
Inference of Sparse Networks with Unobserved Variables. Application to Gene Regulatory Networks

Nikolai Slavov

Princeton University, NJ 08544

nslavov@princeton.edu

Abstract

Networks are becoming a unifying framework for modeling complex systems and network inference problems are frequently encountered in many fields. Here, I develop and apply a generative approach to network inference (*RCweb*) for the case when the network is sparse and the latent (not observed) variables affect the observed ones. From all possible factor analysis (FA) decompositions explaining the variance in the data, *RCweb* selects the FA decomposition that is consistent with a sparse underlying network. The sparsity constraint is imposed by a novel method that significantly outperforms (in terms of accuracy, robustness to noise, complexity scaling and computational efficiency) Bayesian methods and *MLE* methods using ℓ_1 norm relaxation such as K-SVD and ℓ_1 -based sparse principle component analysis (PCA). Results from simulated models demonstrate that *RCweb* recovers exactly the model structures for sparsity as low (as non-sparse) as 50% and with ratio of unobserved to observed variables as high as 2. *RCweb* is robust to noise, with gradual decrease in the parameter ranges as the noise level increases.

1 INTRODUCTION

Factor analysis (FA) decompositions are useful for explaining the variance of observed variables in terms of fewer unobserved variables that may capture systematic effects and allow for low dimensional representation of the data. Yet, the interpretation of la-

tent variables inferred by FA is fraught with problems. In fact, interpretation is not always expected and intended since FA may not have an underlying generative model. A prime difficulty with interpretation arises from the fact that any rotation of the factors and their loadings by an orthogonal matrix results in a different FA decomposition that explains the variance in the observed variables just as well. Therefore, in the absence of additional information on the latent variables and their loading, FA cannot identify a unique decomposition, much less generative relationships between latent factors and observed variables.

A frequent choice for a constraint implemented by principle component analysis (PCA) and resulting in a unique solution is that the factors are the singular vectors (and thus orthogonal to each other) of the data matrix ordered in descending order of their corresponding singular values. Yet, this choice is often motivated by computational convenience rather than by knowledge about the system that generated the data. Another type of computationally convenient constraint applied to facilitate the interpretation of FA results is sparsity, in the form of sparse Bayesian FA (West 2002, Dueck 2005, Carvalho 2008), sparse PCA (d'Aspremont 2007, Sigg 2008) and FA for gene regulatory networks (Srebro 2001, Pe'er 2002). However, papers that introduce and use sparse PCA do not consider a generative model but rather use sparsity as a convenient tool to produce interpretable factors that are linear combinations of just a few original variables. In sparse PCA, sparsity is a way to balance interpretability at the cost of slightly lower fraction of explained variance.

The algorithm described in this paper (*RCweb*) also uses a sparse prior, but *RCweb* explicitly considers the problem from a generative perspective. *RCweb* asserts that there is indeed a set of hidden variables that connect to and regulate the observed variables via a sparse network. Based on that model, I derive a network structure learning approach within explicit theoretical framework. This allows to propose an approach for

Appearing in Proceedings of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS) 2010, Chia Laguna Resort, Sardinia, Italy. Volume 9 of JMLR: W&CP 9. Copyright 2010 by the authors.

sparse FA which is conceptually and computationally different from all existing approaches such as K-SVD (Aharon 2005), sparse PCA and other LARS (Bradley 2004) based methods (Banerjee 2007, Friedman 2008). *RCweb* is appropriate for analyzing data arising from any system in which the state of each observed variable is affected by a strict subset of the unobserved variables. To assign the inferred latent variables to physical factors, *RCweb* needs either data from perturbation experiments or prior knowledge about the factors. This framework generalizes to non-linear interactions, which is discussed elsewhere. Furthermore, I analyze the scaling of the computational complexity of *RCweb* with the number of observed and unobserved variables, as well as the parameter space where *RCweb* can accurately infer network topologies and demonstrate its robustness to noise in the data.

2 DERIVATION

Consider a sparse bipartite graph $\mathcal{G} = (\mathcal{E}, \mathcal{N}, \mathcal{R})$ consisting of two sets of vertices \mathcal{N} and \mathcal{R} and the associated set of directed edges \mathcal{E} connecting \mathcal{R} to \mathcal{N} vertices. Define a graphical model in which each vertex s corresponds to a random variable; N observed random variables indexed by \mathcal{N} ($x_{\mathcal{N}} = \{x_s | s \in \mathcal{N}\}$) whose states are functions of P unobserved variables indexed by \mathcal{R} , $x_{\mathcal{R}} = \{x_s | s \in \mathcal{R}\}$. Since the states of $x_{\mathcal{N}}$ depend on (are regulated by) $x_{\mathcal{R}}$, I will also refer to $x_{\mathcal{R}}$ as regulators. The functional dependencies are denoted by a set of directed edges \mathcal{E} so that each unobserved variable $x_i | i \in \mathcal{R}$ affects (and its vertex is thus connected to) a subset of observed variables $x_{\alpha_i} = \{x_s | s \in \alpha_i \subset \mathcal{N}\}$. Given a dataset $\mathbf{G} \in \mathbb{R}^{M \times N}$ of M configurations of the observed variables $x_{\mathcal{N}}$, *RCweb* aims to infer the edges \mathcal{E} and the corresponding configurations of the unobserved variables, $x_{\mathcal{R}}$.

