

MinkApp: Generating Spatio-temporal Summaries for Nature Conservation Volunteers

Nava Tintarev, Yolanda Melero, Somayajulu Sripada,
Elizabeth Tait, Rene Van Der Wal, Chris Mellish

University of Aberdeen

{n.tintarev, y.melero, yaji.sripada,
elizabeth.tait, r.vanderwal, c.mellish@abdn.ac.uk}@abdn.ac.uk

Abstract

We describe preliminary work on generating contextualized text for nature conservation volunteers. This Natural Language Generation (NLG) differs from other ways of describing spatio-temporal data, in that it deals with abstractions on data across large geographical spaces (total projected area 20,600 km²), as well as temporal trends across longer time frames (ranging from one week up to a year). We identify challenges at all stages of the classical NLG pipeline.

1 Introduction

We describe preliminary work on summarizing spatio-temporal data, with the aim to generate contextualized feedback for wildlife management volunteers. The MinkApp project assesses the use of NLG to assist volunteers working on the Scottish Mink Initiative (SMI). This participatory initiative aims to safeguard riverine species of economic importance (e.g., salmon and trout) and species of nature conservation interest including water voles, ground nesting birds and other species that are actively preyed upon by an invasive non-native species - the American mink (Bryce et al., 2011).

2 Background

Our test ground is one of the world's largest community-based invasive species management programs, which uses volunteers to detect, and subsequently remove, American mink from an area of Scotland set to grow from 10,000 km² in 2010 to

20,600 km² by the end of 2013 (Bryce et al., 2011). Such a geographical expansion means that an increasing share of the monitoring and control work is undertaken by volunteers supported by a fixed number of staff. An important contribution of volunteers is to help collect data over a large spatial scale.

Involving members of the public in projects such as this can play a crucial role in collecting observational data (Silvertown, 2009). High profile examples of data-gathering programmes, labelled as citizen science, include Galaxy Zoo and Springwatch (Raddick et al., Published online 2010; Underwood et al., 2008). However, in such long-term and wide ranging initiatives, maintaining volunteer engagement can be challenging and volunteers must get feedback on their contributions to remain motivated to participate (Silvertown, 2009). NLG may serve the function of supplying this feedback.

3 Related work

We are particularly interested in summarizing raw geographical and temporal data whose semantics need to be computed at run time – so called spatio-temporal NLG. Such extended techniques are studied in data-to-text NLG (Molina and Stent, 2010; Portet et al., 2009; Reiter et al., 2005; Turner et al., 2008; Thomas et al., Published online 2010). Generating text from spatio-temporal data involves not just finding data abstractions, but also determining appropriate descriptors for them (Turner et al., 2008). Turner et. al (2008) present a case study in weather forecast generation where selection of spatial descriptors is partly based on domain specific (weather related) links between spatial descriptors

and weather phenomena. In the current project we see an opportunity to investigate such domain specific constraints in the selection of descriptors over larger temporal and spatial scales.

4 Current Status

Over 600 volunteers currently notify volunteer managers of their ongoing mink recording efforts. Our work is informed by in-depth discussions and interviews with the volunteer managers, as well as 58 (ground level) volunteers' responses to a questionnaire about their volunteering experience. The set of volunteers involves different people, such as conservation professionals, rangers, landowners and farmers with the degree of volunteer involvement varying among them. Most volunteers check for sightings: footprints on a floating platform with a clay-based tracking plate (raft hereafter) readily used by mink, or visual sightings on land or water. Others set and check traps, and (much fewer volunteers) dispatch trapped mink.¹ In terms of feedback, volunteers currently receive regional quarterly newsletters, but tailored and contextualized feedback is limited to sporadic personal communication, mostly via email.²

4.1 Why NLG in this context?

Where the initiative has been successful, mink sightings are sparse. Such a lack of sightings can be demotivating for volunteers and leads to a situation in which negative records are seldom recorded (Beirne, 2011). As one volunteer stated: *"Nothing much happens on my raft so my enthusiasm wanes."* Also, 73% of the volunteers who completed the questionnaire said they checked their raft at the recommended frequency of every two weeks. Similarly, 72% said that they got in touch with their manager rarely or only every couple of months – when they needed more clay or saw footprints. NLG based feedback could motivate volunteers by informing them about the value of negative records. If they were to stop because of a lack of interest, mink are likely to reinvade the area.

