

Inferring large graphs using l1-penalized likelihood

Magali Champion, Victor Picheny, Matthieu Vignes

► **To cite this version:**

Magali Champion, Victor Picheny, Matthieu Vignes. Inferring large graphs using l1-penalized likelihood. 2017. hal-01172745v2

HAL Id: hal-01172745

<https://hal.archives-ouvertes.fr/hal-01172745v2>

Preprint submitted on 24 Jan 2017 (v2), last revised 4 Oct 2017 (v3)

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.

Inferring large graphs using ℓ_1 -penalized likelihood

Magali Champion

Laboratoire MAP5, Université Paris Descartes,
Sorbonne Paris Cité, France

Victor Picheny

INRA, UR875 Applied Mathematics and Computer Science Unit,
Castanet-Tolosan, France

and

Matthieu Vignes

Institute of Fundamental Sciences, Massey University,
Palmerston North, New Zealand.

Abstract

We address the issue of recovering the structure of large sparse directed acyclic graphs from noisy observations of the system. We propose a novel procedure based on a specific formulation of the ℓ_1 -norm regularized maximum likelihood, which decomposes the graph estimation into two optimization sub-problems: topological structure and node order learning. We provide oracle inequalities for the graph estimator, as well as an algorithm to solve the induced optimization problem, in the form of a convex program embedded in a genetic algorithm. We apply our method to various data sets (including data from the DREAM4 challenge) and show that it compares favorably to state-of-the-art methods.

Keywords: Directed Acyclic Graphs Lasso Convex program Optimization

1 Introduction

Revealing the true structure of a complex system is paramount in many fields to identify system regulators, predict its behavior or decide where interventions are needed to disentangle direct relationships (Newman, 2003; Souma et al., 2006; Verma et al., 2014). This problem can often be seen as a graph inference problem. Given observational data, we aim at predicting the presence (or absence) of edges between elements of the system, which form the vertices of a graph. Edges specify the relational structure of the system depicted as a graph or network. As a motivating problem, the reconstruction of Gene Regulatory Networks (GRN), which model activation and inhibition relationships between genes, is one of the main challenges in modern computational biology (Barabási & Oltvai, 2004).

A popular approach consists in assuming that the data are generated by a Directed Acyclic Graph (DAG) (Pearl, 2009). DAGs are made of a collection of vertices, which stand for variables, and directed edges to model the dependency structure among the variables, avoiding self-loops and cycles. However, inferring a DAG is a rather challenging problem. Firstly, the number of nodes p of the graph may be so large that exploring relevant DAG topologies is simply infeasible, since the number of possible DAG structures is super-exponential in p (Robinson, 1973; Koivisto & Sood, 2004; Tsamardinos et al., 2006; Grzegorzczuk & Husmeier, 2008). Another dimension flaw occurs when p , even being reasonable, is larger than the number of observations, and parameter estimation is jeopardized. High-dimensional statistical techniques are then needed to overcome this issue (Bühlmann & van de Geer, 2011; Giraud, 2015). Secondly, even if the ratio between p and the sample size n is not impeding model estimation, the nature of the data can be an additional obstacle (Ellis & Wong, 2008; Guyon et al., 2010; Fu & Zhou, 2013). The available observational data are in general not sufficient to identify the true underlying DAG, and can only determine an equivalence class of DAGs (Verma & Pearl, 1991). This approach relies on the assumption that the joint distribution is Markov and faithful with respect to the true graph (Spirtes et al., 2000).

A large number of methods have been proposed for estimating DAGs, including for instance score-based methods (Bayesian score, Friedman & Koller 2003 or Bayesian Information Criterion, Schwarz 1978), complex space sampling (Zhou, 2011) or the PC algorithm (Spirtes et al., 2000). The latter has been proved to be uniformly consistent in the high-dimensional case, but requires a test of conditional independences that quickly becomes computationally intractable (Kalisch & Bühlmann, 2007).

In this work, we focus on Gaussian structural equation models (Wright, 1921) associated with maximum likelihood estimators (MLE). In the last years, the ℓ_0 -regularization of the MLE drew the attention of a large number of works since it leads to infer sparse graphs. In DAGs, a not necessarily topological ordering of the nodes can always be defined according to edge distribution (Kahn, 1962). Identifying this ordering is known to be a challenging problem (Cook, 1985). When the order of the variables is unknown, van de Geer & Bühlmann (2013) studied the convergence of the ℓ_0 -penalized likelihood. However, the ℓ_0 -regularized approaches (Silander & Myllymäki, 2006; Hauser & Bühlmann, 2012) remain impractical for estimating graphs with more than 20 vertices, either due to an exhaustive exploration of the set of DAGs or overfitting (Chen & Chen, 2008). For a

known order among the variables in the graph, Shojaie & Michailidis (2010) present results for the estimation of high-dimensional graphs based on independent linear regressions using an adaptive Lasso scheme.

Our objective consists in overcoming this drastic dimensional limitation, and find inference strategies for graphs with up to several hundred nodes. Such strategies must ensure a high level of sparsity, be supported by computationally affordable algorithms, while preserving sound theoretical bases. Here, we propose to use the ℓ_1 -regularization, similarly to Fu & Zhou (2013) and Shojaie & Michailidis (2010), to penalize the MLE. From a computational point of view, this regularization makes the criterion to maximize partially convex while ensuring sparse estimates. Our contribution is two-fold: firstly, we provide oracle inequalities that guarantee good theoretical performances of our proposed estimator in the sparse high-dimensional setting. Secondly, we provide an efficient algorithm to infer the true unknown DAG, in the form of a convex program embedded in a genetic algorithm.

The next section covers the model definition and the associated penalized MLE problem. Section 3 details the oracle inequalities, and Section 4 our inference algorithm. Section 5 reports numerical experiments both on toy problems and realistic data sets.

2 The ℓ_1 -penalized likelihood for estimating DAGs

2.1 DAG's modelling and estimation

This work considers the framework of an unknown DAG $\mathcal{G}_0 = (V, E)$, consisting of vertices $V = \{1, \dots, p\}$ and a set of edges $E \subseteq V \times V$. The p nodes are associated to random variables X^1, \dots, X^p . A natural approach, developed by Meinshausen & Bühlmann (2006) to solve the network inference problem is to consider that each variable X^i ($1 \leq i \leq p$) of the DAG can be represented as a linear function of all other variables X^j ($j \neq i$) through the Gaussian Structural Equation Model:

$$\forall j \in \llbracket 1, p \rrbracket, \quad X^j = \sum_{i=1}^p (G_0)_i^j X^i + \varepsilon^j, \quad (1)$$

with $\varepsilon^j \sim \mathcal{N}(0, \sigma_j^2)$ (σ_j^2 known) a Gaussian residual error term. The set of edges E corresponds to the non-zero coefficients of G_0 , *i.e.* $(G_0)_i^j$ encoding the relationship from variable X^i to variable X^j .

Assume that we observe an n -sample consisting of n i.i.d. realizations (X^1, \dots, X^p) of Equation (1), distributed according to a $\mathcal{N}(0, \Sigma)$ law where Σ is non-singular. We denote by $X := (X^1, \dots, X^p)$ the $n \times p$ data matrix. The relations between the variables can then be represented in its matrix form:

$$X = XG_0 + \varepsilon, \quad (2)$$

where $G_0 = ((G_0)_i^j)_{1 \leq i, j \leq p}$ is the $p \times p$ matrix compatible with the graph \mathcal{G}_0 and $\varepsilon := (\varepsilon^1, \dots, \varepsilon^p)$ is the $n \times p$ matrix of noise vectors.

The negative log-likelihood of the model is then (Rau et al., 2013):

$$\ell(G) = \frac{np}{2} \log(2\pi) + n \sum_{j=1}^p \log \sigma_j + \sum_{k=1}^n \sum_{j=1}^p \frac{1}{\sigma_j^2} (X_k(I - G)^j)^2. \quad (3)$$

To recover the structure of the DAG \mathcal{G}_0 and make the estimated graph sparse enough, we focus on a penalized maximum likelihood procedure (Bickel & Li, 2006):

$$\hat{G} = \operatorname{argmin}_{G \in \mathcal{G}_{DAG}} \{ \ell(G) + \lambda \operatorname{pen}(G) \}, \quad (4)$$

where $\ell(\cdot)$ is the negative log-likelihood of Equation (3), $\operatorname{pen}(\cdot)$ is a penalization function, λ is a trade-off parameter between penalization and fit to the data, and \mathcal{G}_{DAG} is the set of $p \times p$ matrices compatible with a DAG over p nodes.