If the state of each observed variable is a linear superposition of a subset of unobserved variables, the data \mathbf{G} can be modeled with a very simple *generative model* (1): The data is a product between $\mathbf{R} \in \mathbb{R}^{M \times P}$ (a matrix whose columns correspond to the unobserved variables and the rows correspond to the M measured configurations) and $\mathbf{C} \in \mathbb{R}^{P \times N}$, the weighted adjacency matrix of \mathcal{G} . The unexplained variance in the data \mathbf{G} is captured by the residual \mathbf{Y} .

$$\mathbf{G} = \mathbf{RC} + \mathbf{Y} \quad (1)$$

This decomposition of \mathbf{G} into a product of two matrices can be considered to be a type of factor analysis with \mathbf{R} being the factors and \mathbf{C} the loadings. Even when $P \ll M$ the decomposition of \mathbf{G} does not have a unique solution since $\mathbf{RC} \equiv \mathbf{RZC} \equiv \mathbf{RQ}^T\mathbf{Q}\mathbf{C} \equiv \mathbf{R}^*\mathbf{C}^*$ for any orthonormal matrix \mathbf{Q} . Thus the identification of a unique decomposition corresponding to

the structure of \mathcal{G} requires additional criteria constraining the decomposition. The assumption that \mathcal{G} is sparse requires that \mathbf{C} be sparse as well meaning that the state of the i^{th} observed vertex $x_i | i \in \mathcal{N}$ is affected by a strict subset of the unobserved variables $x_{\psi_i} = \{x_s | s \in \psi_i \subset \mathcal{R}\}$, the ones whose weights in the i^{th} column of \mathbf{C} are not zeros, $\mathbf{C}_{i\psi_i} \neq 0$ and $\mathbf{C}_{i\bar{\psi}_i} = 0$ where $\bar{\psi}_i$ is the complement of ψ_i . Thus, to recover the structure of \mathcal{G} , *RCweb* seeks to find a decomposition of \mathbf{G} in which \mathbf{C} is sparse. The sparsity can be introduced as a regularization with a Lagrangian multiplier λ :

$$(\hat{\mathbf{C}}, \hat{\mathbf{R}}) = \arg \min_{\mathbf{R}, \mathbf{C}} \|\mathbf{G} - \mathbf{RC}\|_F^2 + \lambda \|\mathbf{C}\|_0 \quad (2)$$

In the equations above and throughout the paper $\|\mathbf{C}\|_F^2 = \sum_{i,j} C_{ij}^2$ denotes entry-wise (Frobenius) norm, and the zero norm of a vector or matrix ($\|\mathbf{C}\|_0$) equals the number of non-zero elements in the array.

To infer the network topology *RCweb* aims to solve the optimization problem defined by (2). Since (2) is a NP-hard combinatorial problem, the solution can be simplified significantly by relaxing the ℓ_0 norm to ℓ_1 norm (Bradley 2004). Then the approximated problem can be tackled with interior point methods (Banerjee 2007). As an alternative approach to ℓ_1 approximation, I propose a novel method based on introducing a degree of freedom in the singular-value decomposition (SVD) of \mathbf{G} by inserting an invertible¹ matrix \mathbf{B} .

$$\mathbf{G} = \mathbf{USV}^T \equiv \underbrace{(\mathbf{US}(\mathbf{B}^T)^{-1})}_{\hat{\mathbf{R}}} \underbrace{(\mathbf{B}^T\mathbf{V}^T)}_{\hat{\mathbf{C}}} \quad (3)$$

The prior (constraint) that $\hat{\mathbf{C}}$ is sparse determines \mathbf{B} that minimizes (2), and thus a unique decomposition. The goal of introducing \mathbf{B} is to reduce the combinatorial problem to one that can be solved with convex minimization. When the factors underlying the observed variance are fewer than the observations in \mathbf{G} there is no need to take the full SVD; if P factors are expected, only the first P largest singular vectors and values from the SVD of \mathbf{G} are taken in that decomposition so that \mathbf{USV}^T is the matrix with rank P that best approximates \mathbf{G} in the sense of minimizing $\|\mathbf{G} - \mathbf{USV}^T\|_F^2$. Since conceivably sparse decompositions may use columns outside of the best ℓ_2 approximation, *RCweb* considers taking the first P^* for $P^* > P$ singular pairs. Such expanded basis is more likely to support the optimal sparse solution and especially relevant for the case when P is not known. Such choice can be easily accommodated in light of the ability of *RCweb* to exclude unnecessary explanatory variables, see section 4.3.

¹ \mathbf{B} is always invertible by construction, see section 4.3 and equation (4)

Next \mathcal{RCweb} computes \mathbf{B} based on the requirement that \mathbf{C} is sparse for the case $N > P$. To compute \mathbf{B} , one may set an optimization problem (4). Once \mathbf{B} is inferred, $\hat{\mathbf{R}}$ and $\hat{\mathbf{C}}$ can be computed easily, $\hat{\mathbf{R}} = \mathbf{U}\hat{\mathbf{B}}^{-1}$ and $\hat{\mathbf{C}} = (\mathbf{V}\hat{\mathbf{B}})^T$.

$$\hat{\mathbf{B}} = \arg \min_{\mathbf{B}} \|\mathbf{V}\mathbf{B}\|_0, \text{ so that } \det(\mathbf{B}) > 1 \quad (4)$$

The constraint on \mathbf{B} in (4) is introduced to avoid trivial and degenerate solutions, such as \mathbf{B} being rank deficient. Thus the introduction of \mathbf{B} reduces (2) to a problem (4) that is still combinatorial and might also be approximated with a more tractable problem by relaxing the ℓ_0 norm to ℓ_1 norm and applying heuristics (Cetin 2002, Candès 2007) to enhance the solution. I propose a new approach, \mathcal{RCweb} , outlined in the next section.

