

Controlling the False Discovery Rate for Feature Selection in High-Resolution NMR Spectra

COSMOS Technical Report 06-03

Seoung Bum Kim*, Victoria C.P. Chen

Department of Industrial and Manufacturing Systems Engineering
University of Texas at Arlington, Arlington, TX, 76019, USA

Youngja Park, Thomas R. Ziegler, Dean P. Jones

Clinical Biomarker Laboratory, Center for Clinical and Molecular Nutrition
Department of Medicine
Emory University, Atlanta, GA, 30322, USA

* Corresponding author: Seoung Bum Kim; (817) 272-3150; sbkim@uta.edu

Abstract

Successful implementation of feature selection in nuclear magnetic resonance (NMR) spectra not only improves classification ability, but also simplifies the entire modeling process and, thus, reduces computational and analytical efforts. Principal component analysis (PCA) and partial least squares (PLS) have been widely used for feature selection in NMR spectra. However, extracting meaningful metabolite features from the reduced dimensions obtained through PCA or PLS is complicated because these reduced dimensions are linear combinations of a large number of the original features. In this paper, we propose a multiple testing procedure controlling false discovery rate (FDR) as an efficient method for feature selection in NMR spectra. The procedure clearly compensates for the limitation of PCA and PLS and identifies individual metabolite features necessary for classification. In addition, we present orthogonal signal correction to improve classification and visualization by removing unnecessary variations in NMR spectra. Our experimental results with real NMR spectra showed that classification models constructed with the features selected by our proposed procedure yielded smaller misclassification rates than those with all features.

Keywords: False discovery rate; Metabolomics; Nuclear magnetic resonance; Orthogonal signal correction; Feature selection

1. Introduction / Background

Metabolomics approaches that use proton nuclear magnetic resonance ($^1\text{H-NMR}$) have been developed in recent years as a way to study the dynamic and time-dependent profiles of metabolic responses that occur in response to pathophysiological stimuli or to genetic modifications in integrated biological systems [1]. A metabolomics approach enables the investigation of hundreds of metabolites simultaneously. In biofluids, metabolites are maintained in dynamic balance with those that are inside cells and tissues, and consequently, abnormal perturbations after toxic insult or disease progression will be reflected in disturbances in the ratios and concentrations of biofluid metabolites. The introduction of $^1\text{H-NMR}$ spectroscopic techniques makes it possible to investigate changes in metabolic composition and to quantify them without complex preparation of samples.

$^1\text{H-NMR}$ depends upon the characteristic spin of the atomic nucleus of hydrogen, which exists in two spin states in the presence of a strong external magnetic field. The lower energy spin state is aligned with the external field, but the higher energy spin state is opposed to the external field. Irradiation of a sample with radio frequency energy corresponding exactly to the spin state separation will cause excitation of those nuclei in the lower spin state to the higher spin state. This radio frequency energy is reflected in the ppm (parts per million) value of the x -axis shown in Figure 1, which displays an example of multiple NMR spectra of human plasma. In different biological molecules, the energy required for this transition, often referred to as the *chemical shift*, is a characteristic for different hydrogen in specific chemicals. Traditionally, chemical shifts along the x -axis are listed from largest to smallest. The y -axis is proportional to the

concentration of the hydrogen nucleus absorbing energy so that the spectrum provides information on the chemicals which are present and also the concentrations of those chemicals.

[Figure 1 about here]

A major limitation in NMR spectroscopic evaluation of tissues and body fluids lies in the complexity of the signals generated by the presence of multiple chemical species, which often have signals that overlap. Because of this overlap, we refer to the amplitude at specific ppm (chemical shift) values as features. To search for chemical species that are useful for classification and prediction, we separate two components, the first being an effort to identify metabolite features linked to potential risk factors (i.e., feature extraction/selection) and the second being a study to explicitly test whether these metabolite features predict disease (i.e., prediction or classification).

The main purpose of feature extraction/selection is to simplify the entire modeling process and to reduce computational and analytical efforts by identifying the important features from the high-dimensional original dataset. There is a distinction between feature extraction and feature selection, although much of the literature fails to make a clear distinction. Feature extraction techniques attempt to create new features based on transformations of the original features to extract the useful information for the model, while feature selection techniques attempt to pick the subset of original features that leads to the best prediction or classification [2, 3]. Despite numerous studies of feature extraction/selection that have been conducted in fields of signal and image processing, feature extraction/selection methods have not been thoroughly studied for the application of NMR spectra.

