

Relational Descriptive Analysis of Gene Expression Data

Igor Trajkovski ^{a,1}, Filip Zelezny ^b Nada Lavrac ^a Jakub Tolar ^c

^a *Department of Knowledge Technologies, Jozef Stefan Institute,
Ljubljana, Slovenia*

^b *Department of Cybernetics, Czech Technical University in Prague,
Prague, Czech Republic*

^c *Department of Pediatrics, University of Minnesota Medical School,
Minneapolis, USA*

Abstract. This paper presents a method that uses gene ontologies, together with the paradigm of relational subgroup discovery, to help find description of groups of genes differentially expressed in specific cancers. The descriptions are represented by means of relational features, extracted from publicly available gene ontology information, and are straightforwardly interpretable by medical/biology researchers. We applied the proposed method to two known data sets: (i) acute lymphoblastic leukemia (ALL) vs. acute myeloid leukemia (AML) and (ii) classification of fourteen types of cancer. Significant number of discovered groups of genes had right description, confirmed by medical expert, which highlighted the underlying biological process that is responsible for distinguishing one class from the other classes. We view our methodology not just as a prototypical example of applying more sophisticated machine learning algorithms to gene expression analysis, but also as an motivation for developing increasingly more sophisticated functional annotations and ontologies, that can be processed by such learning algorithms.

Keywords. Relational learning, Learning from structured data, Learning in bioinformatics, Scientific discovery, Inductive logic programming, Meta-learning

1. Introduction

Microarrays are at the center of a revolution in biotechnology, allowing researchers to simultaneously monitor the expression of tens of thousands of genes. Independent of the platform and the analysis methods used, the result of a microarray experiment is, in most cases, a list of genes found to be differentially expressed. The common challenge faced by the researchers is to translate such gene lists into a better understanding of the underlying biological phenomena. Manual or

¹Correspondence to: Igor Trajkovski, Department of Knowledge Technologies, Jozef Stefan Institute, Jamova 39, 1000 Ljubljana, Slovenia. Tel.: +386 1 477 3125; Fax: +386 1 477 3315; E-mail: igor.trajkovski@ijs.si

semi-automated analysis of large-scale biological data sets typically requires biological experts with vast knowledge of many genes, to decipher the known biology accounting for genes with correlated experimental patterns. The goal is to identify the relevant "functions", or the global cellular activities, at work in the experiment. For example, experts routinely scan gene expression clusters to see if any of the clusters are explained by a known biological function. Efficient interpretation of these data is challenging because the number and diversity of genes exceed the ability of any single researcher to track the complex relationships hidden in the data sets. However, much of the information relevant to the data is contained in the publicly available gene ontologies. Including the ontologies as a direct knowledge source for any algorithmic strategy to approach such data may greatly facilitate the analysis.

Here we present a method to identify groups of genes with a similar signature in gene expression data that also have functional similarity in the background knowledge formally represented with gene annotation terms from the gene ontology. Precisely, we present an algorithm that for given multi-dimensional numerical data set, representing the expression of the genes under different conditions (that define the classes of examples) and ontology used for producing background knowledge about these genes, is able to identify groups of genes, described by conjunctions of first order features, whose expression is highly correlated with one of the classes. For example, one of the applications of this algorithm is to describe groups of genes that were selected as discriminative factor for some classification problem. Medical experts are usually not satisfied with separate description of every discriminative gene, but want to know the processes that are controlled by these genes. With our algorithm we are able to find these processes and cellular component where they are "executed" and which genes from preselected list of discriminative genes are included in these processes.

For doing this we use the methodology of Relational Subgroup Discovery (RSD) [10]. *With RSD we were able to induce set of discrimination rules between the different types (or subtypes) of cancers in terms of functional knowledge extracted from the gene ontology and information about gene interactions.* In other words, we try to explain the differences between types of cancer in terms of the functions of the genes that are differentially expressed in these types.

