社団法人 情報処理学会 研究報告
IPSJ SIG Technical Report

2006－BIO－5 (8)
2006／6／15

ネットワーク型雑音除去によるマイクロアレイデータからの薬剤耐性予測

加藤 毅[†] 三浦 康[††] 村田 幸生[††] 浅井 潔[†,†††] ポールホートン[†††]
藤渕 航[†††] 津田 宏治[†††]

†††産総研 生命情報科学研究センター，〒135-0064 東京都江東区青海 2-43
†東京大学大学院新領域創成科学研究科，〒277-8562 柏市柏の葉 5-1-5
††東北大学大学院生体調節外科，宮城県仙台市青葉区星陵町 1-1
E-mail: †kato-tsuyoshi@cb.k.u-tokyo.ac.jp, ††{k-miura,yukio-m}@surg1.med.tohoku.ac.jp,
†††{asai-cbrc,horton-p,fujibuchi-wataru,koji.tsuda}@aist.go.jp

あらまし マイクロアレイデータから抗がん剤の耐性を予測するのは難しいタスクである．その理由のひとつにマイクロアレイデータが多くの雑音を含んでいることがあげられる．本論文ではオフ部分空間雑音除去法を拡張して配列相同性やたんぱく質相互作用などから得られる非一様なネットワークデータを利用して雑音を除去する方法を提案する．実験により，薬剤耐性予測にこの雑音除去法を組み合わせると，予測性能が向上することを示す．
キーワード マイクロアレイデータ，抗がん剤，ネットワーク，オフ部分空間雑音除去法

Drug-response prediction from microarray data using network-based de-noising

Tsuyoshi KATO[†], Koh MIURA[††], Yukio MURATA[††], Kiyoshi ASAI[†,†††], Paul B. HORTON[†††],
Wataru FUJIBUCHI[†††], and Koji TSUDA[†††]

†††AIST Computational Biology Research Center, 2-43, Aomi, Koto-ku, Tokyo, 135-0064 Japan
†Graduate School of Frontier Sciences, The University of Tokyo, 5-1-5, Kashiwanoha, Kashiwa city, 277-8562, Japan
††Division of Biological Regulation and Oncology, Department of Surgery, Tohoku University Graduate School of Medicine, 1-1, Seiryo-machi, Aoba-ku, Sendai 980-8574, Japan
E-mail: †kato-tsuyoshi@cb.k.u-tokyo.ac.jp, ††{k-miura,yukio-m}@surg1.med.tohoku.ac.jp,
†††{asai-cbrc,horton-p,fujibuchi-wataru,koji.tsuda}@aist.go.jp

Abstract Prediction of human cell response to anti-cancer drugs (compounds) from microarray data is a challenging problem, due to the noise properties of microarrays as well as the high variance of living cell responses to drugs. Hence there is a strong need for more practical and robust methods than standard methods for real-value prediction. We devised an extended version of the off-subspace noise-reduction (de-noising) method to incorporate heterogeneous network data such as sequence similarity or protein protein interactions into a single framework. Experimental results show that prediction performance is improved by combining a prediction method with our de-noising method.
Key words microarray data, anti-cancer drugs, network, off-subspace de-noising method

1. Introduction

Cancer diagnosis based on gene expression data has been widely and extensively explored in the clinical research field since the earlier papers on gene expression arrays were published. Early studies mainly focused on the classification of

cancer types, for example, discrimination of leukemia classes, a field in which powerful classifiers such as support vector machines are applied and the predictions are largely successful [3].

Recent cancer phenotype analysis is shifting from predicting a class to predicting a real-valued response. For example, predicting the effects of anti-cancer drugs(注1) is an important problem in cancer therapy, since a careful choice of proper not only drug but also dosage is required for different cancer cells and patients to maximize effectiveness and minimize deleterious side-effects. Although the drug response itself is a continuous quantity, this problem has often been simplified to the binary classification problem of drug sensitive vs. drug resistant [6], [9], [10]. Among the drug sensitivity classification studies, Staunton et al. [10] selected 232 out of 5,084 compounds or drugs to classify 60 cells(注2) by a sum of vote type classifier. According to their results, the rate of correct classification is significantly better than random classification.

