An Empirical View on IQA Follow-up Questions

Manuel Kirschner and Raffaella Bernardi
KRDB Center, Faculty of Computer Science
Free University of Bozen-Bolzano, Italy
{kirschner,bernardi}@inf.unibz.it

Abstract

In a realistic Interactive Question Answering (IQA) situation, one third of the users pose follow-up questions, i.e., go beyond a single question per dialogue. We identify two different perspectives according to which these follow-ups can be described: informational transitions and context dependency. By understanding exactly how informational transitions occur in IQA dialogues, we propose a method to guarantee that focus tree based IQA systems provide wide coverage of follow-up questions that trigger the respective set of informational transitions.

1 Introduction

This is an empirical study of follow-up questions in Interactive Question Answering (IQA) dialogues that we collected through a previous Wizard-of-Oz study. In this paper, we show that user follow-up questions are an interesting phenomenon because they occur relatively frequently in IQA dialogues, and are potentially difficult for an IQA system to understand. We will look at them from two different perspectives: (i) which informational transition can be identified between the follow-up and the dialogue context, and (ii), how some of the follow-up questions are context-dependent in that they can only be properly understood in combination with information from the dialogue context. In understanding (i), we try to find patterns and regularities in our data that enable us to predict the topics that users of an IQA system will ask about next. This knowledge

will help in improving an IQA system, since we can ensure that the system will be prepared to answer the specific follow-up questions that we predicted for a specific situation in an IQA dialogue. As for (ii), on the other hand, we need to understand also *how* users typically pose follow-up questions: as we show in this paper, many follow-up questions are context-dependent, and need to be combined with information from the previous dialogue in order to be understandable for the IQA system. After analyzing follow-up questions from these two perspectives, we propose a new way of processing a certain class of follow-ups in an actual IQA system on the library domain.

2 Statistics of Follow-ups in IQA Dialogues

We conducted a Wizard-of-Oz experiment where the participants were free to choose their question topic, and the way in which to interact with the system. In this experiment, we collected 63 user-librarian dialogues by letting spontaneous visitors of the library web-site interact with what was announced as a new IQA system, but was in reality a web-based instant-messaging-like interface (Kirschner, 2006).

From the total of 192 user utterances in our corpus (spread across 63 dialogues and 166 user turns), we identified 35 that are both follow-up initiatives (i.e., from the set of 90 questions or assertions that are not from the very first user turn in each dialogue) and that are also about a topic from the library information domain, or some task related to this domain.

While from the set of 90 follow-up initiatives the proportion of user utterances we marked as off-topic is high (56, versus the 35 domain/task-related

43

ones), we can assume that they will not pose a major problem to the IQA system. We conjecture that in most cases these utterances can be easily ignored by the natural-language understanding module, which should robustly spot only questions and assertions about task-related topics. Moreover, the analysis shows that many users do take the opportunity that IQA dialogue offers and do ask follow-up questions. Even more, the latter actually contains some of the most important parts of the dialogues (besides the first user question in each single dialogue), and the most interesting and difficult user utterances for an IQA system to process.

3 Informational Transitions

In some of the literature, the term *thematic relatedness* is used to describe transitions between utterances. We assume this is just a matter of different terminology; for the sake of clarity, we define that those follow-up questions that trigger some informational transition at the same time define the set of thematically related follow-ups. Also, note that throughout this paper, we use the term follow-up (question) to denote *any* user question that is not the very first question in a given IQA dialogue; thus, it does not imply that the follow-up be in some specific way related to the previous dialogue.

The general goal of all the approaches to be presented in this section is to explore specific relations holding between two discourse segments or dialogue turns. This is of primary interest in the context of building an IQA application, since by understanding how the conversation topic evolves via user follow-up questions, we can improve the way the system will understand and answer these follow-ups. In our empirical approach, we want to analyze how informational transitions are used in real IQA dialogues. Thus, a preliminary goal is to find a method of reliably identifying these phenomena in our dialogue data. In what follows, we describe three previous approaches to this problem, focusing on their generalizability and practical applicability for identifying informational transitions in data. At the end of this section, we will then propose a somewhat restricted (but on the other hand more practical and concise) method of identifying (a subset of) informational transitions.

