

# Structural Learning of Probabilistic Graphical Models of Cumulative Phenomena

Daniele Ramazzotti<sup>1</sup>, Marco S. Nobile<sup>2</sup>, Marco Antoniotti<sup>2</sup>, and Alex Graudenzi<sup>2</sup>

<sup>1</sup> Dept. of Pathology, Stanford University, CA, USA  
daniele.ramazzotti@stanford.edu

<sup>2</sup> Dept. of Informatics, Systems and Communication, Università degli Studi di Milano-Bicocca, MI, Italy  
{marco.nobile,marco.antoniotti,alex.graudenzi}@unimib.it

**Abstract.** One of the critical issues when adopting Bayesian networks (BNs) to model dependencies among random variables is to “learn” their structure. This is a well-known *NP*-hard problem in its most general and classical formulation, which is furthermore complicated by known pitfalls such as the issue of *I*-equivalence among different structures. In this work we restrict the investigation to a specific class of networks, i.e., those representing the dynamics of phenomena characterized by the *monotonic accumulation* of events. Such phenomena allow to set specific structural constraints based on Suppes’ theory of probabilistic causation and, accordingly, to define constrained BNs, named *Suppes-Bayes Causal Networks* (SBCNs). Within this framework, we study the structure learning of SBCNs via extensive simulations with various state-of-the-art search strategies, such as canonical local search techniques and Genetic Algorithms. This investigation is intended to be an extension and an in-depth clarification of our previous works on SBCN structure learning. Among the main results, we show that Suppes’ constraints do simplify the learning task, by reducing the solution search space and providing a temporal ordering on the variables, which simplifies the complications derived by *I*-equivalent structures. Finally, we report on tradeoffs among different optimization techniques that can be used to learn SBCNs.

## 1 Introduction

*Bayesian Networks* (BNs) are probabilistic graphical models representing the relations of *conditional dependence* among random variables, encoded in *directed acyclic graphs* (DAGs) [21]. In the last decades, BNs have been effectively applied in several different fields and disciplines, such as (but not limited to) diagnostics and predictive analytics [21].

One of the most challenging task with BNs is that of *learning* their structure from data. Two main approaches are commonly used to tackle this problem.

1. *Constraint-based* techniques: mainly due to the works by Judea Pearl [26] and others, these approaches aim at discovering the relations of conditional independence from the data, using them as constraints to learn the network.
2. *Score-based* techniques: in this case the problem of learning the structure of a BN is defined as an *optimization* task (specifically, *maximization*) where the search space of the valid solutions (i.e., all the possible DAGs) is evaluated via *scores* based on a *likelihood function* [21].

Regardless of the approach, the main difficulty in this optimization problem is the huge number of valid solutions in the search space, namely, all the possible DAGs, which makes this task a

known *NP*-hard one in its most general instance, and even when constraining each node to have at most two parents [9,10]. Therefore, all state-of-the-art techniques solve this task by means of meta-heuristics [21,23,35].

Moreover, the inference is further complicated by the well-known issue of *I-equivalence*: BNs with even very different structures can encode the same set of conditional independence properties[21]. Thus, any algorithm for structural learning can converge to a set of equivalent structures rather than to the correct one, given that the inference itself is performed by learning the statistical relations among the variables emerging from their induced distributions rather than the structure itself [21].

In this paper, we investigate the application of BNs for the characterization of a specific class of dynamical phenomena, i.e., those driven by the *monotonic accumulation of events*. In particular, the process being modeled/observed must imply:

1. a *temporal ordering* among its events (i.e., the nodes in the BN), and
2. a monotonic accumulation over time, which probabilistically entails that the occurrence of an earlier event must be *positively correlated* to the subsequent occurrence of its successors, leading to a *significant temporal pattern* [29].

An example can be found in the dynamics of *cascading failures*, that is a failure in a system of interconnected parts where the failure of a part can trigger the failure of successive parts. These phenomenon can happen in different contexts, such as power transmission, computer networking, finance and biological systems. In these scenarios, different configurations may lead to failure, but some of them are more likely than others and, hence, can be modeled probabilistically [5].

The two particular conditions mentioned above can be very well modelled by the notion of *probabilistic causation* by Patrick Suppes [34,19], and allow us to define a set of *structural constraints* to the BNs to be inferred, which, accordingly, have been dubbed as *Suppes-Bayes Causal Networks* (SBCNs) in previous works [4,29]. SBCNs have been already applied in a number of different fields, ranging from cancer progression inference [24,28,7] to social discrimination discovery [4] and stress testing [15].

We specifically position our work within the aforementioned optimization-based framework for BN structure learning. The goal of this paper is to investigate how structure learning is influenced by different algorithmic choices, when representing cumulative dynamical phenomena. In particular, it is known that given a *temporal ordering* on the variables (i.e., a partially ordered set among the events, *poset* in the terminology of Bayesian networks) of a BN, finding the optimal solution that is consistent with the ordering can be accomplished in time  $O(n^k)$ , where  $n$  is the number of variables and  $k$  the bounded in-degree of a node [6,12]. Thus, the search in the *space of orderings* can be performed way more efficiently than the search in the *space of structures*, as the search space is much smaller, the branching factor is lower and acyclicity checks are not necessary [35,30].

