

Determining Content for Unknown Users: Lessons from the MinkApp Case Study

Gemma Webster, Somayajulu G. Sripada, Chris Mellish, Yolanda Melero, Koen Arts, Xavier Lambin, Rene Van Der Wal

University of Aberdeen

{gwebster, yaji.sripada, c.mellish, y.melero, k.arts, x.lambin, r.vanderwal}@abdn.ac.uk

Abstract

If an NLG system needs to be put in place as soon as possible it is not always possible to know in advance who the users of a system are or what kind of information will interest them. This paper describes the development of a system and contextualized text for unknown users. We describe the development, design and initial findings with a system for unknown users that allows the users to design their own contextualised text.

1 Introduction

Requirements of an NLG system are derived commonly by analysing a gold standard corpus. Other knowledge acquisition (KA) techniques such as interviewing experts and end-users are also frequently employed. However, when these KA studies result in only a partial specification of the system requirements or complications make carrying out a detailed user study in the time available difficult, an initial system for unknown users may need to be developed. The initial system needs to fulfil the known requirements making a number of assumptions to fill the gaps in the requirements. In this paper, we concentrate on the content determination problem for such a system.

We encountered this particular problem when producing an initial NLG system to give feedback to volunteers submitting information about signs of American Mink, an invasive species in Scotland. Our response can be viewed, on one hand, as that of exposing an early prototype for evaluation in real use. On the other hand, it can be viewed as an approach to allowing users to “design their own contextualised text”. We expected that this approach would have a number of advantages. In the paper, we draw our conclusions about how this worked out in our example application.

2 Background - MinkApp

The Scottish Mink Initiative (SMI) project aims to protect native wildlife by removing breeding American Mink (an invasive species) from the North of Scotland. SMI in the form discussed here was launched in May 2011 and ran until August 2013, after which it continued but on a much smaller funding base. SMI’s success and future rely on an ongoing network of volunteers from across Scotland to monitor the American mink population. During the period from 2011 to 2013, these volunteers were coordinated by 4 and later 3 full-time Mink Control officers (MCOs) who had 2.5 year fixed term contracts, had no communal offices and were geographically located across Scotland.

At present volunteers are provided with rafts to monitor American mink. Rafts are simple devices that float on water and are monitored by volunteers who regularly check a clay pad for mink footprints. In the past, volunteers in turn reported signs or lack of signs to their corresponding MCO. Now volunteers can do the same through the MinkApp website, introduced in 2012, though some choose to continue to use the previous reporting method. The data should ideally be entered roughly every 10 days; it concerns either positive or negative records from raft checks, or visual sightings of mink and actual mink captures. The records contain geographical information and a timestamp. MinkApp checks whether this data is complete and then informs the respective mink officer for that volunteer’s area and enters the data into the database.

Volunteers used to receive a quarterly newsletter that had some regional specific content but was not volunteer specific. They could receive sporadic contact from their mink control officer in the form of a phone call or email. MinkApp allowed an infrastructure to be developed to provide volunteers with specific and immediate

feedback upon submission of their observations by means of contextualised feedback text.

SMI's funding base was severely reduced in August 2013 and MinkApp has proven central to its endurance. Volunteer activities of the SMI are now supported by staff from 10 local rivers and fisheries trusts (as one of their many activities). This limited amount of staff time available could make the development of automatic personalised feedback generation vital to allow volunteers to have tailored information on the progress of the project and to keep volunteers engaged.

3 The Problem - SMI Volunteers: The Unknown Users

The nearest to a gold standard for what information to offer was the corpus of newsletters containing information on the project as a whole. However, we learned that these newsletters were often not read and we have no way of judging their level of success. These newsletters, along with emails and discussions conducted with SMI employees on their interactions with volunteers, however, gave us ideas about potential content that could be selected and indication of potential lexical structure and word use when addressing volunteers.

Although some SMI volunteers monitor mink as part of their job (e.g. gamekeepers), they could in fact be anyone with a desire to contribute to nature conservation. Volunteers are located in very disparate geographical locations across Scotland, with no set gender or age range and so volunteers' motivations, computer skills and professions are mostly unknown. Because of the range of types of people who could in principle be volunteers, they can be expected to be very varied.