3 \mathcal{RCweb}

Assume that a sparse \mathbf{c}_i^T corresponding to an optimal $\hat{\mathbf{b}}_i$ (the i^{th} column of $\hat{\mathbf{B}}$) form $\mathbf{V}\hat{\mathbf{b}}_i = \mathbf{c}_i^T$ is known. Define the set of indices corresponding to non-zero elements in \mathbf{c}_i^T with ω_- and the set of indices corresponding to zero elements in \mathbf{c}_i^T with ω_0 . Furthermore, define the matrix \mathbf{V}_{ω_0} to be the matrix containing only the rows of \mathbf{V} whose indices are in ω_0 . If ω_0 and thus \mathbf{V}_{ω_0} are known, one can easily compute $\hat{\mathbf{b}}_i$ as the right singular vector of \mathbf{V}_{ω_0} corresponding to the zero singular value. Since \mathbf{c}_i^T and ω_0 are not known, \mathcal{RCweb} approximates $\hat{\mathbf{b}}_i$ (the smallest² right singular vector of \mathbf{V}_{ω_0}) with \mathbf{v}_s , the smallest right singular vector of \mathbf{V} . This approximation relies on assuming that a low rank perturbation in a matrix results in a small change in its smallest singular vectors (Wilkinson 1988, Benaych-Georges 2009). Thus given that \mathcal{RCweb} is looking for the sparsest solution and the set ω_- is small relative to N , the angle between the singular vectors of \mathbf{V}_{ω_0} and \mathbf{V} is small as well. Therefore, \mathbf{v}_s can serve as a reasonable first approximation of \mathbf{b}_i . Then \mathcal{RCweb} systematically and iteratively uses and updates \mathbf{v}_s by removing rows of \mathbf{V} until \mathbf{v}_s converges to \mathbf{b}_i or equivalently \mathbf{V}_{ω_0} becomes singular for the largest set of ω_0 indices. When \mathbf{V}_{ω_0} becomes singular, all elements of \mathbf{c}_i whose indices are in ω_0 become zero.

\mathcal{RCweb} also has an intuitive geometrical interpretation. Consider the matrix \mathbf{V} mapping the unit sphere in \mathbb{R}^P (the sphere with unit radius from \mathbb{R}^P) to an ellipsoid in \mathbb{R}^N . The axes of the ellipsoid are the left singular vectors of \mathbf{V} . In this picture, starting with $\omega_- = \{\emptyset\}$ and $\omega_0 = \{1, \dots, N\}$, solving (4) requires moving the fewest number of indices from ω_0 to ω_- so that \mathbf{V}_{ω_0}

maps a vector from \mathbb{R}^P to the origin. How to choose the indices to be moved? At each step \mathcal{RCweb} chooses $i_{|max|}$, the index of the largest element (by absolute value) of the smallest axis of the ellipsoid which is the left singular vector of \mathbf{V} with the smallest singular value. \mathcal{RCweb} moves $i_{|max|}$ from ω_0 to ω_- effectively selecting the dimension whose projection is easiest to eliminate and removing its largest component, which minimizes as much as possible the projection in that dimension. \mathcal{RCweb} keeps moving indices from ω_0 to ω_- using the same procedure until the smallest right singular vector of \mathbf{V}_{ω_0} converges to \mathbf{b}_i and the smallest singular value of \mathbf{V}_{ω_0} approaches zero. \mathcal{RCweb} is guaranteed to stop after at most $(N-P+1)$ steps since after removal of $(N-P+1)$ indices from ω_0 , \mathbf{V}_{ω_0} will be at most rank $P-1$. If \mathcal{RCweb} finds a sparse solution it will converge in fewer steps.

1. Task:

$$\hat{\mathbf{b}}_i = \min_{\mathbf{b}_i^T \mathbf{b}_i \geq 1} \|\mathbf{V}\mathbf{b}_i\|_0$$

2. Initialization:

- $\omega_- = \{\emptyset\}$ and $\omega_0 = \{1, 2, \dots, N\}$
- Set $\mathbf{K}_{\omega_0}^{-1} = (\mathbf{V}_{\omega_0}^T \mathbf{V}_{\omega_0})^{-1} = \mathbf{I} \in \mathbb{R}^{P \times P}$
- Set $J = 1$;
- $i_{|max|} = \arg \max_j \left(\sum_j |\mathbf{V}_{ij}| \right)$
 $\omega_- = \{i_{|max|}\}$, $\omega_0 = \{i | i \in \omega_0, i \neq i_{|max|}\}$
- Update $\mathbf{K}_{\omega_0}^{-1} = \text{RankUpdate}(\mathbf{K}_{\omega_0}^{-1}, \mathbf{V}_{i_{|max|}})$

3. Cycle: $J = J + 1$ Repeat until convergence

- Find the eigenvector \mathbf{v} for $\mathbf{K}_{\omega_0}^{-1}$ with the largest eigenvalue λ_{max}
- If $\lambda_{max}^{-1} \approx 0$ or $\mathbf{v}^J \rightarrow \mathbf{v}^{J-1}$, $\hat{\mathbf{b}}_i \equiv \mathbf{v}$; **STOP**
- Compute the left singular vector of \mathbf{V}_{ω_0}
 $\mathbf{u} = \mathbf{s}^{-1} \mathbf{V}_{\omega_0} \mathbf{v}$
- $i_{|max|} = \arg \max [(|u_1|, \dots, |u_i|, \dots, |u_N|)]$;
- $\omega_- = \{\omega_-, i_{|max|}\}$
 $\omega_0 = \{i | i \in \omega_0, i \neq i_{|max|}\}$
- Update $\mathbf{K}_{\omega_0}^{-1} = \text{RankUpdate}(\mathbf{K}_{\omega_0}^{-1}, \mathbf{V}_{i_{|max|}})$

