

Ensemble of Dissimilarity Based Classifiers for Cancerous Samples Classification

Ángela Blanco, Manuel Martín-Merino¹, and Javier de las Rivas²

¹ Universidad Pontificia de Salamanca
C/Compañía 5, 37002, Salamanca, Spain
ablancogo@upsa.es, mmartinmac@upsa.es

² Cancer Research Center of Salamanca (CIC)
Salamanca, Spain
jrivas@usal.es

Abstract. DNA Microarray technology allow us to identify cancerous tissues considering the gene expression levels across a collection of related samples.

Several classifiers such as Support Vector Machines (SVM), k Nearest Neighbors (k -NN) or Diagonal Linear Discriminant Analysis (DLDA) have been applied to this problem. However, they are usually based on Euclidean distances that fail to reflect accurately the sample proximities. Several classifiers have been extended to work with non-Euclidean dissimilarities although none outperforms the others because they misclassify a different set of patterns.

In this paper, we combine different kind of dissimilarity based classifiers to reduce the misclassification errors. The diversity among classifiers is induced considering a set of complementary dissimilarities for three different type of models. The experimental results suggest that the algorithm proposed helps to improve classifiers based on a single dissimilarity and a widely used combination strategy such as Bagging.

1 Introduction

DNA Microarray technology allow us to monitor the expression levels of thousands of genes simultaneously across a collection of related samples. This technology has been applied particularly to the prediction of different type of cancer with encouraging results [12].

A large variety of machine learning techniques have been proposed to this aim such as Support Vector Machines (SVM) [10], k Nearest Neighbors [9] or Diagonal Linear Discriminant Analysis (DLDA) [9]. However the algorithms considered in the literature rely frequently on the use of the Euclidean distance that fails often to reflect accurately the proximities among the sample profiles [8,16,19]. The classifiers mentioned above have been extended to work with non-Euclidean dissimilarities [22]. In spite of this, the resulting algorithms misclassify a different set of patterns and fail to reduce significantly the errors. This can be explained because each dissimilarity reflects different features of the data and they induce different type of errors.

Several authors have pointed out that combining non-optimal classifiers can help to reduce particularly the variance of the predictor [17,24]. In order to achieve this goal, different versions of the classifier are usually built by sampling the patterns or the features [5]. Nevertheless, in our application, this kind of resampling techniques reduce the size of the training set. This may increase the bias of individual classifiers and the error of the combination [24].

In this paper we build the diversity of classifiers considering three different kinds of models such as SVM, k -NN and DLDA. The diversity is increased considering a set of complementary dissimilarities for each model. The classifiers induced will take advantage of the whole sample avoiding the bias introduced by resampling techniques such as Bagging. In order to incorporate non-Euclidean dissimilarities the base classifiers are modified in an appropriate way. Finally, the classifiers are aggregated using a voting strategy [17]. The method proposed has been applied to the prediction of different type of cancer using the gene expression levels with remarkable results.

This paper is organized as follows. Section 2 introduces the dissimilarities considered to build the diversity of classifiers. Section 3 comments how the classifiers can be extended to work from a dissimilarity matrix. In section 4 we present our combination strategy. Section 5 illustrates the performance of the algorithm in the challenging problem of gene expression data analysis. Finally, section 6 gets conclusions and outlines future research trends.

2 Dissimilarities for Gene Expression Data Analysis

An important step in the design of a classifier is the choice of a proper dissimilarity that reflects the proximities among the objects. However, the choice of a good dissimilarity is not an easy task. Each measure reflects different features of the data and the classifiers induced by the dissimilarities misclassify frequently a different set of patterns. Therefore no dissimilarity outperforms the others.

In this section, we comment shortly the main differences among several dissimilarities proposed to evaluate the proximity between cellular samples considering the gene expression levels. For a deeper description and definitions see [8,16,11].

The Euclidean distance evaluates if the gene expression levels differ significantly across different samples. An interesting alternative is the cosine dissimilarity. This measure will become small when the ratio between the gene expression levels is similar for the two samples considered. It differs significantly from the Euclidean distance when the data is not normalized by the L_2 norm.