The determination of the right ordering in complex dynamical phenomena is generally a difficult task, which often requires considerable domain knowledge. However, the representation of cumulative phenomena via SBCNs allows to soften this hurdle, as Suppes' constraints dramatically reduce the search space of valid solutions, also providing a temporal ordering on the variables. This represents a serious theoretical advancement in structure learning of BNs for the modeling of cumulative phenomena, which we investigate in this work with a series of synthetic experiments.

In particular, in this paper we quantitatively assess the performance of learning the structure of a BN when:

- the temporal ordering among variables is given / not given, i.e., when Suppes' constraints are imposed / not imposed (in the former case we deal with SBCNs);

- different heuristic search strategies are adopted, i.e., *Hill Climbing* (HC), *Tabu Search* (TS), and *Genetic Algorithms* (GA);
- different regularization terms are used, i.e., *Bayesian Information Criterion* (BIC) and *Akaike information criterion* (AIC).

## 2 Methods

In this Section we present the foundations of our framework and, specifically, we define the main characteristics of the SBCNs and of some heuristic strategies for the likelihood fit. Without losing in generality, from now on, we consider a simplified formulation of the problem of learning the structure of BNs where all the variables depicted in the graph are Bernoulli random variables, i.e., their support is  $(0, 1)$ . All the conclusions derived in these settings can be also directly applied to the general case where the nodes in the BN describe geneal random variables [21].

More precisely, we consider as an input for our learning task a dataset  $D$  of  $n$  Bernoulli variables and  $m$  cross-sectional samples. We assume the value 1 to indicate that a given variable has been observed in the sample and 0 that the variable had not been observed.

### 2.1 Suppes-Bayes Causal Networks

In [34], Suppes introduced the notion of *prima facie causation*. A prima facie relation between a cause  $u$  and its effect  $v$  is verified when the following two conditions are true.

1. *Temporal Priority* (TP): any cause happens before its effect.
2. *Probability Raising* (PR): the presence of the cause raises the probability of observing its effect.

**Definition 1 (Probabilistic Causation, [34]).** *For any two events  $u$  and  $v$ , occurring respectively at times  $t_u$  and  $t_v$ , under the mild assumptions that  $0 < \Pr(u), \Pr(v) < 1$ , the event  $u$  is called a prima facie cause of  $v$  if it occurs before and raises the probability of  $v$ , i.e.,*

$$\begin{cases} \text{(TP)} & t_u < t_v \\ \text{(PR)} & \Pr(v | u) > \Pr(v | \neg u) \end{cases} \quad (1)$$

The notion of prima facie causality has known limitations in the context of the general theories of causality [19], however, this characterization seems to appropriate to model the dynamics of phenomena driven by the monotonic accumulation of events where a temporal order among them is implied and, thus, the occurrence of an early event positively correlates to the subsequent occurrence in time of a later one. Let us now refer again to systems where cascading failure may occur: some configurations of events, in a specific order, may be more likely to cause a failure than others. This condition leads to the emergence of an observable *temporal pattern* among the events captured by Suppes' definition of causality in terms of statistical relevance, i.e., statistical dependency.

Let us now consider a graphical representation of the aforementioned dynamics in terms of a BN  $G = (V, E)$ . Furthermore, let us consider a given node  $v_i \in V$  and let us name  $\pi(v_i)$  the set of all the nodes in  $V$  pointing to (and yet temporally preceding)  $v_i$ . Then, the joint probability distribution of the  $n = |V|$  variables induced by the BN can be written as:

$$\Pr(v_1, \dots, v_n) = \prod_{v_i \in V} \Pr(v_i | \pi(v_i)) \quad (2)$$

When building our model, we need to constrain the characteristics of the considered relations as depicted in the network (i.e., the arcs in the graph), in order to account for the cumulative process above mentioned, which, in turns, needs to be reflected in its induced probability distribution [29]. To this extent, we can define a class of BNs over Bernoulli random variables named *Monotonic Progression Networks* (MPNs) [29,22,14]. Intuitively, MPNs represent the progression of events monotonically<sup>3</sup> accumulating over time, where the conditions for any event to happen is described by a probabilistic version of the canonical boolean operators, i.e., conjunction ( $\wedge$ ), inclusive disjunction ( $\vee$ ), and exclusive disjunction ( $\oplus$ ).

MPNs can model accumulative phenomena in a probabilistic fashion, i.e., they are also modeling irregularities (noise) in the data as a small probability  $\varepsilon$  of not observing later events given their predecessors.