In constrast to the simplified problem of classification, direct prediction of real-valued responses of a cancer drug from microarray data is not an easy task due to the noisy properties of both microarray technology and living cell experiments. Despite the limitation of available data, Mariadason et al [8] attempted to predict the cell apoptosis response against a chemotherapeutic agent (5-FU) by principal component regression (PCR). The leave-one-out test for 30 different cells in their analysis gives correlation coefficients of predicted and observed responses as low as 0.46. Gruvberger-Saal et al [4] also tried to predict the real-valued response of an estrogen receptor from gene expressions using artificial neural networks used in their earlier study.

In this paper, we focus on the noise and errors in microarray data that potentially degrade prediction performance. De-noising is similar to missing value estimation. Both infer the true values. Typical methods for missing value imputation (e.g. [2], [13]) capture the important dimensions by principal component analysis (PCA). However, they do not exploit the *side information* about genes, such as sequence similarity, GO classification, or protein-protein interactions, though those heterogeneous data sources are expected to be useful for identifying related genes, and furthermore to effectively correct noisy data.

We devise a new de-noising method using the side information represented as a *network*, where the nodes correspond to the genes, and the edges represent relations among the genes. In de-noising the expression data of a certain gene, we only look at its neighborhood genes in the network. A principal subspace is made only from the neighborhood expression data, and the target vector is de-noised by robust projection (Figure 1). Here, we use a so-called "off-subspace" projection method [14] to prevent over-de-noising. This projection algorithm is formulated as a linear program, which can be efficiently solved even for a large number of neighborhood genes. The basic idea of our method is similar to local PCA approaches [12], where de-noising is done by projection to local subspaces. However, the novelty of our method is that the neighborhood relation is determined by the network.

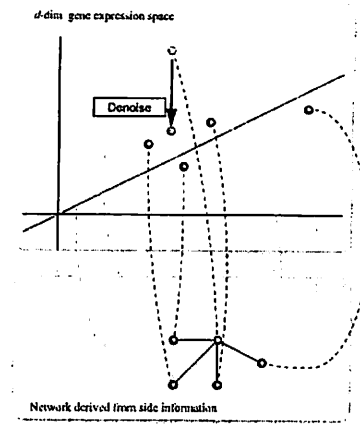

図 1 De-noising based on a network. In the top figure, the target expression vector to be de-noised is depicted as a red point. Black points are the neighbors in the network (below) derived from side information; gray points are vectors which are not directly connected (i.e. related) to the target vector in the network. The edge of the network is depicted by a solid line. A dashed curve indicates the correspondence between data in a network and the expression vector. De-noising is done by robust projection onto the principal subspace made from only the neighbors. In this case, the subspace (gray line) is obtained by PCA of the target and the four neighbors.

Typically we have multiple data sources as the side information, which are represented as *multiple networks*. When a principal subspace is derived from each network, we have a set of subspaces for each expression vector. A simple way is to take the sum of all subspaces, and project each target expression vector onto the combined subspace. However, since some of the networks might not be useful for de-noising, it is preferable to select important subspaces automatically, and then take the sum of those subspaces. To this end, we extend the off-subspace projection method to deal with multiple networks. Network selection is implemented by giving a non-negative weight parameter to each network, optimiz-

(注1) : In this paper, compounds are referred to as drugs.
(注2) : To be exact, the term 'cell' should be called a cell line.

ing the weight vector, and removing the networks with zero weights. The de-noising problem for multiple networks is also formulated as a linear program.

In predicting drug responses, it is often the case that the responses for many drugs are predicted from a microarray gene expression dataset. In this case, our problem is to learn a vector-to-vector mapping, where the output vector is composed of drug responses. Since the output vector provided for training is also noisy, our network-based de-noising method can also be applied to the output vector. As we have no side information about drugs, the correlation coefficients among the output vectors are used to construct a network.

In numerical prediction experiments using the drug response data by Staunton et al. [10], our method with multiple networks outperformed standard prediction methods, PCR and the k-nearest neighbor method, significantly. Note that the output de-noising was also effective to enhance the accuracy of prediction. The improvement of correlation between true responses and predictions was observed for 930 drugs out of 1,427 drugs. The number (930/1,427=0.65) is statistically significant ($p < 10^{-2}$) in a cumulative binomial distribution model under the null hypothesis that half of the drugs (714) are chosen by chance.

