

# Connectionist Approaches for Predicting Mouse Gene Function from Gene Expression

Emad Andrews Shenouda, Quaid Morris, Anthony J. Bonner

Department of Computer Science  
University of Toronto  
Toronto, ON.  
[emad@cs.toronto.edu](mailto:emad@cs.toronto.edu)

**Abstract.** Identifying gene function has many useful applications especially in Gene Therapy. Identifying gene function based on gene expression data is much easier in prokaryotes than eukaryotes due to the relatively simple structure of prokaryotes. That is why tissue-specific expression is the primary tool for identifying gene function in eukaryotes. However, recent studies have shown that there is a strong learnable correlation between gene function and gene expression. This paper outlines a new approach for gene function prediction in mouse. The prediction mechanism depends on using Artificial Neural Networks (NN) to predict gene function based on quantitative analysis of gene co-expression. Our results show that neural networks can be extremely useful in this area. Also, we explore clustering of gene functions as a preprocessing step for predicting gene function.

## 1 Introduction

Gene function prediction is one of the primary goals of Bioinformatics. Identifying gene function can be extremely useful in many ways, especially in Gene Therapy [5]. Identifying gene function in prokaryotes is much easier than eukaryotes due to their lower structural complexity and number of genes. Tissue-specific gene expression is the most widely used predictor for gene function in mammals; for example, genes expressed in tongue are most probably involved in tasting. However, this method has not been scientifically justified. Recent studies showed that there is a strong learnable correlation between gene function and gene expression [16]. According to Zhang et al, predicting gene function based on its expression is more effective than using tissue-specific function as a guide [16]. In this paper, we explore a new approach for gene function prediction in mouse based on quantitative analysis of gene co-expression. We will try to learn this correlation using machine learning techniques to predict gene function category from gene expression.

The machine learning technique used in this work is Neural Networks (NN). As we will describe in the next section, Support Vector Machines (SVM) and graphical models have been used in related work. In spite of their scalability and huge learning and classification capabilities, (NN) has not, to our knowledge, been used to predict gene function in higher order organisms than yeast. According to Baldi, (NN) are

**Emad Andrews Shenouda, Quaid Morris, Anthony J. Bonner**

superior classification machines, and in theory they can approximate *any* reasonable function to *any* degree of required precision [2]. In the same source, authors showed that (NN), with the right modeling, can be extremely useful in various Bioinformatics problems, such as sequence encoding and correlation, prediction of protein secondary structure, and prediction of signal peptides and their cleavage sites.

In this work we used the back-propagation algorithm (BP) to predict gene functions from gene expression. The earliest description of (BP) was presented by Werbos in 1974 [15]; then, it was independently rediscovered by Le Cun et al [7]. Multilayer perceptrons (MLP) is perhaps the most famous implementation of (BP). It has been very successfully used in various domains, such as prediction, function approximation and classification. For classification, MLP is considered a super-regression machine that can draw complicated decision borders between nonlinearly separable patterns [6]. The nonlinearity power of MLP is due to the fact that all its neurons use a nonlinear activation function to calculate their outputs.

Unlike (SVM), MLP can be implemented in two versions: first is the Softmax version, which has only one output neuron to be activated with a certain input, while the other has more than one output activated, which is a very powerful feature of MLP as it allows the possibility that a gene might be involved in more than one function. In this work we used MLP in three different ways: Firstly, we used it as a binary classifier, in which the input is gene expression measurements and the desired output is a single GO-category. Having 922 different GO-categories, this will result in 922 binary classifiers. In the second approach, we unleashed the power of the MLP as a super regression machine by training it with the same input as before but with all 922 GO-categories as the desired output. This approach not only discovers the relation among co-expressed genes but also ties those expressions to many possible functions. Finally, we tried clustering gene annotations into distinct groups using the K-means algorithm; then, we used those groups as the desired output for the network. For a full explanation about K-means the reader is directed to Mitchell [9]

Our results show that NN is a good candidate for predicting gene function from expression. As a binary classifier, MLP is very successful and has the advantage of eliminating the need for the kernel function needed by SVM. As a super regression machine, MLP was able to predict the whole set of 922 possible functions for each gene but with the drawback of lots of training time needed. Due to the nature of the data, as we will see later, grouping gene annotations method did not produce as good results as the first two methods.