It is extremely difficult to contact all volunteers as each SMI catchment is managed and organized in different ways and volunteers are contacted using different media e.g. mail, email, telephone, face-to-face. SMI is also careful to avoid attempting to contact volunteers too often, conscious that they are providing their services for free and should not be bothered unnecessarily. There is also some uncertainty about which volunteers are active, as records are often partial or out of date. It is known anecdotally from MCOs that many volunteers are unwilling to use any kind of computer system and so it is unclear what kind of people will be reached through MinkApp. Finally, most observations of mink

signs that arise are "null records", i.e. records of observing no mink prints on rafts. It is not known which volunteers will be sufficiently motivated to submit "null records" and which will remain apparently inactive because they have nothing positive to report.

So, even though there was a need for automatically generated feedback *now*, there was a real question of who the readers would be and how to select the content to include in the feedback.

4 Related Work

A standard approach to establish user requirements for NLG is to assemble a corpus of human-authored texts and their associated inputs (Reiter & Dale, 2000). This can be the basis of deriving rules by hand, or one can attempt to replicate content selection rules from the corpus by machine learning (Duboue & McKeown, 2003; Konstas & Lapata, 2012). To produce a useful corpus, however, one has to know one's users or have reliable expert authors.

As first pointed out by Levine et al. (1991), an NLG system that produces *hypertext*, rather than straight text, can avoid some content selection decisions, as the user makes some of these decisions by selecting links to follow. A similar advantage applies to other adaptive hypertext systems (Brusilovsky, 2001). Another general possibility is to allow users to design aspects of the texts they receive. For instance, ICONOCLAST (Power, Scott, & Bouayad-Agha, 2003) allows users to make choices about text style. However, relatively little is known about how such approaches work 'in the wild'.

Various previous work has attempted to build models of users through observing interactions with an interface (Fischer, 2001). Alternatively, it is possible to explicitly ask questions to the user about their interests (Tintarev & Masthoff, 2008), though this requires the users to have the time and motivation to take part in an initial activity with no direct reward.

Our approach can be seen to have similarities with hypertext generation, in that we are offering alternative texts to users, and non-invasive approaches to user modelling.

5 Approach to Content Selection

To overcome the ‘unknown’ user and ‘unknown’ feedback problem it was decided to implement a relatively quick exploratory tool that could be used to help understand user requirements, provide initial evaluation of feedback content and build an understanding of user interests. To achieve these aims we developed a tool that allows users to generate their own text, selecting content from a larger set of possibilities. The information on the type of feedback generated by the user would allow us to investigate user stereotypes, their detection and the automatic adaptation of content based on their interests (Zancanaro, Kuflik, Boger, Goren-Bar, & Goldwasser, 2007).

5.1 Exploratory Tool - The Feedback Form

The feedback form (Figure 1) is displayed to users of the MinkApp system once they have submitted a raft check. The form allows the user to select which raft they wish to have their feedback generated on from a list of the rafts they manage. The users have four types of information they can select to have feedback generated on: Signs (information on signs of mink reported through raft checks), Captures (information on mink captures), My Rafts (information on their personal raft checks and submission record) and Mink Ecology (information on mink behaviour and seasonality).

Two of the four options, Signs and Captures, allow the user to select to what geographic scale they would like their feedback based on: the whole of the SMI project area, their river or their catchment – the geographical region that they report to e.g. Aberdeenshire, Tayside etc.

Once the user has made their selection the personalised feedback based on their choices is generated and displayed along with an option to rank how interesting they found this feedback or any comments they wish to make. The user can gen-

erate multiple texts in one session. All data from each click of an option, the generated text and user comments on the text are recorded.

5.2 Generation of the paragraphs

The structure of the text is separated out into self-contained paragraphs to allow analysis of what volunteers regularly view. For each type, the structure of the generated paragraph is determined by a simple schema:

Signs:

Neighbourhood (based on user selection) – *In the Don catchment there have been 6 signs of mink reported over the past 12 months which is higher than the previous 12 months*

Additional Information / Motivation – *Mink are coming into your area to replace captured mink. This shows your area has good ecology for mink and it is important to keep monitoring.*

Personal – *There have been no signs of mink (in the form of either footprints or scat) in the past 30 days. No signs of mink recently does not mean they are gone - remain vigilant.*

Captures:

Neighbourhood (based on user selection) – *In the Spey catchment we have trapped 5 mink over the past 12 months which is lower than the previous 12 months.*

