

Preface

This volume contains the papers accepted for presentation at the workshop on Question Answering in Restricted Domains, which is part of the 20th National Conference on Artificial Intelligence AAAI-05, held on July 9-13, 2005 in Pittsburgh, Pennsylvania.

There are many document sets in restricted domains that are potentially valuable as a source for question answering systems. For example, the documentation pages of Unix and Linux systems would make an ideal corpus for QA systems targeted at users that want to know how to use these operating systems. There is a wealth of information in other technical documentation such as software manuals, car maintenance manuals, and encyclopediae of specific areas such as medicine. Users interested in these specific areas would benefit from QA systems targeted to their areas of interest.

Restricted domains typically have limited data available and therefore conventional open-domain question answering techniques based on data redundancy can simply not be applied in an effective way. The scarcity of data available seems to prompt for a more targeted, NLP-intensive approach to QA. The use of additional corpora such as the WWW raises a number of interesting questions. For instance, will these corpora help or obstruct the proper functioning of NLP-intensive approach to QA? And, how do we find good pockets of information that are appropriate to the chosen domains?

On the other hand, restricted domains (e.g. law, medicine) have specific stylistic conventions that need to be handled. Often these domains use terminology that is not stored in conventional lexica. Consequently NLP approaches devised for open-domain systems may choke on these specific domains, thus raising the question of how portable these systems can be.

In this workshop we aim at answering some of the following questions:

- Are open-domain question answering techniques appropriate for QA in restricted domains?
- What types of question-answering techniques are best for what types of restricted domains?
- Is it easy/possible/worthwhile to develop domain-independent QA systems for restricted domains? What would be the cost of porting a QA system to other domains?
- Are restricted domains more suitable than open domains to drive research in NLP?
- Is evaluation of restricted-domain QA systems different than that of open-domain QA systems?

This workshop is a follow-up of the ACL2004 workshop on Question Answering in Restricted Domains held in Barcelona, Spain, on July 2005.

We are very grateful to our programme committee for the effort they put in reviewing the full papers. We are also grateful to the AAAI-05 conference organisers on whom we could rely for the local organization.

We hope that you find the contents of these proceedings useful.

Diego Mollá & José Luis Vicedo (editors)
May 2005