Templates for Wearables in Context

Sabine Geldof

Vrije Universiteit Brussel
Artificial Intelligence Laboratory
Pleinlaan 2, B-1050 Brussels, Belgium
sabine@arti.vub.ac.be
http://arti.vub.ac.be

1 Introduction

In the COMRIS[1] project we are developing natural language generation (NLG) technology for output on a wearable device: a user, moving freely in the area of a conference receives from her COMRIS parrot (spoken) advice (e.g. on events and encounters not to miss, reminders about commitments and proposals for meetings). This setting requires NLG technology that is fast and context sensitive: otherwise the user will be bored and turn off her device. In contrast to other NLG projects for wearable devices, e.g., HIPS [1], the COMRIS parrot is supported by a mixed reality set-up. A multitude of agents, defending users' interests and aware of their physical context, are continuously interacting in the virtual space, eager to pursue the specific (user) interest they stand for. They try to push the results of their activities to the physical space in the form of short, ad-rem messages to the user. A mechanism of competition for attention [2] ensures that only the relevant messages are passed to the user. The NLG module produces two outputs based on the same text: one version contains basic prosodic annotations (phrase boundaries, accentuation marks) and will be further processed by the speech synthesis module, the other version consists of the same text annotated with html tags for web-based interaction, in case the user wants to explore further related information. We will briefly describe and more extensively demonstrate how we model the different aspects of a user's context (in order to annotate NLG input) and how these data influence the (output of the) NLG process. The underlying idea is that context sensitivity is a prerequisite for NLG technology to evolve with current technological developments and that a user's context encompasses more than discourse history.

2 NLG strategy in COMRIS: focus on context sensitivity

Multi-dimensional context Context is a very complex and multifaceted phenomenon which has been studied in logic for several years [3]. There is a growing awareness and interest to study it also from two other perspectives: engineering and natural language processing [4]. Within the field of NLP, mostly the linguistic (or discourse) context has been studied both for language understanding and generation. We are striving at a more global account of context, specifically for the purpose of generating natural language. In earlier experiments [5] we let NLG output vary according to the minute-by-minute evolving interests of the user. COMRIS adds another dimension: the physical context.

Thus, at least three dimensions of a user's context should influence the output of an NLG component: (a) linguistic context (the discourse model), (b) extra-linguistic context (the physical context) and (c) the user's profile. Consider the following examples (*italics* indicate context-sensitive expressions of the corresponding type):

(a) John Lewis will give a presentation on Robotics. *He* will *also* chair a panel on *the same topic*.
(b) Please note that you have to give a presentation on Monday, June 14th at 2 o'clock in Room B. Please note that you have to give a presentation *within 10 minutes* in Room B *at the other side of the building*.

[1] http://arti.vub.ac.be/~comris

(c) Isabel Baud, *who is also interested in natural language generation and wearable devices*, is currently in the same room as you. She has an interesting demo at the booth of the SIRCOM project.

We developed a simple representation formalism (see Table 1) for these three types of context information in order to annotate the propositional content provided as input to the NLG process. The text generator produces variable output according to the values of these contextual parameters.

Generation phases Ideally, context sensitivity will be taken care of in all three phases of the generation process. In COMRIS content determination and sentence planning take place outside the proper NLG component. Still, each of the three phases somehow contributes to context sensitivity.

The content of a generated utterance is mostly determined by the activities of the personal representative agents (PRA), who are exploring the virtual space in search of interesting information and encounters related to the user's interest which they have to pursue (e.g., robotics, template-based NLG). Their interactions are mainly driven by user profile information. Hence, the objects that are part of the propositional content are automatically choosen w.r.t. the hearer's interest. The outcome of PRAs activities consists of input to the NLG component (e.g., data about a panel discussion on The Future of Agents). Parts of these input structures might not be realised by the NLG component, if that suits the hearer's context better (e.g., include or omit information about the affiliation of a mentioned person might (not) be relevant for the hearer).