The above algorithm can compute a single vector, $\hat{\mathbf{b}}_i$, which is just one column of $\hat{\mathbf{B}}$. To find the other columns, \mathcal{RCweb} applies the same approach to the modified (inflated) matrix, which for the i^{th} column of \mathbf{B} is $\mathbf{V}^{(i)} = \mathbf{V}^{(i-1)} + \mathbf{V}\hat{\mathbf{b}}_i\hat{\mathbf{b}}_i^T$ for $i = 2, \dots, P$. Thus, after the inference of each column of \mathbf{B} \mathcal{RCweb} modifies $\mathbf{V}^{(i-1)}$ to $\mathbf{V}^{(i)}$ ($\mathbf{V}^{(1)} \equiv \mathbf{V}$) so that the algorithm will not replicate its choice of ω_0 . Note that in the i^{th} update of $\mathbf{V}^{(i-1)}$ the $\mathbf{V}\hat{\mathbf{b}}_i\hat{\mathbf{b}}_i^T$ will modify only the ω_- rows in $\mathbf{V}^{(i-1)}$ since the rows of the $\mathbf{V}\hat{\mathbf{b}}_i\hat{\mathbf{b}}_i^T$ whose indices are in ω_0 contain only zero elements. Applying \mathcal{RCweb} to the inflated matrices avoids inferring

²By smallest singular vector I mean the singular vector corresponding to the to the smallest singular value

multiple times the same $\hat{\mathbf{b}}_i$, but a $\hat{\mathbf{b}}_i$ inferred from the inflated matrix is generally going to differ at least slightly from the corresponding $\hat{\mathbf{b}}_i$ that solves (4) for \mathbf{V} . To avoid that, *RCweb* uses the inflated matrices only for the first few iterations until the largest (by absolute magnitude) of the Pearson correlations between the smallest eigenvector of \mathbf{V}_{ω_0} from the current (i^{th}) iteration and the recovered columns of \mathbf{B} is less than $1 - \epsilon$ and monotonically decreasing; ϵ is chosen for numerical stability and also to reflect the similarity between the connectivity of \mathcal{R} vertices that can be expected in the network whose topology is being recovered. A simpler alternative that works great in practice is to use the inflated matrix for the first k iterations that are enough to find a new direction for $\hat{\mathbf{b}}_i$ and then *RCweb* switches back to \mathbf{V} so that the solution is optimal for \mathbf{V} . The switch requires k rank update of $\mathbf{K}_{\omega_0}^{-1} \equiv (\mathbf{V}_{\omega_0}^T \mathbf{V}_{\omega_0})^{-1}$ and thus choosing k small saves computations. Choosing k too small, however, may not be enough to guarantee that $\hat{\mathbf{b}}_i$ will not recapitulate a solution that is already found. Usually $k = 10$ works great and can be easily increased if the new solution is very close to an old one.

There are a few notable elements that make *RCweb* efficient. First, *RCweb* does almost all computations in \mathbb{R}^P and since $P \ll N$, $P < M$, that saves both memory and CPU time. Second, each step requires only a few matrix-vector multiplication for computing the eigenvectors (since the change from the previous step is generally very small) and $\mathbf{K}_{\omega_0}^{-1}$ is computed by a rank-one update which obviates matrix inversion.

The approach that *RCweb* takes in solving (4) does not impose specific restrictions on the distribution of the observed variables (\mathbf{G}), the noise in the data (\mathbf{Y}) or the latent variables $\hat{\mathbf{R}}$. However, the initial approximation of \mathbf{b}_i with \mathbf{v}_s can be poor for data arising from dense networks or special worst-case datasets. As demonstrated theoretically (Benaych-Georges 2009) and tested numerically in the next section, *RCweb* performs very well at least in the absence of worst-case scenario special structures in the data.

4 VALIDATION

To evaluate the performance of *RCweb*, I first apply it to data from simulated random bipartite networks with two different topology types, (1) Erdős & Rényi and (2) scale-free whose corresponding degree distributions are (1) Poisson and (2) power-law. The network topology is encoded in a weighted adjacency matrix \mathbf{C}^{gold} and the values for the unobserved variables are drawn from a standard uniform distribution. The simulations result in data matrices $\mathbf{G} \in \mathbb{R}^{M \times N}$ containing M observations of all N unobserved variables. Accord-

ing to *RCweb*, the optimal sparse adjacency matrix ($\hat{\mathbf{C}}$) and the hidden variables ($\hat{\mathbf{R}}$) can be inferred by the decomposition, $\hat{\mathbf{G}} = \hat{\mathbf{R}}\hat{\mathbf{C}}$ so that $\hat{\mathbf{C}}$ is as sparse as possible while $\hat{\mathbf{G}}$ is as close as possible to \mathbf{G} .

In addition to *RCweb*, such decomposition can be computed by 3 classes of existing algorithms. For a comparison, I use the latest versions for which the authors report best performance: (A) *PSMF* for Bayesian matrix factorization as implemented by the author Matlab function PSMF1 (Dueck 2005); (B) *BFRM 2* for Bayesian matrix factorization as implemented by the author compiled executable (Carvalho 2008); (C) *em-PCA* for maximum likelihood estimate (MLE) sparse PCA (Sigg 2008); (D) *K-SVD* (Aharon 2005). All algorithms are implemented using the code published by their authors, and with the default values of the parameters when parameters are required. The results are compared for various M , N , P , sparsity, and noise levels.

