Technical Perspective: Natural Language Explanations for Query Results

Zachary G. Ives
University of Pennsylvania
zives@cis.upenn.edu

Motivated by conversational agents such as Siri, Cortana, the Google Assistant, and Alexa — there has been a surge of interest in spoken as well as textual natural language interfaces. To this point, such systems have relied on innovations in speech recognition (such as recurrent neural networks, LSTMs, and so on) and in specially encoding specific question-answering strategies via "skills." A "natural" question for the SIGMOD community is how to best connect natural language interfaces systems to DBMSs, ideally in a way that generalizes to any database schema or instance.

In fact, the problem of providing a natural language interface to a database system (i.e., mapping from a spoken or textual question to a structured query) dates back at least to the 1980s [4]. Such efforts had middling success due to issues of accuracy, so the problems were later revisited in the 2000's with an eye towards restricting the space of options in order to improve precision [6]. Nonetheless, such systems did not gain much traction, again due to the challenges of ensuring accuracy for a given database when the user might ask an ambiguous question.

Recent work by Li and Jagadish [5], called `NaLIR`, proposed an *interactive communicator* within the query system, which presents to the user a *query tree* explaining what the system was going to do — such that the user could correct any mistakes. This was helpful in improving reliability, but it required that the user understand tree structured representations of queries.

In "Natural Language Explanations for Query Results," Deutch and his co-authors suggest that a more effective means of helping the user understand and correct results might be through *provenance* information — i.e., giving an explanation for each answer of how and why it exists. Their approach adapts the `NaLIR` system and nicely leverages the recent body of work on provenance semirings [3, 2, 1]. The provenance semiring model has an important property that equivalent query plans (as produced by a query optimizer) will have equivalent provenance expressions.

The innovations in this paper are in three areas. First, the authors use *the structure of the natural language query itself* (and the mappings to structured queries, and then later, from queries to provenance) to present the provenance in a form that matches the natural language query — and thus the user's expectations. Second, they reduce the size (and repetition) of the provenance via factoring. Finally, they incorporate aggregate results (e.g., counts) in place of certain details.

The paper does a great job of clearly identifying and articulating what makes the provenance problem different for natural language query systems, and presenting elegant technical solutions to these new challenges.

1. REFERENCES

[1] Yael Amsterdamer, Daniel Deutch, and Val Tannen. Provenance for aggregate queries. In *PODS*, pages 153–164, 2011.

[2] Todd J. Green, Grigoris Karvounarakis, Zachary G. Ives, and Val Tannen. Update exchange with mappings and provenance. In *VLDB*, 2007. Amended version available as Univ. of Pennsylvania report MS-CIS-07-26.

[3] Todd J. Green, Grigoris Karvounarakis, and Val Tannen. Provenance semirings. In *PODS*, 2007.

[4] S. Jerrold Kaplan. Designing a portable natural language database query system. *ACM Trans. Database Syst.*, 9:1–19, March 1984. Available from http://doi.acm.org/10.1145/348.318584.

[5] Fei Li and HV Jagadish. Constructing an interactive natural language interface for relational databases. *Proceedings of the VLDB Endowment*, 8(1):73–84, 2014.

[6] Ana-Maria Popescu, Oren Etzioni, and Henry Kautz. Towards a theory of natural language interfaces to databases. In *IUI*, pages 327–327, 2003.