In the setting of Gaussian structural equation models with equal noise variance, Peters et al. (2011, 2014) showed that the true DAG was identifiable for respectively discrete and continuous data. In a nutshell, it implies that the true DAG could be inferred, not just the Markov equivalence class of the underlying DAG - a partially directed graph exactly encoding the conditional dependency structure. Using an ℓ_0 -norm regularization in Equation (4) is an attractive option to infer sparse graphs. From a computational point of view, the main difficulty when solving the optimization problem in Equation (4) lies in exploring the set of DAGs \mathcal{G}_{DAG} . (Chickering, 1996) showed it to be an NP-hard problem: an ℓ_0 -regularization does not set a favorable framework for this task. To avoid the whole exploration of \mathcal{G}_{DAG} , a dynamic programming method has been proposed in Silander & Myllymäki (2006), using a particular decomposition of the ℓ_0 -penalized maximum likelihood. The greedy equivalent search algorithm of Chickering (2002) is a hill climbing alternative method. (Hauser & Bühlmann, 2012) rather restricted the search space to the smaller space of equivalence classes, and they provide an efficient algorithm without enumerating all the equivalent DAGs. They showed that they are asymptotically optimal under a faithfulness assumption (i.e. independences in the distribution are those read from \mathcal{G}_0). However, none of the approaches above can be used on high-dimensional data to estimate graphs with a large number of nodes. In this context, we focus on the ℓ_1 -norm convex regularization instead of ℓ_0 for its sparse, high-dimensional and computational properties.

The ℓ_1 -regularization clearly improves the computation of (4). It allows us to write a convex formulation of the problem (see Section 2.2). Given Equation (3) and omitting constant terms, the ℓ_1 -penalized likelihood estimator we consider is:

$$\hat{G} = \operatorname{argmin}_{G \in \mathcal{G}_{DAG}} \left\{ \frac{1}{n} \|X(I - G)\|_F^2 + \lambda \|G\|_1 \right\}, \quad (5)$$

where, for any matrix $M := (M_i^j)_{1 \leq i, j \leq p}$, we denote by $\|M\|_F = \sum_{i,j} (M_i^j)^2$ the Frobenius norm and by $\|M\|_1 = \sum_{i,j} |M_i^j|$ the ℓ_1 -norm.

2.2 A new formulation for the estimator

We propose here a new formulation of the minimization problem of Equation (5). It naturally decouples the initial problem into two steps of the minimisation procedure: node

ordering and graph topology search. A key property is that any DAG leads to a topological ordering of its vertices, denoted \leq , where a directed path from node X^i to node X^j is equivalent to $X^j \leq X^i$ (Kahn, 1962; Cormen et al., 2001) (see Example 1 below for more explanations). This ordering is not unique in general. Proposition 2.1 from Bühlmann (2013) then gives an equivalent condition for a matrix to be compatible with a DAG.

Proposition 2.1 (Bühlmann, 2013) *A matrix G is compatible with a DAG \mathcal{G} if and only if there exists a permutation matrix P and a strictly lower triangular matrix T such that:*

$$G = PTP^T.$$

Graphically, the permutation matrix sets an ordering of the nodes of the graph and is associated to a complete graph. The strictly lower triangular matrix T sets the graph structure, *i.e.* the non-zero entries of G , as illustrated in Example 1.

Example 1 *Consider the DAG \mathcal{G} given in Figure 1 (left). The corresponding matrix G can then be written as the strictly lower-triangular matrix T by permutation of its rows and columns using P :*

$$G = \begin{pmatrix} 0 & 0 & 0 & 7 & 5 \\ 2 & 0 & 1 & 6 & 4 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 3 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} = PTP^T,$$

$$\text{with } T = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 & 0 \\ 5 & 0 & 7 & 0 & 0 \\ 4 & 1 & 6 & 2 & 0 \end{pmatrix} \quad \text{and} \quad P = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Looking at the non-zero values of P column by column, P defines a node hierarchy $X^5 \leq X^3 \leq X^4 \leq X^1 \leq X^2$ compatible with the topological orderings of \mathcal{G} . Graphically, P is associated to the complete graph represented in Figure 1 (right). The dashed edges then correspond to the lower zero entries of T . Note that since X^3 is not connected with X^1 , X^4 and X^5 , four topological orderings are possible ($X^5 \leq X^4 \leq X^1 \leq X^3$, $X^5 \leq X^4 \leq X^3 \leq X^1$, $X^5 \leq X^3 \leq X^4 \leq X^1$ and $X^3 \leq X^5 \leq X^4 \leq X^1$).

Using Proposition 2.1, the estimator in (5) leads to the following equivalent optimization problem:

$$(\hat{P}, \hat{T}) = \underset{(P,T) \in \mathcal{C}}{\operatorname{argmin}} \left\{ \frac{1}{n} \|X(I - PTP^T)\|_F^2 + \lambda \|T\|_1 \right\}, \quad (6)$$

where the optimization space of constraints \mathcal{C} is defined as $\mathcal{C} = \mathbb{P}_p(\mathbb{R}) \times \mathbb{T}_p(\mathbb{R})$, with $\mathbb{P}_p(\mathbb{R})$ the set of permutation matrices and $\mathbb{T}_p(\mathbb{R})$ the set of strictly lower-triangular matrices. Note that a similar formulation has already been proposed by van de Geer & Bühlmann (2013) to ensure good theoretical properties for the ℓ_0 -penalized log-likelihood estimation. However, it has never been exploited from a computational point of view to recover the




Figure 1: An example of DAG \mathcal{G} (top) and the action of P and T on \mathcal{G} : P is associated to a complete graph that orders the nodes of the graph (bottom) and T sets the weights on the edges. The dashed edges correspond to null weight edges (a zero entry in T).

graph structure optimizing problem (5). In the following two sections, we propose a theoretical analysis of the proposed estimator (Section 3) and a computationally efficient algorithm to solve Problem (6) (Section 4).

3 Oracle inequalities for the DAG estimation

The main result of this section deals with convergence rates: in Theorem 1, we provide upper bound for error associated with the ℓ_1 -penalized maximum likelihood estimator considered in Equation (6), both in prediction (Equation 7) and estimation (Equation 8). Following the works of van de Geer & Bühlmann (2013) on the ℓ_0 -penalized maximum likelihood estimator and of Bickel et al. (2009) on the Lasso and the Dantzig Selector, we obtain two convergence results under some mild sparsity assumptions, when the number of variables is large but upper bounded by a function $\varphi(n)$ of the sample size n .

3.1 Estimating the true order of variables

For a known ordering among the variables of the graph (Shojaie & Michailidis, 2010), an unrealistic assumption in many applications, the DAG inference problem can be cast in a convex optimization problem. To provide oracle inequalities of the proposed estimator in the most general case of an unknown order we consider here, we first focus on the problem of estimating the true variable order. Let us denote by Π_0 the set of permutation matrices compatible with the true DAG \mathcal{G}_0 :

$$\Pi_0 = \{P \in \mathbb{P}_p(\mathbb{R}), P^T G_0 P \in \mathbb{T}_p(\mathbb{R})\}.$$

Π_0 contains one or more permutation matrice(s) (see Example 1). We will have to make a decision as to whether the estimated order of variables \hat{P} given by Equation (6) is in Π_0 or not.

To answer this question, we investigate the effect of learning an erroneous order of variables $P \notin \Pi_0$. We introduce the following additional notations: for any permutation

matrix $P \in \mathbb{P}_p(\mathbb{R})$, we denote by $G_0(P)$ the matrix defined as:

$$G_0(P) = PT_0P^T,$$

with $T_0 = P_0^T G_0 P_0$ a lower triangular decomposition of G_0 . From a graphical point of view, while $P \notin \Pi_0$, the graph $\mathcal{G}_0(P)$ associated to $G_0(P)$ is obtained from \mathcal{G}_0 by permuting some of its nodes (see Example 2), otherwise, if $P \in \Pi_0$, $\mathcal{G}_0(P) = \mathcal{G}_0$. We also denote by $\varepsilon(P) := X - XG_0(P)$ the associated residual term. We denote by $\Omega(P)$ the covariance matrix of $\varepsilon(P)$ and $\omega_j(P) := \text{Var}(\varepsilon^j(P))$ the associated noise variances of each node.

With these notations and checking that the assumptions presented in Section 3.2 hold, we ensure that, with large probability, we choose a right order of variables and the estimated graph converges to the true graph when n and p grow to infinity (see Section 3.3).

Example 2 *Let*

$$P = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix} \notin \Pi_0$$

a wrong permutation.

In Figure 2, we represent the permuted graph $\mathcal{G}_0(P)$ (right) associated to the graph \mathcal{G}_0 (left). The latter is obtained from \mathcal{G}_0 after permutation of its nodes using PP_0^T , where P_0 (corresponding to the matrix P in Example 1) defines a right order of variables.




Figure 2: The graph \mathcal{G}_0 (top) and the permuted graph $\mathcal{G}_0(P)$ (bottom) associated to the permutation P .

3.2 Assumptions on the model

For a square matrix $M \in \mathcal{M}_{p \times p}(\mathbb{R})$ and a subset \mathcal{S} of $\llbracket 1, p \rrbracket^2$, we denote by $M_{\mathcal{S}} \in \mathcal{M}_{p \times p}(\mathbb{R})$ the matrix that has the same elements as M on \mathcal{S} and zero on the complementary set \mathcal{S}^C of \mathcal{S} . We now introduce the assumptions we used to obtain statistical properties of our estimator.