4.1 LIMITATIONS

Before comparing the results, consider some of the limitations common to all algorithms and the appropriate metrics for comparing the results. In the absence of any other information, the decomposition of \mathbf{G} (no matter how accurate) cannot associate hidden variables (corresponding to columns of $\hat{\mathbf{R}}$) to physical factors. Furthermore, all methods can infer $\hat{\mathbf{C}}$ and $\hat{\mathbf{R}}$ only up to an arbitrary diagonal scaling or permutation matrix. First consider the scaling illustrated by the following transformation by a diagonal matrix \mathbf{D} , $\hat{\mathbf{G}} = \hat{\mathbf{R}}\hat{\mathbf{C}} = \hat{\mathbf{R}}\mathbf{Z}\hat{\mathbf{C}} = \hat{\mathbf{R}}(\mathbf{D}\mathbf{D}^{-1})\hat{\mathbf{C}} = (\hat{\mathbf{R}}\mathbf{D})(\mathbf{D}^{-1}\hat{\mathbf{C}}) = \hat{\mathbf{R}}^*\hat{\mathbf{C}}^*$. Such transformation is going to rescale $\hat{\mathbf{R}}$ to $\hat{\mathbf{R}}^*$ and $\hat{\mathbf{C}}$ to $\hat{\mathbf{C}}^*$, which is just as sparse as $\hat{\mathbf{C}}$, $\|\hat{\mathbf{C}}^*\|_0 = \|\hat{\mathbf{C}}\|_0$. Since both decompositions explain the variance in \mathbf{G} equally well *RCweb* (or any of the other method) cannot distinguish between them. Thus given $\hat{\mathbf{C}}$, there is a diagonal matrix $\hat{\mathbf{D}}$ that scales $\hat{\mathbf{C}}$ to \mathbf{C}^{gold} , the weighted adjacency matrix of \mathcal{G} .

The second limitation is that in the absence of addition information, *RCweb* can determine $\hat{\mathbf{C}}$ and $\hat{\mathbf{R}}$ up to a permutation matrix. Consider comparing $\hat{\mathbf{C}}$ to the adjacency matrix used in the simulations, \mathbf{C}^{gold} . Since the identity of the inferred hidden variables is not known the rows of $\hat{\mathbf{C}}$ do not generally correspond to the rows of \mathbf{C}^{gold} ; $\hat{\mathbf{C}}_i$ (the i^{th} row of $\hat{\mathbf{C}}$) is most likely to correspond to the \mathbf{C}^{gold} row that is most correlated to $\hat{\mathbf{C}}_i$ and the Pearson correlation between the two rows quantifies the accuracy for the inference of $\hat{\mathbf{C}}_i$. To implement this idea, all rows of $\hat{\mathbf{C}}$ and \mathbf{C}^{gold} are first normalized to mean zero and unit variance resulting in \mathbf{C}_{nor} and \mathbf{C}_{nor}^{gold} . The correlation matrix then is, $\Sigma = \mathbf{C}_{nor}^T \mathbf{C}_{nor}^{gold}$ and the most likely vertex (index of the unobserved variable) corresponding to

$\hat{\mathbf{C}}_i$ is $k = \arg \max_j (|\Sigma_{i1}|, \dots, |\Sigma_{ij}|, \dots, |\Sigma_{iP}|)$, where $k \in \mathcal{R}$. The absolute value is required because the diagonal elements of $\hat{\mathbf{D}}$ can be negative. The accuracy is measured by the corresponding Pearson correlation, $\rho_i = |\Sigma_{ik}|$. An optimal solution of this matching problem can be found by using belief propagation algorithm for the simple case of a bipartite graph even the *LP* relaxed version guarantees optimal solution (Sanghavi 2007). The overall accuracy is quantified by the mean correlation $\bar{\rho} = (1/p) \sum_{i=1}^{i=p} \rho_i$ where p is the number of inferred unobserved variables and can equal to P or not depending on whether the number of unobserved variables is known or not. In computing $\bar{\rho}$, each row of \mathbf{C}^{gold} is allowed to correspond only to one row of $\hat{\mathbf{C}}$ and vice versa.

In addition to the two common limitations of permutation and scaling, some algorithms (d’Aspremont 2007) for sparse PCA require $M > N$ and since this is not the case in many real world problems and in some of the datasets simulated here, those methods are not tested. Instead I chose emPCA, which does not require $M > N$ and is the latest *MLE* algorithm for sparse PCA that according to its authors is more efficient than previous algorithms (Sigg 2008).

4.2 ACCURACY AND COMPLEXITY SCALING

RCweb has a natural way for identifying the mean degree³. However, some of the other algorithms require the mean degree for optimal performance. To avoid underestimating an algorithm simply because it recovers networks that are too sparse or not sparse enough, I assume that the mean degree is known and it is input to all algorithms. First all algorithms are tested on a very easy inference problem, Fig.1. Since PSMF and BFRM have lower accuracy and PSMF is significantly slower than the other algorithms, the rest of the results will focus on the *MLE* algorithms that also have better performance. PSMF gives less accurate results with power-law networks which can be understood in terms of the uniform prior used by PSMF. PSMF has the advantage over the *MLE* algorithms in inferring a probabilistic network structure rather than a single estimate. Special advantage of BFRM is the seamless inclusion of response variables and measured factors in the inference. For the *MLE* algorithms, the accuracy of network inference increases with the ratio of observed to unobserved variables N/P (Fig.2) and with the number of observed configurations M , Fig.3. In contrast, as the noise in the data and the mean out-degree (mean number of edges from \mathcal{R} to \mathcal{N}

³When *RCweb* learns all edges, the smallest singular value of \mathbf{V}_{ω_0} approaches zero and its smallest singular vector converges to $\hat{\mathbf{b}}_1$.

