Knowledge and Reasoning for Medical Question-Answering

Pierre Zweigenbaum
CNRS, LIMSI
Orsay, F-91403 France
pz@limsi.fr

Abstract

Restricted domains such as medicine set a context where question-answering is more likely expected to be associated with knowledge and reasoning (Mollá and Vicedo, 2007; Ferret and Zweigenbaum, 2007). On the one hand, knowledge and reasoning may be more necessary than in open-domain question-answering because of more specific or more difficult questions. On the other hand, it may also be more manageable, since by definition restricted-domain QA should not have to face the same breadth of questions as open-domain QA. It is therefore interesting to study the role of knowledge and reasoning in restricted-domain question-answering systems. We shall do so in the case of the (bio-)medical domain, which has a long tradition of investigating knowledge representation and reasoning and, more generally, artificial intelligence methods (Shortliffe et al., 1975), and which has seen a growing interest in question-answering systems (Zweigenbaum, 2003; Yu et al., 2005; Demner-Fushman and Lin, 2007; Zweigenbaum et al., 2007).

1 Knowledge and Reasoning for Processing Medical Questions

Medical question-answering has to address questions other than the usual factual questions of most QA evaluations. This calls for different question classifications (Ely et al., 2000; Yu et al., 2005), especially to determine whether a given question can be answered using medical knowledge backed with a sufficient level of evidence (Lin and Demner-Fushman, 2005; Kilicoglu et al., 2009). This can also lead to a different representation of questions, for instance using a structured representation such as PICO (Niu et al., 2003; Huang et al., 2006; Demner-Fushman and Lin, 2007) or simple concepts and relations (Lin, 2001; Jacquemart and Zweigenbaum, 2003).

2 Knowledge and Reasoning for Finding Medical Answers

Answers to medical questions should be searched in the most reliable data available. When data exist in structured knowledge bases (*e.g.* a drug compendium), it may be more appropriate to query such knowledge bases directly. Therefore an approach akin to that of Start/Omnibase (Lin and Katz, 2003) may be indicated. When answers are to be found in a collection of documents, as is the case in traditional question-answering systems, a representation of the information contained in these documents can be built, offline (Fleischman et al., 2003; Sang et al., 2005; Delbecque et al., 2005) or dynamically.

In medical QA systems, both document analysis and question analysis nearly always rely on extensive knowledge of domain concepts and relations, *e.g.* as provided by the UMLS knowledge sources (McCray and Nelson, 1995). More than named entities, systems need to detect mentions of concepts (Aronson, 2001) and their relations (Rindflesch et al., 2005). Besides, taking into account the structure of documents such as scientific articles or encyclopedia entries may help focus on more relevant sections (Niu and Hirst, 2004; Sang et al., 2005). Finally, answers to complex medical questions often need to span more than one sentence. Extractive summarization is performed both from single documents (Demner-Fushman and Lin, 2007) and from multiple documents (Fiszman et al., 2008).

References

Alan R. Aronson. 2001. Effective mapping of biomedical text to the UMLS Metathesaurus: The MetaMap program. *Journal of the American Medical Informatics Association*, 8(suppl):17–21.

Thierry Delbecque, Pierre Jacquemart, and Pierre Zweigenbaum. 2005. Indexing UMLS semantic types for medical question-answering. In Rolf Engelbrecht, Antoine Geissbuhler, Christian Lovis, and G. Mihalas, editors, *Proceedings Medical Informatics Europe*, volume 116 of *Studies in Health Technology and Informatics*, pages 805–810, Amsterdam. IOS Press.

Dina Demner-Fushman and Jimmy Lin. 2007. Answering clinical questions with knowledge-based and statistical techniques. *Computational Linguistics*, 33(1):63–103.

John W. Ely, Jerome A. Osheroff, Paul N. Gorman, Mark H. Ebell, M. Lee Chambliss, Eric A. Pifer, and P. Zoe Stavri. 2000. A taxonomy of generic clinical questions: classification study. *BMJ*, 321:429–432. Available at http://bmj.com/cgi/content/full/321/7258/429.

Olivier Ferret and Pierre Zweigenbaum. 2007. Représentation des connaissances pour les systèmes de question-réponse. In Brigitte Grau and Jean-Pierre Chevallet, editors, *La recherche d'informations précises : traitement automatique de la langue, apprentissage et connaissances pour les systèmes de question-réponse*, chapter 4, pages 133–169. Hermès-Lavoisier, Paris.

