

Integration of Hybrid Bio-Ontologies using Bayesian Networks for Knowledge Discovery

Ken McGarry*, Sheila Garfield*, Nick Morris† and Stefan Wermter*

*Department of Computing and Technology, University of Sunderland, UK

†Institute for Cell and Molecular Biosciences, University of Newcastle, UK

{ken.mcgarry,sheila.garfield,stefan.wermter}@sunderland.ac.uk, n.j.morris@ncl.ac.uk

Abstract

This paper describes how high level biological knowledge obtained from ontologies such as the Gene Ontology (GO) can be integrated with low level information extracted from a Bayesian network trained on protein interaction data. We can automatically generate a biological ontology by text mining the type II diabetes research literature. The ontology is populated with the entities and relationships from protein-to-protein interactions. New, previously unrelated information is extracted from the growing body of research literature and incorporated with knowledge already known on this subject from the gene ontology and databases such as BIND and BioGRID. We integrate the ontology within the probabilistic framework of Bayesian networks which enables reasoning and prediction of protein function.

1 Introduction

The large amounts of genomic and proteomic data that are generated by biological experiments is now enabling deeper insights into cellular and molecular function. New technologies such as microarrays and electrophoresis gels are providing vast quantities of experimental data at unprecedented rates. All of this information needs to be stored and carefully annotated. With each new experiment providing details of new protein-to-protein interactions, new biological pathways and new genes it is essential that these discoveries are made available to the scientific community. To this end, online scientific databases are now in place that disseminate these results. These databases such as the popular Gene Ontology (GO) are updated at intervals to reflect the latest developments [Ashburner, 2000].

The updating is done by experts who manually revise each entry by reading the research literature and annotating the database collections accordingly. If necessary, they will contact the experimenters to resolve any ambiguities or problems. In terms of data quality, the databases are quite reliable and robust. Unfortunately, hand annotation is a slow process and the databases are lagging behind the experimental work by a considerable margin. This prevents researchers from immediately accessing the most recent discoveries.

Unless the researchers are familiar with the journals where the new results are published, they would be unlikely to encounter this information. Given, the fragmented and highly specialized nature of biological research, this may seldom occur. Therefore the need for automated extraction of knowledge from the literature is well motivated. However, recent advances in text analytics combines techniques from information retrieval (IR) and information extraction (IE) which allows researchers to explore the relevant literature more effectively [Mack and Henenberger, 2002]. However, these techniques require knowledge discovery methods to uncover complex embedded structures, relationships and connections between seemingly unrelated facts that typically exist in the biomedical literature [Tiffin *et al.*, 2005].

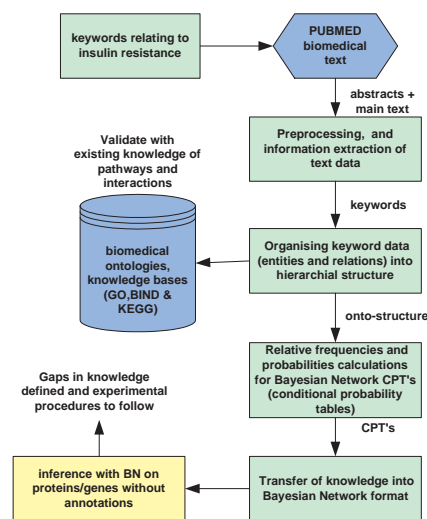


Figure 1: Overview of methodology and information extraction process

Our particular research area is that of diabetes, in particular the effects of insulin resistance on protein expression and insulin regulated protein trafficking in fat cells. In recent years there has been a dramatic worldwide increase of those suffering with diabetes. In the year 2000, there were 171 million cases and by 2030 the World Health Organization (WHO) has predicted there will be 366 million people suffer-

ing from this condition (www.who.int/diabetes/facts/). The WHO data is for diagnosed cases but the undiagnosed cases are estimated by the WHO at 14.6 million alone for the US.

In this paper we present our results of how we automatically generate a viable ontology based on information extraction of keywords from the research literature. The keywords define the entities and relationships of important genes, gene relationships, protein-to-protein interactions operate and co-exist in biological processes related to insulin resistance. Furthermore, the ontology is cast within a probabilistic framework using Bayesian networks which are used for the inferring and prediction of protein function. Figure 1 gives the overall methodology for the extraction of information and construction of the ontology.