Hypotheses

H₁ There exists $\sigma^2 > 0$ such that

$$\forall j \in \llbracket 1, p \rrbracket, \text{Var}(\varepsilon^j) = \sigma^2.$$

H₂ There exists σ_0^2 , independent of p and n , such that

$$\max_{1 \leq j \leq p} \text{Var}(X^j) \leq \sigma_0^2.$$

H₃ There exists $\lambda^* > 0$ such that the minimal eigenvalue of the covariance matrix Σ of X satisfies

$$\lambda_{\min} \geq \lambda^* > 0.$$

H₄ There exists $g_{\max} < \infty$ such that the maximal weight of the DAG \mathcal{G}_0 is bounded

$$\max_{1 \leq i, j \leq p} |(G_0)_i^j| \leq g_{\max}.$$

H₅ The number of nodes p satisfies

$$p \log p = \mathcal{O}(n).$$

H₆ There exists $\kappa(s) > 0$ with $1 \leq s^2 \leq p^2$ such that:

$$\min \left\{ \frac{\|XM\|_F}{\sqrt{n} \|M_{\mathcal{S}}\|_F} \right\} \geq \kappa(s),$$

where the minimum is taken over the set of $p \times p$ matrices satisfying $\|M_{\mathcal{S}^c}\|_1 \leq 3 \|M_{\mathcal{S}}\|_1$, with $\mathcal{S} \subset \llbracket 1, p \rrbracket^2$ and $|\mathcal{S}| \leq s$.

H₇ There exists $0 < \eta \leq C \frac{n}{p \log p} \times \frac{1}{s}$ such that, for all permutations $P \notin \Pi_0$,

$$\frac{1}{p} \sum_{j=1}^p (|\omega_j(P)|^2 - 1)^2 > \frac{1}{\eta}.$$

Assumption **H₁** states that the noise variances are the same among all variables. This assumption is clearly hard to test in practice but makes the problem identifiable and ensures that we can recover the true DAG. Otherwise, minimizing (5) only leads to the identification of one element of the Markov equivalence class of the true DAG (partially directed graph). To simplify the theoretical results and proofs, until the end of this work, we assume that the noise variances σ^2 are equal to one. Our results are still valid even if $\sigma^2 \neq 1$, by small modifications in the constant terms as long as they are all equal.

Assumption **H₅** deserves a special attention since it bounds the high dimensional setting. The considered problem is obviously non-trivial and requires a sufficient amount of information. A more detailed discussion about assumptions **H₃** and **H₅** is proposed in Section 3.4.

Assumption **H₆** is a natural extension of the Restricted Eigenvalue condition of Bickel et al. (2009) to our multi-task setting. More precisely, denoting

$$\tilde{X} = \begin{pmatrix} X & 0 \\ 0 & X \end{pmatrix} \begin{matrix} \uparrow \\ n \times p, \\ \downarrow \end{matrix}$$

$\xleftrightarrow{p^2}$

\mathbf{H}_6 is equivalent to assuming that the Gram matrix $\frac{\tilde{X}\tilde{X}^T}{n}$ is non-degenerate on a restricted cone (Lounici et al., 2009; Bühlmann & van de Geer, 2011). Notice that this condition is very classical in the literature. It yields good practical performance even for small sample sizes, and some recent works discuss an accurate population eigenvalue estimation even in a large dimension setting (Mestre, 2008; El Karoui, 2008; Liu et al., 2014; Ledoit & Wolf, 2015).

The last assumption \mathbf{H}_7 is an identifiability condition needed to ensure that the estimated permutation \hat{P} is in Π_0 . This assumption was introduced by van de Geer & Bühlmann (2013) as the “omega-min” condition. In a sense, it separates the set of compatible permutations from its complement in a finite sample scenario.

3.3 Main result

The result we establish in this section is double-edged: (a) with large probability, we ensure that the estimated \hat{P} belongs to Π_0 , and (b) we provide oracle inequalities both in prediction and estimation for the graph estimated from the minimisation problem (6). This result clearly states the desirable theoretical properties of the derived estimator, assuming reasonable conditions on the complex system embedding the data.

Theorem 1 *Assume that \mathbf{H}_{1-7} are satisfied, with $s^2 \subset \llbracket 1, p^2 \rrbracket$ in \mathbf{H}_6 such that $\sum_{i,j} \mathbf{1}_{(G_0)_i^j \neq 0} \leq s^2$ (G_0 is s^2 -sparse). Let $\lambda = 2C\sqrt{s\frac{\log p}{n}}$. Then, with probability greater than $1 - 5/p$, any solution $\hat{G} = \hat{P}\hat{T}\hat{P}^T$ of the minimization problem (6) satisfies that $\hat{P} \in \Pi_0$. Moreover, with at least the same probability, the following inequalities hold:*

$$\frac{1}{n} \left\| X\hat{G} - XG_0 \right\|_F^2 \leq \frac{16C^2}{\kappa^2(s)} s^3 \frac{\log p}{n}. \quad (7)$$

$$\left\| \hat{G} - G_0 \right\|_1 \leq \frac{16C}{\kappa^2(s)} s^{5/2} \sqrt{\frac{\log p}{n}}. \quad (8)$$

The proof of this result is deferred in Section C of the Supplementary Materials.

Theorem 1 states that with probability at least $1 - 5/p$, we choose a compatible order of variables over the set of permutations. Inequalities (7) and (8) give non-asymptotic upper bounds on the loss under conditions depending on s , p and n (see Section 3.4). They also ensure that the estimated \hat{T} is close to the true T_0 with large probability.

3.4 Discussion on the high-dimensional scenario

Sparsity of the graph Assumption **H₇** and Theorem 1 naturally require a trade-off between signal sparsity, dimensionality and sample size. In the ultra sparse regime (where the sparsity s^2 of the true graph is bounded by $s^* > 0$), Theorem 1 provides oracle inequalities for \hat{G} choosing $\eta \leq \frac{\alpha}{\sqrt{s^*}}$ with $p \log(p) = \alpha n$ in Assumption **H₇**.

In the standard sparsity scenario, if s^2 is at least of the order of p , then η should be of the order of α/\sqrt{p} , which is unrealistic as $p \rightarrow +\infty$. This case thus requires a stronger dimensional assumption **H₅**. Taking at least $p^2 \log(p) = \mathcal{O}(n)$ ensures a good estimation of the graph.

Note however that universal conditions cannot be overcome and the ultra-high dimension settings (e.g. Wainwright (2009); Verzelen (2012)) is an insurmountable limit.

Minimal eigenvalue condition Assumption **H₃** ensures that the minimal eigenvalue of the covariance matrix Σ of X is not too small. In the high-dimensional scenario, this could be hard to verify, λ_{min} decreasing while n, p growing to infinity (Hogben, 2007). A natural bound for λ_{min} is:

$$\lambda_{min} \geq \frac{1}{p \max(1, g_{max}^2) (1 + s)}, \quad (9)$$

with g_{max} and s as in **H₄** and Theorem 1.

Assumption **H₃** can thus be relaxed by allowing λ_{min} to decrease with $1/ps$. The price to pay for this relaxation is a data dimensionality reduction $p^3 \log(p) = \mathcal{O}(n)$, which automatically implies:

$$\frac{3\lambda_{min}}{4} - 2\sqrt{\frac{\log(p)}{n}} - 3\sigma_0\sqrt{\frac{2p \log(p)}{n}} > 0,$$

with Equation (9) (for more details, see Section A of the Supplementary Material *Proof details*).

4 Inference algorithm

4.1 Global algorithm overview

In this section, we present GADAG (*Genetic Algorithm for learning DAG*) a computational procedure devoted to solve Equation (6). Although decomposing the original problem made it more natural to handle, this problem is still a very challenging task from an optimization point of view, due to the different nature of the variables P and T , the non-convexity of the cost function and the high dimension of the search space.

An intuitive approach consists in using an alternating scheme: one of the variables P or T is fixed and the other one is sought so as to optimize the score function, then the roles of P and T are reversed and the procedure is repeated iteratively until convergence for some criterion (Csiszár & Tusnády, 1984). However, the structure of our problem does

not allow us to use such a scheme: looking for an optimal T given a fixed P makes sense, but changing P for a fixed T does not.

In our inference algorithm GADAG, an outer loop is used to perform the global search among the DAGs space, which is driven by the choice of P , while a nested loop is used to find an optimal T for each given fixed P (see Figure 3). As we show in the following, population-based metaheuristics algorithms are a natural and efficient choice for exploring the space of permutation matrices (Section 4.3). The nested optimization problem can be resolved using a steepest descent approach (Section 4.2).