The correlation measure evaluates if the expression levels of genes change similarly in both samples. Correlation based measures tend to group together samples whose expression levels are linearly related. The correlation differs significantly from the cosine if the means of the sample profiles are not zero. This measure is sensitive to outliers. The Spearman rank dissimilarity is less sensitive to outliers because it computes a correlation between the ranks of the gene expression levels. An alternative measure that helps to overcome the problem of outliers is the Kendall- τ index which is related to the Mutual Information probabilistic measure [11].

Due to the large number of genes, the sample profiles are codified in high dimensional and noisy spaces. In this case, the dissimilarities mentioned above are affected by the ‘curse of dimensionality’ [1,20]. Hence, most of the dissimilarities become almost constant and the differences among dissimilarities are lost [15]. To avoid this problem, it is recommended to reduce the number of features before computing the dissimilarities.

3 Dissimilarity Based Classifiers

Classical Support Vector Machines (SVM) [25] and Diagonal Linear Discriminant Analysis (DLDA) [9] are not able to work directly from a dissimilarity matrix. In this section, the classical SVM algorithm is extended to work from a dissimilarity matrix by defining a kernel of dissimilarities. Next DLDA is adapted following a different approach by embedding the patterns in a Euclidean space.

The SVM algorithm looks for a linear hyperplane $f(\mathbf{x}; \mathbf{w}) = \mathbf{w}^T \mathbf{x}$ that maximizes the margin $\gamma = 2/\|\mathbf{w}\|^2$. γ determines the generalization ability of the SVM. The slack variables ξ_i allow to consider classification errors. The figure 1 illustrates the meaning of the SVM parameters.

The hyperplane that minimizes the prediction error is given by the following optimization problem [25]:

$$\begin{aligned}
 \text{minimum}_{\mathbf{w}, \{\xi_i\}} \quad & \langle \mathbf{w}, \mathbf{w} \rangle + C \sum_{i=1}^n \xi_i^2 & (1) \\
 \text{subject to} \quad & y_i (\langle \mathbf{w}, \mathbf{x}_i \rangle + b) \geq 1 - \xi_i \quad i = 1, \dots, n \\
 & \xi_i \geq 0 \quad i = 1, \dots, n
 \end{aligned}$$

where C is a regularization parameter that achieves a balance between the empirical error and the complexity of the classifier. The optimization problem can




Fig. 1. Scheme of the hyperplane generated by the SVM algorithm for a non-linearly separable problem

be solved efficiently in dual space and the discriminant function can be expressed exclusively in terms of scalar products,

$$f(\mathbf{x}) = \sum_{\alpha_i > 0} \alpha_i y_i \langle \mathbf{x}, \mathbf{x}_i \rangle + w_0 \tag{2}$$

The SVM algorithm can be easily extended to the non-linear case substituting the scalar products by a Mercer kernel [25].

Non-Euclidean dissimilarities can be incorporated into the SVM algorithm by defining a kernel of dissimilarities [22,23]. Next we detail the idea.

Let d be a dissimilarity [6] and $R = \{p_1, \dots, p_n\}$ a subset of representatives drawn from the training set. Define the mapping $D(z, R) : \mathcal{F} \rightarrow \mathbb{R}^n$ as:

$$D(z, R) = [d(z, p_1), d(z, p_2), \dots, d(z, p_n)] \tag{3}$$

This mapping define a dissimilarity space where feature i is given by $d(\cdot, p_i)$.

The set of representatives R determine the dimensionality of the feature space. The choice of R is equivalent to select a subset of features in the dissimilarity space. Due to the small number of training samples in our application we have considered the whole sample as a representative set. It has been suggested in the literature that selecting a smaller subset of representatives does not help to improve the resulting classifier [22].

Once the patterns have been represented in the dissimilarity space, a kernel of dissimilarities can be defined as:

$$K_{ij} = \langle D(\mathbf{x}_i, R), D(\mathbf{x}_j, R) \rangle \tag{4}$$

where $\langle \cdot, \cdot \rangle$ denotes the scalar product in the feature space. Thus, for the linear SVM the kernel matrix is written as $K = DD^T$. This matrix is positive definite and keeps the nice properties of the optimization problem in the original SVM algorithm.

