



Spectral analysis of gene expression profiles using gene networks

Franck Rapaport, Andrei Zinovyev, Marie Dutreix, Emmanuel Barillot,
Jean-Philippe Vert

► To cite this version:

Franck Rapaport, Andrei Zinovyev, Marie Dutreix, Emmanuel Barillot, Jean-Philippe Vert. Spectral analysis of gene expression profiles using gene networks. 2006. hal-00021785

HAL Id: hal-00021785

<https://hal.science/hal-00021785>

Preprint submitted on 26 Mar 2006

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.

Spectral analysis of gene expression profiles using gene networks

Franck Rapaport
Center for Computational Biology
Ecole des Mines de Paris
and Service de Bioinformatique
Institut Curie
Franck.Rapaport@curie.fr

Andrei Zinovyev
Service de Bioinformatique
Institut Curie
Andrei.Zinovyev@curie.fr

Marie Dutreix
CNRS-UMR 2027
Institut Curie
Marie.Dutreix@curie.fr

Emmanuel Barillot
Service de Bioinformatique
Institut Curie
Emmanuel.Barillot@curie.fr

Jean-Philippe Vert
Center for Computational Biology
Ecole des Mines de Paris
Jean-Philippe.Vert@ensmp.fr

March 26, 2006

Abstract

Microarrays have become extremely useful for analysing genetic phenomena, but establishing a relation between microarray analysis results (typically a list of genes) and their biological significance is often difficult. Currently, the standard approach is to map *a posteriori* the results onto gene networks to elucidate the functions perturbed at the level of pathways. However, integrating *a priori* knowledge of the gene networks could help in the statistical analysis of gene expression data and in their biological interpretation. Here we propose a method to integrate *a priori* the knowledge of a gene network in the analysis of gene expression data. The approach is based on the spectral decomposition of gene expression profiles with respect to the eigenfunctions of the graph, resulting in an attenuation of the high-frequency components of the expression profiles with respect to the topology of the graph. We show how to derive unsupervised and supervised classification algorithms of expression profiles, resulting in classifiers with biological relevance. We applied the method to the analysis of a set of expression profiles from irradiated and non-irradiated yeast strains. It performed at least as well as the usual classification but provides much more biologically relevant results and allows a direct biological interpretation.

Supplementary material: <http://bioinfo.curie.fr/projects/kernelchip>

1 Introduction

During the last decade microarrays have become the technology of choice for dissecting the genes responsible for a phenotype. By monitoring the activity of virtually all the genes from a sample in a single experiment they offer a unique perspective for explaining the global genetic picture of a variant, whether a diseased individual or a sample subject to whatever stressing conditions.

However, this strength is also the major weakness of microarrays, and has led to "gene list" syndrome. With careful experimental design and data analysis, microarrays often present a list of genes that are differentially expressed between two conditions, or allow samples to be classified according to their known phenotypic features. Once obtained, the meaning behind this list of genes, being typically a few hundreds, must be deciphered. Normally, the biologist relies on experience for making sense out of the results. Microarrays are prone to certain errors, and a single gene showing differential expression may not necessarily be involved in the process under study. More fundamentally,

when it is clear that a gene is involved, the biologist is often more interested in the biological function than in the single gene. The behaviour of a biological function usually involves more than just the expression of a single gene, but involves the entire set of genes that form the regulation pathways governing the function, which have to be examined.

One of the aims of microarray data analysis is to discover co-operating genes that are similarly affected during the experiment. Many databases and tools help verify this *a posteriori*. For example, Gene Ontology [Consortium, 2000], Biocarta (www.biocarta.com), GenMAPP (www.genmapp.org) and KEGG [Kanehisa et al., 2004] all allow a list of genes to be crossed with genetic networks, including metabolic, signalling or other regulation pathways. Basic statistical analysis (see, e. g., Hosack et al., 2003) can then determine whether a pathway is over-represented in the list, and whether it is over-activated or under-activated. However, it is obvious that introducing information on the pathway at this point in the analysis process sacrifices some statistical power to the simplicity of the approach. For example, if we assume that genes in the same pathway show a coherent expression pattern, assessing differential gene expression could be made more efficient; coherent expression patterns in a regulation pathway would be defined as those in which over-expressed inhibitors or activators result in respectively under-expressed or overexpressed targets, or vice versa. Thus, a small but coherent difference in the expression of all the genes in a pathway should be more significant than larger but random differences.

There is therefore a pressing need for methods integrating *a priori* pathway knowledge in the gene expression analysis process, and several attempts have been carried out in that direction so far. An initial approach is to derive a model for gene expression from an *a priori* model of gene networks; for example, by deriving constraints on co-regulated gene expression. Logical discrete formalism [Thomas and Kaufman, 2001] can be used to analyse all the possible steady states of a biochemical reaction network described by positive and negative influences and can determine whether the observed gene expression may be explained by a perturbation of the network. If only the signs of the concentration differences between two steady states are considered, it is possible to solve the corresponding Laplace equation in sign algebra [Radulescu et al., 2005], giving qualitative predictions for the signs of the concentration differences measured by microarrays. Other approaches, such as the MetaReg formalism [Gat-Viks et al., 2004], and boolean and Bayesian networks [Akutsu et al., 2000, Friedman et al., 2000] have also been used to predict possible gene expression patterns from the network structure, although these approaches adhere less to the formal theory of biochemical reaction networks.