2. De-Noising with a Single Network

The dataset we analyzed contains one microarray hybridization experiment [1] for each cell sample. Let d denote the number of hybridizations and N denote the number genes used in the analysis. We consider the d dimensional space of hybridizations populated by N points representing individual genes. In our drug response case, which has expressions of 60 cell samples for each gene, $d = 60$. Unfortunately in our application, the N points are generally quite noisy. Thus our goal is to "correct" the N points in a way which effectively reduces noise while maintaining signal.

2.1 Derivation of Subspaces

Let us define $\boldsymbol{x}_i$ as the d-dimensional gene expression vector of the i-th gene. Our task is to de-noise $\boldsymbol{x}_i$ using other vectors and a network which is represented by the $N \times N$ symmetric matrix W. The (i,j) element w_{ij} represents the strength of the edge between two nodes i and j. If there is no edge, $w_{ij} = 0$. The principal subspace for the i-th gene is computed using the neighborhood nodes only, i.e., the nodes with $w_{ij} \neq 0$. The basis vectors of the subspace are obtained as the principal eigenvectors $\boldsymbol{z}_{is}, s = 1, \ldots, n_i$, of the following covariance matrix,

$$S_i = \frac{\sum_{j=1}^{N} w_{ij} \boldsymbol{x}_j \boldsymbol{x}_j^\top}{\sum_{j=1}^{N} w_{ij}} \qquad (1)$$

The weighting covariance matrix represents the distribution of neighbors, and thereby yields the subspace by taking major eigenvectors as the basis vectors. We determine the number of basis vectors, n_i according to the Kaiser-Guttman rule [5]: the number of basis vectors is set as the number of eigenvalues greater than one in the normalized covariance matrix $\tilde{S}_i$ in which $[\tilde{S}_i]_{kl} = [S_i]_{kl} / \sqrt{[S_i]_{kk}[S_i]_{ll}}$.

2.2 Off-subspace Projection

Most simply, de-noising is done by projecting $\boldsymbol{x}_i$ to the subspace in terms of the least squares error. However, when the number of neighborhood nodes is small, or the non-zero weights w_{ij} are concentrated in only a few neighbors, the dimensionality of the local subspace can be too small. In that case, simple projection may result in an unacceptably large loss of signal, called *over-de-noising*.

Tsuda and Rätsch [14] addressed this by devising an off-subspace projection method (Figure 2) which corrects the data to a point generally closer to, but not necessarily in the subspace. Let $\bar{\boldsymbol{x}}$ denote the de-noised result of $\boldsymbol{x}_i$, which is obtained by solving the following optimization problem,

$$\min_{\bar{\boldsymbol{x}}, v_s} d \|\bar{\boldsymbol{x}} - \sum_{s=1}^{n_i} v_s \boldsymbol{z}_{is}\|_\infty + \beta_0 \|\boldsymbol{x}_i - \bar{\boldsymbol{x}}\|_1, \qquad (2)$$

where β_0 is a non-negative constant parameter. Multiplying the first term by d is a convenient normalization which reduces the dependence of the numerical value of β_0 on d. With a suitable transformation in Eq. 2 can be efficiently minimized with standard linear programming.

As can be seen, the objective function in Eq. 2 is the sum of two distances; the distance between the subspace and the off-subspace point $\bar{\boldsymbol{x}}$, and the distance between $\bar{\boldsymbol{x}}$ and the input vector $\boldsymbol{x}_i$. The two distances are measured with different norms; the former with ℓ_∞, the later with ℓ_1.

If $\beta_0 \to \infty$ and if ℓ_∞-norm is replaced with ℓ_2-norm, this becomes a least square method and the resulting off-subspace solution $\bar{\boldsymbol{x}}$ is equal to $\boldsymbol{x}_i$. As β_0 decreases, $\bar{\boldsymbol{x}}$ becomes close to the subspace.

Now let us describe the ℓ_1-norm and ℓ_∞ norm. Since the ℓ_1-norm regularizer yields a sparse vector as the optimal solution, this algorithm almost corrects only the contaminated elements in $\boldsymbol{x}_i$ without change of the non-contaminated elements. The ℓ_∞ norm has similar behavior to the ℓ_2-norm [11]. Indeed, experiments in [14] show that the ℓ_∞ and ℓ_2 achieve similar de-noising performance. If ℓ_2 is used, the optimization problem can be transformed into a quadratic programming problem. If ℓ_∞ is used, the problem is a linear program, which can be solved more quickly and more stably than quadratic programs.