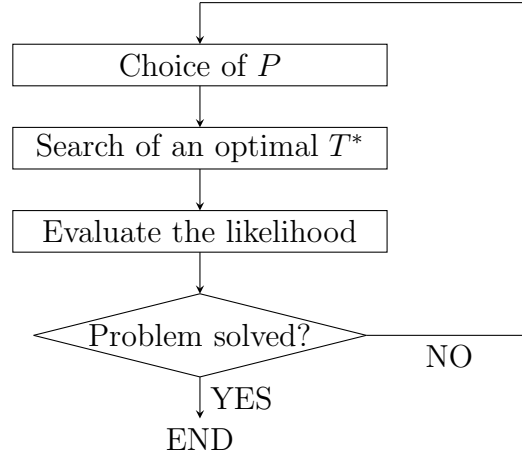


Figure 3: Overview of our hybrid algorithm GADAG.

4.2 Graph structure learning when the variable order is fixed

Assume first that the variable ordering $P \in \mathbb{P}_p(\mathbb{R})$ is fixed. The problem of inferring a graph is then reduced to estimating the graph structure, which can be solved by finding a solution of:

$$\min_{T \in \mathbb{T}_p(\mathbb{R})} \left\{ \frac{1}{n} \|X(I - PTP^T)\|_F^2 + \lambda \|T\|_1 \right\}. \quad (10)$$

The minimization problem given by Equation (10) looks like a well-studied problem in machine learning, as it is closely related to the ℓ_1 -constrained quadratic program, known as the Lasso in the statistics literature (Tibshirani, 1996). Indeed, the ℓ_1 -regularization leads to variable selection and convex constraints that make the optimization problem solvable. We note here that this allows us to always provide a locally optimal solution, i.e optimal weight estimates given a hierarchy between the nodes.

A large number of efficient algorithms are available for computing the entire path of solutions as λ is varied, e.g. the LARS algorithm of Efron et al. (2004) and its alternatives. For example, in the context of the estimation of sparse undirected graphical models, Meinshausen & Bühlmann (2006) fit a Lasso model to each variable, using the others as predictors, and define some rules for model symmetrization as they do not work on DAGs.

The graphical Lasso (or glasso, Friedman et al. 2007) algorithm directly relies on the estimation of the inverse of a structure covariance matrix assumed to be sparse. Improvements were proposed for example by Duchi et al. (2008) (improved stopping criterion) and Witten et al. (2011) (estimation of a block-diagonal matrix). Other authors propose to solve the optimization problem using an adaptation of classical optimization methods, such as interior point (Yuan & Lin, 2007) or block coordinate descent methods (Banerjee et al., 2008; Friedman et al., 2007).

We propose here an original convex optimization algorithm to find the solution in Equation (10) in a form similar to a steepest descent algorithm. Our proposed algorithm is much quicker than a glasso approach, a desirable feature as it will run at each iteration of the global algorithm (see the “Search of an optimal T^* ” box in Figure 3 and the “Evaluate the new individuals” item in Algorithm 2). Moreover, its mechanistic components (see Section B of the Supplementary Material *Proof details*) allowed us to derive the theoretical results of Theorem 1. The proposed scheme can be seen as an adaptation of the LARS algorithm with matrix arguments.

Let $(T_k)_{k \geq 0}$ the sequence of matrices defined for all $i, j \in \llbracket 1, p \rrbracket^2$ as:

$$(T_{k+1})_i^j = \text{sign}((U_k)_i^j) \max\left(0, |(U_k)_i^j| - \frac{\lambda}{L}\right), \quad (11)$$

where for all $k \geq 0$, $U_k = T_k - \frac{\nabla\left(\frac{1}{n}\|X(I - PT_k P^T)\|_F^2\right)}{L}$, L is the Lipschitz constant of the gradient function $\nabla\left(\frac{1}{n}\|X(I - PT_k P^T)\|_F^2\right)$ and $\text{sign}()$ is the sign of any element. Then, a solution of (10) is given by performing Algorithm 1, where:

- the projection $\text{Proj}_{\mathbb{T}_p(\mathbb{R})}(T)$ of any $p \times p$ real-valued matrix $T = ((T_k)_i^j)_{i,j}$ on the set $\mathbb{T}_p(\mathbb{R})$ is given by

$$\left(\text{Proj}_{\mathbb{T}_p(\mathbb{R})}(T_k)\right)_i^j = \begin{cases} 0 & \text{if } i < j, \\ (T_k)_i^j & \text{otherwise.} \end{cases} \quad (12)$$

- the gradient of $\frac{1}{n}\|X(I - PT_k P^T)\|_F^2$ is

$$\nabla\left(\frac{1}{n}\|X(I - PT_k P^T)\|_F^2\right) = -\frac{2}{n}(XP)^T(X - XPT_k P^T)P. \quad (13)$$

The detailed calculations are deferred to Section B of the Supplementary Material *Proof details*.

4.3 A Genetic Algorithm for a global exploration of the permutation matrices space incorporating network topologies

As the optimal T can be calculated for any P using Algorithm 1 and with a very good approximation accuracy according to Theorem 1, the optimization task (6) comes down to exploring the $\mathbb{P}_p(\mathbb{R})$ space of permutation matrices in dimension p and to evaluating the

Algorithm 1: Graph structure learning - minimization of (10)

Input: $\lambda, L, \epsilon > 0$.

Initialization: T_0 the null squared $p \times p$ matrix, $k = 0$ and $e = +\infty$.

while $e > \epsilon$ **do**

 Compute $U_k = T_k - \frac{\nabla\left(\frac{1}{n}\|X(I-PT_kP^T)\|_F^2\right)}{L}$ with Equation (13);
 Using Equation (11), compute the current matrix $T_{k+1} = ((T_{k+1})_i^j)_{i,j}$;
 Project T_{k+1} on $\mathbb{T}_p(\mathbb{R})$ with Equation (12): $T_{k+1} \leftarrow \text{Proj}_{\mathbb{T}_p(\mathbb{R})}(T_{k+1})$;
 Compute $e = \|T_{k+1} - T_k\|_F$;
 Increase k : $k \leftarrow k + 1$;

end

Output: $T_k \in \mathbb{T}_p(\mathbb{R})$ the unique solution of (10).

quality of permutation candidates $P \in \mathbb{P}_p(\mathbb{R})$. We first note that the number of permutation matrices is $p!$, which rules out any exact enumeration method, even for relatively small p . We propose instead to use a meta-heuristic approach, which has proven to be successful for many discrete optimization problems like wire-routing, transportation problems or traveling salesman problem (Michalewicz, 1994; Dréo et al., 2006).

Among the different meta-heuristics (Simulated annealing, Tabu search, Ant Colony,...) we focused on Genetic Algorithms (GA) because, despite limited convergence results (Cerf, 1998; Michalewicz, 1994), they were found much more efficient in problems related to ours than alternatives with more established convergence proofs (e.g. Granville et al. (1994) for simulated annealing), while allowing the use of parallel computation.

GAs mimic the process of natural evolution, and use a vocabulary derived from natural genetics: populations (a set of potential solutions of the optimization problem), individuals (a particular solution) and genes (the components of a potential solution). Each generation/iteration of the algorithm will improve the constituting elements of the population. In short, a population made of N potential solutions of the optimization problem samples the search space. This population is sequentially modified, with the aim of achieving a balance between exploiting the best solutions and exploring the search space, until some termination condition is met.

We use here a classical Genetic Algorithm (GA), as described in Michalewicz (1994) for instance, which is based on three main operators at each iteration: selection, crossover and mutation. The population is reduced by selection; selection shrinks the population diversity based on the individual fitness values. The crossover allows the mixing of good properties of the population to create new composite individuals. Mutations change one (or a few in more general GAs) components of the individuals to allow random space exploration. The complete sketch of algorithm GADAG is given in Algorithm 2. A discussion on parameters to set in Algorithm 2 is found in Section 5.1. The details of the different operators are given in the following.

As we show in Example 3, any $P \in \mathbb{P}_p(\mathbb{R})$ is uniquely defined by a permutation vector of $\llbracket 1, p \rrbracket$. Hence, we use as a the search space \mathfrak{S}_p the set of permutations of $\llbracket 1, p \rrbracket$, which is

a well-suited formulation for GAs.

Example 3 Consider the permutation matrix ($p = 5$):

$$P = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

Then, P is represented by the $\boxed{5}\boxed{3}\boxed{4}\boxed{1}\boxed{2}$ vector, looking at the ranks of non-null values of P column by column. The nodes are ranked according to their topological ordering.

Note that our problem closely resembles the classical Traveling Salesman Problem (TSP), which has been successfully addressed by means of genetic algorithms (Grefenstette et al., 1985; Davis, 1991). Identically to the TSP, we optimize over the space of permutations, which induces specific constraints for defining the crossover and mutation operators. However, unlike the TSP, the problem is not circular (in the TSP, the last city is connected to the first one), and the permutation here defines a hierarchy between nodes rather than a path, which makes the use of TSP-designed operators a potentially poor solution. As we show in the following, we carefully chose these operators in order to respect the nature of the problem at hand. In particular, we emphasize two of their desirable features: their efficiency in exploring the search space and the interpretable aspect they offer in terms of modifications on a given network or the blend of two different networks (crossover).