Unfortunately, methods based on network models are rarely satisfactory because detailed quantitative knowledge of the complete reaction network parameters is often lacking, or there are only fragments of the network structure available. In these cases, more phenomenological approaches need to be used. Pathway scoring methods try to detect perturbed "modules" or network pathways while ignoring the detailed network topology (for recent reviews see Curtis et al., 2005, Cavalieri and De Filippo, 2005). It is assumed that the genes inside a module are co-ordinately expressed, and thus a perturbation is likely to affect many of them.

With available databases containing tens of thousands reactions and interactions (KEGG [Kanehisa et al., 2004], TransPath [Krull et al., 2006], BioCyc [Karp et al., 2005], Reactome [Joshi-Tope et al., 2005] and others), the problem is how to integrate the detailed graph of gene interactions (and not just crude characteristics such as the inter/intra-module connectivity) into the core microarray data analysis. Some promising results have been reported with regard to this problem. Vert and Kanehisa [2003] developed a method for correlating interaction graphs and different types of quantitative data, and Rahnenfuhrer et al. [2004] showed that explicitly taking the pathway distance between pairs of genes into account enhances the statistical scores when identifying activated pathways. Hanisch et al. [2002] reported the co-clustering of gene expression and gene networks, and the PATIKA project [Babur et al., 2004] proposed quantifying the compatibility of a pathway with a given microarray data by scoring individual interactions.

In this paper, we investigate a different approach for integrating genetic networks early in the gene expression analysis. We propose a method for calculating the eigen modes of the response of the gene network to a perturbation and suggest how these can be introduced into supervised and unsupervised microarray data analysis. The method uses a hierarchical approach to automatically divide the network into modules having co-ordinated responses. We are able to filter out high-frequency noisy modes, thus increasing our ability to interpret expression profiles from the underlying genetic network. We illustrate the relevance of our approach by analysing a gene expression dataset that monitors the transcriptional response of irradiated and non-irradiated yeast colonies [Mercier et al., 2004]. We show that by filtering out 80% of the eigen modes of the KEGG metabolic network in the gene expression, we obtain accurate and interpretable discriminative model that may lead to new biological insights.

2 Methods

In this section, we will explain how a gene expression can be decomposed into the eigen modes of a gene network, and how to derive unsupervised and supervised classification algorithms from this decomposition.

2.1 Spectral decomposition of gene expression profiles

We consider a finite set of genes V of cardinality $|V| = n$. The available gene network is represented by an undirected graph $G = (V, E)$ without loop and multiple edges, in which the set of vertices V is the set of genes and $E \subset V \times V$ is the list of edges. We will use the notation $u \sim v$ to indicate that two genes u and v are neighbors in the graph, that is, $(u, v) \in E$. For any gene u , we denote the degree of u in the graph by d_u , that is, its neighbour number. Gene expression profiling gives a value of expression $f(u)$ for each gene u , and is therefore represented by a function $f : V \rightarrow \mathbb{R}$.

The Laplacian of the graph G is the $n \times n$ matrix [Chung, 1997]:

$$\forall u, v \in V, \quad L(u, v) = \begin{cases} d_u & \text{if } u = v, \\ -1 & \text{if } u \sim v, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The Laplacian is a central concept in spectral graph theory [Mohar, 1997] and shares many properties with the Laplace operator on Riemannian manifolds. L is known to be symmetric positive definite and singular. We denote its eigenvalues by $0 = \lambda_1 \leq \dots \leq \lambda_n$ and the corresponding eigenvectors by e_1, \dots, e_n . The multiplicity of 0 as an eigenvalue is equal to the number of connected components of the graph, and the corresponding eigenvectors are constant on each connected component. The eigen-basis of L forms a Fourier basis and a natural theory for Fourier analysis and spectral decomposition on graphs can thus be derived [Chung, 1997]. Essentially, the eigenvectors with increasing eigenvalues tend to vary more abruptly on the graph, as the smoothest functions (constant on each connected component) are associated with the smallest (zero) eigenvalue. In particular, the Fourier transform $\hat{f} \in \mathbb{R}^n$ of any expression profile f is defined by:

$$\hat{f}_i = \sum_{u \in V} e_i(u) f(u), \quad i = 1, \dots, n.$$

Like its continuous counterpart, the discrete Fourier transform can be used for smoothing or for extracting features. Here, our hypothesis is that analysing a gene expression profile from its Fourier transform on an *a priori* given gene network should allow the profile to be further analysed through this prior knowledge. In the next two sections we illustrate the potential applications of this approach by describing how this leads to a natural definition for distances between expression profiles, and how this distance can be used for classification or regression purposes.

