



# New local move operators for Bayesian networks structure learning

Jimmy Vandel, Brigitte B. Mangin, Simon de Givry

## ► To cite this version:

Jimmy Vandel, Brigitte B. Mangin, Simon de Givry. New local move operators for Bayesian networks structure learning. The Sixth European Workshop on Probabilistic Graphical Models, Sep 2012, Granada, Spain. pp.8. hal-02744796

**HAL Id: hal-02744796**

**<https://hal.inrae.fr/hal-02744796>**

Submitted on 3 Jun 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# New Local Move Operators for Learning the Structure of Bayesian Networks

Jimmy Vandel and Brigitte Mangin and Simon de Givry  
UBIA, UR 875, INRA, F-31320 Castanet Tolosan, France  
{jvandel,mangin,degivry}@toulouse.inra.fr

**Abstract.** We propose new local move operators incorporated into a score-based stochastic greedy search algorithm to efficiently escape from local optima in the search space of directed acyclic graphs. We extend the classical set of arc addition, arc deletion, and arc reversal operators with a new operator replacing or *swapping* one parent to another for a given node, *i.e.* combining two elementary operations (arc addition and arc deletion) in one move. The old and new operators are further extended by doing more operations in a move in order to overcome the acyclicity constraint of Bayesian networks. These extra operations are temporally performed in the space of directed *cyclic* graphs. At the end acyclicity is restored and newly defined operators actually lead to a gain in graph score. Our experimental results on standard Bayesian networks and challenging gene regulatory networks show large BDeu score and recall value improvements compared to state-of-the-art structure learning algorithms when the sample size is small.

## 1 Introduction

Learning the structure of Bayesian networks from fully observed data is known to be an NP-hard problem [5] which has received a lot of attention from researchers during the last two decades [8]. Due to its difficulty, heuristic methods have been widely used to learn Bayesian network structures. Two main approaches have been studied: constraint-based methods and score-based methods [18]. Constraint-based methods aim at satisfying as much independence present in the data as possible using conditional independence tests. Unfortunately, these methods can be sensitive to failures in individual independence tests. Score-based methods use a scoring function  $f$  to score a network structure with respect to the data. They score the whole network at once, being therefore less sensitive to individual failures. Different Bayesian and non-Bayesian scoring metrics can be used such as the *Bayesian Information Criterion* (BIC) [26] or the *Bayesian Dirichlet criterion* (BDeu) [16]. Score-based methods explore the space of network structures to find the highest-scoring network. This space being superexponential in the number of variables, local search methods are used such as greedy ascent search (also called hill-climbing), tabu search, simulated-annealing, and other complex metaheuristics like ant colony optimization [7]. In spite of its simplicity, the (repeated randomized or *stochastic*) greedy search method reveals to be a competitive method compared to more complex algorithms [12]. Starting from an initial network structure, it performs a series of local moves until a local optimum is found. Each move selects and applies the best elementary operation(s) on the current network structure. The set of candidate neighboring structures

is called the *neighborhood* in the sequel. A classical neighborhood is composed of single arc additions, deletions, and reversals. Using *larger* neighborhoods efficiently allows to find better local optima, and thus better network structures.

[20] proposed an optimal reinsertion of a target node by removing all its edges and reinserting it optimally. However this approach is limited to small problems only. [17] used a restricted form of look ahead called LAGD, combining several operations in a single move. In this paper, we follow a similar approach by focusing our local operations on a target node and combining several operations guided by the global acyclicity constraint of Bayesian networks. By doing so, we are able to exploit large neighborhoods efficiently. Other approaches use a compact representation of a set of network structures. They explore either the search space of variable orderings (an optimal structure compatible with the order being easier to find) [6, 28, 25, 2], or the search space of Markov-equivalent network classes like *Greedy Equivalence Search* (GES) [3, 21, 7].

In Section 2, we give Bayesian network background. Next, we define a specific stochastic greedy search algorithm and introduce the new local move operators in Section 3. We report experimental results in Section 4 and conclude.

## 2 Bayesian network structure learning

A Bayesian network [18] denoted by  $B = (G, \mathbf{P}_G)$  is composed of a directed acyclic graph (DAG)  $G = (\mathbf{X}, \mathbf{E})$  with nodes representing  $p$  random discrete variables  $\mathbf{X} = \{X_1, \dots, X_p\}$ , linked by a set of directed edges or *arcs*  $\mathbf{E}^1$ , and a set of conditional probability distributions  $\mathbf{P}_G = \{P_1, \dots, P_p\}$  defined by the topology of the graph:  $P_i = \mathbb{P}(X_i | Pa(X_i))$  where  $Pa(X_i) = \{X_j \in \mathbf{X} \mid (X_j \rightarrow X_i) \in \mathbf{E}\}$  is the set of parent nodes of  $X_i$  in  $G$ . A Bayesian network  $B$  represents a joint probability distribution on  $\mathbf{X}$  such that:  $\mathbb{P}(\mathbf{X}) = \prod_{i=1}^p \mathbb{P}(X_i | Pa(X_i))$ .

Conditional probability distributions  $\mathbf{P}_G$  are determined by a set of parameters. Given the network structure  $G$ , and the fully observed data  $D$ , parameters can be estimated by simple counting, following the maximum likelihood principle.

Learning the structure of a Bayesian network consists in finding a DAG  $G$  maximizing the posterior distribution  $\mathbb{P}(G|D)$ . We have  $\mathbb{P}(G|D) \propto \mathbb{P}(D|G)\mathbb{P}(G)$  since  $\mathbb{P}(D)$  is independent of  $G$ . Under specific assumptions, the *marginal loglikelihood*  $\log(\mathbb{P}(D|G))$  can be expressed as a consistent decomposable scoring function  $f$ , such as the BIC and BDeu criteria [3] :

---

<sup>1</sup> In the paper, we use  $G = \mathbf{E}$  when the set of nodes is implicit.

$$f(D, G) = \sum_{i=1}^p f_{X_i}(D, G) = \sum_{i=1}^p f_{X_i}(D, Pa(X_i)) \quad (1)$$

As  $f$  is consistent, maximizing  $f$  when the sample size tends to infinity leads to select the true structure or one among the Markov-equivalent structures set. A set of Bayesian networks are Markov-equivalent if they imply exactly the same set or *map* of independence constraints among variables<sup>2</sup>. Next, we describe a novel greedy search method maximizing  $f$  in the space of DAGs.