Given these premises, in [28] the authors describe an efficient algorithm (named CAPRI, see [28] for details) to learn the structure of constrained Bayesian networks which account for Suppes' criteria and which later on are dubbed *Suppes-Bayes Causal Networks* (SBCNs) in [4]. SBCNs are well suited to model cumulative phenomena as they may encode irregularities in a similar way to MPNs [29]. The efficient inference schema of [29] relies on the observation (see [35]) that a way for circumventing the intrinsic computational complexity of the task of learning the structure of a Bayesian Network is to postulate a pre-determined ordering among the nodes. Intuitively, CAPRI exploits Suppes' theory to first mine an ordering among the nodes, reducing the complexity of the problem, and then fits the network by means of likelihood maximization. In [29] it is also shown that a SBCN, learned using CAPRI, can also embed the notion of accumulation through time as defined in a MPN, and, specifically, conjunctive parent sets; nevertheless SBCNs can easily be generalized to represent all the canonical boolean operators (*Extended Suppes-Bayes Causal Networks*), notwithstanding an increase of the algorithmic complexity [29]. We refer the reader to [29] for further details and, following [4], we now formally define a SBCN.

**Definition 2 (Suppes-Bayes Causal Network).**

Given an input cross-sectional dataset  $D$  of  $n$  Bernoulli variables and  $m$  samples, the Suppes-Bayes Causal Network  $SBCN = (V, E)$  subsumed by  $D$  is a directed acyclic graph such that the following requirements hold:

1. [**Suppes' constraints**] for each arc  $(u \rightarrow v) \in E$  involving the selective advantage relation between nodes  $u, v \in V$ , under the mild assumptions that  $0 < \Pr(u), \Pr(v) < 1$ :

$$\Pr(u) > \Pr(v) \quad \text{and} \quad \Pr(v | u) > \Pr(v | \neg u). \quad (3)$$

2. [**Simplification**] let  $E'$  be the set of arcs satisfying the Suppes' constraints as before; among all the subsets of  $E'$ , the set of arcs  $E$  is the one whose corresponding graph maximizes the log-likelihood  $\mathcal{LL}$  of the data and the adopted regularization function  $R(f)$ :

$$E = \arg \max_{E \subseteq E', G=(V,E)} \mathcal{LL}(G, D) - R(f). \quad (4)$$

Before moving on, we once again notice that the efficient implementation of Suppes' constraints of CAPRI does not, in general, guarantee to converge to the monotonic progression networks as depicted before. To overcome this limitation, one could extend the Algorithm in order to learn, in addition to the network structure, also the logical relations involving any parent set, increasing

<sup>3</sup> The events accumulate over time and when later events occur, earlier events are observed as well.

the overall computational complexity. Once again, we refer the interested reader to the discussions provided in [28,29] and, without losing in generality, for the purpose of this work, we consider the efficient implementation of CAPRI presented in [28].

It is important to remark that the evaluation of Suppes' constraints might be extended to longer serial dependence relations, by assessing, for instance, the statistical dependency involving more than two events. We here decide to evaluate pairwise conditions to keep the overall computational complexity at a minimum. However, we leave the investigation of this issue to further development of the framework.

## 2.2 Optimization and Evolutionary Computation

The problem of the inference of SBCNs can be re-stated as an optimization problem, in which the goal is the maximization of a likelihood score. Regardless of the strategy used in the inference process, the huge size of the search space of valid solutions makes this problem very hard to solve. Moreover, as stated above, the *general* problem of learning the structure of a BN is *NP*-hard [10]<sup>4</sup>. Because of that, state-of-the-art techniques largely rely on heuristics [21], often based on stochastic global optimization methods like Genetic Algorithms (GAs) [23,29]. Methods for BN learning can roughly be subdivided into two categories: single individual or population-based meta-heuristics.

Hill Climbing (HC) and Tabu Search (TS) both belong to the first category. The former is a greedy approach for the structural learning of BNs, in which new edges are attached to the current putative solution as long as they increase the likelihood score and they do not introduce any cycles in the network. TS is a stochastic variant of HC able to escape local minima, in which solutions visited in the past are not repeated by means of a tabu list.

GAs [20], a global search methodology inspired by the mechanisms of natural selection, belong to the second category. GAs were shown to be effective for BN learning, both in the case of available and not available *a priori* knowledge about nodes' ordering [23,29]. In a GA, a population  $\mathfrak{P}$  of candidate solutions (named individuals) iteratively evolves, converging towards the global optimum of a given fitness function  $f$  that, in this context, corresponds to the score to be maximized. The population  $\mathfrak{P}$  is composed of  $Q$  randomly created individuals, usually represented as fixed-length strings over a finite alphabet. These strings encode putative solutions of the problem under investigation; in the case of BN learning, individuals represent linearized adjacency matrices of candidate BNs with  $K$  nodes, encoded as string of binary values whose length<sup>5</sup> is  $O(K^2)$ .

