

Ontology-Based Interactive Information Extraction from Scientific Abstracts

Milward, David¹, Bjärelund, Marcus², Hayes, William³, Maxwell, Michelle⁴, Öberg, Lisa², Tilford, Nick⁴, Thomas, James¹, Hale, Roger¹, Knight, Sylvia¹ and Barnes, Julie⁴

¹ Linguamatics Ltd., St. John's Innovation Centre, Cambridge, CB4 0WS, UK

² AstraZeneca R&D Mölndal, SE-431 83 Mölndal, Sweden

³ AstraZeneca R&D Boston, 35 Gatehouse Drive, Waltham, MA 02451, USA

⁴ BioWisdom Ltd, Babraham Hall, Babraham, Cambridge, CB2 4AT, UK

david.milward@linguamatics.com, marcus.bjareland@astrazeneca.com,
william.s.hayes@astrazeneca.com, julie.barnes@biowisdom.com.

Abstract

Over recent years, there has been a growing interest in extracting information automatically or semi-automatically from the scientific literature. This paper describes a novel Ontology-based Interactive Information Extraction (OBIIE) framework and a specific OBIIE system. We describe how this system enables life scientists to make *ad hoc* queries similar to using a standard search engine, but where the results are obtained in a database format similar to a pre-programmed information extraction engine. We present a case study where the system was evaluated for extracting cofactors from EMBASE and MEDLINE.

Introduction

Information Retrieval (IR) systems are designed to find the highest ranked documents that match a user query, such as a set of keywords. This contrasts with Information Extraction (IE) which returns relationships, e.g. a table of protein-protein interactions or gene-disease relationships, rather than a ranked set of documents. Traditionally these two techniques have been seen as very different: IR is all about finding documents, IE about finding facts within documents. However, from a user perspective the difference is not so great. The users' aim is primarily to get to information as fast as possible: if a system can get them directly to relevant sentences then this will save time. If the system can also get them to structured results, appropriately sorted, this can save further time. Interactive Information Extraction is a new concept which combines the interactive querying style of a web search engine with the structured output that is provided by standard IE. This allows scientists to refine queries and explore a set of texts in a similar way to web search, but with the possibility of much more precise search and results.

In the Linguamatics I2E System, the user can start with a standard search for words within a document, then refine this to require the words to be in the same sentence, or in a particular linguistic pattern. The results are output as HTML tables or in a format suitable for database entry. The system pre-indexes documents such as MEDLINE abstracts to allow fast querying. All linguistic processing is done prior to indexing, including "tokenisation" to split a string of characters into individual words, "sentence splitting" to recognise sentence ends, "tagging" to recognise parts of speech such as nouns or verbs, and "chunking" to group words into meaningful units according to their parts of speech.

The OBIIE Framework and System

Conventional IE systems typically allow extraction of syntactic classes (nouns or verbs) and a few unstructured semantic classes, so called named-entities such as proteins, diseases, or amounts.¹ Until recently there has been little use of richer domain knowledge within such systems. In OBIIE, ontologies provide that domain knowledge, enabling the users to interact with the system on a conceptual level without having to know all possible synonyms for a concept. The relationships “part-of” or “is-a” provide a basic taxonomy allowing the user to choose particular concepts or families of concepts (see Figure 1). The user can now construct queries using the many thousands of different classes found in typical ontologies.

A scientist may start with a query such as two keywords in the same document (equivalent to standard keyword search), then refine this to look for two *classes* e.g. a disease and a protein (co-occurrence within a document), then refine this further to require co-occurrence within a sentence, and finally refine this to look for the protein and disease in a particular syntactic configuration. Incorporation of linguistic constraints is up to the user, and is justified by an increase in precision that is enough to balance any decrease in recall. For example, a query for the word “RAF” followed by the word “phosphorylate” would be improved by putting “phosphorylate” within a so-called “verb group”. This would then match the text string “Raf has been shown to phosphorylate”, without returning the larger number of false hits you would have got by allowing Raf within e.g. 5 words of phosphorylate.

The first OBIIE system was constructed by putting together ontologies from BioWisdom with the I2E System. The ontologies incorporated within OBIIE express ‘is-a’, ‘is-a -part-of’ (i.e. taxonomic) relationships as well as the ‘has synonym’ relationship. Ontological concepts can be selected from up to 25 concept types, ranging from genes, proteins, tissues, cells, clinical disorders, symptoms, processes, pathways, drugs, adverse effects, and techniques, technologies etc. Where appropriate, ontologies can be used as species-specific structures, or as master species-independent ontologies.