1.1. Measuring gene expression

The process of transcribing a gene's DNA sequence into the RNA that serves as template for protein production is known as *gene expression*. A gene's expression level indicates the approximate number of copies of that gene's RNA produced in a cell. This is thought to correlate with the amount of corresponding protein made. While the traditional technique for measuring gene expression is labor-intensive and produces approximate quantitative measure of expression, the new technologies have greatly improved the resolution and the scalability of gene expression monitoring. "Expression chips", manufactured using technologies derived from computer-chip production, can now measure the expression of thousands of genes simultaneously, under different conditions. These conditions may be different time points during a biological process, such as the yeast cell cycle

or drosophila development; direct genetic manipulations on a population of cells such as gene deletions; or they can be different tissue samples with some common phenotype (such as different cancer specimen). A typical gene expression data set is a matrix, with each column representing a gene and each row representing a condition, e.g. a cancer type. The value at each position in the matrix represents expression of a gene under some condition.

1.2. Analysis of gene expression data

Large scale gene expression data sets include thousands of genes measured at dozens of conditions. The number and diversity of genes make manual analysis difficult and automatic analysis methods necessary. Initial efforts to analyze these data sets began with the application of unsupervised machine learning, or clustering, to group genes according to similarity in gene expression [4]. Clustering provides a tool to reduce the size of the dataset to a simpler one that can more easily be manually examined. In typical studies, researchers examine the clusters to find those containing genes with common biological properties, such as the presence of common upstream promoter regions or involvement in the same biological processes. After commonalities have been identified (often manually) it becomes possible to understand the global aspects of the biological phenomena studied. As the community developed an interest in this area, additional novel clustering methods were introduced and evaluated for gene expression data [1,6].

The analysis of microarray gene expression data for various tissue samples has enabled researchers to determine gene expression profiles characteristic of the disease subtypes. The groups of genes involved in these genetic profiles are rather large and a deeper understanding of the functional distinction between the disease subtypes might help not only to select highly accurate 'genetic signatures' of the various subtypes, but hopefully also to select potential targets for drug design. Most current approaches to microarray data analysis use (supervised or unsupervised) clustering algorithms to deal with the numerical expression data. While a clustering method reduces the dimensionality of the data to a size that a scientist can tackle, it does not identify the critical background biological information that helps the researcher understand the significance of each cluster. However, that biological knowledge in terms of functional annotation of the genes is already available in public databases. Direct inclusion of this knowledge source can greatly improve the analysis, support (in term of user confidence) and explain obtained numerical results.

1.3. Gene Ontologies

One of the most important tools for the representation and processing of information about gene products and functions is the Gene Ontology (GO). GO is being developed in parallel with work on a variety of other biological databases within the umbrella project OBO (Open Biological Ontologies). It provides a controlled vocabulary for the description of cellular components, molecular functions, and biological processes. As of January 2006 (www.geneontology.org), GO contains 1681 component, 7386 function and 10392 process terms. Terms are organized

in parent-child hierarchies, indicating either that one term is more specific than another (is_a) or that the entity denoted by one term is part of the entity denoted by another (part_of). Typically, such associations (or 'annotations') are first of all established electronically and later validated by a process of manual verification which requires the annotator to have expertise both in the biology of the genes and gene products and in the structure and content of GO. The Gene Ontology, in spite of its name, is not an ontology in the sense accepted by computer scientists, in that it does not deal with axioms and definitions associated to terms. It is rather a taxonomy, or, as the GO Consortium puts it, a 'controlled vocabulary' providing a practically useful framework for keeping track of the biological annotations applied to gene products.

Recently, an automatic ontological analysis approach using GO has been proposed to help solving the task of interpreting the results of gene expression data analysis [8]. From 2003 to 2005, 13 other tools have been proposed for this type of analysis and more tools continue to appear every day. Although these tools use the same general approach, identifying statistically significant GO terms that cover a selected list of genes, they differ greatly in many respects that influence in an essential way the results of the analysis. A general idea and comparison of those tools is presented in [9]. Another approach to descriptive analysis of gene expression data is [14]. They present a method that uses text analysis to help find meaningful gene expression patterns that correlate with the underlying biology described in scientific literature.

2. Descriptive analysis of gene expression data

The fundamental idea of this paper is as follows. First, we construct a set of discriminative genes, $G_C(c)$, for every class $c \in C$. These sets can be constructed in several ways. For example: $G_C(c)$ can be the set of the k ($k > 0$) most correlated genes with class c , computed by, for example, Pearson's correlation. $G_C(c)$ can also be the set of the best k single gene predictors, using the recall values from microarray experiment (absent/present/marginal) as the expression value of the gene. These predictors can look like this: **If** $gene_i = \text{present}$ **Then** class = c . In our experiments we used a measure of correlation, $P(g, c)$, that emphasizes the "signal-to-noise" ratio in using the gene g as predictor for class c . Definition and analysis of $P(g, c)$ is presented in later section.