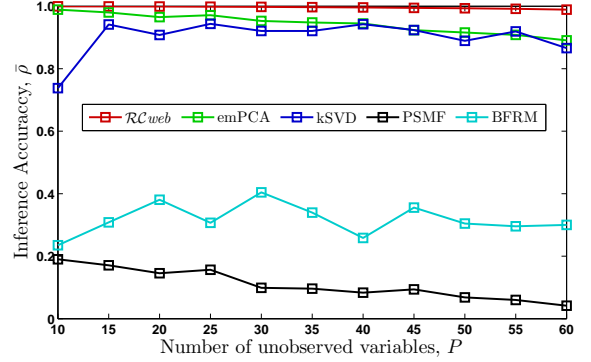


Figure 1: Accuracy of network recovery as a function of the number of unobserved variables, P . Number of observed variables $N = 500$; Number of observations, $M = 280$. All networks are with Poisson mean out-degree $0.10N = 50$, with 10 % noise in the observations.

vertices) are increased, the accuracy of the inference decreases. All algorithms perform better on Poisson networks and the lower level of noise in the data from power-law networks was chosen to partially compensate for that. An important caveat when comparing the results for different algorithms is that K-SVD iteratively improves the accuracy of the solution, and thus the output is dependent on the maximum number of iterations allowed (I_{max}). For the results here, $I_{max} = 20$ and the accuracy of K-SVD may improve with higher number of iterations even though I did not observe significant improvement with $I_{max} = 100$. Even at 20 iterations K-SVD is significantly slower than *RCweb* and emPCA, Fig.4. The scaling of the algorithms with respect to a parameter was determined by holding all other parameters constant and regressing the log of the CPU time against the log of the variable parameter, Fig.4. The scaling with respect to some parameters is below the theoretical expectation since the highest complexity steps may not be speed-limiting for the ranges of parameters used in the simulations.

4.3 INTERPRETABILITY

In analyzing real data the number of unobserved variables (P) may not be known. I observe that if a simulated network having P regulators is inferred assuming $P^* > P$ regulators, the elements of \mathbf{B} corresponding to the excessive regulators are very close to zero, $|B_{ij}| \leq 10^{-10}$ for $i > P$ or $j > P$. Thus, if the data truly originate from a sparse network *RCweb* can discard unnecessary unobserved variables.

The results of *RCweb* can be valuable even without

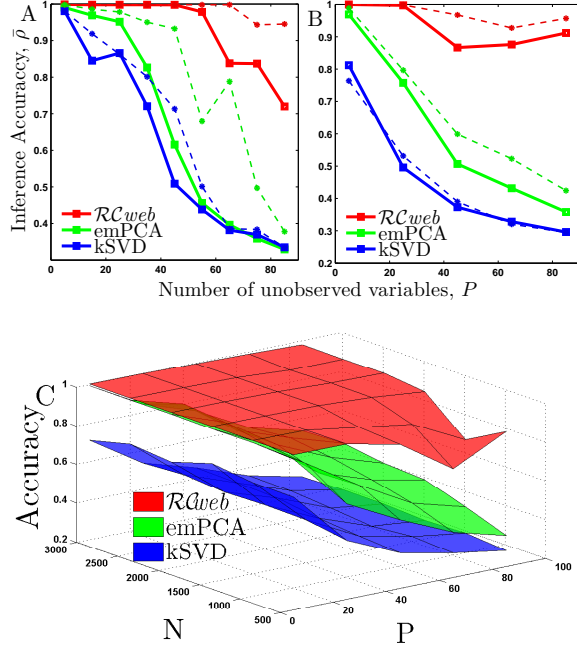


Figure 2: Accuracy of network recovery as a function of the number of unobserved variables, P . The thicker brighter lines with squares correspond to number of observed variables $N = 500$ and the dashed, thinner lines correspond to $N = 1000$. In all cases the number of observations is $M = 2000$. A) Poisson networks with mean out-degree $0.25N$ {125 and 250}, with 50 % noise in the observations and B) power-law networks (with mean out-degree $0.40N$ {200 and 400}, with 10 % noise in the observations. C) The same as (B) except for wider range of the observed variables.

identifying the physical factors corresponding to $x_{\mathcal{R}}$. Yet identifying this correspondence, and thus overcoming the limitations outlined in section 4.1 can be very desirable. One practically relevant situation allowing in-depth interpretation of the results requires measuring (if only in a few configurations) the states of some of the variables that are generally unobserved ($x_{\mathcal{R}}$). This is relevant, for example, to situations in which measuring some variables is much more expensive than others (such as protein modifications versus messenger RNA concentrations) and some $x_{s \in \mathcal{R}}$ can be observed only once or a few times. Assume that the states of the k^{th} physical factor are measured (data in vector \mathbf{u}_k) in n_k number of configurations, whose indices are in the set ϕ_k . This information can be enough for determining the vertex $x_{s \in \mathcal{R}}$ corresponding to the k^{th} factor and the corresponding $\hat{\mathbf{D}}_{ss}$ as follows: 1) Compute the Pearson correlations $\bar{\rho}$ between \mathbf{u}_k and the columns of $\hat{\mathbf{R}}_{\phi_k}$. Then, the vertex of the inferred network most likely to correspond to the k^{th} physical factor is $s = \arg \max_i (|\rho_1|, \dots, |\rho_i|, \dots, |\rho_P|)$. 2)

