

Influencing User Behaviour in Personalised Location Based Services

Srinivasan Janarthanam and Oliver Lemon
School of Mathematical and Computer Sciences,
Heriot Watt University, Edinburgh.
sc445, o.lemon@hw.ac.uk

1. Introduction

We present an exploratory analysis of how dialogue systems providing location based services (LBS) to pedestrian users can influence their behaviour. SPACEBOOK is a EU FP7 project (2011-2014) with an objective of building a speech based LBS system that helps users to a) locate themselves, b) navigate to a destination, c) be informed about points of interest (PoI) near them and d) search for amenities like restaurants and ATMs nearby (Mackness et al 2011). The system will be implemented as a mobile phone app with which a user interacts conversationally (i.e. speech-only, no visual information) as he is walking along the streets.

There are two scenarios where the information provided by the system is likely to influence user behaviour. Firstly, information provided by the system is sought by the user himself to make decisions. For example, the system provides the name and location of the closest restaurant on request. Secondly, the system provides information that is relevant to the current spatial context (i.e. location and viewshed). For example, informing the closest point of interest (say a famous church) when the user is near it or can see it. It is important that in both cases the user finds the information reliable, interesting and useful.

2. Kinds of influences

The main influence the system aims to have on users is to help them to take situated decisions based on information presented to them. For example, the user could choose to visit a museum because the system informs him that it houses artwork that he is interested in. Another example would be for the user to dine in a restaurant the system recommended. Also, every task-based dialogue system implicitly aims to influence the user to keep using the system. This is reflected in the user satisfaction questionnaires that the evaluating users are asked to fill in when these systems are evaluated. Users are typically asked questions such as "Will you use this system in future?" or "Will you recommend this system to your friend?".

Since the SPACEBOOK system also gives information about points of interest, we consider it desirable for users to extend the conversation by asking follow up questions. Users who are interested may also initiate conversations with the system and consult it to verify information that they already know. In order to exert considerable influence on the user, the system must cooperate with the user as much as possible and therefore adhere to the Gricean maxims on cooperative conversation (Grice 1975).

3. Mechanisms to influence users

There are at least three key areas where SPACEBOOK can work towards influencing its users: *user modelling*, *dialogue management* and *natural language generation*. Usually the three modules do not work in isolation. User modelling usually supports both the dialogue manager and the language generation module in their efforts to personalise the conversation to the user. We describe the challenges in each of three areas.

3.1 User Modelling

In order to provide interesting information to the user, SPACEBOOK should combine the user location information with user interests. This information can be augmented with information predicted based other verbal reports by the user and by analysing his movement patterns. Finally, this information can be used effectively to present the most appropriate information.

User modelling techniques can also be used to model the user's knowledge of the city so that effective navigation instructions can be generated. These techniques can also be used to identify user's preferences and constraints to handle tasks like offering amenity information. The user may have different kinds of constraints (e.g. financial, physical, and environmental) which must be taken into consideration. The challenge here is to collect a large amount of information about users that have a bearing on the different SPACEBOOK tasks. The system should try to collect data about the user as unobtrusively as possible so that the user is not frustrated.

3.2 Dialogue Management

The Dialogue management (DM) module is responsible for managing the interaction between the user and the system. It decides *when to speak* and *the high-level content of what to speak* to the user. The system should prioritise among its many tasks and attend to the most important task. For example, a PoI narration can be paused if the user is at a navigation decision point. Another instance where the timing (i.e. when to speak) is important is when finding the appropriate spatial context in order to present PoI information.

The content of information presented by SPACEBOOK should be reliable, useful and interesting. Firstly, correct and reliable information may increase the trust users have in the system. For example, if the system informs that the museum will be open at a certain time and it turns out to be correct, the user may ask for more information and use them in his decision making process. Therefore it is important for the system to identify the source of the information and make a decision to either present it or politely decline saying "I don't know" depending on certainty of information. Secondly, when presenting users with unsolicited information (e.g. PoI narration), it should make efforts to present information that is interesting and novel to the user. This will prevent users becoming bored. Also, novel information can have a positive effect of learning something new. Finally, it is important not to overload the user cognitively during the conversation as he is already engaged in exploring and navigating activities.