3. De-Noising with Multiple Networks

We further introduce a way to incorporate heterogeneous

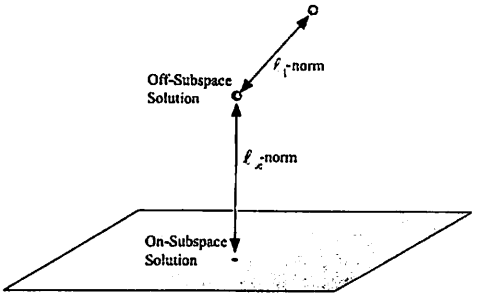

図 2 Off-subspace de-noising method. This figure illustrates how to de-noise an expression vector using a principal subspace. Instead of simple projection, we simultaneously find two points (off-subspace solution and on-subspace solution) which minimize the sum of ℓ_1-norm distance and ℓ_∞-norm distance given in Eq. 2.

data sources into the noise reduction process by weighting the multiple networks W_k, $k = 1, \ldots, m$. The advantage of this method is its ability to incorporate various kinds of biological knowledge into a single framework.

Denote the s-th basis vector obtained from the k-th network by $z_{is}^{(k)}$. To put it more precisely, each network yields the weighted covariance matrix $S_i^{(k)}$ and we take $n_i^{(k)}$ basis vectors $z_{is}^{(k)}, (s = 1, \cdots, n_i^{(k)})$ from $S_i^{(k)}$. To use all the subspaces gained from the networks, one can take the sum of the subspaces and apply the off-subspace projection method. Then, the optimization problem can be described as

$$\min_{\bar{x}, v_{is}^{(k)}} d\|\bar{x} - \sum_{k=1}^{m} \sum_{s=1}^{n_i^{(k)}} v_{is}^{(k)} z_{is}^{(k)}\|_\infty + \beta_0 \|x_i - \bar{x}\|_1. \quad (3)$$

Note that any point in the sum of the subspaces can be represented as $\sum_{k=1}^{m} \sum_{s=1}^{n_i^{(k)}} v_{is}^{(k)} z_{is}^{(k)}$. To automatically select important subspaces, we need to introduce a regularization term to the above optimization problem so that all the coefficients $v_{is}^{(k)}$ of unnecessary networks degenerate to zero.

For that, we introduce the upper bound of the absolute values of the coefficients of the k-th subspace as

$$t_k = \max_{1 \le s \le n_i} |v_{is}^{(k)}| = \| \left[v_{i,1}^{(k)}, \quad \cdots, \quad v_{i,n_i}^{(k)} \right] \|_\infty,$$

and penalize the ℓ_1-norm of the vector of upper bounds $t = (t_1, \cdots, t_m)^\top$ as follows,

$$\min_{\bar{x}, v_{is}^{(k)}} d\|\bar{x} - \sum_{k=1}^{m} \sum_{s=1}^{n_i} v_{is}^{(k)} z_{is}^{(k)}\|_\infty + \beta_0 \|x_i - \bar{x}\|_1 + \beta_1 \|t\|_1 \quad (4)$$

Due to the regularizer, some elements of t are exactly zero at the optimal solution Moreover one can control the number of nonzero elements with the constant parameter β_1. If $t_k = 0$, all the coefficients of the k-th network are zero, implying that that network is not used at all in deriving the de-noised result $\bar{x}$. This optimization problem can also be transformed into a linear program which can be solved efficiently (details not shown).

4. Experimental Settings

In the following experiments, our task is to predict the drug resistance levels of cells based on their gene expression data. By combining our novel de-noising methods and a standard prediction method, we wish to predict the resistance levels of the *test cells* accurately by learning from the input-output relations of the *training cells*.