The most widely used method in NMR applications includes a combination of principal component analysis (PCA) and classification models [4, 5]. Although PCA has been successfully used to facilitate visualization of the complicated spectra and, thus, provide an initial idea to guide subsequent statistical analyses [6-8], PCA is not an efficient method for the interpretation of the features selected. The reduced dimensions, called principal components (PCs) from PCA are each a linear combination of the original features, where the number of original features in the NMR spectra is usually hundreds to thousands. Thus, interpretation of the PCs cannot readily be made, and the extraction of meaningful information is cumbersome. Another major drawback of PCA is that the PCs may not always produce maximum discrimination between classes because the transformation process of PCA relies solely upon the input variables, and does not take into account class information.

Partial least squares (PLS) is another popular method that utilizes class information to help select a transformation better suited for classification [9, 10]. However, the reduced dimensions from PLS also do not provide a clear interpretation with respect to the original features due to the same reason described for PCA. Recently, a two-stage genetic programming (GP) method was developed for identifying important metabolite features for the classification of genetically modified barley [11]. GP is also known as a hierarchical genetic algorithm, introduced by Koza [12]. The first stage of this method searches a whole spectrum to identify the significant features for classification using standard GP. In the second stage, additional GP is employed to find significant features from the features selected in the first stage. This method apparently overcomes the interpretation problem of PCA and PLS since the method does not utilize

any transformations to find the significant features. However, GP is known to be vulnerable to overfitting when the features are noisy or the dimension of the features is high [13]. Furthermore, GP is a purely numerical optimization technique, and it may be hard to interpret the feature selection results statistically.

In this paper, we formulate a feature selection problem in NMR spectra as a multiple hypothesis testing problem controlling false discovery rate (FDR). More precisely, we propose a procedure to determine which metabolite features discriminate spectra among the experimental conditions based on a multiple testing framework. It is well known that applying a single testing procedure to the multiple testing problems leads an exponential increase of the probability that at least one of the tests leads to rejection of a null hypothesis when the null hypothesis holds. To compensate for this problem, the procedures controlling the family-wise error rates (FWE) that control all the tests simultaneously were developed [14]. However, these procedures are too stringent to detect true significant features. In other words, the procedures that control FWE lead to low power, where the power is the proportion of false null hypotheses, which are correctly rejected [14]. The FDR is the error rate in multiple hypothesis tests and is defined as the expected proportion of false positives among all the hypotheses rejected [15]. The advantage of an FDR procedure is to identify as many significant hypotheses as possible while keeping a relatively small number of false positives [15-17].

The FDR-based procedure overcomes a limitation posed by the reduced dimensions in PCA or PLS and can identify the metabolite features necessary for classification without losing any information in the original features. Moreover, the FDR-based procedure is a statistical approach that enables us to control FDR of the features

identified, thus provide the logical interpretation of feature selection results. FDR-based procedures have recently been used in microarray analysis to find co-expressed genes [18] and applied to neuroimaging analysis to identify active voxels in the image data [19]. In microarray/neuroimaging studies, a hypothesis test was performed in each gene/voxel to determine whether the gene/voxel contributes to classification between the different experimental conditions.

It is known that high-resolution $^1\text{H-NMR}$ spectra obtained from multiple samples may contain variability from a variety of sources (e.g., instrumental, environmental, and physiological), which deteriorate classification ability [20]. To improve visualization and discrimination in spectra, we applied orthogonal signal correction (OSC) that enhances the separability of the different classes by removing variability that do not contribute to prediction or discrimination [21]. The OSC technique for NMR spectra has been shown to minimize physical and biological variation, thus, improving the interpretation of statistical results [20, 22].

Throughout the paper, the effectiveness of our procedures is demonstrated with $^1\text{H-NMR}$ spectra of human plasma, where the goal is to examine metabolic perturbation in response to sulfur amino acids (SAA) intake. SAA are highly variable in human food, and either the deficiency or the excess of SAA may lead to potential risk in human health. They are required for physiologic processes in addition to their role in the maintenance of protein synthesis and nitrogen balance.