Additional Information / Motivation – *Information available on this year's captures: An adult female mink was captured on: 2014-02-19.*

My Rafts:

Personal – *You have been very active over the past 60 days with 7 'no mink signs' reported and 2 signs of mink (in the form of either footprints or scat) reported, the last of which was logged on 14 Sep 2013 23:00:00 GMT.*

Additional Information / Motivation – *Please keep checking your raft as this evidence means there are mink in your area.*

Mink Ecology:

Temporal - *We are in the normal mink breeding season!*

Motivation – *During the breeding season female mink will defend an area covering approximately 1.5 miles.*

Additional Information - *Female mink are small enough to fit into water vole burrows which they explore in search of prey. Did you know there can be brown, black, purple, white and silver mink which reflects the colours bred for fur?*

To produce the actual content to fill the slots of the schemas, the system was designed to reason over geographical location to allow examination of the various notions of neighbourhood (Tintarev et al 2012). The system also looks at temporal trends when developing text based on the number of record submissions for a given time. The system initially looks at record submissions in the past week then opens out to a month, season and finally activity between the same seasons on different years. This use of temporal trends ensures volunteers are supplied with the most relevant (recent) mink activity information first in busy periods such as the breeding season but ensures ‘cleared’ areas with little mink activity are still provided with informative feedback.

6 Evaluation of the Feedback Approach

We were initially apprehensive about how much usage the feedback system would get. MinkApp was launched through the SMI newsletters, but we knew that volunteers were not always receiving or reading these. Also it turned out that the initial estimate of active volunteers was over-inflated. Indeed, initially the usage of MinkApp in general was much lower than was expected. So we worked hard to promote the system, for instance asking the fisheries trusts to actively ask any volunteers they had contact with if they had heard of MinkApp and to try to use it. As a result, we did manage to increase the system usage to a level where some initial conclusions can be drawn.

MinkApp and specifically the feedback form use were monitored for 50 days (7 weeks). During this time 308 raft checks were submitted by volunteers for 98 different rafts by 44 unique users. The feedback system was used by volunteers to generate 113 different texts about 36 different rafts. 32 out of the 44 (72.7%) of all MinkApp users requested generated feedback at least once.

In 47% of the feedback form use sessions multiple texts were generated and there are some particularly interesting use patterns:

- “Regular explorer”: One user accessed MinkApp seven times and generated feedback text on every use: 1 text, 3 texts, 5 texts, 5 texts, 4 texts, 2 texts and 1 text

- “Periodic explorer”: One user accessed MinkApp six times and generated at least one feedback text on every second use
- “Try once only”: The user who accessed MinkApp the most with eleven different sessions only generated feedback text on their first use of MinkApp.

These different patterns of use require further investigation as the number of users using MinkApp increases. The patterns can be affected by idiosyncratic factors. For instance, one volunteer informed the project coordinator that they continually selected Captures within their area as they had caught a mink and their capture had not yet been added to the system - the volunteer was using the feedback form to monitor how long it took for mink capture data to appear in MinkApp.

Of the four types of information available to volunteers Signs was the most viewed although Captures was what SMI staff had felt volunteers would be most interested in. Signs had 56.6% of the overall use and catchment was the most widely selected option for geographic area for both Signs and Captures. However there was no clearly predominant second choice for information option with Captures and My Rafts having only 2.7% of a difference within their use. Mink Ecology was the least used category, partly to do with the lack of clarity in the name ‘Mink Ecology’. Signs on a local geographical scale were the most common selection for volunteers but the actual use was not clear enough to support a fixed text type or removing other options.

7 Conclusions

The results of this initial study did support the value of feedback to volunteers (more directly than we would have been able to determine in advance) with 73% of volunteers choosing to generate feedback. The feedback enabled us to offer contextualized information to volunteers quickly, without initial extensive user studies, which was very important for supporting the continuation of SMI.

The fact that the volunteer population was relatively unknown meant that there were some unpleasant surprises in terms of uptake and interest. It was necessary to make special efforts to encourage participation to get larger numbers.

When our system gets used over longer periods we might observe more meaningful patterns of behaviour.

The patterns of interest we observed were noisy and were influenced by many contextual factors meaning there was little potential yet for statistical analysis or machine learning.