### 3 Stochastic Greedy Search

We define the Stochastic Greedy Search (SGS) algorithm<sup>3</sup> for structural learning of Bayesian networks in Algorithm 1. It collects the best DAG found by  $r$  randomized hill climbing algorithms. Stochasticity comes from two random draws. The first one, common in the structure learning community, is due to `InitGraph` that returns a random DAG used by the inner greedy search loop. The second is more original. It is the random DAG drawn among the best neighbors in the neighborhood of the current DAG  $G$  (`SelectRandom` at line 1 of Algorithm 1). The neighborhood of  $G$ , returned by `Neighbors`, is composed of the usual elementary operations on DAGs: arc addition (ADD), arc deletion (DELETE), and arc reversal (REVERSE). This classical neighborhood is denoted  $\mathcal{N}^{ADR}$  in the sequel. Only feasible operations are considered, which do not create cycles. In the next subsections, we are going to extend this set of operations. We empirically observe that the first random draw may be counter-productive and starting from an empty graph is often better than from a random DAG, as also observed in [19].

---

#### Algorithm 1: Stochastic Greedy Search algorithm.

---

**Input** : An iid sample  $D$ , a scoring function  $f$ , number of repeats of different GS  $r$   
**Output**: A directed acyclic graph  
 $G^* \leftarrow \emptyset$  /\* Best DAG initialized with the empty graph \*/;  
 $s^* \leftarrow f(D, G^*)$  /\* score of the empty graph \*/;  
/\* Repeat  $r$  randomized greedy searches \*/;  
**for**  $i \leftarrow 1$  **to**  $r$  **do**  
   $G \leftarrow \text{InitGraph}()$  /\* Choose an initial DAG \*/;  
   $s \leftarrow f(D, G)$ ;  
  **repeat**  
     $\text{improvement} \leftarrow \text{false}$  ;  
     $s^{\max} \leftarrow \max_{G' \in \text{Neighbors}(G)} f(D, G')$ ;  
    **if**  $s^{\max} > s$  **then**  
      /\* Select at random among the best neighbors \*/;  
       $G^{\max} \leftarrow \{G' \in \text{Neighbors}(G) \mid f(D, G') = s^{\max}\}$  ;  
       $G \leftarrow \text{SelectRandom}(G^{\max})$  ;  
       $s \leftarrow s^{\max}$  ;  
       $\text{improvement} \leftarrow \text{true}$  ;  
  **until**  $\neg \text{improvement}$  ;  
  /\* Keep the best DAG of  $r$  greedy searches \*/;  
  **if**  $s > s^*$  **then**  
     $G^* \leftarrow G$  ;  
     $s^* \leftarrow s$  ;  
**return**  $G^*$  ;

---

**Proposition 1.** *Let  $D$  be a dataset of  $n$  records that are identically and independently sampled from some distribution  $\mathbb{P}(\cdot)$ . Let  $f$  be a*

<sup>2</sup> BIC/BDeu give the same score for Markov-equivalent DAGs.

<sup>3</sup> Don't mistake for the SGS algorithm in [27] which is a constraint-based algorithm, but ours is score-based.

locally consistent scoring function. The inner loop of the SGS algorithm returns a minimal independence map of  $\mathbb{P}(\cdot)$  as the sample size  $n$  grows large.

The local consistency of  $f$  ensures that adding any arc that eliminates an independence constraint that does not hold in the generative distribution  $\mathbb{P}(\cdot)$  increases the score. Conversely, deleting any arc that results in a new independence constraint that holds in  $\mathbb{P}(\cdot)$  also increases the score. Hence, starting from any DAG, by selecting strictly improving moves, the algorithm terminates (because of a finite number of DAGs) and returns an independence map of  $\mathbb{P}(\cdot)$  (*I-map* of  $\mathbb{P}(\cdot)$ ): every independence implied by the resulting DAG is verified in  $\mathbb{P}(\cdot)$ , in the worst case it can be a complete graph) which is minimal thanks to arc deletion operations and local consistency.

The main interest of our randomization approach is to simulate a search in the space of score-equivalent networks. Each greedy search moves from a randomly-selected DAG instance of a Markov-equivalence class  $\mathcal{E}(G)$  to another randomly-selected DAG of an *adjacent*<sup>4</sup> Markov-equivalence class  $\mathcal{E}(G')$  thanks to our `SelectRandom` function. It results in a stronger property:

**Proposition 2.** *Let  $D$  be a dataset of  $n$  iid fully observed samples of some faithful distribution  $\mathbb{P}(\cdot)$ . Let  $f$  be a locally consistent scoring function. SGS returns a perfect map of  $\mathbb{P}(\cdot)$  as both the sample size  $n$  and the number of restarts  $r$  grow large.*

Recall that a faithful distribution admits a unique perfect map corresponding to the optimal structure. Compared to the GES algorithm [3], which offers the same optimality guarantee within a two-phase greedy search, SGS chooses the orientation of some *compelled arcs*<sup>5</sup> of the *true* DAG at random, whereas GES waits while no *v-structures* impose orientation constraints. See an example in Figure 1.

Notice that neither GES nor SGS find an optimal structure in polynomial time in the worst case, even when using a constant time consistent scoring function and a faithful distribution<sup>6</sup>. In general, without the faithfulness assumption, learning the optimal structure is NP-hard even when the sample size is large and when each node has at most  $k$  parents, for all  $k \geq 3$  [4].

We observe in the experiments that a small number of restarts  $r$  allows to find DAGs with better scores than GES, especially when the sample size  $n$  is limited, in this case GES found a local optimum and SGS is able to find other better local optima thanks to randomization. This was also observed in [21].

When the sample size is small the learning problem becomes more difficult: the empirical distribution may be far from a perfect map resulting in many local optima and the scoring function is no more consistent, *i.e.* the likelihood does not dominate the penalty term of the scoring function which increases with the parent variable domain sizes [3]. In this complex situation, we propose a new operator to escape from some local optima.


#### 3.1 SWAP operator

Consider the 3-variable example in Figure 2 with observed data  $D$ , scoring function  $f$ , and initial DAG  $G_0 = \{X_2 \rightarrow X_3\}$ . Let assume  $f(D, \{X_1 \rightarrow X_3\}) > f(D, \{X_2 \rightarrow X_3\}) > f(D, \{X_1 \rightarrow X_3, X_2 \rightarrow$


<sup>4</sup> Two equivalence classes  $\mathcal{E}(G)$ ,  $\mathcal{E}(G')$  are adjacent iff  $G$  is an I-map of  $G'$  or vice-versa and the number of edges in the graphs  $G$  and  $G'$  differs by one.

<sup>5</sup> An arc  $X \rightarrow Y$  in  $G$  is *compelled* if that arc exists in every DAG of  $\mathcal{E}(G)$ , otherwise it is said *reversible*.

<sup>6</sup> The number of possible sets  $S$  in the GES operator  $\text{ADD}(E, S)$  is exponential in the maximum degree  $d$  of the current graph in the worst case and  $r$  is unbounded for SGS.