The remainder of this paper is structured as follows; section two outlines our information extraction scheme for identifying the entities and relationships of interest, section three provides an overview of biological ontologies and gives details of how we use Bayesian networks for inference and reasoning. Section four discusses our methodology and experimental results, section five reviews the related work and our claim for novelty and finally section six presents the conclusions.

2 Information Extraction

Unstructured text is a very flexible and powerful means of communication, it allows us to describe quite complex concepts. The semantic meaning of a sentence can be expressed in many different ways but it is this flexibility which is the cause of difficulty for algorithmic sentence analysis by computers. One technique of overcoming this problem is to use information extraction (IE) to seek out the important entities in the text and the relationships between them [Hearst, 1992; Rosario and Hearst, 2004]. The IE process can involve encoding patterns by hand such as *regular expressions* to search for the required entities and relations or to use semi-automated machine learning techniques [Nahm and Mooney, 2002; Krauthammer and Nenadic, 2004]. The algorithm we developed is shown in figure 2.

Inputs: Abstract file A, String str

Outputs: Keyword file B

Load file A

While unprocessed "abstracts" in A

Remove end of line characters

Read each line into str

Search string for concept term

If contains phrase (the | a | an) + 2words(and |) + 2words

write word preceding key phrase and string after key phrase to B

elseif str contains phrase (the | a | an) + 1word(and |) + 2words

write word preceding key phrase and string after key phrase to B

elseif str contains phrase (the | a | an) + 2words

write word preceding key phrase and string after key phrase to B

close A and B

Figure 2: Information extraction algorithm

The algorithm encodes through regular expressions templates for recognizing the types of "action" words that typically occur in biological texts. We discuss this process in more detail in section 4. However, the main problem that our algorithm considers is to discover in advance the kind of information that can be encountered. Rather than attempt to parse the entire corpus we exploit certain linguistic regularities and search for specific semantic relations that need only be defined once. The algorithm takes into account a variable distance between related terms i.e. longer passages of text, and therefore provides a much more reliable identification of the relationships. Seeking up two words difference has empirically shown to be a reasonable trade-off of accuracy versus computational complexity. Examples of relationships include:

- A inhibits B
- A activates B
- A interacts with B
- A suppresses B

3 Biological Ontologies and Bayesian Networks

In this section we briefly motivate the need for ontologies and define their limitations with respect to the biological field and for knowledge discovery. Ontologies describe the concepts and relationships that exist for a particular area of interest. They are very useful for the semantic labeling of concepts or definitions [Grivell, 2002; Bard and Rhee, 2004]. This process ensures that entities which are equivalent to other entities in separate databases are identified as referring to the same concepts. Even if these entities have different names or forms they can still be identified by semantic labeling. The role of semantics therefore is much deeper than matching the co-occurrence of a tag or label, since it defines the relationship that exists between concepts. Figure 3 shows the structure and elements of the gene ontology that are pertinent to our study. The first entry refers to GO:0008150 and is one of the three top level structures (biological process, physiological process and cellular process) in the gene ontology hierarchy; the last number (GO:0015758) defines the relationships for the glucose transport pathway. The numbers in brackets refer to the number of entries at that particular level.

The use of ontologies in biology for the semantic integration of heterogeneous data is receiving increased attention, however problems occur because of the dynamic, changing nature of biological knowledge [McGarry *et al.*, 2006]. These difficulties arise from the highly complex structures that are expensive and problematic to update and maintain [Blaschke and Valencia, 2002]. Another, related problem is that current ontologies have a rather limited vocabulary and cannot express the richness of biological information. Little attention has been paid to defining the *relations*, much of the research effort and complexity of structure has concentrated on defining the *terms*. Other considerations that are important are the spatial and temporal characteristics of the entities.

Furthermore, ontologies such DAML+OIL, OWL and RDF are based on crisp logic and have difficulty managing

Accession: GO:0015758
Ontology: biological process
Synonyms: None
Definition: movement of the hexose monosaccharide glucose into, out of, within or between cells.
GO:0008150 : biological process (127987)
GO:0009987 : cellular process (78769)
GO:0050875 : cellular physiological process (71999)
GO:0006810 : transport (21084)
GO:0008643 : carbohydrate transport (498)
GO:0015749 : monosaccharide transport (206)
GO:0008645 : hexose transport (168)
GO:0015758 : glucose transport (115)

Figure 3: GO structure for Glut4 protein within glucose transport pathway

uncertainty; incomplete data and noisy information that is encountered in many domains, especially the bioinformatic field. Our research is concerned with Type 2 diabetes, in order to develop a suitable ontology it is necessary to identify the relevant entities within the domain, their attributes and the relationships that exist between these entities.