The DLDA is a variant of the Linear Discriminant Analysis (LDA) that considers diagonal and constant covariance matrices along the classes [9]. However, in order to apply this technique, a vectorial representation of the data should be obtained. To this aim, we follow the approach of [22]. First, the dissimilarities are embedded into an Euclidean space such that the inter-pattern distances reflect approximately the original dissimilarity matrix. Next, the test points are added to this space via a linear algebra operation. Finally the DLDA is applied considering the vectorial representation obtained.

We comment briefly the mathematical details of the embedding operation.

Let $D \in \mathbb{R}^{n \times n}$ be the dissimilarity matrix made up of the object proximities for the training set. A configuration in a low dimensional Euclidean space can be found via a metric multidimensional scaling algorithm (MDS) [6] such that the original dissimilarities are approximately preserved. Let $X = [\mathbf{x}_1 \dots \mathbf{x}_n]^T \in \mathbb{R}^{n \times p}$ be the matrix of the object coordinates for the training patterns. Define $B = XX^T$ as the matrix of inner products which is related to the dissimilarity matrix via the following equation:

$$B = -\frac{1}{2}JD^{(2)}J, \tag{5}$$

where $J = I - \frac{1}{n}\mathbf{1}\mathbf{1}^T \in \mathbb{R}^{n \times n}$ is the centering matrix, I is the identity matrix and $D^{(2)} = (\delta_{ij}^2)$ is the matrix of the square dissimilarities for the training patterns. If B is positive semi-definite, the object coordinates in the low dimensional Euclidean space \mathbb{R}^k can be found through a singular value decomposition [6,13]:

$$X_k = V_k \Lambda_k^{1/2}, \quad (6)$$

where $V_k \in \mathbb{R}^{n \times k}$ is an orthogonal matrix with columns the first k eigen-vectors of XX^T and $\Lambda_k = \text{diag}(\lambda_1 \dots \lambda_k) \in \mathbb{R}^{k \times k}$ is a diagonal matrix with λ_i the i -th eigenvalue. Several dissimilarities introduced in section 2 generate inner product matrices B non semi-definite positive. Fortunately, the negative values are small in our application and therefore can be neglected [6] without losing relevant information about the data.

Once the training patterns have been embedded into a low dimensional Euclidean space, the test pattern can be added to this space via a linear projection [22]. Next we comment briefly the derivation.

Let $X_k \in \mathbb{R}^{n \times k}$ be the object configuration for the training patterns in \mathbb{R}^k and $X_n = [\mathbf{x}_1 \dots \mathbf{x}_s]^T \in \mathbb{R}^{s \times k}$ the matrix of the object coordinates sought for the test patterns. Let $D_n^{(2)} \in \mathbb{R}^{s \times n}$ be the matrix of the square dissimilarities between the s test patterns and the n training patterns that have been already projected. The matrix $B_n \in \mathbb{R}^{s \times n}$ of inner products among the test and training patterns can be found as:

$$B_n = -\frac{1}{2}(D_n^{(2)}J - UD^{(2)}J), \quad (7)$$

where $J \in \mathbb{R}^{n \times n}$ is the centering matrix and $U = \frac{1}{n}\mathbf{1}^T \mathbf{1} \in \mathbb{R}^{s \times n}$. The derivation of equation (7) is detailed in [22]. Since the matrix of inner products verifies

$$B_n = X_n X_k^T \quad (8)$$

then, X_n can be found as the least mean-square error solution to (8), that is:

$$X_n = B_n X_k (X_k^T X_k)^{-1}, \quad (9)$$

Given that $X_k^T X_k = \Lambda_k$ and considering that $X_k = V_k \Lambda_k^{1/2}$ the coordinates for the test points can be obtained as:

$$X_n = B_n V_k \Lambda_k^{-1/2}, \quad (10)$$

which can be easily evaluated through simple linear algebraic operations.

4 Combination of Dissimilarity Based Classifiers

In this section we introduce our ensemble of classifiers to reduce the errors and comment briefly the related work.

Our method builds the diversity of classifiers considering three different kind of models such as SVM, k -NN and DLDA. To increase the diversity among




Fig. 2. Aggregation of classifiers using a voting strategy. Bold patterns are misclassified by a single hyperplane but not by the combination.

classifiers we have considered several dissimilarities introduced in section 2. Each dissimilarity reflects different features of the data and the resulting classifiers will produce different errors. Thus, the combination will improve the performance of classifiers based on single dissimilarity [5,18]. Besides, the diversity of classifiers is generated considering the whole training sample. In this way, we avoid to reduce the size of the training set which may induce bias in the individual classifiers. Notice that the combination strategies are not able to reduce the bias of single classifiers [24].