**Figure 1.** Four adjacent Markov-equivalence classes found by GES during its first phase of edge and v-structure insertions. (a) GES and SGS start from the empty graph. (d) The true DAG is found after three moves. The orientation of  $X3 \rightarrow X4$  and  $X1 \rightarrow X3$  edges are chosen at random by SGS, whereas GES waits until its third move to decide on edge orientations based on DAG score comparisons (enforcing the v-structure  $X1 \rightarrow X3 \leftarrow X2$  as stated by the extra ADD parameter  $\{X1\}$ , and forbidding  $X1 \rightarrow X3 \leftarrow X4$  in its second move).



**Figure 2.** The operator  $\text{SWAP}(X2|X1 \rightarrow X3)$  applied to a 3-variable problem.

$X3\}) > f(D, \{X3 \rightarrow X1, X2 \rightarrow X3\}) > f(D, \{X2 \rightarrow X1, X2 \rightarrow X3\}) > f(D, \emptyset)$ . Then  $G_0$  is a local minimum for the classical neighborhood  $\mathcal{N}^{ADR}$ . Our new operator, denoted  $\text{SWAP}(X|Y \rightarrow Z)$ , consists in changing one parent  $X$  to another parent  $Y$  for one target node  $Z$ . This is equivalent to a simultaneous pair of ADD and DELETE operators restricted to the same target node. In our example, applying  $\text{SWAP}(X2|X1 \rightarrow X3)$  corresponds to  $\text{DELETE}(X2 \rightarrow X3), \text{ADD}(X1 \rightarrow X3)$ , resulting in the better DAG  $G_1 = \{X1 \rightarrow X3\}$  as shown in Figure 2. The extended neighborhood using the four operators is denoted  $\mathcal{N}^{ADRS}$  and SGS using  $\mathcal{N}^{ADRS}$  (respectively  $\mathcal{N}^{ADR}$ ) is denoted  $\text{SGS}^2$  (Stochastic Greedy Search with Swap) (resp.  $\text{SGS}^1$ ) in the sequel.


A typical suboptimality problem that we observed in our experiments happens when two nodes have the same parents. Figure 3 shows such an example with four variables. Because child nodes  $X3$  and  $X4$  are highly correlated due to their common parents  $X1$  and  $X2$ , the first arc usually added is either  $X3 \rightarrow X4$  or  $X4 \rightarrow X3$  especially if the conditional probability distributions of  $X3$  and  $X4$  given  $X1$  and  $X2$  are close. Then adding a first parent to  $X3$  and furthermore a second parent to  $X4$  and  $X3$  give the DAG  $G_4$ . Here deleting  $X3 \rightarrow X4$  is going to decrease the score and the same negative effect results when adding  $X1 \rightarrow X4$  because it leads to a three-parent node. For this node, the increase in likelihood may not overcome the penalization term. But doing both operations simultaneously thanks to our SWAP operator results in a better local optimum DAG.

Let  $p$  be the number of variables in the current DAG and  $k$  be the maximum number of parents per node. Assuming a sparse graph,  $p \gg k$ , the number of SWAP operations is bounded by  $O(kp^2)$  in the extended neighborhood  $\mathcal{N}^{ADRS}$ , whereas it is bounded by  $O(p^2)$  for ADD and  $O(kp)$  for DELETE and REVERSE. The complexity

of  $\mathcal{N}^{ADRS}$  is therefore in  $O(kp^2)$ , whereas it is in  $O(p^2)$  for the classical neighborhood  $\mathcal{N}^{ADR}$ . Notice that other approaches using larger neighborhoods such as *h-look ahead in l good directions* (LAGD) has a worst-case complexity in  $O(l^{h-1}p^2)$  [17] and optimal reinsertion (OR) neighborhood is in  $O(2^k p^{k+1})$  [20]. In particular, LAGD complexity does not benefit from sparsity,  $l^{h-1}$  being constant, whereas our approach is faster when  $k$  decreases. Moreover computing the score difference of two DAGs before and after a SWAP operation is easy to do as it remains local to a single target node thanks to score decomposition.


Another source of suboptimality comes from the global acyclicity constraint of Bayesian networks.

### 3.2 Breaking cycles by successive deletions and swaps



**Figure 4.** Applying an extended  $\text{SWAP}^*$  operation breaking a cycle by an additional SWAP operation:  
 $\text{SWAP}^*(X2|X7 \rightarrow X3) = \{\text{SWAP}(X2|X7 \rightarrow X3), \text{SWAP}(X4|X5 \rightarrow X6)\}$ .

Consider the 7-variable DAG example in Figure 4. Swapping the parent  $X2$  of  $X3$  by  $X7$  in DAG  $G$  (Fig. 4.left) introduces a directed cycle  $\{X7 \rightarrow X3, X3 \rightarrow X4, X4 \rightarrow X6, X6 \rightarrow X7\}$  and is therefore forbidden in our  $\mathcal{N}^{ADRS}$  neighborhood. However it may correspond to a large *local* score improvement with respect to variable  $X3$ . Let us denote this improvement by  $\Delta_{X3}(G, \text{SWAP}(X2|X7 \rightarrow X3)) = f_{X3}(D, G') - f_{X3}(D, G)$  with  $G'$  obtained by applying the SWAP operation on  $G$  ( $G'$  is not a valid DAG), and  $D$  and  $f$  being the sample and scoring function. Our idea is to heuristically guide the search for a second (or more) local operator to be applied on  $G'$  in order



**Figure 3.** Problems with the classical neighborhood  $\mathcal{N}^{ADR}$  when two nodes (here  $X3$  and  $X4$ ) have the same parents. True DAG is  $G_6$ . Starting from the empty graph (not shown),  $G_4$  can be locally optimum for  $\mathcal{N}^{ADR}$  when the sample size is small (score of  $G_5$  lower than  $G_4$ ), whereas in  $\mathcal{N}^{ADRS}$ ,  $\text{SWAP}(X3|X1 \rightarrow X4)$  moves directly from  $G_4$  to  $G_6$ .

to restore graph acyclicity ( $G'$  becomes valid) and such that the *true* score of the final DAG is greater than the score of the original one. In Figure 4, it is obtained by applying a second SWAP.