Fitness function Given a potential solution $p_i \in \mathfrak{S}_p$, the fitness function is defined as:

$$J_i = J(p_i) = \frac{1}{n} \|X(I - P_i T_i^* P_i^T)\|_F^2 + \lambda \|T_i^*\|_1, \quad (14)$$

with P_i constructed from p_i as in Example 3 and T_i^* the solution of Equation (10) with $P = P_i$. As mentioned earlier, at each step of the proposed GA, the evaluation of the fitness function thus requires running the nested loop of our global algorithm GADAG.

Selection operator The selection operator (or survival step) consists in generating a population of N individuals from the N existing individuals by random sampling (with replacement, hence some individuals are duplicated and others are deleted). It aims at improving the average quality of the population by giving to the best potential solutions a higher probability to be copied in the intermediate population. We have chosen to use the classical proportional selection of Holland (1975): each individual is selected with a probability inversely proportional to its fitness value of Equation (14).

Crossover operator A crossover operator generates a new set of potential solutions (children) from existing solutions (parents). Crossover aims at achieving at the same time (i) a good exploration of the search space by mixing the characteristics of the parents to create potentially new ones while (ii) preserving some of the desirable characteristics of the parents. By desirable features, we mean features of the network which lead to good fitness values, and which in turn are favored by selection over the generations. The crossover population (set of parents) is obtained by selecting each individual of the population with a probability p_{xo} ; the parents are then paired randomly.

We have chosen the *order-based* crossover, originally proposed for the TSP (Michalewicz, 1994, Chapter 10), which is defined as follows. Given two parents p_1 and p_2 , a random set of crossover points are selected, which we denote Ω . It consists in a permutation of k elements taken from $\llbracket 1, p \rrbracket$, with k uniformly drawn between 0 and p . A first child C_1 between p_1 and p_2 is then generated by:

1. fixing the crossover points of p_1 ,
2. completing C_1 with the missing numbers in the order they appear in p_2 .

Example 4 Consider the two following parents:

| | | | | | | | | | | |
|-------|----------|---|----------|----------|----|----------|----------|----------|---|----------|
| p_1 | 4 | 3 | 10 | 7 | 5 | 9 | 1 | 2 | 6 | 8 |
| p_2 | 6 | 1 | 9 | 4 | 10 | 2 | 8 | 3 | 7 | 5 |

Assume that the crossover points randomly chosen are 4, 9, 2 and 8 (in bold red above). Then, the child C_1 is defined by inheriting those points from p_1 and filling the other points in the order they appear in p_2 :

| | | | | | | | | | | |
|-------|----------|---|--------------|--------------|----|--------------|--------------|----------|---|----------|
| p_2 | 6 | 1 | 9 | 4 | 10 | 2 | 8 | 3 | 7 | 5 |
| C_1 | 4 | * | * | * | * | 9 | * | 2 | * | 8 |
| ⇓ | | | | | | | | | | |
| | 4 | 6 | 1 | 10 | 3 | 9 | 7 | 2 | 5 | 8 |

From a graphical point of view, a crossover between p_1 and p_2 , which encode two complete graphs \mathcal{G}_{P_1} and \mathcal{G}_{P_2} , constructs two new graphs. One of them, \mathcal{G}_{C_1} is composed of the sub-graph of \mathcal{G}_{P_1} induced by the set of crossover points Ω and the sub-graph of \mathcal{G}_{P_1} induced by the complementary set Ω^C of Ω in $\llbracket 1, p \rrbracket$ (see Figure 4). The second child graph \mathcal{G}_{C_2} is obtained in an identical manner by reversing the roles played by the two parental graphs.




Figure 4: Graphical representation of crossover between two 10-node graphs. The two parental graphs are represented on the left. The third graph, on the right, is obtained by combining the blue and red part of its parents using the crossover operator.

Mutation operator Mutation operators usually correspond to the smallest possible change in an individual (unary operator). We thus define it as an alteration of two neighbouring genes (see Example 5). Graphically, a mutation consists in switching the arrowhead of an edge between two nodes. Mutation is applied to each child with probability p_m .

Example 5 A possible mutation for the first child of Example 4 is to swap the genes “1” and “10” (in bold red below):

$$M_1 \quad \begin{array}{|c|c|c|c|c|c|c|c|c|} \hline 4 & 6 & \mathbf{1} & \mathbf{10} & 3 & 9 & 7 & 2 & 5 & 8 \\ \hline \end{array}$$

$\uparrow \quad \uparrow$

Stopping criterion Two quantities are monitored along the iterations: the heterogeneity of the population and the value of the objective function.

For the first indicator, we use the Shannon entropy, defined for each rank position $j \in \llbracket 1, p \rrbracket$ as:

$$H_j = - \sum_{i=1}^p \frac{N_{i,j}}{N} \log \left(\frac{N_{i,j}}{N} \right),$$

where $N_{i,j}$ is the number of times when i appears in position j . $H_j = 0$ if all the individuals “agree” on the position of a node and the population is perfectly homogeneous at this node. On the contrary, it is maximum when we observe a uniform distribution of the different nodes at a given position and the population is in this case at a maximum of heterogeneity or disorder for this position. The algorithm stops if the population entropy value $H = \sum_{j=1}^p H_j$ drops below a threshold since $H = 0$ if all the individuals are identical. A second criterion can terminate GADAG if difference in the average fitness (denoted \bar{J} thereafter) of the population between a given number of consecutive iterations, does not change by more than a predefined threshold.

Algorithm 2: GADAG overview

Input: $p_{xo}, p_m, \epsilon_H > 0, \epsilon_J > 0, k_{max} > 0, i_{max} > 0, \lambda, L$.

Initialization: Generate the initial population \mathcal{P}_0 with N permutations of $\llbracket 1, p \rrbracket$, $k = 0$ and $e_J = +\infty$.

while $H > \epsilon_H$ & $e_J > \epsilon_J$ & $k \leq k_{max}$ **do**

 Generate \mathcal{P}_{k+1} as a random **selection** of N individuals from \mathcal{P}_k ;

 Pick an even subset \mathcal{P}_{xo} of \mathcal{P}_{k+1} (each individual of \mathcal{P}_{k+1} selected with probability p_{xo});

 Perform **crossover** on \mathcal{P}_{xo} by randomly pairing the individuals;

Mutate each obtained individual with probability p_m ;

Evaluate the new individuals \mathcal{P}_m by running Algorithm 1;

 Replace \mathcal{P}_{xo} by \mathcal{P}_m in \mathcal{P}_{k+1} ;

 Compute the Shannon entropy H and the difference in the average fitness

$e_J = \max_{0 \leq i \leq i_{max}} (\bar{J}(\mathcal{P}_{k+1}) - \bar{J}(\mathcal{P}_{k-i}))$;

 Increase k: $k \leftarrow k + 1$;

end

5 Numerical experiments

This section is dedicated to experimental studies to assess practical performances of our method. We demonstrate the ability of our algorithm to analyse data sets, which mimic the activity of a complex biological system, and we compare it to other state-of-the art methods. In Section 5.1, we present the calibration of the Genetic Algorithm parameters. The competing methods are presented in Section 5.2 and Section 5.3 introduces the measures we used to assess the merits of the different methods. Experimental results are then detailed in Section 5.4.

5.1 Algorithm parameters

Running the procedure of Algorithm 2 requires to define parameters of the outer loop, which generates our population of P 's, and of the nested loop to find the optimal T^* . The evaluation of the Lipschitz gradient constant L , used to find the optimal graph structure T^* , is known as a hard established problem in optimization. Some authors propose to choose an estimate of L from a set of possible values (Jones et al., 1993; Sergeyev & Kvasov, 2006), to estimate local Lipschitz constants (Sergeyev, 1995), or to set it a priori to a fixed value (Evtushenko et al., 2009; Horst & Pardalos, 1995). Here, observing Equation (13), a major bound for L is given by:

$$L \leq \frac{2}{n} \|X^T X\|_F.$$

We found that setting L to this bound worked well in practice in all our scenarios.

Five parameters need to be tuned to run the Genetic Algorithm: the crossover rate p_{xo} , the mutation rate p_m , the constant of the stopping criteria ϵ_H and ϵ_J and the size of the population N . For the first four parameters, we observed that their value had a limited

effect on the efficiency, hence we chose commonly used values in the literature (see Table 1). The size of the population has a more complex effect and has been investigated in several prospective papers (e.g. Schaffer et al. 1989; Alander 1992; Piszcz & Soule 2006; Ridge 2007) but without providing a definitive answer to the problem. In our simulation study, we chose as a rule-of-thumb $N = 5p$, which was found as a good compromise between computational cost and space exploration on several experiments.

The complete parameter settings used in our experiments are reported in Table 1.