2. Experimental Data

2.1. Collection of NMR spectra and experimental data structure

We conducted the study within two phases during which plasma samples were obtained from four healthy subjects under controlled metabolic conditions in the Emory General Clinical Research Center (GCRC). The subjects signed an informed consent approved by the Emory Institutional Review Board. During the 12-day GCRC admission, the subjects consumed defined diets at standardized intervals. During the 48-hour equilibration period, the subjects consumed a balanced meal plan with foods selected to ensure adequate energy, protein and SAA intake. The subjects were then placed on constant semi-purified diets designed to alter SAA intake, a 5-day zero sulfur amino acid phase (zero-SAA phase) immediately followed by a 5-day SAA supplementation phase (supplemented-SAA phase). In other words, during the initial 17 time points (first 5 days), blood was collected when the subjects were consuming diets providing zero SAA, and during the latter 17 time points (next 5 days), blood was collected when the subjects were consuming diets providing 117 mg/kg/day SAA, or approximately 10 times the recommended daily/dietary allowance for SAA. ¹H-NMR spectra of human plasma (drawn from blood) were then obtained by a Varian INOVA 600 MHz instrument.

In our experiment, 34 NMR spectra were collected from each subject. The 34 spectra were analyzed in each individual, where the half of them (17 spectra) were collected from the zero-SAA phase (class label=0) and the other half were collected from the supplemented-SAA phase (class label=1). Because there are four subjects, the total number of spectra are 136 (= 4 subjects × 34 spectra), where again the half of them (68 spectra) were collected from the zero-SSA phase (class 0) and the other half (68 spectra) were collected from the supplemented-SAA phase (class 1).

2.2. Preprocessing of NMR spectra

NMR spectra require preprocessing steps before conducting statistical analysis to detect subtle variations from metabolic profiles. In general, the steps involve phase and baseline correction, spectral alignment, elimination of non-informative spectral regions, and normalization. Phase and baseline corrections using NUTS software (Acorn NMR Inc., Livermore, CA) achieved accurate computing of integration of peak regions. Spectral alignment enables direct comparison across different spectra. We utilized a beam search algorithm [23], which by maximizing their correlation determines the best alignment between the reference spectrum and the sample spectra that require alignment. Because they can mask signals from low metabolites, we removed the spectral regions containing signals from water (4.5 – 5 ppm) and artificial signals caused by endogenous plasma calcium and magnesium (in ppm: 2.573616 – 2.589738, 2.721409 – 2.726783, 3.020354 – 3.156727, 3.22189 – 3.260854, and 3.62429 – 3.649818). Further, the spectral regions in 5.4 - 6.7 ppm and 7.8 -10 ppm were not considered in subsequent statistical analyses because they contain no significant metabolite signal. To ensure comparability between spectra, spectra were normalized to the area of the internal standard. An example of original spectrum is displayed in Figure 2 (a). A spectrum after preprocessing and removal of the redundant regions is shown in Figure 2 (b).

[Figure 2 about here]

3. Analytical Methods

3.1. A multiple testing procedure controlling false discovery rate

A multiple testing procedure controlling FDR is proposed to identify the metabolite features that play a significant role in discriminating time points in the zero-

SAA phase from those in the supplemented-SAA phase. FDR is a convenient definition of an error rate, which is defined as the expected proportion of false positives among the entire hypothesis rejected [15]. Furthermore, Benjamini and Hochberg introduced an ordered p -value procedure and proved that the procedure controls the specified FDR (or α) [15]. A multiple testing procedure identifies the metabolite features that are central to discriminating time points by their amount of SAA intake. More precisely, each null hypothesis states that the average intensities of the i^{th} metabolite features are equal between the time points from the zero-SAA phase and the supplemented-SAA phase, and the alternative hypothesis is that they differ. Because we have 8,444 metabolite features of interest, we can construct multiple hypotheses as follows:

$$\begin{aligned}
 H_1 &: m_{1,g1} = m_{1,g2} \\
 H_2 &: m_{2,g1} = m_{2,g2} \\
 &\vdots \\
 H_{8,444} &: m_{8,444,g1} = m_{8,444,g2},
 \end{aligned}$$

where $g1$ and $g2$ are the time groups with the zero-SAA phase and the supplemented-SAA phase, respectively. The metabolite features corresponding to the rejected hypotheses imply strong discrimination between the two SAA phases. Now we can employ a two-sample t test for each metabolite feature using test statistics as follows:

