

Bisociative Knowledge Discovery for Microarray Data Analysis

Igor Mozetič¹, Nada Lavrač^{1,2}, Vid Podpečan¹, Petra Kralj Novak¹,
Helena Motaln³, Marko Petek³, Kristina Gruden³,
Hannu Toivonen⁴, Kimmo Kulovesi⁴

¹ Jožef Stefan Institute, Jamova 39, Ljubljana, Slovenia
{igor.mozetic, nada.lavrac, vid.podpecan, petra.kralj}@ijs.si

² University of Nova Gorica, Vipavska 13, Nova Gorica, Slovenia

³ National Institute of Biology, Večna pot 111, Ljubljana, Slovenia
{helena.motaln, marko.petek, kristina.gruden}@nib.si

⁴ Department of Computer Science, University of Helsinki, Finland
{hannu.toivonen, kimmo.kulovesi}@cs.helsinki.fi

Abstract. The paper presents an approach to computational knowledge discovery through the mechanism of *bisociation*. Bisociative reasoning is at the heart of creative, accidental discovery (e.g., serendipity), and is focused on finding unexpected links by crossing contexts. Contextualization and linking between highly diverse and distributed data and knowledge sources is therefore crucial for the implementation of bisociative reasoning. In the paper we explore these ideas on the problem of analysis of microarray data. We show how enriched gene sets are found by using ontology information as background knowledge in semantic subgroup discovery. These genes are then contextualized by the computation of probabilistic links to diverse bioinformatics resources. Preliminary experiments with microarray data illustrate the approach.

1 Introduction

Systems biology studies and models complex interactions in biological systems with the goal of understanding the underlying mechanisms. Biologists collect large quantities of data from wet lab experiments and high-throughput platforms. Public biological databases, like Gene Ontology and Kyoto Encyclopedia of Genes and Genomes, are sources of biological knowledge. Since the growing amounts of available knowledge and data exceed human analytical capabilities, technologies that help analyzing and extracting useful information from such large amounts of data need to be developed and used.

The concept of association is at the heart of many of today's ICT technologies such as information retrieval and data mining (for example, association rule learning is an established data mining technology, [1]). However, scientific discovery requires creative thinking to connect seemingly unrelated information, for example, by using metaphors or analogies between concepts from different domains. These modes of thinking allow the mixing of conceptual categories and

contexts, which are normally separated. One of the functional basis for these modes is the idea of *bisociation*, coined by Artur Koestler half a century ago [7]:

“The pattern . . . is the perceiving of a situation or idea, L , in two self-consistent but habitually incompatible frames of reference, M_1 and M_2 . The event L , in which the two intersect, is made to vibrate simultaneously on two different wavelengths, as it were. While this unusual situation lasts, L is not merely linked to one associative context but *bisociated* with two.”

Koestler found bisociation to be the basis for human creativity in seemingly diverse human endeavors, such as humor, science, and arts. The concept of bisociation in science is illustrated in Figure 1.




Fig. 1. Koestler’s schema of bisociative discovery in science ([7], p. 107).

We are interested in creative discoveries in science, and in particular in computational support for knowledge discovery from large and diverse sources of data and knowledge. To this end, we participate in a European FP7 FET-Open project BISON⁵ which investigates possible computational realizations of bisociative reasoning. The project is based on the following, somewhat simplified, assumptions:

- A bisociative information network (named BisoNet) can be created from available resources. BisoNet is a large graph, where nodes are concepts and edges are probabilistic relations. Unlike semantic nets or ontologies, the graph is easy to construct automatically since it carries little semantics. To a large extent it encodes just circumstantial evidence that concepts are somehow related through edges with some probability.

⁵ <http://www.bisonet.eu/>

- Different subgraphs can be assigned to different contexts (frames of reference).
- Graph analysis algorithms can be used to compute links between distant nodes and subgraphs in a BisoNet.
- A bisociative link is a link between nodes (or subgraphs) from different contexts.

In this paper we thus explore one specific pattern of bisociation: long-range links between nodes (or subgraph) which belong to different contexts. More precisely, we say that two concepts are bisociated if:

- there is no direct, obvious evidence linking them,
- one has to cross contexts to find the link, and
- this new link provides some novel insight into the problem domain.