The drug resistance dataset by Staunton et al. [10] contains the expression level of 6,817 genes from 60 human cancer cells. Among them, we pre-selected 2,067 highly variant genes (details not shown). For each cell, the drug resistance levels for 5,084 drugs are available as well. Those resistance levels are measured on a continuous scale by growth inhibition score (GI50). Prediction would be too easy, when the drug resistance levels are almost constant among cells. So we chose 1,427 drugs whose gap between the maximum and minimum levels was more than 0.5 after log-normalization.

We applied our network-based de-noising method in two ways. First, the expression profiles, i.e., the input of prediction, are de-noised with various networks. In the following experiments, we built three networks based on the correlation coefficients of the profiles, the gene ontology (GO), and the protein-protein interactions. Here, the vector being the de-noised x_i is the 60-dimensional expression vector of the i-th gene. Second, we also de-noised the vector of drug resistance levels, i.e., the output of prediction. Here, the vector x_i is composed of resistance levels of training cells for the i-th drug, and apply Eq. 4 to de-noise x_i. In this case, we used only one network based on the correlation coefficients of the resistance level vectors. Namely, $m = 1$.

A schematic representation of the entire process is shown in Figure 3.

For predicting drug responses from the de-noised expression data, we tested two kinds of standard prediction algorithms. One is principal component regression (PCR), the other is k-nearest-neighbors (kNN). Both algorithms have one constant parameter to be tuned manually, that is, the number of principal components and the number of nearest neighbors, respectively.

Principal Component Regression

PCR is also used in Mariadason et al. [8]. For training cells, the Pearson correlation between the de-noised expression data of each of the 3,725 genes and (de-noised) responses of a drug of interest were calculated, and the 50 highest absolute value correlations (i.e., corresponding to 50 genes) were

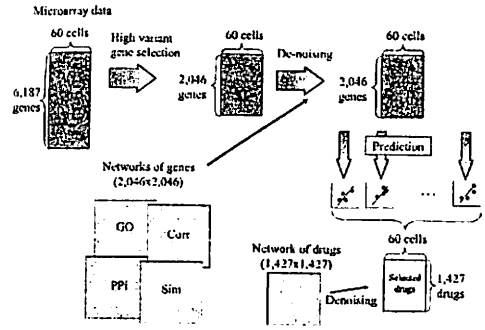

図 3 Method overview

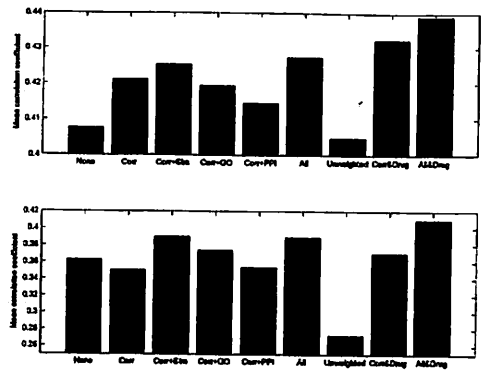

図 4 Improvement of prediction by noise reduction with various combinations of networks. The upper figure plots the results of PCR, the bottom of kNN. The mean correlation coefficients of prediction for 1,427 drugs before and after the off-subspace noise reduction are shown. Abbreviations are Corr: correlation coefficient for gene expressions; Sim: sequence similarity; GO: gene ontology; PPi: protein-protein interaction; All: Corr+Similarity+GO+PPi; Corr&Drug: input de-noising only via Corr and output de-noising; and All&Drug: input de-noising using All and output de-noising.

selected. To reduce the number of genes to a smaller set of variables, PCA was performed. From the PCA, the principal components (PCs) having the d_{pca} largest eigenvalues were selected. Next we get the parameters, $\boldsymbol{a} \in \Re^{d_{\mathrm{pca}}}$ and $b \in \Re$, of a linear regression function $g(\boldsymbol{z}|\boldsymbol{a}, b) = \boldsymbol{a}^\top \boldsymbol{z} + b$ by least square manner, where $\boldsymbol{z} \in \Re^{d_{\mathrm{pca}}}$ is the d_{pca} PCs of a cell. Namely, we find the values of $\{\boldsymbol{a}, b\}$ which minimize the sum-of-square errors:

$$\sum_i (g(\boldsymbol{z}_i|\boldsymbol{a}, b) - y_i)^2$$

where $\boldsymbol{z}_i \in \Re^{d_{\mathrm{pca}}}$ and y_i are the PCs and the drug response of i-th training cell, respectively. Once the regression function $g(\boldsymbol{z}|\boldsymbol{a}, b)$ was derived, the d_{pca} PCs corresponding to the test cell were computed and substituted into the derived regression function to yield a prediction of response of the test cell.

k-Nearest Neighbor Prediction

Using the same manner as the above Mariadason et al.'s method [8], 50 genes are selected. Then each cell has a 50-dimensional vector. For each test cell, we find the k_{nn} closest cells and predict the drug response of the test cell by the average of k_{nn} responses.