The individuals in  $\mathfrak{P}$  undergo an iterative process whereby three genetic operators—selection, crossover and mutation—are applied in sequence to simulate the evolutionary recombination process, which results in a new population of possibly improved solutions. During the selection process, individuals from  $\mathfrak{P}$  are chosen, using a fitness-dependent sampling procedure [3], and are inserted into a new temporary population  $\mathfrak{P}'$ . The crossover operator is then used to recombine the structures of two promising selected parents. Finally, the mutation operator replaces an arbitrary symbol of an offspring, with a probability  $\mathcal{P}_m$ , using a random symbol taken from the alphabet. In the case of BNs, the mutation consists in flipping a single bit of the individual with a certain probability. It is worth noting that in the case of ordered nodes both crossover and mutation are *closed* operators,

<sup>4</sup> We are aware of *special* formulations of the problem that are solvable in polynomial time. Their existence points to interesting questions regarding the “barrier” between *NP* problems and polynomial ones; however, these are questions beyond the scope of the present paper.

<sup>5</sup> Since BNs are DAGs, the representation can be reduced by not encoding the elements on the diagonal, which are always equal to zero. In such case, the strings representing the individual have length  $K \times K - K$ .

because the resulting offsprings always encode valid DAGs. To the aim of ensuring a consistent population of individuals throughout the generations, in the case of unordered nodes the two operators are followed by a correction procedure, in which the candidate BN is analyzed to identify the presence of invalid cycles. For further information about our implementation of GAs for the inference of BNs, including the correction phase, we refer the interested reader to [29].

### 3 Results

We now discuss the results of a large number of experiments we conducted on simulated data, with the aim of assessing the performance of the state-of-the-art score-based techniques for the BN structure inference, and comparing the performance of these methods with the learning scheme defined in CAPRI.

Our main objective is to investigate how the performance is affected by different algorithmic choices at the different steps of the learning process.

*Data Generation.* All simulations are performed with the following generative models. We consider 6 different topological structures.

1. *Trees*: one predecessor at most for any node, one unique root (i.e., a node with no parents).
2. *Forests*: likewise, more than one possible root.
3. *Conjunctive DAGs with single root*: 3 predecessors at most for each node, all the confluences are ruled by logical conjunctions, one unique root.
4. *Conjunctive DAGs with multiple roots*: likewise, possible multiple roots.
5. *Disjunctive DAGs with single root*: 3 predecessors at most for each node, all the confluences are ruled by logical disjunctions, one unique root.
6. *Disjunctive DAGs with multiple roots*: likewise, possible multiple roots.

We constrain the induced distribution of each generative structure by implying a cumulative model for either conjunctions or disjunctions, i.e., any child node cannot occur if its parent set is not activated as described for the MPN in the Method Section 2. For each of these configurations, we generate 100 random structures. Furthermore, we simulate a model of noise in terms of random observations (i.e., false positives and false negatives) included in the generated datasets with different rates.

These data generation configurations are chosen to reflect: (a) different structural complexities of the models in terms of number of parameters, i.e., arcs, to be learned, (b) different types of induced distributions suitable to model cumulative phenomena as defined by the MPNs (see Section 2.1), i.e., conjunction ( $\wedge$ ) or inclusive disjunction ( $\vee$ )<sup>6</sup> and, (c) situations of reduced sample sizes and noisy data.

We here provide an example of data generation. Let now  $n$  be the number of nodes we want to include in the network and let us set  $p_{\min} = 0.05$  and  $p_{\max} = 0.95$  as the minimum and maximum probabilities of any node. A *directed acyclic graph without disconnected components* (i.e., an instance of types (3) and (5) topologies) with maximum depth  $\log n$  and where each node has at most  $w^* = 3$  parents is generated.

<sup>6</sup> Here we stick with the efficient search scheme of CAPRI and, for this reason, we do not consider exclusive disjunction ( $\oplus$ ) parent sets

**Algorithm 1:** Data generation: single source directed acyclic graphs

**Input:**  $n$ , the number of nodes of the graph,  $p_{\min} = 0.05$  and  $p_{\max} = 0.95$  be the minimum and maximum probabilities of any node and  $w^* = 3$  the maximum incoming edges per node.

**Result:** a randomly generated single source directed acyclic graph.

- 1 Pick an event  $r \in G$  as the root of the directed acyclic graph;
- 2 Assign to each node  $u \neq r$  an integer in the interval  $[2, \lceil \log n \rceil]$  representing its depth in the graph (1 is reserved for  $r$ ), ensuring that each level has at least one node;
- 3 **forall** nodes  $u \neq r$  **do**
- 4     Let  $l$  be the level assigned to the node;
- 5     Pick  $|\Pr(u)|$  uniformly over  $(0, w^*]$ , and accordingly define the parents of  $u$  with events selected among those at which level  $l - 1$  was assigned;
- 6 **end**
- 7 Assign  $\Pr(r)$ , a random value in the interval  $[p_{\min}, p_{\max}]$ ;
- 8 **forall** events  $u \neq r$  **do**
- 9     Let  $\alpha$  be a random value in the interval  $[p_{\min}, p_{\max}]$ ;
- 10    Let  $\pi(u)$  be the direct predecessor of  $u$ ;
- 11    Then assign:
 
$$\Pr(u) = \alpha \Pr(x \in \pi(u));$$
- 12 **end**
- 13 **return** *The generated single source directed acyclic graph.*

*Performance Assessment.* In all these configurations, the performance is assessed in terms of:

- accuracy =  $\frac{(TP+TN)}{(TP+TN+FP+FN)}$ ;
- sensitivity =  $\frac{TP}{(TP+FN)}$ ;
- specificity =  $\frac{TN}{(FP+TN)}$ ;

with  $TP$  and  $FP$  being the true and false positives (we mark as positive any arc that is present in the network) and  $TN$  and  $FN$  being the true and false negatives (we mark negative any arc that is not present in the network) with respect to the generative model. All these measures are values in  $[0, 1]$  with results close to 1 indicators of good performance.