$$t_i = \frac{\bar{y}_{i,g1} - \bar{y}_{i,g2}}{\sqrt{\frac{s_{i,g1}^2}{n_{g1}} + \frac{s_{i,g2}^2}{n_{g2}}}},$$

for $i=1,2, \dots, 8,444$. $\bar{y}_{i,g1}$ and $s_{i,g1}^2$ are the sample mean and variance of metabolite feature i among the time points taken from $g1$. Likewise, $\bar{y}_{i,g2}$ and $s_{i,g2}^2$ are obtained

from g_2 . We compute t_i for $i=1, 2, \dots, 8,444$ assuming that the null hypothesis is true. We next calculate p -values using a permutation method due to repeated measurements and thus we cannot clearly assume that each t_i follows a t distribution. Under the assumption that there is no differential spectral intensity between the two classes (two SAA phases), the t statistic should have the same distribution regardless of how we make the permutation of spectra. Therefore, we can permute (shuffle) the labels of the 136 spectra and re-compute a set of t statistics for each individual metabolite feature based on the permuted dataset. If this procedure is repeated M times, we can obtain M sets of t -statistics as follows: $t_1^m, t_2^m, \dots, t_{8,444}^m$, $m = 1, 2, \dots, M$. The p -value for metabolite feature i for $i=1, 2, \dots, 8,444$ and $M=200$ is obtained by

$$p_i = \sum_{m=1}^{200} \frac{\#\{k : |t_k^m| \geq |t_i|, k = 1, 2, \dots, 8,444\}}{8,444 \cdot 200}.$$

Finally, we used the FDR procedure [15] that used ordered p -values ($p(1) \leq p(2) \leq \dots \leq p(8,444)$) to select the significant metabolite features to discriminate the time points in the zero-SAA phase from those in the supplemented-SAA phase. The summary of the procedure is as follows:

- Select a desired FDR level ($=\alpha$) between 0 and 1.
- Find the largest i denoted as w .

$$w = \max \left[i : p(i) \leq \frac{i \alpha}{m \delta} \right],$$

where m is the total number of hypothesis (here $m=8,444$) and δ denotes the proportion of true null hypothesis. Several studies discuss the assignment of δ . We used $\delta=1$, the most conservative choice.

- Let the p -value threshold be $p_{(w)}$, and declare the metabolite feature t_i significant if and only if $p_i \leq p_{(w)}$.

The original FDR procedure [15] assumes that the hypotheses tests are all independent. Later work revealed that the conclusion still holds even if the hypothesis tests are positively correlated [24]. Furthermore, Benjamini and Yekutieli [24] provided a procedure that can handle general dependency of the hypothesis tests through a simple modification of the original FDR procedure. It should be noted that metabolite features in an NMR spectrum are correlated with each other.

3.2. Orthogonal signal correction

To determine whether visualization and classification could be improved by removal of unwanted spectral variations, an OSC method was used. In a seminal paper, Wold *et al.*, [21] introduced OSC, a partial least squares (PLS)-based solution that can selectively remove from \mathbf{X} (i.e., predictor) the unwanted largest variation orthogonal (or unrelated) to \mathbf{Y} (response). The first step of OSC is to calculate the principal component score vector \mathbf{t} from \mathbf{X} . The score vector \mathbf{t} is then orthogonalized with respect to \mathbf{Y} through the following equation:

$$\mathbf{t}^* = \mathbf{t} - \mathbf{Y}(\mathbf{Y}^T\mathbf{Y})^{-1}\mathbf{Y}^T\mathbf{t}.$$

It is not difficult to show that \mathbf{t}^* is orthogonal to \mathbf{Y} [21]. In the next step, a weight vector, \mathbf{w} is calculated such that $\mathbf{X}\mathbf{w} = \mathbf{t}^*$, followed by the calculation of a new score vector from \mathbf{X} and \mathbf{w} , $\mathbf{t} = \mathbf{X}\mathbf{w}$. The above steps are repeated until the difference between the new score vector and the previous score vector is less than a predefined threshold. Finally, a loading vector, \mathbf{p} is calculated, and the correction is performed by subtracting $\mathbf{t}\mathbf{p}^T$ from \mathbf{X} giving a

residual vector ($\mathbf{E} = \mathbf{X} - \mathbf{t}\mathbf{p}^T$). The subsequent components can be computed using \mathbf{E} as a replacement of \mathbf{X} .