Figure 2 shows in an intuitive way how the combination of classifiers reduces the misclassification errors. For instance bold patterns are assigned to the wrong class by one classifier but using a voting strategy the patterns will be assigned to the right class.

Hence, our combination algorithm proceeds as follows: First, the set of complementary dissimilarities introduced in section 2 are computed. As we mentioned earlier each classifier incorporates the dissimilarities in a different way. For the SVM algorithm, the kernel of dissimilarities is computed and the optimization problem is solved in the usual way. k -NN is able to work directly from a dissimilarity matrix but to avoid the ‘curse of dimensionality’ and to increase the diversity among dissimilarities it is recommended to reduce previously the number of features. For the DLDA algorithm, the dissimilarities should be embedded in an Euclidean space via a Multidimensional Scaling algorithm. The ensemble of classifiers is aggregated by a standard voting strategy [17]. The diagram 1 shows the steps of the algorithm.

A related technique to combine classifiers is the Bagging [5,3]. This method generates a diversity of classifiers considering several bootstrap samples as training sets. Next, the classifiers are aggregated using a voting strategy. Nevertheless there are three important differences between Bagging and the method proposed in this section.

First, our method generates the diversity of classifiers by considering the whole sample. Bagging trains each classifier using around 63% of the training set. In our application the size of the training set is very small and neglecting part of the patterns may increase the bias of each classifier. It has been suggested in

Algorithm 1. Aggregation of classifiers based on multiple models and dissimilarities

- 1: For each measure compute the dissimilarity matrix
 - 2: Compute the kernel of dissimilarities using equation (4) for the SVM algorithm
 - 3: Embed each dissimilarity into an Euclidean space through equation (6) for DLDA algorithm
 - 4: Train the classifiers for each dissimilarity
 - 5: Combine the different models using a voting strategy
 - 6: Evaluate the ensemble by ten-fold cross-validation. Test points are embedded for DLDA algorithm using equation (10)
 - 7: End
-

the literature that Bagging does not help to reduce the bias [24] and so, the aggregation of classifiers will hardly reduce the misclassification error.

A second advantage of our method is that it is able to work directly with a dissimilarity matrix.

Finally, the combination of several dissimilarities avoids the problem of choosing a particular dissimilarity for the application we are dealing with. This is a difficult and time consuming task.

5 Experimental Results

In this section, the ensemble of classifiers proposed is applied to the identification of cancerous samples using Microarray gene expression data.

Three benchmark gene expression datasets have been considered. The first one consisted of 72 bone marrow samples (47 ALL and 25 AML) obtained from acute leukemia patients at the time of diagnosis [13]. The RNA from marrow mononuclear cells was hybridized to high-density oligonucleotide microarrays produced by Affymetrix and containing 6817 genes. The second dataset consisted of 49 samples from breast tumors [26], 25 classified as positive to estrogen receptors (ER+) and 24 negative to estrogen receptors (ER-). Those positive to estrogen receptors require a different treatment. The RNA of breast cancer cells were hybridized to high-density oligonucleotide microarrays produced by Affymetrix and containing 7129 genes. Finally the third dataset consists of 40 tumor and 22 normal colon samples, analyzed with an Affymetrix oligonucleotide array complementary to more than 6,500 human genes. The number of genes was reduced in the original dataset to 2000 [2].

Due to the large number of genes, samples are codified in a high dimensional and noisy space. Therefore, the dissimilarities are affected by the 'curse of dimensionality' and the correlation among them becomes large [20]. To avoid this problem and to increase the diversity among dissimilarities we have reduced aggressively the number of genes using the standard F-statistic [11]. The number of genes considered for SVM and DLDA are 14% while for k -NN the number of genes kept is 3% because this technique is more sensible to noise.

The dissimilarities have been computed without normalizing the variables because as we have mentioned in section 2 this operation may increase the correlation among them.




Fig. 3. Curve of eigenvalues for the Multidimensional Scaling algorithm and the χ^2 dissimilarity

The algorithm chosen to train the Support Vector Machines is C-SVM. The C regularization parameter has been set up by ten fold-crossvalidation [21,4]. We have considered linear kernels in all the experiments because the small size of the training set in our application favors the overfitting of the data. Consequently error rates are smaller for linear kernels than for non linear ones.