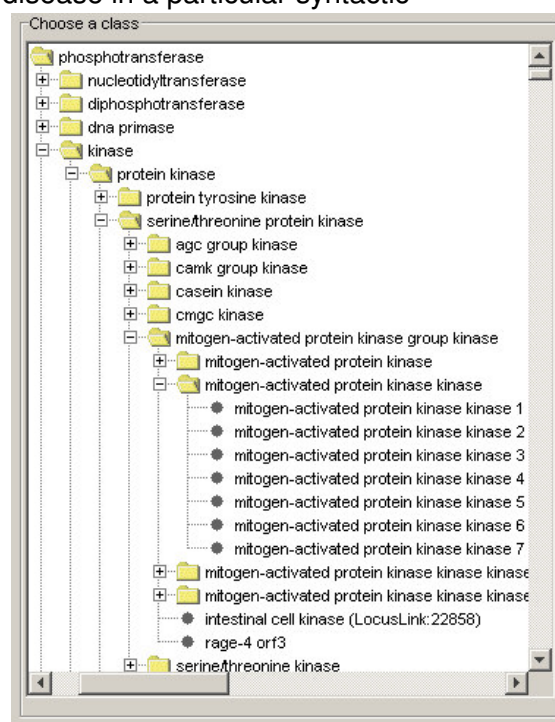


Figure 1: Protein Ontology

As well as defining patterns on the fly, it is possible to reuse existing patterns or pattern templates. Patterns can be organised hierarchically to provide relationship ontologies, so that a user interested in interactions between drugs and proteins can choose the family of “drug interaction relationships”, and obtain alternative phrasings such as the verbs “activates”, “inhibits”, “blocks” (and their morphological variants) or phrases such as “is agonist for”.

¹ Named entities may be closed classes formed from an enumerated list e.g. a list of names, or open classes recognised via patterns (e.g. protein spotting routines such as Fukuda et al. 1998).

Case Study : Text-Mining for Nuclear Receptor Cofactors

In this evaluation we compared the use of two methods. The first consisted of finding a set of abstracts, reading the abstracts, and extracting cofactors from these. The second used the OBIIE tool over a larger set of abstracts, extracting the results, and then filtering these by hand. We were particularly interested in the relative recall between the two methods, and the speed of each method.

Nuclear receptors (NRs) are ligand-dependent transcription factors that typically recruit protein complexes (cofactors) to enhance or repress transcription of target genes. It is believed that several phenomena such as level of transactivation and tissue specificity of NRs heavily depend on the specific recruited cofactors. Since NRs are very important drug targets (18 of the 48 known human NRs are targets for registered drugs) the amount of literature on these proteins is rapidly increasing. The aim of this case study was to generate a comprehensive and annotated lists of cofactors for three NRs: Androgen Receptor (AR) and Liver X Receptors (LXR) α and β . AR abstracts were used for “training” (i.e. query tuning and ontology refinement) and LXR abstracts for testing. The project validated the use of ontology- and linguistic-based text-mining against the previous best practice for obtaining sets of cofactors, based on published or manually generated lists.

For AR we constructed a secondary corpus of 7748 abstracts from MEDLINE and EMBASE by fine-tuning synonym choices. We used two sources for the manual list of cofactors for AR: A list of AR-interacting proteins compiled by Dr. Lenore Beitel (Beitel, 2002), and a list constructed by manually examining a subset of abstracts from the secondary corpus containing ~300 abstracts. The abstracts underlying this sub-corpus were further analyzed in several iterations with regard to sentences containing information on cofactor recruitment. When we reached sufficient recall (~90%) of extracted relationships we stopped the iterations. The relationship ontology includes various alternative phrasings and words within larger patterns. Compiling out the embedded alternatives gives a relationship ontology representing 188 distinct patterns. The results of applying the OBIIE system with the relationship ontology are shown in Table 1. The two top reasons why we failed to reach 100% recall was that the secondary corpus did not contain any abstracts with a reference to those symbols (in fact, no abstracts in MEDLINE or EMBASE contained co-occurrences of those symbols and AR), and that references spanned several sentences. Cross-sentence patterns were not used as they severely affected precision.