2.2 Deriving a metric for expression profiles

The definition of new metrics on expression profiles that incorporate information encoded in the graph structure is a first possible application of the spectral decomposition. Following the classical methodology in Fourier analysis, we assume that the signal captured in the low-frequency component of the expression profiles contains the most biologically relevant information, particularly the general expression trends, whereas the high-frequency components are more likely measurement noise. For example, the low frequency component of an expression vector on the gene metabolic network should reveal areas of positive and negative expression on the graph that are likely to correspond to the activation or inhibition of pathways. We can translate this idea mathematically by considering the following class of transformations for expression profiles:

$$\forall f \in \mathbb{R}^V, \quad S_\phi(f) = \sum_{i=1}^n \hat{f}_i \phi(\lambda_i) e_i, \quad (2)$$

where $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}$ is a non-increasing function that quantifies how each frequency is attenuated. For example, if we take $\phi(\lambda) = 1$ for all λ , the profile does not change that is, $S_\phi(f) = f$. However, if we take:

$$\phi_{\text{thres}}(\lambda) = \begin{cases} 1 & \text{if } 0 \leq \lambda \leq \lambda_0, \\ 0 & \text{if } \lambda > \lambda_0, \end{cases} \quad (3)$$

we produce a low-pass filter that removes all frequencies from f above the threshold λ_0 . Finally, a function of the form:

$$\phi_{\text{exp}}(\lambda) = \exp(-\beta\lambda), \quad (4)$$

for some $\beta > 0$, keeps all the frequencies but strongly attenuates the high-frequency components.

If $S_\phi(f)$ includes the biologically relevant part of the expression profile, we can compare two expression profiles f and g through their representations $S_\phi(f)$ and $S_\phi(g)$. This leads to the following metric between the profiles:

$$\begin{aligned} d_\phi(f, g)^2 &= \|S_\phi(f) - S_\phi(g)\|^2 \\ &= \sum_{i=1}^n (\hat{f}_i - \hat{g}_i)^2 \phi(\lambda_i)^2. \end{aligned}$$

The associated norm is the Euclidean norm associated with the inner product:

$$\begin{aligned} \langle f, g \rangle_\phi &= \sum_{i=1}^n \hat{f}_i \hat{g}_i \phi(\lambda_i)^2 \\ &= \sum_{i=1}^n f^\top e_i e_i^\top g \phi(\lambda_i)^2 \\ &= f^\top K_\phi g, \end{aligned}$$

where $K_\phi = \sum_{i=1}^n \phi(\lambda_i)^2 e_i e_i^\top$ is the positive semidefinite matrix obtained by modifying the eigenvalues of L through ϕ . For example, taking $\phi(\lambda) = \exp(-\beta\lambda)$ leads to $K_\phi = \exp_M(-\beta L)$, where \exp_M denotes the matrix exponential.

2.3 Supervised learning and regression

The construction of predictive models for a property of interest from the gene expression profiles of the studied samples is a second possible application of the spectral decomposition of expression profiles on the gene network. Typical applications include predicting cancer diagnosis or prognosis from gene expression data, or discriminating between different treatments applied to micro-organisms. Most approaches build predictive models from the gene expression alone, and then check whether the predictive model is biologically relevant by studying, for example, whether genes with high weights are located in similar pathways. However the genes often give no clear biological meaning. Here, we propose a method combining both steps in a single predictive model that is trained by forcing some form of biological relevance.

We use linear predictive models to predict a variable of interest y from an expression profile f that are obtained by solving the following optimisation problem:

$$\min_{w \in \mathbb{R}^n} \sum_{i=1}^p l(w^\top f_i, y_i) + C \|w\|^2, \quad (5)$$

where $(f_1, y_1), \dots, (f_p, y_p)$ is a training set of profiles containing the variable y to be predicted, and l is a loss function that measures the cost of predicting $w^\top f_i$ instead of y_i . For example, the popular support vector machine [Boser et al., 1992] is a particular case of equation (5) in which y can take values in $-1, +1$ and $l(u, y) = \max(0, 1 - yu)$ is the hinge loss function; ridge regression is obtained for $y \in \mathbb{R}$ by taking $l(u, y) = (u - y)^2$ [Hastie et al., 2001].

Here, we do not apply algorithms of the form (5) directly to the expression profiles f , but to their images $S_\phi(f)$. That is, we consider the problem:

$$\min_{w \in \mathbb{R}^n} \sum_{i=1}^p l(w^\top S_\phi(f_i), y_i) + C \|w\|^2. \quad (6)$$

Indeed, for any $w \in \mathbb{R}^d$, let $v = K_\phi^{1/2} w$. We observe that for any $f \in \mathbb{R}^n$:

$$\begin{aligned} w^\top S_\phi(f) &= w^\top \sum_{i=1}^n \hat{f}_i \phi(\lambda_i) e_i \\ &= f^\top \sum_{i=1}^n e_i \phi(\lambda_i) e_i^\top w \\ &= f^\top K_\phi^{1/2} w \\ &= f^\top v, \end{aligned}$$

showing that the final predictor obtained by minimizing (6) is equal to $v^\top f$. We also note that:

$$\begin{aligned} \|w\|^2 &= w^\top w \\ &= v^\top K_\phi^{-1} v \\ &= \sum_{i=1}^n \frac{\hat{v}_i^2}{\phi(\lambda_i)^2}, \end{aligned}$$

where the last equality remains valid if K_ϕ is not invertible simply by not including in the sum the terms i for which $\phi(\lambda_i) = 0$. This shows that (6) is the equivalent of solving the following problem in the original space:

$$\min_{v \in \mathbb{R}^n} \sum_{i=1}^p L(v^\top f_i, y_i) + C \sum_{i: \phi(\lambda_i) > 0} \frac{\hat{v}_i^2}{\phi(\lambda_i)^2}. \quad (7)$$

Thus, the resulting algorithm amounts to finding a linear predictor v that minimizes the loss function of interest l regularised by a terms that penalises the high-frequency components of v . This is different to the classical regularisation $\|v\|^2$ that only focuses on the norm of v . As a result, the linear predictor v can be made smoother on the gene network by increasing the parameter C . This allows the prior knowledge to be directly included because genes in similar pathways would be expected to contribute similarly to the predictive model.