3.1 Bayesian networks for Ontology Inference and Integration

The integration of sub-symbolic and symbolic computation has received considerable interest over the years [McGarry *et al.*, 1999]. Within this framework the Bayesian approach can be seen as both a learning mechanism and as a knowledge representation technique.

Bayes theorem is shown in equation 1 and presents the probability of the hypothesis (H) conditionalised on evidence (E).

$$P(H | E) = \frac{P(E | H)P(H)}{P(E | H)P(H) + P(E | \neg H)P(\neg H)} \quad (1)$$

where: $P(H | E)$ defines the probability of a hypothesis conditioned on certain evidence, $P(E | H)$ is the likelihood, $P(H)$ is the probability of the hypothesis prior to obtaining any evidence, is the $P(E)$ evidence. Therefore, according to Bayesian theory we can update our beliefs regarding the hypothesis when provided with new evidence that is conditional upon using probabilities and is called *conditionalization*.

The conditional probability distributions (CPD) are described by $P(X_i | U_i)$, where X_i represents node i and U_i are its parent nodes. We must specify the prior probabilities of the nodes and the conditional probabilities of the nodes given all the combinations of their ancestor nodes. The joint distribution of random variables is given by $X = \{X_1, \dots, X_n\}$ and together with the CPD values is used to calculate the choice of X_i and is given by :

$$P(X_1, \dots, X_n) = \prod_i P(X_i | U_i) \quad (2)$$

The CPD's values are easy enough to calculate and inference but require the number of parameters is dependent upon

the number of parent nodes, they are usually represented in table format. The nodes are assumed to be discrete or categorical values, however, continuous values may be discretised [Korb and Nicholson, 2004].

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_j \pi_j[C_j] \quad (3)$$




Figure 4: The advantages of Bayesian networks include a graphical representation of the structure i.e. the interconnection relationships between the variables of interest and they allow for causal discovery or causal interpretation. The example shows the relationships between the insulin resistance problem and how it is affected by the proteins ACE and GLUT4 and the effects of insulin resistance upon other proteins such as GLUT1 and ADRB3.

In figure 4, the various possibilities for inferencing are shown within the insulin resistance domain. The first network shows the *diagnostic reasoning* approach which enables the relationships between symptoms and causes to be evaluated, thus when given some evidence regarding the presence of *Glut4* we can update our beliefs about the likelihood of IR being present. When using *predictive reasoning* we can derive new information about effects given some new information regarding the causes.

4 Methods and Results

We reviewed the literature associated with Type 2 diabetes, the initial focus associated with protein interaction in diabetes and from this review a list of “events” indicative of protein interactions was identified, eg, activate, inhibit and modulate. This list was used as the starting point to help identify which entities are involved in each type of action or relation. After identifying the names of possible event relations the focus moved to identifying potential entities involved in these relations. In order to complete this task a suitable dataset was required. A search of the PubMed database was conducted and 6113 abstracts, related to Type 2 diabetes were used; this dataset is used throughout each subsequent stage of this work. Initially a count was made of the number of times each of the action words occurred in this sample dataset. Some of the words, eg, “acetylate” and “destabilize” did not occur at all, while other words such as “interaction” and “suppression” occurred more frequently.

We now explain how the various parts of our system function together, the information extraction technique synthesizes the entities and relationships from the literature abstracts and generates the structure for a specific ontology on

Table 1: Biological keywords

Action Word	No	Action Word	No	Action Word	No
acetylate	0	inhibit	109	phosphorylates	5
acetylated	1	inhibited	95	phosphorylation	362
acetylates	0	inhibition	222	regulate	62
acetylation	0	inhibits	59	regulated	62
activate	47	interact	34	regulates	35
activated	69	interacted	0	regulation	333
activates	18	interacting	14	stabilization	6
activation	435	interaction	213	stabilize	3
bind	31	interactions	101	stabilized	3
binding	914	interacts	7	stabilizes	3
binds	16	modulate	74	suppress	56
bound	31	modulated	23	suppressed	116
destabilization	0	modulates	25	suppresses	13
destabilize	1	modulation	59	suppression	386
destabilized	0	phosphorylate	13	target	235
destabilizes	0	phosphorylated	15		

insulin resistance. We then use the ontologies structure to build a Bayesian network for the purposes of inference and prediction of new protein-to-protein interactions. The relative frequencies of the keywords (entities and relationships) are used to construct the conditional probability tables which define the parent/child node relationships.