*Implementation.* In all the following experiments, the adopted likelihood functions (i.e., the fitness evaluations) are implemented using the *bnlearn* package [33] written in the R language, while GA [20] the *inspyred* [16], *networkx* [31] and *numpy* [25] packages.

The framework for the inference of SBCNs is implemented in R and is available in the TRONCO suite for TRanslational ONCOlogy [13,2]. TRONCO is available under a GPL3 license at its webpage: <https://sites.google.com/site/troncopackage> or on Bioconductor.

*Algorithm Settings.* We test the performance of classical search strategies, such as Hill Climbing (HC) and Tabu Search (TS), and of more sophisticated algorithms such Genetic Algorithms (GA)<sup>7</sup>.

<sup>7</sup> Further experiments on multi-objective optimization techniques, such as *Non-dominated Sorting Genetic Algorithm* (NSGA-II), were performed, but are not shown here because of the worse overall performance, and of the higher computational cost, with respect to canonical GA.

For HC and TS, we generate data as described above with networks of 10 and 15 nodes (i.e., 0/1 Bernoulli random variables). We generated 10 independent datasets for each combination of the 4 sample levels (i.e., 50, 100, 150 and 200 samples) and the 9 noise levels (i.e., from 0% to 20% with step 2.5%) for a total of 4,320,000 independent datasets. The experiments were repeated either (i) including or (ii) not including the Suppes' constraints described in CAPRI [28], and independently using 5 distinct optimization scores and regularizers, namely standard (i) log-likelihood [21], (ii) AIC [1], (iii) BIC [32], (iv) BDE [18] and (v) K2 [11], leading to a final number of 86,400,000 different configurations.

Being more precise, given an input dataset of observations  $D$  and a graphical model  $G$ , we can define a function to evaluate the goodness of this structure given the data:

$$f(G, D) = \mathcal{LL}(D|G) - \mathcal{R}(G),$$

where  $\mathcal{LL}(\cdot)$  is the log-likelihood, while  $\mathcal{R}(\cdot)$  is a regularization term with the aim of limiting the complexity of  $G$ . The DAG induced by  $G$  in fact defines a probability distribution over its nodes, namely  $\{x_1, \dots, x_n\}$ :

$$\Pr(x_1, \dots, x_n) = \prod_{x_i=1}^n \Pr(x_i | \pi_i), \quad \Pr(x_i | \pi_i) = \theta_{x_i|\pi_i},$$

where  $\pi_i = \{x_j | x_j \rightarrow x_i \in G\}$  are  $x_i$ 's parents in the DAG, and  $\theta_{x_i|\pi(x_i)}$  is a density function. Then, the log-likelihood of the graph can be defined as:

$$\mathcal{LL}(D|G) = \log \Pr(D | G, \theta).$$

Then, the regularization term  $\mathcal{R}(G)$  introduces a penalty for the number of parameters in the model  $G$  also considering the size of the data. The above mentioned scores that we considered differ in this penalty, with AIC and BIC being Information-theoretic score and BDE and K2, Bayesian scores [8].

While a detailed description of these regularizers is beyond the scope of this paper, we critically discuss the different performances granted by each strategy for the inference of BNs.

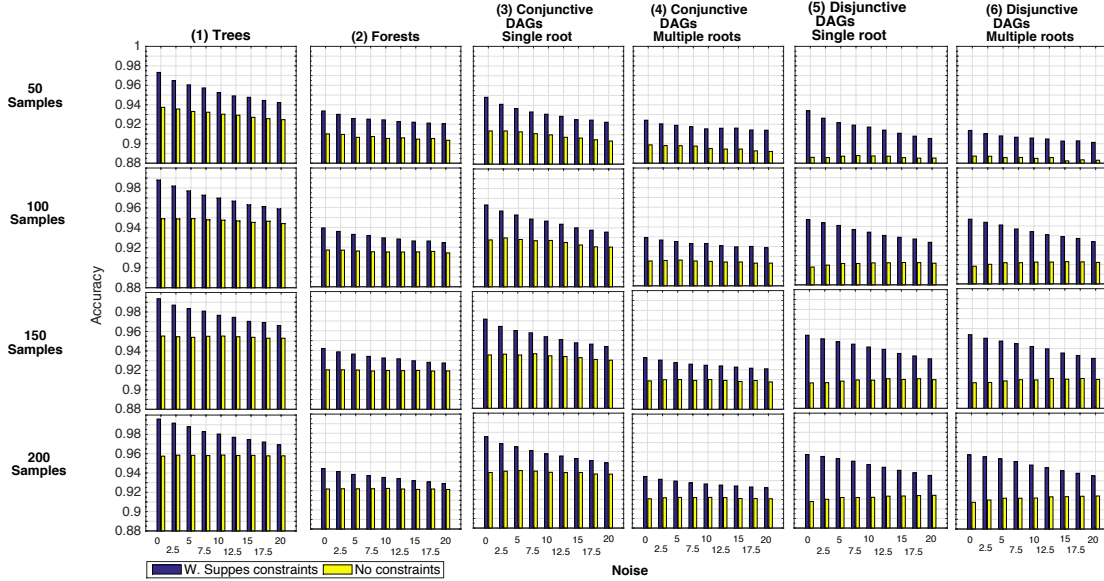
With respect to GA we used a restricted data generation settings, using networks of 15 nodes, datasets of 100 samples and 5 noise levels (from 0% to 20% with step 5%) for a total of 3,000 independent datasets. We tested the GA either (i) with or (ii) without Suppes' constraints, using BIC regularization term, leading to the final total of 6,000 different configurations. Finally, the GA was launched with a population size of 32 individuals, a mutation rate of  $p_m = 0.01$  and 100 generations.