Since the introduction of OSC by [21], a number of OSC algorithms have been proposed in the literature. These include modified versions of the original OSC [25, 26], a direct orthogonal signal correction (DOSC; [27]), and orthogonal projections to latent structures [28]. A comprehensive summary of OSC algorithms can be found in [29].

In the present study a DOSC algorithm was used and implemented using MATLAB codes available from [27]. DOSC was applied to raw spectra for the purposes of classification or discrimination, as in PLS-discrimination analysis. Thus, the response matrix \mathbf{Y} , as categorical variables, contains information about class memberships of spectra (e.g., zero-SAA phase and supplemented-SAA phase).

It is important to determine the proper number of OSC components in a given application because there is a risk of overfitting when too many components are used. Thus, OSC results were cross validated using k -nearest neighbor (kNN), one of the widely used supervised classification algorithms ([9], pp.415-420). Here we used Euclidean distance to determine the neighborhoods and tested different k (i.e., $k = 2, 3, \dots, 10$). To obtain cross-validated error rates, the experimental datasets were split into four groups corresponding to four subjects. Three subjects were used for training, and the one remaining subject was used for testing. This process was repeated three more times. The final classification results from the four different testing samples were then averaged to obtain the cross-validated error rates of the models. It can be seen from Figure 3 that kNN models, processed by one OSC component yield higher cross-validated error rates than other kNN models, processed by more than two OSC components over the entire range of

k ($k = 2, 3, \dots, 10$) considered here. A comparison of models with more than two OSC components shows that they are all comparable in terms of their cross-validated errors. Thus, we used two-component OSC models in this paper. Previous studies also indicated that one or two OSC components are usually sufficient [21, 27].

[Figure 3 about here]

3.3. Classification methods and validation

Classification was performed to obtain the models for separating distinct classes and classifying new spectra into previously defined classes. We considered k-NN, classification trees and support vector machines (SVMs); however, the k-NN classifier is known to be poorly suited for high-dimensional datasets and performed somewhat unreliably. Hence, we present only the misclassification rates for trees and SVMs. Classification trees partition the feature space into disjoint hyper-rectangular regions according to performance measures (e.g., misclassification errors, Gini index, and cross-entropy) and then fit a constant model in each disjoint region (Breiman *et al.*, 1984). We used the Gini index as a performance measure. The number of disjoint regions (or the number of terminal nodes in a tree) should be determined appropriately because a very large tree overfits the training set, while a small tree cannot capture important information in the data. A common approach to determine the right size of a tree is tree pruning that removes the leaves and branches of a full-grown tree to find the right size of the tree [30]. A cost-complexity pruning algorithm was used for tree pruning [31]. SVMs use geometric properties to obtain the separating hyperplane by solving a convex optimization problem that simultaneously minimizes the generalization error and maximizes the geometric margin between the points for different classes (two SAA

groups in our case) [32, 33]. Nonlinear SVM models can be constructed from kernel functions that include linear, polynomial, radial basis functions. In our experiment, we used polynomial (with degrees of two) kernel functions. Classification trees and SVMs were performed using MATLAB (The MathWorks Inc., Natick, MA).

Classification models are frequently evaluated based on their misclassification rate. Four-fold cross validation was used to estimate of the true misclassification rate. As discussed in Section 2 for determining the number of OSC components, we divided the experimental datasets into four groups corresponding to four subjects. In each of four iterations, three subjects were used for training the models, and the one remaining subject was used for testing. This process was repeated three times more. The final classification results from the four different testing samples were then be averaged to obtain the cross-validated error rates (or misclassification rates) of the classification tree and SVMs models.

4. Results and discussions

4.1. Metabolite feature selection

A multiple testing procedure controlling FDR at 0.01 was performed to find the potentially significant metabolite features that differentiate spectra from time points in the zero-SAA phase to those in the supplemented-SAA phase. The procedure identified the region containing 1,458 chemical shifts as potentially significant metabolite features, demonstrating significant dimension reduction. In order to confirm the assignment of these chemical shifts to specific metabolites, an experimental validation is needed using, for example, two-dimensional NMR spectroscopy. In the present study, we focus on a

computational evaluation of features selected. Interpretation of results in the context of FDR implies that there are on average 15 ($15 \approx 1,458 \times 0.01$) false discoveries out of the 1,458 features discovered from the FDR-based procedure. Higher levels of FDR (e.g., FDR level = 0.05 or 0.1) increase the number of significant features, which yields larger power but produces more false significant features (or false discoveries).