4.1 The Extracted Ontology and Bayesian network Mapping

Initially, one of these action words, “interaction” was selected to identify possible entities involved in a relation. The word “interaction” however generally forms part of a phrase such as “interaction between”, “interaction of”, and “interaction with”, and therefore each of these phrases would be used by the algorithm to search for potential entities. The first phrase used was “interaction between”. Examples of the resulting phrases extracted are provided in the table 2.

Table 2: Biological keywords extracted for the ontology for the phrase “interaction between”

Preceding word	Following words
the	thyroid function and insulin sensitivity
the	dysregulated fat and glucose metabolism
strong	insulin resistance and serum
significant	obesity and insulin resistance
possible	BMI and the adiponectin gene

Ultimately, the successful application of Bayesian techniques is dependent on the use of *prior knowledge* to improve the estimation of the posterior. If a prior belief exists about a situation then we can use this information to pre-structure our BN. For example if a particular gene (IPA) is known to regulate several target genes (GDH, GL4, HK2), we would then assign this relationship within the BN by setting the edges between these two entities and setting the values in the conditional probability table to define the structural prior accordingly. This is a powerful strategy, but only when it makes

sense to do so. The application of incorrect beliefs will produce unreliable estimates of the true posterior regardless of the abundance of the likelihood evidence. Equation 4 shows how we modify the BN with prior knowledge (causal intervention) from the extracted ontology [Chrisman *et al.*, 2003].

$$P(X_{i,j} = z \mid par_M(x), M, \theta : X_{i,j} = Z, \dots) = 1 \quad (4)$$

where par_M are the parameters within the model, $X_{i,j}$ are the known effects of the parents of a given node, θ is the conditional probability conditionalized and represents the causal conditions. The biological knowledge is incorporated into the BN by specifying the probability for the existence of each potential connection (edge) between them. We assume independence between edges and the variables in the BN are also assumed to be discrete, this ensures that the calculations are computationally tractable.

Figure 5 shows the structure of a section of our ontology. The nodes are the entities and the arcs determine the relationships between them. The numbers in brackets preceded by “GO:” are the probabilities of the term occurring in the GO ontology, the numbers.

For example the following abstract fragment captures knowledge about several proteins and their interactions:

“Overexpression of the cytosolic domain of syntaxin 6 did not affect insulin-stimulated glucose transport, but increased basal deGlc transport and cell surface Glut4 levels. Moreover, the syntaxin 6 cytosolic domain significantly reduced the rate of Glut4 reinternalization after insulin withdrawal and perturbed subendosomal Glut4 sorting; the corresponding domains of syntaxins 8 and 12 were without effect.”

We encountered difficulties with negative implications, i.e. the “did not” and “without effect” phrases negate the occurrence of the relationship but would be taken by the information extraction algorithm as a positive relationship. A more elaborate NLP technique or further crafting of specific regular expression templates would reduce this effect.

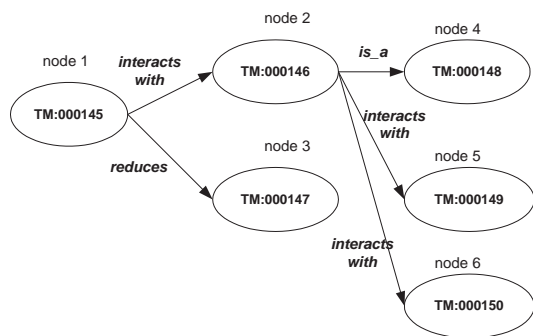


Figure 5: Fragment of the ontology (entities and relations) extracted from the literature

4.2 Validation against Existing Knowledge

We determined a base line accuracy for our system by “re-discovering” known protein-to-protein interactions from the literature and validating the relationships through accessing a number of online database and ontology repositories. The most up to date and complete is the gene ontology (GO), we compare extracted relationships from our ontology with the GO structure. To determine the accuracy, we apply the well known information retrieval measures of recall and precision. We define recall as the percentage of entity relations represented in the GO and correctly identified. We define precision as the the percentage of relations found in GO and returned by our system.

The recall and precision are calculated by:

$$recall = TP / (TP + FN),$$

$$precision = TP / (TP + FP),$$

where: TP=true positives such as , FP= false positives, TN=true negatives and FN= false negatives.