¹Traps are only placed once a sighting has occurred. Once placed, by law a trap must be checked daily.

²In this project, we are using a corpus based on newsletters from the North Scotland Mink Project and the Cairngorms Water Vole Conversation Project.

In addition, volunteers who work alone can be isolated and lack natural mechanisms for information exchange with peers. We postulate that giving the volunteers contextualized feedback for an area gives them a better feeling for their contribution to the project and a better sense of how the initiative is going overall. A need for this has already been felt by volunteers: *"Knowing even more about progress in the catchment would be good - and knowing in detail about water vole returning and latest mink sightings. It would be helpful to learn about other neighboring volunteers captures sightings in 'real time'."*

5 Approach

In this section we describe the generation of text in terms of a classic NLG pipeline, (Reiter and Dale, 2000), while addressing the additional tasks of interpreting the input data (from volunteers) to meaningful messages that achieve the desired communication goals: providing information to, as well as motivating volunteers. The NLG system which will generate these texts is actively under development.

5.1 Gold standard

Our nearest comparison is a corpus of domain specific conservation newsletters containing text such as the one below. These newsletters give us an idea of the type of structure and lexical choice applied when addressing volunteers, using both temporal and spatial summaries. However, these texts are not contextualized, or adapted to a particular volunteer.

"With an ever expanding project area, we are progressing exceptionally well achieving and maintaining areas free of breeding mink through-out the North of Scotland. Currently, the upper Spey, upper Dee and Ythan appear to be free of breeding mink, with only a few transients passing through..."

We would like to improve on these existing texts and aim to generate texts that are tailored and consider the context of the volunteer. The text below is developed from a template supplied from a volunteer manager in the process of corpus collection. In the following sections we describe the steps and challenges involved in the process of generating such a text.

“Thank you for your helpful contribution! You may have not seen any signs this time, but in the last week two people in the Spey catchment have seen footprints on their rafts. This means there might be a female with a litter in your neighborhood – please be on the lookout in the coming weeks! Capturing her could mean removing up to 6 mink at once!”

5.2 Example input

The data we receive from volunteers includes positive and negative records from raft checks (every 14 days), visual sightings, and mink captures. Each record contains a geographical reference (x and y coordinate) and a timestamp. In addition, for trapped mink we may know the sex (male, female, or unknown) and age (juvenile, adult, or unknown).

5.3 Data analysis and interpretation

Spatial trends. The current version of the system can reason over geographical information, defining various notions of neighborhood.³ For a given point the following attributes can be used to describe its neighborhood: geographical region (catchment and subcatchment), Euclidean distance from another point, and relative cardinal direction to another point (north, south, east, west). The system reasons about sightings and captures using facts such as:

- This point (on land or water) is in the Dee catchment.
- Three neighbors have seen footprints (within a given time window).
- One neighbor has caught a mink (within a given time window).
- The nearest mink footprint is 15 km north east of this point.

The definition of neighborhood will differ according to domain specific factors. Euclidean distance appears to be the most likely candidate for use, because sightings may belong to different geographic

³The reasoning is performed using the opensource GIS Java library Geotools, <http://geotools.org>, retrieved Jan 2012

regions (catchments) but be very close to each other. More importantly, the definition of neighborhood is likely to depend on the geographic region (e.g. areas differ in terms of mink population density with mountainous regions less likely to be utilized than coastal regions).