Table 1. Results for extraction of NR cofactors.

NR	# Abstracts in secondary corpus	# Abstracts retrieved by OBIIE	# Cofactors manually retrieved	# Cofactors retrieved by OBIIE	#Cofactors in found in total
AR	7748	564	101	100	110
LXR α/β	N/A	68	9	9	10

Instead of cross-validating the results on the AR secondary corpus, we applied the same query, with the AR synonyms replaced by a selected number of LXR α and β synonyms to the whole of MEDLINE (i.e. we did not construct a secondary corpus in this case). The manual list was generated by manual examination of 240 LXR abstracts, but the relationship ontology was not changed based on any of those abstracts. The results of applying OBIIE to extract LXRs from the whole of Medline (the set of Medline abstracts as of Nov 25th 2003) are displayed in Table 1. One cofactor was missed by the OBIIE system due to a missing pattern in the relationship ontology. The OBIIE system managed to retrieve a cofactor that manual curation missed. For

LXRs, recall was 90%. Recall was calculated by dividing the number of cofactors found by OBIIE, divided by the best figure available for the total number of cofactors (the total number of cofactors discovered by either method). Lists of the retrieved cofactors can be found at <http://www.linguamatics.com/obiie/>.

The results obtained by the OBIIE system are clearly satisfactory, not only for the recall achieved, but also for the amount of time saved by the automated process. Without risking too many human errors we estimate that one person can read 100 abstracts in a day. With the tabular output format from the OBIIE system, a domain expert can increase this by an order of magnitude, since only sentences with the NR symbol, a relationship phrase, and a cofactor symbol are displayed.

Related Work and Discussion

In this extended abstract, we have described the new method of Ontology-Based Interactive Information Extraction. There is prior work on the integration of ontologies with standard Information Extraction. For example, Mädche et al. (1999) discuss the use of ontologies as a way to provide information in a canonicalised format to a user or for input into a database, and Todirascu et al. (2002) use ontologies for identifying concepts. Although there are precedents for the use of ontologies in information extraction, there are a number of necessary features of any ontology applied in this way. The BioWisdom ontologies used in this case study incorporate a comprehensive list of synonyms, and have a detailed hierarchical structure allowing fine-grained concept distinctions. The domain specific nature of the ontologies also ensures that the synonyms are appropriate to the pharmaceutical domain.

Interactive Information Extraction has few precedents. There is some similarity with work on *question answering systems*. Here NLP and IR techniques are used to obtain the best answer to a question expressed in natural language. In contrast, I2E provides all results for a structured query. There is also some similarity with work which provides an interactive front end on top of the output of a fixed IE system (e.g. Gazaiskus et al. 2001). However, in these systems the result of a new query will always be a subset of the results provided by the original fixed IE query patterns. There are other systems which can be said to be positioned somewhere between document search (IR) and relationship search (IE), but they are typically designed for one specific task, e.g. looking for symbol co-occurrences in sentences. In contrast, in Interactive Information Extraction there is a natural gradation from document search, via search within sentences, to search for specific relationships within sentences. It is also possible to perform combined searches e.g. search for relationships, but only in documents containing a particular concept.

An examination of cofactors similar to the case study was performed by Albert et al. 2003 using a co-occurrence based approach. They searched for tri-occurrences, i.e. entity, relation, and entity within the same sentence. Recognition of entities was performed using finite state string matching automata (regular expressions) rather than using linguistic processing. They retrieved fewer cofactors, but it is not appropriate to compare the results directly, since they were working with a smaller corpus, theirs being extracted 10th Sep 2001 and ours during the fall of 2003 (Oct 21st for AR and Nov 25th for LXRs).

The case study presented here highlights the power of OBIIE in performing systematic textual analysis. The use of synonyms, coupled with the interactive nature of the tool makes it very quick to engineer queries to accommodate the variety of linguistic forms that phrases take, and hence allows for high recall coupled with good precision. The results for cofactors were remarkably good, with new cofactors discovered by the OBIIE system that had not been in the

original manually retrieved set. By exploiting the redundancy inherent in a large corpus, we were able to use relatively specific patterns giving high precision, while still getting recall that rivals that achieved manually. Although the case study in this paper focused on extracting cofactors, the incorporation of other large biomedical ontologies covering areas of disorders, symptoms, tissues, cells, compounds, biological processes etc. makes this a flexible tool for use across all parts of pharmaceutical R&D.