Table 3: Recall and Precision of IE on protein-to-protein interaction data

Keyword	TP	TN	FP	FN	Recall	Precision
interact	100	171	20	32	37	83
bind	200	167	17	14	54	92
promote	240	188	17	15	56	93
inhibit	230	178	12	19	56	95

We should note that certain errors in GO have been identified, inconsistencies and even spelling mistakes. We have also identified that certain GO terms are too general and a more specific term would have been more appropriate. Thus entries with low semantic similarity but high functional similarity can be identified. Figure 6 presents the results of a comparison between the semantic richness between GO and our extracted ontology. We define the semantic richness measure to be based on the correlations between functional similarity and semantic content, a detailed description of this approach can be found in [Lord *et al.*, 2003].

The GO ontology structure is extremely limited with total reliance on “*is_a*” type links. This means that a large amount of semantic information that was originally available

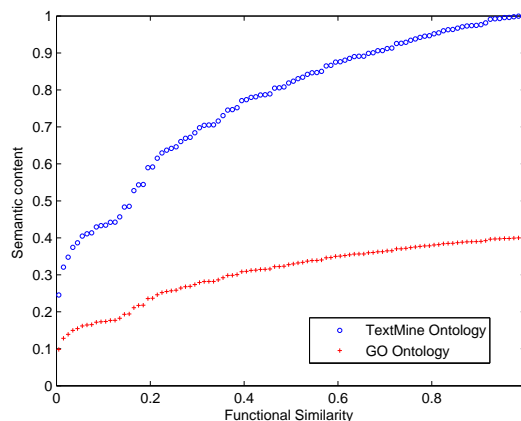


Figure 6: Comparison of the semantic richness of vocabulary of the GO and Text Mine ontologies.

from the research articles is missing. We suspect that as ontologies such as GO increase in the number of entities, the relationships between will take on increased value. However, without incorporating the semantic similarity of the entities any increase in size will reduce the ontology to free text.

5 Related Work

Research into the automatic generation of ontologies from textual data has received limited attention to date, notable exceptions are the work of Blaschke and Valencia, which used clustering techniques at a document level [Blaschke and Valencia, 2002]. The majority of the research attempts to alleviate partial gaps in the knowledge or to repair incorrect annotations in existing ontologies [Missikoff *et al.*, 2003; Wolsencroft *et al.*, 2005]. Using probabilistic techniques to model ontologies is receiving increased attention but this is for manually curated ontologies [Mitra *et al.*, 2005; Smith *et al.*, 2005]. The modeling of biological networks with bayesian networks using genomic data has seen considerable attention in recent years [Ong *et al.*, 2002]. The initial work on integrating heterogeneous data within a bayesian network framework was led by Friedman and Segal [Friedman *et al.*, 2000; Segal *et al.*, 2001]. This work proved that Bayesian networks could be trained on genomic data to reconstruct the relationships between genes. The work by Pan *et al.* is the most similar to ours, however the authors used Bayesian networks to integrate two ontologies from similar problem domains [Pan *et al.*, 2005]. Comparisons between the semantic similarity and genetic sequence similarity of ontologies has been conducted by Lord [Lord *et al.*, 2003]. We found this work particularly useful as motivation for the development of a richer vocabulary to define entity relationships.

6 Conclusions

The fusion of low level information from sub-symbolic techniques with logic or higher order structures is critically dependent on the level of granularity used. The nodes of our

Bayesian networks are robust to semantic topic drift or catastrophic interference which typically occurs when MLP or other neural feed-forward techniques are trained in dynamic situations using heterogeneous data. In the case of our bioinformatics work we use Bayesian networks to learn from data but also to map existing ontological relations to new Bayesian network structures. Clearly, further work is needed, however, we have extended the current knowledge of automatically generating and integrating ontologies from low level data. The utilization of ontologies as a framework for guiding the knowledge discovery process has to date received little attention. The experimental results presented in this paper led us to conclude that a principled approach such as the Bayesian framework can successfully integrate and represent heterogeneous data and knowledge.

7 Acknowledgements

This work was part supported by a Research Development Fellowship funded by HEFCE and the Biosystems Informatics Institute (Bii).