---

**Algorithm 2:**  $\text{SWAP}^*(X|Y \rightarrow Z)$  operator.

---

**Input :** operation  $X|Y \rightarrow Z$ , sample  $D$ , score  $f$ , DAG  $G(\mathbf{X}, \mathbf{E})$   
**Output :** a set of local operations  $\mathbf{L}$   
 $\mathbf{L} \leftarrow \emptyset$  /\* Initialize output operations to the empty set \*/;  
 $\mathbf{X}' \leftarrow \mathbf{X}$  /\* Candidate parent set for future swaps \*/;  
 $G' \leftarrow G$  /\* Copy of input DAG \*/;  
3  $\Delta = \Delta_Z(G', \text{SWAP}(X|Y \rightarrow Z))$  /\* Putative score improvement \*/;  
**if**  $\Delta > 0$  **then**  
     $\mathbf{L} \leftarrow \mathbf{L} \cup \{\text{SWAP}(X|Y \rightarrow Z)\}$ ;  
    Apply  $\text{SWAP}(X|Y \rightarrow Z)$  to  $G'$ ;  
    /\* Repeat deletion or swap operations until no more cycles  
4   /\* **while**  $\Delta > 0 \wedge (\mathbf{C} \leftarrow \text{NextCycle}(G')) \neq \emptyset$  **do**  
5      $\mathbf{X}' \leftarrow \mathbf{X}' \setminus \text{nodes}(\mathbf{C})$ ;  
    /\* Choose the best deletion to break cycle  $\mathbf{C}$  \*/;  
6      $(U^* \rightarrow W^*) \leftarrow \arg\max_{(U \rightarrow W) \in \mathbf{C} \setminus \{Y \rightarrow Z\}} \Delta_W(G', \text{DELETE}(U \rightarrow W))$ ;  
    /\* Test if the sum of local score changes is positive \*/;  
    **if**  $\Delta + \Delta_{W^*}(G', \text{DELETE}(U^* \rightarrow W^*)) > 0$  **then**  
         $\mathbf{L} \leftarrow \mathbf{L} \cup \{\text{DELETE}(U^* \rightarrow W^*)\}$ ;  
         $\Delta \leftarrow \Delta + \Delta_{W^*}(G', \text{DELETE}(U^* \rightarrow W^*))$ ;  
        Apply  $\text{DELETE}(U^* \rightarrow W^*)$  to  $G'$ ;  
    **else**  
        /\* Choose the best swap to get a positive change \*/;  
7        $(U^*|V^* \rightarrow W^*) \leftarrow \arg\max_{(U \rightarrow W) \in \mathbf{C}, V \in \mathbf{X} \setminus \{Y\}} \Delta_W(G', \text{SWAP}(U|V \rightarrow W))$ ;  
         $\Delta \leftarrow \Delta + \Delta_{W^*}(G', \text{SWAP}(U^*|V^* \rightarrow W^*))$ ;  
        **if**  $\Delta > 0$  **then**  
             $\mathbf{L} \leftarrow \mathbf{L} \cup \{\text{SWAP}(U^*|V^* \rightarrow W^*)\}$ ;  
            Apply  $\text{SWAP}(U^*|V^* \rightarrow W^*)$  to  $G'$ ;  
        **else**  
8            $\mathbf{L} \leftarrow \emptyset$  /\* Abort all local operations \*/;  
**return**  $\mathbf{L}$  ;

---

For that purpose, we define an extended SWAP operator, denoted  $\text{SWAP}^*$ , able to break all directed cycles by performing a succession of deletion or swap operations. It can be seen as a kind of greedy descent search in the space of directed cyclic graphs, trying to remove the less important arcs or to swap them in order to compensate for their loss, until a better valid DAG is found. We use local score changes to guide the search:  $\Delta_{X_i}(G, OP) = f_{X_i}(D, G') - f_{X_i}(D, G)$ , with  $G'$  the result of applying the local move operator  $OP \in \{\text{DELETE}, \text{SWAP}\}$  to  $G$ . A negative sum of local changes aborts the search. Recall that finding a minimum number of arc dele-

tions in order to restore acyclicity is NP-hard, see *minimum feedback arc set problem* in [10]. We use a greedy approach instead. The pseudo-code of  $\text{SWAP}^*$  is given in Algorithm 2. The local score improvement of the initial SWAP operation is evaluated at line 3. It corresponds to a putative gain on the current score. If it is positive then this operation is applied to a copy of the input DAG  $G$ , checking next if it creates some directed cycles. Each cycle is retrieved by the `NextCycle` function<sup>7</sup> and the algorithm searches for an arc deletion in this cycle with minimum deterioration of the local score at line 6. If the combined local score change of the SWAP and DELETE operations is positive then it applies the selected arc deletion and continues to test if there are no more directed cycles at line 4. If the combined local score change is negative then it tries to swap an arc of the cycle such that the combined local score change is maximized (line 7) and positive. If it fails to find such an operation then it stops breaking cycles and returns an empty operation set. Finally if it succeeds, breaking all cycles, then it returns a feasible set of SWAP and DELETE operations resulting into a new valid DAG  $G'$  with a better score than  $G$ . The true score improvement is equal to  $\Delta$ .

Algorithm 2 terminates because there are at most  $O(pk)$  directed cycles to break, assuming  $p$  nodes and  $k$  maximum number of parents per node. Arcs newly created by swap operations cannot be swapped again more than  $p$  times thanks to our restricted list of alternative candidate parent nodes  $\mathbf{X}'$  used at line 7, initialized to all the variables at the beginning, and updated by the list of nodes present in the current cycle at line 5.

We apply the same approach to extend the other operators  $\text{ADD}^*$  and  $\text{REVERSE}^*$ , breaking cycles with deletions or swaps. Because  $\text{REVERSE}^*(X \rightarrow Y) = \text{ADD}^*(Y \rightarrow X)$ , we use only  $\text{ADD}^*$ . The resulting neighborhood exploiting these extended operators is denoted  $\mathcal{N}^{A^*DS^*}$  and SGS using this neighborhood is denoted  $\text{SGS}^3$  (Stochastic Greedy Search with Successive Swaps) in the experiments.

## 4 Experimental Results

In this section, we describe a set of experiments aimed at testing the performance of  $\text{SGS}^i$  algorithms compared with state-of-the-art Bayesian network structure learning algorithms on standard Bayesian networks and challenging gene regulatory networks.

<sup>7</sup> We implemented an incremental acyclicity test [14] which returns a shortest directed cycle using breadth first search.

## 4.1 Results on Standard Bayesian Networks

We used four gold-standard networks from the Bayesian Network Repository<sup>8</sup> whose main properties are shown in Table 1. In this table, *Nodes* and *Arcs* specify the number of nodes and arcs respectively in the DAG. The *Max in-degree* is the maximum number of parents of a node. The *Min-max states* are the minimum and maximum number of states in the domain of a variable associated to a node. Finally, *Longest path* is the length of the longest directed path between any pair of nodes in the DAG. 100 samples of size  $n = 500$  and  $n = 5000$  were generated for each network using Causal Explorer [1].