There are two consequences of this procedure. First, if the true predictor really is smooth on the graph, the formulation (7) can help the algorithm focus on plausible models even with very little training data, resulting in a better estimation. As a result, we can expect a better predictive performance. Second, by forcing the predictive model v to be smooth on the graph, biological interpretation of the model should be easier by inspecting the areas of the graph in which the predictor is strongly positive or negative. Thus the model should be easier to interpret than models resulting from the direct optimisation of equation (5).

3 Data

We collected the expression data from a study analysing the effect of low irradiation doses on *Saccharomyces cerevisiae* strains [Mercier et al., 2004]. The first group of extracted expression profiles was a set of twelve independent yeast cultures grown without radiation (not irradiated, NI). From this group, we excluded an outlier that was indicated to us by the author of the article. The second group was a set of six independent irradiated (I) cultures exposed to a dose of 15-20 mGu/h for 20h. This dose produces no mutagenic effects, but induces transcriptional changes. We used the same normalisation method as in the first study of this data (Splus LOWESS function, see Mercier et al., 2004 for details) and attempted (1) to separate the NI samples from the I ones, and (2) to understand the difference between the two populations in terms of metabolic pathways.

The gene network model used to analyse the gene expression data was therefore built from the KEGG database of metabolic pathways [Kanehisa et al., 2004]. The metabolic gene network is a graph in which the enzymes are vertices and the edges between two enzymes indicate that the product of a reaction catalysed by the first enzyme is the substrate of the reaction catalysed by the second enzyme. We reconstructed this network from the KGML v0.3 version of KEGG, resulting in 4694 edges between 737 genes. We kept only the largest connected component (containing 713 genes) for further spectral analysis.

4 Results

4.1 Unsupervised classification

First, we tested the general effect of modifying the distances between expression profiles using the KEGG metabolic pathways as background information in an unsupervised setting. We calculated the pairwise distances between all 17 expression profiles after applying the transformations defined by the filters (3) and (4), over a wide range of parameters. We assessed whether the resulting distances were more coherent with a biological interpretation by calculating the ratio of intraclass distances over all pairwise distances, defined by:

$$r = \frac{\sum_{u_1, v_1 \in V_1} d(u_1, v_1)^2 + \sum_{u_2, v_2 \in V_2} d(u_2, v_2)^2}{\sum_{u, v \in V} d(u, v)^2},$$

where V_1 and V_2 are the two classes of points. We compared the results with those obtained by replacing KEGG with a random network, produced by keeping the same graph structure but randomly permutating the vertices, in order to assess the significance of the results. We generated 100 such networks to give an average result with a standard deviation. Figure 1 shows the result for the function $\phi_{\text{exp}}(\lambda) = \exp(-\beta\lambda)$ with varying β (left), and for the function $\phi_{\text{thres}}(\lambda) = 1(\lambda < \lambda_0)$ for varying λ_0 (right). We found that, apart from very small values of β , the change of metric with the ϕ_{exp} function performed worse than that of a random network. The second method (filtering out the high frequency components of the gene expression vector), in which up to 80% of the eigenvectors were removed, performed significantly better than that of a random network. When only the top 3% of the smoothest eigenvectors were kept, the performance was similar to that of a random network, and when only the top 1% was kept, the performance was significantly worse. This explains the disappointing results obtained with the ϕ_{exp} function: by giving more weight to the small eigenvalues exponentially, the method focuses on those first few eigenvectors that, as shown by the second method, do not provide a geometry compatible with the separation of samples into two classes. From the second plot, we can infer that at least 20% of the KEGG eigenvectors should be given sufficient weight to obtain a geometry compatible with the classification of the data in this case.

4.2 PCA analysis

We carried out a Principal Component Analysis (PCA, Jolliffe, 1996) on the original expression vectors f and compared this with a PCA of the transformed set of vectors $S_\phi(f)$ obtained with the function ϕ_{thres} to further investigate the effect on their relative positions of filtering out the high frequencies of the expression profiles.

Analysis of the initial sample distribution (see fig.2) showed that the first principal component can partially separate irradiated from non-irradiated samples, with exception of the two irradiated samples "I1" and "I2". They have larger projections onto the third principal component than onto the first one. The experimental protocol revealed that these two samples were affected by higher doses of radiation than the four other samples.

Gene Ontology analysis of the genes that gave the biggest contribution in the first principal component showed that "pyruvate metabolism", "glucose metabolism", "carbohydrate metabolism", and "ergosterol biosynthesis" ontologies (here we list only independent ontologies) were over-represented (with p-values less than 10^{-10}). The second component was associated with "trehalose biosynthesis", and "carboxylic acid metabolism" ontologies and the third principal component was associated with the KEGG glycolysis pathway. The first three principal components collected 25%, 17% and 11% of the total dispersion.