The second step aims at improving the interpretability of G_C . Informally, we do this by identifying groups of genes in $G_C(c)$ (for each $c \in C$) which can be summarized in a compact way. Put differently, for each $c_i \in C$ we search for compact descriptions of group of genes which correlate strongly with c_i and weakly with all $c_j \in C; j \neq i$.

Searching of these groups of genes, together with their description, is defined as separate supervised machine learning task. We refer to it as the secondary, or meta-mining task, as it aims to mine from the outputs of the primary learning process where genes of predictive strength are searched. This secondary task is, in a way, orthogonal to the primary discovery process in that the original attributes (genes) now become training examples, each of which has a class label $c \in C$.

To apply a discovery algorithm, information about relevant features of the new examples is required. No such features (i.e., 'attributes' of the original attributes - genes) are usually present in the gene expression microarray data sets themselves. However, this information can be extracted from a public database of gene annotations (in this paper, we use the Entrez Gene database maintained at the US National Center for Biotechnology Information, <ftp://ftp.ncbi.nlm.nih.gov/gene/>). For each gene we extracted its molecular functions, biological processes and cellular components where its protein products are located. Next, using GO, in the gene's background knowledge we also included their generalized annotations. For example, if one gene is functionally annotated as: zinc ion binding, then it will also be annotated as: transition metal ion binding, metal ion binding, cation binding, ion binding and binding. In the gene's background knowledge we also included information about the interactions of the genes, in the form of pairs of genes for which there is a evidence that they can interact.

In traditional machine learning, examples are expected to be described by a tuple of values corresponding to some predefined, fixed set of attributes. Note that a gene annotation does not straightforwardly correspond to a fixed attribute set, as it has an inherently relational character. For example, a gene may be related to a variable number of cell processes, play role in variable numbers of regulatory pathways etc. This imposes 1-to-many relations hard to elegantly capture within an attribute set of fixed size. Furthermore, a useful piece of information about a gene g may for instance be expressed by the feature

*g interacts with another gene whose functions
include protein binding.*

Going even further, the feature may not include only a single interaction relation but rather consider entire chains of interactions. The difficulties of representing such features through attribute-value tuples are evident.

In summary, we are approaching the task of subgroup discovery from a relational data domain. For this purpose we employ the methodology of relational subgroup discovery proposed in [10,16] and implemented in the RSD algorithm ¹. Using RSD, we were able to discover knowledge such as

The expression of genes coding for proteins located in the integral-to-membrane cell component, whose functions include receptor activity, has a high correlation with the BCR class of acute lymphoblastic leukemia (ALL) and a low correlation with the other classes of ALL.

The RSD algorithm proceeds in two steps. First, it constructs a set of relational features in the form of first order logic atom conjunctions. The entire set of features is then viewed as an attribute set, where an attribute has the value true for a gene (example) if the gene has the feature corresponding to the attribute. As a result, by means of relational feature construction we achieve the conversion of relational data into attribute-value descriptions. In the second step, groups of genes are searched, such that each group is represented as a conjunction of

¹<http://labe.felk.cvut.cz/zelezny/rsd/rsd.pdf>

selected features. The subgroup discovery algorithm employed in this second step is an adaptation of the popular propositional rule learning algorithm CN2 [3].

2.1. Relational feature construction

The feature construction component of RSD aims at generating a set of relational features in the form of relational logic atom conjunctions. For example, the feature exemplified informally in the previous section has the relational logic form

$$\text{interaction}(g,G), \text{function}(G,\text{protein_binding})$$

Here, upper cases denote existentially quantified variables and g is the key term that binds a feature to a specific example (here a gene). The user specifies a grammar declaration which constraints the resulting set of constructed features. RSD accepts feature language declarations similar to those used in the inductive logic programming system Progol [12].

The construction of features is implemented as depth first, general-to-specific search where refinement corresponds to adding a literal to the currently examined expression. During the search, each search node found to be a correct feature is listed in the output.