**Table 1.** Bayesian network properties

	ALARM	INSURANCE	HAILFINDER	PIGS
Nodes	37	27	56	441
Arcs	46	52	66	592
Max in-degree	4	3	4	2
Min-max states	2-4	2-5	2-11	3-3
Longest path	11	10	14	6

We compared  $SGS^i$  algorithms with LAGD [17], available in the WEKA software [15] and GES [3] implemented in Tetrad 4.4.0 [24]. LAGD was shown to outperform repeated hill-climbing, order-based search, tabu search, and simulated annealing in [23]. GES, which is considered the reference algorithm for Bayesian network structure learning, was reported having similar performance to the most recent order-based search algorithms in [2]. Recall that  $SGS^1$  is similar to repeated hill-climbing,  $SGS^2$  uses the SWAP operator, and  $SGS^3$  breaks cycles by successive DELETE and SWAP operators. Experiments were performed on a 3 GHz Intel Core2 computer with 4GB running Linux 2.6.

We fixed the maximum number of parents per node at  $k = 5$  for  $SGS^i$  and LAGD. LAGD exploits a  $h = 2$ -look ahead in  $l = 5$  maximum improving directions. GES was restricted on the number of adjacent nodes:  $d = 7$  for Hailfinder and  $d = 10$  for Pigs network as done in [2]. All methods were initialized by an empty graph and optimized the BDeu score with *equivalent sample size*  $\alpha = 1$  and no prior on the network structures. For each sample, we recorded the best score obtained by GES, and by  $r = 10$  randomized greedy searches for  $SGS^i$  (Algorithm 1 at line 2) as for LAGD<sup>9</sup>.

In order to test the statistical significance of the difference in BDeu score between two methods, we applied a non-parametric paired test, the Wilcoxon signed-rank test [32]. Table 2 presents the test results for all the pairs of methods by using an unilateral alternative (no difference versus better), the pairwise type I error is fixed to  $6.25 \cdot 10^{-4}$  which corresponds to a familywise error rate of 5% with Bonferroni correction for multiple comparisons. The above results were summarized in Table 3, in a summarized Wilcoxon score for each method, obtained by summing the positive comparisons and subtracting the negative ones. The non significant comparisons were ignored.

$SGS^3$  was the best method for the four networks, except for Pigs network with  $n = 5000$  which is more accurately estimated by GES. We conjecture that in this case, GES was closed to its asymptotic optimal behavior, which may be due to the smallest network features of Pigs network among all.  $SGS^2$  improved over  $SGS^1$  and reached the second position especially for small sample sizes, probably for

**Table 2.** Wilcoxon test comparing pairs of algorithms (familywise error rate = 5%). For Method1 versus Method2, + means that Method1 is significantly better than Method2, - means that Method1 is significantly worse than Method2 and ~ means there is no significant result

Sample size	ALARM		INSURANCE	
	500	5 000	500	5 000
$SGS^3$ vs $SGS^1$	+	+	+	+
$SGS^3$ vs $SGS^2$	+	+	+	+
$SGS^2$ vs $SGS^1$	+	~	~	~
$SGS^3$ vs GES	+	+	+	+
$SGS^3$ vs LAGD	+	+	+	+
$SGS^2$ vs GES	+	~	+	+
$SGS^2$ vs LAGD	~	~	+	+
$SGS^1$ vs GES	+	~	+	+
$SGS^1$ vs LAGD	~	~	+	+
LAGD vs GES	+	~	+	+
Sample size	HAILFINDER		PIGS	
	500	5 000	500	5 000
$SGS^3$ vs $SGS^1$	~	+	+	+
$SGS^3$ vs $SGS^2$	~	+	+	+
$SGS^2$ vs $SGS^1$	~	~	~	~
$SGS^3$ vs GES	+	+	+	-
$SGS^3$ vs LAGD	~	+	n/a	n/a
$SGS^2$ vs GES	+	+	~	-
$SGS^2$ vs LAGD	~	+	n/a	n/a
$SGS^1$ vs GES	+	+	~	-
$SGS^1$ vs LAGD	~	+	n/a	n/a
LAGD vs GES	+	+	n/a	n/a

**Table 3.** Summarized Wilcoxon scores (the higher the score, the better the method)

	ALARM	INSURANCE	HAILFINDER	PIGS
$SGS^3$	8	8	5	4
$SGS^2$	0	2	2	-3
$SGS^1$	-2	2	2	-3
LAGD	-1	-4	-1	n/a
GES	-5	-8	-8	2

**Table 4.** Number of spurious edges (+) and missing edges (-) to sum for the structural Hamming distance. Both scores are in bold when giving the best distance per configuration

Sample size		ALARM		INSURANCE	
		500	5 000	500	5 000
$SGS^3$	+	<b>8</b>	6	<b>4</b>	<b>2</b>
	-	<b>3</b>	2	<b>20</b>	<b>8</b>
LAGD	+	11	8	<b>4</b>	5
	-	4	2	<b>20</b>	11
GES	+	<b>6</b>	<b>4</b>	2	3
	-	<b>5</b>	<b>2</b>	23	12
Sample size		HAILFINDER		PIGS	
		500	5 000	500	5 000
$SGS^3$	+	17	<b>16</b>	32	41
	-	24	<b>13</b>	0	0
LAGD	+	21	20	n/a	n/a
	-	26	19	n/a	n/a
GES	+	<b>15</b>	11	<b>2</b>	<b>0</b>
	-	<b>24</b>	22	<b>7</b>	<b>0</b>

<sup>8</sup> <http://www.cs.huji.ac.il/site//labs/compbio/Repository/>

<sup>9</sup> We randomly permute the input variables at each run.

reasons commented in Figure 3. LAGD got poor results, the difference with  $\text{SGS}^1$  can be explained by a better randomization in  $\text{SGS}^1$  (Algorithm 1 at line 1). LAGD failed on the Pigs network due to the large number of variables  $p = 441$  that makes the exploration of 2-look ahead neighborhoods infeasible in a reasonable time. GES was the worst method of this evaluation (except for Pigs) due to limited sample sizes.