## **1.1 Related work**

Using gene expression for predicting gene functions has been proven to be effective in prokaryotes. For eukaryote, tissue-specific expression is widely used for predicting gene functions [16]. However, the correlation between gene expression, gene sequence and gene functions has attracted the attention of many bioinformaticians. In [3], the authors used a probabilistic approach to correlate gene sequence and expression. In [5], there is an investigation of gene function by identifying interactions between a protein and other macromolecules. Some bioinformaticians tried to predict gene functions only from the pattern of the GO-annotation categories

## Connectionist Approaches for Predicting Mouse Gene Function from Gene Expression

vocabulary mentioned at the Gene Ontology (GO) database without considering the gene expression. In [11], Oliver et al used Bayesian networks and decision trees to model the relationship among different GO-categories vocabulary, only 41 genes out of 100 manually assessed were judged to be true.

SVMs are the most widely used machine learning approach to predict genes functions from genes sequences and expressions. In [8], the authors used SVM to functionally classify genes by using gene expression data from DNA microarray. They used gene expression of 2,467 genes with known function from the budding yeast *Saccharomyce cerevisiae* measured in 79 different DNA microarrays. Still, this work used yeast genes not higher order organisms; also, the classification is binary and the performance is assessed against other basic classification algorithms like Fisher's linear discriminant and decision trees.

In 2004, Arunachalam et al used SVM to predict gene function based on the nucleotide sequences mentioned in the (GO) database. Although there is no gene expression involved, this research is relevant because it uses the same method to predict gene function from tissues taken from different species ranging from as simple as yeast to as complicated as mouse [1].

The first attempt to use NN to predict gene function class from gene expression was in 2002 by Mateos et al [18]. In that work, the classes used to train the MLP were taken from the Yeast Genome Database at the Munich Information Center for Protein Sequences (MIPS) functional catalog.

The most related work is by Zhang et al in [16], where gene expression is used as SVM input to predict gene function category. This work is very important in many ways. First of all, it is for mouse genes which are homologous to human genes. Also, authors showed the shortcomings of depending on tissue-specific information to predict gene function and the crucial need for another method. The dataset of the microarray generated for this work is a public resource for mammalian functional genomics. The authors stated four reasons that support the integrity and validity of their data. Finally, it uses specific 922 GO-categories mentioned at [10]. Again, authors provided a mechanism to measure the accuracy of their predictions through a recall value. The reader is directed to [13] and [16] for full details about this work.

Instead of using SVM, we used MLP for prediction and K-means for clustering. Our work goes beyond binary classification by predicting all possible functions the gene might be involved in simultaneously, which expresses the relation between the expression level and all possible functions. To decrease the required training time, we used K-means to group genes based on their annotations into distinct annotation groups and then trained the MLP to predict gene group number instead of predicting the whole 922 vector of annotations.

## 2 Data preparation

Data used for this paper is provided by Zhang et al [13]. The reader is directed to [16] for a complete explanation of the data gathering process. The authors designed their own microarray to contain nearly 40,000 genes from 55 different mouse tissues. From those 40,000 only 21,266 confidently detected transcripts were extracted. Those are

**Emad Andrews Shenouda, Quaid Morris, Anthony J. Bonner**

the ones that exceeded the 99<sup>th</sup> percentile of intensities from the negative controls. From those 21,266 genes, only 7,388 genes have at least one specific Go-category; the rest were considered negative genes because their annotations were too general. We will refer to these 7,388 genes as the *annotated genes*.

Our data preparation process has three stages: first, extracting the expression microarray data of the annotated genes; second, constructing a binary matrix that describes the annotations associated with each of those genes, and finally clustering those binary vectors.

As an input to this data preparation process, two matrices were downloaded from [13]: the (21,266 X 55) matrix of the microarray data which is normalized and centered by subtracting the mean, and the two-column annotation matrix which has genes in one column and their annotations in the other column. Then, the two-column annotation file is parsed to extract all distinct genes and all distinct annotations. Finally, a two dimensional binary matrix (7,388 X 922) is created; in which, each annotated gene has a binary vector denoting all of the 922 annotations. If the gene is annotated in a Go-Category, its corresponding bit is set to 1 otherwise to 0.

This binary matrix is the desired output of the MLP. If we chose only one Go-category as a desired output, the MLP will work as a binary classifier; if we considered more than one column as an output, the MLP will have number of output neurons equal to the number of GO-Categories we are considering.

