Generating Surplus Content in a Q/A-Setting

Sebastian Reuße and **Ralf Klabunde**
{reusse,klabunde}@linguistics.rub.de
Sprachwissenschaftliches Institut
Ruhr-Universität Bochum

Jon Stevens and **Anton Benz**
{benz,stevens}@zas.gwz-berlin.de
ZAS Berlin

Abstract

We present a framework for question-answering systems which employ pragmatic reasoning based on domain-level goals inferred from users' queries. Inferred user goals are used to anticipate how generated replies are interpreted. This allows the system to predict when surplus content can be felicitously included for the sake of discourse efficiency, and in cases where surplus content is strictly required to avoid supplying misleading answers.

1 Background

The system presented here is part of the *PragSales* project, which set out to develop a Q/A-system architecture for imperfectly cooperative domains. This paper illustrates an application to the real-estate sales domain, where prospective tenants looking for an apartment to let interact with an automated real-estate agent. Interaction proceeds based on customer's queries, posed in the form of Y/N-questions for attributes of the flat being discussed. Within this setting, the system needs to 1. ensure the *felicity* of generated utterances, and 2. keep interactions *efficient* and non-repetitive. We propose a theoretically grounded framework which is able to address both of these concerns from a unified perspective.

2 Architecture

Domain Model. A model of a real-estate sales domain was implemented using an adapted version of *PyKE*, Horn logic theorem prover. A range of predicates were defined to represent a variety of flat-intrinsic attributes, such as size and pricing, as well as extrinsic attributes pertaining to a flat's surroundings, such as distances to public transit stops, schools and grocery stores.

User Model. We posit that users issue queries in order to ascertain whether a flat fulfills one or more of their underlying *goals*, i. e., sets of preconditions on desirable flats. We formalize these as sets of Horn-logic implications, where a goal term is fulfilled when one or more conjunctions of attribute predicates uphold. E. g., we might define that some customers are seeking a place which allows them to enjoy some time in the sun, and that such a goal may be realized in either of two ways:

```
sunTan ← flat(e) ∧ hasGarden(e)
sunTan ← flat(e) ∧ hasBalcony(e)
```

Goal Inference. A user goal underlying each particular query is inferred by performing a Bayesian update on a flat prior over user goals g, given an observed query term a. Each query term is deemed an a priori equally likely option of finding out about its superordinate goal.

Dialogue Management As opposed to a canonical Q/A-setting, our application setting allows for queries to be posed in an iterative fashion, simulating a dialogical interaction between a sales agent and a customer. To this end, we implemented a minimal dialogue manager using the *PyTrindiKit* toolkit.

3 Anticipating Pragmatic Inference

Whereas canonical Q/A-systems create a relatively constrained interpretation context, the continuous fashion in which users interact with our application induces a context that makes it plausible for users to assume the sales agent to be aware that their overt queries are motivated by implicit goals. Consequently, replies trigger a pragmatically enriched interpretation of system responses (van Rooy, 2003). Specifically, a game-theoretic analysis of our domain predicts two interrelated types of pragmatic inference (Stevens et al., 2014), both of which a dialogue move engine should take into account when generating replies:

Appears in *JerSem: The 20th Workshop on the Semantics and Pragmatics of Dialogue,* edited by Julie Hunter, Mandy Simons and Matthew Stone. New Brunswick, NJ USA, July 16-18, 2016, pages 116-117.

1. Users assume the system to supply the maximally beneficial reply, so that unrealized, better answers are assumed to be negated. This allows No-responses to be addressed indirectly, or in the form of *No, but...* replies, which acts as a license to supply unrequested, but beneficial information.
2. As a corrolary, plain Yes/No-responses can trigger misleading implicatures which an automated system may want to avoid.

In the following, we briefly outline how both of these concerns can be addressed within our framework.

Increasing Dialogue Efficiency. Based on our supposition that what is relevant for the purposes of the conversation is determined by implicit user goals, rather than by overt query terms, representing these goals within the scope of the system allows generating *indirect* response types which give *unrequested, but goal-relevant information*, as in (1-b), or which leave out an overt *No*, thereby increasing dialogue efficiency, as shown in evaluations by Stevens et al. (2015).

(1) Q.Does the flat feature a balcony?
 a. Well, it features a garden.
 $\rightsquigarrow$ It does not feature a balcony.
 b. No, but there's a garden.

Since, given the user's underlying goal, a *garden* is a valid alternative to a *balcony*, our system can predict that (1-a/b) are coherent responses to the original query, which may be preferred for the sake of discourse efficiency.

Blocking Undesirable Inferences. In some contexts, the assumption that answer interpretation is strengthened based on common knowledge about the user's underlying goals allows us to predict answers which, while factually true, implicate false propositions. In (2-a), a literal *No* response negates, by implicature, not only an overt query term, but all manifestations of the underlying goal, since a cooperative or self-interested seller would have included any relevant, true alternatives in his response.

(2) Q.Does the flat feature a balcony?
 a. No. $\rightsquigarrow$ It's not good for getting a tan.
 b. No, but if you're looking to get a tan, there's a garden.

In addition to literal responses, our system generates contrastive responses, such as (2-c), which block the undesirable inference. We define the underlying contrastive message operator $but(m_1, m_2)$ to be formally licensed when: 1. both m_1 and m_2 are sufficient preconditions of some user goal D, 2. the user issues a query for m_1, 3. m_1 fails to uphold, and 4. m_2 upholds.

Vice versa, a literal *Yes* response triggers an expectation for the superordinate user goal to be fulfilled, as in (3-a). This expectation would mislead the user when this is not actually the case.

(3) Q.Does the flat feature a balcony.
 a. Yes. $\rightsquigarrow$ The balcony is a place to work up a tan.
 b. Yes, though it faces north.

In cases such as this, our system is able to generate concessive replies such as (3-b), using a concession operator $although(m_1, m_2)$ which is licensed when: 1. both m_1 and m_2 are necessary preconditions of some user decision problem D, 2. the user issues a query for m_1, 3. m_1 upholds, and 4. m_2 fails to uphold.

4 Discussion

We believe that the way in which our framework allows formalizing the use conditions of a variety of linguistic means, encompassing both indirect and augmented answers, as well as contrastive resp. concessive replies, shows that the theory of van Rooy (2003) forms a workable base on top of which pragmatic phenomena can be handled within the scope of a Q/A-setting.

References

Stevens, J. S., Benz, A., Reuße, S., Klabunde, R., & Raithel, L. (2015). Pragmatic query answering: results from a quantitative evaluation. In C. Biemann, S. Handschuh, A. Freitas, F. Meziane, & E. Métais (Eds.), *Natural language processing and information systems* (pp. 110–123). Lecture Notes in Computer Science. Springer.

Stevens, J. S., Benz, A., Reuße, S., Laarmann-Quante, R., & Klabunde, R. (2014). Indirect answers as potential solutions to decision problems. In *Proceedings of the 18th workshop on the semantics and pragmatics of dialogue* (pp. 145–153).

van Rooy, R. (2003). Questioning to resolve decision problems. *Linguistics and Philosophy*, *26*(6), 727–763.