Sentence planning is controlled by the personal assistant agent (PA) who is on the edge of the virtual and the physical world, representing the user as a person (not one of her specific interests). Indeed, PRAs have to compete for the attention of the user with their coded message contents. The PA rules this competition, taking into account the physical context of the user but mainly by evaluating relevance, competence and performance measures (e.g., when a message, estimated very relevant by the PRA receives negative feedback from the user, the corresponding PRA's performance will decrease, thus lowering its chances in future competition for attention). Further sentence planning is carried out in the NLG component: the presence or absence of particular items in the input structure determines which rules to follow and hence which templates will be realised.

Finally, most of the context sensitive adaptations take place at the level of surface realisation through referring expressions (linguistic context), relative time and spatial expressions (physical context) and insertion of additional information (related to the user's known interests). These surface annotations are guided by the contextual annotations, contributed by the different components along the generation process (PRA: profile values; PA: physical context; NLG-discourse module: linguistic context).

Variable output based on templates We have various reasons for using templates rather than deep generation as basic technique. The expected output is canonical, i.e., it is determined by the limited set of protocols (scenes) PRAs can enter. Moreover, efficiency is a major concern, since real-time processing is required. A third argument is that modular sets of templates allow us to adapt our NLG component easily to eventual extensions of the application (new scenes). A well known limitation of templates, nl. the difficulty in reusing templates across domains and even applications holds also for our case, especially at the macro-level of sentence patterns. However, the relative ease with which templates and inputstructures can be created compensates for this limitation and we managed to reuse at least partly, some specific structures, e.g., for manipulating lists, or rendering dates, from an earlier application.

Our NLG component is developed using the TG/2 tool[2], which allows for template-based NLG as well as deep generation [6]. We currently use templates only, but as we extend our system to more scenes, we are looking for generalities that are worth implementing small subgrammars, similar to the functions in the syntactic templates of [7]. COMRIS text templates are encoded

[2] TG/2 is used in the COMRIS project under a license agreement between VUB and DFKI for the purpose of scientific research.

into TG/2 production rules in a special purpose language (TGL). A first strategy for obtaining variable output according to contextual parameters is to formulate conditions on the presence of particular context values. For instance: if the profile value denotes high interest of the hearer in the topic (of a talk to be announced), then include it in the output message. Table 1 gives an overview of this input-output context sensitive variability for all the sub-mentioned context parameters.

Table 1. Overview of context annotation values in COMRIS

Linguistic context	**lcv (a,**	**b)**
Concept/ Instance (mentioned to user:..)	NUMeric focus on **Concept** **1**: occurred 2/more messages ago **3**: occurred 1 message ago **5**: occurred in previous message	NUMeric focus on **Instance** **1**: occurred 2/more messages ago **3**: occurred 1 message ago **5**: occurred in previous message
Extra-Linguistic context:	**d_elcv (+/- n)**	**l_elcv (”string“)**
(user is in ..) time/ space social implicature: (user is..)	(date) 0: today (time) +1: in one hour (time) +.10: in 10 minutes **s_elcv (n)** **1..2**: standing, wandering around **3**: moving towards a goal **4**: attending an event **5**: talking to someone	(location) close-by, at the other side of the building, on the same floor, ...
Profile Value: user's interest	**t_pv (n)**	
..in a topic:	**1..2**: .. you might be interested in **3**: .. you are interested in **5**: .. your favourite topic	
	p_pv (pnumval n	**{pqval ”string“})**
..in a person:	**1..2**: .. you might want to meet **3**: .. you wanted to meet **5**: .. you absolutely wanted to meet	for dicsussion, for introduction, for lobbying, for socializing, as an expert,...

This strategy requires explicitly acquired and encoded information about the context of the hearer. However, in our experiments it became clear quite soon that the acquisition and management of such data might become a bottleneck, both from the conceptual and efficiency point of view: the NLG component has to rely on other COMRIS components to provide these data. Each interaction among components requires an exchange of agent messages via the COMRIS infrastructure. By putting extra burden on each of the components involved in the NLG process from content determination (by PRA) to linguistic realisation, we might, in the worst case, endanger the timely delivery of a message to the hearer.