4.2. Direct orthogonal signal correction processing

The effect of DOSC preprocessing of original data in terms of discrimination can be found from the PCA score plots (Figures 4(a)-(f)). PCA was performed after normalization of spectra using SIMCA-P11 (Umetrics, Sweden). PCA identifies a lower dimensional space called principal components (PCs) through orthogonal transformations of the original features [34]. The PCs are uncorrelated to each other and generally, the first few PCs are sufficient to account for most of the variability in the original high-dimensional space. Figures 4(b)-(f) display two-dimensional PCA score plots of the first two PCs (PC1 vs. PC2) for the DOSC processed spectra with different numbers of OSC components. Compared to Figure 4(a) (obtained from original spectra), Figures 4(b)-(f) (obtain from DOSC-processed spectra) show a well-defined discrimination of the two classes. It can be seen that the PCA score plots with large numbers of OSC components showed that metabolic variation, because of dietary SAA intake, is largely described by PC1. Different numbers of OSC components were plotted to see their effect on the PCA score plots (Figures 4(b)-(f)). Although using one OSC component improve the separation of group somewhat compared to original spectra, the score plot using more than two OSC components clearly visualizes groupings between the time points from the zero-SAA phase and the supplemented SAA phase. Generally, as the number of OSC

components increase, the score plots provided a slightly better separation in that the groups appeared more compact but not to a significant degree. This graphical analysis from the PCA score plots further justified that two OSC components are sufficient to discriminate two classes.

[Figure 4 about here]

4.3. Development of classification rules and validation

Classification methods construct the decision rules that maximize the mathematical separation between different classes with knowledge of predefined class labels. To see the potential advantage of FDR and DOSC, the datasets with the different numbers of metabolite features indicated in parentheses were used for the classification methods.

- Dataset 1: All metabolite features (8,444).
- Dataset 2: Metabolite features selected by FDR with level = 0.01 (1,458).
- Dataset 3: DOSC-processed dataset 1 (8,444).
- Dataset 4: DOSC-processed dataset 2 (1,458).
- Dataset 5: Metabolite features selected by FDR at $\alpha=0.01$ after DOSC (5,810).

Dataset 1 is the full dataset containing all metabolite features. Dataset 2 contains 1,458 features selected by the FDR procedure (FDR level = 0.01). Datasets 3-5 were preprocessed by DOSC with two OSC components. Dataset 3 consists of all metabolite features (dataset 1) processed by DOSC. Dataset 4 consists of the features selected from FDR (dataset 2) processed by DOSC. Dataset 5 applies FDR to Dataset 3 (post-DOSC), resulting in the selection of 5,810 features. The FDR procedure selected more features after DOSC because the reduced variation in the dataset enabled FDR to distinguish more potentially important features.

Misclassification rates obtained from four-fold CV are shown to evaluate the accuracy of classification trees and SVMs for each dataset (Table 1). For both classification trees and SVMs, the performance of classification using DOSC-processed datasets (Datasets 3, 4, 5) was significantly better than datasets without DOSC (Datasets 1 and 2). More importantly, the FDR-based feature selection method identified 1,458 features among 8,444 features, and classification trees and SVMs constructed with the features selected by FDR (Dataset 2) yielded smaller misclassification rates than those with all features (Dataset 1). Hence, the FDR-based feature selection method achieved a huge dimensionality reduction by eliminating non-informative features for classification, and yielded classification accuracy that is at least as good as using all features. Classification modeling with these reduced features reduces computational and analytic efforts. For the DOSC-processed datasets (Datasets 3, 4, 5), the application of DOSC post-FDR (Dataset 4) takes advantage of the dimensionality reduction (from 8,444 to 1,458) while achieving nearly as good misclassification rates as using all features (Dataset 3), and the application of FDR post-DOSC (Dataset 5) achieves dimensionality reduction (from 8,444 to 5,810) and equivalent misclassification rates compared with all features. The results also indicated that the performance of the different classification methods depended on different datasets. In general, SVMs perform better than classification trees except on Dataset 2. Overall, the FDR-based feature selection method reduced the dimensionality of the original data without degrading classification accuracy

5. Conclusion

¹H-NMR spectroscopy measures extract a considerable amount of the structural information available from a profile of a metabolite signal. Analysis of NMR spectra, accompanied by appropriate multivariate statistical techniques, has the potential to provide the tools necessary to facilitate visualization of inherent patterns by reducing the complexity of such information-rich data and discerning metabolic changes reflective of physiological variation, disease states, toxic stress, and nutrition intake.