Table 1: Algorithm parameter settings

| Parameter | Value |
|-------------------|---------------------------|
| p_{xo} | 0.25 |
| p_m | 0.5 |
| N | $5 \times p$ |
| L | $\frac{2}{n} \ X^T X\ _F$ |
| max. nb. of eval. | $5 \times p$ |
| ϵ_H | 10^{-6} |
| ϵ_J | 10^{-4} |

5.2 Comparison to state-of-the art

We compare GADAG to five state-of-the-art inference methods. Among them, the Genie3 method (Huynh-Thu et al., 2010), based on random forests, was the best performer of one of the DREAM4 sub-challenges, while the BootLasso (Allouche et al., 2013) was one of the key components of the best performing approach of one of the DREAM5 sub-challenges (Allouche et al., 2013). The two methods decompose the prediction of the network into p feature selection sub-problems. In each of the p sub-problems, one of the node is predicted from the other ones using random forests (Breiman, 2001) for Genie3 or a bootstrapped version of the Lasso (Bach, 2008) for BootLasso. For the random forest approach, parents of each node were detected as most significant explanatory variables according to a variance reduction criterion in a regression tree framework. The process was repeated on a randomized set of trees, which made up the so-called random forest. This method allowed us to derive a ranking of the importance of all variables for the target by averaging the scores over all the trees. We used the R package `randomforest` (Liaw & Wiener, 2002) for our results. The Lasso is a ℓ_1 -norm penalization technique for solving linear regression. Following the works of Bach (2008), BootLasso uses bootstrapped estimates of the active regression set based on a Lasso penalty: only those variables that are selected in every bootstrap are kept in the model. In both cases, actual coefficient values are estimated from a straightforward least square procedure. Note that we slightly relax the condition for a variable to be included in the model, a variable was selected at a given penalty level if more than 80% of bootstrapped samples led to selecting it in the model (Allouche et al., 2013).

We also compare our algorithm to three classical methods for Bayesian Networks (BNs)

modelling. BNs are graphical models (Pearl, 2009) defined by a DAG and parameters that set quantitative relationships between variables. Algorithms devoted to structure and parameter learning in BNs either aim at maximising a score that reflects the fit of the data to the learnt structure, or test for independencies between variables. They are often used as references in a gene regulatory network inference context (Tsamardinos et al., 2006), although mainly for moderate size networks. The first compared algorithm we used is the PC-algorithm (Spirtes et al., 2000), a popular constraint-based method that drastically reduces the number of conditional independence tests. It first builds the skeleton of the graph by removing edges from a complete undirected graph before determining the orientation of the edges, when possible. A large number of implementations of the PC-algorithm exists. The numerical results presented here were obtained using the `pcAlg` function of the R-package `pcalg`, based on standard correlation estimates for conditional independence testing. We also ran ARACNE (Margolin et al., 2006), an improved version of minimum-weight spanning tree that uses the information inequality to eliminate the majority of indirect relationships between variables. We used the ARACNE function of the R-package `bnlearn`. We finally compare GADAG to the Greedy Equivalence Search (GSE) algorithm (Chickering, 2002), implemented in the R-package `pcalg`, which heuristically searches in the space of equivalent classes the model with the highest Bayesian score.

5.3 Performance metrics

A classical performance measure for graph inference methods consists in comparing predicted interactions with the known edges in the true graph \mathcal{G}_0 using precision versus recall (P/R) curves. We denote by TP, FP, FN and TN, the true positive (correctly predicted) edges, the false positive (inferred by mistake) edges, the false negative (missed) edges, and the true negative (correctly non-predicted) edges. The recall, defined as $\frac{TP}{TP+FN}$, measures the power (or sensitivity) of reconstruction of non-zero elements of the true matrix G (or equivalently of the true network) for one method, whereas the precision, equal to $\frac{TP}{TP+FP}$, measures the accuracy of the reconstruction. The closer to one the precision and the recall the better.

P/R curves represent the evolution of those quantities when varying the sparsity of the methods. Random forests produce a ranked list of regulatory interactions, which corresponds to the edges of the inferred graph. Edges are then successively introduced with decreasing confidence scores to produce the random forest P/R curve. GADAG and Boot-Lasso are based on penalized optimization: both seek linear dependencies between the variables with a controlled level of parsimony (λ in Equation (5) for GADAG). For λ varying from 0 (complete graph) to $+\infty$ (empty graph), each of these methods produce a list of edges, successively introduced in the model. These lists of edges define the precision versus recall curves for these two approaches. Note that the implementation we used for running ARACNE was only able to produce a final network prediction (interaction ranking is not available). As a summary performance measurement, we also computed the classical area under the P/R curve (AUPR).

5.4 DREAM data analysis

The datasets we used mimic activations and regulations that occur in gene regulatory networks. It is provided by the DREAM4 challenge on “In Silico Network Challenge”. Note that although plausibly simulated, DREAM4 data sets are not real biological data sets. However, the used network structures (five in total) were extracted from *E. coli* and *S. cerevisiae* -two biological model organisms- transcriptional networks. These networks contain cycles, but self-loops were discarded. The gene expression observations were not simulated by an equal noise Gaussian multivariate model, stochastic differential equations were used to mimic the kinetic laws of intricate and intertwined gene regulations. In addition to the biological noise simulated from the stochastic differential equations, technical noises were added to reproduce actual gene measurement noise. All data sets were generated by the GNW software (Marbach et al., 2009).

Working with simulated networks, we are able to quantitatively and objectively assess the merit of competing methods in terms of true predictions (true positive TP and true negative TN) vs. incorrect predictions (false positive FP and false negative FN) edges. While the analysis of a real data set is certainly the final goal of a methodology motivated by a real problem like ours, there are only imprecise ways of validating a method when analysing a real data set. Well known systems are often small and even if knowledge has accumulated on them, these can be noisy and difficult to gather to obtain a fair picture of what can adequately be considered as sets of true positive and true negative sets of edges. Even if the data generation process of the DREAM4 In Silico Network Challenge is completely understood, no existing method is able to predict all regulatory relationships, but at the price of including many false positive predictions. The DREAM4 datasets we considered have $p = 100$ nodes and only $n = 100$ observations making it a very challenging task.

5.5 Numerical results

The P/R curves for the five DREAM problems are shown in Figure 5. Each curve corresponds to one of the five networks used in the challenge. In general, for all the problems the five methods are able to achieve a precision equal to one (that is, to include only true edges), but these correspond to overly sparse graphs (very small recall). Conversely, a recall equal to 1 can only be reached by adding a large number of FP edges, whatever the method we consider, even if some fail earlier than others. The main differences between the methods appear on the leftmost part of the P/R curves, especially those of Figure 5 B, C and D: while the precision of BootLasso, Genie3, PCalg and GSE drops rapidly with a slow increases in recall above 20% recall, it remains higher for GADAG. Hence, its first predicted edges are at least as accurate than those of the four other methods and it produces a larger set of reliable edges. For graphs of lesser sparsity, none of the five methods is really able to identify clearly reliable edges. Large number of FP edges are produced to achieve a recall higher than 60%.

For Networks 1 and to a lesser extent 5 (Figure 5 A and E), GADAG recovers with more difficulty the first true edges than other methods, with a high level of FP edges at the beginning of the curve (low precision and low recall). However, as soon as the recall

exceeds the 10%, *resp.* 15%, for graph A, *resp.* for graph E, GADAG performance is again superior to that other methods.

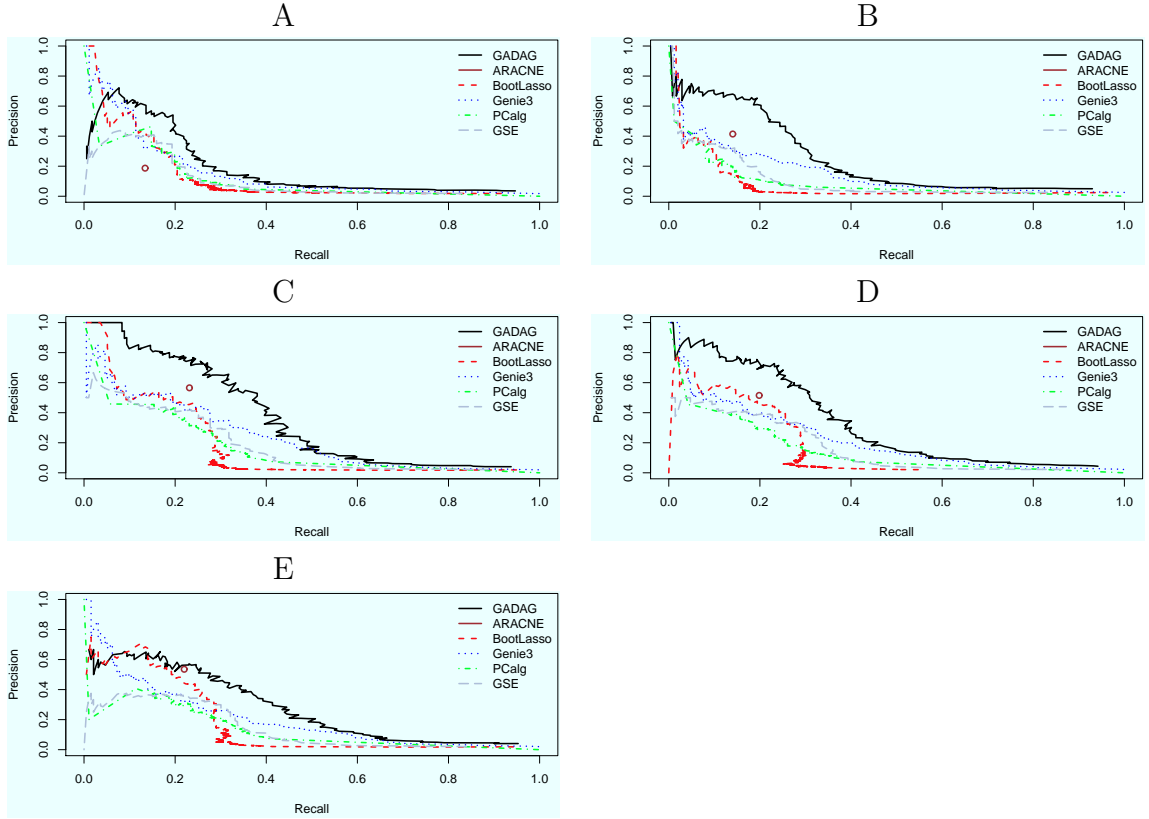


Figure 5: P/R curves for the five Dream networks and the five compared methods.