References

- [Ashburner, 2000] M. Ashburner. Gene ontology: tool for the unification of biology. *Nature Genetics*, 25:25–29, 2000.
- [Bard and Rhee, 2004] J. Bard and S. Rhee. Ontologies in biology: design applications and future challenges. *Nature Reviews Genetics*, 5:213–222, 2004.
- [Blaschke and Valencia, 2002] C. Blaschke and A. Valencia. Automatic ontology construction from the literature. *Genome Informatics*, 13:201–213, 2002.
- [Chrisman *et al.*, 2003] L. Chrisman, P. Langley, S. Bray, and A. Pohorille. Incorporating biological knowledge into evaluation of causal regulatory hypothesis. In *Proceedings of the Pacific Symposium on Biocomputing*, pages 128–139, Kauai, Hawaii., 2003.
- [Friedman *et al.*, 2000] N. Friedman, M. Linial, I. Nachman, and D. Pe'er. Using bayesian networks to analyze expression data. *Journal of Computational Biology*, 7(3-4):601–620, 2000.
- [Grivell, 2002] L. Grivell. Mining the bibliome: searching for a needle in a haystack?: new computing tools are needed to effectively scan the growing amount of scientific literature for useful information. *EMBO Reports*, 3(31):200–203, 2002.
- [Hearst, 1992] M. Hearst. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the 14th conference on Computational linguistics*, pages 539–545, 1992.
- [Korb and Nicholson, 2004] K. Korb and A. Nicholson. *Bayesian Artificial Intelligence*. Chapman and Hall/CRC, 2004.
- [Krauthammer and Nenadic, 2004] M. Krauthammer and G. Nenadic. Term identification in the biomedical literature. *Journal of Biomedical Informatics*, 37:512–526, 2004.
- [Lord *et al.*, 2003] P. Lord, R. Stevens, A. Brass, and C. Goble. Investigating semantic similarity measures across the gene ontology: the relationship between sequence and annotation. *Bioinformatics*, 19:1275–1283, 2003.
- [Mack and Henenberger, 2002] R. Mack and M. Henenberger. Text-based knowledge discovery: search and mining of life-sciences documents. *Drug Discovery Today*, 7:11, 2002.
- [McGarry *et al.*, 1999] K. McGarry, S. Wermter, and J. MacIntyre. Hybrid neural systems: from simple coupling to fully integrated neural networks. *Neural Computing Surveys*, 2(1):62–93, 1999.
- [McGarry *et al.*, 2006] K. McGarry, S. Garfield, and N. Morris. Recent trends in knowledge and data integration for the life sciences. *Expert Systems: the Journal of Knowledge Engineering*, 23(5):337–348, 2006.
- [Missikoff *et al.*, 2003] M. Missikoff, P. Velardi, and P. Fabriani. Text mining techniques to automatically enrich a domain ontology. *Applied Intelligence*, 18:323–340, 2003.
- [Mitra *et al.*, 2005] P. Mitra, N. Noy, and A. Jaiswal. Ontology mapping discovery with uncertainty. In *Fourth International Semantic Web Conference (ISWC)*, 2005.
- [Nahm and Mooney, 2002] U. Nahm and R. Mooney. Text mining with information extraction. In U. Nahm and R. Mooney. *Text Mining with Information Extraction. In Proceedings of the AAAI 2002 Spring Symposium on Mining Answers from Texts and Knowledge Bases.*, 2002.
- [Ong *et al.*, 2002] I. Ong, J. Glasner, and D. Page. Modelling regulatory pathways in *E. coli* from time series expression profiles. *Bioinformatics*, 18(1):241–248, 2002.
- [Pan *et al.*, 2005] R. Pan, Z. Ding, Y. Yu, and Y. Peng. A bayesian network approach to ontology mapping. In *ISWC 2005 4th International Semantic Web Conference*, pages 563–577, Galway, Ireland, 2005.
- [Rosario and Hearst, 2004] B. Rosario and M. Hearst. Classifying semantic relations in bioscience texts. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL2004)*, pages 430–437, 2004.
- [Segal *et al.*, 2001] E. Segal, B. Tasker, A. Gasch, N. Friedman, and D. Koller. Rich probabilistic models for gene expression. *Bioinformatics*, 17(1):243–252, 2001.
- [Smith *et al.*, 2005] B. Smith, W. Ceusters, and J. Kohler. Relations in biomedical ontologies. *Genome Biology*, 6(5):46–58, 2005.
- [Tiffin *et al.*, 2005] N. Tiffin, J. Kelso, A. Powell, H. Pan, V. Bajic, and W. Hide. Integration of text and data-mining using ontologies successfully selects disease gene candidates. *Nucleic Acids Research*, 33(5):1544–1552, 2005.
- [Wolstencroft *et al.*, 2005] K. Wolstencroft, R. McEntire, R. Stevens, L. Taberner, and A. Brass. Constructing ontology-driven protein family databases. *Bioinformatics*, 21(8):1685–1692, 2005.