Although the algorithms are designed to maximize a (BDeu) score, we generally look for a network structure as close as possible to the true one. When the sample size grows large, improving the BDeu score leads to reduce the *structural Hamming distance* (SHD) between the learned and the true network. We report in Table 4 the means over 100 datasets (rounded values to the nearest integer) of the missing and spurious edges without taking into account the edge orientations. The SHD is the sum of the above values. We do not report  $\text{SGS}^1$  and  $\text{SGS}^2$  results that are outperformed by  $\text{SGS}^3$ .  $\text{SGS}^3$  (resp. GES) got the best SHD in 4 (resp. 5) configurations and outperformed LAGD (which performed as  $\text{SGS}^3$  in 1 configuration). GES performed extremely well on the Pigs network, finding the true network with 5,000 samples, whereas  $\text{SGS}^3$  learned too many edges but recovered all the true edges (even with  $n = 500$ ). The spurious edges learned by  $\text{SGS}^3$  are exclusively due to random orientations of compelled arcs in v-structures (see Figure 1). Assuming  $X1 \rightarrow X3 \leftarrow X2$  in the true network (v-structures are very frequent in the Pigs network) and a large sample size, if during its greedy search  $\text{SGS}^3$  adds first  $X1 \leftarrow X3$  and  $X3 \rightarrow X2$  then it will add next a *covering edge*  $X1 \rightarrow X2$  or  $X1 \leftarrow X2$  in order to find a minimal independence map (see Proposition 1). We corrected this fault by comparing for each covered v-structure the 3 possible non-covered v-structures independently from the other variables and selecting the highest score configuration. After this post-processing step on Pigs network, we had only 1 (resp. 2) spurious edges for  $\text{SGS}^3$  with 500 (resp. 5000) samples without increasing the missing edges number. The same post-processing had no effect on the other smaller networks.

## 4.2 Detailed analysis on the Alarm network


We conducted a series of algorithm analyzes on the Alarm network with a fixed sample size  $n = 500$ .

**Table 5.** Single greedy search analysis on the Alarm network ( $n = 500, r = 1$ )

	BDEU SCORE	ITER.	BDEU CACHE	TIME(s)
$\text{SGS}^3$	-5490.53	66	4543	3.2
$\text{SGS}^2$	-5513.89	58	4310	2.3
$\text{SGS}^1$	-5541.55	55	3305	1.5
LAGD	-5544.61	35	4782	3.2
GES	-5659.26	72	5531	2.4

Table 5 shows a typical example of the BDeu score reached by the algorithms for a single greedy search ( $r = 1$ ) on a particular sample. It also reports the number of local moves *Iter.* (Algorithm 1 at line 1), the number of local score  $f_{X_i}(D, G)$  computations for *different* parent configurations, i.e. assuming a *perfect BDeu cache*, and the CPU time in seconds. As expected, the number of iterations for LAGD was less than half that of other methods, since LAGD applies two local operators (ADD, DELETE, or REVERSE) at each move. The swap operations used by  $\text{SGS}^2$  and  $\text{SGS}^3$  slightly increased the number of moves allowing to find better local optima. Comparing CPU times is a difficult task due to the various implementation details (language choice, caching strategy, sample compilation techniques, see

for instance Chapter 18.4.3 in [18]). Nevertheless  $\text{SGS}^i$  algorithms are comparable. In our implementation,  $\text{SGS}^3$  was twice slower than  $\text{SGS}^1$ . A fairer comparison is the number of calls to new local score computations.  $\text{SGS}^1$  needed the fewest number, whereas GES needed the most as it tends to explore larger neighborhoods in the space of Markov-equivalent networks. Notice that  $\text{SGS}^2$  and  $\text{SGS}^3$  required less score computations than LAGD due to graph sparsity (see Section 3.1). The greater BDeu cache size of  $\text{SGS}^3$  compared to  $\text{SGS}^2$  is mainly due to the increased number of local moves.



**Figure 5.** Best BDeu scores, averaged on 30 Alarm samples ( $n = 500$ ), found by  $\text{SGS}^i$  algorithms as the number of restarts  $r$  increases and starting either from an empty (solid line) or a random graph (dashed line). Results of GES (dotted line) and BDeu score of the true network *Gold* (dash-dotted line) are also given. Methods are sorted from the best,  $\text{SGS}^3$  with an empty graph, to the worst,  $\text{SGS}^1$  with a random initial graph, at  $r = 1$ .

We further analyzed the impact on performances of the number of restarts  $r$  and the initial graph used by  $\text{SGS}^i$  algorithms (see *InitGraph* in Algorithm 1). Figure 5 reports averaged BDeu scores on 30 Alarm samples (the length of the 95% confidence interval of the BDeu score mean was around  $\pm 30$ ). The 30 initial random graphs, common for  $\text{SGS}^i$  algorithms, are composed of 71 arcs with at most two parents per node<sup>10</sup>. All  $\text{SGS}^i$  methods reached a better BDeu score than GES in this small sample size situation by the use of less than  $r = 4$  restarts.  $\text{SGS}^i$  methods converged quite rapidly as  $r$  increases. Only  $\text{SGS}^3$  found a better score than the true network. Initial random graphs were counter-productive for all the methods, except for  $\text{SGS}^3$ . It shows that starting from a random graph is useful if the available operators allow to move efficiently in the search space, which is illustrated by the following experiment. On the contrary, using randomness within the greedy selection process (see *SelectRandom* in Algorithm 1) was always beneficial.

Finally, we analyzed how often our new local move operators are used during the search. Table 6 shows the mean number of local moves using a specific operator averaged on  $r = 50$  greedy searches over 30 Alarm samples ( $n = 500$ ). Labels ADD+ and SWAP+ mean that at least two local operators were successively applied at a given move in order to break cycles<sup>11</sup>. The new local move operators are

<sup>10</sup> For each node, we randomly select two parents and remove a parent if it creates a directed cycle.

<sup>11</sup> We found at most 4 (resp. 5) successive operations for the 1500 runs starting from an empty (resp. random) graph.

**Table 6.** Operator usage depending on the initial graph

	EMPTY GRAPH INIT.			RANDOM GRAPH INIT.		
	SGS <sup>1</sup>	SGS <sup>2</sup>	SGS <sup>3</sup>	SGS <sup>1</sup>	SGS <sup>2</sup>	SGS <sup>3</sup>
ADD	53	53	52	53	9.5	11.5
DELETE	0.5	0.1	1.4	64	20	43
REVERSE	0.9	1.1	0.7	5.3	3.4	1.9
SWAP	0	1.1	3.2	0	47	33
ADD+	0	0	2	0	0	14
SWAP+	0	0	1.6	0	0	9.5

mostly used when starting from a random graph (there are more cycles to break), but when starting from an empty graph, only a few applications of the extended moves allow to significantly improve the BDeu score as previously shown in Figure 5 and Table 3.

### 4.3 Results on Gene Regulatory Networks

Gene regulatory network reconstruction from gene expression data using Bayesian network structure learning was first proposed in [11]. We used simulated expression datasets of the Systems Genetics Challenge A from the international reverse engineering competition DREAM5 [9]. Genetics data were not used as they require additional modelling to be taken into account, see *e.g.* [31]. Expression data were generated using the SysGenSIM generator [22] based on ordinary differential equation simulation. Five datasets are available for three different sample sizes ( $n = 100, 300$ , and  $999$ ). The 15 datasets were obtained from *different* known gene networks composed of 1,000 variables and containing directed cycles. For each sample size, the five network structures contain a different number of edges varying from  $\approx 2,000$  (*Net1*) to more than 5,000 (*Net5*). We discretized gene expression levels into 2 to 4 states using an adaptive method based on an adapted  $k$ -means algorithm and the more general framework of Gaussian mixture models as described in [31].

