

A Wizard-of-Oz environment to study Referring Expression Generation in a Situated Spoken Dialogue Task

Srinivasan Janarthanam

School of Informatics
University of Edinburgh
Edinburgh EH8 9AB
s.janarthanam@ed.ac.uk

Oliver Lemon

School of Informatics
University of Edinburgh
Edinburgh EH8 9AB
olemon@inf.ed.ac.uk

Abstract

We present a Wizard-of-Oz environment for data collection on Referring Expression Generation (REG) in a real situated spoken dialogue task. The collected data will be used to build user simulation models for reinforcement learning of referring expression generation strategies.

1 Introduction

In this paper, we present a Wizard-of-Oz (WoZ) environment for data collection in a real situated spoken dialogue task for referring expression generation (REG). Our primary objective is to study how participants (hereafter called users) with different domain knowledge and expertise interpret and resolve different types of referring expressions (RE) in a situated dialogue context. We also study the effect of the system's lexical alignment due to priming (Pickering and Garrod, 2004) by the user's choice of REs. The users follow instructions from an implemented dialogue manager and realiser to perform a technical but realistic task – setting up a home Internet connection. The dialogue system's utterances are manipulated to contain different types of REs - descriptive, technical, tutorial or lexically aligned REs, to refer to various domain objects in the task. The users' responses to different REs are then logged and studied.

(Janarthanam and Lemon, 2009) presented a framework for reinforcement learning of optimal natural language generation strategies to choose appropriate REs to users with different domain knowledge expertise. For this, we need user simulations with different domain knowledge profiles that are sensitive to the system's choice of REs. A WoZ environment is an ideal tool for data collection to build data-driven user simulations. However, our study requires a novel WoZ environment.

In section 2, we present prior related work. Section 3 describes the task performed by partici-

pants. In section 4, we describe the WoZ environment in detail. Section 5 describes the data collected in this experiment and section 6 presents some preliminary results from pilot studies.

2 Related Work

(Whittaker et al., 2002) present a WoZ environment to collect data concerning dialogue strategies for presenting restaurant information to users. This study collects data on strategies used by users and human expert wizards to obtain and present information respectively. (van Deemter et al., 2006) present methods to collect data (the TUNA corpus) for REG using artificially constructed pictures of furniture and photographs of real people. (Arts, 2004) presents a study choosing between technical and descriptive expressions for instruction *writing*.

In contrast to the above studies, our study is novel in that it collects data from users having different levels of expertise in a real situated task domain, and for spontaneous spoken dialogue. Our focus is on choosing between technical, descriptive, tutorial, and lexically aligned expressions rather than selecting different attributes for generating descriptions.

3 The Domain Task

In this experiment, the task for each user is to listen to and follow the instructions from the WoZ system and set up their home broadband Internet connection. We provide the users with a home-like environment with a desktop computer, phone socket and a Livebox package from Orange containing cables and components such as the modem, broadband filters and a power adaptor. During the experiment, they set up the Internet connection by connecting these components to each other. Prior to the task, the users are informed that they are interacting with a spoken dialogue system

that will give them instructions to set up the connection. However, their utterances are intercepted by a human wizard. The users are requested to have a conversation as if they were talking to a human operator, asking for clarifications if they are confused or fail to understand the system's utterances. The system's utterances are converted automatically to speech using the Cereproc Speech Synthesiser and played back to the user. The user follows the instructions and assembles the components. The setup is examined by the wizard at the end of the experiment to measure the percentage of task success. The user also fills in questionnaires prior to and after the task answering questions on his background, quality of the system during the task and the knowledge gained during the task.

4 The Wizard-of-Oz environment

The Wizard-of-Oz environment facilitates the entire experiment as described in the section above. The environment consists of the Wizard Interaction Tool, the dialogue system and the wizard. The users wear a headset with a microphone. Their utterances are relayed to the wizard who then annotates it using the Wizard Interaction Tool (shown in figure 1) and sends it to the dialogue system. The system responds with a natural language utterance which is automatically converted to speech and is played back to the user and the wizard.