The transformation (2) resulting from a step-like attenuation of eigenvalues ϕ_{thres} removing 80% of the largest eigenvalues significantly changed the global layout of data (fig. 2, right) but generally preserved the local neighbourhood relationships. The first three principal components collected 28%, 20% and 12% of the total dispersion, which was only slightly higher than the PCA plot of the initial profiles. The general tendency was that the non-irradiated normal samples were more closely grouped, which explains the lower intraclass distance values shown in fig.1. The principal components in this case allowed them to be associated with gene ontologies with higher confidence (for the first component, the p-values are less than 10^{-25}). This is a direct consequence of the fact that the principal components are constrained to belong to a subspace of smooth functions on KEGG, giving coherence in terms of pathways to the genes contributing to the components. The first component gave "DNA-directed RNA polymerase activity", "RNA polymerase complex" and "protein kinase activity". Fig.4 shows that these are the most connected clusters of the whole KEGG network. The second component is associated with "purine ribonucleotide metabolism", "RNA polymerase complex", "carboxylic acid metabolism" and "acetyl-CoA metabolism" ontologies and also with




Figure 1: Performance of the unsupervised classification after changing the metric with the function $\phi(\lambda) = \exp(-\beta\lambda)$ for different values of β (left), or with the function $\phi(\lambda) = 1(\lambda < \lambda_0)$ which varying λ_0 , that is, by keeping a variable number of smallest eigenvalues (right).. The red curve is obtained with the KEGG network. The black curves show the result (mean and one standard deviation interval) obtained with a random network.

"Glycolysis/Gluconeogenesis", "Citrate cycle (TCA cycle)" and "Reductive carboxylate cycle (CO₂ fixation)" KEGG pathways. The third component was associated with "prenyltransferase activity", "lyase activity" and "aspartate family amino acid metabolism" ontologies and with "N-Glycan biosynthesis", "Glycerophospholipid metabolism", "Alanine and aspartate metabolism" and "riboflavin metabolism" KEGG pathways. Thus, the PCA components of the transformed expression profiles were affected both by network features and by the microarray data.

4.3 Supervised classification

We tested the performance of supervised classification after modifying the distances with a support vector machine (SVM) trained to discriminate irradiated samples from non-irradiated samples. For each change of metric, we estimated the performance of the SVM from the total number of misclassifications and the total hinge loss using a "leave-one-out" (LOO) approach. This approach removes each sample in turn, trains a classifier on the remaining samples and then tests the resulting classifier on the removed sample. For each fold, the regularization parameter was selected from the training set only by minimising the classification error estimated with an internal LOO experiment. The calculations were carried out using the `svmpath` package in the R computing environment.

Figure 3 shows the classification results for the two high frequency attenuation functions ϕ_{exp} and ϕ_{thres} with varying parameters. The baseline LOO error was 2 misclassifications for the SVM in the original Euclidean space. For the exponential variant ($\phi_{\text{exp}}(\lambda) = \exp(-\beta\lambda)$), we observed an irregular but certain degradation in performance for positive β for both the hinge loss and the misclassification number. This is consistent with the result shown in fig.1 in which the change of metric towards the first few eigenvectors does not give a geometry coherent with the classification of samples into irradiated and non-irradiated, resulting in a poorer performance in supervised classification as well. For the second variant, in which the expression profiles were projected onto the eigenvector of the graph with the smallest eigenvalues, we found that the performance remained as accurate as the baseline performance until up to 80% of the eigenvectors were discarded, with the hinge loss even exhibiting a slight minimum in this region. This is consistent with the classes being more clustered in this case than in the original Euclidean space. Overall these results show that classification accuracy can be kept high even when the classifier is constrained to exhibit a certain coherence

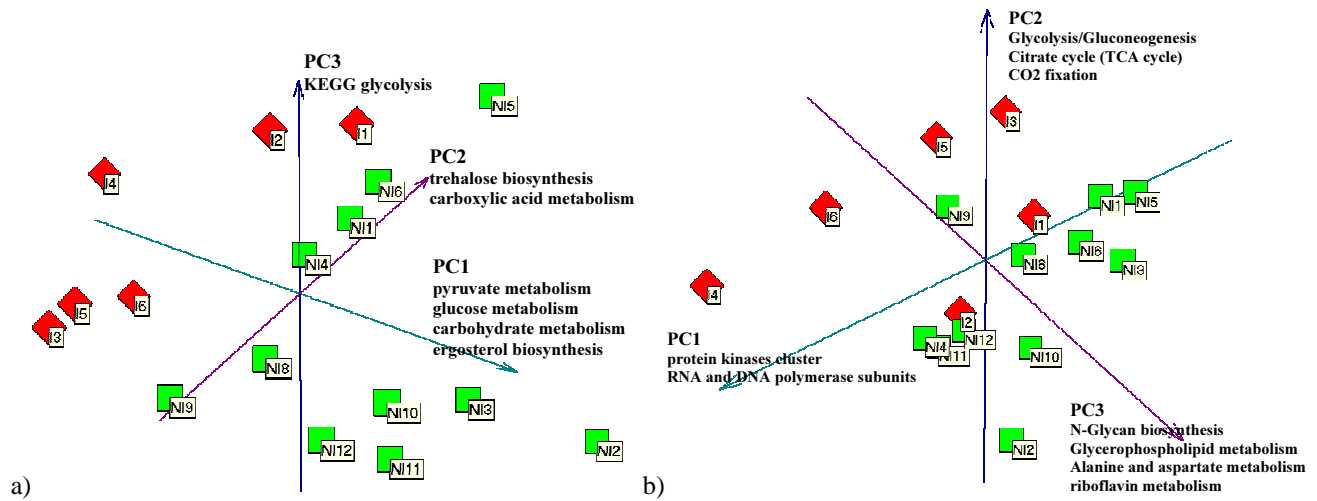


Figure 2: PCA plots of the initial expression profiles (a) and the transformed profiles using network topology (80% of the eigenvalues removed) (b). The green squares are non-irradiated samples and the red rhombuses are irradiated samples. Individual sample labels are shown together with GO and KEGG annotations associated with each principal component.

with the graph structure.