However, the challenge is to manage an interesting and a useful conversation under uncertainty from three sources of information. Firstly, the accuracy of GPS positioning that we use to locate users depends on several factors and is shown to be off by several meters (up to 30 meters) from the true position of the user (Zandbergen et al 2011). Secondly, users' speech input may not be easy to recognise owing to factors like traffic noise and wind. Finally, there may be discrepancies between the real world and the backend database of real world entities.

3.3 Natural Language Generation

The manner in which the information is presented also contributes to the system's influence on the user. NLG systems have been shown to persuade users in recommendation systems (Tintarev 2010), to help them in making decisions (Gatt et al 2009), and also to influence their affective states (van der Sluis and Mellish 2008).

The natural language generation (NLG) module should present the information in a concise manner avoiding excess information. For instance, in giving route directions, it may be unnecessary to give redundant information (e.g. street name) when the system has already identified the next navigation goal using a landmark (e.g. a visible tower).

When presenting uncertain information or making decisions based on uncertain inputs, it is important to decide between being vague and being specific during presentation (van Deemter 2009). With uncertain location and orientation information the NLG module should decide cautiously when to use vague terms like "near", "in a few yards", "around", etc in place of specific terms like "after 7.56 yards", "on your right", etc. The NLG module could also inform the user how much the information is reliable using phrases like "I think.." or "Wikipedia reports that.." etc.

The NLG module should also aim to adapt the information to the user based on the user model. One instance of such adaptation would be to choose referring expressions of structures based on the user's knowledge. For instance, when the user knows how Scott's monument looks like, it is best to call it by name than describe it when being used as a landmark in navigation (e.g. "Walk past Scott monument" instead of "walk past the tall dark gothic tower") (see Janarthnam and Lemon 2010b,c).

4. A Machine Learning approach

In this section, we present the approach we propose to address the challenges described above. We propose a reinforcement learning approach to handle user modelling, interaction management and natural language generation. Reinforcement learning (RL) is a popular machine learning approach where an agent learns through trial-and-error by interacting with its environment (Sutton and Barto, 1998). The agent is presented a learning problem in the form of a Markov Decision Process (MDP) consisting of a set of states S , a set of actions A , transition probabilities T from one state to another (when an action is taken) and rewards R associated with such transitions. The agent learns to solve the problem by learning a policy, $\pi: s \rightarrow a$ that optimally maps all the states to actions that lead to a large long term reward.

We plan to approach the problem of user modelling using dynamic modelling techniques that will allow the system to question the user about his interests at appropriate times and build a user model. Janarthanam and Lemon (2010a) showed how to model an individual user's domain knowledge dynamically during a conversation in a technical support dialogue setting using RL techniques. Similar methods can be applied to learn the user's knowledge of the city to facilitate smart navigation. For instance, the user already knows how to get to A, but he wants to get to B which is beyond A, then it would be logical to not give him detailed route instructions until he reaches A, which he can do by himself. We propose to extend this approach to user modelling in the SPACEBOOK tasks. The challenge here is to model the several dimensions of the user (i.e. preferences, goals, constraints, etc) in addition to his knowledge of the city unobtrusively.

Dethlefs and Cuayahuyitl (2011) showed how hierarchical reinforcement learning techniques can be used to generate route instructions in a virtual indoor environment. We propose to extend this approach to generate instructions in real outdoor environments. The challenge here is to generate appropriate referring expressions and instructions (with different levels of vagueness or precision) in the presence of uncertain information. We also plan to use hierarchical reinforcement learning to handle a hierarchy of tasks that the interaction manager is supposed to handle for effective conversation with the user. The system will learn when to speak and how to prioritise between tasks, and what to say such that it maximizes productivity, user engagement, satisfaction and reduces the cognitive load of the user.