4.1 Building Networks

In this section, we describe the details of network construction, namely, the computation of matrix W. The first four networks are used for de-noising the input (i.e., expression data), and the last one is for the output (i.e., drug resistance levels). Details are referred to our paper [7].

4.2 Parameter Selection

Our de-noising method has the two parameters, β_0 and β_1. In addition PCR and kNN have one constant parameter, d_{pca} and k_{nn}, respectively. For this purpose, we performed a joint grid search over the following values: $\beta_0 = 0.1, 0.2, 0.5, 1.0, 2.0, 3.0, 4.0$, $\beta_1 = 0.00, 0.02, 0.05, 0.1, 0.2, 0.5$, $d_{\mathrm{pca}} = 1, 2, 3, 4, 5, 7, 10, 20, 30, 40, 50$, $k_{\mathrm{nn}} = 1, 2, 3, 4, 5, 6, 7, 8, 9, 10$, and chose the parameter values yielding the best regression performance.

5. Results

Prediction performances of PCR with and without off-subspace noise reductions for different combinations of networks are shown in the upper plot in Figure 4. The accuracy of prediction is measured by the mean correlation coefficients in 12-fold cross validation. The leftmost bar 'None' corresponds to the performance without any de-noising. As anticipated, integration of both input and output de-noising yields the highest mean correlation coefficient ('All&Drug': 0.439). Among the other cases where only the input is de-noised, we obtained the best result when all networks are combined with weights ('All': 0.428). The weightless combination ($w_{ij} = 1$ for$\forall i, j$) was significantly poorer ('Unweighted'). The lower plot of Figure 4 shows the performance of kNN. The kNN also achieves the best performance when both input and output are de-noised. In the case of de-noising only the input, noise reduction without using side-information degrades the prediction performance, but the use of side-information improves prediction. We also counted the drugs achieving statistically-significant predictions (Table 1).

Further experimental results are detailed in our journal paper [7].

6. Concluding Remarks

The prediction of drug response data is critical for the

表 1 The number of drugs given statistically-significant prediction of the responses. Here we define a drug that achieves the correlation coefficient more than 0.33 as a successfully predicted, which is derived from one sample t-test with the probability less than 0.01 examining the null hypothesis of "no correlation."

	None	Corr	Corr+Sim	Corr+GO	Corr+PPi	All	Unweighted	Corr&Drug	All&Drug
PCR	983	1,027	1,043	1,018	995	1,037	988	1,065	1,085
kNN	867	794	925	886	815	937	565	861	994

field of cancer therapeutics, which demands improved diagnostics for determining the appropriate choice and dosage of anti-cancer drugs. Combining gene relations from various biological resources to adjust values of gene expressions or drug response data is a new approach in this field. This approach requires effective methods, such as the one presented here, for utilizing heterogeneous data.

This algorithm is invariant if the network weights are multiplied by a constant, as shown in Eq. 1. However, the change of the ratio among weights may have an influence on the de-noising performance. Although various weighting schemes could be considered, we did not systematically investigate that issue in this work. However we did confirm that off-subspace noise reduction with the continuous weights defined above for sequence similarity, expression correlation, and GO heterogeneous data sources was more effective than using a 0-1 weighting scheme based on some threshold. The current weighting scheme might not yet be optimal and tuning would yield improvement to some extent.

We extended the off-subspace noise reduction method of Tsuda and Rätsch [14] and applied it to the noise reduction of gene expression data in the context of real-value prediction to drug response data. Our results show the method to be robust to noisy data and more effective than the traditional principal component regression, improving the prediction of 868 out of 1,427 drugs. We expect it will prove generally useful for correcting the values of noisy microarray data.