With such large networks, we had to adapt the learning procedure of SGS<sup>i</sup> algorithms<sup>12</sup>. We decided to restrict their lists of candidate parents as done in [13]: we selected for each variable  $X$  the set of parents  $S$  such that each element  $Y$  of  $S$  improves the local BDeu score when it is considered as a unique parent compared to the orphan situation ( $f_X(D, \{Y \rightarrow X\}) > f_X(D, \emptyset)$ ). This filtering process was done prior to the search<sup>13</sup>. In these experiments, SGS<sup>i</sup> algorithms have a maximum number of parents per node fixed at  $k = 5$  and use  $r = 10$  restarts. Instead of LAGD (which was too slow), we used MMHC [30] having two steps similar to SGS<sup>i</sup>. It first selects the skeleton of the graph using mutual information measures (MMPC [29] filtering step) and then orientates edges in a greedy manner. We recorded the best BDeu score of 10 runs for MMHC, by randomly permuting the input variables at each run. The skeleton being built locally for each variable, MMHC can handle large networks. All the methods started from an empty graph and optimized the BDeu score with  $\alpha = 1$  and no prior on the network structures.

It was difficult to perform statistics with the results of this experiment. Indeed, contrary to the standard network experiment, there were no replicated samples of the same network. We decided to pool results per sample size and we performed the Wilcoxon test on the groups. With only five results per group to compare the methods, we are aware that the power of the test is very low. So, we applied a

<sup>12</sup> GES managed in  $\sim 1$ -hour CPU time each network thanks to its better implementation of caching and heap data structure.

<sup>13</sup> It could also be done during the search as in [12].

**Table 7.** Wilcoxon test (error rate = 5%) for different gene network sample sizes

	100	300	999
SGS <sup>3</sup> vs MMHC	+	+	+
SGS <sup>3</sup> vs GES	+	+	+
MMHC vs GES	-	~	+

pairwise type I error of 5% and we did not try to correct for multiple comparisons, see Table 7. However, it is worth noting that SGS<sup>3</sup> was always the best method and that it increased the BDeu score by about 2% in average.

Surprisingly, GES appeared to be better on smaller sample sizes compared to MMHC. As explained below, MMHC was penalized by its filtering process, especially on the smallest sample size, whereas GES had no restrictions on the candidate parent sets.

In these experiments, the structural Hamming distance (SHD) was not informative due to the poor results reached by all tested algorithms for such large networks. SHD, equal to the sum of spurious edges ( $s$ ) and missing edges ( $m$ ), was greater than the number of edges ( $e$ ) of the true graph. In this situation, even the empty structure appears better. For this reason, we computed another aggregated structural quality measure based on the *precision* ( $\frac{e-m}{e-m+s}$ ) and *recall* ( $\frac{e-m}{e}$ ) measures and took the Euclidean distance to the origin to combine them ( $\sqrt{\text{precision}^2 + \text{recall}^2}$ ). Contrary to SHD, a high distance indicates a better structural quality. We observed in Table 8 contrasted performances between the tested methods depending on the sample size: for  $n=100$ , MMHC got the best results, for  $n = 300$ , it was GES, and finally SGS<sup>3</sup> performed the best for the largest sample size. Better BDeu scores are not always synonymous with a better structural quality, the limited sample size in addition to the non faithfulness of the data (generated by ordinary differential equations with cycles) could explain this behavior. We should notice that the good optimization behavior of SGS<sup>3</sup> in terms of BDeu scores resulted in a better structural quality than GES and MMHC as sample size increases.

**Table 8.** Euclidean distances to the origin of the (precision, recall) values. Best distances per sample size are in bold

$n$		SGS <sup>3</sup>	MMHC	GES
100	<i>Net1</i>	0.170	<b>0.218</b>	0.196
	<i>Net2</i>	0.213	<b>0.295</b>	0.232
	<i>Net3</i>	0.214	<b>0.266</b>	0.236
	<i>Net4</i>	0.201	<b>0.265</b>	0.214
	<i>Net5</i>	0.206	<b>0.247</b>	0.243
300	<i>Net1</i>	<b>0.510</b>	0.483	0.464
	<i>Net2</i>	0.342	0.337	<b>0.385</b>
	<i>Net3</i>	0.484	0.488	<b>0.505</b>
	<i>Net4</i>	0.453	0.478	<b>0.498</b>
	<i>Net5</i>	0.419	0.397	<b>0.428</b>
999	<i>Net1</i>	<b>0.578</b>	0.537	0.549
	<i>Net2</i>	<b>0.581</b>	0.510	0.505
	<i>Net3</i>	0.454	0.441	<b>0.484</b>
	<i>Net4</i>	<b>0.476</b>	0.450	<b>0.476</b>
	<i>Net5</i>	<b>0.479</b>	0.471	0.458

Moreover, we compared the quality of the filtering process used by MMPC and SGS<sup>3</sup>. Table 9 reports for the first network (*Net1*) of each sample size, the total number of arcs kept by the filtering process and the recall value, which represents the percentage of true edges in the filter. Our first observation is the poor recall of both filtering processes, which indicates strong structural differences between the



true network and the learned network even with  $n = 999$  sample size. Our filtering approach obtained better recall values with similar compression sizes than MMPC.

Finally, we tried the same methods as in Section 4.1 on 50 small gene regulatory networks (*Web50* Artificial Gene Networks with 50 nodes and 50 arcs) without doing any filtering process. For a sample size  $n = 500$ , SGS<sup>3</sup> was significantly better than GES and LAGD in terms of BDeu scores and slightly better in terms of Euclidean distances to the origin of the (precision, recall) values.

**Table 9.** Total size and recall of candidate parent lists for SGS<sup>3</sup> and MMPC on the most sparse gene network *Net1*

		100	300	999
BDeu test	size	2670	2430	5984
	recall	9%	18%	35%
MMPC	size	2568	3064	3842
	recall	5%	12%	23%

## 5 Conclusion

We have presented in this paper a new greedy search algorithms called SGS<sup>i</sup> exploiting stochasticity from two random draws. We have developed a new local move operator called SWAP and extended versions for ADD and SWAP operators to overcome frequent limitations of local search methods which are local maxima and cyclic situations. We compared SGS<sup>3</sup> using SWAP and extended operators to state-of-the-art methods and we obtained significant BDeu and recall value improvements on classical benchmarks when the sample size is small. The complexity of SGS<sup>3</sup> stays moderate with sparse networks. In case of large networks with many variables we applied a filtering process in preprocessing using the same criterion as for the search. This process kept more edges from the true network than mutual information-based methods with significant reduction of the search space.

