Knowledge-Driven Text Interpretation and Question-Answering: Some Current Activities at Boeing Mathematics and Computing Technology

Peter Clark, Phil Harrison, Tom Jenkins, John Thompson, Rick Wojcik
Boeing Mathematics and Computing Technology, Seattle, WA

At Boeing, we are currently developing methods for knowledge-driven text interpretation and question-answering, based on matching analyzed input text against strong, background expectations from a knowledge-base (Clark *et al*. 2002). Our goal is to be able to answer questions about text which goes beyond facts explicitly stated in the text. For example, the statement (1) "China launched a meterological satellite into orbit Wednesday" suggests that (among other things) there was a rocket launch; China probably owns the satellite; the satellite is for monitoring weather; the orbit is around Earth; etc., although none of these facts are explicit in the text.

Our approach is based on (a) building representations of prototypical scenarios, e.g., launching a satellite, in a large knowledge-base, (b) extracting fragments of information from the text of interest using NLP techniques, e.g., subject-verb-object relations, and then (c) identifying which of the scenario(s) best match the processed input text. The system thus has strong expectations about what sort of things might happen in the world, and we seek to exploit those expectations for interpreting text, including resolving the many types of ambiguity which can arise. After matching, the background representations provide a source of additional facts about the situation described in the text, including the large number of facts assumed but not explicitly mentioned. This additional knowledge can be used for supporting inference, question-answering, and advanced search of a text corpus. For example, a query for rocket launch events would identify (1) as relevant, even though rockets are not explicitly mentioned in that text.

The representations themselves are constructed by identifying the key "participants" (both objects and events) in the scenario, and then creating a graph of relationships that normally exist between those participants. These graphical representations are compositional (Clark & Porter 1997) in two important ways: First, through inheritance, a representation can be combined with representations of its generalizations (e.g., representations of "launching a satellite" and "placing something in position" can be combined). Second, different viewpoints/aspects of a concept such as launching a satellite are encoded as separate representational structures (e.g., the sequence of events; the temporal information; the spatial information; goal-oriented information). During text interpretation, only those representation(s) of aspects/views which the text itself refers to will be composed into the structure matched with the text.

To match the representations with the NLP-processed text fragments, the system searches for matches between objects in the representations and objects mentioned in the text; and relationships in the representations and relationships mentioned in the text. Much of the ambiguity in the source text, e.g., the mapping from syntactic to semantic roles, is resolved at this point (rather than earlier), as the knowledge base structures define allowable/prototypical ways in which various concepts and relationships can be combined.

We are also exploring ways to efficiently build the KB of scenarios in the first place, given that there are probably thousands of scenarios required for a typical application, and are looking at semi-automated ways of doing this. We are currently exploring three avenues: use of NLP-processed glossary definitions as a starting point; seeding of the knowledge-base using a Schubert-style generate-and-filter approach (Schubert 2002) (namely, produce simple, text-derived generalizations for a human to review, e.g., from "China launched a satellite Wednesday" the generalization "countries launch satellites" can be automatically derived); and manual encoding.

This overall approach treats text interpretation, and subsequent question-answering about that text, fundamentally as a modeling activity, in which text suggests scenario models to use, and those models suggest ways of interpreting text. Although this approach is challenging for a number reasons, it offers significant potential for allowing question-answering to go beyond facts explicitly stated in the various text sources used.

References

Clark, P., and Porter, B. 1997. Building concept representations from reusable components. In *AAAI-97*, 369–376.

Clark, P.; Duncan, L.; Holmback, H.; Jenkins, T.; and Thompson, J. 2002. A knowledge-rich approach to understanding text about aircraft systems. In *Proc. HLT-02*, 78–83.

Schubert, L. 2002. Can we derive general world knowledge from texts? In *Proc. HLT-02*, 84–87.