### **3 Neural network topology selection**

Selecting an appropriate network topology is one of the main difficulties of using NN in classification; that is why Wang and Huang used Extreme Learning Machine (ELM) algorithm instead of NN in spite of the potential promising results of NN in classifying protein sequence [14]. In addition to network parameters like learning rate and momentum, NN topology is also determined by its size, synaptic weights connections, and the hidden-units activation function. By network size we mean the number of hidden layers and number of hidden units in each layer. Network size is a measurement of the system complexity and it is directly proportional to the training time required. The less complex the network is, the less its tendency to memorize the training set. That is why MLP size reduction is always recommended when possible [6], [17]. In general, the most common way to reduce network size is weight pruning as mentioned by Zurada and Haykin [6], [17], which suggests removing the ineffective weights during training. Instead of pruning only the weights, we started with number of hidden units equal to twice the number of the input features. Then, without affecting the cross validation (CV) error, we kept dividing this number by 2. The number of input neurons is constant in all our networks and equal to the number of tissues which is 55.

Concerning the synaptic weights connections, there are two types of MLPs: the standard MLP in which each layer is fully connected to its next layer only; the second type is the Generalized Feed Forward (GFF) network which is a generalization of the MLP such that connections can jump over one or more layers. In theory, an MLP can solve any problem that a generalized feedforward network can solve [6]. In practice,

## Connectionist Approaches for Predicting Mouse Gene Function from Gene Expression

however, generalized feedforward networks often solve the problem more efficiently. In our experiments, in the case of group number prediction, GFF required fewer training epochs than the MLP and it showed fewer tendencies to memorize. We used the sigmoidal activation function in all our work because no negative output is desired. Also, sigmoidal activation function is superior in classification [4], [12], [14].

### 3.1 Regularization, convergence criteria and error measurement

Cross Validation (CV) is our main method for regularization and stopping. Before training starts, the data set is divided into 3 subsets: validation set, 5 % of the whole dataset; testing set, 15% of the dataset and the remaining examples are used for training. Once the network starts, its initial state after 50 epochs is recorded and the whole training stops if the CV Mean Square Error (MSE) did not improve for 100 consecutive epochs. The best weights are saved once the CV starts to increase and used for testing.

For error measurement, we based our comparison on the (MSE) because the actual error rate of classification sometimes becomes misleading depending on the output values [4]. The (MSE) is twice the average cost.

## 4 Neural networks as a binary classifier

In this approach we used NN as a binary classifier machine. The input is the gene microarray expression level and the output is 1 or 0 depending on whether the gene is predicted to be in the desired Go-Category or not, respectively. A separate NN is needed for each gene function. So, 922 binary classifiers are needed to predict all GO-Categories using this approach. Instead of using only one output for the network, we used two binary output neurons to well assess the network accuracy. The ordered pair (0, 1) means the same as 0 and (1, 0) means the same as 1.

Due to space limitations, we present only prediction results for the GO-Categories with maximum gene participation. Those categories are the best example of the network performance in this type of problem as they provide the maximum generalization and hence are the hardest to predict. Table 1 summarizes the prediction results for the selected GO-Categories. With some statistical data analysis from the above prepared data, we found most of the categories have gene participation of less than 20; the maximum participation is 455 genes in the “*lipid metabolism [GO:0006629]*” category. It is clear from the way we prepared our data that the minimum participation will be 1 and never 0. Figure 1 shows the overall gene participation distribution. This highly skewed and non-uniform distribution of the desired output might make the network performance questionable because the maximum gene participation in any category is 6.15%. So in many cases the network might tend to sacrifice the positive output for the sake of negative output and still the MSE will be low.




Fig. 1. Overall gene participation distribution for all categories

#### 4.1 More expressive dataset

To prove our hypothesis and provide a fair assessment of the NN performance in this problem without getting affected by the skewed data distribution mentioned above, we prepared a balanced subset from our original data. The number of examples in this dataset is 990, containing nearly equal numbers of positive and negative genes for the “lipid metabolism [GO:0006629]” category. The network recorded a much lower training MSE because of the fewer training examples. The actual output is in table 2.