We have to emphasize that context crossing is subjective, since the user has to move from his ‘normal’ context (frame of reference) to an *habitually incompatible context* to find the bisociative link [2]. In terms of Koestler (Figure 1), a habitual frame of reference (plane M_1) corresponds to a BisoNet subgraph as defined by a user or his profile. The rest of the BisoNet represents different, habitually incompatible contexts (in general, there may be several planes M_2). The creative act here is to find links (m_2) which lead ‘out-of-the-plane’ via intermediate, bridging concepts (L). Thus, contextualization and link discovery are two of the fundamental mechanisms in bisociative reasoning as implemented in BISON.

Finding links between seemingly unrelated concepts from texts was already addressed by Swanson [10]. The Swanson’s approach implements *closed discovery*, the so-called A-B-C process, where A and C are given and one searches for intermediate B concepts. On the other hand, in *open discovery* [16], only A is given. One approach to open discovery, RaJoLink [8], is based on the idea to find C via B terms which are rare (and therefore potentially interesting) in conjunction with A. Rarity might therefore be one of the criteria to select links which lead out of the habitual context (around A) to known, but non-obviously related concepts C via B.

In this paper we present an approach to bisociative discovery and contextualization of genes which should help in the analysis of microarray data. The approach is based on semantic subgroup discovery (by using ontologies as background knowledge in microarray data analysis), and the linking of various publicly available bioinformatics databases. This is an ongoing work, where some elements of bisociative reasoning are already implemented: creation of the BisoNet graph, identification of relevant nodes in a BisoNet, and computation of links to indirectly related concepts. Currently, we are expanding the BisoNet with textual resources from PubMed, and implementing open discovery from texts through BisoNet graph mining. We envision that the open discovery process will identify potentially interesting concepts from different contexts which will act as the target nodes for the link discovery algorithms. Links discovered in this way, crossing contexts, might provide instances of bisociative discoveries.

The currently implemented steps of bisociative reasoning are the following. The *semantic subgroup discovery* step is implemented by the SEGS system [14]. SEGS uses as background knowledge data from three publicly available, semantically annotated biological data repositories, GO, KEGG and NCBI. Based on the background knowledge, it automatically formulates biological hypotheses: rules which define groups of differentially expressed genes. Finally, it estimates the relevance (or significance) of the automatically formulated hypotheses on experimental microarray data. The *link discovery* step is implemented by the Biomine system [9]. Biomine weakly integrates a large number of biomedical resources, and computes most probable links between elements of diverse sources. It thus complements the semantic subgroup discovery technology, due to the explanatory potential of additional link discovery and Biomine graph visualization. While this link discovery process is already implemented, our current work is devoted to the contextualization of Biomine nodes for bisociative link discovery.

The paper is structured as follows. Section 2 gives an overview of five steps in exploratory analysis of gene expression data. Section 3 describes an approach to the analysis of microarray data, using semantic subgroup discovery in the context of gene set enrichment. A novel approach, a first attempt at bisociative discovery through contextualization, composed of using SEGS and Biomine (SEGS+Biomine, for short) is in Section 4. An ongoing experimental case study is presented in Section 5. We conclude in Section 6 with plans for future work.

2 Exploratory gene analytics

This section describes the steps which support bisociative discovery, targeted at the analysis of differentially expressed gene sets: gene ranking, the SEGS method for enriched gene set construction, linking of the discovered gene set to related biomedical databases, and finally visualization in Biomine. The schematic overview is in Figure 2.

The proposed method consists of the following five steps:

1. **Ranking of genes.** In the first step, class-labeled microarray data is processed and analyzed, resulting in a list of genes, ranked according to differential expression.
2. **Ontology information fusion.** A unified database, consisting of GO⁶ (biological processes, functions and components), KEGG⁷ (biological pathways) and NCBI⁸ (gene-gene interactions) terms and relationships is constructed by a set of scripts, enabling easy updating of the integrated database (details can be found in [12]).
3. **Discovering groups of differentially expressed genes.** The ranked list of genes is used as input to the SEGS algorithm [14], an upgrade of the

⁶ <http://www.geneontology.org/>

⁷ <http://www.genome.jp/kegg/>

⁸ ftp://ftp.ncbi.nlm.nih.gov/gene/GeneRIF/interaction_sources




Fig. 2. Microarray gene analytics proceeds by first finding candidate enriched gene sets, expressed as intersections of GO, KEGG and NCBI gene-gene interaction sets. Selected enriched genes are then put in the context of different bioinformatic resources, as computed by the Biomine link discovery engine. The '+' and '-' signs under Microarray data indicate over- and under-expression values of genes, respectively.

