

'Literature-based research is critical in turning an idea into a marketable therapeutic product 12–15 years later.'

editorial



Roger Hale,
Chief Operating Officer, Linguamatics Ltd

Text mining: getting more value from literature resources

► Scientists in the pharmaceutical and biotech industries need to review vast quantities of literature during the discovery and development of a new drug. Pinpointing the most relevant information is hard with current search tools. New developments in text mining are proving effective in getting better value from available literature resources.

Literature-based research: current challenges

The past forty years have seen a fivefold increase in the number of abstracts published annually in Medline. There are now more than 12 million, and most are also available online as full-text articles. In addition to this there are patents, internal reports,

and other potentially valuable in-house and public sources. Although a small proportion of this information is in a structured form that can be managed using database systems, around 80% is unstructured and written in a natural language. This is far too much to review or keep up-to-date with manually.

Literature-based research is critical in turning an idea into a marketable therapeutic product 12–15 years later. Industry estimates suggest that 90% of drug targets are derived from the literature [1]. In fact, most background research for drug discovery is publicly available through published articles and abstracts, patent application data, and other sources. Strategic decisions during R&D need to be supported by information extracted from a wide range of text content, sometimes about unrelated or unfamiliar topics. In particular, it is usually important to extract and understand relationships, for example between genes and diseases or compounds and side effects. Typical questions include:

- Which proteins interact with my target protein?
- Who is patenting on what targets, compounds and diseases?
- Will my lead compound fail due to toxicity?

In the case of the last question, surveys suggest that about 50% of all potentially therapeutic compounds undergo attrition due to safety concerns and that about 50% of them had some indication in the literature already [1].

To answer questions like those above from unstructured information sources using conventional search technologies requires a large amount of human intervention and curation. Studies by International Data Corporation estimate that an enterprise employing 1,000 knowledge workers wastes nearly \$2.5 million per year due to an inability to locate and retrieve information, resulting in lost opportunities and diminished competitiveness [2].

Information retrieval

Information retrieval tools based on traditional keyword search, such as Google or PubMed, are familiar and highly tuned to the task of finding the most relevant documents containing certain keywords, but they have limitations, both on the kinds of query that the user can make and on the kinds of output returned. Input is usually a set of keywords (or MESH terms), and users get documents back instead of answers. Various enhancements to the basic search include query expansion, where your keywords are augmented by related keywords to broaden the search, and categorization, where results are clustered for easier browsing, but the basic proposition is still to input words and return documents.

Shortcomings of keyword search include:

- Although the documents returned can contain relevant information, a laborious and often prohibitively time-consuming manual review is needed to extract that information.
- It is ill-suited to finding relationships, which is often the whole purpose, for example when finding the mechanism of interaction between two proteins.
- It does not cope well with complexity and ambiguity in the source material, for instance where synonyms and homonyms are common.

Information extraction and text mining

Information extraction systems are better suited to finding relationships, using precise queries that look for patterns of concepts and words, for example, 'protein interacts

with protein' to find pairs of proteins that interact [3,4]. The result of a search is structured output (e.g. spreadsheet or database), in which the sought-after information is readily accessible.

Information extraction is based on an understanding of the structure and meaning of the natural language in which documents are written. The ability to search for items occurring within the same sentence or section already gives a greater chance of a relationship between them than if they occur only in the same document. Recognition of grammatical syntax enables nouns and verbs to be distinguished. Entities, such as proteins and diseases, correspond to noun groups, whereas the relationships between them appear as verb groups. This enables specific searches for direct relationships (Figure 1).

The major advantage of information extraction systems is the precision of the queries and the clarity of the output, which can be efficiently reviewed, entered into a database or displayed visually. The term 'text mining' has come to apply to such systems by analogy with conventional data mining systems.

An alternative to text mining using natural language processing is to use statistical co-occurrence methods. Although the occurrence of a particular gene in the same document as a particular disease might not tell us much, the co-occurrence of the same gene and disease in many documents suggests some kind of relationship. Although information extraction and statistical co-occurrence are sometimes seen as competitors, the two can usefully be

combined. Information extraction provides the direction of relationships, and can find information containing negative relationships. Co-occurrence could suggest indirect relationships that cannot easily be found by precise linguistic patterns.

Over the past few years, text mining systems have begun to move from academic research into the market place. A number of factors have driven this move, including market pull due to competitive pressure, the increasing availability of electronic documents, and the falling cost of computing power and storage.

Early text mining systems were aimed at information specialists. They typically require a combination of domain and informatics expertise to configure. Once set up for a particular task they can be run repeatedly over different sets of documents, producing high quality results, but changing to a new task or domain can take weeks and a lot of specialist linguistic expertise. Typical uses include analysis of newsfeeds and constructing pathway databases for subsequent access by end users.