In the context of planning coherent discourse in a natural language generation system, (McCoy and Cheng, 1991) gives a comprehensive account of informational transitions (there called focus shifts). For each node type, they list certain focus shift candidates, i.e., the items that are likely to come into focus in a coherent discourse (cf. Table 1). While their list of focus shift targets for the different node types is comprehensive, this is at the same time a major problem when it comes to a practical implementation: it is not at all clear how to (algorithmically) determine the correct node types, and thus the viable candidate targets for informational transitions.

In a related approach that targets IQA dialogues rather than single-speaker discourse, (Chai and Jin, 2004) define informational transitions between subsequent user questions in IQA dialogues in terms of the question "topic". The topic is either of type `entity` or `activity` and closely resembles the object and activity node types given in Table 1. While the informational state is now described in terms of only two types of elements (entity/object and activity/action) instead of the five postulated by (McCoy and Cheng, 1991), the rich set of discourse roles that these elements can introduce would still render an automatic construction of a representation of the informational state extremely difficult.

A further description of informational transitions in IQA dialogues is given in (Bertomeu et al., 2006). Unlike the two previously mentioned approaches, this work considers also system responses as possible sources for informational transitions. In fact, the authors identify specific thematic relations that may hold between a user follow-up question and the immediately previous user question, some previous user question, the immediately previous system answer or some previous system answer. Interestingly, this approach is based entirely on questions and answers corresponding to (sets of) entities that can be retrieved from a database. Thus, informational transitions are defined here in terms of the extensions of entities that are being referred to in thematically related turns of the dialogue, and in terms of which properties of these entities are being referred to. However, the transitions also lack the generality of the previously introduced approaches, since they are only useful for analyzing similar kinds of (natural language database query) dialogues that contain

Node type	Focus shift targets
object	Attributes of the object, actions the object plays a prominent role in (e.g., is actor of)
action	Actor, object, etc., of the action – any participant (Fillmore) role; purpose (goal) of action, next action in some sequence, subactions, specializations of the action
attribute	objects which have the attribute, more specific attribute
setting	objects involved in the setting; actions which typically occur in this setting
event	actions which can be grouped together into the event

Table 1: Informational transition targets for different focus nodes (from (McCoy and Cheng, 1991, p. 112))

only rather constrained types of questions and answers.

We will base our work on the observations on these three works.

3.1 Coverage vs. Conciseness: Searching for a Definition of Thematic Relatedness

From (McCoy and Cheng, 1991), we adopt the general idea of introducing candidate focus shift targets that represent coherent continuations of the discourse (or in our case, dialogue). To avoid the difficulty of choosing between up to five different node types that could represent the current focus of attention, we restrict ourselves to just action-type nodes. This is advantageous in two ways. On the one hand, actions correspond to verbs, which are inherently connected to some argument structure defining the verb's semantic roles. By querying available lexical resources like PropBank (Palmer et al., 2005), we can retrieve the verb's arguments. The corresponding semantic roles of the verb yield possible topics of follow-up questions. Thus, we can take advantage of existing lexical resources to automatically find focus nodes that represent follow-up questions involving any of the semantic roles of the verb. On the other hand, we conjecture that actions/verbs form a suitable and robust basis for describing the (informational) meaning of utterances in IQA, since most user utterances include a predicate (or an implicit reference to some predicate in the dialogue history), and syntactic parsers can be used to extract the main verbs of sentences. Taking the main verb plus any arguments to represent the core meaning of user questions seems to be an interesting possibility for automatically detecting certain informational transitions.

Once we adopt the action-based paradigm for focus nodes, we can instantiate two of the informational transition relations proposed by (Chai and Jin, 2004). In the following, we define our own set of informational transitions, starting from the definitions in (Chai and Jin, 2004), but addressing their shortcomings mentioned previously.

First of all, we use verbs and their semantic roles, plus a focus marker, as the only elements needed for representing the informational perspective, and for defining our transition types. This allows us to replace the somewhat unclear terms from the original definitions in (Chai and Jin, 2004) with clearly defined ones: verbs and arguments, as defined in PropBank. Secondly, we parametrize the transitions with respect to their origin: last user question (U_{-1}), or last system response (S_{-1}).