4.1 Wizard Interaction Tool (WIT)

The Wizard Interaction Tool (WIT) (shown in figure 1) allows the wizard to interact with the dialogue system and the user. The GUI is divided in to several panels.

a. System Response Panel - This panel displays the dialogue system's utterances and RE choices for the domain objects in the utterance. It also displays the strategy adopted by the system currently and a visual indicator of whether the system's utterance is being played to the user.

b. Confirmation Request Panel - This panel lets the wizard handle issues in communication (for e.g. noise). The wizard can ask the user to repeat, speak louder, confirm his responses, etc using appropriate pre-recorded messages or build his own custom messages.

c. Confirmation Panel - This panel lets the wizard handle confirmation questions from the user. The wizard can choose 'yes' or 'no' or build a custom message.

yes	"Yes it is on"
no	"No, its not flashing"
ok	"Ok. I did that"
req_description	"Whats an ethernet cable?"
req_location	"Where is the filter?"
req_verify_jargon	"Is it the ethernet cable?"
req_verify_desc	"Is it the white cable?"
req_repeat	"Please repeat"
req_rephrase	"What do you mean?"
req_wait	"Give me a minute?"

Table 1: User Dialogue Acts.

d. Annotation panel - This panel lets the wizard annotate the content of participant's utterances. User responses (dialogue acts and example utterances) that can be annotated using this panel are given in Table 1. In addition to these, other behaviours, like remaining silent or saying irrelevant things are also accommodated.

e. User's RE Choice panel - The user's choice of REs to refer to the domain objects are annotated by the wizard using this panel.

4.2 The Instructional Dialogue Manager

The dialogue manager drives the conversation by giving instructions to the users. It follows a deterministic dialogue management policy so that we only study variation in the decisions concerning the choice of REs. It should be noted that typical WoZ environments (Whittaker et al., 2002) do not have dialogue managers and the strategic decisions will be taken by the wizard. Our dialogue system has three main responsibilities - choosing the RE strategy, giving instructions and handling clarification requests.

The dialogue system, initially randomly chooses the RE strategy at the start of the dialogue. The list of strategies are as follows.

1. Jargon: Choose technical terms for every reference to the domain objects.
2. Descriptive: Choose descriptive terms for every reference to the domain objects.
3. Tutorial: Use technical terms, but also augment the description for every reference.

The above three strategies are also augmented with an *alignment feature*, so that the system can either align or not align with the user's prior choice of REs. In aligned strategies, the system abandons the existing strategy (jargon, descriptive or tutorial) for a domain object reference when the user




Figure 1: Wizard Interaction Tool

uses a different expression from that of the system to refer to the domain object. For instance, under the descriptive strategy, the ethernet cable is referred to as “the thick cable with red ends”. But if the user refers to it as “ethernet cable”, then the system uses “ethernet cable” in subsequent turns instead of the descriptive expression. In case of non-aligned strategies, the system simply ignores user’s use of novel REs and continues to use its own strategy.

The step-by-step instructions to set up the broadband connection are hand-coded as a dialogue script. The script is a simple deterministic finite state automaton, which contains execution instruction acts(e.g. Plug in the cable in to the socket) and observation instruction acts(e.g. Is the ethernet light flashing?) for the user. Based on the user’s response, the system identifies the next instruction. However, the script only contains the dialogue acts. The dialogue acts are then processed by a built-in realiser component to create the system utterances. The realiser uses templates in which references to domain objects are changed based on the selected strategy to create final utterances. By using a fixed dialogue management policy and by changing the REs, we only explore users’ reactions to various RE strategies.

The utterances are finally converted to speech and are played back to the user.

The dialogue system handles two kinds of clarification requests - open requests and closed requests. With open CRs, users request the system for location of various domain objects (e.g. “where is the ethernet cable?”) or to describe them. With closed CRs, users verify the intended reference, in case of ambiguity (e.g. “Do you mean the thin white cable with grey ends?”, “Is it the broadband filter?”, etc.). The system handles these requests using a knowledge base of the domain objects.