The number of neighbors for k -NN algorithm is estimated by cross-validation.

Before applying DLDA the dissimilarities should be embedded in an Euclidean space using a Multidimensional Scaling algorithm. An important parameter is the dimensionality of this space generated by the first eigen-vectors of the inner product matrix (5). The number of eigenvectors considered is determined by the curve of eigenvalues.

Figure 3 shows the eigenvalues for the breast cancer data and the χ^2 dissimilarity. The first eleven eigenvalues account for 85% of the variance. Therefore, they preserve the main structure of the data.

The algorithms have been evaluated considering the global errors and the false negative errors. Both have been estimated by ten-fold cross-validation which gives good experimental results for the problem at hand [21].

Table 1 shows the experimental results for the best single classifier for each technique. Table 2 compares the method proposed with Bagging, introduced in section 3. Both, Bagging and the best classifiers based on a single dissimilarity for each model have been taken as a reference.

From the analysis of tables 2 and 2, the following conclusions can be drawn:

- The dissimilarity that minimizes the error depends strongly on the classifier and on the particular dataset considered. No dissimilarity outperforms the others for a wide range of models and datasets. Hence the choice of a proper dissimilarity is not an easy task for human experts.

Table 1. Empirical results for the best single classifier for each technique

Technique	Datasets	Error %	False negative %
SVM (Correlation)	Golub	6.94%	2.77%
SVM (Tau)	Breast	6.12%	2.04%
SVM(Correlation)	Colon	14.5%	6.45%
K -NN (Tau)	Golub	1.38%	1.38%
K -NN(Tau)	Breast	8.16%	2.04%
K -NN (Cosine)	Colon	12.9%	4.83%
DLDA (Tau)	Golub	2.77%	1.38%
DLDA(Spearman)	Breast	8.16%	2.04%
DLDA (Euclidean)	Colon	11.29%	4.83%

Table 2. Empirical results for the combination of classifiers. The Bagging technique has been taken as reference.

Technique	Datasets	Error %	False negative %
	Golub	8.33%	6.94%
Bagging (SVM)	Breast	6.12%	2.04%
	Colon	12.9%	4.83%
	Golub	5.55%	5.55%
Bagging (k -NN)	Breast	14.28%	6.12%
	Colon	14.51%	9.67%
	Golub	6.94%	4.16%
Bagging (DLDA)	Breast	14.28%	2.04%
	Colon	11.29%	3.22%
	Golub	1.38%	1.38%
Combination	Breast	4.08%	2.04%
	Colon	11.2%	3.22%

- The combination strategy proposed outperforms significantly the misclassification errors of the best single classifiers. In particular, the ensemble of classifiers improves significantly the SVM algorithms for the three problems considered. False negative errors are particularly reduced in Golub and Colon datasets. We also report that our method improves the best k -NN classifier for Breast and Colon that are the most complex according to the literature. Finally, DLDA is also improved for Golub and Breast Cancer.
- The ensemble of classifiers proposed improves a widely used combination algorithm such as Bagging. Both kind of errors are particularly reduced for Golub and Breast Cancer. This result supports the idea that our algorithm performs better than the resampling techniques when the sample size is small.

6 Conclusions and Future Research Trends

In this paper, we have proposed an ensemble of classifiers based on a diversity of models and dissimilarities. Our approach aims to reduce the misclassification

error of classifiers based solely on a single measure. The algorithm has been applied to the classification of cancerous samples using gene expression data.

The experimental results suggest that the method proposed improves the misclassification error of classifiers based on a single dissimilarity. We also report that our method compares favorably with a widely used combination algorithm such as Bagging.

As future research trends, we will try to extend the method proposed to improve clustering algorithms.

Acknowledgment

This work has been partially supported by the Junta de Castilla y León grant PON05B06.