A remark is needed concerning the way constants (such as *protein_binding*) are employed in features. Rather than making the user responsible for declaring all possible constants that may occur in features, RSD extracts them automatically from the training data. The user marks the types of variables which should be replaced by constants. For each constant-free feature, a number of different features are then generated, each corresponding to a possible replacement of the combination of the indicated variables with constants. RSD then only proceeds with those combinations of constants which make the feature true for at least a pre-specified number of examples. Finally, to evaluate the truth value of each feature for each example for generating the attribute-value representation of the relational data, the first-order logic resolution procedure is used, provided by a Prolog language engine.

2.2. Subgroup Discovery

A subgroup discovery task is defined as follows: *Given a population of individuals and a property of individuals we are interested in, find population subgroups that are statistically 'most interesting', e.g., are as large as possible and have the most unusual statistical (distributional) characteristics with respect to the property of interest.* [15].

Notice an important aspect of the above definition: there is a predefined property of interest, meaning that a subgroup discovery task aims at characterizing population subgroups of a given target class. This property indicates that standard classification rule learning algorithms could be used for solving the task. However, while the goal of classification rule learning is to generate models (sets of rules), inducing class descriptions in terms of properties occurring in the descriptions of training examples, in contrast, subgroup discovery aims at discovering individual patterns of interest (individual rules describing the target class).




Figure 1. Description of discovered subgroups can cover: a) Individuals from only one class (1), b) Some not classified individuals (2), c) Individuals from other classes (3) and d) Individuals already covered by other subgroups (3,4)

Rule learning typically involves two main procedures: the search procedure that performs search to find a single rule, and the control procedure (the covering algorithm) that repeatedly executes the search in order to induce a set of rules. The two procedures are described in [10]

3. Experiments

This section presents a statistical validation of the proposed methodology. We do not assess here the accuracy of disease classification from gene-expression values itself, as this is a property of the particular method used for the primary mining task, which is not our main concern. We rather aim at evaluating the properties of the secondary, descriptive learning task. Namely, we wish to determine if the high descriptive capacity pertaining to the incorporation of the expressive relational logic language incurs a risk of *descriptive overfitting*, i.e., a risk of discovering fluke subgroups. We thus aim at measuring the discrepancy of the quality of discovered subgroups on the training data set on one hand and an independent test set on the other hand. We will do this through the standard 10-fold stratified cross-validation regime.

The specific qualities measured for each set of subgroups produced for a given class are average *precision* (PRE) and *recall* (REC) values among all subgroups in the subgroup set. See eg. [10] for details of their calculation.

3.1. Materials and methods

We apply the proposed methodology on two problems of predictive classification from gene expression data.

The first was introduced in [5] and aims at distinguishing between samples of acute lymphoblastic leukemia and acute myeloid leukemia from gene expression profiles obtained by the Affymetrix HU6800 microarray chip, containing probes for 6817 genes. The data contains 73 class-labeled samples of expression vectors. The second was defined in [13]. Here one tries to distinguish among 14 classes of cancers from gene expression profiles obtained by the Affymetrix Hu6800 and Hu35KsubA microarray chip, containing probes for 16,063 genes. The data set contains 198 class-labeled samples.

To access the annotation data for every gene considered, it was necessary to obtain unique gene identifiers from the microarray probe identifiers available in the original data. We achieved this by querying Affymetrix site² for translating probe ID's into unique gene ID's. Knowing the gene identifiers, information about gene annotations and gene interactions can be extracted from Entrez gene information database³. We developed a program script in the Python language, which extracts gene annotations and gene interactions from this database, and produces their structured, relational logic representations which can be used as input to RSD. This script is available on request to the first author.

In both data sets, for each class c we first extracted a set of discriminative genes $G_C(c)$. In our experiments we used a measure of correlation, $P(g, c)$, that emphasizes the "signal-to-noise" ratio in using the gene g as predictor for class c . $P(g, c)$ is computed by the following procedure:

Let $[\mu_1(g), \sigma_1(g)]$ and $[\mu_2(g), \sigma_2(g)]$ denote the means and standard deviations of log of the expression levels of gene g for the samples in class c and samples in all other classes, respectively.

