"Hands Busy, Eyes Busy": Generating Stories from Sensor Data for Automotive applications

Joe Reddington, Ehud Reiter, Nava Tintarev
Department of Computing Science
University of Aberdeen
j.reddington, e.reiter, n.tintarev@abdn.ac.uk

Rolf Black, Annalu Waller
School of Computing
University of Dundee
rolfblack,awaller@ computing.dundee.ac.uk

ABSTRACT
This paper examines the potential of using natural language generation to support "hands busy, eyes busy" automotive applications. It outlines a hierarchy of complexity of output text, and the type of sensor data that may be collected. It also suggests a number of ways natural language generation can generate narrative events from sensor data for drivers.

Author Keywords
NLG, AAC, event generation, narrative, story, sensors, automotive applications

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: Miscellaneous

INTRODUCTION
This work examines the potential of using automatically harvested information to generate new phrases automatically, creating support for "hands busy, eyes busy" automotive applications. Of particular interest is a review of how technologies and techniques developed in an assistive technology application (the recent "How was School Today...?" project) can be applied to the automotive domain.

Mobile usage while driving has been identified as a risk factor in road accidents [2, 5]. Reducing both the motivation to use such devices while driving and the length of time for which they are used would potentially reduce the number of road accidents. The position of the authors is that the use of automatic narration techniques can support communication in scenarios such as making regular deliveries or public transportation. Methodologies to enable this type of automatic text generation are under-researched and NLG can aid in this task by creating a story that is structured, relevant and flexible to the current situation, based on sensor data.

IUI 2011, February 13 - 16, 2011, Palo Alto, CA, USA.

It is easy to envisage a system by which buses or delivery vans automatically send an update of location to a home server, and indeed many services offer near real-time tracking of packages from source to destination. In contrast, this work focuses on combining such messages, augmented with information from weather reports, traffic reports and other data, to form a larger message with an overall narrative.

In this paper we situate the work with regard to existing work, then introduce the "How was School Today...?" project that informed this work. We go on to identify potential application areas in the automotive domain, and discuss the possible effects, risks, and advantages.

RELATED WORK
Our existing work sits on the boundary between Natural Language Generation (NLG), which is a subcategory of natural language processing that examines the creation of text from nonlinguistic data such as sensor readings, and Alternative and Augmentative Communication (AAC), an area examining communication for those with restrictions on speech. NLG techniques can dynamically combine and change some output depending on the changing internal state of a system [11]. A popular application area for NLG has been weather forecasting (generating textual weather forecasts from the results of a numerical atmosphere simulation model), and several weather forecast generators have been fielded and used operationally [17, 16]. A number of data-to-text systems have also been developed in the medical community, such as BabyTalk [15], which generates summaries of clinical data from a neonatal intensive care unit, and the commercial Narrative Engine [14] which summarises data acquired during a doctor/patient encounter.

In this paper, we seek to focus the technology away from AAC and on the automotive domain, where natural language processing systems have been used with some success. For example, RoadSafe is an NLG system that has been operationally deployed at Aerospace and Marine International (AMI) to produce weather forecast texts for winter road maintenance. It generates forecast texts describing various weather conditions on a road network [10]. Other systems have focused more on processing language to visualise and animate 3D scenes from car accident reports [3].

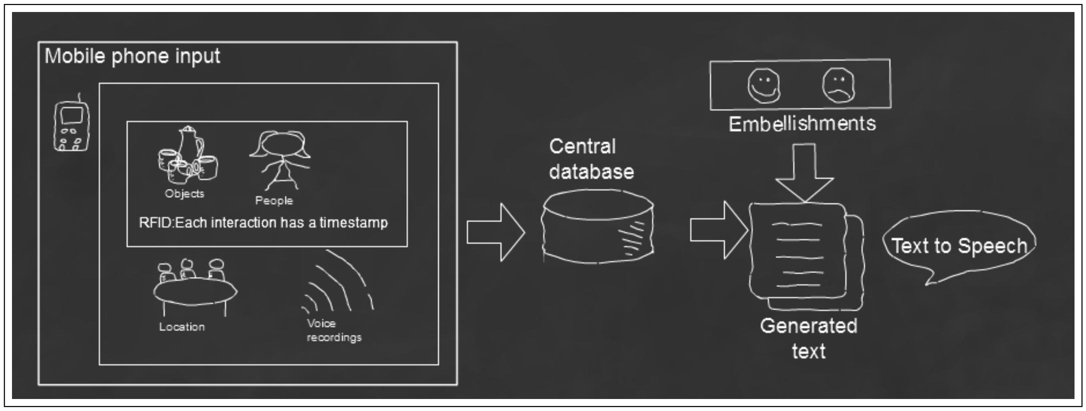

Figure 1. Types of input that can be collected by a mobile device: voice recording, RFID, voice, emotional embellishments