RSD relational subgroup discovery algorithm [3, 4, 13], specially adapted to microarray data analysis. The result is a list of most relevant gene groups that semantically explain differential gene expression in terms of gene functions, components, processes, and pathways as annotated in biological ontologies.

4. **Finding links between gene group elements.** The elements of the discovered gene groups (GO and KEGG terms or individual genes) are used to formulate queries for the Biomine link discovery engine. Biomine then computes most probable links between these elements and entities from a number of public biological databases. These links help the experts to uncover unexpected relations and biological mechanisms potentially characteristic for the underlying biological system.
5. **Gene group visualization.** Finally, in order to help in explaining the discovered out-of-the-context links, the discovered gene relations are visualized using the Biomine visualization tools.

3 SEGS: Search for Enriched Gene Sets

The goal of the gene set enrichment analysis is to find gene sets which form coherent groups and are distinguished from the rest of the genes. More precisely, a gene set is *enriched* if the member genes are statistically significantly differentially expressed as compared to the rest of the genes. Two methods for testing the enrichment of gene sets were developed: Gene Set Enrichment Analysis (GSEA) [11] and Parametric Analysis of Gene Set Enrichment (PAGE) [6]. Originally, these methods take individual terms from GO and KEGG (which annotate gene sets), and test whether the genes that are annotated by a specific term are statistically significantly differentially expressed in the given microarray dataset.




Fig. 3. Schematic representation of the SEGS method.

The novelty of the SEGS method, developed by Trajkovski et al. [12,14] and used in this study, is that the method does not only test existing gene sets for differential expression but it also generates new gene sets that represent novel biological hypotheses. In short, in addition to testing the enrichment of individual GO and KEGG terms, this method tests the enrichment of newly defined gene sets constructed by the intersection of GO terms, KEGG terms and gene sets defined by taking into account also the gene-gene interaction data from NCBI.

The SEGS method has four main components:

- the background knowledge (the GO, KEGG and NCBI databases),
- the SEGS hypothesis language (the GO, KEGG and interaction terms, and their conjunctions),
- the SEGS hypothesis generation procedure (generated hypotheses in the SEGS language correspond to gene sets), and
- the hypothesis evaluation procedure (the Fisher, GSEA and PAGE tests).

The schematic workflow of the SEGS method is shown in Figure 3.

4 SEGS+Biomine: Contextualization of genes

We made an attempt at exploiting bisociative discoveries within the biomedical domain by explicit contextualization of enriched gene sets. We applied two methods that use publicly available background knowledge for supporting the work of biologists: the SEGS method for searching for enriched gene sets [14] and the Biomine method for contextualization by finding links between genes and other biomedical databases [9]. We combined the two methods in a novel way: we used SEGS for hypothesis generation in the form of interesting gene sets, and then formulated queries for Biomine for out-of-the-context link discovery and visualization (see Figure 4). We believe that by forming hypotheses with SEGS, constructed as intersections of terms from different ontologies (different contexts), discovering links between them by Biomine, and visualizing the SEGS hypotheses and the discovered links by the Biomine graph visualization engine, the interpretation of the biological mechanisms underlying differential gene expression is easier for biologists.




Fig. 4. SEGS+Biomine workflow.

In the Biomine⁹ project [9], data from several publicly available databases were merged into a large graph, a BisoNet, and a method for link discovery between entities in queries was developed. In the Biomine framework nodes correspond to entities and concepts (e.g., genes, proteins, GO terms), and edges represent known, probabilistic relationships between nodes. A link (a relation between two entities) is manifested as a path or a subgraph connecting the corresponding nodes.