We summarize the performance evaluation of the distinct techniques and settings in the next Subsections and in Figures 1, 2, 3 and 4.

## Performance Assessment

By looking at Figure 1, one can first appreciate the variation of accuracy with respect to a specific search strategy, i.e., HC with BIC, which is taken as an example of typical behavior. In brief, the overall performance worsens with respect to: (i) a larger number of nodes in the network, (ii) more complex generative structures, and (iii) smaller samples sizes / higher noise rates. Although such a trend is intuitively expected, given the larger number of parameters to be learned for more complex






**Fig. 1.** Performance in terms of accuracy for the 6 considered structures with 15 nodes, noise levels from 0% to 20% with step 2.5% and sample sizes of 50, 100, 150 and 200 samples. Here we use BIC as a regularization scheme and we consider HC as a search strategy both for the classical case and when Suppes’ priors are applied.

models, we here underline the role of statistical complications, such as the presence of *spurious correlations* [27] and the occurrence of *Simpson’s paradox* [17].

For instance, it is interesting to observe a typical decrease of the accuracy when we compare topologies with the same properties, but different number of roots (i.e., 1 root vs. multiple roots). In the former case, we expect, in fact, a lower number of arcs (i.e., dependencies) to be learned (on average) and, hence, we may attribute the decrease of the performance to the emergence of spurious correlations among independent nodes, such as the children of the different sources of the DAG. This is due to the fact that, when sample sizes are not infinite, it is very unlikely to observe perfect independence and, accordingly, the likelihood scores may lead to overfitting. The trends displayed in Figure 1 are shared by most of the analyzed search strategies.

*Role of the Regularization Term.* By looking at Figure 2 one can first notice that the accuracy with no regularization is dramatically lower than the other cases, as a consequence of the expected overfitting (in this case we compare the performance of HC on disjunctive DAGs with multiple roots, but the trend is maintained in the other cases). Conversely, all regularization terms ensure the inference of sparser models, by penalizing the number of retrieved arcs. BDE regularization term seems to be the only exception (see Figure 2), leading to unintuitive behaviors: in fact, while for all the other methods the performance decreases when higher level of noise are applied, for BDE the accuracy seems to improve with higher noise rates. This result might be explained by observing that given a topological structure, structural spurious correlations may arise between a given node and any of its undirected predecessors (i.e., one of the predecessors of its direct parents): with




**Fig. 2.** Performance in terms of accuracy for directed acyclic graphs with multiple sources and disjunctive parents (structure  $vi$ ) of 15 nodes, noise levels from 0% to 20% with step 2.5% and sample sizes of 50, 100, 150 and 200 samples. Here we consider HC as a search strategy both for the classical case and when Suppes’ priors are applied and we show the results for all five regularizers introduced in the text.

higher error rates, and, accordingly, more random samples in the datasets, all the correlations are reduced, hence leading to a lower impact of the regularization term. Given these considerations, one can hypothesize that the overall trend of BDE is due to a scarce penalization to the likelihood fit, favoring dense networks rather than sparse ones.

*Search Strategies.* No significant differences in the performance between the accuracy of HC, TS and GA are observed. However, one can observe a consistent improvement in sensitivity when using GA (see Figures 3 and 4). This suggests different inherent properties of the search schemes: while with HC and TB the regularization terms, rather than the search strategy, account for most of the inference performance, GAs are capable of returning denser networks with better hit rates. This is probably due to GA’s random mutations, which allow jumps into areas of the search space characterized by excellent fitness, which could not be reached by means of greedy approaches like HC.