Marcelo Fiszman, Dina Demner-Fushman, Halil Kilicoglu, and Thomas C Rindflesch. 2008. Automatic summarization of MEDLINE citations for evidence-based medical treatment: A topic-oriented evaluation. *J Biomed Inform*, November.

Michael Fleischman, Abdessamad Echihabi, and Eduard Hovy. 2003. Offline strategies for online question answering: Answering questions before they are asked. In *Proceedings of the ACL Conference*, pages 1–7, Sapporo, Japan.

Xiaoli Huang, Jimmy Lin, and Dina Demner-Fushman. 2006. Evaluation of PICO as a knowledge representation for clinical questions. In *AMIA Annu Symp Proc*, page 359–63.

Pierre Jacquemart and Pierre Zweigenbaum. 2003. Towards a medical question-answering system: a feasibility study. In Robert Baud, Marius Fieschi, Pierre Le Beux, and Patrick Ruch, editors, *Proceedings Medical Informatics Europe*, volume 95 of *Studies in Health Technology and Informatics*, pages 463–468, Amsterdam. IOS Press.

Halil Kilicoglu, Dina Demner-Fushman, Thomas C Rindflesch, Nancy L Wilczynski, and R Brian Haynes. 2009. Towards automatic recognition of scientifically rigorous clinical research evidence. *J Am Med Inform Assoc*, 16(1):25–31.

Jimmy Lin and Dina Demner-Fushman. 2005. "Bag of words" is not enough for strength of evidence classification. In *AMIA Annu Symp Proc*, page 1031.

Jimmy Lin and Boris Katz. 2003. Question answering techniques for the World Wide Web. In *Tutorial at EACL 2003*, Budapest. ACL.

Jimmy Lin. 2001. Indexing and retrieving natural language using ternary expressions. Master's thesis, Massachusetts Institute of Technology.

Alexa T. McCray and Stuart J. Nelson. 1995. The semantics of the UMLS knowledge sources. *Methods of Information in Medicine*, 34(1/2).

Diego Mollá and José Luis Vicedo. 2007. Question answering in restricted domains: An overview. *Computational Linguistics*, 33(1):41–61.

Yun Niu and Graeme Hirst. 2004. Analysis of semantic classes in medical text for question answering. In *Proceedings ACL 2004 Workshop on Question Answering in Restricted Domains*. ACL.

Yun Niu, Graeme Hirst, Gregory McArthur, and Patricia Rodriguez-Gianolli. 2003. Answering clinical questions with role identification. In *ACL Workshop Natural Language Processing in Biomedicine*, pages 73–80. ACL.

Thomas C. Rindflesch, Marcelo Fiszman, and B. Libbus. 2005. Semantic interpretation for the biomedical literature. In H Chen, S Fuller, WR Hersh, and C Friedman, editors, *Medical informatics: Advances in knowledge management and data mining in biomedicine*, pages 399–422, Berlin / Heidelberg. Springer.

Erik Tjong Kim Sang, Gosse Bouma, and Maarten de Rijke. 2005. Developing offline strategies for answering medical questions. In *Proceedings AAAI-05 workshop on Question Answering in restricted domains*, pages 41–45. AAAI.

E H Shortliffe, R Davis, S G Axline, B G Buchanan, C C Green, and S N Cohen. 1975. Computer-based consultations in clinical therapeutics: explanation and rule acquisition capabilities of the MYCIN system. *Comput Biomed Res*, 8(4):303–20, August.

Hong Yu, Carl Sable, and Hai Ran Zhu. 2005. Classifying medical questions based on an evidence taxonomy. In *Proceedings AAAI 2005 Workshop on Question Answering in Restricted Domains*. AAAI.

Pierre Zweigenbaum, Dina Demner-Fushman, Hong Yu, and K. Bretonnel Cohen. 2007. Frontiers of biomedical text mining: current progress. *Briefings in Bioinformatics*, 8:358–375, October. doi:10.1093/bib/bbm045.

Pierre Zweigenbaum. 2003. Question answering in biomedicine. In Maarten de Rijke and Bonnie Webber, editors, *Proceedings Workshop on Natural Language Processing for Question Answering, EACL 2003*, pages 1–4, Budapest. ACL. Keynote speech.