4.4 Interpretation of the SVM classifier

Figure (4) shows the global connection map of KEGG generated from the connection matrix by Cytoscape software [Shannon et al., 2003]. The coefficients of the decision function v of equation (7) for the classifier constructed either in the original Euclidean space or after filtering the 80% top spectral components of the expression profiles are shown in colour. We used a color scale from green (negative weights) to red (positive weights) to provide an easy visualization of the main features of the classifiers. Both classifiers give the same classification error but the classifier constructed using the network structure can be more naturally interpreted, because the classifier variables are grouped according to their participation in the network modules.

Although from a biological point of view, very little can be learned from the classifier obtained in the original Euclidean space (figure 4, left), it is indeed possible to distinguish several features of interest for the classifier obtained in the second case (figure 4, right). First, oxidative phosphorylation was found among the pathways with the most positive weights, which is consistent with previous analyses showing that this pathway tended to be up-regulated after irradiation [Mercier et al., 2004]. An important cluster involving the DNA and RNA polymerases is also found to bear weights slightly above average in these experiments. Several studies have previously reported the induction of genes involved in replication and repair after high doses of irradiation [Mercier et al., 2001], but the detection of such an induction at the low irradiation doses used in the present biological experiments is rather interesting. The strongly negative landscape of weights in the protein kinases cluster has not been seen before and may lead to a new area of biological study. Most of the kinases are involved in signalling pathways, and therefore their low expression levels may have important biological consequences.

Figure 4 shows a highlighted pathway named "Glycolysis/Gluconeogenesis" in KEGG. A more detailed view of this pathway is shown in fig.5. This pathway contains enzymes that are also used in many other KEGG pathways and is therefore situated in the middle and most entangled part of the global network. As already mentioned, this pathway is associated with the first and the third principal components of the initial dataset. The pathway actually contains two alternative sub-pathways that are affected differentially. Over-expression in the gluconeogenesis pathway seems to be characteristic of irradiated samples, whereas glycolysis has a low level of expression in that case. This shift can be observed by changing from anaerobic to aerobic growth conditions (called diauxic shift). The reconstruction of this from our data with no prior input of this knowledge strongly confirms the relevance of our analysis method. It also shows that analysing expression in terms of the global up- or down-regulation of entire pathways as defined, for example, by KEGG, could mislead as there are many antagonist processes that take place inside pathways. Representing




Figure 3: Performance of the supervised classification when changing the metric with the function $\phi_{\text{exp}}(\lambda) = \exp(-\beta\lambda)$ for different values of β (left picture), or the function $\phi_{\text{thres}}(\lambda) = 1(\lambda < \lambda_0)$ for different values of λ_0 (i.e., keeping only a fraction of the smallest eigenvalues, right picture). The performance is estimated from the number of misclassifications in a leave-one-out error.

KEGG as a large network helps keep the biochemical relationships between genes without the constraints of pathway limits.

5 Discussion

Our algorithm groups predictor variables according to highly connected "modules" of the global gene network. We assume that the genes within a tightly connected network module are likely to contribute similarly to the prediction function because of the interactions between the genes. This motivates the filtering of gene expression profile to remove the noisy high-frequency modes of the network.

Such grouping of variables is a very useful feature of the resulting classification function because the function becomes meaningful for interpreting and suggesting biological factors that cause the class separation. This allows classifications based on functions, pathways and network modules rather than on individual genes. This can lead to a more robust behaviour of the classifier in independent tests and to equal if not better classification results. Our results on the dataset we analysed shows only a slight improvement, although this may be due to its limited size. Therefore we are currently extending our work to larger data sets.

An important remark to bear in mind when analyzing pictures such as fig.4 and 5 is that the colors represent the weights of the classifier, and not gene expression levels. There is of course a relationship between the classifier weights and the typical expression levels of genes in irradiated and non-irradiated samples: irradiated samples tend to have expression profiles positively correlated with the classifier, while non-irradiated samples tend to be negatively correlated. Roughly speaking, the classifier tries to find a smooth function that has this property. If more samples were available, better non-smooth classifier might be learned by the algorithm, but constraining the smoothness of the classifier is a way to reduce the complexity of the learning problem when a limited number of samples are available. This means in particular that the pictures provide virtually no information regarding the over- or under-expression of individual genes, which is the cost to pay to obtain instead an interpretation in terms of more global pathways. Constraining the classifier to rely on just a few genes would have a similar effect of reducing the complexity of the




Figure 4: Global connection map of KEGG with mapped coefficients of the decision function obtained by applying a customary linear SVM (**left**) and using high-frequency eigenvalue attenuation (80% of high-frequency eigenvalues have been removed) (**right**). Spectral filtering divided the whole network into modules having coordinated responses, with the activation of low-frequency eigen modes being determined by microarray data. Positive coefficients are marked in red, negative coefficients are in green, and the intensity of the colour reflects the absolute values of the coefficients. Rhombuses highlight proteins participating in the Glycolysis/Gluconeogenesis KEGG pathway. Some other parts of the network are annotated including big highly connected clusters corresponding to protein kinases and DNA and RNA polymerase sub-units.