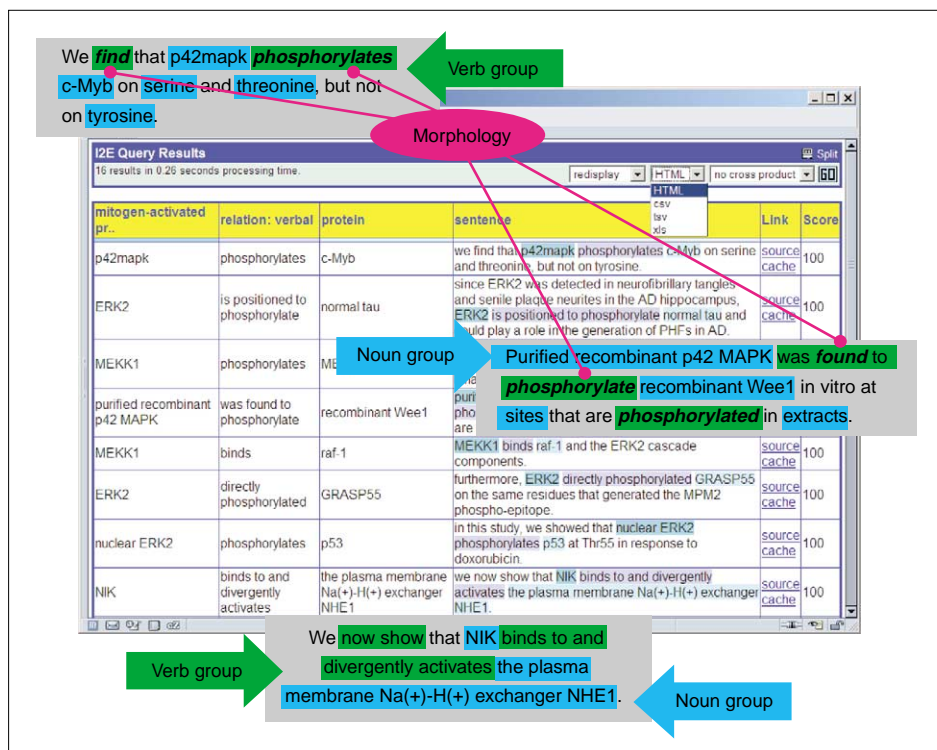


FIGURE 1.

An example output from Linguamatics' interactive text-mining system showing use of linguistics in extracted sentences.

Recent advances: ontologies and usability

The challenge is to make such powerful tools more readily deployable and accessible. Two recent advances have made an important step forward in achieving this. The addition of 'ontologies' enables more-powerful searches, better-defined results, and easier customization to new domains. The development of interactive information extraction enables answers to be found in seconds rather than hours, and for the first time makes information extraction available on demand.

Early text mining systems typically have a few concepts (such as 'company' or 'protein') built in, but the incorporation of ontologies enables semantic search for many thousands of specific concepts (such as 'protein kinase'), and semantic output with standardized naming of concepts (such as the 'LocusLink ID'), which enables results to be sorted and used directly in a database. An ontology is a way of organizing the relationships and entities in a domain. A taxonomy is a simple ontology in which entities or relationships are organised hierarchically so that, for example, 'MEK' is a 'protein kinase', and a search for all protein kinases would also identify all instances of MEK. Similarly, in an ontology of relationships, 'upregulation' might be classified as a direct relationship between proteins. Ontologies can contain synonyms for each entity, again powerful when querying.

Interactive information extraction combines the attractive features of information retrieval and extraction by providing the ease-of-use, scalability and interactivity of information retrieval with the precision and output of information extraction [5]. End users can move naturally and seamlessly from keyword-style searches to the full power and precision of information extraction. Co-occurrence within a document or within a sentence is also possible. Information scientists can now make ad hoc searches and get their results back in seconds.

Case study: mining nuclear receptor cofactors

Validation is a key driver for the adoption of new technologies, and text mining is no exception. Although much good work is being done towards community-accepted validations [6], case studies involving use by real customers can be very persuasive. In a recent case study [7], Linguamatics' I2E software was used with Biowisdom ontologies at AstraZeneca to mine MEDLINE and EMBASE for nuclear receptor cofactors.

Typically, 100 abstracts can be analyzed manually for cofactors in one person-day, whereas in just two person-

hours with I2E, cofactors were extracted with a recall of 91% from around 8000 abstracts, rivalling the quality of hand analysis. This represented an order of magnitude speed-up over existing best practice. Further, the system discovered new cofactors that were not in the standard set.

Summary

In summary, text mining systems:

- Allow more specific queries than keyword search
- Provide structured output for easy review and analysis

Some systems can also:

- Use linguistics to interpret human language and identify relationships
- Use concepts and ontologies to provide domain knowledge
- Mine vast quantities of text interactively

The main benefits of using text mining are:

- Getting to decision points more quickly
- At least 10x speedup over previous methods
- Finding information that wouldn't otherwise be found

References

- 1 Fickett, J. and Hayes, W. (2004) Text mining for drug discovery. *European Pharmaceutical Contractor*. Autumn 2004.
- 2 Feldman, S. and Sherman, C. (2001) The High Cost of not Finding Information. *IDC White Paper*.
- 3 Thomas, J. *et al.* (2000) Automatic extraction of protein interactions from scientific abstracts. In *Proceedings of the Pacific Symposium on Biocomputing 2000*, (Altman, R.B. *et al.* eds), Honolulu, Hawaii. World Scientific Publishing Co., Singapore. pp. 541-552
- 4 Blaschke C, Hirschman L and Valencia A. (2002) *Information extraction in Molecular Biology*, Briefings in Bioinformatics 3: 154-165
- 5 Milward, D. and Thomas, J. (2000) From Information Retrieval to Information Extraction. *Proceedings of the ACL-2000 Workshop on Recent Advances in Natural Language Processing and Information Retrieval*, 8 October 2000, Hong Kong University of Science and Technology, Hong Kong, (Klavans, J and Gonzalo, J., eds). 85-97
- 6 Krallinger, M., Alonso-Allende Erhardt, R. and Valencia, A. (2005) Text Mining approaches in molecular biology and biomedicine. *Drug Discov. Today* 10, 439-445
- 7 Milward, D. *et al.* (2005) Ontology-Based Interactive Information Extraction from Scientific Abstracts. *Proceedings of the BioLINK SIG Text Mining Workshop, ISMB/ECCB 2004. Comparative and Functional Genomics* (in press)

Roger Hale

Linguamatics Ltd,
St John's Innovation Centre,
Cowley Rd,
Cambridge,
CB4 0WS UK
e-mail: roger.hale@linguamatics.com