Vertex Type	Source Database	Nodes	Degree
Article	PubMed	330,970	6.92
Biological process	GO	10,744	6.76
Cellular component	GO	1,807	16.21
Molecular function	GO	7,922	7.28
Conserved domain	ENTREZ Domains	15,727	99.82
Structural property	ENTREZ Structure	26,425	3.33
Gene Entrez	Gene	395,611	6.09
Gene cluster	UniGene	362,155	2.36
Homology group	HomoloGene	35,478	14.68
OMIM entry	OMIM	15,253	34.35
Protein Entrez	Protein	741,856	5.36
Total		1,968,951	

Table 1. Databases included in the Biomine snapshot used in the experiments.

The Biomine graph data model consists of various biological entities and annotated relations between them. Large, annotated biological data sets can be readily acquired from several public databases and imported into the graph model in a relatively straightforward manner. Some of the databases used in Biomine are summarized in Table 1. The snapshot of Biomine we use consists of a total of 1,968,951 nodes and 7,008,607 edges. This particular collection of data sets is not meant to be complete, but it certainly is sufficiently large and versatile for real link discovery.

⁹ <http://biomine.cs.helsinki.fi/>

5 A case study

In the systems biology domain, our goal is to computationally help the experts to find a creative interpretation of wet lab experiment results. In the particular experiment, the task was to analyze microarray data in order to distinguish between fast and slowly growing cell lines through differential expression of gene sets, responsible for cell growth.

Enriched Gene Sets
1. SLOW-vs-FAST $\leftarrow$ GO.Proc('DNA metabolic process') & INTERACT(GO.Comp('cyclin-dep. protein kinase holoenzyme complex'))
2. SLOW-vs-FAST $\leftarrow$ GO.Proc('DNA replication') & GO.Comp('nucleus') & INTERACT(KEGG.Path('Cell cycle'))
3. SLOW-vs-FAST $\leftarrow$. . .

Table 2. Top SEGS rules found in the cell growth experiment. The second rule states that one possible distinction between the slow and fast growing cells is in genes participating in the process of DNA replication which are located in the cell nucleus and which interact with genes that participate in the cell cycle pathway.

Table 2 gives the top rules resulting from the SEGS search for enriched gene sets. For each rule, there is a corresponding set of over expressed genes from the experimental data. Figure 5 shows a part of the Biomine graph which links a selected subset of enriched gene set to the rest of the nodes in the Biomine graph.

The wet lab scientists have assessed that SEGS in combination with Biomine provide additional hints on what to focus on when comparing the expression data of cells. Additionally, such an in-silico analysis can considerably lower the costs of in-vitro experiments with which the researchers in the wet lab are trying to get a hint of a novel process or phenomena observed. This may be especially true for situations when just knowing the final outcome one cannot explain the drug effect, organ function, or disease satisfactorily. Namely, the gross, yet important characteristics of the cells (organ function) are hidden (do not affect visual morphology) or could not be recognized soon enough. An initial predisposition for this approach is wide accessibility and low costs of high throughput microarray analysis which generate appropriate data for in-silico analysis.

6 Conclusions

We presented SEGS+Biomine, a bisociation discovery system for exploratory gene analytics. It is based on the non-trivial steps of subgroup discovery (SEGS) and link discovery (Biomine). The goal of SEGS+Biomine is to enhance the




Fig. 5. Biomine subgraph related to five genes from the enriched gene set produced by SEGS. Note that the gene and protein names are not explicitly presented, due to the preliminary nature of these results.

creation of novel biological hypotheses about sets of genes. A prototype version of the gene analytics software, which enhances SEGS and creates links to Biomine queries and graphs, is available as a web application at http://zulu.ijis.si/web/segs_ga/.

In the future work we plan to enhance the contextualization of genes with contexts discovered by biomedical literature mining. We will add PubMed articles data into the BisoNet graph structure. To this end, we already have a preliminary implementation of software, called Texas [5], which creates a probabilistic network (BisoNet, compatible to Biomine) from textual sources. By focusing on different types of links between terms (e.g., frequent and rare co-occurrences) we expect to get hints at some unexpected relations between concepts from different contexts.