Let $P(g, c) = \frac{\mu_1(g) - \mu_2(g)}{\sigma_1(g) - \sigma_2(g)}$, which reflects the difference between the classes relative to the standard deviation within the classes. Large values of $|P(g, c)|$ indicate a strong correlation between the gene expression and the class distinction, while the sign of $P(g, c)$ being positive or negative correspond to g being more highly expressed in class c or in other classes. Unlike a standard Pearson correlation coefficient, $P(g, c)$ is not confined to the range $[-1, +1]$. The set of informative genes for class c , $G_C(c)$ of size n , consist of the n genes having the highest $|P(g, c)|$ value. If we have only two classes, then $G_C(c_1)$ consist of genes having the highest $P(g, c_1)$ values, and $G_C(c_2)$ consist of genes having the highest $P(g, c_2)$.

For the first problem we selected 50 discriminative genes for ALL and 50 for AML class. In the second problem we selected 35 discriminative genes for each class. The average value of correlation coefficient, $|P(g, c)|$, of selected discriminatory genes for each class/problem are listed in Table 1. The usage of the gene correlation coefficient is twofold. In the first part of the analysis, for a given class it is used for selection of discriminative genes, and in the second part as initial weight of the example-genes for the meta-mining procedure where we try to describe these discriminative genes. In the second mining task RSD will prefer to group genes with large weights, so these genes will have enough weight to be grouped in several groups with different descriptions.

After the selection of sets of discriminatory genes, $G_C(c)$ for each $c \in C$, these sets were merged and every gene coming from $G_C(c)$ was class labeled as c . Now RSD was run on these data, with aim to find as big as possible and as clear (in terms of class labels) as possible subgroups of this population of example-genes, described by relational features constructed from GO and gene interaction data.

3.2. Results

The discovered regularities have a very interesting biological interpretation.

²www.affymetrix.com/analysis/netaffx/

³[ftp://ftp.ncbi.nlm.nih.gov/gene/](http://ftp.ncbi.nlm.nih.gov/gene/)

Table 1. Average(AVG), maximal(MAX) and minimal(MIN) value of $\{|P(g,c)|g \in G_C(c)\}$ for each task and class c .

TASK	CLASS	AVG	MAX	MIN
ALL-AML	ALL	0.75	1.25	0.61
	AML	0.76	1.44	0.59
MULTI CLASS	BREAST	1.06	1.30	0.98
	PROSTATE	0.91	1.23	0.80
	LUNG	0.70	0.99	0.60
	COLORECTAL	0.98	1.87	0.73
	LYMPHOMA	1.14	2.52	0.87
	BLADDER	0.88	1.15	0.81
	MELANOMA	1.00	2.60	2.60
	UTERUS	0.82	1.32	0.71
	LEUKEMIA	1.35	1.75	1.18
	RENAL	0.81	1.20	0.70
	PANCREAS	0.75	1.08	0.67
	OVARY	0.70	1.04	0.60
	MESOTHELIOMA	0.90	1.91	0.74
	CNS	1.38	2.17	1.21

In breast cancer, RSD has identified a group of genes (described as process(G,'regulation of transcription'), function(G,'zinc ion binding')) containing five genes (Entrez Gene id's: 4297, 51592, 91612, 92379, 115426) whose under expression is a good predictor for that class. These genes are simultaneously involved in regulation of transcription and in zinc ion binding. Zinc is a cofactor in protein-DNA binding, via a "zinc finger" domain (id 92379). This property is shared by many transcription factors, which are major regulators of normal and abnormal (e.g., malignant) cell proliferation. Second, zinc is an essential growth factor and a zinc transporter associated with metastatic potential of estrogen positive breast cancer, termed LIV-1, has been described [7]. Less than optimal expression of the factors involved in zinc metabolism can therefore represent either a cause or effect (biomarker) of dysregulated cellular proliferation in breast cancer. A separate group of genes involved in ubiquitin cycle (process(G,'ubiquitin cycle')) was identified in breast cancer, (Entrez id's: 3093, 10910, 23014, 23032, 25831, 51592, 115426). The role of ubiquitin in a cell is to recycle proteins. This is of a paramount importance to the overall cellular homeostasis, since inappropriately active proteins can cause cancer. Subnormal expression of ubiquitin components (id's 3093 and 23032) resulting in subnormal inactivation of proteins active in cell cycle could thus represent a functional equivalent of ectopic oncogene expression [11]. This is the example where one gene, id: 115426, was included in two groups with different descriptions.