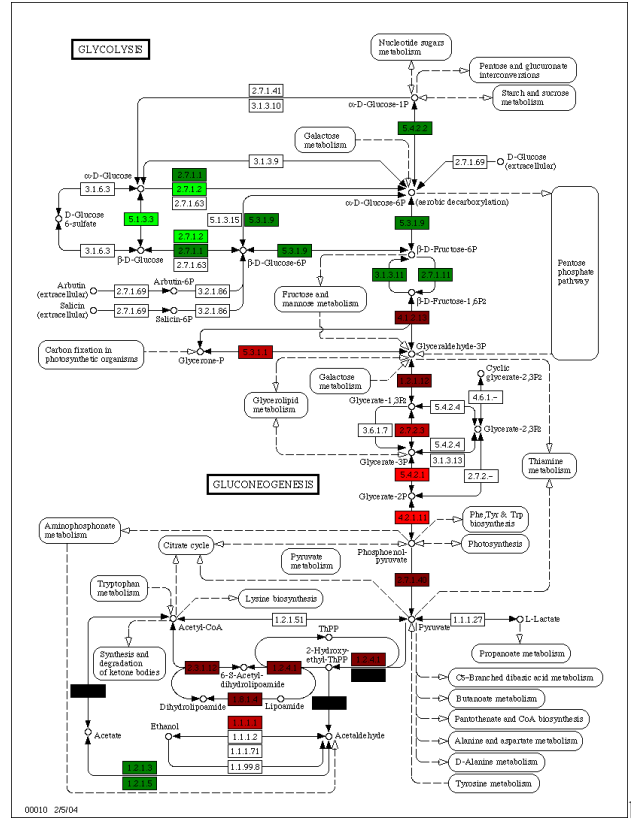
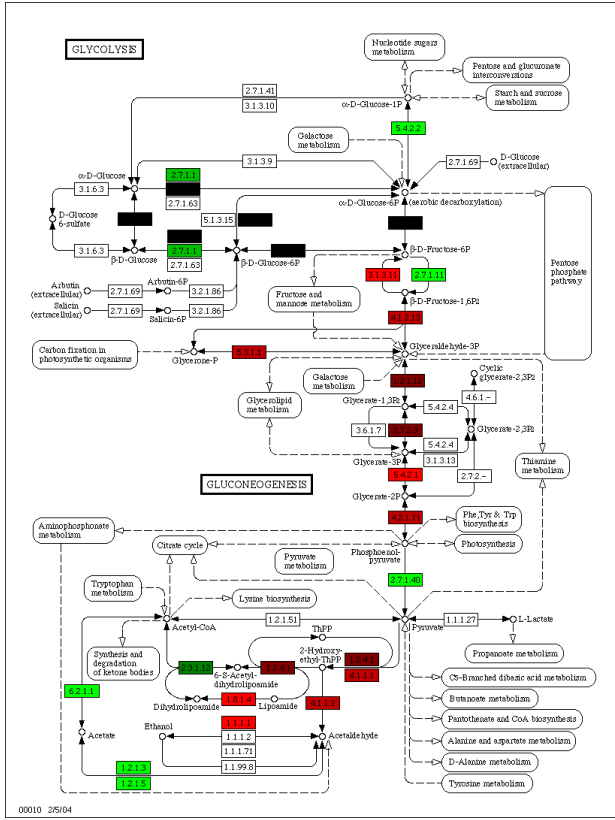


Figure 5: The glycolysis/gluconeogenesis pathways of KEGG with mapped coefficients of the decision function obtained by applying a customary linear SVM (a) and using high-frequency eigenvalue attenuation (b). The pathways are mutually exclusive in a cell, as clearly highlighted by our algorithm.

problem, but would lead to a more difficult interpretation in terms of pathways.

An advantage of our approach over other pathway-based clustering methods is that we consider the network modules that naturally appear from spectral analysis rather than a historically defined separation of the network into pathways. Thus, pathways cross-talking is taken into account, which is difficult to do using other approaches. It can however be noticed that the implicit decomposition into pathways that we obtain is biased by the very incomplete knowledge of the network and that certain regions of the network are better understood, leading to a higher connection concentration.

Like most approaches aiming at comparing expression data with gene networks such as KEGG, the scope of this work is limited by two important constraints. First the gene network we use is only a convenient but rough approximation to describe complex biochemical processes; second, the transcriptional analysis of a sample can not give any information regarding post-transcriptional regulation and modifications. Nevertheless, we believe that our basic assumptions remain valid, in that we assume that the expression of the genes belonging to the same metabolic pathways module are coordinately regulated. Our interpretation of the results supports this assumption.

Another important caveat is that we simplify the network description as an undirected graph of interactions. Although this would seem to be relevant for simplifying the description of metabolic networks, real gene regulation networks are influenced by the direction, sign and importance of the interaction. Although the incorporation of weights into the Laplacian (equation 1) is straightforward and allows the extension of the approach to weighted undirected graphs, the incorporation of directions and signs to represent signalling or regulatory pathways requires more work but could lead to important advances for the interpretation of microarray data in cancer studies, for example.

Acknowledgments

This work was supported by the grant ACI-IMPBIO-2004-47 of the French Ministry for Research and New Technologies. We thank Sabrina Carpentier from the Service de Bioinformatique of the Institut Curie for the help she provided with the normalisation of the microarray data.