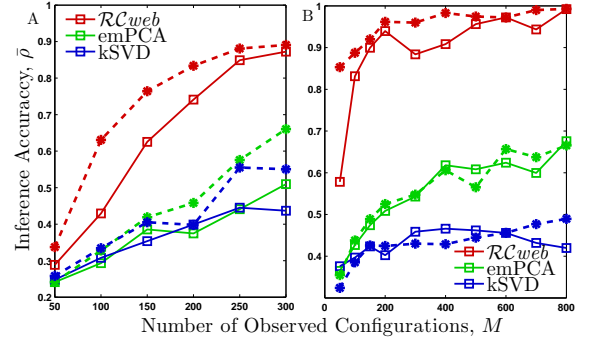


Figure 3: Accuracy of network recovery as a function of the number of observed configurations, M . Continuous lines & squares, $N = 500$; dashed lines & circles, $N = 1000$ A) Poisson networks with mean out-degree $0.20N$ {100 and 200}, with 50 % noise; B) power-law networks (with mean out-degree $0.40N$ {200 and 400}, with 10 % noise. In all cases $P = 30$.

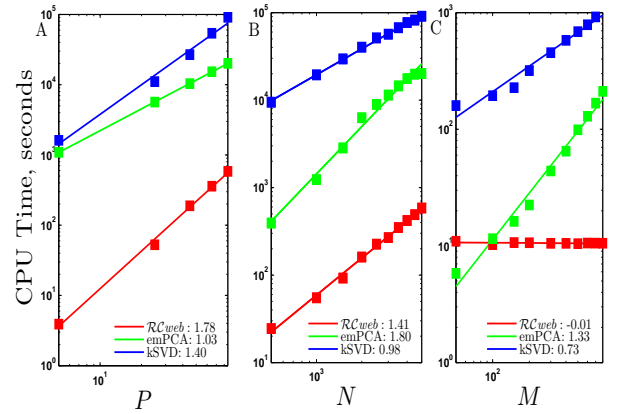


Figure 4: Computational efficiency as a function of P , M and N . The scaling exponents (slopes in log-log space) are reported in the legends. The networks are with power-law degree distribution and 60 % sparsity, with mean out-degree $0.40N$

$\hat{\mathbf{D}}_{ss} = (\hat{\mathbf{R}}_{\phi_k}^T \hat{\mathbf{R}}_{\phi_k})^{-1} \hat{\mathbf{R}}_{\phi_k}^T \mathbf{u}_k$. Similar to section (4.1), weighted matching algorithms (Sanghavi 2007) can be used for finding the optimal solution if there is data for multiple $x_{s \in \mathcal{R}}$.

Partial prior knowledge about the structure of \mathcal{G} can also be used to enhance the interpretability of \mathcal{RCweb} results. Assume, for example, that some of the nodes ($x_{s \in \alpha_k \subset \mathcal{N}}$) regulated by the k^{th} physical factor (which is a hidden variable in the inference) are known. Then the matching approach that was just outlined can be used with $\hat{\mathbf{C}}$ rather than $\hat{\mathbf{R}}$. If the weights are not known all non-zero elements of $\hat{\mathbf{C}}_{\alpha_k}$ can be set to one. The significance of the overlap (fraction of

common edges) of the regulator most likely to correspond to the k^{th} physical and its known connectivity (coming from prior knowledge) can be quantified by a p-val (the probability of observing such overlap by chance alone) computed from the hyper-geometric distribution. This approach is exemplified with gene-expression data in the next section.

4.4 INFERENCE OF A TRANSCRIPTIONAL NETWORK

Simulated models have the advantage of having known topology, and thus providing excellent basis for rigorous evaluation. Such rigorous evaluation is not possible for real biological networks. Yet, existing partial knowledge of transcriptional networks, can be combined with the approach outlined in section 4.3 for evaluation of inference of biological networks. In particular, consider the transcriptional network of budding yeast, *Saccharomyces cerevisiae*. It can be modeled as a two-layer network with transcription factors (TFs) and their complexes being regulators ($x_{\mathcal{R}}$) whose activities cannot be measured on a high-throughput scale because of technology limitations. TFs regulate the expression levels of messenger RNAs (mRNAs) whose concentrations can be measured easily on a high-throughput scale at multiple physiological states and represent the observed variables, $x_{\mathcal{N}}$. A partial knowledge about the connectivity in this transcriptional network comes from ChIP-chip experiments (MacIsaac 2006) which identify sets of genes regulated by the same transcription factors (TFs).

I first apply *RCweb* to a set mRNA configurations followed by the approach of section 4.3 to evaluate the results. The data matrix **G** contains 423 yeast datasets measured on the Affymetrix S98 microarray platform at a variety of physiological conditions. *RCweb* was initialized with $P^* = 160$ unobserved variables (corresponding to TFs) and identified $P = 153$ co-regulated sets of genes. The inferred adjacency matrix $\hat{\mathbf{C}}$ is compared directly to the adjacency matrix identified by ChIP-chip experiments (MacIsaac 2006). Indeed, sets of genes inferred by *RCweb* to be co-regulated overlap substantially with sets of genes found to be regulated by the same TF in ChIP-chip studies, Fig.5. Furthermore, the most enriched gene ontology (GO) terms associated with the *RCweb* inferred sets of genes correspond to the functions of the respective TFs. This fact provides further support to the accuracy of the inferred network topology and suggests that many of the genes inferred to be regulated by the corresponding TFs may be *bona fide* targets that were not detected in the ChIP-chip experiments because of the rather limited set of physiological conditions in which the experiments are performed (MacIsaac 2006). Since the best

