *Proc. of NeurIPS*, pages 1864–1872, 2015.
- G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, pages 461–464, 1978.
- C. Sharma, Z. A. Liao, J. Cussens, and P. van Beek. A score-and-search approach to learning Bayesian networks with noisy-or relations. In *Proc. of PGM*, pages 413–424, 2020.
- T. Talvitie, R. Eggeling, and M. Koivisto. Learning Bayesian networks with local structure, mixed variables, and exact algorithms. *International J. of Approximate Reasoning*, 2019.
- I. Tsamardinos, L. E. Brown, and C. F. Aliferis. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning*, 65(1):31–78, 2006.
- P. van Beek and H.-F. Hoffmann. Machine learning of Bayesian networks using constraint programming. In *Proc. of CP*, pages 428–444, 2015.
- Y. Xiang and N. Jia. Modeling causal reinforcement and undermining for efficient CPT elicitation. *IEEE Transactions on Knowledge and Data Engineering*, 19(12):1708–1718, 2007.
- Y. Yu, J. Chen, T. Gao, and M. Yu. DAG-GNN: DAG structure learning with graph neural networks. In *International Conference on Machine Learning*, pages 7154–7163. PMLR, 2019.
- N. L. Zhang and D. Poole. Exploiting causal independence in Bayesian network inference. *J. of AI Research*, 5:301–328, 1996.
- X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing. Dags with no tears: Continuous optimization for structure learning. *Proc. of NeurIPS*, 31, 2018.