In the future, we would like to test our new operators with other local search methods like tabu search.

## REFERENCES

- [1] C. Aliferis, I. Tsamardinos, A. Statnikov, and L. Brown, 'Causal Explorer: A Probabilistic Network Learning Toolkit for Biomedical Discovery', in *Proceedings of the International Conference on Mathematics and Engineering Techniques in Medicine and Biological Sciences*, (2003).
- [2] J. Alonso-Barba, L. de la Ossa, and J. Puerta, 'Structural learning of bayesian networks using local algorithms based on the space of orderings', *Soft Comput.*, 1881–1895, (2011).
- [3] D. Chickering, 'Optimal structure identification with greedy search', *Journal of Machine Learning Research*, **3**, 507–554, (2002).
- [4] D. Chickering, D. Heckerman, and C. Meek, 'Large-Sample Learning of Bayesian Networks is NP-Hard', *Journal of Machine Learning Research*, **5**, 1287–1330, (2004).
- [5] D. Chickering and D. Heckermann, 'Learning bayesian networks is NP-complete', *In learning from data: AI and Statistics*, (1996).
- [6] G. Cooper and E. Hersovits, 'A Bayesian method for the induction of probabilistic networks from data', *Machine Learning*, **9**, 309–347, (1992).
- [7] R. Daly and Q. Shen, 'Learning Bayesian Network Equivalence Classes with Ant Colony Optimization', *Journal of Artificial Intelligence Research*, **35**, 391–447, (2009).
- [8] R. Daly, Q. Shen, and S. Aitken, 'Learning Bayesian networks: approaches and issues', *The Knowledge Engineering Review*, **26**(2), 99–157, (2011).
- [9] 'The DREAM5 Systems Genetics Challenges'. <http://wiki.c2b2.columbia.edu/dream/index.php/D5c3>, 2010.
- [10] Paola Festa, Panos M. Pardalos, and Mauricio G. C. Resende, 'Feedback Set Problems', in *Encyclopedia of Optimization*, 1005–1016, Springer, (2009).
- [11] N. Friedman, M. Linial, I. Nachman, and D. Pe'er, 'Using bayesian networks to analyse expression data', *Journal of computational biology*, **7**(3/4), 601–620, (2000).
- [12] J. Gámez, J. Mateo, and J. Puerta, 'Learning Bayesian networks by hill climbing: efficient methods based on progressive restriction of the neighborhood', *Data Min. Knowl. Discov.*, **22**, 106–148, (2011).
- [13] A. Goldenberg and A. Moore, 'Tractable learning of large Bayes net structures from sparse data', in *Proc. of ICML'04*, pp. 44–51, (2004).
- [14] B. Haeupler, T. Kavitha, R. Mathew, S. Sen, and R. Tarjan, 'Faster algorithms for incremental topological ordering', in *Proc. of ICALP*, pp. 421–433, (2008).
- [15] M. Hall, F. Eibe, G. Holmes, B. Pfahringer, P. Reutemann, and I. Witten, 'The WEKA Data Mining Software', *SIGKDD Explorations*, **11**, (2009).
- [16] D. Heckerman, D. Geiger, and D. Chickering, 'Learning Bayesian Networks: The Combination of Knowledge and Statistical Data', in *Machine Learning*, volume 20, pp. 197–243, (1995).
- [17] A. Holland, M. Fathi, M. Abramovici, and M. Neubach, 'Competing fusion for bayesian applications', in *Proc. of IPMU 2008*, pp. 378–385, (2008).
- [18] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, 2009.
- [19] G. Melançon, I. Dutour, and M. Bousquet-Mélou, 'Random generation of directed acyclic graphs', *Electronic Notes in Discrete Mathematics*, **10**, 202–207, (2001).
- [20] A. Moore and W.K. Wong, 'Optimal reinsertion: A new search operator for accelerated and more accurate bayesian network structure learning', in *Proc. of ICML '03*, pp. 552–559, (2003).
- [21] J. Nielsen, T. Kocka, and J. Pefia, 'On Local Optima in Learning Bayesian Networks', in *Proc. of UAI-03*, pp. 435–442, (2003).
- [22] A. Pinna, N. Soranzo, I. Hoeschele, and A. de la Fuente, 'Simulating system genetics data with SysGenSIM', *Bioinformatics*, **27**, 2459–2462, (2011).
- [23] E. Salehi and R. Gras, 'An empirical comparison of the efficiency of several local search heuristics algorithms for Bayesian network structure learning', in *Learning and Intelligent Optimization Workshop (LION 3)*, (2009).
- [24] R. Scheines, P. Spirtes, C. Glymour, C. Meek, and T. Richardson, 'The TETRAD Project: Constraint Based Aids to Causal Model Specification', *Multivariate Behavioral Research*, **33**(1), 65–117, (1998).
- [25] M. Schmidt, A. Niculescu-Mizil, and K. Murphy, 'Learning graphical model structure using L1-regularization paths', in *Proc. of AAAI'07*, pp. 1278–1283, (2007).
- [26] G. Schwarz, 'Estimating the dimension of a model', *The annals of statistics*, (1978).
- [27] P. Spirtes, C. Glymour, and R. Scheines, 'Causality from probability', *Evolving knowledge in the natural and behavioral sciences*, 181–199, (1990).
- [28] M. Teyssier and D. Koller, 'Ordering-based Search: A Simple and Effective Algorithm for Learning Bayesian Networks', in *Proc. of UAI'05*, pp. 584–590, (2005).
- [29] I. Tsamardinos, C. Aliferis, and A. Statnikov, 'Time and sample efficient discovery of Markov blankets and direct causal relations', in *Proc. of KDD'03*, pp. 673–678, (2003).
- [30] I. Tsamardinos, L. Brown, and C. Aliferis, 'The max-min hill-climbing Bayesian network structure learning algorithm', *Mach. Learn.*, **65**, 31–78, (2006).
- [31] M. Vignes, J. Vandel, D. Allouche, N. Ramadan-Alban, C. Cierco-Ayrolles, T. Schiex, B. Mangin, and S. de Givry, 'Gene Regulatory Network Reconstruction Using Bayesian Networks, the Dantzig Selector, the Lasso and Their Meta-Analysis', *PLoS ONE*, **6**, (2011).
- [32] F. Wilcoxon, 'Individual Comparisons by Ranking Methods', *Biometrics Bulletin*, **1**, 80–83, (1945).