Automotive research in general is well developed; of particular relevance to this work is the issue of privacy in vehicle-to-vehicle, or vehicle-to-base communication, see e.g. [8, 9].

The "How was School Today...?" project

Our work is informed by the "How was School Today...?" (HWST) project [1, 6] which logged sensor data for students at a special needs school. This data included object and person interactions, voice recordings, and location information (at the room level). It also recorded positive and negative evaluations (e.g. "It was not a good day.") input by the children. This framework has been tested as a proof-of-concept in the context of generating stories for children at the school. The students (who had no, or very limited, speech) could then relay these stories to parents or other conversation partners.

For this particular domain, the types of data recorded for each user are:

- Location data - each time the user entered a new room, this information was recorded. (Pre-processing removed rooms entered for less than three minutes).
- Object interaction - each time the user interacted with an object that had an RFID tag, that interaction was recorded.
- Person interaction - each time the user interacted with a person that had an RFID tag, that interaction was recorded.
- Voice messages - staff and teachers were encouraged to record voice messages, as if the user was speaking in the first person, that described the user's recent activities.

An example set of data would be:

11:36, Location, Tutorial Room
11:36, Object, Money
11:39, Object, Monkey Game

Which is converted into English text to give the story:

> I played with Money and Monkey Game. This happened at a Tutorial Room.

Many of the input sensor data and techniques used in HWST can be applied to the automative domain. Figure 1 outlines the type of input that could be used in such a system and collected with a mobile phone, e.g. voice recordings, location, interactions with people and objects (RFID).

The HWST project is in the process of introducing the Nokia 6212 [1] as a collection device, and may need to be supplemented with an additional system for recording location information on the room level.

Depending on the granularity of location data required, other hardware may supplement a mobile phone. GPS tracking may be more suitable for larger distances while bluetooth or other methods may be preferable for room-level identification. Additional sensor data may be available in a vehicle such as change in light, temperature etc [4], or speed and fuel usage.

TYPES OF PRODUCABLE CONTENT

This section categorises the potential outputs of automatically generated content into a triple-tiered hierarchy of network-based input, sensor-based input, and the creation of narratives from sensor input. This hierarchy can be broadly arranged in terms of invasiveness of the data collection. This and other privacy concerns are key to any implementation.

Network-based input

Network-based input is defined as new utterances that can be determined by access to information over the Internet, or some other large information portal. An example is talking about the weather - phrases such as *"It's very warm today"*, and *"The snow is starting to stick!"*, but this can include *"There was an accident on the M14"*, or *"Traffic is slow around Old Trafford due to the match"*.

[1] `http://europe.nokia.com/find-products/devices/nokia-6212-classic`, retrieved November 2010

Sensor-based input

Sensor-based input is defined as the use of single facts about the user provided by sensor data. Examples might include *"I went to Leeds"* - provided by GPS data, or *"I just handled package 41"* - provided by use of a barcode scanner in combination with an online lookup of the IDs for the packages. Although there is a concern that this sort of data collection can affect both privacy and also workload required to maintain it, messages can be better adapted: *"I got a text message from Jamie this morning, he said 'looking forward to tomorrow' "*. Voice messages are included in this category and can include information that would never be picked up by a sensor - "I helped jump-start a car and was 15 minutes late.".

Creation of narratives from sensor data

This category contains those groups of messages, based on sensor data, that together relate an experience or tell a story, thus adding the problems of creating a narrative structure or consistent style to what has previously been a data-mining exercise. The importance of narrative in exchanging information is well-researched, for an NLG example see [12].