Set	TF	Overlap	p-val	SGD TF Description	Set GO Terms, p-val < 1E-9
1	<i>GCN4</i>	31	7.E-12	Amino acid synthesis	Nitrogen metabolism, amino acid biosynthesis, amine biosynthesis
2	<i>FHL1</i>	37	2.E-11	rRNA processing	Translation
3	<i>RAP1</i>	37	1.E-11	Ribosomal proteins and glycolytic enzymes	Translation, biopolymer biosynthetic process
4	<i>YAP1</i>	23	1.E-11	Oxidative stress	Iron-sulfur cluster assembly, oxidation-reduction
5	<i>HAP4</i>	25	4.E-11	Respiratory gene expression	Oxidative phosphorylation, ATP synthesis coupled electron transport
6	<i>FHL1</i>	23	5.E-10	rRNA processing	Ribosome Biogenesis, ncRNA processing
7	<i>ABF1</i>	344	3.E-08	Chromatin-related events	Regulation, localization, unannotated, p-val < 1E-100

Figure 5: Overlap between the first 7 sets of genes (column 1) that *RCweb* identifies to be regulated by a common regulator and sets of genes whose promoters are bound by the same TF (column 2) in ChIP-chip experiments. The second and third columns list the number of overlapping genes between the *RCweb* set and the ChIP-chip set and the corresponding probabilities of observing such overlaps by chance alone. The fourth column contains a description of the TFs from the second column and the fifth column lists the most enriched GO terms for the *RCweb* gene sets.

“gold standard” is the only partially known topology, comparing the performance of the different algorithms in this setting can be hard to interpret; an algorithm can be penalized for correctly inferred edges if these edges are absent from the incomplete ChIP topology. An objective comparison requires experimental testing of predicted new edges and that is the subject of a forthcoming paper.

The comparison between ChIP-chip gene sets and *RCweb* gene sets allows to identify the likely correspondences between $x_{s \in \mathcal{R}}$ and TFs and to infer the activities of the TFs which are very hard to measure experimentally but crucially important to understanding the dynamics and logic of gene regulatory networks (Slavov 2009).

5 CONCLUSIONS

This paper introduces an approach (*RCweb*) for inferring latent (unobserved) factors explaining the behavior of observed variables. *RCweb* aims at inferring a sparse bipartite graph in which vertices connect inferred latent factors (e.g. regulators of mRNA transcription and degradation) to observed variables (e.g. target mRNAs). The salient difference distinguishing *RCweb* from prior related work is a new approach to attaining sparse solution that allows the natural inclusion of a generative model, relaxation of assumptions on distributions, and ultimately results in more accurate and computationally efficient inference compared to competing algorithms for sparse data decomposition.

Acknowledgements

I thank Dmitry Malioutov and Rodolfo Ros Zertuche for many insightful discussions and extensive, constructive feedback, as well as David M. Blei and Kenneth A. Dawson for support and comments. This work was supported by Irish Research Council for Science, Engineering and Technology (IRCSET) and by Science Foundation Ireland under Grant No.[SFI/SRC/B1155].

References

- West M. (2002) Bayesian factor regression models in the large p , small n paradigm. *Bayesian Stat*, **7**, pp. 723–732
- Carvalho C. et al. (2008) High-dimensional sparse factor modelling: Applications in gene expression genomics. *JASA*, **103**, 484, pp. 1438–1456
- Dueck D., Morris Q., Frey, B. (2005) Multi-way clustering of microarray data using probabilistic sparse matrix factorization. *Bioinformatics*, **21**, pp. 144–151.
- d’Aspremont A, Bach F., Ghaoui L. (2007) Optimal Solutions for Sparse Principal Component Analysis. *JMLR* **9**, pp. 1269–1294
- Sigg C., Buhmann J. (2008) Expectation-Maximization for Sparse and Non-Negative PCA. *ICML*, Helsinki, Finland
- Srebro N. and Jaakkola T. (2001) Sparse Matrix Factorization for Analyzing Gene Expression Patterns, *NIPS*
- Pe’er D., Regev A., and Tanay A. (2002) Minreg: Inferring an active regulator set. *Bioinformatics*, **18**, 258–267
- Bradley E., Trevor H., Iain J., Tibshirani R. (2004) Least Angle Regression. *Annals of Statistics* **32**, pp. 407–499
- Aharon M., Elad M., Bruckstein A. (2005) K-SVD and its non-negative variant for dictionary design. *Wavelets XI* **5914**, pp. 591411–591424.
- Banerjee, O., Ghaoui, L. E. & d’Aspremont, A. (2007) Model selection through sparse maximum likelihood estimation *JMLR*, **9**, pp. 485–516
- Cetin M., Malioutov D., Willsky A. (2002) A Variational Technique for Source Localization based on a Sparse Signal Reconstruction Perspective. *ICASSP*, **3**, pp. 2965–2968, Orlando, Florida
- Candès E., Wakin M. and Boyd S. (2007) Enhancing sparsity by reweighted ℓ_1 minimization. *J. Fourier Anal. Appl.*, **14**, pp. 877–905.
- Wilkinson, J. H. (1988) The Algebraic Eigenvalue Problem, *Oxford University Press, USA*, ISBN-13: 978-0198534181
- Benaych-Georges F., Rao R. (2009) The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices, arXiv:0910.2120v1
- Sanghavi S., Malioutov D., Willsky A. (2007) Linear programming analysis of loopy belief propagation for weighted matching, *NIPS*
- MacIsaac K.D., et al. (2006) An improved map of conserved regulatory sites for *Saccharomyces cerevisiae*. *BMC Bioinformatics*, **7**:113. doi: 10.1186/1471-2105-7-113.
- Slavov N., Dawson KA. (2009) Correlation Signature of the Macroscopic States of the Gene Regulatory Network in Cancer, *PNAS*, **106**, 11, pp. 4079–4084