Temporal trends. Aside from geographic trends, the system will also be used to portray temporal trends. These look at the change in sightings between two time intervals, identifying it as a falling, rising or steady trend in mink numbers. We are primarily observing trends between different years, but also taking into consideration the ecology of the mink including their behavior in different seasons and for quantification. For example, we need to be able to decide if an increase from 0 to 5 mink sightings in an area during breeding is worth mentioning – most likely it is, as this a common size for a litter. Another example is the definition of a ‘cleared’ area - Example 1 below describes a stable zero trend over a longer period of time.

...Currently, the upper Spey, upper Dee and Ythan appear to be free of breeding mink...

(1)

5.4 Document planning

Content determination While useful on its own, the text that could be generated from the data analysis and interpretation described above is much more useful when domain specific rules are applied. Example 2 describes a significant year-on-year increase *for a given definition of neighborhood*, during breeding season.

```
IF ( (month >= 6 AND month <9)
AND sightingsLastYear(area) == 0
AND sightingsThisYear >= 5 )
THEN feedback +=
```

“It looks like the area has been reinvaded. We should get ready to trap them to keep this area mink free.”

(2)

Example rule 2 is applied in the breeding season (ca June-Aug.). It will be given a score which signifies its relative importance compared to other derived content to allow prioritization. For example,

if there are **both** female and male captures in a region, it would be more important to speak about the female capture. This is because the capture of breeding mink has a much larger positive impact on the success of the initiative.⁴ This importance should be reflected in texts such as: ...*Capturing her could mean removing up to 6 mink at once!*...

Document structuring Since our goal is to motivate as well as inform, the structure of the text will be affected. If we consider the example text in Section 5.1, we can roughly divide it into three summary types:

- **Personal** - “Thank you for your helpful contribution! You may have not seen any signs this time.”
- **Neighbor** - “In the last week two people in the Spey catchment have seen small footprints on their rafts.”
- **Biology** - “There might be a female with a litter in your neighborhood ... Capturing her could mean removing up to 6 mink at once!”

If, in contrast to the previous example, a volunteer *would* capture a mink, then the neighborhood summary can be used to emphasize the importance of rare captures.

“IF currentMonth == August AND
capture == true AND nCapturesInSummer == 0”
(3)

The feedback for rule 3 might read something like: “*Well done! So far, this was the only mink captured during the breeding season in the Spey catchment!*”

5.5 Microplanning

Microplanning will need to consider the aggregation of spatio-temporal data that happens on a deeper level e.g., for a given catchment and year. This aggregation is likely to result in a surface aggregation as well deeper data aggregation, such as the catchments in Example 1. In terms of lexical choice, the system will have to use domain appropriate vocabulary. The latter example refers to “breeding mink”,

⁴Established adult females with litters.

which informs the reader that their capture has a large impact on population control. Another example of lexical choice may be “quieter autumn” to denote a decrease in mink for an area.

The best way to communicate neighborhood to volunteers is still an open question. The texts in our corpus describe neighborhoods in terms of geographic regions (catchments and subcatchments, e.g. Spey). However, Euclidean distance may be more informative, in particular close to catchment boundaries.

6 Challenges

There are several key challenges when generating motivating text for nature conservation volunteers, using spatio-temporal NLG.

One challenge is to tailor feedback texts to individuals according to their motivations and information needs. In line with previous research in affective NLG (de Rosis and Grasso, 2000; Belz, 2003; Sluis and Mellish, 2010; Tintarev and Masthoff, 2012; Mahamood and Reiter, 2011), we continue to study the factors which are likely to have an effect on volunteer motivation. So far we have worked together with volunteer managers. We collected a corpus of texts, written by the managers, that are tailored to motivate different volunteer personas, and conducted interviews and a focus group with them. While we found that the mink managers tailored texts to different personas, interviews indicated that the biggest factor to tailor for was the definition of neighborhood. Some volunteers are interested in a local update, while others are interested in a larger overview.