Table 2 gives the areas under the P/R curves for all methods and networks. For this indicator, GA significantly outperforms the state-of-the-art methods for all networks.

| Method | Net 1 | Net 2 | Net 3 | Net 4 | Net 5 |
|-----------|--------------|--------------|--------------|--------------|--------------|
| GADAG | 0.182 | 0.236 | 0.348 | 0.317 | 0.267 |
| Genie3 | 0.154 | 0.155 | 0.231 | 0.208 | 0.197 |
| BootLasso | 0.118 | 0.061 | 0.171 | 0.147 | 0.169 |
| PCalg | 0.116 | 0.089 | 0.171 | 0.149 | 0.130 |
| GSE | 0.101 | 0.089 | 0.170 | 0.153 | 0.133 |

Table 2: Area under the Precision vs. Recall curve for all networks and methods (except ARACNE).

6 Conclusion and discussion

In this paper, we proposed a hybrid genetic/convex algorithm for inferring large graphs based on a particular decomposition of the ℓ_1 -penalized maximum likelihood criterion. We obtained two oracle inequalities that ensure that the graph estimator converges to the true graph under assumptions that mainly control the model structure: graph size (balance between sparsity, number of nodes and maximal degree) and signal-to-noise ratio. From an algorithmic point of view, the estimation task is split into two subproblems: node ordering estimation and graph structure learning. The first one is a non-trivial problem since we optimize over a discrete non-convex large dimensional set. It led us to use a heuristic approach we specifically tailored to achieve the optimization task. The second one is a more common problem, related to the Lasso one, for which we proposed a sound procedure with theoretical guarantees. The potential of such an approach clearly appeared in the numerical experiments, for which the behavior of our algorithm seemed to be very competitive when compared to the state-of-the-art.

Nevertheless, we see many opportunities for further improvements. First, convergence proof for the algorithm, although a challenging task, is worth investigating, for instance using the works of Cerf (1998) on genetic algorithms. An alternative would be to consider other optimization schemes for the node ordering with more established convergence proofs (e.g. simulated annealing (Granville et al., 1994)).

Second, other potential extensions involve algorithmic considerations in order to improve the calculation time, including a finer calibration of the algorithm parameters, an initialization step for the gradient descent, and, in general, improving the interactions between the nested and outer loops. Tackling very large datasets from several thousands of nodes may also require a particular treatment, for instance by adding local search operators to GADAG.

Finally, we would like to emphasize the graph identifiability problem: in our settings, we assume the noise variances of all graph nodes to be equal to ensure graph identifiability (that is no equivalence class of graphs). Such a hypothesis is of course restrictive and likely to be violated for real datasets. In order to infer networks for any noise variances, one solution consists in incorporating interventional data on the model. These data are obtained from perturbations of the biological system (e.g. gene knockouts or over-expressions) and make the equivalence class of graphs smaller (Hauser & Bühlmann, 2012). Then, the use of such new data could be combined with observational on the MLE estimator (as recently proposed by Hauser & Bühlmann (2015) for a BIC-score penalized MLE, or by Rau et al. (2013) for learning Gaussian Bayesian networks in the case of GRN inference) and a modification of our hybrid algorithm could lead to the identification of the true graph.

References

- ALANDER, J. (1992). On optimal population size of genetic algorithms. In *Proceedings of the IEEE Computer Systems and Software Engineering*.
- ALLOUCHE, A., CIERCO-AYROLLES, C., DE GIVRY, S., GUILLERMIN, G., MANGIN, B.,

- SCHIEX, T., VANDEL, J. & VIGNES, M. (2013). A panel of learning methods for the reconstruction of gene regulatory networks in a systems genetics. In *Gene Network Inference, Verification of Methods for Systems Genetics Data*. Springer, Heidelberg, alberto de la fuente ed., pp. 9–32.
- BACH, F. (2008). Bolasso: model consistent lasso estimation through the bootstrap. In *Proceedings of the Twenty-fifth International Conference on Machine Learning*.
- BANERJEE, O., EL GHAOU, L. & D’ASPREMONT, A. (2008). Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J Mach Learn Res* **9**, 485–516.
- BARABÁSI, A. & OLTVAI, Z. (2004). Network biology: understanding the cell’s functional organization. *Nat Rev Genet* **5**, 101–113.
- BICKEL, P. & LI, B. (2006). Regularization in statistics. *Test* **15**, 271–344.
- BICKEL, P. J., RITOV, Y. & TSYBAKOV, A. B. (2009). Simultaneous analysis of lasso and Dantzig selector. *Ann Stat* **37**, 1705–1732.
- BREIMAN, L. (2001). Random forests. *Mach Learn* **45**, 532.
- BÜHLMANN, P. (2013). Causal statistical inference in high dimensions. *Math Method Oper Res* **77**, 357–370.
- BÜHLMANN, P. & VAN DE GEER, S. (2011). *Statistics for high-dimensional data*. Springer Series in Statistics. Springer, Heidelberg. Methods, theory and applications.
- CERF, R. (1998). Asymptotic convergence of genetic algorithms. *Adv Appl Probab* **30**, 521–550.
- CHEN, J. & CHEN, Z. (2008). Extended bayesian information criteria for model selection with large model spaces. *Biometrika* , 759–771.
- CHICKERING, D. M. (1996). Learning Bayesian networks is NP-complete. In *Learning from data (Fort Lauderdale, FL, 1995)*, vol. 112 of *Lect. Notes Stat*. Springer, New York, pp. 121–130.
- CHICKERING, D. M. (2002). Optimal structure identification with greedy search. *J Mach Learn Res* **3**, 507–554.
- COOK, S. (1985). A taxonomy of problems with fast parallel algorithms. *Inform Control* **64**, 2–22.
- CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L. & STEIN, C. (2001). *Introduction to algorithms*. MIT Press, Cambridge, MA; McGraw-Hill Book Co., Boston, MA, 2nd ed.
- CSISZÁR, I. & TUSNÁDY, G. (1984). Information geometry and alternating minimization procedures. *Statist. Decisions* , 205–237Recent results in estimation theory and related topics.

- DAVIS, L., ed. (1991). *Handbook of Genetic Algorithms*. Van Nostrand Reinhold.
- DRÉO, J., PÉTROWSKI, A., SIARRY, P. & TAILLARD, E. A. (2006). *Metaheuristics for hard optimization*. Springer-Verlag, Berlin. Methods and case studies, Translated from the 2003 French original by Amitava Chatterjee.
- DUCHI, J., GOULD, S. & KOLLER, D. (2008). Projected subgradient methods for learning sparse gaussians. In *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*.
- EFRON, B., HASTIE, T., JOHNSTONE, I. & TIBSHIRANI, R. (2004). Least angle regression. *Ann Stat* **32**, 407–499.
- EL KAROUI, N. (2008). Spectrum estimation for large dimensional covariance matrices using random matrix theory. *Ann Stat*. **36**, 2757–2790.
- ELLIS, B. & WONG, W. (2008). Learning causal Bayesian network structures from experimental data. *J Am Stat Assoc* **103**, 778–789.
- EVTUSHENKO, Y., MALKOVA, V. & STANEVICHYUS, A. (2009). Parallel global optimization of functions of several variables. *Comput. Math. Math. Phys.* **49**, 246–260.
- FRIEDMAN, J., HASTIE, T. & TIBSHIRANI, R. (2007). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**, 432–441.
- FRIEDMAN, N. & KOLLER, D. (2003). Being Bayesian about network structure: a Bayesian approach to structure discovery in Bayesian networks. *Mach Learn* **50**, 95–125.
- FU, F. & ZHOU, Q. (2013). Learning sparse causal Gaussian networks with experimental intervention: regularization and coordinate descent. *J Am Stat Assoc* **108**, 288–300.
- GIRAUD, C. (2015). *Introduction to high-dimensional statistics*, vol. 139 of *Monographs on Statistics and Applied Probability*. CRC Press, Boca Raton, FL.
- GRANVILLE, V., KRIVANEK, M. & RASSON, J.-P. (1994). Simulated annealing: A proof of convergence. *IEEE T Pattern Anal* **16**, 652–656.
- GREFENSTETTE, J., GOPAL, R., ROSMAITA, B. & VAN GUCHT, D. (1985). Genetic algorithms for the traveling salesman problem. In *Proceedings of the first International Conference on Genetic Algorithms and their Applications*.
- GRZEGORCZYK, M. & HUSMEIER, D. (2008). Improving the structure MCMC sampler for bayesian networks by introducing a new edge reversal move. *Mach Learn* **71**.
- GUYON, I., ALIFERIS, C. & COOPER, G. (2010). *Causation and Prediction Challenge*, vol. 2. Microtome Publishing. Challenges in Machine Learning.
- HAUSER, A. & BÜHLMANN, P. (2012). Characterization and greedy learning of interventional Markov equivalence classes of directed acyclic graphs. *J Mach Learn Res* **13**, 2409–2464.