Our long term goal is to help biologists better understand inter-contextual links between genes and their role in explaining (at least qualitatively) underlying mechanisms which regulate gene expressions. The proposed approach is considered a first step at computational realization of bisociative reasoning for creative knowledge discovery in systems biology.

7 Acknowledgements

The work presented in this paper was supported by the European Commission under the 7th Framework Programme FP7-ICT-2007-C FET-Open project

BISON-211898, by the Slovenian Research Agency grants Knowledge Technologies and Systems Biology J4-2228, and by the Algorithmic Data Analysis (Algodan) Centre of Excellence of the Academy of Finland. We thank Igor Trajkovski for his previous work on SEGS, and Filip Železný and Jakub Tolar for their earlier contributions leading to SEGS.

References

1. R. Agrawal, R. Srikant. Fast algorithms for mining association rules in large databases. *Proc. 20th Intl. Conf. on Very Large Data Bases, VLDB*, pp. 487–499, Santiago, Chile, September 1994.
2. W. Dubitzky. Personal communication, FP7 BISON project review, Leuven, June 2009.
3. D. Gamberger and N. Lavrač. Expert-guided subgroup discovery: Methodology and application. *Journal of Artificial Intelligence Research* 17:501–527, 2002.
4. D. Gamberger, N. Lavrač, F. Železný, J. Tolar. Induction of comprehensible models for gene expression datasets by the subgroup discovery methodology. *Journal of Biomedical Informatics* 37:269–284, 2004.
5. M. Juršič, N. Lavrač, I. Mozetič, V. Podpečan, H. Toivonen. Constructing information networks from text documents. ECML/PKDD 2009 Workshop "Explorative Analytics of Information Networks", Bled, 2009.
6. S.Y. Kim, D.J. Volsky. PAGE: Parametric Analysis of Gene Set Enrichment. *BMC Bioinformatics* 6:144, 2005.
7. A. Koestler. *The Act of Creation*, The Macmillan Co., New York, 1964.
8. I. Petrič, T. Urbančič, B. Cestnik, M. Macedoni-Lukšič. Literature mining method RaJoLink for uncovering relations between biomedical concepts. *Journal of Biomedical Informatics*, 42(2): 219–227, 2009.
9. P. Sevon, L. Eronen, P. Hintsanen, K. Kulovesi, H. Toivonen. Link discovery in graphs derived from biological databases. *Proc. of 3rd International Workshop on Data Integration in the Life Sciences, (DILS'06)*, July 2006. Springer.
10. D.R. Swanson, N.R. Smalheiser, V.I. Torvik. Ranking indirect connections in literature-based discovery: The role of Medical Subject Headings (MeSH). *JASIST* 57(11): 1427–1439, 2006.
11. P. Subramanian, P. Tamayo, V.K. Mootha, S. Mukherjee, B.L. Ebert, M.A. Gillette. Gene set enrichment analysis: A knowledge based approach for interpreting genome-wide expression profiles. *Proc. of the National Academy of Science, USA*, 102(43): 15545–15550, 2005.
12. I. Trajkovski. Functional interpretation of gene expression data. Ph.D. Thesis, Jozef Stefan International Postgraduate School, December 2007.
13. I. Trajkovski, F. Železný, N. Lavrač, J. Tolar. Learning relational descriptions of differentially expressed gene groups. *IEEE Transactions of Systems, Man and Cybernetics C*, special issue on *Intelligent Computation for Bioinformatics*, 38(1): 16–25, 2008a.
14. I. Trajkovski, N. Lavrač, J. Tolar. SEGS: Search for enriched gene sets in microarray data. *Journal of Biomedical Informatics*, 41(4): 588–601, 2008b.
15. F. Železný, N. Lavrač. Propositionalization-based relational subgroup discovery with RSD. *Machine Learning*, 62(1–2): 33–63, 2007.
16. M. Weeber, H. Klein, L.T.W. de Jong-van den Berg, R. Vos. Using concepts in literature-based discovery: Simulating Swanson's Raynaud-fish oil and migraine-magnesium discoveries. *J. Am. Soc. Inf. Sci. Tech.*, 52(7): 548–557, 2001.