In this paper, we proposed use of a multiple testing procedure controlling FDR to compensate for the limitations of PCA and PLS by achieving better interpretation of the feature selection results. The advantage of using the FDR-based feature selection method is fourfold. First, the identified features are not linear combination of a large number of original features; thus, the interpretation of results is straightforward. Second, the procedure utilizes class information to help select features better suited for classification. Third, the ability of the procedure to control FDR allows us to interpret the results efficiently. Fourth, the effect of correlated features can be automatically accommodated. In addition, the present study utilized a DOSC technique to enhance visualization, which leads to more accurate classification results.

The experimental results from real NMR spectra from human plasma demonstrate the effectiveness of FDR-based feature selection and DOSC to achieve better feature selection and classification. The overall results show that classification accuracy was improved by FDR and DOSC, while keeping only relevant features.

It should be noted that because our procedures are purely statistical methods based upon the available data, the metabolite features selected in the procedure may not be generally applicable to samples collected under other conditions. Further, in the

experimental design, giving 10 times more the RDA of SAA to one of two subject groups might be a lot in reality. Finally, the relatively small number of samples used in the present study may not allow the biological implications of our application to be generalized for practical applications. Nevertheless, we have shown that the procedures used here are potentially useful for the understanding of implicit metabolic patterns in biological systems and stimulate further investigation in the development of better analytic tools for the study of complex high-resolution NMR spectra.

Acknowledgments

We thank the editor and three anonymous referees for their constructive comments and suggestions, which greatly improved the quality of the paper. We are grateful to the nursing and laboratory staff of the Emory General Clinical Research Center for valuable help in collecting samples. S.B. Kim thanks U.T.-Arlington and the Office of the Provost for start-up funds.

References

1. Nicholson, J.K., et al., *Metabonomics: a platform for studying drug toxicity and gene function*. Nature Reviews Drug Discovery, 2002. **1**(2): p. 153-161.
2. Guyon, I. and A. Elisseeff, *An introduction to variable and feature selection*. Journal of Machine Learning Research, 2003. **3**: p. 1157-1182.
3. Jain, A.K., R.P.W. Duin, and J. Mao, *Statistical pattern recognition: A review*. IEEE Transactions on Pattern Analysis and Machine Intelligence, 2000. **20**: p. 4-37.
4. Goodacre, R., et al., *Chemometric discrimination of unfractionated plant extracts analyzed by electrospray mass spectrometry*. Phytochemistry, 2003. **62**(6): p. 859-863.
5. Sumner, L.W., P. Mendes, and R.A. Dixon, *Plant metabolomics: large-scale phytochemistry in the functional genomics era*. Phytochemistry, 2003. **62**(6): p. 817-836.
6. Beckonert, O., et al., *NMR-based metabonomic toxicity classification: hierarchical cluster analysis and k-nearest-neighbour approaches*. Analytica Chimica Acta, 2003. **490**(1-2): p. 3-15.
7. Farrant, R.D., et al., *An Automatic Data Reduction And Transfer Method To Aid Pattern-Recognition Analysis And Classification of NMR-Spectra*. Journal of Pharmaceutical and Biomedical Analysis, 1992. **10**(2-3): p. 141-144.
8. Lindon, J.C., E. Holmes, and J.K. Nicholson, *Pattern recognition methods and applications in biomedical magnetic resonance*. Progress in Nuclear Magnetic Resonance Spectroscopy, 2001. **39**(1): p. 1-40.