A second, related challenge, regards correctly defining the reasoning over spatio-temporal facts e.g., quantifying the magnitude of significant changes (increases and decreases in sightings and captures) for different seasons, regions, and the time frames over which they occur. We believe this will lead to generating text referring to more compound abstractions such as mink free areas, or re-invasion.

A final challenge brought out by the interviews is to supply varied feedback that helps volunteers to continue to learn about mink and their habitat. This is a challenge for both content determination and microplanning.

References

- Christopher Beirne. 2011. Novel use of mark-recapture framework to study volunteer retention probabilities within an invasive non-native species management project reveals vocational and temporal trends. Master's thesis, University of Aberdeen.
- Anja Belz. 2003. And now with feeling: Developments in emotional language generation. Technical Report ITRI-03-21, Information Technology Research Institute, University of Brighton.
- Rosalind Bryce, Matthew K. Oliver, Llinos Davies, Helen Gray, Jamie Urquhart, and Xavier Lambin. 2011. Turning back the tide of american mink invasion at an unprecedented scale through community participation and adaptive management. *Biological Conservation*, 144:575–583.
- Fiorella de Rosis and Floriana Grasso, 2000. *Affective Interactions*, volume 1814 of *Lecture Notes in Artificial Intelligence*, chapter Affective Natural Language Generation. Springer-Verlag.
- Saad Mahamood and Ehud Reiter. 2011. Generating affective natural language for parents of neonatal infants. In *ENLG*.
- Martin Molina and Amanda Stent. 2010. A knowledge-based method for generating summaries of spatial movement in geographic areas. *International Journal on Artificial Intelligence Tools*, 19(3):393–415.
- Francois Portet, Ehud Reiter, Albert Gatt, Jim Hunter, Somayajulu Sripada, Yvonne Freer, and Cindy Sykes. 2009. Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 173:789–816.
- M. Jordan Raddick, Georgia Bracey, Pamela L. Gay, Chris J. Lintott, Phil Murray, Kevin Schawinski, Alexander S. Szalay, and Jan Vandenberg. Published online 2010. Galaxy zoo: Exploring the motivations of citizen science volunteers. *Astronomy Education Review*, 9(1), 010103, doi:10.3847/AER2009036.
- Ehud Reiter and Robert Dale. 2000. *Building natural language generation systems*. Cambridge University Press.
- Ehud Reiter, Somayajulu Sripada, Jim Hunter, Jin Yu, and Ian Davy. 2005. Choosing words in computer-generated weather forecasts. *Artificial Intelligence*, 167:137–169.
- Jonathan Silvertown. 2009. A new dawn for citizen science. *Trends in Ecology & Evolution*, 24:467–471.
- Ielka Van der Sluis and Chris Mellish, 2010. *Empirical Methods in Natural Language Generation*, volume 5980 of *Lecture Notes in Computer Science*, chapter Towards Empirical Evaluation of Affective Tactical NLG. Springer, Berlin / Heidelberg.
- Kavita E. Thomas, Somayajulu Sripada, and Matthijs L. Noordzij. Published online 2010. Atlas.txt: Exploring linguistic grounding techniques for communicating spatial information to blind users. *Journal of Universal Access in the Information Society*, DOI 10.1007/s10209-010-0217-5.
- Nava Tintarev and Judith Masthoff. 2012. Evaluating the effectiveness of explanations for recommender systems: Methodological issues and empirical studies on the impact of personalization. *User Modeling and User-Adapted Interaction*, (to appear).
- Ross Turner, Somayajulu Sripada, Ehud Reiter, and Ian Davy. 2008. Using spatial reference frames to generate grounded textual summaries of georeferenced data. In *INLG*.
- Joshua Underwood, Hilary Smith, Rosemary Luckin, and Geraldine Fitzpatrick. 2008. E-science in the classroom towards viability. *Computers & Education*, 50:535–546.