In CNS (central nervous system) cancer, we discovered two important groups concerning neurodevelopment (description: process(G,'nervous system development'), Entrez Gene id's: 333, 1400, 2173, 2596, 2824, 3785, 4440, 6664, 7545, 10439, 50861), and immune surveillance (Entrez Gene id's: 199, 1675, 3001, 3108,

Table 2. Precision-recall figures and average size of found subgroups, for the ALL/AML and multi-class classification task obtained through 10-fold cross-validation.

TASK	DATA	PRE(st.dev.)	REC(st.dev.)	AVG. SIZE
ALL-AML	Train	0.96(0.01)	0.18(0.02)	12.07
	Test	0.71(0.12)	0.12(0.04)	
MULTI-CLASS	Train	0.51(0.03)	0.15(0.01)	8.35
	Test	0.42(0.10)	0.10(0.02)	

3507, 3543, 3561, 3588, 3683, 4046, 5698, 5699, 5721, 6352, 9111, 28299, 50848, 59307). The genes in the first/second group are over/under-expressed (respectively) in CNS. As for the former, reactivation of genes relevant to early development (i. e., ineffective recapitulation of embryonal or fetal neural growth at a wrong time) is a hallmark of the most rapidly growing tumors (id's 3785 and 10439 specific to neuroblastoma). The latter illustrates the common clinical observation that immune deficiency (subnormal expression of genes active in immune response shown in this work) creates a permissive environment for cancer persistence. Thus, both major themes of malignant growth are represented in this example: active unregulated growth and passive inability to clear the abnormal cells.

In ALL, RSD has identified a group of 23 genes, described as: component(G, 'nucleus') AND interaction(G,B),process(B,'regulation of transcription, DNA-dependent'). The products of these genes, proteins, are located in the nucleus of the cell, and they interact with genes that are included in the process of regulation of transcription. In AML, RSD has identified several groups of overexpressed genes, located in the membrane, that interact with genes that have 'metal ion transport' as one of their function.

In addition, we subjected the RSD algorithm to a 10-fold stratified cross-validation on both classification tasks. Table 2 show the PRE and REC values (with standard deviation figures) results for the two respective classification tasks. Overall, the results show only a small drop from the training to the testing set in terms of both PRE and REC, suggesting that the number of discriminant genes selected (Table 1) was sufficient to prevent overfitting. In terms of *total* coverage, RSD covered more then $\frac{2}{3}$ of the preselected discriminative genes (in both problems), while $\frac{1}{3}$ of the preselected gene were not included in any group. One interpretation of that is that they are not functionally connected with the other genes, but were selected by chance. This information can be used in the first phase of the classification problem, feature selection, by choosing genes that were covered by some subgroup. That will be the next step in our future work, using the proposed methodology as feature (gene) selection mechanism.

4. Discussion

In this paper we presented a method that uses gene ontologies, together with the paradigm of relational subgroup discovery, to help find patterns of expression

for genes with a common biological function that correlate with the underlying biology responsible for class differentiation. Our methodology proposes to first select set important discriminative genes for all classes and then finding compact, relational descriptions of subgroups among these genes.

It is noteworthy that the 'post-processing' step is also a machine learning task, in which the curse of dimensionality (the number of attributes - gene expressions measured) usually ascribed to the type of classification problem considered, actually turns into an advantage. The high number of attributes (important genes), incurring the risk of overfitting, turns into a high number of examples, which on the contrary works against overfitting in the subsequent subgroup discovery task. Furthermore, the dimensionality of the secondary attributes (relational features of genes extracted from gene annotations) can be conveniently controlled via suitable constraints of the language grammar used for the automatic construction of the gene features.

Furthermore, since genes frequently have multiple functions that they may be involved in, they may under some of the conditions exhibit the behavior of genes with one function and in other conditions exhibit the behavior of genes with a different function. Here subgroup discovery is effective at selecting a specific function. The same gene can be selected by multiple subgroups (gene id: 115426 in breast cancer), each emphasizing the different biological process critical to the explanation of the underlying biology responsible for observed experimental results. Unlike other tools for analyzing gene expression data that use gene ontologies, which report statistically significant single GO terms and don't use gene interaction data, we are able to find set of GO terms (the first reported group of genes, for breast cancer, is described with two GO terms), that cover the same set of genes, and we use available gene interaction data to describe features of genes that can not be represented with other approaches (the third reported group, for ALL).