*Suppes’ Structural Constraints.* Finally, the most important result, which can be observed across all the different experiments, is that the overall performance of all the considered search strategies is dramatically enhanced by the introduction of Suppes’ structural constraints. In particular, as one can see, e.g., in Figure 1, there is a constant improvement in the inference, up to 10%, when Suppes’ priors are used. Even though the accuracy of the inference is affected by the noise in the observations,




**Fig. 3.** Performance in terms of accuracy for directed acyclic graphs with multiple sources and disjunctive parents (structure  $vi$ ) of 15 nodes, noise levels from 0% to 20% with step 2.5% and sample sizes of 50, 100, 150 and 200 samples. Here we use BIC as a regularization scheme and we consider both HC and TB as search strategies for the classical case and when Suppes' priors are applied.

in fact, the results with Suppes' priors are consistently better than the inference with no constraints, with respect to all the considered inference settings and to all the performance measures. This is an extremely important result as it proves that the introduction of structural constraints based on Suppes' probabilistic causation indeed simplify the optimization task, by reducing the huge search space, when dealing with BNs describing cumulative phenomena.

## 4 Conclusion

In this paper we investigated the structure learning of Bayesian Networks aimed at modeling phenomena driven by the monotonic accumulation of events over time. To this end, we made use of a subclass of constrained Bayesian networks named Suppes-Bayes Causal Networks, which include structural constraints grounded in Suppes' theory of probabilistic causation.



**Fig. 4.** Performance of HC, TB and GA, in terms of sensitivity and specificity for forests (panels 1A, 1B), directed acyclic graphs with multiple sources and conjunctive parents (panels 2A, 2B) and directed acyclic graphs with multiple sources and disjunctive parents (panels 3A, 3B) (configurations (ii), (iv) and (vi)) of 15 nodes, noise levels from 0% to 20% with step 5% and sample sizes of 100 samples. Here we use BIC as a regularization scheme for HC and TB, and for all the algorithms we either consider (panels A) or not consider (panels B) Suppes' constraints (SC).

While the problem of learning the structure of a Bayesian Network is known to be intractable, such constraints allow to prune the search space of the possible solutions, leading to a tremendous reduction of the number of valid networks to be considered, hence taming the complexity of the problem in a remarkable way.

We here discussed the theoretical implications of the inference process at the different steps, also by comparing various state-of-the-art algorithmic approaches and regularization methods. We finally provided an in-depth study on realistically simulated data of the effect of each inference choice, thus providing some sound guidelines for the design of efficient algorithms for the inference of models of cumulative phenomena.

According to our results, none of the tested search strategies significantly outperforms the others in all the experimental settings, in terms of both sensitivity and specificity.

Yet, we could prove that Suppes' constraints consistently improve the inference accuracy, in all the considered scenarios and with all the inference schemes, hence positioning SBCNs as the new benchmark in the the efficient inference and representation of cumulative phenomena.

**Acknowledgments.** This work was supported in part by the ASTIL Program of Regione Lombardia, by the ELIXIR-ITA network, and by the SysBioNet project, a MIUR initiative for the Italian Roadmap of European Strategy Forum on Research Infrastructures (ESFRI). We would like to thank for the useful discussions our colleagues Giulio Caravagna of ICR, London, UK, Giancarlo Mauri of DISCo, Università degli Studi di Milano-Bicocca, Milan, Italy, and Bud Mishra of Courant Institute of Mathematical Sciences, New York University, NY, USA.