We restrict ourselves to transitions where the main verb either stays the same, or the follow-up question contains a synonymous verb, or no verb at all (to account for fragmentary questions). We now define the resulting three types of informational transitions.

1. TOPIC EXTENSION (U_{-1}):
Example: U1: "Can every student use inter-library loan?" – U2: "Even high-school students?"

   1. Either no verb exists in the follow-up question, or the main verb of the follow-up question is synonymous to the main verb in the last user question.
   2. Either the roles of the verb are filled differently by the follow-up (CONSTRAINT REFINEMENT), or different roles of the verb are filled by the follow-up (PARTICIPANT SHIFT).
   3. The question focus (the expected answer type) stays the same.

2. TOPIC EXPLORATION (U_{-1}):
Example: U1: "Can every student use inter-library

loan?" – U2: "How?"

1. Either no verb exists in the follow-up question, or the main verb of the follow-up question is synonymous to the main verb in the last user question.

2. The question's focus (the expected answer type) changes.

3. TOPIC EXPLORATION (S_{-1}):
Example: U1: "Can high-school students use the library?" – S1: "Yes, if they got a library card." – U2: "So how do I get it?"

1. The main verb of the follow-up question is synonymous to SOME main verb in the system response.

2. Either the roles of the verb are filled differently by the follow-up (CONSTRAINT REFINEMENT), or different roles of the verb are filled by the follow-up (PARTICIPANT SHIFT).

4 Context-dependent User Follow-up Initiatives

Besides studying the thematic relatedness of follow-up questions with respect to previous dialogues, context-dependency yields a new perspective under which to analyze follow-ups. We call a follow-up question context-dependent if it requires any information from the dialogue context in order to be fully understandable. Although this might not generally hold for more complex types of dialogue, we found that in our corpus of IQA dialogues, every user follow-up initiative that we consider context-dependent according to the above definition actually exhibits some discourse phenomena .

In a nutshell, our study shows that (1) discourse phenomena can be resolved without global context (or dialogue history), and (2) the last system response S_{-1} was often the location of the antecedents of discourse phenomena.

5 Conclusions

We showed that in a realistic IQA situation, one third of the users pose follow-up questions, i.e., go beyond a single question per dialogue. We have then introduced two different perspectives according to which the follow-ups can be described and further

categorized: informational transitions and context dependency. For the latter, we have looked at discourse phenomena, and studied how these appear in IQA dialogue data. As for informational transitions, we showed that a rather concise definition is possible if we considerably reduce the scope of the problem, thus limiting the types of informational transitions we deal with. A concise definition is required for letting an IQA system predict informational transitions automatically, given some local dialogue history. The empirical evaluation of this definition shows that it fails in predicting any larger set of specific follow-up initiatives. The problem of concisely identifying informational transitions in IQA seems to be a more complex one, as the variety of different thematic relations found in our corpus alone suggests. While in future work we will try to fine-tune our definitions to further extend the modelling of follow-up initiatives in IQA, on the practical side we have started to extend our baseline IQA system for the library information domain by implementing the three proposed definitions of informational transitions, since they provide a principled way of extending the system.

References

N. Bertomeu, H. Uszkoreit, A. Frank, H.-U. Krieger, and B. Jörg. 2006. Contextual phenomena and thematic relations in database qa dialogues: results from a wizard-of-oz experiment. In *Proc. of the Interactive Question Answering Workshop at HLT-NAACL 2006*, pp. 1–8, New York, NY.

J. Y. Chai and R. Jin. 2004. Discourse structure for context question answering. In *Proc. of the HLT-NAACL 2004 Workshop on Pragmatics in Question Answering*, Boston, MA.

M. Kirschner. 2006. Building a multi-lingual interactive question-answering system for the library domain. In *Proc. of SemDial'06*, Potsdam, Germany.

K. F. McCoy and J. Cheng. 1991. Focus of attention: Constraining what can be said next. In C. L. Paris, W. R. Swartout, and W. C. Mann, eds, *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, pp. 103–124. Kluwer Academic Publishers, Norwell, MA.

M. Palmer, D. Gildea, and P. Kingsbury. 2005. The proposition bank: An annotated corpus of semantic roles. *Computational Linguistics*, 31(1):71–106.