9. Hastie, T., R. Tibshirani, and J. Friedman, *The Element of Statistical Learning*. 2001, New York: Springer.
10. Tapp, H.S., M. Defernez, and E.K. Kemsley, *FTIR spectroscopy and multivariate analysis can distinguish the geographic origin of extra virgin olive oils*. Journal of Agricultural and Food Chemistry, 2003. **51**(21): p. 6110-6115.
11. Davis, R.A., et al., *Novel feature selection method for genetic programming using metabolomic ¹H NMR data*. Chemometrics and Intelligent Laboratory Systems, 2006. **81**: p. 50-59.
12. Koza, J.R., *Genetic Programming: On the Programming of Computers by Means of Nature Selection*. 1992, Cambridge, MA: MIT Press.
13. Gourvénec, S., X. Capron, and D.L. Massart, *Genetic algorithms (GA) applied to the orthogonal projection approach (OPA) for variable selection*. Analytica Chimica Acta, 2004. **519**: p. 11-21.
14. Shaffer, J.P., *Multiple Hypothesis-Testing*. Annual Review of Psychology, 1995. **46**: p. 561-584.
15. Benjamini, Y. and Y. Hochberg, *Controlling the false discovery rate - a practical and powerful approach to multiple testing*. Journal of the Royal Statistical Society Series B, 1995. **57**(1): p. 289-300.
16. Kim, S.B., K.-L. Tsui, and M. Borodovsky, *Multiple Hypothesis Testing in Large-Scale Contingency Tables: Inferring Pair-Wise Amino Acid Patterns in β -sheets*. Int. Journal of Bioinformatics Research and Applications, 2006. **2**: p. 193-217.
17. Storey, J.D. and R. Tibshirani, *Statistical significance for genomewide studies*. PNAS, 2003. **100**(16): p. 9440-9445.
18. Dudoit, S., J.P. Shaffer, and J.C. Boldrick, *Multiple hypothesis testing in microarray experiments*. Statistical Science, 2003. **18**(1): p. 71-103.
19. Genovese, C.R., N.A. Lazar, and T. Nichols, *Thresholding of statistical maps in functional neuroimaging using the false discovery rate*. Neuroimage, 2002. **15**(4): p. 870-878.
20. Holmes, E. and H. Antti, *Chemometric contributions to the evolution of metabonomics: mathematical solutions to characterising and interpreting complex biological NMR spectra*. Analyst, 2002. **127**(12): p. 1549-1557.

21. Wold, S., et al., *Orthogonal signal correction of near-infrared spectra*. Chemometrics and Intelligent Laboratory Systems, 1998. **44**: p. 175-185.
22. Beckwith-Hall, B.M., et al., *Application of orthogonal signal correction to minimise the effects of physical and biological variation in high resolution H-1 NMR spectra of biofluids*. Analyst, 2002. **127**(10): p. 1283-1288.
23. Lee, G.C. and D.L. Woodruff, *Beam search for peak alignment of NMR signals*. Analytica Chimica Acta, 2004. **513**(2): p. 413-416.
24. Benjamini, Y. and D. Yekutieli, *The control of the false discovery rate in multiple testing under dependency*. Annals of Statistics, 2001. **29**(4): p. 1165-1188.
25. Fearn, T., *On orthogonal signal correction*. Chemometrics and Intelligent Laboratory Systems, 2000. **50**(1): p. 47-52.
26. Sjoblom, J., et al., *An evaluation of orthogonal signal correction applied to calibration transfer of near infrared spectra*. Chemometrics and Intelligent Laboratory Systems, 1998. **44**(1-2): p. 229-244.
27. Westerhuis, J.A., S. de Jong, and A.K. Smilde, *Direct orthogonal signal correction*. Chemometrics and Intelligent Laboratory Systems, 2001. **56**(1): p. 13-25.
28. Trygg, J. and S. Wold, *Orthogonal projections to latent structures (O-PLS)*. Journal of Chemometrics, 2002. **16**(3): p. 119-128.
29. Svensson, O., T. Kourti, and J.F. MacGregor, *An investigation of orthogonal signal correction algorithms and their characteristics*. Journal of Chemometrics, 2002. **16**(4): p. 176-188.
30. Huo, X., et al., *A frontier-based tree pruning algorithm (FBP)*. INFORMS Journal on Computing, 2006. **18**(4): p. 494-505.
31. Breiman, L., et al., *Classification and Regression Trees*. 1984: Chapman & Hall.
32. Burges, C.J.C., *A tutorial on support vector machines for pattern recognition*. Data Mining and Knowledge Discovery, 1998. **2**: p. 1-43.
33. Cristianini, N. and J. Shawe-Taylor, *Support Vector Machines and Other Kernel-Based Learning Methods*. 2000: Cambridge University Press.
34. Jolliffe, I.T., *Principal Component Analysis (Second Edition)*. 2002, New York, NY: Springer.