In HWST, stories were generated using additional reasoning, such as giving more importance to events that occurred in locations which were unexpected compared to a timetable. These stories were also augmented by users with positive and negative annotations of utterances "She was nice." (for people) or "It was not a good day." (for the whole story) [1]. The creation of multi-fact, multi-sentence messages with a structured narrative is a step forward in NLG-terms, requiring more sophisticated techniques than previous levels in the hierarchy. In particular, this moves the focus of NLG research to the tasks of document planning and document structuring, compared to text generation on the sentence level.

The analysis of sensor-based data, defining one of these multi-fact and multi-sentence messages as an 'event' is discussed in [6]. While the NLG techniques outlined in [11] can combine facts into plain English, a further challenge lies in defining boundaries between groups of sensor data to define separate events. The goal is to arrange the sensor-based input into a narrative structure that accurately relates events.

Based on a modified version of the data recording in the HWST project, one could assume input data such as that highlighted in Figure 2. The generated text could then be:

"This morning, after picking up two packages, I helped jump-start a car and was delayed by 15 minutes. Later, I arrived at the Leeds depot and delivered the packages to Mr. Roberts. The delivery went fine".

06:27:00, Object, Package1
06:27:07, Object, Package2
07:34:00, Voice Recording, I helped jump-start a car and was delayed by 15 minutes.
09:40:00, Location, Leeds depot.
09:40:00, Object, Package1
09:40:05, Object, Package2
09:40:00, Person, Mr. Roberts
09:43:00, Embellishment, Positive ...

Figure 2. Possible input data

APPLICATION AREAS

The previous section discussed the types of text that can be generated. This section outlines several practical applications of the generated narrative text in automative applications: staying in touch; communication with head office; and accident reports. Privacy is an important consideration in any application; the people on whose behalf the story is generated should always have the possibility to read and edit any text before it is transmitted. Moreover, any generated text can be read aloud by text-to-speech software.

This would also facilitate responses to messages originally sent to a driver, allowing the original sender (which may also be a driver) to hear the response without extra effort and reducing cognitive load.

Staying in touch

Many people keep in touch with mobile texts and an increasing number stay connected using social media such as Facebook and Twitter[2]. Professional drivers may feel that updating their status is important from a social as well as professional prospective. However, while driving, attention should be on the road, and hands and eyes will be occupied by driving. An application that uses NLG to automatically update friends on one's activities may help drivers feel connected in their everyday lives. The necessity to automatically generate such short messages is highlighted in [4] who suggest messages such as *"35 centigrades? It is very hot in here!"*. In particular, the work on structuring narrative produced by HWST technology allows a move from the functional single sentence update to a more expressive longer update.

Work Reports

The key application in this area is the generation of automatic work reports based on a driver's sensor data. This sort of narrative can supply an employer with information about his drivers, such as the hours that they have worked and which deliveries or other tasks have been successfully executed. At the same time, the automatic generation of the text relieves the employee of the task of writing lengthy reports. Of particular use is text informing end-users of the current conditions - rather than a simple "Delayed, new ETA:15:27" message, one can imagine "When coming from a previous delivery at Hogsmeade, there was heavy traffic due to an accident in the town so the delivery has been diverted via Hogwarts and should be with you by 15:27".

Accident Reports

Generative narrative stories from sensor data can also be used to support police and ambulance staff at the scene of the accident. The generated reports can offer a human readable summary of the situation well ahead of arrival on the scene, allowing professionals to be ready once they arrive. This sort of report can help assess the degree of damage

[2] www.facebook.com, www.twitter.com, retrieved November 2010

incurred at an accident by considering road conditions and travel speed. This type of report could also help police (and insurance companies) assess potential accountability for a given accident. Infra-red sensors may help assess how many victims were involved in an accident as well, ensuring that all victims get pulled out of an affected vehicle.

CONCLUSION AND ONGOING RESEARCH

This paper describes the type of text that can be automatically generated to support drivers, and highlighted three application areas: staying in touch, communication with head office, and accident reports. Although a future goal for this research is to integrate with a commercial product, privacy and security of such systems require careful consideration While care has been taken to keep such concerns a key part of the research, the authors welcome any communication from parties with expertise in this area.

ACKNOWLEDGEMENTS

The authors are particularly grateful to the school, staff, and children. This research was supported by the UK Engineering and Physical Sciences Research Council under grants EP/F067151/1, EP/F066880/1, EP/E011764/1, EP/H022376/1, and EP/H022570/1.