However, this approach of translating a list of differentially expressed genes into subgroups of functional categories using annotation databases suffers from a few important limitations. The existing annotations databases are incomplete, only a subset of known genes is functionally annotated and most annotation databases are built by curators who manually review the existing literature. Although unlikely, it is possible that certain known facts might get temporarily overlooked. For instance, [9] found references in literature published in the early 90s, for 65 functional annotations that are yet not included in the current functional annotation databases. In the GO, associations more than half of the terms are inferred exclusively from electronic annotations (i.e. without any expert human involvement). The vast majority of such electronic annotations are reasonably accurate [2]. However, many such annotations are often made at very high-level GO terms, which limit their usefulness.

Despite the current imperfectness of the available ontological background knowledge, the presented methodology was able to discover and compactly describe several gene groups, associated to specific cancer types, with highly plausible biological interpretation. We thus strongly believe the presented approach will significantly contribute to the application of relational machine learning to gene

expression analysis, given the expected increase in both the quality and quantity of gene/protein annotations in the near future.

References

- [1] Ben-Dor, A., Shamir, R., & Yakhini, Z. (1999). Clustering gene expression patterns. *J. Comput. Biol.*, 6:3/4, 281-297.
- [2] Camon, E.B. et al. (2005). An evaluation of GO annotation retrieval for BioCre-AtIvE and GOA. *BMC Bioinformatics*, 6, S17.
- [3] Clark, P. & Niblett T. (1989). The CN2 induction algorithm. *Machine Learning*, pages 261-283.
- [4] Eisen, M. B., Spellman, P. T., Brown, P. O., & Botstein, D. (1998). Cluster analysis and display of genome-wide expression patterns. *Proc. Natl. Acad. Sci. USA*, 95:25, 14863-14868.
- [5] Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M. et al. (1999). Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring. *Science*, 286:5439, 531-537.
- [6] Heyer, L. J., Kruglyak, S., & Yooseph, S. (1999). Exploring expression data: Identification and analysis of coexpressed genes. *Genome Res.*, 9:11, 1106-1115.
- [7] Kasper, G. et al. (2005). Expression levels of the putative zinc transporter LIV-1 are associated with a better outcome of breast cancer patients. *Int J Cancer*. 20;117(6):961-73.
- [8] Khatri, P. et al. (2002). Profiling gene expression using Onto-Express. *Genomics*, 79, 266-270.
- [9] Khatri, P. & Draghici S. (2005). Ontological analysis of gene expression data: current tools, limitations, and open problems. *Bioinformatics*, 21(18):3587-3595
- [10] Lavrac, N., Zelezny F. & Flach P. (2002). RSD: Relational subgroup discovery through first-order feature construction. In *Proceedings of the 12th International Conference on Inductive Logic Programming*, pages 149-165.
- [11] Mani, A. & Gelmann E.P. (2005). The ubiquitin-proteasome pathway and its role in cancer. *Journal of Clinical Oncology*. 23:4776-4789
- [12] Muggleton, S. (1995). Inverse entailment and Prolog. *New Generation Computing*, Special issue on Inductive Logic Programming, 13(3-4):245-286.
- [13] Ramaswamy, S., Tamayo P., Rifkin R. et al. (2001). Multiclass cancer diagnosis using tumor gene expression signatures, *PNAS* 98 (26) 15149-15154.
- [14] Raychaudhuri, S., Shtze, H.S. & Altman, R.B. (2003). Inclusion of textual documentation in the analysis of multidimensional data sets: application to gene expression data. *Machine Learn.*, 52, 119-145
- [15] Wrobel, S. (1997). An algorithm for multi-relational discovery of subgroups. In Jan Komorowski and Jan Zytkow, editors, *Proceedings of the First European Symposium on Principles of Data Mining and Knowledge Discovery (PKDD-97)*, pages 78-87.
- [16] Zelezny, F. & N. Lavrac. (2006). Propositionalization based relational subgroup discovery with RSD. *Machine Learning*. <http://labe.felk.cvut.cz/~zelezny/pubs/mlj06.pdf>.