#### 4.2 Results

We did binary classification for GO-Categories with a gene participation of at least 350. That gave us 10 categories to predict. The following table summarizes the results for the 10 selected Go-Categories: where, A is the number of positive genes, B is the training MSE, C is the CV MSE, D is the testing MSE and E is the training time in seconds. (Note: cells B, C, D are divided by  $10^{-3}$ )

Table 1. Binary prediction results for the selected GO-Categories

Function Name	Go-category	A	B	C	D	E
lipid metabolism	GO:0006629	455	47.726	74.236	43.3793	195
intracellular protein transport	GO:0006886	454	46.539	35.539	42.299	1054
carbohydrate metabolism	GO:0005975	446	45.141	34.291	51.039	19281
cation transport	GO:0006812	410	42.293	42.405	42.220	944
response to abiotic stimulus	GO:0009628	396	38.608	27.702	58.423	38148
response to pest/pathogen/parasite	GO:0009613	394	42.712	65.791	32.425	480
cytoskeleton org. and biogenesis	GO:0007010	392	41.041	29.260	41.503	1212
neurogenesis	GO:0007399	371	46.374	46.374	33.157	384
cell-cell signaling	GO:0007267	361	38.951	45.093	32.468	695
mitotic cell cycle	GO:0000278	350	36.708	21.205	50.843	42360

The above table reflects an acceptable performance for the MLP to predict gene function from its expression level. But, we used a balanced dataset from the lipid

## Connectionist Approaches for Predicting Mouse Gene Function from Gene Expression

metabolism [GO:0006629] category to assess the MLP performance in case of a balanced dataset. In this dataset, the NN achieved a training MSE of  $10^{-7}$  in approximately 10 hours of training. The testing MSE was 0.247147. The actual output of the testing results is summarized in the following table where D and A stand for the desired and actual outputs respectively.

**Table 2.** Actual NN results for the balanced sub-dataset

D 0	D 1	A 0	A 1	D 0	D 1	A 0	A 1	D 0	D 1	A 0	A 1
1	0	0.999215	0.000785	1	0	0.429985	0.570019	1	0	0.999167	0.000833
1	0	1.001267	0.00127	0	1	0.391373	0.60863	1	0	0.00031	1.000306
0	1	0.0012	1.001202	1	0	0.999388	0.000613	0	1	0.994135	0.005865
0	1	0.00018	1.000176	0	1	0.999378	0.000622	1	0	0.998971	0.00103
1	0	1.000075	7.4E-05	0	1	0.09959	0.900412	0	1	0.997675	0.002326
0	1	8E-06	1.000009	1	0	0.0014	1.001399	1	0	0.00956	1.009562
1	0	0.777683	0.222318	0	1	0.999638	0.000363	0	1	0.998558	0.001442
1	0	0.986617	0.013384	1	0	0.999047	0.000953	0	1	0.251832	0.748172
0	1	0.999549	0.000451	0	1	0.00318	1.003182	1	0	0.999172	0.000828
1	0	1.001881	0.00188	0	1	0.310283	0.689721	0	1	0.966519	0.033481
1	0	0.00018	1.000177	0	1	1.003398	0.0034	1	0	0.125626	0.874377
1	0	0.999156	0.000844	1	0	0.999119	0.000882	0	1	0.000196	0.999805
1	0	0.033314	0.966688	0	1	0.999297	0.000703	1	0	0.982307	0.017694
1	0	0.983298	0.016702	0	1	0.00094	1.000937	1	0	0.999131	0.000869
0	1	0.000224	0.999776	0	1	0.000332	0.999668	1	0	0.999056	0.000945
1	0	0.999263	0.000738	1	0	1.002547	0.00255	0	1	0.999296	0.000705
1	0	1.000588	0.00059	1	0	1.000438	0.00044	1	0	0.000139	0.999862
0	1	1.0005	0.0005	0	1	0.000245	0.999756	1	0	0.999218	0.000782
1	0	0.998465	0.001536	0	1	0.994506	0.005494	1	0	0.998925	0.001076
0	1	0.00455	1.004549	1	0	1.000548	0.00055	1	0	0.999724	0.000276
1	0	0.999284	0.000716	0	1	0.999474	0.000527	1	0	0.998531	0.00147
0	1	0.00019	1.000193	1	0	0.999405	0.000596	1	0	0.999211	0.000789
1	0	0.977609	0.022392	1	0	0.99844	0.001561	1	0	0.231151	0.768853
1	0	0.999154	0.000846	1	0	0.999442	0.000559	1	0	0.117864	0.882139
1	0	0.999457	0.000542	0	1	0.971673	0.028326	0	1	1.00252	0.00252
0	1	0.000218	0.999782	1	0	0.999129	0.000871	1	0	1.001117	0.00112
0	1	0.000442	0.999559	1	0	1.000617	0.00062	0	1	0.994494	0.005506
0	1	0.999625	0.000375	1	0	0.983058	0.016942	0	1	0.00215	1.002151