ADDITIONAL AUTHORS

REFERENCES

1. R. Black, J. Reddington, E. Reiter, N. Tintarev, and A. Waller. Using nlg and sensors to support personal narrative for children with complex communication needs. In *Proceedings of the NAACL HLT 2010 Workshop on Speech and Language Processing for Assistive Technologies*, pages 1–9, Los Angeles, California, June 2010. Association for Computational Linguistics.

2. F. A. Drews, H. Yazdani, C. N. Godfrey, J. M. Cooper, and D. L. Strayer. Text messaging during simulated driving. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 51 (5):762–770, 2009.

3. S. Dupuy, A. Egges, V. Legendre, and P. Nugues. Generating a 3d simulation of a car accident from a written description in natural language: the carsim system. In *Proceedings of the workshop on Temporal and spatial information processing - Volume 13*, pages 1:1–1:8, Morristown, NJ, USA, 2001. Association for Computational Linguistics.

4. C. Endres and D. Braun. Pleopatra: A Semi-Automatic Status-Posting Prototype For Future In-Car Use. In *Adjunct proceedings of the 2nd International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI 2010)*, page 7, Pittsburgh, PA, USA, November 2010.

5. S. P. McEvoy, M. R. Stevenson, and M. Woodward. Phone use and crashes while driving: a representative survey of drivers in two australian states. *Medical journal of Australia*, 185(11/12):630–634, 2006.

6. J. Reddington and N. Tintarev. Automatically generating stories from sensor data. In *Intelligent User Interfaces*, 2011 (to appear).

7. E. Reiter, R. Turner, N. Alm, R. Black, M. Dempster, and A. Waller. Using nlg to help language-impaired users tell stories and participate in social dialogues. In *In Proceedings of the 12th European Workshop on Natural Language Generation (ENLG-09*, 2009.

8. F. Schaub, F. Kargl, Z. Ma, and M. Weber. V-tokens for conditional pseudonymity in vanets. In *IEEE Wireless Communications & Networking Conference (IEEE WCNC 2010)*, Sydney, Australia, 04/2010 2010. IEEE, IEEE.

9. F. Schaub, Z. Ma, and F. Kargl. Privacy requirements in vehicular communication systems. *Computational Science and Engineering, IEEE International Conference on*, 3:139–145, 2009.

10. R. Turner, Y. Sripada, and E. Reiter. Generating approximate geographic descriptions. In *Proceedings of the 12th European Workshop on Natural Language Generation*, ENLG ’09, pages 42–49, Morristown, NJ, USA, 2009. Association for Computational Linguistics.

11. E. Reiter and R. Dale. Building natural language generation systems, *Cambridge University Press*, 2000.

12. E. Reiter, A. Gatt, F. Portet, and M. van der Meulen. The importance of narrative and other lessons from an evaluation of an NLG system that summarises clinical data. *INLG ’08*, pp. 147–156, Morristown, NJ, USA, 2008. Association for Computational Linguistics.

13. S. Ashraf, A. Judson, I. W. Ricketts, A. Waller, N. Alm, B. Gordon, F. MacAulay, J. K. Brodie, M. Etchels, A. Warden, and A. J. Shearer. Capturing phrases for ICU-Talk, a communication aid for intubated intensive care patients. In *ACM Conference on Assistive technologies*, pp. 213–217, New York, NY, USA, 2002.

14. M. D. Harris. Building a large-scale commercial NLG system for an EMR. In *INLG ’08: Proceedings of the Fifth International Natural Language Generation Conference*, pages 157–160, Morristown, NJ, USA, 2008. Association for Computational Linguistics.

15. A. Gatt, F. Portet, E. Reiter, J. Hunter, S. Mahamood, W. Moncur, and S. Sripada. From data to text in the neonatal intensive care unit: Using NLG technology for decision support and information management. *AI Commun.*, 22(3):153–186, 2009.

16. E. Reiter, S. Sripada, J. Hunter, J. Yu, and I. Davy. Choosing words in computer-generated weather forecasts. *Artif. Intell.*, 167(1-2):137–169, 2005.

17. E. Goldberg, N. Driedger, and R. I. Kittredge. Using natural-language processing to produce weather forecasts. *IEEE Expert: Intelligent Systems and Their Applications*, 9(2):45–53, 1994.