The system will learn optimal dialogue management, user modelling and language generation choices by interacting with user simulations built based on data from a corpus collected from real users using wizard-of-Oz studies. Our initial rule-based prototype interaction manager handles tasks such as navigation, Poi information and closest amenity. A statistical variant has shown that it is possible to learn an effective IM policy that maps IM states to dialogue actions. However, these prototypes do not yet adapt to different types of users.

Finally, we propose to evaluate the system's influence on the user by having real users interact with the system both in simulated (Janarthanam and Lemon 2011) and real environments. Users will be asked to rate the system in terms of usability, helpfulness, obtrusiveness, trustworthiness, future use, etc that will show how much the system influenced the users.

5. Conclusion

We have presented some considerations in how a LBS dialogue system can influence its users from the perspective of the SPACEBOOK project. We described how user modelling, dialogue management and natural language generation can work together to influence the user to use the information presented to make situated decisions and to improve the future use of the system itself. We also presented preliminary ideas as to how these issues will be addressed using reinforcement learning techniques and corpus driven environment simulations. We are currently implementing these ideas in the SPACEBOOK dialogue system and in the future we hope to evaluate the influence it has on real users.

Acknowledgements

The research has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 216594 (SPACEBOOK project www.spacebook-project.org).

Bibliography

- Dethlefs, N. and Cuay´ahuitl, H. 2011. Hierarchical Reinforcement Learning and Hidden Markov Models for Task-Oriented Natural Language Generation. In Proc. of ACL. 2011.
- Gatt, A., Portet, F., Reiter, E., Hunter, J., Mahamood, S., Moncur, W. and Sripada, S. 2009. From Data to Text in the Neonatal Intensive Care Unit: Using NLG Technology for Decision Support and Information Management *AI Communications* 22:153-186
- Grice, H. P. 1975. Logic and Conversation. In Cole, P. and Morgan 1975 41–58.
- Janarthnam, S. and Lemon, O. "The GRUVE Challenge: Generating Routes under Uncertainty in Virtual Environments", Proceedings of European Natural Language Generation (Generation Challenges) workshop, 2011
- Janarthnam, S. and Lemon, O. 2010a. Learning Adaptive Referring Expression Generation Policies for Spoken Dialogue Systems. In Krahmer, E., Theune, M. (eds.) *Empirical Methods in Natural Language Generation*, LNCS, vol. 5980. Springer, Berlin / Heidelberg (2010).
- Janarthnam, S. and Lemon, O. 2010c. Adaptive Referring Expression Generation in Spoken Dialogue Systems: Evaluation with Real Users. In proceedings of SigDial 2010 (Tokyo, Japan).
- Janarthnam, S. and Lemon, O. 2010b. Learning to Adapt to Unknown Users: Referring Expression Generation in Spoken Dialogue Systems. In proceedings of ACL 2010 (Uppsala, Sweden).
- Mackaness, W., Boye, J., Clark, S., Fredriksson, M., Geffner, H., Lemon, O., Minock, M. and Webber, B. 2011. "The SpaceBook Project: Pedestrian Exploration of the City Using Dialogue Based Interaction over Smartphones", Proceedings of 8th Symposium on Location Based Services, Vienna 2011
- Sutton, R. and Barto, A. 1998. *Reinforcement Learning*. MIT Press.
- Tintarev, N. 2010. *Explaining Recommendations*. PhD Thesis, University of Aberdeen.
- van Deemter, K. 2009. Utility and language generation: the case of vagueness. *Journal of Philosophical Logic* 38 (6).
- van der Sluis, I. and Mellish, C. 2008. "Towards Affective Natural Language Generation: Empirical Investigations", In proceedings of AISB convention 2008.
- Zandbergen, P. A. and Barbeau, S. J. 2011. Positional accuracy of assisted GPS data from high-sensitivity GPS-enabled mobile phones. *Journal of Navigation*, 64(3):381–399.