Bibliography

- Albert, S., Gaudan, S., Knigge, H., Raetsch, A., Delgado, A., Huhse, B., Kirsch, H., Albers, M., Rebholz-Schumann, D. & Koegl, M. (2003) Computer-assisted generation of a protein-interaction database for nuclear receptors, *Molecular Endocrinology*, 2002-0424
- Beitel, L. (2002) List of AR-interacting proteins, available from *The Androgen Receptor Gene Mutations Database World Wide Web Server*: <http://ww2.mcgill.ca/androgendb/>
- Biowisdom, <http://www.biowisdom.com/>
- Blaschke, C., Hirschman, L., and Valencia, A. (2002) Information Extraction in molecular biology. *Briefings in Bioinformatics*. 3(2):1-12.
- Craven, M. and Kumlien, J. (1999) Constructing biological knowledge bases by extracting information from text sources. *ISMB*, AAAI Press, pp. 77-86.
- De Bruin, B. & Martin, J. (2002) Literature mining in molecular biology. *EFMI Workshop on Natural Language Processing in Biomedical Applications*, Baud, R. & Ruch, P. (eds.), Nicosia, Cyprus, 1-5. *EMBASE*, <http://www.embase.com>
- Fukuda, K, Tsunoda, T., Tamura, A., and Takagi, T. (1998) Toward information extraction: identifying protein names from biological papers. *Proc. PSB, Hawaii*, 3, 705-716.
- Gaizauskas, R., Herring, P., Oakes, M., Beaulieu, M., Willett, P., Fowkes, F. and Jonsson, A. (2001). *Intelligent Access to Text: Integrating Information Extraction Technology into Text Browsers*. *Proceedings of the Human Language Technology Conference (HLT2001)*, 189-193
- Hopkins, A.L. and Groom, C.R (2002) The Druggable Genome. *Nature Reviews Drug Discovery*, Vol. 1, 727-730.
- Humphreys, K., Demetriou, G. and Geizauskas, R. (2000) Two applications of information extraction to biological science journal articles: Enzyme interactions and protein structure. *Pac. Symp. Biocomp.*, pp. 502-513.
- Jenssen, T.K, Laegrid, A., Komorowski, J., Hovig, Eivind, (2001), 'A literature network of gene for high-throughput analysis of gene expression', *Nature Genetics*, 28:21-28.
- Linguamatics Ltd., (2003) *Interactive Information Extraction: White Paper*, http://www.linguamatics.com/resources/white_paper_ie.html
- Maedche, A., Staab, S. and Studer, R. *Ontology-Oriented Information Extraction and Integration*. *Workshop on Language Technologies in Information and Knowledge Management (1999)*, 7th *Conference on Computational Linguistics of the German Society for Language Technologies*. Saarbrücken, Germany, October 7-8.
- Medline, <http://www.ncbi.nlm.nih.gov/entrez/query.fcgi?db=PubMed>
- Rindflesch, T., Rajan, J. and Hunter, L. (2000) Extracting Molecular binding relationships from biomedical text. *Proc. of the ANLP-NAACL, ACL*, pp. 188-195.
- Seger, R. and Krebs E.G. (1995) The MAPK Signaling Cascade. *FASEB J*, 9(9), 726-735
- Sekimizu, T., Park, H.S., and Tsujii, J. (1998), Identifying the interaction between genes and gene products based on frequently seen verbs in Medline abstracts. *Genome Inf. Serv.* pp. 62-71.
- Stapley, B.J. and Benoit, G. (2000), *Biobibliometrics: information retrieval and visualization from co-occurrences of gene names in Medline abstracts*. *Pac. Symp. Biocomp.*, pp. 529-540.
- Thomas, J., Milward, D., Ouzounis, C., Pulman, S. and Carroll, M. (2000) Automatic extraction of protein interaction from scientific abstracts. *Pac Symp Biocomp*, 541-552.
- Todirascu, A., Romary L., Bekhouche, D. (2002) Vulcain - An Ontology-Based Information Extraction System. *Natural Language Processing and Information Systems: 6th International Conference on Applications of Natural Language to Information Systems, NLDB 2002*, Stockholm, Sweden, 64-75.