4.3 Wizard Activities

The primary responsibility of the wizard is to understand the participant’s utterance and annotate it as one of the dialogue acts in the Annotation panel, and send the dialogue act to the dialogue system for response. In addition to the primary responsibility, the wizard also requests confirmation from the user (if needed) and also responds to confirmation requests from the user. The wizard also observes the user’s usage of novel REs and records them in the User’s RE Choice panel. As mentioned earlier, our wizard neither decides on which strategy to use to choose REs nor chooses

the next task instruction to give the user.

5 Data collected

Several different kinds of data are collected before, during and after the experiment. This data will be used to build user simulations and reward functions for learning REG strategies and language models for speech recognition.

1. WIT log - The WIT logs the whole conversation as an XML file. The log contains system and user dialogue acts, time of system utterance, system's choice of REs and its utterance at every turn. It also contains the dialogue start time, total time elapsed, total number of turns, number of words in system utterances, number of clarification requests, number of technical, descriptive and tutorial REs, number of confirmations etc.
2. Background of the user - The user is asked to fill in a pre-task background questionnaire containing queries on their experience with computers, Internet and dialogue systems.
3. User satisfaction survey - The user is requested to fill in a post-task questionnaire containing queries on the performance of the system during the task. Each question is answered in a four point Likert scale on how strongly the user agrees or disagrees with the given statement. Statements like, "Conversation with the system was easy", "I would use such a system in future", etc are judged by the user which will be used to build reward functions for reinforcement learning of REG strategies.
4. Knowledge pre-test - Users' initial domain knowledge is tested by asking them to match a list of technical terms to their respective descriptive expressions.
5. Knowledge gain post-test - Users' knowledge gain during the dialogue task is measured by asking them to redo the matching task.
6. Percentage of task completion - The wizard examines the final set up on the user's table to determine the percentage of task success using a form containing declarative statements describing the ideal broadband set up (for e.g. "the broadband filter is plugged in to the phone socket on the wall"). The wizard awards one point to every statement that is true of the user's set up.
7. User's utterances WAV file - The user's utterances are recorded in WAV format for building language models for automatic speech recognition.

6 Results from pilot studies

We are currently running pilot studies (with 6 participants so far) and have collected around 60 minutes of spoken dialogue data. We found that in the jargon strategy, some users take a lot longer to finish the task than others (max 59 turns, min 26 turns). We found that besides requesting clarifications, sometimes novice users assume incorrect references to some domain objects, affecting their task completion rates.

7 Conclusion

We have presented a novel Wizard-of-Oz environment to collect spoken data in a real situated task environment, and to study user reactions to a variety of REG strategies, including system alignment. The data will be used for training user simulations for reinforcement learning of REG strategies to choose between technical, descriptive, tutorial, and aligned REs based on a user's expertise in the task domain.

Acknowledgements

The research leading to these results has received funding from the European Community's Seventh Framework (FP7) under grant agreement no. 216594 (CLASSiC Project www.classic-project.org), EPSRC project no. EP/E019501/1, and the British Council (UKIERI PhD Scholarships 2007-08).

References

- A. Arts. 2004. *Overspecification in Instructive Text*. Ph.D. thesis, Tilburg University, The Netherlands.
- S. Janarathanam and O. Lemon. 2009. Learning Lexical Alignment Policies for Generating Referring Expressions for Spoken Dialogue Systems. In *Proc. ENLG'09*.
- M. J. Pickering and S. Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27:169–225.
- K. van Deemter, I. van der Sluis, and A. Gatt. 2006. Building a semantically transparent corpus for the generation of referring expressions. In *Proc. INLG'06*.
- S. Whittaker, M. Walker, and J. Moore. 2002. Fish or Fowl: A Wizard of Oz Evaluation of Dialogue Strategies in the Restaurant Domain. In *Language Resources and Evaluation Conference*.