## 5 Neural network with 1000 output (Non-Softmax)

In this approach we tried to use the full capability of NN to simultaneously predict all functions the gene might be involved in as a binary vector. The same input matrix is used but the desired output is a binary vector of 922 bits. Although the results are promising, the training time was huge. We concentrated on reducing the network size to the lowest we can without affecting either the network initial state or the CV MSE. The minimum topology was 3 hidden layers with 200, 100, and 50 hidden units (from input to output), and 922 output units. After nearly 60 hours of training, we obtained training MSE of 0.061335, CV MSE of 0.061616 and testing MSE of 0.061778.

Emad Andrews Shenouda, Quaid Morris, Anthony J. Bonner

We used batch learning to reduce the training time; the network then consumed 1059 epochs in 211673 seconds which means about 199 sec/epoch. Those results along with the CV error progress show that this method is effective but requires a large sized network and huge training time which needs specialized hardware to produce the best network results in a reasonable amount of time.

## 6 Neural network with K-means as a classifier

In this approach we tried to overcome the main drawback of the previous approach which is the huge training time required. We used a sequential version of k-means to group the binary vectors denoting the genes possible annotations into k groups. Instead of predicting the whole binary vector, the MLP tries to predict the group number the gene belongs to based on its microarray expression level. Biologically, this is still very useful as it predicts all possible annotations a certain gene might have.

### 6.1 Results

As expected, the problem became an easier one for the NN to solve and the training time required was far less than the previous method. However, the network initial state, in terms of MSE and error percentage, was worse. By monitoring the CV error progress it was easy to see that the network started to memorize the training set after the first few epochs. The best results were obtained by using a 2-hidden layer GFF version, which delayed the CV error increase with a considerable improvement of test MSE. Both test MSE and CV MSE were higher than the training MSE. The following table summarizes the results for 4 different groupings (4 different values of k).

Table 3: Results for 4 different GO-Categories groups

# of Groups(k)	Training MSE	CV MSE	Testing MSE	Training time in sec
500	0.096455	7.350921	7.081350	1660
300	0.110561	0.179318	0.179431	7560
200	0.121474	0.294703	0.252912	790
100	0.096305	6.743814	5.059804	24240

This poor performance can be explained by the grouping distribution. As in the following histograms (figures 2 and 3), the first group always dominates and has nearly 50% of the total number of vectors regardless of the number of means specified. Some other groups have as few as 1 member only. A careful review of [16] should explain this grouping behavior because, as stated, about 4475 genes are detected in one sample. Another clustering algorithm like Kohonen Self-Organizing Maps (SOM) is suggested because the group centers adapt and learn with every member they acquire which gives more chance for a more balanced distribution [19].



## Connectionist Approaches for Predicting Mouse Gene Function from Gene Expression




Fig. 2. Genes distribution (500 groups)




Fig. 3. Genes distribution (300 groups)

## 7 Results Comparison

The closest to compare against this work is Mateos et al [18] and Zhang et al [16]. Although Mateos et al work is for predicting yeast gene functions and among only 100 predefined classes, the prediction methods are similar. Taking the dataset size into consideration, results in our work present a more promising performance of NN. The authors have a consistently poor performance (<60%) in terms of the number of False Positive (FPs). They defined three biological factors, in addition to the high noise to data ratio of their dataset, that cause this poor performance: Class size, Heterogeneity and the high degree of intersection among functional classes. Their conclusion is that when it comes to NN, FPs and False Negative (FNs) computationally do not necessarily mean the same biologically [18].

A quick comparison with Zhang et al work in [16] shows that NN outperforms SVM in this problem. However, a detailed comparison is not easy to perform for many reasons: First of all, NN is used in three different ways in this work while only binary classification is performed in [16]. Also, the authors presented prediction results for both annotated and unannotated genes while we used only the annotated genes for this work. Finally, Zhang et al did not use the SVM discriminant value to assess the SVM performance, they rather processed it to obtain an estimate of the probability of correct prediction of each gene in each category (*the recall value*).