8 Future Work

In-depth analysis is required as more volunteers use MinkApp and the feedback form to fully understand patterns of behaviour. Additionally qualitative studies such as interviews with volunteers could help explain use and preferences. These studies could help us improve the feedback system and text to better suit the user's needs. In the meantime, we have a working system that offers choices to users to 'generate their own text' even though we had hoped to be able to tailor to individual volunteer preferences sooner.

9 Acknowledgments

We would like to thank SMI for their on-going commitment to this research. This work is supported by the Rural Digital Economy Research Hub (EPSRC EP/G066051/1).

Reference

- Arts, K., Webster, G. ., Sharma, N. ., Melero, Y. ., Mellish, C., Lambin, X., & Van der Wal, R. (2013). Capturing mink and data. Interacting with a small and dispersed environmental initiative over the introduction of digital innovation Uploader. Case study for the online platform "Framework for Responsible Research and Innovation in ICT." Retrieved from <http://responsible-innovation.org.uk/torii/resource-detail/1059>
- Beirne, C., & Lambin, X. (2013). Understanding the Determinants of Volunteer Retention Through Capture-Recapture Analysis: Answering Social Science Questions Using a Wildlife Ecology Toolkit. *Conservation Letters*, 6(6), 391–401. doi:10.1111/conl.12023
- Brusilovsky, P. (2001). Adaptive Hypermedia. *User Modeling and User-Adapted Interaction*, 11(1-2), 87–110. doi:10.1023/A:1011143116306
- Bryce, R., Oliver, M. K., Davies, L., Gray, H., Urquhart, J., & Lambin, X. (2011). Turning back the tide of American mink invasion at an unprecedented scale through community participation and adaptive management. *Biological conservation*, 144(1), 575–583. Retrieved from <http://cat.inist.fr/?aModele=afficheN&cpsidt=23779637>
- Duboue, P. A., & McKeown, K. R. (2003). Statistical acquisition of content selection rules for natural language generation. In *Proceedings of the 2003 conference on Empirical methods in natural language processing - (Vol. 10, pp. 121–128)*. Morristown, NJ, USA: Association for Computational Linguistics. doi:10.3115/1119355.1119371
- Fischer, G. (2001). User Modeling in Human–Computer Interaction. *User Modeling and User-Adapted Interaction*, 11(1-2), 65–86. doi:10.1023/A:1011145532042
- Konstas, I., & Lapata, M. (2012). Concept-to-text generation via discriminative reranking, 369–378. Retrieved from <http://dl.acm.org/citation.cfm?id=2390524.2390576>
- Levine, J., Cawsey, A., Mellish, C., Poynter, L., Reiter, E., Tyson, P., & Walker, J. (1991). IDAS: Combining hypertext and natural language generation. In *Proc of the Third European Workshop on NLG (pp. 55–62)*. Innsbruck, Austria.
- Power, R., Scott, D., & Bouayad-Agha, N. (2003). Generating texts with style, 444–452. Retrieved from <http://dl.acm.org/citation.cfm?id=1791562.1791619>
- Reiter, E., & Dale, R. (2000). *Building Applied Natural Language Generation Systems*. clt.mq.edu.au (Vol. 33.). Cambridge: Cambridge university press. Retrieved from <http://clt.mq.edu.au/~rdale/publications/papers/1997/jnle97.pdf>
- Tintarev, N., & Masthoff, J. (2008). Adaptive Hypermedia and Adaptive Web-Based Systems. (W. Nejdl, J. Kay, P. Pu, & E. Herder, Eds.) (Vol. 5149, pp. 204–213). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-540-70987-9
- Tintarev, N., Melero, Y., Sripada, S., Tait, E., Van Der Wal, R., & Mellish, C. (2012). MinkApp: generating spatio-temporal summaries for nature conservation volunteers, 17–21. Retrieved from <http://dl.acm.org/citation.cfm?id=2392712.2392720>
- Zancanaro, M., Kuflik, T., Boger, Z., Goren-Bar, D., & Goldwasser, D. (2007). Analyzing museum visitors' behavior patterns. In C. Conati, K. McCoy, & G. Paliouras (Eds.), *11th International Conference on User Modeling (Vol. 4511, pp. 238–246)*. Berlin, Heidelberg: Springer Berlin Heidelberg. doi:10.1007/978-3-540-73078-1