References

1. Aggarwal, C.C.: Re-designing distance functions and distance-based applications for high dimensional applications. In: Proc. of the ACM International Conference on Management of Data and Symposium on Principles of Database Systems (SIGMOD-PODS), vol. 1, pp. 13–18 (March 2001)
2. Alon, U., Barkai, N., Notterman, D.A., Gish, K., Ybarra, S., Mack, D., Levine, A.J.: Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. Proc. Nat'l. Acad. Sci. USA 96, 6745–6750 (1999)
3. Bauer, E., Kohavi, R.: An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. Machine Learning 36, 105–139 (1999)
4. Braga-Neto, U., Dougherty, E.: Is cross-validation valid for small-sample microarray classification? Bioinformatics 20(3), 374–380 (2004)
5. Breiman, L.: Bagging predictors. Machine Learning 24, 123–140 (1996)
6. Cox, T., Cox, M.: Multidimensional Scaling, 2nd edn. Chapman & Hall/CRC Press, New York (2001)
7. Cristianini, N., Shawe-Taylor, J.: An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press, Cambridge (2000)
8. Drăghici, S.: Data Analysis Tools for DNA Microarrays. Chapman & Hall/CRC Press, New York (2003)
9. Dudoit, S., Fridlyand, J., Speed, T.: Comparison of discrimination methods for the classification of tumors using gene expression data. Journal of the American Statistical Association 97, 77–87 (2002)
10. Furey, T., Cristianini, N., Duffy, N., Bednarski, D., Schummer, M., Haussler, D.: Support vector machine classification and validation of cancer tissue samples using microarray expression data. Bioinformatics 16(10), 906–914 (2000)
11. Gentleman, R., Carey, V., Huber, W., Irizarry, R., Dudoit, S.: Bioinformatics and Computational Biology Solutions Using R and Bioconductor. Springer, Heidelberg (2006)
12. Golub, T., Slonim, D., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J., Coller, H., Loh, M., Downing, J., Caligiuri, M., Bloomfield, C., Lander, E.: Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring. Science 286(15), 531–537 (1999)

13. Golub, G.H., Loan, C.F.V.: *Matrix Computations*, 3rd edn. Johns Hopkins university press, Baltimore, Maryland, USA (1996)
14. Guyon, I., Weston, J., Barnhill, S., Vapnik, V.: Gene selection for cancer classification using support vector machines. *Machine Learning* 46, 389–422 (2002)
15. Hinneburg, C.C.A.A., Keim, D.A.: What is the nearest neighbor in high dimensional spaces? In: *Proc. of the International Conference on Database Theory (ICDT)*, pp. 506–515. Morgan Kaufmann, Cairo, Egypt (2000)
16. Jiang, D., Tang, C., Zhang, A.: Cluster analysis for gene expression data: A survey. *IEEE Transactions on Knowledge and Data Engineering* 16(11) (November 2004)
17. Kittler, J., Hatef, M., Duin, R., Matas, J.: On combining classifiers. *IEEE Transactions on Neural Networks* 20(3), 228–239 (1998)
18. Kuncheva, L.I.: *Combining Pattern Classifiers*. John Wiley, New Jersey (2004)
19. Martín-Merino, M., Muñoz, A.: Self organizing map and sammon mapping for asymmetric proximities. *Neurocomputing* 63, 171–192 (2005)
20. Martín-Merino, M., Noz, A.M.: A new sammon algorithm for sparse data visualization. In: *International Conference on Pattern Recognition (ICPR)*, pp. 477–481. IEEE Press, Cambridge (UK) (2004)
21. Molinaro, A., Simon, R., Pfeiffer, R.: Prediction error estimation: a comparison of resampling methods. *Bioinformatics* 21(15), 3301–3307 (2005)
22. Pekalska, E., Paclick, P., Duin, R.: A generalized kernel approach to dissimilarity-based classification. *Journal of Machine Learning Research* 2, 175–211 (2001)
23. Schölkopf, B., Smola, A.: *Learning with Kernels*. MIT Press, Cambridge, USA (2002)
24. Valentini, G., Dietterich, T.: Bias-variance analysis of support vector machines for the development of svm-based ensemble methods. *Journal of Machine Learning Research* 5, 725–775 (2004)
25. Vapnik, V.: *Statistical Learning Theory*. John Wiley & Sons, New York (1998)
26. West, M., Blanchette, C., Dressman, H., Huang, E., Ishida, S., Spang, R., Zuzan, H., Olson, J., Marks, J., Nevins, J.: Predicting the clinical status of human breast cancer by using gene expression profiles. *PNAS* 98(20) (September 2001)