- HAUSER, A. & BÜHLMANN, P. (2015). Jointly interventional and observational data: estimation of interventional Markov equivalence classes of directed acyclic graphs. *J R Stat Soc. Series B* **77**, 291–318.
- HOGBEN, L., ed. (2007). *Handbook of linear algebra*. Discrete Mathematics and its Applications (Boca Raton). Chapman & Hall/CRC, Boca Raton, FL. Associate editors: Richard Brualdi, Anne Greenbaum and Roy Mathias.
- HOLLAND, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, Mich. An introductory analysis with applications to biology, control, and artificial intelligence.
- HORST, R. & PARDALOS, P. M., eds. (1995). *Handbook of global optimization*. Nonconvex optimization and its applications. Dordrecht, Boston: Kluwer Academic Publishers.
- HUYNH-THU, V. A., IRRTHUM, A., WEHENKEL, L. & GEURTS, P. (2010). Inferring regulatory networks from expression data using tree-based methods. *PLoS ONE* **5**, e12776.
- JONES, D. R., PERTTUNEN, C. D. & STUCKMAN, B. E. (1993). Lipschitzian optimization without the lipschitz constant. *J Optimiz Theory App* **79**, 157–181.
- KAHN, A. B. (1962). Topological sorting of large networks. *Commun. ACM* **5**, 558–562.
- KALISCH, M. & BÜHLMANN, P. (2007). Estimating high-dimensional directed acyclic graphs with the PC-algorithm. *J Mach Learn Res* **8**, 613–636.
- KOIVISTO, M. & SOOD, K. (2004). Exact bayesian structure discovery in Bayesian networks. *J Mach Learn Res* **5**, 549–573.
- LEDOIT, O. & WOLF, M. (2015). Spectrum estimation: A unified framework for covariance matrix estimation and pca in large dimensions. *J Multivariate Anal* **139**, 360384.
- LIAW, A. & WIENER, M. (2002). Classification and regression by randomforest. *R News* **2**, 18–22.
- LIU, H., WANG, L. & ZHAO, T. (2014). Sparse covariance matrix estimation with eigenvalue constraints. *J Comput Graph Stat* **23**, 439–459.
- LOUNICI, K., PONTIL, M., TSYBAKOV, A. & VAN DE GEER, S. (2009). Taking advantage of sparsity in multi-task learning. In *Proceedings of the 22nd Conference on Learning Theory*.
- MARBACH, D., SCHAFFTER, T., MATTIUSSI, C. & FLOREANO, D. (2009). Generating realistic in silico gene networks for performance assessment of reverse engineering methods. *J Comput Bio* **16**, 229–239.

- MARGOLIN, A., NEMENMAN, I., BASSO, K., WIGGINS, C., STOLOVITZKY, G., DALLA FAVERA, R. & CALIFANO, A. (2006). Aracne: An algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics* **7**, S7.
- MEINSHAUSEN, N. & BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *Ann Stat* **34**, 1436–1462.
- MESTRE, X. (2008). Improved estimation of eigenvalues and eigenvectors of covariance matrices using their sample estimates. *IEEE T Inform Theory* **54**, 5113–5129.
- MICHALEWICZ, Z. (1994). *Genetic algorithms + data structures = evolution programs*. Springer-Verlag, Berlin, 2nd ed.
- NEWMAN, M. (2003). The structure and function of complex networks. *SIAM Rev* **45**, 157–256.
- PEARL, J. (2009). *Causality*. Cambridge University Press, Cambridge, 2nd ed. Models, reasoning, and inference.
- PETERS, J., MOOIJ, J., JANZING, D. & SCHÖLKOPF, B. (2014). Causal discovery with continuous additive noise models. *J Mach Learn Res* **15**, 2009–2053.
- PETERS, J., MOOIJ, J. M., JANZING, D. & SCHÖLKOPF, B. (2011). Identifiability of causal graphs using functional models. In *27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*.
- PISZCZ, A. & SOULE, T. (2006). Genetic programming: Optimal population sizes for varying complexity problem. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- RAU, A., JAFFR’EZIC, F. & NUEL, G. (2013). Joint estimation of causal effects from observational and intervention gene expression data. *BMC Syst Biol* **7**, 111.
- RIDGE, E. (2007). *Design of Experiments for the Tuning of Optimisation Algorithm*. Ph.D. thesis, The University of York, Department of Computer Science.
- ROBINSON, R. W. (1973). Counting labeled acyclic digraphs , 239–273.
- SCHAFFER, J. D., CARUANA, R., ESHELMAN, L. J. & DAS, R. (1989). A study of control parameters affecting online performance of genetic algorithms for function optimization. In *Proceedings of the Third International Conference on Genetic Algorithms*.
- SCHWARZ, G. E. (1978). Estimating the dimension of a model. *The Annals of Statistics* **6**, 461–464.
- SERGEYEV, Y. (1995). An information global optimization algorithm with local tuning. *SIAM J Optimiz* **5**, 858–870.

- SERGEYEV, Y. & KVASOV, D. (2006). Global search based on efficient diagonal partitions and a set of lipschitz constants. *SIAM J Optimiz* **16**, 910–937.
- SHOJAIE, A. & MICHAILIDIS, G. (2010). Penalized likelihood methods for estimation of sparse high-dimensional directed acyclic graphs. *Biometrika* **97**, 519–538.
- SILANDER, T. & MYLLYMÄKI, T. (2006). A simple approach for finding the globally optimal bayesian network structure. In *Proceedings of the Twenty-second Conference on Uncertainty in Artificial Intelligence*.
- SOUMA, W., FUJIWARA, Y. & AOYAMA, H. (2006). Heterogeneous economic networks. In *The complex networks of economic interactions*, vol. 567 of *Lecture Notes in Econom. and Math. Systems*. Springer, Berlin, pp. 79–92.
- SPIRITES, P., GLYMOUR, C. & SCHEINES, R. (2000). *Causation, prediction, and search*. Adaptive Computation and Machine Learning. Cambridge, MA: MIT Press, 2nd ed. With additional material by David Heckerman, Christopher Meek, Gregory F. Cooper and Thomas Richardson, A Bradford Book.
- TIBSHIRANI, R. (1996). Regression shrinkage and selection via the lasso. *J R Stat Soc Series. B* **58**, 267–288.
- TSAMARDINOS, I., BROWN, L. & ALIFERIS, C. (2006). The max-min hill-climbing bayesian network structure learning algorithm. *Mach Learn* **65**, 31–78.
- VAN DE GEER, S. & BÜHLMANN, P. (2013). ℓ_0 -penalized maximum likelihood for sparse directed acyclic graphs. *Ann Stat* **41**, 536–567.
- VERMA, T., ARAÚJO, N. & HERRMANN, H. (2014). Revealing the structure of the world airline network. *Scientific Reports* **5**.
- VERMA, T. & PEARL, J. (1991). Equivalence and synthesis of causal models. In *Proceedings of the 6th Annual Conference on Uncertainty in Artificial Intelligence*.
- VERZELEN, N. (2012). Minimax risks for sparse regressions: Ultra-high dimensional phenomena. *Electron J Stats* **6**, 38–90.
- WAINWRIGHT, M. (2009). Information-theoretic limits on sparsity recovery in the high-dimensional and noisy setting. *IEEE T Info Theory* **55**, 5728–5741.
- WITTEN, D., FREIDMAN, J. & SIMON, N. (2011). New insights and faster computations for the graphical lasso. *J CComput Graph Stat* **20**, 892–900.
- WRIGHT, S. (1921). Corelation and causation. *J Agric Res* , 558–585.
- YUAN, M. & LIN, Y. (2007). Model selection and estimation in the Gaussian graphical model. *Biometrika* **94**, 19–35.
- ZHOU, Q. (2011). Multi-domain sampling with applications to structural inference of Bayesian networks. *J Am Stat Assoc* **106**, 1317–1330.