References

- T. Akutsu, S. Miyano, and S. Kuhara. Inferring qualitative relations in genetic networks and metabolic pathways. *Bioinformatics*, 16(8):727–734, 2000.
- O Babur, E Demir, A Ayaz, U Dogrusoz, and O Sakarya. Pathway activity inference using microarray data. *Technical report, Bilkent Center for Bioinformatics (BCBI)*, 2004.
- B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th annual ACM workshop on Computational Learning Theory*, pages 144–152. ACM Press, 1992.
- D. Cavalieri and C. De Filippo. Bioinformatic methods for integrating whole-genome expression results into cellular networks. *Drug Discov Today*, 10(10):727–34, 2005.
- F. R. K. Chung. *Spectral graph theory*, volume 92 of *CBMS Regional Conference Series*. American Mathematical Society, Providence, 1997.
- The Gene Ontology Consortium. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nat. Genet.*, 25(1):25–29, May 2000. doi: 10.1038/75556.
- Keira R. Curtis, Matej Oresic, and Antonio Vidal-Puig. Pathways to the analysis of microarray data. *Trends in Biotechnology*, 23(8):429–435, 2005. doi: 10.1016/j.tibtech.2005.05.011.
- N. Friedman, M. Linial, I. Nachman, and D. Pe’er. Using Bayesian Networks to Analyze Expression Data. *J. Comput. Biol.*, 7(3-4):601–620, 2000. doi: 10.1089/106652700750050961.
- I. Gat-Viks, A. Tanay, and R. Shamir. Modeling and analysis of heterogeneous regulation in biological networks. *J Comput Biol*, 11(6):1034–49, 2004.

- D. Hanisch, A. Zien, R. Zimmer, and T. Lengauer. Co-clustering of biological networks and gene expression data. *Bioinformatics*, 2002.
- T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer, 2001.
- D. Hosack, G.Jr Dennis, B.T. Sherman, H.C. Lane, and R.A. Lempicki. Identifying biological themes within lists of genes with EASE. *Genome Biology*, R70:1–7, 2003.
- I.T. Jolliffe. *Principal component analysis*. Springer-Verlag, New-York, 1996.
- G. Joshi-Tope, M. Gillespie, I. Vastrik, P. D’Eustachio, E. Schmidt, B. de Bono, B. Jassal, G. R. Gopinath, G. R. Wu, L. Matthews, S. Lewis, E. Birney, and L. Stein. Reactome: a knowledgebase of biological pathways. *Nucleic Acids Res*, 33(Database issue):D428–32, 2005. 1362-4962 (Electronic) Journal Article.
- M. Kanehisa, S. Goto, S. Kawashima, Y. Okuno, and M. Hattori. The KEGG resource for deciphering the genome. *Nucleic Acids Res.*, 32(Database issue):D277–80, Jan 2004. doi: 10.1093/nar/gkh063.
- P. D. Karp, C. A. Ouzounis, C. Moore-Kochlacs, L. Goldovsky, P. Kaipa, D. Ahren, S. Tsoka, N. Darzentas, V. Kunin, and N. Lopez-Bigas. Expansion of the BioCyc collection of pathway/genome databases to 160 genomes. *Nucleic Acids Res*, 33(19):6083–9, 2005.
- M. Krull, S. Pistor, N. Voss, A. Kel, I. Reuter, D. Kronenberg, H. Michael, K. Schwarzer, A. Potapov, C. Choi, O. Kel-Margoulis, and E. Wingender. TRANSPATH: an information resource for storing and visualizing signaling pathways and their pathological aberrations. *Nucleic Acids Res*, 34(Database issue):D546–51, 2006.
- G. Mercier, Y. Denis, P. Marc, L. Picard, and M. Dutreix. Transcriptional induction of repair genes during slowing of replication in irradiated *Saccharomyces cerevisiae*. *Mutat. Res.*, 487(3-4):157–172, Dec 2001.
- G. Mercier, N. Berthault, J. Mary, J. Peyre, A. Antoniadis, J.-P. Comet, A. Cornuejols, C. Froidevaux, and M. Dutreix. Biological detection of low radiation doses by combining results of two microarray analysis methods. *Nucleic Acids Res.*, 32(1):e12, 2004. doi: 10.1093/nar/gnh002.
- B. Mohar. Some applications of Laplace eigenvalues of graphs. In G. Hahn and G. Sabidussi, editors, *Graph Symmetry: Algebraic Methods and Applications*, volume 497 of *NATO ASI Series C*, pages 227–275. Kluwer, Dordrecht, 1997.
- O Radulescu, S Lagarrigue, A Siegel, M Le Borgne, and P Veber. Topology and static response of interaction networks in molecular biology. *J.R.Soc.Interface*, Published online, 2005.
- J Rahnenfuhrer, FS Domingues, J Maydt, and T. Lengauer. Calculating the statistical significance of changes in pathway activity from gene expression data. *Statistical Applications in Genetics and Molecular Biology*, 3(1): Article 16, 2004.
- P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res*, 13(11):2498–504, 2003.
- R. Thomas and M. Kaufman. Multistationarity, the basis of cell differentiation and memory. II. Logical analysis of regulatory networks in terms of feedback circuits. *Chaos*, 11(1):180–195, 2001.
- J. P. Vert and M. Kanehisa. Extracting active pathways from gene expression data. *Bioinformatics*, 19 Suppl 2: II238–II244, 2003.