Therefore, we are looking for alternatives using global and numeric context information, as it is handled by the COMRIS agents. For instance, we might rate the social implicature of a user at every moment, indicating how much she is involved in a social interaction (wandering around, listening to a talk, talking to someone) and rate alternatives within the rule-set accordingly. When the user is talking to someone, only extremely short (and important) messages should be generated, (e.g., 'Look for Isabel Baud', as an alternative for example c). Through a mechanism of parametrization, the TG/2 tool allows to influence the order in which rules are considered, e.g., according to how much social implicature they tolerate. When the user is wandering around, the situation is less resource-bound and more elaborate messages are welcome. The social implicature rating is a (extra-linguistic) context parameter that applies globally to the situation of the user and not to one particular aspect of the contents of the message (as with (extra)linguistic context values or profile values). Comparable to Reiter's constraints [8], we are investigating how to use such global context values which are readily available from the other COMRIS components (PA

& PRA). Being the software part of the wearable device, the PA has direct access to information about the user's environment (e.g., location near beacons, proximity to other (parrot-wearing) users), from which social implicature can be derived.

3 Discussion

In this paper we propose an NLG strategy that takes into account multiple dimensions of the user's context. The input to the NLG process is annotated with both local and global information about the context of the user. We are interested in a hybrid approach: combining local context values (e.g., indicating how much a user is interested in a particular topic of the domain, or whether a specific object has been mentioned recently to the user) with global context parameters (esp. how much the user is available for new information). We experiment with different mechanisms for producing variable output within a template based approach. Our parameters for context sensitivity are application independent, but for every application domain it has to be made explicit how they influence the output text. The COMRIS environment is an interesting test-bed for these ideas on context-sensitive NLG.

Acknowledgments We are grateful to Stephan Busemann for his advice on the use of TG/2. This research is funded by the EC as part of the COMRIS project (LTR 25500), within the Intelligent Information Interfaces[3] (I3) programme.

References

1. Not, E., Petrelli, D., Stock, O., Strapparava, C. and Zancanaro, M.: Person-oriented guided visits in a physical museum. In: Proceedings of the Fourth International Conference on Hypermedia and Interactivity in Museums (ICHIM97), Paris, (1997)
2. Van de Velde, W., Geldof, S., Schrooten, R. Competition for attention. In: Singh, M.P., Rao, A.S., and Wooldridge, M.J., (eds.) Proceedings of ATAL: 4th International Workshop on Agent Theories, Architectures and Languages, LNAI Vol. 1365. Springer-Verlag, Heidelberg (1998) 282-296.
3. McCarthey, J.: Notes on Formalizing Context. In Proceedings of 13th IJCAI conf. (1993) 555-560
4. Brézillon, P., Cavalcanti, M.: Modeling and using context. The Knowledge Engineering Review, 1997 12(**4**) 185-194
5. Geldof, S., Van de Velde, W.: An architecture for template based (hyper)text generation. In: Höppner, W. (ed) Proceedings of the 6th European Workshop on Natural Language Generation. Gerhard-Mercator-Universität Duisburg, Germany. (1997) 28-37
6. Busemann, S.: Best-First Surface Realization. In: Scott, D. (ed) Proceedings of the 8th Intl. workshop on Natural Language Generation, Herstmonceux Castle, University of Brigthon, UK, (1996) 101-110
7. van Deemter, K., Krahmer, E., Theune M.: Plan-based vsr. Template-based NLG: a false Opposition? In: Busemann,S. and Becker, T. (eds) Proceedings of the KI workshop on NLG: May I Speak Freely? Bonn, 1999 (this volume)
8. Reiter, E.: Shallow vs. Deep Techniques for Handling Linguistic Constraints and Optimisations. In: Busemann,S. and Becker, T. (eds) Proceedings of the KI workshop on NLG: May I Speak Freely? Bonn, 1999 (this volume)

[3] http://www.i3net.org/