## References

1. Hirotugu Akaike. Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer, 1998.
2. M Antoniotti, G Caravagna, L De Sano, A Graudenzi, G Mauri, B Mishra, and D Ramazzotti. Design of the tronco bioconductor package for translational oncology. *THE R JOURNAL*, 8(2):39–59, 2016.
3. Thomas Back. Selective pressure in evolutionary algorithms: A characterization of selection mechanisms. In *Evolutionary Computation, 1994. IEEE World Congress on Computational Intelligence., Proceedings of the First IEEE Conference on*, pages 57–62. IEEE, 1994.
4. Francesco Bonchi, Sara Hajian, Bud Mishra, and Daniele Ramazzotti. Exposing the probabilistic causal structure of discrimination. *International Journal of Data Science and Analytics*, 3(1):1–21, 2017.
5. Sergey V Buldyrev, Roni Parshani, Gerald Paul, H Eugene Stanley, and Shlomo Havlin. Catastrophic cascade of failures in interdependent networks. *Nature*, 464(7291):1025–1028, 2010.
6. Wray Buntine. Theory refinement on Bayesian Networks. In *Proceedings of the Seventh conference on Uncertainty in Artificial Intelligence*, pages 52–60. Morgan Kaufmann Publishers Inc., 1991.
7. Giulio Caravagna, Alex Graudenzi, Daniele Ramazzotti, Rebeca Sanz-Pamplona, Luca De Sano, Giancarlo Mauri, Victor Moreno, Marco Antoniotti, and Bud Mishra. Algorithmic methods to infer the evolutionary trajectories in cancer progression. *Proceedings of the National Academy of Sciences*, 113(28):E4025–E4034, 2016.
8. Alexandra M Carvalho. Scoring functions for learning bayesian networks. *Inesc-id Tec. Rep.*, 2009.
9. David Maxwell Chickering. Learning bayesian networks is NP-complete. In *Learning from data*, pages 121–130. Springer, 1996.
10. David Maxwell Chickering, David Heckerman, and Christopher Meek. Large-sample learning of Bayesian Networks is NP-hard. *Journal of Machine Learning Research*, 5(Oct):1287–1330, 2004.
11. Gregory F Cooper and Edward Herskovits. A bayesian method for constructing bayesian belief networks from databases. In *Proceedings of the Seventh conference on Uncertainty in Artificial Intelligence*, pages 86–94. Morgan Kaufmann Publishers Inc., 1991.
12. Gregory F Cooper and Edward Herskovits. A bayesian method for the induction of probabilistic networks from data. *Machine learning*, 9(4):309–347, 1992.
13. Luca De Sano, Giulio Caravagna, Daniele Ramazzotti, Alex Graudenzi, Giancarlo Mauri, Bud Mishra, and Marco Antoniotti. Tronco: an r package for the inference of cancer progression models from heterogeneous genomic data. *Bioinformatics*, 32(12):1911–1913, 2016.
14. Hossein Shahrabi Farahani and Jens Lagergren. Learning oncogenetic networks by reducing to mixed integer linear programming. *PloS ONE*, 8(6):e65773, 2013.
15. Gelin Gao, Bud Mishra, and Daniele Ramazzotti. Efficient simulation of financial stress testing scenarios with suppes-bayes causal networks. *Procedia Computer Science*, 108:272–284, 2017.
16. Aaron Garrett. inspyred: Bio-inspired algorithms in python. *URL* <https://pypi.python.org/pypi/inspyred>, 2012.
17. IJ Good, Y Mittal, et al. The amalgamation and geometry of two-by-two contingency tables. *The Annals of Statistics*, 15(2):694–711, 1987.
18. David Heckerman, Dan Geiger, and David M Chickering. Learning bayesian networks: The combination of knowledge and statistical data. *Machine learning*, 20(3):197–243, 1995.

19. Christopher Hitchcock. Probabilistic causation. *Stanford encyclopedia of philosophy*, 2010.
20. John H Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press, 1975.
21. Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
22. Ilya Korsunsky, Daniele Ramazzotti, Giulio Caravagna, and Bud Mishra. Inference of cancer progression models with biological noise. *arXiv preprint arXiv:1408.6032*, 2014.
23. Pedro Larranaga, Mikel Poza, Yosu Yurramendi, Roberto H. Murga, and Cindy M. H. Kuijpers. Structure learning of bayesian networks by genetic algorithms. *IEEE transactions on pattern analysis and machine intelligence*, 18(9):912–926, 1996.
24. Loes Olde Loohuis, Giulio Caravagna, Alex Graudenzi, Daniele Ramazzotti, Giancarlo Mauri, Marco Antoniotti, and Bud Mishra. Inferring tree causal models of cancer progression with probability raising. *PloS one*, 9(10):e108358, 2014.
25. T Oliphant. A guide to numpy, vol. 1. *Spanish Fork: Trelgol Publishing*, 2006.
26. Judea Pearl. *Causality*. Cambridge university press, 2009.
27. Karl Pearson. Mathematical contributions to the theory of evolution.—on a form of spurious correlation which may arise when indices are used in the measurement of organs. *Proceedings of the royal society of london*, 60(359-367):489–498, 1896.
28. Daniele Ramazzotti, Giulio Caravagna, Loes Olde Loohuis, Alex Graudenzi, Ilya Korsunsky, Giancarlo Mauri, Marco Antoniotti, and Bud Mishra. Capri: efficient inference of cancer progression models from cross-sectional data. *Bioinformatics*, 31(18):3016–3026, 2015.
29. Daniele Ramazzotti, Alex Graudenzi, and Marco Antoniotti. Modeling cumulative biological phenomena with suppes-bayes causal networks. *arXiv preprint arXiv:1602.07857*, 2016.
30. Daniele Ramazzotti, Marco S. Nobile, Paolo Cazzaniga, Giancarlo Mauri, and Marco Antoniotti. Parallel implementation of efficient search schemes for the inference of cancer progression models. In *IEEE International Conference on Computational Intelligence in Bioinformatics and Computational Biology*. IEEE, 2016.
31. Daniel A Schult and P Swart. Exploring network structure, dynamics, and function using networkx. In *Proceedings of the 7th Python in Science Conferences (SciPy 2008)*, volume 2008, pages 11–16, 2008.
32. Gideon Schwarz et al. Estimating the dimension of a model. *The annals of statistics*, 6(2):461–464, 1978.
33. Marco Scutari. Learning Bayesian networks with the bnlearn R package. *arXiv preprint arXiv:0908.3817*, 2009.
34. Patrick Suppes. *A probabilistic theory of causality*. North-Holland Publishing Company Amsterdam, 1970.
35. Marc Teyssier and Daphne Koller. Ordering-based search: A simple and effective algorithm for learning bayesian networks. *arXiv preprint arXiv:1207.1429*, 2012.