## 8 Concluding remarks and future work

In this work we showed that NN can effectively predict mouse gene function from gene expression levels. The learnable correlation between gene expression levels and their function categories in eukaryotes have been confirmed by our results. We presented three different ways to use NN to solve this problem. NN solved the binary classification version of the problem effectively. The actual results for a balanced, and hence harder, data set have been presented. Using MLP to predict all the available categories for a certain gene showed very promising results; however, a huge NN size and a lot of training time are needed. Finally, K-means was used to group similar gene annotations as a preprocessing step for the NN. We showed that K-means is not the best candidate for this problem because it has fixed means that got readjusted only at

**Emad Andrews Shenouda**, Quaid Morris, Anthony J. Bonner

the end of the training process. Ways to improve obtained results include using another learning clustering algorithm, like Kohonen SOM, and optimizing the MLP parameters using evolutionary algorithms.

## References

1. Vinayagam, A., König, R., Moormann, J., Schubert, F., Eils, R., Glatting, K., Suhai, S.: Applying Support Vector Machines for Gene Ontology Based Gene Function Prediction BMC Bioinformatics 5(116) (2004)
2. Baldi, P., Brunak, S.: Bioinformatics a Machine Learning Approach. 2<sup>nd</sup> edn. MIT Press, Cambridge, London (2001)
3. Beer, M.A., Tavazoie, S.: Predicting Gene Expression from Sequence. Cell 16 117(2) Apr. (2004)185-98
4. Michie, D., Spiegelhalter, D.J., Taylor, C.C.: Machine Learning, Neural and Statistical Classification. Elis Horwood, London (1994)
5. Morrison, K.E., Daniels, R. J., Campbell, L., McPherson, J., Davies, K.E.: (1993). Human Molecular Genetics 2. (1993).
6. Haykin, S.: Neural Networks A Comprehensive Introduction. Prentice Hall, New Jersey (1999)
7. Le Cun, Y., Touresky, D., Hinton G., Sejnowski, T.: A Theoretical Framework for Backpropagation. In: The Connectionist Models Summer School. (1988) 21-28
8. Brown, M. P. S., Grundy, W. N., Lin, D., Cristianini, N., Sugnet, C., Furey, T. S., Ares, M., Jr., Haussler, D.: Knowledge-Based Analysis of Microarray Gene Expression Data Using Support Vector Machines. In: Proceedings of the National Academy of Sciences 97(1) (1997) 262-267
9. Mitchell, T. M.: Machine Learning. McGraw-Hill Inc., USA (1997).
10. NCBI COGs database [<http://ncbi.nlm.nih.gov/COG/>]
11. King, O.D., Foulger, R.E., Dwight, S.S., White, J.V., Roth F.P.: Predicting Gene Function from Patterns of Annotation. Genome Research 13 (5) (2003) 896-904
12. Shenouda, E.: A Quantitative Comparison of Different MLP Activation Functions in Classification. In: Proceedings of the International Symposium of Neural Networks ISNN06 (2006)
13. The Functional Landscape of Mouse Gene Expression. [<http://hugheslab.med.utoronto.ca/Zhang>]
14. Wang, D., Huang, G.: Protein Sequence Classification Using Extreme Learning Machine. In: IJCNN05, Vol. 3. Montréal (2005) 1406-1411
15. Werbos, P. J.: Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences. Doctoral Thesis, Applied Mathematics, Harvard University. Boston (1974)
16. Zhang W., Morris Q.D., Chang, R., Shai, O., Bakowski, M.A., Mitsakakis, N., Mohammad, N., Robinson, M.D., Zirngibl, R., Somogyi, E., Laurin, N., Eftekharpour, E., Sat, E., Grigull, J., Pan, Q., Peng, W.T., Krogan, N., Greenblatt, J., Fehlings, M., Kooy D. V., Aubin J., Bruneau B.G., Rossant, J., Blencowe, B.J., Frey, B.J., Hughes T.R.: The Functional Landscape of Mouse Gene Expression Journal of Biology 3 Article 21 (2004)
17. Zurada, J. M.: Introduction to Artificial Neural Systems. PWS Publishing, Boston (1999)
18. Mateos, A., Dopazo, J., Jansen, R., Tu, Y., Gerstein, M., Stolovitzky, G.: Systematic Learning of Gene Functional Classes From DNA Array Expression Data by Using Multilayer Perceptrons Genome Research 12 (2002) 1703-1715
19. Kohonen, T.: Self-Organizing and Associative Memory. 3<sup>rd</sup> edn Springer-Verlag, New York (1998)