文　献

[1] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, and P. Walter. *Molecular Biology of the Cell. 4th ed.*, chapter Microarrays Monitor the Expression of Thousands of Genes at Once, III. Methods / 8. Manipulating Proteins, DNA, and RNA / Studying Gene Expression and Function. Garland Publishing, New York, 2002.

[2] T. H. Bo, B. Dysvik, and I. Jonassen. Lsimpute: accurate estimation of missing values in microarray data with least squares methods. *Nucleic Acids Res.*, 32(3):e34, Feb 2004.

[3] T. S. Furey, N. Cristianini, N. Duffy, D. W. Bednarski, M. Schummer, and D. Haussler. Free full text support vector machine classification and validation of cancer tissue samples using microarray expression data. *Bioinformatics*, 16(10):906–914, 2000.

[4] S. K. Gruvberger-Saal, P. Eden, M. Ringner, B. Baldetorp, G. Chebil, A. Borg, M. Ferno, C. Peterson, and P. S. Meltzer. Predicting continuous values of prognostic markers in breast cancer from microarray gene expression profiles. *Mol. Cancer Ther.*, 3(2):161–168, Feb 2004.

[5] L. Guttman. Some necessary conditions for common-factor analysis. *Psychometrika*, 19:149–161, 1954.

[6] Y. Kaneta, Y. Kagami, T. Katagiri, T. Tsunoda, I. Jin-nai, H. Taguchi, H. Hirai, K. Ohnishi, T. Ueda, N. Emi, A. Tomida, T. Tsuruo, Y. Nakamura, and R. Ohno. Prediction of sensitivity to STI571 among chronic myeloid leukemia patients by genome-wide cDNA microarray analysis. *Jpn. J. Cancer. Res.*, 93(8):849–856, Aug. 2002.

[7] T. Kato, Y. Murata, K. Miura, K. Asai, P. B. Horton, K. Tsuda, and W. Fujibuchi. Network-based de-noising improves prediction from microarray data. *BMC Bioinformatics*, 7(Suppl 1):S4, March 2006.

[8] J. M. Mariadason, D. Arango, Q. Shi, A. J. Wilson, G. A. Corner, C. Nicholas, M. J. Aranes, M. Lesser, E. L. Schwartz, and L. H. Augenlicht. Gene expression profiling-based prediction of response of colon carcinoma cells to 5-fluorouracil and camptothecin. *Cancer Res.*, 63(24):8791–8812, Dec 2003.

[9] J. Okutsu, T. Tsunoda, Y. Kaneta, T. Katagiri, O. Kitahara, H. Zembutsu, R. Yanagawa, S. Miyawaki, K. Kuriyama, N. Kubota, Y. Kimura, K. Kubo, F. Yagasaki, T. Higa, H. Taguchi, T. Tobita, H. Akiyama, A. Takeshita, Y. H. Wang, T. Motoji, R. Ohno, and Y. Nakamura. Prediction of chemosensitivity for patients with acute myeloid leukemia, according to expression levels of 28 genes selected by genome-wide complementary DNA microarray analysis. *Mol. Cancer Ther.*, 1(12):1035–1042, Oct 2002.

[10] J. E. Staunton, D. K. Slonim, H. A. Coller, P. Tamayo, M. J. Angelo, J. Park, U. Scherf, J. K. Lee, W. O. Reinhold, J. N. Weinstein, J. P. Mesirov, E. S. Lander, and T. R. Golub. Chemosensitivity prediction by transcriptional profiling. *Proc. Natl. Acad. Sci. U.S.A.*, 98(19):10787–10792, Sept 2001.

[11] R. Tibshirani. Regression selection and shrinkage via the lasso. *Journal of the Royal Statistical Society Series B*, 58(1):267–288, 1996.

[12] M.E. Tipping and C.M. Bishop. Mixtures of probabilistic principal component analyzers. *Neural Computation*, 11(2):443–482, 1999.

[13] O. Troyanskaya, M. Cantor, G. Sherlock, P. Brown, T. Hastie, R. Tibshirani, D. Botstein, and R. B. Altman. Missing value estimation methods for dna microarrays. *Bioinformatics*, 17(6):520–525, Jun 2001.

[14] K. Tsuda and G. Rätsch. Image reconstruction by linear programming. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, Cambridge, MA, 2004. MIT Press.